Head for 64 Channel Tape Recording

Robert G. Biskeborn, David J. Seagle and Trevor Olson

Western Digital Corporation, San Jose, CA 95119, USA, robert.biskeborn@wdc.com

Linear magnetic tape systems have taken many forms over the past 73 years. An important example is the transition from open tape reels to ½" tape cartridges beginning in the 80's. In 2000, tape drives built to the new "Linear Tape Open" (LTO) specification launched. Cartridges held 0.1TB, and tape heads had 8 concurrently operable write/read channels, supporting a data rate of 20MB/sec. LTO is now in its 10th generation, with cartridges holding 30TB, and heads having 32 concurrently operable channels, supporting a data rate of 400MB/sec. The 2024 INSIC Tape Technology Roadmap anticipates continued sustained capacity growth, a doubling to 64 active channels and a data rate of more than 800MB/sec, in a small number of years. This talk discusses head design challenges and concludes with an outlook for a future head and drive.

***Index Terms*—Linear tape recording, multi-channel recording heads, tape dimensional stability, magnetostatic crosstalk**

I. Introduction

MODERN tape libraries may house up to several thousand ½" tape cartridges and store over 100PB. Data access is handled by robots that retrieve the cartridges and hand them off to any available drives. After a tape is queued, it is wound onto the drive take-up reel at high speed, e.g., 10 or more m/s, to its starting location. Data can then be written (appended) or read at up to the drive's maximum data rate, host permitting, e.g. 400MB/s for the "Linear Tape Open" (LTO) 10th generation tape drive. This is more than 1.4x higher than the maximum data rate of a modern HDD at its OD, and unlike HDD, which is nearly 2x slower at its ID, does not diminish as the medium fills. Tape drives achieve this data rate despite having significantly lower head-to-medium velocity of ~2-6m/s, vs approximately a 33.5m/s maximum velocity for a 7200rpm HDD, and lower linear density, as will be discussed.

II. Multiple Active Channels

A. Background

That tape has inherently lower single channel data rate stems from the fact that tape media is particulate, unlike HDD, which transitioned to sputtered media 40 years ago. Accordingly, tape magnetic coating is both thicker (~50nm) and rougher than disk coating and has significantly larger magnetic "grains" (particles). Furthermore, tape is contact recording, so the head elements, particularly sensitive magneto-resistive read sensors, must be recessed after lapping, e.g., by ion-milling; and the tape bearing surfaces must be provided with relatively thick dielectric protective coatings. Today, effective magnetic spacing for tape is roughly 35-45nm [1], compared to less than 2-3nm for HDD. Thus, linear density (bits/um) for tape is more than 10x lower than it is for disk. So, today as in the past, tape drives rely on deployment of multiple, concurrently active transducing elements to achieve target data rates.

B. LTO format

The first LTO drives, which launched in 2000 (with a cartridge capacity of 100GB), had 8 concurrently operating write/read channels supporting a data rate of 20MB/s. The array of write/read channels was flanked by timing-based servo (TBS) readers at each end and spanned slightly less than ¼ of the ½" tape width. Five servo tracks divides the tape into 4 equal regions called "data bands," each being 2859um wide. The consortium anticipated that the number of active channels in the drive would need to increase every few generations to maintain drive performance [2]. Further, it envisioned that, if possible, each new generation would support up to twice the capacity of the previous generation and would also support writing tapes from the previous one, and reading tapes from the previous two, generations. Accordingly, the pitch between legacy channels is fixed, and channel count increases twofold when more are needed. The net is that the pitch between legacy and new channels drops precisely 2x. The width of the servo tracks on tape may change but their centerline spacing does not. Finally, ASIC design constraints mandate design choices compatible with doubling.

TABLE I
Tape drive parameters

gen	year	channels	capacity (TB)	data rate (MB/s)	hours to fill	one-way wraps	tape length (m)	linear density (bits/um)	tracks
1	2000	8	0.1	20	1.38	48	609	4.88	384
2	2003	8	0.2	40	1.38	64	609	7.40	512
3	2005	16	0.4	80	1.38	44	680	9.64	704
4	2007	16	0.8	120	1.85	56	820	13.25	896
5	2010	16	1.5	140	3.17	80	846	15.14	1280
6	2012	16	2.5	160	4.33	136	846	15.14	2176
7	2015	32	6	300	5.55	112	960	19.09	3584
8	2017	32	12	360	9.27	208	960	20.67	6656
9	2021	32	18	400	12.50	280	1035	21.46	8960
10	2025	32	30	400	20.83	462	1035	21.65	14784

Table 1 shows that in 2005, the number of active channels increased from 8 to 16 for the 3rd generation drive, and again in 2014 from 16 to 32 for the 7th generation drive. In both cases, elements were added between the legacy elements. Pitch between was 333um for the 1st generation and is currently 83.35um. Increasing the number of channels, results in fewer one-way "wraps" needed to fill the tape, thus minimizing wear and tear of the tape, head and the drive, and addressing the time needed to fill the tape (capacity/data rate). It also helps to enable data rate growth while keeping tape speed below ~6.5m/s. The strategy of increasing channels by doubling also enables preserving the tape 4-data band format. However, it is to be noted that the 32-channel 10th generation drive does not support backward read or write compatibility due to a modification of the head-tape interface.

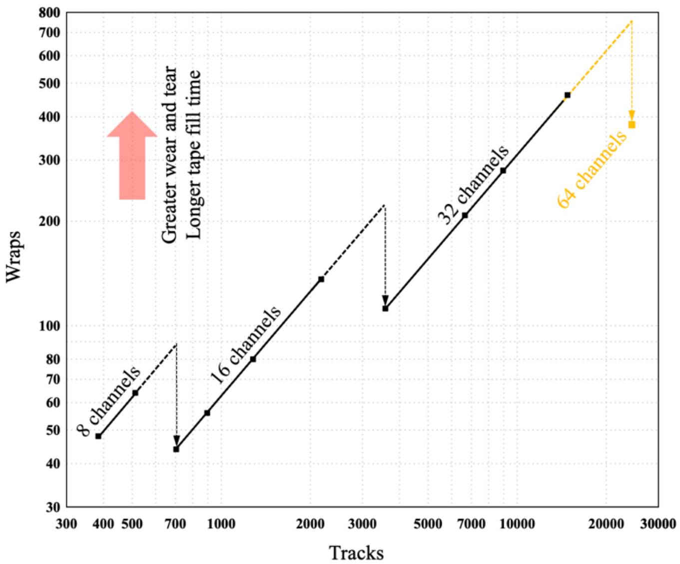

Fig. 1. Number of wraps versus tracks for the 10 tape generations shown in the table. Note that Wraps = Tracks / Channels. The 64 channels data point is anticipated based on the historical trend, which is driven by need to minimize wear and tear, tape fill time and tape velocity.

C. Transition details

First generation LTO heads from IBM deployed 8 write and 8 read elements in an alternating side-by-side arrangement (as opposed to piggyback). In that configuration, each element in the array is on a half pitch of 166.5um. Thus, there was no need to change transducer pitch for the transition to 16 channels. For that, single elements were replaced by piggy-back reader-writer pairs. However, this did bring a substantial increase in wiring complexity. The transition to 32 channels presented new challenges, as it meant adding elements where there were none. And making these piggyback pairs would have resulted in very congested I/O, as well as creating other problems. A decision was made to build the writers and readers on separate wafers. This resulted in 32-channel tape heads having 3 modules (2 writers and a central reader) as described in [3].

III. 64 Channels Strategies

Doubling the number of channels again, from 32 to 64, presents even greater challenges. The question of whether backward compatibility is required is beyond the present scope, and so the authors will explore two basic strategies that they believe potentially support backward compatibility to LTO10.

The first of these is to adopt the 32-channel strategy, namely insert recording elements between the existing elements, thus reducing the element pitch to 83.25 / 2 = 41.625um. This is shown with the current 32-channel head in the top and middle schematic drawings of Fig. 2. For this, the write transducer yoke and coils (presently a 2-layer 14 turn design), would require a major redesign to fit in half the space. A concern is writer-to-writer magnetostatic crosstalk, a phenomenon that was investigated and shown to degrade channel performance [4]. Whether or not the yoke can be designed such that crosstalk is acceptable depends on design details beyond the scope of this digest.

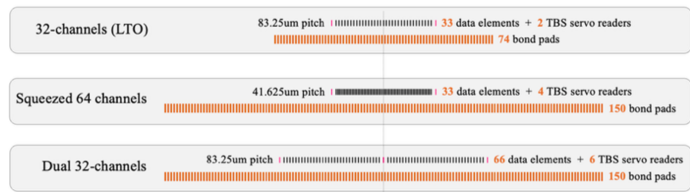

Fig. 2. Two approaches that address the transition to 64 active channels. In the "Squeezed 64 channels," all transducers are in the same span as in the 32-chsnnles head. In the "Dual 32-channels," the 32-channel array is duplicated, and the two span two (of 4) data bands.

A second strategy takes a very different approach. This is shown in the lower drawing of Fig. 2. Here, rather than squeezing the elements to fit into the existing 32-channel head span, the current 32-channel element array is duplicated and re-centered, thus creating a pair of arrays that share the central servo reader(s). This preserves element pitch and transducer dimensions. But it also doubles the span between outermost servo reader(s), potentially affecting alignment between head and recorded tracks due to a phenomenon called Tape Dimensional Stability (TDS). One method [5] for dealing with linear TDS is based on making small changes to a nominally tilted head to align projected head span with the tape. However, linear compensation does not address mis-registration due to non-linear distortion (head and/or tape) which could potentially impact achievable capacity. Presumably, misregistration can be dealt with via mechanical engineering principles, which is preferable to having to face fundamental principles of physics, e.g., magnetostatic crosstalk.

IV. Outlook

The dual 32-channel head records data over a twofold wider (6mm) region than a squeezed 64-channels configuration, thus enabling ECC to better deal with non-functioning read elements. For comparison, in HDD ECC is confined to a 50nm wide region. With non-linear TDS remediation, a 128 channels head comprised of two dual-32 channel arrays may be possible.

References

[1] INSIC International Magnetic Tape Storage Technology Roadmap 2024, White Paper, 2024
[2] lto.org/roadmap
[3] R. G. Biskeborn, "Head and interface for high areal density tape recording," IEEE Transactions on Magnetics, vol. 48, no. 11, Nov. 2012.
[4] R. G. Biskeborn, P. Hergert, and P. O. Jubert, "Crosstalk between write transducers," IEEE Transactions on Magnetics, vol. 44, no. 11, Nov. 2008.
[5] John S. Judge and Robert A. Johnson, "Method of reading recorded information from a magnetic tape that compensates for track pitch change," US Patent 6,141,174, Quantum Corporation, Oct. 2000.