



CHALMERS

Chalmers Publication Library

Private Information Retrieval in Distributed Storage Systems Using an Arbitrary Linear Code

This document has been downloaded from Chalmers Publication Library (CPL). It is the author's version of a work that was accepted for publication in:

IEEE International Symposium on Information Theory (ISIT)

Citation for the published paper:

Kumar, S. ; Rosnes, E. ; Graell i Amat, A. (2017) "Private Information Retrieval in Distributed Storage Systems Using an Arbitrary Linear Code". IEEE International Symposium on Information Theory (ISIT) pp. 1421-1425.

<http://dx.doi.org/10.1109/ISIT.2017.8006763>

Downloaded from: <http://publications.lib.chalmers.se/publication/252123>

Notice: Changes introduced as a result of publishing processes such as copy-editing and formatting may not be reflected in this document. For a definitive version of this work, please refer to the published source. Please note that access to the published version might require a subscription.

Chalmers Publication Library (CPL) offers the possibility of retrieving research publications produced at Chalmers University of Technology. It covers all types of publications: articles, dissertations, licentiate theses, masters theses, conference papers, reports etc. Since 2006 it is the official tool for Chalmers official publication statistics. To ensure that Chalmers research results are disseminated as widely as possible, an Open Access Policy has been adopted. The CPL service is administrated and maintained by Chalmers Library.

(article starts on next page)

Private Information Retrieval in Distributed Storage Systems Using an Arbitrary Linear Code

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 an information-theoretic private information retrieval (PIR) scheme for distributed storage systems where data is stored using a linear systematic code of rate $R > 1/2$. The proposed scheme generalizes the PIR scheme for data stored using maximum distance separable codes recently proposed by Tajeddine and El Rouayheb for the scenario of a single spy node. We further propose an algorithm to optimize the communication price of privacy (cPoP) using the structure of the underlying linear code. As an example, we apply the proposed algorithm to several distributed storage codes, showing that the cPoP can be significantly reduced by exploiting the structure of the distributed storage code.

I. INTRODUCTION

In data storage applications, besides resilience against disk failures and data protection against illegitimate users, the privacy of the data retrieval query may also be of concern. For instance, one may be interested in designing a storage system in which a file can be downloaded without revealing any information of which file is actually downloaded to the servers storing it. This form of privacy is usually referred to as *private information retrieval* (PIR). PIR is important to, e.g., protect users from surveillance and monitoring.

PIR protocols were first studied by Chor *et al.* in [1], which introduced the concept of an n -server PIR protocol, where a binary database is replicated among n non-colluding servers (referred to as nodes) and the aim is to privately retrieve a single bit from the database while minimizing the total upload and download communication cost. The communication cost in [1] was further reduced in [2] and references therein. Since then, coded PIR schemes have been introduced, where the database is encoded (as opposed to simply being replicated) across several nodes [3]. With the advent of distributed storage systems (DSSs), where the database is encoded and then stored on n nodes, there has been an increasing interest in implementing coded PIR protocols for these systems. PIR protocols for DSSs, where data is stored using codes from two explicit linear code constructions (one protocol for each code construction), were presented in [4], and information-theoretic lower bounds on the tradeoff between the storage cost and the retrieval cost were provided in [5]. In [6], the authors introduced PIR codes which when used in conjunction with traditional n -server PIR protocols allow to achieve PIR on

DSSs. These codes achieve high code rates without sacrificing on the communication cost of an n -server PIR protocol. Recently, the authors in [7] proposed a coded PIR protocol for DSSs that use an (n, k) maximum distance separable (MDS) code for storing data on n storage nodes. The proposed protocol achieves privacy in the presence of at most $n - k$ colluding nodes. In addition, when there are no colluding nodes, the protocol achieves the lowest possible amount of downloaded data per unit of stored data, referred to as the *communication price of privacy* (cPoP).

In the storage community, it is well known that MDS codes are inefficient in the repair of failed nodes. Repair is essential to maintain the initial state of reliability of the DSS. To address efficient repair, Pyramid codes [8] and locally repairable codes (LRCs) [9], have been proposed. They achieve low locality, i.e., a low number of nodes need to be contacted to repair a single failed node.

In this paper, for the scenario with no colluding nodes (i.e., a single spy node), we extend the PIR protocol from [7] to a more general case where data is stored using an arbitrary systematic linear storage code of rate $R > 1/2$. We show that the cPoP can be optimized using the structure of the code, and we provide an algorithm to search for an optimal (in terms of the lowest possible cPoP) protocol. We present the optimal cPoP that can be achieved for various linear codes, including LRCs and Pyramid codes. Interestingly, our numerical results show that non-MDS codes can also achieve the lower bound on the cPoP provided in [5]. Our work bears some similarities to the parallel work in [10], where a PIR protocol protecting against multiple colluding nodes for any linear storage code was presented. However, we show that our extended protocol achieves better cPoP for the scenario of a single spy node.

II. SYSTEM MODEL

We consider a DSS that stores f files $\mathbf{X}^{(1)}, \mathbf{X}^{(2)}, \dots, \mathbf{X}^{(f)}$, where each file $\mathbf{X}^{(m)} = [x_{ij}^{(m)}]$, $m = 1, \dots, f$, is a $\beta \times k$ matrix over $\text{GF}(q^{\alpha\ell})$, with β , k , α , and ℓ being positive integers and q some prime number. Each file is divided into β stripes (blocks) and encoded using a linear code as follows. Let $\mathbf{x}_i^{(m)} = (x_{i,1}^{(m)}, x_{i,2}^{(m)}, \dots, x_{i,k}^{(m)})$, $i = 1, \dots, \beta$, be a message vector (corresponding to the i -th row of $\mathbf{X}^{(m)}$) that is encoded by an (n, k) linear code \mathcal{C} over $\text{GF}(q^\alpha)$, having *subpacketization* α , into a length- n codeword $\mathbf{c}_i^{(m)} = (c_{i,1}^{(m)}, c_{i,2}^{(m)}, \dots, c_{i,n}^{(m)})$, where $c_{i,j}^{(m)} \in \text{GF}(q^{\alpha\ell})$.

The work of S. Kumar 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.




Fig. 1. System Model. (a) The colored boxes in each storage node represent the f coded chunks pertaining to the f files. (b) Coded chunk corresponding to the 2nd file in the n -th node. It consists of β code symbols, $c_{i,n}^{(2)}$, $i = 1, \dots, \beta$. (c) The user sends the queries $\mathbf{Q}^{(j)}$, $j = 1, \dots, n$, to the storage nodes and receives responses \mathbf{r}_j .

When $\alpha = 1$, the code \mathcal{C} is referred to as a *scalar* code. Otherwise, the code is called a *vector* code [11]. The βf generated codewords $\mathbf{c}_i^{(m)}$ are then arranged in the array $\mathbf{C} = ((\mathbf{c}_1^{(1)})^\top | \dots | (\mathbf{c}_\beta^{(1)})^\top | (\mathbf{c}_1^{(2)})^\top | \dots | (\mathbf{c}_\beta^{(f)})^\top)^\top$ of dimension $\beta f \times n$, where $(\cdot)^\top$ denotes the transpose of its argument and $(\mathbf{v}_1 | \dots | \mathbf{v}_{\beta f})$ denotes the concatenation of column vectors $\mathbf{v}_1, \dots, \mathbf{v}_{\beta f}$. For a given column j of \mathbf{C} , we denote the vector $(c_{1,j}^{(m)}, c_{2,j}^{(m)}, \dots, c_{\beta,j}^{(m)})$ as a coded chunk pertaining to file m . Then the f coded chunks in column j are stored on the j -th node as shown in Fig. 1(a). We also assume that the (n, k) code \mathcal{C} is systematic and that the first k code symbols of $\mathbf{c}_i^{(m)}$ are message symbols. Accordingly, we say that the first k nodes are systematic nodes and the remaining nodes are parity nodes.

A. Privacy Model

We consider a DSS where any single node may act as a spy node. Let $s \in \{1, \dots, n\}$ denote the spy node in the DSS. The role of the spy node is to determine which file, m , is accessed by the user. We assume that the user does not know s , since otherwise it can trivially achieve PIR by avoiding contacting the spy node. In addition, the remaining non-spy nodes do not collaborate with the spy node. To retrieve file $\mathbf{X}^{(m)}$ from the DSS, the user sends a $d \times \beta f$ matrix query $\mathbf{Q}^{(j)} = [q_{il}^{(j)}]$ over $\text{GF}(q^\alpha) \subseteq \text{GF}(q^{\alpha\ell})$ to the j -th node for all $j \in \{1, \dots, n\}$. Depending on the queries, node j sends the column vector

$$\mathbf{r}_j = (r_{j,1}, \dots, r_{j,d})^\top = \mathbf{Q}^{(j)} (c_{1,j}^{(1)}, c_{2,j}^{(1)}, \dots, c_{\beta,j}^{(1)}, \dots, c_{\beta,j}^{(f)})^\top, \quad (1)$$

referred to as the response vector, back to the user as illustrated in Fig. 1(c). The following definition shows how such a scheme can achieve perfect information-theoretic PIR.

Definition 1. Consider a DSS with n nodes storing f files in which a node $s \in \{1, \dots, n\}$ acts as a spy. A user who wishes to retrieve the m -th file sends the queries $\mathbf{Q}^{(j)}$, $j = 1, \dots, n$, to all storage nodes, which return the responses \mathbf{r}_j .

This scheme achieves perfect information-theoretic PIR if and only if

$$\text{Privacy: } H(m | \mathbf{Q}^{(s)}) = H(m) \quad (2a)$$

$$\text{Recovery: } H(\mathbf{X}^{(m)} | \mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_n) = 0, \quad (2b)$$

where $H(\cdot)$ denotes the entropy function.

Queries satisfying (2a) ensure that the spy node is not able to determine which file is being downloaded by the user. The recovery constraint in (2b) ensures that the user is able to recover the requested file from the responses sent by the DSS.

The efficiency of a PIR scheme is defined as the total amount of downloaded data per unit of retrieved data, since it is assumed that the content of the retrieved file dominates the total communication cost, i.e., ℓ is much larger than f [7].

Definition 2. The cPoP of a PIR scheme, denoted by θ , is the total amount of downloaded data per unit of retrieved data,

$$\theta = \frac{nd}{\beta k}.$$

It was shown in [5, Th. 3] that the cPoP for a DSS with a single spy node is lowerbounded by $\frac{1}{1-R}$ for a special kind of linear retrieval schemes, where R is the rate of the linear code used to store the data in the DSS. In the case of more than one spy node, an explicit lower bound is currently unknown.

III. CONSTRUCTION

In this section, we present a PIR scheme for a DSS where any node may be a spy node. The DSS uses an (n, k) systematic linear code over $\text{GF}(q^\alpha)$, of rate $R = k/n > 1/2$ and subpacketization α . The code is defined by its parity-check matrix, \mathbf{H} , of size $(n-k) \times n$, and its minimum distance is denoted by d_{\min} . Since the code is systematic, \mathbf{H} can be written as $\mathbf{H} = (\mathbf{P} | \mathbf{I})$, where \mathbf{I} is an $(n-k) \times (n-k)$ identity matrix and \mathbf{P} is an $(n-k) \times k$ parity matrix. In the following, let \tilde{d}_{\min} denote the minimum distance of the $(\tilde{n} = k, \tilde{k} \geq 2k - n)$ code, denoted by $\tilde{\mathcal{C}}$, defined by the parity-check matrix $\tilde{\mathbf{H}} = \mathbf{P}$. We choose $d = k$ and design the n queries as

$$\mathbf{Q}^{(l)} = \begin{cases} \mathbf{U} + \mathbf{V}^{(l)}, & \text{if } l = 1, \dots, k \\ \mathbf{U}, & \text{if } l = k+1, \dots, n \end{cases} \quad (3)$$

where $\mathbf{U} = [u_{ij}]$ is a $k \times \beta f$ matrix whose elements u_{ij} are chosen independently and uniformly at random from $\text{GF}(q^\alpha)$, and $\mathbf{V}^{(l)} = [v_{ij}^{(l)}]$ is a $k \times \beta f$ deterministic binary matrix over $\text{GF}(q^\alpha)$. Note that each $k \times \beta f$ query matrix $\mathbf{Q}^{(l)}$ represents k subqueries, where each subquery corresponds to a row of $\mathbf{Q}^{(l)}$, and where $v_{ij}^{(l)} = 1$ means that the j -th symbol in node l is accessed by the i -th subquery of $\mathbf{Q}^{(l)}$.

Let $\mathbf{E} = [e_{ij}]$ be a $k \times k$ binary matrix, where $e_{ij} = 1$ represents the i -th subquery of the j -th query accessing a message symbol. The design of $\mathbf{V}^{(l)}$ depends on the structure of \mathbf{E} , which must satisfy the following conditions.

- 1) The user should be able to recover β unique symbols of the requested file $\mathbf{X}^{(m)}$ from the i -th subquery of all k queries, i.e., each row of \mathbf{E} should have exactly β ones.

- 2) The user should be able to recover β unique symbols of the requested file $\mathbf{X}^{(m)}$ from each query (consisting of k subqueries), i.e., each column in \mathbf{E} should have β ones.
- 3) The user should be able to recover all βk symbols of the requested file $\mathbf{X}^{(m)}$. This means that all rows of \mathbf{E} (considered here as length- k erasure patterns, with a one indicating an erasure) should be correctable by a maximum likelihood (ML) decoder for the $(\tilde{n} = k, \tilde{k} \geq 2k - n)$ code $\tilde{\mathcal{C}}$ on the binary erasure channel (BEC), i.e., the rows of \mathbf{E} considered as erasure patterns are ML-correctable by $\tilde{\mathcal{C}}$.

From conditions 1) and 2) it follows that \mathbf{E} is a regular matrix with β ones in each row and column. Condition 3) ensures the recovery condition (see (2b)). Details are given in the proof of Theorem 1. Given \mathbf{E} , $\mathbf{V}^{(l)}$ has the following structure

$$\mathbf{V}^{(l)} = (\mathbf{0}_{k \times (m-1)\beta} \mid \mathbf{\Delta}_l \mid \mathbf{0}_{k \times (f-m)\beta}),$$

where $\mathbf{0}_{i \times j}$ denotes the $i \times j$ all-zero matrix and $\mathbf{\Delta}_l$ is a $k \times \beta$ binary matrix. For $l = 1, \dots, k$,

$$\mathbf{\Delta}_l = \left(\boldsymbol{\omega}_{\pi(j_1^{(l)})}^\top \mid \boldsymbol{\omega}_{\pi(j_2^{(l)})}^\top \mid \dots \mid \boldsymbol{\omega}_{\pi(j_k^{(l)})}^\top \right)^\top, \quad (4)$$

where $\pi : \{0, \dots, \beta\} \rightarrow \{0, \dots, \beta\}$ is an arbitrary permutation of size $\beta + 1$ with a fixed point at zero, i.e., $\pi(0) = 0$, $\boldsymbol{\omega}_t$, $t = 1, \dots, \beta$, is the t -th β -dimensional unit vector, i.e., a length- β weight-1 binary vector with a single 1 at the t -th position, $\boldsymbol{\omega}_0$ is the all-zero vector of length β , and

$$j_i^{(l)} = \begin{cases} z_i^{(l)}, & \text{if } e_{il} = 1 \\ 0, & \text{otherwise} \end{cases},$$

where $z_i^{(l)} \in \{1, \dots, \beta\}$ and $z_i^{(l)} \neq z_{i'}^{(l)}$ for $i \neq i'$, $i, i' = 1, \dots, k$. In the following lemma, we show that such a construction of the queries ensures that the privacy condition (2a) is satisfied.

Lemma 1. Consider a DSS that uses an (n, k) linear code with subpacketization α to store f files, each divided into β stripes, and assume the privacy model with a single spy node. Then, the queries $\mathbf{Q}^{(j)}$, $j = 1, \dots, n$, designed as in (3) satisfy $H(m|\mathbf{Q}^{(s)}) = H(m)$, where $s \in \{1, \dots, n\}$ is the spy node.

Proof: The queries $\mathbf{Q}^{(j)}$, $j = 1, \dots, k$, are a sum of a random matrix \mathbf{U} and a deterministic matrix $\mathbf{V}^{(j)}$. The resulting queries have elements that are independently and uniformly distributed at random from $\text{GF}(q^\alpha)$. The same holds for the remaining queries as they are equal to \mathbf{U} . Hence, any $\mathbf{Q}^{(j)}$ obtained by the spy node is statistically independent of m . This ensures that $H(m|\mathbf{Q}^{(s)}) = H(m)$. ■

In order to show that the proposed PIR protocol achieves perfect information-theoretic PIR, it remains to be proved that from the responses \mathbf{r}_j in (1), sent by the nodes back to the user, one can recover the requested file, i.e., that the constructed PIR scheme satisfies the recovery condition in (2b). We call each symbol of the response \mathbf{r}_j a subresponse symbol generated from the corresponding subquery.

Theorem 1. Consider a DSS that uses an (n, k) linear code with subpacketization α to store f files, each divided into β stripes. In order to retrieve the file $\mathbf{X}^{(m)}$, $m = 1, \dots, f$, from the DSS, the user sends the queries $\mathbf{Q}^{(j)}$, $j = 1, \dots, n$, in (3) to the n storage nodes. Then, for the responses \mathbf{r}_j in (1) received by the user, $H(\mathbf{X}^{(m)}|\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_n) = 0$.

The proof is given in the extended version of the paper [12]. Theorem 1 generalizes [7, Th. 1] to any linear code with rate $R > 1/2$. We remark that for the theorem to hold there is an implicit assumption that the three requirements for the matrix \mathbf{E} mentioned above are all satisfied (\mathbf{E} is used in the construction of the queries $\mathbf{Q}^{(j)}$). Thus, the parameter β (which is not explicitly mentioned in the theorem) has to be carefully selected such that a β -regular matrix \mathbf{E} (satisfying the third requirement) actually exists. In the following corollary, we provide such a particular value of β .

Corollary 1. For $\beta = \tilde{d}_{\min} - 1$, it holds that

$$H(\mathbf{X}^{(m)}|\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_n) = 0, \quad (5)$$

and the cPoP becomes $\theta = \frac{n}{\beta} = \frac{n}{\tilde{d}_{\min} - 1}$.

Proof: (5) follows directly from Theorem 1, since all erasure patterns of weight less than \tilde{d}_{\min} are ML-correctable, and $\theta = \frac{n}{\tilde{d}_{\min} - 1}$ follows also from $d = k$. ■

In [10], a PIR protocol achieving a cPoP of $\bar{\theta} \geq \frac{n}{\tilde{d}_{\min} - 1}$, with equality when $\tilde{d}_{\min} - 1$ is a divisor of k , was given. Note that $\tilde{d}_{\min} \leq d_{\min}$, and thus $\theta \leq \bar{\theta}$ for our construction.

Example 1. Consider a DSS that uses a $(5, 3)$ scalar ($\alpha = 1$) binary linear code to store a single file by dividing it into β stripes. The code is defined by the parity-check matrix

$$\mathbf{H} = (\mathbf{P}|\mathbf{I}) = \begin{pmatrix} 1 & 1 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 & 1 \end{pmatrix}.$$

To determine the value of the parameter β , we compute the minimum distance \tilde{d}_{\min} of the $(\tilde{n} = 3, \tilde{k} = 1)$ linear code with parity-check matrix $\tilde{\mathbf{H}} = \mathbf{P}$. From $\tilde{\mathbf{H}}$ it follows that $\tilde{d}_{\min} = 3$. Hence, from Corollary 1, $\beta = \tilde{d}_{\min} - 1 = 2$. Let the file to be stored be denoted by the 2×3 matrix $\mathbf{X} = [x_{ij}]$, where the message symbols $x_{ij} \in \text{GF}(2^\ell)$ for a positive integer ℓ . Then,

$$\mathbf{C} = \begin{pmatrix} x_{11} & x_{12} & x_{13} & x_{11} + x_{12} & x_{12} + x_{13} \\ x_{21} & x_{22} & x_{23} & x_{21} + x_{22} & x_{22} + x_{23} \end{pmatrix}.$$

The user wants to download the file \mathbf{X} from the DSS and sends a query $\mathbf{Q}^{(j)}$, $j = 1, \dots, 5$, to the j -th storage node. The queries take the form shown in (3). For $l = 1, \dots, 3$, we construct the matrix $\mathbf{V}^{(l)} = \mathbf{\Delta}_l$ by choosing an appropriate \mathbf{E} . The only condition in the choice of \mathbf{E} is that it is β -regular. We choose

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

and construct $\mathbf{\Delta}_1$ according to (4). Focusing on the first column of \mathbf{E} , we can see that the first two rows have a one in the first position. Thus, we choose $j_1^{(1)} = 2$, $j_2^{(1)} = 1$, and

$j_3^{(1)} = 0$, since $e_{11} = 1$, $e_{21} = 1$, and $e_{31} = 0$, and take the permutation $0 \rightarrow 0$, $1 \rightarrow 2$, and $2 \rightarrow 1$ as π to get

$$\Delta_1 = \begin{pmatrix} \omega_{\pi(2)} \\ \omega_{\pi(1)} \\ \omega_{\pi(0)} \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \end{pmatrix}.$$

In a similar fashion, we construct

$$\Delta_2 = \begin{pmatrix} 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{pmatrix} \text{ and } \Delta_3 = \begin{pmatrix} 0 & 1 \\ 0 & 0 \\ 1 & 0 \end{pmatrix}.$$

The queries $Q^{(j)}$ are sent to the respective nodes and the responses

$$\begin{aligned} \mathbf{r}_1 &= \begin{pmatrix} u_{11}x_{11}+u_{12}x_{21}+x_{11} \\ u_{21}x_{11}+u_{22}x_{21}+x_{21} \\ u_{31}x_{11}+u_{32}x_{21} \end{pmatrix} = \begin{pmatrix} I_1+x_{11} \\ I_4+x_{21} \\ I_7 \end{pmatrix}, \\ \mathbf{r}_2 &= \begin{pmatrix} u_{11}x_{12}+u_{12}x_{22} \\ u_{21}x_{12}+u_{22}x_{22}+x_{12} \\ u_{31}x_{12}+u_{32}x_{22}+x_{22} \end{pmatrix} = \begin{pmatrix} I_2 \\ I_5+x_{12} \\ I_8+x_{22} \end{pmatrix}, \\ \mathbf{r}_3 &= \begin{pmatrix} u_{11}x_{13}+u_{12}x_{23}+x_{23} \\ u_{21}x_{13}+u_{22}x_{23} \\ u_{31}x_{13}+u_{32}x_{23}+x_{13} \end{pmatrix} = \begin{pmatrix} I_3+x_{23} \\ I_6 \\ I_9+x_{13} \end{pmatrix}, \\ \mathbf{r}_4 &= \begin{pmatrix} u_{11} & u_{12} \\ u_{21} & u_{22} \\ u_{31} & u_{32} \end{pmatrix} \begin{pmatrix} x_{11}+x_{12} \\ x_{21}+x_{22} \end{pmatrix} = \begin{pmatrix} I_1+I_2 \\ I_4+I_5 \\ I_7+I_8 \end{pmatrix}, \\ \mathbf{r}_5 &= \begin{pmatrix} u_{11} & u_{12} \\ u_{21} & u_{22} \\ u_{31} & u_{32} \end{pmatrix} \begin{pmatrix} x_{12}+x_{13} \\ x_{22}+x_{23} \end{pmatrix} = \begin{pmatrix} I_2+I_3 \\ I_5+I_6 \\ I_8+I_9 \end{pmatrix}, \end{aligned}$$

where $I_i = \sum_{j=1}^2 u_{tj}x_{js}$ and $i = 3(t-1) + s$, with $s, t = 1, \dots, 3$, are collected by the user. Notice that each storage node sends back $k = 3$ symbols. The user obtains the requested file as follows. Knowing I_2 , the user obtains I_1 and I_3 from the first components of \mathbf{r}_4 and \mathbf{r}_5 . This allows the user to obtain x_{11} and x_{23} . In a similar fashion, knowing I_6 the user gets I_5 from the second component of \mathbf{r}_5 , then uses this to obtain I_4 from the second component of \mathbf{r}_4 . This allows the user to obtain x_{21} and x_{12} . Similarly, knowing I_7 allows the user to get I_8 from the third component of \mathbf{r}_4 . Knowing I_8 allows the user to obtain I_9 from the third component of \mathbf{r}_5 , which then allows to recover the symbols x_{22} and x_{13} . In this way, the user recovers all symbols of the file and hence recovers \mathbf{X} . Note that $\theta = \frac{5-3}{2-3} = 2.5$, which is equal to the lower bound $1/(1-R)$.

IV. OPTIMIZING THE COMMUNICATION PRICE OF PRIVACY

In the previous section, we provided a construction of a PIR scheme for DSSs that use an arbitrary linear systematic code to store data and showed that a cPoP of $n/(\tilde{d}_{\min} - 1)$ is achievable while maintaining information-theoretic PIR (see Theorem 1 and Corollary 1). In this section, we provide an algorithm (based on Theorem 1) to further lower the cPoP taking the structure of the underlying code into consideration. The algorithm is outlined in Algorithm 1.

The main issues that need to be addressed are the efficient enumeration of the set of erasure patterns of a given weight β that can be corrected under ML decoding of $\tilde{\mathcal{C}}$, and the efficient computation of the matrix \mathbf{E} . Such ML-correctable erasure patterns are binary vectors

Algorithm 1: Optimizing the cPoP

Input : Distributed storage code \mathcal{C} of length n
Output: Optimized matrix \mathbf{E}_{opt} and largest possible β

```

1  $\beta \leftarrow \tilde{d}_{\min} - 1$ 
2  $\mathbf{E}_{\text{opt}} \leftarrow \emptyset$ ,  $\beta_{\text{opt}} \leftarrow \beta$ 
3 while  $\beta \leq \tilde{n} - k$  do
4    $\mathcal{L} \leftarrow \text{ComputeErasurePatternList}(\tilde{\mathcal{C}}, \beta)$ 
5   if  $\mathcal{L} \neq \emptyset$  then
6      $\mathbf{E} \leftarrow \text{ComputeMatrix}(\mathcal{L})$ 
7     if  $\mathbf{E} \neq \emptyset$  then
8        $\mathbf{E}_{\text{opt}} \leftarrow \mathbf{E}$ ,  $\beta_{\text{opt}} \leftarrow \beta$ 
9     else
10      return  $(\mathbf{E}_{\text{opt}}, \beta_{\text{opt}})$ 
11    end
12  end
13   $\beta \leftarrow \beta + 1$ 
14 end
15 return  $(\mathbf{E}_{\text{opt}}, \beta_{\text{opt}})$ 

```

(of length $\tilde{n} = k$) which can be ML-decoded on the BEC when the positions corresponding to the 1-entries are erased by the channel. These issues are addressed by the subprocedures $\text{ComputeErasurePatternList}(\tilde{\mathcal{C}}, \beta)$ and $\text{ComputeMatrix}(\mathcal{L})$, in Lines 4 and 6, respectively, of Algorithm 1, and discussed below in Sections IV-A and IV-B.

We remark that the algorithm will always return $\mathbf{E}_{\text{opt}} \neq \emptyset$, since initially $\beta = \tilde{d}_{\min} - 1$. Then, in the first iteration of the main loop, the list \mathcal{L} will contain all length- k binary vectors of weight $\beta = \tilde{d}_{\min} - 1 < \tilde{d}_{\min}$. Thus, any vector that is shift-variant (i.e., the k cyclic shifts are all different) can be chosen for the first row of \mathbf{E} . The remaining rows of \mathbf{E} are obtained by cyclically shifting the first row (the i -th row is obtained by cyclically shifting the first row i times). In the particular case of \mathcal{C} being an MDS code, $d_{\min} = \tilde{d}_{\min} = n - k + 1$, the algorithm will do exactly one iteration of the main loop, and the overall PIR scheme reduces to the one described in [7, Sec. IV].

A. $\text{ComputeErasurePatternList}(\tilde{\mathcal{C}}, \beta)$

Computing a list of erasure patterns that are correctable under ML decoding for a given short code can be done using any ML decoding algorithm. For small codes $\tilde{\mathcal{C}}$, all length- k binary vectors of weight β that are ML-correctable can be found using an exhaustive search, while for longer codes a random search can be performed, in the sense of picking length- k binary vectors of weight β at random, and then verifying whether or not they are ML-correctable. Alternatively, one can apply a random permutation π to the columns of $\tilde{\mathbf{H}}$, apply the Gauss-Jordan algorithm on the resulting matrix to transform it into *row echelon form*, collect a subset of size β of the column indices of *leading-one-columns*, and finally apply the inverse permutation of π to this subset of column indices. The *leading-one-columns* are the columns containing a *leading one*, where the first nonzero entry in each matrix row of a matrix in row echelon form is called a leading one. This will give the support

set of an ML-correctable erasure pattern of weight β that can be added to \mathcal{L} . Finally, one can check whether all cyclic shifts of the added erasure pattern are ML-correctable or not and add the ML-correctable cyclic shifts to \mathcal{L} .

B. ComputeMatrix(\mathcal{L})

Given the list \mathcal{L} of erasure patterns that are correctable under ML decoding, we construct a $|\mathcal{L}| \times k$ matrix, denoted by $\Psi = [\psi_{ij}]$, in which each row is one of these patterns. The problem is now to find a $k \times k$ submatrix of constant column weight β (and constant row weight β). This can be formulated as an integer program (in the integer variables $\eta_1, \eta_2, \dots, \eta_{|\mathcal{L}|}$) in the following way,

$$\begin{aligned} & \text{maximize} && \sum_{i=1}^{|\mathcal{L}|} \eta_i \\ & \text{s. t.} && \sum_{i=1}^{|\mathcal{L}|} \eta_i \psi_{ij} = \beta, \forall j \in \{1, \dots, k\}, \\ & && \eta_i \in \{0, 1\}, \forall i \in \{1, \dots, |\mathcal{L}|\}, \text{ and} \\ & && \sum_{i=1}^{|\mathcal{L}|} \eta_i = k. \end{aligned} \quad (6)$$

A valid $k \times k$ submatrix of Ψ is constructed from the rows of Ψ with corresponding η -values of one in any feasible solution of (6). When $|\mathcal{L}|$ is large, solving (6) may become impractical (solving a general integer program is known to be NP-hard), in which case one can take several random subsets of the list \mathcal{L} of some size, construct the corresponding matrices Ψ , and try to solve the program in (6). Finally, before solving (6), one may check whether there are erasure patterns in \mathcal{L} with all its cyclic shifts (assuming they are all different) also in \mathcal{L} , in which case the corresponding submatrix of Ψ will be a valid $k \times k$ matrix E .

V. NUMERICAL RESULTS

We present optimized values for the cPoP for different systematic linear codes. The results are tabulated in Table I, where $\theta_{\text{LB}} = 1/(1-R)$ is the lower bound on the cPoP taken from [5, Th. 5], θ_{opt} is the optimized value computed from Algorithm 1, and $\theta_{\text{non-opt}} = n/(\tilde{d}_{\text{min}} - 1)$. The code C_1 in the table is from Example 1, C_2 is an (11, 6) binary linear code with optimum minimum distance, while codes C_3 and C_5 are Pyramid codes, taken from [8], of locality of 4 and 6, respectively. C_4 is an LRC of locality 5 borrowed from [9].

In [13], a construction of optimal (in terms of minimum distance) binary LRCs with multiple repair groups was given. In particular, in [13, Constr. 3], a construction based on array low-density parity-check (LDPC) codes was provided. The parity part (or the P matrix) of the parity-check matrix H of the optimal LRC is the parity-check matrix of an array-based LDPC code. The minimum distance of array LDPC codes is known for certain sets of parameters (see, e.g., [14], and references therein). Codes C_6 and C_7 in Table I are *optimal* LRCs based on array LDPC codes constructed using [13, Constr. 3] and having a locality of 11.

For all codes, θ_{opt} is close to the lower bound on the cPoP, θ_{LB} . Remarkably, the codes C_1 , C_3 , and C_5 achieve the lower bound on the cPoP despite the fact that these codes are not MDS codes. The remaining codes (C_2 , C_4 , C_6 , and C_7) achieve a cPoP of $\tilde{n}/(\tilde{n} - k) = k/(k - \tilde{k}) > 1/(1 - R)$, which is the

TABLE I
OPTIMIZED VALUES FOR THE cPOP FOR DIFFERENT CODES

Code	d_{min}	\tilde{d}_{min}	$\theta_{\text{non-opt}}$	θ_{opt}	θ_{LB}
C_1 : (5, 3) (Example 1)	2	3	2.5	2.5	2.5
C_2 : (11, 6)	4	4	3.6667	2.75	2.2
C_3 : (12, 8) Pyramid	4	4	4	3	3
C_4 : (16, 10) LRC	5	5	4	3.2	2.6667
C_5 : (18, 12) Pyramid	5	5	4.5	3	3
C_6 : (154, 121) LRC	4	6	30.8	4.9677	4.6667
C_7 : (187, 121) LRC	7	16	12.4667	3.0656	2.8333

lowest possible value given the parameters of \tilde{C} . The strict inequality is due to the fact that \tilde{H} is not full rank.

VI. CONCLUSION

We generalized the PIR protocol proposed in [7] for a DSS with a single spy node and where data is stored using an MDS code to the case where an arbitrary systematic linear code of rate $R > 1/2$ is used to store data. We also presented an algorithm to optimize the cPoP of the protocol. The optimization leads to a cPoP close to its theoretical lower bound. Interestingly, for certain codes, the lower bound on the cPoP can be achieved.

REFERENCES

- [1] B. Chor, E. Kushilevitz, O. Goldreich, and M. Sudan, "Private information retrieval," in *Proc. Annual IEEE Symp. Foundations Comp. Sci. (FOCS)*, Milwaukee, WI, Oct. 1995, pp. 41–50.
- [2] S. Yekhanin, "Towards 3-query locally decodable codes of subexponential length," *Journal of the ACM*, vol. 55, no. 1, pp. 1–16, Feb. 2008.
- [3] Y. Ishai, E. Kushilevitz, R. Ostrovsky, and A. Sahai, "Batch codes and their applications," in *Proc. 36th Annual ACM Symp. Theory Comput. (STOC)*, Chicago, IL, Jun. 2004, pp. 262–271.
- [4] 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.
- [5] 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.
- [6] A. Fazeli, A. Vardy, and E. Yaakobi, "Codes for distributed PIR with low storage overhead," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Hong Kong, China, Jun. 2015, pp. 2852–2856.
- [7] 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.
- [8] 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.
- [9] 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*, Trento, Italy, Aug. 2013, pp. 325–336.
- [10] R. Freij-Hollanti, O. Gnilke, C. Hollanti, and D. Karpuk, "Private information retrieval from coded databases with colluding servers," Nov. 2016. [Online]. Available: <https://arxiv.org/abs/1611.02062>
- [11] M. Blaum, J. Brady, J. Bruck, and J. Menon, "EVENODD: An efficient scheme for tolerating double disk failures in RAID architectures," *IEEE Trans. Comput.*, vol. 44, no. 2, pp. 192–202, Feb. 1995.
- [12] S. Kumar, E. Rosnes, and A. Graell i Amat, "Private information retrieval in distributed storage systems using an arbitrary linear code," Dec. 2016. [Online]. Available: <https://arxiv.org/abs/1612.07084>
- [13] J. Hao and S.-T. Xia, "Constructions of optimal binary locally repairable codes with multiple repair groups," *IEEE Commun. Lett.*, vol. 20, no. 6, pp. 1060–1063, Jun. 2016.
- [14] E. Rosnes, M. A. Ambroze, and M. Tomlinson, "On the minimum/stopping distance of array low-density parity-check codes," *IEEE Trans. Inf. Theory*, vol. 60, no. 9, pp. 5204–5214, Sep. 2014.