

GigaDORAM: Breaking the Billion Address Barrier

Brett Falk*
University of Pennsylvania

Rafail Ostrovsky*
UCLA

Matan Shtepel*
UCLA

Jacob Zhang*
UCLA

Abstract

We design and implement GigaDORAM, a novel 3-server Distributed Oblivious Random Access Memory (DORAM) protocol. Oblivious RAM allows a client to read and write to memory on an untrusted server, while ensuring the server itself learns nothing about the client’s access pattern. Distributed Oblivious RAM (DORAM) allows a group of servers to efficiently access a secret-shared array at a secret-shared index.

A recent generation of DORAM implementations (e.g. FLORAM, DuORAM) has focused on building DORAM protocols based on Function Secret-Sharing (FSS). These protocols have low communication complexity and low round complexity but linear computational complexity of the servers. Thus, they work for moderate-size databases, but at a certain size, these FSS-based protocols become computationally inefficient.

In this work, we introduce GigaDORAM, a hierarchical-solution-based DORAM featuring poly-logarithmic computation and communication, but with an over $100\times$ reduction in rounds per query compared to previous hierarchical DORAM protocols. In our implementation, we show that for moderate to large databases where FSS-based solutions become computation bound, our protocol is orders of magnitude more efficient than the best existing DORAM protocols. When $N = 2^{31}$, our DORAM is able to perform over 700 queries per second.

1 Introduction

To an outside observer, traditional encryption schemes can effectively hide the *contents* of memory, yet encryption alone does not hide the memory locations being accessed. In many cases, the *access pattern* of a file system can leak sensitive information, even when the contents are encrypted.

Oblivious Random Access Memory (ORAM), introduced by Goldreich and Ostrovsky [GO96] is a cryptographic protocol

that allows a client to read and write from memory while ensuring the *physical access pattern* (which is potentially observable to someone with sufficient access to the machine) is independent of the *virtual access pattern* (the underlying data retrieved by the client). Thus, when memory is accessed using an ORAM protocol, it is *mathematically provable* that an observer learns nothing about the client’s query pattern (beyond the number of queries).

Oblivious RAM was developed in a model where a single client wishes to store and retrieve sensitive data from an untrusted data store. Originally, the untrusted data store was conceptualized as untrusted RAM on the same machine as the client, but today we usually imagine a client storing and retrieving data from an untrusted cloud provider. In this setting, encryption can hide the *data* from the cloud provider, but ORAM is necessary to hide the *access pattern*¹. Thus, ORAM provides the strongest possible guarantee – hiding *both* the data and the access pattern.

Although ORAM was designed in the client-server setting, a slight variant of ORAM is also useful in the context of secure multiparty computation, where a group of servers needs to access a secret-shared array at a secret-shared location. In this setting, the secret sharing hides the *data*, but every participating server observes the physical access pattern. *Distributed Oblivious RAM (DORAM)* provides a method for efficiently accessing a secret-shared array at a secret-shared index, which in turn makes it possible to do secure multiparty computation (MPC) in the RAM model (RAM-MPC). Almost all existing MPC protocols work in the *circuit-model* where the desired computation is first converted to an arithmetic/boolean circuit before being executed securely. This conversion becomes costly (e.g. in the case of a private database) since every “random-access” is replaced with an $O(\text{memorySize})$ “MUX

¹Note that most Searchable Symmetric Encryption (SSE) schemes allow a client to efficiently query encrypted data stored in an untrusted cloud, but they typically do *not* hide the access pattern from the cloud provider. SSE schemes also target a different model, where data payloads may be of drastically different size (in ORAM all entries are of size D) and queries may return different numbers of “matches” [CGK06, Nav15].

*Authors are in alphabetical order.

operation.” RAM-MPC allows random-access programs to be executed securely without this costly conversion, which in turn enables much more flexible and efficient secure multi-party computation protocols [OS97, GKK⁺12].

1.1 Previous DORAMs

The efficiency of a (D)ORAM protocol is usually measured by the communication complexity of each query. Several recent DORAM protocols (e.g. [LO13, FNO22]) achieved ($O((\kappa + D) \log(N))$) communication and computation. These asymptotically-efficient protocols are built using the ORAM’s “*hierarchical solution*”, introduced in [Ost90, Ost92], which we describe in Section 4. The downside of these protocols is that they have high round complexity, requiring $O(\log(N))$ rounds of communication (with large hidden constants) per query. In practice, network latency causes bottlenecks in the performance of these solutions.

It has been noted that the high round complexity of the hierarchical solution makes it unsuitable for practical applications [WHC⁺14], so most DORAM implementations (e.g. [WCS15, SCSL11, GKK⁺12, ZWR⁺16, JW18, WHC⁺14, Ds17, JZLR22, VH22, BKKO20]) take a different approach. These constructions focus on minimizing rounds while compromising on either asymptotic computation or asymptotic communication costs. One common technique (most recently applied by DuORAM [VHG22]) for creating DORAM protocols with low round complexity *and* low communication complexity is Function Secret Sharing (FSS) [GI14, BGI15]. While FSS results in low *communication* complexity, FSS-based protocols require $O(N)$ computation, where N is the size of the database. Thus, protocols like DuORAM shines in high-latency, low-bandwidth networks at a small database size.

By contrast, our DORAM is designed for low-latency environments (such as co-located servers in the same data center). This design choice is present in several current deployments, such as Points of Presence (POP) of mutually distrustful ISP machines present in the same Data Center in support of BGP protocols [ACF⁺12, Smi21]. These mutually mistrustful yet co-located machines offer a viable alternative to data clean rooms, which present a single point of failure (if the security of the clean room is breached). As another example, Cybernetica’s commercial offering “Sharemind” MPC platform is intended to be run on three servers in nearby data centers [BKK⁺16]. In these low-latency environments, FSS-based DORAMs quickly become a bottleneck of performance for large values of N . We focus on this type of network environment because, these types of low-latency networks are necessary for RAM-MPC, which is one of the main motivations for building DORAM protocols in the first place.²

²Although MPC can run in high-latency environments [WRK17a], in a high-latency environment MPC can only compute simple functions that would not benefit from a RAM-model computation.

1.2 Our Contributions

1.2.1 Protocol contributions

In this work, we design and implement a high-performance DORAM protocol in the (3,1)-security model, i.e., where there are three semi-honest servers, and no two of them collude. Two effective techniques for building (D)ORAM protocols are “the hierarchical solution” (e.g. [AKL⁺20, LO13, AKLS21, FNO22]) which yields (D)ORAM protocols with low communication complexity, but high round complexity, and FSS-based DORAMs (e.g. FLORAM and DuORAM) which have very low round complexity, but high computational complexity. In this work, we show how to reduce the round complexity of the hierarchical solution and give round-efficient hierarchical DORAM that scales well beyond the limits of FSS-based (D)ORAM protocols. Specifically, our DORAM requires $O((\kappa^2 + D) \log N)$ communication and computation per query. At the cost of practically significant round-cost, we can easily tweak our protocol to match the best-known (D)ORAM complexity [AKL⁺20, AKLS21, LO13, FNO22] of $O((\kappa + D) \log N)$ (Remark 4.1). Our work is the first (D)ORAM protocol to achieve these asymptotics to be implemented.

1.2.2 Code contributions

Implementing our DORAM protocol required over 9,000 lines of custom C++ code. Additionally, we provide lightweight client implementation in Python and JavaScript, enabling users to utilize DORAM in the client-server setting.

In addition to the implementation of our DORAM protocol, we contribute (1) a from-scratch competitive implementation of the [AFL⁺16] general MPC framework which, in many settings, is the fastest known 3-party MPC protocol.³ (2) A custom (3,1)-garbled circuits protocol built using EMP-toolkit’s 2-party garbled circuits, which can be imported separately from ABY3’s [Rin] large framework, and (3) The first tested circuit files of the LowMC block cipher [ARS⁺15], featuring a novel optimization which reduces cache misses.

2 Preliminaries

Notation. We let N be the number of elements in the DORAM database, D the size in bits of each element, and κ be the computational security parameter (in practice $\kappa = 128$, in theory $\kappa = \omega(\log N)$), σ the statistical security parameter (in practice 2^{-40} or 2^{-80}).

Secret sharing. Our DORAM protocol makes heavy use of secret-sharing. Throughout this work, we use $[\cdot]$ to denote a “replicated” (or CNF [CDI05]) secret sharing. In a 3-party replicated sharing, a secret, x is split into three shares $x =$

³The implementation used to benchmark the results in the [AFL⁺16] is proprietary, and not publicly available.

$x_1 \oplus x_2 \oplus x_3$, and participant x gets *two* of the shares – every share except x_i . We use $[\cdot]^{(i,j)}$ to denote a simple XOR-2-sharing between participants P_i, P_j

Obliviousness: A computation is *data-oblivious* if its control-flow is independent of the input data. An Oblivious RAM protocol is a protocol for accessing an array (indexed by $1, \dots, N$), where the algorithm’s control flow, and in particular, the physical memory accessed, is independent of the index being queried. ORAM protocols are designed to allow a client to make a *any sequence* of queries and are often composed of simpler data structures that are only oblivious on distinct queries. A data structure (e.g. a hash table) is called *distinct-query oblivious* if the control flow between any two sequences of *distinct* queries is indistinguishable, but a sequence with repeated queries might result in a control flow that is distinguishable from a sequence of distinct queries.

Cuckoo Hashing: A cuckoo hash table is a distinct-query oblivious data structure for storing key-value pairs [PR01]. The cuckoo hash table consists of two arrays (“tables”) T_0 and T_1 of size $c \cdot n$, and two hash functions h_0 and h_1 , $h_i : X \rightarrow [cn]$. A key-value pair, (x, y) , can be stored in the table at location $h_0(x)$ in T_0 or $h_1(x)$ in T_1 . It is possible that there is no valid way to store a series of elements (e.g. if there are x_0, x_1, x_2 such that $h_0(x_0) = h_0(x_1) = h_0(x_2)$ and $h_1(x_0) = h_1(x_2) = h_1(x_3)$). In this case, we say there is a “build failure.” The probability of a build failure is $1/\text{poly}(n)$, but if we allow a small “stash” that can hold at most $O(\log(n))$ elements, the probability of a build failure becomes negligible in n [KMW09, Nob21]. In our DORAM (as in most prior hierarchical (D)ORAM protocols) we use cuckoo hash tables as a building block for our oblivious hash tables.

DORAM: A distributed Oblivious RAM protocol is a multiparty protocol that allows a group of participants holding a secret-shared array $[[x_1], \dots, [x_N]]$ to access the array at a secret-shared index, $[[i]]$, and obtain the sharing $[[x_i]]$, without revealing any information about the query, i , or the database x_1, \dots, x_n . The theoretical efficiency of a DORAM protocol is often measured by the amortized communication complexity the servers spend to respond to a single query. In practice, DORAM protocols may be bottlenecked by the amortized communication per query (e.g. [WCS15]), the amortized computation per query (e.g. [Ds17, VHG22]), or the amortized number of communication rounds per query (e.g. [LO13, FNO22]). To compare between these constructions targeting different points in the solution space, *we measure the practical efficiency of a DORAM protocol by the number of queries per second that it can process.*

SISO-PRFs: A core building block of most DORAM protocols is a Shared-Input, Shared-Output PRF (SISO-PRF). A SISO-PRF allows the participants to compute the secret-sharing of a PRF output on a shared input, under a shared key. Any regular PRF can be converted into a SISO-PRF by implementing the PRF under a generic MPC protocol. Several PRF protocols have been designed to be “MPC friendly” (e.g.

LowMC [ARS⁺15]). The basic idea which makes SISO-PRFs useful for DORAM is that servers can generate a random, shared key, $[[k]]$, and build a hash table where the cleartext tags (of secret-shared payloads) are SISO-PRF evaluations of the elements.

3 Secure Multiparty Computation

Secure multiparty computation (MPC) [Yao82, Yao86, GMW87, CCD88] is a protocol that allows a group of participants to securely compute a joint function on their private inputs *without* revealing any information beyond the output of the function. An MPC protocol is called (n, t) -secure if the protocol involves n participants, and remains secure if at most t participants collude (i.e., share private state). Our DORAM protocol works in the standard $(3, 1)$ semi-honest security model, which assumes that there are three semi-honest servers and no collusion between servers. We use the $(3, 1)$ -“replicated” MPC protocol of [AFL⁺16] as a building block. In one crucial place, we also use a custom implementation of the $(3, 1)$ garbled circuit MPC protocol of ABY3 [MR18] to reduce round complexity. Our DORAM protocol is also an important tool in *building* efficient RAM-MPC protocols, because it allows for MPC computation in the *RAM model* of computation, whereas most current MPC protocols work in the *circuit model*.

3.1 The Arithmetic Black-Box Model

Our DORAM protocol makes use of several “basic” operations on secret shared values, e.g. addition, comparisons, and equality tests. In our protocol descriptions, we use the Arithmetic Black Box (ABB) model to abstract away the underlying implementations of these operations. In practice, we use our own implementation of [AFL⁺16]. A formal description of the ABB model can be found in [GRR⁺16, FNO22].

In protocols, we use our ABB by invoking $\mathcal{F}_{\text{ABB}}.\text{FunctionalityName}$, where FunctionalityName makes it obvious what the functionality achieves. For instance, we invoke $[[z]] = \mathcal{F}_{\text{ABB}}.\text{Mult}([x], [y])$ to multiply secret shared values x, y and obtain secret shared value z s.t. $z = x \cdot y$. Although the names are usually self-explanatory, we provide a complete list of our ABB functionalities in Appendix A.

4 Construction Overview

Section 4.1 describes known techniques which enabled building communication and computation efficient “hierarchical” DORAM which takes *many* rounds of communication to execute a single query [FNO22, LO13]. In Section 4.2 we motivate and outline several novel techniques which enable us

to significantly reduce the round complexity of standard “hierarchical” DORAMs. In Section 9, we show that our round-reduced “hierarchical” DORAM, GigaDORAM, is efficient in practice.

4.1 The hierarchical solution

OHTable We call an efficient, oblivious-to-build, distinct-query oblivious, hash table an *OHTable*. Oblivious Hash Tables have three key functionalities: Build, Query, and Extract. $\mathcal{F}_{\text{OHTable}}.\text{Build}(X)$ creates an oblivious hash table storing the elements X , where each element in X is unique. Once the hash table has been built, $\mathcal{F}_{\text{OHTable}}.\text{Query}$ queries the table obliviously, and $\mathcal{F}_{\text{OHTable}}.\text{Extract}$ extracts all elements currently stored in the table which have *not* been queried.

The hierarchical solution. The key idea of the hierarchical solution is that it is fairly easy to build a distinct-query OHTable – for example, reading from a cuckoo hash table is oblivious as long as you never query the same element twice. With this insight, the hierarchical solution [Ost92] can be seen as a compiler for bootstrapping a distinct-query OHTable construction into a full-fledged DORAM (that remains oblivious even if the user makes repeated queries). The hierarchical solution is a powerful tool, used to build many ORAM protocols, e.g. [Ost90, Ost92, GO96, GMOT12, KLO12, LO13, PPRY18, AKL⁺20, FNO22], and is the technique used by many of the most (asymptotically) communication efficient (D)ORAM protocols.

A hierarchical ORAM is made up of a hierarchy of OHTables, $L_0, \dots, L_{\text{numLevels}}$, of geometrically increasing size. Usually, we have $\text{numLevels} = O(\log N)$ and the largest level, $L_{\text{numLevels}}$, has size $O(N)$. The smallest level, L_0 , is a small “cache”. The cache itself needs to remain oblivious even if queries are repeated. But the cache is small, so it can be implemented inefficiently without dramatically increasing the cost of a given ORAM query. For this reason, the cache is often set to be constant sized ($|L_0| = O(1)$). This means that the client can read the *entire* cache with each query (in time $O(|L_0|)$) and this is clearly oblivious. Each larger level, L_i for $i \geq 1$ holds a distinct-query OHTable of size $O(2^i)$.

When a user queries the DORAM, the user queries each level of the hierarchy sequentially. If the item is found at level i , a “dummy” element is queried at subsequent levels $L_{i+1}, \dots, L_{\text{numLevels}}$, and the retrieved item is reinserted into the cache. To maintain obliviousness, if the requested item is not found anywhere in the hierarchy, a dummy item is inserted in the cache. When the cache (or any subsequent level of the hierarchy) is full, all the elements from that level (and smaller levels) are extracted and rebuilt into the next level of the hierarchy. If $|L_{i+1}| \geq 2 \cdot |L_i|$, then level $i+1$, can accommodate all the elements in levels $0, 1, \dots, i$. This periodic rebuild schedule guarantees that no element is ever queried twice at the same level between rebuilds, because once an element has been found at level i , it is moved to the cache, and will always

stay in a level $j < i$, until level i is rebuilt. Since no element will be queried twice in any given distinct-query OHTable, the entire construction is oblivious, even if queries are repeated. The formal proof that the resulting hierarchical data structure is indeed an ORAM is now standard and can be found for example in [Ost92]. With this rebuild schedule, level L_i is rebuilt every $O(|L_i|)$ queries. We outline our instantiation of the hierarchical protocol in Figure 4 in Section 8.

4.2 Reducing numRoundsDORAM

Our protocol is based on a 3-party implementation of the hierarchical solution (similar to [FNO22]). Although the hierarchical solution has low asymptotic communication and computation complexity, it has high *round* complexity which can make it inefficient in practice.⁴ This inefficiency stems from the fact that every query to the DORAM forces a query into *every* level of the hierarchy, $L_0, \dots, L_{\text{numLevels}}$. Critically, these queries must be *sequential*, because if the item is found at L_i , the protocol must query dummy elements at $L_{i+1}, \dots, L_{\text{numLevels}}$. Thus the *round complexity* of a DORAM query⁵, numRoundsDORAM , can be written as

$$\text{numRoundsOHTable} \cdot \text{numLevels} + \text{numRoundsCache} \quad (1)$$

where numRoundsOHTable is the number of rounds to query L_i for $1 \leq i \leq \text{numLevels}$ (which is fixed and independent of i) and numRoundsCache is the number of rounds it takes to query the cache, L_0 . The main technical contribution of our work is to optimize the hierarchical solution to reduce its round complexity.

In light of Equation 1, we can divide our efforts to reduce round complexity into four parts: (1) *reducing numRoundsOHTable via our novel OHTable, ShufTable, and SISO-PRF parallelization*, (2) *generalizing the hierarchical solution with a tuneable parameter, baseAmpFactor, to reduce numLevels*, (3) *designing a new oblivious-cache data-structure, SpeedCache, to reduce numRoundsCache*, and (4) *applying additional engineering/implementation-level optimizations*.

With these design improvements, we are able to significantly reduce the round complexity of the hierarchical solution, while maintaining similar asymptotic overheads to the theoretic state-of-the-art, [FNO22]. In addition to its low asymptotic complexity and low rounds-per-query, we show that our DORAM design is actually quite fast in practice (Section 9).

⁴The (3,1)-DORAM of [FNO22] and the 2-server ORAM of [LO13] use the hierarchical solution to achieve amortized communication complexity of $O((\kappa + D) \log N)$. We remark, however, that both these protocols have high round complexity, and have never been implemented.

⁵DORAM also incurs round costs when building a level L_i and extracting from a level L_i . Since the protocols are only invoked every $|L_i|$ queries and have small, constant, round costs (where our final protocol has $|L_0| = O(\kappa)$), their round cost has negligible impact on the performance of DORAM.

We briefly explain each of our optimizations below and expand on them in Section 5, 6, and 7 respectively.

(1) Reducing numRoundsOHTable: ShufTable & SISO-PRF parallelization. To reduce numRoundsOHTable we present a novel, standalone OHTable, called *ShufTable*, with reduced round complexity. We also devise a method to parallelize the sequentially-dependent SISO-PRF evaluations needed to query each level of the hierarchy.

The key ideas in these optimizations revolve around how to handle queries for elements that *are not in the table*. This type of query happens frequently in a hierarchical ORAM because although each element is only stored at one level of the hierarchy, each ORAM query results in a query to the OHTable at every level of the hierarchy. Prior solutions for handling these “dummy” queries were round-intensive (e.g. the Oblivious Sets in [FNO22]), so we develop a new method.

As in prior works, our OHTable, *ShufTable* inserts dummy elements $(d_1, \perp) \dots, (d_l, \perp)$ along with the real elements, and retrieves a dummy when the queried element is not stored in the table. Assuming the SISO-PRF has already been evaluated, *ShufTable* requires only 5 rounds of server-to-server communication per query. The main ingredient that enables this round-savings is a novel “persistent shuffling” trick (Section 5.1)⁶ which allows to efficiently evaluate a random permutation under MPC. Our OHTable, like those in prior DORAM constructions, requires an equality check on secret-shared values. We implement this equality check using a custom implementation of a 3-party garbled circuit [MR18]. This *increases* the asymptotic communication of each level’s query from $O(\kappa)$ to $O(\kappa^2)$, but *decreases* the round complexity at each level by $O(\log \kappa)$. In practice, when $\kappa \in \{128, 256\}$, this dramatically improves real-world performance.

More details can be found in section 5.2 and the protocol is given in Figure 2.

Remark 4.1. If one were focused on optimizing asymptotic communication complexity, this step could be replaced by a constant-overhead-MPC-based equality check (e.g. using [AFL⁺16]), which would give us best known (same as [FNO22, LO13]) asymptotic communication and computation complexity $O((\kappa + D) \log)$.

In the hierarchical solution, each query into the OHTable at level i , requires evaluating a SISO-PRF, but the query into table L_{i+1} depends on whether the item was found at a lower level. Thus, hierarchical-based (D)ORAMs wait until after L_i . Query is performed to evaluate the SISO-PRF for L_{i+1} .

With *ShufTable* and most other OHTables, there are essentially two types of queries that can be made at any level a “real” query (if the desired element has not yet been found) or a “dummy” query (if the element was found at a lower level). A simple observation we make is that we can evaluate $\llbracket r_i \rrbracket$, the PRF evaluation needed to make a real query at L_i , and

$\llbracket d_i \rrbracket$, the PRF evaluation needed to make a dummy query at L_i , *in parallel for all $i \in [\text{numLevels}]$* . Then, after L_i is evaluated and it is determined under MPC whether the queried element was found, we can multiplex between $\llbracket r_{i+1} \rrbracket$ and $\llbracket d_{i+1} \rrbracket$ under MPC, which takes only one round. Since the outputs of the SISO-PRF are secret-shared and since multiplexing in MPC hides whether the table is queried for a dummy or a real element, the DORAM remains oblivious. This optimization decreases the rounds per query dedicated to evaluating the SISO-PRF from numRoundsPRFEvalnumLevels to numRoundsPRFEval/ + numLevels.

Unfortunately, the above change increases communication per query by $\approx 2\times$. To resolve this issue, we upgrade *ShufTable* with a “just-in-time” mechanism to detect if an element is stored in the table, replacing the retrieved element with a dummy if necessary. Due to this mechanism, only $\llbracket r_i \rrbracket$ is needed to query the OHTable at L_i .

In practice this optimization is noticeable, saving ~ 100 rounds of interaction per query.

To further reduce rounds, in practice, we parallelize the query of the cache, L_0 with the evaluation of the PRFs for $L_1, \dots, L_{\text{numLevels}}$. Note that if we knew multiple queries in advance, we would be able to further batch evaluations, but we do not assume that in this work.

Overall, these optimizations have a significant impact of the efficiency of our protocol. We decrease numRoundsOHTable from the 45 of [FNO22] to amortized $5 + \text{numRoundsPRFEval}/\text{numLevels} \approx 7$. Additionally, we reduce the number of expensive SISO-PRF evaluations necessary by a factor of four when compared to [FNO22] which is the state-of-the-art in low SISO-PRF hierarchical DORAMs.

(2) Generalizing the hierarchical solution to reduce numLevels. Above, we described our techniques for reducing the round complexity of queries to individual OHTables. Yet in the hierarchical solution, every ORAM query requires querying each level of the hierarchy *sequentially*. Thus a hierarchy of depth numLevels immediately adds a multiplicative factor of numLevels to the round complexity of the protocol. Since round complexity is one of the main performance bottlenecks in (D)ORAM protocols, there is a strong motivation to reduce numLevels.

In most hierarchical (D)ORAMs, level i in the ORAM hierarchy had size $2^i \cdot |L_0|$, resulting in numLevels = $O(\log_2 N)$. In Section 6 we show that by introducing a new parameter, baseAmpFactor > 2, and setting $|L_{i+1}| = \text{baseAmpFactor} \cdot |L_i|$, we can reduce the *round complexity* of the protocol with minimal impact on the *communication* complexity, which we find to be less expensive in practice. This simple change immediately reduces numLevels from $\log_2(N)$ to $\log_2(N)/\log_2(\text{baseAmpFactor})$. While this modification is conceptually simple, it requires a more nuanced rebuild schedule.

Empirically, we find that the optimal value for baseAmpFactor is much greater than 2. For instance,

⁶ [LWZ11] presents various efficient shuffling protocols, but none of their shuffles allows for the efficient evaluation of a reverse permutation at a point.

at $N = 2^{30}$ and the network conditions we test, we found that `baseAmpFactor = 128` is optimal. That is, each level is 128 times larger than the previous (smaller) level.

(3) Optimizing the cache: SpeedCache. In the hierarchical ORAM solution, the top level (i.e. L_0 , the “cache”) needs to support oblivious accesses (compared to lower levels in the hierarchy which only need to be *distinct-query oblivious*). For this reason, L_0 is usually implemented as a simple read-all-to-read, append-to-write array. This is obviously oblivious, but its query complexity increases linearly with the size of the cache. In particular, if the cache stores t key-value pairs ($(\llbracket x_1 \rrbracket, \llbracket y_1 \rrbracket), \dots, (\llbracket x_t \rrbracket, \llbracket y_t \rrbracket)$), querying the cache is often implemented by sequentially checking whether the query, $\llbracket x \rrbracket$, is equal to $\llbracket x_i \rrbracket$ (costs $O(\log \log |x_i|)$ sequential rounds) and if so updating return value to $\llbracket y_i \rrbracket$. Unfortunately, this simple implementation has multiplicative depth t , meaning $\text{numRoundsCache} = O(|L_0| \cdot \log \log |x_i|)$.

In Section 7, we outline a simple Cache protocol SpeedCache that allows us to query the cache in $\lceil \log \log |x_i| \rceil + 1$ rounds of communication (which is *independent of* $|L_0|$). Since our SpeedCache protocol has a round complexity that is independent of the cache size, we are somewhat free to increase the cache size, which has other benefits (e.g. reducing numLevels by $\log_2 |L_0|$).

(4) Gadget implementations: Minimizing the round complexity of our SISO-PRF is crucial for the overall efficiency of our protocol, so we provide the first circuit file of the LowMC [ARS⁺15] block cipher within our custom MPC implementation. Our circuit file features a novel optimization we call “wire threading” that allows us to reduce the number of L1-cache misses during evaluation. We multithread the MPC evaluation of LowMC, allowing us to evaluate the PRF 6.7M times per second. See Appendix C for further details.

We adopt the Alibi reinsertion technique [FNO21] for “caching the stash” to the distributed setting, and we provide the first implementation of Alibi. See Appendix D for further details.

5 ShufTable: reduce numRoundsOHTable

In this section, we present ShufTable, a novel, oblivious distinct-query, hash table (OHTable). With ShufTable we have $\text{numRoundsOHTable} \approx 7$ (compared to [FNO22] which required 45 rounds of communication for each OHTable query).

In Section 8 we use the hierarchical solution and [FNO21] to compile ShufTable into an efficient DORAM. The resulting DORAM requires only a constant number of rounds per query, whereas prior hierarchical DORAM protocols (e.g. [FNO22]) had a round complexity that scaled with N .

In Section 5.1 we present the “Persistent shuffle protocol,” a method that enables $\Pi_{\text{ShufTable}}$, the new OHTable protocol we present in Section 5.2. Leveraging features of

$\Pi_{\text{ShufTable}}$. Query, in Section 5.3 we show how to amortize the round cost of SISO-PRF evaluations across $L_1, \dots, L_{\text{numLevels}}$.

5.1 Persistent shuffle Protocol

Like many DORAM protocols, our DORAM construction relies on an efficient oblivious shuffle, which allows players holding a secret shared list, $\llbracket X \rrbracket = \llbracket X_1 \rrbracket, \dots, \llbracket X_n \rrbracket$ to shuffle $\llbracket X \rrbracket$ under some random permutation $\pi \in S_n$ such that *nothing about π is learned by any player*.

As presented in [LWZ11], there is an efficient, linear communication (3,1)-oblivious shuffle that works as follows. The players P_1, P_2, P_3 reshare $\llbracket X \rrbracket$ to P_1, P_2 , who shuffle their secret shares according to some random agreed upon permutation, and reshare the shuffled list to P_2, P_3 , who shuffle and reshare to P_3, P_1 who shuffle and reshare to P_1, P_2, P_3 . Since the composition of permutations is not known to any single player, the final permutation is oblivious. Unfortunately, at the end of the [LWZ11] protocol, information about the permutation, π , is not accessible to the players. In our Persistent shuffle (described in Figure 1), we augment the shuffling protocol to also output a secret sharing of π .

Cost of the Persistent shuffle protocol. The Persistent shuffle protocol requires 4 rounds of communication and each round requires $n(|X_i| + \log n)$ bits of communication. For comparison, shuffling n elements using Persistent shuffle is 6 times less bandwidth than evaluating a SISO-PRF (using LowMC) on those elements, so shuffling contributes only minimally to the overall communication cost of our DORAM.

Lemma 5.1. *Persistent shuffle (described in Figure 1 fulfills the functionality \mathcal{F}_{ABB} .ObliviousShuffle(\circ , DistributeShuffle = True) where the players input $\llbracket X \rrbracket$ and receive as output $\llbracket \pi(X) \rrbracket$ and $\llbracket K \rrbracket = \llbracket K_1 \rrbracket, \dots, \llbracket K_n \rrbracket$ s.t $\pi(i) = K_i$. It also guarantees that π is uniformly sampled from S_n and unknown to all players.*

Proof. In step 1, the 3 participants reshare the shares of the vector X to and the indices $L = (1, 2, \dots, n)$ to two participants. By the security of \mathcal{F}_{ABB} .ReshareReplicatedTo2Sharing, each individual participant learns nothing about X (in fact each player’s shares are uniformly random). In steps 2 & 3, pairs of players locally shuffle their shares and reshare to the next pair. In this case, the security of \mathcal{F}_{ABB} .Reshare2To2Sharing ensures that each player learns nothing about the underlying data (the shuffled values). As in step 1, each player receives uniformly random shares and nothing else. In step 4, pairs of players call \mathcal{F}_{ABB} .Reshare2SharingToReplicated, and again, each player’s view of the protocol is a collection of uniformly random shares.

The final list is a sharing of $\pi(X)$, where $\pi = \pi_{\{3,1\}} \circ \pi_{\{2,3\}} \circ \pi_{\{1,2\}}$. Every player knows two of the three permutations, so the resulting permutation is uniformly random from the perspective of any player. \square

Setup: Each pair of players P_i, P_j agree on a random permutation $\pi_{\{i,j\}} \in S_n$. The players also generate a sharing $\llbracket L \rrbracket = \llbracket 1 \rrbracket, \dots, \llbracket n \rrbracket$.

Protocol:

1. The players reshare to the first shufflers, calling

$$[X]^{(1,2)} = \mathcal{F}_{\text{ABB}}.\text{ReshareReplicatedTo2Sharing}(\llbracket X \rrbracket, \{P_1, P_2\})$$

and

$$[L]^{(3,1)} = \mathcal{F}_{\text{ABB}}.\text{ReshareReplicatedTo2Sharing}(\llbracket L \rrbracket, \{P_3, P_1\})$$

2. Shuffle & Reshare to next shuffling pair #1:

- (a) P_1, P_2 let $[X']^{(1,2)} = [\pi_{\{1,2\}}(X)]^{(1,2)}$. Note that since P_i, P_j hold $[X]^{(i,j)}$ they can obtain $[\pi(X)]^{(i,j)}$ for a known π by locally shuffling their list of secret shares. P_1, P_2 reshare X' to the next shufflers, calling

$$[X'']^{(2,3)} = \mathcal{F}_{\text{ABB}}.\text{Reshare2To2Sharing}([X']^{(1,2)}, \{P_2, P_3\}).$$

- (b) P_3, P_1 let $[L']^{(3,1)} = [\pi_{\{3,1\}}^{-1}(L)]^{(3,1)}$. P_1, P_2 reshare L' to the next shufflers, calling

$$[L'']^{(2,3)} = \mathcal{F}_{\text{ABB}}.\text{Reshare2To2Sharing}([L']^{(3,1)}, \{P_2, P_3\})$$

3. Shuffle & Reshare next shuffling pair #2:

- (a) P_2, P_3 let $[X'']^{(2,3)} = [\pi_{\{2,3\}}(X'')]^{(2,3)}$. P_2, P_3 reshare X'' to the next shufflers, calling

$$[X''']^{(3,1)} = \mathcal{F}_{\text{ABB}}.\text{Reshare2To2Sharing}([X'']^{(2,3)}, \{P_3, P_1\}).$$

- (b) P_2, P_3 let $[L'']^{(2,3)} = [\pi_{\{2,3\}}^{-1}(L'')]^{(2,3)}$. P_2, P_3 reshare L'' to the next shufflers, calling

$$[L''']^{(1,2)} = \mathcal{F}_{\text{ABB}}.\text{Reshare2To2Sharing}([L'']^{(2,3)}, \{P_1, P_2\}).$$

4. Shuffle & Reshare back to all players #3:

- (a) P_3, P_1 let $[X''']^{(3,1)} = [\pi_{\{3,1\}}(X''')]^{(3,1)}$. P_3, P_1 reshare X''' back to the entire group, calling

$$\llbracket X'''' \rrbracket = \text{Reshare2SharingToReplicated}([X''']^{(3,1)}).$$

- (b) P_1, P_2 let $[L''']^{(1,2)} = [\pi_{\{1,2\}}^{-1}(L''')]^{(1,2)}$. P_1, P_2 reshare L''' back to the entire group, calling

$$\llbracket L'''' \rrbracket = \text{Reshare2SharingToReplicated}([L''']^{(1,2)}).$$

Output: The output of the protocol is $\llbracket \pi(X) \rrbracket, \llbracket \pi^{-1}(L) \rrbracket$.

Figure 1: The Persistent shuffle protocol, $\mathcal{F}_{\text{ABB}}.\text{ObliviousShuffle}(\llbracket X \rrbracket, \text{DistributeShuffle} = \text{True})$

5.2 ShufTable construction

In this section, we describe and evaluate $\Pi_{\text{ShufTable}}.\text{Build}$ and $\Pi_{\text{ShufTable}}.\text{Query}$, the component of our new Oblivious Hash Table ShufTable. The pseudocode for Protocol $\Pi_{\text{ShufTable}}$ is presented in Figure 2.

Parameters. $\Pi_{\text{ShufTable}}$ is parameterized by N , the size of the address space, D , the size of each payload, κ , the computational security parameter, n , the number of real elements stored in the table, numDummies , the number of dummy elements stored in the table, and stashSize , the minimum size of a stash in the cuckoo hash table storing n elements such that the probability of a build failure is less than our statistical security parameter, σ .

On each query, we must be able to retrieve either (1) the element being queried or (2) a new dummy element that has never been queried (if the desired element is not in the OHTable). Hence, the number of queries to each ShufTable is bounded by $\min\{\text{numDummies}, n\}$. Thus we set $\text{numDummies} \approx n$.

Since our construction uses a Cuckoo hash table (CHT), it is also implicitly parameterized by c , the number of slots in the CHT table, and t , the number of such tables. Instantiations of these variables are discussed in Section 9.

Input-Output behavior of $\Pi_{\text{ShufTable}}.\text{Build}$. The parties, P_1, P_2, P_3 , input $E = \{(\llbracket X_1 \rrbracket, \llbracket Y_1 \rrbracket), \dots, (\llbracket X_n \rrbracket, \llbracket Y_n \rrbracket)\}$ where $n \leq N$ to $\Pi_{\text{ShufTable}}.\text{Build}$. Protocol $\Pi_{\text{ShufTable}}.\text{Build}$, outputs secret shares of CHT held by P_2, P_3 , $[\text{CHT}]^{(2,3)}$, and $[\hat{X}], [\hat{Y}]$, special shuffling of the elements under a permutation somewhat known to P_2, P_3 (Step 7 of $\Pi_{\text{ShufTable}}.\text{Build}$). The data structure, ShufTable, stores a subset of the elements sent to the build protocol $A \subset E$. $\Pi_{\text{ShufTable}}.\text{Build}$ also outputs $\text{Stash} = E \setminus A$, with $|\text{Stash}| \leq \text{stashSize}$. In the greater DORAM protocol (Section 8), Stash will be inserted into the cache, L_0 . As long as X_1, \dots, X_n are distinct (as assured by the hierarchical solution), the parties learn nothing of $(X_1, Y_1), \dots, (X_n, Y_n)$ or which elements belong to Stash.

Input-Output behavior of $\Pi_{\text{ShufTable}}.\text{Query}$. For the t 'th query where $1 \leq t \leq \text{numDummies}$, the players input $\llbracket x_t \rrbracket$ to $\Pi_{\text{ShufTable}}.\text{Query}$. $\Pi_{\text{ShufTable}}.\text{Query}$, outputs $\llbracket Y_j \rrbracket$ if $x_t = X_j$ for $(X_j, Y_j) \in A$ and $\llbracket \perp \rrbracket$ otherwise (and outputs $\llbracket \text{found} \rrbracket$ accordingly). The security guarantee is that for any sequence of *distinct* queries, x_1, \dots, x_t , P_1, P_2, P_3 do not learn whether Y_j or \perp was outputted, j , or Y_j .

Performance analysis. The cost of $\Pi_{\text{ShufTable}}.\text{Build}$ is dominated by the cost of n SISO-PRF evaluations.⁷ With our implementation of LowMC, this requires a total of $2304n$ bits of communication at $\kappa = 128$ and the computation of many XORs. Despite this cost, $\Pi_{\text{ShufTable}}.\text{Build}$ is very efficient, since via multi-threading LowMC (Section C) we are able to evaluate 6.7M SISO-PRFs/s in our tests.

The time needed to evaluate $\Pi_{\text{ShufTable}}.\text{Query}$ is driven by its round complexity since we must sequentially invoke this

⁷For comparison, [FNO22] requires more than twice as many SISO-PRF evaluations ($2n + \text{numDummies}$) for each OHTable build.

protocol numLevels times during $\Pi_{\text{DORAM}}.\text{Query}$. Each step of the query requires equality checks on secret shared elements, which we execute using 3-party garbled circuits.⁸

Lemma 5.2. $\Pi_{\text{ShufTable}}$ implements the distinct-query oblivious hash table functionality.

Proof. We begin by showing that the view of each player during $\Pi_{\text{ShufTable}}.\text{Build}$ is independent of the input data X, Y . During $\Pi_{\text{ShufTable}}.\text{Build}$, P_2, P_3 only receive secret shares and operate on those secret shares using MPC. Since secret shares are sampled from a (individually) uniformly random distribution, P_2, P_3 's views are independent of underlying data X, Y .

During $\Pi_{\text{ShufTable}}.\text{Build}$, other than operating on secret shares, P_1 receives the list \hat{Q} in the clear. By assumption, X has no repetitions, so the PRF “tags” Q appear uniformly random and independent to a computationally bounded adversary. Since \hat{Q} is a uniformly random shuffle of Q which P_1 does not know, P_1 cannot decipher *any* information about X from seeing \hat{Q} . Moreover, when deciding which elements to place in Stash (by deciding on which indexes to place there), P_1 cannot tell which elements from (X_i, Y_i) he is placing in stash (except that they aren't dummies since $\hat{Q}_i \neq \perp$). Thus, P_1 's view is independent of X, Y , and the stash contains uniformly random real elements.

Next, we show that in $\Pi_{\text{ShufTable}}.\text{Query}$ the view of each player is independent of the query X_t , the output, or X, Y previously stored. During $\Pi_{\text{ShufTable}}.\text{Query}$, other than operating on secret shares, P_2, P_3 learn q in the clear. q is either a random or a pseudorandom element which is thus computationally independent from P_2, P_3 's view (because P_2, P_3 did not see \hat{Q} which P_1 used to build $\text{CHT} \cup \text{Stash}$). Thus, P_2, P_3 do not learn anything by seeing q . Additionally, P_2, P_3 learn l , an index into the \wedge -shuffled list which depends on X_t 's presence at the level (see step 4 of $\Pi_{\text{ShufTable}}.\text{Query}$). Yet, since l is entirely determined by π_\wedge which P_2, P_3 don't know, and moreover, since they do not know which indexes of \hat{Y} correspond to dummies, l is sampled from a distribution indistinguishable from random relative to the individual view of P_2, P_3 .

Finally, during $\Pi_{\text{ShufTable}}.\text{Query}$, other than operating on secret shares, P_1 learns l in the clear. Yet, although P_1 knows which elements from \hat{Y} are dummies, due to an additional oblivious shuffle, he does not know which elements of \hat{Y} are dummies (or were stashed, or were stored in the table, etc). Hence l is independent of P_1 's view.

Thus, on distinct queries and when $i \neq j \implies X_i \neq X_j$, P_1, P_2 , and P_3 's views are independent of the data stored and queried, and hence ShufTable is distinct-query oblivious. \square

⁸We use 3-party garbled circuits rather than the (constant-round) 3-party BMR protocol [BMR90] implemented in EMP because 3-party garbled circuits are significantly more bandwidth efficient than BMR. 3-party garbled circuits require an honest majority, while BMR does not.

$\Pi_{\text{ShufTable}}.\text{Build}$: The players (refers to P_1, P_2, P_3) hold $(\llbracket X_1 \rrbracket, \llbracket Y_1 \rrbracket), \dots, (\llbracket X_n \rrbracket, \llbracket Y_n \rrbracket)$.

1. The players create dummies to satisfy queries that were not found, letting $\llbracket Y_i \rrbracket = \llbracket \perp \rrbracket, \llbracket X_i \rrbracket = \llbracket \perp \rrbracket$ for $i \in \{n+1, \dots, n + \text{numDummies}\}$.
2. The players generate a κ -bit secret-shared PRF key, evaluating $\llbracket k \rrbracket = \mathcal{F}_{\text{ABB}}.\text{RandomElement}(\kappa)$.
3. The players create pseudorandom tags for all the addresses, evaluating $\llbracket Q_i \rrbracket = \mathcal{F}_{\text{ABB}}.\text{PRFEval}(\llbracket k \rrbracket, \llbracket X_i \rrbracket)$ for $i \in \{1, \dots, n\}$ and $\llbracket Q_i \rrbracket = \llbracket \perp \rrbracket$ for $i \in \{n+1, \dots, n + \text{numDummies}\}$.
4. Players obliviously shuffle the lists before revealing Q to P_1 to hide information about the elements in the table/stash, executing $\llbracket \hat{Q} \rrbracket, \llbracket \hat{X} \rrbracket, \llbracket \hat{Y} \rrbracket, \llbracket j \rrbracket = \mathcal{F}_{\text{ABB}}.\text{ObliviousShuffle}(\llbracket Q \rrbracket, \llbracket X \rrbracket, \llbracket Y \rrbracket, \text{DistributeShuffle} = \text{True})$ (c.f. Section 5).
 - (a) The players locally create $\llbracket \text{DI}_i \rrbracket = \llbracket j_{n+i} \rrbracket$ for all $i \in [\text{numDummies}]$. If $\text{DI}_i = k$, that implies $X_k = n+i$ and $Y_k = \perp$. That is, the k th element of (\hat{X}, \hat{Y}) is the i th dummy. We use DI in step 4 of Query to return a dummy if needed.
5. Revealing Q_i to P_1 so that it can build a $\text{CHT} \cup \text{Stash}$ containing X_i using fast, local random accesses without learning X_i , the players call $\hat{Q}_i = \mathcal{F}_{\text{ABB}}.\text{RevealTo}(P_1, \llbracket \hat{Q}_i \rrbracket)$ for $i \in \{1, \dots, n + \text{numDummies}\}$.
6. P_1 locally constructs the Cuckoo hash table and list of stashed indices $\text{CHT} \cup \text{Stash} = \text{BuildCHTWs}(\llbracket \hat{Q}_i \rrbracket \parallel i)_{i \in [n]}, \text{stashSize})$. CHT stores $\hat{Q}_i \parallel i$ for $n - \text{stashSize}$ such different i 's where \parallel represents bit-wise appending.
 - (a) P_1 secret shares CHT between P_2 and P_3 , running $\llbracket \text{CHT} \rrbracket^{(2,3)} = \mathcal{F}_{\text{ABB}}.\text{InputTo2Sharing}(\text{CHT}, P_2, P_3)$. P_2, P_3 will use CHT to satisfy queries.
 - (b) P_1 sends Stash, a cleartext list of stashSize -many indexes, s.t. $i \in \text{Stash}$ indicates that $(\llbracket \hat{X}_i \rrbracket, \llbracket \hat{Y}_i \rrbracket)$ is stashed. If $i \in \text{Stash}$ then $X_i \neq \perp$. **Output** $X^{\text{stash}} = \{\llbracket \hat{X}_i \rrbracket\}_{i \in \text{Stash}}$ and $Y^{\text{stash}} = \{\llbracket \hat{Y}_i \rrbracket\}_{i \in \text{Stash}}$. These will be reinserted into the cache when $\Pi_{\text{ShufTable}}.\text{Build}$ is called as part Π_{DORAM} (Figure 4).
7. The players shuffle the data under a permutation, $\hat{\pi}$, only known to P_2 and P_3 , executing $\llbracket \hat{X} \rrbracket \llbracket \hat{Y} \rrbracket = \mathcal{F}_{\text{ABB}}.\text{ObliviousShuffle}(\llbracket X \rrbracket, \llbracket Y \rrbracket, \text{RevealTo} = \{2, 3\})$. This ‘‘rebalances the information asymmetry,’’ allowing P_2, P_3 to guide P_1 to respond to queries (see step 5 of Query) s.t. P_1 cannot use his privileged information from $\Pi_{\text{ShufTable}}.\text{Build}$, not learning anything about the query.

$\Pi_{\text{ShufTable}}.\text{Query}$: The players hold $\llbracket Q_{\text{query}} \rrbracket = \mathcal{F}_{\text{ABB}}.\text{PRFEval}(\llbracket X_i \rrbracket, \llbracket k \rrbracket)$ (Section 5.3). *Each step corresponds to a single round of communication.* We pack parallelizable/silent instructions into the same step.

1. First, the players compute $\llbracket r \rrbracket = \mathcal{F}_{\text{ABB}}.\text{RandomElement}(\kappa)$ (silent generation of random κ -bit secret share) and then compute $\llbracket q \rrbracket = \llbracket Q_{\text{query}} \rrbracket + \llbracket \text{useDummy} \rrbracket \cdot \llbracket r \rrbracket$ (one MPC multiplication). useDummy is an artifact of the hierarchical solution that indicates if X_i was already found in the previous level.
2. The players reveal enable P_2, P_3 to query $\llbracket \text{CHT} \rrbracket^{(2,3)}$ by revealing q to them, evaluating, $q = \mathcal{F}_{\text{ABB}}.\text{RevealTo}(\llbracket q \rrbracket, P_2, P_3)$.
3. P_2, P_3 input $\llbracket \text{CHT}[h_1(q)] \rrbracket^{(2,3)}, \llbracket \text{CHT}[h_2(q)] \rrbracket^{(2,3)}$ their secret-shares of locations in CHT where q might be stored, calling $\llbracket q'_b \rrbracket \parallel \llbracket i'_b \rrbracket = \mathcal{F}_{2\text{Share}}.\text{Reshare2to3WithoutCheck}(\llbracket T[h_b(Q)] \rrbracket)$ for $b = 1, 2$.
4. The players evaluate QueryCircuit using (3,1) garbled circuits (see [MR18]), evaluating $l = \mathcal{F}_{\text{ABB}}.\text{EvalCircuit}(3,1)\text{GC}(\text{QueryCircuit}, \text{Inputs} = \llbracket q'_1 \rrbracket, \llbracket i'_1 \rrbracket, \llbracket q'_2 \rrbracket, \llbracket i'_2 \rrbracket, \llbracket q \rrbracket, \llbracket \text{DI}_l \rrbracket)$, revealing l only to P_2, P_3 . QueryCircuit returns i'_1 if $q'_1 = q$, returns i'_2 if $q'_2 = q$, and returns DI_l otherwise.
5. P_2, P_3 send $j = \hat{\pi}(l)$ to P_1 . The parties set $\llbracket Y_{\text{output}} \rrbracket = \llbracket \hat{Y}_j \rrbracket$ and append j to list queriedDbhAtIdxs

$\Pi_{\text{ShufTable}}.\text{Extract}$: Output $(\llbracket \hat{X}_i \rrbracket, \llbracket \hat{Y}_i \rrbracket)$ for $i \in ([n + \text{numDummies}] - \text{queriedDbhAtIdxs})$

Figure 2: $\Pi_{\text{ShufTable}}.\text{Build}$ and $\Pi_{\text{ShufTable}}.\text{Query}$.

5.3 Parallelizing sequential SISO-PRF evaluations

In the hierarchical ORAM construction, each query requires searching *every* level in the hierarchy, and the OHTable query at a given level requires evaluating a SISO-PRF.

In previous constructions (e.g. [LO13,FNO22]) these SISO-PRF were evaluated sequentially, because when performing the OHTable query for an index, x , at level i , the SISO-PRF input will be $\llbracket x \rrbracket$ or a dummy element depending on whether x was found in a smaller level of the hierarchy.

Since there are only two possible SISO-PRF inputs at each level, rather than evaluating the PRF sequentially, the players could evaluate *both* the “dummy query”

$$\llbracket a_i \rrbracket = \mathcal{F}_{\text{ABB}}.\text{PRFEval}(\llbracket N + t_i \rrbracket, \llbracket k \rrbracket)$$

and the “real query”

$$\llbracket b_i \rrbracket = \mathcal{F}_{\text{ABB}}.\text{PRFEval}(\llbracket X_{\text{query}} \rrbracket, \llbracket k \rrbracket)$$

for all $i \in \{1, \dots, \text{numLevels}\}$ in parallel before querying L_1 , then multiplex the result using MPC (which costs only a single round) before evaluating L_i . Query.

This trick will reduce the *round* complexity but requires *doubling* the number of SISO-PRF calls per query. Since round complexity is often the bottleneck, this improves practical performance.

In our protocol, however, we can leverage the design of ShufTable to parallelize the SISO-PRF calls *without* increasing communication.

The crucial observation is that we have designed ShufTable to require only $\llbracket b_i \rrbracket$, regardless of found. That is, since previous OHTables stored and retrieved their dummy elements from some data structure, obtaining the index of each dummy from a_i was needed. We observe that it is not necessary: using the (3, 1) garbled circuits of [MR18] we “just-in-time” detect if q_i is stored in the table and output a dummy index if necessary ($\Pi_{\text{ShufTable}}.\text{Query}$, step 4). Thus, if found $_{i-1}$, for ShufTable it is sufficient to set q_i to be some uniformly random κ -bit value r ($\Pi_{\text{ShufTable}}.\text{Query}$, step 1), which, except with negligible probability, will not be stored in the table and will yield the desired dummy-output. This “just-in-time” trick is largely enabled by the Persistent shuffle trick we present in Section 5.1.

Thus to parallelize the SISO-PRF, we evaluate $\llbracket b_i \rrbracket = \mathcal{F}_{\text{ABB}}.\text{PRFEval}(\llbracket X_{\text{query}} \rrbracket, \llbracket k \rrbracket)$ for all $i \in \{1, \dots, \text{numLevels}\}$ in parallel before querying L_1 . Hence we can parallelize the SISO-PRF evaluations across the table without increasing the total number of SISO-PRF calls. This can be seen in Step 1.b of $\Pi_{\text{DORAM}}.\text{ReadAndWrite}$, Figure 4.

This optimization reduces numRoundsDORAM by $(\text{numLevels} - 1) \cdot \text{numRoundsPRFEval}$. In our implementation $\text{numLevels} \approx 5$ and $\text{numRoundsPRFEval} = 9$, so this saves us ≈ 45 rounds per DORAM query.

6 Reducing the depth of the hierarchy

In most hierarchical ORAM solutions, each level is twice as large as the level above it, i.e., $|L_{i+1}| = 2 \cdot |L_i|$. Since we must have $|L_{\text{numLevels}}| = O(N)$ and generally $|L_0|$ is a small constant, this means that $\text{numLevels} = O(\log N)$. In our protocol, we introduce a tuneable “base amplification factor” denoted baseAmpFactor , and set $|L_{i+1}| \approx \text{baseAmpFactor} \cdot |L_i|$. This change reduces numLevels by a factor of $\log_2(\text{baseAmpFactor})$ (but increases communication by a factor of $\text{baseAmpFactor} / \log(\text{baseAmpFactor})$). In our testing, we find that increasing baseAmpFactor dramatically improves practical performance because latency is significantly more time-expensive than bandwidth (Section 9). For instance, for GigaDORAM we found that increasing baseAmpFactor with N to maintain that $|L_0| \approx 2^8$ and $\text{numLevels} \approx 4$ yielded the best performance. For $N = 2^{30}$, this meant setting $\text{baseAmpFactor} = 2^7$, which is much larger than all previous protocols, which implicitly set $\text{baseAmpFactor} = 2$.

In most hierarchical ORAM schemes (where $\text{baseAmpFactor} = 2$), when levels L_0, \dots, L_i are full, they are reshuffled and rebuilt into L_{i+1} . Since

$$\sum_{j=0}^i |L_j| = \sum_{j=0}^i 2^j |L_0| = (2^{i+1} - 1) |L_0|$$

this rebuild schedule works nicely when $\text{baseAmpFactor} = 2$.

In our generalization, we also rebuild levels when they are full, but for $\text{baseAmpFactor} > 2$ we have that

$$\text{baseAmpFactor}^{i+1} \cdot |L_0| \gg \sum_{j=0}^i \text{baseAmpFactor}^j |L_0|$$

thus, we must slightly adjust our rebuild schedule. In particular, we must accommodate for “partial rebuilds.” We formalize this procedure in protocol $\Pi_{\text{DORAM}}.\text{Rebuild}$ presented in Section 8, Figure 4.

Since rebuilding level $i + 1$ costs $O(\kappa \cdot |L_0| \cdot \text{baseAmpFactor}^{i+1})$ communication and computation, the total amortized (re)build cost of the DORAM is at most

$$\begin{aligned} & \sum_{i=1}^{\frac{\log(N)}{\log(\text{baseAmpFactor})}} \frac{O((\kappa + D) \cdot |L_0| \cdot \text{baseAmpFactor}^{i+1})}{|L_0| \cdot \text{baseAmpFactor}^i} \\ &= O\left(\frac{(\kappa + D) \cdot \text{baseAmpFactor} \cdot \log(N)}{\log(\text{baseAmpFactor})}\right) \end{aligned}$$

Since the original amortized (re)build cost of all $O(\log N)$ levels of the DORAM is $O((\kappa + D) \log N)$, the cost of decreasing numLevels by a factor of $\log_2(\text{baseAmpFactor})$ is a factor of $O(\text{baseAmpFactor} / \log(\text{baseAmpFactor}))$ increase in communication and computation.

7 SpeedCache: Larger, optimized cache

Most hierarchical (D)ORAMs use a constant-sized “cache,” i.e., $|L_0| = O(1)$. Since each ORAM query performs a linear scan over the cache (implying a linear communication, computation, and number of rounds), a large cache can significantly hurt both asymptotic and practical performance. In our protocol, we use the Alibi reinsertion technique to “cache-the-stash” from Cuckoo hash tables at each level of the ORAM hierarchy. This means that we must have a moderately large cache size, otherwise reinserting the stashes would fill the stash and trigger an infinite chain of table rebuilds.

Thus we increase the cache size to $|L_0| = \Omega(\text{stashSize})$. This has the additional advantage of eliminating “small” levels from the ORAM hierarchy since $|L_i| > |L_0|$ for every level in the hierarchy.⁹

In order to facilitate our new, larger cache without hurting round complexity, we create, SpeedCache, that can be queried using $O((D + \log N) \cdot |L_0|)$ communication and only requires $\lceil \log \log N \rceil + 1$ rounds of communication (that is, *independent of the number of elements in the cache* – only dependent on the size of the address space). We implement our SpeedCache protocol in the 3-party MPC framework of [AFL⁺16]. We present the protocol, $\Pi_{\text{SpeedCache}}$, in Figure 3.

Since the SpeedCache protocol just uses MPC to compute secret-shared values, the security of the SpeedCache protocol follows immediately from the security of the underlying MPC protocol.

Experimentally we find that it is optimal to maintain $|L_0| \approx 2^{10}$ a constant, rather than scaling $|L_0|$ with N . Note that $2^{10} > \log(N)$ for any conceivable value of N . Concretely, setting $|L_0| \approx 2^9$ as we do in practice allows us to reduce numLevels by $\approx 8/\log_2(\text{baseAmpFactor})$. We found that changing $|L_0|$ useful for adapting the performance of GigaDORAM to different network settings. For instance when latency is high, a larger $|L_0|$ yields better performance.

An alternative approach would be to implement the trivial linear-scan cache (see Appendix B) via garbled circuits in a constant number of rounds. When $|L_0| \approx 2^{10}$ we estimate this would make $L_0.\text{Query}$ up to $100\times$ slower than our implementation of SpeedCache.

8 Full DORAM protocol Π_{DORAM}

We give the GigaDORAM protocol, Π_{DORAM} , in Figure 4.

Round-complexity: With our optimizations the final round

⁹Because cuckoo hash tables have higher build-failure probability when the table size is small, many hierarchical (D)ORAMs had two types of tables, one for the “small” levels and one for the “large” levels (e.g. [LO13, FNO22]). Because our L_0 is sufficiently large, we do not incur this complexity.

complexity of $\Pi_{\text{DORAM}}.\text{Query}$ is

$$\begin{aligned} \text{numRoundsDORAM} &= \\ &\text{numLevels} \cdot \text{numRoundsOHTable} \\ &+ \max(\text{numRoundsCache}, \text{numRoundsPRFEval}) + 1 \\ &= \text{numLevels} \cdot 5 + \max(\log \log(N) + 2, 9) + 1 \\ &\approx 5 \cdot 5 + 9 + 1 = 35 \quad (\text{when } N = 2^{31}) \end{aligned}$$

The term $\max(\log \log(N) + 2, 9)$ occurs because $L_0.\text{Query}$ ($\log \log(N) + 2$ rounds) and the SISO-PRF pre-evaluations (9 rounds – see Section 5.3, Appendix C) can be evaluated in parallel. Hence, compared to the over 2500-round-per-query [FNO22] construction we draw inspiration from, GigaDORAM has a $\approx 70\times$ reduction in round complexity. Our overall *communication* and *computation* complexity is comparable to that of [LO13, FNO22] and previous hierarchical DORAMs.

Lemma 8.1. *Protocol Π_{DORAM} securely realizes the DORAM functionality in the semi-honest (3,1)-setting.*

Proof. Protocol Π_{DORAM} is a straightforward application of the hierarchical solution to our building blocks, $\Pi_{\text{ShufTable}}$ and $\Pi_{\text{SpeedCache}}$ in the distributed setting. The security of the hierarchical solution is standard (see e.g. [Ost92]).

The proof proceeds as follows. Since ShufTable will be oblivious as long as no queries are repeated, we must show that even if the client makes repeated queries into the DORAM, no OHTable is ever queried twice between rebuilds.

To see this, suppose the user makes a series of queries, x_1, \dots, x_t and $x_{i_1} = x_{i_2}$ with $i_1 < i_2$. Suppose x_{i_1} was initially at level ℓ . In this case, after query x_{i_1} , the key-value pair (x_{i_1}, y_{i_1}) will be inserted into the cache. As later queries come in and the cache becomes full (x_{i_1}, y_{i_1}) may get pushed to lower levels of the hierarchy. After sufficiently many queries, levels $i = 1, \dots, \ell - 1$ will be rebuilt into level ℓ . If query x_{i_2} happens *before* level ℓ is rebuilt, x_{i_2} will be found at a level above level ℓ , and at level ℓ a (new) dummy element will be queried. If query x_{i_2} happens *after* level ℓ has been rebuilt, then at this point x_{i_1} has never been queried against this new OHTable. \square

9 Evaluation

In this section, we evaluate the practical performance of GigaDORAM and compare its performance to previous DORAM implementations. To make the comparisons fair, whenever possible we evaluated *on the same hardware setup*. To make our results easily reproducible, we test on AWS, using c5n.metal instances with 72vCPUs and 192GiB memory.

We measure only writes in our experiments because “reads” vary across works (e.g. DuORAM separates “dependent” and “independent” reads) and for all constructions, writes are at least as expensive as reads. For each construction at each N

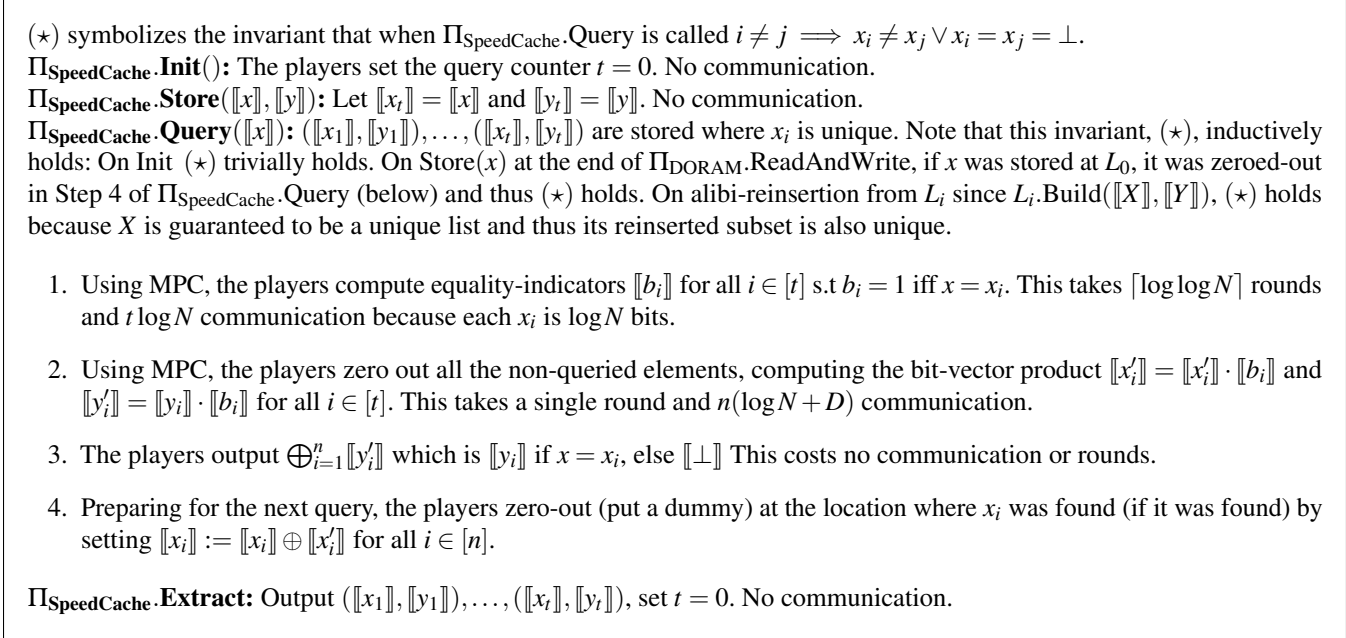


Figure 3: The SpeedCache protocol, $\Pi_{\text{SpeedCache}}$

we averaged performance across multiple writes (up to 100K writes in faster networks). All experiments were on $D = 64$ bit payloads as in DuORAM’s testing.

To compare fairly with previous works, we *re-benchmark all previous DORAMs with existing implementations* (with the exception of proprietary [JZLR22]) *on the same hardware setup we used to test GigaDORAM*. To the best of our knowledge, this is the most comprehensive Distributed ORAM benchmark to date.¹⁰ See Appendix E for details.

9.1 Low-latency tests

One of the primary applications of DORAM is in heavy-compute large-scale MPC computations (where random access could outperform circuit-based computations). Thus we focus on low-latency network environments where large-scale MPC becomes practical.¹¹

With this motivation, we begin by benchmarking GigaDORAM in a realistic low-latency regime. Specifically, we deploy to three cluster-placed AWS c5n.metal instances.¹² In

¹⁰We do not benchmark [Lau14], because it is based on the proprietary Sharemind software, and is focused on batch queries. As noted in that work, in the single-query setting those techniques are not expected to give significant improvements over straightforward “ORAM in MPC” schemes.

¹¹Although MPC is possible under poor network conditions [WRK17b], it is only possible when the circuits are fairly simple and thus would not benefit significantly from random access. To put this in perspective, [WRK17b] benchmark computing the AES circuit under MPC. In GigaDORAM, we compute AES (or LowMC) under a circuit-based MPC *multiple times for each query*. Thus no DORAM that makes use of a SISO-PRF could not possibly improve the performance of simple computations like AES.

¹²For some comparisons (e.g. [VHG22, JW18]) we used existing Docker

this setup, In cluster placement at region us-west-2. Via `ping -c 20` we measure latency of `min/avg/max/mdev = 0.229/0.236/0.244/0.003 ms`.¹³

Figure 5 shows that in this low-latency environment GigaDORAM outperforms all other DORAMs for $2^{12} \leq N \leq 2^{31}$. GigaDORAM perform 700 – 900 queries/sec for all measured N , with oscillations that depend on the size of the hierarchical cache L_0 .¹⁴

9.2 Testing in restricted networks

Previous DORAM papers benchmarked [Ds17, VHG22, JZLR22] under artificial restrictions to the network. For example, [VHG22] benchmarked restricted their bandwidth to 100Mb/s with a latency of 30ms. These network conditions are lower than you would expect from geographically separated machines. For example, we measured 20ms ping from us-east-1 to us-west-2, servers in different AWS regions should enjoy several Gbit connections [Bar18], and it is possible to achieve 1ms ping times *across different cloud providers*¹⁵. Still, we compare the performance of GigaDO-

setups to simulate three parties on a single c5.metal machine in which case we adjust the network via the `tc` command (see Section E).

¹³We observed latency fluctuating according to usage, although AWS does not make such policies public.

¹⁴The “oscillations” seen in our performance are due to the fact $\log_2 |L_0| = N - \text{baseAmpFactor} \cdot (\text{numLevels} - 1)$ and the fact that `baseAmpFactor` is a power of 2, and thus for certain values of N we cannot reasonably set $|L_0|$ as small as we’d like.

¹⁵<https://medium.com/@sachinkagarwal/public-cloud-inter-region-network-latency-as-heat-maps-134e22a5ff19>

$\Pi_{\text{DORAM.Init}}(\llbracket Y \rrbracket)$:

1. The players input the address of each payload $\llbracket X_i \rrbracket = \mathcal{F}_{\text{ReplicatedMPC}}.\text{InputConstant}(i)$ for $i = 1$ to N . The players build the top level $L_{\text{numLevels}} = \mathcal{F}_{\text{OHTable}}.\text{Build}(\llbracket X \rrbracket, \llbracket Y \rrbracket, N)$ and initialize the cache $L_0.\text{Init}()$. The players set the query counter $t = 0$.

$\Pi_{\text{DORAM.ReadAndWrite}}(\llbracket X_{\text{query}} \rrbracket, \llbracket Y_{\text{new}} \rrbracket, \llbracket \text{isWrite} \rrbracket)$: $|L_0|$ is the size of the cache and is tracked externally.

1. The players increment the query counter, t , initialize numLevels-bit Alibi data accumulator $\llbracket \epsilon_{\text{accum}} \rrbracket = \llbracket 0^{\text{numLevels}} \rrbracket$, D-bit payload accumulator $\llbracket Y_{\text{accum}} \rrbracket = \llbracket \perp \rrbracket$, and one-bit flag $\llbracket \text{found} \rrbracket = 0$. Then, the players do the following in parallel (This step takes $\max\{\text{numRoundsCache}, \text{numRoundsPRFEval}\}$ rounds):
 - (a) query the Cache, $\llbracket Y_{\text{accum}} \rrbracket, \llbracket \text{found} \rrbracket = L_0.\text{Query}(\llbracket X_{\text{query}} \rrbracket)$. The players extract $\llbracket \epsilon_{\text{accum}} \rrbracket$ from $\llbracket Y_{\text{accum}} \rrbracket$ (by silently copying the last numLevels-bits).
 - (b) For each level $\ell \in \text{numLevels}$ evaluate $\llbracket a_\ell \rrbracket = \mathcal{F}_{\text{ABB}}.\text{PRFEval}(\llbracket k_\ell \rrbracket, \llbracket X_{\text{query}} \rrbracket)$
2. For each ℓ from 1 to numLevels, if there is an OHTable at level L_ℓ :
 - (a) Set the bit $\llbracket \text{useDummy} \rrbracket = \llbracket \epsilon \rrbracket[\ell] \oplus \llbracket \text{found} \rrbracket$, i.e., locally XOR the ℓ th bit of the share of ϵ with the share of found. Query L_i , calling $\llbracket Y_\ell \rrbracket, \llbracket \text{found}_\ell \rrbracket = L_\ell.\text{Query}(\llbracket a_\ell \rrbracket, \llbracket \text{useDummy} \rrbracket)$. This step takes numRoundsOHTable rounds.
3. Set $\llbracket Y_{\text{new}} \rrbracket \leftarrow \mathcal{F}_{\text{ReplicatedMPC}}.\text{IfThenElse}(\llbracket \text{isWrite} \rrbracket, \llbracket Y_{\text{new}} \rrbracket, \llbracket Y_{\text{accum}} \rrbracket)$. Set the Alibi mask of $\llbracket Y_{\text{new}} \rrbracket, \llbracket Y_{\text{new}}.\epsilon \rrbracket$ to 0. Call $\mathcal{F}_{\text{Linear}}.\text{Store}(\llbracket X \rrbracket, \llbracket Y_{\text{new}} \rrbracket)$. This takes 1 round.
 - (a) Silently extract the numLevels-bit value $\llbracket \epsilon_\ell \rrbracket$ from $\llbracket Y_\ell \rrbracket$ and update $\llbracket \epsilon_{\text{accum}} \rrbracket \leftarrow \llbracket \epsilon_{\text{accum}} \rrbracket \oplus \llbracket \epsilon_\ell \rrbracket$. Update the value of $\llbracket \text{found} \rrbracket$ to $\text{found} \vee \text{found}_\ell$. Additionally, In parallel with the first round of the next iteration of the loop, update $\llbracket Y_{\text{accum}} \rrbracket \leftarrow \mathcal{F}_{\text{ReplicatedMPC}}.\text{IfThenElse}(\llbracket \text{found}_\ell \rrbracket, \llbracket Y_\ell \rrbracket, \llbracket Y_{\text{accum}} \rrbracket)$. This step takes one round.
4. If $t = |L_0|$, run the subroutine $\Pi_{\text{DORAM.Rebuild}}()$.

$\Pi_{\text{DORAM.Rebuild}}()$: We define a level L_i for $1 \leq i \leq \text{numLevels}$ in the DORAM to be *full* if it contains a ShufTable with $(\text{baseAmpFactor} - 1) \cdot \text{baseAmpFactor}^i \cdot |L_0|$ elements. L_0 is full if it has been written to $|L_0|$ times since it was last initialized. Suppose that ℓ is the largest number such that levels L_0, \dots, L_ℓ are full while $L_{\ell+1}$ is not, instead containing $A \cdot \text{baseAmpFactor}^\ell \cdot |L_0|$ elements for some $A \in \{0, \dots, \text{baseAmpFactor} - 2\}$ If $\ell = \text{numLevels}$, then necessarily $A = 1$.

1. By concatenating the output of $L_0.\text{Extract}()$ and $L_i.\text{Extract}()$ for $i \in [\ell + 1]$, the players prepare lists $\llbracket X \rrbracket, \llbracket Y \rrbracket$ of length $(A + 1) \cdot \text{baseAmpFactor}^\ell \cdot |L_0|$ containing all the elements to be (potentially) placed in $L_{\ell+1}$.
2. If $\ell < \text{numLevels}$:
In parallel, relabel the dummies, calling $\llbracket X_j^* \rrbracket \leftarrow \mathcal{F}_{\text{ReplicatedMPC}}.\text{ReplaceIfNull}(\llbracket X_j \rrbracket, \llbracket N + j \rrbracket)$ and $\llbracket Y_j^* \rrbracket = \llbracket Y_j \rrbracket$ (syntactically convenient) for $j \in [(A + 1) \cdot \text{baseAmpFactor}^\ell \cdot |L_0|]$ (this is so P_1 does not learn how many dummies were queried at the previous level).
3. If $\ell = \text{numLevels}$: The players “cleanse out the dummies” by shuffling, $\llbracket \hat{X} \rrbracket, \llbracket \hat{Y} \rrbracket = \mathcal{F}_{\text{ABB}}.\text{ObliviousShuffle}(\llbracket X \rrbracket, \llbracket Y \rrbracket)$, and then revealing which element is dummy by calling $\mathcal{F}_{\text{ABB}}.\text{Reveal}(\mathcal{F}_{\text{ABB}}.\text{Equals}(\llbracket \hat{X}_j \rrbracket, \perp))$ for all j .
The above will reveal exactly N 0’s and N 1’s (because we started with N “real” elements which we have maintained and we made N queries since the last build $L_{\text{numLevels}}$. Since stale elements are relabeled (step 2.A) the invariant holds). Compact the $\llbracket \hat{X} \rrbracket$ and $\llbracket \hat{Y} \rrbracket$ shares corresponding to 0’s into arrays $\llbracket X^* \rrbracket$ and $\llbracket Y^* \rrbracket$.
4. The players reinitialize the cache by calling $L_0 = \mathcal{F}_{\text{Linear}}.\text{Init}(|L_0|)$ and build $L_{\ell+1}$ by calling $\llbracket X^{\text{stash}} \rrbracket, \llbracket Y^{\text{stash}} \rrbracket = L_{\ell+1}.\text{Build}(\llbracket X^* \rrbracket, \llbracket Y^* \rrbracket)$
5. For i from 1 to stashSize, update the Alibi bit $\llbracket \epsilon \rrbracket[\ell + 1] =$ of $\llbracket Y_i^{\text{stash}} \rrbracket$ to 1, then call $\mathcal{F}_{\text{Linear}}.\text{Store}(\llbracket X_i^{\text{stash}} \rrbracket, \llbracket Y_i^{\text{stash}} \rrbracket)$. The players reset $t = 0$.

Figure 4: The complete DORAM protocol, Π_{DORAM}

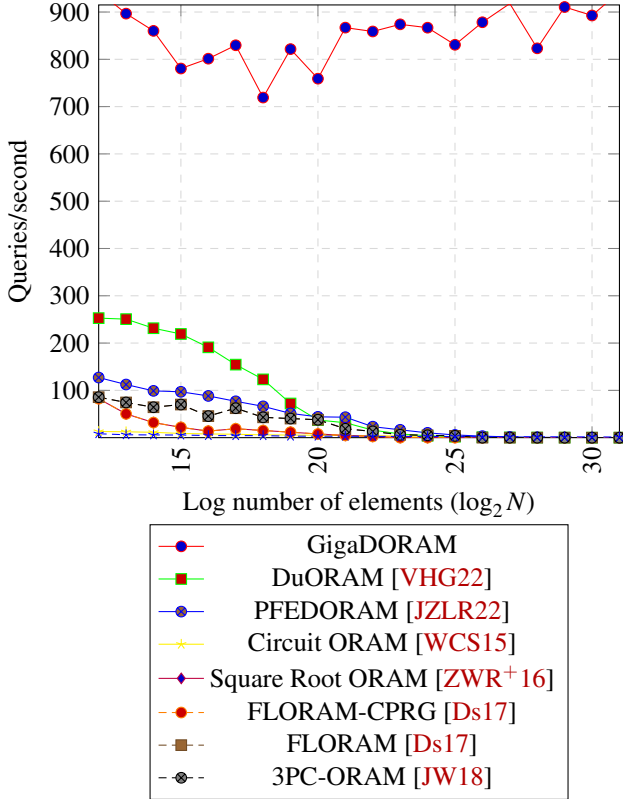


Figure 5: Number of queries/second vs. log number of elements (N). Each query is a write of random data to a random index. Preprocessing costs are accounted for. The number of queries per second is an estimate of (at least) hundreds of queries. All constructions were run in AWS cluster placement network environment.

RAM to other DORAMs under worse network conditions (which may be unavoidable in some applications).

Performance without cluster-placement. Figure 5 shows the performance of GigaDORAM when the machines are cluster-placed. When the machines were not cluster-placed (but still in the same AWS region) the number of queries / sec dropped by about 33% (Table 1).

Performance and breakeven points in varying latency/bandwidth. Using the t_c (as in [VHG22, Ds17]) we vary the latency and bandwidth to search for the breakeven point between GigaDORAM and DuORAM (which is the best previous DORAM in high-latency low-bandwidth environments).

Figure 6.a shows that the performance of GigaDORAM degrades dramatically with latency while the performance of DuORAM degrades dramatically with $\log(N)$. In particular, we see that at $\log(N) = 20$, for latency $\lesssim 8$ ms GigaDORAM outperforms DuORAM while for $\gtrsim 8$ ms latency DuORAM edges out GigaDORAM. At $\log(N) = 25$ GigaDORAM outperforms DuORAM for all sub-40ms latency.

$\log(N)$	Standard Q/s	Cluster Q/s	Slowdown
20	543	812	33%
23	573	910	37%
26	597	877	32%
29	589	891	34%

Table 1: Benchmarking GigaDORAM using 3 different c5n.metal machines standard placement vs cluster placement in us-west-2. We benchmarked over 100,000 queries, using lowMC as our SISO-PRF.

$\log(N)$	AES Q/s	LowMC Q/s	Slowdown
20	393	812	51%
23	405	910	55%
26	401	877	54%
29	417	891	53%

Table 2: Benchmarking GigaDORAM using AES to implement a SISO-PRF instead of LowMC. We run the benchmark in the same setting as Figure 5.

Figure 6.b shows that the performance of neither construction degrades substantially with bandwidth. Our experiments show DuORAM outperforming GigaDORAM for $\log(N) = 20$ and GigaDORAM outperforming DuORAM for $\log(N) = 25$. Again, we see the performance of DuORAM degrading significantly with $\log(N)$ while GigaDORAM’s performance hardly changes.

9.3 The cost of GigaDORAM

Running large computations in AWS can be quite expensive, but due to its high query/sec, GigaDORAM is cost-efficient. Including both network and compute costs, GigaDORAM can handle about over 120,000 queries per dollar with current AWS pricing, which is $10\times$ DuORAM [VHG22], which is to the best of our knowledge, the previously most cost-efficient DORAM. See Appendix F for the calculations.

9.4 Replacing LowMC with AES

For reference, we also evaluated GigaDORAM using AES instead of LowMC as our SISO-PRF. The AES circuit we use has about $10\times$ the ANDs and $10\times$ the AND-depth as the LowMC circuit we use, and is thus less efficient to evaluate under MPC. Table 2 shows approximately a $2\times$ slowdown to GigaDORAM when using AES instead of LowMC.

10 Conclusion

In this work we introduce GigaDORAM, the most efficient and scalable DORAM construction to date. At $N = 2^{31}$, our DORAM can perform over 700 queries per second, making




Figure 6: (a) Number of queries/second vs. latency in milliseconds for fixed 1Gbit network (b) Number of queries/second vs. Bandwidth in Gbit for fixed 8ms latency. Executed on a single c5.metal via multiple processes restricted in communication via the `tc`. Preprocessing costs are accounted for. The number of queries per second is an estimate of 5000 queries for GigaDORAM and 128 for DuORAM.

GigaDORAM orders of magnitude faster than prior DORAM constructions in low-latency environments. We give a custom C++, open-source implementation of GigaDORAM. We hope GigaDORAM will enable the first somewhat practical RAM-MPC applications and open a new realm of possibilities for privacy-preserving cloud data stores.

11 Acknowledgments

We thank AWS solutions engineers, the LowMC [ARS⁺15], the PFEDORAM [JZLR22], the DuORAM [VHG22], and the replicated-MPC [AFL⁺16] authors for their help with benchmarking. We thank Michael Brown, Eric Chen, Matthew Day, Daniel Noble, and Sam Kumar for helpful discussions. We thank the USENIX reviewers for extensive feedback which

improved the quality of the exposition.

Supported in part by DARPA under Cooperative Agreement HR0011-20-2-0025, the Algorand Centers of Excellence programme managed by Algorand Foundation, NSF grants CNS-2246355, CCF-2220450 and CNS-2001096, US-Israel BSF grant 2015782, Amazon Faculty Award, Cisco Research Award and Sunday Group. Any views, opinions, findings, conclusions, or recommendations contained herein are those of the author(s) and should not be interpreted as necessarily representing the official policies, either expressed or implied, of DARPA, the Department of Defense, the Algorand Foundation, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

References

- [ACF⁺12] Bernhard Ager, Nikolaos Chatzis, Anja Feldmann, Nadi Sarar, Steve Uhlig, and Walter Willinger. Anatomy of a large European IXP. In *SIGCOMM*, 2012.
- [AFL⁺16] Toshinori Araki, Jun Furakawa, Yehuda Lindell, Ariel Nof, and Kazuma Ohara. High-throughput semi-honest secure three-party computation with an honest majority. In *CCS*, 2016.
- [AKL⁺20] Gilad Asharov, Ilan Komargodski, Wei-Kai Lin, Kartik Nayak, Enoch Peserico, and Elaine Shi. OptORAMa: Optimal oblivious RAM. In *EUROCRYPT*, 2020.
- [AKLS21] Gilad Asharov, Ilan Komargodski, Wei-Kai Lin, and Elaine Shi. Oblivious RAM with worst-case logarithmic overhead. In *CRYPTO*, pages 610–640. Springer, 2021.
- [ARS⁺15] Martin R Albrecht, Christian Rechberger, Thomas Schneider, Tyge Tiessen, and Michael Zohner. Ciphers for MPC and FHE. In *EUROCRYPT*, pages 430–454. Springer, 2015.
- [Bar18] Jeff Bar. The floodgates are open – increased network bandwidth for EC2 instances. <https://aws.amazon.com/blogs/aws/the-floodgates-are-open-increased-network-bandwidth-for-ec2-instances/>, 2018.
- [BBVY21] Subhadeep Banik, Khashayar Barooti, Serge Vaudey, and Hailun Yan. New attacks on LowMC instances with a single plaintext/ciphertext pair. 2021.
- [BGI15] Elette Boyle, Niv Gilboa, and Yuval Ishai. Function secret sharing. In *EUROCRYPT*, 2015.
- [BKK⁺16] Dan Bogdanov, Liina Kamm, Baldur Kubo, Reimo Rebane, Ville Sokk, and Riivo Talviste. Students and taxes: a privacy-preserving study using secure computation. *Proc. Priv. Enhancing Technol.*, 2016(3):117–135, 2016.
- [BKKO20] Paul Bunn, Jonathan Katz, Eyal Kushilevitz, and Rafail Ostrovsky. Efficient 3-party distributed ORAM. In *SCN*, 2020.
- [BMR90] Donald Beaver, Silvio Micali, and Phillip Rogaway. The round complexity of secure protocols. In *STOC*, pages 503–513, 1990.
- [CCD88] David Chaum, Claude Crépeau, and Ivan Damgård. Multiparty Unconditionally Secure Protocols. In *STOC*, 1988.
- [CDG⁺17] Melissa Chase, David Derler, Steven Goldfeder, Claudio Orlandi, Sebastian Ramacher, Christian Rechberger, Daniel Slamanig, and Greg Zaverucha. Post-quantum zero-knowledge and signatures from symmetric-key primitives. In *CCS*, pages 1825–1842, 2017.

- [CDI05] Ronald Cramer, Ivan Damgård, and Yuval Ishai. Share conversion, pseudorandom secret-sharing and applications to secure computation. In *TCC*, 2005.
- [CGKO06] Reza Curtmola, Juan Garay, Seny Kamara, and Rafail Ostrovsky. Searchable symmetric encryption: Improved definitions and efficient constructions. In *Proceedings of the 13th ACM Conference on Computer and Communications Security*, CCS '06, page 79–88, New York, NY, USA, 2006. Association for Computing Machinery.
- [DLMW15] Itai Dinur, Yunwen Liu, Willi Meier, and Qingju Wang. Optimized interpolation attacks on lowmc. In *ASIACRYPT*, pages 535–560. Springer, 2015.
- [Ds17] Jack Doerner and abhi shelat. Scaling ORAM for secure computation. In *CCS*, 2017.
- [FNO21] Brett Hemenway Falk, Daniel Noble, and Rafail Ostrovsky. Alibi: A flaw in cuckoo-hashing based hierarchical ORAM schemes and a solution. In *EUROCRYPT*, pages 338–369. Springer, 2021.
- [FNO22] Brett Hemenway Falk, Daniel Noble, and Rafail Ostrovsky. 3-party distributed ORAM from oblivious set membership. In *SCN*, pages 437–461. Springer, 2022.
- [GI14] Niv Gilboa and Yuval Ishai. Distributed point functions and their applications. In *EUROCRYPT*, 2014.
- [GKK⁺12] S Dov Gordon, Jonathan Katz, Vladimir Kolesnikov, Fernando Krell, Tal Malkin, Mariana Raykova, and Yevgeniy Vahlis. Secure two-party computation in sublinear (amortized) time. In *CCS*, 2012.
- [GMOT12] Michael T Goodrich, Michael Mitzenmacher, Olga Ohrimenko, and Roberto Tamassia. Privacy-preserving group data access via stateless oblivious RAM simulation. In *SODA*, 2012.
- [GMW87] Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game. In *STOC*, 1987.
- [GO96] Oded Goldreich and Rafail Ostrovsky. Software protection and simulation on oblivious RAMs. *JACM*, 43(3), 1996.
- [GRR⁺16] Lorenzo Grassi, Christian Rechberger, Dragos Rotaru, Peter Scholl, and Nigel P Smart. MPC-friendly symmetric key primitives. In *CCS*, pages 430–443, 2016.
- [JW18] Stanislaw Jarecki and Boyang Wei. 3PC ORAM with low latency, low bandwidth, and fast batch retrieval. In *ACNS*, 2018.
- [JZLR22] Keyu Ji, Bingsheng Zhang, Tianpei Lu, and Kui Ren. Multi-party private function evaluation for RAM. *Cryptology ePrint Archive*, Paper 2022/939, 2022. <https://eprint.iacr.org/2022/939>.
- [KLO12] Eyal Kushilevitz, Steve Lu, and Rafail Ostrovsky. On the (in) security of hash-based oblivious RAM and a new balancing scheme. In *SODA*, 2012.
- [KMW09] Adam Kirsch, Michael Mitzenmacher, and Udi Wieder. More robust hashing: Cuckoo hashing with a stash. *SIAM Journal on Computing*, 2009.
- [Lau14] Peeter Laud. Privacy-preserving minimum spanning trees through oblivious parallel RAM for secure multiparty computation. *IACR ePrint Archive* 2014/630, 2014.
- [LIM21] Fukang Liu, Takanori Isobe, and Willi Meier. Cryptanalysis of full LowMC and LowMC-M with algebraic techniques. In *CRYPTO*, pages 368–401, 2021.
- [LO13] Steve Lu and Rafail Ostrovsky. Distributed oblivious RAM for secure two-party computation. In *TCC*, 2013.
- [LWZ11] Sven Laur, Jan Willemson, and Bingsheng Zhang. Round-efficient oblivious database manipulation. In *ISC*, 2011.
- [MR18] Payman Mohassel and Peter Rindal. ABY3: A mixed protocol framework for machine learning. In *CCS*, pages 35–52, 2018.
- [Nav15] Muhammad Naveed. The fallacy of composition of oblivious RAM and searchable encryption. *IACR ePrint* 2015/688, 2015.
- [NIS21] NIST. Post-quantum cryptography PQC: Round 3 submissions. <https://csrc.nist.gov/Projects/post-quantum-cryptography/post-quantum-cryptography-standardization/round-3-submissions>, 2021.
- [Nob21] Daniel Noble. An intimate analysis of cuckoo hashing with a stash. *IACR ePrint* 2021/447, 2021.
- [OS97] Rafail Ostrovsky and Victor Shoup. Private information storage. In *STOC*, volume 97, 1997.
- [Ost90] Rafail Ostrovsky. Efficient computation on oblivious RAMs. In *STOC*, 1990.
- [Ost92] Rafail Ostrovsky. *Software Protection and Simulation On Oblivious RAMs*. PhD thesis, Massachusetts Institute of Technology, 1992.
- [PPRY18] Sarvar Patel, Giuseppe Persiano, Mariana Raykova, and Kevin Yeo. PanORAMA: Oblivious RAM with logarithmic overhead. In *FOCS*, 2018.
- [PR01] Rasmus Pagh and Flemming Friche Rodler. Cuckoo hashing. In *ESA*, 2001.
- [Rin] Peter Rindal. The ABY3 Framework for Machine Learning and Database Operations. <https://github.com/ladnir/aby3>.
- [SCSL11] Elaine Shi, T-H Hubert Chan, Emil Stefanov, and Mingfei Li. Oblivious RAM with $O((\log N)^3)$ worst-case cost. In *ASIACRYPT*, 2011.
- [Smi21] Philip Smith. Internet exchange point design. <https://nsrc.org/workshops/2021/marwan-cnrt-ixp/networking/peering-ixp/en/presentations/IXP-Design.pdf>, 2021.
- [VHG22] Adithya Vadapalli, Ryan Henry, and Ian Goldberg. DuORAM: A bandwidth-efficient distributed ORAM for 2- and 3-party computation. *IACR ePrint* 2022/1747, 2022.
- [WCS15] Xiao Wang, Hubert Chan, and Elaine Shi. Circuit ORAM: On tightness of the Goldreich-Ostrovsky lower bound. In *CCS*, 2015.
- [WHC⁺14] Xiao Shaun Wang, Yan Huang, T-H Hubert Chan, abhi shelat, and Elaine Shi. SCORAM: oblivious RAM for secure computation. In *CCS*, 2014.
- [WRK17a] Xiao Wang, Samuel Ranellucci, and Jonathan Katz. Global-scale secure multiparty computation. In *CCS*, pages 39–56, 2017.
- [WRK17b] Xiao Wang, Samuel Ranellucci, and Jonathan Katz. Global-scale secure multiparty computation. In *CCS*, pages 39–56, 2017.
- [Yao82] Andrew Yao. Protocols for secure computations (extended abstract). In *FOCS*, 1982.
- [Yao86] Andrew Yao. How to generate and exchange secrets. In *FOCS*, 1986.
- [ZE15] Samee Zahur and David Evans. Obliv-C: A language for extensible data-oblivious computation. *IACR ePrint* 2015/1153, 2015.
- [ZWR⁺16] Samee Zahur, Xiao Wang, Mariana Raykova, Adrià Gascón, Jack Doerner, David Evans, and Jonathan Katz. Revisiting square-root ORAM: efficient random access in multi-party computation. In *S & P*, 2016.

A ABB Functionalities

The basic ABB operations we use are described in Figure 7.

- $[x]^{(i,j)} = \mathcal{F}_{\text{ABB}}.\text{InputTo2Sharing}(x, P_i, P_j)$. A single player shares creates an additive sharing of a secret, x , among participants i and j .
- $[x]^{(i,j)} = \mathcal{F}_{\text{ABB}}.\text{ReshareReplicatedTo2Sharing}(\llbracket x \rrbracket, \{P_i, P_j\})$. Convert a 3-party CNF sharing of a secret, x , to a two-party DNF sharing of the same secret, x , held by participants i and j .
- $[x]^{(i',j')} = \mathcal{F}_{\text{ABB}}.\text{Reshare2To2Sharing}([x]^{(i,j)}, \{P_{i'}, P_{j'}\})$. Convert a two-party DNF sharing of a secret, x , held by participants i and j , to a two-party DNF sharing of the same secret, x , held by participants i' and j' .
- $\llbracket x \rrbracket = \mathcal{F}_{\text{ABB}}.\text{Reshare2SharingToReplicated}([x]^{(i,j)})$. Convert a two-party DNF sharing of a secret, x , held by participants i and j to a three-party CNF sharing of the same secret, x .
- $\llbracket k \rrbracket = \mathcal{F}_{\text{ABB}}.\text{RandomElement}(\kappa)$. Generate a three-party CNF sharing of a uniformly random field element (whose value is unknown to the participants).
- $x = \mathcal{F}_{\text{ABB}}.\text{RevealTo}(P_i, \llbracket x \rrbracket)$. Reveal a secret-shared value, $\llbracket x \rrbracket$, to participant P_i .

We also abstract away a few more sophisticated operations:

- $\llbracket y \rrbracket = \mathcal{F}_{\text{ABB}}.\text{PRFEval}(\llbracket x \rrbracket, \llbracket k \rrbracket)$. Evaluate a SISO-PRF with secret-shared key, $\llbracket k \rrbracket$, on secret-shared input, $\llbracket x \rrbracket$, to obtain a secret-shared output, $\llbracket y \rrbracket$. In our instantiation, we instantiate the $\mathcal{F}_{\text{ABB}}.\text{PRFEval}(\cdot, \cdot)$ functionality by evaluating the LowMC block cipher [ARS⁺15] under MPC. Specifically, in our custom implementation of a (3,1)-MPC protocol (based on [AFL⁺16]).
- $\llbracket \text{QueryCircuit}(x_1, \dots, x_l) \rrbracket = \mathcal{F}_{\text{ABB}}.\text{EvalCircuit}(3,1)\text{GC}(\text{QueryCircuit}, \text{Inputs} = (\llbracket x_1 \rrbracket, \dots, \llbracket x_l \rrbracket))$. Evaluate a 3-party garbled circuit on secret-shared inputs $\llbracket x_1 \rrbracket, \dots, \llbracket x_l \rrbracket$, and returns a *sharing* of the output of the circuit computation. We use a custom implementation of the 3-party Garbled Circuit protocol outlined in ABY3 [MR18].
- $\llbracket Y \rrbracket = \mathcal{F}_{\text{ABB}}.\text{ObliviousShuffle}(\llbracket X \rrbracket)$. This functionality implements a linear-communication, three-party shuffle of secret shared values [LWZ11]. So $Y = \pi(X)$ for some random permutation, unknown to the participants. We also use a modified protocol to output a sharing of the permutation as well. $\llbracket Y \rrbracket, \llbracket L \rrbracket = \mathcal{F}_{\text{ABB}}.\text{ObliviousShuffle}(\llbracket X \rrbracket, \text{DistributeShuffle} = \text{True})$. In this setting, $Y = \pi(X)$ as before, and $L = \pi^{-1}(1, 2, \dots, n)$. We describe how to implement this novel “Persistent shuffle” in Section 5.1.

Figure 7: The ABB functionalities used in our DORAM protocol.

$\log(N)$	# of bytes sent
20	$5.86 \cdot 10^9$
25	$5.64 \cdot 10^9$
30	$6.73 \cdot 10^9$

Figure 8: The number of bytes sent by a single GigaDORAM machine for 100,000 queries, using LowMC as the PRF at varying values of $\log(N)$ at the same values of baseAmpFactor, numLevels, in which our main benchmark was conducted (see Figure 5).

B The naïve cache protocol

The naïve cache protocol works as follows

```
r ← ⊥
for i = 1, ..., t do
  if x = xi then
    r ← yi
  end if
end for
Return: r
```

This protocol has multiplicative depth, t , thus implementing this under the MPC of [AFL⁺16] leads to a protocol with $\text{numRoundsCache} = |L_0|$. It is possible to implement this with a garbled-circuit-based approach, (e.g. 3-party garbled circuits [MR18] or BMR [BMR90]). This results in a constant-round MPC protocol, but the communication complexity is large – too inefficient for our application.

$\Pi_{\text{SpeedCache}}$, (described in Figure 3), parallelizes the equality tests of the naïve protocol leading to a low-depth circuit, that we implement using the 3-party MPC of [AFL⁺16].

C LowMC

In this section, we discuss LowMC, an MPC-friendly block-cipher we use to instantiate $\mathcal{F}_{\text{ABB.PRFEval}}$.

LowMC. LowMC (Low Multiplicative Complexity) [ARS⁺15] is a family of block cipher that is built with MPC, ZK, and FHE in mind. LowMC has a variety of instantiations which trades low number of AND gates and a low circuit AND-depth.¹⁶ The instantiation of LowMC we use has 46837 total gates, out of which 1134 are ANDs, stacked into 9-AND-depth circuit. By contrast, AES has a total of 36663 gates, out of which 6400 are ANDs, stacked into a 60-AND-depth-circuit.¹⁷ Since the AND-depth of the circuit is equal to the number of rounds to evaluate the circuit under MPC, using LowMC instead of AES saves 51 rounds of communication for each query.

We provide the first circuit files for LowMC and encode them in the popular Bristol fashion.¹⁸ The traditional Bristol fashion requires that each virtual “wire” to be only assigned once, requiring more memory for the computation than necessary. Threading wires through the circuit (i.e., reusing memory) we are able speed up the computation of LowMC under the [AFL⁺16] MPC by a factor of 2.

Although LowMC has less analysis than AES, its security has been adapted into the Picnic signature scheme [CDG⁺17], a 3rd round candidate in the NIST post-quantum digital signature contest [NIS21]. Additionally, there have been several thorough cryptanalysis attempts [LIM21,

¹⁶LowMC provides a script to calculate this tradeoff https://github.com/LowMC/lowmc/blob/master/determine_rounds.py.

¹⁷We refer to Bristol Fashion AES circuit file from https://homes.esat.kuleuven.be/~nsmart/MPC/MAND/aes_128.txt

¹⁸<https://homes.esat.kuleuven.be/~nsmart/MPC/>

BBVY21, DLMW15] motivated by an ongoing Microsoft-funded challenge, none discovering an attack which made the community doubt the security LowMC.¹⁹

D Alibi reinsertion

We ensure that $\Pi_{\text{ShufTable.Build}}$ has negligible failure probability (in N) by outputting a small stash, Stash , which we reinsert into the cache. Naïvely, for all $(\llbracket X_i \rrbracket, \llbracket Y_i \rrbracket) \in \text{Stash}$ we could call $L_0.\text{Append}(\llbracket X_i \rrbracket, \llbracket Y_i \rrbracket)$, “reinserting” $(\llbracket X_i \rrbracket, \llbracket Y_i \rrbracket)$ which could not fit in L_i ’s CHT into the cache. Yet, as proved in [FNO21], the naïve reinsertion technique above compromises the obliviousness of (D)ORAM. Intuitively, if we reinsert $(\llbracket X_i \rrbracket, \llbracket Y_i \rrbracket)$ to L_0 from L_j and then query $\llbracket X_i \rrbracket$, in the aggregate, an eavesdropper could tell that we should have queried $\llbracket X_i \rrbracket$ to L_j , but instead queried a dummy because we found $(\llbracket X_i \rrbracket, \llbracket Y_i \rrbracket)$ at some previous level (it was reinserted). Roughly speaking, our goal (*) is to append some data to $(\llbracket X_i \rrbracket, \llbracket Y_i \rrbracket)$ such that when $(\llbracket X_i \rrbracket, \llbracket Y_i \rrbracket)$ is reinserted from L_j and found at L_k for $k < j$, we will know to continue querying L_{k+1}, \dots, L_j as if we did not find $(\llbracket X_i \rrbracket, \llbracket Y_i \rrbracket)$, and query $L_{j+1}, \dots, L_{\text{numLevels}}$ as if we found $(\llbracket X_i \rrbracket, \llbracket Y_i \rrbracket)$ at L_j . To do this, we store $\llbracket \epsilon_{X_i} \rrbracket \in \{0, 1\}^{\text{numLevels}}$ as the last numLevels bits of Y_i (it is assumed that $D > \text{numLevels}$) where $\epsilon_{X_i}[j] = 1$ iff $(\llbracket X_i \rrbracket, \llbracket Y_i \rrbracket)$ was reinserted from L_j . We set $\epsilon_{X_i} = 0^{\text{numLevels}}$ every time we query $(\llbracket X_i \rrbracket, \llbracket Y_i \rrbracket)$. When querying for $\llbracket x \rrbracket$ and finding it at L_k where $\epsilon_x[j] = 1$ from L_j , under MPC we “decide” to query according to (*). The full details of how the Alibi bits are used can be found Figure 4.

E Benchmarking previous DORAMs

We benchmark the (3,1) semi-honest (most efficient) DuORAM variant [VHG22] via their convenient Docker setup on a single c5n.metal machine. We use the `set-networking.sh` script they provide to set .229ms latency and 25Gbit bandwidth simulated network between Docker containers. We do not restrict the number of cores their process can use, and DuORAM used all 96 vCPUs during preprocessing.

We benchmark FLORAM, FLORAM-CPRG, Circuit ORAM, and Square Root ORAM [Ds17, WCS15, ZWR⁺16] via the obliv-c [ZE15] based setup given by [Ds17].²⁰ We benchmark the above 2-party constructions between two cluster-placed c5n.metal machines. For backwards compatibility with obliv-c, we run tests on Ubuntu 18.04.6 LTS. We do not restrict the network via the `tc` command as was done in [Ds17].

Bingsheng Zhang kindly benchmarked the proprietary PFE-DORAM [JZLR22] on a comparable network to ours. Zhang executed the protocol via separate processes on the same Intel(R) Core i7 8700 CPU 3.2 GHz, 6 CPUs, 32 GB Memory,

¹⁹<https://lowmcchallenge.github.io/>

²⁰See <https://gitlab.com/neucrypt/floram/>

1TB SSD machine. Bandwidth between the processes was not limited and latency was restricted to 0.05ms.

We benchmark 3PC-ORAM [JW18] via the dockerization graciously provided by the DuORAM [VHG22] team²¹. Like other constructions, we ran 3PC-ORAM with 0.229ms latency and 25Gbit bandwidth.

F The cost of running GigaDORAM

The c5n.metal machines we rent cost \$3.888 per hour, and running DORAM requires 3 different machines. GigaDORAM. Given that we get ≈ 800 queries per second, (Figure 5), we get $800 \cdot 3600 / (3.888 \cdot 3) = 246914$ queries per USD.

If we benchmark in the same region, *all communication is free*. If we benchmark in different regions, according to AWS the charge per Gigabyte in and out of AWS is \$0.01.²² According to Table 8 it is reasonable to conservatively estimate that GigaDORAM requires $3 \cdot 7 \cdot 10^9 / 100,000 = 210,000$ bytes per query $= 210,000 / 2^{30} = 2 \cdot 10^{-4}$ Gigabytes per query. Multiplying by the dollar cost, we get that GigaDORAM requires $2 \cdot 10^{-4} \cdot .02 = 4 \cdot 10^{-6}$ USD per query.

Summing up communication and computation, we get that we have $1/246914 + 4 \cdot 10^{-6} = 8.05 \cdot 10^{-6}$ which gives approximately 120,000 queries per dollar.

By comparison, DuORAM [VHG22] reports $8900 \cdot 10^{-6}$ dollars for 128 queries (computation cost) giving $128/8900 \cdot 10^{-6}$ queries per dollar. For communication DuORAM gets $5 \cdot 10^{-6}$ dollars per 128 queries, giving $128/5 \cdot 10^{-6}$. This gives $128/8905 \cdot 10^6 = 14,374 \approx 15,000$ queries per dollar.

²¹<https://git-crysp.uwaterloo.ca/iang/circuit-oram-docker>

²²<https://aws.amazon.com/ec2/pricing/on-demand/>