Electronic Colloquium on Computational Complexity, Report No. 104 (2013)

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. Local parities allow to quickly recover any single symbol when it is erased, while 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 recoverable given the code topology. In this paper we present explicit families of maximally-recoverable codes with locality. We also initiate the study of the trade-off between maximal recoverability and alphabet size.

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 has been two lines of work on codes with locality.

In [10] motivated by applications to distributed storage [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 canonical codes of [10], r divides k and $n = \frac{k}{r} + h$. Data symbols are partitioned into $\frac{k}{r}$ groups of size r. For each data group there is a local parity storing the XOR of respective data symbols. There also are h heavy parities where each heavy parity depends on all k data symbols. In what follows we refer to codes above as *data-local* (k, r, h)-codes.

In [4] motivated by applications to data storage on SSDs the authors studied systematic linear $[n, k]$ codes with two extra parameters r and h, where $r \mid (k + h)$. In codes of [4] there are k data symbols and h heavy parity symbols. These $(k + h)$ symbols are partitioned into $\frac{k+h}{r}$ groups of size r. For each group there is a local parity storing the XOR of respective symbols. Thus $n = k + h + \frac{k+h}{r}$. Unlike the codes of [10], codes of [4] provide locality for all symbols data or parity. In what follows we refer to codes above as *local* (k, r, h)-codes.

ISSN 1433-8092

Observe that our descriptions of code families above are so far incomplete. For every parity symbol we specified other symbols that it depends on, i.e., we have fixed codes' topology. To complete defining the codes we need to set coefficients in the heavy parity symbols. Different choices of coefficients lead to codes with different erasure correcting capabilities. Ideally, we would like our codes to correct all patterns of erasures that are correctable for some setting of coefficients in heavy parities. Such codes exist and are called Maximally Recoverable (MR) [5].

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.

1.1 Our results and related work

In this paper we make progress on the problem above. 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 $h = 2$. There have also been few constructions that involved computer search for coefficients [4, 3]. Our codes improve upon the earlier constructions both in concrete settings and asymptotically.

In the asymptotic setting of $h = O(1), r = O(1)$, and growing k our codes use alphabet of size $O\left(k^{h-1}\right)$. In the case of $h \geqslant 2^r$ the alphabet size can be reduced $O\left(k^{\lceil (h-1)\left(1-\frac{1}{2^r}\right)\rceil}\right)$. 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 [14] and is $\Omega(k)$. 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 with random 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 $\Omega\left(k^{h-1}\right)$.

Similarly to [4, 3] we construct our explicit codes via parity check matrices. As in [4] columns of our parity check matrices have the shape $(\alpha_i, \alpha_i^2, \ldots, \alpha_i^{2^{h-1}})$. The key difference from the work of [4, 3] however is that we explicitly specify the sets $\{\alpha_i\}$ used in our constructions.

There are several other models of codes with locality in the literature. The ones most closely related to our work include SD codes [3, 17], locally decodable codes [20], and regenerating codes [6].

1.2 Organization

In section 2 we formally define data-local and local (k, r, h)-codes. We introduce the notion of maximal recoverability, and show that maximally-recoverable local codes yield maximally-recoverable data-local codes. In section 3 we give our two main code constructions. In section 4 we analyze the asymptotic behavior of alphabet size in our codes for large message lengths. We also establish a simple lower bound on the alphabet size of maximally-recoverable local codes. Finally, we compare asymptotic parameters of our codes to asymptotic parameters of random codes. In section 5 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 introduce the notion of locality [10].

Definition 1. *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 $n = \frac{k}{r} + h$;*
- *Data symbols are partitioned into $\frac{k}{r}$ groups of size r. For each such group there is one (local) parity symbol that stores the XOR of respective data symbols;*
- *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 [11, 10, 16, 18, 19, 9]. The importance of this topology was partially explained in [10, Theorem 9]. There it has been shown that in case $h < r + 1$, 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 $\frac{k}{r}$ 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 maximally-recoverable data-local (k, r, h)-code 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.

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 $\frac{k}{r} - t$ 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, $|E'| > \frac{k}{r} + h$. 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* $n = k + h + \frac{k+h}{r}$;
- *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 $\frac{k+h}{r}$ 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 $\frac{k+h}{r}$ local groups, puncturing C in coordinates in E yields a maximum distance separable $[k + h, k]$ code.*

Maximally recoverable local (k, r, h)-codes have been originally introduced in [4] under the name of partial-MDS codes. Similarly to the discussion following definition 3 it is easy to see that these codes correct all erasure patterns that involve erasing one coordinate per local group, and h additional coordinates. Erasure patterns that are not dominated by such patterns are not correctable by any local (k, r, h)-code. The next lemma gives a simple reduction from local MR codes to data-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 $t = \frac{k+h}{r}$. 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 $\{G_i\}, i \leqslant \frac{k}{r}$ as data symbols of C'. The code C' is obtained form the code C by dropping local parity symbols from groups G_i for $i > \frac{k}{r}$. 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 $s = \lfloor \frac{k}{r} \rfloor$. 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. □

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 $A(S, h) = [a_{ij}]$ denote the $h \times n$ matrix where

$$a_{ij} = \alpha_j^{2^{i-1}}$$

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

$$\sum_{i=1}^{n} \alpha_i^{2^{j-1}} x_i = 0 \text{ for } j = 1, \ldots, h. \tag{1}$$

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 $\{\alpha_i\}$ where $d = h+1$, so that the code $\mathcal{C}(\alpha, h)$ 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 $\mathbb{F}'$.*

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 $\{x_i\}_{i \in E}$ can be recovered if we know the values of all $\{x_i\}_{i \in [n] \setminus E}$. This requires solving the following system of equations:

$$\sum_{i \in E} \alpha_i^{2^{j-1}} x_i = b_j, \quad 1 \leqslant j \leqslant h \tag{2}$$

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., [13, Lemma 3.51]) that A_E is invertible if and only if the multi-set $\{\alpha_i\}_{i \in E}$ 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 $\ell = \frac{n}{r+1}$ local groups, with group i containing $r+1$ symbols $x_{i,1}, \ldots, x_{i,r+1}$. Variables in each local group satisfies 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

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

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

We refer to Equations (3) as global constraints and (4) as local constraints. The following proposition is central to our method:

Proposition 10. *Let $\mathcal{C}(S, r, h)$ be the code defined above. 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 $\{i, \mathbf{e}(i)\}_{i=1}^{\ell}$. Then $\mathcal{C}^{-\mathbf{e}}$ is an MDS code if and only if the multi-set*

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

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 $j \in [h]$:

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

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 Basic construction

Lemma 11. *Let $S \subseteq \mathbb{F}, |S| = n$ be a set that is $2h$-wise independent over as subfield $\mathbb{F}'$. Let r be arbitrary such that $\ell = \frac{n}{r+1}$ 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 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

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

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. □

3.2 Optimized construction

In the previous section we used $2h$-wise independence of the set S to ensure h-independence of sets $T(S, \mathbf{e})$. In some cases this is on overkill, and one can ensure h-independence of sets $T(S, \mathbf{e})$ more economically.

Definition 13. *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. The following proposition presents our approach is a general form.

Proposition 14. *Let positive integers k, r, h be such that $\ell = \frac{k+h}{r}$ is an integer. Suppose there exists an $(r+1)$-sized set $S_1 \subseteq \mathbb{F}_{2^a}$. If h is even we require S_1 to be h-weakly independent over $\mathbb{F}_2$; otherwise we require S_1 to be $(h+1)$-weakly independent over $\mathbb{F}_2$. Further suppose that there exists an ℓ-sized set $S_2 \subseteq \mathbb{F}_{2^b}$ that is h-independent over $\mathbb{F}_{2^a}$; then the code $\mathcal{C}(S_1 \cdot S_2, r, h)$ is a maximally recoverable local (k, r, h)-code over the field $\mathbb{F}_{2^b}$.*

Proof. Let $S_1 = \{\xi_1, \ldots, \xi_{r+1}\}$. Let $S_2 = \{\lambda_1, \ldots, \lambda_\ell\}$. For $i \in [\ell], s \in [r+1]$, we set $\alpha_{i,s} = \lambda_i \xi_s$. By Proposition 10 it suffices to show that for all $\mathbf{e} \in [r+1]^\ell$, the set $T(S_1 \cdot S_2, \mathbf{e})$ is h-independent over $\mathbb{F}_2$. Assume the contrary. 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 Proposition 14 with a certain particular choice of independent sets. Our sets come from columns of a parity check matrix of an extended BCH code.

Theorem 15. *Let positive integers k, r, h be such that $\ell = \frac{k+h}{r}$ is an integer. Let m be the smallest integer such that $r \mid m$ and $\ell \leqslant 2^m$; then there exists a maximally recoverable local (k, r, h)-code over the field $\mathbb{F}_{2^t}$ for $t = r + m\left\lceil (h-1)\left(1 - \frac{1}{2^r}\right)\right\rceil$.*

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 $S_2' = \{\beta_1, \ldots, \beta_\ell\}$ 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 $i \in [\ell]$,

$$\lambda_i = (1, \beta_i, \beta_i^2, \ldots, \beta_i^{h-1}) \tag{5}$$

when we treat $\mathbb{F}_{2^t}$ as a linear space over $\mathbb{F}_{2^r}$. The first coordinate in (5) is a single value in $\mathbb{F}_{2^r}$, while every other coordinate is an $\frac{m}{r}$-dimensional vector. In formula (5) we also omit every non-zero power β_i^j whenever $2^r \mid j$. 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

$$\sum_{i \in S} \gamma_i \lambda_i^j, \tag{6}$$

where we assume $0^0 = 1$ and all $\{\gamma_i\} \in \mathbb{F}_{2^r}$. By standard properties of Frobenius automorphisms (6) implies

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

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

Example 16. Instantiating Theorem 15 with $k = 60$, $r = h = 4$, we obtain a $[80, 60, 7]$ maximally recoverable $(60, 4, 4)$ local code over the field $\mathbb{F}_{2^{16}}$. Prior to our work [4, Theorem 4.2] a code with such parameters was not known to exist over any field of size below 2^{80}.

In the proof of Theorem 15 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 $\left[\ell, \ell - \frac{t}{r}, h+1\right]$ code over $\mathbb{F}_{2^r}$ to define the set $S_2 \subseteq \mathbb{F}_{2^t}$ and obtain a MR local (k, r, h)-code over $\mathbb{F}_{2^t}$. While we used columns of the parity check matrix of an extended BCH code, other choices sometimes yield local MR codes over smaller alphabets.

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

In this section we carry out the steps outlined above and present codes that improve upon the codes of Theorem 15 for $h = 3$ or 4 and large k. We replace BCH codes in the construction of Theorem 15 with better codes. The codes we use are not new [7, 21].

Theorem 17. *Let positive integers $k, r, h = 3$ be such that $\ell = \frac{k+h}{r}$ 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 $t = r(\frac{3m}{2} + 1)$.*

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 $S_2' \subseteq \mathbb{F}_{2^r}^{\frac{3}{2}m+1}$ be an arbitrary collection of ℓ columns of the parity check matrix of the code C' from [21, Theorem 5], where we set $q = 2^r$ and $d = 4$. 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 [8] one can get further small improvements upon the theorem above.

Theorem 18. *Let positive integers $k, r, h = 4$ be such that $\ell = \frac{k+h}{r}$ 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 $t = r(2m + \frac{m-1}{3})$.*

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 $S_2' \subseteq \mathbb{F}_{2^r}^{2m+\frac{m-1}{3}}$ be an arbitrary collection of ℓ columns of the parity check matrix of the code U' from [7, Theorem 6], where we set $q = 2^r$. S_2' naturally defines a set $S_2 \subseteq \mathbb{F}_{2^t}$ that is 4-independent over $\mathbb{F}_{2^r}$. □

4 Asymptotic parameters

Unlike data transmission applications, in data storage applications one typically does not need to scale the number of heavy parities linearly with the number of data fragments k to ensure the same level of reliability [12], as the likelihood p a fragment failure during a certain period of time is usually much smaller than $\frac{1}{k}$. Much slower growth in the number of heavy parities suffices. Therefore we find the asymptotic setting of fixed r, h and growing k relevant for practice and analyze the growth rate of the alphabet size in different families of local MR (k, r, h)-codes in this regime.

It is not hard to see that in codes of Theorem 12 the alphabet size grows as $O\left(k^h\right)$. In codes of Theorem 15 the alphabet size grows as $O\left(k^{\lceil (h-1)\left(1-\frac{1}{2^r}\right)\rceil}\right)$. For small values of h one can get some further improvements. MR local $(k, r, h = 3)$-codes of Theorem 17 use alphabet of size $O\left(k^{\frac{3}{2}}\right)$. MR local $(k, r, h = 4)$-codes of Theorem 18 use alphabet of size $O\left(k^{\frac{7}{3}}\right)$.

Obtaining constructions with reduced alphabet size remains a major challenge. The only lower bound we currently have comes from results on the main conjecture for MDS codes and is $\Omega(k)$. In particular the asymptotic lower bound does not depend on h.

Theorem 19. *Let C be a maximally recoverable local (k, r, h)-code with $h \geqslant 2$. Assume C is 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 [1, 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 [1, 2]. In particular, results there allow one to get small non-asymptotic improvements upon Theorem 19.

4.1 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 in this section 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 $\Omega\left(k^{h-1}\right)$. The following theorem is due to Swastik Kopparty and Raghu Meka [15].

Theorem 20. *Let positive integers k, r, h be such that $\ell = \frac{k+h}{r}$ is an integer. Consider a local (k, r, h)-code C, where the coefficients in heavy parities are drawn at random uniformly and independently from a finite field $\mathbb{F}_q$. Suppose $q \leqslant \binom{\lfloor \frac{k}{2} \rfloor}{h-1}$; then the probability that C is maximally-recoverable is at most $\left(1 - \frac{1}{2^h e^{h-1}}\right)^{\frac{k}{2}}$.*

Proof. Let $t \leqslant \lfloor \frac{k}{2} \rfloor$ be the largest integer such that $\binom{t}{h-1} \leqslant q$. Note that for all positive integers x we have

$$\binom{x+1}{h-1} / \binom{x}{h-1} \leqslant \left(\frac{e(x+1)}{h-1}\right)^{h-1} / \left(\frac{x}{h-1}\right)^{h-1} \leqslant (2e)^{h-1}. \tag{7}$$

Let $\varepsilon = \frac{1}{(2e)^{h-1}}$. By (7) and the definition of t we have

$$\varepsilon q \leqslant \binom{t}{h-1} \leqslant q. \tag{8}$$

Consider the $[k+h, k]$ code C' that is obtained from C by deleting all ℓ local parities. Let M be the $h \times (k+h)$ parity check matrix of C'. Columns of M that correspond to heavy parities form the $h \times h$ identity matrix. Other k columns $\mathbf{m}_1, \ldots, \mathbf{m}_k$ are drawn from $\mathbb{F}_q^h$ uniformly at random. In what follows for $S \subseteq [k]$ we denote the span of vectors $\{\mathbf{m}_i\}_{i \in S}$ by $\mathcal{L}(S)$. The code C is maximally recoverable only if C' is MDS. The code C' is MDS only if any h vectors in $\{\mathbf{m}_1, \ldots, \mathbf{m}_t\}$ are linearly independent and 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 any h vectors in $\{\mathbf{m}_i\}_{i \in [t]}$ are indeed independent. 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

$$|U| \geqslant \binom{t}{h-1} q^{h-1} - \binom{\binom{t}{h-1}}{2} q^{h-2} \geqslant \binom{t}{h-1} q^{h-1} \left(1 - \frac{\binom{t}{h-1}}{2q}\right). \tag{9}$$

By the discussion above

$$\begin{aligned}
\Pr\left[C' \text{ is MDS }\right] &\leqslant \prod_{i=t+1}^{k} \Pr[\mathbf{m}_i \notin U] \\
&= \left(\frac{q^h - |U|}{q^h}\right)^{k-t} \\
&\leqslant \left(1 - \frac{\binom{t}{h-1}}{q}\left(1 - \frac{\binom{t}{h-1}}{2q}\right)\right)^{k-t} \\
&\leqslant \left(1 - \frac{\binom{t}{h-1}}{2q}\right)^{k-t},
\end{aligned}$$

where the last bound follows by using the RHS of (8) inside the inner brackets. Finally, using the LHS of (8) in the formula above we obtain

$$\Pr\left[C \text{ is MR }\right] \leqslant \left(1 - \frac{\varepsilon}{2}\right)^{k-t} \leqslant \left(1 - \frac{1}{2^h e^{h-1}}\right)^{\frac{k}{2}}.$$

This concludes the proof. □

One way to interpret Theorem 20 is as saying that random codes cannot offer an asymptotic improvement upon the construction of Theorem 15.

5 Open questions

We studied the trade-off between maximal recoverability and alphabet size in local codes. Most questions in this area remain open. The main challenge is to reduce the field in constructions of Theorems 12 and 15 or to prove that such a reduction is not possible.

1. We are particularly interested in the asymptotic setting of constant r and h and growing k. In this setting can one get local MR codes over a field of size $O(k)$ or local MR codes inherently require a larger field than then their MDS counterparts?

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?

3. While data-local and local codes present two important practically motivated code topologies, constructing MR codes over other topologies is also of interest. Below we sketch the general definitions of code topology and maximal recoverability.

   Assume there are two kinds of characters $\{x_i\}_{i\in[k]}$ and $\{\alpha_j\}_{j\in[t]}$. Characters $\{x_i\}$ represent data symbols and characters $\{\alpha_j\}$ represent free coefficients. An $[n, k]$ systematic code topology is a collection of n expressions $\{E_\ell\}_{\ell\in[n]}$ in $\{x_i\}$ and $\{\alpha_j\}$. Such a collection includes all individual characters $x_1, \ldots, x_k$. Every other expression has the form

   $$E_s = \sum_i L_{i,s}(\alpha_1, \ldots, \alpha_t)x_i,$$

   where $L_{i,s}$'s are arbitrary linear functions of $\{\alpha_j\}$ over a field $\mathbb{F}'$. Specifying code topology allows one to formally capture locality constraints that one wants to impose on the code. Fixing values of coefficients $\{\alpha_j\}$ in a field $\mathbb{F}$ extending $\mathbb{F}'$ turns a code topology into a systematic $[n, k]$ code over $\mathbb{F}$.

   We say that a sub-collection S of expressions implies an expression E_i if E_i can obtained as a linear combination of expressions in S, where the coefficients are rational functions in $\{\alpha_j\}$. We say that an instantiation of a topology is maximally recoverable if every implication as above still holds after we instantiate $\{\alpha_j\}$'s. In other words, instantiating the corresponding rational functions does not cause a division by zero.

Acknowledgements

We would like to thank Swastik Kopparty and Raghu Meka for allowing us to include their Theorem 20 in the current paper.

References

[1] 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*, 14:733–748, 2012.

[2] 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.

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

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

[5] 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.

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

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

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

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

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

[11] 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.

[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.

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

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

[15] Ragu Meka and Swastik Kopparty. Personal communication, April 2013.

[16] 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.

[17] 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.

[18] 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.

[19] 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.

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

[21] 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.

ECCC ISSN 1433-8092

http://eccc.hpi-web.de