Low power computer memory system

by D. E. BREWER
Air Force Avionics Laboratory
Wright-Patterson Air Force Base, Ohio

and

S. NISSIM and G. V. PODRAZA
The Bunker-Ramo Corporation
Canoga Park, California

INTRODUCTION

One of the critical technologies for aerospace missions is electrical power. Increased emphasis is being given to reducing power requirements of electronics as well as toward improvements in power sources themselves. In present day aerospace computers the memory subsystems consume approximately 40 to 60% of the total computer power. This is primarily due to the high drive requirements of present core memories or to the sensing problems associated with thin film memories. Reduction in memory power requirements would, therefore, have a significant effect on extending allowable mission time.

The objective of the effort described in this paper is to demonstrate a capability for the design and construction of an aerospace computer memory having very low power requirements. The device technology selected to accomplish this objective is the metal oxide silicon (MOS) transistor memory array. The MOS transistor is an insulated-gate field effect device formed in a silicon crystal and is suited for low-power applications. The idea of using active devices for storage of data in computers is as old as computer systems. However, until very recently, the power required and the cost of active memories were prohibitive for almost all applications. With the progression of integrated circuits technology from the gate functional level to very large scale integration of system functions, considerable interest has been devoted to development of monolithic memories.[1-4] Predicted power requirement for a random access memory subsystem using monolithics varied over a wide range depending on various proposed designs. This effort primarily emphasizes low power circuit design and memory organization concepts leading to a low power aerospace computer memory that is suitable for economical production.

As the feasibility vehicle for this approach, a 1024-word, 30-bits/word, random access memory was constructed using P channel MOS transistor arrays for storage and bipolar transistors for the peripheral accessing circuitry. This system is an NDRO, electrically alterable, memory with a read access time of 0.7 microseconds, a read or write cycle time of 1 microsecond and has a total power requirement of approximately 3.5 watts. This unit has been successfully constructed and delivered to the Air Force Avionics Laboratory. Significant results were obtained in the areas of low-power circuit design, memory design, and organization concepts using MOS transistor arrays. All of the areas will be covered in detail in this paper.

Outline of technological approach

review of the circuit operation if the basic memory cell

The MOS integrated circuit approach was chosen as a means of developing a low power memory array for two reasons, namely (1) MOS bistable arrays can be operated with very low power by using pulsed operation techniques, and (2) this approach is well suited to large scale integration and reasonable yield, an essential prerequisite from an economy viewpoint.

The peripheral circuitry, which included the address decoder, write drivers, sense amplifier, and control circuits, utilized bipolar transistors and employed commercially available integrated circuits whenever possible. The MOS-to-peripheral system interface circuits were fabricated from discrete components because voltage levels involved in these circuits were not compatible with the relatively low voltages of bipolar inte-

From the collection of the Computer History Museum (www.computerhistory.org)

grated circuits. A savings in power is achieved in the peripheral circuits by employing techniques which minimize power until circuits are actuated by pulse signals. Further power reductions were realized by employing special nonlinear switching circuits with inductive source impedances to decrease the power consumption when driving large capacitive loads in the MOS array and distribution lines.

The system development included a memory exerciser to aid in testing and in demonstrating operation of the memory, and a power supply operated from a 110 volt, 60Hz source to furnish the +13, +12, +4, —4 and —7 supply voltages required by the memory. The exerciser and power supply are of conventional design, were not concerned with power conservation, and therefore, will not be described in detail.

Description of the MOS memory chip

The memory cell circuit, monolithic configuration for the 64-cell chip, and basic principle of low power operation are described in a companion paper.[5] For the sake of completeness in the present paper, a brief review of the circuit operation of the basic memory cell is presented in this section.

Figure 1 is the schematic of the memory cell utilizing P channel enhancement mode, MOS integrated transistor structures, having a nominal threshold voltage of 5 volts The substrate is operated at a 12 volt potential. Transistors Q_1 and Q_2 form the basic bistable storage element of the cell. Transistor Q_3 serves as a "load" for Q_1, while Q_4 is the load for Q_2. Transistors Q_5 and Q_6 are biased off during quiescent memory operation (i.e., neither reading nor writing).

Assuming Q_1 in the bi-stable is on and Q_2 off, the node voltage at the drain of Q_1 will be +12 volts, while the Q_2 node will be near ground potential. During quiescent operation, the capacitance of the Q_2 node will begin to charge toward +12 volts because of the leakage current through the P-to-substrate junctions common to this node. These junctions are constituted by the drains of Q_2 and Q_6, the source of Q_4, and any interconnection crossunders that utilize a P-diffusion onto the substrate. The voltage buildup on the node capacitance by leakage would eventually cause Q_2 to turn on and cause possible loss of the logic state of the bistable. To prevent this, the restore pulse is periodically applied to Q_3 and Q_4, resulting in discharge of the Q_2 node. The Q_1 node voltage is not appreciably affected by restore, since Q_1 is biased on and has a much greater transconductance than Q_3. Periodicity of restore must be sufficient to keep the Q_2 node properly discharged under condition of worst-case leakage current. The restore pulse for each cell

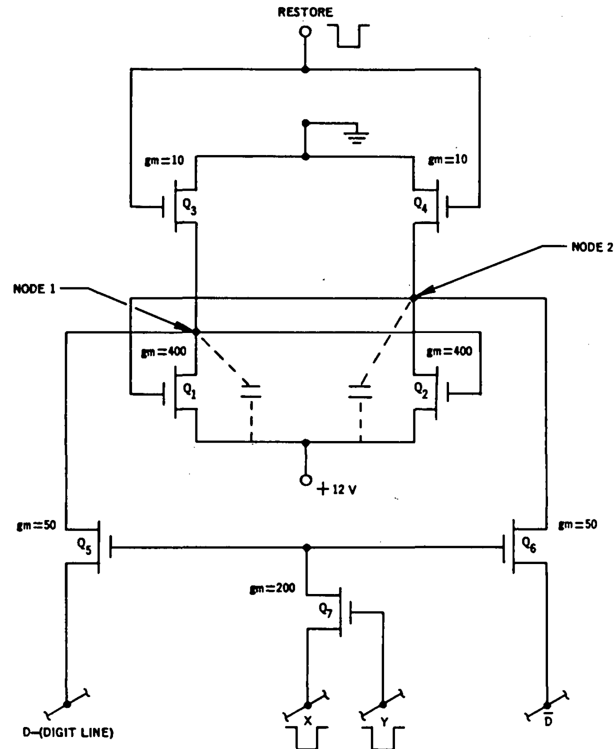

Figure 1—Schematic for basic memory cell circuit

is 18 volts in amplitude, 1.5 microseconds in width and has a repetition rate of approximately 10 KHz.

When reading from the cell, or writing information into the cell, use is made of the Q_5 and Q_6 gating transistors. The cell is addressed by applying the X- and Y-address pulse, as shown in Figure 2 and 3, to the respective source and drain electrodes of Q_7. Coincidence of these address pulses causes Q_7 to conduct, biasing the gates of Q_5 and Q_6 on. The Y-address pulse must have a longer duration than the X-pulse in order that Q_5 and Q_6 be turned off by the transition of the X-address pulse back to +12 volts.

Writing is accomplished by applying a write pulse (Figure 2) to the cell simultaneously with the address pulses. With Q_1 assumed conducting, and Q_2 off, a logic "1" is stored in the cell. If a "O" is to be written into the cell, the write pulse, which rises from ground to +12 volts, is applied to the digit complement while the digit line is held at ground potential.. This causes Q_6 to conduct current into the capacitance at the Q_2 node, resulting in turn-off of Q_1. The Q_1 node is now discharged toward ground through Q5, resulting in the turn-on of Q_2. With termination of the address and write pulse signals, the bistable is then retained in the zero state.

From the collection of the Computer History Museum (www.computerhistory.org)

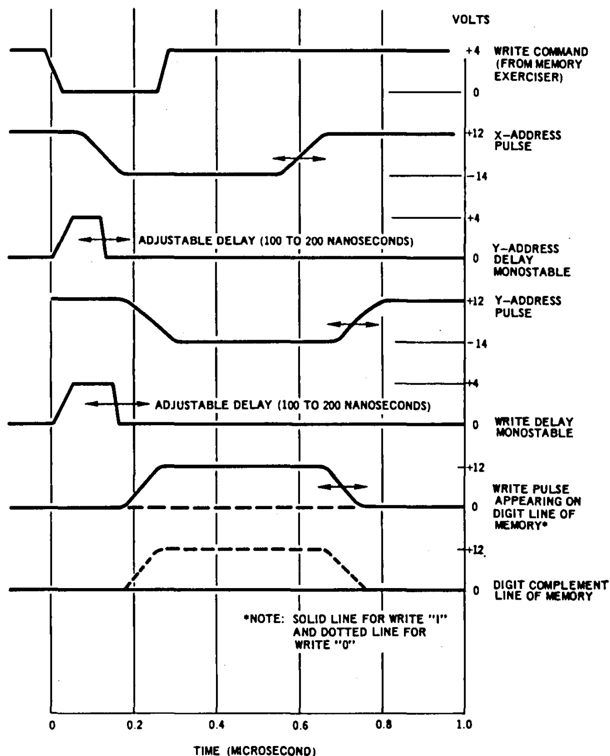

Figure 2—Memory timing diagram for write cycle

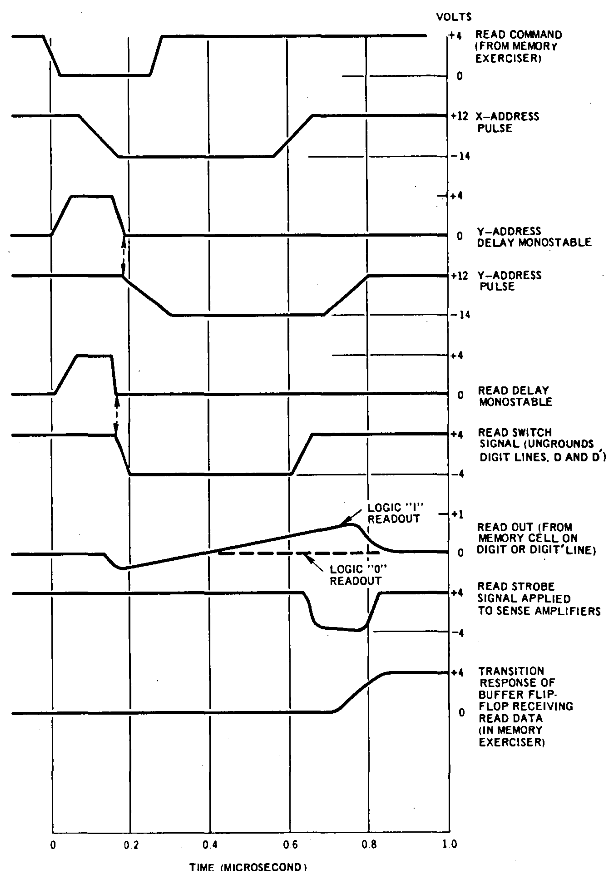

Figure 3—Memory timing diagram for read cycle

Readout of the cell (Figure 3) is accomplished by addressing the cell in the same manner as for writing. Both the digit and digit complement lines are near ground potential for read. When addressed, current from the "on" transistor in the bistable will flow into the respective digit or digit complement line. The node of the "off" transistor will be near ground potential ahd, hence, will not cause a current flow into the corresponding digit line. A differential amplifier, connected to the D and $\overline{D}$ digit lines, senses the readout.

The monolithic chip, shown in the microphotograph of Figure 4, measures 80 mils by 100 mils and contains 64 memory cells interconnected to form one bit of 64 different words. One memory cell occupies an area of approximately 100 $mils^2$ on the chip. The chips were individually housed in 22-lead flatpacks. The monolithic parameters of importance, external to the flatpack, are the capacitances looking into the restore, the digit and the address terminals. These capacitance values are shown in Table I based on averages of several flatpack samples.

Table I

Capacitance to Substrate Measured on 64-cell Monoliths Housed in 22-Lead Flatpacks

Measured Point (to substrate)	Typical Capacitance (picofarads)
D line	17.08
D' line	16.75
X line	7.00
Y line	7.48
R	26.5

System description

Figure 5 shows the basic organization of the system. The memory stack is comprised of the MOS transistor arrays described in the preceding section, while the peripheral circuitry is composed of the address decoder, write drivers, sense amplifiers and timing circuits. In the case of this feasibility model, the memory

From the collection of the Computer History Museum (www.computerhistory.org)

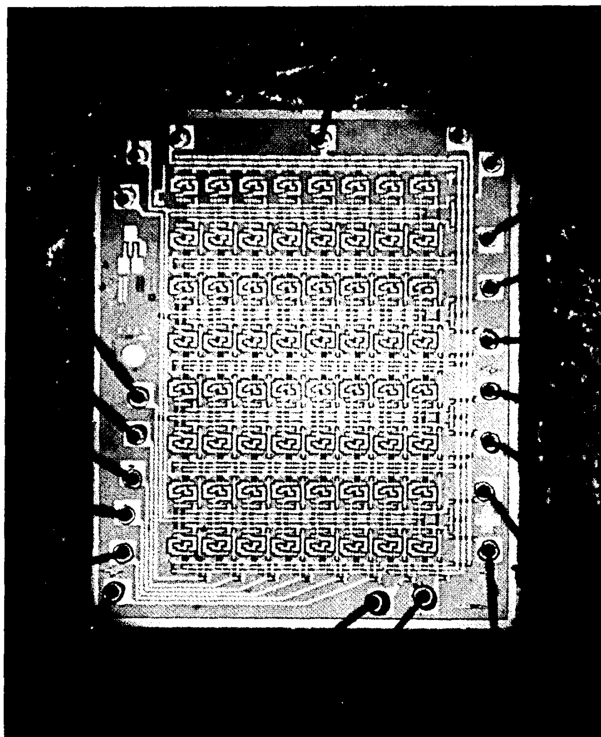

Figure 4—Monolith arrangement of 64 memory cell circuits

exerciser system (not shown) was built into the system to serve as the data processor. The exerciser supplies write information, and accepts read information from the memory, and also has the capability of detecting and locating any digital errors created in the memory system in terms of the word address and specific bits in the word.

As shown in Figure 5, the 10-bit address input is divided into a 5-bit section for decoding in 32 X-lines and a similar 5-bit section for the Y lines. A common pair of digit lines (Data and $\overline{\text{Data}}$) connects all 1024 memory cells corresponding to a specified bit in each word. These digit lines are multiplexed to carry both the write and read information to and from the memory.

Each 64-cell chip in Figure 5 is labeled with the address line designations which it receives. The columns represent the bit or digit positions, while each row represents a 64-word group.

Bipolar transistors were found advantageous for driving the large capacitances associated with the memory arrays. These bipolar peripheral circuits are individually described in the following paragraphs.

Address decoder

The 20 address input lines (10 digits plus complements) supplied from an address register (contained in he exerciser unit) represent a capability for addressing the 1024 words in the memory. Decoding is achieved in a conventional manner for a two coordinate addressing system. Figure 6 presents a block diagram of one of the two identical X or Y decoders. Three of the five digit inputs are decoded by eight NAND gates, labeled 0 through 7, while the remaining two digits are decoded by four NAND gates (A through D). These respective NAND gates drive pulse activated transformer circuits which in turn drive a 4 $\times$ 8 matrix of transformer circuits for the 32 output lines.

Schematics associated with the various blocks are shown by inserts in Figure 6. Commercial integrated circuits were employed for the NAND gates. One high level matrix driver and one low level matrix driver is actuated in correspondence with the address code when the address command pulse occurs. The two actuated circuits in turn select one of the 32 identical address driver circuits in the address matrix corresponding to the circuit path from the high level driver through the transformer winding to the ground return provided by the low level driver. Operation of the address driver circuit is discussed in a subsequent section dealing with power conservation.

Write drivers

Sixty input data lines (30 digits plus complements) representing the information to be written into the memory, are supplied from the memory exerciser. Each input line is supplied to a write driver circuit. Upon receipt of the write command, the lines having a logic "1" state activate the respective write driver for either the D or D digit lines.

The schematic for the write driver is shown in Figure 7. The digit line associated with each write driver is grounded through R_8 and Q4 during quiescent operation. When writing, Q_2 is turned on if the transformer is energized, while Q_3 and Q_4 are driven off thus allowing the twelve volt write pulse to be developed on the digit line. When the pulse terminates, Q_2 is driven off by the reversal of voltage across the secondary windings due to "flyback" action[6] of the transformer, while Q_3 and Q_4 are driven on. This rapidly discharges the digit line capacitanct and clamps the digit line back to ground.

Read sense amplifiers

Similar to the write drivers, each of the 30 sense amplifiers connects to the main digit (and digit complement) line which ties together all the individual digit lines from the 16 flatpacks corresponding to a

From the collection of the Computer History Museum (www.computerhistory.org)

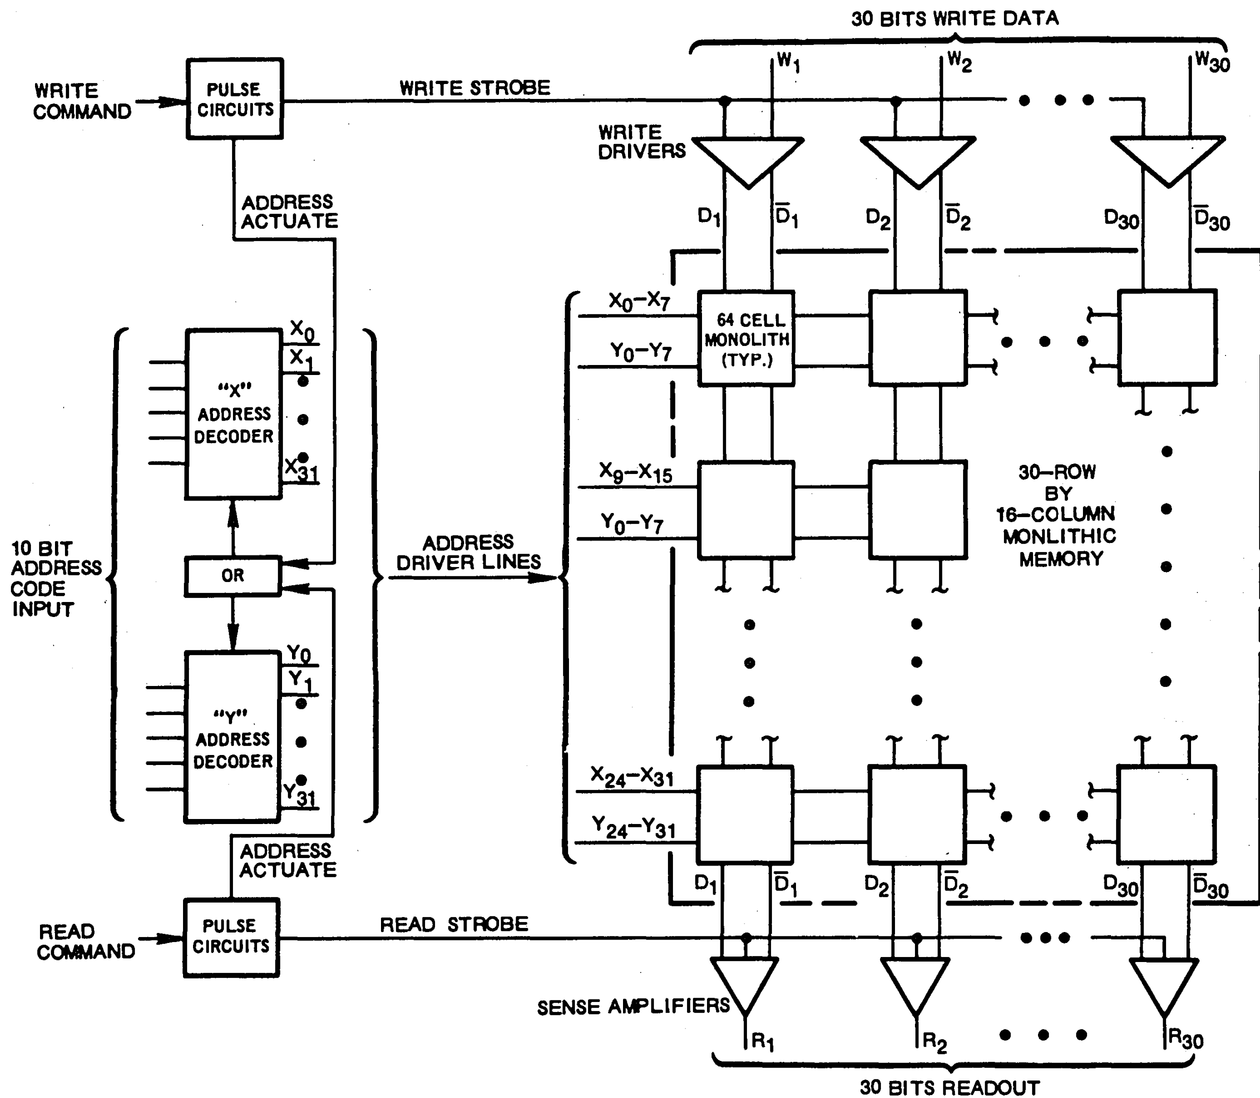

Figure 5—Block diagram of memory system

specified bit position. When a read command is received, the memory cells of the addressed word supply current to the digit lines. The differential between the digit and complement line signals are sensed to produce the corresponding read output when the strobe pulse is applied by the read strobe generator circuit.

The real sense amplifier schematic is shown in Figure 8. In consists of a differential amplifier (Q_5 and Q_6) which "primes" a flip-flop comprised of transistors Q_1 and Q_2. When the negative strobe signal is applied to the flip-flop it is actuated to the state determined by the differential amplifier. The sensitivity of the sense amplifier is improved by isolating the flip-flops from a possible unbalanced output loading by employing buffer transistors Q_3 and Q_4.

Restore circuits

The restore pulse, as previously described, is applied periodically to each cell in the memory to discharge the voltage buildup, due to leakage current, on the intrinsic capacitance at the flip-flop node.

To prevent large surges of current that would be required if the restore signal were applied to the entire memory simultaneously, the memory is divided into eight sections. Figure 9, showing the restore circuitry, consists of a 3-stage counter and associated binary-

From the collection of the Computer History Museum (www.computerhistory.org)

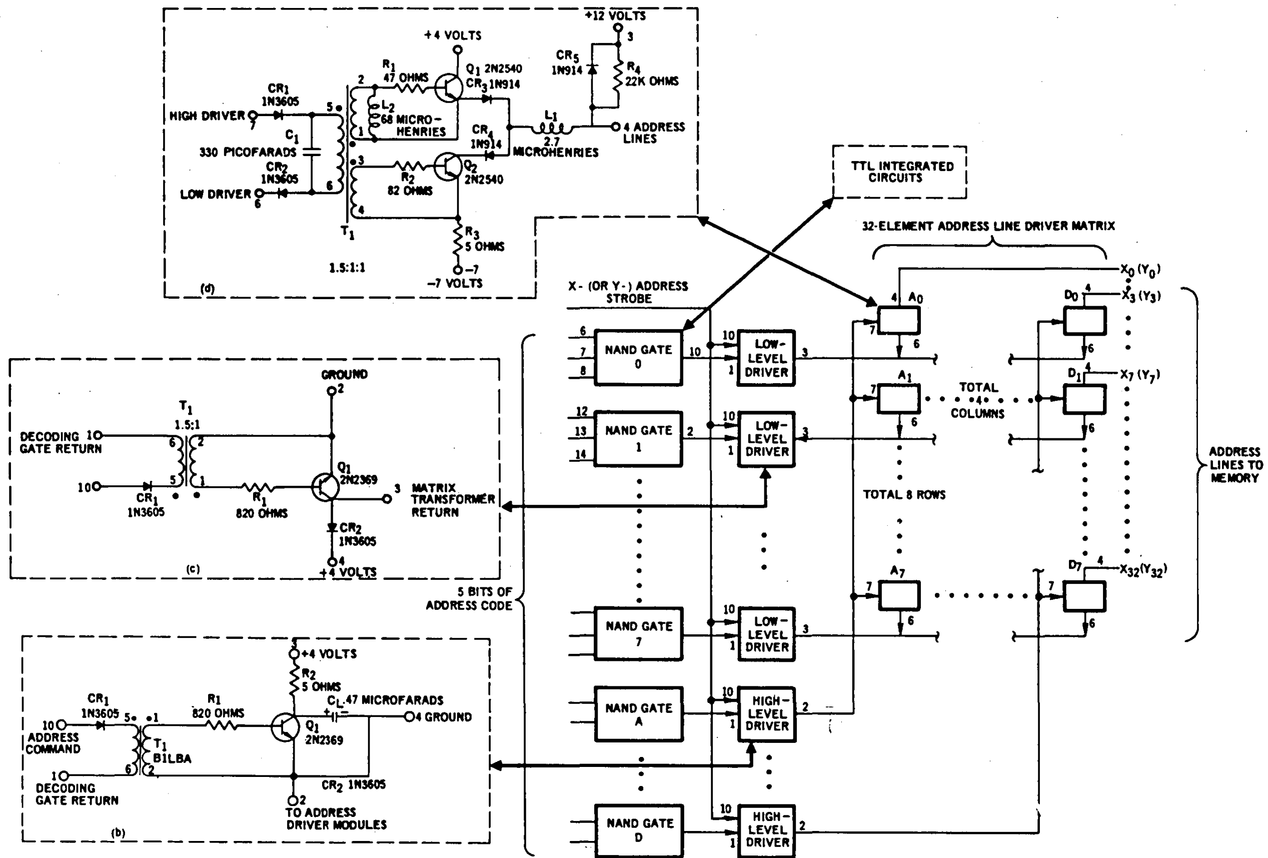

Figure 6—Address decoder block diagram with circuit schematics

to-octal decoder, with each output of the decoder controlling a respective restore pulse generator. A self-contained restore clock, asynchronous with the rest of the system, generates the shift pulses for the counter. Each time a restore clock pulse is generated, the pulse gate corresponding to the actuated line of the octal decoder delivers a pulse to its respective restore driver. The restore driver schematic is shown in the insert of Figure 9.

Timing control circuits

The control circuits cause the memory to be addressed whenever a read or a write command is received, and they generate the corresponding activating signals for the sense amplifier or the write driver circuits. Figure 10 shows a block diagram of the control circuits, with integrated monostables used for timing control. The schematic for the address command pulse driver, shown in Figure 11 is typical of the output driver circuits in Figure 10. Signal timing diagrams for the write and read cycles, respectively, are shown in Figures 2 and 3.

Physical arrangement

The physical packaging arrangement utilizes four multilayer printed circuit boards — two for the memory stack and two for the peripheral circuitry. The X-address decoder and the restore generator are packaged on one of the periphral circuit boards. The Y-address decoder and the control pulse circuitry are packaged on the second peripheral circuit board. No attempt was made to achieve a high degree of miniturization in this developmental effort.

The memory array consists of 480 identical flat-packs arranged in a 16 by 30 matrix and packaged on two multilayer printed circuit boards. One of these boards with the full complement of 120 flat-packs is shown in Figure 12.

The memory, exerciser, and power supplies are packaged in a Samsonite suitcase with exterior dimen-

From the collection of the Computer History Museum (www.computerhistory.org)

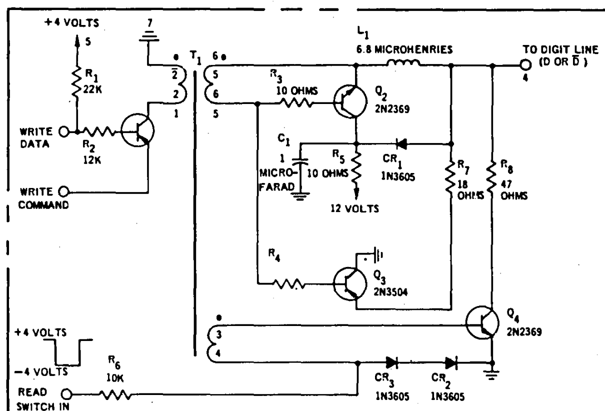

Figure 7—Write driver schematic

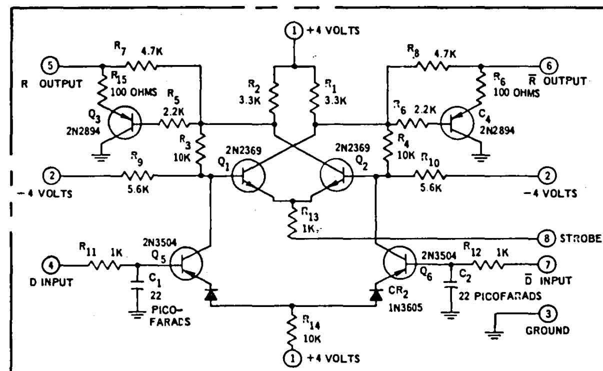

Figure 8—Read sense amplifier schematic

sions of 6 by 15 by 21 inches. The objective of this packaging configuration was to make the feasibility model of the MOS memory completely portable for demonstration purposes.

Power reduction techniques

The primary means of conserving power in the MOS memory cells was the use of the previously discussed restore technique. The power required per memory cell then becomes a function of the restore duty cycle. Using a 10 KHz restore rate resulted in a quiescent dissipation of 30.3 microwatts per cell, thus the 30,720 cells for the memory consume 930 milliwatts.

Table II

Capacitance Loads Presented to Peripheral Drive Circuitry

Circuit	Measured capacitive component (picofarads)
X-address line (120 flatpacks)	1,320
Y-address line (120 flatpacks)	1,170
Digit line, D (16 flatpacks)	349
Digit complement line, D′ (16 flatpacks)	380
Restore line (60 flatpacks)	3,500

Relatively large capacitance loading is presented to the address, write driver and restore cricuits by the memory as shown in Table II. The power consumed in charging and discharging the line capacitance using a resistance driver source is given by

$$P = CV^2 \text{ x (p.r.f.)}$$

where C is the effective capacitance of the memory line, V is the peak pulse voltage and p.r.f. is the pulse repetition frequency at which the capacitance is charged and discharged. For the 1200-picofarad address line capacitance and a 26 volt address pulse, each of the two address drivers would dissipate 0.8 watts, requiring 1.6 watts for address line driving if a resistive source impedance had been used in the line driver design.

The method used for reducing power dissipation is to make the line driver have an inductive source impedance, so that a resonant circuit results, and develops the desired voltage amplitude on the line capacitance in conjunction with nonlinear switching techniques which preverts "ringing" by the resonant circuit. The principle involved is illustrated by the simplified schematic and waveform sketch in Figure 13, where ideal components are assumed. The capacitor in Figure 13a is assumed initially charged to $V_c(O)$. When S_1 is closed at $t = O$, the voltage across the capacitor is given by:

$$V_o = V_1 + [V_c(O) - V_1] \cos \omega_o t$$

where: $$\omega_o = \frac{1}{\sqrt{LC}}$$ and

V_o = output voltage across the capacitor

When $t = \frac{\pi}{\omega_o}$, the voltage, V_o, in the equation would reach a peak value of $2V_1 - V_c(O)$. For example, if $V_1 = -7V$ and $V_c(O) = 12V$, the waveform sketch in Figure 14b would reach a negative peak voltage of —26 volts, as shown in the waveform sketch of Figure 13b. If S_1 were opened at $t_1 = \frac{\pi}{\omega_o}$ the capacitance

From the collection of the Computer History Museum (www.computerhistory.org)

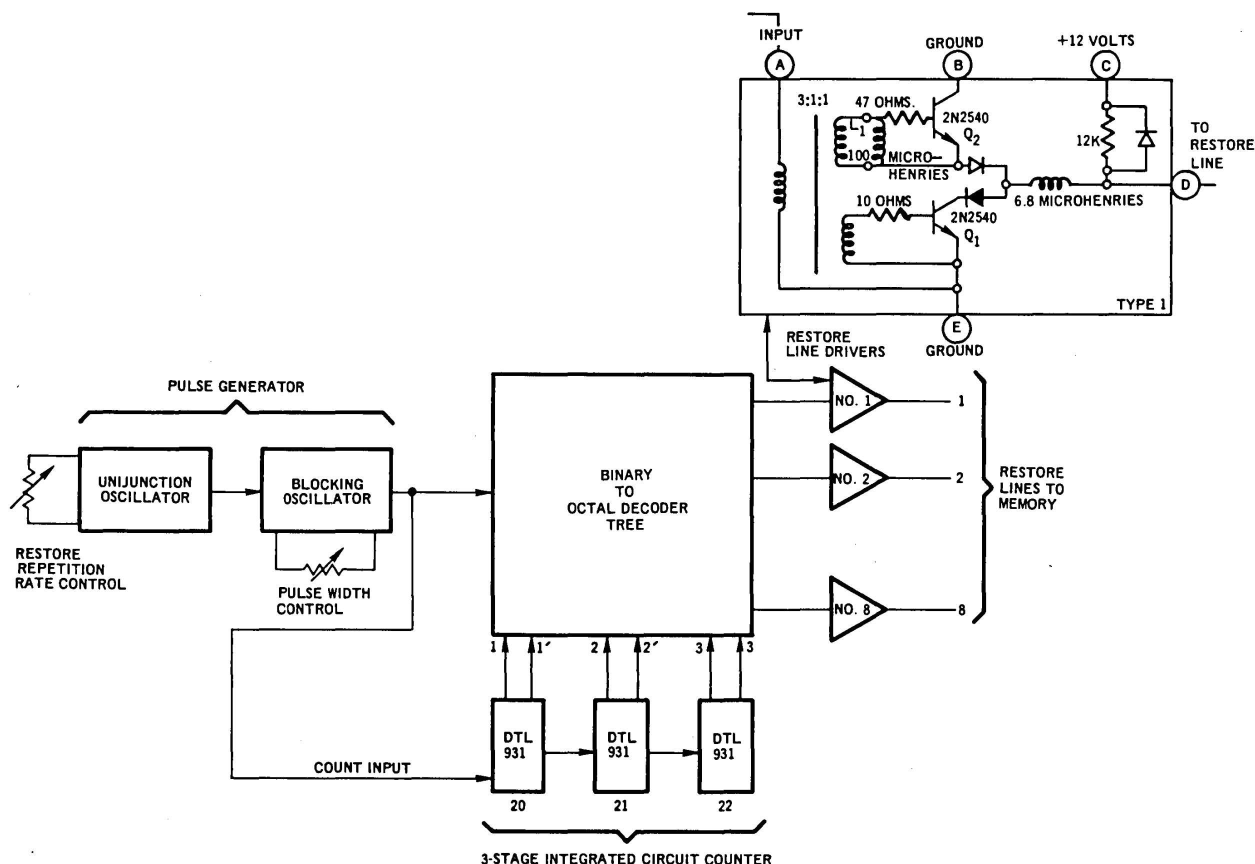

Figure 9—Restore generators using common pulse generator to supply restore line drivers

would remain charged to this value, with no power having been dissipated. If at time t_2, S_2 is closed while S_1 is opened, then the voltage on the capacitor is given by:

$$V_o = V_2 + [V_c(t_2)—V_2] \cos \omega_o t$$

where $V_c(t_2)$ is the voltage on the capacitor prior to closing switch S_2. In the lossless situation under consideration V_2 would equal V_1 in order that the trailing edge from t_2 to t_3 reach a peak value equal to the initial voltage on $V_c(O)$ prior to development of the pulse. With the ideal components that have been assumed the pulse has been developed without power dissipation. When component losses and voltage drops across semiconductor switches are considered, V_1 must be more negative than V_2 in order to accommodate the losses and develop the form of pulse shown in Figure 13b.

Practical application of this principle is illustrated in the case of the address driver by reference to the schematic insert shown in Figure 6. When the primary of the transformer is energized, Q_2 turns on as a saturated switch connecting the —7 volt supply through inductance L_1 to the capacitance of the address line. When the output voltage reaches the negative peak, diode CR_4 prevents reversal of current thus effectively disconnecting the output from Q_2. At the termination of the energizing pulse a high impedance is presented to the primary of the transformer and Q_1 is turned on by the voltage flyback action of the energy stored in the transformer primary[6] and in L_2. This causes Q_1 to turn on and charge the capacitor back to the initial voltage level. Resistor R_4 acts as a bleeder resistor to initially charge the capacitance to +12 volts. When circuits losses and semiconductor voltage drops are considered, the circuit parameters shown resulted in a 26 volt negative address pulse. With $L_1 = 2.7$ microhenries and $C = 1500$ picofarads, ω_o is found to be 15.8×10^6 radians per second, and the rise and fall time of the address waveforms is given by $t = \frac{\pi}{\omega_o}$ which is found to be approximately 0.2 microseconds. Typical X-Y address pulses developed by this circuit are shown in Figure 14.

This technique resulted in a considerable power savings in driving large capacitive loads. For example,

From the collection of the Computer History Museum (www.computerhistory.org)

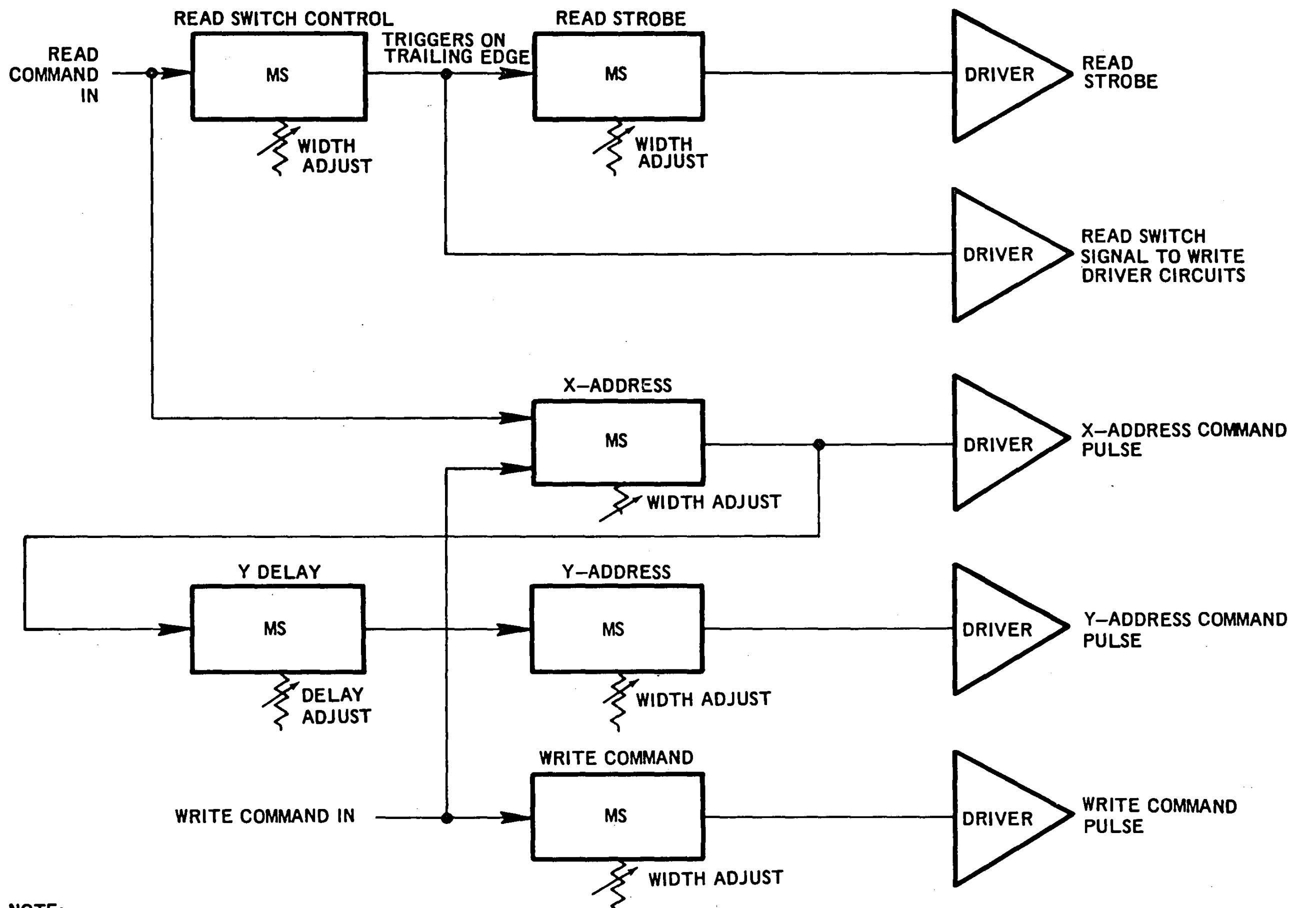

Figure 10—Pulse generator control circuitry block diagram

measured power for the two address driver circuits was about .7 watts, at a 1 mc addressing rate, whereas a resistive driver source would have consumed 1.6 watts in driving the line capacitance. The technique was successfully applied to the address and restore drivers; however, it was not applicable to the write drivers where power consumption due to transient charging of line capacitance was relatively high.

In the sense amplifiers discrete components were employed primarily because of suitable integrated form of low power sense amplifier was not available. Power conservation in the sense amplifier is achieved by use of high-resistance loads and a strobing technique previously described in which very little power is consumed in the circuit during quiescent (nonstrobed) operation.

Power requirements

The power objective for the memory proper was 2.5 watts for a 30% write-to-read ratio at 1 mc clock operation and for operation over a 0°C to 60°C temperature range. This power objective includes the memory stack and all of the peripheral circuitry except the address decoder. Power dissipation in the power supply and memory exerciser are not included in this figure. With 1.0 watts required by the address decoder for continuous addressing, the total power for the memory was to be less then 3.5 watts. Power actually measured was 2.4 watts, excluding the address decoders. With address decoders included, power was measured to be 3.4 watts. Of this the memory array itself consumes slightly less than 1 watt. With a con-

From the collection of the Computer History Museum (www.computerhistory.org)

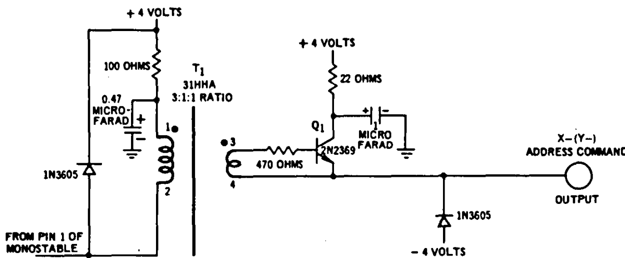

Figure 11—Address command pulse (X or Y) driver schematic

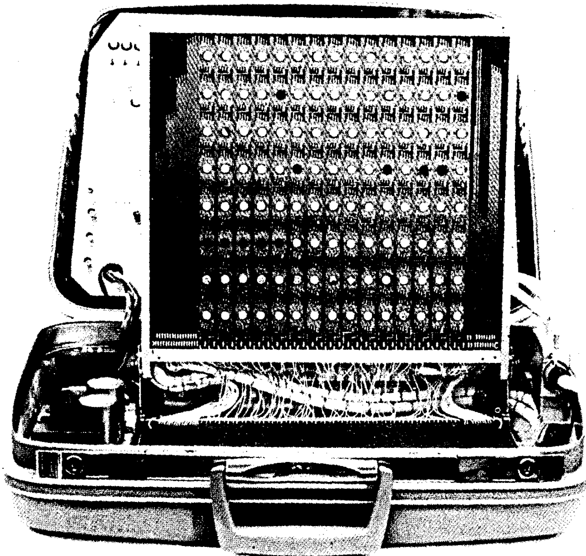

Figure 12—Suitcase configurations of memory and exerciser

certed effort to further reduce power dissipation, by design changes in the memory array and peripheral circuitry, it is estimated that the total power including the address decoding could be reduced below 2.5 watts.

Table III is a summary of the power dissipation in the various circuit sybsystems for 1 megacycle operation. The 30% write-to-read ratio results in the previously quoted power requirement of 3.4 watts.

Performance and evaluation

The entire system was made operational with the speed and power objectives achieved; namely (1) a two microsecond write-read cycle time, and (2) A total power consumption objective of 3.5 watts based on continuous 1 megacycle operation with a 30% write-to-read ratio.

Table III

Power Requirements for Memory Subsystems

	Write Continuous Operation (milliwatts)	Read Continuous Operation (milliwatts)	Standby (milliwatts)
(1) Monolithic memory cells	930	960	930
(2) Write driver circuits	2250	88	88
(3) Read sense amplifiers	120	450	120
(4) Control pulse circuits	240	240	210
(5) Restore generator	112	112	112
(6) Address driver decode circuits	1009	1009	480
Total power:	4661	2859	1940

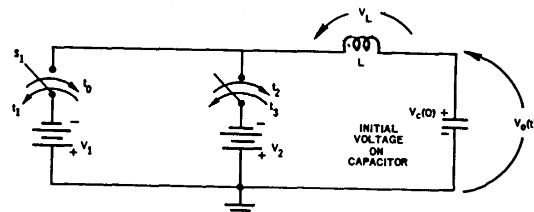

Figure 13a—Simplified schematic for line driver circuit

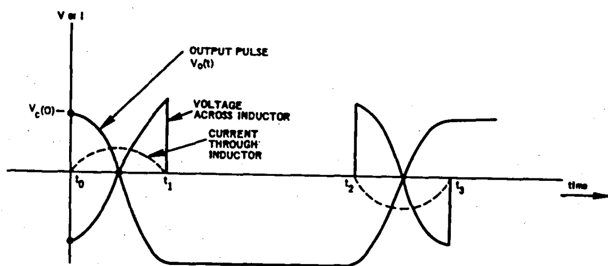

Figure 13b—Wave forms associated with operation of circuit in Figure 13a

In the initial design of the memory cell, a 5 KHz restore rate was estimated based on leakage measurements on sample units of MOS transistors. In order to accommodate the leakage variations in the 480 chips used in the memory, a 10 kc restore rate was found to be necessary. Although the entire system was not operated at 60°C, sample portions of the memory

From the collection of the Computer History Museum (www.computerhistory.org)

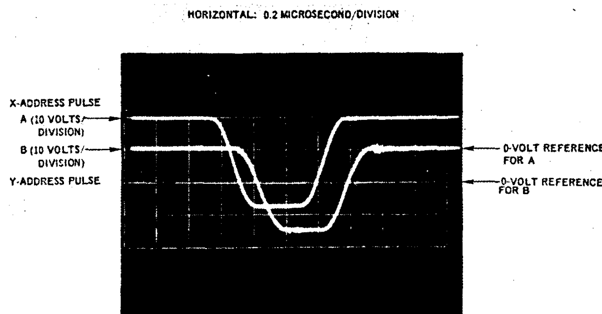

Figure 14—Address waveforms on X- and Y-address lines of memory

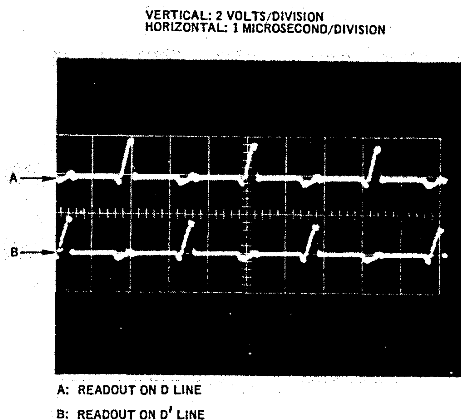

Figure 15—Readout with single flatpack connected to digit line

array were successfully tested for 60°C operation using the 10 kc restore rate.

Of primary importance to the peripheral driving circuitry are the line capacitances of the composite memory array. The measured line capacitances are shown in Table II. Various problem areas involving signal generation and distribution to the memory were anticipated and further investigated when the system was completed. Some of these are discussed in subsequent paragraphs.

To determine the effect of the increased capacitance and crosstalk on the digit line as increasing numbers of flatpacks are multiplexed together, readout waveforms were observed as flatpacks were added to one of the digit lines. Figure 15 shows the readout obtained with one flatpack (64 words); contrasted to this, Figure 16 shows the resultant readout with the full complement of 16 flatpacks, representing the 1024 words connected to the digit line.

Figure 17 shows the gross readout signal envelope over the scan of the 1024 words in the memory as

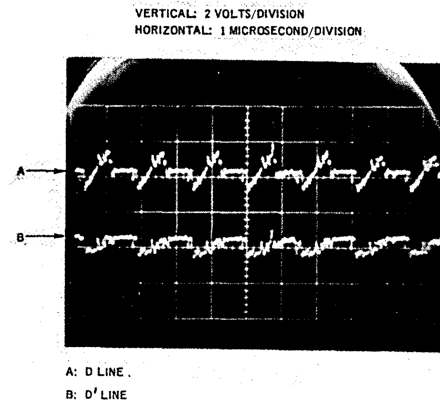

Figure 16—Read with full complement of 16 flatpacks connected

observed on one of the digit line pairs. A single logic "1" readout is shown on the $\overline{D}$ line for sake of comparison with the logic "O" level. The variation of readout amplitude over individual cells is vividly portrayed in this waveform envelope.

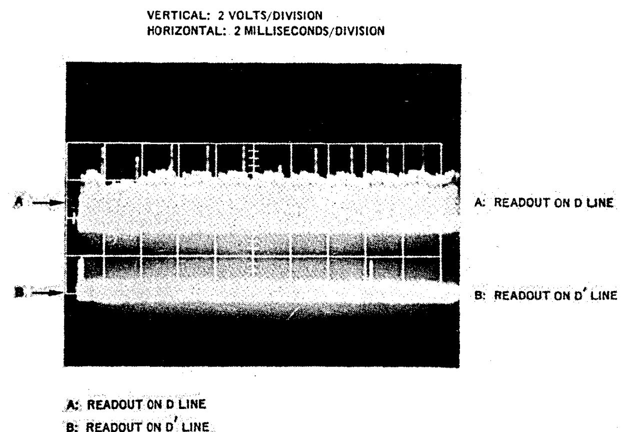

Figure 17—Comparison of readout envelope for the 1024 bits on the D and D′ lines

One concern in the readout operation was whether the crosstalk from the asynchronous restore signal on the flatpacks would be of sufficient magnitude to cause errors in the readout information. Crosstalk was found to result principally from parasitic capacitance coupling with the pulses utilized in the system. Address pulse crosstalk did not present a problem because these are synchronous signals and the readout strobe could be timed to minimize crosstalk from these signals. The

From the collection of the Computer History Museum (www.computerhistory.org)

restore signal, however, being asynchronous, can present a "worst-case" situation when it occurs within the "window" of the read strobe. Fortunately, the use of a differential form of sense amplifier has the advantage of making only the difference in crosstalk between the D and D digit lines of importance. Because of a high degree of symmetry in the monolith and the printed board layout, this differential has been minimized. Figure 16, in the 4th waveform from the left, shows the effect of the restore crosstalk perturbation on the D and $\overline{D}$ readout lines with restore occurring at the peak of the readout signal.

To verify that recovery of the digit line after the 12 volt write pulse would be fast enough to allow correct readout response on a following readout cycle, the "write then read" mode was utilized while monitoring the waveforms appearing on the digit lines. Figure 18 shows typical write pulses followed by readout. In this photograph the write cycle is followed by three read cycles in order to make a comparison of succeeding readouts. It was verified that recovery was sufficient to allow a readout cycle to immediately follow a write cycle at a 1 MHz clock rate.

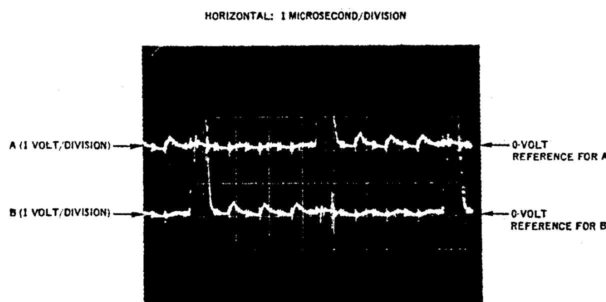

Figure 18—Digit line waveforms for write followed by read operations

A total of 480 memory chips was employed in the system, and thus represents a fairly sizable system using large-scale integration concepts. Although no quantitative reliability study was made as part of this program, after the debugging phases, the memory was operated over a long period of time without failures. However, the debugging phase was a difficult one primarily because of a number of intermittencies in several of the chips, which showed up very infrequently, and were not detected in the unit flatpack testing. The conclusion which can be drawn from this experience is that the unit test on LSI chips should be quite exhaustive if difficulties are to be avoided after the chips are installed in the system. Suitable testing techniques for LSI is presently one of the difficulties faced by the industry.

CONCLUSIONS

The low power MOS memory described in this paper represents a full scale implementation of a complete small-scale memory, not just a partially populated model. The approach has proven the feasibility of achieving low power through the use of pulsed operation of the memory array and peripheral circuits and by use of circuits capable of driving capacitive loads with reduced power dissipation.

The use of the restore technique to conserve power depends on maintaining control of leakage current, and is therefore susceptible to high temperatures and nuclear radiation effects. The present memory was designed for operating environments up to 60°C, and the effects of nuclear radiation was not considered in this program. Increased restore rates can be used for higher temperature and to accommodate leakage current increases by nuclear radiation but at the expense of greater power consumption.

In Table IV power requirements of the MOS memory are compared with those estimated for a number of other memory technologies assuming a 2048 word, 30-bits per word memory size. It is anticipated that power reduction improvements could be made in the MOS as well as other memory technologies by continued design effort with low power objectives.

Table IV

Comparison of Memory Technologies for Aerospace Applications*

	Speed (microseconds)	Power (watts)
MOS arrays	1 (NDRO), 2 (write)	3 to 5
Magnetic thin films	0.75 (DRO)	12 to 16
Plated wires	1.5 (NDRO)	10 to 15
Permalloy Sheet	2 (DRO)	7 to 10
Biax	1 (NDRO), 3 (write)	20 to 25
Magnetic cores	2 (NDRO)	15 to 20
	2 (DRO)	25 to 35

*2048 words, 30 bits per word; equipment to meet MIL-E-5400 and to include complete electronics and power supply. 1969-1970 production assumed.
DRO = destructive readout; NDRO = non-destructive readout; NA = not available

The MOS memory in word sizes up to about 8000 words, does appear practical from the economic viewpoint in that the cost per bit is competitive with other memory technologies.

From the collection of the Computer History Museum (www.computerhistory.org)

ACKNOWLEDGMENT

Grateful acknowledgment is made to the USAF Avionics Laboratory for their sponsorship of the Low Power Computer Memory Project.

The authors are indebted to R. Cole and R. Feuer for their contributions while at Bunker-Ramo, and to T. Kitaguchi for his role in low power circuit design. Walt Wittler's effort in fabrication and testing of the memory is greatly appreciated.

Finally, the authors wish to express their appreciation and acknowledge the efforts of R. Przybylski and D. Farina at the Philco-Ford Microelectronics Division for their role in the successful diffusion of the MOS arrays.

REFERENCES

1 R IGAROSHI T YAITA
An integrated MOS transistor associative memory system with 100 *nanosecond cycle time*
AFIPS Proceedings Spring Joint Computer Conference 1967

2 J R BURNS *et al.*
Integrated memory using complementary field effect transistors
ISSCC Digest of Technical Papers February 1966

3 I CATT E GARTH MURRAY
A high speed integrated circuit scratch pad memory
AFIPS Proceedings Fall Joint Computer Conference 1966

4 H PERKINS J SCHMIDT
An integrated semiconductor memory system
AFIPS Proceedings Fall Joint Computer Conference 1965

5 G PODRAZA S NISSIM D BREWER
A low power MOS monolithic memory array
Presented at the 1st annual IEEE Computer Conference September 1967

6 R LEE
Electronic transformers and circuits
Wiley & Sons New York 1955 pp 298-302

From the collection of the Computer History Museum (www.computerhistory.org)

From the collection of the Computer History Museum (www.computerhistory.org)