

On the Locality of Codeword Symbols

Parikshit Gopalan
Microsoft Research
parik@microsoft.com

Cheng Huang
Microsoft Research
chengh@microsoft.com

Huseyin Simitci
Microsoft Corporation
huseyins@microsoft.com

Sergey Yekhanin
Microsoft Research
yekhanin@microsoft.com

Abstract

Consider a linear $[n, k, d]_q$ code \mathcal{C} . We say that the i -th coordinate of \mathcal{C} has locality r , if the value at this coordinate can be recovered from accessing some other r coordinates of \mathcal{C} . Data storage applications require codes with small redundancy, low locality for information coordinates, large distance, and low locality for parity coordinates. In this paper we carry out an in-depth study of the relations between these parameters.

We establish a tight bound for the redundancy $n - k$ in terms of the message length, the distance, and the locality of information coordinates. We refer to codes attaining the bound as optimal. We prove some structure theorems about optimal codes, which are particularly strong for small distances. This gives a fairly complete picture of the tradeoffs between codewords length, worst-case distance and locality of information symbols.

We then consider the locality of parity check symbols and erasure correction beyond worst case distance for optimal codes. Using our structure theorem, we obtain a tight bound for the locality of parity symbols possible in such codes for a broad class of parameter settings. We prove that there is a tradeoff between having good locality for parity checks and the ability to correct erasures beyond the minimum distance.

1 Introduction

Modern large scale distributed storage systems such as data centers store data in a redundant form to ensure reliability against node (e.g., individual machine) failures. The simplest solution here is the straightforward replication of data packets across different nodes. Alternative solution involves erasure coding: the data is partitioned into k information packets. Subsequently, using an erasure code, $n - k$ parity packets are generated and all n packets are stored in different nodes.

Using erasures codes instead of replication may lead to dramatic improvements both in terms of redundancy and reliability. However to realize these improvements one has to address the challenge of maintaining an erasure encoded representation. In particular, when a node storing some packet fails, one has to be able to quickly reconstruct the lost packet in order to keep the data readily available for the users and to maintain the same level of redundancy in the system. We say that a certain packet has *locality* r if it can be recovered from accessing only r other packets. One way to ensure fast reconstruction is to use erasure codes where all packets have low locality $r \ll k$. Having small value of locality is particularly important for information packets.

These considerations lead us to introduce the concept of an (r, d) -code, i.e., a linear code of distance d , where all information symbols have locality at most r . Storage system based on (r, d) -codes

provide fast recovery of information packets from a single node failure (typical scenario), and ensure that no data is lost even if up to $d - 1$ nodes fail simultaneously. One specific class of (r, d) -codes called Pyramid Codes has been considered in [6].

Pyramid codes can be obtained from any systematic Maximum Distance Separable (MDS) codes of distance d , such as Reed-Solomon codes. Assume for simplicity that the first parity check symbol is the sum $\sum_{i=1}^k x_i$ of the information symbols. Replace this with $\lceil \frac{k}{r} \rceil$ parity checks each of size at most r on disjoint information symbols. It is not hard to see that the resulting code \mathcal{C} has information locality r and distance d , while the redundancy of the code \mathcal{C} is given by

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

1.1 Our results

In this paper we carry out an in-depth study of the relations between redundancy, erasure-correction and symbol locality in linear codes.

Our first result is a tight bound for the redundancy in terms of the message length, the distance, and the information locality. We show that in any $[n, k, d]_q$ code of information locality r ,

$$n - k \geq \left\lceil \frac{k}{r} \right\rceil + d - 2. \quad (2)$$

We refer to codes attaining the bound above as optimal. Pyramid codes are one such family of codes. The bound (2) is of particular interest in the case when $r \mid k$, since otherwise one can improve the code by increasing the dimension while maintaining the (r, d) -property and redundancy intact. A closer examination of our lower bound gives a structure theorem for optimal codes when $r \mid k$. This theorem is especially strong when $d < r + 3$, it fixes the support of the parity check matrix, the only freedom is in the choice of coefficients. We also show that the condition $d < r + 3$ is in fact necessary for such a strong statement to hold.

We then turn our attention to the locality of parity symbols. We prove tight bounds on the locality of parity symbols in optimal codes assuming $d < r + 3$. In particular we establish the existence of optimal (r, d) -codes that are significantly better than Pyramid codes with respect to locality of parity symbols. Our codes are explicit in the case of $d = 4$, and non-explicit otherwise. The lower bound is proved using the structure theorem. Finally, we relax the conditions $d < r + 3$ and $r \mid k$ and exhibit one specific family of optimal codes that gives locality r for all symbols.

Our last result concerns erasure correction beyond the worst case distance of the code. Assume that we are given a bipartite graph which describes the supports of the parity check symbols. What choice of coefficients will maximize the set of erasure patterns that can be corrected by such a code? In [2, 6] the authors gave a necessary condition for an erasure pattern to be correctable, and showed that over sufficiently large fields, this condition is also sufficient. They called such codes maximally recoverable. We show that such codes cannot have any non-trivial parity locality; thus establishing a tradeoff between parity locality and erasure correction beyond the worst case distance.

1.2 Related work

There are two classes of erasure codes providing fast recovery procedures for individual codeword coordinates (packets) in the literature.

Regenerating codes. These codes were introduced in [3] and developed further in e.g., [10, 1, 11, 13]. See [4] for a survey. One crucial idea behind regenerating codes is that of *sub-packetization*. Each packet is composed of few sub-packets, and when a node storing a packet fails all (or most of) other nodes send in some of their sub-packets for recovery. Efficiency of the recovery procedure is measured in terms of the overall bandwidth consumption, i.e., the total size of sub-packets required to recover from a single failure. Somehow surprisingly regenerating codes can in many cases achieve a rather significant reduction in bandwidth, compared with codes that do not employ sub-packetization. Our experience with data centers however suggests that in practice there is a considerable overhead related to accessing extra storage nodes. Therefore pure bandwidth consumption is not necessarily the right single measure of the recovery time. In particular, coding solutions that do not rely on sub-packetization and thus access less nodes (but download more data) are sometimes more attractive.

Locally decodable codes. These codes were introduced in [8] and developed further in e.g., [14, 5, 9]. See [15] for a survey. An r -query Locally Decodable Code (LDC) encodes messages in such a way that one can recover any message symbol by accessing only r codeword symbols even after some arbitrarily chosen (say) 10% of codeword coordinates are erased. Thus LDCs are in fact very similar to (r, d) -codes addressed in the current paper, with an important distinction that LDCs allow for local recovery even after a very large number of symbols are erased, while (r, d) -codes provide locality only after a single erasure. Not surprisingly locally decodable codes require substantially larger codeword lengths than (r, d) -codes.

1.3 Organization

In Section 3 we establish the lower bound for redundancy of (r, d) -codes and obtain a structural characterization of optimal codes, i.e., codes attaining the bound. In Section 4 we strengthen the structural characterization for optimal codes with $d < r + 3$ and show that any such code has to be a canonical code. In Section 5 we prove matching lower and upper bounds on the locality of parity symbols in canonical codes. Our code construction is not explicit and requires the underlying field to be fairly large. In the special case of codes of distance $d = 4$, we come up with an explicit family that does not need a large field. In Section 6 we present one optimal family of non-canonical codes that gives uniform locality for all codeword symbols. Finally, in Section 7 we study erasure correction beyond the worst case distance and prove that systematic codes correcting the maximal number of erasure patterns (conditioned on the support structure of the generator matrix) cannot have any non-trivial locality for parity symbols.

2 Preliminaries

We use standard mathematical notation

- For an integer t , $[t] = \{1, \dots, t\}$;
- For a vector \mathbf{x} , $\text{Supp}(\mathbf{x})$ denotes the set $\{i : \mathbf{x}_i \neq 0\}$;
- For a vector \mathbf{x} , $\text{wt}(\mathbf{x}) = |\text{Supp}(\mathbf{x})|$ denotes the Hamming weight;
- For a vector \mathbf{x} and an integer i , $\mathbf{x}(i)$ denotes the i -th coordinate of \mathbf{x} ;
- For sets A and B , $A \sqcup B$ denotes the disjoint union;
- For vectors \mathbf{p} and \mathbf{x} , $\mathbf{p} \cdot \mathbf{x}$ denotes the dot product.

Let \mathcal{C} be an $[n, k, d]_q$ linear code. Assume that the encoding of $\mathbf{x} \in \mathbb{F}_q^k$ is by the vector

$$\mathcal{C}(\mathbf{x}) = (\mathbf{c}_1 \cdot \mathbf{x}, \mathbf{c}_2 \cdot \mathbf{x}, \dots, \mathbf{c}_n \cdot \mathbf{x}) \in \mathbb{F}_q^n. \quad (3)$$

Thus the code \mathcal{C} is specified by the set of n points $C = \{\mathbf{c}_1, \dots, \mathbf{c}_n\} \in \mathbb{F}_q^k$. The set of points must have full rank for \mathcal{C} to have k information symbols. It is well known that the distance property is captured by the following condition (e.g., [12, theorem 1.1.6]).

Fact 1. *The code \mathcal{C} has distance d if and only if for every $S \subseteq C$ such that $\text{Rank}(S) \leq k - 1$,*

$$|S| \leq n - d. \quad (4)$$

In other words, every hyperplane through the origin misses at least d points from C . In this work, we are interested in the recovery cost of each symbol in the code from a single erasure.

Definition 2. *For $\mathbf{c}_i \in C$, we define $\text{Loc}(\mathbf{c}_i)$ to be the smallest integer r for which there exists $R \subseteq C$ of cardinality r such that*

$$\mathbf{c}_i = \sum_{j \in R} \lambda_j \mathbf{c}_j.$$

We further define $\text{Loc}(C) = \max_{i \in [n]} \text{Loc}(\mathbf{c}_i)$.

Note that $\text{Loc}(\mathbf{c}_i) \leq k$, provided $d \geq 2$, since this guarantees that $C \setminus \{\mathbf{c}_i\}$ has full dimension. We will be interested in (systematic) codes which guarantee locality for the information symbols.

Definition 3. *We say that a code \mathcal{C} has information locality r if there exists $I \subseteq C$ of full rank such that $\text{Loc}(\mathbf{c}) \leq r$ for all $\mathbf{c} \in I$.*

For such a code we can choose I as our basis for \mathbb{F}_q^k and partition C into $I = \{\mathbf{e}_1, \dots, \mathbf{e}_k\}$ corresponding to information symbols and $C \setminus I = \{\mathbf{c}_{k+1}, \dots, \mathbf{c}_n\}$ corresponding to parity check symbols. Thus the code \mathcal{C} can be made systematic.

Definition 4. *A code \mathcal{C} is an (r, d) -code if it has information locality r and distance d .*

For any code \mathcal{C} , the set of all linear dependencies of length at most $r + 1$ on points in C defines a natural hypergraph $H_r(V, E)$ whose vertex set $V = [n]$ is in one-to-one correspondence with points in C . There is an edge corresponding to set a $S \subseteq V$ if $|S| \leq r + 1$ and

$$\sum_{i \in S} \lambda_i \mathbf{c}_i = 0, \quad \lambda_i \neq 0.$$

Equivalently $S \subseteq [n]$ is an edge in H if it supports a codeword in \mathcal{C}^\perp of weight at most $r + 1$. Since r will usually be clear from the context, we will just say $H(V, E)$. If no vertex in H is isolated then all coordinates of \mathcal{C} have locality r . A code \mathcal{C} has information locality r if the set points corresponding to vertices that are incident to some edge in H has full rank.

We conclude this section presenting one specific class of (r, d) -codes that has been considered in [6]:

Pyramid codes. To define an (r, d) -Pyramid code \mathcal{C} encoding messages of dimension k we fix an arbitrary linear systematic $[k + d - 1, k, d]_q$ code \mathcal{E} . Clearly, \mathcal{E} is MDS. Let

$$\mathcal{E}(\mathbf{x}) = (\mathbf{x}, \mathbf{p}_0 \cdot \mathbf{x}, \mathbf{p}_1 \cdot \mathbf{x}, \dots, \mathbf{p}_{d-2} \cdot \mathbf{x}).$$

We partition the set $[k]$ into $t = \lceil \frac{k}{r} \rceil$ disjoint subsets of size up to r , $[k] = \bigsqcup_{i \in [t]} S_i$. For a k -dimensional vector \mathbf{x} and a set $S \subseteq [k]$ let $\mathbf{x}|_S$ denote the $|S|$ -dimensional restriction of \mathbf{x} to coordinates in the set S . We define the systematic code \mathcal{C} by

$$\mathcal{C}(\mathbf{x}) = (\mathbf{x}, (\mathbf{p}_0|_{S_1} \cdot \mathbf{x}|_{S_1}), \dots, (\mathbf{p}_0|_{S_t} \cdot \mathbf{x}|_{S_t}), \mathbf{p}_1 \cdot \mathbf{x}, \dots, \mathbf{p}_{d-2} \cdot \mathbf{x}).$$

It is not hard to see that the code \mathcal{C} has distance d . To see that all information symbols and the first $\lceil \frac{k}{r} \rceil$ parity symbols of \mathcal{C} have locality r one needs to observe that (since \mathcal{E} is an MDS code) the vector \mathbf{p}_0 has full Hamming weight. The last $d - 2$ parity symbols of \mathcal{C} may have locality as large as k .

3 Lower bound and the structure theorem

We are interested in systematic codes with information locality r . Given k, r, d our goal is to minimize the codeword length n . Since the code is systematic, this amounts to minimizing the redundancy $h = n - k$. Pyramid codes have $h = \lceil \frac{k}{r} \rceil + d - 2$. Our goal is to prove a matching lower bound. Lower bounds of k/r and $d - 1$ are easy to show, just from the locality and distance constraints respectively. The hard part is to sum them up.

Theorem 5. *For any $[n, k, d]_q$ linear code with information locality r ,*

$$n - k \geq \left\lceil \frac{k}{r} \right\rceil + d - 2. \quad (5)$$

Proof. Our lower bound proceeds by constructing a large set $S \subseteq C$ where $\text{Rank}(S) \leq k - 1$ and then applying Fact 1. The set S is constructed by the following algorithm:

1. Let $i = 1, S_0 = \{\}$.
2. While $\text{Rank}(S_{i-1}) \leq k - 2$:
3. Pick $\mathbf{c}_i \in C \setminus S_{i-1}$ such that there is a hyperedge T_i in H containing \mathbf{c}_i .
4. If $\text{Rank}(S_{i-1} \cup T_i) < k$, set $S_i = S_{i-1} \cup T_i$.
5. Else pick $T' \subset T_i$ so that $\text{Rank}(S_{i-1} \cup T') = k - 1$ and set $S_i = S_{i-1} \cup T'$.
6. Increment i .

In Line 3, since $\text{Rank}(S_{i-1}) \leq k - 2$ and $\text{Rank}(I) = k$, there exists \mathbf{c}_i as desired. Let ℓ denote the number of times the set S_i has grown. Observe that the final set S_ℓ has $\text{Rank}(S_\ell) = k - 1$. We now lower bound $|S|$. We define s_i, t_i to measure the increase in the size and rank of S_i respectively:

$$s_i = |S_i| - |S_{i-1}|, \quad |S_\ell| = \sum_{i=1}^{\ell} s_i,$$

$$t_i = \text{Rank}(S_i) - \text{Rank}(S_{i-1}), \quad \text{Rank}(S_\ell) = \sum_{i=1}^{\ell} t_i = k - 1.$$

We analyze two cases, depending on whether the condition $\text{Rank}(S_{i-1} \cup T_i) = k$ is ever reached. Observe that this condition can only be reached when $i = \ell$.

Case 1: Assume $\text{Rank}(S_{i-1} \cup T_i) \leq k - 1$ throughout. In each step we add $s_i \leq r + 1$ vectors. Note that these vectors are always such that some nontrivial linear combination of them yields a (possibly zero) vector in $\text{Span}(S_{i-1})$. Therefore we have $t_i \leq s_i - 1 \leq r$. So there are $\ell \geq \lceil \frac{k-1}{r} \rceil$ steps in all. Thus

$$|S| = \sum_{i=1}^{\ell} s_i \geq \sum_{i=1}^{\ell} (t_i + 1) \geq k - 1 + \left\lceil \frac{k-1}{r} \right\rceil \quad (6)$$

Note that $k - 1 + \lceil \frac{k-1}{r} \rceil \geq k + \lceil \frac{k}{r} \rceil - 2$ with equality holding whenever $r = 1$ or $k \equiv 1 \pmod r$.

Case 2: In the last step, we hit the condition $\text{Rank}(S_{\ell-1} \cup T_\ell) = k$. Since the rank only increases by at most r per step, $\ell \geq \lceil \frac{k}{r} \rceil$. For $i \leq \ell - 1$, we add a set T_i of $s_i \leq r + 1$ vectors. Again note that these vectors are always such that some nontrivial linear combination of them yields a (possibly zero) vector in $\text{Span}(S_{i-1})$. Therefore $\text{Rank}(S_i)$ grows by t_i where $t_i \leq s_i - 1$. In Step ℓ , we add $T' \subset T_\ell$ to S . This increases $\text{Rank}(S)$ by $t_\ell \geq 1$ (since $\text{Rank}(S) \leq k - 2$ at the start) and $|S|$ by $s_\ell \geq t_\ell$. Thus

$$|S| = \sum_{i=1}^{\ell} s_i \geq \sum_{i=1}^{\ell-1} (t_i + 1) + t_\ell \geq k + \left\lceil \frac{k}{r} \right\rceil - 2. \quad (7)$$

The conclusion now follows from Fact 1 which implies that $|S| \leq n - d$. \square

Definition 6. We say that an (r, d) -code \mathcal{C} is optimal if its parameters satisfy (5) with equality.

Pyramid codes [6] yield optimal (r, d) -codes for all values of r, d , and k when the alphabet q is sufficiently large.

The proof of theorem 5 reveals information about the structure of optimal (r, d) -codes. We think of the algorithm as attempting to maximize

$$\frac{|S|}{\text{Rank}(S)} = \frac{\sum_{i=1}^{\ell} s_i}{\sum_{i=1}^{\ell} t_i}.$$

With this in mind, at step i we can choose \mathbf{c}_i such that $\frac{s_i}{t_i}$ is maximized. An optimal length code should yield the same value for $|S|$ for this (or any) choice of \mathbf{c}_i . This observation yields an insight into the structure of local dependencies in optimal codes, as given by the following structure theorem.

Theorem 7. Let \mathcal{C} be an $[n, k, d]_q$ code with information locality r . Suppose $r \mid k$, $r < k$, and

$$n = k + \frac{k}{r} + d - 2; \quad (8)$$

then hyperedges in the hypergraph $H(V, E)$ are disjoint and each has size exactly $r + 1$.

Proof. We execute the algorithm presented in the proof of theorem 5 to obtain a set S and sequences $\{s_i\}$ and $\{t_i\}$. We consider the case of $r = 1$ separately. Since all $t_i \leq 1$ we fall into Case 1. Combining formulas (6), (4) and (8) we get

$$|S| = \sum_{i=1}^{\ell} s_i = \sum_{i=1}^{\ell} t_i + \ell = 2k - 2.$$

Combining this with $\sum_{i=1}^{\ell} t_i = k - 1$ we conclude that $\ell = k - 1$, all s_i equal 2, and all t_i equal 1. The latter two conditions preclude the existence of hyperedges of size 1 or intersecting edges in H .

We now proceed to the case of $r > 1$. When $r \mid k$, the bound in equation (6) is larger than that in equation (7). Thus, we must be in Case 2. Combining formulas (7), (4) and (8) we get

$$|S| = \sum_{i=1}^{\ell} s_i = \sum_{i=1}^{\ell} t_i + \ell - 1 = k + \frac{k}{r} - 2.$$

Observe that $\sum_{i=1}^{\ell} t_i = k - 1$ and thus $\ell = \frac{k}{r}$. Together with the constraint $t_i \leq r$, this implies that $t_j = r - 1$ for some $j \in [\ell]$ and $t_i = r$ for $i \neq j$. We claim that in fact $j = \ell$. Indeed, if $j < \ell$, we would have $\sum_{i \leq \ell-1} t_i = k - r - 1$ and $t_\ell = r$, hence we would be in Case 1.

Now assume that there is an edge T with $|T| \leq r$. By adding this edge to S at the first step, we would get $t_1 \leq r - 1$. Next assume that $T_1 \cap T_2$ is non-empty. Observe that this implies $\text{Rank}(T_1 \cup T_2) < 2r$. So if we add edges T_1 and T_2 to S , we have $t_1 + t_2 \leq 2r - 1$. Clearly these conditions lead to a contradiction if $\ell = \frac{k}{r} \geq 3$. In fact, they also give a contradiction for $\frac{k}{r} = 2$, since they put us in Case 1. \square

4 Canonical codes

The structure theorem implies then when d is sufficiently small (which in our experience is the setting of interest in most data storage applications), optimal (r, d) -codes have rather rigid structure. We formalize this by defining the notion of a *canonical* code.

Definition 8. Let \mathcal{C} be a systematic $[n, k, d]_q$ code with information locality r where $r \mid k$, $r < k$, and $n = k + \frac{k}{r} + d - 2$. We say that \mathcal{C} is canonical if the set C can be partitioned into three groups $C = I \cup C' \cup C''$ such that:

1. Points $I = \{\mathbf{e}_1, \dots, \mathbf{e}_k\}$.
2. Points $C' = \{\mathbf{c}'_1, \dots, \mathbf{c}'_{k/r}\}$ where $\text{wt}(\mathbf{c}'_i) = r$. The supports of these vectors are disjoint sets which partition $[k]$.
3. Points $C'' = \{\mathbf{c}''_1, \dots, \mathbf{c}''_{d-2}\}$ where $\text{wt}(\mathbf{c}''_i) = k$.

Clearly any canonical code is systematic and has information locality r . The distance property requires a suitable choice of vectors $\{\mathbf{c}'\}$ and $\{\mathbf{c}''\}$. Pyramid codes [6] are an example of canonical codes. We note that since $r < k$, there is always a distinction between symbols $\{\mathbf{c}'\}$ and $\{\mathbf{c}''\}$.

Theorem 9. Assume that $d < r + 3$, $r < k$, and $r \mid k$. Let $n = k + \frac{k}{r} + d - 2$. Every systematic $[n, k, d]_q$ code with information locality r is a canonical code.

Proof. Let \mathcal{C} be a systematic $[n, k, d]$ code with information locality r . We start by showing that the hypergraph $H(V, E)$ has $\frac{k}{r}$ edges.

Since \mathcal{C} is systematic, we know that $I = \{\mathbf{e}_1, \dots, \mathbf{e}_k\} \subset C$. By theorem 7, $H(V, E)$ consists of m disjoint, $(r + 1)$ -regular edges and every vertex in I appears in some edge. But since the points in I are linearly independent, every edge involves at least one vertex from $C \setminus I$ and at most r from I . So we have $m \geq \frac{k}{r}$. We show that equality holds.

Assume for contradiction that $m \geq \frac{k}{r} + 1$. Since the edges are regular and disjoint, we have

$$n \geq m(r + 1) \geq k + \frac{k}{r} + r + 1 > k + \frac{k}{r} + d - 2$$

which contradicts the choice of n . Thus $m = \frac{k}{r}$. This means that every edge T_i is incident on exactly r vertices $\mathbf{e}_{i_1}, \dots, \mathbf{e}_{i_r}$ from I and one vertex \mathbf{c}'_i outside it. Hence

$$\mathbf{c}'_i = \sum_{j=1}^r \lambda_{i_j} \mathbf{e}_{i_j}.$$

Since the T_i s are disjoint, the vectors $\mathbf{c}'_1, \dots, \mathbf{c}'_{k/r}$ have disjoint supports which partition $[k]$.

We now show that the remaining vectors $\mathbf{c}''_1, \dots, \mathbf{c}''_{d-2}$ must all have $\text{wt}(\mathbf{c}'') = k$. For this, we consider the encoding of \mathbf{e}_j . We note $\mathbf{e}_i \cdot \mathbf{e}_j \neq 0$ iff $i = j$ and $\mathbf{c}'_i \cdot \mathbf{e}_j \neq 0$ iff $j \in \text{Supp}(\mathbf{c}'_i)$. Thus only 2 of these inner products are non-zero. Since the code has distance d , all the $d - 2$ inner products $\mathbf{c}''_i \cdot \mathbf{e}_j$ are non-zero. This shows that $\text{wt}(\mathbf{c}''_i) = k$ for all i . \square

The above bound is strong enough to separate having locality of r from having just information locality of r . The following corollary follows from the observation that the hypergraph $H(V, E)$ must contain $n - \frac{k}{r}(r + 1) = d - 2$ isolated vertices, which do not participate in any linear relations of size $r + 1$.

Corollary 10. *Assume that $2 < d < r + 3$ and $r \mid k$. Let $n = k + \frac{k}{r} + d - 2$. There are no $[n, k, d]_q$ linear codes with locality r .*

5 Canonical codes: parity locality

Theorem 9 gives a very good understanding of optimal (r, d) -codes in the case $d < r + 3$ and $r \mid k$. For any such code the coordinate set $C = \{\mathbf{c}_i\}_{i \in [n]}$ can be partitioned into sets I, C', C'' where for all $\mathbf{c} \in I \cup C'$, $\text{Loc}(\mathbf{c}) = r$, and for all $\mathbf{c}'' \in C''$, $\text{Loc}(\mathbf{c}'') > r$. It is natural to ask how low can the locality of symbols $\mathbf{c}'' \in C''$ be. In this section we address and resolve this question.

5.1 Parity locality lower bound

We begin with a lower bound.

Theorem 11. *Let \mathcal{C} be a systematic optimal (r, d) -code with parameters $[n, k, d]_q$. Suppose $d < r + 3$, $r < k$, and $r \mid k$. Then some $\frac{k}{r}$ parity symbols of \mathcal{C} have locality exactly r , and $d - 2$ remaining parity symbols of \mathcal{C} have locality no less than*

$$k - \left(\frac{k}{r} - 1\right)(d - 3). \quad (9)$$

Proof. Theorem 9 implies that \mathcal{C} is a canonical code. Let $C = I \cup C' \cup C''$ be the canonical partition of the coordinates of \mathcal{C} . Clearly, for all $\frac{k}{r}$ symbols $\mathbf{c}' \in C'$ we have $\text{Loc}(\mathbf{c}') \leq r$. We now prove lower bounds on the locality of symbols in $C' \cup C''$.

We start with symbols $\mathbf{c}'' \in C''$. For every $j \in [k/r]$ we define a subset $R_j \subseteq C$ that we call a *row*. Let $S_j = \text{Supp}(\mathbf{c}'_j)$. The j -th row contains the vector \mathbf{c}'_j , all r unit vectors in the support of \mathbf{c}'_j and the set C'' .

$$R_j = \{\mathbf{c}'_j\} \cup \left(\bigcup_{i \in S_j} \mathbf{e}_i \right) \cup C''.$$

Observe that restricted to $I \cup C'$ rows $\{R_j\}_{j \in [k/r]}$ form a partition. Consider an arbitrary symbol $\mathbf{c}'' \in C''$. Let $\ell = \text{Loc}(\mathbf{c}'')$. We have

$$\mathbf{c}'' = \sum_{i \in L} \mathbf{c}_i, \quad (10)$$

where $|L| = \ell$. In what follows we show that for each row R_j ,

$$|R_j \cap L| \geq r \quad (11)$$

needs to hold. It is not hard to see that this together with the structure of the sets $\{R_j\}$ implies inequality (9). To prove (11) we consider the code

$$\mathcal{C}_j = \{\mathcal{C}(\mathbf{x}) \mid \mathbf{x} \in \mathbb{F}_q^k \text{ such that } \text{Supp}(\mathbf{x}) \subseteq S_j\}. \quad (12)$$

Note that

$$\text{Supp}(\mathcal{C}_j) = \bigcup_{\mathbf{y} \in \mathcal{C}_j} \text{Supp}(\mathbf{y}) = R_j \quad \text{and} \quad \dim \mathcal{C}_j = r,$$

since coordinates of \mathcal{C} outside R_j do not depend on information coordinates in S_j . Observing that the distance of the code \mathcal{C}_j is at least d and $|R_j| = r + d - 1$ we conclude that (restricted to its support) \mathcal{C}_j is an MDS code. Thus any r symbols of \mathcal{C}_j are independent. It remains to note that (10) restricted to coordinates in S_j yields a non-trivial dependency of length at most $|R_j \cap L| + 1$ between the symbols of \mathcal{C}_j .

We proceed to the lower bound on the locality of symbols in C' . Fix an arbitrary $\mathbf{c}'_j \in C'$. A reasoning similar to the one above implies that if $\text{Loc}(\mathbf{c}'_j) < r$; then there is a dependency of length below $r + 1$ between the coordinates of the $[r + d - 1, r, d]_q$ code \mathcal{C}_j (defined by (12)) restricted to its support. \square

Observe that the bound (9) is close to k only when r is large and d is small. In other cases theorem 11 does not rule out existence of canonical codes with low locality for all symbols (including those in C''). In the next section we show that such codes indeed exist. In particular we show that the bound (9) can be always met with equality.

5.2 Parity locality upper bounds

Our main results in this section are given by theorems 15 and 16. Theorem 15 gives a general upper bound matching the lower bound of theorem 11. The proof is not explicit. Theorem 16 gives an explicit family of codes in the narrow case of $d = 4$. We start by introducing some concepts we need for the proof of theorem 15.

Definition 12. Let $L \subseteq \mathbb{F}_q^n$ be a linear space and $S \subseteq [n]$ be a set, $|S| = k$. We say that S is a k -core for L if for all vectors $\mathbf{v} \in L$, $\text{Supp}(\mathbf{v}) \not\subseteq S$.

It is not hard to verify that S is a k -core for L , if and only if S is a subset of some set of information coordinates in the space L^\perp . In other words S is a k -core for L , if and only if k columns in the $(n - \dim L)$ -by- n generator matrix of L^\perp that correspond to elements of S are linearly independent. Equivalently, a k -core for L is a correctable erasure pattern of size k .

Definition 13. Let $L \subseteq \mathbb{F}_q^n$ be a linear space. Let $\{\mathbf{c}_1, \dots, \mathbf{c}_n\}$ be a sequence of n vectors in \mathbb{F}_q^k . We say that vectors $\{\mathbf{c}_i\}$ are in general position subject to L if the following conditions hold:

1. For all vectors $\mathbf{v} \in L$ we have $\sum_{i=1}^n \mathbf{v}(i)\mathbf{c}_i = 0$;
2. For all k -cores S of L we have $\text{Rank}(\{\mathbf{c}_i\}_{i \in S}) = k$.

To get some intuition regarding the definition above observe that for any L and $k = n - \dim L$ columns of a $k \times n$ parity check matrix of L are in general position subject to L . The next lemma asserts existence of vectors that are in general position subject to an arbitrary linear space provided the underlying field is large enough.

Lemma 14. *Let $L \subseteq \mathbb{F}_q^n$ be a linear space and k be a positive integer. Suppose $q > kn^k$; then there exists a family of vectors $\{\mathbf{c}_i\}_{i \in [n]}$ in \mathbb{F}_q^k that are in general position subject to L .*

Proof. We obtain a matrix $M \in \mathbb{F}_q^{k \times n}$ picking the rows of M at random (uniformly and independently) from the linear space L^\perp . We choose vectors $\{\mathbf{c}_i\}$ to be the columns of M . Observe that the first condition in definition 13 is always satisfied. Further observe that our choice of M induces a uniform distribution on every set of k columns of M that form a k -core. The second condition in definition 13 is satisfied as long as all k -by- k minors of M that correspond to k -cores are invertible. This happens with probability at least

$$1 - \binom{n}{k} \cdot \left(1 - \prod_{i=1}^k \left(1 - \frac{1}{q^i}\right)\right) \geq 1 - \binom{n}{k} \cdot \left(1 - \left(1 - \frac{1}{q}\right)^k\right) \geq 1 - n^k \cdot \frac{k}{q} > 0.$$

This concludes the proof. \square

We proceed to the main result of this section.

Theorem 15. *Let $2 < d < r + 3$, $r < k$, $r \mid k$. Let $q > kn^k$ be a prime power. Let $n = k + \frac{k}{r} + d - 2$. There exists a systematic $[n, k, d]_q$ code \mathcal{C} of information locality r , where $\frac{k}{r}$ parity symbols have locality r , and $d - 2$ remaining parity symbols have locality $k - \left(\frac{k}{r} - 1\right)(d - 3)$.*

Proof. Let $t = \frac{k}{r}$. Fix some $t + 1$ subsets P_0, P_1, \dots, P_t of $[n]$ subject to the following constraints:

1. $|P_0| = k - (t - 1)(d - 3) + 1$;
2. For all $i \in [t]$, $|P_i| = r + 1$;
3. For all $i, j \in [t]$ such that $i \neq j$, $P_i \cap P_j = \emptyset$;
4. For all $i \in [t]$, $|P_0 \cap P_i| = r - d + 3$.

For every set P_i , $0 \leq i \leq t$ we fix a vector $\mathbf{v}_i \in \mathbb{F}_q^n$, such that $\text{Supp}(\mathbf{v}_i) = P_i$. We ensure that non-zero coordinates of \mathbf{v}_0 contain the same value. We also ensure that for all $i \in [t]$ non-zero coordinates of \mathbf{v}_i contain distinct values. The lower bound on q implies that these conditions can be met. For a finite set A let A° denote a certain (unspecified) set that is obtained from A by dropping *at most* one element. Note that for all $i \in [t]$ and all non-zero α, β in \mathbb{F}_q we have

$$\text{Supp}(\alpha \mathbf{v}_0 + \beta \mathbf{v}_i) = (P_0 \setminus P_i) \sqcup (P_0 \cap P_i)^\circ \sqcup (P_i \setminus P_0). \quad (13)$$

Consider the space $L = \text{Span}(\{\mathbf{v}_i\}_{0 \leq i \leq t})$. Let $M = P_0 \setminus \bigsqcup_{i=1}^t P_i$. Observe that

$$|M| = k - (t - 1)(d - 3) + 1 - t(r - d + 3) = d - 2.$$

By (13) for any $\mathbf{v} \in L$ we have

$$\text{Supp}(\mathbf{v}) = \begin{cases} \bigsqcup_{i \in T} P_i, & \text{for some } T \subseteq [t] \text{ OR} \\ M \bigsqcup_{i \in [t] \setminus T} (P_0 \cap P_i) \bigsqcup_{i \in T} (P_0 \cap P_i)^\circ \bigsqcup_{i \in T} (P_i \setminus P_0) & \text{for some } T \subseteq [t]. \end{cases} \quad (14)$$

Observe that a set $K \subseteq [n]$, $|K| = k$ is a k -core for L if and only if for all $i \in [t]$, $P_i \not\subseteq K$ and

$$\left[\begin{array}{l} M \not\subseteq K; \text{ OR} \\ M \subseteq K \text{ and } \exists i \in [t] \text{ such that } \left[\begin{array}{l} |P_i \cap P_0 \cap K| < r - d + 2; \text{ OR} \\ |P_i \cap P_0 \cap K| = r - d + 2 \text{ and } P_i \setminus P_0 \not\subseteq K. \end{array} \right. \end{array} \right. \quad (15)$$

Let $I \subseteq [n]$ be such that $M \cap I = \emptyset$ and for all $i \in [t]$, $|I \cap P_i| = r$. By (15) I is a k -core for L . We use Lemma 14 to obtain vectors $\{\mathbf{c}_i\}_{i \in [n]} \in \mathbb{F}_q^k$ that are in general position subject to the space L . We choose vectors $\{\mathbf{c}_i\}_{i \in I}$ as our basis for \mathbb{F}_q^k and consider the code \mathcal{C} defined as in (3).

It is not hard to see that \mathcal{C} is a systematic code of information locality r . All t parity symbols in the set $(\bigsqcup_{i \in [t]} P_i) \setminus I$ also have locality r . Furthermore all $d - 2$ parity symbols in the set M have locality $k - (t - 1)(d - 3)$. It remains to prove that the code \mathcal{C} has distance

$$d = n - k - t + 2. \quad (16)$$

According to Fact 1 the distance of \mathcal{C} equals $n - |S|$ where $S \subseteq [n]$ is the largest set such that vectors $\{\mathbf{c}_i\}_{i \in S}$ do not have full rank. By definition 13 for any k -core K of L we have $\text{Rank}\{\mathbf{c}_i\}_{i \in K} = k$. Thus in order to establish (16) it suffices to show that every set $S \subseteq [n]$ of size $k + t - 1$ contains a k -core of L . Our proof involves case-by-case analysis. Let $S \subseteq [n]$, $|S| = k + t - 1$ be an arbitrary set. Set

$$b = \#\{i \in [t] \mid P_i \subseteq S\}.$$

Note that since $t(r + 1) > |S|$ we have $b \leq t - 1$.

Case 1: $M \not\subseteq S$. We drop $t - 1$ elements from S to obtain a set $K \subseteq S$, $|K| = k$ such that for all $i \in [t]$, $P_i \not\subseteq K$. By (15) K is a k -core.

Case 2: $M \subseteq S$ and $b \leq t - 2$. We drop $t - 1$ elements from S to obtain a set $K \subseteq S$, $|K| = k$ such that $M \not\subseteq K$ and for all $i \in [t]$, $P_i \not\subseteq K$. By (15) K is a k -core.

Case 3: $M \subseteq S$ and $b = t - 1$. Let $i \in [t]$ be such that $P_i \not\subseteq S$. Such i is unique. Observe that

$$|P_i \cap S| = k + t - 1 - (d - 2) - (t - 1)(r + 1) = r - d + 2.$$

Also observe that $|P_i \setminus P_0| = r + 1 - (r - d + 3) = d - 2 \geq 1$. Combining the last two observations we conclude that either

$$\left[\begin{array}{l} |P_i \cap P_0 \cap S| < r - d + 2; \text{ OR} \\ |P_i \cap P_0 \cap S| = r - d + 2 \text{ and } P_i \setminus P_0 \not\subseteq S. \end{array} \right. \quad (17)$$

Finally, we drop $t - 1$ elements from S to obtain a set $K \subseteq S$, $|K| = k$ such that for all $i \in [t]$, $P_i \not\subseteq K$. By (15) and (17) K is a k -core. \square

Theorem 15 gave a general construction of (r, d) -codes that are optimal not only with respect to information locality and redundancy but also with respect to locality of parity symbols. That theorem however is weak in two respects. Firstly, the construction is not explicit. Secondly, the construction requires a large underlying field. The next theorem gives an explicit construction that works even over small fields in the narrow case of codes of distance 4.

Theorem 16. *Let $r < k$, $r \mid k$ be positive integers. Let $q \geq r + 2$ be a prime power. Let $n = k + \frac{k}{r} + 2$. There exists a systematic $[n, k, 4]_q$ code \mathcal{C} of information locality r , where $\frac{k}{r}$ parity symbols have locality r , and 2 remaining parity symbols have locality $k - \frac{k}{r} + 1$.*

Proof. Fix an arbitrary systematic $[r+3, r, 4]_q$ code \mathcal{E} . For instance, one can choose \mathcal{E} to be a (possibly doubly extended) Reed-Solomon code. Let

$$\mathcal{E}(\mathbf{y}) = (\mathbf{y}, \mathbf{p}_0 \cdot \mathbf{y}, \mathbf{p}_1 \cdot \mathbf{y}, \mathbf{p}_2 \cdot \mathbf{y}).$$

Since \mathcal{E} is an MDS code all vectors $\{\mathbf{p}_i\}$ have weight r . Thus for some non-zero $\{\alpha_j\}_{j \in [r]}$ we have

$$\mathbf{p}_1 = \sum_{j=1}^{r-1} \alpha_j \mathbf{e}_j + \alpha_r \mathbf{p}_2, \quad (18)$$

where $\{\mathbf{e}_j\}_{j \in [r]}$ are the r -dimensional unit vectors. To define a systematic code \mathcal{C} we partition the input vector $\mathbf{x} \in \mathbb{F}_q^k$ into $t = \frac{k}{r}$ vectors $\mathbf{y}_1, \dots, \mathbf{y}_t \in \mathbb{F}_q^r$. We set

$$\mathcal{C}(\mathbf{x}) = \left(\mathbf{y}_1, \dots, \mathbf{y}_t, \mathbf{p}_0 \cdot \mathbf{y}_1, \dots, \mathbf{p}_0 \cdot \mathbf{y}_t, \left(\mathbf{p}_1 \cdot \sum \mathbf{y}_i \right), \left(\mathbf{p}_2 \cdot \sum \mathbf{y}_i \right) \right), \quad (19)$$

where the summation is over all $i \in [t]$. It is not hard to see that the first $k+t$ coordinates of \mathcal{C} have locality r . We argue that the last two coordinates have locality $k-t+1$. From (18) we have

$$\left(\mathbf{p}_1 \cdot \sum \mathbf{y}_i \right) = \sum_{j=1}^{r-1} \alpha_j \left(\mathbf{e}_j \cdot \sum \mathbf{y}_i \right) + \alpha_r \left(\mathbf{p}_2 \cdot \sum \mathbf{y}_i \right),$$

where the summation is over all $i \in [t]$. Equivalently,

$$\left(\mathbf{p}_1 \cdot \sum \mathbf{y}_i \right) = \sum_{j=1}^{r-1} \sum_{i \in [t]} \alpha_j \mathbf{y}_i(j) + \alpha_r \left(\mathbf{p}_2 \cdot \sum \mathbf{y}_i \right).$$

Thus the next-to-last coordinate of \mathcal{C} can be recovered from accessing $(r-1)t$ information coordinates and the last coordinate. Similarly, the last coordinate can be recovered from $k-t$ information coordinates and the next-to-last coordinate. To prove that the code \mathcal{C} has distance 4 we give an algorithm to correct 3 erasures in \mathcal{C} . The algorithm has two steps.

Step 1: For every $i \in [t]$, we refer to a subset $(\mathbf{y}_i, \mathbf{p}_0 \cdot \mathbf{y}_i)$ of $r+1$ coordinates of \mathcal{C} as a block. We go over all t blocks. If we encounter a block where one symbol is erased, we recover this symbol from other symbols in the block.

Step 2: Observe that after the execution of Step 1 there can be at most one block that has erasures. If no such block exists; then on Step 1 we have successfully recovered all information symbols and thus we are done. Otherwise, let the unique block with erasures be $(\mathbf{y}_j, \mathbf{p}_0 \cdot \mathbf{y}_j)$ for some $j \in [t]$. Since we know all vectors $\{\mathbf{y}_i\}_{i \neq j, i \in [t]}$, from $\left(\mathbf{p}_1 \cdot \sum_{i \in [t]} \mathbf{y}_i \right)$ and $\left(\mathbf{p}_2 \cdot \sum_{i \in [t]} \mathbf{y}_i \right)$ (if these symbols are not erased) we recover symbols $\mathbf{p}_1 \cdot \mathbf{y}_j$ and $\mathbf{p}_2 \cdot \mathbf{y}_j$. Finally, we invoke the decoding procedure for the code \mathcal{E} to recover \mathbf{y}_j from at most 3 erasures in $\mathcal{E}(\mathbf{y}_j) = (\mathbf{y}_j, \mathbf{p}_0 \cdot \mathbf{y}_j, \mathbf{p}_1 \cdot \mathbf{y}_j, \mathbf{p}_2 \cdot \mathbf{y}_j)$. \square

6 Non-canonical codes

In this section we observe that canonical codes detailed in Sections 4 and 5 are not the only family of optimal (r, d) -codes. If one relaxes conditions of theorem 9 one can get other families. One such family that yields uniform locality for all symbols is given below. The (non-explicit) proof resembles the proof of theorem 15 albeit is much simpler.

Theorem 17. Let n, k, r , and $d \geq 2$ be positive integers. Let $q > kn^k$ be a prime power. Suppose $(r+1) \mid n$ and

$$n - k = \left\lceil \frac{k}{r} \right\rceil + d - 2;$$

then there exists an $[n, k, d]_q$ code where all symbols have locality r .

Proof. Let $t = \frac{n}{r+1}$. We partition the set $[n]$ into t disjoint subsets P_1, \dots, P_t each of size $r+1$. For every $i \in [t]$ we fix a vector $\mathbf{v}_i \in \mathbb{F}_q^n$, such that $\text{Supp}(\mathbf{v}_i) = P_i$. We set all non-zero coordinates in vectors $\{\mathbf{v}_i\}_{i \in [t]}$ to be equal to 1. We consider the linear space $L = \text{Span}(\{\mathbf{v}_i\}_{i \in [t]})$. For every $\mathbf{v} \in L$ we have

$$\text{Supp}(\mathbf{v}) = \bigsqcup_{i \in T} P_i \quad \text{for some } T \subseteq [t].$$

Observe that a set $K \subseteq [n]$, $|K| = k$ is a k -core for L if and only if for all $i \in [t]$, $P_i \not\subseteq K$. Also observe that conditions of the theorem imply $k \leq n - t$. Therefore k -cores for L exist. We use Lemma 14 to obtain vectors $\{\mathbf{c}_i\}_{i \in [n]} \in \mathbb{F}_q^k$ that are in general position subject to the space L . We consider the code \mathcal{C} defined as in (3). It is not hard to see that \mathcal{C} has dimension k and locality r for all symbols. It remains to prove that the code \mathcal{C} has distance

$$d = n - k - \left\lceil \frac{k}{r} \right\rceil + 2. \tag{20}$$

Our proof relies on Fact 1. Let $S \subseteq [n]$ be an arbitrary subset such that $\text{Rank}\{\mathbf{c}_i\}_{i \in S} < k$. Clearly, no k -core of L is in S . Let

$$b = \#\{i \in [t] \mid P_i \subseteq S\}.$$

We have $|S| - b \leq k - 1$ since dropping b elements from S (one from each $P_i \subseteq S$) turns $|S|$ into an $(|S| - b)$ -core. We also have $br \leq k - 1$ since dropping one element from each $P_i \subseteq S$ gives us a br -core in S . Combining the last two inequalities we conclude that

$$|S| \leq k + \left\lceil \frac{k-1}{r} \right\rceil - 1.$$

Combining this inequality with the identity $\left\lceil \frac{k-1}{r} \right\rceil = \left\lceil \frac{k}{r} \right\rceil - 1$ and Fact 1 we obtain (20). \square

7 Beyond worst-case distance

In this section, codes are assumed to be systematic unless otherwise stated. They will have k information symbols and $h = n - k$ parity check symbols.

7.1 Maximally recoverable codes

The supports of the parity check symbols in a code can be described using a bipartite graph. More generally, we define the notion of a set of points with supports matching a graph G .

Definition 18. Let $G([k], [h], E)$ be a bipartite graph. We say that $\mathbf{c}_1, \dots, \mathbf{c}_h \in \mathbb{F}_q^k$ have supports matching G if $\text{Supp}(\mathbf{c}_j) = \Gamma(j)$ for all $j \in [h]$ where $\Gamma(j)$ denotes the neighborhood of j in G .

Given points $\mathbf{c}_1, \dots, \mathbf{c}_h$, consider the $k \times h$ matrix C with columns $\mathbf{c}_1, \dots, \mathbf{c}_h$. For $I \subseteq [k]$ and $J \subseteq [h]$, let $C_{I,J}$ denote the sub-matrix of C with rows indexed by I and columns indexed by J .

Definition 19. Points $\mathbf{c}_1, \dots, \mathbf{c}_h \in \mathbb{F}_q^k$ with supports matching G are in general position if for every $I \subseteq [h]$ and $J \subseteq [k]$ such that there is a perfect matching from I to J in G , the sub-matrix $C_{I,J}$ is invertible.

Standard arguments show that over sufficiently large fields \mathbb{F}_q , choosing $\mathbf{c}_1, \dots, \mathbf{c}_h$ randomly from the set of vectors with support matching G gives points in general position.

Coming back to codes, the supports of the parity checks define a bipartite graph which we will call the support graph. This is closely related to but distinct from the Tanner graph.

Definition 20. Let \mathcal{C} be a systematic code with point set $C = \{\mathbf{e}_1, \dots, \mathbf{e}_k, \mathbf{c}_1, \dots, \mathbf{c}_h\}$. The support graph $G([k], [h], E)$ of \mathcal{C} is a bipartite graph where $(i, j) \in E$ if $i \in \text{Supp}(\mathbf{c}_j)$.

For instance in any canonical (r, d) -code, the support graph is specified up to relabeling. In the right hand side of the graph G there are $\frac{k}{r}$ vertices of degree r corresponding to $C' \subseteq C$, whose neighborhood partitions the set $[k]$ and $d - 2$ vertices of degree k corresponding to $C'' \subseteq C$. The minimum distance of such a code is exactly d , and hence there are some patterns of d erasures that the code cannot correct. However it is possible that the code could correct many patterns of erasures of weight d and higher, for a suitable choice of \mathbf{c}_i s. In general one could ask: among all codes with a support graph G , which codes can correct the most erasure patterns? *A priori*, it is unclear that there should be a single code that is optimal in the sense that it corrects the maximal possible set of patterns. As shown by [2, 6] such codes do exist over sufficiently large fields.

Consider a systematic code \mathcal{C} with support graph $G([k], [h], E)$. Given $I \subseteq [k]$ and $J \subseteq [h]$, let $\Gamma_J(I)$ denote $\Gamma(I) \cap J$ (define $\Gamma_I(J)$ similarly). Consider a set of erasures $S \cup T$ where $S \subseteq [k]$ and $T \subseteq [h]$ are the sets of information and parity check symbols respectively that are erased. To correct these erasures, we need to recover the symbols corresponding to $\{\mathbf{e}_i : i \in S\}$ from the parity checks corresponding to $\{\mathbf{c}_j : j \in \bar{T} = [h] \setminus T\}$. For this to be possible, a necessary condition is that for every $S' \subseteq S$, $|\Gamma_{\bar{T}}(S')| \geq |S'|$. By Hall's theorem [7], this is equivalent to the existence of a matching in G from S to \bar{T} . We say that such a set of erasures satisfies Hall's condition.

Definition 21. A systematic code \mathcal{C} with support graph G is maximally recoverable if every set of erasures satisfying Hall's condition can be corrected.

We can rephrase this definition in algebraic terms using the notion of points with specified supports in general position.

Theorem 22. [2, 6] Let \mathcal{C} be a systematic code with support graph G . \mathcal{C} is a maximally recoverable code iff $\mathbf{c}_1, \dots, \mathbf{c}_h$ are in general position with supports matching G .

7.2 The tradeoff between locality and erasure correction

In this section we study locality of parity symbols in maximally recoverable codes. Locality of information symbols in such codes has been addressed in [2].

For any parity check symbol \mathbf{c}_j , it is clear that $\text{Loc}(\mathbf{c}_j) \leq \text{wt}(\mathbf{c}_j) = \text{deg}(j)$. Below we show that in a maximally recoverable code no better locality is possible. Our result relies on a characterization of the support of the vectors in the space \mathcal{V} spanned by $\{\mathbf{c}_1, \dots, \mathbf{c}_h\}$ in terms of the graph G .

Let \mathcal{V} denote the space spanned by vectors $\{\mathbf{c}_1, \dots, \mathbf{c}_h\}$ which are in general position with supports matching G . Let $\mathcal{F} \subseteq 2^{[k]}$ denote the set of supports of vectors in \mathcal{V} . We give a necessary condition for membership in \mathcal{F} . Our condition is in terms of sets of coordinates that can be eliminated by combination of certain \mathbf{c}_j s.

Definition 23. Let $\mathbf{c} = \sum_{j \in J} \mu_j \mathbf{c}_j$ where $\mu_j \neq 0$ for all $j \in J$. Let $I = \cup_{j \in J} \Gamma(j) \setminus \text{Supp}(\mathbf{c})$. We say that the set I has been eliminated from $\cup_{j \in J} \Gamma(j)$.

Theorem 24. Let $\{\mathbf{c}_1, \dots, \mathbf{c}_h\}$ be vectors with supports matching G in general position. The set I can be eliminated from $\cup_{j \in J} \Gamma(j)$ only if $|\Gamma_J(I')| > |I'|$ for every non-empty $I' \subseteq I$.

Proof. Let $\mathbf{c} = \sum_{j \in J} \mu_j \mathbf{c}_j$ where $\mu_j \neq 0$ for all $j \in J$. Let $I = \cup_{j \in J} \Gamma(j) \setminus \text{Supp}(\mathbf{c})$. Assume for contradiction that there exists $\tilde{I} \subseteq I$ where $|\Gamma_J(\tilde{I})| \leq |\tilde{I}|$. We will show that there exists $I' \subseteq \tilde{I}$ so that $|\Gamma_J(I')| = |I'|$ and that $|\Gamma_J(I'')| > |I''|$ for every non-empty subset $I'' \subsetneq I'$.

It suffices to prove the existence of $I' \subseteq \tilde{I}$ where $|\Gamma_J(I')| = |I'|$; the claim about subsets of I' will then follow by taking a minimal such I' . Observe that every $i \in I$ must have $|\Gamma_J(i)| \geq 2$, since if i occurs in exactly one S_j , then it cannot be eliminated. Hence we must have $|\tilde{I}| \geq 2$.

One can construct the set \tilde{I} starting from a singleton and adding elements one at a time, giving a sequence $I_1, \dots, I_\ell = \tilde{I}$. We claim that for any $l \leq \ell$,

$$|\Gamma_J(I_l)| - |I_l| \geq |\Gamma_J(I_{l-1})| - |I_{l-1}| - 1.$$

This holds since $|\Gamma_J(I_l)|$ can only increase on adding i to I_{l-1} while $|I_l|$ increases by 1. Since $|\Gamma_J(I_1)| - |I_1| \geq 1$ whereas $|\Gamma_J(I_\ell)| - |I_\ell| \leq 0$, we must have $|\Gamma_J(I_l)| - |I_l| = 0$ for some $l \leq \ell$. Thus we have a set where $|\Gamma_J(I_l)| = |I_l|$ as desired.

Since the set I' satisfies Hall's matching condition, there is a perfect matching from I' to $J' = \Gamma_J(I')$ in G . But this means that the sub-matrix $C_{I', J'}$ has full rank. On the other hand, $\mathbf{c} = \sum_j \mu_j \mathbf{c}_j$ and $I' \subseteq I = \cup_j \Gamma(j) \setminus \text{Supp}(\mathbf{c})$. Let $\pi(\mathbf{c})$ denote the restriction of \mathbf{c} onto the coordinates in I' . Then we have

$$\sum_{j \in J} \mu_j \pi(\mathbf{c}_j) = \sum_{j \in J'} \mu_j \pi(\mathbf{c}_j) = \pi(\mathbf{c}) = 0.$$

The first equality holds because $\pi(\mathbf{c}_j) = 0$ for $j \notin \Gamma(I')$, the second by linearity of π and the last since $\pi(\mathbf{c}) = 0$. Hence the vector $\mu'_{J'} = \{\mu_j\}_{j \in J'}$ lies in the kernel of $C_{I', J'}$ which contradicts the assumption that it has full rank.

This shows that the condition $|\Gamma_J(I')| > |I'|$ for all $I' \subseteq I$ is necessary. \square

Corollary 25. If the set I can be eliminated from $\cup_{j \in J} \Gamma(j)$, then $|I| \leq |J| - 1$.

If the field size q is sufficiently large, the necessary condition given by theorem 24 is also sufficient. We defer the proof of this statement to Appendix A and prove our lower bound on the locality of maximally recoverable codes.

Theorem 26. In a maximally recoverable code, $\text{Loc}(\mathbf{c}_j) = \text{deg}(j)$ for all $j \in [h]$.

Proof. Assume for contradiction that $\text{Loc}(\mathbf{c}_t) \leq \text{deg}(t) - 1$ for some $t \in [h]$. Hence there exist $A \subseteq [k]$ and $B \subseteq [h]$ (not containing t) such that

$$\mathbf{c}_t = \sum_{i \in A} \lambda_i \mathbf{e}_i + \sum_{j \in B} \mu_j \mathbf{c}_j.$$

We have $|A| = a$, $|B| = b$ and $a + b \leq \text{deg}(t) - 1$. Hence

$$\mathbf{c}_t - \sum_{j \in B} \mu_j \mathbf{c}_j = \sum_{i \in A} \lambda_i \mathbf{e}_i.$$

Thus we have eliminated at least $\text{deg}(t) - a \geq b + 1$ indices from $\cup_{j \in B \cup \{t\}} \text{Supp}(\mathbf{c}_j)$ using a linear combination of $b + 1$ vectors. By corollary 25 this is not possible for vectors in general position. \square

References

- [1] Viveck R. Cadambe, Syed A. Jafar, and Hamed Maleki. Distributed data storage with minimum storage regenerating codes - exact and functional repair are asymptotically equally efficient. Arxiv 1004.4299, 2010.
- [2] 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.
- [3] 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.
- [4] Alexandros G. Dimakis, Kannan Ramchandran, Yunnan Wu, and Changho Suh. A survey on network codes for distributed storage. *Proceedings of the IEEE*, 99:476–489, 2011.
- [5] Klim Efremenko. 3-query locally decodable codes of subexponential length. In *41st ACM Symposium on Theory of Computing (STOC)*, pages 39–44, 2009.
- [6] Cheng Huang, Minghua Chen, and Jin Li. Pyramid codes: flexible schemes to trade space for access efficiency in reliable data storage systems. In *6th IEEE International Symposium on Network Computing and Applications (NCA 2007)*, pages 79–86, 2007.
- [7] Stasys Jukna. *Extremal combinatorics*. Springer, Berlin, Heidelberg, New York, 2001.
- [8] Jonathan Katz and Luca Trevisan. On the efficiency of local decoding procedures for error-correcting codes. In *32nd ACM Symposium on Theory of Computing (STOC)*, pages 80–86, 2000.
- [9] Swastik Kopparty, Shubhangi Saraf, and Sergey Yekhanin. High-rate codes with sublinear-time decoding. In *43rd ACM Symposium on Theory of Computing (STOC)*, pages 167–176, 2011.
- [10] K. V. Rashmi, Nihar B. Shah, and P. Vijay Kumar. Optimal exact-regenerating codes for distributed storage at the MSR and MBR points via a product-matrix construction. *IEEE Transactions on Information Theory*, 57:5227–5239, 2011.
- [11] Itzhak Tamo, Zhiying Wang, and Jehoshua Bruck. MDS array codes with optimal rebuilding. Arxiv 1103.3737, 2011.
- [12] Michael Tsfasman, Serge Vladut, and Dmitry Nogin. *Algebraic geometric codes: basic notions*. American Mathematical Society, Providence, Rhode Island, USA, 2007.
- [13] Zhiying Wang, Itzhak Tamo, and Jehoshua Bruck. On codes for optimal rebuilding access. Arxiv 1107.1627, 2011.
- [14] Sergey Yekhanin. Towards 3-query locally decodable codes of subexponential length. *Journal of the ACM*, 55:1–16, 2008.
- [15] Sergey Yekhanin. Locally decodable codes. *Foundations and trends in theoretical computer science*, 2012. To appear. Preliminary version available for download at http://research.microsoft.com/en-us/um/people/yekhanin/Papers/LDC_now.pdf.

A Spaces spanned by general position vectors

Lemma 27. *Let $q \geq n$ be a prime power. The set of supports of vectors in any linear space $\mathcal{V} \subseteq \mathbb{F}_q^n$ is closed under union.*

Proof. Consider two vectors \mathbf{a} and \mathbf{b} in \mathcal{V} with $\text{Supp}(\mathbf{a}) = S$ and $\text{Supp}(\mathbf{b}) = T$. We may assume that $|S|, |T| \leq n-1$ and that one set does not contain the other. Now consider $\mathbf{a} + \lambda\mathbf{b}$ for $\lambda \in \mathbb{F}_q^*$. It suffices to find λ such that $\mathbf{a}(i) + \lambda\mathbf{b}(i) \neq 0$ for each $i \in S \cap T$. This rules out at most $|S \cap T| \leq n-2$ values of λ , so there is a solution provided $q-1 > n-2$ or $q \geq n$. \square

It is easy to see that the condition $q \geq n$ is tight by considering the length 3 parity check code over \mathbb{F}_2 , where the set of supports is not closed under union.

Theorem 28. *Let $q \geq n$. Let $\{\mathbf{c}_1, \dots, \mathbf{c}_h\}$ be vectors with supports matching G in general position which span a space \mathcal{V} . Let $\mathcal{F} \subseteq 2^{[k]}$ be the set of supports of vectors in \mathcal{V} . Then \mathcal{F} consists of all sets of the form $\cup_{j \in J} \Gamma(j) \setminus I$ where I satisfies the condition $|\Gamma_J(I')| > |I'|$ for every $I' \subseteq I$.*

Proof. Theorem 24 shows that the condition on I is necessary, we now show that it is sufficient. For $j \notin \Gamma(I)$ the sets $\Gamma(j)$ and I are disjoint. Hence we can write

$$\cup_{j \in J} \Gamma(j) \setminus I = (\cup_{j \in J \cap \Gamma(I)} \Gamma(j) \setminus I) \cup (\cup_{j \in J \cap \overline{\Gamma(I)}} \Gamma(j)).$$

By the closure under union, it suffices to prove the statement in the case when $J \subseteq \Gamma(I)$. Fix $j_0 \in J$ and let $J' = J \setminus \{j_0\}$. Since $|\Gamma_J(I')| > |I'|$ we have $|\Gamma_{J'}(I')| \geq |I'|$ for every $I' \subseteq I$. So there is a matching from I to some subset $J'' \subseteq J'$ where $|J''| = |I|$, and the matrix $C_{I, J''}$ is of full rank since the \mathbf{c}_j s are in general position.

Let $\pi(\mathbf{c})$ denote the restriction of a vector \mathbf{c} onto coordinates in I . Since $C_{I, J''}$ is invertible, the row vectors $\{\pi(\mathbf{c}_j)\}_{j \in J''}$ have full rank. Note that $\pi(\mathbf{c}_{j_0})$ is not a zero vector since $j_0 \in \Gamma(I)$. So there exist $\{\mu_j\}$ for $j \in J$ which are not all 0 and

$$\pi(\mathbf{c}_{j_0}) = \sum_{j \in J''} \mu_j \pi(\mathbf{c}_j).$$

Now consider the vector $\mathbf{c}'_{j_0} = \mathbf{c}_{j_0} - \sum_{j \in J''} \mu_j \mathbf{c}_j$. Note that $\pi(\mathbf{c}'_{j_0})$ is a zero vector, which shows that $\text{Supp}(\mathbf{c}'_{j_0}) \subseteq \cup_{j \in J'' \cup \{j_0\}} \Gamma(j) \setminus I$. We will show that equality holds by using corollary 25. Since we have eliminated $|I|$ vectors, the linear combination must involve at least $|I| + 1$ vectors, which means that $\mu_j \neq 0$ for all j . Further the set of eliminated co-ordinates cannot be larger than I , since this would violate corollary 25. Hence we have

$$\text{Supp}(\mathbf{c}'_{j_0}) = \cup_{j \in J'' \cup \{j_0\}} \Gamma(j) \setminus I. \quad (21)$$

By repeating this argument for every $j_0 \in J$, we will be able to find $J(j_0) \subseteq J$ of size $|I| + 1$ which contains j_0 and a vector \mathbf{c}'_{j_0} such that

$$\text{Supp}(\mathbf{c}'_{j_0}) = \cup_{j \in J(j_0)} \Gamma(j) \setminus I.$$

Using the closure under union of supports, we conclude that \mathcal{F} contains the set

$$\cup_{j_0 \in J} \cup_{j \in J(j_0)} \Gamma(j) \setminus I = \cup_{j \in J} \Gamma(j) \setminus I.$$

\square