Oblivious RAM with *Worst-Case* Logarithmic Overhead

Gilad Asharov[1], Ilan Komargodski[2], Wei-Kai Lin[3], and Elaine Shi[4]

[1] Bar-Ilan University
Gilad.Asharov@biu.ac.il
[2] Hebrew University of Jerusalem and NTT Research
ilank@cs.huji.ac.il
[3] Cornell University
wklin@cs.cornell.edu
[4] CMU
runting@gmail.com

Abstract. We present the first Oblivious RAM (ORAM) construction that for N memory blocks supports accesses with *worst-case* $O(\log N)$ overhead for any block size $\Omega(\log N)$ while requiring a client memory of only a constant number of memory blocks. We rely on the existence of one-way functions and guarantee computational security. Our result closes a long line of research on fundamental feasibility results for ORAM constructions as logarithmic overhead is necessary.

The previous best logarithmic overhead construction only guarantees it in an *amortized* sense, i.e., logarithmic overhead is achieved only for long enough access sequences, where some of the individual accesses incur $\Theta(N)$ overhead. The previously best ORAM in terms of *worst-case* overhead achieves $O(\log^2 N/\log\log N)$ overhead.

Technically, we design a novel de-amortization framework for modern ORAM constructions that use the "shuffled inputs" assumption. Our framework significantly departs from all previous de-amortization frameworks, originating from Ostrovsky and Shoup (STOC '97), that seem to be fundamentally too weak to be applied on modern ORAM constructions.

Keywords: Oblivious RAM · Worst-case overhead · Deamortization.

1 Introduction

Imagine a client that wishes to offload a database containing sensitive information to an untrusted server and later access the database and retrieve parts of it. By now, it is well-known that merely encrypting the entries of the database before uploading them to the server does not guarantee privacy (e.g., [6,21,22,41]). Indeed, the access patterns themselves may reveal non-trivial information about the underlying data or program being executed on the data. To mitigate these kinds of attacks, we would like to be able not only to encrypt the underlying data but also to "scramble" the observed access patterns so that they look unrelated

to the data. The algorithmic tool that achieves this goal is called an Oblivious RAM (ORAM).

An ORAM, introduced in the seminal work of Goldreich and Ostrovsky [16, 17], is a (probabilistic) RAM machine whose memory accesses do not reveal anything about the input—including both program and data—on which it is executed. An ORAM construction accomplishes this by permuting data blocks stored on the server and periodic reshuffling them around. Since their introduction more than 30 years ago, ORAMs have also become a central tool in designing various cryptographic systems, including cloud computing design, secure processor design, multi-party computation protocols, and more [4, 12, 13, 15, 26–29, 31, 34, 35, 38–40].

To be useful, ORAMs have to be "efficient". Whether an ORAM is efficient or not is typically measured by its (asymptotic) *overhead in bandwidth*: that is, how many data items must be accessed in the oblivious simulation as compared to the original non-oblivious implementation. There has been a tremendous effort in designing the most efficient ORAM construction possible [2, 7, 17, 18, 24, 30, 32, 33, 37]. The current record is the OptORAMa scheme by Asharov et al. [2] (building on Patel et al. [30]) who obtained an ORAM with amortized logarithmic overhead. Namely, their ORAM can simulate a RAM of size N so that over the span of T accesses, the total number of accesses would be $O(T \cdot \log N)$. The beautiful lower bound of Larsen and Nielsen [25] (see also [23]) shows that this is essentially the best possible: that is, every ORAM construction must spend *on average* $\Omega(\log N)$ physical accesses per one logical operation.[5]

Worst-case overhead. Much of the recent progress on ORAM constructions focuses on reducing its amortized cost [2, 30], whereas the worst-case overhead of an operation was ignored. Specifically, while achieving logarithmic amortized overhead, these constructions have $\Omega(N)$ worst-case overhead, due to the occasional reshuffling operations. This worst-case behavior renders these schemes much less useful in many applications since every now and then an access will "block" until $\Omega(N)$ physical accesses are complete which is clearly unacceptable.

The first to address this problem were Ostrovsky and Shoup [29] who showed how to spread the reshuffling operations over time, and achieve a worst-case $O(\log^3 N)$ overhead version of the original ORAM of Goldreich and Ostrovsky [17]. Related techniques were later applied on other ORAM schemes [7, 19, 24]. In spite of the recent great progress in ORAM constructions, the best known construction in terms of worst-case overhead is from almost a decade ago due to Kushilevitz, Lu, and Ostrovsky [24] who achieved $O(\log^2 N/ \log\log N)$ worst-case overhead (and their scheme was further clarified in the subsequent work of Chan et al. [7]). Crucially, the techniques of Ostrovsky and Shoup do

[5] The lower bounds of [23, 25] only apply to "online" ORAMs which support operations that come in an online fashion, one by one. These lower bounds even apply to computationally secure constructions. There is a logarithmic lower bound for "offline" ORAMs which see the whole set of operations ahead of time due to Goldreich and Ostrovsky [17], but it only applies to statistically secure constructions in the balls-and-bins model (see Boyle and Naor [5]).

not apply to the recent constructions that are based on "randomness reusing" of [2, 30], as we elaborate below in Section 2.

Thus, the current state of affairs leaves open the following fundamental question (also raised in [2]):

Is there a worst-case *logarithmic overhead ORAM? That is, is there an ORAM construction that can simulate* every *logical operation with $O(\log N)$ physical accesses?*

1.1 Our Contributions

Optimal worst-case overhead ORAM. We propose a new ORAM construction that achieve logarithmic worst-case overhead (in the memory size), while consuming $O(1)$ client-side storage, and $O(N)$ server-side storage. Here, N denoted the memory size. Obliviousness of our construction relies on the existence of one-way functions. Our result answers an important question left open by the recent OptORAMa work [2].

Theorem 1.1. *There is a computationally-secure ORAM with $O(\log N)$* worst-case *overhead assuming that one-way functions exist.*

The construction that achieves Theorem 1.1 is in the most standard model and make the same set of assumptions as all prior computationally-secure ORAM schemes. We assume a standard word-RAM where each memory word has at least $w = \log N$ bits, i.e., large enough to store its own logical address. We assume that word-level addition and boolean operations can be done in unit cost. We assume that the CPU has constant number of private registers. We additionally assume that a single evaluation of a PRF resulting in at least word-size number of pseudo-random bits, can be done in unit cost.

Technically, we significantly depart from all previous "deamortized" ORAM constructions. While all previous works[6] almost directly use the approach of Ostrovsky and Shoup [29], it seems like this approach is fundamentally too weak to be applied on modern ORAM constructions that achieve amortized logarithmic overhead [2, 30]. To this end, we build a new set of novel algorithmic tools from ground up and show how to amend the construction of Asharov et al. [2] so that it could be deamortized.

Linear-time oblivious deduplication. A noteworthy building block that we develop is an algorithm for efficient oblivious *deduplication*. Consider two sets of n elements A and B, where the goal is to obliviously compute $A \cup B$, that is, merge them into one larger set while removing duplicate elements. (Note that we assume that the elements within A are distinct, and ditto for B. Also, we assume that if there is a duplication, we keep the copy coming from A for concreteness.) Oblivious deduplication is a central building block in many previous worst-case

[6] Here we ignore tree-based constructions [32, 33, 37] since it is not known how to use them to get even amortized logarithmic overhead.

efficient ORAM constructions. Before this work, the only known solution was to apply a generic oblivious sort on the concatenation of A and B, followed by a linear scan which "deletes" the duplicates (after the sort, duplicates are adjacent in memory). Using the best known oblivious sort (e.g., AKS [1]), this approach would incur $O(n \cdot \log n)$ time. Eventually the extra log factor propagates into the overhead of the ORAM, resulting in $O(\log^2 N/\log\log N)$ worst-case overhead (at best).

We show that the extra $\log n$ overhead can be avoided by designing a *linear time* algorithm for this task which is oblivious if the input arrays are randomly shuffled. That is, we show that if A and B are shuffled with independent secret permutations, then there is a way to compute $A \cup B$ in time $O(n)$ while maintaining obliviousness. Linear overhead is clearly the best possible (since just reading the input takes linear time), matching the state of the art without obliviousness.

Theorem 1.2. *There is a linear time (probabilistic) algorithm that gets as input two sets A and B and outputs $A \cup B$. The algorithm is further (computationally) oblivious if A and B are independently secretly shuffled and if one-way functions exist.*

Conclusions and Open Problems We see our work as closing a long line of research on theoretical feasibility results for ORAM constructions. Our *worst-case logarithmic overhead ORAM* result, at least asymptotically, is optimal in terms of computational overhead. Unfortunately, the concrete constant, inherited from [2], underlying our construction is rather large. Using better generic building blocks (say the oblivious compaction algorithm of Dittmer and Ostrovsky [11]) we can get a much better constant but we believe that it is still too large for deployment. Whether there is a construction with (worst-case) logarithmic asymptotic overhead and a *small* constant is a major open problem.

Another exciting open problem is to bring down the cost of *statistically secure* ORAMs closer to $O(\log N)$ or prove that it is impossible. Specifically, the best statistically secure ORAMs has overhead $O\left(\log^2 N/\log\log N\right)$ [9] and it relies on the tree-based paradigm due to Shi et al. [32] which was later improved by [9, 10, 13, 33, 37]. Interestingly, tree based ORAMs have so far been more concretely efficient than hierarchical ones and this question is also somewhat related to the previous one.

Lastly, we mention a recent work of Asharov et al. [3] who gives an optimal Oblivious *PRAM* (OPRAM). An OPRAM is an extension of ORAM to the parallel setting where several processors make concurrent accesses to a shared memory. Their main result is that (assuming one-way functions) any PRAM with memory capacity N can be obliviously simulated in space $O(N)$, incurring only *amortized* $O(\log N)$ overhead in work and (worst-case) $O(\log N)$ overhead in depth. We believe that our techniques can be further extended to obtain a *worst-case* logarithmic work and depth overhead OPRAM, but this is left for future work.

2 Technical Overview

2.1 Background: Underlying ORAM Without Deamortization

The hierarchical paradigm and oblivious hash tables. We will build from an underlying (amortized) ORAM scheme which follows the hierarchical paradigm established by Goldreich and Ostrovsky [16, 17]. An ORAM scheme in the hierarchical paradigm can be viewed as a technique to reduce the task of constructing ORAM to the task of constructing an oblivious hash table. Specifically, A hierarchical ORAM typically consists of $\log_2 N + 1$ levels numbered $0, 1, \ldots, n$. Each level i is an *oblivious hash table* that can contain at most 2^i elements. An oblivious hash table is a data structure that supports the following operations:

- Build takes an input array containing (key, value) pairs and creates the data structure (we also say a pair is an element, a block, or an item);
- Lookup receives a key k, and returns the value corresponding to the key k contained in the data structure, or returns $\perp$ if not found or if the key looked up is dummy (denoted $\perp$).
- Extract is called when the data structure is destructed, and returns a list of unvisited items in the data structure.

Almost all known ORAM schemes in the hierarchical paradigm guarantee the following *non-recurrent* invariant: for each oblivious hash table in the hierarchy, the same real (i.e., non-dummy) key must be looked up at most once during the life-cycle of the data structure. Therefore, the data structure only needs to provide obliviousness if this non-recurrent assumption is respected. Finally, an oblivious hash table often has an access budget in the sense that it can only support up to an a-priori fixed number of lookup requests. Typically this budget is at least n which is the size of the array input into Build.

Achieving amortized logarithmic overhead. The original oblivious hash table implementation suggested by Goldreich and Ostrovsky [16, 17] is slow and takes $O(n \log n)$ time to build for an input array of size n. This would result in a non-optimal ORAM scheme. Instead, we adopt the efficient oblivious hash table suggested by Asharov et al. [2] (which is built upon Patel et al. [30]). Asharov et al. showed an oblivious hash table with $O(n)$ build time and $O(1)$ lookup overhead, except with the following input assumptions, output requirement, and caveat:

- *Randomly shuffled requirement:* the input array of Build must be randomly shuffled; moreover, the Extract function outputs the unvisited blocks in a random order. In either case, the randomness is hidden from the adversary.
- *Size assumption:* to obtain negligible in λ failure probability, the construction only works for hash tables that are at least $\mathsf{poly}\log(\lambda)$ in size.
- *Stash:* while the main hash table data structure can indeed be looked up in constant time, the hash table construction actually comes with a stash, and sometimes the element to be looked up actually resides in the stash. Each stash has expected constant size but with noticeable probability, the size can be as large as $O(\log \lambda)$.

We briefly overview how past works [2, 30] deal with the above imperfectness to make the ORAM scheme work. The recent work by Patel et al. [30] and Asharov et al. [2] show that by relying on the "residual randomness", ORAM constructions can respect the randomly shuffled input assumption. Specifically, when each oblivious hash table is destructed in the ORAM, the unvisited elements in the hash table appear in a random order, and the random permutation is hidden from the adversary. The size assumption can be dealt with by using yet another designated, slower, data-structure for smaller levels in the ORAM that are less than $\mathsf{poly}\log(\lambda)$ in size. Finally, the stash issue can be dealt with by merging the stashes of all oblivious hash tables into a single one, and accessing the merged stash once and for all for each ORAM request — one can prove that the merged stash is at most poly logarithmic in size except with negligible probability, and is stored in a designated data structure to allow fast lookup (i.e., taking strictly logarithmic time).

Simplifying assumptions. For ease of understanding, let us first ignore the size assumption and the stash issue mentioned above— these introduce additional technicalities for constructing an optimal, deamortized ORAM as we shall mention shortly. For the time being, we pretend that we can indeed have an oblivious hash table for randomly shuffled inputs can be built in linear time, regardless of the size, and moreover, assuming that lookup need not deal with the stash technicality.

Underlying ORAM scheme without deamortization. With these simplifying assumptions, we can construct an ORAM scheme as follows—our description below matches the rebuild description of Asharov et al. [3] in their optimal OPRAM scheme, and is a variant of Goldreich and Ostrovsky's original hierarchical construction.

Assume that the total memory size N is a power of 2. Imagine that there are $\log_2 N + 1$ levels, where the i-th level is an oblivious hash table (for randomly shuffled inputs) of capacity 2^i. We use $\mathsf{T}_0, \ldots, \mathsf{T}_L$ to denote all the $L + 1$ levels where $L = \log_2 N$.

In the steady state of the ORAM (i.e., ignoring the initial time steps when the levels are not yet populated), every level except level 0 is either *half full* (HF) or *full* (F). A level $i > 0$ is *half full* iff it contains up to 2^{i-1} real blocks. A level i is *full* iff it may contain up to 2^i real blocks[7]. Level 0 is either empty or full, and as we shall see, it is guaranteed to be empty at the end of every ORAM request. Whenever a new ORAM request arrives asking for the block at logical address addr, we do the following:

- *Fetch phase.* From $i = 0$ to L, we look up each oblivious hash table for the logical address addr; once the block is found, for all subsequent levels, we instead look for a dummy block.

[7] The actual number of real blocks may be smaller if the requests keep asking for the same block or a small set of blocks. The maximum load is achieved when the ORAM requests cycle through addresses $1, 2, \ldots, N$ in a round-robin fashion.

– *Maintain phase.* The block at addr just fetched is updated if necessary, and then it is entered into the smallest level (i.e., level 0), which makes the smallest level full. At this moment, let ℓ be either the smallest level that is half full or $\ell = L$ if all levels are full. Regardless of which case, all levels $0, 1, \ldots, \ell - 1$ are full. We now perform the following rebuild procedure:
  1. For each level $i = 0$ to $\ell - 2$ in parallel:
     let $\mathsf{T}'_{i+1} := \mathsf{Build}(\mathsf{Intersperse}(\mathsf{T}_i.\mathsf{Extract}(), D_i))$ — where D_i is an array of size 2^i containing only dummy elements, and $\mathsf{Intersperse}$ merges two randomly shuffled arrays into a randomly shuffled array.
  2. Let $\mathsf{T}'_\ell := \mathsf{Build}(\mathsf{T}_{\ell-1}.\mathsf{Extract}() \cup \mathsf{T}_\ell.\mathsf{Extract}())$ while we also remove dummy elements when unifying the two arrays;
  3. Replace $\mathsf{T}_1, \ldots, \mathsf{T}_\ell$ with the new hash tables $\mathsf{T}'_1, \ldots, \mathsf{T}'_\ell$ and let T_0 be emptied.

  After this rebuild procedure, $\mathsf{T}_0, \ldots, \mathsf{T}_{\ell-1}$ are all half full (and in fact T_0 is empty), and T_ℓ becomes full. In the above Steps 1 and 2, one can also imagine that each level $1, \ldots, \ell - 1$ is "*rebuilding itself down*" into the next level (and we will use this terminology later). For ease of understanding, the maintain phase is depicted in Figure 1.

Fact 2.1. *In the above construction, a level $i \geq 1$ switches state (either from half full to full, or vice versa) every 2^{i-1} requests. Similarly, a level i wants to rebuild itself down every 2^i requests — this also coincides with when level $i + 1$ refreshes.*

Assuming that the oblivious hash table supports Build in linear-time (for randomly shuffled inputs), supports Lookup in constant time, and moreover, its $\mathsf{Extract}$ function outputs unvisited blocks in a random order, then the above ORAM scheme is secure and achieves $O(\log N)$ *amortized* overhead.

2.2 Why Existing Deamortization Techniques Fail

Clearly the scheme mentioned in Section 2.1 cannot give a worst-case efficient ORAM. For instance, when rebuilding the largest level (which happens every N accesses), just reading it requires $O(N)$ work. We describe existing deamortization techniques and explain why they are incompatible with the new generation of amortized optimal ORAM constructions.

Ostrovsky and Shoup [29] proposed a deamortization technique that, roughly speaking, "spread" the rebuild procedures over many accesses rather than performing them atomically. The challenge is how to support accesses to the level while it is being rebuilt.

To do this, first, we modify the access process so that it does *not* delete an element once it is found in some level (while it is still reinserted into the smallest level). With this change, it is immediate that except for the smallest level, the only time the contents of a level changes, is during the rebuild procedure and in the latter we always just "pull" content from its previous (i.e., smaller) level. That is, it takes 2^i accesses until the content of the ith table is modified, and

it will be filled with the content of table $i-1$ which is also fixed and known. Thus, the rebuild process for a level can be done slowly and in advance across many accesses as follows: each level has a table called CurrentActive and has another table that is UnderConstruction, which is being slowly built from CurrentActive and the previous level CurrentActive. When the construction of the level is complete, we just update the pointer of CurrentActive to the new rebuilt table, and start a new UnderConstruction version, thus doing the rebuild in the deamortized sense, by spreading the cost uniformly.

There is one very important technical detail with the above approach: since we never actually delete elements, the same key may (and will) appear multiple times in the structure, perhaps with different values. The important invariant is that the *newest* version of the element is always the one that resides in the smaller level. At some point, these copies will meet in the same level during some rebuild process and then the older copy will be suppressed and discarded. The later task is known as oblivious **deduplication** and it is usually implemented using oblivious sort.

Multiple variants and adaptations of the above deamortization technique were used in previous ORAM constructions [7,19,24]. However, as we shall argue next, there are inherent obstacles one runs into while trying to apply it on the more recent (amortized) logarithmic overhead ORAM constructions [2,30].

Challenge I: Access-while-rebuild breaks security. Recall that in the deamortization technique of Ostrovsky and Shoup, we start rebuilding the table into the next level while we also access it at the same time. However, when applied on the ORAM of [2,30], this completely breaks the input assumption (and therefore security) and is not compatible with the "residual randomness" technique: A lookup that is performed while building reveals the position of an element in the new table—the security of the latter inherently relies on the permutation being completely secret.

To elaborate further, in Ostrovsky and Shoup we know in advance the content of the levels and we start the rebuild process ahead of time, while still allowing accesses to the same levels. In the context of in Ostrovsky and Shoup, this is *secure* as we re-randomize the levels (by obliviously sorting/shuffling it) during the rebuild process. However, in our case we cannot re-randomize the levels as this is too expensive, and we follow the residual randomness technique. Moreover, we cannot know in advance which elements will be looked up in the previous level, i.e., the residual randomness will be consumed and thus for security we are not allowed to reuse randomness.

Challenge II: Deduplication takes quasi-linear time. As mentioned, multiple copies of the same key will appear during the lifetime of the ORAM (some might be with different values) and so a deduplication mechanism is needed. More precisely, the task is to compute $A \cup B$ given two sets A and B of size n. (Assume we prefer the copy in A over B for concreteness.) The best known algorithm uses oblivious sort and takes quasi-linear time. However, if we want to have a logarithmic overhead ORAM construction we must implement it in *linear*

time. Note that, even ignoring obliviousness, it is not immediately obvious how to do this in linear time. The most natural approach is to use sorting, but better algorithms can be achieved using hashing tools.

Challenge III: Cannot "lock" shared memory. As mentioned, each level in [2,30] *effectively* supports lookup in constant time, but this is due to the "shared stash" technique (previously used in [9, 18, 20, 24]). Specifically, in isolation, each level requires $O(1)$ accesses to a main table and an additional scan of a $O(\log N)$-size stash. The shared stash trick utilizes the fact that there are many levels and an access will translate to multiple lookups for the same key in many levels—this allows us to merge all of the stashes into a global one and scan it only once for all levels. Therefore, the number of accesses per level is $O(1)$ *amortized*. This makes the ORAM construction not completely black-box in a hash table and therefore less compatible with Ostrovsky-Shoup [29]. Concretely, while we rebuild a table, we *cannot* "lock" the global stash as we need to allow accesses to it (addition and removal of elements) while handling other ORAM accesses.

2.3 Our Deamortization Approach

We now intuitively describe how to deamortize the ORAM scheme mentioned in Section 2.1. To address Challenge I mentioned in Section 2.2, whenever some level is involved in a rebuild, i.e., its destructor $\mathsf{Extract}$ function is being called, this level no longer can support lookups. Yet the ORAM must continue to serve requests. We propose a new pipelining approach that works with this new constraint that comes with the new amortized logarithmic-overhead ORAMs.

Our key idea is to maintain two copies of each level, henceforth called the A-copy and the B-copy respectively, each uses its own *independent* randomness. At a high level, whenever some level in the A-copy involved in a rebuild and being destructed, the *corresponding* level in the B-copy fills in the lookups; and whenever some level in the B-copy is involved in a rebuilt and being destructed, the *next level* in the A-copy can fill in. This guarantees that we never lookup and rebuild from a table at the same time – while rebuilding all lookups are performed at a different copy that uses independent randomness. That is, we essentially split the queries between A and B, so that we can reuse randomness in each hierarchy by itself. Whenever an element is found in one copy, we update its content and put it in the top of the two hierarchies of A and B.

We next elaborate our idea in the following schedule of deamortization. Henceforth, we use A_i and B_i to denote the i-th level in the A-copy and B-copy, respectively. In our earlier non-deamortized scheme, recall that a level $i \geq 1$ is reconstructed every 2^{i-1} requests. We may imagine that there is some counter ctr that increments upon every request, and thus a level $i \geq 1$ is refreshed whenever the $\mathsf{ctr} = k \cdot 2^{i-1}$ for some integer k. For simplicity, in this overview we ignore the treatment of the first level and the last level, and just look at intermediate levels. To achieve deamortization, the idea is to rely on careful pipelining to maintain the following schedule:

- The level A_i is refreshed (i.e., completes reconstruction) at time step $\mathsf{ctr}_1 = k \cdot 2^i + 2^{i-2}$ where k is an integer. Level A_i is now "half full" and just pulled new content from level $i-1$.
- The level B_i is refreshed at time $\mathsf{ctr}_2 = k \cdot 2^i + 2 \cdot 2^{i-2}$, i.e., just a little later than the corresponding A_i finished reconstruction.
- At time $\mathsf{ctr}_2 = k \cdot 2^i + 2 \cdot 2^{i-2}$, level A_i starts pulling more content from A_{i-1}. This will take 2^{i-2} time, which will finish at $\mathsf{ctr}_3 = k \cdot 2^i + 3 \cdot 2^{i-2}$.
  During this rebuild period $[\mathsf{ctr}_2, \mathsf{ctr}_3]$, A_i is dysfunctional and cannot support accesses; fortunately, B_i contains identical contents as A_i (but rerandomized differently for obliviousness), and therefore B_i can fill in for the lookups.
- At time ctr_3, A_i finishes, and B_i can catch up and pull information from level B_{i-1}. It will start rebuilding and will finish at time $\mathsf{ctr}_4 := \mathsf{ctr}_3 + 2^{i-2}$. At this time range B_i is dysfunctional, however, its content is also in A_i, which also hold in addition some fresher content from A_{i-1}.
- After B_i finishes and becomes full, A_i wants to push all its content down into level A_{i+1} while also pulling new content from level A_{i-1}. However, the rebuild of level A_{i+1} takes twice as much time. Thus, level A_i will have new content already at time $\mathsf{ctr}_5 := \mathsf{ctr}_4 + 2^{i-2}$, while level A_{i+1} will have the old content of A_i only at time $\mathsf{ctr}_6 := \mathsf{ctr}_5 + 2^{i-2}$. Luckily, the content is not lost; The content exists all this time at the B_i copy, which is activated and functional in $[\mathsf{ctr}_4, \mathsf{ctr}_6]$.
  At time ctr_5, we essentially finished a cycle, i.e., A_i is half full and just finished refreshing, and we are essentially back to the first item.

Recall that rebuilding A_i (or B_i, resp.) from A_{i-1} (or B_{i-1}, resp.) takes work $\Theta(2^i)$. This amount of work is spread across $\mathsf{ctr}_2 - \mathsf{ctr}_1$ (or $\mathsf{ctr}_3 - \mathsf{ctr}_2$, resp.) time steps. One can verify that indeed, $\mathsf{ctr}_2 - \mathsf{ctr}_1 = \Theta(2^i)$ and thus in each time step, a constant amount of work is performed associated with the rebuilding of A_i (or B_i, resp.). At most $O(\log N)$ many levels are being rebuilt simultaneously, and thus the total amount of work per time step associated with rebuilds is $O(\log N)$. Observe also that the time in between two adjacent rebuilds of A_i is only 2^{i-1}, and the hash table A_i can support up to 2^i requests, so we will not run out of the access budget. A similar observation holds for B_i.

In our final scheme, we will actually have different designated place for the table A_i when it is half-full, and a different placed for A_i when it is full, denoted as $\mathsf{A}_i^{\mathsf{HF}}$ and $\mathsf{A}_i^{\mathsf{F}}$, respectively (likewise for $\mathsf{B}_i^{\mathsf{HF}}$ and $\mathsf{B}_i^{\mathsf{F}}$). Thus, we have 4 tables at the same level, but (the newest version of) each element has exactly 1 copy in each of A_i and B_i. Moreover, only one copy of $(\mathsf{A}_i^{\mathsf{HF}}, \mathsf{A}_i^{\mathsf{F}})$ is valid at any given time, likewise for B_i. We can view $\mathsf{A}_i^{\mathsf{HF}}$ and $\mathsf{A}_i^{\mathsf{F}}$ as two possible states of A_i, and then we have just two copies in each level, or as independent hash tables, and then we have four tables. The rebuild schedule is depicted in Figure 3.

Why deduplication is necessary in the deamortized scheme. The rebuild procedure of the deamortized scheme is otherwise the same as the underlying non-deamortized scheme except that a deduplication pre-processing step is necessary whenever we are building a half full level A_i and a full level A_{i-1} to create a

new level A_i (and the same for the B-copy). This is because in the deamortized scheme, the rebuilding process is happening while the ORAM is still serving queries. For example, A_i may be rebuilding itself down into A_{i+1}, and the corresponding B_i is taking over as A_i in serving queries. During the rebuild, a memory block at address addr^* may get rebuilt from A_i, into A_{i+1}, but it can also be requested during the same rebuild interim, causing a separate copy of addr^* to be entered into the smallest level. Duplicates will be suppressed when two duplicate copies of the same addr^* "meet" in a future rebuild, while lookup guarantees that the freshest copy (which resides in the level that is smallest compared to any other copy) will be found.

We next describe how to accomplish oblivious deduplication in linear time.

2.4 Linear-Time Oblivious Deduplication

To address Challenge II from Section 2.2, another contribution of our paper is a linear-time oblivious deduplication algorithm that takes advantage of the fact that the input arrays to be deduplicated are randomly shuffled. To see the main idea, let us first describe a *non-oblivious* algorithm that runs in linear time. (This is already not trivial.) While the most natural implementation is to sort the concatenated array and then perform a linear scan while removing duplicates (which reside in adjacent positions after the sort), the cost of this implementation is dominated by the sort. Using the best sorting algorithms in the RAM model, one can achieve $o(n \cdot \log n)$ cost [14,36] but we still do not know of linear time sorting algorithms.

Our idea to get a linear time algorithm is to use hashing tools. Specifically, consider a Cuckoo hash—this is a hash table that supports lookup by just 2 accesses to a main table and another scan of say $O(\log \lambda)$ size stash. (The stash needs to be of this size to guarantee the probability of failure is negligible in λ.) Consider hashing X_1 and X_2 independently into a Cuckoo hash tables T_1, T_2 each having a long enough stash so that the hashing succeeds with all but negligible probability. For $i \in \{1, 2\}$, denote $T_i = (T_i^M, T_i^S)$ where T_i^M is the main table of T_i and T_i^S is the stash of T_i. This costs just $O(n)$. Now, we can perform the deduplication as follows:

1. For each element in the stash of T_1, i.e., T_1^S, we perform a full lookup in T_2. That is, for each of the $O(\log \lambda)$ elements in T_1^S, we touch $O(1)$ elements in T_2^M and then scan T_2^S. Doing so, we remove the elements from T_2 that are also in T_1^S. This costs overall $O(\log^2 \lambda)$.
2. So far, we took care of elements that appear in T_2 and in T_1^S and we need to remove duplicates which appear in T_2 and T_1^M. For this, we scan every remaining element in T_2 and for each we touch the two locations in T_1^M to check if it is there. If so, we delete it from T_2. Crucially, now we do not need to visit the stash T_1^S. When we look for the elements in $T_2 \setminus T_1^S$, we can go directly to the main table and spend $O(1)$ work per lookup, as we know that the element is not in the stash T_1^S!

Eventually, we concatenate the elements of T_1 together with what is left in T_2 after performing the above process. Overall, the cost of this algorithm is $O(n + \log^2 \lambda) = O(n)$ whenever n is large enough compared to λ.

This algorithm is clearly non-oblivious and "naively" replacing the Cuckoo hash with an oblivious version thereof does not meet our goal since known constructions require $\omega(n)$ time for building (since building an oblivious Cuckoo hash uses oblivious sort). This is where the "shuffled inputs" assumptions comes into play: we do not necessarily need the full power of Cuckoo hashing since we are guaranteed that the input lists are shuffled. Therefore, instead of using Cuckoo hash, we use *in a white-box manner* the linear-time oblivious hash table for shuffled inputs from [2] (see also Section 3.4). This hash table has linear build time and is secure only when the input array is randomly shuffled; otherwise, it behaves conceptually in a similar manner to "standard" oblivious Cuckoo hash: lookup is performed by a scan of a stash and $O(1)$ accesses to a "main table". We therefore manage to use it for our purpose, deduplication.

Dealing with the shared stash. So far, we have assumed an idealized oblivious hash table (for randomly shuffled inputs) without stashes. As mentioned earlier, known instantiations of the oblivious hash table with the desired efficiency requirements have an extra stash that must be visited during a Lookup operation, and the constant-time lookup is only possible with a "shared stash" trick, i.e., by merging all logarithmically many stashes into a globally shared one.

Accommodating the shared stashes creates extra technicalities, as we mentioned in Challenge III in Section 2.2. To deal with them, the main idea is to support more fine-grained access to the shared memory area. That is, while in [2], every extract requires to atomically scan the shared memory to retrieve all elements that "belong" to some given level, in our construction we avoid such linear scans by using more efficient data structures. Specifically, we use a version of an oblivious dictionary which supports lookup of elements w.r.t. various auxiliary keys such as the level they came from or the logical address.

To elaborate, recall that each level is associated with a stash of $O(\log \lambda)$ elements, and thus the shared stash consists of $O(\log N \cdot \log \lambda)$ elements. All such elements are store the oblivious dictionary accompanied with the auxiliary keys. Specifically, an element is inserted to the dictionary when a level is newly built, is queried by a logical address when lookups are performed during the fetch phase, and is popped from the dictionary when a level is being extracted. Recall that each operation of the oblivious dictionary takes only $\mathsf{poly}\log(\log N + \log \lambda)$ time, e.g., by instantiating a perfect ORAM [8], and that each level is at least $\mathsf{poly}(\log N + \log \lambda)$ in size. The efficiency follows as inserting or popping elements take only a $o(1)$ fraction of time during build or extract a level, and only $O(1)$ queries are performed during a fetch. The security follows since the dictionary is perfectly oblivious.

Organization. The remaining of the paper is organized as follows. In Section 3 we provide the preliminaries, which includes definition of obliviousness, some basic building blocks we use in our construction, our 2-key dictionary and revisit

and overview the oblivious hash table construction of [2]. In Section 4 we provide our deduplication algorithm, and in Section 5 we provide our deamortized construction. The case of combining the stashes is deferred to the full version.

3 Preliminaries

The security parameter is denoted λ and it is given as input to algorithms in unary (i.e., as 1^λ). A function $\mathsf{negl} : \mathbb{N} \to \mathbb{R}^+$ is *negligible* if for every constant $c > 0$ there exists an integer N_c such that $\mathsf{negl}(\lambda) < \lambda^{-c}$ for all $\lambda > N_c$. Two sequences of random variables $X = \{X_\lambda\}_{\lambda \in \mathbb{N}}$ and $Y = \{Y_\lambda\}_{\lambda \in \mathbb{N}}$ are *computationally indistinguishable* if for any probabilistic polynomial time algorithm $\mathcal{A}$, there exists a negligible function $\mathsf{negl}(\cdot)$ such that $\left|\Pr\left[\mathcal{A}(1^\lambda, X_\lambda) = 1\right] - \Pr\left[\mathcal{A}(1^\lambda, Y_\lambda) = 1\right]\right| \leq \mathsf{negl}(\lambda)$ for all $\lambda \in \mathbb{N}$. For $n \in \mathbb{N}$, denote $[n] = \{1, \ldots, n\}$.

Random-access machines (RAM). A RAM (or *a RAM program*) is an interactive Turing machine that consists of a memory and a CPU. The program maps some input to an output where its computation is performed by the CPU and using the interaction with the memory. The memory is denoted as $\mathsf{mem}[N, w]$ and is indexed by the logical address space $[N] = \{1, 2, \ldots, N\}$. We denote by w to denote the bit-length of each block. The CPU has an internal state that consists of $O(1)$ words. The memory supports read/write instructions $(\mathsf{op}, \mathsf{addr}, \mathsf{data})$ where $\mathsf{op} \in \{\mathsf{read}, \mathsf{write}\}$, $\mathsf{addr} \in [N]$ and $\mathsf{data} \in \{0,1\}^w \cup \{\bot\}$:

- If $\mathsf{op} = \mathsf{read}$ then $\mathsf{data} = \bot$ and the returned value is the content of the block located in the logical address addr in the memory.
- If $\mathsf{op} = \mathsf{write}$ then the memory data in logical address addr is updated to data.

We use the standard setting that $w = \Theta(\log N)$ (so an address can be stored in a word). We follow the standard convention that the CPU performs one word-level operation per unit time, i.e., such that additions or subtraction (arithmetic operations), bitwise operations such as AND, OR, NOT or shift, memory accesses or evaluation of pseudorandom function.

3.1 Oblivious Machines

Intuitively, we say that a machine M is *oblivious* if there exists a simulator that can simulate its access pattern without knowing the input. Specifically, there exists a simulator Sim such that, for all inputs x, all memory accesses in the computation $M(x)$ can be simulated by Sim where Sim just receives the length of x (i.e., $|x|$) and not the input itself.

We say that a RAM program M_f oblivious simulates a (deterministic) RAM program f if for every input x it holds that $M_f(x) = f(x)$, and that M_f is oblivious. For randomized functionalities, we require that the joint distribution of the output of M_f and the access pattern of the simulator is indistinguishable from

the joint distribution of the output of f and its access pattern in the computation. See discussion in [2]. We are now ready for the definition of oblivious simulation:

Definition 3.1 (Oblivious simulation)**.** *Let $f, M_f : \{0,1\}^* \to \{0,1\}^*$ be two RAM machines. We say that M_f* obliviously simulates *f if there exists a probabilistic polynomial time simulator* Sim *such that for every input $x \in \{0,1\}^*$, the following holds:*

$$\{(\mathsf{out}, \mathsf{Addr}) \ : \ (\mathsf{out}, \mathsf{Addrs}) \leftarrow M_f(1^\lambda, x)\}_\lambda \approx \left\{\left(f(x), \mathsf{Sim}(1^\lambda, 1^{|x|}\right)\right\}_\lambda$$

depending on whether $\approx$ refers to computational, statistical, or perfectly indistinguishable we say that M_f is computationally, statistically, or perfectly oblivious, respectively.

Reactive random-access machines. We consider functionalities that are reactive, i.e., proceed in stages, where the functionality preserves an internal state between stages. Such a reactive functionality can be described as a sequence of RAM machines, where each machine also receives as an input a state, updates it, and the output is the input state for the next machine. We use it to capture building blocks such as oblivious hash tables (see Section 3.4).

A reactive machine $\mathcal{F}$ receives commands of the form $(\mathsf{command}_i, \mathsf{inp}_i)$ and produces an output out_i while maintaining some (secret) internal state. Our definition considers an adversary $\mathcal{A}$ (a distinguisher) that participates in either a real execution or an ideal one, and with each command receives the access pattern (resp. simulated access pattern) and the output of the algorithm (resp. output of the functionality). The adversary $\mathcal{A}$ can then adaptively choose the next command to execute.

Definition 3.2 (Oblivious simulation of a reactive functionality)**.** *We say that a reactive machine $M_\mathcal{F}$ is an oblivious implementation of the reactive functionality $\mathcal{F}$ if there exists a PPT simulator* Sim *such that for any non-uniform PPT (stateful) adversary $\mathcal{A}$, the view of the adversary $\mathcal{A}$ in the following two experiments $\mathsf{Expt}^{\mathsf{real},M_\mathcal{F}}(1^\lambda)$ and $\mathsf{Expt}^{\mathsf{ideal},\mathcal{F}}_{\mathcal{A},\mathsf{Sim}}(1^\lambda)$ is computationally indistinguishable:*

$\mathsf{Expt}^{\mathsf{real},M_\mathcal{F}}_{\mathcal{A}}(1^\lambda)$:	$\mathsf{Expt}^{\mathsf{ideal},\mathcal{F}}_{\mathcal{A},\mathsf{Sim}}(1^\lambda)$:
Let $(\mathsf{cmd}_i, \mathsf{inp}_i) \leftarrow \mathcal{A}(1^\lambda)$ *Loop while* $\mathsf{cmd}_i \neq \bot$: $\quad \mathsf{out}_i, \mathsf{Addr}_i \leftarrow M_\mathcal{F}(1^\lambda, \mathsf{cmd}_i, \mathsf{inp}_i)$ $\quad (\mathsf{cmd}_i, \mathsf{inp}_i) \leftarrow \mathcal{A}(1^\lambda, \mathsf{out}_i, \mathsf{Addr}_i)$	*Let* $(\mathsf{cmd}_i, \mathsf{inp}_i) \leftarrow \mathcal{A}(1^\lambda)$ *Loop while* $\mathsf{cmd}_i \neq \bot$: $\quad \mathsf{out}_i \leftarrow \mathcal{F}(\mathsf{cmd}_i, \mathsf{inp}_i)$. $\quad \mathsf{Addr}_i \leftarrow \mathsf{Sim}(1^\lambda, \mathsf{cmd}_i)$. $\quad (\mathsf{cmd}_i, \mathsf{inp}_i) \leftarrow \mathcal{A}(1^\lambda, \mathsf{out}_i, \mathsf{Addr}_i)$

We define statistical or perfect simulation analogously, requiring the two experiments to be either statistically close or identically distributed.

ORAM simulation overhead. We consider the standard ORAM functionality, implementing a logical memory. In this functionality, the user gets to choose the next command (i.e., either read or write) as well as the address and data according

to the access pattern it has observed so far. During its life span, the functionality holds (as an internal state) N memory blocks, each of size w. Denote the internal state $\mathbf{X}[1, \ldots, N]$. Initially, $\mathbf{X}[\mathsf{addr}] = 0$ for every $\mathsf{addr} \in [N]$. The functionality is as follows:

- $\mathsf{Access}(\mathsf{op}, \mathsf{addr}, \mathsf{data})$: where $\mathsf{op} \in \{\mathsf{read}, \mathsf{write}\}$, $\mathsf{addr} \in [N]$ and $\mathsf{data} \in \{0,1\}^w$.
  1. If $\mathsf{op} = \mathsf{read}$, set $\mathsf{data}^* := \mathbf{X}[\mathsf{addr}]$.
  2. If $\mathsf{op} = \mathsf{write}$, set $\mathbf{X}[\mathsf{addr}] := \mathsf{data}$ and $\mathsf{data}^* := \mathsf{data}$.
  3. Output data^*.

Typically, the metric of interest for ORAM construction is known as *computation overhead* and it is is defined as the (multiplicative) blowup in runtime of the compiled program. We distinguish between the worst-case and the amortized variants:

- **Amortized computation overhead**: We say that the amortized computation overhead of the ORAM is $g \colon \mathbb{N} \to \mathbb{N}$ if for every sequence of operations $((\mathsf{op}_1, \mathsf{addr}_1, \mathsf{data}_1), \ldots, (\mathsf{op}_q, \mathsf{addr}_q, \mathsf{data}_q))$ at most $g(N) \cdot q$ computation steps are taken during the execution of the ORAM.
- **Worst-case computation overhead**: We say that the amortized computation overhead of the ORAM is $g \colon \mathbb{N} \to \mathbb{N}$ if every sequence of operations $((\mathsf{op}_1, \mathsf{addr}_1, \mathsf{data}_1), \ldots, (\mathsf{op}_q, \mathsf{addr}_q, \mathsf{data}_q))$, handling each operation in the sequence consumes $g(N)$ computation steps.

It is immediate that worst-case overhead $g(N)$ directly implies amortized overhead of $g(N)$ but the converse is not necessarily true.

3.2 Basic Building Blocks

We briefly describe few functionalities that we use in our construction, and refer to [2] for their implementation:

1. $\mathsf{Intersperse}(X, Y)$ is an algorithm that takes two arrays, each is randomly shuffled, returns a randomly shuffled array, and runs in linear time (i.e., $O(|X| + |Y|)$). It obliviously implements the ideal functionality $\mathcal{F}_{\mathsf{Shuffle}}$–which takes an array and randomly shuffled it.
2. $\mathsf{IntersperseRD}(X)$ takes an array that contains real and dummy elements, and is assumed that all real elements are shuffled among themselves, but there is no guarantee about the locations of the dummies in the array (e.g., they can all reside at the end of the array). The algorithm returns a randomly shuffled array, runs in linear time, and obliviously implements the ideal functionality $\mathcal{F}_{\mathsf{Shuffle}}$.
3. Compaction: Given an array in which some of the elements are distinguished, it moves all distinguished elements to the beginning of the array and runs in linear time.

3.3 Perfectly Oblivious 2-Key Dictionary

In this section we use known results to get a new oblivious dictionary-like data structure. Our structure, termed two-key dictionary, needs to support three operations: Insert, PopKey, and PopTime. Each element has a key k and a label t for "time". The Insert operation takes the key k along with timestamp t and a value v and adds (k, t, v) to the dictionary, the k appears at most once in the dictionary (so that it overwrites if there is a previous tuple (k, t', v') for the same k). One can pop an element with key k by using $\mathsf{PopKey}(k)$—the operation returns and removes the element with the key k. Analogously, $\mathsf{PopTime}(t_1, t_2)$ takes as input a time period $[t_1, t_2]$ and returns an element that is labeled with a timestamp t in the given period. This (reactive) functionality appears as Functionality 3.3.

Functionality 3.3: $\mathcal{F}_{\mathsf{2KeyDict}}$ - Dictionary Functionality

---

- Initialization of the state: let M be an empty list indexed by $k \in [K]$ for the given key space K, where all $M[k]$ are initialized as $\perp$.
- $\mathcal{F}_{\mathsf{2KeyDict}}.\mathsf{Insert}(k, t, v)$:
  - **Input:** a key k, time $t \in \mathbb{N}$, and a value v, where k might be $\perp$, i.e., a dummy insertion.
  - **The procedure:**
    1. If $k \neq \perp$, set $M[k] := (t, v)$.
  - **Output:** The Insert operation has no input.
- $\mathcal{F}_{\mathsf{2KeyDict}}.\mathsf{PopKey}(k)$:
  - **Input:** a key k (that might be $\perp$, i.e., dummy).
  - **The procedure:**
    1. Set $(t^*, v^*) := M[k]$ and then set $M[k] := \perp$.
  - **Output:** The valude v^*.
- $\mathcal{F}_{\mathsf{2KeyDict}}.\mathsf{PopTime}(t_1, t_2)$:
  - **Input:** time $t_1, t_2 \in \mathbb{N}$ such that $t_1 < t_2$.
  - **The procedure:**
    1. Let k be the smallest index such that $M[k] = (t^*, \cdot)$ for some $t^* \in [t_1, t_2]$. If no such k exists, set $v^* := \perp$. Otherwise, set $(t^*, v^*) := M[k]$ and then set $M[k] := \perp$.
  - **Output:** The value v^*.

Theorem 3.4. *Assume the tuple of (k, t, v) (i.e., key, time, and value) can be stored in a constant number of memory words. Assume further that the two-key dictionary needs to support at most n elements. There exists a perfectly oblivious implementation of functionality $\mathcal{F}_{\mathsf{2KeyDict}}$ such that each operation* Insert*,* PopKey*, and* PopTime *takes $O(\log^4 n)$ time in the worst-case.*

Proof. The first step is to obtain a perfectly oblivious ORAM that has $O(\log^3 n)$ worst-case overhead when simulating a memory of size n. Such an ORAM can be obtained by applying the deamortization technique of Ostrovsky and Shoup [29]

on the (amortized) $O(\log^3 n)$ overhead perfect ORAM construction of Chan et al. [8]. It directly works (although never formally stated to the best of our knowledge) since the construction of Chan et al. is based exactly on the original hierarchical framework of Goldreich and Ostrovsky [16, 17] which was deamortized in Ostrovsky and Shoup [29].

Given this ORAM, it is straightforward to prove the theorem by just compiling a non-oblivious implementation of $\mathcal{F}_{\mathsf{2KeyDict}}$. For the latter, we instantiate two balanced binary search trees (e.g., red-black tree), where the first tree orders elements according to the key k, and the second tree orders elements by the given time t. This implementation has logarithmic cost for each operation. Therefore, compiling it using the above worst-case perfect ORAM, we have a perfectly oblivious implementation of $\mathcal{F}_{\mathsf{2KeyDict}}$ taking $O(\log^4 n)$ time in the worst case. $\square$

3.4 Oblivious Hash Table for Shuffled Inputs

Here we recall what an oblivious hash table is and the shuffled inputs assumption (taken from [2, Section 4.4]). An oblivious hash table is a (reactive) functionality that supports three operations: Build, Lookup and Extract that are defines as follows. The Build operation gets as input an array of items, Lookup is used to search for an item and then delete it, and finally Extract returns the "remaining" elements in the table. Obliviousness means, as usual, that the access patterns throughout the life time of the system should be unrelated to the elements in the array nor the values being searched for. We will achieve this guarantee *assuming the input to* Build *is random shuffled.*

Functionality 3.5: $\mathcal{F}_{\mathsf{HT}}$ - Hash Table Functionality for Non-Recurrent Lookups

---

$\underline{\mathcal{F}_{\mathsf{HT}}.\mathsf{Build}(\mathbf{I})}$: The input

- **Input:** an input array $\mathbf{I} = (a_1, \ldots, a_n)$ containing n elements, where each a_i is either dummy or a (key, value) pair denoted $(k_i, v_i) \in \{0,1\}^D \times \{0,1\}^D$ for some $D \in \mathbb{N}$. We assume that both the key and the value can be stored in $O(1)$ memory words, i.e., $D = O(w)$ where w denotes the word size.
- **The procedure:**
  1. Initialize the internal state state to $\mathbf{I}, \mathbf{P}$ where $\mathbf{P} = \emptyset$. $\mathbf{P}$ will store the keys that were already queried.
  2. **Output:** The Build operation has no input.

$\underline{\mathcal{F}_{\mathsf{HT}}.\mathsf{Lookup}(k)}$:

- **Input:** a key $k \in \{0,1\}^D \cup \{\perp\}$.
- **The procedure:**
  1. Parse the internal state as $\mathsf{state} = (\mathbf{I}, \mathbf{P})$.
  2. If $k \in \mathbf{P}$ (i.e., k is a recurrent lookup) then halt and output fail. (Security is only guaranteed if no such recurrent lookup is performed, so in the construction we do not need to check this explicitly.)

3. If $k = \perp$ or $k \notin \mathbf{I}$ then set $v^* = \perp$.
4. Otherwise, set $v^* = v$ where v is the value that corresponds to the key k in $\mathbf{I}$.
5. Update $\mathbf{P} = \mathbf{P} \cup \{(k, v)\}$.

– **Output:** The element v^*.

$\mathcal{F}_{\mathsf{HT}}.\mathsf{Extract}()$:

– **The procedure:**
1. Parse the internal state $\mathsf{state} = (\mathbf{I}, \mathbf{P})$.
2. Define an array $\mathbf{I}' = (a'_1, \ldots, a'_n)$ as follows. For $i \in [n]$ set $a'_i = a_i$ if $a_i = (k, v) \notin \mathbf{P}$. Otherwise, set $a'_i = \mathsf{dummy}$.
3. Shuffle $\mathbf{I}'$ uniformly at random.

– **Output:** The array $\mathbf{I}'$.

The work of Asharov et al. [2, Corollary 8.9] shows a construction of a hash table, denoted as CombHT, with the following properties:

Theorem 3.6. *Assume that one-way functions exist. Then, for any $c \in \mathbb{N}$, there exists a construction, denoted as* CombHT*, that (computationally) obliviously implements the $\mathcal{F}_{\mathsf{HT}}$ functionality (Functionality 3.5) with the following properties:*

1. *The input array is $\log^{9+c} \lambda \leq n \leq 2^\lambda$;*
2. *The input assumption is that the input array is randomly shuffled;*
3. Build *and* Extract *each take $O(n)$ time.* Build *outputs a main table and a stash of size $O(\log \lambda)$;*
4. Lookup *takes $O(1)$ time in addition to linearly scanning a stash of size $O(\log \lambda)$.*

The high-level idea of this construction is given in the full version for completeness, and we refer to [2, Section 8.4] for full details.

4 Oblivious Deduplication in Linear Time

Consider two arrays X_1 of size n and X_2 of size $2n$, where it is guaranteed that at least half of the elements in X_2 are dummies,[8] and the keys in each array are unique. In what follows we describe an algorithm for merging the (real) contents of the arrays while removing duplicates, preferring the ones in X_1. That is, viewing the input arrays as sets, we compute $X_1 \cup X_2$ while preferring duplicate elements from X_1. We start with the abstract functionality $\mathcal{F}_{\mathsf{Dedup}}$ and then give our implementation.

[8] One could easily modify our algorithm to work more generally for a list X_2 of size m which has at least n dummies and result with an array of size m. We chose to be concrete for simplicity.

The functionality $\mathcal{F}_{\mathsf{Dedup}}$. The exact functionality is described next. It can be viewed as a non-efficient non-oblivious implementation. The input consists of an array X_1 of size n and an array X_2 of size $2n$. The array X_2 contains at most n real elements. Every key appears at most once in each input list (a key may appear once in each of the arrays with different associate values). The functionality does the following:

1. Initialize an array Y of size $2n$.
2. Copy to Y all real elements in both arrays X_1 and X_2. If the two arrays contain the same key (with possibly different associated value), then remove the copy from X_2 and prefer the one in X_1. Pad Y with dummies to be of size $2n$.
3. Uniformly shuffle Y and return it.

The main theorem of this section is stated next.

Theorem 4.1. *There is an algorithm that implements the functionality* $\mathcal{F}_{\mathsf{Dedup}}$ *in time* $O(n)$ *for* $n \geq \log^{11} \lambda$ *(and with negligible error probability). The algorithm is computationally oblivious if one-way functions exist and if the input arrays are independently randomly shuffled.*

Recall the $O(n)$ time non-oblivious algorithm from Section 2—it is clearly non-oblivious and "naively" replacing the Cuckoo hash with an oblivious version thereof does not meet our goal since known constructions require $\omega(n)$ time for building. However, we do not necessarily need the full power of Cuckoo hashing since we are guaranteed that the input lists are shuffled. Therefore, instead of using Cuckoo hash, we use the hash table CombHT from Section 3.4 which has linear build time and otherwise behaves conceptually in a similar manner to "standard" oblivious Cuckoo hash: lookup is performed by a scan of a stash and $O(1)$ accesses to a "main table". The idea therefore is conceptually in the same spirit, but we rearrange and then compose the procedures of CombHT (in a non-black-box way) to guarantee obliviousness. Namely in Step 3, for each duplicate that reside in the first hash table, we mark the element by its counterpart in the second hash table; then in Step 5, we are able to emulate identically the lookup procedures on the second table (even we perform no access on its stash). We refer the reader to the construction of CombHT in the full version for a comprehensive construction.

The algorithm $\mathsf{Dedup}(X_1, X_2)$ works as follows:

1. Perform $\mathsf{T}_1 := \mathsf{HT.Build}(X_1)$ and $\mathsf{T}_2 := \mathsf{HT.Build}(X_2)$.
2. Denote $\mathsf{T}_1 = (\mathsf{sk}_1, \mathsf{OBins}_1, \mathsf{CombS}_{1,\mathsf{T}}, \mathsf{CombS}_{1,\mathsf{S}}, \mathsf{OF}_{1,\mathsf{T}}, \mathsf{OF}_{1,\mathsf{S}})$ and $\mathsf{T}_2 = (\mathsf{sk}_2, \mathsf{OBins}_2, \mathsf{CombS}_{2,\mathsf{T}}, \mathsf{CombS}_{2,\mathsf{S}}, \mathsf{OF}_{2,\mathsf{T}}, \mathsf{OF}_{2,\mathsf{S}})$.
3. Initialize an empty array L. Linearly scan $\mathsf{OF}_{2,\mathsf{S}}$ and $\mathsf{Comb}_{2,\mathsf{S}}$, and for each element (k, v), perform the following:
   (a) Perform a real lookup $(k', v') := \mathsf{T}_1.\mathsf{Lookup}(k)$. If k is found in T_1 then:
      i. Mark the element (k', v') at T_1 as "`CombS`" if k comes from $\mathsf{CombS}_{2,\mathsf{S}}$, or "`OF`" if k comes from $\mathsf{OF}_{2,\mathsf{S}}$.

ii. Mark the element (k, v) as "accessed" in T_2.

(b) Write (k', v') to the next slot in L (including the mark when k is found).

In the end of this loop, obliviously shuffle L.

4. Perform $S'_1 := \mathsf{T}_1.\mathsf{Extract}()$. Then, perform $S_1 := \mathsf{Intersperse}(S'_1 \| L)$.
5. Linearly scan S_1, and for each element (k, v), perform the following:
   (a) If k is marked as "OF", then perform a real lookup $(k', v') := \mathsf{T}_2.\mathsf{Lookup}(k)$ while not scanning the stashes $(\mathsf{OF}_{2,\mathsf{S}}, \mathsf{CombS}_{2,\mathsf{S}})$ and proceeding as if k is found in $\mathsf{OF}_{2,\mathsf{S}}$.
   (b) If k is marked as "CombS", then perform a real lookup $(k', v') := \mathsf{T}_2.\mathsf{Lookup}(k)$ while not scanning the stashes $(\mathsf{OF}_{2,\mathsf{S}}, \mathsf{CombS}_{2,\mathsf{S}})$ and proceeding as if k is found in $\mathsf{CombS}_{2,\mathsf{S}}$.
   (c) If k is not marked, then perform a real lookup $(k', v') := \mathsf{T}_2.\mathsf{Lookup}(k)$ while not scanning the stashes $(\mathsf{OF}_{2,\mathsf{S}}, \mathsf{CombS}_{2,\mathsf{S}})$ and proceeding as if k is not found in the stashes.
6. Perform $S_2 := \mathsf{T}_2.\mathsf{Extract}()$ (recall that "accessed" elements in T_2 are not extracted).
7. Run $Z := \mathsf{Intersperse}(S_1 \| S_2)$. Run tight compaction on Z to move all dummy elements to the end. Truncate the array to be of size $2n$. Run IntersperseRD to randomly shuffle Z.
8. Output Z.

The full proof of Theorem 4.1 appears in the full version. Here, we briefly argue that the efficiency is as required. In step 3, we scan the stashes of T_2 and then for each element perform a $O(\log \lambda)$-time lookup, and then we shuffle L. Since the stashes are of size $O(\log \lambda)$, the running time of this step is $O(\log^2 \lambda) \leq O(n)$. Step 5 consist of a linear scan of a list and then an $O(1)$ lookup on each item. Steps 4 and 6 consume $O(n)$ time. Step 7 consumes $O(n)$ time, as well. Overall, the overhead is linear in n, as needed.

5 The ORAM Construction with Worst Case Complexity

In this section we present our deamortized construction.

The combined stash technique. As we saw in Section 3.4, our hash table supports lookup in $O(1)$ time in addition to a lookup in a stash of size $O(\log \lambda)$. To allow faster lookups, constructions use "the combined stash" technique (see [2,7,20]). According to this technique, all stashes of all levels are combined into one global stash. Then, utilizing the fact that we look for the same element in all levels (or dummy lookup once the element is found), we have to search for the element only once in the global stash (instead of searching for it in $O(\log N)$ different stashes), and then spend just $O(1)$ lookup time per level.

As we mentioned in the introduction, the fact that there is a shared memory to all levels introduces some complications in the final construction. We therefore present our deamortized construction in two steps:

1. In Section 5.1 we look at a somewhat idealized construction in which the main building block is a hash table that takes $O(1)$ time per lookup, while it takes linear time for $\mathsf{Extract}$ and Build (on shuffled inputs). That is, "there is no stash". We emphasize that we do not know how to realize such an oblivious hash table. Nevertheless, we describe this construction for aiding understanding and capturing the main ideas behind our deamortized construction.
2. In the full version we proceed to our final construction, in which the hash table is implemented as in Section 3.4, that is, $O(1)$ lookup time in addition to a scan of a stash of size $O(\log \lambda)$. To achieve effectively $O(1)$ time per lookup, we have also to use the combining stash technique as we described above.

5.1 Assuming Hash Table with $O(1)$ Lookup Time

In this section, we assume the existence of a construction of a hash table, denoted as HT, that achieves the following:

Assumption 5.1. *Assume that for any $c \in \mathbb{N}$ there exists a construction, denoted as HT that obliviously implements the $\mathcal{F}_{\mathsf{HT}}$ functionality (Functionality 3.5) with the following properties:*

1. *The input array is $\log^{9+c} \lambda \leq n \leq 2^{\lambda}$;*
2. *The input assumption is that the input array is randomly shuffled;*
3. Build *and* $\mathsf{Extract}$ *each take $O(n)$ time.*
4. Lookup *takes $O(1)$ time.*

This is equivalent to Theorem 3.6 where the construction has no stash and Lookup takes worst-case $O(1)$ time. We proceed to the construction. We first start with the underlying primitives and the memory organization, and proceed to the specification of the construction.

Structure: Let $\ell = \lceil 11 \log\log \lambda \rceil$ and $L = \lceil \log N \rceil$.

1. Each level $i \in \{\ell+1, \ldots, L\}$ consists of four instances of HT as in Assumption 5.1, each of capacity 2^i. We denote the levels as $(\mathsf{A}^{\mathsf{HF}}_{\ell+1}, \ldots, \mathsf{A}^{\mathsf{HF}}_{L})$, $(\mathsf{A}^{\mathsf{F}}_{\ell+1}, \ldots, \mathsf{A}^{\mathsf{F}}_{L})$, $(\mathsf{B}^{\mathsf{HF}}_{\ell+1}, \ldots, \mathsf{B}^{\mathsf{HF}}_{L})$ and $(\mathsf{B}^{\mathsf{F}}_{\ell+1}, \ldots, \mathsf{B}^{\mathsf{F}}_{L})$.
2. Two perfect dictionaries (see Section 3.3), denotes as $\mathsf{A}_\ell, \mathsf{B}_\ell$, each of capacity $2^{\ell+1} + O(\log N \cdot \log \lambda)$. Each dictionary holds elements of the form $(\mathsf{addr}, \mathsf{data})$ where $\mathsf{addr} \in [N]$, $\mathsf{data} \in \{0,1\}^w$.
3. Pointers $(\mathsf{A}_\ell, \ldots, \mathsf{A}_L)$, $(\mathsf{B}_\ell, \ldots, \mathsf{B}_L)$ where each A_i points to either $\{\mathsf{A}^{\mathsf{HF}}_i, \mathsf{A}^{\mathsf{F}}_i, \mathtt{Null}\}$ and each B_i to $\{\mathsf{B}^{\mathsf{HF}}_i, \mathsf{B}^{\mathsf{F}}_i, \mathtt{Null}\}$, where $\mathtt{Null}$ is a null pointer.
4. A global counter ctr, initialized as 0.

Construction 5.2: Oblivious RAM $\mathsf{Access}(\mathsf{op}, \mathsf{addr}, \mathsf{data})$

- **Input:** $\mathsf{op} \in \{\mathsf{read}, \mathsf{write}\}$, $\mathsf{addr} \in [N]$ and $\mathsf{data} \in \{0,1\}^w$.

- **Secret state:** As above.
- **Initialization:** ctr is initialized to 0 as above, and all other data structures are initialized as empty.
- **The algorithm:**
Lookup:
  1. Initialize $\mathsf{found} = \mathsf{false}$, $\mathsf{data}^* = \bot$.
  2. Perform $\mathsf{fetched} := \mathsf{A}_\ell.\mathsf{PopKey}(\mathsf{addr})$.
  3. If $\mathsf{fetched} \neq \bot$: then $\mathsf{B}_\ell.\mathsf{PopKey}(\bot)$.
  Otherwise, $\mathsf{fetched} := \mathsf{B}_\ell.\mathsf{PopKey}(\mathsf{addr})$.
  4. If $\mathsf{fetched} \neq \bot$: set $\mathsf{found} = \mathsf{true}$.
  5. For each $i \in \{\ell+1, \ldots, L\}$ in increasing order, do (if A_i (or B_i resp.) is `Null`, then let the result of Lookup be $\bot$):
     (a) If $\mathsf{found} = \mathsf{false}$:
        i. Set $\mathsf{fetched} := \mathsf{A}_i.\mathsf{Lookup}(\mathsf{addr})$.
        ii. If $\mathsf{fetched} \neq \bot$ then set $\mathsf{found} := \mathsf{true}$ and $\mathsf{data}^* := \mathsf{fetched}$.
     (b) Else, perform $\mathsf{A}_i.\mathsf{Lookup}(\bot)$.
     (c) If $\mathsf{found} = \mathsf{false}$:
        i. Set $\mathsf{fetched} := \mathsf{B}_i.\mathsf{Lookup}(\mathsf{addr})$.
        ii. If $\mathsf{fetched} \neq \bot$ then set $\mathsf{found} := \mathsf{true}$ and $\mathsf{data}^* := \mathsf{fetched}$.
     (d) Else, perform $\mathsf{B}_i.\mathsf{Lookup}(\bot)$.

Write back:
  6. If $\mathsf{found} = \mathsf{false}$, i.e., this is the first time addr is being accessed, set $\mathsf{data}^* = 0$.
  7. Let $(k, v) := (\mathsf{addr}, \mathsf{data}^*)$ if this is a read operation; else let $(k, v) := (\mathsf{addr}, \mathsf{data})$.
  8. Insert (k, v) into A_ℓ and B_ℓ using $\mathsf{Insert}(k, \mathsf{ctr} \bmod 2^{\ell+1}, v)$.

Rebuild:
  9. Increment ctr by 1.
  10. For $i \in \{\ell+1, \ldots, L\}$:
     (a) If $\mathsf{ctr} \equiv 2^{i-2} \bmod 2^i$ then continue to 1-out-of-4 case:

If $\mathsf{ctr} \equiv$	$0 \bmod 2^i$	$2^{i-2} \bmod 2^i$	$2 \cdot 2^{i-2} \bmod 2^i$	$3 \cdot 2^{i-2} \bmod 2^i$
Set $\mathsf{A}_i :=$	`Null`	$\mathsf{A}_i^{\mathsf{HF}}$	`Null`	$\mathsf{A}_i^{\mathsf{F}}$
Set $\mathsf{B}_i :=$	$\mathsf{B}_i^{\mathsf{F}}$	$\mathsf{B}_i^{\mathsf{F}}$	$\mathsf{B}_i^{\mathsf{HF}}$	`Null`
Start	$\mathsf{RebuildHF}(\mathsf{A}_i^{\mathsf{HF}})$	$\mathsf{RebuildHF}(\mathsf{B}_i^{\mathsf{HF}})$	$\mathsf{RebuildF}(\mathsf{A}_i^{\mathsf{F}})$	$\mathsf{RebuildF}(\mathsf{B}_i^{\mathsf{F}})$

  By starting a task we mean to add the relevant task into the list Tasks. The procedures $\mathsf{RebuildHF}$ and $\mathsf{RebuildF}$ are defined below.
  11. In a round robin fashion, for each task $\mathsf{t} \in \mathsf{Tasks}$, execute $\mathsf{t.eachEpoch}$ steps.
  12. Return v.

Before proceeding, we refer the reader to depictions of the rebuilding scheduling in Figures 2 and 3. In Step 10a, the schedule of the rebuild tasks is asymmetric ($\mathsf{A}_i^{\mathsf{HF}}$ and $\mathsf{A}_i^{\mathsf{F}}$ always start earlier than the B_i counterparts). This leads to the asymmetry between the setting of pointers A_i and B_i in Step 10a. Due to the

asymmetry in schedule, there is a period such that both pointers A_i and B_i are available and storing distinct sets of elements (i.e., from $\mathsf{ctr} \equiv 2^{i-2} \bmod 2^i$ to $\mathsf{ctr} \equiv 2 \cdot 2^{i-2} \bmod 2^i$). Hence, Step 5 has to fetch addr in both A_i and B_i as their contents are distinct (our schedule is determinisitic, but fetching addr in both tables is necessary as we do not know which table stores addr).

The procedure RebuildF. In this procedure, we build the table $\mathsf{A}_i^{\mathsf{F}}$ from the two tables $\mathsf{A}_{i-1}^{\mathsf{F}}$ and $\mathsf{A}_i^{\mathsf{HF}}$ (similarly, $\mathsf{B}_i^{\mathsf{F}}$ from $\mathsf{B}_{i-1}^{\mathsf{F}}$ and $\mathsf{B}_i^{\mathsf{HF}}$). This is done by extracting the two tables, running Dedup (see Section 4) on the two tables, and then building the hash table. All those operations take linear work, and therefore we can spend $O(1)$ time per Access to the ORAM and finish the task in linear time. This is formalized in the eachEpoch variable.

In case the level to be rebuilt is $\ell + 1$, we extract all elements from the dictionary of level ℓ. This takes $O(\mathsf{poly}\log\log N)$ per element. This will also be the eachEpoch value. That is, we spend $O(\mathsf{poly}\log\log N)$ work for the rebuilding of level $\ell + 1$ with each Access to the ORAM.

RebuildF($\mathsf{C}_i^{\mathsf{F}}$):

- **Input:** The task has input $\mathsf{C}_i^{\mathsf{F}} \in \{\mathsf{A}_i^{\mathsf{F}}, \mathsf{B}_i^{\mathsf{F}}\}$ for some index $i \in \{\ell + 1, \ldots, L\}$.
- eachEpoch: The total time allocated to this task is 2^{i-2}.
  1. If $i = \ell + 1$: Let $W \in O(2^i \cdot \mathsf{poly}(\log\log N))$ bound the total work of this procedure. Set $\mathsf{eachEpoch} = W/2^{i-2} = \mathsf{poly}\log\log N$.[9]
  2. If $i > \ell + 1$: The total work is $W \in O(2^i)$. Set $\mathsf{eachEpoch} = W/2^{i-2} \in O(1)$.
- **The task:**
  1. If $i = \ell + 1$, run $\mathsf{C}_{i-1}.\mathsf{PopTime}(0, 2^\ell - 1)$ repeatedly for 2^ℓ times. That is, we extract all elements with $\mathsf{ctr} \bmod 2^{\ell+1} \in [0, 2^\ell - 1]$, i.e., all elements that were added to the dictionary while building $\mathsf{A}_i^{\mathsf{HF}}$ and $\mathsf{B}_i^{\mathsf{HF}}$. Let X be the list of popped elements, and then obliviously shuffle X. Run $Y := \mathsf{C}_i^{\mathsf{HF}}.\mathsf{Extract}()$.
  2. Else $i > \ell + 1$, run $X := \mathsf{C}_{i-1}^{\mathsf{F}}.\mathsf{Extract}()$ and $Y := \mathsf{C}_i^{\mathsf{HF}}.\mathsf{Extract}()$.
  3. Run $Z := \mathsf{Dedup}(X, Y)$.
  4. Run $\mathsf{C}_i^{\mathsf{F}} := \mathsf{HT.Build}(Z)$.

The procedure RebuildHF. In this procedure, we rebuild table $\mathsf{A}_i^{\mathsf{HF}}$ from the contents of the table $\mathsf{A}_{i-1}^{\mathsf{F}}$ (or $\mathsf{B}_i^{\mathsf{HF}}$ from $\mathsf{B}_{i-1}^{\mathsf{F}}$). This is performed by adding dummy elements and building the next level. For $i > \ell + 1$ this requires linear work, and therefore we can spend $O(1)$ time per Access to the ORAM and finish the task in linear time. Likewise the case of RebuildF, the level $\ell + 1$ requires some more work but we have to finish also in linear time, so we spend more work with each access to the ORAM.

[9] Note that this implies that we run $\mathsf{poly}\log\log N$ work per each access for the first level.

For the case of $i = L$ we do not simply build $\mathsf{A}_L^{\mathsf{HF}}$ from $\mathsf{A}_{L-1}^{\mathsf{F}}$. Instead, we also have to merge the contents on $\mathsf{A}_L^{\mathsf{F}}$ and $\mathsf{A}_{L-1}^{\mathsf{F}}$ into $\mathsf{A}_L^{\mathsf{HF}}$. This is performed similarly to RebuildF: We first extract the two levels, run deduplication, and build level L.

RebuildHF($\mathsf{C}_i^{\mathsf{HF}}$):

- **Input:** The task gets as input a table $\mathsf{C}_i^{\mathsf{HF}} \in \{\mathsf{A}_i^{\mathsf{HF}}, \mathsf{B}_i^{\mathsf{HF}}\}$ for some index $i \in \{\ell+1, \ldots, L\}$.
- eachEpoch: The total time allocated to this task is 2^{i-2}.
  1. If $i = \ell+1$: Let $W \in O(2^i \cdot \mathsf{poly}(\log\log N))$ bound the total work of this procedure. Set $\mathsf{eachEpoch} = W/2^{i-2} = \mathsf{poly}\log\log N$.
  2. If $i > \ell+1$: The total work is $W \in O(2^i)$. Set $\mathsf{eachEpoch} = W/2^{i-2} \in O(1)$.
- **The task:**
  1. If $i = L$:
     (a) Run $X := \mathsf{C}_{L-1}^{\mathsf{F}}.\mathsf{Extract}()$ and $Y := \mathsf{C}_L^{\mathsf{F}}.\mathsf{Extract}()$.
     (b) Run $Z := \mathsf{Dedup}(X, Y)$.
     (c) Run $\mathsf{C}_L^{\mathsf{HF}} := \mathsf{HT.Build}(Z)$.
  2. Otherwise:
     (a) If $i = \ell+1$, run $\mathsf{C}_{i-1}.\mathsf{PopTime}(2^\ell, 2^{\ell+1}-1)$ repeatedly for 2^ℓ times. That is, we extract all elements with $\mathsf{ctr} \bmod 2^{\ell+1} \in [2^\ell, 2^{\ell+1}-1]$, i.e., all elements that were added to the dictionary while building $\mathsf{A}_i^{\mathsf{F}}$ and $\mathsf{B}_i^{\mathsf{F}}$. Let X be the list of popped elements, and then obliviously shuffle X.
     (b) Else $i > \ell+1$, run $X := \mathsf{C}_{i-1}^{\mathsf{F}}.\mathsf{Extract}()$.
     (c) Initialize an array Y of 2^{i-1} dummies.
     (d) Intersperse X and Y into Z and run $\mathsf{C}_i^{\mathsf{HF}}.\mathsf{Build}(Z)$.

Analysis We next prove the following theorem:

Theorem 5.3. *Let N be the capacity of the ORAM and let $\lambda \in \mathbb{N}$ be a security parameter. Assuming the existence of* HT *as in Assumption 5.1, Construction 5.2 obliviously implements the ORAM functionality, and each* Access *takes $O(\log N + \log^4\log\lambda)$ in the worst case.*

Proof. We start with the efficiency analysis. Each access requires two lookups (PopKey) at the dictionaries $\mathsf{A}_\ell, \mathsf{B}_\ell$ (Steps 2 and 3) and writing back to the two dictionaries (Step 8). Each dictionary contains at most $2^\ell \leq \log^{12}\lambda$, and each access costs $O(\log^4\log\lambda)$ time (see Section 3.3).

Then, we perform one access to each one of the tables $\mathsf{A}_{\ell+1}, \ldots, \mathsf{A}_L$, $\mathsf{B}_{\ell+1}, \ldots, \mathsf{B}_L$, each takes $O(1)$ time by Assumption 5.1, and overall it takes $O(\log N)$ times. So overall, the lookup and write back take $O(\log N + \log^4\log\lambda)$ work.

In the rebuild process, by construction we have exactly one task being rebuilt in each level, and start the next task only when the previous one finishes. It is easy to see that each process takes a linear time in the size of the level, and

therefore we spend $O(1)$ per task with each Access to the ORAM, except for level $\ell+1$. The procedures for level $\ell+1$ require 2^ℓ accesses to the dictionaries, each translates to $O(\log^4 \log \lambda)$ work (total $O(2^\ell \cdot \log^4 \log \lambda)$, and oblivious shuffle of a list of size 2^ℓ which can also be implemented in total $O(2^\ell \cdot \log^4 \log \lambda)$. Therefore, we spend $O(\log N + \log^4 \log \lambda)$ work for the rebuilding of all levels, combined.

Security. Since the ORAM functionality is deterministic, it is enough to separately consider correctness and obliviousness. We show here obliviousness, and then discuss correctness.

We show security in the hybrid model where we invoke $\mathcal{F}_{\mathsf{2KeyDict}}$, $\mathcal{F}_{\mathsf{HT}}$, $\mathcal{F}_{\mathsf{Dedup}}$, $\mathcal{F}_{\mathsf{Shuffle}}$ instead of oblivious dictionary, oblivious hash table, oblivious deduplication and intersperse, respectively. Replacing all ideal functionalities with the corresponding construction is straightforward using the composition theorem.

It is easy to simulate Construction 5.2: We access the two dictionaries, and then access the two hash tables in each level and finally write back to the dictionaries. The rebuild process and which hash table we use has a public schedule known to the adversary. Likewise which tasks are currently running.

We now show how to simulate the two procedures: $\mathsf{RebuildF}$ and $\mathsf{RebuildHF}$. For the case of $i > \ell+1$, in $\mathsf{RebuildF}$: We just have two ideal calls to $\mathsf{Extract}$. Since $\mathsf{Extract}$ returns an oblivious permutation of the element in the hash table, this implies that the input assumption of Dedup is preserved. We then obtain an array of size 2^i which is randomly shuffled and therefore the input assumption of $\mathcal{F}_{\mathsf{HT}}$ is preserved. Simulation is just these three ideal calls. Simulating $\mathsf{RebuildHF}$ is similar for the case of $i = L$, and for the case of $i \in \{\ell+2, \ldots, L-1\}$ it is also just ideal calls to $\mathcal{F}_{\mathsf{HT}}.\mathsf{Extract}$, $\mathcal{F}_{\mathsf{Shuffle}}$ (to intersperse the two arrays) and $\mathcal{F}_{\mathsf{HT}}.\mathsf{Build}$. As for $i = \ell+1$, in both procedures we have ideal calls to $\mathcal{F}_{\mathsf{2KeyDict}}$ and we shuffle the output so that input assumptions are preserved.

Correctness. We also prove the correctness in the hybrid model, and our goal is to show that every Access to an address addr reads the data that was most recently written to addr, i.e., satisfying the ORAM functionalities: Each $\mathsf{Access}(\mathsf{read}, \mathsf{addr}, \perp)$ for a given addr will have the answer data, according to the last operation $\mathsf{Access}(\mathsf{write}, \mathsf{addr}, \mathsf{data})$ that was given to the ORAM (with the same addr). We begin with describing two invariants in Definitions 5.4 and 5.5 and show that they hold.

Definition 5.4 (Vertical invariant)**.** *Fixing any* $\mathsf{addr} \in [N]$*, we say that* $(\mathsf{addr}, \mathsf{data})$ *is the* freshest version *at some given time* ctr *if the pair* $(\mathsf{addr}, \mathsf{data})$ *is the most recent pair having* addr *read or written by* Access *operation to the ORAM. Then, for every* $\mathsf{addr} \in [N]$*, it holds that*

- *every level in the hierarchy* $H_{\mathsf{A}} := \{\mathsf{A}_i\}_{i \in [\ell, L]}$ *consists of at most one version of* addr*, and*
- *Among all levels in* $\{\mathsf{A}_i^{\mathsf{HF}}, \mathsf{A}_i^{\mathsf{F}}\}_{i \in [\ell, L]}$ *that contain* addr *(in which some might be rebuilt and unavailable), the freshest version of* addr *must reside in the smallest level.*

This holds symmetrically for hierarchy $H_{\mathsf{B}} := \{\mathsf{B}_i\}_{i\in[\ell,L]}$.

Definition 5.5 (Horizontal invariant)**.** *For every* $\mathsf{addr} \in [N]$*, it holds that the freshest version of* addr *must fall in one of the following cases:*

1. *It is in the same level of the two hierarchies, i.e., it is in both A_i and B_i for some $i \in [\ell, L]$.*
2. *It is in $\mathsf{RebuildF}(\mathsf{A}_i^{\mathsf{F}})$ and $\mathsf{B}_i^{\mathsf{HF}}$ for some $i \in [\ell+1, L]$, and $\mathsf{B}_i = \mathsf{B}_i^{\mathsf{HF}}$.*
3. *It is in $\mathsf{A}_i^{\mathsf{F}}$ and $\mathsf{RebuildF}(\mathsf{B}_i^{\mathsf{F}})$ for some $i \in [\ell+1, L]$, and $\mathsf{A}_i = \mathsf{A}_i^{\mathsf{F}}$.*
4. *It is in B_i and either in $\mathsf{RebuildHF}(\mathsf{A}_{i+1}^{\mathsf{HF}})$ or in $\mathsf{RebuildF}(\mathsf{A}_{i+1}^{\mathsf{F}})$ for some $i \in [\ell, L-1]$.*
5. *It is in A_{i+1} and either in $\mathsf{RebuildHF}(\mathsf{B}_{i+1}^{\mathsf{HF}})$ or in $\mathsf{RebuildF}(\mathsf{B}_{i+1}^{\mathsf{F}})$ for some $i \in [\ell, L-1]$.*

The invariants imply correctness. Using the above vertical and horizontal invariants (whose proofs are below), it suffices to syntactically check the correctness: we list all possible locations of the newest version below and conclude the correctness of Access.

- In both A_ℓ and B_ℓ: the element from A_ℓ is outputted (following Step 2).
- In B_ℓ while $\mathsf{A}_{\ell+1}$ is rebuilding ($\mathsf{A}_{\ell+1}$ takes elements from A_ℓ but $\mathsf{A}_{\ell+1}$ is still unavailable): the element from B_ℓ is outputted (following Step 3).
- In $\mathsf{A}_{\ell+1}$ while $\mathsf{B}_{\ell+1}$ is rebuilding ($\mathsf{B}_{\ell+1}$ takes elements from B_ℓ but $\mathsf{B}_{\ell+1}$ is still unavailable): the element from $\mathsf{A}_{\ell+1}$ is outputted (Step 5(a)i).
- In both A_i and B_i, $i \in [\ell+1, L]$: the element from A_i is outputted (Step 5(a)i).
- In B_i while A_{i+1} is rebuilding down, for $i \in [\ell+1, L-1]$: there are two cases, either A_{i-1} has finished its rebuild down to A_i, or A_{i-1} has not yet. In both cases, the element from B_i is outputted by Step 5(c)i.
- In A_{i+1} while B_i is rebuilding down, $i \in [\ell+1, L-1]$: the element from A_{i+1} is outputted (Step 5(a)i).

Notice that for two addresses $\mathsf{addr}, \mathsf{addr}'$, it may happen that addr is in Case 4 for some i and that addr' is in Case 5 for $i' = i-1$. This means addr is in B_i while addr' is in A_i, so that both A_i and B_i are available for lookup, but they have disjoint contents addr and addr'. This special case explains the reason we perform lookup on both A_i and B_i in Steps 5(a)i and 5(c)i.

In the full version, we prove two lemmas showing that both the vertical invariant (Definition 5.4) and the horizontal invariant (Definition 5.5) hold in the construction. This concludes the proof. □

Moreover, in the full version we also use the combined stash technique and show how to deamortize it as well. We show:

Theorem 5.6. *Let N be the capacity of the ORAM and let $\lambda \in \mathbb{N}$ be a security parameter. Assuming the existence of one-way functions, the construction described above obliviously implements the ORAM functionality. The construction has $O(\log N + \log^4 \log \lambda)$ worst case overhead.*

Acknowledgments

This work is supported in part by a DARPA Brandeis award, by NSF under the award numbers CNS-1601879, CNS-2044679, by Packard Fellowship, an ONR YIP award, by the Israel Science Foundation (grants No. 2439/20 and 1774/20), by an Alon Young Faculty Fellowship, and by the BIU Center for Research in Applied Cryptography and Cyber Security in conjunction with the Israel National Cyber Bureau in the Prime Minister's Office. This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 891234.

References

1. Ajtai, M., Komlós, J., Szemerédi, E.: An $O(n \log n)$ sorting network. In: ACM STOC. pp. 1–9 (1983)
2. Asharov, G., Komargodski, I., Lin, W., Nayak, K., Peserico, E., Shi, E.: OptORAMa: optimal oblivious RAM. In: Advances in Cryptology - EUROCRYPT. pp. 403–432 (2020)
3. Asharov, G., Komargodski, I., Lin, W., Peserico, E., Shi, E.: Optimal oblivious parallel RAM. IACR ePrint Arch. **2020**, 1292 (2020)
4. Bindschaedler, V., Naveed, M., Pan, X., Wang, X., Huang, Y.: Practicing oblivious access on cloud storage: the gap, the fallacy, and the new way forward. In: ACM CCS. pp. 837–849 (2015)
5. Boyle, E., Naor, M.: Is there an oblivious RAM lower bound? In: ITCS. pp. 357–368 (2016)
6. Cash, D., Grubbs, P., Perry, J., Ristenpart, T.: Leakage-abuse attacks against searchable encryption. In: CCS. pp. 668–679 (2015)
7. Chan, T.H., Guo, Y., Lin, W., Shi, E.: Oblivious hashing revisited, and applications to asymptotically efficient ORAM and OPRAM. In: ASIACRYPT. pp. 660–690 (2017)
8. Chan, T.H., Nayak, K., Shi, E.: Perfectly secure oblivious parallel RAM. In: TCC. pp. 636–668 (2018)
9. Chan, T.H., Shi, E.: Circuit OPRAM: unifying statistically and computationally secure orams and oprams. In: TCC. pp. 72–107 (2017)
10. Chung, K.M., Liu, Z., Pass, R.: Statistically-secure ORAM with $\tilde{O}(\log^2 n)$ overhead. In: ASIACRYPT (2014)
11. Dittmer, S., Ostrovsky, R.: Oblivious tight compaction in $o(n)$ time with smaller constant. In: SCN. pp. 253–274 (2020)
12. Fletcher, C.W., Dijk, M.v., Devadas, S.: A secure processor architecture for encrypted computation on untrusted programs. In: ACM workshop on Scalable trusted computing. pp. 3–8 (2012)
13. Fletcher, C.W., Ren, L., Kwon, A., van Dijk, M., Devadas, S.: Freecursive ORAM: [nearly] free recursion and integrity verification for position-based oblivious RAM. In: ASPLOS. pp. 103–116 (2015)
14. Fredman, M.L., Willard, D.E.: Surpassing the information theoretic bound with fusion trees. J. Comput. Syst. Sci. **47**(3), 424–436 (1993)
15. Gentry, C., Halevi, S., Jutla, C., Raykova, M.: Private database access with he-over-oram architecture. In: CANS. pp. 172–191 (2015)

16. Goldreich, O.: Towards a theory of software protection and simulation by oblivious rams. In: ACM STOC. pp. 182–194 (1987)
17. Goldreich, O., Ostrovsky, R.: Software protection and simulation on oblivious RAMs. J. ACM **43**(3), 431–473 (May 1996)
18. Goodrich, M.T., Mitzenmacher, M.: Privacy-preserving access of outsourced data via oblivious RAM simulation. In: ICALP. pp. 576–587 (2011)
19. Goodrich, M.T., Mitzenmacher, M., Ohrimenko, O., Tamassia, R.: Oblivious ram simulation with efficient worst-case access overhead. In: Proceedings of the 3rd ACM Workshop on Cloud Computing Security Workshop. p. 95–100. CCSW '11 (2011)
20. Goodrich, M.T., Mitzenmacher, M., Ohrimenko, O., Tamassia, R.: Privacy-preserving group data access via stateless oblivious RAM simulation. In: SODA. pp. 157–167 (2012)
21. Grubbs, P., McPherson, R., Naveed, M., Ristenpart, T., Shmatikov, V.: Breaking web applications built on top of encrypted data. In: CCS. pp. 1353–1364 (2016)
22. Islam, M.S., Kuzu, M., Kantarcioglu, M.: Access pattern disclosure on searchable encryption: Ramification, attack and mitigation. In: 19th Annual Network and Distributed System Security Symposium, NDSS (2012)
23. Komargodski, I., Lin, W.: Lower bound for oblivious RAM with large cells. IACR Cryptol. ePrint Arch. **2020**, 1132 (2020)
24. Kushilevitz, E., Lu, S., Ostrovsky, R.: On the (in)security of hash-based oblivious RAM and a new balancing scheme. In: SODA. pp. 143–156 (2012)
25. Larsen, K.G., Nielsen, J.B.: Yes, there is an oblivious RAM lower bound! In: CRYPTO. pp. 523–542 (2018)
26. Liu, C., Wang, X.S., Nayak, K., Huang, Y., Shi, E.: ObliVM: A programming framework for secure computation. In: IEEE S&P (2015)
27. Lu, S., Ostrovsky, R.: Distributed oblivious RAM for secure two-party computation. In: TCC. pp. 377–396 (2013)
28. Maas, M., Love, E., Stefanov, E., Tiwari, M., Shi, E., Asanovic, K., Kubiatowicz, J., Song, D.: PHANTOM: practical oblivious computation in a secure processor. In: ACM CCS. pp. 311–324 (2013)
29. Ostrovsky, R., Shoup, V.: Private information storage. In: ACM STOC. pp. 294–303 (1997)
30. Patel, S., Persiano, G., Raykova, M., Yeo, K.: Panorama: Oblivious RAM with logarithmic overhead. In: IEEE FOCS (2018)
31. Ren, L., Yu, X., Fletcher, C.W., van Dijk, M., Devadas, S.: Design space exploration and optimization of path oblivious RAM in secure processors. In: The 40th Annual International Symposium on Computer Architecture, ISCA. pp. 571–582 (2013)
32. Shi, E., Chan, T.H., Stefanov, E., Li, M.: Oblivious RAM with $O((\log N)^3)$ worst-case cost. In: ASIACRYPT. pp. 197–214 (2011)
33. Stefanov, E., van Dijk, M., Shi, E., Fletcher, C.W., Ren, L., Yu, X., Devadas, S.: Path ORAM: an extremely simple oblivious RAM protocol. In: ACM CCS. pp. 299–310 (2013)
34. Stefanov, E., Shi, E.: Oblivistore: High performance oblivious cloud storage. In: IEEE S&P. pp. 253–267 (2013)
35. Stefanov, E., Shi, E., Song, D.X.: Towards practical oblivious RAM. In: NDSS (2012)
36. Thorup, M.: Randomized sorting in o(n log log n) time and linear space using addition, shift, and bit-wise boolean operations. J. Algorithms **42**(2), 205–230 (2002)
37. Wang, X., Chan, T.H., Shi, E.: Circuit ORAM: on tightness of the goldreich-ostrovsky lower bound. In: ACM CCS. pp. 850–861 (2015)

38. Wang, X.S., Huang, Y., Chan, T.H., Shelat, A., Shi, E.: SCORAM: oblivious RAM for secure computation. In: ACM CCS. pp. 191–202 (2014)
39. Williams, P., Sion, R., Tomescu, A.: Privatefs: A parallel oblivious file system. In: ACM CCS (2012)
40. Zahur, S., Wang, X.S., Raykova, M., Gascón, A., Doerner, J., Evans, D., Katz, J.: Revisiting square-root ORAM: efficient random access in multi-party computation. In: IEEE S&P. pp. 218–234 (2016)
41. Zhang, Y., Katz, J., Papamanthou, C.: All your queries are belong to us: The power of file-injection attacks on searchable encryption. In: USENIX. pp. 707–720 (2016)

A Figures

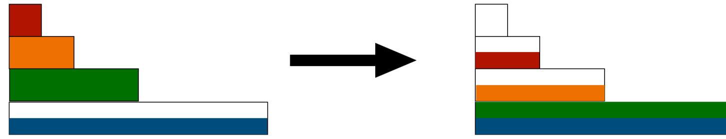

Fig. 1: The rebuild process of [3]: The first three levels are "full" and the forth is the first level which is "half full". Each level is pushed down, while levels 3 and 4 are merged. After this operation, the first level is empty, two levels are "half full" and the last level is full.

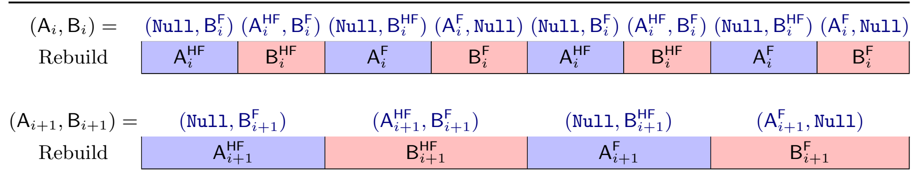

Fig. 2: The Rebuild process (for levels i and $i+1$), demonstrating which table is being rebuilt at each stage and which tables we lookup in with each access. The timeline goes left-to-right, each colored box is rebuilding the enclosed table, and the left/right side of the box denotes the starting/ending time of the rebuild. Notice that the rebuild at level $i+1$ changes the status in both levels i and $i+1$, e.g., the starting of $\mathsf{B}_{i+1}^{\mathsf{F}}$ (on the bottom-right) switches both B_i and B_{i+1} to $\texttt{Null}$, and its ending assigns $\mathsf{B}_{i+1} := \mathsf{B}_{i+1}^{\mathsf{F}}$.

(a) At time (0) (i.e., $\mathsf{ctr} \equiv 0 \bmod 2^i$), we start rebuilding $\mathsf{A}_i^{\mathsf{HF}}$ which pulls elements from level $i-1$ (colored orange next).

(b) At time (1) (i.e., $\mathsf{ctr} \equiv 2^{i-2} \bmod 2^i$), the table $\mathsf{A}_i^{\mathsf{HF}}$ is ready, and we start rebuilding $\mathsf{B}_i^{\mathsf{HF}}$.

(c) At time (2), $\mathsf{B}_i^{\mathsf{HF}}$ is ready, and we start rebuilding $\mathsf{A}_i^{\mathsf{F}}$ – merging elements from $\mathsf{A}_i^{\mathsf{HF}}$ and pulling new elements from level $i-1$ (colored red next).

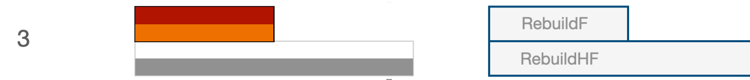

(d) At time (3) $\mathsf{A}_i^{\mathsf{F}}$ is ready, start rebuilding $\mathsf{B}_i^{\mathsf{F}}$.

(e) At time (4) $\mathsf{B}_i^{\mathsf{F}}$ is ready, and we again rebuild $\mathsf{A}_i^{\mathsf{HF}}$, pulling new elements (colored green next). $\mathsf{A}_{i+1}^{\mathsf{HF}}$ start rebuilding, pulling the elements from $\mathsf{A}_i^{\mathsf{F}}$.

(f) At time (5), $\mathsf{A}_i^{\mathsf{HF}}$ is ready with a new content (colored green), while $\mathsf{B}_i^{\mathsf{F}}$ is still active, and $\mathsf{B}_i^{\mathsf{HF}}$ starts to rebuild. $\mathsf{A}_{i+1}^{\mathsf{F}}$ is still rebuilding, as it is bigger.

(g) At time (6), $\mathsf{A}_{i+1}^{\mathsf{HF}}$ is ready with the old content, and $\mathsf{B}_i^{\mathsf{HF}}$ is ready with the new content. We start rebuilding $\mathsf{A}_i^{\mathsf{F}}$ to fetch new content from A_{i-1}.

(h) At time (7), $\mathsf{A}_i^{\mathsf{F}}$ and $\mathsf{A}_{i+1}^{\mathsf{F}}$ are both full, $\mathsf{B}_i^{\mathsf{F}}$ and $\mathsf{B}_{i+1}^{\mathsf{F}}$ are rebuilding.

Fig. 3: The rebuilding process. $\mathsf{A}_i^{\mathsf{HF}}$ and $\mathsf{A}_i^{\mathsf{F}}$ are both shown in the same table, likewise $\mathsf{B}_i^{\mathsf{HF}}$ and $\mathsf{B}_i^{\mathsf{F}}$.