An Upper Bound On the Size of Locally Recoverable Codes

Viveck Cadambe
Research Laboratory of Electronics, MIT
Cambridge, MA 02139
email: viveck@mit.edu

Arya Mazumdar
Department of Electrical and Computer Engineering
University of Minnesota, Minneapolis
MN 55455, email: arya@umn.edu

***Abstract*—In a *locally recoverable* or *repairable* code, any symbol of a codeword can be recovered by reading only a small (constant) number of other symbols. The notion of local recoverability is important in the area of distributed storage where a most frequent error-event is a single storage node failure (erasure). A common objective is to repair the node by downloading data from as few other storage node as possible. In this paper, we bound the minimum distance of a code in terms of of its length, size and locality. Unlike previous bounds, our bound follows from a significantly simple analysis and depends on the size of the alphabet being used.**

I. Introduction

The increased demand of cloud computing and storage services in current times has led to a corresponding surge in the study and deployment of erasure-correcting codes, or simply erasure codes, for distributed storage systems. In the information and coding theory community, this has led to the research of some new aspects of codes particularly tailored to the application to storage systems. The topic of interest of this paper is the locality of repair of erasure codes.

It is well known that an erasure code with length n, dimension k and minimum distance d, or an (n, k, d) code, can recover from *any* set of $d - 1$ erasures. In addition, the code is said to have *locality* r if any *single* erasure can be recovered from some set of r symbols of the codeword. From an engineering perspective, when an (n, k, d) code is used to store information in n storage nodes, the parameter d represents the worst-case (node) failure scenario from which the storage system can recover. The parameter r, on the other hand, represents the efficiency of recovery from a (relatively) more commonly occurring hurdle - a single node failure. It is therefore desirable to have a large value of d and a small value of r. Much literature in classical coding theory has been devoted to understanding the largest possible value of d - the minimum distance - for a fixed (n, k); one of the well-known result from this body [10] of work is the Singleton bound, and code constructions that achieve this bound (such as Reed-Solomon codes). The study of minimizing the locality, r, was pioneered recently in [1] and furthered in [2], [3], [7]–[9], [11]. The key discovery of [1], [7] is that, for any (n, k, d) code with and locality r, the following bound is satisfied:

$$d \leq n - k - \lceil k/r \rceil + 2. \tag{1}$$

The above bound is a generalization of the Singleton bound to include the locality of the codeword, r; when $r = k$, the above bound collapses to the classical Singleton bound. In addition, through an invocation of the *multicast* capacity of wireline networks via random network coding, the reference [7] showed that the above bound is indeed tight for a sufficiently large field size. Intuitively speaking, the above bound implies that there is a *cost* to locality; the smaller the locality, r, the smaller the minimum-distance d. Code constructions that achieve the above bound based on Reed-Solomon codes, among other techniques, have been recently discovered in [3], [4], [7], [9], [12]. Missing from these works is a formal study of the impact of an important parameter of the code - the size of the alphabet (field) of the code. Codes over small alphabets are the central subject of classical coding theory, and are of interest in the application of storage because of their implementation ease. We remove the restriction of the large alphabet size from the study of locality of codeword symbols. Indeed, the impact of the alphabet size on the locality of a codeword forms the object of focus of our work.

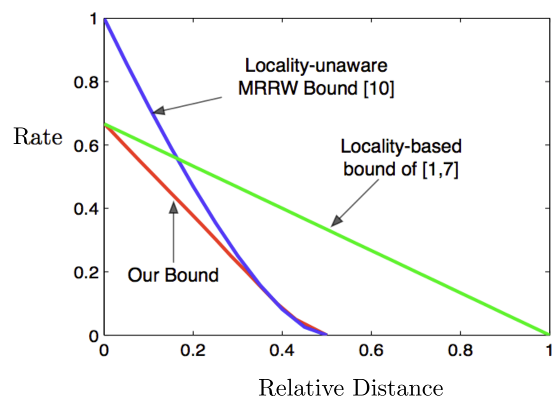

Fig. 1. A depiction of our bound through the trade-off between the *rate*, k/n and *relative distance*, d/n, for binary codes (q=2) for large values of n, with locality $r = 2$. The curves plotted are upper bounds on the achievable rates; the plot clearly demonstrates that our upper bound is better than all the previously known bounds on the rate, for a given relative distance.

A. *Our Contribution:*

The main contribution of this paper is a simple upper bound on the minimum distance on the code that is dependent on the size of the alphabet. While the technical statement of our bound is discussed later (in Theorem 1), it is worth noting here that our bound, which is applicable for any feasible alphabet-size and any feasible (n, k), is least as good as the bound of [1], [7] for all parameters. Recall that even in the absence of locality constraints, the optimum minimum distance of a code over an alphabet of a fixed size remains an open problem in general. Our main result uses this quantity - the largest possible minimum distance of an (n, k) code over a given alphabet size - albeit unknown, as a parameter to characterize a bound under locality constraints. As a consequence, our bound is stringent than the classical (locality-unaware) bounds such as the Mcliece-Rodemich-Rumsey-Welch (MRRW) bounds since they form a special case of our bound (unrestricted locality). The role of the alphabet size on the rate of the code is highlighted in the plot of Fig. 1, where we compare our bounds with existing bounds. Finally, since certain code constructions in previous works are based on multicast codes over networks [7], our result can be interpreted as the demonstration of the impact of alphabet size on the rates of multicast *network codes* for certain networks.

Notation: Sets are denoted by calligraphic letters and vectors are denoted by bold font. Consider an element $\mathbf{X} \in \mathcal{A}^n$, where $\mathcal{A}$ is an arbitrary finite set. The notation $X_i \in \mathcal{A}$ denotes the ith co-ordinate of the tuple $\mathbf{X}$. For any set $\mathcal{R} \subseteq \{1, 2, \ldots, n\}$, the notation $\mathbf{X}_{\mathcal{R}} \in \mathcal{A}^{|\mathcal{R}|}$ denotes the projection of $\mathbf{X} \in \mathcal{A}^n$ on to the co-ordinates corresponding to $\mathcal{R}$. For $\mathbf{X}, \mathbf{Y} \in \mathcal{A}^n$, the Hamming distance $\Delta_H(\mathbf{X}, \mathbf{Y})$ is the cardinality of the set $\{m : X_m \neq Y_m\}$.

II. System Model: Locally Recoverable Codes

A code $\mathcal{C}$ with length n over alphabet $\mathcal{Q}$ consists of $|\mathcal{C}|$ codewords denoted as $\mathcal{C} = \{\mathbf{X}^n(1), \mathbf{X}^n(2), \ldots, \mathbf{X}^n(|\mathcal{C}|)\}$, where $\mathbf{X}^n(i) \in \mathcal{Q}^n, \forall i$. The *dimension of the code*, denoted by k is defined as $k \stackrel{\triangle}{=} \frac{\log |\mathcal{C}|}{\log |\mathcal{Q}|}$, and the rate of the code denoted as R is defined as $R \stackrel{\triangle}{=} \frac{k}{n}$. An (n, k, d)-code over $\mathcal{Q}$ is an n length code $\mathcal{C}$ with dimension k such that the *minimum distance* is d, i.e., with

$$d = \min_{\mathbf{X}^n, \mathbf{Y}^n \in \mathcal{C}, \mathbf{X}^n \neq \mathbf{Y}^n} \Delta_H\left(\mathbf{X}^n, \mathbf{Y}^n\right).$$

We refer to $\delta \stackrel{\triangle}{=} \frac{d}{n}$ as the *relative distance* of the code.

Definition 1: An (n, k, d)-code is said to be *r-locally recoverable* if for every i such that $1 \leq i \leq n$, there exists a set $\mathcal{R}_i \subset \{1, 2, \ldots, n\} \setminus \{i\}$ with $|\mathcal{R}_i| = r$ such that for any two codewords $\mathbf{X}, \mathbf{Y}$ satisfying $X_i \neq Y_i$, we have $\mathbf{X}_{\mathcal{R}_i} \neq \mathbf{Y}_{\mathcal{R}_i}$.

Informally speaking, this means that an erasure of the ith coordinate of the codeword can be recovered by accessing the coordinates associated with $\mathcal{R}_i$. Hence any erased symbol can be recovered by probing at most r other coordinates.

III. Bound on Minimum Distance for Local Recovery

Given parameters n, d, q, let

$$k_{\text{opt}}^{(q)}(n, d) = \max \frac{\log |\mathcal{C}|}{\log q},$$

where the maximization is over all possible n-length codebooks $\mathcal{C}$ with minimum distance d, over some alphabet $\mathcal{Q}$ where $|\mathcal{Q}| = q$. Informally speaking, $k_{\text{opt}}^{(q)}(n, d)$ is the largest possible dimension of an n-length code, for a given alphabet size q and a given minimum distance d. The determination of $k_{\text{opt}}^{(q)}$ is a well known classical open problem in coding theory. It is also well known that k_{opt} satisfies the Singleton bound:

$$k_{\text{opt}}^{(q)}(n, d) \leq n - d + 1, \forall q \in \mathbb{Z}_+.$$

References [1], [7], generalized the above bound under locality constraints as Def. 1. However, it is well known that the Singleton bound is not tight in general, especially for small values of q. The goal of this paper is to derive a bound on the dimension of an r-locally recoverable code in terms of $k_{opt}^{(q)}$. Our main result is the following.

Theorem 1: For any (n, k, d) code over $\mathcal{Q}$ that is r-locally recoverable

$$k \leq \min_{t \in \mathbb{Z}_+} \Big[tr + k_{\text{opt}}^{(q)}(n - t(r+1), d) \Big], \tag{2}$$

where $q = |\mathcal{Q}|$.

Our bound applies to general (including non-linear) codes, as opposed to only linear codes. Note that, the minimizing value of t in (2), t^*, must satisfy,

$$t^* \leq \min \Big\{ \Big\lceil \frac{n}{r+1} \Big\rceil, \Big\lceil \frac{k}{r} \Big\rceil \Big\}.$$

This is true because, 1) for $t \geq \left\lceil \frac{n}{r+1} \right\rceil$, the objective function of the optimization of (2) becomes linearly growing with t; 2) for $t \geq \left\lceil \frac{k}{r} \right\rceil$, the right hand side of (2) is greater than k.

The bound of [1], [7], i.e. (1), is weaker than the bound of Theorem 1. To prove this claim, let us show that, if a (n, k, d, r)-tuple does not satisfy (1), then it will not satisfy (2).

Let, (n, k, d, r)-tuple does not satisfy (1), i.e.,

$$d > n - k - \lceil k/r \rceil + 2.$$

This sets the following chain of implications.

$$\begin{aligned}
&\min_{t \in \mathbb{Z}_+} \Big[tr + k_{\text{opt}}^{(q)}(n - t(r+1), d) \Big] \\
&\leq \lfloor (k-1)/r \rfloor r + \max\{n - \lfloor (k-1)/r \rfloor (r+1) - d + 1, 0\} \\
&= \max\{n - \lfloor (k-1)/r \rfloor - d + 1, \lfloor (k-1)/r \rfloor r\} \\
&< \max\{n - \lfloor (k-1)/r \rfloor - n + k + \lceil k/r \rceil - 2 + 1, k\} \\
&= \max\{-\lfloor (k-1)/r \rfloor + k + \lceil k/r \rceil - 1, k\} \\
&= k,
\end{aligned}$$

which means (2) is not satisfied by this tuple as well.

Notice that the above chain of implications came from plugging in the Singleton bound on $k_{\text{opt}}^{(q)}$. We shall apply bounds that are dependent on q and stronger than the Singleton bound on $k_{\text{opt}}^{(q)}$ to effectively obtain tighter bounds on (1) later in this paper. We shall first present an overview of the proof of Theorem 1. For purposes of the proof, for a given n length code $\mathcal{C}$ we define the function $H(.)$ as follows

$$H(\mathcal{I}) = \frac{\log |\{\mathbf{X}_{\mathcal{I}} : \mathbf{X} \in \mathcal{C}\}|}{\log |\mathcal{Q}|},$$

for any set $\mathcal{I} \subseteq \{1, 2, \ldots, n\}$.

Remark 1: In the language used in [7], $H(\mathcal{I})$ would denote the "entropy" associated with $\mathbf{X}_{\mathcal{I}}$. Here, the above definition is appropriate since our modeling is adversarial, i.e., we do not presuppose any distribution on the messages or the codebook (see, [6] where such assumptions have been made). However, the behavior of the function $H(.)$ is similar to the entropy function; for instance it satisfies submodularity, i.e., $H(\mathcal{I}_1) + H(\mathcal{I}_2) \geq H(\mathcal{I}_1 \cup \mathcal{I}_2) + H(\mathcal{I}_1 \cap \mathcal{I}_2)$

Theorem 1 follows from Lemma 1 and Lemma 2 stated next.

Lemma 1: Consider an (n, k, d)-code over alphabet $\mathcal{Q}$ that is r-locally recoverable. Then, $\forall\ 1 \leq t \leq k/r, t \in \mathbb{Z}\ \exists\ \mathcal{I} \subseteq \{1, 2, \ldots, n\}, |\mathcal{I}| = t(r+1)$ such that $H(\mathcal{I}) \leq tr$.

Lemma 2: Consider an (n, k, d)-code over $\mathcal{Q}$ where there exists a set $\mathcal{I} \in \{1, 2, \ldots, n\}$ such that $H(\mathcal{I}) \leq m$. Then there exists a $(n - |\mathcal{I}|, (k-m)^+, d)$ code over $\mathcal{Q}$.

Proof of Lemma 1: Consider an r-locally recoverable (n, k, d)-code. For any $i \in \{1, 2, \ldots, n\}$, let $\mathcal{R}_i$ denote the corresponding repair-set; by definition $|\mathcal{R}_i| = r$. The key idea is to construct a set $\mathcal{I}$ having the desired properties. Our construction is essentially similar to [7]; we describe our construction here for completeness. We choose

$$\mathcal{I} = \left(\bigcup_{l=1}^{t} \{a_l\} \cup \mathcal{R}_{a_l} \cup \mathcal{S}_l \right)$$

where $a_1, a_2, \ldots, a_t \in \{1, 2, \ldots, n\}$ and $\mathcal{S}_l \subset \{1, 2, \ldots, n\}, l = 1, 2, \ldots, t$ are chosen as follows:

Begin: Choose a_1 arbitrarily from $\{1, 2, \ldots, n\}$. Choose $\mathcal{S}_1$ to be the null set.

Loop: For $m = 2$ to $m = t$

Step 1: Choose a_m so that

$$a_m \notin \bigcup_{l=1}^{m-1} \{a_l\} \cup \mathcal{R}_{a_l} \cup \mathcal{S}_l$$

Step 2: Choose $\mathcal{S}_m$ to be set of $m(r+1) - \left| \{a_m\} \cup \mathcal{R}_{a_m} \cup \bigcup_{l=1}^{m-1} \{a_l\} \cup \mathcal{R}_{a_l} \cup \mathcal{S}_l \right|$ elements, arbitrarily from $\{1, 2, \ldots, n\} - \{a_m\} \cup \mathcal{R}_{a_m} \cup \bigcup_{l=1}^{m} \{a_l\} \cup \mathcal{R}_{a_l} \cup \mathcal{S}_l$.

End

This completes the construction. Note that $\mathcal{I}$ constructed above has cardinality $t(r+1)$. It remains to show that $H(\mathcal{I}) \leq tr$. We now intend to show that $H(\mathcal{I}) = H(\mathcal{I} - \{a_1, a_2, \ldots, a_t\})$ from which the desired bound would follow because of

$$H(\mathcal{I}) = H(\mathcal{I} - \{a_1, a_2, \ldots, a_t\}) \leq t(r+1) - t = tr,$$

where we have used the fact that $H(\mathcal{A}) \leq |\mathcal{A}|$ for any set $\mathcal{A}$.

We therefore intend to show a one-to-one mapping between $\{\mathbf{X}_{\mathcal{I} - \{a_1, a_2, \ldots, a_t\}}\}$ and $\{\mathbf{X}_{\mathcal{I}}\}$. In other words, suppose that $\mathbf{X}_{\mathcal{I}} \neq \hat{\mathbf{X}}_{\mathcal{I}}$, we need to prove that $\mathbf{X}_{\mathcal{I} - \{a_1, a_2, \ldots, a_t\}} \neq \hat{\mathbf{X}}_{\mathcal{I} - \{a_1, a_2, \ldots, a_t\}}$. Equivalently, suppose that $\mathbf{X}_{a_1, a_2, \ldots, a_t} \neq \hat{\mathbf{X}}_{a_1, a_2, \ldots, a_t}$, we need to prove that $\mathbf{X}_{\mathcal{I} - \{a_1, a_2, \ldots, a_t\}} \neq \hat{\mathbf{X}}_{\mathcal{I} - \{a_1, a_2, \ldots, a_t\}}$. Suppose a contradiction, i.e., suppose that $\exists, \mathbf{X}, \hat{\mathbf{X}} \in \mathcal{C}$ such that

$$\mathbf{X}_{\{a_1, a_2, \ldots, a_t\}} \neq \hat{\mathbf{X}}_{\{a_1, a_2, \ldots, a_t\}}$$
$$\mathbf{X}_{\mathcal{I} - \{a_1, a_2, \ldots, a_t\}} = \hat{\mathbf{X}}_{\mathcal{I} - \{a_1, a_2, \ldots, a_t\}}$$

Define $\mathcal{B} = \{j : \mathbf{X}_j \neq \hat{\mathbf{X}}_j\} \subseteq \{a_1, a_2, \ldots a_t\}$. Because of the definition of locality and because $\mathcal{R}_{a_i} \in \mathcal{I}$, the above conditions imply that

$$\mathcal{R}_i \cap \mathcal{B} \neq \phi, \forall i \in \mathcal{B} \tag{3}$$

In other words, the repair set associated with any element, i, in $\mathcal{B}$ should have at least one element in $\mathcal{B}$, because $\mathbf{X}_j = \hat{\mathbf{X}}_j$ for all $i \neq j$. We will show that this is a contradiction to our construction. In particular, we will throw away elements from $\mathcal{B}$ one at a time to obtain, from (3), a relation of the form $j \cap \mathcal{R}_j \neq \phi$ for some $j \in \mathcal{B}$, which is a contradiction. To keep the notation clean, we will show the proof for $\mathcal{B} = \{a_1, a_2, \ldots, a_m\}$, where $m = |\mathcal{B}|$. Our idea generalizes for arbitrary $\mathcal{B}$. By construction (Step 1), note that $a_m \notin \mathcal{R}_{a_i}, i = 1, 2, \ldots, m-1$. Therefore, a_m is not a member of the repair sets of any of the elements of $\mathcal{B}$, and (3) implies that

$$\mathcal{R}_i \cap \{a_1, a_2 \ldots, a_{m-1}\} \neq \phi, \forall i \in \{a_1, a_2, \ldots, a_{m-1}\}$$

Similarly, note that $a_{m-1} \notin \mathcal{R}_{a_i}, i = 1, 2, \ldots, m-2$ and Therefore, a_{m-1} is not a member of the repair sets of any of the elements of $\mathcal{B} - \{a_m\}$. So we get,

$$\mathcal{R}_i \cap \{a_1, a_2 \ldots, a_{m-2}\} \neq \phi, \forall i \in \{a_1, a_2, \ldots, a_{m-2}\}$$

Repeating the above procedure $m-1$ times, we get

$$\mathcal{R}_{a_1} \cap \{a_1\} \neq \phi,$$

which is a contradiction.

Proof of Lemma 2: Without loss of generality, let us assume that $\mathcal{I} = \{1, 2, \ldots, |\mathcal{I}|\}$. Consider any element $\mathbf{Z}$ of the $\mathcal{S} = \{\mathbf{X}_{\mathcal{I}} : \mathbf{X} \in \mathcal{C}\}$. Now, notice that the set of all elements of $\mathcal{C}$ which have $\mathbf{Z}$ as a "prefix" can be used to construct a codebook $\mathcal{C}(\mathbf{Z})$ of length $(n - |\mathcal{I}|)$. In particular denote

$$\tilde{\mathcal{C}}(\mathbf{Z}) = \{\mathbf{X}_{\{|\mathcal{I}|+1, |\mathcal{I}|+2, \ldots, n\}} : \mathbf{X}_{\mathcal{I}} = \mathbf{Z}\}$$

In addition, we can deduce that the codebook $\tilde{\mathcal{C}}(\mathbf{Z})$, has minimum distance d. To see this, consider $\mathbf{U}, \mathbf{V} \in \tilde{\mathcal{C}}(\mathbf{Z})$ and note that

$$\Delta_H(\mathbf{U}, \mathbf{V}) = \Delta_H((\mathbf{Z}, \mathbf{U}), (\mathbf{Z}, \mathbf{V})) \geq d \tag{4}$$

where, above we have used the fact that, by definition of $\tilde{\mathcal{C}}(\mathbf{Z})$, the tuples $(\mathbf{Z}, \mathbf{U})$ and $(\mathbf{Z}, \mathbf{V})$ are elements of $\mathcal{C}$ and therefore have a Hamming distance larger than or equal to d. Now, all we need to show is that there exists at least one $\hat{\mathbf{Z}} \in \mathcal{S}$ such that the dimension of $\tilde{\mathcal{C}}(\hat{\mathbf{Z}})$ is (at least) as large as $k-m$. This can be shown using an elementary probabilistic counting argument. Specifically, by assuming that $\mathbf{Z}$ is uniformly distributed over $\mathcal{S}$, the average value of $|\tilde{\mathcal{C}}(\mathbf{Z})|$ can be bounded as follows.

$$\begin{aligned} |\mathcal{C}| = |\mathcal{Q}|^k &= \sum_{\mathbf{Z} \in \mathcal{S}} |\tilde{\mathcal{C}}(\mathbf{Z})| \\ &= |\mathcal{S}| E\left[\tilde{\mathcal{C}}(\mathbf{Z})\right] \\ \Rightarrow E\left[\tilde{\mathcal{C}}(\mathbf{Z})\right] &= \frac{|\mathcal{Q}^k|}{|\mathcal{S}|} \\ &\geq \frac{|\mathcal{Q}|^k}{|\mathcal{Q}|^m} = |\mathcal{Q}|^{k-m} \end{aligned}$$

where, above we have used the premise of the lemma, namely $|\mathcal{S}| = |\mathcal{Q}|^{H(\mathcal{I})} \leq |\mathcal{Q}|^m$. Therefore, there is at least one $\hat{\mathbf{Z}} \in \mathcal{S}$ such that $\tilde{\mathcal{C}}(\hat{\mathbf{Z}}) \geq |\mathcal{Q}|^{k-m}$ thereby resulting in a $(n-|\mathcal{I}|, k-m, d)$ codebook over $\mathcal{Q}$. This completes the proof.

IV. Applications of Theorem 1 and Discussion

In this section, we apply classical bounds for k_{opt} to Theorem 1. To enable a clean analysis, we look at the regime where $n \to \infty$. In particular we set $R = k/n, \delta = d/n$ and obtain bounds on the trade-off between (R, δ) as r is fixed and $n \to \infty$. We first apply the Plotkin bound on k_{opt} and obtain an analytical characterization of the (R, δ) trade-off with dependence on the alphabet-size, q; in particular, we demonstrate a *distance-expansion* penalty as a result of the limit on alphabet size. Then, we use the MRRW bound for k_{opt} to numerically obtain the plot of Fig. 1.

To begin, observe that dividing the Singleton bound n and letting $n \to \infty$, it can be written as

$$R \leq 1 - \delta + o(1)$$

Similarly, the bound of [1], [7] can be written as:

$$\delta \leq 1 - \frac{rR}{r+1} + o(1).$$

$$\Rightarrow R \leq \frac{r}{r+1}(1-\delta) + o(1) \tag{5}$$

The plot of the above bound is placed in Fig. 1 for $r = 2$. The *cost* of the locality limit above therefore is the factor of $r/(r+1)$ over the Singleton bound. We are now ready to analyze the Plotkin Bound, adapted to Theorem 1.

Application of Plotkin Bound - Distance Expansion Penalty

Let us choose $t = \frac{1}{r+1}(n - \frac{d}{1-1/q})$ in Theorem 1. We have, for any (n,k,d)-code that is r-locally recoverable,

$$k \leq \frac{r}{r+1}\Big(n - \frac{d}{1-1/q}\Big) + k_{\text{opt}}^{(q)}\Big(\frac{d}{1-1/q}, d\Big)$$

It is known, from the Plotkin bound, $k_{\text{opt}}^{(q)}\Big(\frac{d}{1-1/q}, d\Big) \leq \log_q \frac{2qd}{1-1/q}$. See, for example, Sec. 2§2 of MacWilliams and Sloane [5], for a proof of this result for $q = 2$, which can be easily extended for larger alphabets. Hence,

$$k \leq \frac{r}{r+1}\Big(n - \frac{d}{1-1/q}\Big) + \log_q \frac{2qd}{1-1/q}. \tag{6}$$

Generally, this bound is better than (1). Notice that dividing the above by n and taking $n \to \infty$, we have

$$R = \frac{k}{n} \leq \frac{r}{r+1}\Big(1 - \frac{\delta}{1-1/q}\Big) + o(1),$$

whereas, observing the above, it can be noted that the effect of restricting q leads to a distance-expansion penalty of $\frac{1}{1-1/q}$, since the above bound is tantamount to shooting for a distance of $\delta/(1-1/q)$ w.r.t. (5).

Beyond the Plotkin bound

Recall that the MRRW bound is the tightest known bound for the rate-distance tradeoff in absence of locality constraints. We briefly describe an application of this bound for Theorem 1, i.e., when the locality is restricted to be equal to a number r; it is this bound that is plotted in Fig. 1. We restrict our attention to binary codes ($q = 2$) and therefore the dependence on q is dropped in the notation.

Define $R_{\text{opt}}(\delta) \stackrel{\triangle}{=} \lim_{n\to\infty} \frac{k_{\text{opt}}(n, \delta n)}{n}$ Dividing the bound of Theorem 1 by n we can get, as $n \to \infty$,

$$R \leq \min_{0 \leq x \leq r/(r+1)} x + (1 - x(1+1/r))\, R_{\text{opt}}\left(\frac{\delta}{1 - x(1+1/r)}\right)$$

where $x = tr/n$. It is instructive to observe that, setting $x = 0$ above yields classical (locality-unaware) bounds. Setting $x = R$ above and writing out the Singleton bound for R_{opt} yields the bound of (5). Therefore the above bound is superior to all the classical (locality-unaware) bounds on $R(\delta)$ and the bound of (5) since these are special cases. Using the MRRW bound $R(y) \leq \mathrm{H}_2(0.5 - \sqrt{y(1-y)}) + o(1)$ (where $\mathrm{H}_2(.)$ represents the binary entropy function), and numerically solving the optimization problem above (in a brute-force manner) yields our bounds for the rate-distance trade-offs for any given r. Deriving analytical insights for the optimization problem by application of bounds beyond the Plotkin bound is an area of future work.

A. Discussion

In this paper, we have provided upper bounds on on the rate achievable for a fixed locality, distance, and alphabet size. Constructions of locally repairable codes is an interesting open question especially relevant to practice. To understand the related issues (briefly), consider the special case of binary codes ($q = 2$) where, in absence of locality constraints, the best known simple achievable scheme comes via the Gilbert-Varshamov (GV) bound: $R \geq 1 - \mathrm{H}_2(\delta)$. A very simple construction for locally repairable codes is constructed by taking the parity check matrix of a code that achieves the GV bound and add $\lceil \frac{n}{r+1} \rceil$ rows to it; each new row has $r+1$ nonzero values and the support of all the (new) rows are disjoint. It can be readily verified that this code has a locality of r. Note that this new code has rate:

$$R \geq 1 - \mathrm{H}_2(\delta) - \frac{1}{r+1} = \frac{r}{r+1} - \mathrm{H}_2(\delta).$$

For $\delta = 0$, the above clearly meets the outer bound of (1). However, the above achievable scheme does not meet our bound for larger values of δ. For example, in the regime of Fig. 1, i.e., $r = 2$, the above bound implies that $R = 0$ for $\delta \geq \mathrm{H}_2^{-1}(2/3) \approx 0.18$. Clearly, this is not tight with our bound, where $R > 0$ as long as $\delta < 0.5$. Thus, perhaps, the most interesting open question is largest possible (relative) distance of a code with non-zero rate, for a fixed locality and alphabet size.

One of the most promising avenue to pursue to construct a locally repairable code is to consider an LDPC or low density parity check matrix code. In an LDPC code the rows of the parity check matrix have small (constant) number of non zero values. Clearly, the locality of a linear code is upper bounded by the maximum number of nonzero values in any row of the parity check matrix of the code less one. Hence, if one can construct an LDPC code with a specified rate-distance trade-off, that would be a code with small locality value.

V. Acknowledgement

This work was partially supported by a start-up grant from University of Minnesota.

References

[1] P. Gopalan, C. Huang, H. Simitci, and S. Yekhanin. On the locality of codeword symbols. *Information Theory, IEEE Transactions on*, 58(11):6925 –6934, nov. 2012.

[2] C. Huang, H. Simitci, Y. Xu, A. Ogus, B. Calder, P. Gopalan, J. Li, and S. Yekhanin. Erasure coding in windows azure storage. In *USENIX Annual Technical Conference (USENIX ATC)*, 2012.

[3] G. M. Kamath, N. Prakash, V. Lalitha, and P. V. Kumar. Codes with local regeneration. *CoRR*, abs/1211.1932, 2012. http://arxiv.org/abs/1211.1932.

[4] G. M. Kamath, N. Prakash, V. Lalitha, P. V. Kumar, N. Silberstein, A. S. Rawat, O. O. Koyluoglu, and S. Vishwanath. Explicit mbr all-symbol locality codes. *arXiv preprint arXiv:1302.0744*, 2013.

[5] F. MacWilliams and N. Sloane. The theory of error-correcting codes. 2006.

[6] A. Mazumdar, V. Chandar, and G. W. Wornell. Local recovery properties of capacity-achieving codes. *Information Theory and Applications (ITA), San Diego*, Feb 2013.

[7] D. S. Papailiopoulos and A. G. Dimakis. Locally repairable codes. In *Information Theory Proceedings (ISIT), 2012 IEEE International Symposium on*, pages 2771–2775. IEEE, 2012.

[8] D. S. Papailiopoulos, J. Luo, A. G. Dimakis, C. Huang, and J. Li. Simple regenerating codes: Network coding for cloud storage. In *INFOCOM, 2012 Proceedings IEEE*, pages 2801–2805. IEEE, 2012.

[9] A. S. Rawat, O. O. Koyluoglu, N. Silberstein, and S. Vishwanath. Optimal locally repairable and secure codes for distributed storage systems. *arXiv preprint arXiv:1210.6954*, 2012.

[10] R. Roth. *Introduction to coding theory*. Cambridge University Press, 2006.

[11] M. Sathiamoorthy, M. Asteris, D. Papailiopoulos, A. G. Dimakis, R. Vadali, S. Chen, and D. Borthakur. Xoring elephants: Novel erasure codes for big data. *arXiv preprint arXiv:1301.3791*, 2013.

[12] I. Tamo, D. S. Papailiopoulos, and A. G. Dimakis. Optimal locally repairable codes and connections to matroid theory. *arXiv preprint arXiv:1301.7693*, 2013.