Local Reconstruction Codes: A Class of MDS-PIR Capacity-Achieving Codes

Citation for the original published paper (version of record):
Kumar, S., Lin, H., Rosnes, E. et al (2018). Local Reconstruction Codes: A Class of MDS-PIR Capacity-Achieving Codes. 2018 IEEE Information Theory Workshop, ITW 2018. http://dx.doi.org/10.1109/ITW.2018.8613503

(article starts on next page)

Local Reconstruction Codes: A Class of MDS-PIR Capacity-Achieving Codes

Siddhartha Kumar[†], Hsuan-Yin Lin[†], Eirik Rosnes[†], and Alexandre Graell i Amat[‡]
[†]Simula UiB, N–5020 Bergen, Norway
[‡]Department of Electrical Engineering, Chalmers University of Technology, SE–41296 Gothenburg, Sweden

***Abstract*—We prove that a class of distance-optimal local reconstruction codes (LRCs), an important family of repair-efficient codes for distributed storage systems, achieve the maximum distance separable private information retrieval capacity for the case of noncolluding nodes. This particular class of codes includes Pyramid codes and other LRCs proposed in the literature.**

I. Introduction

Private information retrieval (PIR) deals with the scenario where a user wants to retrieve a data item from a database without letting the database know the identity of the requested item. PIR was first introduced in the computer science literature by Chor *et al.* in [1], where the authors considered that the database is replicated across n servers (nodes) and presented a PIR protocol that efficiently achieves privacy in the presence of a single spy node. In [1], the efficiency of the PIR protocol was measured in terms of upload and download cost.

With the advent of distributed storage systems (DSSs), where data is stored in a distributed fashion over a number of nodes using a storage code rather than simply replicated, the concept of PIR has gained traction in the information theory community. As typically the size of the data items stored is much larger compared to the size of the queries sent to the nodes, the upload cost is negligible compared to the download cost [2]. Thus, under the information-theoretic formulation, the efficiency of a PIR protocol, referred to as the PIR rate, is measured in terms of download cost. More precisely, the PIR rate is defined as the ratio between the requested file size and the total amount of downloaded data. The maximum PIR rate over all PIR protocols is the PIR capacity.

The authors in [3] were the first to introduce PIR protocols for DSSs in the information theory community, assuming that data is stored using two explicit linear codes. In [2], an upper bound on the PIR rate for a certain class of linear PIR protocols was given. For the case of replicated data and a single spy node, commonly known as the noncolluding case, Sun and Jafar [4] derived the PIR capacity and presented a PIR capacity-achieving scheme. Also, for the noncolluding case, Banawan and Ulukus [5] derived the maximum achievable PIR rate for the more general scenario where data is stored in the DSS using a maximum distance separable (MDS) code and presented a scheme that achieves it. As the underlying storage code is an MDS code, such a maximum achievable PIR rate is usually referred to as the MDS-PIR capacity.

The MDS-PIR capacity depends on the code rate of the underlying MDS storage code and the number of files stored in the DSS. In [6], a PIR protocol for MDS-coded data that achieves the *asymptotic* MDS-PIR capacity when the number of files tends to infinity was presented. In [7], the authors presented a PIR protocol for the case where the underlying storage code can be an arbitrary linear code and numerically showed that the proposed protocol can achieve the asymptotic MDS-PIR capacity even if the underlying storage code is non-MDS. With some abuse of language, we refer to such codes as MDS-PIR capacity-achieving codes. While the aformentioned protocols assume that nodes in the DSS do not collude, [8]–[11] proposed PIR schemes for the case of colluding nodes.

In a DSS, the storage code is used not just to achieve reliability against node failures, but also to repair failed nodes. Although MDS codes are optimal in terms of storage overhead (for a given rate), they are characterized by a large repair locality, i.e., the repair of a failed node requires contacting a large number of nodes. Thus, with focus on repair locality, several code constructions such as Pyramid codes [12], locally repairable codes [13], and local reconstruction codes (LRCs) [14] have been proposed. Such codes follow a similar design philosophy, and we refer to them globally as LRCs. In [7], it was shown numerically that, interestingly, the asymptotic MDS-PIR capacity for the case of noncolluding nodes can be achieved for some Pyramid codes.

In this paper, we go a step further and formally prove that an important class of repair-efficient storage codes, namely a class of distance-optimal LRCs, are MDS-PIR capacity-achieving codes in the noncolluding case. This implies that one does not need to sacrifice on the repair locality to achieve the MDS-PIR capacity.

II. Definitions and Preliminaries

Throughout the paper we use the following notation. We represent the set of a consecutive integers as $\mathbb{N}_a \triangleq \{1,\ldots,a\}$, while $\mathbb{N}_{a:b} \triangleq \{a,\ldots,b\}$ represents the set of integers from a to b. We use calligraphic upper case, bold upper case, and bold lower case letters to denote sets, matrices, and vectors, respectively. As an example, $\mathcal{X}$, $\boldsymbol{X}$, and $\boldsymbol{x}$ represent a set, matrix, and a vector, respectively. The identity matrix of order a is denoted by $\boldsymbol{I}_a$, and $(\boldsymbol{X}_1|\ldots|\boldsymbol{X}_a)$ denotes the horizontal

The work of S. Kumar, H.-Y. Lin, and E. Rosnes was partially funded by the Research Council of Norway (grant 240985/F20). A. Graell i Amat was supported by the Swedish Research Council under grant #2016-04253.

concatenation of matrices $\boldsymbol{X}_1,\dots,\boldsymbol{X}_a$. A submatrix of $\boldsymbol{X}$ that is restricted in columns by the set $\mathcal{J}$ is denoted by $\boldsymbol{X}|_{\mathcal{J}}$, and the rank of $\boldsymbol{X}$ is denoted by $\operatorname{rank}(\boldsymbol{X})$. $\mathcal{C}$ denotes an $[n,k]$ linear code of block length n, dimension k, and minimum Hamming distance $d_{\mathsf{min}}^{\mathcal{C}}$ over the Galois field GF(q). A generator matrix of $\mathcal{C}$ is denoted by $\boldsymbol{G}^{\mathcal{C}}$, while $\boldsymbol{H}^{\mathcal{C}}$ denotes a parity-check matrix. $\mathcal{C}|_{\mathcal{J}}$ is the punctured code obtained from $\mathcal{C}$ by restricting the code coordinates to the indices in $\mathcal{J}$. A set of coordinates of $\mathcal{C}$, $\mathcal{J} \subseteq \mathbb{N}_n$, of size k is said to be an *information set* if and only if $\boldsymbol{G}^{\mathcal{C}}|_{\mathcal{J}}$ is invertible. With some abuse of language, we sometimes interchangeably refer to binary vectors as erasure patterns under the implicit assumption that the ones represent erasures.

We consider a DSS that stores f files $\boldsymbol{X}^{(1)},\dots,\boldsymbol{X}^{(f)}$, where $\boldsymbol{X}^{(m)} = (x_{i,j}^{(m)})$, $m \in \mathbb{N}_f$, can be seen as a $\beta \times k$ matrix over GF(q^ℓ), with $\beta,k,\ell \in \mathbb{N}$. Let $\boldsymbol{x}_i^{(m)}$ denote the i-th row of $\boldsymbol{X}^{(m)}$. Each $\boldsymbol{x}_i^{(m)}$ is encoded by an $[n,k]$ code $\mathcal{C}$ over GF(q) into a length-n codeword $\boldsymbol{c}_i^{(m)} = \big(c_{i,1}^{(m)},\dots,c_{i,n}^{(m)}\big)$, where $c_{i,j}^{(m)} \in \text{GF}(q^\ell)$, $j \in \mathbb{N}_n$, is stored on the j-th node. The symbols are stored in the order of increasing m and secondly in the order of increasing i (see [10, Sec. III]).

A. MDS-PIR Capacity-Achieving Codes

For a given number of files f stored using an $[n,k]$ MDS code, the MDS-PIR capacity [5, Thm. 1] is $\mathsf{C}_f = \frac{1-k/n}{1-(k/n)^f}$. We refer to C_f as the *finite* MDS-PIR capacity, as it depends on the number of files. When the number of files grows very large, i.e., $f \to \infty$, the MDS-PIR capacity reduces to $\mathsf{C}_\infty = 1-\frac{k}{n}$, which we refer to as the *asymptotic* MDS-PIR capacity.

We denote by $\mathsf{R}_f(\mathcal{C})$ the PIR rate of a PIR scheme that uses code $\mathcal{C}$ as the underlying storage code to store f files. The following theorem gives a condition for the existence of MDS-PIR capacity-achieving codes (under Protocols 1 and 2 presented by the authors in [10]).[1]

Theorem 1: Consider a DSS that uses an $[n,k]$ code $\mathcal{C}$ to store f files. If there exists a binary $n \times n$ matrix $\boldsymbol{E}$ of row and column weight $n-k$ such that each row is an erasure pattern that is correctable by $\mathcal{C}$, then $\mathcal{C}$ achieves the finite MDS-PIR capacity C_f (under Protocol 1 in [10]), i.e., $\mathsf{R}_f(\mathcal{C}) = \mathsf{C}_f$, and the asymptotic MDS-PIR capacity C_∞ (under Protocol 2 in [10]), i.e., $\mathsf{R}_\infty(\mathcal{C}) = \mathsf{C}_\infty$.

In Sections III and IV, we prove that for a class of distance-optimal (r,δ) information locality codes [15], an important class of LRCs, such an $\boldsymbol{E}$ exists, and hence this class of codes is MDS-PIR capacity-achieving.

B. Local Reconstruction Codes

LRCs are a family of codes characterized by their low repair locality, i.e., in order to repair a failed node, only a relatively low number of nodes need to be contacted. In particular, we consider *information locality* codes, which are systematic codes whose focus is to reduce the repair locality of systematic nodes (i.e., nodes that store systematic code symbols) [12]–[15]. Formally, they are defined as follows.

Definition 1 ((r,δ) information locality code [15, Def. 2]): An $[n,k]$ code $\mathcal{C}$ is said to be an (r,δ) information locality code if there exist L_{c} punctured codes $\mathcal{C}_j \triangleq \mathcal{C}|_{\mathcal{S}_j}$ of $\mathcal{C}$ with column coordinate set $\mathcal{S}_j \subset \mathbb{N}_n$ for $j \in \mathbb{N}_{L_{\mathsf{c}}}$. Furthermore, $\{\mathcal{C}|_{\mathcal{S}_j}\}_{j\in\mathbb{N}_{L_{\mathsf{c}}}}$ must satisfy the following conditions:

1) $|\mathcal{S}_j| \leq r+\delta-1$, $\forall\, j \in \mathbb{N}_{L_{\mathsf{c}}}$,
2) $d_{\mathsf{min}}^{\mathcal{C}_j} \geq \delta$, $\forall\, j \in \mathbb{N}_{L_{\mathsf{c}}}$, and
3) $\operatorname{rank}\big(\boldsymbol{G}^{\mathcal{C}}|_{\bigcup_j \mathcal{S}_j}\big) = k$.

In other words, Definition 1 says that there are L_{c} local codes in $\mathcal{C}$ each having a block length of at most $r+\delta-1$, minimum Hamming distance at least δ, and the union of all coordinate sets of the local codes contains an information set. The overall code $\mathcal{C}$ has $d_{\mathsf{min}}^{\mathcal{C}} \leq n-k+1-(\lceil k/r\rceil-1)(\delta-1)$ and can repair up to $\delta-1$ systematic nodes by contacting r storage nodes. Codes that achieve the upper bound on the d_{min} are known as distance-optimal (r,δ) information locality codes and have the following structure.

Definition 2 (Distance-optimal (r,δ) information locality code [15, Thm. 2.2]): Let $r \mid k$ such that $L_{\mathsf{c}} = k/r$. An (r,δ) information locality code $\mathcal{C}$ as defined in Definition 1 is distance-optimal if:

1) Each local code $\mathcal{C}|_{\mathcal{S}_j}$, $j \in \mathbb{N}_{L_{\mathsf{c}}}$, is an $[r+\delta-1,r]$ MDS code defined by a parity-check matrix $\boldsymbol{H}^{\mathcal{C}|_{\mathcal{S}_j}} = (\boldsymbol{P}_j|\boldsymbol{I}_{\delta-1})$ of dimensions $(\delta-1)\times(r+\delta-1)$ and minimum Hamming distance $d_{\mathsf{min}}^{\mathcal{C}|_{\mathcal{S}_j}} = \delta$.
2) The sets $\{\mathcal{S}_j\}_{j\in\mathbb{N}_{L_{\mathsf{c}}}}$ are disjoint, i.e., $\mathcal{S}_j \cap \mathcal{S}_{j'} = \emptyset$ for all $j,j' \in \mathbb{N}_{L_{\mathsf{c}}}$, $j \neq j'$.
3) The code $\mathcal{C}$ has a parity-check matrix of the form

$$\boldsymbol{H} = \left(\begin{array}{ccccccc|c} \boldsymbol{P}_1 & \boldsymbol{I}_{\delta-1} & & & & & & \\ & & \boldsymbol{P}_2 & \boldsymbol{I}_{\delta-1} & & & & \\ & & & & \ddots & & & \\ & & & & & \boldsymbol{P}_{L_{\mathsf{c}}} & \boldsymbol{I}_{\delta-1} & \\ \hline \boldsymbol{M}_1 & \boldsymbol{0} & \boldsymbol{M}_2 & \boldsymbol{0} & \cdots & \boldsymbol{M}_{L_{\mathsf{c}}} & \boldsymbol{0} & \boldsymbol{I}_a \end{array}\right) \tag{1}$$

where the matrices $\boldsymbol{M}_1,\dots,\boldsymbol{M}_{L_{\mathsf{c}}}$ are arbitrary matrices in GF(q) of dimensions $(n-L_{\mathsf{c}}(r+\delta-1))\times r$, and $a \triangleq n-L_{\mathsf{c}}(r+\delta-1)$.

For ease of exposition, we refer to the local parities as the parity symbols that take part in the local codes, while the parity symbols that are not part of the L_{c} local codes are referred to as global parity symbols. According to Definition 2, there exist $n-L_{\mathsf{c}}(r+\delta-1)$ global parities and $L_{\mathsf{c}}(\delta-1)$ local parities. We partition the coordinates of these parities into $L+1$ sets, where $L \triangleq \big\lfloor \frac{n}{r+\delta-1} \big\rfloor$. For $j \in \mathbb{N}_{L+1}$, we have

$$\mathcal{P}_j = \begin{cases} \{(j-1)n_{\mathsf{c}}+r+1,\dots,jn_{\mathsf{c}}\} & \text{if } j \in \mathbb{N}_{L_{\mathsf{c}}}, \\ \{(j-1)n_{\mathsf{c}}+1,\dots,jn_{\mathsf{c}}\} & \text{if } j \in \mathbb{N}_{L_{\mathsf{c}}+1:L}, \\ \{Ln_{\mathsf{c}}+1,\dots,n\} & \text{if } j = L+1, \end{cases} \tag{2}$$

where $n_{\mathsf{c}} \triangleq r+\delta-1$ is the block length of each local code. The set $\mathcal{P}_j$, $j \in \mathbb{N}_{L_{\mathsf{c}}}$, represents the coordinates of the local parities of the j-th local code $\mathcal{C}_j$. The remaining sets $\mathcal{P}_j$, $j \in \mathbb{N}_{L_{\mathsf{c}}+1:L+1}$, represent the coordinates of the global parities of $\mathcal{C}$. As such, the set $\mathcal{P} = \bigcup_{j=1}^{L+1} \mathcal{P}_j$ represents the parity coordinates of $\mathcal{C}$.

[1] Protocol 2 in [10] was originally introduced in [7].

III. Distance-Optimal Local Reconstruction Codes are MDS-PIR Capacity-Achieving

Consider an $[n,k]$ distance-optimal (r,δ) information locality code (see Definition 2) for which the $(n'-k)\times n'$ matrix

$$\left(\begin{matrix} \boldsymbol{P}_1 & \boldsymbol{P}_2 & \cdots & \boldsymbol{P}_{L_\mathsf{c}} \\ \boldsymbol{M}_1 & \boldsymbol{M}_2 & \cdots & \boldsymbol{M}_{L_\mathsf{c}} \end{matrix} \,\middle|\, \boldsymbol{I}_{n'-k} \right) \triangleq \boldsymbol{H}^{\mathsf{MDS}} \tag{3}$$

is the parity-check matrix of an $[n',k]$ MDS code over $\mathrm{GF}(q)$, where $n'=n-(L_\mathsf{c}-1)(\delta-1)$.[2] For such a class of codes, we give an explicit construction of the matrix $\boldsymbol{E}$ in order to design the PIR protocol.

Recall that $L=\left\lfloor\frac{n}{n_\mathsf{c}}\right\rfloor$, $n_\mathsf{c}=r+\delta-1$, and let $\bar{r}\triangleq n \bmod n_\mathsf{c}$. We consider

$$\boldsymbol{E}=\begin{pmatrix} \boldsymbol{E}_{1,1} & \boldsymbol{E}_{1,2} & \cdots & \boldsymbol{E}_{1,L+1} \\ \vdots & \vdots & \vdots & \vdots \\ \boldsymbol{E}_{L+1,1} & \boldsymbol{E}_{L+1,2} & \cdots & \boldsymbol{E}_{L+1,L+1} \end{pmatrix}$$

having $(L+1)^2$ submatrices $\boldsymbol{E}_{l,h}$, $l,h\in\mathbb{N}_{L+1}$. For any $l,h\in\mathbb{N}_L$, the submatrices $\boldsymbol{E}_{l,h}$ have dimensions $n_\mathsf{c}\times n_\mathsf{c}$, $\boldsymbol{E}_{l,L+1}$ has dimensions $n_\mathsf{c}\times\bar{r}$, $\boldsymbol{E}_{L+1,h}$ has dimensions $\bar{r}\times n_\mathsf{c}$, and $\boldsymbol{E}_{L+1,L+1}$ has dimensions $\bar{r}\times\bar{r}$. We denote by $\boldsymbol{e}_i^{(l)}$, $l\in\mathbb{N}_{L+1}$, the i-th row of $\bigl(\boldsymbol{E}_{l,1}|\ldots|\boldsymbol{E}_{l,L+1}\bigr)$. The coordinates of $\boldsymbol{e}_i^{(l)}$ represent the coordinates of the code $\mathcal{C}$ defined by its parity-check matrix in (1). Furthermore, each row vector is subdivided into $L+1$ subvectors $\boldsymbol{e}_{i,j}^{(l)}$, $j\in\mathbb{N}_{L+1}$, as

$$\boldsymbol{e}_i^{(l)}=(e_{i,1}^{(l)},\ldots,e_{i,n}^{(l)})=(\boldsymbol{e}_{i,1}^{(l)},\ldots,\boldsymbol{e}_{i,L}^{(l)},\boldsymbol{e}_{i,L+1}^{(l)}).$$

The subvectors $\boldsymbol{e}_{i,1}^{(l)},\ldots,\boldsymbol{e}_{i,L}^{(l)}$ are of length n_c, while $\boldsymbol{e}_{i,L+1}^{(l)}$ is of length $\bar{r}$. Correspondingly, we can think about $\boldsymbol{E}$ as partitioned into $L+1$ column partitions, where the first L_c partitions correspond to the L_c local codes and the remaining $L+1-L_\mathsf{c}$ partitions correspond to global parities (see also (2)). We can write $\boldsymbol{E}$ as

$$\boldsymbol{E}\triangleq\begin{pmatrix} \boldsymbol{e}_1^{(1)} \\ \vdots \\ \boldsymbol{e}_{n_\mathsf{c}}^{(1)} \\ \vdots \\ \boldsymbol{e}_{n_\mathsf{c}}^{(L)} \\ \boldsymbol{e}_1^{(L+1)} \\ \vdots \\ \boldsymbol{e}_{\bar{r}}^{(L+1)} \end{pmatrix} = \begin{pmatrix} \boldsymbol{e}_{1,1}^{(1)} & \boldsymbol{e}_{1,2}^{(1)} & \cdots & \boldsymbol{e}_{1,L}^{(1)} & \boldsymbol{e}_{1,L+1}^{(1)} \\ \vdots & \vdots & \cdots & \vdots & \vdots \\ \boldsymbol{e}_{n_\mathsf{c},1}^{(1)} & \boldsymbol{e}_{n_\mathsf{c},2}^{(1)} & \cdots & \boldsymbol{e}_{n_\mathsf{c},L}^{(1)} & \boldsymbol{e}_{n_\mathsf{c},L+1}^{(1)} \\ \vdots & \vdots & \cdots & \vdots & \vdots \\ \boldsymbol{e}_{n_\mathsf{c},1}^{(L)} & \boldsymbol{e}_{n_\mathsf{c},2}^{(L)} & \cdots & \boldsymbol{e}_{n_\mathsf{c},L}^{(L)} & \boldsymbol{e}_{n_\mathsf{c},L+1}^{(L)} \\ \boldsymbol{e}_{1,1}^{(L+1)} & \boldsymbol{e}_{1,2}^{(L+1)} & \cdots & \boldsymbol{e}_{1,L}^{(L+1)} & \boldsymbol{e}_{1,L+1}^{(L+1)} \\ \vdots & \vdots & \cdots & \vdots & \vdots \\ \boldsymbol{e}_{\bar{r},1}^{(L+1)} & \boldsymbol{e}_{\bar{r},2}^{(L+1)} & \cdots & \boldsymbol{e}_{\bar{r},L}^{(L+1)} & \boldsymbol{e}_{\bar{r},L+1}^{(L+1)} \end{pmatrix}.$$

We refer to the set of rows $\boldsymbol{e}_1^{(l)},\ldots,\boldsymbol{e}_{n_\mathsf{c}}^{(l)}$ as the l-th row partition of $\boldsymbol{E}$.

For convenience, we divide $\boldsymbol{E}$ into four submatrices $\tilde{\boldsymbol{E}}$, $\boldsymbol{W}$, $\boldsymbol{Z}$, and $\boldsymbol{O}$ defined as

$$\tilde{\boldsymbol{E}}\triangleq\begin{pmatrix} \boldsymbol{e}_{1,1}^{(1)} & \boldsymbol{e}_{1,2}^{(1)} & \cdots & \boldsymbol{e}_{1,L}^{(1)} \\ \boldsymbol{e}_{2,1}^{(1)} & \boldsymbol{e}_{2,2}^{(1)} & \cdots & \boldsymbol{e}_{2,L}^{(1)} \\ \vdots & \vdots & \cdots & \vdots \\ \boldsymbol{e}_{n_\mathsf{c},1}^{(L)} & \boldsymbol{e}_{n_\mathsf{c},2}^{(L)} & \cdots & \boldsymbol{e}_{n_\mathsf{c},L}^{(L)} \end{pmatrix}, \boldsymbol{Z}\triangleq\begin{pmatrix} \boldsymbol{e}_{1,L+1}^{(1)} \\ \boldsymbol{e}_{2,L+1}^{(1)} \\ \vdots \\ \boldsymbol{e}_{n_\mathsf{c},L+1}^{(L)} \end{pmatrix},$$

$$\boldsymbol{W}\triangleq\begin{pmatrix} \boldsymbol{e}_{1,1}^{(L+1)} & \boldsymbol{e}_{1,2}^{(L+1)} & \cdots & \boldsymbol{e}_{1,L}^{(L+1)} \\ \vdots & \vdots & \cdots & \vdots \\ \boldsymbol{e}_{\bar{r},1}^{(L+1)} & \boldsymbol{e}_{\bar{r},2}^{(L+1)} & \cdots & \boldsymbol{e}_{\bar{r},L}^{(L+1)} \end{pmatrix}, \boldsymbol{O}\triangleq\begin{pmatrix} \boldsymbol{e}_{1,L+1}^{(L+1)} \\ \vdots \\ \boldsymbol{e}_{\bar{r},L+1}^{(L+1)} \end{pmatrix},$$

where $\tilde{\boldsymbol{E}}$ is an $n_\mathsf{c}L\times n_\mathsf{c}L$ matrix, having L^2 submatrices $\boldsymbol{E}_{l,h}$, $l,h\in\mathbb{N}_L$.

In the following, we give a systematic construction of $\boldsymbol{E}$ such that it is $(n-k)$-regular.[3] The construction involves two steps.

a) **Initialize matrices $\tilde{\boldsymbol{E}}$, $\boldsymbol{W}$, $\boldsymbol{Z}$, and $\boldsymbol{O}$.** Matrix $\boldsymbol{Z}$ is initialized to the all-zero matrix of dimensions $n_\mathsf{c}L\times\bar{r}$. Matrices $\boldsymbol{W}$ and $\boldsymbol{O}$ are initialized by setting $e_{i,j}^{(L+1)}=1$, $i\in\mathbb{N}_{\bar{r}}$, $j\in\mathcal{P}=\bigcup_{j'=1}^{L+1}\mathcal{P}_{j'}$, where $\mathcal{P}$ corresponds to the parity coordinates of $\mathcal{C}$ and the sets $\mathcal{P}_{j'}$ are defined in Section II-B (see (2)). Let $m=\left\lfloor\frac{n-k}{L}\right\rfloor$, $m_1=m+1$, $\rho_1=\cdots=\rho_t=m_1$, and $\rho_{t+1}=\cdots=\rho_L=m$, where $t=(n-k)\bmod L$. Matrix $\tilde{\boldsymbol{E}}$ is initialized with the structure

$$\tilde{\boldsymbol{E}}=\begin{pmatrix} \boldsymbol{\pi}_1 & \boldsymbol{\pi}_2 & \cdots & \boldsymbol{\pi}_L \\ \boldsymbol{\pi}_L & \boldsymbol{\pi}_1 & \cdots & \boldsymbol{\pi}_{L-1} \\ \vdots & \vdots & \cdots & \vdots \\ \boldsymbol{\pi}_2 & \boldsymbol{\pi}_3 & \cdots & \boldsymbol{\pi}_1 \end{pmatrix}, \tag{4}$$

where each matrix entry $\boldsymbol{\pi}_l$, $l\in\mathbb{N}_L$, is a ρ_l-regular square matrix of order n_c. Notice that due to the structure in (4), $\tilde{\boldsymbol{E}}$ has row and column weight equal to $n-k$, and subsequently each row of $\boldsymbol{E}$ has weight $n-k$. Note also that the columns of $\boldsymbol{E}$ with coordinates in $\mathcal{P}_j$, $j\in\mathbb{N}_L$, have column weight $n-k+\bar{r}$, while the columns with coordinates in $\mathcal{P}_{L+1}$ have weight $\bar{r}$.

b) **Swapping elements between $\tilde{\boldsymbol{E}}$ and $\boldsymbol{Z}$.** The swapping of elements is performed iteratively with $\bar{r}$ iterations. For each iteration, in the i-th row partition and j-th column partition, we consider a set of row coordinates $\mathcal{R}_j^{(i)}$ of size $|\mathcal{P}_j|$ from which $s_j^{(i)}\in\{0,1\}$ ones from columns with coordinates in $\mathcal{P}_j$, $j\in\mathbb{N}_L$, are swapped with zeroes in the corresponding rows of $\boldsymbol{Z}$. For convenience, we define $\boldsymbol{s}^{(i)}=(s_1^{(i)},\ldots,s_L^{(i)})$ and require that $\sum_{j=1}^{L}s_j^{(i)}=1$. Note that $\mathcal{R}_j^{(i)}$ and $\boldsymbol{s}^{(i)}$ depend on the iteration number. We describe the procedure for iteration $j'\in\mathbb{N}_{\bar{r}}$. For the first row partition, select $\boldsymbol{s}^{(1)}$ with $s_j^{(1)}=1$ and $s_z^{(1)}=0$, $\forall\, z\in\mathbb{N}_L\backslash\{j\}$, for some $j\in\mathbb{N}_L$, such that if $j\in\mathbb{N}_{L_\mathsf{c}}$ there exist $\delta-1$ rows in the first row partition and j-th column partition such that their individual weight is strictly larger than $\delta-1$, and otherwise if $j\in\mathbb{N}_{L_\mathsf{c}+1:L}$, all rows in the first row partition and j-th column partition must have weight larger than or equal to $\max(1,m-(\delta-1))$. This will ensure that the resulting erasure patterns after the swap (as described next) are correctable by $\mathcal{C}$ (see Section IV). Such an $\boldsymbol{s}^{(1)}$ will also always exist for all $\bar{r}$ iterations as shown in Section IV below. Next, for all $i'\in\mathcal{R}_j^{(1)}$ and

[2] Examples of codes that satisfy (3) are Pyramid codes, the LRCs in [14], and codes from the parity-splitting construction of [15].

[3] For ease of notation, we will refer to a matrix with constant row weight, constant column weight, and constant row and column weight equal to a as an a-row regular, a-column regular, and a-regular matrix, respectively.

$p \in \mathcal{P}_j$ (where different p's are chosen for different i''s, and index j is such that $s_j^{(1)} = 1$) the one at coordinate (i', p) of $\tilde{\boldsymbol{E}}$ is swapped with a zero at coordinate (i', j') of $\boldsymbol{Z}$ (this corresponds to coordinate $(i', n_{\mathsf{c}}L + j')$ of $\boldsymbol{E}$). Then, for the remaining row partitions $i = 2, \dots, L$, consider $\boldsymbol{s}^{(i)}$ to be the $(i-1)$-th right cyclic shift of $\boldsymbol{s}^{(1)}$ and repeat the swapping procedure for the first row partition. Due to the specific selection of $\boldsymbol{s}^{(1)}$, the corresponding erasure patterns for all row partitions after the swaps are correctable by $\mathcal{C}$ (see Section IV). Note that we have performed $\sum_{j=1}^{L} |P_j| = n - k - \bar{r}$ swaps from the columns of $\tilde{\boldsymbol{E}}$ with coordinates in the set $\cup_{j=1}^{L} \mathcal{P}_j$ to the j'-th column of $\boldsymbol{Z}$. Thus, each column in $\cup_{j=1}^{L} \mathcal{P}_j$ has column weight $n - k + \bar{r} - 1$ and the $(n_{\mathsf{c}}L + j')$-th column has column weight $n - k - \bar{r} + \bar{r} = n - k$. Letting $j' = j' + 1$ and repeating the above procedure $\bar{r}$ times ensures $\boldsymbol{E}$ to be $(n-k)$-regular.

This completes the construction of $\boldsymbol{E}$, which has row and column weight $n - k$. In the following theorem, we show that each row of $\boldsymbol{E}$ (considered as an erasure pattern) can be corrected by any code from the class of distance-optimal (r, δ) information locality codes whose parity-check matrices are as in (1) and are compliant with (3). Thus, this class of codes is MDS-PIR capacity-achieving.

Theorem 2: An $[n, k]$ distance-optimal (r, δ) information locality code $\mathcal{C}$ with parity-check matrix as in (1) and satisfying (3) is an MDS-PIR capacity-achieving code.

Proof: A sketch of the proof is given in Section IV. ■

In the following, we present an example to illustrate the construction of the matrix $\boldsymbol{E}$.

Example 1: Consider an $[n = 7, k = 4]$ Pyramid code $\mathcal{C}$ that is constructed from an $[n' = 6, 4]$ Reed-Solomon code over $\mathrm{GF}(2^3)$ with parity-check matrices

$$\boldsymbol{H}^{\mathcal{C}} = \begin{pmatrix} z^3 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & z^3 & z & 1 & 0 \\ z^4 & 1 & 0 & z^5 & z^5 & 0 & 1 \end{pmatrix}$$

and

$$\boldsymbol{H}^{\mathsf{MDS}} = \begin{pmatrix} z^3 & 1 & z^3 & z & 1 & 0 \\ z^4 & 1 & z^5 & z^5 & 0 & 1 \end{pmatrix},$$

respectively, where z denotes a primitive element of $\mathrm{GF}(2^3)$. It is easy to see that $\mathcal{C}$ is a distance-optimal $(r = 2, \delta = 2)$ information locality code. We have $n_{\mathsf{c}} = 3$, $L = L_{\mathsf{c}} = 2$, and $\bar{r} \triangleq n \bmod n_{\mathsf{c}} = 1$. Since $\rho_1 = 2$ and $\rho_2 = 1$, we get

$$\tilde{\boldsymbol{E}} = \begin{pmatrix} \boldsymbol{\pi}_1 & \boldsymbol{\pi}_2 \\ \boldsymbol{\pi}_2 & \boldsymbol{\pi}_1 \end{pmatrix} = \left(\begin{array}{ccc|ccc} 1 & 1 & 0 & 1 & 0 & 0 \\ 0 & 1 & 1 & 0 & 1 & 0 \\ 1 & 0 & 1 & 0 & 0 & 1 \\ \hline 1 & 0 & 0 & 1 & 1 & 0 \\ 0 & 1 & 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 & 0 & 1 \end{array}\right), \ \boldsymbol{Z} = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix},$$

where $\boldsymbol{\pi}_1$ is a 2-regular 3×3 matrix and $\boldsymbol{\pi}_2$ is picked as the identity matrix. The set of parity coordinates is $\mathcal{P} = \{3, 6, 7\}$, and we set $e_{1,3}^{(3)} = e_{1,6}^{(3)} = e_{1,7}^{(3)} = 1$. As such, we get

$$\boldsymbol{W} = \begin{pmatrix} 0 & 0 & 1 & 0 & 0 & 1 \end{pmatrix} \text{ and } \boldsymbol{O} = \begin{pmatrix} 1 \end{pmatrix}.$$

This completes Step a) of the construction above. Note that each row of $\boldsymbol{E}$ has now weight 3. The second step of the procedure (Step b)) is as follows. Consider the first iteration, $j' = 1$. In the first row partition we choose $\boldsymbol{s}^{(1)} = (s_1^{(1)} = 1, s_2^{(1)} = 0)$. Taking $\mathcal{R}_1^{(1)} = \{2\}$, we do the swap between the coordinates $(i' = 2, p = 3 \in \mathcal{P}_1)$ and $(i', 6 + j')$. For the second row partition we have $\boldsymbol{s}^{(2)} = (0, 1)$ which is a right cyclic shift of $\boldsymbol{s}^{(1)}$. Taking $\mathcal{R}_2^{(2)} = \{6\}$, we do the swap between the coordinates $(i' = 6, p = 6 \in \mathcal{P}_2)$ and $(i', 6 + j')$. Thus, we have

$$e_{2,3}^{(1)} = 0, \ e_{2,7}^{(1)} = 1,$$
$$e_{3,6}^{(2)} = 0, \ e_{3,7}^{(2)} = 1.$$

Since $\bar{r} = 1$, this completes Step b), which results in

$$\boldsymbol{E} = \left(\begin{array}{ccc|ccc|c} 1 & 1 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 & 0 & 1 & 0 \\ \hline 1 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 & 1 \\ \hline 0 & 0 & 1 & 0 & 0 & 1 & 1 \end{array}\right).$$

The entries in red indicate the swapped values within each row. It can easily be verified that each row of $\boldsymbol{E}$ is an erasure pattern that is correctable by code $\mathcal{C}$.

IV. Sketch of Proof of Theorem 2

In the following, we give a sketch of the proof of Theorem 2. A more detailed proof is presented in [10, App. F]. According to Theorem 1, to prove that a distance-optimal (r, δ) information locality code $\mathcal{C}$ is MDS-PIR capacity-achieving, it is sufficient to prove that there exists an $(n-k)$-regular matrix $\boldsymbol{E}$ whose rows represent erasure patterns that are correctable by $\mathcal{C}$. The construction of such a matrix $\boldsymbol{E}$, provided in Section III, involves two steps as follows.

a) The submatrices $\tilde{\boldsymbol{E}}$, $\boldsymbol{W}$, $\boldsymbol{Z}$, and $\boldsymbol{O}$ are systematically constructed such that the row weight constraint is satisfied.

b) Swap elements in certain rows of matrices $\tilde{\boldsymbol{E}}$ and $\boldsymbol{Z}$ in order to meet the column weight constraint of $\boldsymbol{E}$.

The proof is a two-step procedure. First, we prove that all rows in $\boldsymbol{E}$ after Step a) are correctable by $\mathcal{C}$. Secondly, we prove that the swaps in certain rows in Step b) ensure that the resulting rows are correctable erasure patterns. We will make use of the following lemma.

Lemma 1: Let $\mathcal{C}$ be an $[n, k]$ distance-optimal (r, δ) information locality code consisting of L_{c} local codes and with parity-check matrix as in (1). Additionally, it adheres to the condition in (3). Then, $\mathcal{C}$ can simultaneously correct $\delta - 1 + \nu_j$ erasures, $\nu_j \geq 0$, in each local code $\mathcal{C}|_{\mathcal{S}_j}$ provided that the number of global parities available is at least $\nu_1 + \cdots + \nu_{L_{\mathsf{c}}}$.

Proof: The proof is given in [10, App. F]. ■

Consider the erasure patterns in the first row partition of $\boldsymbol{E}$ after Step a). Each of these patterns has $\nu_j = \rho_j - (\delta - 1)$, $j \in \mathbb{N}_{L_{\mathsf{c}}}$, erasures occurring in the coordinates corresponding to the local code $\mathcal{C}|_{\mathcal{S}_j}$ that cannot be corrected locally. Furthermore, the number of nonerased global parities is equal to $\gamma_{\mathsf{tot}} + \bar{r}$, where γ_{tot} is the total number of nonerased global parity coordinates present in the column partitions $L_{\mathsf{c}}+1,\dots,L$. It can be shown that $\sum_{j=1}^{L_{\mathsf{c}}} \nu_j \leq \gamma_{\mathsf{tot}} + \bar{r}$ (see [10, proof of Lem. 8]). From Lemma 1, all erasures in the L_{c} local codes are correctable. This enables the code to correct the remaining erasures at the coordinates of $\mathcal{C}$ in the set $\cup_{j=L_{\mathsf{c}}+1}^{L}\mathcal{P}_j$. Thus, the erasure patterns in the first row partition of $\boldsymbol{E}$ after Step a) are correctable. Through induction, one can prove that the erasure patterns in the remaining $L-1$ row partitions are also correctable. The erasure patterns in $(\boldsymbol{W}|\boldsymbol{O})$ are correctable by $\mathcal{C}$ as they pertain to the local and global parity symbols. This completes the first part of the proof.

We now address the second part of the proof. Note that the columns with coordinates in $\mathcal{P}_j$, $j \in \mathbb{N}_L$, have column weight $n-k+\bar{r}$ after Step a). Step b) involves the swapping of one entries from these coordinates with zero entries in the column coordinates of $\boldsymbol{Z}$. The swapping is done to ensure that the column weight of the columns indexed by $\mathcal{P}_j$, $j \in \mathbb{N}_L$, is reduced to $n-k$, while those of the columns of $\boldsymbol{Z}$ are increased to $n-k-\bar{r}$. Since $\boldsymbol{O}$ is an all-one matrix, the columns of $\boldsymbol{E}$ with indices in $\mathcal{P}_{L+1}$ have also weight $n-k$. It is possible to show that such a swapping always exists. Overall, the resulting matrix $\boldsymbol{E}$ is $(n-k)$-column regular. To ensure that the erasure patterns are correctable, we use Lemma 1. For each row,

$$\sum_{j=1}^{L_{\mathsf{c}}} \nu_j \leq \gamma_{\mathsf{tot}} + \gamma_{L+1}, \tag{5}$$

where γ_{L+1} is number of nonerased parity coordinates in column partition $L+1$, must hold. Clearly, if for a certain row of $(\tilde{\boldsymbol{E}} \mid \boldsymbol{Z})$ a one from a column from a column partition in $\mathbb{N}_{L_{\mathsf{c}}+1:L}$ (corresponding to $\tilde{\boldsymbol{E}}$) is swapped with a zero in a column from partition $L+1$ (corresponding to $\boldsymbol{Z}$), then the resulting erasure pattern is still correctable by $\mathcal{C}$ as (5) is still valid. On the other hand, for $j \in \mathbb{N}_{L_{\mathsf{c}}}$, if for a certain row of $(\tilde{\boldsymbol{E}} \mid \boldsymbol{Z})$ a one from the j-th column partition is swapped with a zero in the $(L+1)$-th column partition, then such a row is still a correctable erasure pattern provided that $\nu_j > 0$ before the swap. This is easy to see as the swapping procedure reduces ν_j and γ_{L+1} by one. Thus, (5) is still satisfied. From the aforementioned arguments and the fact that each row of any row partition of $(\tilde{\boldsymbol{E}} \mid \boldsymbol{Z})$ has at most $\bar{r}$ swaps of ones occurring from the set of $\mathbb{N}_L$ column partitions and zeroes from the $(L+1)$-th partition, it follows that the swaps according to Step b) are valid over all $\bar{r}$ iterations (valid in the sense that the resulting erasure patterns are correctable by $\mathcal{C}$) if

$$\sum_{j=1}^{L_{\mathsf{c}}} \nu_j + \sum_{j=L_{\mathsf{c}}+1}^{L} (m-(\delta-1)) \geq \bar{r}. \tag{6}$$

This is a counting argument, where according to Step b) for each row we restrict swapping ν_j coordinates in the j-th column partition, $j \in \mathbb{N}_{L_{\mathsf{c}}}$, and $m-(\delta-1)$ coordinates in the column partitions $\mathbb{N}_{L_{\mathsf{c}}+1:L}$ to make sure (following the arguments above) that the resulting erasure pattern after the swap is correctable by $\mathcal{C}$. Using that $\nu_j = \rho_j - (\delta-1)$ and $t = n-k-mL$, it can be shown that the left hand side of (6) can be lowerbounded by $n-k-L(\delta-1)$ when $t \leq L_{\mathsf{c}}$. Setting $n = \bar{r} + L(r+\delta-1)$ and $k = L_{\mathsf{c}}r$, it follows that (6) reduces to $L \geq L_{\mathsf{c}}$. By definition, this is always true. When $t > L_{\mathsf{c}}$, the left hand side of (6) is equal to $n-k-L(\delta-1)+L_{\mathsf{c}}-t$, and it can be shown that this is always larger than or equal to $\bar{r}$, since $t \leq L$ (details omitted for brevity). It follows that for all $\bar{r}$ iterations and for all row partitions in the systematic procedure in Step b) there exists a valid swap such that the resulting erasure patterns are still correctable by $\mathcal{C}$.

V. Conclusion

We formally proved that a class of distance-optimal LRCs, an important class of codes used in DSSs, are MDS-PIR capacity-achieving codes. The considered class of codes includes Pyramid codes and other constructions of LRCs given in the literature.

References

[1] B. Chor, O. Goldreich, E. Kushilevitz, and M. Sudan, "Private information retrieval," in *Proc. 36th IEEE Symp. Found. Comp. Sci. (FOCS)*, Milwaukee, WI, Oct. 1995, pp. 41–50.

[2] T. H. Chan, S.-W. Ho, and H. Yamamoto, "Private information retrieval for coded storage," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Hong Kong, China, Jun. 2015, pp. 2842–2846.

[3] N. B. Shah, K. V. Rashmi, and K. Ramchandran, "One extra bit of download ensures perfectly private information retrieval," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Honolulu, HI, Jun./Jul. 2014, pp. 856–860.

[4] H. Sun and S. A. Jafar, "The capacity of private information retrieval," *IEEE Trans. Inf. Theory*, vol. 63, no. 7, pp. 4075–4088, Jul. 2017.

[5] K. Banawan and S. Ulukus, "The capacity of private information retrieval from coded databases," *IEEE Trans. Inf. Theory*, vol. 64, no. 3, pp. 1945–1956, Mar. 2018.

[6] R. Tajeddine and S. El Rouayheb, "Private information retrieval from MDS coded data in distributed storage systems," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Barcelona, Spain, Jul. 2016, pp. 1411–1415.

[7] S. Kumar, E. Rosnes, and A. Graell i Amat, "Private information retrieval in distributed storage systems using an arbitrary linear code," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Aachen, Germany, Jun. 2017, pp. 1421–1425.

[8] R. Freij-Hollanti, O. W. Gnilke, C. Hollanti, and D. A. Karpuk, "Private information retrieval from coded databases with colluding servers," *SIAM J. Appl. Algebra Geom.*, vol. 1, no. 1, pp. 647–664, Nov. 2017.

[9] R. Freij-Hollanti, O. W. Gnilke, C. Hollanti, A.-L. Horlemann-Trautmann, D. Karpuk, and I. Kubjas, "*t*-private information retrieval schemes using transitive codes," Dec. 2017, arXiv:1712.02850v1 [cs.IT].

[10] S. Kumar, H.-Y. Lin, E. Rosnes, and A. Graell i Amat, "Achieving maximum distance separable private information retrieval capacity with linear codes," Dec. 2017, arXiv:1712.03898v4 [cs.IT].

[11] R. Tajeddine, O. W. Gnilke, and S. El Rouayheb, "Private information retrieval from MDS coded data in distributed storage systems," 2018, *IEEE Trans. Inf. Theory*, to appear.

[12] C. Huang, M. Chen, and J. Li, "Pyramid codes: Flexible schemes to trade space for access efficiency in reliable data storage systems," in *Proc. IEEE Int. Symp. Net. Comp. Appl. (NCA)*, Cambridge, MA, Jul. 2007, pp. 79–86.

[13] M. Sathiamoorthy, M. Asteris, D. Papailiopoulos, A. G. Dimakis, R. Vadali, S. Chen, and D. Borthakur, "XORing elephants: Novel erasure codes for big data," in *Proc. 39th Very Large Data Bases Endowment (VLDB)*, Trento, Italy, Aug. 2013, pp. 325–336.

[14] C. Huang, H. Simitci, Y. Xu, A. Ogus, B. Calder, P. Gopalan, J. Li, and S. Yekhanin, "Erasure coding in Windows Azure storage," in *Proc. USENIX Annual Tech. Conf.*, Boston, MA, Jun. 2012.

[15] G. M. Kamath, N. Prakash, V. Lalitha, and P. V. Kumar, "Codes with local regeneration and erasure correction," *IEEE Trans. Inf. Theory*, vol. 60, no. 8, pp. 4637–4660, Aug. 2014.