Reliability of Erasure-Coded Storage Systems with Latent Errors

Ilias Iliadis
IBM Research Europe – Zurich
8803 Rüschlikon, Switzerland
email: ili@zurich.ibm.com

***Abstract*—Large-scale storage systems employ erasure-coding redundancy schemes to protect against device failures. The adverse effect of latent sector errors on the Mean Time to Data Loss** (MTTDL) **and the Expected Annual Fraction of Data Loss** (EAFDL) **reliability metrics is evaluated. A theoretical model capturing the effect of latent errors and device failures is developed, and closed-form expressions for the metrics of interest are derived. The** MTTDL **and** EAFDL **of erasure-coded systems are obtained analytically for (i) the entire range of bit error rates, (ii) the symmetric, clustered, and declustered data placement schemes, and (iii) arbitrary device failure and rebuild time distributions under network rebuild bandwidth constraints. For realistic values of sector error rates, the results obtained demonstrate that** MTTDL **degrades whereas, for moderate erasure codes,** EAFDL **remains practically unaffected. It is demonstrated that, in the range of typical sector error rates and for very powerful erasure codes,** EAFDL **degrades as well. It is also shown that the declustered data placement scheme offers superior reliability.**

Keywords–Storage; Unrecoverable or latent sector errors; Reliability analysis; MTTDL; EAFDL; RAID; MDS codes; stochastic modeling.

I. INTRODUCTION

Today's large-scale data storage systems and most cloud offerings recover data lost due to device and component failures by deploying efficient erasure coding schemes that provide high data reliability [1]. The replication schemes and the Redundant Arrays of Inexpensive Disks (RAID) schemes, such as RAID-5 and RAID-6, which have been deployed extensively in the past thirty years [2-5] are special cases of erasure codes. Modern storage systems though use advanced, more powerful erasure coding schemes. The effectiveness of these schemes has been evaluated based on the Mean Time to Data Loss (MTTDL) [2-11] and, more recently, the Expected Annual Fraction of Data Loss (EAFDL) reliability metrics [12-16]. The latter metric was introduced, because Amazon S3 [17], Facebook [18], LinkedIn [19] and Yahoo! [20] consider the amount of lost data measured in time.

The reliability level achieved depends not only on the particular choice of the erasure coding scheme, but also on the way data is placed on storage devices. The reliability assessment presented in [4] demonstrated that, for a replication factor of three, a declustered data placement scheme achieves a superior reliability than other placement schemes. The declustered placement scheme ensures that codewords are spread equally across devices. This is the scheme that was originally used by Google [21], Facebook [22], and Microsoft® Azure[1] [23], but, to improve data reliability and storage efficiency further, today they use erasure coding schemes that offer higher efficiency [24-26].

The reliability of storage systems is further degraded by the occurrence of unrecoverable sector errors, that is, errors that can be corrected neither by the standard sector-associated error correction code (ECC) nor by the re-read mechanism of hard-disk drives (HDDs). These sector errors are latent, because their existence is only discovered when there is an attempt to access them. Once an unrecoverable or latent sector error is detected, it can usually be corrected by the erasure coding capability. However, if this is not feasible, it is permanently lost, leading to an unrecoverable failure. Consequently, unrecoverable errors do not necessarily lead to unrecoverable failures. Permanent losses of data due to latent errors are quite pronounced in higher-capacity HDDs and storage nodes, because of the higher frequency of their occurrence [27-30]. The risk of permanent loss of data rises in the presence of latent errors.

Previous works have shown that actual latent-error rates degrade MTTDL by orders of magnitude [8][11][28][30]. Does this also apply to the case of the EAFDL metric given that, when a data loss occurs, the amount of sectors lost due to latent errors is much smaller than the amount of data lost due to a device failure? What is the range of error rates that cause EAFDL to deteriorate? This article addresses these critical questions.

Analytical results for the MTTDL and EAFDL metrics in the context of general erasure-coded storage systems, but in the absence of latent errors, were obtained in [12-14]. The first analytical assessment of EAFDL in the presence of latent errors was presented in [15] for the case of RAID-5 systems by presenting a comprehensive theoretical stochastic model that captures all the details of the rebuild process. This model was subsequently extended to a significantly more complex one for the case of RAID-6 systems [16]. Clearly, extending this model further to assess EAFDL in the presence of latent errors for arbitrary erasure coding schemes seems to be a daunting task because of its state explosion. The state of the model developed in this article does not explode, because it takes into account only the most significant details of the rebuild process.

To assess the reliability of erasure-coded systems, we adopt the non-Markovian methodology developed in prior work [12-

[1]Microsoft is a trademark of Microsoft Corporation in the United States, other countries, or both.

14] to evaluate MTTDL and EAFDL of storage systems and extending it to assess the effect of latent errors. The validity of this methodology for accurately assessing the reliability of storage systems has been confirmed by simulations in several contexts [4][9][12][31]. It has been demonstrated that theoretical predictions of the reliability of systems comprising highly reliable storage devices are in good agreement with simulation results. Consequently, the emphasis of the present work is on theoretically evaluating the reliability of storage systems with latent errors.

The reliability results obtained by the model developed here are shown to be in agreement with previous specific theoretical and simulation results presented in the literature. We verify its validity by comparing the results obtained for the cases of RAID-5 and RAID-6 systems and showing that they match with those derived by the detailed models in [16]. Furthermore, we demonstrate that the model developed yields theoretical reliability results that match well with the simulation results obtained in [32], which studies the effect of erasure codes deployed in a realistic distributed storage configuration. This establishes a confidence for the model presented, the results obtained, and the conclusions drawn. The model developed is a practical one that takes into account the characteristics of latent errors observed in real systems. It is realistic, because it considers general device failure distributions including real-world ones, such as Weibull and gamma. It can also be used to assess system reliability when scrubbing is employed by applying the methodology described in [8]. This is the first work to study the effect of latent errors on EAFDL for general erasure-coded storage systems.

Note that the storage model considered in this work is relevant and realistic, because it properly captures the characteristics of erasure coding and of the rebuild process associated with the declustered data placement scheme currently used by Google [24], Microsoft® Azure [26], Facebook [33], and DELL/EMC [34]. Consequently, the theoretical results derived here are important, because they can be used to assess the reliability of the above schemes and also determine the parameter values that ensure a desired level of reliability. It can also be used to assess system reliability when scrubbing is employed by applying the methodology described in [8].

The key contributions of this article are the following. We consider the reliability of erasure-coded storage systems with latent errors that was derived analytically for the MTTDL and EAFDL reliability metrics for the entire range of sector error rates, and for the symmetric, clustered, and declustered data placement schemes [1]. In this article, we extend our previous work by also presenting results for the additional reliability metric of interest $E(H)$, namely, the conditional expected amount of lost user data, given that data loss has occurred. We subsequently demonstrate that, in the range of typical sector-error rates, unrecoverable failures are frequent, which degrades MTTDL. However, in [1], it was shown that the relative increase of the amount of data loss is negligible, which leaves EAFDL practically unaffected in this range. In the present work, we demonstrate that this result holds for moderate erasure codes, but for very powerful erasure codes, it may not be the case. We have confirmed that reliability results obtained by the model developed here are in agreement with previous specific theoretical and simulation results presented in the literature. We also assess the reliability of real-world erasure coding schemes employed by enterprises. The model developed provides useful insights into the benefits of the erasure coding schemes and yields results for the entire parameter space, which allows a better understanding of the design tradeoffs.

TABLE I. NOTATION OF SYSTEM PARAMETERS

Parameter	Definition
n	number of storage devices
c	amount of data stored on each device
l	number of user-data symbols per codeword ($l \geq 1$)
m	total number of symbols per codeword ($m > l$)
(m, l)	MDS-code structure
s	symbol size
k	spread factor of the data placement scheme, or group size (number of devices in a group) ($m \leq k \leq n$)
b	average reserved rebuild bandwidth per device
$B_{\max}$	upper limitation of the average network rebuild bandwidth
X	time required to read (or write) an amount c of data at an average rate b from (or to) a device
$F_X(.)$	cumulative distribution function of X
$F_\lambda(.)$	cumulative distribution function of device lifetimes
P_{bit}	probability of an unrecoverable bit error
s_{eff}	storage efficiency of redundancy scheme ($s_{\text{eff}} = l/m$)
U	amount of user data stored in the system ($U = s_{\text{eff}}\, n\, c$)
$\tilde{r}$	MDS-code distance: minimum number of codeword symbols lost that lead to permanent data loss ($\tilde{r} = m - l + 1$ and $2 \leq \tilde{r} \leq m$)
$f_X(.)$	probability density function of X ($f_X(.) = F'_X(.)$)
C	number of symbols stored in a device ($C = c/s$)
μ^{-1}	mean time to read (or write) an amount c of data at an average rate b from (or to) a device ($\mu^{-1} = E(X) = c/b$)
λ^{-1}	mean time to failure of a storage device ($\lambda^{-1} = \int_0^\infty [1 - F_\lambda(t)]dt$)
P_s	probability of an unrecoverable sector (symbol) error
P_{DL}	probability of data loss during rebuild
P_{UF}	probability of data loss due to unrecoverable failures during rebuild
P_{DF}	probability of data loss due to a disk failure during rebuild
Q	amount of lost user data during rebuild
H	amount of lost user data, given that data loss has occurred during rebuild
S	number of lost symbols during rebuild

The remainder of the article is organized as follows. Section II reviews prior relevant work and analytical models presented in the literature for assessing the effect of latent errors on the reliability of erasure-coded systems. Section III describes the storage system model and the corresponding parameters considered. Section IV presents the general framework and methodology for deriving the MTTDL and EAFDL metrics analytically for the case of erasure-coded systems and in the presence of latent errors. Closed-form expressions for relevant reliability metrics are derived for the symmetric, clustered, and declustered data placement schemes. Section V presents numerical results demonstrating the adverse effect of unrecoverable or latent errors and the effectiveness of these schemes for improving system reliability. The reliability of real-world erasure coding schemes employed by enterprises to protect their stored data is assessed in Section VI. Finally, we conclude in Section VII.

II. RELATED WORK

The adverse effect of latent errors on the MTTDL reliability metric of RAID-5, RAID-6, replication, and erasure-coded systems has been demonstrated in [8][11][27-30]. Analytical reliability expressions for MTTDL that take into account the effect of latent errors have been obtained predominately using Markovian models, which assume that component failure and rebuild times are independent and exponentially distributed

[8][11][28][29]. The effect of latent errors on MTTDL of erasure-coded storage systems for the realistic case of non-exponential failure and rebuild time distributions was assessed in [30][35] for a limited range of error rates. In this article, we consider the entire range of sector error rates and assess the effect of latent errors not only on MTTDL, but also on the amount of lost data for the realistic case of non-exponential failure and rebuild time distributions.

Disk scrubbing has been used to mitigate the adverse effect of latent errors on system reliability [8][36][37][38]. The scrubbing process identifies latent errors at an early stage and attempts to correct them before disk failures occur. This in effect reduces the probability of encountering a latent error during the rebuild process. The resulting latent-error probability was derived in [8] as a function of the scrubbing and workload parameters. Subsequently, it was shown that the reliability level achieved when scrubbing is used can be obtained from the reliability level of a system that does not use scrubbing by adjusting the probability of encountering a latent error accordingly. The methodology presented in [8] for deriving the adjusted latent error probability when scrubbing is employed is also applicable for assessing the efficiency of other scrubbing schemes, such as the adaptive scrubbing schemes proposed in [37][38]. Moreover, this methodology can also be applied in conjunction with the reliability results presented in this article to assess the reliability of erasure-coded systems when scrubbing is used.

A simulation analysis of reliability aspects of erasure-coded data centers was presented in [39]. Various configurations were considered and it was shown that erasure codes and redundancy placement affect system reliability. In [32] it was recognized that it is hard to get statistically meaningful experimental reliability results using prototypes, because this would require a large number of machines to run for years. This underscores the usefulness of the analytical reliability results derived in this article.

III. Storage System Model

To assess the reliability of erasure-coded storage systems, we adopt the model used in [14] and extend it to cover the case of latent errors. The storage system comprises n storage devices (nodes or disks), where each device stores an amount c of data such that the total storage capacity of the system is $n\,c$. This does not account for the spare space used by the rebuild process.

A. Redundancy

User data is divided into blocks (or symbols) of a fixed size s (e.g., sector size of 512 bytes) and complemented with parity symbols to form codewords. We consider (m, l) maximum distance separable (MDS) erasure codes, which map l user-data symbols to a set of m ($> l$) symbols, called a codeword, having the property that any subset containing l of the m symbols can be used to reconstruct (recover) the codeword. The corresponding storage efficiency s_{eff} and amount U of user data stored in the system is

$$s_{\text{eff}} = l/m \quad \text{and} \quad U = s_{\text{eff}}\, n\, c \;=\; l\, n\, c/m\,. \tag{1}$$

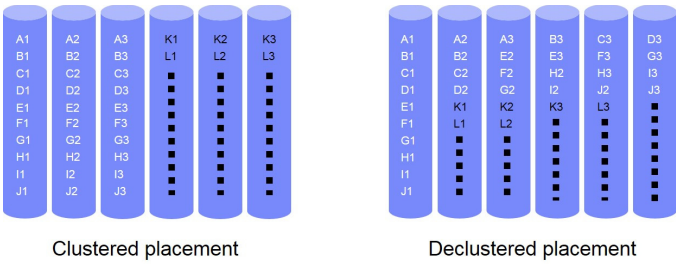

Figure 1. Clustered and declustered placement of codewords of length $m = 3$ on $n = 6$ devices. X1, X2, X3 represent a codeword (X = A, B, C, . . . , L).

Also, the number C of symbols stored in a device is

$$C = c/s\,. \tag{2}$$

Our notation is summarized in Table I. The derived parameters are listed in the lower part of the table. To minimize the risk of permanent data loss, the m symbols of each codeword are spread and stored on m distinct devices. This way, the system can tolerate any $\tilde{r} - 1$ device failures, but $\tilde{r}$ device failures may lead to data loss, with

$$\tilde{r} = m - l + 1\,, \quad 1 \le l < m \quad \text{and} \quad 2 \le \tilde{r} \le m\,. \tag{3}$$

Examples of MDS erasure codes are the replication, RAID-5, RAID-6, and Reed–Solomon schemes.

B. Symmetric Codeword Placement

In a symmetric placement scheme, the system effectively comprises n/k disjoint groups of k devices, and each codeword is placed entirely in one of these groups. Within each group, all $\binom{k}{m}$ possible ways of placing m symbols across k devices are used equally to store all the codewords in that group [40]. In particular, we consider the *clustered* and *declustered* placement schemes, as shown in Figure 1, which are special cases of symmetric placement schemes with k being equal to m and n, respectively. In the case of clustered placement, the storage system comprises n/m independent groups, referred to as *clusters*. Each codeword is stored across the devices of a particular cluster. In the case of declustered placement, all $\binom{n}{m}$ possible ways of placing m symbols across n devices are used equally to store all the codewords in the system.

C. Codeword Reconstruction and Rebuild Process

When storage devices fail, codewords lose some of their symbols, which immediately triggers the rebuild process.

1) Exposure Levels: The system is at exposure level u ($0 \le u \le \tilde{r}$) when there are codewords that have lost u symbols owing to device failures, but there are no codewords that have lost more symbols. These codewords are referred to as the *most-exposed* codewords. Transitions to higher exposure levels are caused by device failures, whereas transitions to lower ones are caused by successful rebuilds. We denote by C_u the number of most-exposed codewords upon entering exposure level u, ($u \ge 1$). Upon the first device failure it holds that

$$C_1 = C\,, \tag{4}$$

where C is determined by (2). In Section IV, we will derive the reliability metrics of interest using the *direct path*

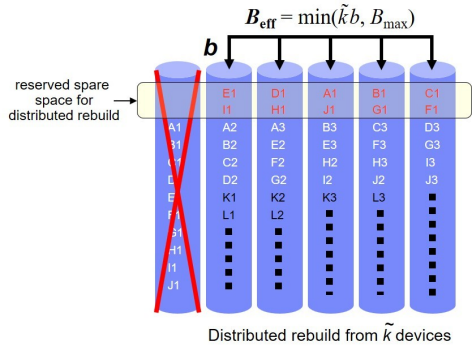

Figure 2. Rebuild under declustered placement.

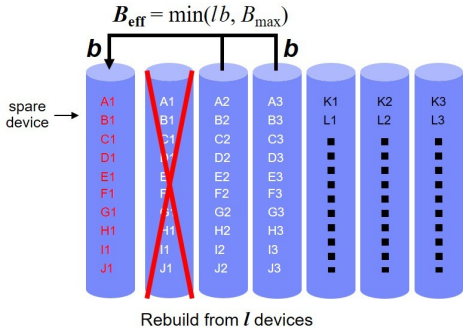

Figure 3. Rebuild under clustered placement.

approximation, which considers only transitions from lower to higher exposure levels [4][9][12][31][40]. This implies that each exposure level is entered only once.

2) Prioritized Rebuild: When a symmetric or declustered placement scheme is used, as shown in Figure 2, spare space is reserved on each device for temporarily storing the reconstructed codeword symbols before they are transferred to a new replacement device. The rebuild process to restore the data lost by failed devices is assumed to be both *prioritized* and *distributed*. A prioritized (or intelligent) rebuild process always attempts first to rebuild the *most-exposed* codewords, namely, the codewords that have lost the largest number of symbols [4][9][14][26][32]. At each exposure level u, it attempts to bring the system back to exposure level $u-1$ by recovering one of the u symbols that each of the C_u most-exposed codewords has lost by reading l of the remaining symbols. In a distributed rebuild process, the codewords are reconstructed by reading symbols from an appropriate set of surviving devices and storing the recovered symbols in the reserved spare space of these devices. During this process, it is desirable to reconstruct the lost codeword symbols on devices in which another symbol of the same codeword is not already present.

In the case of clustered placement, the codeword symbols are spread across all $k\ (= m)$ devices in each group (cluster). Therefore, reconstructing the lost symbols on the surviving devices of a group would result in more than one symbol of the same codeword on the same device. To avoid this, the lost symbols are reconstructed directly in spare devices as shown in Figure 3 and described in [14].

TABLE II. NOTATION OF SYSTEM PARAMETERS AT EXPOSURE LEVELS

Parameter	Definition
u	exposure level
C_u	number of most-exposed codewords upon entering exposure level u
$\tilde{n}_u$	number of devices at exposure level u whose failure causes an exposure level transition to level $u+1$
V_u	fraction of the most-exposed codewords that have symbols stored on any given device from the $\tilde{n}_u$ devices
R_u	rebuild time at exposure level u
α_u	fraction of the rebuild time R_u still left when another device fails, causing the exposure level transition $u \to u+1$
$P_{u\to u+1}$	transition probability from exposure level u to $u+1$
b_u	average rate at which recovered data is written at exposure level u

3) Rebuild Process: A certain portion of the device bandwidth is reserved for read/write data recovery during the rebuild process, and the remaining bandwidth is used to serve user requests. Let b denote the actual average reserved rebuild bandwidth per device. The lost symbols are rebuilt in parallel using the rebuild bandwidth available on each surviving device. Let us denote by b_u $(\leq b)$ the average rate at which the amount of data corresponding to the number C_u of symbols to be rebuilt at exposure level u is written to selected device(s). This rate depends on $B_{\max}$, the upper limitation of the average network rebuild bandwidth [14]. Also, let $1/\mu$, $f_X(.)$, and $F_X(.)$ denote the mean, the probability density function, and the cumulative distribution function of the time X required to read (or write) an amount c of data from (or to) a device, respectively. The kth moment of X, $E(X^k)$, and its mean $E(X)$ are then given by

$$E(X^k) = \int_0^\infty t^k f_X(t)dt\ , \quad \text{for } k = 1, 2, \dots\ , \tag{5}$$

$$\mu^{-1} \triangleq E(X) = c/b\ . \tag{6}$$

4) Failure and Rebuild Time Distributions: The lifetimes of the n devices are assumed to be independent and identically distributed, with a cumulative distribution function $F_\lambda(.)$ and a mean of $1/\lambda$. We consider real-world distributions, such as Weibull and gamma, as well as exponential distributions that belong to the large class defined in [31]. Note that, although the model considered here does not account for correlated device failures, their effect can be assessed by enhancing the model according to the approach presented in [8]. This issue, however, is beyond the scope of this article. The results in this article hold for *highly reliable* storage devices, which satisfy the condition [14][31]

$$\mu \int_0^\infty F_\lambda(t)[1 - F_X(t)]dt \ll 1, \quad \text{with} \quad \frac{\lambda}{\mu} \ll 1\ . \tag{7}$$

This condition expresses the fact that the ratio of the mean time $1/\mu$ to read all contents of a device (which typically is on the order of tens of hours) to the mean time to failure of a device $1/\lambda$ (which is typically at least on the order of thousands of hours) is very small, and in particular the fact that it is very unlikely that a given device fails during a rebuild period.

When the devices are highly reliable, the MTTDL and EAFDL reliability metrics of erasure-coded storage systems tend to be insensitive to the device failure distribution, that is, they depend only on its mean $1/\lambda$, but not on its density $F_\lambda(.)$. They are, however, sensitive to the distribution $F_X(.)$ of the device rebuild times [14].

5) Amount of Data to Rebuild and Rebuild Times at Each Exposure Level: We denote by $\tilde{n}_u$ the number of devices at exposure level u whose failure causes an exposure level transition to level $u+1$ and V_u the fraction of the C_u most-exposed codewords that have a symbol stored on any given such device. Note that $\tilde{n}_u$ depends on the codeword placement scheme. The notation used here is summarized in Table II. Let R_u denote the rebuild time of the most-exposed codewords at exposure level u. At exposure level 1, the amount of data to be recovered is equal to c. Given that this data is recovered at an average rate of b_1 and that the time required to write an amount c of data at an average rate of b is equal to X, it follows that the rebuild time R_1 is given by

$$R_1 = (b/b_1)\, X \ . \tag{8}$$

Let α_u be the fraction of the rebuild time R_u still left when another device fails, causing the exposure level transition $u \to u+1$. In [41, Lemma 2], it was shown that, for highly reliable devices satisfying condition (7), α_u is approximately uniformly distributed in $(0,1)$, that is

$$\alpha_u \sim U(0,1), \quad u = 1, \ldots, \tilde{r}-1 \ . \tag{9}$$

We proceed by considering that the rebuild time R_{u+1} is determined completely by R_u and α_u in the same manner as in [13][14][40]. For the rebuild schemes considered, the fraction of the C_u most-exposed codewords that were not yet considered by the rebuild process upon the next device failure is roughly equal to the fraction α_u of the rebuild time R_u still left. Therefore, upon the next device failure, an approximate number $\alpha_u\, C_u$ of the C_u codewords were not yet considered by the rebuild process. Clearly, the fraction V_u of these codewords that have symbols stored on the newly failed device depends only on the codeword placement scheme. Consequently, the number C_{u+1} of the most-exposed codewords upon entering exposure level $u+1$ is

$$C_{u+1} \approx V_u\, \alpha_u\, C_u \ , \ \text{ for } \ u = 1, \ldots, \tilde{r}-1 \ . \tag{10}$$

Repeatedly applying (10) and using (4) and the convention that for any sequence δ_i, $\prod_{i=1}^{0} \delta_i \triangleq 1$, yields

$$C_u \ \approx \ C \prod_{i=1}^{u-1} V_i\, \alpha_i \ , \ \text{ for } \ u = 1, \ldots, \tilde{r} \ . \tag{11}$$

Unconditioning (11) on $\alpha_1, \ldots, \alpha_{u-1}$ yields

$$E(C_u) = C \left(\prod_{j=1}^{u-1} V_j \right) E \left(\prod_{j=1}^{u-1} \alpha_j \,\middle|\, \text{level } u \text{ was entered} \right), \quad \text{for } \ u = 1, \ldots, \tilde{r} \ . \tag{12}$$

6) Unrecoverable Errors: The reliability of storage systems is affected by the occurrence of unrecoverable or latent errors. Let P_{bit} denote the unrecoverable bit-error probability. According to the specifications, P_{bit} is equal to 10^{-15} for SCSI drives and 10^{-14} for SATA drives [8]. Assuming that bit errors occur independently over successive bits, the unrecoverable sector (symbol) error probability P_s is

$$P_s = 1 - (1 - P_{\text{bit}})^s \ , \tag{13}$$

with s expressed in bits. Assuming a sector size of 512 bytes, the equivalent unrecoverable sector error probability is $P_s \approx P_{\text{bit}} \times 4096$, which is 4.096×10^{-12} in the case of SCSI and 4.096×10^{-11} in the case of SATA drives. In practice, however, and also owing to the accumulation of latent errors over time, these probability values are higher. Indeed, empirical field results suggest that the actual values can be orders of magnitude higher, reaching $P_s \approx 5 \times 10^{-9}$ [42].

IV. Derivation of MTTDL and EAFDL

The MTTDL metric assesses the expected time until some data can no longer be recovered and therefore is lost forever, whereas the EAFDL assesses the fraction of stored data that is expected to be lost by the system annually. The reliability metrics are derived using the methodology presented in [12][13][14] and extending it to assess the effect of latent errors. This methodology uses the direct path approximation [11], does not involve Markovian analysis [4][9][12][31][40], and holds for general failure time distributions, which can be exponential or non-exponential, such as the Weibull and gamma distributions that satisfy condition (7).

At any point in time, the system can be thought to be in one of two modes: normal or rebuild mode. During normal mode, all devices are operational and all data in the system has the original amount of redundancy. A *first device* failure causes a transition from normal to rebuild mode. A rebuild process attempts to restore the lost data, which eventually leads the system either to a data loss (DL) with probability P_{DL} or back to the original normal mode by restoring initial redundancy, with probability $1 - P_{\text{DL}}$. Any symbols encountered with unrecoverable or latent errors are usually corrected by the erasure coding capability. However, it may not be possible to recover multiple unrecoverable errors in a codeword, which therefore leads to data loss.

Let T be a typical interval of a fully operational period, that is, the interval from the time t that the system is brought to its original state until a subsequent first device failure occurs. For a system comprising n devices with a mean time to failure of a device $1/\lambda$, the expected duration of T is [12]

$$E(T) = \frac{1}{n\,\lambda} \ , \tag{14}$$

and MTTDL is

$$\text{MTTDL} \approx \frac{E(T)}{P_{\text{DL}}} = \frac{1}{n\,\lambda\,P_{\text{DL}}} \ . \tag{15}$$

The EAFDL is obtained as the ratio of the expected amount $E(Q)$ of lost user data, normalized to the amount U of user data, to the expected duration of T [12, Eq. (9)]:

$$\text{EAFDL} \approx \frac{E(Q)}{E(T)\cdot U} \overset{(14)}{=} \frac{n\,\lambda\,E(Q)}{U} \overset{(1)}{=} \frac{m\,\lambda\,E(Q)}{l\,c} \ , \tag{16}$$

with $E(T)$ and $1/\lambda$ expressed in years.

The expected conditional amount $E(H)$ of lost user data, given that data loss has occurred, is determined by [12, Eq. (8)]:

$$E(H) = \frac{E(Q)}{P_{\text{DL}}} \ . \tag{17}$$

It follows from (15), (16), and (17) that

$$\text{EAFDL} = \frac{E(H)}{\text{MTTDL}\cdot U} \ , \tag{18}$$

with the MTTDL expressed in years.

TABLE III. NOTATION OF RELIABILITY METRICS AT EXPOSURE LEVELS

Parameter	Definition
u	exposure level
P_u	probability of entering exposure level u
P_{UF_u}	probability of data loss due to unrecoverable symbol errors at exposure level u
P_{UF}	probability of data loss due to unrecoverable symbol errors
P_{DF}	probability of data loss due to $\tilde{r}$ successive device failures
P_{DL}	probability of data loss
q_u	probability that, at exposure level u, a codeword that has lost u symbols can be restored
$\hat{q}_u$	probability that, under instantaneous transitions from exposure level 1 to exposure level u, all of the C_u most-exposed codewords, which have lost u symbols, can be restored

A. Reliability Analysis

At any exposure level u $(u = 1, \ldots, \tilde{r}-1)$, data loss may occur during rebuild owing to one or more unrecoverable failures, which is denoted by the transition $u \to \text{UF}$. Moreover, at exposure level $\tilde{r}-1$, data loss occurs owing to a subsequent device failure, which leads to the transition to exposure level $\tilde{r}$. Consequently, the direct paths that lead to data loss are the following:

$\overrightarrow{UF_u}$: the direct path of successive transitions $1 \to 2 \to \cdots \to u \to \text{UF}$, for $u = 1, \ldots, \tilde{r}-1$, and

$\overrightarrow{DF}$: the direct path of successive transitions $1 \to 2 \to \cdots \to \tilde{r}-1 \to \tilde{r}$,

with corresponding probabilities P_{UF_u} and P_{DF}, respectively. The notation for the probabilities of the events that lead to data loss is summarized in Table III.

1) Data Loss: The probability P_{UF} of data loss owing to unrecoverable failures is

$$P_{\text{UF}} \approx \sum_{u=1}^{\tilde{r}-1} P_{\text{UF}_u} \,, \tag{19}$$

where P_{UF_u} denotes the probability of data loss associated with the direct path $\overrightarrow{UF_u}$. Also, it holds that

$$P_{\text{UF}_u} = P_u \; P_{u \to \text{UF}} \,, \quad \text{for} \;\; u = 1, \ldots, \tilde{r}-1 \,, \tag{20}$$

where P_u is the probability of entering exposure level u, which is derived in Appendix A as follows:

$$P_u \approx (\lambda\, c)^{u-1} \; \frac{1}{(u-1)!} \; \frac{E(X^{u-1})}{[E(X)]^{u-1}} \; \prod_{i=1}^{u-1} \frac{\tilde{n}_i}{b_i} \; V_i^{u-1-i} \,, \tag{21}$$

and $P_{u \to \text{UF}}$ is the probability of encountering an unrecoverable failure during the rebuild process at this exposure level.

In [11], it was shown that P_{DL} is accurately approximated by the probability of all direct paths to data loss. Therefore,

$$P_{\text{DL}} \approx \; P_{\text{DF}} + \sum_{u=1}^{\tilde{r}-1} P_{\text{UF}_u} \; \overset{(19)}{\approx} \; P_{\text{DF}} + P_{\text{UF}} \,. \tag{22}$$

Approximate expressions for the probabilities of data loss P_{UF_u} and P_{DF} are subsequently obtained by the following proposition.

Proposition 1: For $u = 1, \ldots, \tilde{r}-1$, it holds that

$$P_{\text{UF}_u} \;\approx\; -(\lambda\, c)^{u-1} \; \frac{E(X^{u-1})}{[E(X)]^{u-1}} \; \left(\prod_{i=1}^{u-1} \frac{\tilde{n}_i}{b_i} \; V_i^{u-1-i} \right) \cdot \log(\hat{q}_u)^{-(u-1)} \left(\hat{q}_u - \sum_{i=0}^{u-1} \; \frac{\log(\hat{q}_u)^i}{i!} \right), \tag{23}$$

where $$\hat{q}_u \triangleq q_u^{\,C \, \prod_{j=1}^{u-1} V_j} \,, \tag{24}$$

$$q_u = 1 - \sum_{j=\tilde{r}-u}^{m-u} \binom{m-u}{j} P_s^j (1-P_s)^{m-u-j} \,, \tag{25}$$

$$P_{\text{DF}} \approx (\lambda\, c)^{\tilde{r}-1} \; \frac{1}{(\tilde{r}-1)!} \; \frac{E(X^{\tilde{r}-1})}{[E(X)]^{\tilde{r}-1}} \; \prod_{i=1}^{\tilde{r}-1} \frac{\tilde{n}_i}{b_i} \; V_i^{\tilde{r}-1-i}, \tag{26}$$

where $E(X^{\tilde{r}-1})$ is obtained from (5).

Proof: Equation (23) is obtained in Appendix A. Equation (26) is obtained from the fact that $P_{\text{DF}} = P_{\tilde{r}}$ and, subsequently, from (21) by setting $u = \tilde{r}$. ■

The MTTDL metric is obtained by substituting (22) into (15) as follows:

$$\text{MTTDL} \approx \frac{1}{n\,\lambda\,(P_{\text{DF}} + \sum_{u=1}^{\tilde{r}-1} P_{\text{UF}_u})} \,, \tag{27}$$

where P_{UF_u} and P_{DF} are determined by (23) and (26), respectively.

2) Amount of Data Loss: We proceed to derive the amount of data loss during rebuild. Let Q, H, and S be the amount of lost user data, the conditional amount of lost user data, given that data loss has occurred, and the number of lost symbols, respectively. Let also Q_{DF} and Q_{UF_u} denote the amount of lost user data associated with the direct paths $\overrightarrow{DF}$ and $\overrightarrow{UF_u}$, respectively. Similarly, we consider the variables H_{DF}, H_{UF_u}, S_{DF}, and S_{UF_u}. Then, the amount Q of lost user data is obtained by

$$Q \approx \begin{cases} H_{\text{DF}} \,, & \text{if } \overrightarrow{DF} \\ H_{\text{UF}_u} \,, & \text{if } \overrightarrow{UF_u} \,, \;\; \text{for} \;\; u = 1, \ldots, \tilde{r}-1 \\ 0 \,, & \text{otherwise} \,. \end{cases} \tag{28}$$

Therefore,

$$E(Q) \;\approx P_{\text{DF}} \; E(H_{\text{DF}}) + \sum_{u=1}^{\tilde{r}-1} P_{\text{UF}_u} \; E(H_{\text{UF}_u}) \tag{29}$$

$$= E(Q_{\text{DF}}) + \sum_{u=1}^{\tilde{r}-1} E(Q_{\text{UF}_u}) \tag{30}$$

$$\approx \; E(Q_{\text{DF}}) + E(Q_{\text{UF}}) \,, \tag{31}$$

where

$$E(Q_{\text{DF}}) = \; P_{\text{DF}} \; E(H_{\text{DF}}) \,, \tag{32}$$

$$E(Q_{\text{UF}_u}) = \; P_{\text{UF}_u} \; E(H_{\text{UF}_u}) \,, \;\; \text{for} \;\; u = 1, \ldots, \tilde{r}-1 \tag{33}$$

$$E(Q_{\text{UF}}) = P_{\text{UF}} \; E(H_{\text{UF}}) \approx \sum_{u=1}^{\tilde{r}-1} E(Q_{\text{UF}_u}) \,, \tag{34}$$

where Q_{UF} denotes the amount of lost user data due to unrecoverable failures and H_{UF} the conditional amount of lost

user data, given that data loss due to unrecoverable failures has occurred.

Note that the expected amount $E(Q)$ of lost user data is equal to the product of the storage efficiency and the expected amount of lost data, where the latter is equal to the product of the expected number of lost symbols $E(S)$ and the symbol size s. Consequently, it follows from (1) that

$$E(Q) = \frac{l}{m}\, E(S)\, s \overset{(2)}{=} \frac{l}{m}\, \frac{E(S)}{C}\, c\,. \tag{35}$$

Similarly,

$$E(Q_{\text{DF}}) = \frac{l}{m}\, E(S_{\text{DF}})\, s \overset{(2)}{=} \frac{l}{m}\, \frac{E(S_{\text{DF}})}{C}\, c\,, \tag{36}$$

$$E(Q_{\text{UF}_u}) = \frac{l}{m}\, E(S_{\text{UF}_u})\, s \overset{(2)}{=} \frac{l}{m}\, \frac{E(S_{\text{UF}_u})}{C}\, c\,. \tag{37}$$

Approximate expressions for the expected amounts $E(Q_{\text{UF}_u})$ and $E(Q_{\text{DF}})$ of lost user data are subsequently obtained by the following proposition.

Proposition 2: For $u = 1, \ldots, \tilde{r} - 1$, it holds that

$$E(Q_{\text{UF}_u}) \approx\, c\, \frac{l\,\tilde{r}}{m}\, (\lambda\, c)^{u-1}\, \frac{1}{u!}\, \frac{E(X^{u-1})}{[E(X)]^{u-1}}\, \left(\prod_{i=1}^{u-1} \frac{\tilde{n}_i}{b_i}\, V_i^{u-i}\right) \cdot \binom{m-u}{\tilde{r}-u} P_s^{\tilde{r}-u}\,, \quad \text{for } P_s \ll \frac{1}{m-\tilde{r}}\,, \tag{38}$$

$$E(Q_{\text{DF}}) \approx\, c\, \frac{l}{m}\, (\lambda\, c)^{\tilde{r}-1}\, \frac{1}{(\tilde{r}-1)\,!}\, \frac{E(X^{\tilde{r}-1})}{[E(X)]^{\tilde{r}-1}}\, \prod_{i=1}^{\tilde{r}-1} \frac{\tilde{n}_i}{b_i}\, V_i^{\tilde{r}-i}, \tag{39}$$

where $E(X^{u-1})$ and $E(X^{\tilde{r}-1})$ are obtained from (5).

Proof: Equations (38) and (39) are obtained in Appendix B. Note that (39) can also be obtained from (38) by setting $u = \tilde{r}$, which is the same result as Eq. (25) of [14]. ■

The EAFDL metric is obtained by substituting (30) into (16) as follows:

$$\text{EAFDL} \approx \frac{m\,\lambda\,[E(Q_{\text{DF}}) + \sum_{u=1}^{\tilde{r}-1} E(Q_{\text{UF}_u})]}{l\, c}\,, \tag{40}$$

where $E(Q_{\text{UF}_u})$ and $E(Q_{\text{DF}})$ are determined by (38) and (39), respectively.

The conditional amounts $E(H)$, $E(H_{\text{DF}})$, $E(H_{\text{UF}_u})$, and $E(H_{\text{UF}})$ of lost user data, given that data loss has occurred, are obtained from (17), (32), (33), and (34), respectively.

Remark 1: From (26), (32), and (39), it follows that

$$E(H_{\text{DF}}) \approx \left(\frac{l}{m} \prod_{i=1}^{\tilde{r}-1} V_i\right) c\,. \tag{41}$$

Note that when entering exposure level $\tilde{r}$, for each of the $C_{\tilde{r}}$ most-exposed codewords there are $\tilde{r}$ symbols permanently lost. Consequently, the expected number of user-data symbols permanently lost is $C_{\tilde{r}}\,(l/m)\,\tilde{r}$, which implies that, for a symbol size of s, the expected amount $E(H_{\text{DF}}\,|\,C_{\tilde{r}})$ of user data lost is

$$E(H_{\text{DF}}\,|\,C_{\tilde{r}}) = C_{\tilde{r}}\, \frac{l}{m}\, \tilde{r}\, s\,. \tag{42}$$

Unconditioning (42) on $C_{\tilde{r}}$ yields

$$E(H_{\text{DF}}) = E(C_{\tilde{r}})\, \frac{l}{m}\, \tilde{r}\, s\,. \tag{43}$$

Combining (41), (43), and using (2), yields

$$E(C_{\tilde{r}}) \approx \left(\prod_{i=1}^{\tilde{r}-1} V_i\right) \frac{C}{\tilde{r}}\,. \tag{44}$$

From (12) and (44), it follows that

$$E\left(\prod_{j=1}^{\tilde{r}-1} \alpha_j \,\middle|\, \text{level } \tilde{r} \text{ was entered}\right) \approx \frac{1}{\tilde{r}}\,. \tag{45}$$

Remark 2: It turns out that when a data loss has occurred, the variables $\alpha_1, \ldots, \alpha_{\tilde{r}-1}$ are not distributed identically. More specifically, for a rebuild time R_u, the uniform distribution of α_u in the interval $(0, 1)$, given by (9), holds under the assumption that there is a failure during this rebuild period, that is, an exposure level transition $u \to u+1$. However, conditioning on the exposure level transitions $u \to u+1 \to \cdots \to u' \to u'+1$ ($u' > u$), α_u is no longer uniformly distributed in $(0, 1)$. This is due to the fact that, conditioning on the fact that additional failures occur during the rebuild times $R_{u+1}, \ldots, R_{u'}$, it is more likely that the R_{u+1} period is long rather than short. In this case, only α'_u is uniformly distributed in $(0, 1)$. Assuming that the system has entered exposure level u, we deduce from (45) that

$$E\left(\prod_{j=1}^{u-1} \alpha_j \,\middle|\, \text{level } u \text{ was entered}\right) \approx \frac{1}{u}\,, \quad \text{for } \; u = 2, \ldots, \tilde{r}\,. \tag{46}$$

Substituting (46) into (12) and using (4) yields

$$E(C_u) \approx \left(\prod_{i=1}^{u-1} V_i\right) \frac{C}{u}\,, \quad \text{for } \; u = 1, \ldots, \tilde{r}\,. \tag{47}$$

Remark 3: For small values of P_s, it holds that $P_{\text{UF}} \to 0$ and $E(Q_{\text{UF}}) \to 0$. Therefore, from (22) and (31), for small values of P_s, it holds that $P_{\text{DL}} \to P_{\text{DF}}$ and $E(Q) \to E(Q_{\text{DF}})$. Consequently, from (17), (32), and (41), it follows that

$$E(H) \approx E(H_{\text{DF}}) \approx \left(\frac{l}{m} \prod_{i=1}^{\tilde{r}-1} V_i\right) c\,, \quad \text{for } P_s \to 0\,. \tag{48}$$

When P_s is extremely small and an unrecoverable failure occurs, this failure most likely occurs when rebuilding a codeword after it has lost $\tilde{r} - 1$ of its symbols owing to $\tilde{r} - 1$ device failures and an unrecoverable error is encountered. In this case, $\tilde{r}$ of its symbols are lost and therefore the expected number of lost user symbols is equal to the product of the storage efficiency l/m and $\tilde{r}$, which implies that

$$E(H_{\text{UF}}) \approx \frac{l}{m}\, \tilde{r}\, s \overset{(2)}{=} \frac{1}{C}\, \frac{l}{m}\, \tilde{r}\, c\,, \quad \text{for } P_s \to 0\,. \tag{49}$$

Remark 4: For large values of P_s, it holds that

$$P_{\text{DL}} \approx P_{\text{UF}} \approx P_{\text{UF}_1} \approx 1\,, \qquad \text{for } P_s \to 1\,, \tag{50}$$

which, by virtue of (15), implies that

$$\text{MTTDL} \approx \frac{1}{n\,\lambda}\,, \quad \text{for } P_s \to 1\,. \tag{51}$$

It also holds that

$$E(Q) \approx E(H) \approx \; l\,c\,, \qquad \text{for } P_s \to 1\,. \tag{52}$$

From (18), and using (51) and (52), it follows that

$$\text{EAFDL} \approx \; m\,\lambda\,, \quad \text{for } P_s \to 1\,. \tag{53}$$

B. Symmetric and Declustered Placement

We consider the case $m < k \leq n$. The special case $k = m$ corresponding to the clustered placement scheme has to be considered separately for the reasons discussed in Section III-C2. At each exposure level u, for $u = 1, \cdots, \tilde{r}-1$, it holds that [13][14]

$$\tilde{n}_u^{\text{sym}} = \; k - u\,, \tag{54}$$

$$b_u^{\text{sym}} = \frac{\min((k-u)\,b, B_{\text{max}})}{l+1}\,, \tag{55}$$

$$V_u^{\text{sym}} = \; \frac{m-u}{k-u}\,. \tag{56}$$

The corresponding parameters $\tilde{n}_u^{\text{declus}}$, b_u^{declus}, and V_u^{declus} for the declustered placement are derived from (54), (55), and (56) by setting $k = n$. which yields

$$\tilde{n}_u^{\text{declus}} = \; n - u\,, \tag{57}$$

$$b_u^{\text{declus}} = \frac{\min((n-u)\,b, B_{\text{max}})}{l+1}\,, \tag{58}$$

$$V_u^{\text{declus}} = \; \frac{m-u}{n-u}\,. \tag{59}$$

C. Clustered Placement

At any exposure level u $(u = 1, \ldots, \tilde{r}-1)$, it holds that [13][14]

$$\tilde{n}_u^{\text{clus}} = m - u\,, \;\; b_u^{\text{clus}} = \min(\,b\,, B_{\text{max}}/l\,)\,, \;\; V_u^{\text{clus}} = 1\,. \tag{60}$$

Remark 5: It follows from (60) that a system is not bandwidth-constrained when $B_{\text{max}} \geq l\,b$. Then, $b_u^{\text{clus}} = \min(b, B_{\text{max}}/l) = b$. In the case of RAID-5 and RAID-6, it holds that $m - l = 1$ and $m - l = 2$ or, equivalently, $\tilde{r} = 2$ and $\tilde{r} = 3$, respectively, such that (22), (23), and (26) yield

$$P_{\text{DL}}^{\text{RAID-5}} \approx \; (m-1)\,\frac{\lambda\,c}{b} \; + \; 1 - (1-P_s)^{(m-1)\,C}\,, \tag{61}$$

which is the same result as Eq. (85) of [16] (with $\frac{c}{b} = \frac{1}{\mu}$), and

$$P_{\text{DL}}^{\text{RAID-6}} \approx \; 1 - q_1^C \; + \left[1 + \frac{1-q_2^C}{\log(q_2^C)}\right](m-1)\,\frac{\lambda\,c}{b} + \frac{(m-1)(m-2)}{2}\left(\frac{\lambda\,c}{b}\right)^2 \frac{E(X^2)}{[E(X)]^2}\,, \tag{62}$$

where q_1, q_2 are determined by (25). This result is in agreement with Eq. (243) of [16] (with $\frac{c}{b} = \frac{1}{\mu}$). Also, (30), (38), and (39) yield

$$E(Q^{\text{RAID-5}}) \approx \; \frac{l}{m}\,(m-1)\left(\frac{\lambda\,c}{b} + 2\,P_s\right)c\,, \tag{63}$$

TABLE IV. TYPICAL VALUES OF DIFFERENT PARAMETERS

Parameter	Definition	Values
n	number of storage devices	64
c	amount of data stored on each device	20 TB
s	symbol (sector) size	512 B
λ^{-1}	mean time to failure of a storage device	876,000 h
b	rebuild bandwidth per device	100 MB/s
m	symbols per codeword	16
l	user-data symbols per codeword	13, 14, 15
U	amount of user data stored in the system	1.04 to 1.2 PB
μ^{-1}	time to read an amount c of data at a rate b from a storage device	55.5 h

which is the same result as Eq. (105) of [16] (with $\frac{c}{b} = \frac{1}{\mu}$), and

$$E(Q^{\text{RAID-6}}) \approx \frac{l}{m}\,\frac{(m-1)(m-2)}{2} \cdot\left[\left(\frac{\lambda\,c}{b}\right)^2 \frac{E(X^2)}{[E(X)]^2} \; + 3\,\frac{\lambda\,c}{b}\,P_s + 3\,P_s^2\right]c\,. \tag{64}$$

This result is in agreement with Eq. (264) of [16] (with $\frac{c}{b} = \frac{1}{\mu}$).

V. NUMERICAL RESULTS

Here we assess the reliability of the clustered and declustered schemes for a system comprised of $n = 64$ devices (disks) and protected by an erasure coding scheme with $m = 16$, which is the codeword length used by Microsoft® Azure [26], and $l = 13$, 14, and 15. Each device stores an amount of $c = 20$ TB, which is the capacity of the latest generation of Seagate drives, and the symbol size s is equal to a sector size of 512 bytes [43].

Typical parameter values are listed in Table IV. The annualized failure rate (AFR) of HDDs for the year 2021 is in the range of 0.11% to 4.79% [44], which corresponds to a mean time to failure in the range of $180,000$ h to $8,000,000$ h. The parameter λ^{-1} is chosen to be equal to $876,000$ h (100 years) that corresponds to an AFR of 1%, which is the average AFR across all drive models [44]. Considering that 35% of the maximum transfer rate of 285 MB/s [43] is allocated for recovery operations, the reserved rebuild bandwidth b is then equal to 100 MB/s, which yields a rebuild time of a device $\mu^{-1} = c/b = 55.5$ h. Also, it is assumed that the maximum network rebuild bandwidth is sufficiently large ($B_{\text{max}} \geq n\,b =$ 6.4 GB/s), that the rebuild time distribution is deterministic, such that $E(X^k) = [E(X)]^k$. The obtained results are accurate, because (7) is satisfied, given that $\lambda/\mu = 6.3\times10^{-5} \ll 1$.

First, we assess the reliability for the declustered placement scheme ($k = n = 64$). The probability of data loss P_{DL} is determined by (22) as a function of P_s and shown in Figure 4. The probabilities P_{UF_u} and P_{DF} are also shown, as obtained from (23) and (26), respectively. We observe that P_{DL} increases monotonically with P_s and exhibits a number of $\tilde{r}$ plateaus. In the interval $[4.096 \times 10^{-12}, 5 \times 10^{-9}]$ of practical importance for P_s, which is indicated between the two vertical dashed lines, the probability of data loss P_{DL} and, by virtue of (15), the MTTDL are degraded by orders of magnitude. The normalized λ MTTDL measure is obtained from (15) and shown in Figure 5. Increasing the number of parities (reducing l) improves reliability by orders of magnitude.

The normalized expected amount $E(Q)/c$ of lost user data relative to the amount of data stored in a device is

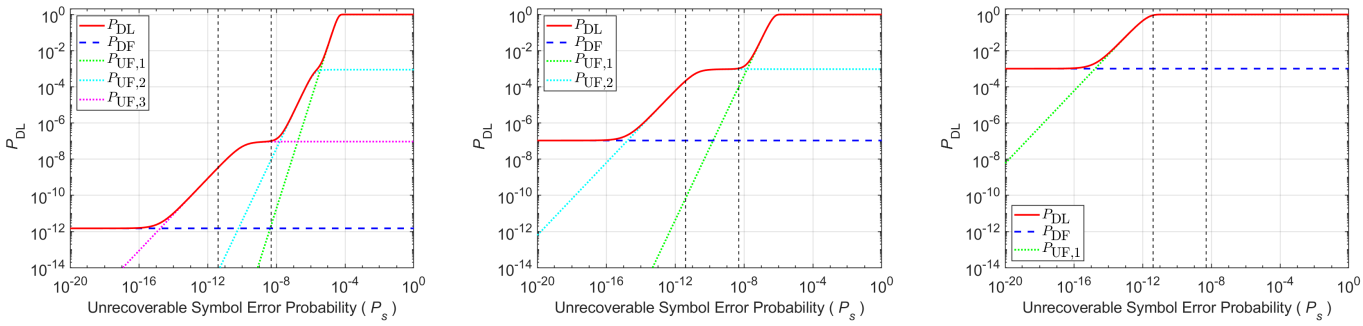

(a) $l = 13 \quad (\tilde{r} = 4)$ (b) $l = 14 \quad (\tilde{r} = 3)$ (c) $l = 15 \quad (\tilde{r} = 2)$

Figure 4. Probability of data loss P_{DL} vs. P_s for $l = 13, 14, 15$; $m = 16$, $n = k = 64$ (declustered scheme).

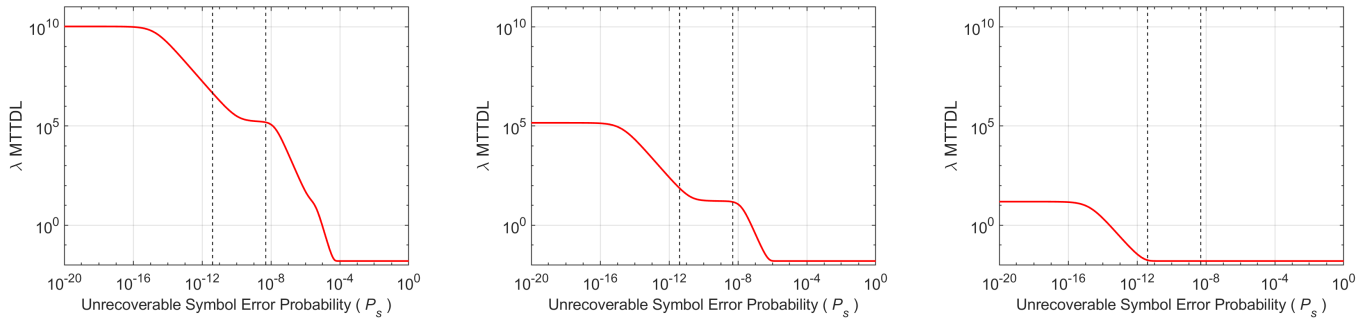

(a) $l = 13 \quad (\tilde{r} = 4)$ (b) $l = 14 \quad (\tilde{r} = 3)$ (c) $l = 15 \quad (\tilde{r} = 2)$

Figure 5. Normalized MTTDL vs. P_s for $l = 13, 14,$ and 15; $m = 16$, $n = k = 64$ (declustered scheme).

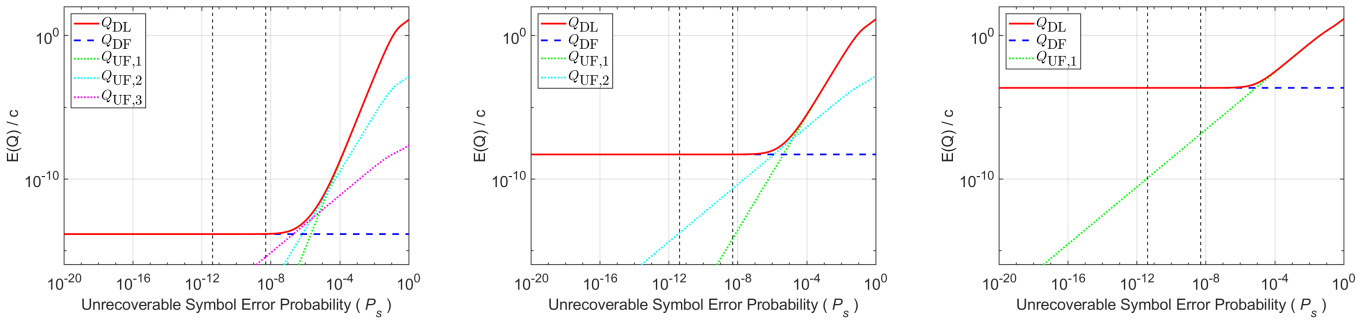

(a) $l = 13 \quad (\tilde{r} = 4)$ (b) $l = 14 \quad (\tilde{r} = 3)$ (c) $l = 15 \quad (\tilde{r} = 2)$

Figure 6. Normalized amount of data loss $E(Q)$ vs. P_s for $l = 13, 14, 15$; $m = 16$, $n = k = 64$ (declustered scheme).

obtained from (30) and shown in Figure 6. The normalized expected amounts $E(Q_{\text{UF}_u})/c$ and $E(Q_{\text{DF}})/c$ are also shown as determined by (38) and (39), respectively. The normalized EAFDL$/\lambda$ measure is obtained from (16) and shown in Figure 7. We observe that $E(Q)$ and EAFDL increase monotonically, but they are practically unaffected in the interval of interest, because they degrade only when P_s is much larger than the typical sector error probabilities. For the EAFDL metric too, increasing the number of parities (reducing l) results in a reliability improvement by orders of magnitude.

The normalized expected amount $E(H)/c$ of lost user data, given that a data loss has occurred, relative to the amount of data stored in a device is obtained from (17) and shown in Figure 8. The conditional amounts $E(H_{\text{DF}})$ and $E(H_{\text{UF}})$ obtained from (32) and (34), respectively, are also shown. In contrast to the P_{DL}, EAFDL, and $E(Q)$ metrics that increase monotonically with P_s, we observe that $E(H)$ does not do so. The reason for that is the following. As shown in Figure 4, for $P_s \gg 10^{-14}$, data loss is more likely to be due to sector errors than to device failures. Given that sector errors result in a negligible amount of data loss compared with the substantial data losses caused by device failures, when P_s increases over the value of 10^{-14}, the conditional amount of lost data decreases. Clearly, this is reversed for high values of P_s, and the conditional amount of lost data increases.

The expected amount $E(H_{\text{UF}})$ of user data lost due to unrecoverable failures, given that such failures have occurred, is shown in Figure 8 by the dotted green line. For extremely small values of P_s, and according to Remark 3, the value of $E(H_{\text{UF}})$ corresponds to a single corrupted codeword that loses

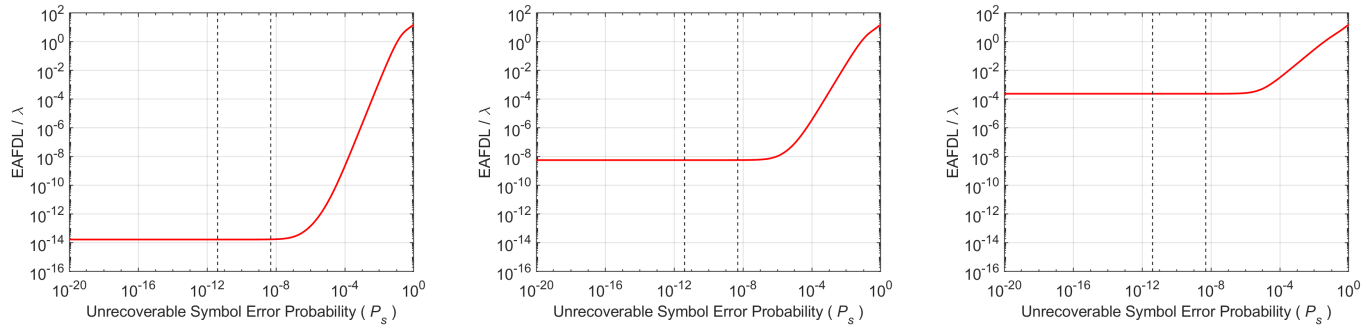

(a) $l = 13 \quad (\tilde{r} = 4)$ (b) $l = 14 \quad (\tilde{r} = 3)$ (c) $l = 15 \quad (\tilde{r} = 2)$

Figure 7. Normalized EAFDL vs. P_s for $l = 13, 14,$ and 15; $m = 16$, $n = k = 64$ (declustered scheme).

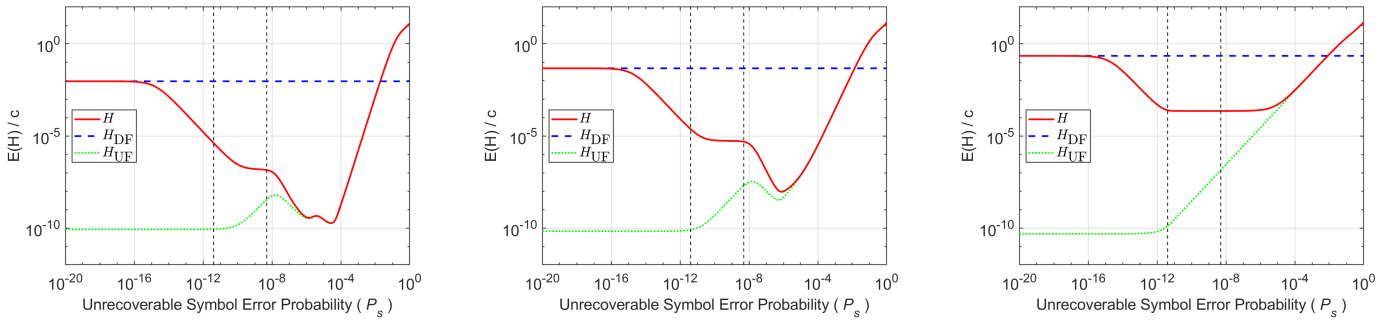

(a) $l = 13 \quad (\tilde{r} = 4)$ (b) $l = 14 \quad (\tilde{r} = 3)$ (c) $l = 15 \quad (\tilde{r} = 2)$

Figure 8. Normalized $E(H)$ vs. P_s for $l = 13, 14,$ and 15; $m = 16$, $n = k = 64$ (declustered scheme).

$\tilde{r}$ of its symbols of which $\tilde{r}-1$ are lost owing to device failures and one is lost owing to an unrecoverable error. Consequently, $E(H_{\rm UF})$ is independent of P_s, as indicated by the horizontal part of the dotted green line. Let us now consider Figure 8(a). When $P_s \gg 10^{-10}$, $E(H_{\rm UF})$ increases, because there are multiple such codewords, each of which loses $\tilde{r}$ symbols. Subsequently, for $\tilde{r} \geq 3$ and when $P_s \gg 10^{-8}$, unrecoverable failures may also be caused by a single corrupted codeword that loses $\tilde{r}$ of its symbols, $\tilde{r} - 2$ of which are lost owing to device failures and two are lost owing to unrecoverable errors. This in turn reduces the amount of lost data in the interval $(10^{-10}, 10^{-8})$, as shown in Figure 8(a). Note that this interval corresponds to that of the second plateau, as shown in Figure 4(a). When $P_s \gg 10^{-6}$, $E(H_{\rm UF})$ increases again owing to the occurrence of multiple such corrupted codewords. Eventually, when $P_s \gg 10^{-5}$, unrecoverable failures are encountered during rebuild prior to a second device failure and are caused by corrupted codewords that lose $\tilde{r}$ of their symbols, one of which is lost owing to the first device failure and $\tilde{r} - 1$ are lost owing to unrecoverable errors. This in turn increases $E(H_{\rm UF})$, which eventually dominates $E(H_{\rm DF})$. Similar observations apply in the cases of Figures 8(b) and 8(c).

The reliability metrics corresponding to the clustered placement scheme ($k = m = 16$) are plotted in Figures 9, 10, 11, 12, and 13. We observe that the reliability achieved by the clustered data placement scheme does not reach the reliability level achieved by the declustered one.

In the cases considered, EAFDL is practically unaffected by the presence of latent errors, as shown in Figures 7 and 12. Note, however, that for larger values of $\tilde{r}$, EAFDL may be affected by the presence of latent errors. For example, when $m = 24$ and $l = 12$, which yields $\tilde{r} = 13$, and for the case of declustered placement scheme, not only MTTDL, but also EAFDL is affected, as shown in Figure 14.

The performance of certain erasure coding schemes was assessed in [32] by obtaining the probability of data loss $P_{\rm DL}$ using a detailed distributed storage simulator. The $P_{\rm DL}$ values corresponding to $P_s = 4.096 \times 10^{-12}$ ($P_{\rm bit} = 10^{-15}$) for two of the configurations considered are indicated by the squares in Figure 15. This figure also shows the probabilities of data loss $P_{\rm DL}$ that correspond to these two configurations and obtained from (22) as a function of P_s. We observe that the theoretical results are in agreement with the simulation results, which confirms the validity of the model and the analytical expressions derived.

VI. Real-World Erasure Coding Schemes

Here we assess the reliability of various practical systems that store an amount of $U = 1.2$ PB user data on devices (disks) whose capacity is $c = 20$ TB. This amount of user data can therefore be stored on $U/c = 60$ devices. The system comprises n devices, where n is determined using (1) as follows:

$$n = \frac{U}{c}\,\frac{m}{l} = 60\,\frac{m}{l}\,. \tag{65}$$

Subsequently, we consider the following real-world erasure coding schemes:

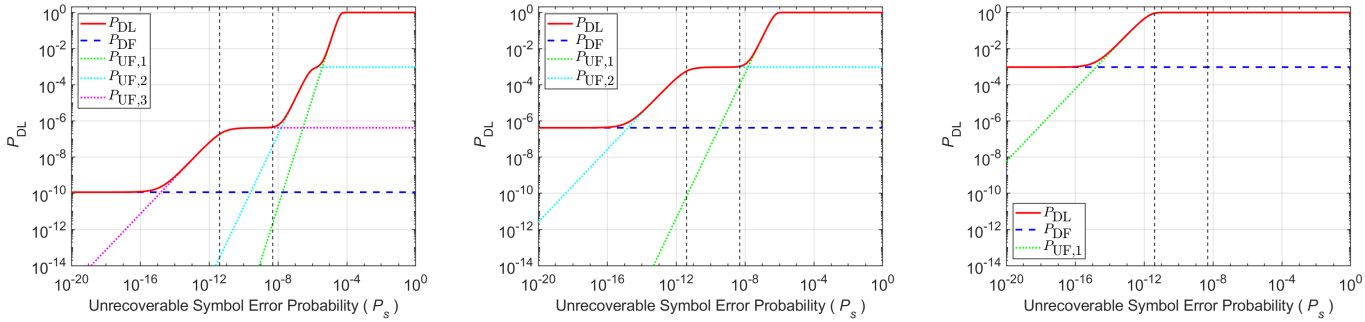

(a) $l = 13 \quad (\tilde{r} = 4)$ (b) $l = 14 \quad (\tilde{r} = 3), \quad$ RAID-6 (c) $l = 15 \quad (\tilde{r} = 2), \quad$ RAID-5

Figure 9. Probability of data loss P_{DL} vs. P_s for $l = 13, 14, 15$; $n = 64$, $k = m = 16$ (clustered scheme).

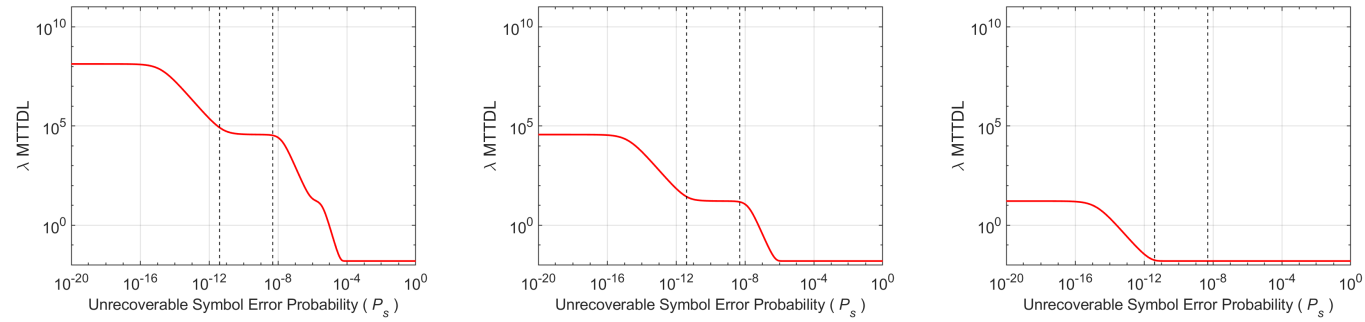

(a) $l = 13 \quad (\tilde{r} = 4)$ (b) $l = 14 \quad (\tilde{r} = 3), \quad$ RAID-6 (c) $l = 15 \quad (\tilde{r} = 2), \quad$ RAID-5

Figure 10. Normalized MTTDL vs. P_s for $l = 13, 14,$ and 15; $n = 64$, $k = m = 16$ (clustered scheme).

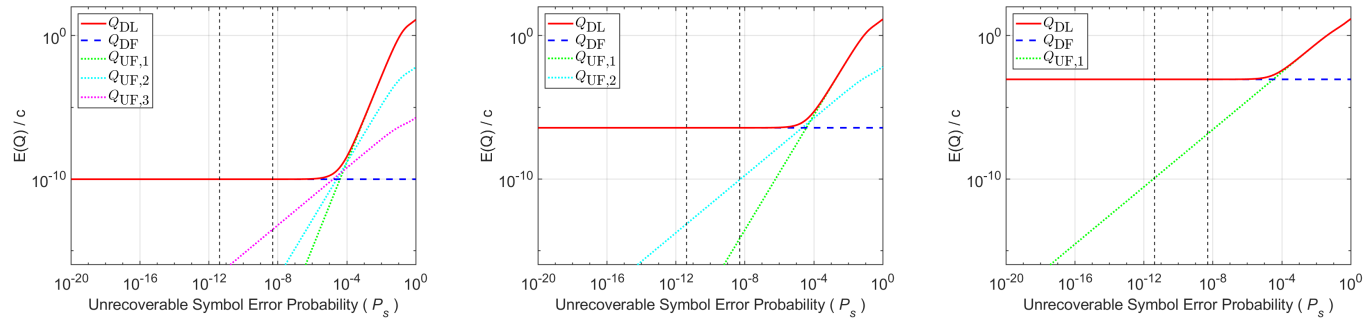

(a) $l = 13 \quad (\tilde{r} = 4)$ (b) $l = 14 \quad (\tilde{r} = 3), \quad$ RAID-6 (c) $l = 15 \quad (\tilde{r} = 2), \quad$ RAID-5

Figure 11. Normalized amount of data loss $E(Q)$ vs. P_s for $l = 13, 14, 15$; $n = 64$, $k = m = 16$ (clustered scheme).

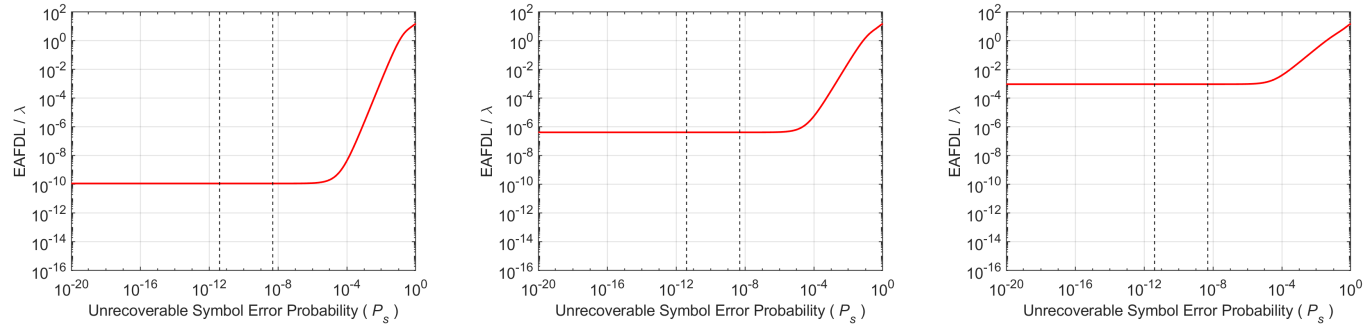

(a) $l = 13 \quad (\tilde{r} = 4)$ (b) $l = 14 \quad (\tilde{r} = 3), \quad$ RAID-6 (c) $l = 15 \quad (\tilde{r} = 2), \quad$ RAID-5

Figure 12. Normalized EAFDL vs. P_s for $l = 13, 14,$ and 15; $n = 64$, $k = m = 16$ (clustered scheme).

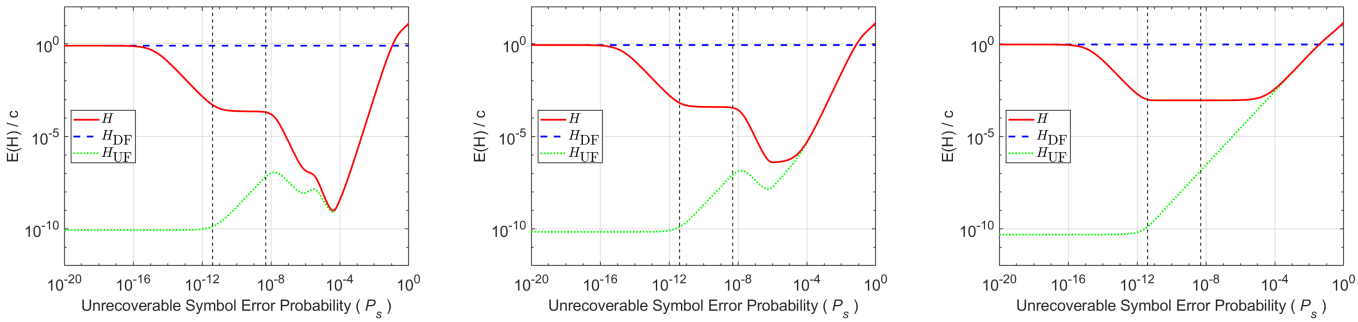

(a) $l = 13$ $(\tilde{r} = 4)$ (b) $l = 14$ $(\tilde{r} = 3)$, RAID-6 (c) $l = 15$ $(\tilde{r} = 2)$, RAID-5

Figure 13. Normalized $E(H)$ vs. P_s for $l = 13, 14,$ and 15; $n = 64$, $k = m = 16$ (clustered scheme).

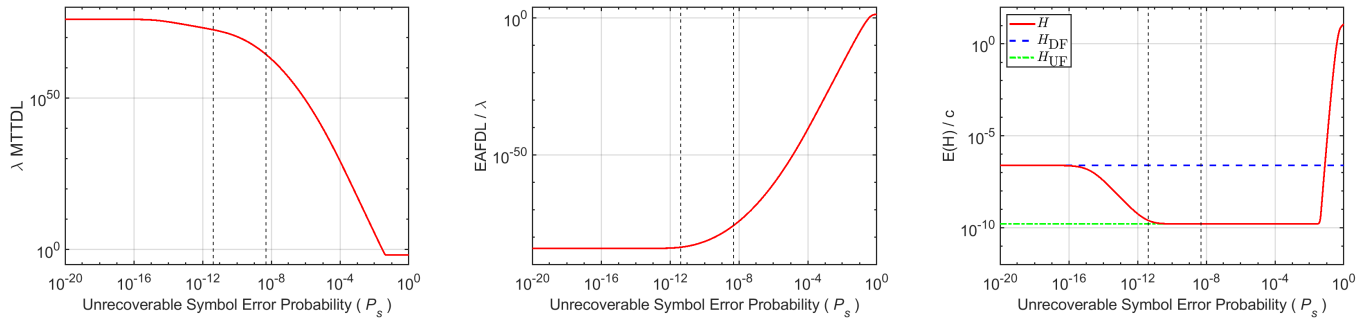

(a) Normalized MTTDL (b) Normalized EAFDL (c) Normalized $E(H)$

Figure 14. Reliability metrics vs. P_s for $n = k = 64$ (declustered scheme), $m = 24$, $l = 12$.

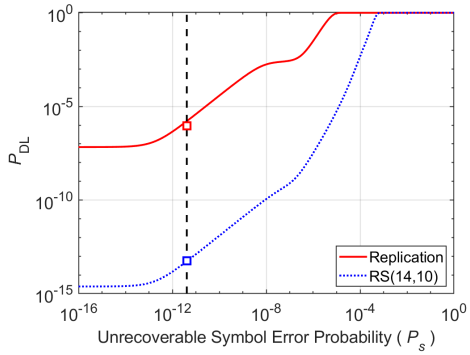

Figure 15. P_{DL} vs. P_s for $n = k = 210$, $\lambda^{-1} = 3$ years, $\mu^{-1} = 34$ hours, $\lambda/\mu = 0.0013$, $c = 15$ TB, and $s = 512$ B. Reliability schemes: 3-way replication and MDS(14,10).

1) the 3-way replication (triplication) scheme that was initially used by Google's GFS, Microsoft® Azure, and Facebook. In this case, $m = 3$, $l = 1$, with a corresponding storage efficiency of $s_{\text{eff}} = 33\%$. According to (65), this scheme requires the employment of $n = 180$ devices.
2) the RS(9,6) erasure coding scheme employed by Google's GFS as well as QFS [24, 45], which for $m = 9$ and $l = 6$ achieves a storage efficiency of $s_{\text{eff}} = 66\%$ and requires a number of $n = 90$ devices.
3) the MDS(16,12) erasure coding scheme akin to the LRC(16,12) code used by Microsoft® Azure [26], which for $m = 16$ and $l = 12$ achieves a storage efficiency of $s_{\text{eff}} = 75\%$ and requires a number of $n = 80$ devices.
4) the RS(14,10) erasure coding scheme employed by Facebook [25], which for $m = 14$ and $l = 10$ achieves a storage efficiency of $s_{\text{eff}} = 71\%$ and requires a number of $n = 84$ devices.

We proceed to assess the reliability of the four erasure coding schemes for two data placement configurations: a symmetric one where the system comprises 2 disjoint groups of k devices, such that $k = n/2$, and a declustered one, such that $k = n$. As we will see next, a superior reliability is achieved by employing the declustered data placement scheme.

A. *Symmetric Data Placement*

First, we assess the reliability of the 3-way replication (triplication) scheme that requires the employment of $n = 180$ devices, which in turn implies that each of the two groups comprises $k = 90$ devices. The reliability measures are obtained for the parameter values listed in Table IV and shown in Figures 16(a) through 20(a). We observe that MTTDL is significantly degraded by the presence of latent errors. In the interval $[4.096 \times 10^{-12}, 5 \times 10^{-9}]$ of practical importance for P_s, which is indicated between the two vertical dashed lines, Figure 17(a) shows that MTTDL is degraded by three to six orders of magnitude, whereas Figure 19(a) reveals that EAFDL is practically unaffected in this range.

Second, we consider the MDS(9,6) erasure coding scheme that requires a number of $n = 90$ devices, which in turn implies that each of the two groups comprises $k = 45$ devices. The corresponding reliability measures are shown in Figures

16(b) through 20(b). Figure 17(b) shows that, in the region of interest for P_s, MTTDL is degraded by three to five orders of magnitude, whereas Figure 19(a) reveals that EAFDL is practically unaffected in this range.

Subsequently, we consider the MDS(16,12) erasure coding scheme that requires a number of $n = 80$ devices, which in turn implies that each of the two groups comprises $k = 40$ devices. The corresponding reliability measures are shown in Figures 16(c) through 20(c). Figure 17(c) shows that, in the interval of practical importance for P_s, MTTDL is degraded by three to five orders of magnitude, whereas Figure 19(c) reveals that EAFDL is practically unaffected in this range.

Finally, we consider the RS(14,10) erasure coding scheme that requires $n = 84$ devices, which in turn implies that each of the two groups comprises $k = 42$ devices. The corresponding reliability measures are shown in Figures 16(d) through 20(d). Figure 17(d) shows that, in the interval of interest, MTTDL is degraded by three to five orders of magnitude, whereas Figure 19(d) reveals that EAFDL is practically unaffected in this range.

From the above, it follows that erasure coding schemes corresponding to higher values of $\tilde{r}$ offer a higher level of reliability. Thus, the MDS(16,12) and MDS(14,10) erasure coding schemes, for which $\tilde{r} = 5$, offer higher levels of reliability compared with the MDS(9,6) and 3-way replication schemes, for which $\tilde{r} < 5$. In particular, MDS(14,10) achieves a higher reliability than that of MDS(16,10), albeit at a lower storage efficiency (71% vs. 75%).

B. Declustered Data Placement

Here, we assess the reliability achieved by the erasure coding schemes considered when the declustered data placement scheme is used, such that $k = n$. The reliability results are shown in Figures 21(a) through 25(a).

First, we assess the reliability of the 3-way replication (triplication) scheme. Comparing Figure 22(a) with Figure 17(a), we deduce that MTTDL is roughly the same. However, comparing Figure 24(a) with Figure 19(a), we deduce that EAFDL improves by one order of magnitude.

Regarding, the reliability of the MDS(9,6) coding scheme, comparing Figure 22(b) with Figure 17(b), we deduce that MTTDL improves slightly by one order of magnitude, especially at smaller values of P_s. However, Figures 24(b) and 19(b) demonstrate that EAFDL improves by two orders of magnitude.

For the MDS(16,12) coding scheme, Figures 22(c) and 17(c) show that MTTDL improves by two orders of magnitude for values around the left vertical dotted line and by one order of magnitude for values around the right vertical dotted line. Also, from Figures 24(c) and 19(c), we observe that EAFDL improves by three orders of magnitude.

Finally, the reliability improvement regarding the MDS(14,10) coding scheme is similar to that of the MDS(16,12) coding scheme. Figures 22(d) and 17(d) show that MTTDL improves by two orders of magnitude for values around the left vertical dotted line and by one order of magnitude for values around the right vertical dotted line. Also, Figures 24(d) and 19(d) show that EAFDL improves by three orders of magnitude.

C. Reliability Improvement

The reliability improvement of the erasure coding schemes considered over the initial 3-way replication is shown in Figures 26 and 27 for the two data placements, respectively. Clearly, in the interval of practical importance for P_s, the MDS(14,10) erasure coding scheme achieves superior reliability for both symmetric and declustered data placement schemes.

In particular, for the symmetric data placement, Figure 26(a) demonstrates that in the interval of interest, the MDS(9,6) erasure coding scheme improves MTTDL by three orders of magnitude for P_s values around the left vertical dotted line and by four orders of magnitude for P_s values around the right vertical dotted line. The MDS(16,12) erasure coding scheme improves MTTDL by six orders of magnitude for P_s values around the left vertical dotted line and by seven orders of magnitude for P_s values around the right vertical dotted line. Also, the MDS(14,10) erasure coding scheme improves MTTDL by seven orders of magnitude for P_s values around the left vertical dotted line and by eight orders of magnitude for P_s values around the right vertical dotted line. On the other hand, Figure 26(b) demonstrates that in the interval of practical importance for P_s, the MDS(9,6), MDS(16,12), and MDS(14,10) erasure coding schemes improve EAFDL by two, five, and six orders of magnitude, respectively.

For the declustered data placement, Figure 27(a) demonstrates that in the interval of interest, the MDS(9,6) erasure coding scheme improves MTTDL by four orders of magnitude, the MDS(16,12) erasure coding scheme improves MTTDL by eight orders of magnitude, whereas the MDS(14,10) erasure coding scheme improves MTTDL by nine orders of magnitude. On the other hand, Figure 27(b) demonstrates that in the interval of practical importance for P_s, the MDS(9,6), MDS(16,12), and MDS(14,10) erasure coding schemes improve EAFDL by three, seven, and eight orders of magnitude, respectively.

Figures 26(c) and 27(c) show the ratios of the $E(H)$ metrics for the MDS(9,6), MDS(16,12), and MDS(14,10) erasure coding schemes to the $E(H)$ metric for the 3-way replication scheme. In the region of interest for P_s, all three erasure coding schemes result in greater amounts of lost user data, given that data loss has occurred, compared to the conditional amount of lost user data in the case of a 3-way replication. In particular, for the symmetric data placement, the MDS(9,6), MDS(16,12), and MDS(14,10) erasure coding schemes result in about 20, 82, and 33 times greater conditional amounts of lost user data for P_s values around the left vertical dotted line and in about 58, 550, and 110 times greater conditional amounts of lost user data for P_s values around the right vertical dotted line, respectively. This is due to the fact that, for small values of P_s and according to Remark 3, $E(H)$ depends not only on the values of the m and l parameters, but also on the number n of devices in the system, which also varies. Accordingly, for the declustered data placement, Figure 27(c) demonstrates that the MDS(9,6), MDS(16,12), and MDS(14,10) erasure coding schemes result in about 9, 18, and 7 times greater conditional amounts of lost user data, respectively.

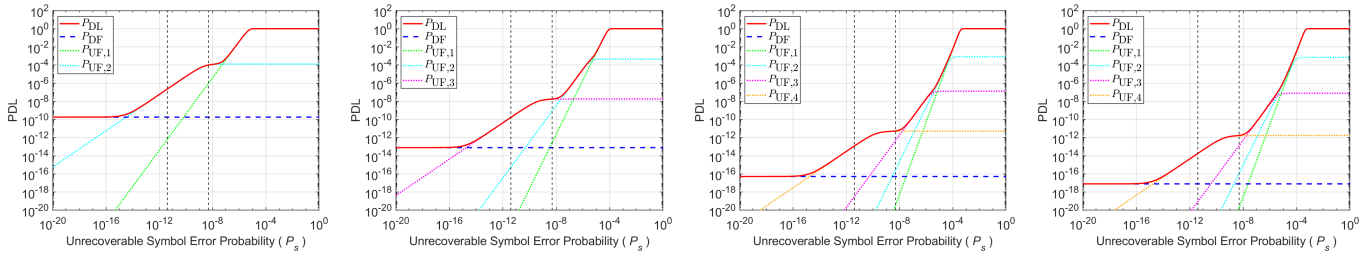

(a) MDS(3,1) ($\tilde{r}$=3) (b) MDS(9,6) ($\tilde{r}$=4) (c) MDS(16,12) ($\tilde{r}$=5) (d) MDS(14,10) ($\tilde{r}$=5)

Figure 16. Probability of data loss P_{DL} vs. P_s for various MDS coding schemes; symmetric data placement with $k = n/2$.

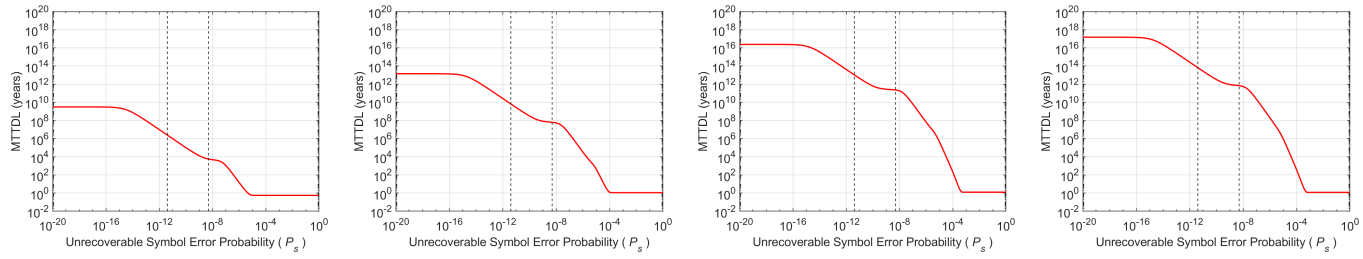

(a) MDS(3,1) (b) MDS(9,6) (c) MDS(16,12) (d) MDS(14,10)

Figure 17. MTTDL vs. P_s for various MDS coding schemes; symmetric data placement with $k = n/2$.

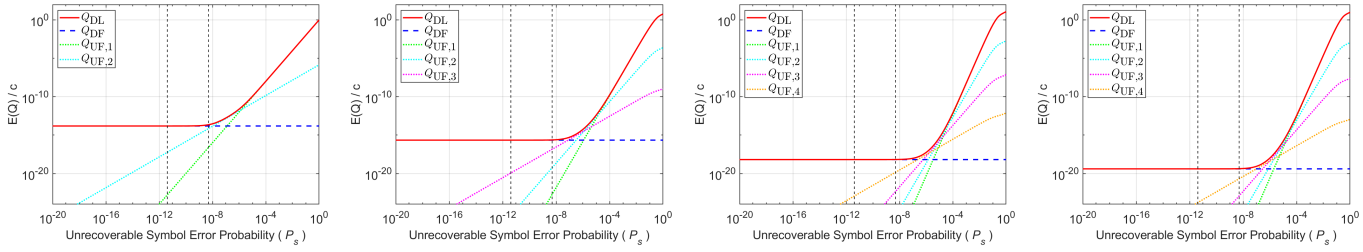

(a) MDS(3,1) (b) MDS(9,6) (c) MDS(16,12) (d) MDS(14,10)

Figure 18. Normalized amount of data loss $E(Q)$ vs. P_s for various MDS coding schemes; symmetric data placement with $k = n/2$.

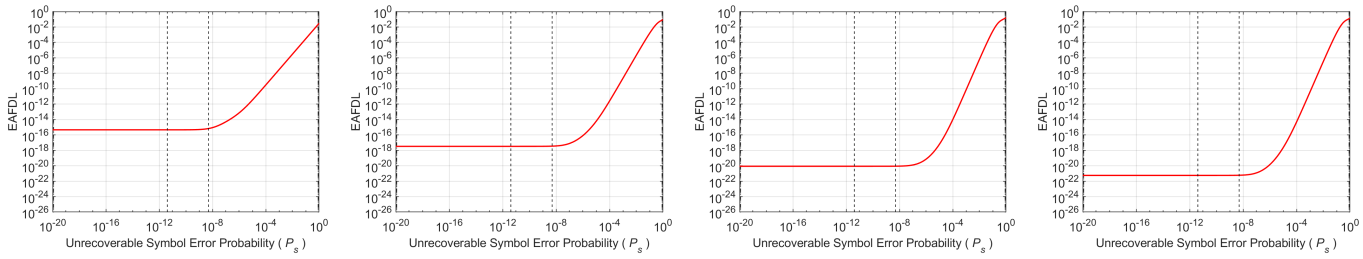

(a) MDS(3,1) (b) MDS(9,6) (c) MDS(16,12) (d) MDS(14,10)

Figure 19. EAFDL vs. P_s for various MDS coding schemes; symmetric data placement with $k = n/2$.

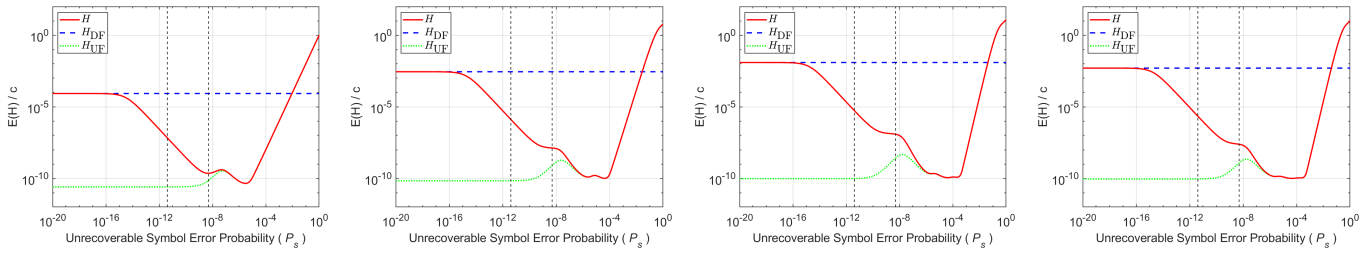

(a) MDS(3,1) (b) MDS(9,6) (c) MDS(16,12) (d) MDS(14,10)

Figure 20. Normalized $E(H)$ vs. P_s for various MDS coding schemes; symmetric data placement with $k = n/2$.

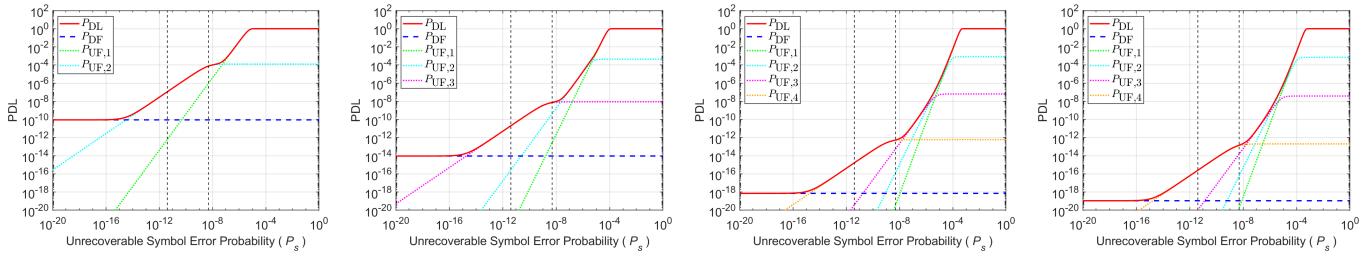

(a) MDS(3,1) ($\tilde{r}$=3) (b) MDS(9,6) ($\tilde{r}$=4) (c) MDS(16,12) ($\tilde{r}$=5) (d) MDS(14,10) ($\tilde{r}$=5)

Figure 21. Probability of data loss P_{DL} vs. P_s for various MDS coding schemes; declustered data placement ($k = n$).

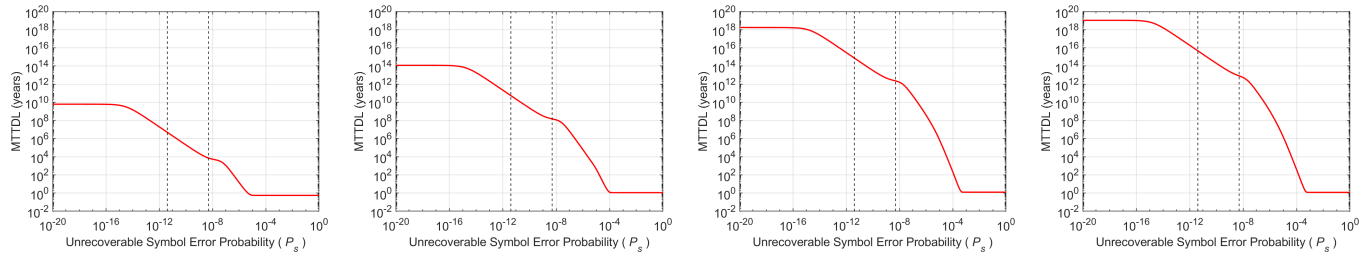

(a) MDS(3,1) (b) MDS(9,6) (c) MDS(16,12) (d) MDS(14,10)

Figure 22. MTTDL vs. P_s for various MDS coding schemes; declustered data placement ($k = n$).

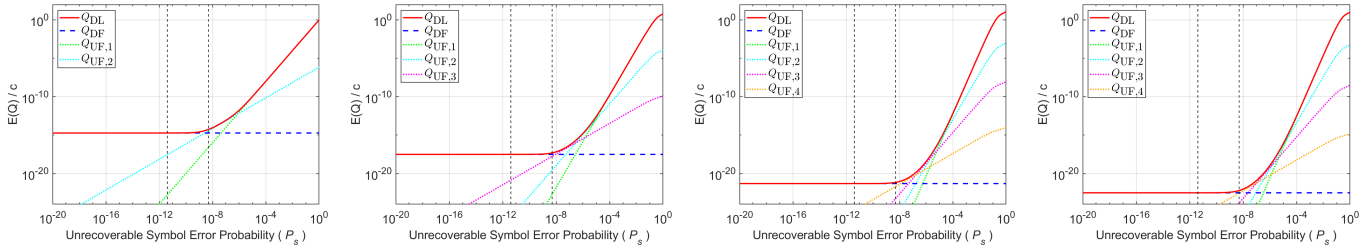

(a) MDS(3,1) (b) MDS(9,6) (c) MDS(16,12) (d) MDS(14,10)

Figure 23. Normalized amount of data loss $E(Q)$ vs. P_s for various MDS coding schemes; declustered data placement ($k = n$).

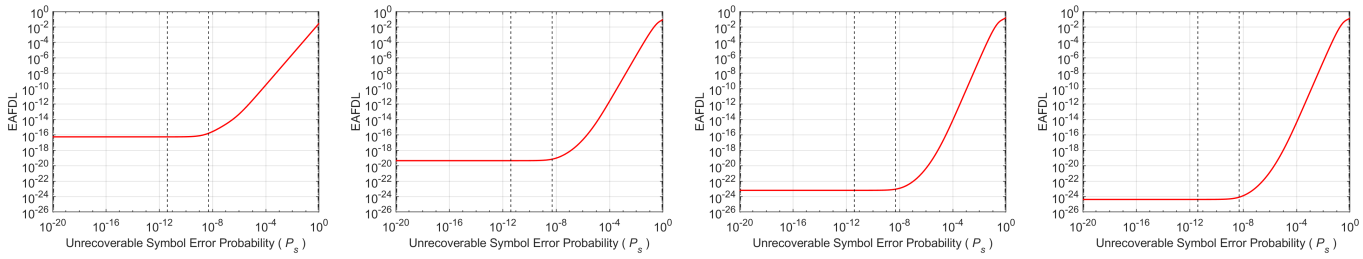

(a) MDS(3,1) (b) MDS(9,6) (c) MDS(16,12) (d) MDS(14,10)

Figure 24. EAFDL vs. P_s for various MDS coding schemes; declustered data placement ($k = n$).

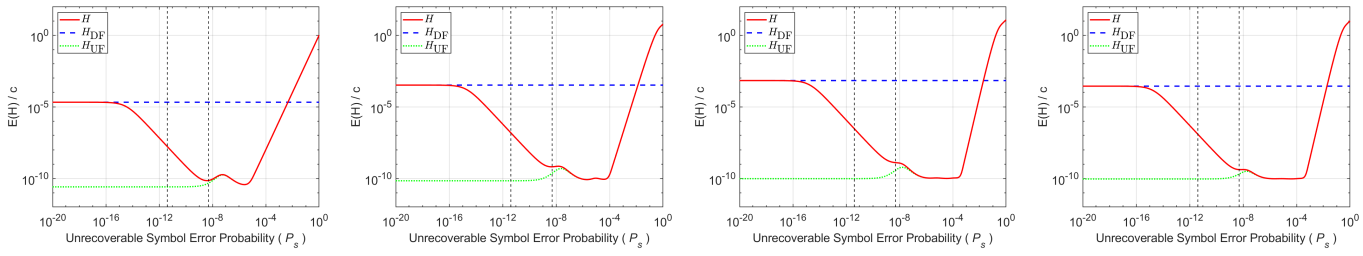

(a) MDS(3,1) (b) MDS(9,6) (c) MDS(16,12) (d) MDS(14,10)

Figure 25. Normalized $E(H)$ vs. P_s for various MDS coding schemes; declustered data placement ($k = n$).

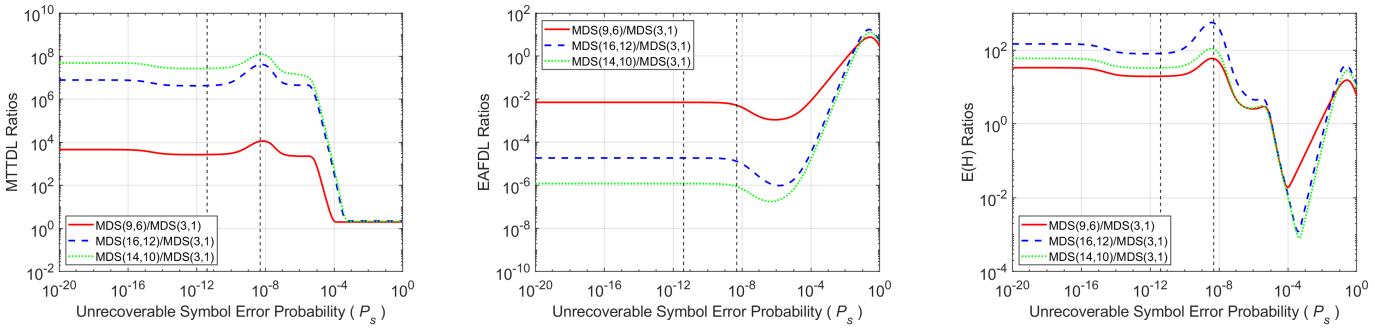

(a) MTTDL ratios (b) EAFDL ratios (c) $E(H)$ ratios

Figure 26. Ratios of the MTTDL, EAFDL, and $E(H)$ metrics for the MDS(9,6), MDS(16,12), and MDS(14,10) schemes to those corresponding to the 3-way replication scheme; symmetric data placement with $k = n/2$.

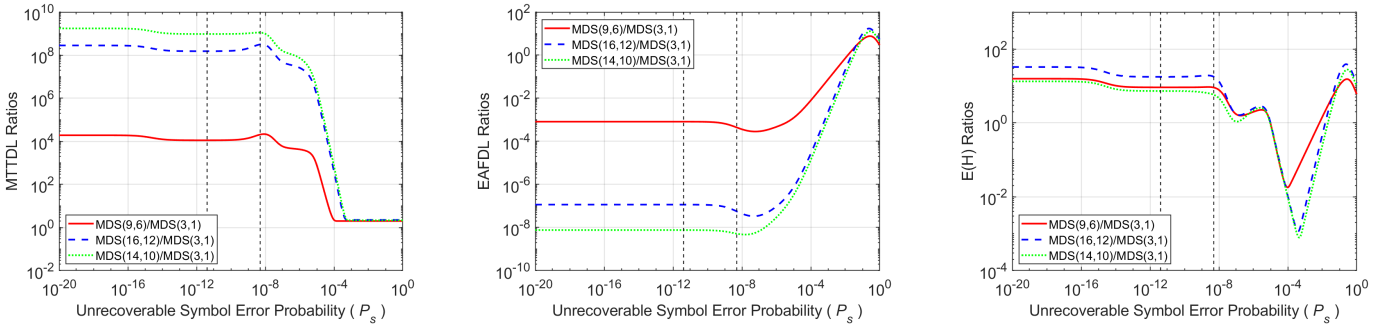

(a) MTTDL ratios (b) EAFDL ratios (c) $E(H)$ ratios

Figure 27. Ratios of the MTTDL, EAFDL, and $E(H)$ metrics for the MDS(9,6), MDS(16,12), and MDS(14,10) schemes to those corresponding to the 3-way replication scheme; declustered data placement ($k = n$).

VII. Conclusions

The effect of latent sector errors on the reliability of erasure-coded data storage systems was investigated. A methodology was developed for deriving the Mean Time to Data Loss (MTTDL) and the Expected Annual Fraction of Data Loss (EAFDL) reliability metrics analytically. Closed-form expressions capturing the effect of unrecoverable latent errors were obtained for the symmetric, clustered and declustered data placement schemes. We demonstrated that the declustered placement scheme offers superior reliability in terms of both metrics. We established that, for realistic unrecoverable sector error rates, MTTDL is adversely affected by the presence of latent errors, whereas EAFDL is not. We considered several real-world erasure coding schemes and demonstrated their efficiency. The analytical reliability expressions derived enable the identification of storage-efficient data placement configurations that yield high reliability.

Applying these results to assess the effect of network rebuild bandwidth constraints is a subject of further investigation. The reliability evaluation of erasure-coded systems when device failures, as well as unrecoverable latent errors are correlated is also part of future work.

Appendix A

We consider the direct path $\overrightarrow{UF_u} = 1 \to 2 \to \cdots \to u \to \text{UF}$ and proceed to evaluate $P_{\text{UF}_u}(R_1, \vec{\alpha}_{u-1})$, the probability of entering exposure level u through vector $\vec{\alpha}_{u-1} \triangleq (\alpha_1, \ldots, \alpha_{u-1})$ and given a rebuild time R_1, and then encountering an unrecoverable failure during the rebuild process at this exposure level. It follows from (20) that

$$P_{\text{UF}_u}(R_1, \vec{\alpha}_{u-1}) = P_u(R_1, \vec{\alpha}_{u-2}) \cdot P_{u\to\text{UF}}(R_1, \vec{\alpha}_{u-1}) \,. \quad (66)$$

It follows from Eq.(111) of [13] by setting $\tilde{r} = u$ that

$$P_u(R_1, \vec{\alpha}_{u-2}) \approx (\lambda b_1 R_1)^{u-1} \prod_{i=1}^{u-1} \frac{\tilde{n}_i}{b_i} \, (V_i \, \alpha_i)^{u-1-i} \,. \quad (67)$$

Given that the elements of $\vec{\alpha}_{u-2}$ are independent random variables approximately distributed according to (9), such that $E(\alpha_i^k) \approx 1/(k+1)$, we have

$$E\left(\prod_{i=1}^{u-1} \alpha_i^{u-1-i}\right) = \prod_{i=1}^{u-1} E(\alpha_i^{u-1-i}) \approx \prod_{i=1}^{u-1} \frac{1}{u-i} = \frac{1}{(u-1)!}. \quad (68)$$

Unconditioning (67) on $\vec{\alpha}_{u-2}$ using (68) yields

$$P_u(R_1) \approx (\lambda b_1 R_1)^{u-1} \, \frac{1}{(u-1)!} \prod_{i=1}^{u-1} \frac{\tilde{n}_i}{b_i} \, V_i^{u-1-i} \,. \quad (69)$$

Unconditioning (69) on R_1 and using (6) and (8) yields (21).

We now proceed to calculate $P_{u\to\text{UF}}(R_1, \vec{\alpha}_{u-1})$. Upon entering exposure level u, the rebuild process attempts to restore the C_u most-exposed codewords, each of which has $m-u$ remaining symbols. Let us consider such a codeword,

and let L_u be the number of symbols permanently lost and I_u be the number of symbols in the codeword with unrecoverable errors. Owing to the independence of symbol errors, I_u follows a binomial distribution with parameter P_s, the probability that a symbol has a unrecoverable error. Thus, for $i = 0, \ldots, m-u$,

$$P(I_u = j) = \binom{m-u}{j} P_s^j \, (1 - P_s)^{m-u-j} \,, \tag{70}$$

$$\approx \binom{m-u}{j} P_s^j \,, \quad \text{for } P_s \ll \frac{1}{m-u-j} \,, \tag{71}$$

such that

$$E(I_u) = \sum_{j=1}^{m-u} j \, P(I_u = j) \; = \; (m-u) \, P_s \,. \tag{72}$$

Clearly, the symbols lost due to the device failures can be corrected by the erasure coding capability only if at least l of the remaining $m-u$ symbols can be read. Thus, $L_u = 0$ if and only if $I_u \leq m-u-l$ or, by virtue of (3), $I_u \leq \tilde{r}-1-u$. Thus, the probability q_u that a codeword can be restored is

$$q_u = P(L_u = 0) = 1 - P(I_u > \tilde{r} - u) \,, \tag{73}$$

which, using (70), yields (25).

Note that if a codeword is corrupted, then at least one of its l user-data symbols is lost. Owing to the independence of symbol errors, codewords are independently corrupted. Consequently, the conditional probability $P_{\text{UF}|C_u}$ of encountering an unrecoverable failure during the rebuild process of the C_u codewords is

$$P_{\text{UF}|C_u} = 1 - q_u^{C_u} \,, \quad \text{for } u = 1, \ldots, \tilde{r} \,. \tag{74}$$

Substituting (11) into (74) and using (24) yields

$$P_{u \to \text{UF}}(R_1, \vec{\alpha}_{u-1}) \approx 1 - q_u^{C \, \prod_{j=1}^{u-1} V_j \, \alpha_j} = 1 - \hat{q}_u^{\prod_{j=1}^{u-1} \alpha_j} \,. \tag{75}$$

Substituting (75) into (66) yields

$$P_{\text{UF}_u}(R_1, \vec{\alpha}_{u-1}) \; \approx \; P_u(R_1, \vec{\alpha}_{u-2}) \left[1 - \hat{q}_u^{\prod_{j=1}^{u-1} \alpha_j} \right] . \tag{76}$$

Unconditioning (76) on $\vec{\alpha}_{u-1}$ and using (67) yields

$$P_{\text{UF}_u}(R_1) \; \approx P_u(R_1) - (\lambda b_1 R_1)^{u-1} \left(\prod_{i=1}^{u-1} \frac{\tilde{n}_i}{b_i} \, V_i^{u-1-i} \right) \cdot E_{\vec{\alpha}_{u-1}} \left[\left(\prod_{i=1}^{u-1} \alpha_i^{u-1-i} \right) \hat{q}_u^{\prod_{j=1}^{u-1} \alpha_j} \right] . \tag{77}$$

LEMMA 1: For $\alpha_i \sim U(0,1)$ and for all $q \in \mathbb{R}$, it holds that

$$E \left[\left(\prod_{i=1}^{u-1} \alpha_i^{u-1-i} \right) q^{\prod_{i=1}^{u-1} \alpha_i} \right] = \frac{1}{(u-1)!} + \log(q)^{-(u-1)} \left(q - \sum_{i=0}^{u-1} \frac{\log(q)^i}{i!} \right) . \tag{78}$$

Proof: It holds that

$$q^{\prod_{i=1}^{u-1} \alpha_i} = e^{\log(q) \, \prod_{i=1}^{u-1} \alpha_i} = \sum_{j=0}^{\infty} \frac{\log(q)^j \, (\prod_{i=1}^{u-1} \alpha_i)^j}{j!} \,, \tag{79}$$

which implies that

$$\left(\prod_{i=1}^{u-1} \alpha_i^{u-1-i} \right) q^{\prod_{i=1}^{u-1} \alpha_i} = \left(\prod_{i=1}^{u-1} \alpha_i^{u-1-i} \right) \left(\sum_{j=0}^{\infty} \frac{\log(q)^j \, (\prod_{i=1}^{u-1} \alpha_i)^j}{j!} \right) = \sum_{j=0}^{\infty} \frac{\log(q)^j \, \prod_{i=1}^{u-1} \alpha_i^{u-1-i+j}}{j!} \,. \tag{80}$$

Consequently,

$$\begin{aligned} & E \left[\left(\prod_{i=1}^{u-1} \alpha_i^{u-1-i} \right) q^{\prod_{i=1}^{u-1} \alpha_i} \right] \\ &= \sum_{j=0}^{\infty} \frac{\log(q)^j \, \prod_{i=1}^{u-1} E(\alpha_i^{u-1-i+j})}{j!} \\ &\approx \sum_{j=0}^{\infty} \frac{\log(q)^j \, \prod_{i=1}^{u-1} \frac{1}{u-i+j}}{j!} \\ &= \sum_{j=0}^{\infty} \frac{\log(q)^j}{(u-1+j)!} = \frac{1}{(u-1)!} + \sum_{j=1}^{\infty} \frac{\log(q)^j}{(u-1+j)!} \\ &= \frac{1}{(u-1)!} + \log(q)^{-(u-1)} \sum_{i=u}^{\infty} \frac{\log(q)^i}{i!} \\ &= \frac{1}{(u-1)!} + \log(q)^{-(u-1)} \left(\sum_{i=0}^{\infty} \frac{\log(q)^i}{i!} - \sum_{i=0}^{u-1} \frac{\log(q)^i}{i!} \right) \\ &= \frac{1}{(u-1)!} + \log(q)^{-(u-1)} \left(e^{\log(q)} - \sum_{i=0}^{u-1} \frac{\log(q)^i}{i!} \right) \end{aligned} \tag{81}$$

■

From (69) and (78), (77) yields

$$P_{\text{UF}_u}(R_1) \; \approx \; -(\lambda b_1 R_1)^{u-1} \left(\prod_{i=1}^{u-1} \frac{\tilde{n}_i}{b_i} \, V_i^{u-1-i} \right) \cdot \log(\hat{q}_u)^{-(u-1)} \left(\hat{q}_u - \sum_{i=0}^{u-1} \frac{\log(\hat{q}_u)^i}{i!} \right) . \tag{82}$$

Unconditioning (82) on R_1, and using (6) and (8), yields (23).

APPENDIX B

At exposure level u, when $I_u \geq m-u-l+1 = \tilde{r}-u$, the number L_u of lost symbols is $I_u + u$. Consequently, the expected number $E(L_u)$ of lost symbols is

$$E(L_u) = \sum_{i=\tilde{r}-u}^{m-u} (i+u) \, P(I_u = i) \,, \tag{83}$$

where $P(I_u = i)$ is given by (70). Considering approximation (71), it follows that

$$E(L_u) \; \approx \; \tilde{r} \binom{m-u}{\tilde{r}-u} P_s^{\tilde{r}-u} \,, \quad \text{for } P_s \ll \frac{1}{m-\tilde{r}} \,. \tag{84}$$

The expected number $E(S_{\rm U}|C_u)$ of symbols lost due to unrecoverable failures during the rebuild of the C_u codewords at exposure level u is equal to $C_u\,E(L_u)$, which yields

$$E(S_{\rm U}|C_u) \overset{(84)}{\approx} C_u\,\tilde{r}\binom{m-u}{\tilde{r}-u}P_s^{\tilde{r}-u},\;\; P_s \ll \frac{1}{m-\tilde{r}}\,. \quad (85)$$

Substituting (11) into (85) yields

$$E(S_{\rm U}|\vec{\alpha}_{u-1}) \;\approx\; C\left(\prod_{j=1}^{u-1} V_j\,\alpha_j\right)\tilde{r}\binom{m-u}{\tilde{r}-u}P_s^{\tilde{r}-u}\,. \quad (86)$$

Subsequently, the expected number $E(S_{{\rm UF}_u}|R_1,\vec{\alpha}_{u-1})$ of symbols lost due to unrecoverable failures encountered during rebuild in conjunction with entering exposure level u through vector $\vec{\alpha}_{u-1}$, and given a rebuild time R_1, is determined as follows:

$$E(S_{{\rm UF}_u}|R_1,\vec{\alpha}_{u-1}) = P_u(R_1,\vec{\alpha}_{u-1})\,E(S_{\rm U}|\vec{\alpha}_{u-1})\,. \quad (87)$$

Substituting (67) and (86) into (87) yields

$$E(S_{{\rm UF}_u}|R_1,\vec{\alpha}_{u-1}) \approx\; (\lambda b_1 R_1)^{u-1}\left[\prod_{i=1}^{u-1}\frac{\tilde{n}_i}{b_i}\,(V_i\,\alpha_i)^{u-i}\right] \cdot\; C\,\tilde{r}\binom{m-u}{\tilde{r}-u}P_s^{\tilde{r}-u}, \quad P_s \ll \frac{1}{m-\tilde{r}}\,. \quad (88)$$

From (68), we have that $E(\prod_{i=1}^{u-1}\alpha_i^{u-i}) = E(\prod_{i=1}^{u}\alpha_i^{u-i}) \approx 1/u!$. Thus, unconditioning (88) on $\vec{\alpha}_{u-1}$ yields

$$E(S_{{\rm UF}_u}|R_1) \approx\; (\lambda b_1 R_1)^{u-1}\left(\prod_{i=1}^{u-1}\frac{\tilde{n}_i}{b_i}\,V_i^{u-i}\right)\frac{1}{u!}\,C\,\tilde{r} \cdot \binom{m-u}{\tilde{r}-u}P_s^{\tilde{r}-u}, \quad P_s \ll \frac{1}{m-\tilde{r}}\,. \quad (89)$$

Unconditioning (89) on R_1, and using (6) and (8), yields

$$E(S_{{\rm UF}_u}) \approx\; (\lambda\,c)^{u-1}\,\frac{E(X^{u-1})}{[E(X)]^{u-1}}\left(\prod_{i=1}^{u-1}\frac{\tilde{n}_i}{b_i}\,V_i^{u-i}\right)\frac{1}{u!}\,C\,\tilde{r} \cdot \binom{m-u}{\tilde{r}-u}P_s^{\tilde{r}-u}, \quad P_s \ll \frac{1}{m-\tilde{r}}\,. \quad (90)$$

Substituting (90) into (37) yields (38).

Remark 6: From (21), (47), (84), and (90), it follows that

$$E(S_{{\rm UF}_u}) \;\approx\; P_u\,E(C_u)\,E(L_u)\,. \quad (91)$$

Upon entering exposure level u, the expected number $E(S_{\rm U}|C_u)$ of symbols lost due to unrecoverable failures during the rebuild of the C_u codewords is equal to $C_u\,E(L_u)$, as determined by (85). Consequently, upon entering exposure level u, the expected number $E(S_{\rm U})$ of symbols lost due to unrecoverable failures during the rebuild of the most-exposed codewords is $E(C_u)\,E(L_u)$. Therefore, the expected number $E(S_{{\rm UF}_u})$ of symbols lost due to unrecoverable failures at exposure level u is obtained by also considering the probability P_u of entering exposure level u, as determined by (91).

Note that when entering exposure level $\tilde{r}$, for each of the $C_{\tilde{r}}$ most-exposed codewords there are $\tilde{r}$ symbols permanently lost. Therefore, the number of data symbols permanently lost is $C_{\tilde{r}}\,\tilde{r}$. Consequently,

$$E(S_{\rm DF}) \;\approx\; P_{\rm DF}\,E(C_{\tilde{r}})\,\tilde{r}\,. \quad (92)$$

Substituting (92) into (36), and using (44), yields (39).

REFERENCES

[1] I. Iliadis, "Reliability assessment of erasure-coded storage systems with latent errors," in Proceedings of the 14th International Conference on Communication Theory, Reliability, and Quality of Service (CTRQ), Apr. 2021, pp. 15–24.

[2] D. A. Patterson, G. Gibson, and R. H. Katz, "A case for redundant arrays of inexpensive disks (RAID)," in Proceedings of the ACM International Conference on Management of Data (SIGMOD), Jun. 1988, pp. 109–116.

[3] P. M. Chen, E. A. Lee, G. A. Gibson, R. H. Katz, and D. A. Patterson, "RAID: High-performance, reliable secondary storage," ACM Comput. Surv., vol. 26, no. 2, Jun. 1994, pp. 145–185.

[4] V. Venkatesan, I. Iliadis, C. Fragouli, and R. Urbanke, "Reliability of clustered vs. declustered replica placement in data storage systems," in Proceedings of the 19th Annual IEEE/ACM International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS), Jul. 2011, pp. 307–317.

[5] I. Iliadis, D. Sotnikov, P. Ta-Shma, and V. Venkatesan, "Reliability of geo-replicated cloud storage systems," in Proceedings of the 2014 IEEE 20th Pacific Rim International Symposium on Dependable Computing (PRDC), Nov. 2014, pp. 169–179.

[6] M. Malhotra and K. S. Trivedi, "Reliability analysis of redundant arrays of inexpensive disks," J. Parallel Distrib. Comput., vol. 17, no. 1, Jan. 1993, pp. 146–151.

[7] A. Thomasian and M. Blaum, "Higher reliability redundant disk arrays: Organization, operation, and coding," ACM Trans. Storage, vol. 5, no. 3, Nov. 2009, pp. 1–59.

[8] I. Iliadis, R. Haas, X.-Y. Hu, and E. Eleftheriou, "Disk scrubbing versus intradisk redundancy for RAID storage systems," ACM Trans. Storage, vol. 7, no. 2, Jul. 2011, pp. 1–42.

[9] V. Venkatesan, I. Iliadis, and R. Haas, "Reliability of data storage systems under network rebuild bandwidth constraints," in Proceedings of the 20th Annual IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS), Aug. 2012, pp. 189–197.

[10] J.-F. Pâris, T. J. E. Schwarz, A. Amer, and D. D. E. Long, "Highly reliable two-dimensional RAID arrays for archival storage," in Proceedings of the 31st IEEE International Performance Computing and Communications Conference (IPCCC), Dec. 2012, pp. 324–331.

[11] I. Iliadis and V. Venkatesan, "Most probable paths to data loss: An efficient method for reliability evaluation of data storage systems," Int'l J. Adv. Syst. Measur., vol. 8, no. 3&4, Dec. 2015, pp. 178–200.

[12] ——, "Expected annual fraction of data loss as a metric for data storage reliability," in Proceedings of the 22nd Annual IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS), Sep. 2014, pp. 375–384.

[13] ——, "Reliability evaluation of erasure coded systems," Int'l J. Adv. Telecommun., vol. 10, no. 3&4, Dec. 2017, pp. 118–144.

[14] I. Iliadis, "Reliability evaluation of erasure coded systems under rebuild bandwidth constraints," Int'l J. Adv. Networks and Services, vol. 11, no. 3&4, Dec. 2018, pp. 113–142.

[15] ——, "Data loss in RAID-5 storage systems with latent errors," in Proceedings of the 12th International Conference on Communication Theory, Reliability, and Quality of Service (CTRQ), Mar. 2019, pp. 1–9.

[16] ——, "Data loss in RAID-5 and RAID-6 storage systems with latent errors," Int'l J. Adv. Software, vol. 12, no. 3&4, Dec. 2019, pp. 259–287.

[17] Amazon Web Services, "Amazon Simple Storage Service (Amazon S3)," 2022. [Online]. Available: http://aws.amazon.com/s3/ [retrieved: December 7, 2022]

[18] D. Borthakur et al., “Apache Hadoop goes realtime at Facebook,” in Proceedings of the ACM International Conference on Management of Data (SIGMOD), Jun. 2011, pp. 1071–1080.

[19] R. J. Chansler, “Data availability and durability with the Hadoop Distributed File System,” ;login: The USENIX Association Newsletter, vol. 37, no. 1, Feb. 2012, pp. 16–22.

[20] K. Shvachko, H. Kuang, S. Radia, and R. Chansler, “The Hadoop Distributed File System,” in Proceedings of the 26th IEEE Symposium on Mass Storage Systems and Technologies (MSST), May 2010, pp. 1–10.

[21] S. Ghemawat, H. Gobioff, and S.-T. Leung, “The Google file system,” in Proceedings of the 19th ACM Symposium on Operating Systems Principles (SOSP), Oct. 2003, pp. 29–43.

[22] D. Borthakur. HDFS and Erasure Codes (HDFS-RAID), Aug. 2009. [Online]. Available: https://hadoopblog.blogspot.com/2009/08 [retrieved: December 7, 2022]

[23] B. Calder et al., “Windows Azure Storage: a highly available cloud storage service with strong consistency,” in Proceedings of the 23rd ACM Symposium on Operating Systems Principles (SOSP), Oct. 2011, pp. 143–157.

[24] D. Ford, F. Labelle, F. I. Popovici, M. Stokely, V.-A. Truong, L. Barroso, C. Grimes, and S. Quinlan, “Availability in globally distributed storage systems,” in Proceedings of the 9th USENIX Symposium on Operating Systems Design and Implementation (OSDI), Oct. 2010, pp. 61–74.

[25] S. Muralidhar et al., “f4: Facebook’s Warm BLOB Storage System,” in Proceedings of the 11th USENIX Symposium on Operating Systems Design and Implementation (OSDI), Oct. 2014, pp. 383–397.

[26] C. Huang, H. Simitci, Y. Xu, A. Ogus, B. Calder, P. Gopalan, J. Li, and S. Yekhanin, “Erasure coding in Windows Azure Storage,” in Proceedings of the USENIX Annual Technical Conference (ATC), Jun. 2012, pp. 15–26.

[27] E. Pinheiro, W.-D. Weber, and L. A. Barroso, “Failure trends in a large disk drive population,” in Proceedings of the 5th USENIX Conference on File and Storage Technologies (FAST), Feb. 2007, pp. 17–28.

[28] A. Dholakia, E. Eleftheriou, X.-Y. Hu, I. Iliadis, J. Menon, and K. Rao, “A new intra-disk redundancy scheme for high-reliability RAID storage systems in the presence of unrecoverable errors,” ACM Trans. Storage, vol. 4, no. 1, May 2008, pp. 1–42.

[29] I. Iliadis, “Reliability modeling of RAID storage systems with latent errors,” in Proceedings of the 17th Annual IEEE/ACM International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS), Sep. 2009, pp. 111–122.

[30] V. Venkatesan and I. Iliadis, “Effect of latent errors on the reliability of data storage systems,” in Proceedings of the 21st Annual IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS), Aug. 2013, pp. 293–297.

[31] ——, “A general reliability model for data storage systems,” in Proceedings of the 9th International Conference on Quantitative Evaluation of Systems (QEST), Sep. 2012, pp. 209–219.

[32] M. Silberstein, L. Ganesh, Y. Wang, L. Alvisi, and M. Dahlin, “Lazy means smart: Reducing repair bandwidth costs in erasure-coded distributed storage,” in Proceedings of the 7th ACM International Systems and Storage Conference (SYSTOR), Jun. 2014, pp. 15:1–15:7.

[33] K. V. Rashmi, N. B. Shah, D. Gu, H. Kuang, D. Borthakur, and K. Ramchandran, “A ”Hitchhiker’s” guide to fast and efficient data reconstruction in erasure-coded data centers,” in Proceedings of the 2014 ACM conference on SIGCOMM, Aug. 2014, pp. 331–342.

[34] DELL/EMC Whitepaper, ”PowerVault ME4 Series ADAPT Software,” Feb. 2019. [Online]. Available: https://www.dellemc.com/ [retrieved: December 7, 2022]

[35] I. Iliadis and V. Venkatesan, “Rebuttal to ‘Beyond MTTDL: A closed-form RAID-6 reliability equation’,” ACM Trans. Storage, vol. 11, no. 2, Mar. 2015, pp. 1–10.

[36] T. J. E. Schwarz, Q. Xin, E. L. Miller, D. D. E. Long, A. Hospodor, and S. Ng, “Disk scrubbing in large archival storage systems,” in Proceedings of the 12th Annual IEEE/ACM International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS), Oct. 2004, pp. 409–418.

[37] A. Oprea and A. Juels, “A clean-slate look at disk scrubbing,” in Proceedings of the 8th USENIX Conference on File and Storage Technologies (FAST), Feb. 2010, pp. 57–70.

[38] B. Schroeder, S. Damouras, and P. Gill, “Understanding latent sector errors and how to protect against them,” ACM Trans. Storage, vol. 6, no. 3, Sep. 2010, pp. 1–23.

[39] M. Zhang, S. Han, and P. P. C. Lee, “SimEDC: A simulator for the reliability analysis of erasure-coded data centers,” IEEE Trans. Parallel Distrib. Syst., vol. 30, no. 12, 2019, pp. 2836–2848.

[40] V. Venkatesan and I. Iliadis, “Effect of codeword placement on the reliability of erasure coded data storage systems,” in Proceedings of the 10th International Conference on Quantitative Evaluation of Systems (QEST), Sep. 2013, pp. 241–257.

[41] ——, “Effect of codeword placement on the reliability of erasure coded data storage systems,” IBM Research Report, RZ 3827, Aug. 2012.

[42] I. Iliadis and X.-Y. Hu, “Reliability assurance of RAID storage systems for a wide range of latent sector errors,” in Proceedings of the 2008 IEEE International Conference on Networking, Architecture, and Storage (NAS), Jun. 2008, pp. 10–19.

[43] Seagate, exos x20, data sheet. [Online]. Available: https://www.seagate.com/products/enterprise-drives/exos-x/x20/ [retrieved: December 7, 2022]

[44] Backblaze drive stats for 2021. [Online]. Available: https://www.backblaze.com/blog/backblaze-drive-stats-for-2021/ [retrieved: December 7, 2022]

[45] M. Ovsiannikov, S. Rus, D. Reeves, P. Sutter, S. Rao, and J. Kelly, “The quantcast file system,” in Proceedings of the 39th International Conference on Very Large Data Bases (VLDB), vol. 6, no. 11. VLDB Endowment, Aug. 2013, pp. 1092–1101.