DATA
STORAGE
HANDBOOK

INSTRUMENTS AND CONTROL SY[redacted]EMS

Rimbach Publications
Division of Chilton Company
845 Ridge Avenue
Pittsburgh, Pennsylvania 15212

Printed in U.S.A.

CONTENTS

COMPUTER STORAGE TECHNIQUES

First in a series of articles describing computing storage technology, this installment introduces the general nature and types of storage and its measurement, as well as its functional relation to the computer. Subsequent articles will consider storage techniques now used in digital computers.

G. A. OLIVER
Burroughs Corporation
Research Center

STORAGE, which is preferred to the term "memory," is the property of a device that permits it to introduce time displacement into the transfer of information. A storage device can accept a quantity of information at some instant and deliver a replica of that information at a later instant.

Storage performs the same function as a scratch pad used by a human when he computes a numerical quantity. Storage also is used to retain the sequence of instructions that the computer will perform. A third use of storage is to alter information transfer rates; for example, information coming from some slow-speed input device would be stored and would later be delivered at a higher speed to the computing system. Storage employed in this fashion is termed "buffer storage."

Fig. 1 depicts an idealized computing system. Note first that almost every major division of the machine contains some type of storage. Further, use of storage elements is not confined to the limits of the computer frame and cabinet.

TERNAL if it is interior to the computer, and *EXTERNAL* otherwise.

The computer's *primary storage* unit constitutes the major portion of internal storage. The main storage is characterized by a medium-to-large capacity and low access-times relative to the operating cycle of the computer. As the primary storage holds both data and instructions, communication with the "arithmetic and control sections" must be provided. Magnetic cores, drums and discs have found wide acceptance in this application.

The cost of large, low-access-time memories is comparatively high. For this reason, as indicated in Fig. 1, the main storage can consist of two or more storage units in cascade. A small, low-access-time storage (e.g., magnetic core) is provided to permit fast computer operations. A larger storage with longer access (e.g., magnetic drum) is then provided to load the fast memory at relatively infrequent intervals. The result of this arrangement is a storage unit of about the speed of the small

TABLE I: SUMMARY OF STORAGE SYSTEM TECHNIQUES AND

	Punched Cards	Punched Paper Tape	Magnetic Tape	Magnetic Drums	Electro-Static	Magneto-Strictive Delay
Type	Static	Static	Static	Dynamic	Static	Dynamic
Access Mode	Sequential	Sequential	Sequential (Random)	Combinational	Random	Combinational
Access Time	0.08 sec. per card	0.01 sec.	0.01 sec. (Seq.) 1-2 min. (Random)	5-20 msec.	10 μsec.	0.5-2 ms.
*Capacity	—	—	6×10^6-30×10^6 bits per reel	2×10^3-1×10^6 bits	5×10^3-5×10^4 bits	5×10^3-5×10^4 bits
Readout	Non-Destructive	Non-Destructive	Non-Destructive	Non-Destructive	Destructive	Non-Destructive
Volatile	No	No	No	No	Yes	Yes
Erasable	No	No	Yes	Yes	Yes	Yes
Cost	\$0.000002 per bit (\$$2 \times 10^{-6}$)	\$0.00002 per bit (\$$2 \times 10^{-5}$)	\$0.000002 per bit (\$$2 \times 10^{-6}$)	\$0.02 per bit (\$$2 \times 10^{-2}$)	\$1-\$2 per bit	\$0.50-\$1.50 per bit

*For Internal Storage Only

Terminology

The two fundamental measures of storage are capacity and access time.

CAPACITY: The upper limit on the amount of information a storage unit can retain. Capacity is expressed in bits (binary digits) or, alternately, words or characters when the number of bits per word or character is known.

ACCESS TIME: The interval between interrogation of the storage unit and receipt of the specified information.

VOLATILITY: Ability to hold the information when power is lost. Some storage media, such as electrostatic delay-lines, lose information when power is lost.

ERASABILITY: Ability to be reused. Storage media that permit reuse by erasing presently stored information include magnetic tape, magnetic drums, etc. Non-eraseable storage is punched paper tape, punched cards, etc.

DESTRUCTIVE READOUT: Ability to hold the information after interrogation. Interrogation of stored information causes some storage media (electrostatic, magnetic core) to return to a neutral state and the information to be lost.

COST: Frequently expressed as dollars/bit.

PHYSICAL PROPERTIES: Size, power, cooling requirements, etc.

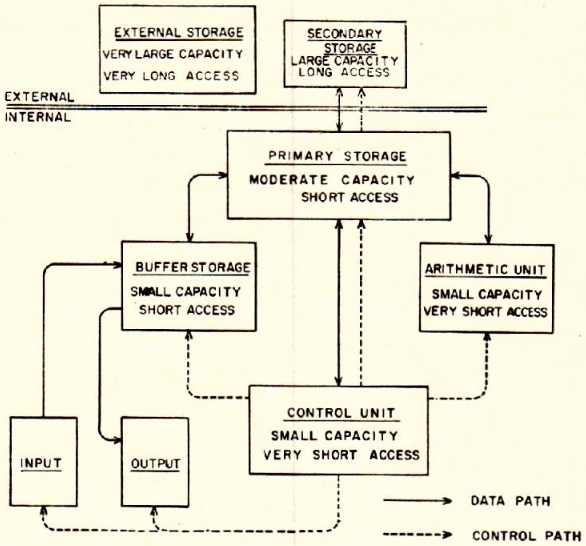

FIG. 1. EXTERNAL STORAGE features large capacity and long access time; internal storage features smaller capacity and access time.

high-speed store, but with the capacity of a much larger store.

Storage requirements in the "arithmetic and control units" are small in comparison to the main storage—usually not more than three to five words. To avoid slowing the machine operation, however, very short access time is desirable. For these reasons and others (e.g., depending on information transfer required for a particular machine operation), the storage elements in the arithmetic and control sections are organized into several single-word storage units. These single-word stores are called *registers*. The description "zero access" has been applied to registers because of their very low access time. Registers usually comprise flip-flops or special magnetic-core circuits, although the same effect can be obtained by organizing data properly on a magnetic drum.

THEIR CHARACTERISTICS

Magnetic Core	Vacuum Tube	Magnetic Thin Film	Random Access Magnetic (Discs)
Static	Static	Static	Dynamic
Random	Random	Random	Combinational
2-10 μsec.	1 μsec.	0.1-0.5 μsec.	0.5-1.0 sec.
5x10^3-5x10^5 bits	50-100 bits	5x10^3-5x10^5 bits	40x10^6 bits
Destructive	Non-Destructive	Destructive	Non-Destructive
No	Yes	No	No
Yes	Yes	Yes	Yes
\$0.50-\$4.00 per bit	\$10/bit	Not Yet Established	\$0.002 per bit (\$2x10^{-3})

The function of *buffer storage*, as mentioned, is to match the high speeds of the computer with the lower speeds of the peripheral input/output units. As large capacity usually is not required, flip-flops or small magnetic core-store can be employed. Buffers that have a capacity in excess of one word generally are arranged to deliver information precisely in the order received. (In some larger computing systems, however, the programmer is given control of the order in which data will be taken from the buffer.)

The *secondary* (*auxiliary*) *storage* is in the awkward position of being external to the computer but of being controlled by the computer. Secondary storage provides extremely large-capacity, long-access, low-cost storage required for files, program libraries, and the like. Magnetic tape usually is employed; magnetic discs are used in some newer equipment.

We come, finally, to *external storage*. This type of storage is neither connected to or controlled directly by the computer. Media used include punched card, paper tape, magnetic tape, photographic film, printed pages, etc. Access is long because the media must usually be selected from some type of file and manually inserted into an appropriate reading device. Only the lowest cost media are suitable for this class of storage because of the very large capacity required.

Primary Storage

A block diagram of a typical *primary storage unit* is shown in Fig. 2. The individual storage elements are organized into a large, interconnected set of one-word registers (storage locations). This set of storage elements is the heart of the unit; the remaining equipment serves only to gain access to a given location and to operate on the information contained in that location.

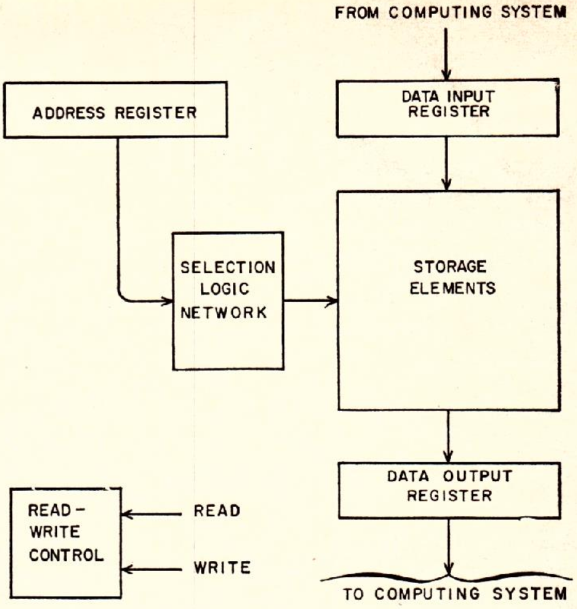

FIG. 2. PRIMARY storage units.

Static or Dynamic

Although many physical phenomena have found application as storage elements and a large number of different hardware arrangements have developed, all storage phenomena can be classed as being either static or dynamic. *Static* storage is that from which the stored information is continuously available and is fixed in space; the flip-flop and the magnetic core storage are examples of this type. Information held in *dynamic storage,* on the other hand, is moving in time and is not continuously available. The magnetostrictive delay line and the magnetic drum are examples of this class of storage. The type of storage used controls the nature of the entire system, as will be seen.

Access to a given word of data is achieved by inserting the address of the selected location into the address register (Fig. 2). Signals from the address register activate the selection network which, in turn, connects the specified group of storage elements to the input and output register.

Storage Access

There are three principal types of access—random, sequential, and combinational. Random access implies that the access time for a word in storage is independent of the location of a previously addressed word.

Some storage techniques (e.g., magnetic core) provide random addressing capabilities, as shown in the second line in Table 1.

The term *sequential access* means that the desired data are located by searching through all locations, beginning with the one presently accessible, until the desired location is found. Access time thus depends on the location of the selected word in the storage. Magnetic tape is a good example of a technique having this type of access.

Combinational access is intermediate to sequential and random and is an access mode typical of dynamic storage techniques. Consider a magnetic drum having a large number of tracks in parallel, each of which contains several words stored *serially*. Access to a particular word requires specification of a track and an interval in time. These are, respectively, random and sequential access modes. Hence it is called combinational.

Note that if the computing system is arranged so that it never calls for information from the storage at a rate in excess of that established by the longest possible combinational access time, access is effectively random. This is why the term "random-access" is applied to certain large-capacity magnetic-disc storage units.

Once access to a selected location is obtained, information can be written into or read out of that location. Again there are three possible modes of operation—(1) parallel, in which all of the bits of the entire word are operated upon simultaneously; (2) serial, in which the word is handled by bits singly and successively; and, (3) combinational, in which, for example, a word can be treated serially by digit and parallel by bits comprising the digit.

References

E. M. Grabbe and D. E. Wooldridge, *Handbook of Automation Computation and Control,* Vol. 2, John Wiley and Sons: New York, N. Y., 1959.

R. K. Richards, *Arithmetic Operations in Digital Computers*, Van Nostrand: Princeton, N. J., 1955.

R. K. Richards, *Digital Computer Components and Circuits*, Van Nostrand: Princeton, N. J., 1958.

J. P. Eckert, Jr., "A Survey of Digital Computer Memory Systems," *Proc. I.R.E.* 41 (10) 1391-1406 (1953).

PUNCHED TAPE

HENRY THOLSTRUP Friden, Inc.

Punched-tape dimensions and hole codes have been standardized, but different program codes (fixed sequential, block address, word address, tab sequential) are in use. A universal program code is being considered by Electronic Industries Association.

PUNCHED TAPE is the most popular method for numerical control of machining and metal-working operations.

The metal-removing industry, or machine-tool industry, has perhaps made the most use of numerical control to date, and has found that it can be applied to many processes and operations in industry such as (1) drilling, (2) boring, (3) milling, (4) turning, (5) shaping, (6) punching, (7) flame cutting, (8) welding and so on.

The basic considerations in use of punched tape include (1) tape dimension, (2) hole code and (3) program code. The first two have been standardized by the Electronic Industries Association, but several program codes are in common use.

Standards

Electronic Industries Association (11 West 42 St., New York 36, N. Y.) has issued the following "One-Inch Perforated Paper Tape Standard, RS-227"—issued Oct. 1959:

"The unpunched tape shall have an overall width of 1.000″ ±0.003″ after the tape has been conditioned to 73° F ±3.5° F and 50% ±2% relative humidity for a period of 24 hours (Fig. 1).

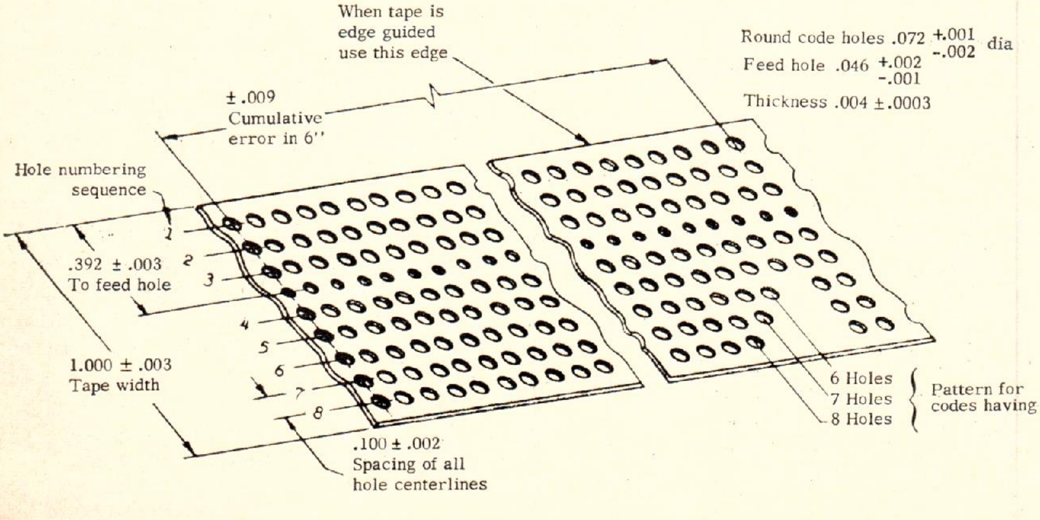

FIG. 1. DIMENSIONAL STANDARDS for punched tape. Track 1 is that track nearest the tape transport desk. (From National Aircraft Standard, tentatively accepted by the Aerospace Industries Association of America, Inc.)

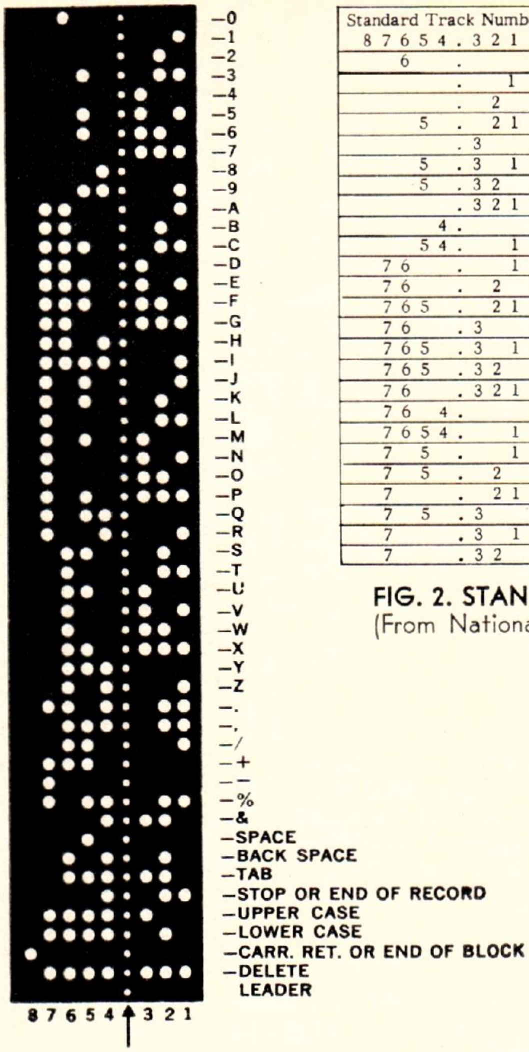

Standard Track Numbers 8 7 6 5 4 . 3 2 1	Digit or Letter Codes	Standard Track Numbers 8 7 6 5 4 . 3 2 1	Digit or Letter Codes
6 .	0	7 5 . 3 2 1	p
. 1	1	7 5 4 .	q
. 2	2	7 4 . 1	r
5 . 2 1	3	6 5 . 2	s
. 3	4	6 . 2 1	t
5 . 3 1	5	6 5 . 3	u
5 . 3 2	6	6 . 3 1	v
. 3 2 1	7	6 . 3 2	w
4 .	8	6 5 . 3 2 1	x
5 4 . 1	9	6 5 4 .	y
7 6 . 1	a	6 4 . 1	z
7 6 . 2	b	7 6 4 . 2 1	. (period)
7 6 5 . 2 1	c	6 5 4 . 2 1	, (comma)
7 6 . 3	d	7 6 5 .	+ (plus)
7 6 5 . 3 1	e	7 .	- (minus)
7 6 5 . 3 2	f	6 5 . 1	/ (check)
7 6 . 3 2 1	g	7 6 5 4 . 3 2 1	Delete
7 6 4 .	h	8 .	End of Block
7 6 5 4 . 1	i	4 . 2 1	End of Record
7 5 . 1	j	5 .	Space
7 5 . 2	k	6 4 . 2	Back Space
7 . 2 1	l	6 5 4 . 3 2	Tab
7 5 . 3	m	7 6 5 4 . 3	Upper Case
7 . 3 1	n	7 6 5 4 . 2	Lower Case
7 . 3 2	o		

FIG. 2. STANDARD hole code. Presence of a track number indicates a hole. (From National Aircraft Standard, tentatively accepted by AIA.)

'The unpunched tape shall have a thickness of 0.004″ ±0.0003″ after the tape has been conditioned to 73° F ±3.5° F and 50% ±2% relative humidity for a period of 24 hours.

"The tape shall be used for recording six, seven or eight levels of information (tracks) across the tape. Track is synonymous with channel, or level, and runs along the tape. Row is across the tape.

"When the tape is edge guided, the guided edge of the tape shall be parallel with the longitudinal centerline of the feed holes and located 0.392″ ±0.003″ from the feed-hole centerline on the three-level side of the tape.

"The feed holes in the tape shall be round with diameter of 0.046″ plus 0.002″ minus 0.001″.

"The code holes in the tape shall be round with a diameter of 0.072″ plus 0.001″ minus 0.002″.

"All holes punched in the tape shall nominally center on the true intersection of longitudinal and perpendicular transverse centerlines spaced 0.100″ apart. Tolerances on locations of code holes in any one transverse row, relative to the center lines of the feed hole in that row, shall be ±0.002″ in the transverse direction and ±0.003″ in the longitudinal direction. Tolerance on distance between centers of adjacent feed holes shall be ±0.001″. Accumulated longitudinal errors between centers of feed holes shall not exceed ±0.009″ within spans of 0.9″ to 6″.

"The measurement of all dimensions of the punched paper tape shall be made with the punched paper tape conditioned to the same environmental conditions under which it was punched."

Fig. 2 shows the standard digit or letter 8-hole code.*

Standardized Programming

Another phase of numerical control that the Electronic Industries Association committee is working on is a standard format of the punched-tape program. The objective is to permit the tape prepared for use on one machine to be used on other machines—that is, a tape prepared for one two-axis numerically-controlled (NC) machine should work on all two-axis NC machines.

Four methods of programming are shown in Fig. 3 for the numerical control tape prepared for the same part. The four systems are: (1) Fixed Sequential, (2) Block Address, (3) Word Address, and (4) Tab Sequential. Also, a format is shown for a Universal Tape which would work with a control system on a machine designed for any one of the previously-shown systems. It is not near a final state yet, but is at present receiving discussion within the EIA committee on numerical control.

Principal objection to the Universal Program is that it requires many more code holes (from 11½% to 31% more) than are required for the other systems. This is an important factor when considering that a program tape for an average part varies from 3′ to 50′. Also, more time to prepare the tape will affect the cost of the part to be produced.

---

* Eight-channel tape is now considered "standard" by most users owning to its greater capacity than 5-, 6-, or 7-channel tape.

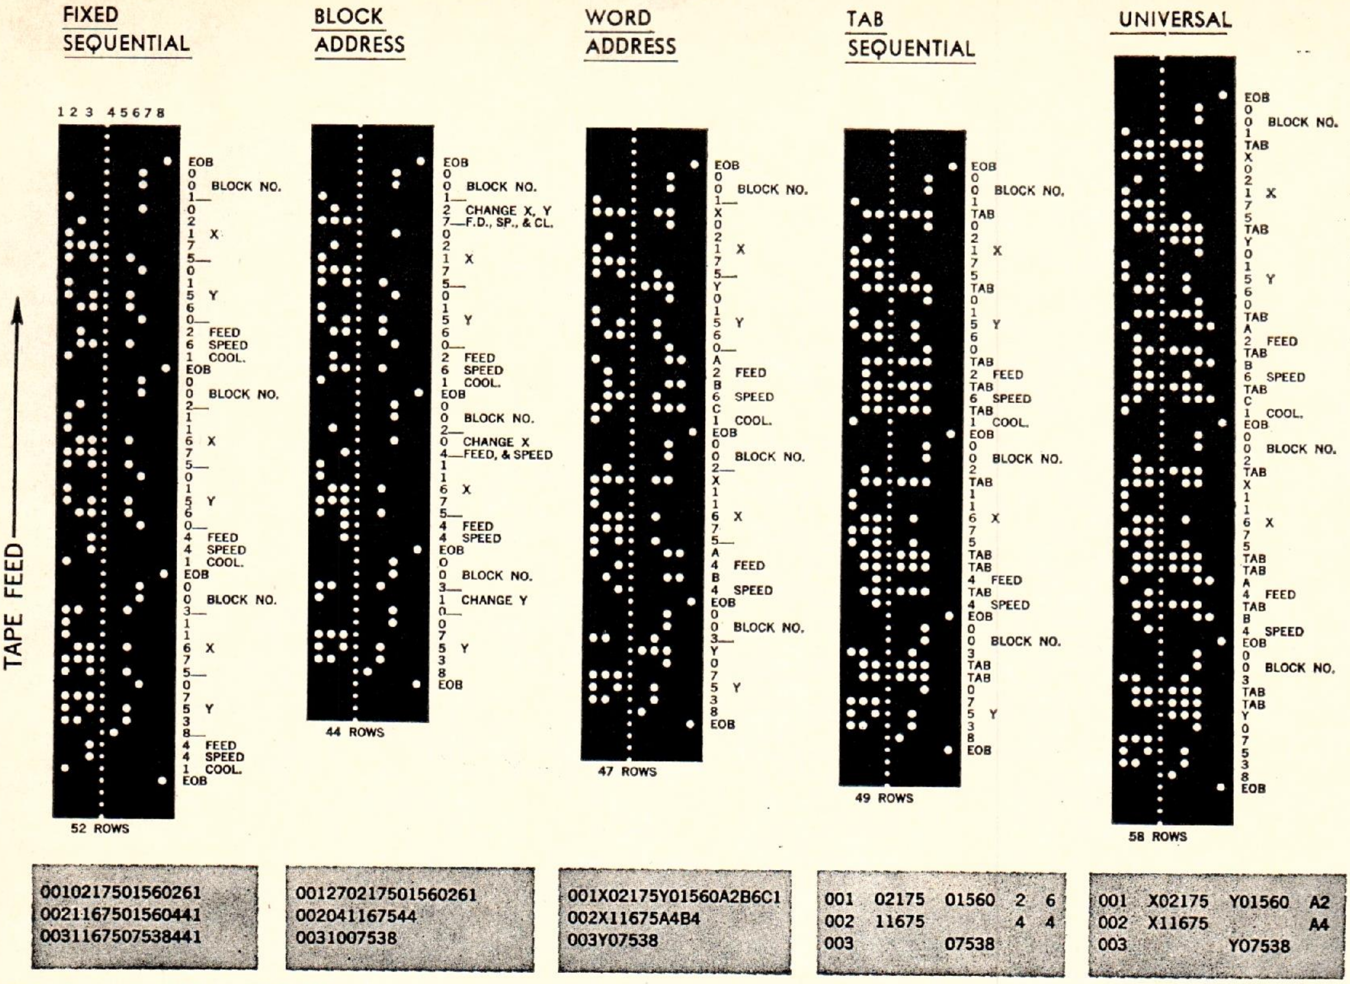

```
0010217501560261
0021167501560441
0031167507538441
```

```
001270217501560261
002041167544
0031007538
```

```
001X02175Y01560A2B6C1
002X11675A4B4
003Y07538
```

```
001  02175  01560  2  6
002  11675         4  4
003         07538
```

```
001  X02175  Y01560  A2
002  X11675          A4
003          Y07538
```

FIG. 3. UNIVERSAL PROGRAM at right can operate a machine designed to accept any of the first four programs shown and now in use. Blocks below the tapes indicate printed copy produced simultaneously with punching the control tape; the letters and numbers indicate feed, speed, coolant, etc.

Steps in Programming

All numerical-control programming formats are not alike; some operations require more information than others. Starting with the raw data, or an engineering drawing of a part (Fig. 4), the process planner studies the part to determine how it is to be made. He checks the standard tools that are in stock—drills, reamers, cutters, holding fixtures, etc. Any additional tools required are requested from the tool designer, who will design them and have them built. These special tools are usually holding fixtures used to mount the part properly on the work table of the machine tool.

The process planner is the focal point for all information pertaining to numerical control processing. He has to know the kind of tool used to cut or drill the specific material, the rate of metal removal (determined by the speed of the cutter), the feed rate or cutter movement into the workpiece, and the maximum allowable temperature of the workpiece or chips. These and other factors affect the life of the tools, the efficiency of operation, and the quality of parts produced.

To machine the casting shown in Fig. 4, the process planner develops the process planning sheet shown in Fig. 5. He enters a considerable amount of information (such as feed rate, slope, etc.) in addition to the X and Y coordinates; this is usually done manually, then checked.

Copy typing is done on the same sheet (Fig. 6); this typing automatically produces the master control tape. The sheet is positioned in the typewriter so that the tab key positions the carriage to the first possible sig-

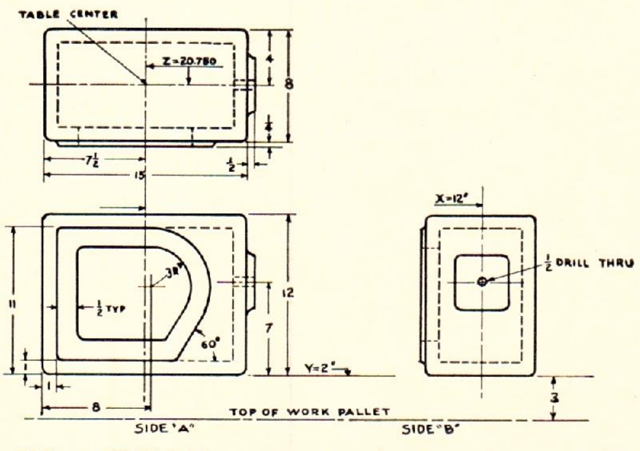

FIG. 4. TYPICAL engineering drawing of a part.

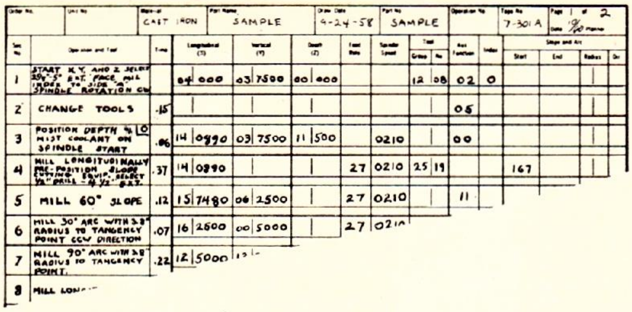
CAST IRON
SAMPLE
4-24-58
SAMPLE
7-301 A
Page 1 of 2

1	START X, Y, AND Z SELECT 29½"-5" EXT. FACE MILL INDEX TO SIDE "A" SPINDLE ROTATION CW		04 000	03 7500	00 000			12	08	02	0				
2	CHANGE TOOLS	.15								05					
3	POSITION DEPTH ½ MIST COOLANT ON SPINDLE START	.06	14 0990	03 7500	11 500		0210			00					
4	MILL LONGITUDINALLY PRE-POSITION SLOPE CUTTING EQUIP. SELECT ½" DRILL - 4½" EXT.	.37	14 0890			27	0210	25	19			167			
5	MILL 60° SLOPE	.12	15 7480	06 2500		27	0210			11					
6	MILL 30° ARC WITH 3.8" RADIUS TO TANGENCY POINT CCW DIRECTION	.07	16 2600	00 5000		27	021								
7	MILL 90° ARC WITH 3.8" RADIUS TO TANGENCY POINT.	.22	12 5000												
8	MILL LON														

FIG. 5. **KEARNEY & TRECKER** Milwaukee-Matic process planning sheet.

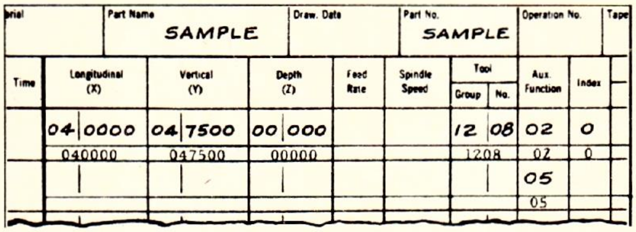
Part Name SAMPLE
Draw. Date
Part No. SAMPLE
Operation No.
Tape

Time	Longitudinal (X)	Vertical (Y)	Depth (Z)	Feed Rate	Spindle Speed	Tool Group	Tool No.	Aux. Function	Index
	04 0000	04 7500	00 000			12	08	02	0
	040000	047500	00000			1208		02	0
								05	
								05	

FIG. 6. **TYPIST** sets tab stops for each data column and types data in space provided directly below manual entry. Machine automatically produces the punch tape.

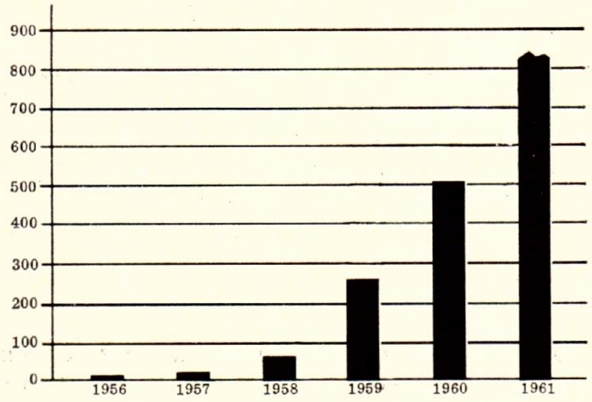

FIG. 7. **NUMBER** of numerical control systems sold for applications to machine tools.

nificant digit. This digit must be typed, even if it is a zero. The tab key also is used to space over to each succeeding column across the page. The operator types in each of those columns where information is to be punched into the control tape.

The typed data is verified by the planner and a duplicate tape punched for use in the shop. The original tape usually is retained as a master control tape by the process planner.

Fig. 8 shows the punched program tape for the first (top) line of information from the process planner's sheet shown in Fig. 5. In this case 33 code holes are punched in the tape for line one, and only 10 for line two. The complete control tape for machining the part shown in Fig. 4 would end up about 38″ or 40″ long, plus a leader.

The machine usually is run from the tape first without any workpiece or fixture on the work table to insure that some step has not been overlooked or programmed incorrectly.

Two readers can expedite processing when the work table is long enough to hold two parts. One is reading Section "A" instructions and the other Section "B." Usually, the same part that has been machined on "A" position is transferred to "B" position, and additional work is done on the part, usually on a surface not attainable from the first or "A" position. This permits the operator to reload one position while machining is being performed at the other position.

Program or numerical control tapes vary in length considerably, depending on the number of operations being performed. Most tapes will range from 18″ to 60″ in length. A few may be as long as 50′; there has been one 1,900′ tape used for a complex aircraft part.

Advantages of Numerical Control

Cost savings per part fabricated using numerical control vary from 20% to 90%, depending on tool cost, lead time, program complexity, etc. Experience with numerical control also indicates possible savings in increased life of cutting tools and reduction in scrap. Fig. 7 shows the growth in numerical control systems for machine tools.

FIG. 8. **CONTROL TAPE** punched from data in first line of planning sheet shown in Fig. 5. and 6.

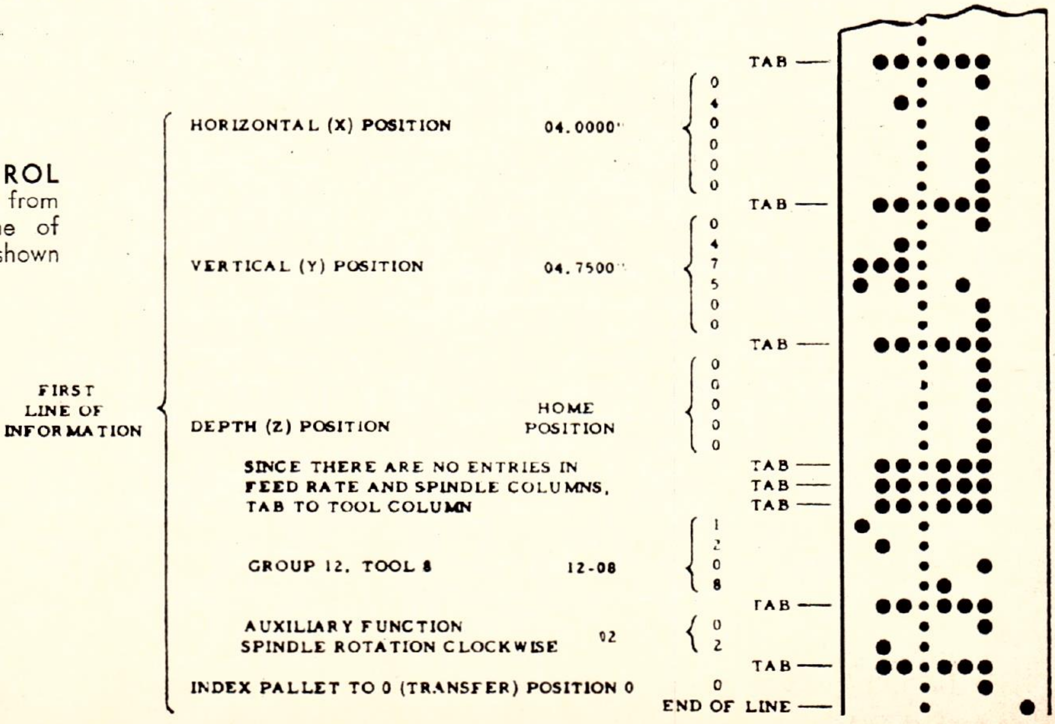

G. W. FLOYD
Rheem Electronics

Modern data applications require error-free high-speed reading of punched tapes. The basic techniques use mechanical or photoelectric sensing. New techniques permit photovoltaic readers to be priced competitively with mechanical types.

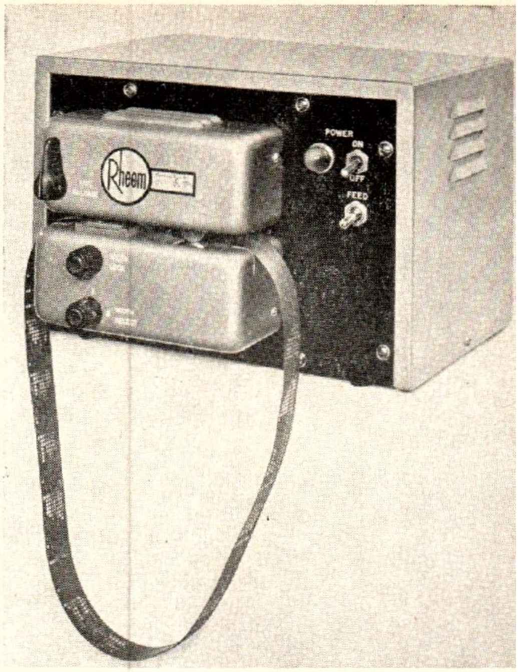

FIG. 1. DESK-TYPE punched-tape reader.

ading Punched Tapes

ONE of the first major uses of punched paper tape was automatic message communication over telephone lines. The tape systems proved more accurate than manual telegraph, but, as message content was normally alphabetic, some error rate could be tolerated.

In recent years more and more numeric-type messages (containing accounting data, sales orders, and other dollar values) are being read and transmitted. Punched tape also provides input to business data processors, scientific digital computers, special-purpose computers, automatic test equipment, process controls, and machine-tool controls. Error rates here must be greatly reduced—and reading speeds must be increased because of the increased volume of data.

Punched-tape-reader requirements for these applications vary, but the basic requirements are reading speed, accuracy, and life. Reading speed usually is stated in lines, or characters, per second. Reading speeds can be divided roughly into two ranges—(1) below 200 lines/second and (2) above 200 lines/second. Message communication, automatic test equipment, process controls, machine-tool controls and some computers use the low-speed readers. Most data processors and other types of digital computers require the high-speed readers.

Error-free reading, desirable in all applications, is essential when punched tape is used to control automatic equipment. Elaborate checking and even correction circuits have been used to provide reliable error-free operation, but this raises the cost of the equipment.

Basic Techniques

Punched-tape readers generally use either *mechanical* or *photoelectric* techniques.

Mechanical readers sense the punched holes by using pins, pronged wheels, brushes, wires, pressurized air, etc., to close an electric contact through the holes in the tape. Mechanical readers have been used in most of the low-speed applications because of cost. However, use of a contact limits the speed to about 20 lines/second (a few models run as high as 100 lines/second). Contacts also reduce reliability and life below that possible with photoelectric technique.

Photoelectric readers sense the punched holes by activating a light-sensitive device, with light passing through the hole or reflecting off non-hole areas of the tape. Speed is limited only by photocell response or the operation of the tape-handling mechanism. Photocell readers are used for high-speed applications (300, 400, 500 and 1000 lines/sec).

Cost

In the past, photocell readers have cost 3 to 10 times more than mechanical readers because (1) the cost of photodiodes, phototransistors, or photomultiplier tubes used for each tape track is high, (2) the optical alignment of individual cells requires accurate, complex holding fixtures, (3) low-drift circuits that compensate the variables and drift of the cells are relatively costly (a-c chopping schemes sometimes are used to reduce drift problems), and (4) tape transport mechanisms often have been copies of costly magnetic-tape transports. However, use of photovoltaic detectors and inexpensive transports now permit photocell readers to be priced competitively with mechanical readers.

Reliability Factors

Reliability and life of mechanical electric contacts are affected by many factors including temperature,

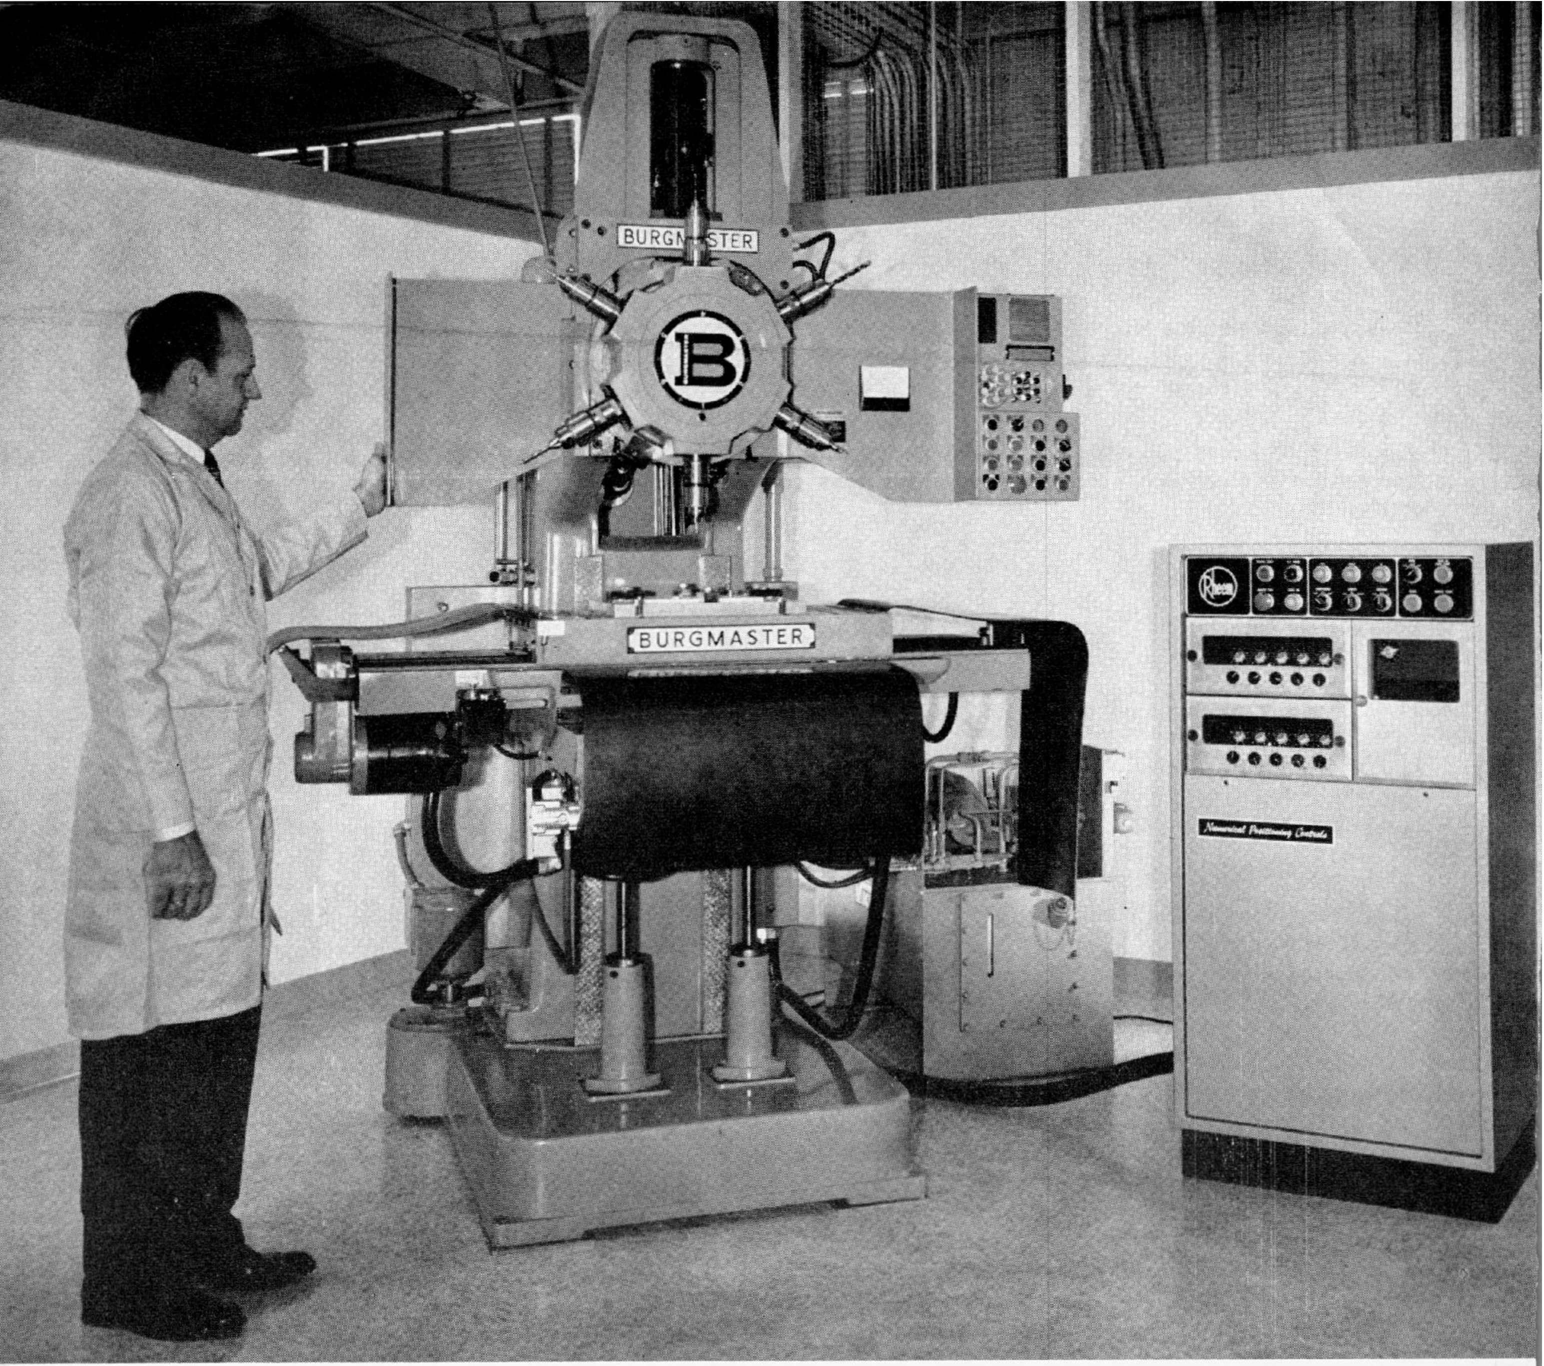

FIG. 2. MACHINE-TOOL numerical positioning control system with punched-tape reader.

humidity, voltage, current, type of load, and type of switching (wet or dry circuit). Contacts must be cleaned and adjusted periodically to operate reliably. The reliability and life of photoelectric systems depend on the short- and long-term drift of the photocells, amplifiers, light source and optical system.

Photodiodes and phototransistors vary considerably from one unit to the next and drift with time and temperature. The more recent photovoltaic cell, which converts light energy directly into electric energy, has low temperature and time drift. A unit for all channels of a reader can be made from the same piece of silicon, making the variation from cell to cell small (typically ±10%).

The drift problem with photovoltaic cells can be solved by designing an amplifier that will operate under minimum cell output. By providing sufficient bias for transistor leakage, it is possible to provide high output currents with a one-stage transistor amplifier.

A reliable low-cost tape transport can be designed for moving tape at speeds up to 150 lines/second. Use of photovoltaic cells and such a transport have reduced the price of photocell tape readers to that of mechanical readers.

Laboratory tests have been made with such a tape transport operated in a start-stop fashion for each line of tape at 60 lines/second for over 1000 hours, providing 3,600,000 operations with no degradation of performance.

Assured error-free operation allows elimination of much error-checking circuitry. Also, tape life is lengthened because tape can be fed with pinch rollers rather than indexed by sprocket holes.

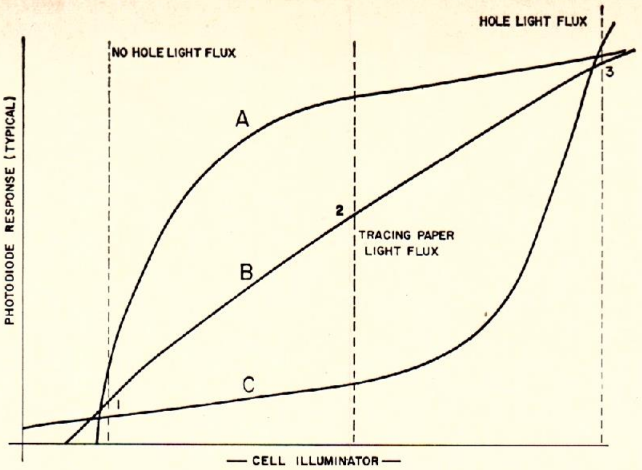

FIG. 1. PHOTODIODE CHARACTERISTIC (electric signal output versus flux input) can be linear or nonlinear. Linear characteristic is most stable and safest.

Photoelectric Paper-Tape Sensing

ANTHONY MAURO
Digitronics Corporation

Photoelectric paper-tape sensing requires consideration of (1) photocell and amplifier characteristics, (2) mechanical considerations and (3) paper-tape handling problems. Over-all system and reliability implications are discussed.

PAPER TAPE is becoming an increasingly important medium for data storage in many applications. Once used only for message communication in 5-channel Teletype systems, paper tape now is being applied to many problems which had been handled previously by punched cards and magnetic tape (both digital and analog).

The advantage of paper-tape systems include:

1. Relatively low cost of paper-tape handling equipment

2. Increase of data rate over corresponding card systems.

3. Ability to physically separate items, together with completely variable item length.

4. Low susceptibility to environmental dirt, humidity and electrical interference.

5. Relative indestructability of master tapes (when proper materials are used).

6. Extremely high reliability (when the proper equipment for the proper speed is chosen).

Paper-tape *preparing equipment* is, at present, completely mechanical. Paper-tape punches now operate at up to 300 characters per second with reasonable reliability, given careful maintenance. (Combined electronic plus mechanical techniques are being developed which will achieve substantial increase in speed and reliability in the near future.)

The present discussion, however, is concerned with the paper-tape *reading* process. Experience has shown good reliability at speeds up to 20 characters per second with mechanical equipment. Although mechanical readers have operated well at speeds up to 100 characters per second, photoelectric sensing units are recommended for speeds between 20 and 2000 characters per second.

Sensing

Photoelectric sensing depends on the difference between the light transmitted through a paper-tape (or in some systems, not considered here, the light reflected from the paper) compared to the direct light through a hole in the paper. A second factor in sensing is the conventional arrangement of a sprocket hole whose diameter is less than, and longitudinally centered in, the data holes. The leading edge of the sprocket hole can be used to allow sampling of the data holes with assurance that all data holes have, by that time, reached a stable "one" or "zero" bit configuration. A third factor in sensing is the transverse alignment of the holes in the paper with respect to the sensing photoelectric elements. (Alignment troubles can result when reading tapes whose sprocket position varies transversely due to punch misadjustment; such tapes read well on mechanical readers which guide on the sprocket but not on photoelectric readers which guide on the edge of the paper.) Finally, the tape must not be allowed to flap excessively over the sensing head because spurious readings can result.

The *transmissivity* of paper tapes varies from complete opacity to 80% transmissivity. There are otherwise relatively opaque tapes which contain defects or, in some cases, oil spots which radically increase the transmissivity. For "digitally safe" paper-tape sensing it is well to assume a maximum paper transmissivity and to design for signal-to-no-signal ratios of no more than 5:1.

Fig. 1 shows photodiode response as a function of illumination. Curves A, B and C show possible responses of a photoelectric transducer. (These curves can be shifted by electrical parameters of the sensing

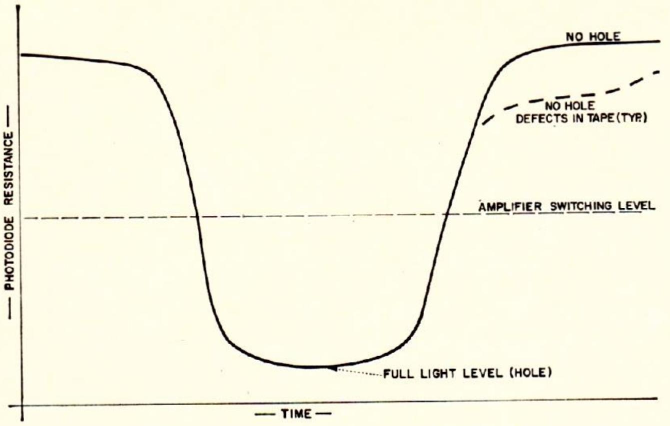

FIG. 2. AMPLIFIER is set to switch on at a diode resistance level well below that which would be due to light transmitted through common defects in paper.

circuit.) Note that although all three curves have similar signal-to-noise-signal ratios at the referenced (no-hole and hole) illumination, curve B affords the most stable (and hence most reliable) characteristic.

Linear characteristics are readily obtainable from reverse-biased silicon photodiodes. Variations between photodiodes can be accommodated by variable circuit parameters so that an inspection procedure can guarantee a properly centered response curve. The practice of inserting a "tracing paper" filter in the inspection jig to establish a third point on the response curve (no hole, hole, tracing paper) is useful, as reference to Fig. 1 shows.

The physical configuration of the photodiode, each with its own rounded lens, has advantages. The diode is mounted with its lens just protruding from the head block so that the paper tape, passing over the lens, continually clears dirt from the optical path. This feature is useful in machine shop or dirty environments.

The amplifiers following the photodiode can be simple. Good digital practice suggests that all stages be completely nonlinear—i.e., completely on or completely off as a function of the hole or no-hole condition of the light to the photodiode. Fig. 2 shows the setting which is used to assure reliable digital operation. Note that the diode resistance level for common defects in the paper is still well above the amplifier switching level. This assures a high degree of immunity to spurious responses. In addition, the full light level affords ample current drive to the amplifier. These conditions can be guaranteed best by a single control for each channel plus rigid stabilization of the lamp voltage. With such a setting, wide variations in environmental conditions can be tolerated safely.

The sprocket channel amplifier for most paper-tape systems should serve as a "strobe" for the data channels. To assure proper waveshape for this strobe it is good practice to include a regenerative circuit as a Schmidt trigger circuit in the sprocket amplifier. The output of the sprocket circuit then can be used optionally to gate the data channel outputs on an "AND" basis. If such logical capabilities are provided it becomes simple to arrange a "Stop character memory." (This means that when the stop command is given the reader, the last character is stored in the amplifiers even though, at high speeds, it is usually desirable to have the tape stop just beyond the "stop" character.)

Chadless Tape

Finally, in sensing, it has been common in Teletype practice to use "chadless" tapes. Such tapes are not punched out completely and are useful where it is desired to print the message on the tape. Mechanical readers have no difficulty in sensing these incompletely punched holes since the sensing pins can readily deflect the "chad" sufficiently to effect a reading.

Photoelectric sensing requires that some other means be used to deflect the chad. Some tape readers have attempted this at high speed by sensing the holes just after the tape has passed over a sharp bend, or as an air blast deflects the chad. Wide experience with such readers is not yet available. However, the desirability of decreasing reliability for the sake of being able to print on the tapes should be questioned. When tape is handling data at up to 1000 characters per second, the need for chadless interpreted tape in the system becomes questionable.

Motion Control

As the reading speed increases, the ability to stop the tape rapidly becomes more important. The command to stop might come from a code on the tape and, if the tape does not stop before the next character, stages of buffer memory become necessary. As the section of tape to be stopped has low inertia, stopping in a few milliseconds does not require much force. However, the associated mass of tape reels can add considerably to the required force. Starting, on the other hand, is not as critical.

Traditional paper-tape transports used a sprocket wheel to drive the tape. As stopping the sprocket wheel in less than one character position at high speeds is difficult, the magnetic-tape transport technique (involving a pinch roller and capstan plus a brake) offers an excellent solution to this problem.

As paper tape is not damaged readily by wear, the braking action need not include mechanical motion of the brake parts. The brake can rest on the paper tape during the motion period. Then, when the stop command is given, it is necessary only to build up the magnetic field. The resulting increase of normal force and corresponding friction can stop the tape quickly. The brake will become operative while the pinch roller and captsan are still disengaging. It is worthwhile to have the brake and pinch roller assemblies on opposite sides of the head so that the resulting tension in the tape during the stopping period helps prevent tape flap.

After the tape has come to a stop the brake pressure should be maintained to avoid accidental tape motion. (This is particularly true of any type of paper-tape-handling equipment associated with the reader.) Consider the possibility that the tape will stop with the light slit edge just beyond a sprocket hole. A small back and forth jitter of the tape as the pinch roller engages will allow a spurious signal to be generated which will throw the next item out of sequence.

TABLE 1—CAPABILITIES OF TAPE-HANDLER TYPES AT MAXIMUM SPEED

Unit	10" Reel *(1000 ft)	8" reel *(500 ft)
Winder	2 IPS** (20 CPS***)	12 IPS (120 CPS)
Spooler	25 IPS (250 CPS)	50 IPS (500 CPS)
Servo	60 IPS (600 CPS)	100 IPS (1000 CPS)

*Standard type NAB reel, using 4.5-mil tape.
**IPS = Inches per second
***CPS = Characters per second.

Fig. 3 shows a bidirectional paper-tape reader. The pinch roller and capstan assemblies are most remote from the head. This assembly could be closer to the head in order to minimize required paper-tape leader. However, this would mean pushing the tape through the brake associated with the opposite direction of motion, and a corresponding possibility of jamming. The design shown allows for a minimum 6″ leader together with bidirectional operation.

Paper-Tape Handling

The relatively large mass of reels and paper tape requires that a mechanical buffer be inserted between the reader and the reel when paper-tape speeds increase over 2 inches (20 characters) per second. Depending on the size of reel and the speed, a number of tape-handling configurations can be used. Listed in order of increasing cost these configurations are:

1. Winders—Unwinders
2. Spoolers
3. Servos

Winder units involve simple tension-arm-controlled drives. So long as the tension arm is away from a home position, the winder motor is operative and pulls tape. The tension arm of the unwinder simply acts to brake the reel motion when the slack loop gets too big.

Spooler units (Fig. 4) use reversible torque motors on both the take-up and feed reels. The tension arms each operate in three zones—*central zone* (de-energizes the motor), *maximum-loop zone* (causes full take-up torque), and *minimum-loop zone* (causes full feed-out torque). Various interacting controls are designed to allow for tape breakage, rewind, etc. As the reader accelerations and decelerations are virtually instantaneous compared to reel motions, a set of straightforward relationships can be derived between tape speed, reel size, tension arm travel and the required motor torques.

Servo units (Fig. 5) have greater control abilities than spoolers because the control function is continuous (or at least multi-zone) and with variable torque.

Table 1 shows comparable capabilities of these various tape handlers.

If the reader and tape handler have been chosen properly, reliability rates in the order of 1 error per 10^{10} bits are attainable. As the popularity of punch-tape equipment increases, lower costs are being announced so that the cost-reliability criteria of paper-tape operation becomes evermore attractive.

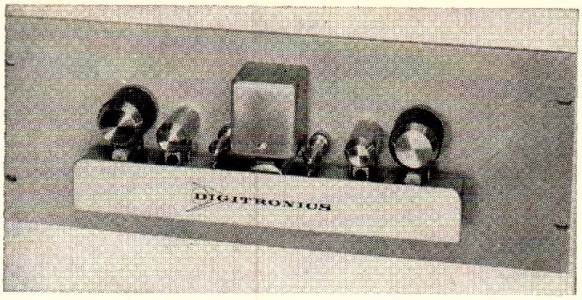

FIG. 3. BIDIRECTIONAL Tape Reader Model B3500.

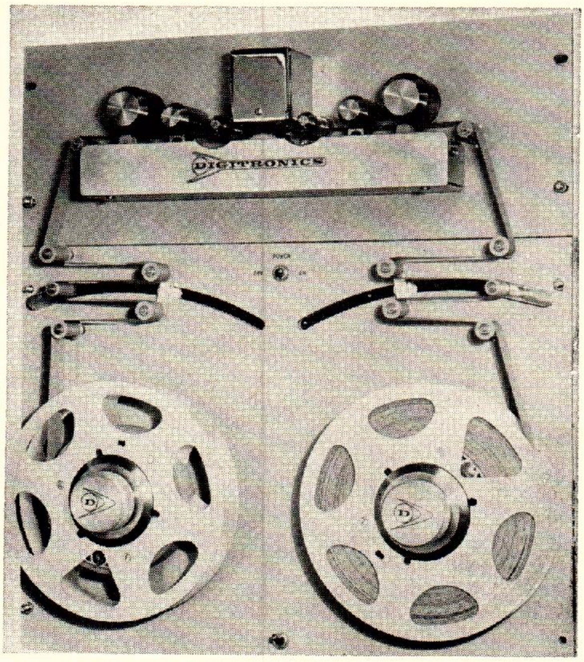

FIG. 4. SPOOLER Model 4588 will handle 500 feet of paper tape or 1000 feet of Mylar tape bidirectionally at 500 characters per second.

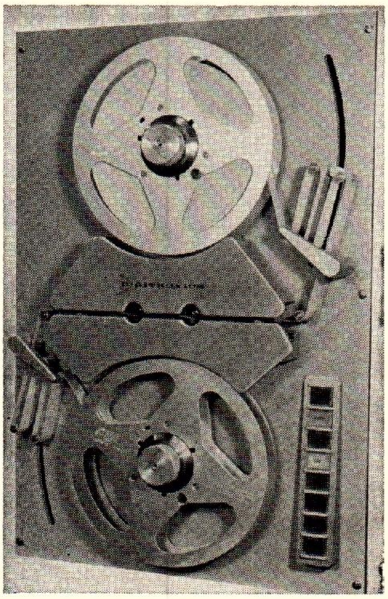

FIG. 5. SERVO-TYPE tape handler Model 4500. This reader plus servo-controlled reels for tape handling will operate at 600 characters per second.

Magnetic Recording Techniques

The most popular recording techniques are (1) direct recording, (2) frequency-modulation recording, (3) pulse-duration-modulation recording, and (4) digital recording. These techniques are described.

D. W. HALFHILL
Ampex Corporation

DIRECT ANALOG RECORDING is the familiar process used in the recording of speech and music. It also can be used for recording instrument signals or data. This is the simplest of all recording processes and usually requires one tape track for each signal channel.

Fig. 1 shows the basic tape recording characteristic—that is, remanent magnetism versus applied magnetic field. If the signal to be recorded is mixed with a high-frequency bias it is possible to record the signal completely within the straight-line portions of the recording characteristic, as shown in Fig. 1.

The bias and recorded signals are combined by a linear mixing process—that is, the signals are mixed in a linear mixer, but not in a nonlinear modulator. This type of mixing is not amplitude-modulation (in which new sum and difference frequencies would be generated) but is a simple mixing. The bias frequency is usually at least 3.5 times the highest signal frequency to be recorded. As the recorded wavelengths of the high-frequency bias are too small to be resolved by the playback head (Fig. 2), the bias frequency is not picked up from the tape—its function is only to insure that the recorded signal remains on the linear portion of the characteristic.

The output voltage of the reproduce head is proportional to the frequency of the signal because the head output is a function of rate-of-change of magnetization. Therefore, the reproduce (decoding) amplifier must have a frequency-response characteristic which is the inverse of the reproduce head characteristic to obtain an over-all flat frequency response. This is referred to as *playback equalization* (Fig. 3).

When signal frequencies are so low that the resulting reproduce head signal output approaches the system noise it is impossible to recover the signal by equalization. This accounts for the principal limitation of direct recording, the lower frequency limit, below which it is impossible to play back successfully. The upper frequency limit is reached where the wavelength on the tape approaches the gap width. This is a function of recording tape speed, as shown in Fig. 4.

The high-frequency end (100 kc at 60″/sec) has

TABLE 1—COMPARISON OF RECORDING PROCESSES

Process	Low-Frequency Limit	High-Frequency Limit	Effect of 1% change in tape speed	Figure of Merit (cycles/inch)	Channels per track	Advantages
Direct 60″/sec 30″/sec 15″/sec	50 cps 50 cps 50 cps	100 kc 50 kc 25 kc	1%	1600	Usually one; several with tone multiplexing	35-db dynamic range; overloads gracefully with gradual distortion; simplest; widest frequency response
FM wide deviation	0 cps	10 kc at 60″/sec	2.5% error	167	One on wide band; up to 18 on one track with narrow-deviation (7.5%) FM	Good to zero frequency; relatively insensitive to tape dropout; good resolution of transients; permits changes in time base; good linearity.
FM narrow deviation	0 cps	Depends on deviation	13% error at 7.5% deviation	167		
PDM	0 cps	Limited by sampling rate up to 5 cps nominal.	1%	40	UP to 86 channels, time multiplexed	Good to zero frequency; less sensitive to flutter than FM; up to 86 channels per track with accuracy better than 1%.
Digital	0 cps	Limited by sampling rate and code used.	None	5—NRZ 2½—RZ	Usually one; can be time multiplexed	Highest accuracy recording in computer language; insensitive to flutter and wow. Simple circuitry; output can go directly to computers.

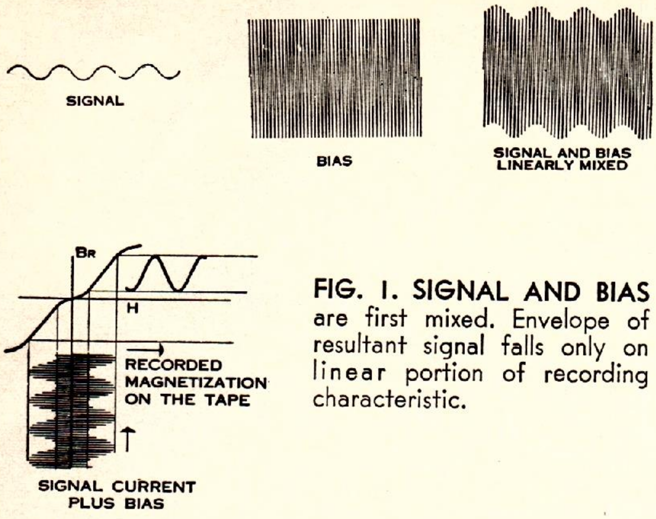

FIG. 1. SIGNAL AND BIAS are first mixed. Envelope of resultant signal falls only on linear portion of recording characteristic.

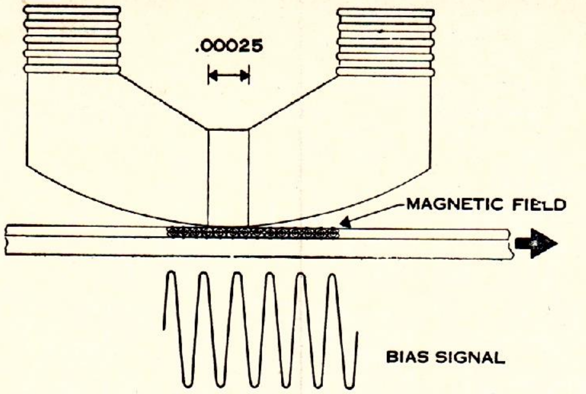

FIG. 2. WAVELENGTH of bias frequency on tape is too short to be resolved by reproduce head.

FIG. 3. REPRODUCE amplifier delivers flat output because of "playback equalization" of reproduce-head output.

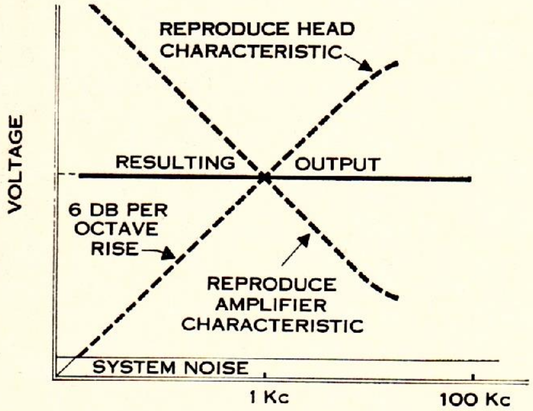

FIG. 4. OVER-ALL frequency response, direct recording, at three tape speeds.

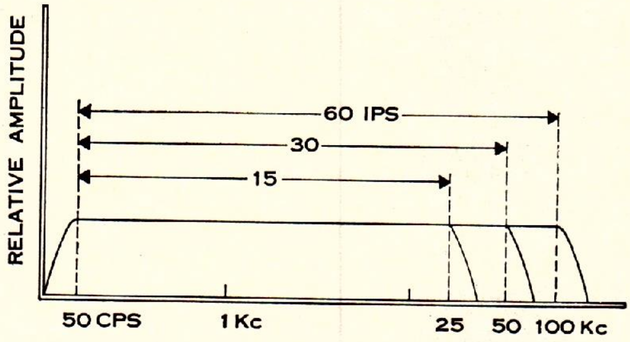

Disadvantages	Applications	Tape Requirements
Audio recorder not suitable for instrumentation; dropout effect serious for instrumentation (not for audio); won't record to 0 cps due to noise.	Audio (with equalization); spectrum analysis of noise and underwater sound; recording any a-m signal with frequency components between 50 cps and 100 kc.	Dropout affects higher frequency signals.
Complex circuitry; severe transport requirements because flutter causes serious error. Tape is used less efficiently than in direct recording.	Shock blasts; explosion; any signals with components from 0 to 10 kc; radio telemetry.	High squareness tape characteristic; low d-c modulation noise; smooth surface.
Limited freq. resp.; complex equipment and circuitry; low tape utilization.	Many channels with low-frequency signals, such as flight and engine testing.	Least demanding of all tape methods.
Poor tape economy; data must be digitized; very sensitive to dropout—requires computer-quality tape and checks for lost or spurious data.	Computers; accounting; numerical applications; processing of flight test and wind tunnel data.	Most demanding because drop-out means an error. NRZ requires tape squareness characteristic; low d-c modulation noise.

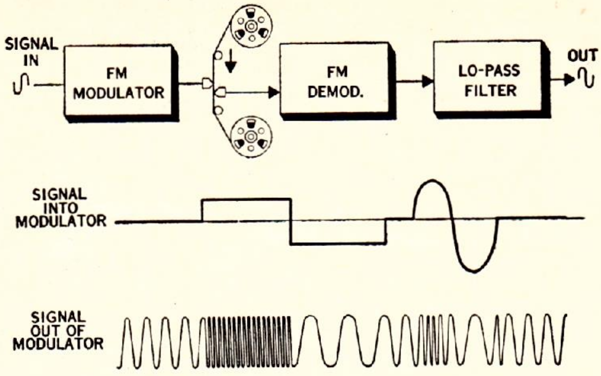

FIG. 5. FM OPERATION. Low-pass filter at output rejects carrier frequency and passes only desired signal.

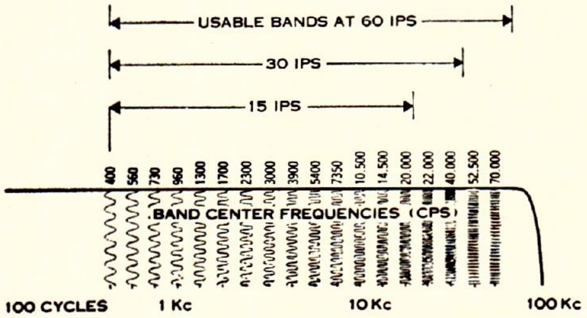

FIG. 6. NARROW-DEVIATION FM is used in frequency multiplexing.

1600 cycles per inch of tape which can be taken as a figure of merit for direct recording and used to compare with other tape recording systems (Table 1).

Audio Recorders

An ordinary audio recorder should not be used for instrumentation recording because the energy content in speech and music signals is not uniformly distributed over the range of signal frequencies and the recording amplifier pre-emphasizes the extreme low and high ends of the frequency spectrum. Thus, when an audio recorder is used for instrumentation, the built-in pre-emphasis distorts the high and low frequencies. This could be overcome by reducing the recording level by a considerable amount, with resulting deterioration of the signal-to-noise ratio.

Frequency-Modulation Recording

Frequency modulation overcomes the two basic limitations or direct recording—inability to record low frequencies and amplitude instability caused by tape dropouts. A particular frequency, selected as the center frequency, corresponds to zero input signal. A d-c signal causes the carrier frequency to vary in one direction; an a-c input signal causes the carrier frequency to vary on both sides of center frequency (Fig. 5). Thus, all information presented to the tape is preserved in the frequency domain and normal amplitude instabilities have little or no effect on the recording.

On playback, the signal is demodulated and fed through a low-pass filter which removes the carrier and other unwanted frequencies generated in the modulation process.

Two FM-carrier systems are in use—wide deviation and narrow deviation.

A widespread application of the latter is in frequency-division multiplexing (Fig. 6). Here a number of individual carrier frequencies are each modulated by a separate input signal. The result is a multiude of signals which are mixed linearly. The composite signal can be reduced on one track using the direct reading process. Thus, the wide bandwidth and linearity of direct recording permits simultaneous recording of many channels of signal information on one track of tape.

Flutter and Wow

FM recording makes stringent demands on the ability of the tape transport to move tape at uniform speed. Any speed variations cause an unwanted modulation of the carrier frequency, causing system noise. This is the limiting factor in the dynamic range and accuracy of FM.

This limitation is particularly acute in the frequency-multiplex scheme (Fig. 6) because a 7.5% frequency deviation usually corresponds to 100% input signal. Hence a 1% deviation in frequency resulting from flutter appears as a noise signal of $100/7.5 = 13.3\%$. Similarly, any speed variation in the transport is multiplied by a factor of 13.3.

Wide-deviation FM has a wider dynamic range and greater over-all accuracy. Here only one channel of signal is recorded on a track of tape and the entire bandwidth is used for this signal (Fig. 7). Hence a deviation of $\pm 40\%$ of carrier frequency is possible. In this case a 1% tape-speed error causes only a $100/40 = 2.5\%$ noise signal, an improvement of better than five times that of the narrow-deviation system.

While it is possible to record data down to dc at any tape speed, the upper-frequency recording limit varies directly with tape travel, going to 10 kc at 60″/sec with standard output filters. This gives a figure of merit of 167 sine-wave cycles per inch of tape, but the upper frequency limit is much higher than those offered by other recording processes which will record dc. The center-carrier frequency varies with tape speed and is usually chosen about midway in the frequency band dictated by this speed. Thus, with wide-deviation FM recording, where a 60″/sec tape movement offers a carrier frequency range from 100 cps through 100 kc, the center-carrier frequency is usually around 54 kc. At 1⅞″/sec the center carrier frequency is 1.7 kc, midway in the range between 100 cps and 3125 cps. At any tape speed, the highest data signal frequency to be recorded is approximately 1/5 the center-carrier frequency.

Time Base Change

If it is not necessary to record frequencies as high as 10 kc, the tape speed can be reduced in direct proportion to the desired upper-frequency limit. When this is done, the center frequency is scaled down in direct proportion to the tape speed, using the same ±40% deviation used in wide-band FM-carrier technique. Thus, the wavelength recorded on the tape corresponding to a given d-c signal voltage is the same, independent of tape speed. This makes it possible to record at one tape speed and reproduce at an entirely different tape speed, permitting time base changes, up or down, one of the important applications of FM recording. Note that the maximum ratio between standard record and playback speeds is 1:32 (ratio of 1⅞ to 60 inches per second).

Pulse Duration Modulation

The basic techniques of direct and FM recording have been described.

The third basic technique is PDM, or pulse-duration modulation, wherein the duration of a pulse is made proportional to the amplitude of the signal to be recorded. This technique is used commonly with commutation to multiplex a large number of channels. This is called *time-division multiplexing* and requires sampling a number of signal channels sequentially.

For example, the signals shown in Fig. 8 can be sampled at uniformly spaced discrete intervals, recording only the instantaneous values at the time of sampling. The original waveform can be reconstructed on playback by passing the discontinuous readings through an appropriate filter. An accurate reproduction of a sine wave can be made using as few as six samples per sinewave cycle. The sampling rate should be at least six times the highest significant frequency-component.

It is possible to space the intervals and use the time between for sampling other data channels by a commutator which samples the outputs of a number of transducers in sequence (Fig. 9).

Two contacts on the commutator are reserved usually for frame reference to allow synchronization with the commutator used on playback. Two other contacts can be used for calibration signals, one for full-scale and the other for zero signal. Thus, the system permits calibration once each revolution. The remaining contacts are connected to the outputs of various transducers.

The signal from the commutator (point A in Fig. 9) is a sequence of short-duration pulses of varying amplitude (Fig. 10, A).

As magnetic tape cannot record these varying-amplitude signals accurately and dependably, the signals are passed through a keyer which converts them from varying-amplitude signals into constant-amplitude signals of varying pulse width, or pulse duration (Fig. 10, B). Pulse durations normally run from 90 microseconds (for zero signal) to 660 microseconds (for 5-volt full-scale signal).

The only significant information in the signals from the keyer is the time each pulse begins and ends. The record amplifier (encoder) differentiates these pulses and presents the record head a positive spike for the begininning of a pulse and a negative spike at the end of a pulse (Fig. 10, C). Constant tape speed is important but flutter and wow is less critical because only the speed error between the instant of pulse start and pulse stop causes error in recorded

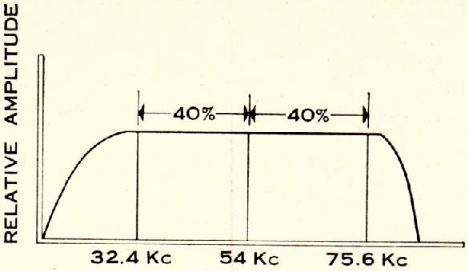

FIG. 7. WIDE-DEVIATION FM carrier bandwidth at 60"/sec tape speed.

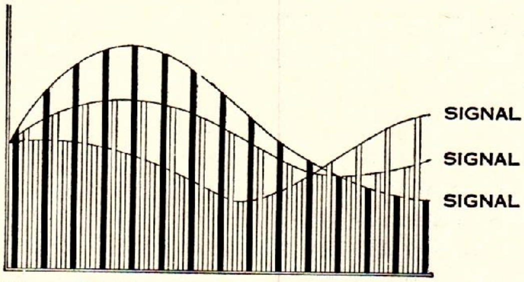

FIG. 8. TIME-DIVISION multiplexing by signal sampling. Time between samples of a given signal is used to sample other signals.

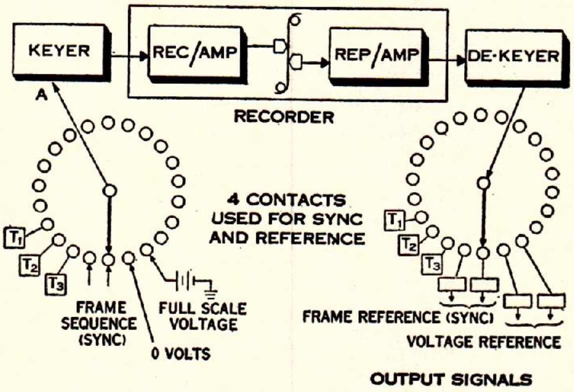

FIG. 9. MULTIPLEX PDM recording technique.

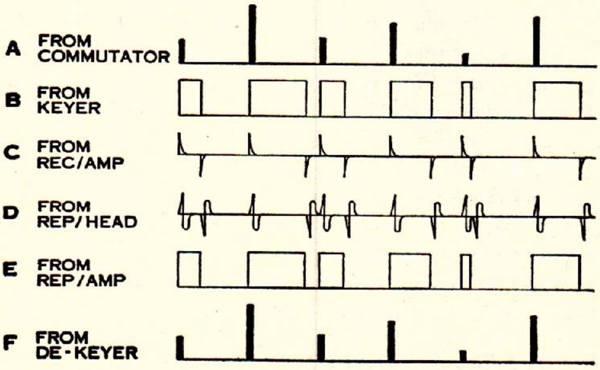

FIG. 10. SIGNAL shapes at various points in PDM.

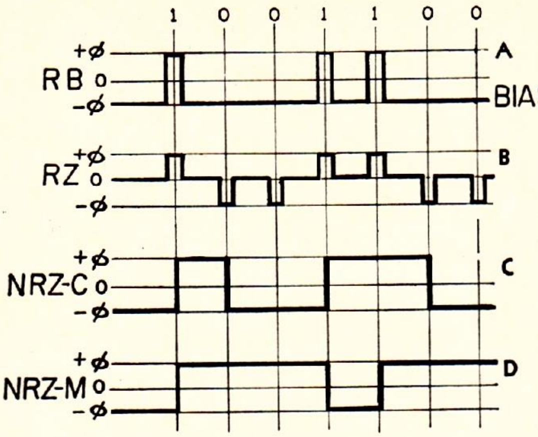

FIG. 11. TAPE magnetization in RB recording, RZ recording, NRZ-C recording, and NRZ-M recording for number 1001100.

data.

Output from the reproduce head is, as in all magnetic-tape systems, a differentiation of the recorded signal (Fig. 10, D). The instant the reproduced wave crosses the axis can represent pulse start and stop. The reproduce amplifier (decoder) contains a multivibrator which recreates the varying-duration pulses (Fig. 10, E). These pulses are then fed through a dekeyer where they are converted back into varying-amplitude pulses of short duration (Fig. 10, F). This is decommutated and the separate data channels sent to individual filters that reconstruct the original data signals.

The chief advantage of PDM is its ability to record many signal channels at better than 1% accuracy overall. One track can carry 86 channels using a 10-rps 90-contact commutator.

Digital Recording

A spot on the surface of the drum or disc is magnetized for each bit of information to be stored.

Any of several techniques can be used for recording binary digits. Four of the most representative methods are shown in Fig. 11, which shows tape magnetization for recording the binary number 1001100 by each system.

A. Return-to-Bias (RB). The tape is magnetized (biased) to saturation in a direction called minus. Zero binary signals have no effect on the tape, which is left magnetized in its biased (minus) direction. Binary 1 signals are recorded by magnetizing the tape in the opposite (positive) direction. After each binary 1 pulse the tape returns to the bias condition. The number 1001100 would be recorded as shown in Fig. 11A. Note that this method requires a clock to read zeros.

B. Return-to-Zero (RZ). The tape is initially in its unmagnetized condition. Binary 1 and 0 are assigned opposite magnetization directions and a pulse is recorded for each bit (Fig. 11B). Return to unmagnetized state occurs after each pulse, and the system is self-clocking.

C. Non-Return-to-Zero—Change (NRZ—C. The tape is always saturated in one direction or the other. The direction of saturation reveals the bit recorded. Reversals in direction occur each time the bit *changes* (Fig. 11C).

D. Non-Return-to-Zero—Mark (NRZ—M). This differs from NBZ-C recording in that magnetization direction is changed each time a binary 1 is recorded (Fig. 11D). Both NRZ systems require an external clock. NRZ-M is the most widely used in digital recording systems, and appears to offer a good compromise between accuracy, simplicity, reliability, and compatibility. Note that this system involves the least number of flux reversals.

Relative timing accuracy is a problem of greatest importance in digital recording because the digits (or bits) making up a given number or character are recorded simultaneously, in parallel across the tape, with each bit on a different track. Obviously the pulse-packing density must not be too great because the closer the pulses, the more the possibility of errors due to tape skew, dropouts, head scatter, etc.

Dropout, Redundancy and Parity Check

Digital recording is sensitive to tape dropout errors; lost or spurious pulses cannot be tolerated. For this reason special precautions are taken in the manufacture, inspection and selection of tape intended for digital recording. Even this does not solve the problem completely.

Safeguards can be built into the digital recording process to provide greater reliability against tape dropouts or other errors. One of these is the use of redundancy technique, in which the same information is recorded twice.

A second scheme is the use of a parity check, in which one track on the tape is reserved for a pulse which is derived from the pulse being recorded simultaneously on the other tracks. A parity pulse is of such polarity that the sum of all bits on playback (including the parity bit) will be an odd number. Thus an error will be indicated if one pulse (or any odd number of pulses) is lost. Parity checking will not detect two simultaneous errors or any even number of errors.

Tape dropouts become most critical at short wavelengths—where the pulse duration approaches the size of the gap in the reproduce head. This demands the use of moderately long pulses on the tape, reducing the pulse packing density. The limitations imposed by dropout and tape skew have resulted in an accepted conservative pulse-packing density of 200 pulses per inch for reliable digital recording. As 7 bits equal one word or sample, and as six samples are required per cycle, 200 pulses per inch are the equivalent of about 5 cycles per inch of tape:

$$\frac{200\ \text{bits/inch}}{(6\ \text{simples/cycle})\ (7\ \text{bits/sample})} = 5\ \text{cycles/inch}$$

Thus the figure of merit is low (Table 1) and it is customary to edit the data while still in analog form to determine those portions of the data which have significance. This is sometimes done using quick-look graphic techniques. The selected data are converted to digital form and recorded for further refinement and computation.

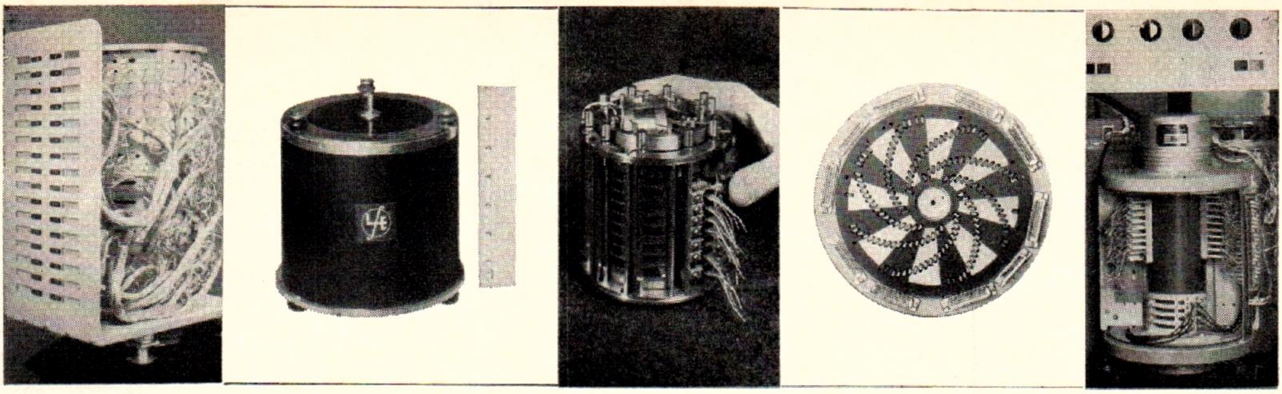

TYPICAL drum and disc storage elements. LFE disc memory (second from left) uses Bernoulli separation principle.

Drum and Disc Storage

J. M. GOTTSCHALK*
W. J. WASYLENKO**
Burroughs Corp. Laboratories
Research Center

This survey of drum and disc recording techniques covers history, techniques (RB, RZ, NRZ-C, NRZ-M) and types.

STORAGE (memory) requirements for most digital computers are:

1. **Capability of storing information indefinitely (in a non-volatile manner).**
2. **Accessibility of information, when called for, with minimum delay (low access-time), and preferably with nondestructive readout.**
3. **Capability of altering or erasing stored data to permit storage of new data.**
4. **Capacity to store a vast amount of information—perhaps 10^8 bits.**
5. **Minimum cost per bit stored.**

Magnetic drums and discs meet most of these requisites. Although their access time is greater than that of cores, these devices are used commonly for the main storage mechanism of most general-purpose digital components. Where information is to be used in a sequential fashion, magnetic tape is used commonly, as in input devices; where smaller access time is essential, the more costly core can be used.

Magnetic Recording

The three essential components of any dynamic magnetic storage device are:

1. **A magnetizable medium with suitable retentivity.**
2. **A write head and a read head (may be the same).**
3. **A means for energizing the write head, and means for converting the read-back signals to a usable form.**

The magnetic *media* most commonly used are iron oxides suspended in a resin binder, or an electroplated cobalt-nickel alloy. The read-write *head* consists of a ferromagnetic core (basically a toroid) with a small gap containing a non-magnetic shim. One or more coils of copper magnet wire are wound about the magnetic core. The core element is mounted in a shell to position the gap close to the magnetic surface. Electronic amplifiers provide the current required to drive the write head and to amplify the signals from the read head to usable levels.

Storing Digital Information

A spot on the surface of the drum or disc is magnetized for each bit of information to be stored.

Any of several techniques can be used for recording binary digits. Four of the most representative methods are shown in Fig. 2, which shows tape magnetization for recording the binary number 1001100 by each system.

---

*Mr. Gottschalk is now with RCA, Advanced Development EDP

**Mr. Wasylenko is now with Remington-Rand Univac.

A. **Return-to-Bias (RB).** The tape is magnetized (biased) to saturation in a direction called minus. Zero binary signals have no effect on the tape, which is left magnetized in its biased (minus) direction. Binary 1 signals are recorded by magnetizing the tape in the opposite (positive) direction. After each binary 1 pulse the tape returns to the bias condition. The number 1001100 would be recorded as shown in Fig. 2A. Note that this method requires a clock to read zeros.

B. **Return-to-Zero (RZ).** The tape is initially in its unmagnetized condition. Binary 1 and 0 are assigned opposite magnetization directions and a pulse is recorded for each bit (Fig. 2B). Return to unmagnetized state occurs after each pulse, and the system is self-clocking.

C. **Non-Return-to-Zero—Change (NRZ—C).** The tape is always saturated in one direction or the other. The direction of saturation reveals the bit recorded. Reversals in direction occur each time the bit *changes* (Fig. 2C).

D. **Non-Return-to-Zero—Mark (NRZ—M).** This differs from NBZ-C recording in that magnetization direction is changed each time a binary 1 is recorded (Fig. 2D). Both NRZ systems require an external clock. NRZ-M is the most widely used in digital recording systems, and appears to offer a good compromise between accuracy, simplicity, reliability, and compatibility. Note that this system involves the least number of flux reversals.

Each *change* in magnetization level on the tape produces an output pulse on readback. The two styles of readback signals are shown in Fig. 3 and 4. Fig. 3 shows the doublet pulse signal developed when reading a binary 1, then a binary 0, which were recorded by the RZ technique (Fig. 2B). The binary 1 signal swings first above, then below, the background level. The binary 0 is reversed; it swings below, then above, the background level. This system is self clocking; each binary bit produces a readback signal. The disadvantage is that the readback signal has a higher frequency content and a smaller readback amplitude than from other recording methods.

The readback signal for data recorded by the RB technique (Fig. 2A) has no 0-bit signal. Only the 1-signal doublet pulse (Fig. 3) will appear—each time a 1 bit was recorded. A 0 bit is presumed when the clock pulse is not accompanied by a readback signal from the tape data track.

The readback signal shown in Fig. 4 is from a binary 1 and 0 recorded by the NRZ-C technique. A readback pulse is present only when the binary bit sequence changes. When a clock pulse is not accompanied by a readback signal pulse, the bit remains unchanged; that is, it maintains its previous value.

The signal shown in Fig. 4 is produced when reading data recorded by the NRZ-M technique. Now each readback pulse means a 1 bit—regardless of the polarity of the readback pulse. When the clock pulse is not accompanied by a readback signal pulse, a bit 0 is presumed.

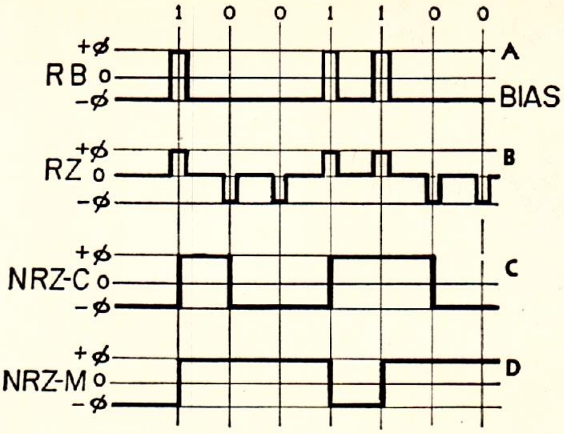

FIG. 2. **TAPE** magnetization in RB recording, RZ recording, NRZ-C recording, and NRZ-M recording for number 1001100.

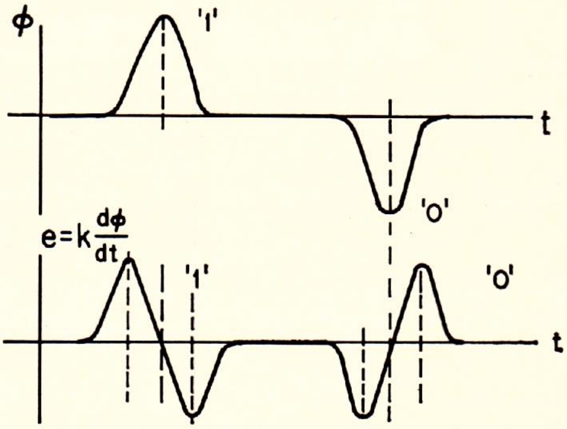

FIG. 3. **RECORDED** bits and readback pulses using RZ technique.

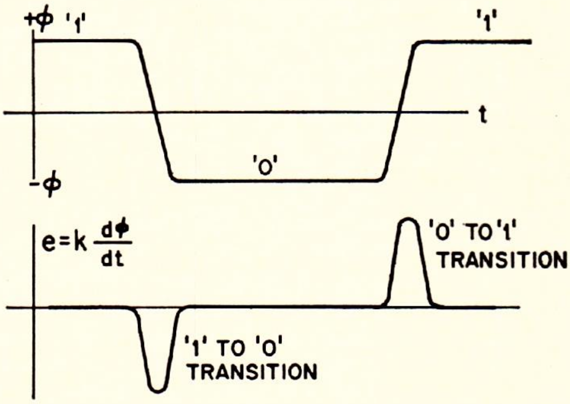

FIG. 4. **RECORDED** bits and readback pulses using NRZ-C technique. Pulses look the same when reading back data recorded by NRZ-M technique, but here any pulse, regardless of direction, always means a 1 bit.

Magnetic Drums

The magnetic-drum storage device is not a recent invention. In 1900, Valdemar Poulsen, a Danish engineer, was issued U. S. Patent No. 661619 on a device consisting of a rotatable brass cylinder, on which was wound a steel wire in a uniform helix, and an associated electromagnet or a recording head. The device was used to record and reproduce audio signals. However, as ear phones were required to hear the reproduced recorded "sound," the invention was not a popular recording device until electronic amplifiers appeared.

When amplifiers became available, tape and wire recorders, rather than drums, were successfully exploited for audio-signal recording and reproducing.

Drums and discs came back into the picture during and shortly after World War II as a medium for computer storage. In the early days of this development, magnetic tape was attached to the cylindrical surfaces of wheels. Today, the device is a true-running cylinder coated with a magnetizable medium, such as iron oxide. Electroplated surfaces also are available.

Types of Drums

Although drums vary in size and speed, they can be classified into one of the following three basic structures:

1. **Direct drive, rotating shaft.**
2. **Belt driven, "dead" (non-rotating) shaft.**
3. **Built-in motor type, using shaftless motors or "inside out" motors.**

The earliest drums were direct driven, rotating shaft. This type is still used in noncritical work. The cylinder is fixed securely to the shaft and mounted between two pairs of preloaded bearings. The cylindrical surface is made true by machining off the runout after assembly. The rotor then is dynamically balanced and the cylindrical surface sprayed with the iron-oxide coating. The head-mounting stator then is placed around the drum and the recording heads positioned properly. A drive motor is either directly coupled (or belt coupled) to the drum shaft. With this structure it is difficult to obtain a dynamic runout of less than 0.0002" total indicator reading.

In the second structure (class 2), the bearings are mounted directly in the cylinder or rotor. The rotor assembly then is axially fixed on a stationary shaft and an axial preload is applied to the bearings. The runout is machined off after assembly and the rotor dynamically balanced. The magnetic surface (iron oxide or electroplate) is then put on and the head-mounting stator (with heads) is placed around the rotor. A belt couples the drive motor to the rotor. Dynamic runout of less than 0.0001" is readily obtainable because the shaft has a high stiffness with both ends fixed. As there is fear of belt failure, the belt life is constructed to have a life in excess of the bearing life.

In the third structure (class 3), the motor is mounted directly on the drum rotor shaft, usually between the supporting bearings. The structure requires an a-c motor of the squirrel-cage type. With a rotating shaft the armature or squirrel cage is fastened to the drum shaft and the motor stator is built into the head-supporting structure. In dead-shaft construction the squirrel cage of the inside-out motor is pressed into the drum cylinder, and the stator is fixed on the stationary shaft. Both structures are fixed-speed devices, and they usually require a warm-up period to attain thermal equilibrium. This compensates for the rotor heating caused by the direct coupling to the drive motor. This structure is compact and is favored where volume is a limiting factor.

Thermal Stability

The thermal stability of drum storage devices depends on two factors—(1) the physical properties of the structural materials and (2) the maximum temperature differential between the drum rotor and the head mounting. Materials can be selected that have similar coefficients of linear expansion or which can compensate for any non-linear system expansion. Also, the materials can be properly heat-treated to minimize the distortion at the maximum operating temperature-differential.

The temperature difference between the rotor and the head mount produces a variation in the head-to-surface separation. Increased separation reduces the writing ability of the recording head.

Magnetic Disc Storage

A disc is defined here as a device with a slenderness ratio (outer radius/thickness) of 20:1 or higher. Thicker discs are called plates or turntables. Thin discs are preferred if a large number of discs are to be stacked on a common shaft.

The use of a disc plated with a magnetic material is suggested as a storage device in U. S. Patent No. 836,339 issued to P. O. Pederson on November 20, 1906. However, only since the mid-forties have discs received serious consideration as general-purpose memories for digital computers. The difficulty in obtaining discs with sufficient flatness held back the development.

Compressed air escaping from a shroud around the head is effective in maintaining a uniform gap or separation between the head and the disc surface. This form of head-guiding is termed hydrostatic because an external source of pressurized air is necessary. This "hydrostatic bearing" permits discs with large surface runout (up to 0.010") to be successfully used for recording.

Discs did not gain approval as storage elements until adequate head-positioning mechanisms were developed.

Kingsbury, in 1897, demonstrated an air-lubricating hydrodynamic slider bearing. The load-carrying capacity of this bearing is low; hence the Kingsbury air-lubricated bearing remained a laboratory curiosity for the student of mechanical engineering. However, the Kingsbury air bearing has proved to be ideally suited as a head-flying device.

To maintain a fixed head-to-surface separation, the low load-carrying capacity of this air bearing favorably limits the head loading force to a small value. This separation is held by a dynamic pressure resulting from a shearing deformation of the air (fluid) molecules wedged between the stationary bearing pad and the moving disc surface. The success of the air-guided head has re-stimulated interest in magnetic disc storage devices, in particular for large-capacity systems.

Air-guided heads are not limited to flat surfaces. Although heads mounted in hydrodynamic bearings were developed to improve the usefulness of discs as storage elements, heads can be air-guided on cylindrical surfaces.

Features which make discs attractive for memory use are:

1. **Expandability by adding discs on a common shaft.**
2. **Higher volumetric efficiency than drums, conservatively estimated at a ratio of 6:1 for medium storage-capacity, higher for large storage-capacity.**
3. **No need for high-precision bearings, nor is a bearing preload required.**
4. **A damaged disc can be replaced; a scored drum requires expensive reworking.**
5. **Simplicity of construction.**
6. **No variable temperature problem (not true of drums having air-guided heads).**

Disc flatness is important in order to increase the pulse-packing density. High-density recording requires head-to-surface separations of the order of 250 microinches or less. To obtain such flying gaps requires discs having a maximum flatness excursion rate of 0.0004" per inch. Other specifications for the disc *substrate* includes: Surface finish of 15 microinches or better, ease of plating, nonmagnetic, dimensional stability, and low cost. Disc materials now meet these specifications; however, cost must be reduced substantially to attain commercial feasibility.

Fig. 5 shows a typical head-disc memory constructed at the Burroughs Research Center. It has its head mounted in a gimbal, which in turn is fastened to a pivoted arm, or head-loading bar. The head is loaded by placing weights on a rod attached to the loading bar. The ratio of the weight to the head force is 6:1. With a head-loading force of 0.7 pound and a surface speed of 950 inches per second, a 500 millivolt peak-to-peak readback signal is obtained, at 500 ppi NRZ (on a cobalt-nickel plating 0.00022" thick having a coercive force of 650 Oersteds). A signal-to-noise ratio of at least 70:1 was measured. The noise level is sufficiently low to be ignored.

To improve on this performance, the head was redesigned using a better ferrite for core material and a smaller gap. This, together with increased head loading for closer spacing, and thinner platings for better high-frequency operation, has made it possible to obtain pulse-packing densities in excess of 1500 ppi with readback signals of 30 millivolts peak-to-peak or better.

Heads can be shifted to serve a number of tracks; however, the access time increases by the time required for head shifting.

To improve the access time of a disc memory, it is possible to limit the area serviced by one head to an annulus on the disc, thereby limiting head travel. In such a multiple-head structure (Fig. 5), each head services a number of annular tracks covering a band equal to the separation of adjacent heads. The entire multiple-head structure is shifted by integral track widths to position any one head on a desired track. Only one head can be brought over a specified track at a given time.

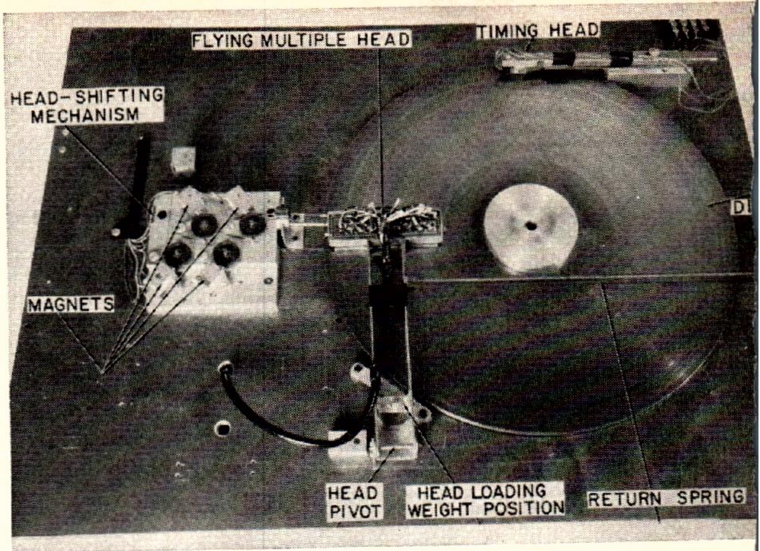

FIG. 5. DISC and multiple-head structure.

The multiple-head assembly of Fig. 5 contains 10 heads, and the entire assembly can be positioned in ten discrete track locations by a combination of electromagnets and a spring return. This gives a total capacity of 100 tracks, the track density being 32 per inch in this design. Disc diameter is 13.25" and thickness 0.260".

Magnetic discs will continue to play an important role as large-scale memory devices for data-processing applications. As recording density increases, the cost of storage per bit will decrease.

Discs working from the 60-cycle power line have an optimum access time of 17 milliseconds, provided no head-shifting is required. It is possible to obtain fairly fast shifting mechanisms at reasonable cost, provided the accuracy for track-positioning is not too demanding.

As in any other problem, the design of magnetic-disc memories involves a compromise between performance and cost. The high-performance features are desirable where discs are to take the place of main memories in small or medium-sized machines. However, in places where discs are considered as replacements for magnetic tape, the high performance may not be required.

Bibliography

S. J. Begun, *Magnetic Recording*, Murray Hill Books, New York, 1949

A. D. and K. H. V. Booth, *Automatic Digital Calculators*, Academic Press, New York, New York, 1953

J. P. Eckert, "A Survey of Digital Computer Memory Systems," *Proceedings of the IRE*, Vol. 41, No. 10, pages 1393-1406, October 1953

J. S. Hill, "Design Features of a Magnetic Drum Information Storage," paper delivered at the Assn. for Comp. Mach. Conf., Rutgers University, March 28-29, 1950

Jacob Rabinow, "The Notched-Disc Memory," *Electrical Engineering*, Vol. 71, No. 8, pages 745-749, August, 1952.

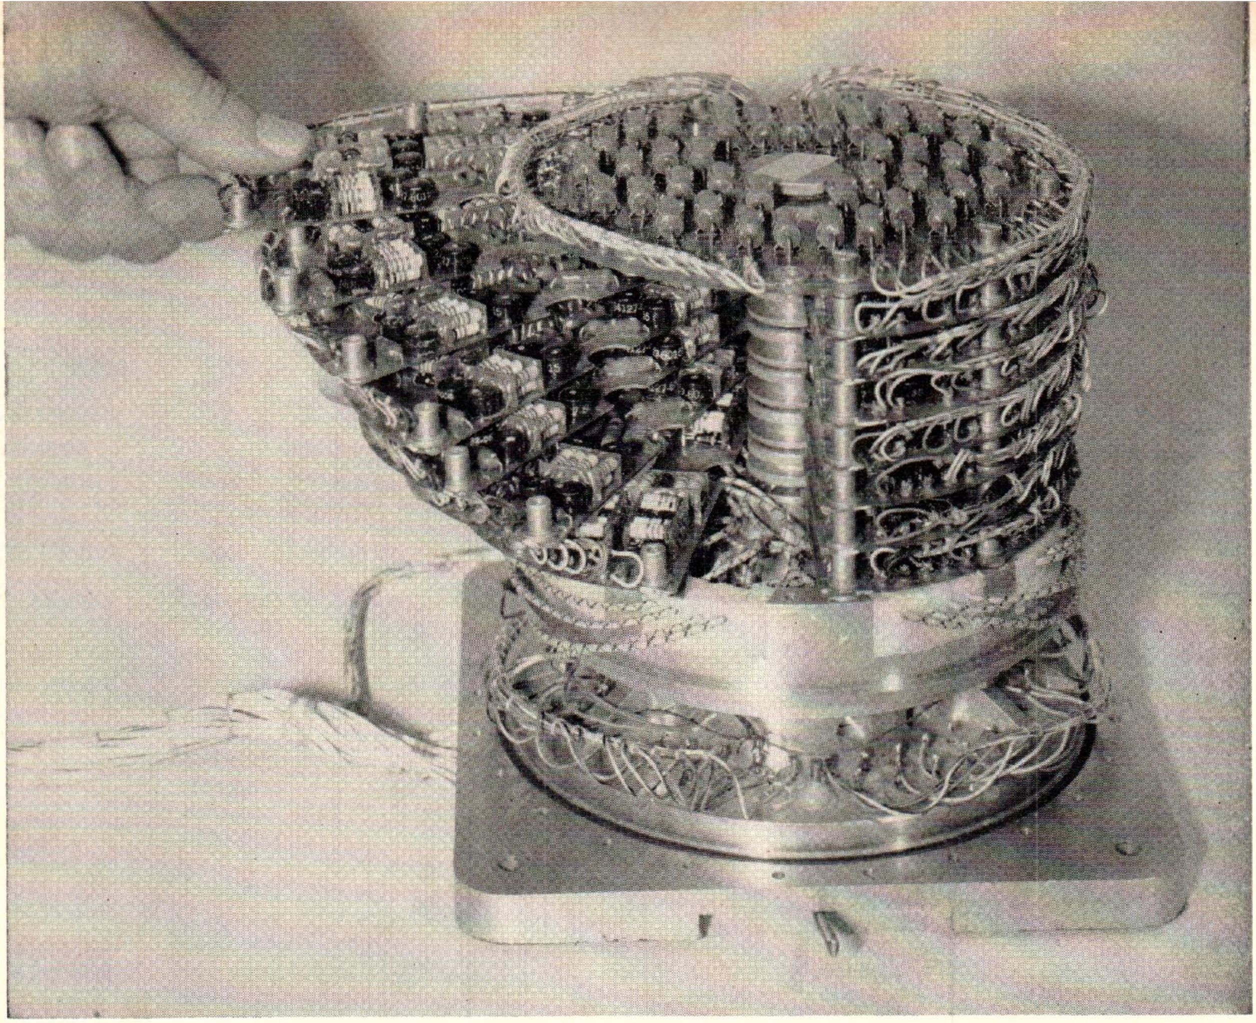

BERNOULLI DISC data storage device weighs only 6 pounds.

ROBERT C. NELSON
Instruments Publishing Co.

MAGNETIC DRUMS AND DISCS

Magnetic drums and discs provide the most popular and economic means for mass storage of digital data requiring access times of milliseconds. Here is a survey of available equipment.

MAGNETIC DRUMS OR DISCS, coated with magnetic oxide and spinning rapidly, provide the memory function for most data-processing equipment. Input-output equipment can link a memory with calculating stages, with remote facilities, and with magnetic-tape files.

Latest advance in the art provides heads that are forced toward the recording surface, but do not touch the surface when at speed due to the resistance to shear of a thin boundary film of air on the rapidly-moving surface. The close proximity of the head gap to the recording surface permits increased bit packing densities and stronger read signals.

One manufacturer fixes heads flush with a flat plate and rotates a thin, flexible disc. The disc cannot touch

TABLE I—SPECIFICATIONS FOR MAGNETIC DRUMS AND DISCS.

Company	Type	Capacity (Bits)	Tracks	Average Access Time (millisec)	Size (dia., length)	Speed (rpm)	Technique
Autonetics							
Recomp II	Disc	160,000+	64	8.7	11" dia	3450	NRZ
Recomp III	Disc	160,000+	64	9.3	11" dia	3450	NRZ
Bendix	Drum	63,104	20	14.5	12" x 4"	1800	RZ
Bryant							
A3000	Drum	to 85,000	to 75	10.0 to 5.0	3" x 3"	6000-12000	Optional
A3000A	Drum	to 120,000	to 50	7.5 to 5.0	3" x 3"	8000-12000	RZ & PM
4000	Disc	to 720,000,000	736/Disc	100	39" Dia.	900- 1200	RZ & PM
5000	Drum	to 570,000	to 280	33.3 to 2.5	5" x 5, 8, 14"	1800-24000	Optional
5000A	Drum	to 1,000,000	to 250	10.0 to 5.0	5" x 5, 8, 14"	6000-12000	RZ & PM
5000HD	Drum	to 2,850,000	to 660	10.0 to 5.0	5" x 5, 8, 14"	6000-12000	PM
7500	Drum	to 860,000	to 280	33.3 to 5.0	7.5" x 5, 8, 14"	1800-12000	Optional
7500A	Drum	to 1,300,000	to 250	16.7 to 7.5	7.5" x 5, 8, 14"	3600- 8000	RZ & PM
7500HD	Drum	to 4,200,000	to 700	16.7 to 7.5	7.5" x 5, 8, 14"	3600- 8000	PM
A8000A	Drum	to 950,000	to 130	16.7 to 7.5	8" x 8"	3600- 8000	RZ & PM
10000	Drum	to 1,720,000	to 420	33.3 to 10.0	10" x 5, 8, 14, 19"	1800- 6000	Optional
10000A	Drum	to 3,570,000	to 420	16.7 to 10.0	10" x 5, 8, 14, 19"	3600- 6000	RZ & PM
10000HD	Drum	to 7,800,000	to 920	16.7 to 10.0	10" x 5, 8, 14, 19"	3600- 6000	PM
18600	Drum	to 6,200,000	to 825	66.7 to 33.3	18.5" x 24, 34"	900 & 1800	Optional
18600A	Drum	to 13,000,000	to 825	33.3 to 16.7	18.5" x 24, 34"	1800 & 3600	RZ & PM
18600HD	Drum	to 30,000,000	to 1800	33.3 to 16.7	18.5" x 24, 34"	1800 & 3600	PM
C-105	Drum	to 150,000	to 50	33.3 to 16.7	10" x 5"	1800 & 3600	Optional
Burroughs							
59279	Drum	11,500	18	1.4	1.75" x 2.3"	21,600	NRZ
51189	Drum	300,000	80	8.5	12" x 16"	3,570	NRZ
Clary	Drum	18,000	14	4.24	6.25" x 1.6"	3,540	NRZ
Consolidated Controls	Drum	to 20,000	to 100	Not Applicable	15" x 36"	Geared to Process	RZ
Ferranti							
200B	Drum	18,000	20	1.33	2" x 1.2"	22,500	Doublet, RZ
217A	Drum	50,000	38	2.67	3" x 1.875"	11,250	Doublet, RZ
IBM		(Characters)					
350	Disc	5,000,000 or 10,000,000	100 or 200	700	24" dia	1,200	NRZI
305	Drum	2,400	24	5	8" x 4"	6,000	NRZ, and biased discrete spot
355	Disc	6,000,000 or 12,000,000	100 or 200	700	24" dia	1,200	NRZI
650	Drum	20,000 or 40,000	40 or 80	2.4	14" x 16"	12,500	Discrete spot
1301	Disc	28,000,000 to 56,000,000	280	160	24" dia	1,800	NRZI
1405	Disc	10,000,000 or 20,000,000	200	600 or 700	24" dia	1,200	NRZI
733	Drum	8,192 36-bit words	256 8-word sectors	12.29	13" x 5½"	2,500	Discrete Spot
734	Drum	60,000	300	8	10.7" x 12.5"	3,750	Discrete Spot
7300	Disc	6,000,000 or 12,000,000	100 or 200	213 min 1013 max	24" dia	1,200	NRZI

the plate at speed because of this film of air.

Another maker does not rotate his drum at speed, but couples it to machinery. An instruction given the memory is read out at a predetermined machine-time later by a head that senses magnetic polarity rather than the change in magnetization.

Company	Type	Capacity (Bits)	Tracks	Average Access Time (millisec)	Size (dia. length)	Speed (rpm)	Technique
LFE							
BD-10	Disc	25,000 to 50,000	12	½ rev.	5.75" x 5.75"	6000-12000	Phase* NRZ or RZ
BD-40	Disc	40,000	32	½ rev.	7.25" x 7"	7700	Ditto
BD-100	Disc	100,000 to 200,000	40	½ rev.	9" x 9"	1800-8000	Ditto
BD-200	Disc	200,000 to 800,000	70 to 140	½ rev.	13.5" x 8"	3600-8000	Ditto
BD-500	Disc	500,000 to 2,000,000	140 to 280	½ rev.	17" x 12"	1800-3600	Ditto
HD	Drum	15 million	300	½ rev.	15" x 14"	180	Double Pulse, RZ
Litton							
C 100	Drum	15,000	10	5	3" x ⅝"	6000	Manchester
C 400	Drum	190,000	112	9	8.4" x 8.5"	4000	Manchester
C7000	Drum	300,000	175	5	6" x 11"	6000	Manchester
C 600	Drum	500,000	160	9	8" x 10"	4000	Manchester
Magnetics, A Division of Idaho Maryland							
MD-0903	Drum	750,000	125	½ rev.	8" x 12"	3600	RZ, NRZ
-0910	Drum	100,000	50	½ rev.	5" x 4"	3600	RZ, NRZ
Ramo-Wooldridge	Drum	300,000	114				
Redmond-Fairchild							
3000	Drum	240,000	120	½ rev.	3" x 6"	25,000	RZ
6000	Drum	960,000	240	½ rev.	6" x 12"	12,000	RZ
8000	Drum	1,600,000	320	½ rev.	8" x 16"	6,000	RZ
10000	Drum	2,400,000	400	½ rev.	10" x 20"	4,800	RZ
12000	Drum	3,600,000	480	½ rev.	12" x 24"	3,600	RZ
Remington Rand							
Randex							
Model I	Double Drum	42,000,000	2,000	385	24" x 44"	870	Pulsed Phase Modulation
Model II	Double Drum	198,620,000	4,000	385	24" x 44"	870	Ditto
FH-880	Drum	28,311,552	768	17	31" x 48.5"	1,800	RZ
File Computer	Drum	1,260,000	300	17		1,800	RZ
LARC	Drum	15,000,000	(Address 100) (Parity and Info 500)	68	24.2 x 27.6"	880	RZ
Shepherd							
204	Drum	to 280,000	to 150	½ rev.	4" x 6"	to 8,000	Phase* NRZ or RZ
205	Drum	to 500,000	to 220	½ rev.	5" x 8"	to 8,000	Ditto
206	Drum	to 1,000,000	to 340	½ rev.	6" x 10½"	to 8,000	Ditto
208	Drum	to 1,350,000	to 360	½ rev.	8" x 10"	to 6,000	Ditto
202	Drum	780,000	to 130	½ rev.	12" x 4"	to 3,600	Ditto
Telex							
Series 31	Disc	10,000,000 to 40,000,000	512	158	31" Dia	1,200	Double Frequency NRZ
I	Disc	155,000,000	512	158	31" dia	1,200	Double Frequency NRZ
IA	Disc	155,000,000	512	53	31" dia	1,200	Double Frequency NRZ
II	Disc	622,000,000	512	125	31" dia	1,200	Double Frequency NRZ
IIA	Disc	622,000,000	512	42	31" dia	1,200	Double Frequency NRZ

*Phase Terminology includes Williams, Manchester or double-pulse.

Autonetics
Div. of North American Aviation, Inc.

Autonetics magnetic disc memory (Fig. 1) consists of a magnetic-oxide-coated disc rotating at 3450 rpm below a headplate containing fixed recording and playback heads. Proper separation between heads and disc is maintained by an autolubricated air bearing. This hydrodynamic air bearing, which forms a boundary layer, does not require pumped air. Spring-loading of the shaft keeps the disk away from the headplate when it is not rotating and during acceleration and deceleration. When revolving at operating speed, a solenoid holds it toward the headplate.

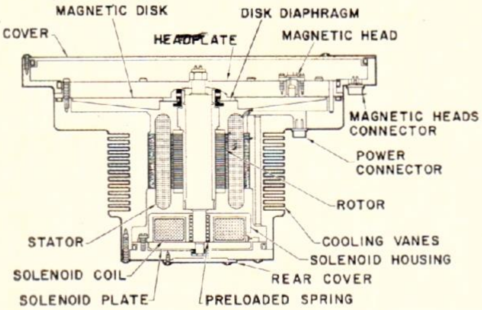

FIG. 1. AUTONETICS magnetic disc recorder in section.

The Bendix Corp.

Information is stored by the Bendix G-15 computer in binary form on a rotating magnetic-drum memory (Fig. 2) with twenty channels of 108 words each. Total capacity is 2176 words; word length is 29 bits. Average access time is 14.5 milliseconds. One drum cycle consists of 108 words or 3,132 bits of information. One word time consists of 29 pulse periods of 9.3 microseconds each. The actual capacity of the drum cylinder is 124 words.

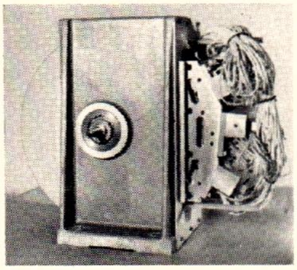

FIG. 2. BENDIX G-15 computer memory drum has twenty channels of 108 words each—a total capacity of 2160 words. Word length is 29 bits.

Fast-access storage consists of 16 words in four channels of 4 words each, with an average access time of 0.54 millisecond. There are also eight words in registers consisting of 1 one-word command register, 1 one-word arithmetic register and 3 two-word arithmetic registers for double-precision operation.

Bryant Computer Products
Division of Ex-Cell-O Corp.

The new Bryant Series 4000 random-access, magnetic-storage disc file (Fig. 3) bridges the gap between high-speed magnetic storage drums and tape memories. It provides 30- to 720-million-bit capacity fast-access memory at 0.1c to 0.025c cost per bit. The disc file is of modular construction, employing both sides of from one to 24 aluminum discs. Each disc face is served by six heads. A rocker arm permits each head to serve 128 tracks. Head positioning takes less than 0.1 second. Average access time after positioning is 34 milliseconds. Maximum total access time is 167 milliseconds. Speed of the disc assembly is 900 or 1200 rpm.

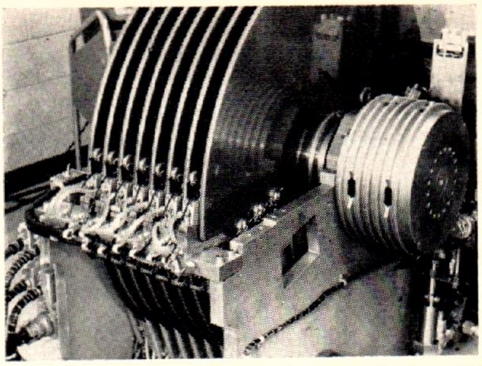

FIG. 3. PROTOTYPE of Bryant's new Series 4000 disc file.

Series 500 drums (Fig. 4) have 5″ diameter and models in this series range from a 5″ to 14″ drum length. All units in the 5000 to 18,500 series are available in various models denoting drum length to encompass a wide range of bit storage. Fig. 5 shows a Model 10,014 with 1,146,800-bit capacity.

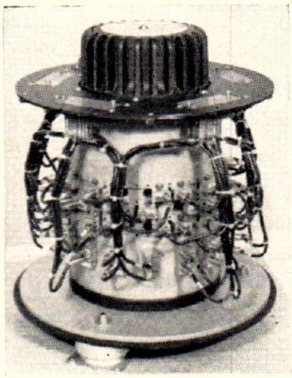

FIG. 4.

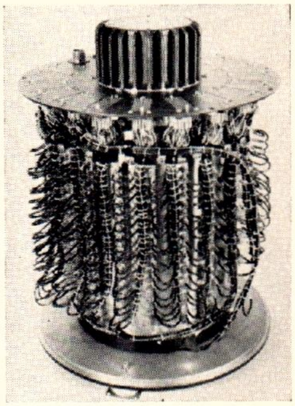

FIG. 5.

The basic drum series features as standard equipment magnetically and electrostatically shielded heads. Recording techniques are optional with pulse densities up to 130 pulses per inch and compatible to either transistorized or vacuum tube circuitry.

The standard series is offered in the—A equipment class, indicating aerodynamic heads. The fail-safe non-contact aerodynamic head provides increased storage capacities with up to 300 pulses per inch recording density at no increase in cost.

Bryant's entry in the economy drum field is the C-105. A conventional-head type having 50 tracks with a diameter of 10″ and 5″ length. Standard speed is 3600 rpm with 1800 rpm optional. The drum is designed to fill commercial requirements and offers data storage at a cost of less than one cent per bit.

Burroughs Corporation

Burroughs Corporation has a small (1.75″ x 2.3″) high-speed magnetic drum (Fig. 6) for intermediate storage. This drum buffers information between the computer and various input-output units and revolves at 21,600 rpm, permitting access to stored data in the average time of 1.4 milliseconds.

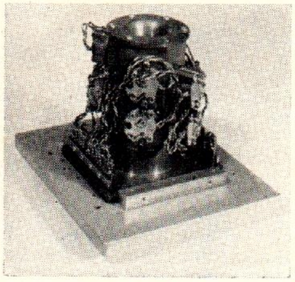

FIG. 6. BURROUGHS Model 59279.

The Model 59279 drum is furnished complete with 20 identical read-write heads. Electrical characteristics permit easy matching with either transistor or vacuum-tube circuitry. Data, clock pulses, and other types of information can be recorded on all channels without restriction.

Drum Model 51189 is a main unit that stores 300,000 bits on a 12″ x 16″ drum. Packing density is 100 bits per inch. A 16-track quick-access memory holds approximately 6000 bits (375 bits per track). Mean access time is 8.5 milliseconds.

Clary Corp.

Clary's DE-60 Memory Drum (Fig. 7) has 6.250″ dia, 1.600″ axial length, and rotates at 3540 rpm. Flux reversals number 130 per inch (standard packing is 65 bits per inch). Record frequency is 50 kc. 14 tracks have nominal track width of 0.075″. Head-to-drum clearance is 0.001″. Minimum register length is 2.8753″; maximum register length is 3.2307″. Magnetic heads are laminated permalloy cores mounted in aluminum cases and potted with epoxy. All heads are shielded with a Mu metal shield surrounding the core to reduce fringing and cross-talk between channels.

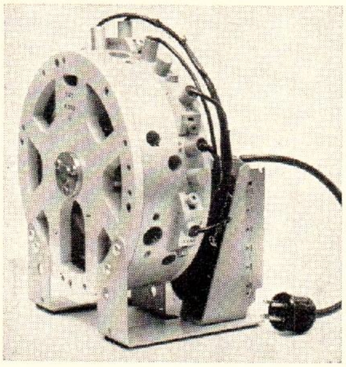

FIG. 7. CLARY's DE-60 drum.

Consolidated Controls Corp.

The Consolidated Dynastat® Digital Magnetic Memory (Fig. 8) uses a different concept than the high-density computer-type drum in which relative motion between the read head and the magnetized record produces an output. The read heads of the Dynastat drum respond solely to the polarity of the recorded magnetic field and yield an output voltage and power which are directly applicable for logic action.

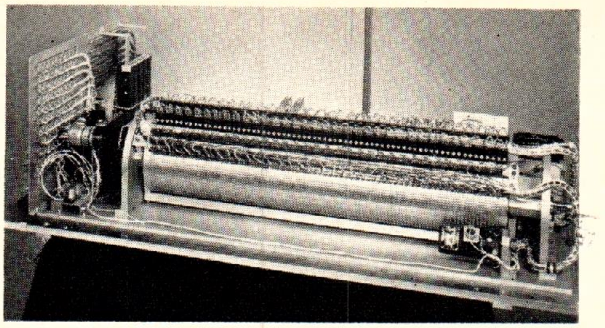

FIG. 8. CONSOLIDATED CONTROL'S slow-speed program drum.

The read heads are self-contained bistable magnetic amplifiers which are triggered by an external magnetic field of a particular polarity. In recording (or writing), the magnetic material on the drum is magnetized with one polarity for "zero" and with the opposite polarity for "one". When the read head perceives a "zero" polarity record, the output remains at minimum noise level. When a "one" polarity spot is presented, the read head snaps to its on condition, and full output voltage is presented.

For operation into transistor circuitry the "one" output is 1.5 vdc into 1000 ohms; the zero output is less than 30 millivolts. For operation into a cold cathode matrix the "one" output is 65 volts ac rms; the zero output is less than 2 v rms.

A multiple-channel drum can be directly driven in synchronism with a conveyor to store parcel addresses keyed in by a dispatcher until the parcel reaches junction points. The drum also can be moved in steps to provide position-to-position control of machine tools.

Ferranti Electric, Inc.

The Ferranti-Packard Type 200B Magnetic Storage Drum (Fig. 9) may be used in small desk-type or other miniature computers. Dimensions are 6″ x 8″ over-all. Rotor surface run-out is less than 0.0001″. A circular glass case allows visual inspection of desiccant in sealed housing. The drum is normally supplied with 20 heads, 18 for data track and two for clock track use. Maximum digit density of 900 bits per track is best obtainable using a drum speed of 11,250 rpm. At full rated speed (22,500 rpm), digit densities of

FIG. 9. FERRANTI-PACKARD Type 200B drum.

450 bits or less, per track, are recommended. Sync tracks can be permanently engraved on the drum surface up to digit densities of 100 engravings per inch. The drum is driven at 22,500 rpm by a 400-cps, 117-volt, 3-phase delta supply. The rotor may be operated over a wide range of speeds by varying the supply frequency from 60 to 400 cps in combination with a proportionate change in voltage.

Type 217 Magnetic Memory Drum has been designed for use in an airborne computer.

Type 364 is a large size magnetic storage drum.

On large-diameter magnetic storage drums, air servos (Fig. 10) are mounted at each end of the head rack to maintain a constant head-to-drum spacing of 0.0005", despite minor drum dimension changes. This provides improved bit density and a 6-db improvement in output signal.

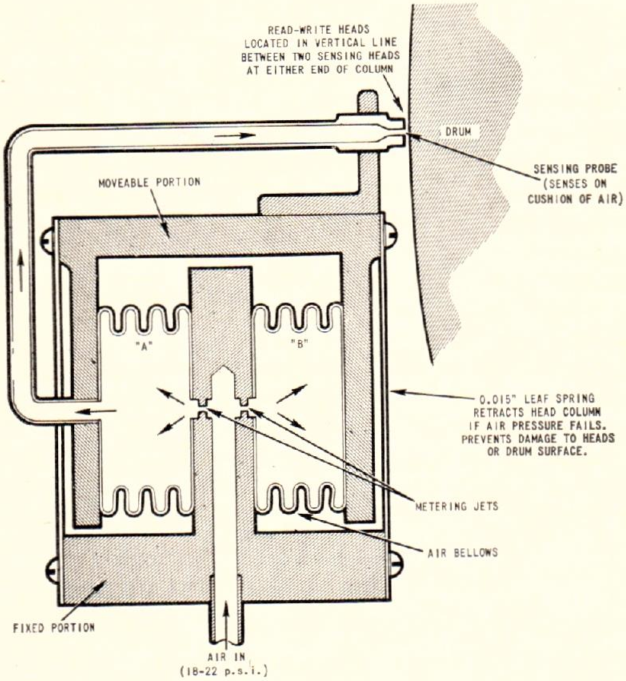

FIG. 10. AIR SERVO maintains 0.0005" head-to-drum spacing.

IBM

IBM's new 1301 disc storage unit (Fig. 11) can be linked to an IBM 1410, 7070, 7074, 7080 or 7090 computers, or shared by any two of them. Depending on the computer employed, each 1301 unit holds from 21,740,000 to 43,480,000 8-bit characters or from 28,000,000 to 56,000,000 6-bit characters. A total of five units can be used with any of the five computers. The 1301 has a read/write head for each disc surface fixed at the end of comb-like access arms.

Related data is stored vertically on separate discs by different heads in the same group providing "data cylinders" that can be reached with little or no horizontal access movement of the head group. Data in the same "data cylinder" can be accessed in an average of 17 milliseconds. Data located anywhere in the file can be transferred to core storage in a maximum of 180 milliseconds. While reading in the same data cylinder, as many as 102,000 characters in a module, 204,000 characters in one disc storage unit, or 1,020,000 characters in a maximum system of five units can be reached with no access motion. These figures vary slightly with the computer used.

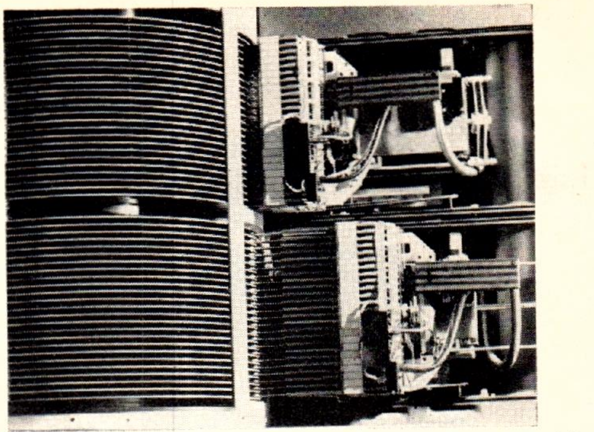

FIG. 11. NEW IBM 1301 disc storage unit. Upper module of 25 discs is being searched for data towards center of discs. Lower bank of access arms has moved outward positioning read/write heads at outer edge of discs.

The usual specification of average access time is altered in describing disc units with movable heads. First time is required to locate the desired track and disc; this is *seek time*. Then time is required before the information is delivered into the processing area; this is *access time*. (These are lumped in the table.)

The 350 disc is used with the RAMAC 305 data processing system to provide up to 20 million characters maximum storage by combining units. Each unit contains 50 discs, with 100, or 200 tracks on each side of each disc. Using a single access arm provides an average access time of 100 milliseconds, a seek time of 600 milliseconds. An optional second arm can make the seek time negligible.

The 355 disc unit can be ganged up to four to provide 48 million characters for a RAMAC 650 data system. Single access arm makes average seek time 600 milliseconds, access time 100 milliseconds. Three access arms are standard to eliminate seek time, access time remains 100 milliseconds.

The 650 drum, used with 650 systems, provides storage capacities of either 20,000 or 40,000 digits on same size drum. Average access time is 2.4 ms.

The 1405 unit can have 25 discs (10 million characters), or 50 discs (20 million characters) in the same housing. The 1401 system uses one 1405 unit; the 1410 system can have up to five, to provide a total storage of 100,000,000 characters. Average access time is 100 milliseconds. Average seek time for one access arm on 25 discs is 500 milliseconds, 600 milliseconds on 50 discs.

The 733 drum unit contains two distinct physical drums, each with a storage capacity of 4,096 words. Each physical drum consists of two logical drums, each of which has a storage capacity of 2,048 words. Access time for the first word is 12.29 milliseconds; each succeeding word requires only 96 microseconds. The 733 is used with the 704 and 709 systems.

The 734 drum is used with the 705 and 705 III systems. Access time for the first character is 8 milliseconds, 40 microseconds for each succeeding character.

The 7300 disc is used with the 7070 and 7074 systems. Four units can be ganged to store 48 million characters. Each unit has 50 discs. Average access time is 113, 138, or 163 milliseconds, depending on prior operation. Seek time disc-to-disc is from 445 to 850 milliseconds, 100 to 235 milliseconds on the same disc.

Laboratory For Electronics

LFE has developed the Bernoulli Disc (Fig. 12), which can be operated at speeds from 1800 to 12,000 rpm with bit frequencies of 90 kc to 400 kc at 3000 bits per track.

FIG. 12. LFE Model BD103 Bernoulli disc memory.

The basic unit consists of a single disc of thin, flexible, instrumentation-grade Mylar magnetic recording tape mounted on a shaft which spins the disc above a smooth headplate in which read/write heads are imbedded flush with the surface (Fig. 13).

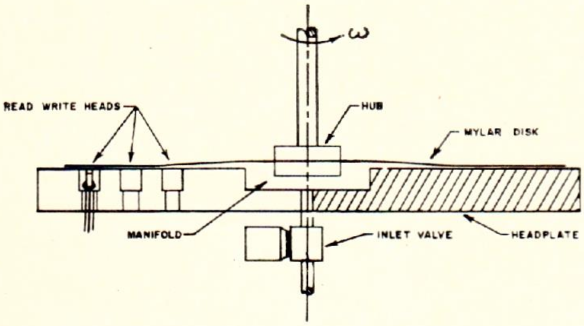

FIG. 13 PARTIAL section of LFE flexible-disc device.

In the non-operating condition, the flexible disc is limp and rests against the headplate. When the drive shaft is rotated, the disc tends to straighten out under the action of centrifugal force and lifts away from the plate. A circumferential velocity is imparted to the fluid in the space between the disc and plate by viscous friction. This circumferential velocity gives rise to a centrifugal force on the fluid and a resulting radial velocity as the air moves out from under the disc. These fluid velocities exhaust the air and reduce the pressure in the separation gap between the rotating disc and the plate. The pressure acting on the open or atmospheric side of the disc forces the disc down. The disc deflection is countered by its own centrifugal and curvature forces in the disc material. The equilibrium operating condition exists as a balance of the pressure forces generated by the fluid flow together with the centrifugal and curvature forces in the disc.

Five separate series of Bernoulli discs are available having the following capacities and sizes. BD-10 Series—25,000 to 50,000 bit capacity, 0.09 cu ft; BD-40 Series—40,000 bit capacity 0.21 cu ft; BD-100 Series—100,000 to 200,000 bit capacity, 0.33 cu ft; BD-200 Series—200,000 to 800,000 bit capacity, 0.66 cu ft; BD-500 Series—500,000 to two-million bit capacity, 2.36 cu ft. All measurements include electronics.

Other specifications of a typical unit, the Model BD-103 (Fig. 9) are as follows:

Storage capacity is 100,000 bits (Max): Bits per track—1000 to 3000: Bit rate is 30 to 400 kc. Typical track layout has 40 total tracks, including 32 data storage tracks, 3 spare tracks, 3 clock and timing tracks, 2 register tracks. Number of registers is 4; register length is 32 bits; register adjustment is ±3 bits. Disc speed is 1800 to 8000 rpm by induction or synchronous motor.

LFE also makes a drum memory. The HD File Drum Unit consists of the file drum, the drive and lubrication systems, a 3-by-10-by-10 track-selection mercury-relay matrix, a linear readout preamplifier, and a final writing amplifier. The file drum, 15" in diameter and 14" tall, is completely enclosed and sealed. Filtered oil is sprayed continuously on the drum for lubrication and to obtain the head-to-drum separation. The fluid motion and action of the oil film obtains a very small and constant head-to-drum separation independent of temperature variations. Any head pair can be removed and replaced without adjustment or loss of information.

The drum is dynamically balanced at 1200 rpm, and ground and lapped by optical techniques to a surface finish better than one microinch rms. The operating faces of each double head are optically ground and hand lapped so that the two surfaces are flat and coplanar to one wavelength of light.

Litton Systems Inc.

Litton's magnetic storage drum (Fig. 14) employs contact read/write heads to achieve large capacity. Heads directly contact the drum until the drum reaches its operating turning speed, at which time surface air movement lifts the heads a few millionths of an inch. Each head is suspended from its mounting bracket by a fine spring that provides constant head pressure. This provides wear characteristics comparable to fixed non-contact heads, but the close proximity of the heads to the drum surface gives efficient electromagnetic coupling. Write currents are reduced by a factor of four to 3.5 to 5 amp turns, compared with 15 ampere turns required with fixed heads. Read signals also are high, with a one-volt signal obtainable.

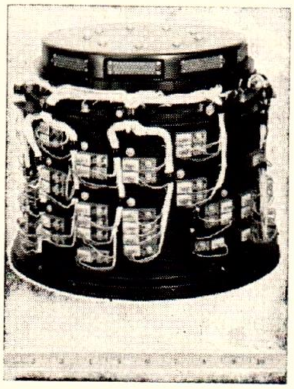

FIG. 14. TYPICAL MAGNETIC storage drum made by Computer Systems Laboratory of Litton Systems Inc. is for Litton's AIDE IV computer; capacity is 190,000 bits.

In acceptance tests, no less than 300 millivolts are accepted.

Greatly improved bit packing densities are achieved, increasing the bit storage capacity of the memory. The close electromagnetic coupling allows the heads to be operated in close physical proximity without objectionable cross-talk. Less than 1% noise is apparent from all sources. Reduction in cross-talk permits one-word recirculating tracks with as few as 18-bit words at 160 bits-per-inch packing density.

Drum end bells and housing are cast from magnesium. Drum is mounted on its selected matched bearings for final turning.

Ramo-Wooldridge Division of Thompson Ramo Wooldridge Inc.

A 2.5″ x 4.5″ magnetic drum memory (Fig. 15), which can store 300,000 bits of information has been developed under contract to the Wright Air Development Center of the USAF. The drum is operated at a bit rate of 546 kc. It stores information on 114 channels, including seven circulating registers of various lengths. The drum assembly is designed to operate to MIL-E-5400. Class 2 specifications, and will withstand shock loads of 16g and continuous vibration of 4g up to 2000 cps.

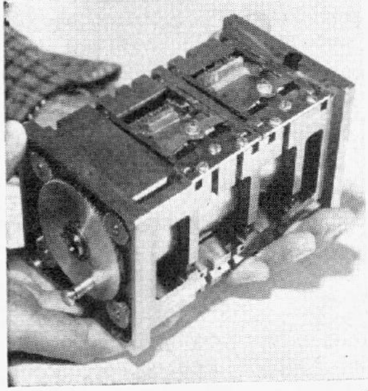
FIG. 15. RAMO-WOOLDRIDGE high-speed drum operates at 550 kc.

Large numbers of record-read heads are mounted in sliding-pad assemblies that are supported on the boundary layer of air above the drum during rotation. 30 record-read heads can be mounted per linear inch side-by-side in a single pad. Head-to-drum spacing is 200-300 microinches.

Redmond-Fairchild, Inc.

Redmond-Fairchild's commercial drum (Fig. 16) has 448 general storage tracks, 6 timing tracks. There are 4 registers with 2 heads each, and 1 register with 8 heads. Average access time is 0.0125 sec. Drum is 10″ in diameter, 24.5″ long, and is rotated at 4800 rpm by a 400-cycle air-cooled motor. Packing density is 200 bits per inch with RZ recording.

Fig. 17 shows another drum for commercial application with provision for 200 tracks. Average access time is 0.0125 sec. Drum is 6¼″ in diameter, 10½″ long. It is rotated at 4800 rpm by a 400-cycle air-cooled motor. Packing density is 200 bits per inch with RZ recording.

A 6¼″-diameter by 4¼″-long drum designed to military specifications has provision for 80 tracks, 90 heads. 6000-rpm motor is totally enclosed.

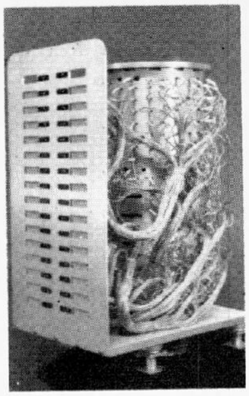
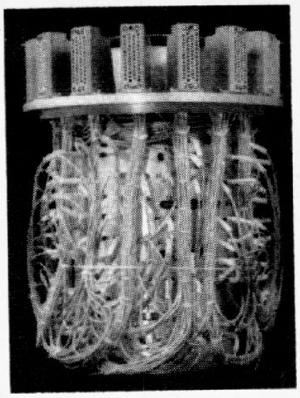
FIG. 16. REDMOND-FAIRCHILD commercial drum has 448 storage tracks (left). FIG. 17. 200-track drum.

Another drum for military application has 250 tracks, 300 heads. Average access time is 0.01 sec. Drum is 12″ in diameter, 12.75″ long, and rotated at 3600 rpm by a 60-cycle air-cooled motor. Packing density is 200 bits per inch with RZ recording.

Remington Rand Univac Division of Sperry Rand Corp.

Univac Randex System Model I provides mass storage for the Univac File-Computer Systems. It consists of a variable number of Randex Drum Units. Each drum unit contains two drums, with each drum having 1,000 tracks dividing into 25 sectors per track. Each sector is capable of storing 120 seven-bit characters. The capacity of each drum is 3,000,000 characters; that of a drum unit is 6,000,000 characters. Randex System capacity may range up to 60,000,000 characters.

Univac Randex System Model II (Fig. 18) provides mass storage for Univac Solid-State Computer Systems. A Randex Storage Unit comprises two drums and is capable of storing over 24,000,000 digits and signs. The system can be expanded to as many as 10 units, providing over 250,000,000 positions of storage. Randex operates completely on-line under program control.

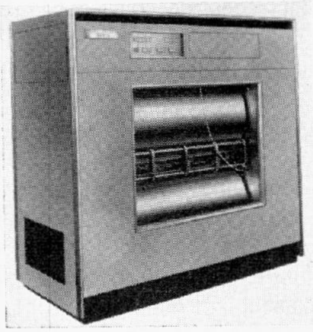
FIG. 18. REMINGTON RAND Randex twin-drum unit stores 24 million digits.

The magnetic drum sub-system incorporating the Univac FH-880 Flying Head Drum is shown in Fig. 19. The FH-880 drum can be used in both the Univac 490 Real Time and the Univac 1107 Thin-Film Memory Computer Systems. The drum read-write heads float on a boundary layer of air, generated by the 1800 rpm rotation of the drum, at 0.0005″ or less from the oxide-coated surface of the drum. There are 128 six-track channels across the length of the drum. Each channel can store 6,144 computer words, allowing a total of 4,718,592 alphanumeric characters to be recorded on the drum. Using the drum's minimum interlace pattern, the word transfer time between the drum and the computer is 16.5 microseconds. Up to eight drums may be connected to a single computer input-output channel.

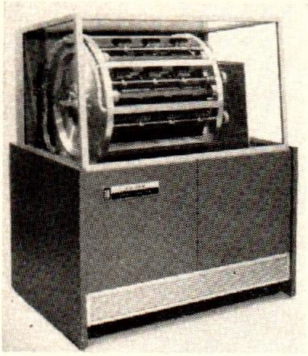

FIG. 19. FH 400 drum storage unit for the Univac 490 real-time computer.

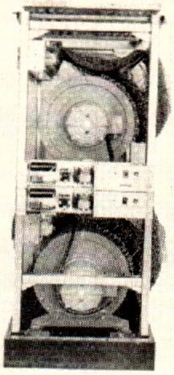

FIG. 20. RANDOM-ACCESS memory for Univac file-computer system Model I.

The Random Access Drum shown in Fig. 20 is used in all Univac File-Computer Systems. From one to ten magnetic drum units can be used. Each unit has a storage capacity of 180,000 characters. A ten-drum system has a capacity of 1,800,000 characters. Average access to any unit record is one-half a drum revolution, 17 milliseconds.

Up to 24 magnetic drums may be included in a Univac LARC System. Each drum is capable of storing 250,000 words of 12 decimal digits each. Synchronizers may be added to the processor for simultaneous control of as many as three reading and two writing operations on the drums, concurrently with input-output operations. Two drums, operating alternately with a single synchronizer, can transfer data at a continuous rate of 330,000 decimal digits per second. Average access time is 68 milliseconds, which includes the time for the lateral movement of the read-write head. A complete 250,000-word drum is serviced by a single six-channel read-write head assembly which is floated on a thin film of air. The "floating" head assembly enables the heads to record and read at a density of 450 pulses per inch.

Shepherd Industries, Inc.

Shepherd Industries' Model DD-200 Magnetic Memory Drum Series (Fig. 21) are designed to store up to 1,350,000 bits of information. Diameters range from 4″ to 12″, with coated length from 1″ to 10″; speeds are from 600 to 8,000 rpm.

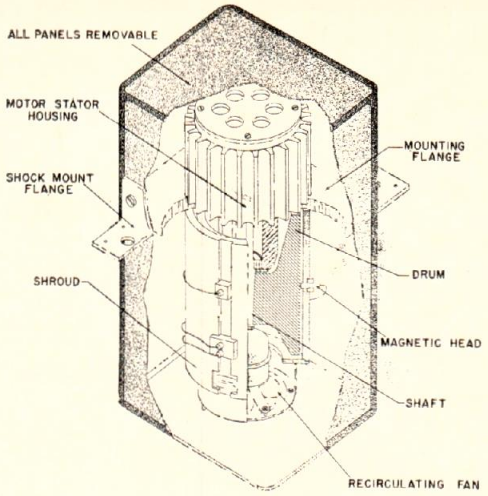

FIG. 21. SHEPHERD Model DD-200 drum.

When used with Shepherd Industries high-performance drum heads, bit densities to 140 phase bits (280 flux reversals) per inch can be obtained. A variety of track widths and head mounts are available. All head mounts permit individual head-to-drum spacing adjustments to insure uniformity of output and resolution. The total number of general storage tracks will depend on the track width and on the number of recirculating loops. Induction or hysteresis synchronous motors can be specified. Some of the units are designed also for airborne use.

Telex, Inc.

The Telex I Disc File (Fig. 22) has a capacity of 155 million characters on 16 discs. Access time is less than 160 milliseconds. The airfoil-like design of the read-record heads allows the heads to ride on the air boundary layer of the rapidly-revolving discs.

The recording discs (Fig. 23) are spaced on a vertical, motor-driven shaft. The read-record heads are mounted in opposed pairs on a two-pronged positioning arm for each disc, so that two pairs of heads serve the upper surface of the disc, the others the lower surface. The positioning arms are driven by motors with the track-locating dimensional references built into the magnetic field of the motors.

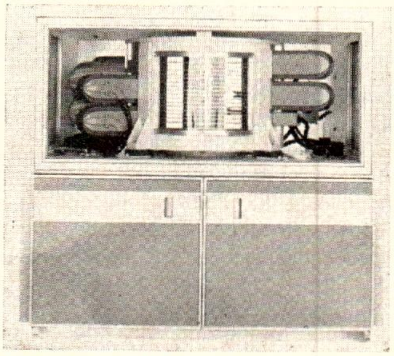

FIG. 22. TELEX I mass memory module.

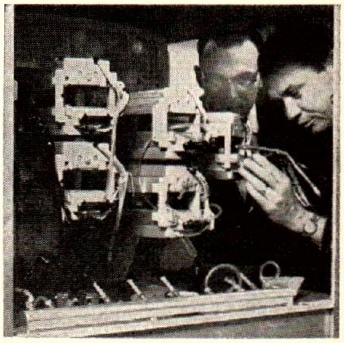

FIG. 23. TELEX head-positioning arms.

Since each disc has its associated head positioner and heads, no time is spent in moving the read-record heads from disc to disc. In actual operation, two or more of the head positioners may be in motion at the same time on the Model IA and IIA versions.

The Telex II, with a capacity of 620 million bits, uses 64 recording discs and head positioners, each with four pairs of read-record heads. All head positioners are connected to common address lines through magnetic reed switches so that only the selected motor is driven.

ADDITION TO I&CS MAGNETIC-DRUM SURVEY

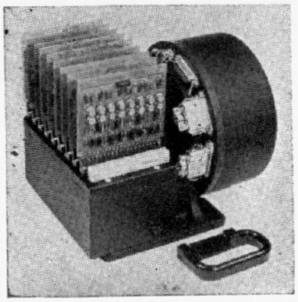
COGNITRONICS 4-20.

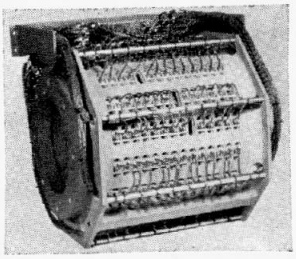
MAGNE-HEAD drum.

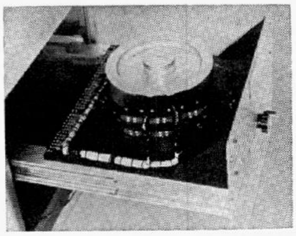
ISC drum.

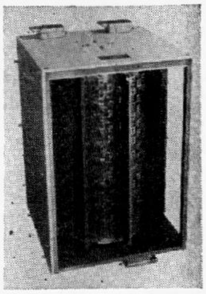
VERMONT drum.

Cognitronics Corp.

The Cognitronics Model 4-20 Magnetic Storage Drum (Figure) provides pulley and belt drive for selecting drum speeds up to twenty-five thousand rpm. The drum has a capacity of up to fifty (50) tracks with individually adjustable heads. The heads are normally spaced 0.0005″ from the drum and employ a record current of about 100 ma through either coil. A normal playback signal of 40 mv is obtained across ½ winding at 2400 ips.

The digital circuitry is designed to operate with an input and output level of ground (0 volt) for a one, and -6 volts for a zero, at frequencies up to 300 kc. Drum circumference is 13″, length is 2″. Bit density is up to 300 ppi with NRZ recording. Capacity is 120,000 bits.

General Instrument Corp.
Magne-Head Div.

Magnetic memories by Magne-Head are available with contact (air floating) or non-contact (fixed spacing) heads, and plated nickel-cobalt or sprayed iron oxide magnetic coatings. Dynamically balanced drum has runout under 0.00005″ and is made from magnesium or aluminum materials. Clock rates to 1 Mc; pulse packing density to 500 bits per inch Manchester.

Modular circuitry can accept and respond to serial or parallel data in return-to-zero (RZ), non-return-to-zero (NRZ), phase modulation or frequency modulation modes.

Instrument Systems Corporation

Instrument Systems' Drum Model X 102 contains 4 channels. The drum is 9″ in dia, with a face width of 1¾″, and rotates at 8⅓ rpm. Surface velocity is 4 ips; a period of revolution is 7.2 seconds. Two heads are adjustable over a range of relative delays extending from slightly less than zero to a maximum of about 6 seconds.

Heads may be operated at spacings from the drum surface of 0.0005″, or less under special conditions. In general, a distance of 0.0007″ has been found satisfactory for electrical performance. Standard head gap is 0.001″, fixed by a silver shim. Inductance of both legs is 15 mh, one leg 4.7 mh at 1000 cps. Q is 2.7 at 1000 cps. Other combinations may be chosen.

Vermont Research Corp.

Vermont Research Corporation offers a line of drums, including a 2,000,000-bit magnetic drum memory (Figure) that provides an average access time of 16 milliseconds to any bit of information. Other features include: Twenty-bit parallel recording, simultaneous input to, and output from, the memory, character transfer rate of 375,000 per sec. Diode matrix boards are included within the drum enclosure. Connections from 500 heads are reduced to three 50-pin connectors. Outputs from all heads are held to a ±10% tolerance. Unit is designed for mounting in a 19″ rack.

TABLE I—SPECIFICATIONS FOR MAGNETIC DRUMS

Company	Type	Capacity (Bits)	Tracks	Average Access Time (millisec)	Size (dia. length)	Speed (rpm)	Technique
COGNITRONICS CORPORATION 4-20	Drum	120,000	50	½ rev.	4" x 2"	to 25,000	NRZ
INSTRUMENT SYSTEMS CORP.	Drum	Pulse delay	4	Delay to 6 sec	9" x 1¾"	8⅓	NRZ, RZ
MAGNE-HEAD DIV./GENERAL INSTRUMENT CORP.	Drum	to 50 million	to 30/in	to 1.25	to 24" x 48"	to 24,000	Manchester (Ferranti) or NRZ
VERMONT RESEARCH CORP. 100 Series	Drum	to 2,500,000	550	5	10" x 4", to 10" x 20"	6000	RZ
500 Series	Drum	to 1,250,000	425	1.25	5" x 4", to 5" x 17"	24,000	RZ

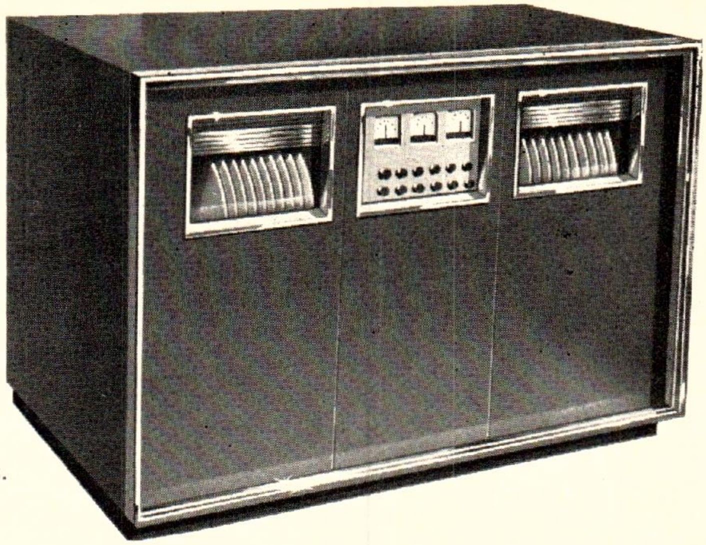

FIG. 1. **BRYANT'S** new Series 4000 Disc Files.

DISC FILE MEMORIES

HAROLD J. McLAUGHLIN

Bryant Computer Products

How do you get data into and out of a disc-file-type memory? This article describes a variety of applications of the disc file type memories to random-access data processing. Serial and parallel operation, as well as discrete clocking, selective alteration and interlacing, are included. Frequency and capacity can be varied over a wide range. This is number 9 in a series on data storage.

DISC FILES are becoming increasingly popular for data storage where access to large amounts of data is required in milliseconds.

The 4000 Series Disc Files (Fig. 1) are memories comprising 39″-diameter magnetic discs rotating on a common shaft at 900-1200 rpm. The maximum-capacity equipment (Model 4240) has 24 discs (48 recording surfaces).

Each surface is serviced by six magnetic heads, each of which can assume any one of 128 positions, giving a total of 768 concentric recording tracks.

A hydraulic positioning mechanism moves all 288 heads (for the 48 surfaces) simultaneously to a selected track within 100 milliseconds. This 100-millisecond positioning operation provides access to 288 tracks of data, or approximately 6,000,000 bits. Switching between the 288 heads is accomplished electronically in 10-50 microseconds.

The length of each of the 768 concentric tracks on a disc surface depends on its radius. Maximum bit loading of the surface would dictate use of the maximum allowable bit density in every track. However, this would result in a separate capacity and frequency for each track. Therefore, some bit capacity is sacrificed to reduce the number of track capacities and frequencies to a reasonable minimum. This is accomplished by grouping the tracks in concentric zones, in each of which the innermost track uses the maximum allowable bit density (or a convenient lower density), and all other tracks have this same capacity and frequency. A maximum of six zones are used on each surface, one for each head.

A typical six-zone arrangement is shown in Fig. 2. Using a maximum bit density of 270 bits per inch and a disc speed of 900 rpm, this arrangement achieves the frequencies shown, and a total surface capacity in excess of 15,000,000 bits.

Another (and perhaps the most popular) disc arrangement is shown in Fig. 3. This arrangement has two zones per surface, each served by three heads. A maximum bit density of 275 bits per inch is used in the inner track in the inner zone, and a maximum of 285 bits per inch in the inner track of the outer zone. This produces an inner-zone track capacity of 11,200

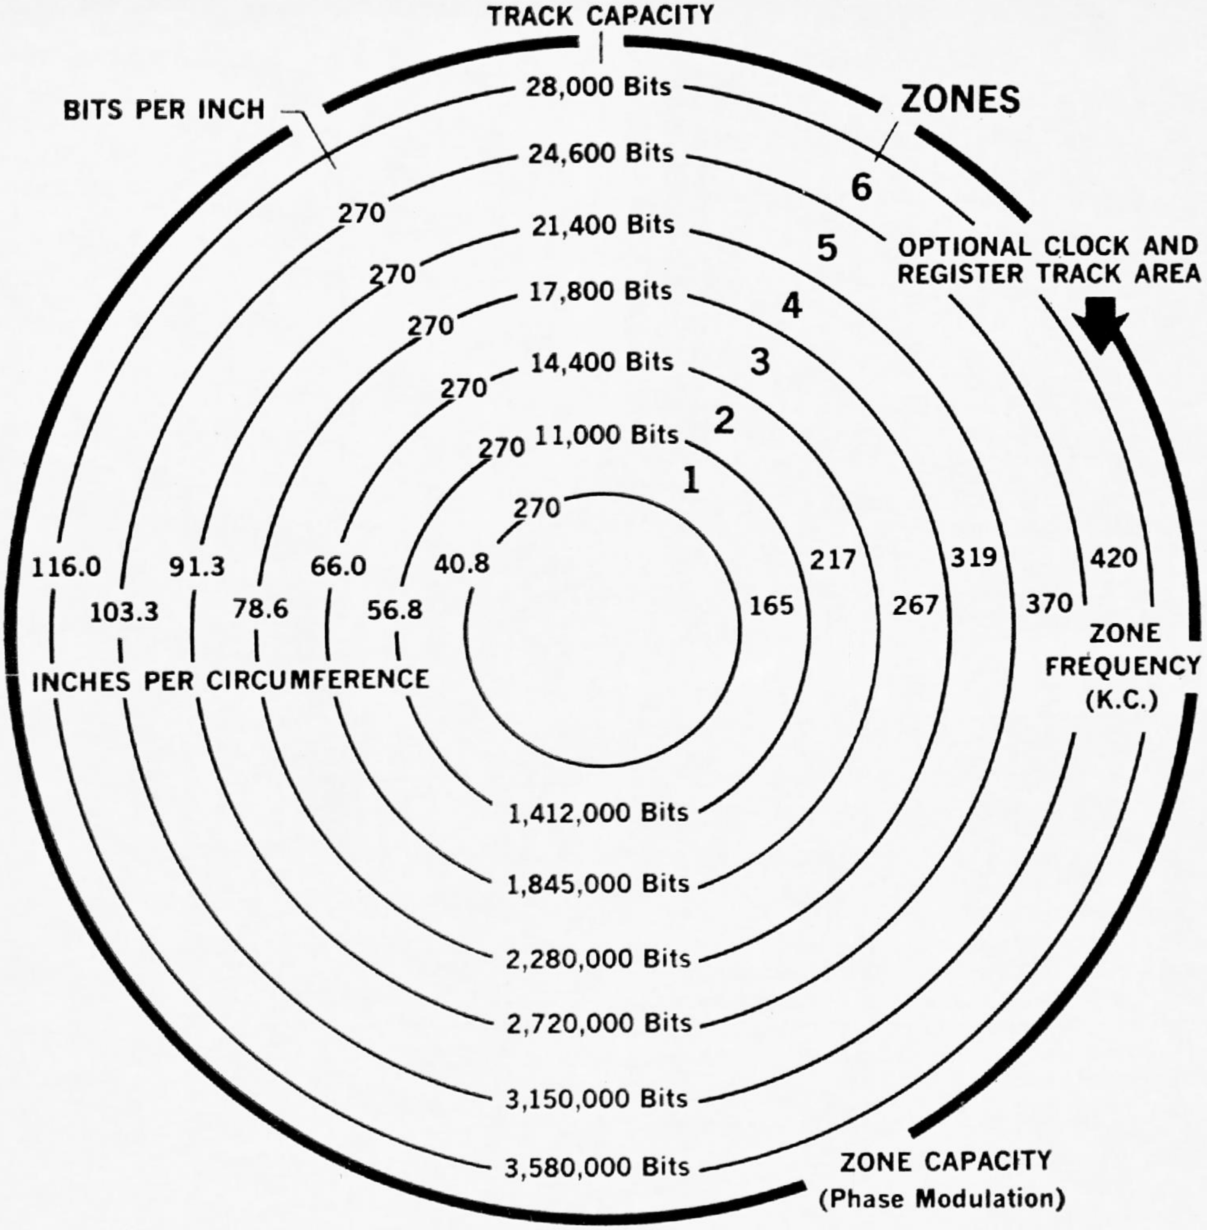

MAXIMUM CAPACITY TRACK LAYOUT

FIG. 2. SIX HEADS are used to service each disc. Each zone, serviced by one head, has its own maximum bit density. Each track in zone 6 has 28,000-bit capacity, giving a zone capacity of 128 × 28,000, or 3,580,000 bits. Each zone has its own timing track (zone frequency).

bits per track. Outer-zone capacity is 22,400 bits per track—exactly double the capacity of the inner zone tracks.

In the outer-zone tracks, 2:1 interlacing is used to produce in each physical track two separately-addressable channels (made up of alternate bits in the track) having the same capacity as the inner-zone tracks. This results in a total of 1,152 separately-addressable tracks per disc surface, each having a capacity of 11,200 bits and a data frequency of 224 kc at 1200 rpm.

Many other track arrangements can be used to satisfy individual system requirements for track length, data frequency, capacity, etc.

Discrete Clocking

All data is recorded in *phase-modulation type** recording, and is discretely clocked—that is, each bit is recorded in a predetermined cell as defined by a prewritten timing track, under control of that timing track. In this manner each bit is precisely positioned and can be erased without affecting adjacent bits, simply by writing a new bit of opposite binary significance in the same cell (under control of the same timing track). This changing of a single bit in a track is known as *selective alteration.*

In addition to eliminating the need for blank spaces between blocks of data, selective alteration simplifies file updating because only changed items, rather than complete data blocks, need be recorded.

Timing Tracks

A timing or clock track, including a precisely recorded signal for each bit position, is supplied for every zone. A single timing-track on one disc serves like zones on all discs.

A timing diagram illustrating the relationships between the timing signals and the data being recorded and played back is shown in Fig. 4. The timing tracks (phase-modulation recorded) control writing of data by gating NRZ-type** information into the recording circuits.

Timing of played-back data is achieved by differentiated timing signals, suitably delayed as indicated in Fig. 4 to compensate for the inherent delays in the playback process. Electronic delays can be used, or separate read heads can be used to produce the read timing pulses. In the latter case the delay is achieved by adjusting the relative positions of the heads.

In addition to the timing tracks, an index track incorporating a single pulse is used to mark the zero (datum position) on each disc. Also, where the data tracks are divided into a number of angular sectors for addressing purposes, sector tracks can be included to mark the boundaries, or can include code words which identify the sectors.

Bit Interlacing

Each track operates at a fixed clock frequency. In order to accommodate binary data rates which are submultiples of this frequency, bit interlacing is used —that is, each track is divided into separately addressable channels, each including every second, third, or Nth bit of the physical track. Each channel, therefore, operates at the frequency of the input system, which is lower than that of the clock.

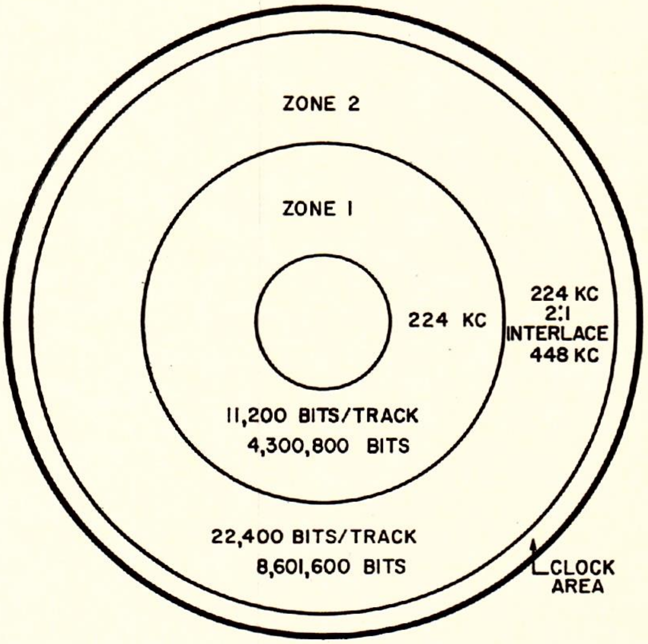

FIG. 3. TWO-ZONE LAYOUT wherein each track in outer zone has twice the bit capacity of the tracks in the inner zone. Both zones have 384 tracks, 3 heads (128 tracks each).

---

*In phase-modulation recording, DC of one polarity is applied to the head during the first half of a pulse period and, at the mid-point of the period, the polarity is reversed for the second half. The order in which the two polarities are applied to the head during each period determines the binary significance.

**NRZ (non-return-to-zero) data utilizes one DC level to represent binary one, and another to represent binary zero.

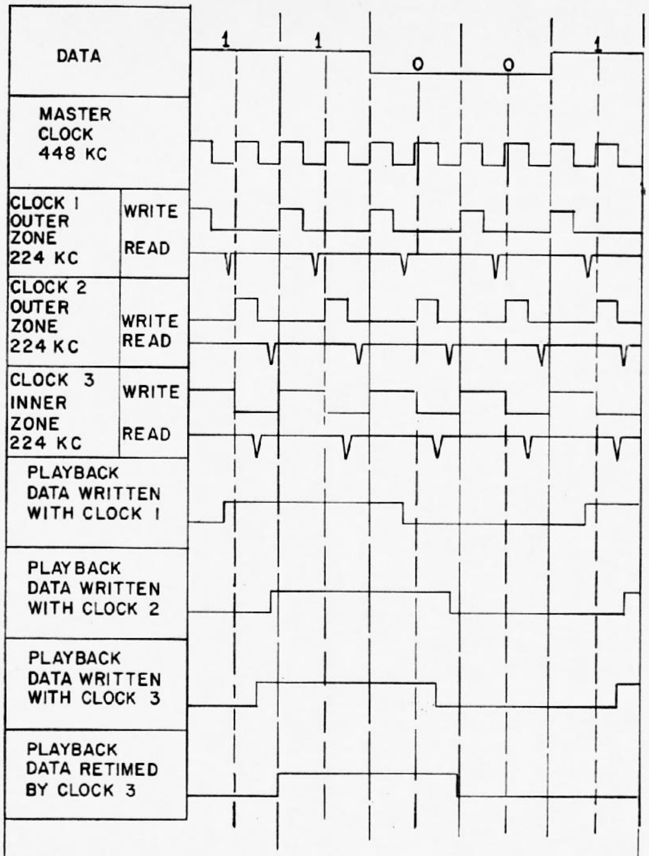

FIG. 4. INTERLACE TIMING diagram for data pattern 11001 as recorded in each interlaced channel (alternate bits of a 48 kc track) and in a non-interlaced, inner zone track (224 kc). Clocks 1 and 2, together, provide a 448-kc clock while clock 3 is at 224 kc. Clock 1 pulses control recording and reading during the first half of each 224-kc period and clock 2 during the second half. Clock 3 controls recording of broader signals across the entire 224-kc period.

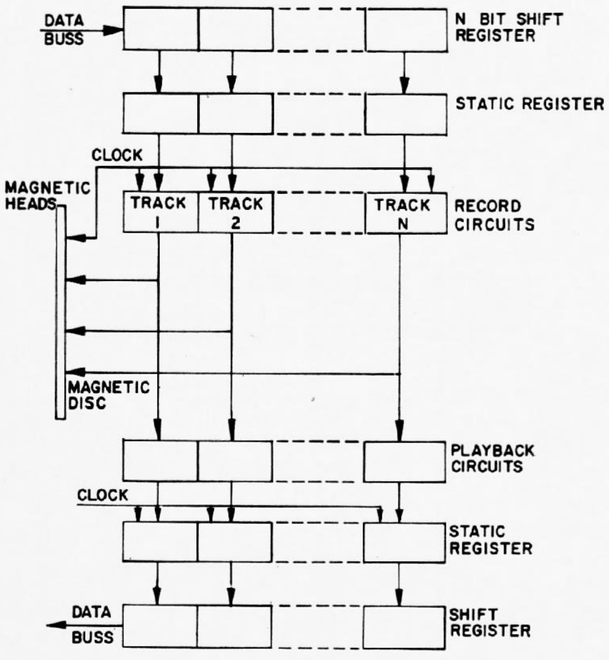

FIG. 5. TRACK INTERLACING block diagram.

Track Interlacing

Data rates which are multiples of the clock frequency are accommodated by using *track interlacing.* This technique uses several tracks, each of which records part of the data. For example, if the incoming data to be recorded has six times the clock frequency, six tracks are used, with one of the six bits being recorded in each track. In other words, N tracks are operated in parallel to record, simultaneously, N bits of a serial data train having a frequency N times that of each track.

Fig. 5 illustrates an input-output arrangement for track interlacing where the frequency of the data buss is N times that of the tracks in which the data is to be recorded. Data is fed into an N-stage shift register at the frequency of the buss. At the end of each Nth-bit period, the contents of the shift register are shifted in parallel to a static register which drives N record circuits in parallel (under control of clock or timing pulses from the disc file). Thus, in every Nth-bit period, N bits are recorded in parallel in N tracks.

Playback (PB) operations are the reverse—that is, the PB amplifiers for the N tracks set the N stages of a static register once every N bit periods. These settings are transferred into a shift register from which they shift onto the data buss at its frequency.

Serial Operation

Fig. 6 is a typical two-zone, single-frequency serial application. This 24-disc system utilizes 2:1 interlacing (Fig. 3) to provide a single track length of 11,200 bits, and a single data frequency of 224 kc throughout the file; 55,296 separately-addressable channels or tracks of 11,200-bit capacity are available. Each track is further divided into ten sectors, each capable of storing 160 seven-bit characters.

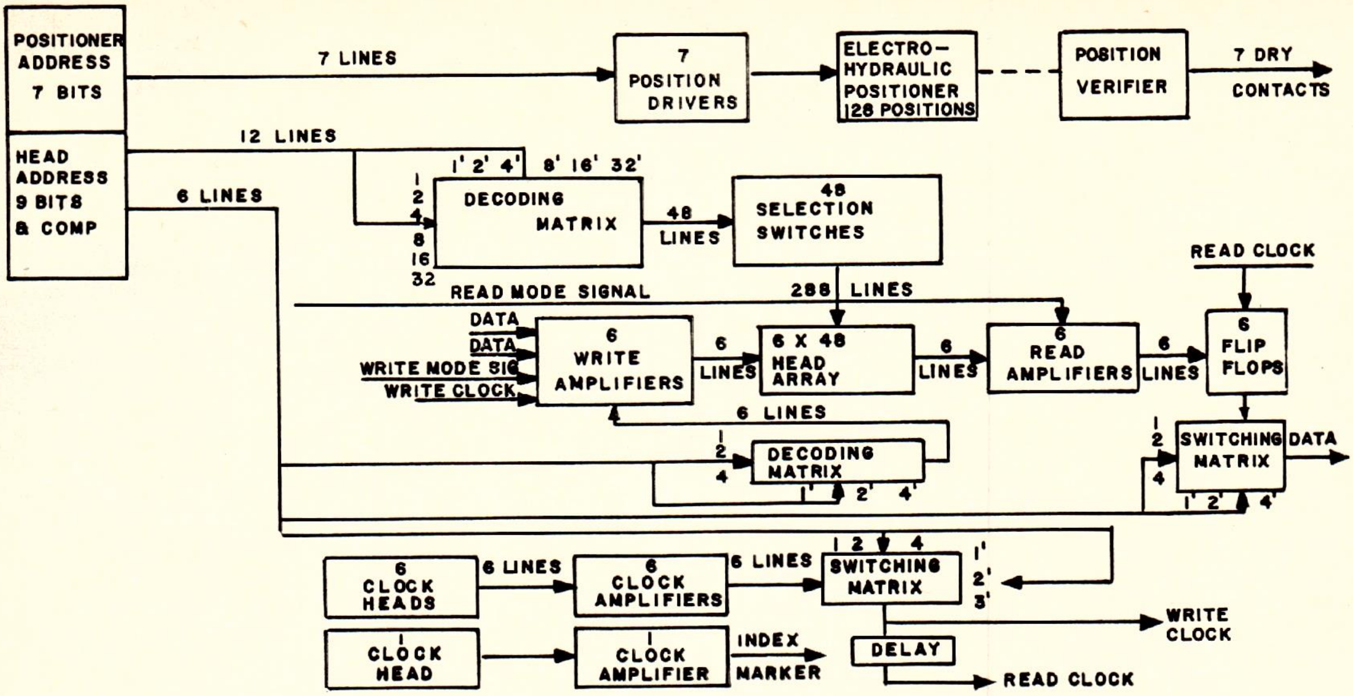

FIG. 6. SERIAL APPLICATION of a disc file. Seventeen bits are used to select proper channel. First seven bits of address direct all heads to proper track (one of 128); next six bits select proper disc; last 3 bits select proper head (one of six). Last bit helps select proper clock.

Addressing is accomplished by a 17-bit word as follows:

1. A seven-bit positioner address selects one of 128 positions for all heads.
2. Six bits of the address select the disc surface by selecting one of 48 groups of six heads. The heads are ordered across the file beginning with the innermost heads. Heads 1 through 48 are the innermost track, and heads 241-288 are the outermost track.
3. Three bits of the address select one of the six write or read circuits in cooperation with "read-mode" and "write-mode" signals. This completes selection of a single head for serial operation.
4. The 17th bit of the address, and one bit of the read-write-amplifier selection group, serve to select the appropriate clock, as illustrated in Fig. 4. Clocks 1 and 2 are used in the interlaced tracks, and clock 3 is used in all of the inner zone tracks. As indicated, clocks 1, 2, and 3 are generated electronically from a master clock.

In order to provide the proper phase relationship between the write and read clocks, a separate read head and gating is used to provide properly delayed read clocks 1 and 2. A third head and gating is used to produce the differently delayed read clock 3.

Data is presented to the file in NRZ form at 224 kc (Fig. 4) under control of timing pulses supplied by the file. This data is "anded" with the selected one of the three operating clocks for recording in the appropriate track. During playback operations the data is read back using the appropriate read clock.

This results in misalignment pulse periods because of the delayed timing of the read strobe pulses. In order to realign the data, it is strobed out of the playback circuit flip-flop into a second flip-flop under control of write clock 3 as shown in Fig. 4.

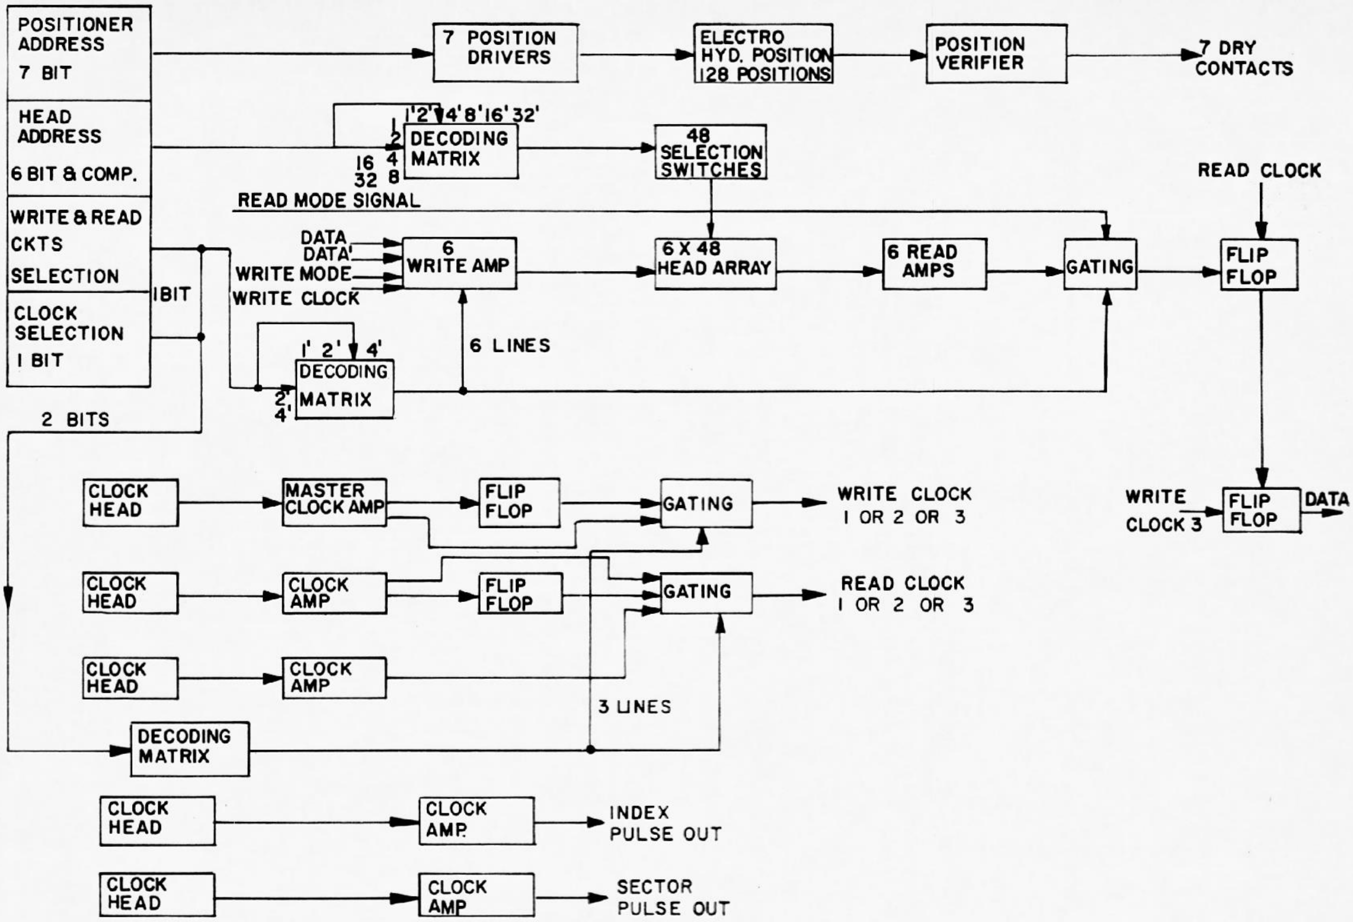

FIG. 7. PARALLEL OPERATION. A 7-bit character is written or read simultaneously. A 15-bit address word is used to select the channel. First 7 bits select 1 of 128 positions for all heads; next 5 bits select 1 of 30 groups of 7 heads; remaining 3 bits select clock track.

Parallel Operation

In parallel operation, N-bit characters are written and read in parallel; that is, all N bits are written or read simultaneously.

Parallel operation of the disc file is illustrated in Fig. 7. This typical seven-bit parallel system uses 35 recording surfaces on 18 discs for storage of 75,000,000, seven-bit characters. Six zones are used on each surface.

To record seven-bits simultaneously, seven heads (selected from the same zone) are used as a group. Thirty such groups of seven heads (five groups per zone) are provided. Any group of seven heads can record seven bits (one character) simultaneously.

Selection of seven tracks for simultaneous read or write operations is accomplished in three stages under control of a 15-bit address word:

1. A seven-bit positioner address selects one of 128 positions for all of the heads.
2. Five bits of the address select one of the 30 groups of seven heads.
3. The remaining three bits of the address are used to select the appropriate clock track.

The programmer and the computer interface designer must compensate for the variations in capacity and frequency from zone to zone, as in serial operation.

Flexible-Disk Storage

ROBERT J. PEARSON

Laboratory For Electronics, Inc.

A flexible-disk magnetic recorder maintains specified head spacing, regardless of vibration.

A flexible-disk magnetic recorder (Fig. 1) has several important advantages. Separations from the headplate of 1 mil or less are readily obtainable without additional peripheral equipment. There is almost complete immunity from shock and vibration because the disk is flexible and does not readily transmit axial movements of the shaft with respect to the headplate. As the disk is light, there is no need for any special procedure or operating sequence when accelerating up to speed or decelerating, even under shock and vibration conditions. The flexible-disk recording medium does not wear when it comes in contact with the headplate during start and stop operation. Over three thousand start and stop cycles have been applied to a test device without any indication of wear on the heads or on the magnetic coating of the disk.

The basic units consist of single disks of thin, flexible, instrumentation-grade Mylar magnetic recording tape mounted on a shaft which spins the disk above a smooth headplate in which read/write heads are imbedded flush with the surface.

In the non-operating condition, the flexible disk is limp and rests against the backplate. When the drive shaft is rotated, the disk tends to straighten out under the action of centrifugal force and lifts away from the stabilizing plate. A circumferential velocity is imparted to the fluid in the space between the disk and stabilizing plate by viscous friction. This circumferential velocity gives rise to a centrifugal force on the fluid and a resulting radial velocity as the air moves out from the under the disk. These fluid velocities exhaust the air and reduce the pressure in the separation gap between the rotating disk and the plate. The pressure acting on the open or atmospheric side of the disk forces the disk down. The disk deflection is countered by its own centrifugal and curvature forces in the disk material. The equilibrium operating condition exists as a balance of the pressure forces generated by the fluid flow together with the centrifugal and curvature forces in the disk. A typical radial section through a spinning disk is shown in Fig. 2.

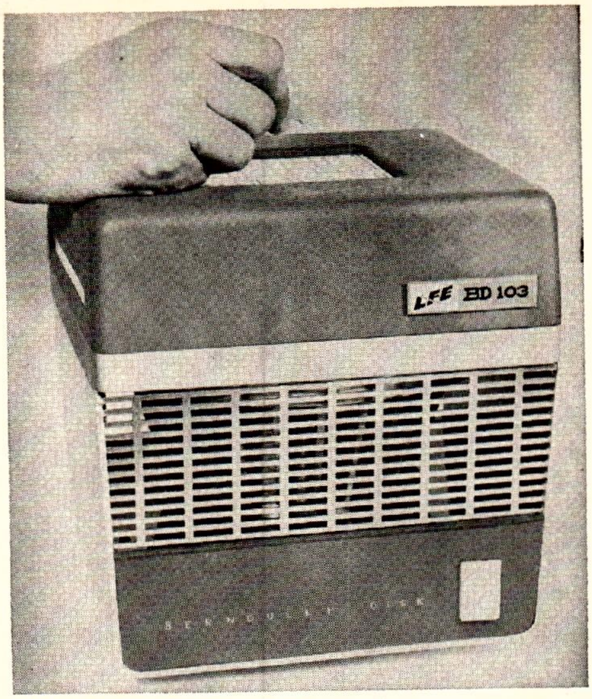

FIG. 1. BERNOULLI DISK recorder.

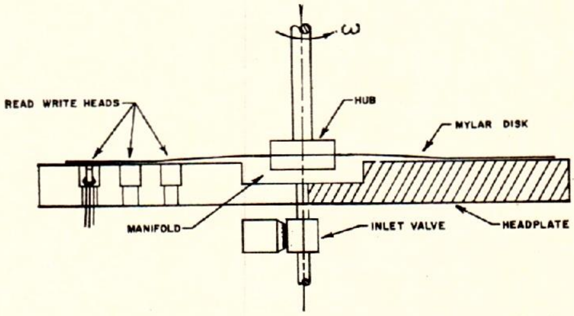

FIG. 2. PARTIAL SECTION of the flexible-disk device.

Head-disk velocity varies with radius. The flexible disk provides opportunity for compensation of the reduction in wavelength as the track radius gets smaller. If the disk-to-headplate separation is increased with radius, and the recording head gaps are similarly increased, it is possible to achieve a constant resolution for all tracks.

Changes in capacity of any Bernoulli disk rotating storage divice can be implemented in the field.

Five series of Bernoulli disks are available: BD-10 Series—25,000 to 50,000 bit capacity in 0.09 cubic feet; BD-40 Series—40,000 to 160,000 bit capacity in 0.21 cubic feet; BD-100 Series—100,000 to 400,000 bit capacity in 0.33 cubic feet, BD-200 Series—200,000 to 800,000 bit capacity in 0.66 cubic feet; BD-500 Series—500,000 to two-million bit capacity in 2.36 cubic feet. All measurements include electronics. Other specifications of a typical unit, the Model BD-103 (Fig. 1) are as follows:

Storage capacity is 100,000 bits (max): Bits per track—1000 to 3000: Bit rate is 30 to 400 kc. Typical track layout has 40 total tracks, including 32 data storage tracks, 3 spare tracks, 3 clock and timing tracks, 2 register tracks. Number of registers is 4; register length is 32 bits; register adjustment is ±3 bits; Disk speed is 1800 to 8000 rpm; induction or synchronous motors are available.

MR. WHITMER is a Government Contracts Administrator with Admiral Corporation where he is responsible for the corporate performance of contract requirements. He has been a Group Leader in the publication department of Admiral and a Technical Editor with Motorola, Inc.

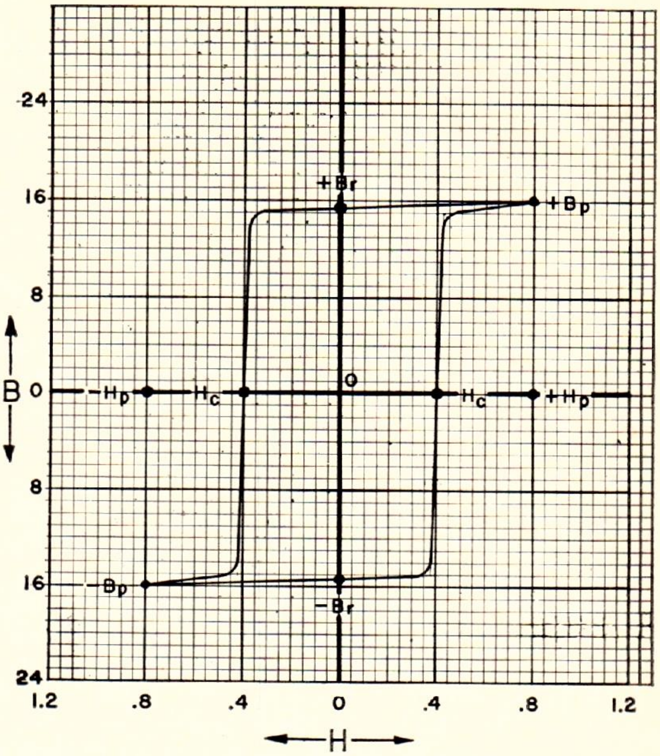

FIG. 1. CHARACTERISTICS of magnetic materials usually are described via the hysteresis loop showing flux (B) versus applied mmf field (H).

Magnetic Memory Cores

The magnet, with its bidirectional field, is admirably suited to binary-type information storage. Here is an introduction to magnetics, core properties, physical characteristics, test procedure and basic wiring.

MELVIN WHITMER

MAGNETISM is closely related to electricity, and magnetic properties can be defined by analogous electrical properties. The "current" in magnetics is *flux,* and the analogy is carried further by speaking of flux flow. The cgs unit of flux is the maxwell (Φ), defined as one line of flux. The "voltage" of magnetics is the *magnetomotive force* (H) usually expressed in ampere-turns. The cgs unit of this force is the gilbert (f), which is $5/2\pi$, or 0.794, ampere-turn. The "resistance" of magnetics is the reluctance offered by a vacuum to magnetic lines. This unit is called the *rel.* These basic magnetic units bear the same relationship as that of Ohm's law:

$$\Phi = f/rel;\ I = E/R$$

One maxwell (flux line) per square centimeter is one gauss (B); one gilbert (mmf) acting through one centimeter is one oersted (H); and the reciprocal of one rel per cubic centimeter is a permeability (μ) of one. These units bear the following relation to each other:

$$B = \mu H$$

A magnet can be "charged" to retain its field flux. In magnetism, the electron orbits of atoms are polarized so that the magnetic field created by each moving electron adds to form a total field within the material. (The electric counterpart is evidenced in superconductivity.)

The charge is placed in a magnet by induction from another magnet or from a d-c coil. When the applied force (field H) creates a field density in excess of the magnet capabilities, the magnet is said to be *saturated.* Removing the external force will result in

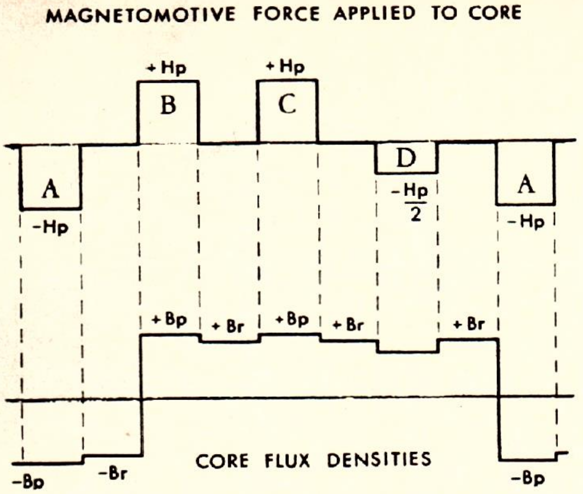

FIG. 2. TYPICAL SEQUENCE of test pulses, showing mmf applied to core at top, and resulting core flux densities (below).

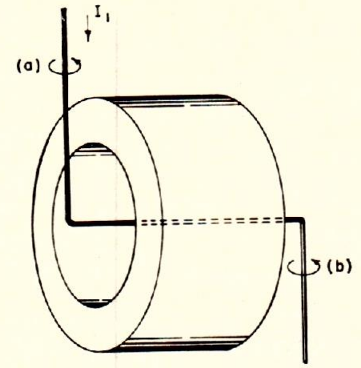

FIG. 3. CORE with wire threading its hole. Current shown will induce counterclockwise flux in core.

a reduction of flux density in the magnet to that of its remanence value (B_r).

The magnetic characteristics most desirable for a *permanent magnet* are (1) high permeability, (2) high coercivity (requires a large power input for magnetization), and (3) high remanence value.

Transformer cores require (1) high permeability, (2) low coercive force, and (3) low remanence.

Properties of a Memory Core

The magnetic-memory core must have slightly different properties from either the permanent magnet or transformer core. High permeability is still required, but moderate coercivity and high remanence field (B_r) are preferred. The memory core must not be influenced by stray magnetic fields; however, the polarity of the field must be reversible by pulse values readily available in a computer. The memory core must retain a large field because the purpose of the core is to generate an output pulse during a polarity reversal of the field, and the larger the field the larger the output.

The action in magnetic materials is easily shown with the hysteresis loop showing flux density (B) versus force field (H). The graph in Fig. 1 is of a core made of Deltamax. The vertical axis (flux density B) and the horizontal axis (designated H for magnetomotive force) are common to all magnetic loop graphs.

B_r shows the residual magnetism which remains after the application and removal of an external field. H_p represents the magnitude of the external magnetic field required to cause the flux to reverse direction. H_c represents the force required to demagnetize the material from a B_r position. These points identify the coercive force of the material. Note that the top and bottom portions of the loop do not run parallel to the H axis. This slope is the result of the permeability of air.

The magnetomotive force H_p creates a magnetic field in excess of the material capabilities. The flux B_p is the result of both the field in the material and the field around the material. When the applied force is removed, the field "slides" down the loop to the zero-H axis. This is the remanence value of the magnetic flux (B_r). The magnetomotive force $-H_p$ will set the flux at the $-B_r$ value, and the magnetomotive force $+H_p$ will set the flux at the B_r value.

An ideal hysteresis loop would be rectangular—that is, B_p would equal B_r. As a perfectly square hysteresis loop is impossible, the degree of squareness serves as a figure of merit.

The energy stored in a magnetic material is usable for developing an output pulse when the core flux changes polarity. The greater the flux density, the greater the output pulse. Designers would like to obtain a large output pulse with small applied force. This of course has definite limitations, but there are several alloys in addition to Deltamax which exhibit these qualities to a high degree.

Testing Cores

Magnetic materials are tested for operational characteristics by applying pulses to a circular core composed of the material under test.

Fig. 2 shows a typical sequence of test pulses applied to a core. Two monostable multivibrators provide pulses of H_p and ½ H_p magnitude. Pulse A sets the core in the $-B_r$ direction, which is generally called the ZERO direction. Pulse B reverses the flux, driving it to the $+B_p$ level, or ONE direction. Pulse C is of the same polarity as B and, therefore, can only drive the flux density to saturation ($+B_p$). A half-amplitude pulse (D) is applied to determine the amplitude of the induced voltage when the flux is momentarily reduced. Finally, the pulse A is applied to change flux direction and measure the output pulse.

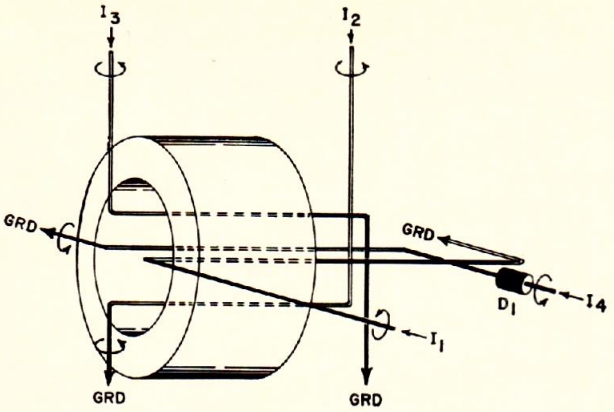

FIG. 4. CORE WIRING includes information wire (I_1), set wire (I_2), timing wire (I_3), and readout wire (I_4).

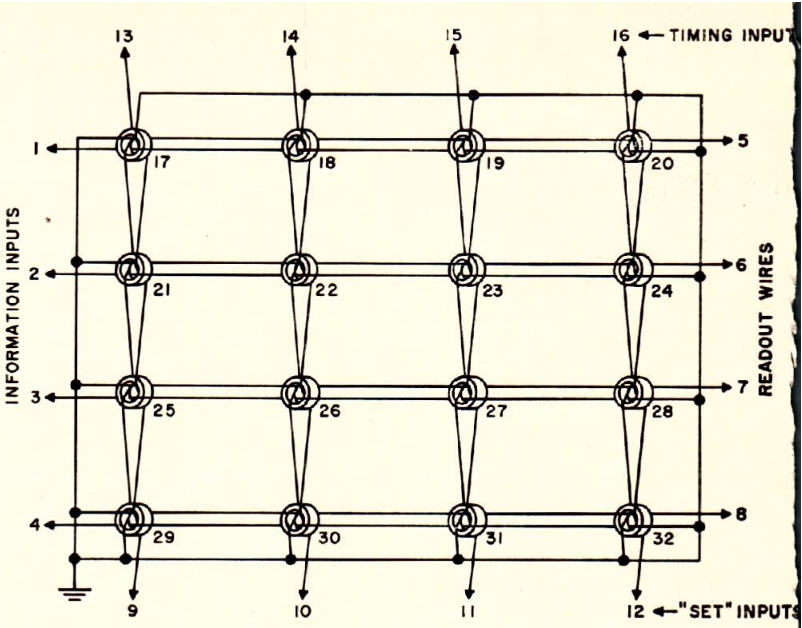

FIG. 5. TYPICAL CORE memory has hundreds of cores.

Physical Characteristics

A memory core usually comprises a doughnut-shaped plastic spool with a thin strip of metal wound on the spool (Fig. 3). Typical dimensions are:

Inside diameter	**0.25 "**
Outside diameter	**0.375"**
Width	**0.198"**

The cores are produced by rolling high-permeability metals into strips only 0.0005″ in thickness and 0.18″ in width, then winding the strips on the spools of plastic until a desired amount of metal is acquired.

With small cores it is not practical to use multiple-turn windings; the wires are simply threaded through the hole of the core. Current through the wire then produces a magnetic field which actuates the core. In Fig. 3 a current in the direction shown by I_1 will generate a field which will set the core flux in the counterclockwise direction.

Core Operation

A complete core wiring is shown in Fig. 4; wire size may vary from number 30 to 40. Arrows identify the current direction. The pulses are applied in this sequence:

I_2 sets the core in the ZERO state, which is the clockwise direction.

I_3 is the timing pulse which supplies $1/2$ of the power necessary to change flux direction.

I_1 is the information pulse to be stored.

I_4 is the readout wire which will receive an induced pulse each time the flux is reversed.

The diode D_1 is included to block the induced readout pulse when the flux is changed from ZERO to ONE; D_1 passes the induced pulse when the flux is changed from ONE to ZERO.

The current pulse I_2 sets the core in the ZERO state before applying the input pulse. I_1 and I_3 must occur simultaneously to change the flux to the ONE state. The core will remain in this state regardless of single $\frac{1}{2}$-H_p pulses (which may be applied to either the I_1 or I_3 wires). When readout is desired, a full H_p pulse applied to the I_2 wire will change the core from ONE to ZERO, and induce a pulse in the I_4 wire, which pulse will pass through the diode.

Memory Rack

As a complete memory may consist of several hundred cores, a method of affecting only one core at a time is required. Cores are selected by applying small pulses through two intersecting wires which both pass through only one core. A simplified version of a memory is shown in Fig. 5. Note that there are four *input wires* at left, four *timing wires* at top, four *readout wires* at right, and four *"set" wires* at the bottom. The input wires are pulsed by the information which is to be stored, while the timing wires are pulsed in sequence, from No. 13 to No. 16, by a pulse-generating "clock" multivibrator. The "set" wires are pulsed at the beginning and end of the storage time by the clock sequence. In Fig. 5, with a pulse applied to input number 3, and the timing pulse applied to number 14, the only core which receives both pulses will be number 26, which contains both number 3 and number 14 wires. All other cores will receive a $\frac{1}{2}$-amplitude pulse, or none at all. Cores 18, 22, 25, 27, 28, and 30 still can be used without disturbing number 26. When readout is desired, the timing pulses are switched to the "set" wires and amplified to full H_p size. As these pulses are applied in sequence, the cores which have been actuated to the one state will reverse, thereby generating a pulse in the readout wire. The "set" pulse also assures that each core is now in the ZERO state, ready to receive and store more information.

In a complete memory, each core is assigned a letter or number to correspond to the information to be stored. Several thousand cores usually are suspended by the wires threaded through them.

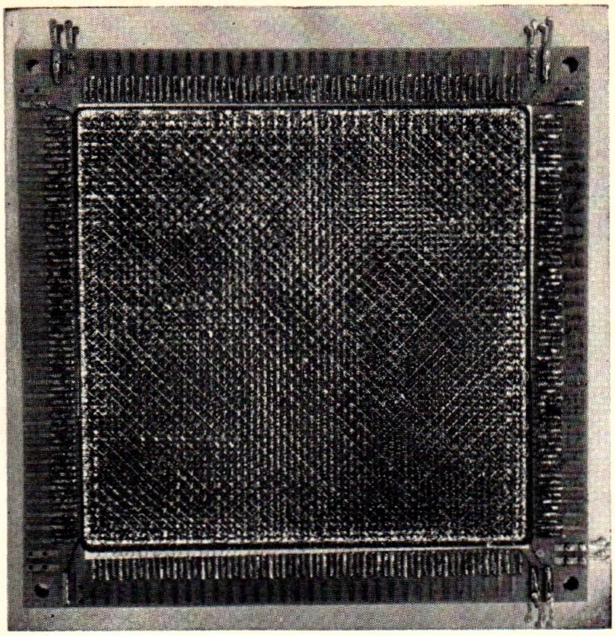

BIT PLANE of memory stack of 4096-word concident-current memory module in the Burroughs B5000 computing system. Bit plane has 4096 cores, with 64 x 64 contacts.

GEORGE H. BARNES
Burroughs Corporation
Burroughs Laboratories

FERRITE-CORE MEMORIES

FERRITE-CORE MEMORIES have filled the bulk storage need for random-access data since the early days of computing development. Although some recent storage techinques, such as magnetic thin films or superconducting memory elements, are potentially superior, ferrite-core memories will continue to dominate the random-access memory market until the other techniques achieve competitive economy and reliability.

Computers and data processing systems require means for providing stored data in quantity. Stored data should be available with random access such that any word of data can be called up by the parent data processing system with the same access time as any other word.

Ferrite-core memories are practical in data capacities beyond 10^5 bits, and have been constructed with capacities as great as 2.5 million bits. Packing densities of 100 bits per cubic inch are typical. Access time of 1 microsecond and read/write cycle times of 4 microseconds are typical of today's high-speed applications. Reliability, except at high temperatures, is generally considerably better than that of the parent system.

Core Properties

The toroidal ferrite core, a bistable permanent magnet, typically 0.050" in diameter, is the basic element in all of the memory systems discussed. These cores exhibit a B-H hysteresis loop (Fig. 1) in which the most significant characteristic, for storage and switching purposes, is a marked threshold for reaction (vertical sides) to a current applied through the aperture of the core.

The driving current produces a *magnetomotive force* (*mmf*). Without an applied mmf, the core rests in one of two opposite polarities of magnetization, or remanent states (A or B in Fig. 1). If an applied mmf is less than the *switching threshold* of the core, no flux change can take place and, after removal of the mmf, the core returns to the prior remanent state. If the mmf is of the appropriate polarity and exceeds the threshold by a sufficient amount, a change of flux results in the core. If an mmf is applied again, in the direction to which flux has already been driven, essentially no flux change takes place.

The minimum operating current necessary to fully switch a core is about 1.6 times the threshold current. Note that the threshold current is the *maximum disturb current* which can be applied to the core without significant flux change. (The term "disturb" is used for those effects which are not supposed to disturb the core, perhaps because in the early days they so often did.) In Fig. 1 the minimum operating (switching) current is designated I_M; the maximum disturb current (maximum current that can be tolerated without disturbing the core) is designated I_T.

Any conductor wound through the core (sense winding) receives a large induced voltage pulse whenever a flux *change* takes place; it receives almost no voltage pulse when insignificant flux change takes place.

Operation

Operation of a memory involves "writing" and "reading" the core. "Writing" into a core means driving it into either the ZERO or ONE state of flux saturation; "reading" means sensing the flux state, which is

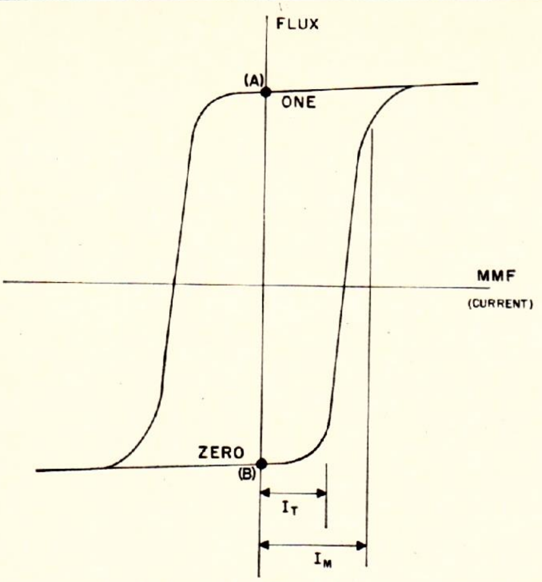

FIG. 1. HYSTERESIS LOOP of ferrite core.

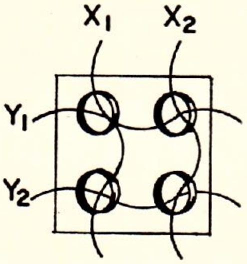

FIG. 2. A BIT PLANE, showing drive windings for the coincident-current (X-Y) type memory.

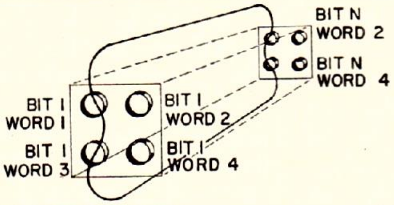

FIG. 3. MEMORY STACK comprises several bit planes.

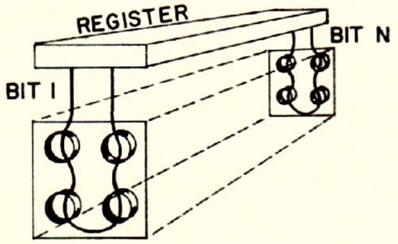

FIG. 4. SENSE WINDING threads all cores in a bit plane.

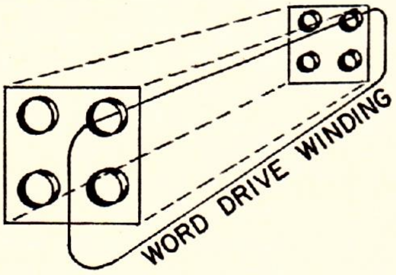

FIG. 5. PRINCIPLE of the linear-select (word-organized) memory.

accomplished by driving the core back to the ZERO state—if the core were in the ONE state, this results in a *change* in flux. A read/write cycle is established by an electronic "clock." During the first half of the cycle (the *read* portion), all cores storing a given word (one core per bit) are cleared to the ZERO state. On the following half of the cycle (the *write* portion), those cores which are to store ONEs are switched to the ONE state, while cores which are not to store ONEs are left in the ZERO state.

This switching, or driving, is accomplished by means of a set of "partial currents." In those cores which are not to switch, the sum of the partial currents is less than I_T; in those cores which are to switch, the partial currents add up to more than I_M. The read/write cycle is used both for reading from memory (in which case the information read out is restored by the following write) or for writing into memory (in which case the first half of the cycle serves only to clear the cores, without sensing, in preparation for write).

Terminology

The terms "extract" and "insert" have been proposed for extracting and inserting information from and into the cores. This leaves terms "read" and "write" to describe the function being performed in any particular memory cycle. Another term, "interrogation," often is used to describe the application of current (read) which specifically causes the stored word to be sensed; the distinction is a subtle one, and possibly creates more confusion than it eliminates. The problem of terminology in distinguishing between circuit function and logic function in digital magnetics is discussed in detail in Chapter 7 of Reference 2.

The term "register" means one word of memory. A register contains one word of information.

A coded address specifies which register (word) in the memory is to be read (interrogated), or written into. In practice, the registers often are referred to not as "registers," but as "words," which sometimes obscures the distinction between a memory and its contents.

The cores are arranged in planes stacked in a three-dimensional array. Each plane of cores (Fig. 2) in the array is called a *bit plane* because it contains one bit of any given word. Thus if a word has 10 bits, the first bit is stored in plane 1, the second bit in plane 2, etc. Thus 10 planes are used for storing this one 10-bit word. The complete set of bit planes is called a *stack* (Fig. 3).

Since only one word is read at a time, a single sense winding, wired through all of the cores in one bit plane, can be used to carry the output of that bit of any word addressed (Fig. 4). Only one bit will be interrogated at any one time in any one plane.

Because information is destroyed in each memory address at the time the address is read, provision must be made for restoring (or replacing) the information on the second (write) part of every cycle of operation. For temporary storage of the information read out (or of a new word to be stored), an external *memory register* is provided (Fig. 4). The outputs of the sense amplifiers feed the memory register in parallel. Connection to the parent data system is us-

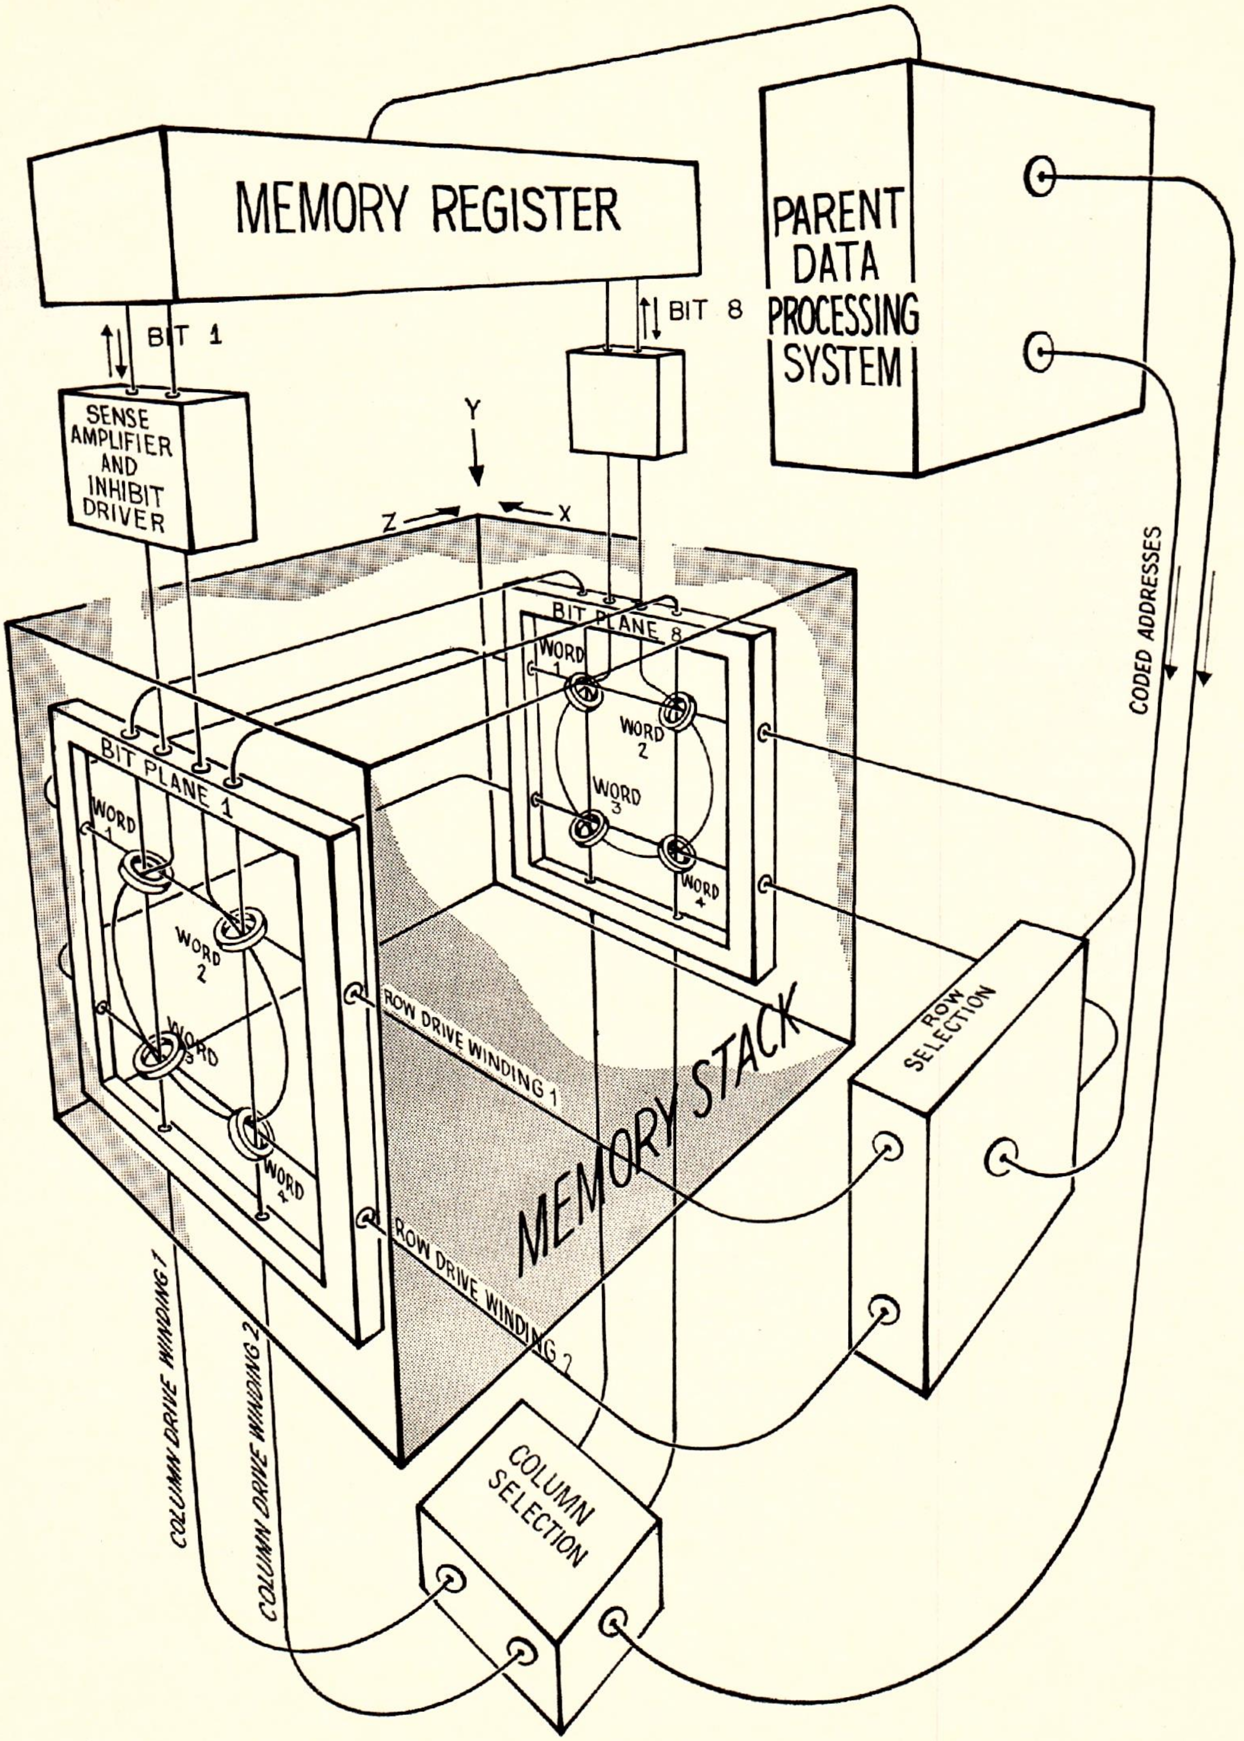

FIG. 6 COINCIDENT-CURRENT memory, showing all basic elements.

ually made to and from the memory register. Such a register is associated with any ferrite-core memory.

Coincident-Current Memory

The classical form of ferrite-core memory is the *coincident-current* memory (Fig. 2). This system achieves read and write by the coincident application of a set of half currents (one on an X line; one on a Y line), each of which is held below I_T. To read the information stored in the core at a given address, a half current is driven in a line in the X direction, another half current in a line in the Y direction.

Each X line (or Y line) passes through the same row (or column) of cores in all bit planes (Fig. 3). The half currents applied are individually less than I_T, but at the one core in each bit plane which is at the intersection of the X and Y lines, the two half currents combine to become more than I_M, and the core is driven toward the ZERO state. The sense line in each bit plane (Fig. 4) detects either a change of flux, (indicating that a ONE was stored in the selected core) or no flux change (indicating that a ZERO was stored).

Reverse currents in the X and Y lines then are applied to accomplish "write." These currents tend to switch all of the cores storing the word back to the ONE state, but in those bit positions where a ZERO is to be stored, an *inhibit current*, again a half current, but in the direction to oppose write, is imposed on the whole plane. Two of the three half currents (X, Y, inhibit) cancel each other, leaving a net half current (less than I_T), so that the core which is to store a ZERO does not switch. Calculations of current[2] tolerances show that for a ratio of I_M to I_T of 1.6, the tolerance on the various drive currents is $\pm 7.5\%$.

Fig. 6 shows a coincident-current memory.

Linear-Select Memory

More recently developed, even though conceptually simpler, is the *word-organized*, or *linear-select*, memory.[4] As shown in Fig. 5, there is a drive line associated with every word. Instead of selecting one core in every bit plane by the indirect process of X and Y coordinate selection (as in the coincident-current type) the cores to be selected are all strung on an individual wire and are selected directly by some external selection means, perhaps a nonlinear switch core or a diode matrix. Design tolerances in word-organized memories are much more generous than in coincident-current memories, but more components are used in the driving circuitry. For example, in a 100-word array, a 100-point switch would be required to select one word in the linear select memory; only two 10-point switches would be needed to select one word in the coincident-current memory.

Fig. 7 shows a linear-select memory.

Limits to Memory Performance

Ferrite-core storage and switching characteristics are notoriously temperature-sensitive. The critical parameter is the temperature coefficient of I_T. In typical cores, the threshold current I_T can vary by 0.4% per °C, and sometimes more. For this reason coincident-current ferrite core memories are exrtemely sensitive to the temperature environment of the memory stack; word-organized versions, on the other hand, have great environmental advantages because the current tolerances are more loose. For a ratio of I_T to I_M of 1.6, the allowable current tolerances are plus or minus 24%. Any increase in temperature decreases I_T. The temperature at which I_T finally drops to zero is called the *Curie point,* or *Curie temperature,* typically 250°C; above this the core is not magnetic.

An impediment to the design and operation of core memories in any environment is noise from the memory stack. The desired signal from the memory may not be large, perhaps from 30 to 100 millivolts. Noise pickup must be kept under close control.

One common source of noise is capacitive coupling between the various drive lines and the sense line. The major portion of the noise so produced can be eliminated by using a sense amplifier which rejects common-mode signal. The capacitive coupling to the sense line will impose approximately equal signals on both ends of the sense line, if the sense line is balanced with respect to ground. Thus, a sense amplifier which rejects common mode at the input will reject the bulk of the signal which is capacitively coupled from drive lines to sense lines. As memories become faster, and the bandwidths of the sense amplifiers correspondingly wider, effects like time-delay differences between the capacitive noise coupled into one terminal and the capacitive noise coupled into the other terminal will become important. The balance required for faster memories will impose additional restrictions on the sense line, or on the voltages permitted on the drive lines.

An additional source of noise is found in coincident-current memories. All of the cores on the selected X line and all of the cores on the selected Y line are half-selected, and will induce a slight amount of noise into the sense line. This noise is negligible in small bit planes; in large bit planes the cumulative effect becomes significant. To reduce this noise, the usual practice is to thread the sense line in each bit plane through the cores in such a manner that half of the cores produce positive noise on the sense line and half produce negative noise. The noise induced in each core, however, does depend slightly on the state of the core. The cancellation is imperfect, and a net noise component remains. This component is known as *delta noise.*[5] In the worst case, the cores that are wired to produce positive noise are all storing ONEs and the cores wired to produce negative noise are all storing ZEROs. This possibility limits the length of the sense line. If the memory is large, the bit planes may require a sense line longer than the maximum length tolerated. Then several sense lines must be employed in each bit plane.[6] Sizes permitting single sense lines are typically in the neighborhood of 4096 words (or addresses).

Coincident-Current vs Linear-Select

Limitations on memory capacity are, in part, a function of the choice between coincident-current memory and linear-select memory. In the coincident-current memory, because of the long X and Y lines, the number of bit planes (or bits per stored word) is limited by the voltage rating of the drivers. With

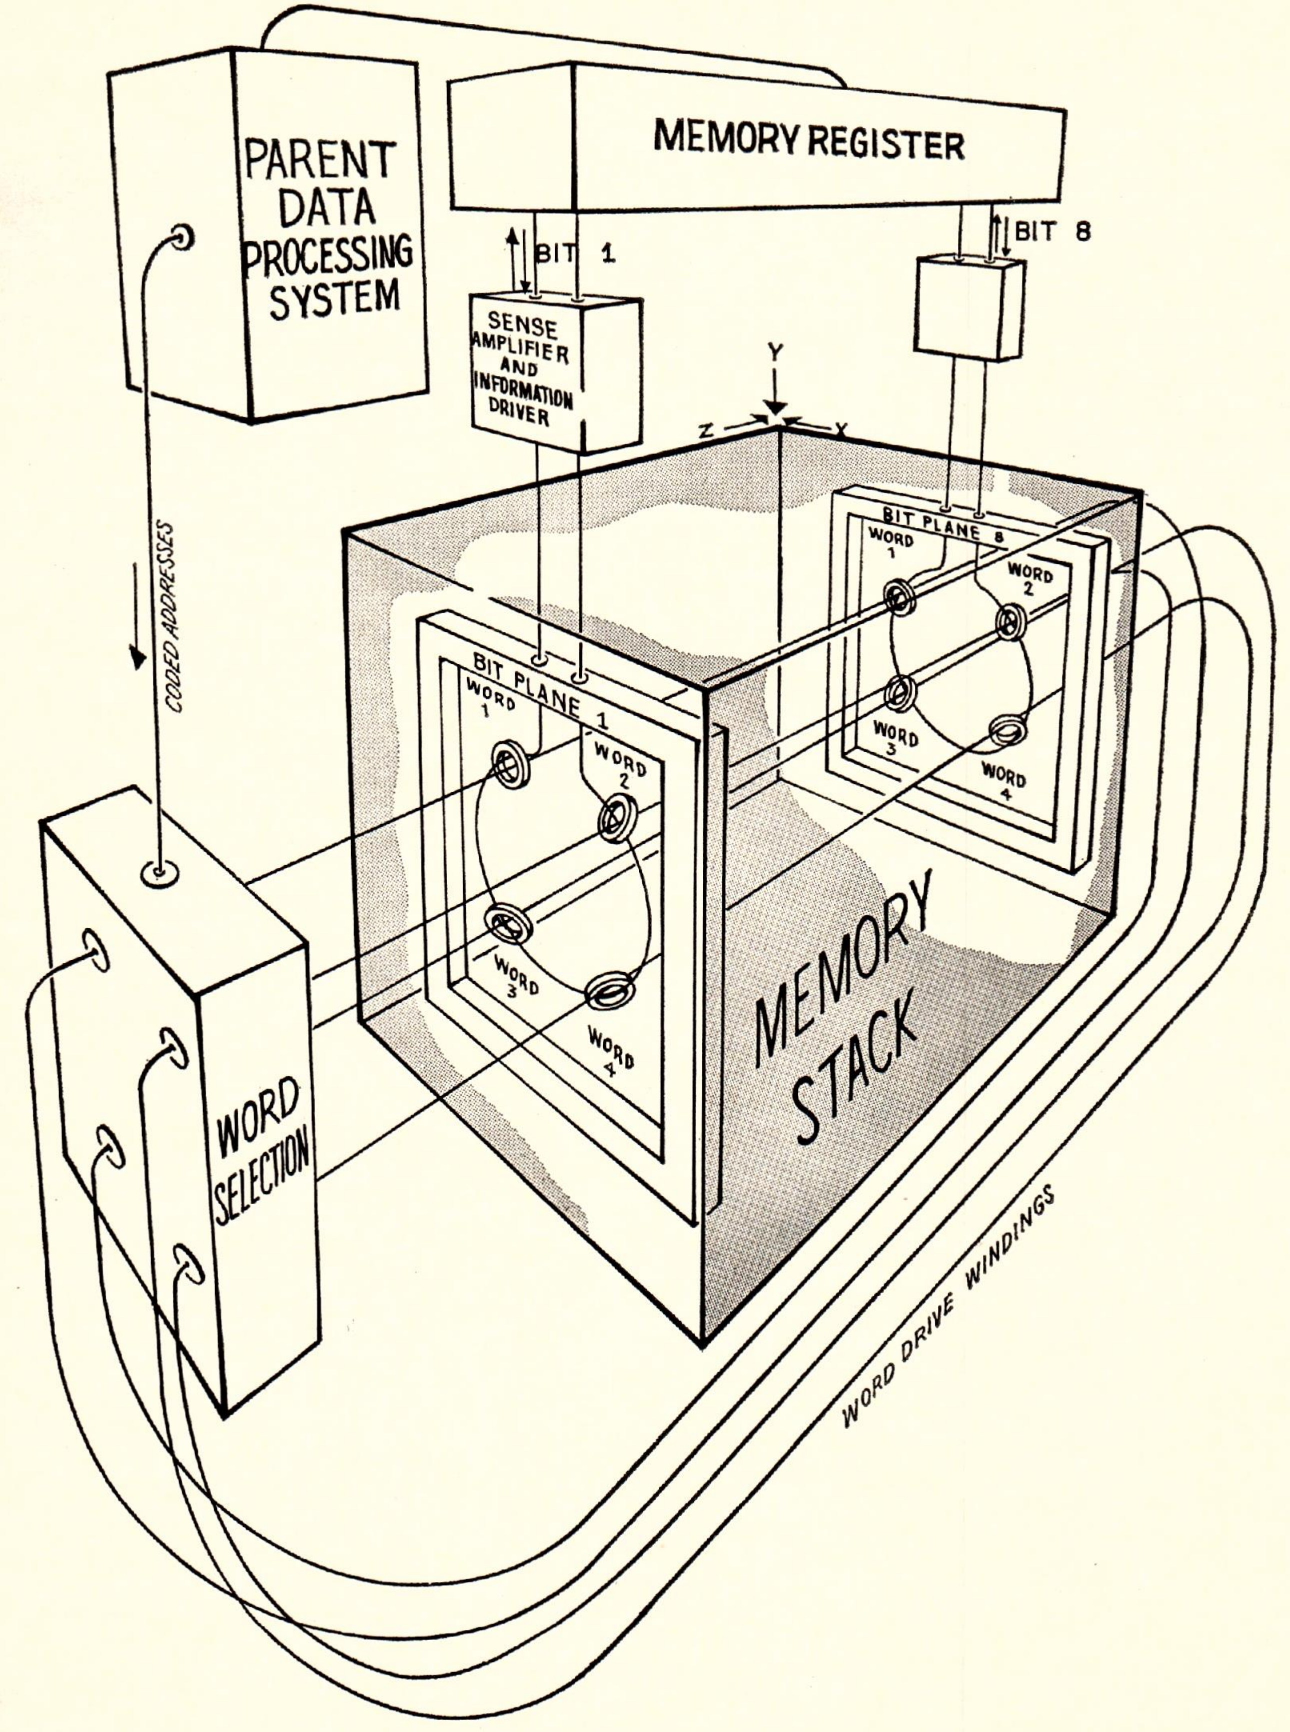

FIG. 7. LINEAR-SELECT memory, showing all basic elements.

reasonable transistor drivers, this limit is perhaps between 40 and 50 bits. Linear-select memories do not feel that limit as sharply,[7] nor are they limited by delta noise to some maximum number of cores per sense line. Very large coincident-current memories have, however, been built in defiance of the delta noise limit.[6]

Cost and complexity comparisons favor the coincident-current memory. At least to the extent that its sensitivity to tolerances can be ignored, the coincident-current memory also has a reliability advantage, in large-capacity applications, by virtue of the much lower component count in the driving and selection circuitry. The linear-select memory is far more capable of high-speed operation although a 2-microsecond coincident-current memory has been reported[8].

A further advantage of the linear-select memory is that the various methods of driving the cores are not nearly so restricted as those for the coincident-current memory, because of greater freedom in choosing current amplitudes. It is possible, for example, to overdrive the cores somewhat, to achieve faster switching time. More notably, the cores may be operated in a partial switching mode, such that a ZERO is represented by a core fully switched toward the ZERO state, while a ONE is represented by a partially switched core.[9] The advantages of this scheme are faster operation for the same drive current amplitude and less local heating of the individual cores which, in turn, permits more overdrive. Disadvantages are a slight loss of control in the amplitude of the ONE signals brought out of the memory, necessity for tighter tolerances on the read current, and a decrease in signal-to-noise ratio.

Future Directions

Limitations on the speed of present-day ferrite-core memory performance are the combined result of two effects. The first of these effects is the temperature sensitivity of the cores, the second the limited range of characteristics in available materials. A material to be used for more rapid switching in a coincident-current memory application is necessarily a material with a higher I_T; thus, the faster cores are also those which dissipate more power in the form of heat. The 2-microsecond memory referred to earlier is oil-cooled.[8] Unfortunately, a different choice of core material does not help. Some newer core types require less drive current for switching, but have higher temperature coefficients, whereas special cores types with low temperature-coefficients require more drive current. The sensitivity to self-heating and, hence, the fastest allowable switching time, has so far resisted improvement by research into materials.

An additional problem associated with attempts to increase the speed of ferrite-core memories is an attendant increase in the severity of noise problems. As speeds increase, the rise time of the current pulse in the drive line and, hence, the voltages induced on these drive lines, are continually higher. Yet the output signal from the core becomes smaller as the memory gets faster, because the cores themselves must be smaller to operate at higher speeds.

Present-day memories are limited in speed by the switching time characteristic of the cores. When faster memories are built, additional factors come into play, such as transmission time delay in the sense line and inhibit line (and also in the drive lines, in the case of the coincident-current memory). Where the transmission time delay effect is significant, high-density packaging of both the memory stack and its associated circuits becomes imperative.

Since the usable speed of the present-day computers is directly limited by the rate of access to memory units, these limitations on available speed are somewhat embarrassing. At this writing, there appear to be two solutions to the basic problem of self-heating, neither of which seems to promise more than an order of magnitude of improvement over present systems.

One of these solutions is the use of partial switching as mentioned above. The only other solution to the speed problem is the use of smaller and smaller cores. The amount of energy generated internally in a memory core is proportional to the volume of the core, if the change in I_T necessary for faster operation is neglected. The amount of heat which can be dissipated by a small device is proportional to its surface area. The maximum permissible speed should, therefore, be inversely proportional to the linear dimension of the memory core. Problems in core and bit plane fabrication soon become the limiting factor, although in large-capacity memories with such speeds, noise may present the final limitation.

The two approaches can, in effect, be combined. Drives that produce partial switching about some small hole in a large piece of ferrite also ensure that only a small amount of material immediately surrounding the hole actually switches. Not only is a small amount of heat generated because of the small size of the material switched, but the surrounding material helps to carry the heat away, as well as simplifying fabrication by permitting many holes to be formed in a single piece of ferrite.

Announcements of new inventions and new breakthroughs are anticipated. Meanwhile, ferrite-core systems still corner the lion's share of the high-speed memory market.

References

1. G. A. Oliver, "Computer Storage Techniques," *Instruments and Control Systems*, Vol. 33, No. 12, December 1960, pages 2070-2072.
2. A. J. Meyerhoff, ed., *Digital Applications of Magnetic Devices*, John Wiley & Sons, Inc., New York 1960.
3. R. K. Richards, *Digital Computer Components and Circuits*, D. Van Nostrand Company, Inc., New York 1957.
4. G. H. Barnes and R. P. Schneider, "High-Speed Word-Organized Memory Techniques," *Electronic Design*, Vol. 7, June 1959, pages 40-43.
5. J. R. Freeman, "Pulse Responses of Ferrite Memory Core," *IRE WESCON Convention Record*, 1954, pages 50-61.
6. W. N. Papian, "High Speed Computer Stores 2.5 Megabits," *Electronics*, Vol. 30, No. 10, October 1957, pages 162-167.
7. R. E. McMahon, "Linear Selection, Core-Memory Techniques Using Transistors," *Digest of Technical Papers, 1958 Solid State Circuits Conference*, Philadelphia, 1958.
8. C. A. Allen, G. D. Bruce, and E. D. Councill, "A 2.18 μsec Megabit Core Storage Unit," *IRE Transactions on Electronic Computers*, Vol. EC-10, No. 2, June 1961.
9. R. E. McMahon, "Impulse Switching of Ferrites," *Digest of Technical Papers, 1959 Solid State Circuits Conference*, Philadelphia, 1959, page 16.
10. M. H. Aronson, editor, *Data Storage Handbook*, Instruments Publishing Co., Pittsburgh, Pa., 1962.

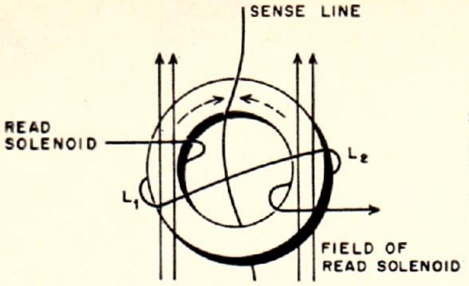

FIG. 1. READ solenoid produces opposing fields in Fluxlok core.

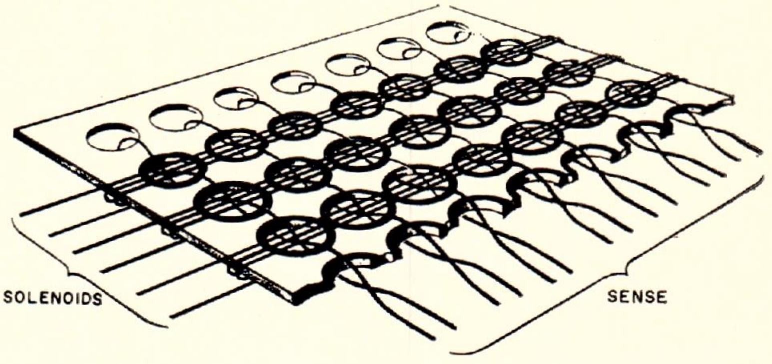

FIG. 3. SEVEN cores (for one 7-bit word) are mounted inside one read-solenoid. Three 7-bit words are shown. Cores fit into holes in phenolic base.

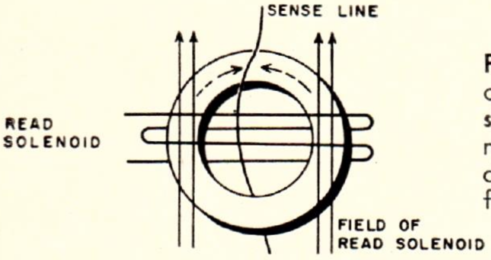

FIG. 2. CORE can be slid inside single solenoid that produces opposing fields.

Fluxlok High-Speed Core Memory

ROBERT M. TILLMAN, Burroughs Corporation Laboratories

The Fluxlok memory provides electrically alterable, random-access, high-speed, nondestructive read, using ferrite cores in a simply wired configuration. Noise or power-supply transients can not change the memory content. Output signals are insensitive to temperature variations from −65°C to +125°C.

FLUXLOK technique employs magnetic fields to (1) establish the state of magnetic cores (write) and (2) read the state of the core at rates to 10 Mc with no apparent core heating effects.

Nondestructive Read Principle

Fluxlok read operation is illustrated in Fig. 1. Two equal windings (L_1 and L_2) of a *read solenoid* are wound on a toroidal magnetic core in series opposition. A *sense line* threads the core. The state of the core can be considered 0 or 1, depending on the direction of the magnetized remanent state of the core (as will be explained in detail). If a current pulse is passed through the opposing read windings, a magnetizing force is produced, but the corresponding circular mmf's in the core cancel one another during the period of the read pulse (as indicated by the dotted arrows), and the remanent state of the core (existing prior to the application of current) remains the same when the read current is removed. Hence the state of the core is not destroyed when it is "interrogated" by a read pulse.

Fig. 2 shows a practical way of placing a single read winding around the outside of a magnetic core so that the mmf from a pulse of current in the winding produces opposing fields, the same effect as in Fig. 1.

The wiring of Fig. 2 is preferred to that of Fig. 1 because of the simple memory plane made possible (Fig. 3). The sense lines can be threaded

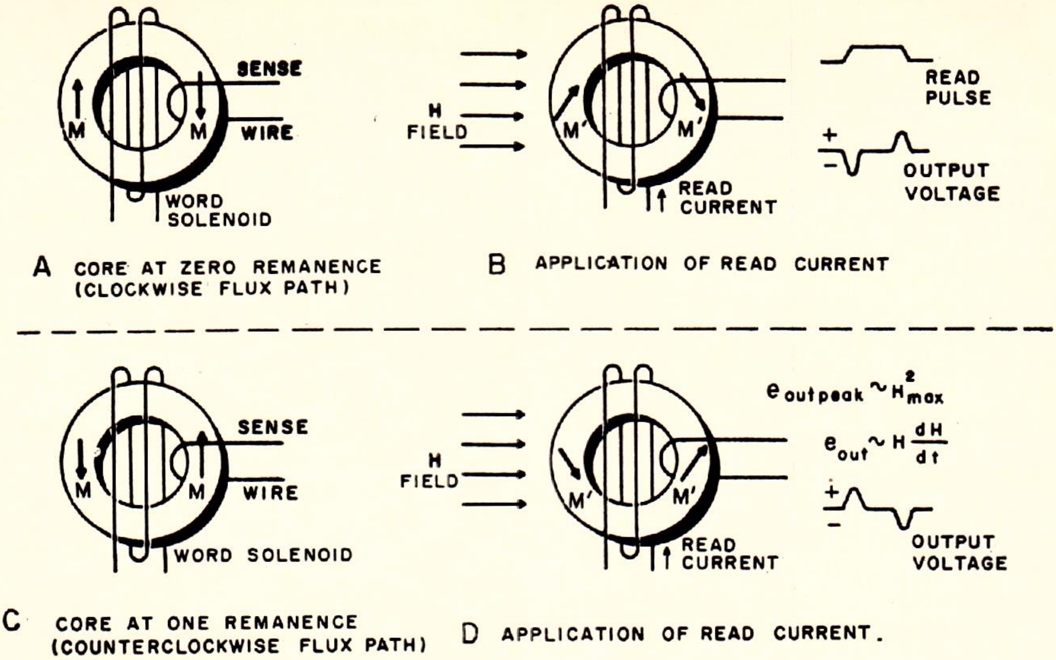

A CORE AT ZERO REMANENCE (CLOCKWISE FLUX PATH)

B APPLICATION OF READ CURRENT

C CORE AT ONE REMANENCE (COUNTERCLOCKWISE FLUX PATH)

D APPLICATION OF READ CURRENT.

FIG. 4. NONDESTRUCTIVE read operation. M is magnetic moment representing core saturation. Core has 0 state in A, and I state in C. A read current causes an H field. Vector sum of fluxes H and M shows that resultant moment M' is rotated. This causes induced signal (read signal) in the sense wire.

through each core twice, in a transposed fashion as shown, to produce larger output signals and to provide common-mode noise cancellation.

Read Operation

Nondestructive read results from current passed through the read solenoid. Consider a core initially magnetized in the clockwise direction (M in Fig. 4A). Application of the H field (due to current in the read solenoid) causes the magnetic moments M to rotate as shown. The moments now complete their flux paths through air. Creation of air paths by the H field produces a net decrease in the flux within the core, and a voltage of fixed polarity is induced in the linking sense wire.

When the H field is removed, the M vectors snap back to their original position, thereby inducing a second output voltage equal in amplitude but of opposite polarity. With the core originally magnetized in the counter-clockwise direction (Fig. 4C), the H field causes rotation of M in the direction shown in Fig. 4D, and the polarities of the two out voltages are reversed.

The actual sense line (Fig. 1) detects a signal from both sides of the core, rather than from one side (as shown in Fig. 4). The two signals are additive because the sense line crosses the top of one side of the core and the bottom of the other side. Note that one output signal appears on the sense line *during the rise time of the read current,* with a polarity directly corresponding to, and solely dependent on, the state of magnetization of the core (0 or 1). Another output signal equal in amplitude but opposite in polarity, is produced *during the fall time of the read pulse.*

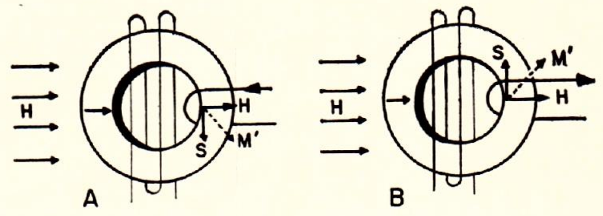

FIG. 5. SIMULTANEOUS currents in the read and sense wire cause moment M' (resultant of H and S) to be in directions shown, depending on polarity of current in sense winding. When currents are removed. M' will relax into the the 0 or I state.

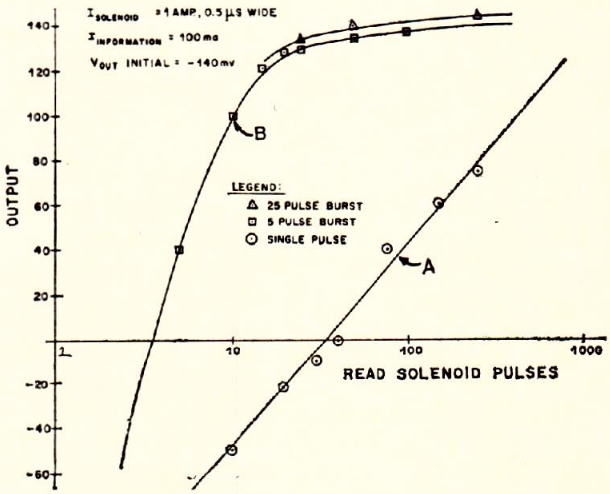

FIG. 6. READ-SIGNAL output (sense signal), in millivolts, versus number of read-solenoid pulses used for write-in. For example, A represent 90 single pulses. B represents ten 5-pulse bursts.

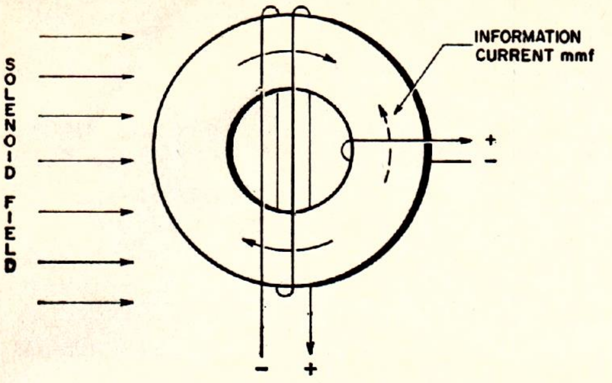

FIG. 7. WRITE action for changing core from 0 (CW) to 1 (CCW).

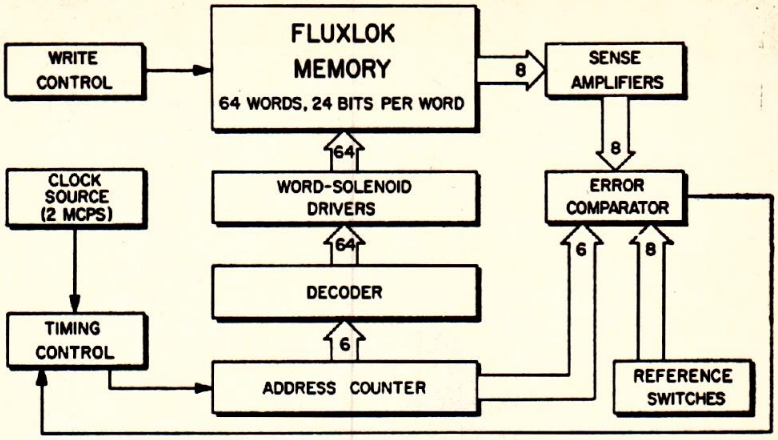

FIG. 8. FLUXLOK memory in test circuit.

Write Operation

Write-in can be performed in conventional coincident-current fashion, or by means of a slower but much simpler orthogonal crossed-field technique. In the orthogonal write operation, the same drivers can be used for read or write because these two operations are identical except for the absence or presence of information current in the sense lines.

Fluxlok orthogonal *write* changes the magnetic state of a core by simultaneous use of currents in the read and sense wires, as shown in Fig. 5. The current in the read winding causes field H; current in the sense winding causes field S. The direction of the resultant M′ depends on the polarity of the current in the sense wire. If M′ is as shown in Fig. 5A, it will relax into the S direction (core in 0 state) when the currents are removed. If M′ is in the direction shown in Fig. 5B, the core will relax into the 1 state when the currents are removed.

Where the bits are to remain in their prior 1 or 0 state, the sense current polarity merely drives the core further into its existing remanent state. Orthogonal write, then, is an overriding action in which it is not necessary to first clear all bits of a word to a reference state, as in destructive read-out core memories.

This Fluxlok write technique might be expected to result in an extremely fast switching of the core from one given state to the opposite state or remanence. However, investigation revealed that a core would *not* change its information state with the application of a single read-solenoid pulse, but that it required a series of such read-solenoid pulses (about 25) to fully change the information state of the core. Increasing the pulse width from a nominal 0.1 to several microseconds had no beneficial effect. Similarly, it was determined that the rise time of the solenoid pulse was not a governing factor.

Fig. 6 shows the effect of varying the number of read-solenoid pulses for a single write operation, starting, in each case, from a saturated reference state of opposite remanence. Additional tests showed that an increase in the read-solenoid pulse amplitude resulted in an expected decrease in the number of pulses required for full switching of the core.

One theory as to the mechanism by which this slow, ratchet-like write operation takes place is concerned with the differing effects of the pulsed solenoid cross-field on the core magnetic moments in the various segments of the core. In Fig. 7, the cross field opposes the direction of the remanence moments in the lower part of the core, while at the same time this cross field is in a direction to aid the remanence moments in the upper part of the core.

Simply stated, during the presence of a solenoid pulse the flux density in the lower part of the core is decreased by the combined solenoid-field and information-current core mmf, and the flux density in the upper part of the core is changed very little because flux changes in this part of the core are opposed by the mmf produced by the information current.

Upon termination of the pulse field, the core assumes a minimum energy state wherein the two flux densities again become equal. In this case the information current decreases the flux in the upper part and opposes an increase in the flux of the lower part. The result is that the net core flux has been reduced *slightly* towards a demagnetized state. Repeated application of solenoid pulses gradually "walks" the core state through the demagnetized state and eventually to the desired opposite state of saturation.

Test System

In the circuit chosen to test the memory (Fig. 8), the contents of the memory are compared (by means of sequential addressing) with the contents of a binary (address) counter performing the addressing of the memory. In addition, provision was made for setting up a comparison word on a series of reference switches so that any pattern could be entered into any word position and compared during the test cycling of the memory. These patterns

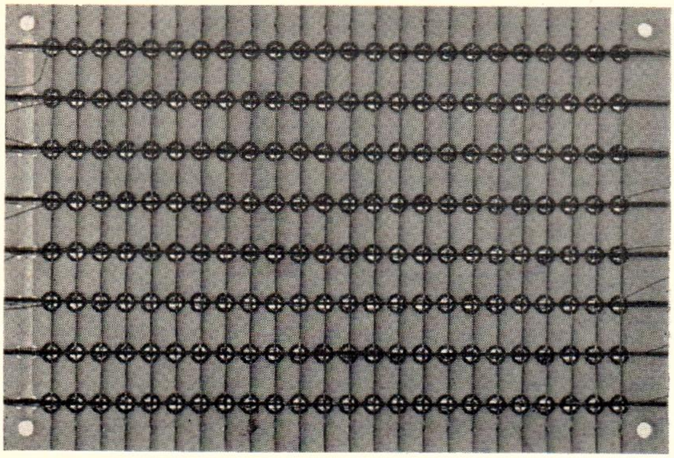

FIG. 9. FLUXLOK memory plane, enlarged.

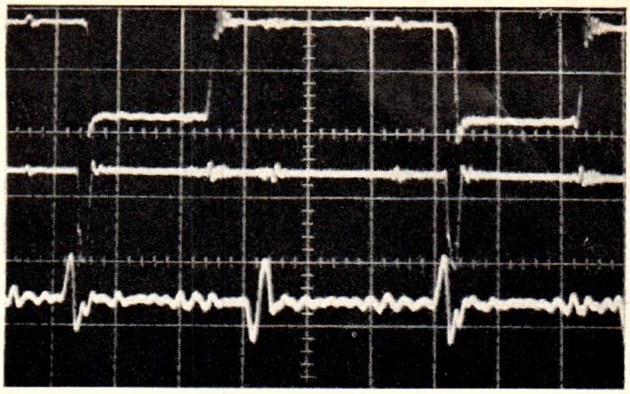

FIG. 10. RELATIVE TIMING of output register (top), sense amplifier (center), and core output (bottom) showing binary 101. Scale of two top graphs is 2 volts/cm; bottom graph scale is 50 mv/cm. Frequency is 1 Mc.

are entered by the orthogonal (crossed field) method of writing.

The memory consisted of eight planes of eight 24-bit words each. Fig. 9 shows one of the memory planes. The storage elements were 50/80-mil ferrite cores. The unit was assembled to simulate a larger memory stack to facilitate studies of noise and crosstalk that might be encountered in a full-sized memory. Additional wires, added in parallel with the word windings for coincident-current write, can be seen in Fig. 9.

A nominal read rate of 2 megacycles (required by the supporting project) was chosen for the test. All circuitry was transistorized except that vacuum-tube drivers were designed to furnish the high-current, fast-rise-time pulses at this frequency. (Reliable solid-state drivers to furnish 1-ampere pulses with 50-nanosecond rise times are limited to about a 500-kc repetition rate with present-day semiconductor devices.) Sense-line signals were fairly clean and uniform, and the sense-amplifier design has proved adequate for the fast-rise-time bipolar output signals.

Fig. 10 is an oscilloscope display of the waveforms taken while reading out binary 101.

Conclusions

The Fluxlok memory technique appears to be promising for those computer applications where its electrical alterability and high-speed nondestructive-read features can be used to advantage. The orthogonal (crossed-field) write feature requires the use of only one diode per word in linear-select memory systems, as opposed to two or even three diodes per word in more conventional designs.

The transition from a laboratory Fluxlok test system to a larger memory system has introduced problems. Proper electrical operation depends on the maintenance of a good orthogonal relationship between the word-windings and the sense wires during packaging. As in any high-speed memory development, the second- and third-order electrical effects become prominent when added together in a large system. A 2048-word 19-bit memory system is under construction. Transistorized drivers and associated memory circuitry have been developed for the operating speeds desired. Circuit capacity and inductance, along with the transistor and diode voltage drops in decoding and selection circuitry created design difficulties. However, the work performed to date indicates that suitable components are now available for this effort, and that the transition from a laboratory test vehicle to a useful, reproducible memory system will be accomplished successfully.

Bibliography

Buck, D. A. and Frank, W. I., "Nondestructive Sensing ing of Magnetic Cores," *AIEE, Communications and Electronics*, January 1954, pages 822-30.

Gianola, U. F., "Nondestructive Memory Employing a Domain-Oriented Steel Wire," *Journal of Applied Physics*, May 1958, pages 849-53.

Newhouse, V. L., "The Utilization of Domain Wall Viscosity in Data Handling Devices," *IRE Proceedings*, November 1957, pages 1484-92.

Papoulis, A., "Nondestructive Read-out of Magnetic Cores," *IRE Proceedings*, August 1954, pages 1283-88.

Rajchman, J. A. and Lo, A. W., "The Transfluxor," *IRE Proceedings*, March 1956, pages 321-32.

Shevel, W. L., Jr., "Millimicrosecond Switching Properties of Ferrite Computer Elements," *Journal of Applied Physics*, Supplement to Vol. 30, April 1959, pages 475-485.

Thorensen, R. and Arsenault, W. R., "A Nondestructive Read for Magnetic Cores," *Western Joint Computer Conference*, March 1955, pages 111-16.

Wanlass, C. L. Wanlass, S. D., "Biax High Speed Computer Element," 1959 *WESCON*, August 1959, San Francisco, California.

Tillman, R. M., "Fluxlok—A Nondestructive Memory Technique Using Standard Memory Cores," *IRE Transactions on Electronic Computers*, Vol. EC-9, Sept. 1960, pages 323-328.

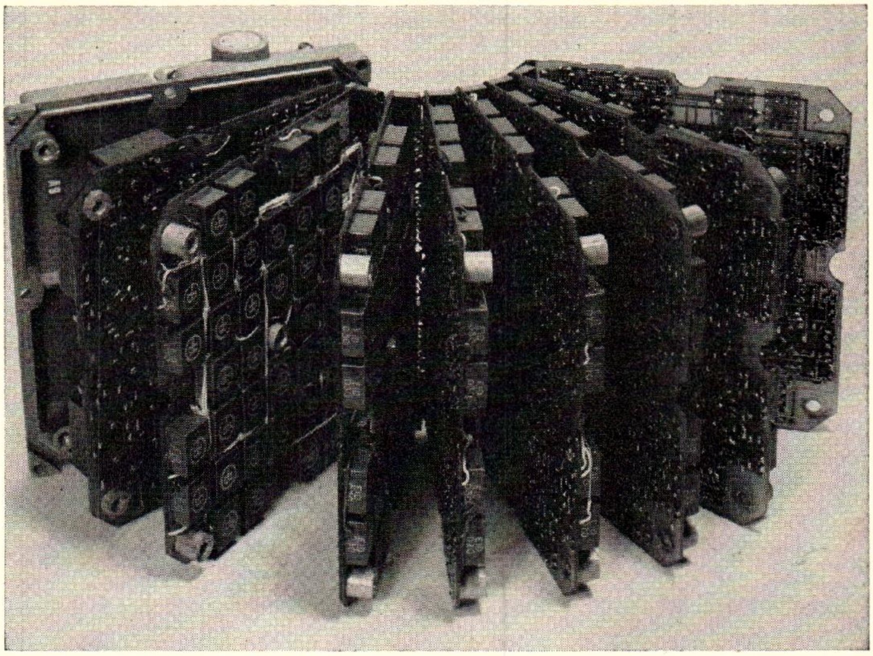

MAGNETIC SHIFT REGISTER

GEORGE E. LUND
EDWARD F. MYERS

Burroughs Corporation
Burroughs Laboratories

Magnetic shift registers, one of the earliest applications of bistable magnetic cores in digital systems, are used widely for information storage (for either fixed or unspecified intervals), and for alteration of information format (as from serial to parallel bit-transfer). Because of low power consumption, extreme reliability, and relative immunity to severe environmental effects, the magnetic shift register has been developed and applied in a variety of configurations. Logic also can be implemented by magnetic-core circuits.

BISTABLE square-loop magnetic cores were first employed in digital applications during the late 1940's as memory devices. Their use soon spread to include all data manipulations required in data-processing systems. One of the first applications was the *magnetic-core shift register* (MSR). Because of extreme reliability, resistance to harsh environments, and low power-consumption at low clock rates, the MSR is used today in many configurations.

Basic Transfer Loop

Several families of digital magnetic circuits have evolved, based on different approaches to the use of the core in performing logic, different methods of timing and wiring, etc. One of the earliest and most broadly applied of these families is the Bimag*, a two-core/bit, parallel magnetic pulse-amplifier technique. The major types of shift registers have been implemented in Bimag circuits.[1, 3] Fig. 1 is the Bimag logical representation for the so-called *single-diode loop,* the basic circuit element of the technique. The symbols are as follows:

1. A *magnetic core* is represented by a circle. *Transfer loops* are represented by straight lines connecting the circles.
2. An *input winding* is represented by an arrow pointing toward and touching the circle. The digit in the circle at the end of the input line denotes the state to which the input brings the magnetic core (1 or 0, saturated positive or negative).
3. An *output winding* is represented by a line without an arrowhead at the circle. The digit in the circle at the beginning of the output lines denotes the remanent magnetic state of the core necessary at the time of the output to produce this output.
4. *Clock driver pulses* are represented by the letter "t", and relative time indication by a subscript; specific clock pulses are thus indicated by symbols such as t_1 and t_2.

In Fig. 1, a pulse input at t_1 sets core A to the 1 state; a pulse on the t_2 input resets the core to 0 and sets core B to 1.

Circuit Operation

A two-core/bit shift register using single-diode loops is shown in Fig. 2. One driver is connected to the A cores and the other to the B cores. The B cores are said to be *clear,* or *empty,* when information is stored in the A cores, and are thus receptive to a transfer from the A cores. This transfer occurs in response to current I_0 in the A drive line. Transfer current I flows when a core storing a 1 is driven to the 0 (in this case, clear) state, which then causes the succeeding core to switch to the 1 state.

Information is shifted in this manner through the register by alternate applications of the two drivers. As an aid to following the operation of core circuits, it is helpful to remember that current entering a dot terminal of a core tends to cause current to exit from the dot terminals of the other windings when switching occurs.

Simple Shift Registers

A series of Bimag loops can be connected together as shown in Fig. 3 to form a simple magnetic shift register. If a 1 is written in Core A, and t_1 and t_2 pulses are applied alternately, the 1 is shifted down the register toward core F. After the first pair of pulses (t_1 and t_2), the 1 has shifted from core A through core B to core C. After the next pair of pulses, the 1 appears in core E.

Additional 1's (or 0's) can be inserted into core A between t_1 pulses while information is being shifted. Any configuration of 1's and 0's can be stored in *alternate cores* of the register and all information shifted simultaneously. 1 and 0 inputs are represented (in Fig. 2) by pulses and absences of pulses, respectively, but the reverse would be equally suitable.

The register in Fig. 3 is referred to as a *serial-serial shift register* because information is inserted bit by bit by bit into core A, and appears at core F as an output, again bit by bit. In a *parallel-parallel register,* information is inserted in parallel and extracted in parallel; such registers are often used for simple delay or static storage.

A one-core/bit shift register using gated transfer loops is shown in Fig. 4. Only one driver is used, although a gating pulse must be provided immediately following each drive pulse. Initially, information is stored in each of the cores, the capacitors are discharged, and the gating device is off. Upon application of the drive pulse, cores storing 1's are switched to 0 (in this case, the clear state), thereby producing a current which charges the associated capacitors. The capacitors temporarily store the information while the cores are being cleared. Immediately following the drive pulse, the gating device is turned on, whereupon the capacitors which have been charged are discharged through the input winding of the succeeding cores. Information is thereby shifted through the register by alternate applications of the drive and gate pulses.

Another one-core/bit circuit is shown in Fig. 5. This circuit is more economical with respect to parts, but is more difficult to design and has narrower operating margins. The circuit operates in much the same way as the previous circuit, except that the information is not gated to the input of the succeeding cores, but rather is delayed in transit by an inductive-capacitive delay network.

A one-core/bit shift register is shown in Fig. 6. The delays are necessary because a core cannot be read out and written into at the same time. The AND gates shown pass the input pulse to the next core only when a shift pulse is present. Designs of this type can be achieved using only passive elements, but are more efficient where gated transfer loops are used.

Reversible Shift Register

Fig. 7 shows a *reversible* one-core/bit shift register. Information is shifted right or left, depending on which shift gate is activated. The register shown in Fig. 3 could be made reversible by the addition of conditional drivers.

Serial-Parallel Register

A *serial-parallel* register makes use of *conditional* transfer. A conditional transfer (one permitted only at particular clock times) is accomplished in a single-diode loop by an overriding technique, as shown in Fig. 8. The dot inside the ZERO input at core C indicates an input of sufficient magnitude to override any other simultaneous input. Transfer of data from core A to core C at t_1, via core B, is prevented by the t_1 pulse at core C, which holds core C in the 0 state, overriding a possible 1 input. The crescent-shaped line between the two 0's in core B implies conditional transfer at time t_2. In this example, the conditional transfer produces at C the negation of the input from A.

---

*Trademark of the Burroughs Corporation.

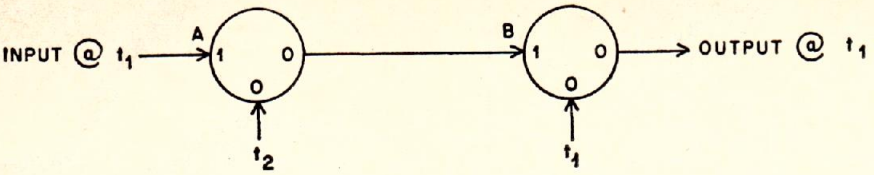

FIG. 1. SINGLE-DIODE transfer-loop. The core is represented by a circle; windings are represented by straight lines with arrows.

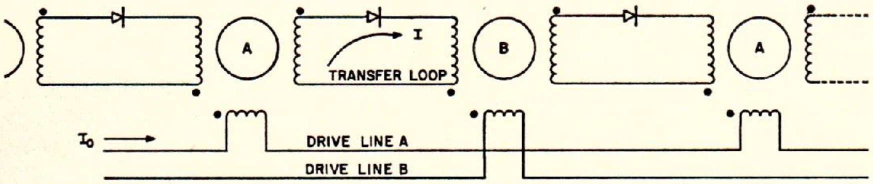

FIG. 2. TWO-CORE/bit MSR.

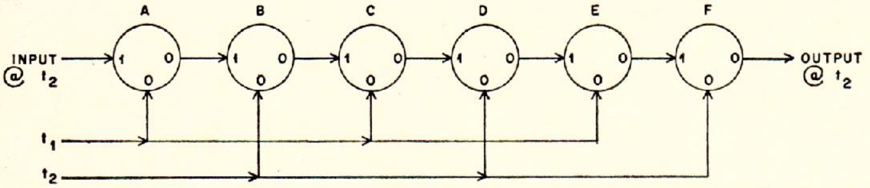

FIG. 3. SIMPLE magnetic-shift register. This is a serial-serial register because both input and output are in serial form.

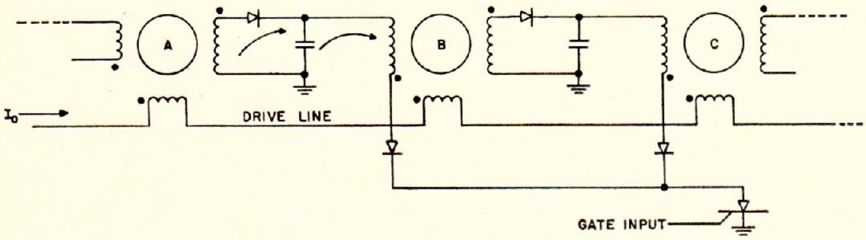

FIG. 4. ONE-CORE/bit MSR.

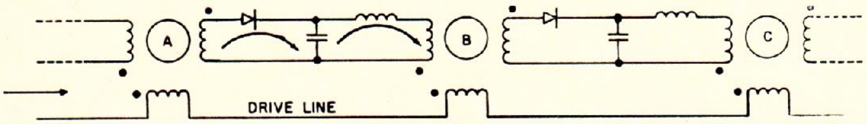

FIG. 5. DELAYED One-Core/bit MSR.

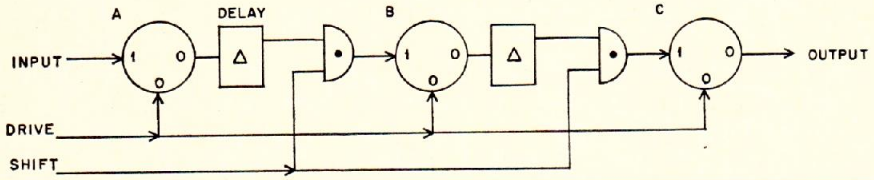

FIG. 6. ONE-CORE/bit MSR must have delay elements between cores because a core cannot be read out and written into at the same time.

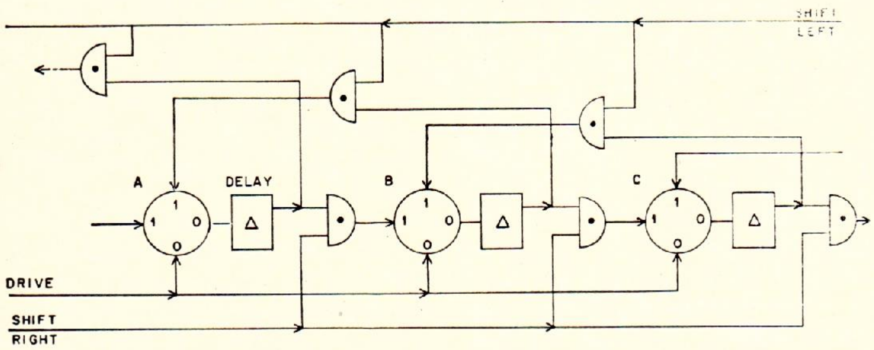

FIG. 7. REVERSIBLE MSR core can shift left or right.

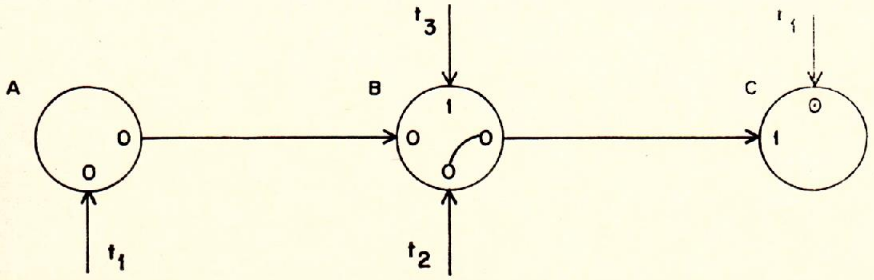

FIG. 8. CONDITIONAL transfer.

The conditional transfer also can be effected by another transfer loop known as the *split-winding* loop. This loop utilizes additional components and a different winding configuration to reduce the total number of cores and windings.

In the serial-parallel register of Fig. 9, information is inserted serially and read out in parallel. Driver D_3 is applied only once after the information has been shifted down the register and appears in cores B, D, and F. An override is applied to the output cores (not shown) at t_1 to prevent parallel noise transfer during serial shifting. In this case the conditional transfer does not involve a negation. A *parallel-serial* register can be made simply by inserting information in parallel to the respective cores and then shifting as in Fig. 3.

Ping-Pong Storage Register

The registers described thus far provide only for delay storage. Fig. 10 illustrates a method, called a *ping-pong,* of providing static storage. The ping-pong may be considered a *one-stage end-around* shift register. As shown, core A is set by inserting a 1 at t_1. An output is produced at t_2 as the 1 is transferred to core B. The information is returned to A at t_1. A conditional transfer is used to allow resetting of the ping-pong at t_2.

Multiple Drivers

Multiple drivers can be used to reduce the core count in long registers as shown in Fig. 11. In this register, information shall be assumed stored initially in cores B, C, D, and E. Cores A and F are cleared. At t_1, the bit stored in E is transferred to F; at t_2, the bit stored in D is transferred to E, and so forth, until one timing cycle is complete. At such time, the information is stored in C, D, E, and F; B is clear; and A stores the new input bit.

The read timing is in reverse order with respect to information transfer. In general, n cores per $(n-1)$ bits (where n equals the number of timing phases) are required for shift registers of this type.

Complex Shift Registers

For special applications in the slow to moderate speed range, there are a variety of more complex MSR arrays which can be used to advantage. One such array (Fig. 12) is a linear-select memory with random access, and includes a memory register. Such a memory is useful where a small-capacity storage is required, particularly in association with magnetic-logic circuitry. The memory core winding configuration is shown at the bottom. The circuits are similar to the single-diode loop, as used in the two-core/bit configurations shown in Figs. 1, 2 and 3.

Information is shifted serially through the memory register and may be stored in the random-access memory in parallel. The selected address input or word line is energized as the input word is being shifted by the A drive line. (The bit position of the input word must, of course, correspond to the bit positions of the memory array.) Information is transferred from the memory to the memory register by energizing the appropriate address devices at the B shift time. The extracted information may be immediately rewritten into the memory, thereby retaining the information for future use, to complete a conventional read/write cycle.

Applications

Chains of magnetic-core circuits (*shift registers*) are used to store information, either for a fixed interval (*delay storage*) or for an unspecified interval (*static storage*) to be terminated upon demand. Shift registers also can be used to alter the format of the information being stored. For example, a shift register may accept the individual bits of a data word one at a time (serially), and transfer the complete word at a single clock time (in parallel).

Core circuits must have the ability to amplify, if they are to transfer information from core to core in an iterative system. Amplification is accomplished in typical MSR circuits by using the timing and power capabilities of a pulse driver. The cores switch relatively large amounts of power from the driver to succeeding cores, while requiring only a small amount of power for an input (at some time prior to the output phase). Power amplification is therefore accomplished, since a relatively low-power input signal produces a relatively high-power output signal.

The pulse power required to drive MSR circuits is directly proportional to the clock frequency. At very low clock rates (below about 10 kc), power consumption is therefore extremely low. Standby power is almost insignificant. Above 10 kc, the decision to use MSR circuits is increasingly dependent upon the needs of the particular application. The reliability, shelf life, and resistance to environment of MSR circuits far exceed that of alternative techniques, and all three advantages become more pronounced in severe environments. Bimag circuits have been successfully employed in high-shock ordnance environments,[2] resisting acceleration shocks as great as 20,000 g's, temperature variations of 230°F, rotational velocities of 15,000 rpm, free-fall water entry from as high as 30,000 feet, and substantial nuclear radiation. Catastrophic failures are extremely rare. For these reasons, the MSR is uniquely suitable to aerospace applications such as long-term navigation, where high clock rates are not required but environmental effects are severe, minimum power is available, and operational use is measured in years. The practical speed limit for the MSR is roughly 100 kc; rates of 500 kc are possible, although power requirements become extreme.

Cores wound with metallic tape, such as Molybdenum Permalloy, are superior to ferrites in MSR applications (except at very high speeds), because of lower power requirements and greatly superior temperature tolerances. Wiring of the individual core-diode elements in an MSR is a fairly complicated process, but packaging is otherwise similar to that of other solid-state circuits, and component counts are generally lower than for corresponding semiconductor circuitry.

Driving and input/output circuits are easily designed. For example, because output power is high, an excellent direct driving input to silicon-controlled switches can be achieved. Because the cores are low-impedance devices, circuits are relatively unaffected

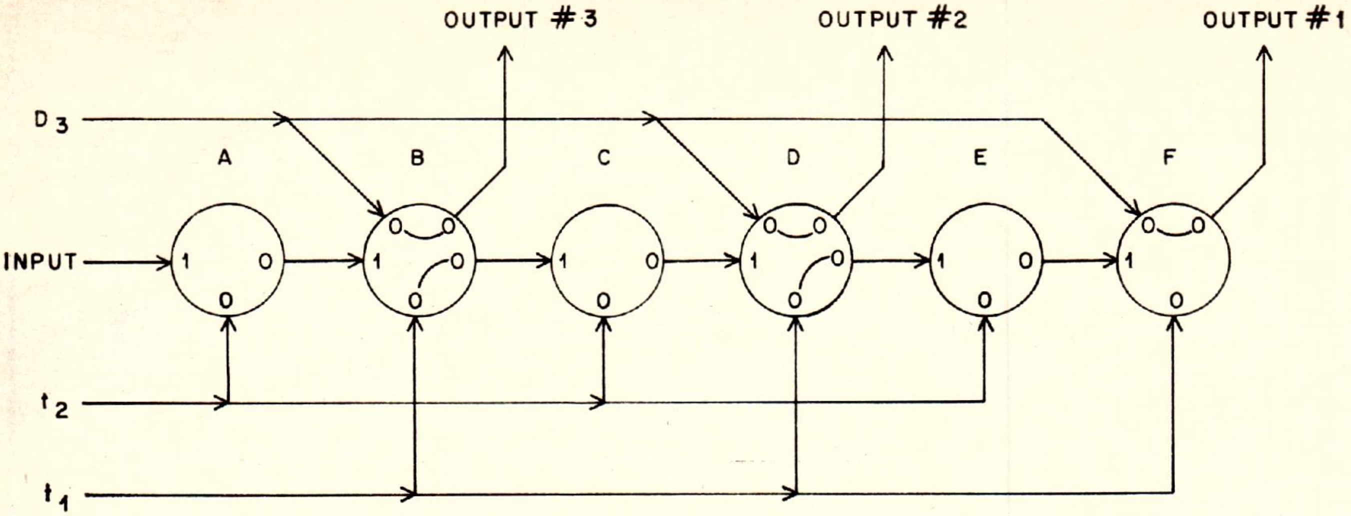

FIG. 9. SERIAL-PARALLEL MSR accepts inputs serially, reads out in parallel.

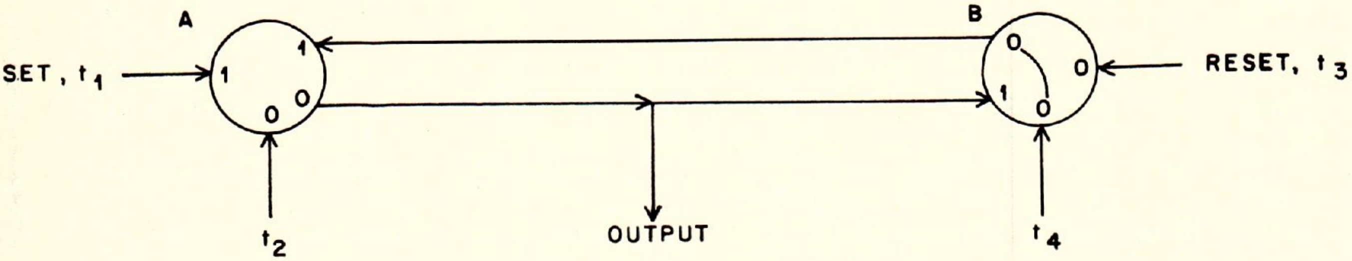

FIG. 10. PING-PONG storage register provides static storage.

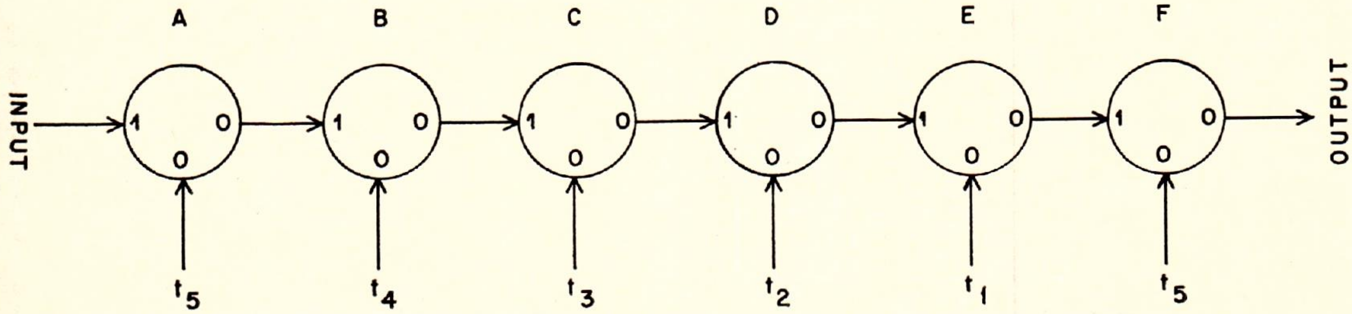

FIG. 11. FIVE-CORE/four bit MSR has multiple drivers.

by either power interruptions or noise transients. Storage is nonvolatile, and although read is inherently destructive, recirculating provisions (as in the ping-pong) are incorporated easily.

Circuit Types

Digital magnetic circuits, as employed in MSR configurations, can be loosely classified in two principal categories—*core-diode* types and *diodeless core* types—and a great many smaller categories of more unconventional and less broadly employed techniques.*

---

*The names of many of the circuits discussed are trademarks of the corporations mentioned.

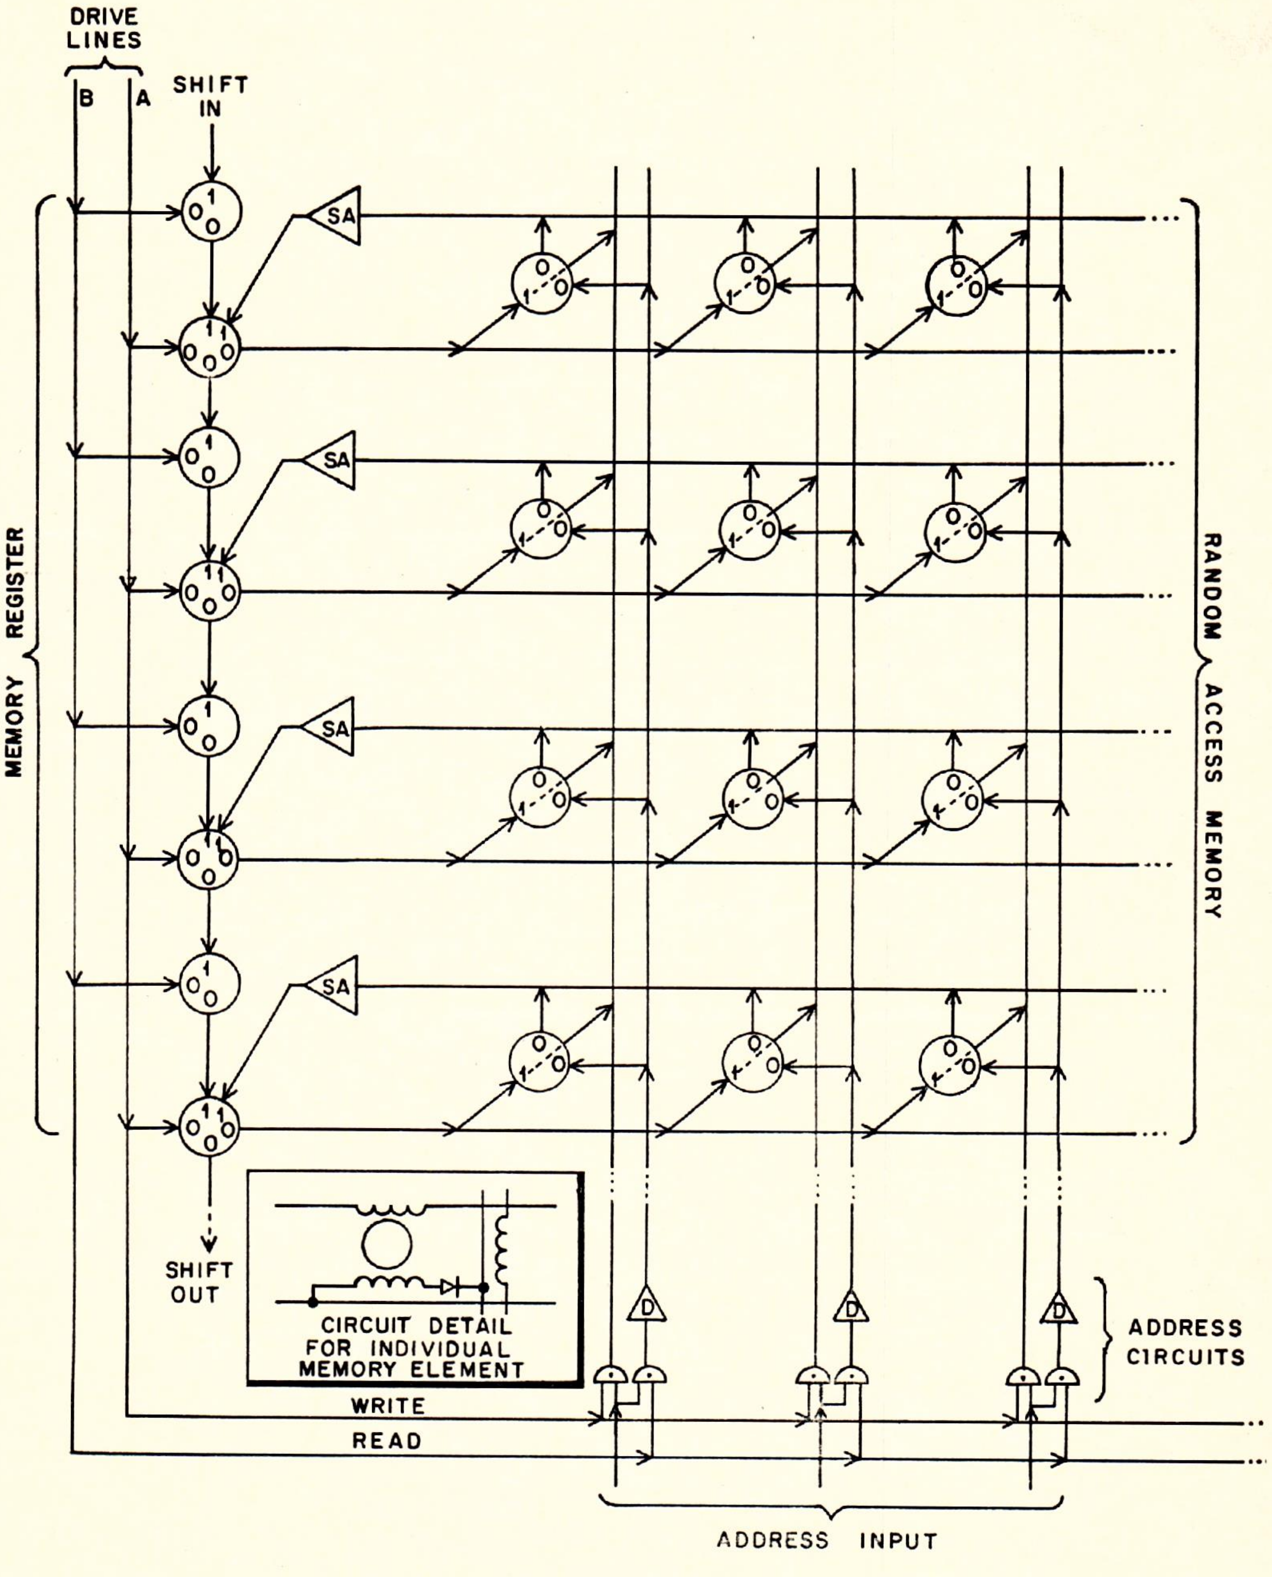

FIG. 12. RANDOM-ACCESS memory.

A comprehensive text on theory, electrical design, circuit design, and application of these circuit types was recently published.[3]

Core-diode circuits, first reported in 1950,[4] have inherent speed advantages over diodeless circuits and a power-consumption advantage on the order of three or four to one. Of the two most widely used core-diode types (*parallel* and *serial* magnetic pulse amplifiers), the parallel two-core/bit circuit takes most advantage of the inherent properties of the core, and is the most thoroughly analyzed. The parallel circuit also makes greater use of the best properties of the diode. The Burroughs *Bimag*,[1,2,3] the example for this article, is the principal example of the parallel magnetic pulse amplifier. The series circuit (the *Ferractor*[3,5] of Remington Rand Univac and the *Cypak*[3,6] of Westinghouse Electric Company, both stemming from the Ramey amplifier[7] of 1952) is more analogous to transistor and vacuum-tube circuitry. Series core-diode circuits operate at higher speeds and higher power levels than the parallel circuits, and are extensively employed in both data processing and industrial control systems. Average component count, however, is higher than that of the parallel circuits.

Another class of core-diode circuits useful in complex logic systems is the one-core/bit delay circuit, such as that of DI/AN Controls, Inc.,[3,8] which employs a delay element between stages to permit read and write effectively at one clock time. The circuit of Fig. 5 is typical of this technique. Other one-core/bit techniques also have been proposed.[9]

The diodeless core circuits have an enormous reliability potential because of the elimination of the diode, but are slower than core-diode circuits and consume more power. Diodeless circuits have been developed with standard toroidal cores and with various types of multiaperture cores. Tolerances, particularly in the *MAD* multiaperture core circuit of Stanford Research Institute,[3,10] are sometimes critical. Nevertheless, because of inherently high reliability and fascinating logic possibilities, many diodeless techniques are being studied. Diodeless types employing standard toroidal cores are those of IBM[11] and a three-core/bit technique developed in France.[12] Another multiaperture type is that described by Gianola,[13] which achieves fairly wide tolerances at the expense of speed and power. The Gianola circuits can be built from RCA *Transfluxor*[14] cores, which differ from the *MAD* cores in having only one minor aperture.

Other circuit types upon which MSRs have been based include parametrons[15] (subharmonic oscillators), ferroresonant carrier magnetic pulse amplifiers (such as the *Ferristor*[16] of Beckman Instruments, Inc.), hybrid transistor-magnetic pulse amplifiers (The Burroughs *Tramag*,[3,*] and circuits of DI/AN[3,17] and Milgo Electronic Corporation), and thin magnetic films.[18] In one interesting version of the latter, developed by American Systems, Inc.,[19] information is advanced by shifting magnetic domains down a continuous film strip, rather than by inducing voltages in transfer loops from element to element. An- excellent review of the many types of square-loop portation[20] permits use of toroidal cores to effect 20-microsecond addition at a 100 kc clock rate. There are, of course, many other examples (and organizations) to be noted. J. Haynes[21] has prepared an other technique recently proposed by Burroughs Cor- core circuits (but not including the ferroresonant and nonsquare-core types).

References

1. D. Loev, W. Miehle, J. Paivinen, and J. Wylen, "Magnetic Core Circuits for Digital Data Processing Systems," *Proc. IRE*, Vol. 44, February 1956, pages 154-162.

2. S. B. Disson and E. F. Myers, "Magnetic Systems for Ordnance and Space Applications," *Proc. National Aerospace Electronics Conference*, May 1961.

3. A. J. Meyerhoff, editor, *Digital Applications of Mag-* 1956, pages 654-665.

6. J. P. Baker, "CYPAK Systems; Application of Logic Functions," *Westinghouse Engineer*, Vol. 16, July 1956, pages 107-111.

7. R. A. Ramey, "Single-Core Magnetic Amplifier as a Computer Element," *Trans. AIEE*, Vol. 71-1, 1952, *netic Devices*, John Wiley & Sons, Inc., New York, 1960.

4. A. Wang and W. D. Woo, "Static Magnetic Storage and Delay Line," *Journal of Applied Physics*, Vol. 21, January 1950, pages 49-54.

5. T. H. Bonn, "The Ferractor," *Conference on Magnetism and Magnetic Materials*, Vol. T-91, October pages 442-446.

8. R. D. Kodis, S. Guterman, and S. Ruhman, "Logical and Control Functions Performed with Magnetic Cores," *Proc. IRE*, Vol. 43, March 1955, pages 291-298.

9. V. L. Newhouse and N. S. Prywes, "High-Speed Shift Registers Using One Core per Bit," *IRE Trans. on Electronic Computers*, Vol. EC-5, September 1956, pages 114-120.

10. H. D. Crane, "A High-Speed Logic System Using Magnetic Elements and Connecting Wire Only," *Proc. IRE*, Vol. 47, January 1959, pages 63-73.

11. L. A. Russell, "Diodeless Magnetic Core Logical Circuits," *1957 IRE Convention Record*, Pt. 4, pages 106-114.

12. M. Dumaire, H. Jeudon, and M. Lilimand, "Logical Circuits Using Magnetic Cores and Capacitors," *Proc. of Special Technical Conference on Nonlinear Magnetics and Magnetic Amplifiers*, Los Angeles, August 1958, pages 243-266.

13. U. F. Gianola, "Integrated Magnetic Circuits for Synchronous Sequential Logic Machines," *Bell System Technical Journal*, Vol. 39, March 1960, pages 295-332.

14. J. A. Rajchman and A. W. Lo, "The Transfluxor," *Proc. IRE*, Vol. 44, March 1956, pages 321-332.

15. E. Goto, "The Parametron, A Digital Computing Element Which Utilizes Parametric Oscillation," *Proc. IRE*, Vol. 47, August 1959, pages 1304-1316.

16. C. Isborn, "Application of Ferristor Type Saturable Reactors to Digital Instrumentation," *Proc. ISA Annual Conference*, Vol. 11-1, September 1956, Paper S6-32-2.

17. S. Guterman and W. Cary, "Transistor-Magnetic Core Circuit: A New Device Applied to Digital Computing Techniques," *1955 IRE Convention Record*, Vol. 3-4, pages 84-96.

18. W. E. Proebster and H. J. Oguey, "High-Speed Magnetic-Film Logic," *Digest of Technical Papers*, 1960 Solid-State Circuits Conference, pages 22-23.

19. K. D. Broadbent and F. J. McClung, "A Thin Magnetic-Film Shift Register," *Digest of Technical Papers*, 1960 Solid-State Circuits Conference, February 1960, pages 24-25.

20. E. T. Walendziewicz, "MODAC: Logic Design of a 100 Kc Magnetic Computer with a 20-Microsecond Add Time," *Proc. 1962 IRE International Convention.*

21. J. L. Haynes, "Logic Circuits Using Square-Loop Magnetic Devices: A Survey," *IRE Trans. on Electronic Computers*, Vol. EC-10, No. 2, June 1961, pages 191-203.

---

[*] Registered trademark of the Burrough Corporation.

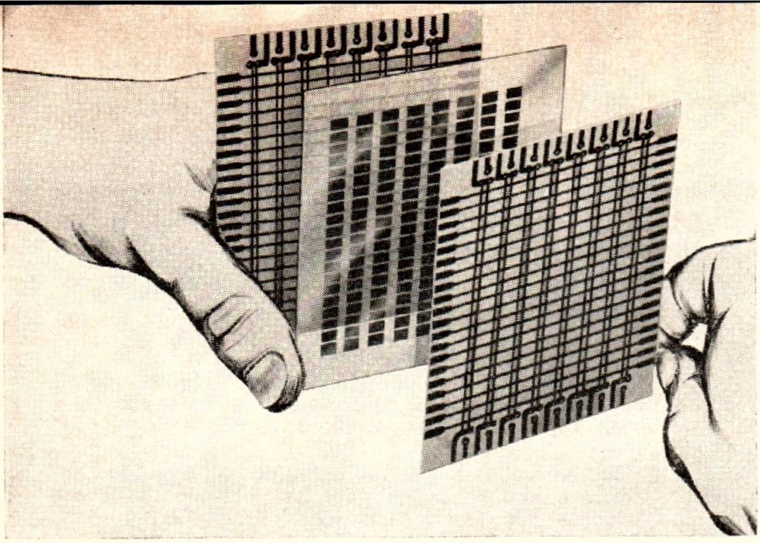

FIG. 1. **EACH PLANE** of a thin-film memory is sandwiched between two printed circuits containing the conductors for the drive (clock), write (information), and read (sense) currents. This plane has 20 x 8 bits, can store 20 8-bit words.

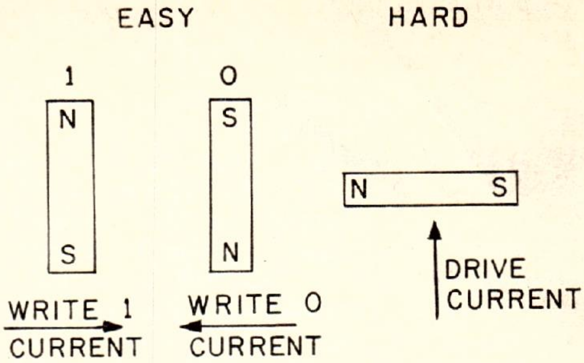

FIG. 2. **EACH THIN-FILM** element can be magnetized to behave like a magnetic dipole. The dipole can be made to have a better B-H characteristic in one direction (preferred or "easy" direction). Drive current in direction shown causes polarity of magnetization to be in "hard" direction. Write current, in directions shown, causes polarity to take the "preferred" 0 or 1 direction.

Thin-Film Memories

Thin-film memories offer advantages of high speed of operation, low cost of fabrication, and simplicity of design. Their versatility in size, shape, and connector arrangements makes them attractive devices for the computer and data storage fields. Here are the principles, operation, features and construction techniques for thin-film memories.

ERIC E. BITTMANN
Burroughs Corporation
Burroughs Laboratories

STORAGE ELEMENTS in the thin-film memories to be discussed are planar 80-20 Ni-Fe films less than 5000 Angstroms thick. They are obtained by vacuum deposition or by electroplating the Ni-Fe alloy onto a suitable base. These thin films can be magnetized rapidly (in the nanosecond range); memories with cycle times of 0.2 μsec, or less, are possible. High speed and low cost make thin films a desirable memory device.

Operating Characteristics

Any magnetic material with a square-loop B-H characteristic can be used as a storage element in a random-access memory. Behavior of thin-film memories can be compared with a magnetic dipole (Fig. 2). The two states of the dipole represent the storage of a "1" or a "0" in each thin-film element.

Films deposited on a substrate while under the influence of a magnetic field of 10 to 50 oersteds have a square B-H characteristic parallel to the axis of the dipole, as shown in Fig. 2. This is called the "easy" or preferred direction of magnetization. When the film is magnetized in a direction perpendicular to this easy direction, the film B-H characteristic is as shown in Fig. 3. This is known as the "hard" direction, shown in Fig. 2.

The principle of operation is as follows: The film element is magnetized in the "hard" direction by a *drive current*. A second current (write or information current) on a separate wire causes the direction of magnetization to swing into one of the two "preferred" polarities (Fig. 2). This is the recording process for a 0 or 1 bit.

The directions of the drive and write currents are shown in Fig. 2. Note that the conductors for these currents must be perpendicular.

To magnetize a selected area of thin-film to the 0 or 1 state, two magnetic fields, perpendicular to each other, are applied—the drive field in the "hard" direction and the information field in the "easy" direction. When both currents are applied, the resultant fields lies between the hard and easy directions, oriented toward either 0 or 1 (depending on the direction of the write current). Upon

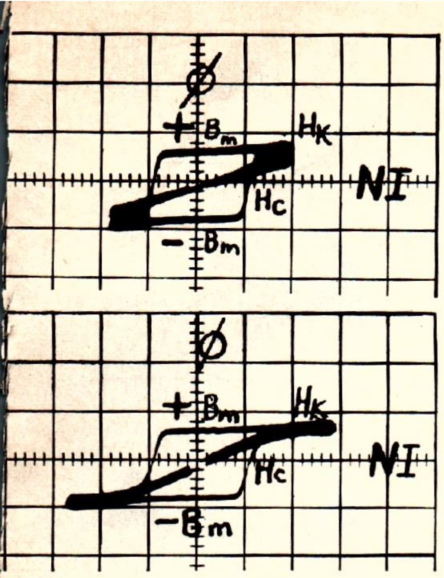

FIG. 3. B-H characteristic in the "preferred" (top) and "hard" directions (bottom). Square and linear loops are superimposed. Note that more drive must be applied to lower system, causing opening of the linear loop.

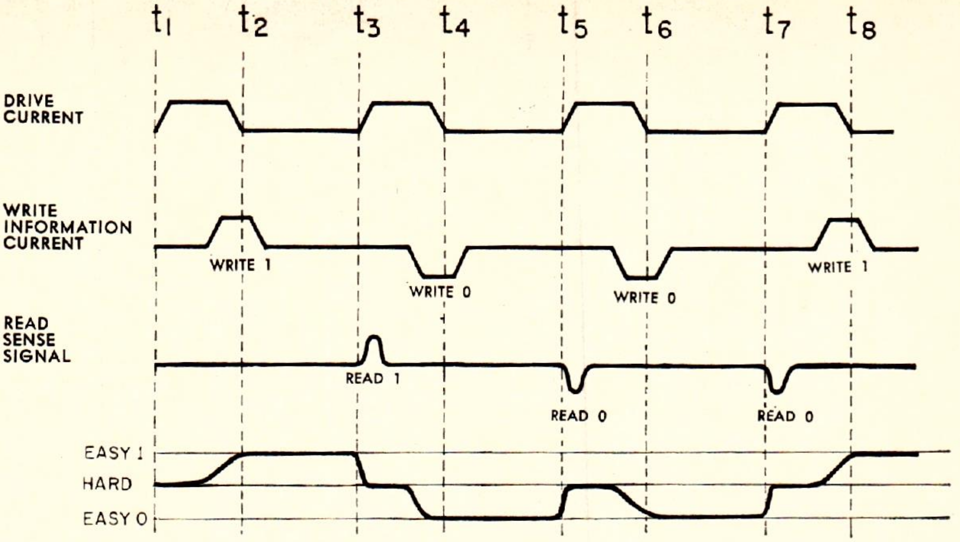

FIG. 4. IN OPERATING a thin-film memory, the drive and write currents must be timed so that they overlap, with the former leading. When drive current is on (t_1), bit is magnetized in "hard" direction. When drive current goes off (t_2), write current causes element to assume the easy "1" polarity.

removal of the drive field, the write current causes the dipole to assume the desired 0 or 1 state, depending on the direction of the write current.

In Fig. 4 the drive current goes on at t_1, off at t_2, on at t_3 etc. When on, it causes the bit to be magnetized in the hard direction.

The write current comes on before the drive goes off, and continues after the drive goes off. Thus, each time the drive goes off, the remaining write current establishes the polarity of the bit—0 or 1. For example, at t_2 the bit goes to the 1 state; at t_4 the bit goes to the 0 state.

For reading, a third (sense) conductor, parallel to the write conductor, senses the polarity of the bit each time the drive (clock) pulse arrives. In effect, each element is interrogated (destructively) by each drive pulse.

This is shown by the third and fourth lines in Fig. 4. At time t_3, the change in polarity of the bit caused by the drive current induces a positive pulse in the sense (read) conductor.

At time t_5, the drive pulse causes the polarity to change in such a way as to induce a negative pulse in the read conductor, indicating that a 0 bit had been in the element.

Similarly, at t_7 the drive current indicated that a 0 bit had been in the element at the previous clock pulse.

Linear-select Thin-film Memories

Available memory planes (Fig. 1) have 20 x 8 bits (rectangles). These planes are sandwiched between two printed-circuit boards (each ten mils thick), which contain 20 drive conductors and eight sense and information conductors. Each rectangle has a dimension of 3/16″ by ⅛″ and gives a signal output of 5 mv when driven by 1-ampere pulses of 0.04-μsec rise time. This memory plane can store 20 words at eight bits per word.

Four planes can be interconnected in a memory that will store 80 words (Fig. 5). The 80-word (drive) lines can be connected to a selection matrix (a diode matrix, for example) so that only one word in each plane is interrogated at any one time. The information lines connect to information drivers, and the sense lines to sense amplifiers.

A word line that has been selected by the diode matrix will receive a full drive-current pulse. This current will interrogate a row of eight thin-film elements. The rotational switching of the magnetic state of each film then will induce a signal in its associated sense (read) conductor. The amplified sense signal is gated with the clock pulse (Fig. 4), and the output of the gate either sets or resets the associated flip-flop in the memory-information register.

Meanwhile, the drive current is still on. Before it can be terminated, an information-current (write) pulse must be sent through each of the eight information conductors. The polarity of each of the eight information currents is controlled from an associated flip-flop. This flip-flop controls the writeback of information into the selected row, or word, that has been read.

The unselected rows of memory bits are subjected only to this information (write) field, and if this field is kept small enough (smaller than H_c, the film's wall coercive force), no stored data will be lost.

The driving scheme described can be modified.

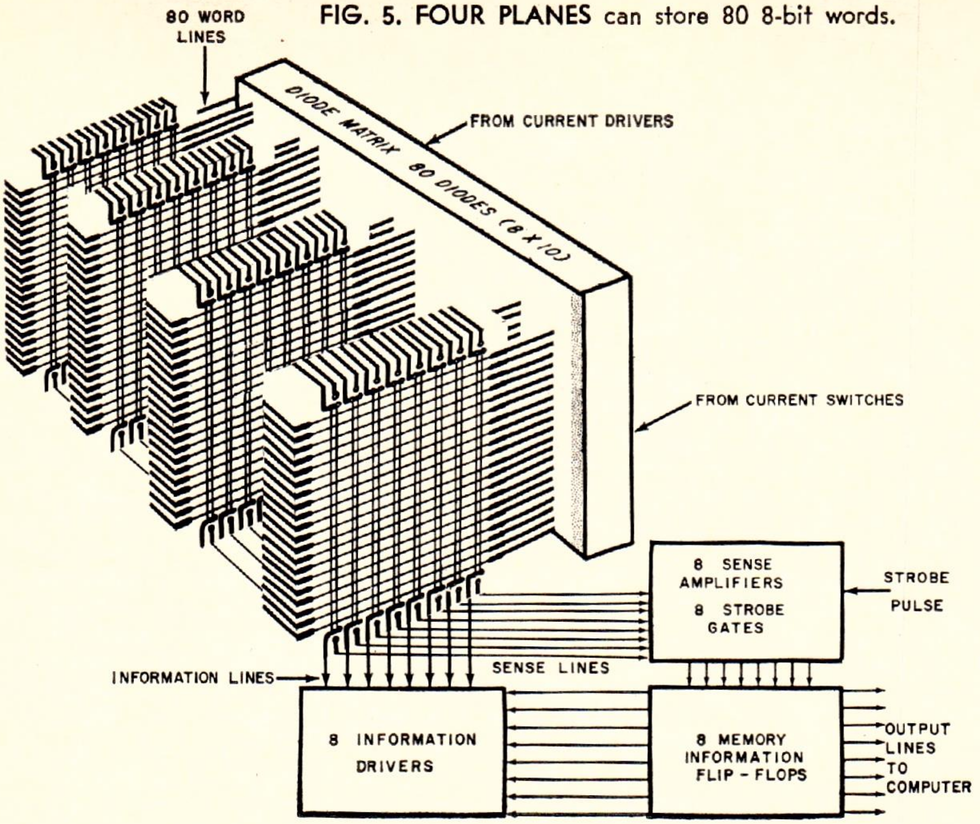
FIG. 5. FOUR PLANES can store 80 8-bit words.

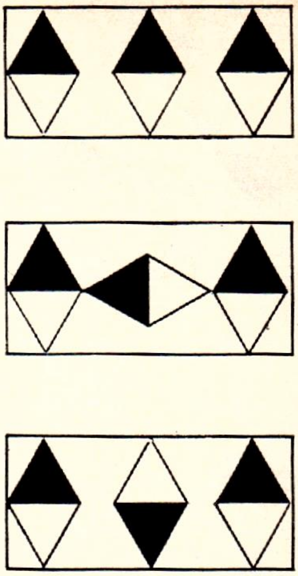
FIG. 6. DIPOLE configuration in nondestructure thin-film readout (NDR). Top shows dipoles in single domain; center, when drive field is applied; bottom, new arrangement after drive is removed.

Because the readout of information occurs during the leading edge of the drive current, and the writing of information during the trailing edge, the drive current can be replaced by two short pulses—one occurring during the leading edge, and one during the trailing edge, of the pulse. This method saves power in the transistors.

Instead of a bipolar information current, a d-c bias and a single-polarity current of twice the amplitude can be used. Absence of information current pulses insures the write-back of a zero into the selected film. (It might be more economical to design single-polarity-current drivers.)

The operating speed of such a memory is limited by the drive-current rise and fall times, and by the delay in all the selection circuitry, sense amplifiers, flip-flops, and drivers. Such memories presently are operated with a 0.2-μsec cycle time.

Coincident-current Thin-Film Memories

In the development of thin-film memories, one of the first ever described was a coincident-current type.[4] Here, operation is identical to that of a coincident-current core memory.

For coincident current, the drive fields are applied parallel to the preferred direction of magnetization. But noise induced in the sense conductor is hard to eliminate because of the parallel relationship of the sense and drive conductors. Also, switching time for the films is short (even if the switching is predominantly a domain-wall type) and the current rise time induces sufficient noise to obscure the signal.

Considering these shortcomings, coincident-current techniques are not presently good enough. Word-select-type memories give more satisfactory results.

Non-destructive-Readout Thin-film Memories

A non-destructive readout scheme, in which only a small drive field in the hard direction is applied to a single-domain thin film, has been described.[5] In this method the field rotates the dipole some degrees away from the preferred (easy) direction, but not the full 90°. After the drive is removed, the dipole rotates back to its original state.

Investigators tested portions of large areas of films (rather than small individual spots) in that mode of operation. They found that stored information cannot be read out an indefinite number of times, even when the drive field is small. It was observed that the readout signal amplitudes decreased after many readouts, and that there was a tendency for the readout signals to reverse their polarity.

This behavior is illustrated in Fig. 6. If a single-domain film is represented by three (instead of one) dipoles, and if the center one is rotated 90° by a drive field and then left on its own, it will fall to its opposite state. (The dipoles on each side of the center one are magnetized to the opposite state.) Readout fields applied only to the center dipole cause it to rotate to the hard direction, but upon removal of the field, the dipole falls back to its original state. This type of non-destructive readout is quite satisfactory.

Domain Structures Observed

Domain structures in thin films can be observed with a Kerr magneto-optic apparatus.[6] Here, polarized light is reflected from a thin film. The plane of polarization is rotated a few degrees to the right or left, depending on the magnetic state of the film. Domains will appear as either white or black areas, and areas magnetized in the hard direction will be grey.

Circuits

Three types of circuits are associated with a thin-film memory—the driver, the information (write), and the sense (read) amplifier.

Drivers. Fast-current drivers, delivering 1-ampere pulses with rise and fall times of 20 nanoseconds, can be constructed with the high-speed transistors presently available. Both PNP and NPN transistors are used, especially when word-selection occurs in a diode matrix. The information driver has to deliver less current than the read drivers, and bipolar drivers with NPN and PNP transistors are easily designed.

Sense Amplifier. The sense amplifier is of the differential type. It rejects the common-mode noise signal. The sense and information conductors of neighboring memory planes are interconnected in a noise-canceling fashion by reversing the sense-conductor connections.

Thin-Film Fabrication

Several techniques have been tried for the fabrication of thin films. Successful results have been obtained in each of the following methods:

Evaporation. In this method,[1] Ni-Fe is transferred to a hot, dielectric substrate in the presence of a magnetic field of 10 to 50 oersteds. The substrates, usually glass, are chemically cleaned before they are inserted into the vacuum chamber. The substrates are placed 5″ to 10″ from the evaporating source, and are heated, by a substrate heater, to 300 to 350°C.

The raw material can be placed either in a crucible or onto a filament. The crucible contains the Ni-Fe alloy and is heated by an induction heater. But uniform depositions, by this method, can be obtained only over a small area. Either Ni-Fe or tungsten materials can be used as filaments. In the latter case, the Ni-Fe material is attached to the tungsten filament. Because of the versatility of filament arrangements, the filament method produces a more uniform evaporation over a larger area.

Evaporation through a mask produces the desired arrays. Insulators and conductors can be deposited in more elaborate systems.

Sputtering. The heated substrate is in an argon atmosphere, at reduced pressure. The Ni-Fe source is some distance away from the substrate and held at a positive electrical potential of about 500 volts. Argon gas ions are attracted to the positive source, and they hit it with sufficient energy to knock Ni-Fe molecules from the source. These molecules strike the surrounding substrate and coat it with a thin layer of Ni-Fe. The magnetic field present causes the desired preferred direction. As the rate of material deposition during sputtering is slow, it is difficult to control the process.

Vapor Deposition (Pyrolysis). Dried nitrogen gas is introduced in a chamber which contains a liquid solution of nickel-iron carbonyls. The gas becomes saturated with these carbonyls and is sent to a plating chamber containing the substrates. This chamber is heated inside an infrared oven. At these elevated temperatures the carbonyls decompose into a metal alloy and carbon monoxide. The carbon monoxide and the nitrogen are fully burned, in a Fisher burner, after leaving the plating chamber. The metal deposits onto the substrates and yields the desired thin films. The process is relatively simple and shows promise, but it is poisonous. It is not recommended in a simple laboratory setup.

Electrodeposition. Thin films can also be fabricated by deposition[2] of Ni-Fe alloy, onto a dielectric substrate, from a plating solution, in the presence of a magnetic field. The glass substrates have to be made conductive before plating occurs. One method is sputtering material on the substrates in sufficient quantity to insure conduction. In another method, silver or gold is sprayed on, or a chemical silvering process is used. The thickness of this conductive layer has an effect on the film; controls are needed to insure uniformity and reproducibility.

After the plating process is over, discrete bits can be arranged by etching out the unwanted material.

All of these fabrication processes—evaporation, sputtering, vapor deposition, and electrodeposition—have yielded usable thin films.

References

1. M. S. Blois, Jr., "Preparation of Thin Magnetic Films and their Properties," *Journal of Applied Physics*, Vol. 26, August 1955, Pages 975-980.
2. I. W. Wolf, H. W. Katz, and A. E. Brain, "The Fabrication and Properties of Memory Elements Using Electrodeposited Thin Magnetic Films of 82-18 Nickel Iron," 1959 *Proceeding of Electronic Components Conference*, AIEE or IRE.
3. C. Kittel, "Theory of the Structure of Ferromagnetic Domains in Films and Small Particles," *Physics Review*, Vol. 70, December 1946, Pages 965-971.
4. A. V. Pohm and S. M. Rubens, "A Compact Coincident Current Memory," 1956 *Proceedings of Eastern Joint Computer Conference*, IRE.
5. K. Chu and J. R. Singer, "Thin Film Magnetization Analysis," *Proceedings* of *IRE*, Vol. 47, No. 7, July 1959, Pages 1237-1244.
6. J. W. Hart, "Domain Behavior in Thin Magnetic Films," *Transactions*, 1960 *Nonlinear Magnetics and Magnetic Amplifier Conference*, AIEE.

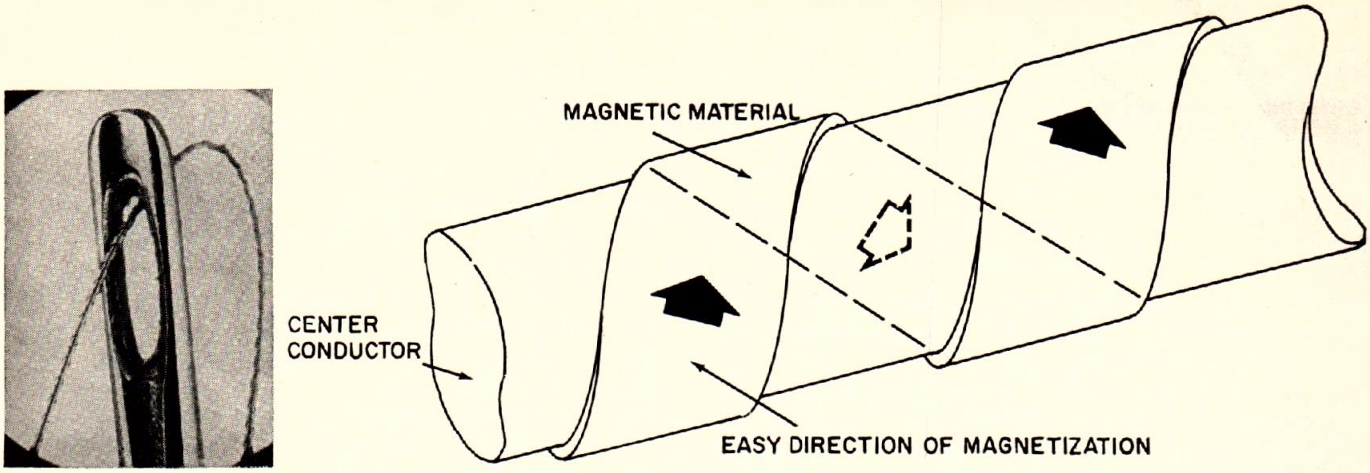

FIG. 1. WRAPPED TWISTOR. Arrows represent easy "1" state. Easy "0" state would be arrows in opposite direction. Element is smaller than eye of needle.

Twistor Memories

Memory based on a continuous low-cost wrapping of magnetic material on a single conductor has great promise for large low-cost memory system.

PHILIP N. McCABE

Burroughs Corporation—Laboratories
Paoli, Pennsylvania

THE DISCOVERY at Bell Telephone Laboratories by Andrew H. Bobeck that a *helical direction of magnetization* could be induced in a nickel wire by *torsion* led to the development of the "twistor."[1] The early embodiments of this device consisted of a small-diameter (0.005″) nickel wire under torsion, and a copper conductor wrapped in the form of a solenoid around the circumference of the nickel wire. Current applied to the solenoid produced an axial field in the twisted nickel wire, causing a voltage to appear at the end of the nickel wire as a result of the changing magnetization. The twisted nickel wire has since been replaced by a wrapped structure (because of the impedance and resultant attenuation of the switching signal). The modern twistor does *not* use a wire under torsion.

Wrapped Twistors

In the twistors' present form a thin, narrow tape (or wire) of magnetic material is wrapped in the form of a helix on a copper conductor (Fig. 1). As the material is anisotropic*, it has an "easy" and a "hard" direction of magnetization. The easy direction of magnetization is along the helix as shown in Fig. 1.

The wrapped twistor is within a solenoid which produces an axial magnetic field H_A (Fig 2).

The wrapped twistor (with associated solenoid) is used in a linear-select memory as follows: A current applied to the twistor center wire produces a circumferential magnetic field. A current in the solenoid produces an axial magnetic field. When these currents are applied coincidentally, they produce a resultant magnetic field whose direction can be aligned with the easy direction of the magnetic tape.

Assume point A in Fig. 2 to be a point *on the surface* of one of the wraps of magnetic material of a twistor. When a circumferential field H_C of insufficient magnitude to change the state is applied in coincidence with an axial field H_A of insufficient magnitude to change the state, the fields add vectorially to give a resultant field which *is* strong enough to change the state—in line with the easy direction. The write-1 operation consists of applying currents as shown in Fig. 2. The write-0 operation is shown in Fig. 3.

The state of the twistor bit is read out by application of a large reversed solenoid field (H_R) which is of sufficient strength to swing the 1 vector counterclockwise from 1, past the hard direction, and toward

---

*i.e., magnetic properties depend on direction.

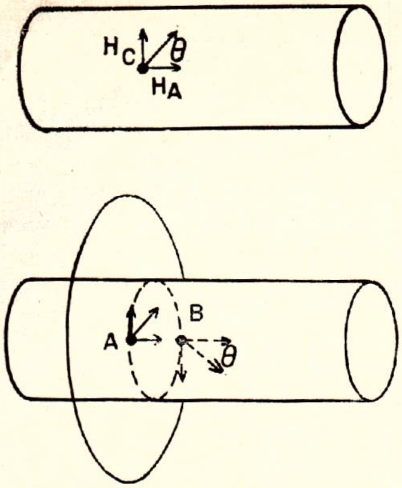

FIG. 2. APPLIED FIELDS H_C and H_A produce easy 1 state. This is the write 1 operation. Point A is point on near surface of conductor. Point B is on far side of conductor.

the easy 0 direction (Fig. 4). During this time a voltage is produced by the changing flux, and sensed across the ends of the twistor center conductor. The read-0 action is shown in Fig. 5.

Characteristics

A typical set of operating characteristics are given for a twistor with a 0.015″-wide magnetic tape wrapped at a 30° angle on a 0.007″-diameter copper wire with a 0.005″ gap. The solenoid is 12 turns of #30 wire:

Write Current (Solenoid)	**120 ma**
Information Current (Twistor)	**112 ma**
Read Current (Solenoid)	**500 ma**
Disturbed One Output	**39 mv**
Disturbed Zero Output	**6 mv**
Switching Time	**0.4 μsec.**

Fabrication

A wrapping machine wraps a twistor by pulling a horizontal copper wire from left to right and at the same time rotating an arm, with a bobbin of magnetic material on it, around the copper wire. By varying the width of the magnetic material and the ratio of the advance of the copper wire to the number of rpm's of the bobbin arm, a wide variety of twistors can be wrapped.

Geometry Effects on Operation

Variation in geometry has two major effects on operational characteristics. The first is to change the cross-sectional area of magnetic material under the solenoid, thereby changing the output signal; the second is to change the angle between the easy direction of magnetization and the applied fields. As this easy direction is not in line with the applied fields, the twistor does not seem to be a square-hysteresis-loop device. As soon as the geometry of the twistor is changed so that the angle θ in Fig. 2 becomes, let us say, less than 45°, the axial field (due to current in the solenoid) sees an apparent increase in squareness and thus this field becomes more effective. In practical terms, this means that less solenoid current is necessary to saturate the twistor.

When the angle θ becomes greater than 45°, the circumferential field sees an apparent increase in squareness and it becomes more effective. Experiments have shown that, for an increase in angle from 30° to 60°, the necessary current changes by a factor of two. From the standpoint of economy, a low angle (30°) is the most economical because the larger current (the read current) is on the solenoid, and any decrease in this current effects a large saving. A high angle (60°) would seem to be the most economical

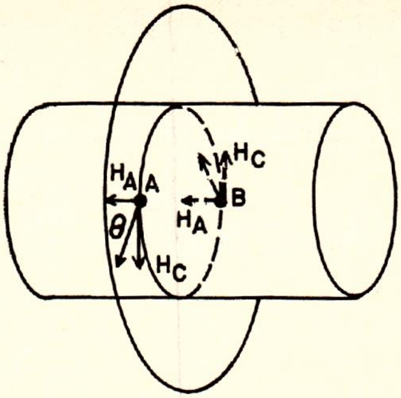

FIG. 3. WRITE 0 operation for point A (near side) and B (far side) on surface of conductor.

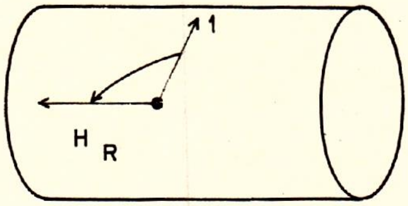

FIG. 4. IN READING a 1, the axial H_R field causes the 1 vector to rotate counterclockwise on near surface.

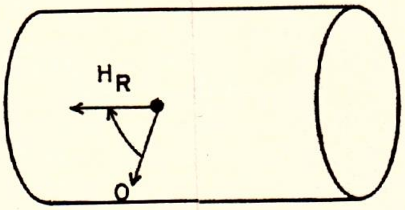

FIG. 5. IN READING a 0, the H_R read field causes point on near surface to rotate clockwise.

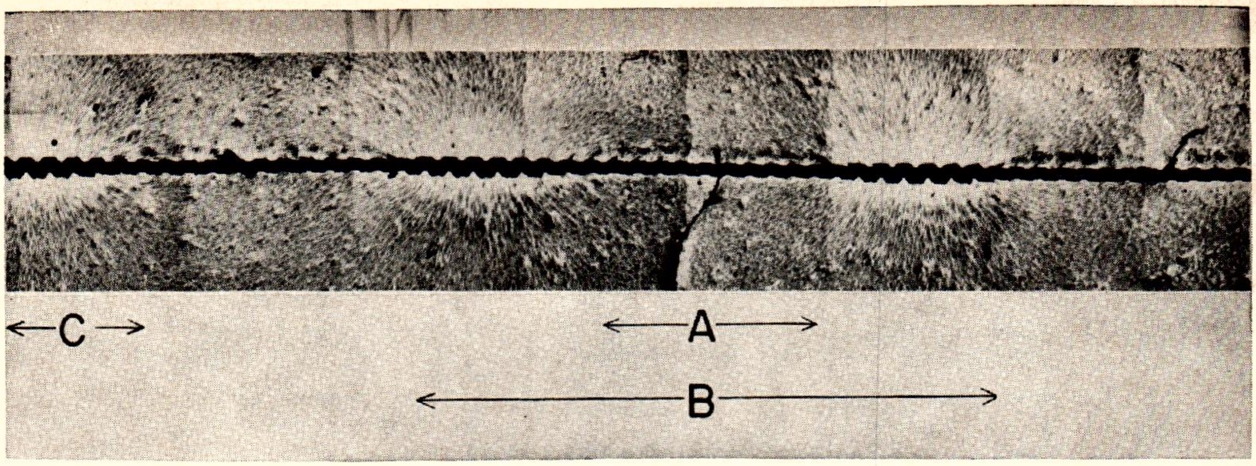

FIG. 6. A IS LENGTH of twistor in writing solenoid. B is effective length of twistor bit due to fringing of writing field. C is beginning of next memory bit.

for the output signal because this would give the greatest number of wraps of magnetic material per linear inch. However, since it is cross-sectional area that is important, a large output signal can be obtained by decreasing the gap between wraps of the magnetic material.

The increase in twistor wire current necessitated by a low angle can be off-set by wrapping the magnetic material on a small-diameter wire. Magnetic materials ranging in width from 0.015″ to 0.003″ have been wrapped on wire diameters from 0.010″ to 0.003″, with gaps from 0.020″ to 0.001″.

Twistor Geometry

A twistor bit in an actual memory is a fairly nebulous thing. This is a result of the fact that the flux path is not closed in the magnetic material (as it is in a toroidal core). Fig. 6 is a photomicrograph of a wrapped twistor. The entire twistor was magnetized from right to left, and then the section labeled B was switched so that it became magnetized from left to right. The twistor then was immersed in a colloidal suspension of iron-oxide particles, which were attracted to the free poles at the bit ends. The solenoid used to write the information into the twistor covered the section labeled A.

The fringing of the solenoid field is evident. The amount of fringing depends on the solenoid field strength and the demagnetizing field of the bit itself. The bit length has a tendency to increase (creep) under repeated partial switching, producing "cross-talk" or interference with adjacent bits. This effect can be reduced by (1) increasing the space between solenoids, (2) effecting a reverse field at each end of the solenoid, or (3) placing a shorting turn at each end of the solenoid.

The output of a twistor bit depends only on the cross-sectional area of magnetic material influenced by the solenoid. Typically, where the calculated cross-sectional area *under* the solenoid should give an output of 10 mv-μsec, the actual output is increased by fringing and creeping to 15 mv-μsec. This 50% increase in output is typical for calculated outputs of from 2 to 10 mv-μsec.

Twistor Encapsulation

The fabrication of long lengths of small-diameter wrapped twistors leads to a problem of handling. This problem has been solved by encapsulating from twenty to fifty lengths of twistor, side by side, in a thin sheet of plastic to give a package approximately 0.010″ thick by 2″ to 5″ wide, by any length. The belt, in which the twistor wires are on 0.100″ centers, is flexible and easy to handle. This form also has the advantage of eliminating the interplane connections.

As twistors are continuous, the belt, rather than being cut into individual planes, is folded back and forth on itself, necessitating connectors only at the top and bottom of the stack.

The magnetic material used in the twistor is unannealed and, therefore, relatively strain-insensitive. Thus the handling and folding of the twistor belt does not affect the magnetic properties. Experience has shown that long lengths of twistor (up to 4000 feet) can be wrapped[2] within a tolerance of ±10%. The variations are such that no bit is completely bad; this eliminates the necessity of cutting out a bit and patching the wire.

The unannealed magnetic material has a relatively high coercivity (from 1 to 4 oersteds). This requires high currents in the word drive on the solenoid. The problem is compounded by the relatively poor coupling between the solenoid and the twistor. As the solenoid does not thread the twistor (as in a toroid), the coupling depends largely on the proximity of the solenoid to the twistor wire. Also, as the solenoid is not a long, thin solenoid, the uniformity of the field is poor and the strength of the field, as seen by the twistor, is sensitive to position. The coupling of the magnetic material to the conductor upon which it is wrapped is almost an order of magnitude better, and the currents here are in the order of 100 milliamperes.

Read Fields

In twistor memories, the high read-fields are obtained in one of two ways. The first is to use a current step-up transformer and a single-turn secondary. In this case, one ferrite switch core is selected from a matrix of cores by the coincidence of two currents. These two currents appear on a multiturn primary of the switch core. The secondary of each switch core is the single-turn word (or read winding) of the twistor array. This single-turn winding is usually a copper strip which encircles the twistor belt, 0.100″ wide by a few thousandths thick.

The other type of word winding is a multiturn solenoid of from 5 to 15 turns of approximately #30 AWG copper wire (Fig. 7). The single-turn copper-strip solenoid can be placed close to the twistor belt so that its air inductance is very low, whereas the multiturn solenoid must be wrapped on a separate form from the twistor belt because the belt is not strong enough to support the solenoid; consequently, its air inductance is high—approximately 10 μhenries.

The air inductance of the multiturn solenoid can be reduced by placing a twistor belt on each side of a conducting plate and then wrapping the solenoid around this form. The conducting plate acts as an eddy-current shield, effectively reducing the cross-sectional area of the air under the solenoid. This method has the added advantage of increasing the package density because a twistor belt is placed on each side of the plate.

A fabrication advantage of a twistor array—namely, that it is only a two-wire system—is partially offset by noise considerations when a memory is assembled. As the sense wire (the twistor wire) is unidirectional throughout the memory, noise must be cancelled by placing a return wire in the twistor belt for each twistor, and treating the wrapped twistor and the return as a transmission line.

Nondestructive Read

The twistor memory operation up to this point has been destructive—that is, the read current always leaves the magnetic material in the zero state. In most memories, this information is automatically reinserted into the memory. However, the twistor can be modified to operate in a nondestructive manner. This requires the return of the twistor to its original state after readout, a condition that can be accomplished in two ways:

The first is to write with a long solenoid, and read with a much shorter solenoid in the center of the long solenoid. The long twistor-bit magnetized by the write solenoid is split into three sections during read time by the short read-solenoid. The two end sections, which are magnetized in the same direction, force the center section to return to their direction after the read field is removed, thus effecting a nondestructive operation.

The second method is to bias each twistor-solenoid intersection with a permanent magnet.[3] The field from the permanent magnet is strong enough to return the twistor to its original state after a read operation. Permanent magnets are placed on a card over each plane of the twistor memory and can be easily removed and replaced with a new pattern.

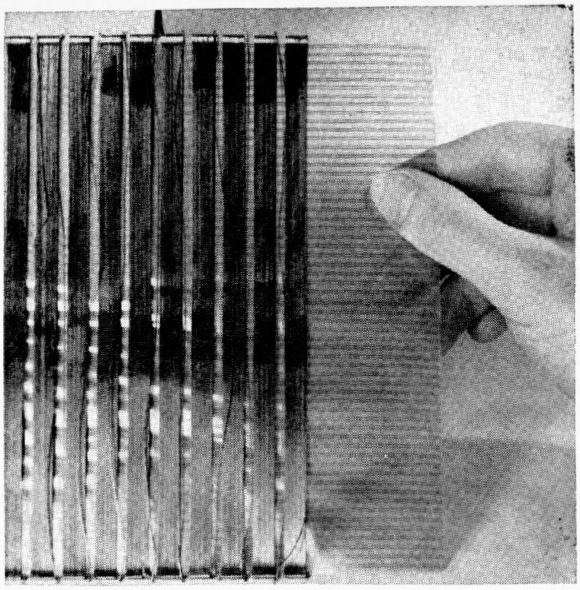

FIG. 7. THESE IMBEDDED TWISTORS form a twistor "belt"—a complete memory plane. Each pair of horizontal wires contains a twistor and a return conductor. The vertical wires are solenoids. A bit of information can be stored at the intersection of a solenoid and each twistor.

Twistors as Logic Devices

As a logic device, the twistor can be used to produce a *shift register* by using the fringing and creeping tendency of the twistor bit.[4] The design takes advantage of the intersection effect between closely spaced bits to slide a magnetized zone down a twistor wire with a multiphased clock pulse. (This is the only report to date of a twistor logic device.)

The twistor has operational characteristics and manufacturing capabilities of great promise. As its magnetic material is a metal, it can operate successfully at a high temperature. The 4-79 molybdenum permalloy from which most wrapped twistors are made decreases in output flux by 0.1%/°C for increasing temperature.

The ability to wrap a continuous rather than a discrete element, and to include a conductor in the same operation, lowers the cost and increases the twistor's reliability by eliminating a large amount of handling.

These advantages make the twistor a good candidate for large low-cost memory systems.

References

1. *The Bell System Technical Journal,* Vol. XXXVI, No. 6, Nov. 1957.
2. Andrew H. Bobeck, "The Twistor—Past, Present and Future," 7th Annual Conference on Computers and Data Processing, July 28, 29, 1960, Estes Park, Colorado.
3. *Ibid.*
4. A. H. Bobeck and Robert F. Fisher, *Journal Applied Physics,* Supplement to Vol. 30, No. 4, April, 1959, 43S-44S.

FIG. 1. MAGNETOSTRICTIVE sonic delay line comprises nickel wire plus two magnetostrictive transducers. Input transducer converts input pulse signals (I_1) into dimensional changes in the wire. Delay results from time required for sonic pulses to travel to output transducer, which reconverts sonic pulses into output signals (E_2).

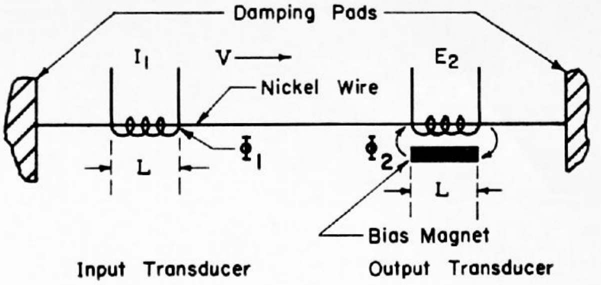

Delay Line Memories

Magnetostrictive sonic delay lines have become a popular type of storage. They can provide a small low-cost memory for computers with a clock rate from 500 kc to 2 Mc.

DAVE ELLIS, Burroughs Laboratories

ANY DEVICE that delays the passage of digital information can be said to store that information for a time equal to the delay time of the device. If the delayed information is fed back from the output to the input, it may be stored indefinitely as a pulse or pulse train moving continuously, or recirculating, around the delay loop. The use of delay lines as storage elements in digital computers is based on this principle. As the data or stored information is in continuous motion, the technique is one of dynamic storage, as opposed to the static storage provided by magnetic cores, tapes, drums, and discs.

Electromagnetic and Sonic Delay Lines

Delay lines are of two general classes—electromagnetic and sonic.

Electromagnetic delay lines are simply transmission lines carrying pulse trains. As the delay time generally available is not sufficient to store a large number of pulses (unless lines of enormous length are employed), this type of delay is not suited for conventional computer memories.

Sonic delay lines convert electrical input pulses into sonic pulses (sound waves), delay these pulses, and reconvert the sonic pulses to electrical pulses at the output. As sound travels much slower than electricity, the sonic delay technique provides a much greater delay time through a given length of delay medium than does the electromagnetic technique.

Several computers use sonic delay lines for storing digital pulses. The early Univac computer[1] used columns of mercury as the delay medium, with quartz piezoelectric crystals to convert electrical energy into sound energy, and *vice versa*. However, this technique is expensive and cumbersome compared with the *magnetostrictive sonic delay line*, a more recent development which has become the principal type of delay-line storage in modern systems.

Magnetostrictive Sonic Delay Lines

Several computers using magnetostrictive sonic delay lines as the storage element[2] have appeared on the market. This type of delay line (Fig. 1) uses acoustic delay in a metal wire. A transducer at the input converts the electrical inputs into sonic signals by using the principle of magnetostriction*; a similar transducer reconverts the sonic signals at the other end of the wire into electrical signals, also by using magnetostriction.

As there is no place to store the pulses in the wire for a long period of time, the data must be recirculated from the output to the input. Recirculation can be continued indefinitely before the data is used. A small

---

* Magnetostriction—the change in dimension of certain metals, such as nickel, when subjected to a magnetic field.

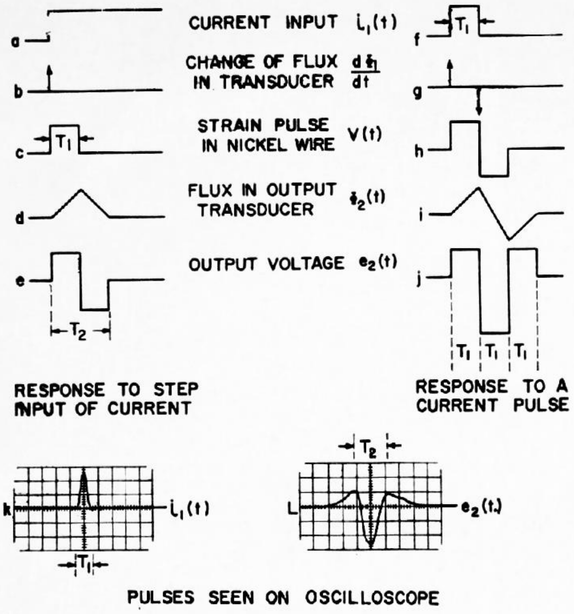

FIG. 2. IDEALIZED waveforms. T_1 is due to the length of the transducer coil and equals length of coil/velocity of strain wave on the wire.

amount of recirculation logic is used to resynchronize the data and reinsert it into the input transducer. Pulses are synchronized with the system clock each time around the recirculation loops. Thus, each bit of information can be located at a definite place in the loop at any specified time. The total delay through the line must be known and held constant to within a close tolerance.

Binary ONEs and ZEROs are represented either as (1) pulses of equal amplitude but opposite polarity, or (2) pulses and absences of pulses.

Environmental effects, particularly temperature, affect the sonic waves traveling down the wire, but the recirculation logic could be designed to recognize and correct the resultant variations in signal amplitude and timing due to environmental factors.

When compared to other storage devices, delay lines of this type offer an interesting compromise in speed, access time and cost-per-bit. They are less expensive but have longer access time than core storage; they have shorter access time but cost more per bit than disc or drum storage. They appear to be most useful where a small, low-cost memory is desired in a digital computer having clock rates of from 500 kc to 2 Mc. A typical magnetostrictive delay unit might operate, for example, at 1 megacycle/second and store 10,000 bits per delay line.

Magnetostrictive Transducers

In the basic form of the delay line (Fig. 1) a piece of magnetostrictive nickel wire is suspended between pads of acoustic damping material. The dimensions of the nickel wire are altered by the application of a magnetic field. For example, if a magnetic field is applied along the axis of the wire, the wire gets shorter. If a magnetic field impulse is applied to a small part of the wire, a dynamic strain pulse is made to travel down the wire with the velocity of sound. The precise velocity of the strain pulse depends on the mechanical properties of the wire material; the shape of the pulse depends on the magnitude and rate of change of the magnetic field.

In a simple delay line, a coil of wire at one end of the delay wire serves as the input tranducer, and a coil at the other end as the output transducer. The input transducer converts a current pulse into a mechanical strain wave on the wire, and the output transducer reconverts this mechanical wave to an electrical signal. The delay time through the line is simply the distance between the transducers divided by the velocity of sound in the wire.

Waveforms

Idealized waveforms for the delay line of Fig. 1 are shown in Fig. 2. The waveforms corresponding to a *step of current* applied to the input transducer are represented by figures 2a through 2e.

The effect of the transducer coil length on the strain wave is indicated in figure 2c: the minimum length of the strain wave is equal to the length of the transducer coil. Thus, the minimum time duration of the strain wave, when observed at a fixed point on the wire, is $T_1 = L/V$, where L is the length of the transducer coil and V is the velocity of the strain wave on the wire. As the physical length of the strain pulse and its time duration at a fixed point in the wire are related by the linear velocity term, the shape of the mechanical strain wave on the wire is similar to the shape of the strain wave.

In the region of the output transducer, a small permanent magnet creates a bias flux in the nickel wire. As the presence of strain changes the permeability of the wire, the flux linking the transducer coil (Φ_2) changes in proportion to the amount of the strain pulse under the coil. If the length of the output transducer is equal to that of the input transducer, the resulting flux change is as shown in figure 2d. The maximum flux change occurs when the strain wave is centered under the output coil. The induced voltage output from the coil (figure 2e) is proportional to $d\Phi_2/dt$.

The waveforms in figures 2f through 2j show that an input pulse of duration $T_1 = L/V$ causes the output shown in figure 2j. This output results from the superposition of two input current steps which are of opposite polarity and displaced from one another in time.

Figures 2k and 2L (bottom) show the actual waveshapes corresponding to the ideal waveshapes of figures 2f and 2j. The departure from the ideal is the combined effect of many non-ideal conditions in the device, including input pulse rise-time, electrical (magnetic) hysteresis in the material under the transducer coils, mechanical hysteresis in the delay wire, and dispersion due to internal reflections.

Several analyses have been made of these waveshapes.[3-7] From these, the optimum conditions for operation of the line have been determined. The time T_2 (figure 2e) is equal to approximately $2T_1$, twice the duration of the input pulse. A separation of $2T_1$

between input pulses represents a compromise between pulse resolution and pulse repetition rate. The optimum transducer coil length can be found by $T_1 = L/V$, where L is the effective length of the magnetic field of the input coil, and V is the velocity of sound in the nickel wire.

Longitudinal Versus Torsional Mode

The basic magnetostriction delay line described transmits strain waves along the delay wire in the form of compressions and elongations along the axis of the wire. This is referred to as the *longitudinal mode.* Solids also can transmit strain waves in the shear mode; a torsional moment in a thin wire causes pure shear mode strain. This is referred to as the *torsional mode.*

Delay-line operation in the torsional mode has several advantages over operation in the longitudinal mode, especially where long delays are required. The velocity of shear strain waves is only 60% of the velocity of longitudinal waves; thus, more delay is available with the same length of delay line. Torsional-mode waves also suffer less dispersion in the delay wire than longitudinal-mode waves. Whereas longitudinal waves are dispersed by any bending or forming of the delay wire, the torsional mode delay wire is almost free from these effects when certain precautions are taken.

The advantages of torsional propagation are offset by one major difficulty—that of producing torsional-mode stress in the delay wire directly by a transducer. This difficulty is overcome at the input by first generating longitudinal-mode waves and then converting them to torsional-mode waves in a "mode converter." At the output of the line, a reconversion to longitudinal-mode waves is performed for the output transducer.

For a typical torsional-mode delay line, both the input and output mode converters take the form shown in Fig. 3. Thin nickel tapes are welded to the top and bottom of the delay wire. Each tape passes through a separate coil in the transducer proper. A permanent magnet in the transducer establishes a bias flux Φ_B in the tapes surrounded by the coils. When current is applied, the flux from the upper coil aids the bias flux; the flux from the lower coil opposes the bias flux. This causes a compression and an elongation in the upper and lower tapes, respectively. These two opposing strain waves travel down the nickel tapes to the delay wire, causing a torsional stress in the wire. In this manner, longitudinal-mode waves are converted into torsional-mode waves. The latter are in turn reconverted to longitudinal waves at the other end of the delay wire by the output mode converter.

Electrical signals in and out of a torsional-mode line are identical to those found in the simple delay line. The mode converters simply change the nature of the strain pulses on the delay wire.

Peripheral Equipment

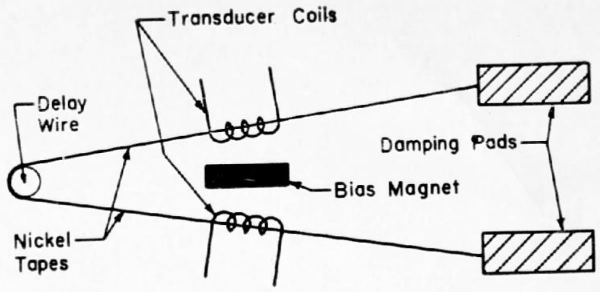

FIG. 3. MODE CONVERTER for developing torsional-mode strain. Longitudinal-mode strain in the tape causes torsion in the delay wire.

The information in a delay-line memory is volatile; it must be continually processed by circulating through a closed delay loop. In addition, as delay lines of this type have from 40- to 60-db attenuation, amplification is necessary. Thus, a certain amount of circuitry must be associated with each line. A sense amplifier, a read/write control, a clocking circuit, and a write driver are required. Fortunately, however, the addressing and buffer storage units can serve a large array of delay lines.

Memory Configuration

The memory configuration depends on the application. A single delay-line gives serial/serial operation, while several lines ganged together provide serial/parallel operation. The two configurations are analogous to single-track or multiple-track tape recorders.

Memory control and addressing schemes also depend on the application. The basic methods used with drum or disc storage are applicable. In most cases the average access time of the line is one-half the delay time. The worst-case access time is equal to the delay time of the line (usually 1 to 5 milliseconds). If very long delay-lines are used and short access-times are desired, short delay-lines can be used as a buffer memory between the main storage unit and the arithmetic unit.

References

1. I. L. Auerbach, J. P. Eckert, R. F. Shaw, and C. B. Sheppard, "Mercury Delay Line Memory Using A Pulse Rate of Several Megacycles," *Proceedings of IRE*, Vol. 37, August 1949, pages 855-861.

2. R. M. Beck, "A High Speed Serial General Purpose Computer Using Magnetostrictive Delay Line Storage," *Proceedings of Eastern Joint Computer Conference*, Vol. 18, December 1960, pages 283-298.

3. G. C. Scarrot and R. Naylor, "Wire-Type Acoustic Delay Lines for Digital Storage," *Proceedings of IRE*, Part B, Supplement 3, 1956, pages 497-508.

4. A. Rothbart and L. Rosenberg, "A Theory of Pulse Transmission Along a Magnetostrictive Delay Line," *Transactions of IRE*, PGUE-6, (Prof. Group on Ultrasonic Eng.), December 1957.

5. A. Rothbart and L. Rosenberg, "Electrical Design of Transducer Networks of a Magnetostrictive Delay Line," *IRE National Convention Record*, Part II, 1958, pages 92-101.

6. T. B. Thompson and J. A. Lyon, "Analyses and Applications of Magnetostriction Delay Lines," *Transactions of IRE*, PGUE-4, August 1956.

7. R. C. Williams, "Theory of Magnetostrictive Delay Lines for Pulse and Continuous Wave Transmission," *Transactions of IRE*, PGUE-7, February 1959.