On Sequential Locally Repairable Codes

Wentu Song, Chau Yuen and Kui Cai
Singapore University of Technology
and Design, Singapore
{wentu_song, yuenchau, cai_kui}@sutd.edu.sg

Kai Cai and Guangyue Han
Department of Mathematics, University
of Hong Kong, Hong Kong
{kcai, ghan}@hku.hk

I. Introduction

Distributed storage systems (DSS) store a large amount of data in a network of distributed storage nodes, and various coding techniques are introduced to maintain data reliability and improve system efficiency [1], [2]. Locally repairable codes (LRC), also known as locally recoverable codes, are a family of erasure codes which aim to reduce the disk I/O complexity for node repair. LRCs have attracted much attention recently, see [2], [3] and references therein. Roughly speaking, an LRC with locality r is an $[n, k]$ linear code where the value of each coordinate (code symbol) can be computed from the values of at most r other coordinates. Correspondingly, a *single* failed node can be repaired by downloading data from a set of at most r other nodes.

However, multiple node failures (erasures) are normal in today's large-scale distributed storage systems [6]. Hence, handling two or more simultaneous node failures is a more important issue in practice and has become a central focus in the research area of LRC [4]−[10].

Most existing works of LRC for multiple erasures consider the so-called *parallel approach*, by which each failed node (erasure) is recovered by resorting to a set of living nodes and hence the erasures can be recovered simultaneously. In this work [10], we consider a more general repair approach, called *sequential approach*, by which the erasures are recovered one by one and the already fixed erasure nodes can be used in the next round of recovering. Potentially, for the same LRC, the sequential approach can fix more erasures than the parallel approach, and hence is a better candidate of repair approach. However, due to a variety of technique difficulties, this more important approach remains far from being thoroughly understood, and there is a severe lack of code constructions and bounds for evaluation of the code performance.

In this work, we focus on high rate LRCs for sequentially recovering any $t \geq 3$ erasures; more specifically, we define the (n, k, r, t)-SLRC (Sequential Locally Repairable Code) as an $[n, k]$ linear code in which any t' ($t' \leq t$) erased code symbols can be sequentially recovered, each by at most r ($2 \leq r < k$) other symbols. We first derived an upper bound on the code rate for an (n, k, r, t)-SLRC with $t = 3$ and any $k > r \geq 2$. Then we constructed two new families of binary (n, k, r, t)-SLRCs. The first family is constructed for any positive integers r (≥ 2) and t. The second family is based on the so-called *resolvable configurations* and is constructed for any $r \geq 2$ and any odd integer $t \geq 3$. For $t \in \{2, 3\}$, the rates of these two families of codes are optimal.

A basic and important fact revealed by our study is: the sequential approach can have much better performance than the parallel approach, e.g, for the direct product of m copies of the binary $[r + 1, r]$ single-parity code, it can recover m erasures with locality r by the parallel approach, but $2^m - 1$ erasures with the same locality by the sequential approach.

This work is supported by SUTD SRG grant SRLS15095 and SUTD-MIT IDC research grant.

This abstract is condensed from [10], which is accepted by IEEE Transactions on Information Theory for publication in May 2017 and now is online available at: http://ieeexplore.ieee.org/document/7940089.

II. Sequential Locally Repairable Codes (SLRC)

For any integer $n \geq 1$, $[n] \triangleq \{1, 2, \cdots, n\}$. For any set A, $|A|$ is the size (the number of elements) of A. If $B \subseteq A$ and $|B| = t$, then B is called a t-subset of A. For any real number x, $\lceil x \rceil$ is the smallest integer greater than or equal to x.

Let $\mathcal{C}$ be an $[n, k]$ linear code over the field $\mathbb{F}$ and $i \in [n]$. A subset $R \subseteq [n]\backslash\{i\}$ is called a *recovering set* of i if there exists an $a_j \in \mathbb{F}\backslash\{0\}$ for each $j \in R$ such that $x_i = \sum_{j \in R} a_j x_j$ for all $x = (x_1, x_2, \cdots, x_n) \in \mathcal{C}$. Equivalently, there exists a $y = (y_1, y_2, \cdots, y_n) \in C^{\perp}$ such that supp$(y) = R \cup \{i\}$, where supp$(y) := \{i \in [n]; y_i \neq 0\}$ is the support of y.

Definition 1 (Sequential Locally Repairable Code): For any $E \subseteq [n]$, $\mathcal{C}$ is said to be (E, r)-recoverable if E can be sequentially indexed, i.e., $E = \{i_1, i_2, \cdots, i_{|E|}\}$, such that each $i_\ell \in E$ has a recovering set $R_\ell \subseteq \overline{E} \cup \{i_1, \cdots, i_{\ell-1}\}$ of size $|R_\ell| \leq r$, where $\overline{E} := [n]\backslash E$; $\mathcal{C}$ is called an (n, k, r, t)-*sequential locally repairable code (SLRC)* (or simply (r, t)-SLRC) if $\mathcal{C}$ is (E, r)-recoverable for each $E \subseteq [n]$ of size $|E| \leq t$, where r is called the locality of $\mathcal{C}$.

Particularly, if for each $E \subseteq [n]$ of size $|E| \leq t$ and each $i \in E$, i has a recovering set $R \subseteq \overline{E}$ of size $|R| \leq r$, then $\mathcal{C}$ is called an (n, k, r, t)-*parallel locally repairable code (PLRC)*.

III. Upper Bound on the Code Rate for $(n, k, r, 3)$-SLRC

Theorem 1: The code rate of any $(n, k, r, 3)$-SLRC satisfies

$$\frac{k}{n} \leq \left(\frac{r}{r+1}\right)^2.$$

The bound in Theorem 1 is derived by using the following graph-theoretical method: We associate each (n, k, r, t)-SLRC with a set of directed acyclic graphs, called repair graphs of the code, where the set of direct predecessors of each non-source vertex is a recovering set of it. Among all repair graphs of an (n,k,r,t)-SLRC, the ones with the least number of sources are called the *minimal repair graph*. The minimal repair graphs have some interesting structural properties, from which we can derive an upper bound on the number of sources of all repair graphs, and in turn derive an upper bound on the code rate of the corresponding (n,k,r,t)-SLRC.

A complete proof of Theorem 1 can be found in [10].

IV. Construction of (n,k,r,t)-SLRC

In this section, we present a summary of the construction of two families of binary (n,k,r,t)-SLRCs. All the details can be found in [10].

Construction 1: The first family of SLRCs is constructed by puncturing the product of m copies of the binary $[r+1,r]$ single-parity code with respect to a proper subset of index, where $r\geq 2$ and $1\leq t\leq 2^m-1$ for some positive integer m. Specifically, for each integer s such that $0\leq s\leq 2^m-1$, let

$$\Gamma_s^{(m)}=\{\alpha\in\mathbb{Z}_{r+1}^m;\mathbf{U}^{(m)}(\alpha)=\text{supp}_m(s)\},$$

where $\mathbb{Z}_{r+1}=\{0,1,\cdots,r\}$ and $\mathbb{Z}_{r+1}^m$ is the Cartesian product of m copies of $\mathbb{Z}_{r+1}$, and

$$\Omega_s^{(m)}=\bigcup_{\ell=0}^{s}\Gamma_\ell^{(m)},$$

where for each $\alpha=(i_m,i_{m-1},\cdots,i_1)\in\mathbb{Z}_{r+1}^m$,

$$\mathbf{U}^{(m)}(\alpha)=\{\ell\in[m];i_\ell=r\},$$

and

$$\mathbf{T}^{(m)}(\alpha)=\{\ell\in[m];i_\ell\in\mathbb{Z}_r=\{0,1,\cdots,r-1\}\}.$$

And $\text{supp}_m(s)$ is the support of the unique m-digit binary representation of s. For $t=2^m-1$, let $\mathcal{C}_{2^m-1}^{(m)}$ be the product of m copies of the $[r+1,r]$ binary code; For $t<2^m-1$, let $\mathcal{C}_t^{(m)}$ be the punctured code of $\mathcal{C}_{2^m-1}^{(m)}$ with respect to $\Omega_t^{(m)}$. Then $\mathcal{C}_t^{(m)}$ is an (n,k,r,t)-SLRC and has code rate

$$\frac{k}{n}=\frac{1}{\sum_{s=0}^{t}\frac{1}{r^{|\text{supp}_m(s)|}}}.$$

Construction 2: The second family of SLRCs is constructed using the so-called *resolvable configurations*.

A (k_{t-1},b_r) *configuration* is a pair $(X,\mathcal{A})$, where X is a set of k elements, called points, and $\mathcal{A}$ is a collection of b subsets of X, called lines, such that

(1) Each line contains r points;
(2) Each point belongs to $t-1$ lines;
(3) Every pair of distinct points belong to at most one line.

The configuration $(X,\mathcal{A})$ is called *resolvable*, if further

(4) All lines in $\mathcal{A}$ can be partitioned into t–1 parallel classes, where a parallel class is a set of lines that partition X.

The basic idea of Construction 2 is as follows. Let $X=[k]$ and $\mathcal{A}_1=\{A_1,\cdots,A_s\}$, $\cdots$, $\mathcal{A}_{t-1}=\{A_{(t-2)s+1},\cdots,A_b\}$ be the $t-1$ parallel classes of $(X,\mathcal{A})$. We partition $[s]$ into $\lceil\frac{s}{r}\rceil$ nonempty subsets $B_1,\cdots,B_{\lceil\frac{s}{r}\rceil}$ such that $|B_i|\leq r$ for all $i\in\{1,\cdots,\lceil\frac{s}{r}\rceil\}$ and let $W=(w_{i,j})$ be a $\lceil\frac{s}{r}\rceil\times b$ matrix with

$$w_{i,j}=\begin{cases}1, & \text{if } j\in B_i\\ 0, & \text{otherwise.}\end{cases}$$

Let M be the incidence matrix of $(X,\mathcal{A})$ and $\mathcal{C}$ be a binary linear code with parity check matrix

$$H=\begin{pmatrix} M & I_b & O_{b\times\lceil\frac{s}{r}\rceil}\\ O_{\lceil\frac{s}{r}\rceil\times k} & W & I_{\lceil\frac{s}{r}\rceil}\end{pmatrix},$$

where I_ℓ denotes the $\ell\times\ell$ identity matrix and $O_{\ell\times\ell'}$ denotes the $\ell\times\ell'$ all-zero matrix for any positive integers ℓ and ℓ'. Then for any positive integer r and any odd integer $t\geq 3$, $\mathcal{C}$ is an (n,k,r,t)-SLRC and has the code rate of

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

V. Discussions and Future Work

We leave it as an open problem to determine the optimal code rate of (n,k,r,t)-SLRCs for general t, i.e., $t\geq 5$. However, we have the following conjecture:

Conjecture: An achievable upper bound of the code rate of (n,k,r,t)-SLRCs has the following form:

$$\frac{k}{n}\leq\left(1+\sum_{i=1}^{m}\frac{a_i}{r^i}\right)^{-1},$$

where $m=\lceil\log_r k\rceil$, and $a_i\geq 0, i=1,2,\cdots,m$, are integers such that $\sum_{i=1}^m a_i=t$.

It is worthwhile to point out that in a recent work [9], an achievable upper bound on the rate of (n,k,r,t)-SLRC matching the above conjecture for any $r,t\geq 3$ is proven.

References

[1] A. G. Dimakis,B. Godfrey,Y.Wu,M. J.Wainwright, and K. Ramchandran, "Network coding for distributed storage systems," *IEEE Trans. Inf. Theory*, vol. 56, no. 9, pp. 4539-4551, Sep. 2010.

[2] D. S. Papailiopoulos and A. G. Dimakis, "Locally repairable codes," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Cambridge, USA, Jul. 2012, pp. 2771-2775.

[3] P. Gopalan, C. Huang, H. Simitci, and S. Yekhanin, "On the locality of codeword symbols," *IEEE Trans. Inf. Theory*, vol. 58, no. 11, pp. 6925-6934, Nov. 2012.

[4] L. Pamies-Juarez, H. D. L. Hollmann, and F. Oggier, "Locally repairable codes with multiple repair alternatives," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Istanbul, Turkey, Jul. 2013, pp. 892-896.

[5] N. Prakash, G. M. Kamath, V. Lalitha, and P. V. Kumar, "Optimal linear codes with a local-error-correction property," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Cambridge, USA, Jul. 2012, pp. 2776-2780.

[6] A. Wang and Z. Zhang, "Repair locality with multiple erasure tolerance," *IEEE Trans. Inf. Theory*, vol. 60, no. 11, pp. 6979-6987, Nov. 2014.

[7] A. S. Rawat, D. S. Papailiopoulos, A. G. Dimakis, and S. Vishwanath, "Locality and Availability in Distributed Storage," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Honolulu, HI, USA, June. 2014, pp. 681-685.

[8] A. S. Rawat, A. Mazumdar, and S. Vishwanath, "Cooperative local repair in distributed storage," *EURASIP Journal on Advances in Signal Processing*, pp. 1-17, 2015:107.

[9] S. B. Balaji, G. R. Kini, and P. V. Kumar, "A Tight Rate Bound and a Matching Construction for Locally Recoverable Codes with Sequential Recovery From Any Number of Multiple Erasures," Available at: https://arxiv.org/abs/1611.08561

[10] W. Song, Kai Cai, C. Yuen, Kui Cai, and G. Han, "On sequential locally repairable codes," *IEEE Trans. Inf. Theory*, 2017, online available at: http://ieeexplore.ieee.org/document/7940089