

# MetaDORAM: Info-Theoretic Distributed ORAM with Less Communication

Brett Hemenway Falk<sup>1</sup>   Daniel Noble<sup>2</sup>   Rafail Ostrovsky<sup>3</sup>

<sup>1</sup> University of Pennsylvania, fbrett@seas.upenn.edu

<sup>2</sup> University of Pennsylvania, dgnoble@seas.upenn.edu

<sup>3</sup> UCLA, rafail@cs.ucla.edu

**Abstract.** This paper presents a Distributed Oblivious RAM (DORAM) protocol, MetaDORAM, that is information-theoretically secure and has lower communication cost than all previous info-theoretically secure DORAM protocols for small block sizes. Specifically, given a memory of  $n$  locations, each of size  $d$  bits, MetaDORAM requires only  $O((d + \log^2(n)) \log(n) / \log(\log(n)))$  bits of communication per query. When  $d = \Theta(\log^2(n))$ , this is a  $\Theta(\log(n) / \log \log(n))$  *overhead*, compared to the cost of reading one memory location directly. By comparison, the only existing statistically secure DORAM with sub-logarithmic overhead has communication cost  $O(\log_a(n)d + a\omega(1)\log^2(n)\log_a(n))$  (Abraham et al. PKC '17), where  $\omega(1)$  is any super-constant function in  $n$  and  $a \geq 2$  is a free parameter.

MetaDORAM obtains sub-logarithmic communication overhead for smaller block sizes than previously achieved (any  $d = \omega(\log^2(n) / \log(\log(n)))$ ) while providing statistical security, i.e., no computational assumptions. We circumvent the Goldreich-Ostrovsky lower bound by allowing servers to perform  $\text{poly}(\log(n))$  work, but without computational assumptions. By a standard transformation, our protocol also implies a 3-server active ORAM, Meta3ORAM, with information-theoretic security and  $O((d + \log^2(n)) \log(n) / \log(\log(n)))$  communication per query. For small  $d$ , this is lower than all previous statistically-secure multi-server ORAMs. MetaDORAM and Meta3ORAM also have low communication costs relative to DORAM and multi-server ORAM protocols which make use of computational assumptions. Even compared to several recent works that make use of  $O(n)$  computation, our protocols have lower communication cost. Our protocols are secure in the semi-honest honest-majority setting. We also show that perfectly secure DORAM/multi-server ORAM with the same efficiency can be obtained using a computationally-expensive once-off setup phase.

## 1 Introduction

Secure Multi-Party Computation (MPC) protocols allow a set of parties in a network, of which some unknown subset are dishonest, to securely simulate a trusted third party. This holds tremendous promise. It allows statistical analysis of sensitive data from different sources without pooling data with any single

party. It allows the creation of secure systems, which do not have a single point of failure; using diversification these systems can therefore tolerate attacks affecting particular operating systems, hardware or local networks. In short, it allows sensitive data to be combined and used without needing to trust any single entity or device with that data.

MPC use cases to date, though, have mostly been restricted to tailor-made protocols for specific applications, such as Private Set Intersection [36], Threshold Signature Schemes [11] and Machine Learning [24]. It was shown in the 1980’s that arbitrary circuits could be evaluated securely [43, 17, 4]. However, many computations cannot be represented efficiently as circuits. Rather, it is often more natural and efficient to represent a computation in the RAM model. For instance, RAM is assumed in many classic algorithms and data structures, such as implementations of dictionaries, pointers, graphs and priority queues. An efficient implementation of a RAM functionality for MPC would therefore enable the adoption of generic efficient MPC.

Distributed Oblivious RAM (DORAM) is a functionality that implements RAM for MPC. It stores  $n$   $d$ -bit blocks of data in a secret-shared memory, and allows accesses to that memory (reads or writes) at secret-shared locations. A DORAM is secure if the views of the parties can be efficiently simulated without knowledge of any private values. See section 5 for a complete definition of the DORAM functionality.

DORAM is closely related to the problem of Oblivious RAM, which solves the problem of a client outsourcing memory to an untrusted server (or servers). In particular a  $w$ -party DORAM can be converted into an ORAM with  $w$  active (i.e. computation-performing) servers and vice-versa, usually without increase in the communication cost. See section 2 for more details. Therefore, efficient DORAM is intrinsically tied to efficient multi-server active ORAM. A primary metric of this efficiency is the total amount of communication per memory access. This is often measured as the *overhead*, that is the number of blocks (of size  $d$ ) of communication require per memory access.

**Our Contribution:** In this work, we present MetaDORAM, a novel statistically-secure 3-party DORAM protocol with sub-logarithmic communication overhead. Our protocol achieves amortized communication cost  $\Theta((\log^2(n) + d) \frac{\log(n)}{\log(\log(n))})$  bits per query. MetaDORAM is the first information-theoretic secure DORAM to achieve sub-logarithmic communication overhead when  $d = O(\log^2(n))$ .

Due to the conversion between DORAMs and ORAMs, MetaDORAM can be converted to a 3-server statistically-secure active ORAM, Meta3ORAM, with communication cost  $\Theta((\log^2(n) + d) \frac{\log(n)}{\log(\log(n))})$ .

Note that a passive (non-active) ORAM has a lower bound of  $\Theta(\log(n))$  overhead even if there are multiple servers [26, 27]. Meta3ORAM, like [1], achieves sub-logarithmic overhead only because the servers perform computation. When the servers are active, the  $\Theta(\log(n))$  lower bound can be circumvented for communication, and instead applies to the amount of memory accessed by the servers. Our result, taken with [1], shows that the asymptotic bounds on the DORAM problem are, as of yet, not well understood, and opens up many inter-

esting questions regarding what lower bounds exist for DORAMs, as well as for active information-theoretic ORAMs in general (see Section 9).

MetaDORAM also has lower communication overhead than most DORAM protocols that use computational assumptions. See Table 2 in section 2 for more details.

**Organization:** Our paper is organized as follows. Section 2 provides a short history of prior ORAM and DORAM protocols. Section 3 provides a technical overview of our results and techniques. Section 4 explains the notation used, in particular the various types of secret-sharing used and how they are represented. The formal DORAM functionality is presented in section 5, as well as the functionalities for Secret-Shared Private Information Retrieval (SSPIR) and secure routing, which are used by our DORAM protocol. Section 6 presents our full DORAM protocol, and analyzes its security and communication costs. Sections 7 and 8 explain how SSPIR and secure routing, respectively, can be implemented using standard techniques. Section 9 concludes by discussing some interesting open questions.

## 2 Prior Works

Distributed Oblivious RAM is closely related to the problem of Oblivious RAM (ORAM), which was first formulated by Goldreich in the 1980’s [16]. Imagine a program which is running in a secure environment with very limited memory. The program wishes to make use of general memory on a device, but an adversary may be able to observe access patterns on this device. For instance, the program may be running in a secure enclave, such as Intel SGX [8], but needs to hide sensitive information even if the operating system is corrupted. The program can encrypt the data; this will hide the data’s contents, but will not hide the memory locations accessed by the program, which might leak sensitive information. An ORAM is an intermediary between the program and the main memory. It provides a RAM functionality to the program using the device’s memory in such a way that the access pattern on the device (the *physical* access pattern) reveals no information about the memory accesses by the program (the *virtual* access pattern), except for the number of accesses.

In a normal RAM, the cost of retrieving a block of data from memory is equal to the size of the block, denoted  $d$ . Adding obliviousness comes at a price; the *overhead* is the multiplicative increase in the number of bits that need to be accessed relative to a normal RAM. Goldreich initially presented an ORAM that had  $O(\sqrt{n} \log(n))$  overhead, where  $n$  is the number of blocks of memory [16]. This was improved to  $O(\log^3(n))$  by Ostrovsky [32], using the hierarchical approach, which we explain in section 3. A series of improvements<sup>4</sup>

<sup>4</sup> ORAM security is very subtle and improvements came both in reducing overhead and in fixing security issues in previous protocols. [19] fixed a security issue in [35]; [6] provided a method for oblivious cuckoo hash table construction that was missing from [25]; [20] observed and fixed a flaw that existed in [19], [25], [34] and [3].

reduced this first to  $O(\log^2(n))$  [35, 19], then  $O(\log^2(n)/\log(\log(n)))$  [25, 6], then  $O(\log(n)\log(\log(n)))$  [34] and finally  $O(\log(n))$  [3]. This final result matches the proven asymptotic lower bounds [18, 26]. These works, and the asymptotic lower bound, are in the setting where the untrusted memory is passive (i.e. performs no computation on behalf of the ORAM) and the ORAM only has enough memory to store  $\Theta(1)$  blocks (and  $\Theta(1)$   $\kappa$ -bit PRF keys).

With the advancement of networking infrastructure in the 1990’s, ORAM was quickly viewed as a solution to another challenge: outsourcing memory in a network. Here, a client with limited memory capacity wishes to store data on an untrusted server such that the access pattern on the servers’ memory leaks no information about the actual memory access pattern by the client. The formalism of the original ORAM use-case was immediately applicable, with the client taking the place of the secure environment and the server replacing the device memory. However, this new application led to various extensions of the model. Firstly, the client could feasibly store much more than  $\Theta(d + \kappa)$  bits of memory—a laptop or modern smartphone has gigabytes of memory available. Secondly, the server is likely to also have significant computational resources, so may be able to perform computation to reduce communication overhead, a variant referred to as *active ORAM*. Thirdly, the client could easily interact with multiple servers, which could reasonably be assumed to not collude. This was referred to as *multi-server ORAM*. This new application also caused metrics to be re-evaluated. Since latency is higher in a network, the number of rounds of execution between the client and server became a very important efficiency metric. Additionally, this setting often involved much larger blocks: if the communication cost had a term that did not depend on  $d$ , this term could become asymptotically irrelevant for sufficiently large  $d$ . A large number of works arose examining various combinations of these new models of super-constant client memory overhead (e.g. [42, 39, 10]), active servers, and multiple servers (e.g. [29, 1, 7]).

At the same time, advances in MPC protocols (e.g. [22], [9], [2]) were dramatically reducing the cost of securely evaluating generic circuits. However, many computations are not efficiently realizable as circuits. It was well-known that ORAM techniques could be applied to create efficient MPC protocols in the RAM model [33, 15, 41]. This is the problem of Distributed Oblivious RAM (DORAM), accessing secret-shared memory at secret-shared locations without leaking anything but the number of accesses. Any client-server ORAM could be transformed into a DORAM by evaluating the client’s circuit inside of a secure computation and allowing one of the parties to act as a server. There was now no trusted client, instead there was a “virtual client” that was simulated by a secure computation.

Some of the extensions to the ORAM model that resulted from the memory outsourcing application were immediately relevant to DORAMs. Since DORAMs already had multiple non-colluding parties, they could trivially take advantage of multi-server ORAMs by having the parties simulate different servers. Furthermore, since the parties were already performing computation as part of the

MPC protocol, they could naturally perform the computation needed by servers in an active ORAM. On the other hand, like in the secure program use-case, the DORAM’s virtual client needed to have very limited memory in order for it to have an efficient circuit representation. Furthermore, in an ORAM, the trusted client can perform local computation essentially for free, including cryptographic operations such as PRF evaluations. In the DORAM setting, every computation performed by the virtual client needs to be evaluated inside of a MPC circuit, which can require significant communication between the parties. On the performance metric side, since the MPC protocol occurs in a network, the number of communication rounds again becomes significant, as with memory outsourcing. However, like the application of secure program evaluation, the size of data blocks is often small (say a single variable), so any terms that do not depend on  $d$  are once again significant.

In general, a  $s$ -server active ORAM tolerating  $t$  corrupt parties can be transformed into a  $(w, t)$ -secure DORAM protocol for any  $w \geq s$  by simulating the client in a  $(w, t)$ -secure MPC and having each server’s role taken by a distinct party. This transformation could, potentially, increase the communication cost, depending on the circuit complexity of the virtual client. Going the other way, any  $(w, t)$ -secure DORAM can also be converted to a  $(w, t)$ -secure multi-server active ORAM by each server acting as one party, and by the client initially secret-sharing their query to the servers, and the servers sending the client the shares to reconstruct the result. This will not lead to any increases in asymptotic communication costs.

Through this client simulation, efficient DORAMs could be produced from ORAMs. In fact the best statistically secure DORAMs prior to our work were achieved by simulating a client in an ORAM. If a client can be represented as a Boolean circuit with  $q$  AND gates, it is possible to simulate this client in a secure computation using only  $\Theta(q)$  communication [4, 2]. However, achieving statistical security for generic MPC requires an honest majority, so simulating the client without introducing cryptographic assumptions requires the use of at least 3 parties, even if the ORAM only uses 1 or 2 servers. Wang et al. [41] created a single-server ORAM specifically designed to have a low client circuit complexity. This resulted in a DORAM with communication complexity  $\Theta(\omega(1) \log(n)d + \omega(1) \log^3(n))$ , where  $\omega(1)$  is any function that is super-constant in  $n$ . In comparison, the MetaDORAM protocol is asymptotically better than [41] by a factor of  $\omega(1) \log(\log(n))$  over all parameter ranges.

Abraham et al. created an efficient 2-server ORAM using PIR [1]. Their ORAM is Tree-based (see section 3), in which the number of children of each vertex is a configurable parameter,  $a \geq 2$ . This protocol achieved a communication cost of  $\Theta(\log_a(n)d + a\omega(1) \log_a(n) \log^2(n))$ . The parameter  $a$  should be set to reduce the amortized cost. For  $d \geq 2\omega(1) \log^2(n)$  the cost is minimized by setting  $a = \frac{d}{\omega(1) \log^2(n)}$ , which results in a cost  $\frac{\log(n)}{\log(d)}d$ . For smaller  $d$ , the cost is minimal when  $a = 2$ . Table 1 shows the communication cost for various choices of  $a$ . Their protocol has a simple client, which can be simulated without increasing the asymptotic cost.

In comparison, MetaDORAM has less communication than [1] for small block sizes, and more communication for large block sizes. When  $d = O(\log^2(n))$  MetaDORAM requires  $\Theta\left(\frac{\log^3(n)}{\log(\log(n))}\right)$  communication whereas [1] is optimized when  $a = 2$ , causing [1] to have cost  $\Theta(\omega(1)\log^3(n))$ . For medium blocks, that is  $\omega(1)\log^{2+\epsilon}(n) < d < \log^c(n)$ , for some constants  $\epsilon > 0$ ,  $c > \epsilon + 2$ , both protocols have communication cost  $\Theta\left(\frac{\log(n)}{\log(\log(n))}d\right)$ . For large blocks,  $d = \Omega(n^\epsilon)$ , [1] has better performance, having only  $\Theta(d)$  communication cost, whereas MetaDORAM still has cost  $\Theta\left(\frac{\log(n)}{\log(\log(n))}d\right)$ .

ORAM protocol	DORAM Communication Cost (bits)	Security
Wang et al. [41]	$\Theta(\omega(1)\log(n)d + \omega(1)\log^3(n))$	Statistical
Abraham et al. [1] $a = \frac{d}{\omega(1)\log^2(n)}$ (if $d \geq 2\omega(1)\log^2(n)$ ) $a = 2$ $a = \log(n)$ $a = \sqrt{n}$	$\Theta(\log_a(n)d + a\omega(1)\log_a(n)\log^2(n))$ $\Theta\left(\frac{\log(n)}{\log(d)}d\right)$ $\Theta(\log(n)d + \omega(1)\log^3(n))$ $\Theta\left(\frac{\log(n)}{\log(\log(n))}d + \frac{\omega(1)\cdot\log^4(n)}{\log(\log(n))}\right)$ $\Theta(d + \sqrt{n}\omega(1)\log^2(n))$	Statistical
MetaDORAM (This work)	$\Theta\left(\frac{\log(n)}{\log(\log(n))}d + \frac{\log^3(n)}{\log(\log(n))}\right)$	Statistical

**Table 1.** Complexity of selected ORAM protocols that can be converted to efficient DORAMs. While these ORAMs require only 1 or 2 servers, achieving statistical security requires an honest-majority MPC protocol, therefore necessitating at least 3 parties. These works all have negligible leakage in  $n$ , that is the advantage of an adversary in distinguishing the real and ideal executions is  $2^{-\omega(\log(n))}$ .

Several works also investigated building Distributed Oblivious RAMs directly without simulating a client-server ORAM. While these works generally did not achieve the same asymptotic efficiency as [41] and [1], many had good concrete efficiency. They took advantage of the existence of multiple non-colluding servers by using Distributed Point Functions [5, 40], secret-shared PIR (SSPIR) [23] and secure shuffles/routing [14]. See Table 2 for details. There has also been work to create DORAMs that are secure against malicious adversaries [13, 21]. These works all depended on computational assumptions. In comparison, the MetaDORAM protocol is statistically secure. It also has strictly better communication cost than all protocols over all parameter ranges with one exception: [14] can have a lower asymptotic communication cost, and that only when  $\kappa = o(\log^2(n)/\log(\log(n)))$  and  $d = o(\log^2(n)/\log(\log(n)))$ .

### 3 Technical Overview

Before explaining our protocol, it is helpful to have an understanding of two general paradigms which are used in constructing (D)ORAMs.

Protocol	Communication Cost (bits)	Security
[12]	$O(\omega(1) \log^2(n)d + \kappa\omega(1) \log^4(n))$	Computational
[23]	$O(\log(n)d + \kappa \log^3(n))$	Computational
[5]	$O(d\sqrt{n})$	Computational
[14]	$O(\log(n)d + \kappa \log(n))$	Computational
DuORAM [40]	$O(\kappa \cdot d \cdot \log n)$	Computational
MetaDORAM (this work)	$\Theta\left(\frac{\log(n)}{\log(\log(n))}d + \frac{\log^3(n)}{\log(\log(n))}\right)$	Statistical

**Table 2.** Complexity of several DORAM protocols.  $\kappa$  is a cryptographic security parameter,  $\omega(1)$  is any super-constant function in  $n$ .

The first is the *Hierarchical* approach, initially proposed by Ostrovsky [32], in which data is stored using hash tables. In the ORAM setting the tables are held (encrypted) by a single server, in the DORAM setting, they can be secret-shared between multiple servers. The hash tables use pseudorandom hash functions, so that the physical locations accessed reveal no information about the corresponding indexes, and are hence called *Oblivious Hash Tables* (OHTables). An OHTable does not solve the ORAM problem, however, because an adversary can typically distinguish repeated queries for the *same element* into an OHTable from queries for *distinct* elements. To solve this, when an item is queried it is cached from an OHTable into a small sub-ORAM. The sub-ORAM is queried first and if the item is found there, random locations are accessed in the OHTables; this is necessary for security since re-querying an item in an OHTable would cause the same locations to be accessed, compromising security. To ensure the sub-ORAM remains small, periodically its contents are extracted and built into a new OHTable. To ensure the number of OHTables remains manageable, periodically the contents of multiple OHTables are extracted and rebuilt into a single new OHTable. Typically, the sub-ORAM and OHTables are envisioned as arranged vertically, with the sub-ORAM at the top and OHTables below it, arranged from smallest to largest, resulting in a pyramid-shaped hierarchy (hence the name) of sub-ORAM/OHTable structures.

Shi et al. [37] proposed the alternative *Tree* approach, which was extended by many subsequent (D)ORAMs (e.g. [38, 41, 12]). In this solution data is arranged in a tree, each item is assigned a path from the root to the leaf and the item must remain on this path until its next access. To query an item, its path is first obtained. All locations on this path are accessed and the item is removed from its location. The item is then assigned a new random path and placed in the root of the tree. The root is typically a small sub-ORAM, whereas each other node in the tree typically has capacity for a constant number of items. To prevent the root sub-ORAM from becoming too large, paths from the root to leaves are periodically accessed and each item in the path moved as far down (leafward) as it can be moved, subject to its own path restriction and congestion from other items. Analysing probability distributions shows that the congestion is unlikely to cause the sub-ORAM at the root to overflow. The assignment of

indices to paths, which is called the *position map*, is stored and updated using a sub-ORAM, which is implemented recursively.

Our DORAM combines ideas from Hierarchical and Tree ORAMs. Data is arranged in a hierarchy of OHTables, but the locations of items within each OHTable is stored on a Tree-ORAM style position map. However, MetaDORAM extends the idea of a position map: rather than storing a path of possible locations in which an item may be located, MetaDORAM stores a position *schedule* of each item’s *exact location* at every point in time. Like in Tree ORAMs, a query first accesses a data-structure to gain information about an item. However, rather than just storing a random path, for each index our data structure stores a tag called a Random Unique NameE (rune), a position schedule and a mask schedule. It does this by storing the rune in a sub-DORAM, and having tables that allow the additional metadata to be obtained using only the rune. By accessing this data-structure, it is possible to obtain (a secret-sharing of) the exact location at which an item is stored. The items themselves are stored in a Hierarchy of OHTables. Unlike a typical OHTable, the potential locations of an item are based on the rune alone; no pseudorandom functions are used. During a query, the rune is revealed, which allows the servers to locally narrow down the possible query location to a small number of locations in each table, or  $t$  locations total. Then, the secret-sharing of the full location is used to securely select the correct item. The last step is possible using a simple information-theoretic PIR approach with only  $\Theta(\sqrt{td} + d)$  communication, rather than  $td$  communication (see section 7).

To allow efficient building of data-structures, the servers have different roles. One server is the *Builder*, while the other two servers are  *HOLDERS*. The Holders hold OHTables of masked blocks. For each rune, the Builder always knows where the block for that rune should be located, however it *does not know*, what index that rune corresponds to. Whenever an item is queried the Holders (and *not* the Builder) learn the rune for that item, which allows them to narrow down where the item can be. At the end of the query, the Builder assigns the item a new rune, without learning the associated index, and without the Holders learning the new rune. Therefore, during an access, the Builder learns (or decides) the item’s new rune and the Holders learn the item’s previous rune, but no single party learns both, which provides obliviousness.

During a rebuild, items must be securely moved from locations in smaller tables to locations in a single larger table. The possible locations in which an item can be located within a given table are determined by its rune. Since the Builder knows all of the runes of items that are being built into the table, the Builder can *locally* compute a satisfying assignment for items in the new table. Furthermore, since the Builder knows the previous locations of all items, the Builder can engage in a secure routing protocol with the Holders to permute the elements from their original locations in the old tables to their new locations in the new table.

A query works as follows. First, the current rune for an item is determined using its index, by querying a sub-DORAM. The rune is revealed to the Hold-



ers. This *dramatically* reduces the possible space in which the Holders need to search for the item. We will later use (secret-shared) PIR to retrieve the item from its exact location. A common trick in many (D)ORAM protocols [31, 1] is to use PIR to access the relevant item once its location has been established. In our construction, the PIR is executed over only  $o(\log^3(n))$  locations, which results in low computational cost, and allows us to use information-theoretic PIR protocols.

Next, the protocol needs to determine (secret-shares of) the exact location in which the item is located. Since the Builder pre-chooses the runes, the Builder can also pre-compute at the very beginning of the protocol the location assignments at every point in time. This schedule will be in terms of the rune, and not the indexes they correspond to, which is not yet necessarily determined even by the environment. The Builder can secret-share the position schedule between the Holders. For each rune, the Builder provides  $\Theta(\log(n)/\log(\log(n)))$  secret-shared time ranges and, for each time range, the position during that time range (from among the possible locations for a given rune). During a query, the protocol uses the secret shares of the timestamp ( $\log(n)$  bits), and compares these with the timestamp ranges in order to determine a secret-sharing of the item’s position. This reduces the search space to a new database of polylog size, which is pairwise replicated. Finally, we use information-theoretic Private Information Retrieval (PIR), to obtain a secret-sharing of the desired masked element.

Recall that the Holders store *masked* items, rather than secret-shared items. This allows the Holders to hold exactly the same tables of masked elements (rather than secret-sharings of tables, as is normal in DORAM). This is necessary for the PIR protocol. To achieve information-theoretic security, the items are masked using One-Time Pads (OTPs). In order to prevent the Holders tracking how blocks travel through the OHTables it is necessary for blocks to use a new OTP after each build. Since the Builder never learns the masked blocks, it is safe for the Builder to pick the OTPs. The Builder therefore knows how to mask each item. The secure routing protocol can be extended to allow the Builder to specify OTPs for each block, thereby both permuting and re-masking each block. Furthermore, the Builder can pick all of the OTPs at the start of the protocol, and can create a mask schedule, akin to the position schedule. In fact, since the time-ranges are the same in both cases, the secret-shared OTP can be stored and retrieved with the secret-shared position. Locally XORing the secret-shared OTP and the secret-shared masked block results in a secret-sharing of the item.

Our protocol also makes use of “balancing” [25]. Instead of combining constant numbers of small OHTables to create larger OHTables, we wait until there are  $\Theta(\log^{0.5}(n))$  OHTables of a given size before rebuilding them into a larger OHTable. This means that the total number of tables is  $\Theta(\log^{1.5}(n)/\log \log(n))$ , but the number of times each item participates in a rebuild is only  $\Theta(\log(n)/\log(\log(n)))$ . This reduces the amortized communication cost of rebuilds.

**Novel contributions:** In addition to providing the communication-efficient DORAM with information-theoretic security, our paper introduces a number

of new techniques. Specifically, we use time-stamping and novel data structures to obtain the precise location of data blocks, we make asymmetric use of the participating compute servers to allow efficient and oblivious construction and querying of these data structures, we use as a subroutine tiny-size PIR protocols where the “databases” are constructed on the fly during query execution, and we show a novel strategy for DORAM that bridges techniques from different ORAM strategies in conjunction with ideas explained above.

## 4 Preliminaries

We use lower-case Latin characters to represent parameters in the protocol.  $n$  is the size of the RAM,  $d$  is the bit-length of each item.  $t$  represents the (maximum) number of OHTables that may exist in the protocol.  $h$  represents the number of hash functions used by the hash tables. For an explicit integer or integral parameter,  $a$ ,  $[1, a]$  denotes the set of integers  $\{1, \dots, a\}$ . We use upper-case Latin characters to represent arrays and matrices, which are indexed using standard subscript notation. Most subscripts are one-indexed; zero-indexing is occasionally used when the first item is somehow special.

We use  $\lg$  to represent the base-2 log. For asymptotic annotation, any constant base is equivalent, in which case we often use  $\log$  to represent some arbitrary constant-base log.

We denote the 3 parties as  $P_0$ ,  $P_1$  and  $P_2$ .  $P_0$  is the Builder.  $P_1$  and  $P_2$  are the Holders. The Adversary  $\mathcal{A}$ , is able to corrupt at most one of the parties. The corruption is semi-honest (passive), that is the corrupted party will still follow the protocol, but  $\mathcal{A}$  is able to view all data visible to the corrupted party. The corruption is static, that is  $\mathcal{A}$  cannot change which party is corrupted during the protocol. The protocol is statistically secure. Concretely,  $\mathcal{A}$  learns no information about the access pattern except with probability that is negligible in  $n$ , regardless of  $\mathcal{A}$ 's computational powers.

We utilize hash functions. Our hash functions are fixed and public. We assume that the hash functions are  $2n$ -wise independent and independent of each other. The hash functions implicitly map to ranges of different sizes (depending on the size of the OHTable). In this cases, the hash functions are calculated modulo the required range. It is assumed that the output of the hash function has sufficient entropy that even when reduced modulo these ranges, the distribution is still essentially uniform.

We use several kinds of secret-sharing, all of which are bit-wise (Boolean) secret-sharings. We can explicitly state the sharing using the following notation:  $(\langle y_0 \rangle_0, \langle y_1 \rangle_1, \langle y_2 \rangle_2)$ , which means that  $P_0$  holds share  $y_0$ ,  $P_1$  holds share  $y_1$  and  $P_2$  holds share  $y_2$ . The most common sharing we use is a 3-party replicated secret sharing (3RSS) [2]. Here,  $x \in \{0, 1\}^\ell$  is secret-shared by having  $x_0, x_1, x_2 \in \{0, 1\}^\ell$  that are uniformly random subject to  $x_1 \oplus x_2 \oplus x_3 = x$ .  $P_i$  holds  $x_i$  and  $x_{i+1}$ . (All such subscript operations are implicitly modulo 3.) When variable  $x$  is held using this secret-sharing, it is represented as  $[x]$ . Using the notation above, we can say  $[x] = (\langle (x_0, x_1) \rangle_0, \langle (x_1, x_2) \rangle_1, \langle (x_2, x_0) \rangle_2)$ . Operations (AND, OR,

NOT, XOR) are performed on this secret-sharing using the methods of Araki et al [2].

We also use a 2-party XOR secret-sharing (2XORS), where 2 parties hold the secret-sharing and the third party is not involved. If  $P_1$  and  $P_2$  hold a 2-party XOR secret-sharing of variable  $x$ , this is denoted as  $[x]_{1,2} = (\langle \rangle_0, \langle x \rangle_1, \langle x \rangle_2)$  where  $x_1, x_2 \leftarrow \{0, 1\}^\ell$  subject to  $x_1 \oplus x_2 = x$ . We also use a variant of XOR secret-sharing in which 2 parties hold one of the shares, and the third party holds the other. For instance, when  $P_0$  holds one share, and  $P_1$  and  $P_2$  hold the other share, this is denoted  $[x]_{0,(1,2)} = (\langle x \rangle_0, \langle x \rangle_1, \langle x \rangle_2)$  where  $x_1, x_2 \leftarrow \{0, 1\}^\ell$  subject to  $x_1 \oplus x_2 = x$ . Sometimes a variable is held privately. If  $x$  is held privately, for instance, by  $P_0$ , we denote this as  $[x]_0 = (\langle x \rangle_0, \langle \rangle_1, \langle \rangle_2)$ . Sometimes a variable is known to 2 parties but not the third. If  $x$  is known to  $P_1$  and  $P_2$ , but not  $P_0$ , this is denoted  $[x]_{(1,2)} = (\langle \rangle_0, \langle x \rangle_1, \langle x \rangle_2)$ .

It can be easily shown that for any  $x \in \{0, 1\}^\ell$  it is possible to convert a sharing of  $x$  from any of the above secret-sharings to a fresh resharing of any other of the above sharings with only  $\Theta(\ell)$  bits of communication. For brevity, we provide here only a single example. To covert  $[x]_{0,(1,2)} = (\langle x \rangle_0, \langle x \rangle_1, \langle x \rangle_2)$  to a fresh  $[x]$ ,  $P_0$  picks  $a_1, a_2 \leftarrow \{0, 1\}^\ell$  and sets  $a_3 = x_1 \oplus a_1 \oplus a_2$ .  $P_1$  picks  $b_1, b_2 \leftarrow \{0, 1\}^\ell$  and sets  $b_3 = x_2 \oplus b_1 \oplus b_2$ .  $P_2$  does nothing.  $P_0$  sends  $a_i, a_{i+1}$  to  $P_i$  and likewise  $P_1$  sends  $b_i, b_{i+1}$  to  $P_i$ . Each party  $P_i$  locally calculates  $c_i = a_i \oplus b_i$  and  $c_{i+1} = a_{i+1} \oplus b_{i+1}$  and the new sharing is  $[x] = (\langle (c_0, c_1) \rangle_0, \langle (c_1, c_2) \rangle_1, \langle (c_2, c_0) \rangle_2)$ .

For conciseness, conversions between types of secret-sharing are typically implicit in our pseudocode, indicated by the sharing-type of the result. For instance,  $[C_j]_{(1,2)} = [v_{new}] \oplus [e]$  means that variables  $[v_{new}]$  and  $[e]$ , both stored using 3RSS, are first XORed to create a result that is shared using 3RSS. This result is then revealed to  $P_1$  and  $P_2$  (but not  $P_0$ ), who store the result and label it  $C_j$ .

## 5 Functionality

We wish to implement the following DORAM functionality:

Functionality $\mathcal{F}_{DORAM}$
$I \leftarrow \mathbf{Init}(\mathbf{n}, \mathbf{d}, [\mathbf{A}]$ ): Store array $A$ containing $n$ items of size $d$ . $[v] \leftarrow I.\mathbf{ReadWrite}([x], [y], f)$ : Given an index $x \in [1, n]$ , set $v$ to $A_x$ . Set $A_x = f([v], [y])$ .

Our definition of a DORAM combines the Read and Write functionalities, allowing for reads, writes, or more complex functionalities. This is done by setting the public function  $f$  appropriately. For a read, define  $f(v, y) = v$ . For a write, define  $f(v, y) = y$ . Allowing the written value to be a function of the input provides additional flexibility, such as writing to only particular bits of the data-value or applying a bit-mask to the memory value. Implicitly,  $f$  must be representable using a Boolean circuit containing  $\Theta(d)$  AND gates.

Security is defined using the simulation paradigm, which is standard for proving the security of MPC protocols [28]. A simulator, given only a party’s inputs and outputs from a protocol must generate a view consistent with the real view of a corrupted party during an execution. The DORAM is *perfectly* secure if the simulated view is from the same distribution as the real view. It is *statistically* secure<sup>5</sup> if the distance between the distributions of the views is negligible in  $n$ . It is *computationally* secure if a computationally-bounded adversary has a negligible advantage in distinguishing views in the simulated and real executions. All of our protocols are statistically secure.

Our DORAM implementation makes use of the following functionalities. We show how to implement these using standard techniques in sections 7 and 8 respectively.

**Functionality  $\mathcal{F}_{SSPIR}$**

$[v] \leftarrow \mathbf{SSPIR}(m, d, [A]_{(1,2)}, [x])$ : Given an array  $A$  held (duplicated) by  $P_1$  and  $P_2$ , containing  $m$  elements of size  $d$ , and a share of  $x \in [1, m]$ , return a fresh secret-sharing of  $A_x$ .

**Functionality  $\mathcal{F}_{Route}$**

$[B] \leftarrow \mathbf{Route}([A], [Q]_0)$ : Given a secret-sharing of array  $A$ , of length  $m$ , and an injective mapping  $Q$ , held by  $P_0$ , of length  $q \geq m$ , create a fresh secret-sharing  $B$  such that  $B_{Q(i)} = A_i$  for all  $i \in [1, m]$  and  $B_j$  is distributed uniformly at random for all  $j \notin \{Q(i)\}_{i \in [1, m]}$ .

## 6 DORAM Protocol

### 6.1 Overview

This section presents the DORAM protocol in full and analyzes its security and communication complexity. The protocol is first presented at a high level in section 6.1 and then presented in detail in sections 6.2 and 6.3. Section 6.4 then demonstrates that the protocol achieves the desired security properties. Finally section 6.5 analyzes the security and complexity of the protocol respectively. We assume the existence of functionalities for secret-shared PIR and secure routing, implementations of which are presented in sections 7 and 8 respectively.

Our protocol maintains a data-structure composed of multiple Oblivious Hash Tables (OHTables). The data-structure is stored by two parties, called the Holders ( $P_1$  and  $P_2$ ). They hold identical copies of the data-structure (this will facilitate the use of PIR to access elements). To provide privacy, all items in

<sup>5</sup> Some works specify a security parameter, say  $\sigma$ , such that the statistical distance should be  $2^{-\Theta(\sigma)}$ . We choose instead for the distance negligible in  $n$ , which is equivalent to saying we set  $\sigma = \omega(\log(n))$ .

these Oblivious Hash Tables are masked. The masks are information-theoretic, based on one-time pads.

We refer to the items in the data structure as blocks. Each block is labelled by a random tag, chosen from  $[1, 2n]$ , which we refer to as a rune (Random Unique Name). Although the Holders store the blocks, they are not aware of the runes for these blocks. The other party, called the Builder ( $P_0$ ), *does* know the rune of every block, and knows the block's location in the data-structure, but never stores any blocks itself.

Every index is assigned a rune. The block for that rune holds the (masked) data value at that index's location in memory. When an index is read or written to, it is assigned a new rune and a new block is created which holds the, potentially updated, value. However, the old block, and the old rune, continue to exist in the system. The Builder does not learn that that rune was used, and the Holders do not learn that that block was accessed. Therefore, the data-structure holds both active and obsolete blocks, and the way that blocks are moved through the data-structure does not depend on whether the block is active or obsolete. These obsolete blocks will periodically be deleted, during a refresh phase which happens every  $n$  accesses, which we explain in more detail later.

The mapping of indices to runes is stored in a sub-DORAM. In the sub-DORAM, adjacent indices are stored together, so there are only  $n/2$  indices in the sub-DORAM. Each data-value in the sub-DORAM therefore stores 2 runes, which needs  $\Theta(\log(n))$  space. The sub-DORAM is implemented recursively, so there are  $\Theta(\log(n))$  levels to the recursion.

We now present the full DORAM protocol. We first describe how writes and rebuilds occur. Reads depend on data-structures for the metadata, so we then show how to create these during the initialization and refresh phases. We then show how reads are achieved. Finally we analyze the communication complexity and security of the protocol.

## 6.2 Writes and Rebuilds

We first show how the data-structure storing the blocks is written to and rebuilt. Initially, all blocks are stored in a single, large, OHTable. When an index is queried, it is assigned a new rune, which is picked from the correct distribution by the Builder, and the sub-DORAM is updated with this information. A new block is then created which holds the new value for that index. This block is placed in an area called the *cache*. The cache is filled sequentially. The cache is of size  $\lg^2(n)/\lg \lg(n)$ . When the cache becomes full, its contents are extracted and built into an OHTable.

We implement the OHTable using cuckoo hashing with many hash functions. The block may be stored in locations corresponding to the output of the hash functions *on the block's rune*. Since the Builder knows the runes of every block, the Builder is able to *locally* compute an assignment from runes to locations. It can then collaborate with the Holders to securely route the blocks to their correct locations. It is important for the Holders not to be able to tell how the blocks were permuted. It is therefore necessary to re-mask them. All masks

are achieved information-theoretically using one-time pads (OTPs). The Builder picks the random OTPs each time a block is masked.

We periodically combine multiple OHTables into a single OHTable. Once there are  $b$  OHTables of a given size, the contents of all of these OHTables are extracted and then are built into a single new OHTable. We refer to the OHTables as being arranged in levels. The first, or top, level,  $L_0$ , contains the cache. The next level,  $L_1$ , contains OHTables that were built by extracting the contents of the cache. We label these tables  $T_{1,1}, \dots, T_{1,b}$ . Since the cache is of capacity  $c$ , each OHTable in  $L_1$  will also be of capacity  $c$ .  $L_1$  will contain at most  $b$  such OHTables; when there are  $b$  such tables, they will be combined into an OHTable of size  $bc$  which will be placed in  $L_2$ , and so on. Note that, once the  $b$ th OHTable in a level is built, it is immediately combined with all other OHTables in that level to construct an OHTable in the next level. Therefore, during queries there are only at most  $b - 1$  tables at any level. Since each level's capacity is  $b$  times larger than that of the level before it, a total of  $\Theta(\log_b(n/c))$  levels will be needed to store the blocks created by  $n$  queries. For the parameters we choose,  $\log_b(n/c) = \Theta(\lg(n)/\lg(\lg(n)))$ . After  $n$  queries, the refresh occurs, the contents of all OHTables and the cache are extracted, and the active blocks are rebuilt into a single, large OHTable of size  $n$ , as at the start of the protocol. The Rebuild protocol is presented in Figure 1, together with the overall DORAM ReadWrite function and the Write function.

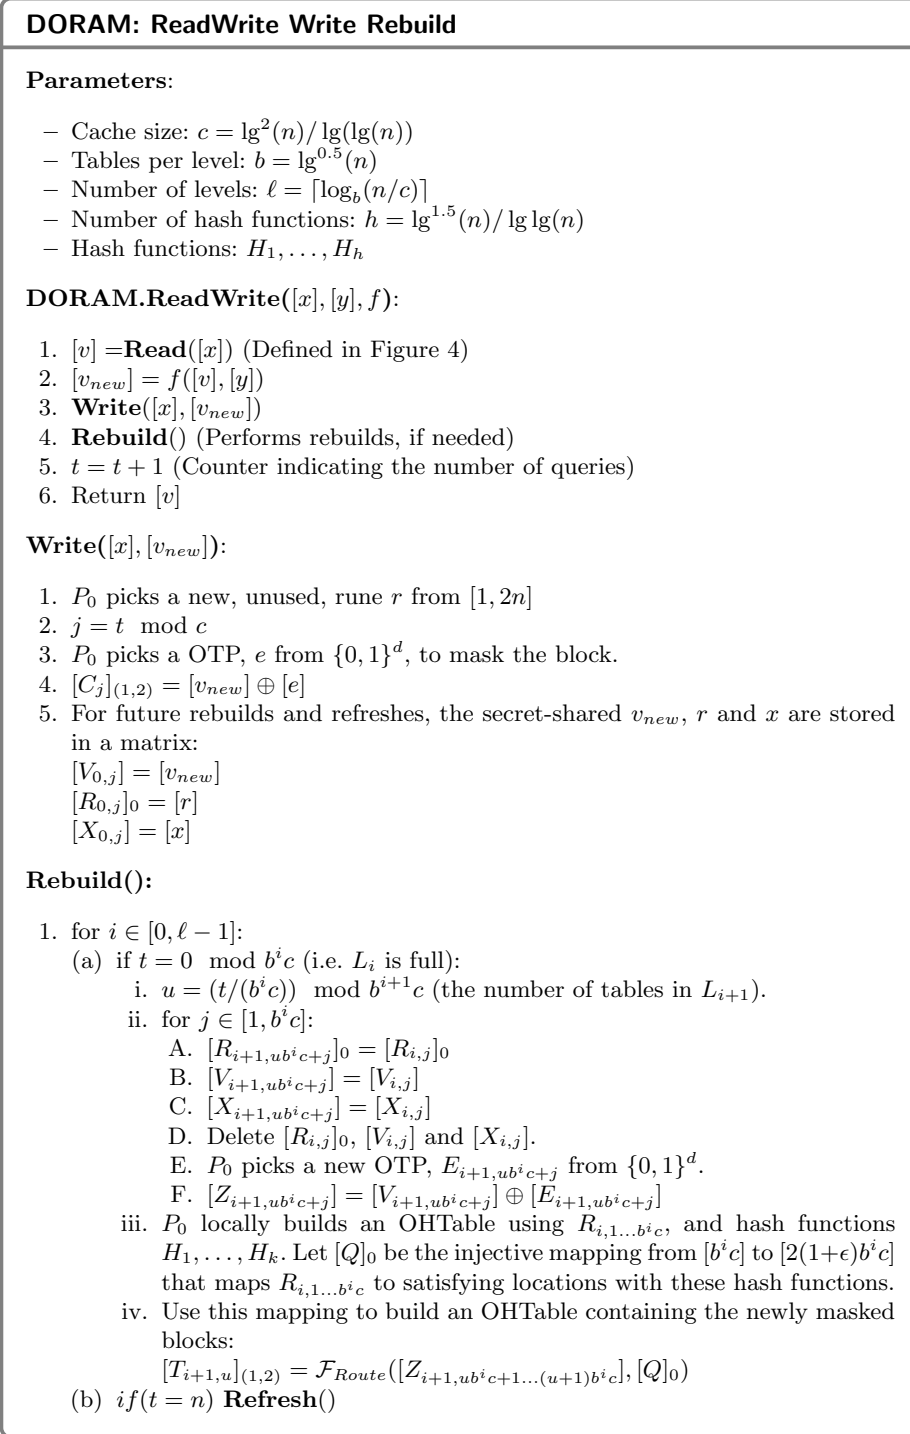
### 6.3 Reads and Refreshes

The question remains as to how the function  $\mathbf{Read}([x])$  can be implemented efficiently. Firstly, we reveal the rune of  $x$  to the Holders, let it be called  $r$ . This greatly simplifies the problem. It is known that the block is stored either in the cache, or in location  $H_k(r)$  of some table  $T_{i,j}$ , for some  $i \in [1, \ell], j \in [1, b-1], k \in [1, h]$ . This reduces the number of possible locations to  $c + \ell(b-1)h = \Theta\left(\frac{\log^3(n)}{\log^2(\log(n))}\right)$ .

This is still a significant number of locations.  $P_0$  knows, for each rune and each time, the location at which each item is stored. However,  $r$  cannot be revealed to  $P_0$  during a read, since  $P_0$  knows when the index with rune  $r$  was last accessed, which would allow  $P_0$  to link access times of indices. In short, the Builder knows the location of each rune, but there seems no way to make use of this without leaking information about the current rune being queried.

Recall that in the description of the DORAM write protocol,  $P_0$  gets to *pick* the rune.  $P_0$  should pick the runes such that each rune is unique, but apart from this runes are chosen uniformly at random from  $[1, 2n]$ . Therefore, the choice of runes does not depend on any other activity in the protocol. Hence,  $P_0$  is able to pick *all of the runes at the beginning of the protocol*. In other words,  $P_0$  can pre-choose the runes that it will assign at each point in time, and during the protocol can assign runes consistently with this original assignment.

Observe, further, that  $P_0$  builds the OHTables based solely on the hash functions and the runes. Since these are both known at the start of the protocol,  $P_0$



**Fig. 1.** DORAM protocol overview, write function and rebuild function

can also pre-calculate all assignments in all hash tables at the beginning of the protocol. This allows  $P_0$  to locally create a *position schedule*, that is a data-structure storing exactly where each rune will be located at each point in time.

This allows us to sidestep the conundrum described above. The Builder can secret-share the position schedule containing all information about the locations of all of the runes, once, at the start of the protocol. The Holders can then access the relevant parts of the position schedule dynamically as they learn the rune of each queried block. Note that this location is the location among all of the possible locations that the block may have been located based on the rune (up to  $c$  cache locations, and up to  $\ell(b-1)h$  table locations).

Given secret-shares of the location of the block, the protocol now engages in a secret-shared PIR (SSPIR) to obtain a secret-sharing of the block. SSPIR can be implemented using a simple modification of any 2-party PIR protocol. The SSPIR protocols we use are explained in more detail in Section 7.

This allows us to obtain secret-sharings of the masked value, but how can this be unmasked?  $P_0$  knows which rune is masked using which OTP, but this information somehow needs to be accessed without revealing to  $P_0$  which rune is being queried. This is the same problem we had with the location mapping, and it can be solved using the same solution! Since the Builder gets to *pick* the OTPs, he can pre-determine, at the initialization of the protocol, which OTPs it will use. He can then secret-share the OTPs that will be used *for all runes at all points in time*. Recall that each time a block is moved, it will be masked using a new OTP. Therefore,  $P_0$  will secret-share a *mask schedule*, analogous to the position schedule, that contains the OTP used to mask each block at each point in time, and which can be accessed dynamically during reads to unmask blocks. This allows us to obtain a secret-sharing of the queried value, performing a read. The read protocol is presented formally in Figure 4.

While we say above that the Builder will pre-determine all runes, locations and ciphertexts at the initialization of the protocol, this is not quite true. A new rune is needed for each query, so an arbitrarily large number of queries would necessitate an arbitrary large number of runes, which cannot be achieved with fixed memory and communication bounds. Instead, we initially set up the system to handle only  $n$  queries. It will therefore have  $2n$  runes ( $n$  initial index runes, and  $n$  which are assigned during the queries). The Builder only predicts the future, so to speak, for the next  $n$  queries, and therefore only creates and shares schedules for locations and masks over  $n$  points in time. After  $n$  queries, the entire system will be refreshed.

We now describe the method for refreshing in more detail. The refresh can be divided into two parts. First, the contents of the up-to-date memory is extracted. This is achieved by randomly permuting all blocks and revealing their runes to the Holders. The Holders know which runes have been queried, so can identify these blocks as obsolete, leaving only the blocks which contain the most recently written value for each index. The extract protocol returns a secret-shared array of the current memory; that is using the same format as that provided for the Init function. The refresh protocol then simply call the Init function using this



secret-shared array to create all of the data-structures necessary for a further  $n$  queries. The Extract functionality is useful in its own right, and may be called by the environment at an arbitrary time (i.e. when there have been fewer than  $n$  queries since the last refresh). The Refresh and Extract protocols are presented formally in Figure 3, while the Init protocol is presented in Figure 2.

**DORAM: Init****Init( $n, d, [V]$ ):**

1.  $P_0$  creates a random permutation which determines the assignment of runes,  $[M]_0 : [1, 2n] \rightarrow [1, 2n]$
2. Assign the first  $n$  of these to be the original runes for the indices. Initialize a new sub-DORAM containing these runes (with adjacent pairs appended together into a single entry).
  - (a) for  $i \in [1, n]$ ,  $[R_i]_0 = [M_i]_0$
  - (b) for  $i \in [1, n/2]$ ,  $[B_i] = [M_{2i-1}]_0 || [M_{2i}]_0$
  - (c) subDORAM =  $\mathcal{F}_{DORAM}.\text{Init}(\frac{n}{2}, 2(\lg(n) + 1), [B])$
3.  $P_0$  locally builds all the OHTables for the next  $n$  queries, based on its knowledge of the runes involved, and the hash functions.  
 If there is *no satisfying assignment* for one of the OHTables,  $P_0$  tells  $P_1$  and  $P_2$  to **abort** the protocol.  
 Otherwise,  $P_0$  can determine where each rune's block will be when, and it creates the *position schedule* which consists of these three matrices:
  - $[S_{i,r}]_0$  contains the time rune  $r$ 's block started to be in its  $i^{th}$  position.
  - $[F_{i,r}]_0$  contains the time rune  $r$ 's block finished to be its  $i^{th}$  position.
  - $[P_{i,r}]_0$  contains the  $i^{th}$  position of rune  $r$ 's block.
4.  $P_0$  creates an mask-schedule. Note that the times will be the same as the position schedule. Therefore all that is needed is one addition matrix containing the OTPs:  
 $[E_{i,r}]_0$  contains the OTP used to mask rune  $r$ 's block when it is in its  $i^{th}$  position.
5.  $P_0$  XOR secret-shares the position schedule and mask schedule between  $P_1$  and  $P_2$ :  $[S]_{1,2}, [F]_{1,2}, [P]_{1,2}, [E]_{1,2}$ .
6.  $P_0$  provides the masks to the blocks, based on his previous selection:  $[E_i]_0 = [E_{0,r}]_0$  for  $M_i = r$ .
7. Based on the Builder's previous assignment of the initial locations of the initial runes, he sets  $[Q]_0$  to be the injection from  $[1, n]$  to  $[1, 2(1 + \epsilon)n]$  that builds the initial table.
8. The parties create the OHTable containing the initial items, and  $P_1$  and  $P_2$  store the masked blocks:  
 $[T_{\ell+1}]_{(1,2)} \leftarrow \mathcal{F}_{Route}([V] \oplus [E]_0, [Q]_0)$
9. The runes, values and indices of the initial items are stored for future reference.  
 That is, for  $i \in [1, n]$ :  $[R_{\ell+1,i}]_0 = [R_i]_0$   
 $[V_{\ell+1,i}] = [V_i]$   
 $[X_{\ell+1,i}] = [i]$
10. Initialize the query counter:  $t = 1$ .

**Fig. 2.** DORAM: Init functionality

### DORAM: Extract and Refresh

$[V] \leftarrow \mathbf{Extract}()$ :

1. Concatenate all (non-deleted)  $R$ ,  $V$  and  $X$  into a single secret-shared array. This will contain all runes that have been used thus far, the index they corresponded to, and the value that was assigned to that index at the time that the rune was assigned:  
 $[R] = [R_0]_0 || [R_1]_0 \dots || [R_{\ell+1}]_0$ ,  $[V] = [V_0] || [V_1] \dots [V_{\ell+1}]$ ,  $[X] = [X_0] || [X_1] \dots [X_{\ell+1}]$
2. Let  $m$  (where  $n \leq m \leq 2n$ ) be the length of these arrays.
3.  $P_1$  picks a random permutation  $S : [1, m] \rightarrow [1, m]$ . Let all items be securely routed according to  $[S]_1$ :  
 $[R] = \mathcal{F}_{Route}([R], [S]_1)$ ,  $[V] = \mathcal{F}_{Route}([V], [S]_1)$ ,  $[X] = \mathcal{F}_{Route}([X], [S]_1)$
4.  $P_2$  similarly picks a random permutation,  $U : [1, m] \rightarrow [1, m]$  which is used to permute all items:  
 $[R] = \mathcal{F}_{Route}([R], [U]_2)$ ,  $[V] = \mathcal{F}_{Route}([V], [U]_2)$ ,  $[X] = \mathcal{F}_{Route}([X], [U]_2)$
5. The values  $R$  are revealed to  $P_1$  and  $P_2$ . Note that  $R$  will contain a random subset of  $m$  items from  $[1, 2n]$ :  $[R]_{(1,2)} \leftarrow [R]$ .
6.  $P_1$  and  $P_2$  identify all runes which have already been revealed to them. The locations of these items in the permuted arrays are made public, and the items are deleted:  
For  $i \in [1, m]$ ,  $I_i = 0$  if  $[R_i]_{(1,2)} \in [D]_{(1,2)}$ , else 1  
If  $I_i = 0$ , delete  $[X_i]$  and  $[V_i]$  (and re-assign indices).
7. Reveal  $[X]$  to all parties. (This will contain all indices in  $[1, n]$  in a random order.) Sort  $[V]$  locally according to  $[X]$ .
8. Return  $[V]$ .
9. Delete all variables and the subDORAM.

**Refresh()**:

1.  $[V] \leftarrow \mathbf{Extract}()$
2.  $\mathbf{Init}(n, d, [V])$

**Fig. 3.** Extract and Refresh functionalities

### DORAM: Read

**Read**( $[x]$ ):

1. Access the subDORAM to learn the rune of  $x$ . Note that indices are stored in the subDORAM in pairs, so the subDORAM will return a share of both  $x$ 's rune and a share of  $x$ 's neighbor's rune. The protocol reveals (only)  $x$ 's rune to  $P_1$  and  $P_2$ . Also, in order to access the subDORAM only once per query, the protocol takes the opportunity to use this access to also write the new rune that is being assigned to  $x$ .
  - (a) Let  $[x_{\ln(n)}]$  be the least significant bit of  $[x]$  (i.e. if  $x$  is odd it is 1, otherwise 0).
  - (b) Set  $[x_{sig}]$  to be the  $\lg(n) - 1$  most significant bits of  $[x]$ , (i.e. drop the last bit).
  - (c)  $P_0$  supplies the new rune  $[r_{new}]_0$  which will be assigned to  $x$  when it is re-written.
  - (d) We define  $f$  to overwrite  $x$  with its new rune, while leaving  $x$ 's neighbor as is. Formally  $f(v, y)$ ,  $v \in \{0, 1\}^{2(\lg n + 1)}$ ,  $y \in \{0, 1\}^{\lg(n) + 2}$  is defined such that if  $y_0 = 0$  (which will happen when  $x$  is even)  $f(v, y) = v_{1, \dots, \lg(n) + 1} || y_{1, \dots, \lg(n) + 1}$  (the second half of the value is overwritten with the remaining bits of  $y$ ) and if  $y_0 = 1$  ( $x$  is odd),  $f(v, y) = y_{1, \dots, \lg(n) + 1} || v_{\lg(n) + 2, \dots, 2\lg(n) + 1}$  (the first half is overwritten).
  - (e)  $[v] \leftarrow \text{subDORAM.ReadWrite}([x_{sig}], [x_{\ln(n)}] || [r_{new}], f)$ .
  - (f) If  $[y_0] = 1$ , securely set  $[r_{old}]$  to be the first half of  $[v]$ , otherwise securely set it to be the second half of  $[v]$ .
  - (g) Reveal  $x$ 's (old) rune to  $P_1$  and  $P_2$ :  $[r]_{(1,2)} \leftarrow [r_{old}]$ .
  - (h) Append  $[r]_{(1,2)}$  to  $[D]_{(1,2)}$ , the set of runes which  $P_1$  and  $P_2$  have already observed.
2.  $P_1$  and  $P_2$  create an array  $Y$  containing all of the (masked) blocks which may hold rune  $r$ 's block:
  - (a)  $[Y_{1, \dots, c}]_{(1,2)}$  contains the blocks from the cache. These are padded to length  $c$  with empty blocks if the cache is not full.
  - (b) For  $i \in [1, \ell + 1]$ ,  $u \in [1, b - 1]$ ,  $k \in [1, h]$ , set  $[Y_{c + (i-1)bh + (u-1)h + k}]_{(1,2)} \leftarrow [T_{i,u,H_k([r]_{(1,2)})}]_{(1,2)}$ . This is, the  $H_k([r])^{th}$  location in table  $T_{i,u}$ . If table  $T_{i,u}$  does not exist, set location to an empty block.
3. Securely determine which time-slot is being used. That is, for  $j \in [0, \ell + 1]$ :
  - (a) Set  $[S_j] \leftarrow [S_{j,[r]_{(1,2)}}]_{1,2} \geq t$
  - (b) Set  $[F_j] \leftarrow [F_{j,[r]_{(1,2)}}]_{1,2} < t$
  - (c) Set  $[J_j]_{1,2} \leftarrow [S_j] \wedge [F_j]$
4. Securely select the correct location and OTP from the position and mask schedules:
  - (a) For  $j \in [0, \ell + 1]$ ,  $[P_j] \leftarrow [P_{j,[r]_{(1,2)}}]_{1,2}$
  - (b) For  $j \in [0, \ell + 1]$ ,  $[E_j] \leftarrow [E_{j,[r]_{(1,2)}}]_{1,2}$
  - (c) For  $j \in [0, \ell + 1]$ , securely set  $[p]$  to  $[P_j]$  if  $[J_j] = 1$
  - (d) For  $j \in [0, \ell + 1]$ , securely set  $[e]$  to  $[E_j]$  if  $[J_j] = 1$
5.  $[v] \leftarrow \mathcal{F}_{\text{BalancedSSPIR}}(c + \ell(b - 1)h, d, [Y]_{(1,2)}, [p]) \oplus [e]$
6. Return  $[v]$

**Fig. 4.** DORAM read protocol

## 6.4 Security Analysis

In this section we show that the DORAM protocol is secure. That is, the views of all participants in the protocol can be efficiently simulated without knowledge of any private values. We show that this security holds in the  $\mathcal{F}_{ABB}$ ,  $\mathcal{F}_{SSPIR}$ ,  $\mathcal{F}_{Route}$ -hybrid model.

All steps of the protocol are one of three cases. Either:

- A secure functionality is being accessed, that only outputs secret-shared results. This can either be a basic ABB functionality, like  $\oplus$ , or a more sophisticated functionality like **SSPIR**.
- The operations are on public, predetermined values (e.g.  $t, u$ ).
- Some value is revealed to some party, or subset of the parties.

We need to examine all revealed values and examine whether they can be simulated without knowledge of the private inputs.

**Init:** No information is revealed to  $P_0$ , rather all private variables it holds are the result of its own random choices (the runes and OTPs) and public parameters (the hash functions).

It is revealed to  $P_1$  and  $P_2$  whether  $P_0$  was able to successfully build all OHTables given his choice of the rune assignment. If  $P_0$  is unable to, the protocol aborts, and no information is leaked as this event happens with fixed probability. If  $P_0$  is able to build all OHTables,  $P_1$  and  $P_2$  learn only this fact, which again occurs with fixed probability. This leaks no information at this point, but has the potential to leak information later, as will be discussed.

$P_1$  and  $P_2$  learn  $T_{\ell+1}$ . All of these blocks have been masked by fresh OTPs, so this is simulatable by generating a uniformly random string.

**Read:** No information is revealed to  $P_0$ .

$P_1$  and  $P_2$  learn the rune queried. The runes are distributed uniformly at random from  $[1, 2n]$ , subject to the fact that they are each unique. Nevertheless, the revealed runes, combined with the knowledge that every OHTable was built successfully, can leak information. This will be analyzed in more detail.

**Write:** No information is revealed to  $P_0$ .

$P_1$  and  $P_2$  learn  $C_j$ . This has been masked using a fresh OTP, so can be simulated by generating a random string.

**Rebuild:** No information is revealed to  $P_0$ .

$P_1$  and  $P_2$  learn  $T_{i,u}$ . This contains blocks which have been masked under fresh OTPs, so can be simulated by generating random strings.

**Extract:**  $P_0$  learns  $X$ . This will contain the items  $[1, n]$  in a randomly permuted order. This can be seen by induction. The protocol maintains the invariant that at each point in time, each index  $x$  has a single rune assigned to it which has not been observed by  $P_1$  and  $P_2$ . In other words, there is a single rune  $R_{i,j}$ , such that  $X_{i,j} = x$  and  $R_{i,j} \notin D$ . Therefore, when the indices corresponding to viewed runes are deleted, a single instance of each index will remain. They will be in a random order because they have been shuffled according to a permutation known to no parties.

$P_0$  also learns  $I$ . This contains  $n$  1s and  $m - n$  0s in a random order, for the

reasons explained above.

$P_1$  and  $P_2$  additionally learn  $R$ . This contains a subset of  $m$  runes from  $[1, 2n]$ . It will necessarily include all  $m - n$  runes from  $D$ , since these runes are definitely stored in the system. The other  $n$  runes are distributed uniformly at random from the set of the remaining  $2n - (m - n)$  runes, so are efficiently simulatable. The ordering must be consistent with  $I$ , that is the  $m - n$  previously observed runes must have  $I_i = 0$ .

Therefore, the only challenging part of the security proof is showing that the distribution of the queried runes (revealed to the Holders), combined with the knowledge that all OHTables were built successfully, does not leak information except with negligible probability. Leakage could occur if some queried set of runes resulted in a set of hash functions that was incompatible with that set being stored in a given OHTable. We show that this does not occur by showing that, for all table capacities  $m \leq n$ , the probability that there exists *any* subset of size  $m$  of the  $2n$  runes that would result in a build failure is negligible.

We prove this making use of Yeo’s analysis of Robust Cuckoo Hashing [44]. Yeo was concerned with an adversary that could pick the indices of items in a hash table, and attempted to pick these such that would cause a build failure, given the predetermined hash functions. His analysis works in general for determining the probability that, given a large set of elements there exists some subset of these that would result in a build failure. Specifically, we can rephrase his Lemma 3 with our notation. Let there be a cuckoo hash table of size  $\Theta(m)$  with  $h$  hash functions. Then the probability that there exists some subset of  $[1, 2n]$  of size  $m$  that results in this cuckoo hash table to have a build failure is at most:

$$\left( \frac{2n}{2^{h-3}} \right)^{h+1}$$

This probability does not depend on  $m$ , except for requiring that  $m \leq 2n$ . We would like this probability to be negligible in  $n$ , i.e.  $\log(1/\epsilon) = \omega(\lg(n))$ . Setting  $h = \lg^{1.5}(n) / \lg(\lg(n)) = \omega(\lg(n))$  achieves this.

This indicates that, for any given OHTable, there is a negligible probability that there exists a set of runes that would be incompatible with this hash table. Since there are  $\text{poly}(n)$  different OHTables ever constructed (in fact only  $\text{polylog}(n)$ , since all tables at all times within a level can re-use the same hash functions), the probability that there is any OHTable in the protocol that has any incompatible set of runes is also negligible in  $n$ . Note that the subDORAMs, even though they have smaller sizes, they should use the same parameter  $h$  as the top level, so that the failure probability remains negligible in the size of the top DORAM,  $n$ .

Therefore, except with negligible probability (over the choice of hash functions), a build failure cannot occur with any choice of runes. This means that, except with negligible probability, there will never be observed a set of runes queried that is incompatible with any allocation of runes to tables.

An interesting corollary of this is that, for most choices of hash functions, the protocol is actually *perfectly* secure, that is *no information* is leaked about the access pattern. Given an  $\exp(n)$ -time setup, it would be possible to test whether a certain choice of hash functions allows for successful builds under all appropriately-sized rune sets, and therefore achieves perfect security. It is not clear whether such a setup for a perfectly-secure protocol can be achieved in  $\text{poly}(n)$  time. Instead this section shows that, over the randomness of the choice of hash functions, the protocol is statistically secure, that is the adversary is unable to distinguish any access patterns, except with probability negligible in  $n$ .

## 6.5 Complexity Analysis

In this section, we show that the amortized communication complexity per access is  $\Theta((\lg^2(n) + d) \lg(n) / \lg(\lg(n)))$  bits. We assume that the cost of  $\mathcal{F}_{\text{BalancedSSPIR}}$  is  $\Theta(\sqrt{md} + d)$  and the cost of  $\mathcal{F}_{\text{Route}}$  is  $\Theta(q(d + \lg(q)))$  as instantiated by our implementations in sections 7 and 8 respectively.

First we analyze the parts of the protocol that have the same cost per-access: reads and writes. We initially analyze only the first level of the recursion. We analyze the number of bits of communication by section, using the same enumeration as the protocols.

### Read:

1. The rune of the index is accessed and a new rune written. Apart from the call to the subDORAM, which will be analyzed later, this involves only operations on runes, each of which requires at most  $\Theta(\log(n))$  AND gates, or revealing  $\Theta(\log(n))$  bits, so  $\Theta(\log(n))$  communication.
2. The Holders arrange the blocks which may hold the rune's block. This requires only local operations and no communication.
3. The time slot is obtained. This requires  $\Theta(\ell) = \Theta(\log(n) / \log(\log(n)))$  comparisons of  $\Theta(\log(n))$ -bit values, which requires  $\Theta(\log^2(n) / \log(\log(n)))$  communication.
4. The correct position and mask is obtained. This requires  $\Theta(\ell) = \Theta(\log(n) / \log(\log(n)))$  secure if-then-else statements on  $\Theta(\log(n))$ -bit and  $\Theta(d)$ -bit values for the positions and masks respectively. The total is therefore  $\Theta((\log(n) + d) \log(n) / \log(\log(n)))$  communication.
5. Finally the SSPIR is executed. The number of locations is  $c + \ell(b - 1)h = \Theta(\log^3(n) / \log^2(\log(n)))$ . Therefore, the cost of the Balanced SSPIR protocol is  $\Theta(\sqrt{\log^3(n)d / \log^2(\log(n))} + d) = \Theta(\log(n) / \log(\log(n)) \sqrt{\log(n)d} + d)$ . For  $d = \Omega(\log(n))$ ,  $\sqrt{\log(n)d} = O(d)$ , so the cost above simplifies to  $O(\log(n)d / \log(\log(n)))$ .

### Write:

1. The first 3 steps are either local to  $P_0$ , or on public values, so require no communication

2. The masked block is created, required  $\Theta(d)$  communication
3. The final steps consist only of re-labelling variables and operations on public values, so require no communication.

Therefore the communication cost of the write is  $\Theta(d)$ .

We next analyze the communication cost of the Rebuild function (excluding the refresh function). The communication cost of this function is variable, so we calculate the average cost per access.

**Rebuild:**

A level of capacity  $m$  is rebuilt every  $m$  accesses. Most steps are simply relabelling of variables, which require no communication. The steps that require communication are:

- The Builder secret-shares the new OTP for each item, which costs  $\Theta(md)$ .
- The Routing protocol, which requires  $\Theta(m(d + \lg(n)))$  communication.

Therefore, the amortized cost per access per level is  $\Theta(\lg(n) + d)$ . Since there are  $\Theta(\lg(n)/\lg(\lg(n)))$  levels, the total communication cost per access is  $\Theta((\lg(n) + d)\lg(n)/\lg(\lg(n)))$ .

**Extract:**

1. Concatenating the arrays requires only local relabelling of variables, except for the runes which are reshared from  $P_0$  to being shared by all parties, at communication cost  $\Theta(n \lg(n))$ .
2. Setting  $m$  is a local operation.
3.  $m = \Theta(n)$  elements are routed, each of size  $\Theta(\log(n) + d)$  resulting in  $\Theta(n(\log(n) + d))$  communication.
4. The same occurs again, resulting in  $\Theta(n \log(n) + d)$  communication.
5. Revealing all runes to Holders requires  $\Theta(n \log(n))$  communication.
6. Holders reveal  $m = \Theta(n)$  bits, hence  $\Theta(n)$  communication.
7. Revealing all (permuted) indices requires  $\Theta(n \log(n))$  communication.
8. The last 2 steps are local operations.

Since this occurs every  $n$  accesses, the cost is  $\Theta(\log(n) + d)$  communication per access.

**Init:**

1. The rune assignment is local, so has no communication.
2. The cost of initializing the subDORAM will be evaluated as part of the cost of recursion.
3. Creating the position schedule is a local operation
4. Creating the mask schedule is a local operation
5. The position schedule has  $\Theta(n)$  columns (for the runes),  $\Theta(\ell) = \Theta(\log(n)/\log \log(n))$  rows (for the levels) and has  $O(\log(n))$  bits per cell, for both the timestamp representations and the position representations. Each cell of the mask schedule is  $\Theta(d)$  bits. Therefore the total cost of secret-sharing the position and mask schedules is  $\Theta((\log n + d)n \log(n)/\log(\log(n)))$  communication.
6. Selecting the pre-chosen OTPs is a local operation.



7. Assigning the mapping to build the OHTable is a local operation.
8. Secret-sharing the mask, and routing the blocks requires a total of  $\Theta((\log(n)+D)n)$  communication.
9. The last step is a local relabelling.

Therefore, the total cost is  $\Theta((\log(n) + d)n \log(n) / \log(\log(n)))$ , or  $\Theta((\log(n) + d) \log(n) / \log(\log(n)))$  per access.

Summing these up, we obtain that the cost at the first level of the recursion is  $\Theta((\log(n) + d) \log(n) / \log(\log(n)))$ . In the first level of the recursion, the block size  $d$  can be arbitrary. However, for the recursively implemented subDORAM, the block size is always  $\Theta(\log(n))$ . Therefore, each level of the recursion has cost  $\Theta(\log^2(n) / \log(\log(n)))$ . There are  $\Theta(\log(n))$  such levels, so the cost of the recursive calls is  $\Theta(\log^3(n) / \log(\log(n)))$ . Hence, the total communication cost per access is  $\Theta((\log^2(n) + d) \log(n) / \log(\log(n)))$ .

While our focus is amortized total communication per query, for completeness we also provide below the performance of our protocol by other metrics. The total memory required by the protocol is  $\Theta(\log(n)dn / \log(\log(n)))$ : this is dominated by the size of the mask matrix (assuming  $d = \Omega(\log(n))$ ) which must be held in memory by  $P_1$  and  $P_2$ . The round-complexity is dominated by the cost of evaluating inequality tests (in step 3 of **Read**) which uses a circuit with AND-depth  $\Theta(\log(\log(n)))$  and therefore needs  $\Theta(\log(\log(n)))$  rounds. This is done sequentially in all  $\Theta(\log(n))$  recursions of the subDORAM, leading to a total round complexity per query of  $\Theta(\log(n) \log(\log(n)))$ . The computation cost depends on the hash function implementation, and in most cases would be dominated by the evaluation of  $\ell h = \Theta(\log^{2.5}(n) / \log^2(\log(n)))$  hash functions per recursion level, or a total of  $\Theta(\log^{3.5}(n) / \log^2(\log(n)))$  hash function evaluations per query. The protocol accesses  $c + \ell(b-1)h = \Theta(\log^3(n) / \log^2(\log(n)))$  memory locations of size  $d$  in the top level, and  $c + \ell(b-1)h$  memory locations of size  $\Theta(\log(n))$  in each of the recursive levels, resulting in a total of  $\Theta(\log^3(n)d / \log^2(\log(n)) + \log^5(n) / \log^2(\log(n)))$  bits of memory accessed per query.

## 7 Secret-Shared Private Information Retrieval

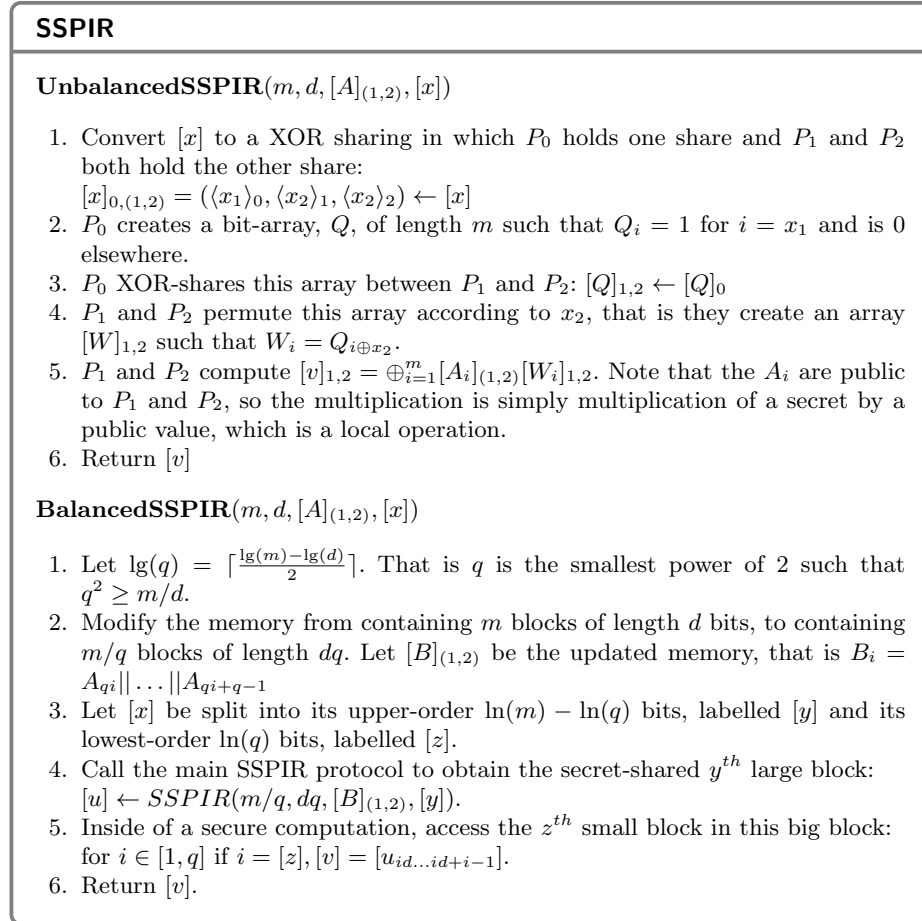
This section presents a simple protocol for secret-shared Private Information Retrieval that is optimized for our use-case.

In general Private Information Retrieval protocols are designed for the case that a single bit is to be retrieved. However in our protocols we need to retrieve  $d$  bits, which all occur in a single location in memory. We therefore use the following “naïve” PIR protocol. Let  $x$  be the secret location, and  $m$  the length of the memory, that is  $1 \leq x \leq m$ . The secret location is represented using a  $m$ -bit array, which is 0 everywhere except for the  $x^{\text{th}}$  bit, which is 1. This array is secret-shared between the two parties, who can locally compute a dot-product of this string with their memory, to obtain a secret-sharing of the desired element. While this PIR protocol has a query of length  $m$ , a single query can be used

regardless of the bit-length  $d$ . That is the same query string is used for all  $d$  bits of the data. The cost is therefore  $\Theta(m + d)$ .

The above protocol assumes that there is a PIR client who can safely learn the location  $x$ . It is possible to apply a transformation to obtain a *secret-shared* PIR protocol. This technique was used, for instance, in the “Data-Rotations” of [12]. The PIR servers ( $P_1$  and  $P_2$ ) are given a location mask  $x_2$ , and locally permute their array according to this mask, such that each item is moved from location  $i$  to location  $i \oplus x_2$ . The PIR client ( $P_0$ ) then searches for a location  $x_1 = x \oplus x_2$ . This will clearly hold the index that was at location  $x$ . The security of the PIR protocol hides the query from the PIR servers. The client only receives  $x_1$  which is a uniform random value.

The protocol is presented in full in figure 5.



**Fig. 5.** Implementation of SSPIR

The SSPIR protocol (figure 5) is secure.  $P_0$  receives only  $x_0$  which is a uniform random value.  $P_1$  and  $P_2$  receive only shares of  $Q$ , which are uniform random bit arrays. The protocol is deterministic and secure.

The protocol presented above has communication cost  $\Theta(m + d)$ . For some situations this is sufficient. However, when  $m = \omega(d)$  it is possible to increase the size of data-blocks to achieve improved complexity, effectively “balancing” the  $m$  and  $d$  terms. We do this by increasing the block size from  $d$  to  $qd$ , for some balancing factor  $q > 1$ , where  $q$  is a power of 2.

We present the balanced PIR protocol in the second part of figure 5. All operations are inside of a secure computation, so the protocol is secure. The cost of the call to the main SSPIR protocol is  $\Theta(m/q + dq)$ . Additionally, there is a cost of  $\Theta(qd)$  to securely select the relevant small block. The total cost is therefore  $\Theta(m/q + dq)$ . Our protocol picks the optimum  $q = \Theta(\sqrt{m/d} + 1)$  which results in a communication cost of  $\Theta(\sqrt{md} + d)$ . (The last term in both equations comes from the case when  $m = O(d)$ .)

## 8 Secure Routing

We here present an implementation of a secure 3-party routing protocol. That is, there is some secret-shared array  $A$  of length  $m$  and one party knows an injective mapping  $Q$  from  $[1, m]$  to  $[1, q]$ , (where  $q \geq m$ ). The items are moved to a new secret-shared array  $B$  such that  $A_i$  is moved to some location  $B_j$  where  $j = Q(i)$ . See section 5 for a formal definition of the functionality. Variants of this protocol have occurred before, for instance as the protocol  $\Pi_{SWITCH}$  in [30]. We include the protocol here for clarity and completeness. The protocol is presented in figure 6, and is analyzed below.

**Security:**  $P_0$ , knowing a desired permutation, secret-shares this permutation between  $P_1$  and  $P_2$ , providing them permutation shares  $R$  and  $S$  respectively. Each of these permutation-shares is distributed as a uniformly random permutation, and leaks no information about the true permutation  $Q$ . Apart from that, parties only receive secret-shares, which are distributed uniformly at random.

**Complexity:** Communicating the permutations requires  $\Theta(q \lg(q))$  communication. There are a constant number of resharings of arrays, each of which contains  $q$  elements of size  $d$  bits, resulting in  $\Theta(qd)$  communication. The total communication cost is therefore  $\Theta((d + \log(q))q)$ .

## 9 Conclusion and Future Work

In this work, we construct an information-theoretic DORAM with  $O((d + \log^2(n)) \log(n) / \log(\log(n)))$  bits of communication per query, which is *below* the lower bound on communication for passive ORAMs.

So what is the correct communication lower bound in DORAM setting? [1] show that constant overhead is possible for polynomial-sized blocks; is it possible for polylogarithmic sized blocks? Also, does the lower-bound introduce

### Routing

**Route**( $[A], [Q]_0, d$ ):

1. Pad  $[A]$  to length  $q$  (if  $q > m$ ) with random values: for  $i \in [1, m]$ ,  $[B_i] = [A_i]$  for  $i \in [m + 1, q]$ ,  $[B_i] \leftarrow \{0, 1\}^d$ .
2.  $P_0$  picks permutations  $R$  and  $S$  which are chosen uniformly at random subject to  $R \cdot S = Q$  (over domain  $[1, m]$  and is an arbitrary permutation elsewhere).  $P_0$  sends  $S$  to  $P_1$  and sends  $R$  to  $P_2$ :  
 $[S]_{(0,1)} \leftarrow [S]_0$   
 $[R]_{(0,2)} \leftarrow [R]_0$
3. Reshare  $B$  to  $P_0$  and  $P_1$ :  $[B]_{0,1} \leftarrow [B]$
4.  $P_0$  and  $P_1$  locally permute  $[B]_{0,1}$  according to  $[S]_{(0,1)}$  to obtain  $[C]_{0,1}$ .
5. Reshare  $C$  to  $P_0$  and  $P_2$ :  $[C]_{0,2} \leftarrow [C]_{0,1}$
6.  $P_0$  and  $P_2$  locally permute  $[C]_{0,2}$  according to  $[R]_{(0,2)}$  to obtain  $[D]_{(0,2)}$
7. Return  $[D]$

**Fig. 6.** Secure Routing protocol

trade-offs between the amount of communication and the number of memory locations accessed, the total memory required, the computation cost, the randomness consumed or the round complexity?

Another open question pertains to deamortization. MetaDORAM achieves *amortized* communication cost  $\Theta((d + \log^2(n)) \log(n) / \log(\log(n)))$ . In particular, the cost of rebuilding the OHTables, and refreshing the namespace is amortized across multiple queries. There are standard techniques for deamortizing the cost of building OHTables, but it seems more challenging to deamortize the cost of refreshing, in particular the cost of refreshing the namespace for runes by reassigning runes to all indices. So the question remains: is it possible to have a DORAM with *worst-case* communication cost  $\Theta((d + \log^2(n)) \log(n) / \log(\log(n)))$ ?

Every DORAM can be used to implement a multi-server active ORAM: the client in the multi-server ORAM can simply secret-share the query between the servers. So our construction implies that  $\Theta((\log^2 n + d) \log(n) / \log \log(n))$  communication can be achieved, without computational assumptions. An interesting final open question raised by this work is whether this is possible for a single-server active ORAM. Due to existing lower bounds, such an ORAM would need to access at least a logarithmic overhead in memory and therefore perform a logarithmic overhead of computation, but this computation could, perhaps, avoid the introduction of computational assumptions. Concretely is it possible to have a single-server active ORAM that has sub-logarithmic communication overhead, but is information-theoretically secure?

## Acknowledgements

This research was 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 2022370, Amazon Faculty Award, Cisco Research Award, Sunday Group, ONR grant (N00014-15-1-2750) “SynCrypt: Automated Synthesis of Cryptographic Constructions” and a gift from Ripple Labs, Inc. Daniel Noble would also like to acknowledge God for supporting him in this research. 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

1. Abraham, I., Fletcher, C.W., Nayak, K., Pinkas, B., Ren, L.: Asymptotically tight bounds for composing ORAM with PIR. In: IACR International Workshop on Public Key Cryptography. pp. 91–120. Springer (2017)
2. Araki, T., Furukawa, J., Lindell, Y., Nof, A., Ohara, K.: High-throughput semi-honest secure three-party computation with an honest majority. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security. pp. 805–817 (2016)
3. Asharov, G., Komargodski, I., Lin, W.K., Nayak, K., Peserico, E., Shi, E.: OptORAMA: optimal oblivious RAM. In: Advances in Cryptology–EUROCRYPT 2020: 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, May 10–14, 2020, Proceedings, Part II 30. pp. 403–432. Springer (2020)
4. Ben-Or, M., Goldwasser, S., Wigderson, A.: Completeness theorems for noncryptographic fault-tolerant distributed computations. In: Proceedings of the 20th Annual Symposium on the Theory of Computing (STOC’88). pp. 1–10 (1988)
5. Bunn, P., Katz, J., Kushilevitz, E., Ostrovsky, R.: Efficient 3-party distributed ORAM. In: Security and Cryptography for Networks: 12th International Conference, SCN 2020, Amalfi, Italy, September 14–16, 2020, Proceedings 12. pp. 215–232. Springer (2020)
6. Chan, T.H.H., Guo, Y., Lin, W.K., Shi, E.: Oblivious hashing revisited, and applications to asymptotically efficient ORAM and OPRAM. In: Advances in Cryptology–ASIACRYPT 2017: 23rd International Conference on the Theory and Applications of Cryptology and Information Security, Hong Kong, China, December 3–7, 2017, Proceedings, Part I 23. pp. 660–690. Springer (2017)
7. Chan, T.H.H., Katz, J., Nayak, K., Polychroniadou, A., Shi, E.: More is less: Perfectly secure oblivious algorithms in the multi-server setting. In: Advances in Cryptology–ASIACRYPT 2018: 24th International Conference on the Theory and Application of Cryptology and Information Security, Brisbane, QLD, Australia, December 2–6, 2018, Proceedings, Part III 24. pp. 158–188. Springer (2018)
8. Costan, V., Devadas, S.: Intel SGX explained. Cryptology ePrint Archive (2016)

9. Damgård, I., Pastro, V., Smart, N., Zakarias, S.: Multiparty computation from somewhat homomorphic encryption. In: Annual Cryptology Conference. pp. 643–662. Springer (2012)
10. Devadas, S., van Dijk, M., Fletcher, C.W., Ren, L., Shi, E., Wichs, D.: Onion oram: A constant bandwidth blowup oblivious ram. In: Theory of Cryptography: 13th International Conference, TCC 2016-A, Tel Aviv, Israel, January 10-13, 2016, Proceedings, Part II 13. pp. 145–174. Springer (2016)
11. Doerner, J., Kondi, Y., Lee, E., shelat, A.: Threshold ecdsa from ecdsa assumptions: The multiparty case. In: 2019 IEEE Symposium on Security and Privacy (SP). pp. 1051–1066. IEEE (2019)
12. Faber, S., Jarecki, S., Kentros, S., Wei, B.: Three-party ORAM for secure computation. In: International Conference on the Theory and Application of Cryptology and Information Security. pp. 360–385. Springer (2015)
13. Falk, B., Noble, D., Ostrovsky, R., Shtepel, M., Zhang, J.: Doram revisited: Maliciously secure ram-mpc with logarithmic overhead. In: TCC (2023)
14. Falk, B.H., Noble, D., Ostrovsky, R.: 3-party distributed ORAM from oblivious set membership. In: International Conference on Security and Cryptography for Networks. pp. 437–461. Springer (2022)
15. Gentry, C., Goldman, K.A., Halevi, S., Julta, C., Raykova, M., Wichs, D.: Optimizing ORAM and using it efficiently for secure computation. In: Privacy Enhancing Technologies: 13th International Symposium, PETS 2013, Bloomington, IN, USA, July 10-12, 2013. Proceedings 13. pp. 1–18. Springer (2013)
16. Goldreich, O.: Towards a theory of software protection and simulation by oblivious RAMs. In: Proceedings of the nineteenth annual ACM symposium on Theory of computing. pp. 182–194 (1987)
17. Goldreich, O., Micali, S., Wigderson, A.: How to play any mental game, or a completeness theorem for protocols with honest majority. In: Providing Sound Foundations for Cryptography: On the Work of Shafi Goldwasser and Silvio Micali, pp. 307–328 (2019)
18. Goldreich, O., Ostrovsky, R.: Software protection and simulation on oblivious RAMs. *Journal of the ACM (JACM)* **43**(3), 431–473 (1996)
19. Goodrich, M.T., Mitzenmacher, M.: Privacy-preserving access of outsourced data via oblivious RAM simulation. In: International Colloquium on Automata, Languages, and Programming. pp. 576–587. Springer (2011)
20. Hemenway Falk, B., Noble, D., Ostrovsky, R.: Alibi: A flaw in cuckoo-hashing based hierarchical oram schemes and a solution. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 338–369. Springer (2021)
21. Ichikawa, A., Komargodski, I., Hamada, K., Kikuchi, R., Ikarashi, D.: 3-party secure computation for rams: Optimal and concretely efficient (2023)
22. Ishai, Y., Kilian, J., Nissim, K., Petrank, E.: Extending oblivious transfers efficiently. In: Annual International Cryptology Conference. pp. 145–161. Springer (2003)
23. Jarecki, S., Wei, B.: 3PC ORAM with low latency, low bandwidth, and fast batch retrieval. In: Applied Cryptography and Network Security: 16th International Conference, ACNS 2018, Leuven, Belgium, July 2-4, 2018, Proceedings 16. pp. 360–378. Springer (2018)
24. Knott, B., Venkataraman, S., Hannun, A., Sengupta, S., Ibrahim, M., van der Maaten, L.: Crypten: Secure multi-party computation meets machine learning. *Advances in Neural Information Processing Systems* **34**, 4961–4973 (2021)

25. Kushilevitz, E., Lu, S., Ostrovsky, R.: On the (in) security of hash-based oblivious RAM and a new balancing scheme. In: SODA (2012)
26. Larsen, K.G., Nielsen, J.B.: Yes, there is an oblivious RAM lower bound! In: Annual International Cryptology Conference. pp. 523–542. Springer (2018)
27. Larsen, K.G., Simkin, M., Yeo, K.: Lower bounds for multi-server oblivious RAMs. In: Theory of Cryptography: 18th International Conference, TCC 2020, Durham, NC, USA, November 16–19, 2020, Proceedings, Part I 18. pp. 486–503. Springer (2020)
28. Lindell, Y.: How to simulate it—a tutorial on the simulation proof technique. Tutorials on the Foundations of Cryptography: Dedicated to Oded Goldreich pp. 277–346 (2017)
29. Lu, S., Ostrovsky, R.: Distributed oblivious RAM for secure two-party computation. In: Theory of Cryptography Conference. pp. 377–396. Springer (2013)
30. Mohassel, P., Rindal, P., Rosulek, M.: Fast database joins and psi for secret shared data. In: Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security. pp. 1271–1287 (2020)
31. Mughees, M.H., Chen, H., Ren, L.: OnionPIR: Response efficient single-server PIR. In: Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security. pp. 2292–2306 (2021)
32. Ostrovsky, R.: Efficient computation on oblivious RAMs. In: Proceedings of the twenty-second annual ACM symposium on Theory of computing. pp. 514–523 (1990)
33. Ostrovsky, R., Shoup, V.: Private information storage (extended abstract). In: Leighton, F.T., Shor, P.W. (eds.) Proceedings of the Twenty-Ninth Annual ACM Symposium on the Theory of Computing, El Paso, Texas, USA, May 4–6, 1997. pp. 294–303. ACM (1997). <https://doi.org/10.1145/258533.258606>, <https://doi.org/10.1145/258533.258606>
34. Patel, S., Persiano, G., Raykova, M., Yeo, K.: PanORAMa: Oblivious RAM with logarithmic overhead. In: 2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS). pp. 871–882. IEEE (2018)
35. Pinkas, B., Reinman, T.: Oblivious RAM revisited. In: Advances in Cryptology—CRYPTO 2010: 30th Annual Cryptology Conference, Santa Barbara, CA, USA, August 15–19, 2010. Proceedings 30. pp. 502–519. Springer (2010)
36. Pinkas, B., Rosulek, M., Trieu, N., Yanai, A.: Spot-light: lightweight private set intersection from sparse set extension. In: Advances in Cryptology—CRYPTO 2019: 39th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18–22, 2019, Proceedings, Part III 39. pp. 401–431. Springer (2019)
37. Shi, E., Chan, T.H.H., Stefanov, E., Li, M.: Oblivious RAM with  $o((\log n)^3)$  worst-case cost. In: Advances in Cryptology—ASIACRYPT 2011: 17th International Conference on the Theory and Application of Cryptology and Information Security, Seoul, South Korea, December 4–8, 2011. Proceedings 17. pp. 197–214. Springer (2011)
38. Stefanov, E., van Dijk, M., Shi, E., Chan, T.H.H., Fletcher, C., Ren, L., Yu, X., Devadas, S.: Path ORAM: an extremely simple oblivious RAM protocol. Journal of the ACM (JACM) **65**(4), 1–26 (2018)
39. Stefanov, E., Shi, E., Song, D.: Towards practical oblivious ram. arXiv preprint arXiv:1106.3652 (2011)
40. Vadapalli, A., Henry, R., Goldberg, I.: DuORAM: A bandwidth-efficient distributed ORAM for 2-and 3-party computation. In: 32nd USENIX Security Symposium (2023)

41. Wang, X., Chan, H., Shi, E.: Circuit ORAM: On tightness of the Goldreich-Ostrovsky lower bound. In: Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security. pp. 850–861 (2015)
42. Williams, P., Sion, R.: Single round access privacy on outsourced storage. In: Proceedings of the 2012 ACM conference on Computer and communications security. pp. 293–304 (2012)
43. Yao, A.: Protocols for secure computations (extended abstract). In: FOCS (1982). <https://doi.org/10.1109/SFCS.1982.88>, <http://dx.doi.org/10.1109/SFCS.1982.88>
44. Yeo, K.: Cuckoo hashing in cryptography: Optimal parameters, robustness and applications. arXiv preprint arXiv:2306.11220 (2023)