This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Tajeddine, Razane; Gnilke, Oliver W.; El Rouayheb, Salim

Private Information Retrieval from MDS Coded Data in Distributed Storage Systems

Published in:
IEEE Transactions on Information Theory

DOI:
[10.1109/TIT.2018.2815607](https://doi.org/10.1109/TIT.2018.2815607)

Published: 13/03/2018

Document Version
Peer reviewed version

Please cite the original version:
Tajeddine, R., Gnilke, O. W., & El Rouayheb, S. (2018). Private Information Retrieval from MDS Coded Data in Distributed Storage Systems. *IEEE Transactions on Information Theory*, 64(11), 7081-7093.
<https://doi.org/10.1109/TIT.2018.2815607>

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.

Private Information Retrieval from MDS Coded Data in Distributed Storage Systems

Razane Tajeddine*, Oliver W. Gnilke*, and Salim El Rouayheb†

*Department of Mathematics and Systems Analysis

Aalto University, School of Science, Finland

Email: {razane.tajeddine, oliver.gnilke}@aalto.fi

†ECE Department, Rutgers University

Email: salim.elrouayheb@rutgers.edu

Abstract—The problem of providing privacy, in the private information retrieval (PIR) sense, to users requesting data from a distributed storage system (DSS), is considered. The DSS is coded by an (n, k, d) Maximum Distance Separable (MDS) code to store the data reliably on unreliable storage nodes. Some of these nodes can be spies which report to a third party, such as an oppressive regime, which data is being requested by the user. An information theoretic PIR scheme ensures that a user can satisfy its request while revealing no information on which data is being requested to the nodes. A user can trivially achieve PIR by downloading all the data in the DSS. However, this is not a feasible solution due to its high communication cost. We construct PIR schemes with low download communication cost. When there is $b = 1$ spy node in the DSS, in other words, no collusion between the nodes, we construct PIR schemes with download cost $\frac{1}{1-R}$ per unit of requested data ($R = k/n$ is the code rate), achieving the information theoretic limit for linear schemes. The proposed schemes are universal since they depend on the code rate, but not on the generator matrix of the code. Also, if $b \leq n - \delta k$ nodes collude, with $\delta = \lfloor \frac{n-b}{k} \rfloor$, we construct linear PIR schemes with download cost $\frac{b+\delta k}{\delta}$.

I. INTRODUCTION

Consider the following scenario. A group of online peers (storage nodes) want to collaborate together to form a peer-to-peer (p2p) distributed storage system (DSS) to store and share files reliably, while ensuring information theoretic private information retrieval (PIR). The PIR [2], [3] property allows a user (possibly one of the peers) to download a file while revealing no information about which file is being downloaded. We are mainly motivated by the following two applications: 1) A DSS that protects users from surveillance and monitoring, for instance from an oppressive regime. The people (peers) collectively contribute to storing the data and making it pervasively available online. But, some peers could be spies

for the regime. They could turn against their “neighbors” and report to the oppressor the identity of users accessing some information deemed to be anti-regime (blogs, photos, videos, etc.), leading to their persecution; 2) A DSS that protects the personal information of users, such as gender, age group, disease, etc., which can be inferred from their file access history. This information can potentially be used to target them with unwanted advertisement, or even affect them adversarially in other areas, such as applications to health insurance or bank loans. In this respect, the studied DSS can provide an infrastructure, at least in theory, over which applications, such as cloud storage and social networking, can be run with a privacy guarantee for the users.

We suppose the DSS is formed of n peers or nodes. Peers can be temporarily offline or can leave the system at any time. The data is stored redundantly in the system to guarantee its durability and availability. We assume that the DSS uses an (n, k, d) maximum distance separable (MDS) code that can tolerate $n - k$ simultaneous node failures. A certain number of nodes in the DSS, say b , whose identities are unknown to the users or the system, are spies that collude and can report the user requests to the oppressor, or sell this information to interested third parties. The user can always achieve PIR by asking to download all the files in the DSS. However, this solution is not feasible due to its high communication cost, and more efficient solutions have been studied in the PIR literature [4]–[10] assuming the data is *replicated* in the system. The next example illustrates our PIR scheme with efficient communication cost that can be run on MDS coded data.

Example 1. Consider a DSS formed of $n = 4$ nodes, as shown in Figure 1, that stores m files (a_i, b_i) , $a_i, b_i \in GF(3^w)$, $i = 1, 2, \dots, m$. The DSS is coded by an $(n, k, d) = (4, 2, 3)$ MDS code over $GF(3)$ to store the files. Nodes 1, $\dots$, 4 store, respectively, $a_i, b_i, a_i + b_i, a_i + 2b_i$, $i = 1, \dots, m$. Suppose the user is interested in retrieving file f , i.e., (a_f, b_f) , which can equally likely be any of the m files. To this end, the user generates a random vector $\mathbf{u} = (u_1, \dots, u_m)$ with components chosen independently and uniformly at random from the underlying base field $GF(3)$. It sends the query vector $\mathbf{q} = \mathbf{u}$ to nodes 1 and 2 and $\mathbf{q} = \mathbf{u} + \mathbf{e}_f$ to nodes 3 and 4, where $\mathbf{e}_f$ is the all zero vector of length m with a 1 in the

This paper is an extension to the work which has been presented in part at the IEEE International Symposium on Information Theory (ISIT) 2016 [1]. This work was supported in part by NSF Grant CCF 1817635. Furthermore, we gratefully acknowledge financial support by the Academy of Finland through grants #276031, #282938, and #303819 to C. Hollanti, Aalto University, Finland. Part of this work was carried out while the first and third authors were with the ECE department at the Illinois Institute of Technology, Chicago. The first two authors would like to thank the Institute for Communications Engineering at the Technical University of Munich for hosting them while part of this work was carried out. Copyright (c) 2017 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

Fig. 1: The user sends queries as specified in Example 1 and receives the responses. From the responses, the user can decode a_f and b_f , thus decoding the desired file privately.

f^{th} entry. Upon receiving the user's request, each node in the DSS returns to the user the projection of all its data on the received query vector. For instance, suppose that the user wants file 1. Then, nodes 1, $\dots$, 4 return the following symbols from $GF(3^w)$, $I_1, I_2, a_f + b_f + I_1 + I_2, a_f + 2b_f + I_1 + 2I_2$, where $I_1 = \sum_{i=1}^m u_i a_i$ and $I_2 = \sum_{i=1}^m u_i b_i$ are thought of as "interference" terms. The returned information forms an invertible linear system and the user can decode a_f and b_f . Assume that the DSS contains no colluding nodes, i.e. $b = 1$. Then, the proposed scheme achieves PIR since the query vector to each node is statistically independent of the file index f . However, if a node, say node 1, knows the query vector of another node, say node 3, it may be able to pin down which file the user wanted, by computing $\mathbf{e}_f = \mathbf{q} - \mathbf{u}$. However, we assume that a node does not have access to the queries coming to any other nodes, and PIR is indeed achieved here. This PIR scheme downloads 4 symbols to retrieve a file of size 2 symbols. We say that the communication price of privacy $cPoP = 4/2 = 2$ for this scheme, which does not depend on the number of files in the system.

Replication-based PIR: PIR was first introduced in the seminal papers of Chor et al. in [2], [3] followed by a significant amount of research in this area [4]–[8], [11], [12]. The classical model considers a binary database of length m and a user that wishes to privately retrieve the value of a bit (a record) in it, while minimizing the total communication cost including the upload (query) and download phase. Chor et al. [3] showed that if there is one server storing the database, the user has to download the whole database in order to achieve information theoretic PIR. However, when the database is replicated on n non-colluding (non-cooperating) servers (nodes), they devised a PIR scheme with total, upload and download, communication cost of $O((n^2 \log n)m^{1/n})$ and $O(m^{1/3})$ for the special case of $n = 2$. In the past few years, there has been significant progress in developing PIR protocols with total communication cost that is subpolynomial in the size of the database [11]–[13]. Moreover, a connection between PIR and blind interference alignment was discussed in [14]. PIR in a computational sense was shown to be achievable with a single server (no replication) in [15] assuming the hardness of quadratic residuosity problem. PIR schemes on databases that are replicated but not perfectly synchronized were studied in [16].

Coded PIR: The original model studied in PIR assumes that the entire data is replicated on each node. PIR on coded data was studied in the literature on Batch Codes [17], where the data is coded to allow parallel processing leading to amortizing the PIR communication cost over multiple retrievals. Recently, the PIR problem in DSSs that use erasure codes was initiated in [9], where it was shown that one extra bit of download is sufficient to achieve PIR assuming the number of servers n to be exponential in the number of files. Bounds on the information theoretic tradeoff between storage and download communication cost for coded DSSs, for arbitrary number of files m , were derived in [10]. The setting when nodes can be byzantine (malicious) was considered in [18] and robust PIR schemes were devised using locally decodable codes. Robust PIR was also studied in [19], [20]. In [21], methods for transforming a replication-based PIR scheme into a coded-based PIR scheme with the same communication cost, up to a multiplicative constant, were studied. PIR array codes with optimal rate were designed in [22].

Following this work in [1], [23], the lowest achievable price of privacy for repetition code on n nodes having m files and b colluding nodes was found in [24], [25] to be $\frac{1-(b/n)^m}{1-(b/n)}$ and that of an (n, k) -code was found in [26] to be $\frac{1-(k/n)^m}{1-(k/n)}$. Also, schemes using GRS codes have been constructed in [27], and they conjectured that the lowest achievable price of privacy is $\frac{1-(\frac{b+k-1}{n})^m}{1-\frac{b+k-1}{n}}$. That conjecture was then disproved using a counter example in [28]. Moreover, in [29], PIR on coded data such that arbitrary sets of servers collude is studied. In [30], PIR schemes for any arbitrary code were discussed. Some work was also done on symmetric PIR, where the objective is to not only protect the privacy of the user, but also the privacy of the server, such that the user should not get information about files other than the one he wants [31]. Also, the capacity of byzantine PIR on replicated storage systems was found in [32].

Contributions: Motivated by the two DSS applications mentioned earlier, we draw the following distinctions with the previous literature prior to this work on coded PIR [1]: (i) To the best of our knowledge, all the previous work on coded PIR, except for [10], assumes that the code is used to encode together data from different files (records). However, the model here is different, since in DSS applications only data chunks belonging to the same file are encoded together (as

done in Example 1); (ii) The work in [10] studies fundamental limits on the costs of coded PIR. Here, we provide explicit constructions of PIR schemes with efficient communication cost.

In comparison with the classical literature on replication-based PIR, we make the following observations: (i) We focus only on the number of downloaded symbols in the communication cost of a PIR scheme. This is since the upload query matrices are dependent only on the number of files and not on the size of the file, and typically in DSSs, the size of a file is relatively larger than the total number of files.; (ii) Up to b nodes may collude and share their queries in the hope of determining the requested file.

In the model we study, we assume that the MDS code parameters (n, k, d) are given and depend on the desired reliability level for the data. Therefore, they are not design parameters that can be chosen to optimize the efficiency of the PIR scheme. However, the code itself may have to be designed jointly with the PIR scheme. A PIR scheme incurs many overheads on the DSS, including communication cost, computations [7], and connectivity; user contacts n instead of k nodes, as seen in Example 1. However, we measure here the efficiency of a PIR scheme only by its total download communication cost, which we refer to as the *communication Price of Privacy (cPoP)*. A more formal definition of *cPoP* is given in Definition 3 after the model we use has been established. The PIR rate is the inverse of the *cPoP*, i.e. the data downloaded from the required file per downloaded symbol. The following questions naturally arise here: (1) What is the minimum achievable cPoP for given n, k and b ? (2) How to efficiently construct codes and PIR schemes that achieve optimal cPoP? (3) Do the code and PIR scheme have to be designed jointly to achieve optimum cPoP? The last question addresses the problem of whether reliability and PIR could be addressed separately in a DSS. Moreover, it may have practical implications on whether data already existing in coded form needs to be re-encoded to achieve PIR with minimum cPoP.

In this paper, we make progress towards answering the last two questions and provide constructions of efficient PIR schemes for querying MDS coded data. Specifically, we make the following contributions: (i) For $b = 1$, i.e., no colluding nodes, we construct a linear PIR scheme with $cPoP = \frac{1}{1-R}$ ($R = k/n$ is the code rate), thus achieving the lower bound on *cPoP* for linear schemes in [10], [26] as $m \rightarrow \infty$; (ii) For $2 \leq b \leq d - 1$, we construct linear PIR schemes with $cPoP = b + k$; (iii) More generally, for $b \leq n - \delta k$, $\delta = \lfloor \frac{n-b}{k} \rfloor$, we construct linear PIR schemes with $cPoP = \frac{b + \delta k}{\delta}$. While the minimum cPoP in this regime is unknown, the constructed schemes have a cPoP that does not depend on m , the number of files in the system. An important property of the scheme for $b = 1$ is its *universality*. It depends only on n, k , and b , but not on the generator matrix of the code. Moreover, both of these schemes can be constructed for any given MDS code, i.e., it is not necessary to design the code jointly with the PIR scheme. This implies that b does not have to be a rigid system parameter. Each user can choose their own value of b to reflect its desired privacy level, at the expense of a higher *cPoP*. The DSS can serve all the users simultaneously

storing the same encoded data, i.e., without having to store different encodings for different values of b . The construction in [27] is a generalized version of the earlier scheme presented here. In both schemes, the parity check matrix of the storage system should be known. The two schemes perform equally well, and are in fact identical, for the case of no-collusion ($b = 1$) and for the case of $(n - k)$ -collusion ($b = n - k$). In the intermediate regime, the generalized scheme in [27] outperforms our scheme.

II. MODEL

Distributed Storage Systems: Consider a distributed storage system (DSS) formed of n storage nodes indexed from 1 to n . The DSS stores m files, $X^1, \dots, X^m$, of equal sizes. The DSS uses WLOG a systematic¹ (n, k, d) MDS code over $GF(q)$ to store the data redundantly and achieve reliability against $d - 1$ node failures. We assume that each file, $X^i, i = 1, \dots, m$, is divided into α stripes, and each stripe is divided into k blocks. We represent the file $X^i = [x_{lj}^i], l = 1, \dots, \alpha, j = 1, \dots, k$, as an $\alpha \times k$ matrix, with symbols from the finite field $GF(q^w)$. We divide the file into stripes to have the number of parts of X^i be divisible by the number of queries and by the number of retrieved symbols per query.

$$X^i = \begin{pmatrix} x_{11}^i & x_{12}^i & \dots & x_{1k}^i \\ x_{21}^i & x_{22}^i & \dots & x_{2k}^i \\ \vdots & \vdots & \ddots & \vdots \\ x_{\alpha 1}^i & x_{\alpha 2}^i & \dots & x_{\alpha k}^i \end{pmatrix}. \quad (1)$$

Define $\mathcal{X}$ to be the $m\alpha \times k$ matrix denoting all the systematic data in the system, i.e.,

$$\mathcal{X}_{m\alpha \times k} = \begin{pmatrix} X^1 \\ X^2 \\ \vdots \\ X^m \end{pmatrix}.$$

Each stripe of each file is encoded separately using the same systematic MDS code with a $k \times n$ generator matrix $\Lambda = [\lambda_{ij}]$ with elements in $GF(q)$. Since the code is systematic, the square submatrix of Λ formed of the first k columns is the identity matrix. The encoded data, $\mathcal{X}\Lambda$, is stored on the DSS as shown in Table I. We assume that the user knows this layout table, i.e., he/she knows the coding coefficients for each node. We denote by $\mathbf{w}_l \in GF(q^w)^{m\alpha}, l = 1, \dots, n$ the column vector representing all the data on node l .

PIR: Suppose the user wants file X^f , where f is chosen uniformly at random from the set $[m] = \{1, \dots, m\}$. To retrieve file X^f , the user sends requests to the nodes, among which there are b colluding nodes. The user does not know which nodes are colluding, else, he/she would avoid them. The goal is to devise a PIR scheme that allows the user to decode X^f , while revealing no information, in an information theoretic sense, about f to the nodes. The colluding nodes can analyze the different requests they receive from the user

¹We focus on systematic codes due to their widespread use in practice. However, our results still hold for non-systematic codes.

	node 1	node 2	...	node k	node $k + 1$	...	node n
file 1	x_{11}^1	x_{12}^1	...	x_{1k}^1	$\lambda_{1,k+1}x_{11}^1 + \dots + \lambda_{k,k+1}x_{1k}^1$	...	$\lambda_{1n}x_{11}^1 + \dots + \lambda_{kn}x_{1k}^1$
	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
	$x_{\alpha 1}^1$	$x_{\alpha 2}^1$	...	$x_{\alpha k}^1$	$\lambda_{1,k+1}x_{\alpha 1}^1 + \dots + \lambda_{k,k+1}x_{\alpha k}^1$	...	$\lambda_{1n}x_{\alpha 1}^1 + \dots + \lambda_{kn}x_{\alpha k}^1$
...	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
	x_{11}^m	x_{12}^m	...	x_{1k}^m	$\lambda_{1,k+1}x_{11}^m + \dots + \lambda_{k,k+1}x_{1k}^m$	...	$\lambda_{1n}x_{11}^m + \dots + \lambda_{kn}x_{1k}^m$
	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
file m	$x_{\alpha 1}^m$	$x_{\alpha 2}^m$	...	$x_{\alpha k}^m$	$\lambda_{1,k+1}x_{\alpha 1}^m + \dots + \lambda_{k,k+1}x_{\alpha k}^m$	...	$\lambda_{1n}x_{\alpha 1}^m + \dots + \lambda_{kn}x_{\alpha k}^m$

TABLE I: The layout of the encoded symbols of the m files in the DSS.

in order to identify the requested file. However, as explained in the introduction, a node has access to the requests coming to at most $b - 1$ other nodes in the system. Under this setting, we are interested in linear PIR schemes.

Definition 1. A PIR scheme is linear over $GF(q)$, and of dimension ρ , if it consists of the following two stages.

1. Request stage: Based on which file the user wants, he/she sends requests to a subset of nodes in the DSS. The request to node l takes the form of a $\rho \times m\alpha$ query matrix Q_l over $GF(q)$.

2. Download stage: Node l responds by sending the projection of its data onto Q_l , i.e.,

$$R_l = Q_l \mathbf{w}_l \in GF(q^w)^\rho. \quad (2)$$

We think of each query matrix Q_l as formed of ρ sub-queries corresponding to each of its ρ rows. Moreover, we think of the response of node l as formed of ρ sub-responses corresponding to projecting the node data on each row of Q_l .

Definition 2 (Information theoretic PIR). A PIR scheme achieves (perfect) information theoretic PIR iff $H(f|Q_j, j \in \gamma) = H(f)$, for all sets $\gamma \in [n]$, $|\gamma| = b$. Here, $H(\cdot)$ denotes the entropy function.

The objective is to design a linear PIR scheme that (i) allows the user to decode its requested file X^f and (ii) achieves information theoretic PIR with a low cPoP that does not depend on m . In the classical literature on PIR, the communication cost includes both the number of bits exchanged during the request and download stages. However, the query vectors depend only on the number of files in the system, while the response vectors depend on the size of the files, i.e. for a single sub-query, the query vector to a node consists of m symbols in $GF(q)$ while the response vector from one node is 1 symbol in $GF(q^w)$. In DSSs, and in the information-theoretic reformulation of this problem, the size of the files are assumed to be arbitrarily large, thus making the number of the files negligible with respect to the size of the files [10], i.e., w is much larger than m . Therefore, the download cost dominates the total communication cost. Hence, we will only consider the download communication cost, which we will refer to as the communication price of privacy (cPoP).

Definition 3. [cPoP] The communication Price of Privacy (cPoP) of a PIR scheme is the ratio of the total number of bits sent from the nodes to the user during the download stage to size of the requested file. This is the inverse of the PIR rate given in the literature.

NOMENCLATURE

n	Number of nodes in an (n, k, d) MDS code
k	Dimension of the codeword in an (n, k, d) MDS code
d	Distance of an (n, k, d) code
b	Number of colluding nodes
m	Number of files
ρ	Dimension of the scheme, number of rounds / subqueries / rows in query matrix
r	Remainder of the division of $n - k$ by k
β	Quotient of the division of $n - k$ by k
α	Number of subdivisions
$\mathbf{u}$	Random vector of size m
$\mathbf{w}_l$	Data on node l
$\mathbf{e}_f$	Indicator vector, the all-zero vector with one 1 in position f
$\mathbf{q}_{l,i}$	Query vector to Node l in sub-query i
$\mathbf{r}_{l,i}$	Response vector from Node l in sub-query i
Q_l	Query Matrix to Node l of dimension $\rho \times m\alpha$
E_l	0-1 matrix of dimension $\rho \times m\alpha$

III. MAIN RESULTS

In this section, we state our two main results. The proof of Theorem 1 is given in Section IV-B, the proof of Theorem 2 is given in Section V-B, and the proof of Theorem 3 is given in Section VI.

Theorem 1. Consider a DSS using an (n, k) MDS code over $GF(q)$, with $b = 1$, i.e. no collusion between the nodes. Then, the linear PIR scheme over $GF(q)$ described in Section IV-A achieves perfect PIR with $cPoP = \frac{1}{1-R}$, where $R = k/n$.

The existence of PIR schemes over large fields that can achieve $cPoP = \frac{1}{1-R}$ for $b = 1$ follows from Theorem 4 in [10]. The scheme in Section IV-A achieves the optimal $cPoP$ given in [26] as $m \rightarrow \infty$. We prove Theorem 1 by providing an explicit construction of the linear PIR scheme. The proposed PIR construction is over same field over which the code is designed and is universal in the sense that it depends only on the parameters n, k and b and not on the generator matrix of the code.

Theorem 2. Consider a DSS using an (n, k) MDS code over $GF(q)$, with b colluding nodes, $2 \leq b \leq d - 1$. Then, there exists an explicit linear PIR scheme over the same field that achieves perfect PIR with $cPoP = b + k$.

The next result is a generalization of Theorem 2 in which we describe a PIR scheme when $b \leq n - \delta k$, for any $\delta \geq 1$. Theorem 2 is a special case of Theorem 3 when $\delta = 1$, but we keep it for a better presentation of the proof. The optimal $cPoP$ for PIR on coded data with colluding nodes is still an open problem.

Fig. 2: cPoP versus rate following the scheme of Theorem 1, the scheme of Theorem 2, and the scheme for Theorem 3 for $n = 16$ and $b = 1$

		Sys. nodes				Parity nodes										
		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Stripes	r	1	1	2	3	4										
		2	2	3	4	1										
		3	3	4	1	2										
		4														
	k	5					1	1	1	1						
		6					2	2	2	2						
		7					3	3	3	3						
		8					4	4	4	4						
	k	9									1	1	1	1		
		10									2	2	2	2		
		11									3	3	3	3		
										4	4	4	4			
				k				k								

TABLE II: Example of the retrieval pattern for $(n, k, d) = (15, 4, 12)$. The $\alpha \times n$ entries of the table correspond to the $\alpha \times n$ coded symbols of the wanted file. All entries with same number, say j (also given the same color) are privately retrieved in the j^{th} sub-query. Note that there are $k = 4$ nodes, including the last $r = 3$ nodes, in every sub-query, that do not have any retrieved symbols. The responses of these nodes are used to decode the “interference” from all the files, needed to confuse the nodes about what is being requested. This interference is then cancelled out from the other sub-responses in order to decode the desired file symbols in each sub-query.

Theorem 3. For $b \leq n - \delta k$ colluding nodes, with $\delta = \lfloor \frac{n-b}{k} \rfloor$, we construct an explicit linear PIR scheme with $cPoP = \frac{b+\delta k}{\delta}$.

To illustrate the performance stated in the above three theorems, the price of privacy versus the rate of the storage code ($R = \frac{k}{n}$) when using the scheme of Theorem 1, the scheme of Theorem 2, and the scheme for Theorem 3 for $n = 16$ and $b = 1$ is shown in Figure 2. We notice that Theorem 1 shows much improvement on Theorem 2 and Theorem 3 for $b = 1$. We can also see that Theorem 3 improves on Theorem 2 when $\delta > 1$.

IV. PIR SCHEME CONSTRUCTION AND PROOF FOR $b = 1$

A. PIR scheme construction for $b = 1$

We describe here the PIR scheme referred to in Theorem 1. We assume WLOG that the MDS code is systematic. The PIR scheme uses the number of stripes $\alpha = d-1$ and the dimension $\rho = k$.² We write $\alpha = \beta k + r$ where, β and r are integers and $0 \leq r < k$ and $\beta \geq 0$.

The scheme consists of the user sending a $\rho \times m\alpha$ query matrix Q_l to each node $l, l = 1, \dots, n$. To form the query matrices, the user generates a $\rho \times m\alpha$ random matrix $U = [u_{ij}]$, whose elements are chosen uniformly at random from $GF(q)$, the same field over which the MDS code is defined. The query matrices have the following structure:

$$Q_l = U + E_{f,l}, \quad l = 1, \dots, n-r, \quad (3)$$

$$Q_l = U, \quad l = n-r+1, \dots, n. \quad (4)$$

U is the random component of the query aimed at confusing the nodes about the request, whereas $E_{f,l}$ is a deterministic matrix that depends on the index f of the requested file. The matrices $E_{f,l}$ add parts of the file X^f that is being retrieved to the responses of the nodes. The user can download $n-k$ symbols privately per sub-query, so the matrices $E_{f,l}$ add a symbol to the responses of $n-k$ of the nodes per sub-query. In this scheme, the user retrieves r symbols from the systematic nodes, and βk symbols from the parity nodes. Moreover, the retrieved symbols should not be redundant. The matrices $E_{f,l}$ are 0-1 matrices of dimensions $\rho \times m\alpha$, every row corresponds to a sub-query and every column corresponds to a stripe of a file. A “1” in the $(i, j)^{\text{th}}$ position of $E_{f,l}$ implies that, during the i^{th} sub-query, the j^{th} symbol on node l is being retrieved privately. The matrices $E_{f,l}$ are designed such that the following conditions hold:

- 1) Each row and column of the matrices $E_{f,l}$ contains at most one 1. The restriction on rows guarantees that we receive one coded symbol from a node, instead of the sum of several symbols. The column condition ensures that every symbol is only retrieved once, and thus, no retrieved symbol is redundant.
- 2) In each sub-query a 1 is added to the queries of exactly $n-k$ nodes, i.e., for $n-k$ of the matrices $E_{f,l}$ the i^{th} row contains a 1. This allows the user to decode a codeword from the MDS storage code, since k symbols are not altered, and subsequently decode $n-k$ symbols of the file X^f .
- 3) If j is the index of a stripe of the requested file f then exactly k of the matrices $E_{f,l}$ contain a 1 in column j . This ensures that we retrieve exactly k MDS coded symbols per row, which are needed to recover the original stripe.

Based on these desired retrieval patterns, we choose

$$E_{f,1} = \begin{bmatrix} \mathbf{0}_{k \times (f-1)\alpha} & I_{r \times r} & \mathbf{0}_{k \times \beta k} & \mathbf{0}_{k \times (m-f)\alpha} \\ \mathbf{0}_{(k-r) \times r} & \mathbf{0}_{(k-r) \times r} & \mathbf{0}_{(k-r) \times \beta k} & \mathbf{0}_{(k-r) \times (m-f)\alpha} \end{bmatrix}, \quad (5)$$

and $E_{f,l}, l = 2, \dots, k$, is obtained from matrix $E_{f,l-1}$ by a single downward cyclic shift of its row vectors.

²The parameters can be optimized to $\alpha = \frac{LCM(k, d-1)}{k}$ and $\rho = \frac{LCM(k, d-1)}{d-1}$, as was done in Example 1. But to simplify notation, we will take $\alpha = d-1$ and $\rho = k$.

	Sys. nodes		Parity nodes		
	1	2	3	4	5
Stripes					
1	1	2			
2			1	1	
3			2	2	

TABLE III: Retrieval pattern for a (5,2,4) code.

We divide the first βk parity nodes into β groups of k nodes each. All nodes in group s , i.e., nodes l where $sk + 1 \leq l \leq sk + k$, receive the same query matrix, such that

$$E_{f,l} = [\mathbf{0}_{k \times (f-1)\alpha + r + (s-1)k}, I_{k \times k}, \mathbf{0}_{k \times (\beta-s)k + (m-f)\alpha}]. \quad (6)$$

For the remaining r parity nodes we let

$$E_{f,l} = 0, \text{ for } l > \beta k + k,$$

and they hence all receive the same matrix U as a query.

Claim 1. Conditions 1, 2, and 3 are satisfied in the choice of the $E_{f,l}$ above.

Proof. 1) $E_{f,l}$ has at most one 1 in each row and column.

2) For the matrices $E_{f,l}$ sent to the parity nodes, all βk of them contain exactly one 1 in row i .

Since the k matrices $E_{f,l}$ for $1 \leq l \leq k$ sent to the systematic nodes are generated by cyclic row shifts of the matrix in (5), and it contains exactly r rows with a single 1, we see that r of these matrices contain a 1 in the i^{th} row. In total we have $\beta k + r = n - k$ matrices $E_{f,l}$ that contain a 1 in their i^{th} row.

3) The columns corresponding to the stripes of file f are in the range $(f-1)\alpha < j \leq f\alpha$. For $(f-1)\alpha < j \leq (f-1)\alpha + r$ we see that the k matrices of the form (5) contain exactly one 1 in column j . For $(f-1)\alpha + r + (s-1)k < j \leq (f-1)\alpha + r + sk$, $s = 1, \dots, \beta$, the k matrices $E_{f,l}$, for $sk + 1 < l \leq sk + k$, contain each one 1 in column j . $\square$

Example 2 (Retrieval pattern). Consider a DSS using an $(n, k, d) = (15, 4, 12)$ MDS code. Therefore, we have $\rho = k = 4$ sub-queries to each node. Also, the number of stripes is $\alpha = d - 1 = 11$. This gives $\beta = 2$ and $r = 3$. Table II gives the retrieval pattern of the PIR scheme, i.e., which file symbols are retrieved in each sub-query. The 11x15 entries in the table represents all the symbols of the desired file with each node being a column. The numbers (alternatively colors) in each entry indicate in which sub-query the specific symbol is retrieved.

Example 3 (Decoding). Now consider another example with $(n, k, d) = (5, 2, 4)$ with generator matrix $\Lambda = \begin{pmatrix} 1 & 0 & 1 & 1 & 1 \\ 0 & 1 & 1 & 2 & 3 \end{pmatrix}$, over $GF(5)$. Suppose the DSS stores $m = 3$ files, X^1, X^2, X^3 . Our goal is to construct a linear scheme that achieves perfect PIR against $b = 1$, with $cPoP = \frac{1}{1-R} = \frac{5}{3}$. The construction above gives $\alpha = d - 1 = 3$ and $\rho = k = 2$. Thus, a file X^i has the following array structure,

$$X^i = \begin{pmatrix} x_{11}^i & x_{12}^i \\ x_{21}^i & x_{22}^i \\ x_{31}^i & x_{32}^i \end{pmatrix}.$$

Therefore, we get $\beta = 1$ and $r = 1$. Suppose WLOG that the user wants file X^1 , i.e., $f = 1$. The user generates an 2×9 random matrix $U = [u_{ij}]$, whose elements are chosen uniformly at random from $GF(5)$. For the nodes $1, \dots, 4$, the query matrix $Q_l = U + E_{1,l}$, and $Q_5 = U$. Therefore, following (3), (4), (5), (6) and Table III we have

$$Q_1 = \begin{bmatrix} u_{11} + 1 & u_{12} & u_{13} & u_{14} & u_{15} & u_{16} & u_{17} & u_{18} & u_{19} \\ u_{21} & u_{22} & u_{23} & u_{24} & u_{25} & u_{26} & u_{27} & u_{28} & u_{29} \end{bmatrix},$$

$$Q_2 = \begin{bmatrix} u_{11} & u_{12} & u_{13} & u_{14} & u_{15} & u_{16} & u_{17} & u_{18} & u_{19} \\ u_{21} + 1 & u_{22} & u_{23} & u_{24} & u_{25} & u_{26} & u_{27} & u_{28} & u_{29} \end{bmatrix},$$

$$Q_3 = \begin{bmatrix} u_{11} & u_{12} + 1 & u_{13} & u_{14} & u_{15} & u_{16} & u_{17} & u_{18} & u_{19} \\ u_{21} & u_{22} & u_{23} + 1 & u_{24} & u_{25} & u_{26} & u_{27} & u_{28} & u_{29} \end{bmatrix},$$

$$Q_4 = \begin{bmatrix} u_{11} & u_{12} + 1 & u_{13} & u_{14} & u_{15} & u_{16} & u_{17} & u_{18} & u_{19} \\ u_{21} & u_{22} & u_{23} + 1 & u_{24} & u_{25} & u_{26} & u_{27} & u_{28} & u_{29} \end{bmatrix}.$$

The added 1s in certain positions of the query matrix are due to the addition of the matrix $E_{1,l}$. This construction achieves perfect privacy, since the only information any node l knows about f is through the query matrix Q_l , which is random and independent of f . Next, we want to illustrate how the user can decode the file symbols. Each node l sends back the length 2 vector, $\mathbf{r}_l = (r_{l1}, r_{l2}) = Q_l \mathbf{w}_l$, $l = 1, \dots, 5$, to the user. Recall that $\mathbf{w}_l$ is the data stored on node l . Consider the sub-responses of the 5 nodes to the first sub-query. They form the following linear system:

$$x_{11}^1 + I_1 = r_{11} \quad (7)$$

$$I_2 = r_{21} \quad (8)$$

$$x_{12}^1 + x_{22}^1 + I_1 + I_2 = r_{31} \quad (9)$$

$$x_{12}^1 + 2x_{22}^1 + I_1 + 2I_2 = r_{41} \quad (10)$$

$$I_1 + 3I_2 = r_{51}, \quad (11)$$

where $I_l = \mathbf{u}_1^T \mathbf{w}_l$, $l = 1, 2$, and $\mathbf{u}_1^T$ is the first row of U .

The user can first decode I_1 and I_2 from (8) and (11). Then, canceling out the values of I_1 and I_2 from the remaining equations, the user can solve for x_{11}^1, x_{12}^1 and x_{22}^1 . Similarly, the user can obtain x_{21}^1, x_{13}^1 and x_{23}^1 from the sub-responses to the second sub-query. This PIR scheme downloads 2 symbols from each server. Therefore, it has a $cPoP = \frac{10}{6} = \frac{5}{3}$, which matches the bound in Theorem 1.

B. Proof of Theorem 1

The following remarks from coding theory will be used on several occasions. For more background and proofs we refer to [33].

Remark 1. A linear $[n, k]$ code C is the set of all vectors $C := \{xG : x \in \mathbb{F}_q^k\} \subseteq \mathbb{F}_q^n$, where G is a generator matrix of the code. Therefore C is a k dimensional subvectorspace of $\mathbb{F}_q^n$ and any linear combination of codewords in C is again a codeword in C .

Remark 2. The following statements are equivalent.

1) A $[n, k]$ code C is MDS

- 2) For any generator matrix G_C of C any k subset of columns is full rank.
- 3) The code C can recover from up to $n - k$ erasures in any coordinates.

We prove Theorem 1 by showing that the scheme described in Section IV-A has the following properties.

Decodability: For any sub-query i , we sort the nodes into two groups to prove decodability. By the properties of the $E_{f,l}$ in Claim 1 exactly k nodes receive only the vector $\mathbf{u}_i^T$, the i^{th} row of U , as a query. And the user is aware of the indices of these nodes. For these nodes l , the received symbols are given by

$$r_{li} = \mathbf{u}_i^T \cdot \mathbf{w}_l.$$

Since every stored stripe is a codeword in C , by Remark 1, any linear combination of stripes will be a codeword too. We notice that r_{li} is indeed the l^{th} component of the codeword $\mathbf{r}'_i = \mathbf{u}_i^T \cdot (\mathbf{w}_1, \dots, \mathbf{w}_n)$. Since we have k of its components, we can recover the whole vector $\mathbf{r}'_i$ by Remark 2.

For the other nodes ℓ , the i^{th} sub-query is of the form $\mathbf{u}_i^T + \mathbf{e}_g$ where $\mathbf{e}_g$ is a standard basis vector, i.e., a single 1 has been added to the vector $\mathbf{u}_i$ in position g . The received symbol $r_{\ell i} = \mathbf{u}_i^T \cdot \mathbf{w}_\ell + \mathbf{e}_g \cdot \mathbf{w}_\ell = \mathbf{u}_i^T \cdot \mathbf{w}_\ell + w_\ell(g)$ therefore is the sum of the ℓ^{th} component of $\mathbf{r}'_i$ and the g^{th} symbol of $\mathbf{w}_\ell$. Since we have recovered $\mathbf{r}'_i$ from the k unaltered components, we can retrieve $w_\ell(g)$. Furthermore, the matrices $E_{f,l}$ are designed such that we retrieve exactly k symbols from every coded stripe of the f^{th} file. Using Remark 2 again allows us to retrieve all stripes of file X^f from these k symbols.

Privacy: Since $b = 1$, the only way a node l can learn information about f is from its own query matrix Q_i . By construction Q_i is statistically independent of f and this scheme achieves perfect privacy.

cPoP: Every node $l \in [n]$ responds with $\rho = k$ symbols. Therefore, the total number of symbols downloaded by the user is kn . Therefore, $cPoP = \frac{kn}{k(n-k)} = \frac{1}{1-R}$.

V. PIR SCHEME CONSTRUCTION AND PROOF FOR $b \leq d - 1$

A. PIR scheme construction for $b \leq d - 1$

In this section, we will describe the general PIR scheme that achieves $cPOP = b + k$ by specifying the query matrices to each node. This scheme requires $b \leq d - 1$. To simplify the description of the scheme, we will assume $b = d - 1$.³ The scheme has dimension $\rho = k$, i.e., it consists of $\rho = k$ sub-queries. Moreover, the scheme requires no subdivisions, i.e., the number of stripes $\alpha = 1$. Since there are no subdivisions, we simplify further the notation and write $x_{i1}^j = x_i^j$ to denote the i^{th} systematic symbol of file X^j , where $j = 1, \dots, m$. Denote by f the index of the file that the user wants, i.e., the user wants to retrieve file X^f . WLOG, we assume the MDS code is systematic.

In the i^{th} sub-query, $i = 1, \dots, k$, the proposed PIR scheme retrieves systematic symbol x_i^f of the wanted file X^f . So, by the completion of the scheme, the user will have all the

k symbols forming the file. In sub-query i , the user creates $d - 1$ random (column) vectors $\mathbf{u}_{1,i}, \dots, \mathbf{u}_{d-1,i}$, of dimension m each, whose elements are chosen uniformly at random from $GF(q)$. Recall that the generator matrix for any systematic (n, k, d) MDS code is of the form

$$G_{k \times n} = [I_{k \times k} \mid P_{k \times (d-1)}], \quad (12)$$

where P is a $k \times d - 1$ matrix describing the parity nodes. A parity check matrix for this code is then given by

$$H_{(d-1) \times n} = [-P^T \mid I_{(d-1) \times (d-1)}].$$

Define U_i to be the $m \times (d - 1)$ matrix with its columns being the $b = d - 1$ random vectors used in sub-query i , i.e.,

$$U_i_{m \times d-1} = [\mathbf{u}_{1,i}, \mathbf{u}_{2,i}, \dots, \mathbf{u}_{d-1,i}].$$

Now for each sub-query the user generates m random codewords in the dual code by multiplying the random matrix $U_i \in GF(q)^{m \times (d-1)}$ by the parity check matrix $H_{(d-1) \times n}$ to calculate

$$U_i H_{(d-1) \times n} = [\mathbf{q}'_{1,i}, \dots, \mathbf{q}'_{n,i}].$$

Note that each row of $U_i H$ is a codeword in the dual of the MDS code used to store the data.

For $i = 1, \dots, k$, let $\mathbf{q}_{l,i}$ be the i^{th} sub-query vector to node l with $l = 1, \dots, n$. These query vectors are chosen as follows:

$$\mathbf{q}_{l,i} = \begin{cases} \mathbf{q}'_{l,i} + \mathbf{e}_f, & \text{if } l = i, \\ \mathbf{q}'_{l,i}, & \text{otherwise,} \end{cases} \quad (13)$$

where $\mathbf{e}_f$ is the standard basis vector with a single 1 in position f .

Therefore, the response of node l to the i^{th} sub-query, denoted by $r_{l,i}$, is given by (2) and can be written as

$$r_{l,i} = \mathbf{q}_{l,i}^T \mathbf{w}_l, \quad (14)$$

where $\mathbf{w}_l$ is the vector representing the data stored on node l .

We will give an example.

Example 4. Next, we illustrate this scheme through an example. Consider a DSS using the following systematic $(5, 3, 3)$ MDS code with generator matrix

$$\Lambda = \begin{bmatrix} 1 & 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 1 & 2 \\ 0 & 0 & 1 & 1 & 3 \end{bmatrix}.$$

Suppose the system is storing $m = 3$ files, $X^1 = \begin{bmatrix} a_1 \\ a_2 \\ a_3 \end{bmatrix}$,

$X^2 = \begin{bmatrix} b_1 \\ b_2 \\ b_3 \end{bmatrix}$ and $X^3 = \begin{bmatrix} c_1 \\ c_2 \\ c_3 \end{bmatrix}$. Then, the data is stored

on the different nodes in the DSS as described in table IV.

Our goal is to construct a linear scheme that achieves perfect PIR against $b = 2$ colluding nodes with $cPoP = k + b = 3 + b$. The scheme will consist of $\rho = k = 3$ sub-queries.

³If $b < d - 1$, only $b + k$ nodes, say the first $b + k$, are queried.

node 1	node 2	node 3	node 4	node 5
a_1	a_2	a_3	$a_1 + a_2 + a_3$	$a_1 + 2a_2 + 3a_3$
b_1	b_2	b_3	$b_1 + b_2 + b_3$	$b_1 + 2b_2 + 3b_3$
c_1	c_2	c_3	$c_1 + c_2 + c_3$	$c_1 + 2c_2 + 3c_3$

TABLE IV: (5,3,3) DSS

Suppose WLOG that the user wants file X^1 , i.e., $f = 1$. We will consider the first sub-query and the remaining sub-queries (i.e. sub-queries 2 and 3) follow similarly. The user creates 2 random vectors $\mathbf{u}_{1,1}, \mathbf{u}_{2,1}$ of dimension $m = 3$ each. $U_1 = [\mathbf{u}_{1,1}, \mathbf{u}_{2,1}]$. The dual code will have a generator matrix

$$H_{(n-k) \times n} = \begin{bmatrix} -1 & -1 & -1 & 1 & 0 \\ -1 & -2 & -3 & 0 & 1 \end{bmatrix}.$$

The sub-query vectors to nodes 1 to 5 are the following respectively

$$\mathbf{q}_{1,1} = -\mathbf{u}_{1,1} - \mathbf{u}_{2,1} + \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \quad (15)$$

$$\mathbf{q}_{2,1} = -\mathbf{u}_{1,1} - 2\mathbf{u}_{2,1}, \quad (16)$$

$$\mathbf{q}_{3,1} = -\mathbf{u}_{1,1} - 3\mathbf{u}_{2,1}, \quad (17)$$

$$\mathbf{q}_{4,1} = \mathbf{u}_{1,1}, \quad (18)$$

$$\mathbf{q}_{5,1} = \mathbf{u}_{2,1}. \quad (19)$$

Next, we want to show that the user can decode its requested file correctly. The nodes send back the length 3 vectors, $\mathbf{r}_l = (r_{l,1}, r_{l,2}, r_{l,3})$, $l = 1, \dots, 5$, to the user. Consider the first symbol in each of the vectors $r_{l,1}$, which form the following linear system:

$$a_1 - I_{11} - I_{12} = r_{1,1} \quad (20)$$

$$-I_{21} - 2I_{22} = r_{2,1} \quad (21)$$

$$-I_{31} - 3I_{32} = r_{3,1} \quad (22)$$

$$I_{11} + I_{21} + I_{31} = r_{4,1} \quad (23)$$

$$I_{12} + 2I_{22} + 3I_{32} = r_{5,1} \quad (24)$$

where $I_{lj} = \mathbf{u}_{j,1}^T \mathbf{w}_l$, for $l = 1, 2, 3$ denoting the node index, and $j = 1, 2$ denoting the random vector. In analogy with the interference alignment literature [34], [35], one can think of a_1 as the signal to be decoded and $I_{11}, I_{12}, I_{21}, I_{22}, I_{31}, I_{32}$ as the interference. And we can notice that if we sum up eqs. (20) to (24), we get a_1 . This PIR scheme downloads 3 packets from each server. Therefore, it has a $cPoP = \frac{5 \times 3}{3} = 5$.

As mentioned for $b < d - 1$ only $b + k$ nodes are queried as shown in the next example. We will revisit this example in the next section and present a more efficient scheme when explaining Theorem 3.

Example 5. Consider the (6,2,5) MDS code in table V, where $\Lambda = \begin{bmatrix} 1 & 0 & 1 & 1 & 1 & 1 \\ 0 & 1 & 1 & 2 & 3 & 4 \end{bmatrix}$. The goal here is to construct a linear scheme that achieves perfect PIR against $b = 2$ colluding nodes with $cPoP = k + b = 4$. Assume WLOG the user wants file X^f . The scheme will consist of $\rho = 2$ sub-queries. We will consider the first sub-query and the second sub-query follows similarly.

node 1	node 2	node 3	node 4	node 5	node 6
a_1	b_1	$a_1 + b_1$	$a_1 + 2b_1$	$a_1 + 3b_1$	$a_1 + 4b_1$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
a_m	b_m	$a_m + b_m$	$a_m + 2b_m$	$a_m + 3b_m$	$a_m + 4b_m$

TABLE V: (6,2,5) DSS

In this case, the user will query only 4 nodes, WLOG the first 4 nodes, with generator matrix $G = \begin{bmatrix} 1 & 0 & 1 & 1 \\ 0 & 1 & 1 & 2 \end{bmatrix}$. As described in section V-A, the user creates 2 random vectors $\mathbf{u}_{1,1}, \mathbf{u}_{2,1}$ of dimension m each, and as in the previous example, forms $U_1 = [\mathbf{u}_{1,1}, \mathbf{u}_{2,1}]$. The dual code of G will have a generator matrix

$$H_{(n-k) \times n} = \begin{bmatrix} -1 & -1 & 1 & 0 \\ -1 & -2 & 0 & 1 \end{bmatrix}.$$

The sub-query vectors to nodes 1 to 4 are the following respectively

$$\mathbf{q}_{1,1} = -\mathbf{u}_{1,1} - \mathbf{u}_{2,1} + \mathbf{e}_f, \quad (25)$$

$$\mathbf{q}_{2,1} = -\mathbf{u}_{1,1} - 2\mathbf{u}_{2,1}, \quad (26)$$

$$\mathbf{q}_{3,1} = \mathbf{u}_{1,1}, \quad (27)$$

$$\mathbf{q}_{4,1} = \mathbf{u}_{2,1}. \quad (28)$$

The nodes will respond to the user by projecting their data on the query matrices. With inspection of the queries, we can see that the user will be able to decode a_f from the first sub-query, and similarly decode b_f from the second sub-query. This achieves a $cPoP = 4$.

B. Proof of Theorem 2

We prove Theorem 2 by showing that the scheme described in Section V-A ensures decodability and privacy. The main ingredient in the proof, which makes it different from the proof of Theorem 1, is that the scheme does not require the user to decode all the interference terms. Recall that the user wants to retrieve file X^f . We will prove that the user can retrieve x_i^f in the i^{th} sub-query. An alternative proof of Theorem 2 is shown in the Appendix.

Decodability:

The response of node $l = 1, \dots, n$ to the i^{th} sub-query is given by

$$r_{l,i} = \mathbf{q}_{l,i}^T \mathbf{w}_l. \quad (29)$$

To decode x_i^f , the user sums the responses of all the nodes to the i^{th} sub-query, i.e., it computes $\sum_{l=1}^n r_{l,i}$.

Claim 2. $\sum_{l=1}^n r_{l,i} = x_i^f$

Proof.

$$\sum_{l=1}^n r_{l,i} = \text{tr}((U_i H)^T \mathcal{X} G) + \mathbf{e}_f^T \mathbf{w}_i \quad (30)$$

$$= \text{tr}(U_i H G^T \mathcal{X}^T) + x_i^f \quad (31)$$

$$= x_i^f. \quad (32)$$

where $\text{tr}(\cdot)$ is the trace operator. Equation (30) follows directly from the scheme, equation (31) follows from the fact

Fig. 3: cPoP versus rate when $n = 16$ and number of colluding nodes $b = 1, 3, 5$, following the scheme in Theorem 2 and Theorem 3. We notice that as the number of colluding nodes increases, the improvement the scheme in Theorem 3 has over the scheme in Theorem 2 grows.

that $\text{tr}(A^T B) = \text{tr}(AB^T)$, and equation (32) follows from the fact that $\text{tr}(U_i H G^T \mathcal{X}^T) = 0$ since $H G^T = 0$. $\square$

Privacy: Recall that $f \in [m]$ is the index of the file wanted by the user. Let S_b be a subset of cardinality b of $[n]$ representing the set of b colluding nodes. We define Q_{S_b} to be the set of query vectors (or matrices) incoming to the b nodes indexed by S_b . We want to show that when b spies collude, they cannot learn any information about f , i.e., $H(f|Q_{S_b}) = H(f)$, for any possible set of colluding nodes $S_b \subset [n], |S_b| = b$.

$$H(f, Q_{S_b}) = H(f, Q_{S_b}) \quad (33)$$

$$H(Q_{S_b}) + H(f|Q_{S_b}) = H(f) + H(Q_{S_b}|f) \quad (34)$$

$$H(f|Q_{S_b}) = H(f) + H(Q_{S_b}|f) - H(Q_{S_b}) \quad (35)$$

$$= H(f) - H(Q_{S_b}) + H(Q_{S_b}|f) \quad (36)$$

$$- \underbrace{H(Q_{S_b}|f, U_i)}_{=0} \quad (37)$$

$$= H(f) - H(Q_{S_b}) + I(Q_{S_b}, U_i|f) \quad (38)$$

$$= H(f) - H(Q_{S_b}) + H(U_i|f) - H(U_i|Q_{S_b}, f) \quad (39)$$

$$= H(f) - H(Q_{S_b}) + H(U_i) \quad (40)$$

$$= H(f).$$

Where the equality in equation (36) follows from the fact that $H(Q_{S_b}|f, U_i) = 0$, since the query vectors are a deterministic function of f and U_i . Equation (39) follows from $H(U_i|f) = H(U_i)$, since the random matrix U_i is independent of the file index f . Moreover, $H(U_i|Q_{S_b}, f) = 0$ since by (48), given f , U_i can be decoded from Q_{S_b} due to the MDS property of the code. Lastly, in (40) $H(Q_{S_b}) = H(U_i) = m$ follows again from (48) and the MDS property of the code.

VI. PIR SCHEME CONSTRUCTION FOR $b \leq n - \delta k$

Let $\delta = \lfloor \frac{n-b}{k} \rfloor$. We can see that for $\delta = 1$, this simplifies to Theorem 2. Figure 3 shows a comparison of the construction of Theorem 2 and Theorem 3.

Example 6. Consider again the $(6,2,5)$ MDS code in table V and $b = 2$. We notice that in example 5, we did not use nodes 5 and 6, and achieved $cPoP = 4$. Now we will show how we can use those nodes and achieve a lower $cPoP = 3$. Assume the user wants file X^f .

We first choose the last $b = 2$ of the parity nodes to be common nodes. Then, we split the rest of the $\delta k = 4$ nodes into $\delta = 2$ groups of $k = 2$ nodes each. We then consider the two punctured codes, each with a 2×4 generator matrix, that intersect in the common nodes. Here, we pick the two subcodes consisting of nodes 1, 2, 5, 6 and 3, 4, 5, 6, respectively. The punctured codes will have the following generator matrices:

$$G_1 = [B_1 | P] = \begin{bmatrix} 1 & 0 & 1 & 1 \\ 0 & 1 & 3 & 4 \end{bmatrix},$$

and

$$G_2 = [B_2 | P] = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & 2 & 3 & 4 \end{bmatrix}.$$

The user can transform the generator matrices of the punctured codes into systematic form by multiplying by the inverse of the $k \times k = 2 \times 2$ matrix formed by the non-common nodes. In this example, we can see that G_2 is not in systematic form, so we multiply by the inverse of the 2×2 sub-matrix formed by nodes 3 and 4, i.e. $\begin{bmatrix} 1 & 1 \\ 1 & 2 \end{bmatrix}$ to get

$$G_2 = \begin{bmatrix} 1 & 0 & -1 & -2 \\ 0 & 1 & 2 & 3 \end{bmatrix}.$$

The parity check matrices of the $(4,2,3)$ MDS codes generated by G_1 , and G_2 are

$$H_1 = [-P_1^T | I_{2 \times 2}] = \begin{bmatrix} -1 & -3 & 1 & 0 \\ -1 & -4 & 0 & 1 \end{bmatrix},$$

$$H_2 = [-P_2^T | I_{2 \times 2}] = \begin{bmatrix} 1 & -2 & 1 & 0 \\ 2 & -3 & 0 & 1 \end{bmatrix}.$$

In this example, we will not subdivide the files into stripes, and one sub-query is required in which we will decode both parts of the file. For this reason, we will remove the subscript for simplicity. Similar to the scheme in section V-A, the user generates 2 random (column) vectors $\mathbf{u}_1$ and $\mathbf{u}_2$, of length m each, whose elements are chosen uniformly at random from $GF(q)$. Define U to be the $m \times 2$ matrix with its columns being the $b = 2$ random vectors $\mathbf{u}_1$ and $\mathbf{u}_2$, i.e.,

$$U_{m \times 2} = [\mathbf{u}_1 \quad \mathbf{u}_2].$$

Now the user generates m random codewords in the dual codes by multiplying the random matrix U by the parity check matrix H_1 to calculate

$$UH_1 = [\mathbf{q}'_1, \mathbf{q}'_2, \mathbf{u}_1, \mathbf{u}_2], \quad (41)$$

and multiplying the random matrix U by the parity check matrix H_2 to calculate

$$UH_2 = [\mathbf{q}'_3, \mathbf{q}'_4, \mathbf{u}_1, \mathbf{u}_2], \quad (42)$$

The query vectors to nodes 1, 2, 3, and 4 are chosen as follows:

$$\mathbf{q}_l = \begin{cases} \mathbf{q}'_l + \mathbf{e}_f, & \text{if } l = 1 \pmod k, \\ \mathbf{q}'_l, & \text{otherwise.} \end{cases} \quad (43)$$

The query vectors to nodes 5, 6 are

$$\mathbf{q}_l = \mathbf{u}_l. \quad (44)$$

Decodability:

Each of the punctured codes is coded as in section V-A. Based on the decodability proved in section V-B, the user can decode a_f from the first code, and $a_f + b_f$ from the second code. Hence, the user retrieves file X^f .

Privacy:

The queries are sent by projecting the random matrix U on the matrix

$$H = \begin{bmatrix} -P_1^T & -P_2^T & I \end{bmatrix}.$$

This is a dual of the code generated by the matrix

$$G = \begin{bmatrix} B_1 & 0 & P \\ 0 & B_2 & P \end{bmatrix}.$$

The dual code is an $(6, 4, 3)$ MDS code, and thus the queries are sent using a $(6, 2, 5)$ MDS code.

This means that any 2 queries are linearly independent and thus if this is private against 2 colluding nodes.

The user contacts 6 nodes, to download 2 information parts, thus the price of privacy of this is $cPoP = \frac{b+\delta k}{\delta} = \frac{6}{2} = 3$.

A. General Proof of Theorem 3

We assume the user wants file X^f . Assume $n = b + \delta k$.⁴ The user uses the last b nodes, $n_{\delta k+1}, \dots, n_{\delta k+b}$ as common nodes. The rest of the nodes will be divided into δ groups, $j = 1, \dots, \delta$, of k nodes each. This forms δ punctured $(b+k, k)$ MDS codes, each with a generator matrix $[B_j \mid P_{b \times b}]$ which can be transformed to systematic form $G_j = [I_{k \times k} \mid P_j]$ by multiplying by the inverse of the $k \times k$ matrix B_j . Here we will use $\alpha = \delta$ subdivisions and k queries.⁵

We calculate the parity check matrix of the δ codes.

$$H_j = \begin{bmatrix} -P_j^T & I_{b \times b} \end{bmatrix}.$$

Now for each sub-query, $i, i = 1, \dots, k$, the user generates m random codewords by multiplying the random matrix $U_i =$

$\begin{bmatrix} \mathbf{u}_{1,i} & \dots & \mathbf{u}_{b,i} \end{bmatrix} \in GF(q)^{m \times b}$ by the parity check matrix H_j of subcode j

$$U_i H_j = \begin{bmatrix} \mathbf{q}'_{1+(j-1)k,i}, \mathbf{q}'_{2+(j-1)k,i}, \dots, \mathbf{q}'_{jk,i}, \mathbf{u}_{1,i}, \dots, \mathbf{u}_{b,i} \end{bmatrix}.$$

For the nodes $l = 1, \dots, \delta k$, the query vectors in sub-query i are as follows:

$$\mathbf{q}_{l,i} = \begin{cases} \mathbf{q}'_{l,i} + \mathbf{e}_{(f-1)\delta+j}, & \text{if } l = k - (j-1) + i, \\ \mathbf{q}'_{l,i}, & \text{otherwise.} \end{cases} \quad (45)$$

For the nodes $l = \delta k + 1, \dots, \delta k + b$, the query vectors in sub-query i are the columns of U_i

$$\mathbf{q}_{l,i} = \mathbf{u}_{l-\delta k,i}. \quad (46)$$

Decodability:

For each subcode j , we follow the scheme of Theorem 2 to obtain the j^{th} stripe of file x^f . Subsequently, the user is able to decode the file x^f .

Privacy: The queries are generating by multiplying the random matrix U on the matrix

$$H = \begin{bmatrix} -P_1^T & -P_2^T & \dots & -P_\delta^T & I \end{bmatrix}.$$

This is a dual of the code generated by the matrix

$$G = \begin{bmatrix} B_1 & 0 & \dots & 0 & P \\ 0 & B_2 & \dots & 0 & P \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & B_\delta & P \end{bmatrix}.$$

We see that the code generated by G is an $(\delta k + b, \delta k, b + 1)$ MDS code, and thus the queries are sent using an $(\delta k + b, b, \delta k + 1)$ MDS code.

This means that any b queries are linearly independent and thus this is private against b colluding nodes.

The user contacts $b + \delta k$ nodes, to download δ information parts, thus the price of privacy of this is $cPoP = \frac{b+\delta k}{\delta}$.

VII. COMPARISON TO FUNDAMENTAL BOUNDS

Our scheme achieves the fundamental bounds currently known for infinite number of files and 1 spy node, i.e. no collusion. The lowest achievable price of privacy of a storage system with replicated databases is given in [25] to be $\frac{1-(1/n)^m}{1-(1/n)}$ which asymptotically approaches $\frac{n}{n-1}$ as $m \rightarrow \infty$. If we apply our PIR scheme for a replicated database, the $cPoP = \frac{1}{1-R} = \frac{n}{n-1}$ which is the limit of the lower bound.

The lower bound for an (n, k, d) MDS-coded database was derived in [26] to be $\frac{1-(k/n)^m}{1-(k/n)}$, which asymptotically approaches $\frac{n}{n-k} = \frac{1}{1-R}$ as $m \rightarrow \infty$, and is again the cPoP achieved by our construction.

⁴If $b < n - \delta k$, only $b + \delta k$ nodes are queried.

⁵The parameters can be optimized to $\alpha = \frac{LCM(\delta, k)}{k}$ and $\rho = \frac{LCM(\delta, k)}{\delta}$, as was done in Example 6. But to simplify notation, we will take $\alpha = \delta$ and $\rho = k$.

VIII. CONCLUSION

We studied the problem of constructing PIR schemes with low communication cost for requesting data from a DSS that uses MDS codes. Some nodes in the DSS may be spies who will report to a third party, such as an oppressive regime, which data is being requested by a user. The objective is to allow the user to obtain its requested data without revealing any information on the identity of the data to the nodes. We constructed PIR schemes against non-colluding nodes that achieve the information theoretic limit on the download communication cost for linear schemes. An important property of these schemes is their universality since they depend on the code rate, but not on the MDS code itself. When there is b -collusion with $2 \leq b \leq n - k$, we devised linear PIR schemes that have download cost equal to $b + k$ per unit of requested data.

REFERENCES

- [1] R. Tajeddine and S. El Rouayheb, "Private information retrieval from mds coded data in distributed storage systems," in *Information Theory (ISIT), 2016 IEEE International Symposium on*, pp. 1411–1415, IEEE, 2016.
- [2] B. Chor, O. Goldreich, E. Kushilevitz, and M. Sudan, "Private information retrieval," in *IEEE Symposium on Foundations of Computer Science*, pp. 41–50, 1995.
- [3] B. Chor, E. Kushilevitz, O. Goldreich, and M. Sudan, "Private information retrieval," *Journal of the ACM (JACM)*, vol. 45, no. 6, pp. 965–981, 1998.
- [4] A. Beimel, Y. Ishai, and E. Kushilevitz, "General constructions for information-theoretic private information retrieval," *Journal of Computer and System Sciences*, vol. 71, no. 2, pp. 213–247, 2005.
- [5] S. Yekhanin, "Private information retrieval," *Communications of the ACM*, vol. 53, no. 4, pp. 68–73, 2010.
- [6] A. Beimel and Y. Ishai, "Information-theoretic private information retrieval: A unified construction," in *Automata, Languages and Programming*, pp. 912–926, Springer, 2001.
- [7] A. Beimel, Y. Ishai, and T. Malkin, "Reducing the servers computation in private information retrieval: PIR with preprocessing," in *Advances in Cryptology—CRYPTO 2000*, pp. 55–73, Springer, 2000.
- [8] A. Beimel, Y. Ishai, E. Kushilevitz, and J.-F. Raymond, "Breaking the $o(n^{1/(2k-1)})$ barrier for information-theoretic private information retrieval," in *The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings.*, pp. 261–270, IEEE, 2002.
- [9] N. Shah, K. Rashmi, and K. Ramchandran, "One extra bit of download ensures perfectly private information retrieval," in *Information Theory (ISIT), 2014 IEEE International Symposium on*, pp. 856–860, IEEE, 2014.
- [10] T. Chan, S.-W. Ho, and H. Yamamoto, "Private information retrieval for coded storage," in *Information Theory (ISIT), 2015 IEEE International Symposium on*, pp. 2842–2846, IEEE, June 2015.
- [11] Z. Dvir and S. Gopi, "2 server PIR with sub-polynomial communication," in *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC '15*, (New York, NY, USA), pp. 577–584, ACM, 2015.
- [12] S. Yekhanin, "Towards 3-query locally decodable codes of subexponential length," *Journal of the ACM (JACM)*, vol. 55, no. 1, pp. 1–16, 2008.
- [13] K. Efremenko, "3-query locally decodable codes of subexponential length," *SIAM Journal on Computing*, vol. 41, no. 6, pp. 1694–1703, 2012.
- [14] H. Sun and S. A. Jafar, "Blind interference alignment for private information retrieval," in *Information Theory (ISIT), 2016 IEEE International Symposium on*, pp. 560–564, IEEE, 2016.
- [15] E. Kushilevitz and R. Ostrovsky, "Replication is not needed: Single database, computationally-private information retrieval," in *FOCS*, p. 364, IEEE, 1997.
- [16] G. Fanti and K. Ramchandran, "Multi-server private information retrieval over unsynchronized databases," in *Communication, Control, and Computing (Allerton), 2014 52nd Annual Allerton Conference on*, pp. 437–444, 2014.
- [17] Y. Ishai, E. Kushilevitz, R. Ostrovsky, and A. Sahai, "Batch codes and their applications," in *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pp. 262–271, ACM, 2004.
- [18] D. Augot, F. Levy-Dit-Vehel, and A. Shikfa, "A storage-efficient and robust private information retrieval scheme allowing few servers," in *Cryptography and Network Security*, pp. 222–239, Springer, 2014.
- [19] A. Beimel and Y. Stahl, "Robust information-theoretic private information retrieval," in *International Conference on Security in Communication Networks*, pp. 326–341, Springer, 2002.
- [20] R. Tajeddine and S. El Rouayheb, "Robust private information retrieval on coded data," in *Information Theory (ISIT), 2017 IEEE International Symposium on*, IEEE, 2017.
- [21] A. Fazeli, A. Vardy, and E. Yaakobi, "Codes for distributed PIR with low storage overhead," in *Information Theory (ISIT), 2015 IEEE International Symposium on*, pp. 2852–2856, 2015.
- [22] S. Blackburn and T. Etzion, "PIR array codes with optimal pir rate," *arXiv preprint arXiv:1607.00235*, 2016.
- [23] R. Tajeddine, S. El Rouayheb, "Private Information Retrieval from MDS Coded data in Distributed Storage Systems (extended version)," 2016. <http://www.ece.iit.edu/~salim/PIRv2.pdf>.
- [24] H. Sun and S. A. Jafar, "The capacity of private information retrieval," *IEEE Transactions on Information Theory*, pp. 4075 – 4088, 2017.
- [25] H. Sun and S. A. Jafar, "The capacity of private information retrieval with colluding databases," in *Signal and Information Processing (GlobalSIP), 2016 IEEE Global Conference on*, pp. 941–946, IEEE, 2016.
- [26] K. Banawan and S. Ulukus, "The capacity of private information retrieval from coded databases," *arXiv preprint arXiv:1609.08138*, 2016.
- [27] R. Freij-Hollanti, O. Gnille, C. Hollanti, and D. Karpuk, "Private information retrieval from coded databases with colluding servers," *arXiv preprint arXiv:1611.02062*, 2016.
- [28] 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 *Information Theory (ISIT), 2017 IEEE International Symposium on*, pp. 1893–1897.
- [29] R. Tajeddine, O. W. Gnille, D. Karpuk, R. Freij-Hollanti, C. Hollanti, and S. E. Rouayheb, "Private information retrieval schemes for coded data with arbitrary collusion patterns," in *Information Theory (ISIT), 2017 IEEE International Symposium on*, pp. 1908–1912.
- [30] S. Kumar, E. Rosnes, and A. G. i Amat, "Private information retrieval in distributed storage systems using an arbitrary linear code," in *Information Theory (ISIT), 2017 IEEE International Symposium on*, pp. 1421–1425, IEEE, 2017.
- [31] Q. Wang and M. Skoglund, "Symmetric private information retrieval for mds coded distributed storage," in *Communications (ICC), 2017 IEEE International Conference on*, pp. 1–6, IEEE, 2017.
- [32] K. Banawan and S. Ulukus, "The capacity of private information retrieval from byzantine and colluding databases," *arXiv preprint arXiv:1706.01442*, 2017.
- [33] J. v. Lint, *Introduction to Coding Theory*. Springer Berlin, 2013.
- [34] Y. Wu and A. G. Dimakis, "Reducing repair traffic for erasure coding-based storage via interference alignment," in *Information Theory, 2009. ISIT 2009. IEEE International Symposium on*, pp. 2276–2280, IEEE, 2009.
- [35] C. Suh and K. Ramchandran, "Exact-repair mds codes for distributed storage using interference alignment," in *Information Theory Proceedings (ISIT), 2010 IEEE International Symposium on*, pp. 161–165, IEEE, 2010.

APPENDIX

Alternative Proof of Theorem 2

To simplify the description of the scheme, we will assume $b = n - k$. The scheme has dimension $\rho = k$, i.e., it consists of $\rho = k$ sub-queries. Moreover, the scheme requires no subdivisions, i.e., the number of stripes $\alpha = 1$. Since there are no subdivisions, we simplify further the notation and write $x_{i1}^j = x_i^j$ to denote the i^{th} systematic symbol of file X^j , where $j = 1, \dots, m$. Denote by f the index of the file that the user wants, i.e., the user wants to retrieve file X^f .

In the i^{th} sub-query, $i = 1, \dots, k$, the proposed PIR scheme retrieves systematic symbol x_i^f of the wanted file X^f . So, by

the completion of the scheme, the user will have obtained all the k symbols forming the file.

In sub-query i , the user creates b random (column) vectors $\mathbf{u}_{1,i}, \dots, \mathbf{u}_{b,i}$, of dimension m each, whose elements are chosen uniformly at random from $GF(q)$. Define U_i to be the $m \times d-1$ matrix with its rows being the $b = n-k$ random vectors used in sub-query i , i.e.,

$$U_i = [\mathbf{u}_{1,i}, \mathbf{u}_{2,i}, \dots, \mathbf{u}_{d-1,i}].$$

Recall that the generator matrix of the MDS code is

$$\Lambda = \begin{bmatrix} I_{k \times k} & \begin{bmatrix} \lambda_{1,k+1} & \dots & \lambda_{1,n} \\ \vdots & \vdots & \vdots \\ \lambda_{k,k+1} & \dots & \lambda_{k,n} \end{bmatrix} \end{bmatrix}_{k \times n}.$$

We write $\Lambda = [I \mid P]$, where

$$P = \begin{bmatrix} \lambda_{1,k+1} & \dots & \lambda_{1,n} \\ \vdots & \vdots & \vdots \\ \lambda_{k,k+1} & \dots & \lambda_{k,n} \end{bmatrix}_{k \times d-1}.$$

We denote by $\mathbf{p}_j^T$ the j^{th} row of P . Let $\mathbf{e}_f^T = [\mathbf{0}_{1 \times (f-1)} \quad 1 \quad \mathbf{0}_{1 \times (m-f)}]$.

For a systematic node l , the user sends the sub-query vector:

$$\mathbf{q}_{l,i} = \begin{cases} \lambda_{l,k+1}\mathbf{u}_{1,i} + \dots + \lambda_{l,n}\mathbf{u}_{d-1,i} + \mathbf{e}_f, & \text{if } l = i, \\ \lambda_{l,k+1}\mathbf{u}_{1,i} + \dots + \lambda_{l,n}\mathbf{u}_{d-1,i}, & \text{otherwise.} \end{cases} \quad (47)$$

This translates to

$$\mathbf{q}_{l,i} = \begin{cases} U_i \mathbf{p}_l + \mathbf{e}_f, & \text{if } l = i, \\ U_i \mathbf{p}_l, & \text{otherwise.} \end{cases} \quad (48)$$

For the parity nodes $l = k+1, \dots, n = k+b$, the i^{th} sub-query vector is given by,

$$\mathbf{q}_{l,i} = \mathbf{u}_{l-k,i}. \quad (49)$$

Therefore, the response of node l to the i^{th} sub-query, denoted by $r_{l,i}$, is given by (2) and can be written as

$$r_{l,i} = \mathbf{q}_{l,i}^T \mathbf{w}_l, \quad (50)$$

where $\mathbf{w}_l$ is the vector representing the data stored on node l .

We prove Theorem 2 by showing that the scheme described in Section V-A ensures decodability and privacy. The main ingredient in the proof, which makes it different from the proof of Theorem 1, is that the scheme does not require the user to decode all the interference terms.

Recall that the user wants to retrieve file X^f . We will prove that the user can retrieve x_i^f in the i^{th} sub-query.

Decodability:

From (48) and (55), the response of systematic node l to the i^{th} sub-query is given by

$$r_{l,i} = \begin{cases} \mathbf{p}_l^T U_i^T \mathbf{w}_l + x_i^f = \mathbf{w}_l^T U_i \mathbf{p}_l + x_i^f & \text{if } l = i, \\ \mathbf{p}_l^T U_i^T \mathbf{w}_l = \mathbf{w}_l^T U_i \mathbf{p}_l & \text{otherwise.} \end{cases} \quad (51)$$

Notice that $\mathbf{w}_l^T U_i^T \mathbf{p}_l$ is the l^{th} diagonal element of $\mathcal{X} U_i^T P^T$, since $\mathbf{w}_l$ is the l^{th} row of $\mathcal{X}$, $l = 1, \dots, k$, due to the assumption that the MDS code is systematic. Thus, the vector representing all the responses of the systematic nodes to the i^{th} sub-query can be written as follows,

$$\begin{bmatrix} r_{1,i} \\ r_{2,i} \\ \vdots \\ r_{k,i} \end{bmatrix} = \text{diag}(\mathcal{X}^T U_i P^T) + \begin{bmatrix} \mathbf{0}_{i-1 \times 1} \\ x_i^f \\ \mathbf{0}_{k-i \times 1} \end{bmatrix}, \quad (52)$$

where $\text{diag}(\cdot)$ is the diagonal of the corresponding matrix.

Denoting by $\mathbf{p}_j'$ the j^{th} column of P , the response of parity node l , $l = k+1, \dots, n$, can be written as

$$r_{l,i} = \mathbf{u}_{l-k,i}^T \mathbf{w}_l \quad (53)$$

$$= \mathbf{w}_l^T \mathbf{u}_{l-k,i} \quad (54)$$

$$= \mathbf{p}_{l-k}'^T \mathcal{X}^T \mathbf{u}_{l-k,i}, \quad (55)$$

where (55) follows from the fact that the coded data stored on parity node l can be written as $\mathbf{w}_l = \mathbf{p}_{l-k}'^T \mathcal{X}$. Thus, similarly to (52), we can write all the responses of the parity nodes in vector form as

$$\begin{bmatrix} r_{k+1,i} \\ r_{k+2,i} \\ \vdots \\ r_{n,i} \end{bmatrix} = \text{diag}(P^T \mathcal{X}^T U_i^T). \quad (56)$$

Next, we want to show that x_i^f can be decoded as follows,

$$x_i^f = \sum_{l=1}^k r_{l,i} - \sum_{l=k+1}^{k+b} r_{l,i}.$$

Indeed, we have

$$\sum_{l=1}^k r_{l,i} = \text{tr}(\mathcal{X}^T U_i P^T) + x_1^f \quad (57)$$

$$= \text{tr}(P^T \mathcal{X}^T U_i) + x_1^f \quad (58)$$

$$= \sum_{l=1}^b r_{l+k,i} + x_1^f \quad (59)$$

$$= \sum_{l=k+1}^{k+b} r_{l,i} + x_1^f, \quad (60)$$

where $\text{tr}(\cdot)$ is the trace operator, (57) follows from (52), (58) follows from the trace property, $\text{tr}(ABC) = \text{tr}(CAB)$, and (59) follows from (56).

This means that the responses of the systematic nodes and those of the parity nodes cancel out to leave the part required, i.e., x_i^f . Therefore, we showed that in the i^{th} , $i = 1, \dots, k$ sub-query, the user can decode x_i^f and by the completion of the k^{th} sub-query the user would have obtained the whole file X^f .

Razane Tajeddine (S'14) is currently a PhD student at the department of Mathematics and Systems Analysis at Aalto University. From 2014 to 2017, she was a PhD student at the Electrical and Computer Engineering Department at the Illinois Institute of Technology. She received her B.Sc. in Electrical Engineering from Notre Dame University-Louaize, Faculty of Engineering, Zouk Mosbeh, Lebanon, in 2012, and her M.Sc. in Electrical and Computer Engineering from the American University of Beirut, Beirut, Lebanon, in 2014. Her research interests lie in information theoretic security and privacy, wireless communications, and renewable energy.

Oliver W. Gnilke received the PhD degree from University College Dublin, Ireland in 2015. Earlier he received a Diplom in Mathematics from the University of Hamburg, Germany. Since 2015 he is a postdoctoral researcher at the Department of Mathematics and Systems Analysis at Aalto University, Finland. His research interests include combinatorial designs, coding theory, and applications of mathematics in security, privacy, and reliability.

Salim El Rouayheb (S'07-M'09) is currently an assistant professor in the ECE department at Rutgers University. From 2013 to 2017, he was an assistant professor at the ECE Department at the Illinois Institute of Technology. He received the Google Research Faculty Award in 2018 and the NSF CAREER award in 2016. He received the Diploma degree in electrical engineering from the Lebanese University, Faculty of Engineering, Roumieh, Lebanon, in 2002, and the M.S. degree from the American University of Beirut, Lebanon, in 2004. He received the Ph.D. degree in electrical engineering from Texas A&M University, College Station, in 2009. He was a postdoctoral research fellow at UC Berkeley (2010-2011) and a research scholar at Princeton University (2012-2013). His research interests lie in the area of information theoretic security and privacy of data in networks and distributed storage systems.