Maximally Recoverable Codes: the Bounded Case

Venkata Gandikota   Elena Grigorescu   Clayton Thomas   Minshen Zhu

***Abstract*— Modern distributed storage systems employ Maximally Recoverable codes that aim to balance failure recovery capabilities with encoding/decoding efficiency tradeoffs. Recent works of Gopalan et al [SODA 2017] and Kane et al [FOCS 2017] show that the alphabet size of grid-like topologies of practical interest must be large, a feature that hampers decoding efficiency.**

To bypass such shortcomings, in this work we initiate the study of a weaker version of recoverability, where instead of being able to correct all correctable erasure patterns (as is the case for maximal recoverability), we only require to correct all erasure patterns of *bounded* size. The study of this notion reduces to a variant of a combinatorial problem studied in the literature, which is interesting in its own right.

We study the alphabet size of codes withstanding all erasure patterns of small (constant) size. We believe the questions we propose are relevant to both real storage systems and combinatorial analysis, and merit further study.

I. Introduction

Modern distributed storage systems need to address the challenge of storing large amounts of data, with small overhead, while providing reliable recovery in the face of failures. Unlike in communication settings, if a failure occurs in a storage system, it is typically easy to detect where it occurs (e.g., rack or data center failure are obvious to the system). Hence, recent trends in practical systems adopt *erasure coding* schemes with fast encoding and decoding capabilities [BHH13], [HSX$^+$12], [SLR$^+$14].

To maximize reliability, the codes employed require large fields, which is not a desirable feature when performing many algebraic computations for encoding and decoding. In addition, such codes need to be able to withstand *correlated failures*, since it is often the case that multiple machines from the same location experience failures at the same time (e.g., data center failure). These design choices result in parameter tradeoffs that have recently triggered active research in the coding theory community [CHL07], [DRWS11], [Yek12], [GHSY12], [HSX$^+$12], [BB12], [BHH13], [TB14b], [GHJY14], [BK15], [LL15], [HY16], [GHK$^+$17].

The notion of Maximally Recoverable (MR) code [CHL07], [GHJY14] captures the design choices that practical distributed storage systems face. Data can be stored as $m \times n$ matrices with entries from a finite field $\mathbb{F}$. Every row satisfies the same set of a parity constraints, and every column satisfies the same set of b parity constraints. In addition, there are h global parity checks constraints, which could involve arbitrary entries from the matrix. This view of the code was recently defined by Gopalan *et al.* [GHK$^+$17] and denoted by the topology $T_{m\times n}(a, b, h)$. They also propose the two-step design: (1) determine the *topology* of the code, by determining the support of the parity check matrix and incorporating the knowledge of the correlated failures– this is possible since the layout of the data is known in advance; (2) specify the finite field $\mathbb{F} = \mathbb{F}_q$ and the coefficients appearing in the parity check matrix. Further, to correct erasures, it is enough to solve a system of linear equations over $\mathbb{F}$. A code is said to be *maximally recoverable* if it corrects every erasure pattern that can be corrected for some fixing of the coefficients in the given topology.

The study of maximally recoverable codes has revealed that even for basic topologies such as $T(1, 1, 1)$, maximal recoverability requires fields of super polynomial sizes [GHK$^+$17], [KLR17]. This is a topology of practical interest and fields of small size are highly desirable [PGM13]. Here we focus on this particular topology, and propose a weaker notion of recoverability that allows us to obtain explicit erasure schemes for polynomial size fields. Specifically, we focus on erasure patterns with a *bounded* number of erasures. It turns out that the difficult patterns are the so-called *irreducible* patterns. An *irreducible* erasure pattern for the $T(1, 1, 1)$ topology is a set of non-trivial erasures, in the sense that each such erasure is not the only erasure in some row/column parity check equation, and hence in order to correct it one needs to solve a system of multiple linear equations. In a general topology, iterative row-column decoders are first used to perform Gaussian elimination locally (using only row/column equations), after which the decoder has to resolve non-local (irreducible)

Purdue University. Research supported by NSF CCF-1649515. Email: vgandiko@purdue.edu

Purdue University. Research supported by NSF CCF-1649515. Email: elena-g@purdue.edu

Purdue University. Email: thoma466@purdue.edu

Purdue University. Email: minshen.zh@gmail.com

erasures patterns. Correcting irreducible erasures is the most expensive part of the decoding process, as it may involve many data disk reads, as well as communication between multiple data centers.

We define an e-*Maximally Recoverable (e-MR)* code to be a code capable of correcting every correctable erasure pattern consisting of at most e erasures. We focus on e-MR codes for the $T(1,1,1)$ topology, with $e = O(1)$. Withstanding a small number of erasures is of practical importance, since at any given time the number of failures in a distributed storage system is typically small. In addition, patterns with few erasures are often more important than those with a large number of erasures [Yek17]. Distributed erasure schemes concerned with recovering from only one, or just a constant number of erasures, have been studied in the active area of local reconstruction and regeneration, where the goal is designing codes with fast (local) decoding from the point of view of disk reads and communication, respectively [IKOS04], [DRWS11], [GHSY12], [SAP+13], [HCL13], [DGRS14], [KPLK14], [WZ14], [TB14a], [FVY15], [BE16], [RPDV16], [TBF16].

A. Our contributions

We initiate the study of e-Maximally Recoverable codes, and obtain explicit constructions, as well as lower bounds, for codes withstanding erasure patterns of size up to $e \leq 12$. Our results improve upon the bounds implied by adapting previous works to our setting.

We restrict our attention to fields $\mathbb{F}$ of characteristic 2 (which is also desired in applications). Let $\mathbb{F}_2 = \{0,1\}$ be the finite field on 2 elements.

[GHK+17] proves that understanding the field size for MR codes reduces to the following combinatorial problem: Let $K_{m,n}$ be a complete bipartite graph and $\gamma : [m] \times [n] \to \mathbb{F}_2^d$ be a labeling of its edges, such that for any simple cycle C, the sum of the edge-labels in the cycle $\sum_{c \in C} \gamma(c)$ is non-zero. What is the smallest (asymptotically, as a function of m, n) value of d for which such a labeling exists? For $m = n$, [GHK+17] shows that $d = \Omega((\log n)^2)$, and [KLR17] improves it to $d \geq n/2 - 2$. [KLR17] further provides an explicit construction for $d = 3n$. Our problem reduces to the following variant of the question:

Question 1.1: Given integer $e > 0$, when does there exist a labeling $\gamma : [m] \times [n] \to \mathbb{F}_2^d$ of the edges of a complete bipartite graph $K_{m,n}$, such that for every simple cycle C *of length at most* e,

$$\sum_{c \in C} \gamma(c) \neq 0?$$

a) Upper bounds: [GHJY14] suggests labeling the mn edges with elements of an e-wise independent set. A set $S \subseteq \mathbb{F}$ is said to be e-wise independent over $\mathbb{F}_2$ if every $T \subseteq S$ with $|T| \leq e$ is linearly independent over $\mathbb{F}_2$. They show that there exists a set $S = \{\alpha_1, \alpha_2, \cdots, \alpha_{mn}\}$, where $\alpha_i \in \mathbb{F}_{2^{\lceil \log mn \rceil e/2}}$, such that S is e-wise independent over $\mathbb{F}_2$ (see full-version for the proof).

A typical setting in applications is $m = n$, in which case their results work over fields of size $O(n^e)$. We show an explicit construction with $O(n^{e/2-1})$ field size, albeit only for $e \leq 12$.

Theorem 1.2: For $e \leq 12$, there exists an e-MR code for $T(1,1,1)$ over fields of size $O(\max\{m,n\}^{e/2-1})$.

We construct an explicit labeling for bounded-length cycles, and analyze it using properties of Moore matrices, linear algebra and basic combinatorial properties of the kernels of such matrices.

b) Lower bounds: A direct corollary of a result of [GHK+17] implies that a labeling $\gamma : [n] \times [n] \to \mathbb{F}_2^d$ of cycles of length at most e must have $d \geq \log\left(\frac{e}{2}\right) \cdot \log\left(\frac{n}{e}\right) + \left(\frac{e}{2} - 1\right) \log\left(1 - \frac{e}{4n}\right)$ (see full version for the proof). We compare our bounds against this bound.

We obtain the following result.

Theorem 1.3: For $e \leq 12$, every e-MR code for $T(1,1,1)$ requires fields of size $\Omega(n^{\lceil \log e \rceil - 1})$.

In particular, for $e = 6$ we improve from $\Omega(n^{\log 3})$ to $\Omega(n^2)$, and for $e = 10$ we improve from $\Omega(n^{\log 5})$ to $\Omega(n^3)$.

The lower bound proof for unbounded cycle length in [KLR17] uses representation theory arguments that appear difficult to adapt to bounded-length cycles. Our proof reduces to elementary combinatorial arguments analyzing the chromatic number of a graph G, whose nodes represent paths in $K_{n,n}$ and the edges connect pairs of paths whose symmetric difference is a simple cycle of bounded length.

We remark that neither our techniques for the upper bounds, nor the ones for the lower bounds generalize to larger values of e, and hence new approaches are needed for further progress.

We believe the notion of bounded recoverability introduced here is well-motivated in practice, and that obtaining tight asymptotics for the combinatorial Question 1.1 is difficult in general and merits further study.

II. Preliminaries

As defined in [GHK+17], a code $C(\{\alpha_j^k\}, \{\beta_i^k\}, \{\gamma_{ij}^k\})$ instantiating the *grid-like topology* $T_{m \times n}(a, b, h)$ requires the following three sets of values:

1) Coefficients of the row constraints: $\{\alpha_j^{(k)}\}_{j\in[n],k\in[a]}$,
2) Coefficients of the column constraints: $\{\beta_i^{(k)}\}_{i\in[m],k\in[b]}$,
3) Coefficients of the global constraints: $\{\gamma_{ij}^{(k)}\}_{i\in[m],j\in[n],k\in[h]}$.

The code symbols $\{x_{ij}\}_{i\in[m],j\in[n]}$ must satisfy the following parity check equations:

$$\forall i \in [m], k \in [a], \qquad \textstyle\sum_{j=1}^{n} \alpha_j^{(k)} x_{ij} = 0, \tag{1}$$

$$\forall j \in [n], k \in [b], \qquad \textstyle\sum_{i=1}^{m} \beta_i^{(k)} x_{ij} = 0, \tag{2}$$

$$\forall k \in [h], \qquad \textstyle\sum_{i=1}^{m}\sum_{j=1}^{n} \gamma_{ij}^{(k)} x_{ij} = 0. \tag{3}$$

A *failure pattern* is specified by a subset of positions in the grid: $E \subseteq [m] \times [n]$. Pattern E is said to be *correctable* for a topology T if there exists a code instantiating T, such that the variables $\{x_{ij}\}_{(i,j)\in E}$ can be recovered from the parity check equations.

A code instantiating the topology T is *Maximally Recoverable (MR)* if it corrects every correctable failure patterns of T. [GHK$^+$17] and [KLR17] show that every MR code for $T_{m\times n}(a,b,h)$ requires exponentially large field size if $a,b,h \geq 1$. As this severely hampers the practicality of MR codes, one naturally asks for more efficient ‘approximate’ recoverability.

To this end, we define e-Maximally Recoverable codes as follows.

Definition 2.1 (e-Maximally Recoverable Codes): A code $\mathcal{C}$ instantiating the topology $T_{m\times n}(a,b,h)$ is e-Maximally Recoverable (e-MR) if $\mathcal{C}$ corrects every correctable pattern of size $\leq e$.

This work initiates the study of optimal field size for e-MR codes and focuses on a few constant values of e, and on the topology $T_{m\times n}(1,1,1)$. This is a topology with practical applicability and has been well-studied in the recent literature [GHK$^+$17], [KLR17]. Moreover, lower bounds on $T_{m\times n}(1,1,1)$ implies general lower bounds for $T_{n\times n}(1,1,h)$, as described in the following theorem implied in previous work. We refer the interested readers to the full version for a proof.

Theorem 2.2: (Implied in [GHK$^+$17]) Let $f : \mathbb{N} \times \mathbb{N} \to \mathbb{N}$ be a function such that any e-MR code for $T_{n\times n}(1,1,1)$ requires field size $f(n,e)$. Then any e-MR code for $T_{n\times n}(1,1,h)$ requires field size $f(\lfloor n/h \rfloor, \lfloor e/h \rfloor)$.

In $T_{m\times n}(1,1,h)$, there is only one parity check equation for every row and column, and h general parity equations. As mentioned before, one can perform an elimination process that iteratively recovers an erasure that appears as the only erasure on its row or column. The process uses only that row/column's parity check equation, and is therefore computationally cheap. The remaining set of erasures form an *irreducible* erasure pattern. We call an irreducible pattern of size $\leq e$ an e-*irreducible* pattern.

As previously alluded to, we will use the following lemma as the starting point of our proofs.

Lemma 2.3: (Corollary to Lemma 15 and Corollary 17 from [GHK$^+$17]) A correctable e-irreducible pattern for $T(1,1,1)$ corresponds to a simple cycle C of length $\leq e$ in $K_{m,n}$. The pattern is correctable by a particular instantiation $\mathcal{C} = \mathcal{C}(\{\gamma_{ij}\}_{(i,j)\in[m]\times[n]})$ iff $\sum_{(i,j)\in C} \gamma_{ij} \neq 0$. (In particular, we may assume that the remaining coefficients α_j's and β_i's of the topology in 1 have value 1.)

Throughout the paper we shall assume that all logs are base 2.

III. Construction

In this section we prove Theorem 1.2. We assume $m = n$ for simplicity of presentation.

Let $\mathbb{F}$ be a field of size $2^{\lceil \log n \rceil}$, and $\ell > 1$ an integer. For sets $\mathcal{P} = \{P_1, \ldots, P_n\} \subseteq \mathbb{F}$, and $\mathcal{Q} = \{Q_1, \ldots, Q_n\} \subseteq \mathbb{F}$, define the labeling $\gamma_{\ell,\mathcal{P},\mathcal{Q}}$ of $K_{n,n}$ by

$$\begin{aligned}\gamma(i,j) &= \gamma_{\ell,\mathcal{P},\mathcal{Q}}(i,j) \\ &= \left(P_iQ_j, P_i^2Q_j, P_i^4Q_j, \ldots, P_i^{2^{\ell-2}}Q_j\right) \in \mathbb{F}^{\ell-1}.\end{aligned}$$

Note that since $\mathbb{F} \simeq (\mathbb{F}_2)^{\lceil \log n \rceil}$, we may view γ as a function $\gamma : [n] \times [n] \to (\mathbb{F}_2)^{(\ell-1)\lceil \log n \rceil}$.

Suppose $\mathcal{P}$ and $\mathcal{Q}$ are each sets of distinct values. We show that γ defines a labeling of $K_{n,n}$ that has no zero-cycles of length $e = 2\ell$, for $e \leq 12$. Since the first $\ell' - 1$ coordinates of γ correspond to $\gamma_{\ell',\mathcal{P},\mathcal{Q}}$, this will imply that $\gamma(C) \neq 0$ for all cycles of length $2\ell' < 2\ell$ as well. This implies the existence of an e-MR code for the topology $T_{n\times n}(1,1,1)$ over a field of size $O(n^{\ell-1})$.

Consider a simple cycle in $K_{n,n}$ of length $e = 2\ell$ given by

$$C = i_1 \to j_1 \to i_2 \to j_2 \to \ldots \to i_\ell \to j_\ell \to i_1$$

We start with a lemma breaking this cycle into a union of 4-cycles. The value of γ on 4-cycles has a simple description, which will allow us to write $\gamma(C)$ as a product of a matrix and a vector that we can analyze using basic linear algebra.

Lemma 3.1: Let γ be any labeling of $K_{n,n}$ over a characteristic two field. Any simple cycle C in $K_{n,n}$ of length $e = 2\ell$ can be decomposed into $\ell - 1$ cycles $C_1, \ldots, C_{\ell-1}$ of length 4 such that $\gamma(C) = \gamma(C_1) +$

$\ldots + \gamma(C_{\ell-1})$. Specifically, if C is given as above, we can take

$$C_m = i_1 \to j_m \to i_{m+1} \to j_{m+1} \to i_1$$

Proof: Induct on ℓ. Note that the result holds trivially for $\ell = 2$. Let $\ell > 2$ and suppose that our result holds for cycles of length $\ell - 1$. Thus, the cycle

$$C' = i_1 \to j_1 \to i_2 \to j_2 \to \ldots \to i_{\ell-1} \to j_{\ell-1} \to i_1$$

decomposes into $\ell-2$ cycles $C_1, \ldots, C_{\ell-2}$. Let $C_{\ell-1} = i_1 \to j_{\ell-1} \to i_\ell \to j_\ell \to i_1$. Because we work over fields of characteristic 2, we get $\gamma(C) = \gamma(C') + \gamma(C_{\ell-1})$. Thus $\gamma(C) = \gamma(C_1) + \ldots + \gamma(C_{\ell-2}) + \gamma(C_{\ell-1})$, as desired. ■

For $k \in [\ell-1]$, denote by $\gamma_k(i,j) \in \mathbb{F}$ the kth coordinate of $\gamma(i,j)$, when identified as a vector of length $\ell - 1$ over $\mathbb{F}$. By lemma 3.1, we have

$$\begin{aligned}
\gamma_k(C) &= \sum_{m=1}^{\ell-1} \gamma_k(C_m) \\
&= \sum_{m=1}^{\ell-1} \gamma_k(i_1, j_m) + \gamma_k(i_{m+1}, j_m) \\
&\quad + \gamma_k(i_{m+1}, j_{m+1}) + \gamma_k(i_1, j_{m+1}) \\
&= \sum_{m=1}^{\ell-1} P_{i_1}^{2^k} Q_{j_m} + P_{i_{m+1}}^{2^k} Q_{j_m} \\
&\quad + P_{i_{m+1}}^{2^k} Q_{j_{m+1}} + P_{i_1}^{2^k} Q_{j_{m+1}} \\
&= \sum_{m=1}^{\ell-1} (P_{i_1} + P_{i_{m+1}})^{2^k} (Q_{j_m} + Q_{j_{m+1}})
\end{aligned}$$

Let $a_m = P_{i_1} + P_{i_{m+1}}$ and let $b_m = Q_{j_m} + Q_{j_{m+1}}$ for $m = 1, \ldots, \ell - 1$. Define

$$M = \begin{pmatrix} a_1 & a_2 & \cdots & a_{\ell-1} \\ a_1^2 & a_2^2 & \cdots & a_{\ell-1}^2 \\ a_1^4 & a_2^4 & \cdots & a_{\ell-1}^4 \\ \vdots & \vdots & & \vdots \\ a_1^{2^{\ell-2}} & a_2^{2^{\ell-2}} & \cdots & a_{\ell-1}^{2^{\ell-2}} \end{pmatrix}, \quad b = \begin{pmatrix} b_1 \\ b_2 \\ \vdots \\ b_{\ell-1}. \end{pmatrix}$$

so that we have

$$\gamma(C) = Mb.$$

A matrix M of the given form is sometimes called a Moore matrix.

The following pair of claims specify useful properties of $\{a_1, \ldots, a_{\ell-1}\}$ and $\{b_1, \ldots, b_{\ell-1}\}$, and follow easily from the distinctness of the P_i's and Q_j's.

Claim 3.2: The set $\{a_1, \ldots, a_{\ell-1}\}$ are distinct and nonzero.

Definition 3.3 (Zero adjacent sums): Consider a vector $b = (b_1, \ldots, b_N)$. Define a zero adjacent sum of b as a set of consecutive indices $i, i+1, \ldots, j$ (for some $1 \le i \le j \le N$) such that $b_i + b_{i+1} + \ldots + b_j = 0$. We denote the set of indices $\{i, i+1, \ldots, j\}$ by $[i,j]$.

We say $[i,j]$ is a zero adjacent sum for a collection of vectors if it is a zero adjacent sum for each vector in the collection.

Claim 3.4: The vector $b = (b_1, \ldots, b_{\ell-1})$ has no zero adjacent sums.

Our goal is now to show that, given that the a_is are distinct, every element of the kernel of M has a zero adjacent sum. Thus, $Mb \neq 0$ for every b as given above.

The next lemma tells us that the kernel of M, i.e. the set of linear dependencies of the columns of M, is determined by the linear dependencies of $a_1, a_2, \ldots, a_{\ell-1}$ (when we treat $\mathbb{F}$ as a vector space over $\mathbb{F}_2$). This will allow us to work with a basis of the kernel of M whose coordinates are all 0 or 1.

Lemma 3.5: Let $T : (\mathbb{F}_2)^{\ell-1} \to \mathbb{F}$ be given by $T(b_1, \ldots, b_{\ell-1}) = b_1 a_1 + \ldots + b_{\ell-1} a_{\ell-1}$. Then

$$\ker M = \operatorname{span}_{\mathbb{F}} \ker T$$

Proof: Because we work over a characteristic two field, $b_1 a_1^{2^m} + \ldots + b_{\ell-1} a_{\ell-1}^{2^m} = (b_1 a_1 + \ldots + b_{\ell-1} a_{\ell-1})^{2^m}$. Thus each $\beta \in \ker T$ also satisfies $M\beta = 0$. Thus any linear combination over $\mathbb{F}$ of elements of $\ker T$ is in $\ker M$, i.e. we get $\ker M \supseteq \operatorname{span}_{\mathbb{F}} \ker T$.

To complete the proof, we will compare the dimensions of $\ker M$ and $\operatorname{span}_{\mathbb{F}} \ker T$. Note that $\dim \ker T = \dim \operatorname{span}_{\mathbb{F}} \ker T$, because a set of linearly independent vectors in $(\mathbb{F}_2)^{\ell-1}$ will still be linearly independent in $\mathbb{F}^{\ell-1}$. Indeed, if $\beta_1, \ldots, \beta_k$ form a basis for $\ker T$, then the matrix whose columns are the first k entries of $\beta_1, \ldots, \beta_k$ will have nonzero determinant (over $\mathbb{F}$ or over $\mathbb{F}_2$).

Let $m = \text{rank } T \le \ell - 1$, i.e. suppose there exist maximally m vectors among $\{a_1, a_2, \cdots, a_{\ell-1}\}$ which are linearly independent over $\mathbb{F}_2$. Suppose without loss of generality that $\{a_1, a_2, \cdots, a_m\}$ is a maximal linearly independent (over $\mathbb{F}_2$) set of vectors among $\{a_1, a_2, \cdots, a_{\ell-1}\}$ It is known (see Lemma 3.51 in [LN86]) that if $\{a_1, a_2, \cdots, a_m\}$ are linearly independent over $\mathbb{F}_2$, then the Moore matrix

$$\begin{pmatrix} a_1 & a_2 & \cdots & a_m \\ a_1^2 & a_2^2 & \cdots & a_m^2 \\ a_1^4 & a_2^4 & \cdots & a_m^4 \\ \vdots & \vdots & & \vdots \\ a_1^{2^{m-1}} & a_2^{2^{m-1}} & \cdots & a_m^{2^{m-1}} \end{pmatrix}$$

is invertible. In particular, the vectors

$$\begin{pmatrix} a_1 \\ a_1^2 \\ a_1^4 \\ \vdots \\ a_1^{2^{\ell-2}} \end{pmatrix}, \begin{pmatrix} a_2 \\ a_2^2 \\ a_2^4 \\ \vdots \\ a_2^{2^{\ell-2}} \end{pmatrix}, \dots, \begin{pmatrix} a_m \\ a_m^2 \\ a_m^4 \\ \vdots \\ a_m^{2^{\ell-2}} \end{pmatrix}$$

are linearly independent over $\mathbb{F}$.

Thus, there are at least m columns of M which are linearly independent over $\mathbb{F}$. Thus rank $M \geq m =$ rank T, so $\dim\ker M \leq \dim\ker T = \dim \operatorname{span}_{\mathbb{F}} \ker T$.

Thus $\ker M = \operatorname{span}_{\mathbb{F}} \ker T$, as desired. ■

Corollary 3.6: There exists a basis of $\ker M$ such that all of the entries of the vectors in the basis are 0 or 1.

We note that a variant of the above lemma holds, with a similar proof, for any finite field $\mathbb{F}_q$.

Let $\mathcal{B}$ be any basis for $\ker T$. We will say that the *weight* of a vector is it's number of nonzero coordinates. Because $\{a_1, \dots, a_{\ell-1}\}$ are distinct and nonzero, no nonzero vector in $\ker T$ has weight 1 or 2. Further, any sum of a subset of elements in $\mathcal{B}$ is also in $\ker T$, so the sum of a subset of elements in $\mathcal{B}$ must have weight > 2 as well.

Consider a basis $\mathcal{B} = \{\beta_1, \dots, \beta_m\}$ for $\ker T$ which is in reduced echelon form. In other words, the matrix whose rows are $\beta_1, \dots, \beta_m$ is in reduced row echelon form as given by (†).

Note that given any basis, we can perform "row operations" to construct a basis of the form given by (†). We will refer to the columns of (†) in which leading 1's appear as *pivot columns* and all other columns, where $*$'s are present, as *non-pivot columns.*

In this form, any sum of three or more basis vectors will automatically have weight at least 3. Thus the condition that all subset sums have weight at least three becomes equivalent to the following two conditions on (†):

- At least two non-pivot positions in every row are equal to 1.
- For every pair of rows, the non-pivot positions are not all equal.

We use the above two facts to bound how large $\dim\ker M$ can possibly be:

Lemma 3.7: $\dim\ker T \leq \ell - 1 - \log\ell$. In other words, rank $M \geq \log\ell$.

Proof: Let $d = \dim\ker T$ be the number of basis vectors in (†). Let $\ell - 1 = d + t$, so that t equals the number of non-pivot columns of (†). There are $2^t - t - 1$ possible bitstrings of length t which have weight at least 2. The starred positions of each row must correspond to distinct such bitstrings. Thus $2^t - t - 1 \geq d$. So we have $\ell - 1 \leq 2^t - 1$ and $d \leq \ell - 1 - \log\ell$. ■

We are now ready to prove that under the conditions above we must have $Mb \neq 0$ for $\ell \leq 6$. Theorem 1.2 will follow immediately from the observation that by Lemma 3.7 we have $\dim\ker M \in \{0, 1, 2\}$ for $\ell \leq 6$, and by the next lemma dealing with these cases.

Lemma 3.8: If $d = \dim\ker M \in \{1, 2\}$, then $b \notin \ker M$.

Proof: Observe that if $[i, j]$ is a zero adjacent sum for every element $\beta \in \mathcal{B}$, then $[i, j]$ is a zero adjacent sum for every element of $\operatorname{span}_{\mathbb{F}}\mathcal{B}$. By Lemma 3.4, this means $b \notin \operatorname{span}_{\mathbb{F}}\mathcal{B} = \ker M$.

If $d = 1$, then $\ker M$ is spanned by a single binary vector β. It is easy to check that there must exist $i \leq j$ such that $[i, j]$ is a zero adjacent sum for β.

Now let $d = 2$ and suppose for the sake of contradiction that no $[i, j]$ is a zero adjacent sum for β_1 and β_2. By lemma 3.7, we need $\ell \geq 6$, so (†) has at least 5 columns. Consider the values of the $*$s in the first row between the first and second pivot. Because there can be no zero columns in (†), each of these $*$s must be 1. But if there are more than 0 such $*$s, then $[1, 2]$ forms a zero adjacent sum for $\mathcal{B}$.

Thus, (†) is of the form

$$\begin{pmatrix} 1 & 0 & * & & * \\ 0 & 1 & * & \cdots & * \end{pmatrix}$$

Now we can apply a case analysis on the third (then fourth) columns of the above. If the third column is $\begin{pmatrix} 1 \\ 1 \end{pmatrix}$ or $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$ then there is a zero adjacent sum. Only the column $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$ does not create a zero adjacent sum among the first three columns of $\mathcal{B}$. Finally, observe that adding any fourth column to

$$\begin{pmatrix} 1 & 0 & 1 & * & & * \\ 0 & 1 & 0 & * & \cdots & * \end{pmatrix}$$

will produce a zero adjacent sum. ■

Remark 3.9: It turns out that this result does not hold for larger values of e. For example, when $\ell = 7$ and the kernel of M has dimension 3, we can have (†) take the form

$$\begin{pmatrix} - & \beta_1 & - \\ - & \beta_2 & - \\ - & \beta_3 & - \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & 1 & 1 & 1 \\ 0 & 1 & 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 & 0 \end{pmatrix},$$

which satisfies the property that no nonempty subset of the basis vectors is zero, yet does not have any zero adjacent sums. It is easy to find a distinct set of elements

$$\begin{pmatrix} - & \beta_1 & - \\ - & \beta_2 & - \\ - & \beta_3 & - \\ & \vdots & \\ - & \beta_d & - \end{pmatrix} = \begin{pmatrix} 1 & * & \cdots & * & 0 & * & & * & 0 & * & & * & 0 & * & \\ & & & & 1 & * & \cdots & * & 0 & * & & * & 0 & * & \\ & & & & & & & & 1 & * & \cdots & * & 0 & * & \\ & & & & & & & & & & & & 1 & * & \cdots \\ 0 & & & & & & & & & & & & & & \ddots \end{pmatrix} \qquad (\dagger)$$

$\{\alpha_1, \dots, \alpha_6\}$ of $\mathbb{F}$ (for instance, take $\mathbb{F} = \mathbb{F}_8$) satisfying all of the linear dependencies given by $\{\beta_1, \beta_2, \beta_3\}$. From there it is easy to find an element b of their span without zero adjacent sums, and from there a set of distict $\{P_1, \dots, P_7\}$ such that $b_m = P_m + P_{m+1}$.

IV. Lower bounds

In this section we prove Theorem 1.3, thus showing our alphabet lower bounds for e-MR codes with $e \leq 12$, for the topology $T_{m\times n}(1,1,1)$. These bounds can also be extended to $T_{m\times n}(1,1,h)$ by Theorem 2.2. We will assume $m = n$ for simplicity of presentation. Any lower bound with $n' = \min\{m,n\}$ for $T_{n'\times n'}(1,1,h)$ trivially holds for $T_{m\times n}(1,1,h)$.

A. High-level strategy

We adopt the same initial strategy used for $e = 2n$ in [KLR17], which was also implicit in [GHK$^+$17]. As in these works, the proof relies on the fact that we are labeling edges with elements in a field of characteristic 2, and the proof easily extends to abelian groups.

Suppose $\gamma : [n] \times [n] \to \mathbb{F}_2^d$ is an edge labeling of the complete bipartite graph $K_{n,n}$ such that for any simple cycle C of length at most e, $\sum_{c\in C} \gamma(c) \neq 0$. One can obtain a lower bounds on $|\mathbb{F}_2^d|$ as a function of the chromatic number of the following graph. Let V be a collection of paths in $K_{n,n}$. Let G be the graph with vertex set V and edges between two paths $p, p' \in V$ whenever $p \oplus p'$ is a simple cycle in $K_{n,n}$ of length at most e. Here the symmetric difference $p \oplus p' = x\{e | e \in p \text{ or } e \in p' \text{ but not both}\}$ is defined by treating p, p' as sets of edges. Denote by $\chi(G)$ the chromatic number of G, i.e. the minimum number of colors in a proper coloring of G.

Lemma 4.1: $|\mathbb{F}_2^d| = 2^d \geq \chi(G)$, and thus $d \geq \lceil \log \chi(G) \rceil$.

Proof: Consider the following coloring $\sigma : V \to \mathbb{F}_2^d$ of G: $\forall p \in V, \quad \sigma(p) = \sum_{c\in p} \gamma(c)$. This coloring is indeed proper because $\forall p, p' \in V, \quad \sigma(p) - \sigma(p') = \sum_{c\in p\oplus p'} \gamma(c)$ is non-zero if $p \oplus p'$ is a simple cycle of length at most e. Therefore adjacent vertices in G receive distinct values from σ. It follows that $|\mathbb{F}_2^d| \geq \chi(G)$. ∎

We will also use the following simple lemma.

Lemma 4.2: Let $\alpha(G)$ be the independence number of G, i.e. the size of a maximum independent set of G. Then $\chi(G) \geq \frac{|G|}{\alpha(G)}$.

Proof: A coloring of G is a partition of G into independent sets. ∎

B. The cases $e = 4, 6, 8$

As a warm-up, we start with cycles of length 4 or 6.

Theorem 4.3: Let $\gamma : [n] \times [n] \to \mathbb{F}_2^d$ be a labeling of the edges of the complete bipartite graph $K_{n,n}$ such that for any simple cycle C of length 4, $\sum_{c\in C} \gamma(c) \neq 0$. Then $d \geq \lceil \log n \rceil$.

Proof: Fix $s \neq t \in [n]$ and consider the set of paths $V = \{(s, a, t) \mid a \in [n]\}$. Clearly $|V| = n$ and any two paths in V form a simple cycle of length 4, so they should receive distinct weights. ∎

Theorem 4.4: Let $\gamma : [n] \times [n] \to \mathbb{F}_2^d$ be a labeling of the edges of the complete bipartite graph $K_{n,n}$ such that for any simple cycle C of length at most 6, $\sum_{c\in C} \gamma(c) \neq 0$. Then $d \geq 2\lceil \log n \rceil - O(1)$.

Proof: Fix $s, t \in [n]$ and consider the set of paths

$$V = \{(s, a, b, t) \mid a \in [n] \setminus \{t\}, b \in [n] \setminus \{s\}\}.$$

Clearly $|V| = (n-1)^2$. For any two different paths $p_1, p_2 \in V$, if they do not share any vertex other than s and t they form a simple cycle of length 6. Otherwise, they share exactly one other vertex and $p_1 \oplus p_2$ is a simple cycle of length 4. Therefore any two paths in V receive distinct weights and $d \geq \lceil \log(n-1)^2 \rceil = 2\lceil \log n \rceil - O(1)$. ∎

For the case $e = 8$ we simply notice that the $\Omega(n^2)$ lower bound for $e = 6$ holds, since in this case more cycles are required to have non-zero weights. We briefly comment why the above proof strategy does not give us a better lower bound for this case.

Following the strategy, we can first build a graph G on a set of paths

$$V = \{(v_1, i_1, j, i_2, v_2) \mid i_1 \neq i_2 \in [n], j \in [n] \setminus \{v_1, v_2\}\}$$

where $v_1, v_2 \in [n]$ are two fixed vertices. If two paths share the same j vertex, then they are connected to each other only when they share some additional vertex(i_1 or i_2). This shows that those paths with some fixed j can be colored with at most $O(n)$ colors, since the maximum degree is $O(n)$. Therefore the chromatic number $\chi(G)$ is at most $O(n^2)$, namely $O(n)$ colors for each of the $(n-2)$ choices of j. In fact we can easily construct a proper coloring of G using $O(n^2)$ colors. Unfortunately that does not directly imply a zero-cycle free labeling either.

C. *The cases* $e = 10, 12$

The case $e = 10$ is the main result of this section. As before, for $e = 12$ the proof strategy does not extend, and for this case we trivially have the $\Omega(n^3)$ lower bound from $e = 10$. The following combinatorial lemma will be useful in the proof.

Lemma 4.5: Let $m, n \geq 2$ be integers. Suppose on a $[m] \times [n]$ grid we select $m + n - 1$ different positions. Then there must exist $a \neq a' \in [m]$ and $b \neq b' \in [n]$ such that the following positions

$$(a, b), \quad (a, b'), \quad (a', b)$$

are all selected.

Proof: We use induction on $m + n$. For the base case $m + n = 4$, the lemma is trivially true when we select 3 positions out of a 2×2 grid. Suppose the lemma is true for all (m, n) pairs with $m+n \leq k$ and $m, n \geq 2$. We will prove the lemma for $m + n = k + 1$.

If $m = 2$ or $n = 2$ the proof follows immediately. Assume $m, n > 2$. Notice that the number of selected positions $m + n - 1$ is more than the number of rows m. By the pigeonhole principle, there is one row m_0 on which there are $t \geq 2$ selected positions $(m_0, n_1), (m_0, n_2), \cdots, (m_0, n_t)$. If there is another selected position whose column falls into the set $\{n_1, n_2, \cdots, n_t\}$ then the proof is done. Otherwise, we consider the subgrid $[m] \setminus \{m_0\} \times [n] \setminus \{n_1, n_2, \cdots, n_t\}$. This subgrid has $m-1$ rows and $n-t$ columns. Notice that in order to fit in the remaining $m + n - 1 - t$ positions, $n-t$ should be at least 2 in this case. Now that $(m-1)+(n-t) = k-t \leq k$ and $m-1 \geq 2, n-t \geq 2$, and we still have $m+n-1-t \geq (m-1)+(n-t)-1$ selected positions remaining, the induction hypothesis will guarantee the existence of a, b, a', b' in the subgrid. ∎

The following theorem is the main result of this section. It gives a $\Omega(n^3)$ lower bound on the field size required to correct every cycle of length up to 10, which slightly improves the previous bound $\Omega(n^{\log 5})$.

Theorem 4.6: Let $\gamma : [n] \times [n] \to \mathbb{F}_2^d$ be a labeling of the edges of the complete bipartite graph $K_{n,n}$ such that for any simple cycle C of length at most 10, $\sum_{c \in C} \gamma(c) \neq 0$. Then $d \geq 3\lceil \log n \rceil - O(1)$.

Proof: Let $n = 2m + 1$. Consider the following set of paths with length 5:

$$V = \{(1, i_1, j_1, i_2, j_2, 1) \mid i_1, j_2 \in \{2, \cdots, m+1\}, \\ i_2, j_1 \in \{m+2, \cdots, 2m+1\}\}.$$

We build a graph G on V as mentioned before Lemma 4.1: we connect by an edge the vertices corresponding to two paths p and p' if $p \oplus p'$ is a simple cycle. Notice that whenever two paths form a simple cycle, its length is at most 10. This graph can be divided into interconnecting cliques in the following way. For $i_1, j_2 \in \{2, \cdots, m+1\}$, let

$$V_{i_1,j_2} = \{(1, i_1, j_1, i_2, j_2, 1) \\ \mid j_1, i_2 \in \{m+2, \cdots, 2m+1\}\}$$

be the set of paths obtained from fixing the 2nd and 5th vertex to be i_1 and j_2, respectively. The induced subgraph of G on V_{i_1,j_2} is a clique, because the symmetric difference between any two paths in V_{i_1,j_2} is always a simple cycle of length 4 or 6, as mentioned in the proof of Theorem 4.4. Therefore, we divide G into m^2 disjoint cliques (indexed by i_1, j_2), each of size m^2.

Now we move on to explore how these cliques are connected to each other. Let $V_{i_1,j_2} \neq V_{i'_1,j'_2}$ be two different cliques, and

$$p = (1, i_1, j_1, i_2, j_2, 1) \in V_{i_1,j_2},$$
$$p' = (1, i'_1, j'_1, i'_2, j'_2, 1) \in V_{i'_1,j'_2}$$

be two paths. We break the analysis into 3 cases.

1) $i_1 \neq i'_1$ and $j_2 \neq j'_2$. In this case p and p' are connected if and only if $j_1 \neq j'_1$ and $i_2 \neq i'_2$.
2) $i_1 = i'_1$ and $j_2 \neq j'_2$. In this case p and p' are *not* connected if and only if $j_1 \neq j'_1$ and $i_2 = i'_2$.
3) $i_1 \neq i'_1$ and $j_2 = j'_2$. In this case p and p' are *not* connected if and only if $j_1 = j'_1$ and $i_2 \neq i'_2$.

Intuitively this is a well-connected graph, and in fact we have an upper bound on the independence number of G:

$$\alpha(G) < 2m - 1.$$

To prove the claim, assume there is an independent set I and $|I| \geq 2m - 1$. Since each V_{i_1,j_2} is a clique, we can pick at most one path out of each V_{i_1,j_2} to form I. By Lemma 4.5 there exists $i_0, j_0, i'_0 \neq i_0, j'_0 \neq j_0 \in \{2, \cdots, m+1\}$ such that I contains one path from each of the following 3 cliques: V_{i_0,j_0}, V_{i_0,j'_0} and $V_{i'_0,j_0}$. Suppose the involving 3 paths are

$$p_1 = (1, i_0, j_1, i_1, j_0, 1) \in I \cap V_{i_0,j_0},$$
$$p_2 = (1, i'_0, j_2, i_2, j_0, 1) \in I \cap V_{i'_0,j_0},$$
$$p_3 = (1, i_0, j_3, i_3, j'_0, 1) \in I \cap V_{i_0,j'_0}.$$

From the discussion above we know it has to be the case that

$$j_2 = j_1 \text{ and } i_2 \neq i_1, \quad j_3 \neq j_1 \text{ and } i_3 = i_1.$$

However, this implies $j_2 \neq j_3$ and $i_2 \neq i_3$. Combining $i_0 \neq i'_0$ and $j_0 \neq j'_0$ shows that p_2 and p_3 are indeed connected. So this contradicts with the fact that I is an independent set.

By Lemma 4.2 we have a thing

$$\chi(G) \geq \frac{|G|}{\alpha(G)} > \frac{m^4}{2m-1} \geq \frac{1}{2}m^3,$$

and by Lemma 4.1 we have

$$d \geq \lceil \log \chi(G) \rceil \geq \left\lceil \log\left(\frac{1}{2}m^3\right)\right\rceil = 3\lceil \log n \rceil - O(1).$$

■

V. Acknowledgments

We thank Schachar Lovett and Sergey Yekhanin for helpful suggestions.

References

[BB12] Simeon Ball and Jan De Beule. On sets of vectors of a finite vector space in which every subset of basis size is a basis II. *Des. Codes Cryptography*, 65(1-2):5–14, 2012.

[BE16] Simon Blackburn and Tuvi Etzion. Pir array codes with optimal pir rate. *arXiv preprint arXiv:1607.00235*, 2016.

[BHH13] Mario Blaum, James Lee Hafner, and Steven Hetzler. Partial-mds codes and their application to RAID type of architectures. *IEEE Trans. Information Theory*, 59(7):4510–4519, 2013.

[BK15] S. B. Balaji and P. Vijay Kumar. On partial maximally-recoverable and maximally-recoverable codes. In *ISIT*, pages 1881–1885. IEEE, 2015.

[CHL07] Minghua Chen, Cheng Huang, and Jin Li. On the maximally recoverable property for multi-protection group codes. In *Information Theory, 2007. ISIT 2007. IEEE International Symposium on*, pages 486–490. IEEE, 2007.

[DGRS14] Alexandros G Dimakis, Anna Gál, Ankit Singh Rawat, and Zhao Song. Batch codes through dense graphs without short cycles. *arXiv preprint arXiv:1410.2920*, 2014.

[DRWS11] Alexandros G Dimakis, Kannan Ramchandran, Yunnan Wu, and Changho Suh. A survey on network codes for distributed storage. *Proceedings of the IEEE*, 99(3):476–489, 2011.

[FVY15] Arman Fazeli, Alexander Vardy, and Eitan Yaakobi. Codes for distributed pir with low storage overhead. In *Information Theory (ISIT), 2015 IEEE International Symposium on*, pages 2852–2856. IEEE, 2015.

[GHJY14] Parikshit Gopalan, Cheng Huang, Bob Jenkins, and Sergey Yekhanin. Explicit maximally recoverable codes with locality. *IEEE Trans. Information Theory*, 60(9):5245–5256, 2014.

[GHK+17] 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, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*, pages 2092–2108, 2017.

[GHSY12] Parikshit Gopalan, Cheng Huang, Huseyin Simitci, and Sergey Yekhanin. On the locality of codeword symbols. *IEEE Trans. Information Theory*, 58(11):6925–6934, 2012.

[HCL13] Cheng Huang, Minghua Chen, and Jin Li. Pyramid codes: Flexible schemes to trade space for access efficiency in reliable data storage systems. *TOS*, 9(1):3:1–3:28, 2013.

[HSX+12] Cheng Huang, Huseyin Simitci, Yikang Xu, Aaron Ogus, Brad Calder, Parikshit Gopalan, Jin Li, and Sergey Yekhanin. Erasure coding in windows azure storage. In *USENIX Annual Technical Conference*, pages 15–26. USENIX Association, 2012.

[HY16] Guangda Hu and Sergey Yekhanin. New constructions of SD and MR codes over small finite fields. In *ISIT*, pages 1591–1595. IEEE, 2016.

[IKOS04] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Batch codes and their applications. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pages 262–271. ACM, 2004.

[KLR17] Daniel Kane, Shachar Lovett, and Sankeerth Rao. The independence number of the birkhoff polytope graph, and applications to maximally recoverable codes. In *FOCS*, page (to appear), 2017.

[KPLK14] Govinda M Kamath, N Prakash, V Lalitha, and P Vijay Kumar. Codes with local regeneration and erasure correction. *IEEE Transactions on Information Theory*, 60(8):4637–4660, 2014.

[LL15] V. Lalitha and Satyanarayana V. Lokam. Weight enumerators and higher support weights of maximally recoverable codes. In *Allerton*, pages 835–842. IEEE, 2015.

[LN86] Rudolf Lidl and Harald Niederreiter. *Finite Fields*. Cambridge University Press, 1986.

[PGM13] James S. Plank, Kevin M. Greenan, and Ethan L. Miller. Screaming fast galois field arithmetic using intel SIMD instructions. In *FAST*, pages 299–306. USENIX, 2013.

[RPDV16] Ankit Singh Rawat, Dimitris S Papailiopoulos, Alexandros G Dimakis, and Sriram Vishwanath. Locality and availability in distributed storage. *IEEE Transactions on Information Theory*, 62(8):4481–4493, 2016.

[SAP+13] Maheswaran Sathiamoorthy, Megasthenis Asteris, Dimitris S. Papailiopoulos, Alexandros G. Dimakis, Ramkumar Vadali, Scott Chen, and Dhruba Borthakur. Xoring elephants: Novel erasure codes for big data. *PVLDB*, 6(5):325–336, 2013.

[SLR+14] Muralidhar Subramanian, Wyatt Lloyd, Sabyasachi Roy, Cory Hill, Ernest Lin, Weiwen Liu, Satadru Pan, Shiva Shankar, Sivakumar Viswanathan, Linpeng Tang, and Sanjeev Kumar. f4: Facebook's warm BLOB storage system. In *OSDI*, pages 383–398. USENIX Association, 2014.

[TB14a] Itzhak Tamo and Alexander Barg. Bounds on locally recoverable codes with multiple recovering sets. In *Information Theory (ISIT), 2014 IEEE International Symposium on*, pages 691–695. IEEE, 2014.

[TB14b] Itzhak Tamo and Alexander Barg. A family of optimal locally recoverable codes. *IEEE Trans. Information Theory*, 60(8):4661–4676, 2014.

[TBF16] Itzhak Tamo, Alexander Barg, and Alexey Frolov. Bounds on the parameters of locally recoverable codes. *IEEE Transactions on Information Theory*, 62(6):3070–3083, 2016.

[WZ14] Anyu Wang and Zhifang Zhang. Repair locality with multiple erasure tolerance. *IEEE Transactions on Information Theory*, 60(11):6979–6987, 2014.

[Yek12] Sergey Yekhanin. Locally decodable codes. *Foundations and Trends in Theoretical Computer Science*, 6(3):139–255, 2012.

[Yek17] Sergey Yekhanin. personal communication, 2017.