Magnetic Tape Storage Technology

MARK A. LANTZ, IBM Research Europe - Zurich, Rüschlikon, Switzerland
SIMEON FURRER, IBM Research Europe - Zurich, Rüschlikon, Switzerland
MARTIN PETERMANN, IBM Research Europe - Zurich, Rüschlikon, Switzerland
HUGO ROTHUIZEN, IBM Research Europe - Zurich, Rüschlikon, Switzerland
STELLA BRACH, IBM Research Europe - Zurich, Rüschlikon, Switzerland
LUZIUS KRONIG, IBM Research Europe - Zurich, Rüschlikon, Switzerland
ILIAS ILIADIS, IBM Research Europe - Zurich, Rüschlikon, Switzerland
BEAT WEISS, IBM Research Europe - Zurich, Rüschlikon, Switzerland
ED R. CHILDERS, retired, IBM, Armonk, United States
DAVID PEASE, Retired, Chico, United States

Magnetic tape provides a cost-effective way to retain the exponentially increasing volumes of data being created in recent years. The low cost per terabyte combined with tape's low energy consumption make it an appealing option for storing infrequently accessed data and has resulted in a resurgence in use of the technology. Magnetic tape as a digital data storage technology was first commercialized in the early 1950's and has evolved continuously since then. Despite its long history, tape has significant potential for continued capacity and data rate scaling. This article strives to provide an overview of linear magnetic tape technology, usage, history, and future outlook. After a short introduction, the article delves into the details of how modern tape drives and media operate, including the basic mechanism and physics of magnetic recording, current tape media technology, state-of-the-art tape head technology, tape layout and encoding, data retrieval, timing-based servo and mechatronics of a tape drive, and the capabilities of current drives. This is followed by a discussion of tape libraries, an overview of tape library performance modeling research, operating system-level and application-level tape support, tape use cases, and the future scaling potential and outlook of tape. The article concludes with a history of tape hardware, media, usage and software.

CCS Concepts: • **Hardware → External storage**; • **General and reference → Surveys and overviews;**

Additional Key Words and Phrases: Magnetic recording, magnetic tape, tape libraries, tape software, tape scaling, tape history

Authors' Contact Information: Mark A. Lantz, IBM Research Europe - Zurich, Rüschlikon, ZH, Switzerland; e-mail: mla@zurich.ibm.com; Simeon Furrer, IBM Research Europe - Zurich, Rüschlikon, ZH, Switzerland; e-mail: sfu@zurich.ibm.com; Martin Petermann, IBM Research Europe - Zurich, Rüschlikon, ZH, Switzerland; e-mail: map@zurich.ibm.com; Hugo Rothuizen, IBM Research Europe - Zurich, Rüschlikon, ZH, Switzerland; e-mail: rth@zurich.ibm.com; Stella Brach, IBM Research Europe - Zurich, Rüschlikon, ZH, Switzerland; e-mail: stb@zurich.ibm.com; Luzius Kronig, IBM Research Europe - Zurich, Rüschlikon, ZH, Switzerland; e-mail: Luzius.Kronig@ibm.com; Ilias Iliadis, IBM Research Europe - Zurich, Rüschlikon, ZH, Switzerland; e-mail: ili@zurich.ibm.com; Beat Weiss, IBM Research Europe - Zurich, Rüschlikon, ZH, Switzerland; e-mail: wei@zurich.ibm.com; Ed R. Childers, retired, IBM, Armonk, New York, United States; e-mail: erc@sabinoerc.com; David Pease, Retired, Chico, California, United States; e-mail: pease@coati.com.



This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2025 Copyright held by the owner/author(s).
ACM 1553-3077/2025/01-ART6
https://doi.org/10.1145/3708997

ACM Reference Format:
Mark A. Lantz, Simeon Furrer, Martin Petermann, Hugo Rothuizen, Stella Brach, Luzius Kronig, Ilias Iliadis, Beat Weiss, Ed R. Childers, and David Pease. 2025. Magnetic Tape Storage Technology. *ACM Trans. Storage* 21, 1, Article 6 (January 2025), 70 pages. https://doi.org/10.1145/3708997

1 Introduction

Despite being one of the oldest data storage technologies, with a history exceeding 70 years, magnetic tape storage continues to thrive unlike many of its early contemporaries such as punch cards or the magnetic drum. Tape's enduring success lies in its sustained evolution. It has consistently scaled its capacity, data rate, and form factor, keeping it relevant for specific applications also today. Moreover, even though tape has been continuously scaled over its long history, tape has significant potential for continued future scaling.

Modern tape systems have a set of attributes that make the technology well suited for storing infrequently accessed data in applications such as cold and active archives, backup and disaster recovery. State-of-the-art enterprise-class tape drives operate with a native cartridge capacity of 50 TB and a native data rate of 400 MB/s [102]. Enterprise class automated tape libraries offer scalable capacities up to many hundreds of PBs and have data center floor space efficiencies of over 50 PB/m^2. The removable nature of tape media enables the capacity and data rate of tape libraries to be scaled independently. Removable media also provides a built-in physical air gap for added security that can be further enhanced by exporting cartridges to an off-site vault. In addition, the relatively high cost of a tape drive can be amortized over the cost of many cartridges, resulting in a low total cost per PB. At any given time, most of the cartridges in a library are stored in slots where they consume no power, leading to much lower running costs and to very low CO_2 emissions. For example, a recent study reported an 86% reduction in total cost of ownership and an 87% reduction in CO_2 emissions for a 10 PB tape archive, growing at 35% a year and stored for ten years compared to HDD [120]. Part of tape's cost advantage is derived from the use of a large area of media for recording. A state-of-the-art tape cartridge contains more than 1.3 kilometers of tape which enables a high capacity but leads to a high data access latency. For example, random seek operations on such a cartridge take on the order of 35 seconds on average. This combination of attributes makes tape well suited for storing large volumes of infrequently accessed data and has resulted in a resurgence in use of the technology particularly amongst hyperscale cloud companies.

The continued use and development of tape systems makes it important to provide a comprehensive overview of current tape technology and use. In the past, several articles and special journal sections have provided overviews of magnetic recording [6, 13, 17, 119, 203] and tape technology [9, 14, 23, 49–52, 56, 139, 177, 180] that provide a snapshot of the state-of-the-art at the time of their publication. This article has several goals: it aims at serving as an introduction and overview of tape technology starting from the basic recording principles and building through tape drive operation to tape libraries, tape software and the applications of current tape technology. In addition, we provide an outlook into the future of tape systems and an overview of the history of tape storage. To these ends, the article is organized as follows. In Section 2, we introduce the basic mechanism and physics of magnetic tape recording. In Section 3, we describe current tape media technology followed by a description of state-of-the-art tape head technology in Section 4. In Section 5, we discuss the tape layout and how data is written to and physically organized on tape as well as how user data is encoded and formatted. In Section 6, we describe data retrieval from tape, and how tape systems provide ultra-reliable data storage. Section 7 discusses the **timing-based servo** (**TBS**) technique and the mechatronic aspects of a tape drive. In Section 8, we describe the

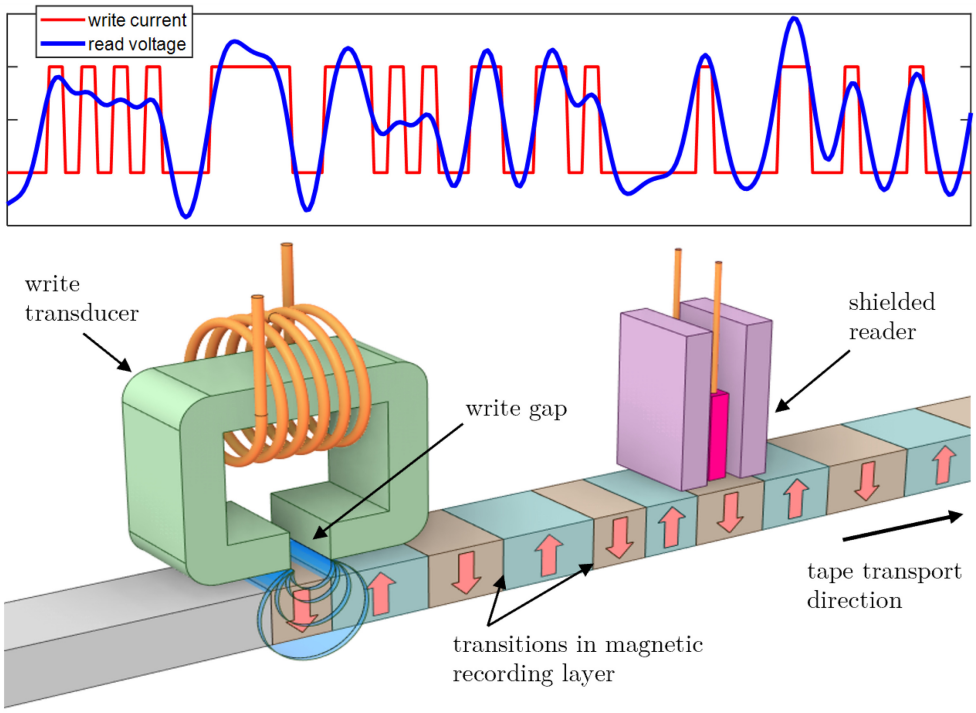


Fig. 1. Schematic of a magnetic tape recording system.

capabilities of current tape drives and discuss some unique aspects of tape systems. Section 9 describes current tape libraries and provides an overview of tape library performance modeling research. Section 10 discusses operating system-level and application-level tape support and special considerations for implementing applications that use tape. Section 11 discusses current and emerging tape use cases and available software options. In Section 12, we discuss the future scaling potential and outlook of tape. Finally, in Section 13, we provide a history of tape hardware, media, usage and software. In Section 3 to Section 7, we use LTO-9 as an example to illustrate the topic covered in each section, however over tape's long history many alternative designs have been implemented as discussed in Section 13 on tape history and the references therein.

2 Magnetic Recording

In magnetic tape storage technologies, digital information is stored in a magnetizable coating of the tape—the *mag layer*—as it is moved past an inductive write transducer. The essence of the recording process, as shown schematically in Figure 1, is to produce a pattern of reversals of magnetization down the length of the track such that information is encoded in the distances between the transition walls delimiting regions of alternating direction of magnetization [18, 143]. Figure 1 illustrates recording on perpendicularly oriented media that is currently used in state-of-the-art tape drives. However, for most of tape's history, longitudinal recording has been used in which the orientation of the magnetization of the mag layer is in the plane of the tape.

The data is subsequently read back using a second transducer comprising a **magnetoresistive** (**MR**) sensor which detects the field fringing out of the media, producing at its output a pattern of voltage pulses which can then be processed by the read channel to reconstruct the pattern of write current reversals originally input to the writer.

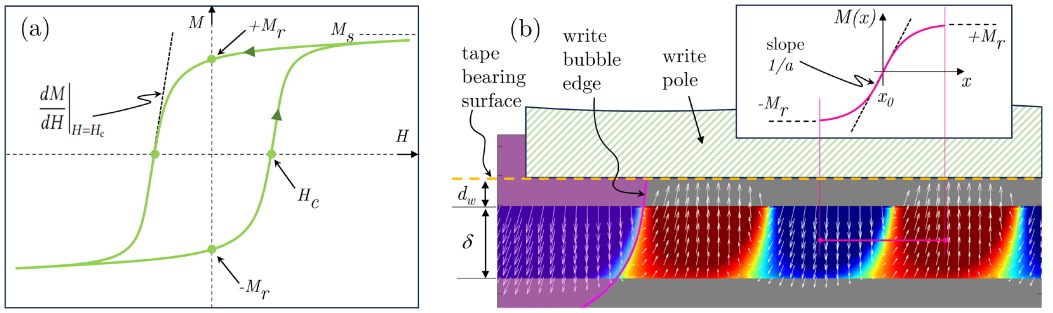


Fig. 2. (a) Typical hysteresis loop of magnetic recording media. M_s, M_r, and H_c are the saturation magnetization, remanent magnetization, and coercivity, respectively. The slope of the curve at the coercive points $\pm H_c$ is one measure of the *squareness* of the loop. (b) Calculated magnetic imprint for realistic writer materials, media particle properties, and spacings, for a periodic pattern of transitions.

Data retention on tape relies on the hysteretic magnetic properties of the particles used in the mag-layer of the tape. Figure 2(a) shows features of the magnetization curve for a sample containing an ensemble of "hard" magnetic particles such as is used in tape media [191, 193, 198]. The horizontal axis H corresponds to the external magnetic field applied to the sample, while the vertical axis shows its average magnetization response M. When the sample is subjected to an increasingly strong external field H the system evolves following the lower path: the particles in the media gradually align in the direction of the applied field until the average magnetization settles at a *saturation magnetization* M_s. When the magnetic field is subsequently decreased, the magnetization follows the upper path, settling at a *remanent magnetization* M_r when the external field returns to zero, and requiring a reverse field of magnitude H_c, called the *coercivity*, to flip half the particles and restore an on-average zero magnetization. Leveraging this hysteresis, data is written on tape by applying an external field with a magnitude that significantly exceeds the coercivity, so that when the field is removed, regions of the media are left in one of two states of remanent magnetization, $\pm M_r$.

The write transducer used in today's tape systems to produce the write field is essentially an electromagnet consisting of a coil and a ring-shaped high-permeability core, as shown schematically in Figure 1. One side of the core ring is polished flat, in order to allow bringing the tape in close contact with it, and discontinuous over a short length called the *write gap* from which an intense magnetic field fringes outwards when a current is supplied to the coil. The region of space in which this gap field exceeds the coercivity of the media is called the *write bubble*, and portions of the media which overlap with it are magnetized along the prevailing direction of the field within. As the tape is streamed by the head, magnetic transitions are created in response to step changes in polarity of the write current, which leave imprints shaped by the trailing edge of the write bubble.

Ideally, those transitions are perfectly abrupt and straight through the entire thickness of the media. In reality, due to the distributions of particle properties and orientations in the media and due to finite gradients of the writing field, magnetic transitions actually occur over a finite length scale a, called the *transition parameter*. Figure 2(b) illustrates the gradual nature of the transition along the direction of travel, also emphasizing that it broadens quickly with distance from the writer and can have curvature through the depth. In tape, both these effects result from a relatively large writer-to-medium spacing d_w and thickness δ of the magnetic layer compared with the size and shape of the write bubble.

A real transition is usually described using a simple analytical model, typically an arctangent function:

$$M(x) = \frac{2M_r}{\pi} \arctan\left(\frac{x - x_0}{a}\right), \tag{1}$$

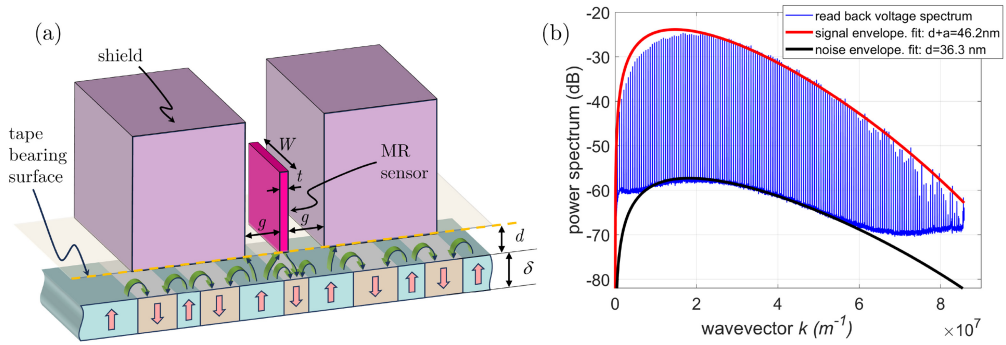


Fig. 3. (a) Schematic of the read back geometry for a shielded MR reader. The sensor slab has a width in the across-track direction W, a free layer thickness t, and is centered within a shield-to-shield gap of length $(2g + t)$. Spacings d and δ designate the distance to the top of the media layer and the media layer thickness, respectively. (b) Power spectrum of the read back signal from a PRBS pattern, with envelope shapes for the signal (SPS) and noise (NPS) spectra. The signal spectrum consists of narrow peaks reflecting the discrete values of frequency content of the 255 bit long PRBS pattern. The spacing d was estimated to be 36.3 nm by fitting Equation (4) to the noise envelope. This value is in agreement with what is expected from the media roughness combined with the pre-recesion and coating thickness of the head as discussed in Sections 3 and 4.

and related to the write field and media magnetization loop using a one-dimensional slope model [18, 198, 202] taken along a line at mid-depth of the media layer and evaluated at the location of its intercept x_0 with the edge of the write bubble:

$$\left.\frac{dM}{dx}\right|_{x=x_0} = \frac{2M_r}{\pi a} = \left.\frac{dM}{dH}\right|_{H=H_c} \left.\frac{dH_{\text{eff}}}{dx}\right|_{x=x_0}, \tag{2}$$

where H_{eff} is the *effective write field* which is essentially the gap field corrected for the angular dependence of the switching process for media with known particle distributions [125]. Equation (2) shows two essential requirements for a narrow transition (or, equivalently, for a small value of the transition parameter a):

— The slope of the hysteresis curve at the coercive point (which is also one measure of the *squareness* of the hysteresis loop) should be as large as possible. The difficulty with achieving large coercive squareness is that this requires very homogeneous properties and orientations of the particles in the media.
— The write field gradient at the location of the transition should be as large as possible. Due to the rapid decay of the magnitude of the field with distance from the write gap, this entails minimizing the spacing d_w (and the media thickness δ). In practice, however, there is a lower limit for this spacing due to the need to maintain a certain roughness of the tape surface in order to manage the tribology of the head-tape interface and prevent stiction [62].

Figure 3(a) shows the principal dimensions affecting the readback process, assuming a shielded MR head. The purpose of shielding the MR sensor slab is to shape its frequency-dependent sensitivity, much in the same way as an aperture in optics, and enable resolving the smallest spacing between transitions required to meet the system's targeted operating linear density.

The read back **signal power spectrum** (*SPS*) representing the envelope of a **pseudo-random binary sequence** (**PRBS**) recorded with a bit length B is derived using the reciprocity principle and analytical approximations for the shape of the magnetic field in the vicinity of the read gap

[18, 125, 127, 181, 198]:

$$SPS(k) \propto \left[M_r W \sin(kB/2) e^{-k(d+a)} \frac{(1-e^{-k\delta})}{k\delta} \frac{\sin(kg/2)}{kg/2} \frac{\sin(k(g+t)/2)}{k(g+t)/2} \right]^2, \tag{3}$$

where the wavevector k corresponds to a spatial frequency (i.e., the inverse of the spacing between transitions). This expression can be seen as a product of three frequency-dependent "loss" terms:

- the *gap loss*, $\frac{\sin(kg/2)}{kg/2} \frac{\sin(k(g+t)/2)}{k(g+t)/2}$, which confers its "aperture-like" shape to the spectrum: roll-offs at low and high frequencies with a maximum at roughly one third the signal bandwidth corresponding to the targeted linear density, see the "envelopes" of Figure 3(b).
- the *spacing loss*, $e^{-k(d+a)}$, an exponential decay term that dictates, through the generalized spacing $(d + a)$, the slope with which the spectrum rolls off towards high frequencies, and
- the *thickness loss*, $\frac{(1-e^{-k\delta})}{k\delta}$, that accounts for additional spacing losses associated with the finite thickness δ of the magnetic layer.

In a similar fashion, the power spectrum for uncorrelated particulate noise, the type of noise which is found to be dominant in tape, is given by the following expression (introducing additional parameters of particle volume fraction in the magnetic layer, p, mean particle volume, V, variance of the distribution of particle volumes, σ_V, and noise bandwidth, Δk) [125]:

$$NPS(k) \propto M_r^2 W k^2 e^{-2kd} \frac{(1-e^{-2k\delta})}{2k\delta} \left(\frac{\sin(kg/2)}{kg/2} \frac{\sin(k(g+t)/2)}{k(g+t)/2} \right)^2 \frac{V}{p} \left(1 + \left(\frac{\sigma_V}{V} \right)^2 \right) \Delta k. \tag{4}$$

Taking the ratio of these spectra highlights the three more critical parameter dependencies impacting the system's key performance metric, the readback **signal-to-noise ratio** (**SNR**):

$$SNR(k) = \frac{SPS(k)}{NPS(k)} \propto W e^{-2ka} \frac{(1-e^{-k\delta})^2}{k(1-e^{-2k\delta})} \left(\frac{\sin(kB/2)}{k} \right)^2 \frac{p}{V} \frac{1}{\left(1 + (\frac{\sigma_V}{V})^2\right)} k. \tag{5}$$

First, considering frequency-dependent terms, the written-in transition parameter a strongly dominates SNR performance at high frequencies, and therefore limits the achievable linear (i.e., along-track) density. While the read back spacing, d, does not appear explicitly, it is present implicitly in that a depends critically on the spacing at the time of writing, d_w, which is generally comparable to d. Furthermore, it can be seen that increasing the media thickness δ boosts SNR slightly, though mainly at lower frequencies. This apparent benefit is offset, in practice, by the regions of stronger transition curvature and degraded a parameter that come with additional depth. The term $(\sin(kB/2)/k)^2$ is the frequency-shaped "kernel" of the PRBS pattern for which the signal spectrum of Equation (3) is constructed. It does not present a means of improvement through changes of geometry.

Second, SNR performance scales proportionally to the across-track width W of the read sensor (providing lateral fringing effects can be neglected). This is an important design consideration when attempting to scale down the width of the track to increase system capacity.

Finally, the media formulation may be optimized for SNR by increasing p, reducing V, and, as is always beneficial, narrowing the size distribution (σ_V/V).

Although of the latter it would seem simplest to continue decreasing the volume of the particles, this eventually raises issues of thermal stability. The process of magnetization reversal of a particle is thermally driven and, in the absence of a magnetic field, follows an Arrhenius law with an energy barrier $K_u V$ and a time constant $\tau(T) = f_0 \exp(-K_u V/k_B T)$, where K_u is the particle's magnetic anisotropy constant, V its volume, k_B is Boltzmann's constant, and T is the absolute temperature in Kelvin (the prefactor f_0 is a materials-dependent thermal attempt frequency). For an ensemble

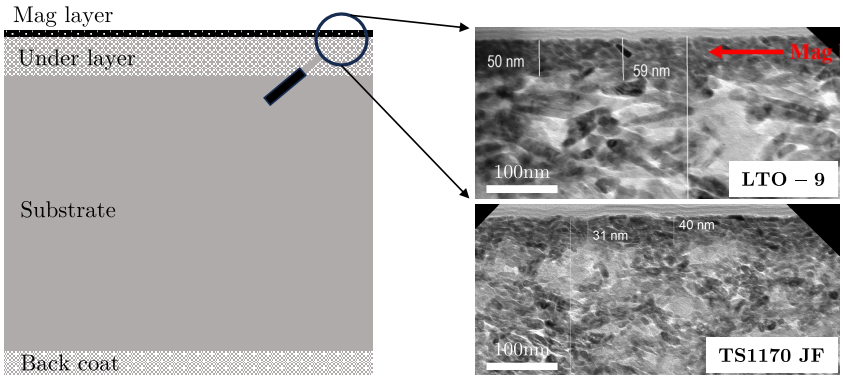


Fig. 4. Left: Schematic of tape media structure. Right: Transmission electron microscope images of cross sections of the mag-layer and under layer of LTO-9 (top) and TS1170 JF (bottom) tape.

of weakly interacting particles, such as in tape media, the net magnetization decays exponentially over time t from an initial saturation value M_0 as [179, 200]:

$$M(t) = M_0 e^{-\tau t} = M_0 e^{-t f_0 \exp\left(-\frac{K_u V}{k_B T}\right)}. \tag{6}$$

The nested exponent in this expression makes the rate of decay extremely sensitive to the energy ratio $K_u V / k_B T$, also known in the industry as the *thermal stability ratio*, which must be kept greater than ≈ 60 to ensure that data is retained over the several decades of archival time commonly expected from tape.

One means of decreasing the particle volume without compromising stability would be to increase the particle anisotropy K_u in compensation. As this entails a larger media coercivity H_c, the strength of the field required to saturate the media increases as well. Historically, tape media coercivity has been progressively increased as new generations of tape have been introduced. For example, the Fe_2O_3 media used in the mid seventies had a coercivity of 350 Oe [51], the metal particle based media used in the first few generations of LTO had a coercivity of 1,850 Oe [51] and the latest LTO-9 media has a coercivity of 2,850 [63]. This strategy, however, has limits because the strongest field that a writer can produce is limited by the saturation magnetization of the write poles on either side of the write gap, and the M_s of the CoFe-based materials of which these are made today are already close to the theoretical Slater–Pauling limit of 2.4 Tesla.

Hence, the need to decrease the particle volume to enhance SNR and enable higher recording densities conflicts either with the thermal stability of the media or with the ability to write it. This illustrates the so-called *trilemma of magnetic recording*, which refers to the fundamental challenge of balancing three conflicting system properties, readback SNR performance, stability, and writability, while pursuing higher data storage densities. The point at which the particles or grains become so small that thermal fluctuations become dominant, causing frequent data loss, is known as the *superparamagnetic limit*.

3 Tape Media

The tape media used in state-of-the-art tape systems is a composite material made up of four layers, as illustrated in Figure 4. The central layer, called the *substrate* or *base film*, is around 3.5 to 4.5 microns thick in recent generations of tape. The other three layers are much thinner, such that the substrate properties largely determine the overall mechanical properties of the tape.

Recent generations of tape have used substrates made of **Polyethylene terephthalate** (**PET**), **Polyethylene naphthalate** (**PEN**) or Aramid (aromatic polyamide). Tape is manufactured by coating the substrate with the other three layers using a roll to roll liquid coating process. The bottom layer, called the *back coat*, consists of a partially conductive polymer film that can dissipate the electrostatic charges that would otherwise accumulate during un-spooling and tape transport. The roughness of the back coat is engineered to reduce adhesion between the tape windings as it is un-spooled and to allow entrained air to escape from the tape pack when the tape is wound onto a reel. Bleeding out entrained air helps to reduce distortions of the tape pack that might otherwise occur.

The two upper layers, referred to as the *mag layer* and the *under layer*, are deposited simultaneously using a dual layer coating technology that enables the deposition of a very smooth, thin and uniform magnetic layer [154]. The under layer is non-magnetic and partially conductive to prevent the build up of charge due to tribo-charging effects as the tape runs over the head. In addition, the under layer serves as a reservoir for lubricant that can diffuse to the surface to reduce tape-head friction. The mag-layer, also referred to as the *recording layer*, is made up of small magnetic particles held together and fixed to the under layer with a binder. The mag-layer also contains a small amount of slightly larger non-magnetic particles that are used to tune the surface roughness of the tape to minimize the tape-head contact area and thus reduce friction. The amount of binder must be as low as possible to ensure a high density of magnetic particles, but still sufficient to hold the particles together and to bind the magnetic coating with the under layer. In order to achieve good recording performance the particles should be isolated from each other and have narrow distributions of particle size and coercivity. Particle isolation is achieved by adding a dispersant to the particle slurry that inhibits the particles from sticking directly together.

Currently there are two tape formats under active development: LTO[1] and IBM[2] TS11xx. The latest generation of LTO media (LTO Gen-9) uses Barium Ferrite (BaFe) magnetic particles [1, 68, 134, 190]. BaFe particles have a hexagonal platelet shape and a magnetization that results from crystalline order rather than shape anisotropy. Early generations of LTO used needle shaped CoFe particles with a magnetization arising from shape anisotropy and that required a protective coating to prevent oxidation. In contrast, BaFe ($BaFe_{12}O_{19}$) is an oxide and therefore does not require a non-magnetic coating to protect against oxidation. As a result, BaFe particles can be scaled to smaller sizes compared to previous generations of magnetic particles. The coercivity of the particles can be tuned by doping the particles with elements such as Co, Zi or Ti. The hexagonal platelet shape of BaFe particles leads to a preferential orientation of the particles relative to the surface of the tape which results in a partial perpendicular orientation of the mag layer. This orientation can be enhanced by applying a magnetic field during the coating process which results in a higher SNR [71]. Figure 4 shows a TEM micrograph of a cross section of LTO-9 media. The mag layer has a thickness of about 50–60 nm, the under layer is about 600 nm thick and the total tape thickness is 5.2 µm. The tape width is 12.65 mm and the length is 1035 m which enables a native LTO-9 cartridge capacity of 18 TB. The latest generation of Enterprise tape, IBM TS1170 JF media, uses a mix of BaFe and SrFe particles. Strontium ferrite is from the same family of hexagonal ferric oxides as BaFe, also has a hexagonal platelet shape but has a higher saturation magnetization and coercivity than BaFe. JF tape has a thickness of 4.0 µm, a width of 12.65 mm and a length of 1337 m which enable a native capacity of 50 TB. Figure 4 also shows a TEM

[1]LTO is a registered trademark of Hewlett Packard Enterprise, IBM and Quantum in the US and other countries.

[2]IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "IBM Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

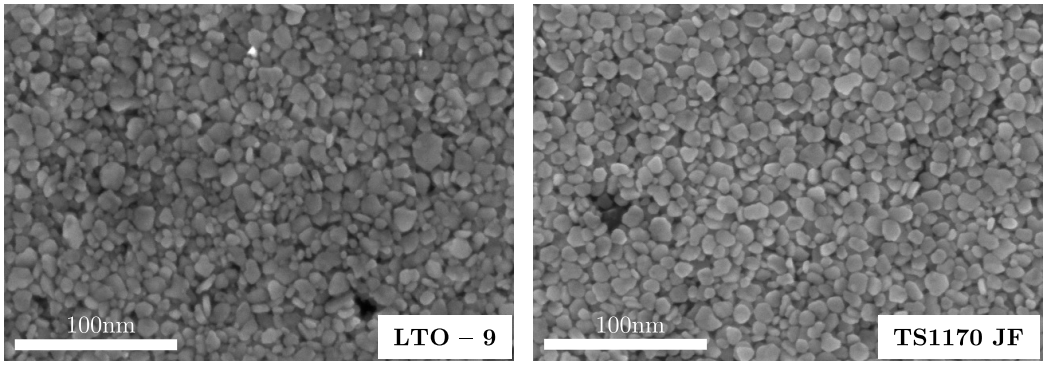


Fig. 5. Left: Top view SEM image of the mag layer of LTO-9 tape. Right: Top view SEM image of the mag layer TS1170 JF tape.

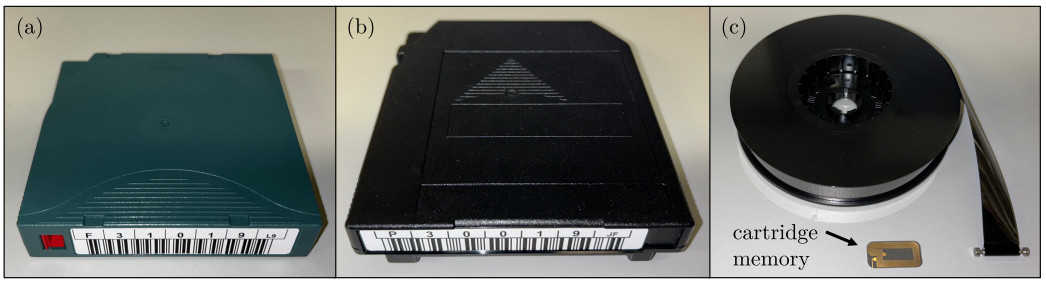


Fig. 6. (a) LTO-9 Cartridge, (b) IBM TS1170 JF Cartridge, (c) tape reel and CM from an TS1170 JF cartridge.

micrograph of a cross section of JF media and Figure 5 shows **scanning electron microscope (SEM)** images of the mag layers of LTO-9 and TS1170 JF tape.

In state-of-the-art linear tape drives, the tape is wound on a single reel which is housed in a plastic cartridge as illustrated in Figure 6. LTO cartridges have dimensions of 102.0 mm × 105.4 mm × 21.5 mm whereas the IBM TS11xx Enterprise tape uses a slightly larger and more robust cartridge design with dimension of 109 mm × 125 mm × 24.5 mm. The same cartridge form factor has been used for all generations of LTO and TS11xx tape media to date. Both cartridge types contain a clutch mechanism that maintains tension on the tape while it is stored in the cartridge. To facilitate the process of threading the tape over the head and onto a second reel housed within a tape drive, a section of a few 10's of cm of thicker tape, called *leader tape* is spliced to the end of the data tape. A metal pin, called a *leader pin*, is attached to the end of the leader tape and is used to pull the end of the tape out of the cartridge and to *thread* it across the head and onto the second reel. The front of the tape cartridge has a region used for a bar code label that can be read by the robotic mechanism of an automated tape library (see Section 9).

LTO and Enterprise tape cartridges contain a non-volatile memory, called a ***cartridge memory* (CM)**, which uses a noncontact passive RF interface. The CM stores a variety of information including a serial number, the media type, manufacturing information, and servo information as well as usage data which is updated by the tape drive, such as a write pass number and a directory of where data has been written on tape. Both cartridge types also contain a write protect switch, which when set to the locked position, prevents data from being written to the cartridge. Recent generations of LTO and TS11xx cartridges are also available in a ***write once read many* (WORM)**

format. Cleaner cartridges that can be used to clean contamination from the tape head are also available in LTO and Enterprise formats. *Cleaner cartridges* contain a slightly more abrasive tape formulation and are typically rated for a fixed number of cleaning cycles. During operation, the tape drive continually monitors the recording performance and will request a cleaner cartridge if potential head contamination is detected.

The continued scaling of tape areal density and capacity requires continued improvements in the SNR of the media. This can be achieved by reducing the size of the magnetic particles which reduces media noise, reducing the **switching field distribution** (**SFD**) of the particles by reducing the variation in particle size and particle coercivity, increasing the saturation magnetization of the particles, improving the dispersion of the particles such that each particle is magnetically independent of the other particles, increasing the degree of perpendicular orientation of the particles which is typically measured as a squareness ratio of the media, reducing variations in the thickness of the mag layer, and decreasing the thickness of the mag layer as the linear density is increased.

Another factor critical to achieving high-density recording capabilities is a reduction in the tape-head spacing [136]. This is primarily determined by the roughness of the tape, although it can also be influenced by tape tension, speed and wrap angles [62]. To reduce the spacing between the magnetic layer and the read elements of the head, the surface of the tape should be as smooth as possible. However, a more intimate physical contact between the tape and the head may result in an increase in friction, which can be detrimental to the overall performance of the storage system. In fact, high friction may excite longitudinal vibrations in the tape that can negatively affect the read channel and the track-following servo system and can thus lead to an increased bit error rate [28, 146, 206]. A smoother tape surface may also reduce the durability and the runnability of the media, due to increased friction and wear.

The friction and wear that result from the contact nature of tape recording is a major challenge in tape system design that increases as the media roughness is reduced. A promising way to reduce friction without compromising spacing is the careful design of both the short-range and long-range roughness of the tape surface. The former is mostly determined by the height and distribution of local asperities in the coating, while the long-range roughness or "waviness" of the tape at a larger length-scale results mostly from the "waviness" of the substrate. The combination of a low long-range and a moderate short-range surface roughness can result in a reduced spacing and hence improved recording, while maintaining low friction and good runnability of the tape [28].

Another method to increase the capacity of tape cartridges is to reduce the thickness of the tape to enable an increased tape length. This approach adds cost to the cartridge and hence is less efficient than scaling the areal density but has still been an important component of tape capacity scaling. Reducing tape thickness makes the tape more fragile and more sensitive to tension variations which has driven improvements in tape handling and tape tension control.

The sensitivity of tape to tension variations is part of a more general challenge referred to as ***tape dimensional stability*** (**TDS**). In state-of-the-art linear tape drives, 32 tracks are written and read back in parallel across slightly less than one quarter of the width of tape. The tape-recording head is fabricated on a hard ceramic material and the pitch between the transducers in the head is relatively constant except for small changes due to thermal expansion. In contrast, tape is a thin strip of polymer and its width changes in response to changes in temperature, humidity and tension. In addition, the width of tape can also change slowly over time due to the pressure in the tape pack that arises from the winding tension; a phenomenon known as *storage creep*. If data is written under one set of environmental conditions and then read back again later under different conditions, the pitch of the transducers in the read head may not match the pitch of the tracks on tape leading to an increased error rate, or in the worst case, to unrecoverable data. In the past, when track widths were on the order of microns to tens of microns, TDS effects could

Table 1. 2023 TDS Targets from the 2019 INSIC Tape Technology Roadmap [115]

TDS Targets (ppm)			
Substrate	PEN	PET	Aramid
Thermal	0	0	50
Hygroscopic	470	300	50
Storage Creep	100	100	50
Tension	160	400	50

be dealt with passively by using a reader that was about 1/2 to 1/3 the width of the track and thus providing a significant margin for track misregistration. As track densities were increased, media manufacturers made incremental improvements in the dimensional stability of the media. In state-of-the-art tape drives, track widths are on the order of one micron or less and TDS effects are actively compensated for as described in Section 7. In this case, the span of the head and the dimensional stability of the media determine how much compensation range is needed. TDS depends on the substrate material and is typically specified in terms of parts per million (ppm) of normalized width change over a specified environmental range. The 2019 INSIC Tape Technology Roadmap [115] lists TDS targets for 2023 for the three substrate materials currently in use that are summarized in Table 1.

4 Tape Head Technology

The recording head is a key component of the tape drive. In early tape drives, write and read operations were performed using a single inductive transducer, whereas in more recent tape drives, these operations are performed using two different types of transducers that are referred to as *writers* and *readers*. In modern linear tape drives, multiple data tracks are written and read back in parallel using linear arrays of writers and readers that are manufactured using thin film microfabrication techniques on ceramic (AlTiC) wafers. AlTiC is a hot isostatic pressed mixture of crystalline aluminum oxide and titanium carbide powders. The use of thin film technology has enabled the size of the readers and writers to be continuously scaled to smaller dimensions, which has been key to enabling tape areal-density scaling.

Tape writers use a so-called *ring head* design that is illustrated conceptually in Section 2. The geometry of a state-of-the-art writer is illustrated in Figure 7. The core (yoke) is made primarily from $Ni_{45}Fe_{55}$. In the region of the write gap, the core narrows to concentrate the flux. The sections of the core on either side of the write gap are called the *write poles*, with the lower and upper pole sections referred to as P1 and P2, respectively. In recent generations of tape drives, additional thin layers of CoFe have been added on either side of the write gap. These high-moment liners have a larger saturation magnetization than the bulk of the poles, which results in stronger stray fields and enables recording on media with increased coercivity. The high-moment liners also lead to stronger field gradients, which results in sharper magnetic transitions in the mag layer. As the tape is streamed over the writer, it first passes over the P1 pole followed by the P2 pole. Hence the width of the P2 pole effectively determines the width of the written track. The P1 pole is *notched* to the same width as the P2 pole in the region of the write gap to minimize side writing/erasing due to fringing fields. The coil is made from copper and has two layers of 7 turns for a total of 14 turns. The use of multiple coil turns reduces the voltage required to drive the writer and also helps to limit the temperature rise in the head when the writers are driven.

In state-of-the-art tape drives, writers and readers are fabricated on separate wafers that are diced into writer and reader chips. Figure 8 shows an optical micrograph of a writer chip that

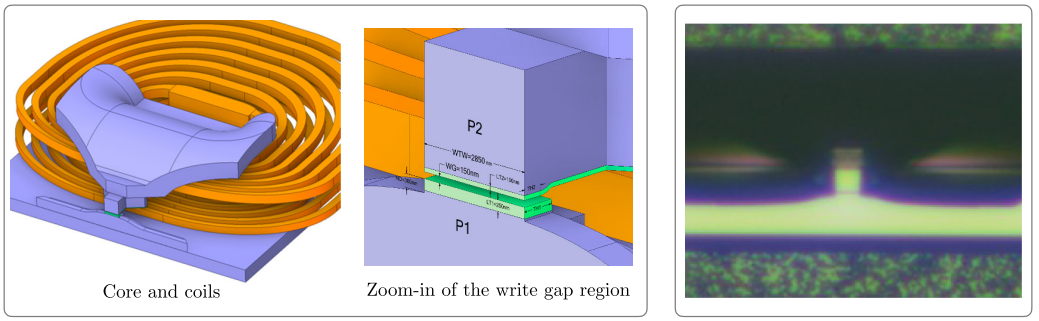


Fig. 7. (a) 3D Schematic illustration of a state-of-the-art tape writer: the core, coils and high-moment liners are shown in purple, orange and green, respectively (colors are shown in the online version of the article). (b) Optical micrograph of a tape bearing surface view of a writer. The AlTiC substrate and closure are visible in the bottom and top of the image, respectively.

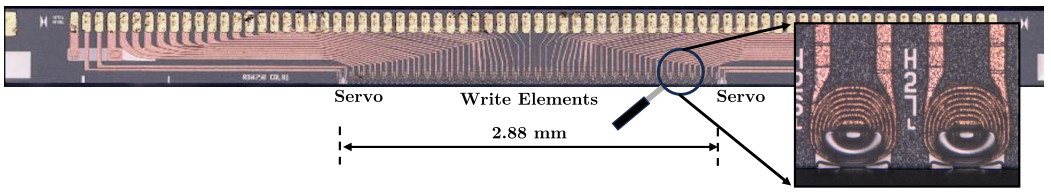


Fig. 8. Optical micrograph of the central part of a writer chip with a zoomed image of two writers shown on the right. The gold bond pads are visible along the top edge of the chip and the writers are centered on the bottom edge and spaced at a pitch of 83.25 microns.

contains 33 writers spaced at a pitch of 83.25 microns, one of which (the first or the last in the array, depending on the tape travel direction) remains unused during operation. This architecture was adopted to enable backwards compatibility with previous 16 channel formats [21]. Servo readers are located on each side of the writer array and are used to read factory-formatted tracks of a TBS pattern, as discussed in Section 7. Each writer and servo reader is wired to a pair of gold pads at the top of the chip that are used to wire bond the chip to a *flex cable*. The flex cable is a flat, flexible, ribbon cable that connects the writers and servo readers in the chip to the main electronics card of the tape drive while enabling the dynamic positioning of the head as described in Section 7.

Figure 9 depicts the structure and operating principle of a tape reader. Recent generations of tape drives use shielded **tunneling magnetoresistance** (**TMR**) sensors to read data written to tape. The magnetoresistance effect refers to a change in electrical resistance due to the application of a magnetic field. A TMR sensor can be viewed as a three-layer component with two electrically conducting ferromagnetic layers separated by a thin, non-magnetic, electrically insulating tunnel barrier as shown in Figure 9. The magnetization direction of one of the ferromagnetic layers, known as the *pinned layer*, is fixed. The magnetic orientation of the second ferromagnetic layer, referred to as the *free layer*, is free to rotate under the influence of the magnetic field emanating from the media. When a bias is applied across the two ferromagnetic layers, electrons tunnel through the barrier from one ferromagnetic layer to the other. The electrical resistance of the sensor decreases when the magnetization direction of the free layer is aligned parallel to that of the fixed layer, whereas it increases as they become more misaligned. The actual construction of a TMR sensor is more complex than this simplified three layer picture and a more detailed description of the materials, structure and operating principle can be found in Maat et al. [142].

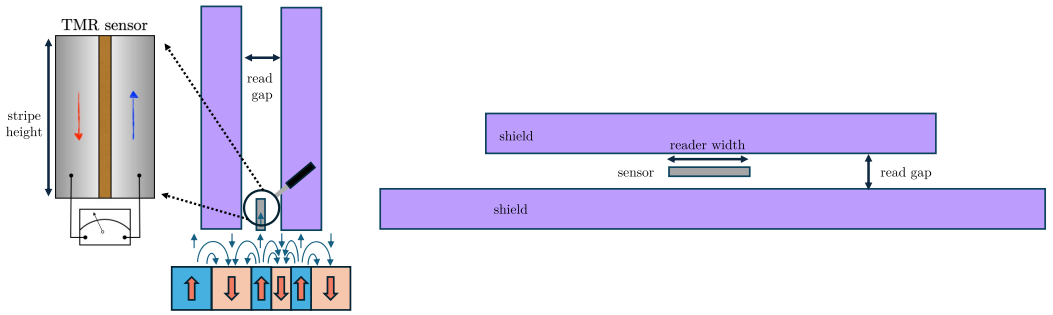


Fig. 9. Illustrations of a shielded tape read sensor. The shields are shown in purple and the sensor in grey (colors are shown in the online version of the article). Left: Cross-section view through the center of the reader and mag layer. Right: Tape bearing surface view of a reader.

The TMR sensor is positioned between two magnetic shields, shown in Figure 9. The shields are made from a high-permeability material such as Sendust or Permalloy that "shields" the sensor from the magnetic fields produced by bits adjacent to the region of tape directly under the sensor. The distance between the two shields is known as the *read gap* and is typically on the order of twice the minimum distance between transitions. The height of the sensor in the direction perpendicular to the surface of the media is called the *stripe height*.

Although based on a similar technology, the TMR readers used for tape have significant differences from those used in **hard disk drives** (**HDD**) [191]. These differences arise primarily from the lower areal density of tape relative to HDD. For example, the TMR readers used in state-of-the-art tape drives have a 3x–4x larger read gap than current HDD readers. In addition, the sensor width is about 10x larger than in HDDs and the sensor area (width x stripe height) is about >400x larger [21, 115]. These differences make it less challenging to fabricate TMR sensors for tape and also result in an electrical resistance that is much lower than HDD readers which facilitates the design of the amplifier in the analog front end of the tape read channel.

The layout of a reader chip is similar to that of a writer chip. The chip contains a linear array of 33 data readers, one of which is not used during operation, similar to the writer modules [21]. Servo readers are also positioned at either end of the array. Reader and writer chips are further processed by gluing a hard ceramic *closure* on top of the array and then the tape bearing surface is lapped to achieve a smooth, flat surface and to achieve the desired stripe height for the readers and the desired throat height for the writers. Processed writer and reader chips are glued onto ceramic *u-beams*, flex cables are then glued to the u-beams and the chips are wire bonded to the cables. Finally these cabled modules are assembled into a tape head.

State-of-the-art tape drives use a three-module head architecture referred to as a *terzetto* head [19, 22, 66]. The design consists of two writer modules with a reader module positioned in between as illustrated in Figure 10. This design enables read-while-write verification functionality while also enabling the independent optimization of reader and writer modules. When tape moves in the forward direction, the left writer module is used to write data which is immediately read-verified as the tape passes over the center reader module. When tape moves in the backwards direction, the right writer module writes data which can then also be read-verified by the center reader module.

In tape storage systems, the media is in physical contact with the head, which causes friction and leads to wear of both the tape and the head. In the absence of countermeasures, the read sensors are vulnerable to damage due to abrasive wear by the hard, non-magnetic particles present in the mag layer, and by media defects. A manifestation of damage is a reduction in the sensitivity of a read sensor due to partial electrical shorting. Pre-recession and coating are used to protect the

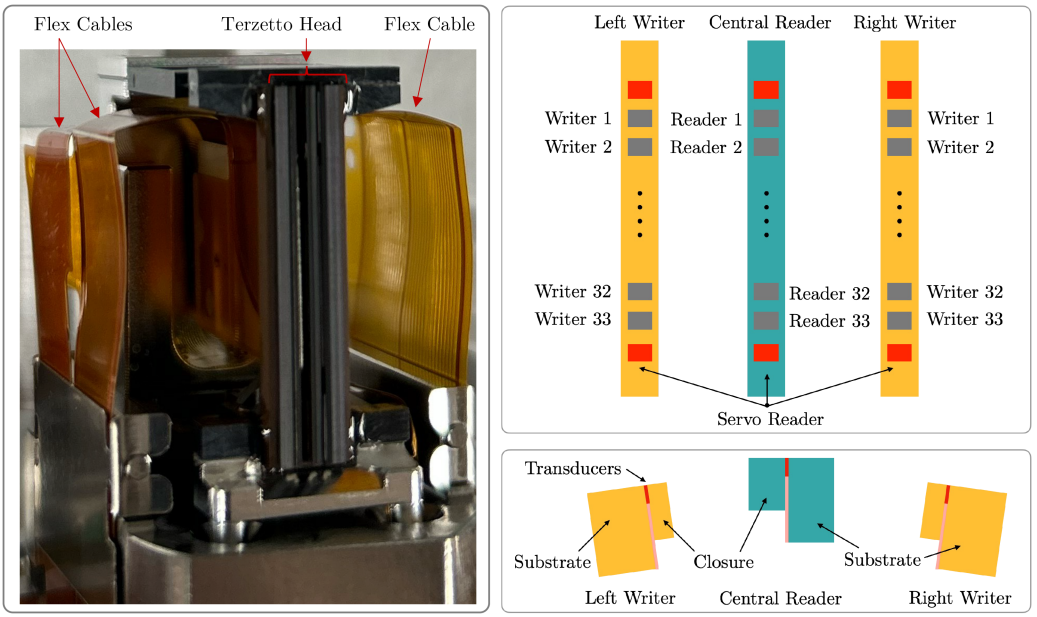


Fig. 10. Left: Photograph of a terzetto head with flex cables mounted in a track-following actuator. Right: View of the tape-bearing surface (top) and cross-section of the head modules (bottom).

active elements of the head [19, 21, 174]. After lapping to a target stripe height, the head chips are subjected to a sputtering/milling process which recesses the transducers relative to the tape-bearing surface. A hard coating material such as crystalline aluminum oxide is then deposited on the head to fill the recessed gap. The coating protects against both corrosion and wear, however the pre-recession and coating both increase the spacing between the mag layer and the surface of the sensor, which may limit the achievable recording density.

In the terzetto head design, each module has a flat profile and so-called *skiving edges* [22, 60]. The top of Figure 11(a) illustrates tape being streamed over a reader module. As the tape runs over the module, the leading skiving edge skives (scrapes) off the air that is dragged along by the tape. A region with sub-ambient pressure forms between the tape and the surface of the module. The resulting pressure difference pushes the tape into contact with the module. A so-called *tape tent* forms at the edge of the module due to the finite bending stiffness of the tape. If the head is symmetrically wrapped, as is the case for the reader module, a second tape tent forms at the trailing edge (see center illustration of Figure 11(a)). The amount of pressure at the skiving edge, the extents of the tape tents, and the resulting length of tape-to-head physical contact, are all strongly dependent on tape tension, speed and on the angle with which the tape is wrapped around the module, called the *wrap angle* [60, 62, 172].

The three modules of a terzetto head are positioned and wrapped with tape as illustrated in the bottom of Figure 11(a). The nominal wrap angles of the outermost edges of the writer modules are ~1° and the inner edges are wrapped at ~0°. The reader module is wrapped symmetrically at approximately ~0.6° on both edges. This configuration was developed to reduce wear of the writer modules and to reduce tape-head friction. To understand how this is achieved we consider the case of forward tape motion, i.e., from left to right as illustrated in the bottom of Figure 11(a). The tape first contacts the left edge of the left writer, a tape tent forms at the leading (left) edge due to the positive wrap angle and tape is pushed into contact with the module. Tape leaves the

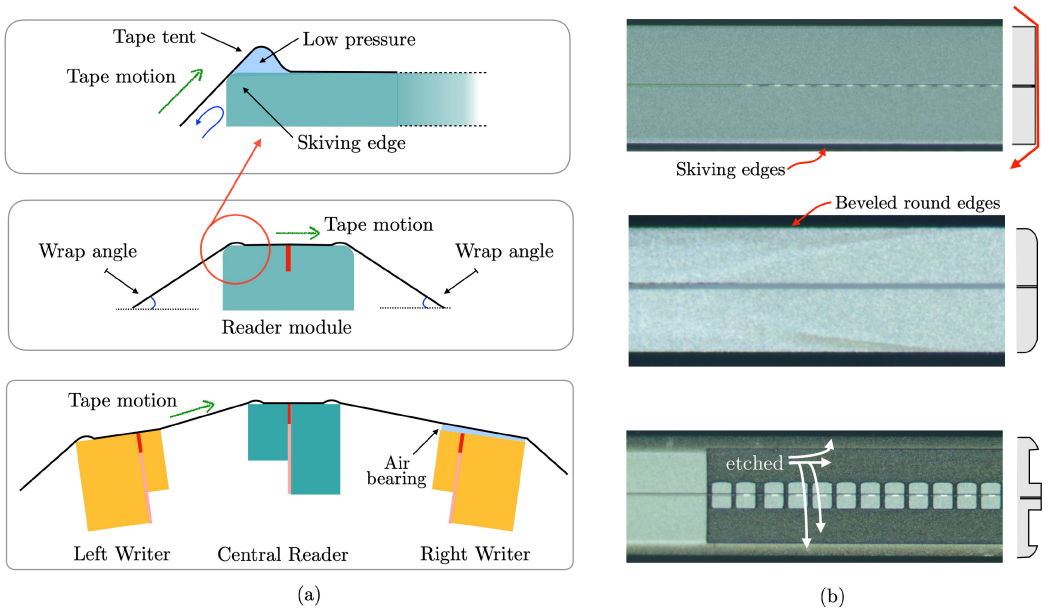


Fig. 11. (a) Top: Illustration of the skiving effect that pushes tape into contact with the head. Center: illustration of forward tape motion over a reader module. Bottom: Illustration of forward tape motion over a terzetto head. (b) Micrographs of the central section of a reader module with sharp skiving edges [61] (top), the left, outer beveled region of a module [28] (center) and an etched-vacuum head design [61] (bottom).

writer module parallel to its surface without forming a tape tent. When the tape reaches the reader module, the skiving effect due to the positive wrap angle on the leading (left) edge results in the tape being pushed into contact with the module as shown in the top of Figure 11(a). When tape reaches the leading (left) edge of the right writer module, the wrap angle is ~0° which is not sufficient for skiving to occur. Consequently, an air bearing—a thin film of entrained air—forms between the tape and the module, preventing direct physical contact. When the tape moves in the opposite direction, the roles of the left and right writer modules are reversed and an air bearing forms over the left writer module. In this manner, moving tape only contacts two of the three modules and three of the six skiving edges, significantly reducing friction and wear compared to previous configurations.

To continue scaling the areal density of tape systems, the spacing between the mag layer of the tape and the recording elements of the head must be reduced. This can be achieved by reducing the pre-recession and coating thickness. However, this will make the head more sensitive to wear damage and therefore will require the development of less abrasive tape. Spacing can also be improved by making the tape smoother (see Section 3), although this tends to increase friction. Most of the friction is caused by the skiving edges [62], hence a simple method to reduce friction is to *bevel* (or round) the skiving edges at the periphery of the head away from the transducer array [28]. This allows the tape to achieve intimate contact with the active elements in the central part of the head (where the skiving edges are preserved), while still supporting the tape by an air bearing in regions where contact is not needed. This concept is illustrated in the top and middle of Figure 11(b). Current generations of tape drives use beveled reader modules and un-beveled writer modules. The tape areal-density recording demonstration reported in Cherubini et al. [28] used a terzetto head in which all three modules were beveled to enable the use of a very smooth prototype media.

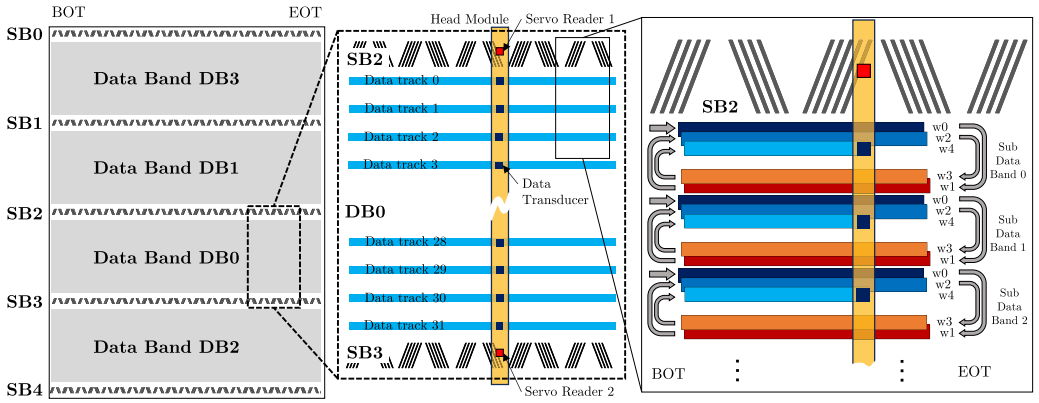


Fig. 12. Left: Tape layout showing SBs and DBs. Middle: Head module with servo and data readers located in DB 0 (zoom-in). Right: Serpentine recording operation showing two-and-a-half forward and two backward wraps/tracks (zoom-in showing three out of 32 sub DBs).

Another alternative head design, called *vacuum head* [61], is shown in the bottom of Figure 11(b). The main differences compared to a conventional or beveled flat head are: (i) the full span of the module is beveled, so there are no skiving edges, and (ii) a cavity is etched into the tape-bearing surface, leaving islands containing the active elements. Based on the same physics governing the air flow at the skiving edges, a sub-ambient air pressure develops inside the surface cavity as the tape is streamed over it. As a result, the tape is pushed into contact with the read/write elements.

5 Tape Layout and Data Recording

In this section, we discuss the tape layout and how data is sequentially written to and physically organized on tape media. Furthermore, we explain how user data is encoded and formatted for reliable storage on tape media.

5.1 Tape Layout and Data Track Recording (Write/Read Access)

In state-of-the-art tape systems such as the latest generation of **Linear Tape Open** (**LTO**) format, LTO-9, the 12.65 mm wide tape is organized in four ***data bands*** (**DBs**) which span the whole length of the tape, i.e., from the ***beginning of tape*** (**BOT**) to the ***end of tape*** (**EOT**). The four DBs are sandwiched between five dedicated ***servo bands*** (**SBs**), which contain a chevron-like TBS pattern [10] as shown in Figure 12. The factory pre-formatted TBS patterns are primarily used to estimate the *lateral position* (y-position) of the tape head relative to the tape [29], but also provide other essential servo parameters such as tape velocity, tape-to-head skew, tape width (TDS), SB identification, and tape ***longitudinal position*** (**LPOS**). The latter is encoded by means of position modulation of 4 out of 18 servo stripes that define a servo frame.

To write/read data to/from tape, the tape drive uses a stepper motor to laterally coarse position a three module tape head over the desired DB, such that the two servo readers are located on the two adjacent TBS patterns. Once the *track-following control system* has fine-positioned the tape head relative to the reference trajectory of the TBS pattern, and a desired tape LPOS has been reached by the *tape transport control system*, data can be written/read by means of the 32 active data transducers located between the servo readers. A set of 32 data tracks that are written in parallel along the full length of tape is called a *wrap*. Note that the 32 parallel tracks that define a wrap are equally distributed across the full width of the DB. Each track is located in one of 32

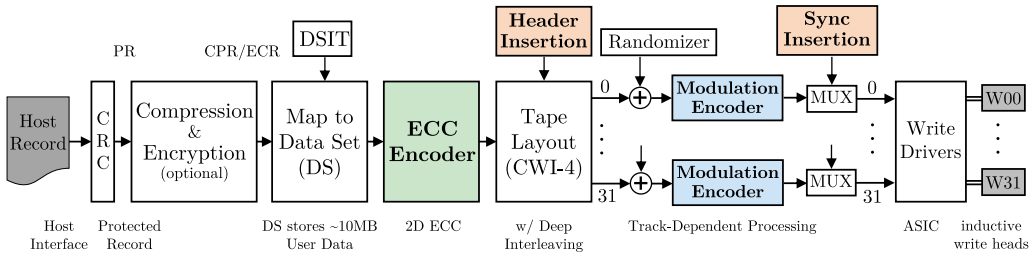


Fig. 13. Data encoding and formatting steps from host records to magnetic transitions on tape.

*sub-DB*s, each having a width of 83.25 µm corresponding to the pitch of the data transducers. The width of the servo patterns is nominally 93 µm and thus provides sufficient range to place the tracks at any desired lateral offset in each sub-DB.

LTO tape drives employ *linear serpentine recording*, where a large number of wraps are sequentially written to each DB, see Figure 12 (Right). Writing to a blank tape starts in DB 0. First, wrap 0 is written in the forward direction (BOT to EOT), placing 32 parallel tracks at the top of each sub-DB. Second, wrap 1 is written in backward direction (EOT to BOT), writing the tracks at the bottom of each sub-DB. Next, wrap 2 is written in forward direction, laterally offset by a track pitch relative to wrap 0 towards the sub-DB center, thereby *shingling* (partially overwriting) the previously written wrap 0 tracks to its final track width. Similarly, backward wrap 3 shingles wrap 1 to its final *shingled* track width, again offset towards the sub-DB center. This process of writing forward and reverse wraps in a spiraling-in manner continues until DB 0 is filled. Writing proceeds in the same manner subsequently in DBs 1, 2, and 3. The LTO-9 format specifies 70 wraps per DB and thus a total of 280 wraps per cartridge, offering a total of 18 TB native cartridge capacity.

The use of shingled magnetic recording enables backward write compatibility. LTO-9 drives can read and write to both LTO-9 and LTO-8 media types. This backward-compatibility allows tape administrators to upgrade to the newest formats immediately without fully migrating all archives. To write the previous generation's (wider) track width, the writer width significantly exceeds the current generation track pitch. On the other hand, the shingled recording nature prevents in-place overwriting of existing data because this would inevitably overwrite adjacent tracks/wraps. Tape systems are designed as *append-only storage*, where new data can be appended to where data writing previously ended, but the existing data is immutable. It is possible to *partition* the tape longitudinally or into groups of wraps, thereby creating independent storage pools and append points. Longitudinal partitioning requires a guard gap between partitions. For wrap-based partitioning, depending on the choice of partition sizes, *guard wraps* may be necessary between partitions. Partitioning using *DB boundaries* does not require guard wraps.

5.2 Data Encoding and Format for Reliable Data Storage on Tape

The host organizes data to be written to tape as *host records* of up to 16 MB in size and *host file marks*. To map a host record to magnetic transitions on tape, the tape drive applies a number of data processing and formatting steps highlighted in Figure 13. The goal of an efficient tape format is to introduce as little redundancy and overhead as possible while ensuring highly reliable data storage and retrieval.

When receiving a write request for a host record, the drive first appends a 4-byte ***cyclic redundancy* check (CRC)** creating a *protected record*, followed by data compression using an algorithm known as ***Streaming Lossless Data Compression* (SLDC)** based on Standard ECMA-321 with extended history buffer. Optionally, the ***compressed protected record* (CPR)** is encrypted using

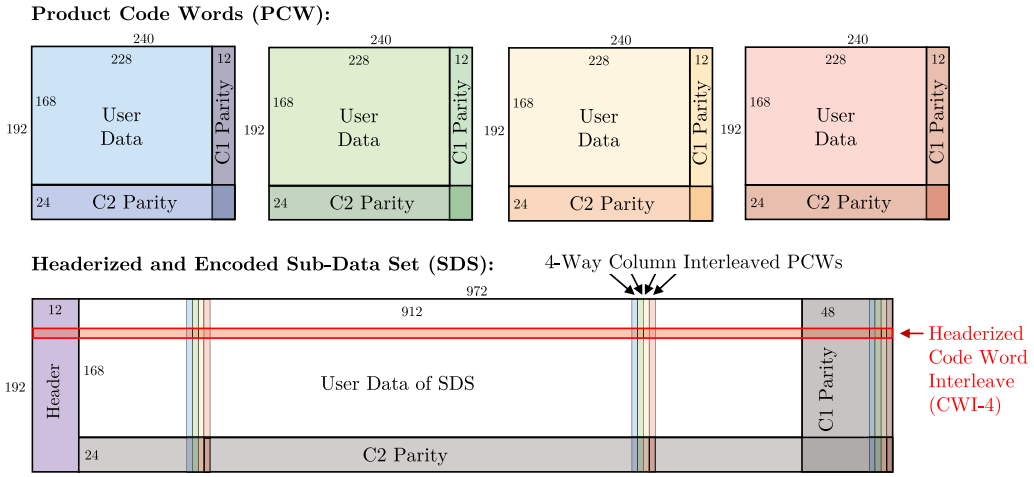


Fig. 14. Structure of PCWs, SDS, and headerized code word interleaves.

AES encryption, leading to an ***encrypted compressed record*** (**ECR**). Next, the stream of formatted records, file marks, and control symbols are broken into ***datasets*** (**DS**) of 9,804,912 user bytes to which a 912-byte ***dataset information table*** (**DSIT**) is appended. A dataset represents the smallest complete unit of information that is written to, or received from, the tape.

Each DS is further broken into 64 ***sub datasets*** (**SDS**). The construction of a SDS including ***error correction code*** (**ECC**) encoding is detailed in Figure 14. The process starts with four matrices of size $K_2 \times K_1$ bytes, where $K_2 = 168$ and $K_1 = 228$ bytes, from the stream of formatted records, labeled as "user data" in Figure 14. Each matrix is subsequently ECC encoded into a ***product code word*** (**PCW**) by means of two orthogonal encoding steps. First, a systematic column encoding using a RS($N_2 = 192, K_2 = 168$) Reed-Solomon code appends $N_2 - K_2 = 24$ rows of so called *C2 parity* bytes. Then a systematic row encoding with a RS($N_1 = 240, K_1 = 224$) code adds $N_1 - K_1 = 12$ columns of *C1 parity* bytes. The resulting ECC PCWs have a total size of $N_2 \times N_1 = 192 \times 240$ data and parity bytes.

The encoded SDS is constructed by means of 4-way column interleaving of the four PCWs as shown in Figure 14. In preparation for the *tape layout* mapping, a 12-byte header field is prepended to each row of the SDS, header protection is applied without overhead, turning each row of the SDS into a 972-byte headerized *code word interleave* (CWI-4), as highlighted in Figure 14. There's a total of 12,288 CWI-4s per dataset.

Referring back to Figure 13, ECC encoding and CWI-4 header insertion is complete. The *tape layout* mapping step defines how the CWI-4s are ordered, assigned to the 32 parallel tracks, and eventually mapped to tape. The chosen mapping provides *deep interleaving* [36] to mitigate spatial bursts errors on the surface of the magnetic tape. Deep interleaving refers to spreading the information both laterally across all parallel recording tracks as well as longitudinally far along the tracks. A first property of deep interleaving aims at rendering the error symbols at the input of the C1 and C2 component decoders to be uncorrelated/independent of each other. A second property is that data can still be decoded correctly when a lateral stripe error affects many millimeters of tape surface in down track direction, for example due to instantaneous tape speed variations. Last but not least, the interleaving ensures that data can be decoded correctly even if up to four of the 32 tape tracks have simultaneous read errors, for example, due to disfunctional readers. As a result of the deep interleaving, the CWI-4s from the same SDS are separated by about 1 mm on tape. Similarly, the bytes of each C2 column code word are uniformly spread across the full dataset.

Referring to Figure 13 again, now that the CWI-4s have been assigned to tracks, each track's data stream passes through a data randomizer, followed by a rate-32/33 *modulation encoder*. The modulation code, often referred to as a ***runlength-limited*** (**RLL**) code, satisfies various run-length constraints on the maximum length of zeros [114], transition-runs [150], but also synchronization patterns used for timing acquisition [34] and twins pattern to limit the path memory length of the sequence detector [35]. The modulation encoded bit streams are extended with synchronization information, and passed to the write driver ASIC, which generates the write currents for the individual inductive tape write heads. The write clock frequency is selected in a tape-speed dependent fashion to achieve a written bit length $B = 46.6$ nm, as specified by LTO-9.

Some of the data encoding and formatting steps discussed above lead to overhead. The largest overhead is due to ECC, i.e., the linear product code with rate $(228/240) \times (168/192) = 0.83$. The modulation code, a nonlinear constrained code, adds a rate $32/33 = 0.97$. Combined with header insertion and sync patterns, the total overhead is about 21%.

5.3 Read-While-Write Verification and Rewrite

Modern tape drives such as LTO-9 use another mechanism to further improve reliability and data integrity: *read-while-write verification* and *rewrites*. This technique uses a set of readers placed downstream of the write transducers to read back data immediately after it has been written to tape. The readback signal from each reader is processed by the read channel, modulation decoder and C1 decoder. If the number of C1 byte errors in a given CWI-4 exceed a threshold level, then the CWI-4 data is rewritten to a new location further down tape. The actual rewrites happen at the end of each dataset, i.e., after the 12,288 first-written CWI-4s have completed.

The number of rewrites per dataset is dynamic, but about 3% of the total cartridge capacity is reserved for rewrites. The average rewrite overhead for a new cartridge is less than 1%. If more than 3% rewrites are necessary, the cartridge may not be able to achieve its advertised capacity.

It may have become clear that the combination of recording techniques discussed in this section mean that modern tape systems are designed using a *write friendly* architecture. For example, the tape drive applies a mechanism called *speed matching* to match the host data rate with the native write speed, so that frequent write stops and back-hitch operations to restart writing can be avoided. In the next section, we discuss data retrieval.

6 Reliable Data Retrieval from Tape

In the previous section, we discussed the tape layout and how data is written to and physically organized on tape media. In this section, we examine data retrieval from tape, and how, despite noise and errors in the detected bits, tape systems provide ultra-reliable data storage and data retrieval from tape.

6.1 Data Retrieval: Readback Signal, Detection, and Decoding

Figure 15 shows the data flow from the 32 MR readers scanning over the magnetic transitions of a data wrap on tape to the retrieved host data record provided via the host interface. Starting on the left side of Figure 15, a set of 32 shielded TMR readers are connected to a pair of analog/mixed-signal ASICs commonly referred to as the analog front- and backend, which implement the electrical biasing of the TMR readers, and act as a parallel bank of readback signal pre-amplifiers, variable gain amplifiers, filters and **analog-to-digital** (**A/D**) converters. To support a range of tape speeds for write and read operations, the analog filters have programmable cut-in and cut-off frequency settings to match the signal bandwidth and further provide variable high-frequency boosting of the readback signals. Furthermore, the programmable band pass filters avoid signal anti-aliasing during subsequent asynchronous A/D conversions at a rate of approximately $4/(5T)$, where T is

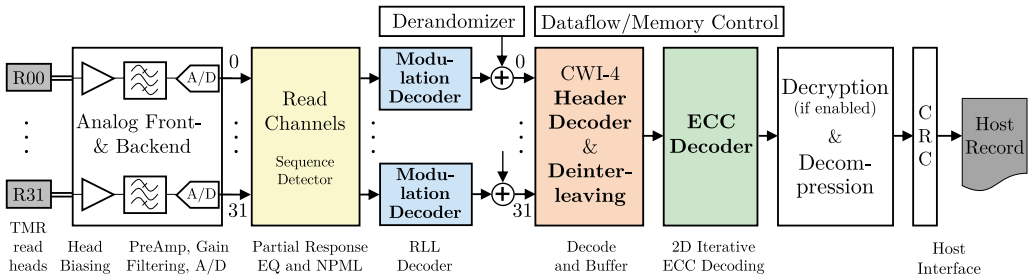


Fig. 15. Readback data flow from the tape read heads to the host interface/records.

the minimal bit temporal duration. This 25% oversampling is typical for tape drives and is a safety margin to cope with variations in tape velocity and linear density. The temporal bit length T and the spatial bit length B are related as $T = B/v_s$, where v_s represents the tape speed.

Shielded MR readers are designed such that the readback signals respond primarily to transitions of the magnetization patterns on tape. At low bit densities, the transition responses had been clearly separated such that reading the recorded data could be accomplished by detecting the peaks of the transition responses by means of a *peak detection* channel. Today, because of the short bit length $B = 46.6\,\text{nm}$ in LTO-9, corresponding to 545 kbpi (kilo bits per inch) linear density, and the finite head resolution, the superposition of isolated transition responses partly cancel each other, causing a known phenomenon called ***intersymbol interference*** (**ISI**). To efficiently combat ISI, magnetic recording channels in tape and HDD products introduced ***partial response*** (**PR**) signaling with ***maximum likelihood*** (**ML**) sequence detection, or PRML detection [34, 64, 132]. A PRML read channel targets a controlled small amount of ISI at the output of the channel **equalizer** (**EQ**), which is left to the sequence detector to combat. This concept avoids full channel equalization and prohibitively strong noise enhancement at high bit densities. Modern tape read channels such as those in LTO-9 implement advanced versions of PRML and **noise-predictive maximum-likelihood** (**NPML**) [5, 37, 58] detection schemes, aiming at minimizing the influence of correlated and data-dependent noise in the detection process.

Read channels for tape drives require a high level of adaptivity to deal with media variations, such as fluctuations in particle dispersion and magnetic layer thickness, as well as variations in head-media spacing, wear of the head and media, head-tape interface related friction and tape velocity fluctuations. To cope with these temporal and spatial variations, the majority of signal processing and detection algorithms in tape read channels operate in a highly dynamic and adaptive manner. Examples include adaptive channel equalization, dynamic multi-channel timing recovery/synchronization, fast gain control, and noise-adaptive NPML detection [58].

Furthermore, the nature of parallel track recording in tape drives, coupled with the need for cartridge interchangeability, demands additional SNR margins and adaptability for robust operation over the lifespan of tape cartridges and drives.

At the output of the read channels, the bit streams of detected bits from the 32 parallel tracks excluding the synchronization patterns are passed to the modulation decoders and derandomizers. At that point, the raw headerized CWI-4s read from tape can be further processed, as indicated in Figure 15, by decoding the CWI-4 header, which provides the ***CWI-4 identifier*** (**CWID**) information necessary to subsequently de-interleave the CWI-4 payloads in memory in order to reconstruct the SDS and ECC PCW matrices discussed in the previous section. The next step, ECC decoding, will be thoroughly examined in the following subsection. However, if error detection and correction through ECC decoding successfully completes for all the data within a dataset, then the

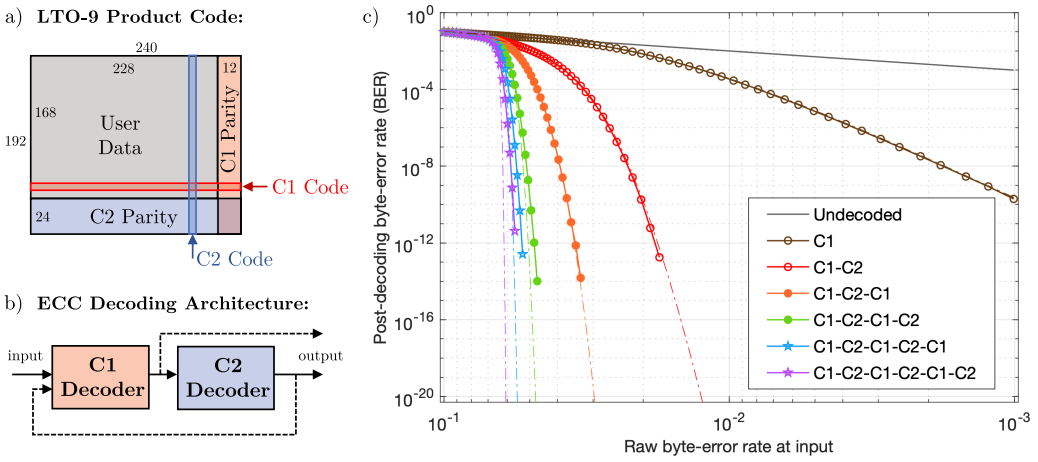


Fig. 16. (a) Structure of LTO-9 product code with RS(240, 224) C1 row code and RS(192, 168) C2 column code. (b) ECC decoding architecture. (c): Iterative ECC decoding performance after a selected number of C1 and C2 decoding iterations. [69].

dataset is considered fully recovered and reconstructed. In case of an ECC decoding failure, i.e., uncorrectable code words, the drive switches to an ***error recovery procedure*** (**ERP**) mode, where additional or alternate ECC decoding steps and/or read retries are attempted.

By combining the user data from subsequent datasets into a stream of symbols, and parsing the symbol stream consisting of user bytes, file marks, and control symbols, the original host records can be derived by means of decryption (if enabled) and decompression. Data integrity is checked by means of the records CRC bytes and the requested host records are transferred back to the host.

6.2 ECC Decoding and Data Reliability

The raw information bits and bytes read from tape, i.e., the data at the output of the read channels or modulation decoders, suffer from both random and permanent errors, typically caused by noise in the readback process and persistent media defects, respectively. ***Error correction coding*** (**ECC**), and specifically ECC decoding, aims to turn these unreliable raw data bytes at the ECC decoder input back into highly reliable user bytes at the output.

From LTO-1 to LTO-9, the storage capacity of LTO cartridges has increased by 180 times from 100 GB to 18 TB. This gain was achieved primarily by increasing the areal density ~70x, but also by increasing the tape length by ~1.7x and the format efficiency by ~1.5x. Increasing the areal density in magnetic recording leads to a loss in SNR, which results in an increase in error rate and hence a reduction in reliability. To avoid such a decrease in reliability, the SNR loss due to areal density scaling needs to be compensated for with advancements in recording technologies such as e.g., improved media, read and write heads, data detection and ECC.

The detailed structure and parameters of the LTO-9 [39] ECC was already introduced in the previous section. As a reminder, Figure 16 highlights the two-dimensional LTO-9 product code with RS(240, 224) C1 row code and RS(192, 168) C2 column code again, and details the ECC decoding architecture and performance. Although the LTO ECC format has always been based on a two-dimensional interleaved dual ECC, it has evolved significantly in terms of performance. For example, LTO-9 has a new C2 format of RS(192, 168) compared to the LTO-8 C2 format of RS(96, 84), delivering superior performance at the same code rate and format overhead.

ECC decoding in the first eight generations of LTO has been based on a single pass through the C1 decoder, followed by a single pass through the C2 decoders. The C1 decoder operates in *"error decoding"* mode, where it detects and corrects byte errors in unknown locations. If the number of byte errors within a C1 code word exceeds the code's error-correction capability, a decoding failure happens. The C2 decoder typically operates in *"error and erasure decoding"* mode: Upon a C1 decoding failure, the C1 decoder provides that failure information to the C2 decoding engine by marking the corresponding byte as erasures, which can be considered as byte errors with known location. The two-dimensional *deep interleaving* discussed in Section 5 spreads the data physically on tape, and therefore randomizes the effects of media defects and burst errors. This randomization results in largely uncorrelated byte errors at the inputs of both the C1 and C2 decoders and enables ECC to work effectively and efficiently.

LTO-9 tape drives, for the first time in the LTO history, introduced iterative ECC decoding in streaming mode. Although the first pass through the C1 and C2 decoder, which still employ hard-decision bounded-distance decoding, is identical to the previous generation's technique discussed above, one or more additional iterations of C1 and/or C2 decoding steps are employed subsequently to further enhance the performance. Figure 16(b) shows the architecture of an ECC decoder with support for iterative decoding.

Figure 16(c) shows the ECC decoding performance for the LTO-9 ECC scheme, evaluated by means of hardware simulations (curves with symbols) and a probabilistic analysis (dash-dotted lines) [69]. Both the simulations and the probabilistic analysis assume uncorrelated byte errors at the input of the ECC decoder. We plot the post-decoding **byte-error rate** (**BER**) performance at the ECC decoder output, after a selected number of m decoding steps, as a function of the raw BER at the input. The "C1-C2" curve illustrates the outcome of a single pass through the C1 and C2 decoders, operating in classic mode. Remarkably, even with just one C1-C2 decoding sequence, the ECC decoder can transform a 1% raw input BER into an exceptionally reliable output BER that is better than 10^{-20}, demonstrating the sheer power of the LTO-9 ECC. By adding an additional C1 decoding step, the curve labeled "C1-C2-C1" shows that with $m = 3$ decoding steps, a $\sim 3\%$ raw input BER is tolerable for a 10^{-20} output BER. With $m = 4$ decoding steps, "C1-C2-C1-C2" decoding can tolerate up to 4.5% raw input BER. Figure 16(c) shows results for up to $m = 6$ decoding steps, indicating that with increasing m, the rate of error correction improvements diminishes, and the probabilistic analysis slightly underestimates the post-decoding BER due to simplifying assumptions [69]. It should be noted that although the plots assume uncorrelated byte errors at the input of the ECC decoder, the LTO-9 ECC scheme is robust against burst errors too. A temporary dysfunctional data reader or large debris on tape can lead to long bursts errors in the C1 code words, but can be corrected by the C2 decoder at the expense of a small increase in post-decoding BER, as discussed in Arslan et al. [7].

LTO tape technology is known to have outstanding data reliability and durability, ensuring stored data remains intact and uncorrupted over time, thus guaranteeing long-term accessibility. This reliability is quantified by the **End-Of-Life** (**EOL**) ***uncorrectable bit error rate*** (**UBER**) of LTO cartridges [7, 115], which represents the likelihood of encountering an uncorrectable data error during readback. LTO has an innate resilience to write mode latent errors, thanks to its read-while-write architecture, where such write-related errors are promptly detected and rewritten during the writing process. The occurrence of uncorrectable error events in LTO tape is exceptionally rare, which makes it very challenging to measure the UBER experimentally. Instead, the LTO's UBER specification is based on the theoretical performance of the ECCs implemented within the tape drive, combined with the assumption that errors are random and uncorrelated. The latest generation, LTO-9, further increased the reliability with an UBER specification of 10^{-20}, which corresponds to the post-decoding BER threshold of 10^{-20} that we used in the LTO-9 ECC

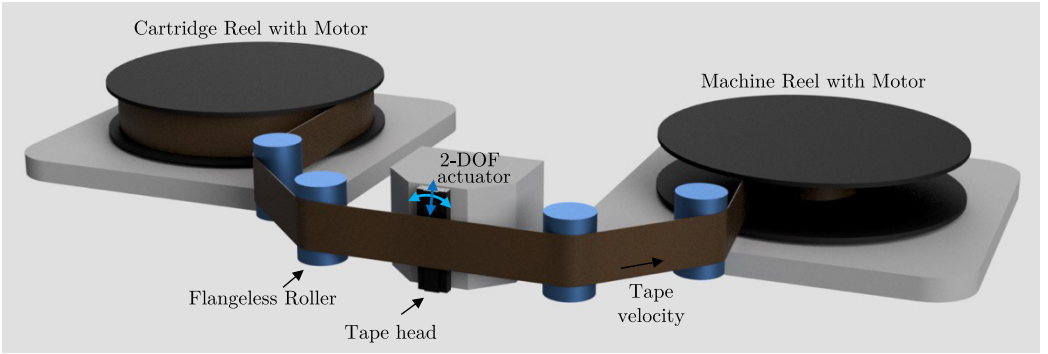


Fig. 17. The tape transport system consisting of two reels which are actuated with brushless DC motors to transport the tape over the read/write head with the help of several guide rollers.

performance evaluation above. An UBER of 10^{-20} corresponds to one unrecoverable read error event for every 12.5 Exabytes of data read. Alternatively stated, on average, only one out of 694,444 LTO-9 tape cartridges will contain an uncorrectable error event due to ECC failure. We also note that the LTO-9 ECC scheme reserves two bytes of parity to detect and correct any one byte error that can occur with very low probability due to a C1 miscorrection. As a result, the probability of an undetected error due to miscorrection is extremely low as discussed in Arslan et al. [7].

7 Mechatronic Aspects of a State-of-the-art Tape Drive

In this section, we describe the mechatronic systems necessary to carry out data write and read operations in state-of-the-art tape drives. We use the LTO-9 drive as an example, however, similar systems are used in recent enterprise drives. The *tape transport system* (Figure 17) is responsible for the movement and longitudinal positioning of the tape within the tape drive. The tape transport system moves tape from the cartridge reel to the machine reel or vice versa, while precisely controlling both the tape speed and the tape tension at the location of the tape head. An accurately controlled tape speed is critically important to write data at a constant linear density. Because of the very thin plastic substrate of modern tape media, tape tension control is crucial for parallel-track recording systems to match the span of the simultaneously written tracks to the width of the reader (and writer) arrays on the tape head.

Although the tape path defines the route that tape follows within the drive, the tape exhibits disturbances such as lateral tape motion and tape skew as it moves from reel to reel, across the head, via tape guide rollers. To enable accurate write track placement including read-while-write verification, and successful data retrieval from tape, the tape head's lateral position and skew relative to the tape have to be continuously estimated and adjusted to follow the tape's motion. Estimation of the head lateral position and skew is accomplished by means of servo patterns pre-formatted on tape, servo readers and servo channels. The head adjustments are accomplished by mounting the tape head on a two degree of freedom *actuator* which is driven by a track-following and a skew-following control systems.

7.1 TBS and Servo Channel

TBS is a powerful servo technology that was developed in the late 1990's specifically for linear tape [10]. All nine generations of LTO to date have used repeating TBS patterns with a basic geometry that is shown in Figure 18(a). A servo frame consists of four bursts of servo stripes designated as A, B, C, and D bursts. Each burst comprises four or five servo stripes having an azimuth angle $\pm\alpha$,

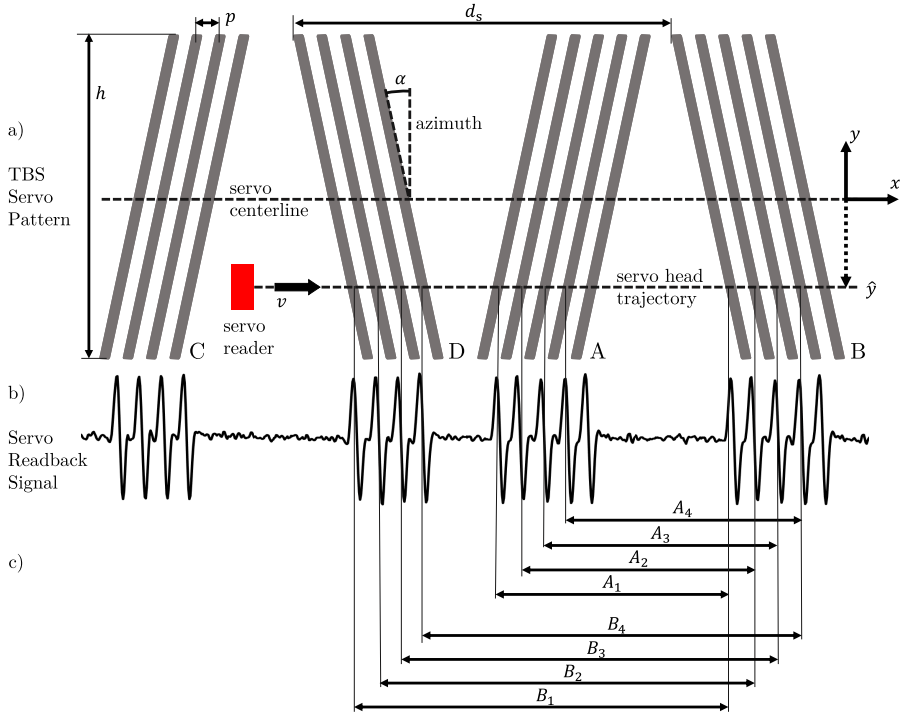


Fig. 18. (a) Geometry of TBS servo pattern and servo reader scanning over the pattern on trajectory $\hat{y}$ below centerline. (b) Readback signal from servo reader. (c) Servo stripe timing differences used to calculate $\hat{y}$.

which allows unambiguous servo frame detection and synchronization. The servo frame geometry is further defined by the servo sub-frame length d_s (or servo frame length $2d_s$), the stripe pitch p and the servo pattern height h. During tape drive operation, a servo reader moves over the repeating TBS patterns on tape and produces a readback signal. The readback signal of an individual servo stripe is called a *dibit* because of its shape [70]. A measured servo readback signal of a TBS frame is shown in Figure 18(b). Note that the timing of the dibits depends on the servo head's trajectory over the servo pattern, which is indicated in Figure 18(a).

The servo readback signal is processed by a servo channel [29, 135] that measures the exact arrival times of the dibit pulses to calculate the time intervals between bursts of stripes with identical and opposite azimuth angle, called B-counts B_i and A-counts A_i, respectively. The servo channel subsequently calculates the estimated lateral position $\hat{y}$ of the servo reader (trajectory) relative to the SB centerline based on the A and B-counts and the TBS pattern geometry according to

$$\hat{y} = \frac{d_s}{2\tan(\alpha)} \left(\frac{1}{2} - \frac{\sum A_i}{\sum B_i} \right). \tag{7}$$

Similarly, the the tape velocity $\hat{v}$ is estimated using the B-counts as

$$\hat{v} = \frac{4 d_s f_s}{\sum B_i}, \tag{8}$$

where f_s represents the sampling rate of the servo signal. The servo channel furthermore extracts additional information, such as the *LPOS information* and an *SB identifier*, which are embedded in

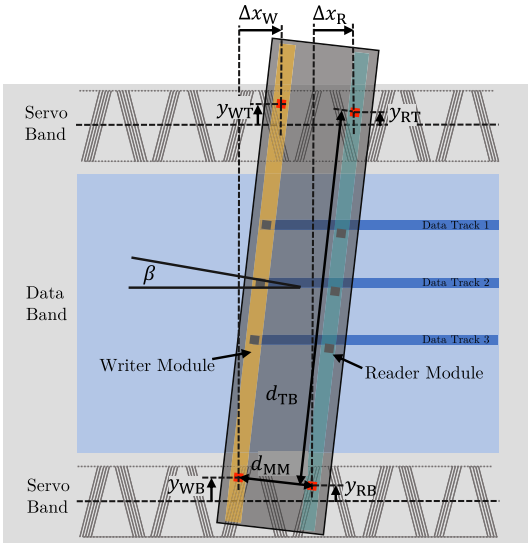


Fig. 19. During a write operation, the head is positioned such that the data writers in the writer module place the data tracks at the desired wrap location. For read-while-write verify, the data readers in the reader module need to follow the data tracks, and thus the head-tape-skew β must be small. The head skew and lateral position are measured with four active servo readers (red) running in the two adjacent SBs.

the TBS pattern by means of *pulse position modulation* on a subset of servo stripes in the A and B-bursts.

As illustrated in Figure 12 (Section 5), LTO media contains five SBs that straddle four DBs spanning the width of tape. At media manufacturing time, repeating TBS patterns are written to all SBs, from the beginning to the end of tape. To write/read data to/from tape, as discussed in Section 5, the tape drive uses a stepper motor to laterally coarse position its head over the desired DB, such that the two pairs of servo readers on both the active writer and reader module of a terzetto head are located on the two adjacent TBS patterns straddling the DB, as depicted in Figure 19. To support lateral position estimates $\hat{y}$ from all four servo readers, indicated as $\hat{y}_{WT}$, $\hat{y}_{WB}$ for the writer module and $\hat{y}_{RT}$, $\hat{y}_{RB}$ for the reader module, the main ASIC in recent generations of LTO drives implements a servo hardware core containing four servo channels running simultaneously.

Furthermore, recent servo cores implement *top-bottom skew estimators* that process the exact arrival times of the dibit pulses from corresponding stripe bursts observed by the servo readers passing over the top and bottom servo patterns. Specifically, the top-bottom skew estimator calculates the distance $\Delta x = \hat{x}_T - \hat{x}_B$ between the LPOSs $\hat{x}_T$ and $\hat{x}_B$ of a module's top and bottom servo reader, respectively, measured in tape travel direction x. The corresponding *head-to-tape skew* estimate, measured as an angle $\hat{\beta}$, can now be expressed as

$$\hat{\beta} = \arcsin\left(\Delta x / d_{TB}\right), \tag{9}$$

where d_{TB} represents the physical distance between the top and bottom servo readers. This skew estimate can be computed independently for both the active writer and the reader module.

7.2 Track-Following and Skew-Following Control System

For a tape drive to write (or read) the data tracks of a specific *wrap*, the tape head's lateral position and skew do not only need to be accurately measured, but they also need to be actively controlled.

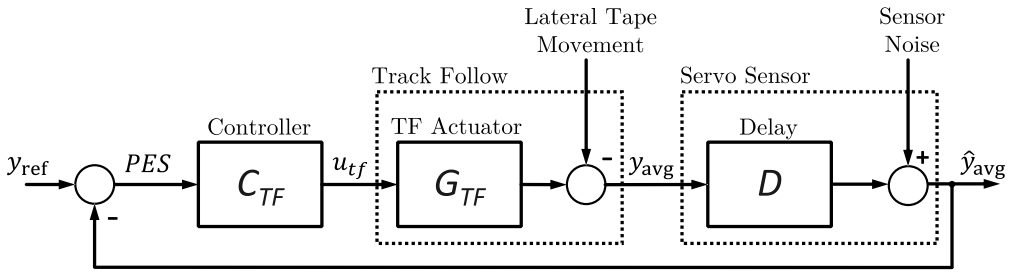


Fig. 20. Block diagram of the track-following control system.

We start by considering a write operation, as shown in Figure 19. Note that for simplicity, only 3 of the 32 data tracks of an LTO-9 drive are shown. The head's lateral position is estimated by the servo channels as $\hat{y}_{WT}$, $\hat{y}_{WB}$ for the writer module and $\hat{y}_{RT}$, $\hat{y}_{RB}$ for the reader module, respectively. The servo reader trajectories are located above the servo pattern centerlines for this example of a write operation in backward direction. For optimal write track placement, the head's lateral position y_{avg} is determined as the average of the top and bottom y-position estimates from the writer module as

$$\hat{y}_{avg} = 0.5(\hat{y}_T + \hat{y}_B). \tag{10}$$

Averaging the top and bottom y-position estimates effectively determines the lateral position in the center of the writer (or reader) module, i.e., in the middle of the write (or read) transducer array. Furthermore, this dual channel averaging improves the position accuracy/resolution by reducing the noise in the position estimate by a factor of $\sqrt{2}$ in standard deviation compared to a single channel estimate [73].

Read-while-write verification is a unique functionality in tape recording, where a set of data readers placed downstream of the data writers immediately read and verify the data after it has been written to tape. A large head-to-tape skew β, as shown in Figure 19, causes the data readers to get "out of the shadow", of the writers, and therefore off track, and consequently would break the read-while-write verification. To avoid such a situation, the measured head-to-tape skew $\hat{\beta}$ also needs to be controlled.

LTO-9 drives use two separate controllers for track following and skew following. Figure 20 shows a block diagram of the track-following control system. The control reference y_{ref} is the desired location of the data tracks (wrap) to be written or read. A **position-error signal** (*PES*) is generated by subtracting the estimated head position $\hat{y}_{avg}$ from the control reference y_{ref} and fed to the controller. The servo controller, in combination with a current driver and a track-following actuator, adjusts the position of the head and thereby closes the track-following servo control loop. Note that instead of the exact head position y_{avg}, only a delayed and noisy estimate $\hat{y}_{avg}$ is available for control [57]. The controller is an H∞ compensator which takes into account the sensor delay D and the transfer function G_{TF} of the track-following actuator. The weighting functions of the controller are shaped to reduce disturbances from the tape path such as roller and reel vibrations [162, 163].

Figure 21 shows a block diagram of the skew-following control system which works in a similar fashion as the track-following control [160], but aims at minimizing the **skew-error signal** (*SES*). The controller aligns the readers and writers perpendicular to the tape by using a skew control reference $\beta_{ref} = 0$.

Tape path disturbances such as reel and roller vibrations, tape pack shifts, and tape path eigenmodes [24, 124] lead to lateral tape motion and tape tilt during tape transport. The skew- and track-following control loops discussed above aim at rejecting these disturbances by minimizing both

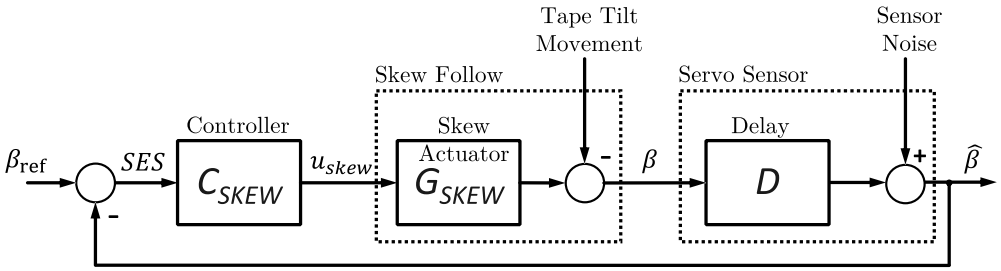


Fig. 21. Block diagram of the skew-following control system.

the skew and lateral position errors. Head actuation for track- and skew-following is implemented using three actuators. The tape head is mounted in a spring guided voice coil actuator responsible for track-following (fine positioning) shown in Figure 10. For coarse positioning of the head to a desired DB and wrap location, a stepper motor controls the movement of the track-following actuator as it traverses along a guide rod. The assembly of tape head, track-following and coarse actuator is attached to a second voice coil actuator that provides rotational capability for skew following. Closing the position and skew feedback loops introduces secondary disturbances from the voice coil actuators or the servo sensor. Secondary disturbances are actuator eigenmodes, sensor noise, sensor delay, and longitudinal compression waves which can propagate through the control loop and must be properly addressed in the design of actuator, servo sensor and controller [57, 74].

7.3 Tape Transport Control System

To access data in a tape cartridge, the cartridge needs to be mounted to a drive, and subsequently loaded by the drive. The load operation includes the process of threading the tape from the cartridge reel over guide rollers and the tape head onto the machine reel located in the drive. At that point, the tape transport system is ready to move tape to carry out seek, read or write operations. Figure 17 depicts the tape's path from one reel to the other via four grooved flangeless guide rollers which ensure a smooth and stable head-tape interface at the tape head. Both reels are driven by brushless DC motors to wind and unwind tape. The reel motors are regulated by the tape transport control system, which aims at moving the tape at a commanded tape speed and tape tension across the read/write head.

As discussed in Section 3, the emergence of tape featuring an ultra-thin polymer substrate, coupled with parallel-track recording at narrow track widths, has posed a growing challenge from the phenomenon known as TDS. TDS refers to changes in tape width in response to changes in temperature, humidity and tension. If data is written under one set of environmental conditions and then read back again later under different conditions, the pitch of the transducers in the read head may not match the pitch of the tracks on tape. To address these TDS effects, LTO-9 drives implement active TDS compensation leveraging tension control. By adjusting the tape tension around a nominal tension value, the width of tape can be changed, as shown in Figure 22. By increasing or decreasing the tape tension, the relative distances between data tracks on tape get smaller or larger, respectively. Figure 22(b) shows a hypothetical example of a write operation, where a (too) large tape tension during write leads to a misalignment between desired and written data tracks. An estimate of the relative *width mismatch* $\hat{w}$ between the *head span* (servo reader pitch) versus the *tape span* (SB pitch) can be derived from the servo position estimates of a writer or reader module as

$$\hat{w} = \hat{y}_{\text{T}} - \hat{y}_{\text{B}}, \tag{11}$$

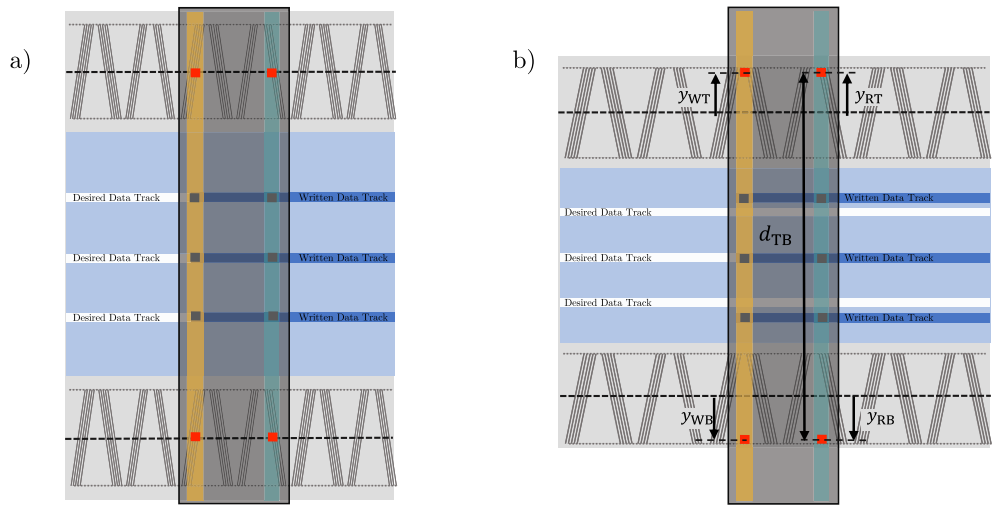


Fig. 22. (a) Tape under tension with the servo readers on their respective band's centerline. (b) Increasing the tape tension stretches it longitudinally and reduces its width due to the Poisson effect. An estimate of the head versus tape span mismatch can be derived from the top and bottom servo readers (11) $\hat{w} = \hat{y}_T - \hat{y}_B$.

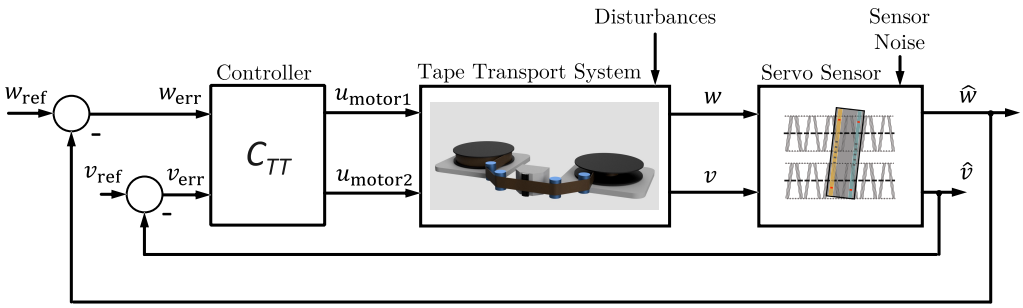


Fig. 23. Block diagram of the tape transport control system with tension-based active TDS compensation.

assuming that the head-to-tape skew β is negligible small.

Figure 23 shows a simplified block diagram of the multiple-input multiple-output tape transport control system [15, 31, 32, 161]. The model based controller is fed with the width mismatch error signal $w_{err} = w_{ref} - \hat{w}$ and tape velocity error $v_{err} = v_{ref} - \hat{v}$ and calculates and applies the motor currents. The controller model takes into account the reel inertias, viscous friction values, and the amount of tape wound up on each reel [30]. Disturbances influencing the width error are tension changes due to roller and reel imbalances, temperature, humidity and tape creep. On the other hand, the tape velocity error is affected by unmodeled friction errors and unmodeled tape transport dynamic effects. The controller performance is also limited by sensor noise and delay.

8 Current Tape Drive Specifications and Tape Specific Considerations

8.1 State-of-the-art Tape Drives

At the time of writing, two families of tape drive were under active development with roadmaps for future generations: IBM TS11xx Enterprise and LTO. The latest generations of both families: the IBM TS1170 and the LTO Gen 9 **full high** (**FH**) and **half high** (**HH**) drives, are shown in



Fig. 24. State-of-the-art tape drives: TS1170 (left), LTO-9 FH (center) and LTO-9 HH (right). Reprint Courtesy of IBM Corporation © (2024).

Table 2. Performance Characteristics of TS1170 [102] and LTO-9 FH [99] and HH [103] Tape Drives

	TS1170	LTO-9 FH	LTO-9 HH
Native capacity (TB)	50	18	18
Compressed capacity (TB)	150	45	45
Native data rate (MB/s)	400	400	300
Compressed data rate (MB/s)	1,000	1,000	750
Host Interface: dual port Fiber Channel (Gb/s)	16	8	8
Host Interface: dual port SAS (Gb/s)	12	12	12
High speed search (m/s)	12.4	9.5	6.4
Form factor	2U	2U	1U

Note: LTO-9 FH peak high speed search speed = 10 m/s, average speed = 9.5 m/s due to reduced speed at the beginning and EOT.

Figure 24. Table 2 presents a summary of their performance characteristics. All three drives have recommended temperature and humidity ranges of 16–25C and 20%–80% relative humidity respectively. LTO-9 drives provide backwards read and write functionality to LTO-8 media. Compared to the FH drive, the HH drive has lower performance in terms of data rate and locate times and is not rated for as many load/unload cycles. Early generations of both families of tape drive are also commercially available and may provide a lower $/TB media price point with the tradeoff of lower data rates and volumetric densities.

State-of-the-art tape drives have a variety of characteristics that are distinct relative to HDD or flash. Knowledge of these characteristics may be useful to users of tape or anyone writing an application to use tape and are discussed in the following sub-sections.

8.2 Append Only

Tape is an append only technology, as discussed in Section 5. As a result, data cannot be deleted or overwritten in place. It is therefore necessary for the tape management software to keep track of what data has been written on which cartridge and to also keep track of what data has been deleted. If a significant fraction of the data stored on a cartridge has been deleted, the remaining data can be migrated to a new cartridge and the initial cartridge can be reused.

8.3 Accessing Data on Tape

LTO tape drives maintain a ***tape directory*** (**TD**) in the CM of each cartridge that contains, among other information, the: (1) record count, (2) file-mark count, (3) write pass number and (4) dataset number at the middle and end of each wrap. When a request to retrieve data from the tape is received, the drive uses the TD to estimate/interpolate the LPOS on tape of the requested data. The drive then high speed locates to an LPOS position several meters before the estimated position, slows down to read velocity and begins to read the DSIT information to determine if it is near the target dataset. If so, the drive continues the read operation until the data has been returned to

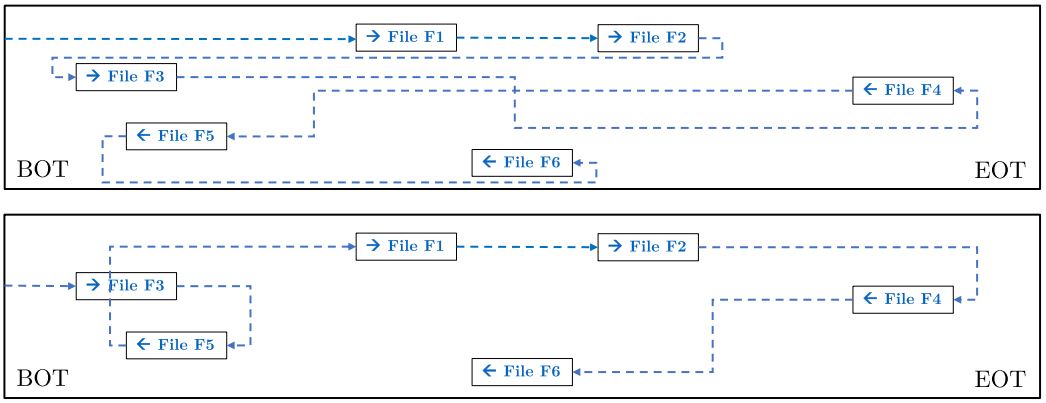


Fig. 25. Illustration of seek trajectory if files accessed in order of block number (top) and RAO access order to minimize seek time (bottom). BOT = beginning of tape, EOT = end of tape.

the host. If the drive significantly over or underestimated the location, the DSIT data is used to re-estimate the target position and a back-hitch or new seek operation will be performed. If the data stored on tape was in-compressible or had a relatively constant compression rate, the initial estimate will be quite accurate and the first seek will likely be successful. However, if there is a lot of variation in the compression rate, the estimate may be inaccurate and the seek operation will take more time. IBM TS11xx drives maintain a ***high-resolution tape directory*** (**HRTD**) with 64 equally spaced entries per wrap. This enables a much more accurate position estimate and hence also enables a better random seek performance. The length of tape used to store user data in an LTO-9 cartridge is about 1,000 m. The average distance travelled for a random seek after load is 1/2 the length of tape and each subsequent random seek is 1/3 the length of tape on average, or about 55 s and 35 s, respectively.

8.4 Recommended Access Ordering (RAO)

Because of the serpentine recording technique used in linear tape drives and built-in compression, there is no clear relation between block number and physical location on tape. Two blocks of data separated by a large interval in block number may be physically very close but on adjacent wraps or in an adjacent DB. Recent generations of TS11xx drives implement a function called RAO which enables faster data retrieval if multiple blocks/files are to be recalled from a cartridge. An application can use the function to request the drive to generate a recommended access order in which blocks should be requested to minimize the seek time. The drive uses the HRTD to estimate the physical start location for each file and then calculates a recall order that minimizes the total seek time as illustrated in Figure 25. If the number of files is large this calculation is very computationally intensive and the drive calculates an approximate solution. The performance gain from RAO depends on the number and size of the files but is typically on the order of a 30–60% reduction in access time [207]. The latest generation of LTO drive, Gen 9, provides a similar functionality called *open RAO,* however the performance gains are lower because of the lower resolution of the TD in LTO.

8.5 Low Tension Unload

The phenomena of TDS and storage creep were discussed in Section 3 of the article. Both the TS1170 and LTO-9 drives use tension control to actively compensate for TDS, as discussed in Section 7. This can result in a significant variation in the tension wound into a tape reel which

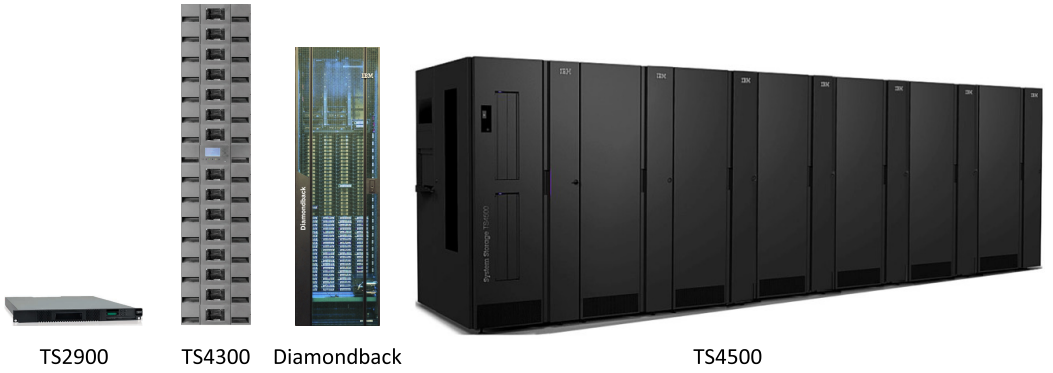


Fig. 26. IBM Tape library portfolio. Reprint Courtesy of IBM Corporation © (2024).

would cause pack distortions over time due to storage creep, if left unaddressed. To minimize storage creep, TS1170 and LTO-9 drives perform a *low-tension unload* operation in which the tape is wound into the cartridge at a low tension immediately before it is unloaded from the drive. To do this efficiently, the drive keeps track of the furthest location down tape accessed during a given mount, and the drive high speed locates to just beyond this point, lowers tension and rewinds the tape into the cartridge reel. In the TS1170 drive this operation is performed at 18 m/s and is significantly faster than in LTO-9 drives that perform the operation at the high speed search speeds listed in Table 2.

9 Automated Tape Libraries

9.1 Introduction and Overview

In the early days of magnetic tape storage, tape reels were stored on shelves in large rooms adjacent to the data center and had to be manually retrieved and loaded onto a tape drive by human operators. Such rooms were referred to as tape libraries and typically housed hundreds to tens of thousands of tape reels. The introduction of tape cartridges and automated robotic tape libraries (see Section 13) reduced the access time to data stored on tape from hours to minutes and significantly improved reliability by removing error prone humans from the loop. Today, an automated tape library typically consists of an enclosure that contains tape cartridges, one or more tape drives for performing read/write operations, an optical mechanism such as a bar code reader to identify cartridges and a robotic and control mechanism used to transport cartridges between storage locations and the tape drives and to mount and unmount cartridges to and from the drives. After the robot mounts a cartridge to a drive, the drive loads the cartridge and performs a seek operation to the desired location on tape to perform a read or write operation. After read/write operations are complete, the drive must rewind and unload the cartridge before the robot can unmount it from the drive. Storage locations are typically referred to as "*slots*" and the robotic mechanism as an "*accessor*" or "*robot arm*". A library also typically includes a mechanism(s) for loading and unloading cartridges to the library referred to as an I/O slot.

Automated tape libraries are available in sizes that range from a few slots and a single drive in a 1U "pizza box", to large modular systems that are expandable to tens of thousands of slots and that support more than one hundred tape drives. Larger systems may also support more than one accessor. We take IBM's current tape library portfolio, shown in Figure 26, as an example to illustrate the range of available systems. The smallest system, the TS2900, has 9 available slots and one HH LTO tape drive in a 1U form factor. The TS4300 is a rack mounted modular system with a starting configuration of 40 slots and up to 3 HH LTO drives or one FH LTO drive and

can be expanded to up to 640 slots and 16 FH or 48 HH LTO drives. The Diamondback library is a rack sized library designed specifically for the unique requirements of hyperscale cloud companies, such as ease and speed of deployment as well as optimal reliability in erasure coded environments and self service. The library can be shipped with pre-installed media and drives and can be installed in a datacenter in less than 30 minutes. The Diamondback library fits in the same floor space as a standard **open compute project** (**OCP**) rack and with LTO-9 technology provides up to 27 PB native capacity in a 0.7 m^2 footprint. At the upper end, the TS4500 is a modular system consisting of 1 to 18 storage frames in a linear architecture that can be configured with up to 23,170 slots and up to 128 FH LTO or TS11xx Enterprise drives. With TS1170 media, the maximum capacity is 877 PB for a density of 51.6 PB/m^2 of floor space. All four library types have an ethernet port for remote library management and have options for fiber channel or SAS for the data path, with the exception of the TS2900 which only supports SAS. A variety of other companies including Dell, Fujitsu, HPE, Quantum, Spectra Logic, and Tandberg also currently offer tape library solutions.

The functionality and features provided by tape libraries vary with the size/type of library and between vendors. To provide some insight into the type of features that are available, we use the IBM TS4500 library as an example. The TS4500 has an integrated management console and provides remote management functionality via a GUI or a **command line interface** (**CLI**) and in addition supports REST API and REST over SCSI commands. Remote monitoring is supported via SNMP, e-mail or syslog. The library has separate and redundant data and control paths with path failover for both and has redundant power supplies. In addition, it supports up to two active accessors and provides automated/library managed media health checking. Mixed media and drive types are supported in the same library and a single physical library can be partitioned into multiple logical libraries each with distinct cartridge and tape drive resources. Partitioning can be used for file management purposes or to give different users or applications access to dedicated, independent resources [8]. Recent generations of tape drives provide built-in hardware-based encryption. There are typically three methods for managing encryption: system managed, library managed and application managed encryption. In all cases a key is required to encrypt and decrypt the data. Details of how this is implemented and managed vary depending on the application and/or library and vendor.

A key performance metric for tape libraries is the time to mount a cartridge, i.e., to fetch it from a slot and insert it into a drive. Performance is often specified as the number of mount cycles that can be performed per hour, where a mount cycle is defined as the steps of removing (unmounting) a cartridge from a drive and replacing it to a storage slot and then fetching a new cartridge and mounting it to the drive. Mount rates can vary considerably and typically depend on the number and speed of the accessors (robot arms) and the architecture of the library. For example, in expandable modular libraries such as the TS4500, the average time for a mount cycle increases with the number of frames in the configuration because on average the accessor needs to travel longer distances between the drives and slots. The design of the storage slots also plays a role. For example, in the TS4500 and Diamondback libraries cartridges are stored in *deep slots* where multiple cartridges are stored one behind the other in "tiers". The accessors are equipped with dual grippers that can manipulate two cartridges such that cartridges in the front two tiers can be accessed quickly by the dual gripper, whereas access to cartridges in deeper tiers requires the cartridges in front to be first moved to other slots, resulting in longer access times. To improve the mount performance of such systems, algorithms have been developed to ensure that frequently accessed cartridges are stored in the front tiers and infrequently accessed cartridges in the back tiers. In Spectra Logic libraries, cartridges are stored in drawers which must first be opened by the accessor to retrieve a cartridge which adds an extra step to the mount process, but can be efficient if multiple cartridges are accessed sequentially from the same drawer.

9.2 Tape Library Performance Modeling

To properly dimension tape library systems, it is imperative to develop models that evaluate the effect of the various parameters on their performance. The results obtained can subsequently be used to provision tape systems to achieve service level agreements for given workloads and data access times. Here we provide a survey of the relevant literature on performance evaluation of tape library systems.

Over the last six decades there have only been a few publications that evaluated the performance of tape library systems. This is attributed to the fact that a theoretical evaluation of their performance is an extremely challenging task and consequently a performance assessment was mainly conducted by simulations [55, 78, 138] and measurements [123]. However, simulations and measurements are time-consuming compared with analytical models that not only provide fast execution times, but also insight into the system dynamics and the effect of the various parameters.

The first analytical attempt resulted in an intractable queueing model, which led the authors to conduct their study by means of simulation [138]. Twenty years later, approximate analyses were presented using $M/M/c$ [77] and $M^X/G/c$ [122] queueing models. Another twenty years later, an improved model that captures the polling nature of operation of tape library systems was presented by Iliadis et al. [113]. This model obtained an approximate closed-form expression of the mean waiting time for independent and identically distributed (i.i.d) seek times and for an abundance of robot arms such that one is always available to mount/unmount a cartridge. System operation and performance was assessed in the light-load and the heavy-load regions using an $M/G/K$ and an $M/G/1$ polling queueing model, respectively, and subsequently in the medium-load region through a tangent interpolation.

An enhanced theoretical analytical model that considered the variability of seek times and captured the principal aspects of tape library operation including the contention for robot arms and its effect on system performance was presented by Iliadis et al. [111, 112]. Using elaborate, non-standard queueing theory models, it evaluated performance measures of interest, such as the mean waiting time and the robot mount rate, as a function of the system parameters, including the number of tape cartridges, the number of tape drives, the number of robot arms, and the tape load, seek, rewind, and unload times. It obtained results that enabled a better understanding of the design tradeoffs and provided useful insights into the behavior of the tape libraries. Also, a theoretical analysis presented by Iliadis et al. [112] yielded a condition for determining whether the robotic mechanism is a bottleneck.

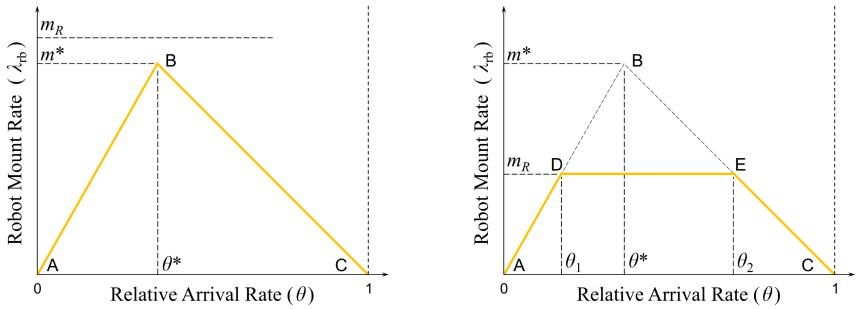
We proceed to briefly review the model presented by Iliadis et al. [112] along with the main parameters and some of the performance results obtained. The tape library system contains c cartridges, d tape drives, and a robot arms with a robot transfer time R. Requests submitted to tapes are queued and subsequently served according to a **first-come-first-served** (**FCFS**) policy. The waiting time of a request is the time interval from its arrival to the tape system until the corresponding data transfer is initiated. To ensure fairness, tapes are mounted according to a cyclic (round-robin) policy. Requests to cartridges are served according to an exhaustive service discipline such that a tape is only unmounted when all its requests have been served. Subsequently, another tape with pending requests is mounted. If, however, there are no pending requests to any other non-mounted cartridge, two policies are considered depending on whether the tape remains mounted.

(a) ***Not-Unmount (NU) policy***: a tape cartridge remains mounted upon completion of all its pending requests, in anticipation of the next request arriving to the same currently mounted, but idle cartridge.

(b) ***Always-Unmount (AU) policy***: a tape cartridge is immediately unmounted upon completion of all its pending requests, in anticipation of the next request arriving to another non-mounted cartridge.

Table 3. Notation of Main System Parameters

Parameter	Definition
c	number of tape cartridges
d	number of tape drives
a	number of arms
R	robot transfer time
λ_{ct}	arrival rate of requests for a cartridge
λ	arrival rate of requests to the tape system ($\lambda = c\,\lambda_{\text{ct}}$)
$\lambda_{\max}$	maximum throughput (Equation (14) of [112])
θ	relative arrival rate ($\theta \triangleq \lambda/\lambda_{\max}$)
ρ	system load (Equation (13) of [112])



(a) Robot is not a bottleneck. (b) Robot is a bottleneck.

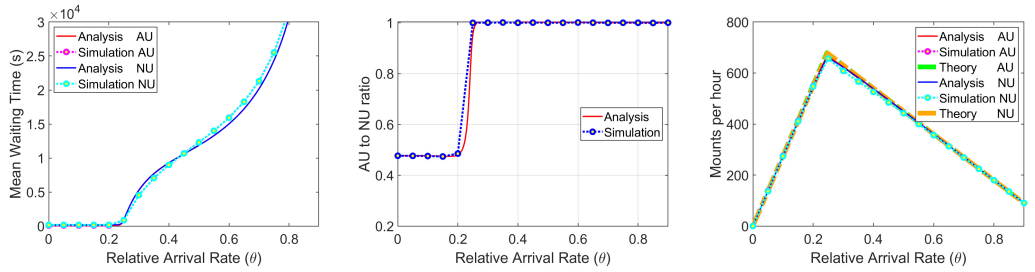
Fig. 27. Theoretical evaluation of the robot mount rate λ_{rb} vs. load.

The workload was assumed to be symmetric with requests arriving to each of the c cartridges according to independent and identical Poisson processes at a rate of λ_{ct} such that requests arrive to the tape system according to a Poisson process with a rate λ, where $\lambda = c\,\lambda_{\text{ct}}$. The request sizes were assumed to be i.i.d. The notation is summarized in Table 3. The parameters are divided according to whether they are independent or derived and are listed in the upper and the lower part of the table, respectively.

Performance was evaluated using a combination of elaborate, non-standard queueing theory models, whereby requests submitted to tapes correspond to jobs arriving to queues and tape drives correspond to servers. System operation and performance was assessed in the light-load and heavy-load regions using an $M/G/K$-based and an $M/G/1$ polling-based queueing model, respectively, based on an iterative procedure. Subsequently, the mean waiting time $\overline{W}$ in the medium-load region was derived through interpolation. Several relevant performance measures including the robot mount rate λ_{rb} were also obtained as a function of the relative arrival rate θ defined as follows: $\theta \triangleq \lambda/\lambda_{\max}$, such that $0 \leq \theta \leq 1$, which reflects the system load.

Next, we review the theoretically obtained closed-form approximation for the robot mount rate λ_{rb} as a function of the load θ [112, Equation (58)]. The maximum robotic mount rate m_R for a number a of robot arms is determined by $m_R = a/(2\,R)$, where R denotes the robot transfer time.

In the light-load region, the robot mount rate λ_{rb} increases linearly with θ, as depicted by segment AB in Figure 27(a). By contrast, in the heavy-load region, λ_{rb} decreases linearly, as depicted by segment BC in Figure 27(a). The maximum robot mount rate m^* is achieved when $\theta = \theta^*$, that is, $m^* = \lambda_{\text{rb}}(\theta^*)$, with θ^* and m^* determined by Equations (56) and (57) of Iliadis et al. [112], respectively.



(a) Mean waiting time, $\overline{W}$ (b) $\overline{W_{\mathrm{AU}}}$ to $\overline{W_{\mathrm{NU}}}$ ratio (c) Robot mount rates, λ_{rb}

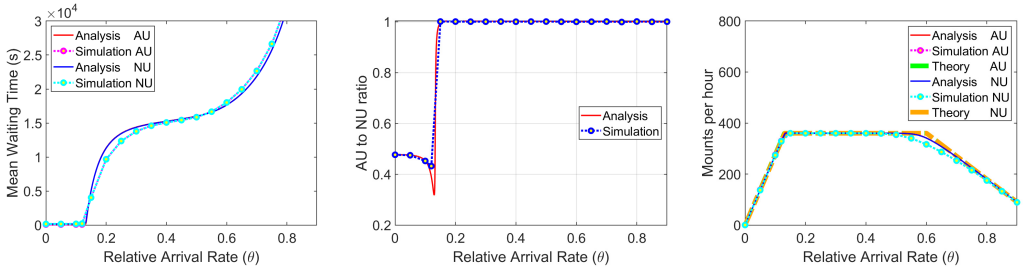
Fig. 28. Performance measures vs. relative arrival rate for a configuration with $c = 3200$, $d = 32$, and $a = 2$.

Note that Figure 27(a) is valid only when the robot mechanism is not a bottleneck, that is, when the robot mount rates do not exceed the maximum possible rate of m_R, or equivalently, when the number of robot arms a exceeds a value, which is determined by Equation (59) of Iliadis et al. [112]. When the number of robot arms a is less than this value, the robot mechanism is a bottleneck and the robot mount rates are truncated, as shown in Figure 27(b) by the segment DE. In this case, it holds that $\lambda_{\mathrm{rb}}(\theta_1) = \lambda_{\mathrm{rb}}(\theta_2) = m_R$.

9.2.1 Numerical Results. The performance of several tape library configurations was assessed in Iliadis et al. [112] by using both theoretical predictions and event-driven simulations. It was verified that the theoretical curves capture the system behavior as they match well with the simulation results.

Next, we present some of the performance results concerning a 4-frame configuration of an IBM TS4500 tape library system that comprises $c = 3200$ cartridges and $d = 32$ drives. The corresponding parameter values are listed in Table 2 of Iliadis et al. [112] and correspond to the LTO-8 tape technology [38]. The workload in each tape is assumed to be symmetric with random requests whose mean is 843 MB, second moment is 8.9 GB^2 and coefficient of variation is 3.39. Performance results, as a function of the load θ, obtained when the library operates using two robot arms ($a = 2$) are shown in Figure 28. The analytical mean waiting times are shown in Figure 28(a) with the simulation results indicated by the circles. The red curve for the AU policy is barely visible because it lies just below the blue one for the NU policy. At light loads the AU policy results in lower waiting times and as the load increases, its performance approaches that of the NU policy. This is illustrated in Figure 28(b) that shows the ratio of the mean waiting time $\overline{W_{\mathrm{AU}}}$ for the AU policy to the mean waiting time $\overline{W_{\mathrm{NU}}}$ for the NU policy. Depending on the system parameters, the AU policy may not always result in lower waiting times at light loads. For instance, when the mount times are much longer than the unmount times, it is preferable to employ the NU policy such that a tape remains mounted in anticipation of future requests arriving to it. However, for large values of the ratio of the number of tape cartridges to the number of tape drives, and according to Remark 1 of Iliadis et al. [112], at light loads the AU policy results in lower waiting times.

The robot activity is shown in Figure 28(c). The analytical and simulation results confirm the validity of the theoretical ones. As the load increases, the robot activity initially increases, but after some point it starts decreasing. Clearly, at high loads there are many requests to be served such that tapes remain mounted for long periods of time. Therefore, the time intervals between unmount operations are long, which implies a reduced robot activity. Note also that the mean waiting time as shown in Figure 28(a) starts increasing drastically at the point at which the robot activity reaches its peak, as shown in Figure 28(b), which occurs at $\theta = \theta^*$.



(a) Mean waiting time, $\overline{W}$ (b) $\overline{W_{\text{AU}}}$ to $\overline{W_{\text{NU}}}$ ratio (c) Robot mount rates, λ_{rb}

Fig. 29. Performance measures vs. relative arrival rate for a configuration with $c = 3200$, $d = 32$, and $a = 1$.

Performance results obtained when the library operates using one robot arm ($a = 1$) are shown in Figure 29. The analytically predicted mean waiting times are shown in Figure 29(a) that demonstrates a dramatic increase of the mean waiting time at $\theta = \theta_1$. This is due to the fact that at this load, the robot arm becomes a bottleneck as it is constantly busy performing mount and unmount operations. This is reflected by the flat part of the robot operation activity that reaches its peak in the range $\theta_1 < \theta < \theta_2$, as shown in Figure 29(c). However, at high loads, that is, when $\theta > \theta_2$ the robot is no longer a bottleneck as its activity is reduced. Also, in Figure 29(b) it is shown that, as the load increases, the ratio of the mean waiting times of the AU to those of the NU policy initially decreases.

The results obtained in Figures 28 and 29 reveal that when the robot mechanism becomes a bottleneck, system performance degrades, but it can be improved by increasing the number of robot arms. Another way to improve performance is to mitigate the burden on the robotic mechanism. This can be achieved by employing the scheme presented in Iliadis et al. [112], which accumulates multiple requests before sending them to the tape library. This causes the robot activity to reach its peak at high loads, which in turn implies that for realistic loads, the mean waiting time is kept low.

The model also captures the effect of the unmount policy deployed. Although at medium and high loads the mean waiting times corresponding to the *AU* and *NU* policies are the same, at light loads and for practical parameter values, the *AU* policy yields mean waiting times that are lower than that of the *NU* policy.

The usefulness of the analysis over simulation is that it provides answers to issues in a way that simulation cannot accomplish. For instance, only analysis enables the determination, as a function of the system parameters, of the number of robot arms that result in the robotic mechanism to become a bottleneck. This function is given in Equation (59) of Iliadis et al. [112].

10 Tape Software Considerations

10.1 OS-level and Application-level Software

Support for tape usage at the operating system and application programming level is very dependent on the operating system being used. As potentially useful examples of tape system support we summarize tape interfaces from a selection of IBM mainframe (IBM Z), Unix, and PC-class operating systems.

10.1.1 IBM z/OS. IBM z/OS [107], the flagship operating system for IBM mainframes, has tightly integrated support for data on tape. As also mentioned in the History section, Section 13, OS/360 (the ancestor of z/OS) was designed to allow programs to be device-independent (to the extent possible based on access requirements) [93]. Programs can be written to access sequential

files, and the choice of tape or disk (or even card) input/output can be left until execution time. In support of this independence, z/OS writes standard volume and file labels on tape volumes, and the system catalog can identify the location of requested datasets (files).

Application programs in z/OS use open and close references to symbolic names; these are in turn linked to **Job Control Language** (**JCL**) which specifies dataset names and optional device information [96]. Reading and writing is managed by Access Methods (somewhat similar to file system drivers), and the application programmer is relieved of the necessity of dealing with device details.

Application programs that need to bypass standard tape processing in z/OS must use very low-level access to the tape device, writing their own channel programs and using the `execute channel program` [93] interface to the I/O subsystem. Depending on the application, they may also need to use an access-restricted JCL option to bypass tape label processing.

10.1.2 IBM z/VM and CMS. IBM z/VM [108] is a virtualization system that allows for the creation of virtual machines, each of which can run an operating system such as z/OS or Linux. The time-sharing component of z/VM is the **Conversational Monitor System** (**CMS**) [104], which also runs in a z/VM virtual machine.

CMS users can attach physical tape drives to their virtual machine, and use utility commands such as the `TAPE` command to manipulate the position and contents of a tape [97]. While CMS commands can read and write standard z/OS volume and file labels, tape label processing is not required in CMS.

CMS application programs written in assembler language can use macros such as `TAPECTL`, `RDTAPE`, and `WRTAPE` to process tape data [98]. The programs have complete freedom as to what to read or write, and are also responsible for low-level control such as tape positioning and tape mark processing.

10.1.3 Unix and Linux. Unix and Linux systems support the attachment of tape drives, and include the `mt` command to manage magnetic tapes [65]. As discussed in the History section, Section 13, Unix/Linux I/O is byte-stream oriented, and this combined with the "device file" concept gives the systems a large degree of device-independence. However, support for tape is intrinsically very low-level, leaving it up to the user to manage the details of tape mounting and read/write location.

For application programmers Linux contains a generic tape device driver, the SCSI Tape Driver. However, IBM has supplied an open-source Linux Tape and Medium Changer device driver called `lin_tape` [101] which makes programming of tape applications much simpler. IBM has also supplied a similar (closed-source) driver for both the AIX and Solaris variants of Unix [100].

The **Linear Tape File System** (**LTFS**) [164] can be installed on Linux systems to allow the system to use a linear tape as a file system device. Once LTFS is installed, files can be listed, read, written, and so on, to and from a tape using file system commands and interfaces.

10.1.4 Microsoft Windows. Modern Windows operating systems provide a generic tape class driver that handles operating system-specific and device-independent tape tasks. The tape class driver is provided as a kernel-mode DLL, and must be extended with a tape miniclass driver to support a specific device [147]. As with the Linux/Unix systems, IBM provides a Tape and Medium Changer device driver for Windows [100].

LTFS is also available for Windows through Windows 9; as with all versions of LTFS, this allows files on tape drives to be managed by standard file system interfaces.

10.1.5 macOS. macOS runs a BSD-derived Unix kernel, and might be expected to have the same low-level support for attaching and controlling tape devices as any other Unix system. However,

since 2009 macOS has not included the `mt` command or the necessary infrastructure to use it [185]. Tape devices can be attached to a macOS system, and appear as generic SCSI devices. But there are no built-in facilities to manage tapes, and programs that wish to support tape must do so at the SCSI **Command Data Block (CDB)** [4] level.

However, LTFS is available for macOS, and installation of LTFS enables the use of linear tape files through standard file system interfaces on a Mac.

10.2 Tape-based Application Considerations

Applications that support tape should consider the unique abilities and concerns of tape devices and media. Any concerns should not deter an application from using tape where appropriate, but rather should inform the implementer's approach to using tape.

10.2.1 Compression and Encryption. Modern tape units are capable of compression and/or encryption of data in the drive. When both compression and encryption are used, the two operations are pipelined, with compression taking place before encryption. (This is important because if data were encrypted first, it would essentially be uncompressible.)

The two capabilities provide different benefits. Encryption allows data to be protected against unauthorized access, while compression can significantly increase the capacity of a tape volume. The latest generation of LTO drive specifies an average compression ratio of 2.5 to 1 [39], while some enterprise drives can provide even higher compression rates.

Both of these features can be enabled in tape software, with no additional programming effort required to take advantage of them.

10.2.2 ***Medium Auxiliary Memory (MAM).*** MAM [86], also referred to as CM, is a small memory area embedded in the cartridge of a linear tape that can be read or written without accessing the actual tape data. MAM can be used by a tape device driver or a tape application to save and retrieve information about the status of a tape cartridge. For example, MAM could be used to store tape creator information, or an indicator of whether a tape was properly closed and unmounted the last time it was written to.

For a detailed example of MAM use by a file system see Section 10 of the LTFS Format Specification version 2.5.1 [183].

10.2.3 Partitioning. It is possible to divide modern linear tapes into two or more partitions which can be managed independently. That is, each partition can be treated as a separate logical tape, and read from, overwritten, and appended to independently of the other partitions. As an example, the LTFS partitions a tape volume into a relatively small index partition and a very large data partition. This allows the tape index to be re-written at the "front" of the tape, where it can be read quickly at mount time, without affecting data written into the data partition.

There are certain things to be kept in mind when considering partitioning a tape volume. One is that while partitioning may reduce seek time within a partition (because related data could possibly be grouped closer together), the time to move between partitions can still be significant.

Another thing to be aware of is "guard wraps". When a tape is arbitrarily partitioned, it is necessary to allocate several wraps of the tape between partitions as "dead" space to insure that writing data in one partition cannot affect data in an adjacent one. For a small number of partitions the amount of space lost (perhaps on the order of 1 percent) is usually of no concern, but too many partitions can begin to affect the usable capacity of the tape.

In the latest generation of LTO and IBM Enterprise tape drives, a new feature called band based partitioning has been implemented that provides a specific exception to the need for guard wraps. In both LTO and IBM enterprise tapes, wraps are written into one of four DBs. These bands are not

visible to applications, but are an internal means of organizing the media. With this new feature, if an LTO-9 or TS1170 tape is partitioned into exactly two or four equal partitions, the partitions can be mapped perfectly to the DBs of the tape, and no guard wraps will be necessary. In this case there will be no capacity loss due to partitioning.

10.2.4 Seek Time. The most obvious issue to be aware of in tape applications is seek time. While modern linear tape is block-addressable, the time to move from one block to another on tape can range from seconds to a couple of minutes. Tape applications should be written to minimize the amount of seeking necessary; most tape applications are expected to read and write data sequentially, in which case seek time is not an issue.

10.2.5 Media Wear. Tape media is generally quite robust, allowing many reuses of the tape and extremely long shelf life (when stored properly) [141]. The tape system's head-tape interface is designed so that the head and media are in intimate contact, and linear tape media is designed to support that intimate contact for the specified life of the media. This contrasts with hard disk drives, where the head is "flying" above the media, and head-media contact can be disastrous.

Data is written to tape in a serpentine fashion (as described in Section 5) and thus completely reading or writing a tape volume requires many passes of the media over the head. Different tape formats or generations require different numbers of passes to completely read or write a volume. For example, LTO-9, the latest generation of LTO tape, has 280 wraps, and therefore takes 280 passes of the tape over the head to read or write a full volume.

Additionally, if the access of data on the tape is random and the access order is not optimized (see the RAO discussion Section 8), the amount of tape that must pass over the head for each random read can be significant. This results in some variability in the specifications and claims for media life. Some vendors publish recommendations for full cartridge passes for each generation of media. Typical recommendations are on the order of 1,000,000 passes over any area of the media and 200–300 full cartridge passes [189]. Sometimes vendors have no specific media life specifications, but provide the ability to query the drive for a percentage of life criteria. (As mentioned above, each cartridge has a memory area where information like the number of uses of the tape can be recorded by the drive.)

In all but the most extreme cases tape media wear is unlikely to be an issue. However, when designing a tape-based application or work flow it may be helpful to keep the characteristics of tape usage and wear in mind. Best practices can be generally simplified to

— tape life is optimal with sequential read or write.
— when performing random reads, use a facility like RAO to optimize the read order.

11 Current Tape Usage and Software

The suitability of a storage technology for a specific use case or application is strongly influenced by its performance characteristics and cost relative to other storage options. In this context, tape's most relevant characteristics are [26, 112]:

— good I/O-performance.
— high latency.
— low total cost of ownership (for large volumes of data).
— high volumetric density.
— low power and low CO_2 footprint.

These characteristics make tape storage particularly well suited for storing data that is rarely accessed and that must be retained for a long time. Two examples that fit into this category are backup and archiving. In fact, tape has been used for these purposes since the era of mainframe

computing [11]. Backup and archiving use cases are similar and indeed an archive can be used as a backup and backup applications can be used to archive data. However, backup typically prioritizes ensuring the most current data is backed up and minimizing recovery time while archive emphasizes long-term storage management and preservation of historical data. In recent years, cloud based backup and archiving have emerged as alternatives to tape, however tape still plays an important role [184]. Cloud storage has some advantages in terms of flexibility, however it has higher costs than tape [116]. Hyperscale cloud companies are also increasingly using tape [169] and cloud-based archive solutions may actually use tape as a back-end.

Tape has a variety of other common use cases and often specific software solutions have been developed for them. The INSIC 2019 Tape Roadmap [116] provides a detailed overview of areas where tape is commonly used. In this section, we briefly review these use cases and provide examples of software solutions currently available for them.

11.1 Use Cases for Tape

11.1.1 Big Data/Analytics. Big Data usually denotes datasets that contain very large quantities of data of a high complexity, and/or for which data is generated at a high rate [75]. The amount of data classified as Big Data is increasing year by year. Cloud based solutions are one option for storing and analyzing Big Data. However, if all of the data does not need to be immediately available, it may make sense to outsource storage to tape to reduce costs [116]. Big Data use cases often have varying access patterns in terms of access frequency, which can make a tiered storage solution a good option. For example, if an analysis is performed using machine learning methods, it is advisable for training data that is rarely used to be moved to tape storage [166]. **Hierarchical Storage Management** (**HSM**) and Active Archive software solutions are possible options for this use case. Examples of HSM software solutions are discussed in Section 11.1.5.

11.1.2 Archive. Archiving is a traditional use case for tape storage. An archive typically denotes a copy of data that is retained for long-term preservation and is not intended to be modified [54]. Unlike backup that usually denotes a duplicate copy of data, an archive is often the only copy of an object or dataset. When data is moved to an archive it is often grouped based on its use, e.g., all of the files related to a completed project. This archive or object can then be named and enhanced with metadata to enable search. To retrieve data from the archive back to primary storage, the corresponding object can be identified by its name or found by searching the meta data. An example of a large tape archive is the *Media archive at the Library of Congress* [129]. Many applications that can be used to perform backup operations can also be used for archiving. Available software options are discussed in Section 11.1.3.

11.1.3 Backup and Recovery. Backup and recovery is another traditional use case for tape. Backups are used to protect against data loss resulting from failed or damaged hardware, accidental or malicious deletion, or cyber threats such as ransomware [54, 170]. As computers have become more critical to businesses, governments, and scientific institutions, the need to protect investments in compute has also grown in importance. The data protection/backup use case has evolved to meet this need. The need for data protection/backup exists in all market segments and generally a priority is placed on the speed of recovery of operations. Strategies are employed to provide various points in time from which data can be recovered and means of minimizing time to recover, time to backup, and the required storage capacity. Examples of strategies to achieve these ends include incremental backups and deduplication. In corporate environments, tape storage is often used for backup and recovery [184]. To improve reliability, data backup should be done regularly and automatically. Automatic tape libraries (see Section 9) were a key technology for enabling automated backups.

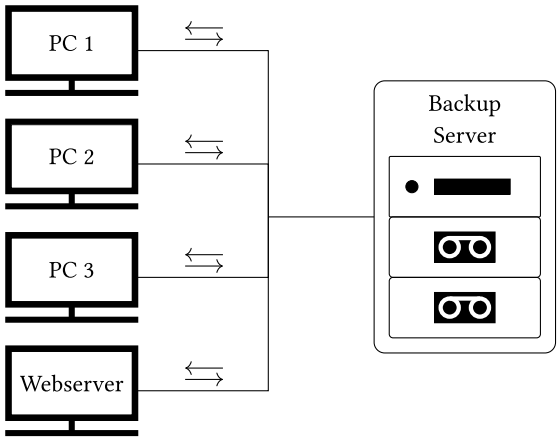


Fig. 30. Illustration of a network based backup solution to tape.

As networks have become more widespread, client-server architectures have been developed to enable access to target devices for backup and archiving over the network. With this approach, tape devices do not need to be physically connected to the backup target, as illustrated in Figure 30.

Gartner describes the current leading backup and recovery software solutions in its *Magic Quadrant* report [87]. Examples of current software solutions that support tape and that can also be used for archiving include: *Arcserve Backup*, *Arcserve UDP*, *Cohesity DataProtect*, *Commvault Platform Release*, *Dell Data Protection Suite*, *HYCU Protégé*, *IBM Storage Protect*, *Rubrik Security Cloud*, *Veeam Data Platform*, and *Veritas NetBackup*.

In the consumer sector, back up and archiving is now typically performed using file hosting services or USB sticks/drives [196] and tape storage is no longer common.

11.1.4 Cold Storage. Cold storage is computer storage for cold data, i.e., data that is rarely used. In some cases, the data may never be used but must be retained for legal compliance reasons or is retained in case of a future need. According to IDC, 60% of corporate data is cold data, meaning that this data will not be accessed for more than 30 days [80]. If higher access latency can be tolerated, tape storage can provide much lower storage costs for cold data compared to other storage technologies. A simple option to reduce costs and free up space in primary storage would be to move cold data to a tape-based archive. However, in this case the data needs to be manually selected and moved by a user. HSM solutions which are described in the following section, provide an automated means for migrating data to a tape tier using policies based on properties such as when the data was last accessed.

11.1.5 HSM/Tiering. The differences in performance, latency and cost of different storage technologies have given rise to the concept of a storage hierarchy that dates back to the era of mainframe computing [128]. Even in the early days of computer development, it was of interest to store data where it cost the least but where sufficient availability was provided. This cost-performance tradeoff is often illustrated as a pyramid as shown in Figure 31.

The idea of HSM is to automatically and seamlessly migrate data to an appropriate storage tier in a tiered storage system. One way to implement this is to initially ingest data to the highest storage tier. Then, if the data is not accessed within a predefined time period, it is automatically migrated to a lower tier.

If data is resident in the SSD or HDD tier when it is accessed, it can be immediately served. If, however the data has been migrated to the tape tier, access latency will be significantly longer. As

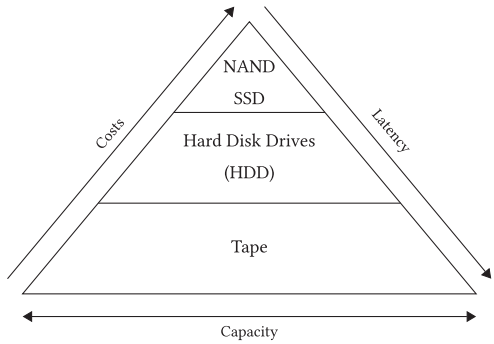


Fig. 31. Illustration of the current storage hierarchy.

a result, the access method for data in the tape tier is different compared to SSD or HDD. There are two implementation options, either data is read directly from tape or it is first moved to a higher tier where an application can access it. In either case, I/O calls are blocked until the data is available. According to Guilleaume [81], systems whose operation corresponds to the first option are called Active Archives, whereas systems that use the second approach are called HSM systems. However, the terms are used synonymously for the IBM product IBM Storage Archive Enterprise Edition [105].

Examples of HSM products include: *IBM Storage Archive Enterprise Edition*, *IBM Storage Protect for Space Management*, *Versity ScoutAM*, *High Performance Storage System*, and *Lustre HSM*.

11.1.6 Bulk Transfer. General purpose file systems and hierarchical file systems date back to the era of mainframe computing [45]. However, unlike other storage technologies, there was no commonly used file system available for tape storage until the advent of LTFS in 2010 [164]. Previously, archiving on a file system-based abstraction of tape technology was only possible within tiered storage with HSM. LTFS is standardized by the **Storage Networking Industry Association** (**SNIA**) [183] and as ISO standard ISO/IEC 20919:2021 [118]. In addition to other capabilities, LTFS enables the bulk transfer of data [182]. There are a variety of implementations of LTFS including: *IBM LTFS - Single Drive Edition*, *IBM LTFS - Library Edition*, *Oracle's StorageTek LTFS - Open Edition*, *Oracle's StorageTek LTFS - Library Edition*, *HP LTFS*, *Quantum LTFS*

11.1.7 Cloud. According to Furthur Market Research, the 15 largest hyperscale cloud companies consumed 65% of the petabytes of enterprise storage capacity shipped in 2023 [149]. Several of these companies have also embraced tape storage and have become the largest consumers of the technology. Cloud storage is most commonly provided as object storage and is typically divided into different tiers that vary in cost, latency, and access rates. A user can select a storage tier to upload to, but usually does not know the storage technology behind it. However, if a storage tier is very low cost and has a high latency, there is a high probability that it uses tape. Many cloud providers offer an object interface that is compatible with the Amazon S3 protocol [3]. The basic interfaces provided by Amazon are an API that is available for various programming languages and a CLI. In addition, other CLIs, applications and file systems have been developed by third parties.

11.1.8 Disaster Recovery/Air Gap Protection. Disaster recovery plans are critical for an enterprise to be able to ensure business continuity, and data backup plays an important role in this area [153]. Data loss can occur due to a variety of reasons, including problems such as software corruption, disasters, ransomware attacks, hardware failures, theft, or human error [59]. A well known recommendation for disaster recovery is the so called the 3-2-1 rule [153]:

3: Keep **three** copies: one primary backup and two copies of your data.
2: Save your backups to **two** different types of media.
1: Keep at least **one** backup copy offsite.

Tape storage is particularly suitable for the last aspect because it is offsite and offline and can also provide half the solution for the second aspect. If a cartridge is on a shelf in storage, it is completely disconnected from the network (air gaping). This prevents cyber attacks such as ransomware [80]. Another strategy to enhance tape's air gap is to export cartridges to a virtual library partition without any drives assigned to it [25]. Examples of current backup software are listed in Section 11.1.3.

11.2 S3 Interface

The previous sections discussed the benefits of tape storage in enterprise IT environments. Cloud storage has become one of the most important storage technologies in recent years [186] and the S3 API developed by Amazon has been widely adopted by other cloud providers [208]. To enable tape use amongst enterprises that are already familiar with the S3 protocol for cloud storage, a variety of products have been developed that provide an S3 interface to tape. Examples include: *PoINT Archival Gateway* [188], *Grau Data XtreemStore* [48], *IBM S3 Deep Archive* [106], *Versity ScoutAM/Gateway* [194, 195], *FUJIFILM Object Archive* [67], *XenData LTO Active Archive* [205], *QStar Archive Storage Manager* [168], and *Spectra On-Prem Glacier* [140].

11.3 Industries

In order to get a comprehensive understanding of how tape storage is used, it is helpful to look at the industries that use tape. According to the 2023 Tape Storage Global Market Report [173], key end users of tape storage include cloud providers, data centers and enterprises. The leading industries are IT, telecommunications, media and entertainment, healthcare, oil and gas, government and defense. A common use case that traverses many industries is data retention on tape for legal compliance reasons [88]. Another common factor motivating the use of tape highlighted in Wu [204] is tape's low energy consumption relative to other storage technologies. Rising energy costs and the impact of global warming affect all enterprises and provide a motivation for reducing IT energy consumption. For enterprises that have a lot of cold data there is an opportunity for significant energy reduction simply by moving this data to tape.

12 Future Tape Scaling, Outlook, and Conclusions

12.1 Tape's Future Scaling Potential and Challenges

Tape storage has undergone a remarkable evolution over its 70-year history, with data rate and capacity experiencing continuous growth. While data rate improvements have been achieved primarily through scaling of the linear density and increasing the number of parallel channels, capacity has increased mostly through scaling of the areal density, complemented by format efficiency improvements and thinner tapes allowing for greater lengths. Currently, a state-of-the-art enterprise class tape drive (IBM TS1170) operates at a native data rate of 400 MB/s, uncompressed capacity of 50TB written at an areal density of 26.1 Gb/in^2 [102]. Compared to the first IBM tape drive, this corresponds to an increase of more than five orders of magnitude in data rate and more than seven orders of magnitude in areal density. However, despite this impressive scaling, the areal density of the TS1170 is still about 48 times lower than the 1,260 Gb/in^2 areal density of a recent 22 TB HDD [178]. Considering that both tape and HDD rely on the same basic magnetic recording principles, this gap implies that from a basic recording physics perspective, tape has considerable potential to continue scaling its areal density. This gap also indicates that there is an opportunity for tape

engineers to continue to leverage and adapt technologies developed for HDD to enable continued tape scaling.

The potential to continue scaling tape areal density has been explored in multiple tape areal density demonstrations [28, 69, 71, 72, 126, 137]. For example, a recent tape demo that used a prototype Strontium Ferrite (SrFe) particulate media showed the potential for tape recording at 317 Gb/in^2 [69], which could enable a cartridge capacity of 580 TB, assuming similar format overheads to current products. Another more recent demonstration on sputtered thin film media showed the potential for recording at 400 Gb/in^2 [126] which could enable even higher cartridge capacities.

The 2024 INSIC Tape Roadmap projects that tape areal density will scale at a 28% **compound annual growth rate** (**CAGR**) over the period 2024–2034 [117]. Combined with a 2.5% CAGR in tape length and small improvements in format efficiency, capacity is projected to scale at a 32% CAGR. Over the same period, data rates are projected to scale at a 15% CAGR. This scaling results in a projected native cartridge capacity of 723 TB recorded at an areal density of 315 Gb/in^2 and a tape drive data rate of around 1.6 GB/s by 2034. Note that the areal density projected at the end of the roadmap is lower than both the aforementioned tape demos and is well below that of current HDD products, indicating the potential to continue scaling beyond the end of the roadmap. The IEEE International Roadmap for Devices and Systems: 2023 Mass Data Storage—Tape Storage Roadmap projects similar scaling rates with predictions out to 2037 at which time native cartridge capacities are predicted to reach 1.5 PB, with an areal density of about 602 Gb/in^2 [33]. Both these roadmaps discuss many of the challenges and the technologies that will likely be required to enable this scaling.

The tape research demos mentioned above used a single recording channel to explore the potential for head and media technologies to support future tape areal density requirements. This aspect arises from the use of HDD heads with a single, very narrow reader to explore these future operating points. Several of these demos also reported other key technologies that will be needed in the future such as technologies to achieve tape track-following accuracies down to the nanometer scale and the potential to write ultra-narrow tracks on tape. However, the additional challenges such as TDS, that arise in tape recording due to the use of multiple parallel channels, were not addressed. Hence, to achieve such areal densities in commercial tape products, continued research and development will be needed in several areas. Foremost is the need for active TDS compensation schemes that can achieve performance levels comparable to those demonstrated for track-following. Additional future research areas include continued improvements in tape media, tape-head tribology, data channel, ECC and servo control, both in terms of refining the technologies developed for these demonstrations for use in commercial tape systems and developing new technologies to enable further scaling.

Another key research topic is adapting state-of-the-art HDD reader and writer technologies for use in tape heads. For example, HDD introduced monopole writers in combination with a media that incorporates a **soft under layer** (**SUL**) in the 2005 time frame [117] at areal densities around 150 Gb/in^2. At current rates of scaling, tape will likely reach this areal density range in the early 2030's time frame. Monopole writers with a SUL enable much stronger write fields that are necessary for writing very high coercivity media. If this technology can be implemented in tape write heads it could enable the use of state-of-the-art HDD media technology. Moreover, if a SUL can be combined with a particulate mag-layer, monopole writers could enable the use of very high coercivity particle technologies such as spherical epsilon-Fe_2O_3 [156]. Another HDD head technology likely to be needed for tape is the combination of TMR readers with soft biasing and side shields that the HDD industry introduced at a reader width of about 40 nm [142]. The use of soft biasing helps to reduce sensor variability and enables the use of side shields that reduce the pickup of signal from adjacent tracks.

To scale tape beyond the limits of conventional perpendicular magnetic recording, tape will likely have to adapt one of the two so-called energy assisted magnetic recording technologies that the HDD industry has been developing. The first of these, called **heat assisted magnetic recording** (**HAMR**), uses a laser and near field transducer integrated on the HDD head to locally heat the recording layer to temporarily lower its coercivity [133]. Implementing this technology for tape would be very challenging. For example, each write transducer in the head would require a laser and near field transducer. Currently tape drives have 64 write transducers and are expected to increase this number to 256 within a decade. HAMR technology also requires the careful design of the heat sinking of the recording layer into the HDD platter to constrain the local heating to the track being written. It seems unlikely that similar heat sinking could be achieved on the 4–5 μm thick polymer substrate used for tape. The second energy assisted recording technology under development is called **microwave assisted magnetic recording** (**MAMR**). In MAMR, a spin torque oscillator is integrated into the write head and used to assist in the recording of very high coercivity media [209]. The spin torque oscillator is fabricated using thin film microfabrication technology during the fabrication of the writer and would, therefore, be more practical to integrate into the manufacture of tape heads. In addition, MAMR does not require the heat sinking needed for HAMR again making it a better candidate for use with tape. Initial research into microwave assisted switching for patriclate based tape media has also been reported [158].

To enable significant future increases in tape data rate, it will be necessary to increase the number of parallel channels. This will most likely occur by first doubling to 64 channels and then later doubling again to 128 channels. This straightforward sounding change impacts many aspects of the drive architecture, including the tape layout, the head and flex cables, the drive card and all of the the custom ASICs needed to drive the parallel channels, i.e., write drivers, analog front and backends, and the main ASIC that hosts the data and servo channels and ECC. The continued scaling of ASIC technology has enabled tape engineers to develop new ASICs with additional parallel channels without significant increases in the ASIC die size. However, as discussed in Sections 5 and 6, operation at variable tape speeds, the removable nature of tape media and backwards compatibility all add significant complexity to the design of these ASICs. Another challenge arises from the rising costs of designing custom ASICs as the technology nodes scale to smaller transistor sizes and more expensive lithography and mask technology. Considering the relatively small volume of ASICs required to meet the demand for tape drives, the tape industry has to amortize the cost of designing and manufacturing custom ASICs by reusing them in multiple generations and across different product families. In spite of these challenges, tape has significant potential and a clear roadmap to continue scaling for multiple future generations. From a technology perspective, we believe tape has the potential to scale to cartridge capacities on the order of a petabyte and beyond. However, whether such capacities will be achieved in commercial products also depends on a variety of additional business factors such as the market demand for such products and the amount of funding available to develop them.

12.2 Outlook and Conclusions

Recent studies have reported that data is growing exponentially and is expected to continue to do so in the coming years. Much of this data is transient and may not need to be retained, however worldwide storage capacity is nevertheless projected to also grow exponentially. For example, a recent study estimated the total installed base of storage capacity to be around 8 zettabytes in 2021 and projected it will grow to 16 zettabytes by 2025 [192].

Multiple factors are responsible for the growth in data and data retention needs, including the growth in smartphone users (currently estimated at 2.87 billion [187]), the growth in connected

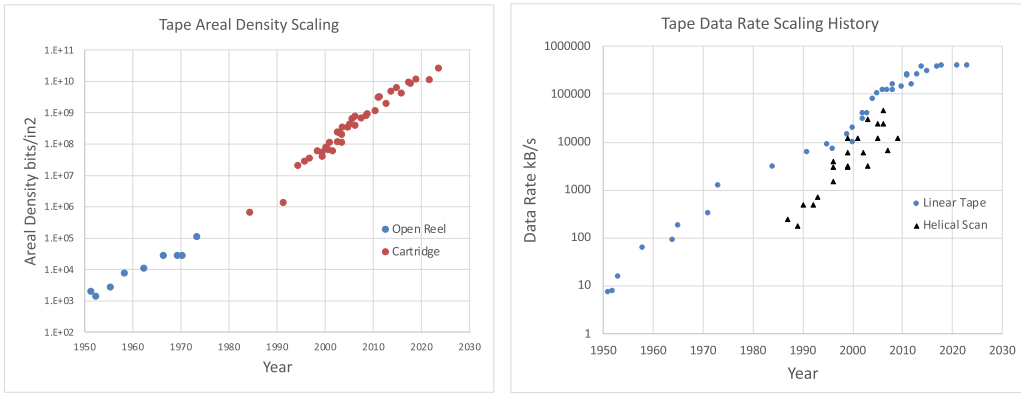


Fig. 32. Left: historical areal density scaling of linear tape systems. Right: historical data rate scaling of tape drives. Note that the data rates of open reel products (1951–1971) are kilo-characters/sec.

devices (projected to grow to 41.6 billion by 2025 [171]), regulatory changes in retention periods for data such as health records, and the growth in the size of training datasets for AI models and the need to preserve the raw data for provenance purposes. Emerging AI technologies also hold the promise of providing new means for extracting value from data, providing an incentive to preserve data for potential future use. A white paper from Horizon Storage Strategies reported that over 60% of all data is archival and projected that this could reach 80% or more by 2024 [187]. Considering all these factors, we conclude that the demand for archival storage capacity will grow exponentially in the coming years.

Magnetic tape storage provides a cost-effective way to retain the exponentially increasing volumes of data currently being created. Tape's low cost per terabyte and enhanced cyber-resiliency combined with its low energy consumption and CO_2 footprint make it an appealing option for storing infrequently accessed data. These factors combined with tape's potential to continue scaling capacity and data rate have resulted in a resurgence in use of the technology and will likely lead to an expanded role to meet the growing need for cost effective, green storage solutions.

13 History and Evolution of Tape Hardware, Media, Software, and Usage

Magnetic tape for data storage was first commercialized more than 70 years ago and remarkably it is still a critical component of modern storage infrastructure, making it the oldest storage technology still in use. The longevity of tape's relevance and success has resulted from continued improvements in the technology, most notably in the scaling of areal density which has increased by more than seven orders of magnitude and of data rate which has increased by more than five orders of magnitude to date, as illustrated in Figure 32. Over the same period the size of a tape drive has been reduced from that of a large refrigerator to the size of a shoe box. The following section presents a history of magnetic tape for data storage, starting with a brief look at its origins in the magnetic recording of analog audio and then focusing on important milestones in the development of tape for data storage. It is not intended as an exhaustive history of all tape products but will rather focus on important milestones and attempt to illustrate how tape storage has evolved and scaled as well as highlighting key technologies that have enabled this scaling. The first subsection focuses on tape hardware development, followed by a subsection on the development of tape media and a subsection on tape usage and software.



Fig. 33. IBM 726 Magnetic Tape Recorders. Reprint Courtesy of IBM Corporation © (2024).

13.1 The Evolution of Tape Hardware

The first practical demonstration of magnetic recording was performed by Valdemar Poulsen in 1898 who recorded analog audio on a steel piano wire [46]. Recording on something resembling modern tape was first demonstrated by Fritz Pfleumer in 1927 who recorded analog audio on thin strips of paper coated with iron oxide powder [47]. Another key step was the 1934 development by Eduard Schüller of the ring recording head which provided a more focused magnetic field and enabled more precise control of the magnetization of small regions of the magnetic media [47]. Over the next two decades, a foundation for the development of tape data storage were laid in the continued improvement of analog audio recording on magnetic tape, led primarily by AEG in Germany and Ampex in the U.S.A. [79]. The first magnetic tape system for data storage, the UNISERVO I, was commercialized by the Remington Rand corporation in 1951 and was the primary I/O device for the UNIVAC I computer. The UNISERVO I used a metal tape made from a very abrasive nickel-plated phosphor bronze alloy. It recorded data using a staggered head in an 8-track format using six tracks for data, one for parity and one for timing. A full reel of tape weighted about 25 pounds requires the use of large motors and a complex pulley and lever mechanism to buffer the tape and enable the rapid start and stop motion required to input small blocks of data to the computer for processing. The block size was limited by the small amount of memory available in early compute systems. Data was recorded at 128 bits/inch at a tape speed of 100 inches/sec. The large inertia of the reels necessitated the use of a relatively large 2.5 inch inter block gap between data blocks of 60 words of 12 characters, resulting in a transfer rate of about 7,200 characters per second. The 1,200 foot long tape of the UNISERVO I had a capacity of about 1.5 MB [201]. Note that early tape systems recorded data as characters rather than bytes, however, to simplify comparisons to later products the capacities of these systems have been converted to an effective capacity in bytes. In 1952, IBM announced the IBM 726 magnetic tape unit for use with the IBM 701 computer. The 726, shown in Figure 33, implemented an innovative vacuum column technology to buffer the tape during start and stop, which enabled the use of a much lighter but more fragile polymer-based tape and a smaller 0.75 inch inter block gap. The 726 had a 7-track format with an in-line head, using 6 tracks for data and one for parity. The 726 recorded at a linear density of 100 bits/inch

in an NRZI encoding format, at tape speeds of 75–100 inches per second. The tape had a length of 1,200 feet and a capacity of about 2 MB. The use of vacuum columns, polymer tape and the 7-track format were widely adopted and became a de facto industry standard. Signal amplification and logic were implemented using vacuum tubes. An important focus in early tape development was on increasing data rates to match the increasing speed of compute systems. Initially this was achieved through a combination of linear density and tape speed scaling. For example, the IBM 727, released in 1955, operated at 15 kB/s and the 729-3, released in 1958 at 63 kB/s. Starting with the 729-2, the vacuum tubes used in earlier drives were replaced with transistors. Another important innovation of the 729 was the introduction of a dual gap head that enabled on the fly read while write verification of the data [83].

The first automated tape library, the IBM 7955 (Tractor), was developed in 1962 and delivered to the NSA as part of the Harvest computer system. The design was very innovative for its time, using 1.75-inch wide tape, housed in dual reel cartridges and the library was equipped with an automated mechanism to fetch, load and unload cartridges from the six tape drives that made up the system [165]. However, only one system was ever made, and it would be another 30 years before automated tape libraries began to see widespread adoption. The next decade saw a continued focus on tape I/O performance as well as innovation in tape handling and system integration. In 1964, IBM launched the IBM 2104 (Models 1-3) for use with the IBM System/360. The 2104 was the first 9 track tape format (8 for data and 1 for parity) and introduced CRC. It had a capacity of about 5 MB per reel and a data rate of up to 90 kB/s. The 2104 Models 4-6 followed in 1965 that had a capacity of about 10 MB and a data rate up to 180 kB/s. Two key innovations that enabled this increased performance were the implementation of electronic skew buffers and self-clocking tracks [83]. The IBM 2420, announced in 1968, operated at a data rate of 320 kB/s and introduced a self-threading mechanism for the tape that provided significant operator time savings. The use of DC motors enabled a tape speed of 200 inch/sec and a remarkable start/stop time of less than 2 ms [83]. Three models of the IBM 3420 were introduced in 1971 for use with the IBM System/370 computer and operated at a linear density of 1,600 bpi and data rates up to 300 kB/s. In 1973, three additional models were introduced operating at a linear density of 6,250 bpi and data rates up to 1,250 kB/s. The 3420 embedded a set of around one thousand hardware-assist instructions that eliminated the need for a separate switching unit to provide more than one CPU access to the tape drive and can be viewed as the beginning of tapemicrocode [110].

At the end of 1964 **Digital Equipment Corporation** (**DEC**) introduced the relatively low cost DECtape system (also called Microtape) with its PDP-7 minicomputer. DECtape was based on the LINCtape system designed at MIT-Lincoln Labs. Compared to the other tape systems discussed above (and those that follow), DECtape had many unique aspects. For example, DECtape was a random access block addressable storage device that essentially behaved like a hard disk with very high latency and could be used as the main storage for the operating system. It used 250 feet of 3/4″ tape wound on an ~ 4″ reel with 10 tracks, (6 data, 2 mark tracks and 2 clock tracks). Mark and clock tracks were pre-formatted during tape manufacturing. To improve reliability, each track was paired with a non-adjacent track that contained the same data. To make the tape durable enough to support random IO, the surface of the magnetic recording layer was coated with mylar to protect the recording layer from wear. In 1978, DEC introduced DECtape II which was a similar block addressable random-access device but used 0.15″ tape housed in a miniature cartridge [16, 40–42].

In 1974, IBM released the first commercially available automated tape library, the IBM 3850 **Mass Storage System** (**MSS**). The 3850, shown in Figure 34, was very innovative for its time, utilizing a variant of helical scan [46, 157] recording technology and cylindrical plastic cartridges with a 1.86-inch diameter and 3.49-inch length. The cartridges held 770 inches (20 m) of tape with a 50

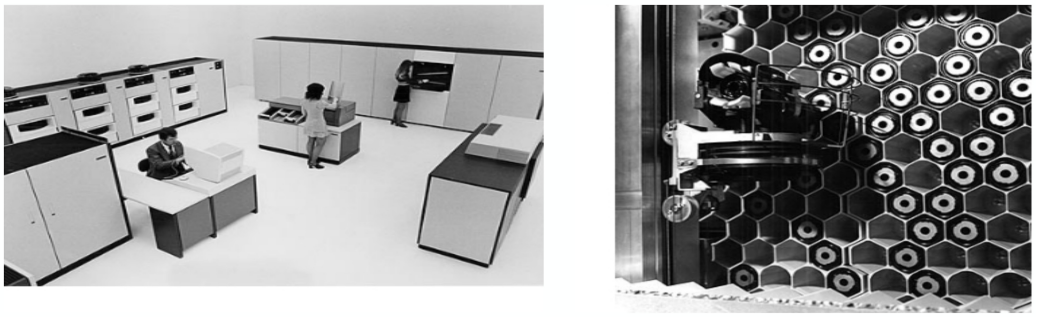


Fig. 34. Left: IBM 3850 MSS, Right: honeycomb storage compartments of the 3850. Reprint Courtesy of IBM Corporation © (2024).



Fig. 35. Left: IBM 3480 Tape Subsystem, Right: 10.5" reel and 3480 cartridge. Reprint Courtesy of IBM Corporation © (2024).

MB capacity and were housed in honeycomb cells along 2 walls of the library that were accessed by two robots (accessors). Data was staged in and out of the system via DASD (HDD) and the host and application treated all data as if it were stored on DASD making the 3850 the first instance of a virtualized storage system. Despite the innovation, the concept was ahead of its time and the 3850 had limited commercial success with no follow-on products [46, 84, 121].

In 1984, IBM released the 3480 magnetic tape subsystem that became a de-facto industry standard; companies including Fujitsu, M4 Data, Overland Data, StorageTek, and Victor Data made tape drives compatible with the 3480 standard. The 3480, shown in Figure 35, was a significant departure from previous open reel tape drive designs and introduced multiple new innovations including a chromium dioxide particle based media housed in a 4" x 5" x 1" cartridge that replaced the previously used 10.5" reels and set the stage for automated tape libraries. In addition, the 3480 introduced an 18-track MR head manufactured using thin film technology. This was the first use of thin film technology in magnetic recording and was a key technology to enable future areal density scaling via the continuous miniaturization of the write/read transducers through improvements in lithography technology. In addition, the 3480 eliminated the vacuum columns used in previous open reel drives enabling a more compact drive that required half the floor space of the 3420. The 3480 also achieved significant improvements in error detection and correction using **adaptive cross parity** (**AXP**) coding. The 3480 had a capacity of up to 400 MB and operated at a data rate of 3 MB/s. In 1986, IBM released the IBM3480 IDRC (improved data recording

capability) which added hardware-based data compression and enabled a 2x increase in capacity and data rate [46, 95]. In the same year autoloaders were added to the 3480. The autoloader held up to seven cartridges and automatically exchanged the current tape for a queued cartridge from the bottom of the loader. The 3490E drive released in 1989 used a 36-track format and provided an increase in native capacity to 800 GB [23, 84].

In the same year the 3480 was released, DEC released the TK50 Compac Tape drive which was intended for use with minicomputers rather than mainframes. The TK50 used a 22-track tape format, recording with a single channel head in a linear serpentine fashion. It had a data rate of 45 KB/s and a cartridge capacity of 94.5 MB [43]. The second generation, the TK70, was released in 1987 and provided 294 MB capacity at the same data rate as TK50 [44]. The third generation, the THZ01, which was later rebranded as the DLT260, was released in 1991 with a capacity of 2.6 GB and a data rate of 800 kB/s. The THZ01/DLT260 introduced the use of cylindrical guide rollers to guide the tape through the tape path. Over the next 8 years there were multiple follow-on versions of DLT (digital linear tape) each providing increases in capacity and data rate. The number of channels was first increased to two in the DLT2000 drive and then 4 channels in the DLT 7000. The final version, the DLT8000, was released in 1999 and had a capacity of 40 GB and a data rate of 6 MB/s. In 2000, Quantum introduced the SuperDLT (SDLT format) that used an optically read servo pattern on the backside of the tape for track following [176]. The first generation SDLT 220 had a capacity of 110 GB and a data rate of 16 MB/s. The final generation was released in 2007 with a capacity of 800 GB and a data rate of 60 MB/s.

In 1987, Storage Tech (STK) launched the Cimarron 4400 ACS (Automated Cartridge System) which became the first tape library to achieve major commercial success and introduced the concept of "Nearline" storage. The 4400 used IBM 3480 compliant tape drives and was based on a modular cylindrical library called a silo. A silo supported up to 16 tape drives, held up to 5,500 200 MB cartridges for a total capacity of over 1TB. The robot accessor had integrated cameras to read barcoded labels on the cartridges. In 1992 STK released the follow-on PowderHorn library that held up to 6,000 cartridges and provided up to 350 cartridges/hour accessor performance as well as smaller libraries called TimberWolf and WolfCreek that held up to 500 and 1,000 cartridges, respectively [2]. In 2005, Storage Tech launched the SL8500 library with a capacity of up to 10,088 cartridges and up to 64 drives but has since discontinued development.

In the late 1980's in Japan, a variety of automated tape libraries were also developed. Examples include the NEC N7645 library announced in 1988 with a capacity of 6,250 cartridges, the Fujitsu F6455 with up to 5152 cartridges and the Hitachi H-6951-1 library with up to 6,560 cartridges [155].

Exabyte Corp. was founded in the mid 80's with the goal of using consumer videotape technology [46, 157] for data storage. In 1987 Exabyte released the EXB-8200 drive, the first commercial data tape drive to use video helical scan technology. The EXB-8200 operated at a data rate of 246 kB/s, had a native capacity of up to 3.5 GB and used 8 mm particulate based consumer video tape media. In 1990, Exabyte released that EXB-8500 that provided an increased data rate of 500 kB/s and in 1992 they released the EXB-8505 that provided an increased native capacity of up to 5 GB. The EXB-8900 Mammoth drive was released in 1996 and used **advanced metal evaporated** (**AME**) media to achieve a 3 MB/s data rate and 20 GB capacity. The final version, the Mammoth-2 was released in 1999 with a data rate of 12 MB/s and capacity up to 60 GB.

In 1989, Sony released the first generation of **Digital Data Storage** (**DDS**) which was based on **digital audio tape** (**DAT**). It used helical scan recording [46, 157] on a 3.81 mm tape and achieved a capacity of 2 GB and a data rate of 183 kB/s. Over the next decade, 5 more generations of the technology were released, scaling capacity and data rate with each generation. The last two generations used 8 mm wide tape and operated at capacities 80 GB and 160 GB and data rates of 6.9 MB/s and 12 MB/s, respectively.



Fig. 36. IBM 3495 robotic tape library. Reprint Courtesy of IBM Corporation © (2024).

In 1993, IBM launched the 3495 robotic tape library that used a large (c.a. 400 kg), bright yellow, 6-axis industrial robot to serve cartridges to tape drives that were housed in a linear string of frames, as illustrated in Figure 36. The base model was 13.4 m long and held 5,660 cartridges. Three larger configurations were also available, the largest of which was 28 m long and held up to 64 tape drives equipped with autoloaders and up to 18,920 cartridges. The robot could mount 120 cartridges per hour but was not fast enough to keep 64 tape drives continuously busy [76, 84]. The internal IBM code name for the product was Caballero, but it was often referred to as Conan, as in Conan the Librarian. In the same year, IBM also released the 3494 library for mid-range/open systems. Cartridges were stored horizontally rather than vertically as in the 3495, which simplified the accessor design considerably. The 3494 was a more compact design with a much lighter custom designed robot that could perform up to 250 mounts per hour. The first version consisted of 2 frames, with two 3480 tape drives without autoloaders. This was extended to 8 frames in 1994 and then 16 frames in 1996 providing a capacity of up to 6420 cartridges. In 1997, dual accessors were added, improving the mount rate to 610 mounts per hour. A further enhancement to the 3494 was the introduction of the **Virtual Tape Server** (**VTS**). The VTS was a large disk cache and server responsible for managing the data and tape media and was placed between the host and tape library. The VTS enabled a much more efficient use of tape capacity and simplified migration to newer tape technology [84, 131]. In 2000, IBM released the 3584 library that could handle both LTO (see below) and DLT technology. The 3584 used a similar linear architecture as the 3494 and was configurable with 1 to 16 frames and up to 192 drives. In 2006 it was renamed the TS3500. In 2008 the capacity was enhanced through the introduction of HD (high density) slots in which multiple cartridges are stored one behind the other, providing a capacity of up to 1,320 LTO cartridges in a single frame [76, 84]. Starting around 2004 IBM also began offering a range of mini and mid-range libraries. Examples of other companies producing tape libraries in this period include Qualstar, ADIC (Advanced Digital Information Corporation), Hewlett Packard, and later additional companies including Overland and Spectralogic. IBM's latest large scale tape library offerings include the TS4500 introduced in 2014 which has a similar architecture to the TS3500 and the single frame TS6000 Diamondback introduced in 2022.

In 1995, IBM released the 3590 Magstar MP (multi purpose) tape drive which was the first tape drive to implement track follow servo control. See Figure 37. Track follow servo was a key innovation that enabled a new era of much faster track density scaling that continues today. The 3590 used an amplitude based servo pattern that is described in reference [159]. The first version of the 3590 (Model B) had 128 tracks and provided a capacity of 20 GB and data rate of 9 MB/s. It

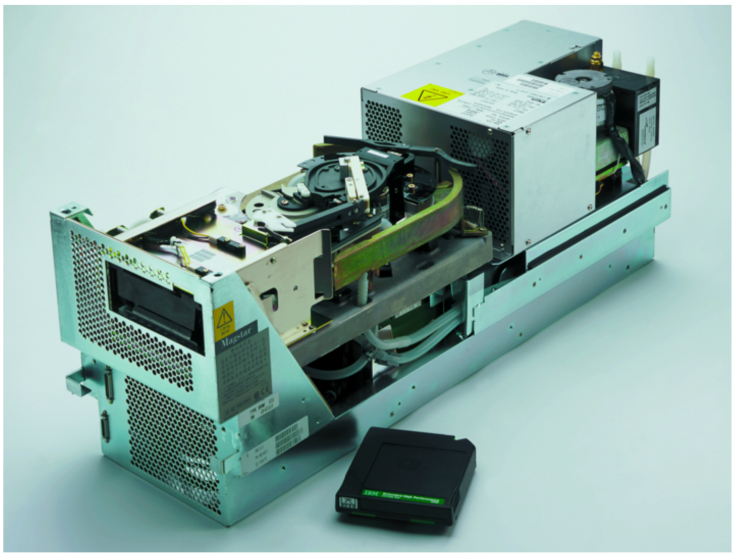


Fig. 37. IBM 3590 Magstar MP Tape Drive. Reprint Courtesy of IBM Corporation © (2024).



Fig. 38. IBM Magstar MP 3570. Reprint Courtesy of IBM Corporation © (2024).

was followed by the Models E and H in 1999 and 2002 which had capacities of 40 GB and 60 GB respectively; both operated at 14 MB/s.

In 1996, IBM introduced the 3570 tape drive targeting mid-range computer systems. See Figure 38. It used a dual reel cartridge and implemented a novel midpoint load design to minimize data access time and boasted an impressive load ready time of 6.7 s and average seek time of only 8 s. The 3570 was the first tape drive to use the TBS technique that has become standard in all recent linear tape drives. The 3570 operated at a data rate of 7 MB/s with an initial capacity of 5 GB that was extended to 7.5 GB in 1999 and 10 GB in 2002 [76].

In 1996, Sony introduced **Advanced Intelligent Tape** (**AIT**) which was based on helical scan recording [46, 157] and used AME media. The first generation had a capacity of up to 35 GB and a data rate of up to 4 MB/s. A total of five generations were released, scaling capacity and data rate to 400 GB and 24 MB/s, respectively, in the 5th generation that was released in 2006. Gen5 of AIT was the first tape drive to use **giant magneto resistive** (**GMR**) reader technology.

In 1999, Ecrix, which later merged with Exabyte corporation, released VXA tape based on an 8mm helical scan technology [46, 157]. The first generation, VXA-1 had a capacity of 33 GB and a



Fig. 39. Left IBM LTO-1 Tape Drive, Right IBM 3592 Tape Drive. Reprint Courtesy of IBM Corporation © (2024).

data rate of 3 MB/s and was followed by two more generations in 2002 and 2005 that had capacities of 80 GB and 160 GB and data rates of 6 MB/s and 12 MB/s, respectively.

In the 1970's, 80's, and 90's there was a proliferation of tape products and formats. In addition to those discussed above, other examples include the **quarter inch cartridge** (**QIC**) format launched in 1972, QIC-Wide launched in 1994, Travan launched in 1995, QIC-EXtra (QIC-EX) launched in 1996, and SLR (scalable linear recording) launched in 1997. Most of the formats of this era were proprietary. Incompatibility between the many formats made it difficult for customers to change technology and vendor and hence tended to lock customers in. In response to this situation, in the late 90's IBM, HP and Seagate formed the LTO consortium with the goal of developing a new more open format that provided interchangeability between drives and media from different manufactures. In 1998, the consortium announced the LTO Roadmap. Initially two formats were planned, a single reel format called Ultrium and a dual reel, mid-point load format based on the IBM 3570 format that was called Accellis. The Accellis Gen 1 format was never commercialized and no follow on formats were developed. As a result the term LTO is currently used to refer exclusively to the Ultrium format. In 2000, IBM, HP, and Seagate each released LTO Ultrium Gen 1 drives (see Figure 39). LTO-1 had a capacity of 100 GB and a data rate of 20 MB/s. It used a linear serpentine recording format with an 8-track head that spanned only a quarter of the width of tape reducing the sensitivity to TDS effects by about a factor of 4. To read and write across the full width of tape, the IBM LTO-1 drive adopted a novel architecture that combined a coarse actuator with a fine actuator for track follow servo and replaced the air bearing tape guides used in early IBM drives with tape guide rollers. Another major innovation was the introduction of flat lapped heads built using HDD head fabrication technology on AlTiC wafers to replace the much larger contoured and slotted nickel-zinc ferrite heads used in earlier tape drives [20]. In addition to other advantages, the light weight flat lapped heads enabled higher bandwidth servo control and hence better track following performance. However, the most revolutionary innovation of LTO was interchangeability, i.e., any LTO1 cartridge provided by the multiple LTO licensed media vendors could be written/read in drives manufactured by any of the LTO drive vendors and then later written/read in a drive from any of the other vendors. LTO-1 drives and follow on generations were made in a higher performance FH form factor and a lower performance/lower cost HH form factor intended for lighter workloads. LTO tape is often referred to as a "mid range" tape solution in terms of performance and cost relative to "Enterprise" solutions often used with mainframes and low-cost solutions used with PCs. Competition between the multiple LTO drive and media providers resulted in more competitive pricing compared to proprietary formats and contributed to LTO becoming the dominant tape format. To date, nine generations of LTO have been brought to market. LTO-2 was introduced in 2002 with a 200 GB capacity and 40 MB/s data rate and introduced a PRML (partial response maximum

likelihood) data channel and a more efficient 16/17 modulation code. LTO-2 drives were also backwards compatible in that they were able to read and write LTO-1 cartridges. Seagate's tape business was renamed Certance in 2003 and then acquired by Quantum in 2004. LTO-3 was also released in 2004 with a new 16 track head format and again doubled capacity and data rate to 400 GB and 80 MB/s, respectively. The potential in linear tape recording to scale the number of parallel channels in a straightforward manner enabled faster data rate scaling compared to helical scan based technologies and was also a contributing factor to LTO becoming the dominant format. LTO-4 was released in 2007 with an 800 GB capacity and 120 MB/s data rate. Tandberg Storage also participated in the LTO consortium for three generations, releasing HH versions of LTO-2 (TS400, 2005), LTO-3 (TS800, 2007), and LTO-4 (TS1600, 2008) drives. LTO-5 was released in 2010 with a 1.5 TB capacity and 140 MB/s data rate followed by LTO6 in 2012 with a 2.5 TB capacity and 160 MB/s data rate. LTO-5 introduced the capability to partition the tape into 2 tape partitions and LTO-6 extended this capability to 4 partitions. LTO-7 was released in 2015 with a 6 TB capacity and used a 32-channel head to enable a data rate of 300 MB/s. The two most recent generations, LTO-8 (2017) and LTO-9 (2021) use the same 32 channel head format and operate at capacities of 12 TB and 18 TB and data rates of 360 MB/s and 400 MB/s, respectively. The latest LTO roadmap describes 5 future generations of the technology, each of which is expected to double the capacity of the previous generation.

SUN Microsystems acquired STK in 2005 and subsequently launched a new family of Enterprise class drives in 2006 that used a linear serpentine format and a single reel cartridge. The first generation, the T10000 had a capacity of 500 GB and a data rate of 120 MB/s. The second generation, the T10000B was released in 2008 with a capacity of 1 TB and a data rate of 120 MB/s. In 2009, Oracle acquired SUN and subsequently released the T10000C in 2011 with a 5 TB capacity and 240 MB/s data rate. In 2013, Oracle launched the T10000D drive with a capacity of 8.5 TB and a data rate of 252 MB/s. A T10000E drive was planned but eventually cancelled when Oracle stopped tape development in 2016.

In 2003, IBM also launched a new family of Enterprise class drives. The first generation was branded 3592 and operated at a capacity of 300 GB and a data rate of 40 MB/s. See Figure 39. The second generation, the TS1120, operated with a capacity of 700 GB and a data rate of 100 MB/s. The TS1120 was the first storage device to provide built-in hardware-based encryption which later also became standard in LTO drives. The TS1130 was released in 2008 and used GMR reader technology to provide a capacity of 1 TB and data rate of 160 MB/s. Compared to LTO drives of the same time frame, Enterprise class drives typically provided higher performance in terms of capacity, data rate, access time and error rate and in addition, enabled "up formatting" of the previous generation of media to a higher capacity using the latest generation of tape drive. The TS1140 was released in 2011 with a 4 TB capacity and a data rate of 250 MB/s. The TS1140 introduced flangeless tape guides that reduced media edge wear as well as active skew control and technology to improve rewrite performance. In addition, the TS1140 introduced a new **data dependent noise predictive maximum likelihood** (**DD-NPML**) data channel, a more powerful ECC scheme that doubled the length of the C2 code, a novel 3 module head design with 32 active channels and a new servo pattern. In 2014, the TS1150 was released with a 10 TB capacity and 360 MB/s data rate. The TS1150 introduced a new read head based on TMR (tunneling magneto resistive) readers, the first use of TMR technology in tape. The TS1155 was released in 2017 with a 360 MB/s data rate and a 15 TB capacity that was partially enabled by the introduction of a new writer with notched poles and a new high moment liner technology. In 2018 the TS1160 introduced a new tension based active TDS compensation scheme that helped to enable a 20 TB capacity and 400 MB/s data rate. The TS1160 introduced new ECC technology that included a new C2 code that again doubled the code length as well as a novel iterative decoding scheme. The latest generation, the TS1170 was released in 2023 with a 50 TB capacity and 400 MB/s data rate. Many of the technologies introduced

in the TS11xx family were also implemented in subsequent generations of LTO drives and were critical to enabling the scaling of IBM's LTO drives. Currently, LTO and IBM TS11xx Enterprise are the only tape formats still under active development. However, despite consolidation in the industry and convergence to two formats, tape has a healthy ecosystem with companies including IBM, HPE, Dell, Quantum, Spectralogic, Overland and Oracle selling tape drives and libraries, two media manufactures and multiple brands of media. Moreover, tape has significant potential for continued scaling as discussed in Section 12.

13.2 The Evolution of Tape Media

Apart from the UNISERVO 1 tape system that used a metal tape made from nickel plated phosphor bronze alloy, all other commercial data tape media have used a polymer substrate coated with a thin layer of magnetic material. Companies who have manufactured magnetic data tape media include: BASF, Datatape/Graham Magnetics, Fujifilm, Imation, Maxell, Memorex, SONY, TDK, and 3M, with both Fujifilm and Sony still actively developing new tape media products. Two types of magnetic coating have been used: particulate and **metal evaporated** (**ME**) coatings. In particulate based tape, the recording layer (mag-layer) is made up of small magnetic particles fixed to the polymer substrate with a binder (i.e., "glued" to the substrate). The tape is manufactured by coating the particles onto the substrate in a liquid slurry that also contains solvent and binder, followed by evaporation of the solvent to "bind" the particles to the substrate. In ME tape the mag-layer consists of a continuous, granular magnetic metal film deposited on the polymer substrate under vacuum by evaporation. Exabyte's Mammoth 1 and 2 as well as Sony's AIT and SAIT tape, used a metal evaporated mag-layer technology that is described by Kawana et al. [130]. All other polymer tape media to date, including state-of-the-art products, are particulate based and, therefore, we will focus on the evolution of particulate tape technology here.

Early generations of particulate tape, such as that used for the IBM 726, had a simple bilayer structure consisting of an approximately 10 μm thick layer of magnetic particles and binder coated directly onto an acetate substrate. The magnetic particles were an acicular form of gamma ferric oxide with particle lengths on the order of 0.2 to 0.8 microns. To achieve the required surface quality, i.e., low defect density, "IBM designed and built the world's most advanced tape coater in Poughkeepsie, and the first clean room used in manufacturing." [23]. In later generations of tape, a thin, carbon loaded back-coat layer was added to control tribo-charging effects and improve the winding/unwinding properties of the tape. Over time, the thickness of the magnetic recording layer and the size of the magnetic particles have been continuously reduced to enable scaling of the areal recording density, as discussed in Section 3. In the 1980's, FeCo particles with a length on the order of 250 nm were used in a mag-layer with a thickness on the order of 3–5 μm. In the 1990's, CrO2 particles were used in the IBM 3480 and then 3490 tape media. In the late 1990's Fujifilm developed a dual coating technique in which a non-magnetic undercoat and a much thinner mag-layer are deposited simultaneously. This technology enabled a significant reduction in the mag-layer thickness to around 300 nm initially with further decreases in subsequent generations to around 100 nm over the next decade. The introduction of a more advanced dual layer coating technology in 2011 enabled a further reduction to around 60 nm [154].

Early generations of LTO and Enterprise tape media (IBM TS11xx and STK T10000) used MP (metal particle) technology which was based on acellular particles with a CoFe core and a (non-magnetic) Yttrium based shell a few nm thick which prevented oxidation of the magnetic core. The magnetic axis of the particles was aligned in the longitudinal direction of the tape, i.e., parallel to the tape transport direction through the application of a magnetic field during the coating process. The first generation of LTO media used particles with a volume on the order of 10,000 nm^3. LTO-1 media had a mag-layer thickness of 220 nm which was reduced to 100 nm for

Gen2 [20]. With each generation of media, the volume of the particles was continually reduced. Reductions were also made in the width of the distribution of particle sizes to improve SNR. LTO Gen4 media, released in 2007 used a particle volume of about 4500 nm^3 [82] and LTO Gen 5 (2010) used particles with a volume of about 2800 nm^3 [28]. Below a volume of around 2800 nm^3 it is difficult to maintain a sufficient particle coercivity to ensure the thermal stability of recorded data [154], as discussed in Section 2.

In 2011, Barium Ferrite (BaFe) particle technology was introduced in Enterprise tape media and then later also adopted in LTO media. LTO Gen6 supported both MP and BaFe media, whereas generations 7 to 9 are based exclusively on BaFe. BaFe particles have a hexagonal platelet shape and a magnetization that results from crystalline order rather than shape anisotropy. The coercivity of the particles can be tuned by doping the particles with elements such as Co, Zi, or Ti. Moreover, BaFe ($BaFe_{12}O_{19}$) is an oxide and, therefore, does not require a non-magnetic shell to protect against oxidation. As a result, BaFe particles can be scaled to much smaller sizes than MP technology. Initial generations of BaFe media (e.g., IBM TS1140 JC tape) utilized a particle volume of 2,100 nm^3 and a mag layer thickness of about 70 nm [71]. In the TS1150 JD media, the particle volume was reduced to around 1950 nm^3 [137]. In both JC and JD media, the particles had an essentially random orientation. The particle volume of the TS1160 JE media was further reduced to 1,700 nm^3 and the particles were partially oriented in the perpendicular direction through the application of a magnetic field during the coating process [69].

The effective distance between the read and write transducers and the magnetic particles, referred to as magnetic spacing, is determined predominantly by the roughness of the surface of the mag-layer. To enable areal density scaling, tape roughness has been continuously reduced as new generations of media were introduced. For example, in the early 90's, the roughness of the mag-layer as measured by optical interferometry, was on the order of Ra 7 nm. The introduction of dual coat technology in the late 90's enabled a reduction to around Ra 5 nm and a continual decrease in subsequent generations to around 1.5 nm by 2012 [154].

Early open reel tape products used an Acetate substrate which was later replaced by PET to improve tape robustness against breakage during the continuous start/stop usage pattern of early tape drives. Later, PEN was also used as a substrate (for example in the IBM 3590 Extended cartridge). Recent generations of LTO tape have used PET, PEN, and Spaltan, which is a blend of PET and Aramid, whereas the most recent generations of enterprise tape media have used Aramid as a substrate. Of these four materials, Aramid is the most robust and has the lowest dimensional change versus environmental conditions, but is the most expensive.

Over the history of tape, the thickness of the tape media has been gradually reduced to enable longer tape lengths and hence higher reel/cartridge capacities. For example, the 8-inch diameter (203 mm) reel used in the IBM 726 held 1,200 ft (365 m) of tape that was about 58 μm thick. The tape cartridge introduced with the IBM 3480 had a diameter of about 96 mm and held 541 ft (165 m) of tape that was about 30 μm thick. LTO cartridges have about the same reel diameter but in Gen 1 they held 609 m of 8.6 μm thick tape and by Gen 9 the length was scaled to 1,035 m of 5.2 μm thick.

The most common tape width format is ½ inch (12.65 mm), introduced with the Uniservo 1 and IBM 726 and still used in LTO and recent Enterprise tape media. However, over tape's 70+ year history a variety of other formats have been used, including 4 inch, 1 inch, ¾ inch, 8 mm, ¼ inch, 4 mm and 1/16 inch.

The Uniservo 1 and IBM 726 used open reels with an 8-inch diameter. In 1953, IBM introduced the 727 magnetic tape unit with 10.5-inch reels that became a de facto industry standard for around 25 years [46]. In 1984, IBM introduced the 3480 that replaced the 10.5″ reels with a 4″ × 5″ × 1″ (101.6 mm × 127 mm × 25.4 mm) cartridge. More recent generations of IBM and

STK/Oracle Enterprise tape have used a similar size cartridge. The LTO tape cartridge format is slightly smaller (102 mm x 105.4 mm x 21.5 mm) and is similar to the dimensions of the DLT cartridge. A variety of dual reel cartridge formats have also been used, such as in the DDS, IBM 3570, and AIT. Dual reel cartridges can provide an advantage in terms of access time, but have a lower volumetric storage efficiency.

In the first generation of AIT, released in 1996, Sony introduced the concept of a CM with their MIC (memory in cassette) technology that was used to store meta data and improve data access performance. In the second generation of AIT, Sony introduced a contactless CM technology called R-MIC and used it to enable the first WORM tape cartridge. LTO and recent Enterprise tape products also use a contactless CM and have also offered WORM capability in all recent generations, starting from STK T9940A, IBM 3592, and LTO-3.

State-of-the-art media technology is described in Section 3 and potential future media technologies are discussed in Section 12.

13.3 Tape Usage and Software History

Over the more than 70 year history of magnetic tape storage there have been profound changes in the way tape is utilized in the computing environment, and corresponding changes in the types of software available for making use of tape storage.

Magnetic tape storage and the wide availability of general-purpose computers both appeared in the early years of the 1950s. However, neither of these technologies entered a vacuum; in fact, both replaced or augmented equipment and processes that had been developed over many decades.

In the 1890s, Herman Hollerith's use of punched-card equipment to process the 1890 U.S. Census ushered in an era of card-based record-keeping and accounting procedures that would form the basis for most early automated business processing. Hollerith's Tabulating Machine Company eventually became a part of the **Computing-Tabulating-Recording** (**CTR**) Company, which was renamed in 1924 to **International Business Machines** (**IBM**) [12].

By the 1930s, IBM and other manufacturers had a wide array of card-based equipment available, including card readers, sorters, printers, collators, and even plugboard-programmable "tabulators" capable of counting, addition and subtraction, and providing multiple levels of accumulators. This equipment was often referred to collectively as **Electric Accounting Machines** (**EAM**) [12].

13.3.1 Early Computing and Tape. By the time the first generation of business-oriented electronic computers became available in the 1950s, EAM equipment was in wide use in academia, business, and government. Early computer systems did not entirely replace this equipment, but instead were often integrated into existing workflows and used to augment the EAM equipment. The card equipment and printers that were developed for EAM were adapted as I/O devices for the early computers, and much of the processing initially remained card-based.

Into this environment came magnetic tape storage. Tape had several benefits over card data storage, including faster I/O time, removal of the 80-character record limit, and much improved volumetric efficiency. However, it also had some drawbacks. Data on tape is intrinsically difficult to re-sort into a different order (a common process in EAM processing) and cannot be accessed or modified manually (like a card in a bin). Additionally, many potential customers were skeptical of the safety of data stored invisibly on a tape reel [12].

The first software specifically written for tape storage would almost certainly have been the **Input/Output Control Systems** (**IOCS**) for accessing tape devices from users' programs. The IOCS provided subroutines that relieved the programmer of having to rewrite the code to perform tape I/O and manage potential errors for each program [89].

Because early computer systems had different architectures (instruction sets, word lengths, etc.) not only between manufacturers but also between different models from the same company, IOCS

routines had to be rewritten for each machine type that could attach to tape devices. These routines were sometimes in the form of an executable card deck appended to an object (executable) program deck, as was the case for tape-only IBM 1401 systems [91]. At other times they were incorporated into a simple disk-based operating system, as was the case for the IBM 1410/7010 [92].

Another early form of tape software provided simple utility functions such as card-to-tape, tape-to-printer, tape-to-punch, and tape-to-tape copies, and record sorting [90]. A classic example of a utility from the System/360 era was called DEBE ("Does Everything But Eat"); such utility functions were important because, as mentioned earlier, much early computer processing remained card-based.

13.3.2 Tape Adoption and Early Usage. Over time the advantages of tape storage over cards became apparent, and tapes began to replace cards as the medium for keeping the master version of an organization's data [12]. Records on the tape "master file" were usually written in order by a key such as employee number or account number. A typical daily or weekly processing run might consist of sorting a set of transactions to be applied to the master file into key order, then processing the transactions against the master tape and writing updated or unchanged master records to a new output tape. The output tape then became the new "generation" of the master file. A system like this had the added benefit of creating point-in-time backups of the master file as it existed before each transaction processing run.

Although smaller computers like the IBM 1401 were often used as the only computing system for a small-to-medium sized business, for larger installations running computers like the IBM 7000 series, smaller systems were frequently used as front-end and back-end processors. Data to be processed by the larger system arrived in the form of cards and was copied to tape by the smaller machine. The tape inputs were then processed on the faster, more powerful system, and the results were written to output tapes. Finally, the smaller system wrote the output reports from tape to a printer. Such a workflow allowed an organization to optimize the use of its larger, more expensive systems [12]. This approach also helps demonstrate the need for simple tape-based utilities.

While disk drives were invented in the mid 1950s, it would be many years before the capacity and cost of disk drives made them an attractive replacement for the large-scale data storage role that tapes filled.

13.3.3 Advanced Operating Systems. The 1960s saw many innovations in the computing landscape, including two very ambitious operating system development projects: Multics [152], the time sharing system project that was a joint venture between MIT, AT&T Bells Labs, and General Electric, and OS/360 [167], the operating system for the new, converged line of System/360 processors from IBM.

Multics was to be MIT's follow-on to its **Compatible Time-Sharing System** (**CTSS**) [197], which ran on an IBM 7090 system using 19 IBM tape units. While CTSS allocated two tape units to each online user, according to one Multics developer tape "was considered a throwback to an earlier age" [151], and was used primarily for backup.

OS/360 became the ancestor of the IBM operating system that was known for many years as MVS and is today the mainframe operating system z/OS [107]. Because OS/360 was initially designed as a batch-oriented operating system, tape support was integral to it. (Time sharing, in the form of the Time Sharing Option, or TSO, was later added to OS/360.)

Both of these projects were very optimistic in their estimates of the amount of time and effort required to build a large-scale, third-generation operating system. At IBM, the OS/360 project was so far behind that it was necessary to build a set of scaled-down, "interim" operating systems. Two of these were eventually named DOS/360 (Disk Operating System), and TOS/360 (Tape Operating System). The two systems shared most of their code, and either could be generated from the same

source files. As their names imply, the difference between them is whether the Operating System is resident on disk or tape. TOS/360 would run on systems having only tape storage and as little as 16K bytes of main memory [167]. (As an interesting aside, DOS/360 became DOS/VSE and eventually z/VSE; IBM support for z/VSE was finally dropped in 2023.)

Due to resource constraints, some managers on the OS/360 project favored requiring disk drives on all System/360 systems. However, a survey of early orders showed that more than 20% of the orders for low-end System/360s were for tape-only systems [167].

13.3.4 Device Independence. When OS/360 eventually became available, one of its notable characteristics was device independence for application programs. Access to datasets (files) was managed by a set of "access methods" (somewhat analogous to modern file system drivers). If an underlying device was capable of supporting a particular access method (for example, sequential access), then the access method took care of all of the device details, and application programmers were isolated from the type of device the dataset resided on [93].

Taking this a step further, the system had a "catalog" that kept track of the volume on which a particular dataset was stored. When a dataset was created, the JCL that initiated the program creating the dataset could specify a device type for the output (e.g., `UNIT=TAPE` or `UNIT=DISK`). When the dataset was later read, the JCL had only to specify the dataset name, and the system catalog would fill in the volume information. Both disk and tape datasets had standard dataset header records that recorded the dataset name and other information. All this allowed disk and tape storage to be used interchangeably for datasets using sequential access [93]. The system also supported stacking multiple datasets on a tape (the catalog would keep track of the dataset number), but in practice this slowed down access and made tape reclamation more difficult, and so was not routinely used.

The device independence of OS/360 meant that many application programs could use either tape or disk storage, with the decision deferred to execution time. Of course, there were applications that used non-sequential (e.g., direct offset) access that could not be mapped to tape.

OS/360 also included a large collection of utility programs [94], most of them supporting tape as well as disk operations. Included were programs for dumping disk volumes to tape as an early form of disk backup. (Data could also be copied to tape at the dataset level for backup, but no backup management software as we know it today was included.) Some utility programs, such as sort/merge, were specifically optimized for tape usage. (One of the most successful early non-IBM software products was an optimized sort/merge program called SyncSort.)

13.3.5 Minicomputers. Another major development of the 1960s was the appearance of minicomputers, small systems with fewer capabilities but a much lower cost than the reigning "mainframe" systems of the day.

Like their larger counterparts, minicomputers had a proliferation of different operating systems. Each of these operating systems had some form of software support for backup to tape, written either by the manufacturer or sometimes by a third-party software company.

The earliest commercially successful minicomputer was arguably the DEC PDP-8, released in 1965. It was just one of a long line of DEC PDP systems.

Many of the PDP-series machines included one or two DECtape devices [16, 40–42]. As mentioned earlier, DECtape was a random access, block addressable device. Machines such as the PDP-8 and PDP-11 could run a small operating system directly from DECtape, sometimes using a second DECtape as a swap device. Of course, DECtape could also be used as data storage, and users often kept their private files on the pocket-sized 4-inch reels.

13.3.6 Unix/Linux. Unix [175] was originally developed on DEC minicomputers (the first being a PDP-7, but quickly reimplemented on the PDP-11). It was developed by AT&T employees who

had earlier worked on Multics, and in fact the name was originally a pun on the name Multics. (Linux is a later reimplementation of all of the major concepts and features of Unix, but as an open-source project.)

One of the key concepts in the Unix operating system is that of a "device file", that is the idea that every device can be identified by a name in the file system tree, and that any device can be read or written (with appropriate authority) by using the name of the device as a file name (for example, `/dev/st0` for a tape in a modern Linux system).

Unix I/O is essentially byte-stream oriented, and any device can be read or written as a byte stream. (Obviously there are devices, such as disk drives, where this would be disastrous, but it is possible. The system is designed so that normal users do not typically have the authority to write directly to system devices.)

For magnetic tape devices, this means that data can be read from or written to tape just by specifying the name of the tape device file as the input or output target for essentially any command. (Once again, with the proper authority.) This gives modern Unix and Linux systems a similar level of device independence for applications as was discussed earlier for OS/360.

There are a set of utility programs in Unix/Linux for managing tape devices. The first is the `mt` (magnetic tape) command, which allows for device control such as rewinding, writing tape marks, skipping to tape marks (files), erasing, and so on. More well-known is the `tar` (tape archive) command, which bundles files into a single archive (similar to a zip file) and (optionally) writes the file to tape.

There are many other utilities which can be used to write and manage backups of data to tape (or other devices) on Unix/Linux systems. One of the best-known is rsync.

13.3.7 Microcomputers. In the later part of the 1970s, a new generation of computer user was introduced to a very different type of tape storage. Early microcomputers such as the TRS-80 and the Commodore PET used cassette tape storage for both programs and data. These systems used commercial audio cassette mechanisms and media to read and write data, and a standard for using audio tapes for computer data called the "Kansas City" standard [53, 144] was developed. Games, assemblers, and utility programs for the systems were released on cassettes during the late 1970s.

Cassette tape storage on these systems worked but was not particularly reliable. Many hours were spent by early microcomputer users waiting for cassette tape programs to load, or trying to debug problems reading back previously written data.

There were some early microcomputer tape systems that attempted to replace the use of cassette drives with more flexible and reliable tape-based devices. One example was the Exatron "Stringy Floppy" [145], released in 1978, which used a tape cartridge about the size of a business card and 3/16 of an inch thick (often referred to as a "wafer"). The tape was in a continuous loop, and moved in only one direction. It was a "direct access" device that could be used in place of a diskette drive. Wafer sizes varied from 5 to 75 feet, which corresponded to about 4 KB–64 KB.

13.3.8 PC Era. The introduction of the IBM Personal Computer (PC) [109] in 1981 signaled the start of a new era in small computers. While microcomputers were already being used in some businesses, the category as a whole was still considered primarily a hobbyist one. The introduction of a small computer with the name IBM changed that, and fairly quickly IBM PCs were visible in businesses of all types and sizes. The term "PC" soon replaced the older term microcomputer.

While the original target price-point for the IBM PC was $1,500, in practice, after adding a monitor, one or two 5¼-inch diskette drives, and perhaps a printer, the price was considerably higher. This made the new IBM system less attractive to hobbyists and home users, and more of a system for business.

The IBM PC did have a cassette tape port in the same form-factor as that on the TRS-80, and cassette-based commands were part of BASIC in ROM. However, given the market for the new system, it is unlikely that many systems ever attached a cassette drive. The cassette port was dropped in the follow-on version, the IBM PC-XT.

By the later half of the 1980s many companies were producing tape backup systems for PCs based on QIC mini-cartridge tape technology. In 1989 PC Magazine tested QIC mini-cartridge drives from 14 different manufacturers, and identified six others [199]. A Computerworld article in 1987 predicted sales of more than 250,000 QIC mini-cartridge drive units for that year alone [148].

One early manufacturer, Colorado Memory Systems, was eventually acquired by HP. HP continued to market tape backup systems using the Colorado brand into the early 2000s. A 1998 version of their product attached to the IDE disk controller of a PC, could use a cartridge with a capacity of up to 8 GB, and included sophisticated backup software [85].

Other tape solutions for PCs based on tape formats such as DLT and DAT were also introduced. However, by the late 1990s most PCs were being shipped with **Compact Disc** (**CD**) drives, and writable CDs with capacities in the 600–700 MB range were becoming common. The ubiquity and ease of use of CD-based optical media soon made it the preferred backup medium for individual PC users.

In the enterprise environment network-based backup software was becoming more common. For example, the Adstar Distributed Storage Manager (ADSM, later re-branded to TSM) [27] was introduced in 1993 and uses a client-server backup architecture. Servers can be run on mainframes, Unix workstations, and even well-equipped PCs, while client software exists for almost every imaginable system from PC-DOS to Cray supercomputers. ADSM/TSM servers typically use a combination of disk and tape devices for storage, and the server supports a wide array of different tape formats and devices.

13.3.9 Workstations. At about the same time that PCs were becoming common, another type of system was entering the marketplace and becoming common in certain environments. This was the workstation, as exemplified by products from companies such as Apollo, Sun, Silicon Graphics, HP, and others. Workstations were usually similar in form factor to PCs: a processor (with memory, etc.) in a case, along with a keyboard and monitor, and perhaps some specialized input/output hardware attached. The difference was primarily one of scale and focus. These systems were often used by a single user but also typically ran operating systems that gave them multi-user capability.

Workstations were usually significantly more powerful in one or more respects than a PC of the same era. Some workstations focused on CPU power, with support for floating-point performance, for example. Others focused on graphics, with high-resolution displays and perhaps 3D graphics acceleration. Still others focused on workgroup support with powerful networking features. Most ran operating systems that were much more powerful than the PC operating systems of the day; often the OS was a derivative of Unix or otherwise Unix-like. As on PCs, tape was commonly used for backup on workstation systems, attached to either the workstation directly or to a central server. DLT tape appears to have been a popular alternative to QIC cartridges for workstation backup.

As PC hardware and software evolved through the 1980s and beyond (often driven by the high-resolution 3D graphics and powerful CPUs required for PC gaming), the differentiation between workstations and PCs blurred, and today there is little to separate the two.

13.3.10 Changing Uses of Tape. As can be seen from the preceding discussions, in the PC, workstation, and even parts of the minicomputer environments the main use case for tape had always been data backup.

However, in the older large systems environment tape had initially been the primary high capacity data storage device. Early disk drives were too small and too expensive to be used to replace tape. But as disk technology advanced, the cost of disk storage became more competitive with tape, driven partly by competition from the many disk drive vendors that sprang up beginning in the 1970s.

Thus during the 1970s and 80s much of the traditional role of tape storage began to shift to disk. There were many factors that drove this change, including:

— the growing capacity and shrinking cost of disk storage.
— the increasing use of real-time access to data through data terminals.
— the use of database systems.
— the growing use of online/time-sharing systems.
— the cost and labor involved in storing, managing, and manually fetching and loading tapes.

During this period, however, new use cases for tape in the enterprise environment evolved and began to grow in importance. The advantages of tape in capacity and long-term cost for nearline and offline data drove its use in new technologies such as HSM, while its value in data protection (backup), and data repository (archive) led to new applications in those areas.

In addition, the transition from large reels of tape to small cartridges allowed the creation of automated tape libraries, which in turn enabled the storage and access of huge amounts of data in a relatively small footprint, and without human intervention.

These technologies and use cases, along with others that have developed more recently, have already been covered in detail in earlier sections.

Acknowledgements

The authors are grateful to Teya Topuria and Eugene Delenia from the IBM Almaden Lab for providing TEM images of tape media, to Lee Randall from the IBM Tape Development Lab in Tucson for providing SEM images of tape media and to Jason Liang of the IBM Tape Head Development team in San Jose for providing the images of tape writers.

References

[1] Fumito Akiyama. 2015. Latest Generation of Magnetic Tape Offers Massive Storage. (2015). Retrieved January 02, 2025 from https://asia.nikkei.com/Business/Biotechnology/Latest-generation-of-magnetic-tape-offers-massive-storage

[2] D. Allan. 2011. STK 4400 Automated Tape Cartridge System (ACS). (2011). Retrieved January 02, 2025 from http://s3.computerhistory.org/groups/stk-4400-20121031.pdf

[3] Amazon. 2024. Amazon Simple Storage Service Documentation. (2024). Retrieved January 02, 2025 from https://docs.aws.amazon.com/s3

[4] Apple. 2024. Set Command CBD. Retrieved March 11, 2024 from https://developer.apple.com/documentation/scsiperipheralsdriverkit/3946455-setcommandcdb

[5] Armando J. Argumedo, David Berman, Robert G. Biskeborn, Giovanni Cherubini, Roy D. Cideciyan, Evangelos Eleftheriou, Walter Häberle, Diana J. Hellman, R. Hutchins, Wayne Imaino, et al. 2008. Scaling tape-recording areal densities to 100 Gb/in^2. *IBM Journal of Research and Development* 52, 4.5 (2008), 513–527.

[6] Thomas C. Arnoldussen. 1986. Thin-film recording media. *Proceedings of the IEEE* 74, 11 (1986), 1526–1539.

[7] Suayb S. Arslan, Mark A. Lantz, Simeon Furrer, Geoff Spratt, and Turguy Goker. 2022. LTO-9 Technology and User Data Reliability Analysis. (2022). Retrieved January 02, 2025 from https://www.lto.org/wp-content/uploads/2022/08/LTO-UBER-Technical-Paper-August-2022.pdf

[8] Ole Asmussen, Robert Beiderbeck, Albrecht Friess, Hans-Gunther Horhammer, Khanh Ngo, Jesus Eduardo Cervantes Rolon, Fabian Corona Villarreal, and Larry Coyne. 2018. IBM Tape Library Guide for Open Systems. (2018). Retrieved January 02, 2025 from https://www.redbooks.ibm.com/redbooks/pdfs/sg245946.pdf

[9] James A. Bain. 1996. Recording heads: Write heads for high-density magnetic tape. In *Proceedings of the High-Density Data Recording and Retrieval Technologies*. SPIE, 165–175.

[10] R. C. Barrett, E. H. Klaassen, T. R. Albrecht, G. A. Jaquette, and J. H. Eaton. 1998. Timing-based track-following servo for linear tape systems. *IEEE Transactions on Magnetics* 34, 4 (1998), 1872–1877.
[11] D. W. Barron, A. G. Fraser, D. F. Hartley, B. Landy, and R. M. Needham. 1967. File handling at Cambridge University. In *Proceedings of the April 18–20, 1967, Spring Joint Computer Conference (AFIPS'67 (Spring))*. Association for Computing Machinery, New York, NY, USA, 163–167. DOI: https://doi.org/10.1145/1465482.1465508
[12] C. Bashe, L. Johnson, J. Palmer, and E. Pugh. 1986. *IBM's Early Computers*. MIT Press, Cambridge, Massachusetts, USA.
[13] Geoffrey Bate. 1986. Particulate recording materials. *Proceedings of the IEEE* 74, 11 (1986), 1513–1525.
[14] Eric Baugh and Frank E. Talke. 1996. Head/tape interface. In *Proceedings of the High-Density Data Recording and Retrieval Technologies*. SPIE, 158–164.
[15] Matthew D. Baumgart and Lucy Y. Pao. 2004. Robust control of tape transport systems with no tension sensor. In *Proceedings of the 2004 43rd IEEE Conference on Decision and Control (CDC)(IEEE Cat. No. 04CH37601)*. IEEE, IEEE, Piscataway, New Jersey, USA, 4342–4349.
[16] Gordon C. Bell, Craig J. Mudge, and John E. McNamara. 1978. *COMPUTER ENGINEERING: A DEC View of Hardware Systems Design*. Digital Press, Beford, Massachusetts, USA.
[17] H. Neal Bertram. 1986. Fundamentals of the magnetic recording process. *Proceedings of the IEEE* 74, 11 (1986), 1494–1512.
[18] H. Neal Bertram. 1994. *Theory of Magnetic Recording*. Cambridge University Press, Cambridge, UK.
[19] Robert G. Biskeborn, W. S. Czarnecki, G. M. Decad, Robert E. Fontana, I. E. Iben, J. Liang, C. Lo, L. Randall, P. Rice, A. Ting, et al. 2013. Linear magnetic tape heads and contact recording. *ECS Transactions* 50, 10 (2013), 19.
[20] Robert G. Biskeborn and James H. Eaton. 2003. Hard-disk-drive technology flat heads for linear tape recording. *IBM Journal of Research and Development* 47, 4 (2003), 385–400.
[21] Robert G. Biskeborn, Robert E. Fontana, Calvin S. Lo, W. Stanley Czarnecki, Jason Liang, Icko E. T. Iben, Gary M. Decad, and Venus A. Hipolito. 2018. TMR tape drive for a 15 TB cartridge. *AIP Advances* 8, 5 (2018), 1–8.
[22] Robert G. Biskeborn, Pierre-Olivier Jubert, Jason Liang, and Calvin Lo. 2012. Head and interface for high areal density tape recording. *IEEE Transactions on Magnetics* 48, 11 (2012), 4463–4466.
[23] Richard Bradshaw and Carl Schroeder. 2003. Fifty years of IBM innovation with information storage on magnetic tape. *IBM Journal of Research and Development* 47, 4 (2003), 373–383.
[24] M. R. Brake and Jonathan A. Wickert. 2010. Lateral vibration and read/write head servo dynamics in magnetic tape transport. *Journal of Dynamic Systems, Measurement, and Control* 132, 1 (2010), 1–11.
[25] Shawn Brume. 2021. Virtual Analyst and Investor Day. (2021). Retrieved January 02, 2025 from https://community.ibm.com/community/user/storage/blogs/shawn-brume1/2021/11/03/tape-its-all-about-that-gap-no-trouble
[26] James Byron, Darrell D. E. Long, and Ethan L. Miller. 2018. Using simulation to design scalable and cost-efficient archival storage systems. In *Proceedings of the 2018 IEEE 26th International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS'18)*. IEEE Computer Society, Washington, DC, USA, 25–39. DOI: https://doi.org/10.1109/MASCOTS.2018.00011
[27] Luis-Felipe Cabrera, Robert Rees, Stefan Steiner, Wayne Hineman, and Michael Penner. 1995. ADSM: A multi-platform, scalable, backup and archive mass storage system. *COMPCON'95. Technologies for the Information Superhighway* (1995), 420–427. DOI: https://doi.org/10.1109/CMPCON.1995.512418
[28] Giovanni Cherubini, Roy D. Cideciyan, Laurent Dellmann, Evangelos Eleftheriou, Walter Haeberle, Jens Jelitto, Venkataraman Kartik, Mark A. Lantz, Sedat Ölçer, Angeliki Pantazi, et al. 2010. 29.5-Gb/in2 recording areal density on barium ferrite tape. *IEEE Transactions on Magnetics* 47, 1 (2010), 137–147.
[29] Giovanni Cherubini, Simeon Furrer, and Jens Jelitto. 2015. High-performance servo channel for nanometer head positioning and longitudinal position symbol detection in tape systems. *IEEE/ASME Transactions on Mechatronics* 21, 2 (2015), 1116–1128.
[30] Giovanni Cherubini, Angeliki Pantazi, and Jens Jelitto. 2013. Identification of MIMO transport systems in tape drives. In *Proceedings of the 2013 IEEE/ASME International Conference on Advanced Intelligent Mechatronics*. IEEE, 597–602.
[31] Giovanni Cherubini, Angeliki Pantazi, and Mark Lantz. 2016. Near-optimal tape transport control with feedback of velocity and tension. *IFAC-PapersOnLine* 49, 21 (2016), 19–25.
[32] Giovanni Cherubini, Angeliki Pantazi, and Mark A. Lantz. 2018. Feedback control of transport systems in tape drives without tension transducers. *Mechatronics* 49 (2018), 211–223.
[33] Ed Childer, Tom Coughlin, Ron Dennison, Simeon Furrer, Roger Hoyt, John Hodman, Dave Landsman, Mark Lantz, Kevin Lu, Niranjan Natekar, et al. 2023. IEEE international roadmap for devices and systems, "Mass Digital Storage." *Institute of Electrical and Electronics Engineers*. DOI: 10.60627/4HS9-2098
[34] Roy D. Cideciyan, Francois Dolivo, Reto Hermann, Walter Hirt, and Wolfgang Schott. 1992. A PRML system for digital magnetic recording. *IEEE Journal on Selected Areas in Communications* 10, 1 (1992), 38–56.

[35] Roy D. Cideciyan, Evangelos Eleftheriou, Brian H. Marcus, and Dharmendra S. Modha. 2001. Maximum transition run codes for generalized partial response channels. *IEEE Journal on Selected Areas in Communications* 19, 4 (2001), 619–634.
[36] Roy D. Cideciyan, Simeon Furrer, and Mark A. Lantz. 2017. Product codes for data storage on magnetic tape. *IEEE Transactions on Magnetics* 53, 2 (2017), 1–10. DOI: https://doi.org/10.1109/TMAG.2016.2614010
[37] Jonathan D. Coker, Evangelos Eleftheriou, Richard L. Galbraith, and Walter Hirt. 1998. Noise-predictive maximum likelihood (NPML) detection. *IEEE Transactions on Magnetics* 34, 1 (1998), 110–117.
[38] LTO Consortium. 2018. Ultrium LTO-8®. (2018). Retrieved January 02, 2025 from https://www.lto.org/lto-8/
[39] LTO Consortium. 2021. Ultrium LTO-9®. (2021). Retrieved January 02, 2025 from https://www.lto.org/lto-9/
[40] Digital Equipment Corporation. 1963. 555/550 Micro-Tape Dual Transport and Tape Control. Retrieved February 12, 2024 from https://archive.org/details/bitsavers_decdectape_2552119/mode/1up?view=theater
[41] Digital Equipment Corporation. 1964. Computing Timeline. Retrieved February 12, 2024 from https://www.vt100.net/timeline/1964.html
[42] Digital Equipment Corporation. 1978. TU58 DECtape II User's Guide. Retrieved February 12, 2024 from http://bitsavers.informatik.uni-stuttgart.de/pdf/dec/dectape/tu58/EK-0TU58-UG-001_TU58_DECtape_II_Users_Guide_Oct78.pdf
[43] Digital Equipment Corporation. 1985. TK50 Tape Drive Subsystem Technical Manual EK-OTK50-TM-001. (1985). Retrieved January 02, 2025 from http://www.bitsavers.org/pdf/dec/magtape/tk50/
[44] Digital Equipment Corporation. 1987. TK70 and TK50 Compac Tape Cartridge Subsystems, ED30721 45/12 02 30.0. (1987). Retrieved January 02, 2025 from http://www.bitsavers.org/pdf/dec/brochures/DEC-TK70+TK50-CompacTapeSubsystem.pdf
[45] R. C. Daley and P. G. Neumann. 1965. A general-purpose file system for secondary storage. In *Proceedings of the November 30–December 1, 1965, Fall Joint Computer Conference, Part I (AFIPS'65 (Fall, part I))*. Association for Computing Machinery, New York, NY, USA, 213–229. DOI: https://doi.org/10.1145/1463891.1463915
[46] Eric D. Daniel, C. Denis Mee, and Mark H. Clark. 1998. *Magnetic Recording: The First 100 Years*. John Wiley and Sons.
[47] F. K. Engel. 1999. The Introduction of the Magnetophon. *Ch. 5 in Magnetic Recording The First 100 Years*, Eric D. Daniel, C. Denis Mee, and Mark H. Clark (Eds.). IEEE Press. https://ieeexplore.ieee.org/servlet/opac?bknumber=5263537
[48] Grau Data. 2024. XtreemStore. (2024). Retrieved from https://www.graudata.com/en/products/tape-object-archive/xtreemstore
[49] Richard H. Dee. 1996. Read heads for magnetic tapes. In *Proceedings of the High-Density Data Recording and Retrieval Technologies*. SPIE, 181–191.
[50] Richard H. Dee. 1998. Magnetic tape recording technology and devices. In *Proceedings of the 7th Biennial IEEE International Nonvolatile Memory Technology Conference*. IEEE, 55–64.
[51] Richard H. Dee. 2002. The challenges of magnetic recording on tape for data storage (the one terabyte cartridge and beyond). In *Proceedings of the NASA CONFERENCE PUBLICATION*. 109–120.
[52] Richard H. Dee. 2008. Magnetic tape for data storage: An enduring technology. *Proceedings of the IEEE* 96, 11 (2008), 1775–1785.
[53] Don Lancaster. 1976. Serial interface. *BYTE* 1 (1976), 22.
[54] Pierre Dorion. 2008. Backup vs. Archive. (2008). Retrieved from https://www.einsteinmed.edu/uploadedFiles/Research/research-computing/Backup%20vs%20Archive-2.pdf
[55] A. L. Drapeau and R. H. Katz. 1993. Striping in large tape libraries. In *Proceedings of the ACM/IEEE International Conference on Supercomputing (SC'93)*. 378–387.
[56] James Eaton. 1996. Magnetic tape trends and futures. In *Proceedings of the High-Density Data Recording and Retrieval Technologies*. SPIE, 146–157.
[57] Patrick Ebermann, Giovanni Cherubini, Simeon Furrer, Mark A. Lantz, and Angeliki Pantazi. 2021. Track-following system optimization for future magnetic tape data storage. *Mechatronics* 80 (2021), 102662.
[58] Evangelos Eleftheriou, Sedat Olçer, and Robert A. Hutchins. 2010. Adaptive noise-predictive maximum-likelihood (NPML) data detection for magnetic tape storage systems. *IBM Journal of Research and Development* 54, 2 (2010), 7–1.
[59] Jessica Elliott. 2024. What Is the 3-2-1 Backup Rule? (2024). Retrieved January 02, 2025 from https://www.uschamber.com/co/run/technology/3-2-1-backup-rule
[60] Johan B. C. Engelen, Simeon Furrer, Hugo E. Rothuizen, and Mark A. Lantz. 2013. Flat-Profile tape–head friction and magnetic spacing. *IEEE Transactions on Magnetics* 50, 3 (2013), 34–39.
[61] Johan B. C. Engelen, V. Prasad Jonnalagadda, Simeon Furrer, Hugo E. Rothuizen, and Mark A. Lantz. 2016. Tape-head with sub-ambient air pressure cavities. *IEEE Transactions on Magnetics* 52, 11 (2016), 1–10.

[62] Johan B. C. Engelen and Mark A. Lantz. 2015. Asymmetrically wrapped flat-profile tape–head friction and spacing. *Tribology Letters* 59 (2015), 1–8.
[63] Hewlett Packard Enterprise. 2024. HPE LTO-9 Media Specification. (2024). Retrieved from https://buy.hpe.com/be/en/storage/storage-media/tape-media/lto-ultrium-cartridges/hpe-lto%E2%80%919-ultrium-45tb-rw-data-cartridge/p/q2079a
[64] G. Forney. 1972. Maximum-likelihood sequence estimation of digital sequences in the presence of intersymbol interference. *IEEE Transactions on Information Theory* 18, 3 (1972), 363–378.
[65] Free Software Foundation. 2014. Linux and Unix Man Pages - mt. Retrieved March 11, 2024 from https://www.unix.com/man-page/Linux/1/mt/
[66] Fujifilm. 2018. Barium Ferrite Technology in Data Storage. (2018). Retrieved January 02, 2025 from https://fujistore.hu/wp-content/uploads/2020/03/THE-ROLE-OF-BARIUM-FERRITE-TECHNOLOGY.pdf
[67] Fujifilm. 2024. Object Archive. (2024). Retrieved January 02, 2025 from https://asset.fujifilm.com/www/de/files/2023-12/a8827832046d899e4b200f14d529ad05/FUJIFILM-Object-Archive-Technical-White-Paper-v001-4.pdf
[68] Fujifilm. Year unknown. Barium Ferrite: Overview. (Year unknown). Retrieved January 02, 2025 from https://www.fujifilm.com/us/en/business/data-storage/fujifilm-technologies/barium-ferrite
[69] Simeon Furrer, Patrick Ebermann, Mark A. Lantz, Hugo Rothuizen, Walter Haeberle, Giovanni Cherubini, Roy D. Cideciyan, Shinji Tsujimoto, Yoshihiro Sawayashiki, Noriko Imaoka, et al.. 2021. 317 Gb/in^2 recording areal density on strontium ferrite tape. *IEEE Transactions on Magnetics* 57, 7 (2021), 1–11.
[70] Simeon Furrer, Pierre-Olivier Jubert, Giovanni Cherubini, Roy D. Cideciyan, and Mark A. Lantz. 2012. Analytical expressions for the readback signal of timing-based servo schemes. *IEEE Transactions on Magnetics* 48, 11 (2012), 4578–4581.
[71] Simeon Furrer, Mark A. Lantz, Johan B. C. Engelen, Angeliki Pantazi, Hugo E. Rothuizen, Roy D. Cideciyan, Giovanni Cherubini, Walter Haeberle, Jens Jelitto, Evangelos Eleftheriou, et al. 2015. 85.9 Gb/in^2 recording areal density on barium ferrite tape. *IEEE Transactions on Magnetics* 51, 4 (2015), 1–7.
[72] Simeon Furrer, Mark A. Lantz, Peter Reininger, Angeliki Pantazi, Hugo E. Rothuizen, Roy D. Cideciyan, Giovanni Cherubini, Walter Haeberle, Evangelos Eleftheriou, Junichi Tachibana, et al. 2017. 201 Gb/in^2 recording areal density on sputtered magnetic tape. *IEEE Transactions on Magnetics* 54, 2 (2017), 1–8.
[73] Simeon Furrer, Angeliki Pantazi, Giovanni Cherubini, and Mark A. Lantz. 2015. Resolution limits of timing-based servo schemes in magnetic tape drives. *IEEE Transactions on Magnetics* 51, 11 (2015), 1–4.
[74] Simeon Furrer, Angeliki Pantazi, Giovanni Cherubini, and Mark A. Lantz. 2018. Compressional wave disturbance suppression for nanoscale track-following on flexible tape media. In *Proceedings of the 2018 Annual American Control Conference (ACC'18)*. IEEE, 6678–6683.
[75] Amir Gandomi and Murtaza Haider. 2015. Beyond the hype: Big data concepts, methods, and analytics. *International Journal of Information Management* 35, 2 (2015), 137–144. DOI: https://doi.org/10.1016/j.ijinfomgt.2014.10.007
[76] Kurt Gerecke and Klemens Poschke. 2010. *IBM System Storage-Kompendium: Die IBM Speichergeschichte Von 1952 Bis 2010*. IBM.
[77] J. J. Gniewek. 1996. Evolving requirements for magnetic tape data storage systems. In *Proceedings of the 5th NASA Goddard Conference on Mass Storage Systems and Technologies (MSST'96)*. 477–491.
[78] L. Golubchik, R. R. Muntz, and R. W. Watson. 1995. Analysis of striping techniques in robotic storage libraries. In *Proceedings of the 14th IEEE Symposium on Mass Storage Systems (MASS'95)*. 225–238.
[79] Beverly Gooch. 1999. Building on the magnetophon. *Magnetic Recording: The First* 100 (1999), 72–91.
[80] Phil Goodwin. 2019. Storage Management in an Age of Minimal Data Deletion. (2019). Retrieved January 02, 2025 from https://asset.fujifilm.com/master/americas/files/2020-03/7bcc575d53cec43c1b68e78e94e7d895/Tape-and-Cloud-Solving-Storage-Problems.pdf
[81] Valéry Guilleaume. 2019. HSM (HIERARCHICAL STORAGE MANAGEMENT) VS ACTIVE ARCHIVE. (2019). Retrieved January 02, 2025 from https://www.nodeum.io/blog/hsm-hierarchical-storage-management-vs-active-archive
[82] Takeshi Harasawa, Ryota Suzuki, Osamu Shimizu, Sedat Olcer, and Evangelos Eleftheriou. 2010. Barium-ferrite particulate media for high-recording-density tape storage systems. *IEEE Transactions on Magnetics* 46, 6 (2010), 1894–1897.
[83] John P. Harris, William B. Phillips, Jack F. Wells, and Wayne D. Winger. 1981. Innovations in the design of magnetic tape subsystems. *IBM Journal of Research and Development* 25, 5 (1981), 691–700.
[84] Diana J. Hellman, Raymond Yardy, and Perry E. Abbott. 2003. Innovations in tape storage automation at IBM. *IBM Journal of Research and Development* 47, 4 (2003), 445–452.
[85] Hewlett Packard. 1998. Colorado Backup User's Guide. Retrieved March 10, 2024 from https://docs.rs-online.com/527f/0900766b8002abeb.pdf

[86] Hewlett Packard. 2024. LTO Cartridge Memory. Retrieved March 11, 2024 from https://support.hpe.com/hpesc/public/docDisplay?docId=sd00001231en_us&docLocale=en_US&page=GUID-860C3D40-27A5-4A64-9DEA-5A0852243BDD.html

[87] Michael Hoeck, Nik Simpson, Jerry Rozeman, and Jason Donham. 2024. Magic Quadrant for Enterprise Backup and Recovery Software Solutions. (2024). Retrieved February 15, 2024 from https://www.gartner.com/document/4605899

[88] Ella Hutchinson. 2024. Rediscovering Tape Storage: The Unconventional Innovation for Modern Data Challenges. (2024). Retrieved January 02, 2025 from https://www.intelligentdatacentres.com/2024/02/09/rediscovering-tape-storage-the-unconventional-innovation-for-modern-data-challenges

[89] IBM. 1961. Principles of Programming, Magnetic Tape Operations. Retrieved March 09, 2024 from https://ibm-1401.info/IBM-PrinProg-08.pdf

[90] IBM. 1963. Sort 7 Specifications and Operating Procedures. Retrieved March 09, 2024 from https://ibm-1401.info/pictures/C24-3317-1_sort7spec-3.pdf

[91] IBM. 1964. IBM 1401 System Summary. Retrieved March 09, 2024 from http://bitsavers.trailing-edge.com/pdf/ibm/1401/A24-1401-1_1401_System_Summary_Sep64.pdf

[92] IBM. 1965. IBM 1410/7010 Operating System. Retrieved March 09, 2024 from https://ia802908.us.archive.org/6/items/bitsavers_ibm1410C28p_1972928/C28-0318-3_1410_basicConcep_text.pdf

[93] IBM. 1966. IBM System/360 Operating System Data Management. Retrieved March 09, 2024 from http://bitsavers.trailing-edge.com/pdf/ibm/360/os/R01-08/C28-6537-1_Data_Management_Apr66.pdf

[94] IBM. 1973. OS Utilities. Retrieved March 11, 2024 from http://www.bitsavers.org/pdf/ibm/360/os/R21.7_Apr73/GC28-6586-15_OS_Utilities_Rel_21.7_Apr73.pdf

[95] IBM. 1987. IBM 3580 Magnetic Tape Subsystem Planning and Migration Guide GC35-0098-5. (1987). Retrieved March 11, 2024 from http://www.bitsavers.org/pdf/ibm/3480/

[96] IBM. 1988. z/OS 3.1 MVS JCL Reference. Retrieved March 11, 2024 from https://www.ibm.com/docs/en/SSLTBW_3.1.0/pdf/ieab600_v3r1.pdf

[97] IBM. 2004. CMS Commands and Utilities Reference. Retrieved March 11, 2024 from https://publibz.boulder.ibm.com/epubs/pdf/hcsd8b00.pdf

[98] IBM. 2018. CMS Application Development Guide for Assembler. Retrieved March 11, 2024 from https://www.vm.ibm.com/library/710pdfs/71625700.pdf

[99] IBM. 2021. IBM LTO 9 Tape Drive Data Sheet. (2021). Retrieved January 02, 2025 from https://www.ibm.com/downloads/cas/4DYRWDGB

[100] IBM. 2021. IBM Tape Device Drivers. Retrieved March 11, 2024 from https://www.ibm.com/docs/en/ts4300-tape-library?topic=tape-device-drivers-diagnostic-tool-users-guide

[101] IBM. 2021. Linux Tape and Medium Changer Device Driver. Retrieved March 11, 2024 from https://www.ibm.com/docs/en/ts4300-tape-library?topic=guide-linux-tape-medium-changer-device-driver

[102] IBM. 2023. IBM TS1170 Tape Drive. (2023). Retrieved January 02, 2025 from https://www.ibm.com/downloads/cas/WQX5N4RJ

[103] IBM. 2023. IBM TS2290 Tape Drive. (2023). Retrieved from https://www.ibm.com/downloads/cas/9WDD3Q1L

[104] IBM. 2023. What is CMS? Retrieved March 11, 2024 from https://www.ibm.com/docs/en/zvm/7.3?topic=environment-what-is-cms

[105] IBM. 2024. Active Archive. (2024). Retrieved January 02, 2025 from https://www.ibm.com/docs/en/storage-archive-ee/1.3.4?topic=edition-active-archive

[106] IBM. 2024. IBM S3 Deep Archive. (2024). Retrieved January 02, 2025 from https://www.ibm.com/products/s3-deep-archive

[107] IBM. 2024. IBM z/OS. Retrieved March 09, 2024 from https://www.ibm.com/products/zos

[108] IBM. 2024. IBM z/VM. Retrieved March 11, 2024 from https://www.ibm.com/products/zvm

[109] IBM. 2024. The IBM PC. Retrieved March 10, 2024 from https://www.ibm.com/history/personal-computer

[110] IBM. Year unknown. IBM Storage. (Year unknown). Retrieved January 02, 2025 from https://www.ibm.com/history/exhibits/storage/storage_3420

[111] I. Iliadis, L. Jordan, M. Lantz, and S. Sarafijanovic. 2021. Performance evaluation of automated tape library systems. In *Proceedings of the 29th IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS'21)*. 1–8. DOI: https://doi.org/10.1109/MASCOTS53633.2021.9614293

[112] Ilias Iliadis, Linus Jordan, Mark Lantz, and Slavisa Sarafijanovic. 2022. Performance evaluation of tape library systems. *Performance Evaluation* 157-158, Article 102312 (2022), 22 pages. DOI: https://doi.org/10.1016/j.peva.2022.102312

[113] I. Iliadis, Y. Kim, S. Sarafijanovic, and V. Venkatesan. 2016. Performance evaluation of a tape library system. In *Proceedings of the 24th IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS'16)*. 59–68. DOI: https://doi.org/10.1109/MASCOTS.2016.37

[114] K. A. Schouhamer Immink. 1990. Runlength-limited sequences. *Proceedings of the IEEE* 78, 11 (1990), 1745–1759.
[115] INSIC. 2019. 2019 INSIC Technology Roadmap. (2019). Retrieved January 02, 2025 from https://www.insic.org/wp-content/uploads/2019/07/INSIC-Technology-Roadmap-2019.pdf
[116] Information Storage Industry Consortium (INSIC). 2019. 2019 Applications and Systems Tape Roadmap. (2019), 23, 49–59 pages. Retrieved January 02, 2025 from https://www.insic.org/wp-content/uploads/2019/07/INSIC-Applications-and-Systems-Roadmap.pdf
[117] Information Storage Industry Consortium (INSIC). 2024. INSIC International Magnetic Tape Storage Technology Roadmap 2024. (2024), 27 pages. Retrieved January 02, 2025 from https://insic.org/roadmap/
[118] ISO. 2016. Linear Tape File System (LTFS) Format Specification. (2016). Retrieved January 02, 2025 from https://www.iso.org/standard/80598.html
[119] Fred Jeffers. 1986. High-density magnetic recording heads. *Proceedings of the IEEE* 74, 11 (1986), 1540–1556.
[120] Brad Johns. 2020. Reducing Data Center Energy Consumption and Carbon Emissions with Modern Tape Storage. (2020). Retrieved January 02, 2025 from https://www.bradjohnsconsulting.com/_files/ugd/8b8555_8c8acce1704045bc8002a63832874ce1.pdf
[121] Clayton Johnson. 1975. IBM 3850: Mass storage system. In *Proceedings of the May 19–22, 1975, National Computer Conference and Exposition*. 509–514.
[122] T. Johnson. 1996. An analytical performance model of robotic storage libraries. *Performance Evaluation* 27-28 (1996), 231–251.
[123] T. Johnson and E. L. Miller. 1998. Performance measurements of tertiary storage devices. In *Proceedings of the 24th International Conference on Very Large Data Bases (VLDB'98)*. 50–61.
[124] Jan Jose, Ryan J. Taylor, Raymond A. De Callafon, and Frank E. Talke. 2005. Characterization of lateral tape motion and disturbances in the servo position error signal of a linear tape drive. *Tribology International* 38, 6-7 (2005), 625–632.
[125] Pierre-Olivier Jubert. 2013. Achieving 100 Gb/in^2 on particulate barium ferrite tape. *IEEE Transactions on Magnetics* 50, 1 (2013), 1–8.
[126] Pierre-Olivier Jubert, Yuri Obukhov, Cristian Papusoi, and Paul Dorsey. 2021. Evaluation of sputtered tape media with hard disk drive components. *IEEE Transactions on Magnetics* 58, 4 (2021), 1–5.
[127] Olle Karlqvist. 1954. Calculation of the magnetic field in the ferromagnetic layer of a magnetic drum. Elanders boktr. (1954).
[128] Harry Katzan. 1971. Storage hierarchy systems. In *Proceedings of the May 18–20, 1971, Spring Joint Computer Conference (AFIPS'71 (Spring))*. Association for Computing Machinery, New York, NY, USA, 325–336. DOI: https://doi.org/10.1145/1478786.1478833
[129] Debra Kaufman. 2014. Media Archiving at the Library of Congress. (2014). Retrieved January 02, 2025 from https://www.smpte.org/blog/media-archiving-library-congress
[130] Takahiro Kawana, Seiichi Onodera, and Tetsuo Samoto. 1995. Advanced metal evaporated tape. *IEEE Transactions on Magnetics* 31, 6 (1995), 2865–2870.
[131] Gregory T. Kishi. 2003. The IBM virtual tape server: Making tape controllers more autonomic. *IBM Journal of Research and Development* 47, 4 (2003), 459–469.
[132] H. Kobayashi and D. T. Tang. 1970. Application of partial-response channel coding to magnetic recording systems. *IBM Journal of Research and Development* 14, 4 (1970), 368–375.
[133] Mark H. Kryder, Edward C. Gage, Terry W. McDaniel, William A. Challener, Robert E. Rottmayer, Ganping Ju, Yiao-Tee Hsia, and M. Fatih Erden. 2008. Heat assisted magnetic recording. *Proceedings of the IEEE* 96, 11 (2008), 1810–1835.
[134] Mark A. Lantz. 2015. Fujifilm's 7th Annual Global IT Executive Summit. (2015). Retrieved January 02, 2025 from https://www.fujifilmsummit.com/wp-content/uploads/2017/04/2015-lantz.pdf
[135] Mark A. Lantz, Giovanni Cherubini, Angeliki Pantazi, and Jens Jelitto. 2011. Servo-pattern design and track-following control for nanometer head positioning on flexible tape media. *IEEE Transactions on Control Systems Technology* 20, 2 (2011), 369–381.
[136] Mark A. Lantz and Evangelos Elefteriou. 2014. Future scaling potential of particulate media in magnetic tape recording. In *Proceedings of the Handbook of Magnetic Materials*. Elsevier, 317–379.
[137] Mark A. Lantz, Simeon Furrer, Johan B. C. Engelen, Angeliki Pantazi, Hugo E. Rothuizen, Roy D. Cideciyan, Giovanni Cherubini, Walter Haeberle, Jens Jelitto, Evangelos Eleftheriou, et al.. 2015. 123 Gbit/in^2 recording areal density on barium ferrite tape. *IEEE Transactions on Magnetics* 51, 11 (2015), 1–4.
[138] S. S. Lavenberg and D. R. Slutz. 1975. Regenerative simulation of a queuing model of an automated tape library. *IBM Journal of Research and Development* 19, 5 (1975), 463–475.
[139] Brahim Lekmine. 1996. Recording channel and data detection in magnetic tape drives. In *High-Density Data Recording and Retrieval Technologies*. SPIE, 176–180.

[140] Spectra Logic. 2024. Spectra On-Prem Glacier. (2024). Retrieved January 02, 2025 from https://spectralogic.com/solutions/on-prem-glacier-solutions
[141] LTO Consortium. 2023. The Benefits of LTO Ultrium Tape Technology. Retrieved March 11, 2024 from https://www.lto.org/benefits-of-lto/
[142] Stefan Maat and Arley C. Marley. 2016. Physics and design of hard disk drive magnetic recording read heads. *Handbook of Spintronics* (2016), 977–1028.
[143] John C. Mallinson. 2012. *The Foundations of Magnetic Recording*. Elsevier.
[144] Manfred Peschke and Virginia Peschke. 1976. Report: BYTE's audio cassette standards symposium. *BYTE* 0, 6 (1976), 72–73.
[145] Mathew Reed. 2007. The Exatron Stringy Floppy. Retrieved March 10, 2024 from http://www.trs-80.org/exatron-stringy-floppy/
[146] Gary M. McClelland, David Berman, Pierre-Olivier Jubert, Wayne Imaino, Hitoshi Noguchi, Masahiko Asai, and Hiroaki Takano. 2009. Effect of tape longitudinal dynamics on timing recovery and channel performance. *IEEE Transactions on Magnetics* 45, 10 (2009), 3587–3589.
[147] Microsoft. 2021. Windows Introduction to Tape Drivers. Retrieved March 11, 2024 from https://learn.microsoft.com/en-us/windows-hardware/drivers/storage/tape-drivers-overview
[148] Catherine D. Miller. 1989. Backup on a personal scale. *PC Magazine* 8, 22 (1989), 190–239.
[149] John Monroe. 2023. Storage Management in an Age of Minimal Data Deletion. (2023). Retrieved January 02, 2025 from https://www.lto.org/wp-content/uploads/2023/07/Storage-Management-in-an-Age-of-Minimal-Data-Deletion_Further-Research.pdf
[150] Jaekyun Moon and Barrett Brickner. 1996. Maximum transition run codes for data storage systems. *IEEE Transactions on Magnetics* 32, 5 (1996), 3992–3994.
[151] Multicians. 2024. Multics Glossary - T. Retrieved March 10, 2024 from https://multicians.org/mgt.html
[152] Multicians. 2024. Multics History. Retrieved March 09, 2024 from https://multicians.org/history.html
[153] Innovation and Employment New Zealand Government Ministry of Business. 2024. Continuity and Contingency Planning. (2024). Retrieved January 02, 2025 from https://www.business.govt.nz/risks-and-operations/planning-for-the-unexpected-bcp/continuity-and-contingency-planning
[154] H. Noguchi, M. Oyanagi, H. Doshita, and Mohamad Ramadan. 2022. Magnetic-particulate recording media: Advanced. In *Encyclopedia of Smart Materials*, Abdul-Ghani Olabi (Ed.). DOI: 10.1016/B978-0-12-815732-9.00135-2
[155] Information Processing Society of Japan Computer Museum. Year unknown. Magnetic Tape Units. (Year unknown). Retrieved January 02, 2025 from http://museum.ipsj.or.jp/en/computer/device/magnetic_tape/index.html
[156] Shin-ichi Ohkoshi, Asuka Namai, Kenta Imoto, Marie Yoshikiyo, Waka Tarora, Kosuke Nakagawa, Masaya Komine, Yasuto Miyamoto, Tomomichi Nasu, Syunsuke Oka, et al. 2015. Nanometer-size hard magnetic ferrite exhibiting high optical-transparency and nonlinear optical-magnetoelectric effect. *Scientific Reports* 5, 1 (2015), 14414.
[157] Haruo Okuda. 2010. The dawn of video tape recording and development of the helical scanning system. In *Proceedings of the 2010 2d Region 8 IEEE Conference on the History of Communications*. IEEE, 1–6.
[158] Eiki Ozawa. 2019. Microwave-assisted magnetization reversal in dispersed nanosized barium ferrite particles for high-density magnetic recording tape. *IEEE Transactions on Magnetics* 55, 7 (2019), 1–4.
[159] Shiba P. Panda. 2003. Application of track following tape drive control. In *Proceedings of the 2003 American Control Conference, 2003.* IEEE, 20–24.
[160] Angeliki Pantazi, Giovanni Cherubini, and Jens Jelitto. 2013. Skew estimation and feed-forward control in flangeless tape drives. *IFAC Proceedings Volumes* 46, 5 (2013), 484–489.
[161] Angeliki Pantazi, Giovanni Cherubini, Eiji Ogura, and Jens Jelitto. 2014. Tape transport control based on sensor fusion. *IFAC Proceedings Volumes* 47, 3 (2014), 6849–6855.
[162] Angeliki Pantazi, Simeon Furrer, Hugo E. Rothuizen, Giovanni Cherubini, Jens Jelitto, and Mark A. Lantz. 2015. Nanoscale track-following for tape storage. In *Proceedings of the 2015 American Control Conference (ACC'15)*. IEEE, 2837–2843.
[163] Angeliki Pantazi, Jens Jelitto, Nhan Bui, and Evangelos Eleftheriou. 2012. Track-following in tape storage: Lateral tape motion and control. *Mechatronics* 22, 3 (2012), 361–367.
[164] David Pease, Arnon Amir, Lucas Villa Real, Brian Biskeborn, Michael Richmond, and Atsushi Abe. 2010. The linear tape file system. In *Proceedings of the 2010 IEEE 26th Symposium on Mass Storage Systems and Technologies (MSST'10)*. 1–8. DOI: https://doi.org/10.1109/MSST.2010.5496989
[165] William B. Phillips. 1998. *Magnetic Recording: The First 100 Years, CH. 17 Data Storage on Tape*. IEEE Press.
[166] Stephen Pritchard. 2022. Storage Requirements for AI, ML and Analytics in 2022. (2022). Retrieved January 02, 2025 from https://www.computerweekly.com/feature/Storage-requirements-for-AI-ML-and-analytics-in-2022
[167] E. Pugh, L. Johnson, and J. Palmer. 1991. *IBM's 360 and Early 370 Systems*. MIT Press.

[168] Qstar. 2024. QStar Archive Storage Manager. (2024). Retrieved January 02, 2025 from https://www.qstar.com/archive-manager
[169] Quantum. 2020. Virtual Analyst and Investor Day. (2020), 20 pages. Retrieved January 02, 2025 from https://static.seekingalpha.com/uploads/sa_presentations/178/74178/original.pdf
[170] Quantum. 2021. Defending Data Against Ransomware. (2021). Retrieved January 02, 2025 from https://quantum.drift.click/Defending-Data-Against-Ransomware
[171] Steve Ranger. 2021. What is the IoT? Everything you Need to know about the Internet of Things Right Now. (2021). Retrieved January 02, 2025 from https://www.zdnet.com/article/what-is-the-internet-of-things-everything-you-need-to-know-about-the-iot-right-now
[172] Peter Reininger, Johan B. C. Engelen, Walter Häberle, and Mark A. Lantz. 2017. A model for head/tape friction for smooth media. *Tribology Letters* 65, 2 (2017), 65.
[173] ReportLinker. 2023. Tape Storage Global Market Report 2023. (2023). Retrieved January 02, 2025 from https://finance.yahoo.com/news/tape-storage-global-market-report-131600055.html
[174] E. Rismani, S. K. Sinha, S. Tripathy, H. Yang, and C. S. Bhatia. 2011. Effect of pre-treatment of the substrate surface by energetic C+ ion bombardment on structure and nano-tribological characteristics of ultra-thin tetrahedral amorphous carbon (ta-C) protective coatings. *Journal of Physics D: Applied Physics* 44, 11 (2011), 115502.
[175] Dennis M. Ritchie and Ken Thompson. 1974. The UNIX Time-sharing system. *Communications of the ACM* 17, 7 (1974), 365–375. DOI: https://doi.org/10.1145/361011.361061
[176] George A. Saliba, Satya A. Mallick, Chan Kim, Carol Turgeon, Leo Cappabianca, and Lewis Cronis. 2006. Multichannel Magnetic Tape System having Optical Tracking Servo. (Sept. 19 2006). US Patent 7,110,210.
[177] Richard C. Schneider. 1996. Design methodology for high-density read equalization. In *Proceedings of the High-Density Data Recording and Retrieval Technologies*. SPIE, 200–209.
[178] Seagate. 2023. Ironwolf Pro Sata Product Manual, Publication Number: 204482900. Retrieved March 20, 2024 from https://www.seagate.com/content/dam/seagate/migrated-assets/www-content/product-content/ironwolf/en-us/docs/204482900b.pdf
[179] Osamu Shimizu, Yuichi Kurihashi, Isamu Watanabe, and Takeshi Harasawa. 2013. Distribution of thermal stability factor for barium ferrite particles. *IEEE Transactions on Magnetics* 49, 7 (2013), 3767–3770.
[180] Gregory L. Silvus and Bhagavatula Vijaya Kumar. 1996. Nonlinear signal model for magnetic-tape recording channels utilizing magneto-resistive heads. In *Proceedings of the High-Density Data Recording and Retrieval Technologies*. SPIE, 192–199.
[181] Boris Slutsky and H. Neal Bertram. 1994. Transition noise analysis of thin film magnetic recording media. *IEEE Transactions on Magnetics* 30, 5 (1994), 2808–2817.
[182] SNIA. 2016. LTFS Bulk Transfer. (2016). Retrieved January 02, 2025 from https://www.snia.org/sites/default/files/technical-work/ltfs/release/SNIA-LTFS-Bulk-Transfer-1-0.pdf
[183] SNIA. 2020. Linear Tape File System (LTFS) Format Specification Version 2.5.1. (2020). Retrieved January 02, 2025 from https://www.snia.org/sites/default/files/technical-work/ltfs/release/SNIA-LTFS-Format-2-5-1-Standard.pdf
[184] Veeam Software. 2023. Data Protection Trends Report 2023. (2023), 5 pages. Retrieved January 02, 2025 from https://www.veeam.com/wp-data-protection-trends-report-2023.html
[185] Stack Exchange. 2020. Is it Possible to Install the GNU mt Tape Drive Command in OSX? Retrieved March 11, 2024 from https://apple.stackexchange.com/questions/380390/is-it-possible-to-install-the-gnu-mt-tape-drive-command-in-osx
[186] Statista. 2023. Share of Corporate Data Stored in the Cloud in Organizations Worldwide from 2015 to 2022. (2023). Retrieved January 02, 2025 from https://www.statista.com/statistics/1062879/worldwide-cloud-storage-of-corporate-data
[187] Horison Information Strategies. 2021. Understanding the Data Lifecycle. (2021). Retrieved January 02, 2025 from https://www.lto.org/wp-content/uploads/2021/02/WP-II-Understanding-the-Data-Lifecyle.pdf
[188] PoINT Software and Systems. 2024. PoINT Archival Gateway. (2024). Retrieved January 02, 2025 from https://www.point.de/en/products/point-archival-gateway
[189] Overland Tandberg. Year unknown. LTO Media. (Year unknown). Retrieved from https://www.overlandtandberg.com/products/neo-tape/lto-media/
[190] Data Tape. 2020. The Technology and Innovation That Keeps Magnetic Tape Alive. (2020). Retrieved January 02, 2025 from https://webuyusedtape.net/2020/10/21/the-technology-and-innovation-that-keeps-magnetic-tape-alive
[191] Aleksander Markovich Taratorin. 2004. *Magnetic Recording Systems and Measurements*. Guzik Technical Enterprises.
[192] Petroc Taylor. 2023. Total Installed Base of Data Storage Capacity in Global Datasphere 2020-2025. (2023). Retrieved from https://www.statista.com/statistics/1185900/worldwide-datasphere-storage-capacity-installed-base/
[193] Gaspare Varvaro and Francesca Casoli. 2016. *Ultra-high-density Magnetic Recording: Storage Materials and Media Designs*. CRC Press.

[194] Versity. 2024. Versity Gateway. (2024). Retrieved January 02, 2025 from https://www.versity.com/products-3/versitygw
[195] Versity. 2024. Versity ScoutAM. (2024). Retrieved January 02, 2025 from https://www.versity.com/products-3/scoutam
[196] Alexis Määttä Vinkler and Sandberg Patrik. 2010. Longterm Storage of Digital Photographs. (2010). Retrieved January 02, 2025 from https://www.csc.kth.se/utbildning/kandidatexjobb/medieteknik/2010/rapport/maatta_vinkler_alexis_OCH_sandberg_patrik_K10078.pdf
[197] David C. Walden and Tom Van Vleck. 2011. *The Compatible Time Sharing System (1961-1973): Fiftieth Anniversary Commemorative Overview*. IEEE Computer Society.
[198] Shan X. Wang and Alex M. Taratorin. 1999. *Magnetic Information Storage Technology: A Volume in the Electromagnetism Series*. Academic Press, San Diego, CA.
[199] Ian Warhaftig and Bruce Polsky. 1987. Maintaining holds on niche storage markets. *Computerworld* 21, 14 (1987), 40.
[200] Dieter Weller and Andreas Moser. 1999. Thermal effect limits in ultrahigh-density magnetic recording. *IEEE Transactions on Magnetics* 35, 6 (1999), 4423–4439.
[201] H. F. Welsh and H. Lukolf. 1952. The uniservo-tape reader and recorder. In *Proceedings of the International Workshop on Managing Requirements Knowledge*. IEEE Computer Society, 47–47.
[202] M. L. Williams and R. L. Comstock. 1972. An analytical model of the write process in digital magnetic recording. In *Proceedings of the AIP Conference*. American Institute of Physics, 738–742.
[203] Roger Wood. 1986. Magnetic recording systems. *Proceedings of the IEEE* 74, 11 (1986), 1557–1573.
[204] Velvet Wu. 2023. Tape Storage Might Be Computing's Climate Savior. (2023). Retrieved January 02, 2025 from https://spectrum.ieee.org/tape-storage-sustainable-option
[205] XenData. 2024. XenData LTO Active Archives. (2024). Retrieved January 02, 2025 from https://xendata.com/lto-archives
[206] Hankang Yang and Sinan Müftü. 2014. Coupling between the in-plane and lateral tape dynamics in high capacity linear tape transport systems. *IFAC Proceedings Volumes* 47, 3 (2014), 5914–5920.
[207] David Yu, Guangwei Che, Tim Chou, and Ognian Novakov. 2019. Best practices in accessing tape-resident data in HPSS. In *Proceedings of the EPJ Web of Conferences*. EDP Sciences, 04022.
[208] Zhang. 2016. Access Azure Blob Storage from Your Apps using S3 Java API. (2016). Retrieved January 02, 2025 from https://devblogs.microsoft.com/ise/access-azure-blob-storage-from-your-apps-using-s3-api
[209] Jian-Gang Zhu, Xiaochun Zhu, and Yuhui Tang. 2007. Microwave assisted magnetic recording. *IEEE Transactions on Magnetics* 44, 1 (2007), 125–131.

Received 15 November 2023; revised 15 July 2024; accepted 13 August 2024