R. B. Mulvany

Engineering Design of a Disk Storage Facility with Data Modules

Abstract: The design of the IBM 3340 Direct Access Storage Facility with IBM 3348 Data Modules incorporates new concepts and required the development of several innovative components, including newly designed magnetic read-write heads. The heads start and stop in contact with the disk and use a tri-rail, air-bearing slider having a low mass. Each data module includes read-write heads, a head carriage, disks, and a disk spindle.

The rationale is discussed for the design concepts and for several components, including the data module, head and arm assembly, and the moving-coil linear actuator. A method of improving data integrity, utilizing a "disk-defect skipping" procedure, is described and its performance implications discussed.

Introduction

Since the introduction of the IBM RAMAC 305 in 1956 [1, 2], disk facilities have been interrelated with the evolution of computer systems. Customer-removable and interchangeable disk packs were introduced as a feature of the IBM 1311 in 1962. Advances in disk areal density and capacity per disk spindle offered customers reduced cost per megabyte of on-line data, a reduction by a factor of two-thirds in the interval from 1967 to 1970. Increases in data rate, reductions in access motion and rotational delay, and improved reliability have enhanced system performance and availability.

High performance and lower cost per Mbyte, however, have been more difficult to achieve in small and intermediate-sized disk drives. Although the number of heads and disks might be reduced, the cost of servo controls, power supplies, motors, etc., does not diminish in proportion to the reduction of data capacity. The IBM 3340 was developed to address this problem and to apply promising new technologies.

In this paper we discuss the development of the design concept of the 3340, several new disk file components, and a new technique to improve data integrity.

Considerations of system requirements

The initial studies to define the disk facility requirements were directed toward small and intermediate systems. The performance and cost per Mbyte were to be similar to those of the IBM 3330.

An important consideration for a new product was to assess the need for data removability. The removable disk pack has offered many advantages to the user, including the ability to change data-set libraries between applications, the availability of data-set backup and security, and the interchange of data among alternate drives and systems. Although the trend is to keep more of the data on line, data removability remains an important function and became an accepted requirement for the 3340 design.

Data removability, however, is not without cost. These costs include disk loading, hardware, added contamination protection, the head-carriage retraction mechanism, and control electronics.

The major cost of removability lies typically in alignment of the drive read-write heads. Head-to-data track alignment must be ensured so that a disk, once written, may be removed and subsequently read on any drive. Radial alignment of the head and track must be accomplished within a fraction of a track width, typically ±20 percent. This radial-alignment tolerance must include static tolerances associated with head alignment differences between drives as well as dynamic tolerances such as spindle-bearing runout, disk-axial runout, carriage-bearing runout, thermal effects, vibration, and tolerance of the access servo system position.

To minimize static radial-position tolerances, head-arm assemblies are adjustably mounted to their carriages and are aligned with the aid of a special disk pack. This disk pack provides radial-position reference tracks for the heads and is used in manufacturing and in customer installations. The requirement for precision adjustment

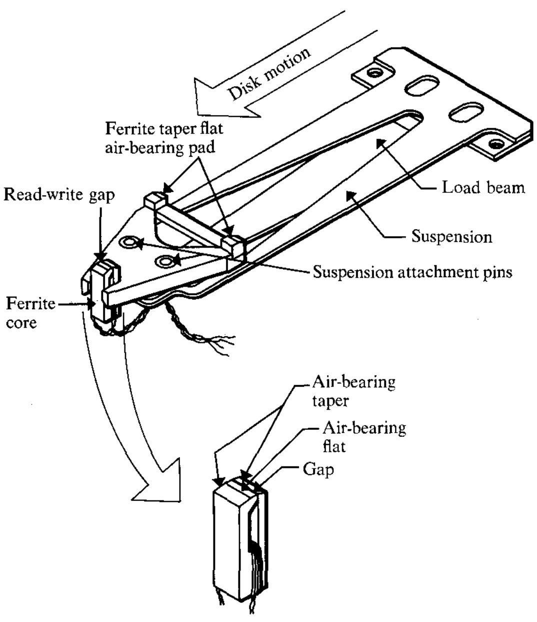

Figure 1 Read-write head used in IBM 3735 Programmable Buffered Terminal.

of head radial position adds significantly to the manufacturing and service costs of conventional removable-disk facilities.

Skew alignment of the data-head gap must also be controlled. Angular misalignment of the reading gap from the writing gap reduces readback amplitude and may introduce data clocking errors. Skew alignment of the head gap is typically measured by optical means and is adjusted to the arm assembly during manufacturing, thus adding to manufacturing cost.

Evaluation of the various requirements suggested that the desired results could be achieved only by new approaches in drive configuration and by new technologies.

In parallel with the requirements study, a review was made of available and potential new developments. One new technology was an air-bearing-slider magnetic head design having a low load and suitable for starting and stopping in contact with a recording disk. Another promising technology was the track-following servo access system then under development in the 3330 program. The utilization of newly developed technologies played an important role in the design of the 3340.

Development of data module concept

As the 3340 developed, performance and capacity improvements extended its application to all IBM System/370 models.

In first considerations of disk removability, it was assumed that the read-write heads would, as in previous drives, remain with the disk drive. This feature required a means to support and align the low-load head and its highly compliant suspension so that the head could be removed from the disk after disk motion stopped. The problem disappeared with the development of the data module concept, which emerged from recognizing the benefits anticipated from a recording head based on the new technology.

The low-load head was under development for other programs at the San Jose laboratory. One program to use the new head involved the development of a single-disk buffer. This application benefited from the high resolution provided by the 0.64-μm (25-μin.) flying height of the head and from the economy gained by eliminating conventional head loading-unloading mechanisms. The IBM 3735 buffer head is illustrated in Fig. 1.

The small air-bearing area of the three taper-flat pads permits the use of relatively low loading forces. This design permits contact between head and disk during disk start and stop, when an epoxy-iron oxide disk coating is used.

Low-load contact heads have been disclosed previously by Kohn [3] and by Hansen and Atamian [4]. These heads were designed to ensure contact with the recording surface during the recording operation. The 3735 head and the subsequent IBM low-load heads have been designed to fly as conventional hydrodynamic air bearings during the recording operation and to contact the disk only during start and stop. The development of the 3348 head and arm is discussed further in a subsequent section.

Another program to use the new slider design was an exploratory development under the direction of IBM Fellow W. S. Buslik. This program was developing a minimum-cost, random access disk drive of five-Mbyte capacity. The drive contained a single head and benefited from the high resolution, small size, low mass, and simplicity of the low-load head. The drive was packaged as a sealed, self-contained unit supported by the spindle drive motor [5].

The data module development was encouraged by the small, self-contained configuration of Buslik's drive, the ability of the new head to start and stop in contact with the disk, and the anticipated low cost of the head. The concept provided a new partitioning of the conventional disk facility with the module containing the mechanical elements required for head-to-track alignment. A range of capacities and performance enhancements was proposed that would use a common basic drive design with a family of data modules.

The basic concept of the data module focused on the utilization of heads dedicated to permanently assigned disk areas. The data module would incorporate a single servo reference disk surface and the data heads would have a fixed relationship to the servo reference head. Each head would read only data that it had written, thus avoiding the need for precise radial or skew alignment of the recording-head gap.

To maintain proper head-to-track alignment, the following components were to be included in the data module: servo head and arm, data heads and arms, head and arm carriage, carriage way, disks, disk hub and spindle, spindle bearings, and a module structure to support the way and spindle bearings.

Two data heads per data disk surface were planned, to minimize access stroke and average access time. The linear actuator moving coil was to be supported and aligned by the data module carriage when coupled for access operation.

The drive interface now includes the following major functions:

1. Access linear actuator and carriage coupling.
2. Actuator coil support and coupling actuation mechanism.
3. Read-write and servo head electrical interface connector.
4. Spindle drive pulley and belt.
5. Data module insertion guides and registration points.
6. Carriage latch operator.

Figure 2 illustrates the functional partitioning as originally proposed and now implemented in the 3340 drive and 3348 Data Module.

Reliability of the disk facility was expected to benefit from the elimination of both high head-load forces (normally 300-400 g) and the conventional head load-unload operation, the contamination protection offered by a sealed head-disk environment, and the elimination of head radial alignment adjustment.

The data module concept offered a means of addressing a number of application requirements by providing a family of data modules based on a common mechanical and head-disk design, all usable on a basic disk drive. Plans were made to offer variable numbers of heads and disk surfaces for capacity options, as well as fixed heads in addition to accessing heads for performance enhancement options.

• *3340 program definition*

Following the configuration proposal, plans for a product program were started while preliminary design of the module continued. Original proposals called for the use of 66 tracks/cm (167 tracks/in.) (similar to the Buslik drive) and five Mbytes per disk surface. Bit density was selected to match the planned System/370 Model 115 and 125 data rate capability of 806 kbytes/s with a rotational speed of 2400 rpm. A range of data module capacities from five to 20 Mbytes was first evaluated. The classic problem of achieving low cost per Mbyte in a small, high-performance disk drive soon increased capacities to the range of 15 to 30 Mbytes. This met the requirements of the smaller System/370s, within the desired cost of the disk facility. The new head could be applied to a number of small-capacity disk applications, but it did not appear cost effective to combine a high-performance servo access system with a removable data module for five- to ten-Mbyte files.

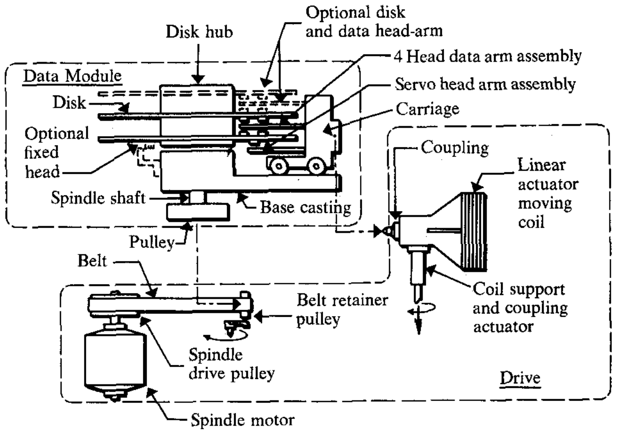

Figure 2 Configurations of IBM 3348 Data Module and IBM 3340, illustrating the functional partitioning.

The product program first called for a dual 30-Mbyte drive unit. Its characteristics are noted in Table 1.

During the completion of the first drive prototype, assessment was made of the feasibility of increasing track density above 66 tracks per cm. A statistical model was used to estimate data-head to data-track misregistration under worst-case conditions. Read-write channel evaluation of off-track capability with a range of head manufacturing tolerance variations was combined

Table 1 Characteristics of 30-Mbyte dual drive unit.

Cartridge capacity	30 Mbytes
Heads per surface	2
Data surfaces per cartridge	4
Track density	66 tracks/cm (167 tracks/in.)
Data tracks per head	188
Access time	
Track to track	15 ms
Average	25 ms
Maximum	60 ms
Data rate	806 kbytes/s
Rotational velocity	2400 rpm
Bit density	2465 bits/cm (6276 bpi)

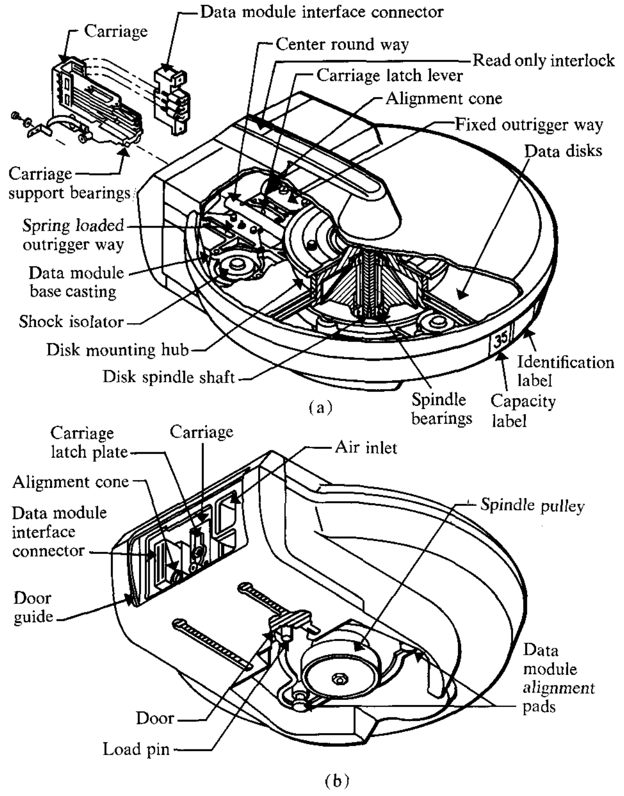

Figure 3 Main design features of IBM 3348 Data Module. (a) Top rear section, exploded and cutaway view; (b) front bottom view, showing pulley, alignment pad, connector, etc.

with the misregistration statistical analysis to determine achievable track density. The results led to a decision to plan for 118 tracks/cm (300 tracks/in.).

Disk rotational delay was reduced by increasing rotational velocity to 2964 rpm. A compromise between maximum bit density and planned data rate for the System/370 was made by increasing the required system data rate to 885 Kbytes/s and reducing the bit density from 2465 to 2180 bits/cm (6276 to 5636 bpi). The new track and bit densities provide a capacity of approximately 12 Mbytes per data surface.

The data module cover, structure, spindle, carriage, and connector were designed for a variety of capacity options, which provided flexibility as application requirements changed during development. Data module capacities varied during the development program and were finally adopted at 35 and 70 Mbytes. The 35-Mbyte capacity also provides a convenient file size for the direct conversion of 2314/2319 data sets. The 70-Mbyte module offers better price and performance, and consolidates data from multiple disk packs and accommodates data sets found in large systems. Another version subsequently announced (Model 70F), provides a capacity of 70 Mbytes, including 0.5 Mbyte available to fixed heads, thus reducing data access time to latency (rotational) delay only for high-activity data such as indices and page data sets.

3348 Data Module

In defining the data module, a number of challenges were recognized. To assure operator convenience, a weight of 6.3 kg (14 lb) was selected as a goal with a maximum limit of 9.1 kg (20 lb). The data module, within these weight limits, would have to provide the required dimensional stability under drive retention force variations, belt load variations, vibration, and thermal gradients to ensure that the planned benefits in head-to-track alignment would be achieved. The high areal density of the 3348 minimized the number of disk surfaces and heads required, thus contributing to the effort to control weight and cost.

The greatest challenge to the data module concept was achieving its cost objective. The typical system user would have multiple data modules per spindle; thus, total data module costs plus drive costs would have to achieve a net reduction in the total cost of the disk facility. Functions normally associated with the drives that were now in the module included the head-arm assembly, the head carriage and way system, and the spindle assembly and drive pulley. These functions, being omitted from the drive, correspondingly simplified its design, and reduced its cost and service requirements.

The main design features of the 3348 Data Module are indicated in Figs. 3(a) and 3(b). Some of the design considerations are included here.

- *Data module structure*

The weight and rigidity requirements noted above were achieved with die-cast aluminum components for disk hub, base plate, and carriage. For the necessary spindle rigidity, a steel shaft is cast in place in the spindle hub. Design simulation using a finite-element method similar to NASTRAN [6] was used to optimize data module rigidity in the presence of vibration caused by access motions and spindle rotation. The end result is a 70-Mbyte data module weighing 8.6 kg (19 lb) and a 35-Mbyte data module of 7.7 kg (17 lb).

- *Carriage bearing way system*

The carriage is supported by two pairs of precision ball bearings designed to roll on a stainless steel way of round cross section. The round cross section provides an economical part with a high surface finish and minimum resistance to carriage motion. The plane of rotation of the carriage bearings in each pair are opposed to one another in a 90° relationship to position the carriage along its horizontal and vertical axes. Rotation of the carriage

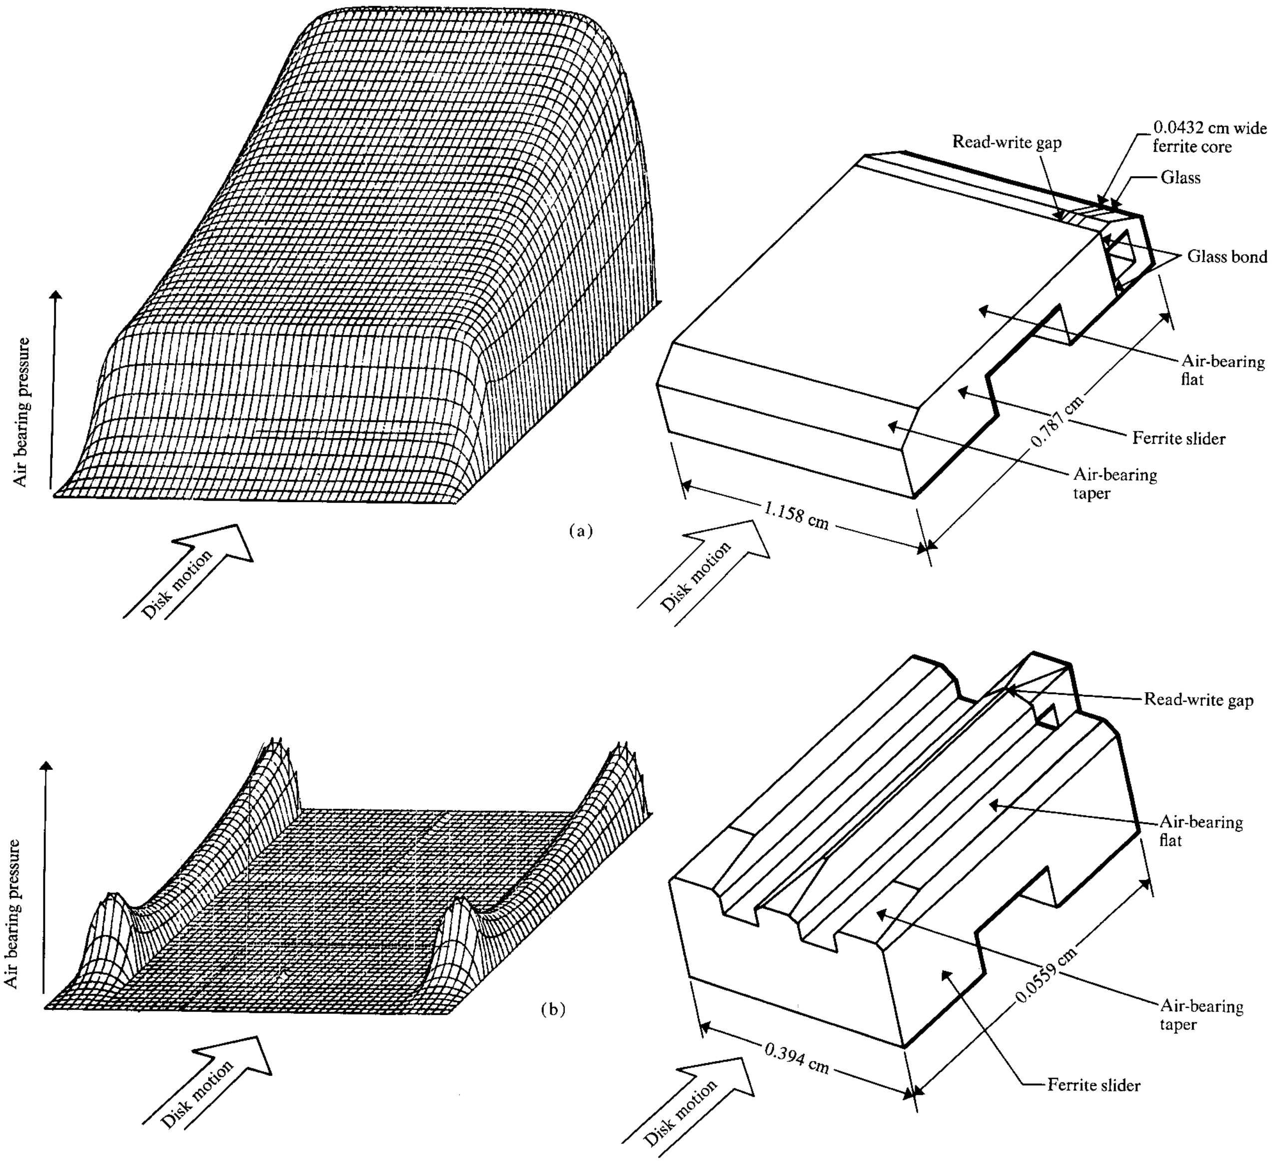

Figure 4 Pressure profiles of read-write heads. (a) Profile for the IBM 2305 head, which is designed with single taper-flat air bearing; (b) profile of IBM 3348 head, with tri-rail ferrite slider.

around the way is restrained by the action of two outrigger ball-bearing rollers.

- *Data module cover*

Mechanical and contamination protection for the data module is provided by a clear, impact-resistant polycarbonate cover mounted to the base plate by means of shock isolators. An elastomer diaphragm seals the cover to the bottom of the base plate in the area adjacent to the data module pulley. The cover incorporates a flexible sliding door similar in concept to a rolltop desk. Low-friction plastic guide strips are provided in the door cover to avoid wear and contamination particles. The door is opened and closed by the drive assembly.

Because the data module cover remains with its disks at all times, its data set identification label (and cover) cannot be inadvertently attached to another set of disks. A read-only interlock (similar in function to the write-protect ring on tape reels) is located in the carrying handle of the data module. The interlock is a small, two-position, rotatable insert that interacts with a sensing switch on the drive.

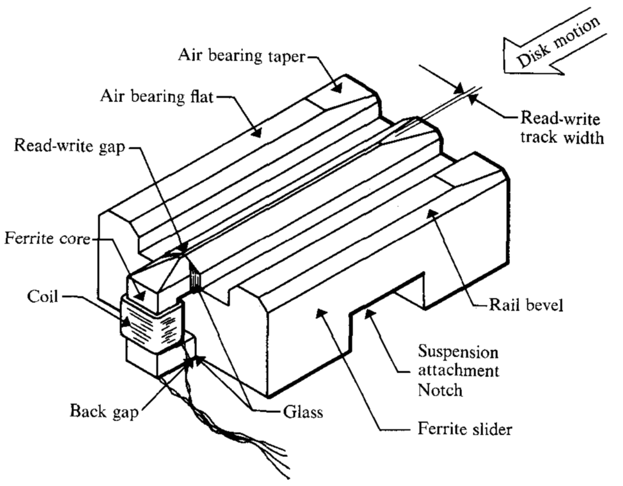

Figure 5 Tri-rail slider, showing core winding and glass-bonded core.

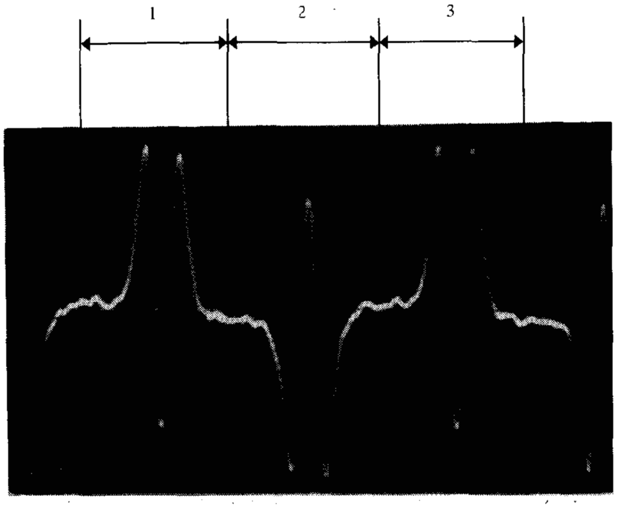

Figure 6 Typical readback signal, 3348 head. For method of determining resolution, see text.

- *Data module identification*

The type of data module (Model 35, 70, or 70F) is identified by means of the data module connector. Three connector pins are selectively interconnected at the data module assembly to provide identification encoding, which is read by the drive at the end of the module loading cycle.

Read-write head and arm

This section discusses the development of the 3348 slider and describes the design of slider, suspension, and head-arm assembly.

The 3348 head utilizes a tri-rail, air-bearing slider having a low load and constructed with two ferrite parts. The head read-write gap is incorporated in the center slider rail. This slider contacts the disk during disk start and stop and is supported by its air bearing during access and read-write operations.

The 3735 low-load head (Fig. 1) utilized a barium titanate slider structure with the ferrite head element and two ferrite pads shaped to provide three taper-flat air bearings. The head required the generation of three precise taper planes and the use of two glass-bonding steps. One glass bond was required for the front and rear gap of the conventional ferrite head core. The second glass bond was required to bond the head element and two ferrite air-bearing pads to the barium titanate slider.

A ferrite slider head assembly requiring only a single glass bond for the transducer gap had been successfully applied in the IBM 2305 [7]. This head utilized a high-density ferrite slider with excellent wear characteristics and suggested a design approach for the new data module that would offer functional and process cost advantages. The 2305 head provided multiple-data-track head gaps 0.431 mm (0.017 in.) wide and operated out of contact with the disk surface at all times. The data module, on the other hand, requires a single head gap less than 0.102 mm (less than 0.004 in.) wide and starts and stops in contact with the disk.

The pressure profile of the 2305 head, using a single taper-flat air bearing, is shown in Fig. 4(a). Because there are no regions of reduced pressure under the slider body and because the air bearing is the entire surface under the slider, a relatively large loading force (>1.2 kg) is required to keep the 2305 slider close to the disk.

- *Tri-rail slider*

A design concept was developed that permitted the adaptation of the single-taper ferrite slider to the low-load head. A taper-flat air bearing of relatively high length-to-width ratio allowed leakage of some pressurized disk boundary air from underneath the slider bearing at its midpoint. A tri-rail ferrite slider was designed with two outer rails approximately 0.635 mm wide and 6.35 mm. long, with a center rail of a width equal to the desired transducer gap width. The pressure profile resulting from this design is illustrated in Fig. 4(b). The four pressure peaks provide a head slider that is very stable in pitch and roll, and requires a small load force. The effective center of slider pitch and roll rotation is near the recording gap, thus providing a stable gap-to-disk flying height.

The ferrite slider forms part of the magnetic path for the read-write gap, as was the case in the 2305 head. The rear half of the transducer is a C-shaped ferrite core section glass-bonded to the slider in alignment with the center rail. The side walls of the slider rails were beveled to provide strength during head-finishing operations and

Table 2 Read-write head comparison.

Head model	*Spacing* (μin.)	*Load* (g)	*Suspended mass* (g)
2314	2.29	350	3.25
3330	1.27	350	5.0
2305	1.27	1200	2.5
3348	0.51	10	0.25

Table 3 3348 Read-write head performance.

Nominal gap flying height	0.51 μm
Bit density at inside track	2216 bits/cm (5636 bpi)
Track density	118 tracks/cm
Data rate	885 bytes/s
Write current, peak to peak	75 mA
Read signal, peak to peak	1.5 mV

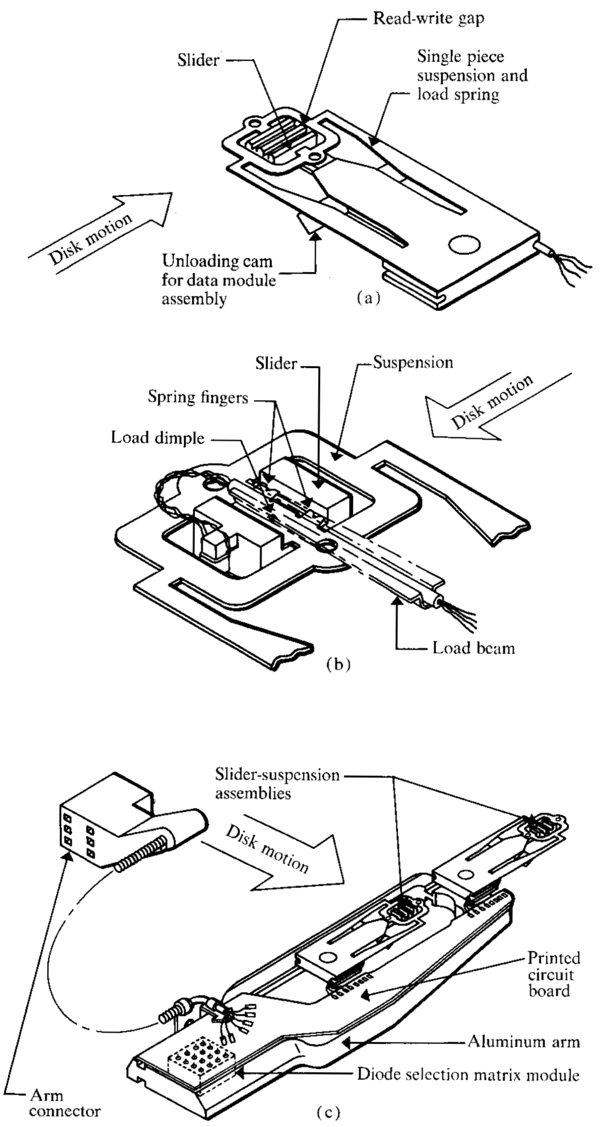

Figure 7 Suspension assembly. (a) Suspension and load spring; (b) snap-in attachment held with spring fingers, showing dimple for load point; (c) read-write assembly arm showing two-head mount.

to avoid edge chipping. The completed slider with glass-bonded core section and core winding is illustrated in Fig. 5.

• *Slider advantages*

Comparison with previous designs shows the following advantages:

1. The design is process oriented and suited for batch fabrication at low cost.
2. Head gap-disk compliance is enhanced by small size, relatively high pitch and roll stiffness and by placement of the rotation axes near the recording gap.
3. The air bearing requires a force of only 10 g, facilitating start-stop in contact with disk.
4. Limiting the area of exposed glass to the recording gap improves wear resistance.
5. A tapered center rail cross-section provides a high-efficiency head with a minimum-reluctance back gap.
6. Dimensional stability of the head is improved by use of a two-material system instead of the more conventional four materials—two bonding glasses, a ferrite, and barium titanate.
7. Flexibility in track width for alternate track densities is available with a single basic slider-core design, in which head-gap width is controlled by the final bevel grinding of the center rail.

The most significant design difference between the 3348 and its two predecessor start-stop-in-contact heads is in suspended mass and load force. Its characteristics, compared with other announced read-write heads, are shown in Table 2. The performance of the 3348 head as used in the data module is shown in Table 3. A typical "triplet" readback signal is given in Fig. 6, which shows an isolated group of three consecutive read voltage pulses used to measure resolution. The base-to-peak height of the two side pulses is typically 75 percent of that of the center pulse.

• *Suspension and load spring*

The suspension and load spring shown in Fig. 7(a) are made from a single piece of nonmagnetic stainless steel sheet. A load dimple is stamped into the symmetric suspension to provide a well-controlled load point relative to the slider. A self-aligning, snap-in suspension attachment to the slider at the center of mass is accomplished by means of spring fingers that are an integral part of the suspension piece [Fig. 7(b)]. A symmetric, stainless-steel-formed load beam is welded to the load spring to transmit the load to the dimple. In the operational position, the suspension and load spring are parallel to the disk to provide high stiffness in the radial-accessing, tangential, and yaw directions. The suspension assembly is welded to a stainless steel mounting block in such a way as to mechanically isolate the suspension from the load spring.

- *Head-arm assembly*

The mounting block is attached to an aluminum arm by swaging a tubular extension of the block to a mounting hole in the arm. The arm is made of aluminum to provide a match with the coefficient of thermal expansion of the disk. Because of the relatively small size of the head and the elimination of the need for head radial or skew adjustments, two or four heads can be mounted on one data arm, two on each side. A two-head arm is illustrated in Fig. 7(c).

To reduce the number of leads going from the accessing arm to the stationary data module connector, a diode selection matrix is mounted on the arm. All interconnections are provided by a conventional printed circuit board mounted on one side of the arm.

The head-arm assembly is protected from corrosion by means of stainless steel for the suspension load beam, an alodyne coating on the aluminum arm, an organic coating on all copper surfaces, and by the development of fluxless coil-wire terminations.

The servo head is mounted on a special single head-arm assembly similar to the data arm.

3340 dual drive

A dual-drive configuration was selected as a reasonable compromise between modularity and cost. Simplification of the drive electronics through the use of newer high-density circuit families and functional packaging provided reliability improvements. These gains in reliability made it practical to share components including power supplies, logic circuits, servo control, and read-write electronics while meeting the goal of improved drive availability.

Industrial design and human factors studies were made to explore alternative dual-drive configurations. Vertical arrangements with one drive placed above the other use minimum floor space. A disadvantage of this configuration, however, is that access for service and disk loading requires a moving-drawer assembly. Horizontal side-by-side dual drives offer the simplicity of top loading without the need for drawer access but increase required floor space.

Several data-module drive orientations were considered with the dual modules' spindles in horizontal and vertical orientations. Operator handling tests were performed with design mockups. A top-loading, side-by-side dual-drive arrangement was chosen because it offered optimum operator convenience, minimum cost, and maximum accessibility for service. Floor space requirements were minimized by limiting the dual-frame width to 107 cm (42 in.) for the 3340 as compared to the 102-cm (40-in.) frame width of the vertically packaged dual-drive 3330. The volumetric efficiency of the 3340 was due to the small size of the data module, the reduced mass and size of the accessing components, and the use of higher density circuit packaging available in the IBM MST-I and MST-E circuit families.

- *Electronics packaging*

The dual drive contains a controller that can attach up to four dual drives to a CPU. Each dual drive contains independent dc power supplies. Custom-integrated circuits were designed for each drive function to allow functional electronics packaging. This approach has minimized the number of electronic components, including circuit modules, cards, connectors and boards, with a significant gain in machine reliability and availability and a reduction in cost. Functional packaging also facilitated early circuit simulation and a shorter time for system debugging during development.

- *Service considerations*

Service time and cost considerations played a role in a number of design decisions. Not only can costs be substantial, but downtime can also affect system availability because of disk-resident programs.

One service objective established was to minimize electrical adjustments in the field. The 3340 requires only one adjustment, which controls servo system velocity and is made with the aid of a microcode diagnostic program and requires no special tools. A "no adjustment" approach for this function had been judged to be more costly because it would require a larger single replaceable electronics card for the velocity function and reduced circuit tolerances.

A second service goal was to avoid the need for field replacement of the air filters that prevent air system contamination. A closed-loop air system was adopted, with a single contamination control and cooling air system for the dual drive. The air filter is 99.97 percent efficient for 0.3-μm particles. Eight percent makeup air helps to maintain a positive pressure in the data module. The closed loop system cools data modules and linear actuators. The design goal of a maximum 40°F temperature rise over ambient temperatures in the air system was achieved by utilizing an air-heat exchanger for the closed-loop system. The heat exchanger is cooled by diverting a part of the open-loop air used to cool the drive electronic circuits.

Moving-coil linear actuator

Moving-coil linear actuators offer the characteristics required for the high-performance 3340 subsystem and have been successfully applied in the IBM 2310 and 3330 disk facilities. The low-mass, moving-coil type actuator is ideally suited to interface with an electronic servo control system for track access and track-following operation [8]. The 3340 servo control system is discussed by R. K. Oswald on page 506 of this issue [9].

• *Actuator functional requirements*

For the performance range desired for the 3340, an objective equal to or less than 25 ms had been established for the average access time of the actuator.

The data module configuration and the newly designed start-stop head permitted some relaxation of the performance demands on the linear actuator as compared with the 3330. Elimination of the need to move the heads off the disk surface reduced the required actuator stroke. The decision to use two data heads per data surface also reduced the required actuator stroke and acceleration, which resulted in the definition of the initial functional requirements as follows:

1. 2.54-cm stroke for data track access.
2. 0.51-cm stroke to move heads from data zone to landing zone.
3. Carriage and head-arm system mass to be 0.22 kg.
4. Total moving mass with moving coil to be less than 2 kg.
5. High radial stiffness of actuator-carriage-head-arm system to facilitate use of high-bandwidth servo control system.
6. Actuator crash stop dynamics to permit minimum crash stop deceleration stroke (<0.445 cm), with impact velocity of 228 cm/s.
7. Low flux leakage from moving-coil actuator to disk area to avoid need for magnetic recording head shielding.
8. Moving-coil assembly to couple and uncouple from data module carriage to allow removability.
9. Moving-coil assembly to be supported and aligned by data module carriage during operation.

The actuator requirements introduced one new area of significant concern: the incorporation of the coupling function in the moving-carriage actuator system. Initially, this requirement was viewed as being possibly inconsistent with the need for high dynamic stiffness.

• *Actuator design decisions*

Several actuator design decisions were made early in the 3340 development program prior to the design of the disk drive:

1. Adopt a short-moving-coil, long-magnetic-gap actuator configuration.
2. Support moving-coil inside diameter with nonmagnetic metallic strut structure and slot actuator center pole to accommodate struts.
3. Support moving coil on nonmagnetic, metallic former.
4. Support bidirectional crash stop in actuator; interact with coil support structure along central axis of actuator.
5. Couple moving-coil structure to carriage with spring preload adequate to resist normal access and crash-stop accelerations.
6. Place magnetic shield in front of actuator magnetic gap with central aperture (to accommodate moving-coil support structure) smaller than gap inside diameter.
7. Contain moving-coil flexible leads behind front face of actuator gap.

• *Discussion of actuator design decisions*

Short-coil, long-gap configuration The typical moving-coil linear actuator utilized in disk facilities has a magnetic structure with a high flux density of 8 000-10 000 gauss, a gap of short axial length (0.8-1.0 T) and a moving coil of long axial length. The coil axial length is long enough so that a portion of coil length equal to the gap axial length always lies within the gap during access motions. The gap magnetic flux, typically provided by permanent magnets, is thus efficiently utilized. The current applied to the long coil and the coil mass is not utilized at peak efficiency because, in typical disk drive actuators using this structure, half or more of the coil length is outside of the gap. The inherent mechanical properties of a long coil may also introduce mechanical dynamic resonances into the carriage system.

The decision to use a short-stroke file configuration with two data heads per disk surface made a short-coil, long-gap actuator configuration more attractive. A short-coil design was expected to improve system dynamic stiffness. A short coil could be made more rigid than a long coil of similar diameter. Reduced coil mass and increased coil rigidity could be achieved while reducing the overhung mass to be supported by the carriage and coupling. The short-coil design permitted the front magnetic shield to be placed between the actuator air gap and the data module with approximately 1.27 cm less overhang distance required from the carriage coupling to the axial center of gravity of the moving coil. Although the gap flux would not be utilized efficiently, the coil current would be used efficiently. This configuration offers reduced power dissipation in the coil, the option of using a lower voltage, a less costly actuator power supply, and reduced cooling requirements. Reduced cooling requirements, in turn, provide reduced air system noise and longer air-filter life.

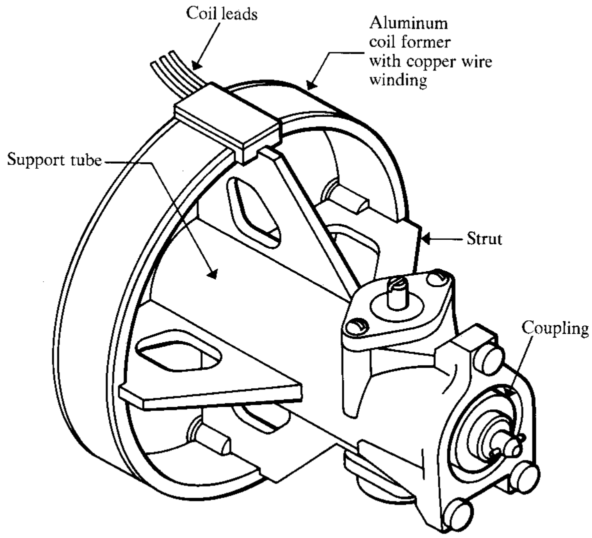

Figure 8 Moving-coil actuator assembly, showing support tube, coil former, and struts.

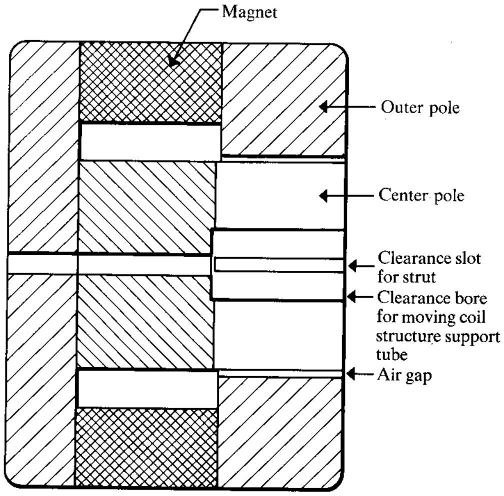

Figure 9 Initial cylindrical design of moving-coil linear actuator.

Adoption of the short coil was thus an engineering judgment made primarily in consideration of dynamic performance and secondarily in consideration of power supply cost, cooling air system cost, and acoustic noise level.

Coil support structure and former A moving-coil actuator with an internal aluminum-strut support structure had been built by the IBM Boulder Development Laboratory. Carter [10] has disclosed a moving-coil assembly with an internal support structure made with a dielectric material. With this type of assembly, the moving coil may have a more rigid construction than a coil of similar mass supported only from one end. An aluminum structure was selected to provide an optimum combination of stiffness, weight and manufacturing cost. The structure was designed with four support struts attached to a central tube. The central tube provided a convenient structure to support the actuator coupling and to interact with the actuator crash stop. The tube and strut structure presented the disadvantage of reduced center pole area; the pole diameter (and moving coil) was increased in diameter to compensate for this area loss.

An aluminum coil former was selected in order to provide maximum rigidity with minimum actuator gap length, the latter being defined as the radial distance between the outer diameter of the center pole and the inner diameter of the outer pole. Eddy current damping effects from the moving coil former were avoided by splitting the former axially at one point on its circumference to avoid creating a shorted turn. Minor eddy current loops between the support struts and the former were avoided by using an insulating anodized finish on the former. The support tube and struts were implemented as a single aluminum sand casting on the first actuator. The production assembly is an aluminum pressure-die casting and is illustrated in Fig. 8.

Actuator bidirectional crash stop A moving-coil actuator requires a crash stop at each end of its stroke to accommodate incidents in which servo control is lost. This stop must decelerate the moving-carriage coil mass in a short distance to minimize the loss of usable disk recording area, because the heads must remain on the disk throughout the crash-stop stroke. The crash-stop should avoid introducing mechanical moments into the carriage system that would cause high vertical accelerations and possible head-disk contact or carriage-bearing damage. A bidirectional crash stop concentrically located in the actuator minimized crash-stop moments and simplified the data-module design. The crash-stop design is discussed in the section entitled "Actuator crash stop."

Coupling preload The actuator coupling was designed to allow the data module carriage to support and align the moving coil concentrically within the actuator air gap. A spring-loaded coupling similar in function to a commercial quarter-turn fastener was adopted. A 36-kg (80-lb) preload was initially selected to prevent separation of the moving-coil structure and carriage during normal access operation and crash-stop engagement. A spring-biased coupling with adequate preload was expected to provide the system rigidity required by the servo system.

Front shield The motor front shield was designed for placement between the actuator magnetic gap and the data module. The concentrically strutted moving-coil structure allowed the use of a single aperture in the front shield smaller than the inside diameter of the actuator gap. This design provided the shielding required to minimize the actuator flux density measured in the data module.

Moving-coil leads To achieve the design goal of minimum moving-coil overhang from the data module carriage, the actuator was placed as close to the data module as permitted by shielding considerations and the data module envelope. To minimize the actuator length, the flexible leads to the moving coil were extended rearward from the coil in a loop contained within the magnetic structure of the actuator. Two opposed slots were placed in the gap surface of the outer pole piece: one to accommodate the rolling flexure-band coil leads, and the other opposed at 180° to maintain the flux symmetry in the gap.

The design features discussed here were successfully implemented in the first development model of the drive and are incorporated in the present version of the 3340. The spring-biased coupling proved to provide a sufficiently rigid connection between the moving-coil assembly and the carriage. The metal moving-coil support structure gave the desired system stiffness.

Optimum designs were derived for the coupling preload and the crash-stop dynamics following an investigation of potential abnormal actuator operations, discussed in the section entitled "Actuator crash stop."

• *Actuator cost optimization*

Coil A simulation program was written to permit the optimization of coil design parameters, i.e., number of turns, wire size, length and number of layers of wire, current, and gap flux density. A decision was made to limit coil voltage to 36 V to minimize power supply cost, and further optimization proceeded under this constraint.

Magnetic configuration The first development model of the actuator used a cylindrical design with two external ring magnet segments, as illustrated in Fig. 9. The cylindrical actuator required complete external shielding and was difficult to align to the drive base plate and the attached data module. An investigation of alternate magnetic configurations was then undertaken.

The selected magnetic structure is illustrated in Figs. 10(a) and 10(b). This design utilizes a concentric magnet positioned behind the actuator air gap and abutting the center pole piece. The magnet is mounted on a rear iron plate that in turn is mounted on a cast-iron housing, the magnetic return path. The housing, which also serves as a shield to external flux leakage, provides mounting reference surfaces to support the front shield and to align the housing assembly to the drive base plate. This is the only shielding required other than the shield plate mounted in front of the actuator magnetic gap. The front shield plate supports one of the data-module reference surfaces (the locating ball) and helps to minimize tolerance buildup between the ball location and the actuator air gap. Precision alignment of the data module to the actuator air gap is thus enhanced, reducing the clearance required in the air gap to accommodate misalignment of the moving coil assembly.

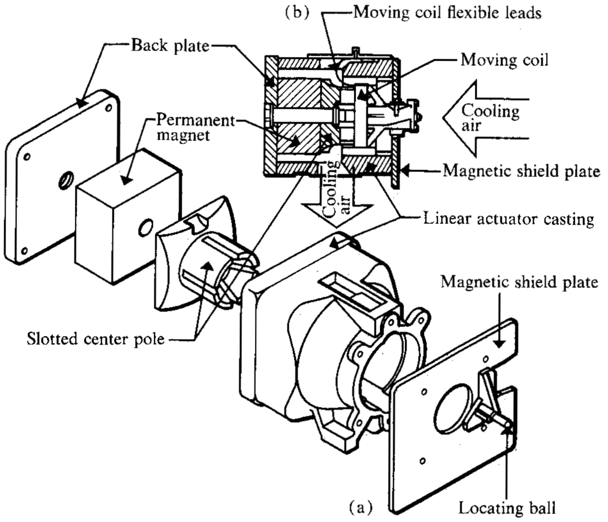

Figure 10 Final design of magnetic structure for actuator. (a) Concentric magnetic assembly located behind air gap, exploded view showing center core, permanent magnet, and housing; (b) view of actuator showing magnetic structure, moving coil, coil leads, etc.

• *Actuator performance*

Coil inductance and performance The development actuator did not initially meet its performance goal because of high moving-coil inductance. This section discusses the significance of coil inductance in the 3340 short-coil, long-gap design and describes the design of an outer-pole shorted turn to minimize effective inductance and improve actuator linearity.

Moving-coil inductance increases the electrical time constant of the moving coil and increases the disk access time. If the moving-coil inductance is very high, short track moves may be impossible because of the inability to reverse the coil current rapidly.

For a given number of coil turns, and a given mean coil diameter, the inductance of the coil varies inversely with the length of the coil. A coil of short axial length has, of course, more layers of coil and shorter length. The inclusion of the entire length of the short coil within

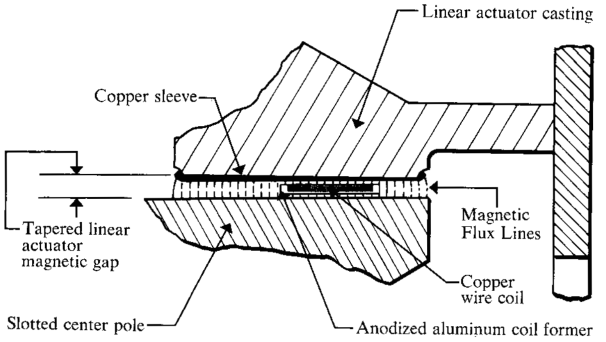

Figure 11 Actuator air gap showing tapered outer pole diameter and copper sleeve that acts as shorted turn.

the gap structure increases the relative coupling of the short-coil field with the magnetic circuit components of the actuator, and this tends to increase coil inductance.

Coil-induced eddy currents in the relatively high resistance surface of the magnetic gap tends to diminish the effective coil inductance. The 3340 slotted center pole and the coil-lead clearance slots in the outer pole significantly increased the surface resistance for eddy currents. The inductance of the 3340 moving coil in air is approximately two mH. The inductance of the coil in the actuator as first built was 20 mH, whereas the design goal was seven mH.

Figure 13 Interface between disk drive and data module. (a) Drive assembly; (b) interface at air shroud area.

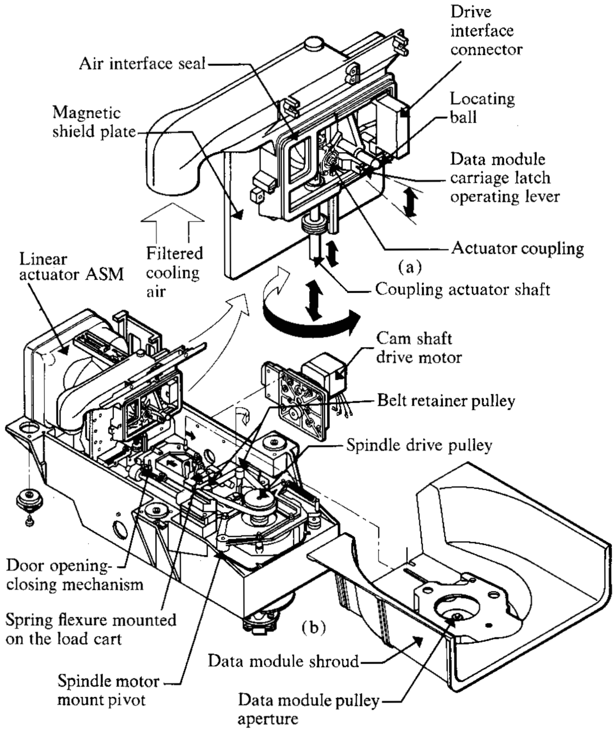

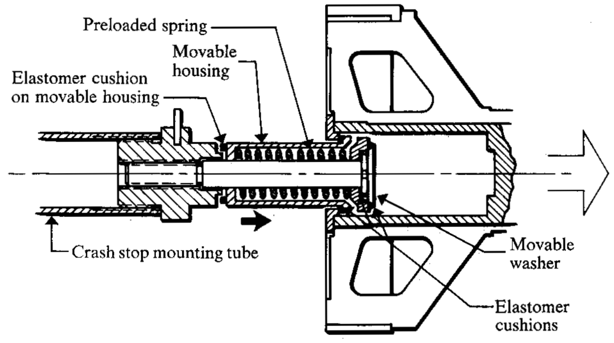

Figure 12 Action of crash stop designed for the IBM 3340, showing crash in direction of spindle with moving housing displaced from normal "home" position.

Reducing coil inductance A conventional approach to reduce effective coil inductance is to place a low-resistance shorted turn close to the actuator moving coil—for example, on the center core. The slotted center pole of the 3340 made a center-core shorted turn difficult to implement and less efficient than in a nonslotted core. An outer-pole shorted turn was designed to fit the irregular shape of the outer-pole gap surface. A copper insert of 0.38-mm wall thickness reduced the effective coil inductance of the first motor to 11 mH, thus permitting satisfactory operation (25-ms average access time) of the first development model of the actuator.

Linearizing actuator moving coil inductance and gap flux An ideal actuator design has constant moving-coil inductance and constant gap flux density through the operating stroke of the actuator. The magnetic structure of the 3340 actuator is axially asymmetric; i.e., the permanent magnet and return path structure are placed at one end of the actuator gap. This configuration causes a nonlinearity in flux density and inductance as the coil is moved axially through the gap. Flux density tends to be higher at the rear of the gap nearest the permanent magnet. The coil field couples more effectively with the actuator components at this position, thus increasing coil inductance.

Inductance and flux density linearization were accomplished by tapering the outer pole prior to the installation of the copper-sleeve shorted turn. The outer-pole diameter closest to the magnet is increased in diameter, thus lengthening the magnetic gap at the magnet end of the actuator gap. The increase in gap length decreases the flux density for this portion of the stroke so that the tapered diameter compensates for the asymmetric shape of the magnetic structure. The final actuator uses a cylindrical copper sleeve that is forced into the shape of the outer pole by the magnetic-pulse metal forming process [11]. The inside surface of the copper sleeve is finished to a constant diameter. The thicker copper section near the magnet end of the air gap compensates for the greater coupling of the field of the coil with the magnetic structure. By providing a magnetic structure with a tapered bore and a copper sleeve with a constant bore (Fig. 11), we were able to linearize both flux density and inductance along the actuator stroke.

Actuator crash stop

In a conventional drive, a runaway actuator moving away from the spindle axis may be accommodated by an operation that unloads heads from the disk surface and withdraws them from the disk pack. A runaway motion in the opposite direction requires that the heads remain over the disks at the inner radii. The data-module design required that both crash-stop motions be accommodated with the heads remaining on the disk; thus, it was important to limit crash-stop motion to a minimum area of disk recording surface.

The main design goal of the crash stop was to prevent head and disk damage. A stopping distance of <0.427 cm was selected. The maximum actuator velocity at impact was estimated to be 228 cm/s. The crash stop was to decelerate the carriage and moving coil with forces less than 250 *g* as measured in the plane of the data arms parallel to the access axis. This acceleration level was used as a design requirement for the head and arm assembly, the arm and carriage attachment, and the actuator coupling. Because the data-module carriage must be latched to a fixed position during operator handling, the crash stop was designed to provide a reference surface to locate the data-module carriage for carriage-latch engagement.

The maximum runaway velocity of a moving coil actuator system is determined by actuator force constant, power supply voltage, available stroke length, coil resistance, and moving mass. Back-emf generated by the coil motion (proportional to the force constant) will limit both runaway velocity and rebound velocity and will assist in bringing the actuator rapidly to rest after crash-stop engagement. The theoretical back-emf-limited velocity of the 3340 actuator is approximately 3.05 m/s; however, the relatively short stroke limits the maximum velocity to less than 2.29 m/s.

- *Design description*

The 3340 crash stop is illustrated in Fig. 12 and incorporates three moving parts: a preloaded spring, a movable washer with two elastomer cushions, and a movable housing with two elastomer cushions. The moving housing provided the means of locating the two crash-stop engagement surfaces in close proximity to one another, thus limiting the moving-coil support structure length required to accommodate the crash stop.

- *Crash-stop performance*

Tests with the first design of the crash stop developed disk damage at crash-stop impact velocities of 2.06 cm/s (85 in/s). The 35.7-kg (80-lb) coupling preload proved inadequate, allowing carriage and actuator to separate during crash-stop engagement in the direction toward the spindle. Access axis loads over 300 *g* and data arm accelerations of 3500 *g* normal to the plane of the disk were determined at the moment the coupling alignment pins recontacted the carriage pins. The coupling preload was increased by about 10 percent and the velocity at which separation occurred was increased to above normal runaway velocity.

Crash-stop cycles in the spindle direction were found to cause higher loads (250 *g*) than crash-stop cycles in the opposite direction (150 *g*). The fact that the mass of the moving housing was higher than the mass of the moving washer was assumed to be the source of the higher acceleration loads. The housing mass was reduced from 35 g to five g by a change to aluminum and a modification of the housing wall. A second elastomer cushion was added to the rear of the moving housing to decelerate the housing at the end of the crash-stop cycle. The final crash-stop design successfully limited radial accelerations to less than 150 *g*.

Interface between 3340 drive and data module

- *Loading motion*

The drive interface [Figs. 13(a) and 13(b)] accepts the data module as loaded by the machine operator. All of the load-unload operations are controlled by the operation of a cam shaft driven by an electric motor. After closure of the drive top cover, the drive moves the data module forward approximately five cm to complete the loading and alignment operation. The drive and data module are ready for normal operation in less than 20 s. During the motion of the data module, the data-module door is opened, and the data module is moved into contact with the actuator-mounted locating ball, the drive interface connector, and the actuator coupling.

The data module pulley extends through an aperture in the moving data module shroud and inside the loop of the spindle drive belt. The drive belt is retained in place by two movable belt-retaining pulleys. During data module movement, the pulley engages the belt and the retaining pulleys are moved out of contact with the belt. Near the end of the load motion, the data-module carriage latch engages an operating lever mounted at the side of the locating ball, the data-module cover is sealed to an air interface, and filtered air is introduced into the module. A spring flexure on the cam driven data-module load cart provides a 64-kg bias force to the load pin of the module at the end of the load motion. This force, acting along the

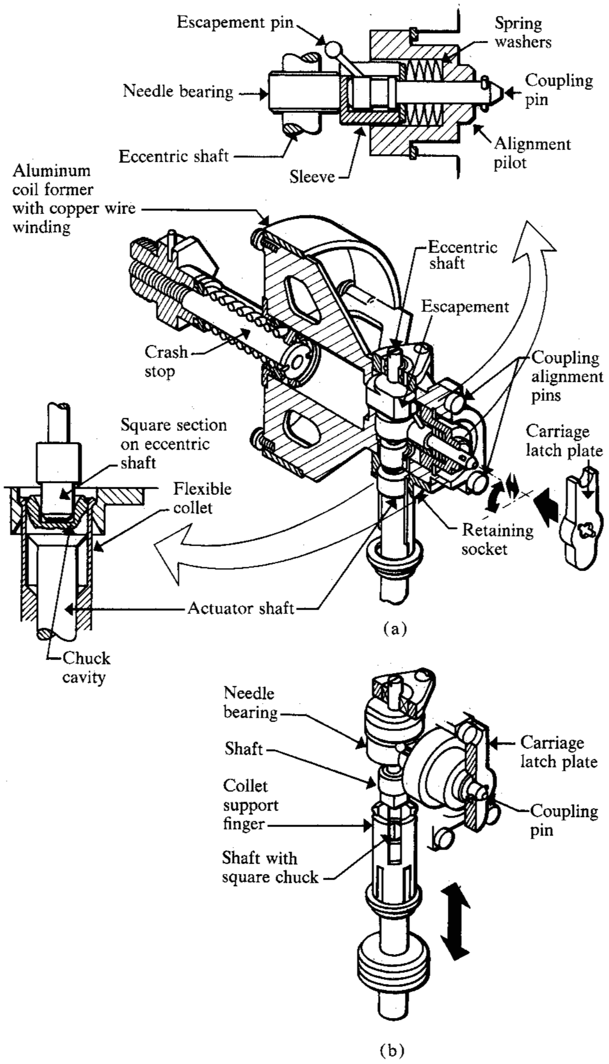

Figure 14 Moving-coil coupling assembly. (a) View showing how "fingers" engage socket and square chuck engages eccentric shaft; (b) simplified view with carriage latch plate coupled and support fingers and square chuck withdrawn.

central axis of the data module in a horizontal plane below the data module registration cone, provides vertical and horizontal moments to register the two alignment pads on the data module.

- *Actuator coupling and carriage latch*

As noted in the section on the actuator, the coupling acts like a quarter-turn fastener with a preloaded spring. A challenge in designing the coupling was to devise a means of operating the coupling with minimum input force or moments. The coupling engages the carriage latch plate and the carriage, in turn, is retained in its way system by a <1.8-kg preload on one outrigger bearing way. High coupling operating forces or moments could move the carriage out of contact with part of the way system or misalign the moving-coil structure during coupling operation so as to prevent engagement of the coupling pin and latch plate.

An eccentric shaft and escapement mounted on the moving-coil structure provided a means to limit input operating force to 0.565 N·m (about five in-lb.). The coupling is illustrated in Fig. 14. The needle bearing on the eccentric shaft releases the preload from the coupling pin and moves the pin forward towards the carriage. At its forward position, the pin may be rotated freely (out of contact with the latch plate) by the action of the escapement.

The coupling is operated by an actuator shaft that is retracted vertically downward at completion of the coupling operation [Fig. 14(b)]. To uncouple the actuator, the data module carriage is put at its home position (in engagement with the spindle-direction crash stop), and the carriage is latched. The coupling actuator shaft is moved upward to engage the eccentric shaft. The motion of the actuator shaft carries a flexible plastic collet chuck upward to engage a retaining socket at the bottom of the moving-coil structure. The collet supports and positions the moving-coil structure when the data module is removed from the drive.

Data integrity

Disk-data integrity has become increasingly important with the advent of on-line data base systems and virtual storage. The single most troublesome problem to overcome, with respect to data integrity, has always been imperfections in the recording medium. Tiny defects in the magnetic coating have become fewer and smaller as disk coating processes have been considerably improved over the past decade, but the recording densities have increased so much that each succeeding disk product has to be carefully designed to maintain an ever-higher data integrity in the presence of occasional disk defects.

- *Data integrity in the IBM 3330*

The data integrity strategy used with the 3330 was to provide an error correction redundancy code (ECC) capable of correcting read data errors of up to 11 bits in length. This ECC also, of course, had the important task of detecting a read data error whether or not it was correctable. With this philosophy of handling disk defects, data are written, ignoring the possibility that a small defect will affect the reading of the data. The assumption is that the ECC redundancy bytes written will provide a correction capability. If a defect is longer than 11 bits or if multiple defects exist on a track, the disk is rejected at the manufacturing stage.

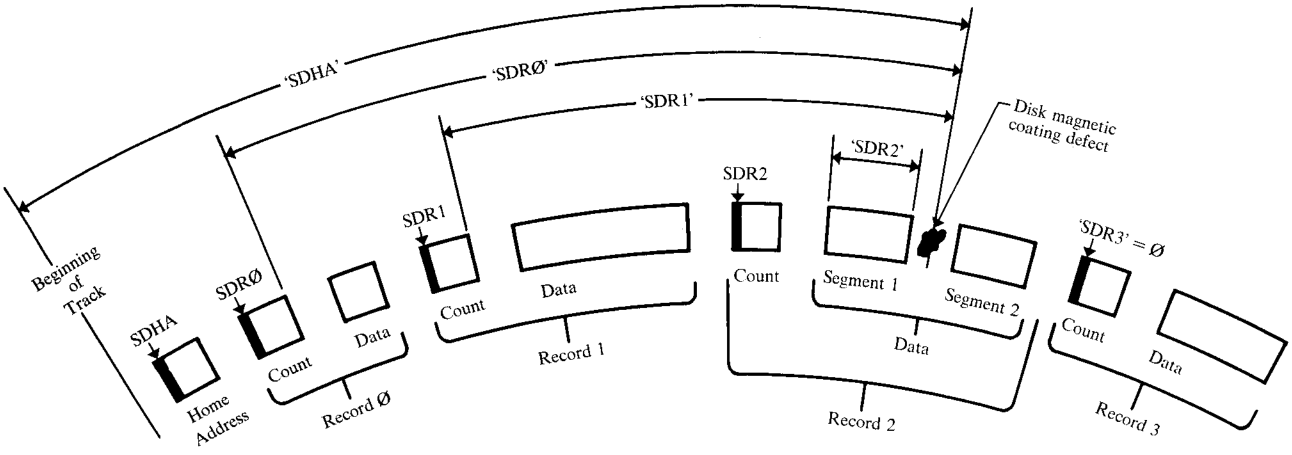

Figure 15 Example of defect skipping. The defect is located SDHA bytes from the beginning of the track. This value is written in the Home Address area during disk manufacture. The value written in the Record 0 Count is the length SDR0. Record 2 is affected by the defect and for ease of later read or update write operations, the value written in SDR2 is the length of segment 1.

- *3340 defect-skipping data integrity design*

For the 3340 a different philosophy of overcoming the disk defect problem was implemented. The procedure uses information from track-by-track tests made during disk manufacture to find possible defects. Thus, during the 3340 development effort, the hardware and microprogram built to attach and control the 3340 were designed so that any single defect on a track is skipped automatically during the read operation. Where there is a defect, its displacement from the beginning of the track is recorded in the home address area at the start of the track. Subsequently, when the customer's records are formatted on the track, the count field of each record is written, including an updated displacement to the defect from the end of that count field. This is shown in Fig. 15.

In formatting the record on the track that would include the defect (if formatted normally), that record is written in a fashion that ensures that a data gap, instead of customer or control data, is written where the defect is located. During subsequent read operations, no attempt is made to read data from this gap section and thus the defect is skipped.

- *Defect-skipping procedure*

Two methods are used to place this gap at the point of defect. The first is to split a normally contiguous field of the record into two segments. Between these two segments is the data gap at the defect area. The second method is to extend a normal gap that separates the fields within a record, to include the defective area. The latter method is used when the defect is near the beginning of a field or when it is located in a normal gap. To prevent track-capacity formulas from being affected by a track containing a defect, the amount of space necessary to split a field or to extend a gap is allocated for every track. This allocation is 128 bytes for each track on a 3348 disk. A maximum size of 16 bytes is allowed for the defect, which, of course, is much larger than the defect areas normally encountered during manufacture. If more than one such defect area, more than 16 bytes apart, is detected in any data track during manufacture, the disk is rejected.

The entire defect skipping procedure is controlled by the attachment microprogram and is entirely independent of the software, i.e., both the system control program and user programs.

- *Performance characteristics*

When a record is formatted (including both count and data) on a track that has a defect, and the record is the one on the track that is to be adjusted to skip the defect, an extra revolution is required to complete the operation. This is so because it cannot be determined whether this record is affected by the defect, because the data length is not known until the count field is written. The extra revolution is used to rewrite the count field with the proper information to skip the defect. No system interrupts occur.

Because the vast majority of tracks have no defects, the extra revolution required when formatting such a record occurs rarely. Once the record is formatted, read or update write operations proceed normally with no performance loss. On the other hand, using conventional error correction to recover from a small defect-caused data error requires software error recovery procedures and extra time.

• *3340 error correction code*

The 3340 recording format also includes an error correcting code, the main purpose of which is to guarantee read data validity; it is equivalent to the 3330 in this respect. When errors are detected with the ECC redundancy, the cause is electrical noise or some other transient, because the physical defects have been eliminated by defect skipping. A simple repeat of the read operation is normally successful. The ECC redundancy is capable of correcting read errors up to three bits long in the data field of a customer's record and is used to correct such errors even if they may be temporary. This approach eliminates the time interval necessary to repeat the entire operation.

Summary

The new technologies utilized in the 3340 Direct Access Storage Facility have enhanced disk file performance and flexibility. Table 4 shows the 3340 characteristics in an evolutionary comparison to its predecessors.

The 3340 and 3348 engineering design is based upon a new magnetic-recording head air-bearing technology originally developed for nonremovable disk applications. A new tri-rail, taper-flat air bearing of large length-to-width ratio was developed for the 3348 head. The tri-rail configuration facilitates the integration of the head element and slider in a two-piece ferrite assembly to simplify construction and to provide improved wear resistance. The low-load force required by the head permits head-disk contact during disk start and stop; the small size, small mass, and air-bearing design provide improved head-disk compliance.

The data module, by incorporating disks, heads, and those mechanical elements that control data head-to-data track alignment, facilitates the use of a high track density and eliminates the head-alignment adjustment costs associated with conventional removable disk pack drives. A family of data modules is usable on a single basic drive.

An average access time of 25 ms and a track density of 118 tracks/cm (300/in.) were achieved by the use of a newly designed moving-coil actuator and a track-following servo system. Actuator design features include the use of a die-cast aluminum, strutted-moving-coil support structure, a cast-iron flux return path and shield housing, and an actuator-mounted, bidirectional crash stop.

The drive interface requirements are more complex than those for a conventional drive because of the number of functions partitioned between drive and data module. A number of new drive interface mechanisms were developed. These mechanisms include the actuator coupling, the coupling actuation and moving-coil support unit, the data-module load cart, and the drive-belt engagement and retention components. The cost of these mechanisms was minimized by the use of precision die castings and molded plastic components. However, the cost of the interface functions and the basic data-module structure appears to limit the cost effectiveness of the removable data-module configuration to capacities greater than approximately 10 Mbytes.

Data integrity in the 3340 is enhanced by the use of disk-defect skipping. This technique utilizes defect location information developed during disk manufacturing. Customer data are formatted to skip disk defect areas. This approach minimizes the disk recording area that is lost to defects, simplifies error correction requirements, and reduces system time required to deal with detected data errors.

Table 4 IBM disk drive performance and capacity.

	Model					
	305	*1311*	*2311*	*2314/2319*	*3330*	*3340*
Year announced	1956	1962	1964	1965	1970	1973
Bits/cm	39	400	434	865	1590	2200
Bits/in.	100	1020	1110	2200	4040	5600
Tracks/cm	8	39	39	39	75.6	118
Tracks/in.	20	50	100	100	192	300
Areal density (kbits/sq. cm)	0.3	17.4	17.4	34.2	120	261
(kbits/sq. in.)	2	51	111	220	775	1680
Drive spindle capacity (Mbytes)	5	2.0	7.25	29.17	100	35/70
Average access time (ms)	600	150	75	60	30	25
Rotational velocity (rpm)	1200	1500	2400	2400	3600	3000
Data rate (kbytes/s)	10	78	156	312	806	885

Acknowledgments

The author gratefully acknowledges the assistance of several individuals who contributed to the preparation of this paper. W. E. Cole prepared the section on data integrity. T. R. Patel was responsible for the engineering design of the linear actuator and the closed-loop air system, and assisted the author in numerous useful discussions. M. W. Warner proposed and developed the 3348 air-bearing slider. Mr. Warner and E. R. Solyst aided in useful discussions and in preparing technical data and illustrations for the head and arm sections. R. K. Turley performed the actuator crash-stop evaluation and design optimization and assisted in preparing the crash-stop section. R. H. Lee, who optimized the design of several drive interface components, coordinated the preparation of illustrations for this paper. S. H. Jacobs offered valuable suggestions on the organization of the paper.

A development program such as that for the 3340 and 3348 requires the conceptual and engineering contributions of many individuals. While some of these contributions have been described and acknowledged, this paper has highlighted only some selected new concepts, components, and contributors.

The important role of the 3735 head and the small file design of W. S. Buslik in encouraging the data-module concept have been noted. Discussions with Mr. Buslik during the early task force effort were particularly helpful in exploring alternate drive configurations. R. W. Lissner participated in the task force and the data module concept development. Mr. Lissner's contribution was recognized by a joint IBM Outstanding Invention Award in 1974. K. E. Haughton initiated the task force effort that developed the data module concept and played an active role in program definition. C. P. Coolures contributed to the program definition and the design considerations discussed in this paper.

A. D. Rizzi contributed to the crash-stop design, the coupling design, and the actuator magnetic configuration. C. P. Barnard also contributed to the coupling design and developed the data-module carriage-way design. The coupling actuator and moving-coil support was designed by W. L. DeJager. E. H. Empson and R. R. Borg contributed to the design of the linear actuator.

The defect skipping concept was first proposed by D. I. Frush. A. B. Johnson developed the microcode algorithms to implement defect skipping.

Cited 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 (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 (1957).
3. M. A. Kohn, "Air-Bearing Magnetic Transducer Assembly," U. S. Patent 3,349,384, October 24, 1967.
4. J. M. Hansen and Z. Atamian, "Self-Seating Contact Head for Magnetic Memory Drum," U. S. Patent 3,022,494, February 20, 1962.
5. W. S. Buslik, "Magnetic Disk Storage File in Sealed Enclosure," U. S. Patent 3,710,357, January 9, 1973.
6. C. W. McCormick, *NASTRAN User's Manual*, NASA SP-222(01), Scientific and Technical Information Office, National Aeronautics and Space Administration, Washington, D.C. 1972.
7. E. R. Solyst, "Multichannel Recording Head," U. S. Patent 3,579,214, May 18, 1971.
8. M. Hiyane, Y. Inoue, and M. Kurusu, "Development of Linear Motion Actuator," *Fujitsu Sci. Tech. J.* **8,** 59 (1972).
9. R. K. Oswald, "Design of a Disk File Head-Positioning Servo," *IBM J. Res. Develop.* **18,** 506 (1974), this issue.
10. T. A. Carter, "Linear Motor-Generator," U. S. Patent 3, 487,241, December 30, 1969.
11. L. L. Seyler, "Electromagnetic Metal Forming," *Tooling and Production* **37,** No. 6, 56 (1971).

Other references

12. R. W. Lissner and R. B. Mulvany, "Magnetic Disk Storage Apparatus," U. S. Patent 3,786,454, January 15, 1974.
13. K. E. Haughton, "Design Considerations in the IBM 3340 Disk File," *Digest of Papers, COMPCON 74*, IEEE, San Francisco, February 26-28, 1974, p. 281.
14. J. M. Harker and H. Chang, "Magnetic Disks for Bulk Storage—Past and Future," *AFIPS Conf. Proc. Spring Joint Computer Conf.* **40,** 945 (1972).

Received June 10, 1974

At the time this manuscript was prepared, the author was located at the IBM General Products Division Laboratory, Monterey and Cottle Roads, San Jose, California 95193. He is now at the IBM Data Processing Division at 525 University Avenue, Palo Alto, California 94301.