Magnetic Recording Storage

ALBERT S. HOAGLAND, FELLOW, IEEE

Abstract—This review article attempts to give a perspective on the evolution of magnetic recording storage devices, highlight the frontiers of current activity, and address the future of this storage technology. Tape, disk (both rigid and flexible), and mass storage systems are covered. It is of interest to note that digital magnetic recording spans not only the twenty-five year history of the IEEE Computer Society which this issue is recognizing but has during this whole period undergone continual technical advances and widespread growth in usage. It is fair to say that magnetic recording is among the most important technologies in the history of the computer age to date.

Index Terms—Computer memory and storage, digital recording hardware, disk files, magnetic recording, mass storage, peripheral devices, storage density, storage trends and projections, tape drives.

INTRODUCTION

IMPROVING the performance and extending the utility of information processing calls for access to ever larger volumes of data. Thus, the need for more data storage with reasonable access time tends to be open ended no matter what progress we make. Capacity, access time, data transfer rate, and cost per bit are key parameters that define the basic performance characteristics of storage devices. Each storage unit has its own particular attributes, and most applications require that a hierarchy or mix of devices be interconnected in a computer system in order to achieve an overall cost-performance balance for the storage subsystem.

The modern electronic computer era is now some 30 years old. We have seen main-memory technology progress from vacuum tubes to mercury delay lines to cathode-ray tubes to magnetic cores and are now witnessing the change to semiconductor technology with LSI.

In auxiliary storage[1] (the storage levels beyond "main" memory) we can trace much less radical changes (not counting the early use of punched cards) but still see dramatic technological progress and while this paper is about such data storage, it is also about magnetic recording for the two have been wedded from the beginning. Its preeminence arises from the fact the medium is inexpensive and the recording process is basically simple and reversible (Fig. 1).

Manuscript received April 15, 1976; revised June 22, 1976.
The author is with the IBM Research Laboratory, San Jose, CA 95193.

[1] There is no clear distinction in the meaning between the terms "memory" and "storage," and we will not attempt to draw one here. Conventionally, the term memory is used in connection with devices directly interfacing with a processor while storage covers slower speed devices whose data are transferred to main memory before being processed.

Although magnetic recording was invented some 70 years ago, it is only in the last two decades that its potential has been realized. A further historical observation is that the magnetic drum emerged in the 1940's as the first widely used exploitation of digital magnetic recording, serving as main memory for medium and small computers.

Magnetic tape was commercially introduced in 1953 in the form of ½ in tape reels at a linear density of 100 bits/in (7 track) and with continuing technical improvements is still a highly used means of off-line storage with the most advanced drives recording at 6250 bits/in (9 track). However, the remarkable feature of magnetic tape devices is their backward compatibility, where the latest drives can still read and write tapes of two decades ago. This requirement has slowed the introduction of technical advances.

However, the role of tape has undergone a change since the introduction of the moving head disk file (RAMAC, 1956) and the significant advances in disk storage density since then. This direct-access storage device, made possible by the innovation of the air bearing head, has caused a shift from batch processing (set by the sequential nature of card and tape processing) to transaction processing and opened up a host of new applications where on-line interactive features were essential. Disk storage density has increased from 2000 bits/in^2 (RAMAC) to approximately 3-million bits/in^2 (IBM 3350), more than three orders of magnitude. The rate of progress in storage density has been essentially constant over this period and it appears will continue so into the 1980's (Fig. 2). While no technology can be economically improved indefinitely, technological potential is not so much a function of newness but intrinsic physical limits and the sophistication we can bring to bear. The volumetric density of storage has likewise improved, significantly reducing floor space for a given capacity, and the advances in LSI have contributed greatly to ease of attachment—both making disk files more attractive.

Well along in the evolution of the technology, the data storage needs of low-cost storage applications such as key entry led to an entirely new product family, the removable floppy disk file, now a major factor in the marketplace. This device emerged specifically to meet the requirements of data storage rather than as in the case of the cassette, a medium adopted from audio recording practice.

Much more recently we have witnessed the emergence of automated tape-type libraries which begin to make possible the storing an entire systems data base on-line. These mass storage systems represent additional evidence of the vitality of the digital magnetic recording field.

Thus, magnetic recording has proven the most enduring

Magnetic Recording as Reference for Alternative Technologies

Advantages: Nonvolatile/update in place.
Flexibility in configuration.
Pervasive (digital, video and audio applications).
Simple transducer (magnetic head) for both write and read.
Numerous media options.
Disadvantages: Involves mechanics.

Fig. 1. Magnetic recording attributes.

as well as the most pervasive of all technologies for digital storage and conversely, by far the largest application of the technology is for data storage.

Magnetic recording is a surface recording phenomenon. A magnetic head magnetizes a thin layer of magnetic material traversing the small region adjacent to the head gap in accordance with an applied write current and provides an induced voltage, reflecting the rate of change of magnetization, when scanning on readback. This basic principle is common to all devices but the design tradeoffs made to secure various storage capabilities lead to a wide diversity in technological directions. We shall cover in more depth the technical nature of the principal computer storage families, namely, rigid disk files, flexible disks, ½ in tape and mass storage systems in the following. Because access is mechanical, the average access time ranges from milliseconds to seconds as we go from disks to tape reels.

Rigid Disk Files

The moving head disk file based on a stack of rigid disks addressed by positionable read–write heads is the predominant data storage device today. Gradually, certain yardsticks have evolved to measure the merit of any given disk file. Some of the parameters often discussed are 1) price per million bytes of storage, 2) average access time, 3) average rotational delay (latency), 4) instantaneous data rate, and 5) capacity per actuator. Typical capacities per spindle for moving head files are in the range of 100–400 Mbytes.

The heart of a disk file resides in a few key technological components. With increasing recording densities the need for improved linear resolution has required a decrease in coating thicknesses from about 25 μm to about 1 μm. Furthermore to achieve these higher linear densities, the head-to-medium spacing has been reduced from about 10 μm to less than 1 μm (Fig. 3). Consequently, recording surface finish requirements have become much more demanding. As the requirements for a thinner coating become more severe as density increases, it is also essential to improve the substrate.

Magnetic recording resolution is essentially established by the geometry in the pole-tip area of the head (head gap, spacing, and medium thickness). Fig. 4 shows how these dimensions have changed over the years. The storage medium now used is a magnetic oxide particulate coating, typically about 1 μm thick. The read–write heads consist

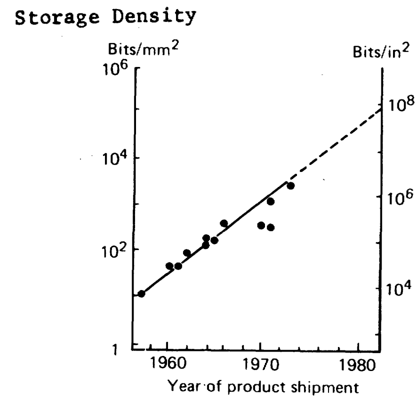

Fig. 2. Storage density versus time.

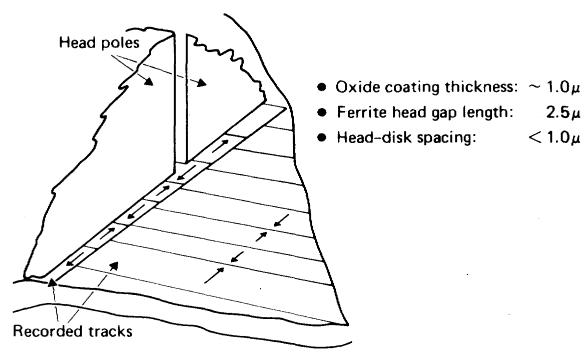

Fig. 3. Current disk recording parameters.

of ferrite cores with a gap of 1–2 μm, spaced 1 μm or less from the disk surface. To date disk products have achieved an areal density of approximately 3×10^6 bits/in^2. Higher densities require a continual scaling down of recording dimensions. The medium appears to be the limiting factor as we push to higher and higher densities for as bit size reduces the medium quality must improve.

The knowledge gained in magnetic film technology suggests a means of improved manufacturing control of the head parameters. That is, the deposit of thin-film magnetic heads on a substrate should, in the near future, become viable as a technique to mass produce magnetic heads with very high density characteristics.

Track densities have increased to the point that open-loop discrete-positioning is no longer adequate and closed-loop track-following systems based on voice coil actuators are utilized. Track density is determined by the accuracy of head positioning and the transverse resolution of the read–write head.

As track densities increase, head to track registration tolerances become more and more critical (Fig. 5). At some track density, the concept of using the same transducer for both reading and writing data as well as sensing servo information (incorporated within the data along each track) will be introduced to eliminate all possible mechanical tolerances. For reasonable disk array sizes, this will probably be in the 1000 tracks/in range. Head performance improvements are also needed for track density advances. Head design must minimize transverse sensitivity to minimize adjacent track pickup and still maintain ade-

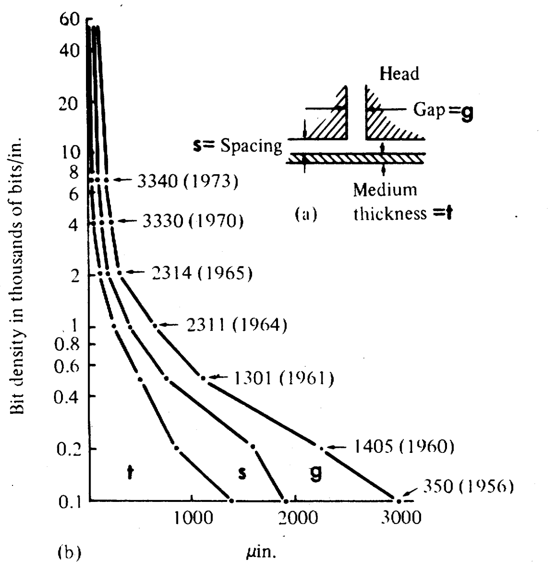

Fig. 4. Bit density as related to physical recording parameters (disk). (a) Head/medium geometry. (b) Density as function of geometry.

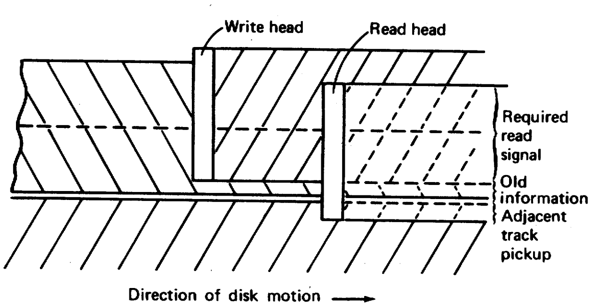

Fig. 5. Read head pickup of old and adjacent track information.

quate SNR since the readback signal decreases with track width and bandwidth increases with linear density.

Actuator performance has steadily improved from an average seek of hundreds of milliseconds to an average seek of tens of milliseconds. While some gains remain, changes of an order of magnitude do not appear feasible nor worthwhile since rotational delay plays a nearly equal part in access time at this point, and further improvements have only a small effect.

Medium flaws make ECC a standard procedure in present files and the use of error correction codes in conjunction with defect demarking (and subsequent skipping) will increase in the future. The most common codes used are the cyclic codes which have burst error correction and detection capabilities.

One can improve the cost per byte by taking advantage of increasing areal density to make a larger capacity file. However, the objective of achieving performance improvements in smaller capacity modules for a low entry cost is a great challenge. One must still pay the costs associated with high technology components to effect engineering tradeoffs that are favorable to the required drive hardware.

As the cost of the disk file function continues to reduce with advances in areal density, the concept of a "fixed disk" array becomes more and more feasible since capacities become such that interchangeability is no longer a necessary feature for economical reasons. Fixing the disks can go a long way toward satisfying the cost requirement that makes fixed disks feasible by minimizing registration tolerances and environmental problems and in turn making the higher densities that are required more realizable.

A recent design trend is to combine fixed-head storage on the same spindle with a moving head and thus reduce the cost premium for having some data available with a very short access time. Previously this capability was available only through separate fixed-head drums or disks.

Therefore, as significant increases can be expected in the size of computer main memory it will be accompanied by large increases in disk storage capacity to give significant improvements in systems performance. A three-level storage system can be postulated—a larger main memory where what is happening now resides, a large disk file as a recording store for all active data, and an automated tape library (tertiary storage) for very low-cost archival-like storage.

In summary the moving head disk file will continue to evolve to an everdecreasing cost per byte with an everincreasing nominal capacity. This will be accomplished by advances in recording density, both through technology and by greater use of error detection and correction codes. Densities to 10^8 bits/in^2 now appear obtainable. To allow increasing recording density with least cost, fixed-disk files will again become common. Performance enhancement through configuration optimization can also be expected.

Considering the evolutionary progress that will continue in disk file technology, the payoff must be very large to justify introducing a competitive technology. Thus, with all factors considered, it is unlikely that any other technology will displace magnetic recording as the principal active data-storage approach in the forseeable future.

Flexible Disks

Another category of disk file uses a flexible disk as the storage medium. The disk is removable and has its own protective container. Capacities up to one megabyte have been realized and designs are aimed at achieving the lowest entry cost possible. Low-cost media with ease of handling is achieved by capitalizing on the mass production capabilities existing for magnetic tape.

The earliest application of the flexible disk was as an alternative to a drum memory developed in the early 1960's with a single-disk riding on an air film over an array of fixed heads.

The "diskette" drive was introduced by IBM in the 1970's for key entry and utilizes *contact* recording at surface speeds in the range of ½ in tape drives. The track density is 48 tpi and the access time currently is greater than 100 ms. The rotational speed is limited by frictional heat generated within the cartridge liner. The "floppy" disk is an example where the application of available technology to a developing need created an entirely new marketplace.

We now see performance increases in the flexible disk exhibiting the same evolutionary trend as for rigid disks. Furthermore, the flexible disk is beginning to be viewed as a low-cost direct access storage device in addition to its role for data entry and load/dump functions.

Ultimately, advances in storage density are the key measure of progress in disk storage. Although the intrinsic hardware required for a drive and read–write channel limits the degree to which areal density gains can be traded for lower cost, the growing need for more on-line capacity provides a ready opportunity to exploit technology advances.

One-Half-Inch Magnetic Tape

The history of magnetic tape data storage has been the history of half-inch tape. The first successful commercial product appeared in 1953 where data were recorded at a density on 100 bits per linear inch (one byte in parallel) of half-wide magnetic tape. Their initial role in providing I/O functions and for sequential processing (a direct extension from punched card procedures) brought about the development of newer, faster, and better drives. Approximately, 100 000 half-inch tape transports are now in existence.

Today there are additional magnetic tape devices which have been introduced to fill a need in the area of low-cost data collection and processing. The most successful has been the modified tape cassette.

The most amazing thing about half-inch tape is the upward/downward compatibility that has been imposed upon this medium. Today's machines can read tapes that were generated 20 years ago at 200 bytes/in and have the capability of regenerating data at 6250 bytes/in on that same reel of tape. In the same time frame, throughput rates have increased from 15 000 to 1 250 000 bytes/s.

In today's drives the oxide-coated surface of the tape touches only two points: the read/write head and the tape cleaner. This development did not take place until the late 1960's. The combination of lower tape prices and higher densities has improved the two figures of merit that are so important to such storage; cost per bit stored and the volumetric efficiency. One-half inch tape is widely used for archival storage and is the cheapest means for storing data off line.

The start-stop feature needed for data storage distinguishes the magnetic tape drive from others. Much higher areal densities have been realized in instrumentation recorders but their access features are not oriented to tape processing. The application of advanced recording technology to tape has been retarded by systems factors, unlike the situation in disk files.

Mass Storage Systems

The Ampex TBM is the first mass storage system (MSS) based on magnetic recording to come on the market.

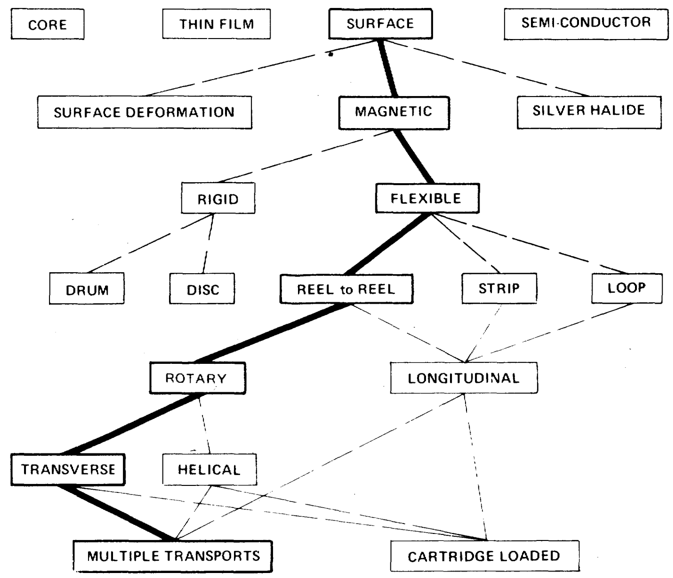

Fig. 6. Decision path in selection of Ampex terabit configuration.

Guidelines for its development were to provide at least 10^{12} bits and sell for no more than one million dollars with a reliability equal to the then current computer peripherals.

Having made the decision that the mass store would be magnetic, a choice had to be made between rigid surfaces (either drum or disk) or flexible surfaces. Flexible was chosen in order to get the most information per unit volume. The cost of the rigid media alone would have exceed cost per bit contemplated for a practical commercial venture. The choice of flexible media was thus easy to make.

In order to package the recording surface (over a million square inches of medium were required) large reels of tape were chosen. To maximize areal density a rotary head read–write arrangement was selected based on prior video recording work. Multiple transports with large tape reels were selected for the configuration to minimize mechanics. Fig. 6 shows the decision tree and the path chosen by Ampex.

A packing density of 1.5 million bits/in^2 was achieved. This packing density on a flexible medium was quite remarkable at the time since conventional computer tape drives were operating at approximately 14 000 bits/in^2. The data reliability was adequate to meet the system performance goals provided error-correcting codes were used in addition to a dual redundancy scheme.

IBM in 1974 announced the 3850 MSS. The 3850 used flexible magnetic media and a rotary head. The IBM 3850 system, however, uses individual spooled cartridges which are moved from a honeycomb-like cartridge library to and from magnetic recording stations (Fig. 7) while the Ampex memory system used large reel transports in parallel. These differences lead to widely different system characteristics.

The IBM 3850 MSS was developed to bring all of a customer's data on-line with disk-access features through

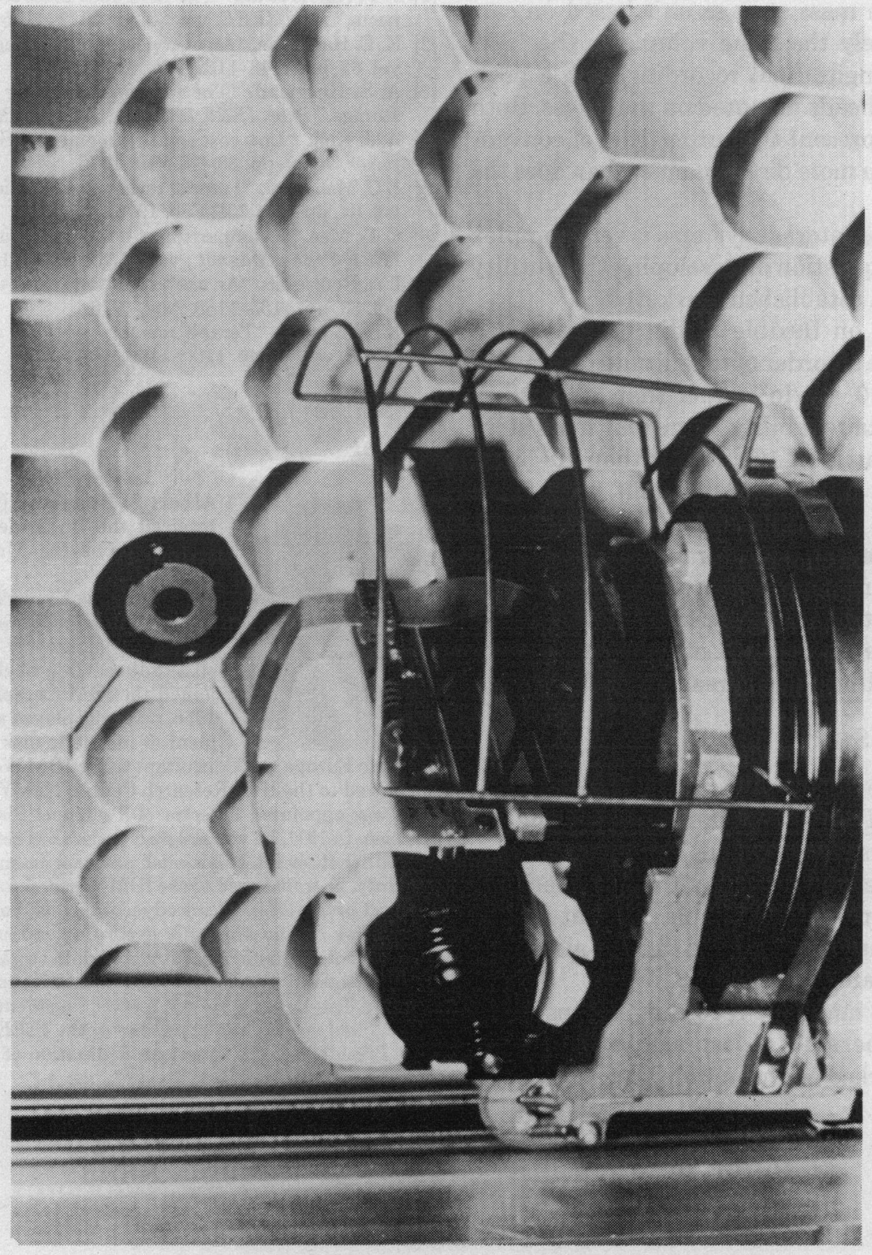

Fig. 7. 3850 cartridge library.

automatic staging. The overall library storage cost can approximate that of today's half-inch tape facilities. A hierarchy of disk and tape-cartridge devices provides the flexibility needed to approach this "virtual storage" objective. The current 3850 provides a capacity configuration from 35-billion to 472-billion bytes.

There are two other entries as of now in the MSS area. CalComp has essentially automated the handling of one-half inch tape reels to provide an on-line library function. CDC has announced a mass storage unit based on cartridges of approximately the same volume as the IBM cartridge but using longitudinal recording with a positionable set of parallel heads mounted on the drives. Both of these approaches represent the automation of conventional tape handling in a more direct manner than does the 3850.

The transition to mass storage systems is very complex and for the time it is a question of developing their utility rather than selling to a established marketplace.

Magnetic recording on flexible media has advanced sufficiently to visualize an order of magnitude increase in packing density, i.e., 10^7 bit/in^2 rather than 10^6 bit/in^2, without any loss in data reliability. A current mass store could thus have a capacity of 10^{13} rather than 10^{12} bits; conversely, a 10^{12}-bit store could be built with fewer modules and therefore at lower cost.

The advent of commercial mass storage systems further extends the role that magnetic recording plays in data processing. As these systems become easier to use and their capacities increase, the possibilities of also putting non-coded information on-line become real.

SUMMARY

The long-term future of digital magnetic recording is not only assured but it will have a continually increasing impact on information systems and thus on our lives.

While it may be overstating to say that magnetic recording is a development comparable to that of the printing press as one author has stated [5]; it is also probably not that much of an exaggeration.

Even though digital magnetic recording is considered a mature technology there is a widespread belief that the next decade will be marked as one in which the application of magnetic recording for information storage will be greatly expanded and the technical advances realized will continue at the rate of the past and provide further major improvements in cost/performance. From the perspective of an alternate technology—the moving target keeps moving!

REFERENCES

[1] J. P. Harris *et al.*, "The IBM 3850 mass storage system: Design aspects," *Proc. IEEE*, vol. 63, pp. 1171–1175, Aug. 1975.
[2] K. E. Haughton, "An overview of disk storage systems," *Proc. IEEE*, vol. 63, pp. 1148–1152, Aug. 1975.
[3] A. S. Hoagland, "The future of magnetic recording for computer storage," *Proc. IEEE Intercon*, vol. 23/2, Mar. 1974.
[4] A. J. Kolk, "Low cost rotating memories: Status and future," *Computer*, vol. 9, pp. 30–34, Mar. 1976.
[5] J. C. Mallinson, "Tutorial review of magnetic recording," *Proc. IEEE*, vol. 64, pp. 196–208, Feb. 1976.
[6] C. D. Mee, "A comparison of bubble and disk storage technologies," *IEEE Trans. Magnet.*, vol. MAG-12, pp. 1–6, Jan. 1976.
[7] I. A. Rodriquez, "An analysis of tape drive technology," *Proc. IEEE*, vol. 63, pp. 1153–1159, Aug. 1975.
[8] M. Wildman, "Terabit memory systems: A design history," *Proc. IEEE*, vol. 63, pp. 1160–1165, Aug. 1975.

Albert S. Hoagland (S'50–A'54–SM'57–F'66) received the Ph.D. degree in electrical engineering from the University of California, Berkeley.

While at the University of California, Berkeley, he participated in the formation of the Computer Laboratory, and in 1954, he accepted a professorship in electrical engineering. He joined the IBM Corporation, San Jose, CA, in 1956, where he played a key role in the development of magnetic disk files. In 1962, he transferred to Europe as a Consultant to the IBM World Trade Corporation. He moved to the IBM Research Center, New York, NY, in 1964, and in 1968 was appointed Director of Technical Planning for the Research Division. In 1971, he was assigned to the IBM Boulder Laboratory to lead a multilocation interdivisional program in mass storage technology. Recently, he transferred to the IBM Research Laboratory, San Jose, CA, to direct programs in advanced recording media for storage. He is author of the book *Digital Magnetic Recording* and numerous publications in the fields of magnetic recording and data storage.

Dr. Hoagland is a member of Phi Beta Kappa, Sigma Xi, Eta Kappa Nu, and Tau Beta Pi. He is a former President of the IEEE Computer Society and is currently a member of the IEEE Board of Directors and Vice President of the American Federation of Information Processing Societies.