Constructions of Maximally Recoverable Local Reconstruction Codes via Function Fields

Venkatesan Guruswami

Computer Science Department, Carnegie Mellon University, Pittsburgh, PA, USA
venkatg@cs.cmu.edu

Lingfei Jin [ID]

Shanghai Key Laboratory of Intelligent Information Processing, School of Computer Science, Fudan University, Shanghai, China
Shanghai Institute of Intelligent Electronics & Systems, Shanghai, China
Shanghai Bolckchain Engineering Research Center, Fudan University, Shanghai 200433, China
lfjin@fudan.edu.cn

Chaoping Xing

School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore
xingcp@ntu.edu.sg

―――― **Abstract** ――――

Local Reconstruction Codes (LRCs) allow for recovery from a small number of erasures in a local manner based on just a few other codeword symbols. They have emerged as the codes of choice for large scale distributed storage systems due to the very efficient repair of failed storage nodes in the typical scenario of a single or few nodes failing, while also offering fault tolerance against worst-case scenarios with more erasures. A maximally recoverable (MR) LRC offers the best possible blend of such local and global fault tolerance, guaranteeing recovery from all erasure patterns which are information-theoretically correctable given the presence of local recovery groups. In an (n, r, h, a)-LRC, the n codeword symbols are partitioned into r disjoint groups each of which include a local parity checks capable of locally correcting a erasures. The codeword symbols further obey h heavy (global) parity checks. Such a code is maximally recoverable if it can correct all patterns of a erasures per local group plus up to h additional erasures anywhere in the codeword. This property amounts to linear independence of all such subsets of columns of the parity check matrix.

MR LRCs have received much attention recently, with many explicit constructions covering different regimes of parameters. Unfortunately, all known constructions require a large field size that is exponential in h or a, and it is of interest to obtain MR LRCs of minimal possible field size. In this work, we develop an approach based on function fields to construct MR LRCs. Our method recovers, and in most parameter regimes improves, the field size of previous approaches. For instance, for the case of small $r \ll \varepsilon \log n$ and large $h \geqslant \Omega(n^{1-\varepsilon})$, we improve the field size from roughly n^h to $n^{\varepsilon h}$. For the case of $a = 1$ (one local parity check), we improve the field size quadratically from $r^{h(h+1)}$ to $r^{h\lfloor (h+1)/2 \rfloor}$ for some range of r. The improvements are modest, but more importantly are obtained in a unified manner via a promising new idea.

46th International Colloquium on Automata, Languages, and Programming (ICALP 2019).
Editors: Christel Baier, Ioannis Chatzigiannakis, Paola Flocchini, and Stefano Leonardi;
Article No. 68; pp. 68:1–68:14

Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1   Introduction

Interest in erasure codes has surged in recent years, with the demands of massive cloud storage systems raising hitherto unexplored, yet very natural and mathematically deep, questions concerning the parameters, robustness, and efficiency of the code. Distributed storage systems need to build in redundancy in the data stored in order to cope with the loss or inaccessibility of the data on one or more storage nodes. Traditional erasure codes offer a natural strategy for such robust data storage, with each storage node storing a small part of the codeword, so that the data is protected against multiple node failures. In particular, MDS codes such as Reed-Solomon codes can operate at the optimal storage vs. reliability trade-off – for a given amount of information to be stored and available storage space, these codes can tolerate the maximum number of erasures without losing these stored information.

Individual storage nodes in a large scale system often fail or become unresponsive. Reconstruction (repair) of the content stored on a failed node with the help of remaining active nodes is important to reinstate the system in the event of a permanent node failure, and to allow access to the data stored on a temporarily unavailable node. The use of erasure codes in large storage systems, therefore, brings to the fore a new requirement: the ability to *very efficiently* reconstruct *parts* of a codeword from the rest of the codeword.

Local Reconstruction Codes (LRCs), introduced in [7], offer an attractive way to meet this requirement. An LRC imposes local redundancies in the codewords, so that a single (or a small number of) erased symbol can be recovered locally from less than r other codeword symbols.[1] Here r is the locality parameter that is typically much smaller than the code length n. In the distributed storage context, an LRC allows for the low-latency repair of any failed node as one only needs to wait for the response from r nodes. LRCs have found spectacular practical applications with their use in the Windows Azure storage system [12].

The challenge in an LRC design is to balance the locality requirement, that allows fast recovery from a single or few erasures, with good global erasure-resilience (via traditional slower methods) for more worst-case scenarios. One simple metric for global fault tolerance is the minimum distance d of the code, which means that any pattern of fewer than d erasures can be corrected. The optimal trade-off between the distance, redundancy, and locality of an LRC was established in [8], and an elegant sub-code of Reed-Solomon codes meeting this bound was constructed in [17].

This work concerns a much stronger requirement on global fault-tolerance, called *Maximal Recoverability*. This requires that the code should simultaneously correct every erasure pattern that is information-theoretically possible to correct, given the locality conditions imposed on the codeword symbols. Let us describe it more formally in the setting of interest in this paper. Define an $(n, r, h, a)_\ell$-LRC to be a linear code over $\mathbb{F}_\ell$ of length n whose n codeword symbols are partitioned into r disjoint groups each of which includes a local parity checks capable of locally correcting a erasures. The codeword symbols further obey h heavy (global) parity checks. With this structure of parity checks, it is not hard to see that the erasure patterns one can hope to correct are precisely those which consist of up to a erasures per local group plus up to h additional erasures anywhere in the codeword.

---

[1] LRCs are also expanded as Locally Repairable Codes or Locally Recoverable Codes, eg. [16, 17, 10].

A *maximal recoverable* (MR) LRC is a *single* code that is capable of simultaneously correcting *all* such patterns. Thus, an MR code gives the most bang-for-the-buck for the price one pays for locality.

This notion was introduced in [2] motivated by applications to storage on solid-state devices, where it was called partial MDS codes. The terminology maximally recoverable codes was coined in [7], and the concept was more systematically studied in [7, 6]. By picking the coefficients of the heavy parity checks randomly, it is not hard to show the existence of MR LRCs over *very large* fields, of size exponential in h. An explicit construction over such large fields was also given in [7], which also proved that random codes *need* such large field sizes with high probability.[2]

Since encoding a linear code and decoding it from erasures involve performing numerous finite field arithmetic operations, it is highly desirable to have codes over small fields (preferably of characteristic 2). Obtaining MR LRCs over finite fields of minimal size has therefore emerged as a central problem in the area of codes for distributed storage. So far, no construction of MR LRCs that avoid the exponential dependence on h has been found. A recent lower bound shows that, unlike MDS codes, for certain parameter settings one cannot have MR LRCs over fields of linear size. This shows that the notion of maximal recoverability is quite subtle, and pinning down the optimal field size is likely a deep question. There remains a large gap between the upper and lower bounds on field size of MR LRCs, closing which is a challenge of theoretical and practical importance.

In this work, we develop a novel approach to construct MR LRCs based on function fields. Our framework recovers and in fact slightly improves most of the previous bounds in the literature in a unified way. We note that since there are at least three quantities of significance – the locality r, the local (intra group) erasure tolerance a, and number of global parity checks h – the landscape of parameters and different constructions in this area is quite complex. Also, depending on the motivation, the range of values of interest of these parameters might be different. For example, if extreme efficiency of local repair is important, r should be small. But on the other hand this increase the redundancy and thus storage requirement of the code, so from this perspective a modest r (say $\sqrt{n}$) might be relevant. If good global fault tolerance is required, we want larger h, but then the constructions have large field size. It is therefore of interest to study the problem treating these as independent parameters, without assumptions on their relative size. We next review the field size of previous constructions, and then turn to the parameters we achieve in different regimes.

1.1   Known field size bounds

For $a \in \{0, r - 1\}$, optimal maximally recoverable local reconstruction codes (MR LRCs, for short) can be constructed by using either Reed-Solomon codes or their repetition. For $h \leqslant 1$, constructions of MR LRCs over fields of size $O(r)$ were given in [2]. For the remaining case: $1 \leqslant a \leqslant r - 2$ and $h \geqslant 2$, there are quite a number of constructions in literature [1, 2, 3, 4, 5, 7, 6, 9, 11, 18].

For the cases of $h = 2$ and $h = 3$, the best known constructions of MR LRCs were given in [9] with field sizes of $O(n)$ and $O(n^3)$ respectively, uniformly for all r, a. (Their field sizes were worse by $n^{o(1)}$ factors compared to these bounds when the field is required to be of

---

[2] This is akin to what happens for random codes to have the MDS property. However, for MDS codes, the Vandermonde construction achieves a linear field size explicitly.

characteristic 2.) For most other parameter settings, the best constructions by [5] provide a family of MR LRCs over fields of sizes

$$\ell = O\left(r \cdot n^{(a+1)h-1}\right) \tag{1}$$

as well as

$$\ell = \max\left\{O(\frac{n}{r}),\; O(r)^{h+a}\right\}^h , \tag{2}$$

The bound (1) outperforms the bound (2) when $r = \Omega(n)$, while the bound (2) is better when $r \ll n$. In both the bounds, the field size grows exponentially with h and a.

Recently, by using maximum rank distance (MRD) codes, the paper [15] (specifically Corollary 14) gives a family of MR LRCs over fields of sizes

$$\ell = O\left(r^{\frac{n(r-a)}{r}}\right). \tag{3}$$

When $r = \Omega(n)$, and a is close to r or h is large, (3) is better than bounds (1) or (2). By using probabilistic arguments, the paper [15] shows existence of a family of MR LRCs over fields of sizes

$$\ell = O\left(\binom{n-1}{k-1}\right), \tag{4}$$

where $k = n\left(1 - \frac{a}{r}\right) - h$ is the dimension of the code.

On the other hand, a lower bound on the field size was presented in [9]. Stating the bound when $h \leqslant \frac{n}{r}$ for simplicity, they show that the field size ℓ of an $(n, r, h, a)_\ell$ MR LRC must obey

$$\ell = \Omega_{a,h}\left(n \cdot r^{\min\{a, h-2\}}\right). \tag{5}$$

The lower bound (5) is still quite far from the upper bounds (1) and (2). In particular, the exponent a or h is to the base growing with n in the known constructions, but only to the base r in the above lower bound. Thus, one can conjecture that there is still room to improve both the constructions and the lower bounds. We note that under more complex structural requirements on the local groups, notably grid-like topologies and product codes, the optimal field size has been pinned down to $\exp(\Theta(n))$ [13].

Several techniques have been employed in literature for constructions of MR LRCs. One prevalent idea is to use a "linearized" version of the Vandermonde matrix, where the heavy parity check part of the matrix consists of columns $(\alpha_i, \alpha_i^q, \ldots, \alpha_i^{q^{h-1}})^T$ where $\alpha_i \in \mathbb{F}_\ell$ for a sufficiently high degree extension field $\mathbb{F}_\ell$ of $\mathbb{F}_q$. This construction is combined with $2h$-wise independent spaces to get an $O(n^h)$ field size in [7], and is also employed in [5]. Another approach is based on rank-metric codes (see, for instance, [4, 15]). Various ad hoc methods have been employed for good constructions of MR LRCs for small h, for example for $h = 2, 3$ in [9].

1.2 Our results

In this work, we develop a new approach to construct MR LRCs based on algebraic function fields. We discuss the key elements underlying our strategy in Section 1.4, but for now state the field sizes of the MR LRCS we can construct for various regimes of parameters. Most of the existing results in literature can be recovered through our methods in a unified way.

In most regimes, the parameters of our codes beat the known ones. For easy reference, we summarize the different possible trade-offs we can achieve in one giant theorem statement below. Since this comprehensive statement may be overwhelming to parse, let us highlight just two of our significant improvements: item (i) for $a = 1$, where we improve r^{h+1} term in (2) quadratically to $r^{\lfloor \frac{h+1}{2} \rfloor}$, and item (vi) for sufficiently large h, where the exponent h in bounds (1) and (2) is improved to εh. Also the exponent h is replaced by $\min\{h, n/r\}$ in the bounds (i)-(iv) that improve (2). In the bounds (vii) and (viii) the factor n/r in the exponent is improved to $\min\{k, n/r\}$; this improvment is less significant as it only applies to the low-rate setting but included for completeness and also to reflect a construction approach based on generator matrices (as opposed to parity check matrices which is a more potent way to reason about MR LRCs that underlies the other parts of the theorem).

▶ **Theorem 1.** *One has a maximally recoverable $(n, r, h, a)_\ell$-local reconstruction code over a field of size ℓ with parameters satisfying any of the following conditions. (Below $\tilde{O}(f)$ denotes $f \log^{O(1)} f$.)*

(i) (see Theorem 10) $a = 1$,$r \geqslant h + 2$ *and*

$$\ell \leqslant \left(\max \left\{ \tilde{O}(\frac{n}{r}), (2r)^{\lfloor \frac{h+1}{2} \rfloor} \right\} \right)^{\min\{h, \frac{n}{r}\}} \quad \text{and } \ell \text{ is even;}$$

(ii) (see Theorem 11) $a = 1$ *and*

$$\ell \leqslant \left(\max \left\{ \tilde{O}(\frac{n}{r}), 2^r \right\} \right)^{\min\{h, \frac{n}{r}\}} \quad \text{and } \ell \text{ is even;}$$

(iii) (see Theorem 13) *for all settings of n, r, h, a and*

$$\ell \leqslant \left(\max \left\{ \tilde{O}(\frac{n}{r}), (2r)^{h+a} \right\} \right)^{\min\{h, \frac{n}{r}\}};$$

(iv) (see Theorem 14) *for all settings of n, r, h, a and*

$$\ell \leqslant \left(\max \left\{ \tilde{O}(\frac{n}{r}), (2r)^r \right\} \right)^{\min\{h, \frac{n}{r}\}};$$

(v) (see Theorem 17) $r = O\left(\frac{\log n}{\log \log n} \right)$ *and* $hr \geqslant \Omega\left(\frac{n^{\frac{2}{3}}}{\varepsilon} \right)$ *for a positive real $\varepsilon \in (0, 0.5)$ and*

$$\ell \leqslant O\left(n^{\frac{2h}{3}(1+\varepsilon)} \right);$$

(vi) (see Theorem 18) $r = O\left(\frac{\varepsilon \log n}{\log \log n} \right)$ *and* $hr = \Omega\left(n^{1-\varepsilon} \right)$ *for a positive real $\varepsilon \in (0, 0.5)$ and*

$$\ell \leqslant n^{\varepsilon h};$$

(vii) (see Theorem 5) *for all settings of n, r, h, a*

$$\ell \leqslant \begin{cases} 2^{\min\{rk, n\}} \leqslant 2^n & \text{if } r \geqslant \log n \\ 2^{\lceil \log n \rceil \min\{k, \frac{n}{r}\}} & \text{if } r \leqslant \log n \end{cases}$$

where $k = \left(1 - \frac{a}{r}\right) - h$ is the dimension of the code;

(viii) (see Theorem 7) $r - a = \Omega(\log n)$ *and*

$$\ell \leqslant 2r^{\lfloor \frac{r-a}{2} \rfloor \min\{k, \frac{n}{r}\}} \quad \text{and } \ell \text{ is even.}$$

The first two bounds, and the bounds in (vii) and (viii) of Theorem 1 are derived from the rational function field $\mathbb{F}_2(x)$. In addition, bounds in (i) and (viii) of Theorem 1 are obtained via a combination with binary BCH codes. Bounds in (iii) and (iv) of Theorem 1 are derived from rational function field $\mathbb{F}_q(x)$, where ℓ is a power of q. The fifth bound is obtained via Hermitian function fields, while the sixth bound is derived from the Garcia-Stichtenoth function field tower. Our codes achieving the trade-offs stated in the above theorem can in fact be explicitly specified. But we note that for MR codes even existence questions over small fields are interesting and non-trivial.

1.3   Comparison

Each of our bounds in Theorem 1 beats the known results in some parameter regimes. Let us compare them one by one.

- The bound in Theorem 1(i) outperforms the bound (2) due to the quadratically better exponent for r.
- The bound in Theorem 1(ii) outperforms even the bound in Theorem 1(i) for $\frac{r}{\log r} < \left\lfloor \frac{h+1}{2} \right\rfloor$.
- The bound in Theorem 1(iii) outperforms the bound (2) for $h > \frac{n}{r}$.
- The bound in Theorem 1(iv) even outperforms the bound in Theorem 1(iv) for $r < h + a$, and hence it beats the bound (2) for $\frac{n}{h} < r < h + a$.
- The bound in Theorem 1(v) outperforms both the bounds (1) and (2) for all parameter settings subject to $r = \widetilde{O}(\log n)$ and $hr = \Omega\left(\frac{n^{\frac{2}{3}}}{\varepsilon}\right)$. It is clear that the bound in Theorem 1(v) is better than (1). As $r = \widetilde{O}(\log n)$, then we have $\left(\frac{n}{r}\right)^h > n^{h(1-o(1))} > n^{2h(1+\varepsilon)/3}$ and hence the bound in Theorem 1(v) beats (2) in this case.
- As the bound in Theorem 1(vi) is even better than the bound in Theorem 1(v), the bound in Theorem 1(vi) beats both the bounds (1) and (2) for all parameter settings subject to $r = \widetilde{O}(\varepsilon \log n)$ and $hr = O\left(n^{1-\varepsilon}\right)$ for a positive real $\varepsilon \in (0, 0.5)$.
- When the dimension k is much smaller than n, then the probabilistic bound (4) gives the field size $O(n^k) = O(2^{k \log n})$ which is the same size as in Theorem 1(vii) for $r \leqslant \log n$. When the dimension k is proportional to n, then the probabilistic bound (4) gives the field size $2^{O(n)}$ which is the same as the bound 2^n in Theorem 1(vii) for $r \geqslant \log n$.
- Finally, the bound in Theorem 1(viii) clearly outperforms the bound (3) when $k < n/r$.

1.4   Our techniques

Note that construction of MR LRCs is equivalent to construction of certain generator or parity-check matrices with requirement of column linear independence (see Section 2.1).

Our construction idea departs from previous approaches and is based on function fields over a finite field $\mathbb{F}_q$. The key in constructing an MR LRC is the choice of the heavy parity checks. We now briefly describe our idea to pick these. We associate with each of the $g = n/r$ local groups a distinguishing (high degree) place P_i, $1 \leqslant i \leqslant g$. The degree of the place is chosen large enough to guarantee the existence of at least g such places. For each local group, we pick functions f_{ij}, $1 \leqslant j \leqslant r$, that have *exactly one pole at P_i*. The coefficients of the h heavy parity checks corresponding to the j'th symbol of i'th local group are chosen to be

$$(f_{ij}(Q), f_{ij}^q(Q), \ldots, f_{ij}^{q^{h-1}}(Q))^T , \tag{6}$$

where Q is a place of sufficiently high degree, so that the evaluations $f_{ij}(Q)$ belong to an extension field $\mathbb{F}_\ell$ which will be the final alphabet size of the MR LRC. By properties of the Moore determinant (Section 2.2) and the large degree of Q, the required linear independence

of columns such as (6) over $\mathbb{F}_\ell$ reduces to a certain linear independence requirement for the f_{ij}'s over $\mathbb{F}_q$. Across different local groups such linear independence follows because a function with one pole at P_i cannot cancel a function with one pole at a different place $P_{i'}$. Within a local group, the required linear independence is ensured by choosing the f_{ij}'s within a group so that any $h + a$ of them (which is the maximum number of erasures we can have within a group) are linearly independent over $\mathbb{F}_q$.

We remark that all our various guarantees of Theorem 1 except Parts (v) and (vi) are obtained using just the rational function field, and can be described in elementary language using just polynomials, as we do in Section 3.

1.5 Organization

The paper is organized as follows. In Section 2, we introduce some preliminaries such as MR LRCs (both the generator and parity check matrix viewpoints) and Moore determinants. In Section 3, we present our constructions of MR LRCs using the rational function field together with a concatenation with classical codes of good rate vs. distance trade-off. We give two constructions, using the generator matrix viewpoint in the first part (yielding Parts (vii) and (viii) of Theorem 1), and then a parity check based construction in the second part which yields Parts (i)-(iv) of Theorem 1. This section is elementary and only uses properties of polynomials. In Section 4, we generalize the construction of MR LRCs via parity-check matrix given in Section 3 by making use of arbitrary algebraic function fields. We then apply this construction to Hermitian function fields and the Garcia-Stichtenoth tower to obtain MR LRCs promised in Parts (v) and (vi) of Theorem 1 respectively.

2 Preliminaries

2.1 Maximally recoverable local reconstruction codes

Throughout this paper, $\mathbb{F}_q$ denotes the finite field of q elements for a prime power q. We use $\mathbb{F}_q^{k \times n}$ to denote the set of all $k \times n$ matrices over $\mathbb{F}_q$.

Consider a distributed storage system where there are g disjoint locality groups and each group has size r and can locally correct any a erasure errors. In addition, the system can correct any h erasure errors together with any a erasure errors in each group. This requires a class of codes called *maximally recoverable local reconstruction codes* or *partial MDS codes* for error correction of such a system. The precise definition of MR LRCs is given below.

▶ **Definition 1.** Let ℓ be a prime power and let a, g, r, h be positive integers satisfying $ga + h < gr$. Put $n = gr$ and $k = n - ga - h$. An ℓ-ary $[n, k]$-linear code with a generator matrix of the form

$$G = (B_1 | B_2 | \cdots | B_g) \in \mathbb{F}_\ell^{k \times n}$$

is called a maximally recoverable $(n, r, h, a)_\ell$-local reconstruction code (or an MR $(n, r, h, a)_\ell$-LRC, for short) if

(i) each B_i has size $k \times r$;

(ii) the row span of each B_i is an $[r, r - a, a + 1]_\ell$-MDS code for $1 \leqslant i \leqslant g$ (note that B_i is not a generator matrix of this MDS code in general);

(iii) after puncturing a columns from each B_i, the remaining matrix of G generates an $[n - ga, k, h + 1]_\ell$-MDS code.

From the definition, an MR $(n, r, h, a)_q$-LRC can correct h erasure errors at arbitrarily positions together with any a erasure errors in each of g groups. To see the recovery procedure for an MR $(n, r, h, a)_\ell$-LRC, we first recover h erasure errors (this can be done from (iii) of Definition 1). We can then correct a errors from each block by (ii) of Definition 1.

The following lemma directly follows from Definition 1.

▶ **Lemma 2.** *A matrix* $G = (B_1|B_2|\cdots|B_g) \in \mathbb{F}_\ell^{k \times n}$ *is a generator matrix of an MR* $(n, r, h, a)_\ell$-*LRC if and only if every* $k \times k$ *submatrix* S *of* G *with at most* $r - a$ *columns per block* B_i *is invertible.*

One can have an equivalent definition via parity-check matrix.

▶ **Definition 2.** Let ℓ be a prime power and let a, g, r, h be positive integers satisfying $ga + h < gr$. Put $n = gr$ and $k = n - ga - h$. An ℓ-ary $[n, k]$-linear code with a parity-check matrix of the form

$$H = \left(\begin{array}{c|c|c|c} A_1 & O & \cdots & O \\ \hline O & A_2 & \cdots & O \\ \hline \vdots & \vdots & \ddots & \vdots \\ \hline O & O & \cdots & A_g \\ \hline D_1 & D_2 & \cdots & D_g \end{array}\right) \in \mathbb{F}_\ell^{(n-k) \times n} \tag{7}$$

is called an MR $(n, r, h, a)_\ell$-LRC if
   (i) each A_i has size $a \times r$ and each D_i has size $h \times r$;
  (ii) each A_i generates an $[r, a, r - a + 1]_\ell$-MDS code for $1 \leqslant i \leqslant g$ (note that the nullspace of A_i is $[r, r - a, a + 1]_\ell$ code);
 (iii) every $ag + h$ columns consisting of any a columns in each group and other arbitrary h columns are $\mathbb{F}_\ell$-linearly independent.

▶ Remark 1.
   (i) To see equivalence between Definitions 1 and 2, we note that each A_i in Definition 2 is actually a parity-check matrix of the code generated by B_i given in Definition 1.
  (ii) In this paper, we will use both Definitions 1 and 2 for constructions of MR LRCs. However, the major results of this paper come from the constructions based one Definition 2, i.e., via parity-check matrices of the required form in (7).

2.2   Moore determinant

Let ℓ be a power of q. For elements $\alpha_1, \ldots, \alpha_h \in \mathbb{F}_\ell$, the Moore matrix is defined by

$$M = \begin{pmatrix} \alpha_1 & \alpha_2 & \cdots & \alpha_h \\ \alpha_1^q & \alpha_2^q & \cdots & \alpha_h^q \\ \vdots & \vdots & \ddots & \vdots \\ \alpha_1^{q^{h-1}} & \alpha_2^{q^{h-1}} & \cdots & \alpha_h^{q^{h-1}} \end{pmatrix} \in \mathbb{F}_\ell^{h \times h}.$$

The determinant $\det(M)$ is given by the following formula

$$\det(M) = \prod_{(c_1, \ldots, c_h)} (c_1 \alpha_1 + \cdots + c_h \alpha_h),$$

where $(c_1, \ldots, c_h)$ runs through all non-zero direction vectors in $\mathbb{F}_q^h$. Thus, $\det(M) \neq 0$ if and only if $\alpha_1, \ldots, \alpha_h$ are $\mathbb{F}_q$-linearly independent.

3 Explicit constructions via rational function fields

In this section, we only introduce constructions of MR LRCs from rational function fields. Our description will be self-contained and elementary in terms of polynomials and we don't require any background on algebraic function fields (we have therefore deferred the background on function fields to Section 4 ahead of our more general construction in the next section).

3.1 Constructions via generator matrix

In this subsection, we present constructions of MR LRCs using Definition 1, i.e., via generator matrices of MR LRCs.

Let $N_q(d)$ denote the number of monic irreducible polynomials of degree d over $\mathbb{F}_q$. Then one has $\sum_{d|m} dN_q(d) = q^m$ for any $m \geqslant 1$ (see [14, Corollary 3.21 of Chapter 3]). This gives $\sum_{d|m} N_q(d) \geqslant \frac{q^m}{m}$. For each monic irreducible polynomial $p(x)$ of degree d with $d|m$, we get a polynomial $p(x)^{m/d}$ of degree m. Thus, for any $g \leqslant \left\lceil \frac{q^m}{m} \right\rceil$, there are g polynomials $p_1(x), p_2(x), \ldots, p_g(x)$ of degree m such that $\gcd(p_i(x), p_j(x)) = 1$ for all $1 \leqslant i \neq j \leqslant g$

Assume that (i) $m \geqslant r$; or (ii) $m < r$ and there is a q-ary $[r, r-m, \geqslant r-a+1]$-linear code, i.e. there exists a subset of $\mathbb{F}_q^m$ of size r such that any $r-a$ elements in this subset are $\mathbb{F}_q$-linearly independent.

Choose $g \leqslant \left\lceil \frac{q^m}{m} \right\rceil$ polynomials $p_1(x), p_2(x), \ldots, p_g(x)$ of degree m such that $\gcd(p_i(x), p_j(x)) = 1$ for all $1 \leqslant i \neq j \leqslant g$. Then for each $1 \leqslant i \leqslant g$, we can form an $\mathbb{F}_q$-vector space $V_i := \left\{ \frac{f(x)}{p_i(x)} : f(x) \in \mathbb{F}_q[x], \deg(f(x)) \leqslant m-1 \right\}$ of dimension m. As there is a q-ary $[r, r-m, \geqslant r-a+1]$-linear code, its parity-check matrix is an $r \times m$ matrix and any r columns of this matrix are linearly independent. This implies that one can find r functions $g_{i1}(x), \ldots, g_{ir}(x) \in V_i$ such that any $r-a$ polynomials out of $\{g_{i1}(x), \ldots, g_{ir}(x)\}$ are $\mathbb{F}_q$-linearly independent. Choose an irreducible polynomial $Q(x) \in \mathbb{F}_q[x]$ such that $Q(x)$ is coprime with every $p_i(x)$ for $1 \leqslant i \leqslant g$. For a function $h(x) \in V_i$, we use $h(Q)$ to denote the residue class of $h(x)$ in the residue class field $\mathbb{F}_q[x]/Q(x) \simeq \mathbb{F}_{q^{\deg(Q)}}$.

▶ **Lemma 3.** *Let T be a subset $\{1, 2, \ldots, g\}$ with $|T| \leqslant \deg(Q)/m$. If $\sum_{i \in T} g_i(Q) = 0$ for some functions $g_i \in V_i$, then $g_i = 0$ for all $i \in T$.*

Proof. Write $g_i = \frac{f_i}{p_i}$ for some polynomials f_i with $\deg(f_i) \leqslant m-1$. The equality $\sum_{i \in T} g_i(Q) = 0$ implies that $\sum_{i \in T} f_i(x) \prod_{j \in T \setminus \{i\}} p_j(x)$ is divisible by $Q(x)$. As the degree of $\sum_{i \in T} f_i(x) \prod_{j \in T \setminus \{i\}} p_j(x))$ is less than $m|T|$, we must have that $\sum_{i \in T} f_i(x) \prod_{j \in T \setminus \{i\}} p_j(x)$ is the zero polynomial. Suppose that $f_t \neq 0$ for some $t \in T$, then we have

$$\sum_{i \in T \setminus \{t\}} f_i(x) \prod_{j \in T \setminus \{i\}} p_j(x) = -f_t(x) \prod_{j \in T \setminus \{t\}} p_j(x) .$$

The l.h.s. of the above equality is divisible by $p_t(x)$, while the r.h.s. is not divisible by $p_t(x)$. This contradiction completes the proof. ◀

Let Q be an irreducible polynomial in $\mathbb{F}_q[x]$ of degree

$$\min\{km, gm\} = \min\left\{km, \frac{nm}{r}\right\} = \min\left\{(n - \frac{an}{r} - h)m, \frac{nm}{r}\right\}.$$

Define the $k \times r$ matrix B_i as follows.

$$
B_i = \begin{pmatrix}
g_{i1}(Q) & g_{i2}(Q) & \cdots & g_{ir}(Q) \\
g_{i1}^q(Q) & g_{i2}^q(Q) & \cdots & g_{ir}^q(Q) \\
\vdots & \vdots & \vdots & \vdots \\
g_{i1}^{q^{k-1}}(Q) & g_{i2}^{q^{k-1}}(Q) & \cdots & g_{ir}^{q^{k-1}}(Q)
\end{pmatrix} \in \mathbb{F}_{q^{\deg(Q)}}^{k \times r}. \tag{8}
$$

The proofs of the remaining results of this Section can be found in the full version of the paper that is available at https://arxiv.org/abs/1808.04539.

▶ **Lemma 4.** *Assume that $m \geqslant r$ or there is a q-ary $[r, r - m, \geqslant r - a + 1]$-linear code. Let B_i be the matrix given in* (8). *Put $\ell = q^{\min\left\{(n - \frac{an}{r} - h)m, \frac{nm}{r}\right\}} = q^{\min\left\{km, \frac{nm}{r}\right\}}$ and $G = (B_1|B_2|\cdots|B_g) \in \mathbb{F}_\ell^{k \times n}$. Then the ℓ-ary code C with the generator matrix G is an MR $(n, r, h, a)_\ell$-LRC.*

By taking $m = r$, we obtain the following result.

▶ **Theorem 5.** *If $r \geqslant \log n$, then there exists an MR (n, r, h, a)-LRC of dimension $k = n - \frac{na}{r} - h$ over a field of size*

$$
\ell \leqslant \begin{cases}
2^{\min\{rk, n\}} \leqslant 2^n & \text{if } r \geqslant \log n \\
2^{\min\left\{k\lceil \log n \rceil, \frac{n}{r}\lceil \log n \rceil\right\}} & \text{if } r \leqslant \log n
\end{cases}
$$

By considering binary BCH codes, we obtain the following binary codes.

▶ **Lemma 6.** *There exists a binary $[r, r - m, \geqslant d]$-linear code with $m = \lfloor \frac{d-1}{2} \rfloor \cdot \lceil \log_2 r \rceil + 1$.*

Combining the binary BCH codes of Lemma 6 with Lemma 4 applied with rational function field $\mathbb{F}_2(x)$ yields the following theorem.

▶ **Theorem 7.** *If $r - a = \Omega(\log n)$, then there exists an MR (n, r, h, a)-LRC of dimension $k = n - \frac{na}{r} - h$ over a field of size*

$$
\ell \leqslant 2r^{\min\left\{k\lfloor \frac{r-a}{2} \rfloor, \frac{n}{r}\lfloor \frac{r-a}{2} \rfloor\right\}} \leqslant 2r^{\frac{n}{r}\lfloor \frac{r-a}{2} \rfloor}.
$$

3.2   Constructions via parity-check matrix

To construct parity-check matrices of MR LRCs, we only need to construct matrices D_i given in (7). The idea of constructing matrices D_i is quite similar to that of constructing matrices B_i in the previous subsection, and leads to the following theorem.

▶ **Theorem 8.** *Let r, g, a, h, m be positive integers with $a \leqslant r$. Suppose that $q \geqslant r$ is a prime power satisfying $q^m \geqslant \frac{mn}{r}$ and there is a q-ary $[r, r - a, a + 1]$-linear code. If (i) $m \geqslant r$; or (ii) $m < r$ and there exists a q-ary $[r, r - m, \geqslant h + a + 1]$-linear code, then there exists an MR (n, r, h, a)-LRC with $n = rg$ over a field of size $\ell = q^{\min\{hm, \frac{nm}{r}\}}$.*

We now instantiate Theorem 8 with suitable choices of parameters to deduce the promises parts (i)–(iv) of Theorem 1.

3.2.1   The case where $a = 1$

Let $r, h \geqslant 2$ be integers. Then there is a q-ary $[r, 1, r]$-MDS code for any prime power q. Rewriting Theorem 8 for $a = 1$ gives the following lemma.

▶ **Lemma 9.** *Suppose that $q^m \geqslant \frac{mn}{r}$. If (i) $m \geqslant r$; or (ii) $m < r$ and there exists a q-ary $[r, r - m, \geqslant h + 2]$-linear code, then there exists an MR $(n, r, h, 1)$-LRC over a field of size $\ell = q^{\min\{hm, \frac{nm}{r}\}}$.*

To apply Lemma 9, we need to find suitable codes and function fields as well. By taking the rational function field $\mathbb{F}_2(x)$ and applying BCH code given in Lemma 6, we obtain the following result.

▶ **Theorem 10.** *If $r \geqslant h + 2$, then there exists an MR $(n, r, h, 1)$-LRC over a field of size*

$$\ell \leqslant \left(\max\left\{\tilde{O}(\frac{n}{r}), (2r)^{\lfloor \frac{h+1}{2} \rfloor}\right\}\right)^{\min\{h, \frac{n}{r}\}} .$$

Proof. Consider the rational function field $F = \mathbb{F}_2(x)$. Put

$$m = \max\left\{\left\lfloor \frac{h+1}{2} \right\rfloor \cdot \lceil \log_2 r \rceil + 1, \left\lceil \log_2\left(\frac{n}{r}\right) + 2\log_2\log_2\left(\frac{n}{r}\right) \right\rceil\right\}.$$

Then $\frac{n}{r} \leqslant \frac{1}{m} 2^m$. This implies that there are $\frac{n}{r}$ places of degree m in $\mathbb{F}_2(x)$. By Lemma 6, there exists a binary $[r, r - m, \geqslant h + 2]$-linear code. It follows from Lemma 9 that there exists an MR $(n, r, h, 1)$-LRC over a field of size $2^{\min\{mh, m\frac{n}{r}\}}$. By choice of our parameters, the desired result follows. ◀

▶ **Theorem 11.** *There exists an MR $(n, r, h, 1)$-LRC over a field of size*

$$\ell \leqslant \left(\max\left\{\tilde{O}(\frac{n}{r}), 2^r\right\}\right)^{\min\{h, \frac{n}{r}\}} .$$

Proof. Consider the rational function field $\mathbb{F}_2(x)$. Put $m = \max\{r, \left\lceil \log_2\left(\frac{n}{r}\right) + 2\log_2\log_2\left(\frac{n}{r}\right) \right\rceil\}$. Then $\frac{n}{r} \leqslant \frac{1}{m} 2^m$. The desired result follows from Lemma 9. ◀

▶ **Remark 2.** Theorem 11 gives a better bound on the field size than Theorem 10 for $h > \frac{2r}{\log_2 r} - 1$, while Theorem 10 gives a better bound on the field size than Theorem 11 for $h < \frac{2r}{\log_2 r} - 1$.

3.2.2 The case where $2 \leqslant a \leqslant r - 1$

▶ **Lemma 12.** *Let $a \leqslant r \leqslant q + 1$ and $m \geqslant h + a$. If $q^m \geqslant \frac{mn}{r}$, then there exists an MR (n, r, h, a)-LRC code over a field of size $\ell = q^{\min\{mh, \frac{mn}{r}\}}$.*

Proof. When $a \leqslant r \leqslant q + 1$ and $m \geqslant h + a$, we have an $[r, r - a, a + 1]_q$-MDS code and an $[r, r - m, h + a + 1]_q$-linear code. The result thus follows from Theorem 8. ◀

▶ **Theorem 13.** *There exists an MR (n, r, h, a)-LRC over a field of size*

$$\ell \leqslant \left(\max\left\{\tilde{O}(\frac{n}{r}), (2r)^{h+a}\right\}\right)^{\min\{h, \frac{n}{r}\}} .$$

Proof. Let q be the smallest prime power such that $q - 1 \geqslant r$. We may take q to be a power of two, so that $q \leqslant 2r$. Consider the rational function field $F = \mathbb{F}_q(x)$ and let

$$m = \max\left\{h + a, \left\lceil \log_q\left(\frac{n}{r}\right) + 2\log_q\log_q\left(\frac{n}{r}\right) \right\rceil\right\}.$$

Then $\frac{n}{r} \leqslant \frac{1}{m} q^m$. The desired result follows from Theorem 8. ◀

▶ **Remark 3.** The field size $\ell \leqslant \tilde{O}\left(\max\left\{\frac{n}{r}, r^{h+a}\right\}^h\right)$ in Theorem 13 was already given in [5, Corollary 11]. Here we provide a better result for $h > \frac{n}{r}$ via a different approach.

▶ **Theorem 14.** *There exists an MR (n, r, h, a)-LRC over a field of size*

$$\ell \leqslant \left(\max\left\{\tilde{O}(\frac{n}{r}), (2r)^r\right\}\right)^{\min\{h, \frac{n}{r}\}} .$$

Proof. Put $q = 2^{\lceil \log_2 r \rceil}$. Then $2r \geqslant q \geqslant r$ and hence we have a q-ary $[r, a]$-MDS code for any $a \leqslant r$. Put $m = \max\{r, \lceil \log_q\left(\frac{n}{r}\right) + 2\log_q \log_q\left(\frac{n}{r}\right) \rceil\}$. Then $\frac{n}{r} \leqslant \frac{1}{m}q^m$. The desired result follows from Theorem 10. ◀

▶ **Remark 4.** Theorem 14 gives a better bound on the field size than Theorem 13 for $h + a > r$, while Theorem 13 gives a better bound on the field size than Theorem 14 for $h + a < r$.

4 Explicit construction via general function fields

The construction via rational function fields given in Section 3 can be easily generalized to arbitrary function fields. We only generalize the constructions of MR LRCs via parity-check matrices given in Section 3.2. The necessary background on algebraic function fields, and specifically Hermitian and Garcia-Stichtenoth tower of function fields, can be found in the full version of this paper. We refer the proofs in this section to the full version of this paper.

Let q be a prime power and let a, r, h, g be integers with $a \leqslant r \leqslant q + 1$. Let $F/\mathbb{F}_q$ be a function field of genus $\mathfrak{g}$. Let $P_1, P_2, \ldots, P_g$ be g positive divisors of degree r whose supports are pairwise disjoint. Let G be a divisor of degree $2\mathfrak{g} - 1$. By Riemann-Roch, $\dim \mathcal{L}(G) = \mathfrak{g}$. Assume that $\{f_1, f_2, \ldots, f_\mathfrak{g}\}$ is a basis of $\mathcal{L}(G)$. For each i, extend this basis to a basis $\{f_1, f_2, \ldots, f_\mathfrak{g}, f_{i1}, f_{i2}, \ldots, f_{ir}\}$ of $\mathcal{L}(G + P_i)$.

Let Q be a place of degree $2\mathfrak{g} + \min\{hr, n\}$ and define the matrix

$$D_i = \begin{pmatrix} f_{i1}(Q) & f_{i2}(Q) & \cdots & f_{ir}(Q) \\ f_{i1}^q(Q) & f_{i2}^q(Q) & \cdots & f_{ir}^q(Q) \\ \vdots & \vdots & \vdots & \vdots \\ f_{i1}^{q^{h-1}}(Q) & f_{i2}^{q^{h-1}}(Q) & \cdots & f_{ir}^{q^{h-1}}(Q) \end{pmatrix} \tag{9}$$

By mimicking the proof of Theorem 8, we have the following result.

▶ **Lemma 15.** *Let $A_i \in \mathbb{F}_q^{a \times r}$ be a generator matrix of an $[r, a]_q$-MDS code for $1 \leqslant i \leqslant g$. Let D_i be the matrix given in (9). Put $\ell = q^{2\mathfrak{g} + \min\{hr, n\}}$. Then the ℓ-ary code C with the matrix H defined in (7) is an MR $(n, r, h, a)_\ell$-LRC.*

Consequently, we have the following theorem.

▶ **Theorem 16.** *Let r, g, a, h be positive integers with $a \leqslant r \leqslant q + 1$. If there is a function field $F/\mathbb{F}_q$ of genus $\mathfrak{g}$ with g positive divisors of degree r whose supports are disjoint, then there exists an MR (n, r, h, a)-LRC with $n = rg$ over a field of size $\ell = q^{2\mathfrak{g} + \min\{hr, n\}}$.*

Finally, let us instantiate the above result with the Hermitian function fields and the Garcia-Stichtenoth tower, to deduce Parts (v) and (vi) promised in Theorem 1 respectively. Note that both the results below kick-in for block lengths which are asymptotically at least $r^{O(r)}$, which is why we have the condition $r \leqslant O(\frac{\log n}{\log \log n})$ in the statement of Theorem 1, Parts (v), (vi).

▶ **Theorem 17.** *Let $a \leqslant r$ be integers. Then there are infinitely many $n \geqslant r^{\Omega(r)}$ such that there is MR (n, r, h, a)-LRC over a field of size at most $n^{\frac{2h}{3}(1+\varepsilon)}$ for any desired $\varepsilon \in (0, 0.5)$ provided $hr \geqslant \Omega\left(\frac{n^{\frac{2}{3}}}{\varepsilon}\right)$.*

We finally state a similar result using the Garcia-Stichtenoth tower of function fields.

▶ **Theorem 18.** *Let $a \leqslant r$ be positive integers and let $\varepsilon \in (0, 0.5)$. Then there are infinitely many $n \geqslant r^{\Omega(r/\varepsilon)}$ such that there is MR (n, r, h, a)-LRC over a field of size at most $n^{\varepsilon h}$ provided $hr \geqslant \Omega\left(n^{1-\varepsilon}\right)$.*

─── **References** ───

1 Mario Blaum. Construction of PMDS and SD Codes extending RAID 5. *arXiv preprint arXiv:1305.0032*, 2013. `arXiv:1305.0032`.

2 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.

3 Mario Blaum, James S Plank, Moshe Schwartz, and Eitan Yaakobi. Construction of partial MDS and sector-disk codes with two global parity symbols. *IEEE Transactions on Information Theory*, 62(5):2673–2681, 2016.

4 Gokhan Calis and O Ozan Koyluoglu. A general construction for PMDS codes. *IEEE Communications Letters*, 21(3):452–455, 2017.

5 Ryan Gabrys, Eitan Yaakobi, Mario Blaum, and Paul H Siegel. Constructions of partial MDS codes over small fields. In *2017 IEEE International Symposium on Information Theory (ISIT)*, pages 1–5. IEEE, 2017.

6 Parikshit Gopalan, Guangda Hu, Swastik Kopparty, Shubhangi Saraf, Carol Wang, and Sergey Yekhanin. Maximally recoverable codes for grid-like topologies. In *28th Annual Symposium on Discrete Algorithms (SODA)*, pages 2092–2108. Society for Industrial and Applied Mathematics, 2017.

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

8 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.

9 Sivakanth Gopi, Venkatesan Guruswami, and Sergey Yekhanin. On maximally recoverable local reconstruction codes. *arXiv preprint arXiv:1710.10322*, 2017.

10 Venkatesan Guruswami, Chaoping Xing, and Chen Yuan. How long can optimal locally repairable codes be? In *Proceedings of RANDOM 2018*, pages 41:1–41:11, 2018.

11 Guangda Hu and Sergey Yekhanin. New constructions of SD and MR codes over small finite fields. In *2016 IEEE International Symposium on Information Theory (ISIT)*, pages 1591–1595. IEEE, 2016.

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

13 Daniel Kane, Shachar Lovett, and Sankeerth Rao. The independence number of the birkhoff polytope graph, and applications to maximally recoverable codes. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 252–259. IEEE, 2017.

14 Rudolf Lidl and Harald Niederreiter. *Finite fields*, volume 20. Cambridge university press, 2003.

15 Alessandro Neri and Anna-Lena Horlemann-Trautmann. Random Construction of Partial MDS Codes. *arXiv preprint arXiv:1801.05848*, 2018.

16 Dimitris S Papailiopoulos and Alexandros G Dimakis. Locally repairable codes. *IEEE Transactions on Information Theory*, 60(10):5843–5855, 2014.

17   Itzhak Tamo and Alexander Barg. A family of optimal locally recoverable codes. *IEEE Transactions on Information Theory*, 60(8):4661–4676, 2014.

18   Itzhak Tamo, Dimitris S Papailiopoulos, and Alexandros G Dimakis. Optimal locally repairable codes and connections to matroid theory. *IEEE Transactions on Information Theory*, 62(12):6661–6671, 2016.