

in2 Recording Areal Density on Strontium Ferrite Tape

Mark Lantz ¹, Simeon Furrer ², Patrick Ebermann ², Hugo Rothuizen ², Walter Haeberle ², Giovanni Cherubini ², Roy D. Cideciyan ², Shinji Tsujimoto ², Yoshihiro Sawayashiki ², Yuto Murata ², Tomohide Ueyama ², Yoichi Akano ², Tetsuya Kaneko ², Hodaka Suzuki ², Masashi Shirata ², Kenji Naoi ², Takashi Koike ², Hiroaki Doshita ², and Noriko Imaoka ²

¹IBM Research

²Affiliation not available

October 30, 2023

Abstract

The recording performance of a new prototype magnetic tape based on perpendicularly oriented strontium ferrite particles is investigated using a 29 nm wide tunneling magnetoresistive reader. At a linear density of 702 kbpi, a post-detection byte-error rate of 2.8×10^{-2} is demonstrated based on measured recording data and a software read channel. The read channel uses a 64-state implementation of an extended version of a data-dependent noise-predictive maximum-likelihood detection scheme that tracks the first and second order statistics of the data-dependent noise. At the demonstrated post-detection byte-error rate, a post-error-correction-coding byte-error rate of less than 1×10^{-20} can be achieved using an iterative decoding architecture. To facilitate aggressive track-density scaling, we made multiple advances in the area of track following. First, we developed a new timing-based servo pattern and implemented a novel quad channel averaging scheme. Second, we developed a new field programmable gate array prototyping platform to enable the implementation of quad channel averaging. Third, we enhanced our low disturbance tape transport with a pair of 20 mm diameter air bearing tape guides and a prototype track-following actuator. Fourth, we developed a novel low friction tape head and finally, we designed a set of tape speed optimized track-following controllers using the model-based H_{∞} design framework. Combining these technologies, we achieved a position error signal (PES) characterized by a standard deviation [?] 3.18 nm over a tape speed range of 1.2 to 4.1 m/s. This magnitude of PES in combination with a 29 nm wide reader enables reliable recording at a track width of 56.2 nm corresponding to a track density of 451.9 ktpi, for an equivalent areal density of 317.3 Gb/in².

"This work has been submitted to the IEEE for possible publication. Copyright may be transferred without notice, after which this version may no longer be accessible."

317 Gb/in² Recording Areal Density on Strontium Ferrite Tape

Simeon Furrer¹, *Senior Member, IEEE*, Patrick Ebermann¹, Mark A. Lantz¹, *Fellow, IEEE*, Hugo Rothuizen¹, Walter Haeblerle¹, Giovanni Cherubini¹, *Fellow, IEEE*, Roy D. Cideciyan¹, *Fellow, IEEE*, Shinji Tsujimoto², Yoshihiro Sawayashiki², Noriko Imaoka², Yuto Murata², Tomohide Ueyama², Yoichi Akano², Tetsuya Kaneko², Hodaka Suzuki², Masashi Shirata², Kenji Naoi², Takashi Koike², and Hiroaki Doshita²

¹ IBM Research - Zurich, CH-8803 Rüschlikon, Switzerland

² Recording Media Research Laboratories, FUJIFILM Corporation, Odawara 250-0001, Japan

The recording performance of a new prototype magnetic tape based on perpendicularly oriented strontium ferrite particles is investigated using a 29 nm wide tunneling magnetoresistive reader. At a linear density of 702 kbp/in, a post-detection byte-error rate of 2.8×10^{-2} is demonstrated based on measured recording data and a software read channel. The read channel uses a 64-state implementation of an extended version of a data-dependent noise-predictive maximum-likelihood detection scheme that tracks the first and second order statistics of the data-dependent noise. At the demonstrated post-detection byte-error rate, a post-error-correction-coding byte-error rate of less than 1×10^{-20} can be achieved using an iterative decoding architecture. To facilitate aggressive track-density scaling, we made multiple advances in the area of track following. First, we developed a new timing-based servo pattern and implemented a novel quad channel averaging scheme. Second, we developed a new field programmable gate array prototyping platform to enable the implementation of quad channel averaging. Third, we enhanced our low disturbance tape transport with a pair of 20 mm diameter air bearing tape guides and a prototype track-following actuator. Fourth, we developed a novel low friction tape head and finally, we designed a set of tape speed optimized track-following controllers using the model-based H_∞ design framework. Combining these technologies, we achieved a position error signal (PES) characterized by a standard deviation ≤ 3.18 nm over a tape speed range of 1.2 to 4.1 m/s. This magnitude of PES in combination with a 29 nm wide reader enables reliable recording at a track width of 56.2 nm corresponding to a track density of 451.9 ktp/in, for an equivalent areal density of 317.3 Gb/in².

Index Terms— magnetic recording, magnetic tape recording, signal detection, tracking

I. INTRODUCTION

THE areal-density scaling of hard disk drives (HDD) has slowed from a compound annual growth rate (CAGR) of over 100% in the late nineties, to about 39% in the early 2000's and to an average of about 7.6% over the last several years [1]. This slow-down in HDD scaling in combination with the ongoing exponential growth in data is driving demand for scalable, cost-effective storage solutions. State-of-the-art tape systems provide high data rates and very low total cost of ownership (TCO) for storing large volumes of data, with the drawback of a higher access latency compared to alternatives such as HDD or flash. This trade-off makes tape well suited for the long-term storage of less frequently accessed data that can tolerate higher latencies, i.e. “cold” data. As a result, tape systems are increasingly used for the storage of cold data in tiered storage infrastructure, both in on-premise data centers as well as in cloud and hybrid cloud solutions. As the amount of data produced continues to grow, the need for cost-effective storage solutions will also grow and hence it is critical to continue scaling the areal density and capacity of tape systems.

State-of-the-art commercial magnetic tape drives provide a native cartridge capacity of 20TB and operate at an areal density of about 11.7 Gb/in², which is today about two orders of magnitude lower than the areal density of HDD [1]. Considering that tape systems utilize the same basic magnetic recording principles as HDD, there is a significant potential for tape to continue scaling before it encounters the challenges currently faced by HDD. Indeed, the 2019-2029 Information Storage Industry Consortium (INSIC) Tape Roadmap projects that the areal density of magnetic tape systems will scale with a

CAGR of about 34%, leading to an expected tape areal density of 278 Gb/in² in 2029 [1].

In this paper, we demonstrate the viability of continuing to scale tape areal density beyond the projections of the INSIC roadmap by performing a single-channel areal density demonstration of 317 Gb/in² on a new strontium ferrite (SrFe) tape medium. This is an increase of more than 2.5x in areal density compared to our previous record for particulate tape media of 123 Gb/in² [2] and more than 1.57x the previous record obtained using sputtered tape media [3]. In the present work, we demonstrate the viability of a 317 Gb/in² operating point by combining a modest increase in linear density with more aggressive track-density scaling, similarly to the INSIC Tape Roadmap projections. Key enabling technologies for this demonstration include the new SrFe medium, a narrow tunneling magnetoresistive (TMR) reader with 29 nm width, an enhanced 64-state data-dependent noise predictive maximum likelihood detector, a new servo pattern, four channel servo averaging, 20 mm rollers, a prototype track-following actuator and a novel low friction tape head.

This paper is organized as follows. In Section II we describe the new SrFe tape media technology. In Section III we evaluate the recording performance of the SrFe media. In Section IV we describe the tape path, servo technologies and track-following performance achieved on the SrFe media. Finally, we present our conclusions in Section V.

II. MEDIA DEVELOPMENT

Enhancing the recording performance of tape media through signal-to-noise-ratio (SNR) improvements is key to enabling

the continued areal-density scaling of tape systems. In this work, we developed a novel prototype tape media that utilizes a SrFe particle-based recording layer. We applied the following three technologies, that are described in detail below, to optimize the recording performance of the new tape media: 1) a reduced magnetic particle volume of 950 nm^3 ; 2) enhanced perpendicular orientation; and 3) a smooth surface design. Table 1 presents the properties of the prototype SrFe tape and for comparison also lists the properties of current best of breed commercial tape (IBM¹ TS1160 JE media).

TABLE 1
Magnetic Particle and Tape Properties

Media type	Prototype SrFe media	TS1160 JE media
Magnetic particle		
Volume (nm^3)	950	1700
Coercivity (kA/m)	208	197
σ_s ($\text{A} \cdot \text{m}^2/\text{kg}$)	46	48
Media		
Magnetic properties		
Longitudinal direction		
Coercivity (kA/m)	115	147
Squareness Ratio	0.17	0.29
Perpendicular direction ¹⁾		
Coercivity (kA/m)	271	222
Squareness Ratio	0.87	0.77
Surface roughness		
Optical interferometry ²⁾		
Ra (nm)	0.7	1.1
AFM ³⁾		
Ra (nm)	1.1	2.0
Rz (nm)	17	42

1) With demagnetization compensation, 2) Measurement area of $474 \mu\text{m} \times 355 \mu\text{m}$, 3) Measurement area of $40 \mu\text{m} \times 40 \mu\text{m}$

1. Fine SrFe Particles

One method to increase the SNR of magnetic tape media is to reduce the volume of the magnetic particles used in the recording layer. The scaling potential of earlier generations of tape that used iron-cobalt-based metal particles (MP) encapsulated with a passivation layer was limited by the difficulty in maintaining a sufficient coercivity to ensure the long-term stability of recorded data if the particle volume was reduced below about 3000 nm^3 . Most current tape media use barium ferrite (BaFe) particles which enable a high coercivity to be maintained as the particle volume is reduced. SrFe is from the same family of hexagonal ferrites as BaFe, with the same crystalline structure except that the element barium is replaced with strontium. Both BaFe and SrFe are chemically stable oxides that do not need the passivation layer that MP media required to prevent oxidation, and therefore both particle types can be scaled to smaller particle volumes than MP. SrFe has the additional benefit of having a higher magnetic anisotropy than BaFe [4] and can be scaled to even smaller particle volumes while maintaining a suitable coercivity for archival stability. This combination of properties makes SrFe an excellent candidate for use in future tape media.

Our new prototype media uses SrFe particles with a volume of 950 nm^3 , about 44 % smaller than the BaFe particles used in JE media. Fig. 1 shows a top view scanning electron microscope (SEM) image of the new media along with an image of JE media for comparison. Fig. 2 presents cross-sectional transmission electron microscope (TEM) images of the two media types. Very fine and uniform-sized particles can be seen in the SEM image of the SrFe tape surface (Fig. 1) and the cross section TEM image of the recording layer (Fig. 2). The 950 nm^3 SrFe particles used here have a thermal stability factor of $K_u V / k_B T = 64$, measured using the method described in [5]. This is sufficient to ensure the long-term archival stability of the media. Here, K_u is the anisotropy energy density, V the particle volume, k_B Boltzmann's constant, and T the absolute temperature.




Fig. 1. Tape surface images measured by SEM (left) Prototype SrFe tape, (right) TS1160 JE media.




Fig. 2. Cross-sectional images of recording layer observed by TEM (top) Prototype SrFe tape, (bottom) TS1160 JE media.

2. Perpendicular Orientation

The recording performance of the media can be enhanced by orienting the anisotropy axes of the particles towards the perpendicular direction [6]. Both SrFe and BaFe particles are platelet-shaped with a magnetic axis perpendicular to the platelet face. During coating, the shape of the platelets naturally leads to a somewhat preferential orientation of their faces parallel to the surface of the tape substrate. The width of their angular distribution can be further narrowed by applying a magnetic field during the coating process, thereby leading to a yet higher degree of magnetic orientation. To further improve the orientation in the perpendicular direction, care must be taken to reduce the aggregation of particles in the coating solution, which also negatively impacts their net magnetization strength as well as impacting recording noise performance. Unfortunately, decreasing the particle volume increases their tendency to aggregate. To achieve a good dispersion of the very fine SrFe particles used here, we optimized the dispersant

¹ IBM is a trademark of International Business Machines Corporation, registered in many jurisdictions worldwide.

formulation and the dispersion conditions. As clearly seen in the cross sectional TEM images of the recording layers in Fig. 2, the degree of perpendicular orientation of the SrFe particles is increased compared with the particles in the JE tape, resulting in a perpendicular squareness ratio (with demagnetization compensation) of 0.87 versus 0.77 for JE media as shown in Table 1.

3. Smooth Surface Design

Another means of improving SNR is to reduce the head-media spacing by decreasing the roughness of the tape surface. In this work, we achieve this by improving the smoothness of the underlayer provided beneath the recording layer in the coating stack. Fig. 3 presents surface profile images, measured using optical interferometry, of the prototype SrFe tape and a JE tape for comparison. In addition, Fig. 4. shows $40\ \mu\text{m} \times 40\ \mu\text{m}$ atomic force microscopy (AFM) images of the surface of the prototype tape and a JE tape. The optical interferometry images enable the characterization of the long-range surface roughness whereas the AFM images provide insight into the short-range surface roughness, both of which contribute to the head-media spacing. On both scales, the prototype SrFe tape is much smoother than the JE tape. The average roughness of the SrFe tape measured by optical interferometry is $R_a = 0.7\ \text{nm}$ compared to $R_a = 1.1\ \text{nm}$ for the JE tape. The SrFe tape, as measured by AFM, is especially smooth; $R_a = 1.1\ \text{nm}$ for SrFe tape compared to $R_a = 2.0\ \text{nm}$ for the JE tape.




Fig. 3. $474\ \mu\text{m} \times 355\ \mu\text{m}$ surface profile images measured by optical interferometry (left) Prototype SrFe tape, (right) TS1160 JE media.




Fig. 4. $40\ \mu\text{m} \times 40\ \mu\text{m}$ surface profile images measured by AFM (left) Prototype SrFe tape, (right) TS1160 JE media.

III. RECORDING PERFORMANCE

1. Experimental Setup and Procedure

We investigated the recording performance of the prototype SrFe tape described in Section II by writing data tracks with an $8.7\ \mu\text{m}$ wide tape write head equipped with a high moment liner [7] and reading the data back using an HDD read sensor. The

experiments were carried out with a commercial tape transport system from Mountain Engineering II Inc. (Longmont, CO, USA) that uses hard edge guiding to limit the amplitude of lateral tape motion (LTM) to the scale of a few micrometers peak to peak. The low amplitude of LTM provided by the hard edge guiding in combination with a precisely positioned narrow read sensor enables reliable readback of the wide tracks without the need for closed-loop track-following. This, in turn, enables the use of HDD heads with very narrow read sensors to explore future operating points for which tape heads have not yet been developed. We conservatively estimate the magnetic width of the side shielded TMR readers used here to be $29\ \text{nm}$ based on measurements of the physical width of the reader by means of an SEM and extrapolating the magnetic width from the plot of physical versus magnetic reader width presented in [8].

The data tracks were composed of repeating 255-bit pseudo-random binary sequences (PRBS) and were written on tape at linear densities ranging from 500 to 768 kbp. The tape was pre-erased using permanent magnets placed upstream of the write head. The PRBS sequence was generated using a linear feedback shift register with the polynomial $x^8 + x^4 + x^3 + x^2 + 1$. Prior to performing readback experiments with HDD heads, the optimal write current was determined using a write saturation experiment in which the writer current was stepped from 13.2 to 33.8 mA and the recorded PRBS data was read using a tape read head positioned immediately downstream of the writer. The SNR values of the captured readback signals were determined by processing the data using a software read channel with an extended partial-response class-4 (EPR4) detector and plotted versus write current as shown in Fig. 5. The optimal write current is taken as the point that produces the maximum SNR. The observation of a clear maximum in Fig. 5 is an indication that the high moment liner write head is able to fully saturate the prototype SrFe media.




Fig. 5. SNR versus writer current at 702 kbp linear density measured on the prototype SrFe media with a tape read head. Maximum SNR is obtained at a write current of 22mA as indicated by the circle.

Readback experiments with HDD readers were performed by gluing the slider of a head gimbal assembly (HGA) onto a custom-made fixture to constrain the gimbal mechanism and rigidly fix the slider to the fixture. The fixture was mounted on a set of micrometer positioning stages in the tape path, as illustrated and described in [7]. The positioning stages facilitate the precise alignment of the position and angle of the HDD slider relative to the tape. The experiments were performed with the head in contact with the tape, with the slider positioned such

that the edge closest to the TMR read element comes into contact with the moving tape first. We denote this edge as the leading-edge, as per convention in tape recording. For the HDD heads used here, the TMR read element is located about $31\ \mu\text{m}$ from the leading edge of the slider and is approximately centered across the width of the slider. Readback experiments were performed while running the tape over the slider with a leading-edge wrap angle of $1.5 - 2.0^\circ$ and a trailing edge wrap angle of about zero degrees. During tape motion, the leading-edge wrap angle in combination with the topography of the surface of the slider results in regions of reduced air pressure between the tape and slider, relative to the air pressure above the tape. This pressure difference pushes the tape into contact with the slider. In our previous experiments with HDD heads on tape [9], we found that the best readback SNR was obtained at a wrap angle of about 0.5° , whereas in the present work we achieved optimal SNR performance with a slightly larger wrap angle of $1.5-2.0^\circ$. The difference in optimal angle most likely results from differences in the air bearing design of the two types of HDD heads.

The 29 nm wide side shielded TMR readers used to characterize the SrFe media were biased with a current of 0.3 mA using the electronics card from an IBM TS1140 tape drive. In addition, the card provided amplification, filtering and digitization of the readback signal. The readback signal was sampled at a rate of $5/(4T)$, where T is the bit duration, and captured using a custom-built data acquisition system. The captured data was processed offline with a software read channel that uses double-precision floating point math to implement all of the signal processing functions of a tape drive read channel.

2. Error Correction Coding Performance Analysis

Compared to the much larger reader widths used in current tape drives, the 29 nm wide HDD reader used here leads to a significant reduction of SNR, which in turn results in higher error rates at the output of the detector. The large loss in SNR can be partially compensated by means of more complex detectors. In addition, an improved error-correction-coding (ECC) / decoding scheme can enable reliable performance in spite of an increase in error rate at the output of the detector.

The error-correction coding scheme employed in state-of-the-art linear tape drives uses product codes in combination with deep interleaving along and across the tape tracks. Each product code word is constructed by encoding a size $K_1 \times K_2$ array of data bytes into a size $N_1 \times N_2$ product code word, where each row is an (N_1, K_1) Reed–Solomon (RS) code word referred to as a C1 code word and each column is an (N_2, K_2) RS code word referred to as a C2 code word.

A simple ECC decoder for these product codes applies an RS decoder sequentially on the two component codes. In previous work [7,10], we have shown that iterative decoding of product codes can provide significant improvements in ECC performance. The iterative decoding technique still employs hard-decision bounded-distance decoding on each of the component codes but employs multiple iterations of the C1 and C2 decoding steps. Here we use a probabilistic analysis and hardware simulations to evaluate the performance of the

product code with parameters $N_1=240$ and $K_1=228$, and $N_2=192$ and $K_2=168$ that is used in state-of-the-art tape drives. We use the results of this analysis to set a target threshold for the detector output error rate for the recording performance evaluation of the SrFe media.

Fig. 6 shows the post-decoding byte-error-rate performance, i.e., the BER at the output of the iterative ECC decoder after a selected number of m decoding steps as a function of the raw byte-error rate, i.e., the byte-error rate at the output of the detector. The symbols in Fig. 6 correspond to numerical results obtained by an field programmable gate array (FPGA) based hardware simulator decoding up to 200 billion 46 080-byte product code words per data point (symbol). The legend for each curve in Fig. 6 contains the sequence of decoding steps applied. For example, C1-C2 denotes the results after a first C1 decoding step followed by C2 decoding as a second step, while C1-C2-C1-C2 represents the results after two C1-C2 iterations corresponding to $m = 4$ decoding steps. The curve “undecoded” corresponds to the raw byte-error rate prior to decoding.




Fig. 6. Performance of iterative decoding of the product code with RS component codes C1 ($N_1=240$, $K_1=228$) and C2 ($N_2=192$, $K_2=168$) after a selected number of C1 and C2 iterations.

To estimate the performance at very low post-decoding byte-error rates, we have extended the probabilistic analysis technique described in [11] to the case of unequal lengths and error-correction capabilities of the component codes. Both the probabilistic analysis and the simulation assume uncorrelated byte errors at the input of the ECC decoder, which leads to random byte-error locations within the product code words and hence the errors in rows and columns are expected to follow a binomial distribution. Fig. 6 shows the results of this probabilistic analysis as dash-dotted lines. For up to three decoding steps, the simulation and probabilistic results show an excellent match. After $m \geq 4$ decoding steps, the probabilistic analysis underestimates the post-decoding byte-error rate due to some simplifying assumptions in the analysis [11]. Extrapolating the simulation results we estimate that a post-decoding byte-error rate of less than $1\text{e-}20$, corresponding to the specified performance of a state-of-the-art enterprise class tape drive, can be achieved with two full decoding iterations ($m = 4$ decoding steps) provided that the raw byte-error rate is less than

4.5e-2. A raw byte-error rate of up to 5e-2 is acceptable if $m \geq 5$ decoding steps are used.

Based on the ECC performance results and the assumption of a reverse concatenation architecture we set a raw byte-error rate (BER) target of less than 4.5e-2 at the output of the detector to evaluate the recording potential of the new media.

3. Experimental Results and Analysis

The spectrum of the timing synchronized PRBS data recorded at a linear density of 702 kbp/s is presented in the upper part of Fig. 7. The middle of the figure presents the dibit response of the recording channel determined using the same data set and the bottom part presents a plot of the SNR of the data versus time at the input of an EPR4 detector. The shape of the dibit response is as expected for a well oriented perpendicular magnetic media. Small echoes are observed in the dibit response at distances of $25T$ and $-58T$ from the main pulse. These perturbations are an indication of non-linearities in the channel such as hard transition shift, reader non-linearity and adjacent-bit non-linear transition shift (NLTS).




Fig 7. Spectrum of the timing synchronized PRBS data (top) and dibit response (middle) of the recording channel and (bottom) SNR versus time at the input of the EPR4 detector of PRBS data recorded at 702 kbp/s.

Fig. 8 summarizes the results of the recording performance evaluation. The upper part of Fig. 8 presents the SNR at the input of an EPR4 detector plotted versus linear density. The middle and lower parts of the figure plot the bit- and byte-error rates versus linear density, obtained by processing the same readback signals using the following five detection algorithms listed in order of increasing complexity [12]-[14]: 1) 8-state EPR4 detection, 2) adaptive 16-state noise-predictive maximum-likelihood (NPML) detection with target polynomial $F(D)=(1-D^2)(1-p_1D-p_2D^2)$, 3) adaptive 16-state data-dependent NPML (DD-NPML) detection where the noise-prediction filter coefficients $p_1(\underline{a})$ and $p_2(\underline{a})$ and the prediction error depend on the data sequence \underline{a} , 4) extended 16-state DD-NPML detection that additionally tracks the mean of the data-dependent noise (D3-NPML), and 5) adaptive 64-state D3-NPML detection with generalized target polynomial $F(D)=(1-D^2)(1-p_1(\underline{a})D-p_2(\underline{a})D^2-p_3(\underline{a})D^3-p_4(\underline{a})D^4)$. The implementation effort for 64-state detection was decreased by halving the number of parallel noise-prediction filters such that they are shared between trellis branches. This reduced the complexity considerably, while only a very minor increase in error rate was observed.




Fig. 8. (Top) SNR versus linear density, (middle) bit-error rate and (bottom) byte-error rate versus linear density and detector type. The target byte-error rate threshold of 4.5e-2 is indicated by a dashed line in the lower plot.

Each data point was obtained by processing approximately 6 million samples, i.e., about 4.8 million bits. Both the bit- and byte-error rates improve with increasing detector complexity, with more pronounced gains observed at higher linear densities. The best performance is achieved using the 64-state D3-NPML detector which achieves a BER of $4.5e-2$ at a linear density of 750 kbp, just at our target BER threshold of $4.5e-2$. At a linear density of 702 kbp the 64-state D3-NPML detector achieves a BER of $2.8e-2$ and all detectors except the EPR4 detector have a BER performance below the target threshold. The average channel SNR at linear densities of 750 kbp and 702 kbp are 9.9 dB and 10.6 dB, respectively. Thus, assuming the use of the 64-state D3-NPML detector, an operating point of 702 kbp provides an SNR margin of about 0.7 dB, similar to our previous demos.

4. Error Event Analysis

As discussed in subsection 3 above, the BER threshold used to determine a reliable operating point was set based on the performance of the selected iterative ECC decoding scheme that was investigated by means of a probabilistic analysis and simulations. Both the probabilistic analysis and the simulations were based on an assumption of uncorrelated byte errors at the decoder input. In order to confirm that this assumption is valid for our data, we analyzed both the error events in the data bit stream at the output of the selected detectors, and the distribution of byte errors at the input of the ECC decoder. For the data recorded at 702 kbp, Table 2 lists the typical bit error events (up to length seven) at the output of the five detectors discussed above, and their relative frequency of occurrence (in percentage) for each of the detectors. The error events are defined as half of the difference between the written sequence $\{a_k\}$ on tape, $a_k \in \{+1, -1\}$, and the detected sequence $\{\hat{a}_k\}$. Here a shorthand notation is used to represent a ternary error event and the corresponding negative error event: e.g. error event “+−” is used to denote both error events “+1, −1” and “−1, +1”. Similarly, “+0+” denotes the error events “+1, 0, +1” and “−1, 0, −1”. Table 2 shows that the distribution of error events gets skewed more to shorter error events when the complexity of the detector is increased. With EPR4 detection, only about 73% of error events are of length two or less, while with 64-state D3-NPML, 88% are of length two or less. Similarly, the relative frequency of occurrence of all the length three and four error events drops from about 16% to 7% for EPR4 and 64-state D3-NPML detection, respectively. Note that for a detector with 2^N states, there can be at most $(N-1)$ consecutive zeros in an error event, because N consecutive zeros would end the error event. For example, only 0.58% of error events at the output of the detector with $N=64$ states are of the type ‘+0000+’, while such an event at the output of all other detectors considered would count as two error events of the type ‘+’. To keep Tables 2 and 3 short, only error events with an occurrence of at least 0.5% are displayed. The row labeled “xxxxxxx” is a sum of all error events up to length seven not listed in the rows above. The last row labeled ‘other’ sums the occurrence of all error events of length greater than seven. Table 3 lists the typical error events and their relative frequency of occurrence (in percentage) for the 64-state D3-NPML detector at the linear densities of {552, 608, 662, 702, 750} kbp. Over this linear

density range, error event “+” is dominant, but with decreasing relative occurrence as the linear density increases. On the other hand, longer error events such as “+−”, “+−+” or “+00+” increase in relative occurrence. At 552 kbp, 98% of error events are of length four or less, and only 0.2% are longer than or equal to the length of one byte (8 bits). The corresponding numbers at 750 kbp are 92% for error event of length four or less, and still only 2.28% are longer than or equal to one byte.

TABLE 2
Percent Occurrence of Bit Error Events at 702 kbp

Error Events	EPR4	NPML	DD-NPML	D3-NPML	D3-NPML 64-state
+	65.31	76.78	74.95	72.85	73.34
+−	7.60	10.10	11.44	14.05	14.83
+0+	3.26	1.98	1.96	1.54	0.95
+−+	8.56	4.11	4.91	5.37	4.31
+00+	1.76	1.79	1.55	1.26	0.77
+−+−	2.01	0.37	0.39	0.45	0.53
+000+	n.a.	1.40	1.33	1.21	0.77
+0+0+	0.94	0.33	0.31	0.18	0.07
+−+−+	2.38	0.32	0.31	0.33	0.49
+0000+	n.a.	n.a.	n.a.	n.a.	0.58
+−+−+	0.98	0.11	0.11	0.09	0.12
+0+0+0+	0.65	0.16	0.20	0.12	0.02
xxxxxxx	3.28	1.61	1.79	1.85	2.20
other	3.27	0.94	0.75	0.70	1.02

Percent occurrence of bit error events for each detector type at a linear density of 702 kbp. Only events with an occurrence of at least 0.5% are displayed. The row labeled xxxxxx is the sum of all error events up to length seven not listed above. The row labeled other sums the occurrence of all error events of length greater than seven.

TABLE 3
Percent Occurrence of Bit Error Events vs. Linear Density for D3-NPML 64-state

Error Events	552 kbp	608 kbp	662 kbp	702 kbp	750 kbp
+	84.78	83.28	79.83	73.34	68.18
+−	9.53	10.79	12.64	14.83	14.42
+0+	0.98	0.91	1.12	0.95	1.19
+−+	2.33	1.67	1.97	4.31	5.97
+00+	0.44	0.42	0.64	0.77	1.02
+−+−	0.33	0.27	0.23	0.53	0.85
+000+	0.35	0.54	0.71	0.77	1.04
+−+−+	0.09	0.08	0.11	0.49	0.87
+0000+	0.13	0.30	0.37	0.58	0.85
+00000+	0.15	0.14	0.35	0.47	0.55
xxxxxxx	0.68	1.25	1.52	1.94	2.78
other	0.21	0.35	0.51	1.02	2.28

To further analyze the distribution of errors, the data bit stream at the output of the detector was converted to a byte stream and then subsequently mapped to 4-way byte-interleaved C1 code words, as is done in state-of-the-art tape

drives. This 4-way byte interleaving breaks up long bit error events by placing them in different C1 code words, and therefore avoids correlated errors in adjacent bytes of a C1 code word unless the bit error events are longer than 25. After interleaving, we plot the probability distribution of k byte errors within a C1 code word of length $N_1 = 240$ bytes and compare the result with a binomial distribution that is expected for uncorrelated errors. Fig. 9 presents examples of the error distributions after processing the data captured at 702 kbp/s with the NPML detector that provides a BER of 4.2×10^{-2} (top) and after processing with the 64-state D3-NPML detector that provides a BER of 2.8×10^{-2} (bottom). The corresponding C1 uncorrectable code word rates (C1uc) are 0.8071 for NPML and 0.4882 for 64-state D3-NPML detection. For the 64-state D3-NPML processed data, comparison with the binomial distribution (red) shows an excellent match, confirming the assumption of uncorrelated errors. For the same data processed with the NPML detector, a relatively good match is also obtained except for a few C1 code words containing 30 or more byte-errors. These events likely result from fluctuations of the signal strength visible in the SNR time-trace presented in Fig. 7 that are probably caused by media imperfections giving rise to short sections with increased error rates. State-of-the-art tape drives implement a read-while-write verification process that detects long burst errors and rewrites affected code words in a defect-free region of tape. Hence, the small number of burst errors observed here that result from media imperfections do not invalidate our assumption of uncorrelated byte errors.

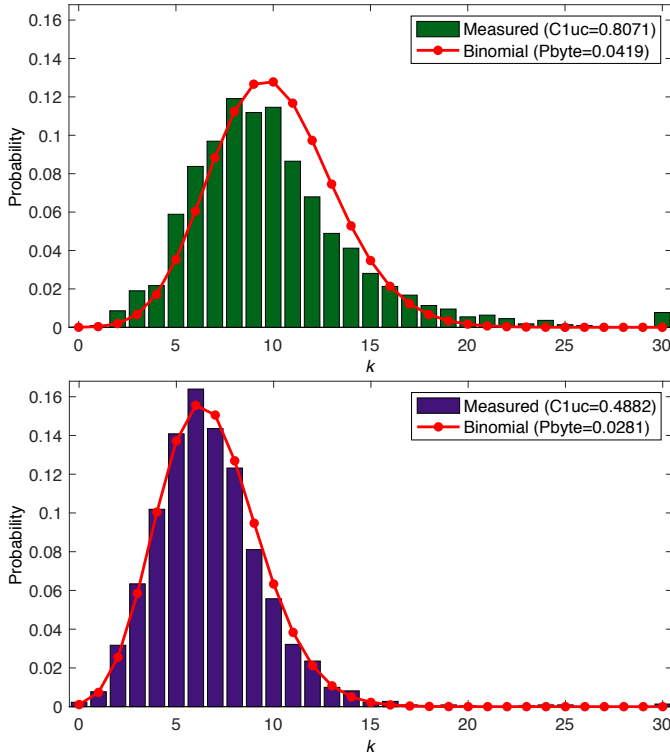


Fig. 9. Probability of k byte errors within a 240 byte C1 code word for the data captured at 702 kbp/s with NPML detection (top) and 64-state D3-NPML detection (bottom). Note that all code words with 30 or more errors are placed in bin 30 of the histogram.

IV. TRACK-FOLLOWING SERVO PERFORMANCE

The latest INSIC Tape Roadmap [1] foresees that future increases in tape areal-density scaling will be achieved by modest increases in linear density combined with aggressive track-density scaling. To explore the track-density scaling potential of our prototype SrFe media and to enable the aggressive track-density scaling projected by the INSIC Tape Roadmap we made multiple advances in the area of track-following servo control. First, we developed a new timing-based servo (TBS) pattern. Second, we implemented a novel four servo channel averaging scheme using a new custom built FPGA prototyping board that enables the implementation of four servo channels. Third, we enhanced our low disturbance tape transport with a pair of 20 mm diameter air bearing tape guides and a new prototype track-following actuator. Fourth, we developed a novel low friction tape head that mitigates friction effects that arise from the use of the smooth prototype SrFe media and finally, we designed a set of tape speed optimized track-following controllers using the model-based H_∞ design framework [15], [16].

1. Servo Pattern and Quad Servo Channel Averaging

To improve the resolution of the head position measurement, we developed a new timing-based servo (TBS) pattern. The geometry of the pattern is shown in Fig. 10 and is characterized by an azimuthal angle of 36° , a stripe pitch of $2.4 \mu\text{m}$, a subframe length of $64.2 \mu\text{m}$ and a height of $25 \mu\text{m}$. The 36° azimuthal angle reduces measurement noise compared to the smaller angle patterns used in current products and for our previous demos. The $25 \mu\text{m}$ pattern height was chosen to be compatible with a potential future 64 channel, 8 data band tape layout projected by the INSIC Tape Roadmap [1]. The choice of height, angle and stripe pitch lead to the $64.2 \mu\text{m}$ subframe length, which is smaller than the $76 \mu\text{m}$ subframe used in current tape products and therefore enables a higher update rate. The $64.2 \mu\text{m}$ subframe length is longer than that used in our previous demos, however, we found that the improved resolution resulting from the 36° angle outweighs the penalty of a lower update rate.




Fig. 10. Illustration of the geometry of the experimental servo pattern.

Current commercial tape drives use a pair of read sensors to simultaneously read two servo patterns that are located above and below each data band. The position estimates from these two servo patterns can be averaged to improve the resolution of the head position measurement by a factor of about $\sqrt{2}$ [2]. In this work, we further improved the measurement resolution by an additional factor of about $\sqrt{2}$ by implementing a quad servo

channel averaging scheme in which the position estimates from four simultaneously read servo patterns are averaged. This scheme requires two servo patterns in each of the nine servo bands of the assumed eight data band tape layout. However, because of the relatively small $25\ \mu\text{m}$ height of our experimental servo pattern, this design would not increase the servo format overhead compared to state-of-the-art commercial linear tape formats.

During tape transport, the four servo patterns are simultaneously read using four $1.0\ \mu\text{m}$ wide TMR readers. Fig. 11 shows the readback signals captured from the four readers at a tape speed of $2.15\ \text{m/s}$. The lateral position of the head and the velocity of the tape are determined by measuring the time intervals between bursts of dibits in the servo readback signal as described in [17]. The 36° azimuthal angle of our prototype servo pattern results in a significant dispersion of the dibit shape in the readback signal that reduces the position estimation accuracy. This dispersion can be mitigated by decreasing the width of the servo reader, albeit at the cost of SNR [17]. The $1.0\ \mu\text{m}$ reader width used here was found experimentally to provide a good compromise between dispersion and SNR as well as providing robustness against media defects compared to a smaller reader width.




Fig. 11. Quad read-back signals (bursts of dibits) from four $1.0\ \mu\text{m}$ wide TMR readers and four servo patterns at $2.15\ \text{m/s}$ tape speed.

2. Hardware Platform and Servo Channel

Closed-loop track-following experiments were performed with the hardware platform depicted in Fig. 12. The platform consists of a low disturbance tape transport system and a prototype track-following actuator equipped with a novel low friction tape head described in subsections 3 and 4 below. In addition, the platform includes the electronics card of an IBM TS1140 tape drive, an FPGA-based prototyping board with high-speed digital to analog converters (DACs), a voltage-controlled current source (VCCS) and a host computer (PC). The TS1140 electronics card was used to provide current biasing for the four servo readers as well as amplification, filtering and digitization of the servo read back signals. In addition, the electronics card was used for reel-to-reel control during tape transport. The servo signals from the four TMR

readers were sampled at a fixed clock rate of $41.66\ \text{MHz}$ and transmitted from the drive card to the FPGA board in LVDS format. Each servo signal is processed using the matched-filter-based servo channel described in [17] which was adapted for the experimental servo pattern. In addition, the coefficients of the matched-filter were tuned to the dibit shape of the readback signal presented in Fig. 11 to optimize the estimation accuracy. The new FPGA board uses an FPGA that is large enough to implement four instances of the servo channel that run in parallel to enable quad servo channel averaging and hosts a soft-core microprocessor running at a clock rate of $240\ \text{MHz}$. The microprocessor is used to run the track-following controllers described in subsection 5. Synchronous track-following control operation is implemented by running the controller in an interrupt-driven mode, where interrupt requests are triggered by the servo channels whenever a new lateral position estimate is available. The computed control effort is output via a high-speed DAC to the VCCS that drives the prototype track-following actuator. Compared to our previous FPGA board [7], the new board enables lower latency due to the faster microprocessor and high-speed DACs and hence enables the use of higher bandwidth track-following controllers.




Fig. 12. Experimental hardware platform used to evaluate the servo performance of the SrFe media.

3. Tape Transport and Prototype Track-Following Actuator

The low disturbance tape path used for closed-loop track-following experiments is depicted in the top part of Fig. 12. The tape is guided from the supply reel to the take up reel over the head by four air bearing tape guide rollers. The two rollers adjacent to the head have a diameter of 20mm whereas the two rollers adjacent to the reels have a diameter of 25mm. The smooth surface of the prototype media investigated here causes significant head-media friction. This friction excites compressional waves that oscillate between the tape-path guide rollers on either side of the head and result in high-frequency periodic disturbances in the servo signal [18]. In previous work we found that the amplitude of these disturbances can be reduced by decreasing the distance between the guide rollers on either side of the tape head, however, the minimum spacing of the guide rollers adjacent to the head was limited by the width of the track-following actuator and the 25mm diameter of the rollers. In order to enable a lower roller spacing in the present work we replaced the rollers adjacent to the tape head with 20 mm diameter rollers and designed a new prototype narrow-body actuator that is depicted in Fig. 13. This combination enabled the use of a 24 mm spacing between the rollers adjacent to the head in the experiments described below.




Fig. 13. Prototype narrow-body track-following actuator.

The actuator consists of an aluminum support beam mounted on a pair of 60 μm thick stainless-steel leaf springs that connect the beam to mechanical ground. A tape head is glued onto the front of the aluminum beam and actuation is provided by a voice coil motor (VCM) mounted directly under the beam. The VCM consists of an elongated ‘racetrack’ coil mounted in an aluminum frame and two pairs of permanent magnets mounted on either side of the coil on a yoke structure attached to mechanical ground. The center of the coil, which is assumed to be coincident with the force vector generated by the VCM, is aligned with the center of mass of the moving parts. Damping of the actuator is provided by a damping gel deposited under the frame of the coil on the yoke structure. The leaf springs guide the head motion in the cross-track direction and suppress undesirable head motion in other directions. The thin and wide

spring design provides a low stiffness in the direction of actuation and a much larger stiffness in the direction of tape travel which helps to suppress resonance modes excited by tape head friction. The low stiffness in the direction of actuation leads to a high DC gain which enables the use of high-bandwidth controllers while maintaining a relatively small current amplitude in the control effort.

The frequency response of the actuator is shown in Fig 14. The measurement was made while running tape under closed-loop servo control with a low-bandwidth controller and applying a frequency sweep to the reference position. The resulting displacement of the actuator was measured using the servo pattern on the tape and hence the phase response data includes an additional measurement delay that results from the length of the servo pattern and the processing of the servo signal [17], [19]. Apart from the additional delay, the transfer function is similar to that of a spring-mass-damper system and is well fit by a second-order model with an additional delay term. The fundamental resonance is at about 80 Hz and is characterized by a broad heavily damped peak. Higher-order resonances are also visible at around 5 kHz, 11 kHz and 18 kHz. The low frequency part of the measured frequency response is distorted by lateral tape motion and is not included in the model fit shown in Fig. 14.

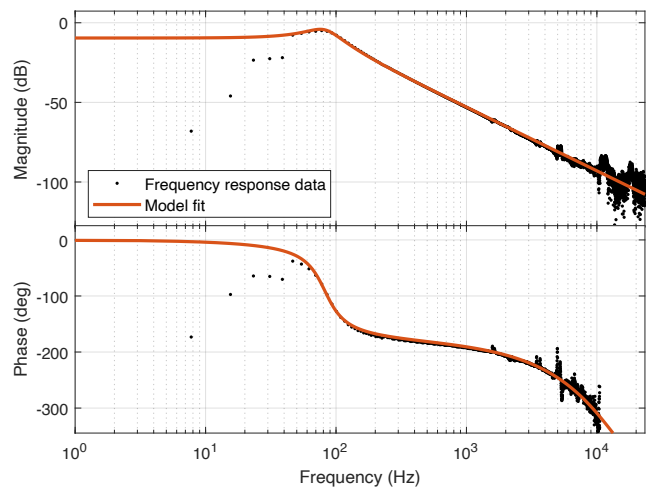


Fig. 14. Transfer function of the prototype track-following actuator measured at a tape speed of 4.08 m/s (symbols) and fit to a second-order model with an additional delay term modelled as a first order Padé approximation (solid line).

4. Low Friction Head

To further mitigate the effects of friction resulting from the smooth surface of the prototype media we developed a new low friction tape head. The new head utilizes the beveling concept described in [20] that creates an air bearing between the tape and the head in the regions away from the transducers by beveling the skiving edges of the head in these regions. Compared to the heads used in our recent demos [2], [3], [7] we reduced the height of the unbeveled region by about 45% and also reduced the width of the tape head in the direction of tape travel by about 45%. This combination resulted in a reduction in friction of about 75%.

5. Track-following Controller Design

The task of the track-following controller is to adjust the position of the head to compensate for lateral tape motion during read and write operations. In this work, the model-based H_∞ control design methodology was used for track-following servo controller design. A block diagram of the control system is shown in Fig. 15. The input to the controller $C(s)$ is the so-called position error signal (PES) which is given by the difference between the average position calculated from the quad servo channels and a target track location or reference position r . The lateral dynamics of the prototype track-following actuator are described by the transfer function $G(s)$. The effects of the loop delay are denoted by $D(s)$ in Fig. 15. This delay includes a tape speed dependent component determined by the servo pattern geometry and constant contributions from the computation time of the control effort, servo channel signal processing, DAC output delay and the frequency characteristics of the VCCS. We used the model fit shown in Fig. 14 to determine the parameters of $G(s)$ and $D(s)$. Because of the interrupt-driven synchronous control mode used here, the sampling frequency of the controller is a function of tape speed. In addition, the frequency components of the LTM also depend on tape speed, i.e., as the tape speed is increased, the LTM disturbances, denoted by d_{LTM} in Fig. 15, shift to higher frequencies while at the same time the servo pattern delay is reduced. To address the speed dependent nature of the disturbances and to take advantage of the reduction in delay with increasing tape speed, a set of four controllers were designed with each controller optimized for a specific tape speed.




Fig. 15. Block diagram of the track-following control system.

The track-following controllers were designed by specifying weighting transfer functions $W_s(s)$ and $W_u(s)$ that describe the control objectives for disturbance rejection and for limiting the control effort, respectively. The sensitivity transfer function $S(s)$ that relates the lateral tape motion to the PES, was shaped by specifying the weighting transfer function $W_s(s)$. Specifically, we used $W_s(s)$ to describe the target level of LTM suppression and the desired controller bandwidth. To enhance the LTM suppression at specific frequencies, e.g. LTM resulting from periodic roller or reel disturbances, second-order bandpass filters were added to $W_s(s)$ for the highest tape speed controller where such disturbances are shifted to relatively high frequencies. The shape of the weighting transfer function $W_u(s)$ was used to limit the control effort u and hence to reduce the amplification of the position estimation noise n . In addition, to limit the excitation of the 5 kHz and 18 kHz actuator resonances, bandpass filters were added to $W_u(s)$ at these frequencies.

The transfer function $T(s)$ relates the head position y to the

desired reference position r , i.e. $T(s) = y/r$. Track-following controllers for tape speeds of 1.23, 2.15, 3.08 and 4.08 m/s were obtained using H_∞ synthesis to solve the optimization problem that incorporates the performance requirements captured by the weighting functions. Defining the controller bandwidth as the zero dB crossing of the $S(s)$ transfer function, the bandwidth of our optimized designs ranged from 700Hz for the lowest tape speed controller to 3100Hz for the highest speed. Fig. 16 presents the transfer functions $T(s)$ and $S(s)$ of the track-following controller designed for a tape speed of 2.15 m/s.




Fig. 16. Magnitude of the closed-loop transfer $T(s)$ and sensitivity function $S(s)$ of a controller designed for operation at 2.15 m/s.

6. Experimental Results

Closed-loop track-following experiments were performed at tape speeds of 1.23, 2.15, 3.08 and 4.08 m/s that span the speed range of the TS1140 drive card used for reel-to-reel control. In order to further suppress the residual compressional wave disturbances resulting from tape-head friction, we used the technique described in [18] in which a fixed skew angle of 1.1° was applied to the head so that the components of the compressional wave disturbance in the position estimates produced from the two pairs of servo readers are in opposition of phase. Fig. 17 plots the standard deviation of the position error signal (σ_{PES}) versus speed and summarizes the results obtained. Fig. 18 presents exemplary details of the closed-loop performance obtained at 2.15 m/s including a time trace (embedded), power spectral density plot of the PES (bottom), and a cumulative plot of the PES versus frequency (top). Note that the data in the top part of the figure is plotted on same the frequency axis as the bottom part of the figure. As shown in Fig. 17, relatively uniform performance was achieved across the speed range with a worst case σ_{PES} of 3.18 nm observed at the highest tape speed and the best σ_{PES} of 3.02 nm observed at 2.15 m/s. This is a significant improvement over our previous best published results of worst case σ_{PES} of 5.9 nm over the same speed range [2].




Fig. 17. Standard deviation of the PES versus tape speed.




Fig. 18. Power spectral density (PSD) of the PES (bottom), cumulative plot of the PES versus frequency (top) and time trace (embedded) of the PES captured during closed-loop track-following at 2.15 m/s.

V. CONCLUSION

Following the approach used in our recent tape areal density demonstrations [2], [3], [7] we calculate the minimum reliable track width enabled by the demonstrated PES performance in combination with the 29nm wide reader used in the recording experiments using the formula described by INSIC [1] as:

$$\text{track width} = 2 \cdot \sqrt{2} \cdot 3 \cdot \sigma_{PES} + \text{reader width}.$$

Assuming the worst case measured σ_{PES} of 3.18 rounded to 3.2 nm and a reader width of 29 nm leads to a calculated track width of 56.2 nm and a track density of 451.9 ktpi. Combining this track density with the linear density of 702 kbpj achieved within the error rate target with the 29 nm reader corresponds to an areal recording density of 317.3 Gb/in². This is about 14% higher than the areal density of 278 Gb/in² projected by the latest INSIC Tape Roadmap for the year 2029 [1]. This work provides further experimental validation of the potential to continue scaling tape at historical rates of 33% CAGR in areal density until at least 2029, in line with the projections of the 2019-2029 INSIC Tape Roadmap.

ACKNOWLEDGMENT

The IBM authors thank their colleagues in the Tucson, San Jose, and Tokyo tape development labs for their support and in particular thank Venus Hipolito for manufacturing the low

friction head prototypes. In addition, we thank Urs Egger of the IBM Zurich Research Lab for support in the development of the FPGA board and Angeliki Pantazi for helpful discussions in the area of control design.

REFERENCES

- [1] Information Storage Industry Consortium (INSIC) 2019-2029 Tape Roadmap, <http://www.insic.org/wp-content/uploads/2019/07/INSIC-Technology-Roadmap-2019.pdf>
- [2] M. A. Lantz *et al.*, "123 Gb/in² recording areal density on barium ferrite tape", *IEEE Trans. Magn.*, vol. 51, no. 11, Nov. 2015, Art. ID 3101304.
- [3] S. Furrer *et al.*, "201 Gb/in² recording areal density on sputtered magnetic tape", *IEEE Trans. Magn.*, vol. 54, no. 2, Feb. 2018, Art. ID 3100308.
- [4] R. C. Pullar, "Hexagonal ferrites: A review of the synthesis, properties and applications of hexaferrite ceramics," *Progress in Materials Science*, vol. 57, pp. 1191–1334, 2012.
- [5] O. Shimizu, Y. Kurihashi, I. Watanabe, and T. Harasawa, "Distribution of thermal stability factor for barium ferrite particles," *IEEE Trans. Magn.*, vol. 49, no. 7, pp. 3767–3770, Jul. 2013.
- [6] O. Shimizu, T. Harasawa, and M. Oyanagi, "Particle orientation effects of barium ferrite particulate media," *IEEE Trans. Magn.*, vol. 46, no. 6, pp. 1607–1609, Jun. 2010.
- [7] S. Furrer *et al.*, "85.9 Gb/in² recording areal density on barium ferrite tape", *IEEE Trans. Magn.*, vol. 51, no. 4, Apr. 2015, Art. ID 3100207.
- [8] S. Maat and A. C. Marley, "Physics and Design of Hard Disk Drive Magnetic Recording Read Heads," *Handbook of Spintronics*, Chapter 25, Springer Science + Business Media, Dordrecht 2016, p. 1012.
- [9] J. B. C. Engelen, S. Furrer, H. E. Rothuizen, and M. A. Lantz, "Where tape and hard-disk technology meet: the HDD head-tape interface", *IEEE Trans. Magn.*, vol. 51, no. 7, July 2015, Art. ID 3300310.
- [10] R. D. Cideciyan, S. Furrer, and M. A. Lantz, "Product codes for data storage on magnetic tape", *IEEE Trans. Magn.*, vol. 53, no. 2, Feb. 2017, Art. ID 3100410.
- [11] J. Justesen and T. Høholdt, "Analysis of iterated hard decision decoding of product codes with Reed-Solomon component codes," in *Proc. Inform. Theory Workshop*, pp. 174–177, Sep. 2007.
- [12] J. D. Coker, E. Eleftheriou, R. L. Galbraith, and W. Hirt, "Noise-predictive maximum likelihood (NPML) detection", *IEEE Trans. Magn.*, vol. 34, no. 1, pp. 110–117, Jan. 1998.
- [13] J. Moon and J. Park, "Pattern-dependent noise prediction in signal-dependent noise", *IEEE J. Sel. Areas Commun.*, vol. 19, no. 4, pp. 730–743, Apr. 2001.
- [14] Z. Wu, P. H. Siegel, J. K. Wolf and H. N. Bertram, "Mean-adjusted pattern-dependent noise prediction for perpendicular recording channels with nonlinear transition shift", *IEEE Trans. Magn.*, vol. 44, no. 11, pp. 3761–3764, Nov. 2008.
- [15] Keith Glover, John Doyle, "State-space formulae for all stabilizing controllers that satisfy an H_∞ -norm bound and relations to risk sensitivity" *System & Control Letters*, 11, pp. 167–172, 1988.
- [16] Sigurd Skogestad, Ian Postlethwaite, *Multivariable Feedback Control: Analysis and Design*, vol. 2, New York, NY, USA: Wiley, 2005.
- [17] G. Cherubini, S. Furrer and J. Jelitto, "High-performance servo channel for nanometer head positioning and longitudinal position symbol detection in tape systems", *IEEE/ASME Trans. Mechatronics*, vol. 21, no. 2, pp. 1116–1128, 2016.
- [18] S. Furrer, A. Pantazi, G. Cherubini, and M. A. Lantz, "Compressional wave disturbance suppression for nanoscale track-following on flexible tape media," 2018 Annual American Control Conference (ACC), pp 6678 – 6683, Jun. 2018.
- [19] A. Pantazi *et al.*, "Nanoscale track-following for tape storage", American Control Conference (ACC), pp. 2837–2843, Jul. 2015.
- [20] G. Cherubini *et al.*, "29.5 Gb/in² recording areal density on barium ferrite tape," *IEEE Trans. Magn.*, vol. 47, no. 1, pp. 137–147, Jan. 2011.