

Daniel Abramovitch
and Gene Franklin

A Brief History of Disk Drive Control

A persistently exciting control application is that of disk drive servos. From the first magnetic drives of the 1950s to the massive-capacity commodity drives of today, the problem of accessing data on rotating disk media has provided a wealth of control challenges. This situation shows no signs of abating as storage densities, capacity, and transfer rates keep rising while costs and size continue to drop (Fig. 1). Although a host of new technologies are poised to challenge the supremacy of hard disks in their primary purpose of providing vast storage at low cost, reports of the latter's demise are consistently and greatly exaggerated.

The purpose of this article is to provide a history of control in disk drives. Although this subject can include both flexible and optical drives, this article will focus on rigid magnetic disks—often called hard disks. While the intent is to focus on the control history, the applied nature of this problem makes this difficult to do in a vacuum. Thus, the structure of the article is as follows. We will start with a short general history of disk drives culled from several recent excellent sources [1]-[3]. From there, we will take a walk around the feedback loop and trace the evolution of the individual components. Finally, we will focus on the control system it-

self, viewed in the context of overall history and the component history. Tutorial material is presented primarily to illuminate the technical history. A more detailed tutorial on the hard disk control problem, as well as a contrast to the problem of controlling optical disks, is found in the *Proceedings of the 2001 American Control Conference* [4], [5].

A Brief History of Drives: From the RAMAC to Here

The history of hard disks is intertwined with the history of computing in an uneasy and stormy alliance. The concept of storing large amounts of data on magnetic media was already in practice in the

early 1950s with magnetic drum memories; however, the volumetric density was limited by the relatively low surface-to-volume ratio of such devices. The concept of storing data on magnetic disks was being discussed in 1953 in several locations [1], [2], but it was IBM's remote research laboratory in San Jose that brought the first disk to market in 1956. The random access

memory accounting machine (RAMAC) is famous as a mechanical marvel of the day; it moved a single pair of read/write heads vertically to access the desired disk and then radially to locate the desired track. The first RAMAC held a grand total of 5 million 7-b characters and was de-



©1994 PHOTODISC, INC.

Abramovitch (danny@labs.agilent.com) is with Agilent Laboratories, 3500 Deer Creek Road, M/S: 25U-9, Palo Alto, CA 94304, U.S.A. Franklin is with the Department of Electrical Engineering, Stanford University, Stanford, CA 94305, U.S.A. This article was originally solicited for the April 2002 special issue on the history of control by the issue's guest editors, L. Bushnell and D. Bernstein. Its publication was deferred to this issue because of page limit constraints.

livered to the offices of Crown Zellerbach, a paper company headquartered in San Francisco.

If one were to read only one or two accounts of the history of disk drives [1], [6], one might conclude that the “cowboys” at IBM San Jose were the only ones to have thought of the concept. In fact, work was being done at Logistics Research, Inc., of Redondo Beach, CA, and Engineering Research Associates, which was eventually folded into Remington Rand Univac [2]. As early as 1956, the Univac operation in St. Paul, MN, was ready to ship a disk drive, but this was delayed because the Univac operation in Philadelphia, which made 18-in magnetic drums, had more political clout. In a story that would be repeated over and over again in the disk drive business, companies would unknowingly mortgage their futures to protect established businesses [7]. The monumentally bad decision by Univac delayed their entry into a market that would change computing forever.

Parts of the story read like a badly written soap opera. From the earliest years of the industry, a company named Telex would attempt to make disk drives, exit the business, then reenter through some merger or acquisition. Influential managers would leave IBM for new disk drive ventures, recruiting large groups of engineers to follow. Notable among these were the departure of Vic Witt for Information Storage Systems (ISS). His successor at IBM, Al Shugart, would leave for Memorex, recruiting up to 200 IBM engineers. Shugart would later leave Memorex to start Shugart Associates, a floppy drive company. He would then get fired by his own venture capitalists, only to return five years later to form a company called Shugart Technology with Finis Conner. They would change the name of the company to Seagate Technology under threat of legal action from Xerox, who had bought Shugart Associates. In a few years, Xerox would shut down Shugart Associates, but Seagate continues today. In between, Finis Conner left to start Conner Peripherals, which was recently bought out by Seagate. Telex would eventually end up as part of Memorex’s disk drive operation before exiting the business for the last time [2].

Before disk drives, computers operated in a mode known as batch sequential; that is, data was read into the computer off a stack of punched cards or a reel of magnetic tape, processed, and then written back out onto punched cards or magnetic tape. The concept of having data continuously available did not exist. However, by the early 1950s, some companies and government agencies wanted something similar to an online database. The team at IBM considered arrays of magnetic drums but settled on a disk structure [1], [2]. With the arrival of the disk drive, interactive computing and continuously online data became a possibility.

Despite all this, mass storage in general and disk drives in particular have held a less than glamorous position in the area of computing. It seems that the concept of having an elegant electronic computation device dependent on moving pieces of metal for storage has always been at odds with the general sentiment of the computer industry. Thus, it is not

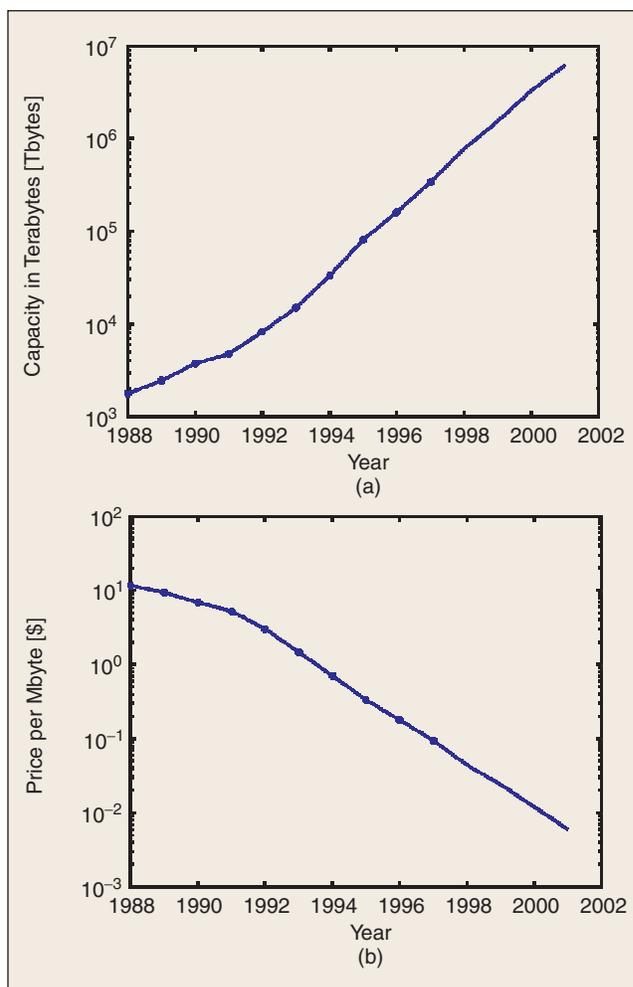


Figure 1. Drive industry total capacity shipped and cost per megabyte. Data taken from Disk Trend [3]. Based on actual data through 1997; estimates are used after that. (a) Total disk capacity shipped by year. (b) Overall price per megabyte by year.

surprising that many computer types have wanted to relegate storage to something done with solid-state memory or at some distant disk farm available over the network. Such an architecture would allow for entirely solid-state machines that would be lighter, faster, quieter, less fragile, and would consume less power. The world would be a more elegant place with diskless computing.

The failing element in this philosophy is that nothing has been able to match the hard disk’s combination of volumetric storage capacity, data availability, transfer rates, and nonvolatility. The death of disk drives has been predicted for a relatively long time, but as with Mark Twain, reports of their demise are greatly exaggerated. Every increase in capacity, storage density, speed, and reliability has opened up new opportunities for computation. The advent of 8-in disk drives made minicomputers a practical matter [7]. The arrival of the 5-Mbyte Seagate ST 506 in 1980 transformed the personal computer into a “real computer” that could do real problems [1], [3], [7]. Disk drives with capacities of more than 10 Gbytes have crossed over into the consumer video

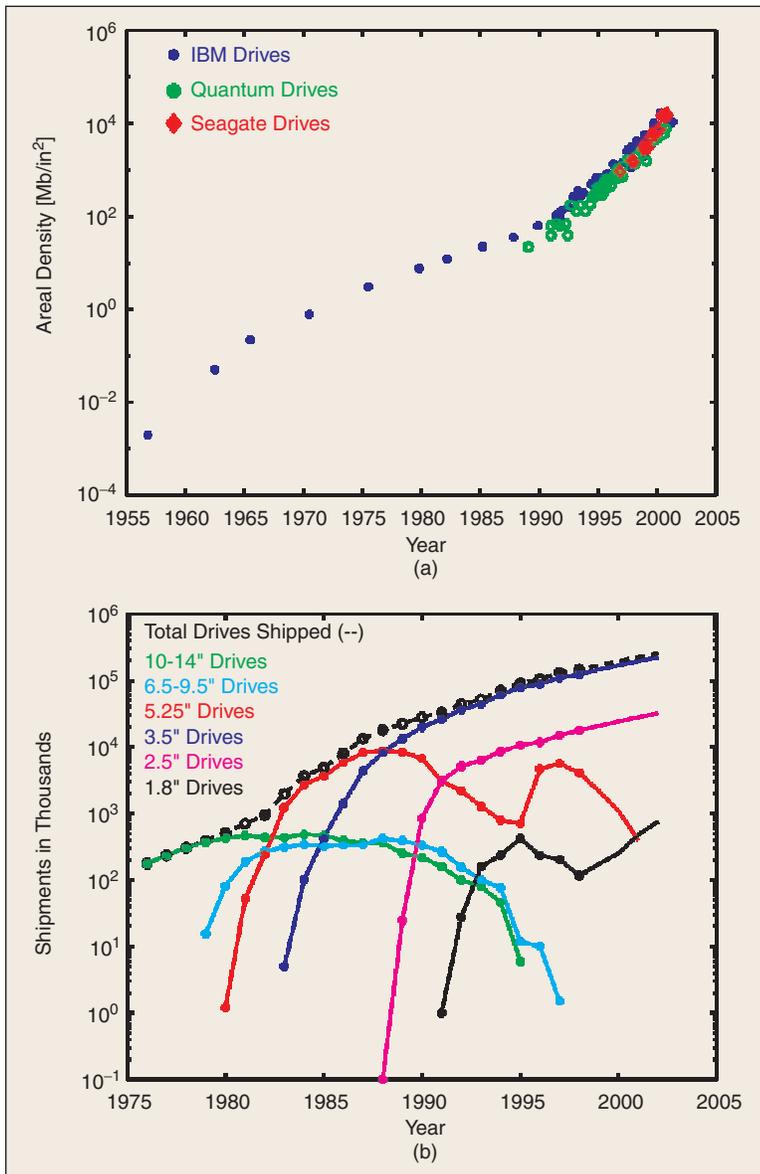


Figure 2. (a) Areal density histories by manufacturer. Data from IBM and Quantum. (b) Disk shipments versus disk diameter. Data from Disk Trend [3]. Based on actual data through 1998; estimates are used after that.

world for use in hard disk recorders that accomplish many of the time-shifting functions that VCRs were supposed to do. Small disk drives may very well revolutionize digital photography, making cameras that hold hundreds of 8-megapixel images (the point at which a digital image has more resolution than most analog photographs on silver halide film) an inexpensive commodity.

Many technology trends have been at work in the history of disk drives, but perhaps none is as important as that of reduced disk size [7]. The power required to rotate a disk scales faster than the fourth power of the diameter [1]. Reducing the size of disks has temporarily lowered the capacity and performance but, in the long term, has enabled much higher storage densities that more than made up for the ca-

capacity losses. Just as important, these small drives have enabled desktop and laptop computers. The effect of opening new computer markets on the number of drives shipped in each size category is evident in Fig. 2, which must be plotted on a vertical log scale just to contain the three-orders-of-magnitude increase in shipments since the mid-1970s. The drive size and the storage technologies that became practical with smaller drives have had a profound effect on the control loops within the drive. The result is the ubiquitous increase in areal density (Fig. 2) through a combination of increases in down-the-track bit density, specified in bits per inch (BPI), and track density, specified in tracks per inch (TPI). The next section will walk around a typical tracking loop of a modern disk drive to illuminate the various components that make it possible.

A Walk Around the Loop

Fig. 3 is a schematic block diagram of a disk drive control loop. The disk loop starts with the disk stack assembly diagrammed in Fig. 4. The magnetic media contains data in concentric circular tracks on both sides.

Modern disk drives read the position of the head relative to the track directly from the disk medium. Virtually all of today's drives use a method called sectored servo, in which user data and position information are multiplexed in space around the disk. As the drive spins, this spatial multiplexing becomes a temporal multiplexing.

The data read heads used to read position and data are universally based on magnetoresistive head technology, which presents some interesting servo challenges. The position information takes the form of a signal modulated into magnetic domains, shown in Fig. 5. From the demodulated signal, the data is digitized and fed into a digital processor for implementation of the control law. The control law typically implements a discrete-time, state-space regulator for track following (when the head is stationary over a single track) and a reference-trajectory-following state-space controller for seeks (moving the head from one track to another). The control law is designed to minimize the effects of internal and external disturbances on the position error signal (PES). The internal disturbances are caused by spindle motion, the airflow over the disk and arms, the noise in sensing the PES, and reactions from seeks. The external disturbances are largely due to shock and vibration. The output of the processor is converted back to an analog signal and sent to a power amplifier, which drives a rotary voice coil actuator. The actuator moves the magnetic

heads through a suspension designed to minimize the effect of the drive mechanics on the servo loop. The suspension also provides a preload to press the sliders down toward the disk in opposition to the air bearing being generated by the spinning disk. At the bottom of the slider are the two magnetic read/write elements. The data is written with a thin-film inductive head and read with a giant magnetoresistive (GMR) head, a descendant of the magnetoresistive (MR) head.

The Evolution of Disk Drive Components

This section gives a brief history of each of the components in the disk servo loop. Understanding this history is critical because of the interplay between the components in the disk drive control problem.

Disks and Media

In a magnetic disk, data is stored on a recording medium (commonly referred to in the industry as simply *the media*), which is responsive to the presence of strong magnetic fields, but stable in their absence. The storage density that a given medium can sustain is determined by a variety of factors, including but not limited to the size and uniformity of the magnetic dipoles in the material, the orientation of the domains, and the coercivity and temperature stability of the media. Since the magnetic field drops off as the cube of the distance between the head and the media, writing and reading smaller spots depends on lowering the distance between the head and the magnetic media. Traditionally, the main component of this has been flying height; however, as the flying height gets lower, other components such as overcoat thickness become more significant factors.

The early disks were physically huge by today's standards. The original RAMAC had 24-in disks. Other designs of the late 1950s had disks as large as 39 in. However, within a few years, the early disk drives settled to an industry standard of 14 in. These disks served primarily the main-frame market and opened up the world of interactive computing. The disks started with a magnesium substrate covered with ferrite media, $\gamma\text{-Fe}_2\text{O}_3$, embedded in a paint similar to that used on the Golden Gate Bridge [1]. However, the substrate evolved to polished aluminum fairly quickly. The ferrite particles were suspended in an epoxy base. The material was applied by spin coating, a process in which material is dropped in the center of a spinning disk and the centrifugal force spreads it evenly across the substrate.

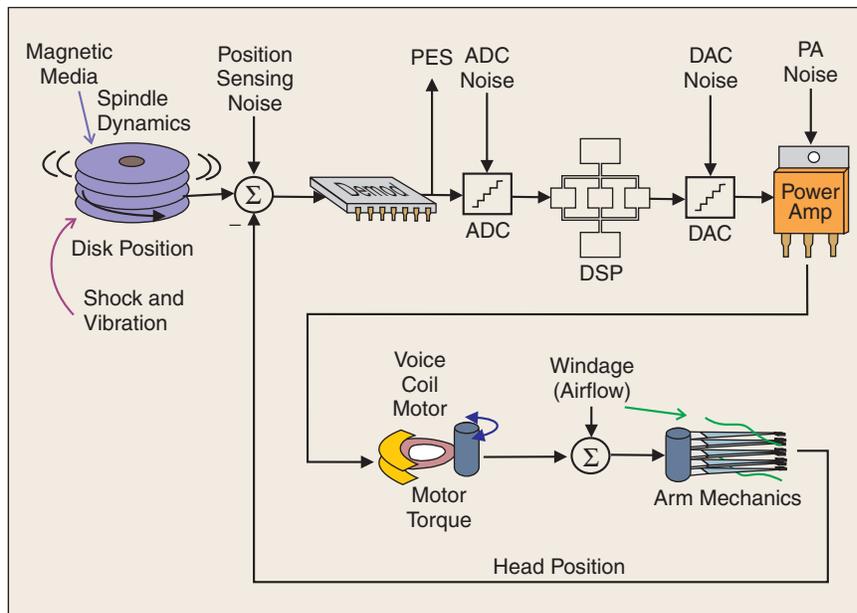


Figure 3. Generalized view of track-following model.

The move to 8-in media was the first sign of the most disruptive technology to hit the disk drive market, that of smaller disk size [7]. Disk shipments by size are shown in Fig. 2. IBM's first 8-in drive, code named Piccolo (IBM 3350), was designed by IBM's British labs and manufactured in Havant, England. This was one of the enabling technologies for the emergence of minicomputers. Although 8-in disks initially held less data than 14-in disks, they allowed several improved features, including smaller power requirements, lower manufacturing costs, reduced airflow, and flatter disks. The smaller size also made possible the use of thin-film sputtering—a chip manufacturing technique—to coat the disks. The new material was smoother and had higher coercivity than the ferrite media. These factors allowed for lower flying heights, which increased bit density and in turn led to a disk capacity that outstripped that of the older 14-in drives.

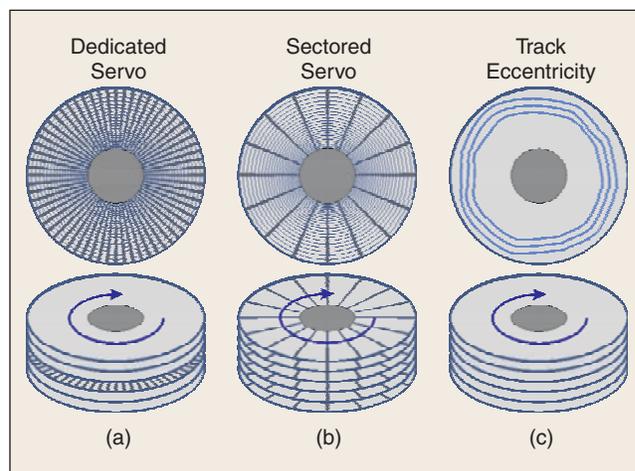


Figure 4. Disk stacks for dedicated and sectored servo systems (left and center) and track eccentricity (right).

One of the key features of MR heads is that they are nonlinear devices, particularly in their offtrack response—a significant issue for any control system.

The 5.25-in drives, pioneered by Seagate [1], [2], were 5-Mbyte full-height drives; they took up an entire drive bay on the original IBM PC. Primitive as they were, they opened up the world of mass storage on PCs. These drives were quickly replaced by half-height 5.25-in drives and then by 3.5-in drives, which opened up another new world of computing—this time for laptops. The moves to 5.25-, 3.5-, and 2.5-in disks (and smaller) have continued to allow more advanced manufacturing methods, smoother surfaces, lower power dissipation, lower flying heights, and higher bit densities. Typical hard disk diameters are 3.5 in for desktop computers, 3 in and smaller for high-speed video and server applications, 2.5 in for laptop computers, and 1.8 or 1 in for PDAs and other mobile applications [7], [8].

Capacities have soared from the original 5 million characters of the RAMAC. A look at the computer advertisements of the *San Jose Mercury News* (in March 2001) shows typical consumer-grade disk drive capacities at 40 to 80 Gbytes priced well below \$200. The current highest capacity disk drive (as of March 2001) is Seagate's Barracuda 180, with 181.6 Gbytes [9].

Media issues pervade the control problem in several ways. First, the same spinning of the disks that produces the air bearing also produces flow-induced vibration—sometimes called windage—an excitation of actuator and disk resonances due to the airflow. Generally, the smaller the disk and the smoother the material, the smaller and less turbulent

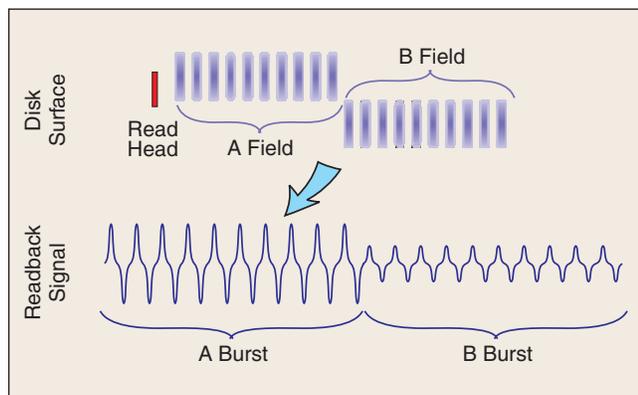


Figure 5. A simplified view of split-field amplitude-encoded servo fields on a hard disk and the resulting readback signals. Here, the A and B fields are separated down the track. Amplitude/area estimates of the A and B fields, A and B , are computed separately and subtracted from each other.

the flow will be. The highest speed drives use 2.5-in diameter disks on a glass substrate [10], [11]. The glass provides a stiffer, smoother disk. The smaller diameter decreases the amount of turbulent airflow, allowing the spindles to be spun at 15,000 revolutions per minute (rpm). The coercivity and stability of the magnetic medium [12], [13] is an issue because servo position is generally not read at the center of a magnetic domain, but at the side edge of a domain, where the edges are often defined by stray fields.

Heads

The magnetic heads are the parts of the drive that use electrical signals to write magnetic flux patterns on the disk and read the flux patterns back to electrical signals. There have been five large steps (and hundreds of small ones) in the evolution of magnetic read/write heads: permalloy heads, ferrite heads, MIG heads, thin-film inductive heads, and finally MR heads and their descendants. No matter what the material, magnetic recording heads have been inductive heads throughout most of the history of magnetic recording.

Inductive heads have one major limiting factor: they are only responsive to the change in magnetic flux that occurs at the transition of magnetic domains. If the domain transitions do not pass frequently enough, the read signal level drops too low for adequate signal-to-noise ratio (SNR). The device that solved this problem is the MR readback head, which is sensitive to the magnetic flux itself, rather than the change in flux. The design and use of MR heads can (and does) fill several textbooks [12], [13]. The first transducers arrived in the 1970s, but it was not until 1991 that IBM first used one in a disk drive [1]. Since then, the MR head is credited with bumping the areal density improvements from a rate of approximately 30% per year before 1991 to a rate of over 60% per year [3].

One of the key features of MR heads is that they are nonlinear devices, particularly in their offtrack response—a significant issue for any control system. Furthermore, MR heads are read-only devices, meaning they cannot be used to write the data. An inductive head is still needed to write the data, and this must reside on the same slider as the MR head. This raises an issue with a rotary actuator. Because the heads must be offset from each other, the skew angle of the slider results in the read and write heads having slightly different angles for which they are over the track center. This means that in between reading position information and writing user data, the servo system must perform a “micro-jog” to shift into a position where the write head is over the track center.

Sliders

One enabling technology for hard disk drives is the concept of floating the magnetic read/write head on a thin air bearing.

Note that the heads themselves are built into a mechanical structure called a slider, which provides the interface between the disk arm and the magnetic medium. This concept was studied as early as 1953 for thin magnetic drums. IBM used this technology in the RAMAC in 1956 [1]. The air bearing had the beneficial property that it maintained a minimum distance from the disk while the Bernoulli principle limited the maximum distance from the disk. However, these early air bearings required that air be pumped through the heads to push them away from the disk. This kept the size of the heads and the flying height large (800 μm for RAMAC).

The next advance was the invention of the self-actuated hydrodynamic air bearing, a bearing generated by the spinning disk and the slider themselves with no need for a pumping mechanism. The work began as part of IBM's Advanced Disk File (ADF) in 1955 and introduced what became known as the comb actuator. This actuator had one suspension, slider, and head per disk surface that were actuated as a unit to a specific radial track position. While actuators have moved from linear to rotary motion, the "comb" structure has remained. This actuator also allowed cylinder mode operation, in which for any radial track position a different disk could be selected via electronic rather than mechanical switching. IBM brought this slider to market in 1962 with the 1301, which dropped the flying height to 250 μm . Successive models cut these heights to the point where today's sliders fly well under 3 μm from the disk.

Flying height is a critical factor for storage density because the strength of the magnetic field varies with the cube of the distance between the head and the disk. Thus, lowering the flying height has been a critical step in increasing storage densities, even as spindle speeds increased. Smaller sliders have dropped the flying height as well. To avoid start-stop friction issues, the bottom surfaces of these sliders had to be textured and lubricated.

Actuators

The original RAMAC actuator had a single pair of heads that moved both vertically and radially to access data (Fig. 6(a)). It used aircraft cable and pulleys, but the next products and on through the IBM 2314—the main product until the IBM 3330 Merlin drive was introduced in 1971—all used hydraulic actuators [14]. The first linear voice coil actuator was developed by IBM in 1965 (Fig. 6(b)). The next major step was a move to rotary actuators, first designed at IBM's Winchester Labs (in Winchester, U.K.) [9].

The rotary actuators for 8-in drives were relatively large and complicated mechanical truss structures. On these actuators, the suspension was turned sideways to most closely mimic the motion of the linear actuator (Fig. 6(c)). As drives shrank once again, the actuator pivot was put in

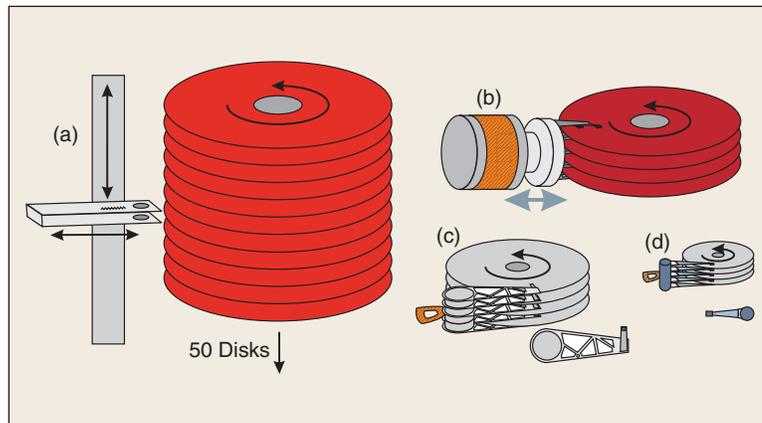


Figure 6. Generalized view of the evolution of actuators. Note the transition from the RAMAC (a) to the comb structure, which had one head per surface (b). The next transition is from linear to rotary (c), then from larger rotary to smaller, stiffer rotary actuators (d).

one corner of the enclosure. The actuator and suspension were being dragged over the disk in a single line (Fig. 6(d)). As actuator sizes have shrunk, the structural resonances have moved to higher frequencies. In contrast, the effects of friction in the pivot bearing have become more noticeable, as discussed in the section on friction.

Spindles

Spindle speed is held fixed with a low sample rate feedback loop. The speed affects not only the flying height of the disk (and therefore the bit density), but is also a major source of disturbances in the tracking loop [15]. The spindle bearings are also a major source of audible noise in disk drives. Because spindle speeds affect the airflow and disturbances in the tracking loop, higher-speed spindles are often implemented only on smaller-diameter drives. Speeds started at 1200 rpm for the RAMAC [2], hovered at 3600/4200 rpm for the 14-in mainframe drives [16], rose to 5400 rpm for 5.25-, 3.5-, and 2.5-in drives, and have increased to 7200, 10,000, and 15,000 rpm for successively faster high-performance server drives [9], [10].

Related to spindle speeds is the issue of enclosures. They were originally sealed in the RAMAC, then opened up in the era of removable disk packs, then sealed again in the Winchester era. With a few exceptions they have remained sealed ever since. The critical factor here has been design of the airflow inside the enclosure to ensure proper heat dissipation, removal of contaminants, and nonturbulent flow around the heads and actuators.

Channels and Electronics

The original RAMAC was built at the beginning of the transistor era. Thus, the designers chose the proven technology of tubes for the early control system [1]. Very quickly, though, the drive electronics became dominated by analog transistors. The use of sophisticated digital electronics be-

came prevalent as drive control systems moved to digital control using digital signal processors (DSPs) [17] and as the data channels went to digital partial response, maximum likelihood (PRML). As drives have become commodity items, the push to lower the chip count has reduced the total number of chips even as the sophistication has increased. The stated goal of many disk drive manufacturers is to package the entire electronics of a disk drive on a single chip. The combination of mixed signals and differing power requirements have slowed this effort.

In the disk drive business, companies would unknowingly mortgage their futures to protect established businesses.

The data channel, the means by which data is actually written to and read from the disk medium, affects the reading of position information in three ways. First, the magnetic domains on a disk drive have maintained a rectangular shape that is roughly 20 times as wide across the track as in the down-the-track direction (although there is increasing pressure to reduce this ratio to four or five). Thus, the bit density often determines the target track density. Second, the head/media combination needed to maintain a given bit density typically has a strong effect on the track density as well. The one caveat is that bit densities are determined by the ability to distinguish transitions at the center of a track, but track densities are determined by reading information at the side of a domain (see Fig. 5). Third, the leading edge of the sophisticated electronics used in a disk drive is in the data channel. As such, a look at the data channels provides a picture of where the servo channel electronics are going.

Servo Signals

The RAMAC “detected” positions by detent marks on the actuators. The early comb actuators were hydraulic until IBM introduced the first voice coil motor in a drive in 1965 [1]. These were open loop until 1971, when IBM introduced the first disk drive that closed the loop with position information read from the loop [18].

Modern disk drives read the position of the head relative to the track directly from the disk media. Over the history of closed-loop control of disk drives, there have been two essential choices for encoding this position information: dedicated and embedded (or sectored). Dedicated servo involves reserving an entire disk surface for position information, leaving the other surfaces free to contain only user data, as shown in the left diagram of Fig. 4. Embedded servo time multiplexes the servo information with the user data on each

surface, as shown in the center diagram of Fig. 4. Dedicated servos have the advantages of higher sample rates and a possible savings in surface area when the total number of disk surfaces is relatively high. On the other hand, they are inappropriate for single-surface systems, poor choices for single-disk systems, and typically are more susceptible to thermal offsets than embedded servos. To minimize the effects of thermal offsets, the servo information on a dedicated servo system is usually encoded on one of the center surfaces, as shown in Fig. 4. Embedded or sectored servos, as shown in the center diagram of Fig. 4, collocate the position sensing with the control but force the servo designers to choose between higher sample rates (desirable) and lower user data density (undesirable). However, as track densities have increased, the thermal offsets in the head stack assembly have become too large a percentage of the track to do anything other than embedded servos (collocated control).

Whether dedicated or embedded, the position information must be placed on the drive at the time of manufacture. The format is created by servowriting, a process in which the read/write head is externally controlled to write position information magnetically on the disk surfaces. The diagrams in Fig. 4 show how a hard disk stack assembly will have individual tracks that are largely concentric with the disk spindle. Due to the imperfect nature of the spindle and the servowriting process, however, the tracks themselves are noncircular. This noncircular position information is encoded at the spindle frequency, yielding noncircular tracks for the disk servo to follow, as diagramed on the right side of Fig. 4. Servowriting on the drive spindle tends to minimize the repeatable runout (RRO). The servowritten track is strongly affected by the runout of the spindle. The servowriter can minimize written errors by writing position information in multiple passes; however, this dramatically increases the cost of servowriting.

Apart from material expenses, the cost of servowriting is one of the largest in the manufacture of a hard disk and is one of the factors pushing magnetic media makers toward methods other than traditional servowriting. The first is to pattern features into the disk, as is done in optical disks, and use these features to read position information. One version of this method involves actually patterning a disk with pits and using a laser diode on the side of the slider to read position [19]. A second version, dubbed pre-embossed rigid magnetic (PERM), involves permanently embossing the disk substrate with discrete tracks. If either of these methods could be made cost-effective, the use of harmonic correctors would make them more practical from a control standpoint by minimizing the effects of the disk-to-spindle misalignment on the control loop [20]. An alternative method for saving on this expense is the concept of a disk drive that can servowrite itself. Such a drive would not only

save on capital equipment, but also cut manufacturing time by allowing the servowriting to be done during the burn-in stage of drive manufacture. It is interesting to note that the technology to do this is considered important enough to be listed in the stock prospectus for the buyout by Maxtor of Quantum's Hard Disk Division [21].

Initiatives to improve the runout of the spindles by replacing ball-bearing spindles with fluid or air-bearing spindles [16] have not caught on for the purpose of runout reduction. However, fluid-bearing spindles have started appearing in drives for acoustic noise reduction [10].

A variety of position-encoding methods have been used to encode the servo position, including amplitude, phase, frequency, and null [22] encoding. Fig. 5 shows a diagram of amplitude-encoded servo patterns. Peak detection channels have traditionally been used to detect position information. Area detection has shown higher noise immunity than peak detection but is still susceptible to the biasing of small signals [23]. A matched filter demodulator would improve noise immunity but would still be susceptible to MR head nonlinearity [23]. Constructing the matched filter out of useful "custom harmonics" may hold the solution here [23]. However, with higher SNRs for current servo channels, largely due to new GMR heads, the natural reluctance of disk drive engineers to add new features will probably keep averaging peak detectors in disk drives for a few more years. It is rumored that several disk drive companies are using discrete fourier transform (DFT) demodulation, largely equivalent to the custom harmonic method when the latter is implemented to demodulate only the first harmonic.

The Control Problem

With the above background and component histories, we can now delve into the history of the disk drive feedback loops themselves. Far from being an isolated problem, the control loops were tightly coupled to all the component technologies that were being developed.

The first commercial hard disk to read the relative position of the read/write head from the disk itself and close a feedback loop around this was IBM's 3330 Merlin drive, which appeared in 1971. However, it was IBM's 3340 Winchester drive, which shipped in 1973, that set the architecture for future disk drive control loops. The Winchester drive is mechanically significant for its use of lightly loaded, lubricated, low-mass sliders, but its control system is most significant in that it is the first drive in which all the pieces of a disk drive control loop were in place. The paper written by Dick Oswald, a lead servo engineer on the project, has been the classic starting point for disk servo engineers [18]. It was common practice at some disk drive companies for engineers to arrive for their first day of work to find a copy of

Oswald's paper on their desks [24]. IBM had also been working on the use of bang-bang control for seeks. However, in Oswald's paper, a modification of the bang-bang scheme, which would later be codified as a proximate time optimal servo (PTOS) [17], was first discussed. With all of these techniques coming together, one can see the architecture of future disk drive control systems in Oswald's paper.

Over the history of closed-loop control of disk drives, there have been two choices for encoding position information: dedicated and embedded.

There was an old saying in the computer industry: if IBM published something, it meant they had decided not to use it. Information available for this section often has a similar flavor. Most disk drive companies are fairly tight-lipped about their servo work, and thus much of the history that follows is from published work, patented inventions, consortium meetings, the work of companies no longer in the industry, work about which the authors have direct knowledge, and/or industry rumors. Although this provides a rich tapestry of work to draw from, the reader is cautioned against thinking that it is complete. There is much that the drive companies are just not telling us. As for the publications, they tend to be either the result of industry-academic collaboration or industry-patented work. By and large, disk drive companies have cross-license agreements on all their patents with their competitors, causing many to shy away from patents (public documents) in favor of keeping information as trade secret (private). Finally, disk drive companies and servo engineers are usually risk averse. This tends to make them extremely slow to pick up new algorithms, much to the chagrin of their academic collaborators. With these caveats in mind, the rest of this section will delve into specific concentrations of work on hard disk servos.

Analog, Sampled Data, and Digital Control

From the beginning, disk drives had digital circuitry to relay the data to and from the computer, although the control loops were analog. Early digital controllers on drives emerged for reasons of either economy or performance. An example of using digital control for economic purposes came from Quantum, which produced some low-end drives with a computer interface microprocessor. By using this processor to do servo control, they were able to save on the cost of analog electronics. IBM, on the other hand, was working to implement advanced algorithms on their drives for minicomputers

[25] and mainframes [17]. Mike Sidman at DEC used an 8085 microprocessor to do harmonic correction in 1978. Sector servo led to sampled data control systems, even when the control laws were implemented using analog electronics. Hewlett-Packard (HP) had drives of this nature through the mid-1980s. It was not until the late 1980s that HP moved to fully digitally controlled drives, largely due to the interaction between Rick Ehrlich (then at HP Labs) and Vernon Knowles at HP's Disk Memory Division [26].

Another interesting start for digital control was relayed by Fred Kurzweil, who has the distinction of being Gene Franklin's first graduate student at Stanford, graduating in 1959 and going to work for IBM Research in San Jose. After 23 years of doing mostly theoretical work on disk drive control, he took early retirement and went to Maxtor, which had just opened its doors. Upon his arrival, he was thrown into the fire of having to make the first in-the-hub spindle controller work. Upon hearing stories circulating about this project, industry pioneer Al Shugart said it would never work.

Putting the motor in the hub raised some serious issues. The motor consumed 10 W of power, and the heat inside the hub built up, changing the viscosity of the ball-bearing grease. Kurzweil solved this problem with a 10-cent microprocessor and some simple digital control. To maintain the spindle speed to one part in a million, he used simple adaptation: correcting for spindle motor changes with a very low order corrector. As the rest of the industry had listened to Shugart, Maxtor had a two-year lead on in-the-hub spindle motors [27].

State Space

With DSPs and digital control on hard disks, the possibility of doing state-space control emerged. Early disk drive digital control systems used classical design methods. The ease of doing classical design from a measured frequency response function was a factor in this. The first use of state-space control seems to have been during seek mode. There are a couple of reasons for this. First, the number of operations for a state-space controller is typically larger than those for a classical controller on the same order problem. Early DSPs were hard pressed to do all the extra operations for state-space control in a single sample interval while track following. The second reason is that a measurement of the back EMF from the voice coil motor—which could be used to estimate velocity—was only useful when the signal was large (i.e., during seek). The use of state-space control to make use of the back EMF sensor during seek while using a classical design in track following continued at Quantum into the late 1980s and at HP until they got out of the drive business in 1996.

On the other hand, IBM Rochester got into the state-space paradigm early on, starting on digital control in 1980. They had obtained a recently released military microprocessor with DSP capability (8196) from Intel. Mike Stich, a servo engineer out of Rochester, took a digital controls

short course from Franklin. On his return to Rochester, he and Hal Ottesen set up a teaching/consulting relationship with Franklin that lasted ten years. An early result of this collaboration was the IBM 9332, which made use of a digital state-space controller that even did online parameter adaptation [25]. IBM Rochester had relatively little pain going into state-space control. There was nobody around who had any experience with digital control, so they were on their own to experiment as they thought best [28].

The case history on disk drive control written by Mike Workman (of IBM San Jose) [17] makes use of state-space control as well, indicating that IBM San Jose had also fully embraced the approach, albeit after Rochester did. In fact, Hal Ottesen recalls that when the researchers from Rochester first proposed digital control to the folks at San Jose, they were laughed out of the room [28]. IBM San Jose's move to digital and state-space control was largely led by Mike Workman, who had been taking classes from Franklin at Stanford and got "the religion." IBM San Jose's first digital control drive was the IBM 3380K.

Quantum's first drive to use state-space control for seeking and tracking was the embedded servo high-end Enterprise drive, which went into mass production in early 1992. The desktop drives went to single-loop state-space control a few years later (when the microprocessor was fast enough that they didn't have to drop to a simple servo in on-track mode to make time for the I/O firmware) [26].

Currently, it is believed that most if not all disk drive controllers are state-space controllers.

Sample Rates and Auxiliary Sensors

The multiplexing of position information with user data on hard disks creates a set of competing objectives. On one side is the desire for maximum data storage, which would push to minimize the number of servo fields within a track. On the other side is the desire for improved performance in the control system, which often requires a higher sample rate. These tradeoffs have limited the achievable sample rates to the range of 6-14 kHz. This in turn has limited the achievable tracking closed-loop bandwidth to the range of 500-1000 Hz.

This has led to a fair amount of work in multirate servos, where the control output is changed at a significantly higher frequency than the sample rate of the PES. Rick Ehrlich had been pursuing the idea of multirate control ever since he was at Hewlett-Packard. Carl Taussig continued this work, coming up with unpublished results quite similar to those reported by W. W. Chiang (out of IBM San Jose) [29]. It turns out that Rick Ehrlich had continued his multirate work at Quantum and was getting similar results. Basically, if the PES sample rate was relatively low—say, less than eight times the open-loop crossover—keeping everything else constant and raising the output sample rate by a factor of three or four over the input sample rate could result in a closed-loop bandwidth improvement of roughly 20%. If, on

the other hand, the PES sample rate was already at roughly 20 times the open-loop crossover, then the improvement was far smaller. Even recent papers on multirate servos have implemented a lot of machinery but have difficulty pointing to quantifiable improvements in the servo loop [30]. Thus, the general use of multirate servos has not been pursued.

The one place where multirate on the main loop has found applicability has been in the use of multirate notches. The ability of designers to design sharp notch filters had been limited by the relatively low sample rate of PES compared to the frequency of the actuator resonances. To combat this, a group of engineers at Quantum (Rick Ehrlich, David Jeppson, and Phil Weaver) came up with a multirate notch filter [31]. A similar approach was used by Don Fasen at HP's Disk Memory Division, although this was never published or patented.

The multirate approach has also been applied to the use of auxiliary sensors in disk drives. In cases where the sensor can sample information independently of the structure of the position information on the disk, performance can be improved by raising the sample rate of the auxiliary sensor and the update rate of the control law [32], [33]. These auxiliary sensors come in three main forms: accelerometers for rejection of internal and external disturbances, extra sensing of the back EMF mentioned above, and instrumented suspension.

Noise Sources

This section will examine the external and internal disturbances that affect hard disks, how these have been measured, and what has been attempted to compensate for them. The external disturbances are typically in the form of environmental shock and vibration, whether due to a moving vehicle, a factory floor environment, a computer under a desk being kicked, or simply the motion of a laptop computer. Internal disturbances are largely stimulated by the spindle's rotation of the disk and the actuator's reaction forces on the drive baseplate and housing.

Repetitive and Spectral Disturbances

One of the concepts to become popular in disk drive control systems at the end of the 1980s was the use of repetitive control to cancel the effects of the spindle eccentricity. Some early unpublished work had been done by Mike Sidman for the DEC RC25 in 1978. However broad use in the drive industry started when Masayoshi Tomizuka's group was working on practical applications of repetitive control as a solution to repetitive disturbances in rotating machinery. Obtaining some disk drives from IBM, but unable to access the drive DSPs to change the code, they used the notion of an add-on controller that would augment the nominal loop to remove the harmonic disturbances [34], [35]. As graduate students branched out from UC-Berkeley (UCB), especially to professorships at Carnegie-Mellon University (CMU) (e.g., Marc Bodson and Bill Messner), these studies also included adaptive feedforward

harmonic cancellers, which were shown to have some equivalence with repetitive controllers.

Despite the presence of repetitive and harmonic correctors, spectral disturbances that were related to the spindle frequency but were not repetitive remained. In a series of studies that would be widely followed, HP's Jeff McAllister characterized these as being driven by an interaction between the spindle rotation and the oscillations of the hard disks themselves [36]. McAllister proposed that the solution lay in either modifying the airflow or changing the substrate with an internal layer of viscoelastic damping material [36]. Since that time, smaller glass disks have emerged as a solution to this problem. The disks themselves are stiffer, and the smaller size reduces the airflow that drives these oscillations. Glass disks also have the advantage of already having been tested on small-form-factor portable drives. In fact, the most recent high-performance drives produced by Seagate and IBM use disk sizes that are roughly between 3.0 and 2.5 in in diameter with a glass substrate [10], [11].

External Shock and Vibration

As drives became smaller and moved to more mobile applications in the early 1990s, the issue of rejecting external disturbances, namely, shock and vibration, became more prominent. An enabling factor was the continuing drop in the cost of accelerometers, to the point where they could reasonably be considered as an option for disk drives.

The earliest examples of accelerometer control in a disk drive, however, go all the way back to the 1970s. Robert White used an accelerometer to sense shock and then minimize the probability of the heads slapping against the magnetic media [37].

Typical use of accelerometers involves sensing the disturbance and moving the actuator before the error ever shows up in the position error signal. As the accelerometer and drive characteristics are subject to change, adaptive methods are often used. Generally, the use of accelerometer feedforward dramatically improves the disturbance rejection capabilities of hard disks. Cost and reliability issues for the accelerometers themselves have limited this practice.

Danny Abramovitch applied an online adaptive, multirate feedforward accelerometer compensation scheme [32], which was an extension of earlier work by Vern Knowles and Mitch Hanks. This work went fallow when HP abandoned the KittyHawk drive in 1995. Davies, an MIT graduate student, and Mike Sidman, of Digital Equipment Corporation (DEC) in Colorado Springs, formulated conditions by which an accelerometer could perfectly cancel external and internal disturbances [38]. There was also considerable activity originating at the Berkeley Sensors and Actuators Center (BSAC) and the UCB ME Department, including work by Matt White, a student of Masayoshi Tomizuka, who developed accelerometer-based feedforward control schemes to detect and reject external

disturbances [39]. He considered two types of controllers: a fixed infinite impulse response filter and an adaptive finite impulse response filter. A survey of these methods was recently published by Charlie Hernden [40].

Broadband Noise

In the mid- to late 1990s, a strong interest in understanding and dealing with broadband noise arose in the disk drive industry. One large body of work was started in 1994 by Danny Abramovitch, Terril Hurst, and Dick Henze at HP Labs [15]. A parallel effort seems to have taken place at Maxtor in the late 1990s. A major realization of the work was that the most significant broadband noise sources in the drive were the air impacting on the disks and actuators (windage) and PES noise [15], [41]-[43]. Rumors seem to indicate that the work has been broadly applied by other disk drive companies.

The airflow affects the servo due to the mechanical interface of the head and the disk through the air bearing. These disturbances affect the dynamics of the actuators and the disks [36] and end up in the PES [15]. The nature of the flow depends on the disk diameter, the position and aerodynamics of the actuator, the enclosure properties, the disk materials, and the air pressure in the enclosure [44]. With the high spindle speeds of modern hard disks (7,200-15,000 rpm) [10], correcting the internal aerodynamics of disk drives has become a critical issue. Unfortunately, it is hard to get recent precise results from most drive companies since the knowledge tends to be mostly empirical, the designs are quite sensitive to small changes, and the companies consider this hard-won information a proprietary advantage.

Noise Sources in Magnetic Servo Signals

Fig. 5 shows a split-field amplitude-encoded pattern. Theoretically, the center portions of the pattern would be enough to determine position within a track. However, the linearity and SNR of the signal can be dramatically increased by adding a set of patterns in quadrature with the first, which is used when the head approaches the edge of a track. If the in-phase (center of track) signal is determined from the A and B fields, then the quadrature signal would be determined from the C and D fields (not shown).

For hard disks, a large amount of the noise comes from media noise at the domains that define the A, B, C, and D fields. Much of this noise is due to the fact that when the head is on track, it is flying over the domain edges of both the A and B fields [22]. There are also noises and effects in the MR heads [22], [23], [45]. Furthermore, the demodulation of these signals maps high-frequency noise into the baseband, where it is seen by the servo loop [23].

MR heads and their descendants have been a boon with respect to raising the bit density, but they cause all sorts of problems for the servo. MR heads have nonlinear behavior across the track [46], [47], which is exactly where the servo signal needs linear behavior. The nonlinear behavior of MR heads means that a tremendous amount of work must be done to make them suitable for use in a feedback loop [22]. Besides the issue of nonlinearity across the track, there are also issues such as baseline shifting, sensitivity to thermal asperities, and baseline popping [23], [47], to which the control loop must be robust.

Resonances

One of the limiting factors for actuators is the resonances of the mechanics. The rotary actuators of hard disks have a large number of resonances, as shown in Fig. 7. In practice, however, the response of the system is dominated by three main resonances, typically the first torsion and sway modes and either the second torsion or the second sway mode. Other resonances are of lesser importance because they are either farther out in frequency or are overlaid in frequency by the main modes. As of 1999, commercial suspensions had first and second torsional modes in the ranges of 1,500-2,500 Hz and 4,800-8,600 Hz, whereas the first sway mode was in the range of 8,000-12,000 Hz [33]. Drive designers would prefer lighter, stiffer suspensions than the aluminum ones they currently use. Some suggest using composite materials [48]; others suggest using beryllium in the arm to stiffen and dampen it.

Actuator resonances limit bandwidth because they are lightly damped and subject to variation from drive to drive and in a single drive over time. Thus, it is extremely difficult to control through the actuator resonances. Instead, control designers have tried to notch these out. The ability to implement these notches has been limited by the relatively low sample rates of hard disks. This has led to such work as the multirate notch filter discussed earlier. The limitations of multirate methods in dealing with resonances have led to a pair of techniques to be discussed later.

Friction

In hard disks, friction in the rotary actuator pivot is an ongoing issue studied by both the mechanical and servo portions of a disk drive team. It was

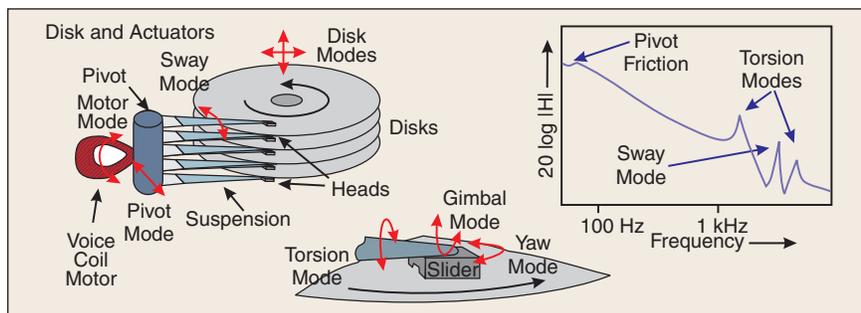


Figure 7. Pictorial view of hard disks and actuators, as well as their modes.

first noticed on the Winchester's linear actuator when a technician inadvertently set the lower frequency of his swept-sine measurement at 10 Hz rather than the customary 100 Hz [49]. While relatively insignificant for the 3-lb Winchester actuator, the issue becomes more noticeable as the actuator inertia drops [50].

Starting in the early 1990s, there were several publicized efforts to analyze and mitigate friction. Quantum's Mike Hatch and Bill Moon used time domain measurements to establish a hysteretic relationship between position and velocity of the actuator, which they reported at a National Storage Industry Consortium (NSIC) meeting in February of 1992. HP's KittyHawk drive (1.3 in) faced what the engineers at HP called a "low-frequency pole." Having witnessed Hatch and Moon's presentation, HP Labs engineers Danny Abramovitch and Dick Henze began to study pivot friction. Feei Wang (a graduate student of Franklin's) and Abramovitch came up with a relationship between swept-sine frequency response function measurements and describing functions that allowed them to develop a friction model that incorporated elements of the Dahl model [51] (well known in tribology circles) with the traditional velocity feedback term. The addition of Terril Hurst to the project led to the use of restoring force methods to characterize the friction [50].

More recently, Jun Ishikawa of Nippon Electric Corporation (NEC) Research Lab, working with Masayoshi Tomizuka while visiting Berkeley, used a disturbance observer that fed back the filtered difference between the sensed voice coil current and the signal from an arm accelerometer to generate a pivot friction cancellation signal [52]. Unfortunately, neither Ishikawa nor NEC pursued the result further since NEC exited from the disk business soon after he returned to Japan. Although it is pretty clear that work on friction continues at drive companies, it does not appear that much more of it has been published.

Dual-Stage Actuators

The success of dual-stage actuators in the tracking loops of optical disks has led to many proposals to add these to hard disks. The reasons for doing so include: lowered inertia and power requirements, mitigation of pivot friction effects, and avoidance of suspension resonances. The proposals along these lines typically consist of either micromachining a second stage onto a slider or at the gimbal [53]-[55] (which achieves all objectives) or including piezoelectric actuation in the suspension [55] (which achieves the third and part of the first two objectives). As micromachined actuators are the only solution that solves all of the above problems, their adoption is widely considered inevitable, although a few experienced researchers believe that cost and reliability factors will limit their use. That being said, predictions on a time scale for their widespread use range from five to ten

years. Piezoelectric actuators are considered an interim method until micromachined actuators become practical [55]. Some predict widespread use of piezoelectric actuators by 2003 and that, by 2007, all hard disks will employ micromachined second-stage actuators [54]. Others be-

Far from being solved, the disk drive control problem should keep engineers busy for many years to come.

lieve that the industry will completely pass over the piezoelectric options and hold out for micromachined actuators.

A more modest approach involves adding sensors to the suspension, allowing for better control through the resonances of the actuator [33] in a manner similar to a flexible robot system. Mike Banther, an HP engineer, did some collaboration with Bob Evans at Hutchinson Technology while the former was a student of Messner at CMU. Each obtained different results for instrumented suspensions, resulting in a patent for Evans [56] and a paper for Banther [33]. Interestingly, this technology has not been pursued much since then, as drive makers have moved closer to using second-stage actuators.

A design technique for such dual-stage actuators that has drawn a fair amount of praise is that of Steven Schroeck and Bill Messner [57]. Their so-called PQ method is a variation on successive loop closures that allows the designer to work in the frequency domain and produces reasonable control designs for dual-input, single-output systems, such as drives with second-stage actuators.

Seeks

The IBM 3340 Winchester [18] used a near-time-optimal seek mechanism that approximated the performance of a bang-bang controller while avoiding the problem of having the control signal "chatter" between maximum and minimum values in response to even small errors. This technique would become standard practice in the industry, but it was not until Mike Workman wrote his doctoral thesis at Stanford that the method would be codified as proximate time optimal servomechanism (PTOS) [17]. The adaptive version was named APTOS (adaptive PTOS), in part because the name reminded Workman's wife Patti of a town in the San Francisco Bay area. Further work to extend PTOS beyond the saturating double integrator and into flexible structures was pursued by Lucy Pao, another student of Franklin's at Stanford [58].

One issue with seeks is that the reference positions given to the control system often excite the flexible modes of the actuator, which take a long time to damp out and effectively

lengthen the seek time. Neil Singer, a graduate student at MIT under Warren Seering, developed a method called command input shaping for use in flexible robotics [59]. The idea is to prefilter the reference command in such a way as to remove spectral components that will stimulate the residual vibration modes of the actuator. It turned out that Singer was a friend of Carl Taussig, who at that time was working on disk drive research at HP Labs. A collaborative effort ensued to apply this technique to disk drive manipulators. Lucy Pao's group got into the act through a chance meeting with Bill Singhose (who worked at Singer's company). Most recently, they have been trying to merge PTOS with input shaping [60].

As a drive initiates a seek, it switches in the seek control algorithm. During deceleration, the drive follows a velocity profile into the target position, then switches to settle mode, and finally into track-following mode. Quite often, the latter involves switching an extra integrator into the control loop. During the early days of disk drives, this mode switching was accomplished by switching in different electronics for each region. As drives went to digital control, they did not switch electronics, but the control law was switched in the microprocessor. Once track-following mode was switched in, an extra integrator was added to the control law computation. Although combined seek and track-following algorithms were codified in PTOS, it seems that, even today, most disk drive servo systems still go through at least three stages of control [26], [28].

Centers of Activity

Over the years, a handful of centers have produced a large amount of work on disk drive servos. This section provides an overview of some of these efforts.

IBM's San Jose facility was largely built around the disk drive business. Many of the people mentioned in this article have worked there at one time or another. It is worth noting that IBM's disk drive operations were originally tied to their various computer operations and thus tended to take on the culture of the computer operations, rather than being monolithic within the drive operation itself. In Rochester's heyday, they made drives for the minicomputer market. As they competed with DEC, they had a relatively open attitude. San Jose, involved with the mainframe industry in which IBM had a virtual monopoly, tended to be more closed.

IBM has had close ties with UC-Berkeley's Mechanical Engineering Department, particularly the students of Masayoshi Tomizuka, Roberto Horowitz, Dave Bogy, and Al Pisano. Many of these were also involved in BSAC, a center for micromachined actuators and accelerometers. Graduates of these groups who have gone into the disk drive control world include Dave Horsely (formerly HP), Matt White (IBM Almaden), Phil Weaver (formerly Quantum), Bill Messner (CMU), Dick Henze (HP), and Carol Wilson (formerly HP). Ho Seong Lee worked at IBM before doing his thesis with Tomizuka and then going to Seagate and then

Maxtor. Lin Guo (of Maxtor) was also a student of Tomizuka, as was Li Yi (now at Western Digital (WD)). Ram Ramakrishnan started out doing disk drive research at HP Labs alongside Rick Ehrlich. He then went to UC-Berkeley to do his doctoral work under Tomizuka in the area of a single-link, flexible robot arm, but then returned to HP and stayed away from storage. Denny Miu (now at Integrated Micromachines) worked on control systems with Tomizuka before doing MEMS work for disk drives at UCLA and Cal Tech.

CMU's ME department has continued to produce excellent work in disk drive control. Much of the work started with Marc Bodson's students but, since his departure for the University of Utah, has been continued by Bill Messner's group. Students in the disk drive industry include Alexi Sacks, Kyle Eddy, P.D. Mathur (Seagate in Minneapolis), and Steven Schroeck (Western Digital). Seagate's site in Minneapolis is originally part of Control Data Corporation (Imprimis).

The influence of Stanford University has largely been through former students of Gene Franklin in the Electrical Engineering Department (Mike Workman at IBM and Conner, Mike Stich at IBM, Mike Sidman at DEC, Lucy Pao at the University of Colorado, Fred Hansen at Conner, Seagate, and Quantum/Maxtor, Dick Curran at Quantum/Maxtor, Danny Abramovitch at HP). Karl A. Belser (IBM, Quinta, and Seagate) was a graduate student under Bernard Widrow at Stanford, but Franklin was his associate adviser and the second reader on his thesis. Fred Kurzweil, Franklin's first graduate student at Stanford, took a job at IBM Research in San Jose before moving on to Maxtor and then Areal Technology.

Among MIT's additions to drive servo work are Carl Taussig and Rich Elder (HP), Rick Ehrlich (HP, Quantum/Maxtor), and Neil Singer (Convolve).

DEC had drive design centers in Maynard and Shrewsbury, MA, and a Servo/Mechanical Advanced Development Group in Colorado Springs. The latter was headed by Mike Sidman, who has continued to consult for the drive industry since leaving DEC. The Shrewsbury site is now part of Maxtor.

Hewlett-Packard's R&D efforts have largely surpassed their success in the disk drive market, causing HP to exit the manufacture of disk drives in July of 1996. Vern Knowles, Mitch Hanks, Bob Davidson, and Jeff McAllister, all from the former Disk Memory Division in Boise, have written numerous publications on disk drive controls. Besides those already mentioned, the HP Labs crew in Palo Alto included Rich Elder, Carl Taussig, and Terril Hurst.

The University of Colorado formed a program in 1998 to work with the storage industry, called the Colorado Center for Information Storage (CCIS), which is completely industry supported. More recently, researchers at Samsung Advanced Institute of Technology in Korea and the Data Storage Institute in Singapore have been prolific in publishing results related to disk drives. Of course, many contributions to drive products have largely gone undocumented, as the disk drive companies have elected to protect these as trade secrets.

Closing Remarks

The purpose of this article has been to trace the history of disk drive control, in the context of the larger disk drive problem itself. Where is the control of disk drives going? While the future is hard to predict, here are some trends that bear watching.

- Dual-stage actuators: Micromachined actuators solve many of the mechanical issues with disk drive actuation. Although their adoption on hard disks is generally considered inevitable, the competing solutions of stiffening and/or instrumenting the actuator may very well be used before second stages, as they can considerably extend the life of single-stage actuation.
- Stiffer, smaller, more aerodynamic disks and actuators are inevitable. Glass substrates should become standard on hard disks for desktops.
- Servowriting as it is now done will be eliminated, probably in one of two ways: self-servowriting or patterning disks. The patterning of the disks may be used for magnetic or optical position detection.
- Synchronous demodulation, a trivial way to improve bandwidth, should become standard.
- Fluid-bearing spindles may become standard, especially for the consumer appliance market.

All these trends will happen in an environment where the increases in bit density are flattening out, leaving servo systems in the critical path for raising areal density through increased track density. Although predictions of the shrinking of bit aspect ratios have persisted for the past ten years, these predictions are most likely to prove true in the next ten years. Far from being solved, the disk drive control problem should keep engineers busy for many years to come.

Acknowledgments

It is impossible to write an article of this type without considerable assistance. P.D. Mathur of Seagate, Dick Curran and Fred Hansen of Quantum/Maxtor, and Ed Grochowski of IBM provided very useful historical data on their companies' drives. Dick Henze and Chuck Morehouse (HP), as well as Bob Evans (Hutchinson Technology), provided useful technical data at key points in the writing. Rick Ehrlich (Quantum/Maxtor), Terril Hurst (HP), Lucy Pao (University of Colorado), Ho Seong Lee (Maxtor), and Masayoshi Tomizuka (UC-Berkeley) provided not only much useful information, but also many helpful suggestions for improving the article. Dick Oswald, Fred Kurzweil, Hal Ottesen, Mike Hatch, and Jim Porter all provided fascinating stories on the early days of disk drive control. Finally, the gathering of old and sometimes obscure references for this work would not have been possible without the tireless assistance of the Agilent Labs Library research staff, particularly Ron Rodrigues, Sandy Madison, and Florence Haas. To all of these people we owe a debt of gratitude.

References

- [1] L.D. Stevens, "Data storage on magnetic disks," in *Magnetic Recording: The First 100 Years*, E.D. Daniel, C.D. Mee, and M.H. Clark, Eds. Piscataway, NJ: IEEE Press, 1997, pp. 270-299.
- [2] G. Rostky, "Disk drives take eventful spin (history of disk drives) (industry trend or event)," *Electron. Eng. Times*, p. 22, July 13, 1998.
- [3] J.N. Porter, "Disk drives' evolution," in *Proc. 100th Anniversary Conf.: Magnetic Recording and Information Storage*, Dec. 14, 1998.
- [4] W. Messner and R. Ehrlich, "A tutorial on controls for disk drives," in *Proc. 2001 American Control Conf.*, Arlington, VA, June 2001, pp. 408-420.
- [5] D.Y. Abramovitch, "Magnetic and optical disk control: Parallels and contrasts," in *Proc. 2001 American Control Conf.*, Arlington, VA, June 2001, pp. 421-428.
- [6] R.L. Comstock and M.L. Workman, "Data storage on rigid disks," in *Magnetic Recording Handbook*, C.D. Mee and E.D. Daniel, Eds. New York: McGraw Hill, vol. 2, 1990, pp. 655-771.
- [7] C.M. Christensen, *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail*. Boston, MA: Harvard Business School Press, 1997.
- [8] T. Reiley, T. Albrecht, K. Kuroki, D. Albrecht, and M. Aoyagi, "Micro Drive—A pluggable one-inch disk drive for portable devices," in *Proc. 7th Biennial IEEE Int. Nonvolatile Memory Technology Conf.*, Albuquerque, NM, June 22-24, 1998, pp. 51-54.
- [9] J. Porter, private communication, Mar. 27, 2001.
- [10] J. Fegreus, "The sound and the fury: New drive firsts from Fujitsu, Seagate," *Data Storage*, vol. 7, pp. 9-10, Apr. 2000.
- [11] IBM Storage Technology Division, "IBM Ultrastar 36Z15 hard disk drive," On-Line White Paper TECHFAX #7016, IBM, San Jose, CA, 2001 [Online]. Available: <http://www.storage.ibm.com/hardsoft/diskrdrl/ultra/36z15data.htm>
- [12] J.C. Mallinson, *The Foundations of Magnetic Recording*, 2nd ed. San Diego, CA: Academic, 1993.
- [13] H.N. Bertram, *Theory of Magnetic Recording*. Cambridge, U.K.: Cambridge Univ. Press, 1994.
- [14] R.K. Oswald, private communication, Mar. 23, 2001.
- [15] D. Abramovitch, T. Hurst, and D. Henze, "An overview of the PES Pareto method for decomposing baseline noise sources in hard disk position error signals," *IEEE Trans. Magn.*, vol. 34, pp. 17-23, Jan. 1998.
- [16] N. Schirle and D.K. Lieu, "History and trends in the development of motorized spindles for hard disk drives," *IEEE Trans. Magn.*, vol. 32, pp. 1703-1708, May 1996.
- [17] G.F. Franklin, J.D. Powell, and M.L. Workman, *Digital Control of Dynamic Systems*, 2nd ed. Menlo Park, CA: Addison-Wesley, 1990.
- [18] R.K. Oswald, "Design of a disk file head-positioning servo," *IBM J. Res. Develop.*, vol. 18, pp. 506-512, Nov. 1974.
- [19] K. Akagi, T. Nakao, Y. Miyamura, T. Munemoto, and K. Mori, "High-density magnetic recording tracking method using a laser diode," *IEEE Trans. Magn.*, vol. 27, pp. 5301-5303, Nov. 1991.
- [20] H.S. Lee, "Implementation of adaptive feedforward cancellation algorithms for pre-embossed rigid magnetic (PERM) disks," *IEEE Trans. Magn.*, vol. 33, pp. 2419-2423, May 1997.
- [21] Maxtor Corporation, "108 117 Notes to Maxtor Unaudited Pro Forma Condensed Combined Financial Statements," SEC Filing Type 424B3, Maxtor Corporation, Mar. 7, 2001 [Online]. Available: http://10kwizard.ccbn.com/file_submis.asp?iacc=1400620&g=#
- [22] A.H. Sacks, "Position signal generation in magnetic disk drives," Ph.D. dissertation, Dept. Elec. Comput. Eng., Carnegie Mellon Univ., Pittsburgh, PA, Sept. 5, 1995.
- [23] D. Abramovitch, "Customizable coherent servo demodulation for disk drives," *IEEE/ASME Trans. Mechatron.*, vol. 3, pp. 184-193, Sept. 1998.
- [24] T. Hurst, private communication, 2001.
- [25] M.C. Stich, "Digital servo algorithm for disk actuator control," in *Proc. Conf. Applied Motion Control '87*, Minneapolis, MN, June 16-18, 1987, pp. 35-41.

- [26] R. Ehrlich, private communication, Mar. 28, 2001.
- [27] F. Kurzweil, private communication, Apr. 13, 2001.
- [28] H. Ottesen, private communication, Apr. 12, 2001.
- [29] W.W. Chiang, "Multirate state-space digital controller for sector servo systems," in *Proc. 29th IEEE Conf. Decision and Control*, vol. 3, Honolulu, HI, Dec. 5-7, 1990, pp. 1902-1907.
- [30] T. Hara and M. Tomizuka, "Multi-rate controller for hard disk drive with redesign of state estimator," in *Proc. 1998 American Control Conf.*, Philadelphia, PA, June 1998, pp. 3033-3037.
- [31] P.A. Weaver and R.M. Ehrlich, "The use of multirate notch filters in embedded-servo disk drives," in *Proc. 1995 American Control Conf.*, Seattle, WA, June 1995, pp. 4156-4160.
- [32] D. Abramovitch, "Rejecting rotational disturbances on small disk drives using rotational accelerometers," *Contr. Eng. Practice*, vol. 5, Nov. 1997.
- [33] M. Banther, Y. Huang, and W.C. Messner, "Optimal strain gauge placement for an instrumented disk drive suspension," in *Proc. 1998 American Control Conf.*, Philadelphia, PA, June 1998, pp. 3023-3027.
- [34] K.K. Chew and M. Tomizuka, "Digital control of repetitive errors in disk drive systems," *IEEE Contr. Syst. Mag.*, vol. 10, pp. 16-20, Jan. 1990.
- [35] C. Kempf, W. Messner, M. Tomizuka, and R. Horowitz, "Comparison of four discrete-time repetitive control algorithms," *IEEE Contr. Syst. Mag.*, vol. 13, pp. 48-54, Dec. 1993.
- [36] J.S. McAllister, "Disk flutter: Causes and potential cures," *Data Storage*, vol. 4, pp. 29-34, May/June 1997.
- [37] R.A. White, "Shock force compensating system," U.S. Patent 4,040,103, Aug. 1977.
- [38] D.B. Davies and M.D. Sidman, "Active compensation of shock, vibration, and wind-up in disk drives," in *Advances in Information Storage Systems, Vol. 5*, B. Bhushan, Ed. New York: ASME Press, 1993, pp. 5-20.
- [39] M. White and M. Tomizuka, "Increased disturbance rejection in magnetic disk drives by acceleration feedforward control," *Contr. Eng. Practice*, vol. 5, no. 6, pp. 741-751, 1997.
- [40] C. Hernden, "Vibration cancellation using rotational accelerometer feedforward in HDDs," *Data Storage*, vol. 7, pp. 22-28, Nov. 2000.
- [41] H.T. Ho, "Noise impact on servo TMR," in *Proc. 1997 American Control Conf.*, Albuquerque, NM, June 4-6, 1997, pp. 2906-2909.
- [42] R. Ehrlich and D. Curran, "Major HDD TMR sources and projected scaling with TPI," *IEEE Trans. Magn.*, vol. 35, pp. 885-891, Mar. 1999.
- [43] H.S. Lee and L. Guo, "Servo performance prediction for high capacity disk drives," in *Proc. 1998 American Control Conf.*, Philadelphia, PA, June 1998, pp. 3050-3054.
- [44] J. Humphrey, C. Chang, H. Li, and C. Schuler, "Unobstructed and obstructed rotating disk flows: A summary review relevant to information storage system," in *Advances in Information Storage Systems, Vol. 1*, B. Bhushan, Ed. New York: ASME Press, 1991, pp. 79-110.
- [45] P.D. Mathur, *Position Sensing and Control for Magnetic Disc Drives*. Ph.D. dissertation, Dept. Elec. Comp. Eng., Carnegie Mellon Univ., Pittsburgh, PA, May 10, 1999.
- [46] D. Cahalan and K. Chopra, "Effects of MR head track profile characteristics on servo performance," *IEEE Trans. Magn.*, vol. 30, pp. 4203-4205, Nov. 1994.
- [47] E. Champion and H.N. Bertram, "Hysteresis and baseline shift in permanent magnetic stabilized (MR/SAL) heads," *IEEE Trans. Magn.*, vol. 32, pp. 13-18, Jan. 1996.
- [48] J. Adley, S. Stole, M. Jones, and R. Fox, "Actuator designs for high areal density with alternate materials," *Data Storage*, vol. 7, pp. 41-46, Mar. 2000.
- [49] M. Hatch, private communication, Apr. 25, 2001.
- [50] D. Abramovitch, F. Wang, T. Hurst, and G. Franklin, "Disk drive pivot nonlinearity modeling Parts I & II," in *Proc. 1994 American Control Conf.*, Baltimore, MD, June 1994, pp. 2600-2607.
- [51] P.R. Dahl, "Solid friction damping of mechanical vibrations," *AIAA J.*, vol. 14, pp. 1675-1682, Dec. 1976.
- [52] J. Ishikawa and M. Tomizuka, "Pivot friction compensation using an accelerometer and a disturbance observer for hard disk drives," *IEEE/ASME Trans. Mechatron.*, vol. 3, pp. 194-201, Sept. 1998.
- [53] S. Aggarwal, D. Horsley, R. Horowitz, and A. Pisano, "Micro-actuators for high density disk drives," in *Proc. 1997 American Control Conf.*, Albuquerque, NM, June 1997, pp. 3979-3984.
- [54] B. McInerney, "Surveying micro-positioning technology for advanced disk drives," *Data Storage*, vol. 7, pp. 24-27, Aug. 2000.
- [55] F. Stevens and J. DeLillis, "Keeping heads on track with dual-stage actuators," *Data Storage*, vol. 7, pp. 39-44, June 2000.
- [56] R.B. Evans and T.A. Krinke, "Head suspension with resonance feedback transducer," U.S. Patent 5,862,015, Jan. 19, 1999.
- [57] S.J. Schroeck, W.C. Messner, and R.J. McNab, "On compensator design for linear time-invariant dual-input single-output systems," *IEEE/ASME Trans. Mechatron.*, vol. 6, pp. 50-57, Mar. 2001.
- [58] L.Y. Pao and G.F. Franklin, "Proximate time-optimal control of third-order servomechanisms," *IEEE Trans. Automat. Contr.*, vol. 38, pp. 560-580, Apr. 1993.
- [59] N.C. Singer and W.P. Seering, "Design and comparison of command shaping methods for controlling residual vibration," in *Proc. 1989 IEEE Int. Conf. Robotics and Automation*, Scottsdale, AZ, May 14-19, 1989, pp. 888-893.
- [60] C. La-orpacharapan and L.Y. Pao, "Control of flexible structures with a projected phase-plane approach," in *Proc. 2001 American Control Conf.*, Arlington, VA, June 2001, pp. 3817-3823.

Daniel Abramovitch is a Senior Member of the IEEE and is with the Communications and Optics Research Lab at Agilent Laboratories, Palo Alto, CA.

Gene Franklin is a Fellow of the IEEE and Professor Emeritus at the Information Systems Lab in the Electrical Engineering Department of Stanford University.