The Capacity of Uncoded Storage Constrained PIR

Mohamed Adel Attia Deepak Kumar Ravi Tandon
Department of Electrical and Computer Engineering
University of Arizona, Tucson, AZ
Email: {*madel, deepakkumar, tandonr*}@email.arizona.edu

***Abstract*—Private information retrieval (PIR) allows a user to retrieve a desired message out of K possible messages from N databases (DBs) without revealing the identity of the desired message. In this work, we consider the problem of PIR from *uncoded storage constrained* DBs. Each DB has a storage capacity of μKL bits, where L is the size of each message in bits, and $\mu \in [1/N, 1]$ is the normalized storage. In the storage constrained PIR problem, there are two key challenges: a) construction of communication efficient schemes through storage content design at each DB that allow download efficient PIR; and b) characterizing the optimal download cost via information-theoretic lower bounds. The novel aspect of this work is to characterize the optimum download cost of PIR with storage constrained DBs for any value of storage. In particular, for any (N, K), we show that the optimal tradeoff between storage (μ) and the download cost ($D(\mu)$) is given by the lower convex hull of the pairs $(\frac{t}{N}, \big(1 + \frac{1}{t} + \frac{1}{t^2} + \cdots + \frac{1}{t^{K-1}}\big))$ for $t = 1, 2, \ldots, N$. The main contribution of this paper is the converse proof, i.e., obtaining lower bounds on the download cost for PIR as a function of the available storage.**

I. Introduction

The classical private information retrieval (PIR) problem involves N non-colluding databases (DBs), each DB stores K messages. The goal of the user is to efficiently retrieve a message without revealing the message identity. Based on the Shannon theoretic formulation, the rate of a PIR scheme is the ratio between the number of desired vs downloaded bits, and PIR capacity is then defined as the maximum achievable rate. Recently in [1], the capacity of PIR (or the inverse of download cost) was characterized as $(1 + 1/N + \ldots 1/N^{K-1})^{-1}$.

Since the appearance of [1], significant recent progress has been made on a variety of variations of the basic PIR problem. We briefly describe some of these advances next. The case of T-colluding PIR (or TPIR in short) was investigated in [2], where any T DBs out of N are able to collude. Robust PIR, in which any subset of DBs fail to respond was also investigated in [2]. In a recent work, [3] characterized the capacity of PIR with byzantine DBs (or BPIR), i.e., a scenario in which any subset of DBs are adverserial (i.e. they can respond with incorrect answers). The problem of PIR with DBs storing MDS coded messages was considered in [4] and the capacity was subsequently characterized by Banawan and Ulukus in [5]. The problem of symmetric PIR (SPIR) was studied in [6]. In this setting, privacy is enforced in both directions, i.e., user must be able to retrieve the message of interest privately, while at the same time the DBs must avoid any information leakage to the user about the remaining messages. The capacity of cache aided PIR (in which the user has a local cache of limited storage) was recently characterized in [7], and it was shown that memory sharing based PIR scheme is information-theoretically optimal (also see recent works [8]–[10] on other variations of the cache aided PIR problem).

Majority of above works, however, assume the presence of replicated DBs, each storing all the K messages. Indeed, exceptions to this statement include the works on the case when DBs store MDS coded data for robust PIR. Furthermore, [11] also investigated the problem of limited storage PIR for the special case of $K = 2$ messages and $N = 2$ DBs. They present interesting lower and upper bounds on the capacity for this special case, and show the optimality of the proposed scheme for the case of linear schemes. However, the generalization to any (N, K) parameters for arbitrary storage constrained PIR problem remains elusive.

Summary of Contribution– In this work, we consider the problem of PIR for *uncoded storage constrained* DBs. Each DB has a storage capacity of μKL bits, where K is the number of messages, L is the size of each message in bits, and $\mu \in [1/N, 1]$ is the normalized storage. On one extreme, when $\mu = 1/N$, then the user has to download all the K messages to achieve privacy. On the other hand, $\mu = 1$ is the replicated DBs case setteled in [1], where the download cost is minimal. Thus, we aim to characterize this trade-off for any value of μ.

In our prior work, [12], we presented an achievable scheme for this problem which works for all (N, K) and all storage parameters μ. The main contribution of this paper is to show that the scheme in [12] is information-theoretically optimal for uncoded storage placement strategies. The key technical challenge is in proving the lower bounds, which go beyond the techniques introduced in [1], to incorporate the fact that PIR must be feasible from storage constrained DBs. The main differences, however, is that we retain the terms of information theoretic capacity which was discarded in [1].

II. Storage Constrained PIR: System Model

We consider the PIR problem with N non-colluding DBs and K independent messages $W_1, \ldots, W_K$, where each message is of size L bits, i.e., $H(W_k) = L, \forall k$. We assume that each DB has a storage capacity of μKL bits. Denoting $Z_1, \ldots, Z_N$ as the contents stored across the DBs, then we have the storage constraints as $H(Z_n) = \mu KL$, for $n \in [1 : N]$. For the scope of this paper, we focus on uncoded storage, i.e., each DB can store uncoded bits of each message subject to

This work was supported by the NSF grant CAREER-1651492.

the storage constraint. Furthermore, we assume that the storage strategy employed by the user is completely public, i.e., each DB knows which contents are stored at all the other DBs.

The normalized storage parameter μ can take values in the range $1/N \leq \mu \leq 1$. The case when $\mu = 1$ is the setting of replicated DBs, with each DB storing all the K messages. The lower bound $\mu \geq 1/N$ is in fact a necessary condition for reliable decoding. To request a message, a user privately selects a number k between 1 and K corresponding to the desired message W_k. Then the user generates N queries $Q_1^{[k]}, Q_2^{[k]}, \ldots, Q_N^{[k]}$, where $Q_n^{[k]}$ is sent to the nth DB (DB$_n$), and the queries are independent of the messages, i.e.,

$$I(W_1, \ldots, W_K; Q_1^{[k]}, \ldots, Q_N^{[k]}) = 0, \quad \forall k \in [1:K]. \tag{1}$$

Upon receiving the query $Q_n^{[k]}$, DB$_n$ returns an answer $A_n^{[k]}$ to the user, which is a function of the corresponding query and the data stored in the DB$_n$, i.e.,

$$H(A_n^{[k]} | Q_n^{[k]}, Z_n) = 0, \quad \forall k \in [1:K], \ \forall n \in [1:N]. \tag{2}$$

From all the answers, the user must be able to correctly decode the desired message W_k, i.e., the following correctness constraint must be satisfied for all $k \in [1:K]$:

$$H(W_k | A_1^{[k]}, \ldots, A_N^{[k]}, Q_1^{[k]}, \ldots, Q_N^{[k]}) = o(L), \tag{3}$$

where $o(L)$ represents a function where $o(L)/L \to 0$ as $L \to \infty$. In order to prevent the DBs from learning the identity of requested message, privacy must be guaranteed through the following statistical equivalence for all $k_1 \neq k_2 \in [1:K]$:

$$\begin{aligned}(Q_n^{[k_1]}, A_n^{[k_1]}, W_1, \ldots, W_K, Z_1, \ldots, Z_N) \\ \sim (Q_n^{[k_2]}, A_n^{[k_2]}, W_1, \ldots, W_K, Z_1, \ldots, Z_N).\end{aligned} \tag{4}$$

For a storage parameter μ, we say that a pair (D, L) is achievable if there exists a PIR scheme with storage, querying, and decoding functions, which satisfy the above constraints. The performance of a PIR scheme is characterized by the number of bits of desired information per downloaded bit. In particular, if D is the total number of downloaded bits, and L is the size of the desired message, then the normalized downloaded cost is D/L. In other words, the PIR rate is L/D. The goal is to characterize the optimal normalized download cost as a function of the DB storage parameter μ:

$$D^*(\mu) = \min \{D/L : (D, L) \text{ is achievable}\}. \tag{5}$$

The storage constrained capacity of PIR is the inverse of the normalized download cost, i.e., $C^*(\mu) = \max \{L/D : (D, L) \text{ is achievable}\}$. In [12], we devised an achievable scheme for storage constrained PIR problem. We next present the main result of this paper which shows that the scheme in [12] is information-theoretically optimal.

III. Main Result and Discussions

Theorem 1: For the uncoded storage constrained PIR problem with N DBs, K messages (of size L bits each), and a per DB storage constraint of μKL bits, the information-theoretically optimal tradeoff between storage and download cost is given by the lower convex hull of the following $(\mu, D(\mu))$ pairs, for $t = 1, 2, \ldots, N$,

$$(\mu, \ D(\mu)) = \left(\frac{t}{N}, \ 1 + \frac{1}{t} + \cdots + \frac{1}{t^{K-1}} \right). \tag{6}$$

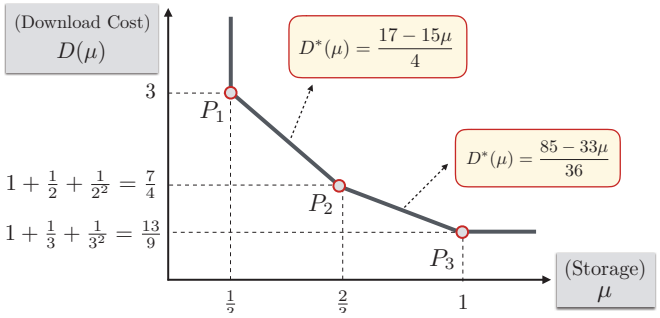

Fig. 1. Optimal Tradeoff between download and storage for $(N, K) = (3, 3)$.

Fig. 1 shows the optimal tradeoff between download cost and storage for $N = 3$ DBs and $K = 3$ messages. We can make the following interesting observations from this figure: a) $\mu = 1$ corresponds to replicated DBs, for which the optimal download cost is the same as that in [1]; b) $\mu = 1/N$ corresponds to the minimum value of storage, for which the optimal scheme is to download all messages to ensure privacy; and c) for intermediate values, i.e., $1/N < \mu < 1$, the optimal trade off is achieved by memory sharing between different PIR schemes, which are designed for specific values of μ (in this example, $\mu = 1/3, 2/3$ and $\mu = 1$).

Achievable Scheme– The general achievable scheme, for any (N, K) and any μ, has been presented in our prior work [12]. We briefly describe the idea here for the example of $N = K = 3$. The three messages are denoted as A, B and C. Following Theorem 1, the tradeoff has three critical points as shown in Fig. 1 (labeled as P_1, P_2 and P_3). As the point P_1 is trivial and P_3 follows from [1], we focus on point P_2, where $\mu = t/N = 2/3$. For this point, we assume that each message is of size $L = 2^3 \times 3 = 24$ bits. In particular, we split each message into $\binom{3}{2} = 3$ sub-messages and label each by a unique subset of $[1:3]$ of size $t = 2$, e.g., $A = \{a_{\{1,2\}}, a_{\{1,3\}}, a_{\{2,3\}}\}$, and each sub-message is of size $2^3 = 8$ bits:

$$\begin{aligned}A &= \{a^1_{\{1,2\}}, \ldots, a^8_{\{1,2\}}, a^1_{\{1,3\}}, \ldots, a^8_{\{1,3\}}, a^1_{\{2,3\}}, \ldots, a^8_{\{2,3\}}\}, \\ B &= \{b^1_{\{1,2\}}, \ldots, b^8_{\{1,2\}}, b^1_{\{1,3\}}, \ldots, b^8_{\{1,3\}}, b^1_{\{2,3\}}, \ldots, b^8_{\{2,3\}}\}, \\ C &= \{c^1_{\{1,2\}}, \ldots, c^8_{\{1,2\}}, c^1_{\{1,3\}}, \ldots, c^8_{\{1,3\}}, c^1_{\{2,3\}}, \ldots, c^8_{\{2,3\}}\}.\end{aligned}$$

Subsequently, DB$_n$ stores those sub-messages (of each message) whose index contains n. For instance, DB$_1$ stores $\{a_{\{1,2\}}, a_{\{1,3\}}, b_{\{1,2\}}, b_{\{1,3\}}, c_{\{1,2\}}, c_{\{1,3\}}\}$ Hence, the total storage required per DB is $6 \times 8 = 48$ bits. The PIR scheme works in three stages as shown in Table I (assuming that the user wants message A).

We can readily verify that from the 42 bits downloaded from all three DBs, the user is able to correctly retrieve the 24 bits of message A. Hence, the download cost of the scheme is $D(\mu = \frac{2}{3}) = \frac{42}{24} = 7/4$, and point P_2 is achieved. The intermediate values of μ, between the points P_1, P_2, and P_3,

TABLE I
STORAGE CONSTRAINED PIR: $(N,K)=(3,3), \mu=\frac{t}{N}=\frac{2}{3}$

DB_1	DB_2	DB_3
$a^1_{\{1,2\}}, b^1_{\{1,2\}}, c^1_{\{1,2\}}$ $a^1_{\{1,3\}}, b^1_{\{1,3\}}, c^1_{\{1,3\}}$	$a^2_{\{1,2\}}, b^2_{\{1,2\}}, c^2_{\{1,2\}}$ $a^1_{\{2,3\}}, b^1_{\{2,3\}}, c^1_{\{2,3\}}$	$a^2_{\{1,3\}}, b^2_{\{1,3\}}, c^2_{\{1,3\}}$ $a^2_{\{2,3\}}, b^2_{\{2,3\}}, c^2_{\{2,3\}}$
$a^3_{\{1,2\}}+b^2_{\{1,2\}}$ $a^3_{\{1,3\}}+b^2_{\{1,3\}}$ $a^4_{\{1,2\}}+c^2_{\{1,2\}}$ $a^4_{\{1,3\}}+c^2_{\{1,3\}}$ $b^3_{\{1,2\}}+c^3_{\{1,2\}}$ $b^3_{\{1,3\}}+c^3_{\{1,3\}}$	$a^5_{\{1,2\}}+b^1_{\{1,2\}}$ $a^3_{\{2,3\}}+b^2_{\{2,3\}}$ $a^6_{\{1,2\}}+c^1_{\{1,2\}}$ $a^4_{\{2,3\}}+c^2_{\{2,3\}}$ $b^4_{\{1,2\}}+c^4_{\{1,2\}}$ $b^3_{\{2,3\}}+c^3_{\{2,3\}}$	$a^5_{\{1,3\}}+b^1_{\{1,3\}}$ $a^5_{\{2,3\}}+b^1_{\{2,3\}}$ $a^6_{\{1,3\}}+c^1_{\{1,3\}}$ $a^6_{\{2,3\}}+c^1_{\{2,3\}}$ $b^4_{\{1,3\}}+c^4_{\{1,3\}}$ $b^4_{\{2,3\}}+c^4_{\{2,3\}}$
$a^7_{\{1,2\}}+b^4_{\{1,2\}}+c^4_{\{1,2\}}$ $a^7_{\{1,3\}}+b^4_{\{1,3\}}+c^4_{\{1,3\}}$	$a^8_{\{1,2\}}+b^3_{\{1,2\}}+c^3_{\{1,2\}}$ $a^7_{\{2,3\}}+b^4_{\{2,3\}}+c^4_{\{2,3\}}$	$a^8_{\{1,3\}}+b^3_{\{1,3\}}+c^3_{\{1,3\}}$ $a^8_{\{2,3\}}+b^3_{\{2,3\}}+c^3_{\{2,3\}}$

can be achieved by memory-sharing (see [12, Lemma 1] for more details), showing that the lower convex hull is achievable.

Converse Proof– Due to space limitations, the main ideas behind the general converse proof are presented through the example of $N=3$ DBs and $K=3$ messages, while the complete general proof can be found in [13]. From Fig. 1, it is clear that we need to show the following two bounds:

$$D^*(\mu) \geq \frac{17-15\mu}{4}; \qquad D^*(\mu) \geq \frac{85-33\mu}{36}. \tag{7}$$

To this end, let us start with the following bound on D:

$$D - L + o(L) \geq I(W_{[2:3]}; Q^{[1]}_{[1:3]}, A^{[1]}_{[1:3]} | W_1), \tag{8}$$

which follows from [1, Lemma 1]. We next show that the term in RHS above can be lower bounded as

$$I(W_{[2:3]}; Q^{[1]}_{[1:3]}, A^{[1]}_{[1:3]} | W_1) \geq \underbrace{I(W_{[2:3]}; Q^{[1]}_1, A^{[1]}_1 | W_1)}_{\triangleq \text{ Term}_1} \tag{9}$$
$$+\underbrace{I(W_{[2:3]}; Q^{[1]}_2, A^{[1]}_2 | W_1, Z_1)}_{\triangleq \text{ Term}_2} + \underbrace{I(W_{[2:3]}; Q^{[1]}_3, A^{[1]}_3 | W_1, Z_{[1:2]})}_{\triangleq \text{ Term}_3}.$$

To prove the above bound, we use the chain rule for mutual information and expand the above term as

$$\begin{aligned} I(W_{[2:3]}; Q^{[1]}_{[1:3]}, A^{[1]}_{[1:3]} | W_1) &= I(W_{[2:3]}; Q^{[1]}_1, A^{[1]}_1 | W_1) \\ &\quad + I(W_{[2:3]}; Q^{[1]}_2, A^{[1]}_2 | W_1, Q^{[1]}_1, A^{[1]}_1) \\ &\quad + I(W_{[2:3]}; Q^{[1]}_3, A^{[1]}_3 | W_1, Q^{[1]}_{[1:2]}, A^{[1]}_{[1:2]}). \end{aligned} \tag{10}$$

We note that the first term in the RHS of (10) is the same as Term_1. We next show that the second term in the RHS of (10) is lower bounded by Term_2 as follows:

$$\begin{aligned} &I(W_{[2:3]}; Q^{[1]}_2, A^{[1]}_2 | W_1, Q^{[1]}_1, A^{[1]}_1) \\ &= H(Q^{[1]}_2, A^{[1]}_2 | W_1, Q^{[1]}_1, A^{[1]}_1) - H(Q^{[1]}_2, A^{[1]}_2 | W_{[1:3]}, Q^{[1]}_1, A^{[1]}_1) \\ &\overset{(a)}{\geq} H(Q^{[1]}_2, A^{[1]}_2 | W_1, Z_1, Q^{[1]}_1, A^{[1]}_1) - H(Q^{[1]}_2, A^{[1]}_2 | W_{[1:3]}, Z_1, Q^{[1]}_1, A^{[1]}_1) \\ &\overset{(b)}{=} H(Q^{[1]}_2, A^{[1]}_2 | W_1, Z_1, Q^{[1]}_1) - H(Q^{[1]}_2, A^{[1]}_2 | W_{[1:3]}, Z_1, Q^{[1]}_1) \\ &= I(W_{[2:3]}; Q^{[1]}_2, A^{[1]}_2 | W_1, Z_1, Q^{[1]}_1) \\ &\overset{(c)}{=} I(W_{[2:3]}; Q^{[1]}_2, A^{[1]}_2, Q^{[1]}_1 | W_1, Z_1) \\ &\geq I(W_{[2:3]}; Q^{[1]}_2, A^{[1]}_2 | W_1, Z_1), \end{aligned} \tag{11}$$

where (a) follows from the fact that conditioning reduces entropy (which allows us to introduce Z_1 in the first term), and the fact that Z_1 is a function of $W_{[1:3]}$ (hence it can be introduced in the second term); (b) follows from the fact that $A^{[1]}_1$ is a function of $(Z_1, Q^{[1]}_1)$; and step (c) follows from (1). The third term in the RHS of (10) can be similarly lower bounded by Term_3. This completes the proof of (9).

Hence, from (8) and (9), we lower bound D as:

$$D + o(L) \geq L + \text{Term}_1 + \text{Term}_2 + \text{Term}_3. \tag{12}$$

Before further lower bounding Term_1, Term_2 and Term_3, we state the symmetric assumption on PIR schemes.

*Remark 1 (**Symmetric Scheme Assumption**):* For symmetric PIR schemes, we assume that the queries, answers, and the storage placement are symmetric across DBs. In particular, for this example, we have the following symmetry condition: for $|\mathcal{K}| = |\mathcal{K}'|$, and $|\mathcal{N}| = |\mathcal{N}'|$, we have

$$\begin{aligned} &I(W_{[1:3]\setminus\mathcal{K}}; Q^{[k]}_n, A^{[k]}_n | W_{\mathcal{K}}, Z_{\mathcal{N}}) \\ &\qquad = I(W_{[1:3]\setminus\mathcal{K}'}; Q^{[k']}_{n'}, A^{[k']}_{n'} | W_{\mathcal{K}'}, Z_{\mathcal{N}'}), \end{aligned} \tag{13}$$

where $\mathcal{N}, \mathcal{N}', \mathcal{K}, \mathcal{K}' \subseteq [1:3]$, $n \in [1:3] \setminus \mathcal{N}$, $n' \in [1:3] \setminus \mathcal{N}'$, $k, k' \in [1:3]$, and $|\mathcal{A}|$ is the cardinality of any set $\mathcal{A}$. For instance, we have $I(W_{[1:2]}; Q^{[3]}_2, A^{[3]}_2 | W_3, Z_2) = I(W_{[1:2]}; Q^{[3]}_3, A^{[3]}_3 | W_3, Z_2)$.

Lemma 1: For symmetric PIR schemes as defined in Remark 1, we can lower bound the three terms in (9) *as follows:*

$$\begin{aligned} \textit{Term}_1 &\geq \frac{4}{9}L + \frac{1}{6}H(W_3|W_{[1:2]}, Z_1) \\ &\quad + \frac{1}{3}H(W_3|W_{[1:2]}, Z_{[1:2]}) + o(L), \\ \textit{Term}_2 &\geq \frac{1}{2}H(W_2|W_1, Z_1) + \frac{1}{4}H(W_3|W_{[1:2]}, Z_1) \\ &\quad + \frac{1}{2}H(W_3|W_{[1:2]}, Z_{[1:2]}) + o(L), \\ \textit{Term}_3 &\geq H(W_2|W_1, Z_{[1:2]}) + H(W_3|W_{[1:2]}, Z_{[1:2]}) + o(L). \end{aligned}$$

The full proof of Lemma 1 is presented in Appendix A.

Remark 2: Due to space limitation, we have specialized the converse proof for symmetric PIR schemes. This assumption can be readily removed by averaging (8) and (9) over all possible permutations of the DBs and the messages indexes.

Using Lemma 1 in (12), we obtain the following bound:

$$\begin{aligned} D \geq \frac{13}{9}L + \frac{1}{2}H(W_2|W_1, Z_1) + \frac{5}{12}H(W_3|W_{[1:2]}, Z_1) \\ + \frac{11}{6}H(W_3|W_{[1:2]}, Z_{[1:2]}) + H(W_2|W_1, Z_{[1:2]}) + o(L). \end{aligned} \tag{14}$$

It is interesting to observe that by trivially lower bounding the entopy terms in the RHS of (14) by zero, we recover the converse result in [1], i.e., $D/L \geq 1 + 1/3 + 1/3^2 = 13/9$. This bound is tight when each DB has the storage capacity to store all messages, i.e., if $\mu = 1$, since $Z_n = (W_1, W_2, W_3)$ for $n \in [1:3]$. However, for storage constrained DBs, these remaining terms contribute to the lower bound and are central to proving optimality. We now specialize the lower bound in (14) for the case of uncoded storage placement as defined next.

*Remark 3 (**Uncoded Storage Assumption**):* For uncoded storage, the DBs store uncoded functions of the K messages. We consider a generic uncoded placement strategy as follows: let us consider message W_k, and denote $W_k^{\mathcal{S}}$ as the set of bits of W_k that are stored at the DBs in the set $\mathcal{S}$. For this example, each message W_k can then be written as $W_k = \{W_k^{\{1\}}, W_k^{\{2\}}, W_k^{\{3\}}, W_k^{\{1,2\}}, W_k^{\{1,3\}}, W_k^{\{2,3\}}, W_k^{\{1,2,3\}}\}$.

For symmetric uncoded schemes (as defined by Remarks 1 and 3), we denote the sizes of the sub-messages of the mesage W_k, $\forall k \in [1:3]$, as follows:

$$\begin{aligned} &H(W_k^{\{1\}}) = H(W_k^{\{2\}}) = H(W_k^{\{3\}}) = x_1 L, \\ &H(W_k^{\{1,2\}}) = H(W_k^{\{1,3\}}) = H(W_k^{\{2,3\}}) = x_2 L, \\ &H(W_k^{\{1,2,3\}}) = x_3 L, \end{aligned} \tag{15}$$

where $x_1, x_2, x_3 \geq 0$. Since each message is of size L bits, the variables (x_1, x_2, x_3) must satisfy $\binom{3}{1}x_1 L + \binom{3}{2}x_2 L + \binom{3}{3}x_3 L = L$, which gives us the following constraint C_1:

$$\text{(Message Size Constraint: } C_1) \quad 3x_1 + 3x_2 + x_3 = 1. \tag{16}$$

Furthermore, each DB must satisfy the storage constraint, i.e., total data stored at each DB cannot exceed $\mu K L = 3\mu L$ bits. For instance, DB_1 can store the sub-messages $W_k^{\{1\}}, W_k^{\{1,2\}}, W_k^{\{1,3\}}, W_k^{\{1,2,3\}}$, $\forall k \in [1:3]$, and hence we must satisfy $(3 \times (\binom{2}{0}x_1 L + \binom{2}{1}x_2 L + \binom{2}{2}x_3 L) = 3\mu L)$, i.e.,

$$\text{(Storage Constraint: } C_2) \quad x_1 + 2x_2 + x_3 = \mu. \tag{17}$$

For uncoded storage, we express the terms appearing in (14) as functions of (x_1, x_2, x_3), defined in (15), as follows:

$$\begin{aligned} &H(W_2|W_1, Z_1) = H(W_2^{\{2\}}, W_2^{\{3\}}, W_2^{\{2,3\}}) = (2x_1 + x_2)L, \\ &H(W_3|W_{[1:2]}, Z_1) = H(W_3^{\{2\}}, W_3^{\{3\}}, W_3^{\{2,3\}}) = (2x_1 + x_2)L, \\ &H(W_2|W_1, Z_{[1:2]}) = H(W_2^{\{3\}}) = x_1 L, \\ &H(W_3|W_{[1:2]}, Z_{[1:2]}) = H(W_3^{\{3\}}) = x_1 L. \end{aligned} \tag{18}$$

Using (18) in (14) and taking the limit $L \to \infty$, we arrive at

$$D^*(\mu) \geq \frac{D}{L} \geq \frac{13}{9} + \frac{14}{3}x_1 + \frac{11}{12}x_2. \tag{19}$$

From constraints C_1 and C_2, we first express x_1 and x_2 in terms of x_3, and bound (19) as

$$\begin{aligned} D^*(\mu) &\geq \frac{13}{9} + \frac{14}{3}\left(\frac{2}{3} - \mu + \frac{x_3}{3}\right) + \frac{11}{12}\left(\mu - \frac{1}{3} - \frac{2x_3}{3}\right) \\ &= \frac{17}{4} - \frac{15\mu}{4} + \frac{17x_3}{18} \overset{(a)}{\geq} \frac{17 - 15\mu}{4}, \end{aligned} \tag{20}$$

where (a) follows since $x_3 \geq 0$. This proves the first bound in (7). Next, we express x_2 in (19), in terms of x_1 to obtain

$$\begin{aligned} D^*(\mu) &\geq \frac{13}{9} + \frac{14x_1}{3} + \frac{11}{12}(1 - \mu) - \frac{11x_1}{6} \\ &= \frac{85}{36} - \frac{11\mu}{12} + \frac{17x_1}{6} \overset{(a)}{\geq} \frac{85 - 33\mu}{36}, \end{aligned} \tag{21}$$

where (a) follows since $x_1 \geq 0$. This proves the second bound in (7), completing the proof of Theorem 1 for $N = K = 3$.

IV. Conclusions

In this paper, we characterized the optimum download cost of PIR from uncoded storage constrained DBs. In particular, for any (N, K), we show that the optimal tradeoff between storage, μ, and the download cost, $D(\mu)$, is given by the lower convex hull of the pairs $(\frac{t}{N}, \left(1 + \frac{1}{t} + \frac{1}{t^2} + \cdots + \frac{1}{t^{K-1}}\right))$ for $t \in [1:N]$. The main technical contribution of this paper is obtaining lower bounds on the download cost for PIR as a function of storage, which matches the achievable scheme in [12], and hence characterizes the optimal tradeoff.

References

[1] H. Sun and S. A. Jafar, "The capacity of private information retrieval," *IEEE Transactions on Information Theory*, vol. 63, no. 7, pp. 4075–4088, 2017.

[2] ——, "The capacity of robust private information retrieval with colluding databases," *IEEE Transactions on Information Theory*, vol. 64, no. 4, pp. 2361–2370, 2018.

[3] K. Banawan and S. Ulukus, "The capacity of private information retrieval from byzantine and colluding databases," *CoRR*, vol. abs/1706.01442, 2017.

[4] R. Tajeddine, O. W. Gnilke, and S. El Rouayheb, "Private information retrieval from mds coded data in distributed storage systems," *IEEE Transactions on Information Theory*, 2018.

[5] K. Banawan and S. Ulukusr, "The capacity of private information retrieval from coded databases," *IEEE Transactions on Information Theory*, vol. 64, no. 3, pp. 1945–1956, 2018.

[6] H. Sun and S. A. Jafar, "The capacity of symmetric private information retrieval," in *2016 IEEE Globecom Workshops*, 2016, pp. 1–5.

[7] R. Tandon, "The capacity of cache aided private information," in *Proceedings of the IEEE 55th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, 2017, pp. 878–885.

[8] Z. Chen, Z. Wang, and S. Jafar, "The capacity of private information retrieval with private side information," *CoRR*, vol. abs/1709.03022, 2017.

[9] Y.-P. Wei, K. Banawan, and S. Ulukus, "Fundamental limits of cache-aided private information retrieval with unknown and uncoded prefetching," *CoRR*, vol. abs/1709.01056, 2017.

[10] S. Kadhe, B. Garcia, A. Heidarzadeh, S. E. Rouayheb, and A. Sprintson, "Private information retrieval with side information," *CoRR*, vol. abs/1709.00112, 2017.

[11] H. Sun and S. Jafar, "Multiround private information retrieval: Capacity and storage overhead," *IEEE Transactions on Information Theory*, 2018.

[12] M. Abdul-Wahid, F. Almoualem, D. Kumar, and R. Tandon, "Private information retrieval from storage constrained databases - coded caching meets PIR," *CoRR*, vol. abs/1711.05244, 2017.

[13] M. A. Attia, D. Kumar, and R. Tandon, "The capacity of private information retrieval from uncoded storage constrained databases," 2018. [Online]. Available: https://www.dropbox.com/s/4m89in7fgeypeyv/paper-dropbox.pdf?dl=0

Appendix A
Proof of Lemma 1

We lower bound Term_1 as follows:

$$\begin{aligned} \mathrm{Term}_1 &= I(W_{[2:3]}; Q_1^{[1]}, A_1^{[1]}|W_1) \\ &\overset{(a)}{=} \frac{1}{3}\big[I(W_{[2:3]}; Q_1^{[1]}, A_1^{[1]}|W_1) + I(W_{[2:3]}; Q_2^{[1]}, A_2^{[1]}|W_1) \\ &\qquad + I(W_{[2:3]}; Q_3^{[1]}, A_3^{[1]}|W_1)\big] \\ &\overset{(b)}{=} \frac{1}{3}[I(W_{[2:3]}; Q_1^{[2]}, A_1^{[2]}|W_1) + I(W_{[2:3]}; Q_2^{[2]}, A_2^{[2]}|W_1) \\ &\qquad + I(W_{[2:3]}; Q_3^{[2]}, A_3^{[2]}|W_1)] \\ &\geq \frac{1}{3}[I(W_{[2:3]}; A_1^{[2]}|W_1, Q_1^{[2]}) + I(W_{[2:3]}; A_2^{[2]}|W_1, Q_2^{[2]}) \\ &\qquad + I(W_{[2:3]}; A_3^{[2]}|W_1, Q_3^{[2]})] \end{aligned}$$

$$\overset{(c)}{=}\frac{1}{3}[H(A_1^{[2]}|W_1,Q_1^{[2]})+H(A_2^{[2]}|W_1,Q_2^{[2]})+H(A_3^{[2]}|W_1,Q_3^{[2]})]$$
$$\overset{(d)}{\geq}\frac{1}{3}[H(A_1^{[2]}|W_1,Q_{[1:3]}^{[2]})+H(A_2^{[2]}|W_1,Q_{[1:3]}^{[2]},A_1^{[2]})$$
$$+H(A_3^{[2]}|W_1,Q_{[1:3]}^{[2]},A_{[1:2]}^{[2]})] \;=\; \frac{1}{3}[H(A_{[1:3]}^{[2]}|W_1,Q_{[1:3]}^{[2]})]$$
$$\overset{(e)}{=}\frac{1}{3}I(W_{[2:3]};A_{[1:3]}^{[2]}|W_1,Q_{[1:3]}^{[2]})$$
$$\overset{(f)}{=}\frac{1}{3}I(W_{[2:3]};Q_{[1:3]}^{[2]},A_{[1:3]}^{[2]}|W_1)$$
$$\overset{(g)}{=}\frac{1}{3}I(W_{[2:3]};W_2,Q_{[1:3]}^{[2]},A_{[1:3]}^{[2]}|W_1)+o(L)$$
$$=\frac{1}{3}[I(W_{[2:3]};W_2|W_1)+I(W_3;Q_{[1:3]}^{[2]},A_{[1:3]}^{[2]}|W_{[1:2]})]+o(L)$$
$$=\frac{1}{3}[L+\underbrace{I(W_3;Q_{[1:3]}^{[2]},A_{[1:3]}^{[2]}|W_{[1:2]})}_{\triangleq \mathrm{T}_1}]+o(L), \tag{22}$$

where (a) follows from the symmetry assumption in Remark 1; (b) follows from privacy constraint (4); (c) follows from (2), i.e., the fact that every answer is a function of query and all messages; (d) follows from the fact that conditioning reduces entropy; (e) follows from (2); (f) follows from (1); and (g) follows from correctness constraint in (3), i.e., W_2 must be decoded from $Q_{[1:3]}^{[2]}$, and $A_{[1:3]}^{[2]}$. Next, the term T_1 in (22) is lower bounded using similar steps to (9) as

$$\mathrm{T}_1 \geq \underbrace{I(W_3;Q_1^{[2]},A_1^{[2]}|W_{[1:2]})}_{\triangleq \mathrm{T}_{(1a)}}+\underbrace{I(W_3;Q_2^{[2]},A_2^{[2]}|W_{[1:2]},Z_1)}_{\triangleq \mathrm{T}_{(1b)}}$$
$$+\underbrace{I(W_3;Q_3^{[2]},A_3^{[2]}|W_{[1:2]},Z_{[1:2]})}_{\triangleq \mathrm{T}_{(1c)}}. \tag{23}$$

$T_{(1a)}$ is similar to Term_1, and hence, we can follow a similar sequence of steps used in (22) to lower bound $T_{(1a)}$ as follows:

$$\mathrm{T}_{(1a)} \overset{(a)}{\geq} \frac{1}{3}I(W_3;W_3,Q_{[1:3]}^{[3]},A_{[1:3]}^{[3]}|W_{[1:2]})+o(L)$$
$$=\frac{1}{3}\left[I(W_3;W_3|W_{[1:2]})+I(W_3;Q_{[1:3]}^{[3]},A_{[1:3]}^{[3]}|W_{[1:3]})\right]+o(L)$$
$$=L/3+o(L), \tag{24}$$

where (a) follows by arguments similar to $(a)\rightarrow(g)$ in (22). We note that $T_{(1b)}$ and $T_{(1c)}$ were trivially lower bounded by zero in the converse proof for replicated DBs in [1]. Carefully lower bounding these terms as a function of the content stored across DBs is one of the key new aspects of the converse proof. We proceed to lower bound $T_{(1b)}$ defined in (23) as follows:

$$\mathrm{T}_{(1b)}=I(W_3;Q_2^{[2]},A_2^{[2]}|W_{[1:2]},Z_1)$$
$$\overset{(a)}{=}\frac{1}{2}[I(W_3;Q_2^{[2]},A_2^{[2]}|W_{[1:2]},Z_1)+I(W_3;Q_3^{[2]},A_3^{[2]}|W_{[1:2]},Z_1)]$$
$$\overset{(b)}{=}\frac{1}{2}[I(W_3;Q_2^{[3]},A_2^{[3]}|W_{[1:2]},Z_1)+I(W_3;Q_3^{[3]},A_3^{[3]}|W_{[1:2]},Z_1)]$$
$$\overset{(c)}{=}\frac{1}{2}[I(W_3;A_2^{[3]}|W_{[1:2]},Z_1,Q_2^{[3]})+I(W_3;A_3^{[3]}|W_{[1:2]},Z_1,Q_3^{[3]})]$$
$$\overset{(d)}{\geq}\frac{1}{2}I(W_3;Q_{[1:3]}^{[3]},A_{[2:3]}^{[3]}|W_{[1:2]},Z_1)$$
$$\overset{(e)}{=}\frac{1}{2}I(W_3;W_3,Q_{[1:3]}^{[3]},A_{[2:3]}^{[3]}|W_{[1:2]},Z_1)+o(L)$$
$$=\frac{1}{2}H(W_3|W_{[1:2]},Z_1)+o(L), \tag{25}$$

where (a) follows from the symmetry assumption in Remark 1; (b) follows from (4); (c) follows from (1); (d) follows by arguments similar to $(c)\rightarrow(e)$ in (22), and (e) follows from (3). We next bound term $T_{(1c)}$ defined in (23) as follows:

$$\mathrm{T}_{(1c)}=I(W_3;Q_3^{[2]},A_3^{[2]}|W_{[1:2]},Z_{[1:2]})$$
$$\overset{(a)}{=}I(W_3;Q_3^{[3]},A_3^{[3]}|W_{[1:2]},Z_{[1:2]})$$
$$\overset{(b)}{=}I(W_3;A_3^{[3]}|W_{[1:2]},Z_{[1:2]},Q_3^{[3]})$$
$$\overset{(c)}{\geq}I(W_3;Q_{[1:3]}^{[3]},A_3^{[3]}|W_{[1:2]},Z_{[1:2]})$$
$$\overset{(d)}{=}I(W_3;W_3,Q_{[1:3]}^{[3]},A_3^{[3]}|W_{[1:2]},Z_{[1:2]})+o(L)$$
$$=H(W_3|W_{[1:2]},Z_{[1:2]})+o(L), \tag{26}$$

where (a) follows from (4); (b) follows from (1); (c) follows by arguments similar to $(c)\rightarrow(e)$ in (22), and (d) follows from (3). We lower bound T_1 from (24), (25) and (26), and then use it in (22) to arrive at the desired bound on Term_1:

$$\mathrm{Term}_1 \geq \frac{4}{9}L+\frac{1}{6}H(W_3|W_{[1:2]},Z_1)$$
$$+\frac{1}{3}H(W_3|W_{[1:2]},Z_{[1:2]})+o(L).$$

We now present the steps to prove the lower bound on Term_2. This term is similar to term $\mathrm{T}_{(1b)}$, and hence, we can follow a similar sequence of steps to arrive at:

$$\mathrm{Term}_2=I(W_{[2:3]};Q_2^{[1]},A_2^{[1]}|W_1,Z_1)$$
$$\overset{(a)}{\geq}\frac{1}{2}I(W_{[2:3]};W_2,Q_{[1:3]}^{[2]},A_{[2:3]}^{[2]}|W_1,Z_1)+o(L)$$
$$\geq\frac{1}{2}[H(W_2|W_1,Z_1)+I(W_3;Q_{[2:3]}^{[2]},A_{[2:3]}^{[2]}|W_{[1:2]},Z_1)]+o(L)$$
$$\overset{(b)}{\geq}\frac{1}{2}[H(W_2|W_1,Z_1)+I(W_3;Q_2^{[2]},A_2^{[2]}|W_{[1:2]},Z_1)$$
$$+I(W_3;Q_3^{[2]},A_3^{[2]}|W_{[1:2]},Z_{[1:2]})]+o(L)$$
$$\overset{(c)}{\geq}\frac{1}{2}H(W_2|W_1,Z_1)+\frac{1}{4}H(W_3|W_{[1:2]},Z_1)$$
$$+\frac{1}{2}H(W_3|W_{[1:2]},Z_{[1:2]})+o(L), \tag{27}$$

where (a) follows by arguments similar to $(a)\rightarrow(e)$ in (25), (b) follows from arguments similar to (9), and (c) follows from the bounds in (25) and (26). This gives the desired bound on Term_2. Finally, Term_3 can be lower bounded as follows:

$$\mathrm{Term}_3=I(W_{[2:3]};Q_3^{[1]},A_3^{[1]}|W_1,Z_{[1:2]})$$
$$\overset{(a)}{\geq}I(W_{[2:3]};W_2,Q_{[1:3]}^{[2]},A_3^{[2]}|W_1,Z_{[1:2]})+o(L)$$
$$=H(W_2|W_1,Z_{[1:2]})+I(W_3;Q_3^{[2]},A_3^{[2]}|W_{[1:2]},Z_{[1:2]})+o(L)$$
$$\overset{(b)}{\geq}H(W_2|W_1,Z_{[1:2]})+H(W_3|W_{[1:2]},Z_{[1:2]})+o(L), \tag{28}$$

where (a) follows by arguments similar to $(a)\rightarrow(d)$ in (26), and (b) follows from the bound in (26).