

PDP-10

System Reference Manual

ORDER NO. DEC-10-HGAA-D FROM PROGRAM LIBRARY, MAYNARD, MASSACHUSETTS

PRICE \$5.00

DIRECT COMMENTS CONCERNING THIS MANUAL TO SOFTWARE QUALITY CONTROL, MAYNARD, MASSACHUSETTS

DIGITAL EQUIPMENT CORPORATION • MAYNARD, MASSACHUSETTS

April 1968
Second printing, May 1968

Changes are indicated by a
triangle (▲) in the outside margin.

Copyright 1968 by
Digital Equipment Corporation

Written and designed for Digital Equipment Corporation by William English, Wayland, Massachusetts
Manufactured in the United States of America

1

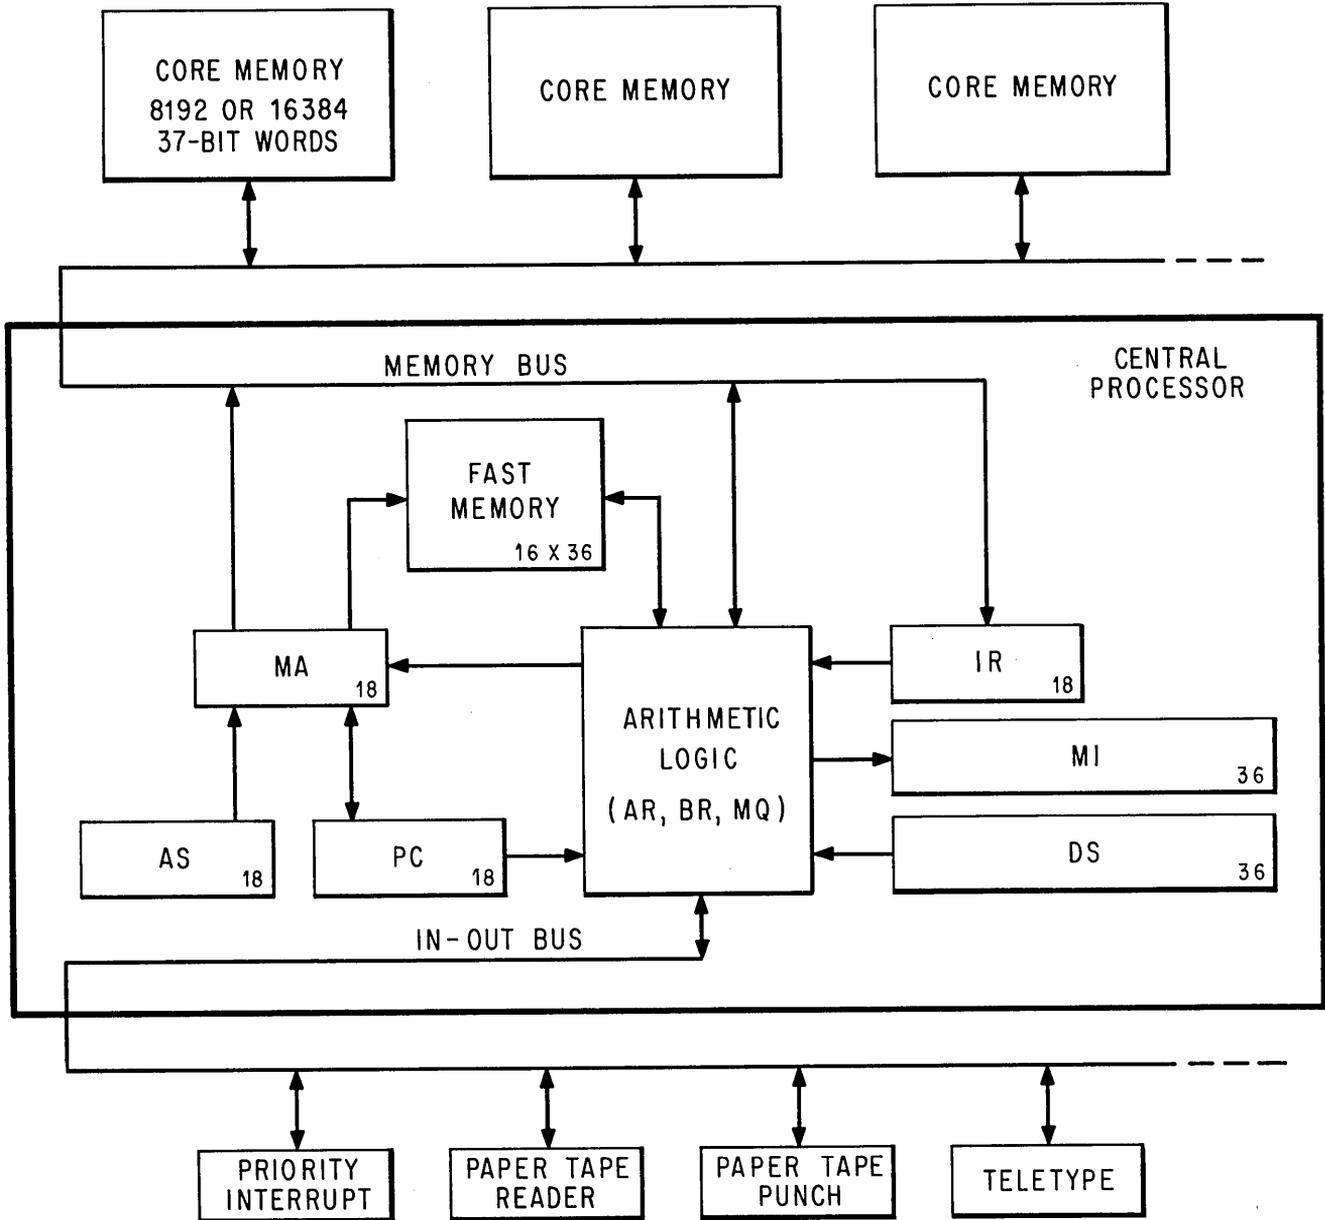
Introduction

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

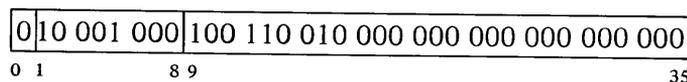
tion is $400000\ 000000_8$ and its magnitude is one greater than the largest positive number).

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.

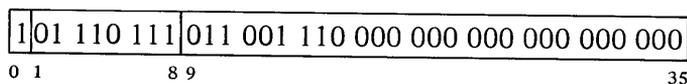
The computer does not keep track of a binary point — the programmer must adopt a point convention and shift the magnitude of the result to conform to the convention used. Two common conventions are to regard a number as an integer (binary point at the right) or as a proper fraction (binary point at the 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 fraction 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

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



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

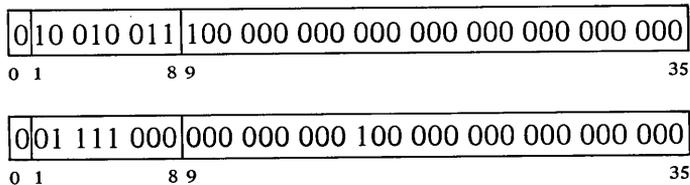


Multiplication produces a double length product, and the programmer must remember that discarding the low order part of a double length negative leaves the high order part in correct twos complement form only if the low order part is null.

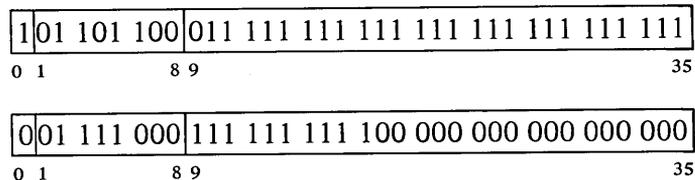
the instruction that scales the exponent, the hardware interprets the integral scale factor in standard twos complement form but produces the correct ones complement result for the exponent.

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 $\frac{1}{2}$ and less than 1. These numbers thus have a fractional range in magnitude of $\frac{1}{2}$ 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:



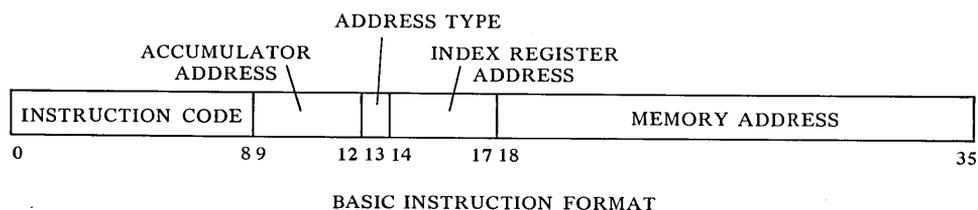
whereas its negative is



1.2 INSTRUCTION FORMAT

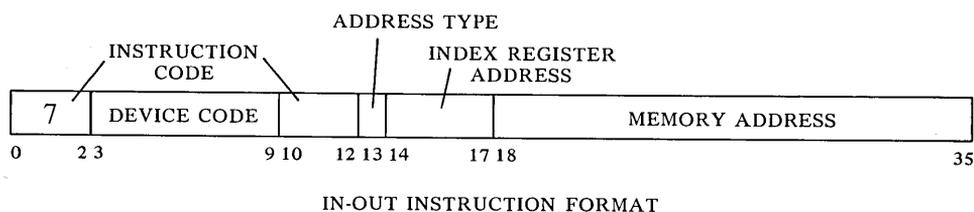
In all but the input-output instructions, the nine high order bits (0–8) specify the operation, and bits 9–12 usually address an accumulator but are sometimes used for special control purposes, such as addressing flags. The

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

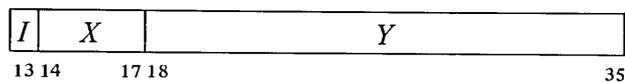


that are not assigned as specific instructions are executed by the processor as so-called “unimplemented operations”, as are the codes for floating point and byte manipulation in any PDP-10 that does not have the optional hardware for these instructions. When the processor encounters one of these unimplemented 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.



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.

	<i>Read</i>	<i>Write</i>
MA10 or MA10A Core Memory (1.00 μ s)	550	200
MB10 Core Memory (1.65 μ s)	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

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)

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

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 0001302, 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.

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.

<i>Operation</i>	<i>Suffix</i>	<i>Effect on Other Half of Destination</i>
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.

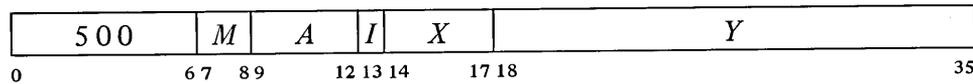
Mode	Suffix	Source	Destination
Basic		<i>E</i>	AC
Immediate	I	The word 0, <i>E</i>	AC
Memory	M	AC	<i>E</i>
Self	S	<i>E</i>	<i>E</i> , but also AC if <i>A</i> is nonzero

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) μ s in self mode; in memory mode the same saving results if no action is taken on the other half, otherwise .47 (.36) μ s is saved.

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

HLL Half Word Left to Left

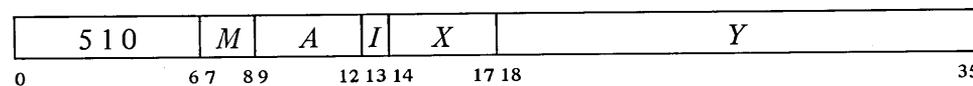


Move the left half of the source word specified by *M* to the left half of the specified destination. The source and the destination right half are unaffected; the original contents of the destination are lost.

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

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

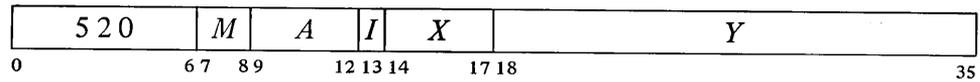


Move the left half of the source word specified by *M* to the left half of the specified destination, and clear the destination right half. The source is unaffected, the original contents of the destination are lost.

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

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

HLLO Half Word Left to Left, Ones

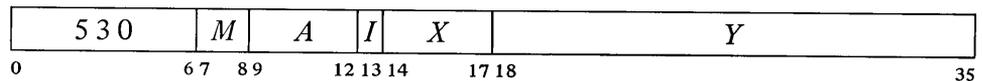


Move the left half of the source word specified by *M* to the left half of the specified destination, and set the destination right half to all 1s. The source is unaffected, the original contents of the destination are lost.

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

HLLO	Half Left to Left, Ones	520 2.21 (2.43) μ s
HLLOI	Half Left to Left, Ones, Immediate	521 1.36 (1.47) μ s
HLLOM	Half Left to Left, Ones, Memory	522 2.47 (2.58) μ s
HLLOS	Half Left to Left, Ones, Self	523 2.76 (2.87) μ s

HLLI Half Word Left to Left, Extend

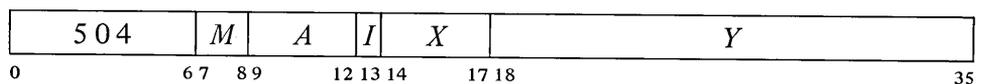


Move the left half of the source word specified by *M* to the left half of the specified destination, and make all bits in the destination right half equal to bit 0 of the source. The source is unaffected, the original contents of the destination are lost.

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

HLLI	Half Left to Left, Extend	530 2.21 (2.43) μ s
HLLI	Half Left to Left, Extend, Immediate	531 1.36 (1.47) μ s
HLLI	Half Left to Left, Extend, Memory	532 2.47 (2.58) μ s
HLLI	Half Left to Left, Extend, Self	533 2.76 (2.87) μ s

HRL Half Word Right to Left

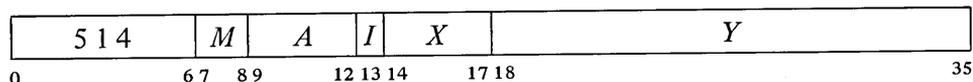


Move the right half of the source word specified by *M* to the left half of the specified destination. The source and the destination right half are unaffected; the original contents of the destination left half are lost.

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

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

HRLZ Half Word Right to Left, Zeros

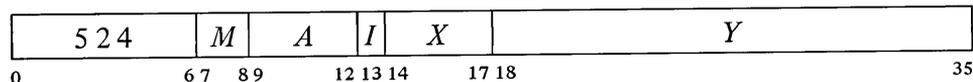


Move the right half of the source word specified by *M* to the left half of the specified destination, and clear the destination right half. The source is unaffected, the original contents of the destination are lost.

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

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

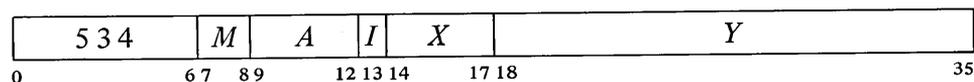
HRLO Half Word Right to Left, Ones



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) μ s
HRLOI	Half Right to Left, Ones, Immediate	525	1.36 (1.47) μ s
HRLOM	Half Right to Left, Ones, Memory	526	2.47 (2.58) μ s
HRLOS	Half Right to Left, Ones, Self	527	2.76 (2.87) μ s

HRLE Half Word Right to Left, Extend

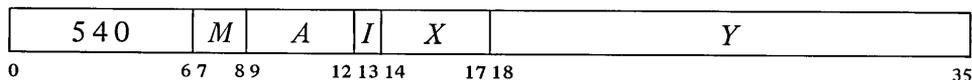


Move the right half of the source word specified by *M* to the left half of the

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) μ s	
HRLEI	Half Right to Left, Extend, Immediate	535	
		1.36 (1.47) μ s	
HRLEM	Half Right to Left, Extend, Memory	536	
		2.47 (2.58) μ s	
HRLES	Half Right to Left, Extend, Self	537	
		2.76 (2.87) μ s	

HRR Half Word Right to Right

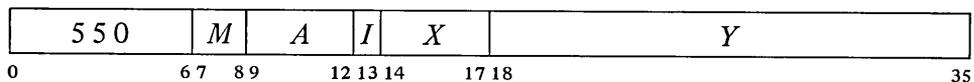


Move the right half of the source word specified by *M* to the right half of the specified destination. The source and the destination left half are unaffected; the original contents of the destination right half are lost.

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

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

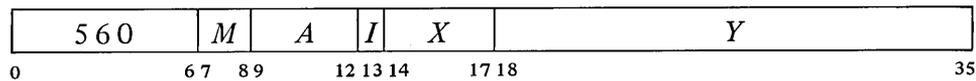
HRRZ Half Word Right to Right, Zeros



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.

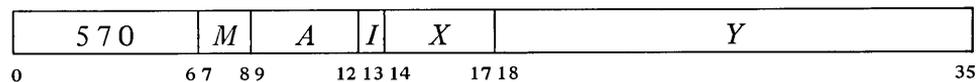
HRRZ	Half Right to Right, Zeros	550	2.21 (2.43) μ s
HRRZI	Half Right to Right, Zeros, Immediate	551	1.36 (1.47) μ s
HRRZM	Half Right to Right, Zeros, Memory	552	2.47 (2.58) μ s
HRRZS	Half Right to Right, Zeros, Self	553	2.76 (2.87) μ s

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

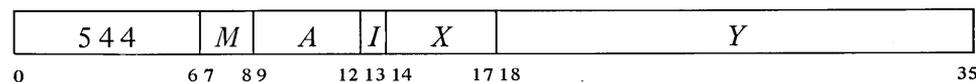
Move the right half of the source word specified by *M* to the right half of the specified destination, and set the destination left half to all 1s. The source is unaffected, the original contents of the destination are lost.

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

HRRE Half Word Right to Right, Extend

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) μ s
HRREI	Half Right to Right, Extend, Immediate	571	1.36 (1.47) μ s
HRREM	Half Right to Right, Extend, Memory	572	2.47 (2.58) μ s
HRRES	Half Right to Right, Extend, Self	573	2.76 (2.87) μ s

HLR Half Word Left to Right

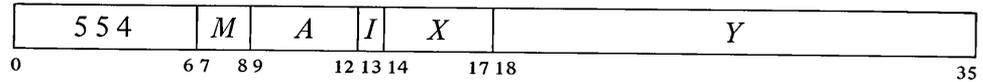
Move the left half of the source word specified by *M* to the right half of the specified destination. The source and the destination left half are unaffected; the original contents of the destination right half are lost.

HLR	Half Left to Right	544	2.70 (2.92) μ s
HLRI	Half Left to Right Immediate	545	1.85 (1.96) μ s

HLRI merely clears AC right.

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

HLRZ Half Word Left to Right, Zeros

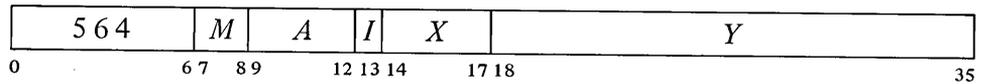


Move the left half of the source word specified by *M* to the right half of the specified destination, and clear the destination left half. The source is unaffected, the original contents of the destination are lost.

HLRZ	Half Left to Right, Zeros	554	2.21 (2.43) μ s
HLRZI	Half Left to Right, Zeros, Immediate	555	1.36 (1.47) μ s
HLRZM	Half Left to Right, Zeros, Memory	556	2.47 (2.58) μ s
HLRZS	Half Left to Right, Zeros, Self	557	2.76 (2.87) μ s

HLRZI merely clears AC and is thus equivalent to HLLZI.

HLRO Half Word Left to Right, Ones

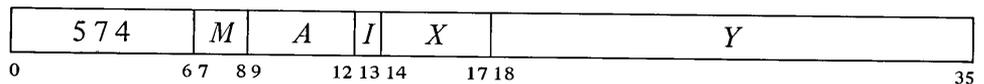


Move the left half of the source word specified by *M* to the right half of the specified destination, and set the destination left half to all 1s. The source is unaffected, the original contents of the destination are lost.

HLRO	Half Left to Right, Ones	564	2.21 (2.43) μ s
HLROI	Half Left to Right, Ones, Immediate	565	1.36 (1.47) μ s
HLROM	Half Left to Right, Ones, Memory	566	2.47 (2.58) μ s
HLROS	Half Left to Right, Ones, Self	567	2.76 (2.87) μ s

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

HLRE Half Word Left to Right, Extend



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) μ s	
HLREI	Half Left to Right, Extend, Immediate	575	
		1.36 (1.47) μ s	
HLREM	Half Left to Right, Extend, Memory	576	
		2.47 (2.58) μ s	
HLRES	Half Left to Right, Extend, Self	577	
		2.76 (2.87) μ s	

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
HRLRI 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.

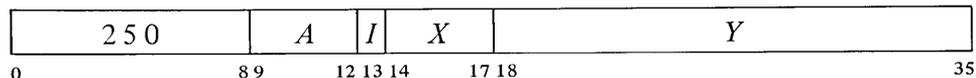
```
HLRZM XR, AC
HLLI XR, ;Clear XR left
```

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) μ s

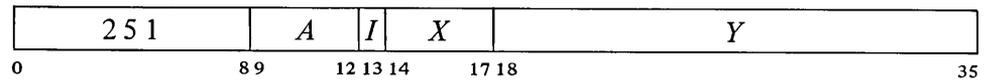


Keeping instructions and operands in different memories saves .20 (.09) μ s.

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) μ s plus 2.26 (2.48) μ s per transfer from fast memory to core and 2.61 (2.83) μ s per transfer from core to fast memory or from one core location to another.

BLT Block Transfer



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,

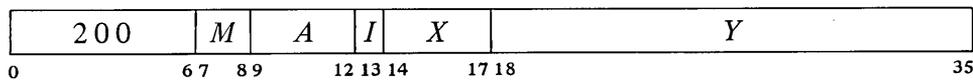
each with four modes that determine the source and destination of the word moved.

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

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

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

MOVE Move

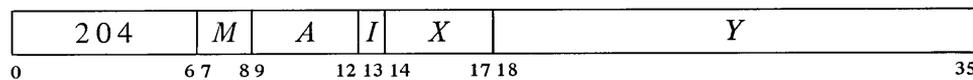


Move one word from the source to the destination specified by *M*. The source is unaffected, the original contents of the destination are lost.

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

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

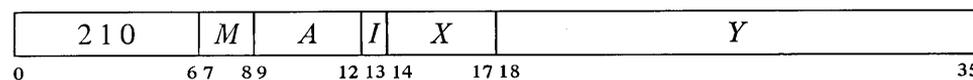


Interchange the left and right halves of the word from the source specified by *M* and move it to the specified destination. The source is unaffected, the original contents of the destination are lost.

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

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

MOVN Move Negative



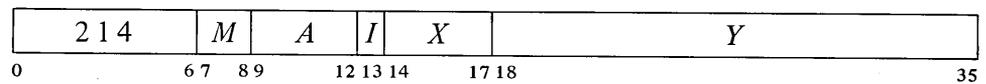
Negate the word from the source specified by *M* and move it to the specified destination. If the source word is fixed point -2^{35} (400000 000000) set the

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) μ s
MOVNI	Move Negative Immediate	211	1.54 (1.65) μ s
▲ MOVNM	Move Negative to Memory	212	2.65 (2.76) μ s
MOVNS	Move Negative to Self	213	2.94 (3.05) μ s

MOVMM Move Magnitude



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.

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

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

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
HRRR   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

```
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;

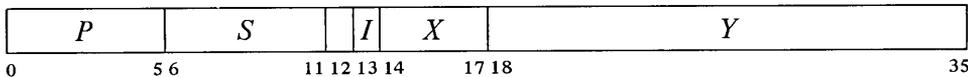
```
POP    P,-2(P)
```

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

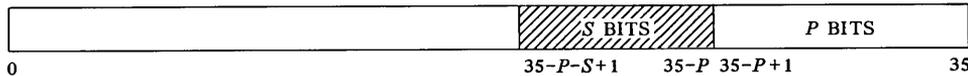
2.3 BYTE MANIPULATION

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



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



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.

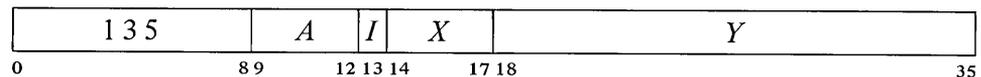
CAUTION

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 μ s. Taking bytes from a fast memory location saves another .34 μ s.

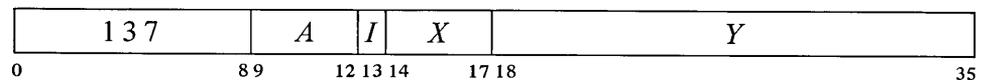
LDB Load Byte 4.02 (4.35) + .15($P + S$) [+ .26] μ s



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

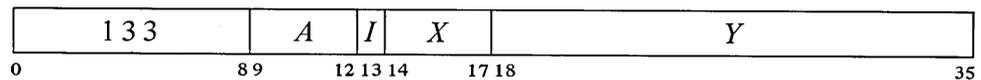
DPB Deposit Byte 4.87 (5.20) + .15($P + S$) [+ .26] μ s



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) μ s; keeping it in a different memory from the instruction saves .20 (.09) μ s.

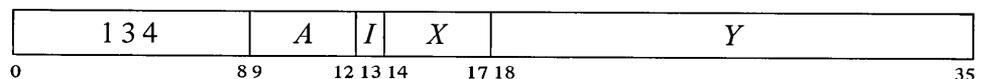
IBP Increment Byte Pointer 2.87 (2.98) [+ .26] μ s



Increment the byte pointer in location E as explained above.

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

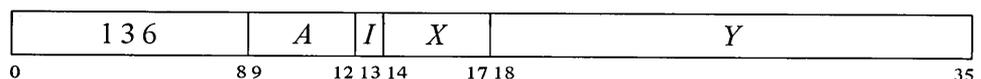
ILDB Increment Pointer and Load Byte 4.24 (4.57) + .15($P + S$) [+ .26] μ s



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

IDPB Increment Pointer and Deposit Byte 5.29 (5.51) + .15($P + S$) [+ .26] μ s



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.

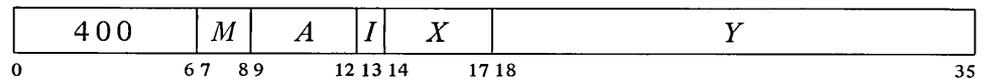
<i>Mode</i>	<i>Suffix</i>	<i>Source of non-AC operand</i>	<i>Destination of result</i>
Basic		E	AC
Immediate	I	The word 0, E	AC
Memory	M	E	E
Both	B	E	AC and E

Keeping instructions and operands in different memories saves .47 (.36) μ s in memory and both modes in the first four of these instructions (those that have no operand or only an AC operand), .20 (.09) μ s 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) μ s 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 μ s less in basic mode, .54 (.43) μ s 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

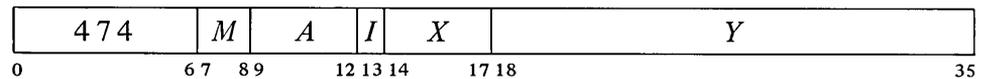


Change the contents of the destination specified by *M* to all 0s.

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) μ s
SETZI	Set to Zeros Immediate	401	1.36 (1.47) μ s
SETZM	Set to Zeros Memory	402	2.33 (2.44) μ s
SETZB	Set to Zeros Both	403	2.33 (2.44) μ s

SETO Set to Ones



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) μ s
SETOI	Set to Ones Immediate	475	1.36 (1.47) μ s
SETOM	Set to Ones Memory	476	2.33 (2.44) μ s
SETOB	Set to Ones Both	477	2.33 (2.44) μ s

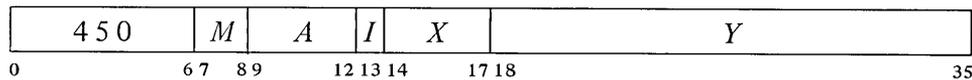
SETA Set to AC



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*).

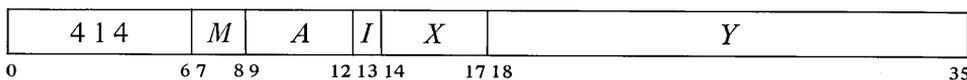
SETA	Set to AC	424	1.50 (1.61) μ s
SETAI	Set to AC Immediate	425	1.50 (1.61) μ s
SETAM	Set to AC Memory	426	2.47 (2.58) μ s
SETAB	Set to AC Both	427	2.47 (2.58) μ s

SETCA Set to Complement of AC

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

SETCA and SETCAI are equivalent (both complement AC).

SETM Set to Memory

Make the contents of the destination specified by *M* equal to the specified operand.

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

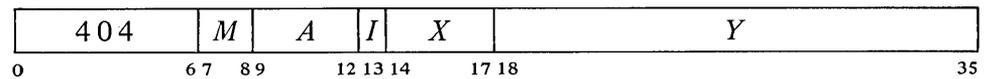
SETM and SETMB are equivalent to MOVE. SETMI moves the word 0,*E* to AC and is thus equivalent to MOVEI. SETMM is a no-op that references memory.

SETCM Set to Complement of Memory

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

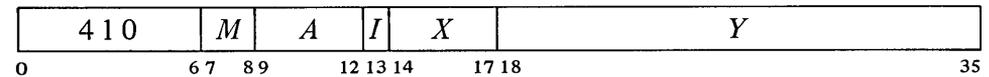
SETCM	Set to Complement of Memory	460	2.21 (2.43) μ s
SETCMI	Set to Complement of Memory Immediate	461	1.36 (1.47) μ s
SETCMM	Set to Complement of Memory Memory	462	2.76 (2.87) μ s
SETCMB	Set to Complement of Memory Both	463	2.76 (2.87) μ s

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

AND And with AC

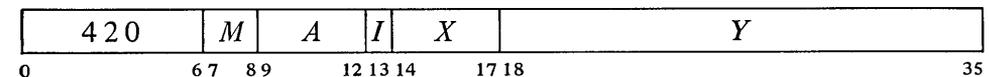
Change the contents of the destination specified by *M* to the AND function of the specified operand and AC.

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

ANDCA And with Complement of AC

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

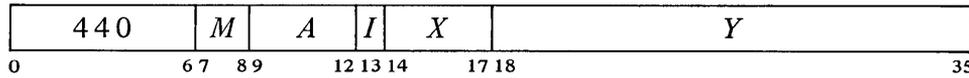
ANDCA	And with Complement of AC	410	2.70 (2.92) μ s
ANDCAI	And with Complement of AC Immediate	411	1.85 (1.96) μ s
ANDCAM	And with Complement of AC to Memory	412	3.52 (3.63) μ s
ANDCAB	And with Complement of AC to Both	413	3.52 (3.63) μ s

ANDCM And Complement of Memory with AC

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) μ s
ANDCMI	And Complement of Memory Immediate	421	1.50 (1.61) μ s
ANDCMM	And Complement of Memory to Memory	422	2.90 (3.01) μ s
ANDCMB	And Complement of Memory to Both	423	2.90 (3.01) μ s

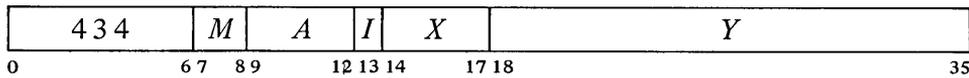
ANDCB And Complements of Both



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) μ s	
ANDCBI	And Complements of Both Immediate	441	
		1.85 (1.96) μ s	
ANDCBM	And Complements of Both to Memory	442	
		3.52 (3.63) μ s	
ANDCBB	And Complements of Both to Both	443	
		3.52 (3.63) μ s	

IOR Inclusive Or with AC

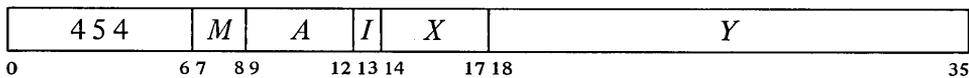


Change the contents of the destination specified by *M* to the inclusive OR function of the specified operand and AC.

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

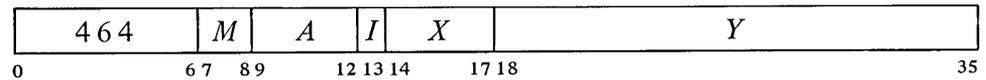
MACRO also recognizes OR, ORI, ORM and ORB as equivalent to the inclusive OR mnemonics.

ORCA Inclusive Or with Complement of AC



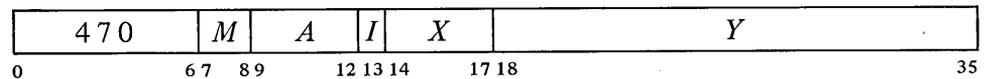
Change the contents of the destination specified by *M* to the inclusive OR function of the specified operand and the complement of AC.

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

ORCM Inclusive Or Complement of Memory with AC

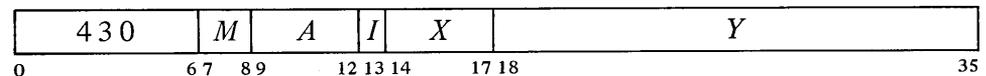
Change the contents of the destination specified by *M* to the inclusive OR function of the complement of the specified operand and AC.

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

ORCB Inclusive Or Complements of Both

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.

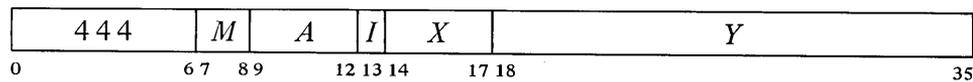
ORCB	Or Complements of Both	470	2.70 (2.92) μ s
ORCBI	Or Complements of Both Immediate	471	1.85 (1.96) μ s
ORCBM	Or Complements of Both to Memory	472	3.52 (3.63) μ s
ORCBB	Or Complements of Both to Both	473	3.52 (3.63) μ s

XOR Exclusive Or with AC

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) μ s
XORI	Exclusive Or Immediate	431	1.50 (1.61) μ s
XORM	Exclusive Or to Memory	432	2.90 (3.01) μ s
XORB	Exclusive Or to Both	433	2.90 (3.01) μ s

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

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) μ s
EQVI	Equivalence Immediate	445	1.50 (1.61) μ s
EQVM	Equivalence to Memory	446	2.90 (3.01) μ s
EQVB	Equivalence to Both	447	2.90 (3.01) μ s

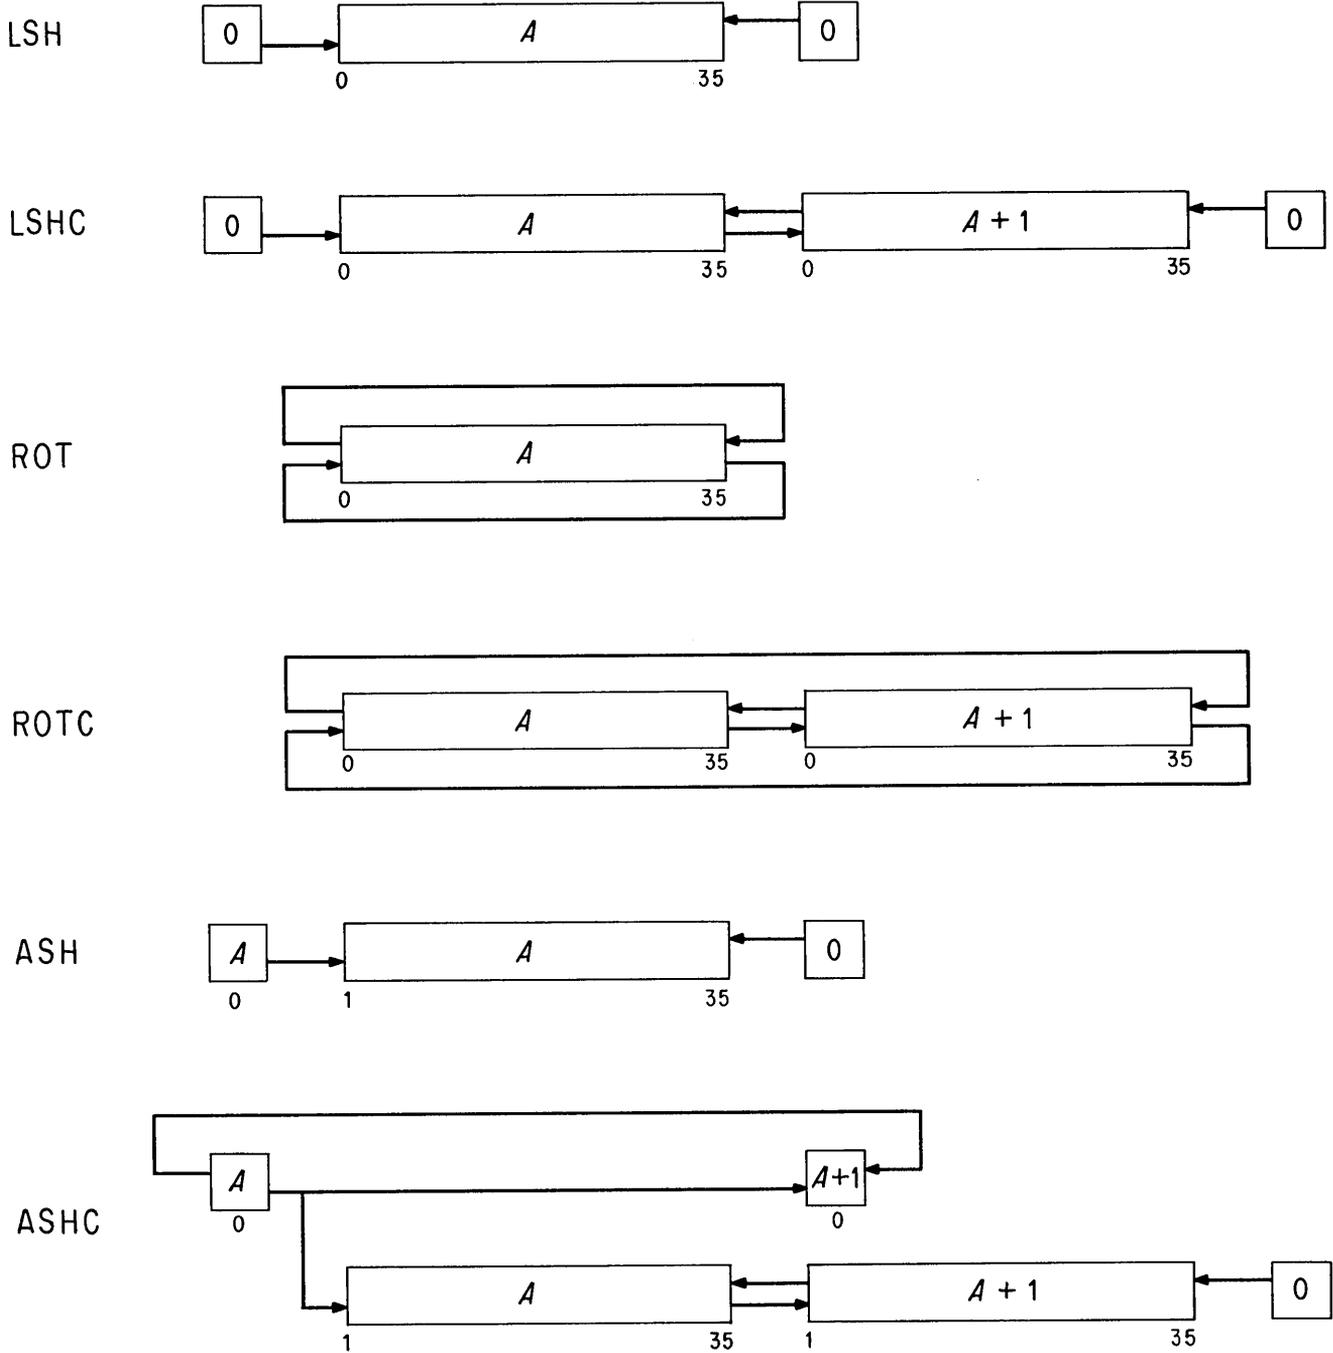
The original contents of the destination can be recovered except in EQVB, where both operands are replaced by the result. In the other three modes the replaced operand is restored by repeating the instruction in the same mode, *ie* by taking the equivalence function of the remaining operand and the result.

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.

	<i>AC</i>	0	1	0	1
<i>Mode Specified Operand</i>		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-



ACCUMULATOR BIT FLOW IN SHIFT AND ROTATE INSTRUCTIONS

ROTC Rotate Combined Left: $2.00 (2.11) + .15|E| \mu s$
 Right: $1.84 (1.95) + .15|E| \mu s$



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 $A+1 \pmod{20_8}$, 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.

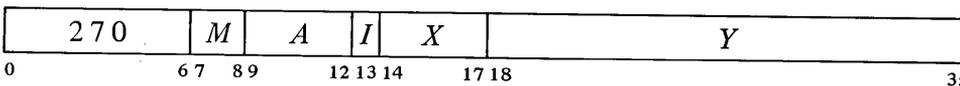
instructions [§2.14]) that can request a priority interrupt if enabled. The conditions detected can only set the flags and the hardware does not clear them, so the program must clear them before an instruction if they are to give meaningful information about the instruction afterward. However, the program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

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

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

ADD Add



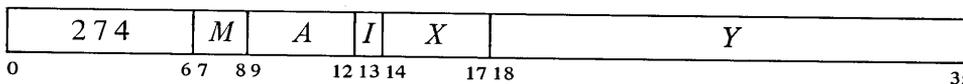
Add the operand specified by *M* to AC and place the result in the specified destination. If the sum is $\geq 2^{35}$ set Overflow and Carry 1; the result stored has a minus sign but a magnitude in positive form equal to the sum less 2^{35} . If the sum is $< -2^{35}$ set Overflow and Carry 0; the result 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) μ s
ADDI	Add Immediate	271	1.68 (1.79) μ s
ADDM	Add to Memory	272	3.08 (3.19) μ s
ADDB	Add to Both	273	3.08 (3.19) μ s

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

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

SUB Subtract



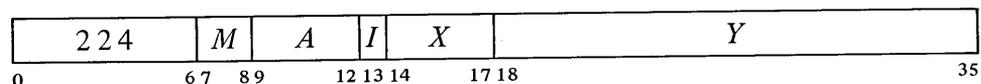
Subtract the operand specified by *M* from AC and place the result in the specified destination. If the difference is $\geq 2^{35}$ set Overflow and Carry 1; the result stored has a minus sign but a magnitude in positive form equal to the difference less 2^{35} . If the difference is $< -2^{35}$ set 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.

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

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

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

MUL Multiply



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} .

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

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

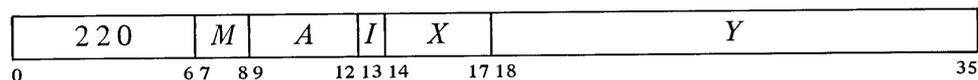
MUL	Multiply	224	10.60 (10.82) μ s
MULI	Multiply Immediate	225	8.58 (8.69) μ s
MULM	Multiply to Memory	226	11.41 (11.63) μ s
MULB	Multiply to Both	227	11.41 (11.63) μ s

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) μ s
MULI	7.41 (7.52) μ s
MULM	9.07 (9.29) μ s
MULB	9.07 (9.29) μ s

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

IMUL Integer Multiply



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 $\geq 2^{35}$ or $< -2^{35}$ (ie if the high order word of the double length product is not null); the high order word is lost.

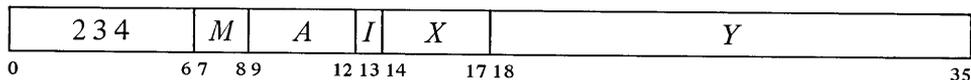
IMUL	Integer Multiply	220	9.59 (9.81) μs
IMULI	Integer Multiply Immediate	221	8.09 (8.20) μs
IMULM	Integer Multiply to Memory	222	10.56 (10.78) μs
IMULB	Integer Multiply to Both	223	10.56 (10.78) μs

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



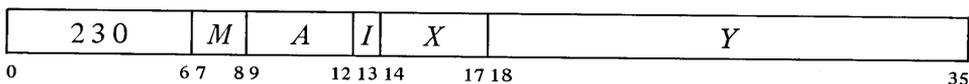
If the magnitude of the number in AC is greater than or equal to that of the operand specified by *M*, set Overflow and No Divide, and go immediately to the next instruction without affecting the original AC or memory operand in any way. Otherwise divide the double length number contained in accumulators *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.

IDIV Integer Divide



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

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

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 $A+1$.

Keeping instructions and operands in different memories saves .5 (.4) μ s in IDIVM, .3 (.2) μ s in IDIVB.

When E addresses a fast memory location, IDIV takes .3 μ s less than the time given, IDIVM takes .8 (.7) μ s less, and IDIVB takes .6 (.5) μ s less.

If the division is not performed, only 3–3.5 μ s are required.

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

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
c f a d g b e h

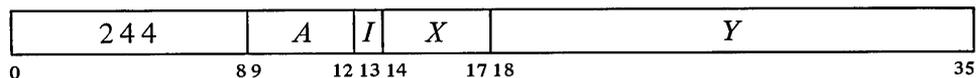
```

Since any digit is at most 1, there can be no carry out of any column with fewer than eight digits. 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

ASHC Arithmetic Shift Combined

Left: $2.00 (2.11) + .15|E| \mu s$
 Right: $1.84 (1.95) + .15|E| \mu s$



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 number 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 positive, 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 negative; 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, subtraction, 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 whether or not the result shall be rounded. Rounding produces the greatest consistent 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, subtraction and multiplication, wherein any bits of significance in the low order part supply information for normalization, and then for rounding if requested. 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 addition 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 containing only one bit of significance.

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 $+255$.

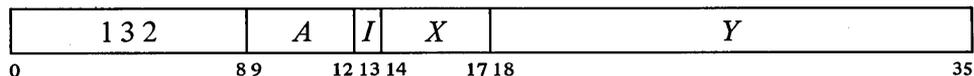
N is the number of left shifts needed to normalize the result.

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 $2.75 (2.86) + .25N \mu s$



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 $< \frac{1}{2}$ LSB. Moreover an extra single-step re-normalization occurs if the rounded word is no longer normalized.

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.

Keeping instructions and operands in different memories saves $.47 (.36) \mu s$ in memory and both modes.

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

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

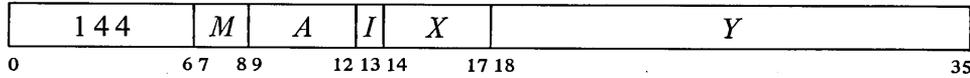
Note however that floating point immediate uses $E, 0$ as an operand, not $0, E$. In other words the half word E is interpreted as a sign, an 8-bit exponent, and a 9-bit fraction.

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

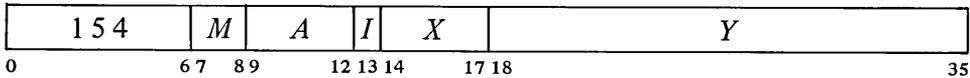


Floating add the operand specified by *M* to AC. If the double length fraction in the sum is zero, clear the specified destination. Otherwise normalize the double length sum bringing 0s into bit positions vacated at the right, round the high order part, test for exponent overflow or underflow as described 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 ≤ 63 . Otherwise $D = 0$.

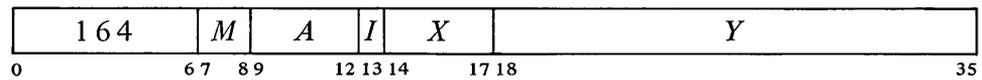
FSBR Floating Subtract and Round



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 ≤ 63 . Otherwise $D = 0$.

FMPR Floating Multiply and Round

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.

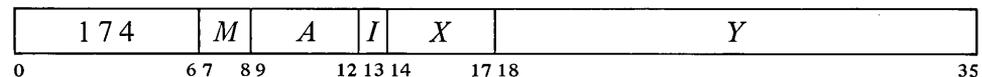
Use of normalized operands requires at most one normalization step for the result. If unnormalized operands are used, all times must be increased by .25*N*.

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

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] μ s
FMPRI	7.71 (7.82) [+ .96] μ s
FMPRM	9.44 (9.66) [+ .96] μ s
FMPRB	9.44 (9.66) [+ .96] μ s

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

FDVR Floating Divide and Round

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 the single-length quotient will already be normalized if the original operands were normalized; in this case, round it using the extra bit calculated. If the quotient is not normalized, do so bringing first the extra calculated bit and then 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.

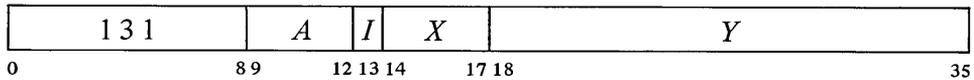
FDVR	Floating Divide and Round	174
		14.1 (14.3) μ s
FDVRI	Floating Divide and Round Immediate	175
		13.3 (13.4) μ s
FDVRM	Floating Divide and Round to Memory	176
		15.1 (15.3) μ s
FDVRB	Floating Divide and Round to Both	177
		15.1 (15.3) μ s

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



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) μ s. ▲

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.

UFA Unnormalized Floating Add 4.62 (4.84) + .15 D μ s



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

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 ≥ 1 , and place the high order part of the result in AC $A+1$. The original contents of AC and E are unaffected.

When E addresses a fast memory location, UFA takes .34 μ s 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) μ s in memory and both modes.

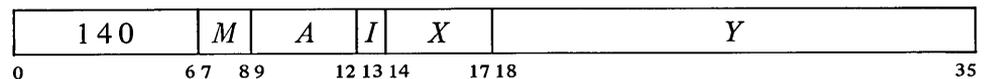
When E addresses a fast memory location, a floating point instruction without rounding takes .34 μ s less than the time listed in basic or long mode, .80 (.69) μ s less in memory or both mode.

<i>Mode</i>	<i>Suffix</i>	<i>Effect</i>
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



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-

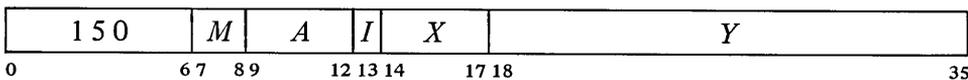
mulators in long mode. Otherwise normalize the double length sum bringing 0s into bit positions vacated at the right, test for exponent overflow or underflow as described above, and place the high order word of the result in the specified destination.

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
		4.46 (4.68) + .15D + .25N μ s
FADL	Floating Add Long	141
		5.31 (5.53) + .15D + .25N μ s
FADM	Floating Add to Memory	142
		5.43 (5.65) + .15D + .25N μ s
FADB	Floating Add to Both	143
		5.43 (5.65) + .15D + .25N μ s

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

FSB Floating Subtract



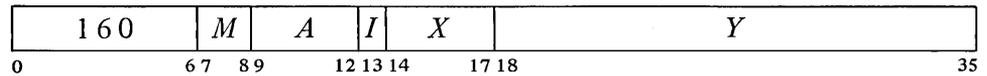
Floating subtract the contents of location *E* from AC. If the double length fraction in the difference is zero, clear the destination specified by *M*, clearing both accumulators in long mode. Otherwise normalize the double length difference bringing 0s into bit positions vacated at the right, test 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 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
		4.64 (4.86) + .15D + .25N μ s
FSBL	Floating Subtract Long	151
		5.49 (5.71) + .15D + .25N μ s
FSBM	Floating Subtract to Memory	152
		5.61 (5.83) + .15D + .25N μ s
FSBB	Floating Subtract to Both	153
		5.61 (5.83) + .15D + .25N μ s

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

FMP Floating Multiply



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.

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

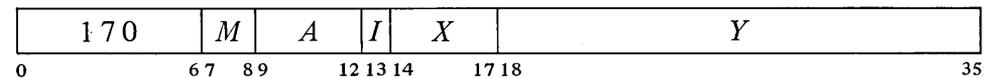
FMP	Floating Multiply	160	10.29 (10.51) μ s
FMPL	Floating Multiply Long	161	11.14 (11.36) μ s
FMPM	Floating Multiply to Memory	162	11.26 (11.48) μ s
FMPB	Floating Multiply to Both	163	11.26 (11.48) μ s

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) μ s
FMPL	9.32 (9.54) μ s
FMPM	9.44 (9.66) μ s
FMPB	9.44 (9.66) μ s

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

FDV Floating Divide



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 non-zero 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.

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

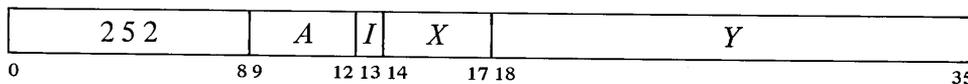
In long mode a nonzero unnormalized dividend whose entire high order fraction is zero produces a zero quotient. In this case the second AC receives rubbish.

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

2.7 ARITHMETIC TESTING

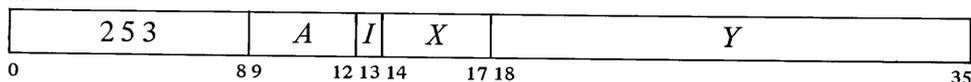
These instructions may jump or skip depending on the result of an arithmetic test and may first perform an arithmetic operation on the test word. Two of the instructions have no modes.

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



Add 1000001_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) μ s



Add 1000001_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.

<i>Mode</i>	<i>Suffix</i>
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) μ s



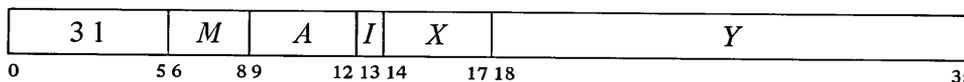
Compare AC with E (ie with the word 0, E) and skip the next instruction in sequence if the condition specified by M is satisfied.

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

CAI is a no-op.

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

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

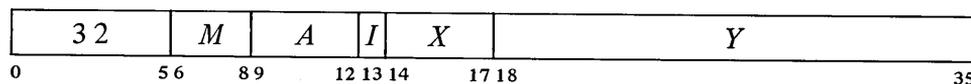


Compare AC with the contents of location *E* and skip the next instruction in sequence if the condition specified by *M* is satisfied. The pair of numbers compared may be either both fixed or both normalized floating point.

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) μ s



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).

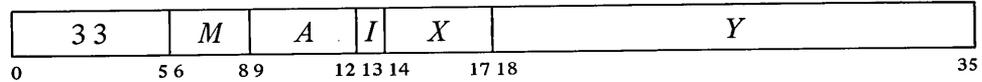
JUMPLE	Jump if AC Less than or Equal to Zero	323
JUMPA	Jump Always	324
JUMPGE	Jump if AC Greater than or Equal to Zero	325
JUMPN	Jump if AC Not Equal to Zero	326
JUMPG	Jump if AC Greater than Zero	327

When E addresses a fast memory location, this instruction takes $.34 \mu\text{s}$ less than the time given.

If A is zero, SKIP is a no-op; otherwise it is equivalent to MOVE. (Instruction code 330 has mnemonic SKIP for symmetry.)

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

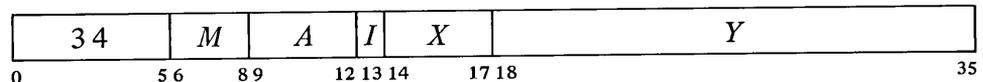
SKIP Skip if Memory Condition Satisfied 2.39 (2.61) μs



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.

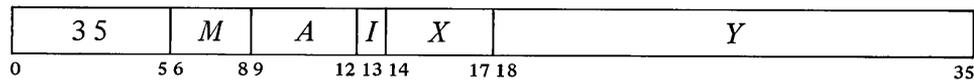
SKIP	Do Not Skip	330
SKIPL	Skip if Memory Less than Zero	331
SKIPE	Skip if Memory Equal to Zero	332
SKIPLE	Skip if Memory Less than or Equal to Zero	333
SKIP A	Skip Always	334
SKIPGE	Skip if Memory Greater than or Equal to Zero	335
SKIPN	Skip if Memory Not Equal to Zero	336
SKIPG	Skip if Memory Greater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

A0J Add One to AC and Jump if Condition Satisfied 1.68 (1.79) μs



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.

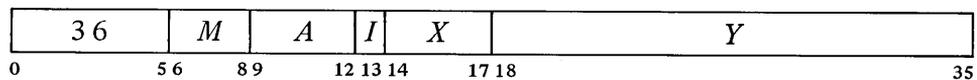
A0J	Add One to AC but Do Not Jump	340
A0JL	Add One to AC and Jump if Less than Zero	341
A0JE	Add One to AC and Jump if Equal to Zero	342
A0JLE	Add One to AC and Jump if Less than or Equal to Zero	343
A0JA	Add One to AC and Jump Always	344
A0JGE	Add One to AC and Jump if Greater than or Equal to Zero	345
A0JN	Add One to AC and Jump if Not Equal to Zero	346
A0JG	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) μ s

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.

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

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

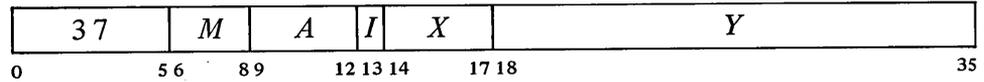
SOJ Subtract One from AC and Jump if Condition Satisfied 1.68 (1.79) μ s

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 instruction 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) μ s; keeping it in a different memory from the instruction saves .20 (.09) μ s.

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



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.

<i>Mask</i>	<i>Letter</i>	<i>Effect</i>
Right	R	AC right is masked by <i>E</i> (AC is masked by the word 0, <i>E</i>)
Left	L	AC left is masked by <i>E</i> (AC is masked by the word <i>E</i> , 0)
Direct	D	AC is masked by the contents of location <i>E</i>
Swapped	S	AC is masked by the contents of location <i>E</i> with left and right halves interchanged

If a direct or swapped mask is taken from a fast memory location, a test instruction takes .34 μ s less than the time listed.

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

<i>Modification</i>	<i>Letter</i>	<i>Effect on AC</i>
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

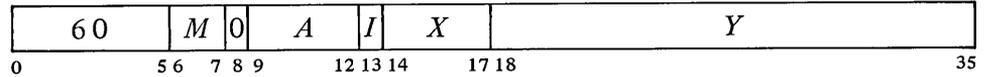
An additional letter may be appended to indicate the mode, which specifies the condition the masked bits must satisfy to produce a skip.

<i>Mode</i>	<i>Suffix</i>	<i>Effect</i>
Never		Never skip
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)

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.

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 unconditional 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 Satisfied** 1.85 (1.96) μ s

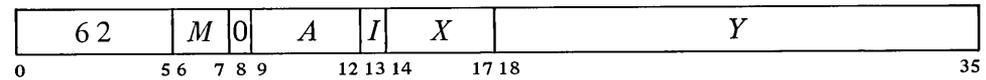


If the bits in AC right corresponding to 1s in *E* satisfy the condition specified by *M*, skip the next instruction in sequence. AC is unaffected.

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

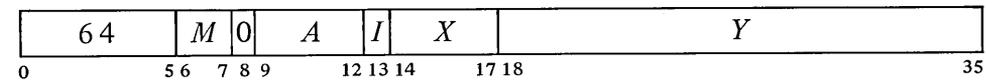
TRZ **Test Right, Zeros, and Skip if Condition Satisfied** 1.85 (1.96) μ s



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

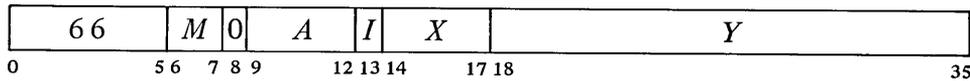
TRC **Test Right, Complement, and Skip if Condition Satisfied** 1.85 (1.96) μ s



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

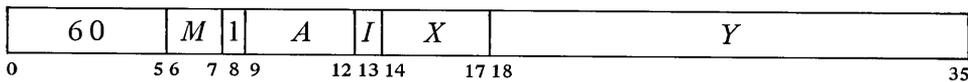
TRO Test Right, Ones, and Skip if Condition Satisfied 1.85 (1.96) μ s



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) μ s

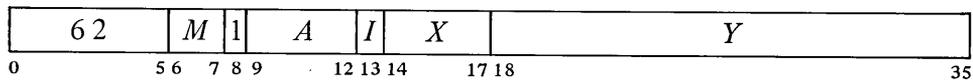


If the bits in AC left corresponding to 1s in *E* satisfy the condition specified by *M*, skip the next instruction in sequence. AC is unaffected.

TLN	Test Left, No Modification, but Do Not Skip	601
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

TLN is a no-op.

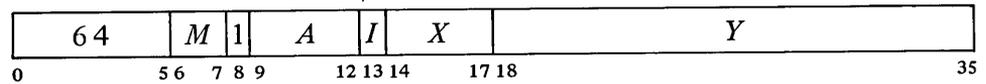
TLZ Test Left, Zeros, and Skip if Condition Satisfied 1.85 (1.96) μ s



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

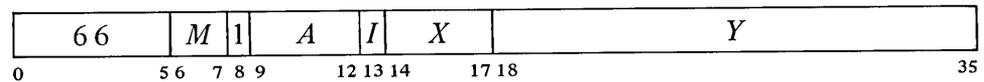
TLC Test Left, Complement, and Skip if Condition Satisfied 1.85 (1.96) μ s



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

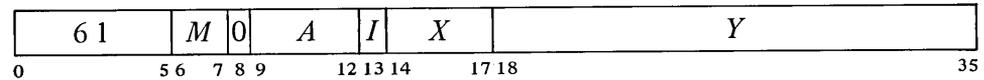
TLO Test Left, Ones, and Skip if Condition Satisfied 1.85 (1.96) μ s



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

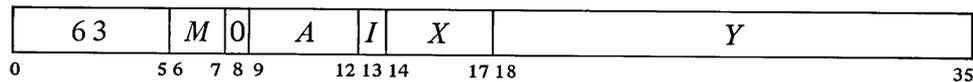
TDN Test Direct, No Modification, and Skip if Condition Satisfied 2.70 (2.92) μ s



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.

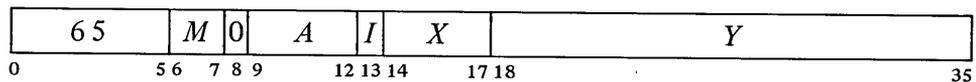
TDN is a no-op that references memory.

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) μ s

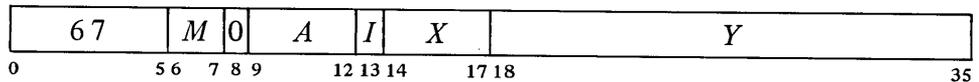
If the bits in AC corresponding to 1s in the contents of location *E* satisfy the condition specified by *M*, skip the next instruction in sequence. Change the masked AC bits to 0s; the rest of AC is unaffected.

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) μ s

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.

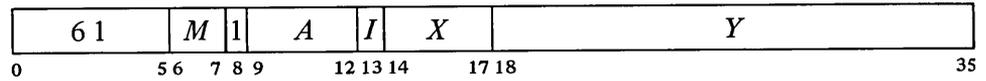
TDC	Test Direct, Complement, but Do Not Skip	650
TDCE	Test Direct, Complement, and Skip if All Masked Bits Equaled 0	652
TDCA	Test Direct, Complement, but Always Skip	654
TDCN	Test Direct, Complement, and Skip if Not All Masked Bits Equaled 0	656

TDO Test Direct, Ones, and Skip if Condition Satisfied 2.70 (2.92) μ s

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 Condition Satisfied** 2.70 (2.92) μ s

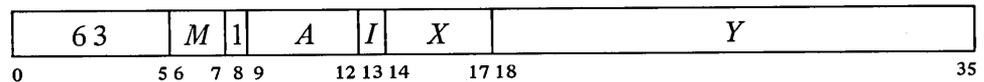


If the bits in AC corresponding to 1s in the contents of location *E* with its left and right halves swapped satisfy the condition specified by *M*, skip the next instruction in sequence. AC is unaffected.

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 Masked Bits Equal 0	613
TSNA	Test Swapped, No Modification, but Always Skip	615
TSNN	Test Swapped, No Modification, and Skip if Not All Masked Bits Equal 0	617

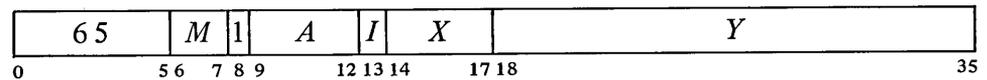
TSZ **Test Swapped, Zeros, and Skip if Condition Satisfied** 2.70 (2.92) μ s



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 Equaled 0	633
TSZA	Test Swapped, Zeros, but Always Skip	635
TSZN	Test Swapped, Zeros, and Skip if Not All Masked Bits Equaled 0	637

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

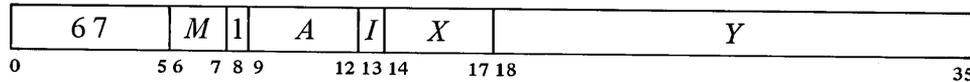


If the bits in AC corresponding to 1s in the contents of location *E* with its left and right halves swapped satisfy the condition specified by *M*, skip the next instruction in sequence. Complement the masked AC bits; the rest of AC is unaffected.

TSC	Test Swapped, Complement, but Do Not Skip	651
TSCE	Test Swapped, Complement, and Skip if All Masked Bits Equaled 0	653

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

TSO **Test Swapped, Ones, and Skip if Condition Satisfied** 2.70 (2.92) μ s



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

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 procedure, 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:

TRO F,*M* TLO F,*M*

and tested and cleared by one of these:

TRZE F,*M* TRZN F,*M* TLZE F,*M* TLZN F,*M*

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

SETCM F,L
TRNE 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 accumulator 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.

<i>Bit</i>	<i>Meaning of a 1 in the Bit</i>
------------	----------------------------------

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 $\geq 2^{35}$ or $< -2^{35}$.

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

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} .

An *MOVNX* or *MOV MX* has negated -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 $\geq 2^{35}$.

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

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

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].

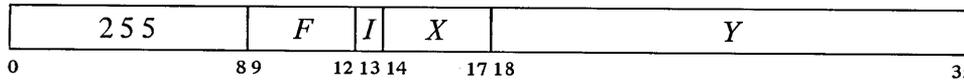
Remember [§2.5], overflow is determined directly from the carries, not from the flags. The carry flags give meaningful information only if no more than one instruction that can set them occurs between clearing and reading them.

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) μ s



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.

<i>Bit</i>	<i>Flag Selected by a 1</i>
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).

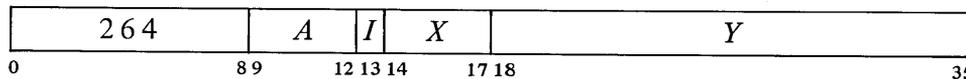
JFCL	JFCL 0,	No-op	25500
JOV	JFCL 10,	Jump on Overflow	25540
JCRY0	JFCL 4,	Jump on Carry 0	25520
JCRY1	JFCL 2,	Jump on Carry 1	25510
JCRY	JFCL 6,	Jump on Carry 0 or 1	25530
JFOV	JFCL 1,	Jump on Floating Overflow	25504

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** 2.21 (2.43) μ s



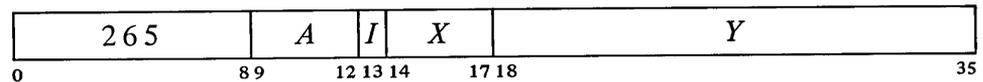
Place the current contents of the flags (as described above) in the left half of location *E* and the contents of PC in the right half (at this time PC contains an address one greater than the location of the JSR instruction). Take the next instruction from location *E* + 1 and continue sequential operation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

If this instruction is executed as a result of a priority interrupt or in unrellocated 41 or 61 while the processor is in user mode, bit 5 of the PC word stored is 1 and the processor leaves user mode.

Interleaving memories saves .47 (.36) μ s.

The *A* portion of this instruction is ignored.

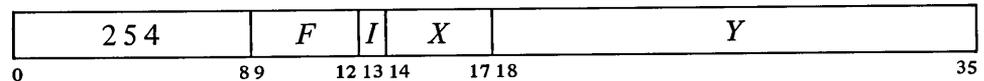
JSP **Jump and Save PC** 1.36 (1.47) μ s



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.

If this instruction is executed as a result of a priority interrupt or in un-relocated 41 or 61 while the processor is in user mode, bit 5 of the PC word stored is 1 and the processor leaves user mode.

JRST **Jump and Restore** 1.36 (1.47) μ s



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].

- 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).

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

This is identical to UWO trapping [§2.10].

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 UWO.

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	25440
		Interrupt Channel	
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) μ s, a JRSTF @PCWORD takes 2.34 (2.56) μ s.

CAUTION

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

If this instruction is executed as a result of a priority interrupt or in unrelocated 41 or 61 while the processor is in user mode, 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.

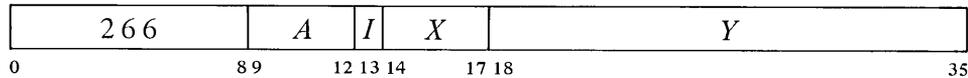
JEN completes an interrupt by restoring the channel and restoring the flags for the interrupted program.


```

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) μ s

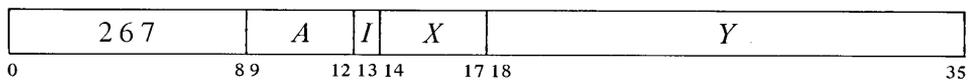
Interleaving memories saves .47 (.36) μ s.



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.

If this instruction is executed as a result of a priority interrupt or in unrelocated 41 or 61 while the processor is in user mode, bit 5 of the PC word stored is 1 and the processor leaves user mode.

JRA Jump and Restore AC 2.92 (3.14) μ s



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.

```

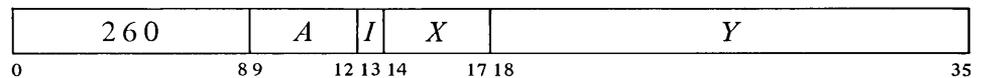
:                               ;Main program
:                               ;
:                               ;
JSA    17,S1                    ;Call to first subroutine (A)
:                               ;
S1:    0                          ;First subroutine starts here
:                               ;
:                               ;
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) μ s.

PUSHJ Push Down and Jump 3.00 (3.11) μ s



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.

If this instruction is executed as a result of a priority interrupt or in unrelocated 41 or 61 while the processor is in user mode, bit 5 of the PC word stored is 1 and the processor leaves user mode.

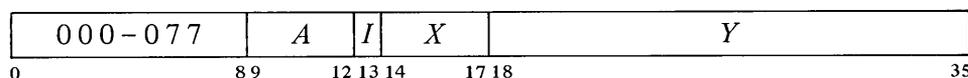
2.10 UNIMPLEMENTED OPERATIONS

An unimplemented user operation is usually referred to as a UO, but this mnemonic means nothing to the assembler. UOs are also sometimes called “programmed operators”.

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.

The total time required is that listed plus the time for the instruction in location 41. Interleaving memories 0 and 1 saves .47 (.36) μ s.

Unimplemented User Operation

2.33 (2.44) μ s

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 UO 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 communication 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 UO 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 UOs 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 equivalent 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 some 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 two's 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

$$\begin{aligned}
 N/b &= q_1 + r_1/b & r_1 < b \\
 q_1/b &= q_2 + r_2/b & r_2 < b \\
 q_2/b &= q_3 + r_3/b & r_3 < b \\
 &\vdots \\
 q_{n-1}/b &= 0 + r_n/b & r_n < b
 \end{aligned}$$

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

$$\begin{aligned}
 61/8 &= 7 + 5/8 \\
 7/8 &= 0 + 7/8
 \end{aligned}$$

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           ;128 = 1010
         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.

```

        PUSH  P,T+1 → HRLM T+1,(P)
        POP   P,T  → HLRZ  T,(P)

```

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

MACRO interprets a number following \uparrow D as decimal.

```

        LSHC   T,- $\uparrow$ 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

```

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.

```

FIND:  MOVE   T,(T)           ;Move next item to T
        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

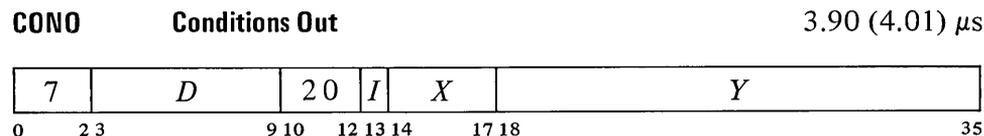
The input-output instructions govern all transfers of data to and from the peripheral equipment, and also perform many operations within the processor. An instruction in the in-out class is designated by 111 in bits 0-2, *ie* its left octal digit is 7. Bits 3-9 address the device that is to respond to 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 in Appendix A lists all devices for which codes have been assigned, and gives their mnemonics and DEC option numbers.

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



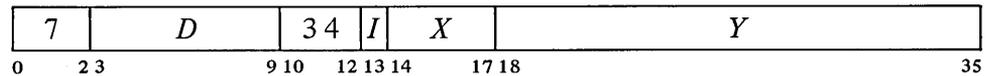
Set up device *D* with the effective initial conditions *E*. The number of condition bits in *E* that are actually used depends on the device.

- ▲ 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).

This is identical to UWO trapping [§2.10].

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.

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

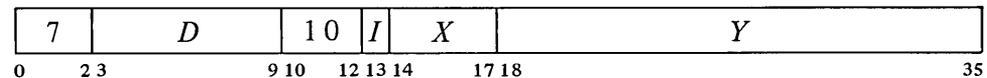


Test the input conditions from device *D* against the effective mask *E*. If any condition bit selected by a 1 in *E* is 1, skip the next instruction in sequence.

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

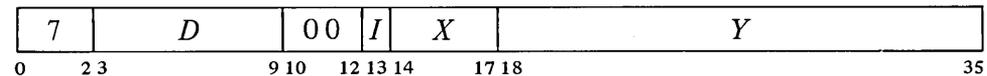
Keeping the pointer in fast memory saves .43 (.34) μ s.

▲ **BLKO** **Block Out** 6.49 (6.71) μ s



Keeping the pointer in fast memory saves .34 μ s. Keeping the instruction and the data block in different memories saves .47 (.36) μ s.

▲ **BLKI** **Block In** 6.49 (6.71) μ s



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 1000001_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.

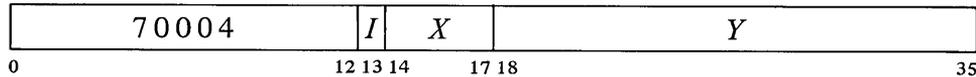
To read in, the operator must set up the device he is using, set its code into the readin device switches, and press the readin key. The processor places the device in operation, brings the first word (the pointer) into location 0, and then reads the data block, placing the words in the locations specified 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.

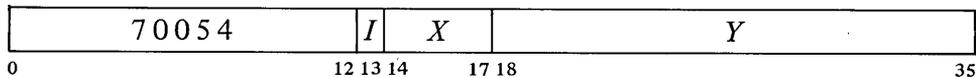
DATAI APR, Data In, Console



MACRO also recognizes the mnemonic RSW (Read Switches) as equivalent to DATAI APR,.

Read the contents of the console data switches into location *E*.

DATAO PI, Data Out, Console



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: unrellocated 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

Interrupt locations for a second processor are $140 + 2N$ and $141 + 2N$.

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 maintenance 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).

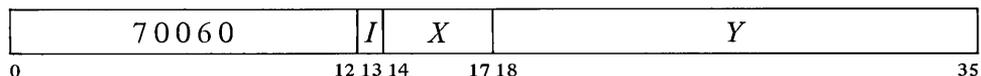
A routine dismisses the interrupt by using a JEN (JRST 12,) to return to the interrupted program (the interrupt system must be active when the JEN is given). This instruction restores the channel on which the interrupt is being held, so it can again accept requests, and interrupts can be started on it 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

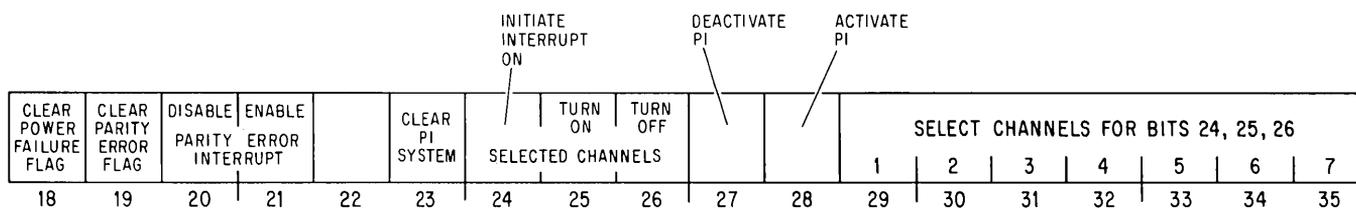
An interrupt routine must dismiss the interrupt when it returns to the interrupted program, or its channel and all channels of lower priority will be disabled, and the processor will treat the new program as a continuation of the interrupt routine.

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



Perform the functions specified by *E* as shown (a 1 in a bit produces the indicated function, a 0 has no effect).



Notes.

- 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.

Bits 18–21 are actually for processor conditions [§2.14].

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 μ s for FDVL, but a ridiculously long shift could take over 35 μ s.

Special Considerations. On a return to an interrupted program, the processor 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.

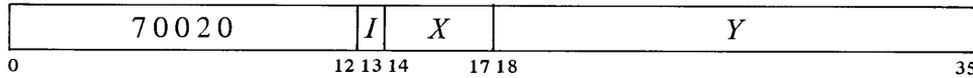
- ◆ No requests can be accepted, not even on higher priority channels, while a break is starting. Therefore do not use lengthy effective address calculations in interrupt instructions.
- ◆ The interrupt instruction that calls the routine must save PC if there is to be a return to the interrupted program. 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 situation 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

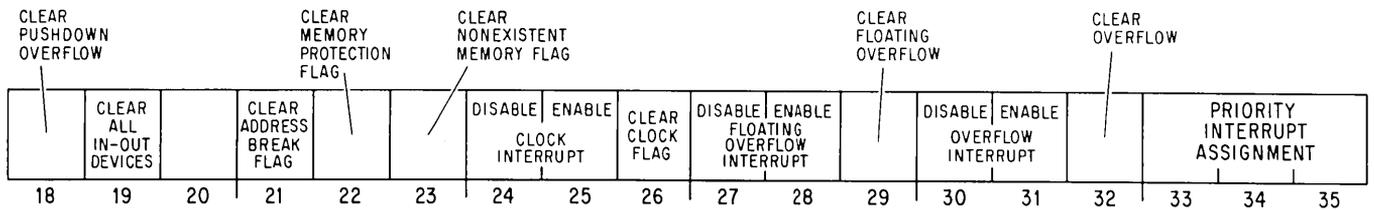
There are a number of internal conditions that can signal the program by requesting an interrupt on a channel assigned to the processor. Flags for

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



Perform the functions specified by *E* as shown (a 1 in a bit produces the indicated function, a 0 has no effect).

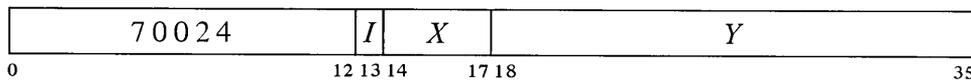


Notes.

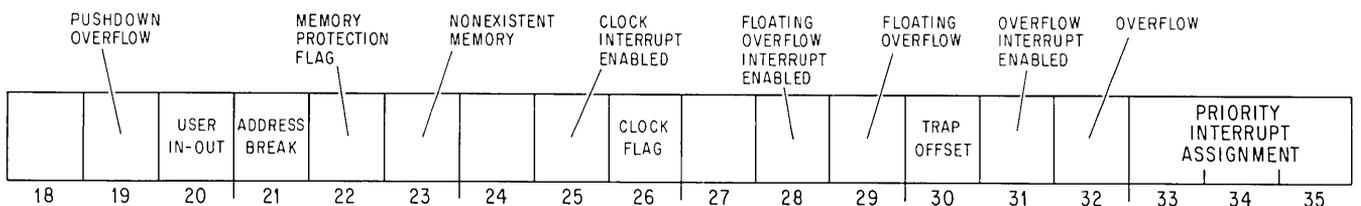
Enabling a particular flag to interrupt means that henceforth the setting of the flag will request an interrupt on the channel assigned (by bits 33-35) to the processor. Disabling prevents the flag from triggering a request.

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



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).



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, the restrictions on user instructions do not apply [§2.15].
- 21 Address Break – while the console address break switch was on, the processor requested access to the memory location specified by the address switches and the memory reference was for the purpose 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 contained 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.
- 22 Memory Protection – a user program attempted to access a memory location outside of its assigned area and the user instruction was terminated at that time. The setting of this flag requests an interrupt, at which time PC points either to the instruction that caused the violation or the one following it.
- 23 Nonexistent Memory – the processor attempted to access a memory that did not respond within 100 μ s. The setting of this flag requests an interrupt, at which time PC points either to the instruction containing 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 overflow occurred.
- 30 Trap Offset – the processor is using locations 140–161 for unimplemented operation traps and interrupt locations.
- 32 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 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.

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.

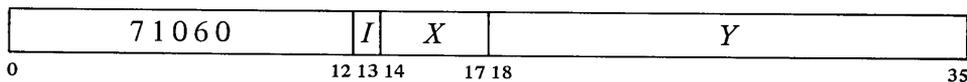
The punch supplies output in the form of 8-channel perforated paper tape in either of two modes. In alphanumeric mode, 8-bit characters are processed; in binary mode, 6-bit characters. The information punched in the tape can be brought into memory by the tape reader, which handles characters in the same 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.

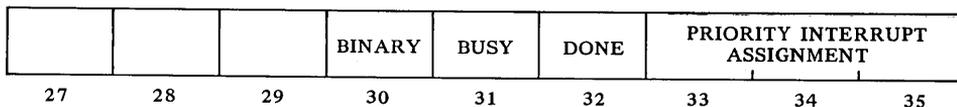
3.1 PAPER TAPE READER

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



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).



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 uncoiled 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).

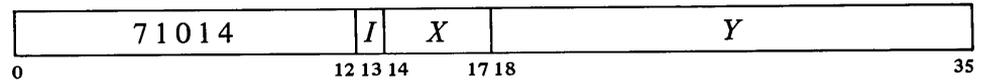
The program below is the RIM10B Loader, which is brought into the accumulators in readin mode, and then continues to read any number of blocks of binary data from the same tape. The tape is formatted as a series of blocks separated by a half-dozen lines of blank tape (tape with only feed holes punched). 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

```


DATAO PTP, Data Out, Paper Tape Punch

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

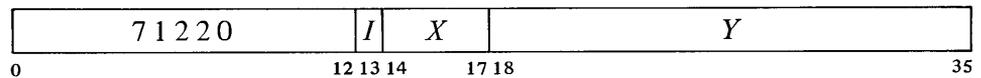
Two teletypewriter models are regularly available with the PDP-10 for use at the console: the KSR 35, which is capable of speeds up to ten characters per second, and the KSR 37, which can handle up to fifteen characters per second. The program can type out characters and can read in the characters produced 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

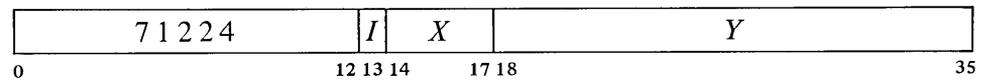


Set up the teletype control register according to bits 24–35 of the effective conditions *E* as shown (a 1 in bit 24 sets Test, a 0 clears it; all other flag functions are produced by 1s, 0s have no effect).

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.

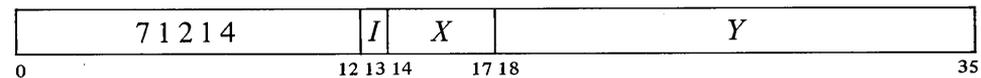
CONI TTY, Conditions In, Teletype



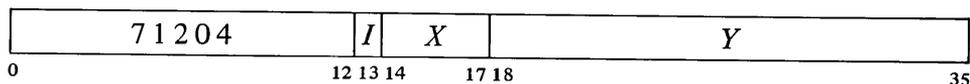
Read the status of the teletype into bits 24 and 29–35 of location *E* as shown.

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



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

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.

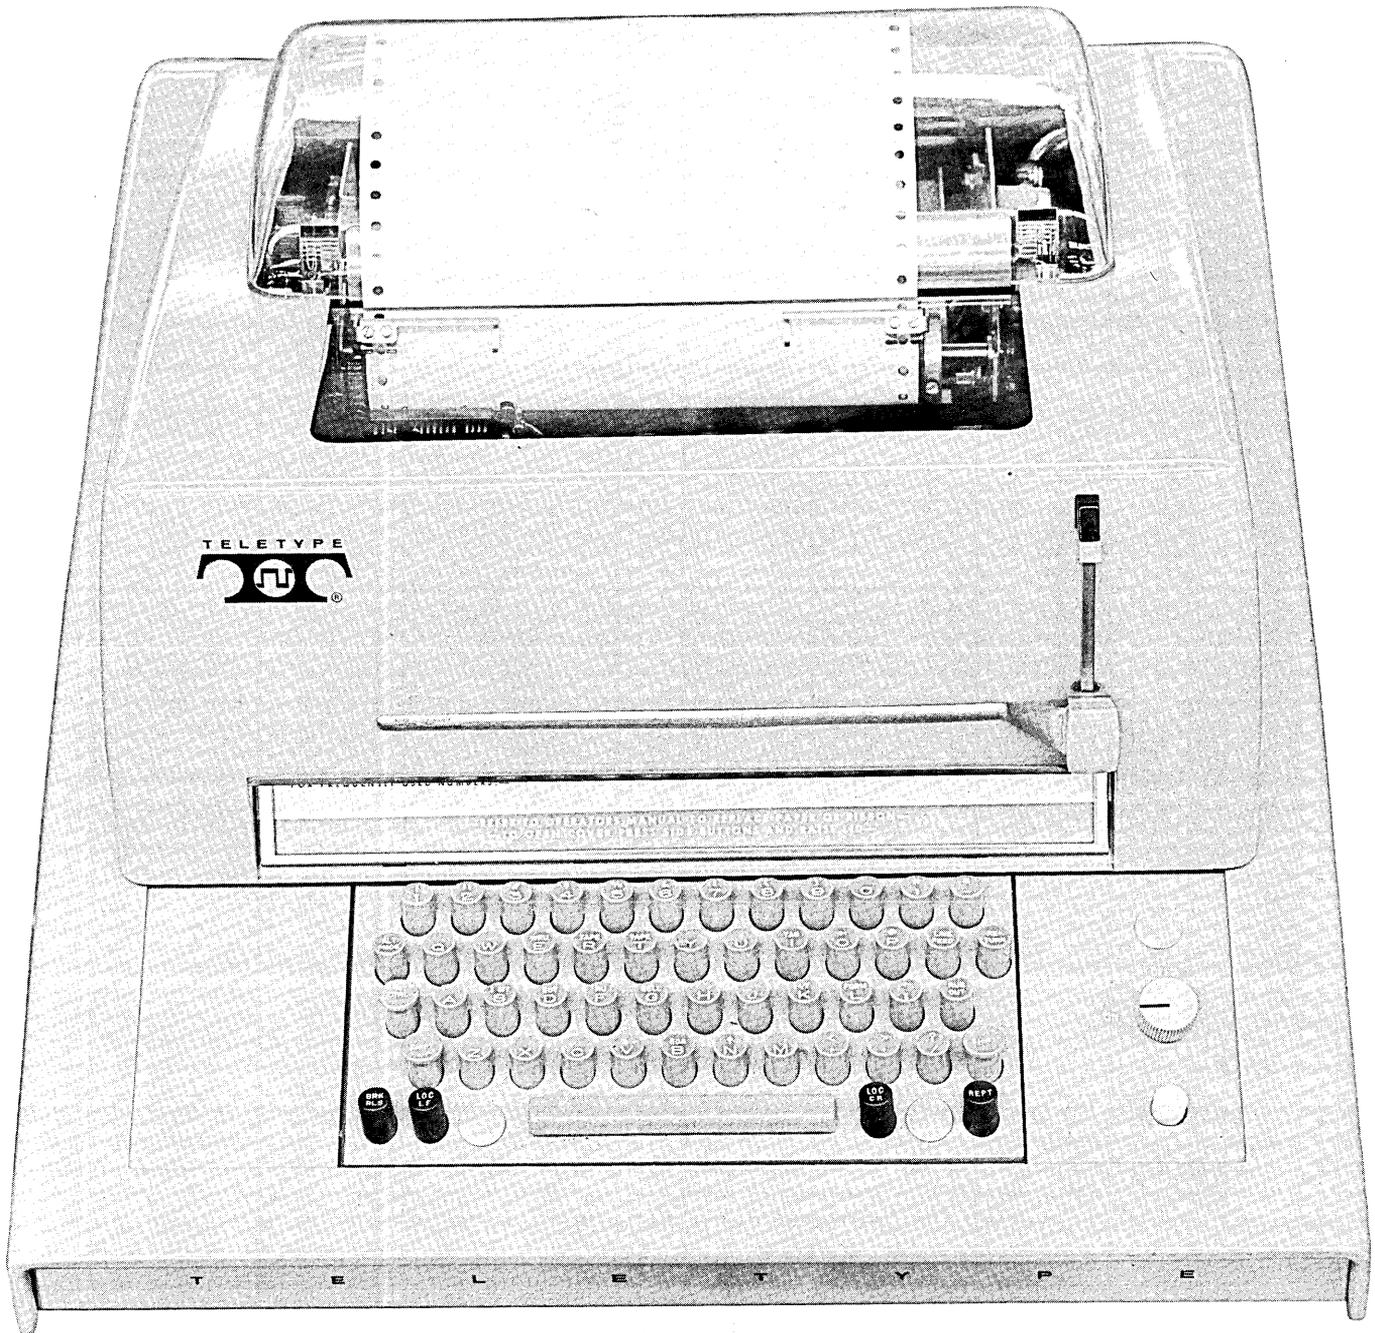
Timing. The Model 35 can type up to ten characters per second. After Output Done is set, the program has 9.09 ms to give a DATAO to keep typing at the maximum rate. After Input Done is set, the character is available for retrieval by a DATAI for 22.73 ms before another key strike can destroy it.

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.

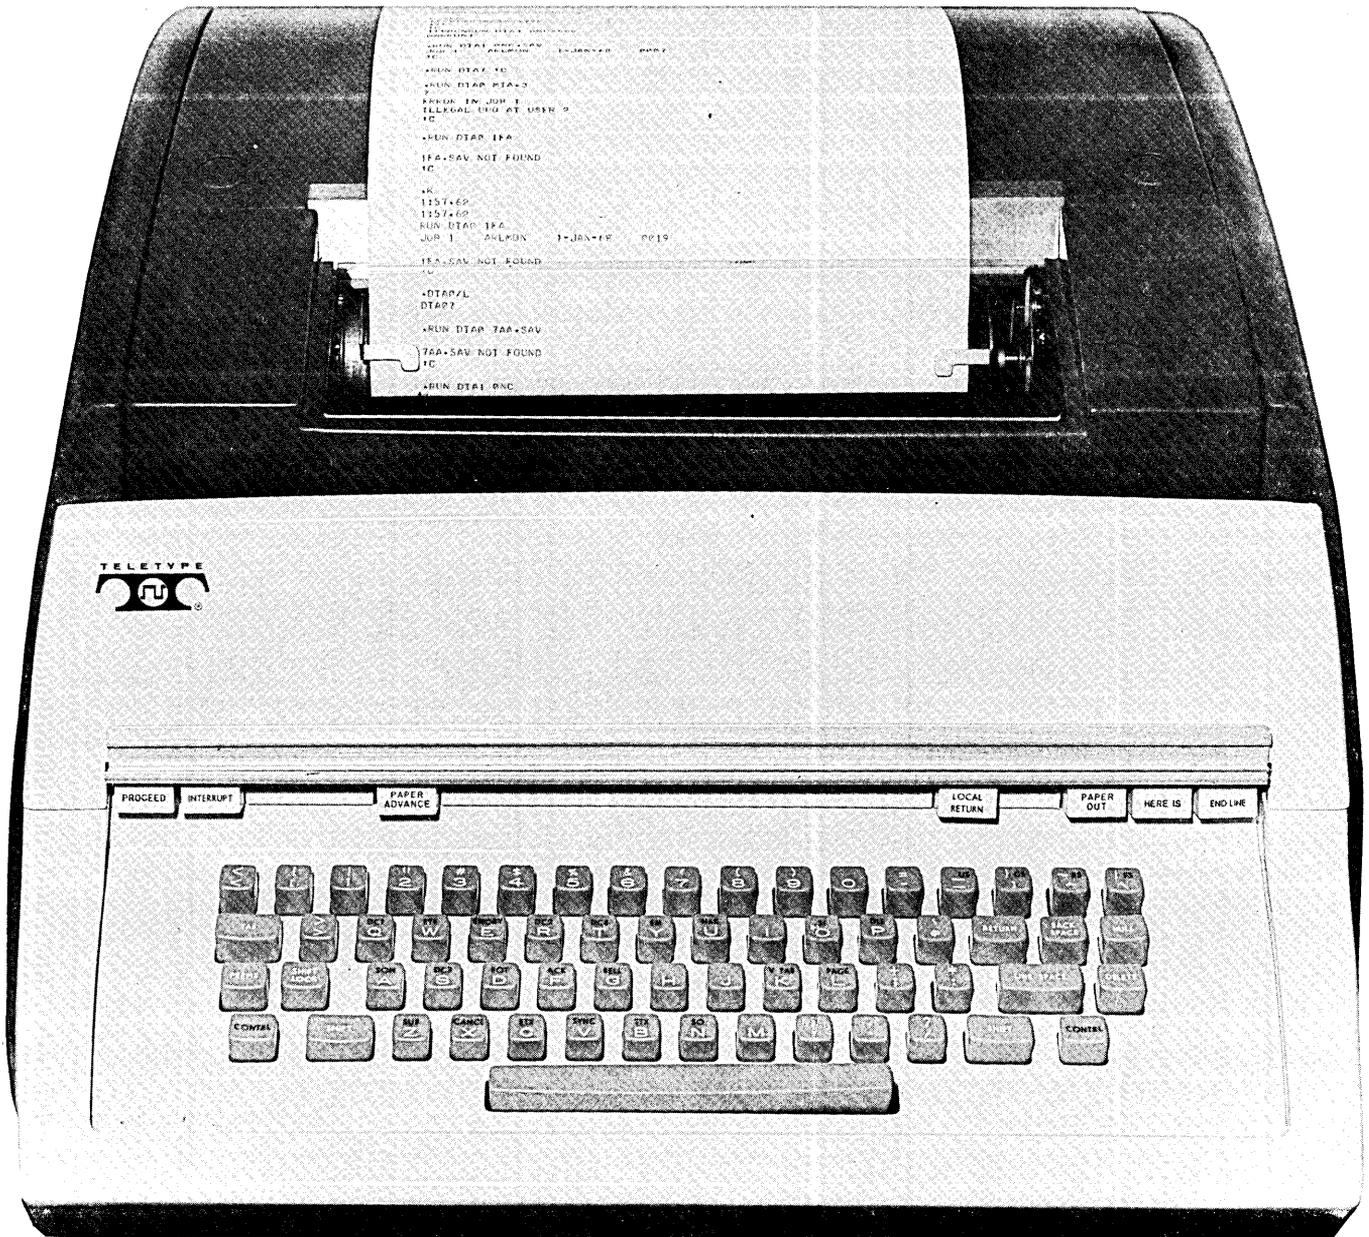
The keyboard resembles that of a standard typewriter. Codes for printable characters on the upper parts of the key tops on the 35 are transmitted by using the shift key; most control codes require use of the control key. Those familiar with the 35 who are using the 37 for the first time should take a 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.

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 UO 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.

INSTRUCTION MNEMONICS

NUMERIC LISTING

000	ILLEGAL	132	FSC	206	MOVSM
001	} USER UUO'S	133	IBP	207	MOVSS
.		134	ILDB	210	MOVN
.		135	LDB	211	MOVNI
037		136	IDPB	212	MOVNM
040	*CALL	137	DPB	213	MOVNS
041	*INIT	140	FAD	214	MOVMM
042	} RESERVED FOR SPECIAL MONITORS	141	FADL	215	MOVMI
043		142	FADM	216	MOVMM
044		143	FADB	217	MOVMS
045		144	FADR	220	IMUL
046		145	FADRI ▲	221	IMULI
047	*CALLI	146	FADRM	222	IMULM
050	*OPEN	147	FADRB	223	IMULB
051	} RESERVED FOR DEC	150	FSB	224	MUL
052		151	FSBL	225	MULI
053		152	FSBM	226	MULM
054		153	FSBB	227	MULB
055		*RENAME	154	FSBR	230
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	AOBN
100	} UNASSIGNED CODES	200	MOVE	254	JRST
.		201	MOVEI	25410	JRSTF
.		202	MOVEM	25420	HALT
.		203	MOVES	25450	JEN
127		204	MOVSI	255	JFCL
130	UFA	205		25504	JFOV
131	DFN				

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

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	†DSK	170
ADDM	272	CAIG	307	†DTC	320
AND	404	CAIGE	305	†DTS	324
ANDB	407	CAIL	301	*ENTER	077
ANDCA	410	CAILE	303	EQV	444
ANDCAB	413	CAIN	306	EQVB	447
ANDCAI	411	*CALL	040	EQVI	445
ANDCAM	412	*CALLI	047	EQVM	446
ANDCB	440	CAM	310	EXCH	250
ANDCBB	443	CAMA	314	FAD	140
ANDCBI	441	CAME	312	FADB	143
ANDCBM	442	CAMG	317	FADL	141
ANDCM	420	CAMGE	315	FADM	142
ANDCMB	423	CAML	311	FADR	144
ANDCMI	421	CAMLE	313	FADRB	147
ANDCMM	422	CAMN	316	FADRI	145
ANDI	405	†CCI	014	FADRM	146
ANDM	406	†CDP	110	FDV	170
AOBJN	253	†CDR	114	FDVB	173
AOBJP	252	CLEAR	400	FDVL	171
AOJ	340	CLEARB	403	FDVM	172
AOJA	344	CLEARI	401	FDVR	174
AOJE	342	CLEARM	402	FDVRB	177
AOJG	347	*CLOSE	070	FDVRI	175
AOJGE	345	CONI	70024	FDVRM	176
AOJL	341	CONO	70020	FMP	160
AOJLE	343	CONSO	70034	FMPB	163
AOJN	346	CONSZ	70030	FMPL	161
AOS	350	†CPA	000	FMPM	162
AOSA	354	†CR	150	FMPR	164
AOSE	352	DATAI	70004	FMPRB	167
AOSG	357	DATAO	70014	FMPRI	165
AOSGE	355	†DC	200	FMPRM	166
AOSL	351	†DCSA	300	FSB	150
AOSLE	353	†DCSB	304	FSBB	153
AOSN	356	†DF	270	FSBL	151
†APR	000	DFN	131	FSBM	152
ASH	240	†DIS	130	FSBR	154
ASHC	244	DIV	234	FSBRB	157
BLKI	70000	DIVB	237	FSBRI	155
BLKO	70010	DIVI	235	FSBRM	156

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	*LOOKUP	076
HLLOI	521	HRROM	562	†LPT	124
HLLOM	522	HRROS	563	LSH	242
HLLOS	523	HRRS	543	LSHC	246
HLLS	503	HRRZ	550	†MDF	260
HLLZ	510	HRRZI	551	MOVE	200
HLLZI	511	HRRZM	552	MOVEI	201
HLLZM	512	HRRZS	553	MOVEM	202
HLLZS	513	IBP	133	MOVES	203
HLR	544	IDIV	230	MOVMM	214
HLRE	574	IDIVB	233	MOVMI	215
HLREI	575	IDIVI	231	MOVMM	216
HLREM	576	IDIVM	232	MOVMS	217
HLRES	577	IDPB	136	MOVN	210
HLRI	545	ILDB	134	MOVNI	211
HLRM	546	IMUL	220	MOVNM	212
HLRO	564	IMULB	223	MOVNS	213
HLROI	565	IMULI	221	MOVS	204
HLROM	566	IMULM	222	MOVSI	205
HLROS	567	*IN	056	MOVSM	206
HLRS	547	*INBUF	064	MOVSS	207
HLRZ	554	*INIT	041	*MTAPE	072
HLRZI	555	*INPUT	066	†MTC	220
HLRZM	556	IOR	434	†MTM	230
HLRZS	557	IORB	437	†MTS	224
HRL	504	IORI	435	MUL	224
HRLE	534	IORM	436	MULB	227
HRLEI	535	JCRY	25530	MULI	225
HRLEM	536	JCRY0	25520	MULM	226
HRLES	537	JCRY1	25510	*OPEN	050
HRLI	505	JEN	25460	OR	434
HRLM	506	JFCL	255	ORB	437
HRLO	524	JFFO	243	ORCA	454
HRLOI	525	JFOV	25504	ORCAB	457
HRLOM	526	JOV	25540	ORCAI	455
HRLOS	527	JRA	267	ORCAM	456
HRLS	507	JRST	254	ORCB	470
HRLZ	514	JRSTF	25410	ORCBB	473
HRLZI	515	JSA	266	ORCBI	471

ORCBM	472	SKIPA	334	TLCN	647
ORCM	464	SKIPE	332	TLN	601
ORCMB	467	SKIPG	337	TLNA	605
ORCM I	465	SKIPGE	335	TLNE	603
ORCMM	466	SKIPL	331	TLNN	607
ORI	435	SKIPL E	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
SETZM	402	TLCA	645	UFA	130
SKIP	330	TLCE	643	*UGETF	073

ALPHABETIC LISTING

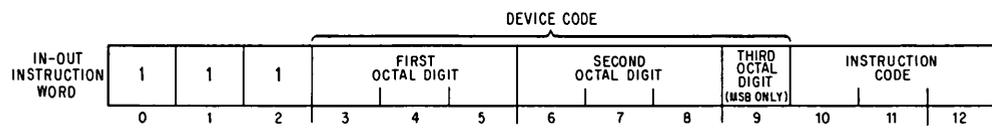
A9

*USETI 074
*USETO 075
†UTC 210

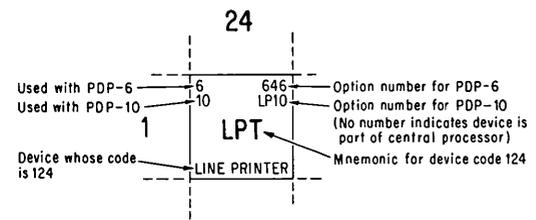
†UTS 214
XCT 256
XOR 430

XORB 433
XORI 431
XORM 432

		SECOND AND THIRD OCTAL DIGITS															
		00	04	10	14	20	24	30	34	40	44	50	54	60	64	70	74
DEC STANDARD DEVICES	FIRST OCTAL DIGIT	6,10	6,10	6	167	10	DA10	10	DA10	10	AD10	10	AD10				
	0	APR CENTRAL PROCESSOR	PI PRIORITY INTERRUPT			CCI PDP-8,9 INTERFACE	CCI2 PDP-8,9 INTERFACE	ADC ANALOG-DIGITAL CONVERTER	ADC2 ANALOG-DIGITAL CONVERTER								
	1	PTP PAPER TAPE PUNCH	PTR PAPER TAPE READER	CDP CARD PUNCH	CDR CARD READER	TTY CONSOLE TELETYPE	LPT LINE PRINTER	DIS DISPLAY	DIS2 DISPLAY	PLT PLOTTER	PLT2 PLOTTER	CR CARD READER	CR2 CARD READER	PDP-7,8 INTERFACE		DSK SMALL DISK	DSK2 SMALL DISK
	2	DC DATA CONTROL	DC2 DATA CONTROL	UTC DECTAPE	UTS DECTAPE	MTC MAGNETIC TAPE	MTS MAGNETIC TAPE	MTM MAGNETIC TAPE		DLS DATA LINE SCANNER	DLS2 DATA LINE SCANNER			MDF MASS DISK FILE	MDF2 MASS DISK FILE	DF DISK FILE	
	3	DCSA DATA COMMUNICATION	DCSB			DTC DECTAPE	DTS DECTAPE	DTC2 DECTAPE	DTS2 DECTAPE	TMC MAGNETIC TAPE	TMS MAGNETIC TAPE	TMC2 MAGNETIC TAPE	TMS2 MAGNETIC TAPE				
	4																
	5																
	6																
7																	



DEVICE MNEMONICS



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.

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	∅	
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

INPUT-OUTPUT CODES

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.

CARD CODES

▲ Character	PDP-10 ASCII	DEC 029	DEC 026	Character	PDP-10 ASCII	DEC 029	DEC 026
<i>Space</i>	040	<i>None</i>	<i>None</i>	@	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
)	051	11 8 5	12 8 4	▲	I	111	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.

CARD CODES

B7

Column Punch	Character	Octal	Column Punch	Character	Octal
<i>None</i>	<i>Space</i>	0000	12 9	I	4001
0	0	1000	11 1	J	2400
1	1	0400	11 2	K	2200
2	2	0200	11 3	L	2100
3	3	0100	11 4	M	2040
4	4	0040	11 5	N	2020
5	5	0020	11 6	O	2010
6	6	0010	11 7	P	2004
7	7	0004	11 8	Q	2002
8	8	0002	11 9	R	2001
9	9	0001	0 1	/	1400
12 1	A	4400	0 2	S	1200
12 2	B	4200	0 3	T	1100
12 3	C	4100	0 4	U	1040
12 4	D	4040	0 5	V	1020
12 5	E	4020	0 6	W	1010
12 6	F	4010	0 7	X	1004
12 7	G	4004	0 8	Y	1002
12 8	H	4002	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

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 2			<i>See note</i>	;]	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				<i>End of File</i>	<i>End of File</i>	7400
12 0 2 4 6 8				<i>Mode Switch</i>	<i>Mode Switch</i>	5252
7 9				<i>Binary</i>	<i>Binary</i>	xx05

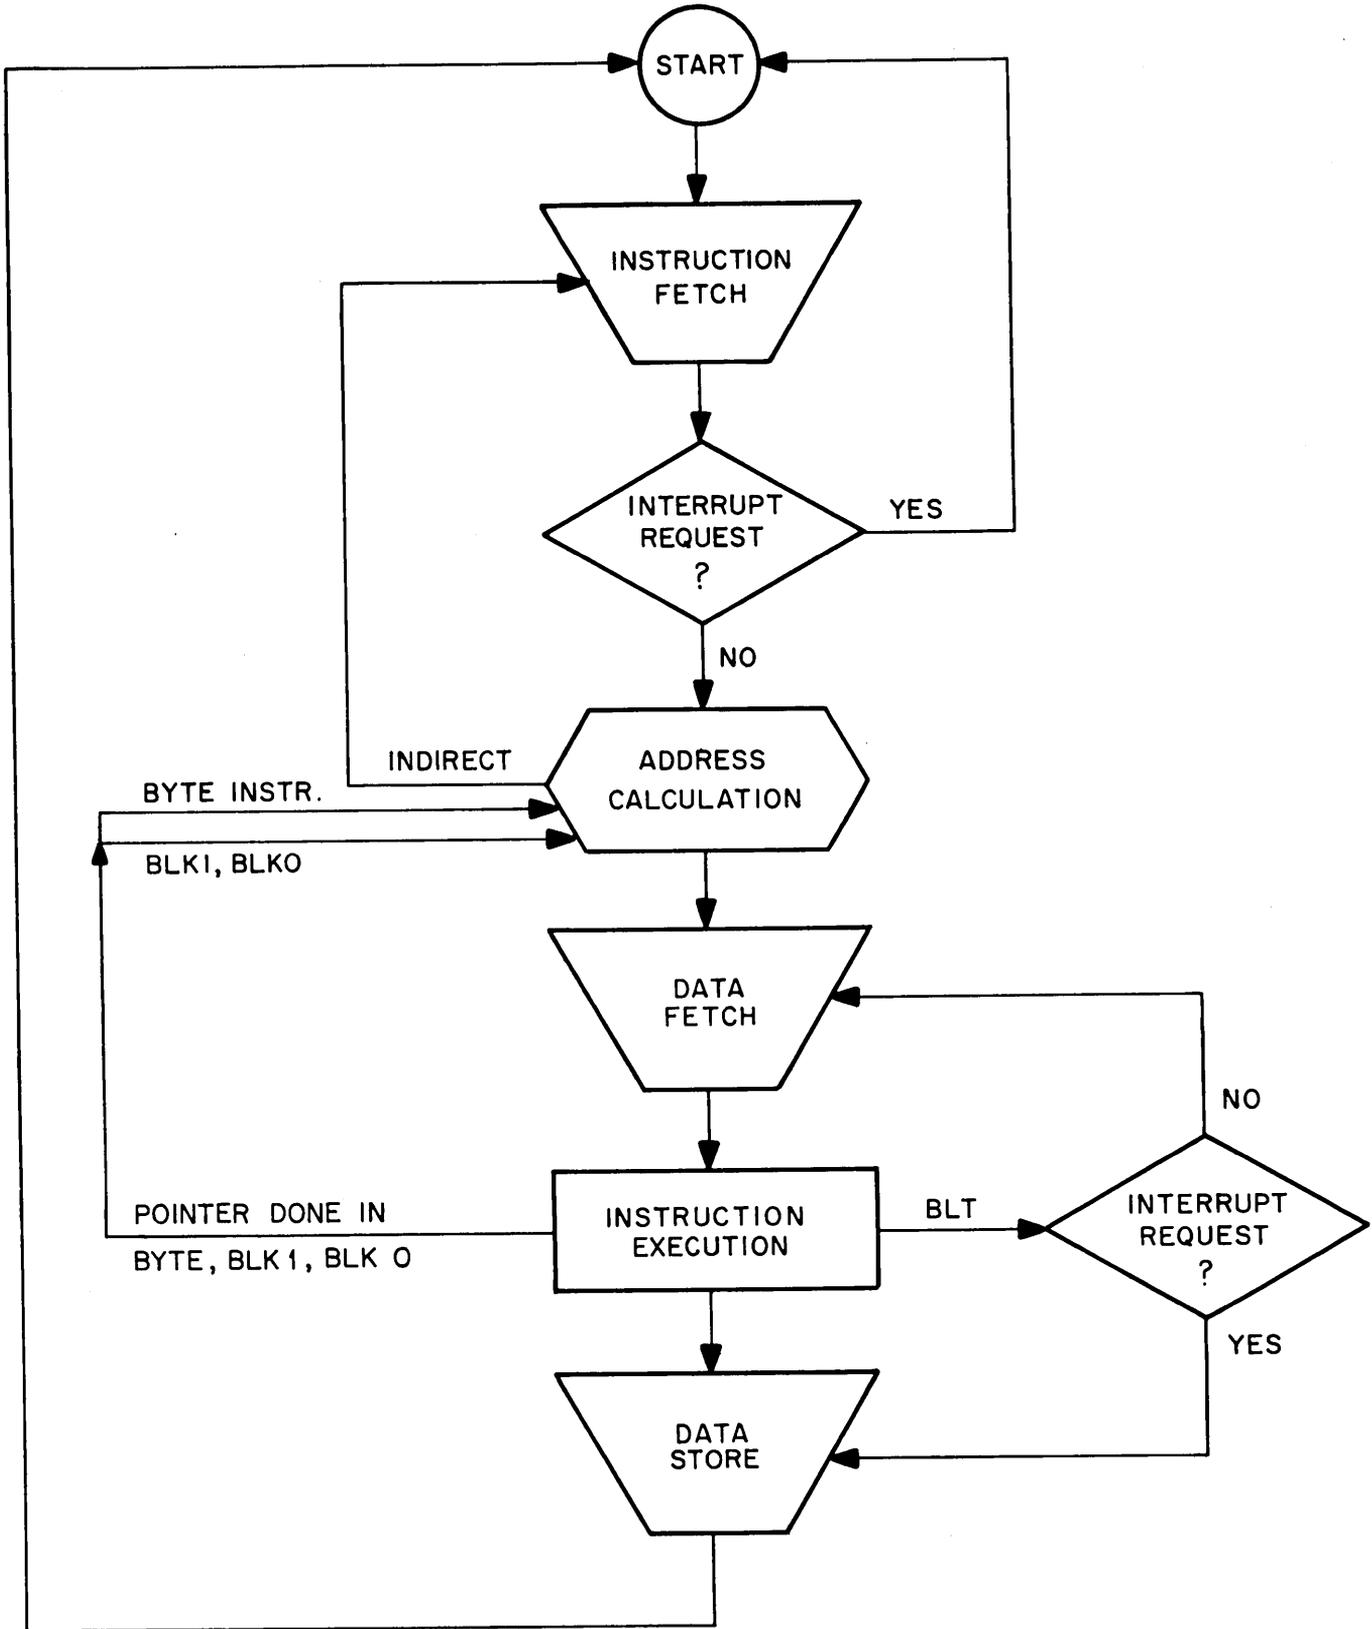
NOTE: There is a single key for the 0 8 2 punch on the 029 but printing is suppressed.

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

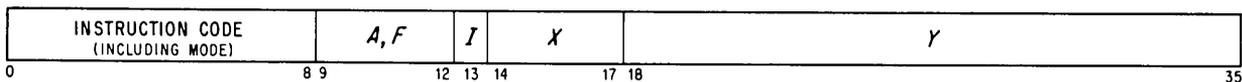
MISCELLANY

Instruction Flow Simplified	C2
Word Formats	C3
Instruction Timing Flow Chart	C4
In-out Device Bit Assignments	C6
Indicator Panels	C8
Powers of Two	C10

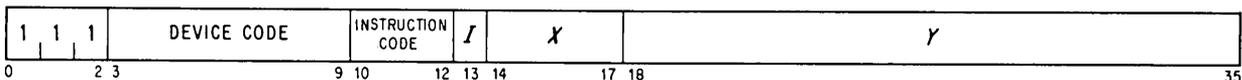


INSTRUCTION FLOW SIMPLIFIED

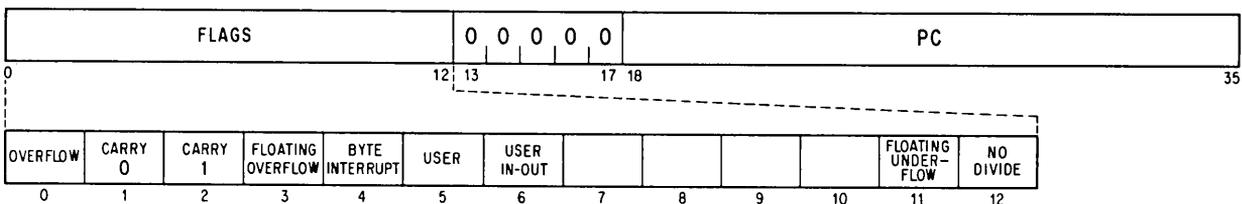
BASIC INSTRUCTIONS



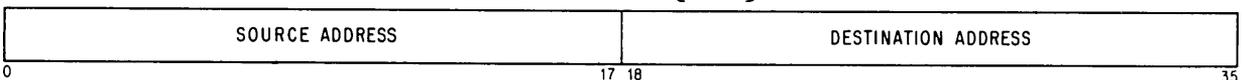
IN-OUT INSTRUCTIONS



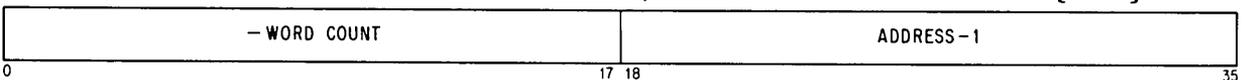
PC WORD



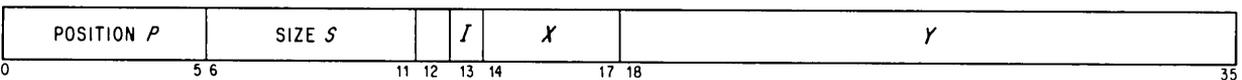
BLT POINTER {XWD}



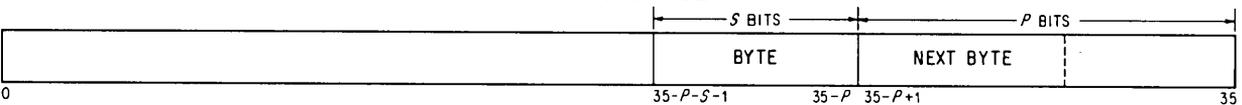
BLKI/BLKO POINTER, PUSHDOWN POINTER, DATA CHANNEL CONTROL WORD {IOWD}



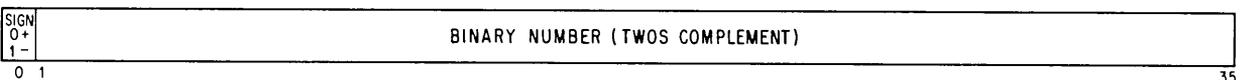
BYTE POINTER



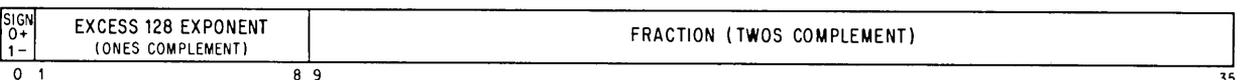
BYTE STORAGE



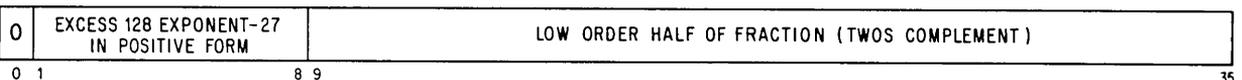
FIXED POINT OPERANDS

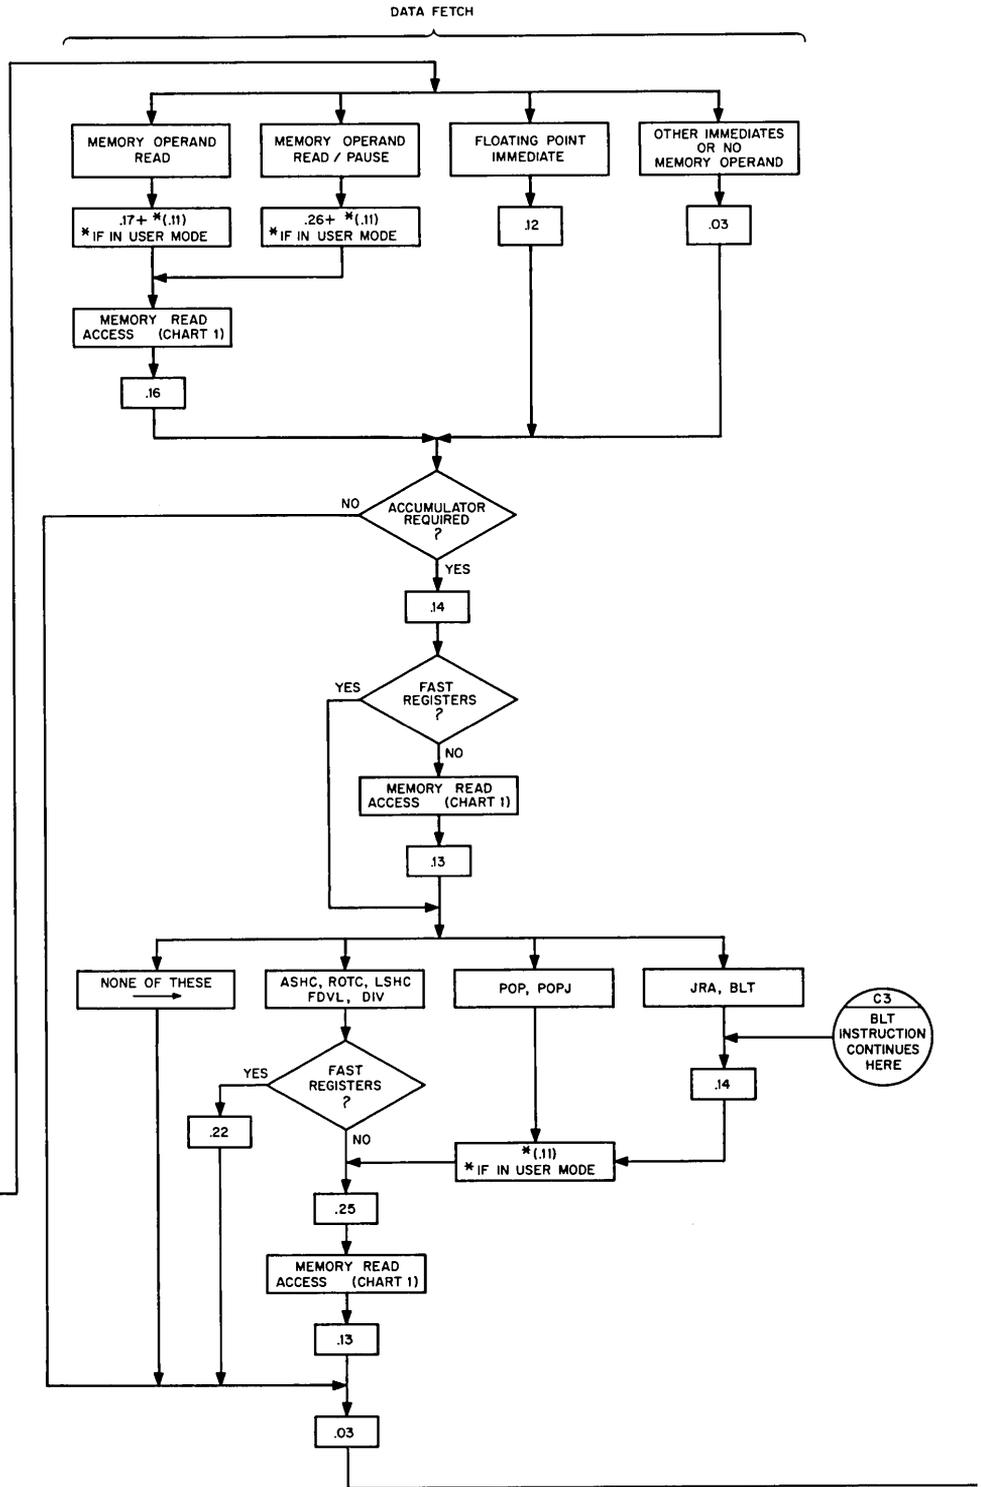
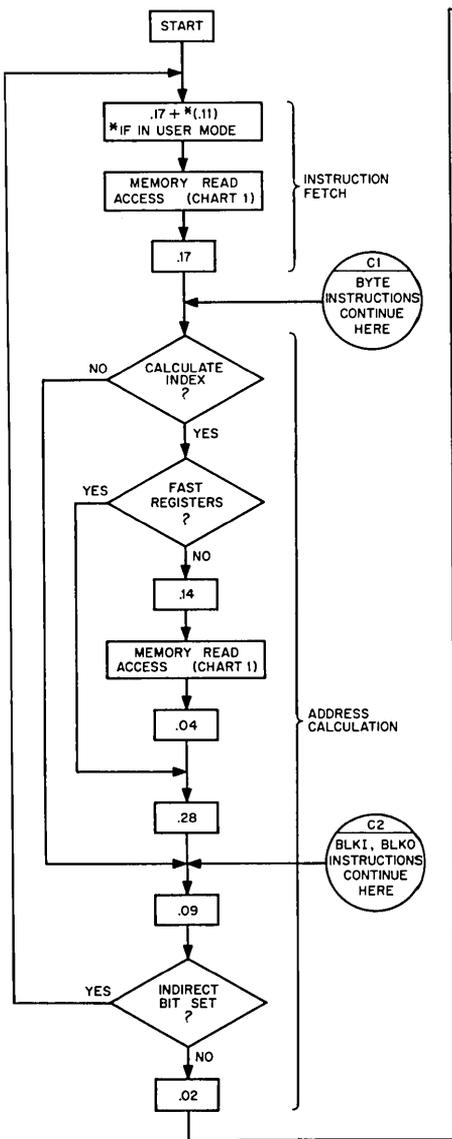


FLOATING POINT OPERANDS



LOW ORDER WORD IN DOUBLE LENGTH FLOATING POINT OPERANDS





INSTRUCTION TIMING FLOW CHART

INSTRUCTION EXECUTION

Boolean (except ANDCA, ANDCB, ORCA, ORCB), Half Words (except HLR, HLRI, HRL, HRLI), MOVE, MOVS, EXCH, JFCL, JRST, JSP, XCT, UUC	.27	
ANDCA, ANDCB, ORCA, ORCB, HLR, HLRI, HRL, HRLI, JSR, JSA, JRA, Test class	.62	
MOVN, MOVN, 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	
CONO, CONI, CONSO, CONSZ, DATAO, DATAI	.12	
CONO, CONI, DATAO, DATAI	+2.69	
CONSO, CONSZ	+2.90	
BLKO, BLKI	.60	Then turn into DATAO, DATAI and go to C2

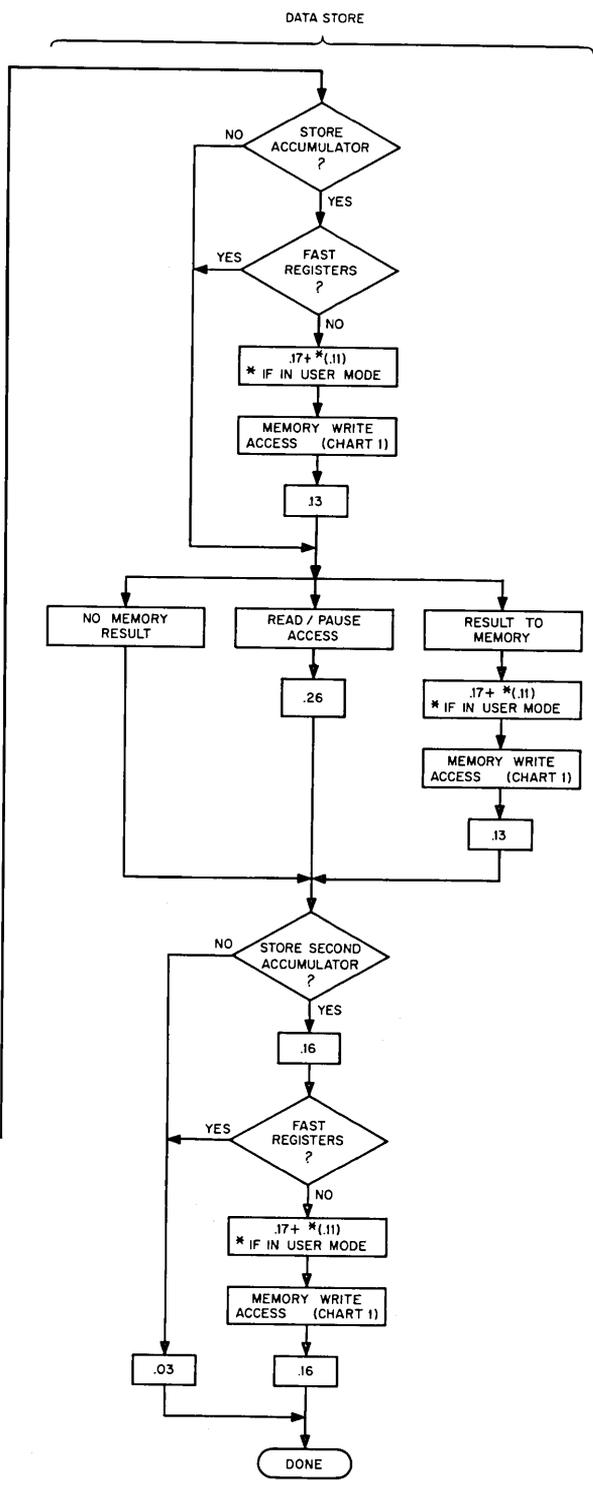


CHART 1

MEMORY	MA10	MB10	MB10	KM10 †
PROCESSORS	SINGLE OR MULTI	SINGLE	MULTI	SINGLE (BUILT IN)
CYCLE	1.00	1.65	1.65	—
READ ACCESS*	.55	.60	.70	.21
WRITE ACCESS*	.20	.20	.30	.21

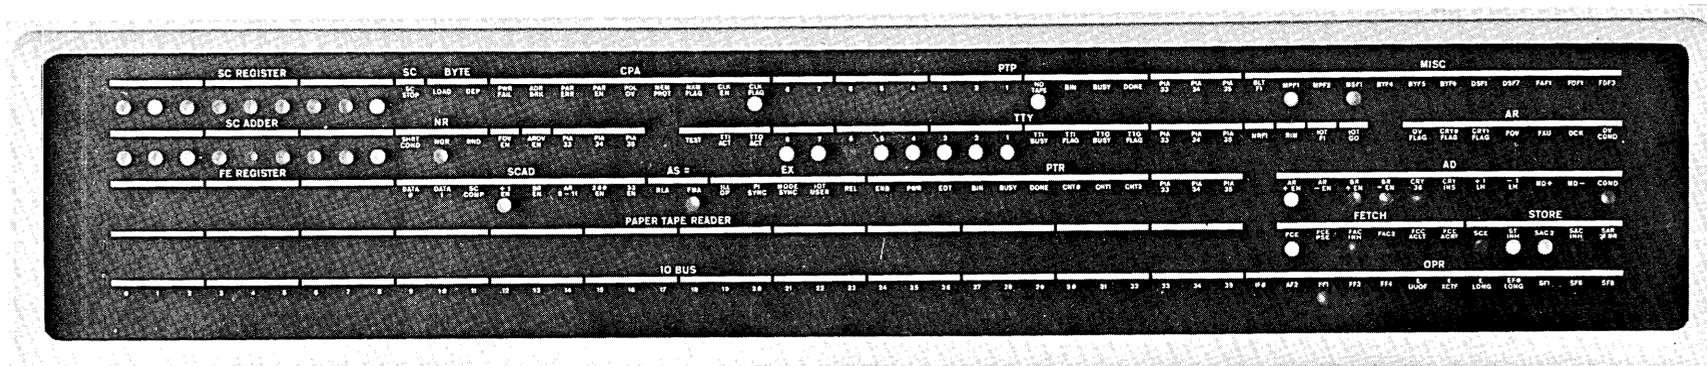
* INCLUDING 20 FT OF CABLE DELAY
 † FAST REGISTERS
 ALL TIMES ARE ±5%

IN-OUT DEVICE BIT ASSIGNMENTS

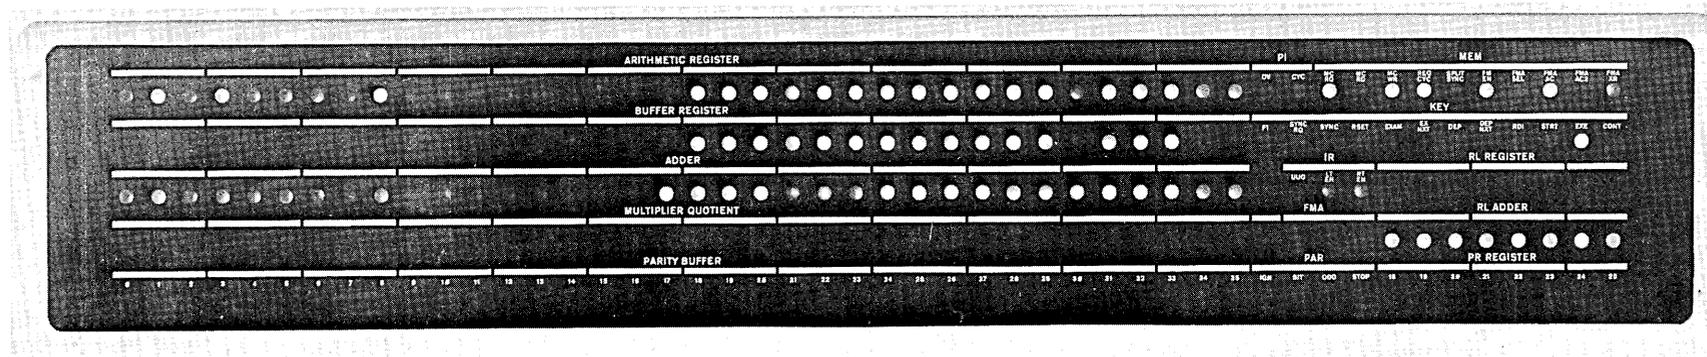
DEVICE	CODE	FUNCTION	0/18	1/19	2/20	3/21	4/22	5/23	6/24	7/25	8/26	9/27	10/28	11/29	12/30	13/31	14/32	15/33	16/34	17/35			
APR	000	CONO	CLR PDL OV	I/O RESET		CLR ADR BREAK	CLR MEM PROT	CLR NXM	CLR CLOCK ENABLE	SET CLOCK ENABLE	CLR CLOCK FLAG	CLR FOV ENABLE	SET FOV ENABLE	CLR FOV FLAG	CLR AROV ENABLE	SET AROV ENABLE	CLR AROV FLAG			PIA			
		CONI		PDL OV FLAG	TOT USER FLAG	ADR BREAK FLAG	MEM PROT FLAG	NXM FLAG		CLOCK ENABLE	CLOCK FLAG		FOV ENABLE	FOV FLAG	TRAP OFFSET	AROV ENABLE	AROV FLAG			PIA			
	DAT0	PROTECTION REGISTER (LH 0-7) RELOCATION REGISTER (RH 18-25)																					
	DATI	← 36 DATA SWITCHES →																					
PI	004	CONO	CLEAR POWER FAIL FLAG	CLEAR PARITY FLAG	CLEAR PARITY ENABLE	SET PARITY ENABLE		CLEAR PI SYSTEM	REQUEST INTERRUPT CHANNELS 1-7	TURN ON CHANNELS	TURN OFF CHANNELS	TURN OFF PIACTIVE	TURN ON PIACTIVE	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								
	DAT0	← 36 BITS TO THE MEMORY INDICATORS →																					
ADC A-D CONVERTER (AD10)	024																						
PTP	100	CONO													BINARY	BUSY	DONE FLAG	PIA					
		CONI													OUT OF TAPE	BINARY	BUSY	DONE FLAG	PIA				
		DAT0													HOLE 8 UNLESS BINARY	HOLE 7 UNLESS BINARY	HOLE 6	HOLE 5	HOLE 4	HOLE 3	HOLE 2	HOLE 1	
PTR	104	CONO													BINARY	BUSY	DONE FLAG	PIA					
		CONI													TAPE FLAG			BINARY	BUSY	DONE FLAG	PIA		
		DATI	← 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				
		DAT0	8 BIT CHARACTER TO TELETYPE																				
		DATI	8 BIT CHARACTER FROM KEYBOARD																				
LPT (LP10)	124	CONO								CLEAR PRINTER			BUSY	DONE FLAG	ERROR PIA	DONE PIA							
		CONI								128 CHAR	96 CHAR		ERROR	BUSY	DONE FLAG	ERROR PIA	DONE PIA						
		DAT0	FIRST CHARACTER					SECOND CHARACTER					THIRD CHARACTER										
		DATI	CHARACTER			FOURTH CHARACTER					FIFTH CHARACTER												
PLT (XY10)	140	CONO													BUSY	DONE FLAG	PIA						
		CONI													POWER ON	BUSY	DONE FLAG	PIA					
		DAT0													PEN RAISE	PEN LOWER	-X DRUM UP	+X DRUM DOWN	+Y CARRIAGE LEFT	-Y CARRIAGE RIGHT			
CR (CR10)	150	CONO								CLEAR READER	OFFSET CARD		READ CARD	TROUBLE	CLEAR DATA MISSED	READY ENABLE	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	READER STOP	CARD IN READER	HOPPER STACKER	READING CARD	READER TROUBLE	DATA MISSED	READY TO READ	END OF FILE	END OF CARD	DATA FLAG	PIA					
		DATI								ROW 12	ROW 11	ROW 0	ROW 1	ROW 2	ROW 3	ROW 4	ROW 5	ROW 6	ROW 7	ROW 8	ROW 9		

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			
DSK (RC10)	170	CONO	SELECT SECTOR CTR		CLEAR DISK DESG ERRORS	CLR TRACK SELECT ERRORS	CLEAR DISK NOT READY	CLEAR PWR SUP FAILURE	CLR DISK DISK ERR	CLR CHAN DATA PAR ERR	CLR CHAN CONTROL PAR ERR	CLEAR NXM	CLEAR ILLEGAL WRITE	CLEAR OVER-EUN	WRITE CHAN CNTL WD	STOP CLEAR BUSY	CLEAR DONE	PIA					
		CONI LH	PROTECT LOWER DISK										BOUNDARY SWITCHES (BCD)										
		CONI RH	DATA XFER IN PROG	SEARCH ERROR	DISK DESG ERROR	TRACK SELECT ERROR	DISK NOT READY	DISK SUPPLY FAIL	DISK PARITY ERROR	CHANNEL DATA PAR ERR	CHANNEL CONTROL PAR ERR	NXM	ILLEGAL WRITE	OVER-RUN	CHAN CNTL WD WRITTEN	BUSY	DONE	PIA					
		DATAO LH	DISK SELECT		100'S				10'S				TRACK (BCD)				SECTOR (BCD)						
		DATAO RH	INITIAL PARITY								DISABLE PAR ERR STOP	DISABLE MEMORY PAR ERR STOP	WRITE	INITIAL CHANNEL CONTROL WORD ADDRESS								WRITE EVEN PARITY	
DATAI RH	PARITY REGISTER								SECTOR CTR SELECTED				SECTOR COUNTER (BCD)										
DTC (TD10)	320	CONO	STOP	GO FORWARD	GO REVERSE	DELAY INHIBIT	SELECT	CLEAR SELECT # UNIT NR	TRANSPORT NUMBER		FUNCTION NUMBER		DATA PIA		FLAGS PIA								
		CONI	STOP	GO FORWARD	GO REVERSE		SELECT	DESELECT	TRANSPORT NUMBER		FUNCTION NUMBER		DATA PIA		FLAGS PIA								
		DATAO	36 BIT WORD																				
DATAI	36 BIT WORD																						
DTS	324	CONO	PARITY ERR ENABLE	DATA MISSED ENABLE	JOB DONE ENABLE	ILLEGAL OPERATION ENABLE	END ZONE ENABLE	BLOCK MISSED ENABLE	DELAY IN REVERSE	ACTIVE	UP TO SPEED	BLOCK NUMBER	REVERSE CHECK	DATA	FINAL DATA	CHECKSUM	IDLE	BLOCK NUMBER READ	STOP ALL TRANSPORT	FUNCTION STOP			
		CONI LH	PARITY ERR ENABLE	DATA MISSED ENABLE	JOB DONE ENABLE	ILLEGAL OPERATION ENABLE	END ZONE ENABLE	BLOCK MISSED ENABLE	DELAY IN REVERSE	WRITE LOCK ON	WRTM SWITCH ON	INCOMPLETE BLOCK	MARK TRACK ERROR	SELECT ERROR								FLAG PI REQUEST	DATA REQUEST
		CONI RH	PARITY ERROR	DATA MISSED	JOB DONE	ILLEGAL OPERATION	END ZONE	BLOCK MISSED	UP TO SPEED TEST												FLAG PI REQUEST	DATA REQUEST	
		DATAO RH																					
		DATAI LH																					
		DATAI RH	REGISTER																				
DLS (DC10)	240	CONO																CLEAR DLS	SET DATA MEMORY	RESET SCANNER	PIA		
		CONI																DTR DISABLED	TRANS FLAG	RCVR FLAG	PIA		
		DATAO LH	LINE NUMBER																				
		DATAI LH	LINE NUMBER																				
TMC (TM10)	340	CONO	UNIT		PARITY 0 = EVEN 1 = ODD	CORE DUMP	FUNCTION			NEXT UNIT ENAB	DENSITY		FLAG PIA		DATA PIA								
		CONI RH	UNIT		PARITY	CORE DUMP	FUNCTION			NEXT UNIT ENAB	DENSITY		FLAG PIA		DATA PIA								
		CONI LH	NEXT UNIT																				
		DATAO	36 BIT WORD																				
		DATAO 7-CHN OR CORE DUMP	(7 CHN: SIX 6 BIT CHARACTERS AS IN DATAI BELOW)																				
		DATAO 9-CHN	1 ST CHARACTER 2 ND CHARACTER 3 RD CHARACTER 4 TH CHARACTER 5 TH CHARACTER 6 TH CHARACTER 7 TH CHARACTER 8 TH CHARACTER 9 TH CHARACTER																				
DATAI 7-CHN	1 ST CHARACTER 2 ND CHARACTER 3 RD CHARACTER 4 TH CHARACTER 5 TH CHARACTER 6 TH CHARACTER 7 TH CHARACTER 8 TH CHARACTER 9 TH CHARACTER																						
DATAI 9-CHN	1 ST CHARACTER 2 ND CHARACTER 3 RD CHARACTER 4 TH CHARACTER 5 TH CHARACTER 6 TH CHARACTER 7 TH CHARACTER 8 TH CHARACTER 9 TH CHARACTER																						
DATAI 9-CHN OR CORE DUMP	1 ST CHARACTER 2 ND CHARACTER 3 RD CHARACTER 4 TH CHARACTER 5 TH CHARACTER 6 TH CHARACTER 7 TH CHARACTER 8 TH CHARACTER 9 TH CHARACTER																						
TMS	344	CONO																CLEAR # DATA PAR ERR	CLEAR # WT CW PAR ERR	WRITE # CTRL WORD	XFER # CRC	STOP	
		CONI LH																CW PAR ERR	NON-EX MEM	DATA PAR ERR	WT CW #	CHARACTER COUNTER	
		CONI RH	XPORT HUNG	REWINDING	LOAD POINT	ILLEG OP	PAR ERROR	END OF FILE	END OF TAPE	READ COMPARE ERROR	RECORD LENGTH INCREASE	DATA LATE	BAD TAPE	JOB DONE	TAPE UNIT RDY (TUE)	CHANNEL ERROR	WRITE LOCK	7-CHAN TAPE	NEXT UNIT	DATA REQ			
DATAO RH	INITIAL CHANNEL CONTROL WORD ADDRESS																						
CCI (DA10)	014	CONO	CLR SELF TEST	SET SELF TEST	CLR FM10 FULL	SET FM10 FULL	CLR FM10 EMPTY	SET FM10 EMPTY	CLR TO 10 FULL	SET TO 10 FULL	CLR TO 10 EMPTY	SET TO 10 EMPTY	PIA										
		CONI		SELF TEST (1)	FM10 FULL (1)	FM10 FULL (1)	FM10 EMPTY (1)	FM10 EMPTY (1)	TO 10 FULL (1)	TO 10 FULL (1)	TO 10 EMPTY (1)	TO 10 EMPTY (1)	PIA										
		DATAI	1ST PDP 8 BYTE (A)						2ND PDP 8 BYTE (B)						3RD PDP 8 BYTE (C)								
		DATAO PDP POPS	36 BITS																				
DATAI PDP POPS	36 BITS																						
DATAO PDP POPS	1ST PDP 9 BYTE (A)																						
DATAI PDP POPS	2ND PDP 9 BYTE (B)																						

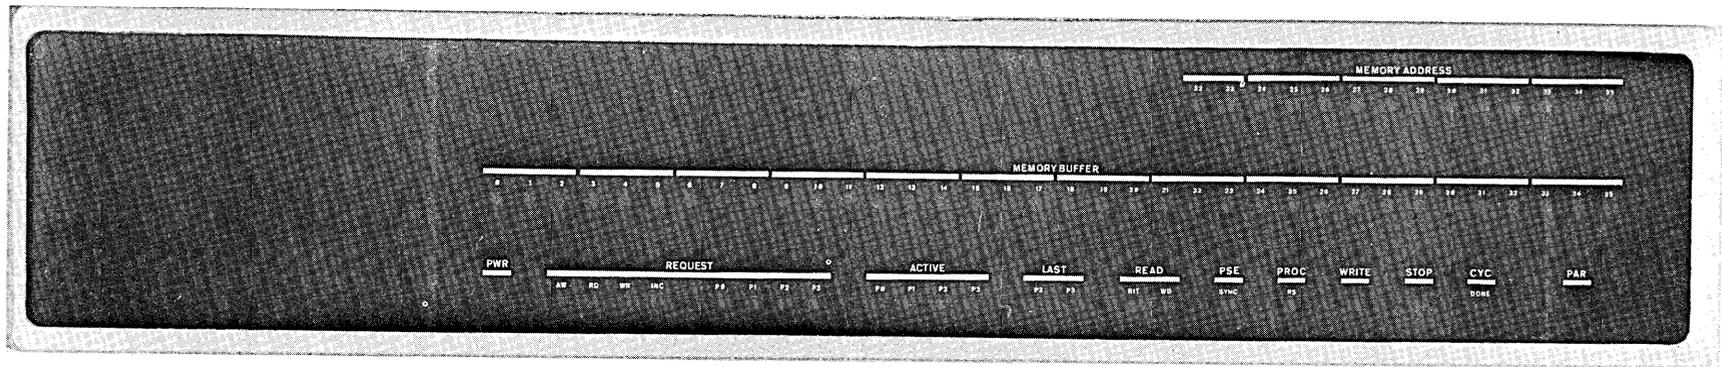
IN-OUT DEVICE BIT ASSIGNMENTS



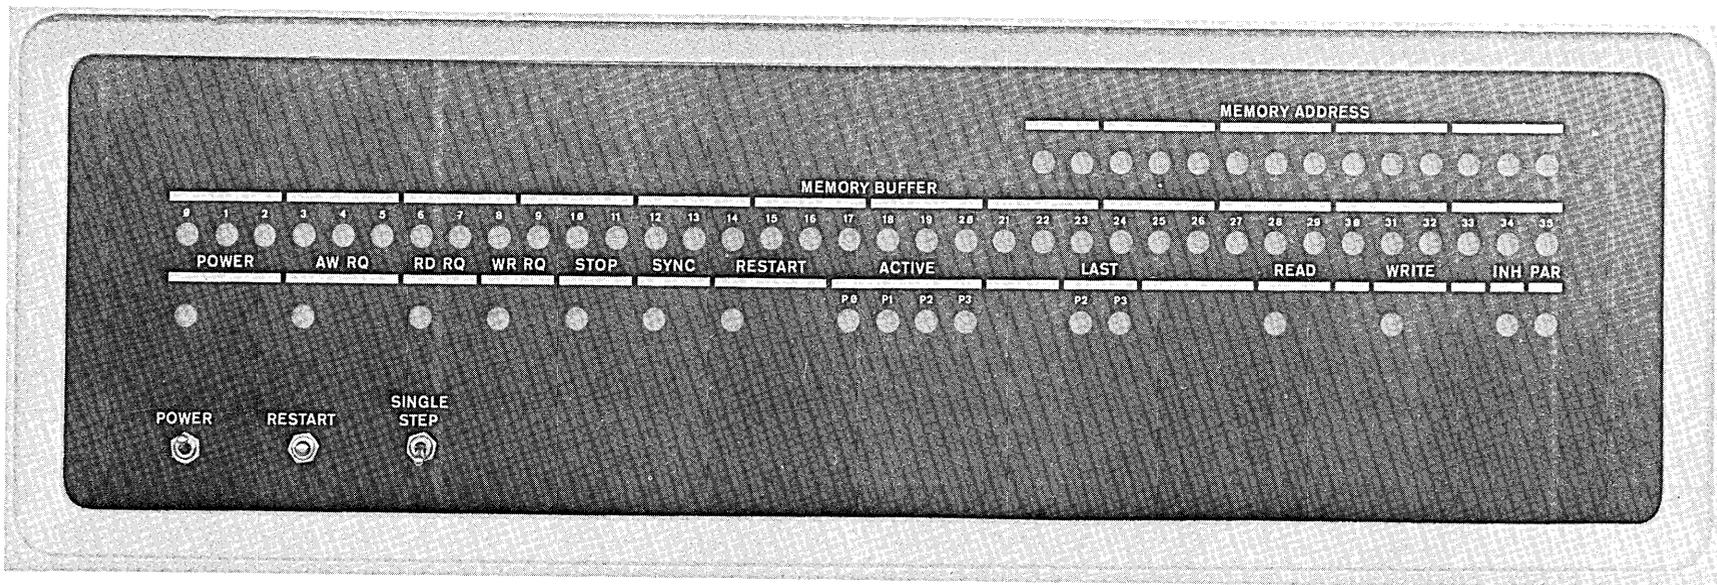
Indicator Panel, KA10 Arithmetic Processor, Bay 1



Indicator Panel, KA10 Arithmetic Processor, Bay 2



Indicator Panel, MA10 Core Memory (1.00 μ s)



Indicator Panel, MB10 Core Memory (1.65 μ s)

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