

G. E. Werner
R. M. Whalen
N. F. Lockhart
R. C. Flaker

A 110-Nanosecond Ferrite Core Memory*

Abstract: The design of a large, very-high-speed ferrite memory is described. The memory has a capacity of 8192 words, 72 bits per word, and has a cycle time of 110 nanoseconds and an access time of 67 nanoseconds. The storage devices are miniature ferrite cores, 0.0075 by 0.0123 by 0.0029 inches, and are operated in a two-core-per-bit destructive read-out mode. A planar array geometry with cores resting on a single ground plane is used to control drive line parameters. Device switching speed and bit line recovery are treated as special problems. The design criteria and operational characteristics of the core, and the approach taken on the bit line recovery problem, are also presented.

Introduction

The memory described in this paper has been developed as part of a study to determine the speed limitations of ferrite technology. In an earlier phase of this program another memory was constructed which had a capacity of 8192 words of 72 bits each and a nominal cycle time of 375 nanoseconds.¹ It used cores with an inside diameter of 0.0075 inches and an outside diameter of 0.0123 inches in a word-organized, one-core-per-bit, partially switched mode.

One of the conclusions drawn from this earlier memory was that, with refinements, the cycle time could be reduced to 250 nsec. At this point, it was apparent that further reduction in cycle time was limited by the switching speed of the storage device and the recovery time of the bit-sense system. As a result, it was decided to undertake the development of a second memory and treat the switching speed of the storage device and the recovery time of the bit-sense system as the two primary problems. This second memory is also word-organized and again has a capacity of 8192 words of 72 bits each. The cycle time is designed to be 110 nanoseconds.

One of the simplest methods of increasing the switching speed of a ferrite core is to employ partial switching. By this means, the switching time can be reduced by switching less flux. The amount of flux switched can be controlled by adjusting the amplitude and duration of the drive currents. However, as the switching time and hence the amount of flux switched is reduced, the 1 to 0 signal ratio obtained from the core becomes progressively worse. In a typical core, the 1 to 0 ratio deteriorates to an unusable

level when the percent flux switched has dropped to about 40%. In order to further reduce switching time, it is necessary to employ some zero-cancellation technique such as a two-core-per-bit array design or to redesign the core itself to permit a greater amount of overdrive. The latter technique restores the 1 to 0 ratio by increasing the percent flux switched for a given switching time at the expense of higher drive fields. Of the two methods, the two-core-per-bit approach was chosen because it offered the greatest facility for adjusting switching speed to fit systems requirements, with a fixed core size and drive system design. Generally, the task of increasing drive fields requires a greater foreknowledge of the speed and power objectives than the two-core-per-bit approach, and entails a broader engineering effort. At the same time, a preliminary "look" at two-core-per-bit indicated a factor-of-three improvement in switching speed over a one-core-per-bit mode for the same core size. This was deemed adequate for this study.

The recovery problem was treated as a function of transmission delay, the amount of energy in the line which must be recovered, and the degree of impedance mismatch introduced by variations in array impedance, cable impedance and resistor tolerances. This approach influences many design choices including the mode of operation of the storage device pair, the bit density in the array, and the interconnections between the array and circuits.

* An earlier version of this paper appears in *IEEE Trans. Magnetics*, Mag-2, No. 3 (1966).

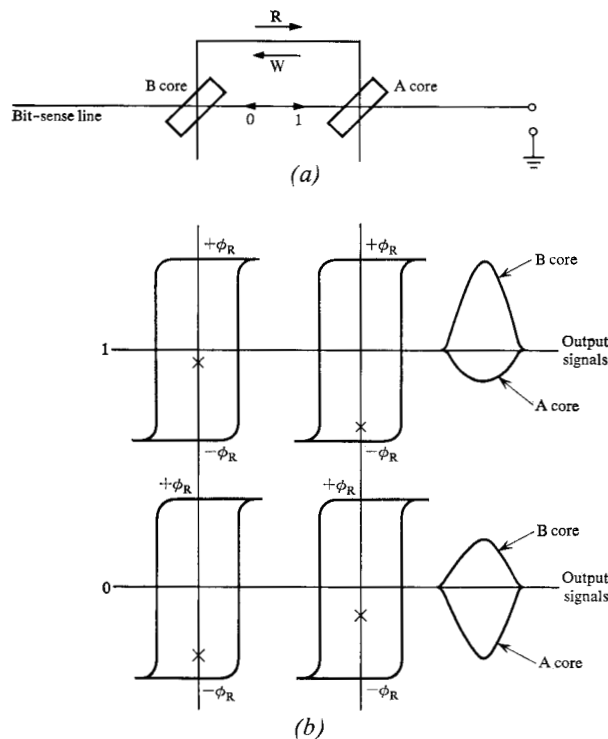


Figure 1 (a) Drive and bit-sense winding configuration; (b) flux states for storage of "one" and "zero."

Core operation and characteristics

There are a number of two-core-per-bit schemes from which to choose, some of which have been used before in memory systems.²⁻⁴ All of them have about equal potential for reducing switching time. Therefore, the choice rests on other advantages a particular scheme may have in relation to the application. The scheme chosen for this application is shown in Fig. 1a. Each core is threaded by two wires, a word line and a bit-sense line. The word line passes through core A, folds and returns through core B. The bit-sense line is a straight wire passing directly through cores A and B.

The bit current is bipolar; one polarity is used for storing a "one" and the opposite polarity is used for storing a "zero." The cores are oriented such that when storing a "one" the bit current pulse aids the write current pulse in core B and opposes it in core A. When storing a "zero" the polarity of the bit current is reversed so that it aids the write current pulse in core A and opposes it in core B. Figure 1b indicates the approximate flux states to which a pair of cores is driven for the storage of a "one" or a "zero." When reading, a current pulse is applied on the word line in the direction opposite to that of the write pulse, resetting both cores to the negative flux state. Cores A and B induce signals of opposite polarity in the bit-sense line thus providing a net positive voltage for a

"one" and a net negative voltage for a "zero." The amplitude of the signal is a function of the difference in flux states between the two cores.

One of the reasons for choosing this mode is that good output signals can be obtained with relatively low-amplitude bit currents. This minimizes the energy to be recovered from the bit-sense line after a write operation. In addition, a pair of cores appears on the bit-sense line as well as the word line. This minimizes bit-sense line impedance variations because the storage cell contribution to the line inductance is almost independent of the information state. Both of these advantages are important factors in dealing with bit-sense line recovery.

The core chosen for this application is a toroid, size 0.0075 by 0.0123 by 0.0029 in. This core has typical coincident-current squareness with a dc threshold of 175 mA as shown in Fig. 2. Figure 2 contains plots of flux switched as a function of drive current amplitude; these are commonly called S-curves. A dc S-curve and three narrow pulse S-curves are shown for a typical core. The notations 35 nsec, 30 nsec and 25 nsec on the narrow pulse curves denote the width of the drive pulse at the 10% points. From these curves it is possible to derive some of the operating parameters of the memory.

The choice of drive current amplitude and duration in a mode of operation where switching is time limited resolves itself into a compromise between the power dissipation in the drive system and the speed of the system. Because of speed considerations, pulse durations of 30 nsec were chosen for read and write. At this speed there is sufficient margin in drive system power dissipation to choose the amplitude of the write current such that switching is centered around the steepest portion of the S-curve. This maximizes the flux difference in a core pair for a given bit current. In this core the minimum point around which there exists a high degree of steepness is point "a" in Fig. 2. Thus a write current amplitude of 500 mA is required. When the write pulse is applied to a pair of reset cores in the presence of a 75-mA bit current pulse, one core is driven to point "b," and the other core is driven to point "c." The distance $\Delta\phi$ represents the net difference in flux, which can be directly correlated to sense signal amplitude.

In a practical memory involving thousands of cores the effect of core parameter distribution on sense signal output must be taken into account. This can be done by superimposing the S-curves of a statistical sample of the core production run. This results in a single plot in which S-curves are represented by bands. Since there is a predictable probability that cores at the extremes of the distribution will be paired in the memory, the S-curves are used to indicate the minimum $\Delta\phi$ or, more appropriately, the minimum sense signal amplitude for a given write and bit current and a given core distribution. By an iterative process of determining the minimum sense signal as a




Figure 2 Flux versus set currents.

function of core distribution, an S-curve band 20 mA wide at the nominal write current amplitude was specified. This yielded a minimum "one" signal of 30 mV.

The minimum read current must be sufficient to completely reset the core in the maximum flux state. In determining the maximum flux state, a variation in write pulse width from 30 to 35 nanoseconds was considered. When this is combined with an allowable 10% increase in the write and bit current amplitude the flux state at point "d" is reached. Since 30 nanoseconds is the minimum read pulse width, the current coordinate of point "e" represents the minimum read pulse amplitude. In practice the minimum read pulse amplitude is increased from the indicated 700 mA to 720 mA in order to take into account core distribution.

The maximum amplitude of the bit current pulses is determined by the disturb threshold of the core. In the bipolar bit system used here, cores are subjected to disturb pulses of both polarities, but since the disturb threshold of a partially switched core in the reset direction is much lower than the disturb threshold in the set direction, it determines the maximum amplitude of bit current.

There are two conditions of storage in a core pair where disturb sensitivity in the reset direction is critical. The first condition occurs when a "zero" polarity bit current disturbs a stored "one"; the second occurs when a "one" polarity bit current disturbs a stored "zero." Of the two, the first is by far the most important since any disturbance here reduces the amplitude of the "one" signal. In the second case it is only necessary to guard against disturbing a core pair to such an extent that a "zero" signal becomes positive.

Figure 3 shows the disturb characteristics of a typical core in the reset direction. This is a plot of the flux state of the core as a function of disturb amplitude. A family of curves is shown representing different flux states set




Figure 3 Flux versus reset-direction disturb current.

into the core. It can be seen from the curves that above 100 mA the core is greatly disturbed, since signal amplitudes become too unstable to operate in this region. In the region below 100 mA, the choice of bit current amplitude is a compromise between a high initial $\Delta\phi$ and high disturbance or lower $\Delta\phi$ and lower disturbance. An optimum arrangement was found when the "one" polarity current was set at a nominal amplitude of 75 mA, and the "zero" polarity current at a nominal amplitude of 35 mA. In this way a large $\Delta\phi$ was obtained for a stored "one" with a minimum amount of disturbance. Although stored "zeros" are disturbed by a 75 mA current they never become less negative than -6mV .

Array design

The choice of a storage device carries with it certain minimum requirements for drive amplitude and wave shape if acceptable operation is to be achieved. One of the principal objectives of array design is to assure that this condition is met once the device is placed in an array environment. A further objective is to preserve the largest signal-to-noise ratio that is possible for the storage cell in the presence of array-generated noise and sense line attenuation.

The planar array first introduced in a 375-nsec memory and now used in this system is in our opinion the best choice of geometry for achieving these objectives. It offers the following advantages:

- 1) The tight electrical coupling which can be achieved between array wires and ground plane results in the lowest line inductance and the smallest line-to-line mutual coupling of any of the geometries considered.
- 2) The ground plane provides a good physical reference for array wires and therefore assures uniformity of line characteristics.

3) At the repetition rates considered in this application, the internally generated heat which results from cycling the flux state of the device must be removed for satisfactory storage device operation. The temperature-controlled ground plane provides a simple heat sink. A thermal bond between ground plane and core is achieved by cementing them with a coating of an epoxy resin.

4) Mechanical simplicity is easier to achieve. The planar geometry provides easy accessibility to any part of the array without disassembly. This simplifies array repair—replacement of cores; correction of wire to ground and wire to wire shorts; etc. In addition to this the array line termination density in a planar array can be closely matched by the terminal density (input/output) of special array circuits. This minimizes cabling distances by minimizing array fan-out between array and circuits, thus reducing a source of noise, electrical load and delay.

In considering the design of the array the standard transmission line characteristics—impedance, delay, and attenuation—are normally considered. Generally speaking the design objective is to minimize all of these. High array line impedance is a source of noise within the array. The higher voltage excursions which occur when these lines are excited aggravate capacitive-coupled noise in addition to placing a burden on the voltage and power requirements of associated circuitry. Array delays add to access time. They also increase the difficulty of timing such things as the sense amplifier strobe. While each of these considerations is important in itself, their combined contribution to bit-sense line noise recovery was the major concern in designing this memory.

During the writing portion of a memory cycle the bit lines are energized with a current pulse. The energy stored on the line at this time is proportional to $I^2 Z_0$. If it is assumed that the bit line is properly terminated in its characteristic impedance, the stored energy will be fully dissipated after transmission down the length of the line. However, if termination is imperfect, a certain percentage of the initially stored energy will be reflected back to the sending end. If the sending end impedance does not match the line a second reflection occurs. This process continues. For each round trip on the bit line the reduction in energy will be inversely proportional to the reflection coefficients at the ends of the line. Recovery time is the time which elapses before the energy level (and therefore the signal level) has fallen to a value which is small compared with the subsequently developed sense signal. It is obvious that this time is dependent upon the initial energy on the bit line, the electrical length of the lines, the method of terminating these lines, and the degree of impedance match which is present between array, cabling and circuits.

The energy on the bit lines is minimized by keeping bit line characteristic impedance to a minimum and by con-


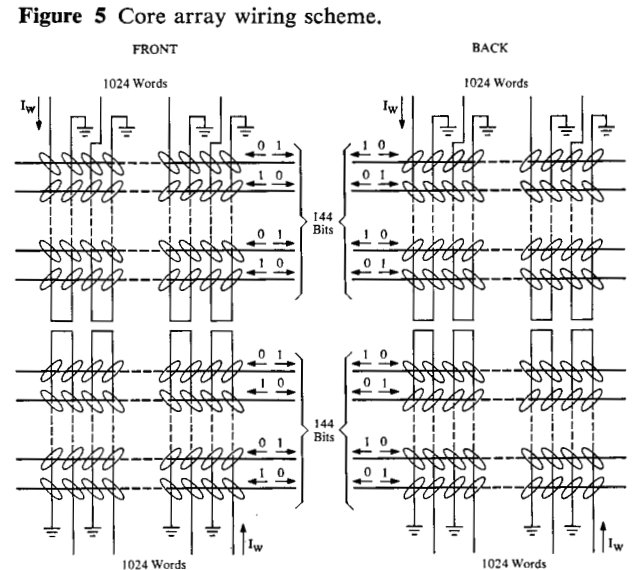


Figure 4 Core plane.



sidering the selection of a storage device and mode of operation in the light of minimum drive currents. Obviously, the best termination is termination at both ends of the line in Z_0 . However, good noise recovery depends upon the ability to choose a terminating resistor which matches the characteristics of the line. This brings up the question of line uniformity. The transmission line characteristics of the array lines will vary as a function of their geometry, line-to-line coupling and device contribution. Geometric variations have been kept to a minimum by providing a good physical reference for the lines—the solid ground plane. Line-to-line coupling has been minimized by providing tight electrical coupling between array wiring and ground plane with adequate separation between lines. Because of the importance of this characteristic in the bit-sense dimension, the bit-sense lines are placed closest to the ground plane. Our experience has indicated that a ratio of line separation to distance above the ground plane of 3 to 1 provides satisfactory isolation.




Figure 6 Ground plane.

The storage device is an important consideration when evaluating line uniformity. Its contribution to the transmission line characteristics of the array lines varies with frequency and with information (flux) state. One can readily recognize the frequency-dependent aspect of the ferrite core storage device by considering the small signal equivalent circuit, a parallel RL network,⁵ which has been previously reported for this device. With this equivalent circuit the core contribution is largely inductive at low frequencies and largely resistive at high frequencies. However, the roll-off in the inductive character of the device occurs at a frequency which is high relative to the frequency components normally encountered in this application. Therefore, signal attenuation due to high frequency mismatch has never been large—possibly of the order of 4%.

The small-signal permeability of a ferrite core can vary as much as 3 : 1 as a function of flux state. The importance of this to array line uniformity can be appreciated when one recognizes that the inductive contribution by the storage device to the bit-sense line in this application represents about 60% of the total line inductance. The solution to this potential problem has been the two-core-per-bit organization. In the particular scheme used, the two cores for each storage cell are adjacent cores on the bit-sense line. Regardless of the information stored, one core is always in a high flux state and the other in a low flux state. Their combined inductive contribution remains approximately constant. It has been estimated that the maximum variation in bit-sense line impedance would be no more than $\pm 1\%$ as a function of information stored. Considering the same information states in a single-core-per-bit organization the variation in Z_0 would be $\pm 10\%$. Insensitivity to information state has allowed high density (0.0075 in.) spacing of devices on the bit-sense lines. By maximizing device density, the inductive and capacitive contribution due to the wire are minimized for a given number of cores. This in turn has reduced recovery time by more than 33% by reducing transmission time by the same amount.

Array description

The array is composed of sixteen core planes. Each core plane (Fig. 4) contains 73,728 cores. These are assembled at one time to form a matrix of 256 words by 144 bits. Cores are spaced on 15-mil centers along a word line and on 7.5 mil centers along a bit line. The double density along the bit line was used in order to shorten the line as much as possible. The core density is over 8500 cores per square inch.

Eight core planes are mounted, side by side, in a 2×4 configuration, on each side of a single aluminum ground plane and the bit lines are interconnected to form two continuous planar structures, 1024×144 each. Thus each side of the array contains two sets of word lines of 1024 lines each and two sets of bit lines of 144 bits each. A wiring schematic of these two structures for each side of the ground plane is shown in Fig. 5.

The ground plane, shown in Fig. 6, contains milled grooves which permit the core plane frames to be counter-sunk sufficiently to allow the cores to rest on the surface of the ground plane. The depth is then finally adjusted by using shim-stock, so that the inner wire (bit-sense) is about 3 mils from the surface of the ground plane.

In order to provide a good thermal bond to the ground plane, which is used as a heat exchanger, the core areas are coated with an epoxy resin with a plasticizer added to prevent hardening. This material also reduces the drive line impedances and dampens magnetostrictive ringing. The temperature of the array is held constant by circulating a temperature-controlled liquid through a tube imbedded in the ground plane. The control apparatus is remotely located and the liquid is circulated to the ground plane through flexible tubes.

Mounting core planes on a single rigid ground plane provides a geometry which yields controllable and uniform drive-line parameters. Typical drive-line parameters for the array are given in Table 1.

Several causes for line parameter variation have been given. Our measurements indicated that those due to geometric variations account for about a 2Ω variation in the bit-sense line impedance.

Table 1 Typical drive line parameters for two-core-per-bit array.

| Parameter | Bit-sense lines Word lines | |
|------------------------------------|----------------------------|--------------|
| | | |
| Length | 21 in. | 6 in. |
| Resistance (20°C) | 3.8 Ω | 1.2 Ω |
| Characteristic impedance | 105 Ω | 110 Ω |
| Transmission delay | 8.7 nsec | 1.4 nsec |
| Inductance | 920 nH | 155 nH |
| Capacitance | 83 pF | 13 pF |
| Mutual inductance to adjacent wire | 60 nH | — |

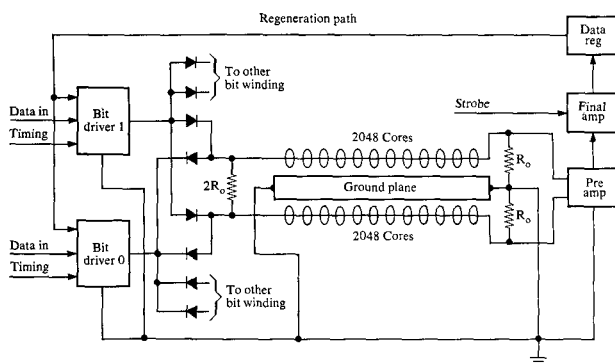


Figure 7 Bit-sense schematic.

Bit-sense system

A common bit-sense system is used in the memory to minimize the number of windings of the core plane. Figure 7 illustrates one bit position of this system. Corresponding bit wires on either side of the ground plane are connected through a diode network to a "one" polarity bit driver and a "zero" polarity bit driver and driven as a pair. Because of the relatively low bit current amplitudes and a desire to economize on bit driver circuits, the same pair of bit drivers is used to drive, in parallel, another pair of bit windings in the other half of the array. This results in a requirement for 144 "one" polarity bit drivers and 144 "zero" polarity bit drivers for an 8K system. A "one" polarity bit driver must supply a maximum of 330 mA and a "zero" polarity bit driver must supply a maximum of 154 mA to the bit windings. Consider a pair of lines for one bit position as illustrated in Fig. 7. When these lines are driven in parallel by the bit driver, the noise voltage developed at the far end (sense amplifier end) is largely common mode. This minimizes sense amplifier recovery by minimizing the difference mode noise voltage presented to it at this time. It also minimizes the base-emitter breakdown requirements on the pre-amplifier transistors. As shown, the far end of the bit-sense lines is terminated in R_0 , the characteristic impedance of the line. At the near end (bit drive end) the pair is bridged by $2R_0$. Both terminations are effective in damping difference mode noise conditions which exist between the lines; however, only the far end provides damping for common mode conditions. Obviously, R_0 terminations to ground at both ends of the lines would be the ideal termination for minimum recovery time. However, that method of termination would require twice as much bit current. The absence of a terminating resistor at the near end of the line which is capable of damping difference mode noise was also considered. This results in recovery times which are inconsistent with the cycle time objective of this memory. Terminating in the manner illustrated here has proven adequate for this system.

The sense signal is detected by a differential amplifier across the end of the bit-sense line pair. Because the line is terminated at both ends only half the signal appears at the sense amplifier, a minimum of 15 mV. Figure 8 shows the sense signal as observed at the input terminals of the sense amplifier.

The bit drivers and sense amplifiers are each connected to the array through a length of twisted pair cable and a "transition" block. A transition block consists of a sheet of 0.004 in. epoxy with printed signal lines on one side bonded to a copper strip as shown in Fig. 9. In order to connect the signal lines on the strip to the bit-sense lines in the array, the ends of the ground plane are expanded so that the copper strip fits into a milled slot where it can be soldered in place. The epoxy sheet is then formfitted around the end of the ground plane. When in place the signal lines on the epoxy sheet are aligned with the printed lands on the core plane frames. A short jumper strap connects the signal lines to the core frames. The epoxy sheet is designed so that when it is in place the impedance of the printed signal lines matches the impedance of the array lines. Twisted pairs connect the transition block to the bit and sense circuits.

Tolerances on cable impedances and terminating resistors inevitably lead to mismatch conditions; there is therefore a noise recovery period in the bit-sense system. Difference-mode noise and reflections arise out of impedance imbalance between the two halves of the bit-sense line. Resistors R_0 and $2R_0$ (Fig. 7) serve to damp this noise after each one-way transmission delay interval on the bit-sense line. Common-mode noise arises out of similar impedance mismatches between these two lines. As previously discussed, $2R_0$ has no value in damping common-mode noise; therefore any common-mode reflection from resistors R_0 must experience a round-trip delay before any additional damping can occur. In considering this characteristic, a worst-case mismatch between the lines which provided for 10 percent cable mismatches and 3 percent line termination mismatches was investigated. Difference-mode noise recovery to a 4 mV level occurred in 35 nsec (two round trips) after the bit current started down the line. The common-mode noise in this same time interval recovered to a 25 mV level. The common-mode rejection capabilities of the sense amplifier are more than adequate to handle this signal condition. Line recovery is illustrated in Fig. 10.

Word drive system

The word drive system used in this memory is of the same type as used previously in a 375-nanosecond memory.¹ The experience gained from that system led to a speeded up version of the drive system which has been adapted to this application.




Figure 8 Sense signal.

Figure 9 Transition block.





Figure 10 Bit-sense line recovery.

In this memory there are 4096 word lines in the array (2048 on each side of the ground plane), one of which must be selected and driven with a bipolar current pulse during each cycle. Each line contains 288 cores and is electrically short enough to permit the line to be grounded at one end and driven as an inductive load.

The fact that cores of some storage cell pairs are placed near opposite ends of a word line means that the read-out pulse starts switching one core before the other. The time difference is nominally proportional to the linear distance separating the cores on the word line. In order to avoid generating any spurious signals in the same direction as a "one" signal, because of this time difference, the core which is in the low flux state for a stored "one" is placed on the word line closest to the current source. For those core pairs which are significantly separated, this core begins to switch first, insuring that cross-wire capacitively coupled noise and bit-to bit inductively coupled noise which may be in the direction of a "one" signal are effectively cancelled.

The drive source must supply 10-nsec current transitions to nominal amplitudes of 720 mA and 500 mA for read and write respectively. In order to achieve this performance and still maintain some degree of economy with a matrix selection system, the word lines are divided into 16 groups of 256 lines each. Each group is driven by 256 linear transformers. A transformer is used for each word line primarily to permit common gating for read and write drives. Another advantage is that it reduces some of the noise voltages that would normally be coupled into the memory array from a matrix selection system.

The primaries of the 256 transformers are connected in a diode matrix and driven by 16 gates, 16 read drivers and 16 write drivers, as shown in Fig. 11. The read and write drivers are in turn driven by a 2×16 matrix consisting of 2 base drivers and 16 emitter drivers. The packaging is

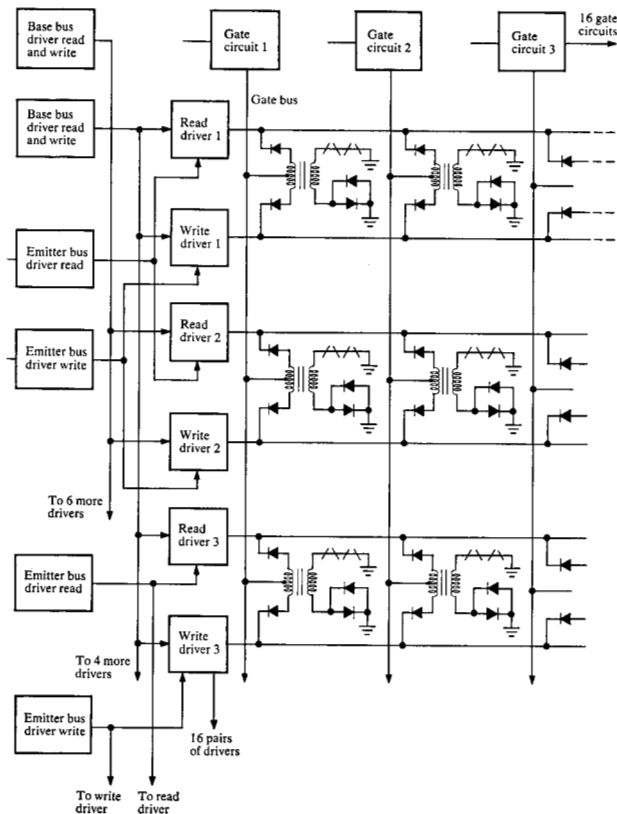


Figure 11 Word drive system.

arranged so that 8 pluggable circuit cards contain a drive system for 256 words. Circuit cards plug into circuit boards which butt against the array ground plane as shown in a cross-section view in Fig. 12. The interconnections between the transformer outputs and the word lines are made with a short printed-wire strap which connects the printed lands at the edge of the circuit board to the printed lands on the core plane frame. The ground planes in the circuit boards are exposed on the inner side of the board and connected to the ground plane by means of a solder-reflow technique.

In the quiescent state, the gate busses are maintained at a potential of -45 volts and the read and write busses at a slightly positive potential, so that the diodes are back-biased and nonconducting. When a word is selected, the appropriate gate bus is switched to ground potential, leaving the diodes in a slightly reverse-biased condition. When a pair of read and write busses are selected and successively driven negative, the diodes of the transformer located at the intersection of the active word drive and gate busses will be forward biased and conduct current.

A driver circuit is activated by driving the base of the output transistor to a clamped position at -45 volts with a base bus driver. When the emitter is driven negative by an




Figure 12 Cross section of array and word drive.

emitter bus driver, the output transistor conducts and supplies current to the primary of the transformer from its collector. This circuit is essentially a grounded-base driver where the amplitude of the word current is controlled by a resistor in the emitter of the output transistor. The back-to-back diode network inserted in series with the secondary of the output transformer reduces to a negligible level the dc shift that would result from the difference between the read and write currents.

Memory organization

The 8K memory unit is self-contained in that it has its own clock, address registers, decoding circuitry, data register, data-in and data-out controls, and byte control. These are well-known and will not be described. The computer need only provide address information and a select-memory pulse to initiate a cycle. The internal timing of the memory can be traced from the timing chart in Fig. 13. In order to achieve a 110-nanosecond cycle, an overlap technique is used. Because of pre-read delays, it is possible to begin a new cycle before the sense line is completely recovered.

Figure 14 is a photograph of a partially assembled 8K memory unit. The overall dimensions are $37.6 \times 22.0 \times 9.7$ in. The unit is shown at an early stage of assembly in order to better show the location of the various circuits. The array is mounted in a vertical plane with word boards positioned along the top and bottom edges. The bit and sense circuits are positioned close to the array at both ends. All the logic circuits are packaged at one end. In order to conserve space, circuit boards are mounted back to back and circuit cards are plugged into both sides.

Conclusions

This memory demonstrates the feasibility of using ferrites in the 10 megacycle performance range. In designing the memory the switching speed of the storage device was considered in relation to the problems of the whole system in order to achieve a proper design balance. Storage device




Figure 13 Timing chart—110 nanosecond cycle.

switching speed, however, still consumes over 60% of the cycle and it appears that further improvement in performance should be based on reducing the device switching speed. Although a two-core-per-bit mode was used in this memory to reduce switching time, this does not exclude the possibility of ultimately reducing core size in addition, for there was no indication that switching time could not be continually reduced by increasing drive fields.

Array and circuit packaging was found to have importance nearly equal to device switching speed in improving system performance. A planar array configuration mounted on a ground plane provided a stable method of controlling line parameters. This was offset somewhat by the need to make a transition, in the bit-sense system, from the array geometry to twisted pair in order to connect the bit and sense circuits to the array. Considerable improvement can be expected in bit line recovery time when circuit densities become high enough to permit these circuits to be integrated into the array geometry.

The experience gained with this system has increased confidence that practical ferrite memory systems can be designed for the 100-nsec range and has helped define the problems requiring solution for overcoming this speed barrier.




Figure 14 Memory unit—8K array.

Acknowledgments

The authors wish to acknowledge the valuable contributions made by the many engineers and technicians who worked directly on the project. In addition, grateful acknowledgment is given to the circuit design and mechanical packaging groups, particularly to J. W. Wyckoff who supervised the circuit design and to R. J. Betz who made important contributions to the design of the mechanical package.

References

1. G. E. Werner and R. M. Whalen, "A 375-Nanosecond Main Memory System Utilizing 7-Mil Cores," *AFIPS Conference Proceedings* 27, 985-993 (Fall, 1965).
2. C. J. Quartly, "A High Speed Ferrite Storage System," *Electronic Engineering* 31, 756 (December, 1959).
3. W. H. Rhodes, L. A. Russell, F. E. Sakalay, and R. M. Whalen, "A 0.7-Microsecond Ferrite Core Memory," *IBM Journal* 5, 174 (1961).
4. H. Amemiya, T. R. Mayhew, and R. C. Pryor, "A 10^5 -Bit High-Speed Ferrite Memory System—Design and Operation," *Proc. FJCC*, 123-164 (1964).
5. J. S. Eggenberger, "Distributed Parameter Aspects of Core Memory Wiring," *IBM Research Memorandum*, October 31, 1960.

Received April 18, 1966