DEC-11-HMFA-D

MM11-F
core memory
manual

DIGITAL EQUIPMENT CORPORATION • MAYNARD, MASSACHUSETTS

1st Edition March 1971

The material in this manual is for information purposes and is subject to change without notice.

The following are trademarks of Digital Equipment Corporation, Maynard, Massachusetts:

DEC	PDP
FLIP CHIP	FOCAL
DIGITAL	COMPUTER LAB
UNIBUS	

CONTENTS

ILLUSTRATIONS

ILLUSTRATIONS (cont.)

TABLES

CHAPTER 1
INTRODUCTION

1.1 SCOPE

This manual provides the user with theory of operation and logic diagrams necessary to understand and maintain the MM11-F Read/Write Core Memory. The level of discussion assumes that tne reader is familiar with basic digital computer theory. Both general and detailed descriptions of the core memory are included.

Although memory control signals and data pass through the Unibus(TM), it is beyond the scope of this manual to describe the operation of the Unibus itself. A detailed description of the Unibus is presented in the *PDP-11 Unibus Interface Manual*, DEC-11-HIAB-D.

A complete set of engineering logic drawings is shipped with each core memory. These drawings are bound in a separate volume entitled *MM11-F Core Memory, Engineering Drawings.* The drawings reflect the latest print revisions and correspond to the specific memory shipped to the user.

This manual is divided into three major sections:

a. Introduction

b. Theory of operation

c. Adjustment and maintenance procedures

1.2 INTRODUCTION

The MM11-F is a 16-bit, 4096-word (4K), read/write core memory designed for the PDP-11. The Unibus concept is used by the PDP-11 System; thus, the central processor does not contain an integral memory. The memory used by the system is an external device and is connected to the processor through the bus interface. The core memory functions as a true peripheral and is compatible with all PDP-11 Systems. It is normally used as the basic memory unit in any system that requires a large core memory.

The MM11-F Core Memory may exist in any one of four configurations, depending on whether parity and/or interleaving capabilities are included. These four configurations are listed in Table 1-1 along with a reference indicating the section of the manual that contains a detailed description of the configuration. Note that the interleaving capability is automatically added whenever an 8K or larger memory is used. The print set in this manual is applicable to all configurations.

(TM) Unibus is a trademark of Digital Equipment Corporation.

Table 1-1
MM11-F Core Memory Configurations

Designation	Description	Reference
MM11-F (4K only)	Standard core memory without parity or interleaving capabilities.	Basic MM11-F manual
MM11-FP (4K only)	Parity capability – an 18-bit core memory, identical in all respects to the MM11-F (4K) except for bits 16 and 17. Bit 16 is byte 0 parity. Bit 17 is byte 1 parity.	Memory covered in basic manual; parity in Appendix A
MM11-F (8K or higher)	Interleaving capability – a basic MM11-F that has address bits 1 and 13 interchanged to allow faster cycle times when alternate adjacent memory banks within an 8K block are addressed.	Memory covered in basic manual; interleaving in Appendix B
MM11-FP (8K or higher)	A memory unit combining both the interleaving and parity capabilities.	Basic Manual, Appendices A and B

1.3 MM11-F GENERAL DESCRIPTION

The standard PDP-11 core memory, designated MM11-F, is a random-access, coincident-current, magnetic read/write core memory; cycle time is 980 ns and access time is 400 ns. The memory comprises ferrite cores wired in a 3-D, 3-wire, planar configuration. The basic unit can store up to 4096 (4K) 16-bit words. The memory can be expanded up to 28K words in 4K increments. The hard-wired address that selects the memory for use as an external device is preselected by the user simply by moving appropriate jumpers on a logic card.

The Unibus concept of PDP-11 System requires that a master/slave relationship exist between the processor and an external device or between two different external devices. Although many devices can function as either master or slave, the core memory always functions as a slave, whether the controlling unit is the processor or another peripheral. The core memory can never function as bus master.

The basic functional components of the core memory are briefly described in the following paragraphs.

1.3.1 Core Array

The ferrite core array consists of 16 core mats, each wired in a 64-by-64 matrix. This arrangement provides a total of 4096 16-bit words of data and/or program storage. The cores are connected by three wires. An X and a Y line

are used for individual core selection. The third line is a shared sense/inhibit line. There is a separate sense/inhibit line for each mat. This single line can function as either a read or write line due to the method of wiring and the control logic circuits.

1.3.2 Memory Control

Memory control circuits acknowledge the requests of the master device, determine which of the four basic operations is to be performed, and set up the appropriate timing and logic sequences to perform the desired read or write operations. The memory control logic also transfers data to or from the Unibus as required.

1.3.3 Address Selection

The core memory receives an 18-bit address from the master device. The address is decoded to determine if the memory is the selected device and to determine the core location specifically addressed. If the operation is a byte operation, bus line A00 L indicates the byte to be used. The actual read or write operation is not selected by the address but is selected by the settings of the bus C lines.

The X and Y portion of the address is decoded through selection switches and a diode matrix to enable passage of read/write current through the selected X and Y drive lines of the memory. The coincidence of these currents selects the specific 16-bit core memory location desired.

1.3.4 Inhibit Drivers

The core memory is so designed that, unless inhibited, all bit locations of the selected memory cell are switched to a logical 1 during the write portion of the memory cycle. To prevent this occurrence, each bit has inhibit drivers that are used during write time to oppose the Y write current, thereby ensuring that, if logic 0 levels are stored in the data register, they are written in the corresponding bit location of the addressed memory cell.

1.3.5 Sense Amplifiers

During the read portion of a memory cycle, sense amplifiers detect analog signals induced in the sense/inhibit windings of the core array. These signals are shaped, amplified, and then time sampled by means of a strobe signal to set corresponding bits of the data register.

1.3.6 Data Register

The data register is a 16-bit flip-flop register (18-bit register if the parity option is included) used to store the contents of a word after it is read out of the destructive memory; the same word can then be written back into memory (restored) when in the data in (DATI) mode. The register is also used to shift data from the Unibus lines to accommodate the loading of incoming data into the core memory during the data out (DATO) or data out, byte (DATOB) cycles.

1.3.7 Switches and Drivers

The switches and drivers direct the flow of current through the magnetic cores to ensure the proper polarity for the desired function. This action is necessary because a single read/write line is used, and the current for a write operation is opposite in polarity to the current required for a read operation. There are separate switches and drivers for the read and write circuits in the selection matrix.

1.3.8 Current Generator

The current generator provides the current necessary to change the state of the magnetic cores. The linear rise time and amplitude of the output-current waveform have been selected to provide optimum switching of the core states and maximum signal-to-noise ratio for a wide range of temperatures.

1.3.9 G103 Levels and Gates Module

The G103 Levels and Gates Module performs address gating functions for X and Y selection; generates the -6V and -5.2V threshold voltages which are derived from the -15V supply; and provides an ac termination for the X-Y switching matrices.

1.4 BASIC MEMORY OPERATIONS

The core memory has four basic modes of operation. The main function of the memory is simply to read or write data. Additional modes are provided, however, to allow for byte operation and to eliminate the restore cycle when it is not needed, thereby increasing overall system efficiency. The four basic memory operations are:

a. Read/restore (DATI)

b. Read only (DATIP)

c. Write (DATO)

d. Write byte (DATOB).

These four modes are discussed briefly in the following paragraphs.

NOTE

In the following discussions, all operations refer to the master (controlling) device. For example, the term data out indicates data flowing out of the master and into the memory.

1.4.1 Data In (DATI) Cycle

The DATI cycle is a read/restore memory cycle. During this operation, the memory reads the information from the selected core location, transfers it to the Unibus, and then writes the information back into the memory location. This last step is necessary because the core memory is a destructive readout device. During the first part of the cycle, the memory loads the data into a register; at the same time, the memory applies the data to the Unibus. Then, during the second part of the cycle, the memory takes the data from the register and writes it back into the memory location.

1.4.2 Data In, Pause (DATIP) Cycle

Normally in reading memory, the information is destroyed in the particular location accessed, and the data must be restored. However, sometimes it is not actually necessary to restore the information after reading, because the location is to have new data written into it. In this instance, eliminating the restore operation decreases the memory cycle time by approximately 50 percent. The DATIP operation is used for this purpose. The data is read from memory and the restore cycle is inhibited. Because no restore cycle is used, a DATIP must always be followed by a write cycle (either DATO or DATOB).

1.4.3 Data Out (DATO) Cycle

The DATO cycle is a write memory cycle used by the master device to transfer data into core memory. To ensure that proper data are stored, the memory unit must first be cleared by reading the cores (thereby setting them all

to zero) before writing in the new data. During a normal DATO, the memory first performs the read operation to clear the cores and then performs a write cycle to transfer data from the bus into the selected core location. If a DATO follows a DATIP (rather than a DATI), the sequence is not the same. The DATIP clears core and generates a pause flag; thus, the memory skips the read cycle and immediately begins the write cycle. This process reduces DATO cycle time by approximately 50 percent.

1.4.4 Data Out, Byte (DATOB) Cycle

The DATOB cycle is similar in function to the DATO cycle, except that during DATOB data is transferred into the core memory from the bus in byte form rather than as a full word. Actually, an entire word is loaded into the selected memory location: the selected byte which is new data from the bus and the non-selected byte, which is restored data from the word previously stored in that memory location. During the read cycle, the non-selected byte is saved by storing it in the data register while the selected byte is cleared. During the write cycle, only the selected byte portion of the word is loaded into the memory location from the bus. At the same time, the non-selected byte is restored from the data register into the memory location. In effect, the memory is first cleared and then simultaneously performs a restore cycle for the non-selected byte and a write cycle for the selected byte.

CHAPTER 2
THEORY OF OPERATION

2.1 INTRODUCTION

This chapter provides a detailed description of the MM11-F Core Memory. The first topic presented is the core array and overall memory operation. Subsequent paragraphs describe X- and Y-line selection and decoding, read/write operation, control and timing logic, read/write drivers and switches, and the current generator.

The information in this chapter is supported by a complete set of engineering drawings which are contained in a separate volume. All calibration and adjustment information is included in Chapter 3.

2.2 CORE ARRAY

The ferrite core memory consists of 16 memory mats; each mat contains 4096 ferrite cores arranged in a 64-by-64 array. Each mat represents a single bit position of a word. The memory is referred to as a planar memory. This planar configuration provides a total of 4096 16-bit word locations. Each ferrite core can assume a stable magnetic state corresponding to either a binary 1 or a binary 0. Even if power is removed from the core, the core retains its state until changed by appropriate control signals. The ferrite cores themselves are 20-mil cores. The outside diameter of each core is 22 mil; the inside diameter is approximately 16 mil. Each core is 5.5 mil thick.

Selection and switching of the cores is provided by only three wires traversing each core in a special selection technique. An X-axis read/write winding passes through all cores in each horizontal row for all 16 mats. A Y-axis read/write winding passes through all cores in each vertical row for all 16 mats. Through the use of selection circuits which control the current applied to specific X-Y windings, any one of the 4096 word locations can be addressed for writing data into memory or reading data out of memory. A third line passes through each core on a mat to provide the sense/inhibit functions. There is one sense/inhibit line per mat. This single sense/inhibit line, as well as the selection circuits, are discussed in subsequent paragraphs.

2.3 MEMORY OPERATION

Figure 2-1 illustrates a typical portion of the core memory. An X and Y winding pass through each core in the mat. The current passing through any one winding is such that no single winding produces a magnetic field strong enough to cause a core to change its magnetic state. Only the reinforcing magnetic field caused by the coincident current of both an X and a Y winding can cause the core located at the point of intersection to change states. It is this principle that allows the relatively simple winding arrangement to select one and only one memory core out of the possible 4096 contained on each mat. The current passing through either an X or Y winding is referred to as the half-select current.

A half-select current passing through the X3 winding (see Figure 2-1) from left to right produces a magnetic field that tends to change all cores in that horizontal row from the 0 to 1 state. The flux produced by the current is, however, insufficient to complete the state transition in any core. Simultaneously passing a half-select current

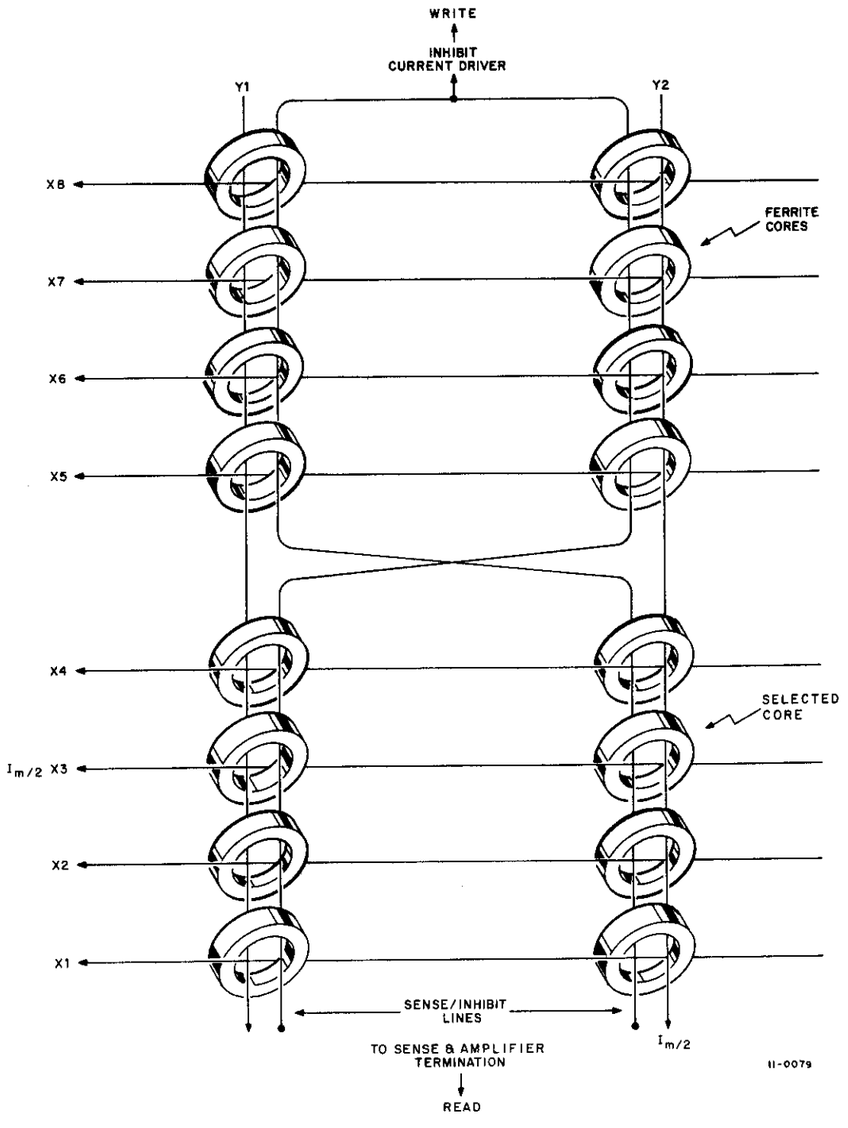

Figure 2-1 Three-Wire Memory Configuration

through the Y1 winding from top to bottom produces the same effect on all cores in that particular vertical row. Note, however, that both currents pass through only one core which is located at the intersection of the X3 and Y1 windings. This is the selected core and the combined current values are sufficient to change the state of the core.

All X and Y windings are arranged in such a manner that whenever a half-select current is passed through each, the resultant magnetic fields combine in the core at the point of intersection. This combined, full-select current ensures that the selected core is left in the binary 1 state. The currents used to select the core are referred to as write currents.

In the MM11-F Core Memory, the X3 windings in all 16 mats are connected in series as are the Y1 windings. Therefore, whenever a full-select current flows through a selected core on one mat, it also flows through an identical core on the other 15 mats. The X3-Y1 cores on all mats switch to a binary 1, causing each of the 16 cores to become one bit of a 16-bit storage cell.

Because of the serial nature of the X-Y windings, a method must be employed to set certain cores to the 0 state; otherwise, every 16-bit word selected would be all 1s. The method used in the MM11-F Core Memory is to first clear all cores to the 0 state by reading and then, using an inhibit winding, inhibit cores on particular mats. The inhibited cores remain 0s even when identical cores on other mats are set to 1s.

In many memories, the inhibit line serves only the inhibiting function; however, in the MM11-F, this line serves as both an inhibit and sense line. As a result, only a three-wire memory is necessary, rather than a four-wire memory. Only the inhibit function is discussed at this point.

The half-select current for the inhibit lines is applied from an inhibit current driver, which is a switch and a resistor between the inhibit line and -15V. The current in the inhibit line flows in the opposite direction from the write current in that line and cancels out the write current. There is a separate inhibit driver for each memory mat, and each mat represents one bit position of a word; thus, selected bits can be inhibited to produce any combination of binary 1s and 0s desired in the 16-bit word. It must be remembered that the inhibit function is active only during write time.

The sense/inhibit lines are also used to read out information in a selected 16-bit memory cell. The specific core is selected at read time in the same manner as during the write cycle with one notable exception: the X and Y currents are in the opposite direction. These opposite half-select currents cause all cores previously set to 1 to change to 0; cores previously set to 0 are not affected. Whenever the core changes from 1 to 0, the flux change induces a current in the sense winding of that mat. This current is detected and amplified by a sense amplifier. The amplifier strobes the signal into the data register for eventual transfer to the Unibus.

Figure 2-2 shows a 16-word by 4-bit planar memory. The MM11-F Core Memory functions in the same manner, except that it has 64 X-lines, 64 Y-lines, and 16 core mats. The core stringing is identical, and the sense windings are strung through all 4096 cores with the interchange between X31 and X32 instead of between X1 and X2. The actual configuration is shown in Figures 2-3 through 2-6. These figures show the core orientation, the sense inhibit windings, the X-Y lines, and word size variations.

2.4 X- and Y-LINE SELECTION

Previous paragraphs have explained the method for selecting a specific core by passing half-select current values through specific X and Y windings. This paragraph is devoted to an explanation of the method used for selecting the specific X and Y lines.

An 8-by-8 decoding matrix is used to select 1 of the 64 X-lines. An identical matrix is used to select 1 out of the 64 Y-lines. For ease of presentation, only X-line decoding is discussed. Decoding of the Y lines is identical.

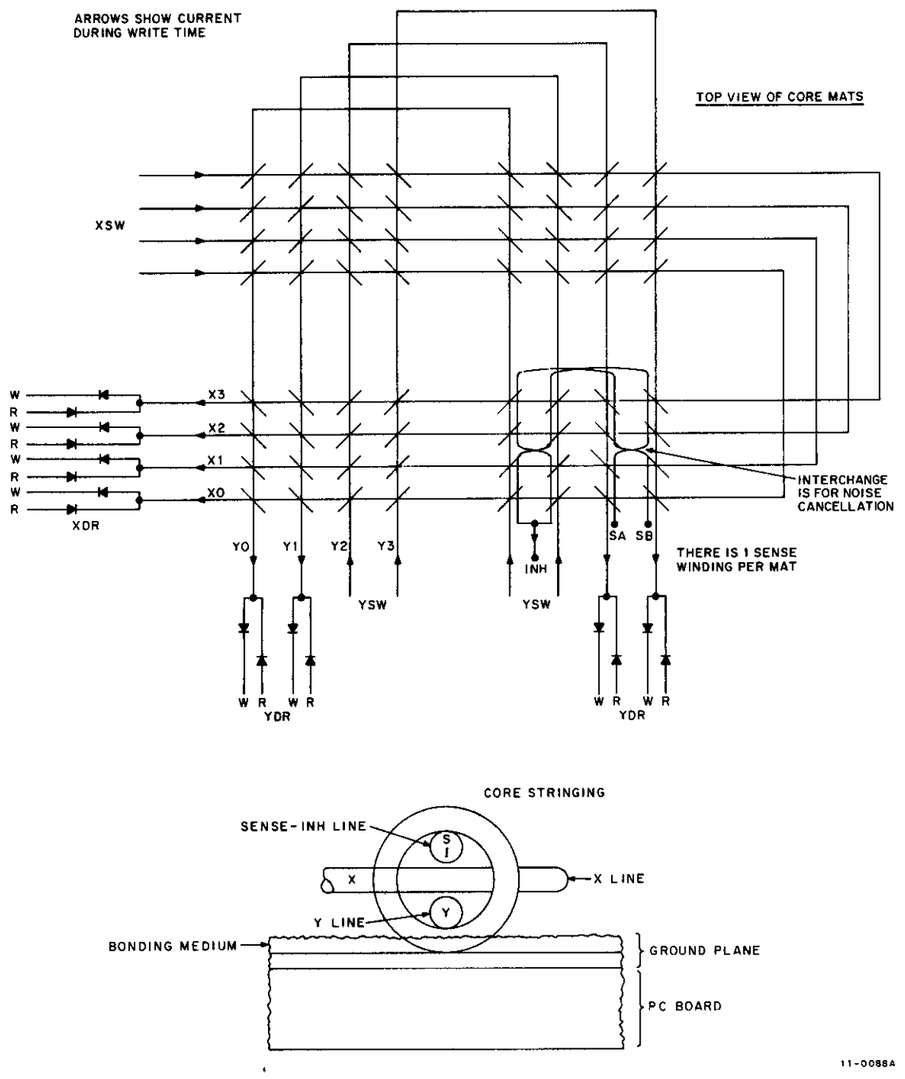

Figure 2-2a Three-Wire 3D Memory, Four Mats Shown for a 16-Word – 4-Bit Memory

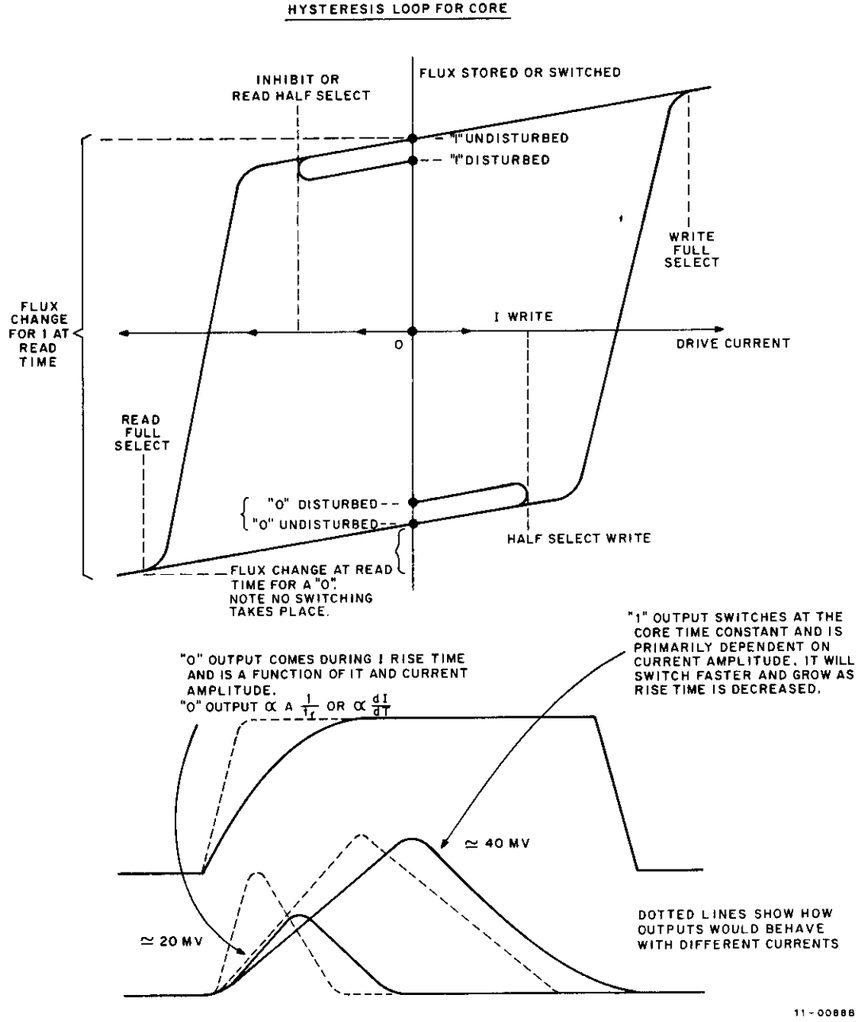

Figure 2-2b Hysteresis Loop for Core

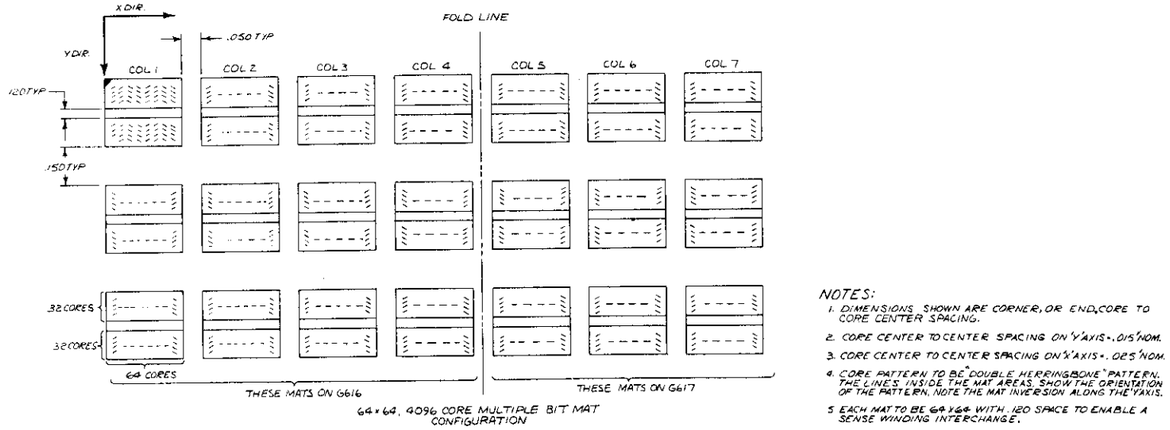

Figure 2-3 Core Orientation

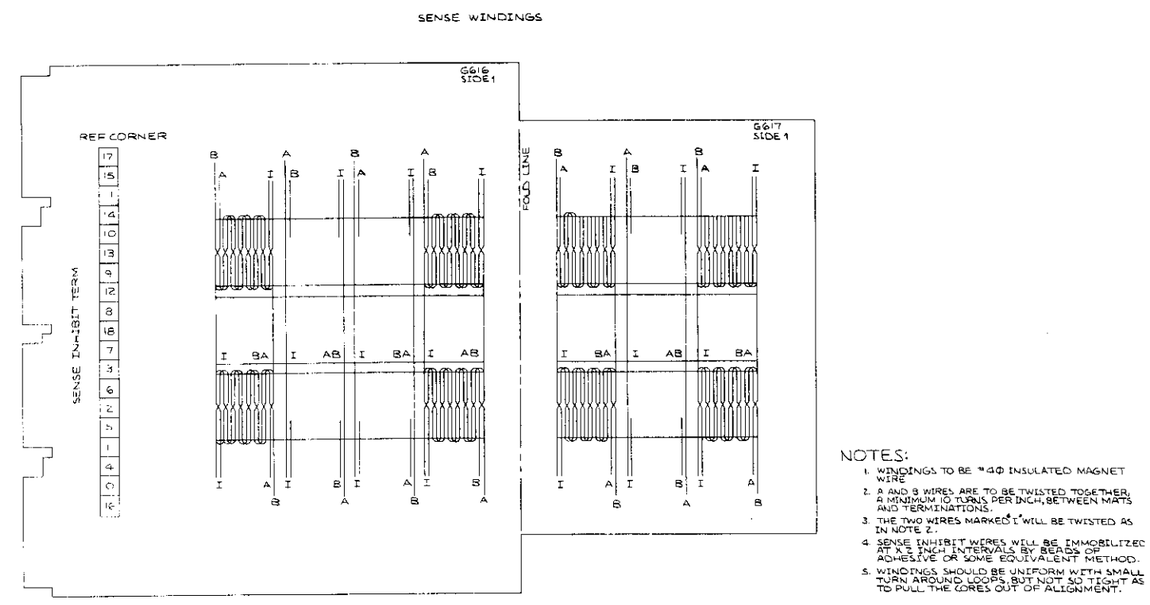

Figure 2-4 Sense Inhibit Windings

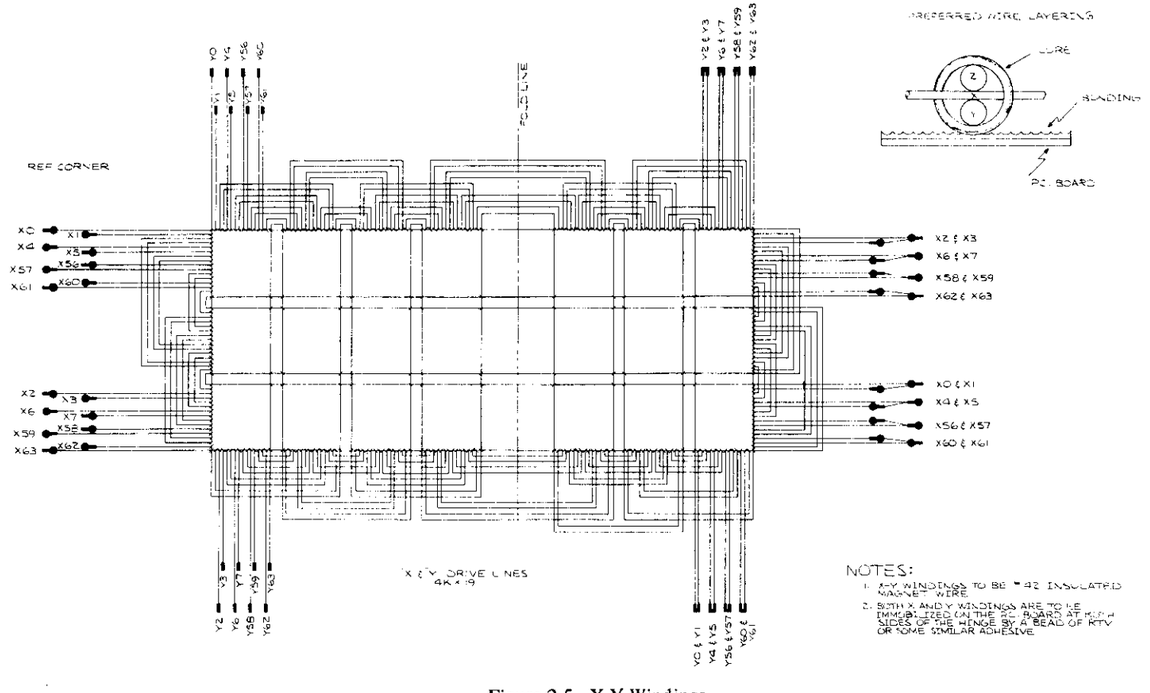

Figure 2-5 X-Y Windings

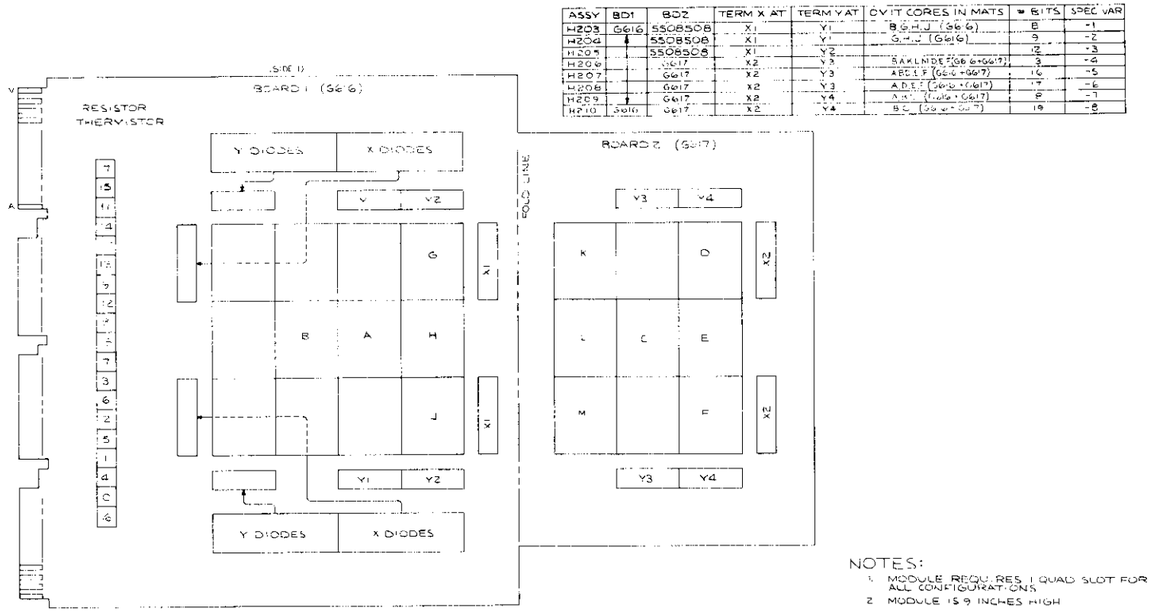

ASSY	BD1	BD2	TERM X AT	TERM Y AT	CW IT CORES IN MATS	= BITS	SPEC VAR
H203	5616	5508508	X1	Y1	B,G,H,J (5616)	8	-1
H204	↑	5508508	X1	Y1	G,H,J (5616)	9	-2
H205	↓	5508508	X1	Y2	↓	12	-3
H206		5617	X2	Y3	SAME AS H203 (5616+5617)	8	-4
H207		5617	X2	Y3	A,B,E,F (5616+5617)	16	-5
H208		5617	X2	Y3	A,B,E (5616+5617)	17	-6
H209		5617	X2	Y4	A,H (5616+5617)	8	-7
H210	5616	5617	X2	Y4	B,C (5616+5617)	19	-8

Figure 2-6 Word Size Variation

The incoming address actually performs four separate functions; however, only X-line decoding is discussed at this point. The four functions and corresponding bit assignments of the address are listed in Table 2-1.

Table 2-1
Address Functions

Address Bit Assignment	Function
0	Select byte only during DATOB
*1, 2 through 6	Y-line selection bits 1 through 3 = positive, negative drivers 0 through 7 bits 4 through 6 = positive, negative switches 0 through 7
7 through 12	X-line selection bits 7 through 9 = positive, negative drivers 0 through 7 bits 10 through 12 = positive, negative switches 0 through 7
*13, 14 through 17	Device selection

*When memory is interleaved, bits 1 and 13 are interchanged.

Figure 2-7a is a representation of one pair of decoders used to select the X-line selection switch. As the method of decoding is slightly unorthodox, consider the A and B decoders as a single 4-line to 16-line decoder with 16 output pins. There are two pins for each switch: one for read operations and one for write operations. Figure 2-7b indicates the 16 possible outputs and the function of each. There are eight read switches and eight write switches: read switches are positive, write switches are negative.

In actuality, bit 12 of the incoming address effectively selects either the A or B decoder. If the bit is binary 0, the A decoder is used; if the bit is binary 1, the B decoder is used. The read signal selects one half of the selected decoder. If the read signal is present, the output comes from the upper half of the decoder, and one of the four read outputs is used. If the read signal is absent, the lower half of the decoder provides one of the four write outputs. The remaining two bits, bits 10 and 11, select one of the four outputs (0 through 3) from the decoder.

As an example, assume that all bits are set and a read signal is present. Bit 12 set selects the B decoder. The read signal selects the upper half of the decoder. Bits 10 and 11 set represent a binary 3 so that output 7 from pin 3 is the selected read switch.

Assume now that only bit 11 is set and the read signal is not present. Bit 12 being 0 causes the A decoder to be used. The absence of a read signal effectively selects the lower (write) portion of the decoder. Bit 11 being set with bit 10 being 0 represents a binary 2. The third number (binary 2) in the write portion of the decoder is represented by pin 6. Therefore, the output is write switch 2.

The read/write driver that corresponds to the other X octal digit is selected in the same fashion using TDRH and bits 9, 8, and 7.

The portion of the address used to select the Y lines is represented by bits *1, 2, through 6. Therefore, bits *1 through 12 of the incoming address are used to select the X and Y lines that determine the specific core where a bit is to be read or written. One core on each bit is selected (refer to Tables 2-1 and 2-2).

*When memory is interleaved, bits 1 and 13 are interchanged.

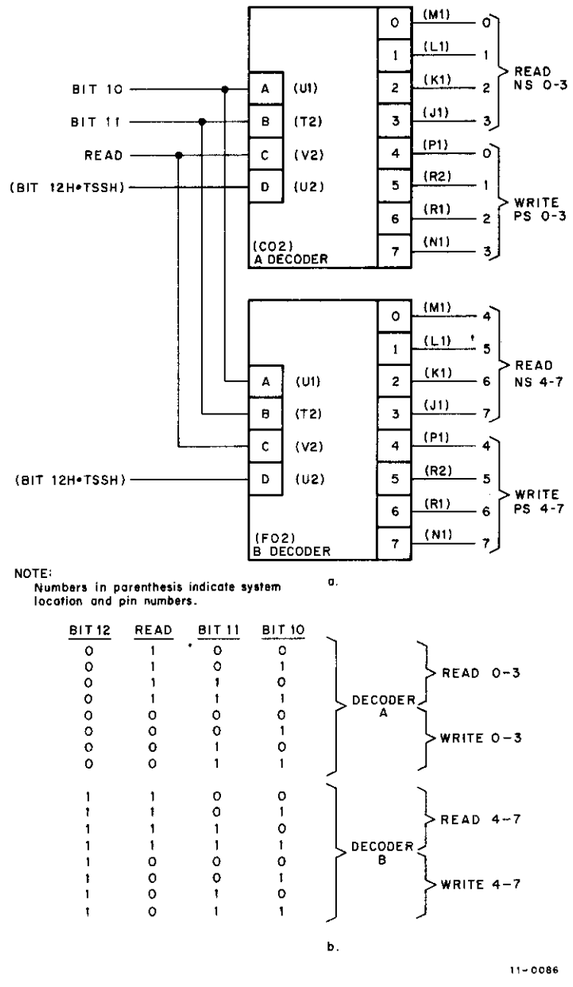

Figure 2-7 X-Line Decoding

The X and Y selection logic is shown on drawing MM11-F-03, sheets 1 and 2. The decoder output numbers are not shown on the drawings in the same manner as described in this section.

Figure 2-8 shows the switch or driver base drive circuit that is connected to the decoder. The timing and address inputs cause the 8251 decoder output to go to ground and current i_1 flows from the +5V supply, through the resistor, through the transformer primary, and into the decoder output. The value of the current is determined by the resistor and the voltage reflected into the transformer primary (approximately 1.0V). An equal current (i_2) is induced into the base emitter circuit connected to the transformer secondary. This current turns on the transistor. Note that all base current is provided from this circuit and the current from circuit i_3 equals the collector current.

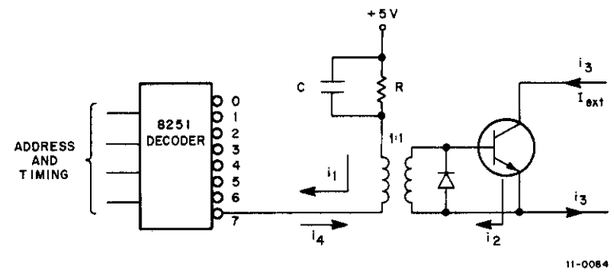

Figure 2-8 Switch or Driver Base Drive Circuit

When the decoder is turned off by the timing input, its pull-up transistor now tries to drive current in the opposite direction (i_4). The capacitor performs two functions: it speeds up the turn-on current i_1 and allows the decoder to pump reverse current i_4 into the transformer primary. The reverse current is necessary to remove the base emitter charge on the output transistor in the transformer secondary to turn it off. The base emitter diode prevents reverse breakdown of that junction during turnoff. Since this diode limits the transformer secondary voltage during turnoff, some of the circuits contain two diodes in series to allow faster transformer recovery. Note that the dc level of the output transistor is irrelevant. This transistor is saturated when on.

2.5 DRIVERS AND SWITCHES

Drivers and switches direct the flow of current through the cores to ensure proper polarity for the desired read or write function. A read driver and switch, as well as a write driver and switch, are provided for groups of eight X lines and groups of eight Y lines in the selection matrix, such that each axis has its 64 lines selected by an 8-by-8 matrix.

Figure 2-9 is a simplified diagram of the switches and drivers. The diagram represents one X line. Note that the flow of read current is shown by a solid line; the write current path is shown by a broken line.

When a read operation is selected, the read driver and read switch are turned on, and the write driver and write switch are turned off. The output current from the current generator flows through the read driver, through the cores on the associated X line, and through the read switch. When a write operation is selected, the write driver and switch are turned on, and the read driver and switch are turned off. In this case, current from the current generator flows through the write switch, through the cores (in the opposite direction from that during read), and through the write driver. The signals that select the read and write operations are given in Table 2-2. Each line has 64 x 16 (or 1024) cores on its as it threads through all 16 mats.

Resistors R1 and R2 are provided as terminations to prevent unwanted currents from affecting the drivers and switches. The circuit that comprises transistor Q1 and resistor R5 is used to make the write current equal the read current and to assist the ferrite core stack in recovering. This method is much faster than allowing the stack to recover by itself. The decoder output and write driver transformer base drive are shown. All base drives are similar.

The drivers and switches in one line share a common line with drivers and switches in other lines. Figure 2-10 (representing 16 lines on a 4-by-4 matrix) shows the interconnection of four drivers and four switches on 16 lines to aid in understanding the interconnection of all 64 lines in the 8-by-8 matrix. Only one line is common between any driver and switch. Note how the diodes prevent sneak paths. Only a unipolar matrix is shown.

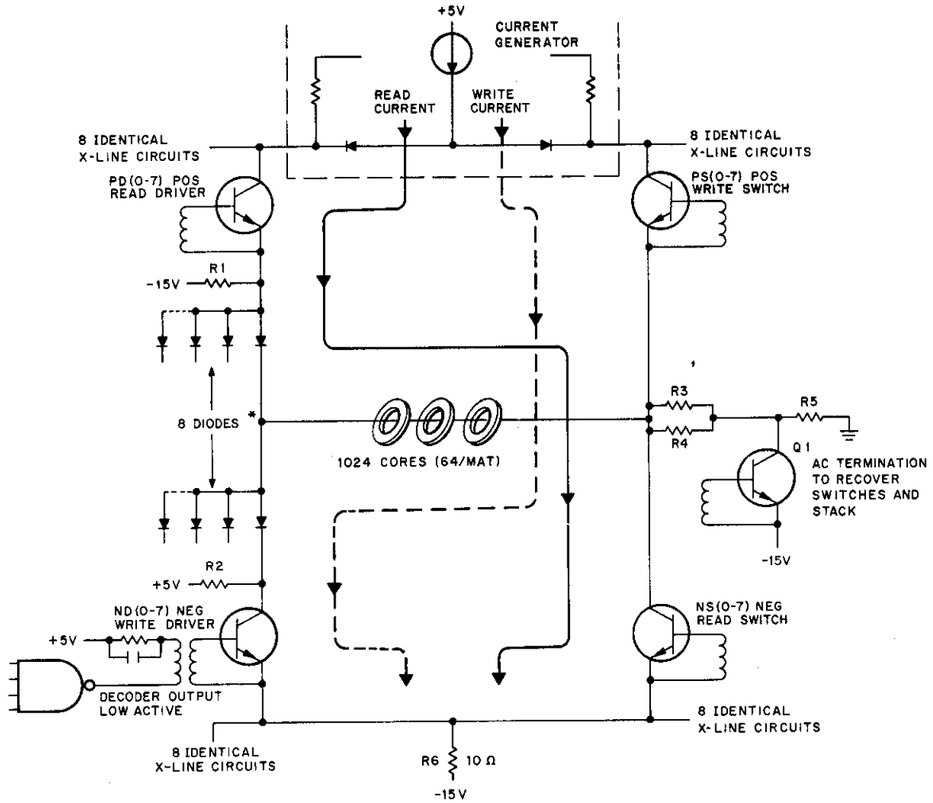

—— READ CURRENT
- - - WRITE CURRENT
* THE DRIVER AND SWITCH EACH FEED 8 LINES. HOWEVER, ANY ONE LINE IS COMMON TO BOTH.

11-0078

Figure 2-9 X-Line Drivers and Switches

Table 2-2
Read and Write Selection

Driver or Switch	Signals Required to Turn On
Read Driver	ADDRESS and TDR and READ H
Read Switch	ADDRESS and TSS and READ H
Write Driver	ADDRESS and TDR and WRITE H
Write Switch	ADDRESS and TSS and WRITE H

NOTE

READ H and WRITE H are both outputs of the same control flip-flop. When the READ flip-flop is set, the output is READ H; when it is cleared, the output is WRITE H.

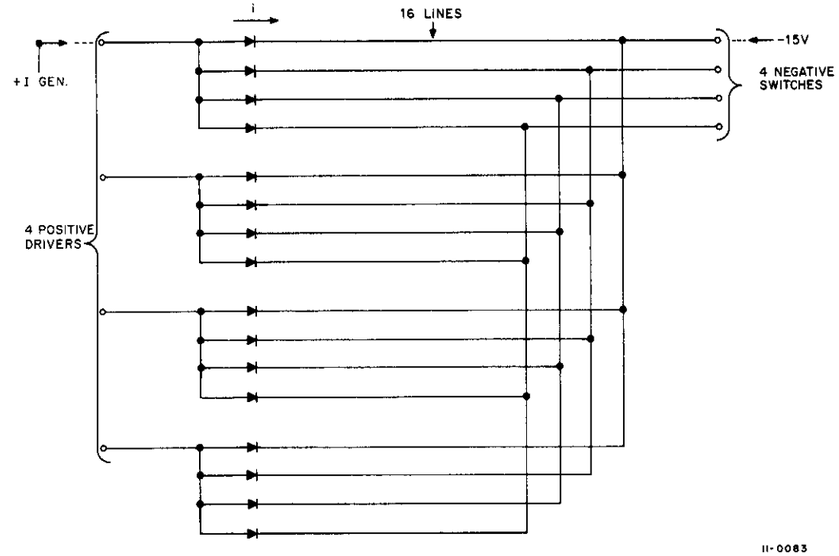

Figure 2-10 Simplified Interconnection Diagram

2.6 CURRENT GENERATOR

The current generator provides the current required to change states of the magnetic core. Generator design is such that the linear rise time and amplitude of the output current waveform ensure optimum switching of core states. Figure 2-11 is a simplified schematic of the current generator; Figure 2-12 illustrates the test point and output waveforms.

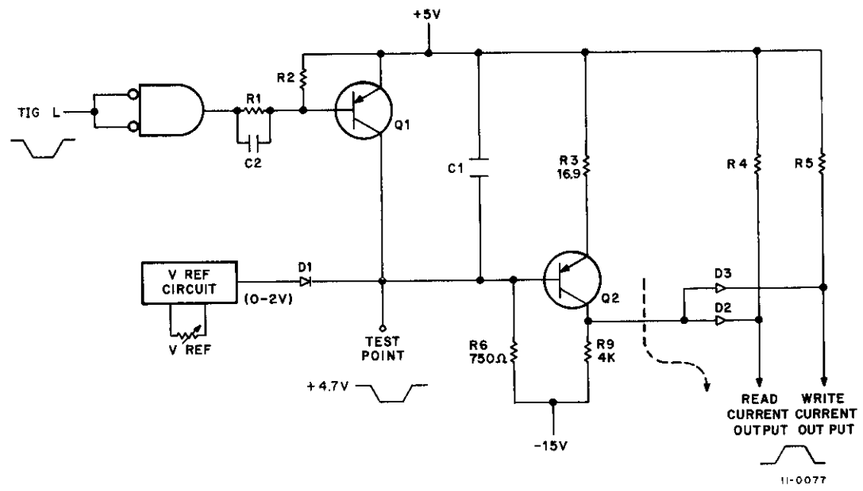

Figure 2-11 Current Generator, Simplified Schematic

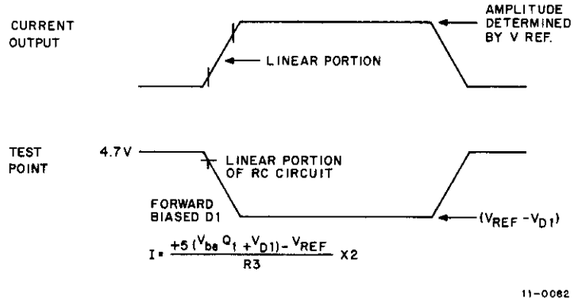

Figure 2-12 Current Generator Waveforms

When the current generator is in the quiescent state, transistor Q1 is saturated, thereby holding the voltage level at the test point to 4.7V. Diode D1 blocks any voltage from the voltage reference (V REF) circuit, and transistor Q2 is cut off. Therefore, the only output at this time is the current through R4 and R5, which is negligible.

Operation of the current generator is triggered by the negative-going TSS L signal from the memory control circuits. This signal is inverted by an AND gate and used to cut off transistor Q1. When Q1 cuts off, it allows the voltage at the test point to start going negative until it reaches the forward bias level of diode D1. This level is equal to the voltage reference minus the voltage drop across diode D1. The voltage reference circuit contains a trimpot to adjust the output between 0V and 2V. The adjustment procedure is given in Chapter 3.

At the same time that transistor Q1 cuts off, capacitor C1 begins charging. The time constant of C1 and resistor R6 determines the rise time as Q2 conducts to provide the read current. The write current is developed in the same manner. There are actually two circuits (R3 and Q2) in parallel.

The value of the output current is:

$$\frac{+5V\ (-V_{beQ2} + V_{beD1}) - V_{ref}}{R3}$$

This value is approximately 250 mA per circuit, or approximately 500 mA for the two circuits. The amplitude of the current output waveform is determined by the setting of the V REF trimpot. Resistors R4 and R5, and diodes D2 and D3 in the output lines, are used to isolate the read switches from the write switches.

Although not shown on the schematic, a resistor and thermistor on the memory stack (see drawing MM11-F-03) are connected to the voltage reference portion of the current generator. These components automatically track the current amplitude with temperature to ensure that the amplitude remains within specified tolerances over the temperature range of 0° C to 50° C.

2.7 CONTROL AND TIMING LOGIC

The memory control and timing logic circuits perform three basic functions:

a. Decode the incoming address (bits *13 through 17) to determine if the memory has been selected to engage in a data transfer

b. Latch the address bits controlling the X and Y lines (bits *A1 through A12) to ensure that the proper core location is selected

*When memory is interleaved, bits 1 and 13 are interchanged.

c. Perform the appropriate data transfer as requested by the master device by providing all internal memory timing.

The following paragraphs discuss device selection, X-Y address latching, and control and timing circuits for each of the four transfer operations (DATI, DATIP, DATO, and DATOB).

The control and timing circuits are shown on drawing MM11-F-05, Sheets 1 and 2. Sheet 2 contains the logic associated with the bus address and data lines, including the device selection logic. The remainder of the control and timing logic circuits are shown on Sheet 1. The timing diagram is shown on drawing MM11-F-08.

Details of the one-shot delays on the M7290 Module are shown on Figure 2-13. This figure provides circuit drawings and waveforms for each of the three types of one-shots used on the M7290 Module.

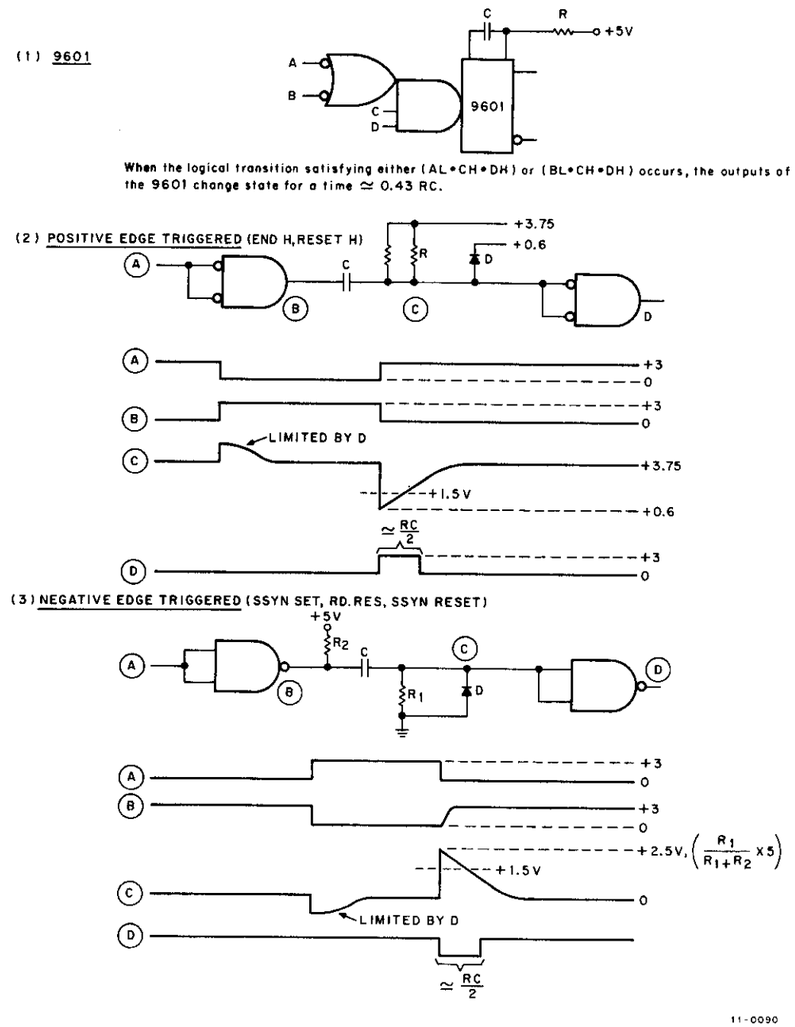

Figure 2-13 M7290 One-Shot Delays

2.7.1 General

All memory operations are started by the generation of a memory select (MSEL) signal. This signal is initiated on receipt of a MSYN (master sync) signal from the controlling device, a device select (D SEL) signal from the memory address (*A13, A14 through A17) decoding circuits, and a DC OK signal. These signals indicate, respectively, that a master device has control of the Unibus, that the master device has selected the core memory as a slave device, and that the dc power level to the memory circuits is within specified tolerances. The three signals set the MSEL flip-flop to produce the memory select signal. This signal remains asserted until either the END WRITE or END PAUSE signal clears the flip-flop to indicate that the selected memory operation is complete. The MSEL flip-flop also acts as a memory busy signal by prohibiting a new cycle from starting until MSEL is cleared.

The bus master places appropriate levels on the two bus C lines to select the type of memory operation to be performed. The data placed on these lines by the master to select a specific operation is shown in Table 2-3.

Table 2-3
Selection of Memory Operation

C Line State	Octal Equivalent	Selected Memory Operation
C0 H, C1 H	0	DATI
C0 L, C1 H	1	DATIP
C0 H, C1 L	2	DATO
C0 L, C1 L, A0 H	3	DATOB 0
C0 L, C1 L, A0 L		DATOB 1

2.7.2 Device Selection

Whenever a master device addresses another peripheral as a slave, the address is applied to all peripherals connected to the Unibus. Therefore, it is necessary for each peripheral to have a method of recognizing its own address and ignoring other addresses. The device selection circuits serve this purpose. The 18-bit address from the master device is used for both device selection and selection of the X-Y matrix. This paragraph covers the memory device selection circuits and Paragraph 2.7.3 covers the X-Y selection.

Bus address lines *13, 14 through 17 are used for device selection. The selection logic is dependent upon jumpers connected to the input of a seven-input AND gate, 74H30. The device address is hard wired. If, for example, the desired bit is to be 0, a jumper is connected to the inverter output of the address buffer and ties this signal to the 74H30 gate. When a 0 is placed on the associated bus line, that leg of the 74H30 gate is qualified. It produces an output if all of the other four inputs to the gate are satisfied. If the bus line has a binary 1, that device selection gate leg is not qualified and the 74H30 gate cannot be qualified. On the other hand, if a binary 1 is desired in the address code, a jumper is connected directly from the SP 380 to the input of the 74H30 gate, bypassing the inverter gate for that line. This causes the 74H30 gate to be qualified whenever a 1 is present on the bus line. In this manner, any combination of 1s and 0s can be preselected as the device address. When the proper signals are present on bus address lines *13, 14 through 17, the 74H30 gate is qualified and produces a low output which, when inverted, is the device select (D SEL) signal. This signal indicates that the master device has selected the core memory as the slave device.

*When memory is interleaved, bits 1 and 13 are interchanged.

2.7.3 X-Y Address Latching

Bus address lines *1 through 12 are used to select the X and Y lines that select the desired core location (refer to Paragraph 2.4). Bits *1 through 6 select Y lines, and bits 7 through 12 select the X lines.

When the bus master generates MSYN, the full 18-bit address appears at the memory control circuits. The five most significant bits are used by the device selection circuits (refer to Paragraph 2.7.2). The next 12 bits are gated into a flip-flop register. If the bit is low on the bus, it qualifies a gate that sets a corresponding flip-flop. If the bit is high on the bus, the gate is not qualified and the corresponding flip-flop remains cleared. The flip-flop is set low rather than high because of the negative logic used by the Unibus. The original X-Y address from the master device is now stored in the flip-flop register to be used by the decoding circuits during the selected memory operation. An asserted address bit is high (+3V) inside the memory.

Bus address line 0 is used for byte selection. The level on line 0 passes through a gate and is stored in a flip-flop in the same manner as the X-Y address lines. The output of the flip-flop represents either byte 1 when set or byte 0 when not set and is applied to the memory control circuits described in Paragraph 2.7.7.

2.7.4 Data In (DATI) Control Circuits

A DATI operation indicates that a selected location in memory is to be read and the information transferred through the Unibus into the master device. Because the core memory is a destructive readout device, the data must be restored to the selected location after the read operation. Therefore, a DATI is basically a read/restore (read/write) operation.

In the following discussion, assume that all flip-flops are initially in the RESET state. When the M SEL flip-flop is in the zero state, MSYN L is low, SSYN L is high, DSEL H is high, and the M SEL flip-flop qualifies an AND gate to produce a CLK 1 H pulse that latches the address and C lines from the bus. The CLK 1 H pulse is also gated through an OR gate to set the DEL flip-flop 0 low, which starts a negative pulse through the delay line. When the pulse comes out of the delay line, it resets the DEL flip-flop. In the interim, outputs from taps on the delay line generate the read timing chain. The TSS H, TIG L, and TDR H signals are gated into the address decoders to enable the current paths through the X and Y switches and drivers. The CLK 1 H pulse also passes through a delay line to set the M SEL flip-flop. The TIG L signal turns on the current driver; the TSS H and TDR H signals enable the decoding matrix so that the selected core location receives the read currents. After the data register flip-flops have been reset, the current flowing through the selected cores produces an output to the sense amplifier for every bit position having a logical 1 (refer to Paragraph 2.8).

When the strobe delay is completed, a STROBE H pulse is generated and applied to the sense amplifier. If, during strobe time, a sense amplifier has an input exceeding the threshold level, the amplifier direct sets the associated flip-flop. The STROBE pulse is 50 ±10 ns wide. This width is fixed and cannot be changed. The STROBE DEL trimpot is factory adjusted so that the STROBE pulse occurs 220 ns after the read current is turned on (measured from the 10 percent point of the current) to ensure that core transitions are compared to the threshold level of the sense amplifiers at the proper time for best overall memory margins in the temperature range from 0° to 50° C.

The STROBE pulse is gated with C1 H and then gated through an OR gate to trigger a positive-edge circuit which produces a negative pulse that direct sets the SSYN flip-flop. The output of the SSYN flip-flop is gated with C1 L through an inverter to produce the DATA OUT H signal. The SSYN output also gates a bus driver to produce the BUS SSYN L signal. The DATA OUT H signal gates the outputs of the data register flip-flops onto the Unibus through the open collector bus drivers.

*When memory is interleaved, bits 1 and 13 are interchanged.

In effect, the memory places data from the selected core location on the Unibus data lines and asserts SSYN to inform the master device that it can now strobe in the data. After the master device strobes the data, it raises MSYN L which causes the memory control to produce a SSYN CLK signal, which resets the SSYN flip-flop. This action completes the read portion of the memory cycle. The memory then immediately enters the restore portion of the DATI cycle.

Between the end of the read cycle and the beginning of the restore cycle, all signals revert to their initial conditions with the exception of the M SEL signal. Although MSYN is dropped, the M SEL flip-flop remains set because it can only be cleared by an M SEL RESET signal which is not generated until the end of the write (restore) cycle. The DEL flip-flop is reset by the DEL R signal at the end of the read cycle.

When the READ flip-flop is reset by ANDing the READ H and END H signals, a WRITE H level is generated. The WRITE H level is combined with PAUSE L through an AND gate and then passes through an OR gate to set the DEL flip-flop 0 low, which starts a pulse through the delay line to generate the write timing chain.

The TSS H and TDR H signals are gated into the address decoder to enable the current paths through the X and Y switches and drivers. The TIG L pulse turns on the current drivers. The INH H pulse functions in the manner described in Paragraph 2.9. The current flowing through the X and Y switching circuits writes 1s into all cores not inhibited and the contents of the memory locations are restored to their original states. After this occurs, the RST H signal is ANDed with WRITE H and PS L to produce the RESET M SEL signal that clears the M SEL flip-flop. While set, the M SEL flip-flop locks out the MSYN L signal so that a new cycle cannot be started until this cycle is completed.

2.7.5 Data In, Pause (DATIP) Control Circuits

A DATIP operation indicates that a selected memory location is to be read out and the information transferred through the Unibus to the master device. Although the memory is a destructive readout device, the restore cycle is not entered because DATIP indicates that new data is to be stored in the memory location.

The read portion of a DATIP is identical to that of a DATI (refer to Paragraph 2.7.4) until the time the SSYN signal is produced. At this point, the output of the SSYN flip-flop and the output of the DATIP gate (from bus line C1) are applied to the PAUSE flip-flop, thereby setting it. This action produces PAUSE L which is the pause flag. The END H pulse is ANDed with READ H to clear the R/W flip-flop. The DP H and READ H levels are ANDed with the RST pulse to clear the M SEL flip-flop; the memory control circuits revert to their original states with the exception of the pause flag, which remains set. The restore (or write) portion of the memory cycle is inhibited. Regardless of which operation is defined, the pause flag prevents the READ flip-flop from being set at the start of the next cycle, thus forcing the next cycle to be either a DATO or DATOB. Therefore, the next time memory is selected, it performs a write only operation.

During normal memory operation, it is always necessary to read the memory before writing into it. This read operation magnetically changes all cores to zero, thereby clearing all cores before new data is written into them. In the case of a write operation (DATO or DATOB) following a DATIP, the normal read operation is not necessary because the cores have already been zeroed by the previous DATIP read operation. Note, however, that the DATO cycle following a DATIP is not identical to the DATO cycle that follows a DATI. In the former case, the pause flag is set; in the latter case, the pause flag is cleared. A discussion of a DATO operation with the pause flag cleared is covered in Paragraph 2.7.6.

The DATO operation associated with a DATIP begins when the master device asserts MSYN in order to start the timing sequence as described previously. The CLK 1 H pulse is generated as before; however, because the pause flag is set, no READ H signal is produced. The READ flip-flop remains cleared and produces a WRITE H output. Data from the bus master device is already applied through the input buffers to the D lines that are connected to the data register D lines. A LOAD H is then generated and applied to the data register C lines. When the LOAD H pulse is generated, the data register flip-flops are set or remain cleared corresponding to the 1s or 0s on the associated bus data lines. This action loads the data from the bus into the data register.

Just after the LOAD H signal is generated, the TSS and TDR signals are produced to enable the switches and drivers in the X and Y selection matrices. The inhibit (INH) signals are produced as described previously. Write currents flow through the selected core locations. The current flowing through the selected core writes 1s into all cores not inhibited. This causes the core memory location to contain the data stored in the data register. Thus, data from the Unibus has been transferred into the selected core memory location.

During a DATO or DATOB operation with the pause flag set, the SSYN signal is generated soon after the CLK 1 H signal is generated. The SSYN flip-flop is reset when the bus master device drops MSYN. The pause flag is reset when SSYN is generated. The M SEL flip-flop is cleared when the PS L, WR H, and RST pulses are ANDed together. This completes the write operation and all memory control circuits are in their original conditions.

2.7.6 Data Out (DATO) Control Circuits

A DATO operation indicates that the master device is to transfer data through the bus into the core memory for storage. In effect, the memory is to write data into the selected core location. It is necessary to clear the selected cores before writing; thus, the memory first performs a read cycle and then writes new data into the location.

The read portion of a DATO operation is similar to the read portion of a DATI operation except that no RESET 0 L, RESET 1 L, STROBE 0 H, or STROBE 1 H pulses are generated. When the master device asserts MSYN, it causes the M SEL, CLK 1 H, CLK 2 H, and READ signals to be produced. The LOAD 0 H and LOAD 1 H signals, together with the data on the bus D lines, are applied to the data register. The data register flip-flops are set or cleared in accordance with the data that is on the bus lines. The CLK 1 H pulse triggers the address flip-flops that produce the enabling signals for the address matrix in order to establish the proper current return paths. After this sequence is complete, the selected cores receive read current and perform the 1 to 0 transition, thereby clearing the cores. This transition does not enter the data register because no STROBE 0 H or STROBE 1 H signals are produced to gate the transitions out of the sense amplifiers.

This completes the read portion of the DATO cycle. All cores are in a 0 state and the data that had been on the bus is now stored in the data register. When the READ signal drops (the pause flag is cleared), a write cycle is initiated. The X- and Y-matrix enabling levels are supplied along with the inhibit signals to gate data from the data register into the selected core location during the write cycle as described previously.

2.7.7 Data Out, Byte (DATOB) Control Circuits

A DATOB operation is the same as a DATO operation except a byte, rather than a full word, is transferred. In the case of a DATOB, a RESET 0 or 1 L and STROBE 0 or 1 H signal are produced for the non-selected byte. This process effectively strobes the byte into the data register for restoration during the write operation, because this byte has not been selected to receive new data. A LOAD 0 or 1 H pulse is not provided for the non-selected byte; therefore, any data on the bus has no effect on the non-selected byte.

The selected byte is controlled in just the opposite manner. Neither RESET 0 or 1 L nor STROBE 0 or 1 H is generated because it is not necessary to restore the byte. A LOAD 0 or 1 H pulse is provided, however, so that data from the bus can be written into the selected byte location.

At the end of the DATOB operation, the selected byte has new information; the non-selected byte retains its previous data.

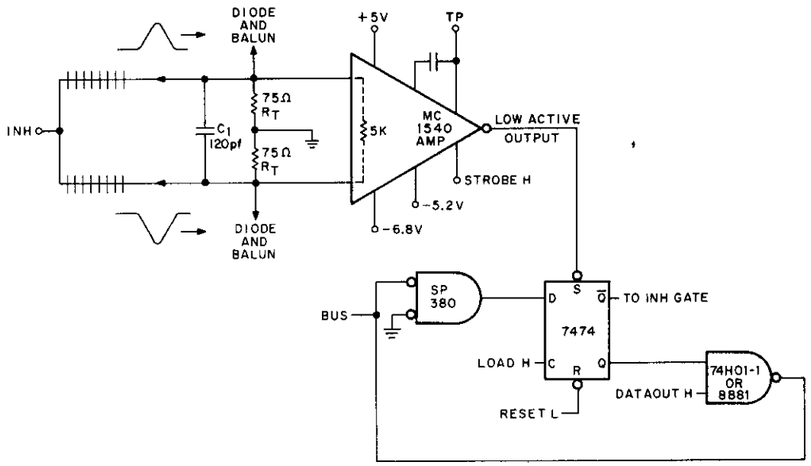

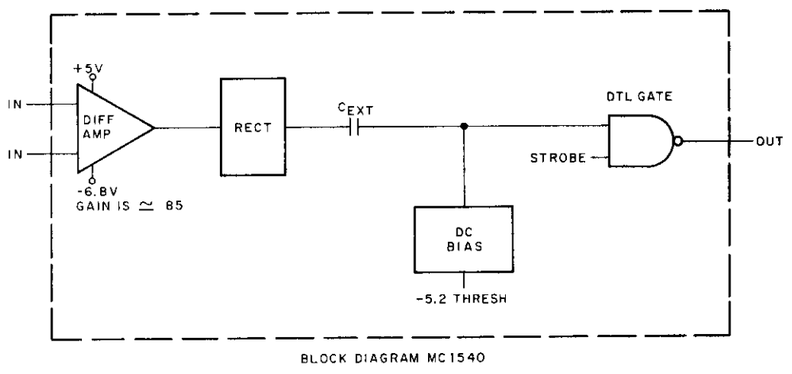

BLOCK DIAGRAM MC1540

11-0089A

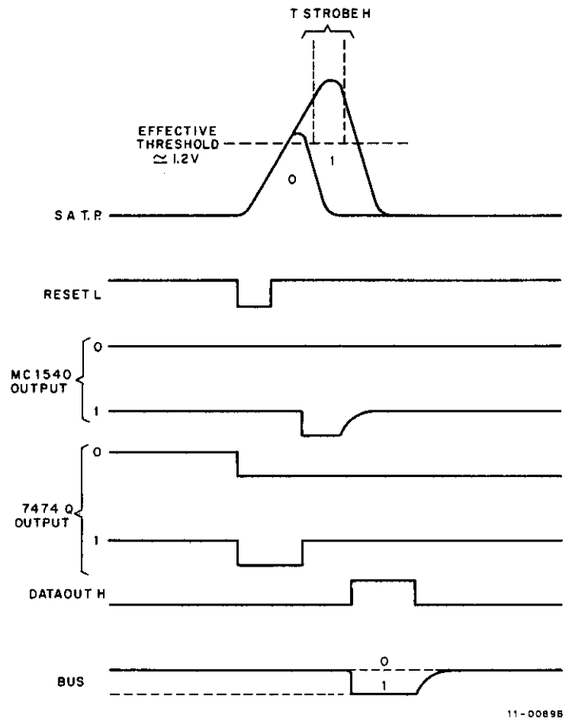

11-0089B

Figure 2-14 Sense Amplifier Detector Circuit

2.8 READ OPERATION

This paragraph is devoted to an explanation of the read operation. The actual control signals that govern the read cycle are covered in Paragraph 2.7. This section covers read operation only from the standpoint of the core memory itself.

Whenever a word is to be read from memory, a read current passes through the address decoder and related X and Y switches and drivers to the selected core. The read current is a half-select current in both the X and Y windings. Since read current flows in the opposite direction of write current, all cores in the selected word previously set to the 1 state are switched to the 0 state; cores previously set to 0 remain unchanged. To make read current opposite to write current, positive drivers and negative switches are used for reading, and negative drivers and positive switches are used for writing. A more detailed explanation of the switches and drivers is given in Paragraph 2.5.

A sense/inhibit termination circuit (referred to as Circuit 1 on logic drawing MM11-F-04) and sense amplifier are provided for each of the 16 mats in the memory. The input to each amplifier comes from a sense/inhibit winding that passes through every core on the associated mat. If, during the read operation, the addressed core in a mat makes the 1 to 0 transition, the flux change induces a current in the sense/inhibit winding of that mat. This current develops a voltage pulse at the input to the amplifier circuit. The input is amplified and (after strobed threshold detection) used to set a flip-flop in the data register when the strobe pulse is applied to the amplifier.

The threshold level (voltage slice) is a factory-set level of -5.2V. Whenever a core makes the 1 to 0 transition, the resultant induced voltage is sufficient to exceed the preset threshold level during strobe time and causes the amplifier to produce an output. Addressed cores that were already in the 0 state induce only a limited amount of current into the sense/inhibit windings when saturated by the full-select current. However, the voltage level produced by this noise is insufficient to reach the threshold level during strobe time and activate the amplifier. Therefore, the data flip-flop for that bit location remains cleared, indicating a logic 0 in that location.

After the sense amplifier outputs have been strobed into the data register, the data register flip-flops are set for each bit position containing a 1. When the control logic generates a data out pulse, the high outputs (1s) from all of the flip-flops that are set in the data register are gated through drivers onto the data bus for transfer to the processor or master device. Note that the bus driver output is low for a logical 1 because the Unibus employs negative logic. Any 0 output from the data register is not gated, and a logical 0 remains on the bus for those bit positions.

This type of readout destroys the contents of the addressed memory location (by switching all cores to 0); thus, the data stored in the data register must be put back into the cores to restore the memory location to its original contents. This is accomplished by immediately following the read cycle with a write (restore) cycle. Note, however, that this write cycle is delayed whenever a DATIP operation is performed, because it is not necessary to restore the contents of the memory. In effect, the memory pauses and waits for new data to be written into the core location. The restore cycle is the same as a normal write cycle as described in Paragraph 2.9.

Figure 2-14 is a circuit schematic of the sense amplifier detector circuit. At sense time, the effect of the inhibit drivers, isolation diodes, and Balun transformer can be ignored due to the high impedance of these components.

The voltage signal is developed on the sense line and propagates through a twisted pair of wires to the sense amplifier module. The input to the MC 1540 Module is an impedance matching network used to terminate the sense line. The MC 1540 amplifies the difference signal approximately 85 times and rectifies it. The signal is ac-coupled through an external capacitor and superimposed upon a dc bias which is controlled by the threshold input. During strobe time, if the dc bias plus the coupled core signal exceeds the DTL gate threshold, its output goes to ground. This sets the 7474 flip-flop by means of a direct set gate. At the start of each read cycle, the 7474 flip-flop is reset by means of the direct reset gate. In a DATO operation, the data is shifted from the bus through the SP 380 inverter to the D input of the 7474 flip-flop and clocked into the flip-flop by a LOAD H pulse.

2.9 WRITE OPERATION

During a DATO operation, the word to be written into core memory is loaded into the data register flip-flops from the data lines when a load pulse is applied to each of the flip-flops in the register. Whenever the input from the data line is a 1, the load pulse (high) and the binary 1 set the flip-flop to the 1 state. If the input from the data line is 0, then the flip-flop remains cleared. The resultant states of the 16 flip-flops represent the word to be written into memory. During a write/restore cycle, the contents of the data register (from reading) are written.

When the write cycle begins, an inhibit signal is applied to an inhibit driver (AND gate) for each line. If the gate is qualified (flip-flop set to 0), the gate enables an inhibit pulse. This inhibit pulse causes a current path through the core opposing the Y write current to prevent a core from being set into a 1 state. If the flip-flop is set, representing a 1, the gate is not qualified, no inhibit signal is generated, and the write currents set the selected core to the 1 state. Thus, the data word stored in the register is written into the appropriate core memory location.

Figure 2-15 is a schematic of the inhibit driver. When it is desired to write a 0 and the inhibit current (I_{inh}) goes high, the 74H40 gate turns on and its output goes to ground. Current i_1 is determined by resistor R1 and the reflected base-emitter voltage (V_{be}) of output transistor Q1. Current i_2 equals i_1 and turns on transistor Q1. When Q1 turns on, current flows from ground, through the Balun transformer (which equalizes the two half-currents), through the isolation diodes and the sense/inhibit winding, to the common inhibit terminal. Current then flows through resistor R2 and transistor Q1 to -15V.

$$I_{inh} = \frac{15V - V_{cc\ sat\ Q} - V_{be\ diodes}}{16.9 + mat^{\Omega} (\approx 2.0)}$$

$$= \frac{15 - 0.8 - 1.2}{16.9 + 2.0} = 690\ mA$$

Each winding sees nearly 345 mA. Capacitor C2 speeds up the rise time of the current. At turn-off time, the inductance of the stack winding would normally push the voltage at the switch (Q1) quite positive; however, the diode at the collector of Q1 clamps this voltage to ground. Capacitor C1 helps turn on Q1 faster and allows the 74H40 gate to provide reverse current to turn off Q1.

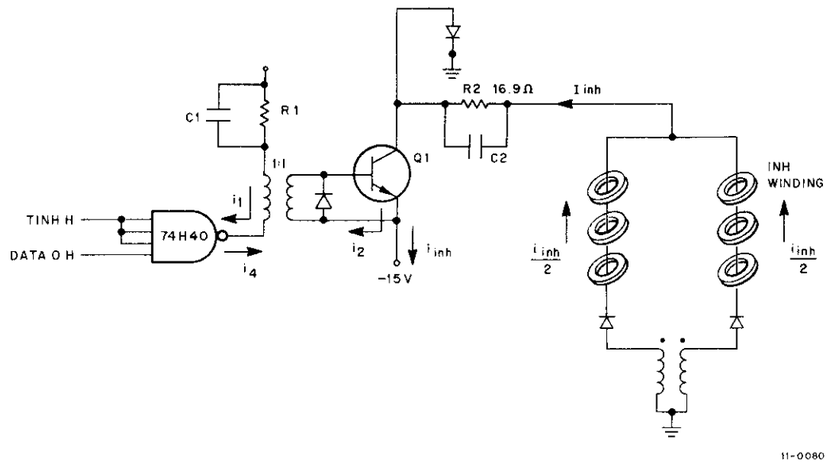

Figure 2-15 Inhibit Driver

CHAPTER 3
MAINTENANCE

3.1 INTRODUCTION

This chapter contains maintenance and adjustment procedures for the MM11-F Core Memory. Adjustment procedures are covered in Paragraph 3.2, and maintenance is covered in Paragraph 3.3.

3.2 ADJUSTMENTS

There are three adjustments that can be performed on the core memory: voltage reference (V REF), voltage slice (V SLICE), and strobe delay (STROBE DEL). It is usually not necessary to make any of these adjustments unless a module has been replaced. If, however, the memory is not functioning properly, check these adjustments before attempting to troubleshoot the memory.

The purpose of each of the three adjustments is given in Table 3-1. Specific adjustment procedures are given in subsequent paragraphs.

Table 3-1
Core Memory Adjustments

Adjustment	Location	Pin	Purpose
V REF	Trimpot on current generator Module G225	B3U2	Sets reference voltage for X-Y current.
V SLICE	Trimpot on V levels Module G103	D2F1	Sets threshold level of sense amplifier.
STROBE DELAY	STROBE DEL trimpot on control and timing Module M729	B3K1 E2S1	Sets delay time of strobe pulse used during read cycle.

3.2.1 Equipment Required

The following test equipments are used when checking or adjusting the MM11-F Core Memory.

a. Tektronix 453 or 547 oscilloscope with dual trace plug in, voltage probes

b. Honeywell 333R Digital Voltmeter (or equivalent 0.5 percent DVM)

NOTE
Although the digital voltmeter is not mandatory, it can be helpful when performing the procedures in this section.

3.2.2 Initial Procedures

Before attempting to check or adjust the core memory, perform the following steps:

NOTE
All tests and adjustments must be performed in an ambient temperature range of 20° C to 30° C (68° F to 86° F).

Step	Procedure
1	Verify that all modules shown in the MUL print MM11-F-06 are properly installed in the backboard.
2	Visually check for broken wires, connectors, modules, or other obvious defects.
3	Verify that power buses are not shorted together.
4	Turn primary power on and make certain that both -15V and +5V (±5 percent) power is present.
5	Start the system. Memory should now cycle without errors. If not, check adjustments in Paragraphs 3.2.3, 3.2.4, and 3.2.5. If memory still does not function properly, a malfunction is indicated.

3.2.3 Voltage Reference (V REF) Adjustment

The voltage reference (V REF) sets the amplitude of the X-Y current and is measured at pin B3U2 of the Current Generator Module G225. This voltage should be in the range of +0.70V to +1.90V at an ambient temperature of 30° C (86° F). The voltage fluctuates with temperature variations, going more positive as the temperature increases. If the reference voltage must be reset, perform the following adjustment procedure:

Step	Procedure
1	Cycle the memory using a worst-case patterns program.
2	Use the V REF trimpot on the current generator module to vary the reference voltage to the points where memory fails (both high voltage and low voltage). The system must have a margin of at least 0.80V, i.e., there must be an interval 0.80V wide in which the system functions properly.
3	Set the V REF at a point midway between the worst-case limits of operation. If the margin is greater than 1.0V, set V REF at 0.5V below the failing high end point. (Examples of margins are 0.4V to 1.8V set at 1.3V; 1.3V equals 1.8V minus 0.5V.) To effect this condition, adjust the trimpot on the current

(continued on next page)

Step	Procedure
3 (cont)	generator for a reference voltage that is midway between the lowest voltage causing the memory to fail at the high end and the highest voltage causing the memory to fail at the low end.

3.2.4 Voltage Slice (V SLICE) Adjustment

The voltage threshold (V SLICE) adjustment must be performed whenever a G103 Discharge Circuit Module is replaced. The adjustment procedure is as follows:

Step	Procedure
1	Measure V SLICE at pin D2F1 of the sense amplifier with respect to ground.
2	Verify that the measured voltage is -5.2V (±0.2V).
3	If the voltage is not within the above tolerance, adjust V SLICE trimpot on the G103 Module until the voltage is within the specified tolerance.

3.2.5 Strobe (STROBE DEL) Adjustment

The STROBE DEL adjustment should be checked whenever an M729 Module (Control and Timing) or a G225 Module (current generator module) is replaced. If necessary, the strobe adjustment should be reset.

NOTE

The adjustment of the STROBE DEL trimpot on the M729 Module is critical. The setting should only be changed when absolutely necessary.

Figure 3-1 illustrates the proper setting of the strobe pulse. If required, adjust the STROBE DEL trimpot so that T STROBE is set to 210 ns ±5 ns.

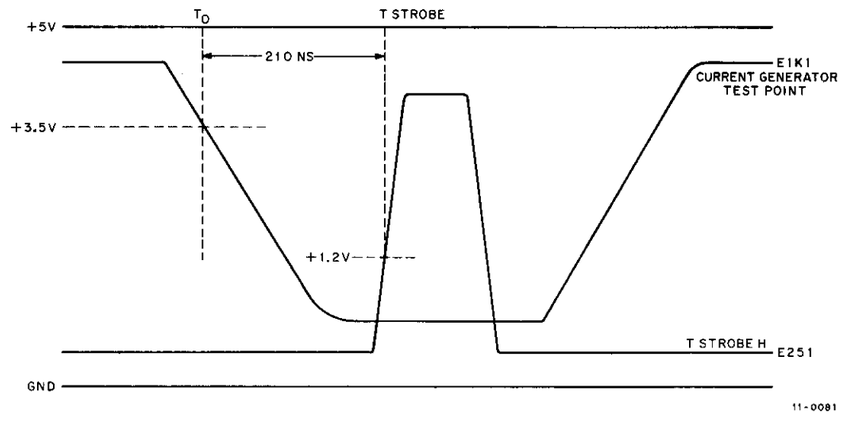

Figure 3-1 Strobe Pulse Waveform

3.3 MAINTENANCE

The basic maintenance philosophy of PDP-11 manuals presents information on normal system operation. The user can utilize this information to analyze trouble symptoms and extrapolate necessary corrective action. This paragraph provides additional maintenance information to aid the user in isolating and correcting malfunctions.

Figure 3-2 illustrates the sense/inhibit waveforms; Figure 3-3 illustrates drive waveforms. Both figures include schematics to indicate the points in the circuit where the waveforms occur. In addition to the normal waveforms, dotted lines are used to indicate the waveform that appears if a component is faulty.

Table 3-2 is a troubleshooting chart that indicates possible malfunctions, probable causes, and corrective actions.

Paragraph 3.4 describes DEC program tests that serve as an aid to maintenance and troubleshooting.

Table 3-2
MM11-F Troubleshooting Guide

Symptom	Affects	Possible Cause(s)	Corrective Action
Picks up bits	All bits	X-Y current too high	Check V REF and reset per set-up procedure.
	All bits	Strobe occurs too early	Reset STROBE DEL to 210 ns on M7290 Module.
	All bits	-15V input low	Reset or repair H720 Power Supply
	One bit	Inhibit driver inoperative; inhibit winding open	Repair G102 Module. Repair G108 Module.
	All bits	Threshold too low (less than -5.2V)	Reset V SLICE to -5.2V on G103 Module.
	Some bits in one byte	Sense amplifier reset; inoperative for that byte	Repair M7290 Module
	8 or 16 bits	TINH H inoperative	Check M7290, G225 Module
	All bits at high temp.	Stack thermistor or resistor	Check components and wiring to G225 Module.
Drop bits	All bits	X-Y Current too low	Reset V REF or repair G225 Module.
	All bits	Strobe too early or too late	Reset STROBE DEL on M7290 Module.
	All bits	Threshold too high	Reset V SLICE on G103 Module
	All bits, 64 locations	Open line in memory	Ring out stack (read or write)
	All bits, 8 locations	Bad switch or driver	Troubleshoot and repair G226 Module.
	One bit, all words	Bad sense line or sense chain (bus receiver, latch, sense amplifier, or bus driver)	Troubleshoot and repair G102 Module.
	Some bits, one byte	STROBE inoperative, LOAD inoperative for that byte	Troubleshoot and repair M7290 Module.

(continued on page 3-5)

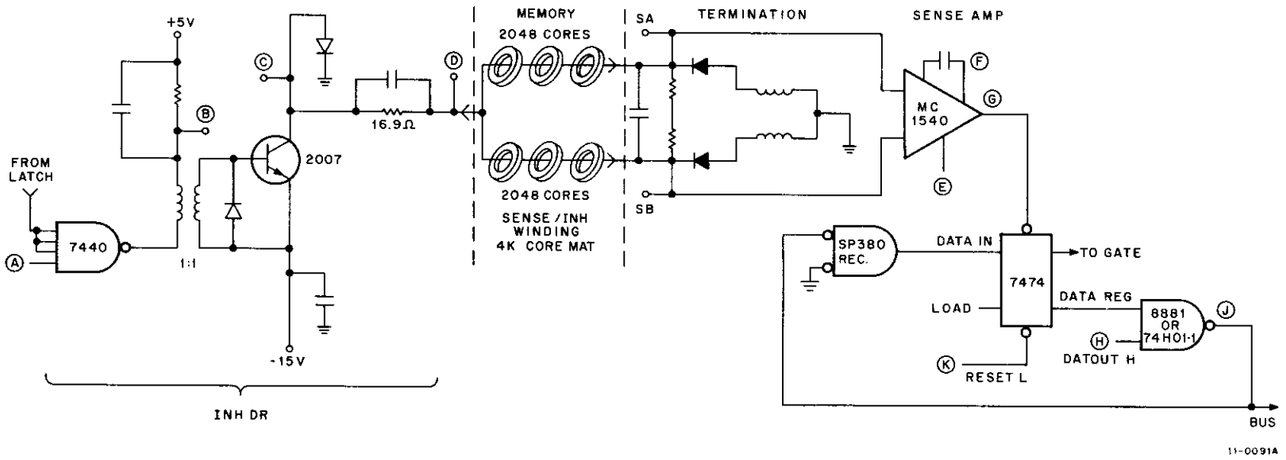

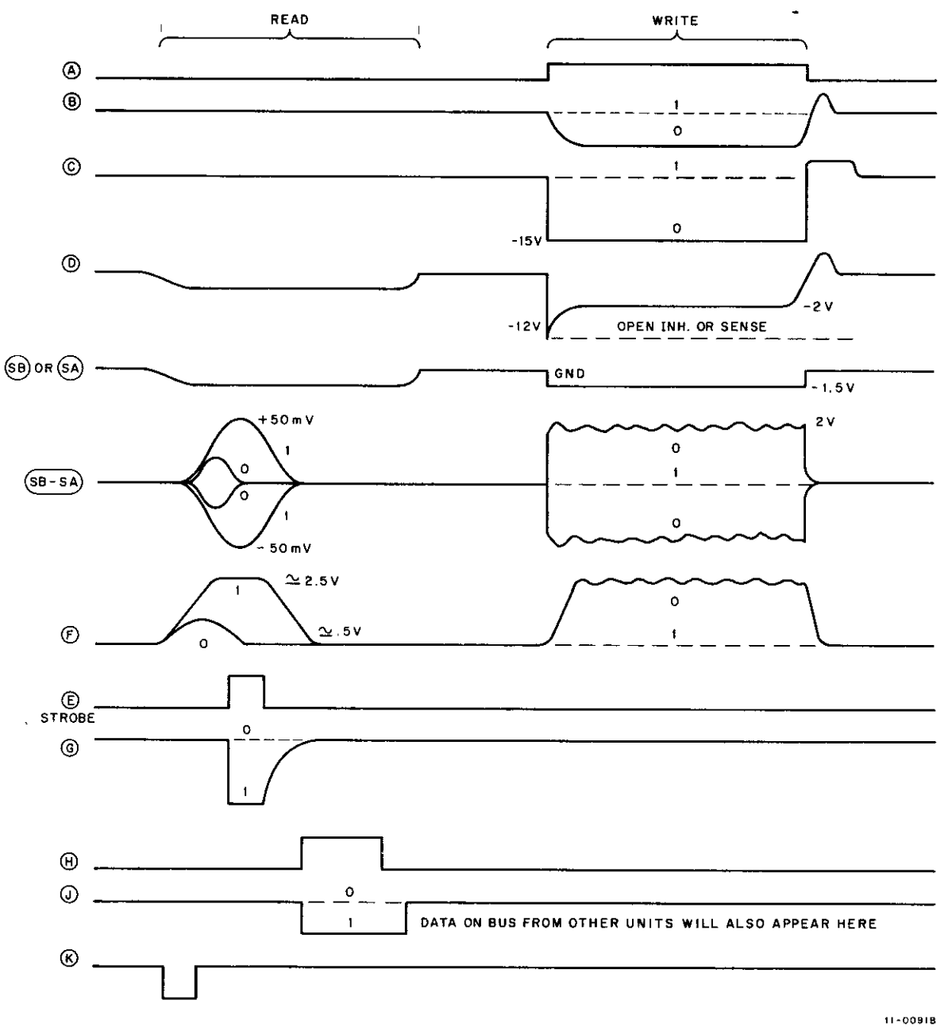

Figure 3-2 MM11-F Sense/Inhibit Waveforms

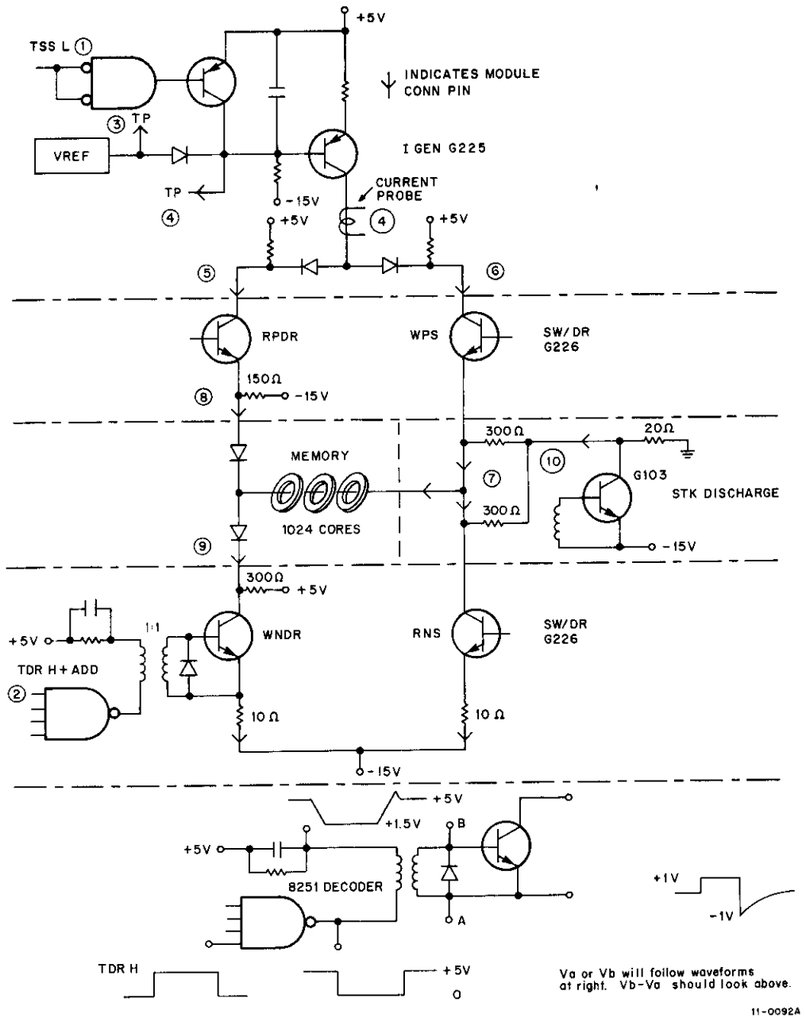

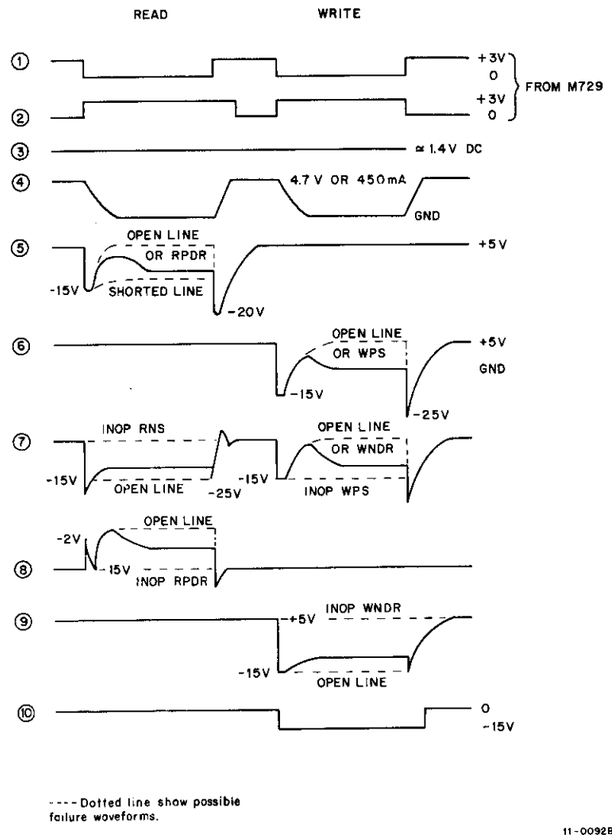

Figure 3-3 Drive Waveforms

Table 3-2 (Cont)
MM11-E Troubleshooting Guide

Symptom	Affects	Possible Cause(s)	Corrective Action
Drop bits (cont)	All bits at low temp.	Stack themistor or resistor	Check components and wiring to G225 Module.
No response to MSYN L		No device select	Check M1091 Module jumpers and gates.
		MSEL not reset on previous cycle	Check M7290 Module circuits.
		SSYN not reset on previous cycle	Check M7290 Module circuits.
		SSYN not setting	Check M7290 Module circuits.
Memory always performs a DATIP	All bits	PROTECT L input to M7290 Module grounded	Check and repair.

3.4 PROGRAMMING TESTS

Certain DEC programs can be used to test various memory operations as an aid to troubleshooting. The purpose of each of these memory-related test programs, as well as the program abstract, is given in the following paragraphs.

3.4.1 Address Test Up (MainDEC-11-DIAA)

The purpose of the Address Test Up program is to demonstrate that the selected memory area is capable of basic read and write operations when address propagation is upward through memory.

This test program writes the address of each memory location (within the test limits) into itself and then increments through memory until the address corresponding to the high limit is reached. After this location has been written, the memory enters the read cycle. The read cycle starts with the high limit location and reads and compares each word location, decrementing down to the low limit location. The program halts on an error.

This program checks that all addresses are selectable and can also be used to isolate bad switches, wiring errors, or address selection errors.

3.4.2 Address Test Down (MainDEC-11-DIBA)

The purpose of the Address Test Down porgram is to demonstrate that the selected memory area is capable of basic read and write operations when address propagation is downward through memory. It is a companion test to the Address Test Up program (Paragraph 3.4.1).

This test program writes the address of each location into itself, downward through memory. After writing down, the program reads and checks back up through the memory test area. The program halts on an error.

The Address Test Down program resides in the high portion of core memory. It does not check memory below address 100, as these locations are reserved for trap and vector locations. The program verifies that all modules can perform their basic functions, checks that all addresses are selectable, and can also be used to isolate faulty switches, wiring errors, or address selection errors.

3.4.3 No Dual Address Test (MainDEC-11-DICA)

The purpose of the No Dual Address Test program is to check the unique selection of each memory address tested.

This test is divided into two parts. The first portion of the test fills the test field with 1s and writes 0s into the first test location. This is followed by a read check from this location. The program then checks each field location to ensure there are no variations from the 1s configuration. Upon completion of this test, the test location pointer is incremented. The next location is then write-read exercised with 0s and the test field rechecked for any change in content. When the selected test field has been tested in this mode, the program sets a flag and the second portion of the test is begun. The program fills the test field with 0s and the field is then tested with a write-read exercised with 1s.

This program checks for faulty switches or wiring errors, checks the complete address selection scheme, and checks all 16 bits in the data field for 1s and 0s operation.

3.4.4 Basic Memory Patterns Test (MainDEC-11-DIDA)

The Basic Memory Patterns Test program has two main purposes:

a. Verify that the selected memory test field is capable of writing and reading fixed data patterns.

b. Verify that the memory plane is properly strung.

This test program writes a specific pattern throughout a given memory zone, then reads the pattern back and compares it with the original for correctness. If the pattern read fails to compare correctly with the original, the program initiates a call to the error subroutine. After completely checking the pattern, the program continues on to the next pattern test.

3.4.5 Worst Case Noise Test (MainDEC-11-DIGA)

The purpose of the Worst Case Noise Test program is to generate the maximum possible amount of plane noise during execution of memory reference instructions to check system operation under worst case conditions.

This test program is designed to produce the greatest amount of plane noise possible during memory read and write cycles. The noise parameters are effected by a number of factors; therefore, the test is designed to test all conditions arising out of the four standard memory construction configurations. The noise generated is distributed across the core plane algebraically and adds to the normal dynamic noise present on the sense lines. This can cause misreading of data (within the plane) that is in the low (1) or high (0) category. The sense windings of most memories are such that worst case patterns can be caused by alternately writing -1 and 0 data configurations throughout memory. Under these conditions, worst case noise is generated by performing a read, write, complement, read, write, complement, operation at each location. The test is repeated after complementing all of the pattern data stored in the memory test zone.

This test deviates from this form of testing only in the distribution of the -1 and 0 data patterns within the memory test zone. The constraint placed on this test requires writing the complement of the data pattern as described by the exclusive OR of the second bit of the X and Y selection lines. These bits correspond to address bits 1 and 8. Therefore, the pattern, or its complement, is written into the memory test zone as determined by the exclusive OR between address bits 1 and 8.

The Worst Case Noise Test program is divided into two parts. Part 1 is run first and, during this part of the program, a -1 configuration is written into all locations having an address with an exclusive OR state between bits 1 and 8. All other locations are loaded with the 0 configuration. After the test zone has been loaded, the memory is rescanned. This time, each location is read, complemented, read, and complemented (RCRC). Any location

detected as being disturbed by a previous RCRC operation is flagged as an error. Upon conclusion of the read scan loop, the program automatically switches to Part 2.

During Part 2 of the program, the data patterns stored in memory are complemented. In other words, 0 patterns are stored in locations having addresses with an exclusive OR between bits 1 and 8. All other locations are loaded with the -1 configuration.

The exclusive OR pattern distribution for Parts 1 and 2 is summarized for reference as follows:

Part 1

Exclusive OR (1 and 8) = -1 pattern
No Exclusive OR (1 and 8) = 0 pattern

Part 2

Exclusive OR (1 and 8) = 0 pattern
No Exclusive OR (1 and 8) = -1 pattern

After memory is loaded, it is scanned again with a read, complement, read, complement (RCRC) loop as described previously. Any location detected as being disturbed by a previous RCRC operation is flagged as an error.

Before writing or reading any location (in either part of the program), the program issues a call to subroutine XORCK (exclusive OR check) which tests bits 1 and 8 and sets the XORFLG if the exclusive OR condition is present.

Subroutine ERRORA is called for any location disturbed from the -1 configuration; subroutine ERRORB is called for any location disturbed from the 0 configuration.

The program prints out errors and repeats when complete without interruption. Upon completion, the program rings the Teletype® bell and then halts if switch 12 is present. A continue from the halt initiates another pass.

If the program indicates an error, the failure may be due to a faulty memory stack, not enough twist in the SA-SB twisted pairs, a faulty sense/inhibit card, incorrect strobe setting, incorrect threshold setting, incorrect current setting, or power inputs that are out of tolerance.

NOTE

In any interleaved memory, the bits that determine the worst case pattern are bits 8 and 13 rather than bits 8 and 1.

®Teletype is a registered trademark of Teletype Corporation.

APPENDIX A
PARITY OPTION

A parity option is available for use with the MM11-F Core Memory. When this option is installed, the memory is an 18-bit memory with bit 16 used for byte 0 parity and bit 17 used for byte 1 parity. Installation of this option consists of replacing the H207 Memory Stack with the H210 Memory Stack and installing the G108 Sense/Inhibit Module into location E1.

All MM11-F Core Memories come wired and documented as 18-bit systems. Therefore, only the H210 stack change and the addition of the G108 are needed to convert the system to an 18-bit system with parity capabilities.

The H210, which is a 19-bit memory stack, is pin compatible with the H207 16-bit memory stack. Note that one bit of the H210 is not used so that it functions as an 18-bit memory stack.

The two-bit sense/inhibit module enables the memory system to read and write parity data into and out of memory. The data bits on the bus are called BUS DPB0 and BUS DPB1. Stack bits 17 and 18 are used for this function. Note that one bit (DPB0) on this module has common strobe, reset, and load with byte 0. The other bit is common to byte 1.

APPENDIX B
INTERLEAVING

When the address register is incremented by one (alternate states of BA01 L) on successive memory cycles, the cycles are performed within a 4K memory bank and cannot be overlapped. If, however, address bit BUS A01 L (inside the memory) is interchanged with BUS A13 L, it causes successive memory cycles to be performed within alternate memory banks. This allows the cycles to be overlapped; that is, the second memory bank can start its cycle before the first memory bank has completed its cycle, provided the bus is free. This effect is called memory interleaving and causes the typical program to be executed faster.

Note that interleaving affects 8K blocks. For example, if a system has a 12K memory, only the first 8K is interleaved. If the system had 16K of memory, the first 8K would be interleaved and the second 8K would also be interleaved.

Interleaving an 8K block in the MM11-F core memory is accomplished by changing the jumpers on the M1091 module on the Device Select Module, M1091, drawing B-CS-M1091-0-1.

It should be noted that any 8K block of memory delivered from DEC is automatically interleaved.