



An MDS-PIR Capacity-Achieving Protocol for Distributed Storage Using Non-MDS Linear Codes

Downloaded from: <https://research.chalmers.se>, 2024-11-09 07:34 UTC

Citation for the original published paper (version of record):

Lin, H., Kumar, S., Rosnes, E. et al (2018). An MDS-PIR Capacity-Achieving Protocol for Distributed Storage Using Non-MDS Linear Codes. IEEE International Symposium on Information Theory - Proceedings, 2018-June: 966-970. <http://dx.doi.org/10.1109/ISIT.2018.8437804>

N.B. When citing this work, cite the original published paper.

© 2018 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, or reuse of any copyrighted component of this work in other works.

(article starts on next page)

An MDS-PIR Capacity-Achieving Protocol for Distributed Storage Using Non-MDS Linear Codes

Hsuan-Yin Lin[†], Siddhartha Kumar[†], 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 propose a private information retrieval (PIR) protocol for distributed storage systems with noncolluding nodes where data is stored using an arbitrary linear code. An expression for the PIR rate, i.e., the ratio of the amount of retrieved data per unit of downloaded data, is derived, and a necessary and a sufficient condition for codes to achieve the maximum distance separable (MDS) PIR capacity are given. The necessary condition is based on the generalized Hamming weights of the storage code, while the sufficient condition is based on code automorphisms. We show that cyclic codes and Reed-Muller codes satisfy the sufficient condition and are thus MDS-PIR capacity-achieving.

I. INTRODUCTION

Private information retrieval (PIR) was first addressed by Chor *et al.* in the computer science community [1]. A PIR protocol allows the users to privately retrieve any requested file stored in a set of servers (referred to as nodes in the sequel) without revealing to the nodes which file is actually being downloaded. The efficiency of a PIR protocol is measured in terms of the overall communication cost, defined as the sum of the upload and the download cost, and the goal is to design a protocol that minimizes it. Recently, PIR for coded distributed storage systems (DSSs) where data is encoded by a linear code and then stored across nodes has attracted a great deal of attention [2]–[7].

Assuming that the size of the files stored in the DSS is much larger than the number of files stored, the upload cost is small compared to the download cost [4], [5], and thus it can be ignored. The PIR rate is then defined as the amount of information retrieved per downloaded symbol and is the measure of efficiency used in the information theory community. For the so-called uncoded PIR problem, where the system can be seen as a coded DSS using a repetition code, the authors in [8], [9] derived an exact expression for the maximum possible PIR rate over all possible PIR protocols, i.e., the *PIR capacity*. A closed-form expression for the coded PIR capacity when no nodes collude was derived in [10] for the case where data is encoded by a maximum distance separable (MDS) code and then stored. The PIR capacity in this case is usually referred to as *MDS-PIR capacity*. Most of the earlier works focus on studying PIR schemes for DSSs where data is stored using an MDS code. A PIR protocol for DSSs where data is stored using an arbitrary linear code was considered in

[6] for the case of noncolluding nodes, and it was shown that the asymptotic MDS-PIR capacity (assuming that an infinite number of files are stored) can be achieved even when the underlying code is non-MDS. PIR with linear codes for the case of colluding nodes was addressed in [7], [11], [12]. A conjecture for the MDS-PIR capacity in the colluding case was stated in [7], but disproved for 2 files in [13]. However, in other cases (e.g., in the asymptotic case), it is still open.

In this paper, we propose a PIR protocol for DSSs where data is stored using an arbitrary linear code for the case of noncolluding nodes. Furthermore, we investigate which classes of codes can achieve the MDS-PIR capacity with the proposed protocol. Specifically, we derive an expression for the PIR rate by exploiting a number of coordinate sets containing information sets of the underlying storage code, and define a class of MDS-PIR capacity-achieving codes, which includes MDS codes. We also provide a necessary and a sufficient condition for a code to achieve the MDS-PIR capacity with the given PIR protocol. The necessary condition is connected to the generalized Hamming weights of the storage code, while the sufficient condition is related to code automorphisms. We show that cyclic codes and Reed-Muller (RM) codes satisfy the sufficient condition and are thus MDS-PIR capacity-achieving codes. In the following, all proofs are omitted due to lack of space. The proofs can be found in the extended version [11].

Notation: We use \mathbb{N} for the set of all positive integers, $\mathbb{N}_a \triangleq \{1, 2, \dots, a\}$, and $\mathbb{N}_{n_1:n_2} \triangleq \{n_1, n_1+1, \dots, n_2\}$ for two positive integers $n_1 \leq n_2$, $n_1, n_2 \in \mathbb{N}$. Vectors are denoted by lower case bold letters, matrices by upper case bold letters, and sets by calligraphic upper case letters, e.g., \mathbf{x} , \mathbf{X} , and \mathcal{X} denote a vector, a matrix, and a set, respectively. \mathcal{C} will denote a linear code over the finite field $\text{GF}(q)$. We use the customary code parameters $[n, k]$ or $[n, k, d_{\min}^{\mathcal{C}}]$ to denote a code \mathcal{C} of blocklength n , dimension k , and minimum Hamming distance $d_{\min}^{\mathcal{C}}$. $(\cdot)^{\top}$ represents the transpose of its argument, while $(\mathbf{X}_1 | \dots | \mathbf{X}_a)$ represents the horizontal concatenation of the matrices $\mathbf{X}_1, \dots, \mathbf{X}_a$, all with the same number of rows. The function $H(\cdot)$ represents the entropy of its argument and $\chi(\mathbf{x})$ denotes the support of a vector \mathbf{x} . Subscripts may be omitted if the arguments we refer to are contextually unambiguous.

II. PRELIMINARIES AND SYSTEM MODEL

In this section, we first review some notions in coding theory and then give the system model and the privacy model.

This work was partially funded by the Research Council of Norway (grant 240985/F20) and the Swedish Research Council (grant #2016-04253).

A. Definitions

Definition 1: Let \mathcal{C} be an $[n, k]$ code with generator matrix $\mathbf{G}^{\mathcal{C}}$, and denote by $\mathbf{G}^{\mathcal{C}}|_{\mathcal{I}}$ the matrix consisting of the columns of $\mathbf{G}^{\mathcal{C}}$ indexed by \mathcal{I} . A set of coordinates of \mathcal{C} , $\mathcal{I} \subseteq \mathbb{N}_n$, of size k is said to be an *information set* if and only if $\mathbf{G}^{\mathcal{C}}|_{\mathcal{I}}$ is invertible.

Definition 2: Let \mathcal{D} be a subcode of an $[n, k]$ code \mathcal{C} . The *support* of \mathcal{D} is defined as

$$\chi(\mathcal{D}) \triangleq \{j \in \mathbb{N}_n : \exists \mathbf{x} = (x_1, \dots, x_n) \in \mathcal{D}, x_j \neq 0\}.$$

Definition 3 (Generalized Hamming weight [14]): The s -th generalized Hamming weight of an $[n, k]$ code \mathcal{C} , denoted by $d_s^{\mathcal{C}}$, $s \in \mathbb{N}_k$, is defined as the cardinality of the smallest support of an s -dimensional subcode of \mathcal{C} , i.e.,

$$d_s^{\mathcal{C}} \triangleq \min\{|\chi(\mathcal{D})| : \mathcal{D} \text{ is an } [n, s] \text{ subcode of } \mathcal{C}\}.$$

B. System Model

We consider a DSS that stores f files $\mathbf{X}^{(1)}, \dots, \mathbf{X}^{(f)}$, where each file $\mathbf{X}^{(m)} = (x_{i,j}^{(m)})$, $m \in \mathbb{N}_f$, can be seen as a $\beta \times k$ matrix over $\text{GF}(q)$ with $\beta, k \in \mathbb{N}$. Each file is encoded using a linear code as follows. Let $\mathbf{x}_i^{(m)} = (x_{i,1}^{(m)}, \dots, x_{i,k}^{(m)})$, $i \in \mathbb{N}_\beta$, be a message vector corresponding to the i -th row of $\mathbf{X}^{(m)}$. Each $\mathbf{x}_i^{(m)}$ is encoded by an $[n, k]$ code \mathcal{C} over $\text{GF}(q)$ into a length- n codeword $\mathbf{c}_i^{(m)} = (c_{i,1}^{(m)}, \dots, c_{i,n}^{(m)})$, where $c_{i,j}^{(m)} \in \text{GF}(q)$, $j \in \mathbb{N}_n$. The βf generated codewords $\mathbf{c}_i^{(m)}$ are then arranged in the array $\mathbf{C} = ((\mathbf{C}^{(1)})^\top \dots | (\mathbf{C}^{(f)})^\top)^\top$ of dimensions $\beta f \times n$, where $\mathbf{C}^{(m)} = ((\mathbf{c}_1^{(m)})^\top \dots | (\mathbf{c}_\beta^{(m)})^\top)^\top$ for $m \in \mathbb{N}_f$. For a given column j of \mathbf{C} , we denote the column vector $(c_{1,j}^{(m)}, \dots, c_{\beta,j}^{(m)})^\top$ as a coded chunk pertaining to file $\mathbf{X}^{(m)}$. The f coded chunks in column j are stored on the j -th storage node, $j \in \mathbb{N}_n$.

C. Privacy Model

We consider a DSS where a node may act as spy. It is assumed that the remaining nonspy nodes do not collaborate with the spy node. The scenario of one spy node is analogous to having a system with no colluding nodes. To retrieve file $\mathbf{X}^{(m)}$ from the DSS, the user sends a $d \times \beta f$ query matrix $\mathbf{Q}^{(l)}$ over $\text{GF}(q)$, for some integer d , to the l -th node for all $l \in \mathbb{N}_n$. In response to the received query, node l sends the response vector \mathbf{r}_l , which is a deterministic function of $\mathbf{Q}^{(l)}$ and the code symbols stored in the node l , back to the user.

Definition 4: Consider a DSS with n nodes storing f files in which a single node acts as spy. A user who wishes to retrieve the m -th file sends the queries $\mathbf{Q}^{(l)}$, $l \in \mathbb{N}_n$, to the storage nodes, which return the responses \mathbf{r}_l . This scheme achieves perfect information-theoretic PIR if and only if

$$\text{Privacy: } \quad \mathbb{H}(m | \mathbf{Q}^{(l)}) = \mathbb{H}(m), \forall l \in \mathbb{N}_n; \quad (1a)$$

$$\text{Recovery: } \quad \mathbb{H}(\mathbf{X}^{(m)} | \mathbf{r}_1, \dots, \mathbf{r}_n) = 0. \quad (1b)$$

Queries satisfying (1a) ensure that the spy node is not able to determine which file is being downloaded by the user. The recovery constraint in (1b) ensures that the user is able to recover the requested file from the responses sent by the DSS.

Definition 5: The PIR rate of a PIR protocol, denoted by \mathbb{R} , is the amount of information retrieved per downloaded symbol, i.e., $\mathbb{R} \triangleq \frac{\beta k}{D}$, where D is the total number of downloaded symbols for the retrieval of a single file.

We will write $\mathbb{R}(\mathcal{C})$ to highlight that the PIR rate depends on the underlying storage code \mathcal{C} . The maximum achievable PIR rate is the PIR capacity. It was shown in [10] that for the noncolluding case and for a given number of files f stored using an $[n, k]$ MDS code, the MDS-PIR capacity, denoted by \mathbb{C}_f , is

$$\mathbb{C}_f \triangleq \frac{n-k}{n} \left[1 - \left(\frac{k}{n} \right)^f \right]^{-1}. \quad (2)$$

III. CAPACITY-ACHIEVING PIR PROTOCOL

In this section, we propose a PIR protocol that achieves the MDS-PIR capacity for the scenario of noncolluding nodes. The protocol is inspired by the protocol introduced in [10].

A. PIR Achievable Rate Matrix

In [8], the concept of exploiting *side information* for PIR problems was introduced. By side information we mean additional redundant symbols not related to the requested file but downloaded by the user in order to maintain privacy. These symbols can be exploited by the user to retrieve the requested file from the responses of the storage nodes. In [10, Sec. V.A], it was shown that for a $[5, 3, 3]$ MDS storage code, the side information is decoded by utilizing other code coordinates forming an information set in the code array. For instance, the authors chose the $\nu = 5$ information sets $\mathcal{I}_1 = \{1, 2, 3\}$, $\mathcal{I}_2 = \{1, 4, 5\}$, $\mathcal{I}_3 = \{2, 3, 4\}$, $\mathcal{I}_4 = \{1, 2, 5\}$, and $\mathcal{I}_5 = \{3, 4, 5\}$ of the $[5, 3, 3]$ MDS code in their PIR achievable scheme. Observe that in $\{\mathcal{I}_i\}_{i \in \mathbb{N}_5}$ each coordinate of the $[5, 3, 3]$ code appears exactly $\kappa = 3$ times. This motivates the following definition.

Definition 6: Let \mathcal{C} be an arbitrary $[n, k]$ code. A $\nu \times n$ binary matrix $\mathbf{\Lambda}_{\kappa, \nu}(\mathcal{C})$ is said to be a *PIR achievable rate matrix* for \mathcal{C} if the following conditions are satisfied.

- 1) The Hamming weight of each column of $\mathbf{\Lambda}_{\kappa, \nu}$ is κ , and
- 2) for each matrix row $\boldsymbol{\lambda}_i$, $i \in \mathbb{N}_\nu$, $\chi(\boldsymbol{\lambda}_i)$ always contains an information set.

In other words, each coordinate j of \mathcal{C} , $j \in \mathbb{N}_n$, appears exactly κ times in $\{\chi(\boldsymbol{\lambda}_i)\}_{i \in \mathbb{N}_\nu}$, and every set $\chi(\boldsymbol{\lambda}_i)$ contains an information set.

Lemma 1: If a matrix $\mathbf{\Lambda}_{\kappa, \nu}(\mathcal{C})$ exists for an $[n, k]$ code \mathcal{C} , then we have

$$\frac{\kappa}{\nu} \geq \frac{k}{n},$$

where equality holds if $\chi(\boldsymbol{\lambda}_i)$, $i \in \mathbb{N}_\nu$, are all information sets.

Example 1: Consider the $[5, 3, 2]$ code with generator matrix

$$\mathbf{G} = \begin{pmatrix} 1 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 1 \end{pmatrix}.$$

One can easily verify that

$$\mathbf{\Lambda}_{2,3} = \begin{pmatrix} 0 & 1 & 1 & 1 & 1 \\ 1 & 0 & 0 & 1 & 1 \\ 1 & 1 & 1 & 0 & 0 \end{pmatrix}$$

is a PIR achievable rate matrix for this code.

Before we state our main results, in order to clearly illustrate the PIR protocol, we first introduce the following definition.

Definition 7: For a given $\nu \times n$ PIR achievable rate matrix $\Lambda_{\kappa,\nu}(\mathcal{C}) = (\lambda_{u,j})$, we define the PIR interference matrices $\mathbf{A}_{\kappa \times n} = (a_{i,j})$ and $\mathbf{B}_{(\nu-\kappa) \times n} = (b_{i,j})$ with

$$\begin{aligned} a_{i,j} &\triangleq u \text{ if } \lambda_{u,j} = 1, \forall j \in \mathbb{N}_n, i \in \mathbb{N}_\kappa, u \in \mathbb{N}_\nu, \\ b_{i,j} &\triangleq u \text{ if } \lambda_{u,j} = 0, \forall j \in \mathbb{N}_n, i \in \mathbb{N}_{\nu-\kappa}, u \in \mathbb{N}_\nu. \end{aligned}$$

Note that in Definition 7, for each $j \in \mathbb{N}_n$, distinct values of $u \in \mathbb{N}_\nu$ should be assigned for all i . Thus, the assignment is not unique in the sense that the order of the entries of each column of \mathbf{A} and \mathbf{B} can be permuted. For $j \in \mathbb{N}_n$, let $\mathcal{A}_j \triangleq \{a_{i,j} : i \in \mathbb{N}_\kappa\}$ and $\mathcal{B}_j \triangleq \{b_{i,j} : i \in \mathbb{N}_{\nu-\kappa}\}$. Note that the j -th column of \mathbf{A} contains the row indices of Λ whose entries in the j -th column are equal to 1, while \mathbf{B} contains the remaining row indices of Λ . Hence, it can be observed that $\mathcal{B}_j = \mathbb{N}_\nu \setminus \mathcal{A}_j, \forall j \in \mathbb{N}_n$.

Definition 8: By $\mathcal{S}(a|\mathbf{A}_{\kappa \times n})$ we denote the set of column coordinates of matrix $\mathbf{A}_{\kappa \times n} = (a_{i,j})$ for which at least one of its entries is equal to a , i.e.,

$$\mathcal{S}(a|\mathbf{A}_{\kappa \times n}) \triangleq \{j \in \mathbb{N}_n : \exists a_{i,j} = a, i \in \mathbb{N}_\kappa\}.$$

The following claim can be directly verified.

Claim 1: $\mathcal{S}(a|\mathbf{A}_{\kappa \times n})$ contains an information set $\forall a \in \mathbb{N}_\nu$. Moreover, for an arbitrary entry $b_{i,j}$ of $\mathbf{B}_{(\nu-\kappa) \times n}$, $\mathcal{S}(b_{i,j}|\mathbf{A}_{\kappa \times n}) \subseteq \mathbb{N}_n \setminus \{j\}$ and it must contain an information set.

We illustrate the previous points in the following example.

Example 2: Continuing with Example 1 and following Definition 7, we obtain

$$\mathbf{A}_{2 \times 5} = \begin{pmatrix} 2 & 1 & 1 & 1 & 1 \\ 3 & 3 & 3 & 2 & 2 \end{pmatrix} \text{ and } \mathbf{B}_{1 \times 5} = (1 \ 2 \ 2 \ 3 \ 3)$$

for $\Lambda_{2,3}$. One can see that $\mathcal{A}_j \cup \mathcal{B}_j = \mathbb{N}_3, \forall j \in \mathbb{N}_5$. Moreover, for instance, take $a = 1$, then $\mathcal{S}(1|\mathbf{A}_{2 \times 5}) = \{2, 3, 4, 5\}$ contains an information set of the $[5, 3, 2]$ systematic code of Example 1.

Now consider the two matrices

$$\begin{pmatrix} c_{\mu+a_{1,1},1}^{(m)} & c_{\mu+a_{1,2},2}^{(m)} & \cdots & c_{\mu+a_{1,n},n}^{(m)} \\ \vdots & \vdots & \vdots & \vdots \\ c_{\mu+a_{\kappa,1},1}^{(m)} & c_{\mu+a_{\kappa,2},2}^{(m)} & \cdots & c_{\mu+a_{\kappa,n},n}^{(m)} \end{pmatrix} \text{ and } \begin{pmatrix} c_{\mu+b_{1,1},1}^{(m)} & c_{\mu+b_{1,2},2}^{(m)} & \cdots & c_{\mu+b_{1,n},n}^{(m)} \\ \vdots & \vdots & \vdots & \vdots \\ c_{\mu+b_{\nu-\kappa,1},1}^{(m)} & c_{\mu+b_{\nu-\kappa,2},2}^{(m)} & \cdots & c_{\mu+b_{\nu-\kappa,n},n}^{(m)} \end{pmatrix}$$

of code symbols of the m -th file, where $\mu \in \mathbb{N}_{\beta-\nu} \cup \{0\}$. Observe that if the user knows the first matrix of code symbols, from Claim 1, since the coordinate set $\mathcal{S}(b_{i,j}|\mathbf{A}_{\kappa \times n}) \subseteq \mathbb{N}_n \setminus \{j\}$ contains an information set and the user knows the structure of the storage code \mathcal{C} , the code symbols $c_{\mu+b_{i,j},j}^{(m)}$ of the second matrix can be obtained. Here, the entries of \mathbf{A} and \mathbf{B} are respectively marked in red and blue. The actual PIR protocol is stated below.

B. PIR Protocol

The proposed PIR protocol generalizes the MDS-coded PIR protocol in [10] to DSSs where files are stored using an arbitrary linear code. Inspired by [8] and [10], a PIR capacity-achievable scheme should follow two important principles: 1) enforcing file symmetry within each storage node, and 2) exploiting side information of undesired symbols to retrieve new desired symbols.¹

The PIR achievable rate matrix $\Lambda_{\kappa,\nu}$ for the given storage code \mathcal{C} plays a central role in the proposed PIR protocol. Moreover, the protocol requires $\beta = \nu^f$ stripes and exploits the corresponding PIR interference matrices $\mathbf{A}_{\kappa \times n}$ and $\mathbf{B}_{(\nu-\kappa) \times n}$. Here, we simply outline the steps of the protocol. Its detailed exposition and corresponding proofs are given in [11, Sec. IV-B and App. B]. A particular example is also given in the extended version [11, Sec. IV-D]. Without loss of generality, we assume that the user wants to download the first file, i.e., $m = 1$. The algorithm is composed of four steps as described below.

Step 1. Index Preparation: For all files, the user interleaves the indices of the rows of $\mathbf{C}^{(m)}$ randomly and independently of each other and generates the interleaved code array $\mathbf{Y}^{(m)} = ((\mathbf{y}_1^{(m)})^\top | \cdots | (\mathbf{y}_\beta^{(m)})^\top)^\top, \forall m \in \mathbb{N}_f$, with rows

$$\mathbf{y}_i^{(m)} = \mathbf{c}_{\pi(i)}^{(m)}, \quad i \in \mathbb{N}_\beta,$$

where $\pi(\cdot) : \mathbb{N}_\beta \rightarrow \mathbb{N}_\beta$ is a random permutation, which is privately known to the user only. Therefore, when the user requests code symbols from each storage node, this procedure is designed to make the requested row indices to be random and independent of the requested file index.

Step 2. Download Symbols in the i -th Repetition: The user downloads the needed symbols in κ repetitions. In the i -th repetition, $i \in \mathbb{N}_\kappa$, the user downloads the required symbols in a total of f rounds. Using the terminology in [10], the user downloads two types of symbols, *desired symbols*, which are directly related to the requested file index $m = 1$, and *undesired symbols*, which are not related to the requested file index $m = 1$, but are exploited to decode the requested file from the desired symbols. For the desired symbols, we will distinguish between round $\ell = 1$ and round $\ell \in \mathbb{N}_{2:f}$.

Undesired symbols. The undesired symbols refer to sums of code symbols which do not contain symbols from the requested file. For every round $\ell, \ell \in \mathbb{N}_{f-1}$, the user downloads the code symbols

$$\begin{cases} \sum_{m' \in \mathcal{M}} y_{((i-1)\mathbf{u}(f-1)+\mathbf{u}(\ell-1)\cdot\nu+a_{1,j},j)}^{(m')}, \\ \cdots, \sum_{m' \in \mathcal{M}} y_{((i-1)\mathbf{u}(f-1)+\mathbf{u}(\ell-1)\cdot\nu+a_{\kappa,j},j)}^{(m')}, \\ \cdots, \sum_{m' \in \mathcal{M}} y_{((i-1)\mathbf{u}(f-1)+\mathbf{u}(\ell-1)\cdot\nu+a_{1,j},j)}^{(m')}, \end{cases}$$

¹In [8] and [10], a third principle was introduced, namely enforcing symmetry across storage nodes. However, this is in general not a necessary requirement for a feasible PIR protocol.

$$\left. \cdots, \sum_{m' \in \mathcal{M}} y_{((i-1)\mathcal{U}(f-1)+\mathcal{U}(\ell)-1)\cdot\nu+a_{\kappa,j},j}^{(m')} \right\}$$

for all $j \in \mathbb{N}_n$ and for all possible subsets $\mathcal{M} \subseteq \mathbb{N}_{2:f}$, where $|\mathcal{M}| = \ell$ and $\mathcal{U}(\ell) \triangleq \sum_{h=1}^{\ell} \kappa^{f-(h+1)}(\nu - \kappa)^{h-1}$.

In contrast to undesired symbols, desired symbols are sums of code symbols which contain symbols of the requested file. The main idea of the protocol is that the user downloads desired symbols that are linear sums of requested symbols and undesired symbols from the previous round.

Desired symbols in the first round. In the first round, the user downloads $\kappa \cdot \mathcal{U}(1) = \kappa \kappa^{f-(1+1)}(\nu - \kappa)^{1-1} = \kappa^{f-1}$ undesired symbols from each storage node. However, these symbols cannot be exploited directly. Hence, due to symmetry, in round $\ell = 1$, the user downloads the κ^{f-1} desired symbols

$$\left\{ y_{\kappa^{f-1}(a_{i,j}-1)+1,j}^{(1)}, \dots, y_{\kappa^{f-1}(a_{i,j}-1)+\kappa^{f-1},j}^{(1)} \right\}$$

from the j -th storage node, $j \in \mathbb{N}_n$, i.e., the user also downloads κ^{f-1} symbols for $m = 1$ from each storage node.

Desired symbols in higher rounds. In the $(\ell + 1)$ -th round, $\ell \in \mathbb{N}_{f-1}$, in order to exploit the side information, i.e., the undesired symbols from the previous round, the user downloads the symbols

$$\left\{ \begin{aligned} & y_{\mathcal{D}(\ell-1)\cdot\nu+a_{i,j},j}^{(1)} \\ & + \sum_{m' \in \mathcal{M}_1} y_{((i-1)\mathcal{U}(f-1)+\mathcal{U}(\ell-1))\cdot\nu+b_{1,j},j}^{(m')} \\ & \cdots y_{(\mathcal{D}(\ell-1)+(\nu-\kappa)-1)\cdot\nu+a_{i,j},j}^{(1)} \\ & + \sum_{m' \in \mathcal{M}_1} y_{((i-1)\mathcal{U}(f-1)+\mathcal{U}(\ell-1))\cdot\nu+b_{\nu-\kappa,j},j}^{(m')} \\ & \cdots y_{[\mathcal{D}(\ell-1)+(\mathcal{U}(\ell)-\mathcal{U}(\ell-1))(\nu-\kappa)-1]\cdot\nu+a_{i,j},j}^{(1)} \\ & + \sum_{m' \in \mathcal{M}_1} y_{((i-1)\mathcal{U}(f-1)+\mathcal{U}(\ell-1))\cdot\nu+b_{\nu-\kappa,j},j}^{(m')} \\ & \cdots y_{(\mathcal{D}(\ell)-1)\cdot\nu+a_{i,j},j}^{(1)} \\ & + \sum_{m' \in \mathcal{M}_{N(\ell)}} y_{((i-1)\mathcal{U}(f-1)+\mathcal{U}(\ell)-1)\cdot\nu+b_{\nu-\kappa,j},j}^{(m')} \end{aligned} \right\}$$

for all distinct ℓ -sized subsets $\mathcal{M}_1, \dots, \mathcal{M}_{N(\ell)} \subseteq \mathbb{N}_{2:f}$, where $j \in \mathbb{N}_n$, $N(\ell) \triangleq \binom{f-1}{\ell}$, and

$$\mathcal{D}(\ell) \triangleq \kappa^{f-1} + \sum_{h=1}^{\ell} \binom{f-1}{h} \kappa^{f-(h+1)}(\nu - \kappa)^h.$$

This indicates that for each combination of files \mathcal{M}_l , $l \in \mathbb{N}_{N(\ell)}$, the user downloads $[\mathcal{U}(\ell) - 1 - \mathcal{U}(\ell - 1) + 1](\nu - \kappa)$ new desired symbols from each storage node, and since there are in total $N(\ell)$ combinations of files, in each round $\mathcal{D}(\ell) - 1 - \mathcal{D}(\ell - 1) + 1$ extra desired symbols are downloaded from each storage node.

Exploiting the side information. Using the fact that for a linear code \mathcal{C} any linear combination of codewords is also

a codeword, and together with Claim 1, it is not too hard to see that by fixing an arbitrary coordinate $j \in \mathbb{N}_n$, there always exist some coordinates $\mathcal{S} \subset \mathbb{N}_n \setminus \{j\}$ (see Claim 1) such that for a subset $\mathcal{M} \subseteq \mathbb{N}_{2:f}$ with $|\mathcal{M}| = \ell$, the so-called *aligned sum*

$$\left\{ \begin{aligned} & \sum_{m' \in \mathcal{M}} y_{((i-1)\mathcal{U}(f-1)+\mathcal{U}(\ell)-1)\cdot\nu+b_{1,j},j}^{(m')} \\ & \cdots, \sum_{m' \in \mathcal{M}} y_{((i-1)\mathcal{U}(f-1)+\mathcal{U}(\ell)-1)\cdot\nu+b_{\nu-\kappa,j},j}^{(m')} \end{aligned} \right\}$$

for $\ell \in \mathbb{N}_{f-1}$ and $i \in \mathbb{N}_\kappa$, can be decoded. Consequently, in the $(\ell + 1)$ -th round, from each storage node j we can collect code symbols related to $m = 1$ from the desired symbols, i.e.,

$$\left\{ y_{\mathcal{D}(\ell-1)\cdot\nu+a_{i,j},j}^{(1)}, \dots, y_{(\mathcal{D}(\ell)-1)\cdot\nu+a_{i,j},j}^{(1)} \right\}$$

is obtained.

Step 3. Complete κ Repetitions: The user repeats Step 2 until $i = \kappa$. We can show that by our designed parameters $\mathcal{U}(\ell)$ and $\mathcal{D}(\ell)$, the user indeed downloads in total $\beta = \nu^f$ stripes for the requested file (see [11, App. B]).

Step 4. Shuffling the Order of Queries to Each Node: The order of the queries to each storage node is uniformly shuffled to prevent the storage node to be able to identify which file is requested from the index of the first downloaded symbol.

C. Achievable PIR Rate

The PIR rate, $\mathcal{R}(\mathcal{C})$, of the protocol proposed in Section III-B for a DSS where f files are stored using an arbitrary $[n, k]$ code \mathcal{C} is given in the following theorem.

Theorem 1: Consider a DSS that uses an $[n, k]$ code \mathcal{C} to store f files. If a PIR achievable rate matrix $\Lambda_{\kappa,\nu}(\mathcal{C})$ exists, then the PIR rate

$$\mathcal{R}(\mathcal{C}) = \frac{(\nu - \kappa)k}{\kappa n} \left[1 - \left(\frac{\kappa}{\nu} \right)^f \right]^{-1} \quad (3)$$

is achievable.

Proof: See [11, App. B]. ■

We remark that from Lemma 1, (3) is smaller than or equal to the MDS-PIR capacity (2) since

$$\begin{aligned} \mathcal{R}(\mathcal{C}) &= \frac{\frac{\nu k}{\kappa n} \left[1 - \frac{\kappa}{\nu} \right]}{\left[1 - \left(\frac{\kappa}{\nu} \right)^f \right]} = \frac{\nu k}{\kappa n} \left[1 + \frac{\kappa}{\nu} + \cdots + \left(\frac{\kappa}{\nu} \right)^{f-1} \right]^{-1} \\ &\leq \left[1 + \frac{k}{n} + \cdots + \left(\frac{k}{n} \right)^{f-1} \right]^{-1}, \end{aligned} \quad (4)$$

and it becomes the MDS-PIR capacity (2) if there exists a matrix $\Lambda_{\kappa,\nu}$ for \mathcal{C} with $\frac{\kappa}{\nu} = \frac{k}{n}$. The inequality in (4) follows from Lemma 1.

Corollary 1: If a PIR achievable rate matrix $\Lambda_{\kappa,\nu}(\mathcal{C})$ with $\frac{\kappa}{\nu} = \frac{k}{n}$ exists for an $[n, k]$ code \mathcal{C} , then the MDS-PIR capacity (2) is achievable.

Definition 9: A PIR achievable rate matrix $\Lambda_{\kappa,\nu}(\mathcal{C})$ with $\frac{\kappa}{\nu} = \frac{k}{n}$ for an $[n, k]$ code \mathcal{C} is called an *MDS-PIR capacity-achieving matrix*, and \mathcal{C} is referred to as an *MDS-PIR capacity-achieving code*.

Note that the largest achievable PIR rate in the noncolluding case where data is stored using an arbitrary linear code is still unknown. Interestingly, it is observed from Lemma 1 and (4) that the largest possible achievable PIR rate for an arbitrary linear code with the proposed protocol strongly depends on the smallest possible value of $\frac{\kappa}{\nu}$ for which a PIR achievable rate matrix $\Lambda_{\kappa,\nu}$ exists. We also remark here that when we say that an MDS-PIR capacity-achieving matrix $\Lambda_{\kappa,\nu}$ exists, it does not necessarily require $(\nu, \kappa) = (n, k)$, but $\frac{\kappa}{\nu} = \frac{k}{n}$. A lower bound on the largest possible achievable PIR rate obtained from Theorem 1 and [11, Lem. 3] is given as follows.

Corollary 2: Consider a DSS that uses an $[n, k, d_{\min}^{\mathcal{C}}]$ code \mathcal{C} to store f files. Then, the PIR rate

$$R(\mathcal{C}) = \frac{\min(k, d_{\min}^{\mathcal{C}} - 1)}{n} \left[1 - \left(\frac{k}{k + \min(k, d_{\min}^{\mathcal{C}} - 1)} \right)^f \right]^{-1}$$

is achievable.

We remark that because every set of k coordinates of an $[n, k]$ MDS code is an information set, we can construct n information sets by cyclically shifting an arbitrary information set n times, hence an MDS-PIR capacity-achieving matrix $\Lambda_{k,n}$ of an MDS code can be easily constructed. In other words, the proposed protocol with MDS codes is MDS-PIR capacity-achieving (see Corollary 1) and MDS codes are a class of MDS-PIR capacity-achieving codes.

IV. MDS-PIR CAPACITY-ACHIEVING CODES

In this section, we provide a necessary and a sufficient condition for an arbitrary linear code to achieve the MDS-PIR capacity C_f (see (2)) with the PIR protocol in Section III-B.

Theorem 2: If an MDS-PIR capacity-achieving matrix exists for an $[n, k]$ code \mathcal{C} , then

$$d_s^{\mathcal{C}} \geq \frac{n}{k} s, \quad \forall s \in \mathbb{N}_k. \quad (5)$$

Proof: See the proof of [11, Th. 3]. \blacksquare

Based on the necessary condition, it can be shown that the code \mathcal{C} in Example 1 is not MDS-PIR capacity-achieving with the PIR protocol in Section III-B, since $d_2^{\mathcal{C}} = 3 < \frac{5}{3} \cdot 2$, i.e., it is impossible to find an MDS-PIR capacity-achieving matrix $\Lambda_{\kappa,\nu}$ for this code.

We have performed an exhaustive search for codes with parameters $k \in \mathbb{N}_n$ and $n \in \mathbb{N}_{11}$ (except for $[n, k] = [10, 5]$ and $[n, k] = [11, 4 \leq k \leq 7]$) and seen that for codes satisfying (5), there always exists an MDS-PIR capacity-achieving matrix. Therefore, we conjecture that (5) is an if and only if condition for the existence of an MDS-PIR capacity-achieving matrix.

Conjecture 1: An MDS-PIR capacity-achieving matrix $\Lambda_{\kappa,\nu}(\mathcal{C})$ with $\frac{\kappa}{\nu} = \frac{k}{n}$ exists for an $[n, k]$ code \mathcal{C} if and only if (5) holds.

In the following, we provide a sufficient condition for MDS-PIR capacity-achieving codes by using the code automorphisms of an $[n, k]$ code [15, Ch. 8].

Theorem 3: Given an $[n, k]$ code \mathcal{C} , if there exist n distinct automorphisms π_1, \dots, π_n of \mathcal{C} such that for every code

coordinate $j \in \mathbb{N}_n$, $\{\pi_1(j), \dots, \pi_n(j)\} = \mathbb{N}_n$, then the code \mathcal{C} is an MDS-PIR capacity-achieving code.

Proof: See the proof of [11, Th. 4]. \blacksquare

Using their known code automorphisms and Theorem 3, it can be shown that the families of cyclic codes and RM codes achieve the MDS-PIR capacity.

Corollary 3: Cyclic codes and RM codes are MDS-PIR capacity-achieving codes.

It can be easily shown that cyclic codes and RM codes satisfy the necessary condition of Theorem 2.

V. CONCLUSION

We presented a PIR protocol for DSSs where data is stored using an arbitrary linear code for the case of noncolluding nodes. By exploiting the information sets of the underlying storage code, an exact expression for the PIR rate of the protocol was derived. Furthermore, a necessary and a sufficient condition for a code to be MDS-PIR capacity-achieving were provided. We proved that cyclic codes and RM codes satisfy the sufficient condition and thus achieve the MDS-PIR capacity with the proposed protocol.

REFERENCES

- [1] B. Chor, O. Goldreich, E. Kushilevitz, and M. Sudan, "Private information retrieval," in *Proc. 36th IEEE Symp. Found. Comp. Sci.*, Milwaukee, WI, USA, Oct. 1995, pp. 41–50.
- [2] Y. Ishai, E. Kushilevitz, R. Ostrovsky, and A. Sahai, "Batch codes and their applications," in *Proc. 36th Annu. ACM Symp. Theory Comput.*, Chicago, IL, USA, Jun. 2004, pp. 262–271.
- [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*, Honolulu, HI, USA, Jun./Jul. 2014, pp. 856–860.
- [4] T. H. Chan, S.-W. Ho, and H. Yamamoto, "Private information retrieval for coded storage," in *Proc. IEEE Int. Symp. Inf. Theory*, Hong Kong, China, Jun. 2015, pp. 2842–2846.
- [5] R. Tajeddine and S. El Rouayheb, "Private information retrieval from MDS coded data in distributed storage systems," in *Proc. IEEE Int. Symp. Inf. Theory*, Barcelona, Spain, Jul. 2016, pp. 1411–1415.
- [6] 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*, Aachen, Germany, Jun. 2017, pp. 1421–1425.
- [7] 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.
- [8] 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.
- [9] —, "The capacity of robust private information retrieval with colluding databases," *IEEE Trans. Inf. Theory*, vol. 64, no. 4, pp. 2361–2370, Apr. 2018.
- [10] 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.
- [11] S. Kumar, H.-Y. Lin, E. Rosnes, and A. Graell i Amat, "Achieving private information retrieval capacity in distributed storage using an arbitrary linear code," Dec. 2017, arXiv:1712.03898v3 [cs.IT].
- [12] 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].
- [13] H. Sun and S. A. Jafar, "Private information retrieval from MDS coded data with colluding servers: Settling a conjecture by Freij-Hollanti et al." in *Proc. IEEE Int. Symp. Inf. Theory*, Aachen, Germany, Jun. 2017, pp. 1893–1897.
- [14] V. K. Wei, "Generalized Hamming weights for linear codes," *IEEE Trans. Inf. Theory*, vol. 37, no. 5, pp. 1412–1418, Sep. 1991.
- [15] F. J. MacWilliams and N. J. A. Sloane, *The Theory of Error-Correcting Codes*. Amsterdam, The Netherlands: North-Holland, 1977.