

W. H. Rhodes
L. A. Russell
F. E. Sakalay
R. M. Whalen

A 0.7-Microsecond Ferrite Core Memory

Abstract: The design and performance of a newly developed magnetic core memory is described. A two-dimensional array organization and partial switching of toroidal cores were employed in the design of this low-power, high-speed memory. The memory features a unique combination of a current-steering diode matrix and a load-sharing magnetic switch for an economical and high-performance drive system. The operating memory has a storage capacity of 73,728 bits and executes instructions reliably up to a repetition rate of 1.47 mc. The discussion will include a description of the organization, the series-parallel delay line clock, the control of critical timing pulses, and the actual measured performance.

Introduction

Recent advances in the techniques of using ferrite cores in digital storage equipment have permitted a significant decrease in memory cycle time. One of the important advances was an operational technique for decreasing the time and energy associated with the switching of the storage array cores. This was accomplished, first, by reducing the amount of flux change and, second, by using a word-organized (two-dimensional) storage array. The first of these techniques, which has been termed *partial switching*,¹ requires the use of the word-organized type array, which in itself also offers substantial advantages in reducing memory cycle time.

The primary disadvantage of using a word-organized array is that there are N array address terminals to be driven, compared to as few as $2\sqrt{N}$ for three-dimensional arrays. This disadvantage has presented a difficult address-drive problem for arrays of large numbers of word addresses, N . For example, the memory described in this paper contains 1024 addresses or terminals which must be selectively driven, one at a time, with bipolar current pulses of approximately 0.5 amp. For economy it was necessary to incorporate high-current pulse-switching techniques in order to reduce the number of drive sources. The magnetic-core matrix switch previously used in lower performance memories² presents design difficulties for this performance range. Recent advances in the silicon semiconductor device art, however, have provided satisfactory high-current, high-speed diodes which are used in a current-steering drive matrix. These diodes, together with the load-sharing matrix switch³ are believed

to be key developments which permit successful memory design for the high speed and high bit capacity reported here.

The 0.7 μ sec core memory, which uses solid-state components throughout, was developed for a special application of an advanced computer system requiring a new high level of memory performance — an instruction execution frequency of 1.33 mc (0.75 μ sec period) and a storage capacity of 73,728 bits. (The operating model executed instructions at 1.47 mc.) Each of the 1024, 72-bit words was to be randomly addressable. A unique feature of this memory system is the variable machine cycle, which allows for various increments of time for data modification within a single instruction-execution cycle. Another feature is the multiplexed and asynchronous operation of the memory on a common data bus in conjunction with other similar and different memory units.

This paper discusses the particular mode of storage-core operation and the reasons for its choice; the array, drive, and sense designs; and finally some comments concerning the over-all logical organization and the evaluation-test results.

Storage core operation

The chosen mode of operation of the storage core (or more accurately the *storage cell*, in which two cores are used to store each bit of information) is most easily explained by first briefly reviewing the requirements that a word-organized array application places on the storage

cells. In this particular memory it is required that all 72 bits of a word be operated on in parallel. When a selected word is operated upon, it is understood that a READ word-current will be first applied and that this current will be followed by a WRITE word-current. No further applications of word-current will be made until this word is again selected. Each of the 72 storage cells of a particular word is simultaneously and individually operated on by BIT currents during the application of a WRITE word-current. Finally, during the application of a WRITE word-current to a particular word, each storage cell of all other words will receive either a ONE or ZERO bit drive. These BIT drives must not affect the information states of the storage cells of nonselected words.

As mentioned above, each storage cell is composed of two ferrite cores.* Figure 1a shows the manner in which these two cores are threaded by drive and sense windings and the respective current and voltage polarities. Figure 1b indicates the approximate flux states to which the pair of cores are driven for the storage of a ONE or ZERO information state. The READ current is conducted through a two-turn winding in order to halve the drive-current requirements. This drive is applied to both cores with a polarity that switches them to the negative flux state. On removal of the READ drive both cores relax to the states indicated by ϕ_R on the hysteresis loops. The cores may now be switched to the appropriate states corresponding to the information to be stored.

*D. J. Crawford, of the IBM Data Systems Product Development Laboratory, is the contributor of the exact two-core technique used.


The switching is accomplished by simultaneously applying the WRITE word-current and a BIT (A) current for storing a ONE, or a BIT (B) current for a ZERO. The amplitudes of the WRITE and BIT currents are adjusted so that a WRITE current alone causes a core to be switched from ϕ_R to ϕ_L and so that WRITE and BIT currents together switch a core from ϕ_R to ϕ_U . Thus, for the case of writing a ONE, BIT (A) is applied in coincidence with WRITE and the two cores attain the ϕ_U and ϕ_L states shown by the left pair of hysteresis loops in Fig. 1b. On the other hand, for the case of writing a ZERO, BIT (B) is applied with WRITE and the cores attain the flux states ϕ_L and ϕ_U shown by the right-hand pair of hysteresis loops. The current amplitude of the BIT drive is maintained sufficiently small so that repeated applications will cause no appreciable change of flux stability in the cores.

At some later time, the application of the READ drive causes the cores to be reswitched to the negative, saturated flux state and, in the process, the core that was previously set to the ϕ_U state induces a relatively large voltage in the sense winding, whereas the other core which was in the ϕ_L state induces a smaller voltage of opposite polarity. Therefore a net positive sense voltage is induced for a ONE state and a net negative voltage for a ZERO state being read from the storage cell.


The above two-core storage cell technique is similar in several respects to one described elsewhere.⁴ The chief difference is that in the technique previously described, a negative BIT drive is applied to one of the cores and simultaneously a positive BIT drive to the other. A

Figure 1 Storage cell for the 0.7-microsecond core memory.

(a) threading of cores by drive and sense windings.



(b) flux states for storage of ONE or ZERO information state.



factor which led to the choice of the unipolar BIT drive technique was that the storage cores used showed a lack of flux-state stability during repeated application of negative BIT current.

A significant reduction in storage-core switching time during the writing operation is obtained through the incorporation of the partial-switching technique.¹ This method had been utilized previously in one of the memory units used in the STRETCH computer. Briefly, the principle is that the drive ampere-turns to a ferrite core may exceed the static or dc threshold for irreversible flux change value without causing an appreciable permanent change in flux state if the duration of the drive is short. This permissible increase in drive may be large if it occurs only once between successive reset or READ drives. Restriction of the WRITE current duration in the 0.7 μ sec memory to 0.25 μ sec permits a nominal amplitude of 450 ma, which is approximately 150 ma in excess of the static threshold value. A final factor contributing to the speed of the storage cell is that the READ drive may be arbitrarily large and thereby may accomplish the change from ϕ_U or ϕ_L to ϕ_R in as short a time as is desired. This hypothesis does not consider practical limitations in drive current capability.


The memory design achieved a READ drive of 600 ma (1.2 ampere-turns) with a rise time of about 0.08 μ sec, which results in a switching interval of 0.15 μ sec for the storage cell. The duration of the READ drive is set to 0.25 μ sec in order to assure stability of the ϕ_R state for the various worst-case operating conditions.


The BIT drives are set to 275 ma, with a duration of 0.20 μ sec. The BIT drive duration is made smaller than that of the WRITE drive in order that its occurrence be within the WRITE current interval, rather than vice versa. This is permissible since the BIT drive is to contribute to, rather than inhibit, the core switching. The advantage is the termination of the BIT drive earlier in the memory cycle, with the result that sense winding and amplifier circuits begin earlier recovery from the large noise voltages induced by BIT drive transitions. In addition, the propagation of information in the regeneration loop may overlap the start of the WRITE sequence.

The amount of flux change occurring in the storage cell with the above drive-current amplitudes may be described in terms of the total difference in remanent flux states for the major hysteresis loop: this quantity is $2\phi_R$. The change from ϕ_R to ϕ_U is approximately 30% of $2\phi_R$ and the change from ϕ_R to ϕ_L is approximately 10% of $2\phi_R$. This flux difference with the 1.2 ampere-turn READ drive resulted in a sense winding signal of ± 75 mv/turn and a READ drive voltage per cell of 150 mv/turn.

From the standpoint of sense amplifier design, the constant total-flux-change technique was chosen because of the equal magnitude and opposite polarity of the sense signal, which reduces the difficulty of signal discrimination. A second and more important reason for choosing this technique is that the back voltage presented by the word winding to the drive circuits is relatively constant for all words, independent of the information stored in each. This property greatly simplified the design of a

Figure 2 64-input, 32-output load-sharing matrix switch.





satisfactory load-sharing magnetic matrix switch used in the drive circuit.

A satisfactory solution to the problem of driving 1024-word address terminals was a key factor in designing the memory with a cycle time of 0.7 μ sec. The solution was to use the recently available high-conductance, high-speed silicon diodes in a current-steering matrix, together with a load-sharing matrix switch.³ A photograph of a load-sharing matrix switch is shown in Fig. 2. The network in Fig. 3 shows how the two types of matrix switches are combined to satisfy the word-drive requirements of the storage array. The drawing is simplified in order to avoid complications which are not necessary for understanding the design. Not shown in the drawing are transformers 3 through 31 of the load-sharing switch and the associated secondary circuits including the READ and WRITE buses;

Each of the rectangular boxes in the diode matrix represents the storage cores and windings of a particular word or address. Two diodes are required for each word, one connected to the READ and the other to the WRITE winding. The opposite ends of both these windings are commoned and connected to a gate bus which is maintained at a positive voltage except when a connected word is to be selected. When a word is selected the appropriate gate bus is switched to ground potential, which reduces the reverse-bias voltage from the diodes to a minimum value. The anode terminals of the READ winding diodes are connected in groups of 32 to the READ buses, which in turn are connected individually to one end of a three-terminal secondary winding. Similarly, the WRITE diode anodes connect through WRITE buses which join these diodes to




Figure 4 Oscilloscope displays of (top to bottom) gate bus voltage, READ and WRITE word-winding currents, and BIT-winding current. Horizontal scale: 200 nsec/div. Vertical scale: Word gate, 30 v/div. current pulses, 300 ma/div.

a terminal at the opposite end of the secondary windings. The separate diodes, resistors and negative-bias voltages associated with each secondary winding are provided in order to avoid unnecessary application of large negative voltage to a WRITE drive bus during a READ operation, and vice versa. The secondary circuits also decouple the secondary windings from the drive buses during inactive periods in order to permit small relaxation flux changes to occur in the matrix-switch cores without inducing load currents. The drive to the load-sharing switch, as described by Constantine,³ causes a positive voltage pulse to appear first on a selected READ bus and later on the corresponding WRITE bus. This voltage will forward-bias the diodes of the word located at the intersection of the active drive and gate buses. The forward-biasing of these diodes causes current pulses to be conducted through the corresponding word windings of the addressed word. These currents will be proportional in amplitude to the primary winding currents of the load-sharing switch. The proportionality constant is individually controlled for READ and WRITE drives by the choice of the number of turns of each section, by the use of three-terminal secondary winding and the values of the shunt resistors in the secondary circuit.

The transistors shown in the primary circuits of the load-sharing switch are high-current, high-speed, medium power transistors developed and manufactured by IBM. Although not shown, these transistors are connected to an input circuit through a nonlinear feedback network, which maintains a low but nonsaturating emitter-collector voltage during conduction. This circuit is essentially the same as described elsewhere.⁵ A primary-winding drive

of 200 ma is required, which is achieved with a rise-time of 0.04 μ sec.

The word gate drive is provided by two parallel-connected transistor circuits of the same type as is used for the load-sharing matrix switch. The collectors of the high-current transistors are connected through small, current-equalizing resistances to the gate bus, which is in turn connected to a +50 volt terminal through a series circuit of a 150-ohm resistance and a 27-mh inductance. This series circuit causes the voltage waveform to be critically damped during the transition from 0 to +50 volts. A final factor concerning the word gate network is that the word buses were each divided into two equal sections, each section connecting 16 words, in order to reduce the capacitive load to each gate driver. A separate set of 32 gate drivers was provided for each of the two sections.

The BIT drive is also provided by the same type of transistor circuit as is used for driving the load-sharing switch. One end of the BIT winding is connected through a current-determining resistor to the transistor collector and the other end through a terminating resistor to +60 volts. The terminating resistance is made equal to the characteristic impedance (140 ohms) of the BIT windings.

The waveforms in Fig. 4 are oscilloscope displays of READ and WRITE word-winding currents, a gate bus voltage and a BIT-winding current. The negative transients following both the READ and WRITE positive current pulses are the result of the reverse recovery characteristics of the diodes. The waveforms show that the diode reverse recovery time is approximately 0.1 μ sec for this circuit application.

The following factors were prime considerations in the particular word selection network shown in Fig. 3. The first is that the diodes were placed on the drive, rather than on the gate terminal of the word windings, in order to reduce the capacitive load to the drive circuit. This arrangement eliminates the capacitance of the 31 nonselected drive windings associated with a given drive bus. This is desirable since the relatively high-impedance drive source provides a pulse of essentially constant current to the drive bus. Therefore, the displacement current that would be conducted by a parallel capacitive load during and immediately after rise and fall of the current pulse is subtracted from the desired current at the selected word. On the other hand, the increased capacitive load to the drivers caused by direct connection of the word lines to the gate buses is not as detrimental, since the source impedance of gate drivers may be low and inductively compensated. Also, transition times for the gate voltage do not determine the waveshapes of the drive currents. The second factor is that both the READ and WRITE diode matrices are operated by positive drive voltages and require identical gate voltage values. This permits a single gate-driver to be used for each column of both matrices and also allows the selected gate to remain in the "on" state during both READ and WRITE drives. Alternative techniques would require time-consuming transitions in gate voltage in the interval between READ and WRITE drives. The requirement for the single


gate-driver technique is that separate word windings be provided for READ and WRITE currents. However, the inconvenience of providing separate word windings was not sufficient to outweigh the advantages of the single gate-drive technique. A third factor is that a high impedance or current source was employed for the word drive because of the inductive nature of the load. Another factor of importance is that this type of drive provides a constant magnetizing current to the word windings regardless of variations in diode voltage drop and storage core back-voltage. This constant-current source is readily

obtained by placing current-regulating resistors in series with the matrix-switch primary windings.

Array organization

Figure 5 shows the storage array for the memory system, with 147,456 ferrite cores for storing 73,728 bits of information. The cores are 30 x 50 x 11 mil toroids. As in three-dimensional memory arrays, the construction is simplified by dividing the array into several separate plane assemblies. However, contrary to the usual three-dimensional arrays in which only one bit or core of each word

Figure 5 A complete array assembly without external connections.







Figure 6 (a) Portion of core plane assembly, showing placement of windings and cores. Dashed line shows division between the two mats.

is contained in a given plane, each plane for the two-dimensional memory contains all bits or cores for many complete words of information. This word-organized plane format provides a favorable arrangement for the interconnection of planes to form the complete array assembly.

The schematic diagram in Fig. 6a illustrates the placement of windings and cores in one of the plane assemblies. Whereas only one pair of cores for one bit of storage is shown with both vertical and horizontal windings, an actual plane contains 128×72 , or 9216, pairs. With this complement of windings and cores, each plane stores 128 words of 72 bits each.

Each vertical set contains two windings: one, a two-turn READ winding and the other, a single-turn WRITE winding. The horizontal set contains two BIT-drive windings, one for each of the two cores per bit of storage, and one sense winding which passes through both cores. As shown, each BIT winding is returned across the plane in the space between two rows of cores. This technique is used in order to reduce the self-inductance and the mutual coupling between sets of horizontal windings. The sense winding is crossed in the center of each plane in

(b) Crossover between plane frames.



order to cancel undesired mutual coupling between BIT and sense windings of a given horizontal set.

In actual practice, the planes in the storage array are made up of two separate mats of cores and windings. Figure 6a schematically shows how each core plane is divided into two mats, with the horizontal dashed line indicating the division. The two separate mats are superimposed in a position that would be obtained by folding the drawing on the dashed line.

Each READ and WRITE winding contains a diode, one end of which is connected to a terminal on the edge of the core-plane frame. Four sets of READ and WRITE drive buses are connected to the other ends of the diodes along one edge of the frame, each bus being connected to 32 diodes.

The array is completed by stacking eight identical planes with every other plane rotated 180° , such that half of the diodes are placed on the top and the other half on the bottom side of the array. The opposite ends of the word-windings from the diodes are then joined within and between planes, the interconnections forming two sets of gate buses, each set containing 32 buses. Physically, one set of gate buses is located on the top side of the array and the other set on the bottom. To complete the array assembly, BIT and sense windings are connected in series from plane to plane on the front and back sides of the array. A second sense crossover is made between plane frames, as shown in Fig. 6b, to further reduce the effects of mutual coupling between BIT and sense windings. The left and right sides of the array are connected to ducts in order to permit forced flow of temperature-controlled air through the array.

BIT and sense system


Figure 7 schematically illustrates the BIT and sense system circuits. Each BIT winding is driven by a single circuit through a current-adjusting resistor, R_1 , and terminated

with a resistor, R_2 , equal to the characteristic impedance of the winding. Each sense winding is segmented into two equal-length sections, the first sensing all bits of address 0 through 511, and the second, all bits of address 512 through 1023. This is done to reduce the amount of noise developed across the input terminals of the sense amplifiers due to application of BIT drives during the WRITE or regenerate operation. This noise, which results from minor flux changes in the cores reversibly switched by the BIT current, is not completely cancelled by the crossovers in the sense windings. There is a lack of complete cancellation because of a difference in amount of reversible flux change for cores in the ϕ_U state vs those in the ϕ_L state. The result is a maximum delta noise voltage (of about 0.75 v) for the case of storing all ONES or all ZEROS on a given sense winding. Ideally, only for the case of storing half ONES and half ZEROS in the bits associated with a given sense winding will a zero noise voltage be developed by application of the BIT drive.

The sense amplifiers are selectively strobed under address control in order that only those amplifiers connected to the array segment being addressed produce output signals. This prevents any spurious noise in the non-addressed segment from interacting with the signal.

Although not shown in Fig. 7, a pulse transformer connects each sense segment to its associated amplifier. The effect of the common-mode noise caused by capacitive coupling between the parallel BIT and sense windings was minimized by incorporation of the pulse transformer.

The waveform photograph in Fig. 8 shows the sense signal as observed at the input terminals to the sense amplifier. The waveform which shows a positive (upward) signal at the left of the photograph is the output obtained when reading a ONE and the one with a negative signal is for a ZERO. The large, bipolar deflections to the right indicate the noise voltage resulting from application of the BIT drive.






Figure 8 Sense signal at input terminals of sense amplifier.

Horizontal scale: 200 nsec/div.

Vertical scale: 250 mv/div.

Logical organization

In designing the 0.7 μ sec memory, considerable time was spent in evaluation of means for reducing logic circuit delays and skew. These two factors will add considerable time to the total memory cycle. Most of the functions performed by the logic circuits, such as address-decode, address and data register storage and selection, et cetera, are well known and therefore will not be described. However, the logical organization of the memory does possess some design sophistications believed to be of interest.

Under computer control through a memory instruction channel the memory will read an address, alter the word read in a prescribed manner, and then write the new information during a single memory cycle. Because the various alterations to be performed on the word require different intervals of time, the information in the four-bit instruction channel also specifies the time interval between READ and WRITE drives. To accomplish the electronic control of driver timing, a delay-line type of timing source (clock) was designed. The clock receives a select-memory pulse whenever a memory cycle is to be executed. This pulse is transmitted through one group of several delay lines, connected in series, with signals extracted from each line to time the sequence properly in the memory cycle. Four groups of series-connected delay lines are incorporated, each group providing different timing intervals.

An overlap technique is used to reduce the effective memory cycle time. The new address is supplied to the memory while regeneration at the old address is still in progress. A second address register is used, which stores the decoded address information so that the primary address register and decode circuits are available to receive new information before completion of the existing memory cycle.

A final item of interest concerning logical organization is that memory operation is asynchronous in that its timing is all self-contained. The computer need only provide address and instruction information and a select-memory pulse. From the start of the memory cycle, however, there is a constant exchange of information

between the memory and the computer. This exchange provides for determining priority, and monitoring the internal operation of the memory.

Summary

Reliable operation exceeding original design objectives has been demonstrated by extensive engineering testing, utilizing a specially developed "memory exerciser." A minimum cycle time of 0.68 μ sec was achieved. To the best of our knowledge this toroidal ferrite memory represents a new capacity level for sub-microsecond READ-WRITE memories.

Over 3000 hours of actual operation in a computer system has demonstrated the utility of this memory to the computer.

The development of the 0.7 μ sec memory has provided a better understanding of the required operating parameters of devices for this type of memory organization in this performance range. For example, the long-standing requirement for square-loop cores with favorable 2:1 selection parameters can be modified in favor of using cores which do not have the square-loop characteristic. This type of information, along with first-hand experience with common-mode noise, environmental control, transmission lines and termination, will prove valuable for future memory developments directed toward higher performance and larger capacities.

Acknowledgments

In a project of this magnitude and complexity, it is impractical to mention all the persons who contributed to its success. However, the over-all program supervision and guidance provided by M. A. Every deserves specific mention. Particular acknowledgments are also given to the several engineers and technicians directly assigned to the project and responsible for the design, construction and testing of the 0.7 μ sec ferrite core memory. In addition, the many supporting activities that contributed specialized talents so necessary to the success of the program are gratefully acknowledged. In this category, the outstanding efforts of the core plane and the array assembly groups of the IBM Poughkeepsie Manufacturing Organization are acknowledged.

References

1. R. E. McMahon, "Impulse Switching of Ferrite," *Digest of Technical Papers, Solid State Circuit Conference*, Philadelphia, Pa., February, 1959. Pp. 16-17.
2. E. Foss and R. S. Partridge, "A 32,000-Word Magnetic-Core Memory," *IBM Journal*, 2, 102-109 (April, 1957).
3. G. Constantine, Jr., "A Load-Sharing Matrix Switch," *IBM Journal*, 2, 204-211 (July, 1958).
4. C. J. Quartly, "A High Speed Ferrite Storage System," *Electronic Engineering*, 31, 756-758 (December, 1959).
5. C. A. Allen, G. D. Bruce, and E. D. Council, "The Design of a 2.18 Mc Core Storage Unit," *IRE Transactions on Electronic Computers*, EC 10, 233-237 (June, 1961).

Received April 17, 1961