

Hard-disk-drive technology flat heads for linear tape recording

R. G. Biskeborn
J. H. Eaton

IBM thin-film tape heads have evolved from the ferrite-based heads first used in the IBM Model 3480 Tape Drive to the hard-disk-drive (HDD) technology flat-profile heads used in IBM Linear Tape-Open® (LTO®) products. This paper describes that transition and discusses the flat tape head manufacturing processes, drive implementation, performance, and outlook. Thin-film head technology for hard-disk drives was first used in tape heads in the early 1990s, when IBM built quarter-inch cartridge head images on HDD-type wafers. This was a springboard for the next step, flat-lapped tape heads, which use not only HDD wafers, but also HDD post-wafer machining technologies. With the emergence of LTO, flat heads entered mainstream tape head production in IBM. These have proven to have high performance and durability.

Introduction

In 1984, IBM introduced the 18-track Model 3480 Tape Drive, the first in IBM to use thin-film, magnetoresistive (MR) heads and half-inch tape cartridges. Prior to this, IBM tape drives, culminating in the Model 3420, employed large, open tape reels that were manually loaded by operators. The advent of the tape cartridge paved the way for automation and advanced tape libraries. The recording heads in the Model 3480 were conceived in the IBM Boulder facility and manufactured at the new IBM Tucson plant site [1]. Those heads are built on magnetic nickel-zinc ferrite wafers, a material that is amenable to post-wafer machining. In the late 1980s, IBM moved tape-head manufacturing from Tucson to San Jose, California, to consolidate tape and HDD manufacturing operations. However, IBM continued to manufacture tape heads in facilities separate from the previous HDD thin-film line. One reason was processing incompatibility. For instance, HDD makes extensive use of plating, which is not used in thin-film ferrite head processing. Ferrite tape-head MR sensors are deposited directly onto the wafer surface, on which a thin first-gap layer has already been deposited. The sensors are then stabilized magnetically by incorporating chevron-shaped depressions, a method not used in HDD thin-film heads. Another incompatibility was

that posed by the magnetic nature of tape-head wafers. Moreover, there was concern over ferrite contamination getting into the HDD thin-film line. On the other hand, converting tape heads to HDD wafers was ruled out largely because of a post-wafer processing difficulty, namely, contouring the tape heads. HDD wafers are extremely hard and would have been much more difficult to machine.

Flat heads

A break with this tradition came in 1993 when engineers built single-track, unidirectional quarter-inch-cartridge (QIC) tape-head wafers on the HDD thin-film wafer line. The QIC heads employed HDD-type write transducers and anisotropic magnetoresistive (AMR) read sensors with thin-film laminated iron magnetic poles and shields, respectively, and these structures obviated the need for magnetic ferrite and chevrons. This established that HDD thin-film head processes could be used to make comparatively long (nearly 150 μm) stable MR sensors with soft adjacent-layer biasing and hard-bias magnetic stabilization. IBM did not actually build finished heads, but rather sold the wafers to a company that supplies original equipment manufacturer (OEM) QIC heads for the desktop backup market. Cylindrical contouring was accomplished by means other than those in place in IBM,

©Copyright 2003 by International Business Machines Corporation. Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free without further permission by computer-based and other information-service systems. Permission to *republish* any other portion of this paper must be obtained from the Editor.

0018-8646/03/\$5.00 © 2003 IBM




Figure 1

Closeup of an early Ultrium[†] 1 chip image showing the side-by-side ("interleaved") write transducers and read sensors, servo reader, and gold-plated connection pads.

and was largely possible due to the comparatively narrow QIC head profile.

It was evident at the time that there would be merit in expanding the scope of QIC activities, possibly into IBM's own Enterprise tape-drive business, which includes, for example, Models 3580 and 3590. Accordingly, early in 1996, the authors began experimenting with writing and reading tape using flat-lapped, HDD thin-film head rows without the HDD slider air-bearing patterning. We set out to prove that HDD head technology could be used to build multi-track, bidirectional, flat tape-recording heads that would be viable alternatives to conventional heads. While flat-tape-head studies had been published [2], structures had not been developed for application in a commercial tape drive. Our goal was to use HDD wafer and post-wafer machining expertise and equipment wherever possible. The QIC experience had shown that substantial reductions in chip size are possible. Using a flat profile would streamline head build and enable high-volume production compared with the relatively labor-intensive, low volumes of the current tape-drive products.

Linear Tape-Open[†]

At about the same time, and in response to demand for interchange standards and improved performance in high-capacity network tape storage, Hewlett-Packard and Seagate formed a consortium and initiated work on what was to become the Linear Tape-Open (LTO[†])

specification, formally defined by ECMA standard 319 [3]. IBM joined the consortium in 1997 and predicted that demand for LTO multi-track heads would tax IBM ferrite tape-head production capabilities. In the meantime, the authors had demonstrated the performance and wear advantages of flat, HDD thin-film technology-based tape heads. This, coupled with projected volumes, led IBM to adopt flat tape heads for LTO drives. This proved to be very successful, and IBM has since established LTO head production facilities in the U.S. and abroad.

This paper describes the new tape heads and focuses on some of the important differences in design and mechanical fabrication methods compared with IBM ferrite Enterprise tape heads. It discusses merging this head into the emerging compact IBM TotalStorage[®] 3580 tape-drive family. This integration has mandated a new actuator design, a new tape path, and a new head-alignment scheme. Performance advances are discussed, and the paper concludes with an outlook for flat-head technology.

Tape-recording basics

Read-while-write and interleaved transducers

Some perspective on tape heads, tape recording, and the emergence of track-following servos elucidates the motivation behind the work discussed in this paper. Tape-recording systems must deal with a seemingly arcane problem set, including imperfect, flexible, and fragile media; thick, particulate, and abrasive coatings; tape-guiding instabilities; omnipresent dust and tape debris; backward compatibility; cartridge interchange and tape brand interchange; etc. Tape-recording scientists and engineers take on these challenges, but they must constantly be aware that the main appeal of tape is its cost-effectiveness when compared with other types of storage systems. A basic requirement is that data must be verified in real time as it is written. So-called *write stops*, when data is not verified and so is rewritten, are a fact of life in modern linear tape recording. Verifying on the fly means that each write transducer must have an active read sensor immediately downstream.

In the Model 3480, write transducers and read sensors were housed in separate, dedicated modules, which were then mounted in tandem. Thus, the Model 3480 operated in read-while-write mode with the tape moving in one direction only. Starting with the 36-track IBM Model 3490E tape drive in 1991, read sensors and write transducers have been placed next to each other in an alternating, side-by-side ("interleaved") arrangement [4]. Two modules are aligned for read-while-write in both tape travel directions. **Figure 1** shows interleaved readers and writers for the LTO head.

Track-following servo

The push to higher cartridge capacity mandated narrower data tracks. In both the 3480 and 3490E, heads are mounted in one fixed position in the drive. Tape guiding imperfections and narrow data tracks precipitated the development of track-following servo control. The Model 3570 eight-mm cartridge tape drive, introduced in 1995, was the first IBM tape product to use head servo control, or “servoing.” The method employed is *timing-based servoing* (TBS), in which the duration between pulses from obliquely written patterns contains the position information. Tapes are factory-formatted with the TBS tracks. The head is mounted on a voice-coil-driven positioning device, and dedicated sensors in the head read the servo-track data. A position error signal controls the voice coil. The actuator performs two functions, moving the head to specific track locations on the tape and maintaining alignment between the head and the servo tracks.

Ferrite-based tape heads evolved to their most advanced form in the IBM Model 3590H, again in an interleaved design. Like the Model 3570, the 3590H drive employs a voice coil actuator and track-following servo. Unlike the Model 3570, the 3590H servo method is based on the read-back amplitude of two tones prerecorded in the servo tracks. With the emergence of LTO, the TBS approach was adopted because it was judged to be extendable to very high track-density recording and it had a proven track record in the Model 3570.

The discussion so far has focused on the need for read-while-write and track-following servoing. These, in turn, influence head design and data organization decisions. The IBM approach to its LTO head was to keep the interleaved, two-module design. Readers familiar with HDD thin-film heads might wonder why an HDD-type *piggyback* design—in which write transducers are fabricated directly on top of read sensors—was not used. In fact, some LTO heads, and even the IBM QIC heads, are piggybacked. Apart from the fact that IBM has years of experience using interleaved heads, another reason for choosing the interleaved design is that it enables track-following via a single pair of servo readers in each module. Heads access tracks in both tape directions at every index position. By way of comparison, piggyback reader-writer designs require two pairs of servos, and heads must index when the tape changes direction [5]. In addition to the interleaved readers and writers, Figure 1 shows a TBS servo reader.

Wafer layout and head design

47-mm rows

With the above background, a discussion of the LTO layout and how HDD thin-film processes are exploited




Figure 2

Chip images of (a) an early LTO wafer and (b) a 3590H product wafer, both photographed at the same magnification. Approximately 12 LTO images fit in the space occupied by two 3590H images. TC bond leads are substantially longer than the compact ultrasonic bond pads. Both products employ interleaved read sensors and write transducers. The recording element span is 2.8 mm for LTO and 11.5 mm for 3590. The reader-to-writer pitch is 0.166 mm for LTO, an eight-channel product, and about twice that for 3590, a 16-channel product. The outermost devices in the LTO image are electronic lapping guides (explained in the text). The LTO contacts span only approximately 7 mm.

follows. The Ultrium[†] 1 and 2 (*Ultrium* is the official consortium name for the first members of the LTO tape-drive product family) are both eight-channel data-recording devices, with eight readers, eight writers, and two servo readers in each Ultrium module, or *chip*. To take full advantage of HDD thin-film wafer processing and post-wafer machining equipment, LTO chip images are arranged on the wafer across 47-mm rows, the same length as HDD thin-film rows. From the start, the LTO chip images were made 22.5 mm long to support the tape at head position extremes. More chips would fit, as can be seen in **Figure 2**, which shows the IBM LTO and 3590H chip images on the same scale. Making the images this long avoided the complexity of assembling outboard tape supports, but resulted in only two chips per row. In anticipation of future possibilities, contact pad pitch was reduced from typical HDD head dimensions to 140 μm to permit six images per row. Although methods have been




Figure 3

Ultrium 1 chip image schematic showing writers, readers, servo elements, and tracks. A 5-in. wafer contains 476 modules (238 heads). Each module is 22.5 mm $\times$ 0.75 mm (two per row). There are 16 interleaved write and read elements on a 166.5- μm pitch.

devised for building partial-span “chipllets” into full-span tape heads [6, 7], none have yet been implemented.

The preceding elucidates the rationale for currently having only two chip images per 47-mm row, but there is more to discuss regarding the arrangement of chip devices within the images themselves. It is evident in Figure 2 that the transducer span for the LTO chip is significantly less than that for the 3590H chip. This has an interesting history and timing relative to 3590H. Starting with the IBM Model 3480 Tape Drive, simultaneously written tracks have spanned nearly the entire tape width. However, hygroscopic swelling and contraction of the tape and other effects cause tracks to shift by up to 0.1% relative to one another. With relatively wide data tracks, this can be tolerated. In the 3590H format, for example, separation between the outermost tracks can creep by nearly 10 μm , and this is only slightly less than the difference between written track and read sensor widths.

Linear tape drives generally employ a *write wide, read narrow* format (as opposed to an *azimuth* recording format), and so it is critical that data readers not straddle adjacent tracks. Since relentless pressure to increase the storage capacity of tape cartridges is driving further reductions in track width, tape creep limits the span of simultaneously written tracks. The LTO consortium members addressed this problem directly: LTO data occupies the four regions between five factory-written, timing-based servo bands, two of which are very close to the edges of the tape. The servo readers are spaced only 2.82 mm apart, or 22% of the tape width. Thus, the read sensors and write transducers located between the two servo readers span less than one-quarter of the tape width. While Ultrium 1 and 3590H cartridges each have 384 data tracks, LTO is easily extendable to significantly higher track density. The device layout within the IBM Ultrium 1 module is shown in **Figure 3**.

Row pitch and ultrasonic bonding

The discussion so far addresses changes and improvements embodied in the LTO head layout within a row, but there is another key advantage of the HDD-technology-based design as implemented in the LTO head; the height of the LTO chip image—and thus the row pitch—is more than sixfold less than its Enterprise counterpart, and this offsets the sacrifice of wafer real estate within a row. In addition to the sensors themselves, the 3590H chip must accommodate a cylindrical contour, air bleed slots, and thermal compression (TC) bonding pads, whereas the LTO chip need accommodate only the shallow flat-profile and compact ultrasonic bonding pads. **Figure 4** shows finished module sections of both types. It does not seem that TC bonding would play an important role in chip layout, but in TC bonding, the pads are taller and farther from the recording elements than are the ultrasonic pads. While TC bonding has the advantage of fusing all cable leads in one step, the chip contacts reach 600°C, and so these are located far enough from the sensors to prevent excessive temperature rise. Because HDD wafer material is a better thermal conductor than ferrite, the pads must be placed even farther from the recording elements. These considerations in part enable the LTO 0.75-mm so-called *nano-wafer row pitch*. Even with only two chips per row, LTO wafers have 476 chip images, compared with 90 for Enterprise wafers. A surprising caveat having to do with the one-at-a-time head-building approach used for 3590H heads is revealed later in the post-wafer fabrication section.

Magnetic considerations

A discussion of leveraging HDD thin-film head technology for tape heads would be incomplete without a description of the thin films themselves. For the greatest possible compatibility, HDD wafer thin films are utilized everywhere and include the Sendust reader shield 1 (S1), the




Figure 4

LTO and 3590E module sections. The 3590E tape-contacting surface is contoured and slotted, while the LTO surface is flat, with sharp, air-skiving edges. The LTO module occupies only one-sixth as much wafer real estate.

Permalloy (81% nickel, 19% iron) reader shield 2 (S2), the writer pole 1 (P1), and the 45/55 nickel-iron (45% nickel, 55% iron) writer pole 2 (P2) [8]. Photographs of the top (tape or disk) bearing surfaces of 3590H, LTO, and HDD heads are shown in **Figure 5**. The shields and poles are easily identified. What is not obvious is the fact that using Permalloy for the S2 reader shield is a departure for IBM tape heads, since both Enterprise and QIC heads use reader shield materials specially engineered for running on tape. The latter are relatively hard, nonductile sputtered films consisting of iron, nickel, and nitrogen. Permalloy, which possesses excellent magnetic shield properties and very low magnetostriction, is ductile and has a propensity for forming electrical contacts when deformed. In the end, the combination of flat rather than cylindrical profile and smooth, dual-coat media enables the use of Permalloy shields. This is discussed further in the performance section of this paper.

IBM LTO heads employ an AMR read sensor fabricated using HDD materials and processes. As such, it is an ion-beam deposited tri-layer with a high-resistance soft adjacent bias layer for current optimization and hard-bias magnet stabilization at the ends (integrated with the leads). This device has very low distortion, as evidenced by near-zero asymmetry. The magnetic shield designs, stripe height, and shield-to-shield spacing are optimized for use with LTO dual-coat magnetic media and Ultrium linear densities. For Ultrium 1 media, the iron particulate magnetic coating is 0.22 μm thick; for Ultrium 2 media, it is 0.1 μm thick. This is approximately ten times thinner




Figure 5

than 3590 single-layer metal-particle media. Accordingly, its linear density of 3660 flux reversals per millimeter (compared with 2500 for the 3590H) is achieved using write transducers having an 0.8- μm gap and a read sensor having a 0.37- μm shield-to-shield separation. These gap dimensions are several-generations-old AMR HDD technology (introduced in 1991), while the 27- μm write-track width and 13- μm read-track width are even older, actually predating the use of MR heads in hard-disk drives.

Ultrium 1 media coercivity is approximately 1850 Oe, and is written with a single-layer copper-plated eight-turn writer coil. Write gap and throat height were selected to give the best balance of transition sharpness and overwrite signal-to-noise. The 45/55 nickel-iron P2 write pole can saturate tapes of this, and higher, coercivity. Tape motion during writing is from the first pole (P1) toward the second pole (P2); i.e., P2 is trailing. This provides better track-

than 3590 single-layer metal-particle media. Accordingly, its linear density of 3660 flux reversals per millimeter (compared with 2500 for the 3590H) is achieved using write transducers having an 0.8- μm gap and a read sensor having a 0.37- μm shield-to-shield separation. These gap dimensions are several-generations-old AMR HDD technology (introduced in 1991), while the 27- μm write-track width and 13- μm read-track width are even older, actually predating the use of MR heads in hard-disk drives.

Table 1 Evolution of IBM tape products.

	3480	3490E	3590B	3590E	3590H	LTO Gen 1	LTO Gen 2
Year introduced	1984	1991	1995	1999	2002	2000	2002
Native capacity (GB)	0.2	0.4	20	40	60	100	200
Number of written tracks	18	36	128	256	384	384	512
Data rate (Mb/s)	3	4.5	9	14	14	15	35
Max. linear density (fc/mm)	972	972	2550	2550	2550	3660	7860
Servoing scheme	—	—	Lf/Hf ratio	Lf/Hf ratio	Lf/Hf ratio	Timing-based	Timing-based
Tape length (m)	165	340	340/600	340/600	340/600	609	609
Tape velocity (m/s)	2	2	2.2	3.14	3.14	4	3.11–6.22
Magnetic particle	CrO ₂	CrO ₂	MP++	MP++	MP++	Fe on underlayer	Fe on underlayer
Coating thickness (μm)	2.6	2.4	2.6/2.4	2.6/2.4	2.6/2.4	0.2	0.2
Coercivity (Oe)	500	500	1600	1600	1600	1850	2100
Particle size (nm)	650	650	120	120	120	20	20
Media base film	PET	PET	PET/PEN	PET/PEN	PET/PEN	PEN	PEN
Media thickness (μm)	30	17.3	16.4/9.4	16.4/9.4	16.4/9.4	8.9	8.9
Surface roughness (nm)	15	12	6	6	6	3–4	2.5–3.5
Transducer layout	Separate R/W modules	Interleaved	Interleaved	Interleaved	Interleaved	Interleaved	Interleaved
Span of active writers (mm)	full	full	11	11	11	2.332	2.332
Writers (per head)	18 (left)	36	32	32	32	16	16
Servo readers (per head)	—	—	12	12	12	4	4
Wafer	Ni–Zn magnetic ferrite	Ni–Zn magnetic ferrite	Ni–Zn magnetic ferrite	Ni–Zn magnetic ferrite	Ni–Zn magnetic ferrite	Al ₂ O ₃ TiC	Al ₂ O ₃ TiC
Reader width (μm)	410 (center-tapped)	195	35	20	12	12.6	10
Writer width (μm)	540	285	80	40	27	27.5	27.5
Servo width (μm)	—	—	55	76	82	7.7	7.7
Reader stripe height (μm)	12	9	3.5	2.5	2	2.25	2.25
Servo stripe height (μm)	—	—	3.5	3.25	5		
MR trilayer (NiFe/Ta/SAL) (Å)	520/—/—	580/250/360	367/210/275	315/210/255	300/210/270	300/60/240	300/60/40
Reader shield spacing (μm)	1.3	1.3	0.65	0.65	0.65	0.37	0.37
Write gap (μm)	1.3	1.3	0.84	0.84	0.84	0.8	0.8
Coil turns	2	3	4	4	4	8	12
First shield	Ni–Zn ferrite wafer	Ni–Zn ferrite wafer	Ni–Zn ferrite wafer	Ni–Zn ferrite wafer	Ni–Zn ferrite wafer	Sendust	Sendust
Second shield	Ni–Zn ferrite closure	Ni–Zn ferrite closure	Fe(N)/NiFe(N) laminate	Fe(N)/NiFe(N) laminate	Fe(N)/NiFe(N) laminate	Permalloy	Permalloy
First pole	Ni–Zn ferrite wafer	Ni–Zn ferrite wafer	Ferrite + laminate	Ferrite + laminate	Ferrite + laminate	Sendust + Permalloy	Sendust + Permalloy
Second pole	Ni–Zn ferrite closure	Ni–Zn ferrite closure	Fe(N)/NiFe(N) laminate	Fe(N)/NiFe(N) laminate	Fe(N)/NiFe(N) laminate	45–55 NiFe	45–55 NiFe

edge definition and sharper written transitions than writing in the opposite direction. Linear density and track pitch combine to give Ultrium 1 100 GB of storage per cartridge. Ultrium 2 media coercivity is 2200 Oe or higher and is written using a 12-turn write head. A slightly higher linear density, track shingling (in which each new track writes over a portion of a previously written track), and partial response maximum likelihood (PRML) detection enable a 200-GB capacity. Readers interested in comparing generations of tape products may refer to the data in **Table 1**.

Post-wafer fabrication

Closure process and lapping

This section describes the unique characteristics of flat heads and the challenges associated with their manufacture. Particular attention is given to the important differences between the HDD-based flat tape heads and IBM's conventional, cylindrically contoured ferrite heads. Unlike traditional cylindrically contoured tape heads, the flat-profile design employs wear-resistant tape supports that have sharp edges for skiving air off the tape [9–11]. Aluminum oxide–titanium carbide ($\text{Al}_2\text{O}_3\text{--TiC}$, or *AlTiC*) HDD thin-film head wafer material is exceptionally hard and well suited for flat heads. By contrast, the nickel–zinc ferrite used in conventional heads is too soft. This is why ferrite heads require cylindrical contours, which in turn use air-bleed slots to tack down the tape.

Tape-head transducer gaps are situated between tape-support surfaces. One of these, the *substrate*, is formed from the wafer substrate itself. The other, the *closure*, is a block of ceramic material bonded directly to the surface of the thin films. IBM ferrite tape-head closures are individually machined and then bonded to chips before lapping. Accordingly, both the chips and closures are large enough to facilitate handling. In stark contrast to this, flat head closures are batch-processed at the quad level [9]. (A *quad*, which is a term borrowed from HDD thin-film wafer processing, is a rectangular wafer section comprising, in this case, 17 rows.) The quads and closures are shown in **Figure 6**. Grooves that are 0.45 mm wide are machined into an AlTiC block, which has dimensions approximately the same as the quad; the grooves are aligned with the electrical contact pads, and the block is bonded [12]; finally, the back web is machined off, thus exposing the contacts and leaving the closure strips in place. Forming the closures *in situ* enables the tight row pitch.

An advantage offered by the HDD post-wafer lapping process is the ability to lap up to four quads (eight modules) simultaneously on HDD row-lapping equipment, compared with cup-lapping ferrite modules one at a time. Row-lapping creates flat, low-gap-recession tape-contacting surfaces with excellent stripe height control, typically within $\pm 0.25\ \mu\text{m}$ or better. Post-lapping closure




Figure 6

(a) LTO mini quad, showing 17 rows and two chip images per row. (b) Finished closure strips after back grind.

creep is less than 15 nm, thus ensuring reliable head–tape contact. Electrical lapping guides (ELGs) monitor stripe height. A difficulty unique to flat HDD-based tape heads was contacting the ELG pads, now recessed in the channels between adjacent closure strips. This difficulty was overcome by outfitting an insulating block with pins having small-diameter ends for contacting the ELG pads. A conventional lapping cable contacts the larger-diameter opposite ends.

Module processing: U-beams and cables

Lapped devices are sawed off the quads immediately after lapping, forming the rows (sometimes called *row bars*). Processing in this fashion is preferable to a common HDD thin-film head practice of sawing quads into row bars and then bonding the row bars to carriers for lapping, as de-bonding exposes the closure bonds to detrimental temperatures and solvents. Finished rows are diced into full-span chips and then bonded to ceramic beams. Tape-bearing surfaces are machined to 0.8-mm width. This controls tape edge curl [11]. The support beams maintain tape-bearing surface flatness and are U-shaped to facilitate cable attachment, module assembly, and module-to-module assembly. A U-beam and a finished chip comprise a module [13, 14], as shown in **Figure 7**.

One might wonder, why go through the trouble of cutting rows this thin only to wind up bonding them to




Figure 7

Ultrium module shown after ultrasonic cable attachment and strain relieving. Reader and writer lines are separated in the cable close to the module to minimize crosstalk. The leads are cut straight and do not contact the closure, which is conductive.




Figure 8

Cabled LTO flat modules joined to form a read-while-write head. Substrate, closures, and U-beams are shown. The modules are angled slightly relative to one another to create the required air-skiving effect.

U-beams? Would it not be simpler to make the row bars 2 mm thick in the first place? The reason is that wafer expense is a primary consideration in the cost of the heads. It is much more economical to bond thin chips to relatively inexpensive U-beams. Another factor, having to do with cabling and assembly, is explained below.

Cables, which are attached prior to assembling module pairs, are manufactured with a short section of the leads exposed and supported by the frame, which is part of the cable itself. This design was originally created for TC bonding, but with a small “twist,” it was ideally suited to ultrasonic bonding. Leads are aligned with chips in the bonding fixture; the twist is chopping the leads to the desired length, thus ensuring that they are all the same length and do not contact the (conductive) closures. Prior to this, the frame protects the extremely fragile leads. Next, the leads are individually ultrasonically bonded directly to the electrical contact pads. Thus, stitch bonding, which would have produced twice as many connections and which is also slower, is averted. Bonded cable ends are strain-relieved to the U-beams.

Read-write heads

Cabled modules are assembled in pairs by ultraviolet (UV) epoxy bonding the U-beam legs together to form bidirectional read-while-write heads. Track-to-track alignment and wrap angle between the two modules are well controlled. A novel system based on optical fringes for aligning the modules to 1- μm track-to-track tolerance and setting the angles to better than 0.1° was devised [15]. A key to preserving these critical alignments during bonding is symmetrical disposition of the U-beam bonds. Curing strains tend to balance one another and thus upset neither the displacements nor the angles. A finished head assembly is shown in **Figure 8**, along with a schematic diagram showing how the tape wraps the lapped surfaces.

In Figure 4, the construction differences between the LTO flat profile head and the slotted, contoured 3590H head can be seen. The LTO modules consume much less wafer material, do not need the air-bleed slots, have batch-processed closures, and are hard and wear-resistant. In spite of the differences, it is interesting to note the similarity of the LTO tape contacting surface width and the 3590H width between inner slots.

Drive integration

Two-stage actuator and tape path

New in IBM is the two-stage actuator and passive tape path used for LTO and flat heads. The Enterprise tape products use a single-stage, voice-coil design. The latter can be used because, as discussed in the design section, the tracks span nearly the entire width of the tape, and so the required head positioning range is less than about a millimeter. In the LTO format, the head indexes to tracks in one of four bands, and the range of head travel is about 9 mm, which is beyond the capability of the IBM single-stage actuators. The LTO actuator is composed of a high-bandwidth, low-mass, spring-mounted voice coil and a heavier stepper-motor-driven carrier. A description of the

head and actuator assembly, shown in **Figure 9**, follows. The left module of the assembled pair is bonded to a molded plastic base. The base becomes part of the actuator voice-coil assembly. The right module remains attached to the left module only to prevent undue stress on the U-beam legs. The cables are relatively wide and stiff and thus are formed into loops to minimize their effect on the actuator spring constant.

The actuator assembly rides up and down on a lead screw, as shown in the exploded view of **Figure 9(b)**. The lead screw is fixed in the drive, and a worm gear attached to the stepper motor drives a geared nut (hidden in the figure) on the lead screw, thus enabling coarse track positioning. The lower lead-screw end is adjustable for setting skew. An anti-rotation arm rides under spring load against a guide and enables decoupling of both skew and rotation from actuator vertical motion. The lead screw, worm gear, and motor provide a coarse step of $3.3\text{ }\mu\text{m}$. The voice coil moves $\pm 150\text{ }\mu\text{m}$. The voice-coil actuator moving mass is just over 3 g. **Figure 10** shows the voice-coil actuator open-loop frequency response. The departure from the mass line near 7 kHz is due to the cables. The loop is closed with a bandwidth of 1 kHz.

The actuator burden is determined by the tape characteristics and tape path. The Ultrium tape path departs significantly from its Enterprise counterparts, where the tape is steered by large positive-pressure D-bearings and multiple compliant button edge guides. In the Ultrium product, four rollers, two on each side of the head-actuator assembly (HAA), guide the tape from the cartridge over the head to the take-up reel. This relatively simple tape path controls tape lateral transients about as well as the D-bearing path. Head position errors are typically less than $3.5\text{ }\mu\text{m}$, worst case. This is in large measure a result of the grooves provided on the surfaces of the rollers adjacent to the head [16]. The head and grooved rollers are shown in a drive in **Figure 11**. Grooved roller modeling, which is discussed in [17, this issue], shows that the operating principle is tape tacking at the high-contact-pressure groove edges. Tracking can be affected by take-up-reel flange hits, scraping inside the cartridge, and roller-flange imperfections, and so all must also be controlled.

Same-gap servo and wrap angle

In the same-gap servo mode, servo readers in the leading (writing) module, and thus in the *same gap* as the writers, detect head position, rather than the servo readers in the trailing (reading), or *opposite*, module. If the drive is viewed from the side, it is evident that the angle the tape makes with the head (the external tape-wrap angle) is determined by the position of the inner rollers relative to the HAA. Setting this angle precisely is important for reliable same-gap-servo-mode drive operation. Same-gap servo, which is used in the IBM Ultrium 2 product, gives




Figure 9

LTO two-stage actuator shown (a) assembled and (b) exploded.




Figure 10

Voice-coil actuator open-loop frequency response displaying a resonance just above 200 Hz and a usable range greater than 1 kHz. The deviation in response near 7 kHz is due to the cables.

more accurate placement of tracks during writing and is less susceptible than opposite-gap servo to misplacement caused by dynamic skewing of the tape.

However, same-gap servo requires better noise isolation between servo readers and writers and is more susceptible to edge loss [10, 11] than opposite-gap servo. Edge loss is worse because of tape cupping, in which the tape curls away from the head. Tapes are manufactured with the magnetic side convex to prevent edge nicking by heads in rotary applications. Tape cupping is significantly more pronounced in the 19-mm span between head and rollers




Figure 11

Head, actuator assembly, and grooved rollers. The LTO head is considerably longer than the tape width. Inner roller positions are adjustable to facilitate setting external wrap angle in the drive. The grooves help limit tape lateral transients, which otherwise consume voice-coil actuator bandwidth.




Figure 12

Read-back amplitude of an 8T pattern recorded at the current required to reach maximum 2T output on Ultrium 1 media as a function of no-tape quasi-static output. There are two populations, but within each the amplitudes are correlated. Courtesy of Leif Kirschenbaum, IBM San Jose.

than in the 0.7- to 0.9-mm span between modules. Increasing wrap minimizes edge loss, but over-wrapping degrades head-tape spacing, particularly at higher tape speeds [11]. Edge loss is a smaller effect, both for read-




Figure 13

Scanning electron micrographs of LTO Ultrium 1 magnetic tape surface at (a) 20,000 $\times$ and (b) 50,000 $\times$. The acicular (needle-shaped) iron particles are oriented along the tape. The large, non-magnetic particles are aluminum oxide and are added to the binder to increase its durability and to reduce friction between head and tape. Courtesy of Garvin Stone, IBM San Jose.

back and opposite-gap servo mode, which is employed for the IBM Ultrium 1 drive.

The method devised to set wrap angle on the outside edges of the head [18] is based on a surprising characteristic of the tape foil bearing—the transition from the skiving and tacked-down state to the flying state is abrupt. Read-back signal vanishes abruptly when the tape wrap is reduced to approximately 0.05°, having shown, at most, some instability immediately prior to vanishing. Thus, the transition provides an accurate indication of tape-wrap angle. The positions of the two rollers adjacent to the HAA are adjustable. These are used for lifting the tape off the head, and so finding the transitions. Servo-track read-back signals facilitate this procedure, which is performed when drives are built.

Skew is also adjusted in the drive using the servo track signals to minimize residual head-build track-to-track misalignment, which is kept to less than approximately $5\text{ }\mu\text{m}$ at module assembly. Drive-level skew adjustment is possible because the opposing module gaps are spaced by 1.5 mm. Bit cell aspect ratio ultimately places a limit on this capability.

Cable design

It is clear from the above discussion that cable design plays an important role in signal integrity. While the head was still on the drawing board, a decision had to be made regarding unscrambling read and write lines, a problem unique to multi-track tape heads. Unscrambling in the chip or in the cable immediately adjacent to the chip is necessary to minimize coupling of radiated noise from the writers into the unused reader lines, a process that introduces unwanted noise into the read electronics. The decision was made to unscramble in the cable rather than on the chip, as wafer row pitch was already being taxed. Reader leads are strip-line arranged (pairs over a ground plane) for shielding and common-mode rejection. An aluminum shield attached to the module end of the left cable further reduces crosstalk between modules during read-while-write. Over-under-writer leads minimize writer noise radiation and are enablers for same-gap servo.

Performance

Quasi-static yield

One of the HDD technologies from which LTO heads borrow is quasi-static testing, which ascertains MR response in an externally applied, uniform magnetic field. LTO rows are 100% quasi-static tested using HDD row-bar procedures and equipment. Fallout at quasi-static screening is typically 5% or less for Ultrium chips. Its value beyond use as a screen has been debated, but quasi-static and tape read-back amplitudes are found to correlate for IBM LTO heads, as shown by the data in **Figure 12**. The appearance of two distributions is believed to result from lower tape read-back amplitude for heads having recessed recording elements. Quasi-static testing also provides hot and cold resistances and is used for screening row bars for telegraph noise. The latter occasionally arises as a result of intermittent metallic bridges formed between the S2 shields and MR sensors. Row bars are finish-lapped parallel to the gaps to minimize the incidence of shield-to-sensor bridges [19, 20]. Obviously, this is increasingly important as read gaps approach present HDD head dimensions ($0.1\text{ }\mu\text{m}$).

Head wear

LTO dual-coat media is smooth and relatively nonabrasive. Even so, it still produces observable wear in




Figure 14

Magnetoresistive sensor resistances for the tracks in an Ultrium 1 head run against IBM media versus time. The upper graph represents the first 175 hours, and the next 1000 hours are shown below. A $0.1\text{-}\Omega$ increase corresponds to approximately $60\text{ }\text{\AA}$ of wear, which is indeed very small.

the head gaps. The magnetic coating has typical peak-to-valley roughness of only 35 nm. **Figure 13** consists of two scanning electron micrographs of an Ultrium 1 tape surface. Visible are the magnetic particles (which clearly are oriented), lubrication pools (darker regions), and embedded aluminum oxide particles. The binder is composed of polyurethane, polyvinyl chloride, and other materials. The aluminum oxide particles are added to increase coating durability and strength and to clean the head. They presumably also play a role in the wear of the head. Tape manufacturers seek a balance in which there is enough aluminum oxide to accomplish the desired effects, but not so much that heads wear out rapidly.

Wear and durability have been assessed using test drives specially equipped with data-acquisition electronics. Heads are stationary. Readers are fully biased, and their resistances read every tape cycle. This is a valuable tool for gaining an understanding of head tribology. Heads are periodically removed and characterized using an atomic force microscope and stylus profilometer and are also tested for write and read losses. Read sensor resistances plotted against time are shown in **Figure 14** for an




Figure 15

Atomic force micrograph of Ultrium 2 head (a) before and (b) after a 24-hour tape run. TiC grains initially are higher than the surrounding aluminum oxide but are recessed after tape wear. Tape removes lapping marks and retextures the head surface. The region shown is part of an AMR read sensor. Hard-bias permanent magnet leads protrude slightly after wear.

Ultrium 1 head run against LTO media streaming at 4 m/s under ambient conditions. Resistance changes typically reach asymptotic limits in less than 1000 hours of tape motion. Shield and sensor recessions produced by Ultrium 1 media have been found to be typically less than approximately 25 nm. Magnetic recording tape has been observed to redistribute Permalloy at the head-tape interface under certain conditions. This can temporarily affect reader resistances. Head processing procedures that minimize susceptibility to the formation of parasitic electrical contacts between the MR leads and S2 shields are employed (see, for example, [21]).

Benign selective wear of the TiC component in the AlTiC material is observed. The change in surface profile is surprising; after lapping, the TiC grains are typically elevated 2 to 6 nm above the surrounding aluminum oxide, which is amphoteric and appears to etch more rapidly than TiC during processing. After tape wear, the grains are recessed by as much as 25 nm. The wear mechanism may be shearing of the TiC grains by tape asperities [22]. There is also informal speculation that running tape oxidizes the TiC grains, forming TiO_2 and CO_2 . Then, the TiO_2 would be quickly removed by the

tape asperities. **Figure 15** shows atomic force microscope (AFM) profiles of an Ultrium 2 head before and after 24 hours of wear. The TiC grains are evident, as are the recording elements.

Error rates

Heads are also tested in actual drives, where they write and read entire data cartridges. Bit (C1) and correlated errors are monitored. C1 error rates in a drive running virgin media tend to improve slightly between 0 and 50 full-file passes (writing an entire cartridge is one full-file pass), as the data in **Figure 16** shows. This is due to a small decrease in head-tape spacing. Media and head skiving edge wear contribute to this effect. After 100 full-file passes, improvement when a new tape is introduced is generally less, because head skiving edge sealing reaches an asymptotic value between 50 and 100 full-file passes and because worn heads do not fully burnish the tapes. The media burnishing effect appears to be more than just wearing of tape surface asperities. AFM scans of worn media indicate an overall flattening of the tape surface. Media have shown as much as 15-nm reduction in head-tape magnetic spacing.




Figure 16

Bit (C1) error rates for an IBM Ultrium 2 production drive. The drive ran an IBM Ultrium 2 cartridge at 6.22 m/s. Errors were counted and reported every full-file pass. Error rates tended to improve slightly during the first 50 full-file passes as a result of subtle improvements in head-tape spacing, as explained in the text. Courtesy of Jason Liang and Stanley Czarnecki, IBM San Jose.




Figure 17

2T amplitudes for individual tracks recorded and read back at 6.22 m/s, normalized by the same at 2.33 m/s for the speed-optimized Ultrium 2 head. The ratios are expressed in decibels and plotted against the 6.22-m/s 2T amplitude. Data for several heads is shown. The amplitude loss of 0.2 to 0.8 dB is the “round trip,” i.e., write plus read change. At the linear density of 3930 flux reversals per millimeter, this is equivalent to a one-way spacing loss of less than 5 nm. Media wear-in can lessen the effect. Data for the 3590H cylindrical contour (not shown) will exhibit substantially greater loss at higher speeds.

Error rates can be negatively affected by a noise source that has plagued metal-shielded MR tape heads, even those used in the 3590H heads where only S2 is metal (the nonconductive, but magnetic, ferrite wafer itself is used for S1). To understand this, one need only remember that, unlike disk-drive media, tape rubs on the metallic sensor shields and so can tribocharge the sensor-shield parallel-plate capacitors. As it happens, charging is typically not disruptive, but modeling of discharging through surface “shorts” shows that intense, single-bit noise voltage spikes appear across the MR leads as up to 50% of the stored charge surges through the sensor. The currents are not destructive, but the pulses can fire at a rate of 1 kHz or more and degrade error-detection margins.

IBM tape heads employ a novel potential clamp which, for LTO heads, is a center-tapped thin-film serpentine tantalum resistor, the ends of which are electrically connected to the MR leads [23, 24]. The center-tap connects to both S1 and S2 shields. Each leg is approximately 70 kΩ. The resistors clamp the shields to the MR mid-potential and so prevent charge buildup. The potential clamp was first employed in the 3590B product as a single resistor connecting to the most positive lead, and it completely eliminated the noise spikes.

Speed matching

Along with the demand for greater cartridge storage capacity comes the desire to maintain or even reduce the time it takes to write a full tape. Adding channels and increasing tape speed are common methods of improving the data rate. Each has its own set of challenges, but it is possible for a drive to get ahead of its host by running at too high a data rate. The IBM Ultrium 2 drive is designed to be compatible with a wide range of host systems and adjusts data rate by changing tape speed from 3 m/s to more than 6 m/s. The Ultrium 2 drive employs PRML detection, which is sensitive to changes in high-frequency roll-off, especially compared with peak detection, used in the Ultrium 1 product. The Ultrium 2 head design was optimized for minimum sensitivity to tape speed changes. **Figure 17** shows the ratio of two-transistor (2T) read-back amplitudes at 2 m/s and 6 m/s, measured using IBM-built equipment. The change is less than 1 dB for signals recorded and read back at each speed. This corresponds to spacing changes of less than 5 nm at the Ultrium 2 linear density of 3900 flux reversals per millimeter, assuming comparable read and write losses [25, 26].

Summary and outlook

HDD thin-film wafer and machining processes have been adapted for the volume production of flat-profile LTO tape heads. Aluminum oxide–titanium carbide wafers and flat tape-contacting surfaces give a durable head–tape interface. The flat profile has been refined for controlling fly height over a range of 2 m/s to more than 6 m/s. (The head has been shown to operate at 8 m/s and higher [11]). A narrow-read-track flat head was used for writing one terabyte of data on an LTO cartridge having advanced-formulation Fujifilm** media [17, this issue]. The HDD thin-film technology has enabled the production of more than 100 000 Ultrium 1 drives in less than two years. The future of the LTO family is dictated in large measure by improvements in media magnetic performance and stability. Increases in track density will require further reduction of the simultaneously written track span. Head track pitch can be reduced by another factor of 3 or so using conventional HDD thin-film head processes. Tighter pitch may require barber-pole or other designs. The future of tape heads based on this technology is promising.

Acknowledgments

The authors acknowledge all who helped pioneer flat-head technology: Managers Erin Keeley and John Teale; David Seagle, and later Peter Koeppel, for head designs and wafer layout; Jane Nealis, for wafer processing; Calvin Lo, for wafer machining; Richard Carlson, for prototype head builds; Reid Anderson, for early work on durability; Leif Kirschenbaum, for tape head quasi-test; Timothy Iben, for MR sensor physics and reliability; David Harper, for

actuator design; Glen Jaquette, for LTO architecture and timing-based servo; Eric Christensen, for head integration; and Judd McDowell, for servo work. They also thank James Smith, David Griesel, and Richard Chacon for head testing support; and Nancy Jubb for empowering the manufacturing team to ramp up LTO head volume production.

*Trademark or registered trademark of International Business Machines Corporation.

**Trademark or registered trademark of Fuji Photo Film U.S.A., Inc.

[†]Ultrium, Linear Tape-Open, and LTO are registered trademarks of International Business Machines Corporation, Hewlett-Packard Corporation, and Seagate Corporation in the United States, other countries, or both.

References

1. D. M. Cannon, W. R. Dempwolf, J. M. Schmalhorst, F. B. Shelledy, and R. D. Silksen, "Design and Performance of a Magnetic Head for a High-Density Tape Drive," *IBM J. Res. & Dev.* **30**, No. 3, 270–277 (May 1986).
2. H. F. Hinteregger and S. Muftu, "Contact Tape Recording with a Flat Head Contour," *IEEE Trans. Magn.* **32**, No. 5, 3476–3478 (September 1996).
3. Standard ECMA-319, 12.7 mm–384 Track Magnetic Tape Cartridges—Ultrium 1 Format, June 2001; see <http://www.ecma-international.org/>.
4. D. S. Fields, Jr., "Two-Module-Read, Read-After-Write, Bi-Directional Tape Drive," U.S. Patent 4,685,005, August 1987.
5. R. Simmons, "HDD Technologies Key to Ultrium Tape," *Data Storage*, pp. 25–28 (June 2000).
6. R. G. Biskeborn and J. H. Eaton, "Partial Width Mass Produced Linear Tape Recording Head," U.S. Patent 5,883,770, March 1999.
7. R. G. Biskeborn and J. H. Eaton, "Method for Mass Producing Partial Width Linear Tape Recording Head," U.S. Patent 6,341,416, January 2002.
8. N. Robertson, H. L. Hu, and C. Tsang, "High Performance Write Head Using NiFe 45/55," *IEEE Trans. Magn.* **33**, No. 5, 2818–2820 (September 1997).
9. R. G. Biskeborn and J. H. Eaton, "Bidirectional Flat Contour Linear Tape Recording Head and Drive," U.S. Patent 5,905,613, May 1999.
10. R. G. Biskeborn and J. H. Eaton, "Flat Profile Tape Recording Head," *IEEE International Magnetism Conference (Intermag Europe 2002)*, 2002, Digest of Technical Papers, AP14.
11. R. G. Biskeborn and J. H. Eaton, "Flat-Profile Tape Recording Head," *IEEE Trans. Magn.* **38**, No. 5, 1919–1921 (September 2002).
12. R. G. Biskeborn, A. P. Deshpande, C. S. Lo, and A. J. Torres, "Precision Bonding Method and Apparatus," U.S. Patent 6,468,376, October 2002.
13. R. G. Biskeborn, "Method and Apparatus for Providing Tape Head Assemblies for Very High Density Recording," U.S. Patent Application Publication, Pub. No. U.S. 2003/0039069 A1, February 2003.
14. R. G. Biskeborn, A. P. Deshpande, C. S. Lo, K. T. Luong, A. J. Taina, and A. J. Torres, "Thin-Film Tape Head Processes," U.S. Patent Application Publication, Pub. No. U.S. 2003/0039070, February 2003.
15. R. G. Biskeborn, "Tape Head Module Assembly System and Method," U.S. Patent Application Publication, Pub. No. U.S. 2003/00339077, February 2003.
16. A. J. Argumedo, R. G. Biskeborn, E. R. Childers, and J. H. Eaton, "Tape Movement Constraint and Method of Constraining Movement of a Tape," World Intellectual Property Organization, International Application, International Publication WO 02/15181 A1 (February 2002).
17. E. R. Childers, W. Imano, J. H. Eaton, G. A. Jaquette, P. V. Koeppe, and D. J. Hellman, "Six Orders of Magnitude in Linear Tape Technology: The One-Terabyte Project," *IBM J. Res. & Dev.* **47**, No. 4, 471–482 (2003, this issue).
18. R. G. Biskeborn, "Signal-Based Tape Wrap Angle Setting for Flat Contour Linear Tape Recording Heads," U.S. Patent Application Publication, Pub. No. U.S. 2003/003045 A1, February 2003.
19. R. G. Biskeborn, C. Y. Inouye, D. J. Seagle, A. J. Wallash, R. O. Barr, G. A. Garfunkel, and S. J. Lewis, "Lapping Process for Minimizing Shorts and Element Recession at Magnetic Head Air Bearing Surface," U.S. Patent 5,603,156, February 1997.
20. R. O. Barr, A. M. Desouches, J. P. Herber, R. G. Biskeborn, C. Y. Inouye, D. J. Seagle, A. J. Wallash, G. A. Garfunkel, and S. J. Lewis, "Lapping Process for Minimizing Shorts and Element Recession at Magnetic Head Air Bearing Surface," U.S. Patent 5,735,036, April 1998.
21. E. Sourty, J. L. Sullivan, and M. D. Bijker, "Numerical Modeling of Sub-Surface Stress in Magnetic Data Tape Heads Due to the Dynamic Contact with a Tape," *Tribology International* **35**, 171–184 (2002).
22. R. G. Biskeborn, "Method and Apparatus for the Manufacture of Thin Film Magnetic Transducers Using a Compliant, Soft Lapping Process," U.S. Patent Application Publication, Pub. No. U.S. 2003/0006213 A1, January 2003.
23. R. G. Biskeborn, "Shielded Magnetoresistive Head with Charge Clamp," U.S. Patent 6,246,553, June 2001.
24. R. G. Biskeborn, T. Y. Chen, M. J. Doscher, P. V. Koeppe, and J. D. Jarratt, "Serpentine Resistive Shunt for MR Heads Having Conductive Shields," U.S. Patent 6,344,952, February 2002.
25. R. L. Wallace, "The Reproduction of Magnetically Recorded Signals," *Bell Syst. Tech. J.* **30**, 1145–1173 (October 1951).
26. H. N. Bertram, *Theory of Magnetic Recording*, Cambridge University Press, Cambridge, England, 1994, pp. 213–219.

Received March 24, 2003; accepted for publication April 14, 2003

Robert G. Biskeborn *IBM Systems Group, 5600 Cottle Road, San Jose, California 95023 (biskebor@us.ibm.com).* Dr. Biskeborn received his Ph.D. degree in physics in 1978 from Columbia University. He joined an IBM packaging technology group in East Fishkill in 1978 and later accepted a position in an advanced bipolar device technology group. Among his accomplishments in East Fishkill were critical contributions to the IBM System/390 thermal packaging technology. He transferred to the IBM Storage Systems Division in San Jose in 1990 and has contributed to both hard disk and tape storage technologies there. Dr. Biskeborn is a member of the American Physical Society.

James H. Eaton *IBM Research Division, Almaden Research Center, 650 Harry Road, San Jose, California 95120 (jimeaton@us.ibm.com).* Dr. Eaton received his B.S. degree in engineering physics in 1958 and his Ph.D. degree in electrical engineering from the University of California at Berkeley in 1962. He was Assistant Professor of Electrical Engineering there until joining the IBM Research Division as a Research Staff Member in 1964. In 1971 he moved to the IBM Thomas J. Watson Research Center, Yorktown Heights, New York, as Director of Applied Research. Before returning to the IBM Research San Jose Laboratory in 1983, Dr. Eaton held several management positions, including Director of the Poughkeepsie Laboratory, Director of the East Fishkill Laboratory, Director of the IBM Corporate Technology Committee, and Director of Research for the Storage Products Division Magnetic Recording Institute. He has worked in magnetic tape recording systems, returning to full-time technical work in 1988. Dr. Eaton holds more than 20 patents; he has received an IBM Corporate Award and IBM Outstanding Contribution Awards for his work in magnetic recording.