Reliability Assessment of Erasure-Coded Storage Systems with Latent Errors

Ilias Iliadis
IBM Research Europe
8803 Rüschlikon, Zurich, 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** EAFDL **remains practically unaffected. 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. 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 [1-4] 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) [1-10] and, more recently, the Expected Annual Fraction of Data Loss (EAFDL) reliability metrics [11-14]. The latter metric was introduced because Amazon S3 [15], Facebook [16], LinkedIn [17] and Yahoo! [18] 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 [3] demonstrated that, for a replication factor of three, a declustered data placement scheme achieves a superior reliability than other placement schemes. This is the scheme that was originally used by Google [19], Facebook [20], and Microsoft Azure [21], but, to improve data reliability and storage efficiency further, today they use erasure coding schemes that offer higher efficiency [22-24].

The reliability of storage systems is further degraded by the occurrence of unrecoverable sector errors, that is, errors that cannot be corrected by the standard sector-associated error-correcting 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 [25-28]. 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 [7][10][26][27]. 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 [28][29] 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. It is known that the actual latent-error rates degrade MTTDL by orders of magnitude [7][10][26][28]. 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 [11-13]. The first analytical assessment of EAFDL in the presence of latent errors was presented in [30] 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 [14]. 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. To assess the reliability of erasure-coded systems, we adopt the non-Markovian methodology developed in prior work [11-13] 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 [3][8][11][31]. It has been

 ISBN: 978-1-61208-841-9

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 [14]. 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 [7]. 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 placement scheme currently used by Google [22], Windows Azure [24], 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 [7].

A simulation analysis of reliability aspects of erasure-coded data centers was presented in [35]. 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 that usefulness of the analytical reliability results derived in this article.

The key contributions of this article are the following. We consider the reliability of erasure-coded storage systems with latent errors and derive closed-form expressions 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. We subsequently demonstrate that, in the range of typical sector-error rates, unrecoverable failures are frequent, which degrades MTTDL. However, the relative increase of the amount of data loss is negligible, which leaves EAFDL practically unaffected in this range.

The remainder of the article is organized as follows. Section II describes the storage system model and the corresponding parameters considered. Section III 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 IV presents numerical results demonstrating the adverse effect of unrecoverable or latent errors and the effectiveness of these schemes for improving system reliability. Finally, we conclude in Section V.

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 an irrecoverable data loss ($\tilde{r} = m - l + 1$ and $2 \leq \tilde{r} \leq m$)
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

II. STORAGE SYSTEM MODEL

To assess the reliability of erasure-coded storage systems, we adopt the model used in [13] 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

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

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

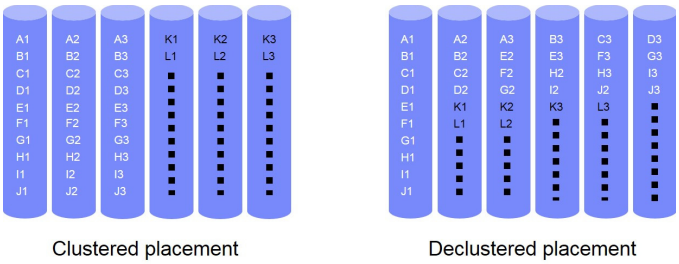

Figure 1. Clustered and declustered placement of codewords of length $m = 3$ on $n = 6$ devices. X1, X2, X3 represent a codeword (X = A, B, C, . . . , L).

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

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

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 [36]. 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

$$C_1 = C\,, \qquad (4)$$

where C is determined by (2). In Section III, we will derive the reliability metrics of interest using the *direct path approximation*, which considers only transitions from lower to higher exposure levels [3][8][11][31][36]. This implies that each exposure level is entered only once.

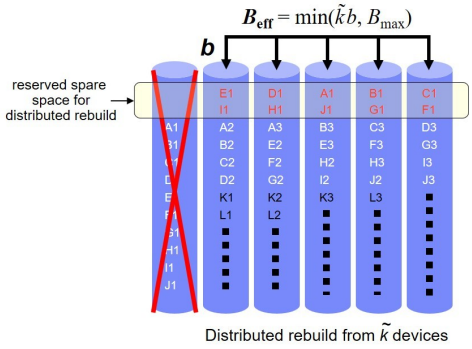

Figure 2. Rebuild under declustered placement.

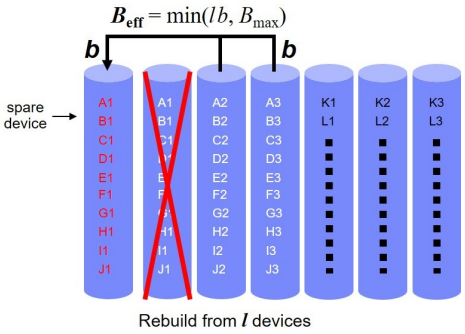

Figure 3. Rebuild under clustered placement.

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 [3][8][13][24][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 [13].

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

 ISBN: 978-1-61208-841-9

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 [13]. Also, let $1/\mu$ and $F_X(.)$ denote the mean 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. Therefore,

$$\mu^{-1} \triangleq E(X) = c/b \; . \tag{5}$$

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 $1/\lambda$. The results in this article hold for *highly reliable* storage devices, which satisfy the condition [13][31]

$$\mu \int_0^\infty F_\lambda(t)[1 - F_X(t)]dt \ll 1, \quad \text{with } \frac{\lambda}{\mu} \ll 1 \; . \tag{6}$$

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. Let R_u denote the rebuild time of the most-exposed codewords at exposure level u. Then, [12, Eq. (34)]

$$R_1 = (b/b_1)\, X \; . \tag{7}$$

Let α_u be the fraction of the rebuild time R_u still left when another device fails, causing the exposure level transition $u \rightarrow u+1$. In [37, Lemma 2], it was shown that, for highly reliable devices satisfying condition (6), α_u is approximately uniformly distributed in $(0, 1)$. 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 [12][13][36]. A moment's reflection leads to the following relation:

$$C_{u+1} \approx V_u\, \alpha_u\, C_u \; , \quad \text{for } \; u = 1, \ldots, \tilde{r} - 1 \; . \tag{8}$$

Repeatedly applying (8) and using (4) and the convention that for any sequence δ_i, $\prod_{i=1}^{0} \delta_i \triangleq 1$, yields

$$C_u \approx C \prod_{i=1}^{u-1} V_i\, \alpha_i \; , \quad \text{for } \; u = 1, \ldots, \tilde{r} \; . \tag{9}$$

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 [7]. Assuming that bit errors occur independently over successive bits, the unrecoverable sector (symbol) error probability P_s is

$$P_s = 1 - (1 - P_{\text{bit}})^s \; , \tag{10}$$

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 $P_s \approx 5 \times 10^{-9}$ [38].

III. Derivation of MTTDL and EAFDL

The reliability metrics are derived using the methodology presented in [11][12][13] and extending it to assess the effect of latent errors. This methodology uses the direct path approximation [10], does not involve Markovian analysis [3][8][11][31][36], and holds for general failure time distributions, which can be exponential or non-exponential, such as the Weibull and gamma distributions that satisfy condition (6).

At any point in time, the system can be thought to be in one of two modes: normal or rebuild mode. 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.

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 $u \rightarrow$ 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 $1 \rightarrow 2 \rightarrow \cdots \rightarrow u \rightarrow$ UF, for $u = 1, \ldots, \tilde{r} - 1$, and

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

with corresponding probabilities P_{UF_u} and P_{DF}, respectively. It holds that

$$P_{\text{UF}_u} = P_u \; P_{u \rightarrow \text{UF}} \; , \quad \text{for } \; u = 1, \ldots, \tilde{r} - 1 \; , \tag{11}$$

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

$$P_u \approx (\lambda\, c)^{u-1} \, \frac{1}{(u-1)!} \, \frac{E(X^{u-1})}{[E(X)]^{u-1}} \prod_{i=1}^{u-1} \frac{\tilde{n}_i}{b_i} \, V_i^{u-1-i} \; , \tag{12}$$

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

In [10], it was shown that P_{DL} is accurately approximated by the probability of all direct paths to data loss. Therefore,

$$P_{\text{DL}} \approx P_{\text{DF}} + \sum_{u=1}^{\tilde{r}-1} P_{\text{UF}_u} \; . \tag{13}$$

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

$$P_{\mathrm{UF}_u} \;\approx\; -\,(\lambda\,c)^{u-1}\;\frac{E(X^{u-1})}{[E(X)]^{u-1}}\;\left(\prod_{i=1}^{u-1}\frac{\tilde{n}_i}{b_i}\;V_i^{u-1-i}\right) \cdot \log(\hat{q}_u)^{-(u-1)}\left(\hat{q}_u - \sum_{i=0}^{u-1}\;\frac{\log(\hat{q}_u)^i}{i!}\right), \quad (14)$$

where $$\hat{q}_u \triangleq q_u^{\,C\,\prod_{j=1}^{u-1} V_j}\,, \quad (15)$$

$$q_u = 1 - \sum_{j=\tilde{r}-u}^{m-u}\binom{m-u}{j} P_s^j (1-P_s)^{m-u-j}\,, \quad (16)$$

$$P_{\mathrm{DF}} \approx (\lambda\,c)^{\tilde{r}-1}\;\frac{1}{(\tilde{r}-1)!}\;\frac{E(X^{\tilde{r}-1})}{[E(X)]^{\tilde{r}-1}}\;\prod_{i=1}^{\tilde{r}-1}\frac{\tilde{n}_i}{b_i}\,V_i^{\tilde{r}-1-i}. \quad (17)$$

Proof: Equation (14) is obtained in Appendix A. Equation (17) is obtained from the fact that $P_{\mathrm{DF}} = P_{\tilde{r}}$ and, subsequently, from (12) by setting $u = \tilde{r}$. ∎

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

$$Q \approx \begin{cases} H_{\mathrm{DF}}\,, & \text{if } \overrightarrow{DF} \\ H_{\mathrm{UF}_u}\,, & \text{if } \overrightarrow{UF_u}\,, \;\text{ for }\; u = 1, \ldots, \tilde{r}-1 \\ 0\,, & \text{otherwise}\,. \end{cases} \quad (18)$$

$$\text{Thus,}\quad E(Q) \approx P_{\mathrm{DF}}\;E(H_{\mathrm{DF}}) + \sum_{u=1}^{\tilde{r}-1} P_{\mathrm{UF}_u}\;E(H_{\mathrm{UF}_u}) \quad (19)$$

$$= E(Q_{\mathrm{DF}}) + \sum_{u=1}^{\tilde{r}-1} E(Q_{\mathrm{UF}_u})\,, \quad (20)$$

$$\text{where}\quad E(Q_{\mathrm{DF}}) = P_{\mathrm{DF}}\;E(H_{\mathrm{DF}})\,, \quad (21)$$

$$\text{and}\quad E(Q_{\mathrm{UF}_u}) = P_{\mathrm{UF}_u}\;E(H_{\mathrm{UF}_u})\,,\;\; u = 1, \ldots, \tilde{r}-1\,. \quad (22)$$

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

$$E(Q) \;=\; \frac{l}{m}\;E(S)\,s \;\overset{(2)}{=}\; \frac{l}{m}\;\frac{E(S)}{C}\;c\,. \quad (23)$$

$$\text{Similarly,}\quad E(Q_{\mathrm{DF}}) \;=\; \frac{l}{m}\;E(S_{\mathrm{DF}})\,s \;\overset{(2)}{=}\; \frac{l}{m}\;\frac{E(S_{\mathrm{DF}})}{C}\;c\,, \quad (24)$$

$$\text{and}\quad E(Q_{\mathrm{UF}_u}) = \frac{l}{m}\;E(S_{\mathrm{UF}_u})\,s \;\overset{(2)}{=}\; \frac{l}{m}\;\frac{E(S_{\mathrm{UF}_u})}{C}\;c\,. \quad (25)$$

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

$$E(Q_{\mathrm{UF}_u}) \approx\; c\;\frac{l\,\tilde{r}}{m}\;(\lambda\,c)^{u-1}\;\frac{1}{u!}\;\frac{E(X^{u-1})}{[E(X)]^{u-1}}\;\left(\prod_{i=1}^{u-1}\frac{\tilde{n}_i}{b_i}\,V_i^{u-i}\right) \cdot \binom{m-u}{\tilde{r}-u} P_s^{\tilde{r}-u}\,, \quad \text{for } P_s \ll \frac{1}{m-\tilde{r}}\,, \quad (26)$$

$$E(Q_{\mathrm{DF}}) \approx\; c\;\frac{l}{m}\;(\lambda\,c)^{\tilde{r}-1}\;\frac{1}{(\tilde{r}-1)\,!}\;\frac{E(X^{\tilde{r}-1})}{[E(X)]^{\tilde{r}-1}}\;\prod_{i=1}^{\tilde{r}-1}\frac{\tilde{n}_i}{b_i}\;V_i^{\tilde{r}-i}. \quad (27)$$

Proof: Equation (26) is obtained in Appendix B. Equation (27) is obtained from (26) by setting $u = \tilde{r}$, which is the same result as Eq. (25) of [13]. ∎

The MTTDL and EAFDL metrics are determined by [11, Eqs. (4), (5), and (9)]:

$$\text{MTTDL} \approx \frac{E(T)}{P_{\mathrm{DL}}} = \frac{1}{n\,\lambda\,P_{\mathrm{DL}}}\,, \quad (28)$$

$$\text{EAFDL} \approx \frac{E(Q)}{E(T)\cdot U} \;=\; \frac{n\,\lambda\,E(Q)}{U} \;\overset{(1)}{=}\; \frac{m\,\lambda\,E(Q)}{l\,c}\,, \quad (29)$$

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

B. Symmetric and Declustered Placement

We consider the case $m < k \leq n$. The special case $k = m$ corresponding to the clustered placement has to be considered separately for the reasons discussed in Section II-C2. At each exposure level u, for $u = 1, \cdots, \tilde{r} - 1$, it holds that [12][13]

$$\tilde{n}_u^{\mathrm{sym}} = \;k - u\,, \quad (30)$$

$$b_u^{\mathrm{sym}} = \frac{\min((k-u)\,b, B_{\max})}{l+1}\,, \quad (31)$$

$$V_u^{\mathrm{sym}} = \;\frac{m-u}{k-u}\,. \quad (32)$$

The corresponding parameters $\tilde{n}_u^{\mathrm{declus}}$, b_u^{declus}, and V_u^{declus} for the declustered placement are derived from (30), (31), and (32) by setting $k = n$.

C. Clustered Placement

It holds that [12][13]

$$\tilde{n}_u^{\mathrm{clus}} = m - u\,,\;\; b_u^{\mathrm{clus}} = \min(\,b\,, B_{\max}/l\,)\,,\;\; V_u^{\mathrm{clus}} = 1\,. \quad (33)$$

Remark 1: It follows from (33) that a system is not bandwidth-constrained when $B_{\max} \geq l\,b$. Then, $b_u^{\mathrm{clus}} = \min(b, B_{\max}/l) = b$. 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 (13), (14), and (17) yield

$$P_{\mathrm{DL}}^{\text{RAID-5}} \approx\; (m-1)\;\frac{\lambda\,c}{b}\;+\;1-(1-P_s)^{(m-1)\,C}\,, \quad (34)$$

which is the same result as Eq. (85) of [14], and

$$P_{\mathrm{DL}}^{\text{RAID-6}} \approx\; 1 - q_1^C\;+\left[1+\frac{1-q_2^C}{\log(q_2^C)}\right](m-1)\;\frac{\lambda\,c}{b} + \frac{(m-1)(m-2)}{2}\left(\frac{\lambda\,c}{b}\right)^2\frac{E(X^2)}{[E(X)]^2}\,, \quad (35)$$

where q_1, q_2 are determined by (16). This result is in agreement with Eq. (243) of [14]. Also, (20), (26), and (27) yield

$$E(Q^{\text{RAID-5}}) \approx\; \frac{l}{m}\;(m-1)\left(\frac{\lambda\,c}{b} + 2\,P_s\right)c\,, \quad (36)$$

which is the same result as Eq. (105) of [14], and

$$E(Q^{\text{RAID-6}}) \approx \frac{l}{m}\;\frac{(m-1)(m-2)}{2} \cdot\left[\left(\frac{\lambda\,c}{b}\right)^2\frac{E(X^2)}{[E(X)]^2} + 3\,\frac{\lambda\,c}{b}\,P_s + 3\,P_s^2\right]c\,. \quad (37)$$

This result is in agreement with Eq. (264) of [14].

TABLE II. TYPICAL VALUES OF DIFFERENT PARAMETERS

Parameter	Definition	Values
n	number of storage devices	64
c	amount of data stored on each device	12 TB
l	user-data symbols per codeword	13, 14, 15
m	symbols per codeword	16
s	symbol (sector) size	512 B
b	rebuild bandwidth per device	50 MB/s
λ^{-1}	mean time to failure of a storage device	**300,000** h to 1,000,000 h
U	amount of user data stored in the system	624 to 720 TB
μ^{-1}	time to read an amount c of data at a rate b from a storage device	66.7 h

IV. NUMERICAL RESULTS

Here, we assess the reliability of the clustered and declustered schemes for a system comprised of $n = 64$ devices (disks), where each device stores an amount $c = 12$ TB, and $m = 16$, $l = 13$, 14, and 15, and symbol size s is equal to a sector size of 512 bytes.

Typical parameter values are listed in Table II. The annualized failure rate (AFR) is in the range of 0.9% to 3%, which corresponds to a mean time to failure in the range of $300,000$ h to $1,000,000$ h. The parameter λ^{-1} is chosen to be equal to $300,000$ h. It is assumed that the reserved rebuild bandwidth b is equal to 50 MB/s, which yields a rebuild time of a device $\mu^{-1} = c/b = 66.7$ h, and that the network rebuild bandwidth is sufficiently large ($B_{\max} \geq n\,b = 3.2$ GB/s). We assume that the rebuild time distribution is deterministic, such that $E(X^k) = [E(X)]^k$. The obtained results are accurate because (6) is satisfied, given that $\lambda/\mu = 2.2 \times 10^{-4} \ll 1$.

First, we assess the reliability for the declustered placement scheme ($k = n = 64$). The probability of data loss $P_{\rm DL}$ is determined by (13) as a function of P_s and shown in Figure 4. The probabilities $P_{{\rm UF}_u}$ and $P_{\rm DF}$ are also shown, as obtained from (14) and (17), respectively. We observe that $P_{\rm DL}$ 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 $P_{\rm DL}$ and, by virtue of (28), the MTTDL are degraded by orders of magnitude. The normalized λ MTTDL measure is obtained from (28) 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 obtained from (20) and shown in Figure 6. The normalized expected amounts $E(Q_{{\rm UF}_u})/c$ and $E(Q_{\rm DF})/c$ are also shown as determined by (26) and (27), respectively. The normalized EAFDL/λ measure is obtained from (29) 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 reliability metrics corresponding to the clustered placement scheme ($k = m = 16$) are plotted in Figures 8, 9, 10, and 11. We observe that the reliability achieved by the clustered data placement scheme does not reach the level achieved by the declustered one.

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 12. This figure also shows the probabilities of data loss $P_{\rm DL}$ that correspond to these two configurations and obtained from (13) 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.

V. 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. The analytical reliability expressions derived here can identify 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 $\overrightarrow{UF_u} = 1 \to 2 \to \cdots \to u \to {\rm UF}$ and proceed to evaluate $P_{{\rm UF}_u}(R_1, \vec{\alpha}_{u-1})$, 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 (11) that

$$P_{{\rm UF}_u}(R_1, \vec{\alpha}_{u-1}) = P_u(R_1, \vec{\alpha}_{u-2}) \cdot P_{u \to {\rm UF}}(R_1, \vec{\alpha}_{u-1}) \,. \quad (38)$$

It follows from Eq.(111) of [12] by setting $\tilde{r} = u$ that

$$P_u(R_1, \vec{\alpha}_{u-2}) \approx (\lambda b_1 R_1)^{u-1} \prod_{i=1}^{u-1} \frac{\tilde{n}_i}{b_i} \, (V_i \, \alpha_i)^{u-1-i} \,. \quad (39)$$

Unconditioning (39) on $\vec{\alpha}_{u-2}$ and using the fact that $E\left(\prod_{i=1}^{u-1} \alpha_i^{u-1-i}\right) = \prod_{i=1}^{u-1} (u-i)^{-1} = [(u-1)!]^{-1}$ yields

$$P_u(R_1) \approx (\lambda b_1 R_1)^{u-1} \, \frac{1}{(u-1)!} \, \prod_{i=1}^{u-1} \frac{\tilde{n}_i}{b_i} \, V_i^{u-1-i} \,. \quad (40)$$

Unconditioning (40) on R_1 and using (5) and (7) yields (12).

We now proceed to calculate $P_{u \to {\rm UF}}(R_1, \vec{\alpha}_{u-1})$. 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

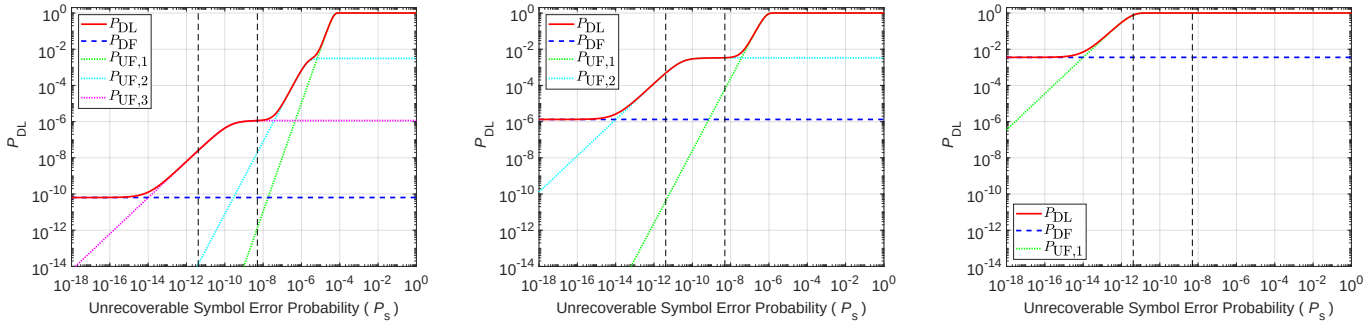

(a) $l = 13 \quad (\tilde{r} = 4)$ (b) $l = 14 \quad (\tilde{r} = 3)$ (c) $l = 15 \quad (\tilde{r} = 2)$

Figure 4. Probability of data loss P_{DL} vs. P_s for $l = 13, 14, 15$; $m = 16$, $n = k = 64$ (declustered scheme), $\lambda/\mu = 0.0002$, $c = 12$ TB, and $s = 512$ B.

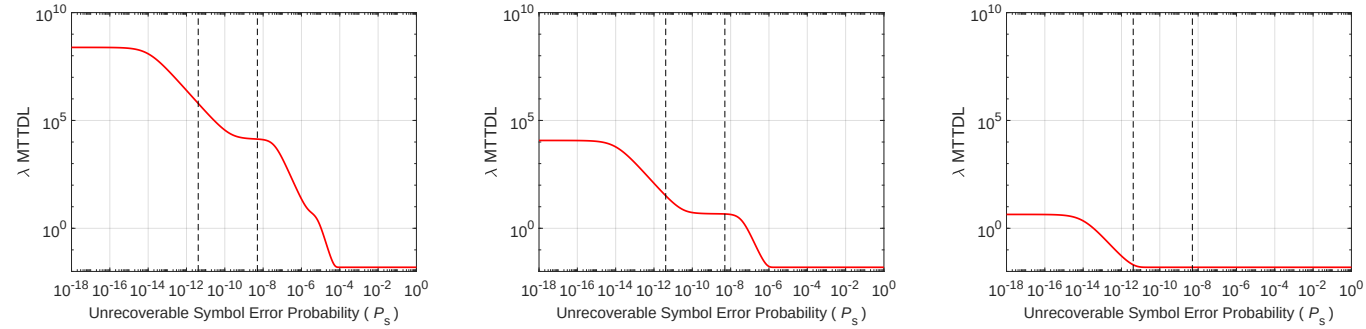

(a) $l = 13 \quad (\tilde{r} = 4)$ (b) $l = 14 \quad (\tilde{r} = 3)$ (c) $l = 15 \quad (\tilde{r} = 2)$

Figure 5. Normalized MTTDL vs. P_s for $l = 13, 14,$ and 15; $m = 16$, $n = k = 64$ (declustered scheme), $\lambda/\mu = 0.0002$, $c = 12$ TB, and $s = 512$ B.

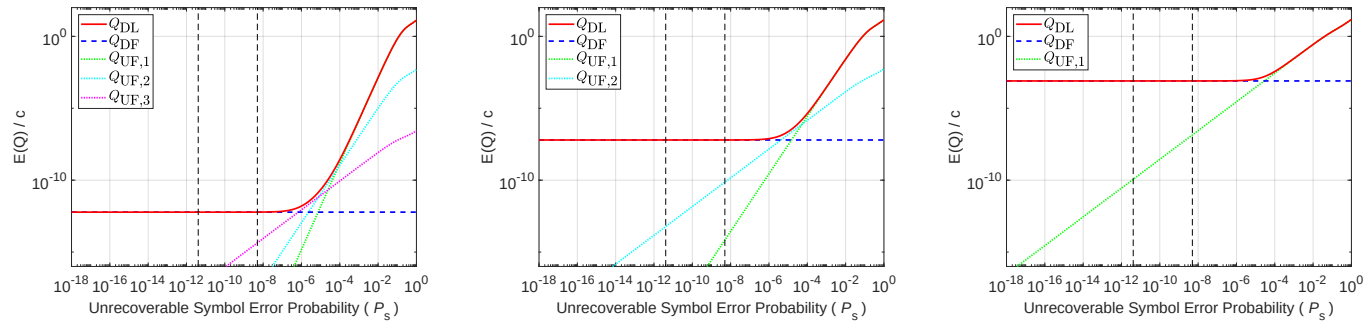

(a) $l = 13 \quad (\tilde{r} = 4)$ (b) $l = 14 \quad (\tilde{r} = 3)$ (c) $l = 15 \quad (\tilde{r} = 2)$

Figure 6. Normalized amount of data loss $E(Q)$ for $l = 13, 14, 15$; $m = 16$, $n = k = 64$ (declustered scheme), $\lambda/\mu = 0.0002$, $c = 12$ TB, and $s = 512$ B.

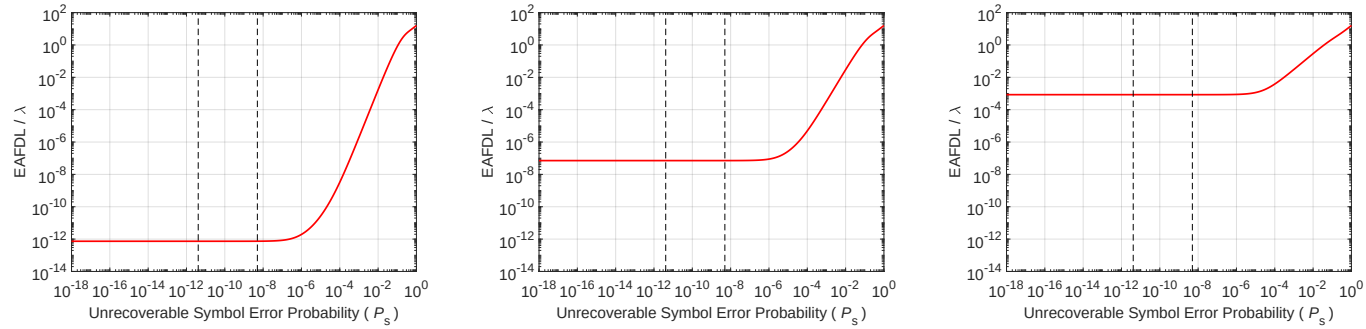

(a) $l = 13 \quad (\tilde{r} = 4)$ (b) $l = 14 \quad (\tilde{r} = 3)$ (c) $l = 15 \quad (\tilde{r} = 2)$

Figure 7. Normalized EAFDL vs. P_s for $l = 13, 14,$ and 15; $m = 16$, $n = k = 64$ (declustered scheme), $\lambda/\mu = 0.0002$, $c = 12$ TB, and $s = 512$ B.

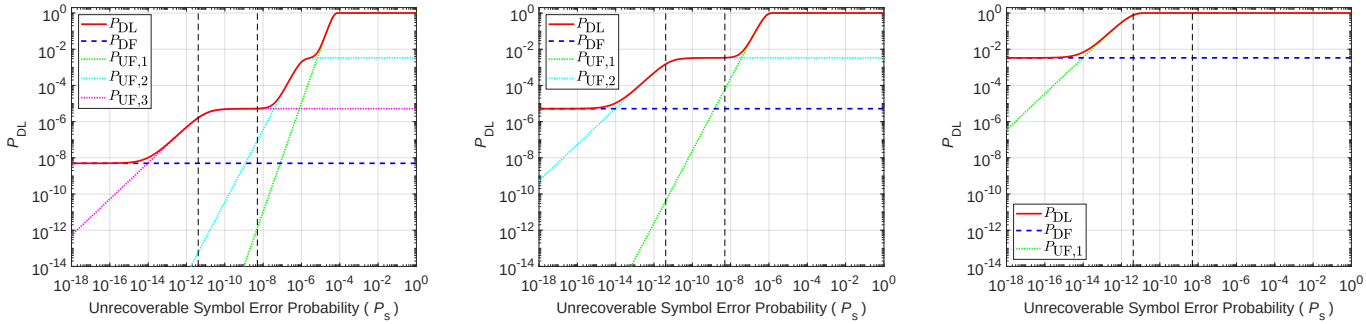

(a) $l = 13 \quad (\tilde{r} = 4)$ (b) $l = 14 \quad (\tilde{r} = 3)$, RAID-6 (c) $l = 15 \quad (\tilde{r} = 2)$, RAID-5

Figure 8. Probability of data loss P_{DL} vs. P_s for $l = 13, 14, 15$; $n = 64$, $k = m = 16$ (clustered scheme), $\lambda/\mu = 0.0002$, $c = 12$ TB, and $s = 512$ B.

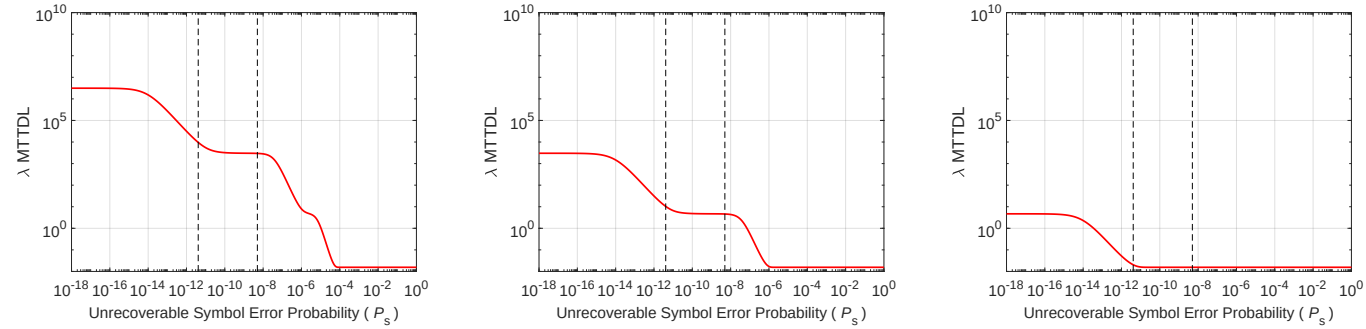

(a) $l = 13 \quad (\tilde{r} = 4)$ (b) $l = 14 \quad (\tilde{r} = 3)$, RAID-6 (c) $l = 15 \quad (\tilde{r} = 2)$, RAID-5

Figure 9. Normalized MTTDL vs. P_s for $l = 13, 14,$ and 15; $n = 64$, $k = m = 16$ (clustered scheme), $\lambda/\mu = 0.0002$, $c = 12$ TB, and $s = 512$ B.

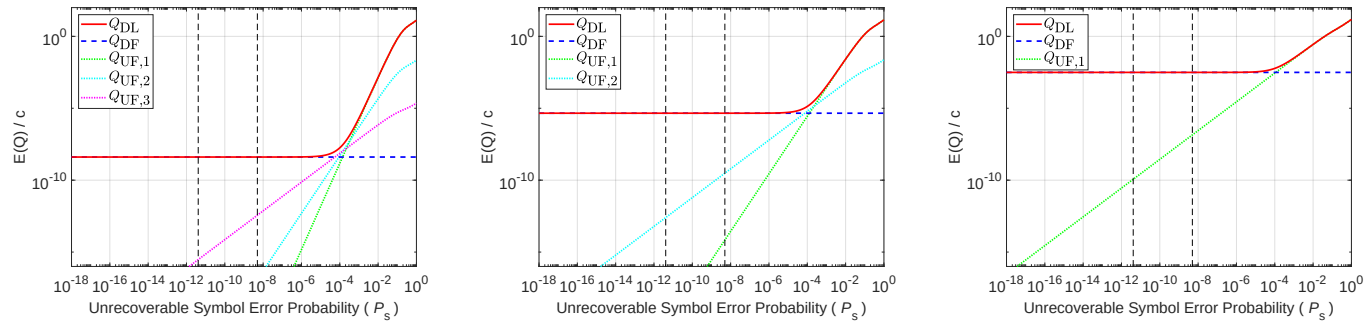

(a) $l = 13 \quad (\tilde{r} = 4)$ (b) $l = 14 \quad (\tilde{r} = 3)$, RAID-6 (c) $l = 15 \quad (\tilde{r} = 2)$, RAID-5

Figure 10. Normalized amount of data loss $E(Q)$ for $l = 13, 14, 15$; $n = 64$, $k = m = 16$ (clustered scheme), $\lambda/\mu = 0.0002$, $c = 12$ TB, and $s = 512$ B.

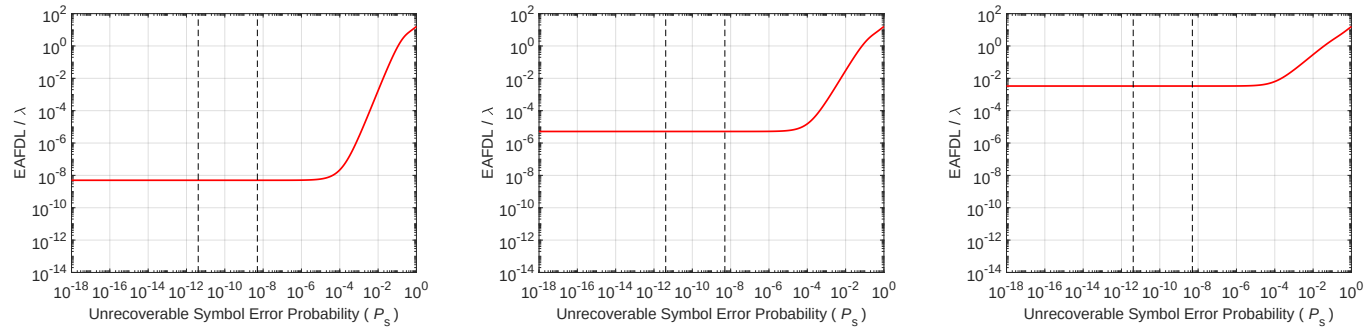

(a) $l = 13 \quad (\tilde{r} = 4)$ (b) $l = 14 \quad (\tilde{r} = 3)$, RAID-6 (c) $l = 15 \quad (\tilde{r} = 2)$, RAID-5

Figure 11. Normalized EAFDL vs. P_s for $l = 13, 14,$ and 15; $n = 64$, $k = m = 16$ (clustered scheme), $\lambda/\mu = 0.0002$, $c = 12$ TB, and $s = 512$ B.

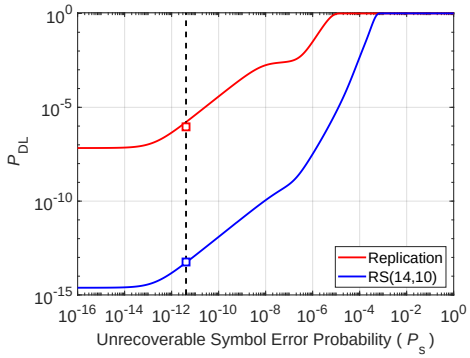

Figure 12. P_{DL} vs. P_s for $n = k = 210$, $\lambda^{-1} = 3$ years, $\mu^{-1} = 34$ hours, $\lambda/\mu = 0.0013$, $c = 15$ TB, and $s = 512$ B. Reliability schemes: 3-way replication and MDS(14,10).

a binomial distribution with parameter P_s, the probability that a symbol has a unrecoverable error. Thus, for $i = 0, \ldots, m-u$,

$$P(I_u = j) = \binom{m-u}{j} P_s^j \, (1-P_s)^{m-u-j} \,, \tag{41}$$

$$\approx \binom{m-u}{j} P_s^j \,, \quad \text{for } P_s \ll \frac{1}{m-u-j} \,, \tag{42}$$

such that $$E(I_u) = \sum_{j=1}^{m-u} j \, P(I_u = j) \; = \; (m-u) \, P_s \,. \tag{43}$$

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

$$q_u = P(L_u = 0) = 1 - P(I_u > \tilde{r} - u) \,, \tag{44}$$

which, using (41), yields (16).

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 $P_{\mathrm{UF}|C_u}$ of encountering an unrecoverable failure during the rebuild process of the C_u codewords is

$$P_{\mathrm{UF}|C_u} = 1 - q_u^{C_u} \,, \quad \text{for } u = 1, \ldots, \tilde{r} \,. \tag{45}$$

Substituting (9) into (45) and using (15) yields

$$P_{u \to \mathrm{UF}}(R_1, \vec{\alpha}_{u-1}) \approx 1 - q_u^{C \, \prod_{j=1}^{u-1} V_j \, \alpha_j} = 1 - \hat{q}_u^{\prod_{j=1}^{u-1} \alpha_j} \,. \tag{46}$$

Substituting (46) into (38) yields

$$P_{\mathrm{UF}_u}(R_1, \vec{\alpha}_{u-1}) \;\approx\; P_u(R_1, \vec{\alpha}_{u-2}) \left[1 - \hat{q}_u^{\prod_{j=1}^{u-1} \alpha_j} \right] . \tag{47}$$

Unconditioning (47) on $\vec{\alpha}_{u-1}$ and using (39) yields

$$P_{\mathrm{UF}_u}(R_1) \;\approx P_u(R_1) - (\lambda b_1 R_1)^{u-1} \left(\prod_{i=1}^{u-1} \frac{\tilde{n}_i}{b_i} \, V_i^{u-1-i} \right) \cdot E_{\vec{\alpha}_{u-1}} \left[\left(\prod_{i=1}^{u-1} \alpha_i^{u-1-i} \right) \hat{q}_u^{\prod_{j=1}^{u-1} \alpha_j} \right] . \tag{48}$$

For $\alpha_i \sim U(0,1)$ and for all $q \in \mathbb{R}$, it can be shown that

$$E\left[\left(\prod_{i=1}^{u-1} \alpha_i^{u-1-i} \right) q^{\prod_{i=1}^{u-1} \alpha_i} \right] = \frac{1}{(u-1)!} + \log(q)^{-(u-1)} \left(q - \sum_{i=0}^{u-1} \frac{\log(q)^i}{i!} \right) . \tag{49}$$

From (40) and (49), (48) yields

$$P_{\mathrm{UF}_u}(R_1) \;\approx\; -(\lambda b_1 R_1)^{u-1} \left(\prod_{i=1}^{u-1} \frac{\tilde{n}_i}{b_i} \, V_i^{u-1-i} \right) \cdot \log(\hat{q}_u)^{-(u-1)} \left(\hat{q}_u - \sum_{i=0}^{u-1} \frac{\log(\hat{q}_u)^i}{i!} \right) . \tag{50}$$

Unconditioning (50) on R_1, and using (5) and (7), yields (14).

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

$$E(L_u) = \sum_{i=\tilde{r}-u}^{m-u} (i+u) \, P(I_u = i) \,, \tag{51}$$

where $P(I_u = i)$ is given by (41). Considering approximation (42), it follows that

$$E(L_u) \;\approx\; \tilde{r} \binom{m-u}{\tilde{r}-u} P_s^{\tilde{r}-u} \,, \quad \text{for } P_s \ll \frac{1}{m-\tilde{r}} \,. \tag{52}$$

The expected number $E(S_{\mathrm{U}}|C_u)$ 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

$$E(S_{\mathrm{U}}|C_u) \overset{(52)}{\approx} C_u \, \tilde{r} \binom{m-u}{\tilde{r}-u} P_s^{\tilde{r}-u}, \;\; P_s \ll \frac{1}{m-\tilde{r}} \,. \tag{53}$$

Substituting (9) into (53) yields

$$E(S_{\mathrm{U}}|\vec{\alpha}_{u-1}) \;\approx\; C \left(\prod_{j=1}^{u-1} V_j \, \alpha_j \right) \tilde{r} \binom{m-u}{\tilde{r}-u} P_s^{\tilde{r}-u} \,. \tag{54}$$

Subsequently, the expected number $E(S_{\mathrm{UF}_u}|R_1, \vec{\alpha}_{u-1})$ 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

$$E(S_{\mathrm{UF}_u}|R_1, \vec{\alpha}_{u-1}) = P_u(R_1, \vec{\alpha}_{u-1}) \, E(S_{\mathrm{U}}|\vec{\alpha}_{u-1}) \,. \tag{55}$$

Substituting (39) and (54) into (55) yields

$$E(S_{\mathrm{UF}_u}|R_1, \vec{\alpha}_{u-1}) \;\approx\; (\lambda b_1 R_1)^{u-1} \left[\prod_{i=1}^{u-1} \frac{\tilde{n}_i}{b_i} \, (V_i \, \alpha_i)^{u-i} \right] \cdot \; C \, \tilde{r} \binom{m-u}{\tilde{r}-u} P_s^{\tilde{r}-u}, \quad P_s \ll \frac{1}{m-\tilde{r}} \,. \tag{56}$$

Unconditioning (56) on $\vec{\alpha}_{u-1}$ and R_1, and using (5), (7), and given that $E(\prod_{i=1}^{u-1} \alpha_i^{u-i}) = \prod_{i=1}^{u-1} (u-i+1)^{-1} = 1/u!$, yields

$$E(S_{\mathrm{UF}_u}|R_1) \approx \; (\lambda b_1 R_1)^{u-1} \left(\prod_{i=1}^{u-1} \frac{\tilde{n}_i}{b_i} \, V_i^{u-i} \right) \frac{1}{u!} \, C \, \tilde{r} \cdot \binom{m-u}{\tilde{r}-u} P_s^{\tilde{r}-u}, \quad P_s \ll \frac{1}{m-\tilde{r}} \,. \tag{57}$$

Unconditioning (57) on R_1, and using (5) and (7), yields

$$E(S_{\mathrm{UF}_u}) \approx (\lambda c)^{u-1} \frac{E(X^{u-1})}{[E(X)]^{u-1}} \left(\prod_{i=1}^{u-1} \frac{\tilde{n}_i}{b_i} V_i^{u-i} \right) \frac{1}{u!} C \tilde{r} \cdot \binom{m-u}{\tilde{r}-u} P_s^{\tilde{r}-u}, \quad P_s \ll \frac{1}{m-\tilde{r}} . \quad (58)$$

Substituting (58) into (25) yields (26).

REFERENCES

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

[2] P. M. Chen, E. K. 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.

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

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

[5] M. Malhotra and K. S. Trivedi, "Reliability analysis of redundant arrays of inexpensive disks," J. Parallel Distrib. Comput., vol. 17, 1993, pp. 146–151.

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

[7] 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, 2011, pp. 1–42.

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

[9] J.-F. Pâris, T. J. E. Schwarz, A. Amer, and D. D. E. Long, "Highly reliable two-dimensional RAID arrays for archival storage," in Proc. 31st IEEE Int'l Performance Computing and Communications Conference (IPCCC), Dec. 2012, pp. 324–331.

[10] 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.

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

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

[13] 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.

[14] ——, "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.

[15] Amazon Web Services, "Amazon Simple Storage Service (Amazon S3)," 2021. [Online]. Available: http://aws.amazon.com/s3/ [retrieved: March, 2021]

[16] D. Borthakur et al., "Apache Hadoop goes realtime at Facebook," in Proc. ACM Int'l Conference on Management of Data (SIGMOD), Jun. 2011, pp. 1071–1080.

[17] R. J. Chansler, "Data availability and durability with the Hadoop Distributed File System," ;login: The USENIX Association Newsletter, vol. 37, no. 1, Feb. 2013, pp. 16–22.

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

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

[20] D. Borthakur. HDFS and Erasure Codes (HDFS-RAID), Aug. 2009. [Online]. Available: https://hadoopblog.blogspot.com/2009/08 [retrieved: March, 2021]

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

[22] D. Ford et al., "Availability in globally distributed storage systems," in Proc. 9th USENIX Symposium on Operating Systems Design and Implementation (OSDI), Oct. 2010, pp. 61–74.

[23] S. Muralidhar et al., "f4: Facebook's Warm BLOB Storage System," in Proc. 11th USENIX Symposium on Operating Systems Design and Implementation (OSDI), Oct. 2014, pp. 383–397.

[24] C. Huang et al., "Erasure coding in Windows Azure Storage," in Proc. USENIX Annual Technical Conference (ATC), Jun. 2012, pp. 15–26.

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

[26] 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, 2008, pp. 1–42.

[27] I. Iliadis, "Reliability modeling of RAID storage systems with latent errors," in Proc. 17th Annual IEEE/ACM Int'l Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS), Sep. 2009, pp. 111–122.

[28] V. Venkatesan and I. Iliadis, "Effect of latent errors on the reliability of data storage systems," in Proc. 21th Annual IEEE Int'l Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS), Aug. 2013, pp. 293–297.

[29] 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.

[30] I. Iliadis, "Data loss in RAID-5 storage systems with latent errors," in Proc. 12th Int'l Conference on Communication Theory, Reliability, and Quality of Service (CTRQ), Mar. 2019, pp. 1–9.

[31] V. Venkatesan and I. Iliadis, "A general reliability model for data storage systems," in Proc. 9th Int'l 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 Proc. 7th ACM Int'l Systems and Storage Conference (SYSTOR), Jun. 2014, pp. 15:1–15:7.

[33] K. V. Rashmi et al., "A "Hitchhiker's" guide to fast and efficient data reconstruction in erasure-coded data centers," in Proc. 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: March, 2021]

[35] 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.

[36] V. Venkatesan and I. Iliadis, "Effect of codeword placement on the reliability of erasure coded data storage systems," in Proc. 10th Int'l Conference on Quantitative Evaluation of Systems (QEST), Sep. 2013, pp. 241–257.

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

[38] I. Iliadis and X.-Y. Hu, "Reliability assurance of RAID storage systems for a wide range of latent sector errors," in Proc. 2008 IEEE Int'l Conference on Networking, Architecture, and Storage (NAS), Jun. 2008, pp. 10–19.

 ISBN: 978-1-61208-841-9