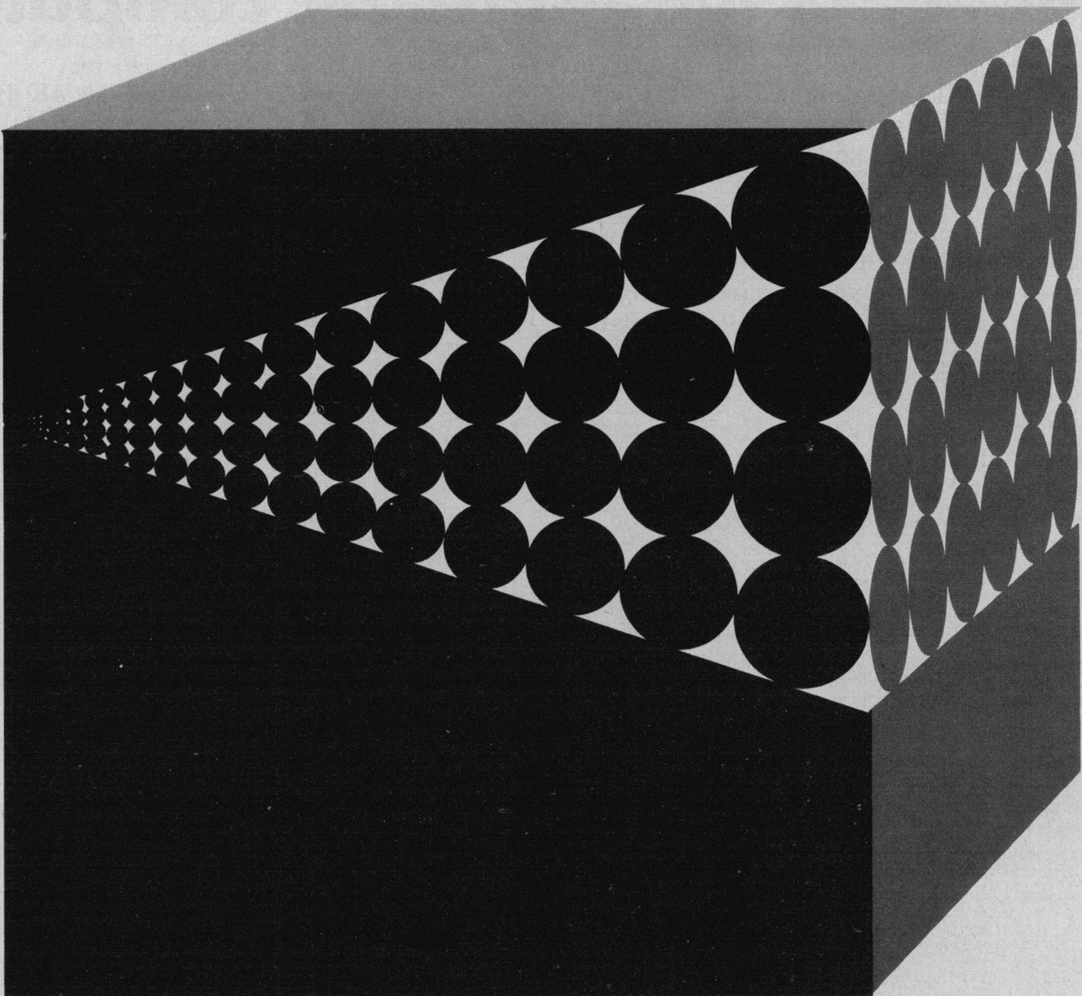

The Development of an Experimental Electron-Beam-Addressed Memory Module

John Kelly
Stanford Research Institute

Introduction

The various technologies needed to build a successful electron-beam-addressed memory (EBAM) have been evolving for over a decade, and are now sufficiently established to build a viable, economical system. Recent work at SRI indicates that EBAM will be an important contender for low-cost, large memory systems with improved random-access capability, high reliability, and high data rates. This will be achieved in small volume and without the use of moving parts. The potential applications of such a mass memory system are widespread, ranging from direct drum and disk replacement to new systems with architectures that would take advantage of improved random access times and increased data rates.

The major attraction of the EBAM concept lies in the ability to produce relatively large capacity, fast access units at low cost. For example, to build a memory of 10^9 to 10^{10} bits, it is desirable to have a module of the order of 10^6 to 10^7 bits, since otherwise the problems of interconnection and packaging become formidable. Neither is it desirable to have each module made up of too many smaller units, again a source of unreliability and expense.

In the case of the EBAM, the tube is the module, which in first production systems is likely to be of the order of 4×10^6 bits with an access time of several microseconds and a system cost of 0.02 cents per bit. Further developments will undoubtedly lead to even greater capacities and sub-microsecond access times.

It is therefore considered that the EBAM is not competing head-on with semiconductor devices, including charge coupled devices, or magnetic devices such as magnetic bubbles, drums, and disks. None of these combine all the desired qualities—high bit capacity, fast access time,

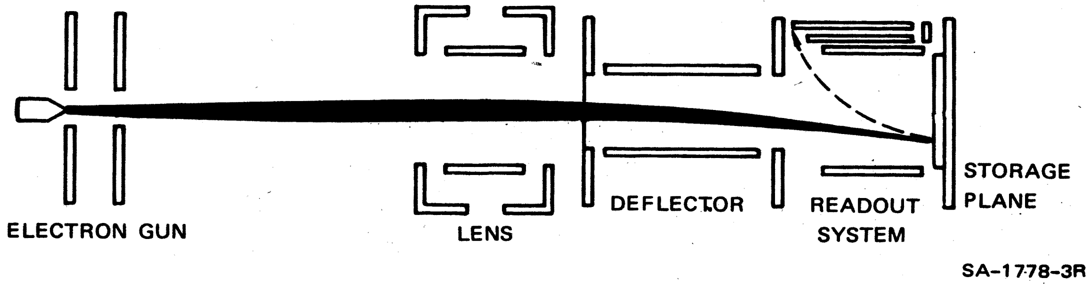

Figure 1. Major Components of EBAM Tube

low cost per bit, wide versatility of data format, low power consumption, small volume, and simplicity of interconnects—that are uniquely achieved in the EBAM. However, any particular specification may be equaled or exceeded by one or the other of these technologies. For example, semiconductor memories are relatively low cost and low volume, and can have high speed and low power consumption. Attempts are being made to increase the capacity. This implies making the current chip sizes larger than at present, or putting more circuitry on the existing area. Eight-kilobit devices are available and larger capacities will appear with time. However, as circuit complexity increases, there will be tremendous problems in achieving sufficiently high yields to keep the costs down. Even with the acceptance of electron beam lithography and refined processing techniques, there are tremendous barriers to be surmounted.

The principal components and operation of an EBAM tube are illustrated in Figure 1. The components are similar to those found in a cathode ray tube. A beam of electrons is generated by an electron gun and focused by means of an electron lens onto the storage target where the beam can cause some local physical change. The bit size corresponds approximately to that of the size of the electron beam at the storage target. Addressing is achieved by deflecting the beam to different locations in the target plane. In the illustration, both the lens and the deflector are electrostatic. Either random access or block access can be achieved, depending on the manner in which the target is scanned. Thus the beam can be used to interrogate individual locations on the target surface (bit access), or it can be scanned across the surface in a TV-like scan to read blocks of data (block access).

In the SRI system, information is stored by a charge deposited by the electron beam on the surface of an electrostatic storage target. In readout, the presence of a stored charge—or lack of it—is detected by analysis of the velocities of the secondary electrons created by the primary beam as they return from the target surface. The energy analyzer permits electrons from ONE-state bits to pass, thereby producing a current pulse at the tube output. To illustrate the general capability of the EBAM concept, a specification for an early production system based on the use of 16 tubes in parallel is given below (see Table 1). This system specification is given as an example of the memory capability that can be produced without significant further research. Naturally higher capacities can be achieved by using larger numbers of tubes in different configurations. The cost of these first systems is estimated to be 0.02¢ per bit for a commercial system. Figure 2 shows the general

Table 1. Sample EBAM System Specification

System capacity:	6.4×10^7 bits
Configuration:	16 tubes of 4×10^6 bits operating in parallel. Block oriented, 1000 $\times$ 16 bit words or less.
Parity:	As required
Access time:	3 μs to any block
Writing speed:	1 Mbit (per channel) 16 Mbits (per system)
Reading speed:	10 Mbit (per channel) 160 Mbits (per system)
Readout:	Nondestructive*
Refresh:	Periodic (every 1 minute)
Power off volatility:	8 hours (no flood beam) Indefinite (with flood beam)
Power consumption:	500 W (7.8 μW per bit)
Volume:	3 ft^3
Weight:	$<$100 lb
Tube life:	$>$2 years

*Information is only partially degraded by the reading process, and full storage level can be restored by refresh.

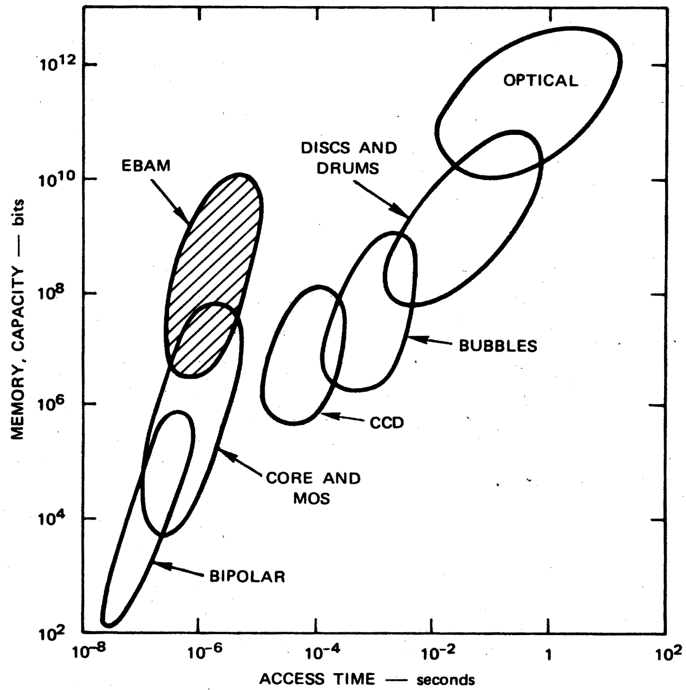

Figure 2. Capacity and Access Time Comparison with Other Memory Technologies

performance of an EBAM memory compared with other memory technologies. In this figure, memory capacity is plotted against access time. In general, this particular type of graph can be misleading, since it is not possible to compare all the types of memory on a bit-by-bit or module-by-module basis without producing a very large number of charts. In this chart each zone indicates the full range of chips and modules. The chart illustrates well that the EBAM technology offers memories of disk or drum-like capacities with essentially the access time of mainframe memory. EBAM systems are not yet available as production items. However, rapid progress is being made at GE, at Micro-Bit, and at SRI. In early 1974, SRI delivered an experimental EBAM module to the USAF Avionics Laboratory at Wright-Patterson AFB.[1,2] Although having somewhat limited specifications, this equipment constitutes a milestone in EBAM system development.

Technology

Electronics An EBAM relies heavily on the availability of highly stable electronic components. Such components have only recently become available. For example, it is now possible to buy the required high-voltage power supplies with stabilities of better than one part in 10^4 for less than \$150 (OEM). It is also possible to obtain sufficiently precise digital-to-analog converters with fast settling times. These are the most expensive components in the system. Specifications are still being improved and costs will be reduced further.

Electron Optics Many changes have occurred in undramatic but highly significant ways. The art of gun and lens design has improved substantially with the use of computers. The factors necessary to achieve reliable, stable performance are now well understood. This knowledge has been gained in the tube industry and from the development of instruments such as the scanning electron microscope. In addition, there have been major improvements in deflector design. For example, an electrostatic octupole deflector was developed by the author under a program sponsored by the U.S. Army in 1969.[3] This deflector enables significantly lower aberrations to be obtained across the deflection field. At the same time good sensitivity and a common center of deflection are achieved.

Cathode Life Dispenser cathodes with a proven life in excess of 10,000 hours and a potential of 50,000 hours are readily available. These cathodes have adequate brightness for the first generation of tubes. Historically, the cathode has been the limiting factor in tube operational life. The availability of these cathodes has changed this situation.

Silicon Technology The availability of a wide range of silicon technology (especially electron lithography and sophisticated etching techniques) has greatly contributed to this changing scene. Storage medium fabrication has shown significant advances.

Tube Technology Steady advances over the past several years in vacuum technology, in vacuum materials, and in assembly techniques ensure that an EBAM memory tube can be made economically and still be very rugged.

The Storage Target

Although all of the above technologies are extremely important, the selection of a suitable storage medium is the most important, since the storage target itself determines most of the operating parameters, including beam energy, current density, beam diameter, electronics requirements, configuration of the readout system, operating temperature range, error rate, and perhaps the life of the tube itself.

Many materials and operating schemes have been proposed for EBAMs; however, all of the groups known to be working on EBAMs appear to be using silicon-based targets and electrostatic charge storage of some form.[4,5] A sectional drawing of the SRI target is shown in Figure 3. It consists of a large MOS capacitor in which a closely spaced array of holes has been etched through the gate metal and partway through the oxide. Metal is then deposited in the bottom of these etched holes to form an array of microcapacitors that are electrically isolated both from the silicon substrate and from the gate metal, but are more closely coupled capacitatively to the silicon than to the gate. These targets have been made with a range of hole sizes and spacings. Quarter-inch targets with 6-μ (micron)

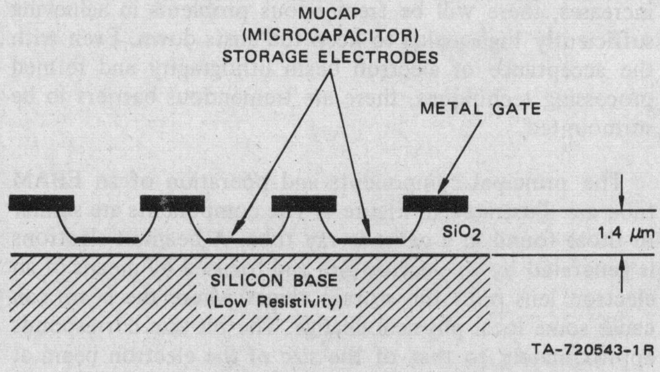

Figure 3. The Mucap (Microcapacitor) Storage Medium

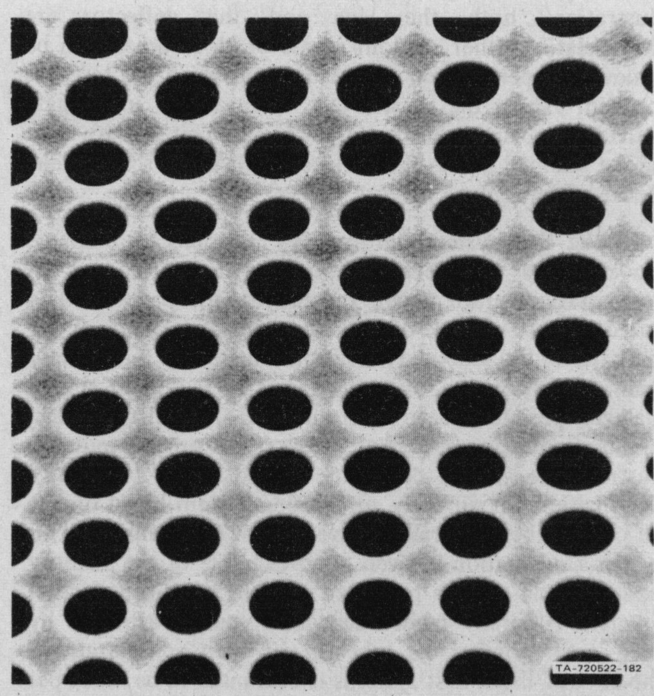

Figure 4. Scanning Electron Micrograph of Storage Target (6μ holes)

diameter holes on 9-μ centers are in regular use. Targets having 0.5-μ holes on 1.2-μ centers have been made over limited areas in the past. A part of a regular storage target is shown in Figure 4. These targets are made using electron lithography and special etching processes developed at SRI. The result is a high degree of uniformity and a low incidence of defects.

In some ways the ideal target would present a structureless surface to the electron beam so that an addressed location would be determined by the impact point of the deflected electron beam and not by the structure on the target. In consequence, and because it is not possible to make a completely perfect target, the storage medium is operated in a "quasi-continuous" mode. The microcapacitor dimensions and spacings are made small compared with the beam diameter, so that the illuminating beam covers several microcapacitors at each address location, a typical number of capacitors being nine in our present work. Information is stored on the microcapacitors as a charge; a negative charge corresponds to a stored ONE, and a zero charge corresponds to a ZERO. The basic reading and writing processes are best understood by reference to a single mucap element.

The reading operation is the easiest to understand and is illustrated in Figure 5. At the top is shown a ONE-state mucap. Note that the mucap is charged negatively to about −40 V and that the base is held at −50 V. Primary electrons striking the mucap surface generate secondary electrons, which are accelerated out of the mucap cavity through the stored potential drop. The energy spectrum of these electrons is illustrated at the right. There are three principal groups of electrons. First are the very slow secondary electrons, created by the primaries striking the gate metal. Second are the accelerated secondary emission electrons from the mucap surface. Third are the elastically scattered electrons with approximately the primary beam energy, V_P. The stored information is represented by the central peak. In the lower part of the figure, a ZERO is illustrated; the stored potential is approximately that of the gate metal. In reality, it is usually slightly positive with respect to the gate. In this case, secondary electrons being emitted from the mucap surface are not accelerated as they leave the storage element and combine with the slow secondaries from the gate metal in the energy spectrum shown at the right. Thus the central peak referred to above is absent during reading of ZERO-state elements. It is

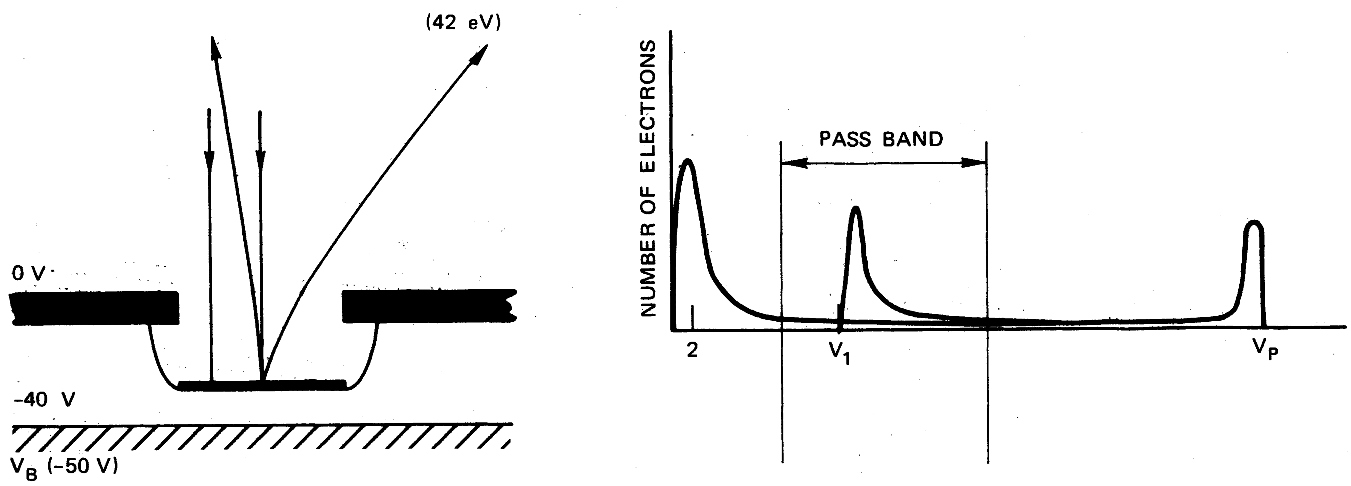

(a) READ ONE

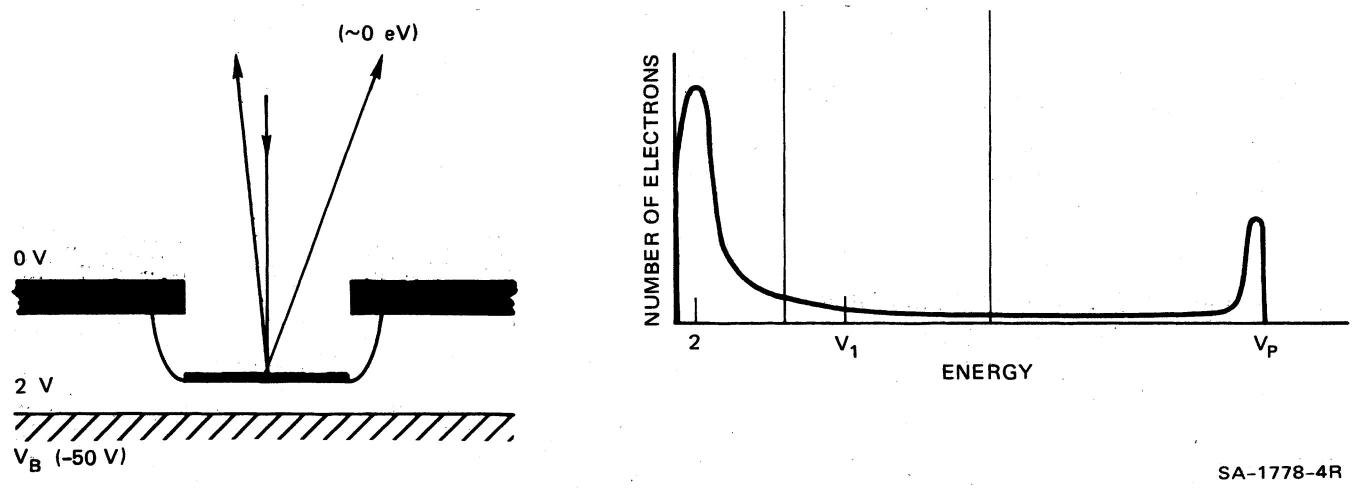

(b) READ ZERO

Figure 5. Zero and One States and Reading Mechanisms

therefore relatively simple to incorporate an energy analyzer into the electron optical system in order to pass only the central band of electrons. Measurements on our present analyzer indicate that a ratio of 100:1 between the two states is readily obtainable.

The writing process is somewhat complex, and a simplified description will be given here. In the present mode of operation, writing entails using a primary beam energy that creates more secondary electrons at the emitting surface than the primary electrons arriving; i.e., the secondary electron emission coefficient is greater than unity. Under these conditions, the mucap potential will always tend toward the gate metal potential under electron bombardment. This will happen whether the mucap is initially positive or negative, since when the mucap is positive the slow secondaries created at the surface cannot escape from the mucap cavity and are returned to the mucap surface, which therefore has an effective secondary emission coefficient of less than unity.

According to the above description, reading a ONE state mucap can be expected to change the ONE to a ZERO. This is indeed the case, as illustrated in Figure 6, although note that readout need only be partially destructive, a significant factor that is discussed below. In writing a ONE, the target base is switched to zero volts, so that all ZERO-state mucaps are made positive with respect to the gate metal. As explained above, bombardment with the primary beam will set these elements to the gate metal potential. Reestablishment of the bias of −50 V on the base electrode will than pull these elements down to the negative condition shown in Figure 5(a). A tabulation of operations is shown in Table 2. Two different times, t_w and t_r, are given for the writing and reading operations respectively, since a writing operation requires more charge to be delivered to the target than the reading operation. The alternative to varying the read and write times is to adjust the beam intensity. Either method may be used. For a platinum secondary emission surface having a secondary emission coefficient of 1.6 at 2 kV, and using an electron beam having a current density of one ampere/cm^2, the writing time should be approximately 0.5 μs per bit. Since reading the target requires only that the stored charge

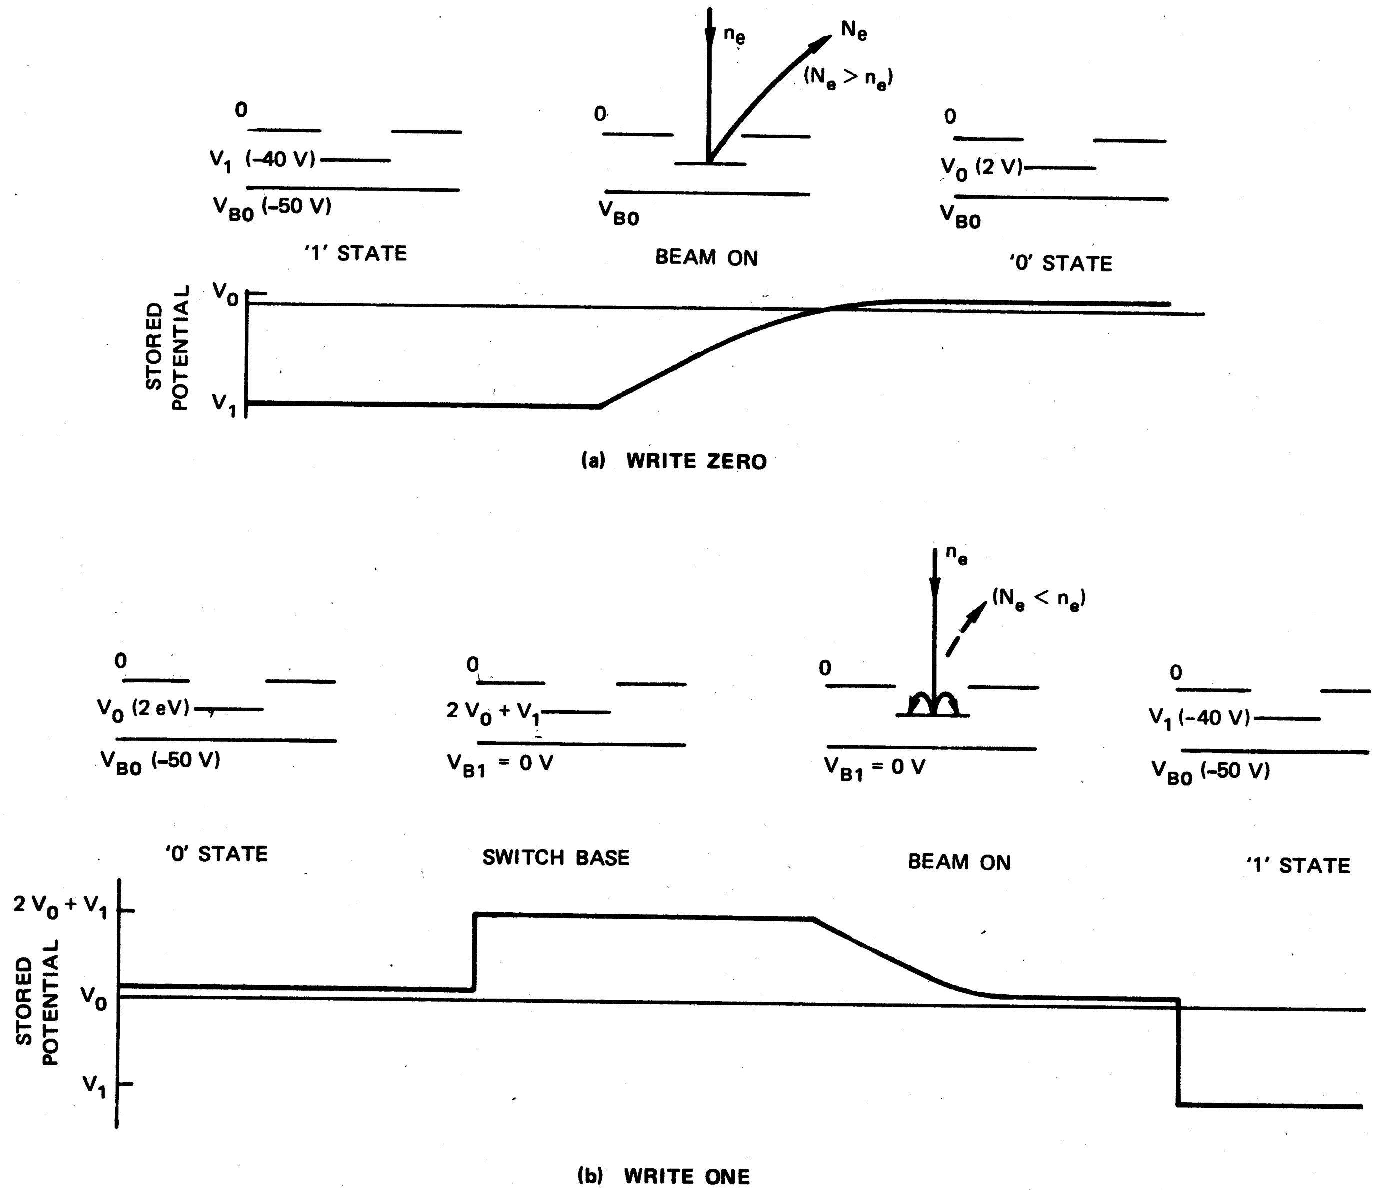

Figure 6. Writing Sequences for ZERO's and ONE's

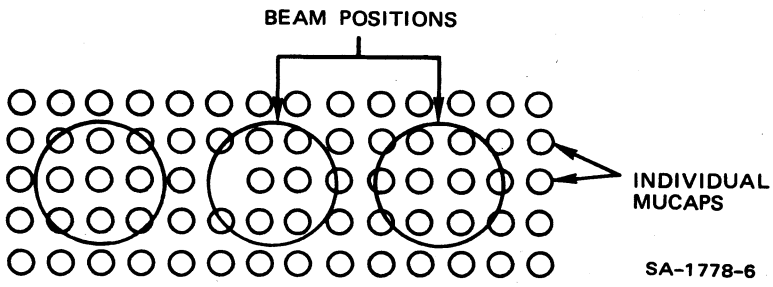

Figure 7. "Quasi-Elements"

Table 2. List of Operations

Operation	*Target Bias*	*Beam On*
Write ONE	0	t_w
Write ZERO	−50 V	t_w
Read	−50 V	t_r

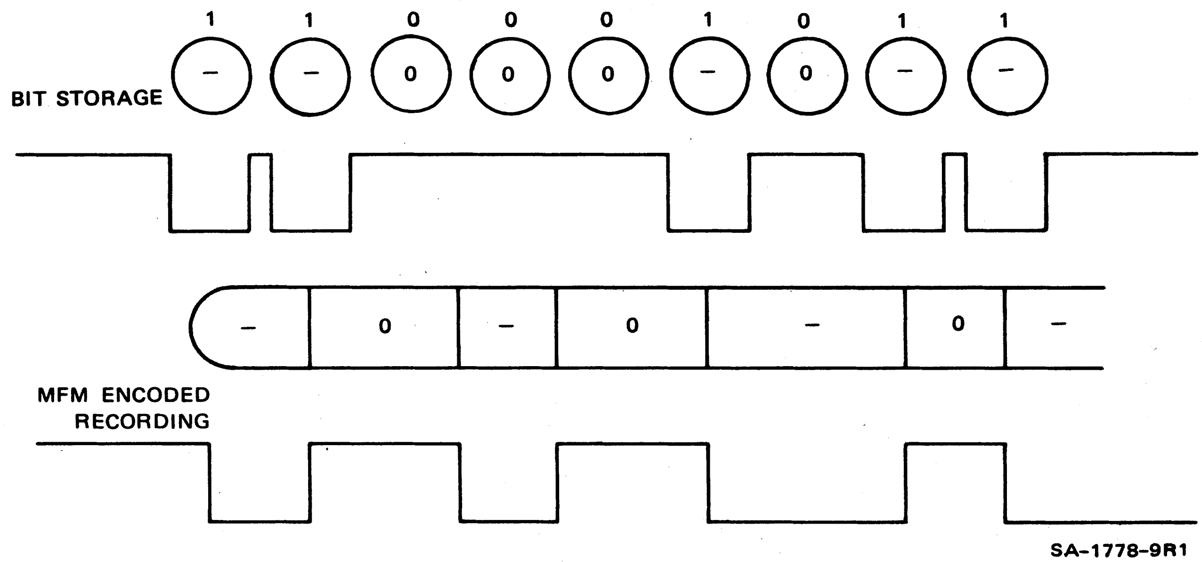

Figure 8. Bit Storage and Encoded MFM Recording

should be sampled, it is possible to use very much shorter pulses than for writing. Although not yet demonstrated in our laboratory, it is believed that reading rates between 10 and 100 MHz will be easily obtained in a single tube.

Target Usage

As mentioned in the introduction, the target is used in the so-called quasi-continuous mode in which each bit is represented by the charge stored in a number of individual mucaps. Bit storage is illustrated in Figure 7. In this case, a group of approximately nine mucaps is used to store each bit. The presence of one or two defective mucaps does not affect target operation. In the system delivered to the Air Force, each bit location is addressed randomly. The deflection system is allowed to settle at a particular location, and then the beam is turned on to interrogate the charge stored at that target position. When operating in this mode, it is desirable to have a reasonable guard area around each stored bit. However, it is highly probable that most EBAM systems will be used in a block-oriented mode and that information will be written as an encoded string of transitions along a continuous path. Random access to the start of the track would be achieved digitally, and the track itself scanned on the fly, with the beam left on continuously. This is illustrated in Figure 8. The upper part of the figure shows a data pattern written as a string of individual bits, and the corresponding charge pattern. The lower part of the figure shows the same data written using MFM encoding. Note that ONEs are represented by a transition (in either direction), and ZEROs by the lack of a transition. Where there is a string of ZEROs, transitions are placed between bits in order to assist in timing the data out.

As mentioned earlier, reading partially destroys the stored information. A succession of reads on ONE-state elements will convert them into the ZERO state. This occurs as illustrated in Figure 9, which shows the discharge curve from the ONE state into the ZERO state as a function of the number of reads. In the present system only a single comparator level is used to distinguish between ONEs and ZEROs. The problem of refreshing the data is solved by incorporating a second comparator level as indicated in the figure. Pulses falling in the region between Levels 1 and 2 would then logically be assigned as a ONE, but would cause refresh of that particular piece of data being read. The incorporation of an automatic refresh technique has several implications for this type of memory tube. First, approximatley 100 reads without rewriting the information has already been demonstrated in the laboratory, and a simple calculation based on a 10-Mbit read rate indicates that several thousand reads should be obtainable for every write. The second factor arises because electrostatic charge storage is somewhat volatile. Early storage targets show several hours of storage and it would appear possible to obtain storage times of several days in

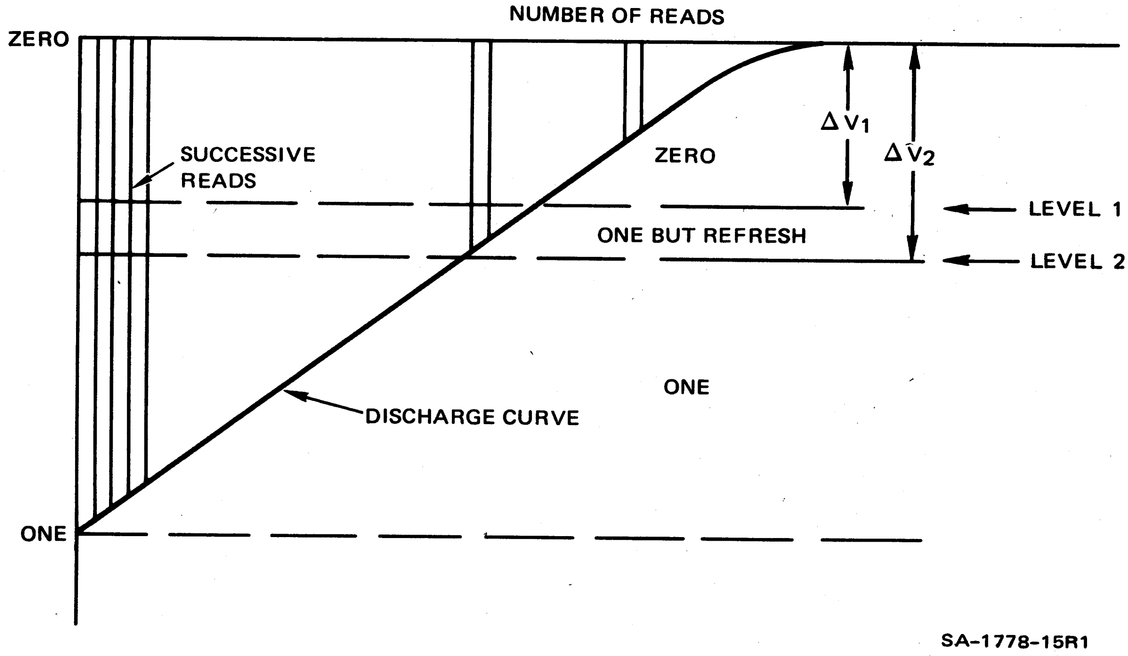

Figure 9. Readout Technique Using Automatic Refreshing

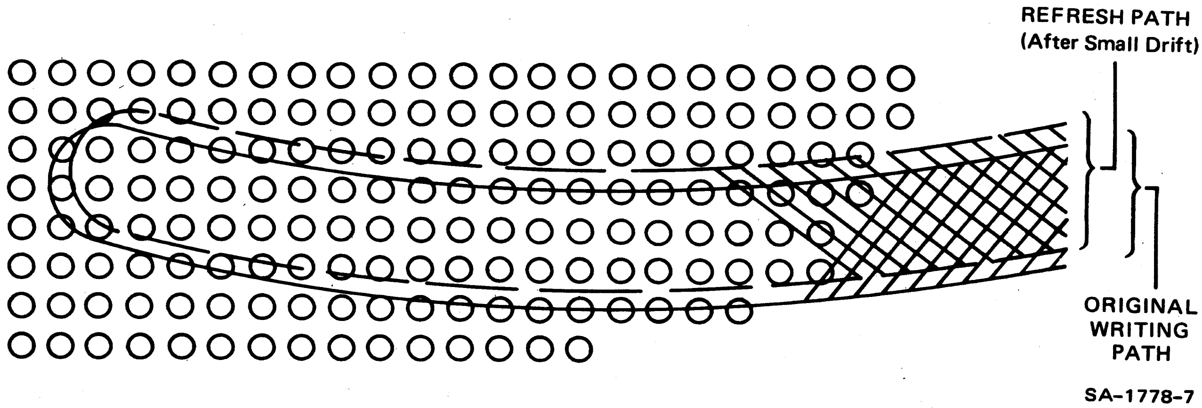

Figure 10. Eliminating Long-Term Drift by Periodic Refreshing

production. Indefinite information retention can be achieved by ensuring that the complete memory is read periodically and therefore is automatically refreshed. The automatic refresh technique also enables us to correct for long-term drifts in the system since misalignment between the reading beam and the stored information will cause a decrease in ONE-level output. Refresh repositions the data under the drifted beam position as illustrated in Figure 10. The actual time required between refresh will be governed by many factors including system drifts, target volatility, and bit interference effects. The latter can be caused when neighboring tracks of information are read repeatedly or rewritten frequently, since the cross section of the writing beam will nearly always be gaussian in shape; the tails on the spatial distribution cause the interference effects. If even larger quantities of information are to be stored in a single tube in the future, it will probably be necessary to incorporate some fiducial mark system for locating the data. The need for this referencing system will probably arise between 10^7 and 10^8 bits per tube. Several schemes have been proposed at SRI, but not implemented. RCA has developed one technique, also under Air Force sponsorship.

The Storage Tube

The design goal for the first fully operational system is 4×10^6 bits per tube on a target 1.8 cm $\times$ 1.8 cm, or a bit spacing of 8 μm. Access time to any address is to be less than 3 μs. The complete tube assembly as delivered to the Air Force is illustrated in Figure 11. It is made using a very precise erector set type of construction that permits a great deal of flexibility for making changes in the laboratory. The electron optical components are mounted on support rings, which are held together by three 0.5-inch diameter rods. The gun is at the bottom; next are the lens, the deflector, the readout analyzer, and the target plane.

Details of the actual storage target plan are shown in Figure 12. This plane can hold two 0.25-inch storage targets, one placed on either side of the center axis. In this figure one target has been removed to show the mounting holes.

The required access time, address accuracy, and freedom from beam distortion are determined by the deflection system. This comprises an octupole deflector[1] and the associated electronics. The octupole deflector (Figure 13) used in present systems offers significant advantages over

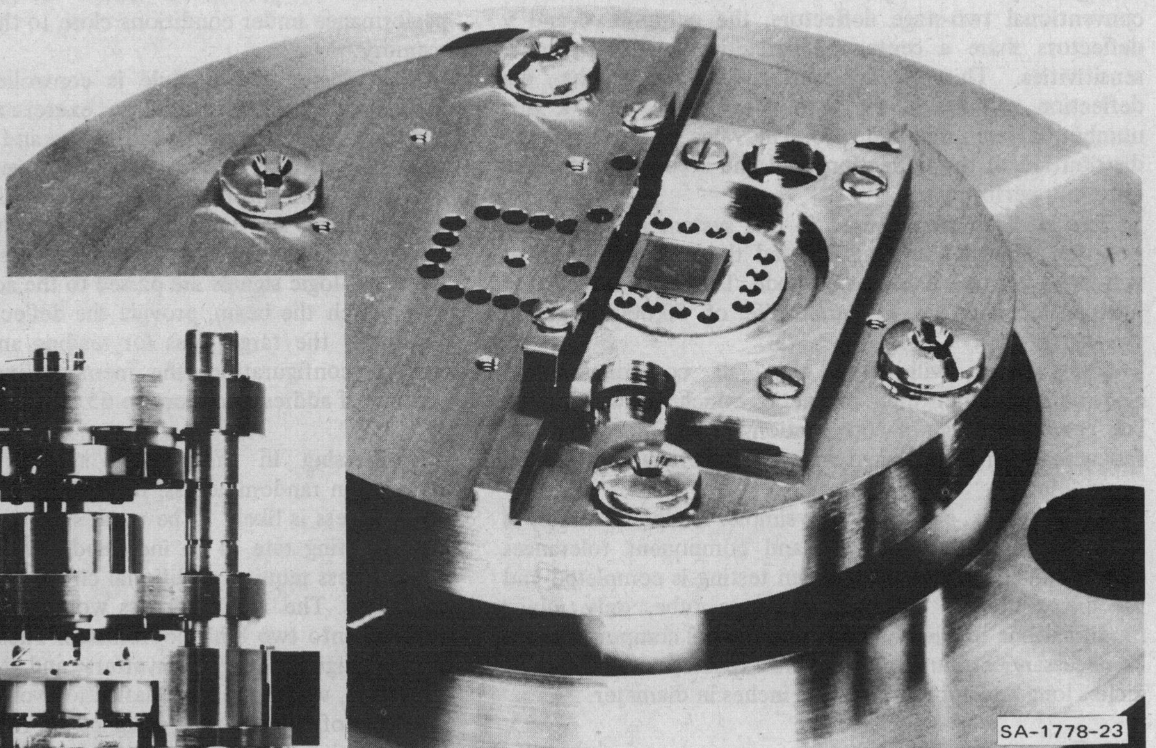

Figure 12. The Storage Target Plane

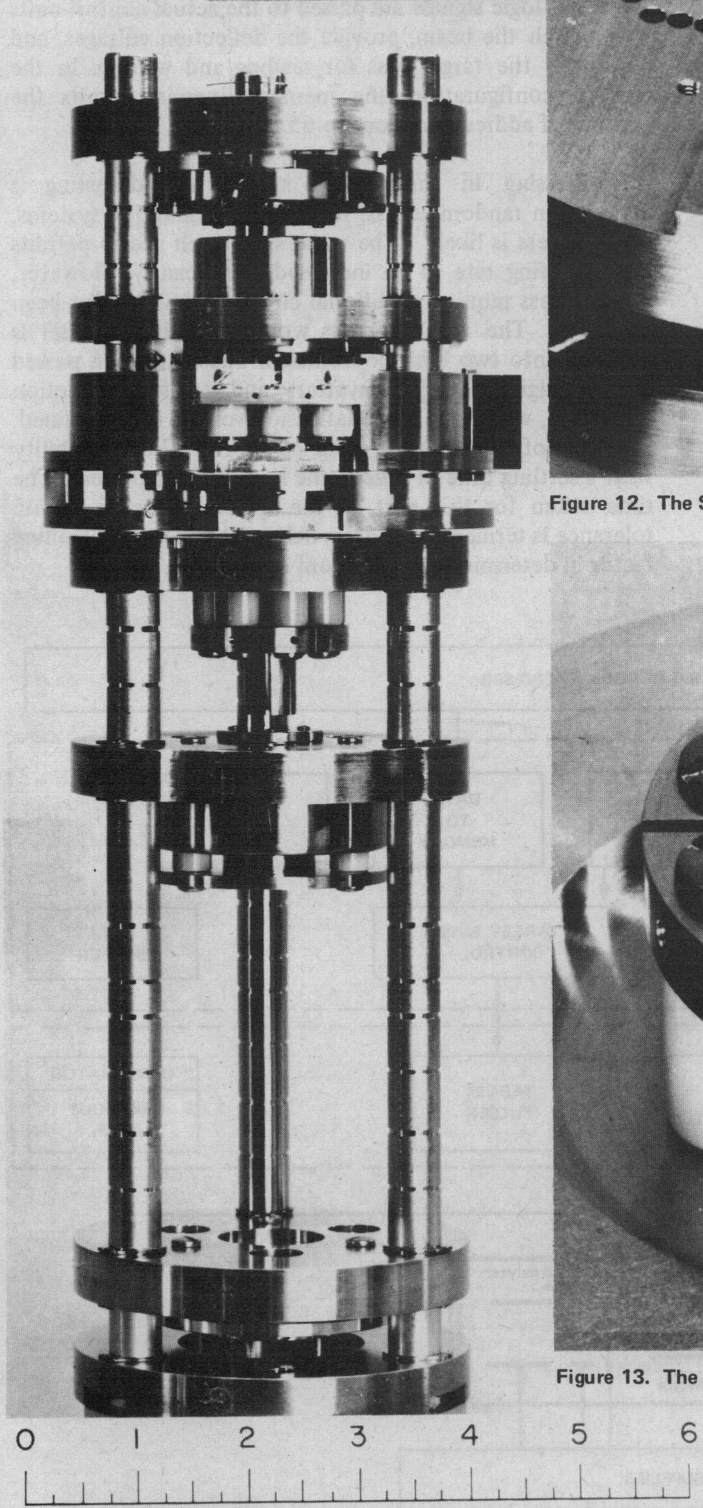

Figure 11. The Electron-Optical Bench

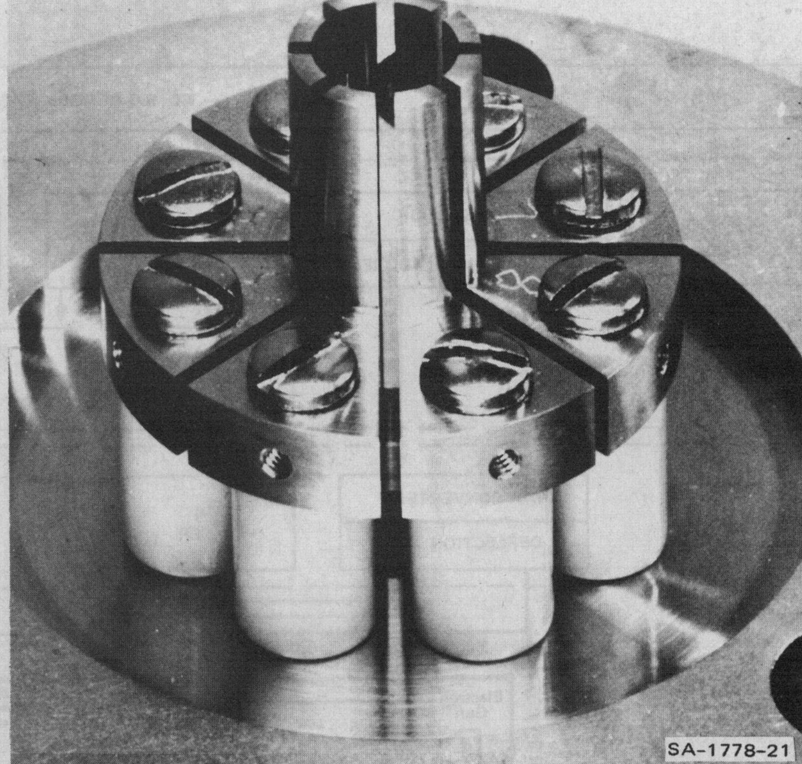

Figure 13. The Octupole Deflector

more conventional deflectors. It consists of eight electrodes arranged symmetrically about the beam axis. Unlike conventional two-stage deflectors, the octupole x and y deflectors share a center of deflection and have equal sensitivities. This is achieved by compounding the deflection voltages in the form $V_x + kV_y$, where k is a number between one and zero. The value of k controls the distribution of the deflection field and is chosen to achieve optimum performance.

The electron source used in this system is a thoriated tungsten cathode. It is anticipated that production tubes will use a barium dispenser cathode. The latter is relatively inexpensive and has a quoted life of 50,000 hours at 2 A/cm^2.

The reader will note that all components are electrostatic rather than magnetic, which allows simpler, less costly, and lighter construction. In addition, accurate, fast access can be obtained without hysteresis in the beam positioning.

In the future, a considerably simpler readout assembly is anticipated. All components and component tolerances have been designed so that when testing is completed and we are ready to build a prototype tube, only minor modifications to the basic electron optical components will be necessary. Production tubes are expected to be about 12 inches long and approximately 2 inches in diameter.

The Demonstration System

The main elements of the experimental unit delivered to the Air Force are shown in Figure 14. This unit combines all the components of a complete prototype system and is intended to provide a means of evaluating EBAM performance under conditions close to that of an operating memory.

The single tube module is controlled by a standard Macrodata MD-100 memory exerciser. This exerciser produces data patterns, addresses, and read/write commands in parallel. It also provides a means of comparing data read from the EBAM with data written originally. The data, address, and operational commands are passed to an interface unit where all the timing is determined. From the interface, logic signals are passed to the actual control units that switch the beam, provide the deflection voltages, and determine the target bias for reading and writing. In the existing configuration, the memory exerciser limits the number of address locations to 65,536.

Addressing In the present system all addressing is achieved in random access. In the first production systems, block access is likely to be used, since such access permits the operating rate to be increased significantly. However, block access requires additional circuitry that has not been developed. The 16-bit address word from the exerciser is decoded into two 8-bit coordinates. These are then passed to the digital/analog converters and to the deflection amplifiers, where the final deflection voltage is determined. This part of the circuit is designed for full 12-bit capability with a settling time of 3 μs to the level of 0.01 percent. The time taken for this part of the system to settle within tolerance is termed the address delay. This is the controlling factor in determining the random access time.

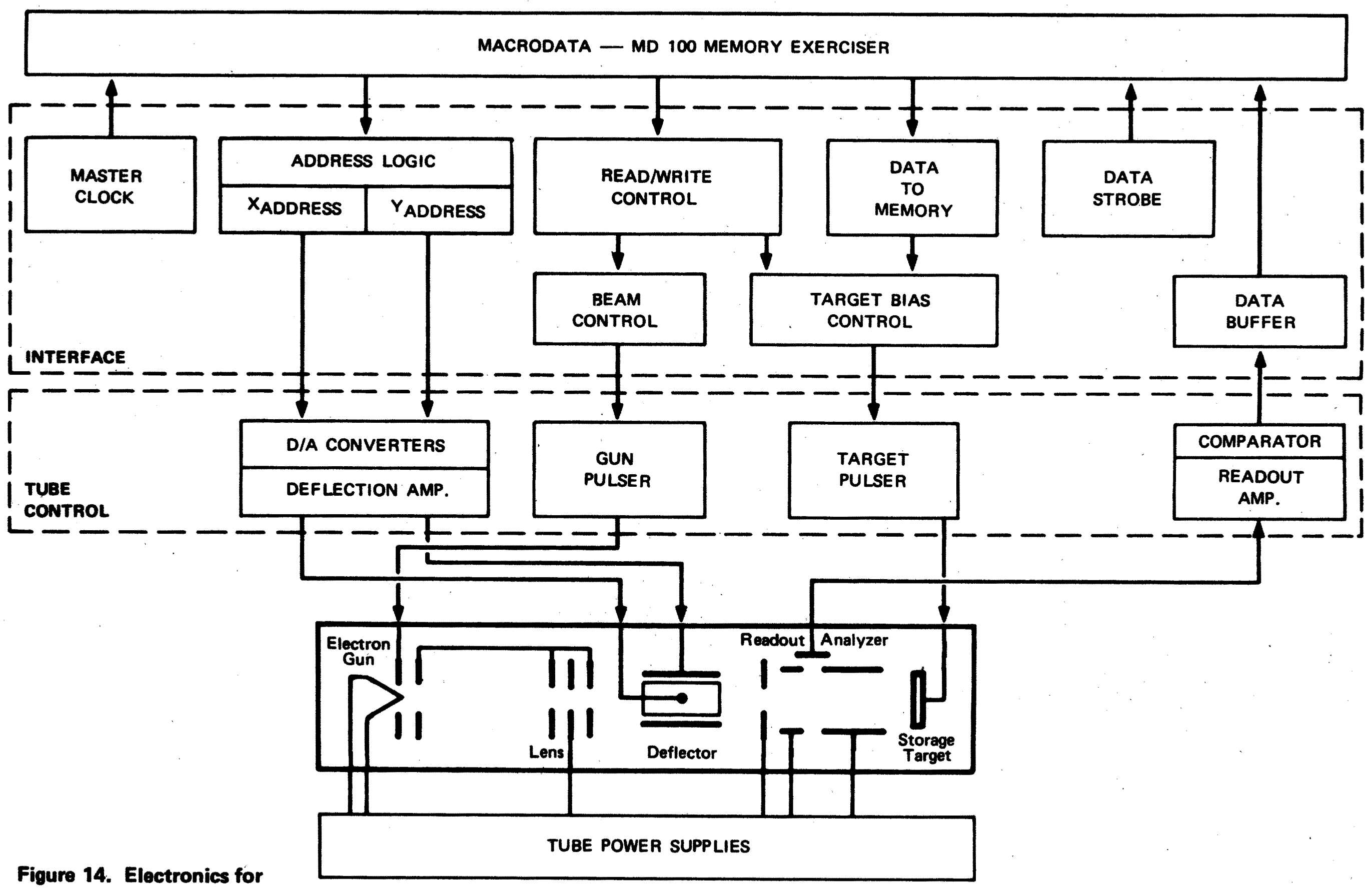

Figure 14. Electronics for EBAM Module System

Reading and Writing The reading and writing operations are carried out by controlling the bias applied to the target base and the length of the beam pulse. Thus, the data to be written determine the bias during a *write* command (−50 V for a ZERO and 0 V for a ONE). During a *read* command the bias is set to −50 V. This bias is set at the start of the memory cycle, but the beam is not turned on until the address delay has been completed.

Reading During a read cycle the signal from the readout analyzer is passed to an amplifier and then to a comparator, which is set to determine the binary condition. The data are then stored in a buffer in the interface until they can be strobed back into the exerciser and compared with what was originally written. If an error is encountered, a pulse is passed to an external connector on the exerciser, from which it can be accessed for counting purposes or used with a suitable display on an X-Y oscilloscope for error mapping. Typical readout signals for a train of alterante ONEs and ZEROs are shown in Figure 15. These signals are as viewed after amplification and just before the comparator circuit.

Performance and Conclusions

The complete demonstration system is shown in Figure 16; the exerciser is at the left, the control unit is in the center, and the memory tube is concealed in its magnetic shield on the right.

During system testing at SRI, the module has operated quite satisfactorily. The most difficult problem was to design a satisfactory readout amplifier. However, the system has now been operated up to 1-MHz read rates, and it appears that 10 MHz will be easily obtainable in the block-oriented mode. Some early measurements of error rates appear extremely encouraging; e.g., an error rate of less than 1 in 10^9 was encountered in the first system tests.

The major research remaining is that of target life. At present, attempts to measure a real lifetime appear to be masked by contamination of the medium surface. This is not surprising since the demountable tubes currently in use are not processed or assembled in an ultra-clean environment. It is also possible that the targets themselves are not sufficiently clean before insertion into the tube. In one experiment an exposure of 15 C/cm^2 caused the signal to decay to its permissible limit. Assuming a 25-μ diameter bit and a writing efficiency of 0.5, and assuming a block access mode of operation with uniform target usage, a 1-MHz write rate and a 10^6 bit target, the life would then be five years. Since an improvement by at least one order of magnitude seems possible, a very respectable life may be anticipated.

Apart from actually proving that the medium can have adequate life, there appears to be no major impediment to the successful construction of a complete operating

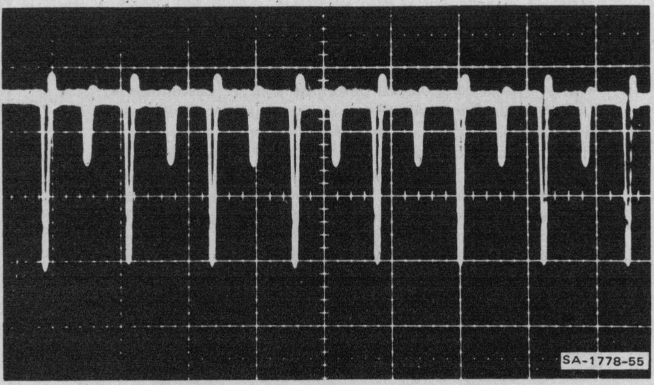

Figure 15. Amplified Readout Signal

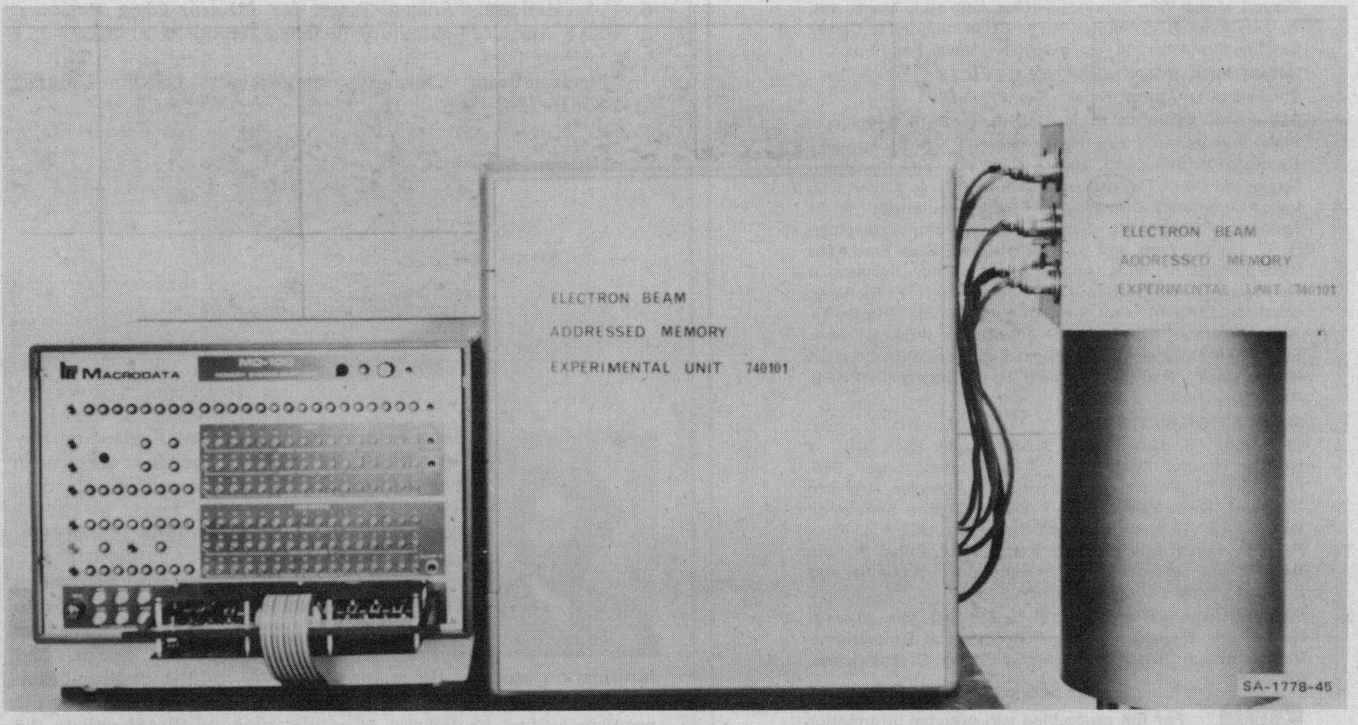

Figure 16. USAF EBAM Module

memory. We believe that a complete memory system, including tubes, can be packaged in a single unit somewhat smaller than the control unit discussed here.

Acknowledgments

The author of this paper acknowledges the support of the Memory Technology Group of the USAF Avionics Laboratory.[7] I would also like to thank J. S. Moore and P. R. Thornton for their help throughout the work, D. Limuti and J. Clark for their work on the electronics, and C. C. Hartelius and E. R. Westerberg for fabricating the storage targets.

References

1. J. Kelly, J. S. Moore, and P. R. Thornton, "The Development of an Experimental Electron-Beam-Addressed Memory Module," *Proceedings of the IEEE 1974 National Aerospace and Electronics Conference*, p. 55.

2. J. Kelly and J. S. Moore, "Electron-Beam Addressed Memory Research," USAF Avionics Laboratory Technical Report AFAL-TR-74-176.

3. "High-Information-Density Storage Surfaces," L. N. Heynick, ed., Seventeenth Quarterly Report, ECOM 01261-17, Contract DA 28-043 AMC-01261(E), SRI Project 5444, Stanford Research Institute, Menlo Park, California (October 1969).

4. G. W. Ellis, G. E. Possin, and R. H. Wilson, "Diode Detection of Information Stored in Electron-Beam-Addressed MOS Structure," *Applied Physics Letters*, Vol. 24, No. 9 (1 May 1974).

5. M. S. Cohen and J. S. Moore, "Physics of MOS Electron-Beam Memory," *J. Appl. Phys.*, Vol. 45, No. 12, pp. 5335-5348 (December 1974).

6. F. J. Marlowe, "Silicon Target for Electron-Beam Memory," USAF Avionics Laboratory Technical Report AFAL-TR-73-253.

7. "Electron-Beam Memory Research," USAF Contract F33615-72-C-1906.

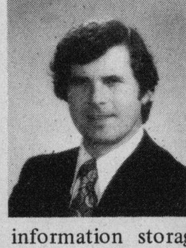

John Kelly was born in Bristol, England on May 18, 1936. He received the B.A. degree in Physics from Cambridge University, England, in 1959. He received an MSc degree in Physics from London University in 1964 and an M.A. from Cambridge in 1966.

In 1966 he joined the Physical Electronics Group at Stanford Research Institute in Menlo Park, California, where he worked on electron optics, electrostatic deflector design, and information storage. He is currently head of the electron-beam-addressed memory (EBAM) program and was responsible for the program leading to delivery of an experimental EBAM unit to the USAF in January 1974.