

Private Read Update Write (PRUW) With Heterogeneous Databases

Sajani Vithana Sennur Ulukus
 Department of Electrical and Computer Engineering
 University of Maryland, College Park, MD 20742
spallego@umd.edu *ulukus@umd.edu*

Abstract—We investigate the problem of private read update write (PRUW) with heterogeneous storage constrained databases in federated submodel learning (FSL). In FSL a machine learning model is divided into multiple submodels based on different types of data used to train it. A given user downloads, updates and uploads the updates back to a single submodel of interest, based on the type of user’s local data. With PRUW, the process of reading (downloading) and writing (uploading) is carried out such that information-theoretic privacy of the updating submodel index and the values of updates is guaranteed. We consider the practical scenario where the submodels are stored in databases with arbitrary (heterogeneous) storage constraints, and provide a PRUW scheme with a storage mechanism that utilizes submodel partitioning and encoding to minimize the communication cost.

I. INTRODUCTION

In private read update write (PRUW) [1]–[8], the users read from and write to selected sections of a data storage without revealing what is read or written. We use PRUW in federated submodel learning (FSL) [1], [2], [9]–[12] to guarantee user privacy in the learning process. FSL is a method developed to reduce the communication cost in federated learning (FL), in which the central FL model is divided into multiple submodels based on different types of data used to train it. In FSL, a given user only downloads the submodel relevant to the user’s local data, updates it, and writes (uploads) the updates back to the submodel. The updating submodel index and the uploaded values of updates leak information about the type of data the user has. With PRUW, the read-write process in FSL is carried out while guaranteeing information-theoretic privacy of the updating submodel index and the values of updates. In PRUW, the submodels are stored in multiple non-colluding databases, similar to private information retrieval (PIR) [13]–[44].

In this paper, we consider the practical scenario where the multiple non-colluding databases have arbitrary storage constraints. PRUW with homogeneous storage constraints is studied in [3]. PIR with homogeneous and heterogeneous storage constrained databases is studied in [45]–[51], which introduce optimum partitioning strategies to fit the messages into the storage constrained databases in such a way that the download cost is minimized. In this work, we extend the ideas used in PIR with storage constrained databases to PRUW. The two main methods used in the PIR literature to handle storage constrained databases are: 1) coded storage, where the




Fig. 1. System model.

model parameters are encoded with a specific (K, R) MDS code, and 2) divided storage, where the model is partitioned, and different partitions are replicated in different subsets of databases. We first observe that there is an advantage of using (K, R) MDS coded storage with higher values of K in PRUW with heterogeneous databases, and propose a PRUW scheme and a storage mechanism for FSL that is based on submodel partitioning and MDS coding to efficiently utilize the available space in all databases to minimize the total communications.

II. PROBLEM FORMULATION

We consider a FSL setting with M independent submodels, each containing L parameters, taking values from a finite field \mathbb{F}_q , stored in a system of N non-colluding databases. Database n has a storage capacity of $\mu(n)ML$ symbols, where $\mu(n) \leq 1$ for $n \in \{1, \dots, N\}$; see Fig. 1. That is,

$$H(S_n^{[t]}) \leq \mu(n)ML, \quad n \in \{1, \dots, N\}, \quad t \in \mathbb{N}, \quad (1)$$

where $S_n^{[t]}$ is the content stored in database n at time t . At a given time t , a user downloads a required submodel without revealing its index to any of the databases. This is the reading phase of the PRUW process. The user then updates the submodel locally, and uploads the updates back to the same submodel in all databases without revealing the values of updates or the submodel index. This is the writing phase.

Privacy of the submodel index: No information on the updating submodel index at time t , $\theta^{[t]}$, is allowed to leak

to any of the databases, i.e., for each n , $n \in \{1, \dots, N\}$,

$$I(\theta^{[t]}; Q_n^{[t]}, U_n^{[t]} | Q_n^{[1:t-1]}, S_n^{[0:t-1]}, U_n^{[1:t-1]}) = 0, \quad (2)$$

where Q_n and U_n are the queries and updates sent by the user to database n at time instances denoted in square brackets.

Privacy of the values of updates: No information on the values of updates is allowed to leak to any of the databases,

$$I(\Delta_\theta^{[t]}; U_n^{[t]} | Q_n^{[1:t]}, S_n^{[0:t-1]}, U_n^{[1:t-1]}) = 0, \quad n \in \{1, \dots, N\} \quad (3)$$

where $\Delta_\theta^{[t]}$ is the update generated by the user at time t .

Security of submodels: No information on the values of submodels is allowed to leak to any of the databases, i.e.,

$$I(W_{1:M}^{[t]}; S_n^{[t]}) = 0, \quad n \in \{1, \dots, N\} \quad (4)$$

where $W_k^{[t]}$ is the k th submodel at time t .

Correctness in the reading phase: The user should correctly download the required submodel in the reading phase, i.e., $H(W_\theta^{[t-1]} | Q_{1:N}^{[t]}, A_{1:N}^{[t]}) = 0$ where $A_n^{[t]}$ is the answer received from database n at time t .

Correctness in the writing phase: The submodels in all databases must be correctly updated as,

$$W_m^{[t]} = \begin{cases} W_m^{[t-1]} + \Delta_m^{[t]}, & \text{if } m = \theta^{[t]} \\ W_m^{[t-1]}, & \text{if } m \neq \theta^{[t]}. \end{cases} \quad (5)$$

The reading and writing costs are defined as $C_R = \frac{\mathcal{D}}{L}$ and $C_W = \frac{\mathcal{U}}{L}$, where \mathcal{D} and \mathcal{U} are the total number of symbols downloaded and uploaded in the reading and writing phases, respectively. The total cost C_T is defined as $C_T = C_R + C_W$. In this paper, we use the notation $C_T(a, b)$ to indicate the total cost when the submodels are encoded with an (a, b) MDS code.

III. MAIN RESULT

Theorem 1 For a given PRUW setting with N non-colluding databases having arbitrary storage constraints $\mu(n)$, $n \in \{1, \dots, N\}$, let $k = \frac{1}{\max_n \mu(n)}$, $p = \sum_{n=1}^N \mu(n)$, $r = kp$ and $s = \lfloor k \rfloor p$.¹ Then, there exist a PRUW storage mechanism that completely fills the N databases and achieves a total communication cost of $\mathcal{C} = \min\{C_1, C_2\}$, where

$$C_1 = (\lceil s \rceil - s)C_T(\lfloor k \rfloor, \lceil s \rceil) + (s - \lfloor s \rfloor)C_T(\lfloor k \rfloor, \lfloor s \rfloor) \quad (6)$$

$$C_2 = \alpha\beta C_T(\lfloor k \rfloor, \lceil r \rceil) + \alpha(1 - \beta)C_T(\lfloor k \rfloor, \lceil r \rceil) \\ + (1 - \alpha)\delta C_T(\lceil k \rceil, \lceil r \rceil) + (1 - \alpha)(1 - \delta)C_T(\lceil k \rceil, \lceil r \rceil) \quad (7)$$

where,

$$\alpha = \begin{cases} \frac{\lfloor k \rfloor (p \lfloor k \rfloor - \lceil r \rceil)}{\lfloor k \rfloor \lceil r \rceil - \lfloor k \rfloor \lceil r \rceil}, & \text{if } r - \lceil r \rceil > k - \lfloor k \rfloor, s \leq \lceil r \rceil \\ \frac{\lfloor k \rfloor}{\lfloor k \rfloor} (\lceil k \rceil - k), & \text{else} \end{cases} \quad (8)$$

$$\beta = \begin{cases} \frac{\lceil r \rceil - r}{\lceil k \rceil - k}, & \text{if } r - \lceil r \rceil > k - \lfloor k \rfloor, s > \lceil r \rceil \\ 1, & \text{else} \end{cases} \quad (9)$$

$$\delta = \begin{cases} 1 - \frac{r - \lceil r \rceil}{k - \lfloor k \rfloor}, & \text{if } r - \lceil r \rceil \leq k - \lfloor k \rfloor \\ 0, & \text{else,} \end{cases} \quad (10)$$

¹Intuitively, r and s represent the maximum number of times a k -coded symbol can be replicated across all databases.




Fig. 2. Example setting with $k = 2.7$ and $p = 4.3$.

if $\lceil r \rceil - \lfloor k \rfloor$ is odd, and

$$\alpha = \begin{cases} \frac{\lfloor k \rfloor}{k} (\lceil k \rceil - k), & \text{if } r - \lceil r \rceil < \lceil k \rceil - k \\ \frac{\lfloor k \rfloor (p \lfloor k \rfloor - \lceil r \rceil)}{\lfloor k \rfloor \lceil r \rceil - \lfloor k \rfloor \lceil r \rceil}, & \text{else} \end{cases} \quad (11)$$

$$\beta = \begin{cases} 1 - \frac{r - \lceil r \rceil}{\lceil k \rceil - k}, & \text{if } r - \lceil r \rceil < \lceil k \rceil - k \\ 0, & \text{else} \end{cases} \quad (12)$$

$$\delta = 1 \quad (13)$$

when $\lceil r \rceil - \lfloor k \rfloor$ is even, with the function $C_T(a, b)$ defined as,

$$C_T(a, b) = \begin{cases} \frac{4b}{b-a-1}, & \text{if } b-a \text{ is odd} \\ \frac{4b-2}{b-a-2}, & \text{if } b-a \text{ is even.} \end{cases} \quad (14)$$

Remark 1 As an illustration of the main result, consider an example with $\max_n \mu(n) = 0.37$ and $p = \sum_{n=1}^N \mu(n) = 4.3$. Then, $k = 2.7$, $r = 11.61$ and $s = 8.6$, which results in $C_1 = 6.6$ as shown by the blue star in Fig. 2. Since $\lceil r \rceil - \lfloor k \rfloor = 9$ is odd, $s = 8.6 < 11 = \lceil r \rceil$ and $r - \lceil r \rceil = 0.61 < 0.7 = k - \lfloor k \rfloor$, for the calculation of C_2 , α , β and δ are given by $\alpha = \frac{\lfloor k \rfloor}{k} (\lceil k \rceil - k) = \frac{2}{9}$, $\beta = 1$ and $\delta = 1 - \frac{r - \lceil r \rceil}{k - \lfloor k \rfloor} = \frac{9}{70}$, respectively. Hence, $C_2 = 5.99$ (blue dot in Fig. 2). Note that C_1 is obtained by storing $\lceil s \rceil - s$ and $s - \lfloor s \rfloor$ fractions of parameters of all submodels using $(\lfloor k \rfloor, \lceil s \rceil)$ and $(\lfloor k \rfloor, \lfloor s \rfloor)$ MDS codes, respectively. Similarly, C_2 is obtained by storing α , $(1 - \alpha)\delta$ and $(1 - \alpha)(1 - \delta)$ fractions of all submodels using $(\lfloor k \rfloor, \lceil r \rceil)$, $(\lceil k \rceil, \lceil r \rceil)$ and $(\lceil k \rceil, \lceil r \rceil)$ MDS codes, respectively, as shown in Fig. 2. The minimum achievable cost for this example is $\min\{C_1, C_2\} = C_2$.

Remark 2 The total cost decreases with the number of replications b and increases with the coding parameter a as shown in (14). However, when the number of replications (r or s) is chosen as a linear function of the coding parameter k , i.e., $r = kp$, $s = \lfloor k \rfloor p$, the total cost decreases with the coding parameter; see Section IV-B2. Therefore, C_2 in the main result uses the maximum k by considering both $\lfloor k \rfloor$ and $\lceil k \rceil$ when $k \notin \mathbb{Z}^+$. However, based on the fluctuating structure of the total cost in (14) (dotted lines in Fig. 2) for special cases of k and p , a lower total cost can be achieved by only considering $\lfloor k \rfloor$. These cases are handled by C_1 .

Remark 3 For a given PRUW setting with arbitrary storage constraints $\{\mu(n)\}_{n=1}^N$, let the total available storage be $p = \sum_{n=1}^N \mu(n)$. Then, there exists an uncoded PRUW scheme (the proposed scheme in Section IV with $k = 1$ and $r = p$) that replicates $\lceil p \rceil - p$ and $p - \lfloor p \rfloor$ fractions of uncoded submodel parameters in $\lfloor p \rfloor$ and $\lceil p \rceil$ databases, respectively, achieving a total cost of $(\lceil p \rceil - p)C_T(1, \lfloor p \rfloor) + (p - \lfloor p \rfloor)C_T(1, \lceil p \rceil)$. Note that this total cost only depends on the sum of all arbitrary storage constraints (not on individual constraints), which is consistent with the corresponding result of PIR with uncoded heterogeneous storage constraints [45]. However, in PRUW, for a given p , the total cost is minimized when $k = \frac{1}{\max_n \mu(n)}$ is maximized, i.e., when all storage constraints are equal.

IV. PROPOSED SCHEME

The proposed scheme for FSL with heterogeneous storage constrained databases comprises two main components, namely, the storage mechanism and the PRUW scheme. In this work, we use the PRUW scheme presented in [3] and [2], and propose a storage mechanism to perform FSL on any given setting with arbitrary storage constraints. In this section, we first state the PRUW scheme briefly for completeness and then present the proposed storage mechanism.

A. PRUW Scheme

The scheme is defined for a single subpacket of all submodels (collection of y coded bits), which is then applied on all subpackets identically.

Storage: The contents of a single subpacket in database n , $n \in \{1, \dots, R\}$, at time t is given by,

$$S_n^{[t]} = \begin{bmatrix} \sum_{i=1}^K \frac{1}{f_{1,i} - \alpha_n} \begin{bmatrix} W_{1,1}^{[i]} \\ \vdots \\ W_{M,1}^{[i]} \end{bmatrix} + \sum_{j=0}^y \alpha_n^j Z_{1,j} \\ \vdots \\ \sum_{i=1}^K \frac{1}{f_{y,i} - \alpha_n} \begin{bmatrix} W_{1,y}^{[i]} \\ \vdots \\ W_{M,y}^{[i]} \end{bmatrix} + \sum_{j=0}^y \alpha_n^j Z_{y,j} \end{bmatrix}, \quad (15)$$

where $W_{m,j}^{[i]}$ is the i th parameter of the j th coded symbol of submodel m , $Z_{i,j}$ are random noise vectors of size $M \times 1$ and $f_{i,j}$ s and α_n s are globally known distinct constants from \mathbb{F}_q . Note that the j th coded data symbol (without the noise) is stored in terms of a (K, R) MDS code.

Reading phase: In the reading phase, the user sends queries to databases to download the required submodel W_θ . The queries sent to database n , $n \in \{1, \dots, R\}$, to download the Ky parameters of each subpacket of W_θ is given by,

$$Q_{n,\ell} = \begin{bmatrix} \frac{\prod_{i=1, i \neq \ell}^K (f_{1,i} - \alpha_n)}{\prod_{i=1, i \neq \ell}^K (f_{1,i} - f_{1,\ell})} e_M(\theta) + \prod_{i=1}^K (f_{1,i} - \alpha_n) \tilde{Z}_{1,\ell} \\ \vdots \\ \frac{\prod_{i=1, i \neq \ell}^K (f_{y,i} - \alpha_n)}{\prod_{i=1, i \neq \ell}^K (f_{y,i} - f_{y,\ell})} e_M(\theta) + \prod_{i=1}^K (f_{y,i} - \alpha_n) \tilde{Z}_{y,\ell} \end{bmatrix} \quad (16)$$

for $\ell \in \{1, \dots, K\}$ where $e_M(\theta)$ is the all zeros vector of size $M \times 1$ with a 1 at the θ th position and $\tilde{Z}_{j,\ell}$ are random vectors of size $M \times 1$. Then, database n generates the answers as,

$$A_{n,\ell} = S_n^T Q_{n,\ell} = \sum_{j=1}^y \frac{1}{f_{j,\ell} - \alpha_n} W_{\theta,j}^{[\ell]} + P_{\alpha_n}(K+y), \quad (17)$$

for each $\ell \in \{1, \dots, K\}$, where $P_{\alpha_n}(h)$ is a polynomial in α_n of degree h . Lemma 1 in [1] is used to obtain (17). The user obtains the Ky parameters of each subpacket of W_θ by downloading the answers in (17) from $R' = 2y + K + 1$ databases, since there are y parameters of W_θ and $K + y + 1$ polynomial coefficients to be solved for, in each set of answers for a fixed ℓ . Since the total number of databases considered is R , $R' = R$ if $R - K$ is odd and $R' = R - 1$ if $R - K$ is even, since $y \in \mathbb{Z}^+$. Therefore, the subpacketization y is given by $y = \frac{R' - K - 1}{2}$, which results in a reading cost of,

$$C_R = \frac{R' \times K}{y \times K} = \frac{2R'}{R' - K - 1}. \quad (18)$$

Writing phase: In the writing phase, the user sends K symbols to each of the R databases. The ℓ th symbol (out of the K symbols uploaded) is the combined update that consists of the updates of W_θ corresponding to the ℓ th parameter of each of the y coded parameters in each subpacket. The K combined updates sent to database n , $n \in \{1, \dots, R\}$, are given by,

$$U_{n,\ell} = \sum_{j=1}^y \prod_{i=1, i \neq j}^y (f_{i,\ell} - \alpha_n) \tilde{\Delta}_{\theta,j}^{[\ell]} + \prod_{i=1}^y (f_{i,\ell} - \alpha_n) \hat{z}_\ell, \quad (19)$$

for $\ell \in \{1, \dots, K\}$ where $\tilde{\Delta}_{\theta,j}^{[\ell]} = \frac{\prod_{i=1, i \neq \ell}^K (f_{j,i} - f_{j,\ell})}{\prod_{i=1, i \neq j}^y (f_{i,\ell} - f_{j,\ell})} \Delta_{\theta,j}^{[\ell]}$ for $j \in \{1, \dots, y\}$, with $\Delta_{\theta,j}^{[\ell]}$ being the update of the ℓ th parameter of the j th coded symbol of the considered subpacket in W_θ and \hat{z}_ℓ is a random noise symbol. Once database n receives all $U_{n,\ell}$, it calculates the incremental update as,

$$\begin{aligned} \bar{U}_{n,\ell} &= U_{n,\ell} \times \tilde{D}_{n,\ell} \times Q_{n,\ell} \\ &= \begin{bmatrix} \frac{1}{f_{1,\ell} - \alpha_n} \Delta_{\theta,1}^{[\ell]} e_M(\theta) + P_{\alpha_n}^{[1]}(y) \\ \vdots \\ \frac{1}{f_{y,\ell} - \alpha_n} \Delta_{\theta,y}^{[\ell]} e_M(\theta) + P_{\alpha_n}^{[y]}(y) \end{bmatrix}, \end{aligned} \quad (20)$$

where $\tilde{D}_{n,\ell}$ is a scaling matrix given by,

$$\tilde{D}_{n,\ell} = \text{diag} \left(\frac{1_M}{\prod_{i=1}^K (f_{1,i} - \alpha_n)}, \dots, \frac{1_M}{\prod_{i=1}^K (f_{y,i} - \alpha_n)} \right) \quad (22)$$

with the all ones vector of size $1 \times M$ indicated as 1_M . Since the incremental updates $\bar{U}_{n,\ell}$ are of the same form as the storage in (15), the submodels are updated by $S_n^{[t]} = S_n^{[t-1]} + \sum_{\ell=1}^K \bar{U}_{n,\ell}$. The resulting writing and total costs are,

$$C_W = \frac{K \times R}{K \times y} = \frac{2R}{R' - K - 1}, \quad (23)$$

$$C_T = \frac{2(R+R')}{R' - K - 1} = \begin{cases} \frac{4R}{R-K-1}, & \text{if } R-K \text{ is odd} \\ \frac{4R-2}{R-K-2}, & \text{if } R-K \text{ is even.} \end{cases} \quad (24)$$

B. Storage Mechanism

The proposed storage mechanism consists of submodel partitioning and submodel encoding, where the parameters are encoded with a specific (K, R) MDS code and stored at different subsets of databases in parts. In this section, we find the optimum coding parameters and fractions of submodels stored in each database to minimize the total cost.

1) *Submodel partitioning*: This determines the fractions of all (K, R) MDS coded submodels to be stored in the N databases with arbitrary storage constraints. Based on the (K, R) MDS coded structure, each coded parameter must be stored in exactly R databases. Let η_i be the fraction of all submodels that are stored in the same subset of R databases (the subset of databases indexed by i). Let \mathcal{B} be the basis containing all $N \times 1$ vectors with elements in $\{0, 1\}$ with exactly R ones, denoted by $\mathcal{B} = \{b_1, b_2, \dots, b_{|\mathcal{B}|}\}$. Note that each b_i corresponds to a specific subset of R databases, indexed by i . Then, it is required to find the η_i s that satisfy,

$$\frac{1}{K} \sum_{i=1}^{|\mathcal{B}|} \eta_i b_i = \mu, \quad \text{where } \mu = [\mu(1), \dots, \mu(N)]^T \quad (25)$$

$$\sum_{i=1}^{|\mathcal{B}|} \eta_i = 1, \quad 0 \leq \eta_1, \dots, \eta_{|\mathcal{B}|} \leq 1, \quad (26)$$

to replicate each coded parameter at exactly R databases, while ensuring that all databases are completely filled. The solution to this problem with uncoded parameters ($K = 1$) is provided in [45] and [46] along with a necessary and sufficient condition for a solution to exist, given by,

$$\mu(n) \leq \frac{\sum_{n=1}^N \mu(n)}{R}, \quad n \in \{1, \dots, N\}. \quad (27)$$

In this work, we find the optimum coding parameters K and R that result in the minimum total cost of the PRUW process while satisfying (27), and solve (25)-(26) to find the partitions η_i of coded submodels to be stored in each database.

2) *Submodel encoding*: For a given set of storage constraints $\{\mu(n)\}_{n=1}^N$, the total available storage is $p = \sum_{n=1}^N \mu(n)$ which is essentially the maximum number of times each uncoded parameter in the model can be replicated in the system of databases. Let the parameters of the model be stored using a (K, R) MDS code. Therefore, the total number of coded parameters in the model is $\frac{ML}{K}$, which allows each coded parameter to be replicated at a maximum of,

$$R \leq \frac{\sum_{n=1}^N \mu(n) ML}{ML/K} = Kp \quad (28)$$

databases, if $Kp \leq N$. At the same time, the storage in each database must satisfy,

$$\mu(n) ML \leq \frac{ML}{K}, \quad n \in \{1, \dots, N\} \implies K \leq \frac{1}{\bar{\mu}}, \quad (29)$$

where $\bar{\mu} = \max_n \mu(n)$, to completely utilize the available space in each database. Therefore, the coding parameter K

case	MDS code	fraction of submodel	space allocated in database n
1	$([k], [r])$	$\alpha\beta$	$\hat{\mu}_1(n)$
2	$([k], [r])$	$\alpha(1-\beta)$	$\hat{\mu}_2(n)$
3	$([k], [r])$	$(1-\alpha)\delta$	$\bar{\mu}_1(n)$
4	$([k], [r])$	$(1-\alpha)(1-\delta)$	$\bar{\mu}_2(n)$

TABLE I
FRACTIONS OF SUBMODELS AND CORRESPONDING MDS CODES.

must satisfy,

$$K \leq \min \left\{ \frac{1}{\bar{\mu}}, \frac{N}{p} \right\} = \frac{1}{\bar{\mu}}, \quad (30)$$

as $p = \sum_{n=1}^N \mu(n) \leq \bar{\mu}N$ implies $\frac{1}{\bar{\mu}} \leq \frac{N}{p}$. Since the total cost of the PRUW scheme decreases with the number of replications R , see (24), for given $K, p \in \mathbb{Z}^+$ satisfying (30), the optimum R for the (K, R) MDS code is given by Kp from (28), and the resulting total cost is,

$$C_T(K, R) = \begin{cases} \frac{4Kp}{Kp-K-1}, & \text{odd } K, \text{ even } p \\ \frac{4Kp-2}{Kp-K-2}, & \text{otherwise,} \end{cases} \quad (31)$$

which decreases with K . Therefore, for a given set of storage constraints $\{\mu(n)\}_{n=1}^N$, the minimum total cost is achieved by the (K, R) MDS code in (15) with K and R given by,

$$K = \frac{1}{\bar{\mu}} = k, \quad R = kp = r, \quad (32)$$

which automatically satisfies (27) since $\sum_{n=1}^N \mu(n) = \bar{\mu} \geq \mu(n)$ for all n . However, since k and r are not necessarily integers, it requires additional calculations to obtain the optimum (k, r) MDS codes with $k, r \in \mathbb{Z}^+$ that collectively result in the lowest total cost. Consider the general case where $k, r \notin \mathbb{Z}^+$. The general idea here is to divide each submodel into four sections and encode them using the four MDS codes given in Table I. The sizes of the four sections are defined by the parameters α, β and δ , where $0 \leq \alpha, \beta, \delta \leq 1$, as shown in Table I. The total spaces allocated for the four MDS codes in the entire system of databases are given by,

$$\sum_{n=1}^N \hat{\mu}_1(n) = \frac{\alpha\beta[r]}{[k]}, \quad \sum_{n=1}^N \hat{\mu}_2(n) = \frac{\alpha(1-\beta)[r]}{[k]} \quad (33)$$

$$\sum_{n=1}^N \bar{\mu}_1(n) = \frac{(1-\alpha)\delta[r]}{[k]}, \quad \sum_{n=1}^N \bar{\mu}_2(n) = \frac{(1-\alpha)(1-\delta)[r]}{[k]} \quad (34)$$

The space allocated for each individual MDS code in each database must satisfy (27) separately, to ensure that all databases are completely filled while also replicating each coded parameter at the respective number of databases, i.e.,

$$\hat{\mu}_1(n) \leq \frac{\alpha\beta}{[k]}, \quad \hat{\mu}_2(n) \leq \frac{\alpha(1-\beta)}{[k]} \quad (35)$$

$$\bar{\mu}_1(n) \leq \frac{(1-\alpha)\delta}{[k]}, \quad \bar{\mu}_2(n) \leq \frac{(1-\alpha)(1-\delta)}{[k]}. \quad (36)$$

All four storage allocations in each database must satisfy,

$$\hat{\mu}_1(n) + \hat{\mu}_2(n) + \bar{\mu}_1(n) + \bar{\mu}_2(n) = \mu(n), \quad \forall n. \quad (37)$$

It remains to find the values of $\hat{\mu}_1(n)$, $\hat{\mu}_2(n)$, $\bar{\mu}_1(n)$, $\bar{\mu}_2(n)$, α , β and δ that minimize the total cost given by,

$$\begin{aligned} \mathcal{C} = & \alpha(\beta C_T(\lfloor k \rfloor, \lceil r \rceil) + (1-\beta)C_T(\lfloor k \rfloor, \lceil r \rceil)) \\ & + (1-\alpha)(\delta C_T(\lceil k \rceil, \lfloor r \rfloor) + (1-\delta)C_T(\lceil k \rceil, \lfloor r \rfloor)) \end{aligned} \quad (38)$$

while satisfying (33)-(37), where C_T is defined in (14). We consider two cases, 1) $\alpha = 1$, 2) $\alpha < 1$. When $\alpha = 1$, only cases 1, 2 in Table I are used in storage, and (27) becomes $R \leq \lfloor k \rfloor p$, which results in $\lfloor k \rfloor p = s$ maximum replications. This replaces all “ r ”s in Table I and (33)-(36) by s when $\alpha = 1$.

Lemma 1 When $\alpha = 1$, β is fixed at $\beta = \lceil s \rceil - s$ to satisfy (33) and (37), and $\hat{\mu}_1(n)$, $\hat{\mu}_2(n)$ satisfying (35) are given by,

$$\hat{\mu}_1(n) = \tilde{m}(n) + (\mu(n) - \tilde{m}(n) - \tilde{h}(n))\tilde{\gamma} \quad (39)$$

$$\hat{\mu}_2(n) = \tilde{h}(n) + (\mu(n) - \tilde{m}(n) - \tilde{h}(n))(1 - \tilde{\gamma}) \quad (40)$$

for all $n \in \{1, \dots, N\}$ where,

$$\tilde{m}(n) = \left[\mu(n) - \frac{s - \lfloor s \rfloor}{\lfloor k \rfloor} \right]^+, \quad \tilde{h}(n) = \left[\mu(n) - \frac{\lceil s \rceil - s}{\lfloor k \rfloor} \right]^+ \quad (41)$$

$$\tilde{\gamma} = \frac{\frac{\lfloor s \rfloor}{\lfloor k \rfloor} (\lceil s \rceil - s) - \sum_{n=1}^N \tilde{m}(n)}{p - \sum_{n=1}^N \tilde{m}(n) - \sum_{n=1}^N \tilde{h}(n)}, \quad (42)$$

with $[x]^+ = \max\{x, 0\}$. The resulting total cost is given in (6).

Lemma 2 The following values of $\hat{\mu}_1(n)$, $\hat{\mu}_2(n)$, $\bar{\mu}_1(n)$ and $\bar{\mu}_2(n)$ for $n \in \{1, \dots, N\}$ satisfy (33)-(37) with any $\alpha < 1$, β and δ that satisfy $\alpha \geq \frac{\lfloor k \rfloor}{k}(\lceil k \rceil - k)$, $\beta \geq \left[1 - \frac{\lfloor k \rfloor}{k\alpha}(r - \lceil r \rceil)\right]^+$ and $\delta \geq \left[1 - \frac{\lfloor k \rfloor}{k(1-\alpha)}(r - \lfloor r \rfloor)\right]^+$.

$$\hat{\mu}_1(n) = \begin{cases} \hat{\mu}(n)\beta, & \text{if } \beta \in \{0, 1\} \\ \hat{m}(n) + (\hat{\mu}(n) - \hat{m}(n) - \hat{h}(n))\hat{\gamma}, & \text{if } \beta \in (0, 1) \end{cases} \quad (43)$$

$$\hat{\mu}_2(n) = \begin{cases} \hat{\mu}(n)(1 - \beta), & \text{if } \beta \in \{0, 1\} \\ \hat{h}(n) + (\hat{\mu}(n) - \hat{m}(n) - \hat{h}(n))(1 - \hat{\gamma}), & \text{if } \beta \in (0, 1) \end{cases} \quad (44)$$

$$\bar{\mu}_1(n) = \begin{cases} \bar{\mu}(n)\delta, & \text{if } \delta \in \{0, 1\} \\ \bar{m}(n) + (\bar{\mu}(n) - \bar{m}(n) - \bar{h}(n))\bar{\gamma}, & \text{if } \delta \in (0, 1) \end{cases} \quad (45)$$

$$\bar{\mu}_2(n) = \begin{cases} \bar{\mu}(n)(1 - \delta), & \text{if } \delta \in \{0, 1\} \\ \bar{h}(n) + (\bar{\mu}(n) - \bar{m}(n) - \bar{h}(n))(1 - \bar{\gamma}), & \text{if } \delta \in (0, 1) \end{cases} \quad (46)$$

for all $n \in \{1, \dots, N\}$, where,

$$\hat{\mu}(n) = m(n) + (\mu(n) - m(n) - h(n))\gamma \quad (47)$$

$$\bar{\mu}(n) = h(n) + (\mu(n) - m(n) - h(n))(1 - \gamma) \quad (48)$$

$$m(n) = \left[\mu(n) - \frac{1 - \alpha}{\lfloor k \rfloor} \right]^+, \quad h(n) = \left[\mu(n) - \frac{\alpha}{\lfloor k \rfloor} \right]^+ \quad (49)$$

$$\gamma = \frac{\frac{\alpha}{\lfloor k \rfloor} (\lceil r \rceil - \beta) - \sum_{n=1}^N m(n)}{p - \sum_{n=1}^N m(n) - \sum_{n=1}^N h(n)} \quad (50)$$

$$\hat{m}(n) = \left[\hat{\mu}(n) - \frac{\alpha(1 - \beta)}{\lfloor k \rfloor} \right]^+, \quad \hat{h}(n) = \left[\hat{\mu}(n) - \frac{\alpha\beta}{\lfloor k \rfloor} \right]^+ \quad (51)$$

$$\hat{\gamma} = \frac{\frac{\alpha\beta}{\lfloor k \rfloor} \lceil r \rceil - \sum_{n=1}^N \hat{m}(n)}{\frac{\alpha}{\lfloor k \rfloor} (\lceil r \rceil - \beta) - \sum_{n=1}^N \hat{m}(n) - \sum_{n=1}^N \hat{h}(n)} \quad (52)$$

$$\bar{m}(n) = \left[\bar{\mu}(n) - \frac{(1 - \alpha)(1 - \delta)}{\lfloor k \rfloor} \right]^+, \quad \bar{h}(n) = \left[\bar{\mu}(n) - \frac{(1 - \alpha)\delta}{\lfloor k \rfloor} \right]^+ \quad (53)$$

$$\bar{\gamma} = \frac{\frac{(1 - \alpha)\delta}{\lfloor k \rfloor} \lceil r \rceil - \sum_{n=1}^N \bar{m}(n)}{\frac{1 - \alpha}{\lfloor k \rfloor} (\lceil r \rceil - \delta) - \sum_{n=1}^N \bar{m}(n) - \sum_{n=1}^N \bar{h}(n)}. \quad (54)$$

Lemma 3 For the case where $\alpha < 1$, the values of α , β and δ that minimize the total cost in (38) while satisfying the constraints in Lemma 2 are specified in Theorem 1, and the corresponding total cost is given in (7).

Once all parameters in Table I are determined from Lemmas 1-3, the methods proposed in [45] and [46] are used to find the solutions to (25)-(26), i.e., η_i s for the four sets of storage allocations $\{\hat{\mu}_1\}_{n=1}^N$, $\{\hat{\mu}_2\}_{n=1}^N$, $\{\bar{\mu}_1\}_{n=1}^N$, $\{\bar{\mu}_2\}_{n=1}^N$ separately. These four sets of solutions determine where each individual coded symbol is replicated in the system of databases, which completes the storage mechanism.

As an example, consider a PRUW setting with $N = 12$ databases, where $\mu(1) = \dots = \mu(5) = 0.37$ and $\mu(6) = \dots = \mu(12) = 0.35$. Note that the values of k , r and s here are the same as what is considered in Remark 1, and therefore, result in the same α , β , δ and total cost stated in Remark 1. Using (43)-(54), we find the storage allocations as,

$$\hat{\mu}_1(n) = \hat{\mu}(n) = \begin{cases} 0.1107, & \text{for } n = 1, \dots, 5 \\ 0.0951, & \text{for } n = 6, \dots, 12 \end{cases} \quad (55)$$

$$\bar{\mu}_1(n) = \begin{cases} 0.033, & \text{for } n = 1, \dots, 5 \\ 0.029, & \text{for } n = 6, \dots, 12 \end{cases} \quad (56)$$

$$\hat{\mu}_2(n) = 0, \quad \bar{\mu}_2(n) = 0.226, \quad \forall n. \quad (57)$$

Note from (57) that the $(\lfloor k \rfloor, \lceil r \rceil)$ MDS code is not used, and the allocated space for the $(\lceil k \rceil, \lfloor r \rfloor)$ MDS code in all databases is the same. Therefore, the optimum storage mechanism for PRUW with homogeneous storage constraints presented in [3] is used to place the $(\lceil k \rceil, \lfloor r \rfloor)$ coded parameters. The storage allocations for $(\lfloor k \rfloor, \lceil r \rceil)$ and $(\lceil k \rceil, \lfloor r \rfloor)$ codes have different storage allocations for different databases as shown by the two different sets of values in (55) and (56). Therefore, for these two cases, we solve (25)-(26) separately, to find the fractions of coded parameters to be stored in each database. For example, for the $(\lfloor k \rfloor, \lceil r \rceil)$ MDS code, we find the solution to (25)-(26) with $N = 12$, $K = \lfloor k \rfloor = 2$, $R = \lceil r \rceil = 11$ and $\mu = [\hat{\mu}_1(1) : \hat{\mu}_1(12)]$ as $\eta_1 = \dots = \eta_5 = 2.73 \times 10^{-4}$ and $\eta_6 = \dots = \eta_{12} = 3.15 \times 10^{-2}$, where the b_i in (25) corresponding to each η_i is the all ones vector of size 12×1 with a zero at the i th position. This essentially means that the storage mechanism assigns an η_i fraction of $(\lfloor k \rfloor, \lceil r \rceil)$ coded parameters to the subset of $\lceil r \rceil = 11$ databases that contain all databases except the i th one.

REFERENCES

- [1] S. Vithana and S. Ulukus. Private read update write (PRUW) in federated submodel learning (FSL): Communication efficient schemes with and without sparsification. Available online at arXiv:2209.04421.
- [2] Z. Jia and S. A. Jafar. X -secure T -private federated submodel learning with elastic dropout resilience. *IEEE Transactions on Information Theory*, 68(8):5418–5439, August 2022.
- [3] S. Vithana and S. Ulukus. Private read update write (PRUW) with storage constrained databases. In *IEEE ISIT*, June 2022.
- [4] S. Vithana and S. Ulukus. Private federated submodel learning with sparsification. In *IEEE ITW*, November 2022.
- [5] S. Vithana and S. Ulukus. Rate distortion tradeoff in private read update write in federated submodel learning. In *Asilomar Conference*, October 2022.
- [6] S. Vithana and S. Ulukus. Efficient private federated submodel learning. In *IEEE ICC*, May 2022.
- [7] S. Vithana and S. Ulukus. Model segmentation for storage efficient private federated learning with top r sparsification. In *CISS*, March 2023.
- [8] S. Vithana and S. Ulukus. Rate-privacy-storage tradeoff in federated learning with top r sparsification. In *IEEE ICC*, May 2023.
- [9] C. Niu, F. Wu, S. Tang, L. Hua, R. Jia, C. Lv, Z. Wu, and G. Chen. Secure federated submodel learning. Available online at arXiv:1911.02254.
- [10] M. Kim and J. Lee. Information-theoretic privacy in federated submodel learning. Available online at arXiv:2008.07656.
- [11] C. Niu, F. Wu, S. Tang, L. Hua, R. Jia, C. Lv, Z. Wu, and G. Chen. Billion-scale federated learning on mobile clients: A submodel design with tunable privacy. In *MobiCom*, April 2020.
- [12] S. Ulukus, S. Avestimehr, M. Gastpar, S. A. Jafar, R. Tandon, and C. Tian. Private retrieval, computing and learning: Recent progress and future challenges. *IEEE JSAC*, 40(3):729–748, March 2022.
- [13] B. Chor, E. Kushilevitz, O. Goldreich, and M. Sudan. Private information retrieval. *Journal of the ACM*, 45(6):965–981, November 1998.
- [14] H. Sun and S. A. Jafar. The capacity of private information retrieval. *IEEE Transactions on Information Theory*, 63(7):4075–4088, July 2017.
- [15] C. Tian, H. Sun, and J. Chen. Capacity-achieving private information retrieval codes with optimal message size and upload cost. *IEEE Transactions on Information Theory*, 65(11):7613–7627, November 2019.
- [16] S. Kumar, H.-Y. Lin, E. Rosnes, and A. G. i Amat. Achieving maximum distance separable private information retrieval capacity with linear codes. *IEEE Transactions on Information Theory*, 65(7):4243–4273, July 2019.
- [17] H. Sun and S. A. Jafar. The capacity of symmetric private information retrieval. *IEEE Transactions on Information Theory*, 65(1):322–329, January 2019.
- [18] Z. Wang, K. Banawan, and S. Ulukus. Private set intersection: A multi-message symmetric private information retrieval perspective. *IEEE Transactions on Information Theory*, 68(3):2001–2019, March 2022.
- [19] K. Banawan and S. Ulukus. Multi-message private information retrieval: Capacity results and near-optimal schemes. *IEEE Transactions on Information Theory*, 64(10):6842–6862, October 2018.
- [20] T. Chan, S. Ho, and H. Yamamoto. Private information retrieval for coded storage. In *IEEE ISIT*, June 2015.
- [21] A. Fazeli, A. Vardy, and E. Yaakobi. Codes for distributed PIR with low storage overhead. In *IEEE ISIT*, June 2015.
- [22] K. Banawan and S. Ulukus. The capacity of private information retrieval from coded databases. *IEEE Transactions on Information Theory*, 64(3):1945–1956, March 2018.
- [23] R. Tajeddine and S. El Rouayheb. Robust private information retrieval on coded data. In *IEEE ISIT*, June 2017.
- [24] H. Sun and S. A. Jafar. Multiround private information retrieval: Capacity and storage overhead. *IEEE Transactions on Information Theory*, 64(8):5743–5754, August 2018.
- [25] H. Sun and S. A. Jafar. Optimal download cost of private information retrieval for arbitrary message length. *IEEE Transactions on Information Forensics and Security*, 12(12):2920–2932, December 2017.
- [26] Q. Wang and M. Skoglund. On pir and symmetric pir from colluding databases with adversaries and eavesdroppers. *IEEE Transactions on Information Theory*, 65(5):3183–3197, May 2019.
- [27] H. Sun and S. A. Jafar. The capacity of robust private information retrieval with colluding databases. *IEEE Transactions on Information Theory*, 64(4):2361–2370, April 2018.
- [28] R. Freij-Hollanti, O. Gnilke, C. Hollanti, and D. Karpuk. Private information retrieval from coded databases with colluding servers. *SIAM Journal on Applied Algebra and Geometry*, 1(1):647–664, 2017.
- [29] H. Sun and S. A. Jafar. Private information retrieval from MDS coded data with colluding servers: Settling a conjecture by Freij-Hollanti et al. *IEEE Transactions on Information Theory*, 64(2):1000–1022, February 2018.
- [30] R. Tandon. The capacity of cache aided private information retrieval. In *Allerton Conference*, October 2017.
- [31] Y.-P. Wei, K. Banawan, and S. Ulukus. Fundamental limits of cache-aided private information retrieval with unknown and uncoded prefetching. *IEEE Transactions on Information Theory*, 65(5):3215–3232, May 2019.
- [32] I. Samy, M. Attia, R. Tandon, and L. Lazos. Asymmetric leaky private information retrieval. *IEEE Transactions on Information Theory*, 67(8):5352–5369, August 2021.
- [33] I. Samy, R. Tandon, and L. Lazos. On the capacity of leaky private information retrieval. In *IEEE ISIT*, July 2019.
- [34] S. Kadhe, B. Garcia, A. Heidarzadeh, S. El Rouayheb, and A. Sprintson. Private information retrieval with side information. *IEEE Transactions on Information Theory*, 66(4):2032–2043, April 2020.
- [35] Q. Wang, H. Sun, and M. Skoglund. The capacity of private information retrieval with eavesdroppers. *IEEE Transactions on Information Theory*, 65(5):3198–3214, May 2019.
- [36] Q. Wang, H. Sun, and M. Skoglund. The capacity of private information retrieval with eavesdroppers. *IEEE Transactions on Information Theory*, 65(5):3198–3214, May 2019.
- [37] K. Banawan and S. Ulukus. The capacity of private information retrieval from Byzantine and colluding databases. *IEEE Transactions on Information Theory*, 65(2):1206–1219, February 2019.
- [38] S. Vithana, K. Banawan, and S. Ulukus. Semantic private information retrieval. *IEEE Transactions on Information Theory*, 68(4):2635–2652, April 2022.
- [39] S. Li and M. Gastpar. Single-server multi-message private information retrieval with side information: the general cases. In *IEEE ISIT*, June 2020.
- [40] Z. Jia and S. A. Jafar. X -secure T -private information retrieval from MDS coded storage with Byzantine and unresponsive servers. *IEEE Transactions on Information Theory*, 66(12):7427–7438, December 2020.
- [41] Z. Jia, H. Sun, and S. A. Jafar. Cross subspace alignment and the asymptotic capacity of X -secure T -private information retrieval. *IEEE Transactions on Information Theory*, 65(9):5783–5798, September 2019.
- [42] J. Xu and Z. Zhang. Building capacity-achieving PIR schemes with optimal sub-packetization over small fields. In *IEEE ISIT*, June 2018.
- [43] H. Yang, W. Shin, and J. Lee. Private information retrieval for secure distributed storage systems. *IEEE Transactions on Information Forensics and Security*, 13(12):2953–2964, December 2018.
- [44] K. Banawan and S. Ulukus. Private information retrieval from multiple access channels. In *IEEE ITW*, November 2018.
- [45] K. Banawan, B. Arasli, Y.-P. Wei, and S. Ulukus. The capacity of private information retrieval from heterogeneous uncoded caching databases. *IEEE Transactions on Information Theory*, 66(6):3407–3416, June 2020.
- [46] N. Woolsey, R. Chen, and M. Ji. Uncoded placement with linear sub-messages for private information retrieval from storage constrained databases. *IEEE Transactions on Communications*, 68(10):6039–6053, October 2020.
- [47] K. Banawan, B. Arasli, and S. Ulukus. Improved storage for efficient private information retrieval. In *IEEE ITW*, August 2019.
- [48] Y.-P. Wei, B. Arasli, K. Banawan, and S. Ulukus. The capacity of private information retrieval from decentralized uncoded caching databases. *Information*, 10, December 2019.
- [49] R. Tandon, M. Abdul-Wahid, F. Almoalem, and D. Kumar. PIR from storage constrained databases - coded caching meets PIR. *IEEE ICC*, May 2018.
- [50] M. A. Attia, D. Kumar, and R. Tandon. The capacity of private information retrieval from uncoded storage constrained databases. *IEEE Transactions on Information Theory*, 66(11):6617–6634, November 2020.
- [51] Y.-P. Wei and S. Ulukus. The capacity of private information retrieval with private side information under storage constraints. *IEEE Transactions on Information Theory*, 66(4):2023–2031, April 2020.