ing system. To arrive at the actual capacity required in a particular file, the amount of wasted space inherent in the design must also be accounted for. As noted previously, unused file space may result from the deletion of records or from the erection of discrete intervals into which the information does not fit precisely. This latter condition arises when the design includes fixed intervals between entry points, fixed record lengths, or, for that matter, fixed word lengths. The inefficiency arising from these factors is unavoidable to a certain degree, and in any case may be preferable to increased access time or additional housekeeping operations, but must be accounted for with additional capacity nevertheless.

High utilization of storage capacity may also be enhanced by designing the file so that it may expand as necessary. As an example, when a magnetic tape file is completely filled, only a new reel of tape is required for expansion. The cost of additional capacity is then only the cost of the medium itself, i.e., the tape, and the additional processing and access time that is now required. This expandability is a particularly valuable characteristic of a general-purpose data file whose future applications cannot all be foreseen in advance. It is especially useful where this feature can be used to insert new records between already existing ones in a sequentially ordered file without requiring that at least the remainder of the store be rewritten to make room for the new insert. Admittedly, file expandability in this sense would not be practical in many of the high-speed file systems currently conceived; on the other hand, it might be possible to achieve at least some expansion if its advantages are clearly recognized and kept in mind.

Connected with file expansion, the ability to interchange the storage media used in a file system contributes to its flexibility. As the total contents of files get larger, there is usually less need to be able to refer to the entire file at one time. Considerable equipment duplication can be avoided in these cases if the storage medium itself is detachable from the associated hardware.

Nonvolatility

A most important requirement for any business data file is that it be able to retain accurately the information stored in it for long periods of time. In general, this means that the stored information should remain unaltered by lack of electric power, preferably permanently, but at least for periods of weeks or months. The need for essentially permanent machine-accessable storage arises from the relatively slow operation of input-output equipment; at the current input-output speeds, it is not feasible to store the data file on more durable paper records, to be read into the processor whenever required. Permanent storage is, of course, the most desirable; however, storage durations of weeks or months are probably acceptable at the additional penalty of requiring complete file restoration at appropriate intervals of time. In any case, the longer the available safe storage period, the more suitable the medium.

General Requirements

It might be pointed out that the techniques to meet the requirements listed in the previous sections are all presently available individually, but their use together results in files of unacceptably high costs. In a sense then, the problem can be considered as one of searching for new means to accomplish the job at reasonable prices, at the same time maintaining the high degree of reliability necessary for any successful processing operation. In this same sense, lower production costs, increased freedom from environmental changes, lower power requirements, etc., can all be considered as part of the goals of new data file development programs.

Engineering Design of a Magnetic-Disk Random-Access Memory

T. NOYES W. E. DICKINSON

THE International Business Machines magnetic-disk random-access memory is a large-capacity storage device with relatively rapid access to any record. Fig. 1 shows the unit.

The information is stored, magnetically, on 50 rotating disks. These disks are mounted, so as to rotate about a vertical axis, with spacing between disks of 0.3 inch. This spacing permits magnetic heads to be positioned to any of the 100 concentric tracks which are available on each side of each disk. Each of these tracks contains 500 alphanumeric characters. Thus, the total storage capacity is 5,000,000 characters.

The reading and writing is accomplished with two magnetic-recording heads. These heads are mounted in a pair of arms which can be moved in a radial direction to straddle a selected disk. The arms are positioned to the selected disk and then moved into the desired track by means of a feedback-control system. A unique arrangement permits one set of magnetic-powder clutches to provide the drive force for both positioning tasks.

The time to position the heads from one track to another depends upon the distance separating the tracks. The average access time, however, is about 0.5 second.

Disks were chosen for the recording surface because they have a good volumetric efficiency for surface storage. In addition, they permit multiple access possibilities to any record. Magnetic recording was chosen because of both its permanence and its ease of modification while still allowing good storage density.

Disk Array

The 50 disks, which make up the disk array, are the key to the layout and to the size of the memory device. A vertical shaft was chosen as this more readily permits multiple access mecha-

T. NOYES and W. E. DICKINSON are with the International Business Machines Corporation, San Jose, Calif.

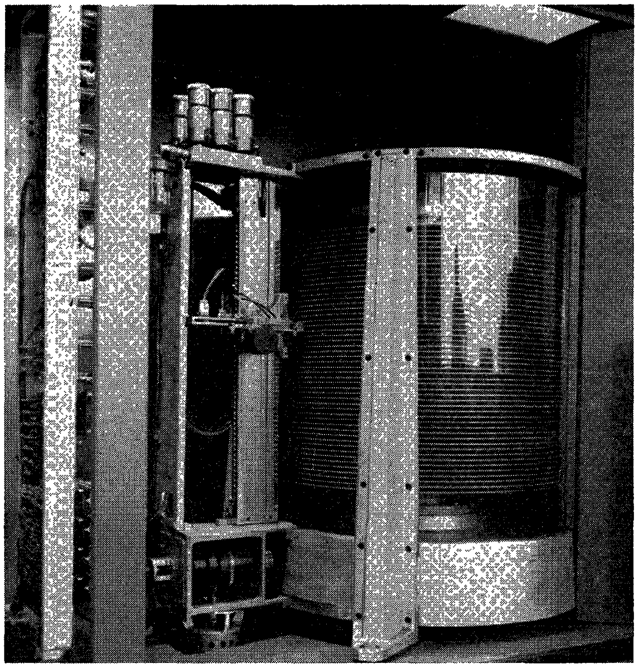

Fig. 1. Random access memory unit

nisms to be used. Theoretically, 20 access mechanisms could be installed. Up to the present time no more than three access mechanisms have been used. Mounting pads are machined on the upper and lower end castings of the array so that the access mechanisms may be bolted in place and be easily removed for periodic servicing.

The disk shaft runs at 1,200 rpm on precision tapered roller bearings mounted on a heavy stationary vertical spindle or axle. The radial runout limitations are less severe than with a magnetic drum because of the read-write design so that 0.001-inch runout is acceptable. The driving power is supplied by a conventional capacitor-start induction-run motor. The large inertia of the disks requires a $1^1/_2$-horsepower motor to accelerate them to running speed in less than a minute but the running power is less than a horsepower. Voltage variations of $\pm$ 10 per cent give less than 1 per cent variation in speed.

The iron-oxide-coated aluminum disk is 24 inches in diameter and 0.1 inch thick. The thickness is required to maintain flatness during assembly and dynamic stability in use. The flatness of the disk is indicated by the requirement that there be no more than 0.0015 inch out-of-flatness in any 2-inch distance on the surface. However, as much as 0.030-inch runout can be tolerated in total axial runout. The uniformity of the magnetic coating is indicated by the tolerance of $\pm$10 per cent on the 60-millivolt peak-to-peak signal at the outside track.

The disk-to-disk spacing is established by the disk thickness and the clearance required for the magnetic heads to go between them. For practical reasons, on this machine, the spacing was set at 0.3 inch space between 0.10-inch-thick disks. Thus, the over-all height of the 50 disks is 20 inches.

A modified non-return-to-zero type of magnetic recording is used. The density of the recording varies inversely with the radius of the track. The density on the inside track is about 100 bits to the inch while on the outside track this drops to about 55 bits to the inch. (The outer-track radius to the inner-track was chosen as near to the optimum ratio of two as was mechanically feasible.) This recording density permits 500 characters composed of eight bits each, to be recorded on each track. With 100 tracks on each of the 100 sides, a total of 5,000,000 characters can be stored. An idea of the quantity of the storage can be conveyed by a few comparisons. This amount of storage is about the same as 60,000 80-column punched cards, or 2,000 feet of magnetic-recording tape recorded at 200 characters to the inch with no space between records or 940 single-spaced typewritten pages. An equivalent storage capacity on a magnetic drum would require a drum 13 inches in diameter and 42 feet long. Such a drum would have about seven times the volume of the disk array.

Magnetic Heads

Closely associated with the disks are the magnetic heads used to record on them. Unlike drum heads, these heads must be of minimum height. Refinements in design and techniques have brought this height down to 0.2 inch. Binocular microscopes must be used in the manufacture of these assemblies which are potted in an epoxy resin after assembly to give durability and shock resistance.

The magnetic element consists of two distinct magnetic circuits. One circuit with its coil erases only, and the other circuit reads and writes. The erase gap erases a wider track than the following write gap records. This design allows reduced precision in radially positioning the heads because there are no magnetically disturbed track edges to contribute noise to the newly recorded track which might not precisely coincide with the track previously written by a different access mechanism.

The spacing of the heads from the disk is maintained by an air bearing obtained from minute air jets in an annular manifold surrounding the magnetic elements. The 0.001-inch spacing is held despite the axial runout in the disk so that there is never physical contact between the heads and the magnetic coating. The head is horizontally constrained in a gimbal socket in the arm.

The use of compressed air in the magnetic heads, and the access positioning detents, requires a small compressor. This unit operates constantly, supplying air to a surge tank or by-passing it to the atmosphere as necessary. Approximately 0.6 cubic feet per minute raised to 50 pounds per square inch is used per access.

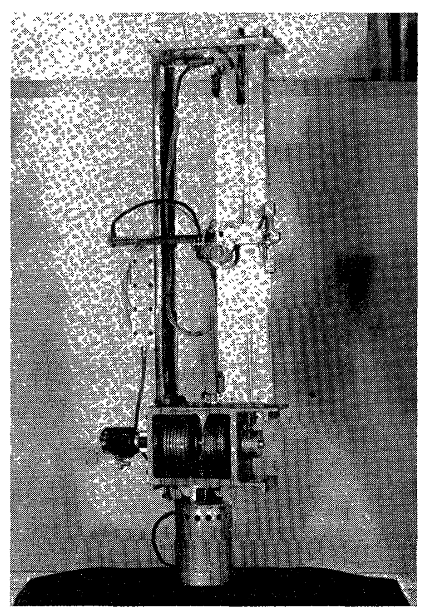

Fig. 2. Access mechanism assembly

Access Mechanism

The access mechanism shown in Fig. 2 is used to position the pair of heads to any track on the disk array. The heads face each other in a pair of arms. The arms move inward to straddle a selected disk when reading or writing. The arms, in turn, are carried on carriage for vertical motion to the desired disk.

The arms are guided, for radial motion, in bearings on the carriage. Within these bearings the arms are capable of about 6 inches of radial motion. The inner 5 inches position the heads over the disk recording area. When the arms are in their outermost position, the arms are completely outside the disks. This is the position the arms are in during vertical-drive motion.

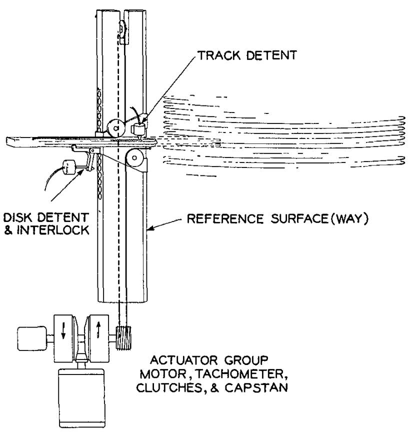

Fig. 3. Mechanical functional schematic

The carriage, during vertical motion, slides on a vertical "way." See Fig. 3, a functional schematic. At each of the 50 disk positions a detent hole is provided in the way. A pneumatic-detent piston is energized upon arrival at the desired disk. This detent, by means of a mechanical linkage, controls an interlock which frees the carriage and locks the arms for vertical drive and frees the arms and locks the carriage for radial drive. The arms are capable of being freed only when the carriage is positioned properly at a disk and the carriage is capable of being freed only when the arms are completely outside of the disks. Thus, a safe interlock is provided to prevent mechanical damage to the disk array.

The driving force is provided by a pair of magnetic-powder, motor-driven counter-rotating clutches. These clutches have a common output shaft on which is located a drive capstan. A small, steel cable connects the drive capstan to the arms through a system of three pulleys. When the arms are locked, the carriage is free and the clutch torques result in vertical-drive motion. Similarly, when the carriage is detented the arms are free and the same clutches control radial motion. In addition to the detent for locking the carriage, a detent is provided to position the arms accurately to the selected track. These detents greatly relieve the positioning requirements of the position-feedback controller.

The clutches are controlled by a null-seeking feedback-control system. Position signals for the radial and the vertical drives are obtained from potentiometers on the carriage and on the way, respectively. Fig. 4 shows a functional schematic which is intended to illustrate either radial or vertical positioning. An electrically floating voltage supply feeds the potentiometer element. Taps are located uniformly along the potentiometer element. The desired address is established by grounding one of these taps with an address-relay tree. An error voltage is seen by the potentiometer wiper unless the wiper is positioned at the selected tap. Through the control amplifier the clutches are controlled to position the wiper so as to make the error voltage zero. The tachometer is used to stabilize the feedback loop. A d-c control system is used with relatively large voltage on the potentiometer to eliminate any serious drift problem.

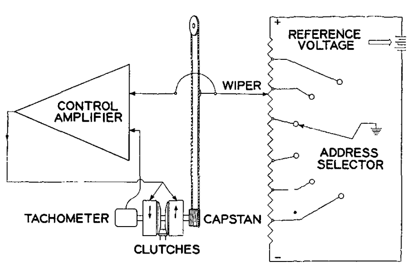

Fig. 4. Electrical functional schematic

A group of relays are used in a logic network to guide the access mechanism in positioning. The logical decisions are based primarily upon the condition of the carriage detent, whether the carriage is at the correct disk and whether the arms are at the correct track.

Read and Write Amplifier

The read and write amplifier circuitry is quite conventional. The only variation from normal is the inclusion of an automatic-gain control in the read amplifier. This feature is included to correct for input variations caused by variations in heads, disks, and surface speeds.

Self-Clocking System

To relieve the accuracy required in positioning the head along the track, a self-clocking system has been incorporated. The tolerance with this system is such that limited movement of the head along the track can occur while reading or writing. Reading or writing with different access mechanisms is also facilitated by this system.

The clocking system is composed of two oscillators. Their operation is such that one oscillator is on while the other is off. In writing, one oscillator runs continuously for timing the entire record. In reading, the oscillators are controlled by the recorded bits. Each bit read switches one oscillator off and the other on, thus resynchronizing on each bit. When an oscillator is turned on it always comes on with a certain phase relationship. The combined output of the two oscillators feeds the bit ring used to determine bit position within the character.

In a system such as this the frequency of the oscillators relative to the disk speed is important. With the present arrangement a tolerance in excess of 1 per cent is permissible. This tolerance is easily met with the large inertia of the disk array, the motor voltage-speed characteristics, and the oscillator stability.

Print I—A Proposed System for the IBM Type 705

R. W. BEMER

Purpose of the Print I System

THE PRINT I (PRe-edited INTerpretive) system has been designed to meet the engineering and scientific computing needs of those Type *705* installations where such work is a secondary computing requirement. It is specifically tailored for this purpose and in general has no similarities to Type *705* systems designed for business and commercial applications. Although provision is made to move freely from abstraction to *705* commands and back, it is not recommended that PRINT I be tied to any business system for the reason that the restrictions of one type of usage will too often hamper the other.

The basic consideration in the planning of this system has been ease of learning and continued operation by personnel with either very modest programming experience or none at all. Pseudo-instructions are simple and straightforward; basic principles may be assimilated quickly. There are no restrictions in the use or combination of pseudo-instructions, although minor increases in operating speeds have been introduced to effect this. Basic logical errors in the written program will be detected automatically and detailed analysis of such errors will be typed out to the operator.

Choice of System Characteristics

Components

Since some Type *705* installations have ordered configurations which do not include a magnetic drum, the only components specified for this system are magnetic core memory and sufficient magnetic tape units to handle expected problem size.

Abstraction Type

The first decision to be made was between the compiling and interpretive modes. The fastest method of computer operation is with a compilation of basic machine commands operating in linear progression without modification. The first drawback to this method is that there are no computers in production with the nearly infinite memory required for this long thin line of instructions. The second is that the time advantage gained is not worth enough to compensate for cost of additional memory if it were available. A more practical variation of this is the type of compiler which still forms machine commands but uses modifying commands to reduce the volume of instructions, just as a conscientious programmer would. The serious drawback to this method is that it requires that extreme complexity be built into the compiler in order for it to function as efficiently as the clever programmer. In addition, the program produced is often much larger than that produced by a compact interpretive system.

An interpretive mode offers quick construction and minimum storage, obtained by the complicated weaving and interlocking that the coder may accomplish. Its drawback, an extremely serious one, is that the command is normally re-interpreted every time it comes up for execution in actual operation. Thus the variable "fetch" and "operate" commands are fabricated a multiplicity of times, whereas a compiler fabricates them only once.

This has led to rather determined stands by the opposing compiling and interpreting adherents. Fortunately there is a compromise available, as demonstrated in the PRINT I system. A pre-editing routine performs a compiling and assembly function to make the commands numerically palatable to the executive routine which is in the interpretive mode. In addition, a repeat command is furnished in the list of pseudo-instructions. This enables certain commands to be executed n times in succession, once with interpretation and command fabrication, and $n-1$ times in the exact fashion that an ultra-efficient compiler would generate the program. Multiaddress commands are used, for although single-address commands are more efficient when the desired characteristics are inherent in the machine language command, multiaddress commands are more efficient in an interpretive abstraction, just as buying more groceries at a time minimizes the cost of trips to the store.

Arithmetic and Format

Floating decimal arithmetic was chosen as most generally acceptable to all scientific users. Since the specification of a mantissa length must, to our way of thinking, specify a corresponding set of subroutines to that accuracy, PRINT I will be packaged initially in 8- and 10-digit versions. A 20-digit version is to be produced after completion of the 8- and 10-digit systems. If enough demand should exist, a 5-digit version may well be produced. Each system will be complete in itself for all operation. There will be complete freedom of interchange between the floating point abstraction and *705* language.

For faster internal operation and convenience, floating point numbers are stored internally as $xxx.\;.\;.\;.\;.x\overset{\pm}{x}P\overset{\pm}{P}$. The x's represent a mantissa which is a proper decimal fraction with a non-zero leading digit. PP is the power of 10 multiplying this fractional number. Externally, the format is $\pm P\overset{\pm}{P} \pm xx\overset{\pm}{x}\ldots$, so that trailing zeros need not be written by the programmer; the pre-edit routine will take care of this automatically.

As far as possible, all arithmetic operations and subroutines will be performed with rounding. A legitimate zero in this system has both power and mantissa equal to zero. To facilitate and accelerate operation with zero operands, the mantissas will be tested for zero in all arithmetic operations. Messages are provided to be typed out in case of error during operation, such as:

Division by zero.
Square root of a negative number.
Power overflow (exceeding +99).
Logarithm of zero or a negative number.
Sine or cosine of angle greater than a prescribed limit.

This list of error messages will be completed as the various elements of the system are finished.

All addressing of operands and locations of pseudo-instructions for this system are in regional form, consisting of one leading alphabetic character and, normally, three numeric characters. Inserts may be made by an additional numeric character on the right, as *G*125 and *G*1251.

Tracing and Diagnostic Routine

There will be a single type of diagnostic routine associated with this system, since the programmer will want to see

R. W. BEMER is with the International Business Machines Corporation, New York, N. Y.