

What Storage Access Privacy is Achievable with Small Overhead?

Sarvar Patel
Google LLC
sarvar@google.com

Giuseppe Persiano
Università di Salerno
Google LLC
giuper@gmail.com

Kevin Yeo
Google LLC
kwlyeo@google.com

ABSTRACT

Oblivious RAM (ORAM) and private information retrieval (PIR) are classic cryptographic primitives used to hide the access pattern to data whose storage has been outsourced to an untrusted server. Unfortunately, both primitives require considerable overhead compared to plaintext access. For large-scale storage infrastructure with highly frequent access requests, the degradation in response time and the exorbitant increase in resource costs incurred by either ORAM or PIR prevent their usage. In an ideal scenario, a privacy-preserving storage protocols with small overhead would be implemented for these heavily trafficked storage systems to avoid negatively impacting either performance and/or costs. In this work, we study the problem of the best *storage access privacy* that is achievable with only *small overhead* over plaintext access.

To answer this question, we consider *differential privacy access* which is a generalization of the *oblivious access* security notion that are considered by ORAM and PIR. Quite surprisingly, we present strong evidence that constant overhead storage schemes may only be achieved with privacy budgets of $\epsilon = \Omega(\log n)$. We present asymptotically optimal constructions for differentially private variants of both ORAM and PIR with privacy budgets $\epsilon = \Theta(\log n)$ with only $O(1)$ overhead. In addition, we consider a more complex storage primitive called key-value storage in which data is indexed by keys from a large universe (as opposed to consecutive integers in ORAM and PIR). We present a differentially private key-value storage scheme with $\epsilon = \Theta(\log n)$ and $O(\log \log n)$ overhead. This construction uses a new oblivious, two-choice hashing scheme that may be of independent interest.

ACM Reference Format:

Sarvar Patel, Giuseppe Persiano, and Kevin Yeo. 2019. What Storage Access Privacy is Achievable with Small Overhead?. In *38th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems (PODS'19)*, June 30–July 5, 2019, Amsterdam, Netherlands. ACM, New York, NY, USA, 18 pages. <https://doi.org/10.1145/3294052.3319695>

1 INTRODUCTION

Privacy-preserving storage primitives consider the outsourcing of the storage of a database to an untrusted server with the ability for clients to retrieve database records while maintaining the privacy of the retrievals. Even though encryption can be used to hide the content of the database, patterns of how the encrypted data is accessed are revealed. The leakage of access patterns has been shown to compromise privacy in many important practical settings [14, 32]. Privacy-preserving storage primitives that guarantee retrieval privacy have been used as a critical component in many systems such as advertisement [30], discovery of identities [8] and publish-subscribe [18]. Therefore, a very important question involves the construction of privacy-preserving storage schemes guaranteeing retrieval privacy while ensuring that record retrieval can be performed efficiently.

A common way to formulate the privacy of retrievals is *obliviousness*. Obliviousness guarantees that for any two fixed sequences of record retrievals of the same length, any adversary that views all accesses to stored data cannot determine which of the two sequences induced the resulting access pattern to stored data. Obliviousness has been considered with both statistical security, providing protection from adversaries with unbounded computational resources, and computational security, where the adversary is assumed to be probabilistic polynomial time (PPT). Oblivious RAM (ORAM) and private information retrieval (PIR) are two oblivious storage primitives that have been the objective of extensive research [12, 19, 25–28, 35, 36, 44, 48]. For ORAM, it has been shown that $\Omega(\log n)$ overhead is necessary [37]. On the other hand, the best constructions for PIR require at least $\Omega(n)$ server computation over the entire outsourced database.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

PODS'19, June 30–July 5, 2019, Amsterdam, Netherlands

© 2019 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-6227-6/19/06.

<https://doi.org/10.1145/3294052.3319695>

For integral storage infrastructure that handle many access requests per second, the increased response time and server costs caused by the usage of ORAM and PIR prevent their implementation in these important systems. In a perfect world, there would exist a storage scheme with strong privacy with only small overhead that could be implemented for these frequently accessed storage infrastructures without negatively affecting the performance and/or expenses. In our work, we address the natural question of the best privacy that may be achieved by any storage scheme with only small overhead compared to plaintext access. To our knowledge, our work is the first that considers the question of maximizing privacy for a specific efficiency goal. Previous works consider minimizing the efficiency for primitives with a specific privacy notion.

To consider this problem, we use another formulation for storage access privacy through *differential privacy* [22–24]. Typically, differential privacy is used in the context of *privacy-preserving data analysis* where global properties of the entire database are disclosed while maintaining the privacy of individual database records. In particular, differential privacy guarantees that any fixed disclosure is just as likely, usually within a multiplicative factor, regardless of whether an individual belonged to the sample population or not. Our work focuses on the notion of *differentially private access* which attempts to maintain privacy for individual record retrievals, but may reveal information about the entire sequence of retrievals. Roughly speaking, differential privacy guarantees that any sequence of accesses to stored data caused by the execution of a sequence of record retrievals will be just as likely, except for a multiplicative factor, caused by another sequence attained by replacing a single retrieval for a database record with a single retrieval for any other database record. Differentially private access is a generalization of oblivious access as oblivious access is achieved by fixing the differential privacy parameters to be $\epsilon = 0$ and $\delta = \text{negl}(n)$.

With this definition, readers might now ask the following two questions. When does differentially private access make sense as a security notion? Or, might there exist a stronger security notion that is achievable with small overhead?

To answer the first question, we revisit the scenario of privacy-preserving data analysis on outsourced databases. In these scenarios, there is no sense in using storage schemes providing obliviousness to hide entire record retrieval sequences when differentially private disclosures only maintain privacy for single records. For example, suppose we wish to disclose a differentially private model trained over a sample from the database. Obliviousness would unnecessarily hide the identity of the entire retrieved sample at a high cost yet the differential privacy would guarantee the privacy

about individuals in the sample. Therefore, differentially private access is the privacy notion that is complementary to differential privacy disclosures on outsourced databases. In general, differentially private access guarantees privacy for individual retrievals. For the second question, recent work by Persiano and Yeo [47] show that $\Omega(\log n)$ overhead is necessary for differentially private RAMs with parameters of $\epsilon = O(1)$ and $\delta \leq 1/3$. Therefore, the task of finding good parameters for differential privacy that provide both meaningful privacy as well as while permitting small overhead seems non-trivial.

We study differentially private variants of both PIR and ORAM. PIR enables clients to obviously retrieve database records outsourced to a server. In PIR, both the client and the server are stateless which means that no information may be maintained between multiple record retrievals beyond the server storing the database. Since both the client and the server are stateless, PIR requires the server to perform an operation on each database record: if a record is not involved in computing the server’s reply, then it cannot be the record retrieved by the client. The majority of PIR constructions use homomorphic encryption [2, 4] or other expensive encryptions with useful properties [12, 25, 36]. Some recent works consider PIR with stateful clients [46] or super-linear server storage [10, 13, 31]. All single-server PIR schemes consider computational security. PIR has also been studied in the multiple, non-colluding server setting where constructions only require servers to perform computation sublinear in the number of database records and provide statistical security [7, 19].

On the other hand, ORAM allows both the client and server to be stateful and maintain information between multiple queries and allows the client to perform both record retrievals and overwrites. Additionally, ORAM allows a setup phase where the untrusted server receives an encrypted version of the database to store and the client is given a secret key. As a result of state, the server is no longer required to perform computation on every database record. With this efficiency, the majority of ORAM schemes consider only simple upload and download operations and forego the use of expensive encryption schemes of PIR that require the server to perform untrusted computation, although some previous works consider ORAM with homomorphic encryption [20]. For a database with n records, it has been shown that ORAM requires overhead of $\Omega(\log n)$ records [27, 37] and that $O(\log n \cdot \log \log n)$ communication suffices [44]. We also examine an extension to ORAM, which we denote as *oblivious key-value storage* (previously also denoted as oblivious storage), where database records are uniquely identified by keys from a large universe and clients might also attempt to retrieve a non-existent key [29].

In our work, we study the differentially private variants of these three primitives, which we denote as *differentially private information retrieval* (DP-IR), *differentially private RAM* (DP-RAM) and *differentially private key value storage* (DP-KVS). We focus on the question the best privacy that can be achieved by each of these primitives that only require small overhead over plaintext access.

Our Contributions. We present both negative and positive results for both DP-IR and DP-RAM as well as other natural variants of these primitives.

Our lower bounds for DP-IR and DP-RAM apply for a wide range of privacy budgets when the database is stored in a natural encoding. In particular, we consider the balls and bins model of storage (previously considered in [11, 15, 27]) where each database record is considered as an opaque ball along with a key containing important metadata. While not covering all database encodings, in our opinion, the balls and bins models encompasses all natural database representations that maximize practical utility and efficiency. In particular, the contents of each record are assumed to be placed together, which is typically done to maximize data locality.

For positive results, we show that a class of simple constructions are optimal for large sets of privacy budgets. These schemes may be viewed as inserting noise into the sequence of record retrievals and/or overwrites. In particular, the real record retrievals and/or overwrites are disguised within a set that also contain fake record retrievals and/or overwrites. These schemes suffice to construct the best privacy-preserving storage protocols with very small overhead for DP-IR, DP-RAM and DP-KVS. While our constructions are simple, we draw attention to the fact that designing differentially private storage systems is delicate even with weak security notions. Some simple constructions are very appealing and, at first, seem to match our lower bounds. However, many variants of simple constructions (including our constructions) end up being completely insecure. As an example, in Section 4, we consider a simple and tempting DP-IR construction and show that it only guarantees differential privacy with $\delta \rightarrow 1$, i.e., no privacy at all. Additionally, while our constructions are simple, the security proofs end up being quite involved (especially for DP-RAM). Furthermore, to handle the additional functionalities of DP-KVS, a more complex algorithm using a novel, oblivious adaptation of the two-choice hashing scheme [41] is required.

DP-IR Results. To our knowledge, previous works on DP-IR consider only the multiple, non-colluding server scenarios [49]. Our work is the first to consider DP-IR in the single server scenario.

We show that for any ϵ -DP-IR that must always output the right answer and any value of $\epsilon \geq 0$, then the server must

operate on all stored n records. This result is very strong and somewhat surprising as there does not exist any weakening of privacy that will improve the server computational costs.

THEOREM 1.1 (INFORMAL). *For any $\epsilon, \delta \geq 0$, any (ϵ, δ) -DP-IR scheme in the balls and bins model must operate on $(1 - \delta)n$ records.*

On the other hand, we show that this strong negative result may be circumvented by considering DP-IR schemes with non-zero error probabilities $0 < \alpha \leq 1$. Here, the error probability is over the internal randomness of the DP-IR scheme and does not depend on the queries and/or stored data. For this case, we present the following weaker lower bound.

THEOREM 1.2 (INFORMAL). *For any $\epsilon, \delta \geq 0$, any (ϵ, δ) -DP-IR scheme in the balls and bins model with error probability $\alpha > 0$ must operate on*

$$\Omega\left(\frac{(1 - \alpha - \delta) \cdot n}{e^\epsilon}\right)$$

records.

As we are focusing on schemes with very small overhead, the above theorem leads credence that there might exist a ϵ -DP-IR scheme with some small, constant error probability $\alpha > 0$ that only performs $O(1)$ operations when $\epsilon = \Theta(\log n)$. We show that there exists a simple construction with these properties.

THEOREM 1.3 (INFORMAL). *There exists an ϵ -DP-IR with $\epsilon = \Theta(\log n)$ and constant error probability $\alpha > 0$ that only operates on $O(1)$ records.*

Due to our lower bounds, the above construction seems to be the best privacy that can be achieved by any DP-IR with only constant overhead compared to plaintext storage access.

We also consider natural extensions of our DP-IR with multiple, non-colluding servers. There exist several multiple-server PIR schemes [7, 19] in literature. In our full paper, we present asymptotically tight lower bounds for the constructions in [49].

DP-RAM Results. For DP-RAM, we once again start with describing our lower bound results. Unlike DP-IR, there is no separation between the best lower bound for perfectly correct DP-RAM and DP-RAM with error probability $\alpha > 0$. We now present our DP-RAM lower bound which applies for all values $0 \leq \alpha \leq 1$.

THEOREM 1.4 (INFORMAL). *For any $\epsilon \geq 0$, any ϵ -DP-RAM with error probability $\alpha \geq 0$ in the balls and bins model and a client that stores at most c blocks must operate on*

$$\Omega\left(\log_c\left(\frac{(1 - \alpha) \cdot n}{e^\epsilon}\right)\right)$$

records.

The above theorem essentially states that there are two ways that one can achieve very efficient ϵ -DP-RAM constructions: either increase the amount of client storage (c) or increase the privacy budget (ϵ). In most scenarios, it is impractical to suppose that the client can store large portions of data. As our desire is to construct schemes that should be easily usable in complex systems, we minimize the requirements of our clients by assuming clients have very small amounts of storage. In [47], an $\Omega(\log(n/c))$ lower bound is given for constant ϵ and $\delta \leq 1/3$. However, their lower bound can be generalized to $\Omega(\log(n/c)/e^\epsilon)$ for any $\epsilon \geq 0$. Their lower bound has an exponentially worse dependence on ϵ compared to our lower bound. For example, the lower bound in [47] does not preclude the existence of a ϵ -DP-RAM with $\epsilon = \Theta(\log \log n)$ and constant overhead. On the other hand, our lower bound improves the privacy budget lower bound showing that an ϵ -DP-RAM that operates on $O(1)$ records must have $\epsilon = \Omega(\log n)$. We show the existence of a constant overhead ϵ -DP-RAM with asymptotically optimal $\epsilon = \Theta(\log n)$ privacy.

THEOREM 1.5 (INFORMAL). *There exists an ϵ -DP-RAM with $\epsilon = \Theta(\log n)$ that only operates on $O(1)$ records.*

Once again, this construction seems to be the best privacy that can be achieved by any DP-RAM scheme with only $O(1)$ overhead over the baseline, unprotected storage access due to our lower bounds. Our scheme improves on previous DP-RAM schemes in [50] which starts from Path ORAM [48] and degrades security to improve efficiency. For their scheme to achieve even client storage of $O(\sqrt{n})$, their construction recursively stores position maps which costs both logarithmic overhead and client-to-server roundtrips. On the other hand, our DP-RAM construction uses both $O(1)$ overhead and roundtrips while achieving small client storage.

DP-KVS Results. Finally, we consider DP-KVS which is an extension of DP-RAM. As a result, all DP-RAM lower bounds also apply to DP-KVS. Therefore, the best construction that is achievable by DP-KVS with $O(1)$ overhead could be $\epsilon = \Theta(\log n)$. Due to the difficulties of handling a larger universe of queries and possibility that clients request on-existent keys, we present a construction with slightly worse than constant overhead. In addition, we can only achieve privacy in the approximate differential privacy framework.

THEOREM 1.6 (INFORMAL). *There exists a ϵ -DP-KVS with $\epsilon = \Theta(\log n)$ that operates on $O(\log \log n)$ records.*

While the scheme has non-constant overhead, the $O(\log \log n)$ overhead is exponentially smaller than the best oblivious key-value storage schemes based on ORAMs.

To construct our DP-KVS scheme, we present an improved, oblivious variant of the power of two choices hashing scheme [41] that may be of separate, independent interest. Traditional power of two choices hashing guarantees that bins do not exceed $O(\log \log n)$ items except with probability negligible in n . Our desired oblivious variant must hide the sizes of bins. One way to hide bin sizes is to pad all bins with dummy items up to a maximum. This technique requires padding all n bins to $O(\log \log n)$ items meaning $O(n \log \log n)$ storage. By using a tree-like structure to allow bins to share storage, we present a new scheme using $O(n)$ storage.

Related Work. DP-RAM was considered previously in [50] which present a construction based on Path ORAM [48]. However, their scheme requires recursively stored position maps which requires $\Theta(\log n)$ client-to-server roundtrips to get client storage of even $O(\sqrt{n})$. We present DP-RAM schemes that only require $O(1)$ overhead with small client storage. A construction of DP-IR in the multiple, non-colluding server scenario was considered in [49], which we show is optimal for certain parameters. Lower bounds for DP-RAM have been considered in [47], which are stronger and weaker in different dimensions. Their lower bound has an exponentially worse dependence on the privacy budget, ϵ . On the other hand, their lower bound applies to general storage encodings whereas our work only apply to the balls and bins model. A variant of DP-RAM that only maintains privacy for database record insertions is considered in [15]. A privacy notion stronger than obliviousness is considered in [34] where accesses are protected using obliviousness while the number of accesses is protected using differential privacy. Obliviousness has been considered for other problems such as sorting [3, 6] and shuffling [43, 45, 51]. The problem of relaxing the security notion of obliviousness to differential privacy has also been studied in the context of multi-party protocols. A simple multi-party protocol for single-bit inputs is shown in [33] that maximizes the accuracy for any privacy budget. The multi-party computation problem for larger input sizes is studied in [39]. Additionally, the problem of differentially private disclosure of specific analysis such as subspace clustering [52], deep learning [1] and many others have also been studied where the database is assumed to be stored on a trusted server.

2 DEFINITIONS

We suppose the database D contains n records denoted B_1, \dots, B_n . We will interchangeably use the terms records and blocks. We will refer to a query as a single operation involving either a record retrieval or overwrite and a query sequence as a list consisting of record retrievals and/or overwrites. We use $[n]$ to refer to the set $\{1, \dots, n\}$. We refer to \mathcal{Q} as the space of all possible queries and let $Q_1, Q_2 \in \mathcal{Q}^l$ be two query sequences

of length l . We define the Hamming distance between Q_1 and Q_2 , which we denote by $d(Q_1, Q_2)$, as the number of queries where Q_1 and Q_2 differ.

2.1 Storage Primitives

The *information retrieval* (IR) primitive stores a database D of n equal sized records where only record retrievals are allowed. The initialization of IR consists of the server receiving $D = (B_1, \dots, B_n)$ and simply processing and storing D . The server is only allowed to store the database and may not keep any other information between multiple queries. Similarly, the client is stateless and may not use any storage between multiple queries. A query to IR is described using an integer $q \in [n]$, which is interpreted as retrieving record B_q . We denote two query sequences $Q_1, Q_2 \in [n]^l$ as adjacent if their Hamming distance is exactly 1. That is, Q_1 and Q_2 retrieve a different record at exactly one query.

The *random access memory* (RAM) primitive will store a database D of n equal sized records where both record retrievals and overwrites are permitted. The initialization of RAM consists of a setup phase consisting of a protocol run between the client and the server. The client will receive a private key, process D using the private key and send the processed database to the server to store. Both the client and server may be stateful and maintain information between multiple queries. A query to RAM is a pair $q = (i, \text{op})$ where $i \in [n]$ refers to record B_i and $\text{op} \in \{\text{read}, \text{write}\}$ describes whether the query is a retrieval or overwrite. Two query sequences $Q_1, Q_2 \in ([n] \times \{\text{read}, \text{write}\})^l$ are adjacent if their Hamming distance is exactly 1. That is, at exactly one query, Q_1 and Q_2 operate on a different record and/or perform a different operation.

Finally, the *key-value storage* (KVS) primitive is an extension to RAM. Queries to KVS consist of a pair $q = (k, \text{op})$ where $k \in U$ is the key and U is universe of all keys and $\text{op} \in \{\text{read}, \text{write}\}$ refers to whether the query is a retrieval or overwrite. Unlike RAM, the universe of keys is very large and might be exponentially larger than the number of operations that will be performed. Furthermore, a retrieval operation $q = (k, \text{read})$ may request a key k that has never been previously inserted into the storage protocol. In this case, the KVS protocol should output \perp . Identical to RAM, two query sequences $Q_1, Q_2 \in (U \times \{\text{read}, \text{write}\})^l$ are adjacent if there exists exactly one query that operates on a different key and/or performs a different operation.

We note that IR, RAM and KVS are the most studied storage primitives to provide oblivious access. There are many other extensions to these primitives that have also been studied [9, 10, 13, 16, 17, 29, 31, 38].

2.2 Differentially Private Access to Data

Our privacy notion for a storage primitive \mathcal{S} storing a database $D = (B_1, \dots, B_n)$ with a query space \mathcal{Q} will consider the random variable of the *view* of the adversarial server for a query sequence $Q \in \mathcal{Q}^l$ of length l . With a slight abuse of notation, we refer to the transcript, $\mathcal{S}(Q)$, as the random variable of the adversary's view on query sequence Q . The transcript contains all movement of records performed by the server, as well as all, possibly encrypted, records that have been uploaded and downloaded and the initial database.

Definition 2.1 (Differentially Private Access). Let \mathcal{S} be a storage primitive with query space \mathcal{Q} . \mathcal{S} provides (ϵ, δ) -differentially private access if for all pairs of query sequences $Q_1, Q_2 \in \mathcal{Q}^l$ of length l that are adjacent, that is $d(Q_1, Q_2) = 1$, and for any subset S of the set of possible views of the adversary, then the following holds:

$$\Pr[\mathcal{S}(Q_1) \in S] \leq e^\epsilon \cdot \Pr[\mathcal{S}(Q_2) \in S] + \delta.$$

The ϵ parameter is referred to as the *privacy budget*. When $\delta = 0$, the above definition is referred to as *pure differential privacy* and the δ is typically dropped from notation. We denote IR, RAM and KVS primitives that provide pure differentially private access with privacy budget ϵ as ϵ -DP-IR, ϵ -DP-RAM and ϵ -DP-KVS. When $\delta > 0$, the definition is referred to as *approximate differential privacy*, which is a weakening of pure differential privacy. We denote IR, RAM and KVS primitives that provide differentially private access with parameters ϵ and δ as (ϵ, δ) -DP-IR, (ϵ, δ) -DP-RAM and (ϵ, δ) -DP-KVS.

3 LOWER BOUNDS

In this section, we present negative results about DP-IR and DP-RAM in the single-server setting. As DP-KVS is an extension of DP-RAM with more functionality requirements, all DP-RAM lower bounds apply directly to DP-KVS.

3.1 Balls and Bins Model

Our lower bounds are presented in the balls and bins model of data manipulation. The n database records are treated as immutable, opaque balls. Each ball is associated with a mutable key containing metadata about the record. The formal definition involves a server $_m$ with storage for m balls and a client $_c$ with storage for c balls.

Definition 3.1 (Balls and Bins Model). A client $_c$ and server $_m$ operate in the balls and bins model if all client memory is initially empty and client-server interactions are restricted to the following:

- (1) Download ball from server $_m$ to client $_c$. For some $i \in [m]$ and $j \in [c]$, store the block at address i at server $_m$ in address j at client $_c$.

- (2) Upload ball from client_c to server_m. For some $i \in [m]$ and $j \in [c]$, store the block at address j at client_c in address i at server_m.

The above definition assumes a passive server acting only as storage. As a result, only lower bounds on communication can be proven directly. However, for algorithms with general computation by the server, lower bounds on communication in the balls and bins model may be modified to provide lower bounds on server computation. The modification simply views the transcript as the balls that must be operated on by the server.

Discussion about storage model. The balls and bins model does not include all possible encodings of databases that may be stored. In particular, we assume that database records are all stored together and the contents do not emit special properties. In general, related information should be stored in nearby memory locations so that all required data can be found with the minimal number of cache misses. Therefore, assuming database records are stored together captures real world scenarios. In addition, opaque balls rule out bothersome corner cases where non-trivial lower bounds cannot hold such as when records are dependent and record contents may be generated using superficial methods.

3.2 DP-IR

We prove our lower bounds directly for approximately differential privacy (in this case, $\delta \geq 0$). However, one can interpret the results for pure differential privacy by setting $\delta = 0$.

To prove our lower bounds, we first define the notion of *transcripts*. For a single-server information retrieval protocol IR in the balls and bins model, the transcript $\text{IR}(i)$ is the random variable denoting the set of blocks requested when retrieving B_i . The query algorithm of an IR algorithm in the balls and bins model only issues download commands to the server.

The main observation to prove the above theorem is that the initialization phase of DP-IR is public which implies that the adversary has knowledge of the identity of each record. By the restriction on correctness, the queried record must always be retrieved. Using the privacy requirement, it turns out all other blocks must also be retrieved. We summarize this in the following main lemma.

LEMMA 3.2. *Let IR be an (ϵ, δ) -DP-IR scheme in the balls and bins model. For any $a, b, c \in [n]$,*

- (1) $\Pr[B_c \in \text{IR}(a)] \leq e^\epsilon \Pr[B_c \in \text{IR}(b)] + \delta.$
- (2) $\Pr[B_c \notin \text{IR}(a)] \leq e^\epsilon \Pr[B_c \notin \text{IR}(b)] + \delta.$

PROOF. Define $S \subseteq \text{Range}(\text{IR})$ to be all transcripts where the block B_c is downloaded. So, for any query q ,

$$\Pr[B_c \in \text{IR}(q)] = \Pr[\text{IR}(q) \in S].$$

By (ϵ, δ) -DP, we know

$$\Pr[\text{IR}(a) \in S] \leq e^\epsilon \Pr[\text{IR}(b) \in S] + \delta$$

completing the first point. The second point follows identically if we choose $S \subseteq \text{Range}(\text{IR})$ to be all transcripts where block B_c is not downloaded. \square

Using the above lemma, we present very strong negative results for DP-IR that are errorless. That is, they always return the correct record. We show the following result:

THEOREM 3.3. *If IR is an (ϵ, δ) -DP-IR scheme in the balls and bins model then IR performs at least $(1 - \delta)n$ operations in expectation.*

PROOF. By Lemma 3.2, we know that

$$\Pr[B_j \notin \text{IR}(i)] \leq e^\epsilon \Pr[B_j \notin \text{IR}(j)] + \delta = \delta$$

since IR is errorless and $\Pr[B_j \notin \text{IR}(j)] = 0$. Therefore, we know that $\Pr[B_j \in \text{IR}(i)] \geq 1 - \delta$ and

$$\mathbb{E}[|\text{IR}(i)|] = \sum_{j \in [n]} \Pr[B_j \in \text{IR}(i)] \geq (1 - \delta)n$$

completing the proof. \square

For algorithms with general server computation, we can interpret the above results as a lower bound on server computation. This result is extremely strong since n server operations must be executed even when increasing the privacy budget. Therefore, the relaxation to ϵ -DP-IR does not result in any gain compared to PIR. For (ϵ, δ) -DP-IR, one could increase δ to decrease costs. However, typical privacy requires $\delta = \text{negl}(n)$ resulting in almost no gain.

To circumvent this result, we move to the case where DP-IR has a non-zero error rate $0 < \alpha \leq 1$. That is, DP-IR will only retrieve the desired record with probability $1 - \alpha$ depending only on the internal randomness of DP-IR. Our hope is that a very small α may significantly improve efficiency of DP-IR and bypass this negative result. We show:

THEOREM 3.4. *If IR is an (ϵ, δ) -DP-IR scheme in the balls and bins model with error probability $\alpha > 0$, then IR performs*

$$\Omega\left(\frac{(1 - \alpha - \delta) \cdot n}{e^\epsilon}\right)$$

operations in expectation.

PROOF OF THEOREM 3.4. Since IR has error probability at most α , $\Pr[B_j \in \text{IR}(j)] \geq 1 - \alpha$, for all $j \in [n]$. By Lemma 3.2, for all $j \neq i$,

$$\Pr[B_j \in \text{IR}(j)] \leq e^\epsilon \Pr[B_j \in \text{IR}(i)] + \delta.$$

Equivalently, this means that $\Pr[B_j \in \text{IR}(i)] \geq \frac{1-\alpha-\delta}{e^\epsilon}$. So,

$$\mathbb{E}[|\text{IR}(i)|] \geq \sum_{j \neq i} \Pr[B_j \in \text{IR}(i)] \geq (n-1) \frac{1-\alpha-\delta}{e^\epsilon}$$

that yields the theorem. \square

For any constant error $\alpha > 0$ and typical privacy budgets of $\epsilon = \Theta(1)$, errorless DP-IR schemes require $\Omega(n)$ server operations. However, it seems possible to bypass the first negative result with very small error. The above lower bound does not preclude the existence of a ϵ -DP-IR with $\epsilon = \Omega(\log n)$ and small error probability $\alpha > 0$ that only requires a constant number of server operations. In Section 5, we present an ϵ -DP-IR scheme that uses $O(1)$ communication when $\epsilon = O(\log n)$. It turns out this ϵ -DP-IR scheme asymptotically matches the lower bound for all values of $\epsilon \geq 0$.

We also extend our negative results for DP-IR in Section C where we present lower bounds when outsourcing storage to multiple, non-colluding servers.

3.3 DP-RAM

We now move to our negative results for DP-RAM. Proving lower bounds for DP-RAM is more challenging than DP-IR due to the private setup phase and the client's private memory. We can no longer directly bound the probability that blocks need to be retrieved and, instead, examine the transcripts seen by the adversary. Let us fix any transcript \mathcal{T} that has non-zero probability of being viewed by the adversary on any query sequence Q of length l . We show that every other query sequence of length l must also induce \mathcal{T} with non-zero probability. In fact, the probabilities that any two fixed query sequences induce \mathcal{T} as the view of the adversary are strongly related by their Hamming distance and the privacy budget, ϵ , as described in the following:

LEMMA 3.5. *Let RAM be a ϵ -DP-RAM scheme for any $\epsilon \geq 0$. For every distribution \mathcal{D} on query sequences of length l , random variable $D \sim \mathcal{D}$ and for any two fixed sequences Q_1, Q_2 such that $\Pr[D = Q_2] > 0$, then*

$$\frac{\Pr[D = Q_1 \mid \text{RAM}(D) = \mathcal{T}]}{\Pr[D = Q_2 \mid \text{RAM}(D) = \mathcal{T}]} \geq e^{-\epsilon \cdot d(Q_1, Q_2)} \frac{\Pr[D = Q_1]}{\Pr[D = Q_2]}.$$

PROOF. By Bayes' law, we have

$$\begin{aligned} \Pr[D = Q \mid \text{RAM}(D) = \mathcal{T}] &= \frac{\Pr[\text{RAM}(D) = \mathcal{T} \mid D = Q] \Pr[D = Q]}{\Pr[\text{RAM}(D) = \mathcal{T}]} \\ &\geq e^{-\epsilon \cdot d_H(Q, Q')} \frac{\Pr[\text{RAM}(D) = \mathcal{T} \mid D = Q'] \Pr[D = Q]}{\Pr[\text{RAM}(D) = \mathcal{T}]} \\ &= e^{-\epsilon \cdot d_H(Q, Q')} \frac{\Pr[D = Q' \mid \text{RAM}(D) = \mathcal{T}] \Pr[D = Q]}{\Pr[D = Q']} \end{aligned}$$

giving us our result. \square

Next, we prove that if a query sequence has positive probability, then it stays so even after seeing a transcript. This holds for all *a priori* query sequences distributions \mathcal{D} .

LEMMA 3.6. *Let RAM be an ϵ -DP-RAM and let \mathcal{T} be a transcript for which there exists at least one query sequence Q such that*

$$\Pr[\text{RAM}(Q) = \mathcal{T}] > 0.$$

Then for all distributions D on the set of query sequences and for all Q' such that $\Pr[D = Q'] > 0$ it holds that

$$\Pr[D = Q' \mid \text{RAM}(D) = \mathcal{T}] > 0.$$

PROOF. Let Q be any query sequence such that $\Pr[D = Q] > 0$ and $\Pr[\text{RAM}(D) = \mathcal{T} \mid D = Q] > 0$. Assume by contradiction that $\Pr[D = Q'] > 0$ but $\Pr[D = Q' \mid \text{RAM}(D) = \mathcal{T}] = 0$ for some query sequence Q' . By Bayes' Law,

$$\begin{aligned} 0 &= \Pr[D = Q' \mid \text{RAM}(D) = \mathcal{T}] \\ &= \frac{\Pr[\text{RAM}(D) = \mathcal{T} \mid D = Q'] \cdot \Pr[D = Q']}{\Pr[\text{RAM}(D) = \mathcal{T}]} \end{aligned}$$

which implies that $\Pr[\text{RAM}(D) = \mathcal{T} \mid D = Q'] = 0$. On the other hand, by ϵ -differential privacy, we have that for all Q such that $\Pr[D = Q] > 0$,

$$\begin{aligned} \Pr[\text{RAM}(D) = \mathcal{T} \mid D = Q] \\ \leq e^{d_H(Q, Q')\epsilon} \cdot \Pr[\text{RAM}(D) = \mathcal{T} \mid D = Q'] = 0 \end{aligned}$$

providing our contradiction. \square

We are now ready to prove our DP-RAM lower bound. For the sake of the completeness, we suppose that the DP-RAM scheme only retrieves the desired block with probability at least $1 - \alpha$ based only on internal randomness. We show that:

THEOREM 3.7. *If RAM is an ϵ -DP-RAM scheme in the balls and bins model with error probability $\alpha \geq 0$ then RAM performs*

$$\Omega\left(\log_c\left(\frac{(1-\alpha) \cdot n}{e^\epsilon}\right)\right)$$

expected amortized operations per query when the client has storage for c balls.

PROOF. Fix $Q = (q_1, \dots, q_l)$ to be any sequence of queries of length $l > c$. Then, client_c can only perform upload/download operations with server_M . Denote the expected amortized bandwidth by k . The server locations of all operations are given by the transcript, $\text{RAM}(q)$. However, which of the c client locations used for each operation remains hidden. For each download and upload, there are c possible execution paths, one path for each of the c client memory locations. Also, after each operation, client_c may use the data stored in any of the c client memory locations to answer a query. Altogether, at most kl blocks of bandwidth are used for all l queries. Therefore, there are at most c^{2kl} different sequences

of blocks that may be returned. Since RAM can only fail with α probability for each query, it must satisfy at least $(1 - \alpha)^l n^l$ different access patterns. So, $c^{2kl} \geq (1 - \alpha)^l n^l$.

Let \mathcal{T} be any transcript with positive probability and consider the uniform distribution U over the set $[n]^l$ of query sequences of length l . By Lemma 3.6, there exists $Q \in [n]^l$ such that

$$\Pr[U = Q \mid \text{RAM}(U) = \mathcal{T}] \geq \frac{1}{n^l} \geq \frac{(1 - \alpha)^l}{c^{2kl}}.$$

For every $Q' \neq Q$, by Lemma 3.5, we have

$$\Pr[U = Q' \mid \text{RAM}(U) = \mathcal{T}] \geq e^{-l\epsilon} \frac{(1 - \alpha)^l}{c^{2kl}}.$$

Finally, we know that

$$1 \geq \sum_{Q' \neq Q} \Pr[U = Q' \mid \text{RAM}(U) = \mathcal{T}] \geq e^{-l\epsilon} (n^l - 1) \frac{(1 - \alpha)^l}{c^{2kl}}$$

completing the proof. \square

Unlike DP-IR, we do not need a non-zero error probability to get a constant overhead scheme. From the above lower bound, it seems like the best one can achieve is a perfectly correct ϵ -DP-RAM with $\epsilon = \Theta(\log n)$ that only requires $O(1)$ overhead with small client storage. In Section 6, we show that such a scheme does exist.

Discussion about lower bounds. It has been shown that previous ORAM lower bounds come with many caveats. The first lower bound [27] in the balls and bins model was for statistical security as pointed out in [11]. The assumption of statistical security is troublesome due to the fact that all ORAM schemes require the use of encryption which only provides computational security. To abstract away this issue, it is assumed that records were opaque balls and were hidden even against computationally unbounded adversaries. However, this abstraction was still not sufficient since the majority of ORAM schemes still required the use of pseudo-random functions. Furthermore, the DP-RAM constructions of [50] and [15] require encryption and are burdened by the same caveats. We explain why these issues do not apply to our results. In our work, we consider typical differential privacy notions (not the computational variants described in [40]) where adversaries are computationally unbounded. Our DP-RAM scheme in Section 6 when only allowing retrievals does not require encryption and provides differentially private access to public data against computationally unbounded adversaries. All ORAM schemes still require encryption even when only retrievals are permitted. In the case that we wish to store encrypted data or protect overwrites in our DP-RAM scheme, we must apply the abstraction of opaque balls to go around encryption. With this abstraction, our DP-RAM scheme is differentially private. Recent

results in [37] present $\Omega(\log(n/c))$ lower bounds for ORAM against computational adversaries with passive servers and general storage schemes. Furthermore, work in [47] extends the lower bound for (ϵ, δ) -DP-RAM where $\epsilon = O(1)$ and $\delta \leq 1/3$. However, the lower bounds in [47] have an exponentially worse dependency on ϵ compared to our results.

4 AN INSECURE CONSTRUCTION

Before presenting our constructions, we consider a simple and tempting, but insecure, construction. The lower bounds presented in Section 3 show that the best privacy any constant overhead storage primitive can achieve will be $\epsilon = \Theta(\log n)$. With such weak privacy requirements, it might seem very easy to construct these primitives at first. We caution that schemes with these weak privacy requirements must be constructed carefully as slight variants of our later schemes could also end up being insecure. To our knowledge, our schemes are the simplest constructions that achieve $\epsilon = \Theta(\log n)$ differential privacy and small overhead.

The main idea of the strawman solution derives from the fact that $\epsilon = \Theta(\log n)$. As a result, the desired block should be queried with probability a multiplicative factor of $\text{poly}(n)$ larger to compared to any other block. To achieve this, one could query the desired block with probability 1 and all other blocks with probability $1/n$. This scheme would have $O(1)$ bandwidth in expectation, perfect correctness and no client storage requirements. However, we show that this scheme is really an (ϵ, δ) -DP-IR with $\epsilon = \Theta(\log n)$ and $\delta = (n - 1)/n$. Denote the above scheme by IR and $\text{IR}(i)$ the set of blocks returned when querying for i . Pick any two queries $i \neq j$. Note, $\Pr[B_i \notin \text{IR}(i)] = 0$ and $\Pr[B_i \notin \text{IR}(j)] = (n - 1)/n$. Then, $(n - 1)/n = \Pr[B_i \notin \text{IR}(j)] \leq e^\epsilon \Pr[B_i \notin \text{IR}(i)] + \delta$ which means that $\delta \geq (n - 1)/n$. Therefore, the above scheme would not provide as much privacy as possible as δ approaches 1 as n increases.

We use the above strawman to show that attentiveness and rigor are required when constructing differentially private storage schemes even with such weak privacy guarantees. The schemes that we will present for DP-IR, DP-RAM and DP-KVS are the simplest algorithms that, to our knowledge, achieve our desired privacy of $\epsilon = \Theta(\log n)$ and small overhead in efficiency.

5 DP-IR CONSTRUCTION

For errorless DP-IR, the asymptotically optimal balls and bins algorithms is required to download the entire database regardless of the privacy budget. With active servers that can perform computation, servers must perform n operations which is identical to PIR. Therefore, there is no reason to use differentially private access when oblivious access has the same efficiency.

We now move to the more interesting case of DP-IR with errors. By introducing a small amount of error $\alpha > 0$, we hope to find algorithms more efficient than PIR for larger privacy budgets where our lower bounds from Section 3 no longer hold. We show that the simplest algorithm ends up being optimal. In particular, the client will download the desired block as well as several other blocks simultaneously. The hope is that the adversary cannot determine the real retrieval from all the fake retrievals. In addition, with probability α , we only perform fake retrievals and error. The pseudocode of this scheme can be found in Appendix G. The proof of the following theorem is found in Appendix B

THEOREM 5.1. *For any $\epsilon \geq 0$, there exists a ϵ -DP-IR that returns $O(n/e^\epsilon)$ blocks for any constant error probability $\alpha > 0$.*

The above upper bound asymptotically matches the lower bound of Theorem 3.4 for all values of $\epsilon \geq 0$. Furthermore, by fixing the privacy budget to be $\epsilon = \Theta(\log n)$, we can achieve a constant overhead DP-IR scheme with the best privacy according to Theorem 3.4.

6 DP-RAM CONSTRUCTION

In this section, we give an errorless construction DP-RAM supporting both retrieval and overwrite operations. Before describing our DP-RAM scheme, we note that one could also use the DP-IR scheme from Section 5 as a DP-RAM scheme without any client storage requirements. However, this DP-IR scheme has non-zero error probabilities which is inherently unavoidable for DP-IR due to our lower bounds in Section 3. On the other hand, our DP-RAM lower bounds do not preclude the existence of a perfectly correct DP-RAM scheme with $\epsilon = \Theta(\log n)$. In this section, we work towards constructing such a DP-RAM scheme.

Our scheme will require the client to store some records in a local *stash*. Our DP-RAM scheme is parameterized by a probability p describing the independent probability that each record is stored in the stash. We assume that (Enc, Dec) is an IND-CPA symmetric-key encryption scheme. The server's storage will consist of an array, A , of n records.

The setup phase of DP-RAM will consist of populating the server-stored array A . For security parameter λ , key $K \rightarrow \{0, 1\}^\lambda$ is randomly selected. A is initialized by setting $A[i] = \text{Enc}(K, B_i)$ for $i \in [n]$. The stash is initialized by independently selecting each record to be in the stash with probability p . In addition, the client keeps the key K in local storage.

The querying (either retrieval or overwrite) for a record B_i consists of two phases: the *download* phase followed by the *overwrite* phase. In the download phase, the client looks for B_i in the stash. If B_i is found, then B_i is removed and returned. The client asks the server for $A[j]$, with j chosen uniformly

at random from $[n]$. If, instead, B_i is not in the stash, the client asks the server for $A[i]$ that contains an encryption of B_i . If the client is performing a write operation, then B_i is updated with the new version. At this point, the client holds the current version of B_i .

In the overwrite phase, the current version of B_i is added to the stash with probability p . If B_i is stored in the stash, then another record is randomly selected, downloaded from the server, decrypted and then re-encrypted with fresh randomness and uploaded to the server. If B_i is not stored in the stash, then the client asks the server for $A[i]$, discards the record received and then uploads to $A[i]$ a freshly computed ciphertext carrying the current version of B_i .

The pseudocode of the algorithms are presented in Appendix H. We note that while the above algorithm is simple, the analysis of privacy is quite complicated. We show the following about this scheme when $p \leq \Phi(n)/n$ for any $\Phi(n) = \omega(\log n)$:

THEOREM 6.1. *There exists an $O(\log n)$ -DP-RAM that returns $O(1)$ blocks. For any function $\Phi(n) = \omega(\log n)$, the client stores $\Phi(n)$ blocks of client storage except with probability $\text{negl}(n)$.*

This scheme is, essentially, the best privacy that can be achieved by an errorless DP-RAM scheme with constant overhead according to Theorem 3.7. Since we use encryption, our DP-RAM satisfies computational differential privacy using simulators (see SIM-CDP in [40]). Our proof uses a simulator that replaces all encryptions of records with randomly generated contents.

Discussion about encryption. The above DP-RAM scheme assumes that both record retrievals and overwrites are permitted. To hide whether queries are overwrites or retrievals, any DP-RAM scheme must use encryption. In the case that we wish to only permit retrievals, we note that the above DP-RAM scheme no longer requires encryption and can provide differentially private access to public data without computational assumptions. In particular, the entire overwrite phase may be skipped. Retrieval-only DP-RAM and DP-IR only differ by their requirements on the client's state. Therefore, using client state is another way to bypass the strong lower bound shown in Theorem 3.3 for errorless DP-IR schemes.

6.1 Roadmap of the Proof

To start, we show that if we choose that $p \leq c/n$ where $c = \omega(\log n)$, then the client will store at most $O(c)$ blocks. The proof is an application of Chernoff Bounds and postponed to Appendix D.

The technical crux of the privacy analysis lies in bounding the following ratio, for every transcript \mathcal{T} seen by the adversary, and for every two neighboring query sequences

Q and Q' ,

$$\frac{\Pr[\text{RAM}(Q') = \mathcal{T}]}{\Pr[\text{RAM}(Q) = \mathcal{T}]} \quad (1)$$

where $\Pr[\text{RAM}(Q) = \mathcal{T}]$ denotes the probability that RAM executing on Q produces the transcript \mathcal{T} . In the first step, we show that the transcript of the adversary at a single query on B_q is only dependent on the most recent query that also queries for B_q . In other words, the ratio is upper bounded by

$$\prod_{i=1}^l \frac{\Pr[\text{RAM}_i(Q') = \mathcal{T}_i \mid \text{RAM}_{\text{pr}(Q',i)}(Q') = \mathcal{T}_{\text{pr}(Q',i)}]}{\Pr[\text{RAM}_i(Q) = \mathcal{T}_i \mid \text{RAM}_{\text{pr}(Q,i)}(Q) = \mathcal{T}_{\text{pr}(Q,i)}]}$$

where $\text{pr}(Q, i)$ is the most recent query for block B_{q_i} before the i -th query. The second step consists of giving an upper bound that holds for each factor of the product in the above equation. The third step shows that the upper bound computed in the second step is too pessimistic and that, all factors, except for 3, in the product in the right hand side of the above equation are 1 when Q and Q' differ in exactly one position.

We now proceed to define notation and terminology used throughout our proof. First, we observe that the transcript of an execution only includes the ciphertexts of the blocks and not the actual content of the blocks transferred. Assuming IND-CPA of the underlying encryption scheme (Enc, Dec), it is straightforward to prove that, for a sequence $Q = (q_1, \dots, q_l)$ of l queries, the transcript generated by Q for blocks B_1, \dots, B_n is indistinguishable from the transcript generated by the same sequence Q for n blocks that are $\mathbf{0}$. For this reason, we shall not consider the ciphertexts of the blocks as part of the transcript and consider a transcript \mathcal{T} for query sequence Q as a sequence $\mathcal{T} = ((d_1, o_1), \dots, (d_l, o_l))$ of l pairs $\mathcal{T}_j = (o_j, d_j)$ of indices of the blocks that are accessed during the download phase and the overwrite phase. We set $\mathcal{T}_{[j]} = ((d_1, o_1), \dots, (d_j, o_j))$. We define $\text{RAM}_j^D(Q)$ and $\text{RAM}_j^O(Q)$ to be the random variables of the indices of the download and of the overwrite block of the j -th query for all $j \in [l]$. Also, for $S \subseteq [l]$ we define $\text{RAM}_S(Q)$ to be the set of random variables $\{\text{RAM}_j^D(Q), \text{RAM}_j^O(Q)\}_{j \in S}$. Also, we set $\text{RAM}(Q) := \text{RAM}_{[n]}(Q)$. It turns out to be convenient to extend the random variable RAM for indices $\pm\infty$ by setting $\text{RAM}_{\pm\infty}^D(Q)$ and $\text{RAM}_{\pm\infty}^O(Q)$ to be random variables that give probability 1 to \perp .

We also let $\text{pr}(Q, j)$ be the index of the most recent *previous* query to block B_{q_j} that happened before the j -th query of Q ; that is, $\text{pr}(Q, j) = \max\{i < j : q_i = q_j\}$. If query q_j is the first query of sequence Q that asks for block B_{q_j} , then $\text{pr}(Q, j) = +\infty$. Similarly, we define $\text{nx}(Q, j)$ to be the index of the nearest *next* query for block B_{q_j} ; that is, $\text{nx}(Q, j) = \min\{i > j : q_i = q_j\}$. If the j -th query of Q is the last query of Q to ask for block B_{q_j} , then $\text{nx}(Q, j) = -\infty$.

6.2 Step I: Reducing dependencies

The next lemmas outline the dependencies of the random variables of the download blocks and of the overwrite blocks. We start by showing that the overwrite block of each query is independent of all previous history and only depends on the current query.

LEMMA 6.2. *For every query sequence Q of length l , for every transcript $\mathcal{T} = ((d_1, o_1), \dots, (d_l, o_l))$ and for every $j \leq l$*

$$\Pr[\text{RAM}_j^O(Q) = o_j \mid \text{RAM}_{[j-1]}(Q) = \mathcal{T}_{[j-1]} \wedge \text{RAM}_j^D(Q) = d_j] = \Pr[\text{RAM}_j^O(Q) = o_j].$$

Moreover, if Q and Q' are two sequences with $q_j = q'_j$, then the distributions $\text{RAM}_j^O(Q)$ and $\text{RAM}_j^O(Q')$ coincide.

PROOF. The lemma follows by observing that the distribution of the j -th overwrite block depends only on whether B_{q_j} is added to the block stash during the j -th overwrite phase which, in turn, depends only on q_j and the random value of r in the overwrite phase of the j -th query. \square

LEMMA 6.3. *For every query sequence Q of length l , for every transcript $\mathcal{T} = ((d_1, o_1), \dots, (d_l, o_l))$ and for every $j \leq l$,*

$$\Pr[\text{RAM}_j^D(Q) = d_j \mid \text{RAM}_{[j-1]}(Q) = \mathcal{T}_{[j-1]}] = \Pr[\text{RAM}_j^D(Q) = d_j \mid \text{RAM}_{\text{pr}(Q,j)}^O = o_{\text{pr}(Q,j)}].$$

Moreover, if Q and Q' are two sequences with $q_j = q'_j$ and $\text{pr}(Q, j) = \text{pr}(Q', j)$ then for all $d \in [n]$ and for all $o \in [n] \cup \{\perp\}$ such that $\Pr[\text{RAM}_{\text{pr}(Q,j)}^O = o] > 0$,

$$\Pr[\text{RAM}_j^D(Q) = d \mid \text{RAM}_{\text{pr}(Q,j)}^O = o] = \Pr[\text{RAM}_j^D(Q') = d \mid \text{RAM}_{\text{pr}(Q',j)}^O = o].$$

PROOF. The j -th download block of sequence Q depends on whether B_{q_j} is in the block stash at the start of the j -th download phase. We distinguish two cases.

If $\text{pr}(Q, j) = +\infty$, then block B_{q_j} has not been queried in the first $j - 1$ queries and its probability of being found in the stash at the start of j -th download phase is equal to the probability of being placed in the stash by RAM.Setup which, obviously, is independent from the previous history.

Suppose instead that $\text{pr}(Q, j) \in [j-1]$. Then the probability that B_{q_j} is in the stash at the beginning of the j -th query depends only on the random variable, $\text{RAM}_{\text{pr}(Q,j)}^O(Q)$, of the overwrite block of query $\text{pr}(Q, j)$. \square

The first part of Lemma 6.2 and the first part of Lemma 6.3 imply

$$\begin{aligned} \Pr[\text{RAM}_j(Q) = (d_j, o_j) \mid \text{RAM}_{[j-1]}(Q) = \mathcal{T}_{[j-1]}] &= \\ \Pr[\text{RAM}_j^D(Q) = d_j \mid \text{RAM}_{[j-1]}(Q) = \mathcal{T}_{[j-1]}] &= \\ \Pr[\text{RAM}_j^O(Q) = o_j \mid \text{RAM}_{[j-1]}(Q) = \mathcal{T}_{[j-1]} \wedge \text{RAM}_j^D(Q) = d_j] &= \\ \Pr[\text{RAM}_j^D(Q) = d_j \mid \text{RAM}_{\text{pr}(Q,j)}(Q) = o_{\text{pr}(Q,j)}] \cdot \Pr[\text{RAM}_j^O(Q) = o_j]. \end{aligned}$$

Thus, we can write

$$\begin{aligned} \frac{\Pr[\text{RAM}(Q') = \mathcal{T}]}{\Pr[\text{RAM}(Q) = \mathcal{T}]} &= \prod_{j=1}^l \frac{\Pr[\text{RAM}_j^O(Q) = o_j]}{\Pr[\text{RAM}_j^O(Q') = o_j]} \\ &\times \prod_{j=1}^l \frac{\Pr[\text{RAM}_j^D(Q') = d_j \mid \text{RAM}_{\text{pr}(Q',j)}(Q) = o_{\text{pr}(Q',j)}]}{\Pr[\text{RAM}_j^D(Q) = d_j \mid \text{RAM}_{\text{pr}(Q,j)}(Q') = o_{\text{pr}(Q,j)}]} \end{aligned}$$

6.3 Step II: Upper bounding factors

In the next two lemmata, we give an upper bound on the contribution of each $j \in [l]$ to the product in the equation above.

LEMMA 6.4. *Let Q and Q' be two query sequences of length l . For every transcript $\mathcal{T} = ((d_1, o_1), \dots, (d_l, o_l))$ and every $j \in [l]$*

$$\frac{\Pr[\text{RAM}_j^D(Q') = d_j \mid \text{RAM}_{\text{pr}(Q',j)}^O(Q') = o_{\text{pr}(Q',j)}]}{\Pr[\text{RAM}_j^D(Q) = d_j \mid \text{RAM}_{\text{pr}(Q,j)}^O(Q) = o_{\text{pr}(Q,j)}]} \leq \frac{n^2}{p}.$$

LEMMA 6.5. *Let Q and Q' be two query sequences of length l . For every transcript $\mathcal{T} = ((d_1, o_1), \dots, (d_l, o_l))$ and every $j \in [l]$*

$$\frac{\Pr[\text{RAM}_j^O(Q') = o_j]}{\Pr[\text{RAM}_j^O(Q) = o_j]} \leq \frac{n}{p}.$$

Both of these lemmata consider the various cases that can occur. As the proofs are case analysis that do not provide better intuition to the problem, we postpone them to Section D.

6.4 Step III: Identifying the many good cases

The bounds given by Lemma 6.4 and Lemma 6.5 would give an $n^{O(l)}$ upper bound on the ratio in Equation 1. This is a very weak bound as it depends on the length l of the sequences. In this section, we tighten the upper bound to $n^{O(1)}$ which is instrumental to prove that RAM is private with $\epsilon = O(\log n)$. Specifically, the next lemma gives sufficient conditions under which the ratio is actually 1 and then we show that, if the two sequences only differ in one position, then there are only three values of j for which the conditions are not satisfied and for those position we use the upper bound of the previous section.

The following lemma follows directly from the second parts of Lemma 6.2 and 6.3.

LEMMA 6.6. *For any two sequences Q and Q' of the same length, every transcript $\mathcal{T} = ((d_1, o_1), \dots, (d_l, o_l))$ and every $j \in [l]$ with $\text{pr}(Q, j) = \text{pr}(Q', j)$ and $q_j = q'_j$,*

$$\begin{aligned} \Pr[\text{RAM}_j^O(Q) = o_j] &= \\ \Pr[\text{RAM}_j^D(Q) = d_j \mid \text{RAM}_{\text{pr}(Q,j)}^O(Q) = o_{\text{pr}(Q,j)}] &= \\ \Pr[\text{RAM}_j^O(Q') = o_j] &= \\ \Pr[\text{RAM}_j^D(Q') = d_j \mid \text{RAM}_{\text{pr}(Q',j)}^O(Q') = o_{\text{pr}(Q',j)}]. \end{aligned}$$

The lemma above says that the distributions of the transcripts associated with two query sequences Q and Q' may differ only at indices j for which $\text{pr}(Q, j) \neq \text{pr}(Q', j)$ or $q_j \neq q'_j$. The next lemma identifies the indices j for which this happens when Q and Q' differ in exactly one position.

LEMMA 6.7. *Let Q and Q' be two query sequences of length l differing only at position $k \in [l]$. If $j \notin \{k, \text{nx}(Q, k), \text{nx}(Q', k)\}$, then $\text{pr}(Q, j) = \text{pr}(Q', j)$.*

PROOF. For $j < k$, we have $(q_1, \dots, q_j) = (q'_1, \dots, q'_j)$ and thus $\text{pr}(Q, j) = \text{pr}(Q', j)$. For $j > k$ such that $q_j \neq q'_k, q'_k$ we have $q_j = q'_j$ and $\text{pr}(Q, j) = \text{pr}(Q', j)$. Next consider the indices $j_1 < \dots < j_l$ such that $q_k = q_{j_1} = \dots = q_{j_l}$. Clearly, if $l > 0$ then $j_1 = \text{nx}(Q, k)$ and, for $i = 2, \dots, l$, we have $q_{j_i} = q'_{j_i}$ and $\text{pr}(Q, j_i) = \text{pr}(Q', j_i)$. A similar argument applies for the indices $j > k$ such $q'_j = q'_k$ thus completing the proof of the theorem. \square

6.5 Wrapping up the proof

PROOF OF THEOREM 6.1. Bounds on bandwidth and server storage are obvious and the one on client storage follows from Lemma D.1.

Consider sequences Q and Q' of length l that differ only in position $k \in [l]$. Let $\mathcal{T} = ((d_1, o_1), \dots, (d_l, o_l))$ be transcript. By Lemma 6.6 and 6.7, we obtain that the ratio $\frac{\Pr[\text{RAM}(Q') = \mathcal{T}]}{\Pr[\text{RAM}(Q) = \mathcal{T}]}$ is

$$\prod_{j \in S} \frac{\Pr[\text{RAM}_j(Q') = \mathcal{T}_j \mid \text{RAM}_{[j-1]}(Q') = \mathcal{T}_{[j-1]}]}{\Pr[\text{RAM}_j(Q) = \mathcal{T}_j \mid \text{RAM}_{[j-1]}(Q) = \mathcal{T}_{[j-1]}]}$$

where $S = \{k, \text{nx}(Q, k), \text{nx}(Q', k)\}$. By Lemma 6.4 and 6.5,

$$\frac{\Pr[\text{RAM}(Q') = \mathcal{T}]}{\Pr[\text{RAM}(Q) = \mathcal{T}]} = \left(\frac{n}{p}\right)^{O(1)}.$$

Since, the above holds for any single transcript, we get that the ratio holds for any set of transcripts. This implies that $\epsilon = O(\log n)$. \square

7 DP-KVS CONSTRUCTION

In this section, we present DP-KVS, our KVS construction with privacy budget $\epsilon = O(\log n)$ and small overhead. Recall that a KVS is an extension of RAM where each of the n blocks

is identified by a unique key taken from a possibly large universe of keys U . In contrast, a block in RAM is uniquely identified by an integer in $[n]$. One could use the DP-RAM construction from Section 6, which results in server storage on the order of $|U| \gg n$. For efficiency, we would like a DP-KVS scheme that stores $O(n)$ blocks on the server.

Our approach to constructing a DP-KVS scheme consists of two steps. First, we show that we can construct a DP-KVS scheme using a DP-RAM scheme and a *mapping scheme* which associates keys in the universe U to subsets of server storage. Next, we present an efficient mapping scheme by constructing a non-trivial *oblivious* variant of two-choice hashing [41] that uses $O(n)$ server storage and which may be of independent interest. Our oblivious two-choice hashing variant hides the number of real items that are stored in each bin at any point in time. Finally, combined with our DP-RAM scheme of Section 6, we present a (ϵ, δ) -DP-KVS with $\epsilon = O(\log n)$ and $\delta = \text{negl}(n)$ using only $O(\log \log n)$ overhead. While non-constant, this is exponentially better than any previous oblivious KVS scheme built from ORAMs. Furthermore, for all practical sizes of n , $O(\log \log n)$ is very small.

7.1 Composing Mapping Schemes and DP-RAM

In this section, we present a generic reduction of DP-KVS to a mapping scheme and our DP-RAM from Section 6. First, we define mapping schemes. Afterwards, we present the reduction.

A mapping scheme is defined as the tuple (Π, \mathcal{S}) where Π is the *mapping function* and \mathcal{S} is the *storing algorithm*. To store n items each uniquely identified by an key from the universe U , a mapping scheme arranges the server storage into $b(n)$ buckets each consisting of at most $s(n)$ blocks. Each bucket is uniquely identified by an index from $[b(n)]$. We note that buckets are not necessarily disjoint and the total number of blocks may be much smaller than $b(n) \cdot s(n)$. In addition, the mapping scheme assumes that the client will hold a *mapping stash* which will contain at most $c(n)$ blocks except with probability $\text{negl}(n)$. For each item $u \in U$, the mapping function maps u to a subset of at most $s(n)$ buckets defined by $\Pi(u) \subseteq [b(n)]$. For convenience, we denote $k(n) := \max_{u \in U} |\Pi(u)|$. When inserting a new item with identifier u , the storing algorithm \mathcal{S} determines whether u is placed into a bucket of $\Pi(u)$ or the mapping stash according to the sizes of buckets in $\Pi(u)$. We now show how to use mapping schemes to construct a DP-KVS.

Our DP-KVS scheme works as follows. We construct the server storage into $b(n)$ buckets as described by the mapping scheme. We build a DP-RAM to be able to query and update the $b(n)$ buckets. In Appendix E, we show that our DP-RAM

construction from Section 6 remains secure and efficient when querying possibly overlapping buckets with minor modifications. When querying our DP-KVS for a key $u \in U$, we perform $k(n)$ DP-RAM queries to retrieve the $s(n)$ blocks from each bucket in $\Pi(u)$. If $|\Pi(u)| < k(n)$, we pick random buckets to pad $\Pi(u)$ to size $k(n)$. Now, we are guaranteed that if u exists in DP-KVS, it appears in a bucket of $\Pi(u)$ or in the mapping stash and can be thus returned. To update an existing key, we can simply update either the bucket or the mapping stash that contains the block associated to key u . For insertion, we can execute the storing algorithm \mathcal{S} as all the contents of buckets $\Pi(u)$ and the mapping stash are available to the client. \mathcal{S} determines the insertion location for the block associated to key u . Finally, we execute $k(n)$ DP-RAM updates to all buckets in $\Pi(u)$. Only the bucket containing the block associated with key u is updated. The other buckets will perform fake updates where the contents remain unchanged. For read operations, none of the contents of buckets in $\Pi(u)$ will be changed. We prove the following theorem about our DP-KVS construction.

THEOREM 7.1. *For n blocks, the above KVS scheme is an ϵ -DP-KVS with $\epsilon = O(k(n) \cdot \log n)$ that returns at most $O(k(n) \cdot s(n))$ blocks. For any function $\Phi(n) = \omega(\log n)$, the client stores $O(s(n) \cdot \Phi(n) + c(n))$ blocks of storage except with probability $\text{negl}(n)$.*

PROOF. The bounds on bandwidth and server storage follow from the mapping scheme properties. The client storage bound follows from Lemma D.1 and the mapping scheme properties. Each query results in at most $2 \cdot k(n)$ queries over the repertoire Σ of the n buckets. By the composition theorem and the discussion in Appendix E, we obtain that $\epsilon = O(k(n) \cdot \log n)$. \square

7.2 Oblivious Two-Choice Hashing

Before presenting our new mapping scheme, we revisit the two-choice hashing scheme [41] (see Section A.1 for more details). This scheme considers n buckets to store up to n keys. The mapping function $\Pi : U \rightarrow [n]$ sets $\Pi(u)$ as $k(n) := 2$ independently and uniformly at random chosen buckets. Typically, Π is succinctly represented using two keys, $\text{key}_1, \text{key}_2$, of a pseudorandom function F and by $\Pi(u) := \{F(\text{key}_1, u), F(\text{key}_2, u)\}$. The storing algorithm \mathcal{S} for $u \in U$, checks which of the two buckets in $\Pi(u)$ is less loaded and places u into the less loaded bucket. There are several different proofs that show that the largest bucket will contain at most $s(n) := O(\log \log n)$ items except with probability $\text{negl}(n)$.

Unfortunately, we are unable to use two-choice hashing directly into our DP-KVS scheme without incurring into a server storage blockup. The DP-KVS scheme from Section 7.1 requires that all buckets are the same size for privacy. The

naive approach is to simply increase all buckets to the worst case size which results in $O(n \log \log n)$ server storage. Instead, we now present a variant of two-choice hashing which will only use $O(n)$ server storage by arranging buckets to share memory.

Our bucket arrangement is best described as $\Theta(n/\log n)$ identical binary trees, each with $\Theta(\log n)$ leaf nodes and $\Theta(\log \log n)$ depth. Leaf nodes are denoted as *height 0* and the height increases going towards the root. This results in a total of $\Theta(n)$ nodes over all binary trees. Each node in the tree will be able to store up to $t = \Theta(1)$ blocks. Furthermore, we pick the binary trees such that there are exactly n leaf nodes overall. Finally, there is a single root node that has $\Theta(n/\log n)$ children corresponding to the roots of the $\Theta(n/\log n)$ binary tree roots. We denote this node as the *super root*. Unlike all other nodes, the super root is stored on the client. We shall show that the probability that the super root holds more than $\Phi(n)$ blocks for any $\Phi(n) = \omega(\log n)$ is negligible in n .

Each of the n buckets is uniquely associated with a leaf node. The memory locations of a bucket consist of all the blocks stored in the nodes on the unique path between the leaf node and the super root. Therefore, a bucket consists of $\Theta(\log \log n)$ server memory locations in addition to the memory locations in the super root. We now describe our new storing algorithm, \mathcal{S} , given this bucket arrangement. When inserting u , \mathcal{S} places u into the node with minimal height (that is, closest to the leaf nodes) in either of the buckets in $\Pi(u)$ with empty space. Note that u might end up being stored in the super root. If all the nodes of both buckets of $\Pi(u)$ are filled then the mapping scheme fails to store u . We show that when inserting any set of at most n keys, if we limit the capacity of the super root to $\Phi(n)$, for some $\Phi(n) = \omega(\log n)$, the above mapping scheme fails with probability $\text{negl}(n)$. The analysis adapts techniques from [5].

THEOREM 7.2. *Let $\Phi(n) = \omega(\log n)$. The probability that, when inserting n keys, mapping scheme \mathcal{S} places more than $\Phi(n)$ keys into the super root is $\text{negl}(n)$.*

PROOF. Our key observation is that a block is stored in a level- $(i+1)$ slot if and only if the two selected buckets are filled up to level i . If we denote by H_i the number of filled node at level i , the probability of selecting a bucket that is filled up to level i is $H_i \cdot 2^i/n$. For a block to be allocated to a slot at level $i+1$, it must be the case that both selected buckets are filled up to level i which has probability $\left(\frac{H_i \cdot 2^i}{n}\right)^2$. For convenience, we define the sequence β_i by setting $\beta_0 = \frac{n}{e \cdot 3^4}$ and $\beta_{i+1} = \frac{e}{n} \cdot \beta_i^2 \cdot 2^{2(i+1)}$, which we will use later in the proof. First, we present a useful property about the sequence and then we show that the probability that $H_i > \beta_i$ is $\text{negl}(n)$. The proof of Lemma 7.3 can be found in Appendix F.

LEMMA 7.3. *For all $i \geq 0$,*

$$\beta_i = \frac{n}{e} \cdot \left(\frac{2}{3}\right)^{2^{i+2}} \left(\frac{1}{2}\right)^{2(i+2)}.$$

LEMMA 7.4. *If $\beta_i = \omega(\log n)$ then, $\Pr[H_i > \beta_i] \leq i/n^{\omega(1)}$.*

PROOF. We proceed by induction on i . The base case $i = 0$ is established by setting $c > 3^4 \cdot e$. We next upper bound the probability that after inserting n blocks there are more than β_{i+1} filled nodes at level $i+1$, given that we start with H_i nodes filled at level i . To do so, we define X_j to be that 0/1 random variable that is 1 iff the j -th block ends up at level $i+1$. Clearly, $\Pr[X_j = 1] = \left(\frac{H_i \cdot 2^{i+1}}{n}\right)^2$ and, if $H_i \leq \beta_i$,

$$\mu_{i+1} := \mathbb{E} \left[\sum_j X_j \right] \leq \frac{\beta_i^2 \cdot 2^{2(i+1)}}{n} = \frac{\beta_{i+1}}{e}.$$

We thus have

$$\begin{aligned} \Pr[H_{i+1} > \beta_{i+1} \mid H_i \leq \beta_i] &\leq \frac{\Pr[H_{i+1} > \beta_{i+1}]}{\Pr[H_i \leq \beta_i]} \\ &= \frac{\Pr[\sum_j X_j > e \cdot \mu_{i+1}]}{\Pr[H_i \leq \beta_i]} \leq \frac{e^{-e \cdot \beta_{i+1}}}{\Pr[H_i \leq \beta_i]} \end{aligned}$$

where the last inequality follows by Theorem A.2. By Lemma 7.3, the sequence β_i decreases with i and thus $\beta_i \geq \beta_{i+1} = \omega(\log n)$. The proof is then completed by observing that

$$\begin{aligned} \Pr[H_{i+1} > \beta_{i+1}] &\leq \Pr[H_{i+1} > \beta_{i+1} \mid H_i \leq \beta_i] \cdot \Pr[H_i \leq \beta_i] + \Pr[H_i > \beta_i] \\ &\leq 1/n^{\omega(1)} + i/n^{\omega(1)} \leq (i+1)/n^{\omega(1)} \end{aligned}$$

using our hypothesis that $\Pr[H_i > \beta_i] = i/n^{\omega(1)}$. \square

We now conclude the proof of Theorem 7.2. Fix a function $\Phi(n) = \omega(\log n)$ and let i^* be the largest index such that $\beta_{i^*} \geq \Phi(n)$. By Lemma 7.3, we obtain that $i^* = \Theta(\log \log n)$. Following a reasoning similar to the one adopted in proof of Lemma 7.4, we can prove that the expected number of filled nodes at level i^*+1 is at most β_{i^*+1}/e , given that no more than β_{i^*} nodes are filled at level i^* . By Lemma 7.4, the condition holds except with negligible probability. From the definition of β_{i^*+1} , we obtain that, for some constant α such that $c \geq \alpha \cdot \beta_{i^*+1}$, the probability that more than c nodes are filled at level i^*+1 is inverse exponential in c by Chernoff Bounds, and thus $\text{negl}(n)$. \square

7.3 Wrapping up the Proof

We complete our DP-KVS construction by observing that the mapping scheme described above has $k(n) = 2$, $s(n) = \Theta(\log \log n)$ and $c(n) = \Phi(n)$ for any $\Phi(n) = \omega(\log n)$.

THEOREM 7.5. *The above KVS scheme is a ϵ -DP-KVS with $\epsilon = O(\log n)$ that returns $O(\log \log n)$ blocks. The server uses $O(n)$ blocks of storage and, for any function $\Phi(n) = \omega(\log n)$,*

the probability that the client stores more than $O(\Phi(n) \cdot \log \log n)$ blocks is $\text{negl}(n)$.

8 CONCLUSIONS

We consider privacy-preserving storage protocols with small overhead that could be implemented with large-scale, frequently accessed storage infrastructures without negatively impacting response times or resource costs. Our main question is to find the best privacy that can be achieved by small overhead storage schemes.

We formulate our privacy notion using differentially private access, which is a generalization of the oblivious access provided by both ORAM and PIR. We present strong evidence that the best privacy achievable by any constant overhead differentially private storage schemes must have privacy budgets $\epsilon = \Omega(\log n)$. For DP-RAM and DP-IR, we present constructions with asymptotically optimal $\epsilon = \Theta(\log n)$ privacy budgets and $O(1)$ overhead. For DP-KVS, we present a scheme with asymptotically optimal $\epsilon = \Theta(\log n)$ privacy budgets and only $O(\log \log n)$ overhead which is exponentially better than previous constructions. Our DP-KVS uses a novel, oblivious variant of two-choice hashing that uses only $O(n)$ server storage that may be of independent interest.

Therefore, we answer that the best privacy achievable by privacy-preserving storage systems with small overhead is differentially private access with $\epsilon = \Theta(\log n)$. On the other hand, any storage scheme achieving stronger privacy most likely must incur non-trivial overhead compared to plaintext access.

REFERENCES

- [1] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang. Deep learning with differential privacy. In *CCS*, pages 308–318, 2016.
- [2] C. Aguilar-Melchor, J. Barrier, L. Fousse, and M.-O. Killijian. XPIR: Private information retrieval for everyone. *Proceedings on Privacy Enhancing Technologies*, 2016(2):155–174, 2016.
- [3] M. Ajtai, J. Komlós, and E. Szemerédi. An $O(N \log N)$ sorting network. In *Proceedings of the Fifteenth Annual ACM Symposium on Theory of Computing*, STOC '83, pages 1–9. ACM, 1983.
- [4] S. Angel, H. Chen, K. Laine, and S. Setty. PIR with compressed queries and amortized query processing.
- [5] Y. Azar, A. Z. Broder, A. R. Karlin, and E. Upfal. Balanced allocations. *SICOMP*, 29(1):180–200, 1999.
- [6] K. E. Batcher. Sorting networks and their applications. In *Proceedings of the April 30–May 2, 1968, Spring Joint Computer Conference*, AFIPS '68 (Spring), pages 307–314. ACM, 1968.
- [7] A. Beimel, Y. Ishai, and T. Malkin. Reducing the servers computation in private information retrieval: PIR with preprocessing. In *CRYPTO*, pages 55–73, 2000.
- [8] N. Borisov, G. Danezis, and I. Goldberg. DP5: A private presence service. *POPET*, 2015(2):4–24, 2015.
- [9] E. Boyle, K.-M. Chung, and R. Pass. Oblivious parallel RAM and applications. In *TCC*, pages 175–204, 2016.
- [10] E. Boyle, Y. Ishai, R. Pass, and M. Wootters. Can we access a database both locally and privately? In *TCC*, pages 662–693, 2017.
- [11] E. Boyle and M. Naor. Is there an oblivious RAM lower bound? In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*, pages 357–368. ACM, 2016.
- [12] C. Cachin, S. Micali, and M. Stadler. Computationally private information retrieval with polylogarithmic communication. In *EUROCRYPT*, pages 402–414, 1999.
- [13] R. Canetti, J. Holmgren, and S. Richelson. Towards doubly efficient private information retrieval. In *TCC*, pages 694–726, 2017.
- [14] D. Cash, P. Grubbs, J. Perry, and T. Ristenpart. Leakage-abuse attacks against searchable encryption. In *CCS*, pages 668–679, 2015.
- [15] T. H. Chan, K.-M. Chung, B. Maggs, and E. Shi. Foundations of differentially oblivious algorithms. Eprint 2017/1033, 2017.
- [16] T.-H. H. Chan, Y. Guo, W.-K. Lin, and E. Shi. Oblivious hashing revisited, and applications to asymptotically efficient ORAM and OPRAM. In *ASIACRYPT*, pages 660–690, 2017.
- [17] B. Chen, H. Lin, and S. Tessaro. Oblivious parallel RAM: improved efficiency and generic constructions. In *TCC*, pages 205–234, 2016.
- [18] R. Cheng, W. Scott, B. Parno, A. Krishnamurthy, and T. Anderson. Talek: A private publish-subscribe protocol. Technical report, Technical Report. University of Washington, 2016.
- [19] B. Chor, O. Goldreich, E. Kushilevitz, and M. Sudan. Private information retrieval. In *FOCS*, pages 41–50. IEEE, 1995.
- [20] S. Devadas, M. van Dijk, C. W. Fletcher, L. Ren, E. Shi, and D. Wichs. Onion ORAM: A constant bandwidth blowup oblivious RAM. In *TCC*, pages 145–174, 2016.
- [21] D. Dubhashi and D. Ranjan. Balls and bins: A study in negative dependence. *Random Struct. Algorithms*, 13(2):99–124, 1998.
- [22] C. Dwork. A firm foundation for private data analysis. *Communications of the ACM*, 54(1):86–95, 2011.
- [23] C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *TCC*, pages 265–284, 2006.
- [24] C. Dwork and A. Roth. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3–4):211–407, 2014.
- [25] C. Gentry and Z. Ramzan. Single-database private information retrieval with constant communication rate. In *ICALP*, pages 803–815, 2005.
- [26] O. Goldreich. Towards a theory of software protection and simulation by oblivious RAMs. In *STOC*, pages 182–194, 1987.
- [27] O. Goldreich and R. Ostrovsky. Software protection and simulation on oblivious RAMs. *Journal of the ACM*, 43:431–473, 1996.
- [28] M. T. Goodrich and M. Mitzenmacher. Privacy-preserving access of outsourced data via oblivious RAM simulation. In *ICALP*, pages 576–587, 2011.
- [29] M. T. Goodrich, M. Mitzenmacher, O. Ohrimenko, and R. Tamassia. Practical oblivious storage. In *Proceedings of the second ACM conference on Data and Application Security and Privacy*, pages 13–24. ACM, 2012.
- [30] M. Green, W. Ladd, and I. Miers. A protocol for privately reporting ad impressions at scale. In *CCS*, pages 1591–1601, 2016.
- [31] A. Hamlin, R. Ostrovsky, M. Weiss, and D. Wichs. Private anonymous data access. ePrint Report 2018/363, 2018.
- [32] M. S. Islam, M. Kuzu, and M. Kantarcioglu. Access pattern disclosure on searchable encryption: Ramification, attack and mitigation. In *NDSS*, volume 20, page 12, 2012.
- [33] P. Kairouz, S. Oh, and P. Viswanath. Secure multi-party differential privacy. In *Advances in Neural Information Processing Systems 28*, pages 2008–2016. 2015.
- [34] G. Kellaris, G. Kollios, K. Nissim, and A. O'Neill. Accessing data while preserving privacy. *arXiv preprint arXiv:1706.01552*, 2017.
- [35] E. Kushilevitz, S. Lu, and R. Ostrovsky. On the (in) security of hash-based oblivious RAM and a new balancing scheme. In *SODA*, pages

143–156, 2012.

- [36] E. Kushilevitz and R. Ostrovsky. Replication is not needed: Single database, computationally-private information retrieval. In *FOCS*, pages 364–373. IEEE, 1997.
- [37] K. G. Larsen and J. B. Nielsen. Yes, there is an oblivious RAM lower bound! In *CRYPTO*, 2018.
- [38] T. Mayberry, E.-O. Blass, and G. Noubir. Multi-client oblivious RAM secure against malicious servers. Cryptology ePrint Archive, Report 2015/121, 2015. <https://eprint.iacr.org/2015/121>.
- [39] S. Mazloom and S. D. Gordon. Differentially private access patterns in secure computation. ePrint Report 2017/1016, 2017.
- [40] I. Mironov, O. Pandey, O. Reingold, and S. Vadhan. Computational differential privacy. In *CRYPTO*, pages 126–142. 2009.
- [41] M. Mitzenmacher. The power of two choices in randomized load balancing. *IEEE Transactions on Parallel and Distributed Systems*, 12(10):1094–1104, 2001.
- [42] M. Mitzenmacher, A. W. Richa, and R. Sitaraman. The power of two random choices: A survey of techniques and results. In *Handbook of Randomized Computing*, pages 255–312, 2000.
- [43] O. Ohrimenko, M. T. Goodrich, R. Tamassia, and E. Uffal. The Melbourne Shuffle: Improving oblivious storage in the cloud. In *ICALP*, pages 556–567, 2014.
- [44] S. Patel, G. Persiano, M. Raykova, and K. Yeo. PanORAMa: Oblivious RAM with logarithmic overhead. In *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS*, pages 871–882, 2018.
- [45] S. Patel, G. Persiano, and K. Yeo. CacheShuffle: A family of oblivious shuffles. In *45th International Colloquium on Automata, Languages, and Programming (ICALP 2018)*, 2018.
- [46] S. Patel, G. Persiano, and K. Yeo. Private stateful information retrieval. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, CCS '18*. ACM, 2018.
- [47] G. Persiano and K. Yeo. Lower bounds for differentially private RAMs. Cryptology ePrint Archive, Report 2018/1051, 2018.
- [48] E. Stefanov, M. Van Dijk, E. Shi, C. Fletcher, L. Ren, X. Yu, and S. Devadas. Path ORAM: an extremely simple oblivious RAM protocol. In *CCS*, pages 299–310.
- [49] R. R. Toledo, G. Danezis, and I. Goldberg. Lower-cost ϵ -private information retrieval. *Proceedings on Privacy Enhancing Technologies*, 2016(4):184–201, 2016.
- [50] S. Wagh, P. Cuff, and P. Mittal. Root ORAM: A tunable differentially private oblivious RAM. *arXiv:1601.03378*, 2016.
- [51] A. Waksman. A permutation network. *J. ACM*, 15(1):159–163, Jan. 1968.
- [52] Y. Wang, Y.-X. Wang, and A. Singh. Differentially private subspace clustering. In *Advances in Neural Information Processing Systems 28*, pages 1000–1008. 2015.

A TOOLS

In this section we briefly review some tools that we use in the design and in the analysis of our constructions.

A.1 Power of Two Choices

The *power of two choices* concept was motivated from the classical *balls and bins* problem. The balls and bins concepts considers n balls and n bins. Each of the n balls chooses a single bin independently and uniformly at random. The *load* of a bin is the number of balls that occupy the bin. It has been shown that with high probability, the load of every bin does not exceed $O(\log n / \log \log n)$ [21]. Consider the case

where each of the n balls now chooses two bins independently and uniformly at random. The ball occupies the least loaded of the two chosen bins. This slight alteration ensures that with high probability, the load of each bin does not exceed $O(\log \log n)$ [41]. This result demonstrates that having two choices significantly improves bounds on the maximum load. Furthermore, it turns out that increasing the number of choices to $d \geq 3$ only improves the maximum load bounds by a constant. It is important to note that the allocation of each is chosen independently from the allocation of all other balls.

THEOREM A.1. *At any time, the load of any bin produced by the power of two choices process exceeds $O(\log \log n)$ with probability at most $1/n^{\Omega(\log \log n)}$.*

A.2 Chernoff Bound

The next theorem gives a bound on the tails of a binomial distribution that we will use to analyze our constructions. See [42] for a proof.

THEOREM A.2. *Let X_i , for $i = 1, \dots, n$ be independent binary random variables with $\Pr[X_i = 1] = p$ and let $\mu := np$. Then for every $t \geq \mu$, it holds that*

$$\Pr \left[\sum_{i=1}^n X_i \geq t \right] \leq \frac{\mu^t}{t^t} \cdot e^{t-\mu}$$

and, in particular,

$$\Pr \left[\sum_{i=1}^n X_i \geq e \cdot \mu \right] \leq e^{-\mu}.$$

B PROOF OF THEOREM 5.1

We remind the reader that, for $i \in [N]$, $\text{IR}(i)$ is the random variable of the set of blocks transferred by the server when the client wishes to access block B_i . We next compute $\Pr[\text{IR}(i) = \mathcal{T}]$ for a subset \mathcal{T} of K blocks.

Case 1: $B_i \in \mathcal{T}$. With probability $1 - \alpha$, we know that $B_i \in \mathcal{T}$ and the remaining $K - 1$ blocks are chosen uniformly at random. On the other hand, with probability α , for $B_i \in \mathcal{T}$, we need to choose B_i as one of the K blocks that are chosen randomly. Therefore,

$$\Pr[\text{IR}(i) = \mathcal{T}] = \frac{1 - \alpha}{\binom{N-1}{K-1}} + \frac{\alpha}{\binom{N}{K}}.$$

Case 2: $B_i \notin \mathcal{T}$. In this case, we know with probability $(1 - \alpha)$, \mathcal{T} is not possible. With probability α , we need to ensure that B_i is not one of the K blocks chosen. So,

$$\Pr[\text{IR}(i) = \mathcal{T}] = \frac{\alpha}{\binom{N}{K}}.$$

Let Q and Q' be any two query sequences of length L for which $d_H(Q, Q') = 1$. Since, the above algorithm is stateless,

we know that its behavior depends solely on the query index. Therefore, it suffices to consider two different sequences queries of length one, $q \neq q' \in [N]$. If we choose a transcript \mathcal{T} such that $B_q, B_{q'} \notin \mathcal{T}$ or $B_q, B_{q'} \in \mathcal{T}$, then we see that $\Pr[\text{IR}(q) = \mathcal{T}] = \Pr[\text{IR}(q') = \mathcal{T}]$. If, instead, $B_q \in \mathcal{T}$ and $B_{q'} \notin \mathcal{T}$ we have

$$\frac{\Pr[\text{IR}(q) = \mathcal{T}]}{\Pr[\text{IR}(q') = \mathcal{T}]} \leq \frac{(1-\alpha)\binom{N}{K}}{\alpha\binom{N-1}{K-1}} + 1 \leq \frac{(1-\alpha)N}{\alpha K} + 1 = e^\epsilon.$$

Since this holds for any single transcript \mathcal{T} , the same ratio holds for any set of transcripts completing the proof.

C LOWER BOUND FOR MULTIPLE-SERVER DP-IR

We extend the single server lower bound to the multiple server IR model with D servers and D_A adversarial servers. If $D_A = D$, then the scenario collapses to the single adversarial server case of Section 3.2. In this section, we assume that $D_A < D$ and there is always at least one honest server. For convenience, we define $t = \frac{D_A}{D}$ to be the fraction of servers corrupted by \mathcal{A} where $0 < t < 1$.

The choice of adversarial servers is modeled as a challenger C with a IR protocol and an adversary \mathcal{A} with the power to corrupt t fraction of servers. Using the knowledge of IR but not the internal randomness of IR, \mathcal{A} picks t fraction of the D servers to corrupt. C then runs IR and \mathcal{A} gets the transcript of downloads sent by IR to the t fraction of corrupted servers. We denote the adversary's transcript on query sequence Q as $\text{IR}^{\mathcal{A}}(Q)$. We now proceed to present our lower bound in the multi-server model.

THEOREM C.1. *If IR is a (ϵ, δ) -DP D -server IR in the balls and bins model for any $\epsilon, \delta \geq 0$ and error probability $\alpha < 1 - (\delta/t)$, then IR performs*

$$\Omega\left(\frac{((1-\alpha)t - \delta) \cdot n}{e^\epsilon}\right)$$

expected operations.

PROOF. We prove our lower bound for the adversary \mathcal{A} that corrupts $D_A \leq tD$ randomly and uniformly chosen servers. We know

$$\Pr[B_i \in \text{IR}(i)] \geq 1 - \alpha.$$

If B_i appears in $\text{IR}^{\mathcal{A}}(i)$, $\Pr[B_i \in \text{IR}^{\mathcal{A}}(i)] \geq (1-\alpha)t$. The rest of the proof follows identically to the proof of Theorem 3.4. \square

As a result, we show that the DP-IR scheme in [49] is optimal when t is considered constant.

D REMAINING DP-RAM PROOFS

In this section we give the proofs that were postponed from Section 6. We start by showing that the client uses $O(c)$ memory except with negligible probability for $c = \omega(\log n)$.

LEMMA D.1. *If $p \leq \frac{c}{n}$ and $c = \omega(\log n)$, then RAM stores at most $O(c)$ blocks in the block stash on client storage at any point in time except with negligible probability.*

PROOF. Pick any point in time. We note that each block B_i has independent probability p of being stored in the client. We let $X_i = 1$ if and only if B_i is stored on the client memory. Let X be the number of blocks stored in client memory, so $X = X_1 + \dots + X_n$. So, we know that $\mathbb{E}[X] = \mathbb{E}[X_1 + \dots + X_n] = pn \leq c$. By Chernoff Bounds, for any $\delta > 0$,

$$\Pr[X > (1 + \delta)c] \leq \exp\left(\frac{-c\delta^2}{2 + \delta}\right)$$

which is $\text{negl}(n)$ when $c = \omega(\log n)$. The lemma follows, by a union bound over all points in time (that is, the length of the queries which is polynomial in n). \square

PROOF OF LEMMA 6.4. The numerator is trivially upper bounded by 1 and the denominator is lower bounded by the conjunctive probability. We prove the lemma by showing that the latter is at least p/n^2 . To this aim, we consider three cases depending on the values of $\text{pr}(Q, j)$ and $o_{\text{pr}(Q, j)}$. Note that, if $\text{pr}(Q, j) \neq +\infty$, $q_{\text{pr}(Q, j)} = q_j$ by definition.

Case 1: $\text{pr}(Q, j) \neq +\infty$ and $o_{\text{pr}(Q, j)} = q_j$.

Consider first the case $d_j = q_j$. If during the overwrite phase of query $\text{pr}(Q, j)$ block B_{q_j} is stored in block stash (and this has probability p), then $\text{RAM}_{\text{pr}(Q, j)}^O(Q) = q_j$ with probability $1/n$. Then at the download phase of query j , B_{q_j} is found in block stash and therefore the probability that $\text{RAM}_j^D(Q) = q_j$ is $1/n$. On the other hand if during the overwrite phase of query $\text{pr}(Q, j)$ block B_{q_j} is not stored in block stash (this has probability $1 - p$), then certainly $\text{RAM}_{\text{pr}(Q, j)}^O(Q) = q_j$ and, since B_{q_j} is not found in block stash at the start of the download phase of query j , $\text{RAM}_j^D(Q) = q_j$ with probability 1. Altogether, we have $\Pr[\text{RAM}_{\text{pr}(Q, j)}^O(Q) = q_j \wedge \text{RAM}_j^D(Q) = q_j] = 1 - p + 1/n^3 \geq 1/n^3$. If, instead, $d_j \neq q_j$ then it means that B_{q_j} must have been stored in the block stash during the overwrite phase of query $\text{pr}(Q, j)$ (this has probability p) and then $\text{RAM}_{\text{pr}(Q, j)}^O(Q) = q_j$ with probability $1/n$. Then, since B_{q_j} is in stash at the start of the download phase of query j $\text{RAM}_j^D(Q) = d_j$ with probability $1/n$. Altogether, we have $\Pr[\text{RAM}_{\text{pr}(Q, j)}^O(Q) = q_j \wedge \text{RAM}_j^D(Q) = q_j] = p/n^2$.

Case 2. $\text{pr}(Q, j) \neq +\infty$ and $o_{\text{pr}(Q, j)} \neq q_j$.

If $o_{\text{pr}(Q, j)} \neq q_j$ then it must be the case that, during the overwrite phase of query $\text{pr}(Q, j)$, block B_{q_j} is stored in block

stash (and this happens with probability p) and $o_{\text{pr}(Q,j)}$ is chosen as overwrite block (and this happens with probability $1/n$). When the download phase of query j starts, block B_{q_j} is found in the stash and so the probability that d_j is selected as download blocks is $1/n$. Whence $\Pr[\text{RAM}_j^D(Q) = d_j \wedge \text{RAM}_{\text{pr}(Q,j)}^O(Q) = o_{\text{pr}(Q,j)}] = p/n^2$.

Case 3. $\text{pr}(Q, j) = +\infty$. In this case, $o_{\text{pr}(Q,j)} = \perp$ and B_{q_j} has not been queried before query j . Therefore, at the start of the download phase of query j , B_{q_j} is found in block stash with probability p . By using arguments similar to the ones of the previous cases, we have

$$\begin{aligned} \Pr[\text{RAM}_j^D(Q) = q_j \wedge \text{RAM}_{\text{pr}(Q,j)}^O(Q) = \perp] \\ = \Pr[\text{RAM}_j^D(Q) = q_j] = (1-p) + p/n \end{aligned}$$

and, for $d_j \neq q_j$,

$$\begin{aligned} \Pr[\text{RAM}_j^D(Q) = d_j \wedge \text{RAM}_{\text{pr}(Q,j)}^O(Q) = \perp] \\ = \Pr[\text{RAM}_j^D(Q) = d_j] = p/n \end{aligned}$$

which completes the proof. \square

PROOF OF LEMMA 6.5. As in the proof of the previous lemma, we upper bound the numerator by 1 and show that the denominator is at least n/p starting from the case $o_j = q_j$. If, in the overwrite phase of the j -th query, block B_{q_j} is added to the block stash (and this happens with probability $1/n$) then the same block is chosen as overwrite block with probability $1/n$. In the remaining $1-p$ probability, q_j is the overwrite block with probability 1. Therefore $\Pr[\text{RAM}_j^O(Q) = q_j] = (1-p) + p/n$. If, instead, $o_j \neq q_j$, then the only case in which o_j is the overwrite block of phase j is when B_{q_j} is added to the block stash and this gives $\Pr[\text{RAM}_j^O(Q) = o_j] = p/n$, for every $o_j \neq q_j$. For any $j \neq i$, $Q_j = Q'_j$ showing the first point. For any o , $\Pr[\text{RAM}_i^O(Q') = o] \geq p/n$ giving the second result. \square

E DP-RAM GENERALIZATION

Our proof of differential privacy of DP-RAM of Section 6 can be seen to give a more general form of differential privacy. Suppose DP-RAM stores n blocks and there exists a repertoire Σ of size $b = O(n)$. We interpret Σ as the specification of b buckets where each bucket contains exactly s blocks. Note, buckets may overlap and two different buckets may contain the same block. Our privacy proof of DP-RAM carries over to the case in which a query retrieves all blocks in a bucket. In other words, we interpret query sequences Q of length $s \cdot l$ as l subsequences each of length s taken from the Σ of size n (that is, each subsequence corresponds to a bucket). In this framework, a query sequence $Q = (q_1, \dots, q_{s \cdot l})$ is associated

with sequence $Q = (\sigma_1, \dots, \sigma_l)$ of length l over Σ ; that is,

$$Q = (\underbrace{q_1, \dots, q_s}_{\sigma_1}, \underbrace{q_{s+1}, \dots, q_{2s}}_{\sigma_2}, \dots, \underbrace{q_{(l-1)s+1}, \dots, q_{l \cdot s}}_{\sigma_l}),$$

and $\sigma_1, \sigma_2, \dots, \sigma_l \in \Sigma$ and $|\Sigma| = b = O(n)$.

The proof of Section 6 handles the case $s = 1$ in which an access sequence is a sequence of n requests taken from the repertoire $\Sigma = [n]$.

To argue the general case let us consider first the case in which the server explicitly stores the subsequences of Σ ; specifically, each element of Σ is a sequence of s of the original n blocks and each is stored using storage equal to s blocks. Then an access sequence Q of length $s \cdot l$ over $[n]$ is simply an access sequence $(\sigma_1, \dots, \sigma_l)$ of length l over Σ and, as it is easily seen, our proof still works and guarantees differential privacy with $\epsilon = \Theta(\log n)$.

This approach, unfortunately, has the drawback that the server storage grows by a factor of s . However, we observe that the blocks that constitute the subsequence of Σ need not to be explicitly stored by the server that instead can store just the n original blocks and, each time it receives a request $\sigma \in \Sigma$, the server fetches the needed blocks. Note that this transformation preserves differential privacy with $\epsilon = \Theta(\log n)$ as this property is independent of the actual layout of the parts of the individual atomic blocks (in this case the subsequences $\sigma \in \Sigma$) on which the RAM is built. The only modification required is that when retrieving or updating a bucket i , the DP-RAM must also check if the any block is stored on the client as part of another bucket $j \neq i$. If so, the block on client storage must be returned as opposed to the one from the server for retrievals. For updates, both the server copy and client copy must be updated.

F REMAINING DP-KVS PROOFS

[Proof of Lemma 7.3] We proceed by induction and we verify the base case by plugging in $i = 0$. For $i \geq 0$, we have

$$\begin{aligned} \beta_{i+1} &= \frac{e}{n} \cdot \beta_i^2 \cdot 2^{2(i+1)} \\ &= \frac{e}{n} \cdot \frac{n^2}{e^2} \cdot \left(\frac{2}{3}\right)^{2^{i+3}} \left(\frac{1}{2}\right)^{4(i+2)} \cdot 2^{2(i+1)} \\ &= \frac{n}{e} \cdot \left(\frac{2}{3}\right)^{2^{i+3}} \left(\frac{1}{2}\right)^{4(i+2)-2(i+1)} \\ &= \frac{n}{e} \cdot \left(\frac{2}{3}\right)^{2^{i+3}} \left(\frac{1}{2}\right)^{2(i+3)}. \end{aligned}$$

G DP-IR PSEUDOCODE

We present the pseudocode for our DP-IR scheme of Section 5.

Algorithm 1 DP-IR.Query: read a block

Input: i, ϵ, α
Set $\mathcal{T} = \emptyset$. {initializing the download set}
Pick random number $r \in [0, 1]$.
if $r > \alpha$ **then**
 $\mathcal{T} \leftarrow \mathcal{T} \cup \{i\}$.
end if
Set $K = \lceil (1 - \alpha)N / (e^\epsilon - 1) \rceil$.
while $|\mathcal{T}| < K$ **do**
 Pick j uniformly at random from $[N] \setminus \mathcal{T}$.
 $\mathcal{T} \leftarrow \mathcal{T} \cup \{j\}$.
end while
Send \mathcal{T} to server and receive $\{B_j\}_{j \in \mathcal{T}}$ from server.
if $r > \alpha$ **then**
 return B_i .
else
 return \perp .
end if

H DP-RAM PSEUDOCODE

We present the pseudocode for the algorithm in Section 6.

Algorithm 2 DP-RAMSetup: initialize client and server storage

Input: $1^\lambda, B_1, \dots, B_N$
Randomly select $K \leftarrow \{0, 1\}^\lambda$.
Initialize array A of size N on server.
Initialize hash table bStash on client.
for $i \leftarrow 1, \dots, N$ **do**
 Set $A[i] \leftarrow \text{Enc}(K, B_i)$.
 Pick r uniformly at random from $[N]$.
 if $r \leq C$ **then**
 Set $\text{bStash}[i] \leftarrow B_i$.
 end if
end for

Algorithm 3 DP-RAM.Query: reading and writing a data block

Input: $i, \text{op}, B_{\text{new}}$
{Download Phase}
if $i \in \text{bStash}$ **then**
 Pick d uniformly at random from $[N]$.
 Download $A[d]$ and discard.
 Set $B \leftarrow \text{bStash}[i]$ and remove i from bStash.
else
 Set $d \leftarrow i$.
 Download $A[d]$.
 Set $B \leftarrow \text{Dec}(K, A[d])$.
end if

if $\text{op} = \text{write}$ **then**
 Set $B \leftarrow B_{\text{new}}$.
end if

{Overwrite Phase}
Draw r uniformly at random from $[N]$.
if $r \leq C$ **then**
 Set $\text{bStash}[i] \leftarrow B$.
 Pick o uniformly at random from $[N]$.
 Download $A[o]$, decrypt it and re-encrypt it obtaining ct .

 Upload ct as $A[o]$.
else
 Set $o = i$.
 Download $A[o]$ and discard it.
 Set $\text{ct} \leftarrow \text{Enc}(K, B)$.
 Upload ct as $A[o]$.
end if
