LTO-9 TECHNOLOGY AND USER DATA RELIABILITY ANALYSIS

Suayb S. Arslan [i,iv] Mark Lantz [ii], Simeon Furrer [ii], Geoff Spratt [iii] and Turguy Goker [i]

i. Quantum Corporation, Irvine, CA, USA
ii. IBM Research Europe, Zürich, Rüschkilon, Switzerland
iii. Hewlett Packard Enterprise, Boise, Idaho, USA
iv. Massachusetts Institute of Technology, Cambridge, MA, USA

August 2022

Abstract

In this paper, we analyze the reliability of the LTO-9 error correction codes to support LTO-9 user data reliability being better than 1 uncorrectable error event in 10^{-20} user bits transferred which is typically referred to as uncorrectable bit error rate (UBER) translating to at least 12 NINES of durability. There are multiple technologies enabling LTO-9 to achieve 10X UBER improvements over LTO-8. We'll show that this is also applicable in the event of any one of the 32-channels becoming degraded.

Table of Contents

Section 1 - LTO-9 Technology

Over the last 2 decades since the introduction of LTO in 2000 with a capacity of 100 GB, the storage capacity of LTO cartridges has increased by 180 times and data rates have increased by 20 times. Over the same period, the specified End-Of-Life (EOL) uncorrectable bit error rate (UBER) of LTO cartridges has also improved by a factor of 1000, 3 orders of magnitude improvement. The recently released Linear Tape Open 9 (LTO-9) format provides a native cartridge capacity of 18 TB, and an uncorrectable bit error rate (UBER) of 10^{-20}. Relative to the previous generation LTO-8 format, this corresponds to a 50% increase in capacity and a 10X improvement in UBER. An UBER of 10^{-20} corresponds to one unrecoverable read error event for every 12.5 Exabytes of data read. The simultaneous increase in capacity and improvement in reliability was enabled by a combination of new technologies implemented in LTO-9. Using rough figures, the 50% capacity increase was achieved by increasing the Areal Density (AD) by 41% through the combination of a 35% higher track density (TD) and a 4% higher linear density (LD), plus longer tape that is greater than 1 km in length (1035 m).

Increasing the areal density in magnetic recording results in a loss in SNR that is described by the following well-known equation: [27]

$$SNR \approx \frac{0.31\gamma}{\alpha^2}\frac{B^2W}{D^3}$$

where, D= grain diameter, B is bit length (Linear Density) and W is bit/reader Width (Track Density). $\frac{0.31\gamma}{\alpha^2}$ is micro-magnetics constant. Variable γ is the normalized density ratio for a given linear density and the pulse shape which is also determined by the TMR head design. Variable α is the transition parameter which is a function of magnetic characteristics of the recording layer and the head to tape spacing. The magnetic particle characteristics, tape-head spacing and TMR head design all play key roles for SNR. The quadradic scaling of SNR with B and linear scaling with W makes it less costly in terms of SNR to scale TD compared to LD and is the motivation for the 35% TD scaling compared to 4% LD scaling implemented in LTO-9. In general, a reduction in SNR will result in an increase in error rate and hence a reduction in reliability. To avoid such a decrease in reliability, the SNR loss due to AD scaling needs to be compensated for with improvements in the recording technology. In LTO-9 this was done using new drive technologies that include a new optimized head-to-tape interface, higher sensitivity TMR (tunnel magnetoresistance) readers and an optimized writer.

In addition, LTO-9 uses a new media with improved BaFe particles and is based on a thinner substrate with optimized stability. The 35% track density scaling implemented in LTO-9 was achieved by reducing the width of the read head in combination with improvements in tracking performance. In LTO-9 Tape Dimensional Variations (TDV) of the new substrate that arise due to variations in the environmental conditions are managed via a unique one-time-per tape calibration algorithm that assists complex drive firmware to control 32-channel magnetic recording using closed-loop servo algorithms. The reliability of LTO-9 was further improved using a new highly efficient C2 error correction code based on a longer block length Reed-Solomon error correction code. Together, these technologies enabled the simultaneous 50% increase in capacity and 10x improvement in uncorrectable bit error rate (UBER) of LTO-9 to 10^{-20}.

Section 2 - LTO-8/9 Format Background

To better understand how this reliability improvement was achieved it is helpful to first review how data is laid out and protected by error correction codes in the LTO-8 and 9 formats. In all LTO generations to date, user data is partitioned into datasets and each data set is further partitioned into multiple sub-data sets. Each sub-data set consists of multiple interleaved rectangular data blocks and are protected by a product code that consists of two Reed-Solomon component codes referred to as C1 and C2 codes. Each rectangular data block can be viewed as a matrix and goes through a C1-C2 product code encoding step before being written to tape. To visualize this encoding process, envisage a rectangle of user data bytes of size k2 × k1. First, the rows of the matrix are encoded using a (n1, k1) Reed-Solomon code (C1 code). Then columns of the new row-encoded matrix (now of size k2 × n1) are encoded with (n2, k2) RS code (C2 code) to generate the final (n2 × n1) coded data block as shown in Figure 1 below.

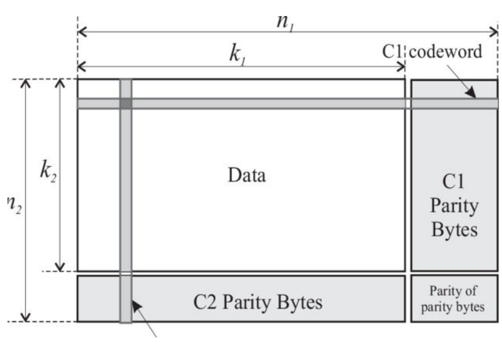

Figure 1: The final (n2× n1) product coded data block

The format of LTO dictates multiple rounds of interleaving and spreading of bytes that belong to the same C1/C2 codeword over different physically separated locations on the tape surface. The separation between the bytes of the same codeword is maximized to guard against correlated errors such as media defects. This deep interleaving results in errors that are essentially independent which helps to improve reliability and is central to the mathematical analysis of the uncorrectable bit error rate.

The parameters that describe the C1 and C2 codes implemented in LTO-8 and LTO-9 are summarized in Table 3 below. Table 3 uses the standard notation of (n, k, d), where n represents the block length, k is the message length and d is the distance. For the Reed Solomon codes used here, d=n-k+1. Both the C1 and C2 codes use 8 bit symbols. Although the format is designed to decorrelate error events to maximize the performance of product decoding, it sometimes fails to do so due to accumulated operation-related errors, environment, or internal mechanics. In this short note though, we'll consider only the completely independent byte error scenario i.e., a discrete memoryless channel (DMC) model with each byte having a probability to be in error of P_{byte}. In LTO-8 and LTO-9, two different decoding approaches can be used in combination with the product codes: 1) All error mode and 2) Erasure mode.

1. **All error:** In this mode, both C1 and C2 use their redundancy to correct byte errors only without any communication channel in between.

2. **Erasure(a=1):** In this mode the C1 code is operated in "error correction" mode and if it fails, it provides that failure information for the C2 decoding engine to mark the corresponding byte as an erasure. The C2 decoding engine is operated in "erasure correction" mode with 2 bytes of the parity reserved to catch and correct one byte of error in the (unlikely) case of a C1 mis-correction. Byte errors due to C1 miscorrections are very rare and occur if the C1 decoder decodes the wrong codeword and does not indicate any failure flag for C2).

The reliability achieved using the erasure mode is analyzed in Section 4 below for the C1 and C2 code parameters summarized in Table 1.

	LTO-7	LTO-8	LTO-9
C1	(249,237,13)	(249,237,13)	(243,231,13)
C2	(96,84,13)	(96,84,13)	(192,168,25)
Efficiency	83.28%	83.28%	83.18%

Table 1: C1 and C2 code parameters for different generations of LTO

Section 3 - What is UBER?

The INSIC 2019 technology roadmap white paper defines UBER (uncorrectable bit error rate) as the number of bits in error divided by the total number of bits transferred, i.e. the real bit error rate after error correction decoding [24]. This is not the same as the definition currently used by the HDD industry which defines UBER as the probability of an unrecoverable sector error in a number of user bits transferred [28].

The definition and calculation of UBER can be complex, especially as it is defined as a bit error rate even though the actual error type is the probability of having a ECC decoder failure which is based on sector ECC for HDD and C2 codeword ECC for tape. Assuming the errors, including bit-level errors, are all random, then we can convert C2 codeword or sector error rate to user bit error rate based on format information.

For instance, an HDD sector has 4KB of user bytes (32,768 bits); an LTO-9 C2 codeword has 168 user bytes (1,344 user bits). With the assumption of random errors, we can convert decoder failures to user bits transferred to represent UBER as a hard read error event per number of user bits transferred such that we can compare LTO to HDD using HDD UBER definitions.

As areal densities increase, the probability of UBER errors is also expected to grow, unless the loss in SNR described in section 1 is compensated for by improvements in recording technology. For tape there appears to be considerable potential to continue scaling [24] whereas for HDD challenges arising from the superparamagnetic limit has resulted in a dramatic slow-down in scaling. UBER is a function of magnetic recording technology, media and head usage/wear, environmental conditions, data format, ECC algorithm, and error characteristics for a given data storage device. Uncorrectable errors are typically latent, meaning unless one reads the data, they remain undetected. For example, with HDD, an UBER event may occur at the sector level where a single latent uncorrectable error in a sector, such as a magnetic defect that cannot be corrected by the HDD's internal ECC format results in an un-useable/un-readable sector. Even with RAID technology, such latent sector errors can lead to data loss when encountered during RAID reconstruction after a disk failure [23]. With LTO tape, latent errors may happen at a data set level. However due to LTO's unique ECC format which is based on orthogonal interleaved 2-dimensional 32 channel Reed-Solomon error correction codes, the probability of an UBER event is orders of magnitude lower than HDD.

With magnetic recording, an uncorrectable error causes individual sectors in HDD or data sets in LTO to become unavailable. Even though the initial media defects are typically mapped out, HDD may still experience latent errors during writing, for example write errors (such as a high-fly write), or by media imperfections (short or long defects), or smeared soft particles. Unless data is fully verified, these errors may go undetected [25]. LTO has built-in resilience to write mode latent errors due to its read-while-write architecture in which such write-related errors are detected and rewritten during writing. In addition to writer errors, there are other kinds of latent errors, e.g. post-write process errors where media magnetics may degrade, or new defects may be generated resulting in latent uncorrectable errors. These are mostly related to capacity scaling and ageing of the magnetic storage devices for both HDD and tape [26]. HDDs requires real time periodic data scrubbing [21] to manage these errors whereas tape can rely on its multi-channel 2-dimensional orthogonal C1-C2 product ECC to provide orders or magnitude better UBER reliability.

Section 4 - LTO-9 UBER Reliability Analysis

Uncorrectable error events in LTO tape are extremely rare which makes it very challenging to measure the error rate experimentally. To address the challenge of quantifying the performance of LTO's error correction codes, both industry and academia use theoretical model-based calculations combined with the assumption that errors are random and uncorrelated. These theoretical models are based on the binomial distribution of raw byte errors, an assumption that has been experimentally verified at controlled high error rate conditions where the assumption continues to hold thanks to the deep-interleaving of the format [1][2]. More complex reliability models which are based on the theory of renewal processes can account for correlated errors and defective header and synchronization fields [3] as well as the removable nature of tape which enables multiple tape - drive combinations [4].

To analyze the power of the ECC codes implemented in LTO-8 and LTO-9, we start by defining the probability that a symbol (typically a byte) error occurs as P_{byte} (and let P_{bit} be the associated bit error probability). Then, for an error correction code (codeword) that is N bytes long which can correct t failures. The probability of decoder failure can be calculated as:

$$DECfail = \sum_{i=t+1}^{N} \binom{N}{i} P_{byte}^{i} \, (1 - P_{byte})^{N-i} \, ,$$

the HDD UBER definition [28] is based on this equation.

$$UBER_{HDD} = \frac{DECfail}{\frac{Bits}{Sector}}$$

$$UBER_{Tape} = \frac{DECfail}{\frac{Bits}{CodeWord}}$$

On the other hand, every decoder failure does not necessarily mean all of the bytes of the codeword are in error. To be able to estimate the uncorrectable byte error rate, we use:

$$UBytERR = \frac{1}{N} \sum_{i=t+1}^{N} i \binom{N}{i} P_{byte}^{i} \, (1 - P_{byte}^{i})^{N-i} \, ,$$

the INSIC UBER definition [24] is based on this equation.

As defined in section 3, UBER is the uncorrectable bit error rate, but it is based on converting the probability of decoder failure to user bits using the format specifications. In fact, assuming bit errors to be independent within a given byte error, we can assume $8 \times UBER \approx UBytERR$ due to the following observation:

$$\frac{Avg.\#of\ bit\ errors}{\#of\ bits\ in\ a\ byte} = \frac{1}{8} \sum_{i=1}^{8} i \binom{8}{i} P_{bit}^{i} \, (1 - P_{bit})^{8-i}$$
$$= \frac{1}{8}(8P_{bit}) \approx \frac{1}{8}(1 - (1 - P_{bit})^{8}) = \frac{1}{8} P_{byte}$$

Hence,

$$UBER = UBytERR/8$$

Figure 2 shows the performance of the LTO-8 and LTO-9 C2 codes with the C2 ECC engine operating in the erasure mode described in section 2. The left plot is for all parities used to correct the erasures, which gives the best error correction performance, however miscorrections may happen and go undetected in this mode, albeit with very low probability. The right plot shows the typical LTO use case where C2 uses all parities except two for error correction and the remaining two are used to detect and correct errors that might result from mis-corrections in the previous C1 decoding step. In this mode, all mis-corrections are for practical purposes eliminated. This unique configuration is the fundamental reason why LTO not only provides extremely low UBER numbers but also nearly eliminates any miscorrection errors from being generated and sent to the user as good data. Figures 3 and 4 show the results of a similar analysis as Figure 2 but with one and two out of 32 read data channels assumed to be unuseable due to typical real-world problems such as debris or clogged heads that may occur if the drive was not cleaned effectively.

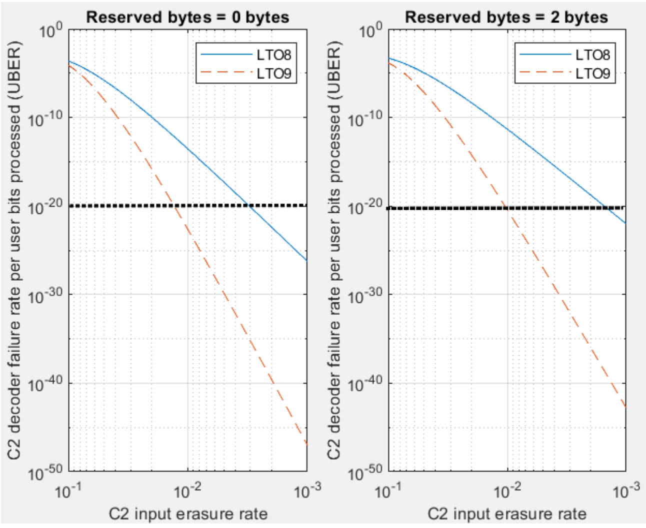

Figure 2: UBER Plot for LTO-8 and 9 with all 32 functional readers

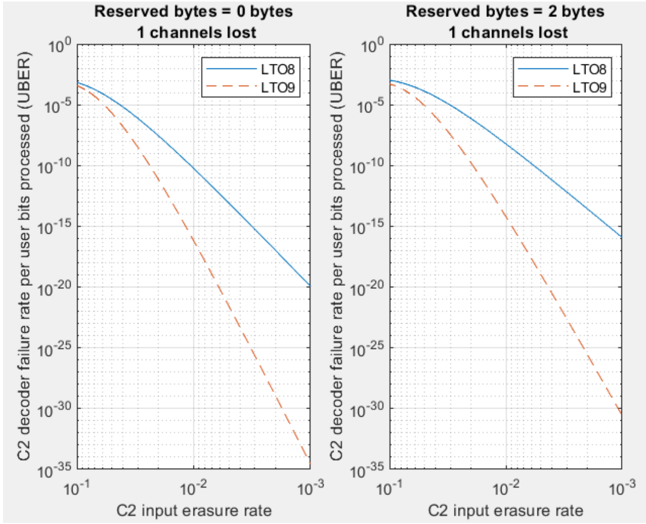

Figure 3: UBER plot for LTO-8 and 9 with one reader impacted by debris during read

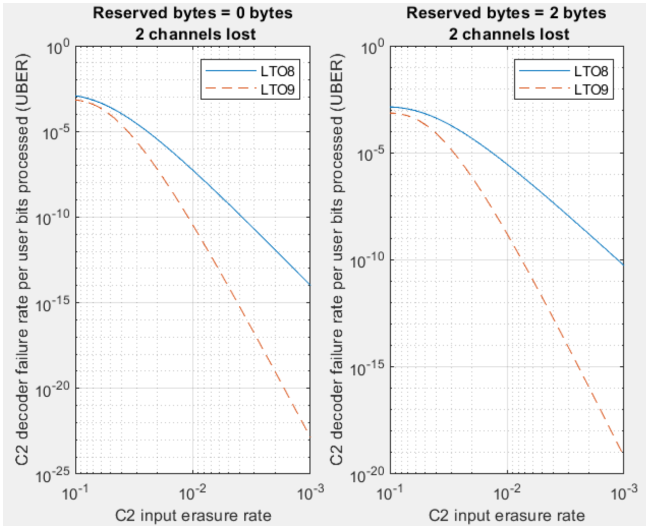

Figure 4: UBER plot for LTO-8 and 9 with two readers impacted by debris during read

Using these results to estimate UBER for a given situation can be a difficult task that is further complicated because tape is a removable medium and where drive firmware uses complex retries when errors are detected. Retries may involve re-reading the same track with different settings as well as using different algorithms and different tunings, including rebuilding data with partial reads. Moreover, the host system can load the media to another (better performing) drive, or clean the head with a cleaner cartridge, all of which can help to reduce input error rates enabling reliable reads free of uncorrectable errors.

In addition to the new more powerful C2 code implemented in LTO-9, another new technology introduced in LTO-9 to reduce the UBER is iterative ECC decoding which can perform additional C1/C2 decoding steps to further improve the error correction performance. Therefore, the actual UBER numbers may be well below the industry published values since these numbers are based on single decoding passes.

In Figures 2 - 4, the x-axis shows the input error rates (C1 decoder output expressed in terms of erasures) to the C2 ECC engine. The input error rate is determined by the drive's tape-to-head interface which may degrade with use, electronics, read channel signal processing, format, and firmware. This is the de-interleaved output of the C1 ECC engines of all 32 channels. The y-axis shows the output error rates which we call the UBER values. Because of the powerful C2 code used in LTO-9 combined with 2-dimensional orthogonal interleaving, the UBER can be extremely low for random media errors or even off-track conditions. As we can see from the figures, for an input error rate of 10^{-3}, the theoretical UBER at the output of C2 decoding can be as low as 10^{-35} .

With LTO-8 and earlier generations, the typical EOL C2 input error rate was about 10^{-3}, a norm for many years [23]. For LTO-8 which uses a (96,84,13) C2 code, this resulted in an UBER of 10^{-19}. LTO-9 uses a new C2 code with a longer codeword (192,168,25) that for the case of random errors assumed here has substantially better error correction power than the previous LTO-8 code, as is clearly indicated by the plots of Figures 1 and 2. The increased length of the C2 code enables this performance improvement with the same efficiency of 87.5% as the LTO-8 code. i.e. the LTO-9 codes have the same overhead cost but provide substantially more powerful user data protection.

Figure 1 shows that when the C2 decoder uses 2 bytes for protection against mis-corrections, the LTO-8 drive will require a C2 input error rate less than 1e-3 to support 10^{-19} UBER. However, LTO-9 with new ECC technology can provide 10^{-20} UBER with C2 input error rates of up to $9x10^{-3}$. This means the LTO-9 drive can provide 10X better user UBER compared to LTO-8 but over a wider range of C2 input error rate conditions, i.e. up to 9X higher C2 input error rates. Therefore LTO-9 supports improved data reliability over a wider range of operating conditions compared to LTO-8. This wider margin helps to enable LTO-9's higher AD by ensuring reliable operation with less SNR. Hence some of the SNR loss discussed in Section 1 that results from the 35% higher TPI and 4% higher LD is compensated by the improved C2 code, while the remainder has been compensated by other improvements in the recording technology discussed in Section 1. Another key feature of LTO-9 is its robustness against correlated errors such as the temporary loss of reading channels due to debris. As can be seen in the right part of Figure 3, LTO-9 achieves an UBER of 10^{-20} with an input error rate of $4x10^{-3}$ even with one dead/unavailable track, whereas even at 4x lower input error rate of 10^{-3}, LTO-8 UBER falls to 10^{-16} with one dead track. The case of two dead tracks is analyzed in Figure 4 where the benefits of the LTO-9 are even more pronounced. For example, in the right panel of Figure 4 we see that at an input error rate of 10^{-3}, the UBER of LTO-8 falls to 10^{-10} whereas LTO-9 achieves an UBER of 10^{-19}. Clearly, LTO-9 offers much lower UBER during real-world error cases such as debris-clogged readers.

Section 5 - LTO-9 NINES of Reliability

UBER represents the probability of an error event and can be thought of as the inverse of the average number of user bits that can be read before an error event is encountered. UBER numbers by themselves are related to average statistics and not particularly related to a time period. In other words, they are only a ratio of numbers of bits. A specific UBER, when related to time, can yield a mean time between failure (MTBF) for an HDD or tape drive and a cartridge, e.g. by calculating the time to read the average number of bits read before an error event assuming the maximum data rate. Table 2 below shows a comparison between the UBER performance of a typical 18 TB HDD with 12 TB LTO-8 and 18 TB LTO-9. In this comparison, the time factor is not included. Instead, we compare the orders of magnitude difference in UBER between HDD and tape. On the right side of the table, we have converted the UBER numbers into the average number of Bytes, TB, PB and EB that can be read before an error event is likely to occur.

Technology	Capacity (Bytes)	UBER*	Orders relative to HDD	Bytes read to error	TB read to error	PB read to error	EB read to error
HDD	18E+12	1E-15		1.25E+14	125	0.125	0.000125
LTO-8	12E+12	1E-19	4	1.25E+18	1,250,000	1250	1.25
LTO-9	18E+12	1E-20	5	1.25E+19	12,500,000	12,500	12.5

Table 2: LTO vs HDD UBER Comparison (*Note that the INSIC Tape UBER definition differs from the HDD definition as discussed in section 4)

One of the important metrics related to reliability is durability. This metric is defined to be the duration of time the system is able to provide access to the user data. If the data is protected by a redundancy mechanism i.e. replication or erasure correction coding, durability refers to the permanent loss of the encoded/replicated user data, i.e. when the amount of data lost exceeds the power of the redundancy mechanism to repair the loss. The unit of durability is usually expressed in terms of days or years. However, in storage industry terminology, durability is expressed in terms of NINES (9's). This refers to the probability of seeing no uncorrectable errors during the operation of the system over a given time period which is typically taken as one year. For example, a system with a 1% probability of failure within a year, has a 99% probability of not failing, or a reliability of 99% = 0.99, which is a reliability of 2 NINES. Similarly, a 0.1% failure probability corresponds to a reliability of 99.9% or 0.999 = 3 NINES of reliability.

The failure rate of a single HDD can be characterized by a metric called the mean-time-to-failure (MTTF). However, in a typical storage scenario, HDDs (as well as tapes) are used in groups, with user data plus some redundancy spread across a set of physically separated storage mediums. In this way, the decorrelation of failures/errors is achieved. In this scenario, a new reliability metric is needed because the added redundancy improves the resilience of the stored data. A known metric is mean-time-to-data-loss (MTTDL). One can compute MTTDL in terms of MTTF of the constituent storage components of the system using a Markov process to model the arrival (occurrence) of failures to the system.

The reliability function R(t) is not, in general, exponentially distributed. However, for simplicity let us assume it is distributed exponentially with the rate 1/MTTDL. The number of 9's is then given by:

$$NINES(durability) = log10\left[\frac{1}{1-R(t)}\right] \; where\ R(t) = e^{-t}/_{MTTDL}$$

For example, for a system with MTTDL of 2,500,000 hours, and an operating time of interest of 1 year (8760 hours), we have the durability number computed as follows, R(8760) = e−8760/2500000 = 0.9965→two 9's which means the system will operate without a failure with probability 0.9965 for the first year of use at a 100% duty cycle [22].

As mentioned above, UBER does not have a time element and hence converting UBER to NINES of reliability requires a modified approach without the time element. To do this we take as inspiration another commonly used storage reliability metric: expected annual fraction of data loss which measures how much data is lost when a data loss event occurs and hence is complementary to the mean time to data loss metric.

To estimate a similar metric for HDD and Tape using UBER, we consider how much data is lost when an UBER event occurs. For HDD, an uncorrectable bit error results in the sector containing that bit becoming un-decodable and hence lost. An HDD sector contains 4096 bytes of user data. For an UBER of 1 error in 10^{-15} bits read, this corresponds to 10^{-15} / (4096 x 8) = 3.05E+10 sectors read on average before a sector is lost, or 10 NINES(UBER). In LTO tape, the UBER is much lower than HDD, however, when such an event occurs the data set containing that error is lost. In LTO-8 a data set contains about 5 Mbytes of user data and in LTO-9 it contains about 9.8 Mbytes of user data. Table 3 below summarizes these parameters and NINES(UBER) for HDD and LTO-8/9. Tables 2 and 3 demonstrate that the probability of encountering an uncorrectable error in tape is much lower than in HDD, however, when such an event does occur, a larger amount of user data is lost. The net result is two orders of magnitude higher NINES(UBER) for LTO-9 compared to HDD.

	HDD	LTO-8	LTO-9
UBER	1 error in 10^{15} bits	1 error in 10^{19} bits	1 error in 10^{20} bits
# Sector /C2 codewords	3.05E+10	2.48E+11	1.28E+12
HDD Sector/LTO Data Set Bytes	4096	5.031E6	9.8E6

Table 3: UBER to NINES Conversion

It is also important to recall that LTO is a removable archival storage medium where host software can manage a pool of drives to read any given tape, especially for mission-critical tasks such as recovery from ransomware where error-free reads are extremely crucial. Note that the NINES(UBER) numbers in Table 3 are based on EOL (End Of Life) values with worn read heads. However, assuming the host is able to manage the drive pool and their conditions, the UBER numbers may be even better than the EOL UBER of 10^{-20} assumed in Table 3. For example, in Figure 2 it can be seen that assuming a factor of two better input error rate for a newer drive, i.e. 4.5E-3 instead of 9E-3, results in an LTO-9 UBER of better than 10^{-25}. This translated to a NINES(UBER) of 17 for LTO-9 which basically means that it is extremely unlikely that users would encounter a hard read error event with LTO-9.

References

[1] S. Furrer, et al., "85.9 Gb/in2 Recording Areal Density on Barium Ferrite Tape," IEEE Trans. on Magnetics, vol. 51, no. 4, April 2015.

[2] M. A. Lantz, et al., "123 Gb/in2 Recording Areal Density on Barium Ferrite Tape", IEEE Trans. on Magnetics, vol. 51, no. 11, Nov. 2015.

[3] 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, Sept. 2014.

[4] S. S. Arslan, J. Peng, and T. Goker, "A data-assisted reliability model for carrier-assisted cold data storage systems," Reliability Engineering & System Safety, vol. 196, p. 106708, 2020.

[5] P. Elias, "Error-Free Coding," IRE Trans. Inform. Theory, vol. 4, pp. 29–37, 1954.

[6] K. A. S. Immink, "Runlength-limited sequences," Proc. IEEE, vol. 78, pp. 1745–1759, Nov. 1990.

[7] J. Moon and B. Brickner, "Maximum transition run codes for data storage systems," IEEE Trans. on Magnetics, vol. 32, pp. 3992–3994, Sep. 1996.

[8] 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.

[9] 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, April 2001.

[10] Expanded Sub Data Set for Magnetic Tape Recording, IPCOM000223249D, Nov. 2012.

[11] 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.

[12] M. Blaum, et al., "High-rate modulation codes for reverse concatenation," IEEE Trans. on Magnetics, vol. 43, no. 2, pp. 740-743, Feb. 2007.

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

[14] T. Mittelholzer, "Enumerative maximum transition run codes," Proc. IEEE Int'l Symp. on Information Theory, ISIT 2009, Seoul, Korea, pp. 1549-1553, June 2009.

[15] 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.

[16] 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), December 2014.

[17] J. Oh, J. Ha, H. Park, and J. Moon, "RS-LDPC Concatenated Coding for the Modern Tape Storage Channel," accepted for publication in IEEE Trans. on Commun.

[18] 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

[19] 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, September 2013, 4998

[20] 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.

[21] Arslan, S. S., Goker, T., Lee, J., & Le, H. (2019). System and method for tape layout optimization, U.S. Patent No. 10,319,406. Washington, DC: U.S. Patent and Trademark Office.

[22] Arslan,S.S. (2014) "Durability and Availability of Erasure-Coded Systems with Concurrent Maintenance". Available online: http://www.suaybarslan.com/Reliability_Systems_14.pdf.

[23] Bianca Schroeder (2010) "Understanding latent sector errors and how to protect against them" https://www.usenix.org/legacy/event/fast10/tech/full_papers/schroeder.pdf

[24] INSIC "Applications, Systems and Technology Roadmap" https://insic.org/insic-application-systems-and-technology-roadmap/

[25] Lakshmi N. Bairavasundaram, "An Analysis of Latent Sector Errors in Disk Drives", https://dl.acm.org/doi/10.1145/1269899.1254917

[26] Saurabh Kadekodi "Cluster storage systems gotta have HeART: improving storage efficiency by exploiting disk-reliability heterogene", FAST '19

[27] Signal to noise ratio scaling and density limit estimates in longitudinal magnetic recording, Berthram https://ieeexplore.ieee.org/document/706722

[28] https://www.seagate.com/files/www-content/product-content/enterprise-hdd-fam/exos-x18-channel/_shared/en-us/docs/100865853c.pdf

About The Authors

Suayb S. Arslan (Senior member, IEEE) received the M.Sc. and Ph.D. degrees in electrical engineering from the University of California, San Diego, CA, USA, in 2009 and 2012, respectively. He was with Mitsubishi Electric Research Laboratory, Boston, MA, USA, in 2009, where he was involved in research and development of image and video processing algorithms for biomedical applications. In 2011, he joined Quantum Corporation, Irvine, CA, USA, where he conducted research on advanced detection and coding algorithms for increased capacity tape storage and cloud systems. He is currently visiting Massachusetts Institute of Technology, Boston, MA, USA as a faculty member. He serves as vice-chair for IEEE ComSoc Turkey and Treasurer for IEEE Data Storage Technical committee, IEEE ComSoc Global. His research interests include information theory, neuroscience, digital communication and storage, cloud and Quantum computing, reliability/system theory and image/video processing.

Mark Lanz (Fellow IEEE) received B.Sc. (91) and M.Sc. (93) degrees in electrical engineering from the University of Alberta, Canada, and a Ph.D. from the University of Cambridge, U.K., in 1997 for work in the field of scanning probe microscopy and tribology. Dr. Lantz is currently a distinguished research staff member in the IBM Research – Zurich laboratory, an IBM Master Inventor and the manager of the Cloud FPGA and Tape Technology Research Group. His current research interests include storage technology, FPGA based accelerators, nanotechnology, and storage system reliability and performance modeling.

Simeon Furrer received the Dipl.Ing. and Ph.D. degrees in Electrical Engineering from ETH Zurich, Switzerland, in 1999 and 2005, respectively. He was a Principal Scientist with Broadcom Corporation, Sunnyvale, CA, USA, from 2005 to 2010, where he was involved in the design and architecture of IEEE802.11n-compliant WLAN chipsets. In 2010, he joined IBM Research Zurich, Switzerland and is currently a member of the Cloud and AI Systems Research department. His research interests include cloud data storage, magnetic tape recording, nanopositioning, signal processing and error-correction coding. He is an IBM Master Inventor and holds over 100 patents.

Geoff Spratt received his BSc in Applied Physics from Coventry University in 1983 and his Ph.D in Fine Particle Magnetism in 1987 from the University of Central Lancashire. He has spent over 30 years in the magnetic recording field in both academia and industry. He was a Visiting Researcher at Hitachi's Central Research Lab in Tokyo from 1989 to 1990 and a Research Associate at Carnegie Mellon University from 1990 to 1997. For the past 25 years he has been with Hewlett Packard Enterprise, predominantly working on LTO from the media perspective.

Turguy Goker is Director of Advanced Development LTO for Quantum Corp., has been in the Tape Drive industry for over 30 years as a magnetic recording technologist specializing in servo and signal processing technologies, designing, and developing LTO drives, media, and tape formats, with over 50 patents. He is a technical team leader representing Quantum in the LTO organization. Turguy is also the current Chairman of the Board for the Information Storage Industry Consortium (INSIC) focusing on research for the worldwide information storage industry, mainly tape technology roadmaps. During his career as an LTO researcher, he was selected for the Technology Hall of Fame with Seagate, Certance, and Quantum. He is a graduate of California State University San Diego with BSEE and MSEE degrees.

About The LTO Consortium

Linear Tape Open (LTO) Ultrium is a high capacity tape storage solution developed and continually enhanced by Hewlett Packard Enterprise, IBM and Quantum and promoted by **the LTO Program.** Linear Tape-Open LTO, the LTO logo, Ultrium and the Ultrium logo are registered trademarks of Hewlett Packard Enterprise, IBM and Quantum in the US and other countries.