Perfectly Secure Oblivious RAM with Sublinear Bandwidth Overhead

Michael Raskin[1] and Mark Simkin[2]

[1] Technical University of Munich, `raskin@mccme.ru`
[2] Aarhus University, `simkin@cs.au.dk`

Abstract. Oblivious RAM (ORAM) has established itself as a fundamental cryptographic building block. Understanding which bandwidth overheads are possible under which assumptions has been the topic of a vast amount of previous works. In this work, we focus on perfectly secure ORAM and we present the *first* construction with sublinear bandwidth overhead *in the worst-case*. All prior constructions with perfect security require linear communication overhead in the worst-case and only achieve sublinear bandwidth overheads in the amortized sense. We present a fundamentally new approach for construction ORAM and our results significantly advance our understanding of what is possible with perfect security.

Our main construction, Lookahead ORAM, is perfectly secure, has a worst-case bandwidth overhead of $\mathcal{O}\left(\sqrt{N}\right)$, and a total storage cost of $\mathcal{O}\left(N\right)$ on the server-side, where n is the maximum number of stored data elements. In terms of concrete server-side storage costs, our construction has the smallest storage overhead among all perfectly and statistically secure ORAMs and is only a factor 3 worse than the most storage efficient computationally secure ORAM. Assuming a client-side position map, our construction is the first, among all ORAMs with worst-case sublinear overhead, that allows for a $\mathcal{O}\left(1\right)$ online bandwidth overhead without server-side computation. Along the way, we construct a conceptually extremely simple statistically secure ORAM with a worst-case bandwidth overhead of $\mathcal{O}\left(\sqrt{N}\frac{\log N}{\log \log N}\right)$, which may be of independent interest.

1 Introduction

More and more sensitive data is stored online. A basic attempt to keep data private while storing it online on an untrusted server is to simply encrypt each data entry. Unfortunately, this is not always sufficient. For instance, Islam et. al. [IKK12] showed that by observing the access patterns induced by encrypted search queries over an encrypted database, an honest-but-curious server storing the database, could learn significant amounts of information about the queries' contents. A more viable approach is to not only encrypt the data, but also hide the access patterns. Oblivious RAM (ORAM) is a cryptographic primitive that allows a client to do exactly this, at the cost of some bandwidth and storage overhead. It enables a client to outsource his data to an untrusted server inside

of an ORAM data structure, and then read and write to his dataset without revealing the position that was accessed or the operation that was performed.

Goldreich and Ostrovsky [Gol87, GO96] first introduced the notion of ORAM, presented the first constructions thereof, and proved the first lower bound of $\Omega(\log n)$ on the bandwidth overhead for a certain type of constructions, where n the maximum number of data elements stored in the data structure. Boyle and Naor [BN16] revisit this lower bound proof and highlight that it only holds for statistically secure ORAMS that behave in a "balls-and-bins" fashion. The same lower bound without these restriction was recently proven by Larsen and Nielsen [LN18].

Understanding which upper bounds can be achieved in which setting has been the topic of numerous works [OS97, GMOT11, SCSL11, DMN11, SSS12, SvS+13, CNS18, PPRY18, AKL+18]. Most commonly, these works measure bandwidth overhead in one of two ways. Either they consider the worst-case overhead, meaning the largest overhead any one operation on the ORAM data structure can incur, or they consider the amortized overhead, meaning the average overhead per operation in a longer sequence of operations. The best known upper bound for worst-case bandwidth overhead is due to Stefanov et al. [SvS+13], who present a statistically secure construction with $\mathcal{O}\left(\log^2 N\right)$ overhead[3]. The best known upper bound for amortized bandwidth overhead is due to Asharov et al. [AKL+18], who present a computationally secure construction with bandwidth overhead of $\mathcal{O}\left(\log N\right)$, which matches the lower. The first perfectly secure ORAM construction is due to Damgård et al. [DMN11] and has an amortized bandwidth overhead of $\mathcal{O}\left(\log^3 N\right)$ and a multiplicative storage overhead of $\mathcal{O}\left(\log N\right)$. This was recently improved upon by Chan et al. [CNS18], who present a perfectly secure construction with the same amortized bandwidth overhead and a storage overhead of $\mathcal{O}\left(1\right)$. In the multi-server setting, where the ORAM data structure is distributed among several non-colluding servers, Chan et al. [CKN+18] present a perfectly secure construction with worst-case bandwidth overhead of $\mathcal{O}\left(\log^2 N\right)$. Since all existing perfectly secure single-server constructions require linear worst-case bandwidth overhead, we pose the natural question:

Can we construct a perfectly secure single-server ORAM with sublinear worst-case bandwidth overhead?

We believe this is an important theoretical question. It is well known that randomization and computational assumptions are powerful tools in algorithm and protocol design. It is a common theme in research to investigate the power of these tools for specific problems by understanding the upper bounds that we can achieve with and without them. In this work we make a significant step towards understanding what ORAM bandwidth overhead can be achieved in the worst-case with perfect security and without relying on randomization[4]

---

[3] For a larger data block size of $\mathcal{O}\left(\log^2 N\right)$ they even achieve an overhead of $\mathcal{O}\left(\log N\right)$ data blocks.

[4] Our main construction is using randomness exclusively for the sake of security, but not for efficiency. We believe this is unavoidable.

or computational assumptions. See [CNS18] for a further discussion about the importance of perfectly secure ORAM.

1.1 Our Contribution

We present the *first* construction of perfectly secure Oblivious RAM with sublinear worst-case bandwidth overhead and, furthermore, we also make the following contributions:

Novel approach to constructing ORAM. We present a fundamentally new approach for constructing ORAM. Somewhat surprisingly, and despite the large amount of research interest that ORAM has received, all existing constructions are based on a handful of conceptually different approaches. We believe it is of theoretical and practical interest to explore new ways to construct this primitive. In this work, we present two new constructions.

Our first construction is conceptually extremely simple. It is statistically secure, meaning that even a computationally unbounded adversary cannot break the obliviousness guarantees and all operations on the data structure succeed with an overwhelming probability. It has a worst-case bandwidth overhead of $\mathcal{O}\left(\sqrt{N}\frac{c+\log t+\log N}{\log\left(c+\log t+\log N\right)}\right)$, where N is the maximum number of data blocks to be stored, t is an upper bound on the number of accesses, and c is the correctness parameter that provides an upper bound of 2^{-c} on the failure probability. To the best of our knowledge, it is one of the conceptually simplest known ORAM constructions to date. The underlying logic is easy to implement and the proof of security is straightforward. Our main construction, called Lookahead ORAM, is loosely based on our first construction. It has a worst-case bandwidth overhead of $\mathcal{O}\left(\sqrt{N}\right)$ and is *perfectly* secure, in the sense that every operation on the ORAM data structure succeeds with probability 1 and obliviousness is guaranteed against an unbounded adversary. The hidden constants behind the big-O notation are small and our construction is significantly faster in the worst-case than the fastest perfectly secure single-server ORAM construction of Chan et al. [CNS18] in the amortized case for any practical parameter range. For instance, for $N = 2^{20}$ our construction has a worst-case bandwidth overhead of less than 7 000 data blocks, whereas their construction has a amortized bandwidth overhead of around 160 000 data blocks[5].

Small concrete storage overhead. Assuming a client-side position map, Lookahead ORAM has the smallest concrete storage overhead among all existing ORAM constructions with sublinear worst-case bandwidth overhead. Our construction only incurs an additive server-side storage overhead of $\sqrt{2N}$ data

---

[5] Our estimate of Chan et al.'s construction is computed by instantiating it with Batcher's Bitonic sort [Bat68] and a hidden constant of 1. For our construction we took the concrete parameters one obtains assuming a server-side position map.

ORAM	Client storage		Server	Bandwidth	Online	Section
	persistent	temporary	storage	overhead	overhead	
Matrix	$\mathcal{O}(N)$	$\mathcal{O}(1)$	$\approx 18N + \frac{N \log N}{2}$	$\mathcal{O}\left(\sqrt{N} \log N\right)$	–	3
Matrix	$\mathcal{O}(1)$	$\mathcal{O}(1)$	$\approx 142N + 2N \log N$	$\mathcal{O}\left(\sqrt{N} \log N\right)$	–	5
Lookahead	$\mathcal{O}(N)$	$\mathcal{O}\left(\sqrt{N}\right)$	$N + 2\sqrt{N}$	$\mathcal{O}\left(\sqrt{N}\right)$	$\mathcal{O}(1)$	4
Lookahead	$\mathcal{O}(1)$	$\mathcal{O}\left(\sqrt{N}\right)$	$6N + 12\sqrt{N}$	$\mathcal{O}\left(\sqrt{N}\right)$	–	5
Lookahead	$\mathcal{O}(1)$	$\mathcal{O}(1)$	$6N + 12\sqrt{N}$	$\mathcal{O}\left(\sqrt{N} \log \sqrt{N}\right)$	–	5.1

Fig. 1. Overview of the different parameter settings for Matrix and Lookahead ORAM. For Matrix ORAM we crudely estimate the parameters for $c = 20$ and $t = 15$. All overheads are stated in data blocks and assume block size $\Theta(\log N)$. Asymptotically line 3 with a client-side storage of $\mathcal{O}(N)$ makes little sense. From a practical perspective, however, the $\mathcal{O}(1)$ online overhead is a powerful feature and the concrete client-side storage is significantly smaller than the concrete amount of data stored on the server-side. The asymptotical behaviour also improves in case of faster growth of block size.

blocks. A small storage overhead is particularly beneficial in outsourced storage settings where data owners have to pay their storage provider for the storage they consume. At the cost of a slightly increased total server-side storage, namely $6N + 12\sqrt{N}$, we can reduce the client-side storage of Lookahead ORAM to $\mathcal{O}(1)$. In this case, our construction has the smallest storage overhead among all statistically and perfectly secure ORAMs and is only a multiplicative factor of around 3 larger than the most storage efficient computationally secure ORAMs [GMOT11] with sublinear worst-case bandwidth overhead. For a more detailed discussion of the concrete server-side storage costs of Lookahead ORAM see Section 4.5 and Section 5. Lookahead ORAM, by default, requires the client to have a temporary client-side storage of $\mathcal{O}\left(\sqrt{N}\right)$ during each access. We show how to reduce the temporary client-side storage to $\mathcal{O}(1)$ at the cost of increasing the worst-case bandwidth overhead to $\mathcal{O}\left(\sqrt{N} \log \sqrt{N}\right)$ in Section 5.1. We illustrate the different parameter options and their efficiency in Figure 1.

Constant online bandwidth overhead. One approach to circumvent the $\Omega(\log N)$ lower bound on the bandwidth overhead was introduced by Boneh et al. [BMP11] and then improved upon in [DSS14, RFK+15]. Their main idea was to split the total bandwidth overhead into two parts. The first part, the so called online overhead, is the amount of data that needs to be transmitted between the client and the server to retrieve a desired data element obliviously. The second part, the offline overhead, is the amount of data that needs to be transmitted between the two parties to ensure obliviousness of future accesses. One can think of the offline overhead as background work that, usually, moves around encrypted data elements in the ORAM data structure to ensure the desired obliviousness guarantees. Splitting the total bandwidth overhead this way and then minimizing the online overhead has practical advantages. It allows

the client to efficiently retrieve data from the server without much latency during bursts of requests and then do the background work during quieter phases.

In [BMP11], Boneh et al. presented a computationally secure construction, for a primitive, which is strongly related to ORAM and has $\mathcal{O}\left(\log N\right)$ online and $\mathcal{O}\left(\sqrt{N}\log N\right)$[6] worst-case overhead. In [DSS14, RFK+15] this idea of splitting the total overhead has been further refined, and computationally secure constructions that achieve an online bandwidth overhead of $\mathcal{O}\left(1\right)$ are presented. However, these constructions require some server-side computation during the online phase, which renders these solutions not applicable for "raw" storage providers that do not support these ORAM constructions explicitly. That is, our construction works in combination with arbitrary storage providers like Dropbox or Google Drive, whereas the constructions from [DSS14, RFK+15] only work with storage providers that explicitly implement their given scheme.

If the client stores the position map and a small client-side storage that can hold up to $\sqrt{2N}$ data blocks locally, then Lookahead ORAM allows the client to obliviously retrieve arbitrary many elements from the ORAM data structure with no bandwidth overhead in the online phase and no server-side computation in the online or offline phase. That is, in the online phase, the client can directly download the desired elements from the server. Our construction is the first to provide such a feature among all computationally, statistically, and perfectly secure ORAMs with sublinear worst-case bandwidth overhead.

To provide a better feeling for how expensive it is to store the stash and the position map locally, consider a 1GB database with a block size of 1KB. To be able to make use of our online overhead feature, the client would need to store a roughly 2MB stash and a 8MB position map locally. As mentioned before, if the client chooses to not use the minimizing online overhead feature, it can reduce its persistent storage to $\mathcal{O}\left(1\right)$.

Implementation and Evaluation. Even though the main focus of this paper is the theoretical feasibility of perfectly secure ORAM with sublinear worst-case bandwidth overhead, we also provide a implementation of Lookahead ORAM for the sake of completeness. A description of our implementation and our performance benchmarks, can be found in Section 6.

Attack on [GMSS16]. We identify a flaw in the ORAM construction of [GMSS16] and outline an attack that breaks the claimed obliviousness guarantees in Section A. We have contacted the authors and they have acknowledged our attack.

---

[6] This worst-case complexity is slightly different from the original paper. The paper has a superlinear worst-case overhead due to an expensive reshuffling phase, but when splitting shuffling over $\sqrt{N}$ accesses, one can achieve the stated complexity

1.2 Other Related Work

A vast amount of works have contributed to our current understanding of ORAM. In this section we merely provide a high-level overview of the works that are directly related to our work.

In order to achieve practical efficiency and overcome the $\Omega(\log N)$ lower bound, several works have looked at different refinements of the classical ORAM notion in the client server model. Path-PIR [MBC14] uses server-side computations to achieve a practically very small, yet still poly-logarithmic bandwidth overhead. In [AKST14], Apon et al. formally define the notion of Verifiable Oblivious Storage, which generalizes the notion of ORAM by allowing the server to perform computations, and show that the ORAM lower bound does not apply to their setting by providing a scheme with constant overhead per access based on Fully Homomorphic Encryption. In [DDF+16] a scheme, called Onion ORAM, is presented that breaks the lower bound, but only relies on additively homomorphic encryption. In this work we will only focus on the classical notion of ORAM that does *not* allow server-side computation.

Demertzis et al. [DPP18] present a *computationally* secure ORAM construction with worst-case bandwidth overhead $\mathcal{O}\left(N^{1/3}\right)$ and perfect correctness. Several recent works cite Demertzis et al. and claim that their construction is perfectly secure. This is *not* correct and this claim is not made by the authors of that paper either. Their construction is a modification of the square-root ORAM construction and requires the client to store a random permutation, which represents the position map. This position map can only be stored succinctly by using a pseudorandom function. Therefore, their construction is either only computationally secure, or requires linear client-side storage, or requires linear bandwidth overhead.

Lastly, a work by Gordon et al. [GMSS16] presents an ORAM construction that may seem *superficially* similar to ours. However, our work significantly differs from theirs in terms of performance guarantees we achieve, underlying ideas we present, and security we obtain. Their construction, called Matrix-ORAM, arranges the data elements in a fixed number of rows. The size of each row linearly depends on the size of the total database and each row has its own stash. Accesses to their data structure are performed in a conceptually and concretely different manner to ours. The authors claim a logarithmic bandwidth overhead. In contrast, our main construction has a rectangular shape [7], and has a bandwidth overhead of $\mathcal{O}\left(\sqrt{N}\right)$.

We discovered a flaw in their construction and present a concrete attack on their scheme. This flaw is discussed in detail in Section A.

2 Preliminaries

On a high level, the ORAM security definition assumes a honest-but-curious server and says that for any two data access sequences, the corresponding access

---

[7] One may even say they look matrix shaped

sequences to the ORAM data structure should be indistinguishable. The security definition is taken almost verbatim from [SSS12].

Definition 1. *(Security Definition) Let*

$$((\mathsf{op}_1, \mathsf{a}_1, \mathsf{data}_1), \ldots, (\mathsf{op}_M, \mathsf{a}_M, \mathsf{data}_M))$$

be a data request sequence of length M, where each op_i is either a $\mathsf{read}(\mathsf{a}_i)$ or a $\mathsf{write}(\mathsf{a}_i, \mathsf{data})$ operation. Let $\mathsf{oram}(\overrightarrow{y})$ denote the (possibly randomized) sequence of accesses to the remote storage given the sequence of data requests $\overrightarrow{y}$. An ORAM construction is said to be secure if for any two data request sequences $\overrightarrow{y}$ and $\overrightarrow{z}$ of the same length, their access patterns $\mathsf{oram}(\overrightarrow{y})$ and $\mathsf{oram}(\overrightarrow{z})$ are computationally indistinguishable and the construction is correct in the sense that it returns on input $\overrightarrow{y}$ data that is consistent with $\overrightarrow{y}$ with probability at least $1 - 2^{-c}$. We call c the correctness parameter.

Position Map All known ORAM schemes need to maintain a position map of size $\mathcal{O}(N)$ that keeps track of the ordering of elements inside the ORAM data structure on the server. For the sake of simplicity we will assume that the client stores the full position map locally. From a practical point of view, this seems to be a reasonable assumption in many client-server settings. For example, the position map of a 1GB database containing 1KB blocks is only around 8MB large. From a theoretical point of view, to reduce the client's persistent storage to $\mathcal{O}(1)$, both of our constructions can be combined with the well-known approach of recursively storing the position map in a sequence of smaller ORAMs, which was first introduced in [SCSL11]. Recursively storing the position map on the server increases the number of round-trips per access to $\mathcal{O}(\log N)$, but it does not change the asymptotic bandwidth overheads of our constructions. We explain how to combine our main construction with the recursive approach in detail in Section 5.

Block Size If we want to use the recursive ORAM approach mentioned above to store the position map on the server-side, then the data blocks need to be $\Omega(\log N)$ large. In the setting, where the client stores the position map locally, we do not make any assumptions about the data block size. However, for the construction to be useful, the data blocks on the server should in total be larger than the position map that the client stores locally.

Integrity The ORAM security definition assumes the server to be honest-but-curious. Similar to previous works [SvS+13], our construction can, at the cost of giving up perfect for computational security, be extended to prevent tampering of an actively malicious server by using a Merkle Tree on top of our ORAM data structures.

3 A Simple Matrix Bucket ORAM

In this section, we will present a very simple oblivious RAM construction with reasonable efficiency and a simple proof of security. To the best of our knowledge this is one of the, arguably, simplest ORAM constructions known to date. Apart from being interesting on its own, it will also serve a stepping stone towards our main construction by introducing some of the ideas behind our main approach.

Initially we are given an array A of length N of data elements. To initialize our scheme, we create a empty $\sqrt{N} \times \sqrt{N}$ matrix C, in which each matrix cell is a bucket of size w. We randomly (and independently) assign each element from A to a bucket in C. Once all elements from A are distributed among buckets in C, we encrypt each bucket separately, and store the matrix on the server. For the sake of simplicity, we assume that the client stores a position map σ that maps indices of elements from A to columns of C locally.

$\mathrm{Init}(A)$	$\mathrm{Access}(\ell, x)$
1 $C \leftarrow \mathtt{initBucketMatrix}()$	**1** $c \leftarrow \mathtt{readColumn}(C, \sigma(\ell))$
2 $\sigma \leftarrow \mathtt{initPosMap}()$	**2** $(i, j) \leftarrow \mathtt{pickRandBucket}()$
3 **for** $\ell = 1 \ldots N$ **do**	**3** $r \leftarrow \mathtt{readRow}(C, i)$
4 $\quad (i, j) \leftarrow \mathtt{pickRandBucket}()$	**4** $r_{\mathsf{dec}} \leftarrow \mathtt{decrypt}(r)$
5 $\quad C[i, j].\mathtt{put}(A[\ell] \| \ell)$	**5** $c_{\mathsf{dec}} \leftarrow \mathtt{decrypt}(c)$
6 $\quad \sigma(\ell) = j$	**6** $data \leftarrow \mathtt{popData}(\ell, c_{\mathsf{dec}})$
7 $C \leftarrow \mathtt{encryptBuckets}(C)$	**7** $\mathtt{putData}(x \| \ell, r_{\mathsf{dec}}, j)$
	8 $\sigma(\ell) = j$
	9 $c^* \leftarrow \mathtt{encrypt}(c_{\mathsf{dec}})$
	10 $r^* \leftarrow \mathtt{encrypt}(r_{\mathsf{dec}})$
	11 $\mathtt{writeBack}(c^*, r^*)$
	12 **return** $data$

Fig. 2. Pseudocode of simple Matrix ORAM construction

To obliviously access some element with index ℓ in A, we need to access column $\sigma(\ell)$ in C. In addition, we need to pick a uniformly random bucket (i, j) in the matrix and obtain row i. We find the element with index ℓ in the retrieved column and perform our desired operation (read or write). We then remove element ℓ from its current bucket, put it into bucket (i, j), re-encrypt all retrieved buckets, and write back the retrieved row and column. Lastly, we update the position map to point to the new column that stores ℓ, i.e., set $\sigma(\ell) = j$. The pseudocode implementing this construction is given in Figure 2

3.1 Security

We prove the following theorem

Theorem 1. *Let* $\mathcal{E} = (\mathsf{gen}, \mathsf{encrypt}, \mathsf{decrypt})$ *be an* IND-CPA *secure encryption scheme. Then the construction in Figure 2 is a statistically secure ORAM scheme with* $\mathcal{O}\left(\sqrt{N}\frac{c+\log t+\log N}{\log\,(c+\log t+\log N)}\right)$ *bandwidth overhead and a total storage cost of* $\mathcal{O}(N\log N)$ *data blocks, where N is the number of data elements, c the correctness parameter, and t is the upper bound on the number of accesses.*

Proof. The key idea of why the proposed scheme is oblivious stems from the basic observation that every column intersects with every row. Intuitively, this means that if we obliviously write an element into some uniformly random position in a row, then, from an adversarial point of view, every column is equally likely a potential candidate for reading that element in a future access. In our scheme, whenever we read an element through a column access, we move it to a new uniformly random bucket and, in particular, a new uniformly random column, through a row access. Importantly, the movement of each element is completely independent of the access history and the other elements residing in the matrix. From these observations it is straightforward to see that the proposed scheme is oblivious.

What remains to show is the relation between the bucket size and the correctness parameter, i.e. we want to pick our buckets sufficiently large such that a bucket overflows with negligible (in the correctness parameter) probability. Towards this goal, we make an observation that simplifies our analysis. Let $\mathsf{Exp}_{\mathsf{move}}^{N,t}$ be the experiment of first throwing N balls into N buckets once and then picking up a random ball from a random bucket and moving it to a new random bucket t times. This experiment expresses the actual movement of data during t many accesses in our oblivious ram construction. Let $\mathsf{Exp}_{\mathsf{throw}}^{N}$ be the experiment of throwing N balls into N initially empty buckets of capacity w. Let $\mathsf{Load}_{i}^{>w}$ denote the event that bucket with index i at some point in time has more than w many elements in it and $\mathsf{Load}^{>w}$ the event that this happens to any of the buckets. We will analyze the probability of the event of one bucket overflowing in $\mathsf{Exp}_{\mathsf{throw}}^{N}$ and use the following lemma to put $\mathsf{Exp}_{\mathsf{throw}}^{N}$ and $\mathsf{Exp}_{\mathsf{move}}^{N,t}$ into relation.

Lemma 1. *Let $t > 0$, then*

$$\Pr[\mathsf{Load}_{i}^{>w}|\mathsf{Exp}_{\mathsf{move}}^{N,t}] \leq t \cdot \Pr[\mathsf{Load}_{i}^{>w}|\mathsf{Exp}_{\mathsf{throw}}^{N}]$$

Proof. Given N balls and N bins, there are N^N different possibilities to distribute the balls among the bins. Let us call each way to distribute the balls a constellation. Let X be one arbitrary but fixed constellation among them and, since all of them are equally probable, we have $\Pr[X|\mathsf{Exp}_{\mathsf{throw}}^{N}] = \frac{1}{N^N}$. Let us now consider constellations, which are one ball move away from X. There are exactly $N^2 - N$ such constellations, because we can select any of the N balls, and pick any of $N-1$ buckets distinct from the current bucket of the selected ball. Selecting a random ball uniformly and moving it to a random bucket yields the original constellation with probability $\frac{1}{N^2}$, and each of the neighbouring constellations with probability $\frac{1}{N^2}$. As all ball moves are reversible, each constellation can be obtained from $N^2 - N$ other constellations. The probability of obtaining a constellation after a uniform selection of constellation and a single random

ball move is therefore equal to $\frac{N^2-N}{N^2}\frac{1}{N^N} + \frac{1}{N}\frac{1}{N^N} = \frac{1}{N^N}$. The lemma follows by induction over t and then applying the union bound.

Using this lemma it is sufficient to upper bound the probability of a bucket overflowing in the experiment $\mathsf{Exp}_{\mathsf{throw}}^{N}$ and then apply the union bound over all buckets. Let us first look at the probability of some single bucket i overflowing by one element after $\mathsf{Exp}_{\mathsf{throw}}^{N}$, i.e. the probability of a bucket containing (exactly) $z = w + 1$ balls after throwing N balls into N buckets at random once. In the analysis we assume N to be sufficiently large, i.e. N should be large enough for our bucket size w to be at least 8, so that our inequalities work out. In the following calculation we will use two inequalities. First, $\forall x \geq 0$ it holds that $\left(1 - \frac{1}{x}\right)^x \leq e^{-1}$ and, secondly, $\forall 0 \leq k \leq n$ it holds that $\left(\frac{n}{k}\right)^k \leq \binom{n}{k} \leq \left(\frac{en}{k}\right)^k$.

$$
\begin{aligned}
\Pr[\mathsf{Load}_i^z | \mathsf{Exp}_{\mathsf{throw}}^{N}] &= \binom{N}{z}\left(\frac{1}{N}\right)^z\left(1 - \frac{1}{N}\right)^{N-z} \\
&\leq \left(\frac{eN}{z}\right)^z N^{-z}\left(\left(1 - \frac{1}{N}\right)^N\right)^{1-\frac{z}{N}} \\
&\leq \left(\frac{eN}{z}\right)^z N^{-z}e^{\frac{z}{N}-1} \\
&= e^z z^{-z}e^{\frac{z}{N}}e^{-1} \\
&= 2^{z(\log e - \log z)+\log e\left(\frac{z}{N}-1\right)} \\
&\leq 2^{z(\log e - \log z)} \\
&\leq 2^{-\frac{1}{2}z\log z}
\end{aligned}
$$

We can provide an upper bound on the event of a single bucket having more than w balls after throwing N balls into N buckets using geometric series as follows

$$\Pr[\mathsf{Load}_i^{>w}|\mathsf{Exp}_{\mathsf{throw}}^N] \leq \sum_{z=w+1}^{N} 2^{-\frac{1}{2}z\log z}$$

$$= \sum_{z=1}^{N-w} 2^{-\frac{1}{2}(w+z)\log(w+z)}$$

$$= \sum_{z=1}^{N-w} 2^{-\frac{1}{2}(w\log(w+z)+z\log(w+z))}$$

$$\leq \sum_{z=1}^{N-w} 2^{-\frac{1}{2}(w\log w+z\log z)}$$

$$= 2^{-\frac{1}{2}w\log w} \sum_{z=1}^{N-w} 2^{-\frac{1}{2}z\log z}$$

$$\leq 2^{-\frac{1}{2}w\log w+1}$$

Applying the union bound over all buckets and using Lemma 1 we obtain

$$\Pr[\mathsf{Load}^{>w}|\mathsf{Exp}_{\mathsf{move}}^{N,t}] \leq 2^{-\frac{1}{2}w\log w+1}tN$$

We want to bound this probability of a bad event happening by some correctness parameter c, i.e. we want this probability to be smaller than 2^{-c}.

$$2^{-\frac{1}{2}w\log w+1}tN \leq 2^{-c}$$

$$\Leftrightarrow -\frac{1}{2}w\log w + 1 + \log t + \log N \leq -c$$

$$\Leftrightarrow w\log w \geq 2(c + \log t + \log N + 1)$$

Hence, the bucket size $w \in \mathcal{O}\left(\frac{c+\log t+\log N}{\log(c+\log t+\log N)}\right)$ and therefore the total bandwidth cost in our construction is $\mathcal{O}\left(\sqrt{N}\frac{c+\log t+\log N}{\log(c+\log t+\log N)}\right)$.

4   Main Construction

In this section we are going to present our main Lookahead ORAM construction. The first difference between our Matrix Bucket ORAM and Lookahead ORAM is that we replace all buckets by cells that can only hold single elements. As a first try to construct a more efficient ORAM we could do the following: Initially, randomly shuffle the data, distribute the data elements among matrix cells, and encrypt each cell separately. To access an element, we retrieve the column

corresponding to that element and a row corresponding to a uniformly random cell. After accessing the desired element, we swap the accessed element with the uniformly random cell, re-encrypt both row and column, and write them back into the matrix.

On an intuitive level, one could hope for this to be a secure ORAM construction, since every element will be swapped into a new uniformly random column at every access. Unfortunately this is not the case. The problem is that the distribution of columns into which elements are swapped is not uniformly random when conditioned on the observed row accesses. In particular, the difference with the simple matrix construction is that, here, the accessed element will change the position of another element, i.e the swap partner. It turns out that the server can infer information about the positions of accessed elements whenever we access the same row twice.

Figure 3 illustrates why the straightforward approach of directly swapping the accessed element with an element from a uniformly random cell fails. In this figure, the root node depicts a 2×2 matrix holding four encrypted entries. Initially the server has no knowledge about the arrangement of data elements in the matrix. Let us assume we access two different data elements. With probability non-negligible in the security parameter the following events will occur. On the first access the server observes the second column and second row being accessed. Edges from the root to the first layer show the possible swaps that could have happened, given the observed row and column accesses. $a \leftrightarrow b$ means element a was accessed and swapped with element b. On the second access the server observes the first column and, again, the second row being accessed. The leaf nodes of this tree represent all possible arrangements of elements in the matrix, given the observed access pattern. Dashed boxes indicate the case, where the first accessed element has changed its column. Solid boxes indicate that the first accessed element is in its original column. Counting the leaf nodes it can be seen that the first accessed element will more likely than not have switched columns after the two accesses. Hence, from the server's point of view, the elements are not distributed uniformly at random and the approach does not provide the desired obliviousness guarantees.

4.1 Intuition for Lookahead ORAM

The main issue with this first approach is that the row accesses reveal too much information. Ideally, we would like to have a swap procedure that allows us to directly access the desired element instead of the whole column and then swap that accessed cell with a new cell without revealing the column or row of that new cell. Observe, that to perform a swap, we have to perform two tasks. We have to remove the accessed element from its cell and put it into the cell of its swap partner. Symmetrically, we have to remove the swap partner from its cell and put it into the cell of the accessed element.

To realize such a swap procedure, we introduce two auxiliary stashes stash_{acc} and stash_{swap}, where stash_{acc} will be storing accessed elements and stash stash_{swap} will store pre-selected swap partners. From a high-level perspective, these stashes
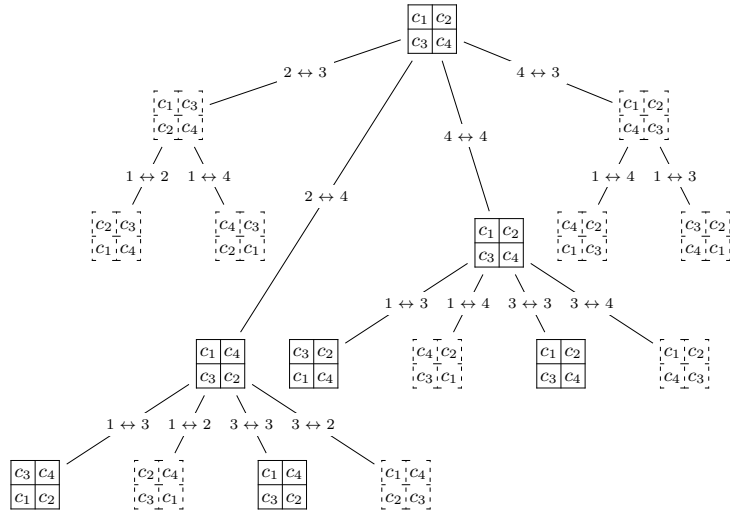
Fig. 3. Illustration of why the naive approach of swapping two random cells via a column and a row access fails. The root node depicts a encrypted 2×2 matrix. The leaf nodes depict all possible arrangements of data elements that are possible after the observed access pattern. Dashed boxes indicate the case, where the first accessed element has changed its column. Solid boxes indicate that the first accessed element is in its original column.

will help us to pretend that we immediately swap accessed elements to an unknown new location in the matrix. From the server's point of view, the client will always read both full stashes, and a uniformly random cell in the matrix, since the client behaves as if accessed elements are immediately magically swapped to their new locations in the matrix. In reality, accessed elements will go to stash_{acc} from where they will be eventually evicted into the cell of their respective swap partner obliviously. Swap partners will be readily waiting in stash_{swap} and upon accessing some element in the matrix, the swap partner will be swapped from stash_{swap} into the accessed cell directly. As an invariant we have that each element is either at its expected location in the matrix or in one of the stashes.

Two issues that need to be addressed are, how do we get swap partners into stash_{swap} before they are used and how to get accessed elements from stash_{acc} into their new cells in the matrix. To solve both these issues, we introduce a (stateful) round-robin column access that will iterate through the columns. Using the round-robin column access, we perform two "maintenance" tasks.

To empty stash_{acc}, we evict all elements from it, whose destination is somewhere in the column of the current round-robin column access. Note that for a matrix of size $\sqrt{N} \times \sqrt{N}$, the round-robin access will have accessed every cell of the matrix in $\sqrt{N}$ steps. This means that no element in stash_{acc} will wait for more than $\sqrt{N}$ steps to be evicted. Since, we add at most one element to

the stash per access, this means that stash_{acc} will never contain more than $\sqrt{N}$ elements.

The second task is to ensure that, whenever we use a swap partner from stash_{swap}, the content of the swap partner's cell must already be available in that stash. Observe, that the swap partner's cell is a uniformly random cell in the matrix, which does not depend on the access pattern and can be selected upfront. Assume stash_{swap} contains $\sqrt{N}$ many preselected swap partners in a queue. These swap partners are sufficient for the next $\sqrt{N}$ accesses. Now upon performing an access, one swap partner from the stash will be used, and we pre-select a uniformly random cell that will be the swap partner once all other swap partners from stash_{swap} are used. At this point in time, the content of the pre-selected swap partner is (likely) not in the stash. However, since we have $\sqrt{N} - 1$ many more accesses before it will be used, we can be certain that the round-robin column access will fetch it in time before it will be used. Since our stashes accommodate accessed elements waiting to be evicted and swap partners waiting to be used, the total stash size is $2\sqrt{N}$.

One detail that we have swept under the rug so far, is the case, where the element we want to access is not in the matrix, but somewhere in the stashes. To get an intuitive feeling for the handling of these cases, it is helpful to keep in mind that at every access, we are basically (virtually) swapping two cells in the matrix. The stashes are just auxiliary data structures that make this process happen. Let (i, j) be the cell in the matrix that is expected to contain the element we are accessing and let the next swap partner from stash_{swap} be some data element v originating from some cell (a, b). If the desired data element is not at position (i, j) in the matrix, but rather in stash_{acc} (waiting to be evicted to cell (i, j)), then take the swap partner from stash_{swap} and place its value v into the matrix at position (i, j). From now on the element in stash_{acc} that was expected to be at (i, j) will be waiting to be evicted to (a, b). If the accessed element, expected at location (i, j), is in stash_{swap}, this means that (i, j) is pre-selected as a swap partner for some future access. Therefore the contents of (i, j) are not supposed to be in the matrix, but rather in the stash_{swap}. In this case, we find (i, j) in the stash_{swap}, put the value of (i, j) into stash_{acc} to be evicted into (a, b), and replace the value of (i, j) in stash_{swap} with v.

4.2 Formal Description

Given this intuition about how our construction work, we are now ready to formally present our construction in Figure 4. Let C be the matrix containing the encrypted data entries and σ be the position map that maps array indices to matrix positions in C. We implement stash_{swap} as a queue and stash_{acc} as an map from positions to values. Init(A) initializes the ORAM data structure, by permuting the elements and storing them in an encrypted matrix C. Initially, both stashes are created empty. stash_{swap}, is filled up with random elements from the encrypted matrix. To write value x or just read at position ℓ in array A inside the ORAM data structure, we use $\mathsf{Access}\,(\ell, \mathsf{x})$, which makes use of the $\mathsf{ReadVirtual}$, $\mathsf{WriteVirtual}$, and $\mathsf{Background}$ subroutines. $\mathsf{ReadVirtual}\,(\mathsf{i}, \mathsf{j})$ reads

Globals

1 C, σ
2 $\mathsf{stash} = (\mathsf{stash}_{swap}, \mathsf{stash}_{acc})$
3 $\mathsf{ind}_{\mathsf{col}}$

Init(A)

1 $C \leftarrow$
    `initEmptyMatrix(`$\sqrt{N} \times \sqrt{N}$`)`
2 $(A', \sigma') \leftarrow$ `shuffle(`A`)`
3 $\sigma \leftarrow$ `fillMatrix(`A', σ', C`)`
4 $(\mathsf{stash}_{swap}, \mathsf{stash}_{acc}) \leftarrow$
    `initStash()`
5 **for** $i = 1 \ldots \sqrt{N}$ **do**
6    $(i, j) \leftarrow$ `pickRandCell()`
7   `enqueue(`$((i, j), \bot), \mathsf{stash}_{swap}$`)`
8 $\mathsf{ind}_{\mathsf{col}} \leftarrow 0$
9 **for** $i = 1 \ldots \sqrt{N}$ **do**
10   `Background();`

Access(ℓ, x)

1 $v.pos \leftarrow \sigma(\ell)$
2 $(v.val, where) \leftarrow$
    `ReadVirtual(`$v.pos$`)`
3 $s \leftarrow$ `dequeue(`stash_{swap}`)`
4 **if** $x \neq \bot$ **then**
5    $v.val = x$
6 `SwapVirtual(`$v, s, where$`)`
7 `swapInPosMap(`$(i, j), s.pos, \sigma$`)`
8 $(i', j') \leftarrow$ `pickRandCell()`
9 `enqueue(`$((i', j'), \bot), \mathsf{stash}_{swap}$`)`
10 `Background()`
11 **return** v

SwapVirtual(v, s, where)

1 **switch** $where$ **do**
2   **case** $matrix$ **do**
3       $C[v.pos] = s.val$
4   **case** $readstash$ **do**
5   `remove(`$v.pos, v.val, \mathsf{stash}_{acc}$`)`
      $C[v.pos] = s.val$
6   **case** $swapstash$ **do**
7       $t \leftarrow$ `find(`$v.pos, \mathsf{stash}_{swap}$`)`
8   `replaceVal(`$t, s.val, \mathsf{stash}_{swap}$`)`

9 `put(`$s.pos, v.val, \mathsf{stash}_{acc}$`)`

Background()

1 $c =$ `readColumn(`$\mathsf{ind}_{\mathsf{col}}, C$`)`
2 **for** $t \in \mathsf{stash}_{acc}$ **do**
3   **if** $t.pos.j = \mathsf{ind}_{\mathsf{col}}$ **then**
4       $c[t.pos.i] == t.value$
5   `remove(`$(i, j), \mathsf{stash}_{acc}$`)`
6 **for** $t \in \mathsf{stash}_{swap}$ **do**
7   **if** $t.pos.j == \mathsf{ind}_{\mathsf{col}}$ **then**
8   **if** $c[t.pos.i] \neq \bot$ **then**
9          $t.value = c[t.pos.i]$
10         $c[t.pos.i] = \bot$
11 `writeColumn(`j, c, C`)`
12 $\mathsf{ind}_{\mathsf{col}} \leftarrow \mathsf{ind}_{\mathsf{col}} + 1 \mod \sqrt{N}$

ReadVirtual((i, j))

1 **if** $C[i, j] \neq \bot$ **then**
2   **return** $(C[i, j], matrix)$
3 $v =$ `lookup(`$(i, j), \mathsf{stash}_{acc}$`)`
4 **if** $v \neq \bot$ **then**
5   **return** $(v, readstash)$
6 $v =$ `lookup(`$(i, j), \mathsf{stash}_{swap}$`)`
7 **return** $(v, swapstash)$

Fig. 4. Pseudocode of Lookahead ORAM

15

the stash and the matrix cell $C[i, j]$ to find the data element that is expected to be at position (i, j) inside the matrix C. SwapVirtual(v, s, where) simulates a swap of accessed value $v.val$ at position $v.pos$ with pre-selected swap partner s with value $s.val$ from position $s.pos$. Background() implements the round-robin column access, which takes care of flushing elements out of stash_{acc} and fetching elements into stash_{swap}.

4.3 Security

Theorem 2. *Let $\mathcal{E} = (\mathsf{gen}, \mathsf{encrypt}, \mathsf{decrypt})$ be an* IND-CPA *secure encryption scheme. Then the construction in Figure 4 is a perfectly secure ORAM scheme with $\mathcal{O}\left(\sqrt{N}\right)$ bandwidth overhead and a total storage cost of $N + \mathcal{O}\left(\sqrt{N}\right)$ data blocks, where N is the number of data elements.*

Proof. Instead of directly arguing about the security of our proposed construction, we will rather argue about the security of an idealized version, which leaks the same amount of information about the access pattern, but is easier to analyze. As previously explained, from a high-level perspective, our construction directly accesses the desired element in the matrix and then swaps it with a random cell. The swap is immediately applied to the position map and we always directly access the cell in the matrix, which should contain a desired element according to the position map. The stash and the round-robin column accesses are there to enable us to (virtually) swap the accessed element into a new cell without leaking anything about that new location.

$\mathsf{IdealizedAccess}_1(\sigma, \ell)$	$\mathsf{IdealizedAccess}_2(\sigma, \ell)$
1 $(i, j) \leftarrow \sigma(\ell)$	1 $(i, j) \leftarrow \sigma(\ell)$
2 $(i', j') \leftarrow \texttt{pickRandCell}(C)$	2 $(i', j') \leftarrow \texttt{pickRandCell}(C)$
3 $v \leftarrow \texttt{readCell}(i, j, C)$	3 $v \leftarrow \texttt{readCell}(i, j, C)$
4 $C \leftarrow \texttt{readMatrix}(C)$	4 $C \leftarrow \texttt{readMatrix}(C)$
5 $\texttt{swapInPosMap}((i, j), (i', j'), \sigma)$	5 $\texttt{swapInPosMap}((i, j), (i', j'), \sigma)$
6 $\texttt{swapInPosMap}((i, j), (i', j'), C)$	6 $(C, \sigma) \leftarrow \texttt{fullReshuffle}(C, \sigma);$
7 **return** v, σ	**return** v, σ

Fig. 5. Idealized access procedures

Since both the stash and the the round-robin column access are always executed independently of the access pattern, they leak no information. Hence, instead of analyzing our construction directly, we can now analyze a construction with the idealized access procedure $\mathsf{IdealizedAccess}_1$ depicted in Figure 5 on the left. The initialization procedure corresponding to $\mathsf{IdealizedAccess}_1$ is a straightforward adaption of our main construction and is not stated explicitly. In $\mathsf{IdealizedAccess}_1$, we directly access the cell that contains our data element and,

next, we retrieve the full matrix to perform the swap operation locally. From an efficiency point of view this is clearly a useless construction, but w.r.t. obliviousness both our main construction and this idealized version thereof leak the same amount of information about the access pattern. More formally, the success probability of any distinguisher D, distinguishing two data access sequences, is the same in our main construction and in the construction with $\mathsf{IdealizedAccess}_1$.

Lemma 2. *Let* oram *be the main construction from Figure 4 and* oram_1^* *the construction using the access procedure* $\mathsf{IdealizedAccess}_1$. *Then, for any distinguisher* D, *for any two data request sequences* $\overrightarrow{y}$ *and* $\overrightarrow{z}$ *we have*

$$|\Pr[\mathsf{D}(\mathsf{oram}(\overrightarrow{y}))] - \Pr[\mathsf{D}(\mathsf{oram}(\overrightarrow{z}))]|$$

$$= |\Pr[\mathsf{D}(\mathsf{oram}_1^*(\overrightarrow{y}))] - \Pr[\mathsf{D}(\mathsf{oram}_1^*(\overrightarrow{z}))]|$$

There are two components that are observable by the server in both our real construction and the idealized access that can leak information about the access pattern. The first component is the swap logic that moves accessed elements to a new position. The second component is the direct accesses to the desired elements, i.e. we need to show that conditioned on previously observed accesses, each new access will fetch a uniformly random cell in the matrix.

Towards showing the first part, let v_k for $1 \leq k \leq N$ be some arbitrary data elements. Let $\mathsf{Exp}_{\mathsf{swap}}^{N,t}$ be the experiment of, initially, distributing the N data elements v_k in a matrix C with N cells uniformly at random and then for t steps repeatedly swapping the contents of two uniformly random cells.

Lemma 3. *Let* C *be a matrix of size* $\sqrt{N} \times \sqrt{N}$ *and let* $v_k \in \mathcal{V}$ *for* $1 \leq k \leq N$ *be arbitrary values from some value space* $\mathcal{V}$. *Then* C, *after running experiment* $\mathsf{Exp}_{\mathsf{swap}}^{N,t}$, *is a uniformly random permutation of the data elements* v_k.

Proof. Initially distributing the data elements v_k uniformly at random in the matrix C corresponds to a uniformly random permutation. For $\mathsf{Exp}_{\mathsf{swap}}^{N,1}$, i.e. distributing the elements and then swapping two uniformly random cells once, the statement holds, since we apply a random permutation of two elements to a uniformly random permutation. The statement for $t > 1$ follows by induction over t.

To conclude the security proof it remains to show that even conditioned on the previously observed accesses the distribution of data elements in C is uniformly random. Assume the server observes an access to position (i, j) in the matrix to fetch some data element v_k. The accessed element is going to be swapped into every position in the matrix with equal probability. Since each element is equally likely, with a probability of $\frac{1}{N}$, to be selected as a swap partner, every element is equally likely to end up in (i, j). From the accessed cells, no other information about the access pattern is leaked. Hence, from the server's perspective all distributions of access patterns are equally likely, no matter what the actual data access sequence is.

For the sake of clarity, let us look at the slightly modified access procedure IdealizedAccess$_2$ depicted in Figure 5 on the right. In IdealizedAccess$_2$ we do not just swap two cells locally, but we fully reshuffle the whole matrix. Due to the full reshuffle, each access is completely independent of the previously observed access pattern. It is straightforward to see that IdealizedAccess$_2$ is a secure ORAM construction. Since in both IdealizedAccess$_1$ and IdealizedAccess$_2$ the access patterns are distributed uniformly at random, IdealizedAccess$_2$ leaks as much information as IdealizedAccess$_1$.

4.4 Online Overhead

For some practical applications it may be of interest to split the total bandwidth overhead into an online and an offline overhead. The point of this is to minimize the online bandwidth, which represents the amount of data that needs to be transmitted, when a client requests an element, and then do some background work, the offline overhead, to ensure the security of the ORAM, when no data requests are actively pending. This way we can minimize the practical latency of user requests despite the inherent lower bound on the overhead shown in [Gol87, GO96].

Looking at our main construction, it is straightforward to split the total bandwidth overhead into online and offline overhead. In the online phase, assuming we download the stash once, we can directly access the desired elements in the matrix on the server without any overhead. In the offline phase we need to do the remaining work, i.e. perform the round-robin column access, fill up the stash of pre-selected swap partners and flush out elements that need to go back into the matrix.

While storing the stash locally can theoretically compromise data integrity in case of a client device failure, the only part of the stash that cannot be randomly reinitialized is the cache of recently accessed elements. If the server is is significantly more reliable than the client device, recently accessed elements can be written to a separate server-side buffer of size $\sqrt{N}$ in a round-robin manner. Each element will be stored in this buffer for $\sqrt{N}$ operations which is enough for the background operations to write it to its long-term server-side storage position.

It is also possible to allow multiple online accesses in a single burst. The simplest way to do it increases local storage requirements by twice the size of the maximum allowed burst length. Note that our implementation faithfully simulates the oram$_1^*$ construction as long as the background work is performed at least $\sqrt{N}$ times between committing to a swap partner and its use, and background work is also performed at least $\sqrt{N}$ times between accessing the element and evicting it from the recently accessed element stash. If there is additional stash space of size b for the potential swap partners and the same amount of space for the recently accessed elements, we only need to have at least $\sqrt{N}$ background work operations performed during every interval when $b + \sqrt{N}$ access operations are performed. Note that it is possible to perform some part of the background work, handle an additional burst and then continue

18

the background work as long as enough background work is done to prevent exhaustion of the swap partners or overflow of the recently accessed elements stash.

The same analysis shows that there is a trade off between stash size and bandwidth overhead: for example, if we have a stash twice as large as needed, we can afford doing only half the background processing step after each access.

4.5 Trading Off Bandwidth and Storage Overhead

For our main construction, the server's storage overhead comes from the two auxiliary stashes that it needs to store in addition to the encrypted data elements. For a matrix C, where the number of columns equals the number of rows, this results in an additive storage overhead of $2\sqrt{N}$. More generally, by considering an arbitrary rectangle C with H rows and W columns, we can trade off the concrete storage and bandwidth overhead costs of our construction. The number of columns W affects the time it takes the round-robin column access to iterate over the whole rectangle C and thus it also affects the stashes which have to be of size W each. The number of rows H, affects the size of the column that we need to download at each access. For example, by setting $H = 2\sqrt{N}$ and $W = \frac{1}{2}\sqrt{N}$, we can, in comparison to a quadratically shaped matrix C, directly reduce the additive storage overhead to $\sqrt{N}$ and maintain a bandwidth overhead of $3\sqrt{N} + 1$. By setting $H = \sqrt{2N}$ and $W = \sqrt{\frac{N}{2}}$, we get a bandwidth overhead of $2\sqrt{2}\sqrt{N} + 1$ and a storage overhead of $\sqrt{2N}$.

5 Constant Client-Side Storage

So far we have assumed that the full position map is stored explicitly on the client side. Following the approach of [SCSL11], we show that our construction can be modified to only require $\mathcal{O}(1)$ client-side storage at the cost of $\mathcal{O}(\log N)$ rounds of interaction per access between the client and the server. To store the position map on the server side, we will create a sequence of ORAM data structures that increase in size, which represent the position map. This means that now the server stores a sequence of position map ORAMs and an ORAM data structure that contains the actual data. To access an element at a certain index, the client will use the position map ORAMs to determine, which index it should query in the ORAM data structure that stores the actual data.

Theorem 3. *Consider an arbitrary ORAM construction with a client-side position map. Assume there exist such constants K and C that for N entries with a block size of $D \geqslant 4 \log N$, the ORAM construction has a multiplicative bandwidth overhead of $C\sqrt{N}$ and total storage of $(N + K\sqrt{N})D$ bits.*

Such an ORAM can be converted into an ORAM with a server-side position map with a multiplicative bandwidth overhead of $2C\sqrt{N} + 4$ and total storage $ND + 2KD\sqrt{N} + 2N \log N$ bits.

Proof. The proof goes by induction over N. In the base case, for size $N \leqslant 4$, we use a linear ORAM, which simply reads and writes all the blocks for each access. This ORAM has a bandwidth overhead of 4 and no storage overhead.

For an arbitrary N, we use the induction hypothesis to replace an ORAM with client-side position map with an ORAM that stores the position map on the server-side. We encode 4 blocks addresses in the position map into one storage block inside the ORAM data structure. Using a block size of $4 \log N$ results in an ORAM data structure that stores $\frac{N}{4}$ blocks, storing 4 addresses each. Now we apply the induction hypothesis.

Since we want to store $\frac{N}{4}$ blocks with $4 \log N$ bits per block, by induction hypothesis our position map ORAM has a total storage cost of $N \log N + 2K(4 \log N)\sqrt{\frac{N}{4}} + 2\frac{N}{4} \log N$ bits. Adding the storage costs of our main ORAM, the total storage cost in bits will be

$$N \log N + 2K(4 \log N)\sqrt{\frac{N}{4}} + 2\frac{N}{4} \log N + ND + KD\sqrt{N}$$

$$= ND + K(4 \log N)\sqrt{N} + KD\sqrt{N} + \frac{3}{2}N \log N$$

$$< ND + 2KD\sqrt{N} + 2N \log N$$

In our construction, every access to the main ORAM requires one access to the position map. Note that we do not need any additional position map accesses for the background work that moves data out of the stash. The position map access adds its bandwidth cost of $2C\sqrt{\frac{N}{4}} + 4$ to the main ORAM bandwidth cost of $C\sqrt{N}$. The total overhead is $C\sqrt{N} + 2C\sqrt{\frac{N}{4}} + 4 = 2C\sqrt{N} + 4$.

This completes the proof of the inductive step.

By setting the parameters as described in Section 4.5, we get

Corollary 1. *Lookahead ORAM with a client-side position map, block size $D = 4 \log N$, multiplicative bandwidth overhead $\mathcal{O}\left(\sqrt{N}\right)$, and total storage of $(N + \sqrt{2N})4 \log N$ bits can be converted into an ORAM with server-side position map and multiplicative bandwidth overhead $\mathcal{O}\left(\sqrt{N}\right)$ and total storage of $\log N(6N + 12\sqrt{N})$ bits.*

5.1 Constant Client-Side Temporary Storage

Using a recursive position map, we reduce the persistent client-side storage to $\mathcal{O}(1)$. However, during any operation on the ORAM data structure, the client still needs to temporarily store $\mathcal{O}\left(\sqrt{N}\right)$ data blocks. If desired, this can be reduced to $\mathcal{O}(1)$ data blocks at the cost of increasing the bandwidth overhead

by a multiplicative factor of $\mathcal{O}\left(\log\sqrt{N}\right)$. We outline our solution here and leave the details to the interested reader.

Recall that during each ORAM operation, we have to access one matrix cell, one matrix column j, and the two stashes. For the sake of simplicity lets assume that the data matrix has height and width $\sqrt{N}$. Upon each ORAM operation we first process the stashes stash_{acc} and stash_{swap} into a temporary stashes $\mathsf{stash}_{final}^{acc}$ and $\mathsf{stash}_{final}^{swap}$, which have the following property: If either stash_{swap} or stash_{acc} contains a cell (i, j') with $j' = j$, then this cell will be in the i-th position of $\mathsf{stash}_{final}^{swap}$ or $\mathsf{stash}_{final}^{acc}$ respectively. For example, if stash_{swap} contains a cell that is associated with the second cell from the top in the current column j, then it will be in the second position in $\mathsf{stash}_{final}^{swap}$. Once we have such stashes, we can iterate over the current column j and both final stashes one cell each at a time and perform the necessary operations. We explain the generation of $\mathsf{stash}_{final}^{acc}$. The generation of $\mathsf{stash}_{final}^{swap}$ is completely analogous:

1. Create an empty stash $\mathsf{stash}_{final}^{acc}$
2. Iterate over stash_{acc} one cell at a time. If the current cell (i, j) is associated with the current column, we append it to $\mathsf{stash}_{final}^{acc}$ with a priority i. If it is not relevant for the current column, then we append it to $\mathsf{stash}_{final}^{acc}$ with a priority of ∞.
3. We append $\sqrt{N}$ dummy elements to $\mathsf{stash}_{final}^{acc}$, where dummy element i has priority $i + \frac{1}{2}$. At this point $\mathsf{stash}_{final}^{acc}$ contains $2\sqrt{N}$ cells.
4. We use an oblivious sorting algorithm [Bat68, AKS83] to sort $\mathsf{stash}_{final}^{acc}$ according to its priorities from smallest to largest.
5. We iterate over $\mathsf{stash}_{final}^{acc}$ one data element at a time. Whenever we read a real cell i, we set the priority of the dummy cell right after it to ∞.
6. We again obliviously sort $\mathsf{stash}_{final}^{acc}$.

A visual illustration of this final stash generation is depicted in Figure 6.

The post-processing of the two final stashes is straightforward. Again we only explain the post-processing of $\mathsf{stash}_{final}^{acc}$, since the post-processing of $\mathsf{stash}_{final}^{swap}$ is completely analogous. We iterate over $\mathsf{stash}_{final}^{acc}$ and assign real each data cell


Fig. 6. Illustration of the generation of $\mathsf{stash}_{final}^{acc}$ with $\sqrt{N} = 4$. Each row represents $\mathsf{stash}_{final}^{acc}$ at a certain stage in the stash generation algorithm. The squares represent data cells and their labels represent their assigned priorities. Shaded rectangles represent dummy elements. (a) depicts $\mathsf{stash}_{final}^{acc}$ at the end of step 3. (b) depicts the stash at the end of step 4. (c) depicts it after step 5.
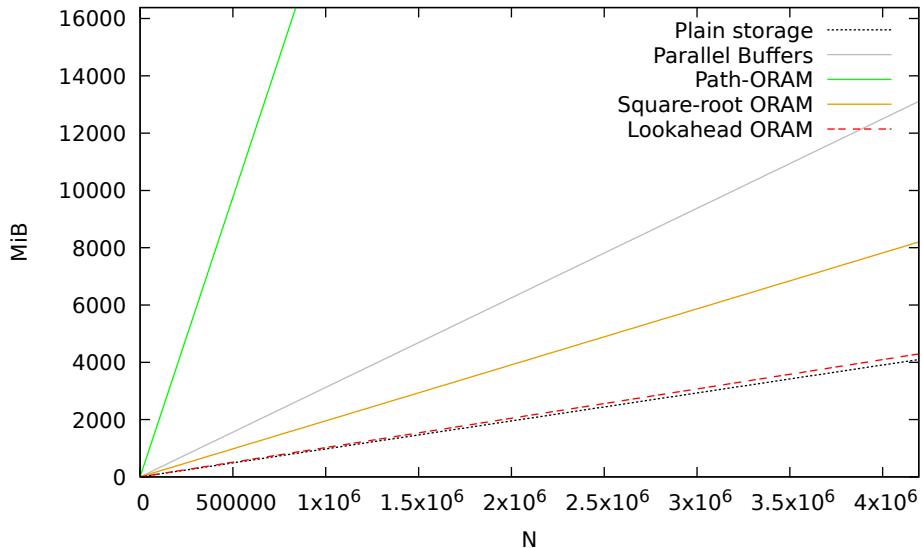
Fig. 7. Comparison of the storage overheads of different ORAM constructions. The x-axis shows different amounts of data blocks. The data block size is fixed to 1024 bytes. The y-axis plots the total required storage on the server-side in MiB. For Lookahead ORAM, the values are taken from empirical measurements of our implementation. For the other schemes, the values are computed based on the concrete formulas and constants that are reported in the respective papers.

priority $-\infty$ and each dummy element priority ∞. We use oblivious sort to sort $\mathsf{stash}^{acc}_{final}$ from small to large. We interpret the first $\sqrt{N}$ as stash_{acc} and delete the remaining last $\sqrt{N}$ (dummy) elements.

Regarding the efficiency of our procedure we observe that we perform a constant amount of oblivious sorts of $\sqrt{N}$ elements per ORAM operation. This can be done with bandwidth overhead of $\mathcal{O}\left(\sqrt{N}\log\sqrt{N}\right)$ data elements with $\mathcal{O}(1)$ temporary storage on the client side [AKS83].

6 Evaluation

To provide a rough idea of the practical performance of our Lookahead ORAM construction, we implemented a prototype with a client-side position map and evaluated it in terms of concrete bandwidth and server-side storage overhead.

We assume that each encrypted block has an extra 40-byte encryption/MAC overhead, and every stash entry has an additional status 20-byte header with status and position information. We assume 4-byte words are also used for denoting the request types.
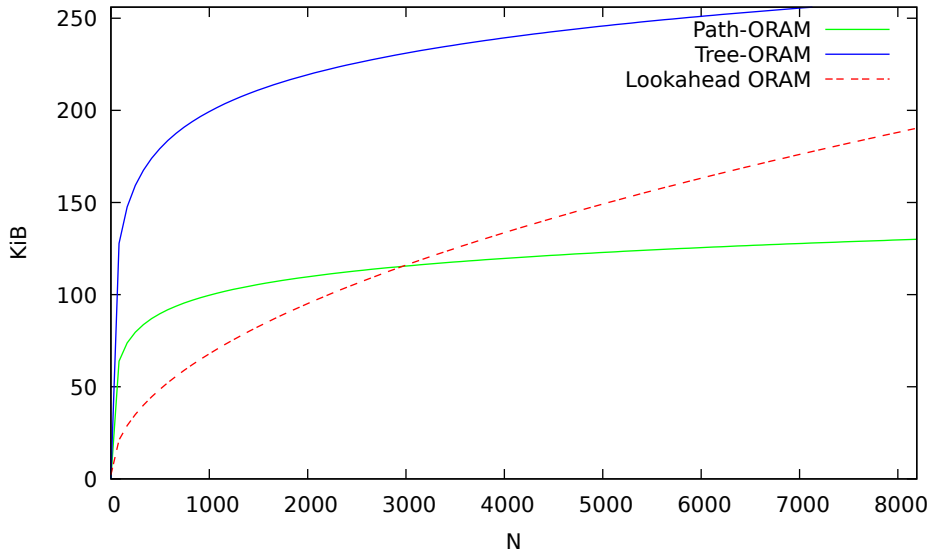
Fig. 8. Comparison of the bandwidth overhead of Lookahead ORAM and Path-ORAM. The x-axis shows different amounts of data blocks. The data block size is fixed to 1024 bytes. The y-axis shows the total amount of transmitted data per access in MiB.

When encrypting N data blocks of size B each, the server's total storage is $N \times (B + 40)$. The corresponding position map is $8N$, and the stash $2\sqrt{N} \times (B + 20)$ bytes large. Whenever the client performs an operation on the ORAM data structure, it needs to download $40 + (B + 40) \times (\sqrt{N} + 1)$ and upload $80 + (B + 40) \times (\sqrt{N} + 1)$ bytes.

During initialization of our ORAM, we fill the storage with zeros and then fill the stash with pre-selected swap partners. During this initalization, we upload $44 + (B + 40) \times N$ bytes to the server.

All of the above formulas have been calculated based on our theoretical construction and validated empirically using our prototype implementation. In the following, the data block size is fixed to 1024 bytes.

Storage overhead We compare our storage overhead on the server-side to the overheads of the most storage efficient related works with sublinear worst-case bandwidth overheads. For our comparison, we measure the total storage requirements on the server-side for varying N. The data points for Lookahead ORAM were obtained through experiments. For the schemes we compare ourselves to, we computed the data points based on the formulas (including constants) given in the respective papers. It should be noted that in [SvS^{+}13] the authors obtain provable security for Path-ORAM with a storage overhead of $20N$, but evaluate their scheme on smaller parameter settings. Since we are interested in provable

security and correctness guarantees, we compare our scheme to theirs with a storage overhead of $20N$. The results are depicted in Figure 7. As expected, assuming a client-side position map, Lookahead ORAM has the smallest storage overhead among all for the tested setting. Storing the position map on the server-side would increase the storage overhead of Lookahead ORAM by a factor of roughly 6.

Bandwidth overhead We compare ourselves to Path-ORAM (with a client-side position map), which is known to be the most efficient construction in terms of asymptotic and practical worst-case bandwidth overhead. For our comparison we use their self-reported bandwidth overhead of $10 \log N$. The results of our comparison are depicted in Figure 8. As expected from the asymptotic behaviour of Path-ORAM and Lookahead ORAM, we can see that Path-ORAM becomes more efficient for large values of N. However, for values of $N < 3000$, Lookahead ORAM is more efficient in terms of concrete bandwidth overhead.

References

AKL+18. Gilad Asharov, Ilan Komargodski, Wei-Kai Lin, Kartik Nayak, Enoch Peserico, and Elaine Shi. Optorama: Optimal oblivious ram. Cryptology ePrint Archive, Report 2018/892, 2018. https://eprint.iacr.org/2018/892.

AKS83. M. Ajtai, J. Komlós, and E. Szemerédi. An 0(n log n) sorting network. In *Proceedings of the Fifteenth Annual ACM Symposium on Theory of Computing*, STOC '83, pages 1–9, New York, NY, USA, 1983. ACM.

AKST14. Daniel Apon, Jonathan Katz, Elaine Shi, and Aishwarya Thiruvengadam. Verifiable oblivious storage. pages 131–148, 2014.

Bat68. K. E. Batcher. Sorting networks and their applications. In *Proceedings of the April 30–May 2, 1968, Spring Joint Computer Conference*, AFIPS '68 (Spring), pages 307–314, New York, NY, USA, 1968. ACM.

BLS12. Daniel J. Bernstein, Tanja Lange, and Peter Schwabe. The security impact of a new cryptographic library. pages 159–176, 2012.

BMP11. Dan Boneh, David Mazieres, and Raluca Ada Popa. Remote oblivious storage: Making oblivious ram practical. 2011.

BN16. Elette Boyle and Moni Naor. Is there an oblivious RAM lower bound? pages 357–368, 2016.

CKN+18. T.-H. Hubert Chan, Jonathan Katz, Kartik Nayak, Antigoni Polychroniadou, and Elaine Shi. More is less: Perfectly secure oblivious algorithms in the multi-server setting. pages 158–188, 2018.

CNS18. T.-H. Hubert Chan, Kartik Nayak, and Elaine Shi. Perfectly secure oblivious parallel RAM. pages 636–668, 2018.

DDF+16. Srinivas Devadas, Marten van Dijk, Christopher W. Fletcher, Ling Ren, Elaine Shi, and Daniel Wichs. Onion ORAM: A constant bandwidth blowup oblivious RAM. pages 145–174, 2016.

DMN11. Ivan Damgård, Sigurd Meldgaard, and Jesper Buus Nielsen. Perfectly secure oblivious RAM without random oracles. pages 144–163, 2011.

DPP18.   Ioannis Demertzis, Dimitrios Papadopoulos, and Charalampos Papamanthou. Searchable encryption with optimal locality: Achieving sublogarithmic read efficiency. pages 371–406, 2018.

DSS14.   Jonathan Dautrich, Emil Stefanov, and Elaine Shi. Burst oram: Minimizing oram response times for bursty access patterns. In *23rd USENIX Security Symposium (USENIX Security 14)*, pages 749–764, 2014.

GMOT11.   Michael T Goodrich, Michael Mitzenmacher, Olga Ohrimenko, and Roberto Tamassia. Oblivious ram simulation with efficient worst-case access overhead. In *Proceedings of the 3rd ACM workshop on Cloud computing security workshop*, pages 95–100. ACM, 2011.

GMSS16.   Steven Gordon, Atsuko Miyaji, Chunhua Su, and Karin Sumongkayothin. M-ORAM: A matrix ORAM with log N bandwidth cost. pages 3–15, 2016.

GO96.   Oded Goldreich and Rafail Ostrovsky. Software protection and simulation on oblivious rams. *Journal of the ACM (JACM)*, 43(3):431–473, 1996.

Gol87.   Oded Goldreich. Towards a theory of software protection and simulation by oblivious RAMs. pages 182–194, 1987.

IKK12.   Mohammad Saiful Islam, Mehmet Kuzu, and Murat Kantarcioglu. Access pattern disclosure on searchable encryption: Ramification, attack and mitigation. 2012.

LN18.   Kasper Green Larsen and Jesper Buus Nielsen. Yes, there is an oblivious RAM lower bound! pages 523–542, 2018.

MBC14.   Travis Mayberry, Erik-Oliver Blass, and Agnes Hui Chan. Efficient private file retrieval by combining ORAM and PIR. 2014.

OS97.   Rafail Ostrovsky and Victor Shoup. Private information storage (extended abstract). pages 294–303, 1997.

PPRY18.   Sarvar Patel, Giuseppe Persiano, Mariana Raykova, and Kevin Yeo. PanORAMa: Oblivious RAM with logarithmic overhead. pages 871–882, 2018.

RFK+15.   Ling Ren, Christopher Fletcher, Albert Kwon, Emil Stefanov, Elaine Shi, Marten van Dijk, and Srinivas Devadas. Constants count: Practical improvements to oblivious ram. In *24th USENIX Security Symposium (USENIX Security 15)*, pages 415–430, Washington, D.C., 2015. USENIX Association.

SCSL11.   Elaine Shi, T.-H. Hubert Chan, Emil Stefanov, and Mingfei Li. Oblivious RAM with $O((\log N)^3)$ worst-case cost. pages 197–214, 2011.

SSS12.   Emil Stefanov, Elaine Shi, and Dawn Xiaodong Song. Towards practical oblivious RAM. 2012.

SvS+13.   Emil Stefanov, Marten van Dijk, Elaine Shi, Christopher W. Fletcher, Ling Ren, Xiangyao Yu, and Srinivas Devadas. Path ORAM: an extremely simple oblivious RAM protocol. pages 299–310, 2013.

A   Attack on [GMSS16]

In a work by Gordon et al. [GMSS16] the authors present an ORAM, called M-ORAM. Their construction has a flaw and does not provide obliviousness. In the following, we give a high-level overview of their scheme and sketch our attack that breaks their obliviousness claims.

The construction partitions the server-side storage into a fixed number of rows and a number of columns that depends on the dataset's size. Every cell in

their rectangular storage layout holds one data element. Additionally, every row has its own separate constant-sized stash.

Initially, all data elements are present in the storage rectangle in a randomly permuted order and the stashes are empty. Simply speaking an access is performed by accessing one element in each row of their data structure. In one of the rows the desired element is accessed and in all other rows a uniformly random cell is selected. More precisely, the authors claim that to achieve obliviousness not all "dummy" cells are selected uniformly at random, instead some of them are random cells from the previous access. After retrieving one cell from each row, the client shuffles the cells and puts one cell into each stash. The client picks one random block from each stash and sends it back to the server as the new content of the retrieved cells.

Let $x_1, \ldots, x_N$ be some data elements stored in the ORAM data structure, the access sequences $(\mathsf{read}(x_1), \mathsf{read}(x_2), \mathsf{read}(x_1))$ and $(\mathsf{read}(x_1), \mathsf{read}(x_2), \mathsf{read}(x_3))$ can be distinguished with a success probability that is non-negligible in the security parameter. From a high-level perspective, every access selects a subset of cells from the data structure and every two subsets corresponding to two consecutive accesses intersect at some random cells. For three accesses the proposed approach breaks down. Looking at our first access sequence, the proposed construction has a slightly higher bias of the first and third access subset intersecting, since we are accessing the same element.

We have contacted the authors and they have acknowledge our attack.