

Lower Bounds for Differentially Private RAMs

Giuseppe Persiano^{1,2} and Kevin Yeo¹

¹ Google LLC, Mountain View, USA kwlyeo@google.com

² Università di Salerno, Salerno, Italy giuper@gmail.com

Abstract. In this work, we study privacy-preserving storage primitives that are suitable for use in data analysis on outsourced databases within the differential privacy framework. The goal in differentially private data analysis is to disclose global properties of a group without compromising any individual’s privacy. Typically, differentially private adversaries only ever learn global properties. For the case of outsourced databases, the adversary also views the patterns of access to data. Oblivious RAM (ORAM) can be used to hide access patterns but ORAM might be excessive as in some settings it could be sufficient to be compatible with differential privacy and only protect the privacy of individual accesses. We consider (ϵ, δ) -Differentially Private RAM, a weakening of ORAM that only protects individual operations and seems better suited for use in data analysis on outsourced databases. As differentially private RAM has weaker security than ORAM, there is hope that we can bypass the $\Omega(\log(n/c))$ bandwidth lower bounds for ORAM by Larsen and Nielsen [CRYPTO ’18] for storing an array of n entries and a client with c bits of memory. We answer in the negative and present an $\Omega(\log(n/c))$ bandwidth lower bound for privacy budgets of $\epsilon = O(1)$ and $\delta \leq 1/3$. The *information transfer* technique used for ORAM lower bounds does not seem adaptable for use with the weaker security guarantees of differential privacy. Instead, we prove our lower bounds by adapting the *chronogram* technique to our setting. To our knowledge, this is the first work that uses the chronogram technique for lower bounds on privacy-preserving storage primitives.

1 Introduction

In this work, we study *privacy-preserving storage* schemes involving a client and an untrusted server. The goal is to enable the client to outsource the storage of data to the server such that the client may still perform operations on the stored data (e.g. retrieving and updating data). For privacy, the client wishes to keep the stored data hidden from server. One way to ensure the contents of the data remain hidden is for the client to encrypt all data before uploading to the server. However, the server can still view how the encrypted data is accessed as the client performs operations. Previous works such as [19, 3] have shown that the leakage of patterns of access to encrypted data can be used to compromise the privacy of the encrypted data. Therefore, a very important privacy requirement is also to protect the access patterns.

The traditional way to define the privacy of access pattern is *obliviousness*. An oblivious storage primitive ensures that any adversary that is given two sequences of operations of equal length and observes the patterns of data access performed by one of the two sequences cannot determine which of the two sequences induced the observed access pattern. The most famous oblivious storage primitive is Oblivious RAM (ORAM) that outsources the storage of an array and allows clients to retrieve and update array entries. ORAM was first introduced by Goldreich [15] who presented an ORAM with sublinear amortized bandwidth per operation for clients with constant size memory. Goldreich and Ostrovsky [16] give the first ORAM construction with polylogarithmic amortized bandwidth per operation. In the past decade, ORAM has been the subject of extensive research [27, 17, 30, 18, 20, 31, 26] as well as variants such as statistically secure ORAMs [7, 6], parallel ORAMs [1, 5, 4] and garbled RAMs [14, 13, 24]. The above references are just a small subset of all the results for ORAM constructions.

Instead, we focus on a different definition of privacy using *differential privacy* [10, 9, 11]. The representative scenario for differential privacy is *privacy-preserving data analysis* which considers the problem of disclosing properties about an entire database while maintaining the privacy of individual database records. A mechanism or algorithm is considered differentially private if any fixed disclosure is almost as likely to be outputted for two different input databases that only differ in exactly one record. As a result, an adversary that views the disclosure is unable to determine whether an individual record was part of the input used to compute the disclosure. We consider the scenario of performing privacy-preserving data analysis on data outsourced to an untrusted server. By viewing the patterns of access to the outsourced data, the adversarial server might be able to determine which individual records were used to compute the disclosure compromising differential privacy.

One way to protect the patterns of data access is to outsource the data using an ORAM. However, in many cases, it turns out that the obliviousness guarantees of ORAM may be stronger than required. For example, let's suppose that we wish to disclose a differentially private regression model over a sample of the outsourced data. ORAM guarantees that the identity of all sampled database records will remain hidden from the adversary. On the other hand, the differentially private regression model only provides privacy about whether an individual record was sampled or not. Instead of obliviousness, we want a weaker notion of privacy for access patterns suitable for use with differentially private data analytics. With a weaker notion of privacy, there is hope for a construction with better efficiency than ORAM.

With this in mind, we turn to the notion of *differentially private access* which provides privacy for individual operations but might reveal information about a sequence of many operations. Differentially private access has been previously considered in [32, 33]. In particular, this privacy notion ensures that the patterns of data access caused by a fixed sequence of operations is almost as likely to be induced by another sequence of operations of the same length with a single different operation. We define (ϵ, δ) -*differentially private RAM* as a stor-

age primitive that outsources the storage of an array in a manner that allows a client to retrieve and update array entries while providing differentially private access. As this privacy notion is weaker than obliviousness, the $\Omega(\log(n/c))$ lower bounds for ORAMs that store n array entries and clients with c bits of storage by Larsen and Nielsen [22] do not apply. There is hope to achieve a differentially private RAM construction with smaller bandwidth. In this work, we answer in the negative and show that an $\Omega(\log(n/c))$ bandwidth lower bound also exists for differentially private RAM for typical privacy budgets of $\epsilon = O(1)$ and $\delta \leq 1/3$. As differential privacy with budgets of $\epsilon = O(1)$ and $\delta \leq 1/3$ provide weaker security than obliviousness, any ORAM is also a differentially private RAM. Therefore, our lower bounds show that the ORAM construction by Patel *et al.* [26] is an asymptotically optimal, up to an $O(\log \log n)$ factor, differentially private RAM with $\epsilon = O(1)$ and $\delta \leq 1/3$.

1.1 Our Results

In this section, we will present our contributions. We first describe the scenarios where our lower bounds apply. Our lower bounds apply to *differentially private* RAMs that process operations in an *online* fashion. The RAM must be both *read-and-write*, that is, the set of permitted operations include both reading and writing array entries. The server that stores the array is assumed to be *passive* in that the server may not perform any computation beyond retrieving and overwriting cells but no assumptions are made on the *storage encoding* of the array. Finally, we assume that the adversary is *computationally bounded*. We now go into detail about each of these requirements.

Differential Privacy. The goal of differential privacy is to ensure that the removal or replacement of an individual in a large population does not significantly affect the view of the adversary. Differential privacy is parameterized by two values $0 \leq \epsilon, \delta \leq 1$. The value ϵ is typically referred to as the *privacy budget*. When $\delta = 0$, the notion is known as *pure* differential privacy while, if $\delta > 0$, the notion is known as *approximate* differential privacy. In our context, an *individual* is a single operation in a sequence (the *population*) of read (also called queries) and write (also called updates) operations over an array of n entries stored on a, potentially adversarial, remote server. For any implementation \mathbf{DS} and for any sequence Q , we define $\mathbb{V}_{\mathbf{DS}}(Q)$ to be the view of the server when sequence Q is executed by \mathbf{DS} . A differentially private RAM, \mathbf{DS} , is defined to ensure that the adversary’s view on one sequence of operations should not be significantly different when \mathbf{DS} executes another sequence of operations which differs for only one operation. We assume that our adversaries are computationally bounded.

Formally, if \mathbf{DS} is (ϵ, δ) -differentially private, then for any two sequences Q_1 and Q_2 that differ in exactly one operation, it must be that $\Pr[\mathcal{A}(\mathbb{V}_{\mathbf{DS}}(Q_1)) = 1] \leq e^\epsilon \Pr[\mathcal{A}(\mathbb{V}_{\mathbf{DS}}(Q_2)) = 1] + \delta$ for any probabilistically polynomial time (PPT) algorithm \mathcal{A} . The notion of computational differential privacy was studied by Mironov *et al.* [25] where various classes of privacy were described. Our lower bounds consider the weakest privacy class and, thus, apply to all privacy classes

in [25]. In the majority of scenarios, differential privacy is only considered useful for the cases when $\epsilon = O(1)$ and δ is negligible. This is exactly the scenario where our lower bounds will hold. In fact, our lower bounds hold for any $\delta \leq 1/3$. We note that differential privacy with $\epsilon = O(1)$ and $\delta \leq 1/3$ is a weaker security notion than obliviousness. Obliviousness is equivalent to differential privacy when $\epsilon = 0$ and δ is negligible. Therefore, our lower bounds also hold for ORAM and match the lower bounds of Larsen and Nielsen [22]. We refer the reader to Section 2 for a formal definition of differential privacy.

Online RAMs. It is important that we discuss the notion of *online* vs. *offline* processing of operations by RAMs. In the offline scenario, it is assumed that all operations are given before the RAM must start processing updates and answering queries. The first ORAM lower bound by Goldreich and Ostrovsky [16] considered offline ORAMs with “balls-and-bins” encoding and security against an all-powerful adversary. “Balls-and-bins” refers to the encoding where array entries are immutable balls and the only valid operation is to move array entries into various memory locations referred to as bins. Boyle and Naor [2] show that proving an offline ORAM lower bound for non-restricted encodings is equivalent to showing lower bounds in sorting circuits, which is a long-standing problem in complexity. Instead, we consider online RAMs where operations arrive one at a time and must be processed before receiving the next operation. The assumption of online operations is realistic as the majority of RAM constructions consider online operations and almost all applications of RAMs consider online operations. Our lower bounds only apply for online differentially private RAMs.

Read-and-Write RAMs. Traditionally, all ORAM results consider the scenario where the set of valid operations include both reading and writing array entries. A natural relaxation would be to consider *read-only* RAMs where the only valid operation is reading array entries. Any lower bound on read-only RAMs would also apply to read-and-write RAMs. However, in a recent work by Weiss and Wichs [35], it is shown that any lower bounds for read-only ORAMs would imply very strong lower bounds for either sorting circuits and/or locally decodable codes (LDCs). Proving lower bounds for LDCs has, like sorting circuits, been an open problem in the world of complexity theory for more than a decade. As differential privacy is weaker than obliviousness, any lower bounds on read-only, differentially private RAMs also imply lower bounds on read-only ORAMs. To get around these obstacles, our work focuses only on proving lower bounds for read-and-write differentially private RAMs.

Passive Server. In our work, we will assume that the server storing the array is *passive*, which means that the server will not any perform computation beyond retrieving and overwriting the contents of a cell as the request of the client. This assumption is necessary as there are ORAM constructions that use server computation to achieve constant bandwidth operations [8]. Therefore, our lower bounds on bandwidth only apply to differentially private RAMs with a passive server. For scenarios with differentially private RAMs and servers with no bounds

on computation, we can reinterpret our results as lower bounds on the amount of server computation required.

We now informally present our main contribution.

Theorem 1 (informal). *Let \mathbf{DS} be any online, read-and-write RAM that stores n array entries each of size b bits on a passive server without any restrictions on storage encodings. Suppose that the client has c bits of storage. Assume that \mathbf{DS} provides (ϵ, δ) -differential privacy against a computational adversary that views all cell probes performed by the server. If $\epsilon = O(1)$ and $0 \leq \delta \leq 1/3$, then the expected amortized bandwidth of both reading and writing array entries by \mathbf{DS} is $\Omega(b \log(n/c))$ bits or $\Omega(\log(n/c))$ array entries. In the natural scenario where $c \leq n^\alpha$ where $0 \leq \alpha < 1$, then $\Omega(\log n)$ array entries of bandwidth are required.*

We note that our lower bounds are tight up to an $O(\log \log n)$ factor since PanORAMA [26] requires only $O(\log n \log \log n)$ array entries of bandwidth whenever $b = \Omega(\log n)$ and any ORAM is also a differentially private RAM with $\epsilon = 0$ and negligible δ . Furthermore, when $b = \Omega(\log^2 n)$, Path ORAM [31] is tight.

1.2 Previous Works

In this section, we present a small survey of previous works on data structure lower bounds. We also describe the first lower bound for data structures that provide privacy guarantees.

The majority of data structure lower bounds are proved using the cell probe model introduced by Yao [36], which only charges for accessing memory and allows unlimited computation. In the case for passive servers that only retrieve and overwrite memory, the costs of the cell probe model directly imply costs in bandwidth. The *chronogram* technique was introduced by Fredman and Saks [12] which can be used to prove $\Omega(\log n / \log \log n)$ lower bounds. Pătraşcu and Demaine [29] presented the *information transfer* technique which could be used to prove $\Omega(\log n)$ lower bounds. Larsen [21] presented an $\tilde{\Omega}(\log^2 n)$ lower bound for two-dimensional dynamic range counting, which remains the highest lower bound proven for any $\log n$ output data structures. Recently, Larsen *et al.* [23] presented an $\tilde{\Omega}(\log^{1.5} n)$ lower bound for data structures with single bit outputs which is the highest lower bound for decision query data structures.

For ORAM, Goldreich and Ostrovsky [16] presented an $\Omega(\log_c n)$ lower bound for clients with storage of c array entries. However, Boyle and Naor [2] showed that this lower bound came with the caveats that the lower bound only for statistical adversaries and constructions in “balls-and-bins“ model where array entries could only be moved between memory and not encoded in a more complex manner. Furthermore, Boyle and Naor [2] show that proving lower bounds for offline ORAMs and arbitrary storage encodings imply sorting circuit lower bounds. In their seminal work, Larsen and Nielsen [22] presented an $\Omega(\log(n/c))$ bandwidth lower bound removing the caveats such that lower bounds applies to any types of storage encodings and computational adversaries. Recently, Weiss

and Wicks [35] show that lower bounds for online, read-only ORAMs would imply lower bounds for either sorting circuits and/or locally decodable codes.

We present a brief overview of the techniques used by Larsen and Nielsen [22], which uses the information transfer technique. We also describe why information transfer does not seem to be of use for differentially private RAM lower bounds. Information transfer first builds a binary tree over $\Theta(n)$ operations where the first operation is assigned to the leftmost leaf, the second operation is assigned to the second leftmost leaf and so forth. Each cell probe is assigned to at most one node of the tree as follows. For a cell probe, we identify the operation that is performing the probe as well as the most recent operation that overwrote the cell that is being probed. The cell probe is assigned to the lowest common ancestor of the leaves associated with the most recent operation to overwrite the cell and the operation performing the probe. Let us fix any node of the tree and consider the subtree rooted at the fixed node. It can be shown that the probes assigned to the root is the entirety of information that can be transferred from the updates of the left subtree to be used to answer queries in right subtree. Consider the sequence of operations where all leaves in the left subtree write a randomly chosen b -bit string to unique array entries and all leaves in the right subtree read an unique, updated array entry. For any **DS** to return the correct b -bit strings asked by the queries in the right subtree, a large amount of information must be transferred from the left subtree to the right subtree. Thus, many probes should be assigned to the root of this subtree. Suppose that for another sequence of operations, **DS** assigns significantly less probes to the root of this subtree. Then, a computational adversary can count the probes and distinguish between the worst case sequence and any other sequence contradicting obliviousness. As a result, there must be many probes assigned to each node of the information transfer tree. Each cell probe is assigned to at most one node. So, summing up the tree provides a lower bound on the number of cell probes required.

Unfortunately, we are unable to use the information transfer technique to prove lower bounds for differentially private RAMs. The main issue comes from the fact that differentially private RAMs have significantly weaker privacy guarantees compared to ORAMs. When $\epsilon = \Theta(1)$, the probabilistic requirements of the adversary's view when **DS** processes two sequences Q_1 and Q_2 degrade exponentially in the number of operations that Q_1 and Q_2 differ in. On the other hand, the privacy requirements of obliviousness do not degrade when considering two sequences that differ in many operations. Larsen and Nielsen [22] use obliviousness to argue that the adversary's view for the worst case sequence of any subtree cannot differ significantly from any other sequence. However, for any fixed sequence of operations, the worst case sequence for the majority of subtrees differ in many operations (on the order of the number of leaves of the subtree). Applying differential privacy will not yield strong requirements for the number of cell probes assigned to the majority of the nodes in the information transfer binary tree. As a result, we could not adapt the information transfer technique for differentially private RAM lower bounds and resort to other techniques.

1.3 Overview of our Proofs

In this section, we present an overview of the proof techniques used in Sections 3 and 4. Our lower bounds use ideas from works by Pătraşcu and Demaine [29] and Pătraşcu [28]. However, we begin by reviewing the original chronogram technique of Fredman and Saks [12]. Let us suppose that t_w is the running time of an update (**write**) and t_r is the running time of a query (**read**). Consider a sequence of operations with $\Theta(n)$ updates followed by a single query. Starting from the query and going backwards in time, updates are partitioned into exponentially increasing epochs at some rate r . Epochs are indexed in reverse time, so the smallest epoch closest to the query is epoch 1. The i -th epoch will have $\ell_i = r^i$ update operations. The goal of the chronogram is to prove that there exists a query that requires information from many of the epochs simultaneously. To do this, we first observe that if each update writes a randomly and independently chosen b -bit entry, an update operation preceding epoch i cannot encode any information about epoch i . Therefore, all information about epoch i can only be found in cells that have been written as part of the update operations of epoch i or any following epochs. Since each update stores b random bits, any epoch i encodes at least $\ell_i \cdot b$ bits in total. If we denote by w the number of bits in a cell and set $r = (t_w w)^2$, it is easy to see that the write operations of epochs $i - 1, i - 2, \dots, 1$ can probe at most $t_w(r^{i-1} + \dots + r) = O(\ell_i/(t_w w^2))$ cells and write at most $O(\ell_i/(t_w w))$ bits. As a result, the majority of the bits encoded by updates in epoch i remain in cells last written in epoch i . Finally, the goal is to construct a random query such that $\Omega(b)$ bits must be transferred from each epoch. If such a query exists, we can prove that $\max\{t_w, t_r\} = \Omega((b/w) \log_r n) = \Omega((b/w) \log n / \log \log n)$.

This lower bound can be improved to $\Omega((b/w) \log n)$ by using an improvement of the chronogram technique by Pătraşcu [28]. In the original chronogram technique, the epochs are fixed since the query's location and the number of updates are fixed. An algorithm may attempt to target an epoch i by having all future update operations encode information only about epoch i . To counteract this, we consider a harder update sequence where epoch locations cannot be predicted by the algorithm. Specifically, we consider a sequence that consists of a random number of update operations followed by a single query operation. For such a sequence, even if an algorithm attempts to target epoch i , it cannot pinpoint the location of epoch i and may only prepare over all possible query locations. We show that any update operation may now only encode $O(t_w w / \log_r n)$ about epoch i where $\log_r n$ is the number of epochs. As a result, future update operations can only encode a $O(1/\log_r n)$ fraction as much information about epoch i as the previous lower bound attempt. This allows us to fix $r = 2$ which increases the number of epochs $\log n$. If we can find a query that requires $\Omega(b)$ bits of information transfer from the majority of epochs, we can prove that $\max\{t_w, t_r\} = \Omega((b/w) \log n)$.

Recall that for differentially private RAM, the update operations enable overwriting a b -bit array entry while the query operations allows retrieving an array entry. We choose our update operations to overwrite unique array entries where

the updated array entry is chosen independently and uniformly at random from $\{0, 1\}^b$. Focus on an epoch i and consider picking a random query from the ℓ_i array indices updated in epoch i . The majority of these queries must read $\Omega(b)$ bits from cells last written in epoch i as future operations cannot encode all $\ell_i \cdot b$ bits encoded by epoch i . As a result, there exists some query such that $\Omega(b)$ bits must be transferred from epoch i for all sufficiently large epochs. We use differential privacy to show that $\Omega(b)$ bits must be transferred from all sufficiently large epochs. Consider two sequences of operations that only differ in the final query operation where one query requires $\Omega(b)$ bits from epoch i for correctness while the other query does not need any information from epoch i . If the latter query transfers $o(b)$ bits from epoch i , the adversary can distinguish between the two sequences with high probability contradicting differential privacy. Note, our privacy guarantees do not degrade significantly as the two sequences only differ at the final query operation. Therefore, we can prove that $\Omega(b)$ bits have to be transferred from most epochs and that $\max\{t_w, t_r\} = \Omega((b/w) \log n)$. The proof of this lower bound is found in Section 3.

A stronger lower bound is presented in Section 4 using more complex epoch constructions. We now describe our main result presented in Section 4. The lower bound outlined above shows that $\max\{t_w, t_r\} = \Omega((b/w) \log n)$ but does not preclude the case where $t_w = \Theta((b/w) \log n)$ and $t_r = O(1)$, for example. We show this cannot be the case. In particular, we show that if $\max\{t_w, t_r\} = O((b/w) \log n)$, then it must be the case that $t_w = \Theta((b/w) \log n)$ and $t_r = \Theta((b/w) \log n)$. The idea is to construct different epoch constructions for the cases when t_w and t_r are small respectively. When $t_w = o((b/w) \log n)$, we know that operations in future epochs cannot encode too much information. We consider an epoch construction where epochs grow by a rate of $r = \omega(1)$ every r epochs increasing the number of epochs to $\omega(\log n)$. In exchange, there are many operations after an epoch i . Since t_w is small, the future operations may not encode too much information about epoch i ensuring most of the information about epoch i remain in cells last written during epoch i . As a result, it can be shown again that $\Omega(b)$ bits must be read from many epochs implying an $t_r = \omega((b/w) \log n)$ lower bound.

On the other hand, consider the case when $t_r = o((b/w) \log n)$. We consider epoch constructions that increase exponentially with rate $r = \omega(1)$. As a result, the number of operations after epoch i is a factor of $O(1/r)$ smaller than the ℓ_i operations in epoch i and there are $\Theta(\log_r n)$ epochs. If $t_r = o((b/w) \log_r n)$, then a query operation may not read $\Omega(b)$ from each of the epochs. Instead, update operations must encode a large amount of information about previous epochs to compensate for t_r being so small. As a result, it can be shown that $t_w = \omega((b/w) \log n)$. Combining the above two statements implies that if $\max\{t_w, t_r\} = O((b/w) \log n)$, then $t_w = \Theta((b/w) \log n)$ and $t_r = \Theta((b/w) \log n)$.

2 Differentially Private Cell Probe Model

We start by formalizing the model for which we prove our lower bounds. We rely on the *cell probe model*, first described by Yao [36], and typically used to prove lower bounds for data structures without any requirements for privacy of the stored data and/or the operations performed. In a recent work by Larsen and Nielsen [22], the *oblivious cell probe model* was introduced and used to prove a lower bound for oblivious RAM. The oblivious cell probe model was defined for any data structures where the patterns of access to memory should not reveal any information about the operations performed. We generalize the oblivious cell probe model and present the (ϵ, δ) -*differentially private cell probe model*. In this new model, all data structures are assumed to provide differential privacy for the operations performed with respect to memory accesses viewed by the adversary. The differentially private cell probe model is a generalization of the oblivious cell probe model as obliviousness is equivalent to differential privacy with $\epsilon = 0$ and $\delta = \text{negl}(n)$, that is, any function negligible in the number of items stored in the data structure.

The cell probe model is an abstraction of the interaction between CPUs and word-RAM memory architectures. Memory is defined as an array of *cells* such that each cell contains exactly w bits. Any *operation* of a data structure is allowed to *probe* cells where a probe can consist of either reading the contents of a cell or overwriting the contents of a cell. The running time or cost for any operation of a data structure is measured by the number of cell probes performed. An algorithm is free to do unlimited amounts of computation based on the contents of probed cells.

The cell probe model does not effectively capture the correct scenario for data structures that provide privacy of the operations performed such as oblivious RAM. The typical scenario considered involves two parties denoted the *client* and the *server*. The client outsources the storage of data to the server while maintaining the ability to perform some set of operations over the data efficiently. In addition, the client wishes to hide the operations performed from the adversarial server that views the contents of all cells in memory as well as the sequence of cells probed in memory. We importantly note that the server does not learn about the contents and sequence of accesses to the client's storage. For this reason, Larsen and Nielsen [22] defined the oblivious cell probe model to prove lower bounds for oblivious RAMs. We note that the differentially private cell probe model is identical to the oblivious cell probe model except for the simple replacement of obliviousness with differential privacy as the privacy requirement. For a full description of the oblivious cell probe model, we refer the reader to Section 2 of [22].

To define the differentially private cell probe model, we first describe a *data structure problem* as well as a *differentially private cell probe data structure* for any data structure problem.

Definition 1. A data structure problem P is defined by a tuple (U, Q, O, f) where

1. U is the universe of all update operations;
2. Q is the universe of all query operations;
3. O is domain of all possible outputs for all queries;
4. $f : U^* \times Q \rightarrow O$ is a function that describes the desired output of any query $q \in Q$ given the history of all updates, $(u_1, u_2, \dots, u_m) \in U^*$.

A differentially private cell probe data structure **DS** for the data structure problem P consists of a randomized algorithm implementing update and query operations for P . **DS** is parameterized by the integers c and w denoting the client storage and cell size in bits respectively. Additionally, **DS** is given a random string \mathcal{R} of finite length r containing all randomness that **DS** will use. Note that \mathcal{R} can be arbitrarily large and, thus, contain all the randomness of a random oracle. Given the random string, our algorithms can be viewed as deterministic. Each algorithm is viewed as a finite decision tree executed by the *client* that probes (read or overwrite) memory cells owned by the *server*. For each $q \in Q$ and $u \in U$, there exists a (possibly) different decision tree. Each node in the decision tree is labelled by an index indicating the location of the server-held memory cell to be probed. For convenience, we will assume that a probe may both read and overwrite cell contents. This only reduces the number of cell probes by a factor of at most 2. Additionally, all leaf nodes are labelled with an element of O indicating the output of **DS** after execution.

Each edge in the tree is labelled by four bit strings. The first bit-string of length w represents the contents of the cell probed. The next w -bit string represents the new cell contents after overwriting. There are two c -bit strings representing the current client storage and the new client storage after performing the probe. Finally, there is a r -bit string representing the random string. The client executes **DS** by traversing the decision tree starting from the root. At each node, the client reads the indicated cell's contents. Using the random string, the current client storage and the cell contents, it finds the edge to the next node and updates the probed cell's contents and client storage accordingly. When reaching a leaf, **DS** outputs the element of O denoted at the leaf.

Note, **DS** is only permitted to use the contents of the previously probed cell, current client storage and the random string as input to generate the next cell probe or produce an output. The running time of **DS** is related to the depth of the decision tree as each edge corresponds to a cell probe. Furthermore, as the servers are passive, the server can only either update or retrieve a cell for the client. As a result, the running time (number of cell probes) multiplied by w (the cell size) gives us the bandwidth of the algorithm in bits. We now define the failure probability of **DS**.

Definition 2. A **DS** has failure probability $0 \leq \alpha \leq 1$ for a problem $P = (U, Q, O, f)$ if for any sequence of updates $u_1, \dots, u_m \in U^*$ and query $q \in Q$:

$$\Pr[\mathbf{DS}(u_1, \dots, u_m, q) \neq f(u_1, \dots, u_m, q)] \leq \alpha$$

where randomness is over the choice of \mathcal{R} .

As \mathcal{R} is finite, it might seem that we do not consider algorithms whose failure probability decreases in the running time but may never terminate. Instead, we can consider a variant of the algorithm that may run for an arbitrary long time but must provide an answer once its failure probability is small enough (for example, negligible in the number of item stored). Therefore, by sacrificing failure probability, we can convert such possibly infinitely running algorithms into finite algorithms with slightly larger failure probabilities. As we will prove our lower bounds for **DS** with failure probabilities at most $1/3$, we may also consider these kind of algorithms with vanishing failure probabilities and no termination guarantees.

We now move to privacy requirements and define the random variable $\mathbb{V}_{\mathbf{DS}}(Q)$ as the *adversary's view* of **DS** processing a sequence of operations Q where randomness is over the choice of the random string \mathcal{R} . The adversary's view contains the sequence of probes performed by **DS** to server-held memory cells. We stress that the view does not include the accesses performed by **DS** to client storage. We now define *differentially private access*.

Definition 3. **DS** provides (ϵ, δ) -differentially private access against computational adversaries if for any two sequences $Q = (\text{op}_1, \dots, \text{op}_m)$ and $Q' = (\text{op}'_1, \dots, \text{op}'_m)$ such that $|\{i \in \{1, \dots, m\} \mid \text{op}_i \neq \text{op}'_i\}| = 1$ and any PPT algorithm \mathcal{A} , then

$$\Pr[\mathcal{A}(\mathbb{V}_{\mathbf{DS}}(Q)) = 1] \leq e^\epsilon \Pr[\mathcal{A}(\mathbb{V}_{\mathbf{DS}}(Q')) = 1] + \delta.$$

Our results focus on *online* data structures where each cell probe may be assigned to a unique operation.

Definition 4. A **DS** is *online* if for any sequence $Q = (\text{op}_1, \dots, \text{op}_m)$, the adversary's view can be split up as:

$$\mathbb{V}_{\mathbf{DS}}(Q) = (\mathbb{V}_{\mathbf{DS}}(\text{op}_1), \dots, \mathbb{V}_{\mathbf{DS}}(\text{op}_m))$$

where each cell probe in $\mathbb{V}_{\mathbf{DS}}(\text{op}_i)$ is performed after receiving op_i and before receiving op_{i+1} .

Finally, we present the definition of an (ϵ, δ) -*differentially private cell probe data structure*. We present a diagram of the model in Figure 1.

Definition 5. A **DS** is an (ϵ, δ) -differentially private cell probe data structure if **DS** has failure probability $1/3$, provides (ϵ, δ) -differentially private access and is *online*.

We comment that the failure probability of $1/3$ does not seem to be reasonable for any scenario. However, proving a lower bound for **DS** with failure probability $1/3$ results in stronger lower bounds as they also hold for more reasonable situations with zero or negligibly small failure probabilities.

We now present the *array maintenance* problem introduced by Wang *et al.* [34], which crisply defines the online RAM problem.




Fig. 1: Diagram of differentially private cell probe model.

Definition 6. The array maintenance problem AM is parameterized by two integers $n, b > 0$ and defined by the tuple $(U_{\text{AM}}, Q_{\text{AM}}, O_{\text{AM}}, f_{\text{AM}})$ where

- $U_{\text{AM}} = \{\text{write}(i, B) : i \in [n], B \in \{0, 1\}^b\}$;
- $Q_{\text{AM}} = \{\text{read}(i) : i \in [n]\}$;
- $O_{\text{AM}} = \{0, 1\}^b$;

and, for a sequence $Q = (u_1, \dots, u_m)$ where $u_1, \dots, u_m \in U^*$, f_{AM} is:

$$f_{\text{AM}}(Q, \text{read}(i)) = \begin{cases} B, & \text{where } j \text{ is largest index such that } q_j = \text{write}(i, B); \\ 0^b, & \text{if there exists no such } j. \end{cases}$$

In words, the array maintenance problem requires that a data structure to store an array of n entries each of b bits. Each array location is uniquely identified by a number in $[n]$. Typically, it is assumed that a cell is large enough to contain an index. In this case, $w = \Omega(\log n)$. However, in our lower bounds, we will only assume that $w = \Omega(\log \log n)$ to achieve a stronger lower bound. An update operation (also called a *write*) takes as input an integer $i \in [n]$ and a b -bit string B and overwrites the array entry associated with i with the string B . For convenience, we denote a write operation with inputs i and B as $\text{write}(i, B)$. A query operation (also called a *read*) takes as input an integer $i \in [n]$ and returns the current b -bit string of the array entry associated with i . We denote a read operation with input i as $\text{read}(i)$. We will prove lower bounds for differentially private RAMs which are differentially private cell probe data structures for the array maintenance problem.

3 First Lower Bound

Let DS be a (ϵ, δ) -differentially private RAM storing n b -bit entries. For any sequence of operations $Q = (\text{op}_1, \dots, \text{op}_m)$, we denote $t_w(Q)$ as the expected

amortized number of cell probes on a write operation and $t_r(Q)$ as the expected amortized number of cell probes on a read operation. Both expectations are over the choice of the random string \mathcal{R} used by **DS**. We write t_w and t_r as the largest value of $t_w(Q)$ and $t_r(Q)$ respectively over all sequences Q . We assume that cells are of size w bits and the client has c bits of storage. In this section, we will prove the following preliminary result. The result will be strengthened in Section 4 where we present the main result of this paper.

Theorem 2. *Let $\epsilon > 0$ and $0 \leq \delta \leq 1/3$ be constants and let **DS** be an (ϵ, δ) -differentially private RAM for n b -bit entries implemented over w -bit cells that uses c bits of local storage. If **DS** has failure probability at most $1/3$ and $w = \Omega(\log \log n)$ then*

$$t_w + t_r = \Omega((b/w) \log(n/c)).$$

In terms of block bandwidth, this implies that at least one of read and write has an expected amortized $\Omega(\log(n/c))$ block bandwidth overhead. The above theorem will be shown when **DS** has to process a sequence Q distributed according to distribution $\mathcal{Q}(0)$, where, for index $\text{idx} \in \{0, \dots, n-1\}$, distribution $\mathcal{Q}(\text{idx})$ is defined by the following probabilistic procedure:

1. Pick m uniformly at random from $\{n/2, n/2 + 1, \dots, n-1\}$.
2. Draw each $\mathcal{B}_1, \dots, \mathcal{B}_m$ independently and uniformly at random from $\{0, 1\}^b$.
3. Construct the sequence $\mathcal{U} = \text{write}(1, \mathcal{B}_1), \dots, \text{write}(m, \mathcal{B}_m)$.
4. Output $Q = (\mathcal{U}, \text{read}(\text{idx}))$.

Thus $\mathcal{Q}(\text{idx})$ assigns positive probability to sequences Q that write independent and uniformly at random chosen b -bit blocks to indices $1, 2, \dots, m$ each for m chosen uniformly at random from $\{n/2, n/2 + 1, \dots, n-1\}$ followed with a single read to idx .

In particular, we prove the above theorem using $\mathcal{Q}(0)$. For convenience, we denote $\mathcal{Q} = \mathcal{Q}(0)$ from now on. Without privacy, the above sequence does not seem to require many probes as index 0 is not overwritten. However, the lower bound will, critically, use the fact that the view of any computational adversary cannot differ significantly from read operations where the last operation attempts to read a previously overwritten index $\text{idx} \in \{1, \dots, m\}$.

We prove the lower bound using the *chronogram* technique first introduced by Fredman and Saks [12] along with the modifications by Pătraşcu [28]. The strategy employed by the chronogram technique when applied to a sequence sampled according to $\mathcal{Q}(0)$ goes as follows. For any choice of m , we consider the $n/2$ **write** operations that immediately precede the **read**(0) operation and we split them into consecutive and disjoint partitions, which we denote as *epochs*. The epochs will grow exponentially in size and are indexed going backwards in time (order of operations performed). That is, the epoch consisting of the **write** operations immediately preceding the **read** operation will have the smallest index, while the epoch furthest in the past will have the largest index. Note that, when the sequence of operations is chosen according to distribution $\mathcal{Q}(\text{idx})$, the set of blocks overwritten by the **write** operations that fall into epoch i is a random variable which we denote by \mathcal{U}^i and depend on the value of m .

To prove Theorem 2, we consider a simple epoch construction. Epoch i consists of $\ell_i = 2^i$ **write** operations and thus there will be $k = \log_2(n/2)$ epochs. We also define s_i to be the total size of epochs $1, \dots, i$. In the epoch construction of this section, we have $s_i = 2^{i+1} - 1$. See Figure 2 for a diagram of the layout of the epochs with regards to a sequence of operations. In Section 4, we will derive stronger lower bounds by considering more complex epoch constructions with different parameters.




Fig. 2: Diagram of epoch construction of Section 3.

Defining random variables. Since we are considering online data structures, each cell probe performed by **DS** while processing a sequence Q can be uniquely associated to a **read** or **write** operation of Q . Random variable $\mathcal{T}_w(Q)$ is defined as the set of cell probes performed by **DS** while processing the **write** operations of the sequence Q . Similarly, we define $\mathcal{T}_r(Q)$ as the random variable of the set of cell probes performed by **DS** when processing the **read** operations of Q . The probability spaces of the two variables are taken over the randomness \mathcal{R} used by **DS**.

The following random variables are specifically defined for sequences Q in the support distribution $\mathcal{Q}(\text{idx})$, for some idx . These sequences first perform m **write** operations for m chosen uniformly at random from $\{n/2, \dots, n-1\}$ followed by a single **read**(idx) operation. The m **write** operations overwrite entries $1, \dots, m$ with random b -bit strings. We denote by $\mathcal{T}_w^j(Q)$ the random variable of the cells that are probed during the execution of a **write** operation of epoch j in Q . We further partition the cells probes in $\mathcal{T}_w^j(Q)$ according to the epoch they were last overwritten. Specifically, for $i \geq j$, we define $\mathcal{T}_w^{i,j}(Q)$ as the random variable of the set of probes $\mathcal{T}_w^j(Q)$ performed when executing **write** operations in epoch j of Q to a cell that was last overwritten by an operation in epoch i . Note that the sets $\mathcal{T}_w^{i,j}(Q)$ for all $i \leq j$ constitute a partition of $\mathcal{T}_w(Q)$. It will be convenient in the proof to define $\mathcal{T}_w^{<i}(Q) = \mathcal{T}_w^{i,1}(Q) \cup \dots \cup \mathcal{T}_w^{i,i-1}(Q)$ as the set of probes that are performed by an operation in any of epochs $\{1, \dots, i-1\}$ to a cell that was last overwritten by an operation in epoch i . Similarly, we denote the random variable $\mathcal{T}_r^i(Q)$ as the set of probes performed by the **read** operation to cells that were last overwritten by an operation in epoch i . In Figure 3, we show a diagram of $\mathcal{T}_w^{<i}(Q)$ and $\mathcal{T}_r^i(Q)$. We extend the definitions above to random distributions of sequences Q in a natural way. For example, $\mathcal{T}_w(Q)$ is defined by first picking sequence Q according to \mathcal{Q} and then sampling a set according to $\mathcal{T}_w(Q)$.

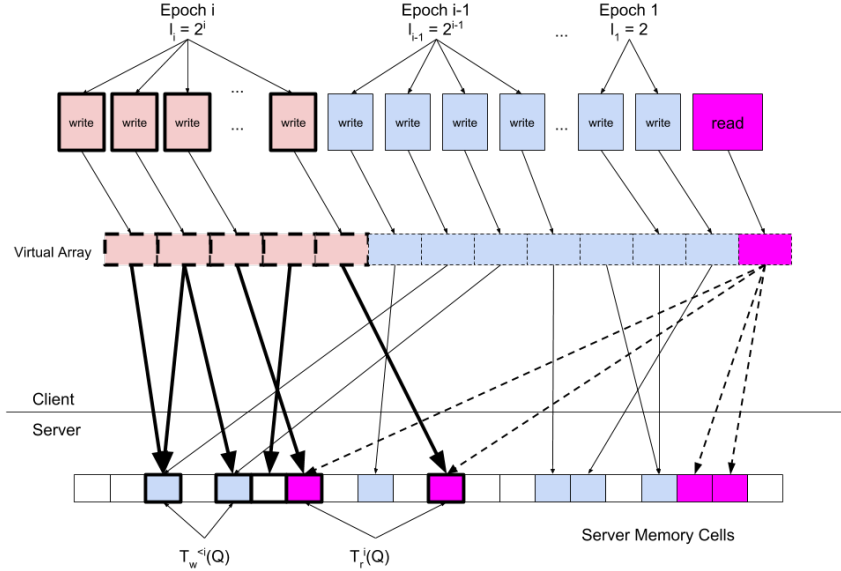


Fig. 3: Diagram of $\mathcal{T}_w^{<i}(Q)$ and $\mathcal{T}_r^i(Q)$.

3.1 A Tradeoff between $\mathcal{T}_w^{<i}(Q)$ and $\mathcal{T}_r^i(Q)$

From a high level, the proof of Theorem 2 is based on the fact that $\mathcal{T}_w^{<i}(Q)$ and $\mathcal{T}_r^i(Q)$ cannot be both small for all epochs i . To see why this must be intuitively true consider distribution Q_i over query sequences where the last **read** operation to the 0-th index is replaced with an index chosen uniformly at random from U^i (remember U^i are the indices of the array entries that are overwritten by **write** operations in epoch i). Since each **write** operation overwrites a distinct entry with a uniformly chosen b -bit string, a sufficiently large number of bits that were encoded by **write** operations in epoch i must be retrieved by the **read** operation. There are only three ways that these bits can be retrieved by the **read** operation. The first way is to probe cells that were last overwritten by any **write** operation of epoch i which corresponds to $\mathcal{T}_r^i(Q_i)$. Another way is to probe cells that were last overwritten by operations that occurred after epoch i ; that is, in any epoch $1 \leq j < i$. However, the total number of bits encoded by operations in epochs $1 \leq j < i$ is upper bounded by the number of probes performed in epochs $1 \leq j < i$ to cells that were last overwritten by an operation in epoch i , which corresponds to $\mathcal{T}_w^{<i}(Q_i)$. The final way to retrieve information from the **write** operations of epoch i is to encode information in the client's storage of c bits. However, if we consider the case when the number of entries overwritten in epoch i , l_i , is significantly larger than c , then the client's storage is too small to encode any significantly large amount of information compared to the total number of

write operations of epoch i . As a result, the total combined size of $\mathcal{T}_w^{<i}(\mathcal{Q}_i)$ and $\mathcal{T}_r^i(\mathcal{Q}_i)$ or, better, a function of the two quantities, can be lower bounded. However, recall that we wish to lower bound the values when processing the random sequence \mathcal{Q} and not \mathcal{Q}_i . The only difference between \mathcal{Q} and \mathcal{Q}_i is the index of the **read** operation performed at the end. By computational differential privacy, any random event that can be verified by a PPT adversary cannot occur with significantly different probabilities when **DS** processes \mathcal{Q} as opposed to \mathcal{Q}_i . Since the sets of cell probes can easily be computed in polynomial time, a lower bound on the sum of $|\mathcal{T}_w^{<i}(\mathcal{Q}_i)| + |\mathcal{T}_r^i(\mathcal{Q}_i)|$ also implies a lower bound on the $|\mathcal{T}_w^{<i}(\mathcal{Q})| + |\mathcal{T}_r^i(\mathcal{Q})|$ for a differentially private **DS**.

As explained above, the technical crux of the lower bound on $|\mathcal{T}_w^{<i}(\mathcal{Q}_i)| + |\mathcal{T}_r^i(\mathcal{Q}_i)|$ is an encoding argument that is captured by the following lemma that shows that a certain random variable $\mathcal{Z}_i(\mathcal{Q}(j))$ is “large” with probability at least $1/2$.

Lemma 1. *Assume that **DS** has failure probability at most $1/3$. Then, for any epoch $i \in \{1, \dots, k\}$ such that $\ell_i \geq \sqrt{n}$, there exists an index $j \in \{1, \dots, n-1\}$ such that*

$$\Pr[\mathcal{Z}_i(\mathcal{Q}(j)) \geq b/8] \geq 1/2$$

where $\mathcal{Z}_i(\mathcal{Q}(j))$ is

$$\frac{1}{\ell_i} \left(|\mathcal{T}_w^{<i}(\mathcal{Q}(j))|w + \log \binom{t_w s_{i-1}}{|\mathcal{T}_w^{<i}(\mathcal{Q}(j))|} \right) + |\mathcal{T}_r^i(\mathcal{Q}(j))|w + \log \binom{t_r}{|\mathcal{T}_r^i(\mathcal{Q}(j))|} + \frac{c}{\ell_i}.$$

The proof of Lemma 1 is found in Section 3.4. $\mathcal{Z}_i(\mathcal{Q}(j))$ can be viewed as the total average information retrieved from the **write** operations of epoch i by the **read**(j) operation at the end of $\mathcal{Q}(j)$. More precisely, let us explain the meaning of each term of the value $\mathcal{Z}_i(\mathcal{Q}(j))$ before showing how running **DS** on the random sequence \mathcal{Q}^i eventually leads to finding a fixed index j . The first term of $\mathcal{Z}_i(\mathcal{Q}(j))$ measures the average amount of information pertaining to each of the ℓ_i **write** operations of epoch i that are read by cell probes performed in epochs following epoch i . Each of the cell probes in $\mathcal{T}_w^{<i}(\mathcal{Q}(j))$ reads exactly w bits in a cell. In addition, the choice of which cell probes performed in epochs following epoch i actually belong to $\mathcal{T}_w^{<i}(\mathcal{Q}(j))$ also encodes some information. As there are s_{i-1} **write** operations epochs following epoch i , there are at most $t_w s_{i-1}$ cell probes and at most $\binom{t_w s_{i-1}}{|\mathcal{T}_w^{<i}(\mathcal{Q}(j))|}$ choices of the cells to probe leading to $\log \binom{t_w s_{i-1}}{|\mathcal{T}_w^{<i}(\mathcal{Q}(j))|}$ bits. Similarly, each probe in $\mathcal{T}_r^i(\mathcal{Q}(j))$ reads w bits in each cell and there are at most $\binom{t_r}{|\mathcal{T}_r^i(\mathcal{Q}(j))|}$ choices of probes when performing **read**(j) that belong to $\mathcal{T}_r^i(\mathcal{Q}(j))$. The last term of $\mathcal{Z}_i(\mathcal{Q}(j))$ considers the average amount of information for each of the ℓ_i operations in epoch i that are encoded in the client’s storage of c bits. Having explained the meaning of $\mathcal{Z}_i(\mathcal{Q}(j))$, let us explain why there must exist such an index j . The total amount of information (in expectation) transferred if all ℓ_i possible **read** operations of \mathcal{Q}^i are performed requires retrieving all ℓ_i random b bit blocks which has $\ell_i \cdot b$ bits of entropy. As a result, j considers the worst case of the ℓ_i possible **read** operations leading us to the above lemma. A formal proof of these ideas are presented in Section 3.4.

3.2 Using Differential Privacy

We note that Lemma 1 does not suffice to prove an $t_w + t_r = \omega(1)$ lower bound for **DS** directly. Typically, chronogram lower bounds will find a single sequence that forces a large amount of information transfer from all epochs simultaneously. Instead, Lemma 1 states there exists some sequence that forces a large information transfer for each epoch and that the sequences are possibly different for each epoch. In fact, without assuming privacy about a data structure, there can be no single sequence that requires large information from many epochs as there are trivial $\Theta(1)$ data structures that solve the array maintenance problem without any privacy guarantees. As Lemma 1 does not assume privacy for **DS**, we will need to incorporate the fact that **DS** is differentially private to achieve a statement that there exists a single sequence that forces large information transfer from many epochs simultaneously.

Let us now assume that **DS** provides differential privacy against computational adversaries with parameters $\epsilon = O(1)$ and $0 \leq \delta \leq 1/3$. For any fixed sequence Q , we consider any probabilistic event $\mathcal{E}(Q)$ over the randomness of the choice of the random string \mathcal{R} such that there exists a probabilistic polynomial time algorithm that can verify $\mathcal{E}(Q)$ being true or false. Then, computational differential privacy implies that for any fixed sequence Q_1 and Q_2 that differ in exactly one operation, then $\Pr[\mathcal{E}(Q_1) \text{ is true}] \leq e^\epsilon \Pr[\mathcal{E}(Q_2) \text{ is true}] + \delta$. In particular, we can consider the event when $\mathcal{E}(Q) = \mathcal{Z}_i(Q) \geq b/8$. Note that $\mathcal{Z}_i(Q)$ can be computed by any computational adversary by simply assigning each cell probe performed by **DS** over Q into one of $\{\mathcal{T}_w^{<i}(Q)\}_{i=1,\dots,k}$ or $\{\mathcal{T}_r^i(Q)\}_{i=1,\dots,k}$ where assigning a cell probe depends only on the last time the cell was overwritten and the current operation of Q . As a result, we know that for any two fixed sequences Q_1 and Q_2 that differ in exactly one operation, then

$$\Pr[\mathcal{Z}_i(Q_1) \geq b/8] \leq e^\epsilon \Pr[\mathcal{Z}_i(Q_2) \geq b/8] + \delta.$$

Note that \mathcal{Q} and $\mathcal{Q}(\text{idx})$ only differ in the input index to the **read** operation at the end of the sequence. We use this fact to prove the following lemma that $\mathcal{Z}_i(\mathcal{Q})$ cannot differ significantly from $\mathcal{Z}_i(\mathcal{Q}(\text{idx}))$ for any idx .

Lemma 2. *Assume that **DS** is differentially private with $0 \leq \delta \leq 1/3$ and has failure probability at most $1/3$. Then, for all epochs $i \in \{1, \dots, k\}$ such that $\ell_i \geq \sqrt{n}$,*

$$\Pr[\mathcal{Z}_i(\mathcal{Q}) \geq b/8] \geq 1/(6e^\epsilon).$$

Proof. Consider any epoch $i \in \{1, \dots, k\}$. By Lemma 3.4, there exists an index idx_i such that $\Pr[\mathcal{Z}_i(\mathcal{Q}(\text{idx}_i)) \geq b/8] \geq 1/2$. We define $\mathcal{Q}(\text{idx}, B_1, \dots, B_m) =$

$\text{write}(1, B_1), \dots, \text{write}(m, B_m), \text{read}(\text{idx})$. Then,

$$\begin{aligned} \Pr[\mathcal{Z}_i(\mathcal{Q}) \geq b/8] &= \sum_{m=n/2}^{n-1} \sum_{B_1, \dots, B_m \in \{0,1\}^b} \frac{1}{m2^{rm}} \Pr[\mathcal{Z}_i(Q(0, B_1, \dots, B_m)) \geq b/8] \\ &\geq \sum_{m=n/2}^{n-1} \sum_{B_1, \dots, B_m \in \{0,1\}^b} \frac{1}{m2^{rm}} \left(\frac{\Pr[\mathcal{Z}_i(Q(\text{idx}_i, B_1, \dots, B_m)) \geq b/8] - \delta}{e^\epsilon} \right) \\ &= \frac{\Pr[\mathcal{Z}_i(Q(\text{idx}_i)) \geq b/8] - \delta}{e^\epsilon} \geq \frac{1/2 - 1/3}{e^\epsilon} = \frac{1}{6e^\epsilon}. \end{aligned}$$

3.3 Completing the proof of Theorem 2

Lemma 2 resembles the typical desired statement for data structure lower bounds as it guarantees existence of a distribution of query sequences, \mathcal{Q} , that forces a large amount of information transfer from all epochs in expectation.

Recall that we consider epoch i consisting of $\ell_i = 2^i$ **write** operation for a total of $s_i = 2^{i+1} - 1$ **write** operation in epochs $1, \dots, i$. We refer the reader to Figure 2 for a visual reminder of our epoch construction. Using Lemma 2, we will show that $\Omega(b/w)$ bits must be transferred from the *majority of large* epochs. In particular, we focus on epochs whose number of **write** operations, ℓ_i , is much larger than c . For otherwise, the **write** operations in epoch i may be entirely encoded into client's storage of c bits and thus no information from epoch i is required to be transferred by cell probes of future operations. Concretely, we say that an epoch is *large* if $\ell_i \geq \max\{\sqrt{n}, c^2\}$ and note that, for our definition of epochs, we have $\hat{k} := \Theta(\log(n/c))$ large epochs. We will show that for many large epochs $\Omega(b/w)$ bits must be transferred by cell probes of either **write** operations of future epochs or the **read** operation.

To achieve our lower bound, we will analyze the expectation of $\mathcal{Z}_i(\mathcal{Q})$ based on our epoch construction. We will provide a high-level overview of the steps of our analysis in this paragraph before performing a formal analysis. Recall that t_w and t_r are an upper bound on the expected amortized number of cells probed per **write** and **read** operation for any sequence. For the majority of epochs i , we cannot expect the **read** operation of \mathcal{Q} to probe more than t_r/\hat{k} cells containing information about the **write** operations of epoch i . This provides an upper bound on $|\mathcal{T}_r^i(\mathcal{Q})|$ for the majority of epochs. We want a similar upper bound on the value of $|\mathcal{T}_w^{<i}(\mathcal{Q})|$. Recall that this number corresponds to the number of probes performed by **write** operations that read cells that encode information about the **write** operations of epoch i . Our argument will critically use the fact that the sequence \mathcal{Q} is chosen at random. Recall that \mathcal{Q} is chosen to have m **write** operations where m is chosen uniformly at random from $\{n/2, n/2+1, \dots, n-1\}$. The data structure **DS** is unable to predict the point in time when the **read** operation will occur. Instead, the best that **DS** can achieve is to prepare for all possible epoch configurations. Since there are \hat{k} epochs with size at least $\max\{\sqrt{n}, c^2\}$, each update should be only able to encode $\frac{t_w \cdot w}{\hat{k}}$ about each of

these epochs. As a result, we can prove the majority of epochs cannot have very large values of $|\mathcal{T}_w^{<i}(\mathcal{Q})|$ in expectation.

As the two bounds above hold for the majority of epochs, we can show there exists at least one large epoch i such that both the values of $|\mathcal{T}_w^{<i}(\mathcal{Q})|$ and $|\mathcal{T}_r^i(\mathcal{Q})|$ are small. In particular, we show the following:

Lemma 3. *There exists $i \in \{1, \dots, k\}$ such that $\ell_i \geq \max\{c^2, \sqrt{n}\}$ and*

$$\mathbb{E}[|\mathcal{T}_w^{<i}(\mathcal{Q})|/\ell_i] = O(t_w/\log(n/c)) \quad \text{and} \quad \mathbb{E}[|\mathcal{T}_r^i(\mathcal{Q})|] = O(t_r/\log(n/c)).$$

Proof. The lemma is derived from the following two statements:

1. There exists $\hat{k}/2 + 1$ large epochs i such that $\mathbb{E}[|\mathcal{T}_w^{<i}(\mathcal{Q})|] = O(s_{i-1}t_w)$.
2. There exists $\hat{k}/2 + 1$ large epochs i such that $\mathbb{E}[|\mathcal{T}_r^i(\mathcal{Q})|] = O(t_r/\log(n/c))$.

Since there are only \hat{k} large epochs, there must exist at least one large epoch where both inequalities hold. We now show the two statements are true.

Let us pick epoch i uniformly at random amongst the \hat{k} large epochs and fix the random string \mathcal{R} as well as the $n - 1$ block values $\mathcal{B}_1, \dots, \mathcal{B}_{n-1}$. We now fix a cell probe **probe** of the execution of **DS** over the **write** operations $\text{write}(1, \mathcal{B}_1), \dots, \text{write}(n - 1, \mathcal{B}_{n-1})$ and consider the probability that **probe** contributes to $\mathcal{T}_w^{<i}(\mathcal{Q})$ from which we derive a bound on $\mathbb{E}[|\mathcal{T}_w^{<i}(\mathcal{Q})|/\ell_i]$. Note that, having fixed \mathcal{R} and the values \mathcal{B}_j 's, the probability space is over the choice of m from $\{n/2, n/2 + 1, \dots, n - 1\}$ and of i . We denote p_r as the index of the **write** operation in \mathcal{U} when **probe** is performed. The value p_w is denoted as the index of the **write** operation in \mathcal{U} when the cell of **probe** was last overwritten. Using p_r and p_w , we can attempt to upper bound the probability that the **probe** belongs to $\mathcal{T}_w^{<i}(\mathcal{Q})$. First, let e be the smallest integer such that $p_r - p_w \leq s_e$. Note that **probe** cannot contribute to $\mathcal{T}_w^{<j}(\mathcal{Q})$ for any epoch $j \leq e - 1$, since there are only s_j operations between the beginning of epoch k and the **read** operation. Since $s_j \leq s_{e-1} < p_r - p_w$, either the **read** operation has to occur after the **read** operation or the last operation to overwrite the cell **probe** occurs before the j -th epoch. We remind the reader that the exact locations of epochs is determined by m . The boundary denoting the end of epoch j has to occur after p_w and before p_r meaning there are at most s_e choices from the position of the **read** operation such that this cell **probe** contributes to $\mathcal{T}_w^{<j}(\mathcal{Q})$. There are $n/2$ choices for m , so the probability is at most $2s_e/n$. We now compute

$$\mathbb{E}[|\mathcal{T}_w^{<i}(\mathcal{Q})|/\ell_i] = \frac{1}{\hat{k}} \sum_{j: \ell_j \geq \max\{\sqrt{n}, c^2\}} \mathbb{E}[|\mathcal{T}_w^{<j}(\mathcal{Q})|/\ell_j].$$

The **probe** only contributes to epochs $j \geq e$. Note, there are at most (in expectation) $t_w \cdot (n - 1)$ cell probes performed when processing the **write** operations of \mathcal{Q} . By linearity of expectation,

$$\sum_{j: \ell_j \geq \max\{\sqrt{n}, c^2\}} \mathbb{E} \left[\frac{|\mathcal{T}_w^{<j}(\mathcal{Q})|}{\ell_j} \right] \leq t_w \cdot n \sum_{j \geq e} \frac{2 \cdot s_e}{n \cdot \ell_j} \leq 2t_w \cdot \left(\frac{s_e}{\ell_e} + \frac{s_e}{\ell_{e+1}} + \dots \right) \leq 4t_w.$$

As a result, there exists $\hat{k}/2 + 1$ fixed epochs i such that their expectation over the m is at most $12t_w$.

We know that $\sum_i \mathbb{E}[|\mathcal{T}_r^i(\mathcal{Q})|] \leq t_r$. Therefore, there exists $\hat{k}/2 + 1$ epochs i with $\ell_i \geq \max\{\sqrt{n}, c^2\}$ such that $\mathbb{E}[|\mathcal{T}_r^i(\mathcal{Q})|] \leq 3t_r$ completing the proof.

We can now achieve our goal of proving Theorem 2 that gives a lower bound on the sum $t_w + t_r$ by plugging the inequalities in Lemma 3 into the expectation of $\mathcal{Z}_i(\mathcal{Q})$ and then using the bound from Lemma 2.

Proof (Theorem 2). First, we analyze the expectation of $\mathcal{Z}_i(\mathcal{Q})$. Note that, for every x, y , $\log\left(\frac{y}{x}\right) = O(x \log(y/x))$. Moreover, for every y , $x \log(y/x)$ is a convex function over x , so we can write the $\mathbb{E}[x \log(y/x)] \leq \mathbb{E}[x] \log(y/\mathbb{E}[x])$ where the expectation is over the choice of x . We now apply this observation to $\mathbb{E}[\mathcal{Z}_i(\mathcal{Q})]$ and get:

$$O\left(\frac{\mathbb{E}[|\mathcal{T}_w^{<i}(\mathcal{Q})|]}{\ell_i} \left(w + \log \frac{t_w s_{i-1}}{\mathbb{E}[|\mathcal{T}_w^{<i}(\mathcal{Q})|]}\right) + \mathbb{E}[|\mathcal{T}_r^i(\mathcal{Q})|] \left(w + \log \frac{t_r}{\mathbb{E}[|\mathcal{T}_r^i(\mathcal{Q})|]}\right) + \frac{c}{\ell_i}\right).$$

By Lemma 2, we know that $\mathbb{E}[\mathcal{Z}_i(\mathcal{Q})] = \Omega(b)$. We now pick our epoch i as the one chosen by Lemma 3 and plug in the inequalities to get

$$\frac{s_{i-1} t_w}{\ell_i \log(n/c)} \left(w + \log \frac{t_w s_{i-1}}{s_{i-1} t_w / \log(n/c)}\right) + \frac{t_r}{\log(n/c)} \left(w + \log \frac{t_r}{t_r / \log(n/c)}\right) = \Omega(b).$$

Here we have used the fact that epoch i is large and thus $\frac{c}{\ell_i} = O(1)$, since $\ell_i \geq c^2$. Also, note that $s_{i-1} = \Theta(\ell_i)$. Therefore, we can simplify and get that $t_w + t_r = \Omega\left(\frac{b}{w + \log \log n}\right) \log(n/c)$. If we assume that $w = \Omega(\log \log n)$, we can simplify and get the following result $t_w + t_r = \Omega\left(\frac{b}{w}\right) \log(n/c)$ which completes the proof.

Therefore, the lower bound of $t_w + t_r$ described in Theorem 2 can be entirely derived from Lemma 1. It remains to prove Lemma 1, which we do next.

3.4 An Encoding Argument using $\mathcal{T}_w^{<i}(\mathcal{Q})$ and $\mathcal{T}_r^i(\mathcal{Q})$

In this section, we formally present the encoding argument that underlies the proof of Lemma 1. We first give a high level description of the proof. The main idea involves converting any **DS** that solves the array maintenance problem into a one-way communication problem between two parties, for which we have a lower bound on the number of bits that must be sent.

Specifically, we consider the case in which for a fixed epoch $i \in \{1, \dots, k\}$ and for a sequence drawn according to \mathcal{Q} , one party, Alice, receives the m values $\mathcal{B}_1, \dots, \mathcal{B}_m$ and a random string \mathcal{R} and the other party, Bob, receives the same random string \mathcal{R} as well as $m - \ell_i$ values; that is, all of $\mathcal{B}_1, \dots, \mathcal{B}_m$ except for the ℓ_i values updated in epoch i of sequence \mathcal{Q} . The goal of the protocol is to send Bob the missing ℓ_i values.

As the ℓ_i b -bit values are generated uniformly and independently at random, Alice's input has $\ell_i \cdot b$ bits of entropy conditioned on Bob's input. By Shannon's

Source Coding Theorem, any protocol for the above problem must have expected communication of at least $\ell_i \cdot b$ bits. We show that if Lemma 1 does not hold, then Shannon's Theorem is contradicted by giving an encoding constructed by simulating **DS** that beats Shannon's bound.

Recall that, for any idx , $\mathcal{Q}(\text{idx})$ is constructed by picking m uniformly at random from $\{n/2, n/2 + 1, \dots, n - 1\}$ and constructing the sequence of m updates $\mathcal{U} = \text{write}(1, \mathcal{B}_1), \dots, \text{write}(m, \mathcal{B}_m)$ where each $\mathcal{B}_1, \dots, \mathcal{B}_m$ is drawn independently and uniformly at random from $\{0, 1\}^b$. We also denote by \mathcal{U}^i the set of **write** of epoch i , for $i = 1, \dots, k$.

Consider the following protocol. Alice and Bob locally execute all **write** operations in epochs $k, k - 1, \dots, i + 1$ using the random string \mathcal{R} . Bob keeps a snapshot snap_B of **DS** at this point. Now Bob can learn each of the ℓ_i values \mathcal{B}_{idx} for $\text{idx} \in \mathcal{U}^i$ written during epoch i , by simulating epoch j , for $j = i, i - 1, \dots, 1$ followed by the **read**(idx) operation. For this to be possible, Bob uses the snapshot snap_B , that gives the state of **DS** before any **write** operations of epoch i are executed, and the following information that can be transferred by Alice.

1. The c bits of client storage of **DS** after the **write** operations of epoch i have been processed.
2. The location and contents of the cells that are probed by the **write** operations of epochs $j = i - 1, \dots, 1$ and by the **read**(idx) operation.

Given this information as well as the random string \mathcal{R} , Bob can simulate **DS** by starting from snap_B and executing all the **write** operations of \mathcal{U} occurring after epoch i as well as **read**(idx) and thus recover \mathcal{B}_{idx} . To encode all ℓ_i block values updated in epoch i , Alice and Bob can repeat the simulation of the **read** operation ℓ_i times with idx ranging over the set of the ℓ_i indices that are updated in epoch i . The number of bits that need to be transferred to Bob by Alice depends on the following three values:

1. The number of bits of the client storage, c .
2. The number of probes performed in epochs $j = i - 1, \dots, 1$ to cells last written in epoch i .
3. The number of probes performed by the ℓ_i **read** operations to the ℓ_i indices updated in epoch i .

By Shannon's source coding theorem, we have a lower bound on the number of bits that can be transferred and, consequently, a lower bound on the number of probes performed by **DS**. The rough description above only works for **DS** that never fails but it only requires some small changes to work for failure probability $1/3$. In particular, Alice can indicate the indices idx for which **DS** fails to return \mathcal{B}_{idx} and explicitly transfer the b bits of \mathcal{B}_{idx} to Bob in addition to the above protocol. We now present the formal proof of Lemma 1.

Proof (Lemma 1). In our proof, we consider **DS** that have failure probability at most $1/512$. Note that any **DS** with failure probability $1/3$ implies the existence of a **DS** with failure probability $1/512$ as one can execute **DS** a constant number

of times with independently chosen randomness and return the most popular result to answer any **read** operation. In fact, proving a lower bound for **DS** with failure probability $1/512$ implies any **DS** with failure probability that is a constant greater than $1/2$ using the above method.

Recall that \mathcal{U}^i denotes the set of all ℓ_i indices that are updated by **write** operations in epoch i . It suffices to prove that $\Pr[\sum_{\text{idx} \in \mathcal{U}^i} \mathcal{Z}_i(\mathcal{Q}(\text{idx})) \geq \ell_i b/8] \geq 1/2$. Since \mathcal{U}^i contains ℓ_i indices, the previous statement implies that there must exist some $\text{idx} \in \mathcal{U}^i$ such that $\Pr[\mathcal{Z}_i(\mathcal{Q}(\text{idx})) \geq b/8] \geq 1/2$ which would complete the proof. Therefore, towards a contradiction, assume that $\Pr[\sum_{\text{idx} \in \mathcal{U}^i} \mathcal{Z}_i(\mathcal{Q}(\text{idx})) \geq \ell_i b/8] < 1/2$ for some data structure **DS** that solves the array maintenance problem with failure probability at most $1/512$. We will present an encoding of $\ell_i \cdot b$ random bits from Alice and Bob using **DS** that uses strictly less than $\ell_i \cdot b$ bits in expectation contradicting Shannon's source coding theorem.

In computing the encoding, Alice receives the m b -bit random values used by the sequence of **write** operations, \mathcal{U} , and a random string \mathcal{R} .

Alice's Encoding

1. Alice executes **DS** on the sequence \mathcal{U} using the random string \mathcal{R} up to the final **write** operation of epoch i . The content of the c bits of client storage after epoch i is completed are added to the encoding.
2. Alice then executes the remaining s_{i-1} **write** operations of \mathcal{U} of epochs $i-1, i-2, \dots, 1$. While processing these **write** operations, Alice records the subset $\mathcal{T}_w^{<i}(\mathcal{U})$ of probes to cells that were last written in epoch i as well as their contents. This information is encoded as follows. First the size $|\mathcal{T}_w^{<i}(\mathcal{U})|$ (at most $\log(t_w \cdot s_{i-1})$ bits) is added to the encoding. Then Alice adds an encoding of which $|\mathcal{T}_w^{<i}(\mathcal{U})|$ probes of the at most $t_w(n-1)$ probes over the entire sequence belong to $\mathcal{T}_w^{<i}$ (this costs $\log \binom{t_w(n-1)}{|\mathcal{T}_w^{<i}(\mathcal{U})|}$ bits). Finally, for each such probe, w bits are added to the encoding to specify the content of the cell probed (for additional $|\mathcal{T}_w^{<i}(\mathcal{U})| \cdot w$ bits).
3. Alice stores a snapshot snap_A of the **DS** after processing all **write** operations of \mathcal{U} . Alice will use this snapshot to simulate the **read** operations for the ℓ_i entries written in epoch i .
4. For each of the ℓ_i indices $\text{idx} \in \mathcal{U}^i$, Alice executes **read**(idx) on snap_A . Let F be the number of **read**(idx) operations that return a wrong value (that is, they return a value other than \mathcal{B}_{idx}). Alice adds the value F to the encoding costing $\log n$ bits and an encoding of the subset of the F *failing* indices costing $\log \binom{\ell_i}{F}$ expected bits.
5. For each of the F failing indices $\text{idx} \in \mathcal{U}^i$, Alice adds \mathcal{B}_{idx} to the encoding costing a total of $F \cdot b$ bits.
6. For each non-failing index $\text{idx} \in \mathcal{U}^i$ (that is, for which **read**(idx) executed on snap_A with \mathcal{R} successfully returns \mathcal{B}_{idx}), Alice adds the subset of probes performed during **read**(idx) to the cells in $\mathcal{T}_r^i(\mathcal{Q}(\text{idx}))$ (these are the cells last written in epoch i) as well as their content to the encoding. This costs w bits for each cell in $\mathcal{T}_r^i(\mathcal{Q}(\text{idx}))$ as well as $\log \binom{t_r}{|\mathcal{T}_r^i(\mathcal{Q}(\text{idx}))|}$ bits to encode the subset $\mathcal{T}_r^i(\mathcal{Q}(\text{idx}))$ of the at most t_r probes in **read**(idx).

7. Alice checks whether either of $\sum_{\text{idx} \in \mathcal{U}^i} \mathcal{Z}_i(\mathcal{Q}(\text{idx})) > \ell_i b/8$ or $F > \ell_i/64$. If either are true, Alice stops and returns an encoding consisting of a 0 bit followed by $\ell_i \cdot b$ bits of the ℓ_i blocks updated by the `write` operations in \mathcal{U}^i .
8. Otherwise, when both $\sum_{\text{idx} \in \mathcal{U}^i} \mathcal{Z}_i(\mathcal{Q}(\text{idx})) \leq \ell_i b/8$ and $F \leq \ell_i/64$, Alice prepends a 1 bit to the encoding computed in Steps 1-6 and returns it.

In decoding the message sent by Alice, Bob receives the random string \mathcal{R} but does not receive the entirety of \mathcal{U} . Instead, Bob receives \mathcal{U} except all block values that are updated in epoch i .

Bob's Decoding

1. Bob checks the first bit of Alice's encoding. If the first bit is a 0, then Bob parses the next $\ell_i \cdot b$ bits as the contents of the ℓ_i block values updated in \mathcal{U} completing the decoding.
2. If the encoding begins with a 1, Bob will execute the `write` operations in epochs $j = k, k-1, \dots, i-1$ using random string \mathcal{R} . Note that this is straightforward as Bob received all the needed values as input and the indices of the `write` are fixed.
3. Note that Bob does not have access to the updated array entries of epoch i , and thus will skip it.
4. Next, Bob sets the client storage as specified in the encoding and starts simulating the `write` operations for epochs $j = i-1, \dots, 1$. As long as the `write` operations do not require probing a cell that was last written in epoch i , Bob can simulate **DS** in the exact same way as done by Alice to compute the encoding (note Bob has access to the same \mathcal{R}). Whenever **DS** requires probing a cell last written in epoch i , Bob will use the encoding of the cell contents found in the encoding to continue simulation. As a result, Bob can simulate all `write` operations of \mathcal{U} after epoch i identically to Alice. Bob will now take a (partial) snapshot of **DS** including all cell locations and contents that Bob is aware of.
5. Next, Bob obtains F , the number of failing `read`, from the encoding along with the indices $\text{idx} \in \mathcal{U}^i$ where `read(idx)` fails to return \mathcal{B}_{idx} . For each of these F indices, Bob obtains the corresponding value \mathcal{B}_{idx} from the encoding.
6. For the remaining $\ell_i - F$ indices $\text{idx} \in \mathcal{U}^i$ such that `read(idx)` returns \mathcal{B}_{idx} , Bob will execute `read(idx)` on the snapshot of **DS**. From the encoding, Bob knows which of the (at most, in expectation) t_r probes performed by `read(idx)` are to cells last written in epoch i . Bob simulates `read(idx)` on his snapshot with \mathcal{R} using the cell contents encoded by Alice to retrieve \mathcal{B}_{idx} .

Analysis. It remains to analyze the expected length of Alice's encoding. Recall that we know from Shannon's source coding theorem that Alice's encoding has to be at least $\ell_i \cdot b$ bits long in expectation.

There are two cases to consider. In the first case, when the first bit is a 0, the encoding will be $1 + \ell_i \cdot b$ bits long. Let us now consider the case in which the

first bit is 1 and thus $F \leq \ell_i/64$ and $\sum_{\text{id}\mathbf{x} \in \mathcal{U}^i} \mathcal{Z}_i(\mathcal{Q}(\text{id}\mathbf{x})) \leq \ell_i b/8$. The encoding of the failed indices has expected length

$$\begin{aligned} \log n + \mathbb{E} \left[\log \binom{\ell_i}{F} \right] + b \cdot \mathbb{E}[F] &\leq \log n + \mathbb{E} \left[\log \binom{\ell_i}{\frac{\ell_i}{64}} \right] + b \cdot \frac{\ell_i}{64} \\ &\leq \log n + \frac{\ell_i}{64} \cdot (b + \log(64 \cdot e)) \leq \log n + \frac{9}{64}(\ell_i \cdot b). \end{aligned}$$

The second inequality uses Stirling's approximation which states that $\binom{x}{y} \leq (ex/y)^y$. We know Alice's encoding of client storage will always be c bits. We know the expected bits of encoding $\mathcal{T}_w^{<i}(\mathcal{U})$ is

$$\log(t_w \cdot s_{i-1}) + \mathbb{E} \left[\log \binom{t_w \cdot s_{i-1}}{|\mathcal{T}_w^{<i}(\mathcal{U})|} + |\mathcal{T}_w^{<i}(\mathcal{U})| \cdot w \right].$$

Note, $\log(t_w \cdot s_{i-1}) \leq 2 \log n$. Similarly, for all $\text{id}\mathbf{x} \in \mathcal{U}^i$ that successfully return $\mathcal{B}_{\text{id}\mathbf{x}}$, we know that the encoding requires

$$\mathbb{E} \left[|\mathcal{T}_r^i(\mathcal{Q}(\text{id}\mathbf{x}))|w + \log \binom{t_r}{|\mathcal{T}_r^i(\mathcal{Q}(\text{id}\mathbf{x}))|} \right].$$

Note that

$$\begin{aligned} &\mathbb{E} \left[\log \binom{t_w \cdot s_{i-1}}{|\mathcal{T}_w^{<i}(\mathcal{U})|} + |\mathcal{T}_w^{<i}(\mathcal{U})|w + \sum_{\text{id}\mathbf{x} \in \mathcal{U}^i} |\mathcal{T}_r^i(\mathcal{Q}(\text{id}\mathbf{x}))|w + \log \binom{t_r}{|\mathcal{T}_r^i(\mathcal{Q}(\text{id}\mathbf{x}))|} \right] + c \\ &\leq \sum_{\text{id}\mathbf{x} \in \mathcal{U}^i} \mathcal{Z}_i(\mathcal{Q}(\text{id}\mathbf{x})) \leq \frac{1}{8}(\ell_i \cdot b). \end{aligned}$$

Summing over all parts of the encoding, we get that

$$3 \log n + \frac{9}{64}(\ell_i \cdot b) + \sum_{\text{id}\mathbf{x} \in \mathcal{U}^i} \mathcal{Z}_i(\mathcal{Q}(\text{id}\mathbf{x})) \leq 3 \log n + \frac{17}{64}(\ell_i \cdot b).$$

The final step of the analysis is to compute the probabilities that Alice places a 0 or a 1 as the first bit of the encoding. By Markov's inequality, $\Pr[F \geq \ell_i/64] \leq 1/8$ and we know that $\Pr[\sum_{\text{id}\mathbf{x} \in \mathcal{U}^i} \mathcal{Z}_i(\mathcal{Q}(\text{id}\mathbf{x})) \geq \ell_i b/8] < 1/2$ by our initial assumption towards a contradiction. As a result, we know that $\Pr[F \geq \ell_i/64 \text{ or } \sum_{\text{id}\mathbf{x} \in \mathcal{U}^i} \mathcal{Z}_i(\mathcal{Q}(\text{id}\mathbf{x})) \geq b/8] < 5/8$. So, Alice's expected encoding size is at most

$$1 + 3 \log n + \frac{5}{8}(\ell_i \cdot b) + \frac{17}{64}(\ell_i \cdot b) < \ell_i \cdot b$$

contradicting Shannon's source coding theorem when $\ell_i \geq \sqrt{n}$.

4 Main Result

In Section 3, we presented a lower bound on the sum of t_w , the worst case expected amortized bandwidth for **write** operations, and t_r , the worst case expected amortized bandwidth for **read** operations that implies that $\max\{t_w, t_r\} =$

$\Omega((b/w) \log(n/c))$. However, this lower bound does not preclude the existence of a differentially private RAM with $t_w = \Theta((b/w) \log(n/c))$ and $t_r = o((b/w) \log(n/c))$ or $t_r = \Theta((b/w) \log(n/c))$ and $t_w = o((b/w) \log(n/c))$. In this section, we strengthen our lower bound and prove the following two statements, for (ϵ, δ) -differentially private RAM, for any constant ϵ and $\delta \leq 1/3$,

1. If $t_w = o((b/w) \log(n/c))$, then $t_r = \omega((b/w) \log(n/c))$;
2. If $t_r = o((b/w) \log(n/c))$, then $t_w = \omega((b/w) \log(n/c))$.

Therefore, since $\max\{t_w, t_r\} = O((b/w) \log(n/c))$, then it must be the case that both $t_w = \Theta((b/w) \log(n/c))$ and $t_r = \Theta((b/w) \log(n/c))$ showing that imbalanced running times for **write** and **read** operations cannot improve the asymptotic efficiency of differentially private RAMs constructions.

To achieve these tradeoffs, we revisit our epoch construction of Section 3. Let us, first, focus our attention on the first statement where we show that $t_r = \omega((b/w) \log(n/c))$ when $t_w = o((b/w) \log(n/c))$. Recall that we constructed epochs that grew exponentially by a factor of 2 for a total of $\Theta(\log n)$ epochs and the number of large epochs (that is with at least $\max\{\sqrt{n}, c^2\}$ **write** operations) is $\Theta(\log(n/c))$. In the techniques used in Section 3, we are only able to show that $\Omega(b/w)$ cells must be probed from the majority of the epochs. As there are only $\Theta(\log(n/c))$ epochs, there is no hope for us to prove a stronger lower bound t_r with this epoch construction.

Instead, we will use a different epoch construction that is suitable for the scenario where we know that t_w is small. In Lemma 3, we show that, on average, for any large epoch i (that is with $\ell_i \geq \max\{\sqrt{n}, c^2\}$) any **write** operations of future epochs $j \in \{1, \dots, i-1\}$ can only encode $O(t_w w / \hat{k})$ bits about epoch i where \hat{k} is the number of large epochs. It is important that future **write** operations cannot encode a lot of information about epoch i as it forces the final **read** operation to read sufficient information from epoch i directly. However, as we are assuming that t_w is already small, we may increase the number of future operations after epoch i while simultaneously ensuring that future epochs cannot encode too much information about epoch i . With this observation, we hope that we can increase the number of total epochs which allows us to prove $\omega((b/w) \log(n/c))$ lower bounds on t_r as desired. We now materialize these ideas in the next section.

4.1 First Epoch Construction

We will fix $r = \omega(1)$ where $r = O(\log n)$ as the rate at which epochs will increase. In this section, we consider an epoch construction where epochs grow by the rate r every r epochs. That is, the first r epochs will each have r **write** operations. The next r epochs will each have r^2 **write** operations, the next r epochs will each have r^3 **write** operations and so forth. See Figure 4 for a diagram of this epoch construction. Once again, we define ℓ_i to be the size of the i -th epoch and s_i to be the total size of epochs $1, \dots, i$. We note that

$$s_{i-1} \leq (r-1)\ell_i + r(\ell_{i-1} + \ell_{i-2} + \dots) \leq 2r\ell_i.$$

In other words, for epoch i , there will be potentially r times more future operations in comparison to the previous epoch construction of Section 3. On the other hand, we note that the number of epochs with at least $\max\{\sqrt{n}, c^2\}$ **write** operations is $\hat{k} = \Theta(r \log_r(n/(rc^2))) = \Theta(r \log_r(n/c))$, which is $\Theta(r/\log r) = \omega(1)$ times larger than the number of epochs in the construction of Section 3. As a result, this epoch construction matches exactly the requirements that we wanted there to be more epochs which are required to be read by the **read** operation while only sacrificing that there are more future **write** operations for any epoch i . We now present a generalization of Lemma 3 which can be applied for the new epoch constructions that are introduced here and in Section 4.2.




Fig. 4: Diagram of epoch construction of Section 4.1.

Lemma 4. *There exists an epoch i such that $l_i \geq \max\{\sqrt{n}, c^2\}$ and both the following inequalities hold:*

$$\mathbb{E}[|\mathcal{T}_w^{<i}(\mathcal{Q})|] = O\left(\left(\max_e \sum_{j \geq e} \frac{s_e}{l_j}\right) \cdot \frac{t_w}{\hat{k}}\right) \text{ and } \mathbb{E}[|\mathcal{T}_r^i(\mathcal{Q})|] = O(t_r/\hat{k}).$$

Proof. Using the same ideas of Lemma 3, we will show that there exists $\hat{k}/2 + 1$ epochs that satisfy the first statement and $\hat{k}/2 + 1$ epochs that satisfy the second statement. As a result, there exists at least one epoch satisfying both statements.

Pick an epoch i uniformly at random from all \hat{k} epochs with at least $\max\{\sqrt{n}, c^2\}$ **write** operations. Fix $\mathcal{B}_1, \dots, \mathcal{B}_{n-1}$ and \mathcal{R} arbitrarily. We will prove an upper bound on $\mathbb{E}[|\mathcal{T}_w^{<i}(\mathcal{Q})|]$ over the randomness of the location of the **read** operation and the randomly chosen i . As a result, the expectation's upper bound will hold over any distribution of $\mathcal{B}_1, \dots, \mathcal{B}_{n-1}$ and \mathcal{R} . Fix any cell probe performed by **DS** when processing $\text{write}(1, \mathcal{B}_1), \dots, \text{write}(n-1, \mathcal{B}_{n-1})$ and suppose that probe occurs when processing the p_r -th **write** operation to a cell that was last written by the p_w -th **write** operation. Once again, we pick the smallest e such that $p_r - p_w \leq s_e$. Consider any epoch j where $j \leq e - 1$. Note that there are only s_j operations between the **read** operation and the beginning of epoch j . But, since $j \leq e - 1$, we know that $s_j \leq s_{e-1} < p_r - p_w$ meaning that either the probe occurs after the **read** operation or the cell was last written before epoch j . When we fix the location of the **read** operation, we fix the epoch construction. As the boundary of the j -th epoch must occur after the p_w -th operation and

before the p_r -th operation, there are at most s_e good locations for the **read** out of $n/2$ total locations. For any $j \geq e$, this cell probe has probability $2s_e/n$ of contributing to $\mathcal{T}_w^{<j}(\mathcal{Q})$. Therefore, by linearity of expectation over the $(n-1)t_w$ expected cell probes:

$$\mathbb{E} \left[\frac{|\mathcal{T}_w^{<i}(\mathcal{Q})|}{\ell_i} \right] = \frac{1}{\hat{k}} \sum_{j: \ell_j \geq \max\{\sqrt{n}, c^2\}} \mathbb{E} \left[\frac{|\mathcal{T}_w^{<j}(\mathcal{Q})|}{\ell_j} \right] \leq \frac{t_w}{\hat{k}} \left(\max_e \sum_{j \geq e} \frac{2s_e}{\ell_j} \right).$$

Therefore, there exists $\hat{k}/2 + 1$ fixed epochs i such that $\mathbb{E}[|\mathcal{T}_w^{<i}(\mathcal{Q})|/\ell_i]$ over the choice of the **read** location is at most 3 times the above bound.

As $\sum_i \mathbb{E}[|\mathcal{T}_r^i(\mathcal{Q})|] \leq t_r$, there exists $\hat{k}/2 + 1$ epochs i where $\mathbb{E}[|\mathcal{T}_r^i(\mathcal{Q})|] \leq 3t_r$ completing the proof.

We can now plug the upper bounds of Lemma 4 into Lemma 2.

Theorem 3. *Let \mathbf{DS} be an (ϵ, δ) -differentially private RAM for n b -bit values implemented over w -bit cells. Assuming that $\epsilon = O(1)$ and $0 \leq \delta \leq 1/3$, \mathbf{DS} has failure probability at most $1/3$ and $w = \Omega(\log \log n)$, then*

$$t_w = o((b/w) \log(n/c)) \implies t_r = \omega((b/w) \log(n/c)).$$

Proof. Recall we get the following inequality by applying convexity to the inequality of Lemma 2 and noting that $c/\ell_i = O(1)$ for our choices i :

$$\frac{\mathbb{E}[|\mathcal{T}_w^{<i}(\mathcal{Q})|]}{\ell_i} \left(w + \log \frac{t_w s_{i-1}}{\mathbb{E}[|\mathcal{T}_w^{<i}(\mathcal{Q})|]} \right) + \mathbb{E}[|\mathcal{T}_r^i(\mathcal{Q})|] \left(w + \log \frac{t_r}{\mathbb{E}[|\mathcal{T}_r^i(\mathcal{Q})|]} \right) = \Omega(b).$$

By applying Lemma 4, we get that $\mathbb{E}[|\mathcal{T}_r^i(\mathcal{Q})|] = O(t_r \log r / (r \log(n/c)))$ and $\mathbb{E}[|\mathcal{T}_w^{<i}(\mathcal{Q})|/\ell_i] = O(t_w \log r / \log(n/c))$ since

$$\left(\frac{s_j}{\ell_j} + \frac{s_j}{\ell_{j+1}} + \dots \right) \leq 2r \sum_{j \geq 0} \frac{1}{r^j} = O(r).$$

Plugging into the inequality above and assuming that $w = \Omega(\log \log n)$,

$$t_w + (t_r/r) = \Omega((b/w) \log(n/c) / \log r) \implies t_r = \Omega((b/w) \log(n/c) r / \log r)$$

as $t_w = o((b/w) \log(n/c))$. Since $r / \log r = \omega(1)$, we complete the proof.

4.2 Second Epoch Construction

In this section, we deal with the opposite scenario when we assume that $t_r = o((b/w) \log(n/c))$ and want to show that $t_w = \omega((b/w) \log(n/c))$. The same intuition from the previous section can be used for this situation. To show that t_w has to be very large, we will need to require that for any epoch i , the total number of future **write** operations in epochs $j \in \{1, \dots, i-1\}$ is small. If for any epoch i , the number of future **write** operations after epoch i is small and

the **read** operation also cannot perform many cell probes into epoch i , then each future **write** operation must encode a large amount of information about epoch i which will be used by the **read** operation. As a result, we can prove a large lower bound on t_w . We now describe the epoch construction which enables us to prove such strong lower bounds on t_w .

Fix the $r = \omega(1)$ where $r = O(\log n)$ as the rate at which epochs will increase once again. This epoch construction will increase each epoch's number of **write** operations by r . So, the first epoch will have r **write** operations, the second epoch will have r^2 **write** operations and so forth. In this case,

$$s_{i-1} = \ell_{i-1} + \ell_{i-2} + \dots \leq \frac{1}{r}(\ell_i + \ell_{i-1} + \dots) \leq 2\ell_i/r.$$

As a result, the number of future operations is $\Theta(1/r)$ times smaller than the epoch construction of Section 3. The number of epochs with at least $\max\{\sqrt{n}, c^2\}$ **write** operations is $\hat{k} = \Theta(\log_r n)$.

Theorem 4. *Let **DS** be an (ϵ, δ) -differentially private RAM for n b -bit values implemented over w -bit cells. Assuming that $\epsilon = O(1)$ and $0 \leq \delta \leq 1/3$, **DS** has failure probability at most $1/3$ and $w = \Omega(\log \log n)$, then*

$$t_r = o((b/w) \log(n/c)) \implies t_w = \omega((b/w) \log(n/c)).$$

Proof. By applying Lemma 4, we get that $\mathbb{E}[|\mathcal{T}_r^i(\mathcal{Q})|] = O(t_r \log r / \log(n/c))$ and $\mathbb{E}[|\mathcal{T}_w^{<i}(\mathcal{Q})|/\ell_i] = O(t_w \log r / r \log(n/c))$ since

$$\left(\frac{s_j}{\ell_j} + \frac{s_j}{\ell_{j+1}} + \dots \right) \leq 2 \sum_{j \geq 0} \frac{1}{r^j} = O(1).$$

Plugging into the inequality of Lemma 2 after applying convexity and noting that $c/\ell_i = O(1)$ and $w = \Omega(\log \log n)$,

$$(t_w/r) + t_r = \Omega((b/w) \log(n/c) / \log r) \implies t_w = \Omega((b/w) \log(n/c) r / \log r)$$

since $t_r = o((b/w) \log(n/c))$. Noting that $r / \log r = \omega(1)$ completes our proof.

5 Discussion

In this section, we discuss two extensions that follow from our lower bound techniques.

*Privacy only for **read** operations.* Our techniques only enforce the requirements of differential privacy for a single **read** operation. Therefore, our lower bounds would also apply *differentially private-read RAMs* where differential privacy is guaranteed only sequences of operations that differ in exactly one **read** operation. This might be important in scenarios where the indices of **write** operations are not sensitive (or may be public) but only the indices of **read** operations need to be protected. Once again, this weakening of security does not suffice to get around the $\Omega(\log(n/c))$ bandwidth overhead lower bounds.

Larger Values of δ . In the proofs of Section 3 and 4, it is assumed that $\delta \leq 1/3$. Most practical scenarios require δ must be negligible in n , so the above results suffice. For theoretical exploration, we note that the proofs can be extended for any constant δ that is strictly smaller than 1. In particular, for any constant ρ such that $\delta < \rho < 1$ and by picking a sufficiently large enough constant C , we can prove that $\Pr[\mathcal{Z}_i(\mathcal{Q}(\text{idx})) \geq b/C] \geq \rho$ which is a variation of Lemma 1 that suffices to prove a lower bound for δ .

6 Conclusion

In this work, we show that the $\Omega(\log(n/c))$ bandwidth overhead lower bound for the array maintenance problem with obliviousness extends to the weaker notion of differential privacy with reasonable privacy budgets of $\epsilon = O(1)$ and $\delta \leq 1/3$. The result is surprising as differentially private RAM seems, at first, to provide significantly weaker privacy compared to the obliviousness guarantees of ORAM. Yet, differential privacy does not allow any asymptotic improvements in efficiency. This leads to the following natural open question: Does there exist a natural, weaker notion of privacy that enables $o(\log(n/c))$ bandwidth overhead for the array maintenance problem?

References

1. E. Boyle, K.-M. Chung, and R. Pass. Oblivious parallel RAM and applications. In *Theory of Cryptography Conference*, pages 175–204. Springer, 2016.
2. 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.
3. D. Cash, P. Grubbs, J. Perry, and T. Ristenpart. Leakage-abuse attacks against searchable encryption. In *Proceedings of the 22nd ACM SIGSAC conference on computer and communications security*, pages 668–679. ACM, 2015.
4. T.-H. H. Chan, Y. Guo, W.-K. Lin, and E. Shi. Oblivious hashing revisited, and applications to asymptotically efficient ORAM and OPRAM. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 660–690. Springer, 2017.
5. B. Chen, H. Lin, and S. Tessaro. Oblivious parallel RAM: improved efficiency and generic constructions. In *Theory of Cryptography Conference*, pages 205–234. Springer, 2016.
6. K.-M. Chung, Z. Liu, and R. Pass. Statistically-secure ORAM with $\tilde{O}(\log^2 n)$ overhead. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 62–81. Springer, 2014.
7. I. Damgård, S. Meldgaard, and J. B. Nielsen. Perfectly secure oblivious RAM without random oracles. In *Theory of Cryptography Conference*, pages 144–163. Springer, 2011.
8. S. Devadas, M. van Dijk, C. W. Fletcher, L. Ren, E. Shi, and D. Wichs. Onion ORAM: A constant bandwidth blowup oblivious RAM. In *Theory of Cryptography Conference*, pages 145–174. Springer, 2016.

9. C. Dwork. A firm foundation for private data analysis. *Communications of the ACM*, 54(1):86–95, 2011.
10. C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography Conference*, pages 265–284. Springer, 2006.
11. C. Dwork, A. Roth, et al. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014.
12. M. Fredman and M. Saks. The cell probe complexity of dynamic data structures. In *Proceedings of the twenty-first annual ACM symposium on Theory of computing*, pages 345–354. ACM, 1989.
13. S. Garg, S. Lu, R. Ostrovsky, and A. Scafuro. Garbled RAM from one-way functions. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pages 449–458. ACM, 2015.
14. C. Gentry, S. Halevi, S. Lu, R. Ostrovsky, M. Raykova, and D. Wichs. Garbled RAM revisited. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 405–422. Springer, 2014.
15. O. Goldreich. Towards a theory of software protection and simulation by oblivious RAMs. In *Proceedings of the nineteenth annual ACM symposium on Theory of computing*, pages 182–194. ACM, 1987.
16. O. Goldreich and R. Ostrovsky. Software protection and simulation on oblivious RAMs. *Journal of the ACM (JACM)*, 43(3):431–473, 1996.
17. M. T. Goodrich and M. Mitzenmacher. Privacy-preserving access of outsourced data via oblivious RAM simulation. In *International Colloquium on Automata, Languages, and Programming*, pages 576–587. Springer, 2011.
18. M. T. Goodrich, M. Mitzenmacher, O. Ohrimenko, and R. Tamassia. Privacy-preserving group data access via stateless oblivious RAM simulation. In *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*, pages 157–167. Society for Industrial and Applied Mathematics, 2012.
19. 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.
20. E. Kushilevitz, S. Lu, and R. Ostrovsky. On the (in) security of hash-based oblivious RAM and a new balancing scheme. In *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*, pages 143–156. Society for Industrial and Applied Mathematics, 2012.
21. K. G. Larsen. The cell probe complexity of dynamic range counting. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 85–94. ACM, 2012.
22. K. G. Larsen and J. B. Nielsen. Yes, there is an Oblivious RAM lower bound! In *Annual International Cryptology Conference*, pages 523–542. Springer, 2018.
23. K. G. Larsen, O. Weinstein, and H. Yu. Crossing the logarithmic barrier for dynamic boolean data structure lower bounds. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 978–989. ACM, 2018.
24. S. Lu and R. Ostrovsky. Black-box parallel garbled RAM. In *Annual International Cryptology Conference*, pages 66–92. Springer, 2017.
25. I. Mironov, O. Pandey, O. Reingold, and S. Vadhan. Computational differential privacy. In *Advances in Cryptology-CRYPTO 2009*, pages 126–142. Springer, 2009.
26. S. Patel, G. Persiano, M. Raykova, and K. Yeo. PanORAMa: Oblivious RAM with logarithmic overhead. Cryptology ePrint Archive, Report 2018/373, 2018. <https://eprint.iacr.org/2018/373>.

27. B. Pinkas and T. Reinman. Oblivious RAM revisited. In *Annual Cryptology Conference*, pages 502–519. Springer, 2010.
28. M. Pătraşcu. *Lower bound techniques for data structures*. PhD thesis, Massachusetts Institute of Technology, 2008.
29. M. Pătraşcu and E. D. Demaine. Logarithmic lower bounds in the cell-probe model. *SIAM Journal on Computing*, 35(4):932–963, 2006.
30. E. Stefanov, E. Shi, and D. Song. Towards practical oblivious RAM. *arXiv preprint arXiv:1106.3652*, 2011.
31. 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 *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*, pages 299–310. ACM, 2013.
32. R. R. Toledo, G. Danezis, and I. Goldberg. Lower-cost ϵ -private information retrieval. *Proceedings on Privacy Enhancing Technologies*, 2016(4):184–201, 2016.
33. S. Wagh, P. Cuff, and P. Mittal. Root ORAM: A tunable differentially private oblivious RAM. *arXiv preprint arXiv:1601.03378*, 2016.
34. X. S. Wang, K. Nayak, C. Liu, T. Chan, E. Shi, E. Stefanov, and Y. Huang. Oblivious data structures. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, pages 215–226. ACM, 2014.
35. M. Weiss and D. Wichs. Is there an Oblivious RAM lower bound for online reads? Cryptology ePrint Archive, Report 2018/619, 2018. <https://eprint.iacr.org/2018/619>.
36. A. C.-C. Yao. Should tables be sorted? *Journal of the ACM (JACM)*, 28(3):615–628, 1981.