utilization of modern digital computers can be realized. The design of this new tape handler, and future extensions of this design, will materially aid in the attainment of the goal sought.

Appendix

For purposes of analysis a unidirectional system is shown in Fig. 5. The contacts have been replaced by a continuously variable resistor, the static tape tension, friction, and the small component inertias have been neglected.

The following forces are acting on the control resistor slider:

$$K_1X = K_2X_2 + L\frac{dx_2}{dt} \quad (1)$$

$$\Theta_m = \left(\frac{K_2}{K_1+K_2}\right)\epsilon + \left(\frac{L}{K_1+K_2}\right)\frac{d\epsilon}{dt} \quad (2)$$

$$\epsilon - Kr(\Theta_i - \Theta_0) \text{ and } \frac{d\epsilon}{dt} = Kr(W_i - W_0) \quad (3)$$

Substituting equations 3 into equation 2

$$\Theta_m = K(\Theta_i - \Theta_0) + K(W_i - W_0) \quad (4)$$

From equation 4, it is apparent that the motor control angle, and hence motor torque is proportional to both input displacement and rate of change of displacement.

Finally, the dynamic forces acting on the motor and tape system are as follows:

$$J_0\frac{d^2\Theta_0}{dt^2} = K\Theta_m = K\left[(\Theta_i - \Theta_0)K_2 + \left(\frac{d\Theta_i}{dt} - \frac{d\Theta_0}{dt}\right)L\right] \quad (5)$$

From

$$\Theta_0 = \Theta_i - \epsilon$$

$$J_0\frac{d^2\Theta_i}{dt^2} = J_0\frac{d^2\epsilon}{dt^2} + L\frac{d\epsilon}{dt} + K\epsilon \quad (6)$$

Equation 6 is typical of a "closed loop" system with "error rate" damping only. In terms of the error E, the solution gives:

$$\epsilon = \frac{wi}{\sqrt{\frac{K}{J}}} \text{SIN}\left(\sqrt{\frac{K}{J_0}}\right)t \quad (7)$$

Or in the critically damped case:

$$\epsilon = w_i t e^{-\frac{L}{2J_0}t} \quad (8)$$

The actual system response may be made to approximate closely either of these cases by adjusting the dashpot damping coefficient L.

Nomenclature

K_1K_2 = centering spring constant
J_0 = Reel and motor inertia
L = Dashpot damping coefficient
$\theta_i\theta_0$ = Input and output tape displacements as referenced to reel radius
ϵ = Take-up arm displacement
θ_m = Motor control resistor angle

Requirements for a Rapid Access Data File

GEORGE EISLER

GENERAL-purpose business-data-processing machines now on the market are limited in their performance mostly by the electronic file systems associated with them. The only file medium that has proved itself acceptable from the standpoint of both cost and speed is magnetic tape. In the quest for better, i.e., faster, cheaper, etc., general-purpose machines, a "random access" file has been held out as perhaps the key to the next step of progress. "Random access," however, is a misleading term if it is used to describe what an ideal file should be, rather than adopted. The considerations which, taken together, point to the requirements for a rapid access data file in a general-purpose data processor will be presented here.

The desirable factors having the greatest influence on system utility are listed as follows:

Speed

High operating speed when large groups of data are to be processed at one time.

Adequate look-up speed when single records are required.

Addressing

Items locatable by the same identifications as in a manual file.

Items locatable by various categories other than normal identifications.

Capacity

Adequate capacity, probably implying both expandability and high utilization of file space.

Interchangeable storage medium.

Nonvolatility

Ability to retain information in the absence of electric power for periods of weeks or months.

Furthermore, these desirable characteristics must be implemented by devices of reasonable cost which, in addition, are within the over-all processor system requirements of reliability, size, weight, and other factors.

The implications of the features listed are considered in detail in the following section; the order, however, in which they are taken up does not necessarily reflect their relative importance since most of them are fairly interrelated.

Speed

One of the most important considerations bearing on the performance of a data-processing system is the speed with which records can be made available by the file for the processing operations.

Data-processing operations calling for some modification of the file can be carried out in two general ways, depending on the particular application. If the time lag is permissible, data for a number of similar operations may be collected and processed at one time as a group; or, if the time lag is not permissible for some reason, individual record modification may be carried out as soon as a new piece of information arrives and the machine is available. The choice of the procedure actually used, however, not only depends on the demands of the processing situation but the characteristics of the file as well.

By way of illustration, the processing of a group of checks which a bank receives from a clearing house may be performed as a single continuous operation at a convenient time in the scheduled daily program. It so happens that with a magnetic tape file in the processor, this is preferable, since it is far faster to process data in a group when both the new and the stored information are in a matched sequence and the tape need be scanned only once. There is no reason, of course, why this operation could not be done by treating the checks separately and in no particular order, if there were a file where any record

GEORGE EISLER is with The National Cash Register Company, Hawthorne, Calif.

might be retrieved in a sufficiently short time. The choice in this case is based on file characteristics.

In the same bank, however, the balance in depositors' accounts must be made available to tellers about to cash large checks over the counter. There is no possibility here of group operation; each inquiry must be answered as it arrives. The example is fairly representative of most applications; that is, the bulk of the processing may be handled in groups with some delay, but a certain amount of immediate action is often necessary.

The more immediate problem in data processor design is to provide reasonable access speed to randomly chosen records. Improvements in the speed of group processing, while certainly to be desired, are not so pressing. On the other hand, the access time requirement alone is relatively modest in the cases where individual record selection is a must.

To illustrate, in a particular bank during an 8-hour workday, there might be something like 1,000 inquiries from both tellers and customers requesting the balance in particular accounts. If further, rather arbitrarily, it is desirable not to devote more than 10 per cent of the processor's time during working hours to the task of retrieving records to answer individual inquiries, and no time sharing is possible, the time allowed within the processor for each inquiry is $\frac{8\times 3{,}600\times 0.1}{1{,}000}$ $=2.88$ seconds. While this is a long time in electronic terms, it must be contrasted with the several minutes that is required to scan a single reel of file tape as specified in most electronic data-processing systems. These figures indicate that the current magnetic-tape files cannot be used directly to answer random record requests in the situation cited, and that, in an installation requiring such operation, either a specialized mechanical arrangement must be made or the information summarized in anticipation of the requests.

On the other hand, the performance of magnetic tape file systems while operating on groups of records is another matter. Considering a representative file, for instance, an average of about 30 records a second will be provided to the processor when alternate records are examined in the processing run (records of 15 words' length are assumed here). This is roughly equivalent to an access time of 35 milliseconds per record; when one out of four records are examined, the access time on this basis is 55 milliseconds. Even at a one-out-of-ten processing density, the equivalent access time is 115 milliseconds.

For a typical application of a general-purpose data processor, the situation is roughly as in the foregoing illustration. The speed of file operation for group processing of data is much more important than the requirements for individual record retrieval simply because most of the processing can be done on a group basis, and for the present this is the most efficient way. This will continue to be the case until files are produced which will provide processing speeds at least approximating those achieved by magnetic tapes, as well as providing individual record retrieval at the somewhat more modest speeds called for.

A discussion of file access time must also include consideration of the sorting question. The group processing speed of the magnetic tape file is based on both the file and modifying data being in the same sequence, so that the tape may be scanned in a single, unidirectional pass. Depending on circumstances and the machine, it has been estimated that about 100 milliseconds per record are required for sorting the new information into sequence before the actual file modification. On the other side of the ledger, however, many business procedures require data in alphabetical or numerical order so that very often storage in sequence is desired apart from file operating needs. In particular where a sorted output such as a stock list is required many times from, essentially, the same file data, sorting at each readout is a wasteful procedure; in these cases, an ordered file storage would be much more efficient. This being the case, it seems that the sorting time should not be considered chargeable only to a magnetic tape file, or for that matter, to any pure scanning-type system. This, especially, since those files which use at least a partial scanning scheme for record location also benefit from ordered record storage.

In assessing the prospects of faster access to records, consider first a scanning-type file. There appear to be two possibilities for increased speed:

1. Scan faster.
2. Scan only the parts of the records which are significant for identification.

However, along with the second possibility, the ability to scan any selected segment of records is desirable in order to maintain the file operating efficiency where, for instance, all items belonging to a given category must be selected.

There are, however, very definite limits to the speed with which scanning can be accomplished. Assuming, strictly for the sake of illustration, that the file contains 10^8 binary digits and that a maximum of one second is allowed for each randomly chosen access, the binary digit rate must be 100 megacycles when the worst case is considered. By way of comparison, digital circuitry is made to work only with great difficulty and expense even at a 10-megacycle rate, so that a completely sequentially scanned file cannot be expected to function very fast for individual location of items unless new techniques in electronic circuitry are developed concurrently.

Scanning only the address portions of records is more promising, but still not clearly feasible since the assumption of file capacity was not realistic but on the low side.

In a discrete (or multiple) entry type of file, the access time to a particular record is reduced by entering the file directly at, or at least near, the desired record's location. The additional requirement for doing this is the knowledge of the record's location in file, so that the proper entry point can be chosen. The greater the number of entry points, the faster the access, but also greater is the amount of information required to choose the proper entry point in the first place.

Addressing

The procedure used in finding records in the data file of a business processor also has great influence on the ultimate performance of the system. The generally desirable features are next considered. These are:

1. A great part of the data handled in the course of business is now filed by classification according to names, stock symbols, nonconsecutive account numbers, etc., and it is greatly desirable that the processor data file be designed to locate information by using these commonly accepted appellations.
2. The amount of information to be filed under a particular identification symbol (often called the "natural name") varies greatly, as, for example, the number of transactions in a particular bank account in a given month; the file should be able, therefore, to handle efficiently records of widely varying lengths.
3. In addition to the ability to select records by natural name, there should also be provision for the selection of file data by types or categories, as would be useful, for instance, in searching the accounts receivable file for overdue accounts.

To make available the desirable features listed, suppose that the approach taken in the main store was utilized and a discrete space in the data file allotted to each possible record identification, i.e., to each possible combination of letters in names, each possible permutation of a stock number, etc. The natural name would then be used directly by the processor as a file address. The result would be a requirement for a file of huge capacity, of which only a relatively small part would be used at any one time, since there are a great number of combinations of letters which are not used by anyone for names; and of those that are, not all are customers of a particular company. About the same situation arises with regard to stock numbers, car registration numbers, and other commonly used identifications, most of which have large amounts of redundancy inherent in them. From the machine design standpoint, at least the file could of course be greatly simplified if business concerns would adopt consecutive numbering systems in their operation. The difficulties encountered in changing to consecutive identification schemes undoubtedly varies greatly from one firm to another. Its advantages for machine processing should, however, be weighed for each application, so that in the future more efficient use of file storage space may be made. In the meanwhile, general-purpose machine design based on the continued use of redundant natural names seems necessary.

In order then to reduce the amount of wasted space in the data file, it is necessary to be able to assign empty locations to records as the particular situation demands. That is, the available file spaces should be filled by the automatic action of the processor as the need arises. The two approaches to this are: (*a*) to store records, in or out of ordered sequence, without reference to location and then scan the file to recover a particular one, or (*b*) to assign a specific file space to the given record, remember this assignment, and use it to locate the record when desired.

Considering the scanning system first, it is noted that the records must contain as part of the entry their own identification symbols, and that a specific record is found by examining, one after another, all of the records on file. However, as previously discussed, with present-day mechanizations, this type of file tends to be slow, particularly when only a single record is to be found. In a group processing operation, the speed is improved by a large margin, but still remains the factor limiting over-all operating performance.

A scanning system handles records of variable size naturally, there being no particular need to have the records be of predetermined, or of equal, length. On the other hand, the closing out of records necessitates the eventual rewriting of the file in a densifying operation, in order to preserve a high space utilization efficiency. Whether rewriting is necessary when new records are to be added depends on whether the records are filed in sequence or not; but, since it is mostly likely that sequential ordering would be used in any completely scanned file, complete file rewriting must be considered a part of the scanning-type file operations. This, of course, is the description of a magnetic tape file. To find ways to overcome some of its shortcomings in the first place, it is best to consider some alternatives.

Selective scanning, that is, scanning only portions of records, as a means of speeding up file access is atractive; the disadvantage here is that most selective scan procedures require that information words be spaced at certain regular intervals so that, for example, every tenth word may be examined for content. With records of varying lengths, this results in wasted space, which can be minimized by the proper choice of intervals, but not completely eliminated. Selection by category is also made correspondingly harder.

Considering a pure discrete entry system next, it is seen that while the access time to a record might be very small the choice of entry point implies prior knowledge of the record location. This prior knowledge can only be stored in a additional memory system if, at the same time, the use of natural names and a reasonable storage utilization is to be retained.

This second memory system, usually referred to as an index, is by no means a small auxiliary. It has the job of picking out a file location (out of possibly 10^5 choices) that contains the particular record having the input natural name. This input natural name itself can be one configuration out of an astronomically large number of possibilities representing the permutations of possibly ten alphanumeric characters. Making provisions for record selection by category complicates the index even further. The magnitude of this indexing operation is such as probably to make the pure discrete entry file impractical.

Between the pure scanning system, which is found to be slow, and the pure discrete entry, which requires too much equipment, might lie a combination system which would provide a suitable compromise. In general, this type of file would be scanned for actual record location, but the scanning would begin at any one of a number of entry points. Minimum average access time would be obtained when approximately equal amounts of information are stored between access points, so that the amount of scanning would be minimized. The indexing mechanism in this arrangement would choose an entry point for a given natural name instead of a specific location, thereby reducing the number of possibilities from which the choice has to be made. Within this framework, a number of alternatives are possible, allowing varying amounts of flexibility to the file system.

One possibility is to divide the entire pertinent range of natural names into intervals corresponding to intervals between file entry points. The actual entry point selection would be done by a matrix which would use little time for operation but would require resoldering, or at best, reconnecting on a plug board, for a change in the interval. An alternate approach would allow the computer to rearrange the intervals, changing the indexing unit appropriately, thus eliminating the waste of memory space in the fixed interval approach because it is necessary to allow spare memory capacity between intervals to account for uneven interval loading. Additional flexibility, however, requires that more and lengthier file housekeeping steps be performed.

It is difficult to classify and examine systematically the various modes of indexing possible; it can be seen, however, that the index, active or passive, is essentially a storage device, and is subject to much the same organizational problems as the data file proper. It is likely then that a successful mechanization will also represent some compromise within the cost-capacity-speed triangle.

Capacity

A major distinguishing feature of business problems is the vast amount of file data that must be handled. A further distinguishing feature seems to be that the file data are growing vaster all the time. Inquiry into the needs of potential users of business data processors indicates that a capacity of 10^8 characters seems to suit the majority of applications. The value of 10^8 characters is then chosen as a reference figure, with the understanding that it is certain that some applications will require a much larger capacity. This additional capacity might be supplied by duplicate equipments, or, if the design and processing flow permit, by additional storage media which are detachable from the file mechanism.

The quoted reference figure, however, represents only the number of characters of information to be filed by the processing 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.