INSIC
Information Storage
Industry Consortium

INSIC International Magnetic Tape Storage Technology Roadmap 2024

WHITE PAPER | MARCH 2024

insic.org

Participants

Leadership team

Bob Biskeborn	Western Digital
Mark Lantz	IBM
Simeon Furrer	IBM
Geoff Spratt	HPE
Eiki Ozawa	Fujifilm
Eiji Nakashio	Sony
Turguy Goker	Quantum
Robert Raymond	Oracle
Mark Hill	IBM
Matt Badger	HPE

2.1

INTRODUCTION

This section discusses the Tape Technology Roadmap as well as the technologies needed to execute the ten-year Roadmap goals. The approach is to determine what goals are needed for tape to remain a competitive storage technology in the markets described in the preceding Applications and Systems chapter by continuing to reduce the cost per GB, i.e. to maintain tape's cost advantage over HDD. It is important to understand that this Roadmap does not represent any specific product roadmap and may not be realized in a product. It is a roadmap of what is believed to be technically possible and does not explicitly consider business conditions and funding constraints that might affect a product roadmap. It is recommended to talk to your tape product supplier to get the latest product roadmaps.

The cost per GB reduction of HDD has been driven in large part by areal density growth. Between 2003 and 2009, areal density growth for HDD was about 39% per year, however, more recently, the scaling rate has decreased to about 2.6% per year, as shown in Figure 1. This slow-down in areal density scaling has been partially compensated for by an increase in the number of platters and heads in an HDD. While this strategy initially enabled continued capacity and cost scaling, each platter and head adds cost and the strategy has become much less efficient as the cost of platters and heads has begun to dominate the cost of the disk.

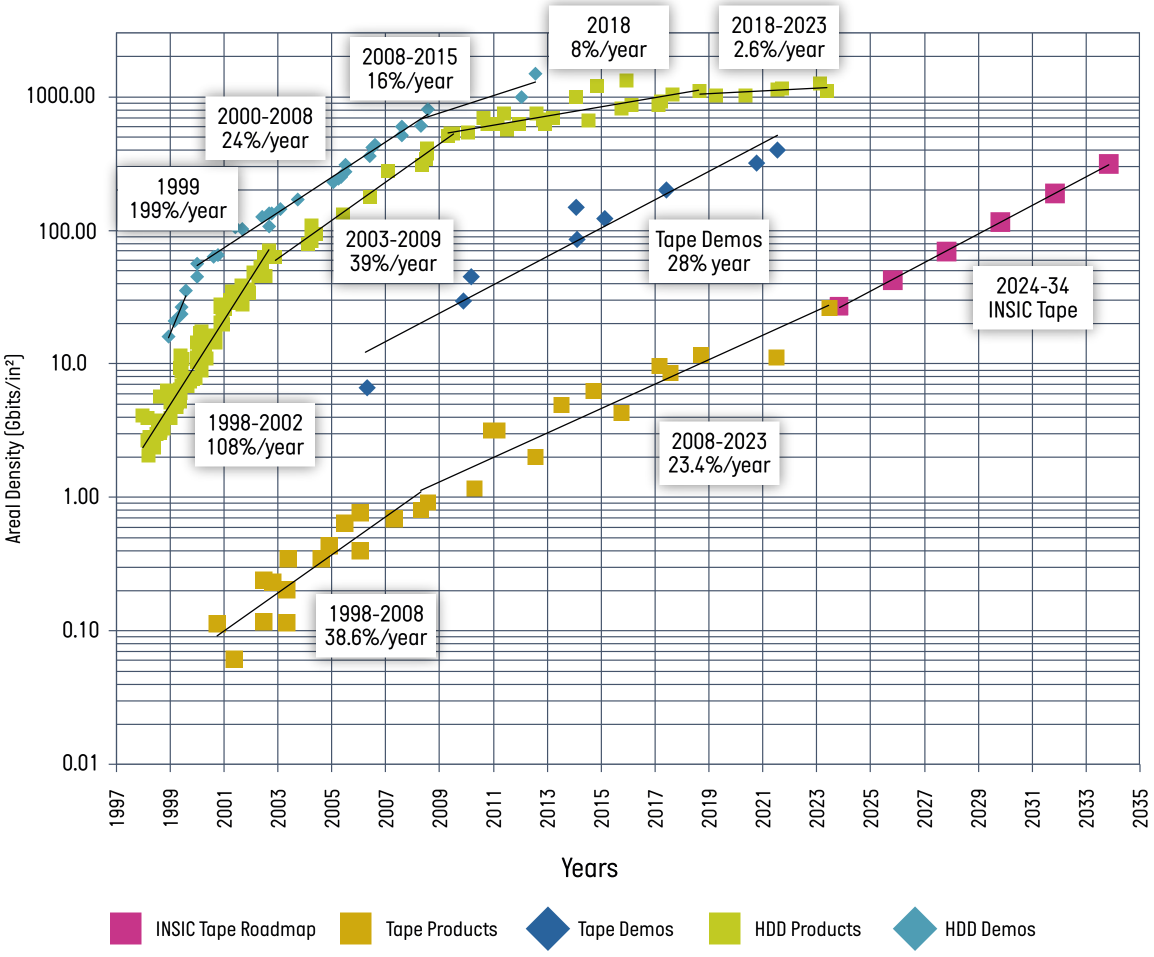

Figure 1: Areal Density Trends. Hard Disk Drive, Tape and Tape Technology Roadmap

As a result, the current rates of HDD cost scaling are much lower than historical rates. There is considerable uncertainty regarding future scaling rates of HDD due to the challenges associated with overcoming the super-paramagnetic effect and the scaling potential of the new technologies needed to continue HDD scaling. For example, commercialization of HAMR (heat assisted magnetic recording) HDDs have been repeatedly delayed. In contrast, state-of-the-art tape drives operate at areal densities that are about fifty times smaller than the latest HDDs. It should therefore be possible to continue scaling tape technology at the 28% per year rate projected by this Roadmap for at least the next decade and beyond, before tape begins to face similar challenges as HDD. Capacities of modern tape cartridges rival, and can exceed, capacities of current generation 3.5" HDD's, despite tape having roughly 50x lower areal densities. This is because a tape cartridge can hold roughly 100x more recording surface area than a 10-platter HDD. And given that the two have roughly similar volumes, tape remains competitive with HDD in terms of TB/unit volume at this "brick" level. This, and many other factors, together help tape achieve its industry-leading total cost of ownership (TCO). In fact, current state-of-the-art tape drives provide a native capacity of 50 TB, which is almost 2x that of the highest capacity HDDs currently available.

As a starting point for the 2024 INSIC Tape Roadmap, we look at the average rate of tape areal density scaling from 2000 to 2023. Over this period, the average scaling was around 29% per year. However, as can be seen in Figure 1, the initial rate of scaling was above this average at around 39% per year and more recently has slowed to around 23% a year. Some of this recent easing was aligned with global economic conditions and workplace conditions. In addition, tape technology recently entered a regime where tape dimension stability (TDS) must be actively compensated to enable continued track density scaling. Development of this technology contributed to an increase in recent development cycles. Current products use a tension-based TDS compensation scheme which is sufficient for current areal densities but has limited compensation range and scalability. The tape industry is currently in the process of implementing a new active TDS compensation scheme with more range and scalability that will enable multiple future generations of tape technology. Taking all of this into account, the current Roadmap projects a 28% per year growth in areal density and a 32% per year growth in capacity. This is slower than the projections of previous roadmaps and corresponds to a doubling in capacity every 2.5 years which is better aligned with the average time between LTO product releases which target doubling capacity. By aligning with these trends, the Roadmap indirectly considers business conditions.

As discussed in the following sections, the projected growth rate is expected to be challenging but feasible with no 'showstoppers' on the horizon. Compared to the 2019 INSIC Tape Roadmap, the track density scaling factor has been reduced from 24.3% to 21.7% and the linear density scaling factor has been reduced from 8% to 5%. The scaling factor for media thickness was reduced from -4% per year in the 2019 Roadmap to -2.5% per year based on input from tape substrate and media manufacturers. The last significant change relative to the 2019 Roadmap is the bit error rate which is set to a constant value, as discussed in Sect. 2.3.4 below. The growth rate for data rate is set to 15% per year which is the same as the 2019 Roadmap.

The Roadmap is summarized in Table 1 below, with the projected capacity and data rate improvements shown, along with the necessary progress required in key technical parameters to achieve these growth rates. The Roadmap requires an areal density in the year 2034 that is about where HDD technology was in 2009 and so has been shown to be feasible from a magnetic recording technology viewpoint. Recent laboratory tape areal density demonstrations also provide evidence of the potential to continue scaling tape technology to areal densities of up to 317 Gb/in^2 [Furrer 2021]. Therefore, from a fundamental areal density perspective, we feel the 2034 areal density target of about 315 Gb/in^2 is achievable. This Roadmap is not a product roadmap but rather a technology roadmap that might represent an average of possible products. As such, there may be no specific products shipping or planned at these exact numbers.

Deeper Dive Into Areal Density Chart

Figures 2 and 3 below show Areal density and recorded bit dimensions illustrating how tape compares to disk (HDD). The amount of data that can be stored in a specific area of a magnetic storage medium is referred to as areal density. This density is commonly measured in bits per square inch (bpsi) or giga bits per square inch (Gb/in^2). It is a representation of the density of magnetic data bits that may be packed onto a magnetic tape or hard disk drive (HDD) surface.

Track pitch and bit length are the two parameters that define the recorded bit, and its area (Track Pitch x Bit Length) is also used to determine Areal Density. The recorded bit can be best described as a rectangular shape. Figure 2 illustrates how magnetic bit shapes chosen for specific Areal Densities are displayed for Tape and HDD products, including Tape Demos and INSIC 2024 Roadmap estimates.

Figure 2 illustrates that, for an 18 TB HDD, 1019 Gb/in^2 Areal Density is required, whereas 18 TB LTO-9 only needs 12 Gb/in^2. This means that, given the same capacities for both technologies, HDD requires an Areal Density that is 85x higher. Similarly, in terms of recorded magnetic bit dimensions, Figure 2 shows that latest 26 TB HDD from WDC with bits dimension is 11.91x40.64 nm, whereas LTO-9 at 18 TB capacity has bits have substantially larger dimensions, 46.6x1156 nm. Basically, when HDD's Areal density increases to make its capacities higher, it may have electromechanical challenges controlling incredibly small bits with high volumes, low cost, and high reliability. In contrast, tape has a lot more space to work with in the Areal Density chart as shown in Figure 2.With this potential and the ability to expand the recordable tape area, tape technology has a promising future.

PARAMETER/Year	**2024**	**2026**	**2028**	**2030**	**2032**	**2034**	
1. Capacity (TB)	45	78	137	238	415	723	32% per year
2. Maximum data rate per channel (MB/sec)	12.5	8.3	10.9	14.5	19.1	12.6	
3. Maximum streaming drive data rate (MB/s)	400.0	529.0	699.6	925.2	1223.6	1618.2	15% per year
4. Minimum streaming drive data rate (MB/s)	90.7	200.0	220.5	243.1	268.0	591.0	1.2 m/s min
5. FC Speed Roadmap (GB/sec*)	128	256	256	512	512	1024	
6. Number of channels	32	64	64	64	64	128	
7. Tape speed (m/sec)	5.3	3.2	3.8	4.6	5.5	3.3	
8. Tape thickness	5.00	4.75	4.52	4.30	4.08	3.88	-2.5% per year
9. Data capacity reserve	3.0%	3.0%	3.0%	3.0%	3.0%	3.0%	
10. Tape length that is recordable (meters)**	1036	1092	1151	1213	1278	1347	40 meters reserve
11. Tape length total (meters)***	1076	1132	1191	1253	1318	1387	2.56% per year
12. Track density (TPI)	45,296	63,677	95,494	143,222	214,824	322,248	21.68% per year
Track pitch = 2.54x10^7/tpi (nm)	561	399	266	177	118	79	
13. Linear bit density (kfci)****	600	662	729	804	886	977	5.00% per year
fcmm = kfci/0.0254	23,622	26,043	28,713	31,656	34,901	38,478	
14. Areal density (Gb/in^2)	27.18	42.12	69.64	115.16	190.44	314.95	27.76% per year
15. Tape width in mm	12.65	12.65	12.65	12.65	12.65	12.65	
16. ECC and formatting overhead	20%	20%	20%	20%	20%	20%	0% per year
17. Servo track and layout overhead****	16.0%	10.4%	10.4%	10.4%	10.4%	10.4%	
18. Number of passes to write a tape	592	444	666	999	1498	1123	
19. Number of passes to end-of-life (media)	34560	37093	39812	42730	45862	49223	3.6% per year
20. Time to fill a tape in hours	31	41	54	71	94	124	14.78% per year
21. Number of data tracks	18,949	28,415	42,613	63,911	95,862	143,799	22.47% per year
22. Bit Aspect Ratio (BAR)	16	12	9	6	5	3	-14.26% per year
23. Uncorrectable Bit Error Rate (UBER)	1e-20	1e-20	1e-20	1e-20	1e-20	1e-20	

Table 1: 2024 Tape Technology Roadmap Detail

* Fibre Channel Industry Association Roadmap V24
** Defined as the length of tape required to store the defined tape capacity. It does not include the reserved space for possible defects.
*** Defined as the total length of tape including length used for attachment and hub covering
****Defined as the 1T Kfci where T is the data cell length
***** On non-capacity reserve overhead only, unused servo area claimed at change to 64 channels

A dimensional comparison of the HDD recorded magnetic bits, INSIC 2024 Tape roadmap, and LTO is shown in Figure 3. It is evident that, in addition to magnetic limitations, HDD bits with incredibly small size may face mechanical constraints. However, tape technology has a good chance of growing its capacities in the upcoming decade.

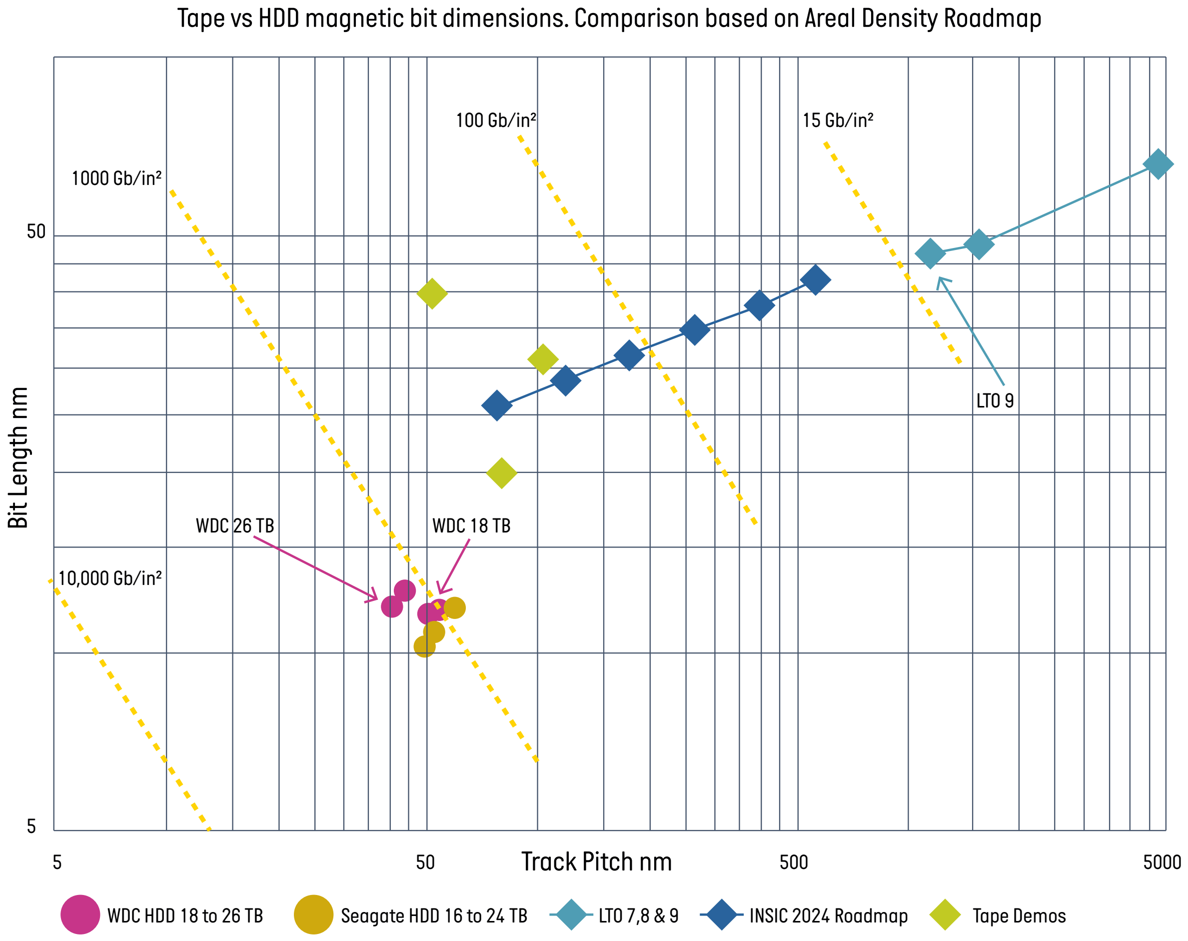

Figure 2 Magnetic Bit Dimensions with Areal Density Contour Lines

Figure 3 below shows at scale comparison of Physical Recorded magnetic Bits for Tape vs HDD Technologies. Physical illustrative Shapes of Recorded Magnetic Bits for 26 TB WDC (light green), INSIC 723 TB (light blue), INSIC 45 TB (dark blue) and light gray is LTO-9 at 18 TB all showing at scale.

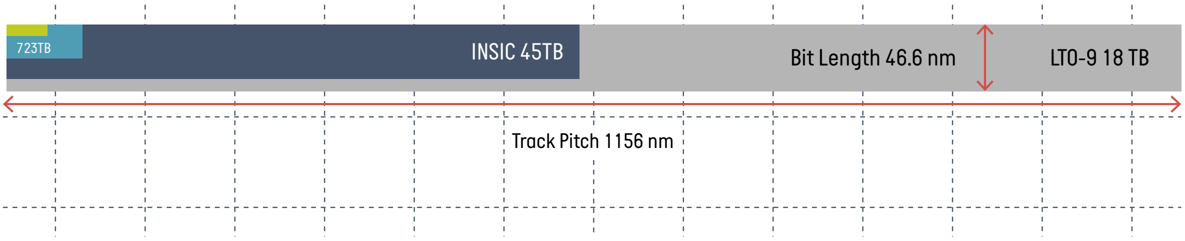

Technology	Bit Length nm	Track Pitch nm	Bit Aspect Ratio (BAR)
WDC 26 TB HDD	11.91	40.64	3.4
LTO-9	46.6	1156	24.8
INSIC 45 TB	42.3	560.8	13.2
INSIC 723 TB	26	78.8	3.0

Figure 3 Comparison of Tape vs HDD Magnetic Bit shapes

2.2

Roadmap Discussion

In reference to Table 1 above, the entries highlighted in orange are inputs into the Roadmap; the other entries are calculated based on these inputs. Additionally:

1) **Capacity** is the user native (uncompressed) capacity in terabytes (TB) of data, with year 2024 of 45 TB being an input as well as the 32% per year capacity growth. The 32% exceeds the anticipated HDD technology growth rate and gives an approximate doubling of capacity every 2.5 years. Tape drives have internal compression hardware that can provide lossless compression if the data has some compressible content. This yields higher data rates and capacities over the native ones discussed in the Roadmap. The Roadmap does not include compression capabilities, so real user capacities and data rates could actually be better than what the Roadmap shows.

2) **Maximum data rate per channel** is the data rate of each channel that writes one track of data. This gives the required capability of each transducer in the head. This growth is calculated by dividing the Total data rate by the Number of channels.

3) **Maximum streaming drive data rate** is the maximum streaming (no drive back-hitches) native (uncompressed) data rate that a user can achieve. This is calculated by taking total data rate with a starting point of 400 MB/s in 2024 and assuming it grows by 15% per year.

4) **Minimum drive streaming data rate** is the minimum streaming data rate calculated by assuming a minimum tape speed of 1.2 meters per second. The drive cannot reliably go below 1.2 meters/second due to difficulties to control servo and tension. This is primarily a limitation of the update rate of the current servo pattern used in LTO and Enterprise drives. A reduction in the frame length of a future servo pattern could enable lower tape speeds and hence lower data rates.

5) **The FC (Fibre Channel Industry Association Speed Roadmap** reference [5] is included to show that the Fibre Channel speed is at least 16x the native data rate of the tape drives out to the year 2034. For midrange products that do not use the Fibre Channel attachment interface, we expect that the interface data rate will not become a limiting factor on overall data transfer rates during the Roadmap timeframe [6]. In addition, the 100 GB/s to 400 GB/s for FCoE serial speed and quad speed technologies should also surpass the interface requirements for tape drives using Ethernet based interfaces [7].

6) **Number of channels** is the number of simultaneous data tracks the drive reads and writes as the tape moves across the head. This parameter is an input that is set for each year without a constant scaling factor, starting at 32 channels in 2024 and then increasing at year 2026 to 64 and at year 2034 to 128. This was felt to be a reasonable increase and characteristic of how products will likely scale.

7) **Tape Speed (m/sec)** is the velocity with which the tape moves over the head when reading and writing at the maximum data rate. In the Roadmap, it is calculated by dividing the data rate per channel by the linear bit density and factoring in the ECC and formatting overhead. Increasing tape speed does not come without risks. Thinning tape will necessitate better tension control and smaller track pitch will necessitate a reduction in allowable Position Error Signal (PES), both of which become even more challenging at high tape speeds. For this reason, we restrict the tape speeds to less than ~7 m/s.

8) **Tape thickness** is the total thickness of the tape and is given as an input into the Roadmap. The thinner the tape, the more tape that can be contained in the cartridge, and hence more capacity. Increasing capacity by thinning the tape results in the cartridge cost increasing somewhat, since each cartridge then contains more tape. Thinning the tape also adds technical challenges to handling of the tape in the drive. The numbers were chosen to represent what is considered achievable.

9) **Data capacity reserve** is how much additional capacity the cartridge has over the user capacity defined in 1. This is used to make sure that tape to tape copies can be done without overflow due to possible defects in the media. It is obtained by having some extra wraps (i.e., extra length of tape). This is assumed to be fixed at 3.0% for the entire 10-year Roadmap.

10) **Tape length that is recordable** is the length of tape required to store the defined tape capacity. It is used to calculate areal density based on the defined capacity. It does not include reserved space for defect mapping or for other uses and is calculated by subtracting the reserve length of 40 meters from the total tape length.

11) **Tape length total** is the length of tape as defined including length used for connection of the tape to a leader mechanism, reserved space for defect mapping, and excess tape used for wrapping around the drive take-up reel hub. This is done to reduce the effects of hub surface imperfections, giving an effectively smoother hub surface for the portion of the tape storing customer data. It is calculated based on the physical dimensions of the cartridge hub and the tape thickness.

12) **Track density** is the number of data tracks per inch (TPI) in the transverse direction of the tape. This is calculated by taking the total areal density and dividing by the linear bit density in the down-track direction within a track. this is a critical parameter in driving the technology as discussed in the following sections. Compared to hard disk, the track density of tape is lower by a factor of around 12 to 25. This has been necessitated by the challenges associated with track following on a flexible tape substrate and a combination of the dimensional instability of the tape substrate and the multi-track recording, which place additional tolerance requirements on tape. Track density improvement has been identified as the area with the greatest potential leverage for advancing tape technology.

13) **Linear bit density** is the number of bits per inch going down a track. It is defined here to be the number of data cells per inch and is expressed in kfci (thousands of flux changes per inch). The highest number of flux changes per inch occurs when one flux reversal is in every data cell, sometimes called the 1T kfci where T is related to the length of the data cell. This is also a critical parameter in driving the technology. When comparing to hard disk, the linear tape density is only about a factor of 4.5 less, so there is not as much opportunity for improvement here as for the track density. The 5% per year growth shown here was chosen to give reasonable linear bit densities in the later years of the Roadmap. Recent technology demos [2] also support future linear density improvement capability.

14) **Areal density** is the raw number of data bits per square inch on the tape. It is calculated by taking the user capacity in item 1, above, and dividing by the usable tape area, factoring in all formatting, ECC (error correction code) and servo overheads shown in 16 and 17, below. Areal density is also equal to the product of track density and linear bit density.

15) **Tape width** is fixed at 12.65 mm for the entire Roadmap.

16) **ECC and formatting overhead** is assumed to start at 20% in 2024, which is in line with current products shipping and is projected to remain constant at 20%, in line with recent trends.

17) **Servo track and layout overhead** is assumed to start at 16% in 2024, drop to 10.4% in 2026 and remain at 10.4% out to 2034. The reduction in overhead is due to an anticipated change in servo pattern height enabled by the transition to 64 parallel channels.

18) **The number of passes to write a tape** is the number of end-to-end passes required to completely fill the tape with data. It is calculated by taking the total number of tracks written on the tape and dividing by the number of channels (simultaneous tracks written by the head).

19) **Number of passes to end of media life** is number of end-to-end passes along the tape for the usable life. A moderate improvement of 3.6% per year is projected.

20) **Time to fill a tape** is how long it takes, in hours, to completely fill a tape when writing at the maximum data rate. This number is increasing because the capacity is increasing faster than data rate. The increase shown in this Roadmap is felt to be acceptable, especially with the shift in the use of tape away from short-term backup and towards long-term archive.

21) **Number of data tracks** is calculated from other parameters and is the total number of data tracks written on the tape to support the user capacity in item 1, above.

22) **Bit Aspect Ratio (BAR)** is the ratio of the track-width to bit length and is used as a comparison to hard disk technology, which is at about 4.9 today. .

23) **Uncorrectable Bit Error Rate (UBER)** is the inverse of the average number of bits transferred (read) before an uncorrectable error is encountered. The starting point in 2024 is 1e-20, the value specified in current LTO and Enterprise products. UBER is kept constant for the duration of the Roadmap because it is expected that improvements in ECC will be used to improve areal density rather than to improve UBER, as discussed in Section 2.3.4.

2.3

Key Technology Challenges

Recent generations of LTO tape drives have used BaFe (barium ferrite) based particulate media in combination with a TMR (Tunneling Magnetoresistive) recording head. However, the most recent particulate tape areal density demonstration [Furrer 2021] showing the feasibility of capacities out to 580 TB, utilized SrFe (strontium ferrite) particulate media rather than BaFe. Moreover, the recently announced 50 TB TS1170 JF media uses hybrid magnetic particles that combine the technologies used in ‘next’ generation Strontium Ferrite (SrFe) particles and Barium Ferrite (BaFe) particles. It therefore seems likely that the industry will transition from BaFe particles to SrFe particles which in combination with TMR recording heads will likely support multiple future generations of tape drive.

One of the challenges highlighted in this Roadmap is to determine how these technologies might be pushed to achieve the 723 TB capacity target in the year 2034 or, alternatively, if something else will be needed to replace these technologies in that timeframe. In addition, given the high track densities, new media substrate technologies, tape drive compensation techniques and/or environmental controls may be required to achieve the dimensional stability goal in the 2034 timeframe.

The following sections of the Roadmap discuss the key technology areas that comprise a tape drive: media, heads, mechanical transport mechanisms, and the recording channel. Each of these areas has been explored in detail by its respective technical teams, and the technology challenges and options are documented in these sections.

2.3.1

Media Technology

The recording media technology chapters of the 2015/2012 INSIC Roadmaps are still very relevant and discuss the key challenges that continue to exist in extending the roadmap.

2.3.1.1

Particulate Media

The most recent technology demonstration for particulate tape media achieved an areal density of 317 Gb/in^2 using ultra-fine magnetic particles of Strontium Ferrite (SrFe) [Furrer 2021]. In this demonstration, the volume of the SrFe particles was reduced to less than 60% compared to that of conventional BaFe. A new dispersant was also used to evenly disperse and arrange the particles on the nanoscale to achieve a high signal-to-noise ratio. Additionally, a newly developed smooth non-magnetic layer (under layer) was used to improve the smoothness of the tape surface even further, reducing the spacing between the magnetic tape and the head for enhanced recording performance. When combined with new track-following technology, signal processing technology and recording/read back devices that support them, the magnetic tape with a servo pattern precisely laid out can deliver an areal density of 317 Gb/in^2, giving a further boost to the potential capacity of magnetic tapes.

2.3.1.2

Sputtered Media

An areal density of 201 Gb/in^2 was demonstrated in 2017 using a prototype sputtered tape [Furrer 2018]. Since that demonstration, the development of sputtered tape has continued, enabling further increases in recording density. For example, the technique of adding Co_3O_4 into magnetic composite targets was successfully used to greatly enhance the magnetic isolation between grains by controlling the oxidation process during sputtering of the magnetic layer.

This technique reduces media noise and improves SNR by +1.8 db. In addition, the introduction of a capping layer further improves SNR by an additional +0.5 db, resulting in total SNR improvement of +2.3 dB. Assuming the combination of the latest track-following technology [Furrer 2021] with this magnetic layer, an areal recording density of 375 Gb/in^2 was estimated to be feasible in a 2023 paper describing this work [Tachibana 2023]. In the future, higher areal recording density can be expected by adding other oxides to the magnetic layer to induce further magnetic isolation, further improving the capping layer technology, and by creating a multilayer structure of magnetic layers that takes into account the magnetization reversal mechanism.

In addition to SNR improvements, there are also prospects for reducing the material costs for sputtered tape, which is one of the challenges to commercialization. For example, the sputtered media used in the 2017 demo mentioned above incorporated an underlayer of the precious metal ruthenium (Ru). The replacement of this Ru layer with a less expensive CoCr/CoCr-oxide multilayer film has been investigated [Tachibana 2021] and it was confirmed that magnetic properties with this multilayer film are equivalent to those with a Ru underlayer. Studies are ongoing to realize sputtered tapes with $/TB prices close to those of particulate tapes.

2.3.1.3

Remarks on Areal Density Demos

The areal density demonstrations discussed in the previous two sections have shown the potential for tape media to support areal densities in the range of the 2034 roadmap target, however, they were also all single channel demonstrations that did not take into account the dimensional stability of the media and the additional challenges that this creates for parallel channel tape recording. As a result, the demo track densities are higher than the roadmap targets at comparable areal densities. On the other hand, current HDD products also use sputtered media and operate at track and linear densities significantly higher than projected in the Roadmap for 2034 and thus provide additional evidence of the feasibility of the roadmap projections for linear and track density. Moreover, strategies to overcome the challenges associated with tape dimensional stability are discussed in section 2.3.3. Such technologies could enable more aggressive track density scaling then projected in Table 1.

2.3.1.4

Substrate

The tables below summarize expected improvements in media properties over the duration of the roadmap. Improvements in TDS will be obtained by tensilization orientation and nano-sized polymer blends. The handling of the thinner tape (to achieve tape lengths in excessive of 1,000m) will provide challenges in both media production and use in the tape drive. The substrate suppliers will continue to reduce the surface roughness of the base-films.

	2024	2026	2028	2030	2032	2034
Substrate Thickness, μm	4.0	4.0	3.8	3.8	3.6	3.6
Thermal*	0	0	0	0	0	0
Hygroscopic	470	470	420	420	390	390
In-Cartridge Creep	100	100	100	100	100	100

Table 2: Dimensional stability goals (in ppm) for Tapes using Modified PEN Substrates

	2024	2026	2028	2030	2032	2034
Substrate Thickness, μm	4.2	4.0	4.0	3.8	3.8	3.6
Thermal*	0	0	0	0	0	0
Hygroscopic	375	350	300	300	250	250
In-Cartridge Creep	100	100	100	50	50	50

Table 3: Dimensional stability goals (in ppm) for Tapes using Advanced PET Substrates

	2024	2026	2028	2030	2032	2034
Substrate Thickness, μm	3.6	3.0	3.0	2.8	2.6	2.4
Thermal*	100	50	0	0	0	0
Hygroscopic*	100	50	0	0	0	0
In-Cartridge Creep	50	50	50	50	50	50

Table 4: Dimensional stability goals (in ppm) for Tapes using Aramid Substrates

* Note: Zero's in the Tables 2, 3 &-4 represent values that are too small to measure with any degree of certainty and as such will have little impact on the overall TDS performance.

From the data in Tables 2, 3 & 4 only the Aramid substrate approaches the TDS needs for the roadmap, in the absence of active TDS compensation. However, Aramid is significantly more expensive than the other substrates. To utilize the other substrate materials, the drive system will need to actively compensate for the changes in track width due to environment or other factors. For example, LTO9, TS1160 and TS1170 drives dynamically adjust tension to maintain a specific track pitch. For this reason, the tension TDS component is not shown in this roadmap. Other TDS compensation schemes may be used in the future. The above tables also show that for the modified PEN and advanced PET substrates the environmental conditions, more specifically the humidity (Hygroscopic component), are by far the largest contributors to the TDS total. A better controlled operating environmental range would allow the tension compensation to be used to counter the irreversible in-cartridge creep TDS component.

2.3.1.5

What Key Improvements are Needed to Exceed 1,000 kbpi?

The 2019 INSIC Tape Roadmap projected an 8% CAGR in linear density scaling such that linear densities greater than 1000 kbpi would be required by 2027. In the current Roadmap, the projected linear density scaling factor was reduced to 5% CAGR which is more in line with the scaling trend of recent products. The smaller scaling factor results in a linear density target at the end of the roadmap of 977 kbpi. Nevertheless, the question of what key improvements are needed to reach and exceed a linear density of 1000 kbpi on tape remains relevant. The linear density of 702 kbpi achieved in the most recent tape recording demo on particulate media [Furrer 2021] is significantly below the current roadmap's target of 977 kbpi. However, the reported demo operating point used a 29 nm wide reader and a calculated track width of 56 nm, which is significantly below the roadmap target track pitch for 2034 of 79 nm. Operating at a wider track pitch with a wider reader would provide additional SNR to enable operation at a higher linear density. Moreover, recent demos on sputtered tape have reported linear densities of 818 kfci [Furrer 2018] and 1035 kfci [Jubert 2022], indicating that the current roadmap target is achievable albeit challenging.

For particulate media to achieve the necessary SNR at the linear densities projected by the current roadmap, a reduction in particle volume is needed. This will require continued improvements relative to current commercial particle technologies and may require new particles to be considered. One promising approach to reduce particle volume while maintaining the required thermal stability is the use of strontium ferrite (SrFe) particles. SrFe was used in the recent areal density demo on particulate media mentioned above [Furrer 2021] where an areal density of 702 kbpi was reported. Additional improvements in SNR could be achieved through a further reduction in particle volume, narrowing the switching field distribution of the particles, and reducing variations in the thickness of the mag layer. Note that the recently announced 50 TB TS1170 JF media uses hybrid magnetic particles that combine the technologies used in 'next' generation Strontium Ferrite (SrFe) particles and the Barium Ferrite (BaFe) particles that are currently used in the latest generation of LTO tapes. Another particle type under consideration is ε-Fe_2O_3. A prior study demonstrated spherical ε-Fe_2O_3 nanoparticles with a diameter of 8.2 nm having an Hc value of 5.2 kOe at room temperature [Ohkoshi 2015]. Recently, ε-Fe_2O_3 magnetic particles combining fine particle size with a tighter particle size distribution and lower Hc have been produced. Improvement in tape write transducer technology to increase the magnitude of magnetic field produced will likely be required to take advantage of particles with such high coercivities. Energy assisted recording paradigms currently under development by the HDD industry, such as microwave assisted magnetic recording (MAMR), could also be adapted to tape recording and used to enable the use of higher coercivity media. Research in this direction was recently reported in a paper that investigated the use of focused millimeter wave energy to reduce the magnetic field required to flip the magnetization of high coercivity ε-Fe_2O_3 particles [Ohkoshi 2020]."

Sputtered media uses similar recording materials to those employed in the hard disk drive industry. State of the art 3.5" HDD products based on perpendicular magnetic recording (PMR) operate at areal densities of up to 1.26 Tb/in^2 and linear densities of up to 2552 kbpi, and with shingled recording operates at areal densities up to 1.322 Tb/in^2 and a linear density of 2133 kbpi. In a recent sputtered tape demo, a linear density of 818 kfci was achieved at tape speeds of meters/second [Furrer 2018] and more recently a linear density of and 1035 kfci was reported using a drag tester operating at lower scanning speeds [Jubert 2022]. More recently, technology to enhance the magnetic isolation of the recording layer in sputtered tape has been shown to increase SNR and is expected to enable further improvements in recording density [Tachibana 2023]. Moreover, the application of additional technologies developed for HDD to sputtered tape, such as a soft magnetic underlayer (SUL), a capping layer and an exchange coupled continuous structure, are expected to enable significant progress in SNR and hence also enable increases in linear density. The use of a SUL will enable perpendicular recording using mono-pole write heads which will improve SNR compared to writing with a ring head and enable the use of higher coercivity media with smaller grain size. Improvements in other aspects of tape read and write transducer technology, as discussed in the section on tape heads, will also help to enable linear density scaling for both particulate and sputtered media.

Another factor critical for exceeding 1,000 kbpi will be the continued reduction in tape-head spacing. Tape recording systems operate with the head in contact with the media and hence tribology issues need to be studied to enable the reduction in tape-head spacing and the smoother tape surfaces that will be required. There are several factors that contribute to tape-head spacing, but the key factors are likely to remain the surface roughness of the media (and substrate) as well as the depth of the pre-recession of the tape head and the thickness of the wear resistant coating used to protect the TMR read sensors. Media manufacturers will need to continue to refine the design of the media surface, via lubricant and surface roughness engineering, to allow closer head-to-tape spacing while reducing abrasivity and maintaining low coefficients of friction to enable good media durability. In addition, new binder technologies may need to be developed to maintain the robustness of the tape surface. In parallel to this effort, tape drive manufacturers will need to develop head technology to reduce friction and wear such as by patterning the surface of the head [Engelen 2016]. Operation at lower tape tension and with lower head wrap angles should also help to reduce friction and wear. In addition, drive manufacturers will need to develop robust, ultra thin head coatings (on the scale of a few nanometers) to further reduce spacing.

2.3.2

Future Technologies for Recording Heads

2.3.2.1

Introduction

This section of the 2024 Technology Roadmap examines present-day recording heads and changes envisioned for enabling the areal density growth shown in Figure 2. Aspects that may change include transducer dimensions and other characteristics, their number and arrangement in the heads and mode of deployment of heads in drives. Going back as far as the 1970’s, there were competing, proprietary tape drive designs and formats, e.g., open reel linear tape, helical scan cassettes, etc. The year 2000 marks the beginning of a what would become a convergence in the digital tape industry, around a single, open platform called “Linear Tape Open” (LTO). To date, LTO’s enduring features include: a cartridge containing a ½” linear recording tape having four data bands, for minimizing track mis-registration due to TDS; and a timing-based servo (TBS) pattern, for accurate and extendible track positioning and LPOS encoding [Jaquette 2003]. Introduction of this platform also marked the launch of a new, HDD-based, flat recording head [Biskeborn 2003]. Flat recording heads have been fabricated largely in the same facilities used to make HDD heads. This has enabled tape to leverage the research, development and manufacturing created for the (high-volume) production of HDD heads. Both HDD and tape wafers are “AlTiC” ceramic, a hot isostatically pressed mixture of crystalline aluminum oxide and titanium carbide powders. This material provides a tough, stable and electrically conductive substrate for both HDD and tape heads, and a wear resistant media-bearing surface (MBS) having durable skiving edges, both of which are required for tape heads. The areal density plots in Figure 1 show that disk is perhaps 15 years ahead of tape, suggesting that tape can continue to use well established HDD materials and processes in the time frame of this Roadmap and beyond.

2.3.2.2

Channels and Layout

The first generation LTO tape drive had 8 concurrent read/write channels for writing to each data band [Biskeborn 2003]. In 2005, LTO-3 launched, marking the transition to 16 concurrent channels. Each head “module” had 16 piggybacked read-write transducer pairs, for a total of 32 transducers (plus two servo readers), arranged on a pitch of 166.5 micrometers (um). Two modules were joined together, creating a head having read-while-write capability between opposite modules. In 2011, the number of active channels doubled again to 32, for an enterprise tape drive, the IBM TS1140, and subsequently for LTO-7, which was introduced in 2015 [Biskeborn 2012]. The transducer pitch was halved to 83.25 um for maintaining backward compatibility with the existing 4-band tape format [Biskeborn 2012]. Piggyback read-write pairs were abandoned in favor of processing read and write heads on their own separate, dedicated wafers (noting that writer heads still have servo readers). Thus, the heads for 32-channel drives are comprised of three modules: two writer modules flanking a central reader module, as shown in Figure 4.

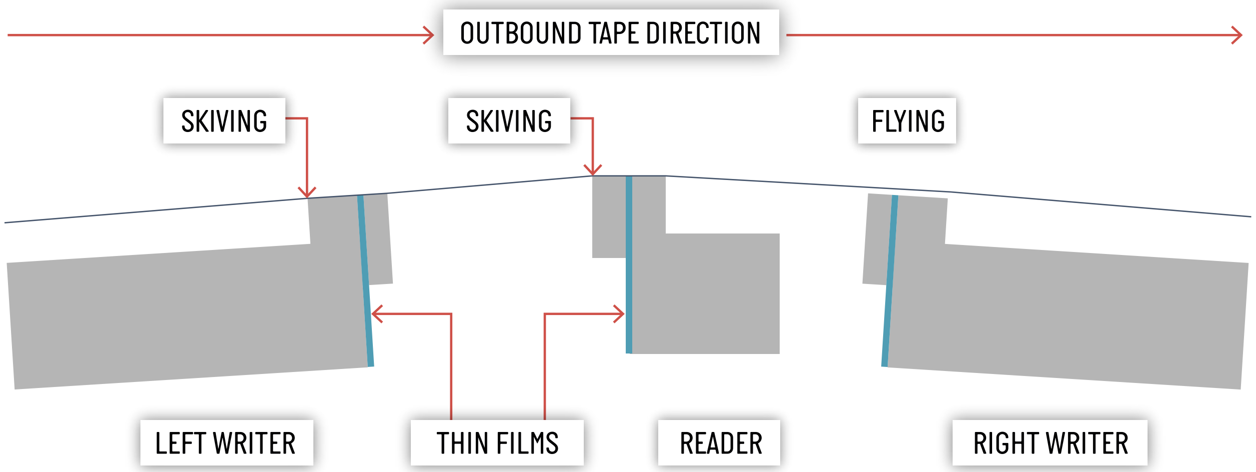

Figure 4. Three-module head with tape, which is shown flying over the downstream module for the indicated (outbound) direction of tape motion.

The center module reads data written by the left module for forward tape motion (shown) and by the right module for reverse motion. This arrangement simplifies the cabling requirements that having 32 piggyback pairs (i.e., 64 transducers, plus servo readers) in each of the two modules would have imposed. Other benefits are lower head-tape friction; independently customizable head-tape interfaces for read and write modules; and narrow mechanical gaps for readers and writers.

The 2024 roadmap anticipates that the number of concurrently operable transducers will double again to 64 within the next 10 years. A question, then, is what would the tape layout for 64 channels become? The 2015 Technology Roadmap projections were for 60 active channels and 16 data bands by 2025, implying a transducer pitch of ~12 um. The 2019 projections were for 64 channels in 8 data bands by 2029, implying a channel pitch of ~22 um; and there was mention that active TDS compensation would likely be required. It is not currently known if reducing channel pitch to this level is practical, for example, because of writer-to writer crosstalk [Biskeborn 2008]. An alternative may be to not shrink channel pitch but instead rely on active TDS compensation to accommodate the resulting ~2x larger head span. This implies, for example, that future tapes would have two data bands rather than four. With this approach, 128 channels by 2034 could become a single data band architecture. Examples is this are shown in Figure 5. It must be noted that this is speculative, and other arrangements may be possible. Furthermore, maintaining backward compatibility may not always be assured. If uncorrectable non-linear mis-registration places restrictions on increasing the number of channels, then, at least until compensation is improved, an approach such as HPSS's (High Performance Storage System) Redundant Array of Independent Tapes (RAIT) collaboration may be utilized, at least for some deployments.

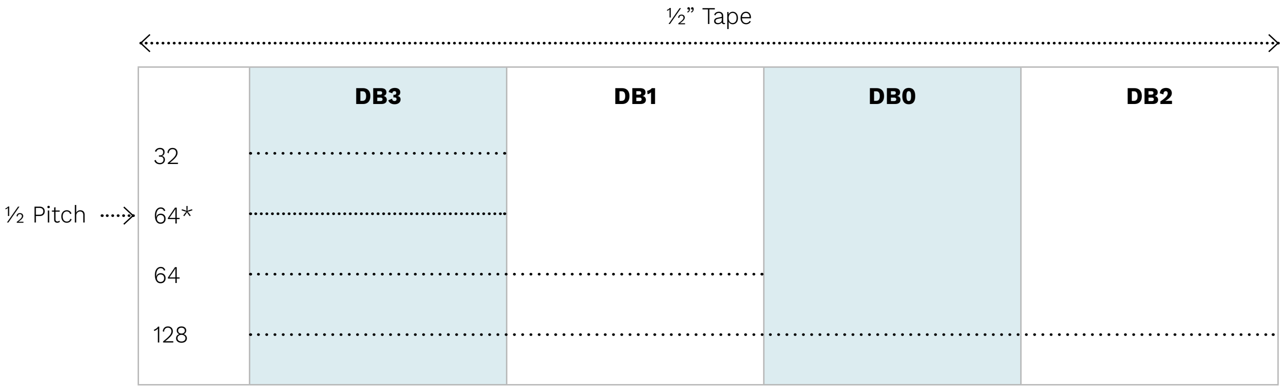

Figure 5. 32 and 64 transducers shown in relation to current LTO data bands (DB0, 1, 2 and 3). The pitch required to squeeze 64 channels into one data band is ~50% less than it is for current LTO channel pitch.

2.3.2.3

The TMR Read Sensor

Introduction

Compared to giant magneto resistive (GMR) sensors, tunneling magneto resistive (TMR) sensors provide several-fold greater amplitude and greater stability, as well as running cooler and exhibiting lower signal loss over their life [Biskeborn 2018], all of which are important not only for disk but also for tape. HDD transitioned from GMR to TMR in 2004.

Tape transitioned to TMR in 2017, first for an enterprise drive and a few months later for LTO-8. This was possible only because of tape's on-going ability to leverage HDD's vast resources. A major question was can TMR survive in continuous contact with abrasive, "dirty" tape recording media, given that a single "smear" can potentially destroy the sensor (vs GMR, which may tolerate multiple smears). Cooperation between the tape and HDD communities was critical. Today, TMR is anticipated to be viable for the timeframe of the current roadmap, if not well beyond.

It was perhaps fortuitous that the first TMR's for tape were huge compared to HDD at the time: trackwidths (TW) >1000 nm vs 40-50 nm for HDD, and stripe heights (SH) of ~350 nm vs 40 nm. This implied that the tape MgO tunnel barrier area (TW x SH) would be nearly 200x greater, implying that tape's tunnel barrier resistivity could be larger while and still providing a sensor resistance in the range 50-150Ω, which was needed for the tape read channels. The thinner tunnel barrier needed for HDD is relatively more challenging to control and may produce higher, and more variable, sensor resistance.

For a given recording medium, optimum AD (and SNR) depends on the interplay between several factors that ultimately determine the optimum bpi/tpi combination, and fine tuning this is left to experimental verification. But for discussion purposes, it is convenient to look at the effects of bpi and tpi on future TMR designs independently. Again, as for HDD, bpi is growing at a slower rate than tpi (~4x), largely because SNR falls off rapidly with increases in bpi vs tpi.

Implications for growth of bit density

SNR of future tape media must grow to enable higher bit and track densities. Increasing bit density (decreasing bit length) will call for a reduction in TMR shield-to-shield (S2S) separation. Generally, S2S needs to scale approximately as the inverse of bit density [Bertram 1994] to maintain optimal recording performance (amplitude, resolution, SNR), but with the caveat that at some point amplitude may decline faster than resolution improves channel error rates. Today S2S is ~90-100 nm for tape at ~550-600 kbpi (i.e., ~2x bit length). Thus, a 1.63x increase in bpi over the next 10 years suggests that S2S will need to approach ~55 nm. Of note is that today S2S for HDD is limited by sensor thickness, which is of the order of 23 nm, thus setting a lower bound for tape many years from now.

TMR's for tape require the use of conductive, non-magnetic spacer layers between the sensor and shields to provide the wider S2S gap needed for the lower linear bit densities of tape. The sensor itself typically occupies about one third of S2S, and the spacer layers each about a third. Nichrome spacers have been replaced by a bi-layer iridium/alumina structure to enhance resistance to smearing and shorting caused by asperities on the moving tape [Biskeborn 2018]. At some point, reducing S2S by thinning the spacers (up to 50%) may increase susceptibility to shorting, implying that other measures may need to be explored in the future.

Increases in bit density will also call for a reduction in head-media separation (HMS), defined as the distance between the tip of the read sensor and the top of the media magnetic layer [Marchon and Olson 2009]. For tape, this is the sum of sensor recession, head coating thickness and height of tape asperities. While difficult to model precisely, trends from HDD indicate that HMS is close to ½ the bit length for products having bit lengths ranging from 300 nm down to 20 nm [Marchon and Olson 2009]. This suggests that in 2024, HMS for tape "should" not be greater than about 42 nm and decrease by a factor of 1.63 to about 26.0 nm in 10 years. Specific values of HMS depend on the product, state of head and media usage, etc. The three contributors each need to decrease on average by about 5 nm. Today, factory pre-recession and alumina head coating are tuned to enable TMR to survive in a contact recording environment while minimizing their contribution to HMS. Note that reducing HMS could go either way insofar as shorting is concerned: smoother media implies smaller asperities but also implies more intimate contact between head and tape. Regarding smaller HMS, reducing the size/protrusion of wear particles in the particulate media coating has a likelihood of increasing both static and running friction of the tape. This is discussed in section 2.3.2.5 entitled, "Static and running friction and HTI." Other concepts for enabling future reduction of HMS are flux guides, or sensors having shields (which exposed at the head's MBS) that are held at the same electrical potential and thus not capable to short the sensor in the presence of a smear.

Implications for growth of track density

The trackwidth (TW) of the narrowest tape TMR sensors today is approximately 300-500 nm. While this is still roughly 10x greater than the width of modern HDD sensors, the roadmap suggests that tape sensor TW will nearly close this gap in the next 10 years, i.e., approach 40 to 50 nm. This is driven in part by tape's need to accommodate a much larger PES than HDD. In general, sensor magnetic stability requires sensor stripe height (SH) to be comparable to or less than TW. Otherwise, unstable magnetic domains may form in the sensor NiFe free layer, and switching of these domains creates noise that can be as large or larger than readback signals, even for sensors having optimized magnetostriction (λ_s) and optimized hard bias magnets. Critical to achieving required SH depends on how well it can be controlled in manufacturing, i.e., its tolerance. Over time, tape has been adopting HDD's state-of-the-art control processes, which provide the dimensions needed for much smaller HDD heads. Stable TMR tape head sensors having the dimensions needed by 2034 are anticipated.

Soft bias

Another sensor improvement created for HDD's, and likely to be required for tape in this timeframe, is "soft bias," which eliminates the hard bias magnets used today. Figure 6 shows MBS views of soft and traditional hard bias sensors. Soft bias addresses two concerns. One is bias variation in a population of heads having SH less than ~200 nm. With short stripe heights, hard bias magnets have fewer, randomly oriented grains. This produces biasing fields that vary significantly from magnet to magnet, leading to bias shifts. The other concern is side reading, where a sensor picks up signal in the erase bands at the recorded track edges, thus degrading SNR: with soft biasing, permalloy replaces the low permeability magnets. However, adapting HDD soft bias to tape will require significant development. This is because the metal spacers needed for setting the wide read gaps required for a tape would block the required coupling between the trailing shield and soft bias film.

Servo readers

Tape heads, unlike HDD heads, have dedicated sensors for reading the servo patterns that are written to the tape by the media manufacturers. The servo heads typically straddle the array of data elements. These sensors are optimized for reading the servo pattern "data," which for TBS appears as a repetitive set of fixed amplitude bursts at frequency of a few megahertz. Since the bandwidth is low, the read gap is much wider than for data readers. Servo reader trackwidth is chosen to provide an optimal balance between TBS resolution (narrower TW is better) and SNR (wider TW is better) and typically changes only when the pattern itself (infrequently) changes.

While not a gate for areal density, there is an interesting consequence of the fact that servo reader TW does not change much over time. As data reader TW is reduced to support higher tpi, RA must follow suit to maintain ~constant data reader resistance. Since servo and data reader TMR's use the same RA, at some point servo reader resistance will become very small. Fortunately, there are several approaches for addressing this: reduce servo reader SH (enabled by improved SH control); utilize HDD TMR shunting technology for the data readers; provide separate RA's for data reads vs servo readers, similar to HDD's two-dimensional magnetic recording (TDMR) readers.

An interesting outcome of having separate wafers for readers and writers is that writer servo readers are not subject to same limitation as reader servos, since there are no data readers in the writer heads.

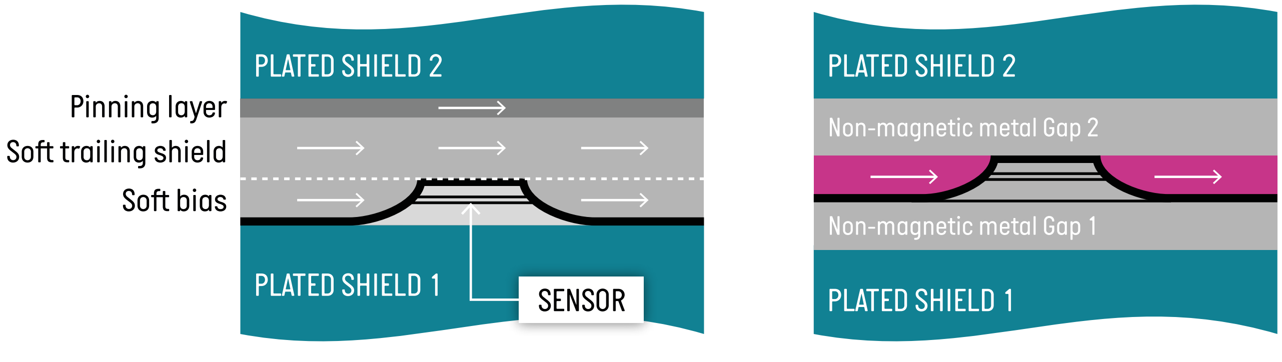

HDD SOFT BIAS

TAPE HARD BIAS

Figure 6. Media bearing surfaces of HDD soft-biased sensor (left) and tape hard-biased sensor (right). The soft bias film is magnetically permeable and thus provides effective side-shielding.

2.3.2.4

Write Transducers

Ring write heads

Today, ring-type (aka "longitudinal") transducers are used for writing modern particulate BaFe (and more recently SrFe) media. A fortuitous characteristic of the plate-like BaFe particles is their anisotropy axes tend to align orthogonal to the tape media surface during manufacturing. This quasi-perpendicular orientation enables the vertical component of the head field near the write gap trailing edge to write the media. Tape demonstrations have shown that magnetically orienting the media during manufacture further enhances readback SNR. A list of recent tape areal density demonstrations is provided in Table 5. Data from these studies suggests that ring heads may be suitable for writing future thin film media, which have strong vertical orientation [Lantz 2015, Furrer 2017, Furrer 2021].

Ring head improvements

Initially, the poles of HDD-based ring heads for tape were plated 45%/55% nickel iron, which has a saturation magnetization of 1.6T. To enable writing to higher coercivity media, an idea was borrowed from HDD, namely providing a CoFe alloy seed layer on the trailing edge side of the write gap. The saturation magnetization of this material is ~2.3T. This enables recording to media having ~35% higher coercivity and creates a larger write field gradient, which sharpens write transitions and improves SNR. The high B_s layer was incorporated in the heads in the TS1155 drive and later used for both top and bottom poles to obtain further improvement in SNR [5].

In the future, media having higher Hc may demand higher write currents, potentially surpassing a given write driver's capabilities. Adding more turns to the write head coil may help address this and limit the increase in power that comes with higher current. If so, then maintaining a ~83 um (or larger) channel pitch for 64 channels helps ensure there will be space for more coils.

An area of critical importance moving forward is side writing, which, roughly speaking, is the width of the portion of a shingled track having less than perfect overwrite (aka the erase band, which is due to the write field not falling off perfectly abruptly) and excessive transition curvature. Typically, SNR is degraded in this region, which today is roughly 100-200 nm. If erase band does not change, it will become a larger fraction of the total shingled trackwidth. Side writing is partly controlled by notching the pole tip edges, shown on right in Figure 7. Side writing also depends on the length of the write gap (~125 nm today) and HMS. Future operating points will likely require higher coercivity media having thinner magnetic layers, as well as lower spacing, which is largely limited by media roughness and friction. Presumably, these trends will demand smaller write gaps. Modeling and empirical observations will be needed for designing write heads that minimize side-writing and thus maximize tpi. At some point, additional measures may be required for achieving the tpi growth predicted in the roadmap.

Write heads for future media

This is the subject of what could be a much longer discussion, so this section will just touch a few key highlights. In the early 2000's, HDD ring heads did not provide sufficient magnetic field and field gradient to utilize the full capabilities of perpendicular sputtered media. It was recognized that a shielded monopole write head having a wide, low reluctance return pole as shown in Figure 7 (left side), combined with highly oriented sputtered media having a magnetically soft under layer (SUL) may address this. It was realized that the low reluctance return pole could be made to wrap around the main pole, and thus contain flux that would otherwise cause significant side writing. In 2005, the first perpendicular magnetic recording (PMR) drives became available. The mono pole write heads were <250 nm wide (compared to today's tape ring heads which are ~3000 nm wide to support read-while-write, which is not an issue for HDD). With PMR, most of the magnetic flux passes straight through the recording layer to the SUL, which conducts it to the low reluctance return pole. Thus, in PMR, head and media magnetics are tightly coupled. A similar situation for PMR tape heads and media seems likely.

This is not to suggest a "free ride" for tape. Among technical challenges are read-while-write and active TDS compensation with a 100-200 nm writer. And while perhaps not optimal, a "wide" (e.g., wider than 200 nm) monopole head might address at least partially the read-while-write requirements. Taking a different tack, the 2018 tape demo (Table 5) achieved a projected areal density of 201 Gb/in^2 on highly oriented sputtered media having an SUL, using a tape ring-type writer. As such, the potential benefits of the SUL were not even fully realized. How well this deals with side writing was not addressed. The question of whether an SUL for tape is even needed for tape was investigated and published in the 2022 400 Gb/in^2 demonstration [Jubert 2022] and again in ref [Tachibana 2023]. The authors found that providing a "cap" layer on the main cobalt recording layer works as well, if not better than, having an SUL (the media was very similar to that used in the 201 Gb/in^2 study). Note also that the head in this demo was a 100 nm HDD side-shielded monopole head. Since measurements were performed using a static tester, tape track following was not addressed, and HMS may be somewhat smaller than is typically achievable with running tape.

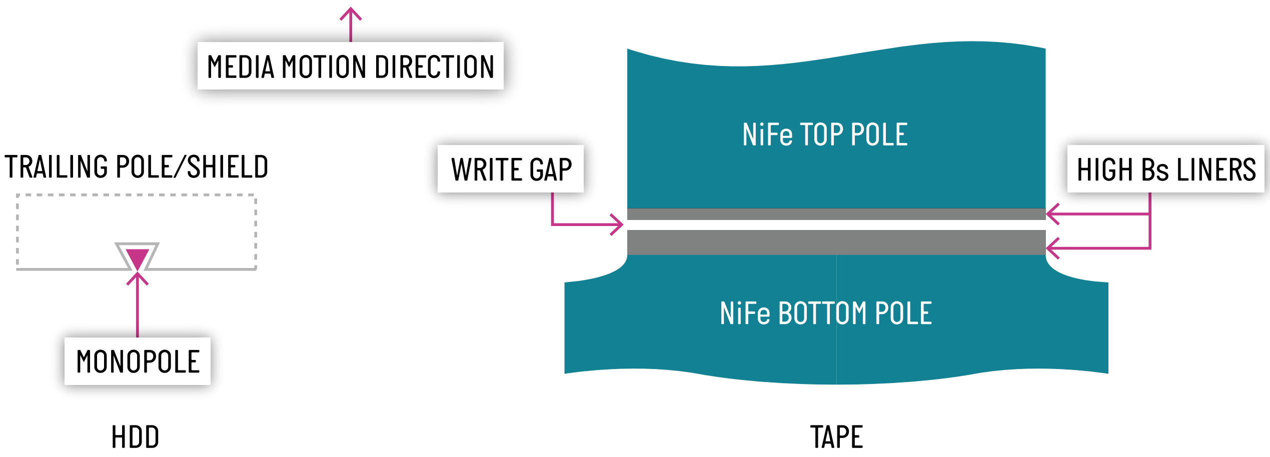

Figure 7. Drawings of typical media bearing surfaces of an HDD shielded monopole writer (left) and a ring tape writer (right).

An area of concern for classic PMR is the potentially higher cost of the sputtered tape media (with or without an SUL), the surface area of tape in a cartridge being roughly 100x greater than in an HDD. The 2021 tape demo examined what is possible with advanced particulate media. This work projected that 317 Gb/in^2 is possible on highly oriented SrFe, indicating that there may be intermediate steps before true PMR is needed. Note however that this paper does not address side writing effects. Other questions for tape are: can monopole write heads achieve higher areal density with particulate media; and is a particulate SUL feasible?

Year	AD Gb/in^2	Media type	Mag layer (nm)	Spacing (nm)	Writer	Method	Write W(nm)/TP(nm)	TMR W (nm)	kfci
2015	123	Oriented BaFe	na	na	Tape/Ring	Tape: Write wide	8,700/140	90(GMR)	680
2018	201	Sputtered + SUL	14	32 (d_{eff})	Tape/Ring/hIB$_S$	Tape: Write wide	8,700/103	48	818
2021	317	Highly oriented SrFe	na	na	Tape/Ring/hIB$_S$	Tape: Write wide	8,700/56.2	29	701
2022	400	Sputtered + cap no SUL	14 + 0-7 cap	10 (W&R)	HDD/Monopole	Static tester	100/66	30	1035

Table 5. Areal density demonstrations

Centerline tolerances

There is yet another area where classic HDD-type PMR could benefit tape. This is regarding position tolerance of the write pole edges responsible for defining the shingled tracks on tape. The top plated pole is used to define the width of a section of the bottom pole adjacent to the recording gap during ion-milling (notching), resulting in potentially more track-to-track variation than a purely photo lithographic process would provide. A small variation, e.g., 25 nm, would be significant when track width is of the order of 100 nm, which is about the width of monopole writers. This will be an area of study for the future.

And finally, current tape drives must accommodate the fact that both write and read head modules are subject to stresses, induced both during and after wafer processing. As a result, the span between the servo readers that flank the transducer arrays could potentially vary from module-to-module, more of a concern as shingled trackwidths shrink. Active TDS compensation may be needed to ensure this does not increase track misregistration effects.

2.3.2.5

Static and Running Friction and HTI

As discussed in the media section of the Roadmap, an on-going goal is reducing media surface roughness enough to help achieve the overall spacing goals but in a way that minimizes friction effects, such as the generation of compression waves in the plane of the moving tape, a phenomenon that can disrupt TBS detection [McClelland et al 2009]. This is chiefly the result of friction at the head skiving edges and was addressed previously by "beveling" the head regions whose only purpose is to provide support for the tape, i.e., have no active transducers. Beveling provides a rounded, slightly roughened (for static friction), surface that eliminates skiving and causes the tape to fly. This is a mechanical finishing process that is applied to the reader module of the 3-module head. In the future, beveling may instead utilize HDD-type patterned MBS and be applied to both reader and writer modules. It may also be possible to reduce friction in the regions containing the recording elements by using novel tape head profiles that reduce skiving edge pressure and minimize the tape-head spacing [Engelen 2016].

It is likely that more wear resistant head coatings, coupled with development of less abrasive media, will be needed to enable on-going reduction of head-tape spacing. This may be jointly addressed by head and media manufacturers working together. As was touched on in Section 2.3.2.3, fabricating read and write heads on separate wafers (as is done today) means that post-wafer head-tape interface processing can be individually tailored for readers and writers. For example, TMR readers may have more pre-recession and/or coating than is needed for writers, for which spacing is more critical. In addition, servo readers in the writer modules can be more recessed because the servo pattern readback bandwidth is significantly lower than for data and so servo spacing loss is much less of a concern.

2.3.2.6

Outlook

Hopefully, this section on future heads has helped to identify some of the challenges and opportunities that media, head, and drive companies will continue to work on together for future tape generations. Heads will continue to leverage HDD head blueprints, including concepts needed for perpendicular heads and media. The outlook is very promising.

2.3.3

Track Mis-registration /Tape Transport Technology

2.3.3.1

Introduction

The mechanical transport area encompasses all technology associated with moving and tensioning the tape in a tape drive and for keeping the read/write transducers registered to the desired track locations on the tape. Modern data storage tape decks must guide the tape accurately while controlling tension as the drive accelerates and decelerates the tape and maintain a constant velocity during write and read operations. Concurrently, they must precisely position recording heads containing multiple read/write elements over the corresponding data tracks on the tape. Modern tape drives also perform on the fly read-while-write verification using a reader module that is positioned downstream of the writer module. This functionality requires that the read transducers are continuously positioned in the 'shadow' of the writers during write operations. Recent generations of LTO and TS11xx tape drives use a three-module head design composed of a 32-channel reader module with 32-channel writer modules mounted on either side of the reader module, with the left module used in the forward direction and the right module in the backwards direction. The introduction of flangeless tape guides to reduce tape edge wear resulted in a significant increase in dynamic tape skew during tape transport and required the introduction of a skew actuator and control loop to follow the dynamic skew of the tape and keep the read transducers in the shadow of the writers during write operations. In addition, recent tape drives, starting with the TS1160 and LTO Gen9 drives, utilize dynamic tension control to compensate for changes in the width of the tape due to tape dimensional stability (TDS) effects. Improvements in the tape transport system as well as the track following, the skew following and the TDS following control systems are critical for achieving the higher track densities necessary to support future capacity requirements. All four of these control systems (reel-to-reel, track follow, skew follow and TDS follow) use information derived from the servo patterns preformatted on the tape during manufacturing. Hence the accuracy and update rate of this information plays a critical role in the performance of these control systems.

The tape is partitioned into alternating servo and data bands with the array of read and write elements in the recording head spanning a single data band. Servo elements at the top and bottom of the transducer arrays sense the position of the head within the data band and enable precise control of the head position relative to the tape. Within a data band, there are many more tracks than head elements. Each head element will write to or read from a subset of the total tracks in a data band called a sub-data band. Data is written and read in a bi-directional serpentine fashion with each forward and reverse pass of the media across the head referred to as a wrap. In recent tape drives, the track following control system uses the average of the position measured from the two servo patterns that straddle the data bands. Hence the track following controller is effectively controlling the position of the head referenced to the center of the data band. During write and read operations, the width of the data band is measured by taking the difference in the lateral position measured from the two servo patterns and the TDS control system adjusts the tension applied to the tape to match the width of the data band to the span of the head. Errors in track following control affect the positioning of all 32 head elements in a similar manor, whereas errors in TDS compensation have a negligible impact on the positioning of transducers near the center of the head and an increasingly large impact on the positioning of transducers towards the top and bottom of the array. The positioning error for any given transducer is impacted by the sum of these two components that can be additive or subtractive.

The roadmap goals laid out in Section 2.1 have several ramifications for tape guiding and motion control. The primary impact will be felt with increasing the track density to 322,248 tracks per inch which is equivalent to a 79 nm track pitch. With multiple head elements reading and writing simultaneously, each read/write element must be accurately positioned over its respective track within a given data band and wrap location. Many factors come into play in ensuring good tracking, including variation in the element pitch, tape expansion or contraction, lateral tape motion, position measurement accuracy, as well as the width of the track and the width of the reader. As track width shrinks, each of these factors must be carefully controlled or compensated for. The consequences of not budgeting for all the factors contributing to track mis-registration could be writing off track, longer store/restore times due to having to reread or rewrite data, or the inability to recover data.

The transport section will discuss multiple solutions that could be implemented in future tape systems. Although not all of the following improvements may be required, many of these will likely be needed over the coming decade. The two dominant factors that contribute to track mis-registration in state-of-the-art tape drives are track following accuracy and TDS compensation accuracy. Both rely on information derived from the servo patterns formatted on tape and hence the accuracy and update rate of the servo information also plays a critical role. Variation in head element pitch between modules and variations in the distance from one head element to another within the recording head are currently less critical but will play an increasingly important role as track pitches are scaled to smaller and smaller dimensions.

2.3.3.2

Tape Tension

As tape gets thinner to enable longer tape lengths, the reels will have to accommodate increasing numbers of winds; therefore, to keep pack stress from increasing unacceptably, tension will have to be reduced. Better methods of managing reel deformation under increasing stress may be needed. Media dimensional changes as the tape is packed, stored, and then unpacked will need to be minimized and compensated for. Recent generations of tape drive that utilize dynamic tension control to compensate for changes in the tape width (TDS) also implement a low-tension unload (LTU) functionality to wind the tape back into the cartridge under a constant low tension immediately before unloading the cartridge from the drive. Storing the tape under low tension reduces changes in tape width due to creep. To minimize the impact of LTU on cycle time, the tape drive keeps track of the furthest point accessed along the length of tape and then before unloading the tape, the drive seeks to just beyond that point, reduces tension to the storage tension, and then rewinds the tape back into the cartridge reel. The tension-based TDS compensation scheme used in state-of-the-art tape drives has a relatively low bandwidth and is effective to compensate environmental and creep effects but does not have sufficient bandwidth to compensate for higher frequency variations such as periodic tension variations arising from imperfections in the tape reels and tape pack. In the future these effects will need to be reduced or compensated for through improvements in the reel-to-reel control system [Cherubini 2018], improvements to the tape reels and reel motors and/or a higher bandwidth TDS compensation scheme.

2.3.3.3

Tape Speed

Some of the operating points in the roadmap will require relatively high tape speeds to achieve the target data transfer rate. Higher tape speeds increase the frequencies at which periodic disturbances appear and may also increase the magnitude of lateral tape motion (LTM) due to dynamic interactions with tape reel flanges and/or tape path guides. One technology under investigation to reduce interactions between the tape and the reel flanges is the use of active tape guiding with actively controlled tilting rollers to steer the tape [Pantazi 2010], [Yang 2015]. The increase in rotation frequency of the tape guide rollers with tape speed increases the frequency and amplitude of disturbances from the rolling guides themselves. This frequency increase can potentially be reduced or prevented by increasing the diameter of the guide rollers, however, improvements will still also likely be needed in the precision of the rolling guides. In order to improve track following at high speeds, the actuator and track following control loop bandwidth will also likely have to increase to compensate for these additional disturbances. The bandwidth of the track following control loop can potentially be increased with synchronous control [Ebermann 2021] and by increasing the update rate of the servo pattern. The accuracy of the position estimated from the servo pattern is also important for track following control and can be improved by increasing the angle of the servo pattern stripes. However, for a fixed pattern height, this has the penalty of decreasing the pattern update rate and hence increasing delay in the control loop [Furrer 2015a], [Ebermann 2021]. Hence future servo pattern design will require careful optimization to balance between improvements in measurement accuracy with increased delay.

2.3.3.4

Tape-Head Friction

Another important consideration is the impact of tape head friction on track following servo performance. To achieve the linear density scaling targets of the road map it will be necessary to reduce magnetic spacing by making the tape surface smoother, which in turn will likely result in increased tape-head friction. Increased tape-head friction can excite compressional waves in the tape as it is streamed over the head, giving rise to high frequency velocity variations that show up as position measurement noise and degrade track following accuracy [Furrer 2018a]. To minimize this effect, it will be necessary to minimize friction increases by careful engineering of the tape surface roughness and through the development of improved tape lubricant and low friction head technologies.

2.3.3.5

Tape Dimensional Stability

The single channel tape areal density demonstrations that are periodically done to demonstrate the scaling potential of tape primarily focus on the recording capability of the media and track following accuracy but largely ignore other issues related to the track mis-registration budget such as TDS and head tolerances. Commercial tape drives with multiple parallel channels must deal with variations caused by head element pitch tolerances as the tape is interchanged from one drive to another. In addition, the tape substrate is not rigid and therefore the width of the tape and relative locations of written data tracks change under different environmental conditions. The tension based active TDS compensation scheme used in state-of-the-art tape drives enables linear TDS and linear head span mismatch effects to be compensated. However, research is required to determine if there are components of these effects that are non-linear, and if so, to develop schemes to deal with these components.

To ensure that the read heads can effectively read the data a long time after it was written, the track mis-registration budget must have sufficient margin to ensure that all the readers are accurately positioned over the written tracks as the tape is streaming by at high speed. Each of the variables that contribute to track mis-registration must be reduced. There are several basic contributors: tolerance variation between the write heads and read heads, tape dimensional stability, lateral tape motion, and the servo systems to minimize both the TDS error and the position error.

One contributor that is becoming increasingly important is the tolerance variation between the write heads in the drive that originally wrote the data on the tape and the read heads in the drive that will eventually read the data. In this context critical tolerances include variations in the width and position of the transducers on the head that arise from the manufacturing process. To keep all the read transducers in a head accurately positioned on tracks that are only 79 nm wide, head manufacturing processes will need to continually improve to improve these tolerances.

2.3.3.6

Track Following

Another contributor to track mis-registration is the residual dynamic error between the actual tape/head position and the desired position. This is known as the position error signal (PES). This error typically increases as the tape speed increases. This is one of the factors that limits the speed at which the tape can be effectively transported during reading and writing. To sufficiently reduce the PES, all disturbances to the tape from the transport system must be minimized. Care needs to be taken to avoid high frequency disturbances. Improvements may be required in the reel flange geometry of both the cartridge and the machine reels to improve their precision. Disturbances from the tape guides themselves must also be further reduced. With rolling element guide systems, the disturbances increase as the tape speed increases so improvements are likely to be required to counter that increase.

The effective bandwidth of the head servo systems will require continued increases to reduce the residual PES for each new generation of tape drive. This is particularly true with the compounding effects of increased tape speed and track pitch reductions. Possible approaches to improvements in this area were discussed in the previous roadmap and remain relevant today.

2.3.3.7

Servo Formatting

The quality of the servo tracks written by the media manufacturers is also a critical factor which contributes to the overall PES. Important aspects include minimizing the position and velocity noise that are written into the servo pattern as well as providing higher signal-to-noise ratio (SNR) for servo signal readback. Significant improvements have been made in servo writing capability over the last several generations of tape drives and this should continue to be addressed with every new generation. To improve the accuracy of the servo position estimation achieved by the servo channel/demodulator, the servo stripe angles were increased in recent tape formats. The spacing between servo frames was also reduced to provide a higher rate of position updates and hence reduce the delay in the control loop. This has enabled higher bandwidth head positioning servo systems. There is also potential to further improve the quality/sharpness of the written servo patterns by adopting thin film technology in the manufacture of servo format heads [Engelen 2012]. Future tape operating points will likely require additional improvements to the servo pattern design to further increase the measurement accuracy and update rate.

2.3.3.8

Technology to Enable High TPI Recording

Recent tape areal density demonstrations have shown progressive improvements in track following accuracy through incremental improvements to the various components of the track following system [Lantz 2015], [Furrer 2018b], [Furrer 2021]. The most recent of these demonstrations [Furrer 2021] reported a track following accuracy characterized by a standard deviation of the PES ≤ 3.2 nm over a tape speed range of 1 - 4 m/s. This result provides confidence that the position measurement accuracy and the track-following accuracy required for sub 100 nm track pitches are possible on flexible tape media. However, it may still be challenging to implement or adapt the lab technologies used for this work for use in commercial tape drives manufactured at scale.

Tape dimensional stability (TDS) is another significant contributor to the track mis-registration budget. Tension based active TDS compensation was introduced in recent products to deal with this, however, the technique has limited compensation range due to the limited range over which tension can be varied. Hence there is a need to develop TDS compensation methods with more compensation range and to investigate countermeasures to reduce TDS. The most effective solution to reduce TDS would be to move to an aramid substrate which could potentially provide a 2x to 4x reduction in TDS compared to conventional PEN and PET substrates. Although this is a proven solution that has been implemented in enterprise class tape drives, the disadvantage is the significant increase in cost and the limited manufacturing capacity. The use of Aramid should continue to be considered as a possible solution.

Another very effective countermeasure would be to reduce the overall span of the write/read transducer array by increasing the number of data bands. However, this creates very difficult challenges for the head manufacturing process; especially if that is coupled with a doubling or quadrupling of the number of channels as projected by the roadmap. The other challenge that this approach presents is the difficulty in reading previous generations of written tapes because the heads are no longer aligned with the written tracks of the previous generation format. To address this, additional heads would be required for backward compatibility which significantly increases head manufacturing cost and head channel interconnect complexity. One possibility that may need to be considered to continue to progress tape technology areal density is to forfeit backward compatibility for a generation to enable a leap forward. Such format changes or breaks have occurred multiple times during the 70+ year history of digital tape but are undesirable from the point of view of some tape customers. However, for other tape customers such as hyperscale cloud companies who have begun to use tape technology on a large scale, backwards compatibility is generally not considered important.

Tape technology development must also account for the growing trend to build green data centers which emphasize optimal energy efficiency, operational cost reduction, and minimal environmental impact. Since tape storage is the most energy efficient data storage technology, one might think this trend is beneficial. Unfortunately, the emphasis on energy efficiency has resulted in data center designs with wider temperature operating ranges and little-to-no humidity controls. Temperature and humidity variation throughout the year, especially between the initial write and later read back on a tape drive, is a concern from a tape dimensional stability (TDS) and reliability standpoint. It is also important to realize that more data centers are keeping their tape archives on-site in tape libraries and not shipping cartridges to off-site, environmentally controlled storage facilities like Iron Mountain.

Future drive and media development need to focus on making tape storage more robust to varying temperature and humidity conditions. Drives are now modulating the tape tension to compensate for environmental expansion and contraction. In the future, the range of compensation will need to increase, either through improvements in tension-based compensation or the development of new active TDS compensation techniques. Media manufacturers have previously tried to minimize the transverse dimensional variation caused by inadvertent tension variation in the drive. However, it is now desirable to increase that sensitivity as much as possible to enable more tape width compensation for the same amount of tension modulation. This approach would however necessitate a reduction in high frequency tension variations. Efforts to reduce the humidity related TDS contribution of the current PET and PEN base films should continue as in previous generations. One area to consider is the modification of the tensilization process, whereby the lateral elasticity modulus is traded off against the machine direction modulus by stretching the substrate as it is manufactured. This can reduce the transverse dimensional change due to humidity variations but increases the transverse dimensional change caused by tension variations in the tape transport.

Other methods for dynamic compensation of media expansion and contraction should also be considered. If tracks are written onto the tape with the head at a small initial azimuth angle, then subsequent head azimuthal rotation can increase or decrease the effective span of the transducer array as the media expands or contracts. One of the challenges of this method is the possible amplitude loss and bit stretching that occurs as the head gaps no longer align with the transitions on the tape. The impact of this problem is reduced as the bit aspect ratio shrinks as the track width is expected to decrease faster than the bit length. Moreover, if the initial rotation angle is increased, the gain of this TDS compensation increases, i.e. a smaller change in angle is required to compensate for a change in tape width. If the angular changes used for TDS compensation are less than +/- 1 degree and the reader width is less than one micron, the impact of dynamic skewing on the read channel is expected to be negligible.

Heads that are capable of expanding or contracting to compensate for TDS have been envisioned but a practical implementation remains elusive. Piezo elements could conceivably expand or contract portions of a head relative to the remaining elements. Alternately, heating elements could achieve the same thing. However, it may take 10mm or more of piezo material in a head to stretch its span by 5-10µm when actuated. Higher voltages could reduce the required length of the piezo material, but this risks damaging the sensitive TMR elements with electrical overstress. Heating elements could be employed, but their role in stain formation and media damage remains largely unknown. More research is required to determine if either of these approaches is feasible.

To achieve the areal density and associated TPI required for the continued data capacity growth of magnetic tape, all of the errors that contribute to the track mis-registration budget must be reduced. This is an interdisciplinary issue that requires continued investment by drive, media, and head manufacturers to be successful.

2.3.4

Error Correction and Channel Technologies

Since the first-generation of linear-tape-open (LTO) cartridges were introduced in 2000 with a capacity of 100 GB, the storage capacity of LTO tape cartridges has increased by a factor of 180 and data rates have increased by 20 times. Over the same period, the specified uncorrectable bit error rate (UBER) of LTO cartridges has also improved by a factor of 1000, i.e. three orders of magnitude improvement.

2.3.4.1

Uncorrectable Bit Error Rate (UBER)

During the first 15 years of LTO drives, up until the seventh generation introduced in 2015, the probability of encountering an error in a cartridge, although very small, has continually increased simply because cartridge capacities steadily increased while the End-Of-Life (EOL) uncorrectable bit error rate (UBER) remained fixed. Calculations of the UBER in a modern tape drive are based on theoretical analysis because such events are so rare. As a result of deep interleaving in tape storage, the theoretical models used for analysis are based on the binomial distribution of raw byte errors [Arslan 2022], an assumption that has been experimentally verified [Lantz 2015, Furrer 2021]. More complex reliability models which are based on the theory of renewal processes can account for correlated errors and defective header and synchronization fields [Arslan 2014]. Furthermore, analytical expressions for the error probability of interleaved block codes with burst errors have been derived in [Cideciyan 2019].

To address the increased probability of encountering an error in a cartridge that results from capacity scaling, the error correction code (ECC) used in LTO-7 was improved to provide enhanced error correction performance resulting in an UBER of 1e-19. The most recent Linear Tape Open 9 (LTO-9) format provides a native capacity of 18 TB and offers another 10x improvement in error correction capability over LTO-7 and LTO-8, which corresponds to an UBER of 1e-20, or equivalently one error per 12.5 Exabytes read [Arslan 2022]. This is comparable to the performance achieved by enterprise class tape drives such as the IBM TS1160 drive with a 20 TB native cartridge capacity and the latest IBM TS1170 product with 50 TB native cartridge capacity. Another improvement introduced in LTO-7 tape drives significantly reduced the area on tape used for rewritten data by collecting all the encoded subunits of a 6 MB Data Set (12 MB Data Set in LTO-9), which satisfy the rewrite condition after decoding the encoded subunits following a read-while-write operation and rewriting them in one batch at the end of the Data Set.

2.3.4.2

Error Correction Code and Format

The error-correction coding scheme used in tape storage is based on a product code with two Reed-Solomon (RS) component codes referred to as C1 (row code) and C2 (column code). In the following, the C1 and C2 RS codes are represented using standard notation of (n, k, d) where n represents the block length, k is the message length and d is the distance. When compared to the LTO-1 format, the main improvements in the LTO-7 format enabling the improved EOL UBER are more powerful RS(249, 237, 13) C1 and RS(96, 84, 13) C2 codes, an increase in the depth of interleaved C1 code words written on tape tracks, and a new 32-track tape layout that increases the decorrelation of byte errors at the inputs of the C1 and C2 decoders. LTO-9 introduced a further improved C2 format of RS(192, 168, 25), where both the code word length and the number of parity symbols have doubled compared to LTO-7. The LTO-9 C2 code maintained the LTO-7/8 C2 code rate of 0.875 to ensure that 4 dead tracks out of 32 tracks can be corrected with deep interleaving in a Data Set, but tolerates significantly increased error rates at the C2 decoder input. For LTO-9, this ECC performance gain was key to achieving both a lower UBER number of 1e-20 and a higher areal density by reducing the requirements for signal-to-noise ratio (SNR) at the input of the data channel.

In the future, we expect continued improvements in the ECC schemes used to encode and decode data and the data channel architecture used in tape storage. These performance gains could be used to provide further reductions in EOL UBER at a constant level of data channel SNR, or, for a constant EOL UBER they could be used to reduce the requirements for data channel input SNR. Although ECC improvements could be used to achieve simultaneous but smaller improvements of both areas, we expect that future scaling in tape storage will focus on the latter, i.e. the ECC and channel gains are expected to be primarily used to achieve higher areal densities and cartridge capacities. Therefore, in the 2024 Roadmap, we take the current state-of-the-art EOL UBER of 1e-20 as the starting point and project a constant error rate in the future. At the 45 TB cartridge capacity assumed at the start of the 2024 Roadmap, on average, only one out of 277,777 tape cartridges will contain an uncorrectable error event due to ECC failure. For the projected 723 TB cartridge capacity in 2034, an EOL UBER of 1e-20 corresponds to one uncorrectable error event in 17,296 cartridges.

Taking full advantage of the ECC improvements to achieve higher areal densities, i.e. the higher error-rates that can be tolerated at the ECC decoder input, requires improvements to the read channel as well. Technologies to enhance the read channel robustness and performance for operation at reduced SNR requirements at the input of the channel are discussed below.

2.3.4.3

Channel Architecture and Technologies

The architecture currently used to encode and decode data in tape storage is forward concatenation wherein compressed and optionally encrypted data is first encoded by a two-dimensional product error correction code and then encoded by a modulation code that satisfies various run-length constraints on the maximum length of zeros [Immink 1990], transition-runs [Moon 1996], synchronization patterns used for timing acquisition (alternating-bit pattern in NRZI notation) [Cideciyan 1992] and twins pattern to limit the path memory of the Viterbi detector [Cideciyan 2001]. Using measured data from a tape drive it has been shown that hard-decision based iterative decoding of product codes allows operation at an SNR of about 10.5 dB [Furrer 2021, Lantz 2015]. For the ECC code rate of 0.83 in LTO-8, computation of the channel capacity of the magnetic tape channel modeled as a discrete symmetric memoryless channel [Cideciyan 2017] has shown that the highest byte-error rate at the input of the ECC decoder allowing error-free retrieval of data at the output of the ECC decoder is about 1e-1 assuming unrestricted ECC codewords length and decoding complexity. For an RS(240, 228, 13) C1 code and an RS(192, 168, 25) C2 code, three full rounds of C1/C2 iterative hard-decision decoding of two-dimensional product codes achieve UBER performance of 1e-20 if the byte-error rate at the input of the ECC decoder is 5e-2 or less [Furrer 2021, Furrer 2015b]. The latest generations of IBM TS1160/70 enterprise and LTO-9 drives implement iterative ECC decoding which can perform additional C1/C2 decoding steps to boost the error correction performance in both streaming mode and error-recovery mode. Increasing the symbol size and the block size of Reed-Solomon (RS) C1 codes [Cideciyan 2019] and C2 codes can further improve the ECC performance to allow operation at lower SNR or optionally lower EOL UBER. Three-dimensional product codes are also a practical approach for providing a modest improvement in data reliability while maintaining the excellent burst error-correction capability required in magnetic tape storage [Cideciyan 2017].

Reverse concatenation (RC) of error-correction coding and modulation coding [Blaum 2007, Mittelholzer 2008] allows the use of very high-rate modulation codes [Mittelholzer 2009, Mittelholzer 2014] which have large block sizes, an overhead of less than 1% and satisfy required constraints of a tape channel. In a RC architecture, error propagation at the output of the modulation decoder is not an issue because modulation decoding is performed after ECC decoding. Moreover, because soft information can be passed between the detector and the ECC decoder, RC schemes allow the use of soft-decoding and iterative detection/decoding techniques to improve the error-rate performance [Galbraith 2010]. The use of partial reverse concatenation [Cideciyan 2014] based on employing a higher rate (≈0.991) modulation code and replacing the C1 code by a lower rate (≈0.94) low-density parity check (LDPC) code, has also demonstrated reliable operation at an SNR of about 10.5 dB using channel measurements from a tape drive without making iterations through the C2 RS decoder. A new reverse concatenation scheme that keeps the inner C1 RS code but replaces the outer C2 RS code by an LDPC code [Oh 2016] promises further performance improvements by performing hybrid erasure/soft decoding. Another promising hybrid approach that provides at least an order of magnitude improvement in UBER by combining a List-Viterbi architecture with noise predictive maximum likelihood detection is described in [Arslan 2013a].

2.3.4.4

Areal-Density Scaling and Low SNR Challenges

The 2024 Roadmap predicts a 27.76% CAGR in areal density scaling, which leads to SNR losses unless they are compensated for by improved recording technologies such as improved write and read heads, higher-SNR media, and further enhancements in read channel signal processing, detection and decoding. As track-widths and bit lengths are decreasing, smaller magnetic particles with tighter distributions, a smaller mag layer thickness, and lower head-media spacings are necessary to maintain acceptable SNR levels. Smaller bit sizes furthermore necessitate smaller non-magnetic spacer and abrasive particles to avoid long signal dropouts and burst errors.

Tape read channels require a high level of adaptivity to cope with media variations such as fluctuations in particle dispersion and magnetic layer thickness, variations in the head-media spacing, head and media wear, head-tape interface related friction and velocity jitter. To deal with these temporal and spatial variations, the majority of read channel signal processing, detection, and decoding algorithms run in a highly adaptive and dynamic manner. Examples include adaptive equalization, timing recovery/synchronization, fast gain control, noise predictive maximum likelihood sequence detection, dynamic error- and erasure decoding. The parallel track recording nature of tape drives in combination with cartridge interchange require both extra SNR margins and adaptability.

Under low SNR operating conditions, timing recovery in the read channel becomes challenging and can lead to cycle slips in which one or more erroneous bits are either inserted into or omitted from the data stream, which in turn leads to long burst errors. Fortunately, several technologies are currently under investigation that could alleviate this challenge. One example utilizes cycle slip detection and correction through classification of modulation code failures [Arslan 2013b]. Another promising approach to cycle slip mitigation involves exploiting the parallel timing information available in a multichannel tape drive [Ölçer 2008]. Finally, a robust timing recovery technique for low-SNR tape read channels, which fully exploits the parallel-track recording nature of linear tape drives and significantly reduces the loss-of-lock rate compared to the conventional second order phase-locked loop approach, with only a small increase in implementation complexity was described in [Furrer 2018c].

Magnetic recording nonlinearities such as non-linear transition shift (NLTS) caused by magnetostatic interactions between adjacent transitions and hard transition shift (HTS) are expected to grow with increasing linear densities. Although the 2024 Roadmap predicts only a modest 5% linear density scaling per year, a combination of write-side pre-compensation of nonlinearities as well as enhanced adaptive signal processing techniques to mitigate non-linear effects in the read channel will become increasingly important.

Read-while-write verification is used in modern tape drives to improve reliability and data integrity. The 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 read-back signals from each reader is processed by the detector and C1 decoder and if the C1 errors in a given channel exceed a threshold level, the data is rewritten to a new location further down tape. This approach enables media defects to be detected on the fly and the data which was written in such areas to be re-written to a new defect free location. Unfortunately, using current rewrite criteria, if the SNR is decreased towards a 10.5 dB operating point, the amount of data that would be rewritten on tape would increase drastically.

For example, the iterative decoding scheme described above can enable an UBER of 1e-20 at detector error rates as high as 5e-2, corresponding to about 10.5 dB of SNR. However, under such high raw error rate conditions, the error rate after the first round of C1 decoding would be large, and hence would result in a drastic increase in rewrites using current rewrite criteria. Various technologies are currently being investigated to alleviate this challenge. Increasing the symbol size, the block size and the interleaving depth of C1 codes results in a stronger C1 code and reduces the probability of rewrite. Rewrite tables with reduced restrictions and dead track detection can be used to increase the rewrite efficiency. Priority and threshold based rewrite concepts can substantially reduce the amount of rewrites, but at the same time must keep the risk of data loss at a minimum.

We believe that the technologies discussed above remain particularly important areas for continued research because lowering the operating point towards an SNR of 10.5 dB could enable an increase in areal density of 2x or more without any changes to the media.

2.4
Acknowledgements

INSIC wishes to express its appreciation to the team members from the tape storage community, and their companies, who contributed to the technogy roadmap discussion. In particular, we would like to thank Toyobo, Toray and Dowa for inputs into the media technology section.

2.5
References

[Furrer 2021] S. Furrer et al., “317 Gb/in^2 recording areal density on strontium ferrite tape,” IEEE Trans. Magn., vol. 57, no.7, Apr. 2021, Art. ID 3101311.

[Furrer 2021] Simeon Furrer et al., “317 Gb/in^2 Recording Areal Density on Strontium Ferrite Tape”, IEEE Transactions on Magnetics, JULY 2021, 57, 7, pp. 1-11, DOI: 10.1109/TMAG.2021.3076868

[Furrer 2018] S. Furrer et al., “201 Gb/in^2 recording areal density on sputtered magnetic tape” IEEE Trans. Magn., vol. 54, no.2, Feb. 2018, Art. ID 3100308.

[Tachibana 2023] J. Tachibana et al., “Challenge of Media Noise Suppression with Oxidation Control in Granular Structure for CoPtCr-based Sputtered Tape towards 400 Gb/in^2,” IEEE Trans. Magn., vol. 59, no. 11, Nov. 2023, Art. ID 2100305.

[Tachibana 2021] J. Tachibana et al., “Ru underlayer elimination for CoPt-based sputtered tape with perpendicular recording toward 330 TB per cartridge product,” IEEE Trans. Magn., vol. 57, no.2, Feb 2021, Art. ID 3200305.

[Jaquette 2003], “LTO: a better format for mid-range tape,” IBM J. Res. & Dev., vol. 47, no. 4, pp. 429-444, July 2003.

[Biskeborn 2003] and J. H. Eaton, “Hard disk drive technology flat heads for linear tape recording,” IBM J. Res. & Dev., vol 47, no. 4, pp. 385-400, July 2003.

[Biskeborn 2012] et al, “Head and interface for high areal density tape recording,” IEEE Transactions on Magnetics, vol. 48, no. 11, pp. 4463-4466, November 2012.

[Biskeborn 2008], et al, “Crosstalk between write transducers,” IEEE Transactions on Magnetics, vol. 44, no. 11, November 2008.

[Biskeborn 2018] et al, “TMR tape drive for a 15 TB cartridge,” AIP Advances 8, 056511 (2018).

[Bertram 1994], “Theory of Magnetic Recording,” Cambridge University Press, Cambridge, England, 1994, pp. 182-186.

[Marchon and Olson 2009], “Magnetic spacing trends: from LMR to PMR and beyond,” IEEE Transactions on Magnetics, vol. 45, no. 10, October 2009.

[Lantz et al 2015], “123 Gb/in^2 recording areal density on barium ferrite tape,” IEEE Transactions on Magnetics, vol. 51, no. 11, November 2015.

[Furrer et al 2017], “201 Gb/in^2 recording areal density on sputtered magnetic tape,” IEEE Transactions on Magnetics, vol. 54, no. 2, July 2017.

[Furrer 2021], “317 Gb/in^2 recording areal density on strontium ferrite tape,” IEEE Transactions on Magnetics, vol. 57, no. 7, April 2021.

[Ohkoshi 2015] T. Shin-ichi Ohkoshi et al. “Nanometer-size hard magnetic ferrite exhibiting high optical-transparency and non-linear optical-magnetoelectric effect” Scientific Reports 5, Article number: 14414 (2015).

[Jubert 2022], “Evaluation of sputtered tape media with hard disk drive components,” IEEE Transactions on Magnetics, vol. 58, no. 4, April 2022.

[Tachibana 2023], “Challenge of media noise suppression with oxidation control in granular structure for CoPt-Cr-based sputtered tape toward 400Gb/in^2,” IEEE Transactions on Magnetics, vol. 59, no. 11, November 2023.

[McClelland et al 2009], “Effect of tape longitudinal dynamics on timing recovery and channel performance,” IEEE Transactions on Magnetics, vol. 45, no. 10, pp. 3585-3589, October 2009.

[Engelen et al 2016], “Tape head with sub-ambient air pressure cavities,” IEEE Transactions on Magnetics, vol. 52, no. 11, November 2016.

[Jubert 2022] P.-O.. Jubert et al., “Evaluation of Sputtered Tape Media With Hard Disk Drive Components,” IEEE Trans. Magn., vol. 58, no. 4, Apr. 2022, Art. ID 3100705.

[Ohkoshi 2020] Shin-ichi Ohkoshi et el., “Magnetic-Pole Flip by Millimeter Wave”, Advanced Materials, vol. 32, no. 48, Dec. 3 2020, 2004897

[Engelen 2016] J. B. B. Engelen et al., “Tape–Head With Sub-Ambient Air Pressure Cavities” IEEE Trans. Magn., vol. 52, no.11, Nov. 2016, Art. ID 3301810.

[Cherubini 2018] Giovanni Cherubini et al., “Feedback control of transport systems in tape drives without tension transducers”, Mechatronics, Volume 49, February 2018, Pages 211-223, DOI: https://doi.org/10.1016/j.mechatronics.2018.01.003

[Pantazi 2010] A. Pantazi et al., “Active Tape Guiding” Proc. 20th ASME Annual Conf. on Information Storage & Processing Systems 2010 “ISPS,” Santa Clara, CA, (June 2010) 304-306.

[Yang 2015] Hankang Yang et al., “Mechanics of Lateral Positioning of a Translating Tape due to Tilted Rollers: Theory and Experiments”, International Journal of Solids and Structures 66 (2015) 88–97 http://dx.doi.org/10.1016/j.ijsolstr.2015.03.029

[Ebermann 2021] Patrick Ebermann et al., “Track-following system optimization for future magnetic tape data storage” Mechatronics 80 (2021) 102662, DOI: https://doi.org/10.1016/j.mechatronics.2021.102662

[Furrer 2015a] S. Furrer et al., “Resolution Limits of Timing-Based Servo Schemes in Magnetic Tape Drives”, IEEE TRANSACTIONS ON MAGNETICS, VOL. 51, NO. 11, NOVEMBER 2015, 3101504, DOI: 10.1109/TMAG.2015.2436718

[Furrer 2015b] S. Furrer et al., “85.9 Gb/in^2 Recording Areal Density on Barium Ferrite Tape,” in IEEE Trans. Magn., vol. 51, no. 4, April 2015, Art no. 3100207.
[Furrer 2018a] Simeon Furrer et al. “Compressional Wave Disturbance Suppression for Nanoscale Track-Following on Flexible Tape Media”, 2018 Annual American Control Conference (ACC), pp 6678 – 6683, DOI: 10.23919/ACC.2018.8431805

[Engelen 2012] Johan B. C. Engelen et al., «Planar Thin-Film Servo Write Head for Magnetic Tape Recording” IEEE TRANSACTIONS ON MAGNETICS, VOL. 48, NO. 11, NOVEMBER 2012, p3539. DOI: 10.1109/TMAG.2012.2200244

[Lantz 2015] Mark A. Lantz et al., “123 Gb/in^2 Recording Areal Density on Barium Ferrite Tape” IEEE Transactions on Magnetics, VOL: 51, Issue: 11, Nov. 2015

[Furrer 2018b] Simeon Furrer et al., “201 Gb/in^2 Recording Areal Density on Sputtered Magnetic Tape” IEEE Transactions on Magnetics, VOL: 54, Issue: 2, Feb. 2018.

[Arslan 2022] S. Arslan et al., “LTO-9 Technology and User Data Reliability Analysis,” LTO Consortium White Paper, Aug. 2022, available online: https://www.lto.org/wp-content/uploads/2022/08/LTO-UBER-Technical-Paper-August-2022.pdf

[Arslan 2014] S. S. Arslan, et al., “MDS Product Code Performance Estimations Under Header CRC Check Failures and Missing Syncs,” IEEE Trans. on Device and Materials Reliability, vol.14, no. 3, Sep. 2014.

[Cideciyan 2019] R. D. Cideciyan, S. Furrer and M. A. Lantz, “Performance of Interleaved Block Codes With Burst Errors,” in IEEE Trans. Magn., vol. 55, no. 3, March 2019, Art no. 3100305.

[Immink 1990] K. A. S. Immink, “Runlength-limited sequences,” Proc. IEEE, vol. 78, pp. 1745–1759, Nov. 1990.
[Moon 1996] J. Moon and B. Brickner, “Maximum transition run codes for data storage systems,” IEEE Trans. Magn., vol. 32, pp. 3992–3994, Sep. 1996.

[Cideciyan 1992] R. D. Cideciyan et al., “A PRML system for digital magnetic recording,” IEEE J. Select. Areas Commun., vol. 10, no. 1, pp. 38–56, Jan. 1992.

[Cideciyan 2017] R. D. Cideciyan, S. Furrer and M. A. Lantz, “Product Codes for Data Storage on Magnetic Tape,” in IEEE Trans. Magn., vol. 53, no. 2, Feb. 2017, Art no. 3100410.

[Blaum 2007] M. Blaum et al., “High-rate modulation codes for reverse concatenation,” IEEE Trans. Magn., vol. 43, no. 2, pp. 740-743, Feb. 2007.

[Mittelholzer 2008] T. Mittelholzer and E. Eleftheriou, “Reverse concatenation of product and modulation codes,” Proc. IEEE Int. Conf. Commun., pp. 1991–1995, May 2008.

[Mittelholzer 2009] T. Mittelholzer, “Enumerative maximum transition run codes,” Proc. IEEE Int’l Symp. On Information Theory, ISIT 2009, Seoul, Korea, pp. 549-1553, June 2009.

[Mittelholzer 2014] T. Mittelholzer and R. D. Cideciyan, “Enumerative modulation codes based on sliding-window substitutions,” Proc. IEEE Int’l Symp. on Inform. Theory, Honolulu, USA, July 2014.

[Galbraith 2010] R. L. Galbraith et al., “Architecture and Implementation of a First-Generation Iterative Detection Read Channel,” in IEEE Trans. Magn., vol. 46, no. 3, pp. 837-843, March 2010.

[Cideciyan 2014] R. D. Cideciyan, R. Hutchins, T. Mittelholzer and S. Ölçer, “Partial reverse concatenation for data storage,” Signal and Information Processing Association Annual Summit and Conference (APSIPA), 2014 Asia-Pacific, Cambodia, 2014.

[Oh 2016] J. Oh, J. Ha, H. Park, and J. Moon, “RS-LDPC Concatenated Coding for the Modern Tape Storage Channel,” IEEE Trans. on Commun., vol. 64, no. 1, Jan. 2016.

[Arslan 2013a] S. S. Arslan, J. Lee, and T. Goker, “Error Event Corrections Using List-NPML Decoding and Error Detection Codes,” IEEE Trans. Magn., vol. 49, no. 7, July 2013.

[Arslan 2013b] S. S. Arslan, J. Lee, and T. Goker, “Cycle Slip Detection and Correction Through Classification of Modulation Code Failures,” IEEE Trans. Magn., vol. 49, no. 9, Sep. 2013.

[Ölçer 2008] S. Ölçer, J. Jelitto, and R.A. Hutchins, “Global Timing Control for Tape Storage Channels,” Digest of IEEE Int’l Magnetics Conf. “Intermag 2008”, Madrid, Spain, May 5-8, 2008.

[Furrer 2018c] S. Furrer and M. A. Lantz, “Timing recovery for low-SNR magnetic tape recording channels,” in Digest of The 29th Magnetic Recording Conference (TMRC 2018), Milpitas, CA, Aug. 2018

INSIC
Information Storage
Industry Consortium

The role of the Information Storage Industry Consortium (INSIC) is to serve as a resource for anyone interested in current and future trends in digital magnetic tape recording technologies and platforms, which are an important component of the data storage landscape. A primary goal is to make this information easily accessible and valuable for existing and prospective tape users. Participants today include corporations involved in creating and producing tape recording hardware, media, systems and infrastructures. INSIC was originally incorporated in April 1991 as NSIC, a nonprofit mutual benefit corporation, and is currently located in Monroe, Virginia, USA.

Read more:
http://www.insic.org/