

J. M. Harker
D. W. Brede
R. E. Pattison
G. R. Santana
L. G. Taft

A Quarter Century of Disk File Innovation

This paper traces the development of disk file technology from the first disk drive to the present. A number of innovative advances are reviewed in the evolution of mechanical design, materials, and processes. These advances constitute the technological base that has permitted almost four orders of magnitude of improvement in areal density; they are discussed from four interrelated aspects: the magnetic head and its air bearing support; the head positioning actuator; the disk substrate and its magnetic coating; and the read/write signal detection and clocking electronics.

Introduction

The first issue of the *IBM Journal of Research and Development* in 1957 featured two papers, one describing the IBM 305 RAMAC system [1] and the other describing the IBM 350 disk file [2]. As Stevens mentions in his overview paper "The Evolution of Magnetic Storage" [3], the 350 was the first production movable-head disk drive. The purpose of this paper is to describe the significant innovations that led to the introduction of that product and to the evolution of the current generation of disk file products.

Inductive magnetic recording was selected as the base technology for disk files because of its advantages of nonvolatility, immediate readback without intermediate processing, unlimited reversibility, and the relative low cost and simplicity of the transducer and recording medium.

Increasing linear bit density has depended upon the scaling down of the three geometric parameters of head-to-disk spacing, read/write gap length, and disk magnetic coating thickness. Progress in reducing these key geometric parameters and the resulting linear bit density of selected products is shown in Table 1. Achieving these reduced parameters has been an iterative process and has required significant innovation. For example, scaling

down the coating thickness from 1200 microinches to 25 microinches has required extensive process development in the area of substrate preparation and magnetic film coating, formulation, application, buffing, and polishing. Reducing the spacing between the head and the disk from over 800 microinches to less than 13 microinches has required an in-depth understanding of air bearing technology through analysis, simulation and testing, development of low-mass sliders with very stiff air bearings, and development of very smooth and flat disk surfaces with few mechanical asperities.

Early recording head cores were made of laminated mu-metal with the gap formed by a copper shim. Later, silicon dioxide was deposited on ferrite to form smaller, more precise gaps, and most recently the use of thin permalloy films and photolithographic techniques has permitted extremely small and accurate dimensions for head gap, pole tip width, and pole tip thickness.

Track density improvements have also been an important part of advances in disk file areal density. These improvements have come through fabrication of smaller gap widths and more accurate head positioning technology. The resulting advances in areal density are shown in Table 1.

Copyright 1981 by International Business Machines Corporation. Copying is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract may be used without further permission in computer-based and other information-service systems. Permission to *republish* other excerpts should be obtained from the Editor.

Table 1 Development of technologies in key areas of magnetic head and its air bearing support, disk substrate and its coating, head-positioning actuator, and read/write electronics.

Year of first ship	1957	1961	1962	1963	1966	1971	1973	1976	1979	1979	1981
Product	350	1405	1301	1311	2314	3330	3340	3350	3310	3370	3380
Recording density											
Areal density (Mb/in. ²)	0.002	0.009	0.026	0.051	0.22	0.78	1.69	3.07	3.8	7.8	>12
Linear bit density (bpi)	100	220	520	1025	2200	4040	5636	6425	8530	12134	15200
Track density (tpi)	20	40	50	**	100	192	300	478	450	635	>800
Key geometric parameters (microin.)											
Head-to-disk spacing	800	650	250	125	85	50	18	**	13	**	<13
Head gap length	1000	700	500	250	105	100	60	50	40	25	**
Medium thickness	1200	900	543	250	85	50	41	**	25	41	<25
Air bearing & magnetic element											
Bearing type	hydrostatic		hydrodynamic	**	**	**	**	**	**	**	**
Surface contour	flat	**	cylindrical	**	**	taper flat	**	**	**	**	**
Slider material	Al	**	stainless steel	**	ceramic ferrite	**	ferrite	**	**	ceramic film	**
Core material	laminated mu-metal	**	**	**	**	**	**	**	**	**	**
Slider/core bond	epoxy	**	**	**	**	glass	integral	**	**	deposited	**
Disk											
Diameter (in.)	24	**	**	14	**	**	**	**	8.3	14	**
Substrate thickness (in.)	0.100	**	**	0.050	**	0.075	**	**	**	**	>0.075
Rpm	1200	**	1800	1500	2400	3600	2964	3600	3125	2964	3620
Fixed/removable	fixed	**	**	removable pack	**	module	fixed	**	**	**	**
Data surfaces/spindle	100	**	**	10	20	19	6	15	11	12	15
Actuator											
Access geometry	x-y	**	linear radial	**	**	**	**	**	rotary	linear	**
Heads	2 heads/actuator	**	1 head/surface	**	**	**	2 heads/surface	**	1 h/s	2 h/s	**
Positioning	motor-clutch	**	hydraulic	**	**	voice coil motor	**	**	**	**	**
Final position	detent	**	**	**	**	servo surface	(+sector)	**	servo surface	**	**
Actuators/spindle (max. no.)	3	**	2	1	**	**	**	**	1	2	**
Avg. seek time (ms)	600	**	165	150	60	30	25	**	27	20	16
Read/write electronics											
Data rate (Kbytes/s)	8.8	17.5	68	69	312	806	885	1198	1031	1859	3000
Encoding	NRZI	**	**	**	2 f	mfm	**	**	mfm	2, 7	**
Detection	ampl	**	**	**	peak	delta	**	**	**	delta clip	**
Clocking	2 osc	**	clk trk	osc	vfo	**	**	**	**	**	**

**Same as in preceding column.

The evolution of the whole technology is given in the overview paper [3], and the progress in disk file manufacturing and in selected innovations in materials, processes, and testing is discussed in the paper by Mulvany and Thompson [4]. The present paper traces the development of each part of the technologies in four key areas:

- The development of the magnetic head and its air bearing support that provides the close spacing between the disk surface and read/write head necessary for high-density magnetic recording.
- The development of the disk substrate and its magnetic coating.
- The mechanical design aspects of the actuator that positions the read/write heads over concentric tracks of a rotating disk.
- The key innovations in logic and electronics required to read and write data reliably and accurately from a disk.

Much of this development has been based upon the work of many individuals who have created, over the past quarter of a century, a technological base that has permitted the improvement of almost four orders of magnitude in areal density shown in Table 1 and also enhancements in performance, function, and reliability. Two individuals in particular deserve mention because their influence was so pervasive through the early days of development. They certainly deserve credit for motivation and for active participation in many of the innovations to be discussed. They are R. B. Johnson, who had the vision that such a device was needed and could be built, and L. D. Stevens, who provided the engineering management that realized the first successful product [5].

Air bearing spacing and magnetic heads

Film bearings, both self-acting and externally pressurized, have been in common use for over a hundred years.

They allow one surface to move relative to another, with a film of either a compressible or incompressible fluid preventing contact and therefore preventing wear. A common example would be the journal bearing in an automobile engine. The use of air as the lubricating film was first reported by Kingsbury in 1897 [6]. The whole concept behind magnetic disk files rests on use of this simple mechanism. The innovation was to use such a bearing to maintain a small and consistent spacing between a magnetic read/write head and a disk surface whose axial runout was far greater than the spacing required.

- *Hydrostatic (pressurized) air bearings*

The event that provided the basis for the first disk file was accomplished on June 2, 1953, when W. A. Goddard [7] demonstrated the operation of an air-fed bearing with a magnetic element incorporated. This head, shown in Fig. 1, had small holes drilled around a circumference on its face so that air flowing out provided a pressure pad to resist the force of a spring-loaded suspension. The magnetic head mounted in the center successfully wrote and read bit patterns on a disk, establishing the feasibility of maintaining a consistent spacing between a disk surface and a read/write head. Of all developments related to the disk file, the ability to maintain a consistent spacing (and in subsequent generations, to diminish it) has been the prime driving force for all other improvements in the recording technology.

N. A. Vogel took the air bearing concept and evolved the configuration used in the first production file. It was found that a bleed hole in the center of the bearing was necessary for stability. The force required to load this head onto the disk was provided by three pins acting as pistons connected to the same air chamber that supplied the holes on the bearing surface. When air pressure was applied, the head was forced against the disk with a load proportional to the support provided by the bearing. Light springs were provided to return the head from the disk to the socket when air pressure was removed or in the event of a failure in the air supply [8, 9].

This head provided a spacing of about 800 microinches for a separate write-wide magnetic element and a read-narrow magnetic element. This combination provided a guard band between recorded tracks to allow for the positioning tolerance of the moving-head assembly.

- *Hydrodynamic (self-acting) air bearings*

While Vogel's air head design was improved through several generations of development, it became apparent that eliminating the disk-to-disk excess motion by placing one head element on each disk surface would be a

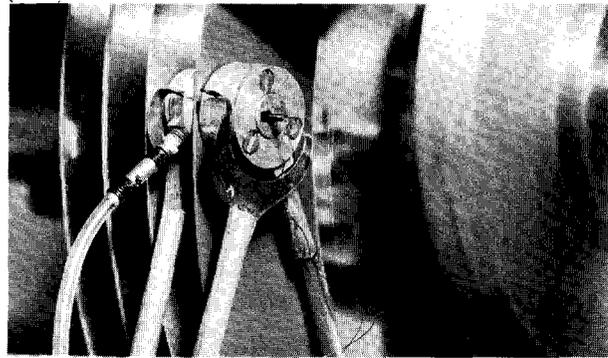


Figure 1 First air bearing magnetic head, tested on June 2, 1953.

significant functional improvement. Indeed, files having several air-fed heads per surface were built for the IBM Stretch computer system. The high consumption of compressed air and the maintenance and reliability associated with compressors, however, made this design of limited applicability.

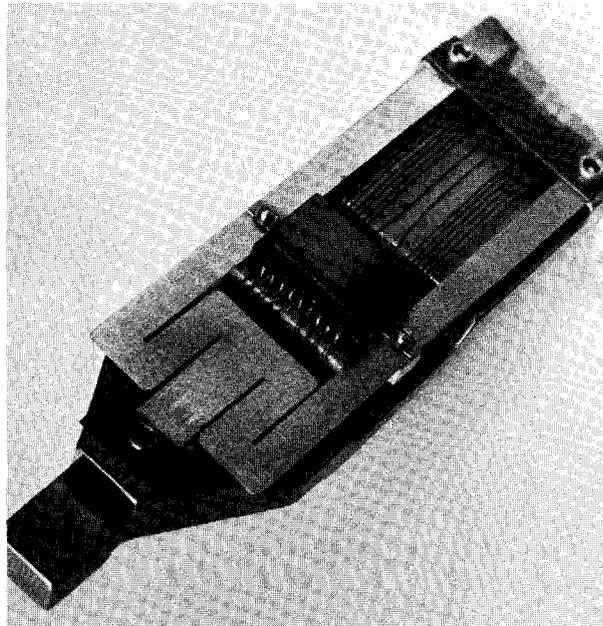
J. J. Hagopian conceived of using a self-acting air bearing in such an array of magnetic head carriers. His experimentation showed that a reasonably flat bearing properly gimballed and loaded would generate a hydrodynamic bearing and provide stable spacing of the magnetic element from the rotating disk. Based upon this concept, an advanced disk file project was initiated. As development proceeded it became apparent that the operation of this bearing was erratic. It would fail, at times, when the slider occasionally touched the disk surface. This would often lead to catastrophic failure of the disk surface and subsequent destruction of the magnetic head in its carrier.

Classical bearing theory dealt only with an incompressible fluid. The effect of compressibility on the operation of such bearings was the subject of an intensive analytical and experimental study conducted in the summer of 1957. The analytical work was done by Dr. W. A. Gross [10] and the experimental work was done principally by Dr. K. E. Haughton and R. K. Brunner [11]. The practical result was that an entry wedge must be provided for the entry of the boundary-layer air film.

This entry wedge was obtained by providing a convex cylindrical surface on the face of the slider, which analysis and experiment indicated should have a radius of curvature of about 250 inches. An innovative way to fabricate such a slider was suggested by Brunner. First, he accurately distorted concavely the stainless steel slider



(a)



(b)

Figure 2 Evolution of the designs for slider and suspension: (a) accessing heads; (b) 2305 nine-track fixed head.

face in the direction of the boundary-layer air flow. The face was then lapped flat while it was distorted. Releasing the distorting force produced a convex cylindrical surface with a precise radius of curvature that provided the necessary entry wedge. The manufacturing process developed to produce this surface economically used the inverse of the technique described above; *i.e.*, the lap was distorted to the desired radius of curvature. Enhancements of this concept were used to produce all cylindrical slider-bearing surfaces.

These bearings would operate as predicted on flat surfaces, but there were still instability problems on the

actual disk surface. Brunner and Haughton found that the local curvature of the disk surface was of the same order of magnitude as the convexity of the slider, producing regions of instability as the disk rotated. They also found that the runout of the disk produced local accelerations that caused excessive change of spacing. Specification of disk flatness in terms of the velocity and acceleration of runout and local curvature has since been basic to the disk manufacturing process [4].

Asperities on the surface of the magnetic coating, even though buffed, could cause instability of the slider. Brunner introduced the technique of running a head made of a hardened material back and forth over the disk at a spacing less than normal, thereby burnishing off any local asperities by repeated contact. This technique is still commonly used in disk manufacture.

Related to the slider-bearing technology is the design of a suspension that is simultaneously rigid in both directions in the plane of the disk but allows freedom to follow the vertical runout of the disk and provides for roll and pitch as the head follows the variations of the disk contour. It is particularly important to have a high radial suspension rigidity because of the high accelerations as the actuator moves the head and suspension from track to track.

Improvements in the self-acting slider-bearing technology, the magnetic elements, and their suspension systems have been key to increasing disk file capacity. The earlier sliders required a load of about 350 grams, as contrasted to the later generation where the required load is on the order of 20 grams. The following section discusses the development of heavy-load slider/head development and a subsequent section reviews the development of the light-load slider head.

Heavily loaded sliders, suspensions, and heads

The first disk drive to use a self-acting slider was the IBM 1301. It used a gimbal ring design that encircled the stainless steel slider and was attached to it and the supporting arm with cone-pointed pivot screws [Fig. 2(a)]. The flying height of the slider was 250 microinches and the magnetic element was composed of laminated mu-metal potted in epoxy in a small can that provided magnetic shielding. The can was in turn epoxied into the slider. A linear density of 520 bits per inch (bpi) and a data rate of 544 Kbits/s was achieved with this combination. The gimbal ring provided freedom for roll and pitch and stiffness for yaw, resisting both radial and circumferential forces. It carried the 350-gram load from a long torsion spring pressing on the arm through the gimbal to the slider. This design was difficult to maintain in that the

pivot screws required adjustment and lubrication. Adjustment was disturbed by temperature variations, and oil sometimes found its way to the disk surface.

This basic suspension design was improved in several steps, as seen in Fig. 2(a). The most significant improvement was made with the design for the IBM 2311. It separated the vertical load requirements from the gimbal system. The load was applied directly to the slider at the center of roll and pitch. The 2311 suspension design used a thin flat flexure ring to control the roll, pitch, and yaw. The flexure ring was cut from stainless sheet stock and the joints were made by spot welding, eliminating screws and lubrication and reducing material, fabrication, and assembly cost. Variants of this design were used on all subsequent sliders.

The increased track density of the 2311 [from 50 to 100 tracks per inch (tpi)] was achieved in part by use of an improved head element configuration. This involved a change to a "tunnel" erase trimming of the newly written track [12] that provided a better signal-to-noise ratio than the previous guard band techniques of write-wide read-narrow and erase-wide write/read-narrow. Also introduced on the 2311 head was a ball staking technique for mounting the magnetic element in the slider [13]. This technique reduced element movement within the slider assembly and maintained the tighter tolerances necessary for the higher track density.

A ferrite magnetic head element was designed for the IBM 2314 disk file, first shipped in 1966, because of its 2.5-megabit/s data rate resulting from increased bit density (1100 to 2200 bpi). Manufacturing process considerations dictated that these ferrite magnetic elements be mounted in a mechanically similar ceramic material. Alumina was selected because it could be formed to the desired circular geometry and had desirable mechanical properties. It provided an improved head/disk interface, and ceramic or ferrite materials have been used in all subsequent products. The ferrite element was bonded to the slider with an epoxy and the face of the assembly lapped to the required cylindrical surface.

As the head-to-disk spacing was reduced from 85 microinches for the 2314 to 50 microinches for the IBM 3330, it became critical that the magnetic element gap be flush with the face of the slider. In addition, the thermal and hydroscopic properties of epoxies used to bond the ferrite core to the slider became inadequate. Consequently, a compatible ferrite-ceramic-glass material combination was developed for the 3330. This required two compatible glasses with enough separation in working temperatures to allow the separate assembly bonding of

the parts of the core and the parts of the slider with a high-temperature glass. Then the final assembly of the previously assembled core and slider was accomplished with a second glass at a lower temperature that did not disturb the dimensions established in the first bonding [14]. A new slider ceramic was developed that had a thermal expansion coefficient compatible with ferrite and glass, and was of sufficient hardness, density, and grain structure to provide a satisfactory bearing surface.

The 3330 slider geometry was changed to rectangular and the bleed holes were replaced by a central slot [Fig. 2(a)]. A spring element in each support arm provided the mechanical load for the associated slider [15]. This design reduced the tolerances associated with the mechanical load and made it possible to set the flying height of each slider to close limits. This combination of materials and geometry produced a dimensionally stable head/slider assembly and made possible a head-to-disk spacing of 50 microinches for the 3330-1, first shipped in 1971, and 35 microinches for the 3330-11, first shipped in 1974. Both products had a bit density of 4040 bpi. The track density of the Model 1 was 192 tpi and that of the Model 11 was 370 tpi.

The head/slider combination developed for the IBM 2305 fixed-head file, a product complementary to the 3330 and also first shipped in 1971, was based upon the concept of a batch manufacturing process developed by E. R. Solyst [16]. In this design, nine head elements and a flat slider with a taper at the leading edge were machined from a single block of ferrite [Fig. 2(b)]. This taper-flat design, with its large air bearing surface, required a loading force of about 1.2 kilograms to achieve a 50-microinch flying height.

Low-mass, lightly loaded sliders, suspensions, and heads

Winchester taper-flat slider The next major step in slider and head development came with the introduction of the IBM 3340, first shipped in 1973. This head combined the taper-flat and batch-fabrication concepts developed for the 2305 with the experience gained at the IBM Los Gatos Laboratory, where a unique file had been developed in the late 1960s using disk technology to provide a refresh buffer for a high-resolution display system. This fixed-head file, developed by Dr. J. T. Ma, was the first to use lubricated iron oxide particulate disks and start-and-stop-in-contact heads. To achieve the linear density required for that buffer, heads developed by Data Disk Corporation for use on a plated disk were used. These heads had three small pads arranged in a triangle with the magnetic element at the apex providing one of

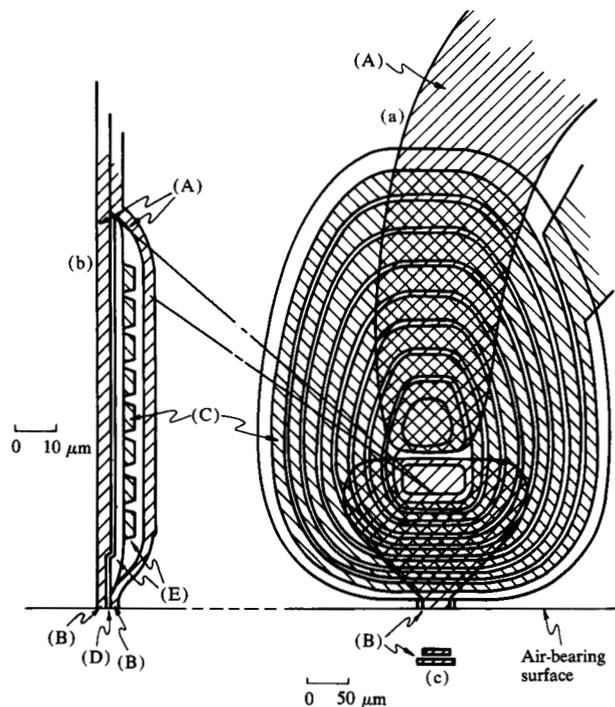


Figure 3 IBM 3370 thin-film head: (a) planar view; (b) center line cross section showing A—magnetic layers, B—pole tips, C—conductor turns, D—gap layer, E—insulation layer; (c) pole tip geometry at the air bearing surface (from Ref. 21).

the pads. They provided a small air bearing area and a low mass, were lightly loaded with about 20 grams by a flexure suspension, and were in contact with the disk surface when it was stopped. In operation they maintained a nominal spacing of 20 microinches. A lubricated particulate disk was developed by IBM to minimize wear during the periods of in-contact operation. IBM obtained rights to produce this tri-pad head design and a new version was developed by W. S. Buslik for use in a disk drive embedded in the IBM 3735 Programmable Buffered Terminal.

Experience had shown that this tri-pad design was difficult and costly to manufacture, and in 1969 a development effort was initiated to design a head that would operate at close spacing but could be produced in quantity with a simple process. The resulting tri-rail design by Warner [17] is known as the "Winchester" head. A small taper-flat bearing area, such as in the 2305, was provided by the outer two rails of the tri-rail design. The center rail defined the width of the magnetic element at the trailing edge where a ferrite core was formed, as in the 2305 design. This head maintained a spacing of about 20 microinches with a load of less than 20 grams. Mulvany [18] gives a more detailed discussion of the evolution of

the Winchester head/slider for the interested reader. The elegant simplicity of this design has given it industry-wide acceptance.

Film magnetic elements Based upon the prior development work of plated magnetic memories, work was begun on film recording heads in the late 1960s. Romankiw, Croll, and Hatzakis [19] built the first operating thin-film head in IBM. The success of these early designs resulted in a development activity to integrate a film magnetic element with a slider bearing in a configuration suitable for product use. Additional contributions by Thompson and Romankiw [20], Jones [21], and others led to the design of the head that was first shipped in the IBM 3370 in 1979. The advantages of this structure, shown in Fig. 3, are the precise control of the gap, the inherently thin pole tips which enhance the sharpness of the readback pulse, and the excellent frequency response of the element. These combine to give a unique potential for enhancement of disk file capacity and data rate.

Design aspects of disk file actuators

The key mechanical components of a moving-head disk file are in the actuator and carriage assembly [22]. Its role is to transport the read/write heads from one track position to another as fast and as accurately as possible. Speed is important because it directly affects seek time, which is a major component of access time. Positioning accuracy is important because it affects the maximum attainable track density. In this section, the progress of actuator and carriage development will be reviewed with emphasis on the improvements made in both seek time and positioning accuracy.

• The IBM 350 actuator

The IBM 350 disk file [2] provides a reference to measure progress. A schematic of the actuator and carriage of the 350 is shown in Fig. 4. Two pressurized air-bearing-supported magnetic heads, gimballed to a pair of arms, were mounted on a carriage that was moved vertically to one of 50 disks 24 inches in diameter and spaced 0.4 inch apart, by a cable connected through two counter-rotating magnetic powder clutches driven by a single motor. A null servo system and then a mechanical detent were used to position the carriage at the selected disk. Detenting the carriage changed the mode of the actuator, permitting the same drive system and a similar null servo to move the two head-arm assemblies to one of 100 concentric tracks in a five-inch radial band. Upon arrival at the selected track, air pressure was applied to detent the arm into its final position and load the air bearing heads against the disk surface preparatory to reading or writing. The detenting system provided a track density of 20 tpi. The average seek time for this two-dimensional actuator was 600 milliseconds.

Hydraulic actuators and comb access structures

The next 24-inch-diameter disk file, the IBM 1301, introduced the self-acting air bearing slider in place of the externally pressurized bearing. This made the development of a comb-type carriage practical, with one self-acting slider for each disk surface. A new hydraulic actuator was developed to increase the access rate, and the comb of heads provided a logical cylinder of data tracks at each access position. This actuator employed two sets of hydraulic cylinders, one each for coarse and fine positioning, and was capable of moving the carriage with its 50 head-arm assemblies from one logical data cylinder to any other with an average seek time of 165 milliseconds. The accuracy of the hydraulic system with detents permitted an initial track density of 50 tpi, later increased to 100 tpi.

Removable disk pack

When it was seen that needs of small systems users could not be met at an acceptable cost using the large components described, a proposal was made for a small disk file, utilizing a removable disk assembly to allow off-line storage and backup [23]. Consequently, a disk drive was developed [24], using six 14-inch-diameter disks mounted in a customer-interchangeable disk pack enclosed in a protective cover that could be removed and replaced only when the pack was in place on the drive hub [25, 26]. The IBM 1311 was the first of a family of files with interchangeable disk packs. This design concept was extended into the early 1970s.

The design of an interchangeable disk pack required tight control of the absolute position and the position tolerances of the data tracks written on each pack. Control of the absolute position of the tracks was established by introducing a special disk pack for use by both manufacturing and field engineering. A cylinder of tracks was very precisely written on this pack using bit patterns chosen for ease in manually adjusting data heads. This special pack permitted all of the heads on all of the drives to be initially aligned to the same absolute cylinder location. The location of all of the other cylinders on the disk pack was then a function of the carriage, disk pack, temperature, and spindle tolerances of the disk file that wrote data on these cylinders.

Table 2 lists the key tolerances affecting track positioning accuracy for several file systems with interchangeable magnetic media. The track positioning accuracy can be determined by combining these tolerances statistically. The result is referred to as the "track misregistration tolerance." Track density in all disk files is determined by the condition that the head, disk, and recording channel off-track capability be greater than or equal to the track

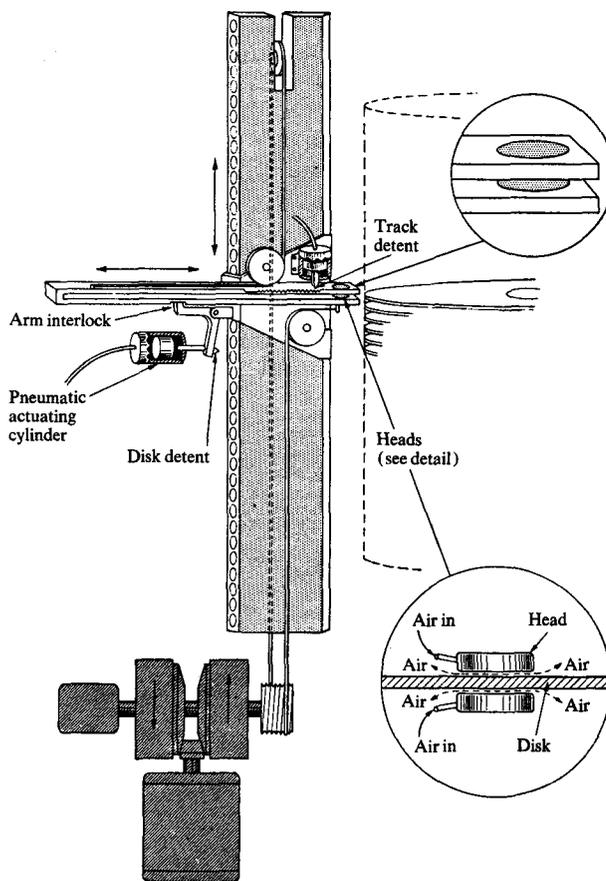


Figure 4 IBM 350 carriage and actuator (from Ref. 2).

Table 2 Position tolerances* affecting track positioning accuracy (microminches).

Item	1311	3330	3340
Spindle eccentricity	250	**	**
Spindle/pack tilt	250	100	15
Carriage—radial position	500	190	} 246
Carriage—circumferential	310	240	
Alignment cylinder	200	100	0
Head adjustment	200	80	0
Vibration	150	100	103
Temperature	500	100	235
Wear	250	50	25
Servo		400	410

*These tolerances are "single-sided"; i.e., a tolerance of 250 means a total error of ± 250 .

**Included under servo.

misregistration tolerance. As already noted, the early disk drives using interchangeable disk packs were able to achieve track densities of 50 tpi.

Disk drives with small removable disk packs required a simple low-cost high-performance actuator and carriage

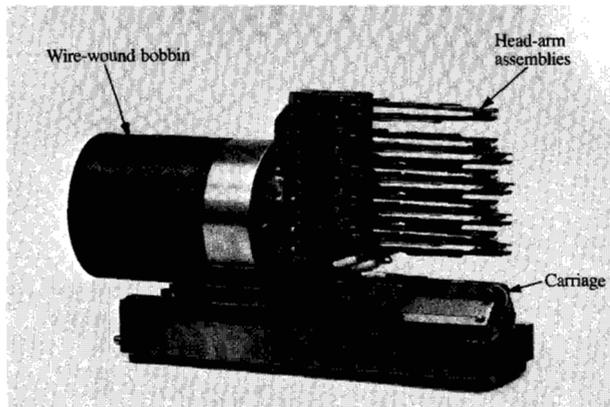


Figure 5 IBM 3330 voice coil actuator.

assembly. An innovative hydraulic actuator was developed that took advantage of the shorter stroke length and reduced carriage mass. The IBM 1311 actuator drove the carriage selectively at one of two speeds determined by fluid flow rates. A high flow rate (and high speed) was used for traveling long distances, and a low flow rate (and low speed) was selected just prior to arriving at the selected data cylinder. At the low speed, a mechanical detent on the carriage could be engaged without creating excessive carriage vibration and overshoot. For short distances, only the low speed was employed. Hydraulic fluid was delivered to this two-speed actuator by a small commercially available gear pump. This combination resulted in a low-cost carriage and actuator assembly with an average seek time of 150 milliseconds.

The basic actuator and carriage concept of the 1311 was also employed on successor products, the 2311 Disk Drive and the 2314 File Facility. Carriage tolerance modifications achieved the more stringent track density requirements of 100 tpi. With the change from a two-speed to a three-speed actuator the average seek time was reduced from 150 milliseconds to 60 milliseconds.

- *Track-following servo actuators*

The 1311 tolerances listed in Table 2 show that the radial positioning accuracy of the carriage and the temperature effects contribute almost 40% of the total. This was so because the radial positioning detent was referenced to the main drive casting rather than a track on the disk pack. There was no compensation for even a portion of the nonuniform and time-varying effects of temperature. By moving the reference from the base casting to the disk pack and providing a feedback loop to control the carriage position, a portion of the tolerance due to temperature could be eliminated.

The IBM 3330 was the first production disk file to incorporate a track-following feedback control system. Feasibility of a similar positioning concept was demonstrated by Hoagland in the early 1960s [27]. The 3330 feedback system consisted of

- One dedicated disk surface with special prerecorded tracks defining the position of each data cylinder.
- A servo head to read the prerecorded tracks.
- An electronic control system that generated a position error signal and translated it into actuator drive current.
- An actuator motor that generated force directly proportional to drive current.

The actuator was made of a wire-wound cylindrical bobbin supported by a carriage (Fig. 5) and moving in a magnetic field structure similar to a loudspeaker (hence the term "voice coil motor"). During seeking, current in the bobbin was controlled by monitoring servo track crossings and resulted in heads arriving at the selected cylinder with very little overshoot. Feedback control was continuously provided in the track following mode, permitting the servo head to follow the radial runout and temperature effects. The signal pattern recorded on the servo tracks [28] was a key factor in attaining accurate track following at a reasonable cost.

Tolerance improvements in other areas are shown in Table 2. The net effect of all tolerance improvements resulted in attaining 192 tpi. The servo control system with the voice coil motor resulted in an average seek time of 30 milliseconds.

With the development of the 3330, the 1311 family had grown from a single spindle disk file of 2.68 million bytes to a dual spindle file of 100 million bytes per spindle. No longer was this file family designed for the first-time file user but rather for the experienced disk file system user, leaving a need for a smaller-capacity and lower-cost entry-level disk file. To meet this need, the IBM 3340 disk file was developed.

The IBM 3340 actuator and data module

The IBM 3340 contained many innovative designs [18]. Key was the low-mass, lightly loaded "Winchester" head and lubricated disks which permitted heads to remain on the disk while it was starting and stopping, thus eliminating the load/unload mechanism. Because of the low cost of the head/slider, two heads per surface could be used. This cut in half the stroke length and lowered the cost of both the carriage and the actuator. The 70-megabyte capacity was obtained with four disks. The disks, the disk spindle and bearings, the carriage, and the head-arm assemblies were incorporated into one removable package called a Data Module [18]. Since each head need only

read the data that it had written—a concept that had been suggested earlier by Buslik [29]—the need for a special disk pack and for adjusting the individual data heads was eliminated. The positioning accuracy tolerances are also shown in Table 2. The IBM 3348 Data Module attained a track density of 300 tpi and an access time of 25 milliseconds.

The IBM 3350, first shipped in 1976, extended the Winchester technology by increasing the number of disks per drive, the bit density to 6425 bpi, the track density to 478 tpi, and the data rate to 1.2 Mbytes/s, resulting in a 4.5 times increase in capacity per spindle to 317.5 Mbytes. One of the key changes that permitted these improvements was the decision not to provide a customer-removable data module. Tolerances associated with removability were eliminated, allowing an increase in track density. Eliminating removability reduced still further the potential for particulate contamination due to handling of the data module.

Rotary actuators

An ingenious actuator [30] developed at the IBM Development Laboratory in Hursley, England, involved a moving-coil rotary actuator, Fig. 6, with pivoting data arms. Although the pivoting data arm concept was not new, this rotary actuator and carriage, with only one moving part, provided most of the advantages of the linear voice coil actuator [31], but at very low product cost. As a result this actuator using the Winchester head technology, combined with a nonremovable disk in a sealed enclosure, has been used in imbedded disk files for several small systems such as the IBM System/32. The heat generated by this small unit was low enough to be dissipated directly through the enclosure, without requiring external air cooling, thus providing low-cost protection from external contamination.

The latest small IBM disk drive developed at the Hursley Laboratory, the IBM 3310, also utilizes a rotary actuator and carriage with a dedicated servo surface. In addition, it employs a sectorized track-following servo concept [32, 33]. Each data surface has prerecorded feedback information located in small radial sectors equally distributed around the disk. During seek operations the dedicated servo surface is used to provide feedback information, but when the desired cylinder is reached, position information from the selected data head is combined with the information from the dedicated servo surface to control the carriage position. The control electronics switches data and servo information to appropriate circuits and samples and holds servo information between sectors. This sector servo control virtually eliminates temperature effects and most of the tolerances between the dedicated servo head and the data head.

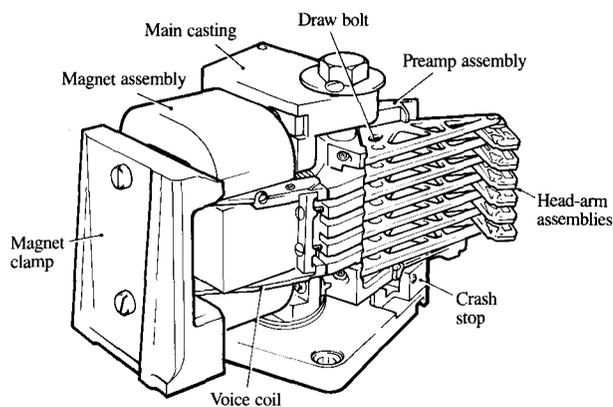


Figure 6 Rotary actuator used in the IBM 62 PC.

Disk substrates and magnetic coating

Substrates In early disk file development the major concern with the substrate was its axial runout. The first disk file used an aluminum substrate of "recording stock" quality and disks were fabricated by laminating two half-disks, 0.051 inch thick, together with an adhesive. The coating on these disks was relatively thick, 1200 microinches, masking the effects of substrate imperfections.

The requirements for higher-density recording demand that the head be closer to the surface and the magnetic coating thickness be reduced. As the thickness of the magnetic coating and the flying height of the head/slider decreased, the effects of the substrate became more important. Various polishing, grinding, and machining techniques were developed to improve flatness and smoothness. A key advance in the fabrication of substrates was the development of a diamond tool-lathe process in conjunction with the use of a much purer aluminum alloy [4]. This process, developed by the Manufacturing Group in Sindelfingen, Germany, established a technology base that has provided substrates for all subsequent generations of IBM disk products.

New and better measurement techniques were developed. The measurements of the first and second derivatives of runout were important in predicting the excursion of the head over local surface imperfections. With the introduction of lower flying heights and increasing rotational speed of the disk, the resonant frequencies of the substrate became important. Rotational speed and substrate thickness were selected to avoid these frequencies.

Magnetic coating The magnetic coating is a dispersion of magnetic particles in a binder. In the early 1950s the magnetic particles were limited to the gamma form of iron

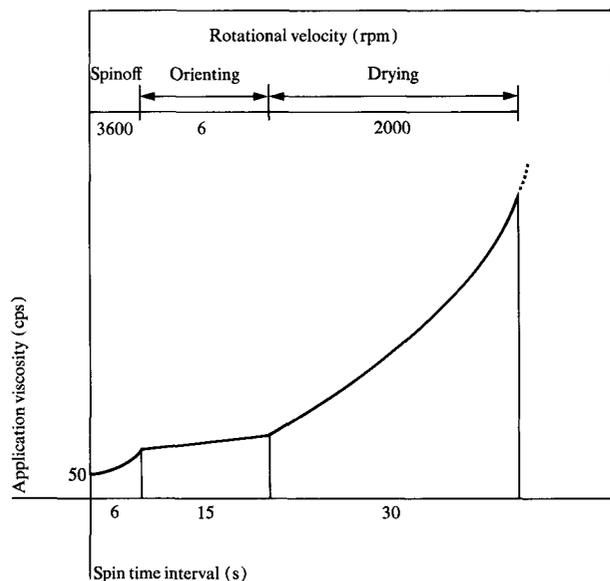


Figure 7 Plot of time required to orient magnetic particles during coating a disk, vs the fluid viscosity of the magnetic medium.

oxide, and polymer chemistry was in its infancy, providing few choices of polymeric materials to satisfy the rigorous demands for a disk recording medium. The key demands on a binder material are

- Ready dispersion of magnetic iron oxide particles.
- Viscosity low enough to spin-coat on an aluminum substrate.
- Curing to form a crosslinked, hard, durable material.
- Good adhesion to the substrate.
- Smooth surface polish with few pinholes or other imperfections.

Literally thousands of combinations were tried through much empirical study and diligent experimentation. The work culminated in a unique binder system described by Johnson, Flores, and Vogel [34]. The choice of a combination of an epoxy, a phenolic, and a polyvinyl methyl ether proved to be an excellent selection and with modification is still the basic formulation in use today [4]. Similarly, the original spin coating used for the first disks continues to be the predominate process for most disk production today.

The magnetic properties of commercially available magnetic particles have continued to improve over the years. Good control in particle geometry yields materials that facilitate particle orientation within the recording medium. Structural defects within particles have also

been improved by minimizing pores and nonmagnetic inclusions between the individual crystallites of a particle, thus leading to better control over the range of particle-switching fields.

Magnetic particle orientation Orientation of the oxide particles as routinely practiced in magnetic tape had long been an objective of disk development. The geometry of a disk and the requirements of performing the orientation on the spin coater proved a difficult hurdle for many years. E. M. Williams found that a new iron oxide particle from Pfizer had high acicularity and a narrow switching field distribution. This oxide and the concurrent work on orienting particles on disks by P. T. Chang, A. W. Ward, and W. N. Johnson led to a new disk product with substantial improvement in recording properties. The oriented disk was first introduced in the 3340 Data Module and has been used in all subsequent disk file products.

The key parameters for orientation are the fluid viscosity of the media, the time required to orient the particles, and the magnetic field applied to the particles. It is necessary to carefully control both the spinoff conditions and duration as well as the application viscosity of the fluid, as indicated in Fig. 7. Because the orienting field is from the gap fringe field of a magnet, both the strength of the field and the time the particles "see" that field become the critical factors for good orientation. With a pair of opposing magnets on each side of the disk so that the fields are essentially in the plane of the disk, it was found that a gap field of 1800–2000 oersteds [$1 \text{ Oe} = (1000/4\pi) \text{ A}\cdot\text{m}^{-1}$] was required. If the field was weaker than that, there was insufficient force to rotate the particles in the viscous medium. If the field was much bigger than that, the particles simply switched their magnetization polarity rather than being rotated. It was also necessary to rotate the disk slowly for maximum orientation. If the disk were rotated too rapidly, the particles switched their magnetization polarity, and if too slowly, it caused coating imperfections.

Lubrication The development of the low-load Winchester slider, designed to be in contact with the disk surface during starting and stopping, put new requirements on the surface characteristics of the magnetic medium. These requirements were met by the development of lubricating techniques for disks similar to those previously used on magnetic tape. The first lubricant used a 0.75% solution of Dow Corning RDC-200 in a freon solvent, and was applied topographically to the disk with a soft absorbent cloth. This disk proved to be very reliable and led to the development of disk file products utilizing the concept of a lubricated surface with start-and-stop-in-contact heads.

In succeeding years, numerous other lubricants were evaluated to meet the demanding needs of the newer high-performance files.

Data encoding and channel electronics

Advances in encoding methods and recording electronics have also contributed to increasing disk file areal densities. As linear bit densities have increased from 100 bpi to 15 000 bpi and rotational velocity from 1200 to 3600 rpm, the serial data rate at the read/write head has increased from 100 kilobits/s to 24 megabits/s. These data rates have, in recent generations, pressed the data rate capability of host system data channels, and in addition have provided a challenge in the design of disk file electronics.

For writing data, the electronics converts logic-level data signals into a current that drives the recording head. Later, when data are being read, the process must be reversed. Analog signals induced in the coils of the read head must be amplified, detected, and clocked. Reliable recovery of stored information involves both time detection and amplitude detection. Encoding methods have been designed to place more emphasis on either one or the other of these two dimensions of detection, depending upon the disk and head magnetic and electrical characteristics.

The first disk file utilized NRZI encoding [35], developed initially for tape drives. Data were recorded by changing the direction of current flow in the head for every 1-bit, with no change in current direction for a 0-bit. Sufficient current was applied to saturate the medium. Since the readback signal is proportional to the rate of change of recorded field, a pulse was read back for every 1-bit. Presence of the signal was detected when the head voltage exceeded a predetermined clipping level. This amplitude detection was capable of discriminating against background noise, base-line overshoot, and extra bit defects. The data-clocking circuits, developed by Seader [36], used two fast-starting oscillators, each started by alternate 1-bits. After starting, the cycles of one oscillator were used to clock any following 0-bits. The second oscillator was started when the next 1-bit was detected. The use of an odd redundancy bit and a space bit at the end of each character limited the free-running cycles of an oscillator to at most one character.

The next disk file used a combination of techniques for developing a clock. A separate clock track provided a stream of periodic pulses with the rotation of the disk. Each bit time was divided into four phases by delaying the base clock three times. At the beginning of each byte or character of data, one of the four phases was selected to clock the remaining data bits. Since the clock track

signal was locked to minor disk speed variations, and with the time reference re-established at the beginning of each byte, the expected time variation between a clock phase and data from a head was small. Furthermore, mechanical variations due to vibration and temperature differences between the clock head and data head, which could cause clocking errors, were low in frequency relative to the data stream. The major drawback of this system was the cost of a clock head and read amplifier together with logic to perform the phase selection. The amplitude detection system was improved by the addition of automatic gain control to eliminate the wide range of signal amplitude variation caused by different head and disk configurations. As bit densities increased, so did data rates, and it became more important to generate a clock that was synchronized with the data to provide additional clocking tolerance.

The variable-frequency oscillator (VFO) was a circuit which offered the potential of reducing many tolerances experienced by previously described circuits. The VFO had been developed by Newman [37] for tape but was redesigned for the high data rate of disk files. The circuit was an oscillator whose frequency could be electronically adjusted with a feedback loop. Recorded pulse rates were compared with an internal oscillator. If the oscillator was operating at too low a frequency, an error signal was generated to increase the oscillator frequency. The circuit eventually locked onto the periodicity of data pulses, but lagged a slight amount to generate a continuous error signal. An improvement of this circuit by Lang, LaPine, and Vaughn [38] integrated the error signal and drove the phase as well as the frequency error to zero. The averaging properties of this circuit reduced to almost zero those errors caused by setting the clocking window. The primary tolerance of clocking was now associated with the data bit rather than the location of the clocking window. Initially, double-frequency encoding was used, providing a clock pulse with every bit cell as a point of reference.

The clocking accuracy of a VFO operating at high frequencies permitted utilization of an encoding method with reduced transitions but required two clock cycles per bit cell. The structure of the code made it slightly more susceptible to noise, and a new detection system known as the delta-V detector was developed by Graham, Scovmand, and Swartz [39] to maintain and improve reliability. The delta-V detector used a series of tests to improve the signal rejection capability over simple amplitude discrimination. A minimum leading edge slope was required, and the signal peak was detected by finding the derivative crossover point, and a minimum trailing slope was sensed. Finally, the amplitude of the signal associated with the pulse had to exceed a minimum level.

Through the use of a delay line and logic, these four conditions were tested, and if they were satisfied, an output pulse was generated.

More recently, Franaszek-type run-length-limited codes [40] have been implemented by Eggenberger and Hodges [41]. These codes limit the minimum and maximum distance between transitions. For example, the 2,7 code has a minimum of two zeros and a maximum of seven zeros between transitions. At the same time they use the transitions with such efficiency that on the average 1.5 bits per transition is achieved. A new method of detection was developed by D. B. Chapman for use on recent IBM disk products with such a 2,7 run-length-limited code. The signal is tested to see that the polarity of the pulses alternates and that a minimum amplitude is achieved. The signal peak is found through differentiation, and finally, a minimum signal change, occurring after the peak, must occur within a specific period of time. These four criteria are logically brought together to make a detection decision.

As track densities increased, the readback signal amplitude was reduced. Furthermore, increasing data rates led to introduction of an electronics module on each access arm. This module provided the write current as well as amplification for the read signal. Since every arm assembly was designed to use one module containing a total of 150 high-speed bipolar circuits, low cost was important to make this system practical.

Summary

When the first disk file was conceived in IBM's newly established development laboratory in California, no one visualized how far the technology would be extended. The first disk file introduced the externally pressurized air bearing as a support for the magnetic head, the 24-inch disk, the coating technology, and a two-dimensional access mechanism. The starting point of areal density was 100 bpi linearly and 20 tpi radially, or 2000 bits per square inch.

The next generation introduced the use of self-acting bearings as a head carrier and the high-performance comb actuator. Areal density was increased to 100 000 bits per square inch (100 tpi and 1000 bpi), a factor of 50 improvement. This file, part of the IBM Sabre system for airline reservations, was key to the start of interactive processing as it is known today.

The disk pack files extended usage to low-end systems. It made the disk drive the primary storage for such systems. The interchangeable pack permitted it to serve

both tape and disk functions. A number of the small systems were supported by disks exclusively.

Further increases in capacity resulted from continuing work in mechanics, electronics, materials, and processes, which further improved storage density. This continuing evolution lowered the cost of on-line disk storage to a point where it became the work storage for all data processing systems, large and small alike, and permitted the expansion of on-line applications in the 1960s and early 1970s. In addition to substantial improvements in heads, disks, and recording electronics, the introduction of the track-following servo improved both performance and track density significantly.

Introduction of the Data Module drive with heads in contact with the disk as it was started and stopped provided a basis for yet another step increase of density.

With the introduction of the IBM 3350, a full circle of disk file concepts had been made. Nonremovable data disks had first been introduced in the IBM 350 design. Next came the "unlimited capacity" family with customer-removable disk packs or data modules. As the storage technology improved over the years, a removable disk file product which held two million bytes in 1963 could, in 1974, hold 200 million bytes, greatly reducing the customer needs for off-line storage. With the significant increases in reliability and availability made during the same period, users no longer had a need for disk interchangeability to meet data availability requirements. Therefore, the IBM 3350, with over 300 million bytes per spindle, was the first file of the 1970s to have nonremovable disks. That this way of reducing tolerances led to further increases in density is witnessed by the fact that the most recently announced disk file product, the IBM 3380, shipped in 1981, is also a fixed-disk product with 1250 million bytes of storage per spindle—a vast improvement over the five million characters first introduced a quarter century earlier.

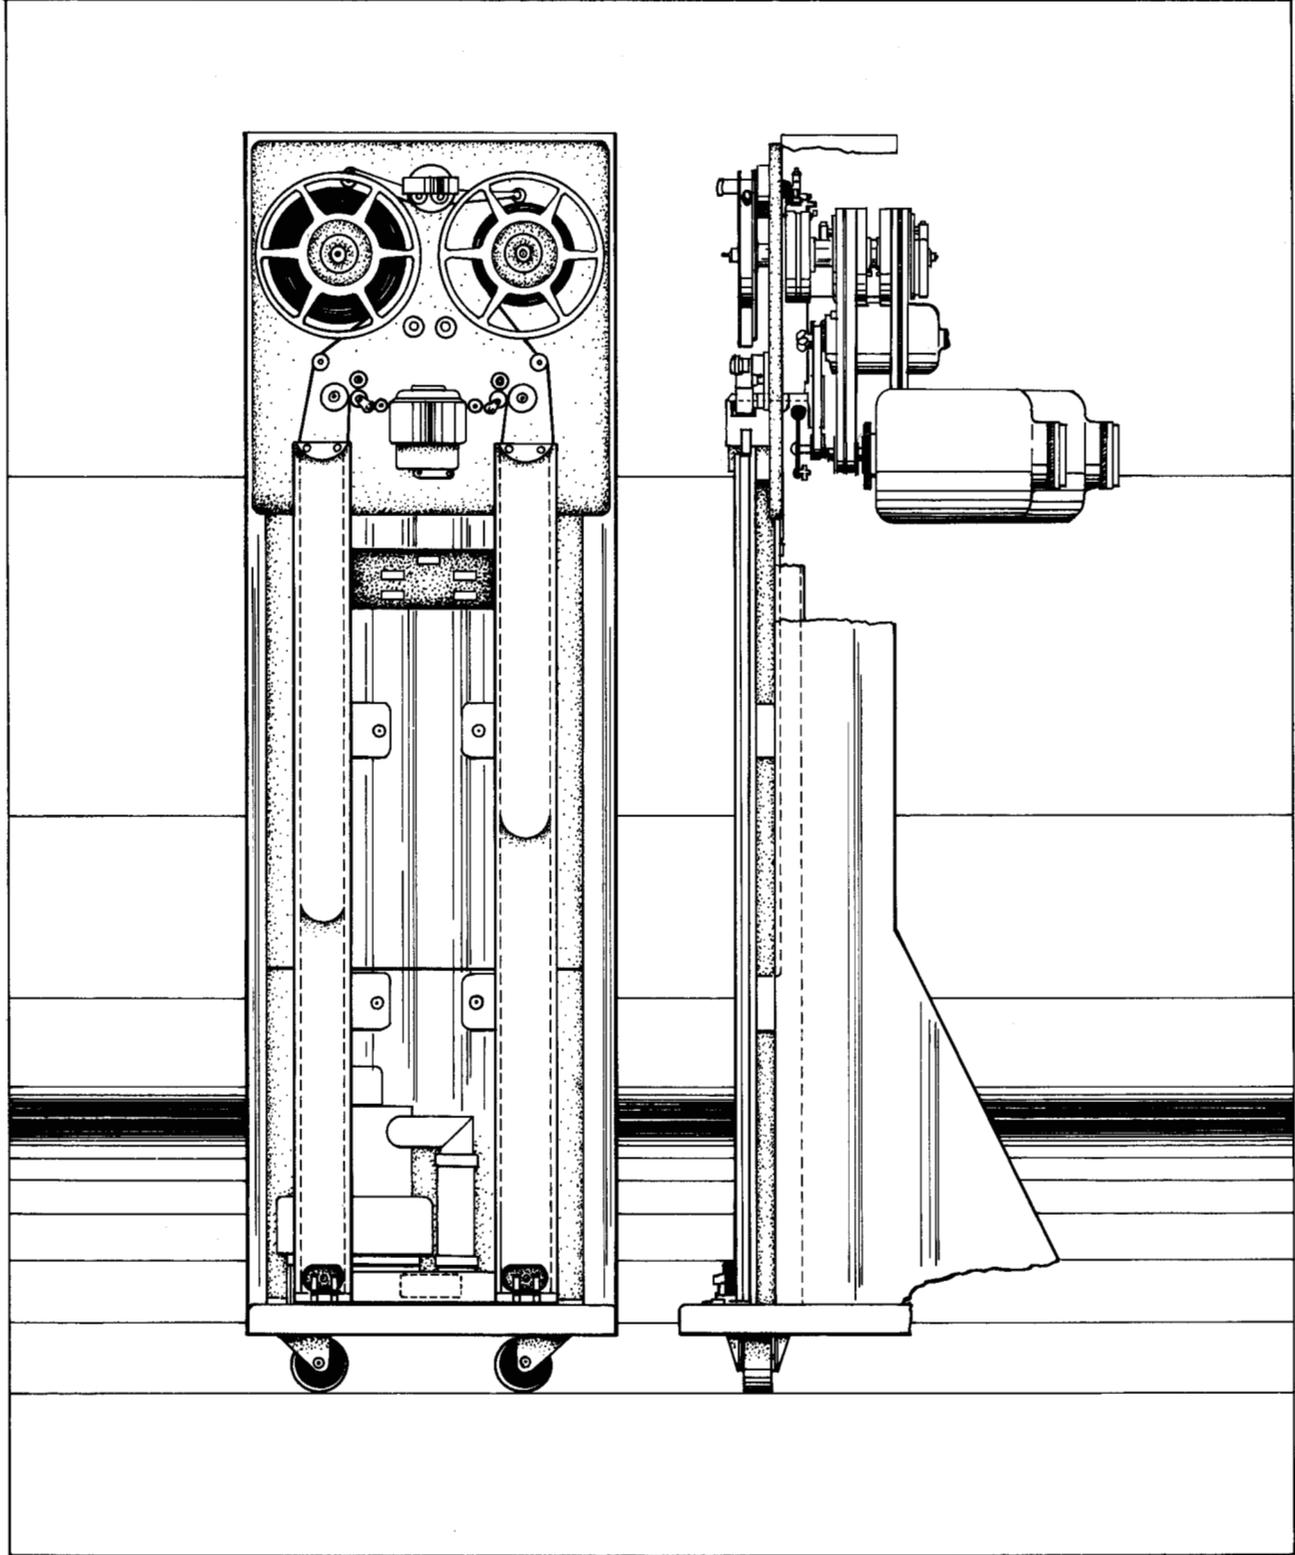
References

1. 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).
2. 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).
3. L. D. Stevens, "The Evolution of Magnetic Storage," *IBM J. Res. Develop.* **25**, 663-675 (1981, this issue).
4. R. B. Mulvany and L. H. Thompson, "Innovations in Disk File Manufacturing," *IBM J. Res. Develop.* **25**, 711-723 (1981, this issue).
5. M. E. Wolf, "The R & D Bootleggers: Inventing Against the Odds," *IEEE Spectrum* **12**, 38-45 (1975).
6. A. Kingsbury, "Experiments With an Air-Lubricated Journal," *J. Amer. Soc. Naval Engineers* **9**, 267-292 (1897).

7. W. A. Goddard, "Air Bearing Head," U.S. Patent 3,213,461, 1965.
8. N. A. Vogel, "Air Head," U.S. Patent 2,886,651, 1955.
9. T. G. Cowan, "Mounting and Operating Means for an Air Head," U.S. Patent 3,018,152, 1958.
10. W. A. Gross, "A Gas Film Lubrication Study, Part I: Some Theoretical Analyses of Slider Bearings," *IBM J. Res. Develop.* 3, 237-255 (1959).
11. 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).
12. W. E. Tibbetts, "Tunnel Erase Magnetic Transducer," U.S. Patent 3,155,949, 1961.
13. E. R. Solyst, "Mounting of a Transducer," U.S. Patent 3,329,379, 1967.
14. D. W. Brede, M. H. Cook, M. E. Freeman, and H. L. Turk, "Glass Bonded Head Assembly," U.S. Patent 3,610,837, 1971.
15. R. A. Applequist, "Magnetic Head Spring Bias," U.S. Patent 3,579,213, 1971.
16. E. R. Solyst, "Multichannel Head," U.S. Patent 3,579,214, 1971.
17. M. W. Warner, "Flying Magnetic Transducer Assembly Having Three Rails," U.S. Patent 3,823,416, 1974.
18. R. B. Mulvany, "Engineering Design of a Disk Storage Facility with Data Modules" *IBM J. Res. Develop.* 18, 489-505 (1974).
19. L. T. Romankiw, I. M. Croll, and M. Hatzakis, "Batch Fabricated Thin-Film Magnetic Recording Heads," *IEEE Trans. Magnetics* MAG-6, 597-601 (1970).
20. David A. Thompson and Lubomyr T. Romankiw, "Film Head Development," *Disk Storage Technology*, 3-5 (1980); Order No. GA26-1665-0; available through IBM branch offices.
21. Robert E. Jones, Jr., "IBM 3370 Film Head Design and Fabrication," *op. cit.* Ref. 20, pp. 6-9.
22. W. A. Goddard and J. J. Lynott, "Direct Access Storage Device," U.S. Patent 3,503,060, 1970.
23. J. M. Harker, "Objectives for a Low Cost Random Access File Program," Internal IBM report, February 9, 1959.
24. 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 Conference Proceedings* 24 (1963 Fall Joint Computer Conference, Las Vegas), 327-340 (1963).
25. R. E. Pattison, "Portable Memory for Data Processing Machine," U.S. Patent 3,176,281, 1965.
26. T. G. Leary, "Transporting and Protecting Cases for Drum and Disk Records," U.S. Patent 3,206,214, 1965.
27. A. S. Hoagland, "A High Track-Density Servo-Access System for Magnetic Recording Disk Storage," *IBM J. Res. Develop.* 5, 287-296 (1961).
28. G. R. Santana, "Method and Apparatus for Recording and Detecting Information," U.S. Patent 3,534,344, 1970.
29. W. S. Buslik, "Magnetic Disk Storage File in Sealed Enclosure," U.S. Patent 3,710,357, 1973.
30. D. E. Cuzner, C. O. R. Dodman, J. S. Heath, and L. J. Rigby, "Magnetic Disk Apparatus," U.S. Patent 3,849,800, 1974.
31. J. S. Heath, "Design of a Swinging Arm Actuator for a Disk File," *IBM J. Res. Develop.* 20, 389-397 (1976).
32. R. J. Sippel, "Transducer Positioning System Utilizing Record With Interspersed Data and Positioning Information," U.S. Patent 3,185,972, 1965.
33. Robert D. Commander and J. R. Taylor, "Servo Design for an Eight-Inch Disk File," *op. cit.* Ref. 20, pp. 90-98.
34. D. D. Johnson, R. Flores, and M. J. Vogel, "Composition of Epoxide Resin, Methylol Phenol Ether, Polyvinyl Methyl Ether, and Acid Anhydride Catalyst, and Metal Substrate Therewith, Especially a Magnetic Signal Storage Device," U.S. Patent 3,058,844, 1962.
35. B. E. Phelps, "Magnetic Recording Method," U.S. Patent 2,774,646, 1956.
36. L. D. Seader, "A Self-Clocking System for Information Transfer," *IBM J. Res. Develop.* 1, 181-184 (1957); also U.S. Patent 2,864,078, 1958.
37. E. G. Newman, "Magnetic Recording System," U.S. Patent 3,197,739, 1975.
38. P. C. Lang, A. N. LaPine, and J. E. Vaughn, "Random Binary Data Signal Frequency and Phase Compensation Circuit," U.S. Patent 3,701,039, 1972.
39. I. H. Graham, H. A. Skovmand, and J. S. Swartz, "Signal Detection Circuit," U.S. Patent 3,631,263, 1971.
40. P. A. Franaszek, "Run-Length-Limited Variable Length Coding with Error Propagation Limitation," U.S. Patent 3,689,899, 1972.
41. J. S. Eggenberger and P. Hodges, "Sequential Encoding and Decoding of Variable Word Length, Fixed Rate Data Codes," U.S. Patent 4,115,768, 1978.

Received November 5, 1980; revised March 5, 1981

The authors are located at the IBM General Products Division laboratory, 5600 Cottle Road, San Jose, California 95193.



Tape feed mechanism—Patent issued to J. A. Weidenhammer, October 9, 1962