

Scaling tape-recording areal densities to 100 Gb/in²

We examine the issue of scaling magnetic tape-recording to higher areal densities, focusing on the challenges of achieving 100 Gb/in² in the linear tape format. The current highest achieved areal density demonstrations of 6.7 Gb/in² in the linear tape and 23.0 Gb/in² in the helical scan format provide a reference for this assessment. We argue that controlling the head-tape interaction is key to achieving high linear density, whereas track-following and reel-to-reel servomechanisms as well as transverse dimensional stability are key for achieving high track density. We envision that advancements in media, data-detection techniques, reel-to-reel control, and lateral motion control will enable much higher areal densities. An achievable goal is a linear density of 800 Kbit/in and a track pitch of 0.2 μm, resulting in an areal density of 100 Gb/in².

A. J. Argumedeo
D. Berman
R. G. Biskeborn
G. Cherubini
R. D. Cideciyan
E. Eleftheriou
W. Häberle
D. J. Hellman
R. Hutchins
W. Imaino
J. Jelitto
K. Judd
P.-O. Jubert
M. A. Lantz
G. M. McClelland
T. Mittelholzer
C. Narayan
S. Ölçer
P. J. Seger

Introduction

Historically, the cost and the size of the media recording area of both tape cartridges and disk drives have remained fairly constant. The areal recording density is the main factor determining the cost per gigabyte, which is fundamental to the commercial viability of these technologies. Figure 1 compares the evolution of areal densities of laboratory demonstrations of linear tape technology and commercially available tape drives and hard-disk drives (HDDs). As can be seen, the previously wide gap in areal density values between helical scan and linear scan tapes has almost disappeared. With an areal density two orders of magnitude lower than that of HDDs, tape drives maintain a lower cost per gigabyte only because tape media can be produced at a very low cost per area and because the media is removable, allowing the cost of the tape drive to be amortized over many cartridges. The cost-per-gigabyte difference is the key reason that tape technology has remained viable, even though the research and development investment in tape drives is far lower than that for HDDs. From another point of view, this difference presents a great opportunity for achieving a much lower cost per gigabyte for tape if the areal density can be brought closer to that of HDDs by skillful engineering.

In this paper, we analyze the feasibility and technologies required to achieve a target operating point of 100 Gb/in² in a linear magnetic tape drive. Previous demonstrations of record areal densities (including helical scan and linear tape technologies) and the state of current HDD technology provide important insight into key technological choices and highlight critical parameters that have to be considered for further advances in areal density [1–8]. Figure 2 illustrates the track density versus linear density of HDDs (square symbols) and tape drives (circles), where the linear density takes into account the rate loss due to modulation coding. The product of linear and track densities yields the areal density. The diagonal lines in Figure 2 thus indicate the points of constant areal density. Current tape drives, as specified by the Linear Tape-Open (LTO**) standard for generation 4 (LTO-4) [9], operate at a linear density of ~300 Kbit/in, which is not very far from that of current HDDs. Today, although the gap in linear density between HDDs and tape drives is rather small, tape drives have much lower track density, indicating that there is room for significant improvements. Therefore, in order to reach the operating point of 100 Gb/in², the bit aspect ratio of tapes, usually defined as the ratio of linear density and track density, would need to approach that of current HDDs.

Media considerations
 Areal recording density on tape is limited by medium magnetic stability, maximum achievable write-head fields, the minimum bit length that can be recorded in the medium, and the signal-to-noise ratio (SNR). Broadband SNR varies with the number of magnetic particles per bit volume of the medium. Model calculations show that in order to achieve a reasonable SNR, about 100 particles per bit are required [10]; thus,

Figure 2
 Linear density versus track density.





Figure 1
 Areal density of tape drives and HDDs (reprinted with permission from [1]; ©2007 IEBE). Linear Tape-Open (LTO) is a magnetic tape data storage technology.



$$a_{wc} \approx \frac{\pi Q}{(1 - S^*)(d + t/2)} + \sqrt{\frac{\pi Q}{(1 - S^*)(d + t/2)^2} + \frac{\pi Q H_c}{M_r(d + t/2)}}$$

Here, d is the head-media spacing, t the medium thickness, M_r the remanent magnetization, S^* the hysteresis squareness, Q the head-field gradient, and H_c the medium coercivity. H_c is proportional to the medium anisotropy field H_a and can be approximated for an imperfectly oriented medium at recording timescales by $0.5H_a$ or K_u/M_s , where M_s is the saturation magnetization [15].

Using the above equations, we can estimate the minimum bit length for longitudinal recording as a function of medium thickness, head-media spacing, coercive field, and medium saturation magnetization. Figure 3 presents the minimum bit length achievable, assuming a maximum coercive field of 5,300 Oe (corresponding to B_s of 20 kG, 80% usable deep gap field, and $H_c =$ one third of the deep gap field to account for head-media spacing), and a saturation magnetization of 550 emu/cc. These values correspond to a required particle anisotropy constant of 2.9×10^6 erg/cc and a particle diameter of at least 12 nm. It also indicates that a media can sustain a linear density of 800 Kb/in if the medium is no thicker than 15 nm.

For such a linear density and a medium thickness of 15 nm, the minimum bit area is $32 \text{ nm} \times 380 \text{ nm}$, assuming 100 particles each of 12-nm diameter per bit

as the areal density increases, particles must become smaller. However, in order to avoid thermally induced switching of the particles, the anisotropy energy $K_u V$ of each particle should be much larger than thermal energies, $k_B T$. Here, K_u is the uniaxial anisotropy energy density, V the volume of the particle, k_B the Boltzmann's constant, and T the absolute temperature. A widely used rule to facilitate data retention is $K_u V > 60 k_B T$. Unfortunately, higher-anisotropy particles require larger magnetic fields for magnetization reversal, and head fields are limited by the saturation magnetization (B_s) of the head-pole materials. High- B_s materials allow about 20 kG saturation magnetization (e.g., 24 kG for $\text{Co}_{55}\text{Fe}_{45}$). Following the work in [11] and [12], the minimum bit length, B_{min} , in a thin longitudinal medium is approximately

$$B_{min} \approx \pi \sqrt{0.35 a_{wc}^2 + \left(\frac{D}{2}\right)^2}$$

where D is the particle diameter and a_{wc} is the Williams-Comstock transition width parameter. In SI units, a_{wc} can be expressed as [13, 14]

with a 50% volume packing fraction. Adding a 50% margin to the track pitch to account for servoing error and dimensional stability yields a track density of 45 kilo-tracks/in, and a maximum areal density of 36 Gb/in², in agreement with earlier projections [16]. To achieve 100 Gb/in² areal density with a linear density of 800 Kb/in, the track density needs to be 125 kilo-tracks/in, corresponding to a minimum bit area of 32 nm × 135 nm (having removed 50% of the track width to account for track-following error) or N = 36 particles of 12-nm diameter with a 50% volume packing fraction. This corresponds to a loss in SNR of 4.5 dB compared to the N = 100 particle case, in which SNR is defined as 20 log(N^{1/2}). This loss in SNR would need to be compensated for by improving signal processing. Other routes to an areal density of 100 Gb/in² with longitudinal media have been proposed, and these place even more constraints on track density or head-medium spacing [3, 10, 17].

Though theoretically possible, achieving a linear density of 800 Kb/in appears to be extremely challenging with longitudinal recording. In order to achieve such high linear densities, migration to perpendicular recording may be a more attractive alternative. In perpendicular recording, the magnetic orientation of the data bits is aligned vertically, that is, perpendicular to the tape. In this orientation, smaller particles can be used. Also, demagnetizing fields for perpendicular recording actually favor high linear densities, and thicker films could be used for thermal stability. For HDDs, 1,500-Kb/in linear densities have been demonstrated, using pole heads and perpendicular media with a soft underlayer [18].

Achieving 800-Kb/in linear density on tape will require very thin coatings, probably containing only a monolayer of particles. In order to achieve low noise for small bit sizes, the particles must be magnetically isolated but packed closely and uniformly with their neighbors. Uniformly sized spherical particles would enable the easiest packing, and recently there has been much success in synthesizing nanoparticles and arranging them uniformly on surfaces [19, 20]. Using these particles in magnetic tape will require expensive synthesis as well as a rapid method for coating them uniformly on the tape, which is now approaching 1 km in length. One alternative to particles requires the formation of thin magnetic films using vacuum deposition. Evaporated metal films have already been introduced in tape products [2] and have the advantage of being deposited at a high rate and of using inexpensive metals. Sputtered metal films have been successfully developed for HDD and have the potential for much larger areal density. Such technology would need to be transferred to tape media in an economically viable solution, which may be possible but remains challenging [21, 22]. In summary, sputtered magnetic films are a promising approach for tape, but optimized


particulate and evaporated metal films may still have the potential to achieve a 100-Gb/in² areal density.

Write and read heads

Unlike heads in disk drives, magnetic recording heads in modern tape drives employ arrays of transducers operating simultaneously to enable high-speed writing and reading. A multiplicity of transducers is needed because compared to disk drives, tapes move slower (3–6 m/s for tape versus >14 m/s for disk), and data is recorded to tape at 2 to 3 times fewer bits per inch due to the magnetic coating characteristics described in [23]. Multiple parallel transducers also facilitate writing or reading of an entire 1/2-inch tape cartridge within 2 to 3 hours. This has not changed significantly since tape cartridges were introduced more than 20 years ago. However, a paradigm in which the data rate does not scale with capacity necessarily increases fill times (i.e., time to write the full capacity of the cartridge) in the future. The large investment in cartridges and automation in modern tape libraries and the need to maintain a cost advantage over disk-drive libraries provide a strong incentive to preserve the existing 1/2-inch-wide tape and cartridge formats. Thus, a challenge involves increasing areal density and 1/2-inch tape cartridge capacity while preserving or molding customer expectations for backup and restore times.

A cartridge containing a 1,000-m-long tape, recorded at 100 Gb/in² has a capacity of more than 100 TB, more than 100 times that of LTO-4 tapes. Given 32 active write heads, which is the current maximum in a drive (Sun

Figure 3 Minimum bit length versus medium thickness calculated for three different values of head-medium spacing. $H_c = 5,300$ Oe, $M_s = 550$ emu/cc, $\bar{Q} = 0.8$, $S^* = 0.8$. With these values, the maximum particle anisotropy equals 2.9×10^6 erg/cc, and the minimum particle diameter for magnetic stability equals 12 nm.



on top of the read transducer), as in disk heads, and the pitch (i.e., the space between adjacent write-read pairs) is $166.5 \mu\text{m}$. The span between outermost pairs is $(16 - 1) \times 166.5 \mu\text{m}$, or $\sim 2,500 \mu\text{m}$. A cross-section schematic of a pair of write transducers is shown in Figure 4. A transducer-to-transducer spacing that is closer than the spacing in LTO was chosen for this figure in order to display more than one writer and thus emphasize a key difference between tape and HDD heads, which have only a single write transducer (*transducer* and *writer* are used interchangeably in this context). A 400-ppm environmentally induced change in tape width between writing and subsequent read back could produce up to $0.5\text{-}\mu\text{m}$ misregistration between an ideal head and the tape, or ~ 2.5 times the track spacing projected for 100-Gb/in^2 areal densities. Thus, a threefold active-element span reduction will be needed. Since at least 64 active channels will be required, transducer centerline spacing must decrease to $\sim 14 \mu\text{m}$, as shown schematically in Figures 5(a) and 5(b). Figure 5(c) shows an exemplary LTO write transducer. At this spacing, writers may interfere, or produce *crossstalk*, a phenomenon in which current in one write transducer alters the current in its nearest neighbors. A design requirement should ensure that writer yokes must not saturate, particularly in the portion wound by the coils (called the *back-gap region*), as this can divert stray flux into neighboring write gaps. Thus, high- B_c pole materials and wide back gaps (wider than top pole) are preferred. Writer poles in the range of $0.5\text{-}1.0 \mu\text{m}$ may be used to *single-write* a $0.2\text{-}\mu\text{m}$ residual track. With a single-write, the final written track width is not determined by the writer width but by the slight overlap of the adjacent written track. A $1.0\text{-}\mu\text{m}$ -wide write pole would have approximately a $2\text{-}\mu\text{m}$ back gap. This leaves a space between neighboring write yokes of $\sim 12 \mu\text{m}$, enough to accommodate two coil turns on a $1.75\text{-}\mu\text{m}$ pitch, which is within the capabilities of current thin-film microfabrication technology [see Figures 5(a) and 5(b)]. Crossstalk effects due to coil proximity may be tolerable. In any case, traditional design methodology would suggest that coils in heads at this pitch will be limited to between two (one-layer) turns and six (three-layer) turns.

To help orient the reader in studying Figure 5, note that the current channel pitch shown in Figure 5(c) is $166.5 \mu\text{m}$, as compared to $14 \mu\text{m}$ for future heads. The current write width of $12 \mu\text{m}$ is contrasted with $1 \mu\text{m}$ for future heads. The current yoke length of $25 \mu\text{m}$ is reduced significantly to $8 \mu\text{m}$ in future heads. The back-gap width of $24 \mu\text{m}$ in current heads is contrasted with $2 \mu\text{m}$ in future heads.

Restricting tape-head coils to two to six turns poses challenges. Whereas disk heads may have only 3 to 5 coil turns, tape heads typically have 8 to 14 turns. The reason

transducers are piggybacked (i.e., the write transducer is servo-track read transducers. The write and read

contains 16 write, 16 read, plus 2 or more dedicated in real time as it is written, see [23]). Each module *modules* (two modules are needed for read-verifying data which is composed of two symmetrical halves called Consider the span of transducers in an LTO head, at the other end of the head array.

element at one end of the head array to the read element smaller span. The "span" is the distance from the read today. Thus, more transducers will have to fit into a may improve only to about 400 ppm from 800 ppm accommodate transverse dimensional stability, which span occupied by the transducers must shrink to smaller tracks and higher linear densities. In addition, the transducer dimensions must shrink to enable writing sufficient for scaling to 100 Gb/in^2 . As explained below, simply increasing the number of active transducers is not of active channels regardless of tape speed. However,

rate. There is still an ongoing need to increase the number serious of which may be the burgeoning per-channel data rate. Increasing the tape speed has its own challenges, the most time to 25 hours, still outside what is acceptable today. Doubling force a change in backup management strategy. Doubling corresponds to more than 100 hours. This challenge could be required to fill this cartridge. A tape speed of 5 m/s (i.e., moving the tape from beginning to the end) would simple calculation shows that 1,875 one-way passes

Microsystems T10K tape drive), and 60,000 data tracks, a

Schematic of a magnetic write head, showing 2 of 16 write transducers. The yokes are sectioned to show front and back gaps, yoke length, and coils. The spacing between transducers in this exemplary view is less than the LTO spacing of $166.5 \mu\text{m}$, which would permit only one device to be displayed in this figure. (Coils are red, and yokes are in green, blue, and purple. The back gap occurs where the green and purple meet inside the coil. The front gap occurs where the green and purple meet in front of the coil.)




Figure 4

for this is as follows. In disk drives, current-mode write drivers are mounted on the actuator arm in close proximity to the head. Cable impedance is not critical, and more than 100 mA can be supplied to the head, enabling fewer turns (the writing field is proportional to the number of turns times the current). In tape systems, write driver chips are located at the stationary end of the cables, that is, on the circuit boards with other drive electronics, due to heat, space, row-access-strobe latency, and other restrictions. Cables must be sufficiently long to allow heads to access the entire tape, and so cable impedance becomes critical. As a result, a voltage mode must be used, reducing the maximum write current and raising the minimum number of coil turns. Operating in the lowest possible write driver voltage range helps. This enables use of smaller series write resistors but requires lower cable impedance, which ultimately affects the minimum acceptable turns count. Impedance matching also helps, and while a 10 to 90 rise time is smallest in this case, up to 50% overshoot has been used in tape and disk products, even though the reasons for overshoot are different. The relationship between overshoot and error rates, especially for thinner and non-oriented tapes, is a subject of active discussion [24]. Another consideration is that lower write voltage enables lower power dissipation, especially important for 32 or more active write transducers. This lowering of write voltages is a likely future trend, regardless of the number of coil turns. In summary, six coil turns is an aggressive lower bound for the turn number. Accordingly, though challenging for technologists, the 14- μm scaled transducer pitch is still possible. Writer crossstalk may affect minimum achievable pitch. Another factor that could affect write-head coil turns is electromigration, which places a limit on maximum acceptable current. Over time, higher media coercivity will demand larger head fields, which for a given design demands larger magnetomotive force, usually specified in units of amp-turns. However, media coatings will become thinner, enabling the use of smaller write gaps, thus resulting in higher head efficiency and helping to alleviate the current demands.




Figure 5 Future and current write heads compared. Channel pitch, write-pole width, yoke length, and back-gap width are shown in microns. (a) A depiction of a future write head at the same scale having three layers of coils, with two turns in each layer. The pitch between transducers is only 14 μm , thus enabling 64 writers (four times the number of writers in LTO drives) to fit in one-third the LTO span. (b) Cross-section of (a). (c) Schematic drawing of an exemplary LTO write transducer, having two layers (one shown) of coils, with seven turns in each layer. The pitch between transducers is 166.5 μm ; thus, only one transducer fits in this figure.

pitch otherwise required. However, this approach has device yield and tier-to-tier alignment challenges. Further, tiers are associated with a wider mechanical gap, that is, the gap of thin films residing between the hard ceramic portions of the head [23], leading to greater susceptibility to wear-induced gap recession. Alternatively, helical ("barber pole") coil transducers have a spacing advantage, but a 1.8- μm coil pitch would accommodate only five helical coil turns with a 10- μm yoke. This is problematic for the reasons cited above.

Scaling read heads for 100-Gb/in² areal densities is less about achieving small centerline pitches, which have been successfully demonstrated as small as ~15 μm , than it is about maintaining sufficient output signal. Read-head output is preferably at least 1 mV, or about ten times above any non-medium-related electronic noise, including the Johnson noise voltage in the sensor itself ($V_n \sim [4k_B T R \Delta f]^{1/2}$, where T is temperature, R is resistance, and Δf is the frequency bandwidth over which the noise is measured). At 100 Gb/in², the width of the read heads approaches 0.1 μm or less, and the read heads have an even smaller shield-to-shield spacing. For anisotropic magnetoresistive (AMR) heads, sensing layer resistance change, measured as $\delta R/R$, is ~0.023. R is

proportional to $W_{MR}/(t_{MR}/h_{MR})$, where W_{MR} denotes the sensor width, t_{MR} the thickness, and h_{MR} the height, and thus the decline of W_{MR} may be opposed by reducing the sensor-free layer thickness t_{MR} . However, for thickness less than ~ 13 nm, the output no longer rises inversely with thickness, because of the relative contribution of electron scattering at the free layer surfaces; resistance increases, but δR does not by the same proportion. Thinner media (even assuming we can overcome the challenge of making M_r small) and reduced shield-to-shield spacing, which is required for reading shorter tape wavelengths that accompany high linear density, lead to reduced signal amplitude. Increasing sensor bias current to offset declining amplitude raises temperature. As temperature increases, tape asperity cooling noise also increases. Also, the ratio of signal to sensor Johnson noise becomes reduced. Giant magnetoresistive (GMR) sensors have about five times higher sensitivity and are required for reader widths of less than ~ 4 μm for the reasons cited. GMR sensors are expected to appear in the first terabyte cartridge products. However, GMR sensors do not have enough sensitivity for scaling to 100 Gb/in^2 . Tunneling magnetoresistive sensors, currently in use in HDDs, will be required for maintaining 1-mV output for 0.1 - μm -wide read heads.

Planar head technology, in which transducers are fabricated with the critical recording gaps orthogonal to the wafer surface, is promising. Accordingly, planar magnetic heads are not sliced and lapped like conventional disk heads. Planar technology enables staggering writer arrays for reducing the effective pitch between elements. This may lead to fabricating adjacent track bundle writing arrays, which is not feasible with conventional processes. In the absence of dynamic tape skewing, adjacent track bundle writing does not experience an unwanted wavy track effect, in which residual shingled track widths vary down the length of the tape. To date, planar write heads have been successfully fabricated, but planar read-head fabrication processes have yet to be developed. This relates to another challenge. A requirement for high-density recording is accurate track placement. Ideally, this is achieved using servo readers proximate to the writing transducers. For example, servo tracking via readers in the writing module was first implemented in the IBM LTO-2 product and is referred to as the "same-gap servo," as the servo readers are in the same physical gap as the active writers. Since there are no planar readers, this writing mode is not possible yet. Another challenge for writing using staggered arrays is detecting and controlling tape dynamic skewing during writing. Again, planar readers will be needed.

In the quest for scaling to 100 Gb/in^2 , there is opportunity for improving the reading process itself. For

example, arrays of readers may be configured for electronic track following, in which readers may bridge into neighboring tracks during read back. Note that by having several readers per track, such that a few bridge into the neighboring tracks, the signal may be deconvoluted electronically rather than requiring precise mechanical control over the path of the reader. Another opportunity relates to the separation between head modules [23]. This spacing enables isolating writers from downstream readers, which are on for read verification during writing. However, similar to the case of staggered induced mis-tracking between readers and writers. Simply reducing module separation does not work well, not only because of the isolation requirement but also because of head-building constraints. Thus, some form of dynamic skew compensation will be required for achieving 100 - Gb/in^2 areal density.

Head-tape interaction

The head-tape interface must be designed to achieve a small magnetic spacing while minimizing head-tape forces, wear, and chemical degradation of the magnetic elements.

In tape recording, a critical problem is signal degradation due to spacing loss as the magnetic elements wear down as much as 40 nm below the plane of the bearing surface. This is unacceptable for 100 - Gb/in^2 areal density. A protective diamond-like carbon (DLC) coating is used on HDDs and metal-evaporated tapes, but not on metal-particle tapes or linear tape drive heads, where the coating would be worn away. Because head material choices are limited by magnetic, chemical, and physical requirements, designing for wear resistance has largely been unfruitful. Perhaps smoother future media will lead to less head recession. Reducing the head closure gap can further reduce head recession.

If head recession can be reduced to 10 nm, the 20 -nm magnetic spacing required for 800 - Kb/in linear density can be achieved only if the tape roughness is less than 10 nm. Decreasing surface roughness increases real contact area, leading to higher adhesion and frictional forces [25]. The static friction to start the tape, called *stiction*, is often much higher than the friction during operation. When the tape is stopped, stiction increases as the tape deforms into local contact with the head, and liquids condense or migrate to the contact. Of course, stiction cannot be so high that it damages the tape or stalls the tape drive. More fundamentally, variations in a high running frictional force can distort the tape laterally, imparting tracking and distorting the track spacing, or longitudinally, distorting the signal timing. To improve volumetric storage density for a given areal density, the thickness of the tape should be reduced, which further

increases its sensitivity to these problems. Thus, there is a fundamental conflict between the requirements of high linear density and unperturbed tape motion.

In tape drives, air bearings are not used at the read and write elements, because the head-tape interface cannot be controlled as well as the head-disk interface. The tape surface cannot be made as smooth as a disk, and the fast-moving tape generates many particles that can clog the air bearing. In modern tape drives, the tape is in partial contact with the head, supported by the tips of the asperities associated with the tape roughness. The sharp leading corner of the head keeps particles from the interface. Loose particles are brushed from the head during tape loading. A mildly abrasive cleaning tape occasionally removes any adherent debris. Excess mobile lubricant is absorbed in the tape underlayer, forming a reservoir for replenishing the tape surface. On the smooth flexible tape surface, the friction can become large even at zero-applied normal load, because the surface energy of the interface pulls the head and tape into contact.

Protruding asperities tend to hold most of the tape surface away from the head, but the overall attraction of the tape surface partially flattens these asperities, in turn leading to greater attraction. Thus, adhesion and attraction increase quite rapidly as the roughness is decreased to maintain a small spacing between the head and the tape.

Mobile liquids on the tape surface also contribute to attraction and increased friction. Liquids with contact angles less than 90 degrees tend to wet the area around an asperity contact, bridging the several-nanometer-wide gap between the head and tape near the asperity. The tape is attracted to the head because the total surface energy decreases as the surfaces move closer and wetting increases. Mobile lubricants contribute to this problem, as does water condensed from the atmosphere. Between surfaces wet by water, at 80% humidity, condensation can fill a 4-6-nm gap. This humidity sensitivity can be prevented by selecting tape, head, and lubricant systems having water contact angles more than 90 degrees.

Friction thus becomes problematic with increasing areal density. The fundamental problem is that as the surface profile is scaled down to decrease the magnetic spacing, the energy required (per nominal surface area) to deform the surface to flatness and to increase friction decreases, whereas the energy available from surface interactions remains constant.

One strategy to reduce the head-surface contact and resulting friction is to minimize the size of the head. For example, for the head design used by IBM, the entire length of tape contact is only 1.2 mm [26]. However, this length cannot be reduced much further while maintaining close head-tape spacing. Careful microscopic design of the head-tape interface is thus essential. For example, to ensure that the valleys of the tape surface are no further than the 10 nm (specified above) from the contacting tape peaks, the RMS roughness of the tape surface must be no larger than 1.7 nm. The bulk Young's modulus of a particulate tape coating is about 10^{10} N/m², which makes the tape sufficiently flexible for the surface topology to deform completely into contact with the head, even for small surface energies, in the absence of liquids.

This problem can be avoided by incorporating hard (e.g., alumina) particles, which are ~ 0.2 μ m, into the medium. These particles protrude from an otherwise flat medium to support local head contact. Until recently, the main role of such particles was to clean debris from the head. The sharp particle radius ensures a weak attraction to the head, but the particles are sufficiently large to distribute a large head force over the surrounding medium. An average spacing between contacting asperities of 10 μ m imposes a low enough load/particle to prevent plastic deformation of the medium around the head. The protrusion heights of the particles need to be uniform so that a high fraction of the particles support the head. Otherwise, the required high particle density will cause many local magnetic droplets. This scheme can successfully support the magnetic media within 10 nm of the head without generating strong attraction and friction.

We conclude that it is feasible to design a head-tape interface that will support 100-Gb/in² areal density.

Data detection

Achieving the areal recording densities discussed in this paper poses significant challenges in terms of read-channel design. The main challenge is to ensure highly reliable operation of all front-end analog and digital signal-processing functions, including adaptive equalization and gain and timing control, despite significant reductions in the available SNR values.

At high areal densities, read channels also need to rely on powerful data-detection methods to guarantee post-detection symbol error rates that are sufficiently low. Noise-predictive maximum-likelihood (NPML) sequence detection [27] is well suited to address these performance requirements. For example, the class of NPML detectors, with target polynomials $(1 - D^2)W(D)$, where $W(D)$ represents a noise whitening filter and D a delay operator, has been implemented successfully in HDDs and is also attractive for tape systems. At high linear densities, NPML schemes achieve a better match between the detector target and physical channel characteristics than, for example, extended partial-response class 4 (EPR4) schemes. Moreover, the NPML schemes whiten the noise process (i.e., make the noise have a flatter frequency spectrum) at the detector input and also reduce its power. Another important aspect of tape systems is the inherent

519

IBM J. RES. & DEV. VOL. 52 NO. 4/5 JULY/SEPTEMBER 2008

A. J. ARGUMEDO ET AL.

variability of the recording channel, which is due to characteristics, variations in the read- and write-head factors. This aspect can best be dealt with by employing a detector target that automatically adapts itself to the current channel characteristics. Such a feature is enabled by the class of NPML targets mentioned above.

NPML detection can be extended further to take into account the data-dependent nature of the noise process. It is, for example, well known that surface roughness in tape media introduces a type of noise that is colored (i.e., with a non-flat frequency spectrum) and data dependent. Data-dependent NPML detection allows one to achieve the best detection performance in the presence of such noise processes, which can be the predominant contributors to the total channel noise. In conjunction with advanced coding, these techniques could provide the additional 4.5-dB noise margin that is needed to achieve the 100-Gb/in² operating point.

Format efficiency

To make efficient use of the magnetic recording channel and thus achieve reliable read-back operation of the user data, the bit stream written onto the magnetic medium includes redundancy and synchronization patterns. The specification of this overhead data is referred to as *formatting* or *data format*. The goal of an efficient format is to introduce as little overhead as possible while ensuring proper operation of the data acquisition and timing loops and meeting the 10⁻¹⁷ bit error rate requirement of tape recording systems. Clearly, improvements in format efficiency directly lead to higher cartridge capacity.

The format efficiency of current tape-recording systems such as LTO-4 is about 71.6% [9]. The 28.4% overhead of LTO-4 can be broken down into about 17% for error-correction coding (ECC), 5.5% for modulation coding, and about 5.9% for sync patterns and encoded data headers. There are several ways to improve format efficiency. One approach is based on using longer ECC codes of higher rates, resulting in about 4% gain without sacrificing error-correction performance. Higher efficiency, however, is achieved by another approach, which relies on reverse concatenation (RC), and has already been effectively implemented in HDD products. In a standard forward-concatenation scheme, user data is first ECC encoded and then passed through a modulation encoder to enforce predetermined modulation constraints for timing and efficient data detection purposes. In an RC scheme, the order of the ECC encoder and modulation encoder is reversed [28]. This reversal of the encoding order provides three major benefits: 1) There is no error propagation through the modulation decoder; 2) because error propagation is not

a concern, the first modulation code can be taken to be very long, allowing the use of capacity-efficient and high-rate modulation codes and thereby resulting in code rate gains; 3) in the read-back path, the ECC decoding block comes immediately after the channel-detection block, which can readily pass tentative decisions to the decoder on a bit-by-bit basis. The placement of the ECC decoding block creates the appropriate framework for using novel ECC techniques, which are based on turbo and low-density parity-check (LDPC) codes and may provide significant performance improvements [29].

These three benefits can also be exploited in the framework of tape recording. However, the ECC used in HDDs has a different structure from that used in tape recording. Previously, RC has been proposed for one-dimensional ECC architectures, for which the ECC typically consists of a single code such as a Reed-Solomon or an LDPC code [30, 31]. These RC architectures cannot be directly applied to two-dimensional ECC in tape systems, which are based on Reed-Solomon product codes with a C1 code along rows and a C2 code along columns. To overcome this problem, a novel RC scheme is proposed, which is illustrated in Figure 6. The main steps in the write path are as follows: 1) user data is reorganized into a stream of N_2 rows by the serial-to-parallel function (represented by the S/P block); 2) modulation encoding of each row is accomplished by the first modulation encoder ME-1; 3) formatting is performed for partial symbol interleaving; 4) C2-column-dependent encoding is performed; 5) C1 encoding is performed along rows; and 6) modulation coding of the C1 parity is performed by a systematic modulation encoder ME-2.

A key component in the proposed RC scheme is the first modulation code, for which one may select a very high-rate $n/(n+1)$ Fibonacci code, typically $n > 200$ [31]. These codes involve simple enumerative encoders and achieve very restrictive modulation constraints, which are comparable to those of the LTO-4 standard. Modulation constraints avoid the writing of sequences for which the overall detection process would be less reliable (i.e., they avoid unfavorable timing patterns, reduce path-memory length in sequence detectors, and avoid quasi-catastrophic error propagation). Another specific feature of the RC scheme is the formatting block, which transforms the modulated user data array into an array with "empty" components in each column, which are the locations in which the parity symbols of the C2 code will be introduced.

The new RC scheme has a modulation method with less than 1% redundancy while maintaining essentially unaltered modulation constraints. This improvement in rate is more than 5% above the rate-16/17 code of the LTO-4 standard and, together with the 4% potential gain

from longer C2 codes, leads to an overall format with an overhead of about 18% rather than the 27% of LTO-4. This represents a substantial reduction in redundancy. Furthermore, by using an LDPC code or turbo code for C1, the new format supports novel ECC techniques based on iterative decoding. The C1/C2-based ECC structure is an ideal setting for LDPC or turbo codes because the typical error floor issue (in which the correctable error rate is not sufficiently low) of these codes is resolved by the C2 Reed-Solomon code, which can reduce the error rates to the desired 10^{-17} level.

Track density limits

Track misregistration (TMR) limits the allowable track pitch that can be written on tape. To achieve 100 Gb/in² track density of 125 kilo-tracks per inch or 0.2- μ m written tracks. To a first-order approximation, if we assume the reader width to be 0.13 μ m, the TMR, which is defined here as the difference between the track pitch and the reader width, must be less than 70 nm. Various factors contribute to TMR, including the dimensional variability of the head, track-following fidelity, and the transverse dimensional stability (TDS) of the tape.

Because multiple heads are used in parallel, lateral expansion and contraction of the tape contributes to TMR. Achieving 100-Gb/in² areal density, or equivalently 100 TB in a tape cartridge, will require a tape that can be written in one environment, appended to in a second environment, and read in a third environment. This imposes severe dimensional-stability constraints. Currently, the substrate is stretched in the machine direction and transverse direction to thin the tape to facilitate high volumetric density and to increase its modulus for easier coating and better performance. However, this process complicates and degrades the TDS. Current media exhibit a lateral dimensional change of 750-800 ppm over the full environmental variation, and improvement beyond 500 ppm is unlikely. At this level, dimensional stability can be accommodated by reducing the head span by a factor of 3, as discussed above.

Tape paths for high track density

Lateral tape motion (LTM) must be significantly improved to achieve higher track densities. In high-performance drives, rolling elements transport the tape between reels and limit LTM. Continued use of rollers requires the prevention of debris accumulation on the roller flanges, which strikes the tape edges, causing LTM. This LTM disturbance often exceeds the bandwidth and slewing capability of the track-following actuator. Debris accumulates because the spacing between the two reel




Figure 6 Reverse concatenation architecture. (S/P: serial to parallel; Enc: encoder; MUX: multiplexer.)

flanges is much larger than the roller flange spacing. Tape tends to stack against the reel flanges, so that when it is transported from a reel to the first roller, a large force develops between the tape edge and the roller flange, causing wear and debris accumulation. One obvious solution to this problem is to remove the flanges, but this introduces other challenges. First, without the constraint of the flanges, LTM increases as the tape moves up or down between the widely spaced reel flanges. Second, the angle of the tape with respect to the head can become skewed. These additional challenges can be addressed by constructing a more advanced actuator, capable of following a larger LTM and of servoing its rotation angle to keep the head perpendicular to the tape. In addition, actively controlled tilting elements elsewhere in the path may be used to reduce skew and lateral excursion. By implementing these and other advancements, track-following can be improved significantly. Experimental paths incorporating flangeless grooved rollers have achieved a position error signal (PES) with a standard deviation of PES as low as 61 nm, using a legacy servo channel. Although this represents a substantial improvement, we anticipate that an even larger reduction of PES will result from improved detection of the position and velocity information, larger actuator bandwidth, and improved control of the reel-to-reel servomechanism.

implemented in a field-programmable gate array (FPGA) and tested in a prototype environment. The performance

The synchronous servo-channel concept was output to monitor the quality of the servo channel. In addition, a reliability measure is assigned to the detector of estimates of tape velocity and lateral position. In signals in the presence of noise and for the computation approach is employed for optimal detection of PPM independent of velocity. A digital dibit correlator fixed number of signal samples per unit of length of tape, channel that relies on a digital interpolator to generate a Figure 7 shows a block diagram of a synchronous servo realized in a prototype drive system.

implements the aforementioned functions, has been experimental fully synchronous servo channel [3], which of length of tape, irrespective of tape velocity. An (3) generation of a fixed number of signal samples per unit bursts, 2) optimal demodulation of LPOS symbols, and functions: 1) optimal matched-filter detection of servo are required that include the following three main σ_{PES} values, advances in the servo-channel architecture information and PES can be derived. To achieve small channels are normally available, from which LPOS modulation (PPM). In tape systems, two dedicated servo their normal pattern position, using pulse-position encoding is obtained by properly shifting transitions from affecting the generation of the transversal PES. This longitudinal position (LPOS) information without TBS patterns also enable the encoding of additional pulses generated by a narrow head reading the pattern. The lateral position is derived from the relative timing of consist of transitions with two different azimuthal slopes. systems, recorded patterns to aid track-following servo and it has been adopted as an LTO standard. In TBS specifically for linear tape drives in the mid 1990s [32]. Timing-based servoing (TBS) is a technology developed Synchronous servo channel

Block diagram of a synchronous servo channel. (ADC: analog-to-digital converter; LPOS: longitudinal position; PES: position error signal.)




Figure 7

important steps in defining the reel-to-reel control substantially reduce LTM. which in conjunction with an optimized tape path can enables simultaneous control of tension and velocity, tape-transport system in [35]. A MIMO architecture introduced in [34] and has been applied to a prototype notion of a MIMO control system for tape transport was the sampling frequency of the digital controller. The measurements from the sensors is not commensurate with capability of handling designs in which the rate of important feature of a state-space-based system is its estimated parameters need to be utilized. Another required if information from multiple sensors and several system. A MIMO control design allows multiple inputs as the design of a multi-input multi-output (MIMO) control of a state-space-based control system is its suitability for space-based methods are required. The main advantage derivative (PID) controllers and the introduction of state-tension, a departure from standard proportional-integral requires simultaneous control of tape velocity and

For the tape transport system mentioned above, which control becomes increasingly important. the reel-to-reel servo system for tape velocity and tension In order to counteract these effects, an adequate design of reduced TDS and a larger susceptibility to tape damage. coating and tape substrate, which in turn may lead to higher volumetric densities require thinner magnetic affects the data write and read performance. Moreover, control because the quality of the tape transport directly areal densities require excellent tape motion and tension and volumetric storage densities must be improved. High further increase the tape cartridge capacity, both areal density by winding a very long tape on a single reel. To systems is their ability to achieve a very high volumetric One of the main advantages of tape-based storage

Reel-to-reel control


One of the main advantages of tape-based storage systems is their ability to achieve a very high volumetric density by winding a very long tape on a single reel. To further increase the tape cartridge capacity, both areal and volumetric storage densities must be improved. High areal densities require excellent tape motion and tension control because the quality of the tape transport directly affects the data write and read performance. Moreover, higher volumetric densities require thinner magnetic coating and tape substrate, which in turn may lead to reduced TDS and a larger susceptibility to tape damage. In order to counteract these effects, an adequate design of the reel-to-reel servo system for tape velocity and tension control becomes increasingly important.

Example of high-areal-density recording

To demonstrate the potential extensibility of tape technology density, we have constructed an experimental apparatus for measuring the recording properties of tape media at high areal density. This apparatus precisely positions an HDD head with a narrow writer and reader to probe tape media at extremely small dimensions. As the head is moved along the tape, the bearing surface of the head is in contact with the tape surface, with the magnetic spacing being limited only by the tape surface roughness. Positioning is achieved with a Physik Instrumente (PI) P-587 six-axis piezoelectric nanopositioning stage, which achieves sub-10-nm accuracy, so that we have precise control over all degrees of freedom of the head-tape interface. The head writer width is 0.3 μm , and the giant magnetoresistive (GMR) reader width is 0.15 μm . In the example below, the tape particle length is on the order of 35 nm, resulting in a reader width of approximately five particles, so that we can probe the magnetic recording properties to the level of a few particles. To demonstrate the relative difference between HDD density and the large potential for improvement in tape technology, we have used a piece of LTO-4 tape in an experiment in this setup. The tape was prewritten in an LTO tape drive with a standard LTO-4 track width of 11.3 μm . Figure 8 shows a grayscale image of the read-back signal from the GMR head on our high-density recording apparatus. White stripes correspond to positive transitions, and dark stripes correspond to negative transitions. A section of this tape was overwritten by the HDD write head in our high-density recording system to spell out the letters *IBM* and to fit them into a single tape track. Each letter consists of tracks parallel to the tape track direction, with transitions positioned to correspond to the horizontal stripes in the *IBM* logo. The tracks in the letters are written at a pitch

techniques to gather the necessary information to provide feedback of tension and velocity.

IBM logo written by high-density tape recording.




Conclusions

While it is clear that there will be significant challenges in scaling linear magnetic tape technology to areal densities comparable to current HDD technology, there appears to be a viable path toward this goal. Controlling the tape-head interaction will be key to achieving high linear density, while improvements in track-following and reel-to-reel servomechanisms as well as improvements in TDS and reduced-span heads will be key to achieving high track densities. In addition, advanced head- and data-detection technologies as well as improved LTM control will have an impact on both linear density and track pitch and, therefore, will be key enablers to achieving ultrahigh areal densities in a linear magnetic tape system. Through the combination of these technologies, a linear density of 800 Kb/in and a track pitch of 0.2 μm appears feasible, leading to an areal density of 100 Gb/in².

*Trademark, service mark, or registered trademark of International Business Machines Corporation in the United States, other countries, or both.

Large magnification image of the top of the letter *I* in "IBM."



19. J. Pyun, "Nanocomposite Materials from Functional Polymers and Magnetic Colloids," *Polymer Rev.* 47, No. 2, 231-263 (2007).

20. X.-M. Lin and A. C. S. Samia, "Synthesis, Assembly and Physical Properties of Magnetic Nanoparticles," *J. Magn. Mater.* 305, No. 1, 100-109 (2006).

21. H.-S. Lee, L. Wang, J. A. Bain, and D. E. Laughlin, "Thin-Film Recording Media on Flexible Substrates for Tape Applications," *IEEE Trans. Magn.* 41, No. 2, 654-659 (2005).

22. K. Moriwaki, K. Usuki, and N. Nagao, "CoPtCr-SiO₂/Ru Longitudinal Media with C Underlayer for High-Density Flexible Disk," *IEEE Trans. Magn.* 41, No. 10, 3244-3246 (2005).

23. R. G. Biskoborn and J. H. Eaton, "Hard-Disk-Drive Technology Flat Heads for Linear Tape Recording," *IBM J. Res. & Dev.* 47, No. 4, 385-400 (2003).

24. International Storage Industry Consortium (INSIC), "International Magnetic Tape Storage Roadmap," April 2005; see <http://www.insic.org/ippdmpsm.pdf>.

25. B. Bhushan, "Adhesion and Sliction: Mechanisms, Measurement Techniques, and Methods of Reduction," *J. Vac. Sci. Technol. B* 21, No. 6, 2262-2296 (2003).

26. R. G. Biskoborn and J. H. Eaton, "Flat-Profile Tape Recording Head," *IEEE Trans. Magn.* 38, No. 5, 1919-1921 (2002).

27. J. D. Coker, E. Elieheriou, R. L. Galbraith, and W. Hirt, "Noise-Predictive Maximum-Likelihood (NPMML) Detection," *IEEE Trans. Magn.* 34, No. 1, 110-117 (1998).

28. W. G. Bliss, "Circuitry for Performing Error Correction Calculations on Baseband Encoded Data to Eliminate Error Propagation," *IBM Tech. Disc. Bull.* 23, No. 10, 4633-4634 (1981).

29. A. Dhoulakia, E. Elieheriou, T. Mittelholzer, and M. P. C. Fossorier, "Capacity-Approaching Codes: Can They be Applied to the Magnetic Recording Channel?" *IEEE Commun. Mag.* 42, No. 2, 122-130 (2004).

30. A. J. van Wijngaarden and K. A. S. Immink, "Maximum Runlength-Limited Codes with Error Control Capabilities," *IEEE J. Select. Areas Commun.* 19, No. 4, 602-611 (2001).

31. M. Blaum, R. D. Cideciyan, E. Elieheriou, R. Galbraith, K. Lakovic, T. Mittelholzer, T. Oeming, and B. Wilson, "High-Rate Modulation Codes for Reverse Concatenation," *IEEE Trans. Magn.* 43, No. 2, 740-743 (2007).

32. R. C. Barrett, E. H. Klaassen, T. R. Albrecht, G. A. Jaquette, and J. H. Eaton, "Timing-Based Track-Following Servo for Linear Tape Systems," *IEEE Trans. Magn.* 34, No. 4, 1872-1877 (1998).

33. G. Cherubini, E. Elieheriou, J. Jelitto, and R. Hutchins, "Synchronous Servo Channel Design for Tape Drive Systems," *Proceedings of the 17th Annual ISPS Conference*, Santa Clara, CA, 2007, pp. 160-162.

34. G. F. Franklin, J. D. Powell, and M. Workman, *Digital Control of Dynamic Systems*, Third Edition, Addison-Wesley, Menlo Park, CA, 1997.

35. P. D. Mathur and W. C. Messner, "Controller Development for a Prototype High-Speed Low-Tension Tape Transport," *IEEE Trans. Control Syst. Technol.* 6, No. 4, 534-542 (1998).

Received September 20, 2007; accepted for publication October 9, 2007

1. D. Berman, R. Biskoborn, N. Bui, E. Childers, R. D. Cideciyan, E. Elieheriou, D. Hellman, et al., "6.7 Gb/in² Recording Area Density on Barium Ferrite Tape," *IEEE Trans. Magn.* 43, No. 8, 3502-3508 (2007).

2. K. Motohashi, T. Sato, T. Samoto, N. Ikeda, H. Ono, and S. Onodera, "Investigation of Higher Recording Density Using an Improved Co-CoO Metal Evaporated Tape with a GMR Reproducing Head," *IEEE Trans. Magn.* 43, No. 6, 2325-2327 (2007).

3. R. H. Dea, "Magnetic Tape: The Challenge of Reaching Hard-Disk-Drive Data Densities on Flexible Media," *MRS Bull.* 31, No. 5, 404-408 (2006).

4. T. Ozne, M. Kondo, Y. Soda, S. Fukuda, S. Onodera, and T. Kawana, "11.5-Gb/in² Recording Using Spin-Valve Heads in Tape Systems," *IEEE Trans. Magn.* 38, No. 1, 136-140 (2002).

5. M. Xiao, H. Do, W. Weresin, Q. Dai, Y. Ikeda, K. Takano, A. Moser, et al., "Recording Studies of Perpendicular Media Leading to 230 Gb/in²," *J. Appl. Phys.* 99, No. 8, 08E712-1-08E712-3 (2006).

6. E. R. Childers, W. Imatino, 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).

7. T. Nagata, T. Harasawa, M. Oyamagi, N. Abe, and S. Saito, "A Recording Density Study of Advanced Barium-Ferrite Particulate Tape," *IEEE Trans. Magn.* 42, No. 10, 2312-2314 (2006).

8. A. Matsumoto, Y. Endo, and H. Noguchi, "The Feasibility of +15 Gb/in² High-Density Recording with Barium-Ferrite Particulate Media and a GMR Head," *IEEE Trans. Magn.* 42, No. 10, 2315-2317 (2006).

9. LTO Ultrium Technology; see <http://www.ultrium.com/>.

10. J. C. Mallinson, "On Extremely High Density Recording," *IEEE Trans. Magn.* 10, 368-374 (1974).

11. D. Weller and A. Moser, "Thermal Effect Limits in Ultrahigh-Density Magnetic Recording," *IEEE Trans. Magn.* 35, No. 6, 4423-4439 (1999).

12. H. Zhou and H. N. Bertram, "Scaling of Hysteresis and Transition Parameter with Grain Size in Longitudinal Thin-Film Media," *J. Appl. Phys.* 85, No. 8, 4982-4984 (1999).

13. M. L. Williams and R. L. Comstock, "An Analytical Model of the Write Process in Digital Magnetic Recording," *17th Annual AIP Conference Proceedings*, Part I, No. 5, 1971, pp. 738-742.

14. H. N. Bertram, *Theory of Magnetic Recording*, Cambridge University Press, Cambridge, U.K., 1994.

15. M. P. Sharrock, "Measurements and Interpretation of Magnetic Time Effects in Recording Media," *IEEE Trans. Magn.* 35, No. 6, 4414-4422 (1999).

16. S. H. Charap, P.-L. Lu, and Y. He, "Thermal Stability of Recorded Information at High Densities," *IEEE Trans. Magn.* 33, No. 1, 978-983 (1997).

17. H. N. Bertram and M. Williams, "SNR and Density Limit Estimates: A Comparison of Longitudinal and Perpendicular Recording," *IEEE Trans. Magn.* 36, No. 1, 4-9 (2000).

18. World Magnetic Technology (September 15, 2006). Seagate Breaks Square Inch Equivalent to Storing 4,000 Hours of Digital Video on Your PC. Press release; see http://www.seagate.com/world/index.jsp?loc=en-US&name=Seagate_Breaks_World_Video_on_Your_PC&age=939597d83e010VgnVCM100000dd04090aRCRD.

References

**LTO is a registered trademark of Hewlett-Packard, International Business Machines Corporation, and Quantum in the United States, other countries, or both.

Roy D. Cideciyan IBM Research Division, Zurich Research Laboratory, Säumerstrasse 4, 8803 Rüschlikon, Switzerland (cid@zurich.ibm.com). Dr. Cideciyan received the Dipl.-Ing. degree from Aachen University of Technology, Germany, in 1981 and the M.S.E. and Ph.D. degrees from the University of Southern California, Los Angeles, in 1982 and 1985, respectively. He joined the IBM Zurich Research Laboratory in 1986, where he initially worked on partial-response maximum-likelihood channels for hard disk drives and later on coding for high-speed local area networking, and wireless communications. Since 1997, he has been

Giovanni Cherubini IBM Research Division, Zurich Research Laboratory, Säumerstrasse 4, 8803 Rüschlikon, Switzerland (gchi@zurich.ibm.com). Dr. Cherubini received a Laurea degree in electrical engineering (summa cum laude) from the University of Padova, Italy, and M.S. and Ph.D. degrees in electrical engineering from the University of California, San Diego. He joined the research staff of IBM in 1987. His research interests include high-speed data transmission and data storage systems. He was coeditor of the 100BASE-T2 Standard for Fast Ethernet transmission over voice-grade cables. More recently, he contributed to the realization of the first fully functional atomic force microscope-based data-storage prototype system. Dr. Cherubini was a co-recipient of the 2003 IEEE Communications Society Leonard G. Abraham Prize Paper Award. He is a Fellow of the IEEE and Distinguished Lecturer of the IEEE Communications Society.

Robert G. Biskeborn IBM Systems and Technology Group, Almaden Research Center, 650 Harry Road, San Jose, California 95120 (rbiskebor@us.ibm.com). Dr. Biskeborn received his Ph.D. degree in physics from Columbia University in 1978. The same year, he joined IBM in East Fishkill, New York. He initially worked in advanced packaging technology and later transferred to an advanced bipolar device design and technology group. Among his accomplishments is reinventing module thermal packaging, 1990, he transferred to the IBM Storage Systems Division in San Jose, California, where he has been working in both disk and tape storage technologies. Among his contributions is work that revolutionized the design and implementation of magnetic heads that are a critical component of high-end tape recording systems. Dr. Biskeborn is a member of the American Physical Society and is a Distinguished Scientist at the IBM Almaden Research Center.

David Berman IBM Research Division, Almaden Research Center, 650 Harry Road, San Jose, California 95120 (bermand@almaden.ibm.com). Dr. Berman received his B.S. degree in electrical engineering from the University of Cincinnati in 1992. He received his M.S. and Ph.D. degrees from MIT in electrical engineering in 1994 and 1998, respectively. Since 1998, he has worked in various fields in magnetic storage at IBM. His areas of interest include magnetic head technology, signal processing for storage read channels, and magnetic media research.

Armando J. Argumedo IBM Systems and Technology Group, 9000 South Rita Road, Tucson, Arizona 85744 (argo@us.ibm.com). Mr. Argumedo received his B.S. and M.S. degrees in energy engineering from the University of Illinois at Chicago. He joined IBM in 1977 and began development work on recording head tribology and tape transport design. Mr. Argumedo's recent assignment areas have included mechanical development of the IBM Linear Tape-Open and Enterprise tape transports.

Robert Hutchins IBM Systems and Technology Group, 9000 South Rita Road, Tucson, Arizona 85744 (bohutchins@us.ibm.com). Mr. Hutchins received a B.S. degree in biology (1977) and an M.S. degree in electrical engineering (1982) both from the University of Wyoming. In 1982, he joined the advanced tape technology organization in Tucson and has worked exclusively in the area of removable data storage throughout his career at IBM. Since 1990, most of his focus has been on digital data detection and signal processing for magnetic tape drives.

Diana J. Hellman IBM Systems and Technology Group, 9000 South Rita Road, Tucson, Arizona 85744 (dianahel@us.ibm.com). Dr. Hellman is a Media Tribologist 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. thesis on developing a novel membrane fabrication process, the thermally assisted evaporative phase separation (TAFPS) process. Prior to her current work in media, Dr. Hellman worked on tape library systems, focusing on the materials selection in the libraries.

Walter Häberle IBM Research Division, Zurich Research Laboratory, Säumerstrasse 4, 8803 Rüschlikon, Switzerland (wha@zurich.ibm.com). Mr. Häberle is a development engineer in the Science and Technology department of the IBM Zurich Research Laboratory. In 1974, he joined the IBM Semiconductor Plant in Sindelfingen, Germany, for his education as physics technician and was later promoted to group leader in the Quality Control department of the semiconductor line. In 1987, he joined the IBM Physics Group of Professor G. K. Binnig at the University of Munich, Germany, where he contributed to the development of atomic force microscopes (AFMs) for applications in a liquid environment for investigating living cells. In 1993, he transferred to the IBM Zurich Research Laboratory, where he was responsible for prototyping a new low-cost AFM instrument. Since 1997, he has been involved in various aspects of the probe-based data-storage project. In 2003, he spent a sabbatical year at the Wayne State University in Detroit, Michigan, where he investigated heat-transfer rates in biological samples using thermomechanical lever technologies. Since 2006, the focus of his research has been on the tape path of tape drive systems.

Evangelos Eleftheriou IBM Research Division, Zurich Research Laboratory, Säumerstrasse 4, 8803 Rüschlikon, Switzerland (ele@zurich.ibm.com). Dr. Eleftheriou received a B.S. degree in electrical engineering from the University of Patras, Greece, in 1979, and M.Eng. and Ph.D. degrees in electrical engineering from Carleton University, Ottawa, Canada, in 1981 and 1985, respectively. He joined the IBM Zurich Research Laboratory in 1986, where he worked on various projects related to transmission technology, magnetic recording, and probe storage. He currently manages the advanced storage technologies group. In 2005, Dr. Eleftheriou was co-recipient of the Eduard Rhein Technology Award. He was co-recipient of the 2003 IEEE Communications Society Leonard G. Abraham Prize Paper Award. In January 2002, he was elected Fellow of the IEEE, and in 2005, he became an IBM Fellow and was elected to the IBM Academy of Technology.

working on signal processing and coding problems for data storage on hard disks and tape.

Wayne Imano IBM Research Division, Almaden Research Center, 650 Harry Road, San Jose, California 95120 (waynei@almaden.ibm.com). In 1980, Dr. Imano received his Ph.D. degree in physics from Purdue University focusing on solid-state spectroscopy and was awarded the Karl Lark-Horowitz prize. He currently manages the Advanced Storage Concepts in Tape Research Laboratory. He has been employed by IBM since 1980 when he joined the staff the San Jose Center. During this time, Dr. Imano has worked in electrophotography, optical storage, dynamics of actuator mechanisms, and most recently, tape technology. He has published 26 technical papers and holds 24 patents.

Jens Jelitto IBM Research Division, Zurich Research Laboratory, Säumerstrasse 4, 8803 Rüschlikon, Switzerland (jje@zurich.ibm.com). Dr. Jelitto received his M.Sc./Dipl.-Ing. and Ph.D. degrees from the Dresden University of Technology, Germany, in 1995 and 2001, respectively. From 1995 to 1996, he worked in the field of speech recognition at the Institute for Acoustics and Speech Communication in Dresden. In July 1996, he joined the Mannesmann Mobilfunk Chair for Mobile Communications Systems at the Dresden University of Technology, Germany, to work toward his Ph.D. degree. At this time, his main research interests included digital signal processing, smart antennas, and spatial dimension reduction problems. He joined the IBM Zurich Research Laboratory, Rüschlikon, Switzerland, as a Research Staff Member in March 2001. He has worked on several research topics in the field of digital signal processing for wireless LANs and for magnetic recording. Currently, his research work is focused on advanced signal processing techniques for the tape read channels and on servo-control aspects to improve the storage capacity and reliability of tape systems.

Kevin Judd IBM Systems and Technology Group, 9000 South Rita Road, Tucson, Arizona 85747 (kevjudd@us.ibm.com). Mr. Judd received both B.S. and M.S. degrees in mechanical engineering from Brigham Young University in 2000 and 2001, respectively. In 2001, he joined IBM as part of the tape drive development group, working on the mechanical design of tape drives. Since 2001, he has been working on all aspects of the mechanics of tape drives. The main focus of his work is on the development of the tape path.

Pierre-Olivier Jubert IBM Research Division, Almaden Research Center, 650 Harry Road, San Jose, California 95120 (pjubert@us.ibm.com). Dr. Jubert is a Research Staff Member at the IBM Almaden Research Center. He received his Ph.D. degree from the University of Grenoble in 2001, in the field of epitaxial growth and magnetism of nanoparticles. Next, he joined the IBM Zurich Research Laboratory as a Postdoctoral Fellow, working on magnetism in nanoscale systems and more particularly on spin-forced-induced magnetic domain wall motion. He joined the Advanced Storage Concepts Group at the IBM Almaden Research Center in 2005, focusing on magnetic recording physics.

Mark A. Lantz IBM Research Division, Zurich Research Laboratory, Säumerstrasse 4, 8803 Rüschlikon, Switzerland (mla@zurich.ibm.com). Dr. Lantz received B.Sc. and M.Sc. degrees in electrical engineering from the University of Alberta, Canada, in 1991 and 1993, and a Ph.D. degree from the University of Cambridge, U.K., in 1997 for work in the field of scanning-probe microscopy. He then spent two years as a postdoctoral

researcher at the Joint Research Center for Atom Technology in Japan, investigating the application of scanning probes in biophysics, followed by two years of research in the area of low-temperature scanning force microscopy at the Physics Institute in University of Basel, Switzerland. In 2001, he joined the Micro/Nanomechanics group of the IBM Zurich Research Laboratory as a Research Staff Member. His current research activities are focused on micro-mechanical and nanomechanical devices and systems for scanning-probe-based data storage, tribology, and magnetic tape drive technology.

Gary M. McClelland IBM Research Division, 650 Harry Road, San Jose, California 95120 (mcclell@almaden.ibm.com). Dr. McClelland received his B.S. degree in chemistry from the University of Illinois and his Ph.D. degree in chemical physics at Harvard University. After research at Stanford, he became an Assistant Professor at Harvard and then joined IBM in 1985. Before his current work on the tribology of tape recording, his research interests have included molecular dynamics, surface science, ultrastatic dynamics, atomic force microscopy, novel solid state storage, and micro-mechanical systems.

Thomas Mittelholzer IBM Research Division, Zurich Research Laboratory, Säumerstrasse 4, 8803 Rüschlikon, Switzerland (tmit@zurich.ibm.com). Dr. Mittelholzer received his diploma and Ph.D. degree in mathematics and his diploma in computer science from the Swiss Federal Institute of Technology (ETH) Zurich, Switzerland, in 1981, 1987, and 1988, respectively. From 1989 to 1993, he was a Research Associate with the Signal and Information Processing Laboratory at the ETH. He spent 1994 as a Visiting Post-graduate Research Scientist at the Center for Magnetic Recording Research at the University of California in San Diego. After working three years as a consultant for security and cryptography, he joined the IBM Zurich Research Laboratory in 1999, where he has focused on coding for hard disk and tape recording. His research interests, primarily in coding and magnetic recording theory, include signal processing and coding for communication theory, include signal processing and coding for

Chandrasekhar (Spike) Narayan IBM Research Division, Almaden Research Center, 650 Harry Road, San Jose, California 95120 (narayan@us.ibm.com). Dr. Narayan is a Senior Manager of nanoscale science and technology in IBM Research and is responsible for developing nanoscale storage devices, processes to fabricate nanostructures, storage technology on flexible media, and molecular electronics. In addition, he is a Master Inventor within IBM Research. Previously, Dr. Narayan was the Technical Assistant to the Vice President of Science and Technology in IBM Research, where he was responsible for working with executives to develop the science and technology strategy for the IBM worldwide research laboratories and was also responsible for developing the technology aspect of the Global Technology Outlook. Additionally, Dr. Narayan has managed IBM high-performance CMOS logic integration, where he led the team responsible for the CMOS interconnect technology transfer from research to development and for integrating copper with the next-generation ultralow-k dielectrics. Throughout his career, he has held a variety of positions in research and research management. He has received several IBM awards for his work in cobalt barrier layer for thin-film metallization, thermal conduction module failure analysis, high-information-content display prototyping and eFuse development. He holds 50 U.S. patents. In addition, Dr. Narayan has contributed to the external engineering community by serving as the general and program chair for the IEEE/EMT Symposium in 1999 and 2000, respectively, and he chaired the DRAM Development Alliance Invention Board in 1999.

Sedar Oiger IBM Research Division, Zurich Research Laboratory, Säumerstrasse 4, 8803 Rüschlikon, Switzerland (oel@zurich.ibm.com). Dr. Oiger received a diploma of electrical engineering and a Ph.D. degree from the Swiss Federal Institute of Technology, Lausanne, Switzerland, in 1978 and 1982, respectively. From 1982 to 1984, he was a research associate at the Information Systems Laboratory of Stanford University and at Yale University. He then joined the IBM Zurich Research Laboratory, where he initially worked on digital transmission techniques for magnetic recording channels and later on high-speed data communications for local area networking and broadband network access. Since 2003, his work has focused on signal processing and coding problems for data storage on tape.

Paul J. Seger IBM Systems and Technology Group, 9000 South Rita Road, Tucson, Arizona 85747 (pseger@us.ibm.com). Mr. Seger is a Senior Engineer in the Tape Integration and Electronics department in the IBM Tucson Laboratory. In 1969, he received a B.S. degree in electrical engineering from Villanova University. In 1969, he also joined the IBM Federal Systems Division in Owego, New York, specializing in analog and digital circuit design for memory systems used aboard aircraft, submarines, and spacecraft. In 1974, he received M.S. degree in electrical engineering from Syracuse University specializing in feedback and control systems. In 1980, he joined the IBM Tucson Laboratory working on advanced recording channel technology. His current work includes modeling and developing logical formats for tape systems. Mr. Seger's participation in the standards community includes chairmanship of ECMA TC17 and Head of U.S. Delegation to ISO SC11. Mr. Seger has 12 patents and applications and has contributed chapters to two texts on magnetic recording.