Six orders of magnitude in linear tape technology: The one-terabyte project

E. R. Childers
W. Imaino
J. H. Eaton
G. A. Jaquette
P. V. Koeppe
D. J. Hellman

For the last 50 years, tape has persisted as the media of choice when inexpensive data storage is required and speed is not critical. The cost of tape storage normalized per unit capacity (dollars per gigabyte) decreased steadily over this time, driven primarily by advances in areal density and reduction of tape thickness. This paper reports the next advance in tape storage—a demonstration of a tenfold increase in capacity over current-generation Linear Tape-Open® (LTO®) systems. One terabyte (1 TB, or 1000 GB) of uncompressed data was written on half-inch tape using the LTO form factor. This technical breakthrough involves significant advances in nearly every aspect of the recording process: heads, media, channel electronics, and recording platform.

Introduction

The key players in the computer memory and storage hierarchy (from fast and expensive to slow and inexpensive) are semiconductor memory (SCM), hard disk drives (HDDs), and tape. In low-end and consumer products, optical storage plays a major and increasing role in storage, publishing, and interchange. Since the commercial introduction of SCM chips, their capacity has increased five orders of magnitude, HDD areal density has increased seven orders of magnitude, and tape areal density has increased five orders of magnitude. These improvements tracked one another well enough for all three technologies to remain key players for 30 to 50 years. During that time, other technologies disappeared: drums were superseded by lower SCM cost; punch cards were made obsolete by the advantages of tape storage, by terminals for data input, and by improved HDD performance; and floppy disks are being supplanted by optical devices for data interchange. The role once played by tape included data interchange, processing, and storage; today it serves primarily as a media for archiving and backup and restore. Its narrowing role is being driven by changing relationships in cost and performance among other members of the hierarchy. For the last few years, HDD cost has declined at a faster rate than tape, driven largely by rapid improvements in areal density. To prevent a further narrowing of the role played by tape, tape system costs must keep pace with HDD costs.

Relative costs are driven largely but not exclusively by areal density. While areal density has increased by about two and a half orders of magnitude from the mid-1980s to today, the price of a half-inch tape cartridge has increased from $20 at the time of its introduction in the 1980s to $100 today. This is partly because consumers are turning to optical media for audio and video, thereby reducing the demand for tape, and also because of the growing technical requirement for high-density data tape.

Historically, tape has had to bear the burden of legacy products; it had to remain compatible with previous generations. SCM and HDD faced no such burden. This has slowed the rate of improvement in tape. For example, nine-track compatibility restricted density advances solely to linear density. It took an order-of-magnitude improvement in density as well as improvements in reliability to make the 18-track IBM 3480 drive a success. It took a similar order-of-magnitude improvement to introduce the

0018-8646/03/$5.00 © 2003 IBM

Table 1 LTO Gen 1 operating point and related scaling factors to achieve the LTO Gen 2 and 1-TB demonstration operating points.

		Gen 1	*Gen 2*	*1-TB demonstration*
Linear density (bits/mm)		**4880**	**7398**	**12 330**
	Gen 1 multiplier	1.00	1.5	2.53
	Gen 2 multiplier		1.00	1.67
Tracks		**384**	**512**	**1216**
	Gen 1 multiplier	1.00	1.33	3.17
	Gen 2 multiplier		1.00	2.38
Tape thickness (μm)		**8.9**	**8.9**	**6.9**
Tape length (m)		**600**	**600**	**775**
	Gen 1 multiplier	1.00	1.00	1.29
	Gen 2 multiplier		1.00	1.29
Capacity (GB)‡		**100**	**200**	**1000**
	Gen 1 multiplier	1.00	1.99	10.2
	Gen 2 multiplier		1.00	5.11
Min track width (mm)		**0.0275**	**0.0202**	**0.0082**
	Gen 1 multiplier	1.00	0.735	0.297
	Gen 2 multiplier		1.00	0.405
Max raw areal density (Mb/in.^{2})		**114.49**	**236.27**	**970.10**
	Gen 1 multiplier	1.00	2.06	8.5
	Gen 2 multiplier		1.00	4.12

‡ Minimum specified capacity; actual capacity is somewhat greater (e.g., 5%).

noncompatible IBM 3590 drive. In view of the legacy burden of tape, the LTO† layout was designed to allow backward-compatible improvements up to ten times the capacity at introduction.

The progress in tape has resulted from substantial channel, head, drive, and tape advances. The IBM 3480 introduced magnetoresistive (MR) read heads [1]. The IBM 3570 and 3590 introduced high-bandwidth tracking and broke the channel-per-track-on-tape relationship to allow many tracks on a half-inch tape. These products featured high-density metal-particle tape. The IBM 3570 also introduced a timing-based servo [2] in order to achieve high reliability and ease of migration to higher track densities. The LTO format reduced the active head span to less than one quarter of the tape width to permit future advances in track density while retaining backward compatibility. The increasing trend to tape automation can ease backward-compatibility requirements. The terabyte demonstration shows that the rate of progress in tape density can be accelerated. The next advance will require solving substantial engineering problems, but it will not require fundamental breakthroughs in the physics of magnetic recording. Tape is well positioned for another order-of-magnitude improvement beyond the 1-TB demonstration.

Achieving the operation point

The 1-TB-capacity operation point is defined as the combination of track pitch, linear density, and tape length that provides a native capacity of 1 TB of data in a half-inch tape cartridge. Owing to the wide variety of half-inch tape cartridge formats, the 1-TB-capacity operating point can best be discussed by referencing an existing standard: the LTO Ultrium† Generation (Gen) 1 format, which, for this discussion, is 100 GB.

The 1-TB operating point is achieved by appropriately scaling the linear density, track pitch, and tape length from LTO Gen 1. **Table 1** shows the LTO Gen 1 operating point and related scaling factors to achieve the LTO Gen 2 and the 1-TB demonstration operating points. The table includes the LTO Gen 2 operating point to show the jump in linear density between LTO Gen 1 and the LTO Gen 2. This jump is achieved primarily by implementing a partial response maximum likelihood (PRML) channel. PRML detection is common in HDD but has arrived in advanced linear tape drives in only the last year or two. Because a PRML channel is a base assumption for the 1-TB operating point, the table includes the LTO Gen 2 operating point and scaling multipliers to help clarify how the technology is advancing.

The 1-TB operating point was evaluated and demonstrated by first defining an acceptable tape length so that the options for linear density and track pitch could be assessed. Because the goal of this effort was to demonstrate a technology operating point that could be reasonably achieved, it was important to select the tape length conservatively. For this reason, we chose a tape thickness currently being produced by media manufacturers—6.9 μm. The tape length was then determined by assuming use of the LTO Gen 1 reel, the smallest reel of any current half-inch tape format. Tape length could have been increased by increasing the size of the tape reel, but this would have obscured the objective of the technology demonstration.

After a tape length was selected, the operating point was determined by appropriately scaling track pitch and linear density. Table 1 shows the relationship between the 1-TB operating point used in the demonstration and the Gen 1 and Gen 2 operating points. The 1-TB demonstration had a linear density more than 2.5 times that of Gen 1 and more than 1.6 times that of Gen 2. The 1-TB system recorded more than three times the number of tracks compared with Gen 1 and more than twice the number of tracks of Gen 2.

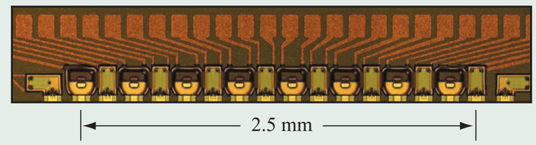

Figure 1

Side-by-side configuration of eight pairs of writer and reader transducers, spanned by two servo reader transducers for track following. Two such eight-channel modules placed face to face comprise a head providing bidirectional read-after-write capability.

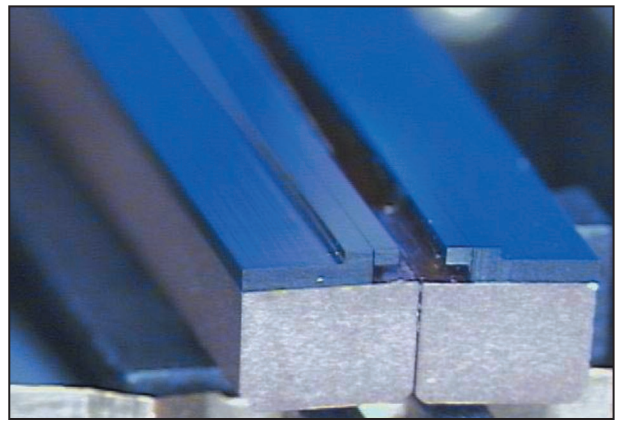

Figure 2

End view of the recording head used in this study. Two face-to-face modules with flat tape-bearing surfaces provide eight-channel data flow in both tape directions.

Head design

To support the 1-TB-per-cartridge capacity, a matrix of heads targeting multiple combinations of linear and track densities were designed and fabricated. The heads were designed to accommodate magnetic linear densities up to 15000 bits/mm and/or track densities up to 150 tracks/mm (tr/mm), using as selectable design parameters combinations of reader gap lengths and reader track widths, respectively. The heads are eight-channel devices, appropriate for use in LTO Ultrium-format tape systems. The heads, which span approximately one quarter of the tape width at a time, comprise a side-by-side configuration of interleaved writer and reader transducers (eight each, spanned by two servo reader transducers, as shown in **Figure 1**) on each of two identical face-to-face modules. As shown in **Figure 2**, a flat profile is used at the head–tape interface [3].

Writer transducers are 12-turn thin-film inductive elements. Writer track widths were designed to provide for writing a continuous range of track densities beginning with a track density of approximately 36 tr/mm in the standard LTO Gen 1 (100-GB capacity) and extending to the much higher track densities (up to 150 tr/mm) for this study, which were achieved by a technique called *shingling*—overwriting a fraction of the previously written track. The writer transducers were fabricated with the standard thin-film processes used for heads in HDDs. The track-defining top pole (P2) is made of plated $Ni_{45}Fe_{55}$ [4]. It is 4 μm thick and easily capable of writing on the 2400-oersted (Oe) tape used in this study (33-dB overwrite at 14800 bits/mm).

Reader transducers are conventional anisotropic magnetoresistive (AMR) sensors [5, and references therein]. They are transverse-biased using soft-film biasing and are longitudinally stabilized using hard biasing in the tail regions. To achieve the desired linear density range, reader gaps were fabricated in the range of 0.17 to 0.25 μm shield-to-shield. These were matrixed with reader track widths in the range of 2.0 to 5.5 μm based on modeling of track misregistration for the track densities of interest. The various reader-transducer designs were tuned to the magnetic characteristics of the high-performance tape used for this study, providing high-amplitude, low-distortion signals. All reader transducers incorporated resistive shunts between the MR head and its shields, eliminating unwanted electrical interference from head–tape interactions (e.g., electrical discharges) [6].

Media

The media used for the 1-TB demonstration was an experimental tape from Fujifilm** constructed with advanced metal particles possessing improved particle size distribution and coercivity. Metal particles with an effective coercivity (the opposing magnetic force necessary to remove residual magnetism of a material) of more than

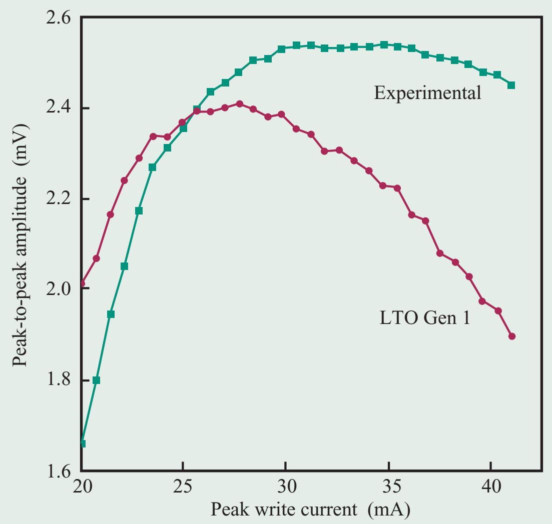

Figure 3

Write saturation response of the media used for the 1-TB demonstration, labeled “Experimental.”

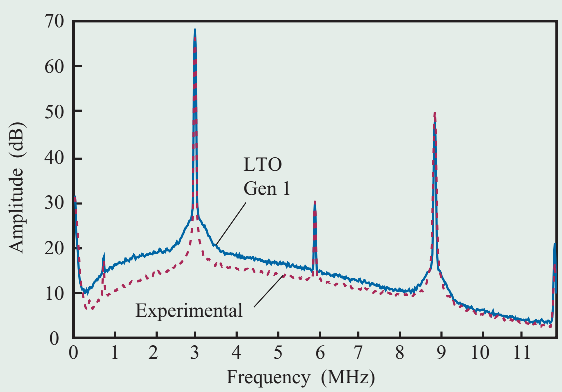

Figure 4

Spectral response of the media used for the 1-TB demonstration, labeled “Experimental.”

2300 Oe and a mean particle length of 60 nm (8000 nm^3 volume) were dispersed in a glassy polymer binder system of modified polyester–polyurethane. The magnetic layer was coated uniformly to a coating thickness of 0.1 μm as a dual-layer tape construction with a nonmagnetic underlayer, resulting in a final arithmetic mean surface roughness (R_a) of less than 2 nm. This surface roughness is slightly more than half that of the LTO Gen 1 media. The tape employs a conductive backcoat that produces a total tape thickness of 6.9 μm when coated on a 5.2-μm-thick polyethylene naphthalate (PEN) substrate. Conventional fatty-acid ester lubricants were used to provide acceptable durability and “runability.”

Mechanical behavior is also important. The tape exhibited a composite machine direction storage modulus (E') of seven gigapascals (GPa) at 20°C with a weak glass transition temperature (T_g) at 30°C, similar to other media with PEN substrates. The subtracted magnetic coating (both recording and underlayer) modulus was determined to be 9.5 GPa at 20°C, with a better-defined T_g at 25°C.

The recording behavior is shown in **Figures 3** and **4**. To determine the optimum recording current for the media, the write saturation response (Figure 3) was determined by an electromagnetic recording test. In this test, the readback response is measured as the write current is changed. This test was conducted at the LTO Gen 2 operating point (7400 bits/mm) with an IBM LTO Gen 2-prototype head having MR read elements. The write saturation response is directly related to the coercivity and remanence (amount of magnetic induction remaining in a magnetic material after the material has been removed from an external magnetic field) of the tape particles and the thickness of the tape’s magnetic coating. The results of this test indicate that the amplitude is low enough that the response will remain linear in the MR performance regime. After the optimal write current was determined, the spectral response (Figure 4) was used to determine the signal-to-noise ratio (SNR). The results shown here used the LTO Gen 2 method of recording and measurement. These results show that there is a 4.5-dB broadband SNR increase between the LTO Gen 1 media and the media used for the 1-TB demonstration. This is important because it shows that with the 1-TB demonstration media it is possible to halve the read track pitch relative to the LTO Gen 2 drive, yet have no loss in SNR.

Deck design

To demonstrate the viability of extending current drive technology, the 1-TB system drive uses many components directly from the IBM LTO tape drive. These include reel-to-reel motors and servo electronics, head actuator and electronics, surface-control-guiding grooved rollers, and cartridge clutch and reels. However, there are key differences. For example, the 1-TB track pitch (8.2 μm) is 3.36 times smaller than the LTO Gen 1 track pitch, and the written-track width minus the reader width is 2.8 μm—five times smaller than LTO Gen 1. These differences required better tracking than that provided by the LTO Gen 1 drive. Tracking improvement was achieved by modifying the tape path. **Figure 5(a)** is a schematic view of the tape path in the 1-TB system. Rollers R1 and R6 have flanges but not grooves, rollers R2 and R5 have neither flanges nor grooves, and rollers R3 and R4 have both flanges and grooves. **Figure 5(b)** is a photograph of the 1-TB system. Several of the modifications shown in the figure are described below.

Grooved rollers

The grooved rollers allow the air in the self-acting roller air bearing to bleed out so that the tape rolls on the roller. Measurements show that a drag torque of approximately 1 newton-mm (N-mm) is required for an appropriately designed grooved roller to drop out of synchronism with the tape, while the non-grooved rollers generally turn much more slowly than synchronous velocity.[1] The lateral motion of tape rolling on the grooved roller is given by

$$\textit{lateral motion} = V \sin \varphi,$$

where φ is the angle between the plane normal to the axis of the roller and the edge of the tape. Thus, high-frequency lateral disturbances generated by the reels or outer rollers are highly damped in passing over the grooved roller. The rollers must be well aligned because, if φ is too large, the tape is forced against the flange on the grooved rollers and can generate disturbances. Tension, wrap angle, groove, and land width are key considerations in grooved roller design. The well-known foil-bearing formula [7] describes the height (h) at the nip of an infinitely wide foil bearing and, without side leakage, the height remains fixed until near the exit:

$$h = KR(12\mu VW/T)^{2/3}.$$

Here h is air-film thickness (the film of air trapped between the tape and the bearing), R is roller radius, μ is air viscosity, T/W is tension per unit width, and K is a constant, 0.643. The grooves provide side leakage to let the air escape, reducing flying height so that the tape can contact asperities on the roller surface. Modeling the roller as a long narrow land yields

$$h(\theta) = h(0)/(1 + 2h(0)^2 T\theta/\mu W^3 V)^{1/2},$$

where θ is the angle from the tape tangency point. For sufficiently large θ, the air-film thickness is inversely proportional to land width. These equations provided guidance for choosing the 175-μm land width. **Figure 6** presents the tape–roller spacing surrounding one of the grooves, as simulated by R. G. Biskeborn using TapeLab2.** Note the broadening of the low spacing band as the tape progresses from the entrance to the exit region. The rollers used have a slight rounding of the land edges to keep the contact pressure from growing too large. Referring to Figure 5, rollers R2 and R5 were added to provide a sufficiently large wrap angle for good tape-to-roller contact while providing added distance between the supply reel and the first flanged roller.

[1] Tangential velocity of roller equals linear velocity of tape.

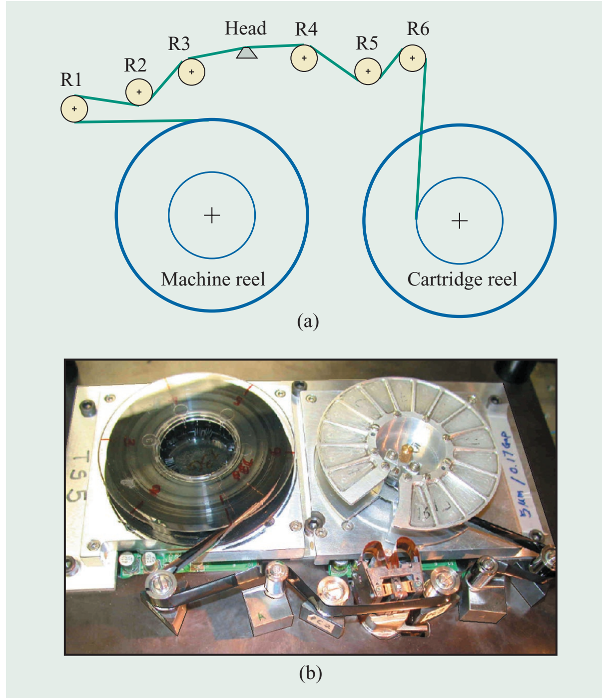

Figure 5

(a) Schematic of the tape path in the 1-TB system. Rollers R1 and R6 have flanges, but not grooves, rollers R2 and R5 have neither flanges nor grooves, and rollers R3 and R4 have both flanges and grooves. (b) Photograph of the 1-TB drive.

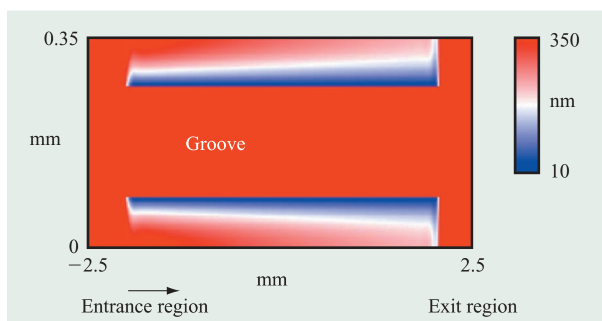

Figure 6

Tape–roller spacing surrounding one of the grooves, shown as simulated by R. G. Biskeborn using TapeLab2. Note the broadening of the low spacing band as the tape progresses from the entrance to the exit region. This fly height is important because asperities on the roller are approximately 30 nm tall; when the spacing is less than 30 nm, the tape contacts the roller, and tape damage can occur.

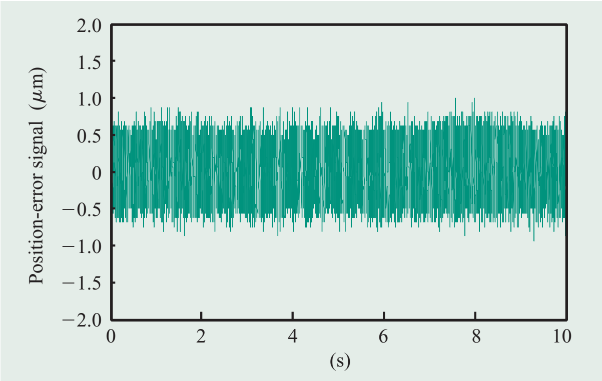

Figure 7

Ten-second segment of the position-error signal at the head during closed-loop operation, demonstrating the tracking achievable with the 1-TB drive system.

Machine reel

At its periphery, the flanges of the standard IBM LTO Gen 1 machine reel can be 0.5 mm wider than the tape. Even with perfect alignment, the tape stacking against one flange or the other causes a 250-μm height misalignment with the flanges of the first roller. Using the Carnegie Mellon University Tape Path Guidance Analysis Program [8] and a tape path with 37 mm from machine-reel tangency to first-roller tangency, this offset produces a force of 0.7 N on the first roller flange and a tape axial compression of 12.7 MPa at the top of the tape at the first roller (compared with an axial tension of 8 MPa at the center of the tape). The model fails when an edge of the tape is in compression and buckles. Increasing the tangency-to-tangency span to slightly more than 60 mm allows the tape to accommodate a 250-μm offset, but the load on the first flange is still high (150 mN, or about 15 g force). For these reasons, a machined aluminum machine reel with improved runout tolerances and flange-to-flange spacing of 220 μm at the periphery was installed. The minimum span from either reel pack to the nearest roller was also increased to 60 mm or more.

Tracking performance

Figure 7 illustrates the tracking achievable with this system. It shows a ten-second segment of the position-error signal (PES) at the head during closed-loop operation running the 6.9-μm-thick 1-TB tape. Some large, physically unrealizable, single-sample errors that would normally be detected and replaced when using a second PES channel have been replaced with the average of adjacent samples. The remaining large PES values cause the system to stop writing until the PES is back in range. Once anomalies have been removed, the 1σ standard deviation is 0.23 μm.

Channel and format

An essential requirement of this technology demonstration was to write the data accurately to the media and be able to read it back. To achieve a 1-TB capacity, several aspects were increased over the LTO Gen 2 operating point of 200 GB. For example, the linear density was increased by a factor of 5/3, the number of tracks was increased by a factor of 19/8, and tape length was

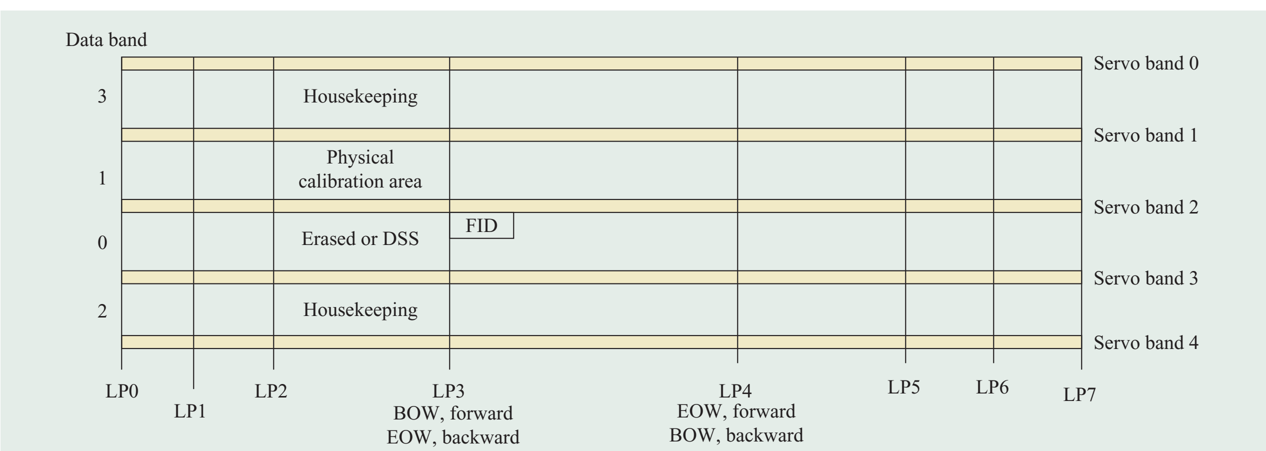

Figure 8

Diagrammatic view of the LTO tape layout into four data bands, each band straddled on each side by a servo band. DSS is *data-set separator*. FID is *format identification data set*, the first data set within the user data area. LP is *linear position*. BOW and EOW are *beginning of wrap* and *end of wrap*, respectively.

increased by a factor of 1.28 by decreasing the thickness from 8.9 μm to 6.9 μm. Multiplying the LTO Gen 2 operating point by these additional factors (200 GB · 5/3 · 19/8 · 1.28) leads to an uncompressed (native) capacity of 1013 GB, a figure greater than the 1-TB goal.

Figure 8 is a diagrammatic view of the tape layout. The tape is divided laterally into four data bands. Each data band is straddled on either side by servo bands for a total of five servo bands. The active part of a write head has to span only a data band and the two servo bands that straddle it, a distance of just under one quarter of the tape width (about 0.125 inch). Each write head has two servo heads, one for each of the two servo bands it spans, and eight write heads equally spaced in the data band. **Figure 9** is a schematic of a closer look at one data band. As the figure demonstrates, the data band is actually composed of eight data sub-bands, one for each of the eight write heads. Each sub-band is written by a given write-head position in a technique called a *linear serpentine*, which means that the tape moves back and forth longitudinally while the head is indexed up or down laterally at each pass. This makes it possible to write multiple distinct tracks in a given data sub-band.

Figure 10 is a schematic of an even closer look at a data band. It shows how 12 data tracks can be written in a converging spiral, which is the linear serpentine form used. The 12 data tracks are written by a given write-head position in 12 passes: six down (tape outbound from cartridge) and six back (tape inbound to cartridge), each at a different lateral offset. Because the heads are eight-channel heads, eight tracks of the data are written simultaneously in this linear serpentine pattern, each track in separate data sub-bands. Once a given data band is full, a coarse-actuator motor moves the head to another quarter of the tape. This process continues until all four data bands are filled.

The data processing scheme described here, substantially an LTO Gen 2 data channel, was used to record 1 TB. **Figure 11** represents this process in the form of a flowchart.

In LTO Gen 2, data records are first compressed using a dual-scheme data-compression algorithm, LTO-DC. This algorithm uses a primary scheme on compressible data, but passes incompressible data through a secondary scheme. This technique prevents the substantial expansion of incompressible data seen in most data-compression algorithms. This type of compression creates a continuous compressed data stream. For the investigation reported here, no compression algorithm was used, which is equivalent to incompressible data being processed by scheme 2 of LTO-DC. Thus, a native capacity of 1 TB was demonstrated, which, if data had been 2:1-compressible, would have given a compressed capacity of 2 TB. (English text, for example, typically compresses at about 2.3:1.) The compressed data stream is then cut into 403884-byte pieces

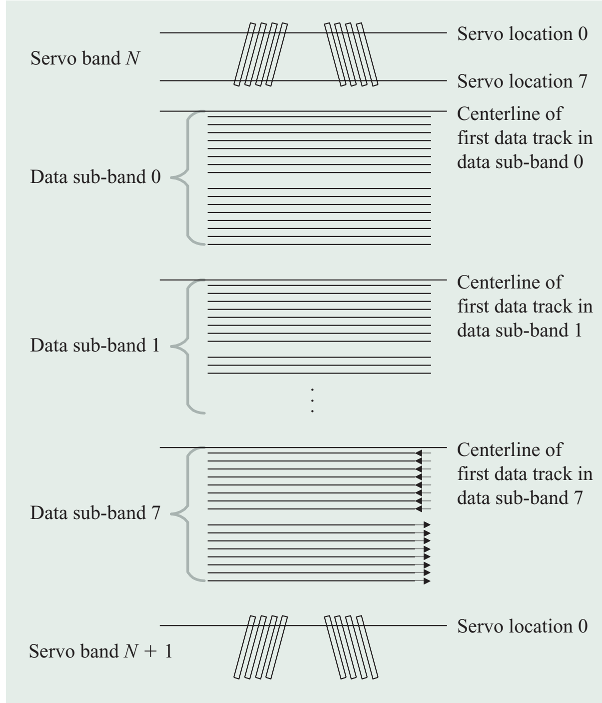

Figure 9

Schematic of a section of tape showing one data band and its surrounding servo bands.

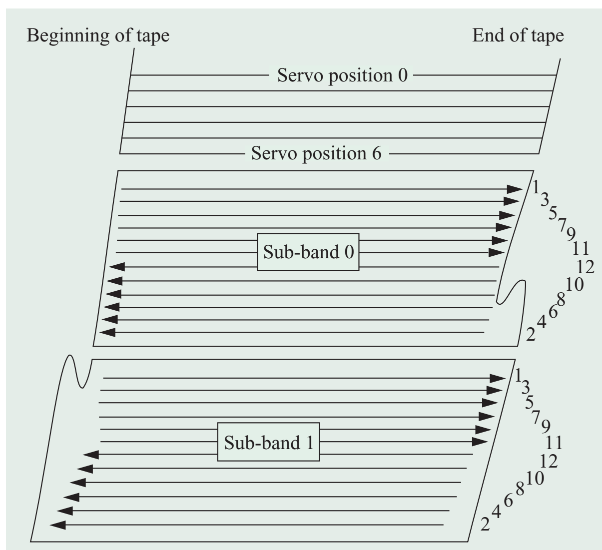

Figure 10

Schematic of a close-up view of two sub-bands and one servo band. This schematic demonstrates the serpentine method used to write data. The numbers on the right side indicate the tracks, which are written simultaneously.

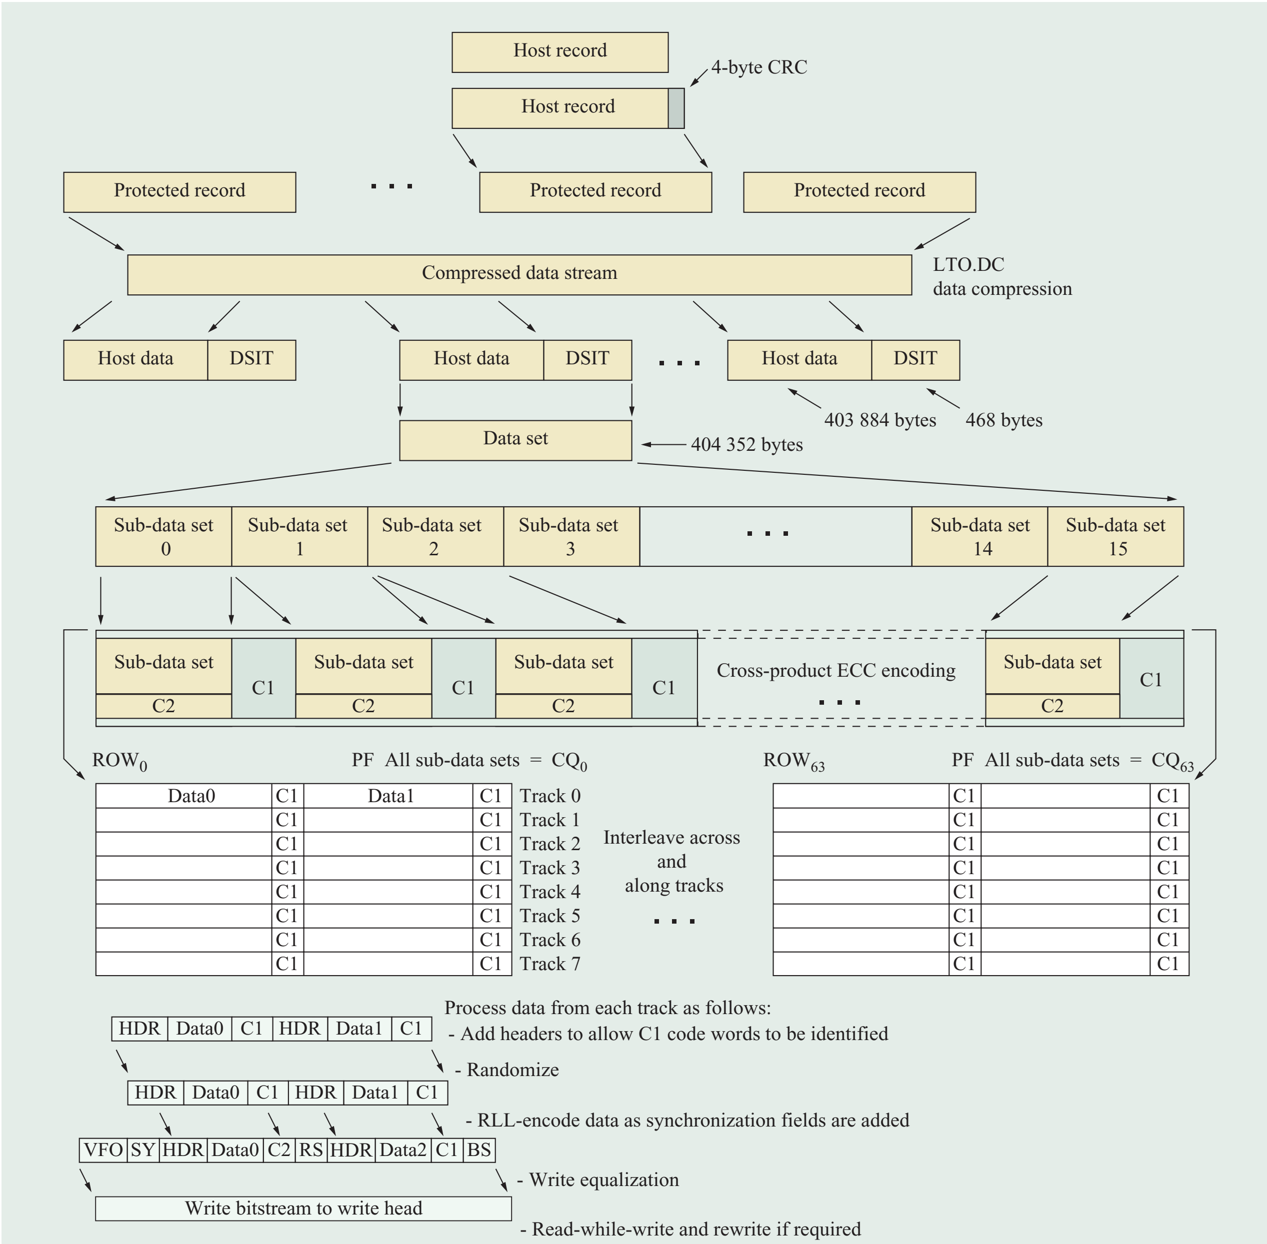

Figure 11

Diagrammatic view of LTO data-flow processing. DS is *data set*, a data region with 403 884 bytes. DSIT is *data-set information table,* which has 468 bytes and describes the content of a *data region* (DR). When a DR is appended to a DS, its length is 404 352 bytes. A *data-set separator* (DSS) appears before each DS and after the last DS on each wrap. C1 and C2 are different Reed–Solomon algorithms for *error-correction code* (ECC). LP is *logical point*. HDR is *header*. BS is *back synchronization*. RS is *ready synchronization*. VFO is *variable frequency oscillator*.

which correspond to the minimum recording entities stored to tape, known as *data sets*. Each data set is error-correction-code (ECC) encoded with a true cross-product ECC code. This ECC code implements an inner code, designed to handle a relatively high background error rate, and an outer code, designed to handle defects, scratches, and any random errors that are beyond the ability of the inner code to correct. The ECC-protected entities are interleaved laterally and longitudinally as they are broken into 64 recording units known as *codeword quads* (CQs).

Each data set is written to tape in 64 CQs. Format features are inserted with each CQ as it is written. The first of these format fields is a variable-frequency oscillator (VFO) field that enables a data-handling VFO to lock to it. A synchronization mark that enables a run-length-limited (RLL) decoder to synchronize its decode to the incoming data stream is then written. After that, a header is written that enables a reading drive to identify which CQ of data is being read and at what lateral rotation.

Finally, the CQs are written to tape. Both the headers and data are first randomized and then RLL encoded with a 16/17 rate (0, 13/11) code. The RLL-encoded bitstream is then write-equalized by a write equalizer, which uses frequency multiplication to create a write clock five times higher than the RLL bit clock, guaranteeing two write-clock periods between current transitions sent to the head. The net effect is to produce a write-current waveform with minimum pulse widths equal to 40 percent of the RLL-bit clock period. While each CQ of data is written by the write head, a read head reads it and checks whether it was written properly or is obstructed by some media defect or scratch. If the ECC finds an error of sufficient magnitude in the detected signal, the CQ is dynamically rewritten farther down tape, rotated laterally so that each CQ is now on a different track. This is called *dynamic rewrite of defective data*. The ability to do this enables LTO Gen 2, and therefore the 1-TB demonstration drive, to perform streamed backup operations even in the midst of media defects, scratches, and high background-error rates.

Writing the uncompressed terabyte

Experimental setup

Figure 12 shows a schematic diagram of the experimental arrangement used in this investigation. A pseudorandom bit-sequence pattern of 8364 bits with the required synchronization and data-set separator (DSS) fields consistent with the LTO format is provided by the SyntheSys BA622/BA16CH1F BitAlyzer** to the prototype electronics card and written to tape. In readback, the head signal is processed by the analog electronics, detected by the PRML channel, and compared with the original data by the BitAlyzer to determine the error rate. The prototype electronics also detect the timing-based servo pattern written on the tape according to the LTO format and servo-controls the actuator for track-following operation. An onboard microprocessor controls the overall operation of the electronics, and a serial RS-232 interface with an IBM ThinkPad* Model 600 laptop computer provides a means of setting various parameters of the channels and/or servo controller and, if necessary, enables real-time adjustments and optimization of the behavior of the electronics. If required, a microprocessor emulator

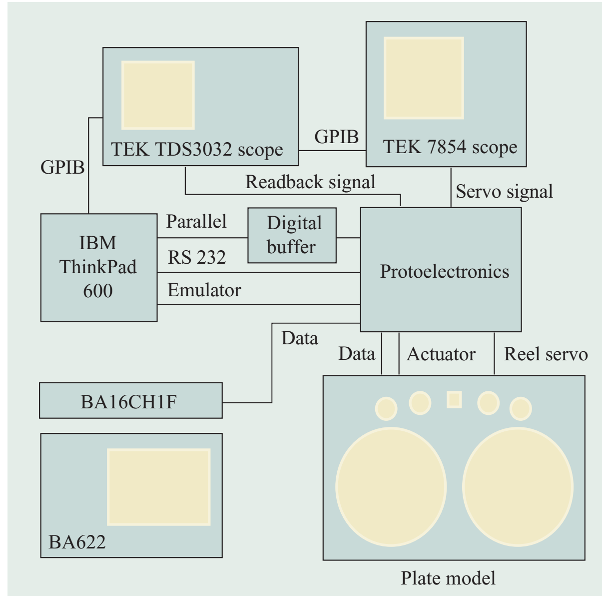

Figure 12

Schematic diagram of experimental setup to read/write 1 TB.

cable may be used to monitor the real-time execution of the processor, and an IEEE 488 interface to the digital scopes that display the readback signals from the data and servo heads can be used to provide a means of data acquisition. In addition, an external digital buffer is used to directly capture sampled readback signals from the data or servo channels.

Results

The 1-TB trace in **Figure 13** shows the sampled readback signal written with a bit-cell length of 76 nm compared with the nominal LTO Gen 2 bit-cell length of 127 nm. This illustrates the increased spatial frequency response of the head–tape combination under study. **Figure 14** shows the decorated image of the written tape compared with a standard Gen 1 tape, graphically verifying the placement of the 8.2-μm tracks relative to the servo bands and showing the relative track pitch compared with that of a Gen 1 tape.

Figure 15 shows the measured error rate as a function of track position. The data show a 9×10^{-5} on-track nonburst error rate with about a 2-μm allowance for track misregistration. The data rate was 30 MB/s, limited by the prototype electronics speed.

Summary

Technology improvements in heads, media, channel electronics, and recording platform that permit a tenfold

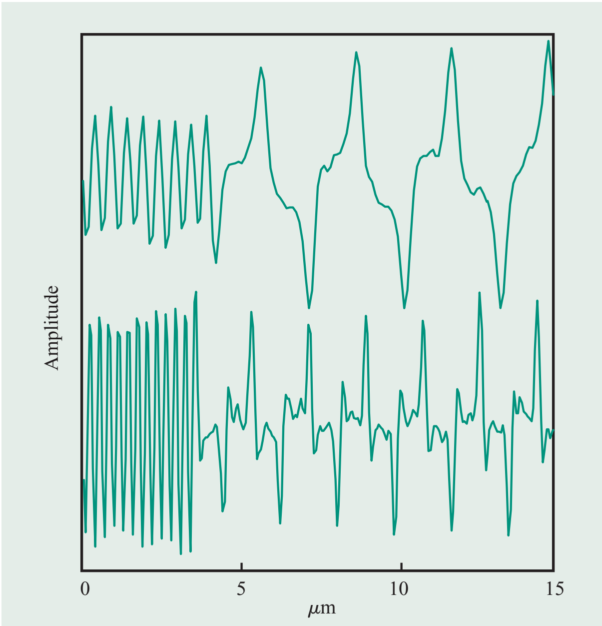

Figure 13

Plot of readback signal with respect to tape position. The lower trace is the data from the 1-TB tape, and the upper trace is the LTO Gen 2 reference.

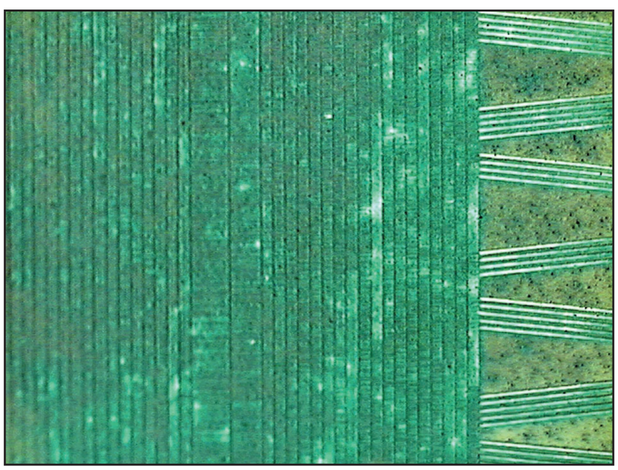

Figure 14

Decorated (ferrofluid) image of written tape.

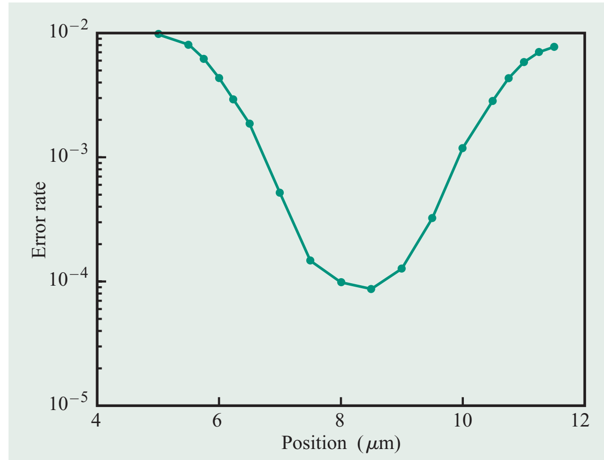

Figure 15

Nonburst error rate as a function of track position.

storage density and capacity increase over LTO Gen 1 have been described. This unprecedented linear tape system achievement illustrates the viability of magnetic tape storage and its continued reduction of cost-per-unit capacity. The evolutionary path to 1 TB is now evident and feasible, and while serious engineering challenges must be overcome to advance tape storage capacity beyond 1 TB, fundamental magnetic recording limits will allow at least another order-of-magnitude improvement.

Acknowledgment

The authors thank J. Teale and R. Bradshaw for their support and contributions to this paper, E. Christensen for his work on the recording behavior of the media, and B. Axtman for his work on the electronics for this endeavor.

*Trademark or registered trademark of International Business Machines Corporation.

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

†Linear Tape-Open, LTO, and Ultrium 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. Silkensen, “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. T. R. Albrecht, R. C. Barrett, E. H. Klassen, J. H. Eaton, and G. A. Jaquette, “Timing-Based Track Following Servo for Linear Tape Systems,” *IEEE Trans. Magn.* **34,** No. 4, 1872–1877 (July 1998).
3. R. G. Biskeborn and J. H. Eaton, “Flat-Profile Tape Recording Head,” *IEEE Trans. Magn.* **38,** No. 5, 1919–1921 (September 2002).
4. N. Robertson, H. L. Hu, and C. Tsang, “High Performance Write Head Using NiFe 45/55,” *IEEE Trans. Magn.* **33,** 2818–2820 (September 1997).
5. C. Tsang, T. Lin, S. MacDonald, M. Pinarbasi, N. Robertson, H. Santini, M. Doerner, T. Reith, L. Vo, T.

Diola, and P. Arnett, "5 Gb/in.^{2} Recording Demonstration with Conventional AMR Dual Element Heads and Thin Film Disks," *IEEE Trans. Magn.* **33,** 2866–2871 (September 1997).
6. R. G. Biskeborn, T. Y. Chen, M. J. Doscher, P. V. Koeppe, and J. E. Jarratt, "Serpentine Resistive Shunt for MR Heads Having Conductive Shields," U.S. Patent 6,344,952, February 2002.
7. A. Eshel and H. G. Elrod, Jr., "The Theory of the Infinitely Wide, Perfectly Flexible, Self-Acting Foil Bearing," *J. Basic Eng.* **87,** 831–836 (1965).
8. J. Wickert and M. Takahiro, NSIV-CMU Tape Path Guiding Analysis Program, Data Storage System Center, Carnegie Mellon University, Pittsburgh, PA 15213.

Received May 16, 2002; accepted for publication October 14, 2002

Edwin R. (Ed) Childers *IBM Systems Group, 9000 South Rita Road, Tucson, Arizona 85744 (erc@us.ibm.com).* Mr. Childers is a Senior Technical Staff Member working in the Tucson Tape Technology and Drives area. He graduated from the University of Texas at Austin in 1978 with a B.S. degree in mechanical engineering and received his M.S. degree in mechanical engineering from the University of Arizona in 1983. He joined IBM in 1978 and worked in printer and optical products prior to moving to tape technology and development. Mr. Childers holds more than ten U.S. patents and has received an IBM Outstanding Technical Achievement Award for his contributions to the creation of the Linear Tape-Open (LTO) product family. He is currently responsible for media and tape mechanism development for IBM tape products.

Wayne Imaino IBM Research Division, Almaden Research Center, 650 Harry Road, San Jose, California 95120 (wayne1@almaden.ibm.com). Dr. Imaino joined the IBM Research Division in 1980 after receiving a Ph.D. degree in physics from Purdue University. He began work on electrophotography and later developed optical instrumentation and tools for hard disk storage, for which he received several IBM Outstanding Technical Achievement Awards. Prior to his current efforts in tape, he worked on various aspects of optical and hard disk storage.

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.

Glen A. Jaquette *IBM Systems Group, 9000 South Rita Road, Tucson, Arizona 85744 (jaquette@us.ibm.com).* Mr. Jaquette is a Senior Technical Staff Member working in the Tape Technology and Drives area. He graduated from University of California at Davis in 1983 with a B.S.E.E. degree and received his M.S.E.E. degree in 1988 from the University of Arizona, with major emphasis in signal processing. He joined IBM in 1983, working on 3380 hard disk drives at IBM San Jose. He transferred to the IBM Tucson facility in 1986 to work on optical disk drives in the read/write channel group. Since 1996 he has been working on tape drives. Mr. Jaquette holds 39 U.S. patents and has eight patents pending. He received an IBM Outstanding Technical Achievement Award for his contributions to the creation of the Linear Tape-Open (LTO) logical format and the electronic architecture of the IBM first-generation LTO tape drive, Ultrium-1. He is currently responsible for the architecture of future IBM tape drives.

Peter V. Koeppe *IBM Systems Group, 5600 Cottle Road, San Jose, California 95193 (pkoeppe@us.ibm.com).* Dr. Koeppe received the B.S. and B.A. degrees in physics and mathematics, respectively, from Hope College in 1982, the M.S. degree in electrical and computer engineering from the University of Cincinnati in 1986, and the Ph.D. degree in electrical and computer engineering from Carnegie Mellon University in 1991, studying and imaging micromagnetic activity in recording heads by means of the Kerr magneto-optic effect. He was a Carnegie Institute of Technology Fellow and an IBM Pre-Doctoral Fellow while at Carnegie Mellon. From 1991 to 1998, he was with the Research Division of Eastman Kodak Company, San Diego, California, where he developed magnetoresistive sensors for a wide variety of applications. In 1998 Dr. Koeppe joined IBM, where he is a Senior Technical Staff Member working in the Tape Head Development group in San Jose, responsible for magnetic head and thin-film wafer designs used in IBM tape drive products. Dr. Koeppe is a member of the Institute of Electrical and Electronics Engineers and the American Physical Society.

Diana J. Hellman *IBM Systems Group, 9000 South Rita Road, Tucson, Arizona 85744 (dianahel@us.ibm.com).* Dr. Hellman is a mechanical and materials engineer working in the Tape Technology and Drives area with an emphasis in media. She received her B.S., M.S., and Ph.D. degrees in 1996, 1998, and 2000, respectively, all from the University of Colorado at Boulder. She joined IBM in 2000 after completing her Ph.D. on developing a novel membrane fabrication process, the Thermally Assisted Evaporative Phase Separation (TAEPS) process. Prior to her current work in media, Dr. Hellman worked on tape library systems, focusing on the materials selection in the libraries.