
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; Karpuk, David; Freij-Hollanti, Ragnar; Hollanti, Camilla
Private Information Retrieval from Coded Storage Systems with Colluding, Byzantine, and Unresponsive Servers

Published in:
IEEE Transactions on Information Theory

DOI:
[10.1109/TIT.2018.2890285](https://doi.org/10.1109/TIT.2018.2890285)

Published: 01/06/2019

Document Version
Peer reviewed version

Please cite the original version:
Tajeddine, R., Gnilke, O., Karpuk, D., Freij-Hollanti, R., & Hollanti, C. (2019). Private Information Retrieval from Coded Storage Systems with Colluding, Byzantine, and Unresponsive Servers. *IEEE Transactions on Information Theory*, 65(6), 3898-3906. [8598994]. <https://doi.org/10.1109/TIT.2018.2890285>

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.

This is the accepted version of the original article published by IEEE.

© 2019 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Private Information Retrieval from Coded Storage Systems with Colluding, Byzantine, and Unresponsive Servers

Razane Tajeddine*, Oliver W. Gnilke*, David Karpuk†, Ragnar Freij-Hollanti‡, Camilla Hollanti,*

* Department of Mathematics and Systems Analysis, School of Science, Aalto University, Espoo, Finland
Emails: {firstname.lastname}@aalto.fi

† Departamento de Matemáticas, Universidad de los Andes, Bogotá, Colombia
Email: da.karpuk@uniandes.edu.co

‡ Department of Electrical and Computer Engineering, Technical University of Munich, Munich, Germany
Email: ragnar.freij@tum.de

Abstract—The problem of Private Information Retrieval (PIR) from coded storage systems with colluding, Byzantine, and unresponsive servers is considered. An explicit scheme using an $[n, k]$ Reed-Solomon storage code is designed, protecting against t -collusion and handling up to b Byzantine and r unresponsive servers, when $n > k + t + 2b + r - 1$. This scheme achieves a PIR rate of $\frac{n-r-(k+2b+t-1)}{n-r}$. In the case where the capacity is known, namely when $k = 1$, it is asymptotically capacity-achieving as the number of files grows. Lastly, the scheme is adapted to symmetric PIR.

I. INTRODUCTION

Private Information Retrieval (PIR) is concerned with designing schemes for a user to retrieve a certain file from a storage system without revealing the identity of the file to the servers. This problem was introduced by Chor et al. in [2], where the database was viewed as an M -bit binary string $x = [x^1 \dots x^M] \in \{0, 1\}^M$ from which the user wants to retrieve one bit x^i while keeping the index i hidden from the server. In this work, we consider files encoded and stored on n servers, and assume that the user wants to retrieve some file x^i from the storage system, without revealing the index i . We assume a system with t -collusion, wherein any t of the n servers may collude in an attempt to deduce the desired file index. We additionally assume the presence of b Byzantine servers, which return erroneous information, and r unresponsive servers, which do not return any information at all.

The download rate, PIR rate, or simply *rate* of a PIR scheme in this model is measured as the ratio of the size of the downloaded file to the total amount of downloaded data; upload costs of the requests are ignored. As with the large majority of information-theoretic work on PIR, the rate will be our primary metric for judging the efficiency of a PIR scheme.

Initially, PIR constructions served to reduce the total download cost from a storage system with data replicated on multiple servers [3]–[8]. More recently, PIR schemes were constructed on coded data. The authors in [9] show that downloading one extra bit is enough to achieve privacy, if the number of servers is exponential in the number of files. In [10], the authors derive bounds on the tradeoff between storage cost and download cost for linearly coded data. The optimal upper bounds on PIR rate for maximum distance separable (MDS) coded data were derived in [11]. PIR schemes for MDS coded storage that achieve the asymptotic optimal download cost for specific numbers of colluding servers were presented in [12]. For the case of any number of colluding servers, the authors in [13] constructed a new family of PIR schemes on Reed-Solomon (RS) coded data achieving a lower download cost than the ones in [12]. PIR schemes on arbitrary linear storage codes were constructed in [14]. The notion of symmetric PIR, where the user is allowed to gain no information about the non-requested files, was studied in [15], [16].

In [8], it is shown that the asymptotic PIR capacity for replicated data, as the number of files $M \rightarrow \infty$, for a fixed number of colluding servers t , is $\frac{n-t}{n}$, where n is the number of nodes. When the data is coded using an $[n, k]$ MDS code, it was shown in [11] that the asymptotic capacity is $\frac{n-k}{n}$. Codes achieving this PIR rate were first presented in [12].

The problem of constructing PIR schemes on replicated data in which some servers can be Byzantine (malicious) was considered in [17]–[19]. The asymptotic capacity of PIR on replicated storage systems with t colluding servers and b Byzantine servers was found in [20] to be $\frac{n-(2b+t)}{n}$. In [21], the authors investigate the problem of providing symmetric PIR from a replicated system with colluding servers and adversaries in the system. A PIR scheme on coded data with colluding and either Byzantine or unresponsive servers was constructed in [22]. PIR from unsynchronized servers was studied in [23]. In this model, the files are stored on multiple servers and some servers might not be updated to the latest version. An adaptive PIR scheme is constructed for the user to privately retrieve the file they require. The setting

This paper was presented in part at the 2018 IEEE International Symposium on Information Theory (ISIT) [1]. 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 permissions@ieee.org.

of unsynchronized servers in [23] is similar to the Byzantine servers since in both cases some servers are giving erroneous responses. Compared to the present paper, the work in [23] is more restrictive and uses an adaptive scheme.

Main Contributions: In this paper, we construct a PIR scheme with servers storing data coded using a Reed-Solomon code, with up to t colluding servers, b Byzantine servers, and r unresponsive servers. We improve the PIR rate from [1] in the case where k does not divide $n - 2b - r - t + 1$. Provided that $n > k + t + 2b + r - 1$, our scheme achieves a PIR rate of

$$\frac{n - r - (k + t + 2b - 1)}{n - r}. \quad (1)$$

Specializing to the case of $r = b = 0$, we achieve a rate $\frac{n - (k + t - 1)}{n}$, which coincides with the rate achieved in [13]. Finally, the scheme is adapted to symmetric PIR, where the rate is optimal in the known cases in terms of achieving the bounds given in [15], [16], [21].

II. SYSTEM MODEL

A. Basic Definitions

We consider a storage system with n servers storing M files f^1, \dots, f^M , where each file is a matrix of size $L \times k$ over a finite field \mathbb{F} . We refer to L as the number of stripes in the file.

$$f^i = \begin{pmatrix} f_1^i \\ \vdots \\ f_L^i \end{pmatrix} = \begin{pmatrix} f_{1,0}^i & \cdots & f_{1,k-1}^i \\ \vdots & \ddots & \vdots \\ f_{L,0}^i & \cdots & f_{L,k-1}^i \end{pmatrix}. \quad (2)$$

Each file f^i is encoded using an $[n, k]$ linear code \mathcal{C} over \mathbb{F} in the following way. Let $G_{\mathcal{C}} \in \mathbb{F}^{k \times n}$ be a generator matrix of \mathcal{C} . Then the encoded file y^i is given by $y^i = f^i \cdot G_{\mathcal{C}}$, a matrix of size $L \times n$. The encoded files are distributed across the n servers by defining

$$\begin{pmatrix} f^1 \\ \vdots \\ f^M \end{pmatrix} \cdot G_{\mathcal{C}} = \begin{pmatrix} y^1 \\ \vdots \\ y^M \end{pmatrix} \\ = (y_1 \ \cdots \ y_n), \quad y_j \in \mathbb{F}^{LM \times 1}.$$

The vector y_j is then stored on server j , for $j = 1, \dots, n$.

The PIR problem for the above encoded storage system can be described as follows. A user wishes to download a file f^i without revealing the index i to any server. To do this, the user generates, according to some distribution, queries $q_j : \mathbb{F}^{LM} \rightarrow \mathbb{F}^S$ for some S , and sends q_j to server j . Considering schemes for which $S > 1$ is necessary so that the total file size is divisible by the amount of data that can be privately retrieved with a single query. The server responds with the value $r_j = q_j(y_j)$, and the desired file f^i can be computed as a deterministic function of the set r_j , $j = 1, \dots, n$.

In line with previous work on PIR, we refer to each of the S coordinate functions $q_j^{(s)}$ of the query $q_j = (q_j^{(1)}, \dots, q_j^{(S)})$ as a *subquery*, which takes place during a single *round* or *iteration* of the scheme. Thus, during the s^{th} round, the user sends the function $q_j^{(s)} : \mathbb{F}^{LM} \rightarrow \mathbb{F}$ to server j , who responds

with $r_j^{(s)} = q_j^{(s)}(y_j)$. During each round, we assume the presence of b *Byzantine* servers, who will instead respond with an arbitrary element of \mathbb{F} , as well as the presence of r *unresponsive* servers, whose responses are replaced with an erasure symbol $?$. The erasure symbol $?$ is absorbing with respect to addition, in the sense that $? + x = ?$ for all $x \in \mathbb{F}$. The identities of the servers which are Byzantine and unresponsive is allowed to change from round to round. During round s , we can write the *total response vector* as

$$r^{(s)} + \varepsilon = (r_1^{(s)}, \dots, r_n^{(s)}) + \varepsilon$$

where ε is a vector containing at most r erasure symbols and at most b non-zero elements of \mathbb{F} .

A PIR scheme as described above *protects against t -collusion*, or is *t -private*, if for every subset $T = \{j_1, \dots, j_T\} \subseteq \{1, \dots, n\}$ of servers of size $|T| = t$, we have

$$I(i; q_{j_1}, \dots, q_{j_T}) = 0$$

Our principal metric of efficiency of a PIR scheme is the *download rate*, or simply *rate*, defined as

$$R = \frac{Lk}{S(n - r)}$$

which is the ratio of the size of the desired file to the amount of total downloaded data. We assume that exactly r of the servers are unresponsive during each round of the scheme, hence the number of downloaded symbols in total in one round will be $n - r$.

B. Reed-Solomon Codes

From now on, we assume that $|\mathbb{F}| \geq n$. Let $\alpha_1, \dots, \alpha_n$ be n distinct elements of \mathbb{F} . Let $k \leq n$ and consider the space $\mathbb{F}[z]^{<k}$ of all single-variable polynomials of degree $< k$. We define an evaluation map

$$\text{eval} : \mathbb{F}[z] \rightarrow \mathbb{F}^n, \quad \text{eval}(f(z)) = (f(\alpha_1), \dots, f(\alpha_n)) \in \mathbb{F}^n.$$

The *Reed-Solomon code* $RS[n, k]$ is the image of $\mathbb{F}[z]^{<k}$ under this map:

$$RS[n, k] = \{\text{eval}(f(z)) : f(z) \in \mathbb{F}[z]^{<k}\}.$$

The code $RS[n, k]$ is MDS. If we write $f(z) = f_0 + f_1 z + \dots + f_{k-1} z^{k-1}$, then

$$\text{eval}(f(z)) = (f_0 \ f_1 \ \cdots \ f_{k-1}) \cdot G_{RS[n,k]}, \\ \text{where } G_{RS[n,k]} = (\alpha_j^i)_{\substack{0 \leq i \leq k-1 \\ 1 \leq j \leq n}}$$

hence the Vandermonde matrix $G_{RS[n,k]}$ is a generator matrix of $RS[n, k]$. We denote the inverse of eval on $RS[n, k]$ by dec , which performs polynomial interpolation to recover $f(z)$ from the vector of evaluations:

$$\text{dec} : RS[n, k] \rightarrow \mathbb{F}[z]^{<k}$$

The maps eval and dec are vector space isomorphisms between $RS[n, k]$ and $\mathbb{F}[z]^{<k}$.

NOMENCLATURE

n	Number of servers
\mathcal{C}	$[n, k]$ storage code
\mathcal{D}	$[n, t]$ query code
M	Number of files
t	Number of colluding servers
b	Number of Byzantine servers
r	Number of unresponsive server
S	Number of rounds
L	Number of stripes / rows in file matrix
ρ	Number of symbols retrieved per round

$$\begin{array}{c}
 f^i \\
 \hline
 \overbrace{f_{1,0}^i + f_{1,1}^i z + f_{1,2}^i z^2 + f_{1,3}^i z^3} \\
 \hline
 \begin{array}{|c|c|}
 \hline
 f_{1,0}^i + f_{1,1}^i z & z^2 (f_{1,2}^i + f_{1,3}^i z) \\
 \hline
 \end{array} \\
 \hline
 h^{i,(2)} \quad \quad z^2 \cdot h^{i,(1)}
 \end{array}$$

Fig. 1: Divisions for Example 1. The file f^i is split into two parts $h^{i,(1)}$ and $h^{i,(2)}$, which are retrieved in rounds 1 and 2, respectively.

C. Storage Systems from Reed-Solomon Codes

We will consider storage codes \mathcal{C} which are Reed-Solomon codes: $\mathcal{C} = RS[n, k]$. Given a file f^i , the k information symbols of row $f^{i,\ell}$ are encoded as coefficients of a polynomial

$$f_\ell^i(z) = f_{\ell,0}^i + f_{\ell,1}^i z + \dots + f_{\ell,k-1}^i z^{k-1} \quad (3)$$

of degree $< k$. This polynomial is evaluated at n different points $\alpha_1, \dots, \alpha_n \in \mathbb{F}$, and the evaluations of these polynomials at α_j are stored on server j . Therefore, the matrix $y^i = f^i \cdot G_{RS[n,k]}$ is of the form

$$y^i = \begin{pmatrix} f_1^i(\alpha_1) & \dots & f_1^i(\alpha_n) \\ \vdots & \ddots & \vdots \\ f_L^i(\alpha_1) & \dots & f_L^i(\alpha_n) \end{pmatrix}$$

and the contents y_j of server j are the length LM column vector

$$y_j = (f_1^1(\alpha_j), \dots, f_L^1(\alpha_j), \dots, f_1^M(\alpha_j), \dots, f_L^M(\alpha_j))^T.$$

We remark that all schemes presented in this work can be used for storage systems using generalized Reed-Solomon codes as well, but we restrict the description to RS codes for simplicity.

III. A PIR SCHEME FOR COLLUDING BYZANTINE SERVERS

A. A Simple Example

We start with an example which shows how the scheme works. In the following subsection we will describe the scheme in its full generality, but the present example suffices to convey the basic ideas.

Example 1. Suppose we have a system with parameters $n = 9$, $k = 4$, $t = b = r = 1$. Our storage code \mathcal{C} is an $[9, 4]$ Reed-Solomon code. Let α be the evaluation vector of

the Reed-Solomon code. To correct the errors and erasures, we require that the responses live in a code with a minimum distance of at least $2b + r + 1 = 4$, and thus can tolerate b errors and r erasures. Among the $n = 9$ total dimensions of a single response, $k + t - 1$ are spanned by randomness required to preserve privacy and introduced by the storage code. Additionally, $2b + r$ dimensions are required to perform erasure and error correction. This forces us to use no more than the remaining $\rho := n - (k + t + 2b + r - 1)$ dimensions to retrieve information. Our files will have $L = 1$ row and our scheme will require $S = 2$ rounds.

Suppose that we want to download f^i . Based on the above discussion, in each round we will recover exactly $\rho = 2$ symbols of file f^i . We see that in this example f^i only has one row, hence

$$f^i(z) = f_{1,0}^i + f_{1,1}^i z + f_{1,2}^i z^2 + f_{1,3}^i z^3 \quad (4)$$

$$= h^{i,(2)}(z) + h^{i,(1)}(z)z^2 \quad (5)$$

where $h^{i,(1)}(z) = f_{1,2}^i + f_{1,3}^i z$ and $h^{i,(2)}(z) = f_{1,0}^i + f_{1,1}^i z$, as depicted in Figure 1. In round s of the scheme we will retrieve the coefficients of $h^{i,(s)}(z)$, which comprise ρ symbols of the desired file, see Figure 2.

In round 1 we choose a random coefficient/constant polynomial $d^{m,(1)}$ for every file f^m , $1 \leq m \leq M$. For the requested file f^i we add the monomial z^2 to $d^{i,(1)}$. In other words, define

$$q^{(1)}(z) = (q^{1,(1)}(z), \dots, q^{M,(1)}(z)),$$

$$\text{where } q^{m,(1)}(z) = \begin{cases} d^{m,(1)} & \text{if } m \neq i \\ d^{m,(1)} + z^2 & \text{if } m = i \end{cases}$$

and the query q_j sent to the j^{th} server is given by

$$q_j^{(1)} = q^{(1)}(\alpha_j) = (q^{1,(1)}(\alpha_j), \dots, q^{M,(1)}(\alpha_j))$$

$$= \begin{cases} d^{m,(1)} & \text{if } m \neq i \\ d^{m,(1)} + \alpha_j^2 & \text{if } m = i. \end{cases}$$

Before the addition of errors and erasures induced by the Byzantine and unresponsive servers, the response from server j is

$$r_j^{(1)} = \langle q_j^{(1)}, y_j \rangle = \sum_{m=1}^M q^{m,(1)}(\alpha_j) f^m(\alpha_j) \quad (6)$$

$$= \sum_{m=1}^M d^{m,(1)} f^m(\alpha_j) + \alpha_j^2 f^i(\alpha_j). \quad (7)$$

The vector $r^{(1)} = (r_1^{(1)}, \dots, r_n^{(1)})$ consisting of all of the responses from the n servers is therefore

$$r^{(1)} = \text{eval}(r^{(1)}(z)),$$

$$\text{where } r^{(1)}(z) = \sum_{m=1}^M d^{m,(1)} f^m(z) + z^2 f^i(z).$$

Since $\deg(r^{(1)}(z)) \leq 5$, we see that $r^{(1)}$ is an element of a Reed-Solomon code with parameters $[9, 6]$ with minimum distance 4, and can thus tolerate the one error and one erasure introduced by the Byzantine and unresponsive servers.




Fig. 2: Retrieval scheme construction for Example 1. The file f^i , consisting of 4 coefficients, is retrieved over 2 rounds.

Expanding the polynomial $r^{(1)}(z)$, we have

$$\begin{aligned} r^{(1)}(z) &= \sum_{m=1}^M d^{m,(1)} f^m(z) + z^2 f^i(z) \quad (8) \\ &= \underbrace{\sum_{m=1}^M d^{m,(1)} f^m(z) + h^{i,(2)}(z)z^2 + h^{i,(1)}(z)z^4}_{=:g(z), \deg \leq 3}. \quad (9) \end{aligned}$$

Since $\deg(g(z)) \leq 3$, we have that $\text{eval}(g(z)) \in \mathcal{C}$. Therefore, we can recover the coefficients of $h^{i,1}(z)$, namely $f_{1,2}^i$ and $f_{1,3}^i$, from the response vector $r^{(1)}$.

In round 2, we again sample M constant polynomials $d^{m,(2)}$ and set

$$\begin{aligned} q^{(2)}(z) &= (q^{1,(2)}(z), \dots, q^{M,(2)}(z)), \\ \text{where } q^{m,(2)}(z) &= \begin{cases} d^{m,(2)} & \text{if } m \neq i \\ d^{m,(2)} + z^4 & \text{if } m = i. \end{cases} \end{aligned}$$

A similar calculation as in the first round shows that the total response vector $r^{(2)}$ in the second round, before the addition of errors and erasures, is $r^{(2)} = \text{eval}(r^{(2)}(z))$, where

$$r^{(2)}(z) = \underbrace{\sum_{m=1}^M d^{m,(1)} f^m(z) + h^{i,(2)}(z)z^4}_{\deg \leq 3} + \underbrace{h^{i,(1)}(z)z^6}_{\text{known from round 1}}. \quad (10)$$

The user receives $\text{eval}(r^{(2)}(z))$ plus a vector consisting of errors and erasures. From round 1, the user knows $h^{i,(1)}(z)$, and before correcting for errors and erasures can subtract $\text{eval}(h^{i,(1)}(z)z^6)$ from what they receive. What is left is a codeword of a Reed-Solomon code which can correct the errors and erasures as in round 1. Again as in round 1, the user reads off the coefficients of $h^{i,(2)}(z)$, namely $f_{1,0}^i$ and $f_{1,1}^i$.

The user now has all of the coefficients of f^i and can reconstruct the entire file. The rate of this scheme is clearly seen to be

$$\frac{n-r-(k+t+2b-1)}{n-r} = \frac{\rho}{n-r} = \frac{2}{8} = \frac{1}{4}.$$

B. The General Scheme

Let $\rho := n - (k + t + 2b + r - 1)$. We choose the parameters L and S to be minimal such that $Lk = S\rho$, that is,

$$L = \frac{\text{lcm}(\rho, k)}{k}, \quad S = \frac{\text{lcm}(\rho, k)}{\rho}.$$

Our scheme requires that

$$n > k + t + 2b + r - 1, \quad (11)$$

and the rate of our scheme will be given by

$$\frac{\rho}{n-r} = \frac{n-k-t-2b-r+1}{n-r}, \quad (12)$$

in other words, ρ is the number of symbols retrieved during each round of the protocol.

To correct the b errors and r erasures from the Byzantine and unresponsive servers, our scheme will be constructed so that the response vector $r^{(s)}$ during round s is an element of an affine shift of a Reed-Solomon code with minimum distance $d = 2b + r + 1$. The vector by which the Reed-Solomon code is shifted is known to the user during each round, allowing one to correct for the b errors and r erasures.

For every round s , we choose i.i.d. uniform codewords from the query code $\mathcal{D} = RS[n, t]$, as the evaluation of polynomials $d_\ell^{m,(s)}(z)$, for every row ℓ of every file m . For the rows of file i we add the evaluation of another polynomial to the query, such that all ρ coefficients of a polynomial $h^{i,(s)}(z)$ (defined formally in eq. (24)) can be retrieved in round s , as represented in Figure 3. Explicitly, we let

$$e_x(z) = \begin{cases} z^x & \text{if } x \geq t \\ 0 & \text{otherwise.} \end{cases}$$

and add the evaluation of

$$e_{s\rho - \ell k + k + t - 1}(z) \quad (13)$$

to the entries of the query corresponding to the ℓ^{th} row of file i in round s . Our query polynomials are then defined as

$$q_\ell^{m,(s)}(z) := \begin{cases} d_\ell^{m,(s)}(z) + e_{s\rho + (1-\ell)k + t - 1}(z) & \text{if } m = i \\ d_\ell^{m,(s)}(z) & \text{if } m \neq i. \end{cases} \quad (14)$$




Fig. 3: A PIR scheme from an $[n, k]$ storage code with t colluding, b Byzantine, and r unresponsive servers for $\rho < k$.

Note that $\deg(q_\ell^{m,(s)}(z)) \leq t - 1$ for $m \neq i$, and hence $\deg(f_\ell^m(z)q_\ell^{m,(s)}(z)) \leq k + t - 2$. Furthermore, in round 1, $\deg(q_\ell^{i,(1)}(z)) \leq n - 2b - r - k$. During round s , the user sends the vector q_j to server j , where

$$q_j^{(s)} = (q_1^{1,(s)}(\alpha_j), \dots, q_L^{1,(s)}(\alpha_j), \dots, q_1^{M,(s)}(\alpha_j), \dots, q_L^{M,(s)}(\alpha_j)) \in \mathbb{F}^{LM},$$

thus the user sends to server j all the evaluations of the query polynomials $q_\ell^{m,(s)}(z)$ at $z = \alpha_j$.

In round s , before the addition of the errors and erasures due to the Byzantine and unresponsive servers, the response $r_j^{(s)}$ from server j is given by

$$r_j^{(s)} = \langle q_j^{(s)}, y_j \rangle \quad (15)$$

$$= \sum_{m=1}^M \sum_{\ell=1}^L q_\ell^{m,(s)}(\alpha_j) f_\ell^m(\alpha_j) \quad (16)$$

$$= \sum_{m=1}^M \sum_{\ell=1}^L d_\ell^{m,(s)}(\alpha_j) f_\ell^m(\alpha_j) + \sum_{\ell=1}^L e_{s\rho - \ell k + k + t - 1}(\alpha_j) f_\ell^i(\alpha_j). \quad (17)$$

Note that $s\rho - \ell k + k + t - 1 \geq t \Leftrightarrow \ell \leq \lceil s\rho/k \rceil$, therefore the above is equal to

$$= \sum_{m=1}^M \sum_{\ell=1}^L d_\ell^{m,(s)}(\alpha_j) f_\ell^m(\alpha_j) + \sum_{\ell=1}^{\lceil s\rho/k \rceil} \alpha_j^{s\rho - \ell k + k + t - 1} f_\ell^i(\alpha_j). \quad (18)$$

The second summand in the above illustrates that during round s , the response involves rows $1, \dots, \lceil s\rho/k \rceil$. The total

response vector (before the addition of errors and erasures) during round s is therefore $r^{(s)} = (r_1^{(s)}, \dots, r_n^{(s)}) = \text{eval}(r^{(s)}(z))$, where

$$r^{(s)}(z) = \underbrace{\sum_{m=1}^M \sum_{\ell=1}^L d_\ell^{m,(s)}(z) f_\ell^m(z)}_{=:g^{(s)}(z), \deg(g^{(s)}) < k+t-1} + \sum_{\ell=1}^{\lceil s\rho/k \rceil} z^{s\rho - \ell k + k + t - 1} f_\ell^i(z). \quad (19)$$

We will refer to $r^{(s)}(z)$ as the *response polynomial* during round s .

To best illustrate why we can recover ρ information symbols from the above response, we first consider what happens in round $s = 1$. The response polynomial during round 1 is

$$r^{(1)}(z) = g^{(1)}(z) + \sum_{\ell=1}^{\lceil \rho/k \rceil} z^{\rho - \ell k + k + t - 1} f_\ell^i(z) \quad (20)$$

$$= g^{(1)}(z) + z^{k+t-1} \sum_{\ell=1}^{\lceil \rho/k \rceil - 1} z^{\rho - \ell k} f_\ell^i(z) + z^{\rho - \lceil \rho/k \rceil k + k + t - 1} f_{\lceil \rho/k \rceil}^i(z). \quad (21)$$

Expressing the last summand as two terms, one of which has degree $< k + t - 1$, and one which has degree $\geq k + t - 1$ yields

$$\begin{aligned}
&= g^{(1)}(z) + z^{k+t-1} \sum_{\ell=1}^{\lceil \rho/k \rceil - 1} z^{\rho - \ell k} f_{\ell}^i(z) \\
&+ \underbrace{\sum_{\kappa=0}^{\lceil \rho/k \rceil k - \rho - 1} f_{\lceil \rho/k \rceil, \kappa}^i z^{\rho - \lceil \rho/k \rceil k + k + t - 1 + \kappa}}_{=: \gamma(z), \deg(\gamma(z)) < k + t - 1} \\
&+ z^{k+t-1} \sum_{\kappa=\lceil \rho/k \rceil k - \rho}^{k-1} f_{\lceil \rho/k \rceil, \kappa}^i z^{\rho - \lceil \rho/k \rceil k + \kappa} \\
&= \underbrace{g^{(1)}(z) + \gamma(z)}_{\deg < k + t - 1} + z^{k+t-1} h^{i,(1)}(z) \quad (22)
\end{aligned}$$

where

$$\begin{aligned}
h^{i,(1)}(z) &= \sum_{\kappa=\lceil \rho/k \rceil k - \rho}^{k-1} f_{\lceil \rho/k \rceil, \kappa}^i z^{\rho - \lceil \rho/k \rceil k + \kappa} \\
&+ \sum_{\ell=1}^{\lceil \rho/k \rceil - 1} z^{\rho - \ell k} f_{\ell}^i(z) \quad (23)
\end{aligned}$$

From (20) we see that $\deg(r^{(1)}(z)) < \rho + k + t - 1 = n - (2b + r)$, therefore $\text{eval}(r^{(1)}(z))$ is a codeword in a Reed-Solomon code with minimum distance $2b + r + 1$. Hence the user can correct up to b errors and r erasures introduced by the Byzantine and unresponsive servers.

After correcting errors and erasures, the user obtains from the above expression the k coefficients of the polynomials $f_{\ell}^i(z)$ for $\ell = 1, \dots, \lceil \rho/k \rceil - 1$, and when $\ell = \lceil \rho/k \rceil$, we obtain from the polynomial $f_{\lceil \rho/k \rceil}^i(z)$ the coefficients $f_{\lceil \rho/k \rceil, \kappa}^i$ for $\kappa = \lceil \rho/k \rceil k - \rho, \dots, k - 1$. Thus ρ information symbols are downloaded in the first round.

Now consider round s . We define polynomials $h^{i,(s)}(z)$ of degree $< \rho$ by the following recursive formula. The polynomial $h^{i,(1)}(z)$ is defined as in (23), and for $s > 1$ we define them recursively via the formula

$$r^{(s)}(z) = g^{(s)}(z) + z^{k+t-1} \sum_{\sigma=1}^s z^{\rho(s-\sigma)} h^{i,(\sigma)}(z). \quad (24)$$

where $\deg(g^{(s)}(z)) < k + t - 1$. Note that the $h^{i,(s)}(z)$ depend only on the file f^i and not the randomness present in the queries. Picking off the polynomials whose coefficients we know from previous rounds, we can write

$$\begin{aligned}
r^{(s)}(z) &= g^{(s)}(z) + z^{k+t-1} h^{i,(s)}(z) \\
&+ z^{k+t-1} \sum_{\sigma=1}^{s-1} \underbrace{z^{\rho(s-\sigma)} h^{i,(\sigma)}(z)}_{\text{known from rounds } 1, \dots, s-1}.
\end{aligned}$$

The user receives $\text{eval}(r^{(s)}(z))$ plus errors and erasures. First subtracting off the evaluation of the known summands in the above, we arrive at a codeword in the same Reed-Solomon code as in round 1, which allows us to correct errors and

erasures. The user then recovers the ρ coefficients of the polynomial $h^{i,(s)}(z)$.

To prove that this suffices to download the whole file, consider the response polynomial $r^{(S)}(z)$ in round S . Equating the expressions (19) and (24) for $r^{(S)}(z)$, we arrive at the equality

$$\begin{aligned}
r^{(S)}(z) &= g^{(S)}(z) + z^{k+t-1} \sum_{\ell=1}^L z^{(L-\ell)k} f_{\ell}^i(z) \\
&= g^{(S)}(z) + z^{k+t-1} \sum_{s=1}^S z^{(S-s)\rho} h^{i,(s)}(z)
\end{aligned}$$

or equivalently,

$$\sum_{\ell=1}^L z^{(L-\ell)k} f_{\ell}^i(z) = \sum_{s=1}^S z^{(S-s)\rho} h^{i,(s)}(z)$$

which shows that the $h^{i,(s)}(z)$ determine the polynomials $f_{\ell}^i(z)$ completely, as depicted in Fig. 4.

The achieved rate by this scheme is easily seen to be

$$R = \frac{Lk}{S(n-r)} = \frac{\rho}{n-r} = \frac{n-r-(k+t+2b-1)}{n-r}. \quad (25)$$

The scheme is t -private since the retrieval code D is MDS of dimension t , the proof is identical to the proof of privacy in [13].

In order to illustrate some of the subtleties that appear in the scheme construction for more general sets of parameters, we present the following example. This serves to illustrate how the scheme works when we require $L > 1$ rows, $S > 1$ rounds, and $t > 1$ colluding servers. Moreover, note that for the given parameters we have $\rho > k$, which forces us to retrieve coefficients from more than one polynomial $f^{i,\ell}(z)$ at a time.

Example 2. Suppose we have the parameters $n = 14, k = 4, t = 2, r = 1, b = 1$. We construct a scheme that achieves a PIR rate of $6/13$. We can retrieve $\rho = 6$ symbols per round, and require $L = 3$ rows per file and $S = 2$ rounds of the scheme. Thus, we decompose the file in two ways as

$$\sum_{\ell=1}^3 z^{4(3-\ell)} f^{i,\ell}(z) = \sum_{s=1}^2 z^{6(2-s)} h^{i,s}(z) \quad (26)$$

as shown in Figure 5.

We pick the retrieval code $\mathcal{D} = RS[14, 2]$. In round $s = 1$, we sample i.i.d. uniform codewords $\text{eval}(d_{\ell}^{m,(1)}(z)) \in \mathcal{D}$. The query polynomials $q_{\ell}^{i,(1)}(z)$ are given by

$$q_{\ell}^{i,(1)}(z) = \begin{cases} d_1^{i,(1)}(z) + z^7 & \text{if } \ell = 1 \\ d_2^{i,(1)}(z) + z^3 & \text{if } \ell = 2 \\ d_3^{i,(1)}(z) & \text{if } \ell = 3. \end{cases}$$

The response polynomial $r^{(1)}(z)$ is of the form

$$\begin{aligned}
r^{(1)}(z) &= g^{(1)}(z) + f_{2,2}^i z^5 + f_{2,3}^i z^6 + f_1^i(z) z^7 \\
&= g^{(1)}(z) + h^{i,(1)}(z) z^5
\end{aligned}$$

where $\deg(g^{(1)}(z)) < k + t - 1 = 5$, which allows us to download the four coefficients of $f_1^i(z)$ plus the two additional




Fig. 4: The coefficients of the polynomials $f^{i,\ell}$ are lined up and partitioned in two different ways, into L blocks of size k each (above) and into S blocks of size ρ each (below).

coefficients of $f_2^i(z)$. In round two, the response polynomial $r^{(2)}(z)$ is of the form

$$\begin{aligned} r^{(2)}(z) &= g^{(2)}(z) + f_3^i(z)z^5 + f_{2,0}^i z^9 + f_{2,1}^i z^{10} + h^{i,(1)} z^{11} \\ &= g^{(2)}(z) + h^{i,(2)}(z)z^5 + \underbrace{h^{i,(1)} z^{11}}_{\text{known from first round}} \end{aligned}$$

from which we obtain the coefficients of $h^{i,(2)}(z)$, as shown in Figure 6.

IV. COMPARISON WITH PREVIOUS PIR SCHEMES AND A SYMMETRIC VARIANT

A. Comparison with Other Work on PIR with Byzantine and Unresponsive Servers and Coded Data

Recently, Zhang and Ge [22] constructed a PIR scheme for coded data and colluding servers, which is adaptable for unresponsive and Byzantine servers (but not for both simultaneously). In this section we briefly compare the rates obtained in this paper with those of [22] in the asymptotic regime as $M \rightarrow \infty$. The scheme of [22] only achieves positive rates assuming certain inequalities in the basic system parameters are satisfied, namely the obvious inequalities which guarantee that the expressions below in (27) and (28) are positive. To compare the two schemes at their best, we grant this assumption.

When $b = 0$ and $r > 0$, the asymptotic rate as $M \rightarrow \infty$ from [22] can be expressed as

$$\bar{R} = \frac{n}{n-r} \left(\frac{\binom{n-r}{k} + \binom{n-t}{k} - \binom{n}{k}}{\binom{n}{k}} \right). \quad (27)$$

An elementary calculation shows that $\bar{R} < \frac{n-r-(k+t-1)}{n-r}$, the rate obtained for the scheme described in the previous sections. In the case where $b > 0$ and $r = 0$, the asymptotic rate obtained in [22] is

$$\bar{R} = \frac{2 \left(\binom{n-b}{k} - \binom{n}{k} \right) + \binom{n-t}{k}}{\binom{n}{k}} \quad (28)$$

which, again by a simple argument, is less than $\frac{n-(k+t+2b-1)}{n}$, the rate obtained by the proposed scheme in this case.

Lastly, we remark that the rates obtained in [22] decrease with an increasing number of files, while the rates we obtain are constant in the number of files. As noted in [22], the rates therein outperform those of [13] for a small number of files. We can see from Figure 7 that the same holds here for these example parameters.

B. A Symmetric Variant

A PIR scheme is *symmetric* if the user, while retrieving the requested file f^i , gains no information about any of the other files $f^{i'}$ for $i \neq i'$. To construct a symmetric variant of our scheme, we assume the servers have access to a joint source of randomness. Each round, the joint source of randomness outputs a uniform random codeword $\pi^{(s)} = \text{eval}(\pi^{(s)}(z))$, where $\pi^{(s)}(z) \in \mathbb{F}[z]^{<k+t-1}$.

The scheme proceeds exactly as before, though all servers compute $r_j^{(s)} = \langle q_j^{(s)}, y_j \rangle + \pi_j^{(s)}$, which the responsive, non-Byzantine servers transmit back to the user. As before, the user receives an erasure symbol from the unresponsive servers, and an arbitrary element of \mathbb{F} from the Byzantine servers. Since $\deg(\pi^{(s)}(z)) < k+t-1$, it is absorbed into the ‘randomness’ term $g^{(s)}(z)$ and therefore does not affect how the user recovers the ρ information symbols. Since $\pi^{(s)}(z)$ is uniformly chosen, there is clearly no information leaked about any files $f^{i'}$ for $i' \neq i$.

We note that this is the same amount of randomness needed to symmetrize the scheme of [15], which deals with the case of $b = r = 0$, as well as in the scheme of [21], which deals with the case $k = 1$.

C. Conjectures

We venture the following conjectures regarding the asymptotic and symmetric capacities of this PIR problem.

Conjecture 1. *The asymptotic capacity (as $M \rightarrow \infty$) of Private Information Retrieval for an $[n, k]$ MDS storage code with t -collusion, b Byzantine servers, and r unresponsive servers is $\frac{n-r-(k+t+2b-1)}{n-r}$. That is, the current scheme is asymptotically capacity-achieving.*

Conjecture 2. *The capacity of Symmetric Private Information Retrieval for an $[n, k]$ MDS storage code with t -collusion, b Byzantine servers, and r unresponsive servers is $\frac{n-r-(k+t+2b-1)}{n-r}$. That is, the symmetrization of the current scheme is capacity-achieving.*

Furthermore, to guarantee symmetry, the minimum entropy per round required for the shared randomness amongst the servers is $k+t-1$ (in q -ary units).

V. CONCLUSION

A PIR scheme was presented in this paper which can simultaneously handle coded data and colluding, unresponsive servers, and Byzantine servers. In the current work, the response from the servers is an element of a linear code which allows the user to correct for the erasures and errors produced by the unresponsive and Byzantine servers.

$$\begin{array}{c}
\overbrace{f_3^i} \qquad \qquad \qquad \overbrace{z^4 \cdot f_2^i} \qquad \qquad \qquad \overbrace{z^8 \cdot f_1^i} \\
\hline
\begin{array}{|l|l|l|}
\hline
f_{3,0}^i + f_{3,1}^i z + f_{3,2}^i z^2 + f_{3,3}^i z^3 & f_{2,0}^i z^4 + f_{2,1}^i z^5 + f_{2,2}^i z^6 + f_{2,3}^i z^7 & f_{1,0}^i z^8 + f_{1,1}^i z^9 + f_{1,2}^i z^{10} + f_{1,3}^i z^{11} \\
\hline
f_{3,0}^i + f_{3,1}^i z + f_{3,2}^i z^2 + f_{3,3}^i z^3 + f_{2,0}^i z^4 + f_{2,1}^i z^5 & z^6 (f_{2,2}^i + f_{2,3}^i z + f_{1,0}^i z^2 + f_{1,1}^i z^3 + f_{1,2}^i z^4 + f_{1,3}^i z^5) & \\
\hline
\end{array} \\
\hline
h^{i,(2)} \qquad \qquad \qquad z^6 \cdot h^{i,(1)}
\end{array}$$

Fig. 5: Coefficients lined up and partitioned into $L = 3$ blocks of size $k = 4$ each (above) and into $S = 2$ blocks of size $\rho = 6$ each (below) in Example 2.




Fig. 6: Retrieval scheme for Example 2. Note how in the last round, the lined up polynomials correspond to Fig. 5, with the addition of a shift by z^5 .




Fig. 7: The PIR rate versus number of files M when $n = 12$, $k = 2$, and $t = 3$ following the scheme in [22] and the scheme in this paper.

The scheme has rate $\frac{n-r-(k+2b+t-1)}{n-r}$, which is equal to the asymptotic capacity (as the number of files goes to infinity) in all cases where the capacity is known. The scheme compares favorably to previous schemes which account for unresponsive and Byzantine servers. Additionally, the scheme is easily symmetrizable.

ACKNOWLEDGMENTS

This work is supported in part by the Academy of Finland, under grants #276031, #282938, and #303819 to C. Hollanti, and by the Technical University of Munich – Institute for

Advanced Study, funded by the German Excellence Initiative and the EU 7th Framework Programme under grant agreement #291763, via a *Hans Fischer Fellowship* held by C. Hollanti.

R. Freij-Hollanti was supported by the German Research Foundation (Deutsche Forschungsgemeinschaft, DFG) under Grant WA3907/1-1.

O. W. Gnilke and R. Tajeddine were visiting the group of Professor Antonia Wachter-Zeh at the Technical University of Munich while this work was carried out, and are thankful for the hospitality of the LNT Chair and the COD Group.

O. W. Gnilke is partially supported by the Finnish Cultural Foundation.

REFERENCES

- [1] R. Tajeddine, O. W. Gnilke, D. Karpuk, R. Freij-Hollanti, and C. Hollanti, “Robust private information retrieval from coded systems with byzantine and colluding servers,” in *2018 IEEE International Symposium on Information Theory*, IEEE, 2018.
- [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] A. Beimel and Y. Ishai, “Information-theoretic private information retrieval: A unified construction,” in *Automata, Languages and Programming*, pp. 912–926, Springer, 2001.
- [4] 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.
- [5] 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.
- [6] S. Yekhanin, “Private information retrieval,” *Communications of the ACM*, vol. 53, no. 4, pp. 68–73, 2010.
- [7] H. Sun and S. A. Jafar, “The capacity of private information retrieval,” *IEEE Transactions on Information Theory*, pp. 4075 – 4088, 2017.