DISK DRIVE CONTROL: THE EARLY YEARS

Daniel Abramovitch * **and Gene Franklin** **

* *Agilent Laboratories, Communications and Optics Research Lab, 3500 Deer Creek Road, M/S: 25U-9, Palo Alto, CA 94304, Phone: (650) 485-3806, FAX: (650) 485-4080, E-mail: danny@labs.agilent.com*
** *Stanford University, Information Systems Lab, 252 Packard Electrical Engineering Building, 350 Serra Mall, Stanford, CA 94305, Phone: (650) 723-4837, E-mail: franklin@isl.stanford.edu*

Abstract: One of the persistently exciting control applications is that of disk drive servos. From the start in the early 1950s to the massive capacity commodity drives of the early 2000s, the problem of accessing data on rotating disk media has provided a wealth of control challenges to be solved. This survey paper traces the early history of disk drive control from the first disk drive in 1956 to the first commercial drive with Magneto-Resistive heads in 1990. Rather than the approach used in (Abramovitch and Franklin, 2002*a*) in which the histories of the components were outlined first, we will focus on the feedback loop itself in those early days. The paper will survey the different areas of the disk drive control problem and how they evolved.

Keywords: Disk drive control, early history, servomechanisms, computer hardware

1. INTRODUCTION

One of the persistently exciting control applications is that of disk drive servos. From the start in the early 1950s to the massive capacity commodity drives of the early 2000s, the problem of accessing data on rotating disk media has provided a wealth of control challenges to be solved. This shows no signs of abating as storage densities, capacity and transfer rates keep rising while costs and size keep dropping (Porter, 1998). Although a host of new technologies (Magnetic RAM, Atomic Resolution Storage, optical storage) are persistently 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 paper is to provide a history of control in the early days of disk drives. While this subject can include both flexible and optical drives, this paper will focus on rigid magnetic disks – often called hard disks. For a short general history of disk drives culled from several recent excellent sources (Stevens, 1997; Rostky, 1998; Porter, 1996; Porter, 1998) can be found in (Abramovitch and Franklin, 2002*a*). In the latter we presented an outline of the history of disk drive control. In this paper, we will focus on the early years of the drive control problem. The question of what constitutes "early" is a natural one at this point. We have chosen 1991, the year that IBM came out with the Corsair drive (Figure 1),

Fig. 1. IBM Corsair Disk Drive *(Courtesy IBM.)*

the first commercial drive based on the Magneto-Resistive (MR) head (Daniel *et al.*, 1997).[1]

This drive is chosen as a breakpoint not due to its technical or commercial success (which was considered quite low a the time), but because as the first commercial MR head drive it for-shadowed the future of disk drives. While many drive companies had MR head drives in the lab, the leap to a commercial MR head drive was at once necessary and risky. Seagate would not produce a MR head drive until late 1996. HP would exit the drive business trying to produce their first MR head drives. However, as MR heads changed the slope of areal density improvement from 30% per year to 60% per year, they dramatically changed what the disk drive market by the year 2001. Thus, the first MR head commercial drive is our break point.

2. A WALK AROUND THE LOOP

A schematic block diagram of a disk drive control loop is shown in Figure 2. The disk loop starts with the disks stack assembly diagrammed in Figure 3: a stacking of magnetic disks on a spindle with an internal spindle motor. The magnetic media contains data in concentric circular tracks on both sides of the media.

Modern disk drives read the relative position of the head to the track directly from the disk media. Virtually all of today's drives use a method called sectored servo[2], 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 are 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 Figure 4 – and this signal must be demodulated. From there, the data is digitized and fed into a digital processor, most

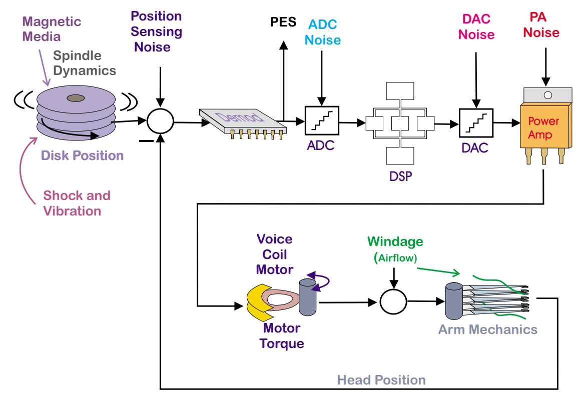

Fig. 2. Generalized view of track following model.

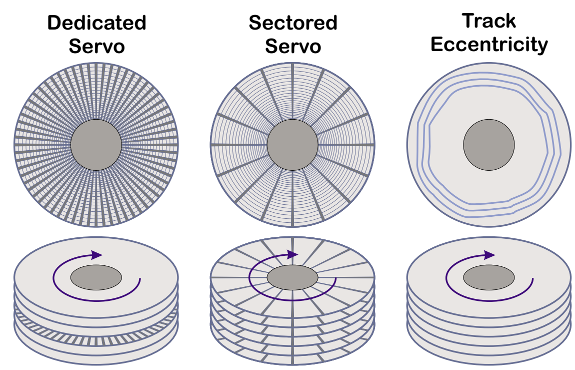

Fig. 3. Disk stacks for dedicated and sectored servo systems. The dedicated servo on the left allows the entire surface to be covered with position information, thus enabling a higher sample rate. The sectored servo in the center multiplexes the position information with the user data, which puts a practical limit on the sample rate. However, it does co-locate the position information with the user data making the servo system largely immune to mechanical offsets in the actuator assembly. Both types of formatting are subject to track eccentricity, shown on the right, where the tracks can be non-circular and not properly centered.

---

[1] The conference version of this paper (Abramovitch and Franklin, 2002*b*) incorrectly identified the IBM Redwing, which came out in 1990, as the first MR Head disk drive. This, in fact, was the first drive to use a Partial Response, Maximum Likelihood (PRML) read channel.

[2] Also commonly called embedded servo. However, there are several ways to embed servo information, so the terms are technically not synonyms.

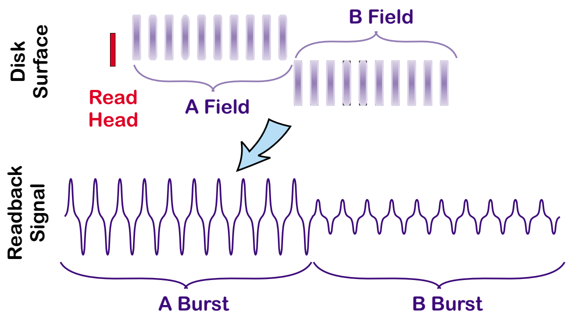

Fig. 4. A simplified view of split field amplitude encoded servo fields on a hard disk and the resulting readback signals. In split field encoding the A fields and B fields are separated down the track. Amplitude/area estimates of the A and B fields, $\bar{A}$ and $\bar{B}$, are computed separately and subtracted from each other.

often a DSP, for implementation of the control law. 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, shown in Figure 5d. The rotary voice coil 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 is the magnetic read/write element. In current disk drives, this is a pair of heads. 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.

3. ACTUATORS

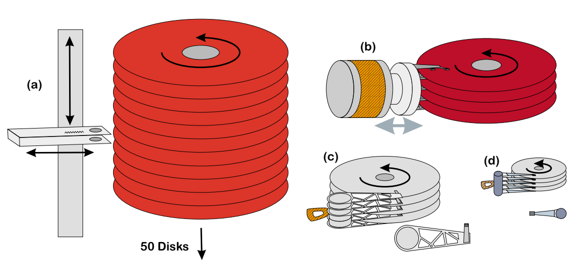

Fig. 5. Generalized view of the evolution of actuators. Note the transition from the RAMAC (a) to the comb structure which had one head per surface (b). The next transition is from linear to rotary (c), then from larger rotary to smaller stiffer rotary actuators (d).

The original RAMAC (Figure 6) actuator used aircraft cable and pulleys, as did the system described by Hagopian (Hagopian, 1961), 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. Hydraulic actuators offered far better seek performance than magnetics, but could not provide a track following servo and also had a problem with contamination due to oil leaks (Oswald, 2001).

Fig. 6. IBM RAMAC *(Courtesy IBM.)*

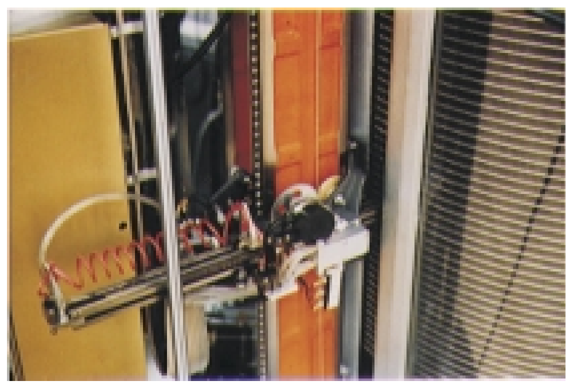

Fig. 7. RAMAC Actuator *(Courtesy IBM.)*

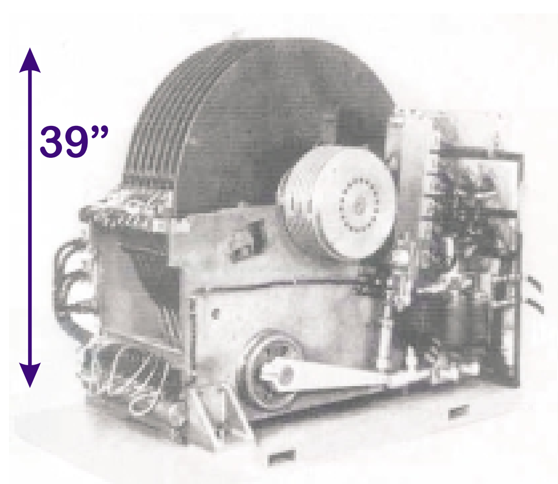

Fig. 8. Bryant Computer Systems' 39" Disk Drive *(Courtesy* Electronic Design*)*

The original RAMAC actuator had a single pair of heads that moved both vertically and radially to access data (Figures 5a and 7). The heads floated on a static air bearing which required a separate pumping mechanism. The invention of the self actuating air bearing allowed there to be one slider per surface and led to the comb actuator. By 1965, these were driven by a linear voice coil motor (Figure 5b).

The next major step was a move to rotary actuators. The earliest of these belonged to the drive made by Bryant Computer Systems (Figure 8). Designed for factory environments, this drive was distinguished by its 39” disk platters. While the drive was not portable, it did tend to “walk” across the floor during seeks. Rivals joked that Bryant solved this problem with a “short cord solution”: when the drive went far enough, it pulled the power cord out of the wall and shut itself off. Actually, Bryant alleviated the problem by balancing the actuator and bolting the drive to the factory floor (Rostky, 1998).

The actual industry move toward rotary actuators started with the one designed at IBM’s Winchester Labs (in Winchester, England) (Porter, 2000). The IBMers in England were associated with IBM’s Rochester disk drive group – working on drives for minicomputers and smaller systems, so they had a different approach than the folks in San Jose. The early Rochester drives (Spartan and the IBM 9332) had linear actuators. However, people at IBM’s Fujisawa operation were looking at rotary actuators. When they put the rotary actuator drives on shaker tables, they found that those drives had much better immunity to linear shock and vibration.

The rotary actuators for 8” drives were relatively large and complicated mechanical truss structures (Figure 5c) . They were large due to the size of the disk, but the trusses minimized the inertia of the actuator. The suspension was turned sideways on these as to most closely mimic the motion of the linear actuator. As drives shrank once more, the actuator pivot was put in one corner of the enclosure. The sideways turn in the suspension was removed so that the actuator and suspension were in a single line being dragged over the disk (Figure 5d). 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 Section 5.6.

4. SERVO SIGNALS

The RAMAC had a servo system to move the actuator vertically and radially. The positions were “detected” 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 (Stevens, 1997). These were open loop until 1971 when IBM introduced the first disk drive that closed the loop with position information read from the loop (Oswald, 1971; Oswald, 1974), although the idea was studied much earlier (Hoagland, 1961).

A variant of the off-disk position information was found in early drives made by Quantum which used an optical encoder (stuck to the side of the VCM) as a servo mechanism (up to track-densities of around 2100 TPI). They worked well, and did not require a servowriter. They of course suffered all of the disadvantages of dedicated servo systems, plus the disadvantage of having a position reference that was completely off the disk (Ehrlich, 2001).

Modern disk drives read the relative position of the head 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 side of Figure 3. Embedded servo time multiplexes the servo information with the user data on each surface as shown in the center diagram of Figure 3. Dedicated servos have the advantage 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 have more susceptibility to thermal offsets than embedded servos. In order 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 Figure 3. Embedded or sectored servos, as shown in the center diagram of Figure 3, co-locate 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 (co-located control).

A variety of position encoding methods have been used to encode the servo position including amplitude, phase (Boutaghou *et al.*, 1994) , frequency, and null (Wilson, 1994; Sacks, 1995; Sacks *et al.*, 1996) encoding. Figure 4 shows a diagram of amplitude and null encoded servo patterns. The servo signals must be demodulated into the baseband to be useful. Over the course of disk drives’ history, the methods have trailed those of the data channels. Peak detection channels have been used to detect position information up until very recently (Silicon Systems Staff, 1994).

5. THE CONTROL PROBLEM

With all the above background plus that in (Abramovitch and Franklin, 2002*a*), 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. Limits on the ability to improve the drive's track density would prompt changes in the magnetics or mechanics of the drive. Through all of this, the bit aspect ratio, that is the track width of a bit compared to the down the track length of a bit, has stayed centered around 20 to 1 (Abramovitch and Franklin, 2002*a*).

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 came out 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 (Oswald, 1974). It was common practice at some disk drive companies for an engineer to go to their first day at work only to find a copy of Oswald's paper on their desk (Hurst, 2001).

The 3340 Winchester drive put together many ideas that had been floating around for a while. The idea of reading a signal from the magnetic medium and using this for fine positioning appears to be discussed in a paper by R. Tickell at the 1966 Joint American Conference (Tickell, 1966) and in an earlier paper by Al Hoagland (now at Santa Clara University) in the IBM Journal (Hoagland, 1961). Tickell worked at Sperry Rand in the Univac Division which had put off coming out with a disk drive product in 1956 in favor of magnetic drum storage (Abramovitch and Franklin, 2002*a*). Thus it was that Tickell's paper closes the loop using signals read back from a magnetic drum – not a disk. Still, the paper introduces some concepts that would be used by Merlin and future disk drives. Rather than use each of the data heads to read the track position, this system uses a single pair of heads to read a dedicated position track – the first dedicated servo head. A pair of heads is used to get common mode rejection of media variations. Furthermore, the positioning is segmented between track numbers (detected digitally) and the sub-track position (detected with the servo head).

Hoagland's paper (Hoagland, 1961) discusses a team project to build a laboratory setup that did fine positioning of a hydraulic actuator on a magnetic disk by reading magnetically encoded position fields off the interior area of the disk. It also introduces a set of concepts that would be used in the IBM 3340. A set of data heads is ganged together with a servo head and a sector head to determine the fine position. Multiple position pulses are averaged together to provide some immunity from noisy or bad pulses. The concept of using a dual element head, one to read data and one to write data shows up here as well, allowing the write wide, read narrow procedure that is fairly common in today's drives that use MR heads.

The actual position encoding method in the 3340 was a tri-bit encoding method proposed by Mueller (Mueller, 1972). It won out over the method used in the IBM 3330 Merlin, in which alternating position information tracks had magnetization patterns with opposite polarity (Santana, 1970). IBM had also been working on average seek time studies (Hertrich, 1965) and on the use of bang-bang control for seeks (Brown and Ma, 1968). However, in Oswald's paper a modification of the bang-bang scheme, which would later be codified as a Proximate Time Optimal Servo (PTOS) (Workman, 1987*b*) was first discussed. With all of these coming together, one can see the architecture of future disk drive control systems in Oswald's paper.

There is an old saying in the computer industry which dates back to the days when IBM had a virtual monopoly on computing. In those days it was said that if IBM published something, it meant they had decided not to use it[3]. 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 with 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 this is complete. There is a lot 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.

[3] Thanks go to Dennis Bernstein for triggering this memory.

Another phenomenon that keeps the servo work secret is that by and large disk drive companies have cross license agreements on all their patents. As their competitors have ready access to their patents, many disk drive companies shy away from patents in favor of keeping information as a trade secret. Trade secrets don't get published while patents are public information.

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. The general feeling is that a new piece of technology gets added slowly as the risk of that technology goes down and when some change in a component around the loop changes the balance of the PES composition.

With this caveat in mind, the rest of this section will delve into specific concentrations of work on hard disk servos.

5.1 Analog, Sampled Data and Digital Control

Fig. 9. DEC RL01 (on left) and RK07 (on right) Disk Drives *(Courtesy Ákos Varga)*

While the RAMAC used tube electronics (Stevens, 1997), drives progressed rather quickly to transistors and integrated circuits. From the beginning, the drives had digital circuitry to relay the data to and from the computer. 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 had some low end drives with 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 (Ottesen, 1988; Stich, 1987) and mainframes (Franklin *et al.*, 1990).

It is worth understanding that the very nature of encoding position on a disk drive implies a sampling process. Thus, even when the control laws were implemented using analog electronics, the system was a sampled data system. For dedicated servo drives, this sample rate was often high enough to masquerade as continuous position information passed through a low pass filter. However, sectored servo pushes sample rates lower, forcing the control system designer to deal with the sampling of their system. HP had drives of this nature through the mid 1980s. Some of the quirky aspects of this type of a system are touched on in a paper on frequency response function measurements (Ehrlich *et al.*, 1989). It was not until the late 1980s that HP moved to fully digitally controlled drives, in large part due to the interaction between Rick Ehrlich (then at HP Labs) and Vernon Knowles at HP's Disk Memory Division.

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. This was particularly evident in different levels of openness with the rest of the industry between San Jose and Rochester. Rochester, Minnesota is most famous for the Mayo Clinic, but it was also the site of a major IBM hard disk facility. In its heyday, this site made drives for the minicomputer market. Rochester, being tied to the minicomputer business which competed with DEC, had a relatively open attitude. San Jose, involved with the mainframe industry in which IBM had a virtual monopoly, tended to be more closed.

Rochester was competing in the small form factor market, so they were more outward looking and open than San Jose. They would buy cameras and consumer electronics and take them apart to see how those products did packaging. They were always looking for ways to get things into smaller form factors. It's likely that if the Rochester folks had been in San Jose, they wouldn't have been able to do what they were doing. Their distance from San Jose allowed them to be the "rebels in a cornfield" (Ottesen, 2001). This is an interesting irony, in that it was San Jose's distance from IBM headquarters in Armonk that had given them the freedom to work on the RAMAC almost 3 decades earlier. By the mid-1980s San Jose was the entrenched big brother; Rochester the strident kid brother. Thus, Rochester management was somewhat antagonistic to San Jose management, whereas the latter tended to ignore the upstarts. Thus, IBM San Jose was producing 14" disk drives late into the 1980s, long after the rest of the industry had moved to smaller form factors.

When Rochester folks proposed digital control, they got laughed out of the room by San Jose folks. Thus, it was the gang at Rochester that came up with IBM's first digital control drive, code named Spartan. This was an 8" drive intended for mini-computers. They had an advanced development lab in Hursley, England and a development lab next door in Winchester, England. It was the English team that came up with the sectored servo approach at IBM. This is currently used in virtually all disk drives. While the Winchester team may have lead IBM into the sectored servo domain, they were not the first in the industry. DEC came out with the 125 TPI RL01 disk drive in 1978, which was the first mass produced full sectored (embedded) servo drive on the market (see the left side of Figure 9).

Another interesting start for digital control was relayed by Fred Kurzweil. Fred, who has the distinction of being Gene Franklin's first graduate student at Stanford, graduated in 1959 and went 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 in June of 1982. Upon arriving at Maxtor, 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. Folks around the industry listened to Shugart eventually giving Maxtor a 2 year head start in this technology.

Arriving at Maxtor in October of 1982, Fred was one of the few people there who knew all about electric motors. Putting the motor in the spindle hub raised some serious issues. The motor consumed 10 watts of power, so being inside the hub, it tended to heat up. This meant that conditions were different at a cold start when the grease around the spindle ball bearings was cold than once the hub heated up and the grease flowed more freely. Fred solved this problem with a 10 cent microprocessor and some simple digital control. In order to maintain the spindle speed to 1 part in a million, he used simple adaptation: he would measure the speed over a revolution and perform a very low order correction so that as temperature affected the grease in the bearing (it started running faster as things got warmer) the speed stayed constant. Another issue was that the motor was a 3 phase induction motor. They had to get the motor up to speed before they could get a feedback signal from it to close the loop. They ran into stiction problems at start up. Often the motor didn't start and they had to dither the input to break it free (Kurzweil, 2001).

All these neat control ideas were tried on the spindle controller. However, this microprocessor was not suitable for high speed digital control and thus the rest of the drive used analog control loops.

However, the first use of digital signal processing in a disk drive was probably in the DEC RC25 fixed-plus-removable drive servo system. Much like modern optical drives (Abramovitch, 2001), the removable disk pack lead to inexact centering of the disks on the spindle. This in turn leads to large harmonic components in the PES, described further in Section 5.4. To address this problem, Mike Sidman developed the Adaptive Runout Correction System (ARCS) at DEC/Maynard in 1978, which used digital circular convolution to eliminate repeatable runout. Because the algorithm originally ran on an Intel 8085 microprocessor, which didn't have hardware multiply, Sidman hand coded a Booth's multiplier in 8085 assembly code. ARCS simultaneously corrected any desired number of runout harmonics and worked so well that it was used in DEC's removable and fixed media drives after that point (Sidman, 1985; Sidman, 1991).

5.2 State Space

With the emergence of DSPs and digital control on hard disks, the possibility of doing state space control became more real. Early disk drive digital control systems used classical design methods. The first use of state space control seems to have been during seek mode. There are a couple of reasons for this. First of all, the number of operations for a state space controller are typically larger than those for a classical controller on the same order problem. Early DSPs were hard pressed to do all the extra operations 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 in seek mode when the signal was large, i.e. during seek.

Classical design techniques lend themselves to being able to directly use the measured frequency response functions, without having to compute a parametric model. Parametric models that closely match these frequency response functions often have on the order of 30 modes (Hanselmann and Engelke, 1988). Thus, if a designer wants to use state space techniques, they must either design with a high order model and controller, or go through an often lengthy model reduction step to capture the essence of the dynamics.

The use of state space control to make use of the back EMF sensor during seek while using a classical design in track following was used by Quantum into the late 1980s and at HP until they got out of the drive business in 1996 (Knowles, 1991).

On the other hand, IBM Rochester got into the state space paradigm early on. Rochester started on digital control in 1980. They had gotten a microprocessor from Intel that was just released. A military processor that had DSP capability (the Intel 8196). Mike Stich, a servo engineer out of Rochester, took the digital controls class at Stanford from Gene Franklin through fellow professor Marty Hellman's company. Franklin and Powell's first digital control book, the "little blue book" (Franklin and Powell, 1980), had just come out. After Stich went back to Rochester, he contacted Franklin to see if he was interested in teaching at Rochester. Ottesen had a copy of the little blue book and handled the logistics of having Franklin come to IBM Rochester to do some consulting and teach some digital controls classes.

The IBM 9332 out of IBM's Rochester operation was not only the drive with the first 1-7 RLL code (IBM Storage Technology Division, 2001), but also made use of a digital state space controller that even did on-line parameter adaptation (Ottesen, 1988; Stich, 1987; Stich *et al.*, 1987). IBM Rochester, had relatively little pain going into state space. There was nobody around that had any experience with digital control so they were on their own to experiment as they thought best. Back then they used APL, a cryptic computer language at best, to model their systems (Ottesen, 2001).

The case history on disk drive control written by Mike Workman (of IBM San Jose) for the digital control book by Franklin, Powell, and Workman (Franklin *et al.*, 1990) makes use of state space control as well, indicating that IBM San Jose had fully embraced the state space approach as well, but after Rochester did. In fact, Hal Ottesen, a long time servo engineer at IBM Rochester recalls that when the folks from Rochester first proposed digital control to the folks at San Jose, they were laughed out of the room (Ottesen, 2001). IBM San Jose's move to digital and state space control was largely led by Mike Workman, who had been taking classes at Stanford and gotten "the religion". It was largely through the influence of Mike Workman that IBM San Jose moved into the digital control and state space world. Their first digital control drive was the IBM 3380K.

Hal Ottesen would travel around IBM sites teaching internal classes on digital and state space control of disk drives. His hand written notes evolved over time, but have never been published.

According to Rick Ehrlich, Quantum's first drive to use state-space for seeking and tracking, and with embedded servo was the high-end Enterprise drive which went into mass production in early 1992. The desktop drives went to single-loop state-space control a couple of years later (when the microprocessor was fast enough that they didn't have to drop to a simple servo in ontrack mode to make time for the I/O firmware) (Ehrlich, 2001).

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

5.3 Sample Rates

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 multi-rate 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 multi-rate control ever since he was at Hewlett-Packard (Ehrlich *et al.*, 1989). Carl Taussig had continued this work, coming up with unpublished results that were quite similar to those reported by W. W. Chiang (out of IBM Almaden Research Center in San Jose) in 1990 (Chiang, 1990). It turns out that Rick Ehrlich had continued his multi-rate work at Quantum and was getting similar results. Basically, if the Position Error Signal (PES) sample rate was relatively low – say under 8 times the open loop crossover – then keeping everything else constant and raising the output sample rate by a factor of 3 or 4 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.

5.4 Repetitive and Spectral Disturbances

One of the concepts to hit disk drive control systems at the end of the 1980s was that of repetitive control to cancel the effects of the spindle eccentricity. However, the earliest work on this was by

Fig. 10. DEC RK05 Disk Pack *(Courtesy John Holden)*

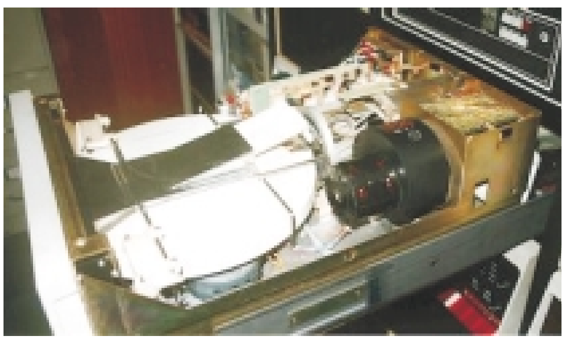

Fig. 11. DEC RK05 Disk Pack in Drive *(Courtesy John Holden)*

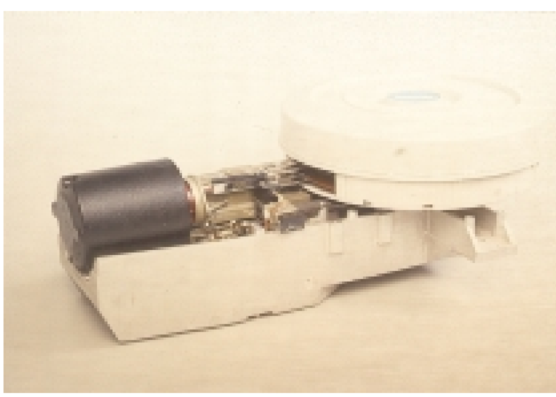

Fig. 12. DEC RK07 Disk Drive *(Courtesy Mike Sidman)*

Mike Sidman at Digital Equipment Corporation (DEC). DEC sold minicomputers and so they had a considerable effort in removable disk packs, such as the RK05 shown in Figures 10 and 11 and the RK07 shown in Figures 9 and 12. The 385 TPI RK07 was a dual 14” platter removable cartridge drive that motivated Sidman’s work on what he called the Adaptive Runout Correction System (ARCS). Mike was the RK07 servo and project engineer and had to deal with cartridge seating and interchangeability issues. He used the RK07 as the breadboard for the first operational ARCS, though ARCS was first released in a production machine in the RC25 fixed-plus-removable disk drive (Sidman, 1985).

This hit broad use after work in Masayoshi Tomizuka’s group in the Mechanical Engineering Department at UC Berkeley. 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 (Tomizuka *et al.*, 1988; Chew and Tomizuka, 1989; Kempf *et al.*, 1993; Chew and Tomizuka, 1990). As graduate students branched out from Berkeley, especially it seems to professorships at CMU (e.g. Marc Bodson who later moved to Utah and Bill Messner), these studies also included adaptive feedforward harmonic cancellers, which were shown to have some equivalence with repetitive controllers. As the 1990s would progress, the use of harmonic correctors would become standard throughout the industry.

5.5 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 in this was the continuing drop in the cost of accelerometers to the point at which one could reasonably consider them as an option for disk drives.

However, the earliest examples of accelerometer control in a disk drive go all the way back to the 1970s. White used an accelerometer to sense shock and then minimize the probability of the heads slapping against the magnetic media by 1) increasing the air pressure within the drive to stiffen the air bearing and 2) simply unload the heads from the disk surface (White, 1977). There was also some work by Robert Smith at Seagate (Smith, 1993).

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 speaking, the use of accelerometer feedforward dramatically improves the disturbance rejection capabilities of hard disks. Cost and reliability issues for the accelerometers themselves limited this practice to the lab through the early 1990s.

An exception to this occurred in the late 1980s. Hewlett-Packard was making use of a center of percussion rotary actuator as opposed to the more common balanced actuator. This left the drive susceptible to translational shocks and so two of the servo engineers, Vern Knowles and Mitch Hanks made use of a linear accelerometer to generate a feedforward compensation signal (Knowles and Hanks, 1987). However, each of these had to be calibrated to the disk in the factory, thus raising the manufacturing cost.

Davies, an MIT graduate student, and Mike Sidman, an engineer at Digital Equipment Corporation's Advanced Storage Lab in Colorado Springs, formulated conditions by which an accelerometer could perfectly cancel external and internal disturbances (Davies and Sidman, 1991; Davies and Sidman, 1993).

5.6 Friction

Friction is an issue whenever two surfaces in contact move relative to each other. 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. The issue becomes more noticeable as the actuator inertia drops and thus is more of an issue for small disk drives.

The story goes that friction in disk drives was discovered accidentally when a lab technician was running a swept sine measurement. Apparently, the technician accidentally set the lower sweep frequency to 10 Hz instead of 100 Hz. The resulting frequency response function did not look like anything described by their existing models and the study of disk drive friction was on.

Starting in the early 1990s, there were several publicized efforts to analyze and mitigate friction. Work was being done at Quantum in 1991 to characterize the nature of the ball bearing actuator pivots of their 3.5" drives. 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 an early National Storage Industry Consortium (NSIC) meeting in La Jolla in February of 1992. Hewlett-Packard would struggle with friction issues on the KittyHawk project, a 1.3" drive (Abramovitch *et al.*, 1994; Wang *et al.*, 1994). Friction mitigation work would continue sporadically around the industry through the 1990s.

5.7 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 this are:

- Lowered inertia for higher bandwidth actuation.
- Lowered energy power requirements for high bandwidth.
- Mitigation of pivot friction effects.
- Putting high bandwidth actuation at the end of the actuator beyond the effect of the suspension resonances.

However, most of these projects would not get going until the middle of the 1990s. By the year 2000, there would be predictions of universal acceptance of dual stage actuators even though none had shown up in any products (Stevens and DeLillis, 2000). The burgeoning field of micromachining would be the chief technology for most of these designs.

5.8 Seeks

The first drive that closed the loop based on position information read from the disk was IBM's 3330 Merlin drive. The Merlin drive attempted use of a bang-bang controller. However, the control system was superseded by that of the IBM 3340 Winchester, described in the seminal (but short paper) by Dick Oswald (Oswald, 1974). In this, he describes not only the track following operation, but also a near time-optimal seek mechanism. It turns out that IBM had been working on this for a while (Brown and Ma, 1968), and this technique would become standard practice in the industry. However, it was not until Mike Workman (an IBM engineer turned manager) wrote his doctoral thesis at Stanford while in Gene Franklin's group that this method would be codified as Proximate Time Optimal Servomechanism (PTOS) (Workman, 1987*b*; Workman, 1987*a*; Workman *et al.*, 1987*a*; Workman *et al.*, 1987*b*). The adaptive version was named APTOS, in part because it reminded his 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 (Pao and Franklin, 1990; Pao and Franklin, 1992; Pao and Franklin, 1993; Pao, 1994; Pao and Franklin, 1994).

One of the issues with seeks is that the reference positions given to the control system often excite the flexible modes of the actuator, even in closed-loop. This residual vibration takes a long time to damp out and effectively lengthens the seek time, since data cannot be read or written until the head is settled to within some small fraction of a track (typically under 10%). Neil Singer, a graduate student at MIT under Warren Seering, developed a method called command input shaping, for use in flexible robotics (Singer and Seering, 1988; Singer and Seering, 1989). 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. This results in a response with slightly slower rise times, but much faster settling times and was applied to NASAs space shuttle manipulator arm. In any event, it turned out that Neil Singer was a friend of Carl Taussig, who was at that time working on disk drive research at HP Labs. A collaborative effort ensued to apply this technique to disk drive manipulators.

5.9 Mode Switching Control

One of the aspects common to virtually all disk drive control systems is the notion of mode switching control, where the control system switches from one mode of operation to another. As a drive initiates a seek, it switches in the seek control algorithm. Typically these mimic bang-bang control at least during the acceleration phase. 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. A salient feature of then DEC disk drive servo engineer Mike Sidman's doctoral thesis was a smooth transfer between seek, settle, and track following mode (Sidman, 1986).

While the seek algorithms were codified by Mike Workman in his original work on PTOS (Workman, 1987*a*), it seems that even today, most disk drive servo systems still go through at least 3 stages of control. HP's disk drives of the early 1990s had 5 modes during seek, then a flare, gross settle, and settle mode, before moving to track following mode (Knowles, 1991).

The early IBM digital control drives mimicked the mode switching of analog controllers. The seek control followed a velocity profile. Settle mode used fairly stable poles. Once track following mode was switched in, an extra integrator was added to the control law computation. Furthermore, computational delay from position measurement to control output was modeled. Finally, they waited longer in settle mode if the drive was about to write data than if it was going to read data (Ottesen, 2001).

6. THE 1990S AND BEYOND

The 1990s would see several major trends. The use of MR heads would increase through the decade to become standard. The heads themselves would go through several generations to the present day Giant Magneto-Resistive Heads (GMR).

Drives would get smaller as areal densities increased at 60% a year resulting in small drives with massive capacities. Unfortunately for the industry, hard disks would become commodity items, limiting profits of the companies.

DSPs would become standard in hard disk servo systems. Servo system would move to state space. Harmonic correctors would become standard practice. Spindle speeds would rise, pushing bandwidths and mechanics. This in turn would put pressure on improving the servo signals (Abramovitch and Franklin, 2002*a*).

In the 1990s, most drive companies would move their production offshore. Most of these chased cheap labor in Southeast Asia. These moves did not hold off the consolidation in the industry. DEC would sell off what was left of their disk drive business to Quantum, who would sell off their disk business to Maxtor. Conner Peripherals and Imprimis would be absorbed into Seagate. IBM would shut down its Rochester operations. HP would exit the business completely. By the spring of 2002, IBM would announce the sale of their disk drive operations into a joint venture, mostly owned by Hitachi (Diaz, 2002).

ACKNOWLEDGMENTS

It is impossible to write such a paper as this without soliciting and receiving a considerable amount of assistance. P.D. Mathur of Seagate, Dick Curran and Fred Hansen of Quantum/Maxtor, and Ed Grochowski of IBM for very useful historical data on their companies drives. Mike Sidman (DEC) provided historical information on Digital Equipment Corporation's drive efforts. 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 (U. of Colorado), Ho Seong Lee (Maxtor), and Masayoshi Tomizuka (UC Berkeley) provided not only lots of useful information, but also many helpful suggestions on improving the paper. Dick Oswald, Fred Kurzweil, Hal Ottesen, 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

and Sandy Madison. To all of these people, we owe a debt of gratitude.

REFERENCES

Abramovitch, Daniel, Feei Wang and Gene Franklin (1994). Disk drive pivot nonlinearity modeling Part I: Frequency Domain. In: *Proceedings of the 1994 American Control Conference*. AACC. IEEE. Baltimore, MD. pp. 2600–2603.

Abramovitch, Daniel Y. (2001). Magnetic and optical disk control: Parallels and contrasts. On-line full draft of paper. Agilent Labs. Palo Alto, CA 94304. http://www.labs.agilent.com/-personal/Danny_Abramovitch/-pubs/hd_vs_od_servo.pdf.

Abramovitch, Daniel Y. and Gene F. Franklin (2002*a*). A brief history of disk drive control. *IEEE Control Systems Magazine* **22**(3), 28–42.

Abramovitch, Daniel Y. and Gene F. Franklin (2002*b*). Disk drive control: The early years. In: *Proceedings of the 2002 IFAC World Congress*. IFAC. IEEE. Barcelona, ES.

Boutaghou, Zine-Eddine, Dana H. Brown, Kevin J. Erickson and Richard Greenberg (1994). Digital servo signal demodulation method and apparatus utilizing a partial-response maximum-likelihood (PRML) channel in a disk file. United States Patent 5,343,340. International Business Machines Corporation. Armonk, NY USA.

Brown, C. J. and J. T. Ma (1968). Time-optimal control of a moving-coil linear actuator. *IBM Journal of Research and Development* **12**, 372–379.

Chew, Kok-Kia and Masayoshi Tomizuka (1989). Digital control of repetitive errors in disk drive systems. In: *Proceedings of the 1989 American Control Conference*. AACC. IEEE. Pittsburg, PA. pp. 540–548.

Chew, Kok Kia and Masayoshi Tomizuka (1990). Digital control of repetitive errors in disk drive systems. *IEEE Control Systems Magazine* **10**(1), 16–20.

Chiang, W. W. (1990). Multirate state-space digital controller for sector servo systems. In: *Proceedings of the 29th IEEE Conference on Decision and Control*. Vol. 3. IEEE. IEEE. Honolulu, HI. pp. 1902–1907.

Daniel, Eric D., Mee, C. Denis and Clark, Mark H., Eds.) (1997). *Magnetic Recording: The First 100 Years*. IEEE Press. New York, NY.

Davies, David B. and Michael D. Sidman (1991). Active compensation of shock, vibration, and wind-up in disk drives. In: *Proceedings of the 1991 ASME Winter Annual Meeting*. ASME.

Davies, David B. and Michael D. Sidman (1993). Active compensation of shock, vibration, and wind-up in disk drives. In: *Advances in Information Storage Systems, Vol. 5* (Bharat Bhushan, Ed.). pp. 5–20. ASME Press. New York, NY.

Diaz, Sam (2002). IBM to shed S.J. disk business. *The San Jose Mercury News* p. 1A.

Ehrlich, Richard M., Carl Taussig and Daniel Abramovitch (1989). Identification of sampled data systems at frequencies beyond the nyquist rate. In: *Proceedings of the 28th IEEE Conference on Decision and Control*. IEEE. IEEE.

Ehrlich, Rick (2001). Personal correspondence.

Franklin, Gene F. and J. David Powell (1980). *Digital Control of Dynamic Systems*. Addison-Wesley. Menlo Park, California.

Franklin, Gene F., J. David Powell and Michael L. Workman (1990). *Digital Control of Dynamic Systems*. second ed.. Addison-Wesley. Menlo Park, CA.

Hagopian, J.J. (1961). Data storage apparatus. U.S. Patent 3,007,144. IBM.

Hanselmann, Herbert and A. Engelke (1988). LQG-Control of a highly resonant disk drive head positioning actuator. *IEEE Transactions on Industrial Electronics* **35**(1), 100–104.

Hertrich, F.R. (1965). Average motion times of positions in random access devices. *IBM Journal of Research and Development* **9**, 124–133. Cited by Oswald:74. Analysis of average seek times for hard disk actuator type mechanisms.

Hoagland, Al (1961). A high track density servo access system for magnetic recording disk storage. *IBM Journal of Research and Development* **5**, 287–296.

Hurst, Terril (2001). Personal correspondence.

IBM Storage Technology Division (2001). IBM history of firsts. Web site. IBM. San Jose, CA 95193. http://www.storage.ibm.com/firsts/.

Kempf, Carl, William Messner, Masayoshi Tomizuka and Roberto Horowitz (1993). Comparison of four discrete-time repetitive control algorithms. *IEEE Control Systems Magazine* **13**(6), 48–54.

Knowles, Vernon (1991). An introduction to the wolverine servo system. Very instructive notes on HP's hard disk servo architecture.

Knowles, Vernon L. and Darwin M. Hanks (1987). Shock and vibration disturbance compensation system for disc drives. European Patent Application 871065555.3. Hewlett-Packard Co.. Corporate Patent Department, M/S 20B-O, 3000 Hanover Street, Palo Alto, CA 94304 USA.

Kurzweil, Fred (2001). Personal correspondence.

Mueller, F. E. (1972). Positioning system including servo track configuration and associated demodulator. U.S. Patent 3,691,543. IBM.

Oswald, Richard K. (1971). Position detection for a track following servo system. United States Patent 3,593,333. International Business Machines. San Jose, CA USA.

Oswald, Richard K. (1974). Design of a disk file head-positioning servo. *IBM Journal of Research and Development* **18**, 506–512.

Oswald, Richard K. (2001). Personal correspondence.

Ottesen, Hjalmar (2001). Personal correspondence.

Ottesen, Hjalmar H. (1988). Digital servo control for disk units. In: *IBM Application System/400 Technology*. pp. 156–160. International Business Machines. Rochester, MN.

Pao, Lucy Y. (1994). Characteristics of the time-optimal control of flexible structures with damping. In: *Proceedings of the Third IEEE Conference on Control Applications*. IEEE. IEEE. pp. 1299–1304.

Pao, Lucy Y. and Gene F. Franklin (1990). Time-optimal control of flexible structures. In: *Proceedings of the 29th IEEE Conference on Decision and Control*. IEEE. IEEE. pp. 2580–2581.

Pao, Lucy Y. and Gene F. Franklin (1992). Robustness studies of a proximate time-optimal controller. In: *Proceedings of the 31st IEEE Conference on Decision and Control*. IEEE. IEEE. pp. 3559–3564.

Pao, Lucy Y. and Gene F. Franklin (1993). Proximate time-optimal control of third-order servomechanisms. *IEEE Transactions on Automatic Control* **38**(4), 560–580.

Pao, Lucy Y. and Gene F. Franklin (1994). The robustness of a proximate time-optional controller. *IEEE Transactions on Automatic Control* **39**(9), 1963–1966.

Porter, James (2000). Personal correspondence.

Porter, James N. (1996). Where do we go from here?. *IDEMA Insights*.

Porter, James N. (1998). Disk drives' evolution. In: *Proceedings of the 100th Anniversary Conference: Magnetic Recording and Information Storage*. Santa Clara University.

Rostky, George (1998). Disk drives take eventful spin (history of disk drives) (industry trend or event). *Electronic Engineering Times* (1016), 22.

Sacks, Alexei Hiram (1995). Position Signal Generation in Magnetic Disk Drives. PhD thesis. Carnegie Mellon University. Data Storage and Systems Center, Department of Electrical and Computer Engineering, Pittsburgh, PA.

Sacks, Alexi H., William C. Messner and Marc Bodson (1996). MR Head effects on PES generation: Simulation and experiment. *IEEE Transactions on Magnetics* **32**(3, Part: 2), 1773–1778.

Santana, G. R. (1970). Method and apparatus for recording and detecting information. U.S. Patent 3,534,344. IBM.

Sidman, Michael (1986). Adaptive Control of a Flexible Structure. PhD thesis. Stanford University. Department of Electrical Engineering.

Sidman, Michael D. (1985). Adaptive misposition correcting method for magnetic disk servo system. United States Patent 4,536,809. Digital Equipment Corporation. Maynard, MA.

Sidman, Michael D. (1991). Convergence properties of an adaptive runout correction system for disk drives. In: *Advances in Information Storage Systems, Vol. 1* (Bharat Bhushan, Ed.). pp. 61–78. ASME Press. New York, NY.

Silicon Systems Staff (1994). SSI 32P4730 Read Channel with 1,7 ENDEC, 4-burst Servo. Technical report. Silicon Systems, Incorporated. http://www.ti.com/sc/docs/storage/products/data/4730.pdf.

Singer, Neal C. and Warren P. Seering (1988). Using acausal shaping techniques to reduce robot vibration. In: *Proceedings of the 1989 IEEE International Conference on Robotics and Automation*. IEEE. IEEE. Scottsdale, AZ. pp. 1434–1439.

Singer, Neal C. and Warren P. Seering (1989). Design and comparison of command shaping methods for controlling residual vibration. In: *Proceedings of the 1989 IEEE International Conference on Robotics and Automation*. IEEE. IEEE. Scottsdale, AZ. pp. 888–893.

Smith, Robert F. (1993). Apparatus for sensing operating shock on a disk drive. United States Patent 5,235,472. Seagate Technology, Inc.. Scotts Valley, CA.

Stevens, Fred and Joe DeLillis (2000). Keeping heads on track with dual-stage actuators. *Data Storage* **7**(6), 39–44.

Stevens, Louis D. (1997). Data storage on magnetic disks. In: *Magnetic Recording: The First 100 Years* (Eric D. Daniel, C. Denis Mee and Mark H. Clark, Eds.). Chap. 18, pp. 270–299. IEEE Press. New York, NY.

Stich, M. C. (1987). Digital servo algorithm for disk actuator control. In: *Proceedings of Conference on Applied Motion Control '87*. Minneapolis, MN. pp. 35–41.

Stich, Michael C., John B. Resman, Hjalmar H. Ottesen and Todd B. Anderson (1987). Adaptive control technique for a dynamic system. U.S. Patent 4,697,127. IBM.

Tickell, R. (1966). A high-performance position control. In: *Proceedings of the Joint Automatic Control Conference*. JACC. University of Washington, Seattle, Washington.

Tomizuka, Masayoshi, Tsu-Chin Tsao and Kok-Kia Chew (1988). Discrete-time domain analysis and synthesis of repetitive controllers. In: *Proceedings of the 1988 American Control Conference*. AACC. IEEE. Atlanta, GA. pp. 860–866.

Wang, Feei, Terril Hurst, Daniel Abramovitch and Gene Franklin (1994). Disk drive pivot nonlinearity modeling Part II: Time Domain. In: *Proceedings of the 1994 American Control Conference*. AACC. IEEE. Baltimore, MD. pp. 2604–2607.

White, Robert A. (1977). Shock force compensating system. United States Patent 4,040,103. Sperry Rand Corporation. New York, NY.

Wilson, Rosser S. (1994). Sample data position error signal detection for digital sector servo. United States Patent 5,301,072. Mastor [sic.] Corporation. San Jose, CA USA.

Workman, Michael L. (1987*a*). Adaptive Proximate Time-Optimal Servomechanisms. PhD thesis. Stanford University. Stanford, CA.

Workman, Michael L., Robert L. Kosut and Gene F. Franklin (1987*a*). Adaptive proximate time-optimal servomechanisms: Continuous-time case. In: *Proceedings of the 1987 American Controls Conference*. AACC. IEEE. pp. 589–594.

Workman, Michael L., Robert L. Kosut and Gene F. Franklin (1987*b*). Adaptive proximate time-optimal servomechanisms: Discrete-time case. In: *Proceedings of the 26th IEEE Conference on Decision and Control*. IEEE. IEEE. pp. 1548–1553.

Workman, M.L. (1987*b*). *Applications of Magnetic Recording*. Chap. 2: The Head Positioning Servomechanism, edited by D. Mee and E Daniels. Mc Graw Hill.