2019 INSIC

Technology Roadmap

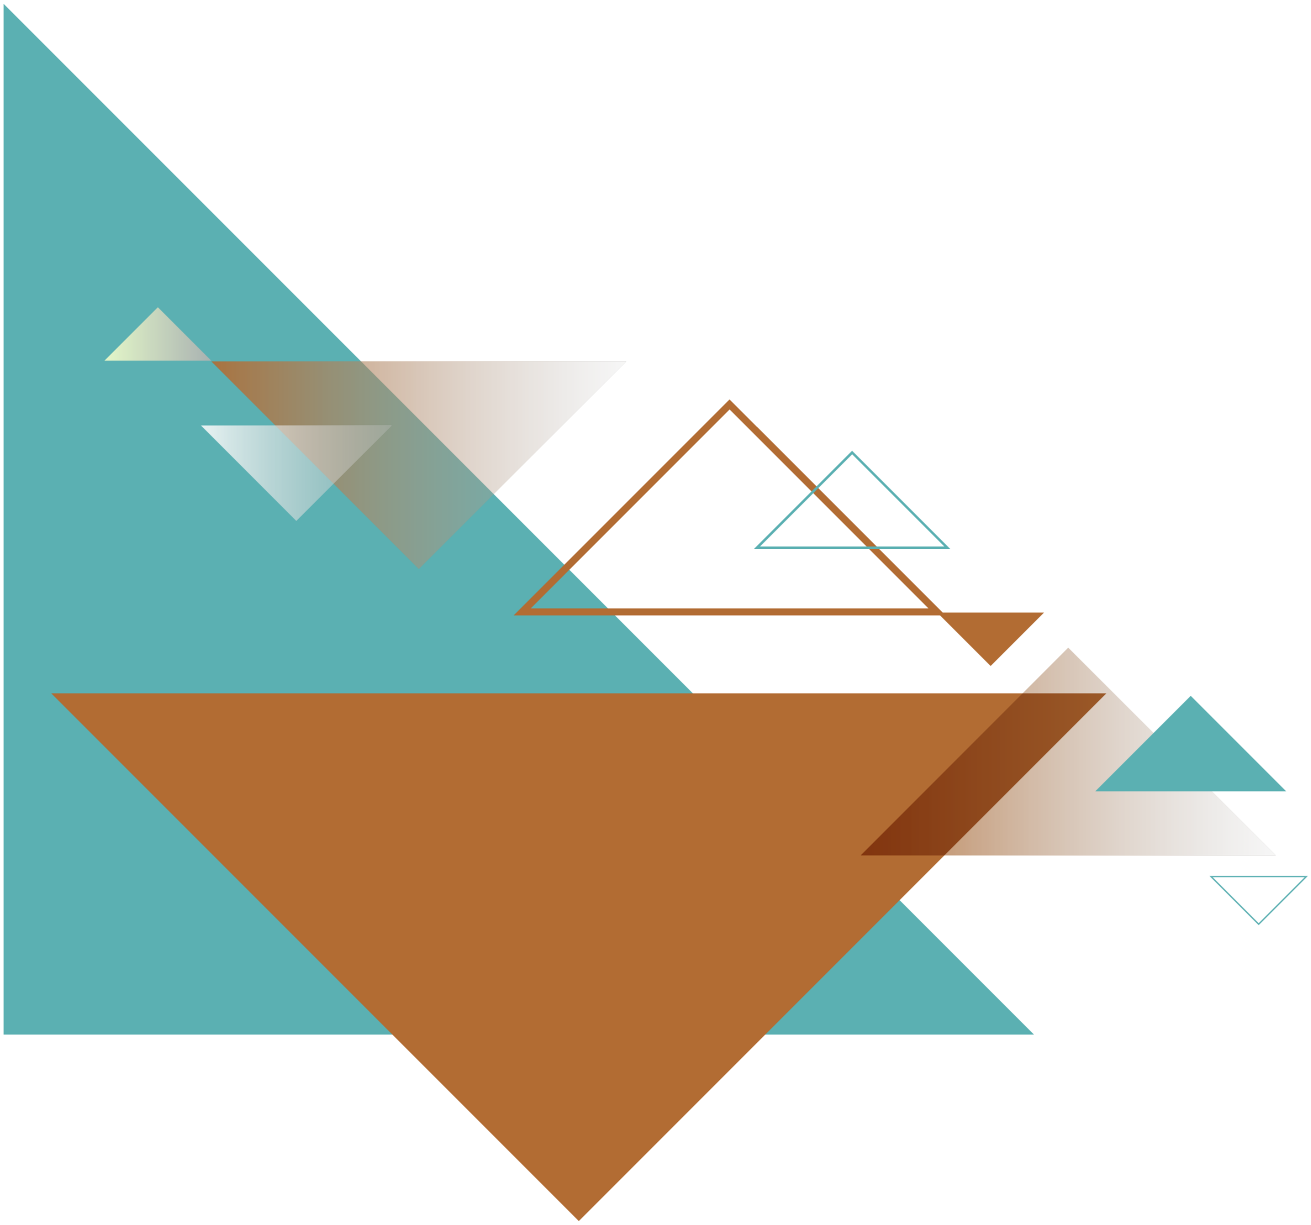

2.0 | TECHNOLOGY

2.0 | PARTICIPANTS

Leadership Team

Robert Raymond, Oracle (Team Lead)
Mark Lantz, IBM
Jeff McAllister, Hewlett Packard Enterprise
Geoff Spratt, Hewlett Packard Enterprise
Turguy Goker, Quantum
James Peng, Quantum
Suayb Arsian, Quantum

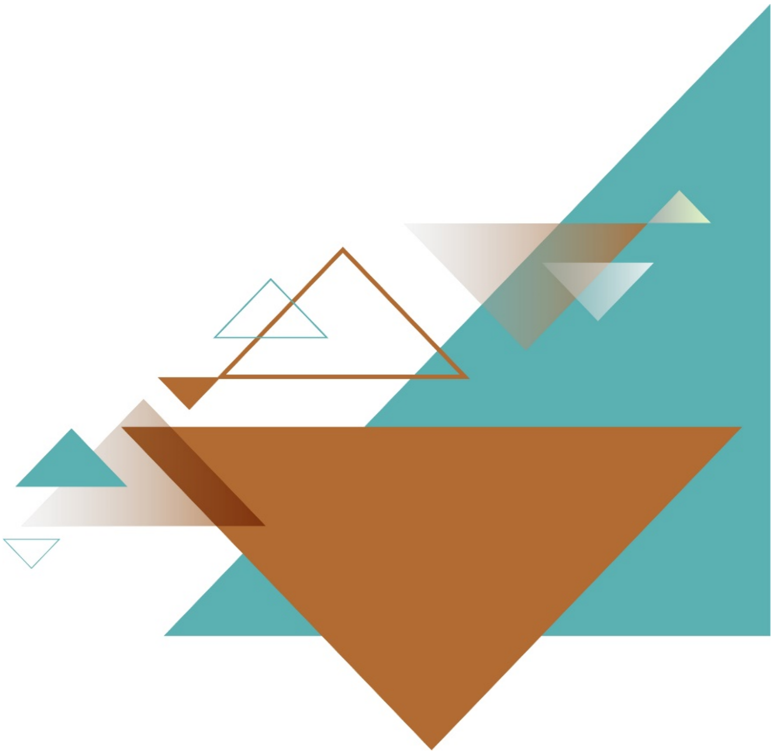

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

2.1 | INTRODUCTION

This section discusses the Tape Technology Roadmap as well as the technologies and challenges needed to execute the ten-year Technology Roadmap goals. The approach is to determine what Roadmap goals are needed for tape to remain a viable storage alternative in the markets described in the preceding Applications and Systems section. The Roadmap was developed to preserve the cost advantage that tape has over hard disk drive (HDD) technology in these applications. The metric used is the cost per gigabyte (GB) of data stored. The method to reduce the cost per GB of tape is to push the technology to enable storing more data in the same half-inch form factor cartridge without significant increase in cost of the cartridge. The Roadmap discussed here is a tape roadmap, which has its base primarily in the enterprise and midrange tape markets. 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 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 for HDD has been driven in large part by the areal density growth of hard disk technology. Between 2003 and 2009, the areal density growth for HDD was about 39% per year, as shown in Figure 1 below. More recently, between 2009 and 2018, the average rate of areal density scaling has decreased to about 7.6% per year. This is a significant reduction from the 16% per year in the 2015 INSIC Tape roadmap. 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; however, the current rate of HDD capacity scaling is still much lower than historical rates. Looking to the future, there is considerable uncertainty regarding future scaling rates of HDD due to the challenges associated with overcoming the super-paramagnetic effect and the uncertainty over the timing and the eventual success of the introduction of new technologies needed to continue HDD scaling. In contrast, state of the art tape drives operate at areal densities that are about two orders of magnitude smaller than the latest HDDs. It should therefore be possible to continue scaling tape technology at historical rates for at least the next decade, before tape begins to face similar challenges related to the super-paramagnetic effect. Even though the areal density of tape is much lower than HDD, tape gets its capacity advantage over hard disk by having a much larger recording surface in a cartridge, with about 1000X the area of a 3 ½ inch disk platter, and subsequently does not need as high a recorded areal density to achieve its cost per GB advantage.

The historical rate for the tape areal density growth of 34% per year is set as the target for the 2019 roadmap and corresponds to a 40% per year capacity growth which is slightly less than the 41% per year capacity growth rate in the 2015 INSIC International Magnetic Tape Storage Roadmap [1]. This slight reduction in capacity growth is due to less aggressive goals in tape thickness reduction and ECC overhead reduction to better reflect what is expected to be possible in the future. The growth rate for data rate is set to 17% per year which is significantly lower than the 22.5% growth rate for the 2015 roadmap. This reduction in data rate is due to tape speed limitations with a reasonable number of channels in the later years of the roadmap.

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

The Tape Technology 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 Tape Technology Roadmap shown here requires an areal density in the year 2029 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 201 Gb/in^2 [2]. Therefore, from a fundamental areal density perspective, we feel the 2029 areal density target of about 300 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.

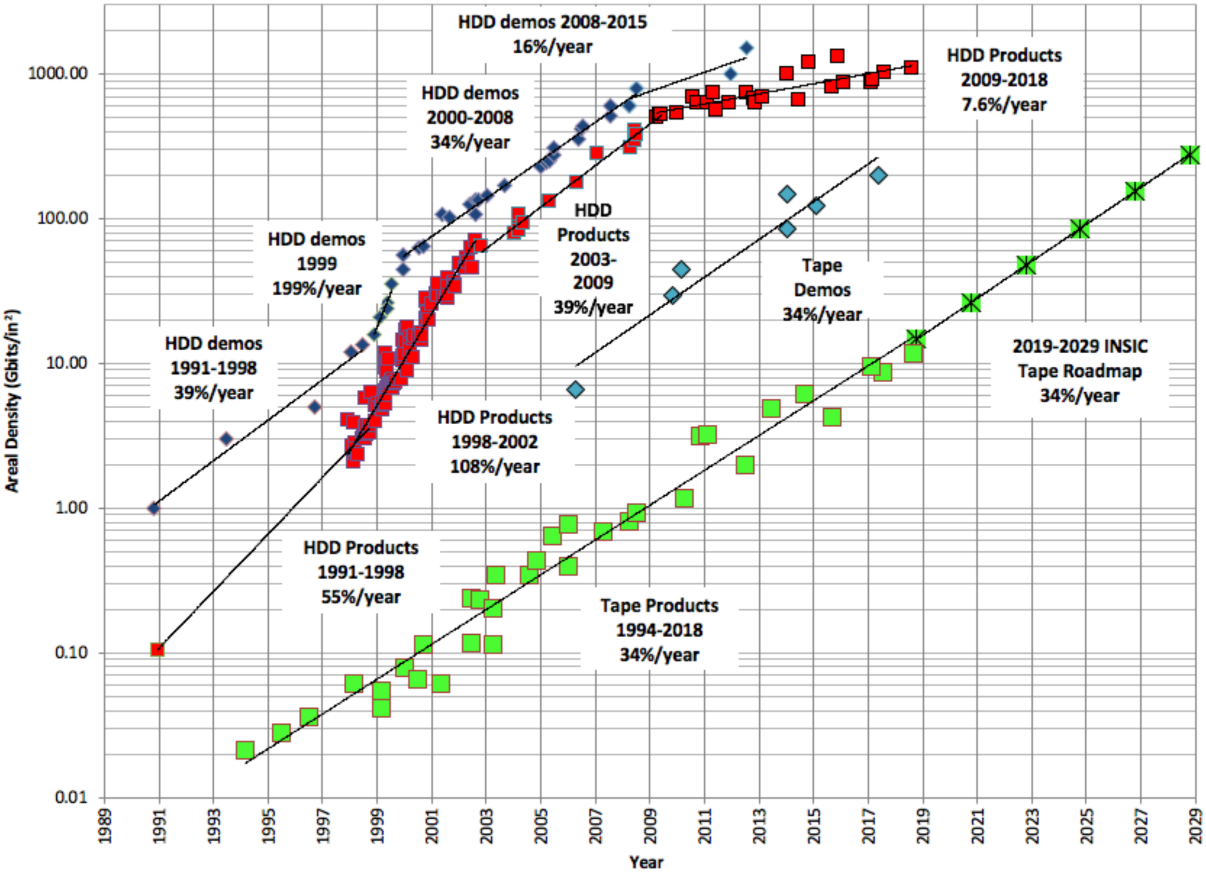

Figure 1: Areal Density Trends. Hard Disk Drive, Tape Product and Tape Technology Roadmap

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

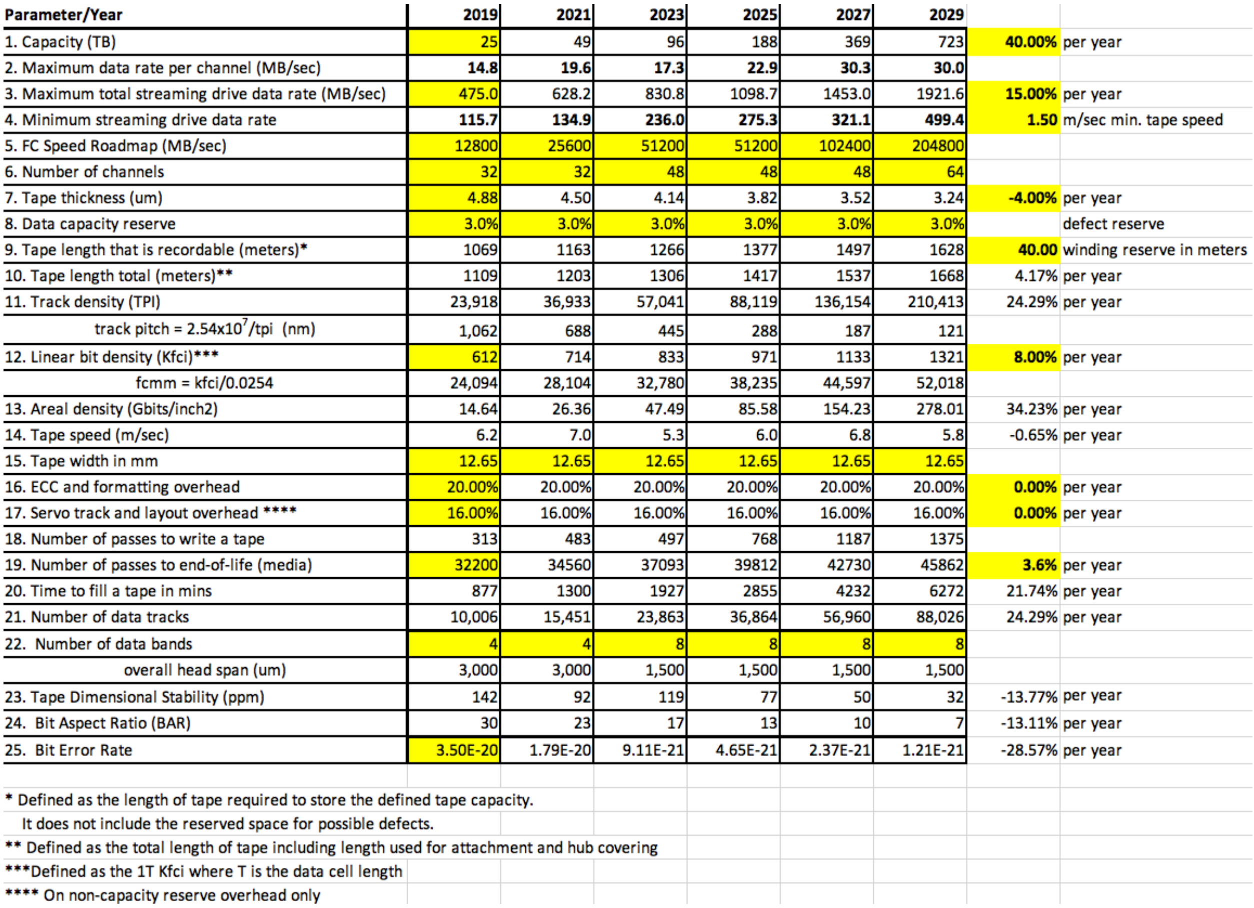

Parameter/Year	**2019**	**2021**	**2023**	**2025**	**2027**	**2029**		
1. Capacity (TB)	25	49	96	188	369	723	**40.00%**	per year
2. Maximum data rate per channel (MB/sec)	**14.8**	**19.6**	**17.3**	**22.9**	**30.3**	**30.0**		
3. Maximum total streaming drive data rate (MB/sec)	475.0	628.2	830.8	1098.7	1453.0	1921.6	**15.00%**	per year
4. Minimum streaming drive data rate	**115.7**	**134.9**	**236.0**	**275.3**	**321.1**	**499.4**	**1.50**	m/sec min. tape speed
5. FC Speed Roadmap (MB/sec)	12800	25600	51200	51200	102400	204800		
6. Number of channels	32	32	48	48	48	64		
7. Tape thickness (um)	4.88	4.50	4.14	3.82	3.52	3.24	**-4.00%**	per year
8. Data capacity reserve	3.0%	3.0%	3.0%	3.0%	3.0%	3.0%		defect reserve
9. Tape length that is recordable (meters)*	1069	1163	1266	1377	1497	1628	**40.00**	winding reserve in meters
10. Tape length total (meters)**	1109	1203	1306	1417	1537	1668	4.17%	per year
11. Track density (TPI)	23,918	36,933	57,041	88,119	136,154	210,413	24.29%	per year
track pitch = 2.54×10^7/tpi (nm)	1,062	688	445	288	187	121		
12. Linear bit density (Kfci)***	612	714	833	971	1133	1321	**8.00%**	per year
fcmm = kfci/0.0254	24,094	28,104	32,780	38,235	44,597	52,018		
13. Areal density (Gbits/inch2)	14.64	26.36	47.49	85.58	154.23	278.01	34.23%	per year
14. Tape speed (m/sec)	6.2	7.0	5.3	6.0	6.8	5.8	-0.65%	per year
15. Tape width in mm	12.65	12.65	12.65	12.65	12.65	12.65		
16. ECC and formatting overhead	20.00%	20.00%	20.00%	20.00%	20.00%	20.00%	**0.00%**	per year
17. Servo track and layout overhead ****	16.00%	16.00%	16.00%	16.00%	16.00%	16.00%	**0.00%**	per year
18. Number of passes to write a tape	313	483	497	768	1187	1375		
19. Number of passes to end-of-life (media)	32200	34560	37093	39812	42730	45862	**3.6%**	per year
20. Time to fill a tape in mins	877	1300	1927	2855	4232	6272	21.74%	per year
21. Number of data tracks	10,006	15,451	23,863	36,864	56,960	88,026	24.29%	per year
22. Number of data bands	4	4	8	8	8	8		
overall head span (um)	3,000	3,000	1,500	1,500	1,500	1,500		
23. Tape Dimensional Stability (ppm)	142	92	119	77	50	32	-13.77%	per year
24. Bit Aspect Ratio (BAR)	30	23	17	13	10	7	-13.11%	per year
25. Bit Error Rate	3.50E-20	1.79E-20	9.11E-21	4.65E-21	2.37E-21	1.21E-21	-28.57%	per year

* 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

Table 1: 2019 Tape Technology Roadmap Detail.

Figure 2 below shows two data plots to further illustrate how tape compares to disk (HDD). These plots show various areal density based operating points for both Tape and HDD including current products, future roadmap predictions, and technology demos [2], [9]. The operating points for these plots show bit length on the y axis for the first plot and track pitch for the second plot, areal density on the x axis and capacity illustrated as the size of the data point bubble. The figures show the clear advantage of tape over HDD since it can increase capacities using moderate areal densities compared to disk. A recently published paper [3] predicts 30TB HAMR HDD using 1970 Gbits/inch2 areal density operating point which is nearly 40X higher than the 96TB Tape based on INSIC roadmap. Basically, an increase of 40X the areal densities over future tape product still can only achieve 1/3 the capacity of the tape making tape still the best option for bulk archival cold data storage for many years to come.

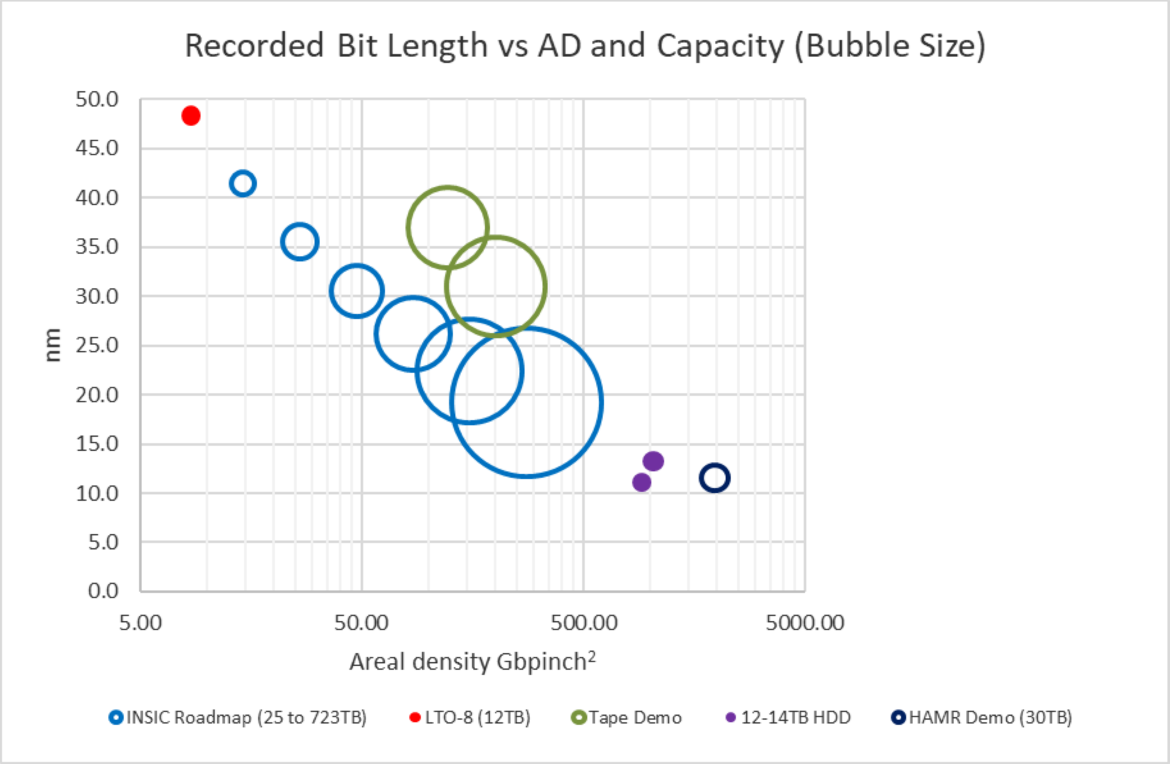

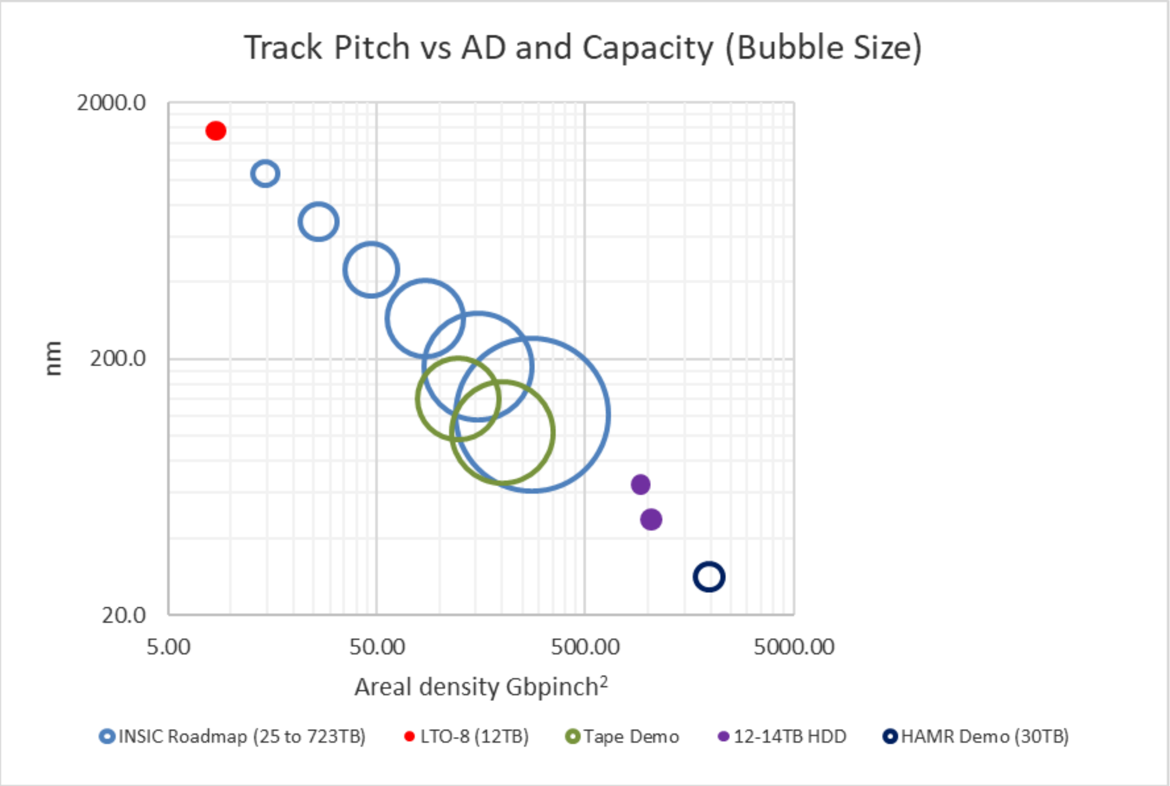

Figure 2: Recording Technology comparison plot for Tape and HDD

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

2.2 | ROADMAP DISCUSSION

In reference to Table 1 above, the entries highlighted in yellow 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 2019 of 25 TB being an input as well as the 40% per year capacity growth. The 40% exceeds the anticipated HDD technology growth rate and gives an approximate doubling of capacity every two 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 total 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 475 MB/s in 2019 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 1.5 meters per second. The drive cannot reliably go below 1.5 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 more than 80X the data rate of the tape drives out to the year 2029. 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 2019 and then increasing at year 2023 to 48 and at year 2029 to 64. This was felt to be a reasonable increase and characteristic of how products will likely scale.

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

7) **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.

8) **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.

9) **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.

10) **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.

11) **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 50 to 100. 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.

12) **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 less, so there is not as much opportunity for improvement here as for the track density. The 8% 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.

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

13) **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 15 and 16, below. Areal density is also equal to the product of track density and linear bit density.

14) **Tape speed** 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.

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 2019, which is in line with current products shipping and is projected to remain constant at 20%, in line with recent trends and to enable the projected improvements in uncorrectable bit error rate (UBER).

17) **Servo track and layout overhead** is assumed to be at 16% for the entire 10-year Roadmap, which is in line with current products shipping.

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 minutes, to completely fill a tape when writing at the maximum data rate. This number is increasing since it is easier to increase tape capacity 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) **Number of data bands** is the number of sections on tape separated by servo patterns. The distance between servo pattern bands is known as the overall head span. Increasing the number of bands reduces the tape dimensional stability requirement.

23) **Tape Dimensional Stability** (TDS). In this roadmap the TDS in row 23 is the residual TDS after all compensations in the tape drive have been implemented and allows for

20% worse case track overlap on the outside elements of the head. The total pre-compensated TDS is the change in tape width in ppm due to changes in temperature, humidity and in-cartridge creep. We assume that tension is controlled to compensate for dimensional changes resulting from these factors and hence tension is not included as an additional TDS term as in previous roadmaps (see discussion in section 2.3.1 and 2.3.3) We define the total pre-compensated tape dimensional stability as the sum of all these factors relative to the head (i.e., a worst case).

24) **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 3.5 today. At the end of the Roadmap in year 2029 tape still has a BAR of about 2X of where disk is today.

25) **Uncorrectable Bit Error Rate (UBER)** is the number of bits in error divided by the total number of bits transferred. The starting point in 2019 is taken as the 3.5-20 UBER, a value intermediate between current LTO and Enterprise products and the scaling rate is set such that the product of the end-of-life (EOL) UBER and the tape cartridge capacity is a constant.

2.3 | KEY TECHNOLOGY CHALLENGES

The combination of BaFe (barium ferrite) particulate media with TMR (Tunneling Magnetoresistive) recording heads has emerged as the technology recipe that will most likely carry tape through the next several generations. Three high-capacity products [8] that implemented this combination of technologies are currently shipping, with the first launched in May of 2017. In addition, technology demonstrations have been done showing tape capacities out to at least 220 terabytes (TB) in a cartridge with these technologies [9]. 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 2029 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 2029 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 barium ferrite media areal density [9] paper achieved an areal density of 123 Gbit/in^2 using perpendicularly oriented media. A linear recording density of 680 kfci with a written track width of 140nm and a 90nm GMR reader lead to the 123 Gb/in^2 areal density.

2.3.1.2 | Sputtered Media

A more recent technology demonstration for sputtered tape media [2] has extended the areal density from 148Gb/in^2 published in 2014 [4] to 201Gb/in^2. The earlier demonstration used a drag tester in combination with conventional HDD write/read head and estimated the track density as double the reader width, without taking account of track-following performance. In contrast, the track density for the recent 201Gb/in^2 areal density was estimated using the track-following performance characterized by the standard deviation of the position error signal (σ-PES). Using the 48 nm reader width and the worst case measured σ-PES of 6.5 nm leads to an estimated track width of 103 nm and the resultant track density of 246.2 ktpi. Recording measurements made with a prototype high-moment tape write head and a 48nm TMR disk drive read head demonstrated a post-detection byte-error rate < 0.023 at a linear density of 818kbpi.

The spacing between the magnetic tape and the head is critical to achieving high-density recording capabilities for tape storage media. To realize both small spacing and low friction, a new lubricant was applied. This low-friction lubricant reduced the friction between the tape surface and head and featured a highly durable bond between the lubricant and magnetic layer of the tape [10]. Furthermore, this new technology made it possible to create a nano-grained magnetic layer with microscopic magnetic particles in extended tape length. With these advancements in processing technology, the foundation has been laid for production of sputtered tape storage cartridges that can hold more than 1,000 meters of tape.

Although these demonstrations have shown media to be capable of supporting areal densities approaching the roadmap target, they both make one crucial assumption. These areal densities do not take into account the dimensional stability of the media and hence assume higher track densities than the Roadmap is targeting. 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 2029 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.

2.3.1.3 | Substrate

Tape substrate suppliers have provided the following extensions to the 2015 Roadmap. Improvements in TDS will be obtained by tensilization orientation and nano-sized polymer blends. The handling of the thinner tape 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.

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

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

	2019	2021	2023	2025	2027	2029
Thickness, µm	3.6	3.6	3.2	3.2	2.8	2.8
Thermal Hygroscopic In-Cartridge Creep	0 470 100	0 470 100	0 470 100	0 470 100	0 350 100	0 350 100
Total (TDS)	570	570	570	570	450	450
Tension	130	130	160	160	200	200

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

	2019	2021	2023	2025	2027	2029
Thickness, µm	4.4	4.2	4.0	4.0	3.8	3.8
Thermal Hygroscopic In-Cartridge Creep	25 375 100	0 350 100	0 300 100	0 300 100	0 275 100	0 250 100
Total (TDS)	500	450	400	400	375	350
Tension	150	150	150	100	75	75

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

	2019	2021	2023	2025	2027	2029
Thickness, µm	3.6	3.3	3.0	2.8	2.6	2.4
Thermal Hygroscopic In-Cartridge Creep	100 100 50	50 50 50	50 50 50	0 0 50	0 0 50	0 0 50
Total (TDS)	250	150	150	50	50	50
Tension	50	50	50	50	50	50

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

From the above data, only the Aramid substrate approaches the roadmap target for TDS. However, Aramid is significantly more expensive than the other substrates. So 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 changes. One such method could dynamically adjust tension to maintain a specific track pitch. For this reason the tension TDS component is not counted towards the total TDS value, but is included to show a possible compensation range of the TDS. 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.

2.3.1.4 | What key improvements are needed to exceed 1,000 kbpi?

The linear densities achieved in the areal density demos (680 kfci and 818 kfci) are well below the current roadmap's target of 1,342 kfci.

For particulate media to achieve the necessary SNR at the roadmap linear densities a reduction in particle volume is needed. This may require other particles to be considered. One promising avenue for reduction in particle volume, while maintaining the required thermal stability, is fine strontium ferrite (SrFe) particles. These particles can be scaled to almost half the volume of the BaFe particles used in current products [11]. Another particle being considered is the nanometer sized $\varepsilon\text{-Fe}_2\text{O}_3$. A prior study showed spherical nanoparticles with a diameter of 8.2 nm having an H_c value of 5.2 kOe at room temperature. [12] Recently, $\varepsilon\text{-Fe}_2\text{O}_3$ magnetic particles combining features of fine particles, tighter distribution of particle size and lower H_c have been obtained. High moment writers will be able to take advantage of the benefit of $\varepsilon\text{-Fe}_2\text{O}_3$ even further as the feature of the particle is high H_c.

Sputtered media uses similar recording materials to those employed in the hard disk drive industry. Current 3.5" HDD products based on PMR operate at areal densities of up to 1.1Tb/in^2. In the case of sputtered tape, when technologies such as a soft magnetic underlayer (SUL), a capped layer and an exchange coupled continuous structure are applied, significant progress in areal density is expected. The use of the SUL will enable perpendicular recording using single pole write heads which will improve SNR compared to writing with a ring head.

Another factor critical for exceeding 1,000 kfci will be the continued reduction in head-tape spacing. Since tape recording systems result in contact between the head and media, tribology issues need to be studied for the more aggressive reduction in head to tape spacing and smoother surfaces that will be required. There are several factors involved in the head tape spacing, but the key factors are likely to remain the surface roughness of the media (and substrate) and head wear. The media manufacturers will need to continue the refinement of the media surface, via lubricant and surface design, to allow closer head-to-tape spacing while reducing abrasivity and maintaining low coefficients of friction to enable good media durability.

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

2.3.2 | Future Technologies for Recording Heads

Tape recording head technology has benefitted significantly over the last few decades by adopting and adapting materials and technologies developed by the HDD industry. Notable examples are the application of GMR reader technology in tape drives and the more recent introduction of TMR reader technology, initially in the IBM TS1155 drive in early 2017, followed by the IBM LTO8 drive in late 2017. Compared to GMR readers, TMR readers provide approximately 3x more read-back amplitude and exhibit much smaller signal loss over their lifetime [13]. In addition, well designed TMR readers also have better magnetic stability and operate at lower temperatures due to reduced bias current requirements [13].

TMR technology is also well suited for use with tape for several other reasons. Compared to recent HDD products, state of the art TMR tape readers have much larger widths and stripe heights, resulting in a tape read sensor area that is > 400x larger than for HDD sensors [13]. This large sensor area enables the use of much thicker MgO tunnel barriers compared to HDD, which provides improved reliability, consistency and higher yield [13]. Yield is particularly important for tape heads which host many transducers that are operated in parallel. State of the art heads currently operate 32 data and 2 servo readers in parallel and by 2029 this is expected to increase to 64 parallel data readers and 2 or more servo readers. In current TMR tape heads, the dimensions of the data and servo readers are similar, resulting in a comparable sensor resistance for both transducers. This resistance is ideally in the range of 50 – 100 Ohms to achieve a good bandwidth in the amplification stage of the signal processing chain. In the future, as track pitch is scaled from around 1um in 2019 to about 120nm by 2029 the reader width and stripe height will have to be scaled to much smaller dimensions, leading to a smaller sensor area and hence thinner tunnel barrier to achieve the target sensor resistance. Over the same period, the dimensions of the servo reader are expected to change much less, leading to a divergence in the optimal tunnel barrier thickness for data and servo readers.

The larger dimensions of state-of-the-art tape heads relative to HDD heads provides additional advantages in terms of reduced complexity and manufacturing challenges. Although track sizes are continually shrinking, the written bits are still large enough for conventional pole definition through a combination of photolithography, sputter deposition and/or plating, and ion milling. More importantly, the magnetic coercivity of the media is not so high as to require “booster” measures such as heat assisted magnetic recording (HAMR). Conventional inductive coil writers with notched poles to suppress side writing can continue to be utilized. Recently, write transducers with a sputter-deposited high-moment CoFeNi alloy seed layer on the trailing edge side of the write gap where introduced in the TS1155 drive [13]. The seed layer has a saturation magnetization of B_s = 2.2T compared to 1.6T for 45/55 NiFe used for the bulk of the poles and enables stronger write fields which in turn enables recording on media with an increased coercivity of about 35% relative to current state of the art commercial tape media [14]. In addition, the high-moment liner creates stronger field gradients which help sharpen write transitions and provide an improvement in SNR. In the future, we expect the write gap to be scaled to smaller dimensions as linear density is increased and magnetic layer thicknesses are reduced. In addition, we expect to see further optimizations of the geometry and materials of write transducers to enable the use of media based on yet higher coercivity recording layers.

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

Looking further into the future we will likely require the head and media to be co-designed, initially with conventional ring writers optimized for a specific media design point. Later, if media with a soft under-layer is introduced in combination with monopole writers, all three components of media, soft under-layer, and writer pole will need to be adjusted in their magnetic properties in order achieve optimal recording performance.

Monopole writers combined with a soft under-layer were introduced by the HDD industry when perpendicular recording was introduced in the mid 2000's. In addition to the improved SNR performance provided by the perpendicular orientation of the media, the use of a monopole writer and soft under-layer provides several other enhancements. First, it effectively places the magnetic recording layer in the write gap, enabling much larger write fields and stronger field gradients that in turn enable the use of increased coercivity media and provide sharper transitions, respectively. Second, the soft under-layer enhances the low frequency components of the read back spectrum and can provide an increase in SNR and bit error rate if the analog front end and read channel are adapted to take advantage of it. Past commercial linear tape drives have used either a longitudinal orientation of the particles with MP media or more recently with the introduction of BaFe, a quasi random orientation of the particles. The 2014/15 tape areal density demonstrations of 85.9Gb/ in^2 and 123 Gb/ in^2 on prototype particulate BaFe media used a conventional ring head in combination with a more perpendicular orientation of the particles to improve the recording performance [14,9]. To date, it has not been possible to incorporate a soft under-layer with particulate media. However, the 2017 demonstration of tape recording at 201Gb/in^2 used a prototype highly-oriented perpendicular sputtered media that included a soft under-layer in the media stack [2]. The demonstration was performed using a ring-type writer, and as such the potential benefits of the soft under-layer were not fully realized. However, if in the future this media technology can be combined with a suitably designed monopole writer, even higher areal densities should be possible. In fact, by continuing to adapt the head and media technologies developed by the HDD industry, it should be possible to continue scaling to the areal density regime at which HDDs currently operate.

Another key challenge in tape head design is the projected continued increase in the number of parallel transducers required to enable continued data rate scaling. Current commercial tape drives use 32 parallel recording channels. As mentioned above, by 2029 this is projected to increase to 64 parallel channels. The larger number of channels per head module has implications for cable design, bonding pad fan-out, writer crosstalk and transducer yield. In addition, the use of more channels requires additional electronics on the drive card. One way to alleviate all but the last of these challenges is the use of multiple heads and multiple actuators, albeit at increased manufacturing cost. The potential problem of writer crosstalk that is expected to occur at small transducer pitch is further compounded by increasing the number of data bands to alleviate TDS issues. For example, doubling the number of data bands halves the span of the active part of the head resulting in half the transducer pitch. A 64 channel, 8 data band design point implies a transducer pitch of about 21.5 microns compared to about 86 microns in the 32 channel 4 DB design used in LTO8. In addition to potential issues of crosstalk, it will be very challenging to design a multi-turn write head compatible with this pitch that does not suffer from electromigration and heating issues. Because of these challenges, the number of data bands in the 2019 roadmap was restricted to 8 rather than the 16 data bands projected

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

by the 2015 roadmap, and we further assume that TDS challenges will be addressed through active compensation schemes, as discussed in sections 2.3.1 and 2.3.3. Finally, we note that an alternative or complementary approach to increasing data rate is the use of RAIT architectures that can enable much higher data rates at the system level.

Another critical area concerns the head-tape interface. With increasing linear density, head-tape separation becomes exponentially more important. To continue reducing tape-head separation, it is necessary to increase the smoothness of the media, which in turn creates challenges for tape-head friction. One way to address this is through the use of novel tape head contour designs that reduce or eliminate skiving edges and minimize the tape-head contact area [15,16]. Research will also be needed to control the electrochemistry and wear mechanisms at the interface. The buildup of deposits from the media onto the heads must be minimized. Additionally, work to create more wear resistant heads coupled with the development of less abrasive media must continue in order to enable the continued reduction of head-tape spacing. This is an issue that needs to be jointly addressed by head and media manufacturers.

2.3.3 | Track Mis-registration /Tape Transport Technology

The mechanical transport area encompasses all technology associated with moving and tensioning the tape in a linear tape drive, including the technology for keeping the read/write head registered to the desired tracks on the tape. Modern data storage tape decks must guide the tape accurately while holding tension steady as the drive accelerates and decelerates the tape. Concurrently, they must precisely position recording heads containing multiple read/write elements over the corresponding data tracks on the tape. Improvements in the tape transport and tracking capability are critical for achieving the higher track densities necessary to support the future capacity requirements.

The tape is partitioned into alternating servo and data bands with a recording head spanning a single data band. Servo elements sense the position of the head within the data band and enable precise control of the head tracking. Within a data band, there are many more tracks than head elements, so each head element will write to or read from a subset of the total tracks in a band called a sub-band. Data is read and written in a bi-directional fashion with each forward and reverse pass of the media across the head referred to as a wrap.

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 210,413 tracks per inch which is equivalent to 121 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. Many factors come into play in ensuring good tracking, including variation in the element pitch, tape expansion or contraction, lateral tape motion, and reader width. As track width shrinks and tracks are crowded closer together, each of these factors must be carefully controlled or compensated for. The consequences of not budgeting for all of the tracking contributors could be writing off track, longer store/restore times due to having to reread or rewrite data, or inability to recover customer data.

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

The transport section will detail many solutions that could be implemented, although not all may be required. However, many of the following improvements will be needed. Reducing tape expansion and contraction or compensating for it has the largest effect on improving tracking margin. A smaller, but still significant factor is lateral tape motion. Last in terms of tracking is reducing the variation in head element pitch, the distance from one head element to another within the recording head.

As tape gets thinner to accommodate longer lengths, the reels will have to accommodate increasing numbers of wraps; 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 either need to be minimized or compensated for.

Some of the operating points in the roadmap will require higher tape speed to achieve the necessary transfer rate. Higher tape speeds increase the frequencies at which servo tracking variations appear, and may also increase the magnitude of lateral tape motion (LTM) due to dynamic interactions with tape path reels and/or guides. The rolling guides will rotate at higher speeds which will also increase the frequency and amplitude of disturbances from the rolling guides themselves. Improvements will likely be required in the precision of the rolling guide elements. Actuator bandwidth must also increase to compensate for these additional disturbances. Written-in servo variation on the tape and position sensing noise must also be reduced.

The extremely high areal density technology demonstrations that are done periodically to demonstrate the capabilities of magnetic media largely ignore many issues related to the track mis-registration budget. Tape drives have to deal with variations caused by head pitch tolerances as the tape is interchanged from one drive to another. 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 flexible media, which will continue to get thinner with each generation, does not track perfectly across the heads at high speed, and the ability of the head servo system to follow that lateral motion becomes more difficult as tape speed increases as predicted by the roadmap.

In order 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 of 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 four basic contributors: tolerance variation between the write heads and read heads, tape dimensional stability, lateral tape motion, and the servo system to minimize the Position Error Signal (PES).

One significant contributor 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 order to keep a read head accurately positioned on a track that is only 121 nm wide, head manufacturing processes will need to improve dramatically to reduce these tolerances.

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

Another significant 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 increases as the tape speed increases. This is one of the factors that limits the speed that the tape can be effectively transported during reading and writing. In order to sufficiently reduce the PES, all disturbances to the tape from the transport system must be minimized. Particular care needs to be taken to remove all 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 in order to dramatically reduce the residual PES. 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. Adaptive servo systems which sense the resonances and possible instabilities in the closed loop system and modify the loop accordingly to enable higher gain and bandwidth may be required.

The accuracy of the servo tracks as written by the media manufacturers is also a critical factor which contributes to the overall PES. Dramatic improvements have been made in this servo writing capability over the last several generations of tape drives and this should continue to be addressed with every new generation. In order to improve the accuracy of the servo position estimation achieved by the servo channel/demodulator, the servo stripe angles have been increased. The spacing between servo frames has been reduced as well to provide higher frequency position update rate and improve phase margin in the closed loop servo system. This has enabled higher bandwidth head positioning servo systems. An additional method that is effective in reducing the contribution of written-in variation in the servo track, as well as random position sensing noise due to variations in the servo frame peak detection, is to simultaneously use all the servo heads available. The typical implementation uses only the two servo elements on the module that is reading or writing. Utilizing the additional servo elements effectively averages out noise and written-in PES.

The largest contributor to track mis-registration budget is the effect of tape dimensional stability (TDS). There are several possible countermeasures that should be investigated and developed to reduce this variation. The most effective solution would be to move to an aramid substrate which could potentially produce a 4x reduction in this contributor 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. This should continue to be considered as a possible solution.

Another very effective countermeasure would be to reduce the overall span of the data heads by increasing the number of data bands. Such an increase is shown on the roadmap in 2023. This creates very difficult challenges for the head manufacturing process; especially if that is coupled with a doubling of the number of channels. The other challenge that this approach presents is the difficulty in reading previous generation written tapes because the heads are no

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

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 in order to continue to progress tape technology areal density is to forfeit backward compatibility to enable a leap forward. This is very undesirable from a customer point of view and could be a serious roadblock to adopting such a new tape format.

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 translates into 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, especially cloud providers, 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 in order to compensate for environmental expansion and contraction. This is a technology that needs to be further developed to increase the available range of compensation. Media manufacturers have previously been required 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. Efforts to reduce the humidity related TDS contribution of the current PET and PEN basefilms should continue as in previous generations. One area to consider is 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 track pitch 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.

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 GMR 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 will be feasible.

Another potential concept is the use of PES target to compensate for TDS. In this method, the goal is to vary the PES target to reduce track squeeze due to TDS, thus guaranteeing minimum track width during write mode. The basic methodology is to use the TDS range from the previous track to calculate a new PES target. The impact of this concept is the capacity loss if the tape experiences multiple instances of extreme environment change.

In order 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 a multi-discipline issue that requires continued investment by drive, media, and head manufacturers in order 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 and enterprise tape cartridges has increased by a factor of 120 to 200 and data rates have increased by more than an order of magnitude. During the first 15 years of LTO drives, up until the seventh generation was 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 uncorrectable bit error rate (UBER) remained fixed. Calculations of the UBER in a modern tape drive, a.k.a. the hard error rate, 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, an assumption that has recently been experimentally verified [9,14]. More complex reliability models which are based on the theory of renewal processes can account for correlated errors and defective header and synchronization fields [17]. In order to address the increased probability of encountering an error in a cartridge that results from capacity scaling, the error correction capability used in LTO-7 was improved to provide enhanced error correction performance. Theoretical analysis of the error correction scheme implemented in LTO-7 leads to an UBER over the lifetime of the drive of not more than one error per 1e19 bits read. This is comparable to the performance achieved by enterprise class tape drives such as the TS1150/55/60 or T10000C/D that utilize a 32 track format and provide an UBER of 1e-20. Another improvement in LTO-7 tape drives significantly reduced the rewritten area on magnetic tape by collecting all the encoded subunits of a 6MB Data Set written on 32 tracks, 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.

The product error correction codes used in tape storage have two component codes which are usually referred to as C1 (row code) and C2 (column code). When compared to the LTO-1 format, the main improvements in the LTO-7 format enabling the improved end-of-life (EOL)

UBER are more powerful C1 and C2 Reed-Solomon (RS) 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 input of the C1 and C2 decoder. In the future, error correction schemes used to encode and decode data and the data channel architecture used in tape storage are expected to improve significantly. For example, the block size of C2 codewords in LTO-9 tape drives is expected to increase significantly, thus enabling the EOL UBER to be improved by several orders in magnitude. These performance gains could be used to provide a constant probability of hard error while reducing the requirements for signal-to-noise ratio (SNR) at the input of the data channel and hence to enable higher areal density, or, for a constant level of SNR, they could be used to provide a reduced error probability. In the future, we project that the improvements in ECC and the data channel will enable a combination of both of these goals to be achieved. For this reason, a projected scaling rate for the EOL UBER was added to the Roadmap in 2015. In the 2019 Roadmap, we take as a starting point an UBER of 3.5 e-20 approximately midway between LTO8 and current enterprise drives and then the scaling rate is set such that the product of the EOL UBER and the tape cartridge capacity is a constant, corresponding to the product of the projected tape cartridge capacity of 25 TB and the EOL UBER of 3.5e-20 in 2019. Therefore, on average, only one out of 89,285 tape cartridges will contain an uncorrectable error event due to ECC failure. As the tape cartridge capacity is expected to grow at a 40% compound annual growth rate (CAGR), corresponding to approximately doubling cartridge capacity every two years, the EOL UBER is specified to shrink at the commensurate rate of 28.57% (see Figure 3). It is worth mentioning that the magnitude of the improvements in encoding and decoding data within tape drives are expected to be more than sufficient to allow for a shrinkage of the EOL UBER at the foreseen rate of 28.57% assuming a constant channel SNR. These improvements will therefore also enable the requirements for SNR at the input of the channel to be reduced and hence will contribute to future areal density and capacity gains. Examples of technologies currently under investigations that could enable such gains are described below.

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

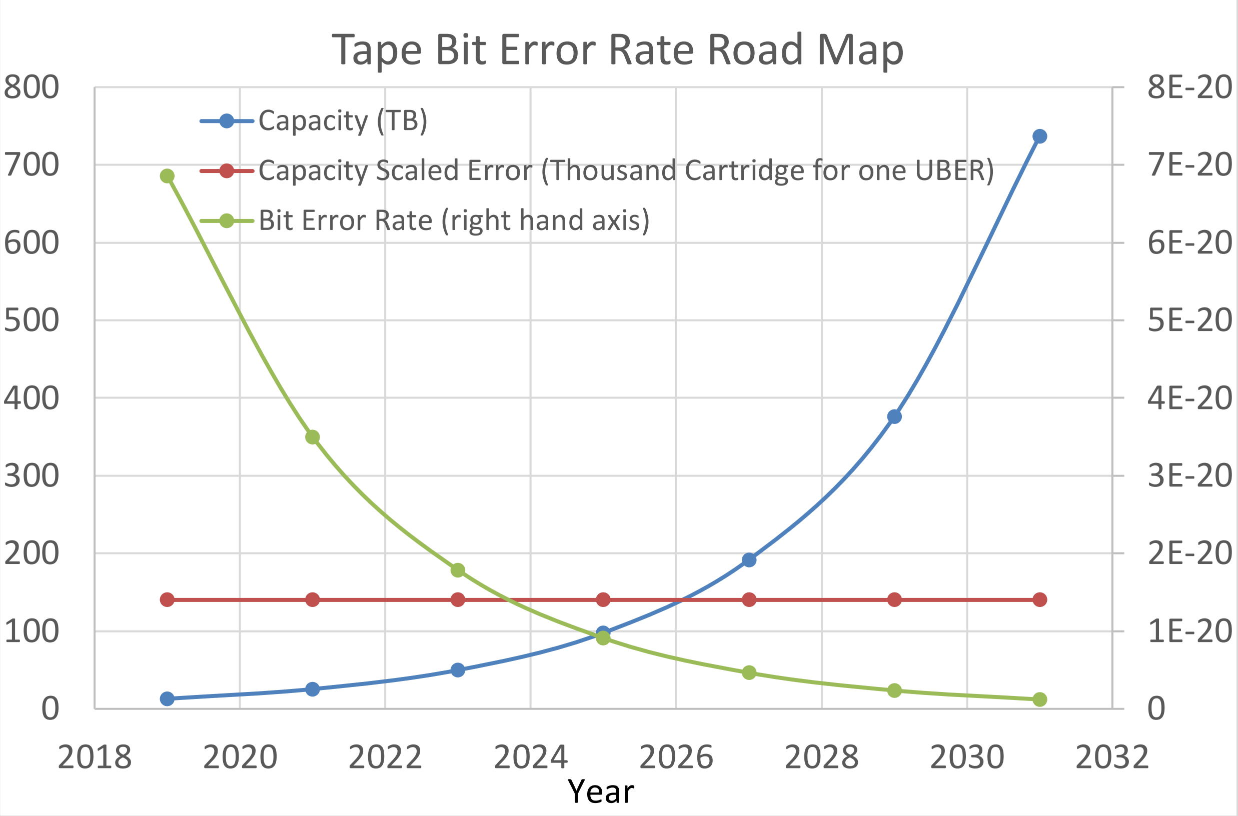

Figure 3: Scaling of end-of-life uncorrectable bit error rate

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 [18] and then encoded by a modulation code that satisfies various run-length constraints on the maximum length of zeros [19], transition-runs [20], synchronization patterns used for timing acquisition (alternating-bit pattern in NRZI notation) [21] and twins pattern to limit the path memory of the Viterbi detector [22]. Increasing the symbol size and the block size of Reed-Solomon (RS) C1 codes [34] and the block size of RS C2 codes [23] can already significantly improve the EOL UBER by several orders in magnitude and allow operation at lower SNR. 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 [9]. For the ECC code rate of 0.83 in LTO-8, the computation of the channel capacity of the magnetic tape channel modeled as a discrete symmetric memoryless channel [35] 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 the length of the ECC codewords and the ECC decoding effort/complexity is not limited. For an RS(240,230) C1 code and an RS(192,168) 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 [14]. 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 [35]. Reverse concatenation of error-correction coding and modulation coding [24-26] allows the use of very high rate modulation codes [27,28] which have large block sizes, an overhead of less than 1% and satisfy required constraints of a tape channel. Error propagation at the output of the modulation decoder with large block size is not an issue because modulation decoding is performed after decoding the error correction code. Moreover, because soft information can directly be passed from the detector to the error-correction decoder, reverse concatenation schemes allow the use of soft-decoding and iterative detection/decoding techniques to improve the error-rate performance [36]. The use of partial reverse concatenation [29] based on employing a higher rate (≈0.991) modulation code [27,28] 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 [30] 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 [31].

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 [32]. Another promising approach to cycle slip mitigation involves exploiting the parallel timing information available in a multichannel tape drive [33]. 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 [37].

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 10.5dB 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,

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

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. Although LTO-7 and LTO-8 tape drives increased the efficiency of the rewrite scheme, significant additional improvements will be needed in order for tape drives to operate at lower SNR operating points in the future. 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. Furthermore, rewrite tables with reduced restrictions and dead track detection in conjunction with rewriting on live tracks that are not dead can be used to increase the rewrite efficiency. We believe that this is a particularly important area for continued research because lowering the operating point towards an SNR of 10.5dB could enable an increase in areal density of 2x or more without any changes to the media.

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

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 technology roadmap discussion. In particular, we would like to thank the engineers at Fujifilm, Sony, Teijin, Toray and Dowa for the media technology inputs.

Additionally, we would like to thank Roy Cideciyan (IBM) for his contributions to the error correction and channel section and Hugo Rothuizen (IBM) for this contributions to the head section.

2.5 | REFERENCES

1. "2015 International Magnetic Tape Storage Roadmap," Information Storage Industry Consortium, http://www.insic.org
2. 201 Gb/in² Recording Areal Density on Sputtered Magnetic Tape: Simeon Furrer et al. IEEE Transactions on Magnetics, VOL: 54, Issue: 2, Feb. 2018.
3. https://ieeebicop.com/wp-content/uploads/2019/01/BICOP-Mark-Gubbins.pdf
4. J. Tachibana et al., "Exploratory Experiments in Recording on Sputtered Magnetic Tape at an Areal Density of 148 Gb/in^2", IEEE Trans. Magn., Vol. 50, No. 11, 3202806, Nov. 2014.
5. "Fibre Channel Speed Roadmap," Fibre Channel Industry Association, V21. http://www.fibrechannel.org/roadmap/
6. "Serial Attached SCSI (SAS) Master Roadmap," SCSI Trade Association. http://www.scsita.org/library/2015/10/serial-attached-scsi-technology-roadmap.html
7. "The 2016 Ethernet Roadmap," Ethernet Alliance Organization. http://www.ethernetalliance.org/roadmap.
8. IBM TS1155 shipped with TMR heads on 5/9/17, LTO8 shipped with TMR heads on 10/10/17, and the TS1160 shipped on 12/01/18 with TMR heads, based on input from IBM.
9. 123 Gb/in^2 Recording Areal Density on Barium Ferrite Tape: Mark A. Lantz et al. IEEE Transactions on Magnetics, VOL: 51, Issue: 11, Nov. 2015

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

10. T. Iwano and K. Kobayashi, “Improving the corrosion resistance of metal-evaporated tape using dicarboxy acid as a lubricant”, *IEEE Trans. Magn.*, vol. **41**, no. 10, pp. 3010-3012, 2005.

11. M. Oyanagi, “Particulate magnetic tape for data storage and future technologies,” HEPiX Fall/Autumn 2017 Workshop, Book of Abstracts page 20.

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

13. R. G. Biskeborn et al. “TMR tape drive for a 15TB cartridge, AIP Advances **8**, 056511 (2018)

14. S. Furrer et al., “85.9 Gb/in2 Recording Areal Density on Barium Ferrite Tape.” IEEE Trans. Magn., Vol. 51, no. 4, April 2015.

15. G. M. McClelland et al., “Effect of tape longitudinal dynamics on timing recovery and channel performance,” IEEE Trans. Magn., vol. 45, no. 10, pp. 3585–3589, Oct. 2009

16. Johan B. C. Engelen et al. “Tape Head with Sub-Ambient Air Pressure Cavities” IEEE Trans. On Magnetics, Vol. 52, No 11, Nov 2016, 3301810, DOI 10.1109/TMAG.2016.2588

17. 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.

18. P. Elias, “Error-Free Coding,” IRE Trans. Inform. Theory, vol. 4, pp. 29–37, 1954.

19. K. A. S. Immink, “Runlength-limited sequences,” Proc. IEEE, vol. 78, pp. 1745–1759, Nov. 1990.

20. J. Moon and B. Brickner, “Maximum transition run codes for data storage systems,” IEEE Trans. on Magnetics, vol. 32, pp. 3992–3994, Sep. 1996.

21. R. D. Cideciyan, F. Dolivo, R. Hermann, W. Hirt, and W. Schott, “A PRML system for digital magnetic recording,” IEEE J. Select. Areas Commun., vol. 10, no. 1, pp. 38–56, Jan. 1992.

22. R. D. Cideciyan, E. Eleftheriou, B. H. Marcus, and D. S. Modha, ”Maximum transition run codes for generalized partial response channels,” IEEE J. Select. Areas Commun., vol. 19, no. 4, pp. 619–634, Apr. 2001.

23. Expanded Sub Data Set for Magnetic Tape Recording, IPCOM000223249D, Nov. 2012.

24. W. G. Bliss, “Circuitry for performing error correction calculations on baseband encoded data to eliminate error propagation,” IBM Tech. Discl. Bull., vol. 23, pp. 4633–4634, Mar. 1981.

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

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

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

27. T. Mittelholzer, “Enumerative maximum transition run codes,” Proc. IEEE Int'l Symp. on Information Theory, ISIT 2009, Seoul, Korea, pp. 1549-1553, June 2009.

28. 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.

29. R. D. Cideciyan, R. Hutchins, T. Mittelholzer, and S. Ölçer, "Partial Reverse Concatenation for Data Storage," Proceedings of Asia-Pacific Signal and Information Processing Association, 2014 Annual Conference and Summit (APSIPA), Dec. 2014.

30. 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, pp. 59-69.

31. S. S. Arslan, J. Lee, and T. Goker, “Error Event Corrections Using List-NPML Decoding and Error Detection Codes,” IEEE Trans. on Magnetics, vol. 49, no. 7, July 2013, 3775.

32. S. S. Arslan, J. Lee, and T. Goker, “Cycle Slip Detection and Correction Through Classification of Modulation Code Failures,” IEEE Trans. on Magnetics, vol. 49, no. 9, Sep. 2013, 4998.

33. 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, pp. 1519-1520.

34. R. D. Cideciyan, S. Furrer, and M. A. Lantz, “Performance of Interleaved Block Codes with Burst Errors,” accepted for publication in IEEE Trans. on Magnetics, vol. 55, no. 3, Mar. 2019.

35. R. D. Cideciyan, S. Furrer, and M. A. Lantz, “Product Codes for Data Storage on Magnetic Tape,” IEEE Trans. on Magnetics, vol. 53, no. 2, Feb. 2017.

36. R. L. Galbraith, T. Oenning, M. Ross, B. Wilson, I. Djurdjevic, and J. Park, “Architecture and Implementation of a First-Generation Iterative Detection Read Channel,” IEEE Trans. on Magnetics, vol. 46, no. 3, Mar. 2010, pp. 837-843.

37. S. Furrer and M. A. Lantz, “Timing Recovery for Low-SNR Magnetic Tape Recording Channels,” Proceedings of The 29th Magnetic Recording Conference (TMRC 2018), pp. 104-105, Aug. 2018.

© 2019 Information Storage Industry Consortium – All Rights Reserved
Reproduction Without Permission is Prohibited

ABOUT THE INFORMATION STORAGE INDUSTRY CONSORTIUM

The Information Storage Industry Consortium (INSIC) is the research consortium for the worldwide information storage industry. INSIC membership consists of corporations, universities and government organizations with common interests in the field of digital information storage. Corporate membership includes major information storage product manufacturers and companies from the storage industry infrastructure. INSIC has its headquarters in San Diego and was originally incorporated in April 1991 as NSIC, a non-profit mutual benefit corporation.

For further information or to contact INSIC please go to http://www.insic.org/

Copyright ©2019 INSIC. All rights reserved worldwide. All other trademarks and registered trademarks are property of their respective owners. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, without prior written permission. All opinions in this paper are those of INSIC and based on information from company material, industry reports, and interviews.