James T. Engh

The IBM Diskette and Diskette Drive

The diskette and diskette drive have had a major influence on data processing. They provide a low-cost, compact, high-performance solution to the need for a reusable magnetic medium and have largely replaced the punched card in many applications. Early applications were simple program-load functions. Today these have expanded to a wide range of medium exchange, information storage, and data processing applications. This paper examines the history of the development of these products within IBM. The discussion includes some of the alternatives considered and some of the problems encountered during these developments.

Introduction

In 1971, IBM started producing computer systems using a new form of magnetic storage medium. The medium, ultimately named "diskette," would in a few short years find worldwide usage and contribution to the design of smaller and lower-cost systems. The diskette was a flexible magnetic disk enclosed in a jacket measuring eight inches square and one-sixteenth inch thick and weighing a few ounces. The original application was for loading control microprograms and diagnostics. Since then, applications have broadened so that it now serves as a distribution, exchange, and archival medium for data, programs, microcode, diagnostic procedures, and other digital information, and also as a removable disk in system files. This growth in applications has resulted from the economy and versatile function of the diskette together with the low cost and DASD-like performance of the diskette drive. These attributes may be traced back to the early development work on the IBM 23FD [1] diskette drive.

Early development

The introduction of semiconductor memory created a need for a low-cost device to load programs and information. The requirement was for a read-only device and, at minimum, a serial tape-like medium would have been acceptable. It was recognized, however, that direct access to a section of data and the ability to read a record repetitively would be useful in some applications. This encouraged a search for an alternative form of medium. Initially, a capacity of 65 kilobytes (65 kbytes, where k = 1000) was sought. Other characteristics that were requested included a medium that could be replaced without tools within fifteen seconds and that was small and light-weight to permit inexpensive shipment. A data transfer rate of 15–30 kilobits per second (kbps) was proposed, along with no more than five seconds for access to a record.

A number of technologies were considered, including rigid disk files, magnetic belts and disks similar to those found in dictating machines, a stretched membrane file then being developed in the IBM Los Gatos Development Laboratory, standard 45-revolution-per-minute phonograph records, and audio tape cartridges. None of these approaches was completely satisfactory. D. L. Noble, a senior engineer in the IBM San Jose Development Laboratory, performed these evaluations during 1967 and, after recognizing the individual shortcomings, proposed a file having a removable flexible disk. He and others [2] contributed to the development that followed.

The early concept called for an nonenclosed flexible disk clamped at its center to a vertical turntable, similar to a phonograph. This disk would use a substrate or base material of suitable stiffness to allow it to be inserted and clamped. However, material of this stiffness was not readily available and so more flexible material, used for half-inch magnetic tape, was selected. This was a 1.5×10^{-3}-inch substrate coated on one side with $\approx 0.5 \times 10^{-3}$ inch of magnetic material. The magnetic particles were randomly oriented rather than magnetically oriented in a linear direction as is normally done with half-inch tape.

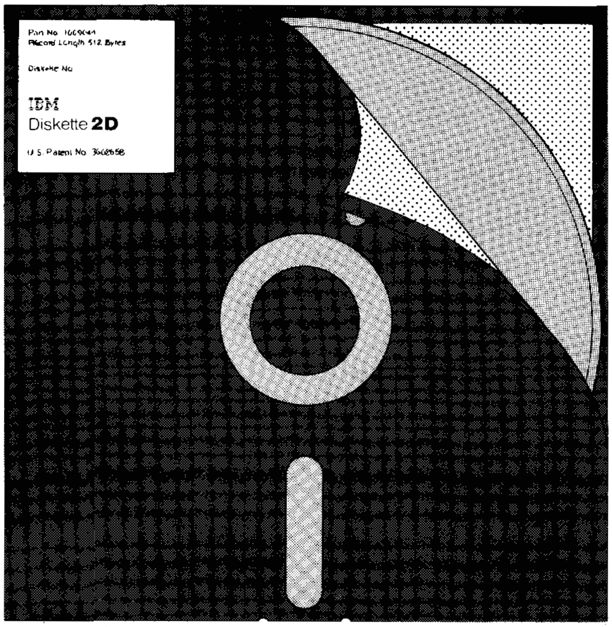

Figure 1 The diskette consists of a flexible plastic jacket with a liner of a nonwoven fabric and a magnetic disk sealed in the jacket. The head contacts the disk through the slot in the jacket. (The figure shows the jacket partially cut open.)

Disks cut from this material were affixed to an eighth-inch foam pad. This pad provided a stiffness to the extremely flexible disk and a resilience for the head-to-disk compliance.

The goal of low cost pointed toward a combination of low rpm, simple read-channel electronics, and a self-clocking encoding algorithm together with a contact recording system where a light contact pressure exists between the disk and the head. The choice of the encoding algorithm was influenced by prior work in IBM on rigid disk files and the decision was made to use frequency modulation recording. Each bit cell begins with a clock bit, and data bits are written between the clock bits. This provides a self-clocking system and allows design of a simple, low-cost data separator.

The design of the data handling system considered the worst-case analysis of bit shift tolerances from variations in components such as the write turntable speed, the read turntable speed, the write oscillator, read electronics, head-medium effects, head azimuth, center-of-rotation misalignment between turntable and disk, and other contingencies. It must be remembered that the diskette read in a 23FD would be written on a factory-controlled writer, and thus the accumulation of tolerances involved is not as great as in a read/write interchange circulation with a larger group of machines. As a result of this analysis, the data handling system was designed with a fixed-delay data separator with the data window set at a specific time beyond the leading clock pulse of each bit cell.

Initial development was based on an inner track bit density of approximately 1100 bits per inch (bpi) and a total of 32 tracks. Because of the success with the contact recording and data handling system at this density and requests from users for higher data rates, the density was increased to 1594 bpi. This increased the formatted disk capacity to 81.6 kbytes.

The early design used a lead screw accessing mechanism, rotated by two linear solenoids and an eccentric crank, to move the head across the disk to any selected position. This system provided the ability to move across tracks of information in either direction, as in rigid disk files. In addition, a third solenoid was used to load the head into contact with the disk. This permitted the disk to rotate continuously while the head was loaded or unloaded from the disk, thus minimizing wear on the head and disk and eliminating rotational start-up time before data could be read. Reliability problems with the linear solenoids and eccentric crank led to the substitution of a rotary solenoid which provided the required reliability and reduced the track-to-track movement time to 333 ms.

Diskette jacket

A search for a mailing package led to the concept of a plastic jacket to permanently enclose the disk. This eliminated the need for the foam pad and the turntable. The inside of the jacket was lined with a nonwoven fabric which protected the disk from abrasion and served as a wiping surface to clean the disk [3]. It can now be seen that the adoption of this enclosure was key in the widespread acceptance of the diskette. It provided the ease of handling and protection that are essential in the variety of applications that have developed. Figure 1 illustrates the configuration of a diskette.

Along with development of the jacket came a change in the disk itself. The substrate thickness was increased to 3×10^{-3} inch and the magnetic coating was applied to both sides even though recording was on one side only. This provided a stiffer disk and minimized curl. In addition, the head carrier was modified to include a spring-loaded resilient pressure pad opposite the head so that the disk, between the head and the pad, could be pressed lightly against the head to maintain head-to-disk compliance.

Track registration, or location, on the disk is a function of drive and disk materials, component and adjustment tolerances, and operating environment. In a low-cost device these may exhibit substantial variation. It was recognized that recording on the disk had to allow for the differences that could be expected between the write unit in the factory and the read unit in the field. Analysis of these variations led to a decision to use a write-wide and read-narrow system where disks would be recorded with a write head producing a track 31×10^{-3} inch wide and read with a read head containing a read core 12×10^{-3} inch wide. Track-to-track center spacing was 31×10^{-3} inch. This allowed the recorded track and read head to be misaligned as much as 9×10^{-3} inch without obtaining magnetic interference from the adjacent track.

The early development work had used available read heads from tape applications. These heads, however, were not optimum for this purpose. In a contact recording system, the contour of the head surface must ensure that a minimum and uniform recording gap-to-disk separation be maintained at both the inner- and outer-diameter track surface velocities. Disk velocity had been established at 90 rpm and this resulted in track velocities of 21 and 30 inches per second (ips), respectively. Experimental data involving wear, environment, signal reliability, and other factors led to a selection of a 2.5-inch spherical radius for the surface of the head that contacted the disk. As other parameters were fixed, a design was completed for a read head with a stainless steel contact surface, and a laminated metal assembly for the 12×10^{-3}-inch-wide core containing a 0.15×10^{-3}-inch read gap.

This same experimental work showed that the disk surface quality had to be improved in order to avoid permanent errors. The coating process for magnetic tape produces a long web from which the disk is punched. Minor mechanical or magnetic variations can occur across the width of the coated web that normally will not affect a narrow length of tape. However, when a disk is cut from this material and rotated in a suitable recording system, these linear variations can be observed. Investigation led to the decision to burnish the surface of the recording band. The burnishing process involves moving a cutting edge across the disk to reduce surface roughness and remove high spots. This process improved the modulation of the read signal envelope and decreased the disk defects found during disk surface analysis.

Addition of write capability

The development of the 23FD was based on a requirement for a read-only device. However, there were no technological limitations that prevented it from being a read/write device and, in fact, the factory writers that were built were modifications of the 23FD. During the development, proposals were being considered for a next-generation device which would provide read/write capability together with improvements in nearly all functional parameters.

In 1969, increased bit recording density on diskettes was experimentally demonstrated utilizing modified frequency modulation (MFM) encoding, in which clock bits are written at the beginning of a bit cell only when data bits are not recorded in the previous and present bit cells. (Kalstrom has summarized encoding alternatives and their advantages and disadvantages [4].) This approach reduced the bit-cell size by one-half and doubled the bit capacity, but reliability problems were encountered because of excessive bit shift; MFM encoding was not pursued further at that time. However, the potential for increased capacity was proven and provided encouragement for a new development.

A new diskette and drive were soon proposed that would provide 256 kbytes of formatted data per disk and a data rate of 250 kbps. In order to achieve this, the recording density had to be increased from 1594 to 3268 bpi on the innermost track, the track separation narrowed from 32 to 48 tracks per inch (tpi) and the number of tracks increased to 77, and the disk velocity quadrupled from 90 to 360 rpm. The FM encoding was retained. Other proposed enhancements included reduction of the track-to-track movement time by one-half to 167 ms and provision for more than one record sector length.

This next version of the flexible disk medium and drive was first included in IBM product announcements in 1971 and shipped in 1973. The drive was identified as the 33FD [1] and the medium was designated "Type 1." The design was based on the 23FD but applications and users did not require medium exchange with the 23FD. Therefore, the 33FD could change design where it appeared that these changes would better satisfy the requirements.

The basic diskette was not changed but its recording area was enlarged. The 23FD disk had eight sector holes uniformly spaced around the outer edge. These holes were sensed with a detector and indicated the start of a record sector. All users of the disk were limited to this same record length architecture. The 33FD designers needed more room at the outer edge of the disk for additional recorded tracks and, in addition, wanted to allow for other record length architectures in the future. Therefore, the outer sector holes were eliminated and a single index hole was located inside the innermost track. When this hole was detected, it indicated the beginning of a track. Individual sectors were located at uniform angular positions from this beginning of track and were identified by recorded markers on the track. The use of the eight sector holes had been known as *hard sectoring*. The use of the magnetic markers is known as *soft sectoring*.

When magnetic media are used for read/write exchange, old data are erased during the recording of new data and no remnant of old data can be permitted which would interfere with reading. A variety of alternatives were available for arrangement of the head pole-tip elements including write-wide/read-narrow, erase-wide and read-narrow/write-narrow, and tunnel erase. The tunnel-erase system was selected for the 33FD [5]. In this method, the new information is written over the old information and then tunnel-erased with a dc current in the tunnel-erase coil to create clean guard bands adjacent to the new information. Models of this arrangement of erase cores and laminated read/write cores were built with stainless steel housings for use in the compliance and read/write analysis.

The increase in disk velocity from 90 to 360 rpm required that additional studies be performed on the head-to-disk compliance. Inner and outer track velocities were now 76 and 136 ips, respectively, and there was concern that head-disk separation would occur because of an increase in air bearing pressure. The testing that ensued resulted in reducing the spherical head radius to two inches. Satisfactory compliance was demonstrated under all environmental conditions.

Media exchange

The design of the 33FD had to consider the variations encountered in media exchange among a group of diskettes and drives. The accumulation of tolerances from parts and adjustments plus the effects of environment will affect the location of a recorded track on the diskette. (Greenberg *et al.* have studied effects of environment on flexible disk media [6].) When the diskette is taken to another drive to be read, the second drive may not be able to read the track if its location, or registration, is too far from the nominal position. Therefore, an important study was conducted to understand what variations would be encountered and what limits could be tolerated.

One part of the study involved a measure of the limits of the read head/read electronics track reading capability. Tracks were recorded at nominal spacing and read while the read head was being moved off nominal location in very small increments toward an adjacent track until errors occurred. Tracks were then re-recorded at slightly closer spacing and the test repeated. At the greatest track spacing, the failure first occurs because the width of the track under the head has been reduced so that an adequate signal is no longer available but no adjacent track magnetic interference has been detected. The minimum track spacing is established when the read head cannot be moved off the track position because adjacent track recording then introduces magnetic interference into the read signal. A study of the convergence of these numbers defines the magnetic tracking limit.

A Monte Carlo simulation was used to study the probability of the head-track relationship exceeding the magnetic tracking limit. All drive and medium contributions to head-to-track registration were tabulated, and a population of drives was assumed as simulation input. The frequency of exchange failure was then predicted. This kind of simulation allowed proposed design changes to be analyzed in order to obtain a satisfactory probability of successful media write/read exchange.

One of the areas identified for improvement was the head positioning accuracy provided by the access mechanism. The 23FD lead screw was driven directly by a rotary solenoid. Any angular variation in the driver was seen directly on the lead screw and, ultimately, on head position. The 33FD required better control of this variable and so a geneva mechanism was used to couple the lead screw to a low-cost 90° stepper motor. Track-to-track movement time was reduced to 50 ms. The geneva mechanism could accept substantial angular variation at its input pin and still provide the required accuracy on head position. This combination was an important element of the design.

The increase in bit density from 1594 to 3268 bpi was accompanied by a reduction in the disk magnetic coating thickness and improvements in the head as well as a change in the data separator electronics. The new coating thickness was 100 microinches. At this dimension, problems were encountered when substrate irregularities occasionally exceeded the coating thickness, when holes or streaks occurred in the coating, or when other imperfections were identified. Surface finishing of the disk was again used but new techniques and specifications had to be developed for process control to ensure that the coating was not damaged due to removal of too much material.

The higher disk surface velocities, together with the expectation of increased usage, created a concern about wear on both the medium and the head. To forestall this problem, extensive testing was carried on to evaluate medium manufacturing and finish, head contact pressure, the effects of environment, and the benefits of a lubricant. Ultimately, a lubricant was selected which was compatible with the disk, did not contaminate the head, and provided greater life.

The increase in data rate required a new worst-case analysis of bit shift and the 23FD fixed delay data separator. Bit shift is the variation between the expected, or nominal, timing of a data bit and the actual timing. The bit shift increases with increased bit density. The 23FD data rate was 33 333 bits per second (bps), or 30 microseconds (μs) per bit. The 33FD data rate was 7.5 times greater, which reduced the bit cell to 4 μs per bit. With increased bit shift, the fixed delay of the data separator was no longer adequate and a phase-locked oscillator [7, 8] was introduced for the 33FD, in which frequency and phase were synchronized to the frequency of data being read.

This oscillator established the timing window used to decode the signals being read. The oscillator was phase-locked to the read signal and special data were encoded on the track for this purpose. This scheme provided a dynamic adjustment of the data window timing and permitted the clocking of worst-case bit shift.

The implementation of the 33FD recording electronics was one of the first applications of large-scale integrated analog circuits within IBM. A custom masterslice technology intended for digital applications was used to implement many of the analog functions required for the 33FD read electronics. This technology significantly reduced the cost and overall size of the electronics and later became the IBM standard for analog circuit implementation.

During the final stages of development of the 33FD, an unexpected problem developed with the diskette enclosure. Under environmental extremes, a contaminant transferred from the nonwoven liner to the disk surface. The contaminant was not easily dislodged and caused hard errors. In addition, the jacket warped and permanently distorted at elevated storage temperatures. A comprehensive search finally identified replacement materials. A liner material was obtained which was demonstrated to be free of contaminants, stable at all environments, capable of bonding to the jacket, and free of lint or loose fibers that might wedge between the head and the disk. The original jacket material was found to have a nonuniform glass transition temperature range that could not meet the storage temperature requirement because of variations in coloring pigment. After many evaluations, it was determined that material colored with carbon black had a repeatable glass transition temperature range and performed satisfactorily at these storage temperatures. Thereafter black was the only color used.

The user capacity of the diskette was established at 242 944 bytes on 73 tracks with 26 sectors on each track. In addition, one track was reserved as an index track, and three tracks were provided as spares to be used if errors were later found requiring the relocation of a sector or a track. Later, other combinations of capacity, sector size, and tracks were released as more IBM systems established usage. As these new versions were released, care was taken to provide a logical method of compatibility so that earlier versions could be read on later systems.

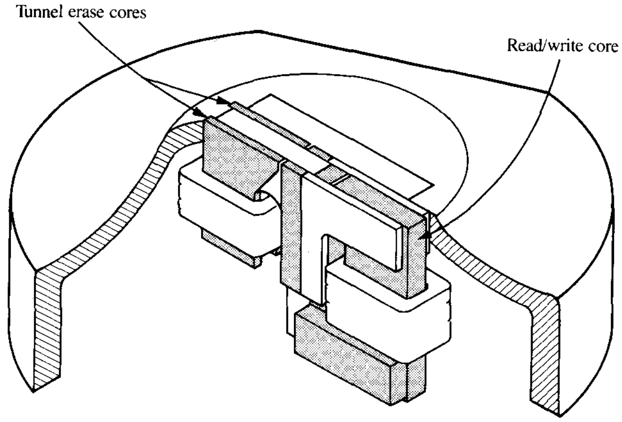

Figure 2 The 33FD head button supports the tunnel erase and read/write cores. The button is mounted in a head carriage.

Ferrite head

Soon after 33FD shipments started, questions of production capacity and cost spurred a search for an alternative design for the metal read/write head. One requirement was basic; the new head had to be interchangeable both physically and functionally with the metallic head. Ferrite was chosen for the magnetic cores in the read/write and tunnel-erase sections, and barium titanate ceramic was selected for the magnetically inert structure which supports the cores [9]. These components were assembled in the structure shown in Fig. 2, which is then mounted in the head carriage. Heads of this style were tested to even greater life requirements than the metallic head and were found satisfactory in all functional and life tests.

While the 33FD was under development, comparisons were made in several IBM programs between the diskette and tape. These comparisons involved the commercially available cassettes as well as a unique cassette design then being considered by IBM. Most of the comparison parameters were obvious ones such as cost, data rate, reliability, etc., but the ultimate decision in favor of the diskette was due to the performance improvement provided by the random access capability. This, together with a growing need among IBM systems for a common magnetic exchange medium, cemented the choice of the diskette.

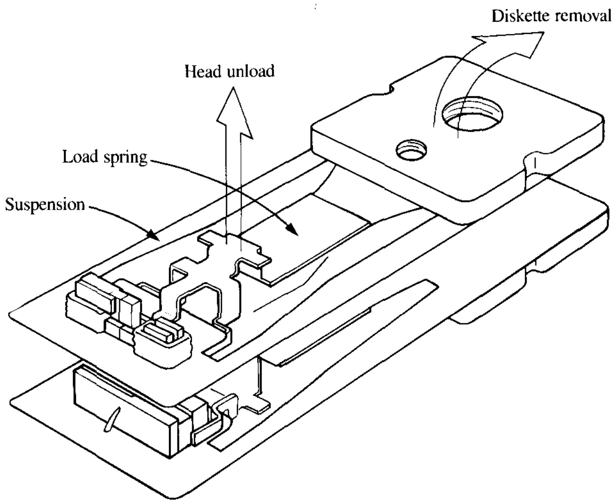

Figure 3 The pair of 43FD head suspensions and load springs are symmetrical. The prestressed load springs press the heads against the media until a head unload force is applied. A diskette removal operation causes the heads to separate further.

The 33FD was the first drive that provided diskette read/write exchange in IBM, and its use soon became widespread. In many of these applications, the diskette replaced the punched card. It provided faster data rates, reusable medium, much smaller storage volume, and lower system cost. Equipment designers were now able to make substantial reductions in system size, power, and cooling requirements by eliminating the punched card readers and punches. The utility of the diskette was similarly recognized by other users and manufacturers, and work was undertaken by the American National Standards Institute, the European Computer Manufacturers Association, and the International Organization for Standardization to develop diskette standards.

Dual head drive

The success of the 33FD and diskette soon created a need within IBM for greater medium capacity. System designers wanted the ability to read and record more information before the operator was required to change a diskette. Because some of these new applications would be extensions of the original ones, an important criterion was placed on the design options: they had to provide an ability to read and record existing Type 1 diskettes in addition to any new diskette. It was not required, however, that a new diskette be processed on the earlier 33FD. This upward compatibility served to establish the new medium dimensions identical to those of Type 1. The options considered were changes in linear bit density and the addition of recording on the second side. Changes in track density were ruled out in order to maintain the same simple open-loop track-accessing method. Considerable merit existed for each option, but the decision was made to use the second side and maintain the existing data separator and FM encoding logic. The new disk would be tested and factory-formatted on both sides. This would double the capacity to approximately 0.5 megabytes (Mbytes). The new drive would be known as 43FD [1, 10, 11] and would be announced and shipped by IBM in 1976. The new medium was designated "Type 2."

A new two-head carriage design was needed to permit reading and recording on each side of the diskette while it remained in the drive. The single 33FD head was rigidly mounted on a carriage and located close to the plane of the rotating disk. A felt pad on the opposite side lightly pressed the disk to the head. For the new design, attempts were made to use two opposing, rigidly mounted heads and to press one against each side of the disk for two-sided operation; however, satisfactory head-to-disk compliance could not be obtained. Any minute irregularities in the disk caused excessive bounce, and wear of the medium was unsatisfactory. The next approach used one head rigidly mounted and the second head on a flexible suspension. The second head closed on the disk, and the suspension system provided a pressure on the disk and a resilience to absorb disk irregularities. Again, a stable compliance was not achieved. The principal problem was the unpredictability of disk dynamic movements over a wide range of environments with a variety of diskettes.

The designers then turned to rigid disk technology in search of a solution to the problem. The IBM 3348 Data Module head suspension (Winchester) had encountered similar design problems. It provided a reference for development of the two-head carriage and suspension system.

Both the head carriage and the head contact surface were redeveloped for the 43FD. New analysis techniques and measurement systems were employed. The use of transparent glass heads and white light interferometry were vital for determining the nature and stability of head-to-disk compliance. The design quickly centered on a pair of opposing flat surfaces with a central relief running their length. The effect was much like two sleds with their runners in contact. Careful design of the relief dimensions and edges of the flat surface was found necessary in order to maintain a proper air bearing beneath the head for signal stability and good wear performance, and to prevent damage to the disk during head loading. The read/write cores were radially offset from each other by four tracks to avoid magnetic coupling through the disk. This reduced the radius of the innermost track on one side and compressed the bit density to 3408 bpi.

The head mounting system had to maintain accurate head position and azimuth alignment, control the pitch and roll accelerations contributed by the disk, provide high stiffness in the radial, tangential, and yaw directions, and allow the heads to be separated by at least 6×10^{-2} inch to remove the diskette. It also had to control head bounce during load and unload cycles so that the disk was not damaged and so that uniform signals could be quickly established. A symmetrical stainless steel suspension and load spring were designed to control the head displacement and loading within the region close to the disk surface. Stiffer secondary flexure arms provided the greater head separation needed for diskette removal. Figure 3 illustrates the suspension and load spring members.

The 43FD used the same ferrite read/write and tunnel-erase cores, FM encoding algorithm, and data separator found in the 33FD. To assist in studying their performance, a new laboratory measurement system was used that divided the data window into sixteen parts having read intervals of 125 nanoseconds (ns) each. Each part had an associated counter. The system detected the time position of a bit and incremented the appropriate counter. In this manner, a bit-timing histogram could be created for a pattern of bits, providing an important measure of bit shift. This analysis system measured the total bit shift resulting from all components of drive, medium, head, and electronics rather than from a single component and allowed further refinement of the read channel electronics. It would prove to be extremely valuable when used in subsequent diskette and drive developments.

Early in the development of the 43FD, a new head positioning system was designed that reduced track-to-track movement time from 50 to 5 ms, improved accessing accuracy and reliability, and lowered the acoustical noise. This system replaced the lead screw-stepper motor with guide rods and the band-capstan-stepper motor combination shown in Fig. 4. It provided a low-cost method for translating rotary motion from a fast, economical stepper motor into linear motion.

Double density encoding

The earlier work done with MFM encoding and the availability of large-scale analog integrated circuits now combined to encourage development of a new drive and an even larger-capacity diskette. A goal was set for 1.2 Mbytes of formatted capacity. The new drive would again need the upward compatibility feature which would permit it to read and write on Type 1 and Type 2 diskettes. This requirement affirmed the same diskette dimensions for the new medium. Again the track locations were left unchanged to preserve the simple open-loop accessing method. The capacity increase would be obtained by doubling the linear bit density to 6816 bpi.

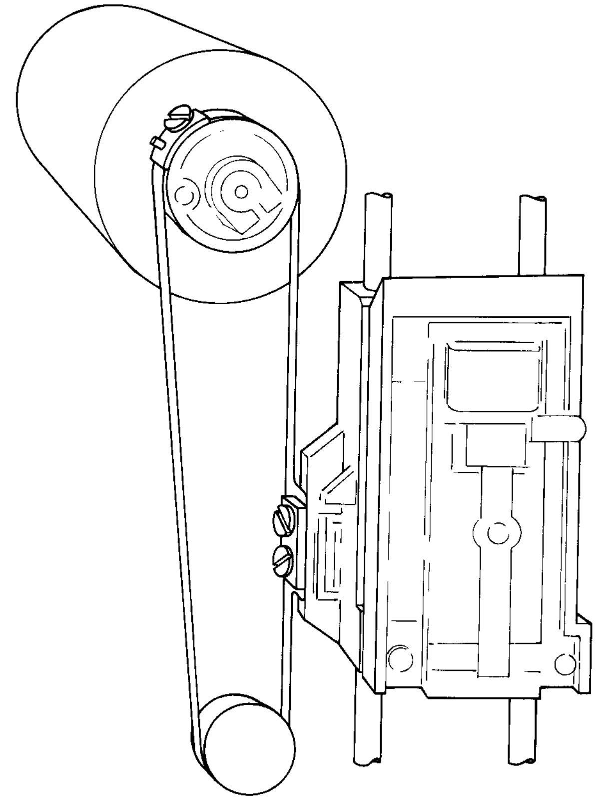

Figure 4 The 43FD head positioning system uses a flexible stainless steel band to translate the stepper motor displacement into linear motion of the read/write head carriage. Guide rods provide precision support for the carriage.

This next diskette drive was identified as the 53FD [1] and was announced and shipped by IBM during 1977. The new diskette was called "Type 2D."

The 43FD head carriage and suspension were reexamined and used in the 53FD with only minor modification. The head retained the tunnel-erase core arrangement, but the erase core gaps were now placed so they were no longer parallel to the read/write core gap. This minimized the coupling of adjacent-track magnetic transitions into the read/write core. Further improvements were made to reduce the magnetic reluctance of the erase core arrangement and improve the erase efficiency. The read/write core gap was reduced to 0.075×10^{-3} inch in order to record the smaller flux transitions at this higher bit density.

The MFM provided a doubling of capacity and also doubled the data rate to 0.5 Mbps. This reduced the bit cell duration to 2 μs. Bit shift problems were immediately experienced but the analysis and correction were readily accomplished with the shift sampling system described earlier. Modified to divide the data window into sixteen parts of 62 ns each and using worst-case or random patterns, this system was able to quickly examine modifications to the recording channel and head. The benefits of bit shift compensation methods were compared and eventually a peaking filter technique was adopted to compensate for loss in recording channel bandwidth due to head and medium effects. The filter consisted of an underdamped second-order low-pass stage with a resonant frequency just above the all-ones frequency of the recording channel. This provided a high-frequency boost in the channel response, resulting in broader bandwidth and reduced bit shift.

With the use of MFM encoding, the disk surface condition became more important. Narrow scratch-like irregularities, say 50 microinches wide, running in the direction of the web coating process, were seen on some disks. If this scratch were coincident with a radius of the disk, it would be coincident with the head read gap at that position on the disk, and it could create a noise transition which might be detected as a false, or extra, bit. On either FM or MFM disks, extra bit defects are not permitted, but the sensitivity of the MFM read electronics required new specifications for Type 2D diskette magnetic quality.

Diskette Type 2D and the 53FD provided new diskette capacity and performance to IBM systems and permitted additional users to further optimize or extend their systems. Systems such as the IBM 5110 and 5120 have enhanced their applications with dual diskette drives which provide an economical processing DASD file capability as well as the exchange medium function of the diskette.

Automatic diskette handling

Systems with nonremovable DASD storage require some form of removable medium for archival and backup purposes. In the interest of economy, they often use the same drive and medium that are used for data exchange. IBM systems such as System/32 and System/34 have used the diskette for these purposes. As the fixed DASD capacity of systems is increased, more data is saved via diskette and the time and handling required to create these diskettes becomes a concern.

These concerns generated a requirement to develop a device using diskettes but with faster data rate, the ability to process multiple diskettes with limited operator handling, and the provision for single diskette insertion and removal by the operator as with the 53FD. Under system control, this device would provide either the save-restore function or single or multiple diskette data entry. An earlier IBM device had provided the capability to process a group of diskettes sequentially [12]. What was now sought was the ability to process a group of diskettes in an intermixed order. The device would also have to be capable of reading and writing on Type 1, Type 2, and Type 2D diskettes in order to maintain the medium exchange disciplines that existed with other systems.

The objective required only a faster data rate, not a larger-capacity diskette. Therefore, no attempt was made to develop a new diskette. Instead, it was proposed that the disk velocity be doubled to 720 rpm. The unit which was developed was announced by IBM in 1978 and shipped in 1979. It was designated 72MD [1, 13, 14] and is known as the IBM Diskette Magazine Drive.

The initial work focused on the head-disk compliance and it was found that the 53FD head suspension would not perform satisfactorily at 720 rpm. Inner and outer track surface velocities were now 147 and 272 ips, respectively. The disk frequency content was doubled from 6 to 12 Hz and the energy transmission from the disk into the suspension increased. The result was that a suspension system designed for 360 rpm would, at 720 rpm, become unstable, causing erratic signal output and greatly accelerated medium wear.

Both the head suspension and bearing surface had to be redesigned to stiffen the suspension. The bearing surface has two rails which function as slider bearings on the surface of the disk. A film of air is built up between the slider and the disk and, if not controlled, causes excessive separation between the recording gap and the disk and may cause oscillation of the head. The profile of the rails was shaped to reduce the pressure buildup at the leading edge and maintain the desired separation at the gap. It had been established on the 43FD that when the head is loaded, contact first occurs with an edge of the bearing surface, and careful edge preparation was required to avoid disk damage. With the stiffened suspension, these edges had to be newly defined to prevent disk damage during the load conditions anticipated.

Also added to the 72MD was an automatic gain control circuit to permit the read electronics to tolerate the greater envelope modulation resulting from the higher disk velocity. Further optimization of the electronic peaking filter, for reduced bit shift, was provided by the addition of a threshold detector. Analysis of these changes was again performed by the shift sampling system. At this data rate, the data window was now divided into sixteen parts of 31 ns each. Again, worst-case data patterns were measured and statistical analyses were performed to ensure that bit-shift compensation was satisfactory. This

Table 1 Characteristics of IBM diskettes and drives.

Diskette type	*Formatted capacity [Note II]* (bytes)	*Encoding*	*Inner track* (bpi)	*Recorded sides*	*Total tracks*
[Note I]	81 664	FM	1594	1	32
1	242 944	FM	3268	1	77
2	568 320	FM	3408	2	154
2D	1 212 416	MFM	6816	2	154

Drive	*Velocity (rpm)*	*Heads*	*Data rate* (bps)	*Bit cell duration* (μs)	*Track-to-track movement* (ms)
23FD	90	1	33 333	30	333
33FD	360	1	250 000	4	50
43FD	360	2	250 000	4	5
53FD	360	2	500 000	2	5
72MD	720	2	1 000 000	1	5

Note I The 23FD diskette did not have a designation.
Note II Capacities shown are typical. Other capacities may be obtained.

analysis demonstrated the capability of the redesigned head and suspension system and read electronics to operate satisfactorily at 720 rpm.

In order to minimize operator handling of individual diskettes, a magazine was designed that held ten diskettes, each located in a slot and separated from its neighbor. The magazine measures 2.25 inches wide by 8.5 inches high by 8.25 inches deep and is easily carried or stored. Two magazines can be inserted into the 72MD. In addition, three positions are available on the drive where individual diskettes may be inserted alongside the magazines. In operation, a load device can grasp any one diskette and position it for read/write operation, then return it and move any other selected diskette into position.

This combination of faster rpm, dual magazines, and automatic handling of the diskette now provides system designers with a device that can operate in either background or foreground mode with minimal operator attention to the save-restore function and with less system time for this task. The data rate has been doubled to 1 megabit per second (Mbps). Adjacent diskettes can be exchanged within three seconds. With two magazines in place, 24 Mbytes are available to the system. The 72MD has provided a significant enhancement to the function and performance of the diskette.

Other developments

Many of the developments reported here have also been reported by other companies; they may exhibit differences in their specific drive implementations but typically they offer the same diskette read/write characteristics, and thus diskettes can be exchanged among a wide community of users. This has encouraged the adoption of the diskette as a nearly universal medium. In addition, some companies have adopted other methods of encoding such as group coded recording (GCR) and modified-modified-frequency modulation (M^2FM), thus offering encoding alternatives for system designers.

A related development is a smaller diskette measuring 5.25 inches square and a proportionately smaller drive, both of which are offered by some of these same companies. While most versions of this smaller diskette have reduced capacity, it is notable that some use higher track densities of 96 and 100 tpi and, when combined with recording on both sides, could achieve capacities up to $\approx$1 Mbyte unformatted. No attempt will be made to identify the large number of companies offering the small diskette.

A large body of literature has been published in the United States and other countries discussing the diskette and drive and related technology. A very modest summary of these references is cited here [15–25].

During 1979, IBM announced additional models of diskette drives and shipments started in 1980. They are not reported here since they are primarily alternative versions of the 33FD and 53FD.

Summary

The IBM diskette and drive developments have been traced since they began in 1967; Table 1 summarizes their characteristics. Diskette capacity has increased nearly fifteen times through improvements in recording electronics, read/write heads, and media. A thirtyfold increase in data rate has resulted from these same improvements plus increases in disk velocity. The application utility of the diskette has been strongly enhanced with greater reliability in the medium and drive, reduced track-to-track movement time, and the development of automatic diskette handling capability. Because of its excellent cost performance and as a result of these and other improvements, the diskette has been accepted and established as a principal magnetic medium throughout the computer industry. Technology potential exists for future advances. The realization of these potentials will be influenced by the rate of convergence of cost, function, and performance of competing technologies.

Acknowledgments
The author recognizes the contribution of V. R. Witt for the direction, support, and encouragement he provided during the development of the 23FD. The author also expresses his thanks and appreciation to D. O. Castrodale, R. D. Chenoweth, K. A. Shidler, and W. S. Wentink for their extensive consultation and contributions during the preparation of this paper.

References and notes

1. 23FD, 33FD, 43FD, 53FD, and 72MD are designations used internally by IBM.
2. W. L. Dalziel, J. B. Nilson, and D. L. Wartner, "Data Storage Apparatus Implying a Single Magnetic Disk," U.S. Patent 3,678,481, July 18, 1972.
3. R. Flores and H. E. Thompson, "Magnetic Record Disk Cover," U.S. Patent 3,668,658, June 6, 1972.
4. D. Kalstrom, "Simple Encoding Schemes Double Capacity of a Flexible Disk," *Computer Design* **15,** 98 (1976).
5. M. F. Garnier and R. D. Nelson, "Trim Erase Head Assembly For Providing a Uniform Erase Field," U.S. Patent 3,769,469, October 30, 1973.
6. H. J. Greenberg, R. D. Stephens, and F. E. Talke, "Dimensional Stability of 'Floppy' Disks," *IEEE Trans. Magnetics* **MAG-13,** 1397 (1977).
7. P. C. Lang, A. N. La Pine, and J. E. Vaughn, "Random Binary Data Signal Frequency and Phase Compensation Circuit," U.S. Patent 3,701,039, October 24, 1972.
8. A. M. Gindi, "Self-Clocking System Having a Variable Frequency Oscillator Locked to Leading Edge of Data and Clock," U.S. Patent 3,599,110, August 10, 1971.
9. W. W. Childers, K. H. Elser, and P. J. Peterson, "Read/W e and Longitudinal Edge Erase Head Assembly Having Multiple Similarly Shaped Layers," U.S. Patent 3,846,840, November 5, 1974.
10. D. O. Castrodale, W. J. Pendy, Jr., and W. S. Wentink, "Data Storage Apparatus Using a Flexible Magnetic Disk," U.S. Patent 4,089,029, May 9, 1978.
11. D. O. Castrodale, S. Steele, W. S. Wentink, and S. Wheeler, "Recording Apparatus For Magnetic Disks Using Both Sides of the Disk," reissue patent, U.S. Patent Re. 29,380, August 30, 1977.
12. The IBM 3747 Data Converter and the IBM 3540 Diskette Input/Output Unit included a device that would sequentially feed individual diskettes from one stack into a diskette drive and eject them into another stack.
13. D. O. Castrodale and T. R. Fournier, "Drive Machine Moving Magnetic Disks Only Partially Therein for Data Transfer," U.S. Patent 4,170,030, October 2, 1979.
14. W. E. Beuch, C. A. Larson, M. N. Zell, and T. R. Fournier, "Loading Mechanism for Magnetic Diskettes," U.S. Patent 4,170,031, October 2, 1979.
15. D. Noble, "Some Design Considerations for an Interchangeable Flexible Disk File," Intermag Conference, Toronto, Ontario, May 1974.
16. *IBM Diskettes 1, 2, and 2D, Original Equipment Manufacturers Information*, Order No. GA21-9388, IBM Corporation, Old Orchard Rd., Armonk, NY 10504.
17. *The IBM Diskette General Information Manual*, Order No. GA21-9182, IBM Corporation, Old Orchard Rd., Armonk, NY 10504.
18. G. Sollman, "The Future Growth of Floppy Disk Technology," *1978 WESCON Technical Papers* **22,** Paper 16/5, 1 (1978).
19. C. Lin and R. F. Sullivan, "An Application of White Light Interferometry in Thin Film Measurements," *IBM J. Res. Develop.* **16,** 269 (1972).
20. P. Sidnu, "Group-Coded Recording Reliably Doubles Diskette Capacity," *Computer Design* **15,** 84 (1976).
21. J. Hoeppner and L. Wall, "Encoding/Decoding Techniques Double Floppy Disk Capacity," *Computer Design* **19,** 127 (1980).
22. D. Bowers, "Floppy Disk Drives and Systems," *Mini-Micro Syst.* **10,** 36 (1977).
23. D. B. Bogy, N. Bugdayci, and F. E. Talke, "Experimental Determination of Creep Functions for Thin Orthotropic Polymer Films," *IBM J. Res. Develop.* **23,** 450 (1979).
24. P. Kinnucan, "Suddenly, The Double-Sided Floppy Is In," *Mini-Micro Syst.* **10,** 22 (1977).
25. D. Resnick, "Drive Mechanism Design Reduces Errors in Mini-Floppies," *Computer Design* **19,** 160 (1980).

Received April 3, 1980; revised October 3, 1980

The author is located at the IBM Information Systems Division laboratory, Highway 52 and Northwest 37 Street, Rochester, Minnesota 55901.