CS 860 Topics in Coding Theory Nov. 10, 2025

Lecture 20

Lecturer: Elena Grigorescu *Scribe: Jacob Skitsko*

1 Local Reconstruction Codes (LRCs)

The idea is to have local codes for distributed storage. These codes were introduced by a mix of theory and applied people [HSX+12], and was implemented for significant savings in scalable cloud systems, such as Windows Azure systems.

$$\underbrace{(\text{Data}_1 \mid \text{Data}_2 \mid \ldots \mid \text{Data}_k)}_{\text{“Terabyte” Sized Data centers}}$$

Some issues these systems need to deal with:

1. Many servers crash simultaneously, e.g. when servers reach the 3 year expected lifetime. This is a rare event, but we cannot afford to lose data.
2. One or two (i.e. $O(1)$) servers get busy with updates or reboots. This happens frequently.

Historically, there have been three solutions:

1. Replicate data three times, which is clearly inefficient for the amount of data to be copied over
2. Use Reed Solomon codes. For example a $[9, 6]$ code with distance 3 replicates the data 1.5 times, so has lower storage cost with the same distance. Improving storage cost further, e.g. with a $[16, 12]$ code would require querying 12 servers to recover one entry, which is too slow for this frequent type of error.
3. Local Reconstruction Codes. Observe we know which servers fail, so we are dealing with *erasures*. This is a type of code that should protect from many simultaneous erasures, and give fast recovery from a small number of erasures.

Definition 1 *A Local Reconstruction Code (LRC) with locality ℓ is an error correcting code $C : \mathbb{F}_q^k \to \mathbb{F}_q^n$ that allows the recovery of* any *erased codeword symbol by reading $< \ell$ other codeword symbols. More generally: can locally recover even when $a \geq 1$ codeword symbols are erased.*

Example Consider *data symbols* $x_0, x_1, x_2, y_0, y_1, y_2$. We can encode *local parities* $p_x = x_0 + x_1 + x_2, p_y = y_0 + y_1 + y_2$ and *global parities* $p_0 = x_1 + y_1 + y_2, p_1 = x_0 + x_2 + y_2$ (which check over some subsets of both the x_i and the y_j data symbols). There are

$$\begin{aligned} k &= 6 \quad \text{data fragments } x_i, y_i \\ t &= 2 \quad \text{local parities } p_x, p_y \\ r &= 2 \quad \text{global parities } p_0, p_1 \end{aligned}$$

Claim 2 *Assume one erasure of a data fragment, say x_0. Then we need only three fragments to reconstruct (instead of 6, which is how many a comparable Reed Solomon code needs).*

Proof For x_0 simply read p_x, x_1, x_2 and compute what x_0 should be. ■

Question: What tradeoff can we attain for n, k, ℓ?

Theorem 3 (Singleton bound for LRC) *If $C : \mathbb{F}_q^k \to \mathbb{F}_q^n$ is an LRC with locality ℓ and minimum distance d, then*

$$n \geq k + d + \left\lceil \frac{k}{\ell} \right\rceil - 2 \,.$$

Observe this recovers the classical Singleton bound when $\ell = k$.
Question: Is Singleton for LRCs achievable in general?
Answer: It turns out yes! We'll see the easy case next, when message symbols have local recovery. Afterwards we'll consider the more interesting case, when parity symbols can also be locally recovered.

Claim 4 *The following "pyramid" codes are ℓ-locally recoverable for message symbols, and achieve the LRC Singleton bound.*

We'll define a pyramid code as follows. Start with a $[n, k, d]$ Reed Solomon code. Note $k = n - d + 1$, and we have $d - 1$ parity checks. We may assume one of these checks contains all data, say

$$p_0 = \sum_{i \in [k]} x_i \,.$$

We also have $d - 2$ parity checks p_i,

$$p_1 = x_1 + \cdots + x_\ell, \quad p_2 = x_{\ell+1} + \cdots + x_{2\ell}, \quad \ldots$$

Consider splitting our single global parity check into k/ℓ many parity checks

$$p_i' = \sum_{j \in \text{chunk } i} x_j \,.$$

You may convince yourself that the resulting code has distance d, and observe

$$N = k + (d - 2) + \frac{k}{\ell} \,.$$

Further observe each message bit can be recovered from the parity check of the corresponding chunk and the remaining symbols in the chunk.

This leads us into the following:

1.1 The Tamo-Bary Construction

Theorem 5 *There are ℓ-LRCs that recover any codeword symbol from ℓ other symbols which meets the LRC Singleton bound, over fixed field sizes of $O(n)$ size.*

Proof We'll explicitly build the parity check. Define the following parameters

$$\begin{gathered} r = \ell + 1\,, \\ q \geq n + 1 = p^t \quad \text{such that } r \mid q-1\,, \\ \gamma^{q-1} = 1\,, \quad \gamma \text{ a primitive in } \mathbb{F}_q^*\,, \\ \alpha = \gamma^{(q-1)/r}\,. \end{gathered}$$

We'll make simplifying assumptions that $r \mid n$ and $r \mid d-2$. We'll also start with some intuition: to recover a coordinate from ℓ others we need all $\ell+1$ of them to satisfy a parity check. This is said to form a "local group" of size $r = \ell + 1$.

Define the code $C := \{y : Hy = 0\}$ where

$$H = \begin{array}{l} \\ \text{local checks} \\ \\ \\ \text{global checks} \end{array} \left\{ \begin{array}{l} \\ \left(\begin{array}{cccc} \overbrace{1 \ldots 1}^{r} & 0 & \ldots & 0 \\ 0 & 1 \ldots 1 & \ldots & 0 \\ & & \ldots & \\ 0 & 0 & \ldots & 1 \ldots 1 \\ B_0 & B_1 & \ldots & B_{n/r-1} \end{array} \right) \end{array} \right\} \begin{array}{l} \\ \\ \\ \\ \frac{n}{r} \text{rows} \end{array}$$

where

$$B_i = \underbrace{\left(\begin{array}{cccc} \beta_{i,1} & \beta_{i,2} & \ldots & \beta_{i,r} \\ \beta_{i,1}^2 & \beta_{i,2}^2 & \ldots & \beta_{i,r}^2 \\ & & \ldots & \\ \beta_{i,1}^{d-1} & \beta_{i,2}^{d-1} & \ldots & \beta_{i,r}^{d-1} \end{array} \right)}_{r} \Bigg\} d-2$$

with

$$\beta_{i,j} = \gamma^i \alpha^j = \gamma^{i+j((q-1)/r)}\,, i \leq \frac{n}{r} \leq \frac{q-1}{r}\,.$$

Observe by definition that all $\beta_{i,j}$ are distinct.

Claim 6 *C has minimum distance d.*

Proof It suffices to show that every $d-1$ columns of H are linearly independent. To show that, it suffices to consider

$$H' = \begin{pmatrix} \overbrace{1 \ldots 1}^{r} & 1 \ldots 1 & \ldots & 1 \ldots 1 \\ B_0 & B_1 & \ldots & B_{n/r-1} \end{pmatrix}$$

rather than H, since that only makes showing columns are linearly independent more challenging. Since row operations don't change the column rank, and the $\beta_{i,j}$ values are all distinct, we have that every $d-1$ columns gives a $(d-1) \times (d-1)$ Vandermonde matrix. Such a matrix has full rank, as desired. ■

Claim 7 *C has locality $\ell = r - 1$*

Proof From the "$1 \dots 1$" rows of length r. ■ We'll see how to recover $a > 1$ errors later, the idea will be similar.

Claim 8 *If* $\dim(C) = k$ *then*

$$n - k = \frac{n}{r} + (d-2) - \frac{d-2}{r}\,.$$

i.e.

$$n = k + \frac{n-d+2}{r} + d - 2 \geq k + \frac{k-1}{r} + d - 2\,.$$

Proof

$$\dim(C) = n - \mathrm{rank}(H)$$

H has $n/r + (d-2)$ rows. We'll show that $(d-2)/r$ rows of H are linearly independent of other rows. Hence, $\mathrm{rank}(H) \leq n/r + (d-2) - (d-2)/r$ as desired. Indeed, let's look at row t of B_i:

$$\left(\beta_{i,1}^t, \beta_{i,2}^t, \dots, \beta_{i,r}^t\right)$$

and observe for any i

$$\beta_{i,j}^t = \gamma^{it}\alpha^{jt} = \gamma^{it} * (\gamma^{jt/r})^{q-1} = \gamma^{it}$$

where we recall we assumed $r \mid t$. So whenever r divides t, the $-t$th row of $(B_0 \dots B_{n/r-1})$ is spanned by the first n/r rows of H. This occurs for all multiplicities of r in $[1, d-2]$, i.e. for $(d-2)/r$ rows. ■

■

Question: Are some LRCs better than others, assuming the same parameters (i.e. assuming the same $d, k, n, \ell, \dots$)?

They could both correct from any pattern of $d-1$ erasures, but maybe the can correct from other patterns of $> d-1$ erasures also. These other patterns could be different for each LRC ...

1.2 Maximally Recoverable LRCs (MR LRCs)

Definition 9 *A MR LRC is an "optimal" LRC that corrects from* every *pattern correctable by any other LRC with the same parameters n, k, ℓ, d and same local groups.*

What is a correctable pattern? Note this is an interesting question beyond the minimum distance.

Observation 10 *Say $Hy = 0$, and $y_{i_1}, y_{i_2}, y_{i_3}$ are erased. Then can correct (i_1, i_2, i_3) if and only if columns i_1, i_2, i_3 in H are linearly independent (i.e. no matter what values 2 of them have, can find the third).*

Now we consider how to construct a MR LRC. Consider the following

- each group has size r
- if there are only a crashes in a local group, then can recover the group locally
- overall: can recover from a crashes per local group locally, and h crashes anywhere.

Theorem 11 *There exists MR LRCs for large enough fields.*

For general LRCs with locality ℓ, correcting from a erasures. Let $r = a + \ell$, and assume $r \mid n$. Subset $\{r(i-1)+1, \ldots, ri\} \subseteq [n]$ are local groups of size r. So, there are $g = n/r$ local groups. We define a parity check

$$H = \begin{pmatrix} A_1 & 0 & \ldots & 0 \\ 0 & A_2 & \ldots & 0 \\ & & \ldots & \\ 0 & 0 & \ldots & A_g \\ B_1 & B_2 & \ldots & B_g \end{pmatrix}$$

where $A_1, \ldots, A_g$ are $a \times r$ sized matrices over $\mathbb{F}_q$ and $B_1, \ldots, B_g$ are $h \times r$ sized matrices over $\mathbb{F}_q$. We have local parity checks: $\{A_i : i \in [g]$ is a parity check of an $[r, r-q, a+1]$ code$\}$. We'll take each A_i to be a Maximum Distance Separable Code (e.g. a Reed Solomon code). Then, we can recover from any a erasures in a local group (from $\ell = r - a$ symbols in that local group).

We have rows $(B_1, \ldots, B_J)$ giving global parity checks. $\dim(C) = a \cdot g + h$, and $\text{dist}(C) \leq a + h + 1$ if $a + h < r$, since we cannot correct $a + h + 1$ erasures in a local group. But we can correct many patterns beyond the minimum distance! Namely, exactly those patters obtained by erasing a coordinates in each local group, and h additional coordinates anywhere. A code is MR LRC if we can correct from a set of erasures as above, where the number of erasures is at least $g \cdot a + h$.

There exist such codes by the probabilistic method if $q > (ag + h)\binom{r}{ag+h}$.

References

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