DO NOT REMOVE FROM LIBRARY

A COPY FROM
MITRE ARCHIVES

THEORY OF OPERATION

OF

CENTRAL COMPUTER SYSTEM

FOR

AN/FSQ-7
COMBAT DIRECTION CENTRAL

AND

AN/FSQ-8
COMBAT CONTROL CENTRAL

VOLUME II

1 February 1959

This document contains information of a proprietary nature. Any use or reproduction of this document for other than government purposes is subject to the prior consent of International Business Machines Corporation.

MILITARY PRODUCTS DIVISION
INTERNATIONAL BUSINESS MACHINES CORPORATION
KINGSTON, NEW YORK

PART 7
MEMORY ELEMENT

CHAPTER 1
INTRODUCTION

1.1 GENERAL

The memory element of the Central Computer is a computer-controlled, large-capacity, random-access, high-speed storage facility which provides for the semipermanent storage of all information (operating program, raw data, and processed results) required for or resulting from the normal operation of the Central Computer. Since the memory storage circuits do not differentiate between instruction words and data words (raw or processed), the theory of operation presented in this part is based on the two types of memory cycles that affect memory operations; that is, the memory readout and memory store cycles. The memory readout cycle (equivalent to PT, OTA, and BO machine cycles) results in the transfer of a memory word out of a specific memory location; the memory store cycle (equivalent to OTB and BI machine cycles) results in clearing the specified memory location and storing the desired memory word in the cleared location. Each type of memory cycle requires 6.0 μsec for execution.

1.2 BLOCK DIAGRAM ANALYSIS

A block diagram of the memory element is shown in figure 7–1. Although the three memory address registers are shown as individual blocks, each register is actually an integral part of the associated memory device since it is used to condition the internal address selection circuits within that memory. Each of the three memory devices is capable of storing different quantities of information: core memory 1 has a storage capacity of $65{,}536_{(10)}$ memory words, core memory 2 has a storage capacity of $4{,}096_{(10)}$ memory words, and test memory has an effective storage capacity of $16_{(10)}$ memory words. The storage medium used in each of the core memory devices is a 3-dimensional ferrite core array. Test memory storage is accomplished by the use of a manually wired control panel, two toggle switch registers, and a flip-flop register.

During computer operation, each memory cycle is initiated by the transfer of the desired memory address to the three memory address registers, and the subsequent application of a start-memory pulse to initiate the internal operations in the selected device. Concurrently, the memory buffer registers are cleared to prepare them for the temporary storage of data to be transferred from either the selected memory location or from an external register, depending upon the type of cycle in process. During the execution of a memory

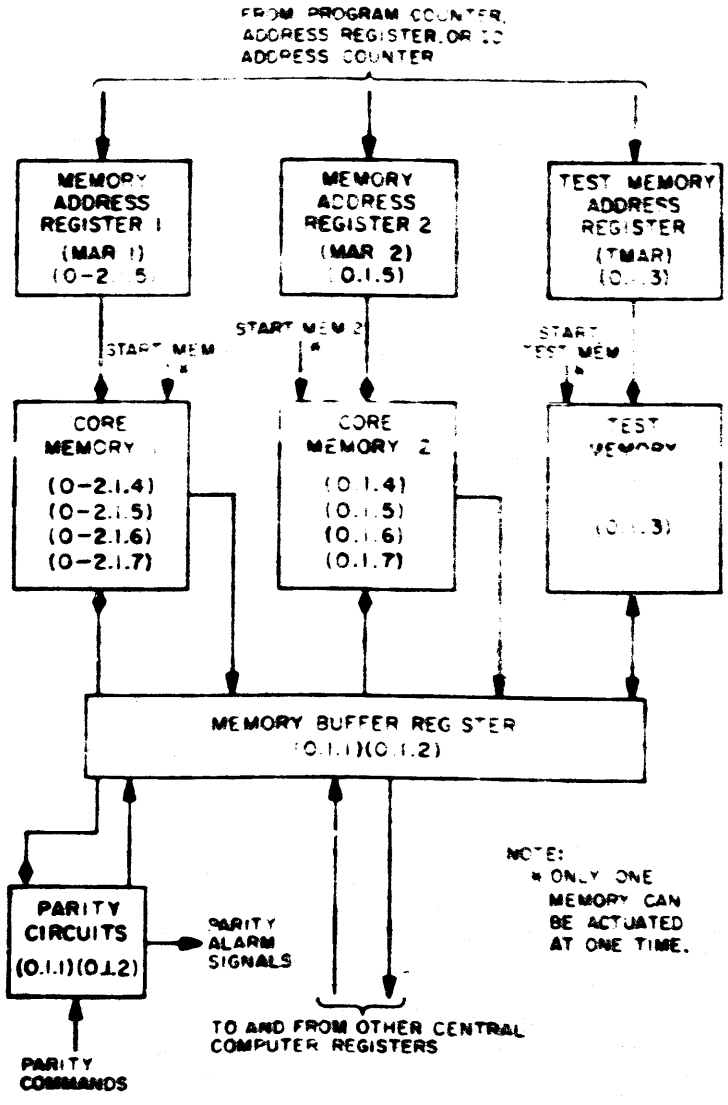

Figure 7–1. Memory Element, Block Diagram

readout cycle, the contents of the specified memory location are transferred to the memory buffer register approximately 3.0 μsec after the cycle was initiated. The memory buffer register contents are then transferred to a specific computer register as directed by a computer command. Since the core memory devices operate on the principle of destructive readout (the information contained in the selected core register is erased), the latter portion of the memory readout cycle is used to store the memory buffer contents back into the specified core memory location. During the execution of a memory store cycle, the first 3.0 μsec are used to erase the content of the specified memory location and to transfer the new word from one of the computer registers to the memory buffer register. During the latter portion of this memory cycle, the new information is stored in the cleared memory location.

The parity circuits associated with the memory buffer register provide a means of checking the accuracy of information transfer into and out of either of the core memory devices. During each memory store cycle, a parity bit is assigned to the memory word (33rd bit of memory word) before it is stored into the specified core memory location. During each memory readout cycle, the parity bit is checked to determine whether the initial transfer into and the present transfer out of the core memory were accurately accomplished. If the parity check shows that the memory word is in error, a parity-alarm signal is generated to inform the operating personnel of the malfunction.

Since the memory storage devices and the parity circuits operate independently of each other, the detailed analysis of each of these circuit groups is presented in separate independent chapters of this part. Because the operating characteristics of the two core memory devices are identical, the principles of core memory storage are discussed separately in Chapter 2. This chapter, which analyzes in detail the characteristics of ferrite cores and ferrite array wiring, contains only a minimum amount of core memory circuit analysis. Conversely, Chapters 3 and 4, which contain the detailed theory of operation of the 64^2 and 256^2 core memory devices, respectively, do not include an analysis of the principles of ferrite core storage and are limited to an analysis of the logic circuits required to produce the desired results.

CHAPTER 2

PRINCIPLES OF CORE MEMORY OPERATION

2.1 ANALYSIS OF FERRITE CORES AND FERRITE CORE ARRAYS

2.1.1 Characteristics of Ferrite Cores

The principal component of a magnetic core memory storage device is a ferrite core which possesses a square hysteresis loop, a low coercive force, and a short flux reversal or switching time. The ferrite core used in the two core memory devices is composed of a mixture of ferric and manganese oxide powders which are bonded together in the form of a toroid having an inside diameter of 0.050 inch (0.127 cm), an outside diameter of 0.080 inch (0.203 cm), and a thickness of 0.025 inch (0.0635 cm). A carefully controlled sintering or firing process imparts the desired characteristics to the core.

The usefulness of a ferrite core as a binary storage device depends upon four important characteristics: the ability of the core to remain in one of two stable magnetic states, the squareness ratio, the switching time, and the ratio of coercive force to applied field required to produce the major or saturation hysteresis loop. The major hysteresis loop of a typical ferrite core, which is a plot of flux density (B) versus applied magnetic field intensity (H), is shown in figure 7-2.

Storage of binary information in a ferrite core is dependent on the ability of the core to retain a relatively large value of residual or remnant flux upon the termination of a driving or switching pulse of current. As shown in figure 7-2, two primary states of residual or remnant flux density (B_r and $-B_r$) are possible, and these have been defined as the zero and one states respectively. To switch a core from one state to the other, it is necessary to apply a current pulse of I_w or I_r ma (depending on the initial content of the core) to the drive line that links the core. The resultant applied field of 820 ma-turns will saturate the core in the desired direction, causing the flux state to reverse by traversing the loop path a-b-c or c-d-a.

During core memory operation, the state of a core is determined by applying a read current pulse (I_r) to the drive line, and detecting the resultant changes in flux density by measuring the voltage induced in the sense winding. If the core was initially in the one state, application of the read current pulse (I_r) will cause the core to be switched to the zero state (traversing path a-b-c in fig. 7-2), with the result that a relatively large change in flux density ($2B_r$) will be detected by the sense winding (one output signal). If the core was initially in the zero state, application of the read current pulse (I_r) will not switch the core (path c-b-c in fig. 7-2 is traversed), with the result that a relatively small change in flux density will be detected by the sense winding

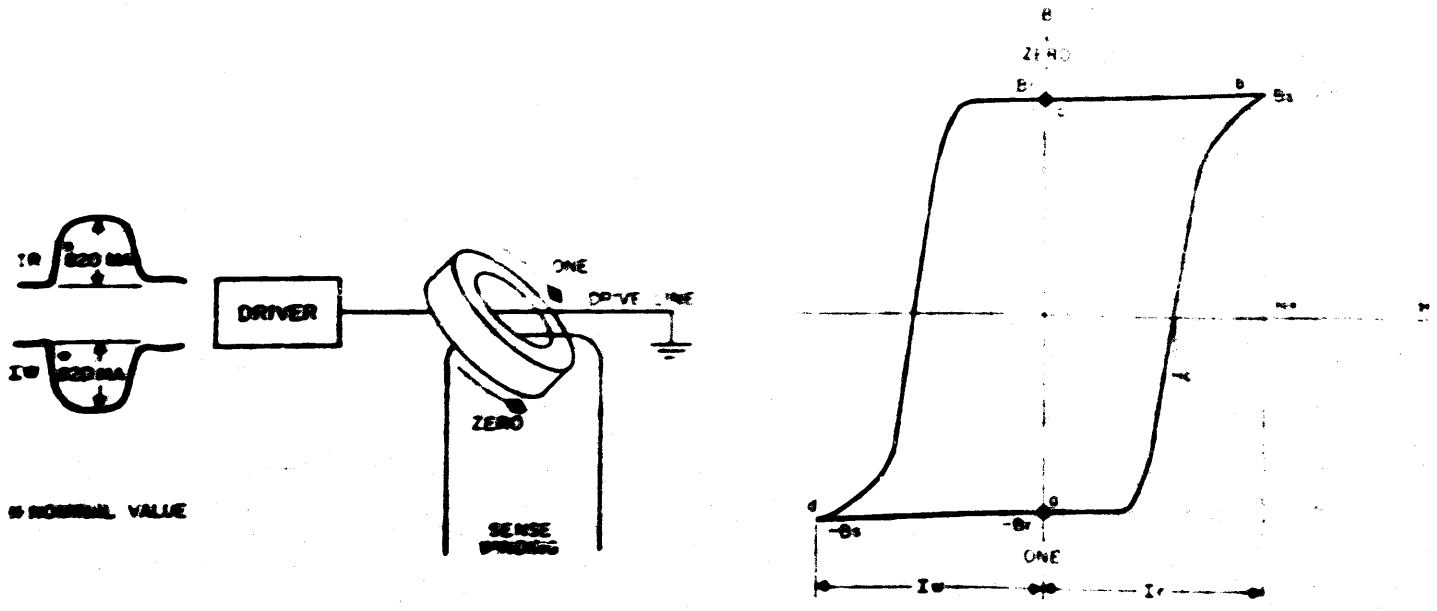

Figure 7-2. Major Hysteresis Loop

ing (zero output signal). The amplitude of the zero output signal is dependent upon the ratio of the residual flux density (B_r) to the saturation flux density (B_s), which is defined as the squareness ratio ($R_s = B_r/B_s$) of the core. This ratio is important in that it determines the relationship between the amplitudes of the one and zero output signals. The magnitude of the zero output signal decreases as the value of R_s approaches 1, with the result that the difference between the amplitudes of the one and zero output signals becomes greater. Ferrite cores used in the two core memories have a squareness ratio greater than 0.95.

The time required to produce a flux change or flux

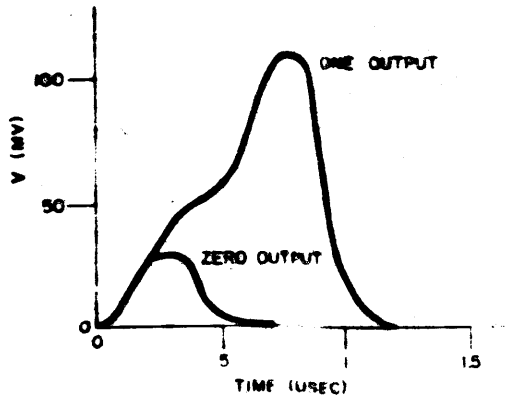

Figure 7–3. Typical Response of a Ferrite Core

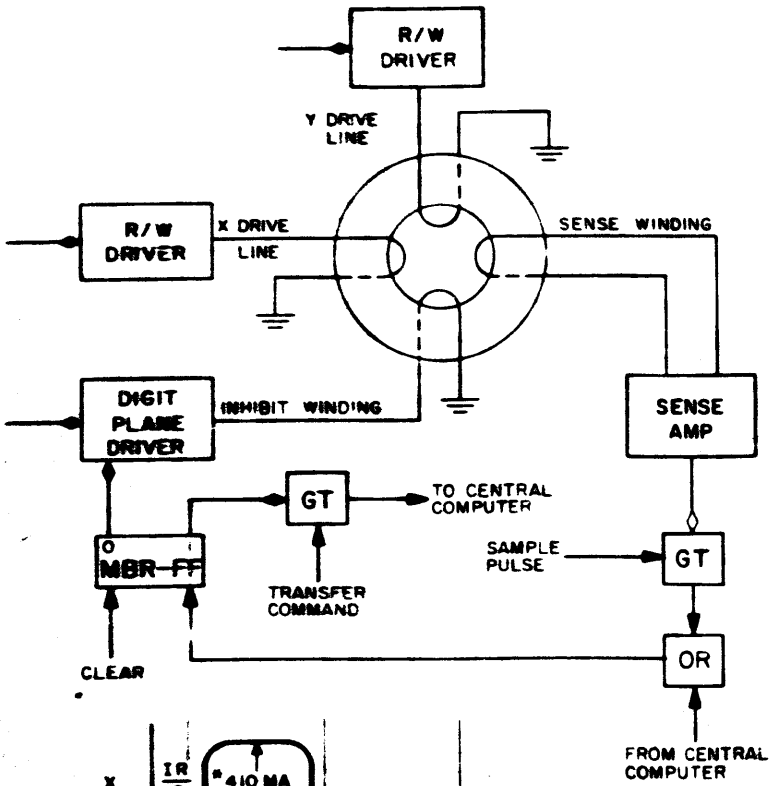

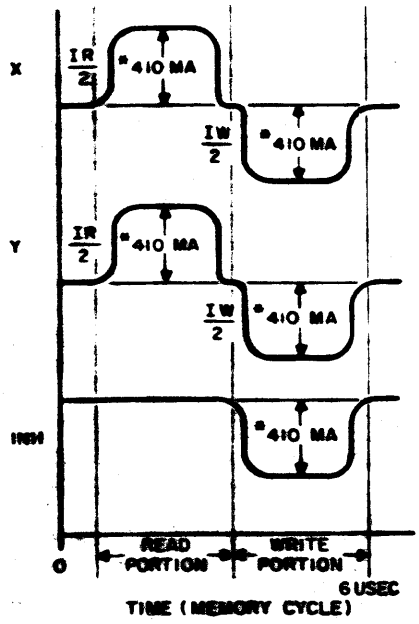

* NOMINAL VALUE

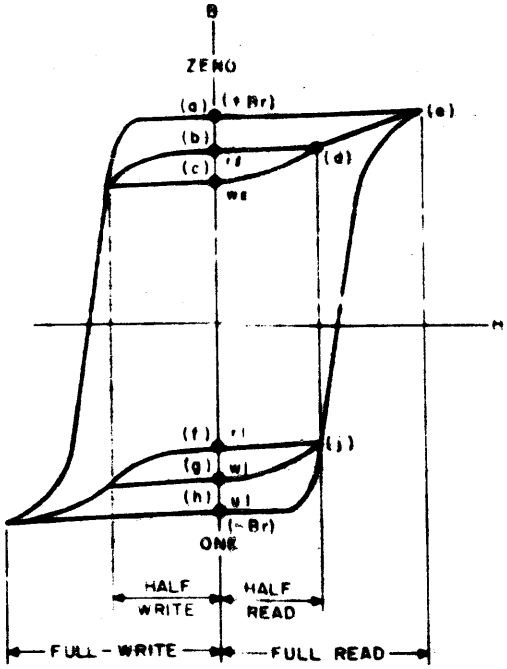

rz - READ DISTURBED 0
wz - WRITE DISTURBED 0
r1 - READ DISTURBED 1
w1 - WRITE DISTURBED 1
u1 - UNDISTURBED OR NEWLY WRITTEN 1

reversal is of equal or greater importance than the squareness ratio. Figure 7–3 shows the typical output signals produced as the result of a read current pulse (I_r) being applied to a core in the one and zero states. As noted in the figure, the zero output signal is not only smaller in amplitude than the one signal, but also peaks earlier and is of shorter duration. Since these differences exist, it would appear that information detection could be accomplished by either a time-amplitude or a difference-amplitude sampling technique. However, in a practical core memory, a time-amplitude sampling technique must be used because the noise signals generated in the memory (discussed in 2.4) make the difference-amplitude sampling technique unfeasible. The time required for a core to switch from one state to the other is defined as the switching time, T_s. This time is actually measured as the elapsed time between the time that the driving current pulse reached 50 percent of its amplitude, and the time at which the one output signal has dropped to 10 percent of its peak amplitude. The switching time for cores used in the memories is approximately 1.2 μsec for an applied field of 820 ma through a 1-turn drive winding.

The coercive force (H_c) of a magnetic material is defined as the maximum value of field intensity that can be applied without causing a change in flux polarity. The ratio of coercive force (H_c) to saturation field intensity (H_s) is as important to proper memory operation as the squareness ratio (R_s) and the switching time (T_s). Since the coercive force (H_c) of the memory core in use is greater than one-half of the saturation field intensity (equivalent to $I_s/2$ and $I_{-s}/2$ in fig. 7–4), it is not sufficient to cause the core to switch from one state to the other. As a result of this characteristic, the core switching function can be and is, performed by two separate 1-turn drive lines, each actuated with $\frac{1}{2}$ the drive current required to switch a core. As noted in figure 7–4, the two single-turn drive lines, labeled X and Y, are wound through the core in the same direction so that the individual magnetic fields set up by these two windings will add directly when the drive lines are energized by current pulses of the same polarity. Since each drive line is supplied with half-amplitude current pulses ($\frac{1}{2}$ the current required to switch the core), both drive lines must be pulsed in coincidence in order to switch a core. A study of the hysteresis loop reveals that application of individual half-amplitude current pulses of either polarity does produce small changes in the flux density of a core (in fact, alternate polarity pulses produce the four remnant states of the two minor hysteresis loops); however, under no condition will a half-amplitude current pulse cause the core to switch from one state to the other. It is this important characteristic of the ferrite core that makes it possible to design a ferrite memory array that will operate on a coincident current method of address selection.

2.1.2 Characteristics of a Ferrite Core Array

Figure 7–5 shows the square arrangement of ferrite cores in a memory plane (one of the 33 horizontal layers of a ferrite core array) which contains the simi-

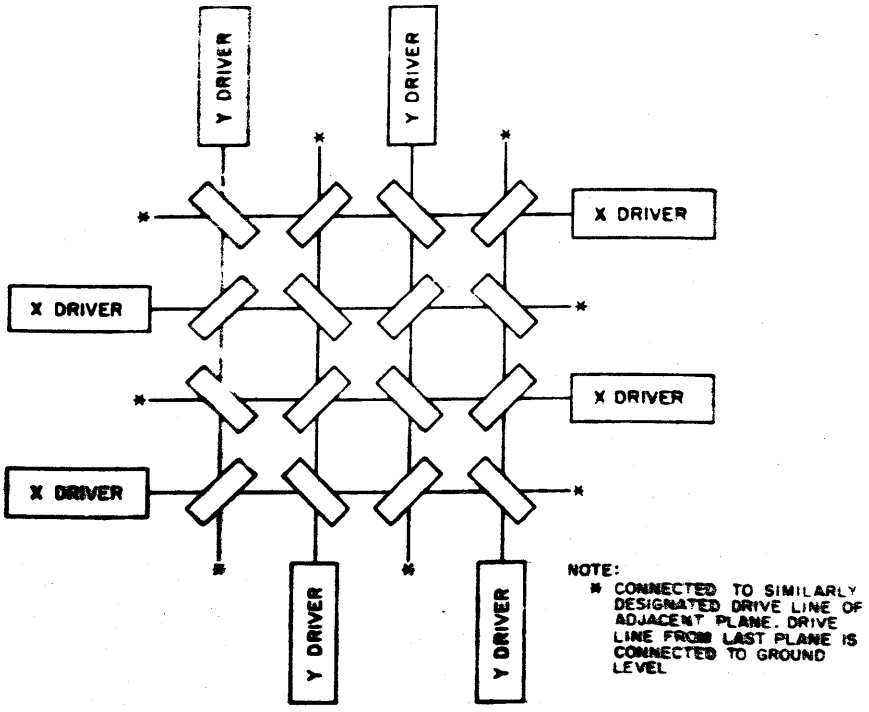

larly designated bit of all the memory words stored in a core memory device. Although the drawing shows only the X and Y drive line wiring of 16 cores in a 4 x 4 matrix, it is representative of the core memory 2 plane, which contains 4,096 cores in a 64 x 64 matrix, and the core memory 1 plane, which contains 65,536 cores in a 256 x 256 matrix. In addition to being linked by the individual X and Y drive lines, all the cores of a memory plane are linked by two additional wires (not shown), a sense winding and an inhibit winding.

In constructing a ferrite core array, the 33 memory planes that comprise the array are stacked vertically and the X and Y drive lines of each plane are connected in series with the similarly positioned drive lines of adjacent planes so that one current driver can be used to control the associated X or Y drive line of the 33 planes. That is, the individual X and Y drive lines are series-connected so that one current driver will affect the similarly positioned row or column of cores in each of the 33 planes. As a result of this wiring scheme, each pair of mutually perpendicular drive lines (one X and one Y) intersects at the same point in each of the 33 planes to mutually link the similarly positioned core in each of the 33 planes of the array. That is, each pair of X and Y drive lines mutually links a specific vertical column of 33 cores (one in each plane).

During core memory operation, selection of a specific memory register (a vertical column of 33 cores) is accomplished by conditioning one X and one Y current driver. Subsequently, the two conditioned drivers are actuated so that they will simultaneously generate half-amplitude read and half-amplitude write current pulses in the sequence noted in figure 7–4. Since the two drive lines that link a core (fig. 7–5) are wound through the core in the same direction (with reference to the associated driver), the individual magnetic fields set up by these two selected drivers will add at the point of intersection. Because each current driver can only supply half-amplitude current pulses, only the cores at the 33 intersections of the two selected drive lines (one core in each plane) will be affected by sequentially applied magnetic fields of plus and minus 820 ma-turns (fully selected); all the other cores on these two drive lines will be affected by sequentially applied magnetic fields of plus and minus 410 ma-turns (half-selected). The sequential application of magnetic fields of plus and minus 410 ma-turns will not produce any change in the status of any of the half-selected cores (fig. 7–4). However, the sequential application of magnetic fields of plus and minus 820 ma-turns to the 33 selected cores (one core in each plane) will cause each of these cores to be switched first to the zero state and then to the one state. Since normal memory operation may require that individual cores of the selected memory register (one core in each plane) remain in the zero state at the end of a memory cycle, a separate control is provided for each plane which when actuated will prevent the minus 820 ma-turn applied field from switching the associated selected core to the one state. This control to inhibit the writing of a 1 in the selected core of an individual memory plane is accomplished by a digit plane driver. When conditioned, the digit plane driver can be actuated to generate a negative current pulse of 410 ma on its associated one-turn inhibit winding in coincidence with the half-amplitude write current pulses that are applied to the selected X and Y driver lines. The inhibit winding, which links all the cores of a single memory plane, is wound so that its applied field opposes the individual magnetic fields set up by the X and Y drive lines when the latter are energized by half-amplitude write current pulses.

As a result, the coincident application of these three current pulses to the selected core of a plane will produce a resultant magnetic field of minus 410 ma-turns, which is not sufficient to switch the selected core from the zero to the one state. Although the inhibit current pulses also affect all the other cores of the plane, it will not produce any change in the status of any of these cores. The detailed analysis of how the digit plane driver is controlled to perform its function is given in 2.2.

2.2 CORE MEMORY CYCLE ANALYSIS

During normal computer operation, either of the two core memory devices can be individually selected to perform either of two distinct functions: to supply a previously stored word to the computer (readout cycle) or to store a new word supplied by the computer (store cycle). For either of these two cycles, the Central Computer must specify which particular memory register (vertical column of cores) is to be involved in the subsequent memory operation and in which direction information is to be transferred. In addition, the Central Computer must also generate a start-memory pulse to actuate the internal memory circuits. As a result, at the beginning of every core memory cycle, the following sequence of events is performed by computer commands:

a. At TP 0, the memory address register, the memory buffer register, and the internal memory control circuits are reset to a neutral or starting condition.

b. At TP 1 (approximately), the desired address information is transferred to the memory address register, which functions to condition the pair of current drivers (one X and one Y) that are to be actuated during the subsequent portion of the cycle.

c. At TP 1 (delayed), a start-memory pulse is applied to the memory pulse distributor (a tapped 5.5-μsec delay line) which functions to control all the internal memory operations required to complete the memory cycle.

In performing its function, the memory pulse distributor provides delayed start-memory pulses to actuate the conditioned X and Y current drivers so that they will both generate a half-amplitude read and a half-amplitude write current pulse in the sequence noted in figure 7—4. In addition, delayed start-memory pulses are also supplied to the digit plane driver controls so that each conditioned digit plane driver will generate an inhibit current pulse in the sequence noted in figure 7—4. Because the read, write, and inhibit current pulses are generated in the same sequence for every core memory cycle, the analysis of internal memory operations consists of defining the function of the read and write portions of the cycle.

Since all memory cores function in the same manner, the following discussion on the details of the memory readout and store cycles will be based on the single core shown in figure 7—4, which represents the core at the intersection of the selected X and Y drive lines of one memory plane.

2.2.1 Memory Readout Cycle

The read portion of a core memory readout cycle is used to determine the status of the selected core. Since the magnetic fields set up by the selected X and Y drive lines add at their point of intersection, a magnetic field of +820 ma-turns will be applied to the selected core upon the generation of the coincident (X and Y) half-amplitude read current pulses ($I_r/2$). If the selected core initially contains a 1, the applied field will cause the core to switch its magnetic state and the resultant flux change (approximately $2B_r$) will induce a one output signal (fig. 7—3) in the 1-turn sense winding that links all the cores of the associated memory plane. This signal, when amplified by the sense amplifier, will condition the associated gate tube. The gate tube is sensed by a sample pulse (delayed start-memory pulse) at a specific time (that is, at the peak of the amplified 1-output signal), and the resulting output pulse is used to set the associated memory buffer register flip-flop to the 1 state.

If the selected core initially contains a 0, the applied field cannot cause the core to switch, although a small flux change will result which induces a zero output signal (fig. 7—3) in the sense winding. However, since this induced signal is small, the associated gate tube will not be conditioned when sampled, and the memory buffer register flip-flop will remain in the zero state.

Thus, at the end of the read portion of the memory readout cycle, the selected core is in the zero state regardless of the initial contents of the core (destructive readout), and the associated memory buffer register flip-flop contains the initial contents of the selected core. The contents of the memory buffer flip-flop are then transferred to the specified Central Computer register by the appropriate computer-generated transfer command. Since a memory readout cycle requires that the initial contents of the selected core be preserved for future reference, the write portion of the cycle is used to restore the selected core to its initial state.

Since the memory buffer register flip-flop contains the information to be written into the selected core, this flip-flop is used to control whether the selected core will be switched to the one state or be made to remain in the zero state. This control is accomplished by supplying either a conditioning or a deconditioning level to the associated digit plane driver which, when actuated, will generate a negative current pulse of 410 ma on the 1-turn inhibit winding that links all the cores of the associated memory plane. If the memory buffer register flip-flop contains a 1, its zero side output (−30V level) will decondition the digit plane driver so that an inhibit current pulse will not be generated during the write portion of the cycle. Under this condition, the coincident (X and Y) half-amplitude write current pulses ($I_w/2$) supplied to the selected X and Y drive lines, which result in the application of a magnetic field of −820 ma-turns to the selected core, will cause the selected core to be switched to the one state.

If the memory buffer register flip-flop contains a 0, its zero side output (+10V level) will condition the digit plane driver so that an inhibit current pulse will be generated in coincidence with the half-amplitude write current pulses. Since the inhibit winding is wound so that its applied field will oppose the individual magnetic fields set up by the X and Y drive lines when the latter are energized by half-amplitude write current pulses, the inhibit current pulse, although actually negative, has the same effect as a half-amplitude read current pulse which is positive. Under this condition, the magnetic fields set up in the selected core by the three coincident current pulses will add algebraically to produce a net applied field of −410 ma-turns (equivalent to half-write), which is not sufficient to switch the core to the one state. As a result, the selected core remains in the zero state.

2.2.2 Memory Store Cycle

The memory store cycle is very similar to the memory readout cycle in that the read, write, and inhibit current pulses are controlled and generated in exactly the same manner. In fact, these two types of memory cycles differ in only two respects. During the execution of a memory store cycle, the sense-amplifier-gate-tube sample pulse is not generated (inhibited by a computer command), and the memory buffer register flip-flop is loaded from an external source prior to the write portion of the cycle. Since the memory buffer register flip-flop cannot be made to reflect the status of the selected core, the read portion of the cycle effectively erases the initial contents of the core by ensuring that the selected core will be in the zero state prior to the write portion of the cycle. During the write portion of the store cy-

cle, the new information contained in the memory buffer register flip-flop is written into the selected core in exactly the same manner described for the memory readout cycle.

2.3 ANALYSIS OF FERRITE CORE OUTPUT SIGNALS

To simplify the previous discussions of the principles of ferrite core operation, the simple waveforms of figure 7—3 were purposely used as examples of the output signals of both the selected ferrite core and the sense winding of a memory plane. However, since the sense winding links all the cores of a memory plane, the sense winding output voltage actually represents the summation of the selected core output voltage and the noise voltages generated by all the other cores in that plane. Since the output voltage characteristics (amplitude, duration, and polarity) of each core depend on (1) the flux state of the core, (2) the amplitude, duration, rise and fall time, and polarity of the current pulses that affect the core, and (3) the wiring geometry of the sense winding, the resultant amplitude and timing of the noise voltages generated during a specific memory cycle depend on many variable operating conditions. As a result, the sense winding output voltage waveform is a complex variable in which the amplitude of noise spikes can exceed (depending on the operating conditions) the amplitude of a ONE output signal by a factor of 3 to 1. The present discussion is limited to an analysis of the individual ferrite core output signals produced under various operating conditions; the manner in which the individual cores affect the sense winding output voltage is discussed in 2.4.

Since the output signal of a selected core is sampled during the read portion of a memory readout cycle, the following analysis of ferrite core output signals is based on the core response to read-current pulses only. Because the noise voltage generated during this portion of the memory cycle is produced by the output signals from the half-selected cores (to be discussed in 2.4), this analysis includes the core response to both full-selected and half-selected cores.

An examination of the hysteresis loop of figure 7—4 shows that under normal core memory operation a ferrite core can be made to remain in one of five basic remnant flux states. (Actually, the ferrite core exhibits nine distinct flux states; however, only the five basic flux states will be considered in this discussion. The four flux states that are omitted are actually minor deviations of four of the five basic flux states.) The five remnant flux states shown in figure 7—4 are identified by symbols which designate the information status of the core (one or zero) and the polarity and amplitude of the last current pulse that was applied to the core prior to the memory cycle in process. (The manner in which a core is made to exhibit a specific remnant flux state is discussed in 2.4.) If the last current pulse applied to a core was equivalent to a full-amplitude current pulse (summation of X and Y drive line current pulses applied to a selected core), the core is said to be in an undisturbed state (only the undisturbed one state is possible). If the last current pulse applied to a core was equal to a half amplitude current pulse (read, write, or inhibit), the core is said to be in a disturbed state (rz, wz, r1, and w1), and the half-amplitude current pulse is said to be a disturbing pulse. The disturbed flux states of a core are referred to as write disturbed or read disturbed, depending on whether the last disturbing pulse was a half-amplitude write or a half-amplitude read current pulse. Since the inhibit current pulse is equivalent to a half-amplitude read current pulse, a core that is disturbed by an inhibit current pulse is considered to be read disturbed.

The output voltage generated by a core during the read portion of a memory cycle depends upon the initial magnetic state of the core and the characteristics of the read current pulse. Since each of the five remnant flux states can be affected by either half-amplitude or full-amplitude (fully selected) read current pulses, 10 distinct output voltages are possible. Figure 7—6 lists these output voltages by name and symbol and shows the general waveshapes and voltage amplitudes that will be obtained for a typical core. In each case, the output voltage symbol subscripts fully define the initial state of the core as well as the amplitude of the read current pulse that produced the voltage. As an example, consider the symbol V_{rz}. The subscripts r and z specify that the core was in the read disturbed zero state, and the subscript h specifies that the core was actuated by a half-amplitude read current pulse. Absence of the subscript specifying current pulse amplitude indicates that the core is actuated by a full-amplitude (selected core) read current pulse.

A study of the hysteresis loop of figure 7—4 reveals that the remnant flux of a core in the read disturbed one (r1) or read disturbed zero (rz) state is not permanently changed by the application of a half-amplitude read current pulse in that the respective flux path f-j-f or b-d-b is traversed. As a result, the core output voltage that is produced under either of these two half-selected conditions is a result of completely reversible flux changes. Because these flux changes are reversible, the amplitude and duration of these two output voltages are directly related to the rise and fall time of the applied current pulse and the flux density of the core.

Since the flux density of a core in the read disturbed zero state (rz) is slightly greater than the flux density of a core in the read disturbed one state (r1), the amplitude of the half-select read disturbed one (V_{r1h}) output voltage is slightly greater than the half-select read disturbed zero (V_{rzh}) output voltage (fig. 7—6).

	INITIAL STATE OF CORE	OUTPUT WHEN I_r IS APPLIED	AMP. IN MV	OUTPUT WHEN I_r 2 IS APPLIED	AMP IN MV
1	wz	$_wV_z$ - FULL SELECTED WRITE DISTURBED ZERO	22	$_wV_{hz}$ - HALF SELECTED WRITE DISTURBED ZERO	10
2	rz	$_rV_z$ - FULL SELECTED READ DISTURBED ZERO	16	$_rV_{hz}$ - HALF SELECTED READ DISTURBED ZERO	5
3	ol	$_oV_l$ - FULL SELECTED UNDISTURBED ONE	131	$_oV_{hl}$ - HALF SELECTED UNDISTURBED ONE	4
4	rl	$_rV_l$ - FULL SELECTED READ DISTURBED ONE	131	$_rV_{hl}$ - HALF SELECTED READ DISTURBED ONE	7
5	wl	$_wV_l$ - FULL SELECTED WRITE DISTURBED ONE	135	$_wV_{hl}$ - HALF SELECTED WRITE DISTURBED ONE	12

NOTE

THE AMPLITUDES SHOWN WERE OBTAINED UNDER TEST CONDITIONS AND THEREFORE MAY DIFFER FROM THOSE OBTAINED UNDER ACTUAL OPERATING CONDITIONS. HOWEVER, THE RELATIONSHIPS BETWEEN THESE VOLTAGES WILL REMAIN THE SAME UNDER BOTH CONDITIONS.

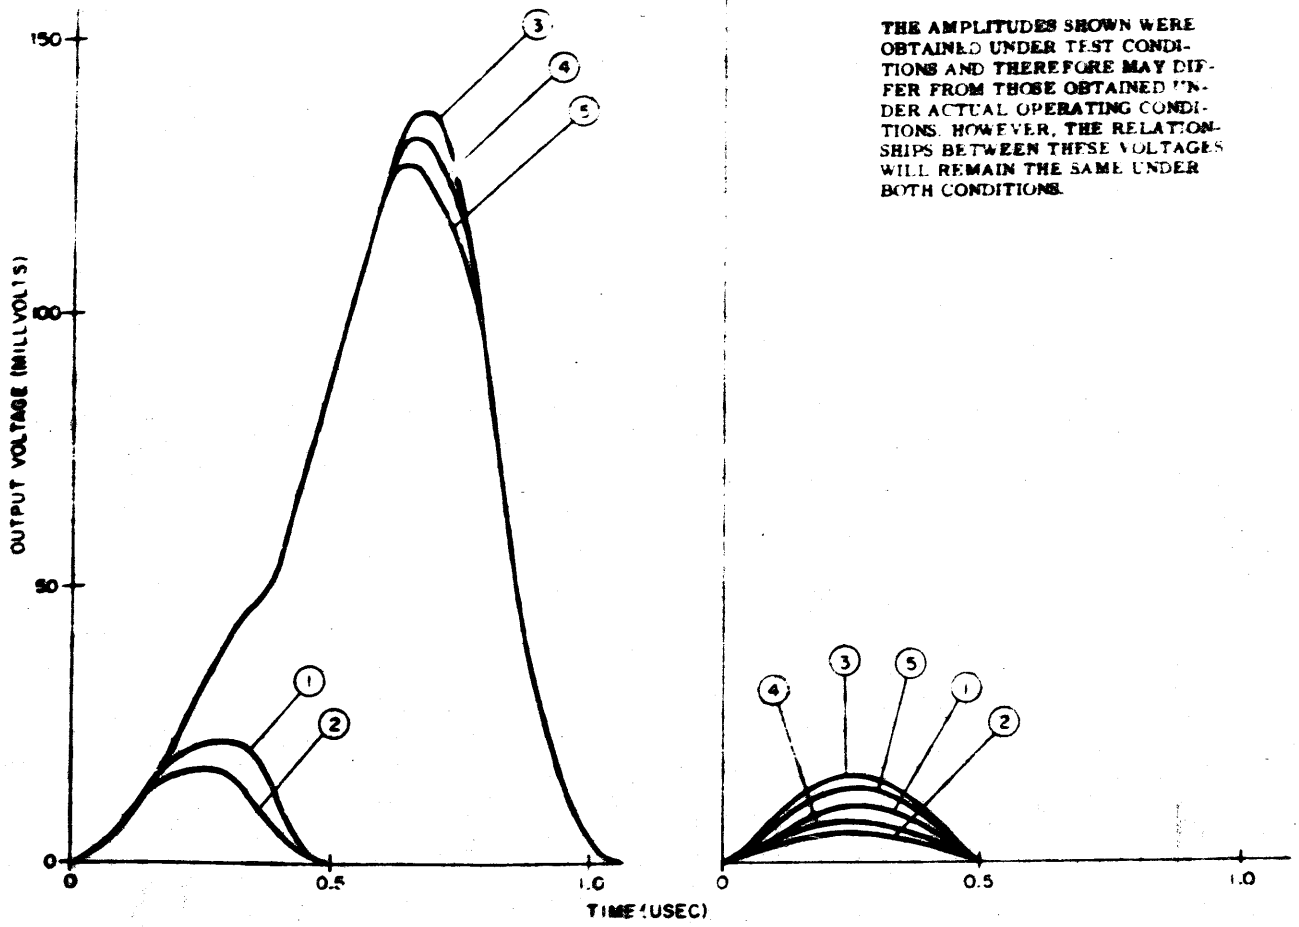

Figure 7–6. Outputs of a Typical Ferrite Core

A further study of the hysteresis loop of figure 7–4 reveals that the application of a full-amplitude read current pulse to a core in any of the five remnant flux states, or the application of a half-amplitude read current to a core in the undisturbed one (u1), write disturbed one (w1), or write disturbed zero (wz) flux state will result in an irreversible flux change since the core flux does not return to its initial state upon the termination of the driving current pulse. Table 7–1, which refers to the hysteresis loop of figure 7–4, specifies the flux path that is traversed for each of the eight selected conditions that result in an irreversible flux change. Since a reversible component of flux change is also produced under the abovementioned selected conditions, the output voltage generated in the sense winding is actually a summation of reversible and irreversible components of voltage. Because the duration of the irreversible component of the output voltage is dependent on the total amount of flux change and the switching time of the ferrite material (refer to 2.1), the duration of the three full-selected 1-output signals ($_sV_1$, $_sV_1$, and $_sV_1$ in fig. 7–3) is considerably longer in time and larger in amplitude than the five other irreversible output signals ($_sV_{n1}$, $_sV_{n1}$, $_sV_{n1}$, $_sV_{n1}$, and $_sV_1$ in fig. 7–6). Examination of the hysteresis loop of figure 7–4 shows that the same amount of irreversible flux change is produced when a half-amplitude read current pulse is applied to a core in either the write disturbed one state (w1) or the write disturb zero state (wz). However, since the flux switching rate is slightly slower when the flux density of a core is being decreased than when the flux density is being increased, the write disturbed half-selected zero ($_sV_{hz}$) output signal is slightly greater in amplitude and shorter in duration than the write disturbed half selected one ($_sV_{hz}$) output signal (fig. 7–6).

2.4 SENSE WINDING CONFIGURATION AND OUTPUT SIGNALS

As stated in 2.1, the sense winding of a memory plane in either of the two ferrite core arrays links all the cores in that plane. In the core memory 2 array this is accomplished by passing a single winding through all the cores of the 64 x 64 core matrix. Since the memory plane of the core memory array (fig. 7–7) is essentially composed of 16 core memory 2 planes (designated as subplanes), the sense winding of a core memory 1 plane is composed of 16 individual sense windings, one in each subplane. In a memory 1 plane, the individual subplane sense windings are connected in a series-parallel manner to form a single memory plane sense winding. Since the sense windings of both the memory 2 plane and the memory 1 subplane are essentially identical, the following discussion, which is based on the memory 2 plane, is applicable to both core memory arrays. The analysis of the sense winding output voltage is preceded by a discussion of the sense winding geometry and the manner in which the memory plane cores are affected by the read, write, and inhibit current pulses.

TABLE 7–1. IRREVERSIBLE FLUX CHANGES

INITIAL REMNANT STATE	READ CURRENT PULSE	FLUX PATH TRAVERSED
u1	Full amplitude	h-j-e-a
w1	Full amplitude	g-j-e-a
r1	Full amplitude	c-d-e-a
wz	Full amplitude	c-d-e-a
rz	Full amplitude	b-d-e-a
u1	Half-amplitude	h-j-f
w1	Half-amplitude	g-j-f
wz	Half-amplitude	c-d-b

2.4.1 Sense Winding Geometry

The sense winding of a memory plane is wired in a specific pattern to provide for the optimum cancellation of the noise voltages generated during the execution of a memory cycle. These noise voltages must be minimized so that the associated sense amplifier can reliably distinguish whether the selected core contained a 0 or a 1. The present discussion is limited to an analysis of the sense winding geometry; the manner in which noise voltages are minimized is presented in 2.4.3.

Figure 7–8, A, which illustrates an 8 by 8 portion of the 64 x 64 memory 2 plane matrix, shows the arrangement of the ferrite cores and the manner in which the X, Y, and sense windings are wound through these cores. Although the inhibit winding, which is wound in parallel to the Y drive lines to link all the cores in the plane, is referred to in subsequent discussions, it has been omitted in the diagram to improve drawing clarity. The arrowheads on the X and Y drive lines indicate the direction of current flow for the half-amplitude read current pulses. The polarity sign associated with each core is used to indicate the polarity of the core output signal that is induced in the sense winding as a result of a read current pulse. The polarity of the core output signal is determined by using the right hand rule to determine the direction of flux change that results from the application of a read current pulse and then, using Lenz's law, to determine the direction of the induced voltage in the sense winding. In all cases, the polarity of the core output signal is referenced to the start terminal of the sense winding.

To illustrate the manner in which the sense winding polarity is determined, consider the core shown in figure 7–8, B, which is used to represent the core in the upper left hand corner of figure 7–8, A. To determine the direction of flux change produced in this core by a read current pulse, grasp the X or Y drive line in the

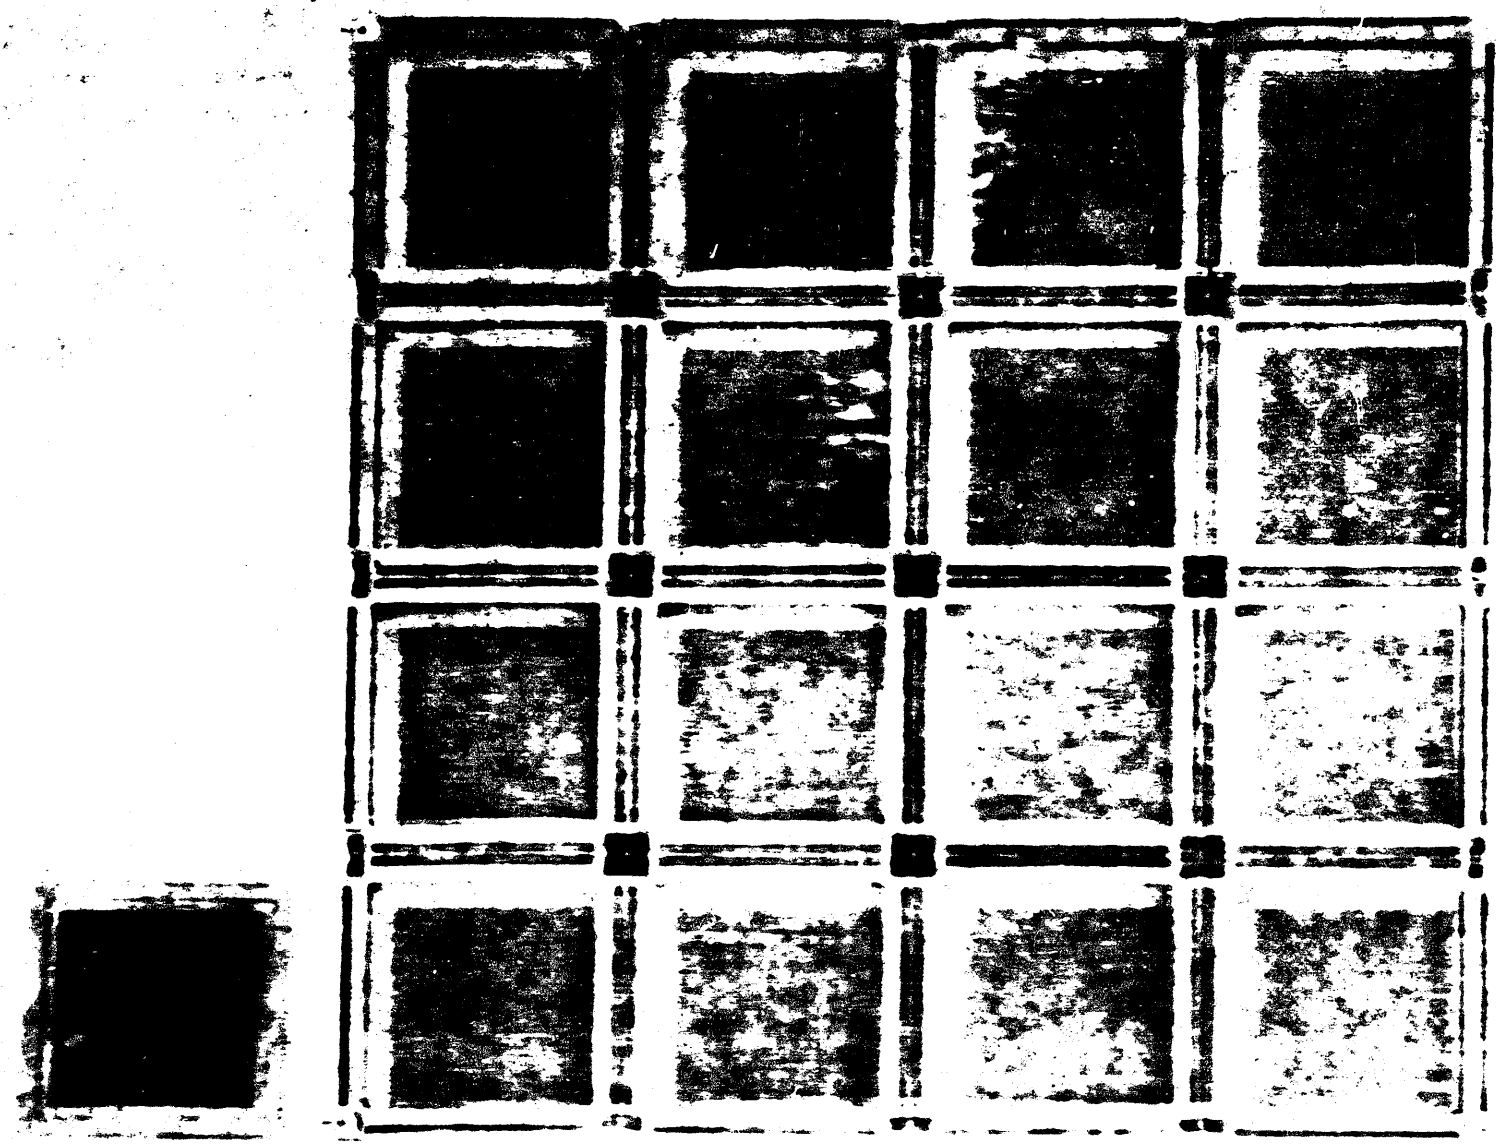

Figure 7–7. Comparison of 64² and 256² Core Memory Planes

right hand with the thumb pointing in the direction of current flow. The fingers of the right hand indicate that the resultant flux change in the core is in a counterclockwise direction. Lenz's law states that the current which flows as a result of an induced voltage is such that it will set up a field to oppose the original change of flux. Since, in this example, the direction of the flux change has been determined to be counterclockwise, the field set up by the sense winding must be in the clockwise direction. To determine the direction of the resultant current flow in the sense winding, grasp the sense winding in the right hand so that the fingers point in the direction of the magnetic field set up by the sense winding (clockwise). The thumb of the right hand points toward the positive terminal of the sense winding. In this specific example, the end of the sense winding is the positive terminal and the start of the sense winding is the negative terminal. Since the polarity sign associated with a core is referenced to the start terminal of the sense winding, the core in this example is assigned a negative polarity. The polarity of all the other cores can be determined in the same manner. As indicated in figure 7–8, A, the cores induce voltages in the sense winding according to a definite polarity pattern in which the number of positive signs is equal to the number of negative signs. Although the figure only shows an 8 x 8 core matrix, this polarity pattern also exists in the 64 x 64 core matrix.

2.4.2 Development of Disturbed Flux States

Prior to analyzing the sense winding output voltages produced during the execution of a memory readout cycle, it is first necessary to consider the manner in which the read, write, and inhibit current pulses will affect the cores of a plane. As previously stated, during

the execution of either type of memory cycle, a specific core is selected for operation by the application of coincident half-read and half-write current pulses to the mutually perpendicular X and Y drive lines that link the selected core (to effectively produce the full-read and full-write current pulses required to switch the selected core). Reference to figure 7—8, A, shows that, because of the coincident current method of core selection, 7 cores (63 on an actual plane) on each of the selected lines will be affected by half-read and half-write current pulses; that is, they will be half-selected. The remainder of the cores that are not linked by the selected X and Y drive lines are not affected by read and write current pulses; therefore, these cores are designated as nonselected. Thus, during the execution of a memory cycle, the individual cores of a 64^2 memory plane are designated as fully selected (1 core), half-selected (126 cores), or nonselected (3,969 cores).

If the selected core is to contain a 1 at the end of a memory cycle, then during the execution of the memory cycle the selected core will effectively be actuated by full-read and full-write current pulses, and the 126 half-selected cores will be actuated by half-read and half-write current pulses. The nonselected cores will not be affected under this condition. However, if the selected core is to contain a 0 at the end of a memory cycle, then an inhibit current pulse (equivalent to a half-read current pulse), which is applied to the inhibit winding in coincidence with the write current pulses, will also be generated during the execution of the memory cycle. As a result, all cores (full-, half-, and nonselected) will be affected during the write portion of this memory cycle. Since the magnetic field set up by an inhibit current pulse is equal and opposite to the magnetic field set up by a half-amplitude write current pulse, the simultaneous application of these two current pulses to a core will result in the effective cancellation of the two fields; that is, the core will behave as if the two current pulses were not applied. Thus, if the selected core is to contain a 0, the selected core will be effectively actuated by a full-read and a half-write current pulse, the 126 half selected cores will be effectively actuated by a half-read current pulse only, and the 3,969 nonselected cores will be actuated by an inhibit current pulse (effective half-read).

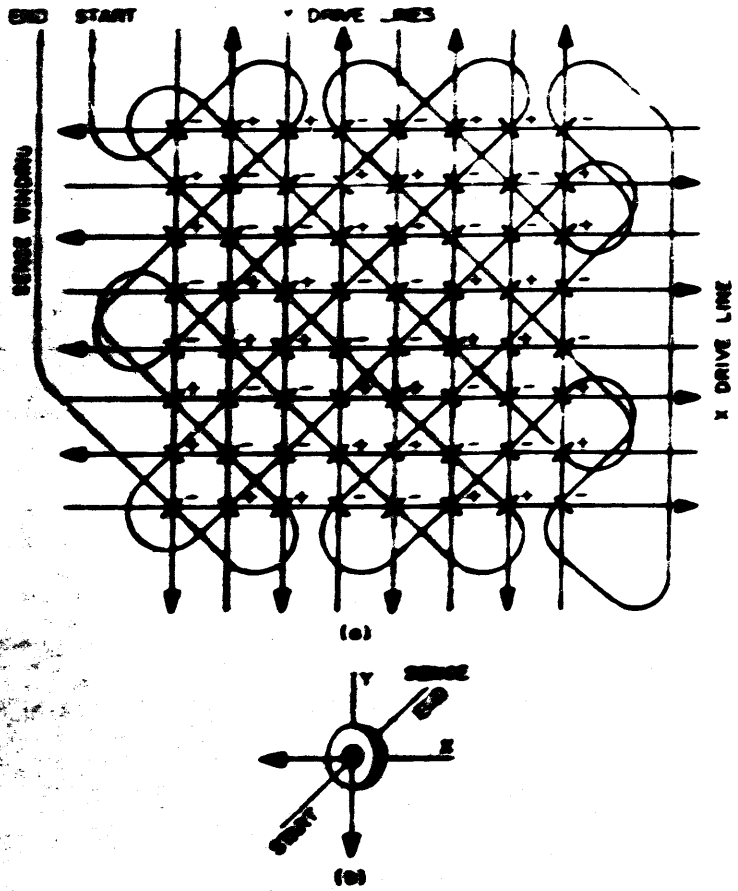

Figure 7—8. Subplane Wiring Geometry

Table 7—2 provides a summary of the sequence of current pulses that are applied to the memory plane for the two conditions of selected core content. Since the table is based on the final content of the selected core, the current pulse sequences noted in the table are applicable for both the memory readout cycle and the memory store cycle.

Since the remnant flux state of a core is determined by the polarity and amplitude of the last applied current pulse, a study of table 7—2 and figure 7—4 will show how each core is affected under the two conditions of selected core content. If the selected core is to contain a 1 at the end of the memory cycle, the selected core will be left in the undisturbed one (r1) flux state, the individual half-selected cores will be left in the write disturbed one (w1) or write disturbed zero (wz) flux state (depending upon whether the individual core contains a 1 or a 0), and the 3,969 nonselected cores will remain static; that is, in their original flux state. If the selected core is to contain a 0 at the end of the memory cycle, the selected core will be left in the write disturbed zero (wz) flux state and the individual half-selected cores and nonselected cores will be left in either the read disturbed zero (rz) or read disturbed one (r1) flux state (depending upon whether the individual core contains a 1 or a 0). From the above discussion, it can be seen that any cores in a plane can be made to exhibit any one of the five remnant flux states.

2.4.3 Analysis of Sense Winding Output Voltage

A study of the 64 x 64 core matrix (a part of which is shown in fig. 7—8, A) reveals that, when the selected X and Y drive lines are energized with half-amplitude read current pulses (to read out the contents of the selected core), a total of 127 cores are involved in producing the sense winding output voltage. That is, when a particular core is selected, the sense winding output represents the summation of the selected core output voltage and the noise voltages generated by the 126

half-selected cores (63 cores on the selected X and Y drive lines). A study of figure 7–8, A, shows that each X and Y drive line contains a number of positive and negative cores; therefore, in a 64 by 64 core matrix, each drive line will contain 32 positive and 32 negative cores. Thus, because of this sense winding polarity pattern, the application of coincident half-read current pulses to the selected X and Y drive lines of a memory plane will result in the generation of 62 (31 cores on the selected X and Y drive lines) half-selected output voltages whose polarity is the same as the polarity of the selected core output voltage, and 64 (32 cores on the selected X and Y drive lines) half-selected output voltages whose polarity is opposite to the polarity of the selected core output voltage. As a result, the half-selected output voltages of 124 of the 126 half-selected cores (62 positive and 62 negative output voltages) will tend to cancel each other. However, complete cancellation of these 62 pairs of half-selected cores (a pair consists of one positive and one negative half-selected core) can only occur under specific conditions since the amplitude of the individual half-selected output voltages depends upon which flux state the core exhibits prior to the application of the half-read current pulse. The half-selected output voltages of 2 of the 64 half-selected cores whose polarity is opposite to the polarity of the selected core will not be cancelled; therefore, these two half-selected output voltages will subtract from the selected core output voltage. The equation to express the sense winding output voltage produced during the read portion of the memory cycle may be written as follows:

$$V_{ss} = V_s - 2V_h \pm 62V_d$$

where:

V_{ss} = the output voltage of the sense winding

V_s = the output voltage of the selected core

V_h = the average half-selected output voltage whose polarity is opposite to the selected core output voltage

V_d = the difference between the average positive half-selected output voltage and the average negative half-selected output voltage.

The last term of the equation, which deals with the summation of the half-selected output voltages that tend to cancel each other, is defined as the delta voltage of the sense winding output. The delta voltage may be either positive, negative, or zero, depending upon whether the average positive half-selected output voltage is larger than, smaller than, or equal to the average negative half-selected output voltage. Since the delta voltage is obtained from 62 pairs of half-selected cores, it follows that, if each core of a pair of cores is in the same flux state (fig. 7–4) and if the characteristics of all the cores are considered to be exactly the same, the application of a half-read current pulse to the half-selected cores will cause each core of a pair of cores to generate a half-selected output voltage that is exactly equal to, but of opposite polarity to, its mate. If each of the 62 pairs of half-selected cores is thus balanced, the delta voltage will be equal to zero to provide for the maximum cancellation of half-selected output signals. In memory maintenance programming, this condition of maximum cancellation of half-select output voltages is obtained by using either a 1's or a 0's test pattern wherein all the cores of a memory plane are in either the one or zero flux state.

A study of the equation reveals that the maximum value of delta voltage is obtained when all the half-selected cores of one polarity produce the smallest half-selected output voltages, while all the half-selected cores of the opposite polarity produce the largest half-selected output voltages. Reference to figure 7–6 shows that the smallest half-selected output voltage is produced by a core in the read disturbed zero (rz) flux state, and that the largest half-selected output voltage is produced by a core in the undisturbed one (u1) flux state. However, since it is only possible for one core on an X or Y drive line to exhibit the undisturbed one (u1) flux state at any particular time (an attempt to cause a second core to exhibit the undisturbed one flux state will cause the existing undisturbed core to be write disturbed, refer to 2.3), this flux state does not enter into the determination of the maximum value of delta voltage. Instead, the maximum value of delta voltage is obtained when all the half-selected cores of one polarity are in the read disturbed zero (rz) flux state, and all the half-selected cores of the opposite polarity are in the write disturbed one (w1) flux state.

Since the delta voltage can be of either polarity, the maximum value of delta voltage can be produced in

TABLE 7–2. EFFECTIVE CURRENT PULSE SEQUENCE APPLIED TO CORES OF ONE MEMORY PLANE

	SELECTED CORE IS TO CONTAIN A	
	1	0
Selected Core (1)	Full-Read, Full-Write	Full-Read, Half-Write
Half-selected Cores (126)	Half-Read, Half-Write	Half-Read, No Pulse
Nonselected Cores (3,969)	No Pulse, No Pulse	No Pulse, Half-Read

two different ways; that is, by two different test patterns. The maximum positive delta voltage is produced when all the plus half-selected cores are in the write disturbed one (w1) flux state and all the minus half-selected cores are in the read disturbed zero (rz) flux state. If all the cores of the plane contain the pattern just described, the plane is said to contain a regular checkerboard pattern. The maximum negative delta voltage is produced when all the plus half-selected cores are in the read disturbed zero (rz) flux state and all the minus half-selected cores are in the write disturbed one (w1) flux state. The test pattern that will produce this condition is defined as the inverted checkerboard pattern. Since the polarity of the delta and selected core output voltages can be individually positive or negative (fig. 7–8, A), eight separate conditions exist whereby the maximum value of delta voltage will either add to or subtract from the selected core output voltage to produce the maximum distortion of the selected core output voltage. A summary of the conditions under which the delta voltage affects the selected core output voltage to produce a smaller or a larger sense winding output voltage is given in table 7–3.

Reference to figure 7–6 shows that, if all the cores of a plane have identical characteristics, the maximum value of delta voltage produced under the abovementioned conditions will have a peak amplitude in excess of 400 mv and a duration of approximately 0.5 μsec. Since the delta voltage and the selected core output voltage both start at the same time, the delta voltage always distorts the selected core output signal. If the selected core initially contained a 0, the delta voltage will completely mask the 0 output voltage since both voltages have the same duration. However, if the selected core initially contains a 1, the delta voltage will only distort the first 50 percent of the 1 output signal, and the 135-mv (approximate) peak of the 1 output signal will not be affected. It is because of this important timing factor that a time-amplitude sampling technique can be used (during the execution of a memory readout cycle) to reliably determine whether the selected core contained a 0 or a 1. That is, although the peak amplitude of the delta voltage can be much greater than the amplitude of the selected core output signal, reliable memory operation is obtained by sampling (fig. 7–4) the sense amplifier output voltage at a specific time, namely, at the peak of the amplified 1 signal. Because the content of the selected core is sampled at a specific time, reliable discrimination between a 0 and a 1 output signal depends upon the ability of the sense amplifier to faithfully reproduce the relative timing of the sense winding waveform.

TABLE 7–3. MAXIMUM DELTA VOLTAGE COMBINATIONS

SELECTED CORE POLARITY AND CONTENT	NEGATIVE CORES	POSITIVE CORES	RESULTANT POLARITY	RESULTANT OUTPUT SENSE WINDING VOLTAGE
Positive 1	rz	w1	+	Larger positive 1
Positive 1	w1	rz	−	Smaller positive 1
Negative 1	rz	w1	+	Smaller negative 1
Negative 1	w1	rz	−	Larger negative 1
Positive 0	rz	w1	+	Larger positive 0
Positive 0	w1	rz	−	Smaller positive 0
Negative 0	rz	w1	+	Smaller negative 0
Negative 0	w1	rz	−	Larger negative 0

CHAPTER 3
THEORY OF OPERATION OF CORE MEMORY 2 (64^2)

SECTION 1
64^2 FERRITE CORE ARRAY

1.1 GENERAL

The 64^2 ferrite core array contained in unit 11 (fig. 1–10) is the principal component of core memory 2 since it is the information storage center of this memory device. This section describes the physical characteristics of the ferrite core array to show the arrangement of its cores and windings, and to provide a definition of the terms used in the subsequent sections of this chapter.

The storage capacity of the 64^2 ferrite core array is equal to 4,096 words of 34 bits each. Since a single core can store one bit of information, the array contains 4,096 x 34 or 139,264 ferrite cores. These cores are arranged in a 3-dimensional array in which each horizontal layer or digit plane contains 4,096 cores arranged in a 64 x 64 square formation. The 34 digit planes of this array are stacked vertically, and the X and Y selection windings of these planes are interconnected to form the X and Y selection windings of the array. Actuation of the current drivers associated with one X and one Y selection winding will mutually affect the vertical column of 34 cores (one core in each plane) that represent the selected memory register. Since only 33 of these digit planes are connected to sense amplifiers and digit plane drivers, only 33 planes can be active at any time. The 34th plane of the array, which is required to provide for the symmetrical wiring of the array, is used as a spare plane.

1.2 64^2 PLANE WIRING CONFIGURATION

All the digit planes used in the 64^2 ferrite core array are identical. Each plane contains 4,096 cores in a 64 x 64 formation and all the windings necessary for memory operation. Figure 7–9 shows an abbreviated version of a 64^2 digit plane using 16 cores in a 4 x 4 formation. Since the inhibit windings of an odd-even pair of planes are interconnected (although not interrelated), this drawing also shows a portion of the associated even-numbered plane. In all cases, the even-numbered plane of an odd-even pair is located beneath the odd-numbered plane and is turned upside down in relation to the odd-numbered plane. In general, the odd-numbered planes of the array represent the individual bits of the left half memory word, and the even-numbered planes represent the individual bits of the right half memory word. In all cases, the digit plane drivers and sense amplifiers associated with the left half-word are connected to the digit planes on the left side of the array. Conversely, the digit plane drivers and sense amplifiers associated with the right half-word are connected to the digit planes on the right side of the array.

1.2.1 X and Y Selection Windings

As noted in figure 7–9, each X selection winding is wired through a single column of cores, while each Y selection winding is wired through a single row of cores. Selection of a particular core is accomplished by applying coincident half-amplitude, read-write current pulses to the appropriate X and Y selection windings (those that link the desired core) in such a direction that the resultant magnetic fields will add at the intersection of the two windings. As a result of this requirement, and because of the manner in which the cores are positioned (fig. 7–9), the direction of the read-write current pulses that are applied to an X or Y selection winding must be of the opposite polarity (applied in the opposite direction) when compared with the current pulses that are applied to the adjacent X or Y selection windings of a memory plane. Since all the X and Y current drivers produce read-write current pulses of the same polarity pattern (fig. 7–4), the required reversal of current flow in adjacent selection windings is accomplished by connecting the current drivers to consecutive selection windings on alternate sides of the array.

The manner in which the X and Y drivers are connected to the 64^2 core array to produce the desired result is discussed in 1.3. Included in 1.3 is a discussion of the manner in which the similarly numbered X and Y selection windings of the 34 planes are interconnected to form the selection windings of the array.

1.2.2 Inhibit Winding

The inhibit winding of each active digit plane of the 64^2 ferrite core array is assocated with an individual digit plane driver. During the execution of a memory cycle, each digit plane driver is controlled so that a negative current pulse will be applied to the associated

inhibit winding during the write portion of the cycle if it is required to inhibit the writing of a 1 in the selected core of the plane.

As noted in figure 7–9, the inhibit winding of each digit plane consists of a 1-turn winding which is wound through all the cores of a plane in parallel with the Y selection windings. Because the direction of this winding is alternated in adjacent rows of cores, the inhibit current pulse effectively flows in opposite directions through alternate rows of cores. Since write current pulses (negative) are applied to adjacent Y selection windings in opposing directions, the proper connection of the digit plane driver will cause the individual magnetic fields set up by the inhibit current pulse and the Y selection winding write current pulse in each core of the selected row of cores to completely cancel each other. To illustrate this point, consider the following example, which assumes that the digit plane shown in figure 7–9 represents the topmost plane of the array. Under this condition, a read-write current driver will be connected to the Y-O selection winding on the right side of the array so that the direction of this winding through this plane will be from right to left. The digit plane driver for this plane is connected to pin 11 of the next lower plane. As a result of the internal array wiring (jumper wires serially connected from pin 11 of the even-numbered plane to pin 18 of the odd-numbered plane), the direction of the inhibit winding through the Y-O row of cores will be from left to right. Since the Y selection winding write current pulse and the inhibit current pulse are both of the same polarity (negative), the magnetic fields set up by these two windings will be in opposite directions for each core in the Y-O row of cores. The net result is that these two fields cancel each other in each of the mutually affected cores.

Because the X and Y array selection windings are formed by serially connecting the similarly numbered drive lines of adjacent planes, the direction of read-write current flow in the similarly numbered drive lines of adjacent digit planes will always be in mutually opposing directions. To compensate for the effects of the winding reversal of the similarly numbered X and Y drive lines of adjacent planes, the direction of the inhibit winding must also be reversed in adjacent planes. In the 64^2 ferrite core array, the inhibit winding reversal is accomplished by the physical inversion of alternate planes of the array. To illustrate the manner in which this is accomplished, consider the inhibit winding con-

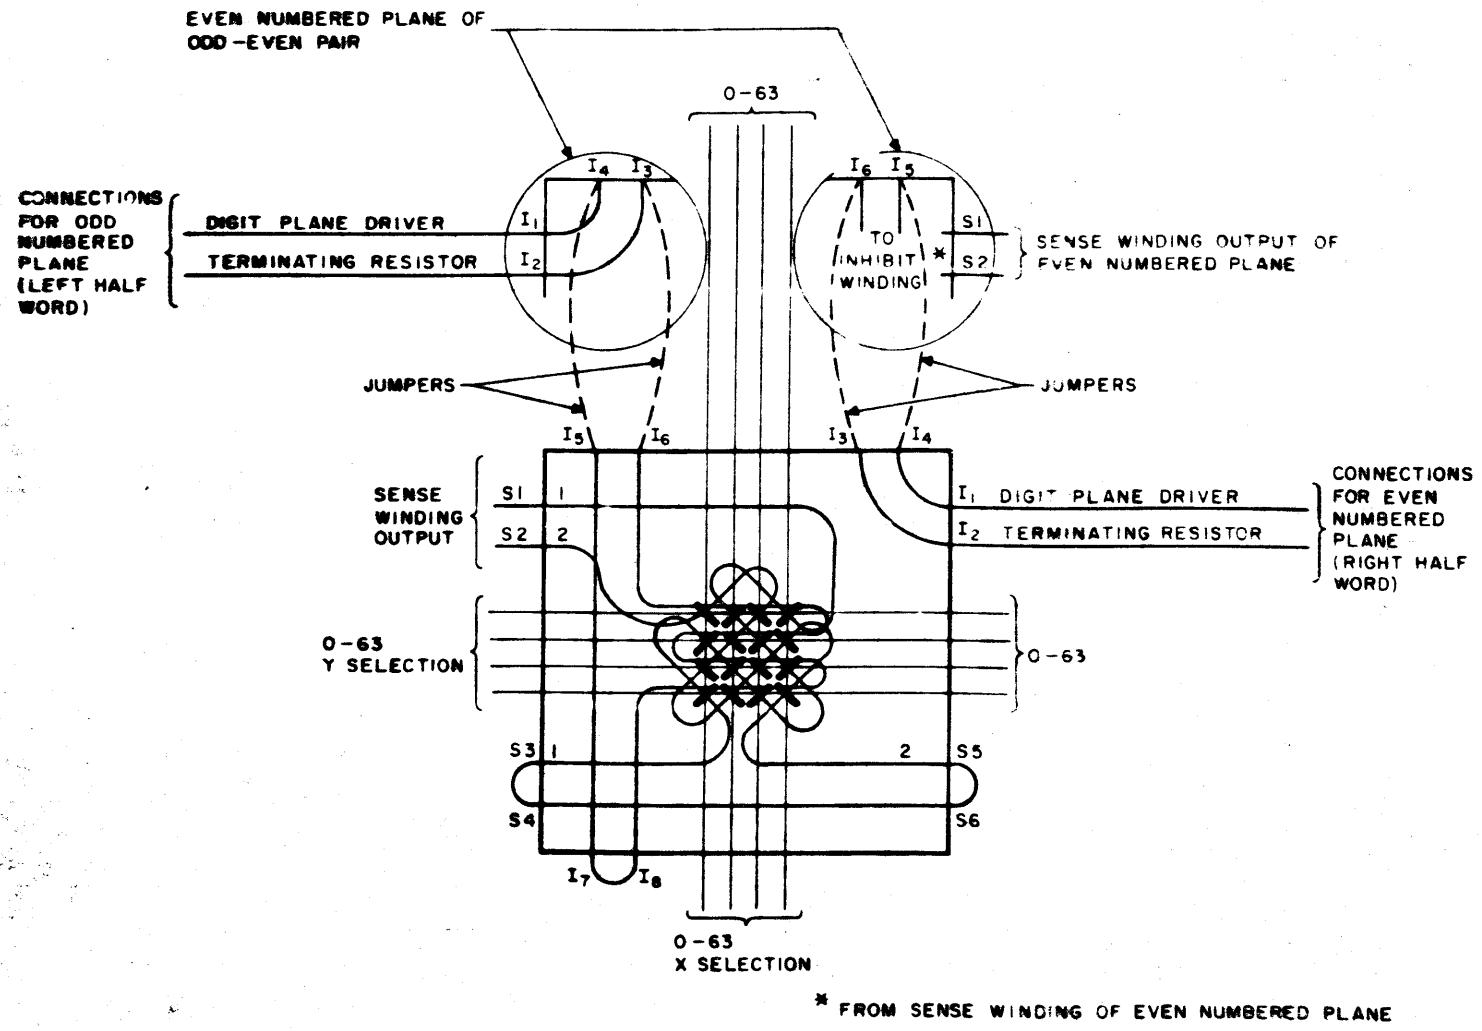

Figure 7–9. Top View of 64^2 Digit Plane (Odd)

nections of the odd-even pair of planes shown in figure 7–9. From the previous example, it was determined the inhibit winding of an odd-numbered plane is wound through the Y-O row of cores in a left to right direction. Thus, for an even-numbered plane, the inhibit winding must be wound through its Y-O row of cores in a right to left direction. Figure 7–9 shows that the digit plane driver for an even-numbered plane is connected to pin 11 of the associated odd-numbered plane. As a result of the internal array wiring (jumper wires serially connected from pin 11 of an odd-numbered plane to pin 18 of an even-numbered plane) and because the plane is relatively upside down, the direction of the inhibit winding through the Y-O row of cores of the even-numbered plane will be from right to left.

1.2.3 Sense Winding

The sense winding of each active digit plane is associated with a differential input sense amplifier which functions to amplify the induced voltages produced by the switching action of each core in the plane. As noted in figure 7–9, each digit plane has two separate sense windings (labeled 1 and 2) which are connected in series by means of internal and external jumper wires to form one long winding. Each individual sense winding which passes through half of the cores of the plane follows a diagonal path in order to minimize the capacitive and inductive coupling between itself and the other windings of the plane. The two ends of the sense winding (S1 and S2) are connected to the differential input sense amplifier so that only the induced voltage will be amplified. That is, since capacitive-coupled voltages do not produce a difference voltage between these two terminals, they will be rejected by the amplifier.

1.3 X AND Y CURRENT DRIVER CONNECTIONS

The complete 64^2 ferrite core array contains 34 digit planes (of the type described above) and 2 dummy planes. Figure 7–10 shows the overall arrangement of these planes to form the array and the word bit assignment of each plane. Although two dummy planes are not involved in the array wiring, they are included in this figure because they affect the plane identification numbering sequence. The two dummy planes do not contain any cores or windings and are only used as spacers in the array assembly. A study of figure 7–10 shows that the odd-numbered planes of the array (3, 5, 7, . . . , 35) are associated with the left half memory word, and the even-numbered planes (4, 6, 8, . . . , 36) are associated with the right half memory word. Plane number 18 is designated the spare plane because, being in the approximate center of the array, it requires the shortest jumpers (for sense amplifier and digit plane driver connections) if it is used to replace one of the active planes.

Because selection of a memory word involves the simultaneous application of read-write current pulses to similarly addressed cores in each plane, the corresponding X and Y drive lines in each of the 34 planes of the array will be involved in the operation. To provide a control so that one current driver can supply read-write current pulses to the corresponding X or Y drive line of each plane, the similarly numbered X and Y drive lines of all the planes are connected in series (by means of jumpers) so that common selection windings will be formed. One end of each of the 64 X and 64 Y selection windings is connected to a read-write current driver; the other end of each winding is connected to ground through a terminating resistor. In this ferrite core array, all the X selection line drivers (64) are connected to the front and back sides of the array and all the Y selection line drivers (64) are connected to the left and right sides of the array. Since the input and output connections of the X and Y selection windings are exactly the same, the following discussion will deal with Y selection windings only.

As noted in figure 7–10, the 64 Y line current drivers are divided into four equal groups. Each group of 16 drivers is contained in a pluggable driver panel which is connected to the array in a specific manner. The 16 Y line drivers of driver panel YA are connected (on the right side of the array) to every fourth Y selection line of plane 3, starting with line Y-O. Since the similarly numbered Y lines of each of the 34 digit planes are connected in series (by connecting the similarly numbered terminals of adjacent planes on alternate sides of the array), each of these Y lines will terminate on the right side of plane 36. To complete the circuit of each of these windings, the associated terminals of this plane are connected to ground through individual terminating resistors and a common 240-ohm resistor and a circuit breaker arrangement.

The 16 current drivers of each of the other Y driver panels are connected to the array in a similar manner; that is, the current drivers are connected to one end of the array and the associated terminating resistors are connected to the other end. Because the Y line current drivers are connected to the array in the sequence noted in figure 7–10, the direction of read-write current flow in the adjacent Y lines of each plane will be in mutually opposing directions.

The 64 X selection winding current drivers and terminating resistors are connected to the front and back sides of the array in exactly the same manner as described above. The current drivers contained in driver panels XA and XD are connected to the front side of the array, to every fourth X selection line, and the current drivers contained in driver panels XB and XC are connected to the rear side of the array.

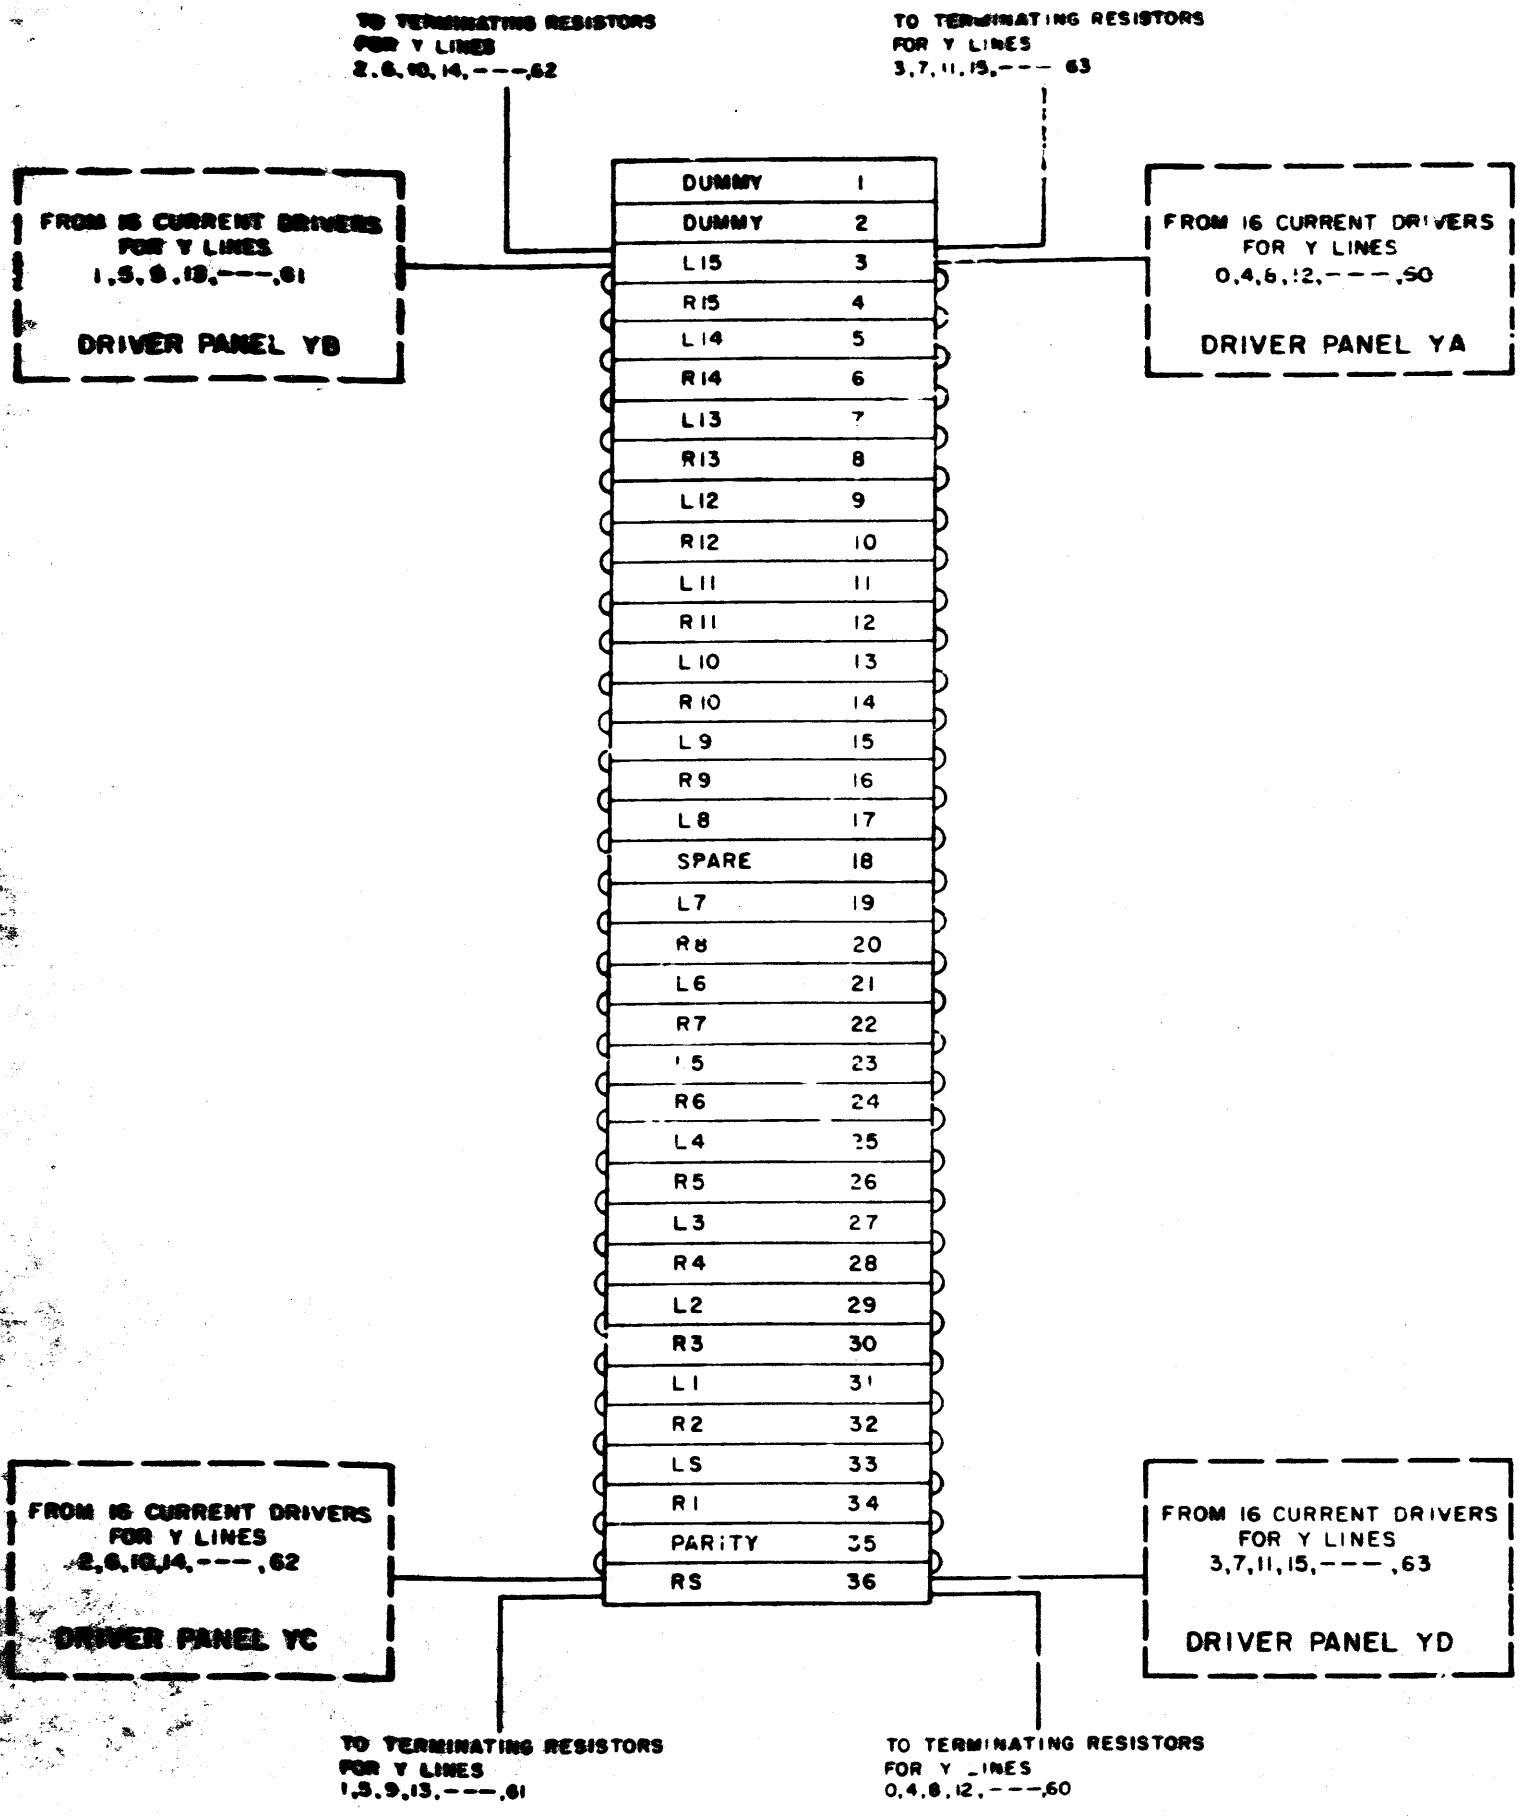

Figure 7–10. Front View of 64² Ferrite Core Array

SECTION 2
BLOCK DIAGRAM ANALYSIS

A simplified block diagram of the 64^2 core memory pulses is shown in figure 7–11. As noted in the figure, various input and output pulses and levels are generated during the execution of a memory cycle. The source and destination of these pulses and levels are tabulated in table 7–4.

As indicated in figure 7–11, the 64^2 core memory is divided into five sections: the selection section, the array section, the sense section, the digit plane driver section, and the timing and gating section. The memory buffer register, although not exclusively a part of the 64^2 memory, is also shown in this simplified block diagram since it is the only path by which core memory can communicate with the rest of the Central Computer System.

The selection section of the 64^2 core memory is composed of two identical portions, an X portion and a

TABLE 7–4. SOURCE OR DESTINATION OF CONTROL

CONTROL PULSE	TO OR FROM
Clear memory control	Instruction control element
Set memory address register	Program element
Start memory	Program element
Inhibit sample	Instruction control command generator
Sense amplifier output	Memory buffer register
Inhibit control levels	Memory buffer register

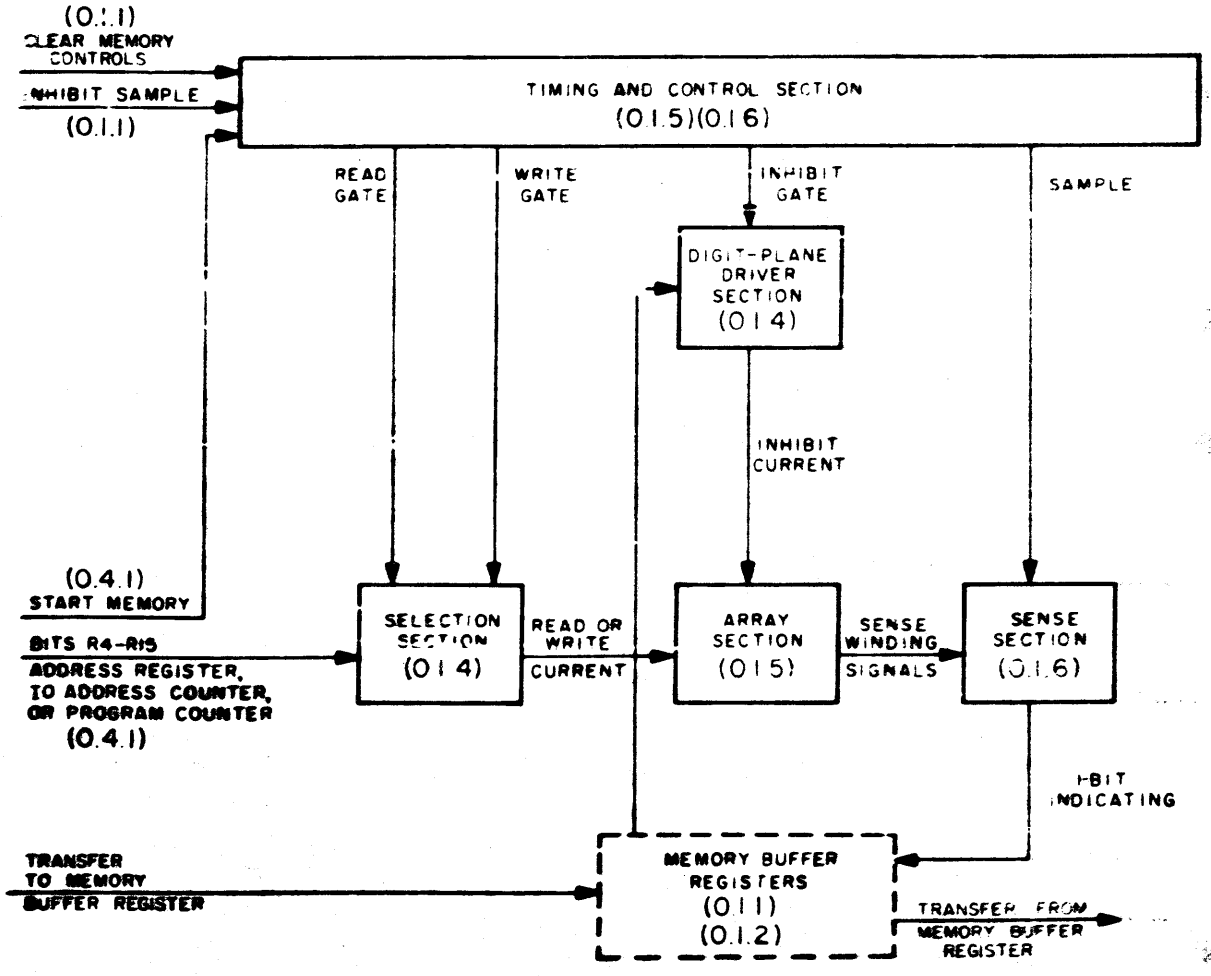

Figure 7–11. 64^2 Memory Device, Block Diagram

Y portion. The selection section contains circuits which decode information received from the program counter, address register, and the IO address counter. It also contains the memory address register, the memory gate generators, and the 64 X and 64 Y core memory drivers which generate the read and write current pulses that are applied to the selected X and Y lines in the array section.

As previously stated, the array section contains 33 active digit planes. Each digit plane contains a sense winding and an inhibit winding in addition to the 64 X and the 64 Y drive lines. The sense and inhibit windings of each active plane are connected to the sense amplifier and digit plane driver associated with a specific bit of the memory buffer register.

The sense section contains 33 individual sense amplifiers, one for each of the 33 active planes. Each sense amplifier amplifies the signal picked up from its associated sense winding and applies it to a sample gate tube which is sensed at the peak of an amplified 1-core output signal. The sampling pulse is developed in the timing and gating section and simultaneously samples all 33 gate tubes. If a specific gate tube is conditioned when it is sampled, the resulting output pulse will set the associated memory buffer register flip-flop to the 1 state. If a sense amplifier does not condition its sample gate tube, no output is produced when the tube is sampled; therefore, the associated buffer register flip-flop remains in the 0 state. During the execution of a memory store cycle, the sample pulse is inhibited so that none of the sample gate tubes can be sensed.

The digit plane driver section contains 33 digit plane drivers. Each digit plane driver is associated with a digit plane of the core memory array section and a particular flip-flop of the memory buffer register. Each individual driver is conditioned by its associated memory buffer register flip-flop and is controlled to generate an inhibit current pulse if the flip-flop contains a 0. If a memory buffer register flip-flop contains a 1, the associated digit plane driver is not conditioned and therefore does not generate an inhibit current pulse.

The timing and gating section contains the memory pulse distributor, which is composed of series-connected delay lines interspersed with pulse amplifiers to amplify the signal between the output of one delay line and the input to another. These three types of control flip-flops contained in this section are used to generate the gates which control the operation of the other sections of the 64 memory. These three gates which are obtained from the 1 side output of the control flip-flops are identified as the read gate, the write gate, and the inhibit gate. The input pulses which control the setting and clearing of these flip-flops are obtained from the memory pulse distributor.

SECTION 3
ADDRESS SELECTION

Selection of the 64^2 core memory register from which information is to be read, or into which information is to be stored, is controlled by the selection section. This section contains the flip-flop memory address register which is composed of 12 flip-flops. Six of these flip-flops and a diode matrix decoder are associated with the X portion of the selection section; the other six flip-flops and a second decoder are associated with the Y portion. The two portions operate similarly and are controlled simultaneously to select one X and one Y core memory current driver. Since the two portions of the selection section are exactly the same, the following discussion will deal only with selection of an X driver.

At the beginning of each memory cycle, the memory address register flip-flops are cleared by a clear-memory-controls pulse from the instruction control element. Approximately 0.6 μsec later, new address information is transferred to the memory address register from either

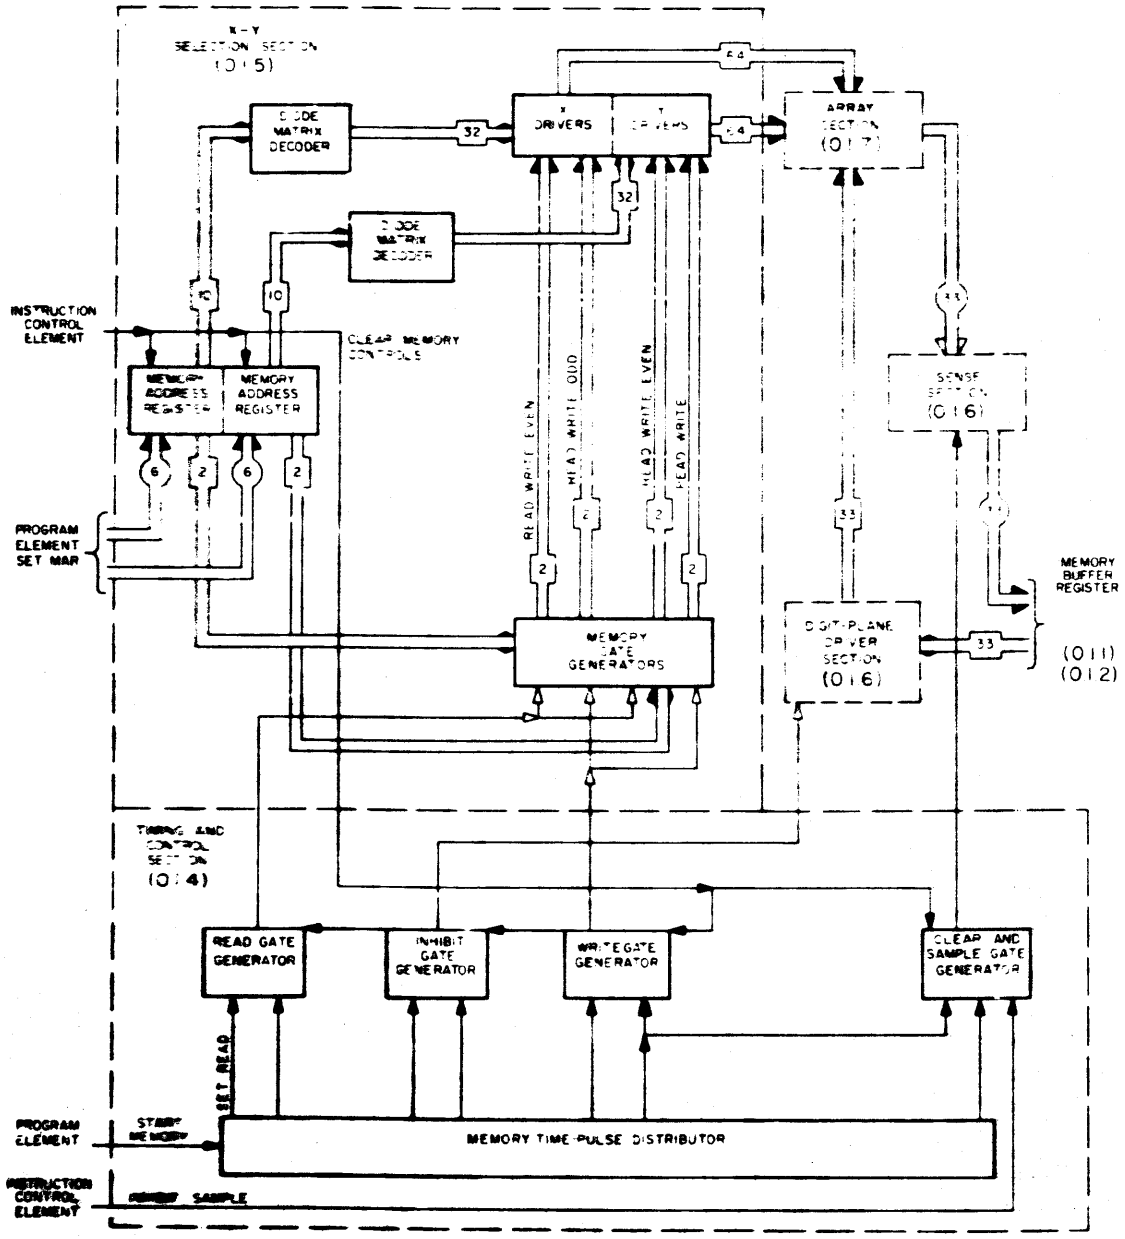

Figure 7–12. 64^2 Memory Selection Section

the program counter, the address register, or the IO address counter (fig. 7–12). The output levels from both the 1 and 0 sides of five of the flip-flops (bits R10 through R14) are supplied to the diode matrix decoder, and the output of the sixth flip-flop (R15) is fed to the memory gate generator circuits. The information which the decoder accepts from the five flip-flops is decoded in a diode negative AND circuit to select one of 32 output lines. (See fig. 7–13.) The 31 non-selected lines have an output level of +10V; the selected line has an output level of −30V. The selected output level is amplified to partially condition two adjacent core memory drivers, one for an even address and one for an odd address. One of these two drivers is then further selected by the proper memory gate generator.

Note

The memory gate generators, shown in figure 7–12, are logical AND circuits, each requiring both a d-c level from the sixth memory address register flip-flop and a read or a write gate from the timing and gating section. Four memory gate generators are used, two for reading (read odd, read even) and two for writing (write odd, write even).

If the sixth memory address register flip-flop is in the cleared state, the d-c level from its 0 side is at +10V and conditions the read-even and write-even memory gate generators (fig. 7–14). When pulsed by a read or a write or a write gate from the timing and gating section, these generators activate the read or control lines of the 32 even-core memory drivers and cause the selected current driver to generate a read or write current pulse. If the sixth flip-flop is set, it conditions the read-odd and write-odd memory gate generators, which, in turn, activate the read or write control lines of the 32 odd core memory drivers. When these gate generators are pulsed by a read or a write gate from the timing and gating section, the partially selected current driver will generate a read or write current pulse.

Note

In the Y selection section, the outputs of bits R4 through R8 are supplied to the diode ma-

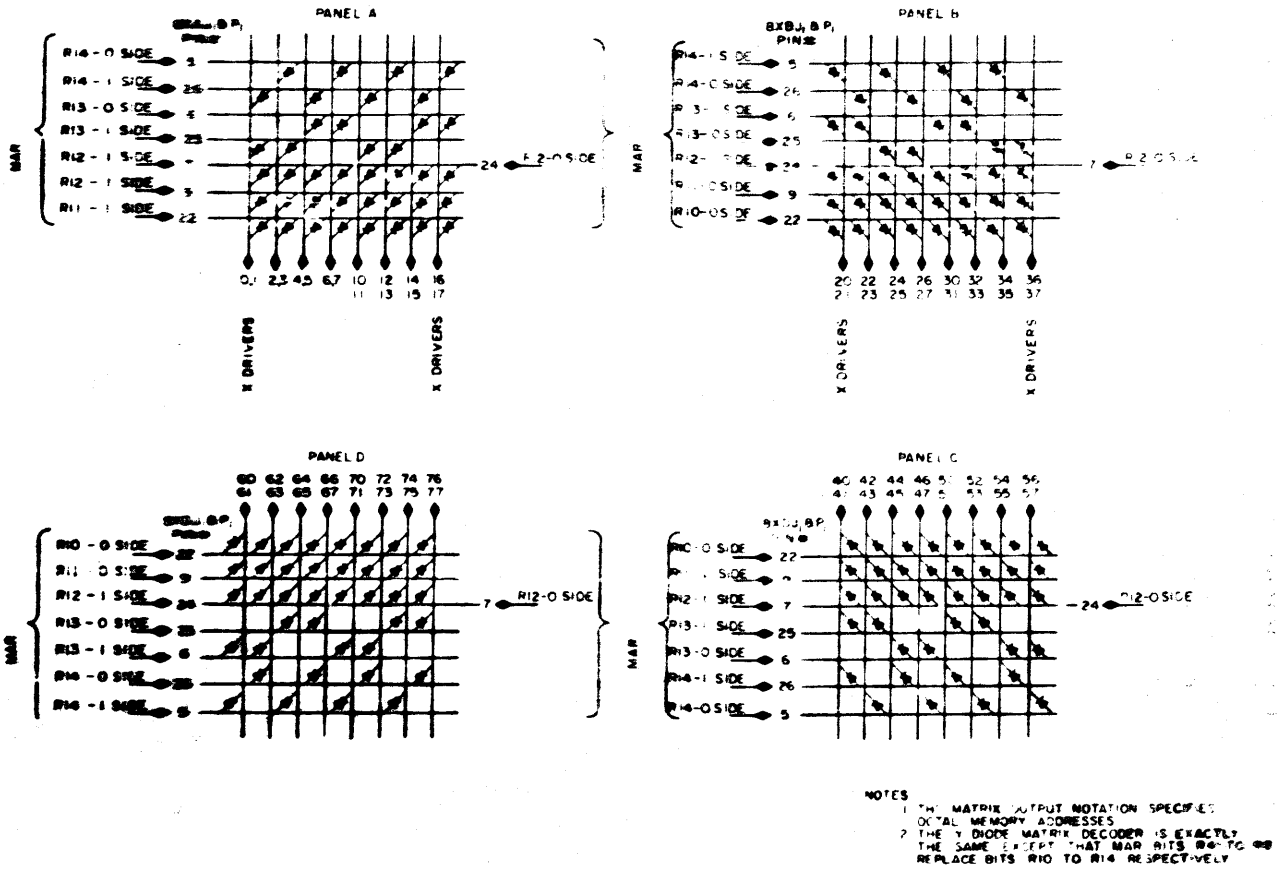

Figure 7–13. Diode Matrix Decoders

trix decoder and the outputs of bit R9 are used to select the odd or even pair of memory gate generators.

The X and Y core memory drivers, which are partially conditioned by the diode matrix decoders through the matrix output amplifiers, are further conditioned at the proper time by a read or a write gate from the selected memory gate generators described above. The selected X and Y drivers supply current pulses to an X and a Y co-ordinate line of the core memory array. These current pulses are of sufficient amplitude and duration to half-select the cores of the co-ordinate lines. At the intersection of the two co-ordinate lines, the individual half-select current pulses are algebraically added to apply a full-amplitude current pulse to the selected core. Overall operation of the core memory drivers is the same during the read and the write functions, except that the polarity of the output is reversed.

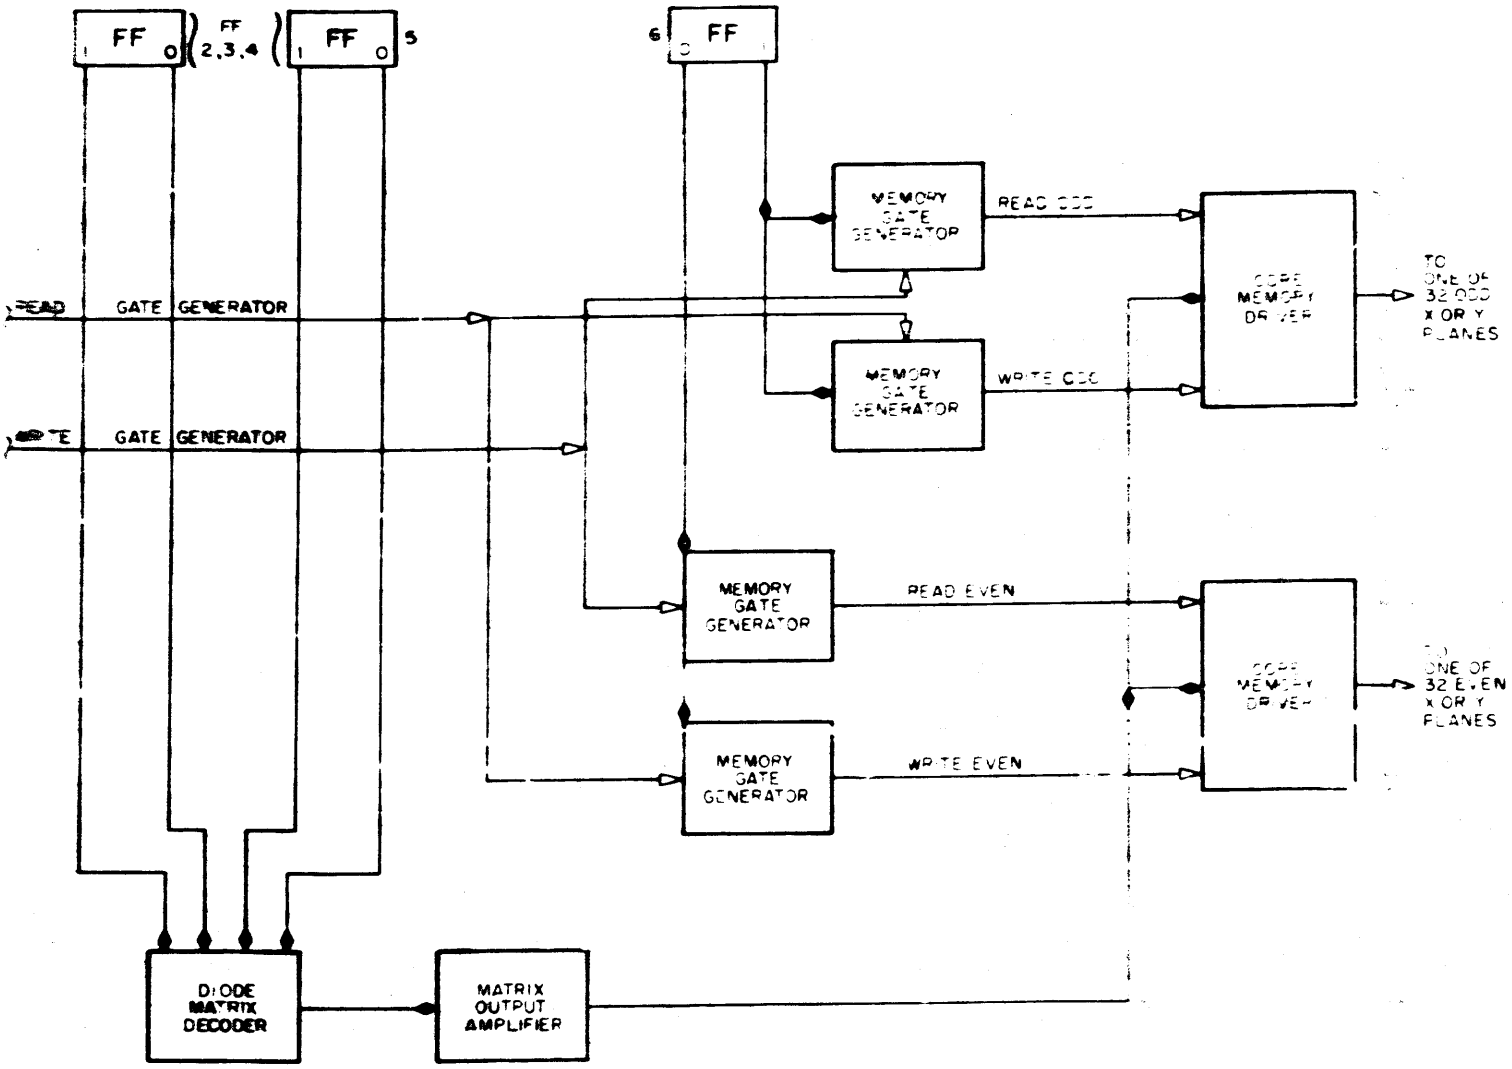

Figure 7—14. 64² Memory Driver Selection

SECTION 4
TIMING AND CONTROL

The timing and gating control section of the 64^2 core memory receives the start-memory pulse from the program element for both types of memory cycles and, during the execution of a memory store cycle, this section also receives the inhibit-sample pulse from the instruction control element. The timing and gating section provides the three gate signals and the sample pulse to operate the various portions of the core memory in proper sequence. A block diagram of the timing and gating section showing typical timing notations is shown in figure 7–15.

As shown in figure 7–15, the start-memory pulse is applied to the memory time pulse distributor, which is a chain of delay lines. The start-memory pulse delayed 0.1 μsec is used to set the read-gate generator. This is subsequently cleared by the start-memory pulse delayed approximately 2.3 μsec. The 1-side output of the read gate generator, which is a positive pulse, is applied to the selection section and determines the timing and duration of the read-current pulse that is supplied to the selected X and Y driver lines.

The sample gate generator of the timing and gating section differs from the other gate generators in this section in that the 0-side output controls a gate tube which is sampled by the sample pulse from the memory time pulse distributor. During a memory readout cycle, the inhibit-sample pulse is not generated, and the clear-and-sample gate generator remains in the 0 state. In this condition, the gate passes the sample pulse to the sense section approximately 1.9 μsec after the start-memory pulse is initiated. During the memory store cycle, an inhibit-sample pulse from the instruction control element is supplied to the sample-gate-generator flip-flop, setting the flip-flop to the 1 state. Thus, the gate tube is deconditioned when sensed by the sample pulse.

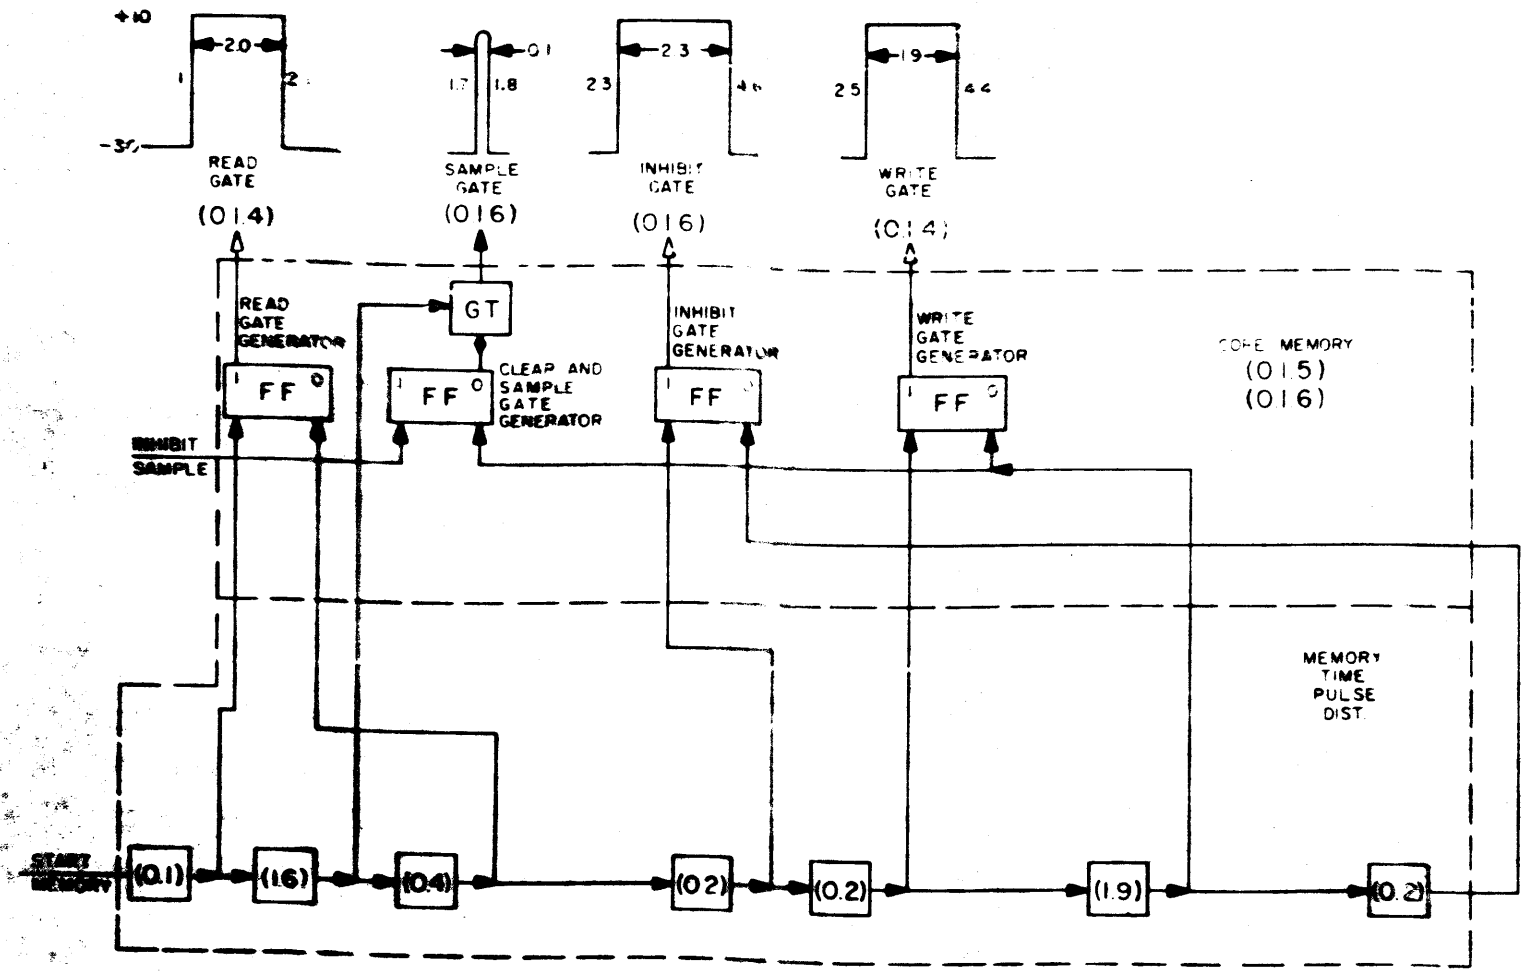

Figure 7–15. 64^2 Timing and Gating Circuits

The write-gate-generator flip-flop receives a set-write pulse from the memory time pulse distributor approximately 2.7 μsec after the start-memory pulse is initiated and receives a clear-write-gate-generator pulse approximately 1.9 μsec later. The resultant output pulse from the write-gate-generator flip-flop, which is approximately 1.9 μsec in duration, is applied to the selection section and determines the timing and duration of the write current pulse that is supplied to the selected X and Y driver lines.

The inhibit gate generator, shown in figure 7—13, is similar to the read and write gate generators in operation and function. The set-inhibit pulse is supplied by the memory time pulse distributor approximately 2.5 μsec after the start-memory pulse is initiated. The clear-inhibit pulse (from the same source) is received approximately 4.8 μsec after the start-memory pulse is initiated. The 1-side output pulse of the inhibit-gate-generator flip-flop is approximately 2.3 μsec in duration and overlaps the write-gate-generator output. The inhibit gate is applied to the 33 digit plane drivers; however, only the drivers which are conditioned by the associated memory buffer register flip-flops can generate an inhibit current pulse.

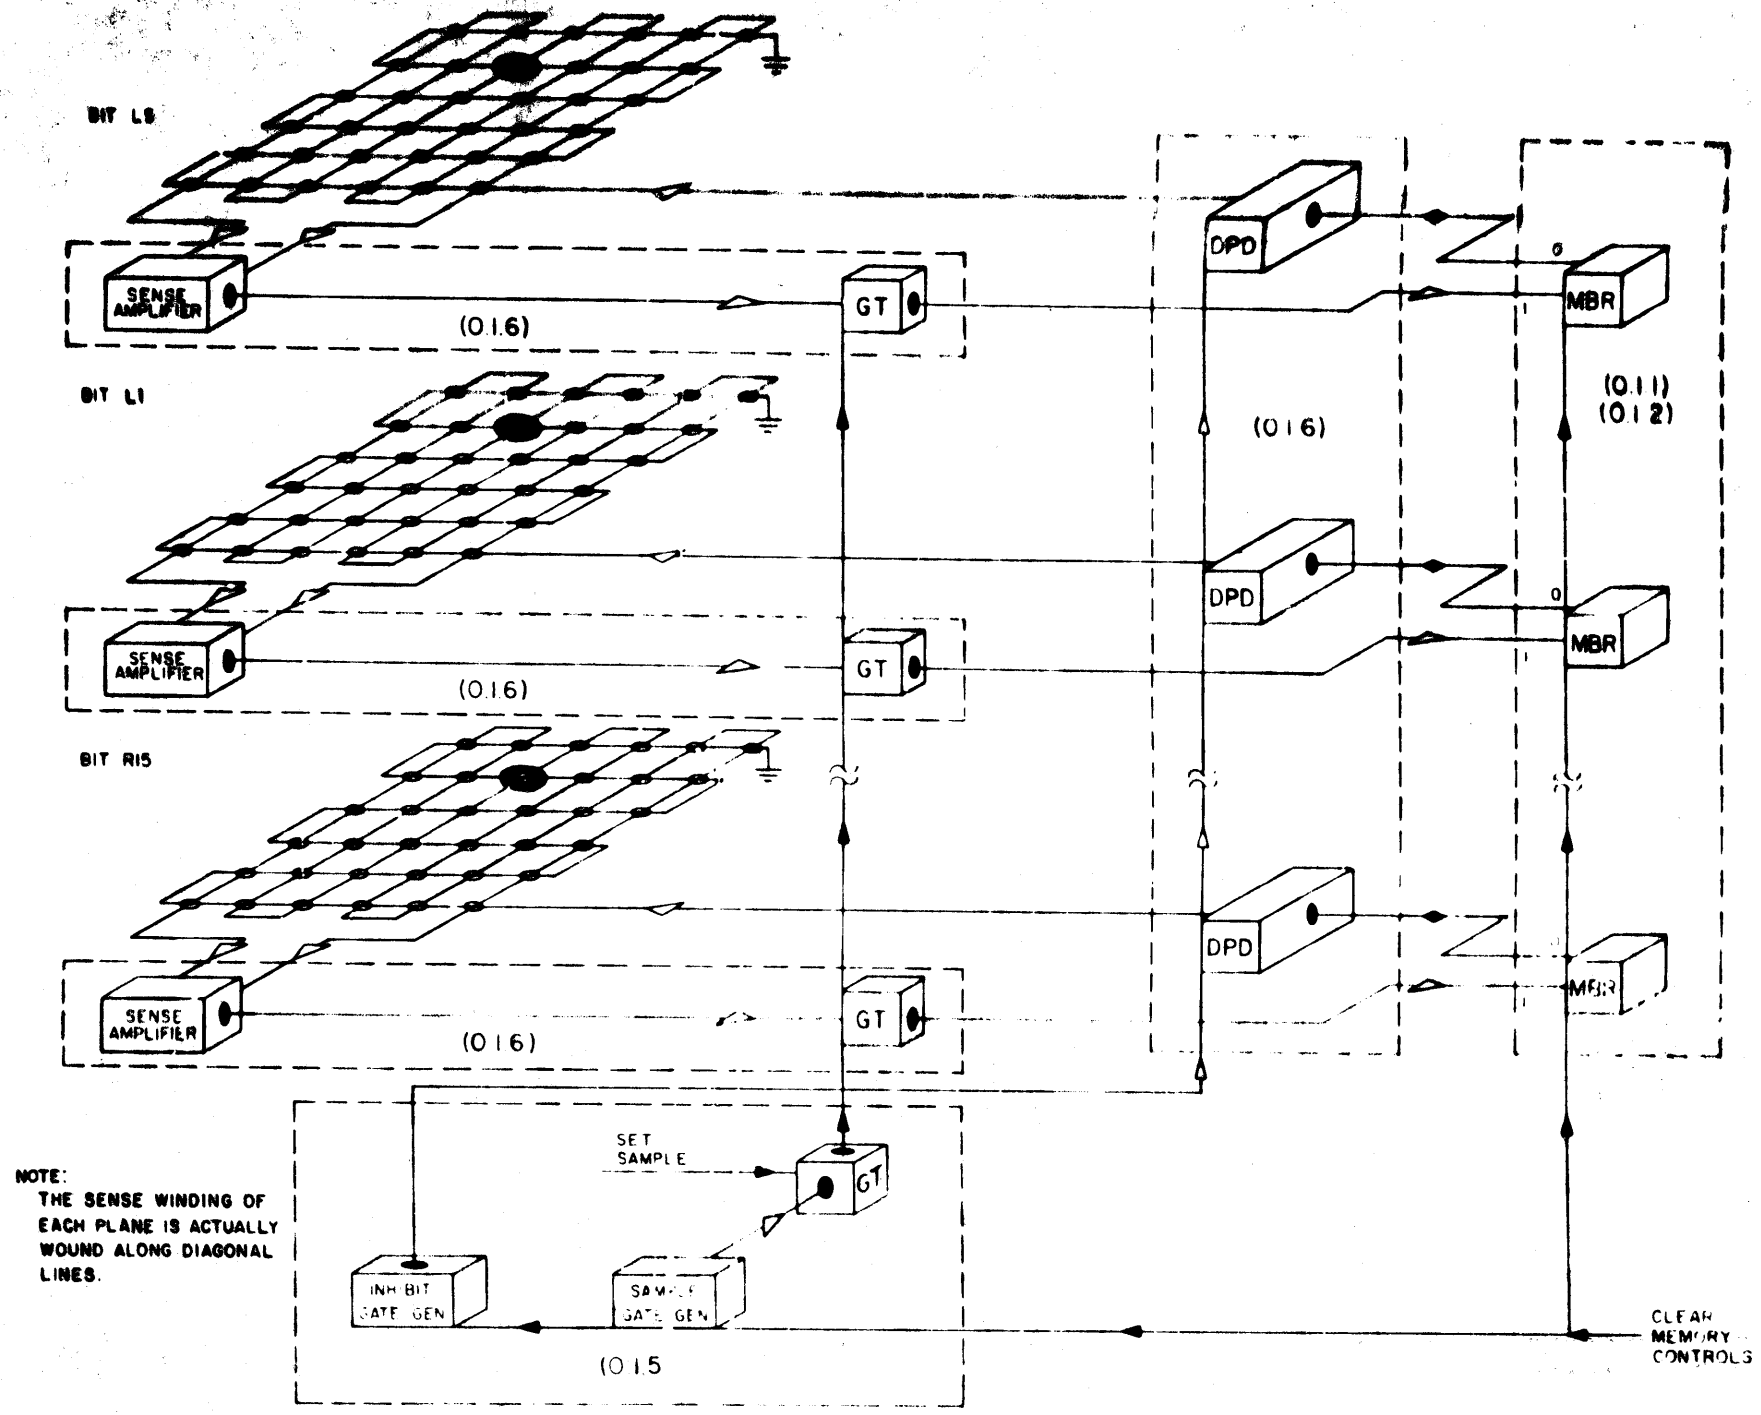

Figure 7–16. 64² Sense and Inhibit Winding Pictorial

SECTION 5
DIGIT PLANE CIRCUITRY

Figure 7–16 illustrates the relationship of the sense section, the memory buffer register, and the digit plane driver section, which constitute the digit plane circuitry.

5.1 SENSE SECTION AND MEMORY BUFFER REGISTER

The sense section of the 64² core memory consists of 33 amplifiers and 33 gate tubes, one set for each plane of the core memory array. Each individual sense amplifier amplifies the output voltages induced in the sense winding of the associated plane.

The output of a sense amplifier – a nonstandard pulse which is positive regardless of the polarity of the input signal – is applied to and conditions a gate tube circuit. If the selected core contains a 1 prior to being read out, the gate tube, when sampled, provides the standard pulse required to activate the memory buffer register. The gate tube is sampled by a standard pulse gated sample gate generator. The time relationship between the sense amplifier output and the sample pulse is adjusted so that the sample pulse occurs at approximately the peak of an amplified 1 output signal. Thus, if a core containing a 1 is read out during a readout cycle, the sampled gate tube develops an output which sets the associated memory buffer register flip-flop to the 1 state. If the core contains a 0, the flip-flop remains in the 0 state.

The output of the 0 side of the memory buffer register flip-flop conditions an associated digit plane driver in order to provide the necessary control to subsequently restore the original content to the selected core. If the core contains a 1 prior to readout, the X and Y write current pulses write a 1 back into the core. However, if the core contains a 0 prior to readout, the associated digit plane driver is conditioned, resulting in the generation of an inhibit current pulse (coincident with the X and Y write current pulses) which prevents the writing of a 1, thus effectively writing a 0.

5.2 DIGIT PLANE DRIVER SECTION

The digit plane driver section consists of 33 functionally identical circuits, one corresponding to each active plane of the core memory array. These circuits function to supply the inhibit current pulses to the associated planes.

The output of a digit plane driver is a negative current pulse having an amplitude of approximately 400 ma. Because of the inhibit winding geometry, the inhibit current pulse has the same effect as a half-select current pulse in the read direction. To inhibit the writing of a 1, the inhibit current pulse (effectively +I/2 amp) adds algebraically to the X-write current pulse (–I/2 amp) and the Y-write current pulse (I/2 amp) to produce an effective field of –I/2 amp turns [(+I/2) + (–I/2) + (–I/2) = –I/2].

The input stage of each digit plane driver consists of a d-c gate tube which is conditioned by the 0 side of the associated memory buffer register flip-flop when the selected core contains a 1. The other input to this gate tube is supplied by the inhibit gate generator, which is located in the timing and gating section. The output of the d-c gate tube supplies a cathode follower, a differential amplifier, and a d-c power amplifier which drives the low-impedance, high-current inhibit winding of the core memory array plane.

CHAPTER 4
THEORY OF OPERATION OF CORE MEMORY 1 (256^2)

SECTION 1
256^2 FERRITE CORE ARRAY

1.1 GENERAL

The 256^2 ferrite core array contained in Unit 66 (fig. 1–9) is the principal component of core memory 1 since it is the information storage center of this memory device. This section describes the physical characteristics of the ferrite core array to show the arrangement of its cores and windings. This description provides a definition of the terms that are used in the core memory logic analysis presented in the subsequent sections of this chapter.

The storage capacity of the 256^2 ferrite core array is equal to 65,536 words of 33 bits each. Since a single core can store one bit of information, the array contains 65,536 x 33 or 2,162,688 ferrite cores. These cores are arranged in a 3-dimensional array in which each horizontal layer or digit plane contains 65,536 cores arranged in a 256 x 256 square formation. The 33 digit planes of this array are stacked vertically and the X and Y selection windings of these planes are interconnected to form the X and Y selection windings of the array. Actuation of the current drivers associated with one X and one Y selection winding will mutually affect the vertical column of 33 cores (one core in each digit plane) that represents the selected memory register.

Each digit plane of the 256^2 ferrite core array is composed of 16 identical subplanes. Each subplane contains 4,096 cores arranged in a 64 x 64 square formation and all of the windings required for memory operation. Two distinct types of subplanes are used in the construction of this array in order to provide for wiring symmetry. The two types of subplanes are very similar and differ only in the manner in which the inhibit winding is wound. Since each digit plane is composed of identical subplanes, the 256^2 array is actually composed of two types of digit planes. One type of subplane is used in the construction of 17 digit planes, while the second type is used in the construction of the other 16 digit planes of the array.

The detailed analysis of array wiring is presented under the subsequent headings of this section. This analysis consists of separate discussions of: the internal wiring of each type of subplane; the interconnection of subplanes to form the two types of digit planes; and the interconnection of the X and Y digit plane windings to form the X and Y selection windings of the array.

1.2 SUBPLANE WIRING

The two types of subplanes used in the 256^2 ferrite core array are similar in that each contains 4,096 cores in a 64 x 64 formation and all the windings necessary for memory operation. Figures 7–17 and 7–18, respectively, show an abbreviated version of each type of subplane using 16 cores in a 4 x 4 formation. A comparison of the two figures reveals that the only difference between the two types of subplanes is that in the Type 1 subplane the inhibit winding is wound parallel to the X windings, while in the Type 2 subplane the inhibit winding is wound parallel to the Y windings.

1.2.1 X and Y Selection Windings

As noted in figures 7–17 and 7–18, each X selection winding is wired through a single row of cores, while each Y selection winding is wired through a single column of cores. In either type of subplane, selection of a particular core is accomplished by applying coincident half amplitude, read-write current pulses to the appropriate X and Y selection windings (those that link the desired core) in such a direction that the resultant magnetic fields will add at the intersection of the two windings. As a result of this requirement, and because of the manner in which the cores are positioned (refer to figs. 7–17 and 7–18), the direction of the read-write current pulses that are applied to an X or Y selection winding must be of the opposite polarity (applied in the opposite direction) when compared to the current pulses that are applied to an adjacent X or Y selection winding. Since all of the X and Y current drivers produce read-write current pulses of the same polarity pattern (fig. 7–4), the required reversal of current flow in adjacent selection windings is accomplished by connecting read-write current drivers to consecutive selection windings on alternate sides of the subplane. The manner in which the X and Y drivers are connected to the 256^2 core array to produce the desired result will be discussed in 1.4 of this section.

1.2.2 Inhibit Windings

The inhibit winding of each subplane of the 256² ferrite core array is associated with a digit plane driver (DPD). During the execution of a memory cycle, the associated DPD is controlled so that an inhibit-current pulse (negative) will be applied to the subplane inhibit winding (during the write portion of the cycle) if it is required to inhibit the writing of a 1 in the selected core of the subplane.

As noted in figures 7–17 and 7–18, the inhibit winding of each type of subplane consists of a single winding which is wound through all of the cores of the subplane in parallel with either the X or the Y selection windings. Although the two inhibit windings differ in wiring configuration, they both perform exactly the same function in that they each affect the subplane cores in exactly the same manner. Reference to the two figures shows that the direction of the inhibit winding is alternated in adjacent rows or columns of cores; therefore, the inhibit current pulse effectively flows in opposite directions through adjacent rows or columns of cores. Since write-current pulses (negative) are applied to adjacent X and Y selection windings in opposing directions (for both types of subplanes) the proper connection of the DPD will cause the individual magnetic fields to be set up by the inhibit-current pulse and the selected X and Y write-current pulses (in each core of the selected row and column of cores). In each type of subplane, all of the nonselected cores are affected by the inhibit current pulse in exactly the same manner.

1.2.3 Sense Winding

The sense winding of each subplane of the array is associated with a differential input sense amplifier which functions to amplify the induced voltages that are produced by the switching action of each core in the subplane. As noted in figures 7–17 and 7–18, each subplane has two separate sense windings, labeled 1 and 2. (In the digit plane, these individual windings are con-

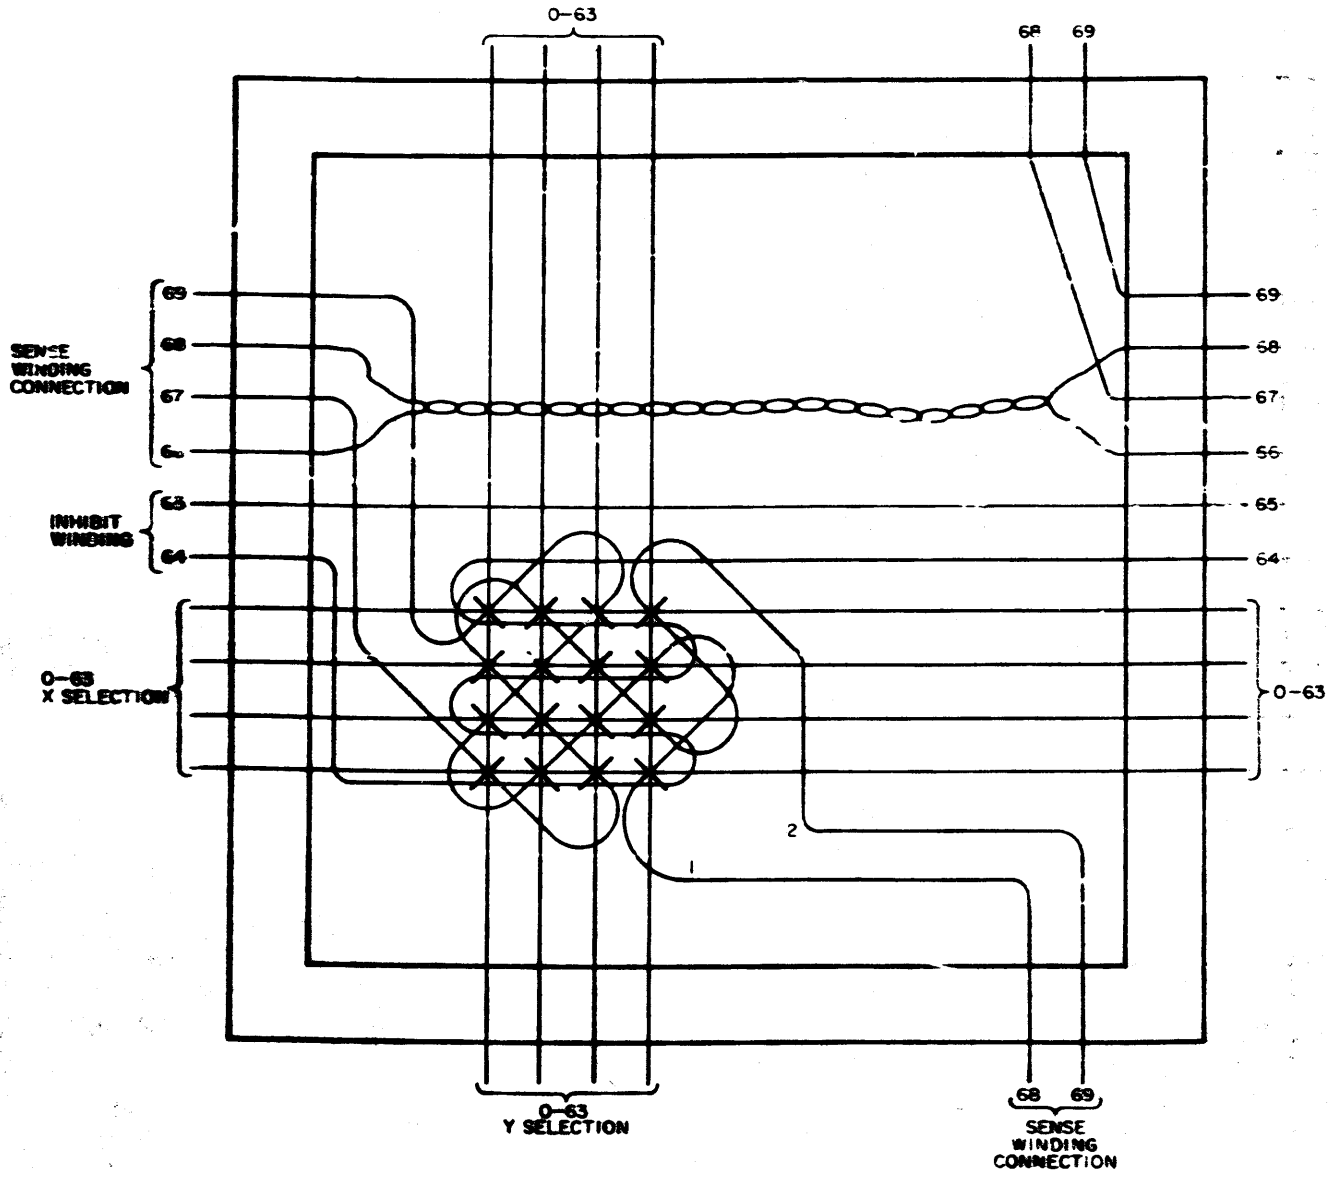

Figure 7–17. 256² Subplane Wiring, Type 1

nected in series by means of internal and external jumper wires to form one long winding.) Each winding passes through half of the cores of the subplane following diagonal paths in order to minimize the capacitive and inductive coupling between itself and the other windings of the subplane.

1.3 DIGIT PLANE WIRING

As previously stated, each 256^2 digit plane is composed of 16 identical subplanes. The arrangement and interconnections of these subplanes to form the two types of digit planes are shown in figures 7–19 and 7–20, foldouts. As noted, the 16 subplanes of each digit plane are numbered from 0 to 17 in octal notation. Since the subplanes of each digit plane are interconnected in the same manner, the following discussions are applicable to both types of digit planes.

1.3.1 Selection Windings

As shown in figures 7–19 and 7–20, the 64 x 64 Y lines of each subplane are connected to the similarly numbered lines of adjacent subplanes. In this manner, each of the X and Y lines becomes a continuous winding through the entire digit plane.

1.3.2 Inhibit Windings

As noted in figures 7–19 and 7–20, the inhibit windings of four subplanes in a row are connected in series to form one winding. As a result, four such independent windings are obtained, one for each of the four inhibit regions into which the digit plane is divided. The four inhibit region windings of each digit plane are associated with individual DPD's. These drivers are controlled so that inhibit current, if required, is only generated in the inhibit region which contains the selected core. The subplane groups which comprise the four inhibit regions of a digit plane are tabulated in table 7–5.

As shown in figures 7–19 and 7–20, the subplane feedthrough wires (internally connected between pin 65 on the left and right sides) of each inhibit region are

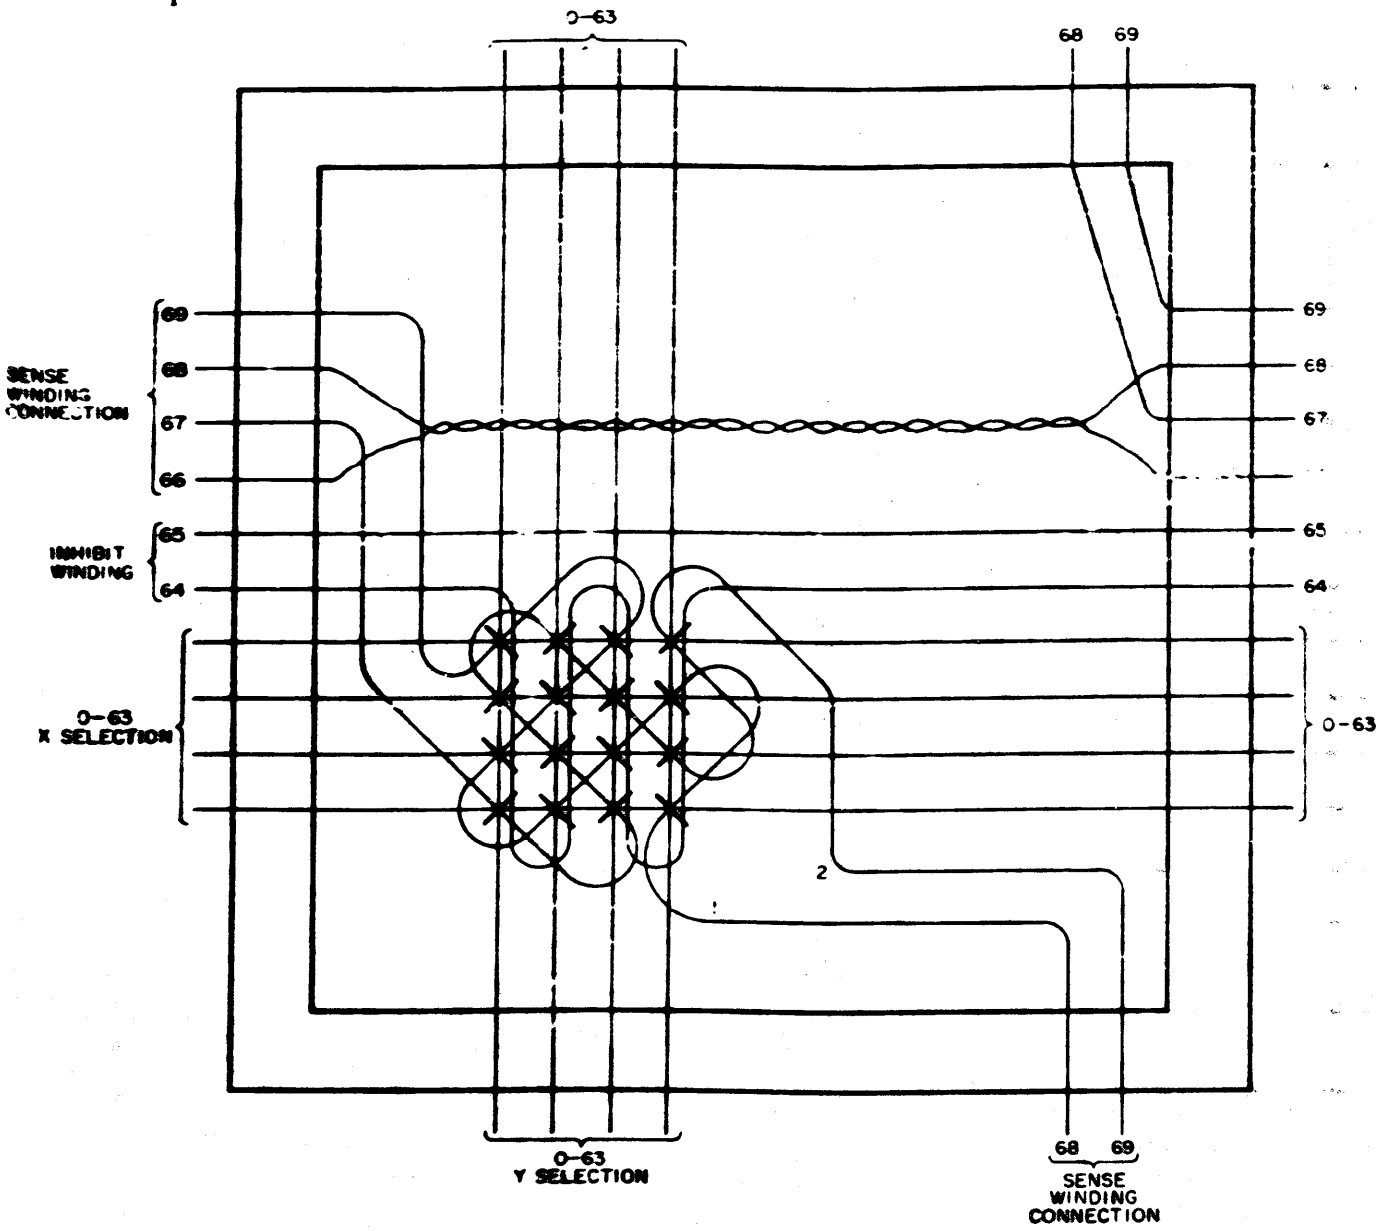

Figure 7–18. 256^2 Subplane Wiring, Type 2

TABLE 7–5. SUBPLANE GROUPING IN INHIBIT REGIONS

REGION	SUBPLANES
0	0, 4, 10, 14
1	1, 5, 11, 15
2	2, 6, 12, 16
3	3, 7, 13, 17

also connected in series to form a long wire. This feedthrough wire is connected to the inhibit region winding on one side of the digit plane (jumper between pins 64 and 65) so that the DPD input and terminating connections can both be made on the opposite side of the digit plane.

The external connections to the inhibit region windings readily identify the digit plane as being in either the left or right half-word. For example, figures 7–19 and 7–20 both typify a digit plane of the left half-word because the input and terminating connections (pins 64 and 65) for the inhibit region windings are located on the left side of the plane, while the feedthrough and inhibit windings are jumpered together (pins 64 and 65) on the right side of the plane. Conversely, the inhibit winding input and terminating connections for right half-word planes (not shown) are made on the right side of the digit planes, while the feedthrough and inhibit windings are jumpered together (pins 64 and 65) on the left side of the plane.

Because the X and Y array selection windings are formed by serially connecting the similarly numbered drive lines of adjacent digit planes, the direction of read-write current flow in the similarly numbered drive lines of adjacent digit planes will always be in mutually opposing directions. Since the inhibit-current pulse must always oppose the write-current pulses that are applied to the selected X and Y windings of each digit plane, the direction of the inhibit winding must also be reversed in adjacent digit planes. The required reversal in the inhibit winding direction of adjacent planes is accomplished by alternately connecting the DPD (associated with each digit plane) to pins 64 and 65 for each successive digit plane of the array. To illustrate the manner in which these connections are made, the following discussion refers to figure 7–21 which indicates the overall arrangement of the 33 digit planes of the array. Since the DPD's associated with the left half-word are all connected to the left side of the array, each DPD of the topmost or parity plane will be connected to pin 64 on the left side of the plane. Each DPD of the next lower or LS plane will be connected to pin 65 on the left side of the plane. This alternate pattern is repeated for the remainder of the left half-word with the result that each DPD of the L15 plane will be connected to pin 64 on the left side of the plane. Since the DPD's associated with the right half-word are connected to the right side of the array, each DPD of the RS plane will be connected to pin 64 on the right side of the plane. The alternate pattern of connecting the DPD's is resumed for each successive digit plane comprising the remainder of the right half-word.

1.3.3 Sense Windings

As noted in figures 7–19 and 7–20, the subplanes of each digit plane are grouped into four sense sections with four subplanes in each section. The sense windings in the four subplanes of each sense section are connected in a series parallel arrangement to provide common output points. These common output points are connected to the front side of the digit plane by means of the vertical twisted pair wires that were added to the subplanes of inhibit regions 0 and 1.

The grouping of subplanes to form the four sense sections of the digit plane is shown in table 7–6. As noted, sense section A consists of subplanes 2, 5, 10, and 17. The sense windings of subplanes 5 and 10 are connected in series, and the sense windings of subplanes 2 and 17 are also connected in series. These two series-connected groups of windings are connected in parallel (by means of the twisted pair wires in subplanes 0 and 1) to supply one input to the associated sense amplifier. The winding connections for sense section A may be traced on figures 7–19 and 7–20.

Pin 69 at the back of subplane 1 is connected to pin 69 on the left side of the jumpered connection into subplane 5, then through sense winding 1 of subplane 5 and sense winding 2 of subplane 10 to pin 69 at front of subplane 10. Pin 69 is connected to pin 68. From pin 68 the path goes through sense winding 1 of subplane 10 and sense winding 3 of subplane 5 to pin 68 at the back of subplane 1. The series connection of the sense windings of subplanes 2 and 17 is similar to that for subplanes 5 and 10 with the addition of the twisted pair connecting link running through subplanes 2, 6, 12, and 16.

From pins 68 and 69 at the back of subplane 1, the series-parallel connected subplane sense windings

TABLE 7–6. SUBPLANE GROUPING IN SENSE SECTIONS

SENSE SECTION	SUBPLANES
A	2, 5, 10, 17
B	3, 6, 11, 14
C	0, 7, 12, 15
D	1, 4, 13, 16

are connected through twisted pair links in subplanes 1 and 0 to pins 66 and 67 at the front of subplane 0. These terminals are the external connections of the sense section A sense windings.

The connection of the sense windings of the other sense sections is similar to that just described for section A. Signals from the four sense sections of one digit plane are applied to the four input preamplifier channels of one sense amplifier. Thus, as a result of this wiring scheme, a single sense amplifier can detect and amplify the output signal from any core in a digit plane.

1.4 ARRAY WIRING

The 256^2 ferrite core array, which contains a total of 33 digit planes, is composed of 16 Type 1 digit planes and 17 Type 2 digit planes. Figure 7–21 indicates the overall arrangement of these digit planes by identifying the bit designation of each digit plane of the array, and the inhibit regions and sense sections of each digit plane. Although not noted in the figure, the 17 odd planes (parity, L1, L3, . . . R15) of the array belong to the Type 2 group of planes, and the 16 even planes (L5, L2, L4, . . . R14) belong to the Type 1 group of planes.

1.4.1 Interconnection of X and Y Selection Lines

Selection of a memory register involves the simultaneous application of read-write current pulses to the similarly addressed core in each digit plane. Therefore, the corresponding X and Y drive lines in each of the 33 digit planes of the array will be involved in the operation. In order to provide a control so that one current drive can supply read-write current pulses to the corresponding X or Y drive line of each digit plane, the similarly numbered X and Y drive lines of all digit

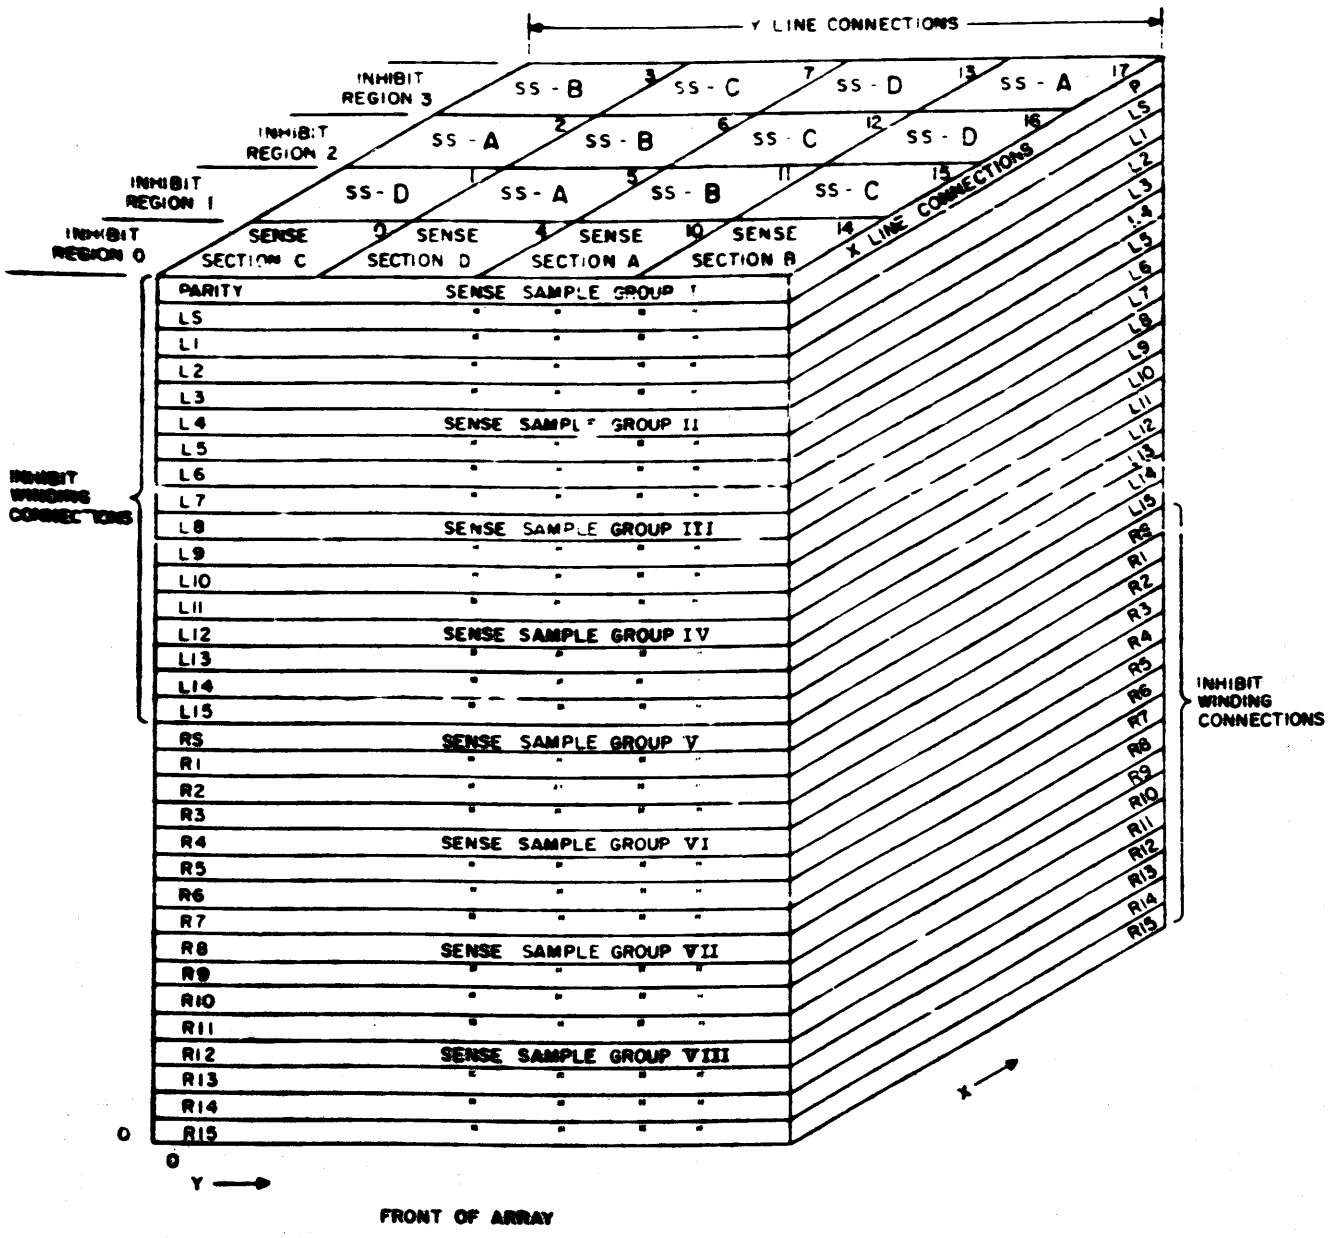

Figure 7–21. 256^2 Memory Array

planes are connected in series (by means of jumpers) so that common selection windings will be formed. One end of each of the 256X and 256Y selection windings is connected to a read-write current driver; the other end of each winding is connected to ground through a terminating resistor. In this ferrite core array all of the X selection line drivers (256) are connected to the left and right sides of the array while all of the Y selection line drivers (256) are connected to the front and rear sides of the array. In each case, the selection line drivers are connected to the parity plane while the selection line terminating resistors are connected to the bit R15 plane. Since the input and output connections of the X and Y selection windings are exactly the same, the following discussion will deal with the X selection windings only.

The 256 X read-write current, one for each X selection line, are connected to the array as shown in figure 7–22. As noted, current drivers for even numbered subplane lines (0, 2, 4, etc.) are connected to the even numbered pins on the right side of the parity plane. The even numbered pins on the left side of the parity plane are jumpered to similarly numbered pins on the left side of the bit LS plane, the bit LS plane is similarly connected to plane L1 on the right side of the array, etc. Alternate planes are similarly connected to include the 33 planes of the array. The terminating resistors for these even numbered lines are connected to the left side of the bit R15 plane and mounted on the lower left side of the array.

As noted in figure 4–6, drivers for odd-numbered subplane lines (1, 3, 5, . . . 63) are connected to the odd-numbered pins on the left side of the parity plane. The odd-numbered pins on the right side of the parity plane are jumpered to similarly numbered pins on the right side of the LS plane. Alternate planes are similarly connected to include all 33 planes of the array. The terminating resistors for these odd-numbered lines are connected to the right side of the R15 plane. These terminating resistors are mounted on the lower right side of the array.

The Y selection windings are connected on the front and back sides of the array in exactly the same manner as described above. The even-numbered Y line drivers are connected to the even-numbered subplane pins (0, 2, 4, etc.) on the front side of the parity plane; the odd-numbered Y line drivers are connected to the odd-numbered subplane pins (1, 2, 5, etc.) on the rear side of the parity plane. The terminating resistors for even-numbered Y lines are mounted on the lower rear

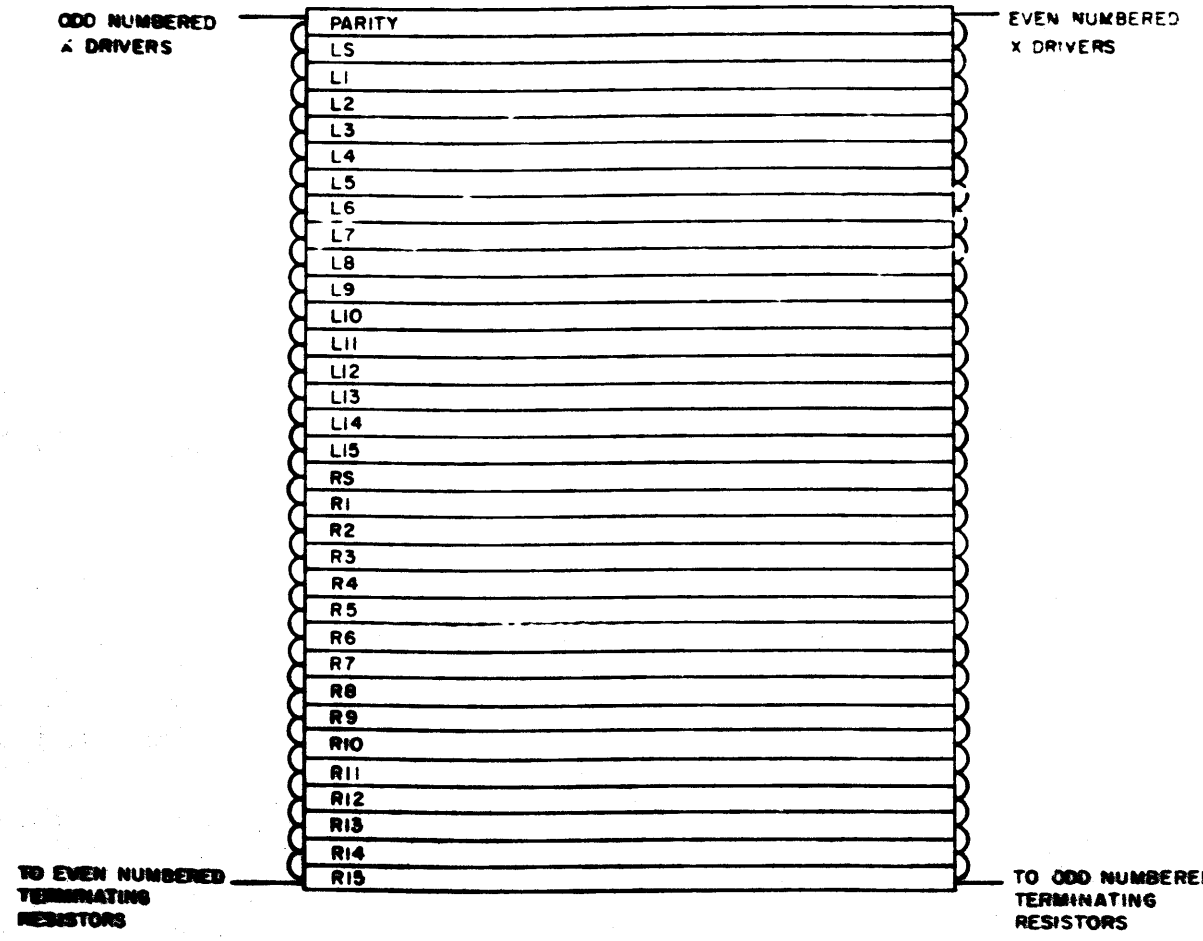

Figure 7–22. Front View of 256² Ferrite Core Array

side of the array; the odd-numbered Y line terminating resistors are mounted on the lower front side of the array.

1.4.2 Sense Winding Sample Groups

Because the X and Y selection windings are electrically long, a definite amount of time is required for the read-write current pulses to appear at any specified point on the line; that is, these current pulses are delayed slightly as they travel through the array. This delay in the selection-line-current pulses affects the array readout timing, since all of the selected cores are not being switched at the same time. For optimum results, when a core is switched, the sense amplifier output is at a maximum. To insure optimum results the sense amplifiers are grouped to be sampled progressively. Figure 7–21 indicates the sample time grouping of the digit planes of the array. The time difference between these sense group sample pulses is approximately 0.04 μsec.

SECTION 2
BLOCK DIAGRAM ANALYSIS

A block diagram of the 256^2 memory is shown in figure 7–23. As noted in this figure, the 256^2 memory is divided into five sections: the selection section, the array section, the sense section, the DPD section, and the timing and gating section. Figure 7–23 also includes the MBR since it is the only path by which information can be transferred into or out of the memory. This register is not a part of the 256^2 memory since it performs a common function for all three of the computer memories.

Normal operation of the 256^2 memory consists of transferring information into or out of specific registers of the core array. As a result, two specific types of memory cycles are possible; i.e., a readout cycle during which old information is obtained from selected cores, and a store cycle during which new information is stored in the selected cores.

The following sequence of events occurs during execution of a readout cycle:

a. The desired address information is transferred to the selection circuits to condition the desired X and Y line drivers.

b. A start memory pulse is generated.

c. The MBR is cleared to prepare it for the subsequent information transfer.

d. Read-current pulses are applied to the selected X and Y lines to switch all of the selected cores to the 0 state.

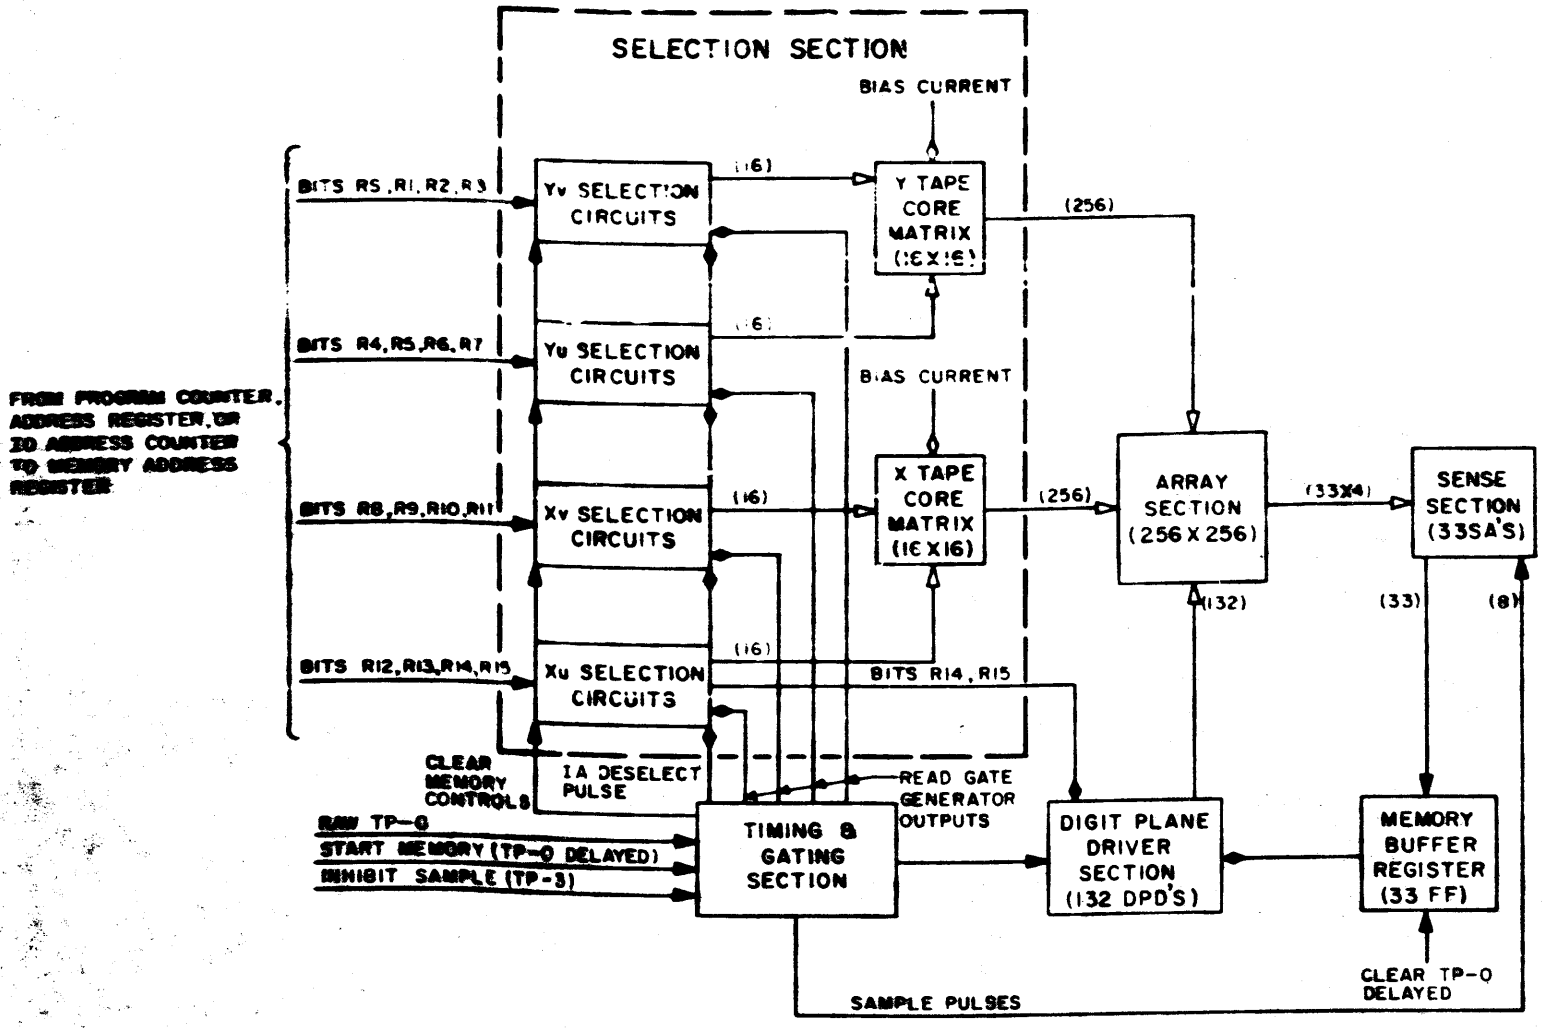

Figure 7–23. 256^2 Memory Device, Block Diagram

e. A sample pulse is generated to sample the output of the sense amplifiers. Individual outputs are used to set associated bits of the MBR.

f. Write-current pulses are applied to the selected X and Y lines to rewrite the information just read out. An inhibit-current pulse is applied to selected inhibit region of each plane, if required to prevent the writing of a 1 in the selected core.

g. The memory selection circuits are reset to prepare them for the next memory cycle.

The following sequence of events occurs during execution of the store cycle:

a. The desired address information is transferred to the selection circuits to condition the desired X and Y line drivers.

b. A start-memory pulse is generated.

c. The MBR is cleared.

d. New information is transferred to the MBR from an external source and an inhibit-sample pulse is generated.

e. Read-current pulses are applied to the selected X and Y lines to switch all of the selected cores to the 0 state, thus in effect erasing the old content.

f. Write-current pulses are applied to the selected X and Y lines to write the new information into the selected cores. An inhibit-current pulse is applied to the selected inhibit region of each plane, if required to prevent the writing of a 1 in the selected core.

g. The memory selection circuits are reset to prepare them for the next memory cycle.

Actually, during the operation of the 256^2 memory, a third type of memory cycle is possible, namely, the memory not selected cycle. The only operation performed in the 256^2 memory during such a condition (i.e., a machine cycle is executed but does not select the 256^2 memory) is that the memory selection circuits are reset to prepare them for the next memory cycle.

As noted above, the operations performed during the execution of either a readout cycle or a store cycle are very similar, the only difference being that in the readout cycle a sample pulse is generated, while in the store cycle, the MBR is loaded from an external source. Since the two cycles are similar, the following analysis of figure 7–23 will be based on a readout cycle. This analysis will not discuss the various timing details since this information is covered in subsequent sections.

At TP 0, the content of the program counter, address register, or IO address counter, which specifies the desired memory address, is transferred to the MAR (which were reset during the preceding memory cycle), four bits of information being transferred to each of the four selection circuit groups. The portion of the desired address contained in each selection circuit group is then decoded to condition one of its 16 tape core drivers; thus one tape core driver is conditioned in each of the four selection circuit groups.

Also at TP 0, a clear-memory-controls pulse (row TP 0) is applied to a 1.5-μsec delay line in the timing and gating section. If the 256^2 memory is not selected for operation, then this delayed pulse will be gated to clear the MAR to prepare it for the next memory cycle. If the 256^2 memory is selected for operation, then this delayed TP 0 pulse is suppressed and the MAR is cleared later in the memory cycle by a delayed start memory 1 pulse.

If the 256^2 memory is selected for operation, then a start-memory-1 pulse is generated at approximately TP 0 +0.4 μsec and applied to the timing and gating section. This pulse is applied to the memory pulse distributor (delay line clock) and also sets the Xv and Yv read gate generators. The output levels of these read gate generators are applied to their associated selection circuit groups to activate the conditioned tape core driver in each group. As a result, a read current pulse is generated on one of the 16 Xv and one of the 16 Yv matrix selection lines to partially select one line of 16 tape cores in each tape core matrix.

Note

The tape core matrix contains 256 tape cores connected in a 16 by 16 square formation. Selection of a specific core is based on the coincident current principle. However, because the cores are affected by a biasing current (a separate biasing winding is used), the selected tape core will generate a positive output when the coincident read-current pulses are applied and a negative output when the read-current pulses are terminated. Details of tape core selection together with a discussion of the resultant output pulses are contained in Section 3 of this chapter.

If the 256^2 memory is selected for operation, then the remaining internal operations are controlled by delayed-start-memory-control pulses which are obtained by tapping the delay line at various points. The first clock pulse is used to set the Xu and Yu read gate generators to initiate a read-current pulse on the selected Xu and Yu matrix selection lines. Since the Yv and Vv read gate generators have already been set, coincident read-current pulses are applied to one U and one V selection line of each tape core matrix. The tape core at the intersection of the selected U and V lines is switched to produce a positive output current pulse of sufficient amplitude and duration to read half select the

tape core matrices operate in unison, coincident read-ferrite cores on one selection line of the array. Since the current pulses are supplied to the 256^2 array to cause the 33 selected cores to be switched to the 0 state. The selected core output signals are applied to the sense section where they are amplified and applied to condition 33 gates. These gates are sensed by a clock-sample pulse; the conditioned gates will pass the sample pulse to set the associated bits of the MBR which was cleared at TP 1.

Clock pulses now clear the read gate generators, terminating the read-current pulses applied to the selected lines of the tape core matrices. Termination of the read-current pulses causes the two selected tape cores (one X and one Y) to be switched back to their original state. This action produces negative output current pulses of sufficient amplitude and duration to write half select the ferrite cores on the selected lines (one X and one Y) of the array. Since the two current pulses are generated in unison, the ferrite cores at the intersection of the selected X and Y line will be switched to the 1 state. However, the outputs of the MBR combined with the inhibit gate generator output are used to control the action of the 33 DPD's associated with the selected inhibit region of the array. If a particular bit of the original memory word contained a 0, the associated memory buffer flip-flop will contain a 0. Under this condition the associated DPD will generate an inhibit-current pulse in the selected inhibit region of the array. Since the inhibit-current pulse is timed to actually overlap the write-current pulses applied to the array, the selected core of the associated plane will not be switched to the 1 state, but will remain in the 0 state.

During the latter portion of the memory cycle, a clock pulse is used to reset the memory address register (each selection circuit group contains 4 of the MAR flip-flops) in preparation for the next memory cycle. Since the output levels of the MAR are only used during the first half of the memory cycle (to specify which tape core drivers are to be activated), this reset action does not affect the remaining internal operations of the memory cycle being executed.

SECTION 3
ADDRESS SELECTION

3.1 ADDRESS DESIGNATION

The memory address register (MAR) of the 256^2 memory element, which specifies the selected memory address, has been previously identified (refer to Part 4 of this manual) as a 16-bit register consisting of bits RS through R15. Since the digit plane used in this memory consists of a 256 by 256 array of cores, which are selected by a coincident current method (the selected core is defined as the core at the intersection of mutually perpendicular X and Y drive lines), the MAR content actually specifies one X and one Y drive line to select the desired core. Because the digit plane is symmetrical, the 16 binary bits of the MAR are divided into two equal groups, bits RS through R7 and bits R8 through R15. These two groups of binary bits are used to specify the selected Y and X drive lines respectively. Since the MAR content is usually designated by a 6-digit octal number, a conversion scheme is employed to convert the octal address designation into a short symmetrical system of X and Y drive line designation.

As shown in figure 7–24, the MAR bits have been relabeled according to an X and Y notation, and regrouped to form two 3-digit numbers, based on an octal-quadral system of notation. The first two digits of each of these numbers are in octal notation (radix of 8), and the third digit is in quadral notation (radix of 4). To determine the X and Y line octal-quadral designation (ocquad) of an octal address, or vice versa, the binary equivalent of the given number must be regrouped to form the six digits of the desired notation.

Example:

Octal Address	0	1	2	3	4	5
Binary Equivalent	0	0	0	1	0	1	0	0	1	1	1	0	0	1	0	1
Octal-Quadrant Equivalent	0		5	0		7	1		1
			Y				X		

Note

Since the third digit of an ocquad (octal-quadral) number represents two binary bits, the maximum value of this digit is 3. If a 1 is added to an ocquad number ending in 3, this maximum is exceeded, and the third digit develops a carry that is added to the next significant digit; namely, digit 2.

Example:

$543_{(\text{ocquad})} + 1 = 550_{(\text{ocquad})}$

In all cases, numbers to be added to an ocquad number must also be in ocquad notation.

As noted in figure 7–24, the physical X and Y drive lines of each subplane are no longer referenced to the individual subplanes; rather, for addressing purposes, the X and Y drive lines are numbered in consecutive decimal order from 0 to 255. It should also be noted (from the ocquad designation shown for the associated physical lines) that consecutive X lines are contained in consecutive inhibit regions, and consecutive Y lines are contained in consecutive columns of subplanes. The relationship of the physical line notation to the ocquad line designation obtained from the octal address designation is shown in table 7–7. Since the MAR is grouped in a symmetrical system, the table is applicable for both X and Y line designation.

In using the table to determine the physical line specified by an ocquad number, the least significant digit (3rd) of the ocquad number should be examined first to pin-point the applicable columns as follows:

CONTENT OF 3RD DIGIT	PHYSICAL LINE COLUMNS
0	1
1	2
2	3
3	4

Since the first two digits of the ocquad numbers in each pair of columns are listed in consecutive order from 00 to 77, a brief search will yield the desired results.

Example:

Ocquad number 612 = physical line 177
The third digit of the example specifies that the ocquad number is contained in column 3 of table 7–7.

Reference to the table shows that the ocquad number specifies physical line 177.

A study of the new addressing scheme reveals that selected cores can also be located by first identifying the subplane specified, and then determining which X and Y drive lines are specified within the selected subplane. Reference to figure 7–24 shows that the individual subplanes within a digit plane are numbered in consecutive octal order from 0 through 17. It should also be noted that the third digit of the X line octaquad designation is always the same within an inhibit region, and the third digit of the Y line octaquad designation

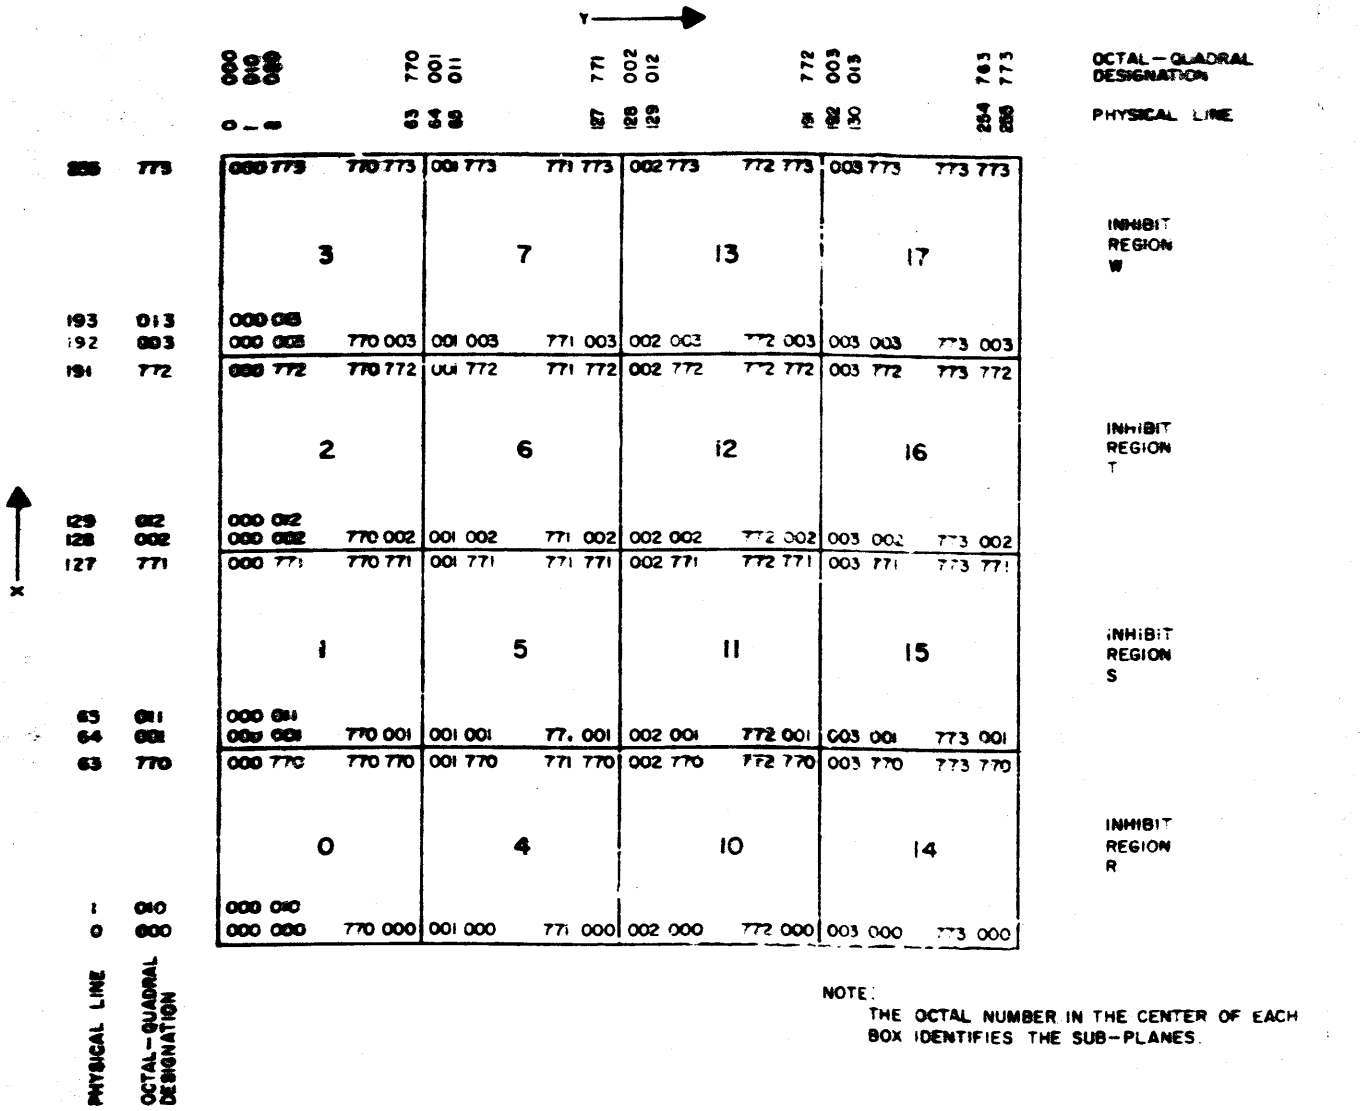

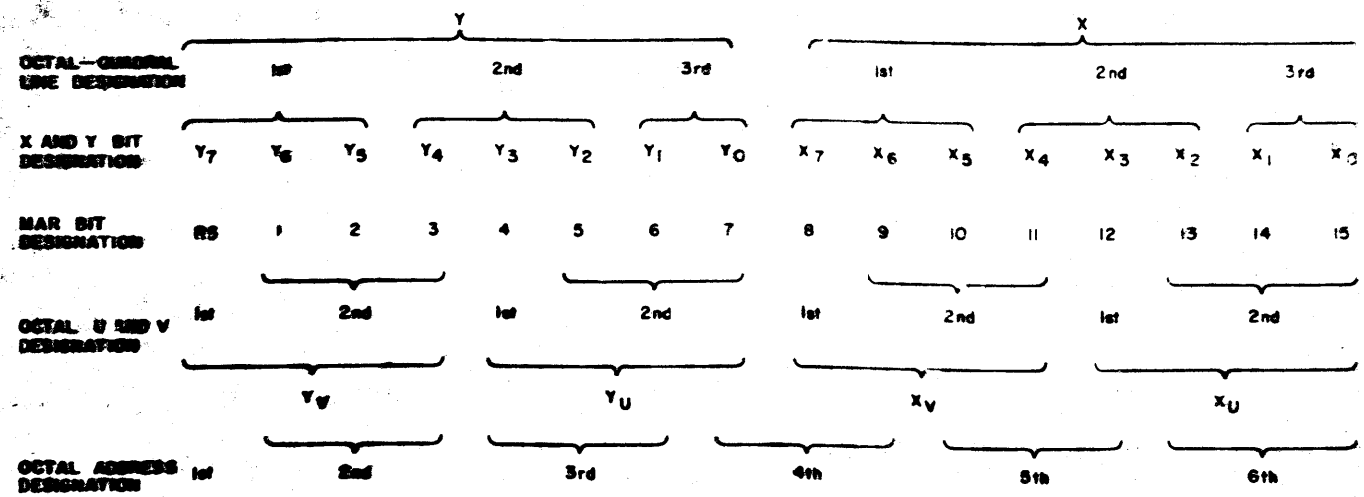

Figure 7–24. Addressing of 256² Memory

TABLE 7–3. RELATIONSHIP OF PHYSICAL LINES TO OCQUAD LINE DESIGNATION

PHYS LINE	X-Y OCQUAD ADR	PHYS LINE	X-Y OCQUAD ADR	PHYS LINE	X-Y OCQUAD ADR	PHYS LINE	X-Y OCQUAD ADR
0	000	64	001	128	002	192	003
1	010	65	011	129	012	193	013
2	020	66	021	130	022	194	023
3	030	67	031	131	032	195	033
4	040	68	041	132	042	196	043
5	050	69	051	[illegible]	052	197	053
6	060	70	061	134	062	198	063
7	070	71	071	135	072	199	073
8	100	72	101	136	102	200	103
9	110	73	111	137	112	201	113
10	120	74	121	138	122	202	123
11	130	75	131	139	132	203	133
12	140	76	141	140	142	204	143
13	150	77	151	141	152	205	153
14	160	78	161	142	162	206	163
15	170	79	171	143	172	207	173
16	200	80	201	144	202	208	203
17	210	81	211	145	212	209	213
18	220	82	221	146	222	210	223
19	230	83	231	147	232	211	233
20	240	84	241	148	242	212	243
21	250	85	251	149	252	213	253
22	260	86	261	150	262	214	263
23	270	87	271	151	272	215	273
24	300	88	301	152	302	216	303
25	310	89	311	153	312	217	313
26	320	90	321	154	322	218	323
27	330	91	331	155	332	219	333
28	340	92	341	156	342	220	343
29	350	93	351	157	352	221	353
30	360	94	361	158	362	222	363
31	370	95	371	159	372	223	373
32	400	96	401	160	402	224	403

TABLE 3-3. RELATIONSHIP OF PHYSICAL LINES TO OCQUAD LINE DESIGNATION (cont'd)

PHYS LINE	X-Y OCQUAD ADR	PHYS LINE	X-Y OCQUAD ADR	PHYS LINE	X-Y OCQUAD ADR	PHYS LINE	X-Y OCQUAD ADR
33	410	97	411	161	412	225	413
34	420	98	421	162	422	226	423
35	430	99	431	163	432	227	433
36	440	100	441	164	442	228	443
37	450	101	451	165	452	229	453
38	460	102	461	166	462	230	463
39	470	103	471	167	472	231	473
40	500	104	501	168	502	232	503
41	510	105	511	169	512	233	513
42	520	106	521	170	522	234	523
43	530	107	531	171	532	235	533
44	540	108	541	172	542	236	543
45	550	109	551	173	552	237	553
46	560	110	561	174	562	238	563
47	570	111	571	175	572	239	573
48	600	112	601	176	602	240	603
49	610	113	611	177	612	241	613
50	620	114	621	178	622	242	623
51	630	115	631	179	632	243	633
52	640	116	641	180	642	244	643
53	650	117	651	181	652	245	653
54	660	118	661	182	662	246	663
55	670	119	671	183	672	247	673
56	700	120	701	184	702	248	703
57	710	121	711	185	712	249	713
58	720	122	721	186	722	250	723
59	730	123	731	187	732	251	733
60	740	124	741	188	742	252	743
61	750	125	751	189	752	253	753
62	760	126	761	190	762	254	763
63	770	127	771	191	772	255	773

is always the same within a column of subplanes. As a result, the octal subplane designation corresponds to and is readily identified by regrouping the binary equivalent of the third and sixth digits of the ocquad address designation. These two ocquad digits, which consist of MAR bits Y_2, Y_0, X_1, and X_0, are regrouped so that the first digit of the octal subplane designation is specified by the content of bit Y_1, while the second digit is specified by the sum of bits Y_0, X_1, and X_0. Table 7–8 lists the four bit combinations required to select each of the 16 subplanes. Since the first two digits of the X and Y line ocquad designation specify consecutive physical X and Y lines (in octal notation) within a subplane, it is only necessary to determine which subplane is selected in order to completely identify the selected X and Y drive lines. The following example will serve to illustrate this fact.

Ocquad address	123 451
Binary equivalent of third and sixth digits	1101
Ocquad equivalent of third and sixth digits	15

Thus, ocquad address 123 451 selects the core at the intersection of $Y = 12_{(8)}$ (10_{10}) and $X = 45_8$ (37_{10}) of subplane 15_8.

TABLE 7–8. SUBPLANE SELECTION

MAR BIT GROUPING				SELECTED SUBPLANE
Y_1	Y_0	X_1	X_0	(OCTAL)
0	0	0	0	0
0	0	0	1	1
0	0	1	0	2
0	0	1	1	3
0	1	0	0	4
0	1	0	1	5
0	1	1	0	6
0	1	1	1	7
1	0	0	0	10
1	0	0	1	11
1	0	1	0	12
1	0	1	1	13
1	1	0	0	14
1	1	0	1	15
1	1	1	0	16
1	1	1	1	17

Numbers added to or subtracted from an ocquad number must of necessity be in ocquad notation. Thus, if it is desired to add the index register content to a memory address specified in ocquad notation, it is necessary to first convert the index register content into ocquad notation. In performing the actual addition, it must be noted that the third and sixth digits of the ocquad number are actually in quadral notation (radix of 4); therefore, the maximum number that can be expressed by these digits is 3. If this value is exceeded during the addition of quadral digits, a 4 is subtracted from the digit sum to develop a carry of 1 to the next most significant digit. The numerical difference represents the value of the quadral digit. When subtracting ocquad numbers (e.g., modification of the index register content by the index interval), these same precautions must be observed. The following examples serve to illustrate the addition and subtraction processes.

	OCTAL NOTATION	OCQUAD NOTATION
(1) Memory Address	0.03163	012 343
Index Register	0.01234	002 470
Sum	0.04417	021 033
(2) Index Register	0.01234	002 470
Index Interval	0.00452	001 122
Difference	0.00562	001 342

3.2 ADDRESS SELECTION CIRCUITS

The address selection process consists of decoding the contents of the memory address register and determining, on the basis of the information thus obtained, which X line and which Y line in the ferrite core array is to be driven in the subsequent reading and writing process. When read- and write-current pulses are applied to the selected array lines, the read and write processes take place in the cores of the selected address; that is, in the cores at the junctions of the selected X and Y lines.

3.2.1 Block Diagram Analysis

A simplified block diagram of the 256^2 memory address selection and X-Y driving circuitry is shown in figure 7–25. As noted in the figure, the MAR content – which is transferred from either the program counter, address register, or IO address counter – is decoded by four groups of selection circuits. Two of these circuit groups are used to drive the U and V selection lines of one tape core (CCD) matrix. Each circuit group decodes four bits of the address to select one of its 16 switch drivers (SWD). When a read-gate signal is applied to the associated current regulator (CR), the selected switch driver delivers a current pulse

to one U or one V selection line of its associated tape core (CCD) matrix. Each matrix receives a current pulse on both a U and V line and, as a result, a current pulse is generated on one of its 256 output lines. This current pulse drives one selection line of the ferrite core array. One tape core matrix drives an X line, and the other drives a Y line to complete the selection and driving process.

3.2.2 Tape Core Matrix

Figure 7–26 shows the electrical characteristics of the tape core used in the tape core matrices. The U, V, and bias windings produce approximately equal applied

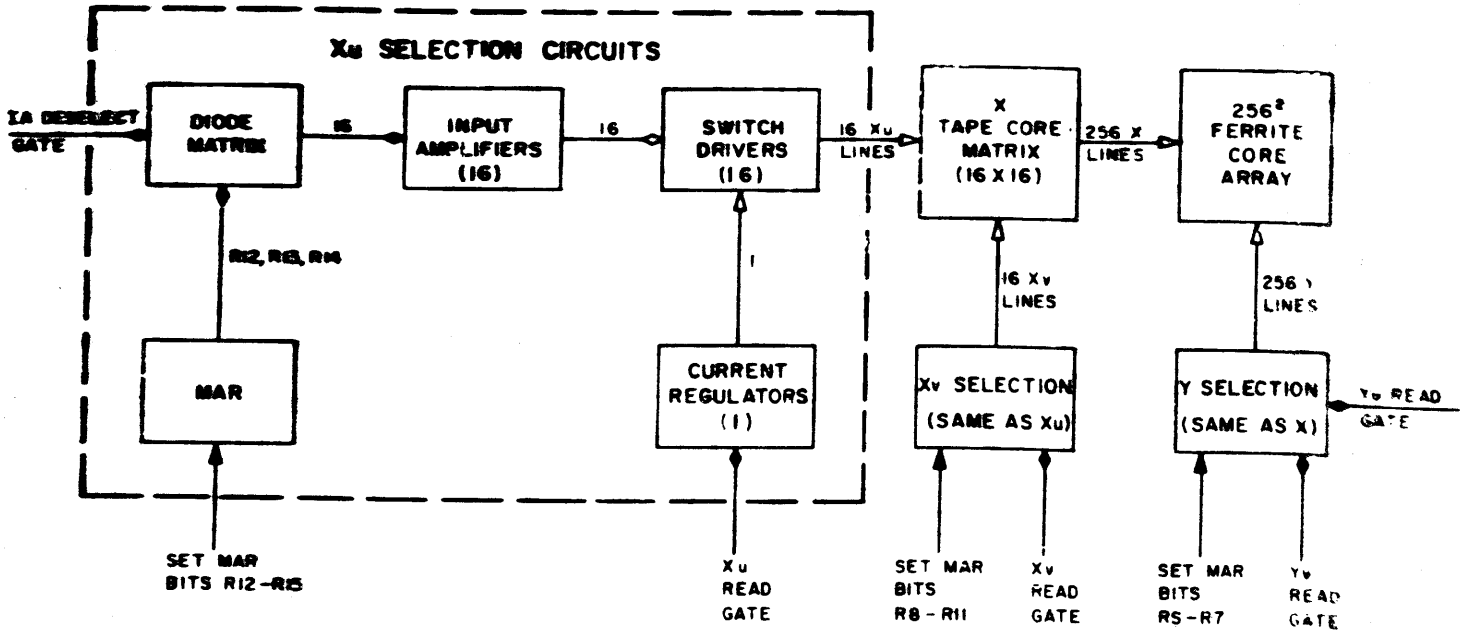

Figure 7–25. 256² Memory Selection Circuits, Block Diagram

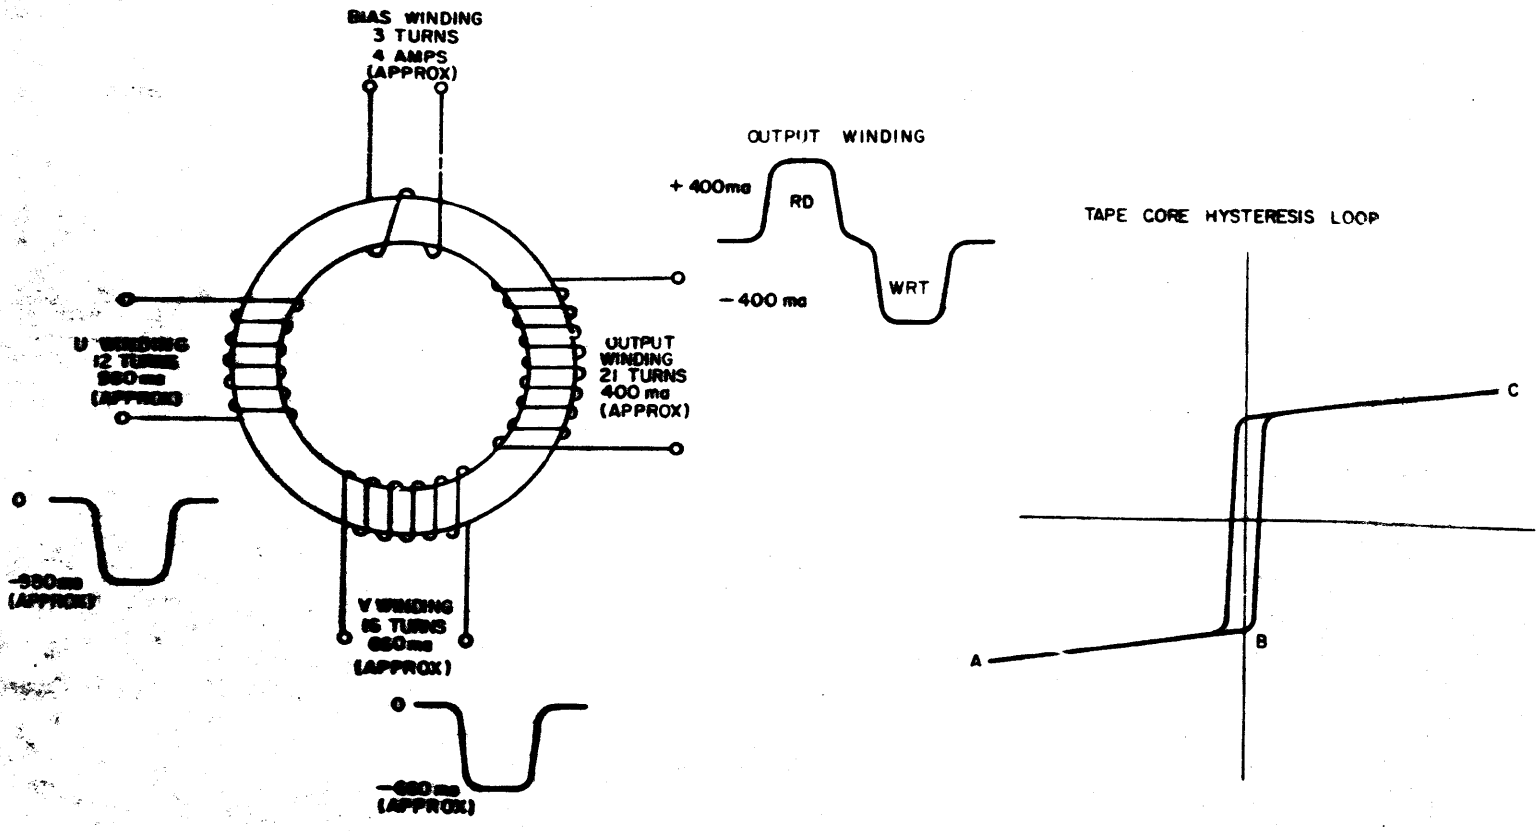

Figure 7–26. Tape Core Characteristics

fields; however, the windings are connected in such a manner that the bias field opposes the U and V fields. Under nonselected conditions, the U and V windings are not energized and the bias field will cause the core to settle at point A of the hysteresis loop. If a tape core is half selected, that is, if either the U or V windings is energized singly, the resultant field will cancel the bias field and the core will settle at point B. Since only a relatively small change in flux is produced, the output winding signal will be an early peaking positive pulse of approximately 40 ma at its peak. For half selected tape cores, a similar negative output signal is obtained when the U or V current pulse is terminated. If a tape core is fully selected, that is, if both the U and V windings are energized simultaneously, the resultant field will cause the core to switch to point C, producing a relatively large change in flux to yield approximately a 400-ma signal on the output winding. This signal is the read-current pulse used to half read select a ferrite core. When the U and V current pulses are terminated the bias field will cause the tape core to switch back to point A, again producing a relatively large change in flux to yield approximately a 400-ma signal on the output winding. This latter signal is the write-current pulse used to half write select a ferrite core.

Figure 7–27 shows the arrangement and connection of the tape cores to form a portion of the tape core matrix. As previously stated, only one U and V line can be

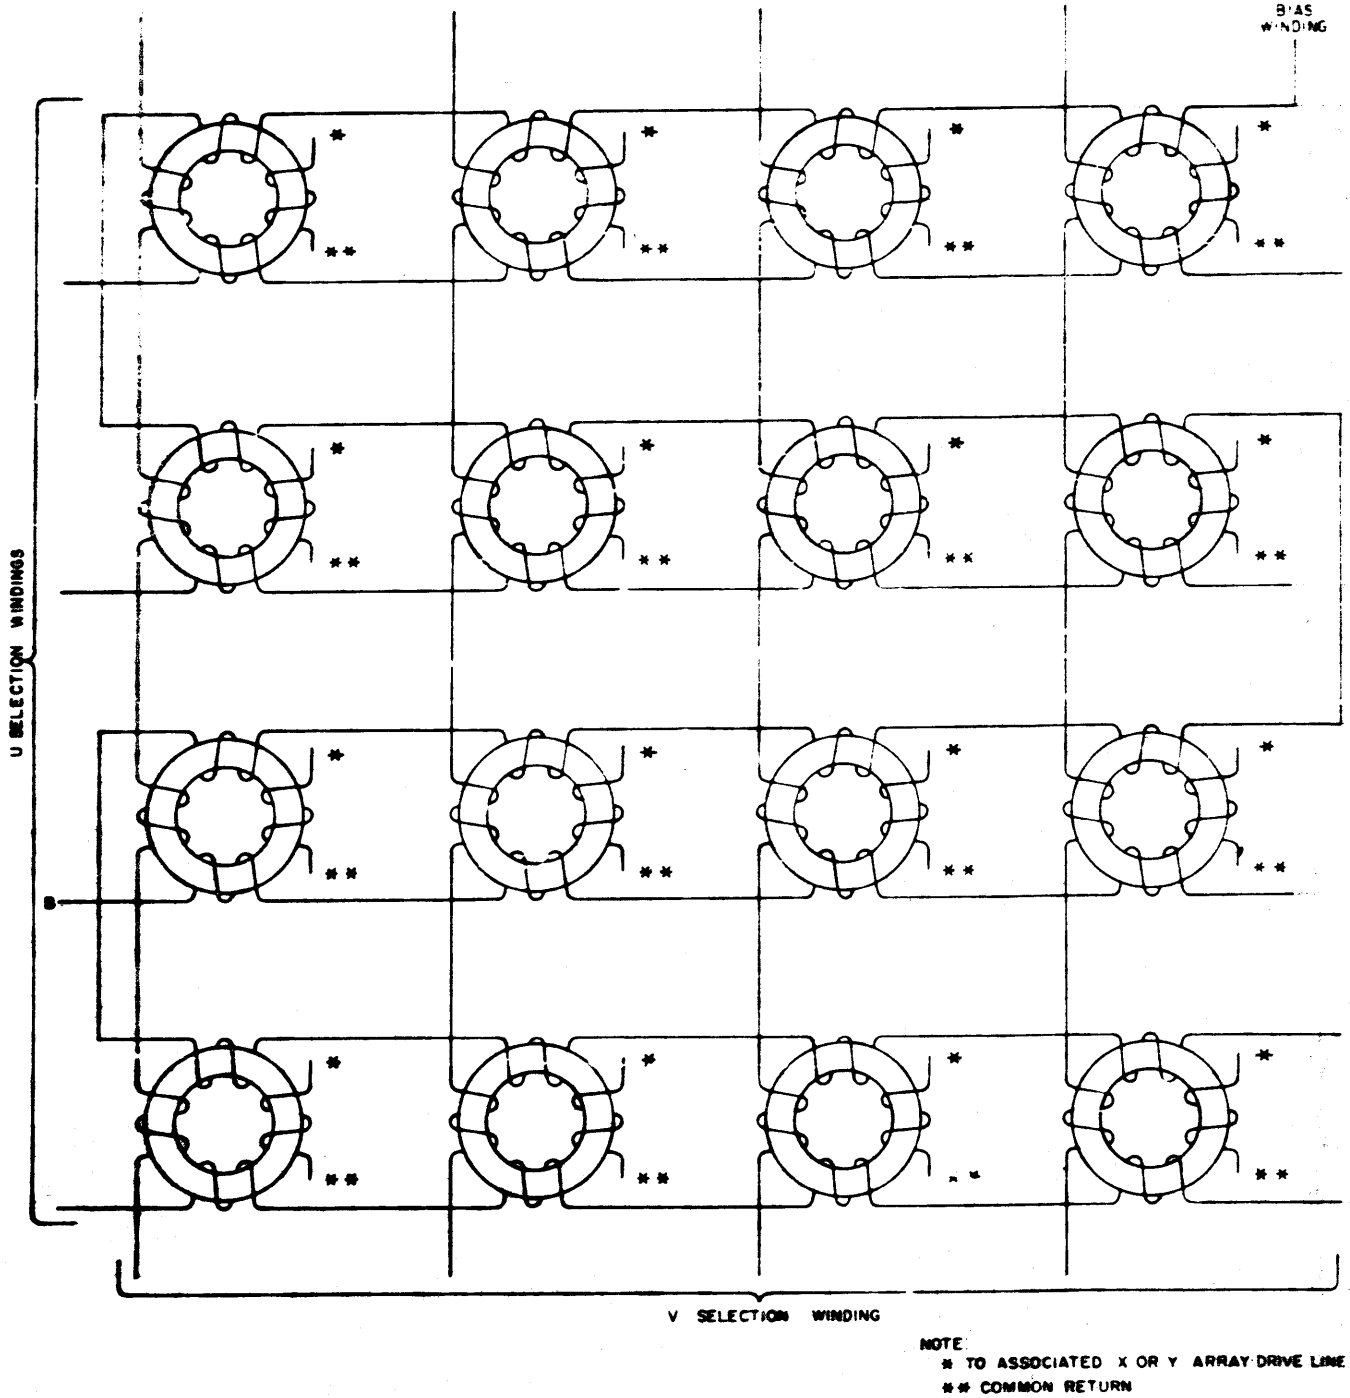

Figure 7–27. Tape Core Matrix, Schematic Diagram

energized at any one time; therefore, only the core at the intersection of the selected U and V line can be switched to produce read-write currents. Each tape core output winding is connected to one X or one Y winding; thus, only the selected array drive line will receive the required half-amplitude current required by the ferrite cores. The half-selected tape cores produce small noise pulses.

3.2.3 Selection Circuit Analysis

As noted in figure 7–25, the four groups of selection circuits are physically alike, and the only functional difference between them is the source of input signals and the destination of the output signals. The input signals are representative of information contained in the MAR. The distribution of these signals is shown in table 7–9. Since the circuit groups are alike, the following discussion will be based on only one of these; namely, the Xv selection circuit group which is shown in figure 7–28.

As noted in figure 7–28 the clear-memory-controls pulse (delayed TP 0 or start-memory pulse) resets the MAR flip-flops and the input amplifier (IA) deselect

TABLE 7–9. DISTRIBUTION OF MAR BITS

CIRCUIT GROUP	DIODE MATRIX DECODER INPUTS
Yv	R5-R3
Yu	R4-R7
Xv	R8-R11
Xu	R12-R15

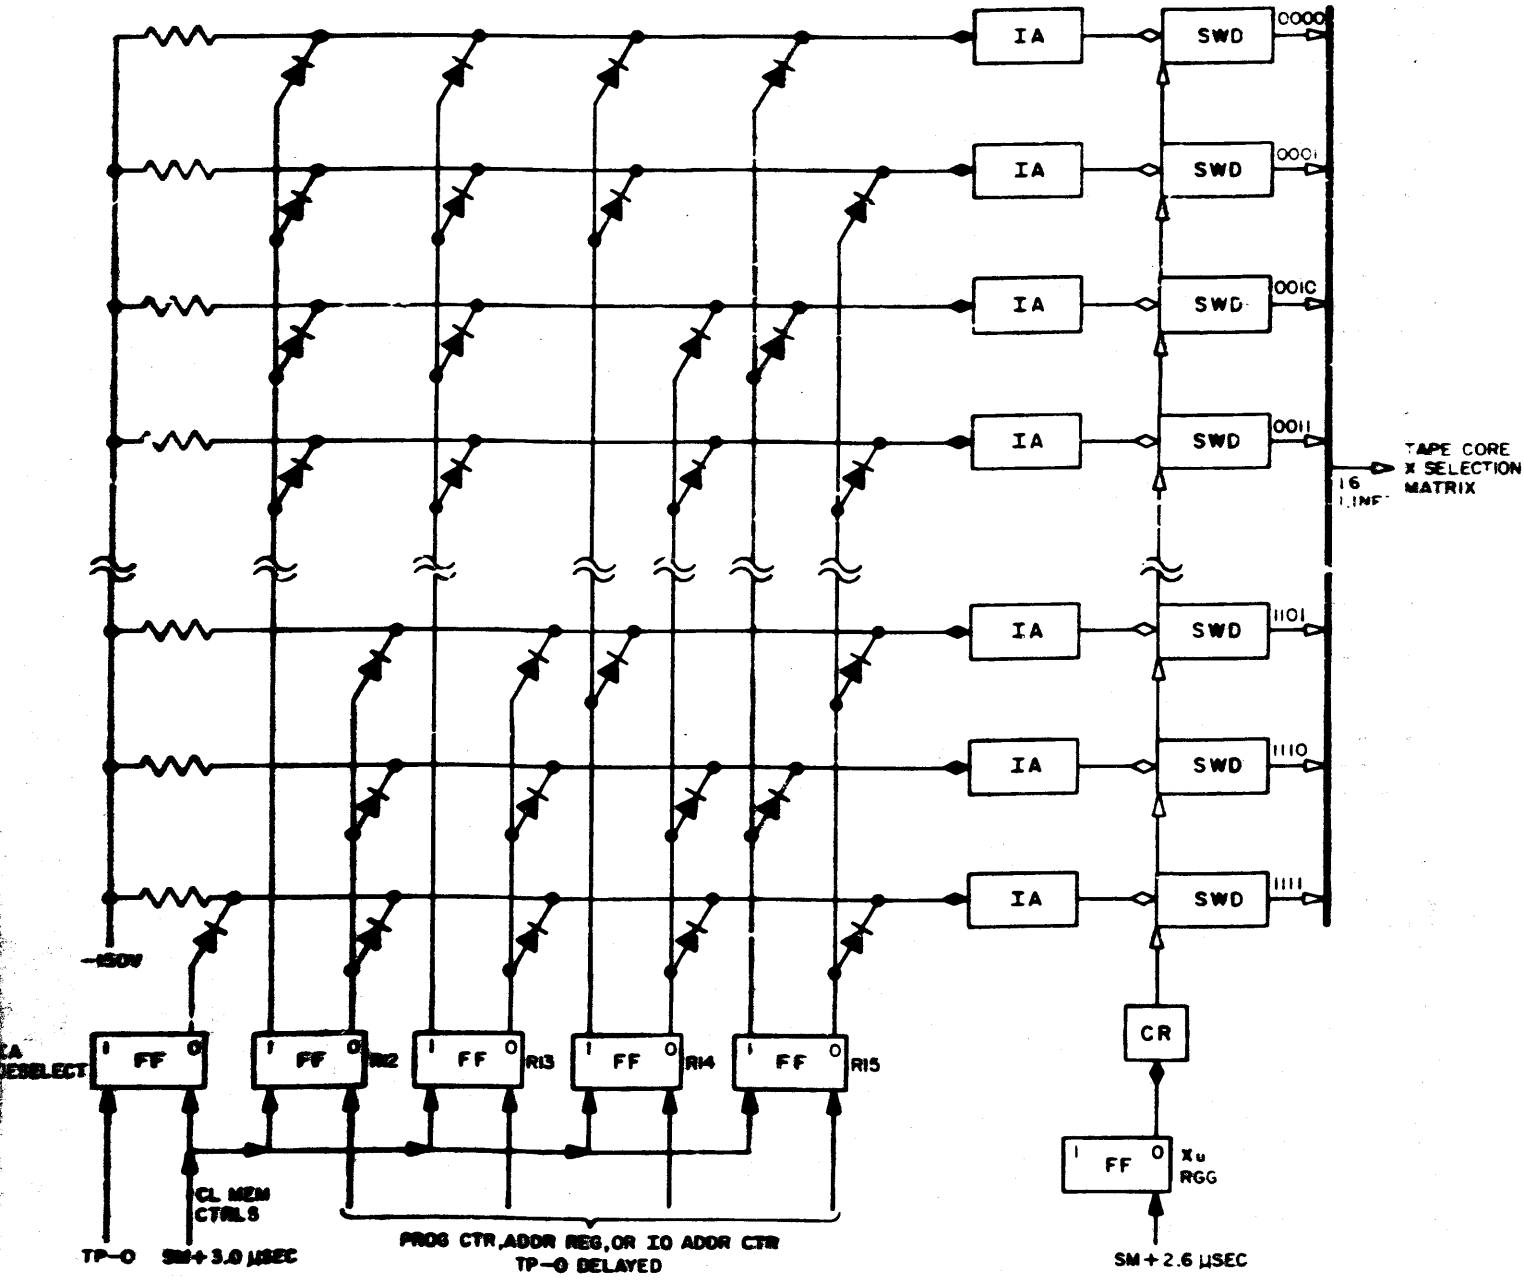

Figure 7–28. Xu Selection Circuits

flip-flop. At TP 0, the IA deselect flip-flop is set and the desired address is transferred to the MAR. The bit R12, R13, R14, and R15 flip-flop content is decoded by a 16-way-AND diode matrix decoder (DMD) to condition one of the 16 input amplifiers. The 4-digit binary numbers associated with each of the matrix output lines represent the contents of MAR bits R12, R13, R14, and R15 respectively. Only one of the diode matrix outputs can be selected (−30-volt level) at any one time; the remaining 15 outputs are at a +10-volt level. The outputs of the DMD are applied to the input amplifiers which invert, amplify, and change the level of the input signal. The nonselected input amplifier outputs are at a −240-volt level (+10-volt input), and the selected input amplifier is at −150-volt level (−30-volt input).

It should be noted that the −AND circuit whose output is designated by 1111 has an additional input from the IA deselect flip-flop. This additional input is required because the input amplifiers have capacitive coupling between stages and clamping circuits to hold their outputs at the nonselect level of −240 volts. When an input amplifier is selected, it can maintain its selected output level of −150 volts for only a short time, determined by the time constant of the capacitive coupling network. The output level will fall to the −240-volt level even though the input signal remains constant at a −30-volt level.

This means that when the same input amplifier is to be used during a number of consecutive cycles, it must be deselected and then selected again for each memory cycle. Deselection of 15 of the 16 input amplifiers is performed automatically when the MAR is cleared and then set to the new address. The 16th input amplifier is driven by the −AND circuit whose inputs are from the 0 sides of the bit R12, R13, R14, and R15 memory address register flip-flops. When each of these four flip-flops contains 1, designating the cleared condition, four −30-volt conditioning levels are applied to the −AND circuit. If a number of consecutive memory cycles are executed in which address bits R12, R13, R14, and R15 remain cleared (contain 1's), the MAR flip-flop input levels to the −AND circuit will remain constant. To provide for IA deselection in this case, the −AND circuit has five inputs instead of four. The fifth input comes from the 0 side of the IA deselect flip-flop. The IA deselect flip-flop is set at TP 0 and applies a −30-volt level to the −AND circuit, enabling it to act as any other −AND circuit in the matrix. When the memory address register is cleared, the IA deselect flip-flop is also cleared. Its output to the −AND circuit becomes +10V, effectively deselecting the −AND circuit and the input amplifier it drives. Thus, in all cases a −30-volt gate is actually applied to the selected input amplifier, thereby insuring proper action.

As shown in figure 7−28, the output of each input amplifier is applied to an associated SWD. A switch driver is a gating circuit which, when conditioned by a −150-volt level, will pass a current pulse generated by a current regulator (CR) to drive the selected U line of the X tape core. The current regulator is a power amplifier which is activated by a negative read gate signal to generate a current pulse in one of the 16 associated SWD's. Figure 7−28 shows that only one SWD can have an output at any one time. The active SWD is the one receiving both a −150-volt conditioning level from an input amplifier and a current pulse from a current regulator.

Since all four of the selection circuits groups are actuated by read gates at approximately the same time, current pulses will be applied to the selected U and V lines of each tape core matrix. At the point of coincidence, a tape core is selected in each matrix which, in turn, supplies the required coincident read-write current pulses to its associated X or Y array drive line. Table 7−10, which constitutes the tape core matrix selection gridwork, identifies the array drive line (in physical and ocquad notation) associated with the output winding of each tape core. Since the ocquad address notation provides symmetry in specifying the selected X and Y drive lines, this table represents both the X and Y tape core matrices. As noted in the table, the 2-digit octal numbers in the leftmost column specify the U input lines, and the 2-digit octal numbers in the topmost row specify the V input lines (fig. 7−24). During memory operation, the tape core at the intersection of the selected U and V lines will generate the required read-write current pulses. The individual tape core output windings are connected to the array X or Y drive lines as shown by the identification numbers contained in each box. The upper number in each box specifies the physical X or Y drive line, and the lower number, which is obtained by regrouping the combined binary equivalent of the selected U and V lines, designates the ocquad coding of that line.

As noted in table 7−7, consecutive ocquad addresses are contained in consecutive columns; thus one method to determine the selected U and V lines (switch drivers) for a given ocquad address consists of searching the table column by column for both the X and Y portions of the address. A second method consists of regrouping the binary equivalent of the ocquad address into V and U line notation for each tape core matrix directly (fig. 7−24). If the selected X or Y line is specified by physical line notation, then reference to table 7−10 will yield the associated ocquad line designation, which can then be converted according to one of the two methods noted above.

TABLE 7—IC. X AND Y TAPE CORE MATRIX OUTPUT

V (Octal)

U (Octal)	00	01	02	03	04	05	06	07	10	11	12	13	14	15	16	17
	0	4	8	12	16	20	24	28	32	36	40	44	48	52	56	60
00	000	040	100	140	200	240	300	340	400	440	500	540	600	640	700	740
	64	68	72	76	80	84	88	92	96	100	104	108	112	116	120	124
01	001	041	101	141	201	241	301	341	401	441	501	541	601	641	701	741
	128	132	136	140	144	148	152	156	160	164	168	172	176	180	184	188
02	002	042	102	142	202	242	302	342	402	442	502	542	602	642	702	742
	192	196	200	204	208	212	216	220	224	224	232	236	240	244	248	252
03	003	043	103	143	203	243	303	343	403	443	503	543	603	643	703	743
	1	5	9	13	17	21	25	29	33	37	41	45	49	53	57	61
04	010	050	110	150	210	250	310	350	410	450	510	550	610	650	710	750
	65	69	73	77	81	85	89	93	97	101	105	109	113	117	121	125
05	001	051	111	151	211	251	311	351	411	451	511	551	611	651	711	751
	129	133	137	141	145	149	153	157	161	165	169	173	177	181	185	189
06	012	052	112	152	212	252	312	352	412	452	512	552	612	652	712	752
	193	197	201	205	209	213	217	221	225	229	233	237	241	245	249	253
07	013	053	113	153	213	253	313	353	413	453	513	353	613	653	713	753
	2	6	10	14	18	22	26	30	34	38	42	46	50	54	58	62
10	020	060	120	160	220	260	320	360	420	460	520	560	620	660	720	760
	66	70	74	78	82	86	90	94	98	102	106	110	114	118	122	126
11	021	061	121	161	221	261	321	361	421	461	521	561	621	661	721	761
	130	134	138	142	146	150	154	158	162	166	170	174	178	182	186	190
12	022	062	122	162	222	262	322	362	422	462	522	562	622	662	722	762
	194	198	202	206	210	214	218	222	226	230	234	238	242	246	250	254
13	023	063	123	163	223	263	323	363	423	463	523	563	623	663	723	763
	3	7	11	15	19	23	27	31	35	39	43	47	51	55	59	63
14	030	070	130	170	230	270	330	370	430	470	530	570	630	670	730	770
	67	71	75	79	83	87	91	95	99	103	107	111	115	119	123	127
15	031	071	131	171	231	271	331	371	431	471	531	571	631	671	731	771
	131	135	139	143	147	151	155	159	163	167	171	175	179	183	187	191
16	032	072	132	172	232	272	332	372	432	472	532	572	632	672	732	772
	195	199	203	207	211	215	219	223	227	231	235	239	243	247	251	255
17	033	073	133	173	233	273	333	373	433	473	533	573	633	673	733	773

SECTION 4

SENSE SECTION

The sense section of the 256² memory element consists of 33 sense amplifiers, one for each memory plane. As already discussed in Section 1 of this chapter, the sense windings of the 16 subplanes of a memory plane are connected in four sense sections of four subplanes each.

This grouping of subplane sense windings was made in order to reduce the effects of half-select noise that is generated during the read portion of the memory cycle. In selecting any given address in a plane, 255 X and 255 Y cores are disturbed by half-select current pulses. If the resultant noise generated by these 510 disturbed cores was allowed to add during the read portion of the memory cycle, the resultant noise level could not be rejected by the sense amplifier.

To minimize the effects of noise signals, the subplane sense windings of each sense section have been connected in a series-parallel manner. Table 7–11 lists the subplanes in each group and also shows the series parallel grouping.

As noted in the memory plane shown in figure 7–29, if the core shown at the intersection of the X and Y lines is selected, then half-selected noise will be picked up on sense winding C7, A5, D4, A2, C12, and D16. It should be noted that the only noise generated in sense section B is that which is generated in subplane B6. Thus, even though 510 cores are half selected, sense section B will only pick up the noise from the 63X and 63Y half-selected cores in sense winding B6. Under these conditions, the amount of half-select noise added to or subtracted from the selected core output does not differ from the amount of noise generated in the existing 64² memories (refer to Ch 2).

The output of each sense section (see fig. 7–29) is connected to one of the four input sections of a sense amplifier; thus each sense amplifier consists of four separate input sections and a common section. The common section of the sense amplifier contains a gating circuit which is sensed by a sample pulse to determine the output of the selected core. During a readout cycle, if the selected core contained a one, the associated MBR flip-flop, which is cleared at the beginning of the cycle, will be set to the 1 state. If the core contained a 0 the associated MBR flip-flop will remain in the 0 state.

As noted in figure 7–29, the sample pulse does not sense all 33 sense amplifiers at the same time, but rather is delayed in steps of approximately 0.04 μsec to sample groups of sense amplifiers. This delay in sample time between the first group and subsequent groups of sense amplifiers (table 7–12) is required because of the delay inherent in the X and Y selection lines. For optimum results if a 1 is read out of a plane, it is necessary to sample the amplified output pulse when it is at its peak. The use of eight sample pulses compensates for the selection line delays and permits maximum memory reliability.

During a store cycle, the sample pulse generated by the memory clock must not be applied to the sense amplifiers. When such a cycle is executed, an inhibit sample pulse is supplied to the memory element at the beginning of the cycle. This pulse clears the sample gate generator, thereby deconditioning the sample gate so that the sample pulse will be inhibited during this cycle. The sample gate generator is set at the beginning of each memory cycle; thus a readout cycle is executed if an inhibit sample pulse is not supplied to the memory.

TABLE 7–11. SENSE WINDING CONNECTIONS

SENSE SECTION	SUBPLANES
A	5 & 10; 2 & 17
B	3 & 6; 11 & 14
C	0 & 15; 7 & 12
D	1 & 4; 13 & 16

Note: The sense windings of subplanes 5 and 10 are connected in one series loop, and the sense windings of subplanes 2 and 17 are connected in a separate series loop. The two series loops (5 and 10; and 2 and 17) are connected in parallel to the sense section output terminals. The other sense sections are connected in a similar manner.

TABLE 7–12. SAMPLE PULSE GROUPING

SAMPLE GROUP	PLANES
I	P-L3
II	L4-L7
III	L8-L11
IV	L12-L15
V	R5-R3
VI	R4-R7
VII	R8-R11
VIII	R12-R15

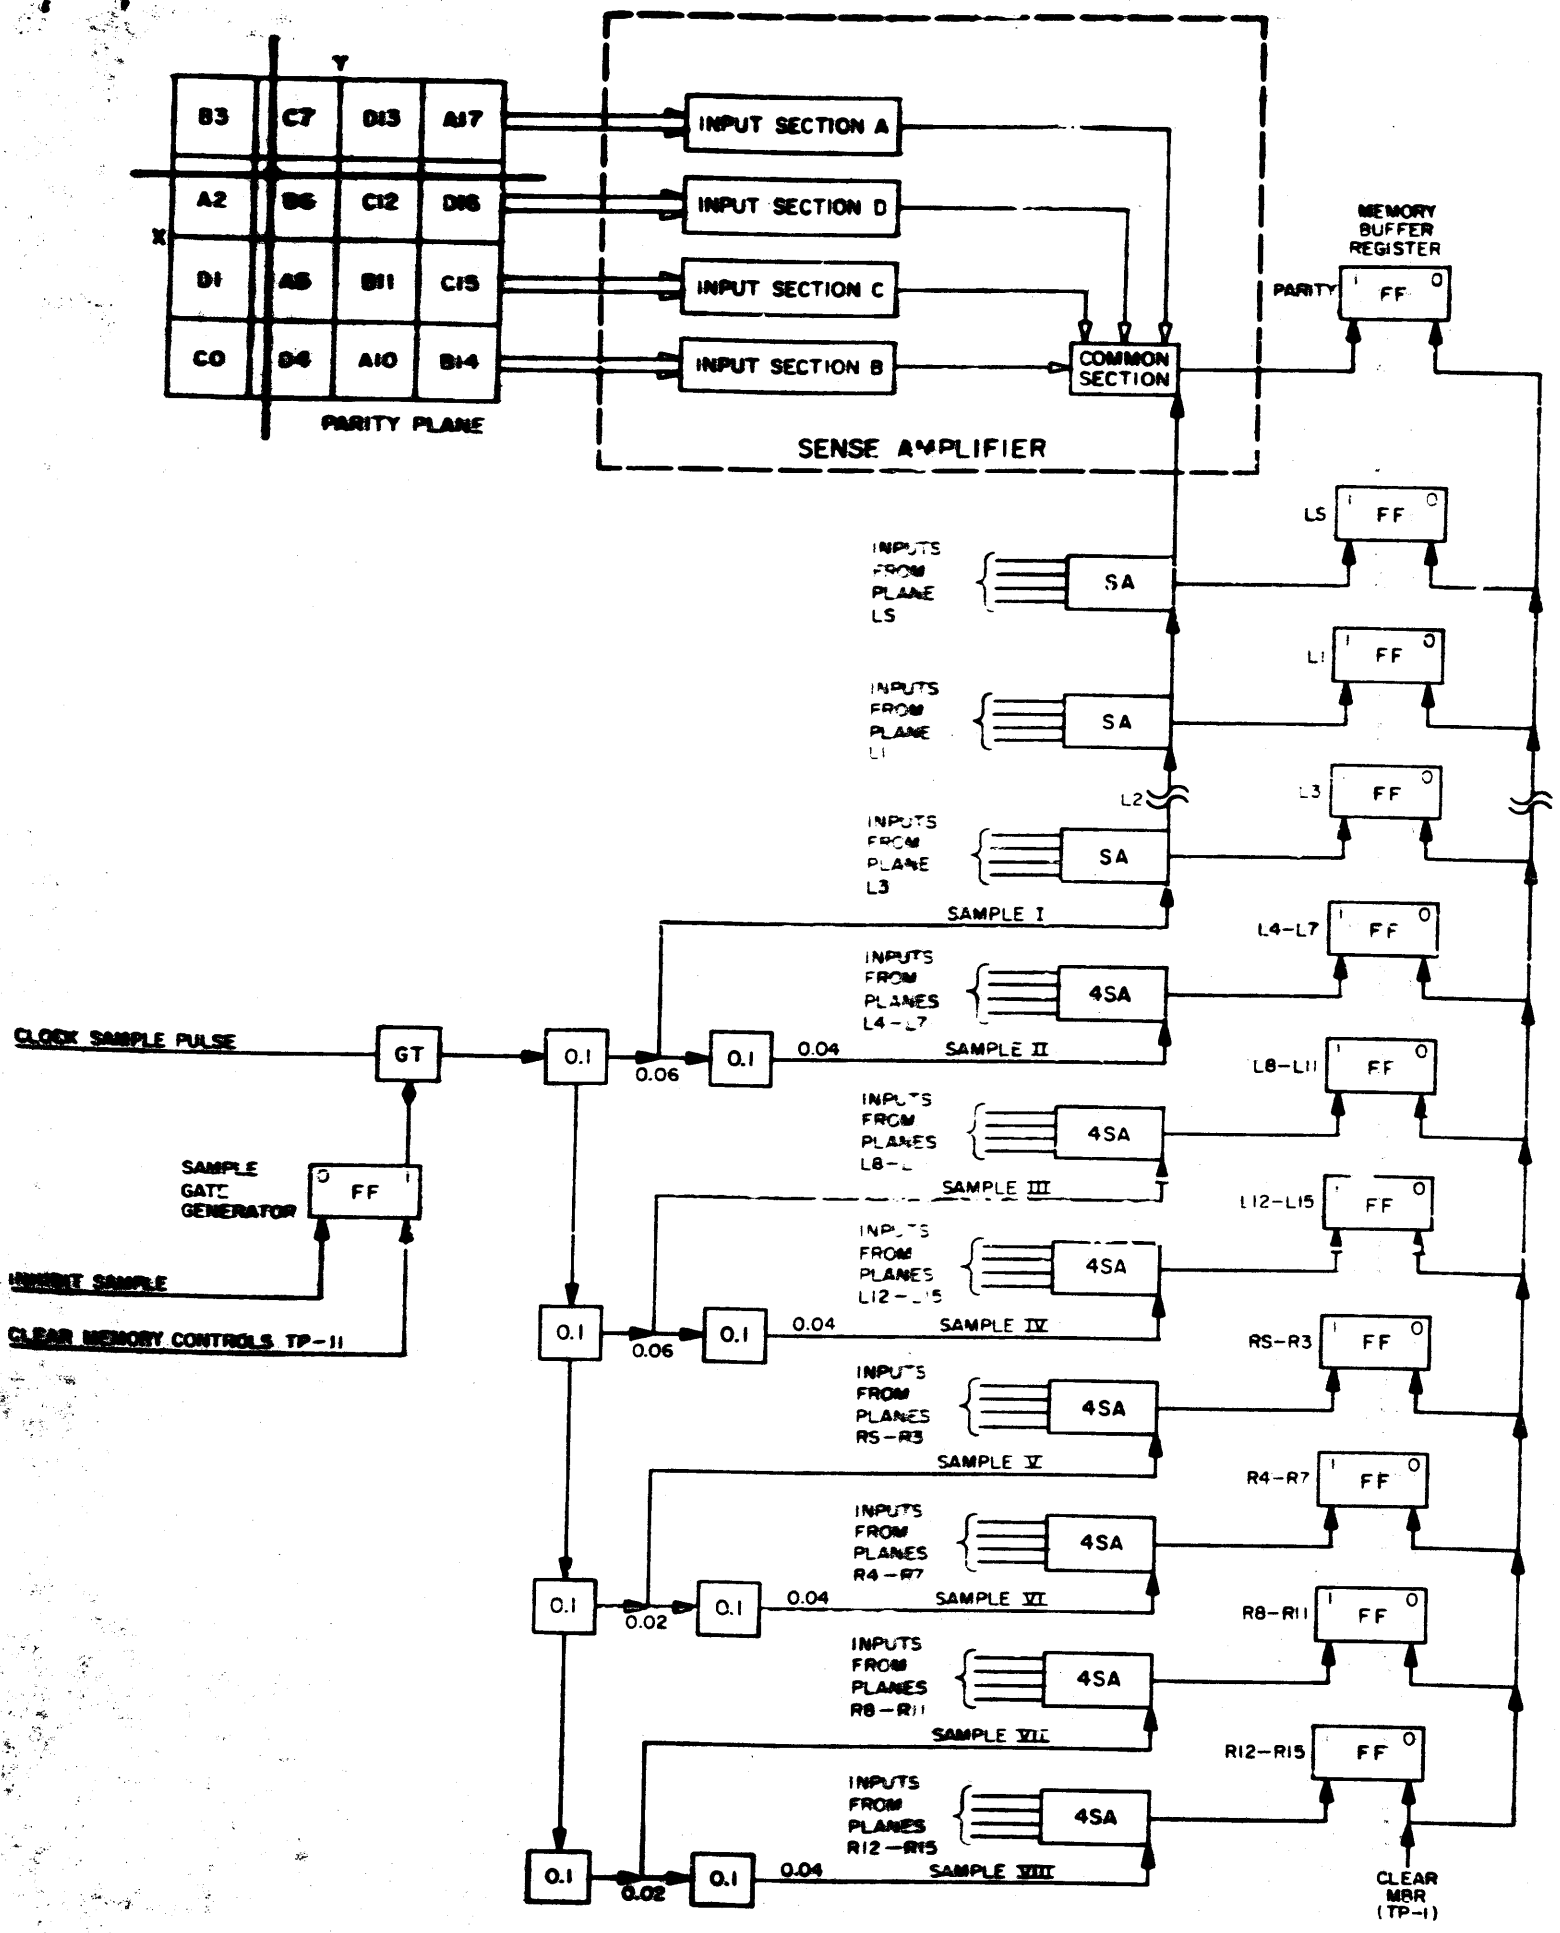

Figure 7-29. 256² Sense Amplifier Information Flow

SECTION 5
INHIBIT OPERATION

The DPD section of the 256² memory element contains 132 DPD's, which are used to supply inhibit current pulses to the memory planes of the array. Although the output of a DPD is a negative current pulse of approximately 400 ma, the winding geometry of the array ensures that this output will have the same effect as a half-select read current pulse. An inhibit current pulse is used during the write portion of a memory cycle to prevent the writing of a one by cancelling the effect of one of the write-current pulses applied to the selected core of the plane.

As discussed in Section 1 of this Chapter, the memory plane is divided into four inhibit regions, each region containing four subplanes. The individual subplane digit windings of an inhibit region are connected in series to form a common inhibit region digit winding. One DPD is used to supply an inhibit-current pulse to one inhibit region digit winding; thus, four DPD's are

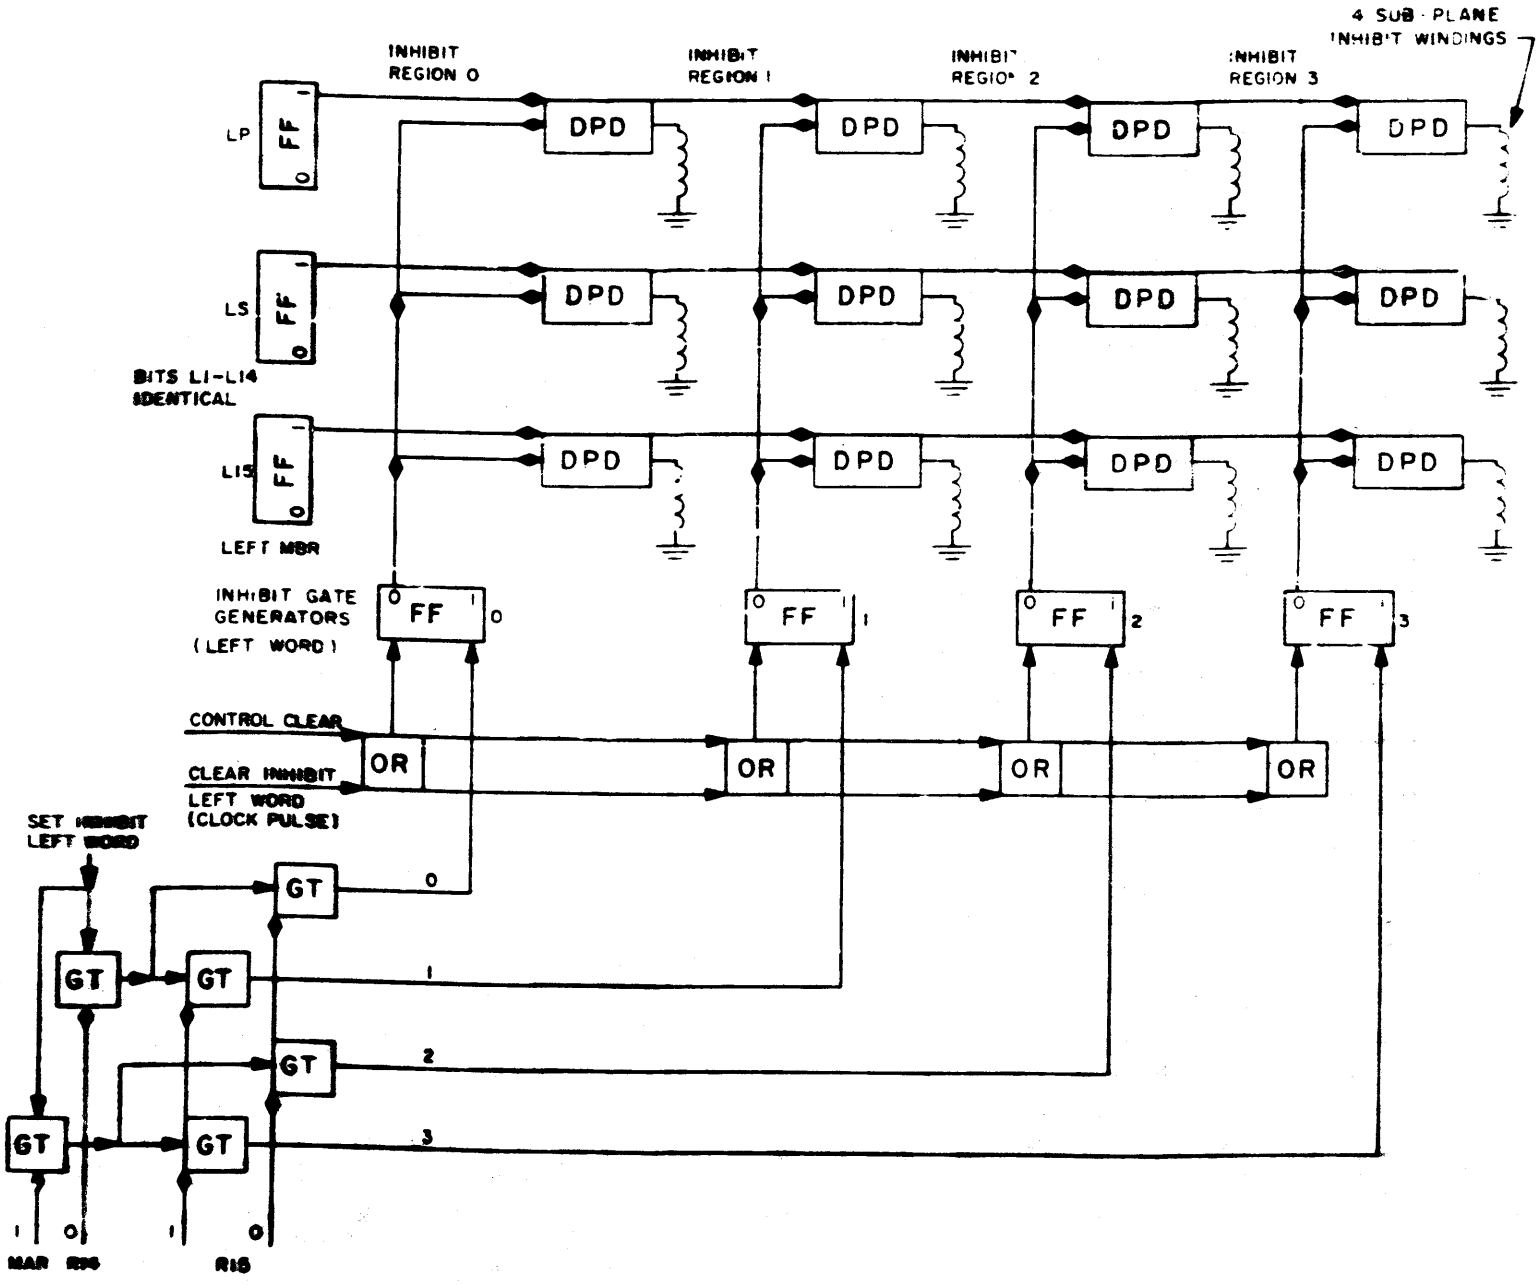

Figure 7–30. 256² Digit Plane Driver Control

required per memory plane. The four DPD's of each plane are controlled to supply an inhibit-current pulse, if required, to the inhibit region which contains the selected core.

Because of the method of addressing used in the 256^2 memory, consecutive memory addresses are contained in consecutive inhibit regions (see Sect. 2 of this Chapter). Since four inhibit regions are available, the two least significant bits of a memory address contained in MAR bits R14 and R15 are used to determine which inhibit region is selected. Table 7—13 lists the four component subplanes of each inhibit region as well as the binary code of MAR bits R14 and R15 that will designate the inhibit region selected.

Figure 7—30 is a block diagram of the DPD control circuits that are used to control the generation of inhibit-current pulses for the left half-word. Since the right half-word is controlled in exactly the same manner by a second set of control circuits, the following circuit analysis is applicable to both half-words. As noted in figure 7—30 the inhibit gate generators are in the cleared state at the beginning of each memory cycle. If the 256^2 memory is selected for operation, a start-memory pulse is applied to the memory clock to initiate memory operation. During the read portion of the memory cycle, information is transferred to the MBR from either the selected memory address (readout cycle) or from an external source (store cycle). In either case, during the write portion of the memory cycle, the information contained in the MBR is written into the selected address. If specific bits of the MBR contain 0's then an inhibit current must be generated in the selected inhibit region of the associated memory planes to overlap the write-current pulse, thereby preventing the writing of a 1 in the selected core. As noted in figure 7—20 a set-inhibit pulse is generated by the memory clock which senses the gate tubes controlled by MAR bits R14 and R15 to set the inhibit gate generator for the selected inhibit region. The selected inhibit gate generator develops a negative pulse (+10 to −30) on its 0-side output which is applied to the —AND input circuit of each associated DPD. If a particular MBR flip-flop contains a 0 its 1-side output will be at a −30-volt level and the negative inhibit gate will activate the associated DPD circuit. An inhibit-current pulse will thus be generated in the selected inhibit region of the associated plane.

The selected inhibit gate generator is cleared approximately 2.0 μsec after it was set; thus the inhibit-current pulse (2.0 μsec in duration) will properly overlap the write-current pulses.

TABLE 7—13. INHIBIT REGION SELECTION

INHIBIT REGION	SUBPLANES	MAR BITS R14 AND R15
0	0, 4, 10, 14	00
1	1, 5, 11, 15	01
2	2, 6, 12, 16	10
3	3, 7, 13, 17	11

SECTION 6
TIMING AND GATING

The various operations performed during a memory cycle must be executed according to a very stringent schedule. It is the function of the timing and gating section of the 256^2 memory to provide this rigid control in the form of control pulses and gate signals.

The timing and control section consists of a memory pulse distributor (MPD) and the entire complement of gate generators controlled by it. The MPD is actually a long delay line (approximately 6.0 μsec) which receives a start-memory pulse at the beginning of each 256^2 memory cycle. The delay line is tapped at several points, and the various delayed pulses are used to control the gating circuits. These gating circuits consist of flip-flops, whose output levels control the various memory circuit functions.

A simplified block diagram of the timing and gating circuits of the 256^2 memory is shown in figure 7–31, foldout. As noted in the figure, four different input signals are supplied from the Central Computer System to condition and activate these circuits. These inputs are identified as follows:

a. TP 0 pulse
b. MAR bit R14 and R15 output levels
c. Start-memory pulse (TP 0 delayed)
d. Inhibit sample pulse (TP 2)

The first two signals are always supplied to the 256^2 memory, regardless of its selection status. If the 256^2 memory is not selected, the TP 0 pulse initiates action to reset the MAR and IA deselect flip-flop. If the 256^2 memory is selected for operation, the third pulse, start memory, is also generated to control the timing of the read-, write- and inhibit-current pulses required during the execution of both the readout cycle (OTB or BI) and the store cycle (PT, OT or BO). During execution of the store cycle, the fifth pulse, inhibit sample, is also generated to inhibit the sampling of the sense amplifier, thereby erasing the content of the specified memory location during the read portion of the cycle. (Refer to Sect. 4 of this Chapter.)

Figure 7–32 is a timing chart showing the sequence of events for the memory-selected condition. As noted, the read and inhibit gate generators are set for this condition to control the generation of read, write, and inhibit-current pulses. Figure 7–31 shows that the V gate generators are set by the start-memory-delayed pulse (approximately TP 0 + 0.4 μsec) to supply a negative gate to the conditioned Xv and Yv current regulators, which in turn control the generation of the current pulses supplied to the specified V selection lines of the tape core matrices. The TP 0 pulse is also applied to MPD IV; however, in this sample the MPD IV gate tube is conditioned when sensed, so that no action results. Approximately 0.2 μsec after the start memory pulse was generated, the U read gate generators are set to supply current pulses to the specified U selection lines of the tape core matrices. As a result, each tape core matrix is now supplied with U and V current pulses and the

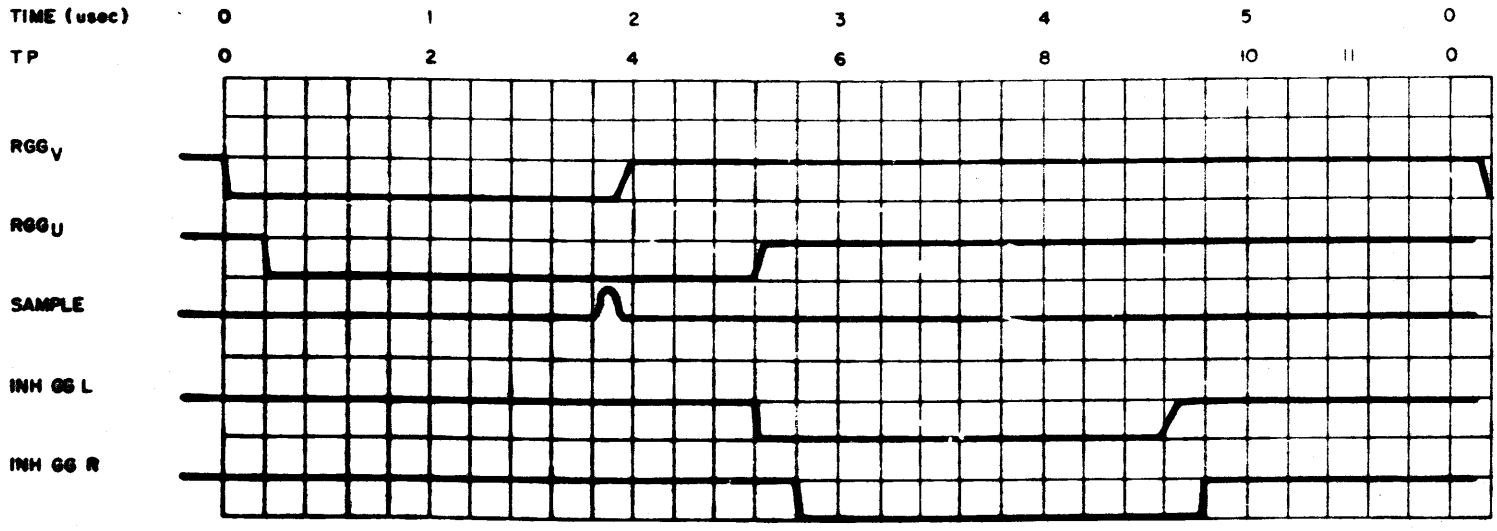

Figure 7–32. 256^2 Memory Selected, Timing Chart

selected tape core of each matrix will switch to provide read-current pulses to the associated X and Y drive lines of the array.

Approximately 1.4 μsec after the start-memory pulse was generated, MPD I develops a sample pulse which is applied to the sample gate generator tube. If a memory readout cycle is being executed, the gate tube will be conditioned so that the eight sample pulses will be applied to the 33 sense amplifiers. Staggered sample timing is required because of the delay characteristics of the X-Y drive lines.

Approximately 1.5 μsec (MPD I) after the start-memory pulse, the Xv and Yv read gate generators are cleared to terminate the read-current pulses.

Approximately 1.1 μsec later, MPD II generates a pulse to clear the Xu and Yu read gate generators to initiate the write portion of the cycle. The U and V read gate generators are set and cleared at separate times in order to speed up selection and also to reduce the noise generated in the sense winding. The action of clearing the read gate generators causes the two selected tape cores (one in each tape core matrix) to be switched back to the original state, thus generating the write-current pulses required by the associated array drive lines.

Approximately 2.2 μsec after the start memory pulse the selected left word inhibit gate generator is set by a MPD II pulse which senses the six gate tubes controlled by the MAR bit R14 and R15 flip-flops. About 0.2 μsec later the selected right word inhibit gate generator is set in exactly the same manner. The selected inhibit-gate pulses (negative pulses from 0 side output) sense the 33 digit plane drivers of the associated inhibit region; only those DPD's that are conditioned (at a −30-volt level) by the 1 side of the associated MBR flip-flop will pass the inhibit gate to activate the associated DPA, which in turn generates an inhibit-current pulse.

The left and right word inhibit gate generators are cleared at approximately 4.2 μsec and 4.4 μsec, respectively, after the start-memory pulse. Since the inhibit-current pulse of each plane must overlap the write-current pulse, staggered setting and clearing of the left and right word inhibit gates is required because of the delay characteristics of the X and Y drive lines.

During a memory cycle in which core memory 1 is selected, the MAR and 1A deselect flip-flop are cleared by a MPD II pulse which is developed approximately 3.0 μsec after the start-memory pulse was generated. Since the read gate generators are cleared prior to this time, the clearing of these controls does not affect the operation of the selected tape core in each of the tape core matrices. That is, the write portion of the memory cycle is not affected by this clearing action. The MAR is cleared at this time in order to provide more time for the selection of the next memory address during the next memory cycle. If core memory 1 is not selected during a memory cycle, then the MAR and the 1A deselect flip-flop are cleared by a delayed TP 0 pulse. As shown in figure 7-32, a raw TP 0 pulse is applied to MPD IV during every memory cycle. This pulse is delayed 1.5 μsec and applied to the MPD IV gate which is controlled by the Yu read gate generator. Since this gate generator is set only by a delayed start-memory pulse, the gate tube (MPD IV) is conditioned only when core memory 1 is not selected. If this latter condition exists, a core-memory-not-selected pulse is routed through MPD II to clear the MAR and the 1A deselect flip flop approximately 1.5 μsec after TP 0. The clear inhibit pulses which are generated by MPD III do not perform any useful function since none of the inhibit gate generators were set.

CHAPTER 6
PARITY CHECKING CIRCUITS

6.1 GENERAL

The basic definition of parity is equality with respect to an established quantity or standard. More applicably, parity may be defined as uniformity in either the oddness or evenness of number. The latter concept, in a slightly modified version, is used in the AN/FSQ-7 equipment for detecting errors incurred during the transfer or storage of binary information. Although parity is sometimes employed to verify the results of computations, it does not serve this particular function in the AN/FSQ-7 Central Computer System.

In an error detection scheme based on the more appropriate definition of parity, all words entering, leaving, or circulating within a system must have uniform parity; that is, every binary word either naturally exhibits or must be made to exhibit a parity in common with that of all other words in the system. Before elaborating on this principle, a further clarification of parity is necessary from the standpoint of its application in the Central Computer System. First, the established parity for all words in the Central Computer is odd. Second, the parity does not hinge on the absolute numerical value of a word of information; rather, it is based on the number of binary 1's contained in a word, be it an instruction word or a data word. Since the established parity is odd, each word in the Central Computer (which has parity) must have an odd number of 1's. Obviously, the parity of many words will naturally be odd. Those words having an even parity are given odd parity. This is accomplished, without changing the information contained in the normal 3-bit word, by prefixing to it a parity bit. When the natural parity of a word is even, the parity bit is made a 1; it is left at 0 when the original parity is odd. Having assigned an odd parity to all words, every word can then be checked for conformance with this characteristic. If, upon subsequent examination, a word is found to have even parity, it can rightfully be assumed that an error was generated either in storing or in transferring the word from a particular source to a specific destination. For purposes other than checking the accuracy with which information is transferred, the parity bit is meaningless and is discarded before a data word is entered into calculation or before an instruction word is decoded.

As a result, a practical application of parity requires circuits that are capable of (1) determining whether the number of binary 1's in a word of information is odd or even, (2) assigning a uniform parity to all words, (3) checking the parity of words to which parity has been assigned, and (4) generating a signal when a word with incorrect parity is detected. The first requirement is integral to both parity assigning and checking.

6.2 MEMORY PARITY CIRCUIT

The memory parity circuit performs a dual function: it assigns a uniform parity to all words originating within and without the Central Computer to which no parity has been previously assigned, and it checks the parity of words to which parity has been assigned. The latter function constitutes the test on the accuracy with which data is transmitted from the core memory array to the memory buffer register (MBR), and from several of the input sources to the MBR. Since the MBR is situated at the crossroads of all information paths in the Central Computer System, both parity assigning and checking take place at the MBR.

Of the several sources of words which are transferred through the MBR, only tapes and specific drum fields have a parity bit assigned. Therefore, words from these sources need only be checked for correct parity; all other input sources require that the parity circuits assign a parity bit. All words read out from memory, of course, have parity bits and require only parity checking.

A simplified logic diagram of the memory parity control circuit is shown in figure 7–34. Its principal elements are a series of gates and OR circuits that constitute a counting circuit, a parity write flip-flop, a parity check flip-flop, and a parity check control flip-flop.

The counting circuit senses the parity of the word contained in the MBR. The parity write flip-flop conditions a gate by means of which the parity bit flip-flop in the MBR may be set at 1 in the event that the count indicates a word of even parity. The parity check flip-flop controls a pair of gates through which the parity alarm circuits are activated. The parity check control flip-flop controls the operation of the parity circuits during BI cycles. Note that the input pulses to clear this flip-flop come from an input source not having parity.

It was mentioned earlier that the parity operations performed at the MBR are assign and check. An assign operation may be considered to consist of a parity count followed by the writing of a 1 or 0 in the parity bit flip-flop. The checking operation also begins with a count but is followed by a parity check. It will be noted that the counting procedure is common to both parity assign

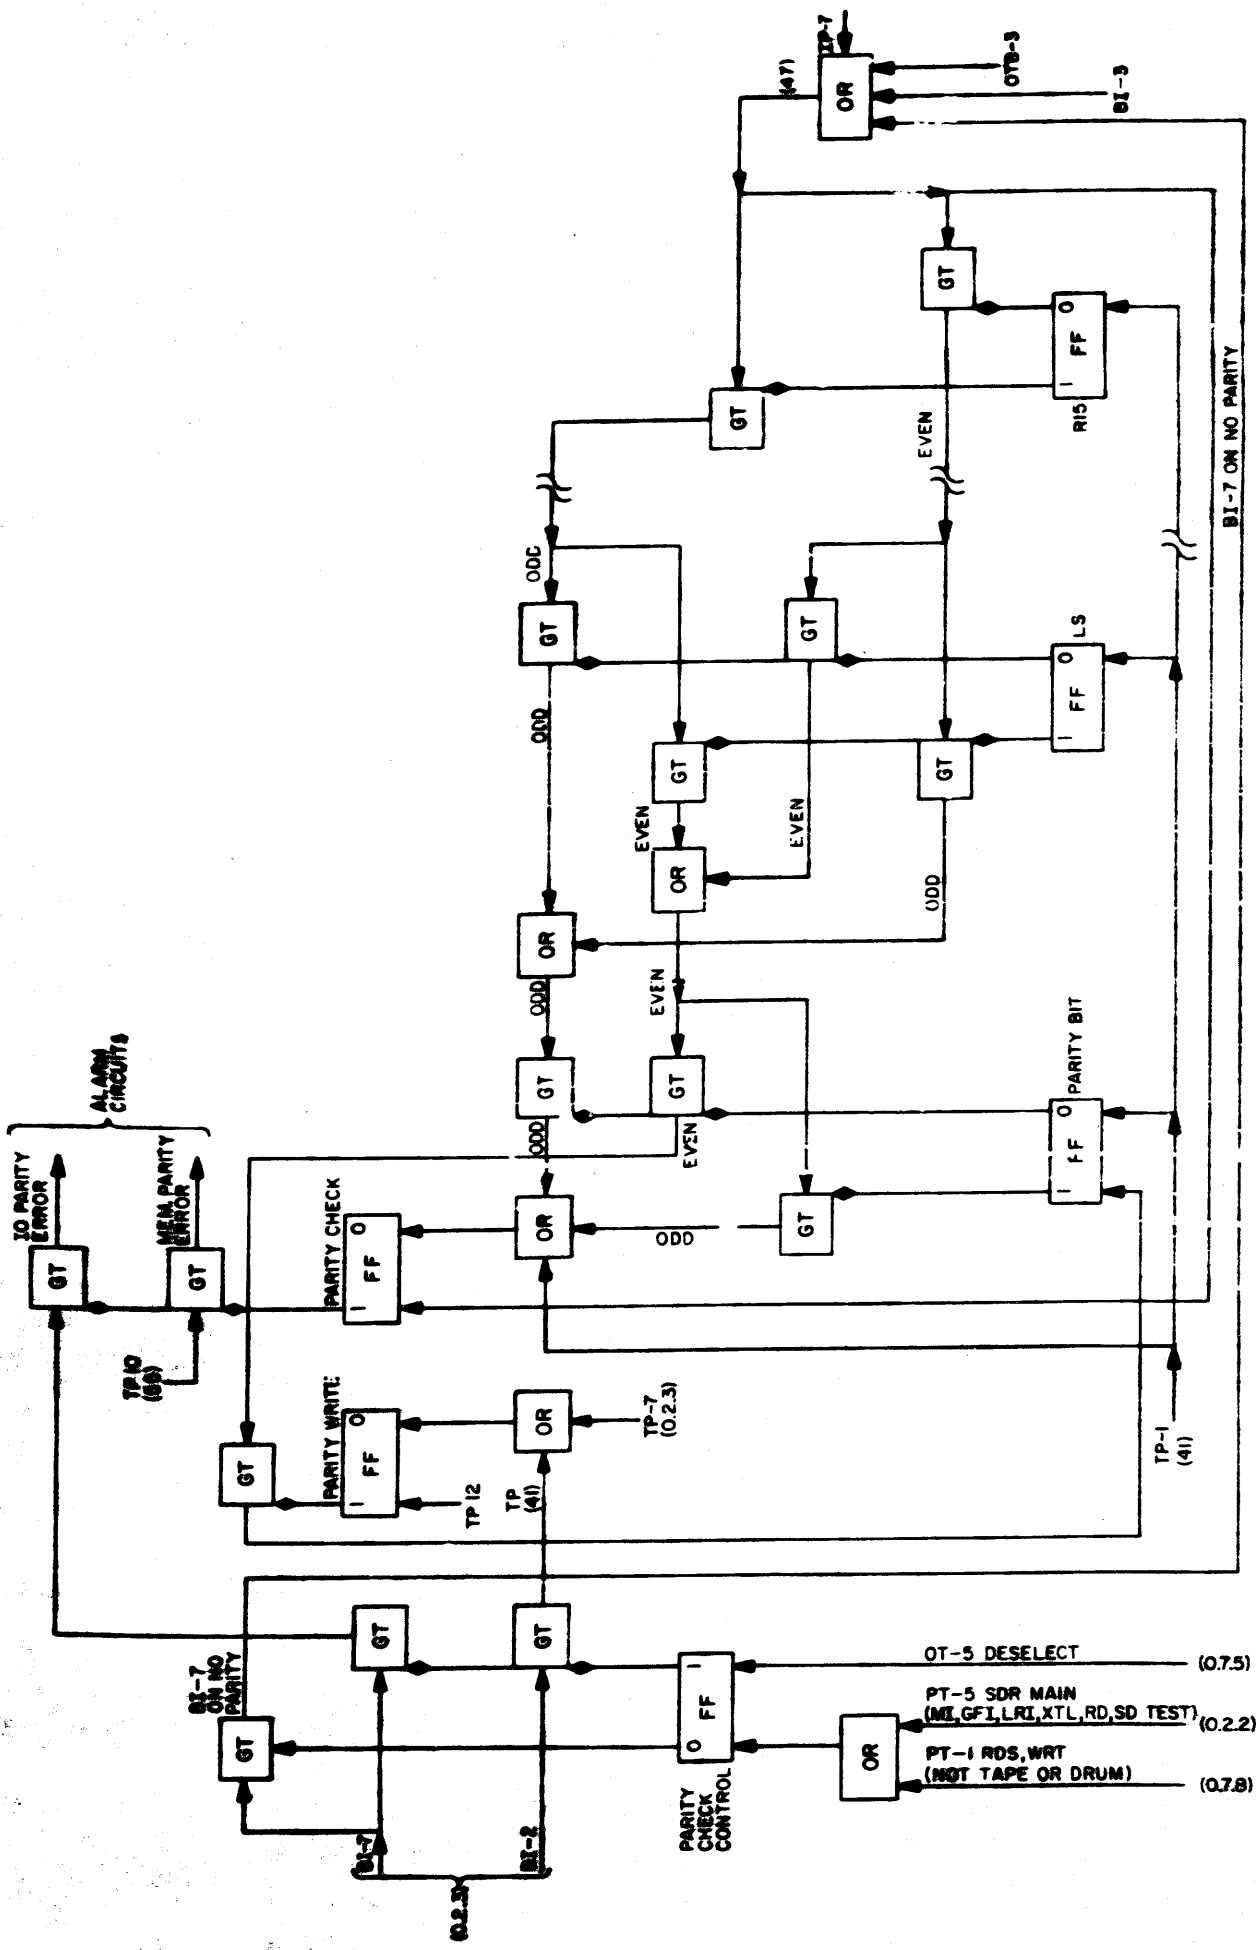

Figure 7–34. Parity Checking Circuits, Simplified Logic Diagram

and check. For convenience, the counting operation will be treated first and will be followed by the overall operation of the parity control circuits.

6.3 PARITY COUNTING

The counting operation neither counts nor totals the number of 1's in a binary word but only ascertains whether the total is odd or even. A count is begun with the application of a pulse to the two gates associated with the R15 flip-flop in the MBR (fig. 7–34). (Henceforth, a register flip-flop and its associated gates and OR circuits will be termed a stage.) One of the gates is conditioned by the 0 side of the flip-flop, the other by the 1 side. If R15 happens to be a 1, the counting pulse will pass through the gate conditioned by the 1 side and will emerge on the odd line. Conversely, if R15 is 0, the count pulse will pass through the gate conditioned by the 0 side and will exit on the even line. Although the remaining stages in the register are each complicated by the inclusion of two additional gates, the operation at each stage is identical with that described for R15 and its associated gates. The additional gates are necessary because the count pulse may arrive at any subsequent stage on either an odd or even line, whereas at the first stage only a single input line was encountered.

Upon leaving the first stage, the count pulse successively samples the gates at each of the remaining stages. The line on which it enters a particular stage is dependent upon the parity count at the preceding stage. The line on which it leaves that stage depends on the content of the associated flip-flop. Since the flip-flop may be set at either 1 or 0, and since the count pulse may arrive on either of two lines, the combination of these considerations gives rise to four possible conditions. The four conditions and the resultant output line (parity count) produced by each are listed below and are applicable to all but the first stage.

It will be noted that a pulse sampling the gates of a flip-flop containing a 0 leaves the stage on the same line it entered on. A count pulse arriving on either an odd line or an even line, upon sampling a 1, exits on a line opposite to that on which it entered a stage. In all cases, the action is equivalent to a progressive determination of the parity at each and every stage.

INPUT LINE (PARITY)	CONTENT OF MBR FF	OUTPUT LINE (PARITY)
Odd	0	Odd
Odd	1	Even
Even	0	Even
Even	1	Odd

In summary, a parity count pulse determines the overall parity of a word by traversing the entire MBR and sampling the bit at every stage to sense the parity of the word up to and including that stage. The final parity is obtained when the pulse samples and leaves the last stage in the MBR.

6.4 PARITY ASSIGNING AND CHECKING

The parity assigning and checking operations differ, depending on the type of word transfer:

a. All words being read out of core memory.
b. Words from the arithmetic element being read into core memory.
c. Words from IO devices being transferred into core memory with (1) parity or (2) no parity.

The discussion of the parity assigning and checking operations is divided accordingly.

6.4.1 Words Read Out of Core Memory

In this discussion, the parity check control flip-flop is of no consequence since its gates are strobed only with BI pulses.

The operation begins with a TP 1 pulse clearing the MBR and the parity check flip-flop and setting of the parity write flip-flop. The parity write flip-flop is set at TP 1 for every type of memory transfer operation in anticipation of the necessity to assign a parity bit. In this case, however, parity is not to be assigned and the flip-flop will be subsequently cleared. The start memory pulse is also issued at approximately TP 1 and, by TP 6, the selected word has been transferred from a core register to the MBR.

It is now necessary to check for correct parity and, accordingly, the parity count is begun with a command 47 generated by an IP 7. At the same time, the parity write flip-flop is cleared and the parity check flip-flop is set with a TP 7 pulse. The parity check flip-flop is set on the premise that a parity error exists in the MBR. It can only be cleared by a pulse on one of the odd lines leaving the gates on the parity bit stage, which indicates that the parity is correct. The parity check flip-flop is sensed at TP 10 (command 55), and if found set the pulse is channeled to the alarm circuits to set a memory parity alarm condition.

6.4.2 Words Read Into Memory from Arithmetic Element

Computer words are transferred from either the accumulators or the A registers into memory through the MBR. Since no parity bit exists in these words, it is necessary to assign parity. Such a transfer always takes place during an OTB or store cycle.

As before, the MBR and the parity check flip-flop are cleared and the parity write flip-flop is set with a TP 1. The parity count is begun by an OTB 3 pulse. If the output pulse from the LS stage is even, it is necessary

to set the parity bit flip-flop. Accordingly, the pulse on the even line strobes the gate on the 1 side of the parity write flip-flop and is fed back to set the parity bit flip-flop. If the output pulse from the LS stage was odd, the parity bit flip-flop is properly left in the cleared state.

At OTs 7, the parity count is started again, and the parity write flip-flop is cleared to prevent a parity write operation. The second parity count and check is performed even though parity was just assigned to check the parity circuits themselves. Sensing of the parity check flip-flop is done as explained in 6.4.1.

6.4.3 Words Read into Memory from IO Devices

The parity check control flip-flop is set at OT 5 of every SELECT instruction (deselect pulse). It remains set unless one of the IO devices which has no parity bit is to be read. It is cleared at PT 5 of the *SDR* (main) instruction for certain drum fields and at PT 1 of the *RDS* or *WRT* instruction if the tape or Drum Systems are not selected. The gates on the parity check control flip-flop channel BI pulses involved in the parity check.

Assume that an IO device having parity is being read (the parity check control flip-flop is set). The MBR and the parity check flip-flop are cleared, and the parity write flip-flop is set with a TP 1 pulse. A BI 2 finds the parity check control flip-flop clear and goes on to clear the parity write flip-flop. (This will prevent the writing of parity.) At BI 3, the parity count is begun. At BI 7, the parity check flip-flop is sensed. If found set, the pulse is channeled to the alarm circuits as an IO parity alarm.

Now assume the reading of an IO device which does not have parity. (The parity check control flip-flop is cleared.) As in all cases, the MBR's and the parity check flip-flop are cleared and the parity write flip-flop is set at TP 1. The BI 2 pulse cannot clear the parity write flip-flop as before so that parity is to be assigned at the completion of the count. The count is begun with a BI 3. A BI 7 pulse senses the gate on the clear side of the parity check control flip-flop and, passing, goes on to begin the second count. At TP 10, the parity check flip-flop is sensed. If an incorrect parity is detected, a memory parity error is generated.