L. D. Stevens

The Evolution of Magnetic Storage

Since delivery of the first vacuum-column magnetic-tape transport in 1953 and the first movable-head disk drive in 1957, tape and disk devices in many configurations have been the principal means for storage of the large volumes of data required by data processing systems. Magnetic drums and other device geometries have also been important system components, but to a lesser extent. Over the past twenty-five years significant developments have been made that increase the capacity, reduce the cost, and improve the performance and reliability of these devices. With each improved device the range and nature of the applications undertaken have expanded and, in turn, led to a need for further device improvement. This paper gives a general review and historical perspective of magnetic storage development within IBM and is an introduction to the subsequent papers on disk, diskette, and tape technology and on disk manufacturing.

Introduction

Data processing applications of computers have grown over the past twenty-five years from incidental significance to a point where they have now become a pervasive influence in our society. Early data processing systems used magnetic tape as the principal storage medium for large data files. Processing was batch sequential on a job-by-job basis and the application focus was accounting. These systems had only secondary impact on the operational aspects of business. Those early computers are in sharp contrast to the data processing systems of today, which allow many different jobs to run concurrently with very-large-capacity on-line magnetic storage (*i.e.*, directly accessible without human intervention), data-base-oriented transaction processing, and an application focus on making more efficient use of operational resources.

Improvements in the cost, capacity, and performance of on-line magnetic storage have fueled these growing systems and their application capability. Three distinct periods can be identified in this evolution. During the first—the early years from 1953 to 1962—limited on-line storage was provided by the tape drives (with mounted reels of tape) attached to the system. Disk storage was a scarce resource, found only in those systems where the high cost, limited capacity, and difficulty of use could be justified by its capability for direct access of data. In the next period—the transition years from 1963 to 1966—rapid development of disk technology and systems software removed many of these constraints. Disk storage and on-line processing began to be an important part of most systems although tape storage and batch processing were still dominant. During the third period—the growth years from 1967 to 1980—the cost per Mbyte of disk storage was reduced twentyfold and with further improved systems software, new terminals, communication facilities, and on-line application development, substantial growth occurred in the on-line storage capacity of the average system; see Fig. 1. Here, the main memory capacity of the average IBM data processing system is compared with its disk and tape storage capacity during this period. Disk capacity per system increased by a factor of forty from a base of 23 Mbytes, attached-tape capacity increased by a factor of seven from a base of 47 Mbytes, and main-memory capacity per system increased by a factor of nineteen from a base of 50 Kbytes (where K = 1024). Disk capacity per system has been about 1000 times greater than that of main memory since 1973, and combined disk and tape capacity has grown to 1600 times that of main memory.

After a brief review of the basic capacity, cost and performance aspects of magnetic storage devices, this paper gives a historical perspective of device developments within IBM during each of these periods.

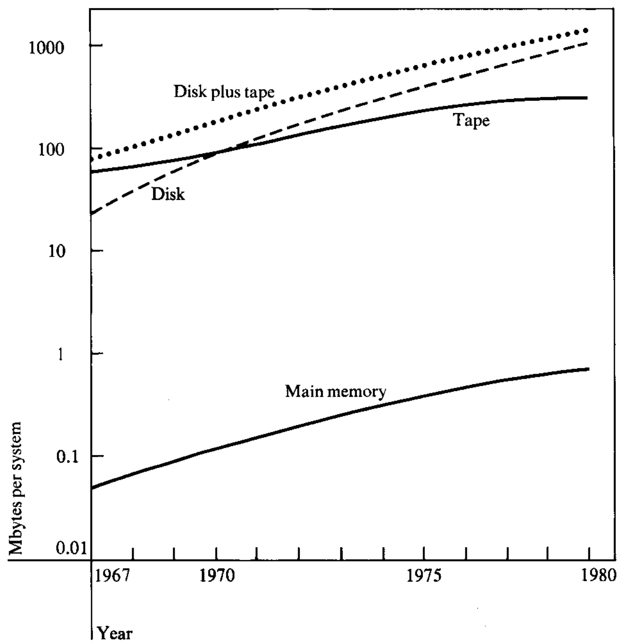

Figure 1 On-line storage capacity: disk, tape, and main memory capacity installed on the average IBM data processing system from 1967 to 1979.

Basic aspects of magnetic storage

Improvement of storage device cost, capacity, and performance has been achieved by a continuing development of magnetic recording and electromechanical access technology. Subsequent papers in this issue describe the innovative details of these developments for devices based on half-inch magnetic tape [1], rigid magnetic disks [2, 3], and flexible magnetic diskettes [4]. The following discussion provides a general introduction to these topics.

Magnetic recording technology includes the magnetic storage medium, the read and write heads, the recording channel electronics, the data encoding and clocking logic, and the technology that controls the head-to-medium spacing. The key parameter is areal density, the product of the linear bits per inch along a track (bpi) and the number of tracks per inch (tpi). It affects both capacity and performance and is the key parameter that determines the cost per Mbyte of storage.

Performance and capacity are also affected by access technology. For tape devices, it is the technology that controls the tape acceleration, velocity, and deceleration. The key parameters are the time required to start and stop, and the length of the resulting gap between blocks of data. For moving-head disk devices, it is the technology that controls the radial positioning of one or more read/write heads to selected concentric storage tracks. The key parameters are the average seek time (positioning time) and the accuracy of positioning.

- *Cost, cost/capacity, and cost/performance*

A large fraction of the cost of a tape transport or disk drive can be attributed to the fixed costs of the basic hardware required to support, protect, and transport the relatively inexpensive storage medium. The cost per Mbyte of storage is reduced with increasing areal-density capability because these fixed costs can be shared with more storage capacity. In addition, as areal-density capability has increased, the cost associated with the heads, actuators, media, recording channel electronics, and other technology components has remained nearly constant because of design improvements. Consequently, most designs have focused on providing substantial improvements in the cost per Mbyte for disks and the cost per Kbyte per second for tape by increasing the capacity and performance for equivalent or slightly increased cost.

- *Capacity*

Increased areal density has been the primary means of improving the capacity of both tape and disk devices. The companion papers [1-4] discuss the details of how this has been achieved through thinner particulate coatings having improved magnetic properties, the use of improved head materials, better fabrication techniques for smaller recording head gap length, reduced spacing between the head gap and the magnetic surface, and more accurate head positioning for disk drives.

Reducing the head-to-surface spacing has been a significant factor in increasing linear bit density. Contact between head and medium, with the proper head design, is acceptable for low-velocity diskette devices, as discussed by Engh [4]; contact is, however, not acceptable for high-velocity tape and disk devices. For these devices, a thin film of air is used to provide a lubricating air bearing that in turn determines the spacing. With magnetic tape and high-speed diskettes, the intrinsic boundary-layer airfilm forms a hydrodynamic air bearing as the flexible tape moves by the head. Control of this bearing is achieved by design of the head surface contour, as discussed by Engh [4] for diskettes, and by Harris, Phillips, Wells, and Winger [1] for half-inch magnetic tape. Spacing in the range of 5-10 microinches has been achieved. With magnetic disks, the spacing is accomplished by a separate air bearing support for the head. Progress in disk air bearing technology has reduced the spacing from 800 microinches in 1957, with hydrostatic (pressurized) air bearings, to the current 10-20 microinches today, with lightly loaded hydrodynamic (self-acting) slider bearings, as discussed by Harker, Brede, Pattison, Santana, and Taft [2]. These authors also discuss the evolution of improved disk read/write heads from laminated mu-metal, to ferrite, to thin films; the improvements in fabrication technologies that have reduced the gap length from 1000 to 40 microinches;

and the improvements in particulate magnetic coating and processing technologies that have reduced magnetic film thickness on disks from 1200 to 20 microinches. The manufacturing aspects of these disk technologies are further discussed by Mulvany and Thompson [3].

Increased track density (tpi) has been of minor significance for half-inch tape devices. At first it was limited by head fabrication technology for the construction of seven, and later nine, accurately aligned parallel gaps across the half-inch tape. After the track pitch was established, the use of tape for system data interchange created a compatibility requirement that has constrained change. A similar data interchange constraint has influenced the design of diskettes, as discussed by Engh [4].

Track density improvement for moving-head disks has been an important contributor to increased areal density, but less significant than linear density. The track density of disks has been limited by the transverse resolution of the read/write head(s) and by the accuracy of their positioning. An error in positioning that exceeds a small fraction of the track width can cause either partial erasure of data on an adjacent track or failure to erase (or overwrite) previously written data. Either, upon readback, will result in a decreased signal-to-noise ratio. Mechanical detents were used to establish the final head position on early disks, and track density was limited to about 100 tpi. Separate write-wide/read-narrow or tunnel erase heads were required to establish guard bands between each data track. Positioning accuracy was improved by the development of closed-loop track-following servo systems. The result was a significant increase in tpi and the ability to use the same gap for reading and writing, as discussed by Harker *et al.* [2].

- *Performance*

The time required to obtain access to the start of a block of data plus the time required to transfer it to main memory determines the performance of both tape and disk devices. The relative importance of these two time components is a function of the data processing environment. The effective data rate—that fraction of the intrinsic data rate that can be realized—is of primary importance for either disk or tape used for sequential processing of a large number of records. But the rate of access (*i.e.*, accesses per second) deliverable with a reasonable response time is the important measure for disk storage used by interactive systems that generate essentially random requests for relatively short records.

Improvement of sequential processing performance has been significant. The intrinsic data rate of tape has been improved from 7.5 Kbytes per second (Kbyps) to 1.25 Mbytes per second (Mbyps) by means of increased linear density and velocity. The effective data rate of tape has been improved from about 50% of the intrinsic value to over 80% of the intrinsic value by increasing data block size and reducing the interblock-gap (IBG) length. Similar progress has been made in the sequential throughput from disks by improving the intrinsic data rate from 4.4 Kbyps to 3 Mbyps and in the effective data rate by developing of the multiple-head access structure, which provides a conceptual "cylinder" of data tracks at each access position and allows an essentially continuous stream of data by electronic switching between tracks.

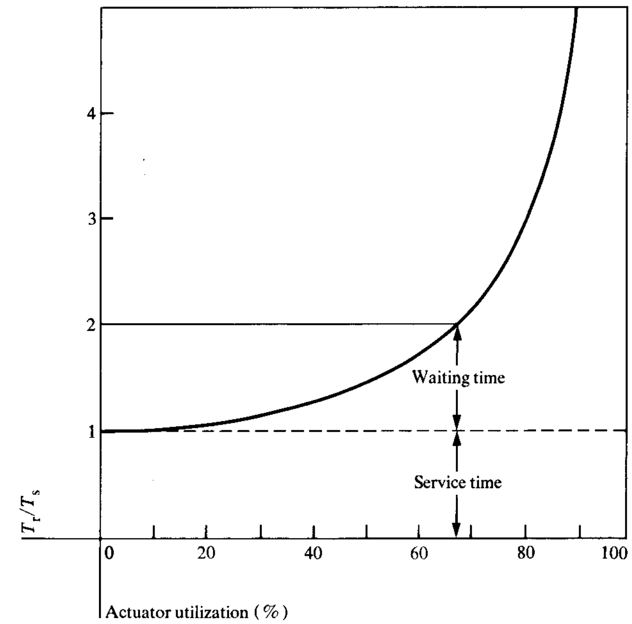

Figure 2 Disk response time and service time relationship: ratio of mean disk response time T_r to the average actuator service time T_s as a function of percent actuator utilization.

Comparison of the improvement in disk access rate achieved by actuator development leads to a discussion of the basic components of disk service time and the influence of actuator utilization on response time.

Disk service time is composed of three components: average seek time to position the heads, rotational latency to locate the start of the desired record (on the average, half of a revolution), and data transfer time (typically about one-tenth of a revolution). The maximum random access rate given by the reciprocal of this service time is not achievable when the constraint of a reasonable response time is imposed, because response time is composed not only of the service time but also of the time a request must wait while the actuator is servicing previous requests. Waiting times become large when an actuator is working close to 100% of its capacity and grow without bound when saturation is reached. This is shown in Fig. 2, which plots the ratio of mean response time to the average service time as a function of actuator utilization, assuming random arrival of requests and constant service time [5, 6]. It can be seen that the waiting time component of the ratio begins to increase rapidly as the actuator is utilized more than two-thirds of the time. Here the ratio is two and the waiting time is equal to the service time. Given these assumptions, the random access throughput rate for disks has improved from about one access per second for the first disk drive to about 26 accesses per second for the most recently announced disk drive.

In addition to improvements in the access rate from individual disk actuators, the number of actuators that can be operating in parallel with overlapping seek, data transfer, and processing has been increased by improvements in the architecture of I/O subsystems [7], including the development of operating systems with multiprogramming capability, data management software, data channels, and intelligent device control units.

Historical perspective

This section reviews the development of magnetic storage devices and control units in the context of their systems software and application environment during each period of evolution.

- *The early years from 1953 to 1962*

Highlights of the changing environment during this period are as follows:

- Magnetic tape replaced punched cards as the principal storage medium for large data files; batch processing techniques continued to be used but at a much higher speed.
- Magnetic drums were used as main memory early in the period and as extensions of main memory throughout the period.
- Direct access to limited amounts of data became possible with the delivery of the first moving-head disk drives in 1957.
- Overlap of processing with I/O operations grew as the concepts of multiprogramming, CPU interrupt, data channels, and independent device control units were developed.

Magnetic tape (1953–1962) IBM delivered a high-performance digital magnetic tape drive in 1953 with the IBM 701. Its unique vacuum-column design isolated the high-inertia reel drive from a high-performance pinch roller and continuously rotating capstan used to control tape velocity [8]. The design concepts of this tape drive, the IBM 726, were the foundation for technological improvement for a decade.

Using an IBM-developed version of nonreturn-to-zero encoding (NonReturn-to-Zero-Inverted or NRZI), data recording was done on seven parallel tracks across a half-inch plastic substrate [acetate at first and later Mylar tape coated with magnetic iron oxide (Mylar is a trademark of E. I. du Pont de Nemours and Co., Inc., Wilmington, DE.)]. Each of six tracks represented a weighted bit in a six-bit binary coded decimal (BCD) character; the seventh was a redundancy bit. Recording density increased from 100 to 800 bpi during the period and tape velocity increased from 75 to 112.5 inches per second (ips), while the interblock gap (IBG) remained at 0.75 inch.

Improving start/stop time and thus reducing the necessary IBG length has been one objective of tape drive development. In addition, the blocking of a number of logical records to increase the size of each physical block between gaps has also been important to increased tape storage efficiency and effective data rate. In early systems, however, blocking was limited by the lack of memory for buffering and the lack of blocking and deblocking software. By the end of the period, main memory capacity had increased, and software had been developed to allow blocking factors of ten or more logical records per physical block. Blocking of logical records has continued to be an important factor relating to the efficiency and effective data rate of magnetic storage devices [9].

Processing of the data on magnetic tape followed a pattern similar to that of punched cards. Inputs were grouped together in a collection called a "batch," sorted into the same sequence as the files of data and processed job-by-job until the entire batch was completed.

To make space for new records and use the space of deleted records, the updating of a tape file required rewriting the entire tape. For example, transactions to a file would be read from one transport, the old file read from a second, and a new updated file written on a third. Therefore, the updatable unit for tape processing was a complete reel of tape.

This process was very efficient [10] if the file had a high activity rate. But if only a few percent of the records in a file were to be changed, as was typical of applications such as large life insurance policy files, maintenance of the master file was a major problem.

Another related problem was inquiry. The status of a record could be determined only through listings made each time the file was updated. Many applications, such as bank demand deposit and inventory control, require frequent inquiry, and the need for a storage medium with better access to individual records began to grow.

In early tape processing systems the CPU was not busy much of the time. It had to wait during access to each block of data and the writing of the new block. Much of the actual process time was spent inspecting records, with little useful work being accomplished. At first, the user had to program all of the I/O operations explicitly, including blocking and deblocking of records and recognition of special tape marks such as end-of-file and end-of-tape. Simple I/O monitor routines were soon developed, however, and stored on a system tape for common use [11]. In addition, tape control units were developed that independently managed the tape drive.

The IBM 777 Tape Record Coordinator, delivered in 1956 with the IBM 705 II, was the first such control unit and was also a precursor of the data channel. This buffered tape control unit allowed the CPU to overlap tape operation and processing by use of primitive "multiprogramming" [12]. Rochester, who was the first to use this term, describes the operation in an insurance policy file maintenance application [13]:

"The Tape Record Coordinator . . . runs its master tape units for the file maintenance job almost autonomously while passing over inactive records. Then when an active record is found, the calculator briefly interrupts the work it was doing (on another job) to deal with this active record. . . . When it is passing over inactive records the calculator needs to spend less than 1/2 percent of its time supervising the search. All the rest it can devote to the other job."

These primitive hardware and software concepts were expanded and generalized in later evolution periods and were destined to have a significant impact on the evolution of systems design as they were generalized to include all types of I/O including a person at a terminal [14].

Magnetic drums (1953–1962) Magnetic drums could not achieve as high a recording density as magnetic tape because of mechanical limitations of the head-to-medium spacing. A nominal spacing was established by machining with a final adjustment to about 0.001 inch by some type of differential screw. In addition to the spacing limitations, many early drums used a less efficient return-to-zero (RZ) recording technique to allow easy selective alteration of individual bits [15]. Linear densities were about 50 bpi and track spacing was about 20 tpi.

Two magnetic drums, typical of the early part of the period, were developed by IBM, one as an extension of electrostatic memory on the IBM 701 [16] and the other as main memory for the IBM 650 [17]. Use of magnetic drums as main memory began to diminish with the advent of magnetic cores, but their use as a main memory extension persisted well into the next period of evolution.

Toward the end of the period IBM developed an innovative drum for a version of the SAGE air defense computer. It was the first magnetic storage device to use a hydrodynamic slider bearing to establish head-to-medium spacing [18, 19]. The use of this technology permitted linear densities approaching those on tape with no mechanical adjustments.

Magnetic disks (1957–1962) Development of a hydrostatic air bearing to accurately space a magnetic head close to the surface of a rotating disk made possible the shipment in 1957 of the first movable-head disk drives with the IBM 305 [20] and 650 RAMAC systems. Its use allowed the magnetic head to follow minor axial runout of the disk while maintaining a constant spacing of about 800 microinches. A recording density of 100 bpi and 20 tpi [2 Kbits per square inch (Kbpsi)] was achieved [21]. This first disk drive used a pair of air-bearing-supported heads mounted on an access arm that could be moved under servo control to one of fifty 24-inch-diameter disks mounted on a vertical shaft and rotating at 1200 revolutions per minute (rpm). When positioned at a disk, the head pair could be moved radially to any of 100 tracks. Average seek time was 600 milliseconds (ms). Each track stored 500 BCD characters formatted into five fixed-length 100-character records. Each disk surface stored 50 000 characters; the 100 surfaces, a total of five million characters. Additional capacity was sacrificed for the simplicity of consistent track capacity and a single data rate; the maximum bpi was used only on the innermost track.

By the end of the period the hydrodynamic slider [22] had replaced the hydrostatic air bearing, and head-to-disk spacing in the range of 250 microinches was achievable. A comb of 50 access arms, one per disk surface, each with a slider and read/write head, was now made practical by the elimination of the large air compressor requirement for a similar array of hydrostatic bearings. Two such access structures, each with its own hydraulic actuator, were used, together with advanced magnetic technology, on the IBM 1301 disk drive, which was first installed in 1962. It provided 56 million characters of storage on fifty 24-inch-diameter disks. This capacity improvement of more than a factor of ten was achieved by an increased recording density of 520 bpi and 50 tpi (26 Kbpsi). The 1301 had an improved average seek time of 165 ms, achieved by a high-performance hydraulic actuator and the elimination of the disk-to-disk motion.

Magnetic disks provided a combination of direct and sequential access. With the comb access mechanism a conceptual cylinder of 50 data tracks was formed at each track position. Once a position was selected by mechanical motion, any one of the 50 tracks could be selected by electronic switching. A data record on the selected track could be read on one disk revolution and an updated version written on the next revolution without affecting any other record. This in-place update and direct access capability of disk storage offered significant advantages.

Fixed-length records were used in these early disk drives. Record formats, however, were different for each application file and some flexibility was desirable. The first disk drive control unit, the IBM 7631, used a track associated with each cylinder of data for format control and allowed the specification of a different record format for each cylinder. It also provided for seek overlap in conjunction with the Input/Output Control System (IOCS). The IOCS provided for the management of buffering which allowed the use of multiple devices without the user being concerned with the synchronization of the hardware. The IOCS also blocked and deblocked the user's logical records to the physical tracks of the disk. It allowed a user to process a sequential data set from disk by use of high-level macro instructions.

The early disk storage devices presented other problems. These devices were not very reliable and there were no well-developed techniques to address a disk record directly, except in the very limited case that its physical address could be obtained by a linear transformation of the record identifier. As a result, the disk data management functions of the IOCS were, for the most part, an extension of those provided for tape. Direct-access processing required the application program to provide the physical address of the record (or block of records) by transformation of the logical record identifier to the physical device address before issuing calls to the IOCS. This required a detailed knowledge of the physical device. Soon, however, generalized indexing systems and nonlinear transformation techniques were developed to provide skip sequential and direct access processing [23, 24], and on-line data processing [20] began to grow [25].

With disk storage as an element of the system, jobs could be entered into the system in arrival sequence and queued on disk. If I/O units required by a job were busy, the processor could proceed with another task from the job queue on disk. The disk also provided residence for system software, allowing only the most frequently used programs to reside in main memory.

Disk storage remained a scarce system resource throughout the period because it was expensive. Although the cost per Mbyte had been reduced by a factor of two, disks were still limited to applications such as programming systems residence, reservations systems, and inventory control, where the cost of disk storage was still justified.

• *The transition years from 1963 to 1966*

Highlights of the changing environment during this period are as follows:

- The cost of a Mbyte of disk storage was reduced to approximately the cost of attached tape, and, as noted by Bonn [26], the design of computer systems entered a period of transition from tape to disk storage with online processing. Batch processing with tape, however, was still dominant.
- Introduction of the removable disk pack in 1963 provided off-line shelf storage of disks, but at the end of the period they were still twenty times more expensive than tape.
- IBM System/360 was introduced with an I/O architecture and with programming systems that required and enhanced disk storage capability.
- The role of magnetic tape began to shift from primary storage medium to systems interchange and archival storage.
- The first intelligent microprogrammed storage control units were introduced, and system attachment was simplified with the definition of a common I/O interface.

Magnetic tape (1963-1966) Two tape drives were announced with System/360. One, a new model of the IBM 7340 Hypertape drive [27], had a remarkable density (for the time) of 3022 bpi using one-inch tape in a cartridge with a two-cartridge autoloader, but was incompatible with existing half-inch tape libraries and was not widely accepted. The other, the IBM 2401, used a nine-track format for the eight-bit byte of System/360, was compatible with standard half-inch tape, and found wide acceptance. Recording density was at 800 bpi with NRZI encoding and later at 1600 bpi phase-encoded. Various models of the 2401 offered both nine- and seven-track formats. The interblock gap on nine-track models was 0.6 inch, but remained 0.75 inch on the seven-track models for compatibility.

Magnetic strip direct access storage Magnetic strip direct access devices were introduced in this period by IBM and others [28]. These devices offered very high capacity, moderate seek time, and low cost for both on-line and off-line storage. They suffered from high mechanical complexity compared to disk drives and were displaced by disk drives as disk storage capacity and cost improved.

Though differing in detail, the devices were all similar in that they consisted of a cartridge or cell containing a group of magnetic strips, each about twice the size of a punched card, one of which could be selected and rotated past a movable multitrack read/write head assembly. Spacing between the strip and head bar was determined by a film of air that created a hydrodynamic air bearing.

The IBM 2321 Data Cell Drive [29], delivered in 1966, had an on-line capacity of 400 Mbytes. It consisted of ten removable data cells, for off-line storage, and each data cell contained 200 13 $\times$ 2.25 $\times$ 0.005-inch Mylar strips with magnetic coating on one side and an antistatic carbon coating on the other. The cell was divided into twenty subcells and each of the ten strips in a subcell had a latching slot for picking the strip and a unique coding tab for selection. Each strip had chamfered edges and a "swallow" tail for control of anticlastic curvature and strip dynamics.

Recording density was 1750 bpi and 50 tpi on 100 data tracks on each strip. Strip seek time was 550 ms, and the selected strip was rotated at 1200 rpm to provide an intrinsic data rate of about 55 Kbyps and a maximum latency of 50 ms.

Magnetic disks (1963-1966) This was a very eventful period in the evolution of magnetic disk storage. The hydrodynamic slider developed for the 24-inch fixed-disk configuration led indirectly to the development of a small removable disk pack consisting of six 14-inch-diameter disks that stored 2.68 Mbytes. This innovation provided the first off-line storage capability for disks. The IBM 1311 disk drive [30] was first shipped in 1963 with the IBM 1441. Although a factor of two more expensive per Mbyte than attached tape and 60 times more expensive than tape off-line, it was a significant milestone, for it turned disk drive development in a fruitful new direction.

As the recording density improved from 1025 bpi and 50 tpi (51.25 Kbpsi) in 1963 to 2200 bpi and 100 tpi (220 Kbpsi) in 1966, the cost of disk storage was reduced to slightly less than attached tape, but still a factor of twenty more expensive than off-line tape. Average seek time was reduced by improvements in the hydraulic actuators from 150 to 75 ms. The disk rotational period was reduced from 57 to 25 ms and data rate increased from 69 to 312 Kbyps.

The first microprogrammed intelligent direct access storage control unit, the IBM 2841, was announced with System/360. It contained a processing unit with a transformer read-only storage for microcode control words. Through device-unique microcode, it could provide device-dependent interpretation of channel command words (CCWs) and provide real-time logical and electrical signaling to control devices with widely varying characteristics. It was attached to System/360 via a common (electrical and protocol) I/O interface and a data channel. This control unit design approach allowed implementation of the System/360 disk record format architecture [30] without separate logic circuits for each device type. In this architecture, called count-key-data (CKD), each record contains a count area and an optional key area as a header to the data record. The count area contains several fields of data; among them are the relative record number (on this track) and the length in bytes of the data record. The optional key field contains the record identifier and is used by the control unit to search for a record automatically, using its key as the argument.

Disk storage was no longer a scarce system resource after the delivery of the IBM 2314 File Facility in 1966. It had a new disk pack that stored 29 Mbytes. The packaging was innovative; it consisted of nine disk drives—eight on-line and one spare—and an improved version of the 2841 control unit housed in a single frame. The 233 Mbytes of on-line capacity and the relatively high performance that resulted from overlapping the seek time of its eight actuators [32] mark shipment of this file facility as the turning point for on-line system and application development [26].

New operating systems and data management software were developed in this period to use and enhance the growing capabilities of disks [33, 34]. These new complex systems required substantially more storage capacity than that provided by main memory and their development was made practical by direct access disk storage.

Data management capability [35, 36] provided by OS/360 introduced a level of device independence heretofore unavailable. A new set of access methods replaced IOCS. Two of these, Sequential Access Method (SAM) and Basic Direct Access Method (BDAM), were extensions of IOCS. The other two, Indexed Sequential Access Method (ISAM) and Basic Partitioned Access Method (BPAM), were innovative in that they handled for the user the translation of his logical record identifier into the physical address on disk.

System/360 data management also introduced space management for the storage subsystem. Disk storage space in the past had been managed individually by each user or installation and it was possible for one user to destroy another's data by inadvertently writing over it. While the label processing provided by IOCS did preclude this to an extent, there was no central facility with which to determine the disposition of space within the storage subsystem.

Storage space was divided into subunits called "extents" by System/360. Data management allowed the user to request space in terms of records (or physical attributes such as tracks if he wished) which were then converted into physical extents by the data management routines. A volume table of contents (VTOC) which contained the disposition of all space on a volume (typically a disk pack)—used as well as free space—was recorded on each device.

The first widely accepted multiprogramming system was introduced with OS/360. With its capability for isolating programs from one another through use of memory protection, privileged instructions, priority interrupts, and application programs, as well as its ability to queue jobs for input and output, several unrelated jobs or tasks could occupy the same system. The objective of multiprogramming was to use all of the system resources as heavily as possible. There was little impact on a user if another unrelated program made use of disks when his task was in the compute state. Given several unrelated job streams, the relatively long access time to disk was masked and system throughput improved.

Disks now provided storage for the total on-line data base of systems programs, application programs, and application data files. As the transition from serial batch processing to on-line processing accelerated, file indexing and addressing techniques were improved, and the seeds of modern data base management systems were sown as techniques began to be developed [37] to use a common data base efficiently for a range of purposes.

- *The growth years from 1967 to 1980*

Highlights of the changing environment during this period include the following:

- The cost of a Mbyte of disk storage was reduced by more than a factor of twenty and on-line data processing became the dominant mode in most systems.
- Disk drive design returned to the early fixed-disk configuration; the use of removable disk packs began to diminish as the superior reliability of the new fixed-disk technology was proven and as transaction-oriented processing against the systems-managed data bases increased.
- Streaming tape drives were developed in response to a new role of tape in fixed-disk systems: tapes used as disk save/restore, in addition to systems interchange and archival storage.
- On-line mass storage using cartridges of wide tape in automatically managed libraries was introduced with on-line capacity greater than disks, lower costs than disks, but slower retrieval time.
- Small flexible disks or "diskettes" were introduced in 1971 for microprogram load and evolved into a new medium of system data interchange and the bulk storage medium for small computer systems.
- Smaller rigid disks about eight inches in diameter were introduced to provide low-cost on-line storage for small systems requiring moderate capacity, high performance, and small size.
- Intelligent storage control units were significantly improved in function and reduced in cost by the application of LSI microprocessors.
- Data base management software and new operating systems were introduced that were dependent on and allowed widespread use of disk storage in interactive data base applications.

Magnetic tape (1967-1980) Magnetic tape recording density increased from 1600 to 6250 bpi during this period. Maximum velocity increased from 112.5 to 200 ips and start/stop times improved with simpler mechanical designs. An innovative low-inertia high-torque motor driving a single capstan, with tape supplied from an additional "stubby" vacuum column, provided a new level of start/stop capability with the IBM 2420, first shipped in 1969. Tape velocity of 200 ips could be attained on this drive in less than 2 ms from a dead stop.

A model of the IBM 3420, shipped in 1973, provided further improvements in the drive that reduced the IBG to 0.3 inch, decreased the start/stop time to less than 1 ms, increased the linear bit density to 6250 bpi, and increased the intrinsic data rate to 1.25 Mbyps. This high bit density was achieved by the use of new recording technology and run-length-limited group-coded recording (GCR) encoding [1], while maintaining compatibility with existing tape [38].

Further simplification of the tape-drive mechanics was made possible by the elimination of the high-speed start/stop mechanism in a low-cost, low-performance, servo-controlled reel-to-reel tape drive announced in 1979 with the IBM 4331 and 8100 systems. This drive had two modes of operation: start/stop at 12.5 ips and streaming mode at 100 ips. In streaming mode the tape runs without stopping, provided the CPU can continue to accept the data. This mode of operation recognizes the importance of tape for the functions of disk save/restore as well as those of journaling, backup and recovery, and archival storage.

Mass storage systems (1975-1980) By the mid-1970s the management of very large magnetic tape libraries had become a major problem in terms of cost and space for installations with very large collections of data on tape. These library management costs in many cases exceeded the cost of tape and tape drives. This led to the need for automated tape library storage. Since the standard 10.5-inch tape reel was not convenient for automatic handling, a new medium—the IBM Data Cartridge—was shipped in 1975 with a Mass Storage System (MSS) [39, 40].

The data cartridge was about two inches in diameter and four inches long. It contained 770 inches of 2.7-inch-wide magnetic tape and could store 50 Mbytes. One of many cartridges stored in a honeycomb-like library could be automatically selected and transported in about 10-15 seconds to a data recording device, where the cartridge cover was removed and the tape wrapped around a mandrel containing a rotating read/write head. The tape is recorded in diagonal tracks called "stripes." A unique track-following servo was incorporated to locate a stripe and maintain its position relative to the read/write head.

The MSS combined the low cost of tape with the flexibility of disks by staging large blocks of data (250 Kbytes) on demand into disk storage where normal disk access was available. It thus provided virtual disk storage of from 32 to 472 Gbytes (gigabytes) of on-line data at a cost per Mbyte somewhat less than disks and with staging time in the range of 10-15 seconds.

Rigid magnetic disks (1967-1980) The low cost and the high-performance disk storage required by the IBM System/370 were provided by the development of an improved moving-head disk file (IBM 3330)—for data base storage—and a new fixed-head disk file (IBM 2305) replacing drum storage. Both of these files were announced in 1970 and each had a new microprogrammed control unit. These storage subsystems were also key components in the first widespread implementation of virtual storage techniques [41] announced on System/370 in 1972. In virtual storage systems each application considers itself the occupant of the addressable limits of the system, and pages or segments of data are automatically moved from disk storage to main memory by a combination of hardware and software so that only those segments actually in use will occupy main memory. These systems are heavily dependent on disk storage capacity and performance.

Disk pack capacity was increased from 29 to 100 Mbytes, using densities of 192 tpi and 4040 bpi (776 Kbpsi). Average seek time was reduced from 75 to 30 ms and disk rotational period from 25 to 16.7 ms. Improvements in flying height to 50 microinches and improved magnetic recording technology allowed the increase in bpi; development of the first track-following servo with a voice-coil motor provided the improved track density and seek time. This disk drive, the 3330-1, reduced the costs of on-line disk storage by nearly a factor of three and of off-line disk storage by a factor of two. They were both reduced another factor of two in 1974 when the 3330-11 was delivered with double the track density and double the capacity of the original 3330.

The heads in the fixed-head disk drive represented an advance in ferrite head and slider bearing technology. For the first time, the magnetic element and the slider were designed in an integrated structure using the same magnetic ferrite for the slider and the magnetic core [2]. Each slider contained nine read/write elements. Each of two different models of the IBM 2305 used multiple heads with a total of 768 of these elements: one model with each head on its own track, and one model with two heads per track. While they both had a rotation speed of 6000 rpm, the model with two heads per track reduced the average access time from 5 to 2.5 ms and reduced the storage capacity from 11.3 to 5.4 Mbytes. Data transfer rates were 1.5 and 2.0 Mbyps, respectively.

The two new microprogrammed control units had much in common and contained many design innovations [42] for improved performance, including a writable control memory. In addition to storing control microcode, this memory was used to buffer the count and key fields for error correction; to save counter contents, etc., for error logging and statistics; to accept multiple requests for data on the 2305 and execute them in the sequence of the shortest latency first; and, in conjunction with block multiplexer channels [43] and improved systems software, to allow the implementation of rotational position sensing (RPS). This innovation allowed the disk subsystem to monitor the progress of a data access and remain disconnected from the channel until just before the data passed under the read/write head, thus allowing a significant decrease in channel time required to locate a record.

Another innovation of the 3830 and the 2835 was a small flexible disk known as the 23FD [4] that was used in a read-only mode to load microcode into the writable control memory. Use of this little diskette was destined to grow beyond all expectations, as will be discussed.

The next step in disk evolution came with the development of a lightly loaded slider bearing [44] that carried the read/write head at a flying height of about 20 microinches, eliminated the mechanism necessary to supply the 300-to-400-gram slider preload required by previous designs, and allowed starting and stopping while in contact with the disk. This slider was one of the innovations that led to the announcement of the IBM 3340 "Winchester" disk drive [45] in 1973. Removability of the disk pack in the conventional manner was not possible with these sliders because the sliders depended upon the disk to support them even when stopped; the disks and sliders had to stay together at all times. After some initial studies of techniques to load and unload the new low-mass heads from the disk, the decision was made to feature the capability of the head to start-stop in contact and its planned low cost. The complete head and disk assembly was made a removable unit. The IBM 3348 Data Module was developed incorporating heads, disks, spindle, head arms, and moving carriage. All of the key elements relating to critical tolerances associated with interchangeability of conventional disk packs were incorporated into the removable module—each head read only the data it had written. Significant savings were achieved by elimination of the head alignment procedures in manufacturing and the field [2]. The Data Module was sealed in a shock-mounted plastic enclosure that incorporated an access door, configured like a rolltop desk, for the drive actuator, air system, and electrical connection. Two heads (sliders) per surface reduced the stroke length required, and with improved servo technology [46], the average seek time was reduced from 30 to 25 ms.

The reliability of this new Winchester technology proved to be such an improvement that for the first time no scheduled maintenance was required for a disk drive. Other innovations in the 3340 worthy of note were the following: use of a single integrated circuit located on the access arm to provide the read/write electronics, automatic disk defect skipping, oriented magnetic particles for improved resolution of the magnetic medium, and a fixed-head feature as a part of the basic drive. Data module storage capacity of 35 Mbytes with two disks and 70 Mbytes with four disks was obtained by an areal density of 1.7 Mbpsi with 300 tpi and 5636 bpi.

The removability feature of disks diminished in importance as on-line capacity and reliability increased. Since the inception of the removable disk pack a trend toward fewer disk packs per drive had emerged. For example, with the 1311 the average number of disk packs per drive was greater than 12, with the 2314 it was down to 4, and with the 3330-11 only 1.2. These factors led to the development of a new disk drive with a nonremovable eight-disk spindle storing 317.5 Mbytes. First shipped in 1976, the IBM 3350 had a recording density of 3 Mbpsi (478 tpi and 6425 bpi) using improved Winchester technology. Average seek time was 25 ms. The cost per Mbyte of disk storage was reduced from that of the 3330-11 by a factor of two and by more than a factor of 70 from the original disk drive of twenty years previous.

Another factor of two reduction in disk cost per Mbyte was achieved with the announcement of the IBM 3370 in 1979. Film head technology, a new improved slider with flying height of less than 13 microinches, a new run-length encoding technique [2], and an improved track-following servo allowed 7.7 Mbpsi recording density. Average seek time of 20 ms was achieved with an innovative voice-coil actuator that used a single magnetic assembly for two independent actuators, each providing access to one-half of the data on a fixed seven-disk spindle containing 571 Mbytes of fixed-block storage [47, 48].

IBM announced another innovative disk drive in 1979 that used 210-mm disks (about eight inches). The IBM 3310 disk drive used a simple swing-arm actuator and had a unique track-following servo [49] that obtained its track-following position error information from the data head (by using samples taken between sectors of data) both for data recording and servo data detection.

Average seek time was reduced to 16 ms in 1980 with the announcement of the IBM 3380. Storage capacity was increased to 625 Mbytes per actuator and 1250 Mbytes per spindle by increased areal density. The rental cost of a Mbyte of disk storage was reduced to under one dollar per month.

A new LSI microprocessor [50–52] was at the heart of the new control unit, the IBM 3880, announced in 1979. This control unit provides both fixed-block and CKD (count-key-data) format control, significantly enhances maintenance facilities (through a special maintenance device) [53], and reduces the required number of channel reconnects by means of channel-command stacking. In addition, one model provides speed-match buffering from the high disk data rate to slower data channels.

Flexible magnetic disks (1967–1980) IBM shipped the first flexible disk drive in 1971 for use as a diagnostics and microprogram load device for the 3830 and 2835 Storage Control Units and the System/370 Model 145. The first drive, identified internally as 23FD, had a capacity of 81.6 Kbytes and was a read-only unit. Data and programs were written on a factory-controlled writer, and thus the tolerances involved in interchange between a large number of machines presented no problem. Interchange compatibility, however, was a major design consideration in all later flexible disk developments, as discussed by Engh [4].

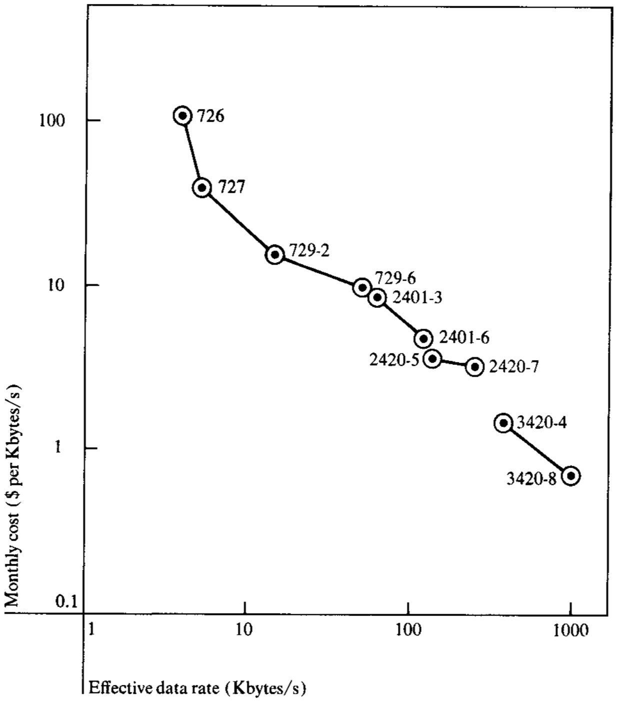

Figure 3 Magnetic tape cost and performance progress: Cost is the ratio of the monthly rental of one device to its intrinsic data rate, given in dollars per month per Kbyte/s. Performance is the effective data rate.

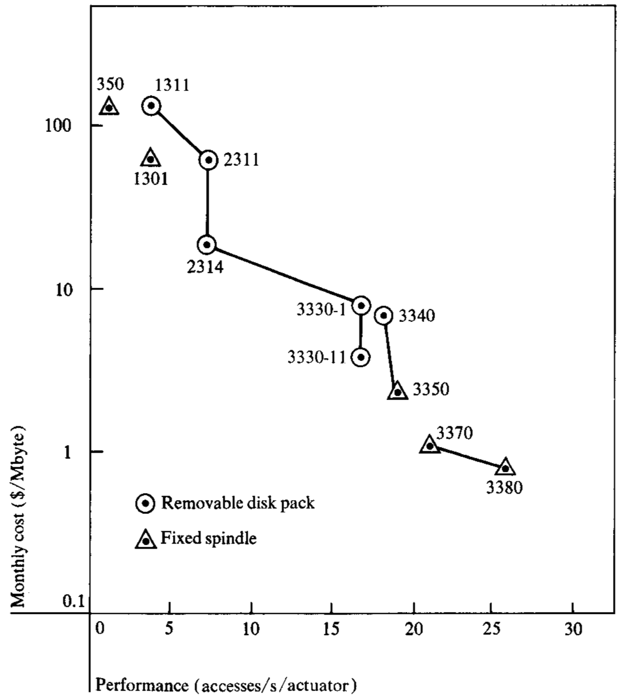

Figure 4 Magnetic disk cost and performance progress: Cost is the ratio of the monthly rental of one device to its capacity in Mbytes. Performance is the number of direct accesses (per actuator) per second deliverable with a device utilization of 67%.

The first diskette with write and read capability was the IBM 33FD. This drive, with a capacity of 243 Kbytes, was first shipped in 1973 as the output medium of the IBM 3740 data entry station. This increase in capacity was achieved through an increase in linear bit density from 1594 to 3268 bpi, an increase in the number of tracks from 32 to 77 (73 usable, 1 index, and 3 spares). Performance was also improved by increasing the speed of the disk from 90 to 360 rpm and the data rate from 33.3 to 250 Kbps. Track-to-track move time was reduced from 333 to 50 milliseconds per track, which reduced the average random access time (one-third of the tracks plus settle time) from 3.6 to 1.3 seconds.

The 33FD found wide acceptance in the industry, and since its introduction a large number of products which use the "floppy" disk have been announced by IBM and other suppliers. Zschu [54] makes the following observation on the importance of these devices:

"For large computer systems it means that keypunches, card handling equipment, key-to-tape and key-to-disk may be replaced with equipment using diskettes. For many minicomputer systems, it means that the cost of the peripherals, which today amount to the lion's share of the system cost, can be substantially reduced. Most importantly, it establishes . . . a medium that can be used as a multipurpose miniperipheral and which, when combined with LSI microprocessors, can do many general and special purpose data processing tasks at a fraction of their cost today."

With the announcement and shipment of the 43FD in 1976, the capability to record on both sides of the disk doubled the capacity to 568 Kbytes. In addition, this new drive reduced the track-to-track move time from 50 to 5 ms and the average random access time from 1.3 seconds to 170 ms. The capacity was again doubled in 1977 with the 53FD by means of a doubling of the linear bit density to 6418 bpi through use of improved heads and MFM encoding. The data rate was also doubled to 500 Kbps because of the increase in linear density. In 1979, the data rate was doubled again by increasing the disk speed to 720 rpm in the 72MD. This drive further increased the amount of on-line storage possible with diskettes by the addition of automatic diskette magazines. Each magazine contained ten diskettes and each drive, two magazines, for an on-line storage capacity of about 24 Mbytes.

Since their introduction in 1973, the read/write capacity of IBM diskettes has been increased by a factor of fifteen, the data rate by a factor of thirty, and the average seek time reduced by nearly a factor of eight. Engh [4] describes the innovations in air bearing control, magnetic head design, electronic circuits, and mechanical design that have permitted these substantial improvements.

Summary

IBM has developed several magnetic storage products that were the first of their kind in that they provided a significant new functional capability for data processing systems. Progress in the evolution of each of these products can be identified with one or more technological innovations which go beyond improvements strictly in basic magnetic recording technology. These include

- The vacuum column tape drive
  - High-torque, low-inertia motor
  - Stubby vacuum column
  - Encoding techniques: NRZI and run-length-limited codes
- The moving-head disk file
  - Hydrostatic air bearing and hydrodynamic slider bearing
  - Mechanical design for disk pack removability
  - Voice-coil motor and track-following servos
  - Data encoding, detection, and clocking circuits
- The flexible diskette
  - Low-cost actuator design
  - Low-wear in-contact head design
  - Diskette materials and packaging for ease of use
- The magnetic cartridge mass store
  - Rotating-head design for digital recording
  - Flexible-media track-following servos
  - Cartridge library storage facility
  - Data staging algorithms

Figures 3 and 4 summarize graphically the overall progress in tape and disk cost and performance. They are plots in cost-performance space of the key devices discussed in the body of the paper and in the subsequent papers on tape and disk storage technology. The points connected by lines are members of the same product family and show the improvements made on a similar technological base.

In addition to the specific developments cited in this paper, some general references are included [55-67] for background on the thread of development of the technologies.

Acknowledgments

The author gratefully acknowledges the assistance of several individuals who contributed to the preparation of this paper. Helpful comments on the contents and organization were made by C. J. Bashe, J. T. Engh, J. M. Harker, J. P. Harris, F. Kurzweil, R. B. Mulvany, M. I. Schor, and R. W. Shomler. C. R. Holleran contributed to the discussion of the disk control units and A. J. Bonner to the sections concerned with systems and data management software development.

References

1. J. P. Harris, W. B. Phillips, J. F. Wells, and W. D. Winger, "Innovations in the Design of Magnetic Tape Subsystems," *IBM J. Res. Develop.* **25**, 691–699 (1981, this issue).
2. J. M. Harker, D. W. Brede, R. E. Pattison, G. R. Santana, and L. G. Taft, "A Quarter Century of Disk File Innovation," *IBM J. Res. Develop.* **25**, 677–689 (1981, this issue).
3. R. B. Mulvany and L. H. Thompson, "Innovations in Disk File Manufacturing," *IBM J. Res. Develop.* **25**, 711–723 (1981, this issue).
4. James T. Engh, "The IBM Diskette and Diskette Drive," *IBM J. Res. Devleop.* **25**, 701–710 (1981, this issue).
5. H. Kobayashi, *Modeling and Analysis—An Introduction to System Performance Evaluation Methodology*, Addison-Wesley Publishing Co., Reading, MA, 1978.
6. P. H. Seaman, R. A. Lind, and T. L. Wilson, "On Teleprocessing System Design, Part IV: An Analysis of Auxiliary Storage Activity," *IBM Syst. J.* **5**, 158–170 (1966).
7. J. P. Buzen, "I/O Subsystem Architecture," *Proc. IEEE* **63**, 871–879 (1975).
8. W. S. Buslik, "IBM Magnetic Tape Reader and Recorder," *Proceedings of the Joint AIEE-IRE-ACM Computer Conference (Review of Input and Output Equipment Used in Computer Systems)*, March 1953, pp. 86–90.
9. C. W. Bachman, "The Evolution of Storage Structures," *Commun. ACM* **15**, 628–634 (1972).
10. V. P. Turnburke, Jr., "Sequential Data Processing Design," *IBM Syst. J.* **2**, 37–48 (1963).
11. R. F. Rosin, "Supervisory and Monitor Systems," *ACM Computing Surv.* **1**, 37–54 (1969).
12. E. F. Codd, "Multiprogramming," *Advances in Computers*, Vol. 3, Academic Press, Inc., New York, 1962, pp. 77–153.
13. N. Rochester, "The Computer and its Peripheral Equipment," *Proceedings of the Eastern Joint Computer Conference*, Boston, MA, 1955, pp. 64–69.
14. M. V. Wilkes, "Historical Perspective—Computer Architecture," *AFIPS Conf. Proc., Fall Joint Computer Conf.* **41**, 971–976 (1972).
15. A. A. Cohen and W. R. Keye, "Selective Alteration of Digital Data in a Magnetic Drum Computer," Contract N6ONR-240 Task 1; Engineering Research Assoc., Office of Naval Research; December 1947.
16. C. E. Frizzell, "Engineering Description of the IBM Type 701 Computer," *Proc. IRE* **41**, 1257–1287 (1953).
17. F. E. Hamilton and E. C. Kubie, "The IBM Magnetic Drum Calculator Type 650," *J. ACM* **1**, 13–20 (1954).
18. T. Leverett, "Floating Drum Head With Retractable Shoe," *Technical Report PH23-00012IM-306*, IBM Corporation, Kingston, NY, 1957.
19. W. A. Stetler, V. G. Spiegel, and R. C. Braen, "Vertical Magnetic Drum With Floating Head," *Technical Report PH23-08001TR-330*, IBM Corporation, Kingston, NY, 1959.
20. M. L. Lesser and J. W. Haanstra, "The Random-Access Memory Accounting Machine—I. System Organization of the IBM 305," *IBM J. Res. Develop.* **1**, 62–71 (1957).
21. T. Noyes and W. E. Dickinson, "The Random-Access Memory Accounting Machine—II. The Magnetic-Disk, Random-Access Memory," *IBM J. Res. Develop.* **1**, 72–75 (1957).
22. R. K. Brunner, J. M. Harker, K. E. Haughton, and A. G. Osterlund, "A Gas Film Lubrication Study, Part III, Experimental Investigation of Pivoted Slider Bearings," *IBM J. Res. Develop.* **3**, 260–274 (1959).

23. W. W. Peterson, "Addressing for Random Access Storage," *IBM J. Res. Develop.* **1**, 130-146 (1957).
24. Werner Buchholz, "File Organization and Addressing," *IBM Syst. J.* **2**, 86-111 (1963).
25. J. D. Aron, "Information Systems in Perspective," *ACM Computing Surveys* **1**, 213-236 (1969).
26. T. H. Bonn, "Mass Storage: A Broad Review," *Proc. IEEE* **54**, 1861-1870 (1966).
27. R. A. Barbeau and J. I. Aweida, "IBM 7340 Hypertape Drive," *AFIPS Conf. Proc., Fall Joint Computer Conf.* **24**, 591-602 (1963).
28. K. R. London, *Techniques for Direct Access*, Auerbach, Philadelphia, 1973, pp. 70-77.
29. A. F. Shugart and Y. H. Tong, "IBM 2321 Data Cell Drive," *AFIPS Conf. Proc., Spring Joint Computer Conf.* **28**, 335-345 (1966).
30. J. D. Carothers, R. K. Brunner, J. L. Dawson, M. O. Halfhill, and R. E. Kubec, "A New High Density Recording System: The IBM 1311 Disk Storage Drive With Interchangeable Disk Packs," *AFIPS Conf. Proc., Fall Joint Computer Conf.* **24**, 327-340 (1963).
31. "Introduction to IBM Direct Access Storage Devices and Organization Methods," *Publication No. GC20-1649*, IBM Corporation, San Jose, CA.
32. J. Abate, H. Dubner, and S. B. Weinberg, "Queueing Analysis of the IBM 2314 Disk Storage Facility," *J. ACM* **15**, 557-589 (1968).
33. C. W. Bachman and S. B. Williams, "A General Purpose Programming System for Random Access Memories," *AFIPS Conf. Proc., Fall Joint Computer Conf.* **26**, 411-422 (1964).
34. C. B. Poland, "Advanced Concepts of Utilization of Mass Storage," *Proc. IFIPS Congress 65* **1**, 249-254 (1965).
35. W. A. Clark, "The Functional Structure of OS/360, Part III: Data Management," *IBM Syst. J.* **5**, 30-51 (1966).
36. G. G. Dodd, "Elements of Data Management," *ACM Computing Surveys* **1**, 117-132 (1969).
37. J. Martin, *Computer Data Base Organization*, Prentice-Hall, Inc., Englewood Cliffs, NJ, 1975.
38. J. A. Rodriguez, "An Analysis of Tape Drive Technology," *Proc. IEEE* **63**, 1153-1159 (1975).
39. C. T. Johnson, "The IBM 3850: A Mass Storage System with Disk Characteristics," *Proc. IEEE* **63**, 1166-1170 (1975).
40. J. P. Harris, R. S. Rohde, and N. K. Arter, "IBM 3850 Mass Storage System: Design Aspects," *Proc. IEEE* **63**, 1171-1176 (1975).
41. P. J. Denning, "Virtual Memory," *ACM Computing Surveys* **2**, 153-189 (1970).
42. G. R. Ahearn, Y. Dishon, and R. N. Snively, "Design Innovations of the IBM 3830 and 2835 Storage Control Unit," *IBM J. Res. Develop.* **16**, 11-18 (1972).
43. D. T. Brown, R. L. Eibsen, and C. A. Thorn, "Channel and Direct Access Device Architecture," *IBM Syst. J.* **11**, 186-199 (1972).
44. M. W. Warner, "Flying Magnetic Transducer Assembly Having Three Rails," U.S. Patent No. 3,823,416, July 9, 1974.
45. R. B. Mulvany, "Engineering Design of a Disk Storage Facility with Data Modules," *IBM J. Res. Develop.* **18**, 489-505 (1974).
46. R. K. Oswald, "Design of a Disk File Head-Positioning Servo," *IBM J. Res. Develop.* **18**, 506-512 (1974).
47. A. D. Rizzi and J. S. Makiyama, "The IBM 3370 Direct Access Storage," *Disk Storage Technology*, 31-33 (1980); Order No. GA26-1665-0, available through IBM branch offices.
48. D. L. Nelson, "The Format of Fixed Block Architecture in the IBM 3370 DAS," *op. cit.* Ref. 47, pp. 34-35.
49. R. D. Commander and J. R. Taylor, "Servo Design for an Eight-Inch Disk File," *op. cit.* Ref. 47, pp. 89-97.
50. J. I. Norris, "A High Performance Microprocessor," *op. cit.* Ref. 47, pp. 27-30.
51. H. Bardsley, "The IBM 3880 Microprocessor Control," *op. cit.* Ref. 47, pp. 69-72.
52. C. R. Holleran, "An Overview of the IBM 3880 Storage Control," *op. cit.* Ref. 47, pp. 62-64.
53. S. J. Duchak, "Maintenance of the IBM 3880 Storage Control," *op. cit.* Ref. 47, pp. 78-82.
54. E. V. W. Zschau, "The IBM Diskette and Its Implications for Minicomputer Systems," *Computer* **6**, No. 6, 21-26 (1973).
55. F. P. Brooks, "Mass Memory in Computer Systems," *IEEE Trans. Magnetics* **MAG-5**, 635-639 (1969).
56. E. W. Pugh, "Storage Hierarchies; Gaps, Cliffs, and Trends," *IEEE Trans. Magnetics* **MAG-7**, 810-814 (1971).
57. R. E. Matick, "Review of Current Proposed Technologies for Mass Storage Systems," *Proc. IEEE* **60**, 266-289 (1972).
58. J. M. Harker and Hsu Chang, "Magnetic Disks for Bulk Storage—Past and Future," *AFIPS Conf. Proc., Spring Joint Computer Conf.* **40**, 945-955 (1972).
59. F. G. Withington, "Five Generations of Computers," *Harvard Bus. Rev.* **52**, 99-103 (1974).
60. U. O. Gagliardi, "Trends in Computer System Architecture," *Proc. IEEE* **63**, 858-862 (1975).
61. K. E. Haughton, "An Overview of Disk Storage Systems," *Proc. IEEE* **63**, 1148-1152 (1975).
62. C. D. Mee, "A Comparison of Bubble and Disk Storage Technologies," *IEEE Trans. Magnetics* **MAG-12**, 1-6 (1976).
63. A. S. Hoagland, "Magnetic Recording Storage," *IEEE Trans. Computers* **C-25**, 1283-1288 (1976).
64. R. E. Matick, *Computer Storage Systems and Technology*, John Wiley & Sons, Inc., New York, 1977.
65. N. C. Wilhelm, "A General Model for the Performance of Disk Systems," *J. ACM* **24**, 14-31 (1977).
66. A. S. Hoagland, "Storage Technology: Capabilities and Limitations," *Computer* **12**, No. 5, 12-18 (1979).
67. A. S. Hoagland, "Trends and Projections in Magnetic Recording Storage on Particulate Media," *IEEE Trans. Magnetics* **MAG-16**, 26-29 (1980).

Received April 3, 1980; revised November 3, 1980

The author is located at the IBM General Products Division laboratory, San Jose, California 95150.