J. P. Harris
W. B. Phillips
J. F. Wells
W. D. Winger

Innovations in the Design of Magnetic Tape Subsystems

This paper explores the selection of magnetic tape as an input/output medium for electronic computers and reviews the innovations in the development of magnetic tape and tape handling machines since their introduction in the early 1950s.

Introduction

In the late 1940s and early 1950s, IBM was engaged in work directed toward the development of stored-program computers. A major consideration in that work was the choice of methods to store and access data characterized by a customer's collection of information files that required processing.

Both magnetic wire and acetate-based magnetic tape were used in the audio industry. Welsh and Lukoff [1] describe the early use of metal tape for digital magnetic recording. Acetate-based tape was used for digital recording by Raytheon at a relatively low tape velocity, as reported by Bloch [2]. At IBM it was decided to pursue an acetate-based plastic magnetic tape for computer input/output because of several attractive characteristics: high data rate, infinite reversibility, low power requirement, low cost, long-term storage capability, compactness, and ease of handling. The use of acetate-based plastic tape was made practical for high-performance systems by the invention of the vacuum column.

Initial development

A small group of engineers working at the Kenyon House at the IBM Poughkeepsie Laboratory in 1949 and 1950 provided the foundation of a long list of products beginning with the IBM 726 tape drive, which was first shipped to customers in 1953.

Prior to the development of the IBM 726, several engineering models of tape drives were built and operated. This work resulted in the development, for later application, of moving coil actuators, the vacuum column, magnetic particle clutches for reel drive, NRZI recording and detection methods, multitrack recording/reading heads, leaders for threading tape, splicing of tape, reliability and defect analysis of tape, rewind methods, and many other items. It is not possible in this short paper to treat all of these areas of technical development as they have progressed over the years. We will attempt to emphasize the most significant developments and to provide references for more detailed information.

In the magnetic tape drive used on computing systems, the primary function of the tape transport is to accept a reel of tape from the library and to spool this medium from the file reel onto a machine reel. The data on the reel of tape are generated in a linear fashion. Ideally those data would be transmitted instantaneously on demand; in real life this is impossible. Since the data are moved in relatively short bursts, the magnetic medium must start and stop as it progresses through the tape transport.

The physical size of the reel and the mass of the full reel of tape dominate the engineer's choice of techniques in designing the mechanisms that transport this medium from one reel to the other.

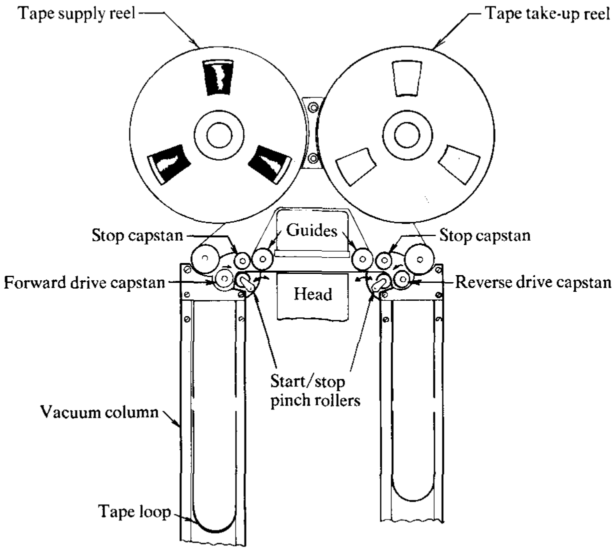

Figure 1 Tape drive utilizing pinch roller.

Given the mass of the components involved, a method of mechanically absorbing the transient condition of starting and stopping this tape was required. IBM was first in using a configuration of columns from which air was partially evacuated. The low-inertia vacuum column buffering technique is described by Weidenhammer [3, 4] and Buslik [3].

IBM 726: The starting point

The IBM 726 tape drive [5], announced in 1952 as part of the IBM 701 system, was the first IBM tape drive product and is used in this paper as the base reference in discussing significant changes over time. It incorporated the vacuum-column and pinch-roller techniques, and recorded digital data in seven parallel tracks on magnetic tape. Six of these tracks represented data bits; one was a parity bit. The six data bits represented a "copy group" on the IBM 701. The terminology for groups of six bits was changed to "character" with the advent of the IBM 727 and 702. The recording density was 100 bytes per inch (bpi) and the velocity 75 inches per second (ips), yielding a data rate of 7500 characters per second (cps). The tape was driven by two capstans (forward/reverse) which were belt-driven and rotated constantly. One or the other of these capstans was engaged by activating a moving coil forcing a pinch roller against the capstan through a mechanically biased four-bar linkage.

The medium chosen for the IBM 726 tape drive was a derivative of the magnetic tape being used in the audio industry at that time, an acetate film coated with fine iron oxide particles in a binder solution. A one-half-inch-wide version of this acetate film tape was obtained from the Minnesota Mining & Manufacturing Corporation for use with the IBM 726.

Splicing tape was used to attach the end of the tape to a steel leader when a new reel was loaded on the unit. The splicing tape was also used to mend tapes in case of breakage.

The reels were mounted on hubs which were driven by an ac motor and a magnetic particle clutch mechanism located behind each reel hub. The tape was spooled from the reel hub into the vacuum column, where a series of vacuum-sensing switches were activated, depending on the position of the loop of tape in the associated column. These switches served as the control for engaging or disengaging the clutch mechanism. Thus, each vacuum column had an independent motor and clutch system.

As the tape moved across the single write/read head located between the two vacuum columns, information was written on or read from the tape, as indicated in Fig. 1. The write/read head contained seven 0.032-inch-wide tracks, each with a gap of 250 microinches. This head was preceded by a two-gap dc erase head. The erase head gaps were designed to erase the medium by bringing it close to zero magnetization to minimize distortion.

The head-to-tape interface (HTI) was formed by a recording head contoured with a small radius (one-sixteenth inch) between two flat surfaces with a 168-degree included angle. The tape was guided over the head by two rods, one on either side of the head. A pressure pad, spring-loaded with six ounces across the tape width, pressed down on the tape at the write/read gap. This pad was needed during writing/reading operations over spliced tape to ensure that the tape stayed in contact with the head. The spliced areas were less pliable and tended to pull the tape away from the head.

Vacuum tubes were used for logic as well as for powering the write circuits and for amplifying the read signal.

The bits were recorded using NRZI code (Non-Return to Zero IBM) [6], illustrated in Fig. 2. The Phelps patent shows a "yes" track and a "no" track which guaranteed a transition per bit to allow self-clocking. Later, this was modified for the IBM 726 by recording only "yes" signals and using odd parity across seven parallel tracks to ensure at least one bit transition per copy group. A record started with the write heads at maximum current in a bias position. A "one" in any of the bit positions reversed the current, and thus reversed the tape magnetization in that track. Any subsequent "one" caused a reversal of current. Thus, "ones" were represented by a change in polarity of magnetization on the tape.

During the read operation, the first pulse from any of the seven tracks started a clock (a single-shot multivibrator). The number of pulses in each character varied from one to seven. All pulses did not occur at the same time when read because of both mechanical and electrical skew. All bits detected during the clock interval were assumed, therefore, to be associated with one copy group.

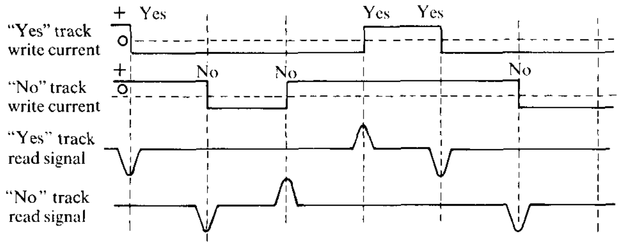

Figure 2 The NRZI code, in which the "yes" track and "no" track guarantee a translation per bit to permit self-clocking.

Errors were detected by checking proper parity for each character across the tape. This check occurred only during a read operation.

System performance improvements

As system development has progressed from the IBM 701 and the associated 726 to today's complex systems, the requirements and applications of magnetic tape storage have changed. In the machines of the mid-1950s, magnetic tape was used primarily as a data processing medium in a batch mode. As time passed, magnetic tape assumed the function of data interchange among systems and that role remains of primary importance today.

As random access storage was introduced and has continued to improve in cost and capacity, the data processing function has migrated toward use of random access storage, while the functions of backup and recovery, archival storage, and interchange have increased in relative importance as uses for magnetic tape.

In the early processors of the 1950s, central memory was very costly (the IBM 701 memory stored only 2000 36-bit words). The system performed one job at a time. Typically only one input/output device could operate at a time. The tape control functions could be incorporated in the central processor (IBM 701) or in simple separate controllers (IBM 702/752). These separate controllers could interpret the macro instructions of the processor and do some minor overlap operations (such as rewind), but were limited to doing a read or a write operation on only one of a string of drives attached to each controller.

Even though magnetic tape provided a relatively high data rate, because of parallelism a speed mismatch still existed between those data rates and the ability of the central processor to use data. An early attempt to relieve this mismatch and to provide some overlap operation was the IBM 777 Tape Record Coordinator [7]. This product, delivered in 1956 along with the IBM 705 Model II, performed the functions of a tape controller. In addition, it contained a memory which buffered the main input/output channel from the tape drive read or write operations, thus freeing the processor to do other things. The buffer memory was small (1024 characters) and programming was difficult, but the IBM 777 improved overall system performance and was an effective product.

As main memories became more economical and processors were equipped with multiple channels, it became attractive to perform the equivalent of buffering in the main memory. This type of buffering was much more flexible than control unit buffering, but it added to complexity of software and required the central processor to spend more of its time managing data input and output.

The IBM System/360 provided for standardization of input/output channels and software control over many central processors. With the software, the increased size of memory, and the expanded capabilities of the central processor, several jobs could run simultaneously, thus adding to the requirements of the input/output system. The flexibility of system configuration offered by the stand-alone controller made it popular over the middle to upper range of central processors. In addition, a new function became important—that of switching tape drives among controllers. It is characteristic of magnetic tape that one tape unit is required for each open data set. As central processors became capable of doing several jobs simultaneously, switching enhanced the load-balancing capability and minimized the number of drives required.

Magnetic tape system performance improvements

As system performance improved with the introduction of new, faster systems, there were corollary requirements for improvement of magnetic tape system performance. The primary areas of improved performance have been data rate, data reliability, and tape start time.

Table 1 Improvements in data rate in magnetic tape systems, 1953-1973.

Tape system, IBM model	*Year*	*Data rate*		
		(in./s)	(characters/in.)	(characters/s)
726	1953	75	100	7 500
727	1955	75	200	15 000
729-III	1958	112.5	556	62 550
729-VI	1962	112.5	800	90 000
2401-6	1966	112.5	1 600	180 000
3420-7	1971	200	1 600	320 000
3420-8	1973	200	6 250	1 250 000

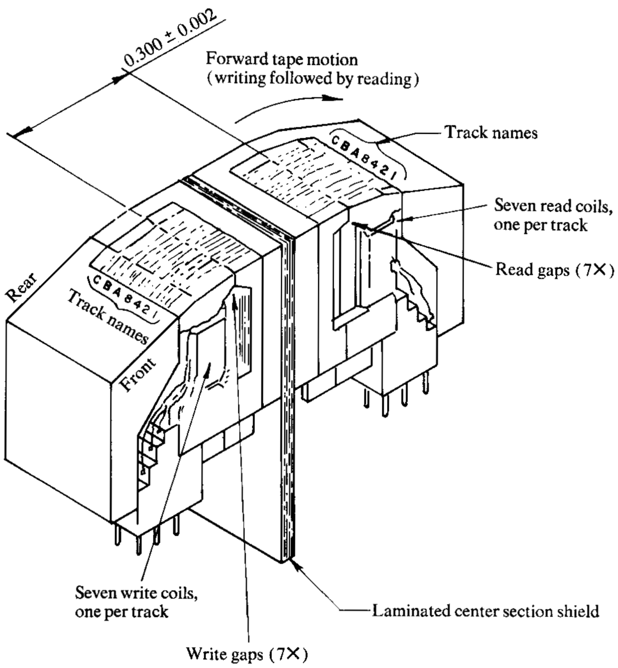

Figure 3 The two-gap, seven-track head used in the IBM 729. The combination of separate read and write heads permitted read checking immediately after writing.

Data rate

The most effective means of improving data rate has been to increase linear density. The IBM 726 was shipped in 1953 with a linear density of 100 characters per inch. The IBM 3420 Models 4, 6, and 8 were first shipped in 1973 with a linear density of 6250 characters per inch. Each improvement in density required a major technical effort. Each was a major advance in performance capability. Beginning with the doubling of density to 200 bpi, each led the industry at that time in providing the highest data rate available with half-inch tapes. The improvements in data rate are indicated in Table 1. Beginning with the advance to 556 bpi, each increase in density was accompanied by backward compatibility. The tape drive and controller which produced a new, higher density could both read and write the older, lower density, thus providing a smooth transition for the user to the new density.

The first density increase came with the First Customer Ship (FCS) in 1955 of the IBM 727 at 200 characters per inch. This density increase was made possible by decreased tape skew in the mechanical transport.

The IBM 727 design also used a reel that was approximately ten and one-half inches in diameter and contained 2400 feet of tape. This reel size was to achieve widespread use in the industry for the next twenty-five years. Tape was threaded by the operator through the tape path and onto the machine reel, eliminating the leader/splicing process. The capstans were attached to the retractable shafts of two ac motors (one forward, one reverse). The capstans were retracted for ease of threading. The pinch rollers were driven by a single moving coil assembly which was biased forward or reverse by a solenoid.

Concurrent with the increases in linear density, there have been a succession of significant developments in the reliability of recording through verification, checking, coding, and recording techniques.

A longitudinal redundancy check was added with the IBM 727. This check ensured that an even number of bits was recorded and read back in each track, thus substantially enhancing error detection capability over the single-character parity system. The concept of write checking by checking the presence of the "inductive kick" on each write head was also introduced; however, bit-for-bit checking was done only in a read-back mode.

It was not uncommon for a user to verify his most critical data by a write, back-space, read operation before he wrote the next record, effectively reducing the performance of the writing drive to one-third normal when verification was required.

The two-gap head, an IBM innovation [8], was introduced in 1958 as a part of the IBM 729 tape drive. Separate read and write heads (Fig. 3) were packaged together to allow read checking immediately after writing. This "on-the-fly" verification capability provided a three-to-one improvement in performance on the writing drive where verification was important. The head was designed as a single unit with intricate track-to-track and gap-to-gap shielding and separate cabling and circuits for writing and reading.

The key limitation in recording NRZI self-clocking information was the variable arrival times of bits within a byte caused by mechanical and electrical skew. The increase in density from 200 to 556 bpi was made possible by a further reduction of mechanical skew in the tape transport [9] and the introduction of delay lines to the write and read processes to minimize electrical and mechanical head tolerances. Improved detection circuits reduced the effect of random noise and increased the effectiveness of the read-after-write checking. The introduction of transistors in 1959 further reduced the noise, the circuit tolerances, and package size.

The use of pulse narrowing or slimming [10] to reduce bit interaction pushed the NRZI bit density to 800 bpi in 1962 on the 729 models. Mechanical improvements including spring-loaded tape guides helped to solve the skew problems at this density.

NRZI recording had always been in a saturation mode; *i.e.*, the tape was fully magnetized to maximum surface signal output of its magnetic material. Even though the tape was saturated, the read signal was derived predominantly from surface magnetization. The read output detection system experiences exponentially decreasing contributions from distant signals.

By limiting the write head flux output, the tape could not be saturated but higher resolution could be obtained. This process of surface recording [11] resulted in reduced intersymbol interference (phase shift or peak shift).

The cyclic redundancy check (CRC) was introduced in 1965 with the IBM 2401 Models 1, 2, and 3. It resulted in dramatically improved error detection over that achieved by parity checking. The CRC also allowed error correction on the reread operation when combined with parity checking [12]. In the 800-bpi mode, the CRC contains eight bits. All error patterns of seven bits or less will be detected. The parity redundancy check, vertical across the byte (VRC) and longitudinal along the track (LRC), would detect all odd-bit errors and many even-bit errors. The combination of the VRC, LRC, and CRC allows limited single-error correction and gives very powerful error detection. On readback, a single error would be corrected and checked via CRC. Although CRC was previously used in serial systems such as disk and drum recorders, the implementation in a parallel system was unique.

Pushing linear density to 1000 bits per inch and above looked formidable indeed. The problems associated with skew caused by the tape path and the tolerances of the electromechanical system seemed very costly to solve.

The breakthrough in linear density for parallel-track systems was made possible by the use of self-clocking tracks with electronic de-skewing [13]. Self-clocking tracks have redundant bits of a periodic nature in the data that generate a clock per track. Bits from each track of data are stored and when a byte is complete it is transmitted. Mechanical and electrical skew are compensated. The designer can then trade off code redundancy and electronic buffer cost with mechanical tolerances. In self-clocking systems, the limitations of density become intersymbol interference, clock stability, and signal-to-noise ratio. Densities of 3000 up to 9000 flux changes per inch have been implemented with this technology. Self-clocking tracks with electronic de-skewing were used in IBM with 1600-bpi phase encoding recording shipped in 1966 on the IBM 2401 Models 4, 5, and 6. Data density of 1600 bpi requires a recording capability of 3200 flux changes per inch because of the less efficient nature of phase encoding recording.

- *Reliability*

The quest for higher density is focused on high-resolution recording. Head-to-tape separation is minimized at the expense of head and tape wear. Because surface recording allowed use of thick tape, tape wear has never been a problem. The air-film-lubricated head was the key to uniformly small spacings with a minimum of wear. Slotted air bearing/heads [14] allow stable contours for minimum head-to-tape spacings.

As the density increases, smaller defects on the tape cause head-tape separation, resulting in signal loss and distortion. Conventional error correction encoding (for example, Hamming code) using several redundant bits per byte would allow single-error correction and double-error detection in each byte. More efficient codes allowed significant improvements in reliability.

Group Coding Recording (GCR) allows an error in any one track and in most cases errors in any two tracks to be transparent to the user [15, 16]. The GCR technique uses redundancy to correct burst multitrack errors. Because of the redundancy, a data density of 6250 bpi requires in excess of 9000 flux changes per inch. GCR has periodic resynchronization of each track's clock. The resynch allows multiple errors in a record to be corrected. The GCR is capable of correcting any single track in error and any two tracks in error when code conditions or signal anomalies indicate track-in-error. Within a typical record of 4096 bytes, the error correction is reset or resynched every 128 bytes. This results in high reliability even with degraded tapes. The GCR technique allowed the 6250-bpi system, as used on IBM 3420 Models 4, 6, and 8, to use the tapes in existing customer libraries and still provide improved reliability.

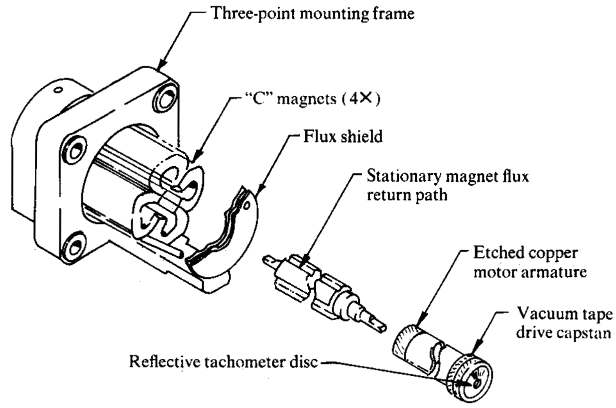

Figure 4 Exploded view of the high-torque, low-inertia motor.

• *Media*

A very important aspect of reliability improvement is associated with the tape itself. The acetate-based magnetic tape introduced for digital recording with the IBM 726 was adapted from the audio recording industry. It quickly became apparent that improvements in the medium were required both for recording reliability and for mechanical handling.

Two important limitations in acetate tape produced a growing concern: first, its excessive sensitivity to temperature and humidity changes; and second, its low ultimate strength, which, in high-speed rewind, could cause tape breakage. In 1956, IBM changed its substrate material from an acetate film to a polyethylene terephthalate film (du Pont trademark Mylar), which provided a substantial improvement in both of these problem areas.

In February 1959, IBM announced a new tape coating binder system called Durexcel, which markedly improved the durability of the magnetic coating material. This new coating proved to be very successful with the relatively gentle tape handling of the moving coil and linkage actuator system of the IBM 727. When the IBM 729 Model III was introduced with its prolay actuator device and relatively more aggressive tape handling and acceleration, the durability of the Durexcel tape was not adequate. Therefore, a more durable binder system was incorporated into Heavy Duty Tape, which was introduced in February 1961.

In 1967, IBM introduced for the first time a tape manufactured and tested in its own facility. This tape, called Series 500, offered improved binder durability and improved write error check rates due to total surface testing. Multi System Tape (MST) followed in 1974 with a new urethane binder system which showed further improvements in reliability.

• *Start time*

Along with the improvements in recording and media, there have been many innovations in the tape transport mechanism.

The increase in data rate demanded by the systems required an improved response in the electromechanical system. The 729 tape drive, announced in 1957, provided both increased data rate and improved electromechanical performance.

The moving coil and biasing mechanism originally used to drive the tape at the required velocity was replaced by a pair of three-position magnet assemblies that moved the pinch rollers to engage the tape to the appropriate capstans. These simplified mechanisms, called prolays, provided two enhancements: a much-improved response time and a significant reduction in mechanical components and complexity in the tape path. This actuator technology, first incorporated in the 729 III, was also used in the following 2400 family, the tape family introduced with the IBM System/360 in 1964.

The most significant technical development in tape drive technology since the vacuum column, however, was the development of the high-torque, low-inertia (HTLI) motor coupled directly to the drive capstan.

In 1965, a development effort was started at IBM directed toward the goal of providing the most outstanding tape drive technology attainable at that time. The result of that technology effort was a family of two tape drives, the IBM 2420 Models 5 and 7. In them, the belt, capstan, and pinch roller systems were eliminated by using dc motors as the driving components for the reel hubs and a unique, high-performance dc HTLI motor that had a drive capstan directly attached to the tubular moving coil armature which was constructed with a stationary magnetic return flux path.

The challenge for engineers at that time was to obtain a dc motor that could directly drive the tape at 200 ips and start and stop in a matter of a few milliseconds. IBM designed and developed a high-torque low-inertia motor which has become a standard in tape drive designs since its inception in the late 1960s. This motor, shown in Fig. 4, was designed to accelerate tape from a linear velocity of zero to 200 ips in less than two milliseconds. Successive improvements in acceleration are shown in Table 2.

Table 2 Improvements in effective acceleration of various tape drives.

Tape system, IBM model	*Year*	*Terminal velocity* (in./s)	*Effective acceleration* (G)	*Actuator*
726	1953	75	25	Moving coil and 4-bar linkage
729-I	1957	112.5	55	Prolay and pinch roller
729-III	1958	112.5	62	Prolay and pinch roller
2420-7	1965	200	400	Low-inertia motor with capstan as an extension of the armature
3420-8	1973	200	520	

The mass of the tape being accelerated from one vacuum column to the other became a deterrent to achieving this very rapid start-and-stop acceleration. To solve that problem, "stubby" columns were introduced. "Stubbies" are an additional set of very small vacuum columns placed in the tape path on either side of the capstan motor and tapered in a fashion that provides self-positioning of the tape in the stubby. With this concept, the additional element of buffering provided by the stubbies allowed for very rapid acceleration of the lower mass of tape in the area of the read/write head.

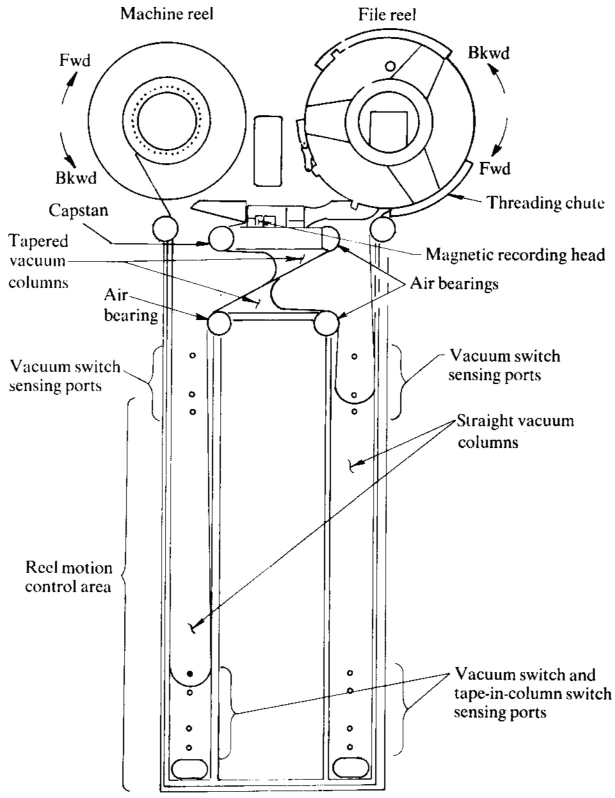

Figure 5 Tape path for the IBM 2420 Model 7.

The 2420 family of tape drives, of which the tape path is illustrated in Fig. 5, also provided almost double the linear velocity, and with this a very significant improvement in the data rate to the system. The increase in linear velocity required refinement in the reel control system. Each reel was driven directly by dc motors. The power to these motors was varied continuously depending on the position of the tape in the columns. (The clutches in the previous 2400 family were either "on" or "off" depending on the position of the tape in the columns.) In addition, a pneumatic system threaded the drive automatically when the operator loaded a new reel.

To summarize, the 2420 tape family introduced in 1969 eliminated the clutch and pulley systems for driving the reel motors; eliminated the separate capstan motors and pinch roller mechanism used to drive the tape from one vacuum column to the other; expanded the role of the vacuum columns by providing an additional stage of buffering; and introduced self-threading for operator convenience.

- *Subsystem*

A new magnetic tape subsystem was designed and introduced in 1971 as the IBM 3420 family of tape units with the IBM 3803 control/switch unit [17]. The three tape drives operate at 75, 125, and 200 ips. The controller operates under microprogram control using a read-only memory.

The innovations in the controller include solid state switching among drives and controllers. Drive switching was enhanced by a reduction in the number of lines between the drive and the controller. This was made possible by time multiplexing. Serviceability and reliability were enhanced by many diagnostic and service aids. For example, the number of sense bytes (status indicators) was increased by a factor of four. Also, the read signals are digitized in each tape drive as opposed to using an analog read bus.

First shipped in 1971 with linear densities of 800 and 1600 bpi, the subsystem was improved in 1973 with the introduction of 6250 bpi and group-coded recording. At the same time, the interblock gap (IBG) was reduced from 0.6 inch to 0.3 inch with a corresponding improvement in access time and reel capacity. The IBG reduction was made possible by further refinements in the high-torque, low-inertia motor/capstan design to provide one-millisecond acceleration time to attain 200 ips.

Table 3 Tape drive performance improvements.

Tape system, IBM model	*Performance* (characters/s/dollar of monthly rental)
726	17.6
727	27.3
729-III	69.5
729-IV	94.8
2401-6	216
2420-7	304
3420-7	478
3420-8	1 440

• *Mass storage system*

As users became more sophisticated and system hardware and software offered greater capability, the requirements for large-scale storage of data became more complex. With the advent of MVS (Multiple Virtual Storage), many jobs could be run simultaneously. In turn, all the data sets necessary for those jobs were required on line. The volume of data required on line grew very rapidly. Customers faced a formidable task in managing tape reels as reels were moved from the library to the processing area and back to the library. Remote job entry applications faced long delays as the correct reels were sought out and loaded on drives.

Conventional magnetic tape can provide the large capacity needed but has two fundamental deficiencies. Manual intervention is required to load each reel on a drive at an appropriate time and a tape drive is required for each open data set.

Direct Access Storage Devices (DASD) can provide on-line storage so that less manual intervention is required (none if disk packs are not removable). DASD may also contain multiple data sets, but storage of large amounts of data costs more with DASD than with tape.

In 1975, the IBM 3850 Mass Storage System was shipped to customers [18]. Although the MSS continues the use of magnetic tape as the medium for data storage, the format of the tape differs from that previously discussed here. The tape is 2.7 inches wide and 794 inches long, and is stored in a cartridge that is 1.86 inches in diameter and 3.49 inches long. The tape was developed for this application and is manufactured specifically for it.

Cartridges are stored in a honeycomb-like structure and a controller catalogs the volumes and their locations. The cartridges are placed in and fetched from their cells by a mechanism which transports them to a read/write station where they are automatically loaded into a drive. The time to fetch and load a cartridge is several seconds.

Data are recorded serially along diagonal stripes across the tape. Each stripe contains 4096 bytes. At the beginning and end of each stripe is a permanent record of servo and stripe address information. The diagonal stripe is created by wrapping the tape in a helix around a cylindrical mandrel. A read-write head rotates in a plane perpendicular to the axis of the mandrel. The tape is stationary while the head rotates along the stripe. The tape is then stepped along a helical path to the next position, where the adjacent stripe is recorded.

The average data rate per revolution from one drive is 374 000 bytes per second. The head-to-tape interface is noncontact so that the tape in a cartridge can be reused indefinitely.

Cartridge capacity is 50 million bytes and the capacity of the system ranges from 35 to 470 billion bytes, depending on its configuration.

The system is architected such that it operates as a virtual DASD. That is, the central processor calls for a data set from DASD through the 3850 mass storage controller. If the data set is not found in the index to be already loaded on DASD, the appropriate cartridge is fetched and loaded, and data are transmitted to DASD and thence to the central processor. Data are also staged down to the 3850 either on command or according to selectable storage algorithms. Most of the operation is transparent to the user.

With this system, increased on-line storage capacity can be provided in a manner which is advantageous for many applications in terms of both performance and cost. The system solves the library-maintenance and multiple-operator problems. It saves the user time and people, it places the management of data and scheduling under system control, and it serves to supplement both conventional tape and DASD capabilities.

• *Streaming mode*

In 1979, the first IBM 8809 tape drive was delivered. This is the first drive designed for streaming mode *vis-à-vis* conventional start/stop processing. In streaming mode, the drive operates at a velocity of 100 ips at 1600 bpi. The data rate is 160 000 bytes per second. In start/stop mode, the drive operates at 12.5 ips and produces a standard 0.6-inch interblock gap. The user can use start/stop mode for batch processing, or switch to streaming mode and achieve eight times the data rate for dumping large blocks of data onto magnetic tape.

A further innovation is that the tape is driven directly from reel to reel with no vacuum columns. A system of transducers and servos is used to control tape velocity and tension.

Summary

The IBM 726 magnetic tape system was introduced along with the IBM central processor in 1953. Since then, the linear density has been improved by 62 times, the data rate has been increased by 167 times, and reel capacity has increased 32 times. Table 3 illustrates the improvement in performance over the years. In the same time frame, the use of magnetic tape has evolved from primarily input and output for batch processing to a much greater role for data interchange, dump/restore, and archival storage.

There is every reason to believe that the next decade will see further significant advances. The limits of magnetic recording technology have not been approached. Laboratory experiments have yielded linear densities in excess of 100 000 bpi, 16 times the present 6250 bpi used in products. Head technology is capable of track densities many times the 18 per inch now used. While the form of further advancement is not yet clear, the technology can clearly carry us much further.

References

1. H. F. Welsh and H. Lukoff, "The Uniservo-Tape Reader and Recorder," *Proceedings of the Joint AIEE-IRE-ACM Computer Conference* (Review of Input and Output Equipment used in Computer Systems), New York, Dec. 10-12, 1952, pp. 47-53.
2. Richard M. Bloch, *Proceedings of the Second Symposium on Large-Scale Digital Calculating Machinery*, Harvard Computation Laboratory, September 13-16, 1949, Harvard University Press, 1951.
3. J. A. Weidenhammer and W. Buslik, "Tape Feed Mechanism," U.S. Patent 3,057,568, 1962.
4. J. A. Weidenhammer, "Tape Feed Mechanism," U.S. Patent 3,057,569, 1962.
5. W. S. Buslik, "IBM Magnetic Tape Recorder," *Proceedings of the Joint AIEE-IRE-ACM Computer Conference* (Review of Input and Output Equipment used in Computer Systems), New York, Dec. 10-12, 1952, pp. 86-99.
6. B. E. Phelps, "Magnetic Recording Method," U.S. Patent 2,774,646, 1956.
7. W. D. Winger, P. W. Jackson, and V. R. Witt, "Data Transmission System," U.S. Patent 2,978,678, 1961.
8. V. R. Witt and R. C. Bradford, "Magnetic Transducer," U.S. Patent 2,922,231, 1960.
9. R. A. Skov, "Pulse Time Displacement in High-Density Magnetic Tape," *IBM J. Res. Develop.* **2,** 130-141 (1958).
10. H. M. Sierra, "Increased Digital Magnetic Recording Readback Resolution by Means of a Linear Passive Network," *IBM J. Res. Develop.* **7,** 22-33 (1963).
11. R. A. Barbeau, R. C. Bradford, N. R. Fraim, and D. E. Lockett, "Gap Length-Limited Saturation Depth Recording," U.S. Patent 3,449,756, 1969.
12. D. T. Brown and F. F. Sellers, Jr., "Error Correction for IBM 800-bit-per-inch Magnetic Tape," *IBM J. Res. Develop.* **14,** 384-389 (1970).
13. W. B. Phillips, "A Systems Approach to Magnetic Tape Drive Design," *Proceedings of the IEEE International Convention*, New York, March 24-27, 1969, pp. 273-274.
14. H. K. Baumeister and V. Nejezchleb, "Hydrodynamically Air Lubricated Magnetic Tape Head," U.S. Patent 3,170,045, 1965.
15. A. M. Patel and S. J. Hong, "Optimal Rectangular Code for High Density Magnetic Tapes," *IBM J. Res. Develop.* **18,** 579-588 (1974).
16. S. J. Hong and A. M. Patel, "Plural Channel Error Correcting Apparatus and Methods," reissue patent, U.S. Patent RE30,187, 1980.
17. J. W. Irwin, J. V. Cassie, and H. C. Oppeboen, "The IBM 3803/3420 Magnetic Tape Subsystem," *IBM J. Res. Develop.* **15,** 391-400 (1971).
18. J. P. Harris, R. Rhode, and N. Arter, "The IBM 3850 Mass Storage System: Design Aspects," *Proc. IEEE* **63,** 1171-1176 (1975).

Received May 8, 1980; revised November 24, 1980

J. P. Harris is located at the IBM General Products Division headquarters, 5600 Cottle Road, San Jose, California 95193. W. B. Phillips and W. D. Winger are located at the IBM General Products Division facility, 9000 South Rita Road, Tucson, Arizona 85744. J. F. Wells is located at the IBM System Products Division laboratory, 1001 W. T. Harris Boulevard, Charlotte, North Carolina 28257.

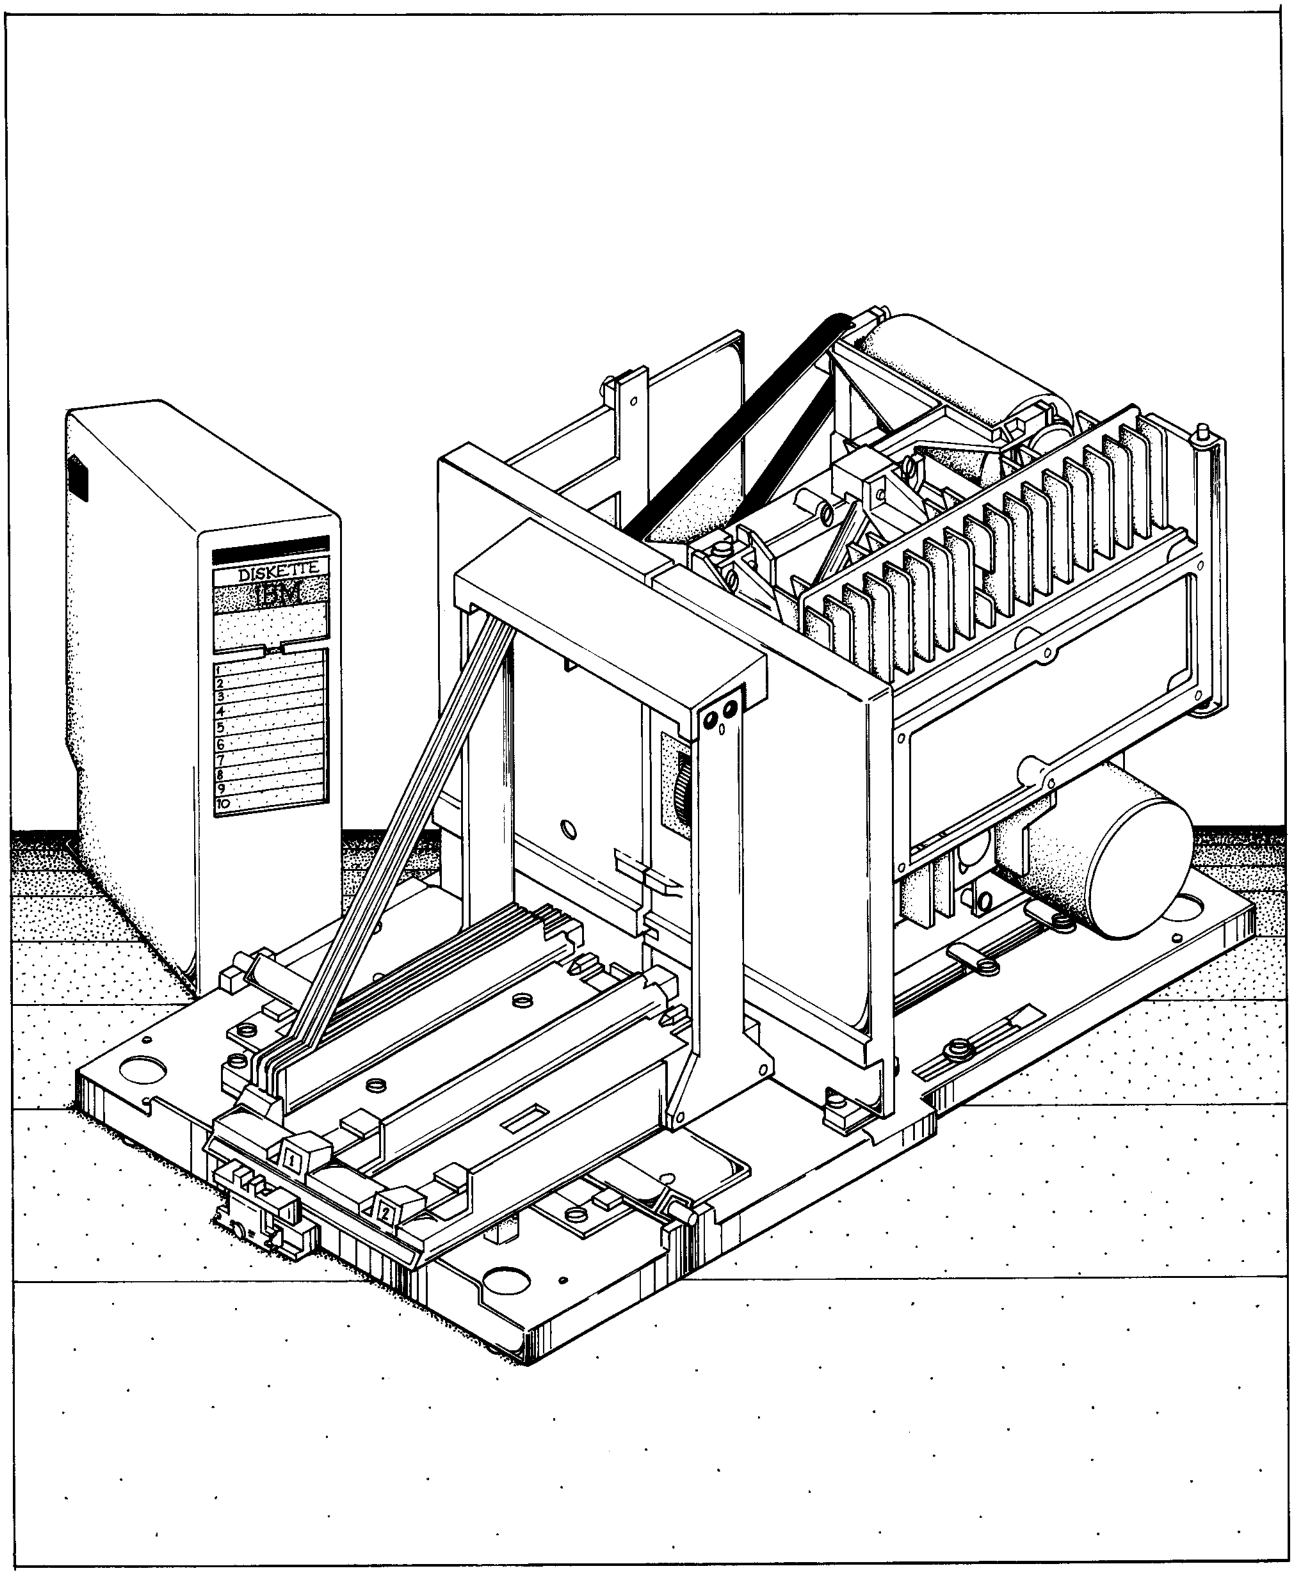

Diskette assembly