Data Loss in RAID-5 Storage Systems with Latent Errors

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

***Abstract*—Storage systems employ redundancy and recovering schemes to protect against device failures and latent sector errors, and to enhance reliability. The effectiveness of these schemes has been evaluated based on the Mean Time to Data Loss (MTTDL) and the Expected Annual Fraction of Data Loss (EAFDL) metrics. The reliability degradation due to device failures has been assessed in terms of both these metrics, but the adverse effect of latent errors has been assessed only in terms of the MTTDL metric. This article addresses the issue of evaluating the amount of data losses caused by latent errors. It presents a methodology for obtaining MTTDL and EAFDL of RAID-5 systems analytically in the presence of unrecoverable or latent errors. A theoretical model capturing the effect of independent latent errors and device failures is developed, and closed-form expressions are derived for the metrics of interest.**

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 use data redundancy schemes to recover data lost due to device and component failures, and to enhance reliability [1]. Erasure coding schemes are deployed that provide high data reliability as well as high storage efficiency. Special cases of erasure codes are the replication schemes and the Redundant Arrays of Inexpensive Disks (RAID) schemes, such as RAID-5 and RAID-6, that have been deployed extensively in the past thirty years [2-5]. 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 [1][12][13]. The introduction of the latter metric was motivated by the fact that Amazon S3 considers the durability of data over a given year [14], and, similarly, Facebook [15], LinkedIn [16] and Yahoo! [17] consider the amount of data lost in given periods.

The reliability of storage systems is also degraded by the occurrence of unrecoverable or latent sector errors, that is, of 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). The effect of latent errors is quite pronounced in higher-capacity HDDs and storage nodes because of the high frequency of these errors [18-22]. The risk of irrecoverable loss of data rises in the presence of latent errors.

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][20][21][23]. The effect of latent errors on MTTDL of erasure coded storage systems for the practical case of non-exponential failure and rebuild time distributions was assessed in [22].

In this article, we consider the effect of latent errors not only MTTDL, but also on the amount of data lost for the case of non-exponential failure and rebuild time distributions. Clearly, when a data loss occurs, the amount of data lost due to a device failure is much larger than the amount of sectors lost due to latent errors. We present a non-Markovian methodology for deriving the MTTDL and EAFDL metrics analytically for the case of RAID-5 systems. We extend the methodology developed in prior work [12][13] to assess MTTDL and EAFDL of storage systems in the absence of latent errors. The validity of this methodology for accurately assessing the reliability of storage systems was confirmed by simulations in several contexts [4][9][12][24]. It was demonstrated that theoretical predictions for the reliability of systems comprised of highly reliable storage devices are in good agreement with simulation results. Consequently, the emphasis of the present work is on theoretically assessing the effect of latent errors on system reliability. This is the first work to study the effect of latent errors on EAFDL.

The remainder of the article is organized as follows. Section II describes the storage system model and the corresponding parameters considered. Section III considers the unrecoverable or latent errors and the frequency of their occurrence. Section IV presents the general framework and methodology for deriving the MTTDL and EAFDL metrics analytically for the case of RAID-5 systems and in the presence of latent errors. Closed-form expressions for relevant reliability metrics, such as the probability of data loss and the amount of data loss, are derived. Section V presents numerical results demonstrating the effectiveness of the RAID-5 scheme for improving system reliability and the adverse effect of unrecoverable or latent errors on the probability of data loss and on the MTTDL and EAFDL reliability metrics. Section VI provides a discussion concerning the results obtained. Finally, we conclude in Section VII.

II. Storage System Model

The storage system considered here comprises n storage devices (nodes or disks), with each device storing an amount c of data, such that the total storage capacity of the system is $n\,c$. 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.

 ISBN: 978-1-61208-699-6

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
N	number of devices in a RAID-5 array ($N = m$)
b	average reserved rebuild bandwidth per device
R	time required to read (or write) an amount c of data at an average rate b from (or to) a device
$F_R(.)$	cumulative distribution function of R
$F_\lambda(.)$	cumulative distribution function of device lifetimes
$se^{(\text{RAID-5})}$	storage efficiency of redundancy scheme ($se^{(\text{RAID-5})} = l/m$)
U	amount of user data stored in the system ($U = se^{(\text{RAID-5})}\, n\, c$)
C	number of codewords stored in a RAID-5 array ($C = c/s$)
μ^{-1}	mean time to read (or write) an amount c of data at an average rate b from (or to) a device ($\mu^{-1} = E(R) = c/b$)
λ^{-1}	mean time to failure of a storage device ($\lambda^{-1} = \int_0^\infty [1 - F_\lambda(t)]dt$)

A. Redundancy

We consider an $(m, l) = (N, N-1)$ maximum distance separable (MDS) erasure code, which is a mapping from $N-1$ user-data symbols to a set of N symbols, called a codeword, having the property that any subset containing $N-1$ of the N symbols of the codeword can be used to decode (reconstruct, recover) the codeword. A single parity symbol is computed by using the XOR operation on $l = N-1$ user-data symbols to form a codeword with $m = N$ symbols in total. Such a scheme can tolerate a single erasure anywhere in the codeword. The N symbols of each codeword are stored on N distinct devices. More specifically, this scheme is used by the popular RAID-5 system, in which the n devices are arranged in groups (or arrays), each with N devices, one of which is redundant [2][3]. The storage system therefore comprises n/N RAID-5 arrays with each array having the ability to tolerate one device failure. The storage efficiency $se^{(\text{RAID-5})}$ of the system is given by

$$se^{(\text{RAID-5})} = \frac{l}{m} = \frac{N-1}{N} \,. \tag{1}$$

Consequently, the amount of user data U stored in the system is given by

$$U = se^{(\text{RAID-5})}\, n\, c = \frac{l\, n\, c}{m} \,. \tag{2}$$

Also, the number C of codewords in a device is given by

$$C = \frac{c}{s} \,. \tag{3}$$

Our notation is summarized in Table I. The parameters are divided according to whether they are independent or derived, and are listed in the upper and lower part of the table, respectively.

B. Codeword Reconstruction

When a storage device of an array fails, the C codewords stored in the array lose one of their symbols. Subsequently, the system starts to reconstruct the lost codeword symbols using the surviving symbols of the affected codewords. We assume that device failures are detected instantaneously, which immediately triggers the rebuild process. A certain proportion of the device bandwidth is reserved for data recovery during the rebuild process, where b denotes the actual average reserved rebuild bandwidth per device. This bandwidth is usually only

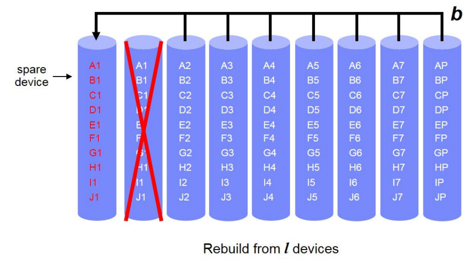

Figure 1. Rebuild for a RAID-5 array with $N = m = 8$ and $l = 7$.

a fraction of the total bandwidth available at each device, the remaining bandwidth being used to serve user requests.

The rebuild process attempts to restore the codewords of the affected array sequentially. The lost symbols are reconstructed directly in a spare device as shown in Figure 1. Decoding and re-encoding of data are assumed to be done on the fly, so the reconstruction time is equal to the time taken to read and write the required data to the spare device. Consequently, the time required to recover the amount c of data lost is equal to the time R required to read (or write) an amount c of data from (or to) a device. In particular, $1/\mu$ denotes the average time required to read (or write) an amount c of data from (or to) a device, which is given by

$$\frac{1}{\mu} \triangleq E(R) = \frac{c}{b} \,. \tag{4}$$

C. Failure and Rebuild Time Distributions

We adopt the model and notation considered in [13]. 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 [24]. The storage devices are characterized to be *highly reliable* in 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 on the order of thousands of hours) is very small, that is,

$$\frac{\lambda}{\mu} = \frac{\lambda\, c}{b} \ll 1 \,. \tag{5}$$

We consider storage devices whose cumulative distribution function F_λ satisfies the condition

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

where $F_R(.)$ is the cumulative distribution function of the rebuild time R. Then the MTTDL and EAFDL reliability metrics tend to be insensitive to the device failure distribution, that is, they depend only on its mean $1/\lambda$, but not on its density $F_\lambda(.)$ [13].

III. DATA LOSS FROM UNRECOVERABLE ERRORS

The reliability of RAID-5 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 $1{\times}10^{-15}$ for SCSI drives and

 ISBN: 978-1-61208-699-6

$1{\times}10^{-14}$ for SATA drives [8]. Assuming that bit errors occur independently over successive bits, the unrecoverable sector (symbol) error probability P_s is given by

$$P_s = 1 - (1 - P_{\text{bit}})^s \ , \tag{7}$$

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 $4.096{\times}10^{-12}$ in the case of SCSI and $4.096{\times}10^{-11}$ in the case of SATA drives. However, empirical field results suggest that the actual values can be orders of magnitude higher reaching $P_s \approx 5{\times}10^{-9}$ [25].

IV. Derivation of MTTDL and EAFDL

The MTTDL metric assesses the expected amount of time until some data can no longer be recovered and therefore is irrecoverably lost whereas the EAFDL assesses the fraction of stored data that is expected to be irrecoverably lost by the system annually. We briefly review the general methodology for deriving the MTTDL and EAFDL metrics presented in [12]. This methodology does not involve Markovian analysis 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 mode or rebuild mode. During normal mode, all devices are operational and all data in the system has the original amount of redundancy. Any symbols encountered with unrecoverable or latent errors are corrected through the RAID-5 capability. However, multiple unrecoverable errors encountered in a codeword can no longer be recovered and therefore lead to data loss. A transition from normal mode to rebuild mode occurs when a device fails; we refer to the device failure that causes this transition as a *first-device* failure. During rebuild mode, an active rebuild process attempts to restore the lost data in a spare device, which eventually leads the system either to an irrecoverable data loss (DL) with probability P_{DL} or back to the original normal mode by restoring initial redundancy, which occurs with probability $1 - P_{\text{DL}}$.

Let T be a typical interval of a fully operational period, that is, the time 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 equal to $1/\lambda$, the expected duration of T is given by [12]

$$E(T) = 1/(n\,\lambda) \ , \tag{8}$$

and MTTDL by

$$\text{MTTDL} \approx \frac{E(T)}{P_{\text{DL}}} = \frac{1}{n\,\lambda\,P_{\text{DL}}} \ . \tag{9}$$

The EAFDL is obtained as the ratio of the expected amount of user data lost, normalized to the amount of user data, to the expected duration of T [12, Equation (9)]:

$$\text{EAFDL} \approx \frac{E(Q)}{E(T)\cdot U} \overset{(8)}{=} \frac{n\,\lambda\,E(Q)}{U} \overset{(2)}{=} \frac{m\,\lambda\,E(Q)}{l\,c} \ , \tag{10}$$

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

The expected amount $E(H)$ of data lost, given that a data loss has occurred, is given by [12, Equation (8)]:

$$E(H) = \frac{E(Q)}{P_{\text{DL}}} \ . \tag{11}$$

From (9) and (10), it follows that the derivation of the MTTDL and EAFDL metrics requires the evaluation of P_{DL} and $E(Q)$, respectively. These quantities are derived using the direct path approximation [4][24][26], which, under conditions (5) and (6), accurately assesses the reliability metrics of interest [11][12][24][27].

A. Rebuild Process

When a storage device of an array fails, the C codewords stored in the array lose one of their symbols. Using the direct-path-approximation methodology, we proceed by considering only the subsequent potential data losses and device failures related to the affected array.

1) Unrecoverable Failure: The rebuild process attempts to restore the C codewords of the affected array sequentially. Let us consider such a codeword and let L be the number of symbols irrecoverably lost and I be the number of symbols encountered with unrecoverable errors in the codeword. As P_s is the probability that a symbol has a latent (unrecoverable) error, $1 - P_s$ is the probability that a symbol can be read successfully and, owing to the independence of symbol errors, it therefore holds that

$$P(I = i) = \binom{m-1}{i} P_s^i (1-P_s)^{m-1-i}, \text{ for } i = 0, \ldots, m-1 \ , \tag{12}$$

such that

$$E(I) = \sum_{i=1}^{m-1} i\,P(I = i) \ = \ (m-1)\,P_s \ . \tag{13}$$

Clearly, the symbol lost due to the device failure can be corrected through the RAID-5 capability only if the remaining $m-1$ symbols can be read. Thus, $L = 0$ if and only if $I = 0$. Using (12), the probability q that a codeword can be restored is given by

$$q = P(L = 0) = P(I = 0) = (1 - P_s)^{m-1} \ . \tag{14}$$

Note that if a codeword cannot be restored, then at least one of its l user-data symbols is lost. We now deduce that the probability P_{UF} of encountering an unrecoverable failure (UF) during the rebuild process of the C codewords is given by

$$P_{\text{UF}} = 1 - q^C \overset{(14)}{=} 1 - (1 - P_s)^{(m-1)\,C} \ . \tag{15}$$

Furthermore, such an unrecoverable failure entails the loss of user data. Let us denote by N_{UF} the number of codewords that cannot be recovered owing to unrecoverable failures. Then it holds that

$$E(N_{\text{UF}}) = C\,(1 - q) \ . \tag{16}$$

Remark 1: For very small values of P_s, it holds that $(1 - P_s)^{(m-1)\,C} \approx 1 - (m-1)\,C\,P_s$. Consequently, it follows from (15) that

$$P_{\text{UF}} \approx \begin{cases} (m-1)\,C\,P_s \,, & \text{for } P_s \ll P_s^{(2)} \\ 1 \,, & \text{for } P_s \gg P_s^{(2)} \ . \end{cases} \tag{17}$$

where $P_s^{(2)}$ is obtained from the approximation (17) as follows:

$$P_{\mathrm{UF}} \approx (m-1)\, C P_s^{(2)} = 1 \quad \Rightarrow \quad P_s^{(2)} \triangleq \frac{1}{C} \cdot \frac{1}{m-1} \,. \tag{18}$$

Note also that from (14) and (16), it follows that

$$E(N_{\mathrm{UF}}) \approx C\,(m-1)\,P_s \,, \quad \text{for } P_s \ll \frac{1}{m-1} \,. \tag{19}$$

In particular, for $P_s = P_s^{(2)}$, it holds that $E(N_{\mathrm{UF}}) \approx 1$ and this, combined with the fact that $P_{\mathrm{UF}} \approx 1$, implies that almost surely one of the C codewords cannot be recovered owing to an unrecoverable failure.

If $I > 0$, the number L of symbols lost is equal to $I + 1$. Consequently, the expected number $E(L)$ of symbols lost is given by

$$E(L) = \sum_{i=1}^{m-1} (i+1)\, P(I = i) = E(I) + 1 - P(I = 0) \,, \tag{20}$$

and using (12), (13), and (14) yields

$$E(L) = 1 - q + (m-1)P_s = 1 - (1-P_s)^{m-1} + (m-1)P_s \,. \tag{21}$$

Remark 2: For small values of P_s, it holds that $q = (1 - P_s)^{m-1} \approx 1 - (m-1)\,P_s$. Consequently, it follows from (21) that

$$E(L) \approx 2\,(m-1)\,P_s \,, \quad \text{for } P_s \ll \frac{1}{m-1} \,. \tag{22}$$

In particular, the expected number $E(L|L > 0)$ of symbols lost, given that the codeword cannot be restored, is given by

$$E(L|L>0) = \frac{E(L)}{P(L>0)} = \frac{E(L)}{1 - P(L=0)} \overset{(14)}{=} \frac{E(L)}{1-q} \overset{(22)}{\approx} 2 \,, \qquad \text{for } P_s \ll \frac{1}{m-1} \,. \tag{23}$$

2) Device Failure: A subsequent device failure (DF) may occur during the rebuild process triggered by the initial device failure. The probability $P_{\mathrm{DF|R}}$ that one of the $m-1$ remaining devices in the array fails during the rebuild process depends on the duration of the corresponding rebuild time R and the aggregate failure rate of these $m-1$ highly reliable devices, and is given by [24]

$$P_{\mathrm{DF|R}} \approx (m-1)\,\lambda\, R \,. \tag{24}$$

In particular, it was shown in [28, Lemma 2] that, for highly reliable devices satisfying conditions (5) and (6), the fraction of the rebuild time R still remaining when another device fails is approximately uniformly distributed between 0 and 1. This implies that the probability $P_{\mathrm{DF}}(j|R)$ that a device failure occurs while reconstructing the jth $(1 \le j \le C)$ codeword during the rebuild process, and given a rebuild time of R, is equal to $P_{\mathrm{DF|R}}/C$, which, using (24), yields

$$P_{\mathrm{DF}}(j|R) \approx \frac{(m-1)\,\lambda\, R}{C} \,, \quad \text{for } j = 1, 2, \ldots, C \,. \tag{25}$$

The probability P_{DF} of a device failure during the rebuild process is obtained by unconditioning (24) on R, that is,

$$P_{\mathrm{DF}} = E(P_{\mathrm{DF|R}}) \approx (m-1)\,\lambda\, E(R) \overset{(4)}{=} (m-1)\,\frac{\lambda}{\mu} \,. \tag{26}$$

Similarly, the probability $P_{\mathrm{DF}}(j)$ of a subsequent device failure during the reconstruction of the jth $(1 \le j \le C)$ codeword is obtained by unconditioning (25) on R, that is,

$$P_{\mathrm{DF}}(j) = E(P_{\mathrm{DF}}(j|R)) \approx \frac{(m-1)\,\lambda\, E(R)}{C} \overset{(4)}{=} \frac{(m-1)}{C}\,\frac{\lambda}{\mu} \,. \tag{27}$$

B. Data Loss

Data loss may occur because of another device failure or an unrecoverable failure of one or more codewords, or a combination thereof. Note that in all cases, data loss cannot involve only parity data, but also loss of user data. Let P_{DL} denote the probability of data loss. Then, the probability $1 - P_{\mathrm{DL}}$ of the rebuild being completed successfully is equal to the product of $1 - P_{\mathrm{DF}}$, the probability of not encountering a device failure during a rebuild, and $1 - P_{\mathrm{UF}}$, the probability of not encountering an unrecoverable failure during the rebuild, namely, $1 - P_{\mathrm{DL}} = (1 - P_{\mathrm{DF}})(1 - P_{\mathrm{UF}})$. Consequently,

$$P_{\mathrm{DL}} = P_{\mathrm{DF}} + (1 - P_{\mathrm{DF}})\, P_{\mathrm{UF}} \,. \tag{28}$$

Substituting (15) and (26) into (28) yields

$$P_{\mathrm{DL}} \approx (m-1)\,\frac{\lambda}{\mu} + \left[1 - (m-1)\,\frac{\lambda}{\mu}\right] \left[1 - (1-P_s)^{(m-1)\,C}\right] . \tag{29}$$

Remark 3: It follows from (17) and (26) that the region $[0, P_s^{(1)}]$ of P_s in which the probability P_{UF} is much smaller than the probability P_{DF} of encountering a device failure during the rebuild process is obtained by

$$\begin{aligned} P_{\mathrm{UF}} \ll P_{\mathrm{DF}} \quad &\Leftrightarrow \quad (m-1)\,C\,P_s \ll (m-1)\,\frac{\lambda}{\mu} \\ &\Leftrightarrow \quad P_s \ll P_s^{(1)} \triangleq \frac{1}{C} \cdot \frac{\lambda}{\mu} \,. \end{aligned} \tag{30}$$

C. Amount of Data Loss

Depending on whether a subsequent device failure occurs during the rebuild process, two cases are considered:

1) No Device Failure during Rebuild: The probability of this event is equal to $1 - P_{\mathrm{DF}}$. The expected number of symbols lost due to unrecoverable errors, $E(S_{\mathrm{U}}^{\odot} \,|\, \text{no DF})$, is given by

$$E(S_{\mathrm{U}}^{\odot} \,|\, \text{no DF}) = C\,E(L) = C\,[1 - q + (m-1)P_s] \,. \tag{31}$$

Unconditioning on the event of not having a device failure during the rebuild process, and using (14) and (26), we get

$$\begin{aligned} E(S_{\mathrm{U}}^{\odot}) &= E(S_{\mathrm{U}}^{\odot} \,|\, \text{no DF}) P(\text{no DF}) = E(S_{\mathrm{U}}^{\odot} \,|\, \text{no DF})(1 - P_{\mathrm{DF}}) \\ &\approx C\,[1 - (1-P_s)^{m-1} + (m-1)P_s\,] \left[1 - (m-1)\,\frac{\lambda}{\mu}\right] . \end{aligned} \tag{32}$$

2) Device Failure during Rebuild: Suppose a subsequent device failure occurs while reconstructing the jth $(1 \le j \le C)$ codeword. The probability of this event, denoted by $P_{\mathrm{DF}}(j)$, is given by (27). In this case, the two symbols of this codeword that are stored on the two failed devices can no longer be recovered and are lost. Furthermore, each of the remaining $m-2$ symbols may be lost owing to unrecoverable errors with probability P_s. The same applies for the remaining $C - j$ codewords. Thus, the total number $S_{\mathrm{D}}(j)$ of symbols that are stored on the two failed devices and are lost is given by

 ISBN: 978-1-61208-699-6

$$S_{\rm D}(j) = 2\,(C + 1 - j)\ . \tag{33}$$

Also, the expected total number $E(S_{\rm U}^{+}\,|\,\text{DF at j})$ of symbols stored in these $C - j + 1$ codewords that are lost owing to unrecoverable errors is given by

$$E(S_{\rm U}^{+}\,|\,\text{DF at j}) = (C + 1 - j)\,(m - 2)\,P_s\ . \tag{34}$$

Furthermore, each of the $j - 1$ codewords considered for reconstruction prior to the subsequent device failure loses an expected number of $E(L)$ symbols. Consequently, the expected total number $E(S_{\rm U}^{-}\,|\,\text{DF at j})$ of symbols stored in these $j - 1$ codewords that are lost owing to unrecoverable errors is given by

$$E(S_{\rm U}^{-}\,|\,\text{DF at j}) = (j - 1)\,E(L)\ . \tag{35}$$

Unconditioning (33), (34), and (35) on the event of a device failure during the reconstruction of the jth codeword, and using (27), yields

$$E(S_{\rm D}) \approx \sum_{j=1}^{C} 2\,(C + 1 - j)\,\frac{(m-1)}{C}\,\frac{\lambda}{\mu} \tag{36}$$

$$= (C + 1)\,(m - 1)\,\frac{\lambda}{\mu}\ , \tag{37}$$

$$E(S_{\rm U}^{+}) \approx \sum_{j=1}^{C} (C + 1 - j)\,(m - 2)\,P_s\,\frac{(m-1)}{C}\,\frac{\lambda}{\mu} \tag{38}$$

$$= \frac{C+1}{2}\,(m - 1)\,(m - 2)\,P_s\,\frac{\lambda}{\mu}\ , \tag{39}$$

and using (21)

$$E(S_{\rm U}^{-}) \approx \sum_{j=1}^{C} (j - 1)\,E(L)\,\frac{(m-1)}{C}\,\frac{\lambda}{\mu} \tag{40}$$

$$= \frac{C-1}{2}\,[1 - (1 - P_s)^{m-1} + (m - 1)P_s](m - 1)\,\frac{\lambda}{\mu}\ . \tag{41}$$

Combining the two cases, and using (32), (39), and (41), the expected number $E(S_{\rm U})$ of symbols lost due to unrecoverable errors is obtained as follows:

$$\begin{aligned} E(S_{\rm U}) &= E(S_{\rm U}^{\odot}) + E(S_{\rm U}^{+}) + E(S_{\rm U}^{-}) \\ &\approx C\,[\,1 - (1 - P_s)^{m-1} + (m - 1)P_s\,] \\ &\quad - \frac{C+1}{2}\,[\,1 - (1 - P_s)^{m-1} + P_s\,](m - 1)\,\frac{\lambda}{\mu}\ . \end{aligned} \tag{42}$$

Remark 4: From (32), (39), and (41), it follows that $E(S_{\rm U}^{\odot}) \gg E(S_{\rm U}^{-}) > E(S_{\rm U}^{+})$ because $E(S_{\rm U}^{-})$ and $E(S_{\rm U}^{+})$ are of the order $O(\lambda/\mu)$, which is very small, whereas $E(S_{\rm U}^{\odot})$ is not. Moreover, for large C, we have $E(S_{\rm U}^{-})/E(S_{\rm U}^{+}) \approx [\,1-(1-P_s)^{m-1}+(m-1)P_s\,]/[(m-2)P_s] > 1$. In particular, for small P_s, we have $(1 - P_s)^{m-1} \approx 1 - (m - 1)P_s$, which implies that $E(S_{\rm U}^{-})/E(S_{\rm U}^{+}) \approx 2(m-1)/(m-2) > 1$. Consequently, the symbols lost due to unrecoverable errors are predominately encountered during a rebuild that is completed without experiencing an additional device failure.

From (37) and (42), it follows that the expected total number of symbols lost $E(S)$ is given by

$$E(S) = E(S_{\rm D}) + E(S_{\rm U}) \tag{43}$$

$$\begin{aligned} &\approx C\,[\,1 - (1 - P_s)^{m-1} + (m - 1)P_s\,] \\ &\quad + \frac{C+1}{2}\,[\,1 + (1 - P_s)^{m-1} - P_s\,](m - 1)\,\frac{\lambda}{\mu}\ . \end{aligned} \tag{44}$$

Remark 5: For small values of P_s, it follows from (44) that

$$E(S) \approx 2\,C\,(m-1)P_s + \frac{C+1}{2}\,(2-mP_s)(m-1)\,\frac{\lambda}{\mu}\ , \tag{45}$$

which implies that for $P_s = 0$, $E(S) = E(S_{\rm D}) = (C{+}1)\,(m{-}1)\,\lambda/\mu$.

Remark 6: When P_s increases and approaches 1, it follows from (44) that $E(S)$ approaches $C\,m$. This is intuitively obvious because when $P_s = 1$, all the $C\,m$ symbols stored in the system are lost because of unrecoverable errors.

We now proceed to derive $E(Q)$, the expected amount of user data lost. First, we note that the expected number of user symbols lost is equal to the product of the storage efficiency to the expected number of symbols lost. Consequently, it follows from (44) that

$$E(Q) = \frac{l}{m}\,E(S)\,s \overset{(3)}{=} \frac{l}{m}\,\frac{E(S)}{C}\,c\ , \tag{46}$$

where s denotes the symbol size. Similar expressions for the expected amounts $E(Q_{\rm D})$ and $E(Q_{\rm U})$ of user data lost due to device and unrecoverable failures are obtained from $E(S_{\rm D})$ and $E(S_{\rm U})$, respectively. Thus, from (37), (42), and (44), it follows that

$$E(Q_{\rm D}) \approx \frac{l}{m}\,\frac{C+1}{C}\,(m - 1)\,\frac{\lambda}{\mu}\,c\ , \tag{47}$$

$$\begin{aligned} E(Q_{\rm U}) \approx \frac{l}{m}\Big\{ &1 - (1 - P_s)^{m-1} + (m - 1)P_s \\ &- \frac{C+1}{2\,C}\,[\,1 - (1 - P_s)^{m-1} + P_s\,](m - 1)\,\frac{\lambda}{\mu}\Big\}\,c\ , \end{aligned} \tag{48}$$

and

$$E(Q) = E(Q_{\rm D}) + E(Q_{\rm U}) \tag{49}$$

$$\begin{aligned} \approx \frac{l}{m}\Big\{ &1 - (1 - P_s)^{m-1} + (m - 1)P_s \\ &+ \frac{C+1}{2\,C}\,[\,1 + (1 - P_s)^{m-1} - P_s\,](m - 1)\,\frac{\lambda}{\mu}\Big\}\,c\ . \end{aligned} \tag{50}$$

Remark 7: For small values of P_s, and using (5), it follows from (48) that

$$E(Q_{\rm U}) \;\approx\; 2\,\frac{l}{m}\,(m - 1)\,c\,P_s\ . \tag{51}$$

Remark 8: From (47) and (51), it follows that the region $[0, P_s^{(3)}]$ of P_s in which $E(Q_{\rm U})$ is much smaller than $E(Q_{\rm D})$

is obtained by

$$
\begin{aligned}
E(Q_{\mathrm{U}}) &\ll E(Q_{\mathrm{D}}) \\
&\Leftrightarrow \quad 2\,\frac{l}{m}\,(m-1)\,c\,P_s \ll \frac{l}{m}\,\frac{C+1}{C}\,(m-1)\,\frac{\lambda}{\mu}\,c \\
&\Leftrightarrow \quad P_s \ll P_s^{(3)} \triangleq \frac{1}{2}\cdot\frac{C+1}{C}\cdot\frac{\lambda}{\mu}\,. \qquad (52)
\end{aligned}
$$

Remark 9: When P_s increases and approaches 1, it follows from (50) that $E(Q)$ approaches $C\,l$. This is intuitively obvious because when $P_s = 1$, upon the first-device failure, all the $C\,l$ user-data symbols stored in the RAID-5 array are lost owing to unrecoverable errors.

D. Reliability Metrics

The MTTDL normalized to $1/\lambda$ is obtained by substituting (29) into (9) as follows:

$$
\lambda\,\mathrm{MTTDL} \approx \frac{1}{n\left\{(m-1)\,\frac{\lambda}{\mu} + \left[1-(m-1)\,\frac{\lambda}{\mu}\right]\,\left[1-(1-P_s)^{(m-1)\,C}\right]\right\}}\,, \qquad (53)
$$

where C and λ/μ are given by (3) and (5), respectively.

The EAFDL normalized to λ is obtained by substituting (50) into (10) as follows:

$$
\begin{aligned}
\mathrm{EAFDL}/\lambda \approx\ & 1-(1-P_s)^{m-1}+(m-1)P_s \\
&+ \frac{C+1}{2\,C}\,[\,1+(1-P_s)^{m-1}-P_s\,](m-1)\,\frac{\lambda}{\mu}\,, \qquad (54)
\end{aligned}
$$

where C and λ/μ are given by (3) and (5), respectively.

The $E(H)$ normalized to c is obtained by substituting (50) and (29) into (11) as follows:

$$
\begin{aligned}
E(H)/c \approx\ & \frac{l}{m}\left\{1-(1-P_s)^{m-1}+(m-1)P_s \right. \\
&\left. + \frac{C+1}{2\,C}\,[\,1+(1-P_s)^{m-1}-P_s\,](m-1)\,\frac{\lambda}{\mu}\right\} \\
&\Big/ \left\{(m-1)\,\frac{\lambda}{\mu} + \left[1-(m-1)\,\frac{\lambda}{\mu}\right]\,\left[1-(1-P_s)^{(m-1)\,C}\right]\right\}, \qquad (55)
\end{aligned}
$$

where C and λ/μ are given by (3) and (5), respectively.

Similarly to (11), expressions for $E(H_{\mathrm{D}})$ and $E(H_{\mathrm{U}})$, the expected amounts of user data lost due to device and unrecoverable failures, given that such failures have occurred, are obtained as follows:

$$
E(H_{\mathrm{D}}) = \frac{E(Q_{\mathrm{D}})}{P_{\mathrm{DF}}}\,, \quad \text{and} \quad E(H_{\mathrm{U}}) = \frac{E(Q_{\mathrm{U}})}{P_{\mathrm{UF}}}\,, \qquad (56)
$$

respectively.

From (11), (49), and (56), we deduce that the following relation holds

$$
E(H) = \frac{P_{\mathrm{DF}}}{P_{\mathrm{DL}}}\,E(H_{\mathrm{D}}) + \frac{P_{\mathrm{UF}}}{P_{\mathrm{DL}}}\,E(H_{\mathrm{U}})\,. \qquad (57)
$$

Note that this is not a weighted average of $E(H_{\mathrm{D}})$ and $E(H_{\mathrm{U}})$ because the events of a subsequent device failure and of unrecoverable failures are not mutually exclusive, and therefore, and according to (28), the sum of weights is not equal to 1.

Remark 10: The normalized $E(H)/c$ exhibits two plateaus. According to (26), (30), (47), and (52), the first plateau is in the region $[0, P_s^{(1)}]$ of P_s, that is,

$$
\frac{E(H)}{c} \approx \frac{l}{m}\,\frac{C+1}{C}\,, \qquad \text{for } P_s \ll P_s^{(1)}\,. \qquad (58)
$$

For the second plateau, depending on the value of λ/μ, the following two cases are considered:

Case 1: $\lambda/\mu \gg 2/[(m-1)(C+1)]$. From (18) and (52), it holds that $P_s^{(2)} \ll P_s^{(3)}$. According to (17), (18), (47), and (52), the second plateau is in the region $[P_s^{(2)}, P_s^{(3)}]$ of P_s, that is,

$$
\frac{E(H)}{c} \approx \frac{l}{m}\,\frac{C+1}{C}\,(m-1)\,\frac{\lambda}{\mu}\,, \quad \text{for } P_s^{(2)} \ll P_s \ll P_s^{(3)}\,. \qquad (59)
$$

Case 2: $\lambda/\mu \ll 2/[(m-1)(C+1)]$. From (18) and (52), it holds that $P_s^{(3)} \ll P_s^{(2)}$. According to (17), (18), (51), and (52), the second plateau is in the region $[P_s^{(3)}, P_s^{(2)}]$ of P_s, that is,

$$
\frac{E(H)}{c} \approx \frac{l}{m}\,\frac{2}{C}\,, \quad \text{for } P_s^{(3)} \ll P_s \ll P_s^{(2)}\,. \qquad (60)
$$

Also, it follows from (51) that

$$
\frac{E(H)}{c} \approx 2\,\frac{l}{m}\,(m-1)P_s\,, \quad \text{for } P_s \gg \max(P_s^{(2)}, P_s^{(3)})\,. \qquad (61)
$$

Substituting (15), (26), (47), and (48) into (56) yields

$$
E(H_{\mathrm{D}})/c \approx \frac{l}{m}\,\frac{C+1}{C}\,, \qquad (62)
$$

and

$$
\begin{aligned}
E(H_{\mathrm{U}})/c \approx\ & \frac{l}{m}\left\{1-(1-P_s)^{m-1}+(m-1)P_s \right. \\
&\left. - \frac{C+1}{2\,C}\,[\,1-(1-P_s)^{m-1}+P_s\,](m-1)\,\frac{\lambda}{\mu}\right\} \\
&\Big/ \left[1-(1-P_s)^{(m-1)\,C}\right]\,, \qquad (63)
\end{aligned}
$$

where C and λ/μ are given by (3) and (5), respectively.

Remark 11: For small values of P_s, substituting (17) and (51) into (56) yields

$$
E(H_{\mathrm{U}})/c \approx \begin{cases} 2\,\frac{l}{m}\,\frac{1}{C}\,, & \text{for } P_s \ll P_s^{(2)} \\ 2\,\frac{l}{m}\,(m-1)P_s & \text{for } P_s \gg P_s^{(2)}\,. \end{cases} \qquad (64)
$$

Remark 12: When P_s increases and approaches 1, it follows from (55) that $E(H)$ approaches $C\,l$. This is intuitively obvious because when $P_s = 1$, all the $C\,l$ user-data symbols stored in the system are lost because of unrecoverable errors.

V. Numerical Results

We consider a RAID-5 system comprised of $n = 8$ devices with $N = m = 8$, $l = 7$, $\lambda/\mu = 0.001$, capacity $c = 1$ TB, and symbol size s equal to a sector size of 512 bytes, such that the number of codewords stored in a device is given by $C = c/s = 1.9{\times}10^9$.

The probability of data loss $P_{\rm DL}$ is obtained by (15), (26), and (29) as a function of the unrecoverable error probability P_s of a symbol (sector), and shown in Figure 2. It follows from (17) that, for small values of P_s, the probability $P_{\rm UF}$ of encountering an unrecoverable failure during the rebuild process increases linearly with P_s, as indicated by the dotted green line in Figure 2. According to (26), the probability $P_{\rm DF}$ of encountering a device failure during the rebuild process is independent of the unrecoverable symbol error probability, as indicated by the horizontal dotted blue line in Figure 2. It follows from (30) that, when P_s is in the region $[0, P_s^{(1)}]$, the probability $P_{\rm UF}$ of encountering an unrecoverable failure is much smaller than the probability $P_{\rm DF}$ of encountering a device failure during the rebuild process. From (30), and for the parameters considered, it follows that $P_s^{(1)} = 5{\times}10^{-13}$, as shown in Figure 2. Subsequently, for $P_s > P_s^{(1)}$, the probability $P_{\rm UF}$ of encountering an unrecoverable failure is much greater than that of encountering a device failure. In particular, it follows from (17) that, when $P_s \gg P_s^{(2)}$, $P_{\rm UF}$ and, in turn, $P_{\rm DL}$ approach 1 and are essentially independent of P_s. From (18), and for the parameters considered, it follows that $P_s^{(2)} = 7{\times}10^{-11}$, as shown in Figure 2. As expected, the total probability of data loss $P_{\rm DL}$ is monotonically increasing in P_s.

The normalized λ MTTDL measure is obtained by (53) and is shown in Figure 3 as a function of the unrecoverable symbol error probability. The various regions and plateaus are also depicted and correspond to the regions discussed above regarding the probability of data loss.

The normalized expected amount of user data lost to the amount of data stored in a device, $E(Q)/c$, is obtained by (47),

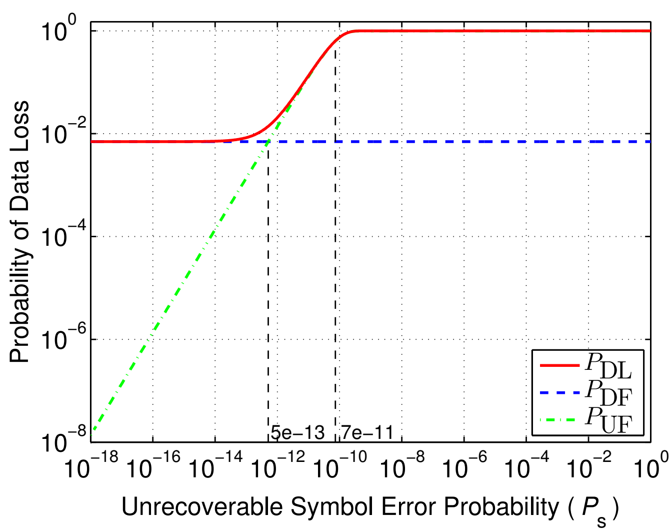

Figure 2. Probability of data loss $P_{\rm DL}$ for a RAID-5 array under latent errors ($\lambda/\mu = 0.001$, $m = N = 8$, $l = 7$, $c = 1$ TB, and $s = 512$ B).

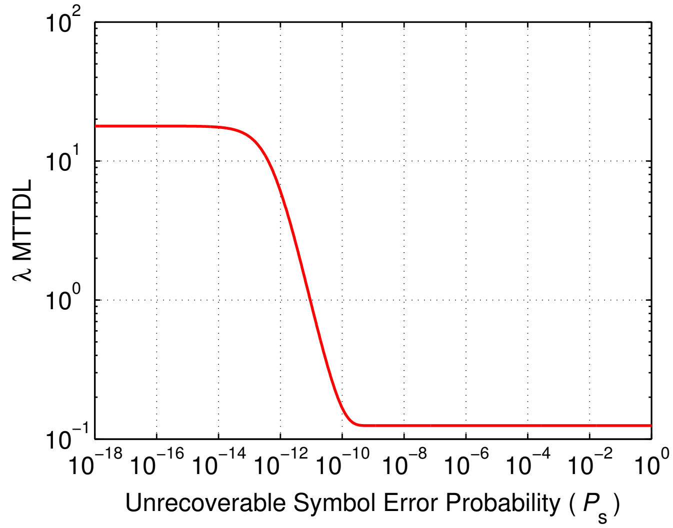

Figure 3. Normalized MTTDL for a RAID-5 array under latent errors ($\lambda/\mu = 0.001$, $m = N = 8$, $l = 7$, $c = 1$ TB, and $s = 512$ B).

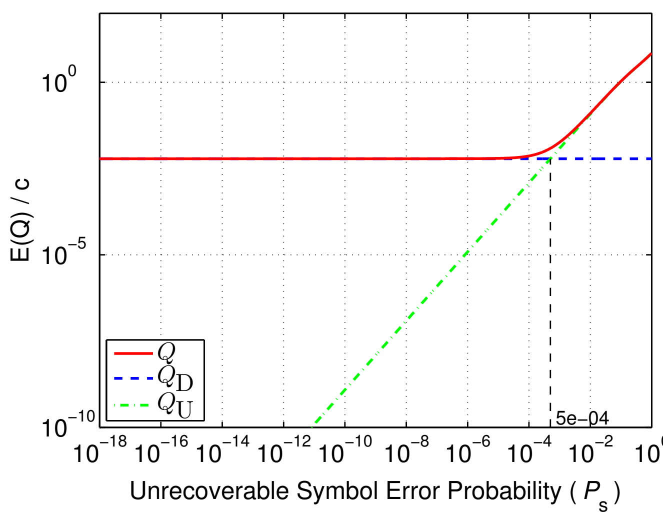

Figure 4. Normalized amount of data loss $E(Q)$ for a RAID-5 array under latent errors ($\lambda/\mu = 0.001$, $m = N = 8$, $l = 7$, $c = 1$ TB, and $s = 512$ B).

(48), and (50) as a function of the unrecoverable symbol error probability P_s, and shown in Figure 4. It follows from (51) that, for small values of P_s, the normalized expected amount $E(Q_{\rm U})/c$ of user data lost due to unrecoverable failures increases linearly with P_s, as indicated by the dotted green line in Figure 4. According to (47), the normalized expected amount $E(Q_{\rm D})/c$ of user data lost due to a subsequent device failure during the rebuild process is independent of the unrecoverable symbol error probability, as indicated by the horizontal dotted blue line in Figure 4. As anticipated, the total expected amount $E(Q)$ of user data lost increases monotonically with P_s. In particular, when P_s approaches 1 and according to Remark 9, the normalized expected amount $E(Q)/c$ of user data lost approaches $l = 7$, as all user data is lost.

The normalized EAFDL/λ measure is obtained by (54) and is shown in Figure 5 as a function of the unrecoverable symbol

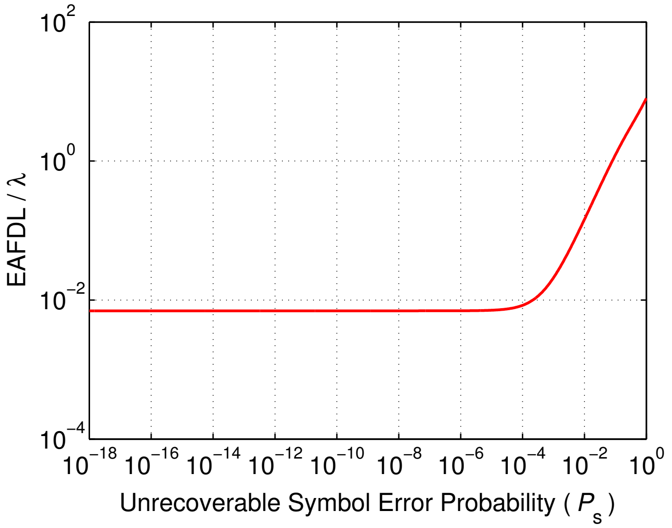

Figure 5. Normalized EAFDL for a RAID-5 array under latent errors ($\lambda/\mu = 0.001$, $m = N = 8$, $l = 7$, $c = 1$ TB, and $s = 512$ B).

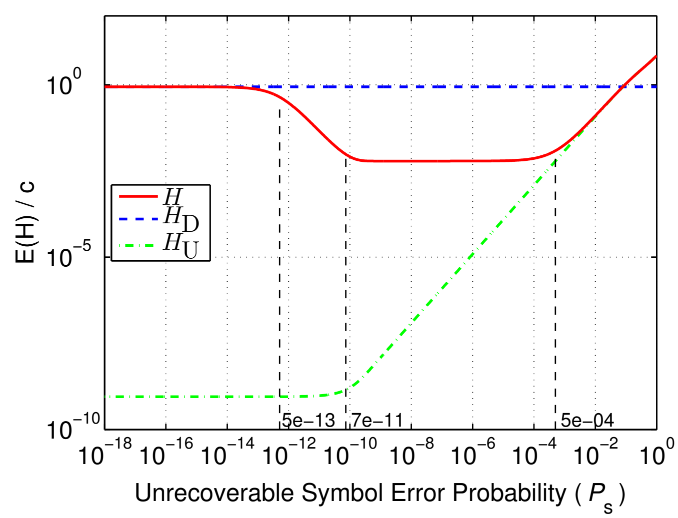

Figure 6. Normalized $E(H)$ for a RAID-5 array under latent errors ($\lambda/\mu = 0.001$, $m = N = 8$, $l = 7$, $c = 1$ TB, and $s = 512$ B).

error probability. Equation (10) suggests that this measure is proportional to $E(Q)$, which implies that the preceding discussion regarding the behavior of $E(Q)$ also holds here. Note also that, although the fraction of data loss never exceeds 1, EAFDL can exceed 1 because it expresses the annual fraction of data loss, which also takes into account the frequency of data losses.

The normalized expected amount $E(H)/c$ of user data lost, given that a data loss has occurred, to the amount of data stored in a device is obtained by (55), (62), and (63) as a function of the unrecoverable symbol error probability P_s, and shown in Figure 6. In contrast to the $P_{\rm DL}$, EAFDL, and $E(Q)$ measures that increase monotonically with P_s, we observe that $E(H)$ does not.

Data losses occur because of a subsequent device failure, unrecoverable failures of codewords, or a combination thereof. According to (62), the expected amount $E(H_{\rm D})$ of user data lost associated with a subsequent device failure, given that such a device failure has occurred during the rebuild process, is independent of the unrecoverable symbol error probability, as indicated by the horizontal dotted blue line in Figure 6. Such a device failure causes the loss of many symbols as opposed to a small number of additional symbols that may be lost owing to unrecoverable failures. In particular, according to Remark 2 and (23), each of the codewords that cannot be restored loses approximately two symbols. When P_s is extremely small, an unrecoverable failure is very unlikely, but when this occurs, it is caused by encountering a single codeword that cannot be recovered, which in turn results in the loss of two symbols. Consequently, and according to (64), for P_s such that $P_s \ll P_s^{(2)} = 7\times10^{-11}$, the expected amount $E(H_{\rm U})$ of user data lost due to unrecoverable failures, given that such unrecoverable failures have occurred, is independent of P_s, as indicated by the horizontal part of the dotted green line shown in Figure 6. Also, the amount of data lost corresponding to the two symbols lost is negligible compared with the amount of data lost due to a subsequent device failure, that is, $E(H_{\rm U}) \ll E(H_{\rm D})$.

The combined expected amount $E(H)$ of user data lost, given that a data loss has occurred, is an average of $E(H_{\rm D})$ and $E(H_{\rm U})$ with weights expressed by (57). For $P_s \ll P_s^{(1)} = 5\times10^{-13}$, a data loss is most likely attributed to a device failure, which results in the first plateau expressed by (58). However, for values of P_s in the region $[5\times10^{-13}, 7\times10^{-11}]$, this is reversed, in that it becomes more likely to encounter an unrecoverable failure than a device failure, and this causes $P_{\rm DL}$ to increase as shown in Figure 2. Consequently, as the weight of the $E(H_{\rm D})$ component decreases, so does $E(H)$. Subsequently, as P_s increases further, this weight can no longer decrease because $P_{\rm DL}$ has reached its maximum value of 1. Also, according to (19) and (64), the number of codewords with unrecoverable failures and the corresponding amount of data lost $E(H_{\rm U})$ increase linearly in P_s, but, although $E(H_{\rm U})$ increases, as indicated by the dotted green line in Figure 6, it still remains negligible compared with $E(H_{\rm D})$. Consequently, $E(H)$ no longer decreases and stabilizes at the second plateau level given by (59). As P_s increases further and exceeds $P_s^{(3)} = 5\times10^{-4}$, the increasing amount of data lost due to unrecoverable failures $E(H_{\rm U})$ exceeds $E(H_{\rm D})$, which in turn leads to an increase of the $E(H)$ metric. In particular, when P_s approaches 1, and according to Remark 9, the amount $l\,c$ of user data stored in the RAID-5 array is lost owing to unrecoverable errors, which in turn implies that the normalized expected amount $E(H)/c$ of user data lost approaches $l = 7$.

VI. Discussion

As discussed in Section III, field results suggest that the probability of unrecoverable sector errors lies in the range $[4.096\times10^{-11},\ 5\times10^{-9}]$. Figure 3 shows that MTTDL is significantly degraded by the presence of latent errors, whereas Figure 5 reveals that EAFDL is practically unaffected in this range. When the probability of unrecoverable sector errors lies in the region of practical interest, the probability of encountering an unrecoverable failure is much larger than that of encountering a device failure, which degrades MTTDL. However, the amount of sectors lost due to latent errors is negligible compared with the amount of data lost due to a device failure, which in turn implies that EAFDL remains unaffected. In contrast, Figure 6 reveals that the expected amount $E(H)$ of data lost, given that a data loss has occurred, decreases in the region of practical interest. This is due to the fact that when a data loss occurs, it is more likely caused by a unrecoverable failures that involve the loss of a small number of sectors rather than by a device failure that results in a significantly larger amount of data lost.

It follows from (30) and (52) that

$$P_s^{(1)} = \frac{1}{C} \cdot \frac{\lambda}{\mu} \ll \frac{1}{2} \cdot \frac{C+1}{C} \cdot \frac{\lambda}{\mu} = P_s^{(3)} \ . \qquad (65)$$

Consequently, increasing P_s first affects P_{DL}, MTTDL, and $E(H)$ and then $E(Q)$ and EAFDL.

VII. Conclusions

The effect of latent sector errors on the reliability of RAID-5 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. We established that the reliability of storage systems is adversely affected by the presence of latent errors. The results demonstrated that the effect of latent errors depends on the relative magnitudes of the probability of encountering a latent error versus the probability of encountering a device failure. It was found that, for actual values of the unrecoverable sector error probability, MTTDL is adversely affected by the presence of latent errors, whereas EAFDL is not.

Extending the methodology developed to derive the MTTDL and EAFDL reliability metrics of erasure coded systems in the presence of unrecoverable latent errors is a subject of further investigation.

References

[1] I. Iliadis, “Reliability of erasure coded systems under rebuild bandwidth constraints,” in Proceedings of the 11th International Conference on Communication Theory, Reliability, and Quality of Service (CTRQ), Apr. 2018, pp. 1–10.

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

[3] 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.

[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, 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] Amazon Simple Storage Service. [Online]. Available: http://aws.amazon.com/s3/ [retrieved: January 2019]

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

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

[17] 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.

[18] Hitachi Global Storage Technologies, Hitachi Disk Drive Product Datasheets. [Online]. Available: http://www.hitachigst.com/ [retrieved: January 2019]

[19] 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.

[20] 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.

[21] 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.

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

[23] 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.

[24] V. Venkatesan and I. Iliadis, “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.

[25] 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.

[26] 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.

[27] I. Iliadis and V. Venkatesan, “An efficient method for reliability evaluation of data storage systems,” in Proceedings of the 8th International Conference on Communication Theory, Reliability, and Quality of Service (CTRQ), Apr. 2015, pp. 6–12.

[28] V. Venkatesan and I. Iliadis, “Effect of codeword placement on the reliability of erasure coded data storage systems,” IBM Research Report, RZ 3827, Aug. 2012.