



Practical Design Considerations for Wide Locally Recoverable Codes (LRCs)

Saurabh Kadekodi, Shashwat Silas, David Clausen, and
Arif Merchant, *Google*

<https://www.usenix.org/conference/fast23/presentation/kadekodi>

This paper is included in the Proceedings of the
21st USENIX Conference on File and
Storage Technologies.

February 21–23, 2023 • Santa Clara, CA, USA

978-1-939133-32-8

Open access to the Proceedings
of the 21st USENIX Conference on
File and Storage Technologies
is sponsored by

**NetApp**[®]

Practical Design Considerations for Wide Locally Recoverable Codes (LRCs)

Saurabh Kadekodi*, Shashwat Silas*, David Clausen, Arif Merchant
Google

Abstract

Most of the data in large-scale storage clusters is erasure coded. At exascale, optimizing erasure codes for low storage overhead, efficient reconstruction, and easy deployment is of critical importance. *Locally recoverable codes (LRCs)* have deservedly gained central importance in this field, because they can balance many of these requirements. In our work we study wide LRCs; LRCs with large number of blocks per stripe and low storage overhead. These codes are a natural next step for practitioners to unlock higher storage savings, but they come with their own challenges. Of particular interest is their *reliability*, since wider stripes are prone to more simultaneous failures.

We conduct a practically-minded analysis of several popular and novel LRCs. We find that wide LRC reliability is a subtle phenomenon that is sensitive to several design choices, some of which are overlooked by theoreticians, and others by practitioners. Based on these insights, we construct novel LRCs called *Uniform Cauchy LRCs*, which show excellent performance in simulations, and a 33% improvement in reliability on unavailability events observed by a wide LRC deployed in a Google storage cluster. We also show that these codes are easy to deploy in a manner that improves their robustness to common maintenance events. Along the way, we also give a remarkably simple and novel construction of distance optimal LRCs (other constructions are also known), which may be of interest to theory-minded readers.

1 Introduction

Large-scale storage clusters currently house exabytes of data, the bulk of which is encoded with erasure codes. With storage devices (hard-disk drives, or just disks) routinely becoming unavailable (due to maintenance or even failures), using erasure coding of some variety is essential to provide acceptable data durability. But this durability comes with the additional storage overhead incurred by erasure codes. At a time when data corpus size is growing exponentially [9, 43], reducing this storage overhead is essential. One way to accomplish this is to utilize wider stripes for encoding i.e. codes that have a higher ratio of data blocks to coded/redundancy blocks. Such codes are sometimes called ‘wide codes’ since they require the overall width of the stripe to be larger (width is also referred to as *blocklength* in the coding theory literature). To unlock large storage savings without compromising reliability,

*Equal contribution




Figure 1: We captured a sample of 278 unavailable stripes captured from four Google storage clusters, along with information about the exact block failures in each sample. The deployed code had total width $n \approx 50$, and always succeeded in recovering data when there were ≤ 6 failures. We then test these failure scenarios with the Uniform Cauchy LRC of the same width and overhead. The deployed code could not recover any of the 278 stripes before restoration, whereas Uniform Cauchy LRC simulation was successful in recovering 92 stripes prior to restoration; a success ratio of 33%.

codes of larger widths like 20 [3] have been deployed, and in this work we present some data from a Google storage cluster using an erasure code of width ≈ 80 blocks. The drawback of most wide codes with low overhead is that they may require a large amount of IO to reconstruct any unavailable or lost data. This is why wide codes are usually designed as *locally recoverable codes (LRCs)* (Definition 5.3), which help mitigate this reconstruction cost in most cases. Wide LRCs can be utilized to balance the challenging storage needs of today (low storage overhead, competitive reliability, competitive average reconstruction cost), but unique challenges arise when designing *wide* LRCs for deployment in storage clusters.

LRCs are optimized to shine in the case when there is a single erasure in a stripe, but much effort has gone into designing LRCs with various other desirable properties [17, 26, 27, 33, 47]. One obvious direction that has received much attention is to design *distance-optimal LRCs* [33, 49] – that is, LRCs with the best possible *distance*, given their width, dimension, and locality. A code with distance d guarantees recovery from any pattern of up to $d - 1$ failures. Indeed, maximizing distance is especially important for wide codes, since they are more likely (simply due to their width) to encounter a larger number of failures simultaneously, as we show using data from Google storage clusters in Figure 2. But we find that even if we use distance-optimal LRCs, storage clusters using wide codes encounter a meaningful number of events where the number of failures is *larger* than the distance of the

distance-optimal LRC (see data from a Google storage cluster in Figure 1). So it becomes important to study LRCs that can successfully reconstruct even a significant fraction of erasure patterns *beyond the optimal distance*. The theory community has been tackling this problem by studying *maximally recoverable locally recoverable codes (MR-LRCs)*, which take the erasure correction capability of an LRC to the information theoretic limit implied by its design parameters [16, 18, 20]. However, to the best of our knowledge, current constructions of MR-LRCs do not yield codes fitting the limitations of current hardware. For example, wide MR-LRCs with storage overhead $<20\%$ require orders of magnitude larger field sizes than the computationally efficient field sizes of up to \mathbb{F}_{256} , which is explained in detail in Section 4.

Optimizing LRC reliability in practical parameter settings is a valuable and open problem, and we tackle aspects of it in this work. We highlight some of our main contributions.

Practical measurement of reliability. We study wide LRCs with a distinctively practical lens. One of our contributions is the curation of a set of *robust and practical* measures of LRC reliability. These include performance against random erasures, comparing the reliability of explicit LRC generator matrices (see Definition 5.1) against the information theoretic limit provided by MR-LRCs, and calculating *mean time to data loss (MTTDL)* using *observed* reliability metrics. As mentioned earlier, going wide creates new reliability challenges for LRCs, and these tools provide a clearer and more realistic set of tests to tackle the new challenges. Using these tools, we meaningfully compare popular deployed LRCs and showcase novel highly performant LRCs.

New distance-optimal LRC constructions. In Section 6, we provide a novel construction of distance-optimal LRCs (Definition 5.4), which we call *Optimal Cauchy LRCs*. This continues a long line of work [47, 49] for constructing distance-optimal LRCs, and while our construction is not the most general, it is remarkably simple and yields many codes (even wide LRCs) in practically useful parameter settings.

Insights into reliable wide LRC design. Our practical measures of reliability provide several insights. While it is true that higher distance gives guarantees about fixing up to a certain number of erasures, it is not enough by itself to guarantee strong empirical results. The first improvement that can be made in this direction is to find codes that are *approximately MR-LRCs* (since it is not yet possible to construct true MR-LRCs in our parameter regimes). Here we show good news: the coefficients used in our simple constructions get us over 99% of the reliability possible with MR-LRCs. But we find that just being (close to) maximally recoverable is not the end of the journey for reliability. Indeed, two codes that have the same width and same storage overhead, and are both MR-LRCs (true or approximate) can have significantly different resilience to *random patterns of erasures*. This is because erasure recovery is affected by how failures are distributed across the local repair groups (see Definition 5.5) of

the LRC, so it is not enough to just optimize the coefficients of an LRC, but also the *size* of its local repair groups. To the best of our knowledge, this fact has not been considered in the literature, even though our experiments show that it can have a significant impact on reliability. In general, codes whose local groups are *evenly sized* have better performance (see discussion in Section 8). This has the additional perk of also lowering the degraded mode read cost and the reconstruction cost of the code.

Novel LRCs that excel in practice. Using some of these insights, we modify our construction of distance-optimal codes to create *Uniform Cauchy LRCs*. We find that these codes truly shine in most practically relevant reliability (and performance!) measures. Figure 1 shows data unavailability events from a deployed wide LRC of width ≈ 50 blocks along with their erasure patterns captured from four large storage clusters at Google with a total disk population of over 1.7 million disks, over a period of one year. For the same code width and storage overhead, our Uniform Cauchy LRC simulation recovered more than 33% of these stripes without the need for restoration. Indeed, further experiments confirm this observation by showing that Uniform Cauchy LRCs outperform many popular (and deployed) LRCs across our metrics. A comprehensive experimental evaluation of LRCs is provided in Section 8, along with the main observations.

Maintenance-robust deployment of wide LRCs. In Section 9 we highlight the importance of maintenance-robust deployment of wide LRCs. Even though a code may have many desirable properties, its exact layout in a cluster affects its robustness to common maintenance events such as kernel upgrades. It is desirable to construct codes that are easier to deploy in a maintenance-robust manner (not all codes are equal here), and we show that the design of Uniform Cauchy LRCs is ideal in this regard.

Our work shows that a myriad of design choices need to be considered in order to optimize wide LRCs. Indeed, accounting for these factors can lead to more reliable deployments of wide LRCs in practice.

2 Background

Large-scale storage clusters. Large-scale storage clusters typically constitute of public cloud offerings or high-performance computing systems. Hard-disk drives (HDDs) make up the primary storage tier in these clusters. It is common for the disk population in a large-scale cluster to be above 100K [2, 32], and some large ones are also reported to have close to 500K disks [30, 31]. Data being stored in large-scale storage clusters is increasing at an alarming rate [7, 9, 10, 41]. Data redundancy using erasure coding is the reliability mechanism of choice for bulk of the data.

Erasur coding. Erasure coding is a more space-efficient alternative to data replication. Usually described as an (n, k) code, an erasure coding *stripe* encodes k data blocks (typi-

cally one, or a few megabytes in size) along with $n - k$ ‘parity’ blocks (of the same size) to form an n block *stripe*. The storage overhead is calculated as $\frac{n}{k}$. Maximum Distance Separable (MDS) codes (like Reed-Solomon codes) are popular erasure codes used in practice because they provide the maximum erasure correction capability for any fixed value of n and k . MDS codes have the property that up to $n - k$ blocks missing from a stripe can be reconstructed using *any* k of the remaining blocks. As data has grown and space-efficiency has become more critical, *wider* MDS codes (i.e. codes with larger values of k) with less storage overhead have become more popular in practice. Indeed, even (20, 17) MDS codes have been deployed and studied [3], in place of the once ubiquitous schemes like RAID-6 (which is a (6, 4)-Reed-Solomon code), (9, 6) [13] and (14, 10) [42]. But wider MDS codes have their own drawbacks because they require all the data from k blocks to reconstruct/repair even a single missing block, leading to a high *reconstruction cost*. A need to use wider encoding schemes with less overhead, combined with the very high reconstruction cost of MDS codes has motivated the study of Locally Recoverable Codes.

Data reconstruction process. Large-scale cluster storage systems have a background process that monitors the redundancy level of all stripes stored in the cluster. Whenever a disk becomes unavailable (either due to server unavailability, or disk failure), the background daemon flags the under-redundant stripe and starts a timeout of a few tens-of-minutes. When the timeout expires, the stripe is marked for reconstruction. Storage clusters often set a soft threshold on the bandwidth used for background activity such as reconstructions (except client-initiated degraded mode reads). In order to balance the reconstruction workload while maintaining highest data safety standards, the reconstructions are processed via a priority queue in which stripes that are more vulnerable are reconstructed before less vulnerable stripes.

Locally Recoverable Codes (LRCs). LRCs [17, 27, 33, 42, 47, 49] (also known as Local Reconstruction Codes) are erasure codes designed to mitigate the high reconstruction cost of MDS codes. An (n, k, ℓ) LRC code divides an n block stripe into local groups, each with at most $\ell < k$ blocks and a local parity¹. In addition to the local parities, the stripe also has *global parities* which cover all k data blocks. One may think of a *typical* LRC in this way: the data blocks and the global parities together form an MDS code, and the local parities are added on top of this code, so as to mitigate the cost of reconstruction (reduce it from k to ℓ) in the case when there is *exactly* one failure in a local group. One may note that LRCs are not MDS codes since they cannot satisfy the Singleton bound unless $\ell = k$ (see Definition 5.6). If there is more than one failure in the same local group, the underlying MDS code can be used to reconstruct the data. Several different LRC constructions have been proposed over the years with different

¹Theoretically each local group can have any number of parities, but in practice the most common configuration involves 1 parity per local group.

trade-offs [27, 33, 47, 49].

Distance of a code. Distance of a code (denoted by d) is the minimum number of failures/erasures that may render the stripe potentially non-recoverable (i.e. all patterns of $< d$ failures are always recoverable). For example, the distance of an (n, k) MDS code would be $n - k + 1$ since an MDS code can recover from any $n - k$ failures. In fact, for an MDS code, any failure beyond $n - k$ failures is strictly non-recoverable. However, for a code with distance d which is *not* MDS (such as an LRC), some patterns of $\geq d$ failures could be recovered. Maximizing this capability to recover as many erasure patterns as possible *beyond the distance* leads us to the notion of *maximally recoverable codes*.

Maximally Recoverable LRCs (MR-LRC). For any code, simply specifying whether each entry in its generator matrix (see Definition 5.1) is zero or non-zero imposes an information theoretic limit on which patterns of erasures could be recoverable. The study of *maximally recoverable codes* is concerned with finding coefficients for the non-zero entries such that this limit is reached [16]. LRCs can also be optimized in this way: once we have specified which entries of the generator matrix are non-zero (this will fix various code parameters like n , k , number of local parities, number of global parities, and the size of the local groups), it is possible to find coefficients that maximize the number of recoverable erasure patterns (including potentially many patterns of $\geq d$ erasures). LRCs which are *maximally recoverable* are known as MR-LRCs. In terms of their reliability, MR-LRCs are the gold standard among LRCs.

3 Motivation for studying wide LRCs

Reducing storage overhead is critical. Storage overhead because of data redundancy is a major component of a cluster’s storage cost. With 3-way replication, the overhead is $3\times$, which is prohibitive for large-scale clusters storing exabytes (EBs) of data on hundreds-of-thousands of disks. Even the popular MDS codes such as (9, 6) [13] or (14, 10) [42] which have an overhead of $1.5\times$ and $1.4\times$ respectively, are considered too expensive for exascale [9, 10, 29, 41]. Large-scale storage clusters are actively adopting *wide MDS codes* to minimize the storage overhead. Backblaze reported the use of a (20, 17) MDS code [3] which has a reasonably low overhead of $1.17\times$ (a rate of $\frac{17}{20} = 0.85$, see Definition 5.1). With every percent reduction in storage overhead resulting in savings of millions in capital, operational and energy costs, lowering storage overhead continues to be a lucrative problem.

Wide MDS codes are costly. Using wide MDS codes has several challenges. The reconstruction IO cost of a wide (n, k) MDS code scales linearly with k . For example, while a (9, 6) MDS code requires reading 6 data blocks for reconstructing a missing block, a wide MDS code such as (20, 17) requires reading 17 data blocks. Wide MDS codes also have a higher unavailability because they have a higher probability




Figure 2: A 24 hour trace with the number of stripes with at least 4 failures in two deployed LRCs captured from a single storage cluster at Google. The wider LRC ($n \approx 80$) has many more stripes with at least 4 failures compared to the relatively narrower LRC ($n \approx 50$).

of having blocks stored on devices that are unreachable due to maintenance (since no two blocks of a stripe can be on the same disk, server, rack, etc.). Thus, not only does a wide MDS code have a higher degraded mode read cost (Definition 5.10), but the frequency of performing degraded-mode reads is also higher due to higher unavailability. Due to a larger number of disks, degraded reads and reconstructions in wide MDS codes are also more likely to suffer from high tail latency due to stragglers. So although wide MDS code are good for durability and reduced storage overhead, they are not good for reconstruction costs, availability, and degraded read performance, all of which are major performance concerns in large-scale storage clusters. Storage overhead minimization is critical, but cannot come at the expense of aforementioned problems. Indeed, this is the reason for the continued popularity of LRCs, which mitigate this drawback.

Real-world failure patterns favor LRCs. We empirically observe the failure patterns of stripes stored in three large-scale storage clusters totalling over 1.5 million disks for a period of 6 months. These clusters have multiple different erasure coding and replication schemes deployed simultaneously, and have over 1.2 trillion stripes. From among the stripes that have > 0 failures, we observe that $\approx 99.2\%$ of stripes have just a single failure. Similar data has been observed by others in [40], where they found that $\approx 98.08\%$ stripes had a single failure on a Facebook warehouse cluster. A single block failure in a stripe is a scenario where LRCs shine (in contrast to MDS codes). However, this comes at a cost of higher storage overhead (as we mentioned, typical LRCs are just MDS codes with the additional overhead of local parities, and they cannot lie on the Singleton bound).

Wide LRCs. In an ideal world, we would like to use *wide* LRCs which have significantly lower overhead than the LRCs showcased in [33, 42] (and many other works). Wide LRCs could reduce the overhead from the 30 – 60% range (which is common in deployed LRCs and MDS codes) to the less than 20% range, while still maintaining many advantages over MDS codes. But wide LRCs create novel challenges. For example, wider stripes are much more likely to result in a larger percentage of stripes with more than a single failure. This makes it critical to study *robust and practical* measures of LRC reliability if we are to utilize wide codes in practice. Further, practical issues like deployment and ease-of-use also need to be addressed.

4 Practical challenges of wide LRCs

Many simultaneous failures are more common in wide LRCs. While most stripes (which have failures) have single-block failures, this does not mean that the tiny fraction of more than one block failures can be ignored. Moving to wider stripes increases the possibility that multiple blocks of a stripe need to be repaired at the same time, since there is a higher chance that devices or servers storing multiple blocks of the same stripe may be undergoing maintenance simultaneously. This is consistent with the observation that most data unavailability events in large-scale storage clusters are as a result of planned outages [13]. Another reason for increased number of block failures in a wide stripe is due to prioritization of reconstructions as explained in Section 2. Figure 2 shows the number of stripes that have at least 4 failures throughout a 24 hour period. There are two LRCs whose stripe failures are being compared, one with width approximately 50 blocks and the other with width approximately 80 blocks². Note that this trace is collected from a single storage cluster, and so the failures seen in this trace are in response to the same machines failing. As is clearly visible, the wider LRC has a significantly higher number of stripes with at least 4 failures compared to the relatively narrower LRC.

Constructing MR-LRCs is hard. As mentioned above, a lot of research has looked at MR-LRCs in recent years [5, 12, 16, 18, 20]. However, even with recent advancements such as [12, 18], to the best of our knowledge it is not possible to construct MR-LRCs in the following regime: where the field size (search-space for coefficients to construct the LRC) is fixed to be 256, n is between 25-150, and the rate of the code (see Definition 5.1) is at least 0.85 (i.e. storage overhead is at most $1.17\times$ and $\ell < k$). This setting is particularly useful for applications, and is the natural next phase for practical LRCs which reduce overhead and maximize reliability.

Typically, maximally recoverable codes are easier to construct when the field is large, and much research is focused on finding explicit codes over small field sizes. In practical settings, most computations are done over individual bytes which restricts the field size of the LRCs we can use to 256 (the field \mathbb{F}_{2^8}). To the best of our knowledge, the current state of the art constructions are in [18], and they require that the

²We cannot disclose the exact configuration due to confidentiality.

field size q be at least $\sim (\ell + 1)^r$, where $\ell + 1$ is the size of a local repair group (Definition 5.5), and r is the number of global parity checks. For our purposes, we can think of $\ell + 1 \sim 16$, and $r \geq 3$, thus $q = 4096$, making the required field size $16 \times$ greater than 256.

Measurements of reliability of wide LRCs are inadequate. While narrow LRCs have received a lot of attention in systems research, wide LRCs have not [25]. Existing systems research on LRCs optimizes metrics such as distance, degraded read cost and reconstruction cost [6, 27, 33, 42, 47, 49]. We enhance this rich literature with a larger suite of *empirical* measures, which give more realistic comparisons in the practical realm, particularly in the case when there are $\geq d$ erasures in a code with distance d . Further, since explicit MR-LRCs are not available in our parameter regimes, it leaves open the question of evaluating whether we can construct codes which offer us *approximately* the same advantages as MR-LRCs could offer (the answer is yes!). Finally, some design choices such as creating evenly sized local groups have not been considered in the literature at all, which we show can have a large impact on performance against random patterns of erasures.

Deployment of wide erasure codes is non-trivial. Another practical hurdle that gets worse with code width is the deployment (placement of blocks) of an erasure coded stripe. Recall that no two blocks of an erasure coded stripe can reside on the same disk, rack, fault-domain, power-source, etc. As the code width increases, it becomes progressively harder to fulfill these placement constraints. The placement problem is further exacerbated because it is common for sets of servers/racks (which comprise a *maintenance zone*) to be down at the same time for maintenance events, potentially causing data unavailability. For example, if too many blocks of a stripe are in the same maintenance zone, then even a planned maintenance event could make the stripe unavailable. We comment on some ways that code design can make it easier to achieve maintenance robust layouts of data.

5 Definitions

We begin by defining important terms that are going to be useful when describing the wide LRC construction in Section 6.

Definition 5.1 (Linear error correcting code). *A linear error-correcting code is simply a subspace $C \leq \mathbb{F}_q^n$ where q is a prime power and $n > 0$. It is customary to think of C as the image of an encoding map $Enc: \mathbb{F}_q^k \rightarrow \mathbb{F}_q^n$ for some $k \leq n$. This encoding may be expressed in matrix form as,*

$$Gx = y$$

where G is an $n \times k$ matrix called the generator matrix, x is the message and y is the codeword. The fraction $\frac{k}{n}$ is the rate of the code, k is the dimension, and n is the blocklength. The symbols $y_i \in \mathbb{F}_q$ of a codeword y are called codeword symbols.

Definition 5.2 (Distance of a code). *The minimum distance of a linear code (often just called the distance) is simply,*

$$\min_{x \in C \setminus \{0\}} wt(x)$$

where $wt(x) = \sum_{i=1}^n 1_{x_i \neq 0}$. An error correcting code with distance d can always correct $d - 1$ erasures.

Definition 5.3 ((n, k, ℓ) -Locally recoverable code (LRC)). *An (n, k, ℓ) -LRC is a linear error correcting code of dimension k and blocklength n . It has the additional property that any codeword symbol can be recovered from at most ℓ other codeword symbols. The parameter ℓ is called the locality parameter of the LRC. Note that $1 \leq \ell \leq k$.*

Definition 5.4 (Generalized Singleton bound and distance optimal LRCs). *Any locally recoverable code must satisfy the following bound [17],*

$$n \geq k + \left\lceil \frac{k}{\ell} \right\rceil + d - 2$$

where d is the distance of the code. Any LRC that meets this bound with equality is called a distance optimal LRC. Note that when $\ell = k$, this reduces to the well-known Singleton bound $n \geq k + d - 1$.

Definition 5.5 (Local repair group). *Given an (n, k, ℓ) -LRC, choose any codeword y . Due to the local recovery property, any codeword symbol y_i can be recovered using at most ℓ other codeword symbols $\{y_{r_1}, \dots, y_{r_\ell}\}$ where $r_j \neq i$ for any j . These at most $\ell + 1$ indices in $\{i, r_1, \dots, r_\ell\}$ form a local repair group of the (n, k, ℓ) -LRC.*

Definition 5.6 (Maximum distance separable code). *Any linear error-correcting code which satisfies the Singleton bound with equality, i.e. satisfies $n = k + d - 1$, is known as a Maximum Distance Separable (MDS) code.*

Definition 5.7 (Cauchy matrix). *Let $\{x_1, \dots, x_n\}$ and $\{y_1, \dots, y_m\}$ be two disjoint sequences of distinct elements from \mathbb{F}_q . The $n \times m$ matrix defined as*

$$C_{ij} = \frac{1}{x_i + y_j}$$

is a Cauchy matrix. It is well-known that every square submatrix of a Cauchy matrix has non-zero determinant (and therefore, is invertible).

Fact 5.8 (Generator matrix of an MDS code). *Cauchy matrices are commonly used to design generator matrices for MDS codes. Indeed, an $n \times k$ matrix over \mathbb{F}_q , whose first k rows form I_k (the $k \times k$ identity matrix), and whose last $n - k$ form an $(n - k) \times k$ Cauchy matrix generates an MDS code of blocklength n and dimension k (see Thm 2.2 and 5.2 in [8]).*

We can visualize several of our definitions using the example shown in Figure 3.

$$G_C = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ c_{11} & c_{12} & c_{13} & c_{14} \\ c_{21} & c_{22} & c_{23} & c_{24} \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ c_{11} + c_{21} & c_{12} + c_{22} & c_{13} + c_{23} & c_{14} + c_{24} \end{pmatrix} \begin{matrix} d_1 \\ d_2 \\ d_3 \\ d_4 \\ c_1 \\ c_2 \\ l_1 \\ l_2 \\ l_g \end{matrix}$$




Figure 3: Example of an LRC with 4 data blocks (d_1, d_2, d_3, d_4), 2 global parities (c_1, c_2), and 3 local parities (l_1, l_2, l_g). This figure shows how each row of the generator matrix corresponds to a ‘block’ in the picture representation of an LRC (the reader may note that the matrix representation is more informative as the specific choice of coefficients can affect reliability in practice). The first 4 rows of the matrix correspond to the data blocks which are encoded systematically; the next two rows correspond to the global parities (these two rows form a 2×4 Cauchy matrix in our implementation); the last 3 rows correspond to 3 local parities, all of which contain the XOR of the data they cover (as is evident by the coefficients of the matrix G_C). The last local parity, l_g , only covers the global parities. This design of LRC may be familiar to some readers as an Azure-LRC+1. The *local groups* of this code are (d_1, d_2, l_1) , (d_3, d_4, l_2) , and (c_1, c_2, l_g) and therefore $\ell = 2$. In this code, *ADRC* (see Def. 5.10) = ARC_1 (see Def. 5.11) = 2, since each data or parity block can be reconstructed by reading the two other blocks in its local repair group.

Definition 5.9 (Maximally recoverable locally recoverable code (MR-LRC)). *An (n, k, ℓ) -LRC is called a MR-LRC if it can recover from any pattern of $n - k$ erasures as long as there is at least one erasure in each local repair group of the code. To clarify, suppose the (n, k, ℓ) -LRC has local repair groups L_1, \dots, L_p for some $p > 1$ where each $L_i \subseteq \{1, \dots, n\}$. Recall that $\cup_i L_i = \{1, \dots, n\}$ and that the L_i are not necessarily disjoint. Let $E = \{e_1, \dots, e_{n-k}\} \subseteq \{1, \dots, n\}$ be any pattern of $n - k$ erased symbols. Then as long as $L_i \cup E \neq \emptyset$ for any i , then the erasure pattern E can be recovered by the code.*

Definition 5.10 (Average degraded read cost (ADRC)). *The ADRC is the average of the cost of reconstructing any of the data blocks. We can measure $cost(b_i)$ as the number of blocks which need to be read in order to reconstruct block i .*

$$ADRC = \frac{\sum_{i=1}^k cost(b_i)}{k}$$

Definition 5.11 (Average repair (or reconstruction) cost (ARC)). *The average repair cost is defined identically to the ADRC, with the addition of the global and local parity blocks to the computation.*

$$ARC_1 = \frac{\sum_{i=1}^n cost(b_i)}{n}$$

For MDS codes, $\forall b_i, cost(b_i) = ARC = k$. For LRCs, $cost(b_i)$ may not be the same as ARC. The above ARC is defined for one block failing in a stripe. In our evaluation we also show the ARC of reconstructing two failed blocks defined as

$$ARC_2 = \frac{\sum_{i=1, j \neq i}^n cost(b_{i,j})}{\binom{n}{2}}$$

where $1 \leq i \leq n, 1 \leq j \leq n$ and $i \neq j$, and $cost(b_{i,j}) = cost$ to reconstruct blocks i and j which can result in a combination of local and global reconstructions depending on i and j .

6 (n, k, r, p) -Optimal Cauchy LRCs

In this section we construct (n, k, r, p) -Optimal Cauchy LRCs and discuss their advantages. Here, n is the blocklength of the code, k is the dimension, r is the number of global parity checks, and p is the number of local parity checks. We make two design choices which are worth pointing out. First, the p local parities in our code are uniformly distributed across the k data, and they are all XOR-ed with the r global parity checks (this helps with proving the distance optimality of the code). Second, we restrict ourselves to the case where each of the data symbols of the code is covered by exactly *one* local parity. The second restriction is common in constructions proposed in the literature, and indeed, important for our goal of minimizing the storage overhead. For simplicity of exposition we assume that p is even (we will mention how this condition can be removed), and $p|k$ where we denote $\frac{k}{p} = t$. It will be clear that the *locality* parameter ℓ of the (n, k, r, p) -Optimal Cauchy LRCs we will construct is $\frac{k}{p} + r$.

In Section 6.2 will show that (n, k, r, p) -Optimal Cauchy LRCs have the best possible distance for LRCs with their dimension and locality (under some modest restrictions). In Appendix A.2 we show how some of these restrictions can be relaxed. We begin with explaining the construction of (n, k, r, p) -Optimal Cauchy LRCs.

6.1 Code construction

The generator matrix for an (n, k, r, p) -Optimal Cauchy LRC is derived naturally from the generator matrix for an $(k + r + 1, k)$ -MDS code which has dimension k and blocklength $n = k + r + 1$. Specifically, we derive our code from the generator matrix of a $(k + r + 1, k)$ -Cauchy MDS code.

The generator matrix $G_{(k+r+1),k}$ of a $(k + r + 1, k)$ -Cauchy MDS code is an $k + r + 1 \times k$ matrix, which simply consists of a $k \times k$ identity matrix stacked on top of a $r + 1 \times k$ Cauchy matrix. The resulting matrix is well-known to generate an MDS code with distance $r + 2$ (See Thm 2.2 and 5.2 in [8]).

$$G_{(k+r+1),k} = \left[\begin{array}{ccc|ccc} 1 & & & & & \\ & \ddots & & & & \\ & & & & & 0 \\ & & & & \ddots & \\ & 0 & & & \ddots & \\ & & & & & 1 \\ \hline c_{11} & & \dots & & & c_{1k} \\ & & & \ddots & & \\ & c_{r1} & & & & c_{rk} \\ c_{(r+1)1} & & & & & c_{(r+1)k} \end{array} \right]$$

When we do not need to refer to r and k are specifically, we refer to $G_{(k+r+1),k}$ as G , and the i th row of $G_{(k+r+1),k}$ as g_i .

To create the generator matrix for an (n, k, r, p) -optimal cauchy LRC, we first partition this last row of G , g_{k+r+1} into p rows r_1, \dots, r_p as follows,

$$\begin{aligned} r_1 &= (c_{(r+1)1}, c_{(r+1)2}, \dots, c_{(r+1)t}, 0, \dots, 0) \\ r_2 &= (0, \dots, 0, c_{(r+1)(t+1)}, \dots, c_{(r+1)2t}, 0, \dots, 0) \\ &\vdots \\ r_p &= (0, \dots, 0, c_{(r+1)(p(t-1))}, \dots, c_{(r+1)pt}) \end{aligned}$$

It is clear that $r_1 + r_2 + \dots + r_p = g_{k+r+1}$. We compute the rows \tilde{r}_i as follows,

$$\tilde{r}_i = r_i + g_{k+1} + g_{k+2} + \dots + g_{k+r}$$

Note that \tilde{r}_i are simply r_i with the addition of the first r ‘cauchy rows’ of G . Since p is even (and the underlying field has characteristic 2) we have,

$$\begin{aligned} \tilde{r}_1 + \tilde{r}_2 + \dots + \tilde{r}_p &= p(g_{k+1} + g_{k+2} + \dots + g_{k+r}) + g_{k+r+1} \\ &= g_{k+r+1} \end{aligned}$$

Finally, the generator matrix $O_{n,k,r,p}$ for the (n, k, r, p) -Optimal Cauchy LRC is given by,

$$O_{n,k,r,p} = \left[\begin{array}{ccc|ccc} 1 & & & & & \\ & \ddots & & & & \\ & & & & & 0 \\ & & & & \ddots & \\ & 0 & & & \ddots & \\ & & & & & 1 \\ \hline c_{11} & & \dots & & & c_{1k} \\ & & & \ddots & & \\ & c_{r1} & & & & c_{rk} \\ & & & \tilde{r}_1 & & \\ & & & \vdots & & \\ & & & \tilde{r}_p & & \end{array} \right]$$

This matrix is simply the $k+r$ rows of $G_{k+r+1,k}$ stacked on top of the $\tilde{r}_1, \tilde{r}_2, \dots, \tilde{r}_p$ vectors. We remark that $n = k + r + p$, and that $O_{n,k,r,p}$ generates an LRC with locality parameter $\ell = \frac{k+r}{p} + r$. This is clear since each row among $\tilde{r}_1, \tilde{r}_2, \dots, \tilde{r}_p$ provides a local parity check on exactly ℓ other rows.

6.2 Distance

We note that the distance of an (n, k, r, p) -Optimal Cauchy LRC is exactly $r+2$. As long as $k, r, p > 0$, $p|k$, $rp^2 > k+rp$, and p is even, it is a distance optimal LRC (Definition 5.4).

Theorem 6.1. For $r > 0$, $O_{n,k,r,p}$ generates an error correcting code with distance exactly $r+2$. Proof details are in Appendix A.1

Lemma 6.2. Given $k, r, p > 0$, $rp^2 > k+rp$, and p even, the locally recoverable code formed by $O_{n,k,r,p}$ is distance optimal. Proof details are in Appendix A.1

We show how some of the conditions in Lemma 6.2 can be relaxed in Appendix A.2.

Remark 6.3. Constructing distance optimal LRCs is a well-studied problem. At first, only constructions whose alphabet size was exponential in the blocklength were known [45, 49]. These do not yield codes which may be used in practical settings where the alphabet size q is fixed at 256. The most well-known construction of distance optimal LRCs may be found in [33, 47]. To get truly optimal codes the codes of [47] require that $\ell+1|n$, and none of the codes we use in our experiments meet this restriction. Distance optimal LRCs with very general parameters were finally shown in [33]. In our experimental analysis in the later sections, we use Optimal Cauchy LRCs (rather than other known constructions such as [33]) as examples of codes which lie on the generalized Singleton bound (essentially, distance optimal LRCs. See Definition 5.4). This is because Optimal Cauchy LRCs are easier to construct in our parameter regime.

7 (n, k, r, p) -Uniform Cauchy LRCs

In the previous section we give a very simple construction of distance optimal LRCs which we will use in our experimental analysis as examples of distance optimal codes. In this section, we provide a simple heuristic modification of these codes which has several practical advantages (though they are not shown to be optimal with regards to distance).

The later experimental analysis of these codes highlights the point that distance optimal (i.e. on the generalized Singleton bound) does not mean most-durable or most cost-efficient from a practical perspective. For example, a code that has the same locality and distance can have different durability (as measured by mean-time-to-data-loss) and robustness against random patterns of erasures. Distance-optimality simply indicates the best distance for fixed values of n, k and ℓ .

These codes are constructed in much the same way as optimal cauchy LRCs, except that each local parity check covers $\frac{k+r}{p}$ of the data blocks and global parity blocks. Following the same notation as the previous section, the generator matrix for a uniform cauchy code has the form (assuming $p|k+r$ for simplicity),

Locality	Azure-LRC	Azure-LRC+1	Optimal Cauchy LRC	Uniform Cauchy LRC
$n = 28, k = 24, r = 2, p = 2$	$\ell = 24$	$\ell = 24$	$\ell = 12 + 2 = 14$	$\ell = \mathbf{13}$
$n = 55, k = 48, r = 3, p = 4$	$\ell = 48$	$\ell = 16$	$\ell = 12 + 3 = 15$	$\ell = \mathbf{13}$
$n = 80, k = 72, r = 4, p = 4$	$\ell = 72$	$\ell = 24$	$\ell = 18 + 4 = 22$	$\ell = \mathbf{19}$
$n = 105, k = 96, r = 5, p = 4$	$\ell = 96$	$\ell = 32$	$\ell = 24 + 5 = 29$	$\ell = \mathbf{26}$
Avg. degraded read cost ($ADRC$)	Azure-LRC	Azure-LRC+1	Optimal Cauchy LRC	Uniform Cauchy LRC
$n = 28, k = 24, r = 2, p = 2$	$\frac{24 \times 12}{24} = \mathbf{12}$	$\frac{24 \times 24}{24} = 24$	$\frac{12 \times 14}{12} = 14$	$\frac{26 \times 13}{26} = 13$
$n = 55, k = 48, r = 3, p = 4$	$\frac{48 \times 12}{48} = \mathbf{12}$	$\frac{48 \times 16}{48} = 16$	$\frac{48 \times 15}{48} = 15$	$\frac{13 \times 12 + 35 \times 13}{26} = 12.72$
$n = 80, k = 72, r = 4, p = 4$	$\frac{72 \times 18}{72} = \mathbf{18}$	$\frac{72 \times 24}{72} = 24$	$\frac{72 \times 22}{72} = 22$	$\frac{72 \times 19}{72} = 19$
$n = 105, k = 96, r = 5, p = 4$	$\frac{96 \times 24}{96} = \mathbf{24}$	$\frac{96 \times 32}{96} = 32$	$\frac{96 \times 29}{96} = 29$	$\frac{78 \times 25 + 18 \times 26}{96} = 25.18$
Avg. repair cost 1 failure (ARC_1)	Azure-LRC	Azure-LRC+1	Optimal Cauchy LRC	Uniform Cauchy LRC
$n = 28, k = 24, r = 2, p = 2$	$\frac{26 \times 12 + 2 \times 24}{28} = \mathbf{12.85}$	$\frac{25 \times 24 + 3 \times 2}{28} = 21.64$	$\frac{28 \times 13}{28} = 13$	$\frac{28 \times 13}{28} = 13$
$n = 55, k = 48, r = 3, p = 4$	$\frac{52 \times 12 + 3 \times 48}{55} = 13.94$	$\frac{51 \times 16 + 4 \times 3}{56} = 15.05$	$\frac{55 \times 15}{55} = 15$	$\frac{13 \times 12 + 42 \times 13}{55} = \mathbf{12.76}$
$n = 80, k = 72, r = 4, p = 4$	$\frac{76 \times 18 + 4 \times 72}{80} = 20.7$	$\frac{75 \times 24 + 5 \times 4}{81} = 22.75$	$\frac{80 \times 22}{80} = 22$	$\frac{80 \times 19}{80} = \mathbf{19}$
$n = 105, k = 96, r = 5, p = 4$	$\frac{100 \times 24 + 5 \times 96}{80} = 27.42$	$\frac{99 \times 32 + 6 \times 5}{81} = 30.45$	$\frac{105 \times 29}{105} = 29$	$\frac{81 \times 25 + 24 \times 26}{105} = \mathbf{25.22}$
Avg. repair cost 2 failures (ARC_2)	Azure-LRC	Azure-LRC+1	Optimal Cauchy LRC	Uniform Cauchy LRC
$n = 28, k = 24, r = 2, p = 2$	30.66	43.46	32.12	27.92
$n = 55, k = 48, r = 3, p = 4$	35.49	39.22	36.93	33.85
$n = 80, k = 72, r = 4, p = 4$	52.80	59.38	54.82	49.22
$n = 105, k = 96, r = 5, p = 4$	70.68	79.73	74.50	67.69
Normalized MTDL comparison	Azure-LRC	Azure-LRC+1	Optimal Cauchy LRC	Uniform Cauchy LRC
$n = 28, k = 24, r = 2, p = 2$	0.64×	0.14×	0.50×	1.00
$n = 55, k = 48, r = 3, p = 4$	0.99×	0.97×	1.01 ×	1.00
$n = 80, k = 72, r = 4, p = 4$	0.99×	0.97×	0.49×	1.00
$n = 105, k = 96, r = 5, p = 4$	0.99×	0.96×	0.96×	1.00

Table 2: This table captures all the analytical metrics used to compare the different LRC constructions across the different wide LRC parameters described in Table 1. Details of this comparison are described in Section 8. The takeaways are that for all metrics except average degraded mode read cost (in which it is < 9% worse than the best LRC), Uniform Cauchy LRCs outperform other LRCs (including the Optimal Cauchy LRC). Another surprising result was that Azure-LRC outperforms Azure-LRC+1 despite its global parities having no local parities protecting them.

ever, we include this code in our analysis due to its popularity and because it is actually used in practice (and our analysis shows that it gives great performance!).

Azure-LRC+1 is an optimization to the Azure-LRC proposed in [33]. Azure-LRC suffers from an expensive MDS-level reconstruction on the failure of any of the global parities. In order to prevent this, Azure-LRC+1 forms a local group of the global parities and protects them using a local parity. Figure 4b captures the Azure-LRC+1 construction. Note that the Azure-LRC+1 construction in [33] introduces an additional local parity to protect the global parities without removing any of the existing parity blocks. This construction changes the storage overhead of Azure-LRC+1 in comparison to other constructions. We deliberately choose to keep the numbers k -of- n the same across all schemes, otherwise we are comparing codes with different redundancy, which in our view is not a fair comparison. Thus, in our construction of Azure-LRC+1, we choose to remove one local parity protecting data blocks (as compared to removing a global parity since removing a local parity has a less adverse effect on the reliability), and

add a local parity protecting global parity blocks.

Figure 4c and Figure 4d are representations of the Optimal Cauchy and Uniform Cauchy constructions which have been described comprehensively in previous sections.

Parameter regimes for comparison. We select four different widths representing wide LRCs, details of which are in Table 1. For an apples-to-apples comparison, we freeze the size of the data, k , the size of the code n , the number of global parities r , and the number of local parities, p . Since a (20, 17) MDS code is reportedly in use [3], we explore schemes where $24 \leq n \leq 105$ (deliberately starting from n close to 20, and codes with rate strictly higher than 0.85). The entire set of results are presented in Table 2, but we highlight the main points below.

Uniform Cauchy LRC has the smallest locality. We first compare the locality of the different LRC constructions. Recall (Definition 5.3) that locality refers to the maximum number of blocks to be read for reconstruction of a single block. Since Azure-LRC requires reading all data blocks to reconstruct any failed global parity, its locality is the high-




Figure 5: This plot compares the durability of the different LRC constructions (Azure-LRC, Azure-LRC+1, Optimal Cauchy and Uniform Cauchy) by measuring their ability to recover from random failures (only values where some recovery ratios were less than one are shown). We evaluate all wide LRCs discussed in evaluation ranging from 24-of-28 to 96-of-105. Except in 48-of-55, we see that Uniform Cauchy has the best durability. Optimal Cauchy has the best durability in 48-of-55. Surprisingly, Azure-LRC has better durability compared to Azure-LRC+1 even though global parities of Azure LRCs are not covered by a local parity.




Figure 6: Comparing performance with MR-LRC. In this plot we show the results of a Monte-Carlo experiment where p failures were forced (one in each local group) and another $n - k - p$ failures were distributed randomly across remaining blocks. MR-LRCs can recover from all such failure patterns. Azure-LRC+1, Optimal Cauchy and Uniform Cauchy are all $> 99\%$ as durable as an MR-LRC in this scenario for our choice of coefficients. Azure-LRC is not shown because it does not have a local group covering its global parities, making it unsuitable for this comparison.

est. This is followed by Azure-LRC+1, whose local groups are larger than the other constructions (since we have constrained p). In fact for 24-of-28, Azure-LRC+1 has the same locality as Azure-LRC due to having only 1 local group that spans all data blocks. Optimal Cauchy LRC evenly divides the data blocks, but requires each local group to contain *all* the global parities (we need this to prove distance optimality as explained in Section 6.2), making its locality $\left\lceil \frac{k}{p} \right\rceil + g$. Uniform Cauchy LRC has the lowest locality $\left\lceil \frac{n}{p} \right\rceil$ since it uniformly divides n blocks into p groups.

Azure-LRC has the lowest average degraded read cost (ADRC), closely followed by Uniform Cauchy LRC. The ADRC is calculated only for the data blocks (Definition 5.10), and therefore is directly proportional to the size of local groups. Therefore, since Azure-LRC has the smallest local groups, it also has the smallest ADRC. Although, as the stripes become wider, the difference between the local group sizes of Azure-LRC and Uniform Cauchy LRC starts reducing. This reflects in the reduction of the ratio of the ADRC of Uniform Cauchy LRCs compared to Azure-LRCs. In particular, the ADRC of Uniform Cauchy LRC is about 8% more than Azure-LRC for 24-of-28, but is only about 5% more for 96-of-105.

Azure-LRC+1 has the highest ADRC followed by Optimal Cauchy LRC, both owing to their larger local repair groups.

Average repair cost (ARC) is lowest for Uniform Cauchy LRC. Very similar to the ADRC, the ARC_1 and ARC_2 (Definition 5.11) are the degraded read cost for all n blocks of a stripe for 1 block failure and 2 block failures respectively. We choose to showcase both ARC_1 and ARC_2 because the wider the LRC, the higher the likelihood that more than one failure can occur in a stripe as explained in Section 3. Here, the disadvantage of Azure-LRC turns into the advantage for Uniform Cauchy LRC. Specifically, the degraded read cost for the global parities involves reading all data blocks in the case of Azure-LRC. For Uniform Cauchy LRC, the degraded read cost of global parities is the same as a data block, both of which are only slightly higher than the data block reconstruction cost of an Azure-LRC. The exception is 24-of-28 where Azure-LRC has a slightly lower ARC than Uniform Cauchy LRC, owing to these specific parameters, but this too becomes favorable for Uniform Cauchy LRCs in the case of ARC_2 ³. Across the schemes, as the number of data blocks increase, the ARC of Azure-LRC reduces, but at the same time, any additional global parity increases the ARC significantly. The difference between the Azure-LRC+1 and Optimal Cauchy LRC is lower than their difference in ADRC. This is due to Azure-LRC+1 having a very low-cost global parity reconstruction. Nevertheless, the Azure-LRC+1 and Optimal Cauchy LRCs have the highest, and second highest ARCs among the four constructions.

Uniform Cauchy LRCs have the best random failure tolerance. We compare the LRC constructions empirically by evaluating their durability against random failures (including when the number of failures are $\geq d$ when d is the distance of the code). For each code, we choose various values of i and conduct a Monte-Carlo experiment in which i blocks are removed uniformly at random (without replacement) from the n blocks of stripe. Then data recovery is attempted. The more times recovery can succeed, the more durable the code. For each random failures experiment at least 1 million recover-

³We refrain from detailing the calculation of ARC_2 because it significantly more complex than ARC_1 .

ies from unique block failure combinations were attempted. The exception was 4 random failures for 24-of-28, where all combinations (20475) were exhaustively checked.

Figure 5 shows the results of the various failure scenarios on each of the four schemes across all four LRCs. Uniform Cauchy LRC outperforms all other LRC constructions in each scenario, for each scheme, except 48-of-55. In 48-of-55, the Optimal Cauchy LRC has the highest recoverability ratio. The intuition behind Uniform Cauchy LRC’s superior performance is that when a high number of failures happen uniformly at random, it’s more likely that at least one failure occurs in each local repair group. This results in all local parities contributing to the reconstruction process, leading to a higher success rate.

Simple choices of coefficients give almost MR-LRCs in our parameter regime. We construct all our generator matrices using Cauchy matrices. We then conduct an experiment to compare our codes to a *hypothetical MR-LRC* of the exact same design. We say ‘hypothetical’ because although there is proof that such an MR-LRC can exist [18, 22], currently there does not exist a deterministic way to construct such a code. Nevertheless, we can precisely characterize the erasure patterns that can be recovered by such MR-LRCs. Simply put, an MR-LRC with p local parities and r global parities can recover *any* pattern of $r + p$ failures, *as long as* there is at least one failure in each local repair group (so that each local parity may contribute towards the reconstruction). An LRC which has this property is an MR-LRC.

We conduct an experiment in which we *plant* one failure in each local group (i.e. p total planted failures), and then add another r failures at random. Then we attempt recovery. We find that at least with coefficients we chose (all derived from Cauchy matrices), *all* the codes we tested were very close to being MR-LRCs. So even if it is hard to construct MR-LRCs, in practical parameter settings, it is not hard to realize many of their benefits using common code constructions.

We do point out that even though all code constructions were close to being MR-LRCs, this does not mean that they were all equally durable against random erasures. This is because the shape of the code affects the probability that $r + p$ random failures will spread out so that each local repair group gets at least one failure. Indeed, this provides some intuition as to why Uniform Cauchy LRCs perform the best against random erasures. It is because the evenly sized local repair groups maximize the probability that each local repair group will see at least one failure (recall MR-LRC Definition 5.9).

Mean-time-to-data-loss (MTTDL). MTTDL (for a stripe) has been a canonical metric for reliability in the coding community [14]. It is modeled using a continuous time Markov chain in which each state represents the number of erasures in a stripe. There is a final absorbing state, denoted *fail*, which represents a data loss/unavailability event. The MTTDL for a coding scheme is simply the mean time until a Markov chain starting at state zero reaches the absorbing state




Figure 7: MTTDL Markov Chain for a 48-of-55 code. The probabilities $p_{5 \rightarrow fail}$, $p_{6 \rightarrow fail}$, and $p_{7 \rightarrow fail}$ are added from Fig. 5 (b, c, d) where they are determined empirically. The probabilities $p_{i \rightarrow fail}$ for $i < 5$ are not pictured because they are zero.

i.e. the expected time until a healthy stripe experiences a data unavailability event. Each state representing e failures, denoted s_e , may transition to the states s_{e-1} , s_{e+1} or *fail*. These transition probabilities are usually modelled using exponential distributions representing *repair time* (which captures the transition from s_e to s_{e-1}), *failure probability* (which captures the transition from s_e to s_{e+1}) and *complete failure probability* (which captures the transition from s_e to *fail*). These theoretical models are ubiquitous in modeling stripe reliability for MDS codes, but they have also been used to study the reliability of LRCs [27, 42, 44]. In our work, replace the *modelled* transition probabilities with *observed* ones, in order to get a better estimate of the MTTDL of our codes.

We accomplish this by supplementing our theoretical model with empirical data in the following way: for each state s_e , we add a transition to the irrecoverable data loss state with probability $p_{e \rightarrow fail}$, where $p_{e \rightarrow fail}$ is probability that a stripe with e randomly distributed failures cannot recover the data, as shown in Figure 7. To be clear, in order to model MTTDL this way, we need the explicit generator matrix of our code, and then to conduct the random failures experiment to generate data. While no model is perfect, we believe that this is a more realistic evaluation of MTTDL of an LRC, since we meaningfully account for the observed durability of the explicit code (i.e. accounting for coefficients).

Table 2 shows the MTTDL comparisons for different LRC constructions. Since we are interested in the relative comparison of MTTDLs, we show the ratio of each LRC construction with the Uniform Cauchy LRC MTTDL (where MTTDL is calculated as a function of p_{fail} and p_{repair} whose values can differ significantly based on specific cluster size, architecture, disk makes/models, etc.). Although p_{repair} can differ based on the architecture of the code, those differences are negligible in comparison to the wait-time for the detection of the missing blocks in a stripe as explained in Section 2. We find that Uniform Cauchy LRC achieves the highest MTTDL (up to a factor of 100× in our evaluation), while Azure-LRC+1 has the lowest MTTDL values. This is in line with our observations that Uniform Cauchy LRCs and Azure-LRC+1 have respectively the highest and lowest performance in reconstructing random erasures, and also the lowest and highest average locality. This is unsurprising since MTTDL of an LRC is inversely proportional to its average locality [44].

9 Maintenance-robust deployment

When deploying an erasure code in a large-scale storage cluster, there are several placement constraints that need to be met in order to ensure adequate data reliability. For example, usually no two blocks of the same stripe are put on the same disk, server, rack or at times even racks powered by the same power source to improve data reliability. Placement restrictions have been studied recently in [30], and in this work we comment that placement constraints can be considered during code design as well.

In a storage cluster, the smallest unit in which maintenance such as kernel/firmware/hardware upgrades can be performed is known as a *maintenance zone*. The design of maintenance zones affects reliability because all servers/devices in a single maintenance zone can be turned off simultaneously during a maintenance event. So we would like to have each stripe intersect a maintenance zone as little as possible, making it robust to the maintenance events in a real storage cluster. In an ideal world, we could make maintenance zones as small as possible (e.g. a single rack) to limit the impact of maintenance events on reliability. But this is not practically feasible, as it would increase the operational toil of maintenance tasks. One reason for this is: after every maintenance task (belonging to a single maintenance zone) cluster-wide (and therefore time-consuming and expensive) data reachability and reliability tests are usually conducted to ensure that the maintenance activity was completed without unforeseen events. If maintenance zones are too small (say an individual server), then a large-scale cluster of tens-of-thousands of servers would be in perennial testing activity. This severely limits the number of maintenance zones, and in practice, storage administrators have informed us that the total number of maintenance zones in a single cluster are typically restricted to below 20. This means that a maintenance event can affect a large number of stripes, and hence it is important to deploy stripes in a manner that is robust to maintenance events.

One desirable property of a code deployment might be that no maintenance event that affects a single zone should render any stripe unrecoverable. We name this *maintenance-robust deployment*. Since maintenance events are foreseeable, and usually performed one zone at a time, a maintenance-robust deployment would ensure high availability.

In the context of LRCs, we can further optimize deployment strategies to increase the likelihood of cost-efficient local repairs. For example, if we guarantee that every maintenance zone has at most one block from each local group, all degraded reads during a maintenance event would only require local repairs. We term such a deployment *maintenance-robust-efficient deployment*. Ideally we want all LRC deployments to be maintenance-robust-efficient deployments.

Consider the $(9, 4, 2)$ -LRC from Figure 3. If the number of maintenance zones, say z is 3. We can assign each block of data from every local repair group to a different maintenance

zone (i.e. a maintenance-robust-efficient deployment). Here, during *any* maintenance event, the degraded read cost for the data blocks stored in that maintenance zone is still $\ell = 2$. However, if $z = 2$, we cannot guarantee local repairs in all cases. As a best case, we could have $z_1 = d_1, d_3, c_1, l_g$ and put the rest of the data in z_2 . Suppose z_2 undergoes maintenance, the data can be repaired by first repairing c_2 using c_1, l_g , and subsequently performing a global repair using d_1, d_3, c_1 and c_2 to reconstruct d_2 and d_4 . Suppose z_1 undergoes maintenance, d_1 can be repaired using d_2, l_1 , and d_3 can be repaired using d_4, l_2 . Thus, with $z = 2$, a $(9, 4, 2)$ -LRC can only be deployed in a maintenance-robust manner, where global repairs are required for reconstructing data. In general, when $z \geq \ell + 1$ for an LRC, it can guarantee maintenance-robust-efficient deployment. The parameter $\ell + 1$ in Table 2 shows that minimum number of maintenance zones required for maintenance-robust-efficient deployment to be possible for the schemes we have evaluated in this work. Thus, with $z \approx 20$, all schemes except 96-of-105 can be deployed in a maintenance-robust-efficient manner for Uniform Cauchy LRCs.

10 Related Work

LRCs have widely studied and deployed in the last decade due to their practical improvements over MDS codes. While Microsoft Azure and Facebook first published the commercial use of LRCs [27, 42], much academic research has also been conducted to find LRC constructions with desirable properties [1, 6, 11, 17, 39, 39, 44, 46, 48]. One area of research has been explicit constructions of distance optimal codes (Definition 5.4) [6, 24, 28, 35, 37, 38, 47, 49, 51]. Tamo and Barg [47] were among the first to provide a general construction of distance-optimal LRCs (with small field sizes), but with some constraints on the allowable parameters. This construction was further generalized by Kolosov et al. [33] for a wide range of parameters. Recent years have seen a surge in research on MR-LRCs [4, 5, 12, 15, 16, 18–22, 36]. Recently, Gopi et al. [18] provided an explicit construction of MR-LRCs which is broad, but does not cover our parameter regime.

Wide codes are known to be deployed in two commercial settings, VAST [50] and Backblaze [3], and recent works have studied wide codes in academia. Haddock et al. analyze wide codes' performance using GPUs [23]. Li et al. show local erasures codes in the context of hard-disk drives [34]. More recently, Hu et al. [25] study the performance issues in repair, encoding, and update performance of wide codes. Our practically inspired work adds to this literature by studying the trade-offs of wide LRCs from construction to deployment.

11 Conclusion

In this work we show that many subtle factors can affect LRC reliability in real world scenarios: coefficients in the generator

matrix, design of local repair groups, and maintenance-robust deployment. We show the value of an experiment-driven understanding of reliability as it provides us novel insights into design choices which have a meaningful impact on reliability. Indeed, this culminates in our construction of Uniform Cauchy LRCs which outperform popular (and provably distance optimal) codes in practice.

12 Acknowledgements

We thank our shepherd Patrick P. C. Lee and the anonymous reviewers for their valuable feedback and suggestions. We extend special thanks to Larry Greenfield and numerous other researchers and engineers at Google.

A Optimal Cauchy LRCs

A.1 Distance optimality of Optimal Cauchy LRCs

Theorem A.1. *For $r > 0$, $O_{n,k,r,p}$ generates an error correcting code with distance exactly $r + 2$.*

Proof. We will simply show that the distance of the code generated by $O_{n,k,r,p}$ is equal to the distance of the code generated by $G_{(k+r+1),k}$. Recall that $G_{(k+r+1),k}$ generates a code with distance exactly $r + 2$, which is equivalent to saying that this code can correct *any* pattern of $r + 1$ erasures, i.e. $G_{(k+r+1),k}$ has row rank at least k if any $r + 1$ rows are deleted.

We will show that $O_{n,k,r,p}$ has row rank at least k whenever any $r + 1$ rows are deleted. Note that whenever some $r + 1$ rows are deleted from $O_{n,k,r,p}$, they may or may not contain any rows among the $\tilde{r}_1, \dots, \tilde{r}_p$.

First we consider the case when at least one of $\tilde{r}_1, \dots, \tilde{r}_p$ are deleted. In this case, we may consider *all* of them lost, and note that we have exactly the same rows remaining as if g_{r+k+1} , and at most r other rows were lost in $G_{(k+r+1),k}$. Since, $G_{(k+r+1),k}$ generates an MDS code with distance $r + 2$, it is clear that the remainder of the rows will have rank k .

Now consider the case when the $r + 1$ deleted rows do not contain any of the rows $\tilde{r}_1, \tilde{r}_2, \dots, \tilde{r}_p$. In this case, we can compute $g_{k+r+1} = \tilde{r}_1 + \tilde{r}_2 + \dots + \tilde{r}_p$, which reduces to the case when $r + 1$ rows are deleted from $G_{(k+r+1),k}$. \square

Lemma A.2. *Given $k, r, p > 0$, $rp^2 > k + rp$, and p even, the LRC formed by $O_{n,k,r,p}$ is distance optimal.*

Proof. Recall that for a LRC to be distance optimal we must have,

$$n = k + \left\lceil \frac{k}{\ell} \right\rceil + d - 2 \quad (1)$$

where ℓ is the locality parameter and d is the distance. We know that $\ell = \frac{k}{p} + r$ and $n = k + r + p$ for an (n, k, r, p) -optimal

cauchy code. Substituting ℓ and $d = r + 2$ (from Theorem A.1), equation 1 becomes,

$$n = k + \left\lceil \frac{k}{\frac{k}{p} + r} \right\rceil + r = k + \left\lceil \frac{pk}{k + pr} \right\rceil + r$$

So an (n, k, r, p) -optimal cauchy code will be distance optimal whenever $\left\lceil \frac{pk}{k + pr} \right\rceil = p$, i.e. $p - 1 < \frac{pk}{k + pr} \leq p$. Both inequalities hold as long as $r, p, k > 0$ and $rp^2 < k + rp$. \square

A.2 Relaxing constraints

We have shown how to construct distance optimal LRCs with blocklength $n = k + r + p$, where k is the dimension of the code, r is the number of global parity checks, and p is the number of local parity checks as long as p is even, $r, p, k > 0$ and $rp^2 < k + rp$. We will briefly discuss some extensions of these parameter regimes in this section.

Claim A.3. *We may modify our construction of distance optimal codes to work with p odd.*

If $p = 1$, we may modify the construction so that $r_1 = r_p = \tilde{r}_p = g_{k+r+1}$, and so $G_{(k+r+1),k} = O_{n,k,r,p}$. This just reduces to the code defined by $G_{(k+r+1),k}$.

If $p \geq 3$, we may modify the construction so that $\tilde{r}_i = r_i$ for $1 \leq i \leq p - 2$, $\tilde{r}_{p-1} = r_{p-1} + g_{k+1} + \dots + g_{k+r}$, and $\tilde{r}_p = r_p + g_{k+1} + \dots + g_{k+r}$. In this case,

$$\begin{aligned} \tilde{r}_1 + \tilde{r}_2 + \dots + \tilde{r}_p &= 2(g_{k+1} + g_{k+2} + \dots + g_{k+r}) + g_{k+r+1} \\ &= g_{k+r+1} \end{aligned}$$

From here, the same reasoning as the previous sections implies that these codes are distance optimal. Notice that when $p \geq 3$, these modifications do not change ℓ because \tilde{r}_p is a local parity check on $\frac{k}{p} + r$ data blocks.

Claim A.4. *If $\frac{rp^2}{2} < k + rp < rp^2$, our construction gives codes whose blocklength is at most one greater than a distance optimal code.*

This is easily observed by following the proof of Lemma 6.2 and considering the cases where $p - 2 < \frac{pk}{k + pr} \leq p - 1$. Previous works have also found such ‘off by one’ LRCs in broader parameter regimes than truly optimal LRCs [47]. This allows us more freedom to explicitly construct almost distance optimal and wide LRCs.

Claim A.5. *We may construct codes with alphabet size $q \geq k + r + 1$.*

Note that we can construct our LRCs starting from the generator matrix of any MDS code in place of $G_{k+r+1,k}$. In particular, taking $G_{k+r+1,k}$ to be the generator matrix of any Reed-Solomon code, we only need $q \geq k + r + 1$ (since Reed-Solomon codes need $q \geq n$). For our settings q is fixed to be 256, so we have the flexibility to construct many wide explicit codes.

References

- [1] Abhishek Agarwal, Alexander Barg, Sihuang Hu, Arya Mazumdar, and Itzhak Tamo. Combinatorial alphabet-dependent bounds for locally recoverable codes. *IEEE Transactions on Information Theory*, 2018.
- [2] Backblaze. Disk Reliability Dataset. <https://www.backblaze.com/b2/hard-drive-test-data.html>, 2013-2018.
- [3] Backblaze. Erasure coding used by Backblaze. <https://www.backblaze.com/blog/reed-solomon/>, 2013-2018.
- [4] SB Balaji and P Vijay Kumar. On partial maximally-recoverable and maximally-recoverable codes. In *IEEE International Symposium on Information Theory (ISIT)*, 2015.
- [5] Alexander Barg, Zitan Chen, and Itzhak Tamo. A construction of maximally recoverable codes. *Designs, Codes and Cryptography*, 2022.
- [6] Alexander Barg, Itzhak Tamo, and Serge Vlăduț. Locally recoverable codes on algebraic curves. *IEEE Transactions on Information Theory*, 2017.
- [7] Shimrit Ben-Yair. Updating Google Photos' storage policy to build for the future. <https://blog.google/products/photos/storage-changes/>, 2020.
- [8] Johannes Bloemer, Malik Kalfane, Richard Karp, Marek Karpinski, Michael Luby, and David Zuckerman. An xor-based erasure-resilient coding scheme. 1995.
- [9] Eric Brewer. Spinning Disks and Their Cloudy Future. <https://www.usenix.org/node/194391>, 2018.
- [10] Eric Brewer, Lawrence Ying, Lawrence Greenfield, Robert Cypher, and Theodore T'so. Disks for data centers. Technical report, Google, 2016.
- [11] Viveck Cadambe and Arya Mazumdar. An upper bound on the size of locally recoverable codes. In *2013 International Symposium on Network Coding (NetCod)*. IEEE, 2013.
- [12] Han Cai, Ying Miao, Moshe Schwartz, and Xiaohu Tang. A construction of maximally recoverable codes with order-optimal field size. *IEEE Transactions on Information Theory*, 2021.
- [13] Daniel Ford, François Labelle, Florentina I Popovici, Murray Stokely, Van-Anh Truong, Luiz Barroso, Carrie Grimes, and Sean Quinlan. Availability in Globally Distributed Storage Systems. In *USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, 2010.
- [14] Garth Alan Gibson. *Redundant disk arrays: Reliable, parallel secondary storage*. PhD thesis, University of California, Berkeley, 1991.
- [15] Parikshit Gopalan, Guangda Hu, Swastik Kopparty, Shubhangi Saraf, Carol Wang, and Sergey Yekhanin. Maximally recoverable codes for grid-like topologies. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, 2017.
- [16] Parikshit Gopalan, Cheng Huang, Bob Jenkins, and Sergey Yekhanin. Explicit maximally recoverable codes with locality. *IEEE Transactions on Information Theory*, 2014.
- [17] Parikshit Gopalan, Cheng Huang, Huseyin Simitci, and Sergey Yekhanin. On the locality of codeword symbols. *IEEE Transactions on Information Theory*, 2012.
- [18] Sivakanth Gopi and Venkatesan Guruswami. Improved maximally recoverable lrcs using skew polynomials. *IEEE Transactions on Information Theory*, 2022.
- [19] Sivakanth Gopi, Venkatesan Guruswami, and Sergey Yekhanin. On maximally recoverable local reconstruction codes. In *Electron. Colloquium Comput. Complex.*, 2017.
- [20] Sivakanth Gopi, Venkatesan Guruswami, and Sergey Yekhanin. Maximally recoverable lrcs: A field size lower bound and constructions for few heavy parities. *IEEE Transactions on Information Theory*, 2020.
- [21] Matthias Grezet, Thomas Westerböck, Ragnar Freij-Hollanti, and Camilla Hollanti. Uniform minors in maximally recoverable codes. *IEEE Communications Letters*, 2019.
- [22] Venkatesan Guruswami, Satyanarayana V Lokam, and Sai Vikneshwar Mani Jayaraman. -msr codes: Contacting fewer code blocks for exact repair. *IEEE Transactions on Information Theory*, 2020.
- [23] Walker Haddock, Purushotham V Bangalore, Matthew L Curry, and Anthony Skjellum. High performance erasure coding for very large stripe sizes. In *2019 Spring Simulation Conference (SpringSim)*.
- [24] Kathryn Haymaker, Beth Malmskog, and Gretchen Matthews. Locally recoverable codes with availability $t \leq 2$ from fiber products of curves. *arXiv preprint arXiv:1612.03841*, 2016.
- [25] Yuchong Hu, Liangfeng Cheng, Qiaori Yao, Patrick PC Lee, Weichun Wang, and Wei Chen. Exploiting combined locality for {Wide-Stripe} erasure coding in distributed storage. In *USENIX File and Storage Technologies (FAST)*, 2021.

- [26] Cheng Huang, Minghua Chen, and Jin Li. Pyramid codes: Flexible schemes to trade space for access efficiency in reliable data storage systems. *ACM Transactions on Storage (TOS)*, 2013.
- [27] Cheng Huang, Huseyin Simitci, Yikang Xu, Aaron Ogus, Brad Calder, Parikshit Gopalan, Jin Li, Sergey Yekhanin, et al. Erasure Coding in Windows Azure Storage. In *USENIX Annual Technical Conference (ATC)*, 2012.
- [28] Lingfei Jin. Explicit construction of optimal locally recoverable codes of distance 5 and 6 via binary constant weight codes. *IEEE Transactions on Information Theory*, 2019.
- [29] Saurabh Kadekodi. *DISK-ADAPTIVE REDUNDANCY: tailoring data redundancy to disk-reliability heterogeneity in cluster storage systems*. PhD thesis, Carnegie Mellon University, 2020.
- [30] Saurabh Kadekodi, Francisco Maturana, Sanjith Athlur, Arif Merchant, KV Rashmi, and Gregory R Ganger. Tiger: {Disk-Adaptive} redundancy without placement restrictions. In *USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, 2022.
- [31] Saurabh Kadekodi, Francisco Maturana, Suhas Jayaram Subramanya, Juncheng Yang, KV Rashmi, and Gregory R Ganger. PACEMAKER: Avoiding heart attacks in storage clusters with disk-adaptive redundancy. In *USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, 2020.
- [32] Saurabh Kadekodi, K V Rashmi, and Gregory R Ganger. Cluster storage systems gotta have HeART: improving storage efficiency by exploiting disk-reliability heterogeneity. In *USENIX File and Storage Technologies (FAST)*, 2019.
- [33] Oleg Kolosov, Gala Yadgar, Matan Liram, Itzhak Tamo, and Alexander Barg. On fault tolerance, locality, and optimality in locally repairable codes. *ACM Transactions on Storage (TOS)*, 2020.
- [34] Yin Li, Hao Wang, Xuebin Zhang, Ning Zheng, Shafa Dahandeh, and Tong Zhang. Facilitating magnetic recording technology scaling for data center hard disk drives through {Filesystem-Level} transparent local erasure coding. In *USENIX File and Storage Technologies (FAST)*.
- [35] Jian Liu, Sihem Mesnager, and Lusheng Chen. New constructions of optimal locally recoverable codes via good polynomials. *IEEE Transactions on Information Theory*, 2017.
- [36] Shu Liu and Chaoping Xing. Maximally recoverable local reconstruction codes from subspace direct sum systems. *arXiv preprint arXiv:2111.03244*, 2021.
- [37] Gaojun Luo and Xiwang Cao. Constructions of optimal binary locally recoverable codes via a general construction of linear codes. *IEEE Transactions on Communications*, 2021.
- [38] Giacomo Micheli. Constructions of locally recoverable codes which are optimal. *IEEE Transactions on Information Theory*, 2019.
- [39] D.S. Papailiopoulos and A.G. Dimakis. Locally repairable codes. In *IEEE International Symposium on Information Theory (ISIT)*, 2012.
- [40] K V Rashmi, Nihar B Shah, Dikang Gu, Hairong Kuang, Dhruva Borthakur, and Kannan Ramchandran. A Solution to the Network Challenges of Data Recovery in Erasure-coded Distributed Storage Systems: A Study on the Facebook Warehouse Cluster. In *USENIX Workshop on Hot Topics in Storage and File Systems (HotStorage)*, 2013.
- [41] David Reinsel-John Gantz-John Rydning, J Reinsel, and J Gantz. The digitization of the world from edge to core. *Framingham: International Data Corporation*, 16, 2018.
- [42] Maheswaran Sathiamoorthy, Megasthenis Asteris, Dimitris Papailiopoulos, Alexandros G Dimakis, Ramkumar Vadali, Scott Chen, and Dhruva Borthakur. Xoring elephants: Novel erasure codes for big data. In *International Conference on Very Large Data Bases (VLDB)*, 2013.
- [43] Seagate. The Digitization of the World From Edge to Core. <https://www.seagate.com/files/www-content/our-story/trends/files/idc-seagate-dataage-whitepaper.pdf>, 2018.
- [44] Mostafa Shahabinejad. Locally repairable linear block codes for distributed storage systems. 2018.
- [45] Natalia Silberstein, Ankit Singh Rawat, O Ozan Koyluoglu, and Sriram Vishwanath. Optimal locally repairable codes via rank-metric codes. In *IEEE International Symposium on Information Theory (ISIT)*, 2013.
- [46] Itzhak Tamo and Alexander Barg. Bounds on locally recoverable codes with multiple recovering sets. In *IEEE Transactions on Information Theory*, 2014.
- [47] Itzhak Tamo and Alexander Barg. A family of optimal locally recoverable codes. *IEEE Transactions on Information Theory*, 2014.
- [48] Itzhak Tamo, Alexander Barg, and Alexey Frolov. Bounds on the parameters of locally recoverable codes. *IEEE Transactions on Information Theory*, 2016.

- [49] Itzhak Tamo, Dimitris S Papailiopoulos, and Alexandros G Dimakis. Optimal locally repairable codes and connections to matroid theory. *IEEE Transactions on Information Theory*, 2016.
- [50] VAST. Providing Resilience, Efficiently. <https://www.usenix.org/node/194391>, 2019.
- [51] Guanghui Zhang. A new construction of optimal (r, δ) locally recoverable codes. *IEEE Communications Letters*, 2020.