A Thin-Film Rod Memory for the NCR 315 RMC Computer

PAUL HIGASHI

Abstract—**This paper describes a thin-film Rod memory that is used as the main store in the NCR Rod Memory Computer. The production procedure for the Rod storage element and the fabrication and assembly of the memory stack is outlined. The mode of operation of the memory is discussed. Line characteristics, design criteria, as well as a description of word, sense, digit and timing circuits are also presented.**

I. Introduction

THE BASIC requirement for any high speed digital computing system is a fast random access memory. Many of today's data processing problems also demand a large storage capacity. Added to these requirements is the need to keep costs within commercially feasible limits.

The NCR 315 Rod Memory Computer (RMC) requires a memory with an access time of approximately 300 ns and a complete cycle time of 800 ns. A single computer is capable of addressing a minimum of 20 000 and a maximum of 80 000, 13-bit words. A thin-film Rod memory was designed to provide these features.

II. Storage Element

The Rod is an electrodeposited thin-film element which has been developed in the Research Department of the NCR Electronics Division in Hawthorne, Calif. Detailed descriptions of Rod characteristics have been presented in previous papers [1], [2]. A brief discussion, sufficient to understand the application of the Rod in this memory system, is presented in this paper.

The thin-film material is 97 percent iron and 3 percent nickel. A film approximately 4000 Å thick is continuously deposited on a 10-mil diameter beryllium copper wire. The film is essentially isotropic and has a coercivity of approximately 16 Oe. When the Rod is driven with a field equal to twice the coercive force, switching occurs in about 25 ns. This condition assumes a field rise time of less than 10 ns. Under normal memory operating conditions, the drive current rise time is approximately 120 ns, and the switching speed is also about 120 ns. An axial switching mode is used with the Rod. The magnetic material is continuous along the Rod, so that the bit is defined by the field driving it. The field is produced by a current traveling through a ten-turn solenoid of 20-mil inner diameter, into which a Rod is inserted. This mode of switching produces an open flux path and leads to some disadvantages. The element could be disturbed by stray magnetic fields; also, demagnetizing fields produced by the discontinuity of the flux limit the bit density. However, the large coercivity of the material minimizes both of these effects.

Manuscript received February 18, 1966.
The author is with the Electronics Division, National Cash Register Company, Hawthorne, Calif.

III. Rod Production Process

A commercially usable memory element must be easy to produce and test. Therefore, a highly automated processing and testing procedure is necessary. In addition, the yield must be high to permit low production costs. The Rod compares favorably with other magnetic elements in these respects. In the present memory system, the finished Rod, prior to stack assembly, is a 100 percent tested 6-inch section with a 96-turn flat ribbon helical winding. The winding is soldered to the substrate at one end.

Rod production consists of three steps. First, a large reel of 10-mil BeCu wire is fed at a constant rate through various plating, cleaning, and coating tanks. The B-H loop of the plated Rod is continuously monitored during this process. The plated wire also passes through a pulse-testing fixture, where a pulse pattern checks the material to specified uVl (min) and dVz (max) limits. Bad sections of Rod are detected by threshold devices and are automatically recorded on a visual readout as failures per feet tested. The B-H loop and pulse testing are used for process control. When the proper temperature, plating rate and plating solutions are established, the process remains stable for long periods of time and with periodic monitoring of the process, a high yield can be maintained. After the plated wire has passed through the test stations, a thin coating of urethane is applied. The wire continues through a curing oven and then is wound on a large diameter reel.

The second step of the process includes the placing of a helical winding around the plated wire, 100 percent pulse testing, and cutting into 6-inch sections. When the plating run is completed, the take-up reel is placed in another line. The wire is then fed through an automatic winder which places a 96-turn-per-inch helix around the Rod. (See Fig. 1.) The wound wire is again coated with urethane, cured, and then passed through a test fixture which tests to prescribed acceptance limits.

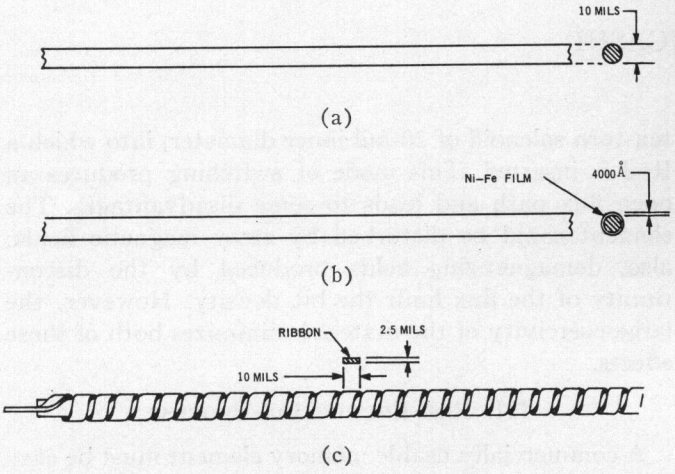

Fig. 1. Thin-film Rod element. (a) Beryllium-copper wire. (b) Wire electroplated with Ni-Fe alloy. (c) Copper ribbon wound wire.

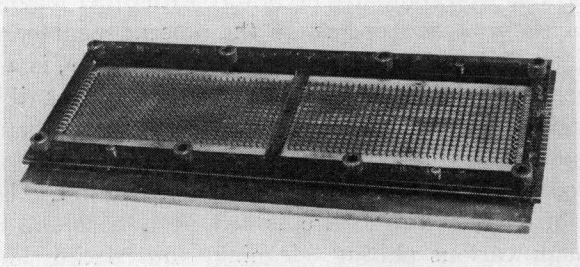

Fig. 2. Solenoid wound on pin jig.

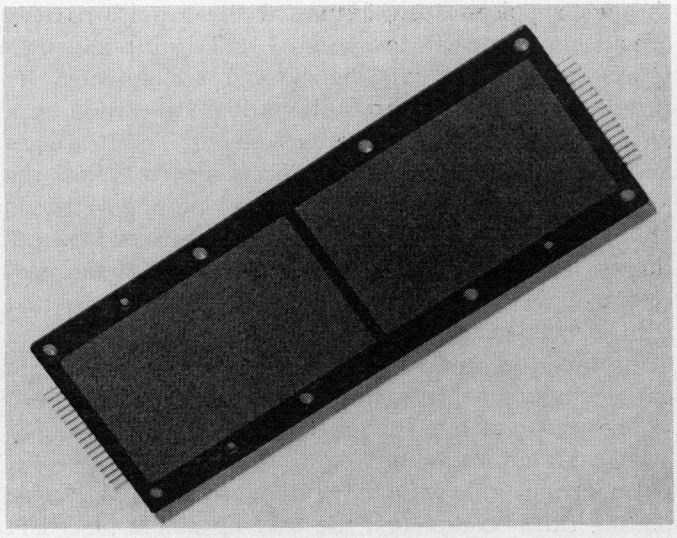

Fig. 3. Potted solenoid plane.

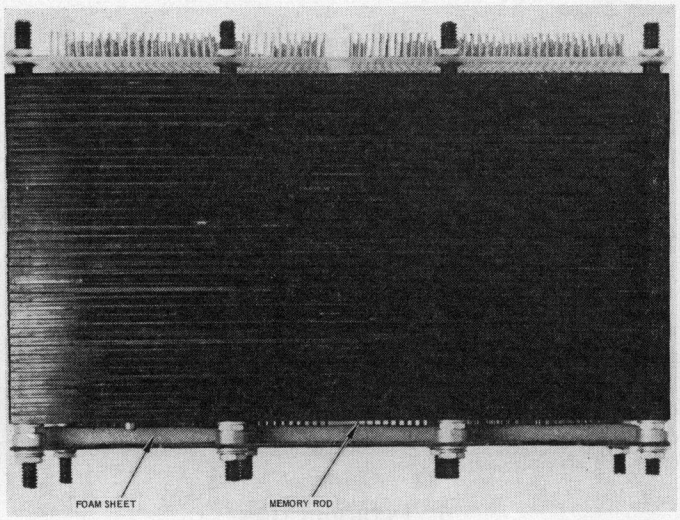

Fig. 4. Basic module.

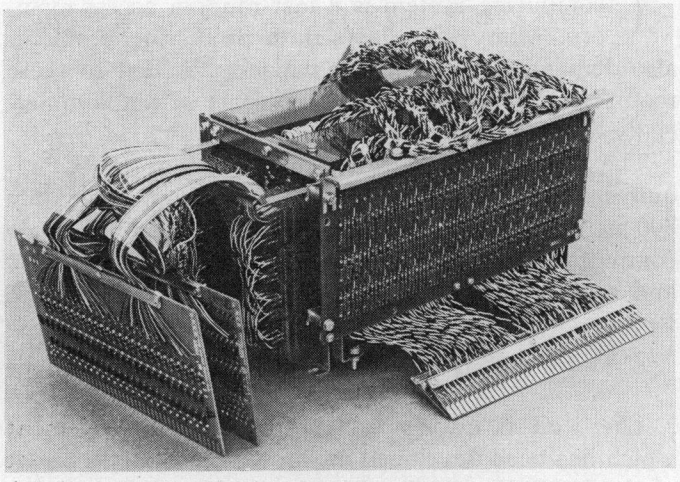

Fig. 5. 5000 word stack.

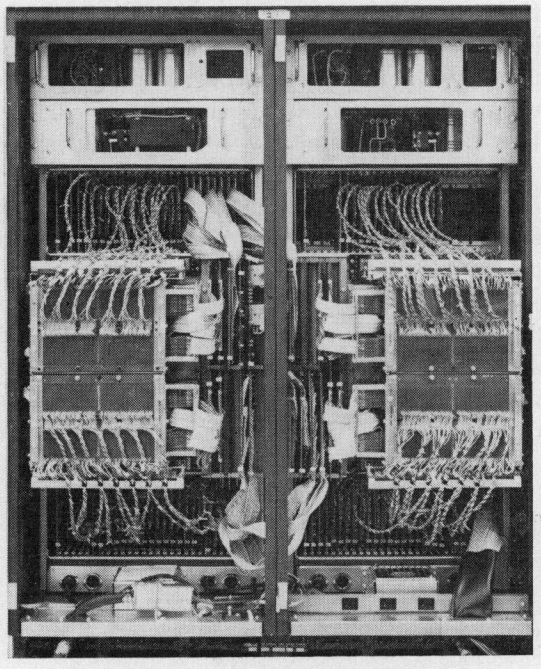

Fig. 6. 315 RMC memory bay.

Automatic test equipment records failures in each successive 6-inch length. At the end of the line, the wire is cut to 6-inch lengths, and a mechanism controlled by the test equipment directs the Rod to either an accept or reject bin, depending on the result of the acceptance test of the section.

During the third step of the process, one end of the Rod is dip-soldered so that the winding is shorted to the substrate. Each Rod represents 40 bits. The 40-bit sections are then ready to be assembled into a stack.

IV. Stack Fabrication

The read-write drive windings are pre-formed 10-turn solenoids, which are potted to form a plane. A plane has 16 rows of solenoids. There are 65 series solenoids in each row.

The windings are formed around a precisely tooled pin jig as shown in Fig. 2. The winding process is mechanized and the complete plane can be wound in a few minutes. The windings are then potted within a molded frame (Fig. 3). One end of each word line is terminated on a pin which is molded into the frame. The other end of the lines are anchored to pins at the opposite side of the frame. The lines are then returned directly back along the rows of solenoids and terminated at the other end. All the return wires are terminated on one pin. The solenoids have a 0.02-inch inner diameter and are approximately 0.062 inch high. The planes are 0.125 inch thick. Forty planes are stacked and Rods are inserted in the aligned solenoid openings. This combination forms the basic module (Fig. 4). Two basic modules are mounted side by side to form a 5000 word unit (Fig. 5). It will be noted that there are 16 word lines per plane and 40 planes per basic module. Two of these modules provide 1280 word lines. There are five 13-bit computer words per line. The unit, therefore, could provide 6400 words. However, only 25 word lines per plane are used in the main memory. Some of the spare lines are used for an auxiliary memory and for special table storage. There are four 5000 word modules in one memory frame. Common access circuitry is used for the entire 20 000 word memory. Up to three additional 20 000 word units may be added to a system. Figure 6 is a photograph of a 20 000 word 315 RMC memory bay.

V. Mode of Operation

The memory uses a two-wire linear select mode of operation. The solenoid windings serve as read-write lines and the winding on the Rod has the dual function of a sense and digit line. During reading, a full select current is applied to the selected word line. The sense line can then transmit the Rod signal to the sense amplifier without the presence of any digit current. During writing, a two-thirds full select write current is applied to the word line and a one-third current (either inhibit or enable) is applied to the digit line. Actual current magnitudes used in the memory are approximately 480 mA read, 280 mA write, and 160 mA digit.

VI. Word Selection Circuitry

In a linear select system, it is inherently expensive to have a short word length. In addition, only a limited number of bits can be placed on one sense-digit line due to impedance and delay considerations; therefore, the decision was made to place multiple words on a drive line and to perform partial word selection at the output of the sense amplifiers. Consequently, five words (or 65 bits) were put on each drive line. The word lines are selected by means of a conventional row-column factoring scheme. For the 20 000 word memory, the factoring is 50×80 on the word lines. The 50 and 80 row-column transistors are driven by matrices of tape-wound switch cores. The primary windings on the switch cores are also factored so that another decoding level is obtained. Diode gates on the input to the switch core drivers provide the first level of decoding [Fig. 7(a)]. The addressing is binary-coded decimal. Each switch core has four windings: a primary, a bias, and two secondaries [Fig. 7(b)]. The bias winding carries a constant current which opposes the primary drive. The selected core is driven in a constant voltage mode against the bias current. The secondary windings are coupled to read and write drive transistors. The polarities of the secondary windings are such that when the core is driven, the read transistor is turned on. When the drive is removed, the constant current bias resets the core and the write transistor turns on. The reset of the core is clamped to a voltage equal in magnitude to the set voltage, so that in both the read and write cycles, a constant voltage drives the transistor.

One of the problems encountered in a word-organized memory is caused by the many word lines that are common to each driver. The capacitance of the lines must be isolated to avoid unreasonable delays in current rise times. The back-biasing scheme shown in Fig. 8 minimizes the capacitance that must be charged by the word-line drivers. If transistors Q1 and Q3 are the selected read driver and switch, the voltage at point B will be slightly below 30 volts. All the diodes on the selected plane are back biased except diode D2. Only the capacitance of the word lines on the selected plane must be charged through the column switch Q3. The word to digit line capacitance is appreciable (0.4 $\mu\mu$F per bit), however, the windings are isolated from system ground and each other through common mode transformers. Essentially, only the capacitance from word plane to word plane must be charged. The digit winding becomes a coupling element in the charging path. Since the switch Q3 is connected to a low impedance voltage source, the capacitance is charged with little delay. The driver Q1 connects the read current source to point A. The voltage at A will fall from 30 volts to the lower clamp level (12 volts). The unselected planes are held at

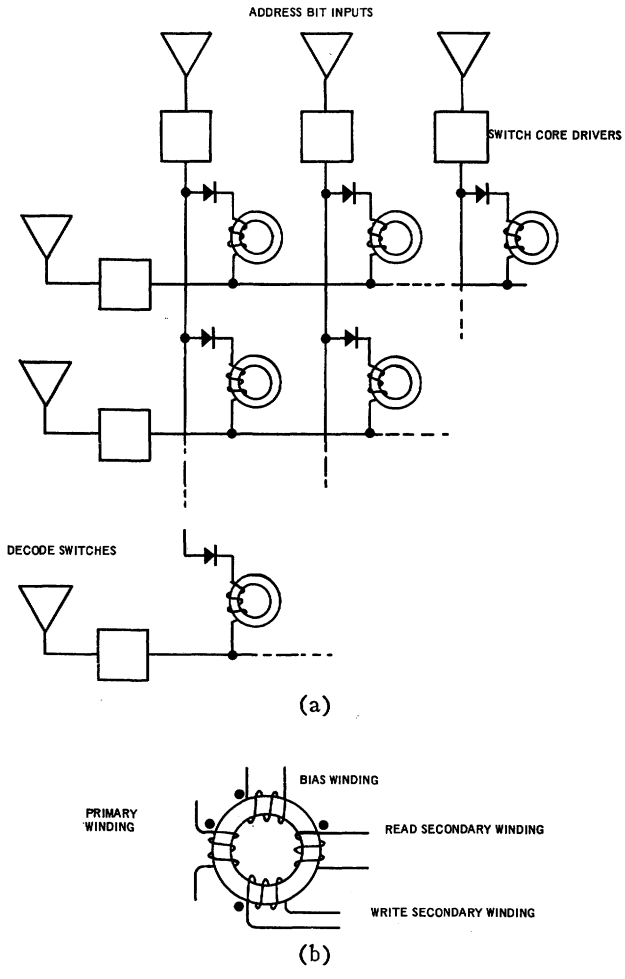

Fig. 7. Switch core address decoding. (a) Switch core matrix. (b) Switch core winding configuration.

12 volts through resistors R10 through R13. The diodes of the unselected lines are never forward biased. Therefore, the current source need only charge the junction capacitances of the unselected diodes that are common to the word driver. An identical back-biasing scheme is employed in the write portion of the cycle. One 30 volt power supply is used for both read and write current sources. The 12-volt clamp and reference voltage is also derived from the same supply. Note that approximately 18 volts are available across the line during reading and 12 volts during writing. The division of voltage was dictated by the larger magnitude and more stringent rise time requirement of the read current.

VII. Sense-Digit Line

The helical windings on the rods are interconnected to form the sense-digit line. The winding is split into two parallel paths, which form a balanced transposed pair as illustrated in Fig. 9. The digit planes are perpendicular to the direction of the word lines. Each stack comprises 65 sense-digit planes spaced on $\frac{1}{8}$ inch centers.

There are 1000 bits on each sense-digit line. The delay and attenuation characteristics of the line dictate the number of bits when the line is used for sensing. The delay along a sense line pair is approximately 40 ns and the characteristic impedance of the line is about 200 ohms. The parallel inductance of the halves of the digit line is approximately 2.5 μH. Mutual inductance between digit lines, usually a severe problem in linear select core memories, is practically negligible in a Rod memory. As mentioned previously, the Rod windings are soldered to the substrate at one end of the Rod. The

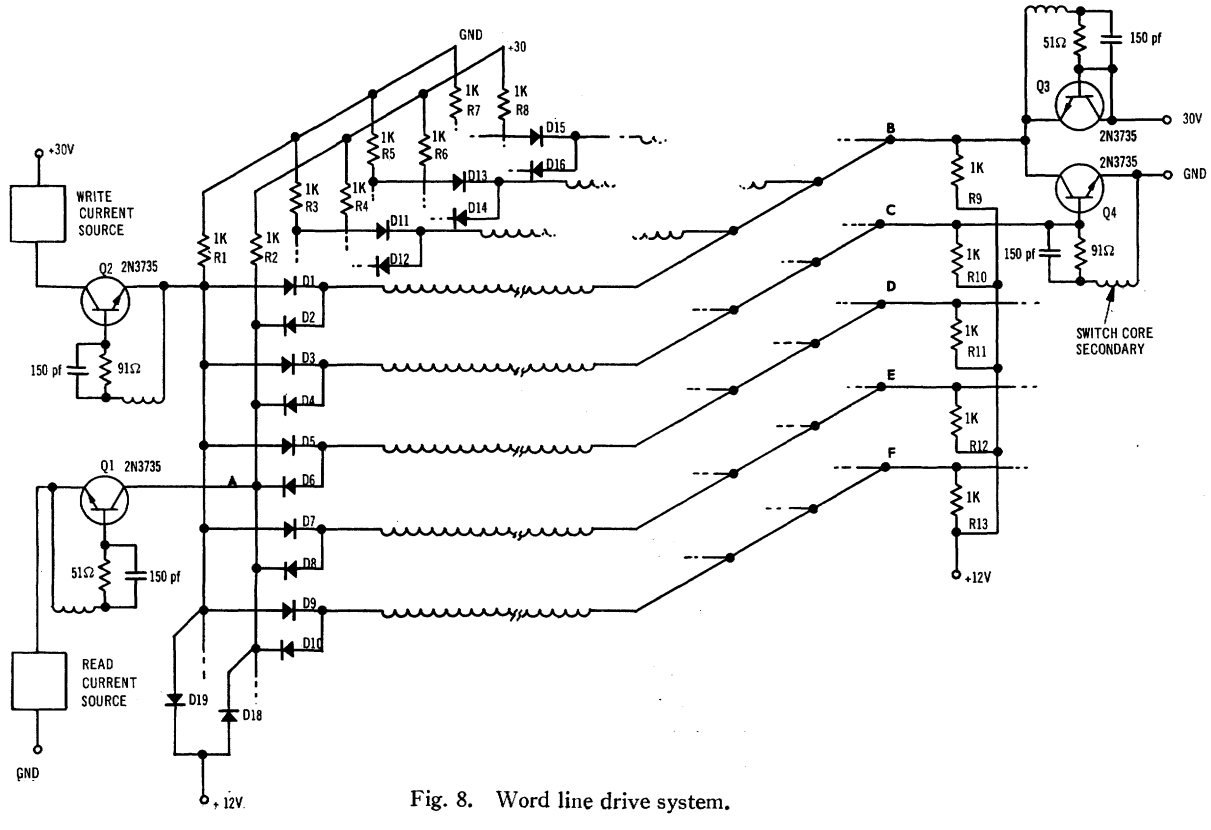

Fig. 8. Word line drive system.

Rod serves as a return path for digit current and effectively cancels the axial component of current due to the finite pitch in the helical Rod winding. The main magnetic field is axial. The coupling of this field to adjacent windings is negligible because of the small diameter (0.01 inch) of the helix.

Capacitive coupling between digit windings, however, proved to be more of a problem. If currents of opposite polarities are driven through adjacent digit lines, the voltage difference between adjacent rods at the ends of each line is maximum. The capacitive interaction between digit lines was significant and resulted in varying degrees of current waveform distortion, depending on the pattern of ones and zeros being written. An effective solution to the problem was to isolate each line from the other by connecting the digit voltage to each digit circuit through a common mode transformer (Fig. 10). The common mode transformer (T5) presents a low impedance to circulating currents, but a high impedance to common mode currents. This high series impedance negates the effect of interwinding capacitance by making the charging time constant much longer than the pulse width of the digit current. Since the line appears to be mainly inductive, terminating resistor values are chosen from a time constant consideration only. The resistance is equally distributed in each leg of the H-switch, so the load is symmetrical. Digit current propagation time is also reduced to a few nanoseconds; thus, no phase-compensating adjustments are necessary in the write timing.

Although a large signal-to-noise ratio is of prime importance in establishing good discrimination margins, fast sense line recovery after writing is essential to reduce cycle times. The digit current rise time is exponential and is determined by the L/R time constant of the line. When the H-switch transistors are cut off, the digit current falls rapidly, since the loop resistance is now very large. If the line could be critically damped, a minimum recovery time could be realized. In the actual case, however, oscillations occurred which prolonged the recovery time considerably. The main ringing circuit was formed by coupling between the digit and word lines. The digit lines are capacitively coupled to the word line at each bit intersection. Each group of word lines on a plane are common at the column switch end. The net ringing current was observed to be flowing between each half of the common word lines. The oscillations were damped by separating the halves of the common word lines and connecting each half separately to the column switch through diodes.

Sensing is done across the center of the sense loop, as shown in Fig. 9. By sensing across the center of the line, signal propagation delays were halved, and digit drive common mode voltage at the sense amplifier input became negligible regardless of digit current polarity. The ends of the sense lines are shorted. Theoretically, a better signal-to-noise ratio is obtained by terminating the line in its characteristic impedance. However, in the Rod memory the signal-to-noise ratio was improved with a shorted termination. Signals of bits next to the sense amplifier input increased by approximately 60 percent, whereas the worst case noise increased only slightly. This is probably due to the fact that most of the noise is caused by capacitive coupling, which is affected not by terminations across the sense line, but by terminations to ground.

There are two sources of noise in any magnetic memory; namely, electrostatic and electromagnetic coupling. In a linear select memory, the electrostatic coupling presents the more serious noise problem. There is no partial select or delta noise during the read portion of the cycle. Inductive noise is merely the 0 output of the selected bit, plus the air coupling of one word solenoid to the sense line. In addition, the inductive noise is always of the same polarity as the 1 signal, so that signal and noise are additive rather than cancelling. This is a significant point because the Rod material switches in the rise time of the read current. The noise resulting from reversible flux switching as well as air coupling also occurs during the current rise time. Phase discrimination between noise and signal is therefore impossible.

The voltage applied to the word lines to obtain a given current rise time is capacitively coupled to the sense line. If the voltage and capacitance to each half

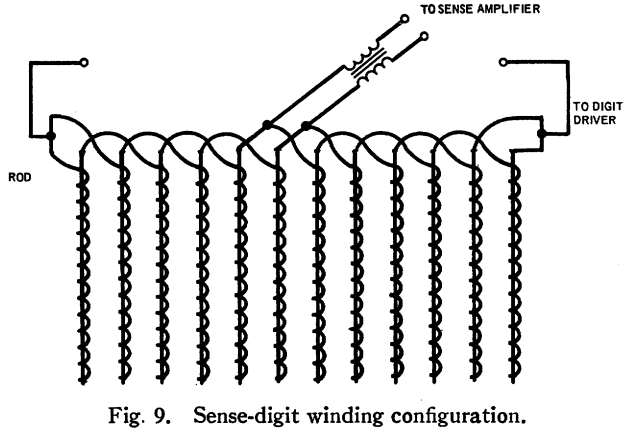

Fig. 9. Sense-digit winding configuration.

of the sense line were exactly equal, all the coupling would appear as common mode voltage. However, in any large system, some degree of imbalance exists. Winding configurations are designed to minimize this imbalance. In the Rod memory, the sense-digit lines are formed so that cancelling bits on either half of the sense winding are physically adjacent. Word-line factoring was also selected so that the minimum number of word lines experience a voltage excursion. The lines common to the selected driver at the diode end are isolated from voltage changes. Only the word lines that are common to the selected switch at the other end undergo a voltage swing. A voltage drop occurs across the selected word line, but not across the other common lines. Therefore, during the read current rise time, a different voltage is impressed on the bits of the selected line than on the cancelling bits on the adjacent line. A noise cancelling plane in the middle of the stack was proposed to offset this one-bit capacitive and inductive imbalance in the system; however, the overall system noise level was found to be low enough so that the added expense of noise-cancelling circuitry was not warranted.

VIII. Digit Circuitry

As was previously mentioned, the Rod memory uses a one-third enable or inhibit current to write a one or zero. Similar systems use a constant one-third read bias. In this case, the read, write and enable current magnitudes are each two-thirds full select. The read bias was not used because it presented a large dc power drain. It also required the fast switching of a two-thirds rather than one-third full select digit current. Furthermore, a well-balanced sense-digit line would be required, since a dc current would be flowing in the sense line during reading.

The one-third enable-inhibit mode of writing necessitates the switching of bi-directional currents. The use of an H-switch proved to be an efficient way to mechanize the function. The memory is required to perform either a read-restore or a clear-write cycle at the request of the computer. The memory does not include a data register; therefore, writing is accomplished either from the computer memory register or directly from the output of the sense amplifiers. Figure 10(a) shows the circuit configuration chosen to mechanize the function. The two AND gates supply the information to be written. The input to one of the gates (SA1) is the output of its corresponding sense amplifier. The other input (SC1) is a logical term which signifies that a read-restore operation is to be performed. The other AND gate inputs (M1 and IC1) are the output of the corresponding bit of the memory register and a term signifying a clear-write operation. The logical propositions SC1 and IC1 also contain address information. Each word line includes five words; all five words are cleared and rewritten during a memory cycle. Three bits of the address register are decoded to select one of the five words. For a read-restore cycle, all bits are rewritten from their corresponding sense amplifier outputs. For a clear-write cycle, new information from the memory register is gated into the proper digit drivers by the decoded address bits. The other four words on the line are restored as in a read-restore cycle.

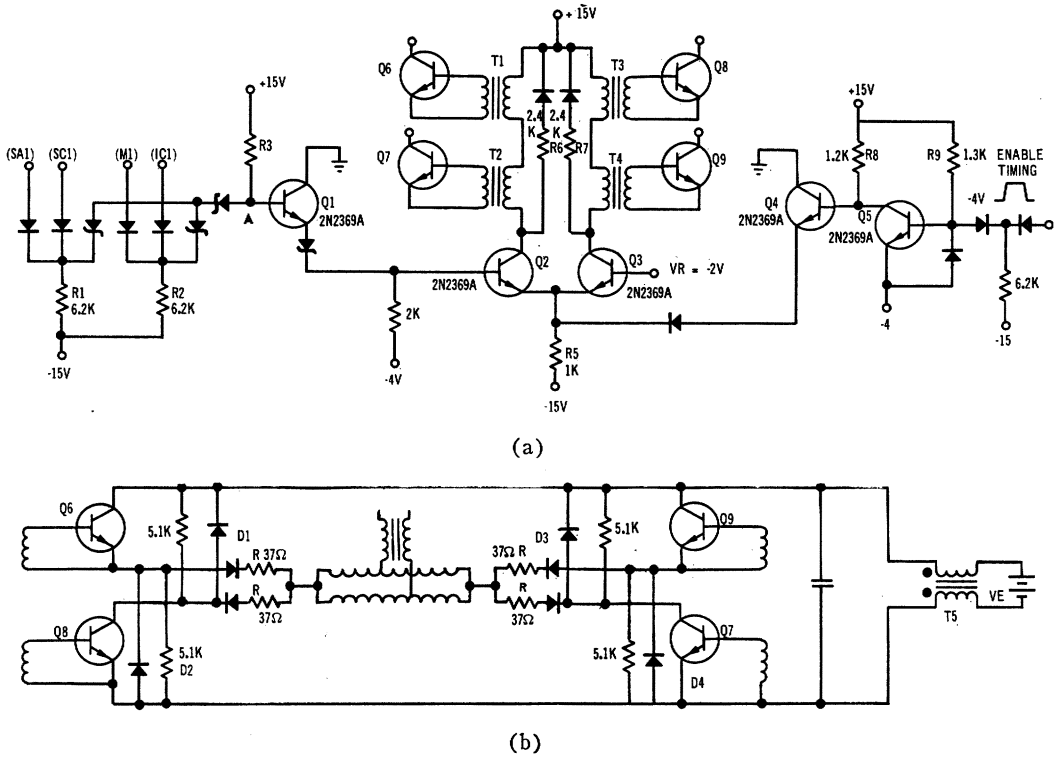

Fig. 10. Digit driver. (Undesignated transistors are 2N3735.) (a) Circuit schematic. (b) H-switch.

All logical inputs swing nominally from 0 to −4 V. The lower level represents a logical one. The diode gates in Fig. 10(a) are coupled to a difference switch through an emitter follower. If either product is true, the voltage at the base of transistor Q2 will be below the reference voltage (V_r) on the base of Q3. If the input logic is false, the base voltage of Q2 will be above the reference voltage. The emitters of Q2 and Q3 are connected to a current source (R5 to −15 V). Transistor Q4 is also connected to the current source. Normally Q4 is saturated and neither Q2 or Q3 can turn on regardless of the state of the input logic. An enable timing pulse turns off Q4 and either Q2 or Q3 will turn on depending on the input logic. The collectors of Q2 and Q3 are transformer-coupled to the base emitters of transistors Q6, Q7, Q8, and Q9. These transistors form the H-switch that provides the bi-directional current drives through the digit lines as shown in Fig. 10(b).

IX. Sense Amplifier

The Rod memory requires a fast, stable sense amplifier that can amplify and detect bipolar signals. The primary stage is a transformer-coupled differential voltage amplifier. This is followed by a current amplifier, which feeds a pair of tunnel diode discriminators. The design principles for a similar sense amplifier were presented in a previous publication [3]. Many of the design features were incorporated in this sense amplifier.

Since tunnel diodes are current-threshold devices, the amplifier must amplify and convert the Rod signal to a current pulse that is proportional to the signal voltage. The tunnel diode current-voltage characteristics are shown in Fig. 11. The peak current (I_p) is nominally 5 mA. The load line represents the inverted base-emitter characteristics of the output transistors Q5 and Q6 (Fig. 12). When the diode is used to discriminate between a one or a zero, it is biased to a point A. The one

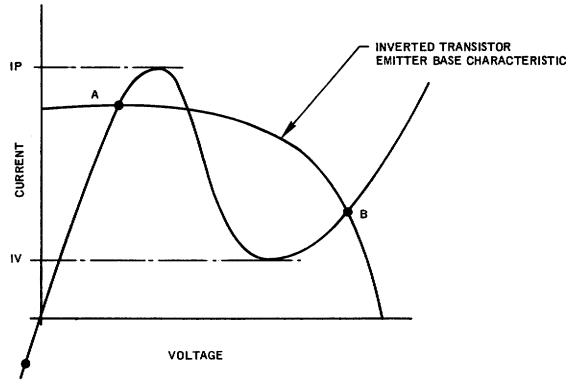

Fig. 11. Tunnel diode characteristics.

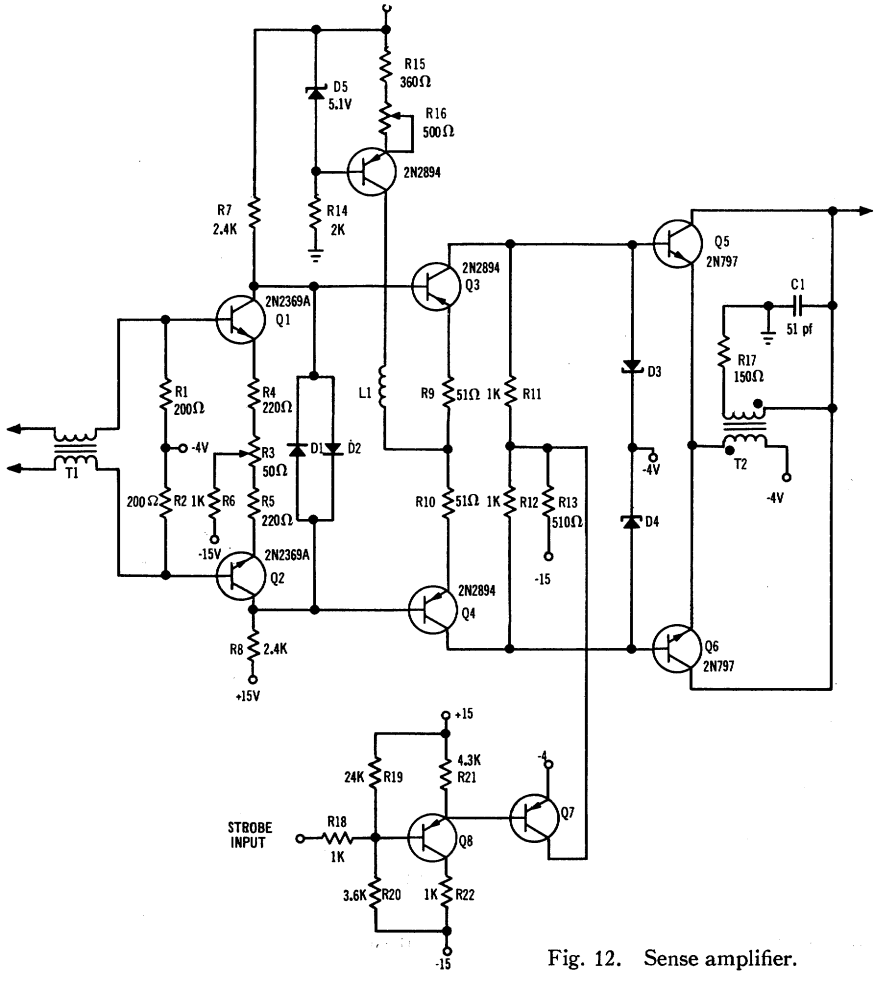

Fig. 12. Sense amplifier.

signal must add sufficient current to the bias so that the peak current is reached and the diode is switched to its high voltage state (point B). Conversely, a zero noise must not cause the diode to trigger.

The bias current is established by the current source in the second stage of the amplifier. The current divides through transistors Q3 and Q4. The tunnel diodes (D3 and D4) are connected to the collectors of the transistors. A strobe circuit also connects to the collectors, and when transistor Q7 is off, all the collector current of Q3 and Q4 is absorbed by the collector resistor R13. During the read time, a timing pulse activates the strobe circuit and the coupling resistors (R11 and R12) are shorted out, thus allowing the bias current to flow through the tunnel diodes. The division of current in the tunnel diodes will be proportional to the differential voltage on the bases of transistors Q3 and Q4. If the differential voltage is sufficiently large, one of the diodes is triggered. The high voltage state of the diode is sufficient to forward-bias the emitter-base junction of its corresponding output transistor. The strobe must be removed before digit noise is introduced into the amplifier. The removal of the strobe resets the tunnel diodes. The transformer in the emitter and collector circuits of the output transistors makes the stage a blocking oscillator, and the output pulse is stretched sufficiently, so it can be used directly in the restore portion of the memory cycle.

Two potentiometer adjustments are provided in the amplifier. One (R3) is a balance adjustment in the emitter circuit of the first stage. The other (R16) is in the current source and allows the setting of the tunnel diode bias currents. Normally, careful component matching must be achieved to minimize offsets in differential amplifiers. However, by using a potentiometer to balance the amplifier, closely matched components become unnecessary. The criterion of balance is the division of bias current in the tunnel diodes. To maintain the widest possible overall margin, the division of current to the tunnel diodes should be set so that the margin for either polarity signal is maximum. The amplifier can be balanced by adjusting the balance potentiometer so that the bias plus the worst-case noise of either polarity are equal. This can be accomplished most effectively by balancing the amplifier dynamically with zeros being recirculated in every address of a particular bit position. The tunnel diode bias current is increased until a zero triggers one of the tunnel diodes. The balance potentiometer is then adjusted until triggering ceases. The bias is again increased and the process is repeated until further adjustments fail to prevent diode triggering.

After balancing has been performed, the bias is at a maximum; that is, the bias plus a maximum zero just equals the peak current of the tunnel diode. If ones are recirculated, and the bias is decreased until the first one fails, the minimum bias condition is reached. The bias can now be set halfway between the two extremes, or offset to favor a one or zero, depending on which output is more susceptible to system variations.

It should be noted that the bias current must be greater than 1/2 I_p. Since differential signal current is merely the offset bias current, this must exceed the peak current in order to trigger either tunnel diode. Conversely, the bias current must not exceed I_p. Therefore, no matter how large the signal or noise, the differential current cannot exceed the maximum bias current. This relaxes the requirements for the strobe circuit since the maximum current that must be absorbed is limited.

The signal is coupled to the sense amplifier through a common mode transformer (Balun). This maintains sense-digit line isolation and provides good high frequency common mode rejection. Signal amplitudes range from 50 to 75 mV. A worst-case pattern of ones and zeros on a sense line produces a digit current transition noise that could saturate the amplifier. Back-to-back silicon diodes (D1 and D2) across the collectors of the voltage amplifier prevent the differential voltage from exceeding the forward drop of the diodes. Voltage gain of the stage is approximately 5; therefore, the minimum signal will not be clipped.

X. Memory Timing

In sub-microsecond cycle time memories, the timing pulses which define the time relationship of the various functions within the cycle must be quite precise. Figure 13 is a block description of the timing generator employed in this memory. A memory cycle request signal from the computer initiates the cycle. This request signal may either be a read request (JMR) or a write request (JMW). When there is no request signal, the delay line is at a -4 V level. T1, T2, T3, and T4 are taps off the delay line. The output of amplifier A1 is true if all the inputs to the product driving it are true. When either memory request signal goes true, inverter I1 will introduce a positive signal (0 volts) to the delay line. When the pulse reaches tap T1, the product C will go false and the output of A1 will make both AND gates A and B false. This negative feedback cuts off the positive pulse which was introduced to the delay line. The tap T1, therefore, determines the pulse width of the signal propagation down the delay line. Subsequent feedback taps from the delay line inhibit the initiation of another memory cycle until the original signal has reached the end of the delay line. Therefore, the maximum allowable repetition rate for the memory is established by the delay to the last feedback tap.

The delay line has taps at 10 ns intervals. Each timing pulse required in the memory cycle is formed by two AND gates which are summed into either an amplifier or an inverter. The delay of each timing signal with respect to the initiate signal and the width of the signal can be set by summing two appropriate taps on the delay line.

The relationship of the various timing pulses are shown in Fig. 14. R_t is the read timing signal. It deter-

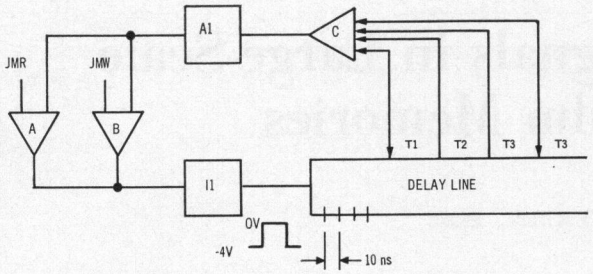

Fig. 13. Block diagram of timing generator.

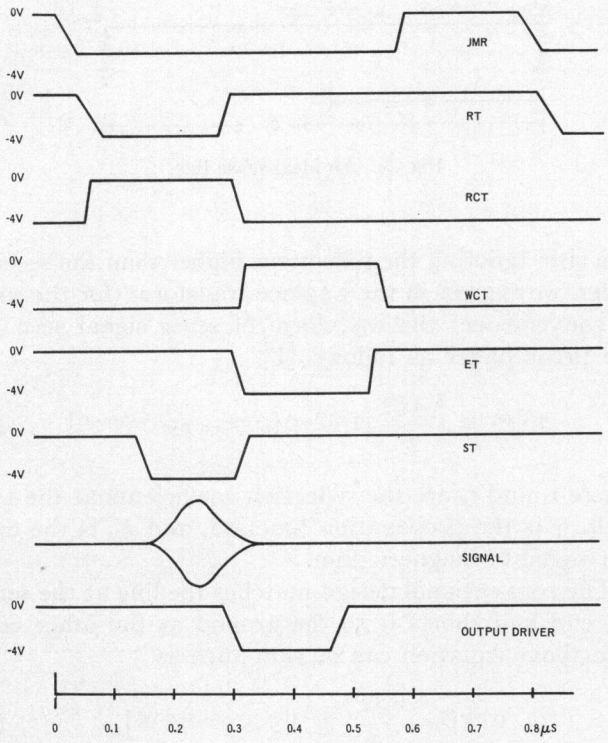

Fig. 14. Memory timing diagram.

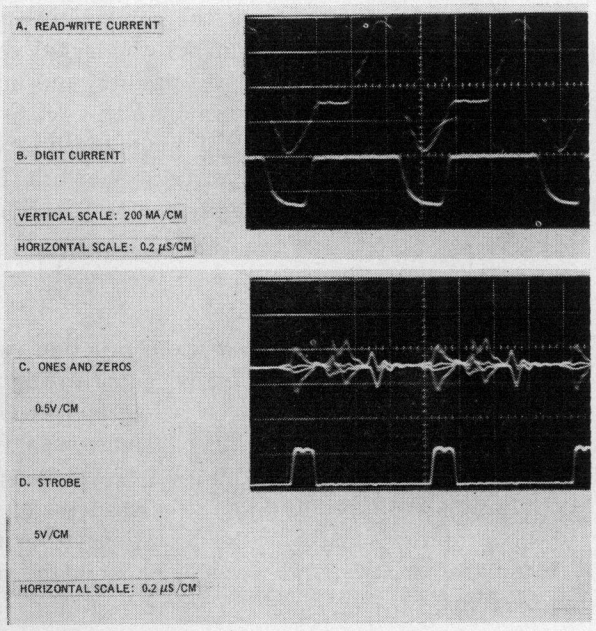

Fig. 15. Current and signal waveforms.

mines the time interval that the address bits are decoded and the selected switch cores are driven. S_t is the strobe timing. It determines the time and duration of the sense amplifier strobe pulse. R_{ct} is the read current source timing. It establishes the time that read current may flow through the selected word line. E_t is the enable timing signal. It gates the digit driver so that either enable or inhibit current flows only during the write time. W_{ct} is the write current source timing signal. It turns on the write current source at the proper time. Figure 15 shows some typical current and signal waveforms.

XI. Conclusion

Thin films have long shown promise as a storage element in fast random access memories. The realization of this potential has been slow in evolving. However, with continued effort and added experience, the rapid development of a thin-film technology, which is competitive with ferrite cores, can be foreseen. The first commercial thin-film Rod memory is a step in this direction. The most encouraging aspect of the program was reproducibility and high yield of the rod storage element. Certainly the uniformity and cost of the element compares favorably with cores. A high degree of mechanization in the stack fabrication also points towards lower costs.

Improvements must be made in reducing line impedances and propagation delays for future higher speed Rod memories. This would, of course, lead to a reduction in the amount of circuitry required. However, the cost of the present Rod memory was within the limits of economic feasibility and the experience gained in the program will certainly point the way for future improvements.

Acknowledgment

Individuals in our Research, Packaging, Product Engineering, and Manufacturing Engineering groups have contributed greatly to the development and production of the first thin-film Rod memory. Since the list of individuals is lengthy, I would like to acknowledge their efforts as a group. However, I would like to give special acknowledgment to the members of the Rod memory development project: G. Yamamoto for his work on sense amplifier design; H. Lee for his contribution to the word line selection and drive circuitry design; G. Solodjagin for his work on the timing and digit circuits; and J. Wood for his untiring efforts in procuring and assembling parts for the development model.

References

[1] D. A. Meier, "Magnetic rod memory," *Proc. Electronic Components Conf.*, pp. 122–128, May 1960.
[2] D. A. Meier and A. J. Kolk, "A magnetic rod, a cylindrical thin film memory element," in *Large Capacity Memory Techniques for Computing Systems*, M. C. Yovitz, Ed. New York: Macmillan, pp. 195–212, 1962.
[3] B. A. Kaufman and J. S. Hammond, III, "A high-speed direct-coupled magnetic memory sense amplifier employing tunnel-diode discriminators," *IEEE Trans. on Electronic Computers*, vol. EC-12, pp. 282–295, June 1963.