5. High-Speed Memory Development at the National Bureau of Standards

R. J. Slutz, A. W. Holt, R. P. Witt, and D. C. Friedman

1. INTRODUCTION

The problem of obtaining reliable and economical storage of information at high access rates has consistently been one of the most difficult in the field of computer development. During the past several years, the NBS Electronic Computers Laboratory has been active in the development and evaluation of promising memory techniques, particularly the acoustic delay line, the Williams method of cathode-ray storage, and the diode-capacitor system.

When the initial plans for SEAC were being developed in the latter part of 1948, there was no reasonably economical high-speed computer memory system available. However, acoustic delay lines of mercury had been used successfully in radar application, and a memory system based on acoustic delay in mercury was under development for the EDVAC computer at the Moore School of Engineering, University of Pennsylvania. This type of mercury design appeared to be well suited to the immediate need for a reasonably reliable high-speed memory for the NBS Computer.

With the completion of SEAC in early 1950, work was begun on a memory improvement program with the chief objective of increasing the speed of computation through the use of higher speed memories. At this time, the Williams system seemed to be the most promising of the memories under development. This laboratory, as well as other groups in this country, designed a Williams type memory, using the parallel mode of operation, which is basically faster than the original Williams-Kilburn serial system.[1]

A more recent NBS development in the search for rapid-access memories is the diode-capacitor memory.[2] This system was suggested by A. W. Holt and is being developed by a group under his direction. It is a striking example of the inadequately recognized situation that the limiting part of a rapid access memory is not the memory itself but the access. This scheme uses the simplest of storage devices, an ordinary capacitor, and gets its importance from an efficient access scheme in which combinatorial tricks reduce the access circuitry to only a little more than two diodes per bit.

2. THE NBS MERCURY MEMORY

The physical structure of the mercury delay line memory for use with SEAC followed that of the EDVAC memory system with only minor modifications. However, entirely original electronic circuit designs were developed. The system provides memory storage capacity of 512 words, each word requiring 48 pulse times and eight words stored in each of 64 mercury lines.

The delay line used in both the EDVAC and SEAC acoustic memories consists of a mercury-filled Pyrex tube with a quartz crystal at each end to serve as electroacoustical transducers. These quartz crystals are backed up by glass end-cells filled with mercury. The delay line proper is approximately 20 in. long and provides 384 μsec of acoustic delay when operated at 50° C. The central body of mercury used for the transfer of acoustic energy is electrically grounded, and the two backing pools of mercury act as the electric input and output terminals.

The information pulses that are to be stored are energy packets resulting from on-off modulation of an 8-Mc carrier. About 200-v peak-to-peak is applied to the sending crystal, which produces

---

[1] F. C. Williams and T. Kilburn, A storage system for use with binary digital computing machines, Proc. IEE (British) part II, **96**, 81-100 (1949).

[2] A. W. Holt, An experimental rapid access memory using diodes and capacitors, Proc. ACM (December 1952).

305790 O - 54 - 7

acoustic waves whose amplitude nearly causes cavitation. This crystal sets up waves from each face; one travels down the line, and the other is scattered in the end-cell. As the receiving crystal closely matches the acoustic impedance of mercury, most of the energy that reaches it passes through and is scattered in the receiving end-cell. A voltage of several hundred millivolts is created by the receiving crystal.

Although there is a total attenuation of about 60 db between the input and output circuits, only about 3 db are lost in acoustic attenuation in the mercury. One-side-loaded double-tuned coupling transformers are used primarily to provide electrical isolation and incidentally to minimize the attenuation that results from the low resistances necessary to achieve a 4-Mc bandwidth in the face of the comparatively large capacitances of the crystal transducers.

The 60-db attenuation makes special precautions necessary to prevent the acoustically delayed signal from being overwhelmed by direct electric coupling between the end circuits. This coupling results from the impedance of the ground connections to the central mercury column, which is common to both circuits, and is aggravated by the fact that the mercury is contained in glass tanks that permit only restricted contact areas. The wire or ribbon leads used may have a considerable contact resistance with the mercury. A solution to these difficulties was found in the circuit arrangement shown in figure 5.1. Double ground leads, insulated end-cell shields, and floating windings of the coupling transformers permit each end circuit to be a self-contained closed circuit from which almost no current escapes to flow in any common ground path. In this way the direct electric feed-through has been made negligible even when contact resistances have reached the order of thousands of ohms.

Platinum, tungsten, and stainless-steel grounding leads were tested in the original SEAC memory installation, together with 10 variations of cleaning procedures for the glass assemblies. The question of effective treatment of the glass tanks arose because of early speculation that the mercury might become contaminated, but no difficulty from contamination has occurred during the 3 years the SEAC mercury memory has been in use, perhaps because of the circuitry used to prevent electrical feed-through. All delay lines were filled dry rather than by displacement of some other liquid by the mercury. In the 3 years of operation of this computer only two lines have been removed from position. One was removed as soon as it was discovered that its delay differed from the other lines by more than a tenth of a microsecond. The other was unnecessarily removed because of trouble that was subsequently traced to a faulty plug-in connection.

In order to store signals indefinitely in an acoustic delay line, it is necessary to perform the various circuitry functions shown in figure 5.2. These circuits permit a signal introduced into the amplifier to modulate a carrier frequency, which is then sent into the delay line where it is converted into an acoustic signal. After 384 μsec it is picked up from the delay line and amplified, detected back to a video signal, reshaped, retimed, and recirculated back through the delay line. This process may keep on indefinitely or until the computer replaces these pulses with other information.

The input and output of the mercury delay line consists of transformer-coupled circuits which are double-tuned to the basic carrier frequency of 8 Mc and are one-side-loaded only for an over-all bandwidth of approximately 4 Mc. These double-tuned circuits are used primarily for electrical isolation, although they also make increased bandwidth possible. Interstage coupling circuits in the IF amplifier are single-tuned. Because the output of the last IF stage must drive into the relatively low-impedance nonlinear detector, a double-tuned circuit is used for this impedance-matching function. Because the secondary of this tuned circuit is completely floating, a relatively high-impedance voltage-doubler detector could be utilized. This detector feeds into a low-pass filter of approximately 100-ohm impedance.

Following the detector output are a series of germanium diode gates which permit a narrow time sample of the detector output to be taken. This sample is then fed into a transformer-coupled amplifier stage and broadened by regenerative broadening. Because the output of this transformer-coupled stage commences with the time of the narrow interrogation pulse which may be held to a fixed timing relative to the main computer, this circuit serves to resynchronize the delay line information with the computer operating rate. The other diode gates into this stage are used to intercept

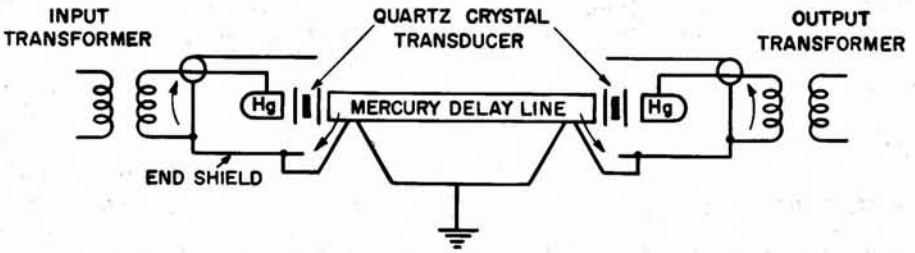

FIGURE 5.1. *Schematic diagram of circuit connections to the acoustic delay line used in NBS mercury memory.*

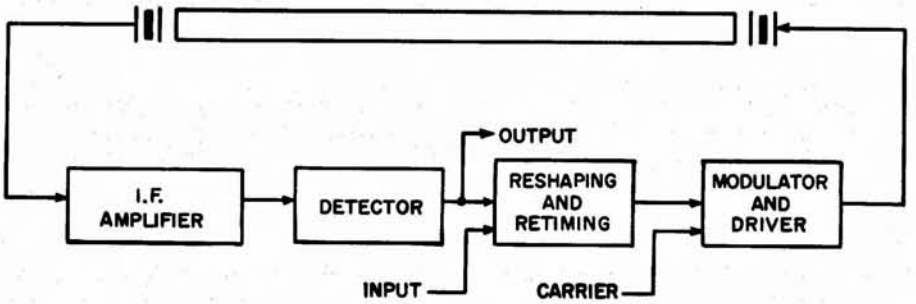

FIGURE 5.2. *Block diagram of the mercury memory system.*

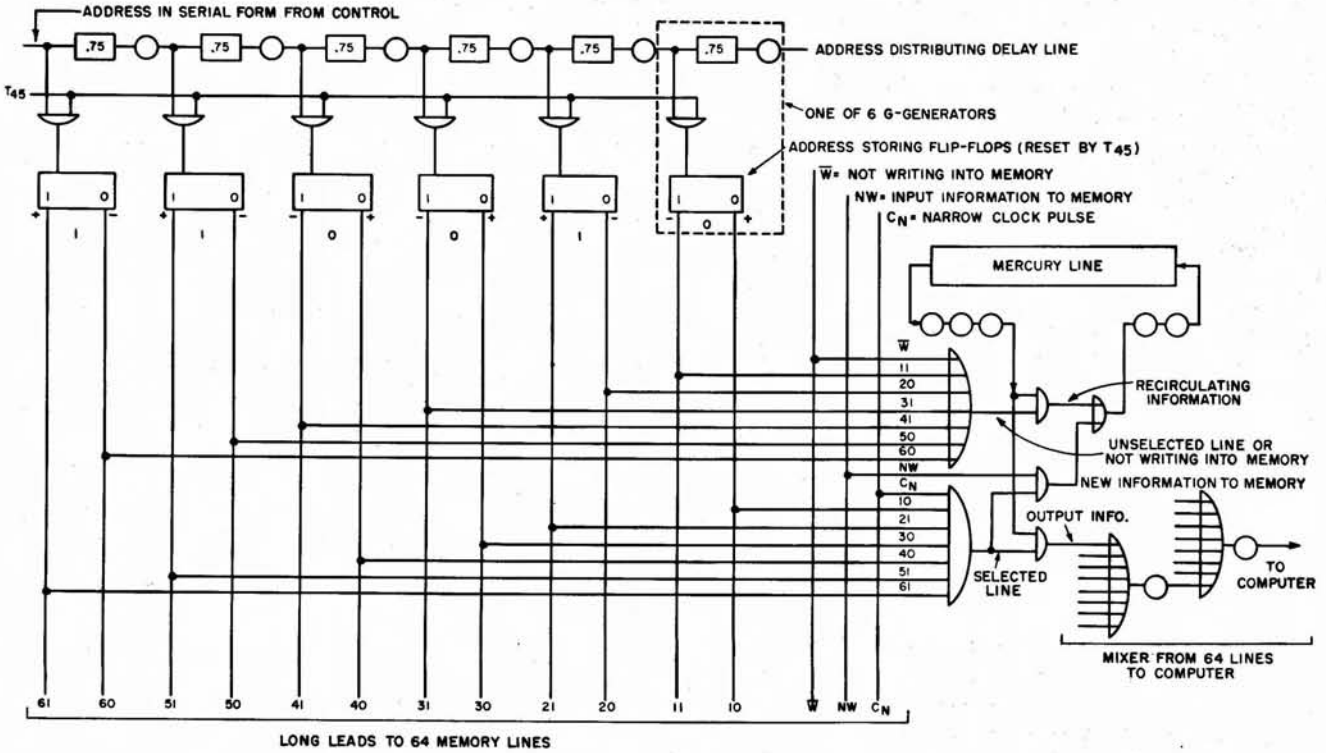

FIGURE 5.3. *Mercury line selection system.*

information that is recirculating in the delay line if a word is to be cleared out in order to insert new information.

The output of the transformer-coupled stage consists of a half-microsecond pulse of very low impedance and approximately 20-v amplitude. This signal is broadened to approximately 0.75 μsec, using distributed electrical delay lines. This signal is then modulated by an 8-Mc carrier signal in a second germanium diode gate which feeds into the driver stage for the mercury delay line. The pulse transformer, all germanium diodes, and the detector circuit are completely plug-in to facilitate maintenance. The amplifier unit as a whole is also plug-in to the main memory cabinet.

Selection of one out of the 64 mercury delay lines in the memory cabinet requires six binary digits, since $2^6=64$. A line may be selected either to read stored information from the memory into the computer, in which case the information is also retained in the memory, or else to write new information from the computer into the memory, in which case the information previously stored there must be cleared out.

The six line selection digits are delivered serially over one wire from the computer to the line selection circuits in the memory cabinet, shown in figure 5.3, where a six-stage distributor line makes them available on six different wires at the time of a pulse called T45. This T45 pulse gates the selection digits into six dynamic flip-flops which have special high-powered output stages to drive the large diode matrix which actually selects a memory line. The signals driving the selection matrix are transformer-coupled and reverse polarity every half microsecond, but all that is significant for selection is their polarity during the narrow clock pulse C_n, which controls the admission and recirculation of information pulses in the mercury memory.

Actually, there are two diode selection matrices driven from the same input leads. One matrix, comprised of or-gates, supplies pulses to every memory line except the selected line in order to gate the previously stored information on around the recirculation path. If the operation is to read from rather than to write into the memory, then circulation must be maintained even in the selected line. This is done by pulses on the line $\overline{W}$ (meaning "not writing"), which is wired into every or-gate.

The other matrix is comprised of and-gates so that it supplies pulses only to the selected memory line. These pulses, timed by C_n, serve to gate information either into or out from the selected memory line.

Three additional information digits are needed to select a particular word of the eight words in a delay line. A counter in the computer records which memory position is currently available at the detector output, and comparison of this counter with the three time-selection digits is used to initiate gating of the desired word.

The memory was considered adequate for use in the transportable DYSEAC computer with only minor structural changes. The same design has been used relatively unchanged by several computer groups throughout the country. Seven models have been constructed with only minor modifications from the original NBS design. Successful operation has been reported from all groups having this memory.

3. CATHODE-RAY TUBE MEMORY

With the decision in 1950 to develop a Williams-type memory, work at the Electronic Computers Laboratory was planned along five major lines: (1) circuitry for a parallel memory system; (2) equipment to determine the suitability of tubes for storage purposes; (3) methods to circumvent the deficiencies of standard cathode-ray tubes; (4) program for improving storage tubes; and (5) development of the theory of Williams-type storage. An early result of the memory improvement program was the construction of a full-scale experimental electrostatic memory of 48 tubes, each storing 512 binary digits of information. This equipment was completed and placed in experimental operation in SEAC by February 1951 and was one of the earliest full-scale Williams-type systems constructed in this country.

In the SEAC Williams memory,[3] each tube is assigned one particular digit of the 45 binary digits making up a word that can be either an instruction or a number. (The other three of the 48 positions in the SEAC word are spares.) All of the digits of one word are assigned the same relative positions in the array of 512 stored spots on each tube. Figure 5.4 is a block diagram of the SEAC electrostatic memory system.

Action during a reading operation is as follows: The address (numerical designation of location) of a word is taken from the Address Register and translated into X and Y coordinates of its location on the tube face by the Staticizer and Deflection Generator. All deflection operations are done in parallel. In 3 μsec, the deflection transients disappear, and the beam is turned on in all tubes for 0.5 μsec (see fig. 5.5). The outputs from the signal plates are amplified and sensed during the STROBE period which occurs within this time. A positive output represents a binary one, and a negative output represents a binary zero. This information is transferred to the Shift Register in parallel, from which it is fed serially to the arithmetic unit or to an output device.

Since the act of reading causes all spots to be written to "zero," the "ones," must be restored. To accomplish this, at the end of the DOT pulse all tubes are provided with a sawtooth deflection pulse which displaces the beam about one spot-diameter. If the output from any tube is positive at STROBE time, the beam is held on by the DASH pulse which occurs during the TWITCH. The spot charged during reading is then discharged by secondaries from under the moving beam and a one is thus restored. If the output of the tube is negative, a zero has been sensed and rewritten during DOT so that no other action is needed. The absence of the positive output prevents the beam from being held on during the TWITCH and the spot is left charged. The total time allowed for this read-write operation is three microseconds. An additional time of 6 μsec is required to allow the staticizer pulse to swing back, making a total operation time of 12 μsec, 3 for deflection, 3 for read and rewrite, and 6 for flyback. Writing is done in similar fashion. The information is stored in the shift register and is used in place of the output of the amplifier to control the holding on of the beam during TWITCH.

Since the charges stored tend to leak away or are discharged by secondary electrons and stray primary electrons during reference to neighboring spots, it is necessary to regenerate (read and rewrite) each word periodically. The operation is the same as consultative reading, except that the source of the address to be rewritten is the regeneration counter rather than the address register, and no information is gated into the shift register. The regeneration counter advances consecutively through the electrostatic storage addresses as each word is regenerated. Since the minimum time to perform an operation in the arithmetic unit of SEAC is 48 μsec, three regenerations can be carried out between each useful consultation. If there is no consultation of the electrostatic memory at the allotted time, an additional regeneration occurs. This may happen when SEAC is using both memory systems and an address in the acoustic memory is called for, or during a long operation such as multiplication.

In operation with SEAC, over 1,500 hr of useful operation have been obtained, but it has never reached the long-term reliability of the SEAC mercury memory.

The problem of supplying storage tubes for a Williams-type system has proved more difficult than was originally anticipated. During the early planning of Williams-type memory systems in this country, standard cathode-ray tubes were expected to provide adequate data storage, and all but one of the currently operating computers which have Williams memories do use standard types of tubes. However, standard tubes which were commercially developed for other purposes have certain limitations when they are used in high-speed memory systems. These limitations include nonuniformity of surface, which causes the blemish problem; interaction between storage locations, which causes the read-around-ratio problem; and interaction between electrodes of the tube, which causes the cross-talk problem.

Performance of the standard type tubes in operating computer memories has not been as good as has been hoped. In order to use these tubes reliably, either the number of bits per tube has had to

---

[3] A. W. Holt and W. W. Davis, Computer memory uses conventional C-R tubes, Electronics (December 1953).

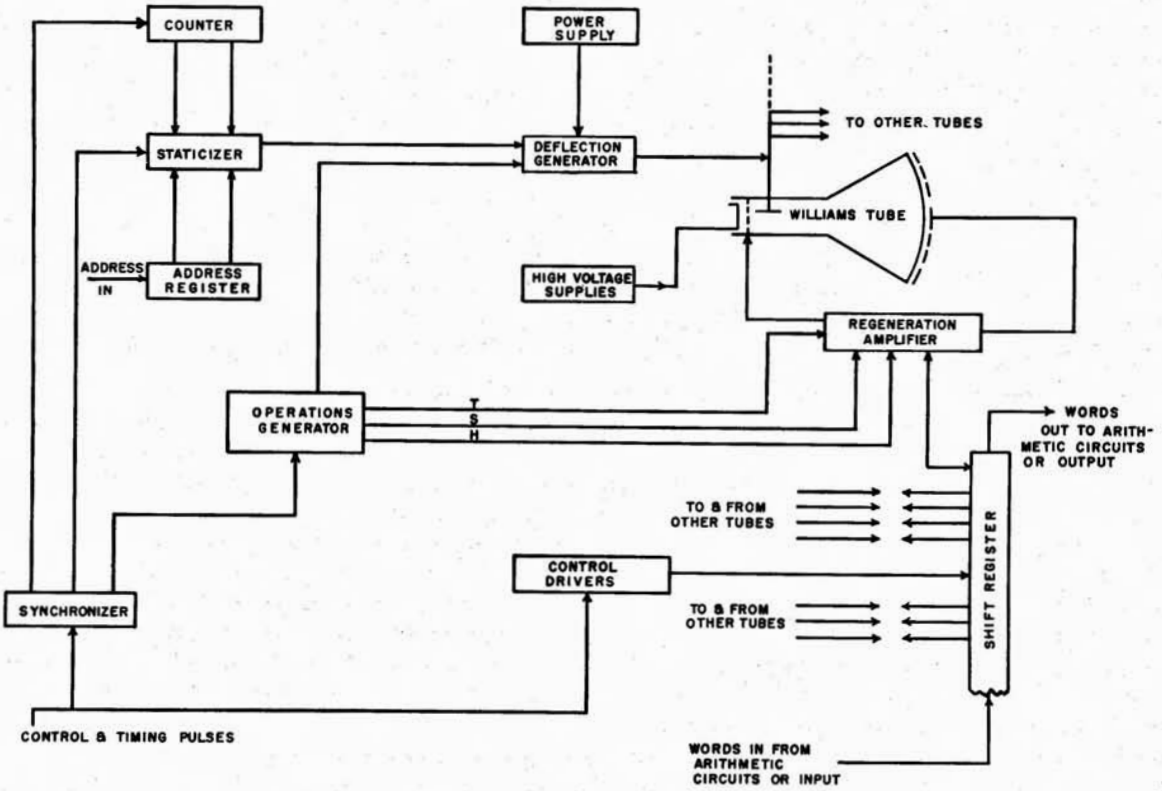

FIGURE 5.4. *Block diagram of electrostatic memory system.*

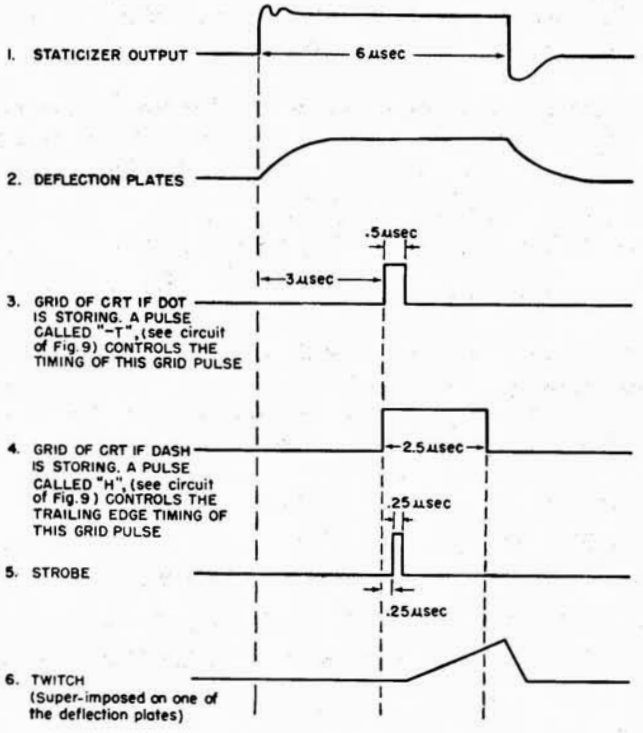

FIGURE 5.5. *Timing cycle: reading, writing, regenerating.*

be reduced below the number originally planned, or the ratio of references to the memory per regeneration has had to be limited. It has also been necessary to reject about four out of every five tubes received in production lots; acceptable tubes are therefore more expensive than had been expected.

Two lines of attack on the problem of improved tubes have been emphasized at the Electronic Computers Laboratory. The first approach was to try to obtain an improved tube which, although special purpose, would nevertheless be less expensive than intricate tubes proposed for other systems of electrostatic storage. The second approach was to try to circumvent the faults of available tubes. Both lines of attack have recently been quite successful.

On of the major problems is the presence of minute areas of low secondary emission ratio on the storage surface of the tubes. These spots are variously called blemishes, flaws, or (in England) phonies. Storage is impossible on these areas, and they must be avoided or the tube must be discarded. Unfortunately, the dodging of blemishes is difficult when many tubes are operated in parallel. In this case, computer operation may become unreliable unless very careful engineering prevents motion of the spot over the face of the tube.

Blemishes appear in about 80 percent of the tubes produced commercially, largely because extreme cleanliness is required to produce blemish-free surfaces. Blemish-free surfaces have been successfully produced in this and other laboratories when care is taken to maintain a high degree of cleanliness. Manufacturers of cathode-ray tubes interested in good surface quality are cooperating in the work of improving the tubes, as shown in table 1.

TABLE 1. Work at commercial tube plants

Manufacturer		Success	Direction of effort	Remarks
A--------	Smaller spot size; improved spot profile; improved deflection defocusing characteristics.	A likely design was produced.	Further testing of this model.	Design originated at NBS Computer Laboratory.
B--------	Smaller spot size; better beam profile.	Tube with half spot size, twice storage capacity of 5UP11.	Possible further testing of this gun in 3-in. bulb.	Finished sample tube delivered one week after initial discussion.
C--------	Over-all improved storage tube.	Final sample tubes were blemish-free at 2KV. Read around ratio was very high compared to that of standard CRT used.	Small-scale of production of further improved tubes, some refinements to be considered.	Work done under BuShips contract. Testing done at many cooperating government laboratories.
D--------	Experimental tubes for surface studies.	Study produced. No exceptional surfaces. Good bulb design for gun-signal plate shielding.	Small production with standard phosphor for tests in SEAC.	For experimental use only. Tube design originated at NBS Computer Laboratory.
E--------	Fine spot, prefocused gun.	----------	----------	In early stages. Work being done under BuShips contract.

The investigation of tube blemishes undertaken at NBS was first directed to the determination of the nature of the blemishes so that some control might be attempted and to correlating the location of these blemishes with the location of physical blemishes in the phosphor. Specially coated plates were checked in a demountable system under a microscope for this purpose. Correlation was low. An attempt to determine the size of a blemish was made by using India ink markings on a glass

plate. It was found that a positive blemish (defect in an area coated with ink) or a negative blemish (spot of ink on a preponderantly glass area) could be identified when one dimension was equal or greater than the beam diameter.

The first method of finding blemishes was to sweep over a rectangle on the face of the tube with a Lissajous pattern produced by sinusoidal X and Y deflection voltages and to apply the same deflection voltages to a monitor tube. The output at the signal plate of the storage tube was amplified and applied to the grid of the monitor. As long as there was no change in the storage surface, the brightness of the monitor was constant. However, when a blemish was struck by the electron beam, it caused a change in brightness on the monitor, and its relative position was marked by a bright or dark spot. The system located the blemishes but did not tell whether they would affect storage. This test was very sensitive but only qualitative.

A test called the Line-of-Dots test was developed at this Laboratory. It was not as sensitive but was able to determine whether a blemish would affect storage. Results of this test correlate very well with the performance of tubes in actual operation, and the test is now used as a specification test by a commercial manufacturer of a developmental Williams tube.

The Line-of-Dots test operates as follows: The beam is turned on and swept across the horizontal axis of the tube under test during the first half of the cycle. During the second half-cycle the beam is swept across the same line, but the beam is turned on only during 10 very short periods. During each one of these short periods the screen on the front of the tube picks up a positive-going signal which is very similar to the usual "dash" signal. These periods are nonsynchronous with the sweep, so that the entire line is interrogated during successive sweeps. When displayed on the vertical axis of a synchronized monitor, the envelope of the dash signals will show a notch or dip at a point where there is a blemish. The block diagram of this equipment is shown in figure 5.6.

Because blemishes afflicted 80 percent of the tubes, attempts were made to remove them. The use of infrared, ultraviolet, and electron bombardment failed to affect them. Rapping with a soft mallet helped to remove some. The best tool seemed to be a spark coil such as that used for checking vacuum leaks. The high-voltage brush discharge across the face of the tube charged the blemish and the surface alike, and the electrostatic repulsion removed many of the blemishes. This method was refined to the extent that 80 percent of a batch of tubes which were treated in this manner were acceptable.

Another aspect of the storage surface problem is the aging of the phosphor during use. A tube which has been in use for some time frequently develops a set of blemishes corresponding to the raster position. These raster burns usually do not affect storage seriously until the tube has had a reasonable life. Los Alamos, however, reports that in MANIAC it is necessary to replace tubes at an undesirable rate because of raster burns. Although our studies of this phenomenon are not complete, it does appear that the difficulty can be prevented by baking the phosphor at higher temperatures than are normally used. This indication has been tentatively confirmed by British work.

The second major difficulty encountered in Williams-type storage is the obliteration of the stored charge at one spot by the splash of the secondary electrons emitted from neighboring spots. The number of consultations of neighboring spots which may be made before regeneration is required is called "read-around ratio," "splash number," or "repetitive consultation number." Low read-around ratio may be improved by one or more of the following methods: by reducing the total number of spots in the tube and separating them further, by systematically forcing sufficient regenerations per useful consultation of the memory, by taking the ratio into account during coding of problems, or by improving the intrinsic properties of the tube in this respect.

Low read-around ratio appears to be mostly a function of gun design. Stray electrons not in the focussed beam, poor current distribution in the beam, and large beam diameter contribute to this effect. It was felt that a relatively small effort to obtain smaller spot size and a more rectangular current profile might result in considerably improved read-around ratios. Indeed, one company (manufacturer B, table 1) provided the laboratory with an improved tube a week after the first consultation. This tube, called the 51CUP11, had a gun in which the beam was heavily masked. It would store 1,024 bits in SEAC, whereas a 5UP11 usually would store only 512. A sample lot of these tubes was produced, but they were unsatisfactory for other reasons, and no more were ordered. However, the original tube was used in SEAC for about 1 1/2 years alongside 3KP and 5UP types.

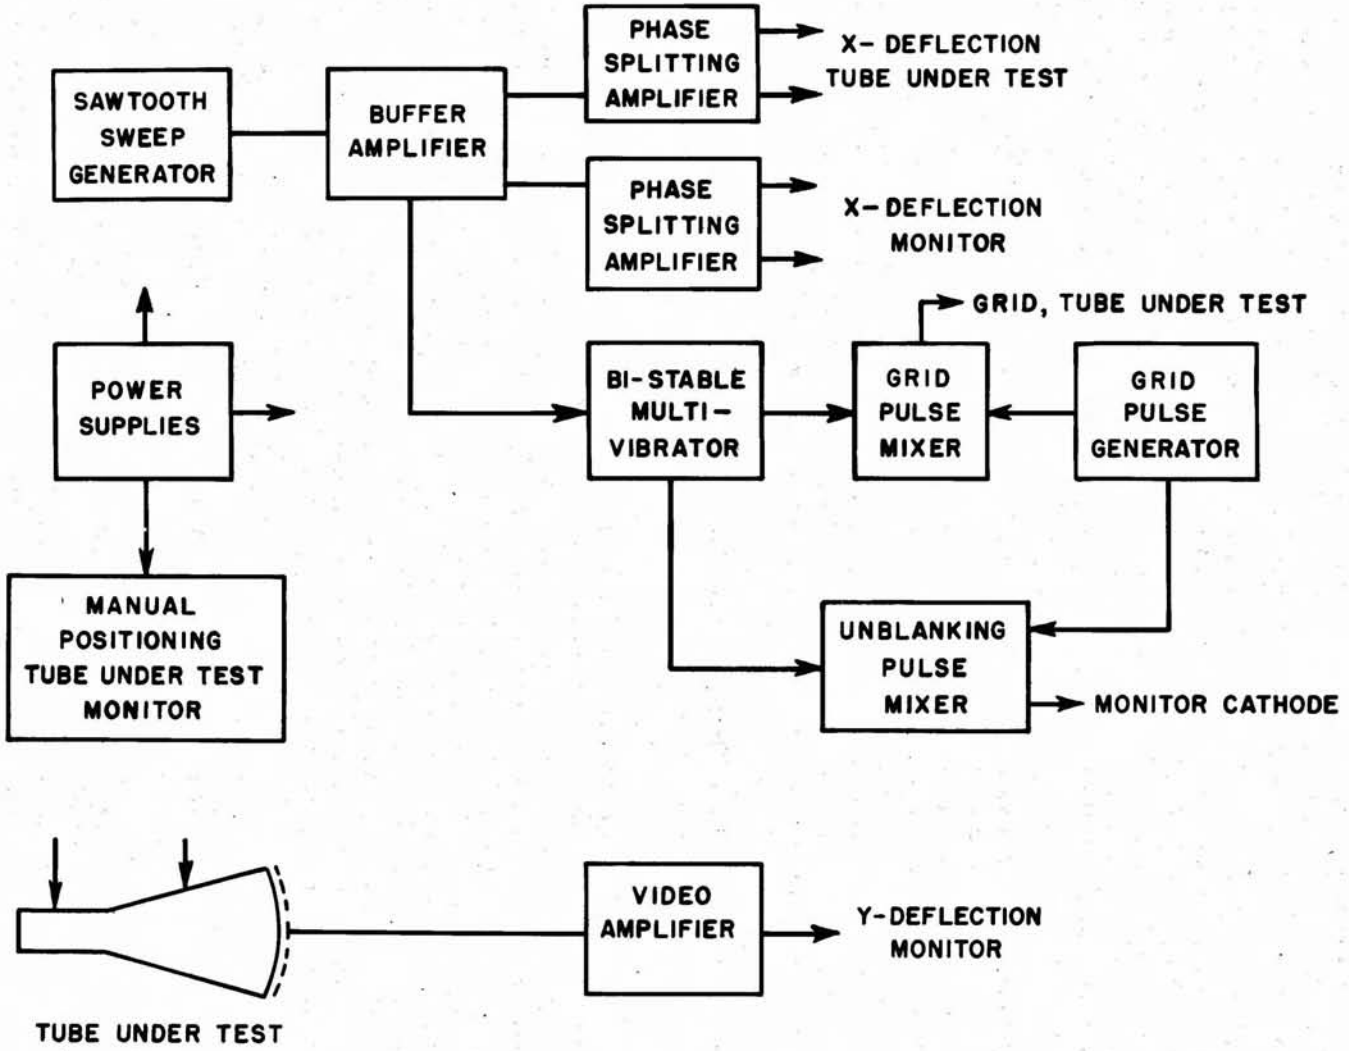

FIGURE 5.6. *Block diagram of line-of-dots blemish test equipment.*

Later the Bureau of Ships approached Manufacturer C to do development work on storage tubes. A contract was let that provided for the production of several trial designs to be followed by small-lot production of the best of these designs. In cooperation with the National Bureau of Standards, which was acting as technical consultant on the contract, the Argonne National Laboratories, the Institute for Advanced Studies, and a commercial laboratory tested these trial tubes. The purpose was twofold; first, to determine the best design of the developmental lot for use in the final sample production, and second, to determine how well the various tests used in the different laboratories agreed among themselves and with operation in a computer. The agreement among the tests in arranging the tubes in order of quality was quite good, considering the variety of test conditions. This order also agreed well with their performance in SEAC.

A second lot of tubes, all alike, are now in process of being tested in digital computers at various laboratories. These computers include ORDVAC at the Ballistics Research Laboratory, Aberdeen Proving Ground; ILLIAC at the University of Illinois; AVIDAC at Argonne National Laboratories; the Institute for Advanced Study's computer; and SEAC. Results to date are very encouraging. On the basis of these results, a conference held at the National Bureau of Standards led to agreement on specifications for a computer tube similar to this lot, and small-scale developmental production of the tube has begun.

This improved tube will have a gun in which the beam is heavily masked. Shielding is carefully designed to decrease the number of stray electrons. Distortion of the beam because of deflection is reduced by reducing the beam diameter in the deflection region and shaping the deflection plates. A storage surface of magnesium tungstate is used because of its good secondary-emission properties and because better control of screening in manufacture is expected. A 3 in. diameter was agreed upon because it was felt little could be gained by going to a larger tube, and considerations of space requirements and of safety favor as small a tube as possible. The specifications are such that no significant blemishes should appear at 2,000-v accelerating voltage, and the read-around ratio should be at least 256 in a field of 1,024 spots.

In the meantime, also under a Bureau of Ships Contract, another company (Manufacturer E) is developing a fine-spot gun with a magnetic fixed focus. The elimination of any electrode or electrode voltage is desirable because it reduces the number of parameters affecting the operation of the gun.

The third major fault with standard cathode-ray tubes when used as storage devices is the cross-talk among elements of the tube and between the signal plate and tube electrodes. For example, in the 3KP-5UP gun a deflection lead serves as a support for the grid cylinder, and the sharply rising grid pulse is capacitively coupled more to this deflection plate than to its mate. If the unbalanced coupled signal is large enough, the spot is in motion during writing or reading, which tends to spoil "dots." Again, the deflection plates in a 3-in. tube are rather near the signal plate. As these are pulsed with rapidly rising pulses of the order of 100 v and the normal signal level is about a millivolt, it is not hard to couple in signals that might cause the amplifier to block or ring so that false data are stored. A study was made of these effects, and it was found that grounding a conductive coating painted on the outside of the tube gave a substantial improvement in tubes without third anodes. A third anode gave the best shielding between gun and signal plate.

4. DIODE- CAPACITOR MEMORY

The basic storage element of the diode-capacitor memory is shown in figure 5.7. The connection E is used for both reading and writing, while the two diodes between A and D are used as a "squeezer" to connect the capacitor to the reading-writing circuits. During holding, both diodes are biased in their back direction. For example, A might be held at −4 v with respect to ground, and D held at +4 v. Then if the capacitor has a charge of, say, 2 v, both diodes will be biased in their back direction, and only small currents will flow into or out of the capacitor. For reading, suppose that both points A and D are forced to ground potential ("squeezed"). This will cause one or the other diode to conduct and a voltage will appear across the resistor R. If C is charged with 2 v of such polarity as to make its lower terminal more negative than its upper terminal, then when the squeeze occurs there will appear at E a pulse of −2 v, which then dies out with the time constant RC. This would be recognized by the reading circuits at E as a binary zero. If the polarity of the charge on C had been in the opposite direction, the squeeze would have produced a positive pulse, which would be recognized as the binary one. Thus the contents of the storage element has been read, but in

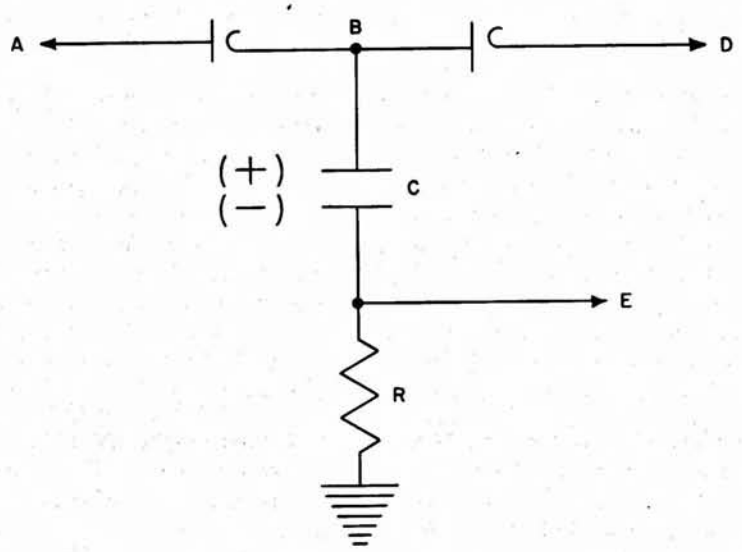

FIGURE 5.7. *Basic circuit of the diode-capacitor memory.*

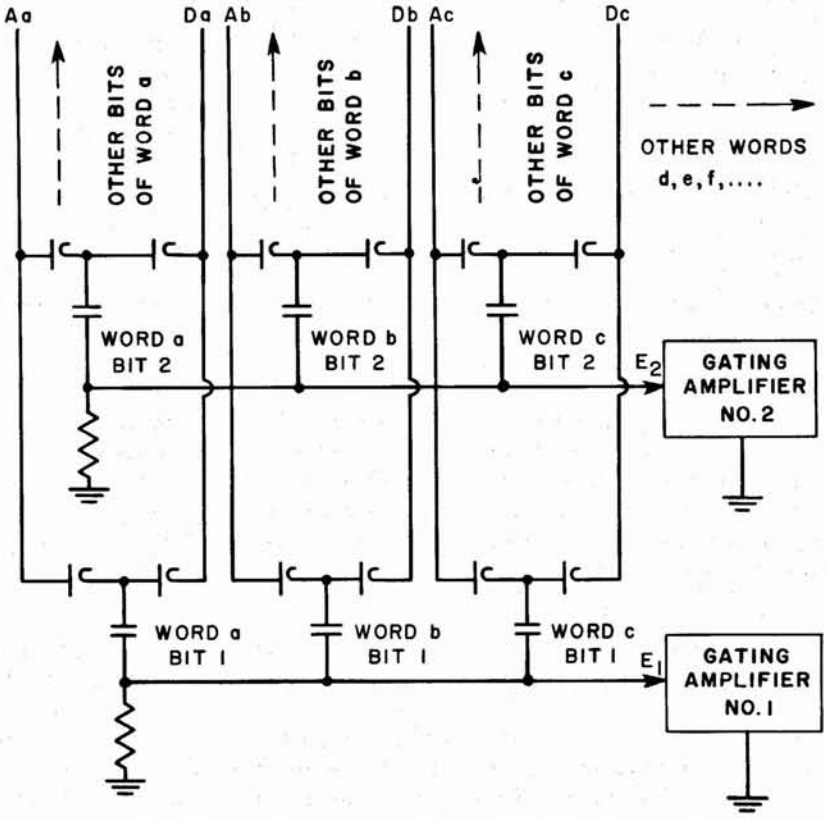

FIGURE 5.8. *Part of a full diode-capacitor memory, showing the manner in which the words, bits and gating amplifiers are interconnected.*

doing so it has been discharged (at least partially) and the information lost from the storage element. The information must be rewritten to continue the storage beyond the reading operation.

In order to write (or rewrite) information it is necessary merely to force the lead E to the desired state during the squeeze and hold it there until the squeeze is over. While A and D are at zero volts, suppose that E is forced to +2 v and held there at least until A and D are returned to their normal voltages of −4 and +4, respectively. Then the capacitor is left with a charge of 2 v, and upon the next squeeze it will produce a positive pulse at E. That is, a one is written. Obviously the opposite is equally possible: forcing E negative until the end of the squeeze will write a zero. Note that once A and D have returned to their normal voltages, the charge on the capacitor will be undisturbed by later changes of E, provided the magnitude of the voltage on E never exceeds 2v. Thus E can have other pulses on it, either positive or negative, and the charge stored on C will remain unaffected because both diodes will remain with backward bias. This is important for organizing many basic storage elements into an efficient memory assembly and is the reason for charging the capacitor to only +2 v while biasing the diodes twice as much.

In the discussion so far the diodes have been implicitly assumed to be ideal, having practically infinite forward conductance and practically zero backward conductance. The effect of finite forward conductance is modest: it will reduce somewhat the output pulse amplitude, and it will determine how long a writing pulse must last to charge the capacitor adequately. The effect of finite backward conductance, however, is critical. During the holding operation relatively long times will elapse, and even minute currents through the diodes will disturb the capacitor charge. The unit would gradually leak toward a condition of no charge on the capacitor, or even a condition in which the sign of the charge is reversed. Therefore, the permissible duration of the holding operation is determined by the rate at which the capacitor charge leaks through the back current of the diodes. Arbitrarily long storage of information is achieved through regeneration. Before the capacitor charge can change to a point where there is danger of losing the information, the memory control circuits read in a routine manner the contents of each cell and rewrite it accordingly.

What is needed at point E, then, is an amplifier which will sense the polarity of E during the early part of the squeeze period, together with a gate structure that will force E to the desired polarity during the latter part of the squeeze period. For reading or regeneration, E is forced to the same polarity that was read; for writing new information the polarity to which E is forced is independent of what was read but is determined by the new information being written. Such a gating amplifier is easy to construct. The amplification required is very modest, since its input is a pulse whose amplitude is of the order of 1 or 2 v. The gating can be accomplished with standard techniques, utilizing 2 or 3 vacuum tubes and several diodes.

An interesting point is that the memory in this form permits the ready incorporation of powerful self-checking features. The input to the gating amplifier is expected to be bipolar. That is, a definite pulse should be received every time a storage element is read. This pulse may be either positive or negative, depending on the information content of the storage element, but it should not be zero. If a signal approaching zero amplitude is received, it is a direct indication that the operation of that particular storage element is marginal. Thus, at the expense of some complication of the gating amplifier, it can be made to recognize three different input levels: acceptably positive, unacceptable, and acceptably negative. An unacceptable input need not, of course, be restricted to being very close to zero. A pulse of anything less than, say, one third of the normal amplitude might be sensed as being unacceptable. This would give a very prompt indication of incipient failure.

In order to achieve acceptable efficiency, it is essential that one such gating amplifier serve many basic storage elements. Figure 5.8 shows how this is done. The busses A and D are made common to all of the bits of a particular computer word, and a particular gating amplifier serves the same bit on each of many words. For 256 words of 40 bits each there might be 256 pairs of leads A and D, and 40 gating amplifiers. For reference to word b, the busses A_b and D_b would be squeezed to zero voltage, while all of the other pairs would be held at their normal values of −4 and +4 v. In this way each gating amplifier receives a pulse from its bit of the selected word, so the word is available in parallel at the gating amplifiers. These amplifiers can then write into this word, or rewrite it, without affecting the other words, since all diodes in the other words remain with backward bias as already described. After the squeezing busses on word b are returned to normal, any

other word may be referred to in the same way. In this way it is possible to have a fully parallel random-access memory. Regeneration is handled by having the memory control intersperse regeneration cycles between the computer access cycles. For the regeneration cycles, the words are read, one after the other, and rewritten to their former state.

At present the quantitative aspects of the regeneration problem appear to be the greatest limitation on this entire memory scheme. As a rough approximation consider the following argument: The rate of discharge of the capacitor during holding is proportional to I_b, where I_b is the back current of the diode at a voltage of about 4 to 6 v. Similarly, the rate of charging during writing and rewriting is proportional to I_f, where I_f is the forward current at something like 0.5 to 1 v. The safe holding time and writing time are inversely proportional to these rates, so the ratio of the permissible holding time to the writing time is I_f/I_b. This ratio of permissible holding time to writing time indicates how many writing operations can be done before it is necessary to come back and rewrite a particular bit. It is an approximate measure of the number of bits that can be served by one gating amplifier. With the inclusion of safety factors, reading time, and possible selection times, this figure comes in the range of about 0.1 (I_f/I_b) to 0.01 (I_f/I_b).

For actual diodes, this means that with the customary germanium whisker diodes, only some 32 to 64 bits can be served by each gating amplifier. It is, of course, possible to have multiple sets of gating amplifiers, but having many such sets would seriously increase the cost of the system. On the other hand, miniature selenium diodes give a much better figure of merit. It would appear possible to operate safely with 256 to 512 bits per amplifier. These diodes have much greater capacitance than the germanium, but the balanced construction of the squeeze circuit overcomes much of the difficulty. An ideal diode for this application is the new silicon junction diode, which has a simple thermally diffused junction. Laboratory models of these diodes have been able to withstand rather less back voltage than the germanium or selenium, but they have a fantastic I_f/I_b ratio. The back voltage that these units will stand is of the order of 20 v, which is entirely acceptable in this memory circuit. On the other hand the ratio I_f/I_b is even greater than some thermionic vacuum diodes. Only two such diodes have so far been available to this laboratory for test (through the courtesy of the Bell Telephone Laboratories), but they formed a basic storage element with writing times of a few microseconds and holding times of two to three seconds. These diodes appear capable of operating in a memory with 10,000 words per amplifier, with safety factors of 10 in the forward direction and 100 in the backward direction. Right now these units are rarities, but there is hope that they will be available in quantity and at reasonable cost in a few years.

A possibility that should be mentioned for the future is the use of capacitors which exhibit strong voltage-charge hysteresis. Such capacitors could be used in this system without requiring tight limits on their characteristics. This system would permit much looser specifications for the capacitors than present alternative proposals for their use. Using them in this system would eliminate the I_f/I_b restriction on the number of memory elements served by each gating amplifier. The specifications for the diodes could also be greatly relaxed, but there would be no decrease in the number of diodes needed.

The system described so far achieves reasonable efficiency for the gating amplifiers but requires a selection circuit capable of squeezing the appropriate pair of buses for a particular word. This could be accomplished by the customary diode matrix, but the usual form of such a matrix has large standby currents. In this memory the squeezing buses require relatively large currents; the resultant selection matrix is feasible but draws large amounts of standby power. To avoid this situation a selection matrix using transformers and diodes is used as shown in figure 5.9. This matrix has no standby power requirement, although it does require more input drivers than would be necessary with a multidimensional diode matrix. For the transformer-diode matrix, 2n inputs are required to select from among n^2 words. The matrix is made up of two sets of crossing buses (X and Y in fig. 5.9). At each crossing a diode and transformer are connected as shown. Normally all of the X buses are held at, say, +10 v, and all of the Y buses are held at −10 v. This puts backward bias on the diodes associated with each transformer, so no current flows through any transformer. If one X bus is dropped to −10 v, still no current will flow; but if simultaneously one of the Y buses is raised to +10 v, then just the one transformer at the crossing of these two buses will receive a signal. If X_2 is lowered to −10 v, and Y_1 raised to +10 v, the transformer secondary connected to buses A_c and D_c will squeeze the voltage on these two buses together. This will select the desired word.

other word may be referred to in the same way. In this way it is possible to have a fully parallel random-access memory. Regeneration is handled by having the memory control intersperse regeneration cycles between the computer access cycles. For the regeneration cycles, the words are read, one after the other, and rewritten to their former state.

At present the quantitative aspects of the regeneration problem appear to be the greatest limitation on this entire memory scheme. As a rough approximation consider the following argument: The rate of discharge of the capacitor during holding is proportional to I_b, where I_b is the back current of the diode at a voltage of about 4 to 6 v. Similarly, the rate of charging during writing and rewriting is proportional to I_f, where I_f is the forward current at something like 0.5 to 1 v. The safe holding time and writing time are inversely proportional to these rates, so the ratio of the permissible holding time to the writing time is I_f/I_b. This ratio of permissible holding time to writing time indicates how many writing operations can be done before it is necessary to come back and rewrite a particular bit. It is an approximate measure of the number of bits that can be served by one gating amplifier. With the inclusion of safety factors, reading time, and possible selection times, this figure comes in the range of about 0.1 (I_f/I_b) to 0.01 (I_f/I_b).

For actual diodes, this means that with the customary germanium whisker diodes, only some 32 to 64 bits can be served by each gating amplifier. It is, of course, possible to have multiple sets of gating amplifiers, but having many such sets would seriously increase the cost of the system. On the other hand, miniature selenium diodes give a much better figure of merit. It would appear possible to operate safely with 256 to 512 bits per amplifier. These diodes have much greater capacitance than the germanium, but the balanced construction of the squeeze circuit overcomes much of the difficulty. An ideal diode for this application is the new silicon junction diode, which has a simple thermally diffused junction. Laboratory models of these diodes have been able to withstand rather less back voltage than the germanium or selenium, but they have a fantastic I_f/I_b ratio. The back voltage that these units will stand is of the order of 20 v, which is entirely acceptable in this memory circuit. On the other hand the ratio I_f/I_b is even greater than some thermionic vacuum diodes. Only two such diodes have so far been available to this laboratory for test (through the courtesy of the Bell Telephone Laboratories), but they formed a basic storage element with writing times of a few microseconds and holding times of two to three seconds. These diodes appear capable of operating in a memory with 10,000 words per amplifier, with safety factors of 10 in the forward direction and 100 in the backward direction. Right now these units are rarities, but there is hope that they will be available in quantity and at reasonable cost in a few years.

A possibility that should be mentioned for the future is the use of capacitors which exhibit strong voltage-charge hysteresis. Such capacitors could be used in this system without requiring tight limits on their characteristics. This system would permit much looser specifications for the capacitors than present alternative proposals for their use. Using them in this system would eliminate the I_f/I_b restriction on the number of memory elements served by each gating amplifier. The specifications for the diodes could also be greatly relaxed, but there would be no decrease in the number of diodes needed.

The system described so far achieves reasonable efficiency for the gating amplifiers but requires a selection circuit capable of squeezing the appropriate pair of buses for a particular word. This could be accomplished by the customary diode matrix, but the usual form of such a matrix has large standby currents. In this memory the squeezing buses require relatively large currents; the resultant selection matrix is feasible but draws large amounts of standby power. To avoid this situation a selection matrix using transformers and diodes is used as shown in figure 5.9. This matrix has no standby power requirement, although it does require more input drivers than would be necessary with a multidimensional diode matrix. For the transformer-diode matrix, 2n inputs are required to select from among n^2 words. The matrix is made up of two sets of crossing buses (X and Y in fig. 5.9). At each crossing a diode and transformer are connected as shown. Normally all of the X buses are held at, say, +10 v, and all of the Y buses are held at −10 v. This puts backward bias on the diodes associated with each transformer, so no current flows through any transformer. If one X bus is dropped to −10 v, still no current will flow; but if simultaneously one of the Y buses is raised to +10 v, then just the one transformer at the crossing of these two buses will receive a signal. If X_2 is lowered to −10 v, and Y_1 raised to +10 v, the transformer secondary connected to buses A_c and D_c will squeeze the voltage on these two buses together. This will select the desired word.

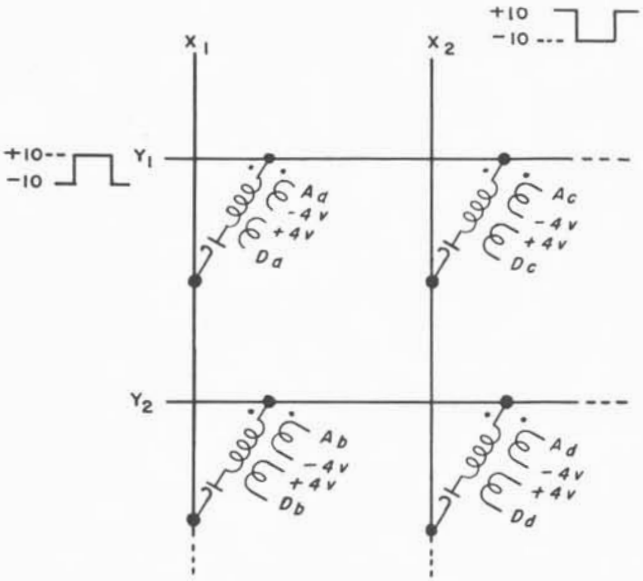

FIGURE 5.9. *A matrix switch using the transformer and-gate.*

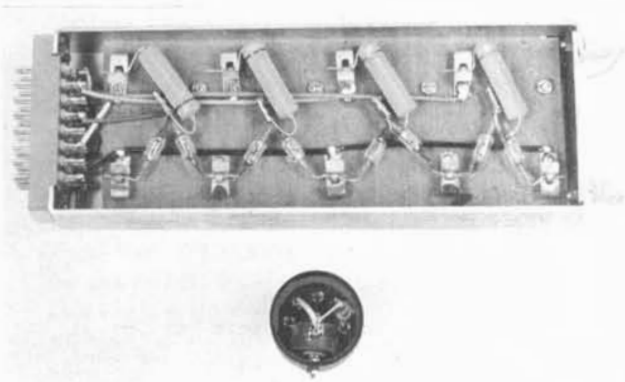

FIGURE 5.11. *Eight-bit memory package and transformer and-gate.*

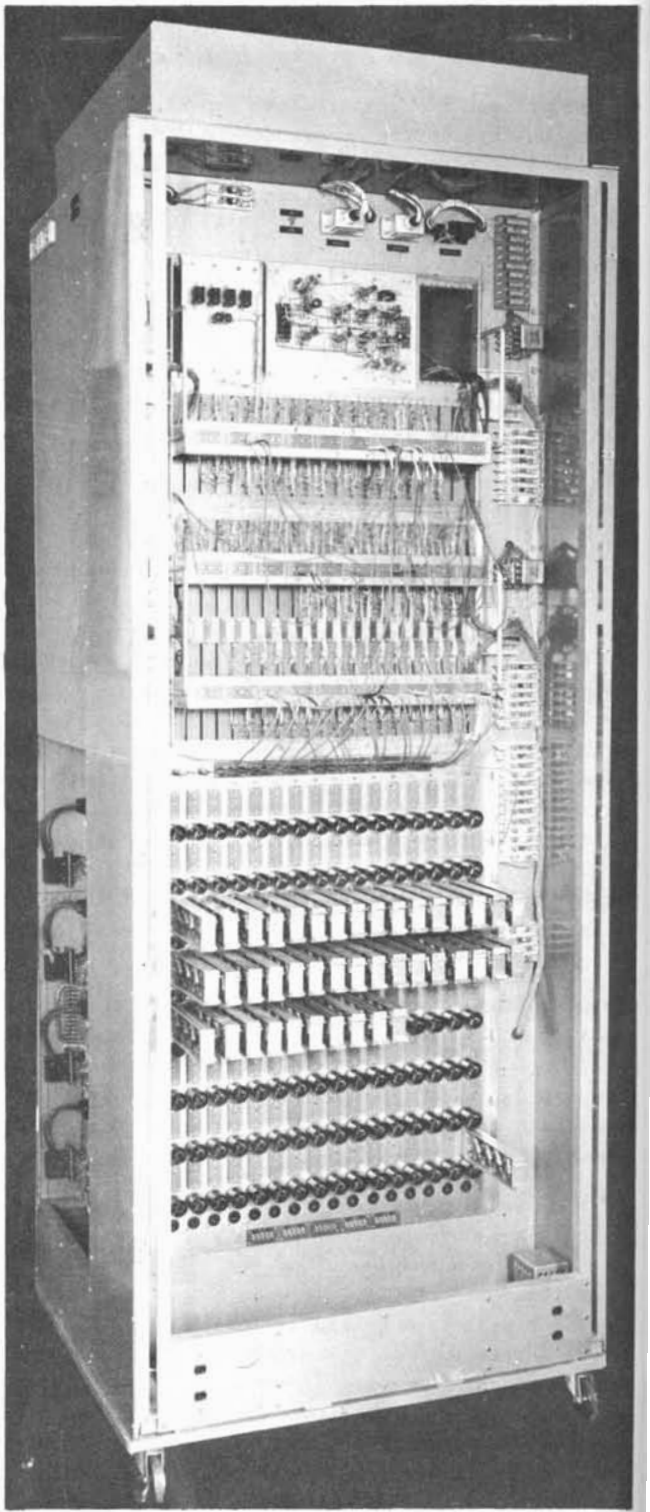

FIGURE 5.10. *Front view of diode-capacitor memory rack.*

After the individual elements of the system were tested, a laboratory model was built containing 16 words of 4 bits each. With this model in its final form several successful lengthy tests of storage were carried out. On five occasions the unit was left running for 3-day periods and found to have the correct information at the end of that time.

Results with the laboratory model have been sufficiently promising to make it desirable to test something more nearly approaching a full memory. A prototype is now in construction which will be attached to the SEAC and tested in the same way that the electrostatic memory prototype has been tested. For the diode-capacitor memory, the unit is designed for a capacity of 256 words of 45 bits each, but only 128 words of 8 bits each are being built as a start. Since the words have only 8 of the customary 45 bits, it will not be possible to operate the SEAC exclusively from this trial memory unit. However, since the SEAC can operate from both the acoustic and the diode-capacitor memories in integrated fashion, it is possible to do extensive testing of the new memory by using test routines stored in the acoustic memory. If all goes well with these tests, the memory will undoubtedly be expanded to a useful size.

Figure 5.10 shows this unit as it is now being assembled, and figure 5.11 is a view of the 8-bit memory package. It will be noticed that very little attention has been paid to compactness in this construction. Quite the opposite, the units have been deliberately separated to permit access to them during experimental runs. Bracketing estimates on a full-scale memory assembly indicate that 1,000 words could be packaged in 20 to 50 ft^3.

In describing the operation of the memory no mention was made of the access rate that can be achieved. That is because the access rate is primarily limited not by the memory elements but by the external circuitry. The characteristics of the diodes in the memory unit determine the ratios that were discussed, but within wide limits the operating rate can be selected by selecting the capacitor size. This generalization becomes more limited if diodes are used which have large capacitances in themselves, such as the selenium diodes, but for low-capacitance diodes the generalization is reasonable. In the experimental equipment being built the active part of the basic cycle will be a 3-μsec period, during which reading and writing occur. This cycle is repeated every 6 μsec, the remaining 3 μsec being used for recovery of the transformers in the selection matrix.

The diode-capacitor memory has several advantages. An especially nice one is that there are no very weak signals or sensitive leads. The minimum signal is of the order of a volt across a few hundred ohms, and all of the selection is truly digital. There are no analog voltages to be derived, and all characteristics of the materials are bounded on only one side. That is, there is no limit on the upper end of the diode characteristic: it does not have to be matched to the other diodes in the circuit. The memory is very rapid, access to random information is possible at well over 100,000 words per second, and it can be very rugged in construction where this is an important attribute.

On the other hand, it has several disadvantages too. A large number of diodes is required for a large memory. A memory of 25,000 bits requires 50,000 diodes, and it is an open question whether 50,000 diodes will give reliable operation even when the design allows them wide tolerances. Still, three years' experience with some 15,000 diodes in SEAC indicate that such operation is not entirely out of line. In addition, schemes have been worked out for rapid maintenance testing of such a large memory by a form of marginal checking which should permit replacement of drifting units before they cause trouble in computing. However, so many individual elements to be assembled will of necessity make for higher fabrication cost. Present estimates indicate that for the same capacity the cost would be approximately twice that of a mercury acoustic memory or a Williams-type electrostatic memory (the units constructed at NBS indicate roughly equal cost for these two types).

It would appear at present as though the proper balance among cost, performance, and serviceability is something that only more experience can indicate. It is hoped that the prototype construction will provide the experience necessary to determine this balance.

5. FACTORS AFFECTING MEMORY DEVELOPMENT

The *reliability* of a memory system is hard to define, since the reliability of a completed piece of equipment is so much a function of the engineering. In this laboratory it is believed that the

simplicity of the device is related in a most fundamental way to the reliability and this concept has proved to be a useful yardstick. Of the three memory systems discussed here, the diode-capacitor is the simplest, the acoustic is median, and the Williams is the most complicated. On the basis of sensitivity to voltage changes and external interference, the diode-capacitor is least sensitive, the acoustic is median, and the Williams is most sensitive. The log book for the acoustic memory in operation with SEAC shows occasional runs of nearly a week without adjustment or complaint from the mathematicians; the longest run on record for the Williams (with SEAC) is 11 hours. The most recent model of diode-capacitor memory has made a run of 269 hours without error in the laboratory. These figures should not be taken as indicative of the limits of reliability, since the engineering is constantly being improved on all these memories.

The *access rate* is more easily defined and is given in terms of the maximum number of randomly selected words per second. These speeds are, of course, variable over a considerable range at the designer's choice. The acoustic memory, being serial, is basically slower than the parallel memories, but it is perhaps inadequately realized today that high computation rates can still be accomplished using the acoustic memory, by careful organization of the machine. How fruitful this approach is, of course, depends greatly on the type of service for which the computer is designed. In SEAC, the acoustic memory is designed to deliver an average of about 6,000 random words per second, but with optimum programming it can deliver about 20,000 words per second. The figure for the Williams memory in SEAC is about 20,000 random words per second, although systems have been built that are as high as 60,000. The basic speed limitation is the amount of current which can be supplied to the storage spot without causing serious interaction between spots; a circuitry problem is to increase the speed of deflection and still maintain accuracy of final point. The model B prototype of the diode-capacitor memory is designed for 80,000 random words per second in laboratory operation and may be converted to 125,000. In operation with SEAC, however, the speed will be 20,000 words per second, since this is the fastest rate at which the machine can accept information. In order to change the speed, one need only change the capacitor; the limits in speed for the diode-capacitor memory are not being pushed as yet.

Figures for the relative cost of these three memories can only be given for the initial installation, although it is apparent that maintenance costs ought to be considered. The price for acoustic memory used in SEAC is about \$1.50 per binary digit, including all access circuitry, amplifiers, and hardware. The cost of the Williams memory in SEAC is about the same, although extra engineering features, such as extensive shielding and elaborate power supplies, would raise the cost. The diode-capacitor memory will probably cost in the range of 2 to 4 dollars per bit, more than half of the cost being in the price of diodes.

Briefly comparing the relative size, weight, and ruggedness of these three memory systems, they are all more or less in the same class of size, running between 20 and 100 ft^3 for 50,000 bits of memory proper. Compact etched-circuit designs have been made for the diode-capacitor memory. In weight, the memories are again in about the same class, although the acoustic and Williams need, in general, more mechanical structure. In ruggedness, however, the Williams memory appears to be basically less desirable than the other two systems, since the cathode-ray tube and gun is a relatively fragile device. Under severe vibration the deflection accuracy is almost certain to be adversely affected.

It should be apparent that each of these memories has its own attractive features. It may be, for example, that the Williams will turn out to be most suitable for very high-speed fixed installations, the diode-capacitor most suitable for very high-speed mobile machines, and the acoustic best for serial high-speed general-purpose equipment. Development and evaluation of all these memories are therefore continuing at this laboratory, with the SEAC acting in a most useful capacity as a tool for evaluation.