

RELIABILITY OF ERASURE CODED STORAGE SYSTEMS: A COMBINATORIAL-GEOMETRIC APPROACH

Vinay A. Vaishampayan, *Fellow, IEEE* and Antonio Campello, *Member, IEEE*

Abstract—We consider the probability of data loss, or equivalently, the reliability function for an erasure coded distributed data storage system under worst case conditions. Data loss in an erasure coded system depends on probability distributions for the disk repair duration and the disk failure duration. In previous works, the data loss probability of such systems has been studied under the assumption of exponentially distributed disk failure and disk repair durations, using well-known analytic methods from the theory of Markov processes. These methods lead to an estimate of the integral of the reliability function.

Here, we address the problem of directly calculating the data loss probability for general repair and failure duration distributions. A closed limiting form is developed for the probability of data loss and it is shown that the probability of the event that a repair duration exceeds a failure duration is sufficient for characterizing the data loss probability.

For the case of constant repair duration, we develop an expression for the conditional data loss probability given the number of failures experienced by a each node in a given time window. We do so by developing a geometric approach that relies on the computation of volumes of a family of polytopes that are related to the code. An exact calculation is provided and an upper bound on the data loss probability is obtained by posing the problem as a set avoidance problem. Theoretical calculations are compared to simulation results.

I. INTRODUCTION

Distributed data storage systems are growing in popularity, driven by demand and enabled by the availability of broadband networks, and declining costs of storage devices. Erasure coding represents a practical method for building highly reliable storage systems using low cost, less reliable storage drives. In an erasure coded storage system, a block of k information symbols from some finite set is encoded into a block of n coded symbols by an (n, k) erasure code and the n code

symbols are placed on separate disks. When a disk fails, it is *repaired*, i.e. redundant information in the code is used to recompute the erased symbol which is then placed on a replacement disk. Repair is essential for the reliability of the overall system. Data loss occurs or the system fails when the total number of failed disks at any time exceeds the erasure correcting capability of the code. If disks are repaired swiftly, the number of failed disks can be kept small on average, reducing the probability of data loss.

An important metric is the reliability function $R(t)$, defined to be the probability that data is not lost in the time window $[0, t]$. In previous works [1], [7], it is assumed that the repair and failure durations are exponentially distributed random variables and the mean time to data loss (MTTDL) is determined by analyzing a state transition diagram, where the system state at a given time is defined as the number of working disks at that time, see e.g. [14], [1], [7]. The reliability function $R(t)$ is then estimated by the formula $\exp(-t/\text{MTTDL})$. Exponentially distributed and independent failure durations are critical for this analysis to proceed. Several disk failure and disk repair modeling and measurement studies have been reported in the literature, e.g. [16], [13], [9], [18]. It is concluded that real world storage devices do not exhibit exponentially distributed lifetimes and that the Weibull distribution with appropriately chosen parameters is a more appropriate model for failure and repair durations. Simulation is a valuable tool for evaluating reliability of disk storage systems, see e.g. [11] which also includes a comprehensive review of previous modeling studies. In a recent contribution [17], it is shown that the reliability analysis based on the above exponential model is robust to changes in the disk failure time distribution. It is worth noting that [17] also points out that the analysis of MTTDL is *not* robust to changes in the repair duration distribution.

The main contributions of this work are summarized below.

- 1) We derive a formula for the data loss probability $P(\mathcal{D}_t)$ for small G and large t , for general independent and identically distributed (iid) failure and repair distributions, (16), restated here for convenience

$$\frac{P(\mathcal{D}_t)}{(G/n)^{(n-k)}} \approx \frac{(n-1)!}{(k-1)!} \frac{t}{E(Y)}, \quad (1)$$

where random variable Y represents a failure duration, G is the probability that $Y < Z$, where Z is a random repair duration and an (n, k) MDS erasure code is used. This is obtained by conditioning on a specific sequence of binary events, to be described later.

This work was partially presented at the IEEE International Conference on BigData 2013, and at the IEEE Information Theory Workshop 2014, Tasmania, Australia.

Vinay A. Vaishampayan, formerly with AT&T's Shannon Laboratory, is now with the City University of New York, College of Staten Island, Department of Engineering Science and Physics. His work was supported in part by CNPq Grant 400441/2014-4 and PSC-CUNY Award # 68631-00 46

Antonio Campello is currently with Télécom-ParisTech, France and University of Campinas, Brazil. His work was supported by São Paulo Research Foundation (FAPESP) grants 2013/25219-5, 2014/20602-8 and was initiated during a short-term visit to AT&T Shannon Laboratory, NJ, USA.

Copyright (c) 2014 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

Our derivation shows that for general distributions, the data loss probability is characterized in terms of the probability that the failure duration is smaller than the repair duration, and supports the fact (already known in the literature) that failure and repair rates are insufficient characterizations for determining the data loss probability. Our contribution is to show that the above mentioned probability is *sufficient* for characterizing the limiting data loss probability.

- 2) For constant repair duration, by conditioning on the number of failure events on each node, we arrive at another expression for the data loss probability, as well as a lower bound i.e. we derive an expression for $P_{\mathbf{m}}(\mathcal{D}_t)$, the data loss probability conditioned on $\mathbf{m} = (m_1, m_2, \dots, m_n)$, where m_i is the number of failures for disk i in time window $[0, t]$. Specifically, we prove that $P_{\mathbf{m}}(\mathcal{D}_t)$ has the following asymptotic behavior as $\tau = t_{\text{rep}}/t \rightarrow 0$:

$$\lim_{\tau \rightarrow 0} \frac{P_{\mathbf{m}}(\mathcal{D}_t)}{\tau^{n-k}} = (n-k+1)! \sum_{\substack{(i_1, \dots, i_{n-k+1}) \\ \text{distinct}}} m_{i_1} m_{i_2} \dots m_{i_{n-k+1}}. \quad (2)$$

This analysis holds for exponential failure distributions and provides a finer analysis of the system, not addressed by previous Markov chain approaches. It also has some implications to non-homogeneous Poisson processes.

- 3) By viewing the data loss probability calculation for constant repair duration as a problem of set avoidance by the Cartesian product of random sets, we derive an *upper* bound on the data loss probability,

$$P_{\mathbf{m}}(\mathcal{D}_t) \leq 1 - \left(1 - \frac{\text{vol } \mathcal{R}}{t^n}\right)^{m_1 \dots m_n}, \quad (3)$$

where $\mathcal{R} \subset [0, t]^n$ is a suitably defined error region associated with the code. The simplicity of the upper bound and the fact that its asymptotic behavior is comparable to the closed forms in some regimes makes it useful in practice. Methods for sharpening this bound remain as an open question.

- 4) We explore the connection between the erasure code and a family of polytopes that determine the error region. This connection is, in our opinion, interesting in its own right, even though it comes from an error probability calculation. Our contribution here is to develop a systematic approach for calculating the volume of a set of ordered points with constrained differences between successive elements. This method underlies the calculations for constant repair duration in this paper.

The paper is organized as follows. Sec. II contains a problem statement and states the assumptions that underlie our analysis. The data loss probability for general distributions is derived in Sec. III. For constant repair durations, we explore the combinatorial and geometric aspects of the problem of evaluating

the data loss probability in Sec. IV. Sec. V presents a method for upper bounding the data loss probability for constant repair duration by viewing the problem as a set avoidance problem. Volume calculations that underlie both the direct calculation as well as the set avoidance upper bound are presented in Sec. VI. Numerical and simulation results that explore some of the implications of the theory developed are presented in Sec. VII. The paper is summarized and suggestions for future research are presented in Sec. VIII. Some mathematical details and proofs are contained in the appendix.

II. ASSUMPTIONS, PROBLEM STATEMENT AND AN EXAMPLE

Code symbols from an MDS (n, k) erasure code are written to n disks¹. We assume that disk failures occur independently and that the disk failure process is modeled by an independent increment process with known probability distribution.

When a disk fails, data is downloaded from other disks and used to repair the lost symbols on the failed disk. We refer to these disks as *helper* disks, and to the set of helper disks as the *helper set*. The probability distribution of the repair duration, Z , is known, and repair durations are assumed to be independent and identically distributed. Since the codes are MDS, we consider that data is available as long as at least k disks are working (alternatively, if there was no instant of time at which less than k disks were working). Thus a data loss event occurs in the interval $[0, t)$ if the number of failed disks exceeds $(n - k)$ the erasure correcting capability of the code.

Characterization of a data loss event is subtle and depends on the system architecture, as well as on characteristics of the erasure code. An example is shown in Fig. 1 for constant repair duration t_{rep} . Disk 1 has failed and the helper set consists of disks 3 and 4. However, prior to disk 1 being restored, disk 2 fails. With a traditional MDS code, replacement symbols for disk 2 would be computed and the repair of disk 2 would begin without interrupting the repair of disk 1. On the other hand, in systems that perform functional repair [8], it is possible that the symbols for disk 1 would need to be recomputed as well, which implies that the repair process for disk 1 would need to be restarted. As a consequence, this sequence of failures and repairs results in a data loss event.

In our analysis we consider a disk to be repaired if and only if that disk repairs successfully, or a subsequently failed disk repairs successfully before the total number of failed disks exceeds the erasure correcting capability of the code.

III. ANALYSIS FOR GENERAL FAILURE AND REPAIR TIME DISTRIBUTIONS

We make the following assumption about our failure process. The i th inter-failure duration (hereafter referred to as the

¹To be precise, in the modern terminology it is said that the information is stored in a *node*. Throughout the paper we use the looser term disk instead, in analogy to classical storage systems.




Figure 1: Sequence of disk failures (shaded dots) that causes a data loss for a $(4, 2)$ coded system and helper set of size 2. Here the helper set for disk 1 is $\{3, 4\}$. When disk 2 fails, even though the helper set remains unchanged, the symbols for disk 1 must be recomputed (for functional repair). Since disk 3 fails prior to the repair being completed, a data loss event has occurred.




Figure 2: Illustration of setup for failure times and repair durations.

failure duration) for the system is denoted Y_i . The process $\{Y_i, i = 1, 2, \dots\}$ is an i.i.d. process with known probability density function (pdf) $f_Y(\cdot)$, cumulative distributive function (CDF) $F_Y(\cdot)$, where $f_Y(y) = 0$ for $y < 0$. Let $X_0 = 0$ and $X_i = X_{i-1} + Y_i, i = 1, 2, \dots$. Here $X_i, i = 1, 2, \dots$, is the instant at which the i th disk failure in the system occurs. Note that X_0 is not regarded as a failure instant. Let the random variable N_t count the number of failures in time interval $[0, t)$, i.e. $X_{N_t} < t$ and $X_{N_t+1} \geq t$. For an (n, k) MDS code an error cannot occur in $[0, t)$ if $N_t < n - k + 1$.

Our failure process associates with each X_i , a disk label drawn independently and uniformly from the set $\{1, 2, \dots, n\}$, where n is the block length of the (n, k) MDS erasure code being used. The amount of time taken to repair a disk after the i th failure instant X_i is denoted Z_{i+1} . The process $\{Z_i, i = 2, 3, \dots\}$ is assumed to be i.i.d with pdf $f_Z(\cdot)$ and CDF $F_Z(\cdot)$. We also define the indicator random variable $B_1 = 1$, and $B_i = 1$ if $Y_i < Z_i$, $B_i = 0$, otherwise, for $i = 2, 3, \dots$. We will use the notation $\mathbf{B}_{i:j} = (B_i, B_{i+1}, \dots, B_j)$. Our calculation is based on runs of ones ('1-runs') in $\mathbf{B}_{2:s}$. Let $\mathbf{U}_s = (U_1, U_2, \dots, U_R)$ denote the vector of runs of 1's in $\mathbf{B}_{2:s+1}$ for $s \geq 1$, where R denotes the number of 1-runs in $\mathbf{B}_{2:s}$. Thus, for the sequence $\mathbf{b}_{1:11} = 10001111011$, $u_s = \phi$ and $r = 0$, for $s = 1, 2, 3$, $u_s = (1)$ and $r = 1$, for $s = 4$ and

$u_s = (4, 2)$, $r = 2$, for $s = 10$.

The probability of data loss is given by

$$P(\mathcal{D}_t) = \sum_{s=n-k+1}^{\infty} \sum_{\mathbf{u}} P(\mathcal{D}_t | N_t = s, \mathbf{U}_s = \mathbf{u}) P(N_t = s, \mathbf{U}_s = \mathbf{u}). \quad (4)$$

In (4), the term $P(\mathcal{D}_t | N_t = s, \mathbf{U}_s = \mathbf{u})$, is the conditional probability of the event that in one of the 1-runs of a failure vector $\mathbf{B}_{2:s+1}$ with 1-run vector \mathbf{u} , the number of distinct disk failures exceeds $(n - k)$. The probability that exactly l distinct disks fail during a run of length u , denoted $\Pi_n(u + 1, l)$, is given by

$$\Pi_n(u + 1, l) = \frac{1}{n^{u+1}} \binom{n}{l} \sum_{a_1 > 0, a_2 > 0, \dots, a_l > 0} \binom{u + 1}{a_1, a_2, \dots, a_l}, \quad (5)$$

and $\Pi_n(u + 1, l) = 0$ for $l > u + 1$. In terms of $\Pi_n(u, l)$ we obtain

$$P(\mathcal{D}_t | N_t = s, \mathbf{U}_s = \mathbf{u}) = 1 - \prod_{j=1}^r \sum_{i=1}^{n-k} \Pi_n(u_j + 1, i), \quad (6)$$

where r is the number of 1-runs in the sequence of failures. Note that

$$\sum_{a_1 > 0, a_2 > 0, \dots, a_l > 0} \binom{u}{a_1, a_2, \dots, a_l} = l! S_2(u, l),$$

where $S_2(u, l)$ is the Stirling number of the second kind.

The second term in (4), $P(N_t = s, \mathbf{U}_s = \mathbf{u})$ is given by

$$P(N_t = s, \mathbf{U}_s = \mathbf{u}) = P(N_t = s | \mathbf{U}_s = \mathbf{u}) P(\mathbf{U}_s = \mathbf{u}) = P(N_t = s | \mathbf{U}_s = \mathbf{u}) N(s, \mathbf{u}) (1 - G)^{s-w(\mathbf{u})} G^{w(\mathbf{u})}, \quad (7)$$

where $w(\mathbf{u}) = \sum_{i=1}^r u_i$, $N(s, \mathbf{u})$ is the number of binary sequences $\mathbf{b}_{2:s+1}$ (of length s) with a 1-run vector \mathbf{u} , $G := \int_0^\infty F_Y(z) f_Z(z) dz$ is the probability that $Z_i > Y_i$ and $P(\mathbf{U}_s = \mathbf{u})$ denotes the probability of selecting a vector $\mathbf{B}_{2:s+1}$ with 1-run vector \mathbf{u} .

We thus obtain the following expression for the data loss probability,

$$\begin{aligned} P(\mathcal{D}_t) &= \sum_{s=n-k+1}^{\infty} \sum_{\mathbf{u}} P(\mathcal{D}_t | N_t = s, \mathbf{U}_s = \mathbf{u}) \\ &\quad \times N(s, \mathbf{u}) (1 - G)^{s-w(\mathbf{u})} G^{w(\mathbf{u})} P_r(N_t = s | \mathbf{U}_s = \mathbf{u}) \\ &= \sum_{s=n-k+1}^{\infty} \sum_{\mathbf{u}} \left(1 - \prod_{j=1}^r \sum_{i=1}^{n-k} \Pi_n(u_j + 1, i) \right) \\ &\quad \times N(s, \mathbf{u}) (1 - G)^{s-w(\mathbf{u})} G^{w(\mathbf{u})} P_r(N_t = s | \mathbf{U}_s = \mathbf{u}). \end{aligned} \quad (8)$$

1) *A Lower Bound:* A lower bound is obtained by writing

$$P(\mathcal{D}_t) = \sum_{s=n-k+1}^{\infty} \sum_{\mathbf{b}} P(\mathcal{D}_t | N_t = s, \mathbf{B}_{2:s+1} = \mathbf{b}) \times P(N_t = s, \mathbf{B}_{2:s+1} = \mathbf{b}) \quad (9)$$

and restricting the sum in (9) to $\mathbf{b} \in \mathcal{B}$ where

$$\mathcal{B} := \{1^{n-k}0^{s-(n-k)}, 01^{n-k}0^{s-1-(n-k)}, \dots, 0^{s-1-(n-k)}1^{n-k}0\}. \quad (10)$$

Observe that the cardinality $|\mathcal{B}| = s - (n - k)$. Also

$$P(\mathcal{D}_t | N_t = s, \mathbf{B}_{2:s+1} = \mathbf{b}) = \Pi_n(n - k + 1, n - k + 1)$$

for $\mathbf{b} \in \mathcal{B}$ and

$$\Pi_n(n - k + 1, n - k + 1) = \frac{(n - 1)!}{n^{n-k}(k - 1)!}.$$

Thus we obtain

$$\begin{aligned} P(\mathcal{D}_t) &\geq \frac{(n - 1)!}{n^{n-k}(k - 1)!} \sum_{s=n-k+1}^{\infty} \sum_{\mathbf{b} \in \mathcal{B}} P(N_t = s, \mathbf{B}_{2:s+1} = \mathbf{b}) \\ &= \frac{(n - 1)!}{n^{n-k}(k - 1)!} \sum_{s=n-k+1}^{\infty} \sum_{\mathbf{b} \in \mathcal{B}} P(N_t = s | \mathbf{B}_{2:s+1} = \mathbf{b}) \\ &\quad \times P(\mathbf{B}_{2:s+1} = \mathbf{b}) \\ &= \frac{(n - 1)!}{n^{n-k}(k - 1)!} \sum_{s=n-k+1}^{\infty} \sum_{\mathbf{b} \in \mathcal{B}} P(N_t = s | \mathbf{B}_{2:s+1} = \mathbf{b}) \\ &\quad \times (1 - G)^{s-(n-k)} G^{n-k} \\ &= \frac{G^{n-k}(n - 1)!}{n^{n-k}(k - 1)!} \sum_{s=n-k+1}^{\infty} (1 - G)^{s-(n-k)} \\ &\quad \sum_{\mathbf{b} \in \mathcal{B}} P(N_t = s | \mathbf{B}_{2:s+1} = \mathbf{b}) \\ &\stackrel{(a)}{=} \frac{G^{n-k}(n - 1)!}{n^{n-k}(k - 1)!} \sum_{s=n-k+1}^{\infty} (1 - G)^{s-(n-k)} \\ &\quad \times (s - (n - k)) P(N_t = s | \mathbf{B}_{2:s+1} = 1^{n-k}0^{s-(n-k)}) \\ &\stackrel{(b)}{=} \frac{G^{n-k}(n - 1)!}{n^{n-k}(k - 1)!} \sum_{s=n-k+1}^{\infty} (1 - G)^{s-(n-k)} \\ &\quad \times (s - (n - k)) P(N_t = s | \mathbf{B}_{2:*} = 1^{n-k}0^*) \end{aligned} \quad (11)$$

where in (a) we have used the facts (i) $|\mathcal{B}| = s - (n - k)$, and (ii) $P(N_t = s | \mathbf{B}_{2:s+1} = \mathbf{b})$ is a constant for $\mathbf{b} \in \mathcal{B}$. Observe that in (b) the conditioning event is that the infinitely long sequence $\mathbf{B}_{2:*} = 1^{n-k}000\dots =: 1^{n-k}0^*$. Henceforth, we denote $\mathbf{B}_{2:*}$ by \mathbf{B} .

2) *Limiting Behavior as $G \rightarrow 0, t \rightarrow \infty$:* The lower bound in (11) is tight in the limit as $G \rightarrow 0$ since it accounts for all terms except ones that are $o(G^{n-k})$. Thus

$$\begin{aligned} &\lim_{G \rightarrow 0} \frac{P(\mathcal{D}_t)}{(G/n)^{(n-k)}} \frac{(k - 1)!}{(n - 1)!} \\ &= \sum_{s=n-k+1}^{\infty} (s - (n - k)) P(N_t = s | \mathbf{B} = 1^{n-k}0^*) \\ &= \sum_{i=n-k+1}^{\infty} P(N_t \geq i | \mathbf{B} = 1^{n-k}0^*). \end{aligned} \quad (12)$$

In order to estimate (12), we follow the approach taken in [10], Ch. 3. Define indicator random variable I_i , $i = 1, 2, \dots$

$$I_i := \begin{cases} 1, & i = 1 \text{ or } (i > 1 \text{ and } \sum_{j=1}^{i-1} Y_j < t), \\ 0, & \text{otherwise.} \end{cases}$$

Note that the events $\{I_n = 1\}$ and $\{N_t \geq n - 1\}$ are identical. Now consider

$$\begin{aligned} &E(Y_1 + Y_2 + \dots + Y_{N_t+1} | \mathbf{B} = 1^{n-k}0^*) \\ &= E\left(\sum_{i=1}^{\infty} Y_i I_i | \mathbf{B} = 1^{n-k}0^*\right) \\ &= E(Y_1 I_1) + E\left(\sum_{i=2}^{n-k+1} Y_i I_i | \mathbf{B} = 1^{n-k}0^*\right) \\ &\quad + E\left(\sum_{i=n-k+2}^{\infty} Y_i I_i | \mathbf{B} = 1^{n-k}0^*\right) \\ &\stackrel{(a)}{=} \underbrace{E(Y) + E(Y | Y < Z) \sum_{i=1}^{n-k} P(N_t \geq i | \mathbf{B} = 1^{n-k}0^*)}_A \\ &\quad + E(Y | Y \geq Z) \sum_{i=n-k+1}^{\infty} P(N_t \geq i | \mathbf{B} = 1^{n-k}0^*), \end{aligned} \quad (13)$$

where in (a) we have used the independence of Y_i and I_i . The left hand side in (13) is larger than t by the mean residual time $\Delta := E(X_{N_t+1} - t | \mathbf{B})$. Thus

$$\sum_{i=n-k+1}^{\infty} P(N_t \geq i | \mathbf{B} = 1^{n-k}0^*) = \frac{t + \Delta - A}{E(Y | Y \geq Z)}. \quad (14)$$

Upon assuming that $E(Y^2 | Y > Z)$ is finite it follows that $E(\Delta | \mathbf{B})$ is finite. Thus for large t

$$\lim_{t \rightarrow \infty} \frac{1}{t} \sum_{i=n-k+1}^{\infty} P(N_t \geq i | \mathbf{B} = 1^{n-k}0^*) = \frac{1}{E(Y | Y \geq Z)}. \quad (15)$$

Since $E(Y | Y \geq Z) \rightarrow E(Y)$ as $G \rightarrow 0$ we obtain

$$\lim_{\substack{G \rightarrow 0, t \rightarrow \infty \\ tG^{n-k} \rightarrow 0}} \frac{P(\mathcal{D}_t)/t}{(G/n)^{(n-k)}} = \frac{(n - 1)!}{(k - 1)!} \frac{1}{E(Y)}. \quad (16)$$

Equation (16) leads to an *interesting and useful* qualitative conclusion about worst case repair duration distributions for the probability of data loss. If G is sufficiently small, the approximation $P(\mathcal{D}_t) \approx \frac{(n-1)!}{(k-1)!} \frac{t}{E(Y)} (G/n)^{(n-k)}$ becomes sharp. If, in addition, we assume that $F_Y(y)$ is *convex* (which holds for practical failure distributions such as exponential and a subset of Weibull distributions), by Jensen's inequality we have $G = P(Y \leq Z) \leq P(Y \leq E[Z])$, and thus

$$P(\mathcal{D}_t) \lesssim \frac{(n-1)!}{(k-1)!} \frac{t}{E(Y)} (P(Y \leq E[Z])/n)^{(n-k)}. \quad (17)$$

This means that in the limiting regime, the highest probability of data loss over a large class of failure distributions is when the time of repair is *constant*.

IV. PROBABILITY OF DATA LOSS CONDITIONED ON THE NUMBER OF FAILURES FOR CONSTANT REPAIR DURATION

We now turn our attention to the probability of data loss conditioned on the number of failures of each disk. The calculations hereafter consider the special yet important case of exponentially distributed failure durations and constant repair duration. We give the reader a glimpse of the main results for the case of a $(2, 1)$ erasure correcting code. Analyses for general (n, k) codes are presented in the subsequent subsection.

A. Motivating Example: $(2, 1)$ code

Suppose that we have one symbol of information stored in two disks. Let m_1 and m_2 denote the number of failures of disks 1 and 2, respectively. Let X_{11}, \dots, X_{1m_1} and X_{21}, \dots, X_{2m_2} be their random failure instants. By analyzing the failure timeline of both disks, we see that an error event occurs if and only if, for some failure instant X_{1i} of disk 1 and X_{2j} of disk 2, we have $|X_{1i} - X_{2j}| \leq t_{\text{rep}}$. Alternatively, there is no data loss if the random vector $\mathbf{X} = (X_{11}, \dots, X_{1m_1}, X_{21}, \dots, X_{2m_2})$ lies in the region:

$$\mathcal{R}^c = \{\mathbf{x} : 0 \leq x_{1i}, x_{2j} \leq t, |x_{1i} - x_{2j}| > t_{\text{rep}}, \forall i, j\}.$$

The probability that $\mathbf{X} \in \mathcal{R}^c$ can be calculated exactly, as outlined next. Consider the permutation π on the set $\{1, 2, \dots, s\}$, $s = m_1 + m_2$, which sorts \mathbf{X} in ascending order. A corresponding failure pattern \mathbf{f} is obtained by applying π to the vector $(1^{m_1} 2^{m_2})$. Given a permutation π , a *transition* $(i, i+1)$ is defined as a pair of consecutive positions of the failure pattern for which $f_i \neq f_{i+1}$, i.e. a transition identifies consecutive failure instants that correspond to distinct disks. Let $\xi(\pi)$ denote the number of transitions for a given permutation π .

Proposition 1. *The probability of data loss $P_{\mathbf{m}}(\mathcal{D}_t)$ of a $(2, 1)$ -code given $\mathbf{m} = (m_1, m_2)$, m_i the number of failures for disk i , is given by*

$$P_{\mathbf{m}}(\mathcal{D}_t) = 1 - \sum_{j=1}^{s-1} (1 - jt_{\text{rep}}/t)^s \Pr(\xi(\pi) = j), \quad (18)$$

where $s = m_1 + m_2$.

Proof:

$$\begin{aligned} P_{\mathbf{m}}(\mathcal{D}_t^c) &= \sum_{\pi} P_{\mathbf{m}}(\mathcal{D}_t^c | \pi) \Pr(\pi) \\ &= \sum_j \sum_{\pi : \xi(\pi)=j} P_{\mathbf{m}}(\mathcal{D}_t^c | \pi) \Pr(\pi) \\ &\stackrel{(a)}{=} \sum_j \sum_{\pi : \xi(\pi)=j} (1 - jt_{\text{rep}}/t)^s \Pr(\pi) \\ &= \sum_j (1 - jt_{\text{rep}}/t)^s \Pr(\xi(\pi) = j). \end{aligned} \quad (19)$$

In (a) we have used the fact that $P_{\mathbf{m}}(\mathcal{D}_t^c | \pi) = \text{vol } \mathcal{R}_{\pi}^c$, where $\mathcal{R}_{\pi}^c = \{(x_1, x_2, \dots, x_s) : 0 \leq x_1 \leq x_2 \leq \dots \leq x_s \leq t, x_{m+1} - x_m > t_{\text{rep}} \text{ for every transition } (m, m+1)\}$, (20)

and the fact that $\text{vol } \mathcal{R}_{\pi}^c = (1 - jt_{\text{rep}}/t)^s$, when $\xi(\pi) = j$ as will be shown in Sec VI. ■

The following corollary provides the asymptotic behavior when t_{rep}/t is small.

Corollary 1. $\lim_{t_{\text{rep}}/t \rightarrow 0} \frac{P_{\mathbf{m}}(\mathcal{D}_t)}{t_{\text{rep}}/t} = 2m_1 m_2$.

Proof: We have

$$\begin{aligned} \lim_{t_{\text{rep}}/t \rightarrow 0} \frac{P_{\mathbf{m}}(\mathcal{D}_t)}{t_{\text{rep}}/t} &= \lim_{t_{\text{rep}}/t \rightarrow 0} \sum_{j=1}^{s-1} \sum_{i=1}^s \binom{s}{i} \frac{(-1)^{i+1} j^i (t_{\text{rep}}/t)^i \Pr(\xi(\pi) = j)}{t_{\text{rep}}/t} \\ &= s \sum_{j=1}^{s-1} j \Pr(\xi(\pi) = j). \end{aligned}$$

The summation in the last term—the average number of transitions in a permutation—is shown to be equal to $2m_1 m_2 / s$ in Thm. 2 in Sec. IV-B. ■

B. The Reliability of (n, k) MDS Codes: Direct Approach

To state the probability of data loss of an (n, k) code we need some initial definitions. Let X_{i1}, \dots, X_{im_i} be the random failure instants of disk i and let

$$\mathbf{X} = (X_{11}, \dots, X_{1m_1}, X_{21}, \dots, X_{2m_2}, \dots, X_{n1}, \dots, X_{nm_n}). \quad (21)$$

Denote the total number of disk failures in $[0, t]$ by $s := \sum_{i=1}^n m_i$. Given a sample \mathbf{x} drawn from the distribution of \mathbf{X} , we define the *failure pattern* \mathbf{f} as the vector obtained by applying the permutation which sorts \mathbf{x} in ascending order to $(1^{m_1} 2^{m_2} \dots n^{m_n})$ (the ties are broken arbitrarily and associated to events with zero probability). Note that the number of possible orderings of \mathbf{x} , $s!$, is the number of possible failure patterns \mathbf{f} times $m_1! \dots m_n!$. For example, the failure pattern for Fig. 1 would be $(1, 2, 3, 4, 1, 2, 1)$.

Let $b \geq a \geq 1$ be integers. We denote by $[a, b]_{\mathbb{N}}$ the integer interval $\{i \in \mathbb{N} : a \leq i \leq b\}$, define its length to be $b - a$, and make the following definitions:

Definition 1. Cluster $[a, b]_{\mathbb{N}}$: An interval $[a, b]_{\mathbb{N}}$ such that $\{f(i), a \leq i \leq b\}$ contains exactly $n - k + 1$ distinct entries. The **length** of a cluster is the length of the interval $[a, b]_{\mathbb{N}}$.

Definition 2. Tight Cluster: A cluster that does not contain a cluster of shorter length.

Note that a *transition* (in the sense of Section IV-A) corresponds to a tight cluster for a $(2, 1)$ code, which by definition is of length 2.

Definition 3. Minimal Cluster: A cluster of length $n - k$.

A minimal cluster is tight, but not every tight cluster is minimal. Furthermore, a cluster $[a, b]_{\mathbb{N}}$ is *tight* if and only if $f(a)$ and $f(b)$ are distinct, and $\{f(i) : a < i < b\}$ has exactly $n - k - 1$ distinct entries which are distinct from $f(a)$ and $f(b)$.

Example 1. Consider the failure pattern $(1, 2, 3, 4, 1, 1, 2, 1)$ for a $(4, 2)$ code. In this case, $[1, 3]_{\mathbb{N}}, [2, 4]_{\mathbb{N}}, [3, 5]_{\mathbb{N}}, [3, 6]_{\mathbb{N}}, [4, 7]_{\mathbb{N}}$ and $[4, 8]_{\mathbb{N}}$ are clusters. All but $[3, 6]_{\mathbb{N}}$ and $[4, 8]_{\mathbb{N}}$ are tight clusters, while $[1, 3]_{\mathbb{N}}, [2, 4]_{\mathbb{N}}$, and $[3, 5]_{\mathbb{N}}$ are minimal clusters.

Tight clusters correspond to critical successive failures that may cause data loss.

Definition 4. Let $\mathbf{b} = (b_1, \dots, b_l), l \leq s - 1$, be a binary vector. The **restriction** of \mathbf{b} to an interval $[u, v]_{\mathbb{N}}$ is $\mathbf{b}([u, v]_{\mathbb{N}}) = (b_u, \dots, b_{v-1})$.

Definition 5. Region associated with \mathbf{b}

$$\mathcal{R}_{\mathbf{b}} = \left\{ (x_1, \dots, x_s) : \begin{array}{ll} 0 \leq x_1 \leq \dots \leq x_s \leq t & \\ x_{i+1} - x_i < t_{\text{rep}} & \text{if } b_i = 1 \\ x_{i+1} - x_i \geq t_{\text{rep}} & \text{if } b_i = 0 \end{array} \right\}.$$

Remark 1. Often some of the successive differences are unconstrained. For example, if $x_2 - x_1 > t_{\text{rep}}$, and $x_5 - x_4 < t_{\text{rep}}$ and $s = 6$, then \mathbf{b} should be written as $0 * 1 *$, where $*$ in position i indicates that no constraint is imposed between x_{i+1} and x_i . As we will see later, as far as volume calculations are concerned nothing is lost by considering \mathbf{b} to be 01 , i.e. omitting the $*$'s and writing \mathbf{b} as $0^i 1^j$ where i is the number of \geq constraints and j is the number of $<$ constraints.

Definition 6. Fundamental Simplex \mathcal{S} : $\{x : 0 \leq x_1 \leq x_2 \leq \dots \leq x_s \leq t\}$.

Definition 7. Volume Polynomial. Given a subregion $\hat{\mathcal{S}}$ of the fundamental simplex \mathcal{S} , we define volume polynomial $v(\rho) = (s! / t_{\text{rep}}^s) \text{vol } \hat{\mathcal{S}}$, where $\rho := t / t_{\text{rep}}$. If $\hat{\mathcal{S}} = \mathcal{R}_{\mathbf{b}}$, then we will use the notation $v_{\mathbf{b}}(\rho)$. As will be seen later, the volume polynomial depends on \mathbf{b} through the number of constraints. Thus if \mathbf{b} contains i zeros and j ones, corresponding to $i + j$ constraints, we will write $v_{ij}(\rho)$ interchangeably with $v_{\mathbf{b}}(\rho)$.

C. Characterization of data loss event

When there are no consecutive repeated elements in \mathbf{f} , we consider that data loss occurs if there is an ordered sequence of failures $x_i, \dots, x_{i+n-k+1}$ from $n - k + 1$ different disks such that $x_{j+1} - x_j < t_{\text{rep}}$, for all $j = i, \dots, i + n - k + 1$. When

there is at least one repeated number in the failure pattern (for example $\mathbf{f} = (1, 1, 2, 2, 4, 4, 3)$) we assume that there is a data loss event if there exists an ordered sequence (x_i, \dots, x_{i+l}) from more than $n - k + 1$ disks such that $x_{j+1} - x_j < t_{\text{rep}}$.

We have two equivalent characterizations of an error event, given a failure pattern \mathbf{f} :

- (i) A binary vector \mathbf{b} is a **no-error vector** if the restriction of \mathbf{b} to every tight cluster of \mathbf{f} has weight at most $(l - 1)$, where l is the length of that tight cluster.
- (ii) The vector \mathbf{b} is an **error vector** if its restriction to at least one tight cluster of length l has weight l .

Let us call $B_{\mathbf{f}}$ the set of all error vectors \mathbf{b} for a given failure pattern \mathbf{f} .

Example 2. Consider a $(4, 2)$ MDS code with $\mathbf{m} = (2, 2, 1, 1)$ and suppose the failure pattern is 121234. Then $B_{\mathbf{f}}$ consists of the error vectors $**110$, $**011$ and $**111$. Following our convention of dropping the $*$'s and writing \mathbf{b} as $0^i 1^j$ we write $B_{\mathbf{f}} = \{2(0^1 1^2), 0^0 1^3\}$.

From simple observations, one can find the following expression for $P_{\mathbf{m}}(\mathcal{D}_t)$.

Theorem 1. The probability of data loss satisfies

$$P_{\mathbf{m}}(\mathcal{D}_t) = \frac{1}{\rho^s \binom{s}{m_1, m_2, \dots, m_s}} \sum_{\mathbf{f}} \sum_{\mathbf{b} \in B_{\mathbf{f}}} v_{\mathbf{b}}(\rho). \quad (22)$$

Proof: Let $\hat{\mathbf{X}}$ be the random vector associated to the ordered failure times. Let $P(\mathbf{f})$ be the probability that $\hat{\mathbf{X}}$ has pattern \mathbf{f} .

$$\begin{aligned} P_{\mathbf{m}}(\mathcal{D}_t) &= \sum_{\mathbf{f}} P_{\mathbf{m}}(\mathcal{D}_t | \mathbf{f}) P(\mathbf{f}) \\ &= \binom{s}{m_1, \dots, m_n}^{-1} \sum_{\mathbf{f}} P_{\mathbf{m}}(\mathcal{D}_t | \mathbf{f}) \\ &\stackrel{(a)}{=} \binom{s}{m_1, \dots, m_n}^{-1} \sum_{\mathbf{f}} \sum_{\mathbf{b} \in B_{\mathbf{f}}} P_{\mathbf{m}}(\hat{\mathbf{X}} \in \mathcal{R}_{\mathbf{b}}) \\ &\stackrel{(b)}{=} \frac{m_1! m_2! \dots m_n!}{t^s} \sum_{\mathbf{f}} \sum_{\mathbf{b} \in B_{\mathbf{f}}} \text{vol } \mathcal{R}_{\mathbf{b}} \end{aligned}$$

where (a) is due to the characterization of a data loss event, given \mathbf{f} , and (b) follows from the fact that the set of ordered vectors $\hat{\mathbf{X}}$ has volume $t^s / s!$. ■

Thus, to give explicit forms for $P_{\mathbf{m}}(\mathcal{D}_t)$, we need two elements

- (i) Computations of the volume of the error regions $\mathcal{R}_{\mathbf{b}}$, or equivalently, computation of the volume polynomial $v_{\mathbf{b}}(\rho)$.
- (ii) Enumeration of the set of error vectors $B_{\mathbf{f}}$.

The volume computation is addressed in Sec. VI. We address the problem of enumerating the error vectors in this section and use Thm. 8, Sec. VI in order obtain the asymptotic behavior of $P_{\mathbf{m}}(\mathcal{D}_t)$ as $t_{\text{rep}} / t \rightarrow 0$.

Thm. 8, Sec. VI gives a formula for computing $v_b(\rho)$. In particular, it shows that if \mathbf{b} is some permutation of $0^i 1^j$, i.e. $w(\mathbf{b}) = j$, then

$$v_{ij}(\rho) = \frac{s!}{(s-j)!} \rho^{s-j} + O(\rho^{s-j-1}), \quad (23)$$

where s is the number of failures in $[0, t]$. This means that the dominant terms in $P_m(\mathcal{D}_t)$ are when $w(\mathbf{b}) = n - k$. In this case

$$v_{i, n-k}(\rho) = \frac{s!}{(s - (n - k))!} \rho^{s - (n - k)} + O(\rho^{s - (n - k + 1)}). \quad (24)$$

Note also that dominant terms correspond to minimal failure clusters (i.e., of length $(n - k)$). This characterization suffices to prove the asymptotic behavior of $P_m(\mathcal{D}_t)$ as $t_{\text{rep}}/t \rightarrow 0$. Let $j_{f, n-k}$ be the number of minimal failure clusters in \mathbf{f} . We have

$$P_m(\mathcal{D}_t) = \frac{s!}{(s - (n - k))!} \rho^{-(n - k)} \sum_{\mathbf{f}} \frac{j_{f, n-k}}{\binom{s}{m_1, m_2, \dots, m_n}} + O(\rho^{-(n - k + 1)}) \quad (25)$$

Thus

$$\lim_{\rho \rightarrow \infty} P_m(\mathcal{D}_t) \rho^{n-k} = \frac{s!}{(s - (n - k))!} \sum_{\mathbf{f}} \frac{j_{f, n-k}}{\binom{s}{m_1, m_2, \dots, m_n}}. \quad (26)$$

As will be shown later in Corollary 2,

$$\begin{aligned} A_{n-k} &:= \sum_{\mathbf{f}} \frac{j_{f, n-k}}{\binom{s}{m_1, m_2, \dots, m_n}} = \\ &= \frac{(n - k + 1)! (s - (n - k))!}{s!} \sum_{\substack{(i_1, \dots, i_{n-k+1}) \\ \text{distinct}}} m_{i_1} \dots m_{i_{n-k+1}}, \end{aligned} \quad (27)$$

which leads to

$$\boxed{\lim_{\rho \rightarrow \infty} P_m(\mathcal{D}_t) \rho^{n-k} = \frac{(n - k + 1)!}{s!} \sum_{\substack{(i_1, \dots, i_{n-k+1}) \\ \text{distinct}}} m_{i_1} m_{i_2} \dots m_{i_{n-k+1}}.} \quad (28)$$

Remark 2. The contribution to $P_m(\mathcal{D}_t)$ from data loss events related to non-minimal clusters is negligible in the limit $t_{\text{rep}}/t \rightarrow 0$.

Remark 3. A result with very similar flavor was proved in [17, Ch. 6], in spite of the difference between the models. The approximations in [17, Sec. 6.3.2] show that the dominant term in the mean time to data loss is due to a “direct path” of failures from $n - k + 1$ different disks. This is completely analogous to the fact that the dominant term in $P_m(\mathcal{D}_t)$ is due to minimal clusters (i.e., to the probability associated to a succession of failures from exactly $n - k + 1$ disks).

D. Upper Bounding the Error Term

By enumerating all failure patterns, we calculate $P_m(\mathcal{D}_t)$ explicitly. However, combinatorial upper bounds for the error terms may be useful. We derive an asymptotically optimal bound in this subsection.

Given a failure pattern \mathbf{f} , there is an error if the restriction of the vector \mathbf{b} to at least one tight cluster of length l has weight l (see characterization (ii) at the start of Sec. IV-C). Let I_1, \dots, I_p be the tight clusters of \mathbf{f} ($I_j = [a_j, b_j]_{\mathbb{N}}$). Let l_j be the length of the j -th tight cluster.

$$\begin{aligned} P_m(\mathcal{D}_t | \mathbf{f}) &= P(b(I_1) = 1^{l_1} \text{ or } b(I_2) = 1^{l_2} \text{ or } \dots b(I_p) = 1^{l_p}) \\ &\leq \sum_{j=1}^p P(b(I_j) = 1^{l_j}) = \frac{1}{t^s} \sum_{j=1}^p \text{vol } R_{1^{l_j}} \\ &= \sum_{l=n-k}^s j_{f, l} \text{vol } R_{1^l}, \end{aligned} \quad (29)$$

where we define $j_{f, l}$ to be the number of tight clusters of length l in \mathbf{f} . From the above inequality:

$$\begin{aligned} P_m(\mathcal{D}_t) &\leq \frac{m_1! \dots m_n!}{t^s} \sum_{\mathbf{f}} \sum_{l=n-k}^s j_{f, l} \text{vol } R_{1^l} \\ &= \frac{m_1! \dots m_n!}{t^s} \sum_{l=n-k}^s \sum_{\mathbf{f}} j_{f, l} \text{vol } R_{1^l} \\ &= \frac{s!}{t^s} \sum_{l=n-k}^s \text{vol } R_{1^l} \underbrace{\left(\sum_{\mathbf{f}} j_{f, l} / \binom{s}{m_1, \dots, m_n} \right)}_{:= A_l}. \end{aligned} \quad (30)$$

But A_l is the average number of tight clusters of length l . Also note that $l = n - k$ is the dominant term. Hence this upper bound collapses with exact calculation for vanishing time of repair.

The following theorem gives a closed form expression for A_l .

Theorem 2. Let \mathcal{I}_{n-k+1} be the set of all $(n - k + 1)$ -tuples of distinct numbers (i_1, \dots, i_{n-k+1}) , $1 \leq i_j \leq n$.

$$\begin{aligned} A_l &= (s - l) \binom{s}{l + 1}^{-1} \times \\ &\quad \sum_{(i_1, \dots, i_{n-k+1}) \in \mathcal{I}_{n-k+1}} \sum_{\substack{q_i \geq 1 \\ \sum q_i = l - 1}} m_{i_1} m_{i_{n-k+1}} \prod_{j=2}^{n-k-1} \binom{m_{i_j}}{q_j}. \end{aligned} \quad (31)$$

Proof: Given a failure pattern \mathbf{f} , let Y_j , $j = 1, \dots, s - l$, be indicator random variables which are 1 if $[j, j + l]_{\mathbb{N}}$ is a tight cluster and 0 otherwise. We would like to calculate $A_l = E[Y_1 + \dots + Y_{s-l}] = (s - l)E[Y_1]$. But $E[Y_1]$ is the probability

that $[1, l+1]_{\mathbb{N}}$ is a tight cluster of \mathbf{f} . We use Definition 2 (and the corresponding lemma) to calculate this probability. Pick a random pattern (F_1, \dots, F_{l+1}) (there are $\binom{s}{l+1}$ ways of doing so). A tight cluster is formed by choosing two different numbers for endpoints F_1 and F_{l+1} (say i_1 and i_{n-k+1}), and then choosing $(n-k-1)$ other numbers (i_2, \dots, i_{n-k}) to fill the remaining $(l-1)$ positions. If $l > n-k$, some of the numbers will appear more than once in f_2, \dots, f_l . Suppose that i_j appears q_j times (there are $\binom{m_{i_j}}{q_j}$ ways in which this happens). Since i_1 and i_{n-k+1} appear only once, the total choices for the pattern are the product between $m_1 m_{n-k+1}$ and the choices for i_2, \dots, i_{n-k} . Summing over all possible q_j gives us the final answer. ■

Corollary 2. A_{n-k} is given by Equation (27)

We now estimate the probability of data loss (or equivalently, the reliability) of an erasure coded storage system with Poisson failures.

Theorem 3. For Poisson distributed failures with rate parameter λ and constant repair time t_{rep}

$$\lim_{t_{\text{rep}} \rightarrow 0} \frac{P(\mathcal{D}_t)}{t_{\text{rep}}^{n-k}} = \frac{n!}{(k-1)!} \lambda^{n-k+1} t \quad (32)$$

Proof: Let $M_i \sim \text{Poisson}(\lambda t)$ be the random variable associated to the number of failures until time t .

$$\begin{aligned} \lim_{t_{\text{rep}} \rightarrow 0} \frac{P(\mathcal{D}_t)}{t_{\text{rep}}^{n-k}} &= \lim_{t_{\text{rep}} \rightarrow 0} \sum_{\mathbf{m}} \frac{P_{\mathbf{m}}(\mathcal{D}_t)}{t_{\text{rep}}^{n-k}} P(\mathbf{M} = \mathbf{m}) \\ &\stackrel{(a)}{=} \sum_{\mathbf{m}} \lim_{t_{\text{rep}} \rightarrow 0} \frac{P_{\mathbf{m}}(\mathcal{D}_t)}{t_{\text{rep}}^{n-k}} P(\mathbf{M} = \mathbf{m}) \\ &\stackrel{(b)}{=} \frac{(n-k+1)!}{t^{n-k}} \sum_{\mathbf{m}} \sum_{\substack{(i_1, \dots, i_{n-k+1}) \\ \text{distinct}}} m_{i_1} \dots m_{i_{n-k+1}} P(\mathbf{M} = \mathbf{m}) \\ &= \frac{(n-k+1)!}{t^{n-k}} \sum_{\substack{(i_1, \dots, i_{n-k+1}) \\ \text{distinct}}} E[M_{i_1}] \dots E[M_{i_{n-k+1}}] \\ &= (n-k+1)! \binom{n}{n-k+1} \lambda^{n-k+1} t. \end{aligned}$$

Interchanging the limit and summation in (a) is justified by bounded convergence (since $P_{\mathbf{m}}(\mathcal{D}_t)/t_{\text{rep}}^{n-k}$ is naturally uniformly bounded). Step (b) follows from the asymptotics derived in (26). ■

E. Possible Generalizations

The machinery developed in this section has some implications to the reliability of other failure point processes. In a fairly general setup, suppose that the failure mechanism is such that the joint probability density between the random failure instants of disks 1 to n (cf Eq. 21), conditioned on the number of failures, is given by $g(\mathbf{x})$. If $g(\mathbf{x})$ is bounded for all \mathbf{x} , we have the following qualitative result:

Theorem 4. As $t_{\text{rep}} \rightarrow 0$, the probability of data loss $P_{\mathbf{m}}(\mathcal{D}_t)$ of an erasure coded system cannot decay slower t_{rep}^{n-k} .

Proof: The characterization of a data loss event in Sec IV-C does not depend on the statistics of the system. The calculations in Theorem 1 can be thus carried out replacing $P(\tilde{\mathbf{X}} \in \mathcal{R}_b)$ by

$$\int_{\mathcal{R}_b} \hat{g}(\mathbf{x}) d\mathbf{x}, \quad (33)$$

where $\hat{g}(\mathbf{x})$ is the pdf of the order statistics obtained by sorting \mathbf{x} in ascending order. Since we assumed that $g(\mathbf{x})$ is bounded, so is $\hat{g}(\mathbf{x})$, and hence the above integral can be upper bounded by a constant times $\text{vol } \mathcal{R}_b$. The result now follows from the asymptotic analysis of $\text{vol } \mathcal{R}_b$, provided by Eq. (26). ■

Notice that this result does not assume independence on the failure time or invariance under time.

Example 3 (Non-Homogeneous Poisson Processes). *This type of process can model situations such as aging effects and reliability growth. Let $\lambda(x)$ be a function of time (referred to as the rate function). For this process, the probability that there are m failures of one disk in the interval $[a, a+h]$ is given by:*

$$P(M_i(a, a+h) = m) = \frac{(\Lambda(a, a+h))^m \exp[-\Lambda(a+h, a)]}{m!}, \quad (34)$$

where $\Lambda(a, a+h) := \int_a^{a+h} \lambda(x) dx$. Define the normalized rate function as

$$\tilde{\lambda}(x) = \frac{\lambda(x)}{\Lambda(0, t)}. \quad (35)$$

It is not hard to see that the failure times of a disk are independently, identically distributed with pdf $\tilde{\lambda}(x)$ (see, for example, [12, p. 64], adapted to the non-homogeneous case). This way, the joint pdf of the ordered failure times, conditioned on the number of failures is m is given by.

$$g(\hat{x}_1, \dots, \hat{x}_m) = s! \tilde{\lambda}(x_1) \tilde{\lambda}(x_2) \dots \tilde{\lambda}(x_m), \quad (36)$$

Now let $C(t) = \max_{x \in [0, t]} \tilde{\lambda}(x)$. We can bound $P(\tilde{\mathbf{X}} \in \mathcal{R}_b)$ by $s! C(t)^s \text{vol } \mathcal{R}_b$, and thus

$$P_{\mathbf{m}}(\mathcal{D}_t) \leq s! C(t)^s \sum_{\mathbf{f}} \sum_{b \in \mathbf{f}} (\text{vol } \mathcal{R}_b) (P_{\mathbf{m}}(\mathbf{f})). \quad (37)$$

A special case of this process are the “Power-Law Processes”, where $\lambda(x) = \lambda x^\beta$, $\lambda > 0$. In this case, $C(t) = (\beta+1)^{\frac{1}{\beta}} t^{\frac{\beta}{\beta+1}}$.

V. SET AVOIDANCE PROBABILITIES FOR CARTESIAN PRODUCTS OF RANDOM SETS

The closed form calculations performed in the previous sections are particularly useful to characterizing the asymptotic behavior of the system. The objective of this section is to provide a simple upper bound based on Jensen’s inequality. The proofs of the theorems, as well as an upper bound based on the inclusion-exclusion principle, along with a geometric characterization of situations when these bounds are tight, can be found in Appendix B and in [6]. The set avoidance lower bound is used to derive a lower bound on the reliability function in Sec. V-B, some examples are presented in Sec. V-C and general results for (n, k) MDS codes are presented in Sec. V-D.

A. Lower Bounds

Given sets $S_1, S_2, \mathcal{R} \subset S_1 \times S_2$ and $x_1 \in S_1$ we define the shadow of a section of \mathcal{R} as $\mathcal{R}_1(x_1) = \{x_2 \in S_2 : (x_1, x_2) \in \mathcal{R}\}$. In the following, the operator \times has precedence over set operations such \cap and \cup .

Lemma 1. Let $\mathcal{X} := \{X_1, X_2, \dots, X_{m_1}\}$ and $\mathcal{Y} := \{Y_1, Y_2, \dots, Y_{m_2}\}$, where the X_i 's are i.i.d on a set S_1 and the Y_i 's are i.i.d on a set S_2 . Let X, Y be generic random variables distributed as X_i and Y_i , resp. Let $\mathcal{R} \subset S_1 \times S_2$. Then

$$P(\mathcal{X} \times \mathcal{Y} \cap \mathcal{R} = \emptyset) \geq \left(P(\{X\} \times \mathcal{Y} \cap \mathcal{R} = \emptyset)\right)^{m_1} \quad (38)$$

and equality holds iff $P(X \in \bigcup_{i=1}^{m_2} \mathcal{R}(y_i))$ is a constant for $(y_1, y_2, \dots, y_{m_2})$ with positive pmf.

Corollary 3.

$$P(\mathcal{X} \times \mathcal{Y} \cap \mathcal{R} = \emptyset) \geq P((X, Y) \notin \mathcal{R})^{m_1 m_2}. \quad (39)$$

B. Application of Set Avoidance Calculations to the Data Loss Probability Calculation

We first apply the bounds developed in the previous section to derive a lower bound on the reliability function $R(t)$. The bound is given in terms of the volume of the error region associated with a given code. A systematic method for calculating the volume of the error region is then presented along with an overview of some of the theoretical results related to the calculation of an error polynomial associated with the code. Proofs are presented in the next session and in the appendix.

For the avoidance upper bound, we need a different definition of a data loss event. Let \mathbf{f} be a failure pattern. We consider that data loss occurs if there is an ordered sequence of failures $x_i, \dots, x_{i+n-k+1}$ from $n-k+1$ different disks such that $x_{j+1} - x_j < t_{\text{rep}}$, for all $j = i, \dots, i+n-k+1$, even when the failure pattern has repeated consecutive elements. From Remark 2, this characterization is asymptotically the same as the one in IV-C.

Let \mathcal{R} be the region

$$\begin{aligned} \mathcal{R} := \{ & (x_1, x_2, \dots, x_n) \in [0, t]^n : |x_{i_1} - x_{i_2}| < t_{\text{rep}}, \\ & |x_{i_2} - x_{i_3}| < t_{\text{rep}}, \dots, |x_{i_{n-k+1}} - x_{i_n}| < t_{\text{rep}} \\ & \text{for some } i_1, i_2, \dots, i_{n-k+1} \}. \end{aligned} \quad (40)$$

Note that $\mathcal{R} \subset [0, t]^n$ contains the error regions of a code when there is precisely one failure of each disk ($s = m_1 + \dots + m_n = n$ and $m_i = 1$). Suppose that in the interval $[0, t]$, disk i fails $m_i > 0$ times. Let $\mathbf{m} = (m_1, m_2, \dots, m_n)$. The failure instants of the i -th disk are denoted by $\mathcal{X}_i = \{X_{i1}, \dots, X_{im_i}\}$, where the X_{ij} are independently and uniformly drawn on the time interval $[0, t]$. A data loss event occurs if and only if $\mathcal{X}_1 \times \mathcal{X}_2 \times \dots \times \mathcal{X}_n \cap \mathcal{R} \neq \emptyset$, where \mathcal{R} is error region for the code as defined in (40). Let $X_i, i = 1, 2, \dots, n$ be i.i.d. random variables, uniformly distributed on $[0, t]$ and let

$\mathbf{X} = (X_1, X_2, \dots, X_n)$. The following proposition follows immediately from Cor. 3.

Theorem 5. The probability that there is no data loss in the interval $[0, t]$, given m_i , the number of failures for disk i in $[0, t]$, $m_i > 0, i = 1, 2, \dots, n$ satisfies

$$P_m(\mathcal{D}_t^c) \geq P_m(\mathbf{X} \in \mathcal{R}^c)^{m_1 m_2 \dots m_n} = \left(1 - \frac{\text{vol } \mathcal{R}}{t^n}\right)^{m_1 \dots m_n}. \quad (41)$$

Proof: The quantity on the left is $P(\mathcal{X}_1 \times \mathcal{X}_2 \times \dots \times \mathcal{X}_n \cap \mathcal{R} = \emptyset | \mathbf{M} = \mathbf{m})$. Thus the inequality in (41) follows directly from Cor. 3. The equality in (41) follows from the fact that X_i 's are uniform random variables iid over $[0, t]$. ■

We proceed to calculate the volume of \mathcal{R} for a few example codes, and then state a general result.

C. Graphical Representation of Constraints, Some Example Error Regions

In order to help calculate the volume of the error region \mathcal{R} defined in (40) we consider a binary vector $\mathbf{b} = (b_1, b_2, \dots, b_l), l \leq n-1$ and to define $\mathcal{R}_{\mathbf{b}} \in [0, t]^n$ as the region

$$\mathcal{R}_{\mathbf{b}} = \left\{ (x_1, \dots, x_n) \in [0, t]^n : \begin{array}{ll} x_1 \leq \dots \leq x_n & \\ x_{i+1} - x_i \leq t_{\text{rep}} & \text{if } b_i = 1 \\ x_{i+1} - x_i > t_{\text{rep}} & \text{if } b_i = 0 \end{array} \right\}.$$

Note that except for the dimension of the binary vector this definition coincides with Def. 5.

The vector \mathbf{b} is conveniently visualized as a graph $G_{\mathbf{b}}$ with n vertices such that there is an edge between i and $i+1$ iff $b_i = 1$. The region \mathcal{R} can be decomposed into a disjoint union of regions $\mathcal{R}_{\mathbf{b}}$, the union being over all edges \mathbf{b} that are **error vectors**.

In cases where there are no constraints between successive failure instants, the dimension of the vector is reduced and the corresponding graph has fewer nodes. As an example consider an ordered vector of failure instants $\mathbf{x} = (x_1, x_2, x_3, x_4)$ with the constraints $x_2 - x_1 > t_{\text{rep}}, x_3 - x_2 < t_{\text{rep}}$. This constraint is represented by the vector $\mathbf{b} = 01$, and is shown as a graph with three vertices.

A systematic method for calculating the volume of $\mathcal{R}_{\mathbf{b}}$ and hence of \mathcal{R} is presented in Sec. VI. Here we show by example, the error vectors that correspond to specific codes.

Example 4. $(n, n-1)$ -single parity code. In general, if $k = n-1$, any simultaneous two disk failures (within an interval of length t_{rep}) will cause data loss. Therefore

$$\mathcal{R} = \{(x_1, \dots, x_n) \in [0, t]^n : \exists i \neq j \text{ s.t. } |x_i - x_j| \leq t_{\text{rep}}\}, \text{ and}$$

$$\mathcal{R}^c = \{(x_1, \dots, x_n) \in [0, t]^n : |x_i - x_j| > t_{\text{rep}} \text{ for all } i, j\}.$$

Fig. 3 is an illustration of region \mathcal{R}^c in three dimensions ($n = 3, k = 2$). The fact that the above region is a simplex is proved in Appendix A, Lemma 3. It is also proved that

$$\text{vol } \mathcal{R}^c = (t - (n-1)t_{\text{rep}})^n.$$




Figure 3: Illustration of the error region \mathcal{R}^c and no-error region \mathcal{R}^c for the $(3, 2)$ code. Also shown is a single simplex, corresponding to one of the orderings of (x_1, x_2, x_3) . The error region is the ‘star-shaped’ region that is unshaded. The error region is the disjoint union of the polytopes \mathcal{R}_b . The corresponding error graphs are G_b , with error vectors 10, 01 and 11.

Also shown in Fig. 3 is a graphical representation for the error and non-error vectors.

Remark 4. For the analysis above, we require that $t \geq (n-1)t_{\text{rep}}$.

For general codes the no-error regions are not elementary simplices as in a $(n, n-1)$ -code. However, a systematic method for calculating the volume of an error region is presented in Sec. VI.

Example 5. $(4, 2)$ -Code:

The error graphs of a $(4, 2)$ code are represented in Fig. 4. For $t \geq 3t_{\text{rep}}$, the volume of the error region is given by

$$\text{vol } \mathcal{R} = 24t^2t_{\text{rep}}^2 - 72t_{\text{rep}}^3t + 64t_{\text{rep}}^4.$$

Details of the volume calculation are presented in Sec. VI.

D. Set Avoidance Bounds for (n, k) MDS codes

For an (n, k) MDS code, let $\alpha_j(n, k)$ denote the number of error graphs labeled by error vectors \mathbf{b} with Hamming weight j . We define the error polynomial as:

$$e(\rho) = \sum_{j=n-k}^{n-1} \alpha_j(n, k) v_{n-1-j, j}. \quad (42)$$




Figure 4: No-error graphs (green, left) and the error vectors (red, right) for the $(4, 2)$ code. G_{000} , G_{100} , G_{010} , G_{001} and G_{101} represent the no-error region \mathcal{R}_w^c and G_{110} , G_{011} and G_{111} represent the error region, \mathcal{R}_w .

Let $\beta_j(n, k)$ be the number of no-error graphs of weight j for an (n, k) code,

$$\beta_j(n, k) = \binom{n-1}{j} - \alpha_j(n, k),$$

where the term $\binom{n-1}{j}$ is the total number of binary strings of length $n-1$ and Hamming weight j . Analyzing the labels $b_1 \dots b_{n-1}$, it follows that $\beta_j(n, k)$ is the number of binary strings of length $n-1$ and weight j that has no runs of $(n-k)$ or more 1s. This number and its relation with generalizations of the Pascal Triangle was thoroughly studied in [2], [3]. It follows immediately that $\alpha_j(n, k) = 0$, for $j = 0, 1, \dots, (n-k-1)$.

Combining two results from [2, Thm. 3.3] and [3, Eq. 3], we have $\beta_j(n, n-k) = C_{n-k}(n-j, j)$, where $C_m(l, s)$ is the coefficient of x^s in the expansion of the polynomial generating function $(1+x+\dots+x^{m-1})^l$. This leads to the following

Lemma 2. The number of no-error graphs of Hamming weight j of an (n, k) -code is given by

$$\begin{aligned} \beta_j(n, k) &= C_{n-k}(n-j, j) = \\ &= \sum_{i=0}^a (-1)^i \binom{n-j}{i} \binom{n-1-(n-k)i}{n-j-1}, \end{aligned} \quad (43)$$

where $a = \min\{n-j, \lfloor j/(n-k) \rfloor\}$.

We are now in position to prove:

Theorem 6. The error polynomial for an (n, k) MDS code satisfies

$$e(\rho) = \frac{n!}{(k-1)!} \rho^k + O(\rho^{k-1}). \quad (44)$$

Proof: When expressed as a polynomial in ρ , the volume polynomial is given by

$$e(\rho) = \sum_{j=0}^n b_s \rho^s.$$

From Remark 7 which follows Thm. 8, $b_s = 0$ for $s = k+1, \dots, n$, i.e. each volume polynomial in Eq. (42) has degree at most k . In fact, the only polynomial that has degree k is $v_{k-1, n-k}(\rho)$. Also from Remark 7, the coefficient of ρ^k in $v_{k-1, n-k}(\rho)$ is $k! \binom{n}{k}$, whereas from Lemma 2, $\alpha_{n-k}(n, k) = k$. Thus, the highest degree term of $e(\rho)$ is $\alpha_{n-k}(n, k) k! \binom{n}{k} \rho^k = n! / (k-1)! \rho^k$, from where the theorem follows. ■

Corollary 4. *The volume of \mathcal{R} for an (n, k) -code satisfies*

$$\text{vol } \mathcal{R} = \frac{n!}{(k-1)!} t_{\text{rep}}^{n-k} + \sum_{s=0}^{k-1} a_s t_{\text{rep}}^{n-s}, \quad (45)$$

where a_s , $s = 0, \dots, k-1$, are constants.

Remark 5. *When t_{rep}/t is small ($t_{\text{rep}}/t \rightarrow 0$),*

$$\text{vol } \mathcal{R} \approx \frac{n!}{(k-1)!} t_{\text{rep}}^{n-k}. \quad (46)$$

E. Averaging the Set Avoidance Bound for Poisson Failures the Multiplicative Gap

We now evaluate the bound for Poisson failures with rate λ i.e. inter failure durations that are iid exponential with mean $1/\lambda$. We also evaluate the multiplicative gap between the set avoidance upper bound and the asymptotic result (32).

Theorem 7. *Let \mathcal{R}_j be the error region of a $(j, j - (n - k))$ -code, $j \geq n - k + 1$. The probability of data loss of an (n, k) coded is bounded by*

$$P(\mathcal{D}_t) \leq \sum_{j=n-k+1}^n \binom{n}{j} e^{-\lambda t(n-j)} (\lambda t)^j \left(\frac{\text{vol } \mathcal{R}_j}{t^j} \right). \quad (47)$$

Proof: Let $w(t)$ denote the random variable associated to the weight, i.e. the number of disks that failed at least once within $[0, t]$. We have:

$$P(\mathcal{D}_t) = \sum_{j=n-k+1}^n \binom{n}{j} (1 - e^{-\lambda t})^j e^{-\lambda t(n-j)} P(\mathcal{D}_t | w(t) = j).$$

and the RHS of the above equation can be bounded by using lower-dimensional versions of Thm. 5:

$$P(\mathcal{D}_t^c | w(t) = j) \geq \left(1 - \frac{\text{vol } \mathcal{R}_j}{t^j} \right)^{\lambda^j t^j / (1 - e^{-\lambda t})^j}. \quad (48)$$

The proof now follows by bounding (48) using the fact that $(1-x)^r \geq 1-rx$ for any real numbers r, x such that $r \geq 1$ and $0 \leq x \leq 1$. ■

Corollary 5. *Let $P^{(u)}(\mathcal{D}_t)$ be the upper bound in (48). The multiplicative gap between $P^{(u)}(\mathcal{D}_t)$ and $P(\mathcal{D}_t)$ satisfies*

$$\lim_{t_{\text{rep}} \rightarrow 0} \frac{P^{(u)}(\mathcal{D}_t)}{P(\mathcal{D}_t)} = (e^{-\lambda t} + \lambda t)^{k-1} \quad (49)$$

In particular, when $k = 1$ the bound is asymptotically tight.

Proof: From Equation (48) and (46), the ratio $P^{(u)}(\mathcal{D}_t)/t_{\text{rep}}^{n-k}$ is well approximated by

$$\begin{aligned} &\approx \sum_{j=n-k+1}^n \binom{n}{j} e^{-\lambda t(n-j)} \left(\frac{j! t^{j-(n-k)}}{(j-(n-k)-1)! t^j} \right) (\lambda t)^j, \\ &= \frac{n!}{(k-1)!} \lambda^{n-k+1} t \\ &\times \left[\sum_{j=n-k+1}^n \binom{k-1}{n-j} e^{-\lambda t(n-j)} (\lambda t)^{j-(n-k+1)} \right] \end{aligned}$$

and the approximation is tight as $t_{\text{rep}} \rightarrow 0$. Using Theorem 3 and after some algebraic manipulation, we conclude (49). ■

VI. VOLUME CALCULATIONS FOR ORDERED SETS WITH CONSTRAINED DIFFERENCES

Both of the approaches presented for estimating the data loss probability, the direct approach of Sec. IV-B and the bounds based on set avoidance presented in Sec. V ultimately rely on the methods for volume calculation presented in this section. The calculations presented here are for an ordered s -tuple, where s is a dummy variable, no longer necessarily associated with the number of disks failures in the interval $[0, t]$. In order to apply the results to Sec. IV-B, s will indeed represent the total number of failures that occur in the interval $[0, t]$, whereas in order to apply the results to Sec. V, s will be replaced by n , the number of disks in the system. The results in this generalizes the formulas in [5] for any (n, k) and provides the exact behavior of such formulas.

The volume of the error region can be determined by splitting it into disjoint simplices. Since, by definition, the region \mathcal{R} is symmetric with respect to different orderings of the failures, we have $\text{vol } \mathcal{R} = s! \text{vol } (R \cap S)$. We can thus restrict our analyses to ordered vectors $x_1 \leq x_2 \leq \dots \leq x_s$. The volume of the regions restricted to the ordered simplex is now presented.

We first observe that $\text{vol } \mathcal{R}_b$ only depends on the weight (number of nonzero entries) of b (see Lemma 3 and the remarks that follow in the Appendix). Thus it suffices to study graphs of the form $G_{0^i 1^j}$, where j is the weight of the vector b . We will work with volume polynomials $v_{ij}(\rho)$, a scaled version of the volume of the region $\mathcal{R}_{0^i 1^j}$, where for convenience we repeat that $\rho = t/t_{\text{rep}}$ and $v_{ij}(\rho)$ associated with $\mathcal{R}_{0^i 1^j}$ is given by $v_{ij}(\rho) = s! \text{vol } G_{0^i 1^j} / t_{\text{rep}}^s$.

We prove in Appendix A that \mathcal{R}_{0^i} is a simplex with volume $(t - it_{\text{rep}})^s / s!$, provided $t \geq it_{\text{rep}}$. Alternatively, $v_{i0} = (\rho - i)^s$. For instance $v_{00}(\rho) = \rho^s$ is the volume polynomial of the region with no constraints on the differences $x_{i+1} - x_i$. Since the union of a region such that $x_{i+1} - x_i \leq t_{\text{rep}}$ and another one such that $x_{i+1} - x_i \geq t_{\text{rep}}$ gives a region with no

constraints on $x_{i+1} - x_i$, we have the following “difference” identity:

$$v_{i,j}(\rho) = v_{i,j-1}(\rho) - v_{i+1,j}(\rho).$$

Summarizing, the following rules provide a systematic method for calculating the volume polynomials associated with any node in the supergraph.

- (Shift) $v_{i+1,j}(\rho) = v_{i,j}(\rho - 1)$, $j = 0, 1, 2, \dots$,
- (First Difference) $v_{i,j}(\rho) = v_{i,j-1}(\rho) - v_{i+1,j}(\rho)$,
- (Initial Condition) $v_{00} = \rho^s$.




Figure 5: Supergraph representation of the set of graphs $G_{0^{i1j}}$ for the $(4, 2)$ code. A vertex with label ij represents the graph $G_{0^{i1j}}$ (depicted on the left of each node of the supergraph). Note that the number of constraints increases from zero at the top layer of the supergraph to three at the bottom layer of the supergraph.

The graphs $G_{0^{i1j}}$ are conveniently organized into a *supergraph*, as illustrated in Fig. 5, in order to facilitate computation of the volume polynomials. In this graph, each node is associated with a volume polynomial. For example, the top or root node in Fig. 5 is associated with the volume polynomial $v_{00}(\rho)$ and the polynomial associated with the graph G_{011} is v_{12} .

We revisit the $(4, 2)$ MDS code and compute the volume of the error region.

Example 6 ($(4, 2)$ -Code). *The error vectors of a $(4, 2)$ code are represented in Fig. 4. Summing the volume polynomial corresponding to all error vectors and considering all orderings of the vector (x_1, \dots, x_n) we obtain*

$$\begin{aligned} \frac{1}{t_{rep}^4} \text{vol } \mathcal{R}_w &= 2v_{12}(\rho) + v_{03}(\rho) \stackrel{(a)}{=} v_{00} - v_{10} - v_{20} + v_{30} \\ &= \rho^4 - (\rho - 1)^4 - (\rho - 2)^4 + (\rho - 3)^4 \\ &= 24\rho^2 - 72\rho + 64, \end{aligned} \quad (50)$$

where in (a) we applied the first difference rule and (b) is a combination of the shift rule and the initial condition. This gives us, for $t \geq 3t_{rep}$,

$$\text{vol } \mathcal{R} = 24t_{rep}^2 t^2 - 72t_{rep}^3 t + 64t_{rep}^4.$$

In the following two examples, we calculate $\sum_{b \in \mathcal{B}_f} v_b(\rho)$ in (22), related to the direct calculation of the data loss probability.

Example 7. Suppose $\mathbf{m} = (1, 1, 1, 1)$ and a $(4, 2)$ MDS code is used. Consider the failure pattern $\mathbf{f} = 1234$. Then $\mathcal{B}_f = \{2(0^1 1^2), 0^0 1^3\}$ and $s = 4$. From this we can write down the volume polynomial as $2v_{12}(\rho) + v_{03}(\rho)$. Upon simplification we obtain $v_{00}(\rho) - v_{10}(\rho) - v_{20}(\rho) + v_{30}(\rho) = \rho^4 - (\rho - 1)^4 - (\rho - 2)^4 + (\rho - 3)^4 = 24\rho^2 - 72\rho + 64$.

Remark 6. Observe that the volume polynomials for Ex. 6 and Ex. 7 are identical.

Another example related to the direct calculation.

Example 8. Suppose $\mathbf{m} = (2, 2, 1, 1)$ and a $(4, 2)$ MDS code is used. Consider the failure pattern $\mathbf{f} = 121234$. Then $\mathcal{B}_f = \{2(0^1 1^2), 0^0 1^3\}$ and $s = 6$. From this we can write down the volume polynomial as $2v_{12}(\rho) + v_{03}(\rho)$. Upon simplification we obtain $v_{00}(\rho) - v_{10}(\rho) - v_{20}(\rho) + v_{30}(\rho) = \rho^6 - (\rho - 1)^6 - (\rho - 2)^6 + (\rho - 3)^6 = 60\rho^4 - 360\rho^3 + 960\rho^2 - 1260\rho + 664$.

The following lemma uses the aforementioned rules to provide closed form expressions for $v_{ij}(\rho)$.

Theorem 8. The volume polynomial $v_{ij}(\rho) = \sum_r a_r \rho^r$ satisfies the following properties (i)

$$v_{ij}(\rho) = \sum_{l=0}^j (-1)^{j-l} \binom{j}{l} (\rho - i - j + l)^s. \quad (51)$$

(ii)

$$a_r = \binom{s}{r} j! (-1)^{s-r+j} \left(\sum_{m=0}^{s-j} \binom{s-r}{m} i^m S(s-r-m, j) \right), \quad (52)$$

where $S(l, m)$ is a Stirling number of the second kind (see, e.g., [4]).

Proof: Given a function $f(x) : \mathbb{R} \rightarrow \mathbb{R}$, define the shift operator $S(f(x)) := f(x - 1)$ and the first difference operator $\Delta(f(x)) := f(x) - f(x - 1)$. Then (i) follows from the observation that $v_{ij}(\rho) = S^i \Delta^j(\rho^s)$. Write $S = (1 - \Delta)$ in order to express the operator in terms of powers of Δ . This gives

$$v_{ij}(\rho) = \left(\sum_{l=0}^i (-1)^{i-l} \binom{i}{l} \Delta^{i+j-l} \right) (\rho^s). \quad (53)$$

To prove (ii), expand the last term in (51) and interchange the order of summation, so that

$$v_{ij}(\rho) = \sum_{m=0}^n \rho^m \binom{s}{m} (-1)^{s-m} \sum_{l=0}^j (-1)^l \binom{j}{l} (i+l)^{s-m}. \quad (54)$$

The result follows directly by further expanding the last term in the above equation and from an identity for Stirling numbers of the second kind (e.g. Prop. 5.3.5, [4]). ■

Remark 7. $\deg(v_{ij}(\rho)) = s - j$ and $a_{s-j} = j! \binom{s}{j}$.

n	k	t	κ_f	κ_r	$1/\lambda$	$1/\mu$	$P(D_t)$ (16)	$P(D_t)$ (Simulation)	Standard Deviation
4	2	1.0	1.5	2.0	0.1	0.001	3.343×10^{-6}	3.429×10^{-6}	4.07×10^{-7}
4	2	1.0	0.75	2.0	0.1	0.001	0.0044	0.0044	5.38×10^{-4}
4	2	1.0	0.75	0.75	0.1	0.001	0.0035	0.0036	1.94×10^{-4}
4	2	1.0	0.75	0.75	0.1	10^{-6}	1.185×10^{-7}	1.221×10^{-7}	1.22×10^{-8}
8	5	1.0	0.75	1.25	0.001	10^{-6}	8.9289×10^{-5}	8.8383×10^{-5}	1.2397×10^{-5}
8	5	1.0	2.0	2.0	0.01	0.001	3.981×10^{-5}	4.012×10^{-5}	1.548×10^{-6}
8	5	1.0	0.5	2.0	0.01	10^{-6}	1.013×10^{-4}	1.008×10^{-4}	2.766×10^{-6}

Table I: Validation of Eqn. (16) through simulation. t is the observation time window, $(\kappa_f, 1/\lambda)$ and $(\kappa_r, 1/\mu)$ are the (shape, mean) values for the Weibull distributed failure and repair durations, resp.




Figure 6: $G = P(Y < Z)$ for Weibull distributed failures and repairs for $(1/\lambda, 1/\mu) = (0.25, 10^{-4})$ (left), and $(1/\lambda, 1/\mu) = (461386, 12)$ (right). The horizontal line is the value of G for exponentially distributed repair and failure distributions.

VII. NUMERICAL AND SIMULATION RESULTS

The data loss probability expression (16) for general distributions was validated through simulation for the family of Weibull distributions, with pdf given by

$$f(x) = \frac{\kappa}{\alpha} \left(\frac{x}{\alpha} \right)^{\kappa-1} e^{-(x/\alpha)^\kappa}, \quad (55)$$

where κ is the *shape parameter* and α is the *scale parameter*. Its mean is given by $\alpha\Gamma(1 + 1/\kappa)$, where Γ is the Gamma function. The exponential distribution is a special case of this family, for $\kappa = 1$ and $\alpha = 1/\lambda$. A few results are tabulated in Table I. The agreement between theory and simulation is

seen to be close. The simulator used here generates a sequence of failure and repair durations according to the specified distribution, calculates runs of the event $Y < Z$ and for each run generates disk labels drawn uniformly and iid on $\{1, 2, \dots, n\}$. The standard deviation of the estimates reported in Table I was varied by changing the number of independent experiments.

The probability $G = P(Y < Z)$ (defined immediately following (7)) that a failure will occur prior to the completion of a repair is sensitive to the shape of the distribution. This is illustrated, in Fig. 6, for the family of Weibull distributions. In the plots below, it is assumed that failures are Weibull distributed with mean failure duration set to $1/\lambda = 461386$ and mean repair duration $t_{\text{rep}} = 1/\mu$ set to 12.0, values that were obtained by Elerath and Pecht in an experimental study of disk failures [9]. For a fixed value of κ_r , the shape parameter for the Weibull repair duration distribution with mean t_{rep} , the shape parameter κ_f of the failure duration is varied, and the scale parameter is adjusted to keep the mean constant at $1/\lambda$. As κ_f varies the value of G is seen to change significantly. This plot also shows the value of G for exponentially distributed failure and repair durations. The value of G is seen to be less sensitive to the shape parameter of the repair distribution.




Figure 7: Probability of data loss for time window of duration $t = 1$, when (failures, repairs) are both Weibull with $(\kappa_f, \kappa_r) = (0.5, 2.0)$ and when (failures, repairs) are exponentially distributed. In all cases the mean repair duration is 10^{-4} . The scale parameters of the Weibull distributions are adjusted to keep the mean failure duration $1/\lambda$ and mean repair duration constant.

In Fig. 7 the performance of a $(4, 2)$ code is compared for Weibull and exponential distributions for fixed mean repair and failure durations. As already observed, the value of G is seen to depend on the shape parameter, and the impact on the data loss probability is magnified by redundancy $(n-k)$ of the code. The predicted gap in reliability is verified by the simulation. The impact on the reliability of the system especially for a powerful code can be quite dramatic—an order of magnitude difference in G becomes ten orders of magnitude for a code with 10 check symbols. It is also clear that the mean time between failures for individual disks is not a sufficient determinant of overall system reliability.




Figure 8: Simulation results for exponential, constant and Weibull distributions, for $(n, k) = (4, 2)$. Here $t = 1$, the mean repair time is set to 0.1, and the mean time to failure is assumed to be exponential with parameter λ between 0.1 and 0.4. For the Weibull distribution, $\kappa = 0.5$ and $\beta = t_{\text{rep}}/2$ (so its mean is t_{rep}).

In Fig. 8 we show that constant repair duration represents the worst case among the distributions considered. The simulation was carried out for a $(4, 2)$ -code, for fixed mean repair time ($t_{\text{rep}} = 0.1$), and exponential failure times with mean $1/\lambda$. For the Weibull distribution, the shape parameter was chosen to be equal to 0.5 and the scale parameter $t_{\text{rep}}/2$, so that the mean equals t_{rep} . Simulations were based on 10^7 samples for each value of λ , using the algorithm of [15] for the failure process.

Also in Fig. 8, for the constant-repair-duration case, we compared the simulation results and the upper bound, and a good agreement between both in cases where λt is close to unity. It is to be noted that this is the case in many practical situations. We stress the fact that in this case we have a *true* upper bound on the reliability, whereas other methods proposed in the literature only provide estimates. We also caution the reader that the gap between the upper bound and the simulation results can be large when the product λt is large, as we have observed earlier in our discussion about the multiplicative gap.

For exponential failure and repair duration distributions it is shown in [14], based on the results in [7] (see also [17, Eq. 6.69]), that the average time until a data loss event for an

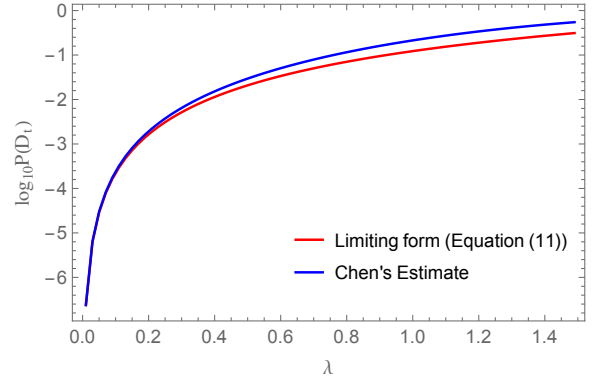


Figure 9: Comparison between the limiting form (16) for the exponential distribution and Chen's estimate with mean repair time $1/\mu = t_{\text{rep}} = 0.1$, $t = 1$ and $(n, k) = (4, 2)$.

(n, k) MDS erasure code is given by

$$MTTDL_c = \frac{1}{\lambda^{n-k+1} t_{\text{rep}}^{n-k}} \frac{(k-1)!}{n!}. \quad (56)$$

From this we get the approximation $\hat{P}(D_t) = 1 - \hat{R}(t) = 1 - e^{-t/MTTDL_c}$. A first order approximation of $P(D_t)$ is

$$\hat{P}(D_t) \approx \frac{n!}{(k-1)!} \lambda^{n-k+1} t_{\text{rep}}^{n-k} t \quad (57)$$

A comparison of (16) for exponential failure and repair distributions and (57) is provided in Fig 9. While there is close agreement in the limit, a deviation is observed for moderately large probabilities.

VIII. SUMMARY, CONCLUSIONS AND FUTURE WORK

We have addressed the problem of directly evaluating the probability of data loss in an erasure coded distributed data storage system. A formula is derived for general iid failure and repair duration distributions using combinatorial methods. For the case where the repair duration is constant, we develop a combinatorial-geometric approach that enables us to directly calculate and bound the data loss probability, in contrast to widely used methods that estimate the integral of the reliability function. Further, our analysis is more refined, in the sense that we are able to derive expressions for the data loss probability conditioned on the number of failures in a given time window.

Our analytic results for general distributions indicate that G , the probability that a failure duration is smaller than a repair duration, is sufficient for characterizing the data loss probability for highly reliable systems. In particular, distributions with the same mean failure and repair durations are seen to exhibit a wide range of values for G and hence data loss probability. This provides motivation for studying, in addition to erasure coding strategies, the impact of networking technologies, and network design, as well as the impact of physical component design (mechanical components in disk drives) with G as a figure of merit.

Finally, we mention that the set avoidance bound misses some of the subtle correlations between n -tuples of failure times. Analytic methods for taking these correlations into account should help to improve this bound. We leave this to future work.

IX. ACKNOWLEDGEMENT

We thank the reviewers and the AE for carefully reading the manuscript and for the numerous suggestions that have resulted in many improvements. This work was begun while the second author was visiting AT&T Labs-Research in 2013. The authors acknowledge with appreciation Prof. Sueli Costa for enabling the visit of the first author (VV) to Campinas in 2014.

APPENDIX A COMBINATORIAL PRELIMINARIES

A. Properties of the Error Polytope (Sec. VI)

We now provide formal justification of the general rules for calculating the volume polynomial $v_{ij}(\rho)$. We start with some observations on the error region and error graphs.

Lemma 3. *Let j_0, \dots, j_i be integers such that $0 = j_0 < j_1 < j_2 < \dots < j_i < s$. Consider the region*

$$\mathcal{R} = \left\{ (x_1, \dots, x_s) : \begin{array}{l} 0 \leq x_1 \leq \dots \leq x_s \leq t \\ x_{j_l+1} - x_{j_l} \geq t_{\text{rep}}, \quad l = 1, \dots, i \end{array} \right\}. \quad (58)$$

We have $\text{vol } \mathcal{R} = (t - it_{\text{rep}})^s / s!$.

Proof: Consider the translation $\phi(\mathbf{x}) = \mathbf{x} - \mathbf{u}$, where $\mathbf{u} = (u_1, u_2, \dots, u_s)$ defined as

$$u_i = lt_{\text{rep}}, \text{ if } i = j_l + 1, \dots, j_{l+1}.$$

Let $y = \phi(\mathbf{x})$. The translated region $\phi(\mathcal{R})$ is given by:

$$\phi(\mathcal{R}) = \left\{ (y_1, \dots, y_s) : \begin{array}{l} 0 \leq y_1 + u_1 \leq \dots \leq y_s + u_s \leq t \\ y_{j_l+1} - y_{j_l} \geq 0, \quad l = 1, \dots, i \end{array} \right\}.$$

Eliminating redundant inequalities we obtain

$$\phi(\mathcal{R}) = \{(y_1, \dots, y_s) : 0 \leq y_1 \leq y_2 \leq \dots \leq y_s \leq t - it_{\text{rep}}\}.$$

This last set of inequalities corresponds to a well-known regular simplex whose volume is $(t - it_{\text{rep}})^s / s!$, concluding the proof. ■

In particular, Lemma 3 shows that the volume of a polytope $\mathcal{R}_{\mathbf{b}}$ defined by an vector \mathbf{b} , depends only on its weight.

Lemma 4. *Let $1 \leq i \leq s - 1$. The volume polynomial associated with the i th node along the left boundary of the super-graph is given by:*

$$v_{i0}(\rho) = (\rho - i)^s. \quad (59)$$

Proof: Recall that, by definition, $v_{i0}(\rho) = s! \text{vol } G_{0^i} / t_{\text{rep}}^s$, where $\rho = t / t_{\text{rep}}$. Thus the statement is equivalent to $\text{vol } G_{0^i} = (t - it_{\text{rep}})^s / s!$, which, in turn, is a special case of Lemma 3, for $j_l = l, l = 1, \dots, i$. ■

APPENDIX B SET AVOIDANCE BOUNDS

Proof of Lemma 1:

$$\begin{aligned} P(\mathcal{X} \times \mathcal{Y} \cap \mathcal{R} = \emptyset) &= E \left(P \left(\mathcal{X} \cap \bigcup_{i=1}^{m_2} \mathcal{R}(y_i) = \emptyset \mid \mathcal{Y} \right) \right) \\ &= E \left(P \left(X \notin \bigcup_{i=1}^{m_2} \mathcal{R}(y_i) \mid \mathcal{Y} \right)^{m_1} \right) \\ &\stackrel{(a)}{\geq} E \left(P \left(X \notin \bigcup_{i=1}^{m_2} \mathcal{R}(y_i) \mid \mathcal{Y} \right) \right)^{m_1} \\ &= \left(P(\{X\} \times \mathcal{Y} \cap \mathcal{R} = \emptyset) \right)^{m_1}, \end{aligned}$$

where in (a) we have used Jensen's inequality. The condition for equality follows directly from the condition for equality in Jensen's inequality. □

In general we do not expect the condition for equality to hold, except in the case where one of the random sets has a single element.

For the next upper bound, we use the following generalized version of the union bound: if A_1, A_2, \dots, A_m are m events, then the probability of $\bigcup_{i=1}^m A_i$ is lower bounded by

$$P \left(\bigcup_{i=1}^m A_i \right) \geq \sum_{i=1}^m P(A_i) - \sum_{j=i+1}^m \sum_{i=1}^m P(A_i \cap A_j). \quad (60)$$

In what follows we denote the event $\{(X, Y) \in \mathcal{R}\}$ by $\varepsilon(X, Y)$.

Theorem 9. *Let $Q_1(x) = P(\varepsilon(x, Y))$ and $Q_2(y) = P(\varepsilon(X, y))$. The set avoidance probability is upper bounded by*

$$\begin{aligned} P(\mathcal{X} \times \mathcal{Y} \cap \mathcal{R} = \emptyset) &\leq 1 - m_1 m_2 P(\varepsilon(X, Y)) + \\ &+ 2 \binom{m_1}{2} \binom{m_2}{2} P(\varepsilon(X, Y))^2 + \\ &+ m_2 \binom{m_1}{2} E[Q_1(X)^2] + m_1 \binom{m_2}{2} E[Q_2(Y)^2] \end{aligned}$$

Proof: First note that

$$P(\mathcal{X} \times \mathcal{Y} \cap \mathcal{R} = \emptyset) = 1 - P \left(\bigcup_{j=1}^{m_2} \bigcup_{i=1}^{m_1} \varepsilon(X_i, Y_j) \right). \quad (61)$$

From Eq. (60), the RHS of (61) can be lower bounded

$$1 - m_1 m_2 P(\varepsilon(X, Y)) + \sum P(\varepsilon(X_i, Y_j) \cap \varepsilon(X_{i'}, Y_{j'})),$$

where the summation is over all $\binom{m_1 m_2}{2}$ distinct choices of cross terms $\varepsilon(X_i, Y_j) \cap \varepsilon(X_{i'}, Y_{j'})$. Now, for the probability of the cross terms, we have three cases. If $i \neq i'$ and $j \neq j'$

j' then, due to independence, $P(\varepsilon(X_i, Y_j) \cap \varepsilon(X_{i'}, Y_{j'})) = P(\varepsilon(X, Y))^2$. On the other hand, if $i = i'$ (and $j \neq j'$) let $f(x_i) = P(\varepsilon(X_i, Y_j) \cap \varepsilon(X_i, Y_{j'}) | X_i = x_i) = Q_1(x_i)^2$. Then:

$$P(\varepsilon(X_i, Y_j) \cap \varepsilon(X_i, Y_{j'})) = E[f(X)] = E[Q_1(X)^2].$$

The case $j = j'$ is analogous. Counting the number of occurrences of the three cases leads us to the theorem. ■

If X and Y are uniformly distributed over a set $\mathcal{S} = \mathcal{S}_1 = \mathcal{S}_2$, the functions Q_1, Q_2 have a natural geometric interpretation, as can be seen in the next example.

Example 9. Let $\mathcal{R} = \{(x, y) \in [0, 1]^2 : |x - y| \leq t_{\text{rep}}\}$ be the error region of a $(2, 1)$ -code, and consider X and Y uniformly distributed over $[0, t]$. Then $P(\varepsilon(X, Y)) = (2t_{\text{rep}}t - t_{\text{rep}}^2)/t^2$. The function $Q_1(x)$ corresponds to the probability that Y belongs to the shadow of $\mathcal{R}_1(x)$ on the y -axis, which, in this case, is the length of $\mathcal{R}_1(x)$. One can easily see that $Q_1(x) \leq 2t_{\text{rep}}$, thus $E[Q_1(X)^2] \leq 4t_{\text{rep}}^2$. By symmetry, $E[Q_2(X)^2] \leq 4t_{\text{rep}}^2$.




Figure 10: Region \mathcal{R} and the shadow of $\mathcal{R}_1(x)$ on the y -axis

Assume that X_1, \dots, X_m and Y_1, \dots, Y_m are iid with the same distribution as X and Y . Applying Corollary 3, we have

$$P(\mathcal{X} \times \mathcal{Y} \cap \mathcal{R} = \emptyset) \geq (1 - t_{\text{rep}}/t)^{2m_1m_2} \geq 1 - 2m_1m_2 \frac{t_{\text{rep}}}{t}.$$

On the other hand, Thm. 9 together with $E[Q_1(X)^2] = E[Q_2(Y)^2] \leq 4t_{\text{rep}}^2$, provides us an upper bound of the type

$$P(\mathcal{X} \times \mathcal{Y} \cap \mathcal{R} = \emptyset) \leq 1 - 2m_1m_2 \frac{t_{\text{rep}}}{t} + \frac{2m_1m_2(m_1m_2 - 1)t_{\text{rep}}^2}{t^2} + o((t_{\text{rep}}/t)^2)$$

From this, we can estimate the gap between upper and lower bounds, and obtain the same asymptotic result as in Cor. 1.

A more general upper bound can be found in [6]. However the upper bound is not optimal, in the sense that it does not collapse with the lower bound for small t_{rep} .

REFERENCES

- [1] J.E. Angus. On computing MTBF for a k-out-of-n:G repairable system. *IEEE Transactions on Reliability*, 37(3):312–313, 1988.
- [2] R. C. Bollinger. Fibonacci k-sequences, Pascal-T triangles, and k-in-a-row problems. *Fibonacci Quarterly*, 2(22):146–151, 1984.

- [3] R. C. Bollinger. Extended pascal triangles. *Mathematics Magazine*, 66(2):pp. 87–94, 1993.
- [4] P. J. Cameron. *Combinatorics: Topics, Techniques, Algorithms*. Cambridge University Press, 1994.
- [5] A. Campello and V.A. Vaishampayan. Reliability of erasure coded storage systems: A geometric approach. In *2013 IEEE International Conference on Big Data*, pages 12–16, Oct 2013.
- [6] A. Campello and V.A. Vaishampayan. Set avoidance probabilities and bounds on the reliability of erasure coded storage systems. In *2014 IEEE Information Theory Workshop (ITW)*, pages 616–620, Nov 2014.
- [7] P. M. Chen, E. K. Lee, G. A. Gibson, R. H. Katz, and D. A. Patterson. RAID: High-performance, reliable secondary storage. *ACM Computing Surveys (CSUR)*, 26(2):145–185, 1994.
- [8] A. G. Dimakis, P. B. Godfrey, Y. Wu, M. J. Wainwright, and K. Ramchandran. Network coding for distributed storage systems. *IEEE Transactions on Information Theory*, 56(9):4539–4551, 2010.
- [9] J. G. Elerath and M. Pecht. Enhanced reliability modeling of raid storage systems. In *Dependable Systems and Networks, 2007. DSN'07. 37th Annual IEEE/IFIP International Conference on*, pages 175–184. IEEE, 2007.
- [10] R. G. Gallager. *Discrete stochastic processes*. Kluwer Academic Publishers, 1996.
- [11] K. M. Greenan. *Reliability and power-efficiency in erasure-coded storage systems*. PhD thesis, University of California, Santa Cruz, 2009.
- [12] L. Kleinrock. *Queueing Systems. Volume 1: Theory*. Wiley-Interscience, 1975.
- [13] E. Pinheiro, W.-D. Weber, and L. A. Barroso. Failure trends in a large disk drive population. In *FAST*, volume 7, pages 17–23, 2007.
- [14] J. K. Resch and I. Volvovski. Reliability models for highly fault-tolerant storage systems. Technical report, Cleversafe Corp., Chicago, IL, USA, 2011.
- [15] B. Sasidharan and P. V. Kumar. High-rate regenerating codes through layering. *arXiv preprint arXiv:1301.6157*, 2013.
- [16] B. Schroeder and G. A. Gibson. Disk failures in the real world: What does an mttf of 1, 000, 000 hours mean to you? In *FAST*, volume 7, pages 1–16, 2007.
- [17] V. Venketasan. *Reliability Analysis of Data Storage Systems*. PhD thesis, Ecole Polytechnique Federale De Lausanne, September 2012.
- [18] Q. Xin, E. L. Miller, T. Schwarz, D. D. E. Long, S. A. Brandt, and W. Litwin. Reliability mechanisms for very large storage systems. In *Mass Storage Systems and Technologies, 2003.(MSST 2003). Proceedings. 20th IEEE/11th NASA Goddard Conference on*, pages 146–156. IEEE, 2003.