

PanORAMa: Oblivious RAM with Logarithmic Overhead

Sarvar Patel^{*1}, Giuseppe Persiano^{†1,2}, Mariana Raykova^{‡1,3}, and Kevin Yeo^{§1}

¹Google LLC

²Università di Salerno

³Yale University

Abstract

We present PanORAMa, the first Oblivious RAM construction that achieves communication overhead $O(\log N \cdot \log \log N)$ for database of N blocks and for any block size $B = \Omega(\log N)$ while requiring client memory of only a constant number of memory blocks. Our scheme can be instantiated in the "balls and bins" model in which Goldreich and Ostrovsky [JACM 96] showed an $\Omega(\log N)$ lower bound for ORAM communication.

Our construction follows the hierarchical approach to ORAM design and relies on two main building blocks of independent interest: a *new oblivious hash table construction* with improved amortized $O(\log N + \text{poly}(\log \log \lambda))$ communication overhead for security parameter λ and $N = \text{poly}(\lambda)$, assuming its input is randomly shuffled; and a complementary *new oblivious random multi-array shuffle construction*, which shuffles N blocks of data with communication $O(N \log \log \lambda + \frac{N \log N}{\log \lambda})$ when the input has a certain level of entropy. We combine these two primitives to improve the shuffle time in our hierarchical ORAM construction by avoiding heavy oblivious shuffles and leveraging entropy remaining in the merged levels from previous shuffles. As a result, the amortized shuffle cost is asymptotically the same as the lookup complexity in our construction.

*sarvar@google.com

†giuper@gmail.com

‡mariana.raykova@yale.edu

§kwlyeo@google.com

1 Introduction

The cryptographic primitive of *Oblivious RAM* (ORAM) considers the question of how to enable a client to outsource its database to an untrusted server and to query it subsequently without any privacy leakage related to the queries and the database. While encryption can help with hiding the outsourced content, a much more challenging question is how to hide the leakage from the access patterns induced by the queries' execution. This is leakage that is not only ruled out by the strong formal definition of privacy preserving data outsourcing but has also been proven to be detrimental in many practical outsourcing settings [27, 6]. Hiding access patterns is only interesting when coupled with efficiency guarantees for the query execution in the following sense: there is a trivial access hiding solution which requires linear scan of the whole data at every access. Once such efficiency properties are in place, ORAM constructions also have another extremely important application as they are a critical component for secure computation solutions that achieve sublinear complexity in their input size [32, 26, 40].

Thus, the study of ORAM constructions has been driven by the goal of improving their bandwidth overhead per access while providing hiding properties for the access patterns. This study has been a main research area in Cryptography for the past thirty years, since the notion was introduced by Goldreich [20], and it has turned ORAM into one of the classical cryptographic concepts. The seminal works of Goldreich and Ostrovsky [20, 31, 21] introduced the first ORAM constructions achieving square root and polylogarithmic amortized query efficiency. These early works also considered the question of a lower bound on the amortized communication complexity required to maintain obliviousness. They presented a lower bound result of $O(\log_C N)$ blocks of bandwidth overhead for any ORAM construction for a database of size N blocks and a client with memory that can store C blocks. However, this result came with a few caveats, which were clarified and more carefully analyzed in the recent work by Boyle and Naor [5].

ORAM Communication Lower Bound and Its Restrictions. The main restriction for the lower bound result of Goldreich and Ostrovsky [21] is the computation model assumed for the server. This work assumes that the server operations are restricted to fetching blocks from memory, writing new blocks into memory, and moving blocks between different memory positions. At first sight, this could be all one expects from an outsourcing server to do, however, this rules out the server performing computation on the outsourced data, which can be easily enabled using homomorphic encryption [17, 33, 18]. The model for which the lower bound holds, was later formalized as the “balls and bins” storage model in the work of Boyle and Naor [5]. Moreover, Boyle and Naor gave much more insight into understanding the lower bound of Goldreich and Ostrovsky. They provided evidence that extending the lower bound beyond the restricted model in the original result will be a challenging task by showing a reduction from sorting circuits to ORAM, which essentially means that extending the ORAM lower bound will imply new lower bounds for sorting circuits.

The second caveat for the applicability of the lower bound is that it holds only for constructions that provide statistical security. The rough reason is that the proof is based on a counting argument and uses the fact that, in a statistically secure ORAM, the observed access pattern is independent of the actual addresses of the query sequence and depends only on the randomness used in the algorithms. This property does not necessarily hold if the security is computational.

Technically, all known ORAM constructions use encryption to hide the content of the outsourced data, which makes their security guarantees only computational. However, the bins and balls

model abstracts away the concern that encryption only computationally hides the underlying blocks by treating the blocks as indistinguishable balls. In fact, the original lower bound result also allows the client to have oracle access to a private random function, which can be used in place of pseudorandom functions.

Finally, the lower bound result of $\Omega(\log_C N)$ applies to algorithms that work for any block size. And while there are results [38, 40] that meet, and even beat, the lower bound, this holds only in cases when they are instantiated with appropriately big block sizes. In particular, there is no known result that meets the lower bound for every block size $\Omega(\log N)$.

ORAM Constructions and Their Efficiency. While the lower bound of Goldreich and Ostrovsky has its caveats, it has become the measure to compare with for every new construction with improved complexity. We next overview the known ORAM constructions, their models and efficiency with respect to the requirements for the lower bound. Existing ORAM schemes could be roughly divided into two categories: constructions [35, 1, 13, 24, 28] that follow the hierarchical blueprint introduced in the work of Goldreich and Ostrovsky [31, 21], and constructions [37, 19, 38, 12, 40, 36, 7] that follow the tree-based template with a recursive position map, which was introduced in the work of Shi et al. [37].

The idea underlying the first class of constructions is to divide the data in levels of increasing size that form a hierarchy and to instantiate each level with an oblivious access structure that allows each item to be accessed obliviously only once in that level. The smallest level in the hierarchy is linearly scanned at each access and each block that is accessed is subsequently moved to this level. To prevent overflowing of the top levels, there is a deterministic schedule, independent of the actual accesses, that prescribes how blocks move from the smaller to the bigger levels.

Within the general hierarchical framework, the main optimization question considered in different constructions is how to instantiate the oblivious structure in each level. The original Goldreich-Ostrovsky construction [21] used pseudorandom functions (PRFs) resulting in $O(\log^3 N)$ amortized communication overhead. Later, the work of Pinkas and Reinman [35] proposed the use of Cuckoo hash tables. This work suffered from a subtle issue related to the obliviousness of the Cuckoo hash tables, which was later fixed in the work of Goodrich and Mitzenmacher [24] who showed an elegant algorithm to obliviously construct a Cuckoo hash table. Their main argument was a reduction of the Cuckoo hash table construction to oblivious sorting that resulted in a $O(\log^2 N)$ -overhead ORAM. Subsequently, Kushilevitz et al. [28] devised a balancing scheme that further improved the bandwidth to $O(\log^2 N / \log \log N)$. The work of Chan et al. [7] presented a unified framework for hierarchical ORAM constructions and made explicit the notion of an oblivious hash table in order to capture the properties of the oblivious structure needed for each level. All the above constructions require that the client’s memory can hold only a constant number of blocks.

The known hierarchical ORAM constructions do not make any assumptions about the block sizes used to store data in memory and can be instantiated with any block size $B = \Omega(\log N)$. Even the construction with best asymptotic efficiency among existing schemes, when instantiated with a private random function in the balls and bins model, does not meet the logarithmic lower bound in the case of client’s memory that holds only a constant number of blocks.

The other main construction template for ORAM schemes leverages the idea of mapping blocks to random position map (PMAP) indices and then arranging the data in a binary tree with leaves indexed according to the position map, where each block can reside only in a node on the path to its corresponding PMAP leaf. Thus, in order to access a block it is sufficient to read the path indexed

by its PMAP value. In order to access efficiently the PMAP for each query, the construction stores recursively the position map by partitioning it into blocks each of which contains at least two PMAP indices. After the recursion the construction consists of a logarithmic number of trees of decreasing size. Every time a block is accessed, it is assigned a new PMAP index and is moved to the root of the tree. In order to prevent overflowing of nodes, the constructions periodically evict blocks down their corresponding tree paths. The main difference between different tree-based ORAMs is related to the concrete eviction algorithms they use, which have been evolving and improving the ORAM access overhead. The work of Shi et al. [37] that pioneered the tree-based approach achieved $O(\log^3 N)$ communication. Gentry *et al.* [19] improved the overhead to $O(\frac{\log^3 N}{\log \log N})$. Currently the most efficient tree-based construction is the Path ORAM construction [38], which achieves $O(\log^2 N)$ bandwidth overhead for general block sizes. If instantiated with blocks of size $\Omega(\log^2 N)$, Path ORAM has bandwidth overhead of $O(\log N)$ blocks. The same efficiency holds also for Circuit ORAM [40], which optimizes the size of the circuits used to implement the ORAM access function in the context of secure computation.

The only computational assumption of the above tree-based constructions is related to the encryption used to hide the content and thus they do offer statistical guarantees in the balls and bins model. Several works [14, 29] also demonstrate how to bypass the lower bound of communication complexity, if the server is allowed to do computation on the data it stores, which is enabled by homomorphic encryption.

Our Contributions. In this paper we present PanORAMa, a computationally secure oblivious RAM with $O(\log N \cdot \log \log N)$ bandwidth overhead and constant client memory¹. Our construction works for any block size $B = \Omega(\log N)$. We note that this assumption on the block size is very natural since all known ORAM constructions including ours require that the blocks store their own addresses for correctness and this already takes $\Theta(\log N)$ bits. In addition, PanORAMa is in the balls and bins model of Boyle and Naor [5] as it treats each data block as an atomic piece of data and the server only fetches blocks from memory, writes blocks to memory and moves blocks between different memory positions. Thus, PanORAMa achieves currently the best asymptotic communication overhead among constructions that work with general block sizes, operate in the balls and bins model and require constant number of blocks client memory.

Our construction can be modified in a straightforward way to obtain statistical security in the balls and bins model if the client is provided with access to a private random function, which matches the assumptions in the original lower bound (see Theorem 6 in [21]). In this case, we obtain a construction that is only $O(\log \log N)$ away from the lower bound overhead proven by Goldreich and Ostrovsky [21].

The PanORAMa construction relies on two main building blocks, which can have applications outside ORAM of independent interest: a *new oblivious hash table construction* and a *new oblivious random multi-array shuffle algorithm*.

- *Oblivious Hash Table (OHT)*. Our oblivious hash table construction offers an amortized access efficiency of $O(\log N + \text{poly}(\log \log \lambda))$ blocks assuming the starting data is randomly shuffled, where λ is the security parameter.
- *Oblivious Random Multi-Array Shuffle*. Complementary to the OHT primitive is our efficient oblivious random multi-array shuffle algorithm that shuffles together data which initial order

¹We measure bandwidth and client memory using the size B of a block as a unit.

has partial entropy with respect to the adversary. More precisely, our algorithm shuffles together A_1, \dots, A_L arrays of total size N , each of which is independently and randomly shuffled. For convenience suppose that the L arrays are arranged in decreasing size. Then, our shuffle requires $O(N \log \log \lambda + \frac{N \log N}{\log \lambda})$ blocks of communication when there exists $\text{cutoff} = O(\log \log \lambda)$ such that $|A_{\text{cutoff}}| + \dots + |A_L| = O(\frac{N \log \log \lambda}{\log N})$.

Technical Overview of Our Result. Our ORAM construction follows the general paradigm of hierarchical ORAM constructions as laid out by Goldreich and Ostrovsky [21] (see Chan et al. [7] for a formalized presentation of the framework). As we discussed above, the hierarchical constructions distribute the data in several levels, which are instantiated with oblivious structures that provide access obliviousness for non-repeating queries. When querying an item, the smallest level is linearly scanned and then each other level’s corresponding oblivious hash table is queried either for the searched item, if it has not been found yet, or for a dummy value, otherwise. Every 2^i accesses, all levels of size less or equal than 2^i are merged and shuffled together and placed in an oblivious data structure in the level of capacity 2^{i+1} .

At a high level of description, the dominant cost in the communication complexity of existing hierarchical ORAM schemes comes from the step that obviously merges the content of multiple levels and initializes a new level. This step usually involves several oblivious sorts that mix all items together. The best known data-oblivious sorting algorithms [39, 22, 23, 8] that are used in these constructions have asymptotic complexity $O(N \log N)$. Such oblivious sorting constructions achieve very strong hiding guarantees even against an adversary that knows the initial order of the data. However, in the case of the shuffling steps in ORAM constructions, a lot of the inputs at each level have not been queried yet and remain mixed together in the eyes of the server from the previous shuffling step. Thus, the input data has more entropy than an input that is completely known to the server, which is the assumed input to oblivious sorting algorithms.

The main insight for our construction is to devise a way to leverage the entropy left in the items in each level at the start of each shuffle in order to obtain a more efficient shuffling algorithm for the content of the ORAM levels that are being merged. In order to achieve this, we construct a new shuffling algorithm that works on inputs that are partitioned into arrays which are independently shuffled. However, this on its own is not enough. We also need an oblivious hash table to store items in each level that both offers efficient query complexity, efficient oblivious initialization as well as a way to extract the non-queried items from the OHT in order to use them as input for the multi-array shuffle.

Oblivious Hash Table. Our oblivious hash table construction is inspired by the two-tier hash scheme proposed by Chan et al. [9], however, with some significant changes that enable constructing the OHT without using an oblivious sort on all the data blocks. The idea of Chan et al. [9] is to allocate the database items into bins on the first level of the hash table using a PRF, where the size of the bins is set to be $O(\log^\delta \lambda)$, for some constant $0 \leq \delta < 1$, which does not guarantee non-negligible overflow probability. All overflow items are allocated to a second level where they are distributed using a second PRF. In order to initialize this two-tier hash scheme, the authors use an oblivious sort which comes at a cost $O(N \log N)$ for N blocks of data.

Our goal is to obtain a construction of an oblivious hash table that allows more efficient oblivious initialization assuming that the database items are already randomly shuffled. The assumption of the randomly shuffled input is not arbitrary. We will use our oblivious random multi-array shuffle

to construct this random shuffle of the input in the context of our ORAM construction.

We will assign items into buckets according to a PRF non-obliviously taking linear time in the size of the data. Then, we sample from a binomial distribution loads for all bins that correspond to randomly distributing only ϵ fraction of the total number of items. We choose a cutoff point, thrsh , such that with overwhelming probability, it is larger than any bin load sampled from the binomial distribution in the second step and, at the same time, is smaller than any load from the distribution of items induced by the PRF in the first step. We cut the size of each bin to exactly thrsh items, among which there will be as many real items as the loads sampled from the binomial distribution and the rest will be dummy items. The remaining overflow items are distributed recursively in following levels of the OHT.

The items assigned to each bin in each level of the OHT are instantiated with another oblivious hash table construction which we call *oblivious bin*. The main difference between a general OHT and an oblivious bin comes from the fact that we will use the oblivious bin only on small inputs of size $O(\text{poly}(\log \lambda))$ for which we can afford to use an oblivious sort for initialization without incurring prohibitive efficiency cost. We present two instantiations for the oblivious bin using a binary tree and a Cuckoo hash table.

Last but not least, our oblivious hash tables have an oblivious extraction procedure that allows to separate the unqueried items in the OHT with just an overhead of $O(\log \log \lambda)$ per item. Additionally we guarantee that the extracted items are randomly shuffled and, thus, we can use them directly as input for our multi-array shuffle. The extraction procedure for our OHT can be done by implementing the extraction on each of the oblivious bins and concatenating the outputs, since the items were distributed to bins using a secret distribution function. We obliviously extract each bin using an oblivious sort.

Oblivious Random Multi-Array Shuffle. Our multi-array shuffle algorithm starts with several independently shuffled arrays and outputs a random shuffle of all items in the arrays. The final output shuffle can be described by a random permutation. As we discuss in Appendix A, one way to sample a random permutation is to partition the input in L parts and sample a random permutation for each of the parts as well as a random `Assign` function that merges the elements from the separate arrays. The oblivious shuffle applying the permutation sampled in the above form, can be broken down to applying the oblivious shuffles on each of the corresponding partition arrays and then applying the `Assign` function in an oblivious manner. However, our assumption on the input is essentially equivalent to the assumption that the input partition arrays are already shuffled according to their respective local permutations. Thus, the functionality of our random multi-array shuffle amounts to sampling a random `Assign` and applying it obliviously on the shuffled input arrays.

The main idea for our multi-array random shuffle relies on the observation that we do not need to hide the access pattern within each of the input arrays since they are already shuffled. Since our shuffle algorithm will be accessing each entry of each input array only once, its initial random shuffle suffices for the obliviousness of these accesses. However, we still need to hide the interleaving accesses to the different input arrays. One way to achieve this is to obliviously shuffle the accesses to different input arrays. If we do this, in general, we will end up doing an oblivious shuffle on the whole input data, which is too expensive. Instead, we will partition the output array, which will contain the result of our multi-array random shuffle, in small segments of size $O(\log^3 \lambda)$, and we will fill each of its parts doing an oblivious sort on all the items destined for that part.

In more detail, we expect that each partition of the output array will contain a number of items from each input array that is close to the expectation (i.e., proportional to the fraction of the total number of items assigned to that array). Our algorithm retrieves from each input array the items that are assigned to the particular output array partition. More precisely, our algorithm retrieves a prefixed number of items that, except with negligible probability, is an upper bound on the actual number of *real* items needed thus hiding the exact number of items retrieved from the array. These accesses are already oblivious since we do not need to hide access patterns within each array and each item is accessed exactly once. In fact, since we are generating a random permutation on the fly, we can just retrieve the first items in each of the input arrays since all input arrays are randomly shuffled.

Once we have retrieved these items, we will sort them obliviously to separate the exact number of items that should be taken from each input array without revealing these numbers. All remaining items will be considered leftovers, which will be separated according to the input arrays they came from and padded to hide their exact sizes. The items which were separated (i.e., that were not assigned to leftover bins) are obviously sorted to be assigned to their exact positions in the output array partition currently considered. We will apply the multi-array random shuffling algorithm recursively on the arrays containing leftover items from the executions filling different output partitions.

The intuition why the above approach helps us to improve our efficiency is that we are using oblivious sort on small arrays that are of size $O(\log^3 \lambda)$ and we perform $O(\frac{N}{\log^3 \lambda})$ of these shuffles. Thus, the total shuffle cost remains $O(N \log \log \lambda)$.

Our resulting shuffling procedure achieves such efficiency that it is no longer the dominant cost in the amortized query complexity for our final ORAM construction. As a result, the optimization technique presented in the work of Kushilevitz *et al.* [28], which balances the cost of the lookups and the cost of the shuffle by splitting each level into several disjoint oblivious hash tables that get shuffled into separately, does not result in any efficiency improvement when applied to our scheme.

Paper Organization. We present in Section 2 the definitions of the primitives we use in the rest of the paper. Section 3 describes our construction of an oblivious random multi-array shuffle, which relies on permutation decomposition lemmas that provide alternative ways to sample randomly a permutation, which are presented in Appendix A. In Section 4 we describe our general oblivious hash table construction and its building block the oblivious bin primitive, for which we provide two different constructions. Finally, we present our ORAM construction in Section 5. In Appendix B we present an additional overview of related work.

2 Definitions

In this section we present the definitions for the existing primitives that we use for our constructions as well as definitions of the new primitives that we introduce in our work.

Notation. We denote $\text{Binomial}[n, p]$ the binomial distribution with parameters: n trials each of which with success probability p . We use $X \leftarrow \text{Binomial}[n, p]$ to denote that the variable X is sampled from the binomial distribution according to its probability mass function $\Pr(k; n, p) = \Pr[X = k] = \binom{n}{k} p^k (1 - p)^{n-k}$.

In a setting where an algorithm Alg is executing using external memory, we denote by $\text{Addrs}[\text{Alg}]$ the memory access pattern that consists of all accessed addresses in the memory. We use PPT as

a shorthand for “probablistic polynomial-time.”

In analyzing our constructions, we express bandwidth and client memory using the size B of a block as a unit.

2.1 Oblivious RAM

Definition 1 (Oblivious RAM). An oblivious RAM scheme $\text{ORAM} = (\text{ORAM.Init}, \text{ORAM.Access})$ consists the following two algorithms:

- $(\tilde{D}, \text{st}) \leftarrow \text{ORAM.Init}(1^\lambda, D)$: the initialization algorithm takes as input a database D and outputs a initialized memory structure \tilde{D} .
- $(v, \text{st}') \leftarrow \text{ORAM.Access}(\text{st}, \tilde{D}, l)$: the ORAM access algorithm takes as input the ORAM database \tilde{D} , the current state st as well as an instruction $l = (\text{op}, \text{addr}, \text{data})$, where $\text{op} \in \{\text{read}, \text{write}\}$ and if $\text{op} = \text{read}$, then $\text{data} = \perp$. If $\text{op} = \text{read}$, the access algorithm returns as v the data stored at address addr in the database. Else, if $\text{op} = \text{write}$ the access algorithm writes data in location addr in the database. Furthermore, an updated state st' is returned.

The resulting construction is oblivious if there exists a PPT simulator $\text{Sim} = (\text{Sim}_{\text{Init}}, \text{Sim}_{\text{Access}})$ such that for any PPT adversarial algorithm \mathcal{A} and for any $n = \text{poly}(\lambda)$,

$$\left| \Pr [b = 1 \mid b \leftarrow \text{Expt}_{\mathcal{A}}^{\text{Real,ORAM}}(\lambda, n)] - \Pr [b' = 1 \mid b' \leftarrow \text{Expt}_{\text{Sim}, \mathcal{A}}^{\text{Ideal,ORAM}}(\lambda, n)] \right| < \text{negl}(\lambda),$$

where the real and ideal executions are defined as follows:

$\text{Expt}_{\mathcal{A}}^{\text{Real,ORAM}}(\lambda, n)$	$\text{Expt}_{\text{Sim}, \mathcal{A}}^{\text{Ideal,ORAM}}(\lambda, n)$
$(D, \text{st}_{\mathcal{A}}) \leftarrow \mathcal{A}_1(1^\lambda)$	$(D, \text{st}_{\mathcal{A}}) \leftarrow \mathcal{A}_1(1^\lambda)$
$\chi \leftarrow \text{Addrs}[(\tilde{D}, \text{st}) \leftarrow \text{ORAM.Init}(1^\lambda, D)]$	$\chi \leftarrow \text{Addrs}[(\tilde{D}, \text{st}_{\text{Sim}}) \leftarrow \text{Sim}_{\text{Init}}(1^\lambda, D)]$
for $j = 1$ to n	for $i = 1$ to n
$(l_j, \text{st}_{\mathcal{A}}) \leftarrow \mathcal{A}_2(\text{st}_{\mathcal{A}}, \tilde{D}, \chi)$	$(l_j, \text{st}_{\mathcal{A}}) \leftarrow \mathcal{A}_2(\text{st}_{\mathcal{A}}, \tilde{D}, \chi)$
$\chi_{\text{Access}} \leftarrow \text{Addrs}[(\text{st}, \tilde{D}) \leftarrow \text{ORAM.Access}(\text{st}, \tilde{D}, l_j)]$	$\chi_{\text{Access}} \leftarrow \text{Addrs}[(\text{st}, \tilde{D}) \leftarrow \text{Sim}_{\text{Access}}(\text{st}_{\text{Sim}})]$
$\chi \leftarrow \chi \cup \chi_{\text{Access}}$	$\chi \leftarrow \chi \cup \chi_{\text{Access}}$

2.2 Oblivious Random Multi-Array Shuffle

In oblivious shuffling, we have one array of N data blocks and the task of the algorithm is to shuffle the blocks into a destination array so that an adversary observing the accesses of the algorithm to the blocks does not obtain any information regarding the final permutation. The security guarantee holds even if the initial arrangement of the blocks is known to the adversary. This is sufficient to be used in the design of Oblivious RAM. An oblivious random multi-array shuffle instead offers a weaker security guarantee, which we show also suffices for the design of an Oblivious RAM. For a range of parameters of interest, we present an implementation with improved efficiency in terms of bandwidth overhead compared to oblivious shuffling. Roughly speaking, in an oblivious random multi-array random shuffle, we have N blocks partitioned into L arrays $\mathbf{A}_1, \dots, \mathbf{A}_L$ and the task is to shuffle the N blocks into a destination array according to a permutation chosen uniformly at random. The associated security notion still guarantees that no information regarding

the final permutation of the blocks is leaked. However, the adversary that only has knowledge of the distribution of blocks to each array, but not of each block's specific location within the array. Let us now proceed more formally.

2.3 Oblivious Bin

We introduce a slight modification of our oblivious hash table definition, which we call *oblivious bin*. As the name implies, we will use this oblivious structure to store and access data in oblivious manner within bins which will be used as building blocks in our oblivious hash table scheme. We will use the oblivious bin for data of smaller size.

Definition 2 (Oblivious Bin). An oblivious hash table scheme $\text{OblivBin} = (\text{OblivBin.Build}, \text{OblivBin.Lookup}, \text{OblivBin.Extract})$ consists of the following algorithm:

- $(\tilde{D}, \text{st}) \leftarrow \text{OblivBin.Init}(D)$: an algorithm that takes as input an array of key-value pairs $D = \{(k_i, v_i)\}_{i=1}^N$ and outputs a processed version of it \tilde{D} .
- $(\tilde{S}, \tilde{H}, \text{st}') \leftarrow \text{OblivBin.Build}(\tilde{D}, \text{st})$: an algorithm that takes as input a processed database \tilde{D} and a state and initializes the hash table \tilde{H} and an additional array \tilde{S} and updates the state st .
- $(v, \tilde{S}', \tilde{H}', \text{st}') \leftarrow \text{OblivBin.Lookup}(k, \tilde{H}, \tilde{S}, \text{st}')$: an algorithm that takes as input the oblivious hash table, the additional array, the state produced in the build stage and a lookup key, and outputs the value v_i corresponding to the key k_i together with updated hash table \tilde{H}' and state st' . If the k is not found, then $v := \perp$.
- $(\tilde{D}, \text{st}') \leftarrow \text{OblivBin.Extract}(\tilde{H}, \tilde{S}, \text{st}')$: this is an algorithm that takes the hash table, the auxiliary array and the state after the execution of a number of queries and outputs a database, which contains only the unqueried items $(k_i, v_i) \in D$ and is padded to size N .

The oblivious property for a bin is identical to that of the oblivious hash table.

2.4 Oblivious Random Multi-Array Shuffle

We start with the definition of a random multi-array shuffling algorithm that obviously shuffles together the content of several input array each of which is independently shuffled (see Appendix 2.2 for more detailed discussion of the functionality).

Definition 3 (Oblivious Random Multi-Array Shuffle). A *random multi-array shuffle* algorithm is an algorithm $D \leftarrow \text{OblMultArrShuff}(A_1, \dots, A_L)$ that takes as input L arrays A_1, \dots, A_L containing a total of N blocks and outputs a destination array D that contain all N blocks. Each of the L arrays are assumed to have been arranged according to a permutation chosen uniformly at random. The blocks in D should be arranged according to a permutation chosen uniformly at random.

A random multi-array shuffle algorithm OblMultArrShuff is *oblivious* if there exists a PPT simulator Sim , which takes as input $(|A_1|, \dots, |A_L|)$, such that for any PPT adversary algorithm $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$:

$$\left| \Pr \left[b = 1 \mid b \leftarrow \mathbf{Expt}_{\mathcal{A}}^{\text{Real, OblMultArrShuff}}(\lambda) \right] - \Pr \left[b' = 1 \mid b' \leftarrow \mathbf{Expt}_{\text{Sim}, \mathcal{A}}^{\text{Ideal, OblMultArrShuff}}(\lambda) \right] \right| < \text{negl}(\lambda),$$

where the real and ideal executions are defined as follows:

$\mathbf{Expt}_{\mathcal{A}}^{\text{Real, OblMultArrShuff}}(\lambda)$	$\mathbf{Expt}_{\text{Sim}, \mathcal{A}}^{\text{Ideal, OblMultArrShuff}}(\lambda)$
$(A_1, \dots, A_L, \text{st}_{\mathcal{A}}) \leftarrow \mathcal{A}_1(1^\lambda)$ for $i = 1$ to L $\tilde{A}_i \leftarrow \tau_i(A_i)$ where τ_i is a random permutation on $ A_i $ elements $\chi \leftarrow \text{Addrs}[D \leftarrow \text{OblMultArrShuff}(A_1, \dots, A_L)]$ Let π be the induced permutation of the elements in A_1, \dots, A_L mapped to D Output $b \leftarrow \mathcal{A}_2(D, \pi, \text{st}_{\mathcal{A}}, \chi)$	$(A_1, \dots, A_L, \text{st}_{\mathcal{A}}) \leftarrow \mathcal{A}_1(1^\lambda)$ $\chi \leftarrow \text{Addrs}[D \leftarrow \text{Sim}(A_1 , \dots, A_L)]$ Let π' be a random permutation on $\sum_{i=1}^L A_i $ elements Output $b \leftarrow \mathcal{A}_2(D, \pi', \text{st}_{\mathcal{A}}, \chi)$

Intuitively, the definition above captures the security with respect to an adversary that does not know the original order of the items in the input arrays (i.e., each array is randomly shuffled) but is allowed to pick the partition of the N blocks across the L levels. We require that the adversary does not obtain any information about the final permutation of the N blocks and we formally capture this requirement by providing the adversary with the actual induced permutation in the real execution and a completely random permutation in the ideal execution. Thus, if the adversary cannot distinguish the real and the ideal experiment, it follows that it has not learned anything about the permutation by observing the access pattern leaked from the shuffle algorithm.

2.5 Oblivious Hash Table

Next we define our oblivious hash table, which provides oblivious access for non-repeating queries. In addition it had an extraction algorithm which allows to extract obliviously the unqueried items remaining in the OHT in a randomly permuted order.

Definition 4 (Oblivious Hash Table). An oblivious hash table scheme $\text{OblivHT} = (\text{OblivHT.Build}, \text{OblivHT.Lookup}, \text{OblivHT.Extract})$ consists of the following algorithm:

- $(\tilde{D}, \text{st}) \leftarrow \text{OblivHT.Init}(D)$: an algorithm that takes as input an array of key-value pairs $D = \{(k_i, v_i)\}_{i=1}^N$ and outputs a processed version of D , which we denote \tilde{D} .
- $(\tilde{H}, \text{st}') \leftarrow \text{OblivHT.Build}(\tilde{D}, \text{st})$: an algorithm that takes as input a processed database \tilde{D} and a state and initializes the hash table \tilde{H} and updates the state st .
- $(v, \tilde{H}', \text{st}') \leftarrow \text{OblivHT.Lookup}(k, \tilde{H}, \text{st})$: an algorithm that takes as input the oblivious hash table, \tilde{H} , the state produced in the build stage, st , and a lookup key, k , and outputs the value v_i corresponding to the key k_i together with an updated hash table, \tilde{H}' , and updated state, st' . If the k is a dummy query, then $v := \perp$.
- $(\tilde{D}, \text{st}') \leftarrow \text{OblivHT.Extract}(\tilde{H}, \text{st})$: an algorithm that takes the hash table, \tilde{H} , and the state after the execution of a number of queries, st , and outputs a database \tilde{D} , which contains only the unqueried items $(k_i, v_i) \in D$ and is padded to size N with dummy items, and the content of \tilde{D} is randomly permuted.

The resulting hash scheme is oblivious if there exists a PPT simulator $\text{Sim} = (\text{Sim}_{\text{Init}}, \text{Sim}_{\text{Build}}, \text{Sim}_{\text{Lookup}}, \text{Sim}_{\text{Extract}})$, which takes as input the size of the database $|D|$, such that for any PPT adversary algorithm $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3)$ and for any $n = \text{poly}(\lambda)$,

$$\left| \Pr \left[b = 1 \mid b \leftarrow \mathbf{Expt}_{\mathcal{A}}^{\text{Real, OblivHT}}(\lambda, n) \right] - \Pr \left[b' = 1 \mid b' \leftarrow \mathbf{Expt}_{\text{Sim}, \mathcal{A}}^{\text{Ideal, OblivHT}}(\lambda, n) \right] \right| < \text{negl}(\lambda),$$

where the real and ideal executions are defined as follows:

$\text{Expt}_{\mathcal{A}}^{\text{Real, OblivHT}}(\lambda, n)$	$\text{Expt}_{\text{Sim}, \mathcal{A}}^{\text{Ideal, OblivHT}}(\lambda, n)$
$(D, \text{st}_{\mathcal{A}}) \leftarrow \mathcal{A}_1(1^\lambda)$ $\chi_{\text{Init}} \leftarrow \text{Addr}[(\tilde{D}, \text{st}) \leftarrow \text{OblivHT.Init}(D)]$ $\chi_{\text{Build}} \leftarrow \text{Addr}[(\tilde{H}, \text{st}) \leftarrow \text{OblivHT.Build}(D, \text{st})]$ $\chi_{\text{Lookup}} \leftarrow \perp, \chi \leftarrow (\chi_{\text{Init}}, \chi_{\text{Build}})$ for $i = 1$ to n $(k_i, \text{st}_{\mathcal{A}} \mid \{k_j \neq k_j\}_{1 \leq j < i}) \leftarrow \mathcal{A}_2(\tilde{H}, \text{st}_{\mathcal{A}}, \chi, \chi_{\text{Lookup}})$ $\chi_{\text{Lookup}} \leftarrow \text{Addr}[(\tilde{H}, v_i, \text{st}) \leftarrow \text{OblivHT.Lookup}(k_i, \tilde{H}, \text{st})]$ $\chi_{\text{Extract}} \leftarrow \text{Addr}[(\tilde{D}, \text{st}) \leftarrow \text{OblivHT.Extract}(\tilde{H}, \text{st})]$ Let π be the permutation induced on the items $D/\{k_i\}_{i=1}^n$ and the padding dummies in \tilde{D} Output $b \leftarrow \mathcal{A}_3(\tilde{D}, \pi, \text{st}_{\mathcal{A}}, \chi_{\text{Extract}})$	$(D, \text{st}_{\mathcal{A}}) \leftarrow \mathcal{A}_1(1^\lambda)$ $\chi_{\text{Init}} \leftarrow \text{Addr}[(\tilde{D}, \text{st}_{\text{Sim}}) \leftarrow \text{Sim}_{\text{Init}}(D)]$ $\chi_{\text{Build}} \leftarrow \text{Addr}[(\tilde{H}, \text{st}_{\text{Sim}}) \leftarrow \text{Sim}_{\text{Build}}(D)]$ $\chi_{\text{Lookup}} \leftarrow \perp, \chi \leftarrow (\chi_{\text{Init}}, \chi_{\text{Build}})$ for $i = 1$ to n $(k_i, \text{st}_{\mathcal{A}} \mid \{k_j \neq k_j\}_{1 \leq j < i}) \leftarrow \mathcal{A}_2(\tilde{H}, \text{st}_{\mathcal{A}}, \chi, \chi_{\text{Lookup}})$ $\chi_{\text{Lookup}} \leftarrow \text{Addr}[(\tilde{H}, \text{st}_{\text{Sim}}) \leftarrow \text{Sim}_{\text{Lookup}}(\tilde{H}, \text{st}_{\text{Sim}})]$ $\chi_{\text{Extract}} \leftarrow \text{Addr}[\tilde{H} \leftarrow \text{Sim}_{\text{Extract}}(\tilde{H}, \text{st}_{\text{Sim}})]$ Let π' be a random permutation on N items Output $b \leftarrow \mathcal{A}_3(\tilde{D}, \pi', \text{st}_{\mathcal{A}}, \chi_{\text{Extract}})$

Although it may not be apparent at this stage why we have separate `OblivHT.Init` and `OblivHT.Build` algorithms, the reason is that in the context of our ORAM construction we will instantiate the `OblivHT.Init` algorithm with our oblivious random multi-array shuffle.

We want to guarantee that the output of extract algorithm is randomly shuffled from the point of view of the adversary. We formalize this similarly to the oblivious random multi-array shuffle, by providing the adversary with the real induced permutation in the real execution and a random independent permutation in the ideal execution. If the adversary cannot distinguish which is the real permutation, it means that it did not learn anything from the access patterns it observed.

3 Oblivious Random Multi-Array Shuffle

In this section, we present a novel oblivious random multi-array random shuffling algorithm that realizes the idea of leveraging entropy in the input. In particular, we assume that each input array has been previously shuffled in an order that is not known by the adversary. This assumption allows us to achieve better efficiency. Our algorithm improves on general oblivious sorting algorithms achieving bandwidth of $O(N \log \log \lambda + \frac{N \log N}{\log \lambda})$ blocks for N data blocks. Our entropy requirement for the input comes in the following form: the N input blocks are divided in L arrays each of which is randomly permuted (with a permutation not known to the server holding the arrays). The arrays have sizes N_1, \dots, N_L where $N_i \geq N_{i+1}$ for $i = 1, \dots, L - 1$. While our multi-array random shuffle algorithm is general and works on any distribution of the array sizes, we achieve the above improved efficiency under the restriction that there exists $\text{cutoff} = O(\log \log \lambda)$ such that $|A_{\text{cutoff}}| + \dots + |A_L| = O(\frac{N \log \log \lambda}{\log N})$.

Construction 5. [Oblivious Random Multi-Array Shuffle] We define our oblivious random multi-array shuffle algorithm `ObIMultArrShuff`, which will also use algorithms `ObIMultArrShuff.BinShuffle` and `ObIMultArrShuff.Shuffle` as building blocks.

`ObIMultArrShuff`. This algorithm takes as input L shuffled arrays and outputs array D , which contains a random permutation of the input array elements.

$D \leftarrow \text{OblMultArrShuff}(A_1, \dots, A_L)$:

1. Initialize the output array D to be an empty array of size $N := |A_1| + \dots + |A_L|$ blocks.
2. Choose the largest cutoff such that $A_{\text{cutoff}} \geq \frac{N}{\log \lambda}$ and then randomly permute the entries of the arrays $A_{\text{cutoff}}, \dots, A_L$ into $A_{0,\text{cutoff}}$ using an oblivious random shuffle.
3. Initialize $A_{0,1}, \dots, A_{0,\text{cutoff}-1}$ with the content of $A_1, \dots, A_{\text{cutoff}-1}$ respectively.
4. Sample a random assignment function $\text{Assign} : [N] \rightarrow [\text{cutoff}]$ such that $|\{b \in [N] : \text{Assign}(b) = i\}| = |A_{0,i}|$ for every $i \in [\text{cutoff}]$. Since we assume only constant local memory, which does not fit the description of Assign , we use the following oblivious algorithm for sampling Assign at random:
 - (a) Store an encryption of $\text{cnt}_i := |A_{0,i}|$, for $i \in [\text{cutoff}]$ and an encryption of $\text{cnt} = N$ on the server.
 - (b) For each block $b \in [N]$, set $\text{Assign}(b) = i$ with probability $\frac{\text{cnt}_i}{\text{cnt}}$, for $i \in [\text{cutoff}]$. We do this obliviously as follows: choose a random value $r_b \in [\text{cnt}]$ and set $\text{Assign}(b)$ to be equal to the minimal s such that $\text{cnt}_1 + \dots + \text{cnt}_s \geq r_b$ and s is computed by scanning the encrypted integers $\text{cnt}_1, \dots, \text{cnt}_{\text{cutoff}}$. During the scanning cnt_s and cnt are each decreased by 1. Store $(b, \text{Enc}(\text{Assign}(b)))$ at the server.
5. Let E_0 be the array $\{(b, \text{Enc}(\text{Assign}(b)))\}_{b \in [N]}$ describing Assign computed and stored at the server in the previous step. Note, the server knows the first indices of each pair in E_0 as they are unencrypted.
6. Set $L := \text{cutoff}$.
7. Execute $\text{OblMultArrShuff.Shuffle}(A_{0,1}, \dots, A_{0,L}, D, E_0, N, 0)$.
8. Return D .

$\text{OblMultArrShuff.Shuffle}$. This algorithm takes as input L randomly shuffled arrays, an output array D , which might be partially filled, the set of empty indices b_i in D together with their corresponding encrypted $\text{Assign}(b_i)$ values stored in E_ℓ , and an index ℓ corresponding the current level of recursion. The algorithm fills D through several recursive steps using the encrypted values of Assign in E_ℓ and the input arrays $A_{\ell,1}, \dots, A_{\ell,L}$.

$\text{OblMultArrShuff.Shuffle}(A_{\ell,1}, \dots, A_{\ell,L}, D, E_\ell, \ell)$:

1. If $|A_{\ell,1}| + \dots + |A_{\ell,L}| \leq \frac{N}{\log \lambda}$, assign the items in $A_{\ell,1}, \dots, A_{\ell,L}$ to the remaining open positions in D using the Assign mappings stored in E_ℓ by running:

$$\text{OblMultArrShuff.BinShuffle}(A_{\ell,1} \cup \dots \cup A_{\ell,L}, E_\ell, D).$$

2. Set $m := (2\epsilon)^\ell \frac{N}{\log^3 \lambda}$ and $\tilde{m} := (1 - 2\epsilon)m$.

3. Initialize \tilde{m} input bins $\text{Bin}_1^{\text{in}}, \dots, \text{Bin}_{\tilde{m}}^{\text{in}}$ with random subsets of blocks from the inputs arrays as follows. For each input array $A_{\ell,i}$, $i \in [L]$ distribute its blocks across $\text{Bin}_1^{\text{in}}, \dots, \text{Bin}_{\tilde{m}}^{\text{in}}$ assigning each block a bin at random and recording an encryption of the source array index i . The items of $\text{Bin}_1^{\text{in}}, \dots, \text{Bin}_{\tilde{m}}^{\text{in}}$ are stored encrypted on the server. Note $\text{Bin}_1^{\text{in}}, \dots, \text{Bin}_{\tilde{m}}^{\text{in}}$ will have different sizes. Furthermore, the above is done in a non-oblivious manner and the server knows the distribution of the blocks from each $A_{\ell,i}$ across the input bins.
4. Initialize m output bins $\text{Bin}_1^{\text{out}}, \dots, \text{Bin}_m^{\text{out}}$ with random subsets of pairs from E_ℓ by assigning each pair from E_ℓ to a randomly selected bin. The items of $\text{Bin}_1^{\text{out}}, \dots, \text{Bin}_m^{\text{out}}$ are stored encrypted on the server. Note the sizes of $\text{Bin}_1^{\text{out}}, \dots, \text{Bin}_m^{\text{out}}$ will be different. Furthermore, the above is done in a non-oblivious manner and the server knows the distribution of the blocks from each E_ℓ across the output bins. Therefore, the server knows the subset of positions from D assigned to each output bin - these are the b values in each pair $(b, \text{Enc}(\text{Assign}(b)))$ of E_ℓ .
5. Initialize $A_{\ell+1,1}, \dots, A_{\ell+1,L}$ to be empty block arrays.
6. For $j = 1, \dots, \tilde{m}$:
 - (a) Distribute the blocks from Bin_j^{in} in D according to the positions specified by Assign in the pairs in $\text{Bin}_j^{\text{out}}$ by running
$$(\text{LeftoverBin}_1, \dots, \text{LeftoverBin}_L) \leftarrow \text{ObIMultArrShuff.BinShuffle}(\text{Bin}_j^{\text{in}}, \text{Bin}_j^{\text{out}}, D).$$
 - (b) Append LeftoverBin_1 to $A_{\ell+1,1}, \dots, \text{LeftoverBin}_L$ to $A_{\ell+1,L}$.
7. Collect all uninitialized indices in D together with their corresponding mappings under Assign , which have been distributed in output bins $\text{Bin}_{\tilde{m}+1}^{\text{out}}, \dots, \text{Bin}_m^{\text{out}}$, and set $E_{\ell+1} := \text{Bin}_{\tilde{m}+1}^{\text{out}} \cup \dots \cup \text{Bin}_m^{\text{out}}$.
8. Execute recursively the shuffling functionality on the remainders of the input arrays, which have not been placed in D so far but were returned as leftovers from the $\text{ObIMultArrShuff.BinShuffle}$ executions above, by running:

$$\text{ObIMultArrShuff.Shuffle}(A_{\ell+1,1}, \dots, A_{\ell+1,L}, D, E_{\ell+1}, \ell + 1).$$

BinShuffle. This is an algorithm that takes as input a bin Bin^{in} that contains items, a bin Bin^{out} that contains mappings under Assign of a subset of indices in the input D . BinShuffle distributes all but an 2ϵ fraction of the items in Bin^{in} into D according to the mappings and positions of D specified in Bin^{out} . The items of Bin^{in} that are not placed in D are returned in leftover bins separated according to their input arrays.

$$(\text{LeftoverBin}_1, \dots, \text{LeftoverBin}_L) \leftarrow \text{ObIMultArrShuff.BinShuffle}(\text{Bin}^{\text{in}}, \text{Bin}^{\text{out}}, D)$$

1. For $i \in [L]$, create $\text{NumLeftover}_i := (4\epsilon) \frac{N_i}{N} \log^3 \lambda$ dummy blocks tagged with an array index i and append them encrypted to Bin^{in} .
2. Obviously sort Bin^{in} according to array index of the blocks placing real blocks before dummy blocks with the same array index.

3. Let k_i^{out} be the number of pairs $(b, \text{Assign}(b)) \in \text{Bin}^{\text{out}}$ such that $\text{Assign}(b) = i$ for all $i \in [L]$. We compute the values k_i^{out} privately in the following oblivious manner:
 - (a) Initialize all $k_1^{\text{out}}, \dots, k_L^{\text{out}}$ to 0 and store them encrypted on the server.
 - (b) For each pair $(b, \text{Assign}(b)) \in \text{Bin}^{\text{out}}$, scan all the ciphertexts of $k_1^{\text{out}}, \dots, k_L^{\text{out}}$ and only increment $k_{\text{Assign}(b)}^{\text{out}}$.
4. Tag all items in Bin^{in} with moving, if they are a real item that will be placed in D using the Assign mapping from Bin^{out} ; leftover if they are real or dummy items that will be returned as leftover; or unused if they are dummy items that will be discarded. We tag the items in the following oblivious manner: for each block j in Bin^{in} :
 - (a) Let i be the input array index of the j -th block in Bin^{in} .
 - (b) For $t \in [L]$, download the counter k_t^{out} . If $t \neq i$, reencrypt and upload back the counter. If $t = i$, upload an encryption of a decremented counter $\text{Enc}(k_t^{\text{out}} - 1)$.
 - (c) Tag the j -th block as follows: if $k_{\ell,i}^{\text{out}} > 0$, then the block is marked as real. If $-\text{NumLeftover}_i < k_{\ell,i}^{\text{out}} \leq 0$, then the block is marked as leftover. Otherwise, the block is marked as unused.
5. Obviously sort Bin^{in} according to the tags computed in the previous step in a manner where all blocks with tag moving precede all blocks with tag leftover and both of these precede blocks with tags unused. All blocks with the same tag are sorted according to their input array index.
6. The blocks in Bin^{in} are separated in the following way:
 - (a) Blocks that will be placed in D - these are the first $|\text{Bin}^{\text{out}}|$ blocks in the sorted Bin^{in} , which are moved to TempD;
 - (b) Blocks that will be returned as leftover blocks: these are the next L groups of blocks of sizes $\text{NumLeftover}_1, \dots, \text{NumLeftover}_L$ blocks, which are placed in $\text{LeftoverBin}_1, \dots, \text{LeftoverBin}_L$ respectively.
7. Obviously sort the pairs $(b, \text{Assign}(b))$ in Bin^{out} according to input array index $\text{Assign}(b)$. Before sorting, encrypt the b value of each pair of Bin^{out} to ensure obliviousness. Recall the $\text{Assign}(b)$ value is already encrypted. Note that TempD and Bin^{out} now contain the same numbers of items tagged with each of the input array indices, which are also sorted according to these indices. Thus, for each position $i \in [|\text{Bin}^{\text{out}}|]$, the block in position i in TempD has input array index equal to $\text{Assign}(b_i)$, where $(b_i, \text{Assign}(b_i))$ is the i -th pair in the sorted Bin^{out} .
8. Assign to each block in TempD its corresponding location in D as follows: for $i \in [|\text{Bin}^{\text{out}}|]$, tag the block in position i in TempD with an encryption of b_i , where $(b_i, \text{Assign}(b_i))$ is the i -th pair in Bin^{out} .
9. Obviously sort TempD according to tags computed in the previous step and copy the content of TempD in the positions denoted by their tags in D. Note these positions of D are public since the positions from D assigned to Bin^{out} are known to the server and only encrypted in Step 7.

How the input shrinks over calls. To get an understanding of `ObIMultArrShuff`, we will analyze the sizes of the inputs to recursive calls of `ObIMultArrShuff.Shuffle`. Let us start with the very first execution of `ObIMultArrShuff.Shuffle`. Initially, $A_{0,1}, \dots, A_{0,L}$ collectively contain exactly N real blocks and E_0 contains N pairs (since all indices of D are still free). The indices of E_0 are uniformly and independently assigned at random to $m := \frac{N}{\log^3 \lambda}$ output bins and each is expected to have $\log^3 \lambda$ indices. The blocks of $A_{0,1}, \dots, A_{0,L}$ are uniformly and independently distributed into $\tilde{m} := (1 - 2\epsilon)m$ input bins and each is expected to have $\log^3 \lambda / (1 - 2\epsilon) \approx (1 + 2\epsilon) \log^3 \lambda$ blocks. We note that E_1 consists of only indices of the last $2\epsilon \cdot m$ groups (see Step 7). Therefore, E_1 will contain $2\epsilon N$ indices in expectation.

Each execution of `ObIMultArrShuff.BinShuffle` outputs L leftover bins of sizes $\{(4\epsilon) \frac{N_i}{N} \log^3 \lambda\}_{i=1, \dots, L}$. So, after m executions of `ObIMultArrShuff.BinShuffle`, a total of $4\epsilon N_i$ will be placed into each $A_{1,i}$, for $i = 1, \dots, L$. As a result of `ObIMultArrShuff.Shuffle`, E_1 is only a 2ϵ fraction of the size of E_0 . Similarly, each $A_{1,i}$ is only a 4ϵ fraction of the size of $A_{0,i}$. This reduction continues as `ObIMultArrShuff.Shuffle` is executed more times. In particular, $A_{\ell,i}$ will contain exactly $(4\epsilon)^\ell N_i$ blocks and E_ℓ contains $(2\epsilon)^\ell N$ indices in expectation.

The above analysis considers counting both real and dummy blocks. We turn our attention to strictly real blocks. At each level of `ObIMultArrShuff.Shuffle`, an 2ϵ fraction of the blocks are unassigned and will be dealt with by the later executions of `ObIMultArrShuff.Shuffle`. As we have seen, the unassigned blocks are kept partitioned by array index and, in order to hide the actual number of blocks that are still to be assigned from each array, dummy blocks are introduced. The next lemma shows that the number of real blocks in $A_{\ell,i}$, $N_{\ell,i}$, is binomially distributed.

Lemma 6. The random variable $N_{\ell,i}$ is distributed according to $\text{Binomial}[N_i, (2\epsilon)^\ell]$.

Proof. We will use the observation that for any $i \in [L]$,

$$N_{\ell,i} = |\{(b, \text{Assign}(b)) \in E_\ell : \text{Assign}(b) = i\}|.$$

The lemma follows by the fact that the events that $(b, \text{Assign}(b)) \in E_\ell$ are independent and have probability $(2\epsilon)^\ell$, for every $\ell \geq 0$.

For $\ell = 0$, this is trivially true. By inductive hypothesis, assume $\Pr[(b, i) \in E_\ell] = (2\epsilon)^\ell$, for some $\ell \geq 0$. All pairs of E_ℓ are distributed uniformly and independently at random into m bins in Step 4 of `ObIMultArrShuff.Shuffle`. A single index $(b, \text{Assign}(b)) \in E_\ell$ appears in $E_{\ell+1}$ if and only if b is assigned to one of $\text{Bin}_{\tilde{m}+1}^{\text{out}}, \dots, \text{Bin}_m^{\text{out}}$. Therefore,

$$\Pr[(b, i) \in E_{\ell+1} \mid (b, i) \in E_\ell] = \frac{m - \tilde{m}}{m} = 2\epsilon.$$

Since the assignment of E_ℓ is done independently of all previous events, we see that

$$\Pr[(b, i) \in E_{\ell+1}] = \Pr[(b, i) \in E_\ell] \cdot 2\epsilon = (2\epsilon)^{\ell+1}.$$

□

In the next lemma, we bound the number of blocks, $k_{\ell,i,j}^{\text{in}}$ of array of index i , A_i , that are assigned to the j -th input bin Bin_j^{in} of level ℓ of the algorithm.

Lemma 7. For every i and j and for every level ℓ ,

$$\Pr \left[(1 + \epsilon) \cdot \frac{N_i}{N} \log^3 \lambda \leq k_{\ell,i,j}^{\text{in}} \leq (1 + 3\epsilon) \cdot \frac{N_i}{N} \log^3 \lambda \right] \geq 1 - \text{negl}(\lambda).$$

Proof. By Lemma 6, we know that $N_{\ell,i}$ is distributed according to $\text{Binomial}[N_i, (2\epsilon)^\ell]$. Moreover, each block in $\mathbf{A}_{\ell,i}$ is independently and uniformly assigned to one input bin at Step 3 and thus $k_{\ell,i,j}^{\text{in}}$ is distributed according to

$$\text{Binomial} \left[N_i, \frac{(2\epsilon)^\ell}{\tilde{m}} \right] = \text{Binomial} \left[N_i, \frac{\log^3 \lambda}{(1-2\epsilon)N} \right].$$

Therefore

$$\mu_{\ell,i,j} := \mathbb{E}[k_{\ell,i,j}^{\text{in}}] = (1+2\epsilon) \cdot \frac{N_i}{N} \cdot \log^3 \lambda$$

where we used the approximation $(1+2\epsilon) \approx \frac{1}{1-2\epsilon}$. By Chernoff Bounds, we get the following:

$$\begin{aligned} \Pr \left[(1+\epsilon) \cdot \frac{N_i}{N} \cdot \log^3 \lambda \leq k_{\ell,i,j}^{\text{in}} \leq (1+3\epsilon) \cdot \frac{N_i}{N} \cdot \log^3 \lambda \right] &= \Pr \left[(1-\epsilon)\mu_{\ell,i,j} \leq k_{\ell,i,j}^{\text{in}} \leq (1+\epsilon)\mu_{\ell,i,j} \right] \\ &\geq 1 - 2^{-\Omega(\mu_{\ell,i,j})}; \end{aligned}$$

where we used the approximations $(1+2\epsilon) \cdot (1+\epsilon) \approx (1+3\epsilon)$ and $(1+2\epsilon) \cdot (1-\epsilon) \approx (1+\epsilon)$. The lemma follows by observing that, by Step 2 of `ObIMultArrShuff`, we have $\frac{N}{N_i} < \log \lambda$ and therefore, $\mu_{\ell,i,j} = \Omega(\log^2 \lambda)$. \square

A similar lemma holds for the number of blocks, $k_{\ell,i,j}^{\text{out}}$, from \mathbf{A}_i that are assigned to a location of array \mathbf{D} that belongs to output bin $\text{Bin}_j^{\text{out}}$.

Lemma 8. For every i and j and for every level ℓ ,

$$\Pr \left[(1-\epsilon) \cdot \frac{N_i}{N} \cdot \log^3 \lambda \leq k_{\ell,i,j}^{\text{out}} \leq (1+\epsilon) \cdot \frac{N_i}{N} \cdot \log^3 \lambda \right].$$

Proof. Each pair of \mathbf{E}_ℓ is independent and uniformly assigned to one output bin at Step 4 of `ObIMultArrShuff.Shuffle`. So, $k_{\ell,i,j}^{\text{out}}$ is distributed according to

$$\text{Binomial} \left[N_i, \frac{(2\epsilon)^\ell}{m} \right] = \text{Binomial} \left[N_i, \frac{\log^3 \lambda}{N} \right]$$

and $\mu_{\ell,i,j} := \mathbb{E}[k_{\ell,i,j}^{\text{out}}] = \frac{N_i}{N} \cdot \log^3 \lambda$. By Chernoff bounds, we obtain that

$$\begin{aligned} \Pr \left[(1-\epsilon) \cdot \frac{N_i}{N} \log^3 \lambda \leq k_{\ell,i,j}^{\text{out}} \leq (1+\epsilon) \cdot \frac{N_i}{N} \log^3 \lambda \right] &= \Pr \left[(1-\epsilon)\mu_{\ell,i,j} \leq k_{\ell,i,j}^{\text{out}} \leq (1+\epsilon)\mu_{\ell,i,j} \right] \\ &\geq 1 - 2^{-\Omega(\mu_{\ell,i,j})}. \end{aligned}$$

The lemma follows by observing that, by Step 2 of `ObIMultArrShuff`, we have $\frac{N}{N_i} < \log \lambda$ and therefore, $\mu_{\ell,i,j} = \Omega(\log^2 \lambda)$. \square

Lemma 9. The construction `ObIMultArrShuff` aborts with probability negligible in λ .

Proof. `ObIMultArrShuff` aborts only if at Step 4 of `ObIMultArrShuff.BinShuffle` one of the following two events occurs, for some i, j, ℓ ,

1. The j -th input bin contains too few real blocks from array $A_{\ell,i}$ and thus a dummy block is marked moving. This happens if $k_{\ell,i,j}^{\text{in}} < k_{\ell,i,j}^{\text{out}}$. But by Lemma 7 and Lemma 8 we have that

$$k_{\ell,i,j}^{\text{in}} \geq (1 + \epsilon) \cdot \frac{N_i}{N} \cdot \log^3 \lambda \geq k_{\ell,i,j}^{\text{out}}$$

except with negligible probability.

2. The j -th input bin contains too many real blocks from array $A_{\ell,i}$ and thus a real block is marked unused. This happens if $k_{\ell,i,j}^{\text{in}} > k_{\ell,i,j}^{\text{out}} + \text{NumLeftover}_i = k_{\ell,i,j}^{\text{out}} + 4\epsilon \frac{N_i}{N} \cdot \log^3 \lambda$. Again by Lemma 7 and Lemma 8 we have that

$$k_{\ell,i,j}^{\text{in}} \leq (1 + 3\epsilon) \cdot \frac{N_i}{N} \log^3 \lambda \leq 4\epsilon \cdot \frac{N_i}{N} \log^3 \lambda + k_{\ell,i,j}^{\text{out}}$$

except with negligible probability.

□

Theorem 10. The construction `ObIMultArrShuff` is an oblivious random multi-array shuffle according to Definition 3.

Proof. We construct a simulation `Sim` that outputs D as follows. `Sim` will fill each A_i with encryptions of random values. `Sim` executes the honest `ObIMultArrShuff` algorithm, which will construct D filled with the encryptions of random values. The output array generated in the real and simulated execution are indistinguishable since they contain encryptions of the same number of items.

The access pattern while running `ObIMultArrShuff` involve either linear scans, oblivious sorts and movements of blocks determined by random coin flips (see Step 3 and 4 of `ObIMultArrShuff.Shuffle`). However, the random coin flips are revealed publicly and indistinguishable from any other truly random coin flip sequence. Therefore, the access patterns from the real and simulated execution are indistinguishable.

Finally, we have to show that final permutation given to the adversary is indistinguishable in conjunction with the access pattern. In the real experiment, each A_i is permuted according to a τ_i hidden from the adversary. `ObIMultArrShuff` applies an `Assign` function chosen uniformly at random. We show in Lemma 21 the result is a uniformly random permutation. Since `Assign` is chosen before revealing any random values in the access pattern, `Assign` is independent of access patterns in the real execution. Therefore, the final permutations in the real and simulated execution are indistinguishable completing the proof. □

Efficiency. For efficiency, we focus on the situation when there exists `cutoff` = $O(\log \log \lambda)$ such that $|A_{\text{cutoff}}| + \dots + |A_L| = O(\frac{N \log \log \lambda}{\log N})$. In this case, we show that `ObIMultArrShuff` uses $O(BN \log \log \lambda + B \frac{N \log N}{\log \lambda})$ bandwidth except with probability negligible in λ .

Note, `ObIMultArrShuff` performs an oblivious sort on $|A_{\text{cutoff}}| + \dots + |A_L|$ blocks in Step 2 of `ObIMultArrShuff`. Since $|A_{\text{cutoff}}| + \dots + |A_L| = O(\frac{N \log \log \lambda}{\log N})$, a total of $O(BN \log \log \lambda)$ bandwidth is required. Constructing `Assign` in Step 4 requires $O(BN \cdot \text{cutoff}) = O(BN \log \log \lambda)$. The remaining steps of `ObIMultArrShuff` require $O(N)$ bandwidth.

Let us now focus on the bandwidth of `ObIMultArrShuff.BinShuffle`. An oblivious sort is applied to at most $|\text{Bin}^{\text{in}}| + 4 \log^3 \lambda$ blocks in Steps 2 and 5. An oblivious sort is applied to at most

$|\text{Bin}^{\text{out}}|$ blocks in Steps 7 and 9. Step 3 requires $O(|\text{Bin}^{\text{out}}| \cdot \text{cutoff}) = O(|\text{Bin}^{\text{out}}| \cdot \log \log \lambda)$ blocks of bandwidth. Step 4 requires $O(|\text{Bin}^{\text{in}}| \cdot \text{cutoff}) = O(|\text{Bin}^{\text{in}}| \cdot \log \log \lambda)$ blocks of bandwidth. All other steps of `ObIMultArrShuff.BinShuffle` require $O(|\text{Bin}^{\text{out}}| + |\text{Bin}^{\text{in}}|)$ blocks of bandwidth. Altogether, the total required bandwidth of `ObIMultArrShuff.BinShuffle` is $O((|\text{Bin}^{\text{in}}| + \log^3 \lambda)(\log |\text{Bin}^{\text{in}}| + \log \log \lambda) + |\text{Bin}^{\text{out}}|(\log |\text{Bin}^{\text{out}}| + \log \log \lambda))$.

Finally, we consider `ObIMultArrShuff.Shuffle` now. If $|A_{\ell,1}| + \dots + |A_{\ell,L}| \leq \frac{N}{\log \lambda}$, then `ObIMultArrShuff.Shuffle` executes `ObIMultArrShuff.BinShuffle` at Step 5 requiring $O\left(\frac{N(\log N + \log \log \lambda)}{\log \lambda}\right)$ blocks of bandwidth. In the other case, `ObIMultArrShuff.BinShuffle` is executed \tilde{m} times. By Lemma 7, we know that, for all $j \in [\tilde{m}]$, $|\text{Bin}_j^{\text{in}}| = \sum_{i=1}^L k_{\ell,i,j}^{\text{in}} \leq (1 + 3\epsilon) \log^3 \lambda$. By Lemma 8, we know that, for all $j \in [m]$, $|\text{Bin}_j^{\text{out}}| = \sum_{i=1}^L k_{\ell,i,j}^{\text{out}} \leq (1 + \epsilon) \log^3 \lambda$. Therefore, each execution of `ObIMultArrShuff.BinShuffle` requires $O(\log^3 \lambda \log \log \lambda)$ blocks of bandwidth. All executions require $O((2\epsilon)^\ell N \log \log \lambda)$ blocks of bandwidth. The cost of all executions of `ObIMultArrShuff.Shuffle` is

$$\sum_{\ell \geq 0} O((2\epsilon)^\ell N \log \log \lambda) = O(N \log \log \lambda)$$

blocks of bandwidth, for large enough N , when $\epsilon < 1/4$ thus completing the proof.

4 Oblivious Hash Table

In this section we present our *oblivious hash table* construction which achieves bandwidth overhead of $O(\log N + \log \log \lambda)$ blocks, amortized per query. It uses as a building block the notion of an *oblivious bin*, which provides the same security properties as a general OHT but the main difference is that the oblivious bin structure will be used only for small inputs, which can be obviously shuffled without violating out overall efficiency requirements. The largest bins will be of size $\frac{N}{\log^c \lambda}$ for $c > 1$. We start with our oblivious bin constructions.

4.1 Oblivious Bin Tree

In this section we present two instantiations for the oblivious bin structure which will be used for inputs of size $O(\text{poly}(\log \lambda))$. The first instantiation uses tree structures and has higher access cost, logarithmic in the input size, but could be instantiated on any input size. The second construction uses Cuckoo hashes and has better constant time access, however, it requires input of size at least $\log^8 \lambda$ because of the best known oblivious initialization algorithm for a Cuckoo hash 3.

Construction 11 (Oblivious Bin Tree). Let F be a pseudorandom function and $(\text{Gen}, \text{Enc}, \text{Dec})$ be a symmetric key encryption. Let $\mathcal{U} = \mathcal{U}_{\text{real}} \cup \mathcal{U}_{\text{dummy}}$ where $\mathcal{U}_{\text{real}}$ and $\mathcal{U}_{\text{dummy}}$ are two non-intersecting sets representing the real and dummy key identifies. We define an oblivious hash construction `OblivBin = (OblivBinTree.Init, OblivBinTree.Build, OblivBinTree.Lookup, OblivBinTree.Extract)` as follows

- $(\text{st}, \tilde{D}) \leftarrow \text{OblivBinTree.Init} (D = (k_i, v_i)_{i=1}^N)$:
 1. Generate N dummy items $(k'_i, \perp)_{i=N+1}^{2N}$ where $k'_i \in \mathcal{U}_{\text{dummy}}$ and append them to D .

2. Generate an encryption key SK and set $\text{st} \leftarrow \text{SK}$.
 3. Set $\tilde{D} = \{\text{ct}_i = \text{Enc}(\text{SK}, (b_i, k_i, v_i))\}_{i=1}^{2N}$ where $b_i = 0$ if k_i is real and $b_i = 1$ if k_i is dummy. The real and dummy items are in any, not necessarily random, order.
- $(\text{st}, \tilde{H}_{\text{tree}}) \leftarrow \text{OblivBinTree.Build}(D = \{(b_i, k_i, v_i)\}_{i=1}^{2N})$:
 1. Generate a PRF key K and encryption key $\text{SK} \leftarrow \text{Gen}(1^\lambda)$, and set $\text{st} \leftarrow (\text{K}, \text{SK})$.
 2. Use an oblivious sort to sort the $2N$ items according to their tags $\text{F}(\text{K}, k_i)$.
 3. Set \tilde{H}_{tree} to be a binary search tree on top of the sorted array from the previous step.
 4. In this oblivious structure the additional array $\tilde{S} \leftarrow \perp$ is not used.
 - $(v, \tilde{H}'_{\text{tree}}) \leftarrow \text{OblivBinTree.Lookup}(i, \tilde{H}_{\text{tree}}, \text{st})$:
 1. Compute $d_i = \text{F}(\text{K}, i)$ and run a binary tree search to retrieve the node on a path of \tilde{H}'_{tree}
 2. Decrypt using SK the leaf node on the retrieved path to obtain (b_j, j, v_j) . If $i = j$, set $v \leftarrow v_j$ and write back a ciphertext in the leaf node $\text{Enc}(\text{SK}, (2, j, v_j))$. Otherwise, set $v \leftarrow \perp$ and write back $\text{Enc}(\text{SK}, (b_j, j, v_j))$.
 - $(\tilde{D}, \text{st}') \leftarrow \text{OblivBinTree.Extract}(\tilde{H}_{\text{tree}}, \text{st})$:
 1. Use an oblivious sort to sort all leaves of \tilde{H}_{tree} according their flag bit b_i to obtain an array \tilde{H} .
 2. Set \tilde{D} to be the first N items of \tilde{H} .
 3. Use an oblivious sort to permute \tilde{D} at random.

Theorem 12. The oblivious tree bin construction `OblivBin` presented in Construction 11 is an oblivious hash table according to Definition 4.

Proof. We construct a simulator Sim that outputs \tilde{H}_{tree} as follows. Sim_{Init} outputs $2N$ ciphertexts encrypting random values. $\text{Sim}_{\text{Build}}$ executes the honest oblivious build algorithm, which creates a binary search tree. $\text{Sim}_{\text{Lookup}}$ chooses a search path at random to access in the binary tree \tilde{H}_{tree} . Finally, $\text{Sim}_{\text{Extract}}$ runs the real extraction algorithm. The oblivious bin structure generated after the initialization and build phase in the real and the simulated execution are indistinguishable since all items in \tilde{H}_{tree} are encrypted. The view of the access pattern during lookups in the real execution reveals the relative ordering of the PRF values corresponding to the searched keys, which cannot reveal more than the actual PRF values. Since the query sequences do not contain any repeated entries a sequence of pseudorandom values is indistinguishable from a sequence of truly random value. Thus, the access patterns during the real and the ideal executions are also indistinguishable. Finally, $\text{Sim}_{\text{Extract}}$ runs the real extraction algorithm and the access patterns in the real and the simulated case are indistinguishable since they consist only of the access pattern from the oblivious sort. Hence, we conclude that the views of the adversary in the real and the ideal execution are indistinguishable. In the final step of `OblivBinTree.Extract`, \tilde{D} is permuted using an oblivious sort to a random permutation chosen independently and uniformly at random. Therefore, the final permutation given to the adversary in the real execution is indistinguishable from a random permutation chosen in the ideal execution. \square

Efficiency. The above bin construction, when instantiated on n blocks, has initialization build and extraction time with bandwidth of $O(n \log n)$ blocks. The bandwidth of a lookup is $O(\log n)$.

4.2 Oblivious Cuckoo Hash Bin

Construction 13 (Oblivious Cuckoo Hash Bin). Let F be a pseudorandom function and $(\text{Gen}, \text{Enc}, \text{Dec})$ be a symmetric key encryption. We define an oblivious hash construction $\text{CuckooBin} = (\text{CuckooBin.Init}, \text{CuckooBin.Build}, \text{CuckooBin.Lookup}, \text{CuckooBin.Extract})$ as follows

- $(\text{st}, \tilde{D}) \leftarrow \text{CuckooBin.Init} (D = (k_i, v_i)_{i=1}^N)$:
 1. Generate N dummy items $(k'_i, \perp)_{i=N+1}^{2N}$ where $k'_i \in \mathcal{U}_{\text{dummy}}$ and append them to D .
 2. Generate an encryption key SK and set $\text{st} \leftarrow \text{SK}$.
 3. Set $\tilde{D} = \{\text{ct}_i = \text{Enc}(\text{SK}, (b_i, k_i, v_i))\}_{i=1}^{2N}$ where $b_i = 0$ if k_i is real and $b_i = 1$ if k_i is dummy. The real and dummy items are in any, not necessarily random, order.
- $(\text{st}, \tilde{H}_{\text{cuckoo}}, \text{S}) \leftarrow \text{CuckooBin.Build} (D = \{(b_i, k_i, v_i)\}_{i=1}^{2N})$:
 1. If $2N < \log^8 \lambda$, then return $\text{OblivBinTree.Build}(D)$.
 2. Generate PRF keys K_1, K_2 and encryption key $\text{SK} \leftarrow \text{Gen}(1^\lambda)$, and set $\text{st} \leftarrow (K_1, K_2, \text{SK})$.
 3. Run an oblivious algorithm to construct a Cuckoo hash table that consists of two tables T_1, T_2 and a stash S with parameters K_1, K_2 , which contain the items of \tilde{D} . Set $\tilde{H}_{\text{tree}} \leftarrow (T_1, T_2, \text{S})$.
- $(v, \tilde{H}'_{\text{cuckoo}}) \leftarrow \text{CuckooBin.Lookup}(k_i, \tilde{H}_{\text{cuckoo}} = (T_1, T_2), \text{S}, \text{st})$:
 1. If $\tilde{H}_{\text{cuckoo}}$ is \tilde{H}_{tree} , then return $\text{OblivBinTree.Lookup}(k_i, \tilde{H}_{\text{tree}}, \text{st})$.
 2. Access all items in S as well as items $T_1[\text{F}_{K_1}(k_i)]$ and $T_2[\text{F}_{K_2}(k_i)]$, decrypting them using SK and checking whether the stored items matches the search index k_i . If the item with index k_i is found, set v to be the corresponding data, and otherwise set $v \leftarrow \perp$.
 3. If either $T_1[\text{F}_{K_1}(i)]$ or $T_2[\text{F}_{K_2}(i)]$ is of the form (b_i, k_i, v_i) write back an encryption of the value $\text{Enc}(\text{SK}, (2, k_i, v_i))$. Otherwise, write back a new encryption of $\text{Enc}(\text{SK}, (b_i, k_i, v_i))$. Set \tilde{H}'_{tree} to be the table with the resulting changes.
- $(\tilde{D}, \text{st}') \leftarrow \text{CuckooBin.Extract}(\tilde{H}_{\text{cuckoo}}, \text{S}, \text{st})$:
 1. If $\tilde{H}_{\text{cuckoo}}$ is \tilde{H}_{tree} , then return $\text{OblivBinTree.Extract}(\tilde{H}_{\text{tree}}, \text{st})$.
 2. Use an oblivious sort together all items in T_1, T_2 and S according their flag bit b_i to obtain an array \tilde{H} .
 3. Set \tilde{D} to be the first N items of \tilde{H} .
 4. Use an oblivious sort to permute \tilde{D} at random.

Theorem 14. The oblivious Cuckoo bin construction CuckooBin presented in Construction 13 is an oblivious hash table according to Definition 4.

Proof. Similarly to the oblivious tree bin, the simulator will initiate the Cuckoo hash table using a random input data and our Sim_{Init} will outputs $2N$ ciphertexts encrypting random values. We will use the result of Goodrich and Mitzenmacher [24], which was later given a complete proof in the work of Chan et al. [7], which says that there existing a simulator that can simulate the initialization of a Cuckoo table on $n = \Omega(\log^7 \lambda)$ elements given just the number of elements with a stash of size $O(\log \lambda)$. We will use this simulator as our $\text{Sim}_{\text{Build}}$. The $\text{Sim}_{\text{Lookup}}$ will just access two random locations in T_1 and T_2 . The indistinguishability of this simulated access from real lookups follows again from the result of [24]. Finally, $\text{Sim}_{\text{Extract}}$ just runs the regular extraction algorithm on the simulated oblivious Cuckoo table. The indistinguishability of the real world, where the adversary gets the real permutation, and the ideal world, where he obtains a random independent permutation, follows from the fact that the only leakage is the access pattern from the oblivious sort, which does not reveal any information about the permutation. At the end of CuckooBin.Extract , \tilde{D} is randomly permuted using an oblivious sort. The random permutation is chosen independently and uniformly at random. Therefore, the final permutation of \tilde{D} is indistinguishable from a uniformly at random chosen permutation. \square

Efficiency. The initialization, building and extraction algorithms for Cuckoo bin have bandwidth of $O(n \log n)$ blocks. The lookup time has instead bandwidth of $O(\log n + \log \log \lambda)$. While at first sight the Cuckoo bin does not seem to bring any efficiency advantage and it also needs much larger minimal size of the blocks in order to use Cuckoo hash tables, the majority of the lookup time $O(\log n)$ is spent on reading the stash. When we use several levels of Cuckoo hashes in our general oblivious hash table construction, we will have a more efficient way to deal with the stashes on multiple Cuckoo hash bins.

Construction 15. [Oblivious Hash Table] Let F be a pseudorandom function and $(\text{Gen}, \text{Enc}, \text{Dec})$ be a symmetric key encryption. Let $\mathcal{U} = \mathcal{U}_{\text{real}} \cup \mathcal{U}_{\text{dummy}}$ where $\mathcal{U}_{\text{real}}$ and $\mathcal{U}_{\text{dummy}}$ are two non-intersecting sets representing the real and dummy key identifiers. We define an oblivious hash construction $\text{OblivHT} = (\text{OblivHT.Init}, \text{OblivHT.Build}, \text{OblivHT.Lookup}, \text{OblivHT.Extract})$ as follows:

OblivHT.Init. This algorithm permutes the items of the input database and adds the necessary dummy items.

$(\text{st}, \tilde{D}) \leftarrow \text{OblivHT.Init}(D = \{(k_i, v_i)\}_{i=1}^N)$:

1. Generate encryption key $\text{SK} \leftarrow \text{Gen}(1^\lambda)$ and set $\text{st} \leftarrow \text{SK}$.
2. Generate real items of the form $\{\text{Enc}(\text{SK}, (0, k_i, v_i))\}_{i=1}^N$ and additional N dummy items of the form $\{\text{Enc}(\text{SK}, (1, k_i, \perp))\}_{i=N+1}^{2N}$.
3. Compute an oblivious shuffle on the above $2N$ real and dummy items and set \tilde{D} to be the result

OblivHT.Build. This algorithm builds an oblivious hash table from the output of OblivHT.Init .

$(\text{st}, \tilde{H}, \tilde{S}) \leftarrow \text{OblivHT.Build}(\tilde{D}, \text{st})$: the data \tilde{D} contains at most half of it real items and the rest are dummy items.

1. Return $(\text{st}, \tilde{H}, \tilde{S}) \leftarrow \text{OblivHT.BuildLevel}(\tilde{D}, N, N, 1, \text{st})$.

OblivHT.BuildLevel. This algorithm constructs one level in the OHT structure.

$(\text{st}, \tilde{H}, \tilde{S}) \leftarrow \text{OblivHT.BuildLevel}(\tilde{D}, R^{\text{real}}, R^{\text{dummy}}, \text{ctr}, \text{SK})$:

1. Let $R = R^{\text{real}} + R^{\text{dummy}}$.
2. If $R \leq \frac{N}{\log \lambda}$, set

$$(\text{st}^{\text{ctr}}, \tilde{H}^{\text{ctr}}, \tilde{S}^{\text{ctr}}) \leftarrow \text{OblivBin.Build}(\tilde{D}).$$

Return $(\text{st}^{\text{ctr}}, \text{d}), \tilde{H}^{\text{ctr}}, \tilde{S}^{\text{ctr}}$.

3. Otherwise, construct a level in the oblivious hash table as follows:

- (a) Let F be a PRF with output range $[\frac{R}{\log^c \lambda}]$. Generate a PRF key K_{ctr} .
- (b) Initialize $\frac{R}{\log^c \lambda}$ empty bins.
- (c) For each $i \in [R]$ append $\text{Enc}(\text{SK}, (b_i, i, v_i))$ to bin $F(K_{\text{ctr}}, i)$.
- (d) Let $\text{thrsh} = (1 - \epsilon)(\log^c \lambda)$.
- (e) Set $T = (1 - \delta)R^{\text{real}}$ and $t = \frac{R}{\log^c \lambda}$. For $1 \leq j \leq \frac{R}{\log^c \lambda}$ do as follows:
 - i. Sample from the binomial distribution $R_j \leftarrow \text{Binomial}(T, \frac{1}{t})^2$.
 - ii. If there are less than R_j real items B_j or the total number of items assigned to B_j were less than thrsh , abort.
 - iii. Linearly scan the items assigned to B_j in Step 3c, leaving the tags $b_i = 0$ to the first R_j of the real items and setting the tag to rest to be $b_i = 2$; the dummy items stay with tag $b_i = 1$.
 - iv. Obviously shuffle the items assigned to B_j according to their assigned tag. Move all items at the end of array starting from position $\text{thrsh} + 1$ to array D_{over} changing the tag of the real items back to $b_i = 0$.
 - v. Initialize a bin oblivious structure on the items left in B_j :

$$(\text{st}^{(\text{ctr},j)}, \tilde{H}^{(\text{ctr},j)}, \tilde{S}^{(\text{ctr},j)}) \leftarrow \text{OblivBin.Build}(B_j).$$

- vi. Append the real items from $\tilde{S}^{(\text{ctr},j)}$ to D_{over} .
- vii. Set $T = T - R_j$ and $t = t - 1$.
- (f) Let $\tilde{H}^{(\text{ctr})} \leftarrow \{\text{Enc}(\text{SK}, \text{st}^{(\text{ctr},j)}), \tilde{H}^{(\text{ctr},j)}\}_{j=1}^{R/\log^c \lambda}$.
- (g) Call recursively the level building functionality on the remaining items in D_{over} :

$$(\text{st}', \tilde{H}', \tilde{S}) \leftarrow \text{OblivHT.BuildLevel}(D_{\text{over}}, |D_{\text{over}}|, \text{ctr} + 1, \text{SK}).$$

- (h) Return $\tilde{H} \leftarrow (\tilde{H}^{(\text{ctr})}, \tilde{H}')$ and $\text{st} \leftarrow (\text{SK}, K_{\text{ctr}}, \text{st}')$

²This can be done in constant time using the Splitting algorithm for binomial random variates in Section 4.4., X4 [15].

OblivHT.Lookup. This algorithm retrieves an item stored in the OHT table.

$(v, \tilde{H}', \tilde{S}', \text{st}') \leftarrow \text{OblivHT.Lookup}(k, \tilde{H}, \tilde{S}, \text{st})$:

1. Let d recorded in st be the number of levels in the OHT.
2. Set $\text{found} = 0$, for $\text{ctr} = d$ to 1, do the following:
 - (a) If $\text{ctr} = d$, then do

$$(v', \tilde{H}', \tilde{S}') \leftarrow \text{OblivBin.Lookup}(k, \tilde{H}^{(\text{ctr})}, \tilde{S}^{(\text{ctr})}, \text{st}^{(\text{ctr})}),$$

If $v' \neq \perp$, set $v \leftarrow v'$ and $\text{found} = 1$.

- (b) If $\text{found} = 0$, set $j = F(K_{\text{ctr}}, k)$, else choose j at random among the bins at level ctr .
- (c) Run

$$(v', \tilde{H}', \tilde{S}') \leftarrow \text{OblivBin.Lookup}(k, \tilde{H}^{(\text{ctr}, j)}, \perp, \text{st}^{(\text{ctr}, j)}).$$

If $v' \neq \perp$, set $v \leftarrow v'$ and $\text{found} = 1$.

OblivHT.Extract. This algorithm returns a fixed size data array that contains only the unquired items in the OHT padded with dummy items.

$(\tilde{D}, \text{st}') \leftarrow \text{OblivHT.Extract}(\tilde{H}, \text{st})$:

1. Let d recorded in st be the number of levels in the OHT.
2. For $1 \leq \text{ctr} \in [d]$ and $j \in [\frac{\epsilon^{i-1}N}{\log^c \lambda}]$, let

$$(\tilde{D}_{\text{ctr}, j}, \text{st}') \leftarrow \text{OblivHT.Extract}(\tilde{H}^{(\text{ctr}, j)}, \tilde{S}, \text{st}^{(\text{ctr}, j)}),$$

for each bin B_j in level ctr , where $\tilde{S} = \perp$ for $\text{ctr} < d$ and otherwise, \tilde{S} is the auxiliary array stored in \tilde{H} . Append $\tilde{D}_{i, j}$ to \tilde{D} .

Efficiency. The initialization for the oblivious hash table is an oblivious sort, which we can do with bandwidth $O(N \log N)$. The building of the oblivious hash table from a shuffled array will be proportional to the the cost of running the build algorithm for each bin at each level plus on oblivious shuffle per bin. The size of each bin in each level except the last is $O(\log^c \lambda)$. The total number of bins of this size is

$$\sum_{i=1}^{\log \log \lambda - 1} \frac{1}{\log^c \lambda} \epsilon^{i-1} N = \frac{1}{\log^c \lambda} (1 - \epsilon^{\log \log \lambda}) N = O\left(\frac{N}{\log^c \lambda}\right).$$

Since the build algorithms for both tree and Cuckoo bin are asymptotically the same, the build for the oblivious has table will take:

$$O\left(\frac{N}{\log^c \lambda} (\log^c \lambda) (\log \log^c \lambda) + \frac{N}{\log \lambda} \log N\right) = O\left(N \log \log \lambda + \frac{N}{\log \lambda} \log N\right).$$

A lookup in the oblivious hash table consists of $O(\log \log \lambda)$ lookups on bins of size $\log^c \lambda$ plus the lookup on the single bin in the smallest level which is of size $N/(\log^c \lambda)$. In the case of binary tree bins the cost per bin lookup is logarithmic in the size of the bin, thus we obtain cost:

$$O(\log N + (\log \log \lambda)^2).$$

In the case when we instantiate the bin OHT with a binary tree bin, we can use any $c \geq 1$.

The cost of a lookup when we instantiate the bins with Cuckoo bins changes since a lookup in a Cuckoo bin without a stash has $O(\log \log \lambda)$ cost. The cost for the lookup on the single bin in the smallest level is $O(\log N + \log \log \lambda)$ since this level has logarithmic stash. Thus, the total lookup cost is

$$O(\log N + \log \log \lambda).$$

However, we need to have a restriction of $c > 7$ in order to satisfy the input size required by the oblivious algorithm constructing a Cuckoo hash table. Thus the OHT lookup can be performed at cost $O((\log \log \lambda)^2)$.

Security. Before stating the main lemma proving that the above construction is an oblivious hash table, we prove some lemmas that we will consequently use. First we argue what is the number of levels in our OHT construction.

Lemma 16. The depth of the oblivious hash table is $d = O(\log \log \lambda)$.

Proof. The size of level i is $\epsilon^{i-1}N$ and we recurse until we reach level of size $\frac{1}{\log \lambda}N$. Therefore it will take $O(\log \log \lambda)$ steps to reach this size. \square

Next we analyze the probability that the algorithm constructing the oblivious hash table fails and aborts.

Lemma 17. When $0 < \delta < \epsilon < 1$ and $c > 1$, algorithm `OblivHT.Build` in the above construction aborts with negligible probability in Step 3(e)ii.

Proof. Let X_j be random variable counting the total number of items assigned to bin B_j after Step 3c. We have that the expectation of X_j is $E[X_j] = \log^c \lambda$,

Using Chernoff bounds we estimate the probability that the total number of items assigned to B_j in Step 3c is less than $\text{thrsh} = (1 - \epsilon)(\log^2 R)$ value:

$$\Pr[X_j \leq \text{thrsh}] = \Pr[X_j \leq (1 - \epsilon) \log^c \lambda] \leq e^{-\frac{\epsilon^2 \log^c \lambda}{2}} = e^{-\omega(\log \lambda)} = \text{negl}(\lambda).$$

Next we estimate the probability that the number R_j sampled at Step 3(e)i is larger than the number of real items assigned to bin B_j during Step 3c. Let X_j^{real} and Y_j^{real} be the variables counting the number of real items assigned to bin B_j in Step 3c and Step 3(e)i. Then, we have $E[X_j^{\text{real}}] = \frac{R^{\text{real}}}{R} \log^2 \lambda$ and $E[Y_j^{\text{real}}] = (1 - \delta) \frac{R^{\text{real}}}{R} \log^2 \lambda$. We remind the reader that R^{real} is the number of real items among the total R items.

We first show that after d recursive calls of `OblivHT.BuildLevel` from within `OblivHT.Build` we can lower bound the fraction of real items as follows $\frac{R^{\text{real}}}{R} \geq \frac{1}{2} \left(\frac{\delta}{\epsilon}\right)^{d-1}$. This holds for the first invocation of `OblivHT.BuildLevel`. Let us assume that $\frac{R^{\text{real}}}{R} \geq \frac{1}{2} \left(\frac{\delta}{\epsilon}\right)^{d-1}$ for the input of a `OblivHT.BuildLevel`

call. We know that the size of D_{over} is fixed to ϵR . Additionally, the number of real items in D_{over} is at least δR^{real} since we distribute only $(1 - \delta)R^{\text{real}}$ items in Step 3(e)i, which will be the maximum number of real items being allocated to the current level of the hash table. Therefore, if $(R^{\text{real}})'$ and R' are the numbers of real items and the total number of items in D_{over} , then $1 \geq \frac{(R^{\text{real}})'}{R'} \geq \frac{\delta R^{\text{real}}}{\epsilon R} \geq \frac{1}{2} \left(\frac{\delta}{\epsilon}\right)^d \geq O\left(\left(\frac{\delta}{\epsilon}\right)^{\log \log \lambda}\right)$ using Lemma 16.

Using Chernoff bounds and the bound on the fraction of real items in each level, we estimate the following probabilities:

$$\Pr \left[X_j^{\text{real}} \leq (1 - \alpha) \frac{R^{\text{real}}}{R} \log^c \lambda \right] \leq e^{-\frac{\epsilon^2 R^{\text{real}} \log^c \lambda}{2}} = e^{-\omega(\log \lambda)} = \text{negl}(\lambda).$$

$$\Pr \left[Y_j^{\text{real}} \geq (1 + \alpha)(1 - \delta) \frac{R^{\text{real}}}{R} \log^c \lambda \right] \leq e^{-\frac{\epsilon^2 (1 - \delta) R^{\text{real}} \log^c R}{2}} = e^{-\omega(\log \lambda)} = \text{negl}(\lambda).$$

If $\delta \geq \frac{2\alpha}{1+\alpha}$ then $(1 - \alpha) \frac{R^{\text{real}}}{R} \log^2 R^{\text{real}} \geq (1 + \alpha)(1 - \delta) \frac{R^{\text{real}}}{R} \log^2 R^{\text{real}}$, hence $\Pr[X_j^{\text{real}} \leq Y_j^{\text{real}}] = \text{negl}(N)$. Since $0 \leq \frac{2\alpha}{1+\alpha} \leq 1$ for $0 \leq \alpha \leq 1$, we do not have any additional restrictions on the value of δ . \square

With the above lemmas, we are ready to state our main security theorem about the oblivious hash table construction.

Theorem 18. Construction 15 is an oblivious hash table according to Definition 4.

Proof. We construct a simulator $\text{Sim} = (\text{Sim}_{\text{Build}}(N), \text{Sim}_{\text{Lookup}}(), \text{Sim}_{\text{Extract}}(N))$ for the above construction as follows:

- $\text{Sim}_{\text{Build}}(N)$ runs the `OblivHT.Build` algorithm using random elements as real elements for the database, and using the simulator for the construction of the oblivious bins.
- $\text{Sim}_{\text{Lookup}}()$ chooses a random bin at each level for the oblivious hash table and runs $\text{Sim}_{\text{Lookup}}^{\text{B}}()$ for that bin.
- $\text{Sim}_{\text{Extract}}(N)$ run the extraction simulator for each bin with input the corresponding thrsh value for that level.

The access pattern in the above simulated distribution is indistinguishable from the real execution for the following reasons:

- The distributions of real items vs random values across bins is indistinguishable since the distributions of the corresponding PRF values are indistinguishable. Also the sampled distributions of real items are indistinguishable when generated from the same number real and dummy items.
- The simulated and the real lookups are indistinguishable as long as the protocol does not abort, which happens with all but negligible probability as we proved in Lemma 17. If we build protocol has completed without abort, the distributions of the real items in bin across each level is independent from the initial distributions of all items across all bins using the PRF and this distributions remains hidden after the completion of `OblivHT.Build` since we ensure that each bin contains fewer items than the threshold and we do not reveal how

many fewer. Also this distribution of real items across bins is indistinguishable from the distributions where the real items were assigned at random among the bins. The former is the distribution revealed from the real execution lookups, and the later is the distribution revealed by the simulated queries.

- The indistinguishability of the real and the ideal extraction follows from the indistinguishability of the bin extractors, which are executed on public values dependent only on the size of the database. We note that the extract procedure of the OHT returns in total N elements: the extract procedure in each level returns total of $(1 - \epsilon)R$ elements from the extractions across all bins, and the size of the $D_{\text{over}} = \epsilon R$ for the next recursive call. If the bin in the final level is initialized with R' items, then the extract algorithm for that bin will also return R' items. Thus, the total size of the data returned from extract is N elements.
- We will argue the items in the output of the extraction algorithm are shuffled under a permutation chosen uniformly at random. We can consider the decomposition of this permutation as defined in Lemma 22 of Appendix A. The unqueried items are distributed according to a PRF into different bins in `OblivHT.BuildLevel`. Since `OblivHT.Extract` only deals with unqueried items, the distribution of unqueried items corresponds to the unrevealed outputs of PRF evaluations with a secret key. Therefore, the distribution of unqueried items is indistinguishable from sampling a random `Assign` function. Finally, the bin extraction methods will apply the local permutations to unqueried items within each bin. Therefore, by Lemma 22, the final permutation of unqueried and dummy items after `OblivHT.Extract` is indistinguishable from a permutation chosen uniformly at random.

□

5 Oblivious RAM Construction

In this section we present our ORAM construction, which follows the hierarchical blueprint, and uses our oblivious hash table to store the data assigned to each level as well as our oblivious random multi-array shuffle to merge the content of consecutive levels when moving data from smaller to bigger levels in the ORAM.

Construction 19. Let `OblivHT` = (`OblivHT.Init`, `OblivHT.Build`, `OblivHT.Lookup`, `OblivHT.Extract`) be an oblivious hash table and `ObIMultArrShuff` be an oblivious random multi-array shuffle. Let $\mathcal{U} = \mathcal{U}_{\text{real}} \cup \mathcal{U}_{\text{dummy}} \cup \mathcal{U}_{\text{query}}$, where the items of the database come from $\mathcal{U}_{\text{real}}$, items used for padding in the construction come from $\mathcal{U}_{\text{dummy}}$ and dummy queries use values from $\mathcal{U}_{\text{query}}$. We construct an oblivious RAM scheme `ORAM` = (`ORAM.Init`, `ORAM.Access`) as follows:

`ORAM.Init`. This algorithm takes as input the database and outputs an ORAM memory structure initialized with the data.

$(\tilde{D}, \text{st}) \leftarrow \text{ORAM.Init}(1^\lambda, D)$:

1. Create a hierarchy of $\log N$ levels $\{\ell_i\}_{i=\alpha}^{\log N}$ of size 2^i where N is the size of D and $2^\alpha = O(\log N)$.

2. Initialize an oblivious hash table that contains the whole database by running

$$\begin{aligned} (\text{st}_{\text{Init}}, \tilde{D}') &\leftarrow \text{OblivHT.Init}(D) \\ (\text{st}_{\text{Build}}, \tilde{H}, \tilde{S}) &\leftarrow \text{OblivHT.Build}(\tilde{D}, \text{st}_{\text{Init}}). \end{aligned}$$

3. Set the first level to be the array representing the stash for the oblivious hash table $\ell_\alpha \leftarrow \tilde{S}$; the last level to be the new oblivious hash table $\ell_{\log N} \leftarrow (\text{st}_{\text{Build}}, \tilde{H})$; and all other levels to be empty $\{\ell_i \leftarrow (\perp, \perp)\}_{i=\alpha+1}^{\log N-1}$.

4. Set $\tilde{D} \leftarrow \{\ell_i\}_{i=\alpha}^{\log N}$. Output (\tilde{D}, st) where st contain the secret encryption key in st_{Build} .

ORAM.Access. This algorithm takes as input a database access instruction and executes it in an oblivious way having access to an ORAM memory structure.

$(v, \text{st}) \leftarrow \text{ORAM.Access}(\text{st}, \tilde{D}, \text{l}, \text{cnt}) :$

1. If $\text{cnt} = 2^j$ for some j , invoke $\text{ORAM.Shuffle}(\{\ell_i\}_{i=\alpha}^j)$ defined below.
2. Parse $\tilde{D} \leftarrow \{\ell_i\}_{i=k}^{\log N}$ where $\ell_i = (\text{st}_i, \tilde{H}_i)$ and $\text{l} = (\text{op}, \text{addr}, \text{data})$.
3. Set flag $\text{found} = 0$.
4. Do a linear scan on the items (k_j, v_j) stored in the array stored in the smallest level ℓ_α looking for item with $k = \text{addr}$. If such an item is found, set $\text{found} = 1$.
5. For $i = \alpha + 1$ to $\log N$, do the following:
 - (a) If $\text{found} = 1$, then sample a random query from $k \leftarrow \mathcal{U}_{\text{query}}$ ³. Otherwise, set $k \leftarrow \text{addr}$.
 - (b) Do a lookup in the oblivious hash table at that level:

$$(v, \tilde{H}'_i, \perp, \text{st}'_i) \leftarrow \text{OblivHT.Lookup}(k, \tilde{H}_i, \perp, \text{st}_i).$$

(c) Update $\ell_i \leftarrow (\text{st}'_i, \tilde{H}'_i)$

(d) If $\text{op} = \text{read}$, the set $d = \text{Enc}(\text{SK}, (\text{addr}, v))$, else if $\text{op} = \text{write}$, set $d = \text{Enc}(\text{SK}, (\text{addr}, \text{data}))$. Append d to the array stored in ℓ_α where SK is stored in st .

ORAM.Shuffle. This algorithm shuffles together the data from a number of consecutive levels in the ORAM hierarchical memory structure.

$\tilde{D} \leftarrow \text{ORAM.Shuffle}(\{\ell_i, \text{st}\}_{i=\alpha}^j):$

1. Append to the smallest level ℓ_α , $2^\alpha + 1$ dummy items from $\mathcal{U}_{\text{dummy}}$ and then apply an oblivious shuffle on the resulting set of real and dummy items and set \tilde{A}_α to be the output of the oblivious shuffle.

³We assume that dummy queries do not repeat. We can enforce this either by setting the universe $\mathcal{U}_{\text{query}}$ to be large enough or by keeping a counter for the dummy queries.

2. For $i = \alpha + 1$ to j . parse $\ell_i = (\text{st}_i, \tilde{H}_i)$

$$(\tilde{A}_i, \text{st}'_i) \leftarrow \text{OblivHT.Extract}(\tilde{H}_i, \text{st}_i).$$

3. Run the multi-array shuffle algorithm on the data extracted from each level oblivious hash table

$$\tilde{D} \leftarrow \text{OblMultArrShuff} \left(\{\tilde{A}_i\}_{i=\alpha}^j \right).$$

4. Construct an oblivious hash table using the output from the multi-array shuffle

$$(\text{st}_j, \tilde{H}_j, \tilde{S}_j) \leftarrow \text{OblivHT.Build} \left(\tilde{D}, \text{st} \right).$$

5. Set $\ell_j \leftarrow (\text{st}_j, \tilde{H}_j)$, $\ell_\alpha \leftarrow \tilde{S}_j$ and all intermediate levels to be empty $\{\ell_i \leftarrow \perp\}_{i=\alpha+1}^{j-1}$.

6. Return $\tilde{D} \leftarrow \{\ell_i\}_{i=\alpha}^{\log N}$.

Correctness. The `OblivHT.Build` algorithm expects input that has equal numbers of real and dummy items. Thus, we need to argue that the input array to the multi-array shuffle contain equal numbers of real and dummy value. The `OblivHT.Extract` functionality returns an array of size equal to the capacity for real items of the oblivious hash table. The size of \tilde{A}_α is $2^{\alpha+1} + 1$ and the size of each $\tilde{A}_{\alpha+i}$ for $1 \leq i \leq t-1$ is $2^{\alpha+i}$. Thus, total size of $\sum_{i=1}^t |\tilde{A}_{\alpha+i}|$ will be $2^{\alpha+1} + \sum_{i=\alpha+1}^{\alpha+t} 2^i = \sum_{i=1}^{\alpha+t} 2^i - \sum_{i=1}^{\alpha} 2^i = 2^{\alpha+1} + 2^{\alpha+t+1} - 2^{\alpha+1} = 2^{\alpha+t+1}$. At the same time the total number of real items across the first t levels is at most $2^{\alpha+t}$. Hence, the number of real items is at most half of the size of the array \tilde{D} from the multi-array shuffle, which is correct input for `OblivHT.Build`.

Theorem 20. Construction 19 is an access pattern hiding oblivious RAM scheme according to Definition 1.

Proof. We construct a simulator for the ORAM scheme as follows. For the initialization phase, `SimInit` just invokes the OHT simulators `SimInit` and `SimBuild`. The query simulator `SimAccess` uses the query simulators `SimLookup` for each oblivious hash table in each level of the ORAM. To simulate the shuffling procedures `SimAccess` uses the OHT `SimExtract` and then the multi-array shuffle simulator `SimMultArrShuff`. Note that the inputs for both the extraction and the shuffle simulators are deterministically depending only on the size of the database. The indistinguishability of the simulation from the real execution follows from the simulators that we used and the fact that their inputs are completely defined by the size of the database. \square

Efficiency. We evaluate the amortized query complexity over N ORAM accesses on database of size N , assuming that $N = \text{poly}(\lambda)$. After each 2^i accesses we reshuffle the top i levels using the oblivious random multi-array shuffle. Thus, we have the following total cost of the shuffle:

$$O \left(N' \log \log \lambda + \frac{N' \log N'}{\log \lambda} \right) = O(N \log \log \lambda), \text{ where } N' = \sum_{i=\log \alpha}^{\log N} 2^i,$$

which results in amortized cost of $O(\log \log \lambda)$ blocks per access from the shuffle. The cost for a look up across all levels consists of the scan on the smallest level which takes $O(\log N)$ plus the cost of the OHT without stash access for each level, which is $O(\log N(\log \log \lambda)^2)$ in the case when the OHT is instantiated with binary tree bins, and $O(\log N \log \log \lambda)$ when we use Cuckoo bins in the OHT.

Therefore, the amortized complexity of a lookup is $O(\log N(\log \log \lambda)^2)$ blocks or $O(\log N \log \log \lambda)$ blocks depending on the instantiation of the oblivious bin in the oblivious hash table as tree bin or Cuckoo bin.

5.1 Optimization Discussion

In our hierarchical ORAM construction each level has capacity to hold all preceding smaller levels and each level stores all blocks using a single oblivious hash table. Kushilevitz *et al.* [28] presented an optimization to hierarchical ORAM constructions where each level has capacity to hold K times the joint capacity of all smaller levels. Each level is implemented using K disjoint oblivious hash tables. When a level gets full, then the level's contents are moved to an empty oblivious hash table of the next larger level of the hierarchical ORAM. This optimization decreases the amortized shuffling cost by a factor of $O(\log K)$ at the cost of increasing the online cost of ORAM queries by a factor of $O(\frac{K}{\log K})$ and it is effective only when the amortized shuffling cost is larger than the online query costs (like for example in [25]).

For our construction, the two effects resulting from this modification are: first, an ORAM query must perform K oblivious hash table accesses per level and, secondly, the number of levels in the hierarchy decreases from $O(\log N)$ to $O(\frac{\log N}{\log K})$. The online bandwidth of an ORAM query becomes $O(\frac{K \log N \log \log \lambda}{\log K})$ or $O(\frac{K \log N (\log \log \lambda)^2}{\log K})$ blocks depending on the underlying oblivious bin scheme used. The amortized bandwidth of shuffling becomes $O(\frac{\log N \log \log \lambda}{\log K})$. We observe that the bandwidth is minimized when $K = O(1)$, which results in the same asymptotic overhead as the construction in Section 5. Therefore, our ORAM construction does not benefit from this optimization.

References

- [1] M. Ajtai. Oblivious rams without cryptographic assumptions. In *STOC '10*, pages 181–190, 2010.
- [2] M. Ajtai, J. Komlós, and E. Szemerédi. An $O(n \log n)$ sorting network. In *Proceedings of the fifteenth annual ACM symposium on Theory of computing*, pages 1–9. ACM, 1983.
- [3] K. E. Batcher. Sorting networks and their applications. In *Proceedings of the April 30–May 2, 1968, spring joint computer conference*, pages 307–314. ACM, 1968.
- [4] E. Boyle, K.-M. Chung, and R. Pass. Oblivious parallel ram and applications. In E. Kushilevitz and T. Malkin, editors, *Theory of Cryptography*, pages 175–204, Berlin, Heidelberg, 2016. Springer Berlin Heidelberg.
- [5] 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*, ITCS '16, pages 357–368, 2016.

- [6] 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, CCS '15*, 2015.
- [7] T. H. Chan, Y. Guo, W. Lin, and E. Shi. Oblivious hashing revisited, and applications to asymptotically efficient ORAM and OPRAM. In *Advances in Cryptology - ASIACRYPT 2017 - 23rd International Conference on the Theory and Applications of Cryptology and Information Security, Hong Kong, China, December 3-7, 2017, Proceedings, Part I*, pages 660–690, 2017.
- [8] T. H. Chan, Y. Guo, W.-K. Lin, and E. Shi. Cache-oblivious and data-oblivious sorting and applications. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2201–2220. SIAM, 2018.
- [9] T.-H. H. Chan, K.-M. Chung, and E. Shi. On the depth of oblivious parallel ram. In T. Takagi and T. Peyrin, editors, *Advances in Cryptology – ASIACRYPT 2017*, pages 567–597, Cham, 2017. Springer International Publishing.
- [10] T.-H. H. Chan and E. Shi. Circuit OPRAM: Unifying statistically and computationally secure orams and oprams. In *Theory of Cryptography Conference*, pages 72–107. Springer, 2017.
- [11] B. Chen, H. Lin, and S. Tessaro. Oblivious parallel ram: Improved efficiency and generic constructions. In E. Kushilevitz and T. Malkin, editors, *Theory of Cryptography*, pages 205–234, Berlin, Heidelberg, 2016. Springer Berlin Heidelberg.
- [12] K. Chung, Z. Liu, and R. Pass. Statistically-secure ORAM with $\tilde{O}(\log^2 n)$ overhead. In P. Sarkar and T. Iwata, editors, *ASIACRYPT 2014*, pages 62–81, 2014.
- [13] I. Damgård, S. Meldgaard, and J. B. Nielsen. Perfectly secure oblivious ram without random oracles. In *Proceedings of the 8th Conference on Theory of Cryptography, TCC'11*, pages 144–163, Berlin, Heidelberg, 2011. Springer-Verlag.
- [14] 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, 2015.
- [15] L. Devroye. *Non-Uniform Random Variate Generation*. Springer-Verlag, 1986.
- [16] J. Doerner and A. Shelat. Scaling oram for secure computation. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS '17*, pages 523–535, New York, NY, USA, 2017. ACM.
- [17] T. Elgamal. A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE Transactions on Information Theory*, 31(4), 1985.
- [18] C. Gentry. Fully homomorphic encryption using ideal lattices. In *Proceedings of the Forty-first Annual ACM Symposium on Theory of Computing, STOC '09*, 2009.
- [19] C. Gentry, K. A. Goldman, S. Halevi, C. Junta, M. Raykova, and D. Wichs. Optimizing oram and using it efficiently for secure computation. In E. De Cristofaro and M. Wright, editors, *Privacy Enhancing Technologies*, pages 1–18, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.

- [20] O. Goldreich. Towards a Theory of Software Protection and Simulation by Oblivious RAMs. In *STOC '87*, pages 182–194, 1987.
- [21] O. Goldreich and R. Ostrovsky. Software Protection and Simulation on Oblivious RAMs. *J. ACM*, 43(3), 1996.
- [22] M. T. Goodrich. Randomized shellsort: A simple oblivious sorting algorithm. In *Proceedings of the Twenty-first Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '10, pages 1262–1277, 2010.
- [23] M. T. Goodrich. Zig-zag sort: A simple deterministic data-oblivious sorting algorithm running in $o(n \log n)$ time. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 684–693. ACM, 2014.
- [24] M. T. Goodrich and M. Mitzenmacher. Privacy-preserving access of outsourced data via oblivious ram simulation. In *Proceedings of the 38th International Conference on Automata, Languages and Programming - Volume Part II*, ICALP'11, pages 576–587, Berlin, Heidelberg, 2011. Springer-Verlag.
- [25] M. T. Goodrich, M. Mitzenmacher, O. Ohrimenko, and R. Tamassia. Oblivious RAM Simulation with Efficient Worst-case Access Overhead. In *Proceedings of the 3rd ACM Workshop on Cloud Computing Security Workshop*, CCSW '11, pages 95–100, 2011.
- [26] S. D. Gordon, J. Katz, V. Kolesnikov, F. Krell, T. Malkin, M. Raykova, and Y. Vahlis. Secure two-party computation in sublinear (amortized) time. In *CCS'12*, pages 513–524, 2012.
- [27] M. S. Islam, M. Kuzu, and M. Kantarcioglu. Access pattern disclosure on searchable encryption: Ramification, attack and mitigation. In *in Network and Distributed System Security Symposium (NDSS)*, 2012.
- [28] 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*, SODA '12, pages 143–156, Philadelphia, PA, USA, 2012. Society for Industrial and Applied Mathematics.
- [29] T. Moataz, T. Mayberry, and E.-O. Blass. Constant communication ORAM with small block-size. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, pages 862–873. ACM, 2015.
- [30] O. Ohrimenko, M. T. Goodrich, R. Tamassia, and E. Upfal. The melbourne shuffle: Improving oblivious storage in the cloud. In *International Colloquium on Automata, Languages, and Programming*, pages 556–567. Springer, 2014.
- [31] R. Ostrovsky. Efficient computation on oblivious rams. In *STOC '90*, pages 514–523, 1990.
- [32] R. Ostrovsky and V. Shoup. Private information storage (extended abstract). In *Proceedings of the Twenty-ninth Annual ACM Symposium on Theory of Computing*, STOC '97, 1997.
- [33] P. Paillier. Public-key cryptosystems based on composite degree residuosity classes. In *Proceedings of the 17th International Conference on Theory and Application of Cryptographic Techniques*, EUROCRYPT'99, 1999.

- [34] S. Patel, G. Persiano, and K. Yeo. Cacheshuffle: An oblivious shuffle algorithm using caches. *CoRR*, abs/1705.07069, 2017.
- [35] B. Pinkas and T. Reinman. Oblivious ram revisited. In *Proceedings of the 30th Annual Conference on Advances in Cryptology, CRYPTO'10*, pages 502–519, Berlin, Heidelberg, 2010. Springer-Verlag.
- [36] L. Ren, C. Fletcher, A. Kwon, E. Stefanov, E. Shi, M. van Dijk, and S. Devadas. Constants count: Practical improvements to oblivious ram. In *USENIX Security 15*, pages 415–430, 2015.
- [37] E. Shi, T. H. H. Chan, E. Stefanov, and M. Li. Oblivious ram with $o(\log^3 n)$ worst-case cost. In D. H. Lee and X. Wang, editors, *ASIACRYPT 2011*, pages 197–214, 2011.
- [38] E. Stefanov, M. van Dijk, E. Shi, C. Fletcher, L. Ren, X. Yu, and S. Devadas. Path ORAM: An Extremely Simple Oblivious RAM Protocol. In *CCS '13*, pages 299–310, 2013.
- [39] A. Waksman. A permutation network. *Journal of the ACM (JACM)*, 15(1):159–163, 1968.
- [40] X. Wang, H. Chan, and E. Shi. Circuit ORAM: On Tightness of the Goldreich-Ostrovsky Lower Bound. In *Proceedings of the 22Nd ACM SIGSAC Conference on Computer and Communications Security, CCS '15*, pages 850–861, New York, NY, USA, 2015. ACM.
- [41] X. S. Wang, Y. Huang, T.-H. H. Chan, A. Shelat, and E. Shi. Scoram: Oblivious ram for secure computation. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security, CCS '14*, pages 191–202, New York, NY, USA, 2014. ACM.
- [42] S. Zahur, X. S. Wang, M. Raykova, A. Gascón, J. Doerner, D. Evans, and J. Katz. Revisiting square-root ORAM: efficient random access in multi-party computation. In *IEEE Symposium on Security and Privacy, SP 2016, San Jose, CA, USA, May 22-26, 2016*, pages 218–234, 2016.

A Permutation Decomposition

Oblivious shuffle algorithms are a central building block in many ORAM constructions. The oblivious properties of these algorithms are very strong and they guarantee that the output is shuffled under a random permutation that the adversary server does not learn anything about even when he knows in the input in the clear. However, in many cases, oblivious shuffles are used as intermediate steps of ORAM constructions when the server does not know the input for the shuffles in the clear. Often, parts of the input are still in a shuffled state from a previous step in the ORAM algorithm. This means that the input has more entropy to the adversary compared to the setting where the input is given in the clear by the server. In our constructions, we will leverage this input entropy in order to achieve oblivious properties for a random shuffle with better efficiency.

In the first decomposition, we assume that we are given a fixed partition of $[N]$ into A_1, \dots, A_L . To generate a random permutation, one can permute each of A_1, \dots, A_L according to randomly generated τ_1, \dots, τ_L and then randomly mix the elements of L into a single array. If the permutations of the partitions and mixing are performed obliviously, the result permutation is indistinguishable from a random permutation by any PPT adversary even if the adversary knows the initial partition A_1, \dots, A_L . This decomposition is used by our multi-array shuffle algorithm, which presents an

oblivious shuffle algorithm on input that consists of several smaller arrays each of which are shuffled according to a random permutation hidden from the adversary. We can think of this input as an intermediate state of the shuffle where the τ_i permutations have been already applied and the algorithm completes the second part of the permutation by obviously mixing the different arrays.

Our second decomposition instead simply receives a set of integers N_1, \dots, N_L that sum up to N . One can generate a random permutation by randomly partitioning $[N]$ into the sets A_1, \dots, A_L each with size N_1, \dots, N_L respectively, each of the sets are permuted using randomly chosen permutations τ_1, \dots, τ_L and the permuted arrays are, simply, concatenated. Our oblivious hash table constructions will mirror this decomposition. In the initialization of oblivious hash tables, items are distributed according to a random partitioning. Items remain in this state while the hash table is processing queries. Unqueried items remain randomly partitioned independent of the access patterns seen during the initialization and queries execution, and therefore remains hidden from the adversary. This intermediate state of the shuffle is given as input to the extract function, which applies τ_1, \dots, τ_L to each set completing the random permutation for unqueried items.

First Decomposition. Fix any N_1, \dots, N_L such that $N = N_1 + \dots + N_L$ and any partition of $[N]$ into sets $A_1 \cup \dots \cup A_L = [N]$. We consider the following decomposition of a permutation π into functions $(\text{Assign}, \{\tau_i\}_{i \in [L]})$, where $\text{Assign} : [N] \rightarrow [L]$ and $|\text{Assign}^{-1}(j)| = N_j$ for all $j \in [L]$ and $\tau_j : [N_j] \rightarrow [N_j]$ for $j \in [L]$. Assign is a random mixing of the L arrays where Assign explicitly states that $\pi(i) \in A_{\text{Assign}(i)}$. Let $\text{Assign}^{-1}(i)$ be ordered in an increasing order and let $\text{Assign}^{-1}(i)[j]$ be the j -th largest item. We define π as follows:

$$\pi(i) = \text{Assign}^{-1}(l)[\tau_l(k)] \text{ where } i \in A_l \text{ and } k = |\{j : j < i \text{ and } j \in A_l\}|.$$

In the above decomposition, Assign allocates exactly N_i locations for items of A_i , while τ_i describes the arrangement of the items of A_i within the N_i locations. In words, π assigns the k -th largest item of A_l to the $\tau_l(k)$ -th largest index of $\text{Assign}^{-1}(l)$. This process is well defined since $|\text{Assign}^{-1}(l)| = |A_l|$.

We argue in the next lemma that we can sample a uniform permutation π by sampling uniform assignment function, Assign and uniform permutations $\{\tau_i\}_{i \in [L]}$.

Lemma 21. For all $L > 0$, $N_1, \dots, N_L = N$ and a partitions $A_1 \cup \dots \cup A_L = [N]$ such that $|A_i| = N_i$ for all $i \in [L]$, the permutation output by the process above is uniformly distributed over the set of permutation over N elements.

Proof. Our proof will proceed in two steps. First, we show that the process has $N!$ possible choices of $(\text{Assign}, \tau_1, \dots, \tau_L)$ and any two such choices give different permutations as output.

The number of possible choices $(\text{Assign}, \tau_1, \dots, \tau_L)$ is easily seen to be

$$\binom{N}{N_1, \dots, N_L} \cdot \prod_{i \in [L]} N_i! = N!.$$

Let $(\text{Assign}, \tau_1, \dots, \tau_L) \neq (\text{Assign}', \tau'_1, \dots, \tau'_L)$ and let π and π' be the associated permutations. Suppose $\text{Assign}(i) \neq \text{Assign}'(i)$ for some $i \in [N]$. Then $\pi(i) \in A_{\text{Assign}(i)}$ while $\pi'(i) \in A_{\text{Assign}'(i)}$, which implies that π and π' are different. Suppose instead Assign and Assign' coincide on all inputs. Then, there exists $j \in [L]$ and $k \in [N_j]$ such that $\tau_j(k) \neq \tau'_j(k)$. In π , the k -th largest item in A_j will be placed into $\text{Assign}^{-1}(l)[\tau_j(k)]$ while the same item will be placed into $\text{Assign}^{-1}(l)[\tau'_j(k)]$ meaning π and π' are different. \square

Second Decomposition. We can also decompose permutations in another way. Fix any N_1, \dots, N_L such that $N_1 + \dots + N_L = N$. First, choose a random partitioning $\text{Distribute} : [N] \rightarrow [L]$ of the integers in $[N]$ into sets A_1, \dots, A_L of sizes N_1, \dots, N_L respectively. In other words, $i \in A_j$ if and only if $\text{Distribute}(i) = j$. Each set A_j is permuted according to a randomly selected permutation τ_j over $[N_j]$. Finally, the permutation π is obtained by concatenating of permuted A_1, \dots, A_L . Formally, π is defined as

$$\pi(i) = \sum_{1 \leq j < \text{Distribute}(i)} N_j + \tau_{\text{Distribute}(i)}(k) \text{ where } k = |\{j : j < i \text{ and } \text{Distribute}(i) = \text{Distribute}(j)\}|.$$

The next lemma proves that π is uniformly distributed over the set of permutation of $[N]$.

Lemma 22. For all $L > 0$, the permutation output by the above process is uniformly distributed over the set of permutations over N elements.

Proof. Again, the number of possible $(\text{Distribute}, \tau_1, \dots, \tau_j)$ is easily seen to be $N!$ following the same process as Lemma 21.

To show injectivity, pick $(\text{Distribute}, \tau_1, \dots, \tau_L) \neq (\text{Distribute}', \tau'_1, \dots, \tau'_L)$ and let π and π' be the associated permutations. Suppose that $\text{Distribute}(i) < \text{Distribute}'(i)$ without loss of generality. We know that $\pi(i) \leq N_1 + \dots + N_{\text{Distribute}(i)}$ while $\pi'(i) > N_1 + \dots + N_{\text{Distribute}'(i)}$ meaning that π and π' differ. On the other hand, suppose $\text{Distribute} = \text{Distribute}'$ and $\tau_j(k) \neq \tau'_j(k)$ for some $j \in [L]$ and $k \in [N_j]$. Denote A_1, \dots, A_L as the sets induced by Distribute as well as $\text{Distribute}'$. In this case, the k -th largest item of A_j is placed into different locations meaning π and π' differ. \square

B Related Work

Since its inception in [20], the concept of an Oblivious RAM has been object of intense study and several constructions with different properties have been proposed. As we have already pointed out above, our construction is in the most general model in which no assumption is made on the block size and the server is considered a storage device and not required to perform any computation on the blocks. The constructions of [21, 35, 24, 28] are all in the same general model as ours. Weakening any of these assumptions leads to more efficient, albeit less general, constructions. Specifically, Circuit ORAM [40] has overhead $O(\alpha(N) \cdot \log N)$ for blocks of size N^ϵ , for any $\alpha(N) = \omega(1)$ and constant ϵ . If the server is also considered to have some computational ability then the overhead can be reduced even further. Onion Oblivious RAM of [14] only requires a constant overhead provided that the blocks have size $\Omega(\log^6 N)$. This requirement was reduced to $\Omega(\log^4 N)$ by [29]

The logarithmic lower bound of Goldreich and Ostrovsky [21] only applies to statistically secure construction in balls and bins model, as pointed out by Boyle and Naor [5]. Specifically, this model only allows data to be shuffled around while not allowing sophisticated data encodings. Furthermore, Boyle and Naor show that without obtaining currently unknown superlinear lower bounds in sorting circuit sizes, there is no hope of getting lower bounds for Oblivious RAM in the general setting.

The client of a hierarchical Oblivious RAM typically uses a Pseudo-Random Function to remember where each block is stored and this is the main reason why they manage to guarantee only computational security. In the original construction of Goldreich and Ostrovsky [21], the need for a PRF can be removed, thus obtaining statistical security in the balls and bins model where blocks are modeled as indistinguishable blobs of data, provided that the client has access to a private random

function. Damgård *et al.* [13] and Ajtai [1] gave statistically secure constructions in the standard model with a overhead of $O(\sqrt{N} \cdot \log^2 N)$. The tree-based approach, pioneered by Shi *et al.* [37] instead does not need to remember where each block is stored; for example, Path ORAM [38] stores the position map recursively in smaller Oblivious RAMs. Thus, they offer statistical security with polylogarithmic overhead. Gentry *et al.* [19] improved the overhead to $O(\frac{\log^3 N}{\log \log N})$. Currently, the most efficient construction with statistical security is Path ORAM with $O(\log^2 N)$ overhead. The construction of Chung *et al.* [12] has a slightly worse overhead of $O(\log^2 N \cdot \log \log N)$ but a much simpler proof of security and performance. Further work by Ren *et al.* [36] improve the hidden constants of tree-based constructions.

Boyle *et al.* [4] introduced the notion of Oblivious Parallel RAM, which modeled the scenario of multiple processes accessing shared memory in parallel. Further works [11, 10, 7] have improved the overhead of Oblivious Parallel RAMs. Chan *et al.* [9] present lower bounds on the depth of Oblivious Parallel RAMs.

In recent years, the problem of using Oblivious RAM in multiparty computation has been heavily studied. Gordon *et al.* [26] showed that Oblivious RAM could be used to perform two-party computation in sublinear amortized time. Many works improving the asymptotics of Oblivious RAM for multiparty computation were presented [41, 40]. Further works have shown that considering ORAM constructions with worse asymptotics but smaller hidden constants can lead to better practical efficiency [42, 16].

Data-oblivious sorting is an important building block of many previous Oblivious RAM constructions. Batcher [3] presented an $O(N \log^2 N)$ -sized sorting circuit. Ajtai *et al.* [2] presented an $O(N \log N)$ -sized sorting circuit, however with very large constants. Goodrich [22] presented an $O(N \log N)$ randomized sorting algorithm with smaller constants, but larger depth. Afterwards, Goodrich [23] presented a deterministic $O(N \log N)$ sorting algorithm with similarly smaller constants, but large depth. Chan *et al.* [8] present a cache-oblivious sorting algorithm. Data-oblivious shuffling is another primitive regularly used in Oblivious RAMs. Waksman [39] presented an $O(N \log N)$ permutation network. Ohrimenko *et al.* [30] present an oblivious shuffle using $O(N \log_C N)$ bandwidth when the client has C blocks of storage available. Patel *et al.* [34] improve the hidden constants of the oblivious shuffle. In the same work, the notion of K -oblivious shuffling is introduced, which generalizes the knowledge of the adversary with respect to the input. In particular it assumes that out of the N blocks that are input for the shuffle the adversary knows the identities of only K . Leveraging the entropy of the remaining $N - K$ blocks the K -oblivious shuffle uses $2N + O(K \log_C K)$ bandwidth for any client with C blocks of storage available. Our multi-array shuffle uses ideas from the K -oblivious shuffle as starting point but achieves better asymptotic complexity and handles much more complex input entropy, which is crucial for its application to ORAM that achieves the new best asymptotic communication complexity.