Explicit Maximally Recoverable Codes with Locality*

Parikshit Gopalan
Microsoft Research
parik@microsoft.com

Cheng Huang
Microsoft Research
chengh@microsoft.com

Bob Jenkins
Microsoft Corporation
bob.jenkins@microsoft.com

Sergey Yekhanin
Microsoft Research
yekhanin@microsoft.com

Abstract

Consider a systematic linear code where some (local) parity symbols depend on few prescribed symbols, while other (heavy) parity symbols may depend on all data symbols. Such codes have been studied recently in the context of erasure coding for data storage, where the local parities facilitate fast recovery of any single symbol when it is erased, while the heavy parities provide tolerance to a large number of simultaneous erasures.

A code as above is maximally recoverable, if it corrects all erasure patterns which are information theoretically correctable given the prescribed dependency relations between data symbols and parity symbols. In this paper we present explicit families of maximally recoverable codes with locality. We also initiate the general study of the trade-off between maximal recoverability and alphabet size.

*This paper was presented in part at the Joint Mathematics Meeting of the American Mathematical Society, 2013, San-Diego.

1 Introduction

We say that a certain coordinate of an error-correcting code has locality r if, when erased, the value at this coordinate can be recovered by accessing at most r other coordinates. Recently there have been two lines of work on codes with locality.

In [GHSY12], motivated by applications to distributed data storage [HSX+12], the authors studied systematic linear $[n, k]$ codes that tolerate up to $h + 1$ erasures, but also have locality r for all information coordinates. In the *canonical codes* of [GHSY12], the locality r divides k and the length is $n = k + k/r + h$. Data (information) symbols are partitioned into k/r groups of size r. For each data group there is a local parity storing the XOR of respective data symbols. In addition, there are h heavy parities, each of which could depend on all k data symbols. In what follows we refer to codes above as *data-local* (k, r, h)-codes.

In [BHH13] motivated by applications to data storage on SSDs the authors studied systematic linear $[n, k]$ codes with two extra parameters r and h. In codes of [BHH13] there are k data symbols and h heavy parity symbols. In the case when $r \mid (k + h)$ these $(k + h)$ symbols are partitioned into $(k + h)/r$ groups of size r. For each group there is a local parity storing the XOR of respective symbols. Thus $n = k+h+(k+h)/r$. Unlike the codes of [GHSY12], codes of [BHH13] provide locality for all symbols data or parity. In what follows we refer to codes above as *local* (k, r, h)-codes.

Our descriptions of code families above are not complete. For every parity symbol we specified other symbols that it depends on, in other words, we have fixed the topology of the codes. To completely define the codes we need to specify coefficients in the heavy parity symbols. Different choices of coefficients could lead to codes with different erasure correcting capabilities. The best we could hope for is to have an optimal choice of coefficients which ensures that our codes can correct all patterns of erasures that are correctable for some setting of coefficients in heavy parities. Such codes do exist and are called Maximally Recoverable (MR) [CHL07]. The existence of those codes is proved by choosing coefficients randomly. Not only does this not yield explicit constructions, it only shows the existence of these codes over fairly large finite fields. Having small finite fields is an important consideration in practice, where coefficient sizes of a few bytes are preferred (see [PGM13, Section 2] for a detailed discussion).

An important problem left open by earlier work has been to come up with explicit maximally recoverable data-local and local codes over small finite fields.

There are several other models of codes with locality in the literature. The ones most closely related to our work include SD codes [Bla13, PBH13], locally decodable codes [Yek12], and regenerating codes [DGW+10].

1.1 Explicit maximally recoverable codes with locality

In this paper we present the first explicit families of maximally recoverable data-local and local codes for all values of k, r and h. Prior to our work infinite explicit families of maximally recoverable local codes were known only for $h = 1$ and $h = 2$ [Bla13, BHH13]. There have also been few constructions that involved computer search for coefficients [BHH13]. It was known that one can construct maximally recoverable local codes by picking coefficients in heavy parities at random from a large enough finite field [CHL07]. But for this randomized construction to succeed

with any constant probability, the coefficients must be drawn from a finite field of size $\Omega\left(k^{h-1}\right)$ (see Theorem 27).

Our codes improve upon the earlier results (explicit and existential) both in concrete settings and asymptotically. To keep the statements simple, we focus here on the asymptotics of our constructions when $h = O(1), r = O(1)$, and k goes to infinity. We present an explicit construction of local (k, r, h)-codes over an alphabet of size

$$q = O\left(k^{h-1}\right)$$

which is at least as good as the randomized construction. For $h \geqslant 2^r + 1$, the alphabet size can be reduced to

$$q = O\left(k^{\left\lceil (h-1)\left(1-\frac{1}{2^r}\right)\right\rceil}\right)$$

which beats the randomized construction. We also obtain further improvements in the special cases of $h = 3$ and $h = 4$.

The only lower bound for the alphabet size known currently comes from results on the main conjecture for MDS codes [MS77] and is $\Omega(k)$. Closing the gap between this lower bound and our upper bounds is an intriguing open question.

As in the work of [BHH13, Bla13] we construct our explicit codes via parity check matrices. As in [BHH13] columns of our parity check matrices have the shape $(\alpha_j, \alpha_j^2, \ldots, \alpha_j^{2^{h-1}})$. The key difference from the work of [BHH13, Bla13] however is that we explicitly specify the α_js. A common theme in our constructions is to reduce the problem of explicitly constructing these sets to the classical coding theory problem of constructing high-rate linear codes with a prescribed minimum distance d. We present two constructions which we call the Basic construction and the Product construction. The Basic construction derives the α_js from high rate binary linear codes and is simpler to describe. The Product construction uses linear codes over larger fields. The alphabet size bounds quoted above are obtained via the Product construction.

Observe that our upper bound of $O(k^{h-1})$ meets the lower bound of $\Omega(k)$ when $h = 2$. Constructions like that have been previously known [Bla13]. Also, let us remark that when $h = 1$ one can have (k, r, h)-local MR codes over a field of size $r + 1$ [BHH13].

1.2 Maximally recoverability for other topologies

The original definition of maximal recoverability [CHL07] does not accommodate arbitrary topologies, for instance, it does not allow for locality within the heavy parities. In this paper we present a generalized definition. Roughly speaking, we start from parity check matrices P (representing topologies) whose entries are $\mathbb{F}_2$-affine functions in variables $z_1, \ldots, z_m$. An assignment to these variables from a field $\mathbb{F}_q$ of characteristic 2 instantiates P and yields a parity check matrix and a code over $\mathbb{F}_q$. An instantiation of a topology is maximally recoverable if the resulting code corrects every erasure pattern that is correctable by some other instantiation of the same topology.

Having generalized the definition of maximal recoverability, we extend our Basic construction of local (k, r, h)-codes to obtain explicit MR codes for a broad class of topologies that in particular includes local codes and data-local codes.

1.3 Organization

In Section 2 we formally define data-local and local (k, r, h)-codes. We introduce the notion of maximal recoverability (restricted to the setting of interest), and show that maximally recoverable local codes yield maximally recoverable data-local codes. In Section 3 we give our two main code constructions: the Basic construction and the Product construction. In Section 4 we establish a simple lower bound on the alphabet size of maximally recoverable local codes. We also present a lower bound due to Kopparty and Meka [KM] on the alphabet size needed for random codes to be maximally recoverable. In Section 5 we give a general definition of maximal recoverability. We extend our Basic construction of local (k, r, h)-codes to obtain explicit MR codes for a broad class of topologies. In Section 6 we conclude with open questions.

2 Preliminaries

We use the following notation

- For an integer n, $[n] = \{1, \ldots, n\}$;
- An $[n, k]$ code is a linear code encoding k-dimensional messages to n-dimensional codewords. Equivalently, one can think of an $[n, k]$ code as a k-dimensional subspace of an n-dimensional space over a finite field;
- An $[n, k, d]$ code is an $[n, k]$ code whose minimal distance is at least d;
- Let C be an $[n, k]$ code and $S \subseteq [n]$. Puncturing C in coordinates in S means restricting C to coordinates in $[n] \setminus S$. It yields a $[k', n - |S|]$ code C', where $k' \leqslant k$.

We proceed to formally define the notion of locality.

Definition 1. *[GHSY12] Let C be a linear $[n, k]$ code. We say that the i-th coordinate of C has locality r, if there exists a set $S \subseteq [n] \setminus \{i\}$, $|S| \leqslant r$, such that across all codewords $\mathbf{c} \in C$, the value of the coordinate $\mathbf{c}(i)$ is determined by values of coordinates $\{\mathbf{c}(j)\}, j \in S$. Equivalently, the i-th coordinate has locality r, if the dual code $C^{\perp}$ contains a codeword $\mathbf{c}$ of Hamming weight at most $r + 1$, where coordinate i is in the support of $\mathbf{c}$.*

Definition 2. *Let C be a linear systematic $[n, k]$ code. We say that C is a (k, r, h) data-local code if the following conditions are satisfied:*

- *$r \mid k$ and $n = k + k/r + h$;*
- *Data symbols are partitioned into k/r groups of size r. For each such group there is one (local) parity symbol that stores the XOR of respective data symbols;*
- *The remaining h (heavy) parity symbols may depend on all k data symbols.*

In what follows we refer to a group of r data symbols and their local parity defined above as a *local group*. Data-local codes have been studied in [HCL07, GHSY12, PD12, PKLK12, SAP+13, FY13]. The importance of this topology was partially explained in [GHSY12, Theorem 9]. There it has been shown that in case $h < r + 1$ and $r \mid k$, any systematic $[n, k]$ code that corrects all patterns of $(h + 1)$ erasures, provides locality r for all data symbols, and has the lowest possible redundancy has to be a data-local (k, r, h)-code. The class of data local-codes is fairly broad as there is a lot of flexibility is choosing coefficients in heavy parities. Below we define data-local codes that maximize reliability.

Definition 3. *Let C be a data-local (k, r, h)-code. We say that C is maximally recoverable if for any set $E \subseteq [n]$, where E is obtained by picking one coordinate from each of k/r local groups, puncturing C in coordinates in E yields a maximum distance separable $[k + h, k]$ code.*

A $[k+h, k]$ MDS code obviously corrects all patterns of h erasures. Therefore a code satisfying Definition 3 corrects all erasure patterns $E \subseteq [n]$ that involve erasing one coordinate per local group, and h additional coordinates. We now argue that any erasure pattern that is not dominated by a pattern above has to be uncorrectable. Thus the codes defined above correct every possible erasure pattern given their topology, which justifies calling them Maximally Recoverable codes.

Lemma 4. *Let C be an arbitrary data-local (k, r, h)-code. Let $E \subseteq [n]$ be an erasure pattern. Suppose E affects t local groups and $|E| > t + h$; then E is not correctable.*

Proof. Suppose E is correctable. We extend E to a larger pattern of erasures E' erasing one arbitrary coordinate in each of $k/r - t$ local groups that are not affected by E. Observe that E' is correctable if E is correctable since each local group has a local parity. Note that the size of E' exceeds redundancy of the code C, $|E'| > k/r + h$. Thus the dimension of C restricted to coordinates outside of E' is below k, and there are codewords in C with identical projections on $[n] \setminus E'$. Therefore E' is not correctable. □

We now proceed to define local codes.

Definition 5. *Let C be a linear systematic $[n, k]$ code. We say that C is a (k, r, h) local code if the following conditions are satisfied:*

- *$r \mid (k + h)$ and $n = k + h + (k + h)/r$;*
- *There are k data symbols and h heavy parity symbols, where each heavy parity may depend on all data symbols;*
- *These $k + h$ symbols are partitioned into $(k + h)/r$ groups of size r. For each such group there is one (local) parity symbol that stores the XOR of respective symbols.*

We refer to a group of r symbols and their local parity as a *local group*. As above we now introduce local codes that maximize reliability.

Definition 6. *Let C be a local (k, r, h)-code. We say that C is maximally recoverable if for any set $E \subseteq [n]$, where E is obtained by picking one coordinate from each of $(k + h)/r$ local groups, puncturing C in coordinates in E yields a maximum distance separable $[k + h, k]$ code.*

One can use similar arguments as before to show that these codes correct all erasure patterns that involve erasing one coordinate per local group, and h additional coordinates and further that any erasure patterns that are not dominated by such patterns are not correctable by any local (k, r, h)-code. Maximally Recoverable local (k, r, h)-codes have been originally introduced in [BHH13] under the name of partial-MDS codes.

2.1 Data-local codes from local codes

The next lemma shows how one can derive constructions of data-local MR codes from constructions of local MR codes.

Lemma 7. *Suppose there exists a local maximally recoverable (k, r, h)-code C over a finite field $\mathbb{F}$; then there exists a data-local maximally recoverable (k', r, h)-code C' over the same field, where $k' \leqslant k$ is the largest integer that is divisible by r.*

Proof. Let $t = (k + h)/r$. Let $G_1, \ldots, G_t \subseteq [n]$ be the local groups. $\cup_i G_i = [n]$. We refer to data symbols and heavy parity symbols of C as *primary* symbols. Altogether primary symbols form a $[k + h, k]$ MDS code. Note that any k symbols of an MDS code can be treated as information symbols. Next we consider two cases:

- $r \mid k$. We treat k primary symbols of C that belong to local groups $\{G_i\}, i \leqslant k/r$ as data symbols of C'. The code C' is obtained form the code C by dropping local parity symbols from groups G_i for $i > k/r$. The code C' clearly satisfies definition 2. Observe that C' also satisfies definition 3 as any code that can be obtained by dropping one coordinate per local group in C' can also be obtained by dropping one coordinate per local group in C.

- $r \nmid k$. Let $s = \lfloor k/r \rfloor$. We refer to local groups $\{G_i\}, i \leqslant s$ as data groups. We refer to group G_{s+1} as special. We treat k' primary symbols of C that belong to data groups $\{G_i\}, i \leqslant s$ as data symbols of C'. We fix some arbitrary $k - k'$ primary symbols in the special group, and refer to them as special symbols. We denote the collection of special symbols by S.

  The code C' is obtained from the code C by dropping all special symbols and $t - s$ local parities in groups other than data groups. Given an assignment of values to k' data symbols of C', we determine the values of heavy parities using the code C assuming that all special symbols are set to zero.

  The code C' clearly satisfies definition 2. It also satisfies definition 3 as any codeword that can be obtained by dropping one coordinate per local group in $C'(x')$ can also be obtained by dropping one coordinate per local group in $C(x' \circ 0^{k-k'})$ restricted to $[n] \setminus S$. The latter restriction does not affect the erasure correcting capability of the code as we are dropping coordinates that are identically zero.

This concludes the proof. □

2.2 On the asymptotic analysis of our constructions

Unlike data transmission applications, in data storage applications one typically does not scale the number of heavy parities linearly with the number of data fragments k to ensure the same level

of reliability [HSX+12], much slower growth in the number of parities suffices. The underlying reason for this is that the values of k which are relevant here are fairly small (of the order of tens, rarely above one hundred) and the likelihood p of a fragment failure during a certain window of time is usually substantially smaller than $1/k$. Note that since p is an absolute constant dependent only on the device characteristics, this also implies k is bounded above by a constant. In appendix A, we attempt to make rigorous some of such assumptions under which it suffices for h to be a very slowly growing function of k to achieve a desired level of reliability.

Thus (with the caveat that explicit bounds are the best for practice) we feel the asymptotic setting of fixed r and h and growing k to be a relevant and natural starting point theoretically. Therefore, for all our constructions we first bound the field-size as an explicit function of k, r and h, and then analyze the asymptotic behavior in the regime where r, h are constants and k goes to infinity.

3 Code constructions

In this section we give our two main constructions of local codes. We restrict our attention to finite fields of characteristic two. Let $\mathbb{F}$ be such a field. Let $S = \{\alpha_1, \ldots, \alpha_n\} \subseteq \mathbb{F}$ be a multi-set of n elements. Let $A(S, h) = [a_{gj}]$ denote the $h \times n$ matrix where

$$a_{gj} = \alpha_j^{2^{g-1}}$$

Let $\mathcal{C}(S, h) \subset \mathbb{F}^n$ be the linear code whose parity check matrix is A. Equivalently, $\mathcal{C}(S, h)$ contains all vectors $\mathbf{x} = (x_1, \ldots, x_n)$ which satisfy the equations

$$\sum_{j=1}^{n} \alpha_j^{2^{g-1}} x_j = 0 \quad \text{for} \quad g = 1, \ldots, h. \tag{1}$$

Let $\mathcal{C}(\alpha, h)$ be an $[n, k, d]$ code. It is easy to see that $k \geqslant n - h$, hence by the Singleton bound, $d \leqslant h + 1$. We are interested in sets $\{\alpha_j\}$ where $d = h + 1$, so that the code $\mathcal{C}(S, h)$ is maximum distance separable. The following lemma characterizes such sets.

Definition 8. *We say that the multi-set $S \subseteq \mathbb{F}$ is t-wise independent over a field $\mathbb{F}' \subseteq \mathbb{F}$ if every $T \subseteq S$ such that $|T| \leqslant t$ is linearly independent over $\mathbb{F}'$.*

The following lemma is standard.

Lemma 9. *The code $\mathcal{C}(S, h)$ has distance $h+1$ if and only if the multi-set S is h-wise independent over the field $\mathbb{F}_2$.*

Proof. Let $\mathbf{x} = (x_1, \ldots, x_n) \in \mathcal{C}(S, h)$ be a codeword. The code $\mathcal{C}(S, h)$ has distance $h+1$ iff every pattern of h erasures is correctable. In other words, for any $E \subseteq [n]$, the values $\{x_j\}_{j \in E}$ can be recovered if we know the values of all $\{x_j\}_{j \in [n] \setminus E}$. This requires solving the following system of equations:

$$\sum_{j \in E} \alpha_j^{2^{g-1}} x_j = b_g, \quad 1 \leqslant g \leqslant h \tag{2}$$

which in turn requires inverting the $h \times h$ matrix A_E which is the the minor of A obtained by taking the columns in E. It is easy to see (e.g., [LN83, Lemma 3.51]) that A_E has a non-zero determinant if and only if the multi-set $\{\alpha_j\}_{j \in E}$ is linearly independent over $\mathbb{F}_2$. □

Lemma 9 describes the effect of adding parity check constraints to n otherwise independent variables. We now consider the effect of adding such constraints to symbols that already satisfy some dependencies. We work with the following setup. The n coordinates of the code are partitioned into

$$\ell = \frac{n}{r+1} = \frac{k+h}{r}$$

local groups, with group i containing $r+1$ symbols $x_{i,1}, \ldots, x_{i,r+1}$. Variables in each local group satisfy a parity check constraint $\sum_{s=1}^{r+1} x_{i,s} = 0$. Thus all code coordinates have locality r. Let

$$S = \{\alpha_{i,s}\}_{i\in[\ell],s\in[r+1]} \in \mathbb{F}^n.$$

We define the code $\mathcal{C}(S,r,h)$ by the parity check equations

$$\sum_{i=1}^{\ell}\sum_{s=1}^{r+1} \alpha_{i,s}^{2^{g-1}} x_{i,s} = 0 \quad \text{for } g \in \{1,\ldots,h\}, \tag{3}$$

$$\sum_{s=1}^{r+1} x_{i,s} = 0 \quad \text{for } i \in \{1,\ldots,\ell\} \tag{4}$$

We refer to Equations (3) as global constraints and (4) as local constraints.

Let $\mathbf{e} \in [r+1]^\ell$ be a vector. Let $\mathcal{C}^{-\mathbf{e}} = \mathcal{C}^{-\mathbf{e}}(S,r,h)$ be the code obtained by puncturing $\mathcal{C}(S,r,h)$ in positions $\{i,\mathbf{e}(i)\}_{i=1}^{\ell}$. In this notation, Definition 6 says that $\mathcal{C}(S,r,h)$ is a maximally recoverable (k,r,h)-code if $\mathcal{C}^{-\mathbf{e}}$ is an MDS code for every $\mathbf{e} \in [r+1]^\ell$. The following proposition is central to our method:

Proposition 10. *The code $\mathcal{C}(S,r,h)$ is a maximally recoverable (k,r,h)-code iff for every $\mathbf{e} \in [r+1]^\ell$,*

$$T(S,\mathbf{e}) = \{\alpha_{i,s} + \alpha_{i,\mathbf{e}(i)}\}_{i\in[\ell],s\in[r+1]\setminus\{\mathbf{e}(i)\}}$$

is h-wise independent.

Proof. Note that $\mathcal{C}^{-\mathbf{e}}$ is a $[k+h,k]$ code. To prove that it is MDS, we will use the local parity constraints to eliminate the punctured locations and then use Lemma 9. Firstly, by renumbering variables (and coefficients $\{\alpha_{i,s}\}$) in each local group, we may assume $\mathbf{e}(i) = r+1$. By the local parity check equations,

$$x_{i,r+1} = \sum_{s=1}^{r} x_{i,s}.$$

We use these to eliminate $x_{i,r+1}$ from the global parity check equations for $g \in [h]$:

$$\begin{aligned}
0 &= \sum_{i=1}^{\ell} \left(\sum_{s=1}^{r+1} \alpha_{i,s}^{2^{g-1}} x_{i,s} \right) \\
&= \sum_{i=1}^{\ell} \left(\left(\sum_{s=1}^{r} \alpha_{i,s}^{2^{g-1}} x_{i,s} \right) + \alpha_{i,r+1}^{2^{g-1}} \left(\sum_{s=1}^{r} x_{i,s} \right) \right) \\
&= \sum_{i=1}^{\ell} \left(\sum_{s=1}^{r} (\alpha_{i,s}^{2^{g-1}} + \alpha_{i,r+1}^{2^{g-1}}) x_{i,s} \right) \\
&= \sum_{i=1}^{\ell} \left(\sum_{s=1}^{r} (\alpha_{i,s} + \alpha_{i,r+1})^{2^{g-1}} x_{i,s} \right).
\end{aligned}$$

Let $T = \{\alpha_{i,s} + \alpha_{i,r+1}\}_{i \in [\ell], s \in [r]}$. By Lemma 9, the code $\mathcal{C}^{-\mathbf{e}}$ is MDS if and only if T is h-wise independent. □

Proposition 10 reduces constructing local MR codes to obtaining multi-sets $S \subseteq \mathbb{F}$ such that all sets $T(S, \mathbf{e})$ are h-wise independent. In what follows we give two constructions of such multi-sets.

3.1 The Basic construction

Lemma 11. *Let $S \subseteq \mathbb{F}, |S| = n$ be a set that is $2h$-wise independent over a subfield $\mathbb{F}'$. Let r be such that $\ell = n/(r+1)$ is an integer. Then for all $\mathbf{e} \in [r+1]^{\ell}$ the set $T(S, \mathbf{e})$ is h-wise independent over $\mathbb{F}'$.*

Proof. Assume the contrary. To simplify notation, we relabel variables and assume that $\mathbf{e}(i) = r + 1$ for every $i \in [\ell]$. Let $D = \{i_j, s_j\}_{j=1}^{d}$ be a set of $d \leqslant h$ indices of T such that

$$\sum_{j=1}^{d} \left(\alpha_{i_j, s_j} + \alpha_{i_j, r+1} \right) = 0$$

We can rewrite this as

$$\sum_{j=1}^{d} \alpha_{i_j, s_j} + \sum_{j=1}^{d} \alpha_{i_j, r+1} = 0$$

We claim that this gives a non-trivial relation between the coefficients $\{\alpha_{i,s}\}$. The relation is non-trivial because the terms in the first summation occur exactly once (whereas terms in the second summation can occur multiple times depending on the set D and could cancel). □

Observe that the task of constructing n-sized subsets of $\mathbb{F}_{2^t}$ that are $2h$-wise independent over $\mathbb{F}_2$ is equivalent to the task of constructing $[n, n-t, 2h+1]$ binary linear codes, as elements of a $2h$-wise independent set can be used as columns of a $t \times n$ parity check matrix of such a code, and vice versa. Therefore any family of binary linear codes can be used to obtain maximally recoverable local codes via Lemma 11 and Proposition 10. The next theorem gives local MR codes that one gets by instantiating the approach above with columns of the parity check matrix of a binary BCH code.

Theorem 12. *Let positive integers k, r, h be such that $r \mid (k+h)$. Let m be the smallest integer such that*

$$n = k + h + \frac{k+h}{r} \leqslant 2^m - 1.$$

There exists a maximally recoverable local (k, r, h)-code over the field $\mathbb{F}_{2^{hm}}$.

Proof. Let $S' = \{\beta_1, \ldots, \beta_n\}$ be an arbitrary subset of non-zero elements of $\mathbb{F}_{2^m}$. Consider $S = \{\alpha_1, \ldots, \alpha_n\} \subseteq \mathbb{F}_{2^{mh}}$ where for all $i \in [n]$, $\alpha_i = (\beta_i, \beta_i^3, \ldots, \beta_i^{2h-1})$ when we treat $\mathbb{F}_{2^{mh}}$ as an h-dimensional linear space over $\mathbb{F}_{2^m}$. It is not hard to see that the set S is $2h$-wise independent over $\mathbb{F}_2$. Thus by Lemma 11 and Proposition 10 the code $\mathcal{C}(S, r, h)$ is a maximally recoverable local (k, r, h)-code. □

Corollary 13. *For constants r and h and for all k such that $r|k+h$, there exists a maximally recoverable local (k, r, h)-code over a field of size $O(k^h)$.*

3.2 The Product construction

The Basic construction uses $2h$-wise independence of the set S to ensure h-independence of sets $T(S, \mathbf{e})$. The Product construction shows that one can ensure h-independence of sets $T(S, \mathbf{e})$ more economically.

Definition 14. *We say that the set $S \subseteq \mathbb{F}$ is t-wise weakly independent over $\mathbb{F}_2 \subseteq \mathbb{F}$ if no set $T \subseteq S$ where $2 \leqslant |T| \leqslant t$ has the sum of its elements equal to zero.*

Unlike independent sets, weakly independent sets may include the zero element.

Recall that our goal is specify the multi-set $S = \{\alpha_{i,s}\}$ for $i \in [\ell], s \in [r+1]$ where $\ell = n/(r+1)$. The code $C(S, r, h)$ is then specified by Equations (3) and (4). We now define the Product construction.

- Let $a|b$ so that $\mathbb{F}_2 \subseteq \mathbb{F}_{2^a} \subseteq \mathbb{F}_{2^b}$.
- Let $S_1 = \{\xi_1, \ldots, \xi_{r+1}\} \subseteq \mathbb{F}_{2^a}$ be h-weakly independent over $\mathbb{F}_2$ if h is even, and $(h+1)$-weakly independent over $\mathbb{F}_2$ if h is odd.
- Let $S_2 = \{\lambda_1, \ldots, \lambda_\ell\} \subseteq \mathbb{F}_{2^b}$ be h-independent over $\mathbb{F}_{2^a}$.
- Let $S = S_1 \times S_2 \subseteq \mathbb{F}_{2^b}$, where for $i \in [\ell]$ and $s \in [r+1]$, $\alpha_{i,s} = \lambda_i \xi_s$.

Proposition 15. *Let $S = S_1 \times S_2$ be as described above. Then for all $\mathbf{e} \in [r+1]^\ell$ the set $T(S, \mathbf{e})$ is h-wise independent over $\mathbb{F}_2$.*

Proof. Assume there exists $\mathbf{e}$ so that $T(S, \mathbf{e})$ is not h-wise independent. To simplify the notation we relabel variables and assume that $\mathbf{e}(i) = r+1$ for every $i \in [\ell]$. Let $D = \{i_j, s_j\}_{j=1}^d$ be a set of $d \leqslant h$ indices of T such that

$$\sum_{j=1}^{d} \left(\alpha_{i_j, s_j} + \alpha_{i_j, r+1}\right) = 0$$

We can rewrite this as

$$\sum_{t\in[\ell]} \lambda_t \cdot \sum_{j\,:\,i_j=t} \left(\xi_{t,s_j} + \xi_{t,r+1}\right) = 0.$$

Observe that after cancelations each non-empty inner sum above involves at least 2 terms. When h is even it involves at most h terms; when h is odd it involves at most $h+1$ terms. Therefore each inner sum is non-zero by the properties of the set S_1. Also note that the outer sum involves at most h terms λ_t with non-zero coefficients from $\mathbb{F}_{2^a}$ and thus is also non-zero by the properties of the set S_2. □

We now instantiate the construction with a certain particular choice of independent sets. Our sets come from columns of a parity check matrix of an extended BCH code.

Theorem 16. *Let positive integers k, r, h be such that $\ell = (k+h)/r$ is an integer. Let m be the smallest integer such that $r \mid m$ and $2^m \geqslant \ell$. There exists a maximally recoverable local (k, r, h)-code over the field $\mathbb{F}_{2^t}$ for*

$$t = r + m\left\lceil (h-1)\left(1 - \frac{1}{2^r}\right)\right\rceil. \tag{5}$$

Proof. Our construction uses $a = r$ and $b = t$. Let $\{\xi_1, \ldots, \xi_r\}$ be an arbitrary basis of $\mathbb{F}_{2^r}$ over $\mathbb{F}_2$. We set $\xi_{r+1} = 0$ and $S_1 = \{\xi_1, \ldots, \xi_{r+1}\}$. Clearly, S_1 is $(h+1)$-weakly independent over $\mathbb{F}_2$ for all h.

Let $\{\beta_1, \ldots, \beta_\ell\}$ be an arbitrary subset of $\mathbb{F}_{2^m}$. Consider $S_2 = \{\lambda_1, \ldots, \lambda_\ell\} \subseteq \mathbb{F}_{2^t}$ where for all $i \in [\ell]$,

$$\lambda_i = (1, \beta_i, \beta_i^2, \ldots, \beta_i^{h-1}) \tag{6}$$

where we omit every non-zero power β_i^j where $2^r \mid j$. We treat $\mathbb{F}_{2^t}$ as a linear space over $\mathbb{F}_{2^r}$. The first coordinate in (6) is a single value in $\mathbb{F}_{2^r}$. There are

$$\left\lceil (h-1)\left(1 - \frac{1}{2^r}\right)\right\rceil$$

more coordinates, each of which is an (m/r)-dimensional vector over $\mathbb{F}_{2^r}$. As an extension of $\mathbb{F}_{2^r}$, $\mathbb{F}_{2^t}$ has dimension

$$s = 1 + \frac{m}{r}\left\lceil (h-1)\left(1 - \frac{1}{2^r}\right)\right\rceil.$$

Hence it has dimension

$$rs = r + m\left\lceil (h-1)\left(1 - \frac{1}{2^r}\right)\right\rceil$$

over $\mathbb{F}_2$.

We claim that the set S_2 is h-independent over $\mathbb{F}_{2^r}$. Assume the contrary. Then for some non-empty set $S \subseteq [\ell]$, $|S| \leqslant h$ for all $0 \leqslant j \leqslant h-1$ whenever $2^r \nmid j$ we have

$$\sum_{i\in S} \gamma_i \lambda_i^j = 0, \tag{7}$$

where we assume $0^0 = 1$ and all $\{\gamma_i\} \in \mathbb{F}_{2^r}$. By standard properties of Frobenius automorphisms, Equation (7) implies

$$\sum_{i \in S} \gamma_i \lambda_i^j = 0,$$

for all $0 \leqslant j \leqslant h-1$ which contradicts the proprieties of the Vandermonde determinant.

Hence Propositions 15 and 10 imply that the code $\mathcal{C}(S_1 \times S_2, r, h)$ is a maximally recoverable local (k, r, h)-code with the claimed parameters. □

Corollary 17. *For constants r and h and for all k such that $r \mid (k+h)$, there exists a maximally recoverable local (k, r, h)-code over a field of size $O\left(k^{\lceil (h-1)\left(1-\frac{1}{2^r}\right)\rceil}\right)$.*

Combining Corollary 17 with Lemma 7 we get similar asymptotic result for data-local codes.

Corollary 18. *For constants r and h and for all k such that $r \mid k$, there exists a maximally recoverable data-local (k, r, h)-code over a field of size $O\left(k^{\lceil (h-1)\left(1-\frac{1}{2^r}\right)\rceil}\right)$.*

Example 19. Instantiating Theorem 16 with $k = 60$, $r = h = 4$, we obtain a $[80, 60, 7]$ maximally recoverable $(60, 4, 4)$ local code over $\mathbb{F}_{2^{16}}$. Prior to our work [BHH13, Theorem 4.2] a code with such parameters was not known to exist over any field of size below 2^{80}.

3.3 Further improvements for $h = 3$ and $h = 4$

In the proof of Theorem 16 we set S_1 to be a basis of $\mathbb{F}_{2^r}$ augmented with a zero. After that we could use columns of a parity check matrix of any linear code of co-dimension s and distance $h+1$ over $\mathbb{F}_{2^r}$ to define the set $S_2 \subseteq \mathbb{F}_{2^{rs}}$ and obtain a MR local (k, r, h)-code over $\mathbb{F}_{2^{rs}}$. While we used columns of the parity check matrix of an extended BCH code, other choices sometimes yield local MR codes over smaller alphabets. In this section we implement this to get codes that improve upon the codes of Theorem 16 for $h = 3$ or 4 and large k. We replace BCH codes in the construction of Theorem 16 with known constructions of codes that improve on BCH codes for distance 4 and 5.

Our first construction relies on the following Theorem due to Dumer [Dum95] that also follows from the Hartman-Tzeng bound [MS77]. See also [YD04]. Due to the lack of a good reference pointer we include a new self-contained proof in Appendix B.

Theorem 20. *[Dum95] Let q be prime power and m be an even integer. There exists a q-ary linear code with parameters $[q^m, q^m - 3m/2 - 1, 4]$.*

Theorem 21. *Let positive integers $k, r, h = 3$ be such that $\ell = (k+h)/r$ is an integer. Let m be the smallest even integer such that $\ell \leqslant 2^{rm}$; then there exists a maximally recoverable local $(k, r, 3)$-code over the field $\mathbb{F}_{2^t}$ for $t = r(3m/2 + 1)$.*

Proof. Let $\{\xi_1, \ldots, \xi_r\}$ be an arbitrary basis of $\mathbb{F}_{2^r}$ over $\mathbb{F}_2$. We set $\xi_{r+1} = 0$ and $S_1 = \{\xi_1, \ldots, \xi_{r+1}\}$. Clearly, S_1 is $(h+1)$-weakly independent over $\mathbb{F}_2$ for all h. Let $S_2' \subseteq \mathbb{F}_{2^r}^{t/r}$ be an arbitrary collection of ℓ columns of the parity check matrix of the code from Theorem 20 where $q = 2^r$. S_2' naturally defines a set $S_2 \subseteq \mathbb{F}_{2^t}$ that is 3-independent over $\mathbb{F}_{2^r}$. □

We remark that using results in [EB99] one can get further small improvements upon the theorem above.

Corollary 22. *For any constant r and for all k such that $r \mid (k+3)$, there exists a maximally recoverable local $(k, r, 3)$-code over a field of size $O\left(k^{3/2}\right)$.*

This improves on the $O(k^2)$ bound implied by Corollary 17. Our second construction relies on the following

Theorem 23. *[Dum95, Theorem 6] Let q be an even prime power and $m \geqslant 2$ be an integer. There exists a q-ary linear code with parameters $[q^{m-1}, q^{m-1} - 2m - \lceil (m-1)/3 \rceil, 5]$.*

Theorem 24. *Let positive integers $k, r, h = 4$ be such that $\ell = (k+h)/r$ is an integer. Let m be the smallest integer such that $3 \mid (m-1)$ and $\ell \leqslant 2^{r(m-1)}$; then there exists a maximally recoverable local $(k, r, 4)$-code over the field $\mathbb{F}_{2^t}$ for $t = r(2m + (m-1)/3)$.*

Proof. As before let $\{\xi_1, \ldots, \xi_r\}$ be an arbitrary basis of $\mathbb{F}_{2^r}$ over $\mathbb{F}_2$. We set $\xi_{r+1} = 0$ and $S_1 = \{\xi_1, \ldots, \xi_{r+1}\}$. Clearly, S_1 is $(h+1)$-weakly independent over $\mathbb{F}_2$ for all h. Let $S_2' \subseteq \mathbb{F}_{2^r}^{t/r}$ be an arbitrary collection of ℓ columns of the parity check matrix of the code of Theorem 23, where we set $q = 2^r$. S_2' naturally defines a set $S_2 \subseteq \mathbb{F}_{2^t}$ that is 4-independent over $\mathbb{F}_{2^r}$. □

Corollary 25. *For any constant r and for all k such that $r \mid (k+4)$, there exists a maximally recoverable local $(k, r, 4)$-code over a field of size $O\left(k^{7/3}\right)$.*

This improves on the $O(k^3)$ bound implied by Corollary 17.

4 Lower bounds

Identifying the right alphabet size for the maximally recoverable local codes seems to be a challenging problem. It is possible that one could improve on the constructions in this work. Indeed, we do not know if the alphabet size needs to grow with h. The only lower bound we currently have comes from results on the main conjecture for MDS codes and is $\Omega(k)$.

Theorem 26. *Let $h \geqslant 2$. Let C be a maximally recoverable data-local (k, r, h)-code (or a maximally recoverable local (k, r, h)-code) defined over the finite field $\mathbb{F}_q$; then $q \geqslant k+1$.*

Proof. Consider the code C' that is obtained from C by deleting all local parities. Clearly, C' is a $[k+h, k, h+1]$ MDS code. Consider the $h \times (k+h)$ parity check matrix of the code C' with entries in $\mathbb{F}_q$. By [Bal12, Lemma 1.2], $k + h \leqslant q + h - 1$. □

Details regarding the recent progress on the main conjecture for MDS codes can be found in [Bal12, BB12]. In particular, results there allow one to get small non-asymptotic improvements upon Theorem 26.

4.1 Lower bounds for Random codes

One way to construct maximally recoverable local codes is by picking coefficients in heavy parities at random from a large enough finite field. In order to compare our constructions in Section 3 with random local codes, we show that random codes are not maximally recoverable (except with probability $o(1)$) unless the size of the finite field from which the coefficients are drawn exceeds $\Omega\left(k^{h-1}\right)$. This follows as a corollary of a result saying that a random $[k+h,k]_q$ code is unlikely to be an MDS code unless $q = \Omega\left(k^{h-1}\right)$. The following theorem is due to Swastik Kopparty and Raghu Meka.

Theorem 27. *[KM] Consider a random $[k+h,k]$ linear code C over a finite field $\mathbb{F}_q$ where*

$$q \leqslant \binom{\lfloor \frac{k}{2} \rfloor}{h-1}.$$

Then the probability that C has distance $h+1$ is at most

$$\left(1 - \frac{1}{2^h e^{h-1}}\right)^{\frac{k}{2}}.$$

Proof. We can assume without loss of generality that C is systematic and it comprises of k information symbols and h parity checks which are chosen to be independent random linear combinations of the information symbols. Let M be the $h \times (k+h)$ parity check matrix of C. The columns corresponding to parity checks give an $h \times h$ identity matrix while the other k columns $\mathbf{m}_1, \ldots, \mathbf{m}_k$ are drawn from $\mathbb{F}_q^h$ uniformly at random.

For $S \subseteq [k]$, let us denote the span of vectors $\{\mathbf{m}_i\}_{i \in S}$ by $\mathcal{L}(S)$. The code C is MDS only if every h of the $\mathbf{m}_i$s are linearly independent. In particular, this requires that for all $t \leqslant k$,

1. Any $h-1$ vectors in $\{\mathbf{m}_1, \ldots, \mathbf{m}_t\}$ are linearly independent.
2. For all $S \subseteq [t], |S| = h-1$ and $i \in [k] \setminus [t]$, $\mathbf{m}_i \notin \mathcal{L}(S)$.

In what follows, we assume that condition (1) holds for a suitable choice of t and argue that condition (2) is very unlikely to be satisfied. Let t be the largest integer such that

$$\binom{t}{h-1} \leqslant q.$$

The bound on q implies that $t \leqslant \lfloor k/2 \rfloor$. Note that for all positive integers x we have

$$\binom{x+1}{h-1} \Big/ \binom{x}{h-1} \leqslant \left(\frac{e(x+1)}{h-1}\right)^{h-1} \left(\frac{h-1}{x}\right)^{h-1} \leqslant (2e)^{h-1}. \tag{8}$$

Let $\varepsilon = 1/(2e)^{h-1}$. By (8) and the definition of t we have

$$\varepsilon q \leqslant \binom{t}{h-1} \leqslant q. \tag{9}$$

Let $U \subseteq \mathbb{F}_q^h$ denote the union of $\mathcal{L}(S)$ over all $S \subseteq [t], |S| = h-1$. By inclusion-exclusion we have

$$\begin{aligned}
|U| &\geqslant \binom{t}{h-1} q^{h-1} - \binom{\binom{t}{h-1}}{2} q^{h-2} \\
&\geqslant \binom{t}{h-1} q^{h-1} \left(1 - \frac{\binom{t}{h-1}}{2q}\right) \\
&\geqslant \frac{1}{2} \binom{t}{h-1} q^{h-1} \qquad (10)
\end{aligned}$$

where we used the RHS of (9).

By the discussion above

$$\begin{aligned}
\Pr\left[C' \text{ is MDS }\right] &\leqslant \prod_{i=t+1}^{k} \Pr[\mathbf{m}_i \notin U] \\
&= \left(\frac{q^h - |U|}{q^h}\right)^{k-t} \\
&\leqslant \left(1 - \frac{1}{2q}\binom{t}{h-1}\right)^{k-t} \quad \text{By Equation (10)} \\
&\leqslant \left(1 - \frac{\varepsilon}{2}\right)^{k-t} \quad \text{By Equation (9)} \\
&\leqslant \left(1 - \frac{1}{2^h e^{h-1}}\right)^{\frac{k}{2}} \quad \text{Since } \ t \leqslant k/2.
\end{aligned}$$

This concludes the proof. □

By the reduction used in the proof of Theorem 26, this gives a lower bound for maximally recoverable local codes.

Corollary 28. *Let positive integers k, r, h be such that $\ell = (k+h)/r$ is an integer. Consider a local (k, r, h)-code C, where the coefficients in the heavy parities are drawn uniformly at random and independently from a finite field $\mathbb{F}_q$ where*

$$q \leqslant \binom{\lfloor \frac{k}{2} \rfloor}{h-1}.$$

Then the probability that C is maximally recoverable is at most

$$\left(1 - \frac{1}{2^h e^{h-1}}\right)^{\frac{k}{2}}.$$

We can interpret Corollary 28 as saying that random codes cannot offer an asymptotic improvement upon the construction of Theorem 16[1]. But in the setting of MDS codes, Reed-Solomon codes give an alphabet size of $O(k+h)$. Thus this leaves open the intriguing possibility that there are similar explicit constructions of maximally recoverable local codes.

5 Maximal recoverability for general topologies

The original definition of maximal recoverability [CHL07] assumes that the code we are trying to construct has two kinds of parities: light and heavy. The light parities are simple XORs and the goal is to pick coefficients for the heavy parities. This definition has some limitations: for instance, it does not allow for locality within the heavy parities. We now present a general definition.

Let $z_1, \ldots, z_m$ be variables over a field of characteristic 2. Consider an $(n-k) \times n$ matrix $P = \{p_{ij}\}$ where each $p_{ij} \in \mathbb{F}_2[z_1, \ldots, z_m]$ is an affine function of the z_is over $\mathbb{F}_2$:

$$p_{ij}(z_1, \ldots, z_m) = b_{ij0} + \sum_{k=1}^{m} b_{ijk} z_k, \quad b_{ijk} \in \mathbb{F}_2.$$

In what follows we refer to a matrix P as a topology. Fix an assignment $\{z_i = \alpha_i\}_{i=1}^m$ where $\alpha_i \in \mathbb{F} \supseteq \mathbb{F}_2$. Viewing the resulting matrix $P(\alpha_1, \ldots, \alpha_m)$ as a parity check matrix defines a linear code of length n and dimension (at least) $n-k$ which we denote by $\mathcal{C}(\alpha_1, \ldots, \alpha_m)$. We say that this code instantiates P. We say that a set $S \subseteq [n]$ of columns of P is *potentially independent* is there exists an assignment $\{z_i = \alpha_i\}_{i=1}^m$ where $\alpha_i \in \mathbb{F} \supseteq \mathbb{F}_2$ such that the columns of $P(\alpha_i, \ldots, \alpha_m)$ indexed by S are linearly independent. Equivalently, one can say that columns in S are potentially independent if they are linearly independent over the field of rational functions $\mathbb{F}_2(z_1, \ldots, z_m)$.

Definition 29. *We say the code $\mathcal{C}(\alpha_1, \ldots, \alpha_m)$ instantiating a topology P is maximally recoverable if every set of columns that is potentially independent in P is linearly independent in $P(\alpha_1, \ldots, \alpha_m)$.*

It is easy to see by standard probabilistic arguments that for all topologies P maximally recoverable codes exist over sufficiently large finite fields.

Lemma 30. *Let $P \in (\mathbb{F}_2[z_1, \ldots, z_m])^{(n-k)\times n}$ be an arbitrary topology. Let $\mathbb{F}$ be a finite field of size more than $(n-k) \cdot \binom{n}{\leqslant n-k}$, $\operatorname{char} \mathbb{F} = 2$. There exists an MR instantiation of P over $\mathbb{F}$.*

Proof. Observe that any set of more than $(n-k)$ columns of P is not potentially independent since in that case the number of vectors exceeds their dimension. Thus the total number of potentially independent sets is at most $\binom{n}{\leqslant n-k}$. For each potentially independent set S there is a collection R_S of $|S|$ rows of P such that the determinant of the minor indexed by R_S and S is a non-zero polynomial $p_S \in \mathbb{F}_2[z_1, \ldots, z_m]$. Note that

$$\deg p_S \leqslant |S| \leqslant n-k.$$

[1]Of course, there is always a possibility that some carefully chosen random ensemble could yield MR codes over much smaller alphabets. However we are not aware even of any candidate ensembles like that.

Consider the polynomial $p(z_1, \ldots, z_m) = \prod_S p_S$. Observe that $\deg p \leqslant (n-k) \cdot \binom{n}{\leqslant n-k}$. Therefore over any field $\mathbb{F}$ of characteristic 2 such that $|\mathbb{F}| > (n-k) \cdot \binom{n}{\leqslant n-k}$, there is an assignment $\{z_i = \alpha_i\}_{i=1}^m$ such that $p(\alpha_1, \ldots, \alpha_m)$ is non-zero [LN83]. Thus for all potentially independent sets S we have $p_S(\alpha_1, \ldots, \alpha_m) \neq 0$ and thus all potentially independent sets are in fact linearly independent. □

Our goal is to find explicit MR instantiations that minimize the field size. Before we proceed we would like to make two observations.

1. In our Definition 29 above, we have not explicitly specified which are the data symbols and which are the parity check symbols. But one can do this from the topology P alone, independent of the specific choice of $\alpha_1, \ldots, \alpha_m$ as long as the resulting code is maximally recoverable. Take a set of indices whose columns form the largest potentially independent set and designate these as the parity check symbols. One can obviously make the choice canonical. The property of being the largest potentially independent set ensures that (assuming the MR property) for any choice of values of data symbols, one can satisfy all the parity check equations by assigning the parity checks suitably.

2. The locality of every coordinate is fixed by the choice of P provided that we consider maximally recoverable instantiations. Recall that the locality of i is the smallest weight of a dual codeword in $\mathcal{C}(\alpha_1, \ldots, \alpha_m)$ whose support contains i. Consider the smallest set S of column indices containing i, so that the corresponding set of columns of P is not potentially independent. We claim that the locality of i is $|S| - 1$. This holds since any smaller subset of columns containing i is potentially independent, and thus linearly independent over $\mathbb{F}$ in any MR instantiation $P(\alpha_1, \ldots, \alpha_m)$. So they cannot support dual codewords.

We present some examples of maximally recoverable codes fitting the Definition 29:

- Define the $h \times n$ topology P^{mds} as $p_{ij}^{\text{mds}} = \{z_{ij}\}$. An instantiation of P^{mds} is maximally recoverable iff the resulting code is an MDS code.

- Assume that $(r+1)|n$ and divide the n symbols into groups of size $(r+1)$. Augment the topology P^{mds} above with $n/(r+1)$ rows by adding the constraint that the XOR of each group is 0. Call the resulting topology P^{loc}. MR codes corresponding to P^{loc} are the same as local MR codes. Such codes are constructed in Theorems 12 and 16.

We do not know how to explicitly construct MR codes for an arbitrary topology P. This problem appears to be very challenging even in narrow special cases. However, in the following section we show how to generalize our Basic construction of MR local codes to accommodate a certain large class of topologies that one gets by starting with a parity check matrix of an arbitrary binary linear code over and adding few rows of generic constraints.

5.1 Maximally recoverable extensions of binary linear codes

Consider an $[n, k, d]_2$ binary linear code $\mathcal{C}$ given by a parity check matrix P over $\mathbb{F}_2$. Note that these equations define a linear code over any field $\mathbb{F}$ of characteristic 2. Assume that the resulting

code is capable of correcting a few erasures with good locality. Our goal is now to add a specified number h of parity check equations that will boost the erasure correction. Formally, define the matrix

$$P^+ = \begin{pmatrix} P \\ P^{\text{mds}} \end{pmatrix}. \tag{11}$$

Recall that P^{mds} is the $h \times n$ matrix where $p_{ij}^{\text{mds}} = \{z_{ij}\}$. We refer to instantiations of the topology P^+ as extensions of the code $\mathcal{C}$, and denote them by $\mathcal{C}^+$.

We start by identifying when such an extension $\mathcal{C}^+$ is maximally recoverable. An erasure pattern is specified by a set of unknown variables V. By substituting values for all other variables in the equations P, we get a system of equations $P|_V$. The solution space has dimension $\dim(V) = |V| - \text{Rank}(P|_V)$. As long as $\dim(V) \leqslant h$, we can hope to solve for all the unknowns using h additional equations. The code is maximally recoverable if it is able to recover from all such erasure patterns. If $\dim(V) > h$, this erasure pattern is not correctable with just h additional equations, regardless of the coefficients.

In an MR code, the distance d^+ of $\mathcal{C}^+$ will satisfy $d^+ \geqslant d + h$. The inequality might be strict for specific choices of $\mathcal{C}$, but starting from $\mathcal{C} = \mathbb{F}_2^n$ gives an example where $d = 1$ and $d^+ = h + 1$. Hence one cannot hope for a better bound for all codes.

We will present a generic construction where by increasing the field size, one can increase the number of erasure patterns that are corrected. By choosing a sufficiently large field, we can in fact get a maximally recoverable code.

Let $S = \{\alpha_1, \cdots, \alpha_n\}$ where the α_is are chosen from some extension field $\mathbb{F}$ over $\mathbb{F}_2$. We define the linear code $\mathcal{C}^+$ over the field $\mathbb{F}$ by adding the parity check equations

$$\sum_{i=1}^{n} \alpha_i^{2^{g-1}} x_i = 0 \quad \text{for} \quad g \in \{1, \ldots, h\} \tag{12}$$

in addition to the parity checks given by P. We refer to these as the global parity checks (although the parity checks in P need not be local).

The following is our main result.

Theorem 31. *If the set S is chosen to be ℓ-wise independent over $\mathbb{F}_2$, then the code $\mathcal{C}^+$ can correct any erasure pattern V where $|V| \leqslant \ell$ and $\dim(V) \leqslant h$.*

Proof. Assume that $\dim(V) = f \leqslant h$. We can fix a basis of f variables in V which we denote $x_1, \ldots, x_f$. Relabel the remaining variables as $y_1, \ldots, y_t$ where $t \leqslant \ell - f$. Using the equations in P, the y_js can be expressed as $\mathbb{F}_2$ linear combinations of $x_1, \ldots, x_f$ and a constant term from $\mathbb{F}$. Let $S_j \subseteq [t]$ be the set of variables that are required to express y_j, so that

$$y_j = \sum_{i \in S_j} x_i + c_j \quad c_j \in \mathbb{F} \tag{13}$$

Note that S_j could be empty if y_j is fixed by the equations in P.

We use these to eliminate the ys from the global parity checks. Rename the coefficients assigned by S to $x_1, \ldots, x_f$ as $\alpha_1, \ldots, \alpha_f$ and to $y_1, \ldots, y_t$ by $\alpha'_1, \ldots, \alpha'_t$. Substituting the values for known

variables in the global parity checks gives equations of the form

$$\sum_{i=1}^{f} \alpha_i^{2^{g-1}} x_i + \sum_{j=1}^{t} {\alpha'}_j^{2^{g-1}} y_j = c'_g \tag{14}$$

Using Equations (13) to eliminate the y_js, we get

$$\begin{aligned}
\sum_{i=1}^{f} \alpha_i^{2^{g-1}} x_i + \sum_{j=1}^{t} {\alpha'}_j^{2^{g-1}} \left(c_j + \sum_{i \in S_j} x_i \right) &= c'_g \\
\sum_{i=1}^{f} x_i \left(\alpha_i^{2^{g-1}} + \sum_{j : i \in S_j} {\alpha'}_j^{2^{g-1}} \right) &= c''_g \\
\sum_{i=1}^{f} x_i \left(\alpha_i + \sum_{j : i \in S_j} {\alpha'}_j \right)^{2^{g-1}} &= c''_g
\end{aligned}$$

Now applying Lemma 9, this system of equations is invertible iff the set $\{\beta_i = \alpha_i + \sum_{j:i \in S_j} {\alpha'}_j\}_{i=1}^{f}$ is linearly independent over $\mathbb{F}_2$. But note that for any non-empty $T \subseteq [f]$ we have

$$\sum_{i \in T} \beta_i = \sum_{i \in T} \left(\alpha_i + \sum_{j : i \in S_j} {\alpha'}_j \right) = \sum_{i \in T} \alpha_i + \sum_{j \in T'} {\alpha'}_j$$

where $T' \subseteq [t]$. Note that T' could be the empty set, and its size is at most $|t| \leqslant \ell - f$. Thus we get a non-empty sum of at most $|T| + |T'| \leqslant f + \ell - f = \ell$ coefficients from the set S. Assuming that the coefficients in S are ℓ-wise independent, the sum is non-zero which implies that the β_is are indeed linearly independent over $\mathbb{F}_2$. □

Corollary 32. *If the set S is chosen to be $(d+h-1)$-wise independent over $\mathbb{F}_2$, then the code $\mathcal{C}^+$ instantiating the topology (11) has distance $d^+ \geqslant d+h$. Hence there exist codes $\mathcal{C}^+$ with distance $d^+ \geqslant d+h$ over a field $\mathbb{F}$ where $|\mathbb{F}| = O_{d,h}\left(n^{(d+h)/2}\right)$.*

Proof. It suffices to show that every erasure pattern V where $|V| \leqslant d+h-1$ is corrected by $\mathcal{C}^+$. Applying Theorem 31 with $\ell = d+h-1$, this holds provided $\dim(V) \leqslant h$ for all such erasure patterns.

$P|_V$ is the restriction of the parity check matrix of $\mathcal{C}$ to the columns corresponding to V. Since $\mathcal{C}$ has minimum distance d, every set of $d-1$ columns is linearly independent. Thus $\mathrm{Rank}(P|_V) \geqslant \min(|V|, d-1)$. Since $\dim(V) = |V| - \mathrm{Rank}(P|_V)$, it follows that

$$\dim(V) \leqslant \begin{cases} 0 & \text{for } |V| \leqslant d-1 \\ |V| - (d-1) & \text{for } |V| \geqslant d. \end{cases}$$

Hence $\dim(V) \leqslant h$ as long as $|V| \leqslant d+h-1$.

For the second part of the claim, we choose αs as the columns of a BCH code of distance $d+h$ and length n. This requires $1 + \lceil (d+h-1)/2 \rceil \lceil \log(n) \rceil$ parity checks. The resulting field size is

$$2^{1+\lceil (d+h-1)/2 \rceil \lceil \log(n) \rceil} = O_{d,h}\left(n^{(d+h)/2}\right).$$

□

We introduce the parameter $\ell(\mathcal{C}, h)$ that bounds the cardinality of any set V of unknowns for which the space of solutions in $\mathcal{C}$ is of dimension at most h. Let

$$\ell(\mathcal{C}, h) = \max_{V:\dim(V)\leqslant h} |V|.$$

The following claim follows from Theorem 31, and by choosing αs using a suitable BCH code, exactly the we way we do this in the proof of Corollary 32.

Corollary 33. *There exist maximally recoverable codes $\mathcal{C}^+$ instantiating the topology (11) over a field $\mathbb{F}$ where $|\mathbb{F}| = O\left((2n)^{(\ell(\mathcal{C},h)/2)}\right)$.*

Proof. It suffices to choose αs which are $\ell(\mathcal{C}, h)$-wise independent. We choose αs as the columns of a BCH code of distance $\ell(\mathcal{C}, h) + 1$ and length n. This requires $1 + \lceil (\ell(\mathcal{C}, h) - 1)/2 \rceil \cdot \lceil \log(n) \rceil$ parity checks. The resulting field size is at most

$$2^{1+\lceil (\ell(\mathcal{C},h)-1)/2 \rceil \cdot (\log(n)+1)} = O\left((2n)^{\ell(\mathcal{C},h)/2}\right).$$

□

We present an alternate view of $\ell(\mathcal{C}, h)$ that might be useful. For a codeword $c \in \mathcal{C}$, let $\mathsf{Supp}(c) \subseteq [n]$ denote its support. For a set of indices $I \subseteq [n]$, let $\mathcal{C}(I) = \{c \in \mathcal{C} \mid \mathsf{Supp}(c) \subseteq I\}$. It is easy to see that $\mathcal{C}(I)$ is a subspace of $\mathcal{C}$. Let $\dim_{\mathcal{C}}(I)$ denote its dimension. It follows that

$$\ell(\mathcal{C}, h) = \max_{I:\dim_{\mathcal{C}}(I)\leqslant h} |I|.$$

To see why this is true, take I to be the indices corresponding to the variables V. For any setting of the other variables, the kernel of the resulting system of equations is exactly $\mathcal{C}(I)$. Hence it follows that $\dim(V) = \dim_{\mathcal{C}}(I)$.

Thus for instance $\ell(\mathcal{C}, 0)$ is size of the largest set of indices that does not contain the support of a non-zero codeword. One can view this as a dual notion to the minimum distance, which is the size of the smallest set of indices that supports a codeword. Similarly, $\ell(\mathcal{C}, h)$ for larger h is a dual of the notion of generalized Hamming weights. The h^{th} generalized Hamming weight is the size of the smallest set I such that $\dim_{\mathcal{C}}(I) \geqslant h$, whereas $\ell(\mathcal{C}, h-1)$ is the size of the largest set I such that $\dim_{\mathcal{C}}(I) < h$. While these parameters seems fairly natural, we are unaware of prior work that studies them. For specific codes $\mathcal{C}$, it might be possible to get good bounds on them and via Corollary 33, get good constructions of maximally recoverable codes $\mathcal{C}^+$.

6 Open questions

We studied the trade-off between maximal recoverability and alphabet size in codes with locality. Lots of questions in this area remain open. The main immediate challenge is to reduce the field in constructions of Theorems 12 and 16 or to prove that such a reduction is not possible.

1. In the asymptotic setting of constant r and h and growing k, can one get local MR codes over a field of size $O(k)$? Or do local MR codes inherently require a larger field than MDS codes?

2. In the setting of $h = O(1)$, $r = \Theta(k)$, and growing k, can one get a lower bound of $\omega(k)$ for the field size of local MR codes?

Looking more broadly we are interested in explicit constrictions of maximally recoverable codes (or lower bounds for the field size of MR codes) in other basic topologies that generalize local codes. We expect many such topologies to become practically relevant.

Given that known constructions of maximally recoverable codes require fairly large finite fields it is also interesting to explore codes that provide weaker guarantees than maximal recoverability, but require smaller field sizes. One family of such codes has been considered in [PBH13].

Acknowledgements

We would like to thank Swastik Kopparty and Raghu Meka for allowing us to include their Theorem 27 in this paper.

References

[Bal12] Simeon Ball. On sets of vectors of a finite vector space in which every subset of basis size is a basis. *Journal of European Mathematical Socisety*, 14(1-2):733–748, 2012.

[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). *Designs Codes and Cryptography*, 65(1-2):5–14, 2012.

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

[Bla13] Mario Blaum. Construction of PMDS and SD codes extending RAID 5. Arxiv 1305.0032, 2013.

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

[DGW+10] Alexandros G. Dimakis, Brighten Godfrey, Yunnan Wu, Martin J. Wainwright, and Kannan Ramchandran. Network coding for distributed storage systems. *IEEE Transactions on Information Theory*, 56(9):4539–4551, 2010.

[Dum95] Ilya Dumer. Nonbinary double-error-correcting codes designed by means of algebraic varieties. *IEEE Transactions on Information Theory*, 41(6):1657–1666, 1995.

[EB99] Yves Edel and Juergen Bierbrauer. Recursive constructions for large caps. *Bulletin of Belgian Mathematical Society*, 6:249–258, 1999.

[FY13] Michael Forbes and Sergey Yekhanin. On the locality of codeword symbols in non-linear codes. Arxiv 1303.3921, 2013.

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

[HCL07] Cheng Huang, Minghua Chen, and Jin Li. Pyramid codes: flexible schemes to trade space for access efficiency in reliable data storage systems. In *Sixth IEEE International Symposium on Network Computing and Applications (NCA 2007)*, pages 79–86, 2007.

[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 *Proceedings of the 2012 USENIX conference on Annual Technical Conference*, pages 2–2, 2012.

[KM] Swastik Kopparty and Raghu Meka. Personal communication, April 2013.

[LN83] Rudolf Lidl and Harald Niederreiter. *Finite Fields*. Cambridge University Press, Cambridge, 1983.

[MS77] F. J. MacWilliams and N. J. A. Sloane. *The Theory of Error Correcting Codes*. North Holland, Amsterdam, New York, 1977.

[PBH13] James S. Planck, Mario Blaum, and James Lee Hafner. SD codes: erasure codes designed for how storage systems really fail. In *Proceedings of the 2013 USENIX conference on File and Storage Technologies*, 2013.

[PD12] Dimitris S. Papailiopoulos and Alexandros G. Dimakis. Locally repairable codes. In *Proceedings of the 2012 IEEE International Symposium on Information Theory (ISIT)*, pages 2771–2775, 2012.

[PGM13] J. S. Plank, K. M. Greenan, and E. L. Miller. Screaming fast Galois Field arithmetic using Intel SIMD instructions. In *FAST-2013: 11th Usenix Conference on File and Storage Technologies*, San Jose, February 2013.

[PKLK12] N. Prakash, Govinda M. Kamath, V. Lalitha, and P. Vijay Kumar. Optimal linear codes with a local-error-correction property. In *Proceedings of the 2012 IEEE International Symposium on Information Theory (ISIT)*, pages 2776–2780, 2012.

[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. *arXiv*, abs/1301.3791, 2013.

[YD04] Sergey Yekhanin and Ilya Dumer. Long non-binary codes exceeding the Gilbert-Varshamov bound for any fixed distance. *IEEE Transactions on Information Theory*, 10(50):2357–2362, Oct. 2004.

[Yek12] Sergey Yekhanin. Locally decodable codes. *Foundations and trends in theoretical computer science*, 6(3):139–255, 2012.

A Reliability analysis for MDS codes

In this section, we formalize a set of assumptions that justify the observation that in data storage applications one typically does not scale the number of heavy parities linearly with the number of data fragments k to ensure the same level of reliability [HSX+12]. We make the following assumptions:

1. There are k disks worth of data that we wish to store reliably. In practice k is typically of the order of tens or (at most) hundreds.

2. Each disk has a failure probability p in a certain time window. For instance, we may take a time window of a day, and $p = \frac{1}{1500}$.

3. The failure probability p is smaller than $1/k$. We will assume that $kp \leqslant 1/2$. Note that p is an absolute constant which depends on the device characteristics. So this assumption also bounds k by an absolute constant $(1/2p)$. But this is consistent with the numbers observed in practice.

4. The goal is to achieve at least the same level of reliability as 3-way replication (similar calculations work for any c in place of 3).

We will compute the number of heavy parities required for a $[k+h, k, h+1]_q$ MDS code to achieve the desired level of reliability. Note that local MR (k, r, h)-codes are at least as reliable as such $[k+h, k]_q$ MDS codes, since one gets such a code by ignoring the local parities.

Let us first analyze 3-way replication. The failure probability within a time window which we denote by $f(3)$ can be bounded as

$$p^3 \leqslant f(3) \leqslant kp^3.$$

Now consider the $[k+h, k]_q$ MDS code where $k \geqslant 5$. For data loss, some $h+1$ machines must fail in that window. Hence we can bound the failure probability which we denote $f(\text{MDS})$ by

$$f(\text{MDS}) \leqslant \binom{k+h}{h+1} p^{h+1} \leqslant \left[\left(\frac{k+h}{h+1} \right) e \right]^{h+1} p^{h+1} \leqslant (2kp)^{h+1}.$$

The latter inequality above follows since when $k \geqslant 5$ the expression in square brackets is upper bounded by $2k$. Since our goal is to have $f(\mathrm{MDS}) \leqslant f(3)$, it suffices that

$$(2kp)^{h+1} \leqslant p^3$$

which (by taking logarithms) is equivalent to requiring

$$h+1 \geqslant \frac{3\log(1/p)}{\log(1/2pk)} = 3\left(1 + \frac{\log(2k)}{\log(1/2pk)}\right) \tag{15}$$

Since by our assumption, $2pk \leqslant 1$, it is sufficient to take

$$h \geqslant 2 + 3\log(2k) = 5 + 3\log(k)$$

Indeed, if we make stronger assumptions like $pk \ll 1$, then the growth as a function of k is even slower than $\log(k)$.

We stress that this is a over-simplification that should not be taken too literally. But it does suggest that for current values of k and p, scaling h linearly with k is overly redundant.

B Proof of Theorem 20

Our goal here is to prove

Theorem 20. *[Dum95] Let q be prime power and m be an even integer. There exists a q-ary linear code with parameters $[q^m, q^m - \frac{3m}{2} - 1, 4]$.*

Proof. Our code is defined by the means of a $\left(\frac{3m}{2}+1\right) \times q^m$ parity check matrix M. Columns of M correspond to elements $x \in \mathbb{F}_{q^m}$. For all x we have, $M(x) = (1, x, x^{q^{m/2}+1})$. Observe that each column is indeed a $\left(\frac{3m}{2}+1\right)$-dimensional vector over $\mathbb{F}_q$ as $x^{q^{m/2}+1} \in \mathbb{F}_{q^{m/2}}$. We need to argue that any three distinct columns

$$\begin{pmatrix} 1 & 1 & 1 \\ x_1 & x_2 & x_3 \\ x_1^{q^{m/2}+1} & x_2^{q^{m/2}+1} & x_3^{q^{m/2}+1} \end{pmatrix} \tag{16}$$

of M are linearly independent. Assume the contrary. Suppose there exist $\alpha, \beta, \gamma \in \mathbb{F}_q$ not all zero, such that $\alpha M(x_1) + \beta M(x_2) + \gamma M(x_3) = 0$. We need two observations to get a contradiction. Our first observation is that x_1, x_2, and x_3 cannot all lie in $\mathbb{F}_{q^{m/2}}$ (in this case (16) is a Vandermonde matrix). Our second observation is that the algebraic variety

$$\begin{cases} \alpha & + & \beta & + & \gamma & = & 0, \\ \alpha x_1 & + & \beta x_2 & + & \gamma x_3 & = & 0, \\ \alpha x_1^{q^{m/2}+1} & + & \beta x_2^{q^{m/2}+1} & + & \gamma x_3^{q^{m/2}+1} & = & 0; \end{cases} \tag{17}$$

defined in $\mathbb{F}_{q^m}^3$ over x_1, x_2, x_3 is invariant under the affine transformations $L(x_i) = Ax_i + B$, applied simultaneously to all x_i, where $A \neq 0$ and B are arbitrary elements of $\mathbb{F}_{q^m}$. To see this let

us replace each x_i in (17) with $L(x_i)$. We obtain

$$\begin{cases} \alpha+\beta+\gamma & =0, \\ A(\alpha x_1+\beta x_2+\gamma x_3) & + \\ B(\alpha+\beta+\gamma) & =0, \\ A^{q^{m/2}+1}(\alpha x_1^{q^{m/2}+1}+\beta x_2^{q^{m/2}+1}+\gamma x_3^{q^{m/2}+1}) & + \\ A^{q^{m/2}}B(\alpha x_1^{q^{m/2}}+\beta x_2^{q^{m/2}}+\gamma x_3^{q^{m/2}}) & + \\ AB^{q^{m/2}}(\alpha x_1+\beta x_2+\gamma x_3) & + \\ B^{q^{m/2}+1}(\alpha+\beta+\gamma) & =0. \end{cases} \tag{18}$$

One can easily see that (17) implies (18). (To demonstrate that the second summand of the last equation of (18) is zero one should raise the second equation of (17) to the power of $q^{m/2}$). Note also that L is a one-to-one mapping. Now let us choose A and B in such a way that $L(x_1)$ and $L(x_2)$ are distinct elements of $F_{q^{m/2}}$. The second equation of (18) asserts that $L(x_3)$ is also in $F_{q^{m/2}}$. Thus we get a contradiction with our first observation. □