Mass storage revisited

by ALBERT S. HOAGLAND
IBM Watson Research Center
Yorktown Heights, New York

INTRODUCTION

Mass Storage as a functional need in computer systems is continually increasing in importance with the growing trend to interactive terminal-oriented systems, serving as peripheral or external on-line memory for storing a systems data base and resident programming systems. The associated capacity, plus the ever expanding magnitude of such data, far exceeds the range where "electronic" memory is economically competitive. Included in the product category defined as mass storage are drum, disk, tape, card, strip, and chip recording structures. Direct access storage is becoming a standard feature of computer systems, with much the same type of distinctiveness as the CPU and main memory have achieved.

Five years ago I made a "tour d'horizon" of mass storage for the 50th Anniversary Issue of the Proceedings of the IRE.[1] It is of interest to preface this review by certain observations on the situation as seen then contrasted with now.

I felt then and still feel that magnetic recording will remain the technological base for mass storage for the foreseeable future. While this prediction has held true, the amount of exploratory activity in so-called beam addressable storage (optical and electron beam recording) has greatly increased. This current review (circa 1967) reflects a large relative increase in the material devoted to new technological approaches to mass storage.

Five years ago the highest storage density in a commercial file was 25,000 bits per square inch, while today this figure is approximately 250,000 bits per square inch—an order of magnitude increase (giving two orders of magnitude improvement over the last decade). However, at that time I felt eventually magnetic recording storage densities in excess of 10^6 bits per square inch would be realized and will restate the same conviction, reinforced by the advances already achieved.

An unanticipated factor which emerged to play a major role in the accelerating importance of development activity in mass storage is the challenge posed by programming systems residence. The complexity and size of operating systems is a fact of life that was little appreciated in 1962. The storage needs here serve to emphasize fixed head files with their capacity-access trade-off favoring short access time.

The success of the replaceable disk pack file has been remarkable and placed this device in such a predominant position that the long range role of strip-type storage structures has not emerged nearly as clearly as was expected.

In the last decade significant advances have also been made in tape drive performance as well as in the elegance of their design for reliability and serviceability. However, we are concerned here with the "image" of mass storage and these product advances do not essentially change the perspective of tape devices in an overall sense.

Overall, in the last five years, technology appears to have advanced as fast as suggested while the systems organization and use of mass storage hierarchies retains almost as much fertile ground for sophisticated design and application as appeared then.

Functional considerations

Computers that manipulate data are primarily limited by the number and size of files that can be made readily accessible for processing. Further, mass storage devices are now called on to fulfill important systems functions. Random (or direct) access units are used in compiling and assembling programs where their ability to reach large directories and subroutines rapidly is necessary for responsive program-preparation and execution. For time sharing, programs and data of many users can be stored on-line with a mass storage, to be run as requested in accord with some "optimum" allocation of facilities.

Capacity, access time, "latency," data transfer rate, and cost per bit are the basic performance characteristics of mass storage devices. Each mass storage

From the collection of the Computer History Museum (www.computerhistory.org)

unit has its own particular attributes, and many applications require that several different devices, or a hierarchy of devices, be connected within the same computer system. The short access times and high data rates of fixed head drums or disks save valuable processor and internal memory time when it is necessary to swap programs, as in time-sharing. However, the higher cost per bit of these devices generally makes them too expensive for file storage, and other structures of higher capacity (and longer access time) but lower cost per bit are used. Cost of storage (in superficial terms) in inversely related to the number of bits stored per independent read/write transducer.

Systems-derived performance factors such as throughput depend not only on the mass storage specifications but also on indexing procedures, memory allocation and chaining provisions, file activity (ratio of records actually processed to total), provisions for queuing and ordering of access requests, checking techniques, etc. Thus, the associated control logic is an integral facet of any mass storage subsystem.

Technological considerations

The continuing need for high capacity storage has required the intensive exploitation of *recording* technologies for the economic implementation of mass storage. Thus, storage units involve the physical integration of recording media, transducers, precision mechanics, servo systems, and electronic encoding and decoding techniques to achieve a meaningful set of capacity access time tradeoffs.

Access to any data location is provided by relative motion between the storage surface and an associated transducer able to record signals on and sense the state of the storage medium. A single transducer may service many data "tracks" by a positioning mechanism operating normal to the direction of scanning. The recording density (bits per square inch) is principally a function of the registration tolerances (three-dimensional) that can be realized between the storage film and the coupling transducer. The access time variability to memory locations arising from the requisite motion necessary to scan large areas, makes the data organization of a mass store a key factor to effective systems utilization.

The greater the bit storage density (and hence the number of bits associated with a given read/write transducer) the lower the cost/bit and, correspondingly, the shorter the average access time for any given capacity. The average random-access time will range from milliseconds to seconds because mechanical motion is required for accessing masses of data, i.e., the larger the memory capacity, the greater the required surface area that must be accessible to a read/write station.

Although it is only in recent years that magnetic recording has come into wide general use, its invention by the Danish engineer, Valdemar Poulsen, dates back to 1898. The paramount functional advantage of magnetic recording surfaces is their unlimited reusability. This property permits the direct modification of stored information. Additional advantages of magnetic recording for mass storage of data over other potential storage film media are: the simplicity of recording transducer (a magnetic head); the flexibility in mechanical structure possible (and hence, choice of performance specifications) due to the ability to place the storage film on almost any supporting surface in conjunction with ease of mounting a magnetic head; the high bit storage densities and read-write transfer rates obtainable with magnetic recording; and great ruggedness with respect to handling and environmental conditions. The further features of replaceability and off-line shelf storage (e.g., tape reels and disk packs) make this mass storage technology extremely attractive and very economical.

Magnetic recording represents the integration of several basic engineering fields and has been generally characterized by rapid progress achieved by evolutionary advances rather than dramatic innovations. The one "breakthrough" that can be identified with the computer field is the "air-floated" head. Otherwise, advances in the magnetic recording art have largely emanated from increasingly higher precision and quality in components.

Historical growth and present status

The original work which ushered in mass data storage was firmly under way by 1947. This activity was associated and concurrent with the explosive "take-off" of the digital computer field at that time. Early work was oriented to the needs of the scientific computer market. The principal mass memory device was the magnetic tape unit, to provide both an auxiliary "back-up" storage for main memory and terminal buffering (data rate "matching" between input/output equipment and the central processor) in large-scale scientific systems. The later emergence of commercial data processing brought with it a wider variety of functional usages and mass storage hardware. Magnetic drum memory development for small and medium speed processors served to significantly add to the technological base of digital magnetic recording.

Commercial or business data processing, as it was evolving as a main facet of activity in the electronic

From the collection of the Computer History Museum (www.computerhistory.org)

computer field in the early 1950's, gave a tremendous impetus to mass storage development and had a major impact on its direction. File storage for records maintenance was the central requirement.

Magnetic tape was first exploited for mass records storage. The only practical way to use tape is to address by record content. In updating a file, for example, the master and transaction tape reels are serially read (information is arranged and maintained in ordered sequence) and a "new" tape is created, on another transport, with the unmodified as well as the altered records being transferred. It was easy to make insertions and deletions in this process as well as to handle variable length records, as physical sections of tape have no specific identify; although to modify or insert a single record, the entire tape must be rewritten. A tape reel is relatively cheap and therefore low-cost, off-line, archival storage is attractive. For low file activity, tape devices are very inefficient, due to the constraint of sequential access. Further, effective file inquiry operations are not possible.

The character of much business data processing indicated the need for an entirely different type of mass storage. The desirability of storing large volumes of information with any record available rapidly gave stimulus to the development of a mass random access memory.

The air bearing supported head (using an air cushion to control head-to-surface spacing) was the innovation which, associated with the above memory concept, brought about this entirely new type of mass storage. An air-bearing head can follow considerable surface fluctuation—up to 100 times the spacing. Since the readback amplitude wavelength dependence on separation is given by $e^{-2\pi d/\lambda}$ (where d = separation and λ = wavelength) high recording densities would be impractical without such a method of maintaining close and accurate spacing. By this air-bearing spacing technique, it was possible to develop a high-capacity rotating disk array since a head could closely follow the appreciable runout of large disks.

The first version (the RAMAC, announced in 1956) could store five million characters with an access time to any record of less than a second, having one head mechanism servicing the entire disk array. Secondary technical features of note were the use of self-clocking and a wide erase narrow read-write head unit. These design approaches, combined with the use of an air-supported head, provided techniques that compensate for the head-to-track registration tolerances of such a gross mechanical structure, and thus permitted the high track density and high bit density necessary for large capacity.

Random access memory involves addressing by physical location to a single record or a particular block of records (one track), which is then scanned. Any record can be read, written or modified without affecting any other record. The "set of keys" (record identifiers) of a file will, in general, bear no direct relation to the closed set of machine addresses. Various randomizing techniques (key transformations) are used to convert scattered keys covering an extensive range to a dense and relatively uniform distribution of numbers to obtain automatic addressing capabilities.

Initially pressurized air was fed into the spacing gap to maintain separation. Around 1960 a significant advance was achieved as self-lubricating air bearings came into general use on both disks and drums, bringing great simplicity and cost advantages, and in particular making head-per-surface disk arrays feasible.

Many approaches to chip, or strip, or card type mass storage devices have been undertaken over the last decade and NCR has been the leader with their CRAM.[2,3] However, the simplicity of rotational motion over card shuffling has favored disks, and no strip-type storage structure has emerged as a clearly identifiable industry symbol. Indications are that such devices must be complementary to disk storage and it is still premature to detail such configurations in the context of their evolutionary significance.

The next major innovation in mass storage (1962-63) was the replaceable disk pack concept, which emerged as a practical alternative with the rapid progress achieved in storage density—from 2000 bits per square inch (1956) to 50,000 bits per square inch (1962), permitting adequate capacity to be stored on a few small disks.[4]

Three discernible themes have emerged in the area of mass storage.

1. The replaceable disk pack file is clearly appearing as the principal mass storage structure characterizing third generating computer systems. The disk pack file appears well on its way to becoming as ubiquitous as were tape drives on the computer systems of the late 1950's. With the major trend toward high performance disk files, there is concurrently much less forward momentum in magnetic "strip" storage as a competitive alternative, contrary to some earlier projections.

2. A second aspect is the vastly increased discussion and presentation of advanced work in new technologies that may become significant for computer storage.

3. Fixed head files, either drum or disk, have entered a stage of greatly renewed interest due to the demands now being generated for high capacity storage with fast access to store extremely large and

From the collection of the Computer History Museum (www.computerhistory.org)

sophisticated programming systems as well as meet the data and program swapping needs of time-sharing.

The disk pack cluster demonstrated another answer to very large on-line storage and could become a principal approach to serving requirements generally met today by large disk array files. For example, the IBM 2314 consists of eight disk packs (plus one spare) each with a completely independent access mechanism, all in one assembly. With this concept, by overlapping accesses, the mean access time can go down with increasing capacity at a relatively fixed cost per bit. The performance region of several hundred million bytes of data with access in milliseconds is a key to many on-line "data bank" type information systems and the appeal of such facilities spurs the rapid trend to direct access file-oriented systems.

The "state of the art" magnetic recording density for production disk files (1966) moved to the region of 250,000 bits per square inch (2500 bpi by 100 tpi), a factor of 125 greater than the density of the first commercial disk file announced only a decade ago. Further, major advances in density and thereby capacity are clearly indicated based on many current laboratory investigations in the range of 10,000 bits per inch. Techniques that may make equivalent strides in reducing access time are still undefined. The trend to large on-line data banks will undoubtedly accentuate reliability considerations far beyond traditional experience—a question with electromechanical devices of constantly growing concern.

New beam addressable concepts for storage are also being described periodically that may be promising for the future. Numerous groups revealed activities looking towards developing reversible magneto-optic storage techniques. Writing is accomplished by thermal heating (e.g., with a laser beam) in the presence of a magnetic field and reading by sensing the rotation (caused by the magnetic state of the film) of a polarized beam of light.

Among other mass storage approaches recently presented were: a trillion bit storage system for the Livermore AEC Laboratory based on electron beam recording on photographic film—with a flying spot scanner for reading; an electron beam multiaperture recording structure; several concepts using holography as a means for very high density optical data storage, a recording scheme which stores information in a plastic film with a laser by "drilling" holes (whose presence or absence can be optically sensed), as well as variants on these methods. These latter schemes do not allow direct updating of recorded information.

Most of these development activities are in a very preliminary stage and one would not anticipate commercial products that capitalize on such technologies to appear for several years. Nevertheless, these releases do dramatize a recognition of the growing importance of achieving higher performance storage products to meet the growing system demands posed by the trend towards on-line multiuser systems oriented to remote terminals.

Future trends

Each user wants all data immediately available to him. The appetite for more mass storage capacity with high speed access appears as insatiable as the appetite for more computer speed.

A major problem in file-centered systems is determining the proper balance of mass storage devices. A hierarchy of storage modules (of varied performance parameters) closely integrated can best optimize overall systems performance, recognizing that cost trends to be inversely related to access time for a given capacity. We need to better understand the flow of data and develop design guidelines here, recognizing that the access "gap" from electronic memory through electromechanical storage devices will cover a relative access speed differential of 10^3 to 10^5. Mass storage control logic is now beginning to extend to the functional ability to automatically arrange records within the hierarchy of memory/storage according to their activity, significantly improving access utilization. Much sophisticated systems development work is still necessary if we are to fully exploit our technology capability.

A major advance in hardware reliability is urgently needed if the full potential of mass storage devices is to be realized. Capacity-access time improvements must be at the price of higher reliability because of the nature of on-line systems.

In this regard the mechanisms of strip-type mass storage devices are considerably more complex than those of disk files or drums (the proper surface must be selected at random from many hundreds, accelerated, guided repeatedly past a read/write station, decelerated, and returned to the file), and there is consequently a higher probability of device malfunction.

Beam addressable storage technology

A read-only photographic rotating disk memory was developed several years ago to serve as a dictionary for langauge translation.[5] The model stored 30 million bits (approximately 1500 bpi × 1500 tpi) with an access time of about 50 milliseconds. Optical sensing using a CRT flying spot scanner permitted "servo tracking" with extreme sensitivity, allowing a much

From the collection of the Computer History Museum (www.computerhistory.org)

closer balance to be obtained between bit and track densities than in magnetic recording.

A number of new electron beam and optical techniques have since been proposed which eliminate the need for mechanical motion for access, at least within a large block of information. These schemes have an appeal since beam addressable techniques have the unique capability for microsecond access to densely recorded information within the field of view of the beam. An electron beam can readily be deflected and modulated. A laser provides a light source whose output can be focused to a small size at high power densities. There is confidence that means for modulating and deflecting laser beams will be found that are practical.

Electron beam spot sizes less than one micron in diameter have permitted writing (in a vacuum) at a density of several million bits per square inch (approximately 2000 bits/in. by 2000 tracks/in.) on silver film. Readback must be done optically at a separate station after film processing. Future improvements may give even smaller beam sizes over a wider field of view. However, as in magnetic recording, we can now write information at much higher densities than we can effectively readback because we encounter still unconquered signal detection problems.

Optical techniques offer greater possibilities for high speed *parallel* data flow as well as the prospect of "distributed" bit storage (e.g., by holography) where the effects of media defects and registration misalignment at high densities may be minimized. A hologram is created from the bit pattern and recorded. The original image is reconstructed for readout by photodetectors. The greater possibilities of developing suitable "reversible" optical media (e.g., magneto-optical films) provide a further incentive for pursuing exploratory studies in optical recording.

Interest in photographic film stems exclusively from the potential it offers for high storage density. Photographic emulsions, which can be written with either light or high-energy electrons, are capable of resolutions of the order of 100 million spots per square inch. High speed recording is possible at a reasonable power level. But silver film is not reusable and a long delay occurs between the writing of information and its availability for write verify or reading. A significant capacity/cost advantage must be achieved in a read-only type store to justify the acceptance of the systems performance limitations on updating, posting, and immediate write-checking.

Organic photochromic materials fatigue with use and are relatively slow. Thermoplastic type media (which can be recorded upon by an electron beam) while reversible, involve a distinct developing phase, and the "deformation" storage mechanism limits resolution.

There are severe problems in locating and tracking information stored at very high density. Servo-techniques (also being pursued for higher track density in magnetic recording) based upon track-seeking principles are essential for beam scanning approaches.

Magnetic recording technology

Magnetic head design skills have advanced to a point where the principal concerns are to evolve batch fabrication techniques that will provide precision assemblies at a much lower cost. Magnetic heads can yield in excess of 10,000 cycles per inch resolution with frequency bandwidths extending considerably beyond ten megacycles.

No physical phenomenon for magnetic recording appears competitive with the simplicity and flexibility of the conventional magnetic head. Rather, the next major state of progress in the transducer area will be automated head fabrication (rather than watchmaker-like assembly), resulting in dramatic cost reductions. Such a breakthrough would have an impact upon the hardware composition of mass storage mechanisms, by making it economical to radically increase the overall ratio of heads to tracks.

The mainstream effort in magnetic surface work is to achieve higher quality, thinner recording surfaces. (Magnetic films in use today are oxide coatings formed from a dispersion of either Fe_3O_4 or $\gamma - Fe_2O_3$ in an organic binder, and Co-Ni platings.) Thinner magnetic recording films are indispensable for higher resolution whenever "bulk" erasure is not feasible and in any event must be correspondingly higher in quality (uniformity, finish, etc).

Reduced head-to-surface spacings (from approximately 100 μinches today down to approximately 25 μinches) will bring about a further large increase in "noncontact" bit density. Track seeking and following servo access techniques, which can circumvent the track density limitations imposed by the build-up of cascaded mechanical dimension tolerances, will be applied to effect significant increases in track density. Thus, recording mechanics still offer considerable room for further advances in storage density.

Improvements in the speed of head positioning actuators appear possible. However, progress in this area is not rapid and even a factor of two in speed will be difficult to achieve. Rather, if proposed batch fabrication processes can be successfully applied to achieve a drastic reduction in the cost of magnetic heads, and LSI (large-scale integration) enables a

From the collection of the Computer History Museum (www.computerhistory.org)

corresponding decrease in the cost of head switching and read/write electronics, there will be a substantial increase in heads per file module, minimizing the access dependence on head positioning.

Recording methodology is as yet an inadequately explored means to further upgrade storage density. Since the advent of digital magnetic recording for mass data storage, the primary avenue taken to increase density and thus performance has been through improved recording mechanics. This work emphasizes improving the "resolution" of a recording system and is approaching basic physical limitations. A relatively untapped potential exists for increasing magnetic recording density by the application of more sophisticated communication concepts to the encoding and decoding of the information to be stored.

Recent communications work exploiting adaptive equalization techniques and advanced coding theory has successfully demonstrated dramatic increases in data rates on transmission lines. In a similar vein the use of advanced communications techniques for the magnetic recording "channel" may yield significant gains in bit density relative to pulse width.

The calculated density limit based on a typical pulse response with standard recording techniques under ideal conditions is extremely high. However (even on paper), a large reduction in this potential figure results from a simple consideration of variability between magnetic heads, tolerances on magnetic coating thickness, radial variations in spacing and speed (with disks), head-to-track registration stability, surface speed fluctuations, etc. In addition we must deal with "noise" sources, such as surface imperfections, etc. All these factors must be accounted for in current file systems, which generally are designed around a single read/write channel serving a multiplicity of positionable transducers. Thus the actual operational density falls far short of that projected based only on a typical pulse waveform.

If adaptive techniques and advanced coding methods permit us to compensate for a number of these factors and operate anywhere near the "theoretical" limit of an idealized magnetic recording channel, the gain (through electronic rather than "mechanical" means) realized could be very significant.

Despite the bright projections into the near future for magnetic recording (the next generation of magnetic recording data storage devices may consolidate around 5000 bits/in. and 200 tracks/in.) the outlook for this recording process in the next ten years is not nearly so reassuring because it does not seem to lead as far as we would like to go.

At the present time there is no approach to mass data storage that could obviate the use of physical motion, using a continuous recording medium and transducer. While a completely electronic mass data store is an obvious goal, such memories seem possible only on the low capacity, high-speed side of mass storage (now filled by drums and disks with fixed heads).

There are substantial reasons to believe magnetic recording will undergo another decade of progress comparable to the past one, although there is something in man that rebels against the mechanical shuffling of data to recall information. We want instant memory and we desire this capability in the computers that we use.

For the foreseeable future, however, it is clear that mass storage based on magnetic recording has a vital and increasingly important role in information processing.

REFERENCES

1 A S HOAGLAND
Mass storage
Proc IRE 50 1087 1962

2 A F SHUGART YU TONG
IBM 2321 data cell drive
Proc SJCC Spartan Books Washington D C vol 28 p 33 1966

3 L BLOOM I PARDO W KENTING E MAYNE
Card random access memory (CRAM): functions and use
Proc EJCC Washington D C 1961 p 147

4 J D CAROTHERS R K BRUNNER J L DAWSON M O HALFHILL R E KUBEC
A new high density recording system: the IBM 1311 disk storage drive with interchangeable disk packs
Proc FJCC Spartan Books Las Vegas Nevada 1963 p 327

5 G W KING G W BROWN L N RIDENOUR
Photographic techniques for information storage
Proc IRE 41 1421 1953

BIBLIOGRAPHY

T H BONN
Mass storage: a broad review
Proc IEEE 54 1861 1966

W W CARVER
Comparing storage methods
Electronic Industries 21 120 1962

J S CRAVEN
A review of electromechanical mass storage
Datamation 12 22 1966

L C HOBBS
Review and survey of mass memories
Proc FJCC Spartan Books Washington D C p 195 1963

W W PETERSON
Addressing for random-access storage
IBM J Res and Dev 1 130 1957

C B POLAND
Advanced concepts of utilization of mass storage
Proc IFIP Spartan Books Washington D C p 249 1965

From the collection of the Computer History Museum (www.computerhistory.org)