

Distributed Oblivious RAM for Secure Two-Party Computation*

Steve Lu¹ and Rafail Ostrovsky²

¹ Stealth Software Technologies, Inc., USA

Email: steve@stealthsoftwareinc.com

² Department of Computer Science and Department of Mathematics, UCLA.

Work done with consulting for Stealth Software Technologies, Inc.

Email: rafail@cs.ucla.edu

Abstract. We present a new method for secure two-party Random Access Memory (RAM) *program* computation that does not require taking a program and first turning it into a circuit. The method achieves logarithmic overhead compared to an insecure program execution.

In the heart of our construction is a new Oblivious RAM construction where a client interacts with two non-communicating servers. Our two-server Oblivious RAM for n reads/writes requires $O(n)$ memory for the servers, $O(1)$ memory for the client, and $O(\log n)$ amortized read/write overhead for data access. The constants in the big- O notation are tiny, and we show that the storage and data access overhead of our solution concretely compares favorably to the state-of-the-art single-server schemes. Our protocol enjoys an important feature from a practical perspective as well. At the heart of almost all previous single-server Oblivious RAM solutions, a crucial but inefficient process known as oblivious sorting was required. In our two-server model, we describe a new technique to bypass oblivious sorting, and show how this can be carefully blended with existing techniques to attain a more practical Oblivious RAM protocol in comparison to all prior work.

As alluded above, our two-server Oblivious RAM protocol leads to a novel application in the realm of secure two-party RAM program computation. We observe that in the secure two-party computation, Alice and Bob can play the roles of two non-colluding servers. We show that our Oblivious RAM construction can be composed with an extended version of the Ostrovsky-Shoup compiler to obtain a new method for secure two-party *program* computation with lower overhead than all existing constructions.

Keywords: Oblivious RAM, Cloud Computing, Multi-Server Model, Software Protection, Secure Computation.

1 Introduction

While circuit-based secure two-party computation has a rich literature starting with Yao and Goldreich-Micali-Wigderson [42, 13] and has garnered recent interest due to fast

* Expanded version of this paper appears on ePrint [27]. Protected by patent application [30]. Supported in part by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior National Business Center (DoI/NBC) contract number D11PC20199. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation therein. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsement, either expressed or implied, of IARPA, DoI/NBC, or the U.S. Government.

implementations (e.g. [20], [24], etc.), modern algorithms are typically represented as Random Access Memory (RAM) programs rather than circuits, that contain multiple branches, recursion, while loops, etc. Unrolling these into circuits it often incredibly costly. The alternative was proposed by Ostrovsky and Shoup (in STOC 1997) [33]: to utilize Oblivious RAM machinery inside two party computation and “simulate” the client by two players. More specifically, Ostrovsky and Shoup [33] suggests simulating the CPU of an oblivious RAM machine using off-the-shelf secure computation to perform CPU execution steps with atomic instructions implemented by circuits (that are executed securely) to simulate a “virtual” client in the Oblivious RAM and rely on one of the players to implement encrypted memory of Oblivious RAM. The simulation of each CPU is done through circuit-based secure two-party computation, thus CPU size in the Oblivious RAM simulation must be minimized, as otherwise it impacts simulation of each step of the computation. Luckily, there are multiple Oblivious RAM solutions that require $O(1)$ CPU memory in the security parameter.

The Ostrovsky-Shoup compiler suffers from two drawbacks: (1) The best running time of Oblivious RAM simulation with $O(1)$ memory requires $O(\log^2 n / \log \log n)$ overhead for running programs of length n due to Kushilevitz, Lu, and Ostrovsky (in SODA 2012) [26] and (2) The most problematic part of this approach is that most Oblivious RAM simulations with small CPU size, starting with Goldreich and Ostrovsky require “Oblivious Sorting” that introduce a huge constant into Oblivious RAM simulation that essentially kills all practicality. In this paper we eliminate both drawbacks stated above.

At the heart of our construction is a model known as “Distributed (or two-Server) Oblivious RAM” where an Oblivious RAM client is allowed to interact with two or more **non-communicating** databases. This model is analogous to the multi-server Private Information Storage model introduced by Ostrovsky and Shoup [33]. The critical difference is that in [33], while using two servers, as a building block they use single-server Oblivious RAM solution as a building block. The principle difference in the current paper is to take a closer look that the Oblivious RAM technology itself and to show how two non-communicating servers can lead to significant efficiency improvements in the Oblivious RAM itself. Further, we argue that two-server model of the Oblivious RAM is very natural in multiple applications. For example, as already mentioned our Oblivious RAM construction critically uses the two servers in obtaining an efficient two-party computation (where two players naturally implement two servers) since we show how completely skip the expensive “oblivious sorting” step of Oblivious RAM. We then show how two players can act as two servers in distributed ORAM and sort directly on the pseudo-random keys. This contribution may also be of independent interest in the realm of practical Oblivious RAM, as well as theoretical constructions that use Oblivious RAM.

We highlight two practical and realistic scenarios in which our solution is especially important: as already mentioned, complicated programs with branching, loops, recursion, and multiple execution paths, and secondly, programs which access during run-time only a small number of bits from large public inputs. If the size of the random access program (which could be dramatically smaller than the corresponding circuit) and inputs are bounded by Λ and runs in time T , our solution is the first practical method with $O((T + \Lambda) \log(T + \Lambda))$ communication and computation complexity. We also remark that if we allow preprocessing on large inputs (such as on a large database, maps, graphs, networks, etc.), the online work can be reduced to $O((T + \varepsilon) \log(T + \Lambda))$, where ε bounds the size of online inputs (such as in a query). This precomputation

model was used in the work of Gordon et al. [19] to achieve *polylogarithmic* online overhead, which we improve to *logarithmic* overhead. We now remind the reader of the motivation behind Oblivious RAM.

The concept of outsourcing data storage or computation is wide-spread in practice. This raises the issue of what happens to the privacy of the data when the outsourcing service is only semi-trusted or untrusted. Encryption can be employed to protect the *content* of the data, but it is apparent that information might be revealed based on how the data is accessed. Simply put, encryption by itself alone does not entirely address the issue of data privacy at hand.

The sequence of reads and writes a client makes to the remotely stored data is known as the *access pattern*. Even if the content of the data is protected by encryption, the server storing the data can deduce information about the encrypted data just by observing and analyzing the access pattern. For instance, the server can correlate this pattern with public information about the client’s behavior, such as the purchase or sale of stock. Over time, the server may learn enough information to predict the behavior of the client or the underlying semantics of the data, thereby defeating the purpose of encrypting it in the first place.

A trivial solution would be for the client to access the entire stored database every single read or write. This clearly hides the access pattern, but the per-access overhead is linear in the size of stored data. The question remains:

Is it possible to hide the access pattern with less than linear overhead?

In the model where the client is a Random Access Machine (i.e. RAM model), Goldreich and Ostrovsky [11, 32, 31, 14] introduced the concept of hiding the access pattern in the context of software protection. A small protected CPU would run on a machine with large unprotected RAM. The goal was to obviously simulate access to RAM, so that the set of instructions ran by the CPU would be protected against an outsider monitoring the RAM. In this manner, an adversary observing the RAM would learn nothing about what instructions were executed except the total number of instructions. The work of Goldreich [11] featured two solutions using constant client memory: a “square-root” solution and a “recursive square-root” solution. The amortized time overhead of executing a program in the former scheme was $O(\sqrt{n})$, and $O(2^{\sqrt{\log n \log \log n}})$ in the latter. Ostrovsky [32, 31] then discovered what is known as the “hierarchical solution” which had amortized overhead $O(\min\{\log^3 n, \log^3 t\})$, where t is the running time. The subsequent work of Goldreich and Ostrovsky [14] contains the merged results of [32, 31, 11] and featured a simpler method of reshuffling. The work described a way of simulating oblivious RAM with $O(\log^3 n)$ amortized overhead per access for n data items, using constant client storage³ and $O(n \log n)$ server storage. A simple way to modify the solution to make it *worst-case* $O(\log^3 n)$ overhead per access was pointed out by Ostrovsky and Shoup in 1997 [33].

While the asymptotic behavior of $O(\log^3 n)$ overhead might seem efficient, this only gives a practical advantage over the trivial solution when $n > \log^3 n$ (without even considering the constants hidden in the O). A database of size $n = 2^{20}$ results in an overhead factor of roughly 8000, and such a large overhead would seem to cast oblivious RAM as outside the realm of practicality. Making oblivious RAM practical would be of great impact, as it can be applied to software protection and several other

³ We count storage as the number of records or data items stored in memory. We do not count small variables such as counters or loop iterators toward this amount as these typically are tiny compared to the size of a data item, nor the private-key encryption/decryption cost.

important problems such as cloud computing, preventing cache attacks, etc. as we discuss later.

A highly interesting and powerful application of Oblivious RAM is in the problem of efficient, secure two-party *program* computation. While there are many ways to model computation, such as with Turing Machines, (Boolean) Circuits, Branching Programs, or Random Access Machines, one representation might be more natural than another, depending on the program. Nearly all secure two-party computation protocols require the program to be specified as a *circuit* between the two parties. Due to a classic result by Pippenger and Fischer [37], any Turing machine running in time T can be transformed into a circuit with only $O(\log T)$ blowup, but it is not known in the RAM model of computation whether there exists such an efficient transformation to circuits. Therefore, even using the most efficient secure two-party protocols for circuits (e.g. IKOS [21] or IPS [22] protocols), there is no clear path on how to apply these to efficiently perform secure RAM computation. We consider the question:

How can one efficiently perform secure two-party computation in the natural RAM model?

RELATED WORK. Subsequent to Goldreich and Ostrovsky [32, 31, 11, 14], works on Oblivious RAM [40, 41, 36, 16, 17, 38, 18, 26, 39] looked at improving the concrete and asymptotic parameters of oblivious RAM. We give a full summary of these schemes in Section 2. The major practical bottleneck of all these works on Oblivious RAMs is a primitive called *oblivious sorting* that is being called upon as a sub-protocol. Although the methods for oblivious sorting have improved, it still remains as both the critical step and the primary stumbling block of all these schemes. Even if new methods for oblivious RAM are discovered, there is an *inherent limit* to how much these schemes can be improved. It was shown in the original work of Goldreich and Ostrovsky [14] that there is a lower bound for oblivious RAM in this model.

([14], Theorem 6): *To obliviously perform n queries using only $O(1)$ client memory, there is a lower bound of $O(\log n)$ amortized overhead per access.*

We mention several results that are similar to Oblivious RAM but work in slightly different models. The works of Ajtai [1] and Damgård et al. [10] show how to construct oblivious RAM with information-theoretic security with poly-logarithmic overhead in the restricted model where the Adversary can not read memory contents. That is, these results work in a model where an adversary only sees the sequence of accesses and not the data. The work of Boneh, Mazieres and Popa [7] suggests ways to improve the efficiency of the “square-root” solution [11, 14] when memory contents are divided into larger blocks.

Finally, the notion of *Private Information Storage* introduced by Ostrovsky and Shoup [33] allows for private storage and retrieval of data. The work was primarily concentrated in the information theoretic setting. This model differs from Oblivious RAM in the sense that, while the communication complexity of the scheme is sub-linear, the server performs a *linear* amount of work on the database. The work of Ostrovsky and Shoup [33] gives a multi-server solution to this problem in both the computational and the information-theoretic setting and introduces the Ostrovsky-Shoup compiler of transforming Oblivious RAM into secure two-party computation. Directly quoting from their STOC 1997 [33] paper (with citations cross-referenced):

Both databases keep shares of the state of the CPU, and additionally one of the databases also keeps the contents of the Oblivious RAM memory. The main reason why we can allow one of the constituent databases to keep both the “share” of the CPU and the Oblivious RAM memory and still show that the view of this constituent database is

computationally indistinguishable for all executions is that the Oblivious RAM memory component is kept in an encrypted (and tamper-resistant) form (see [14]), according to a distributed (between both databases) private-key stored in the CPU. For every step of the CPU computation, both databases execute secure two-party function evaluation of [42, 13] which can be implemented based on any one-way trapdoor permutation family (again communicating through the user) in order to both update their shares and output re-encrypted value stored in a tamper-resistant way in Oblivious RAM memory component.

The current work can be viewed as a generalization of the [33] model where servers must also perform sublinear work. The notion of single-server “PIR Writing” was subsequently formalized in Boneh, Kushilevitz, Ostrovsky and Skeith [6] where they provide a single-server solution. The case of amortized “PIR Writing” of multiple reads and writes was considered in [8].

Also along the lines of oblivious simulation of execution, the result of Pippenger and Fischer [37] shows that a single-tape Turing machine can be obliviously simulated by a two-tape Turing machine with logarithmic overhead.

With regard to secure computation for RAM programs, the implications of the Ostrovsky-Shoup compiler was explored in the work of Naor and Nissim [29] which shows how to convert RAM programs into so-called circuits with “lookup tables” (LUT). This transformation incurs a poly-logarithmic blowup, or more precisely, for a RAM running in time T using space S , there is a family of LUT circuits of size $T \cdot \text{polylog}(S)$ that performs the same computation. The work then describes a specific protocol that securely evaluates circuits with lookup tables. [29] also applies to the related model of securely computing branching programs.

The Ostrovsky-Shoup compiler was further explored in the work of Gordon et al. [19] in the case of amortized programs. Namely, consider a client that holds a small input x , and a server that holds a large database D , and the client wishes to repeatedly perform private queries $f(x, D)$. In this model, an expensive initialization (depending only on D) is first performed. Afterwards, if f can be computed in time T with space S with a RAM machine, then there is a secure two-party protocol computing f in time $O(T) \cdot \text{polylog}(S)$ with the client using $O(\log S)$ space and the server using $O(S \cdot \text{polylog}(S))$ space.

Efforts on making Oblivious RAM perform well in the worst case were initiated by Ostrovsky and Shoup [33] with follow-up works by Shi et. al [38] and Stefanov-Shi-Song [39]. These results independently discover a way to avoid oblivious sorting, though have worse asymptotics than our new scheme. Although these results asymptotically perform worse (such as having $O(\log^3 n)$ overhead), their focus was on reducing the worst-case overhead as well as minimizing actual constants and was effective for improving the running time for reasonably sized databases. Other works that considered worst-case overhead include Goodrich et. al [17] and Kushilevitz et. al [26].

OUR RESULTS. In this paper, we consider a model for oblivious RAM in which we can achieve far better parameters than existing single-server schemes. We mention that our model, like most existing schemes, focuses on computational rather than information-theoretic security, and we only make the mild assumption that one-way functions exist. Instead of having a single server store our data, similar to Ostrovsky-Shoup [33] we consider using multiple⁴ servers to store client’s data. These servers are assumed to not communicate or collude with each other, but only communicate with the client. The

⁴ In general, we can consider multiple servers. For our purposes, two servers suffice.

main difference compared with [33] is that [33], while having several servers implemented virtual “off-the-shelf” Oblivious RAM that has a single server. We, in contrast, examine what two servers can bring to the Oblivious RAM world, where we allow these servers to perform simple computations such as hashing and sorting.

Note that although this Oblivious RAM model differs from the original model of the server only having read/write, we mention that it is still applicable in many of the interesting applications, such as cloud computing, tamper-proof CPUs interacting with other CPUs, and our main application, secure two-party computation. From a theoretical point of view, this model has been used in the past to much success such as in the seminal works in the areas of multi-prover Interactive Proof Systems [4] and multi-server Private Information Retrieval [9]. As already mentioned, this model is also directly applicable to the Ostrovsky-Shoup compiler for the construction of secure two-party RAM computation protocols.

In our two-server model, we introduce a new approach for Oblivious RAM **that completely bypasses oblivious sorting**, which was the inhibiting factor of practicality in most previous schemes (we give a comparison in Section 2.3). To perform a sequence of n reads or writes, our solution achieves $O(\log n)$ amortized overhead per access, $O(n)$ storage for the servers, and constant client memory. This matches the lower bound in the single-server model [14], and thus *no* single-server solution that uses constant client memory can asymptotically outperform our solution.

Theorem 1 (Informal) *Assume one-way functions exist. Then there exists an Oblivious RAM simulation in the two-server model with $O(\log n)$ overhead.*

In the work of [18], the notion of *Stateless* Oblivious RAM was introduced due to the fact that many Oblivious RAM solutions using super-constant client memory also required the client to maintain state between queries. All previous schemes with *constant* client memory were stateless and our new construction follows this trend.

As mentioned above, this new Oblivious RAM protocol leads to a novel application to secure *RAM program* computation. We then show how to perform secure two-party RAM computation by adapting our multi-server Oblivious RAM solution to fit the Ostrovsky-Shoup compiler [33]. This allows us (under cryptographic assumptions) to achieve the most efficient logarithmic communication complexity overhead for secure RAM computation as opposed to the *poly-logarithmic* overhead of all prior schemes [33, 29, 19, 28].

Theorem 2 (Informal) *Given any secure circuit computation protocol with $O(1)$ communication overhead, for any RAM program Π with upper bound T on its running time and Λ on its size (including $|\Pi|$) and length of inputs, there exists a two-party secure computation protocol for executing Π with $O((T + \Lambda) \cdot \log(T + \Lambda))$ communication and computation complexity. If large inputs are pre-processed and the online inputs sizes are bounded by ε , the online cost becomes $O((T + \varepsilon) \cdot \log(T + \Lambda))$.*

As an additional important remark, if players are willing to reveal to each other the actual running time on specific (private) inputs, we can replace in the above theorem T with the exact running on specific inputs, even though this may leak additional information. In some applications this information may be harmless, but in other applications with vastly different average-case/worst-case performance, this leads to a natural and interesting question on the tradeoff between efficiency and privacy.

1.1 Applications

SECURE TWO-PARTY AND MULTIPARTY COMPUTATION. In the case of MPC, we can apply oblivious RAM by letting the participants jointly simulate the client, and have the contents of the server be stored in a secret-shared manner. This application was originally described by Ostrovsky and Shoup [33] in the case of secure two-party computation. As described above, several other subsequent works [29, 10, 19] also investigated the application of Oblivious RAM to the area of secure computation. This can be beneficial in cases where the program we want to securely compute is more suitable to be modeled by a RAM program than a circuit. In particular, we demonstrate in this paper how our two-server ORAM construction can be applied to the case of secure two-party RAM program computation in the semi-honest model to obtain more efficient protocols for this purpose.

SOFTWARE PROTECTION. Original works of Goldreich [11] and Ostrovsky [32] envisioned protecting software using oblivious RAM. A small tamper-resistant CPU could be incorporated in a system with a large amount of unprotected RAM. A program could be run on this CPU by using oblivious RAM to access the large memory. Because this RAM could be monitored by an adversary, the benefit of oblivious RAM is that it hides the access pattern of the program that is running, thus revealing only the running time of the program to the adversary.

CLOUD COMPUTING. With the growing popularity of storing data remotely in the cloud, we want a way to do so privately when the data is sensitive. As mentioned before, simply encrypting all the data is insufficient, and by implementing oblivious RAM in the cloud, a client can privately store and access sensitive data on an untrusted server.

PREVENTING SIDE-CHANNEL ATTACKS. There are certain side-channel attacks that are based on measuring the RAM accesses that can be prevented by using oblivious RAM. For example, an adversary can mount a cache attack by observing the memory cache of a CPU. This poses a real threat as it can be used for cryptanalysis and has even been observed in practice in the work of Osvik-Shamir-Tromer [34].

PRIVATE DATA STRUCTURES. Rather than protecting an entire program, we can consider the middle ground of data structures. Data structures typically fit neatly into the RAM model, where each read or write is a sequence of accesses to memory. Performing these operations will leak information about the data, and we can use oblivious RAM to mitigate such issues. For example, commercial databases typically offer encryption to protect the data, but to protect the access pattern we can replace the data structures with oblivious ones.

2 Background

2.1 Model

We work in the RAM model, where there is a tiny machine that can run a program that performs a sequence of reads or writes to memory locations stored on a large memory. This machine, which we will refer to as the client, can be viewed as a stateful processor with a special data register v that can run a program Π . From a given state Σ of the client and the most recently read element x , $\Pi(\Sigma, x)$ acts as the next instruction function and outputs a read or write query and an updated state Σ' .

Because we wish to hide the type of access performed by the client, we unify both types of accesses into an operation known as a *query*. A sequence of n queries can be viewed as a list of (memory location, data) pairs $(v_1, x_1), \dots, (v_n, x_n)$, along with a sequence of operations op_1, \dots, op_n , where op_i is a READ or WRITE operation. In the case of READ operations, the corresponding x value is ignored. The sequence of queries, including both the memory location and the data, performed by a client is known as the *access pattern*.

In our model, we wish to obviously simulate the RAM machine with a client, which can be viewed as having limited storage, that has access to multiple servers with large storage that do not communicate with one another. However, the servers are untrusted and assumed to only be, in the best case, *semi-honest*, i.e. each server follows the protocol but attempts to learn additional information by reviewing the transcript of execution. For our model, we assume that the servers can do slightly more than just I/O, in that they can do computations locally, such as shuffle arrays, as well as perform hashing and basic arithmetic and comparison operations.

An oblivious RAM is *secure* if for any two access patterns in the ideal RAM, the corresponding views in the execution of those access patterns of any individual server are computationally indistinguishable. Another way of putting it is that the view of a server can be simulated in a way that is indistinguishable from the view of the server during a real execution.

We also briefly state the model of secure two-party RAM computation which we work in (see, e.g. [12], for a more in-depth treatment of general models of secure computation). Let $f(A, B)$ be a function that can be efficiently computed by a RAM machine, that is to say, there exists a program Π that a client can execute starting with A and B stored in the appropriate input memory locations and halting with the result $f(A, B)$ in the appropriate output memory location on the server. We usually denote the running time $T(n)$ and the space used $S(n)$ which depend on the size of the input n .

We use an ideal/real simulation-based definition of security and also work in the setting of semi-honest adversaries. There are two parties, Alice and Bob that receive inputs A and B respectively and they wish to compute $f(A, B)$. In the ideal world, there is an ideal functionality \mathcal{F}_f that on inputs A and B simply computes $f(A, B)$ and sends the output to Alice and Bob. In the real world, we can think of the Alice and Bob executing a protocol π_f that computes $f(A, B)$. Roughly speaking, we say that π_f *securely realizes* the functionality \mathcal{F}_f if there exists an efficient simulator \mathcal{S} playing the role of the corrupted party in the ideal world can produce an output that is computationally indistinguishable from the view of the corrupted party in the real world.

2.2 Tools

HASHING. In our scheme and in previous schemes, hashing is a central tool in storing the records. For our purposes, the hash functions used for hashing will be viewed as either a random function or a keyed pseudorandom function family F_k . Recall the standard hashing with buckets data structure: there is a table of m buckets, each of size b , and a hash function $h : \mathcal{V} \rightarrow \{1 \dots m\}$. A record (v, x) is stored in bucket $h(v)$.

CUCKOO HASHING. A variant of standard hashing known as Cuckoo Hashing was introduced by Pagh and Rodler [35]. In this variant, the hash table does not have buckets,

but now two hash functions h_1, h_2 are used. Each record (v, x) can only reside in one of two locations $h_1(v)$ or $h_2(v)$, and it is always inserted into $h_1(v)$. If there was a previous record stored in that location, the previous record is kicked out and sent to its other location, possibly resulting in a chain of kicks. If the chain grows too long or there is a cycle, new hash functions are chosen, and it was shown that this results in an amortized $O(1)$ insertion time. A version of cuckoo hashing with a stash was introduced by Kirsch et al. [23] where it was shown that the probability of having to reset drops exponentially in the size of the stash.

OBLIVIOUS SORTING. A key ingredient in most previous schemes is the notion of oblivious sorting. This is a sorting algorithm such that the sequence of comparisons it makes is independent of the data. For example, the schemes of Batchier [3] and Ajtai et al. [2] are based on sorting networks, and recently a randomized shell sort was introduced by Goodrich [15].

2.3 Comparison to Prior Work

We briefly overview the relevant key techniques used in previous schemes:

Scheme	Comp. Overhead	Client Storage	Server Storage	# of Servers	Dist. Prob. ⁵
[14]ORAM _{GO1}	$O(\sqrt{n} \log n)$	$O(1)$	$O(n + \sqrt{n})$	1	<i>negl</i>
[14]ORAM _{GO2}	$O(\log^4 n)$	$O(1)$	$O(n \log n)$	1	<i>negl</i>
[14]ORAM _{GO3}	$O(\log^3 n)$	$O(1)$	$O(n \log n)$	1	<i>negl</i>
[40]ORAM _{WS}	$O(\log^2 n)$	$O(\sqrt{n})$	$O(n \log n)$	1	<i>negl</i>
[41]ORAM _{WSC}	$O(\log n \log \log n)$	$O(\sqrt{n})$	$O(n)$	1	poly
[36]ORAM _{PR}	$O(\log^2 n)$	$O(1)$	$O(n)$	1	poly
[16]ORAM _{GM1}	$O(\log^2 n)$	$O(1)$	$O(n)$	1	<i>negl</i>
[16]ORAM _{GM2}	$O(\log n)$	$O(n^\nu)$	$O(n)$	1	<i>negl</i>
[18]ORAM _{GMOT}	$O(\log n)$	$O(n^\nu)$	$O(n)$	1	<i>negl</i>
[17]ORAM _{GMOT2}	$O(\log n)$	$O(n^\tau)$	$O(n)$	1	<i>negl</i>
[38]ORAM _{SCSL}	$O(\log^3 n)$	$O(1)$	$O(n \log n)$	1	<i>negl</i>
[39]ORAM _{SSS}	$O(\log n)$	$cn, c \ll 1$	$4n + o(n)$	1	<i>negl</i>
[26]ORAM _{KLO}	$O(\log^2 n / \log \log n)$	$O(1)$	$O(n)$	1	<i>negl</i>
Our Scheme	$O(\log n)$	$O(1)$	$O(n)$	2	<i>negl</i>

Table 1. Comparison of oblivious RAM schemes.

SQUARE ROOT SOLUTION. In the work of Goldreich [11] and subsequently Goldreich-Ostrovsky [14], a “square root” solution (which we label ORAM_{GO1}) was introduced for oblivious RAM. This solution was not hierarchical in nature, and instead had a permutation of the entire memory stored in a single array along with a cache of size \sqrt{n} which was scanned in its entirety during every query. After every \sqrt{n} queries, the entire array was obliviously sorted and a new permutation was chosen. This results in an amortized communication overhead of $O(\sqrt{n} \log n)$ per access.

HIERARCHICAL SOLUTION. In the work of Ostrovsky [32] and subsequently [14], a hierarchical solution was given for oblivious RAM. In this solution, the server holds a

⁵ Due to flaws in the way hash functions are used, the security of these schemes could be only polynomially secure. For further discussion on the security analysis of these schemes, see [16, 26, 19].

hierarchy of bucketed hash tables, growing geometrically in size. New records would be inserted at the smallest level, and as the levels fill up, they would be reshuffled down and re-hashed by using oblivious sorting. A query for v would scan bucket $h_i(v)$ in the hash table on level i . By using the oblivious sorting of Batcher [3], the scheme achieves an $O(\log^4 n)$ amortized query overhead (ORAM_{GO2}), and with AKS [2], an $O(\log^3 n)$ query overhead is achieved (ORAM_{GO3}).

BUCKET SORTING. In the work of Williams-Sion [40], the client was given $O(\sqrt{n})$ working memory instead of $O(1)$. By doing so, it was possible to achieve a more efficient oblivious sorting algorithm by sorting the data locally in chunks of size \sqrt{n} and then sending it back to the server. This resulted in a solution (ORAM_{WS}) with $O(\log^2 n)$ query overhead. This idea of using the client to sort was continued in the work of Williams et al. [41] in which a Bloom filter [5] was introduced to check whether or not an element was stored in a level before querying upon it. This solution (ORAM_{WSC}) was suggested to have $O(\log n \log \log n)$ overhead, but the a more careful analysis of [36] shows that this depends on the number of hash functions used in the Bloom filter.

CUCKOO HASHING. Pinkas and Reinman [36] suggested a solution in which cuckoo hashing is used instead of standard bucketed hashing. The oblivious sorting algorithm used the more practical one of [15]. This resulted in a scheme (ORAM_{PR}) that only used constant client memory, $O(n)$ server storage, and only $O(\log^2 n)$ query overhead where the constant was empirically shown to be as small as 150. The work of Goodrich and Mitzenmacher [16] also made use of cuckoo hashing, although the stashed variant of cuckoo hashing was used for their scheme (ORAM_{GM1}), which resulted in similar parameters. They also suggested a solution where the client has $O(n^\nu)$ memory (ORAM_{GM2}), in which case they are able to achieve $O(\log n)$ query overhead. A stateless version of this scheme is featured in [18] with similar asymptotic behavior. The best known overhead for schemes with constant client memory come from the work of Kushilevitz, Lu, and Ostrovsky [26] (full version appears on ePrint [25]), where they introduce a new balancing technique for their scheme (ORAM_{KLO}) to achieve an overhead of $O(\log^2 n / \log \log n)$.

WORST-CASE. Recent schemes considered also the worst-case overhead per client query. The first paper to address worst-case overhead was [33] which showed a poly-log worst case overhead. Works such as Goodrich et. al [17] (which we label ORAM_{GMO2}), Shi et. al [38] (ORAM_{SCSL}), Kushilevitz et. al [26], and Stefanov-Shi-Song [39] (ORAM_{SSS}) also featured schemes that provide worst-case guarantees. All these schemes additionally either bypass or amortize oblivious sorting, albeit in independent and different manners than our new construction in this paper.

3 Our Scheme

3.1 Overview

Our new scheme uses the hierarchical format of Ostrovsky [32]. The general principle behind protocols using this technique can be stated as: the data is encrypted (under semantically secure encryption) and stored in hierarchical levels that reshuffle and move into larger levels as they fill up. To keep track of the movement, for each level we logically divide different time periods into *epochs*, based on how many queries the

client has already performed. All parties involved are aware of a counter t that indicates the number of queries performed by the client.

In hierarchical schemes, the reshuffling process is the main bottleneck in efficiency, specifically the need to perform “oblivious sorting” several times. We identify the purposes that oblivious sorting serves during reshuffling and describe methods on how to replace oblivious sorting in our two-server model.

The first purpose of oblivious sorting is to separate real items from “dummy” items. Dummy items are records stored in the levels to help the client hide the fact that it may have already found what it was looking for prior to reaching that level. For example, if the client was searching for virtual memory location v , and it was found on level 3, the client still needs to “go through the motions” and search on the remaining levels to hide the fact that v had already been found. On all subsequent levels in this example, the client would search for “*dummy*” $\circ t$ instead of v .

The second purpose of oblivious sorting is to identify old records being merged with new records. New records are always inserted at the topmost level, and as the levels are reshuffled down, there is the possibility that an old record will run into a new one on some lower level. Because they both have the same virtual memory location v , a collision will occur. To resolve this, when records are being reshuffled, an oblivious sort is performed to place old records next to new ones so that the old records can be effectively erased (re-encrypted as a dummy record).

Finally, oblivious sorting is used to apply a pseudorandom permutation to the records as they are being reshuffled. A permutation is necessary to prevent the server from being able to track entries as they get reshuffled into lower levels.

The key ingredient to our new techniques is the idea of “tagging” the records and letting the two servers do most of the work for the client. A typical record looks like (v, x) where v is the index of the record (virtual memory location), and x is the data stored at that index. In most previous schemes, a hash function was applied to v to determine where the record would be stored in the data structure. Because the client cannot reveal v to the servers, and yet we wish for the servers to do most of the work, the client needs to apply tags to the records. Later, when the client needs to retrieve index location v , the client first computes the tag and then looks up the tag instead of v in the data structure located on the servers.

Note that this tagging must be performed carefully. We want the client to use only $O(1)$ working memory, so it cannot simply keep a list of all the tags it has generated in the past. Instead, the tags must be deterministic so that the client is able to re-create the tag at a future point in time when needed. However, if the tags depend only on v , a server can immediately identify when two encrypted records have the same index location v .

To resolve the apparent tension between these two requirements, we use a pseudo-random function (PRF) applied to v , the level it is stored on, as well as the period of time which it is stored at that level, known as the *epoch*. We describe this in greater detail in our construction. In the expanded version [27] we first present a warm-up construction to demonstrate the utility of tagging and using two servers. For a sequence of client queries of length n , this *insecure* strawman construction will have the servers storing $O(n)$ data, the client having $O(1)$ working memory, and the amortized overhead of queries being $O(\log n)$.

3.2 Full Construction

A recent result [26] points out that hash overflows leads to an adversary distinguishing different access patterns. Plain cuckoo hashing and the variant of cuckoo hashing with a constant stash [16] yield a polynomial chance of this event occurring. The work of Goodrich-Mitzenmacher [16] shows that cuckoo hashing with *logarithmic* size stash yields a superpolynomially small chance (in n) of overflow, under the assumption that the size of the table is $\Omega(\log^7 n)$. Thus, as a starting point, we use the level layout of [16], where smaller levels are standard hash tables with buckets and larger levels are cuckoo hash tables with a stash of size $O(\log n)$. Furthermore, we use the idea of “caching the stash” that was previously used in [26] that can be viewed as an alternative to a “shared stash” [18]. We emphasize that this is where the similarities end with existing schemes and that significant modifications must be diligently balanced to yield a scheme with our desired parameters. Before we begin describing our full construction, we take a quick glance at the balancing dynamics involved in choosing the right parameters for our scheme. Our goal is to achieve $O(\log n)$ amortized overhead per query, while maintaining that the hash tables do not overflow with all but negligible probability.

Recall that the hybrid construction in [16] uses standard hashing with buckets for lower levels, up until the point where a level contains $\log^7 n$ elements, where it switches to cuckoo hashing with a stash of size $\log n$. For the probability of overflow to be negligible for standard hashing, the buckets must be of size $\log n$. To perform a read query, a bucket is scanned at each of the smaller levels, and the entire stash is scanned along with 2 elements of the cuckoo hash table at the larger levels. This operation already incurs a total of $O(\log n \log \log n)$ reads for the small levels and $O(\log^2 n)$ for the larger levels. We now summarize the series of modifications that need to be made to the structure of the scheme:

Reduce Bucket Size. The standard hash tables will now use buckets of size $3 \log n / \log \log n$. This causes the total amount of reads for the small levels to drop down to $O(\log n)$. This produces a negative side effect: a bucket will now overflow with $\frac{1}{n^2}$ probability.

Standard Hash with Stash. We introduce a stash of size $\log n$ to the standard hash tables to hold the overflows from the now reduced bucket sizes. We prove in expanded version [27] that the probability of overflowing the stash is negligible. This produces a negative side effect: each stash must be read at the smaller levels, bringing us back to $O(\log n \log \log n)$ reads for the smaller levels.

Cache the Stash. [18, 16, 26] For both the smaller levels and larger levels, the stash of size $\log n$ will not be stored at that level, but the entire stash is instead re-inserted into the hierarchy. In fact, by choosing the top level to be of size $O(\log n)$, we can fit the entire stash into the top level. We show how this step is done during a reshuffle. Now, because there is no longer a stash at any level, the total amount read from all the levels combined will be $O(\log n)$. This will cause the levels to be reshuffled more often, but we show that it is at most by a constant factor.

We now give the full details of our scheme.

Let $c = 2 \log n$, where c is taken to be the size of the top level ($i = 1$). We split the top level in half so that each server holds half of the top level, and for subsequent levels, server $\mathcal{S}_{i \bmod 2}$ holds level i . Let ℓ_{cuckoo} be the level such that $c \cdot 2^{\ell_{cuckoo}-1}$ is $\Omega(\log^7 n)$, e.g. $7 \log \log n$. For levels $i = 2, \dots, \ell_{cuckoo} - 1$, level i will be a standard

hash table consisting of $c \cdot 2^{i-1}$ buckets, each of size $3 \log n / \log \log n$, along with a “mental”⁶ stash of size $\log n$. For levels $i = \ell_{cuckoo}, \dots, N$, level i will be a cuckoo hash table that can hold up to $c \cdot 2^{i-1}$ elements, which is of size $c \cdot 2^i$, along with a “mental” stash of size $\log n$.

The client keeps a local counter t of how many queries have been performed so far, as well as a counter s to indicate how many dummy stash elements were created. We describe how a query is performed in Figure 1. To reshuffle levels $1 \dots i$ into level $i+1$, suppose \mathcal{S}_b holds level $i+1$ and let $\mathcal{S}_a = \mathcal{S}_{1-b}$ be the other server. The steps in Figure 2 are performed.

1. The client allocates temporary storage m , large enough to hold a single record, initialized to a dummy value “*dummy*”.
2. Read each entry of the entire top level from both servers one at a time. If v is found as some entry (v, x) then store x in m .
3. For small levels $i = 2 \dots \ell_{cuckoo} - 1$, perform the following with the server holding level i :
 - (a) If v has not already been found, compute the tag for v at this level as $z = F_s(i, e_i, v)$. Else, set $z = F_s(i, e_i, \text{“dummy”} \circ t)$.
 - (b) Fetch into local memory the bucket corresponding to $h(z)$ one element at a time, i.e. fetch (v_j, x_j) for $j = 1, \dots, 3 \log n / \log \log n$ from bucket $h(z)$ one element at a time.
 - (c) If v is found in some record (v_i, x_i) , then replace v_i with “*dummy*” $\circ t$ and store x_i in m .
 - (d) Re-encrypt the fetched records and store them back to their original locations, releasing them from local client memory.
4. For large levels $i = \ell_{cuckoo} \dots N$, perform the following with the server holding level i :
 - (a) If v has not already been found, compute the tag for v at this level as $z = F_s(i, e_i, v)$. Else, set $z = F_s(i, e_i, \text{“dummy”} \circ t)$.
 - (b) Fetch into local memory the records (v_0, x_0) and (v_1, x_1) from locations $h_0(z)$ and $h_1(z)$.
 - (c) If v is found at one of these locations, i.e. $v = v_b$ for some $b = 0, 1$, then replace v_b with “*dummy*” $\circ t$ and store x_b in m .
 - (d) Re-encrypt the fetched records and store them back to their original locations, releasing them from local client memory.
5. In the case of a write query, here we overwrite $m = y$.
6. Read each entry of the entire top level one at a time, and re-encrypt each record with the following exception: If the record is of the form (v, x) , then overwrite it with (v, m) before re-encrypting it.
7. If (v, x) was not overwritten at the top level, write (v, m) in the first available empty spot (even if m is “*dummy*”), otherwise write a dummy value (“*dummy*” $\circ t$, “*dummy*”).
8. The client increments the local query counter t . If t is a multiple of $c/2$, then a reshuffle step is performed as described below.

Fig. 1. Main Construction: Query

⁶ There will be no physical stash at this level, but during reshuffles a temporary stash is created for the purpose of hashing which will subsequently be re-inserted back to the top level.

1. \mathcal{S}_a allocates a temporary array and inserts every (encrypted) record it holds between levels 1 and i . \mathcal{S}_a applies a random permutation to this temporary array and sends its contents one by one to the client.
2. The client re-encrypts each record and sends it to \mathcal{S}_b . In this step, both empty and dummy records are treated as real records.
3. \mathcal{S}_b allocates a temporary array and inserts every record it holds between levels 1 and i as well as the records it received from the client in the previous step. \mathcal{S}_b applies a random permutation to this temporary array and sends its contents one by one to the client.
4. The client re-encrypts each record and sends it to \mathcal{S}_a , announcing that it is empty if the record is empty, and tagging remaining records (v, x) with the output of the PRF $F_s(i + 1, e_{i+1}, v)$, where e_{i+1} is the *new* epoch of level $i + 1$. Note that v may be a virtual memory address, a dummy value, or a stash dummy value. In this step, dummy records are treated as real records and we are only concerned with eliminating empty records.
5. \mathcal{S}_a now holds $c \cdot 2^{i-1}$ tagged records. It allocates a temporary hash table (standard or cuckoo, depending on the level), with a stash of size $\log n$ and it uses the hash functions corresponding to level $i + 1$ and epoch e_{i+1} to hash these records into this temporary table. If the insertion fails, new hash functions are selected (we will show this happens with negligible probability). \mathcal{S}_a then informs the client the number of elements inside the stash, σ , then sends both the temporary table and the stash one record at a time to the client.
6. As the client receives records from \mathcal{S}_a one at a time, it re-encrypts each record and sends them to \mathcal{S}_b without modifying the contents except:
 - (a) The first σ empty records in the table the client receives from \mathcal{S}_a are encrypted as $(\text{"stashdummy"} \circ s, \text{"empty"})$, incrementing s each time. Note that a table is always more than half empty, and therefore we can always find σ empty slots.
 - (b) Subsequent empty records from the table are encrypted as $(\text{"empty"}, \text{"empty"})$.
 - (c) Every empty record in the stash is re-encrypted as $(\text{"stashdummy"} \circ s, \text{"empty"})$, incrementing s each time.
7. \mathcal{S}_b stores the table records in level $i + 1$ in the order in which they were received, and stores the stash records at the top level.

Fig. 2. Main Construction: Reshuffle

3.3 Analysis of Main Construction

Theorem 1 *For a sequence of n queries, the main construction uses $O(n)$ memory for each server, $O(1)$ working memory for the client, and $O(\log n)$ amortized overhead for queries.*

Proof. Computing the sizes of the levels, level 1 is of size $c = 2 \log n$, split between the servers, levels $i = 2, \dots, \ell_{\text{cuckoo}} - 1$ are of size $c \cdot 2^{i-1} \cdot 3 \log n / \log \log n$ each, giving a total of $O(\log^9 n)$ size, since $\ell_{\text{cuckoo}} = 7 \log \log n$. Levels $i = \ell_{\text{cuckoo}}, \dots, N$ are of size $c \cdot 2^i$ each, where $c \cdot 2^N = n$, hence there is a total of $O(n)$ size. Note that the additional elements added in by the stash dummy elements can be counted as follows: every $c/2$ steps, we insert another $\log n$ stash dummy records into the hierarchy. Therefore, after n steps, at most $2n \log n / c = n$ stash dummy records have been inserted, and we can simply accommodate this by adding one extra level at the bottom.

Clearly, the client uses constant working memory as it only transmits records one at a time.

When the client performs the read operation, it reads $2 \log n$ records from the top level, $3 \log n / \log \log n$ elements from each level $i = 2, \dots, \ell_{cuckoo} - 1$, and 2 elements from each level $i = \ell_{cuckoo}, \dots, N$. Since $\ell_{cuckoo} = 7 \log \log n$ and $N = \log n - \log \log n - 1$, this gives a total of roughly $25 \log n$ elements read.

Because we re-insert the stash (which is half the size of the top level), we need to reshuffle twice as often. Note that each reshuffle only moves an element in the level at most 3 times. We sketch the analysis of the amortized overhead:

- For levels $2, \dots, \ell_{cuckoo} - 1$, each level contains $c \cdot 2^{i-1} 3 \log n / \log \log n$ elements and needs to be reshuffled every $c \cdot 2^{i-1} / 2$ steps. This incurs an amortized overhead of:

$$3 \sum_{i=2}^{7 \log \log n - 1} \frac{c \cdot 2^{i-1} 3 \log n / \log \log n}{c \cdot 2^{i-2}} = O(\log n)$$

with a constant of roughly 125.

- For levels ℓ_{cuckoo}, \dots, N , each level contains $c \cdot 2^i$ elements and needs to be reshuffled every $c \cdot 2^{i-1} / 2$ steps. This incurs an amortized overhead of:

$$3 \sum_{i=7 \log \log n}^{\log n - \log \log n - 1} \frac{c \cdot 2^i}{c \cdot 2^{i-2}} = O(\log n)$$

with a constant of roughly 10.

Before we prove the security of our construction, we state a few important lemmas.

Lemma 1. *At all times during the execution of the scheme, any record of the form $(v, *)$ will appear at most once in the hierarchy unless $v = \text{“empty”}$.*

Proof. An index v must be either a virtual memory location, a dummy element, a stash dummy element, or empty. Virtual memory locations are only introduced into the hierarchy either from a read query that found v at a lower level and moved it to the top, or from a write query that did not find v in the hierarchy. A dummy element $\text{“dummy”} \circ t$ can only be introduced during query t , and it can be introduced at most once. Similarly, stash dummy elements can only be introduced once as s is incremented after every such entry.

Lemma 2. *The same v will not be queried upon twice between reshuffles at any level.*

Proof. Once v is queried upon at a level, i , either it is a $\text{“dummy”} \circ t$ value (in which case it will trivially never be queried again, as t is incremented at the next query), or it is some virtual memory location. In the latter case, v will be written to the top level after the query, and subsequent queries to v will find v before it reaches level i , and the only way v can reach a deeper level is if i is reshuffled.

Lemma 3. *Every level except the top will always be empty or half-full (a half-full standard hash contains a number of records equal to half the number of buckets) and this state depends only on t .*

Proof. The proof is straightforward and we refer the reader to the full version [27].

Lemma 4. *Any time a level i is reshuffled, its stash is included in the shuffle.*

Proof. We observe that the only way a level is shuffled is if all previous levels are shuffled as well and become empty. Because the stash of level i was stored in the hierarchy above level i , no elements of the stash will fall below level i unless caused by a reshuffle, in which case it will be shuffled with level i .

Theorem 2 *Under the assumption that one-way functions exists, the main construction is a secure two-server oblivious RAM.*

Proof. One-way functions allow private-key encryption and authentication. We use method of [31] to prevent tampering and thus must only show how to protect the access pattern.

We show how to simulate the view of a server’s access pattern during the execution of the protocol upon any sequence of queries q_1, \dots, q_n knowing only the length n . We begin by first making the observation that every record is encrypted and will be re-encrypted whenever it is accessed. By the semantic security of the encryption, we can assume that all these data contents are computationally indistinguishable from the encryption of any other contents. We also replace both the hash functions (which are modeled as PRFs) and the tagging PRF by random functions.

We first consider the view of each server during a reshuffle. If the server is playing the role of S_a , after its initial message out, it sees a random sequence of encrypted records (real or dummy) with tags, and announced empty records. By Lemma 1, all the hidden records will contain elements with unique v ’s, and hence their tags will also be unique with overwhelming probability. The tags came from a random function that had not been previously used, and so the tags that the server sees are independent from its view. Furthermore, because of Lemmas 3 and 4, the number of empty records revealed will be deterministic and will not reveal any additional information. Thus, we can simulate this view by calculating the number of pre-determined items of each type, and use encryptions of 0 for all of them and tagging the appropriate records with completely random tags.

If the server is playing the role of S_b during a reshuffle, it will receive a sequence of encrypted records which reveals no information. Next, after it shuffles these records and sends them out, it receives back another sequence of encrypted records which also reveals no information. This view can be trivially simulated.

Finally, we argue that the sequence of reads can also be simulated. By the above arguments, we see that what each server holds at level i is nearly independent of its view, except for the fact that the tags of the records stored at that level are consistent with the hash function used at that level. By Lemma 2, between two reshuffles, the sequence of queries made to level i will all be distinct, but they may arbitrarily intersect the elements contained in level i . However, because only a negligible fraction of hash functions do not agree with the records in level i (i.e. would cause an overflow), the distribution of the outputs of the hash function applied to any sequence of distinct queries is statistically close to uniform⁷. Thus, we can simulate the probes to level i between reshuffles by a random sequence of probes.

⁷ Note that this does not hold true for plain cuckoo hashing, where there is a noticeable difference between a uniform hash function and one that makes a consistent cuckoo hash table.

4 Application to Secure Two-Party RAM Computation

In this section, we describe how our multi-party Oblivious RAM simulation can be applied to the setting of secure two-party computation on RAM programs. The idea of using Oblivious RAM for the purpose of secure computation has been suggested in the literature [33, 29, 10, 19] and we outline the high-level idea of its use.

Consider the setting of two (semi-honest) parties, Alice and Bob, who wish to securely compute some function f (computed as some RAM program Π that runs in time $T = T(n)$ and uses $S = S(n)$ space) on their inputs A and B . Observe that in an Oblivious RAM, the view of the server can be simulated, so the idea is to let Alice or Bob play the role of the server (or in the case of our construction, two servers). However, in the case of Oblivious RAM, the privacy of the data is not protected from the client, so in order to securely run Π , we need to somehow simulate the client as well. In order to do so, we let the state of the client be *shared* between Alice and Bob so that neither party learns what is going on until the end of the computation when their outputs are revealed. In order to compute on this shared state, each *fixed instruction* of Π is encoded as a circuit. We emphasize that rather than unrolling the entire program into a circuit, which may be quite inefficient, we are only representing each atomic instruction as a circuit.

Because the joint state secure computation occurs at each step in the program, we want to minimize the amount of computation and communication overhead incurred by this step. In particular, in order for Alice and Bob to jointly compute Π and simulate the state of the client efficiently, the client state should be as small as possible. This means that even if an ORAM solution is efficient in terms of computation or communication overhead, we cannot use it if the footprint of the client is too large. In particular, works that require the memory of the client to be $O(\sqrt{n})$ (e.g. [40, 41]) or $O(n^\nu)$ (e.g. [16, 18]) will incur too much overhead per step of the program. The currently most efficient (single-server) ORAM protocol that is suitable for this purpose comes from the work of [26].

We point out that when modeling the client, we can either treat it as operating on bits or on “words”. By this we mean the client may need, for example, pointers of $O(\log S)$ bits so that it can index into memory. The notion of a client having constant memory can implicitly mean that we are operating on words and these can each hold sufficiently many bits to perform the necessary instructions. However, when simulating the client steps using a circuit, we need to operate on bits rather than words.

Because of this, the client state may in fact be larger than a constant number of bits despite having only a constant number of words. In order not to gain any additional overhead when performing the simulation of the client state, we need to use an efficient MPC that has only *constant* overhead. For example, the protocols of IKOS [21] or IPS [22] suit this situation.

By using our two-server ORAM solution in the Ostrovsky-Shoup compiler, we are able to achieve lower overhead for secure RAM computation than any known single-server ORAM solution. We have:

Theorem 3 *Suppose there exists a symmetric-key encryption scheme and a hash function modeled as a random function or an efficient PRF (e.g. [21, 19]). Suppose there exists a two-party secure circuit computation protocol with constant overhead (e.g. [21, 22]). Then to securely compute a RAM program that runs in $T(n)$ time with access to $S(n)$ space with the size of the program (including inputs) bounded by $\Lambda(n)$, there*

exists a two-party secure RAM computation protocol in the semi-honest model with $O(\log(T + \Lambda))$ multiplicative overhead in communication and computational complexity, and an additive one-time cost of $O(\Lambda \log(\Lambda))$ for setup (that can be amortized over multiple secure evaluations on small online inputs). If the client computes on bits instead of words, there is an additional implicit $O(\log S)$ multiplicative overhead.

Proof Sketch. We give a construction of such a scheme and argue that it is secure.

We follow the construction of the Ostrovsky-Shoup [33] compiler and let Alice and Bob hold inputs A and B respectively, and let Π be the program they wish to securely compute. Initialize the two-server ORAM as follows: let Alice play the role of one server and let Bob play the role of the other server. They jointly simulate the state of the client in our two-server ORAM protocol initialized to the secret sharing of the initial state. The parties then proceed by secret sharing A and B with each other. The two parties run the MPC protocol on the instructions that tells the client to obliviously insert (via ORAM) A and B into the locations inside the RAM where the program Π expects to read them as input. At the end of this process, Alice and Bob hold their respective encrypted server data as well as the shared state of the client.

Then, Alice and Bob begin to jointly execute the instructions of Π . Namely, they start with a shared state Σ and a shared value x and they perform the secure two-party computation on the circuit representing the step $\Pi(\Sigma, x)$ to receive a new shared state Σ' and a read or write operation op . The operation is converted into a sequence of oblivious instructions op'_1, \dots, op'_ℓ by running the MPC on the two-server ORAM protocol steps. When the operation involves reading or writing from the server Alice is holding, Bob sends Alice his share of that instruction and Alice reconstructs the instruction and executes it on her server before re-sharing the result. Similarly, Alice reveals her share to Bob when the operation involves reading or writing from his server. At the end of execution of Π , Alice and Bob recombine shares to retrieve the output.

We follow the (standard) proof technique of composition of simulation of CPU and simulation of Oblivious RAM in which we invoke the simulatability of both the underlying MPC and ORAM (see also [19]). To simulate the view of one party, say Alice, we begin by generating a uniformly random share of the initial state for her view. As her input and Bob's input are being stored on the servers obliviously, we simulate the intermediate state shares Σ as random shares as well. To simulate the instruction execution via $(\Sigma', op) \leftarrow \Pi(\Sigma, x)$, again we generate uniform random shares for the intermediate state as well as the values retrieved. In a real execution, the resulting operation op is then converted into a sequence of oblivious instructions op'_1, \dots, op'_ℓ , and by the simulatability of the underlying oblivious RAM, we can in fact simulate the sequence by replacing op with a dummy operation. The simulator runs the sequence of oblivious instructions induced by this dummy operation and writes the sequences of Alice's memory probes to the simulated view.

Finally, when the output is about to be reconstructed, the simulator (which knows the result via interaction with the ideal functionality) sets the revealed share to be $r \oplus f(A, B)$ where r is the random share of the data for Alice during this final step.

5 Conclusion and Open Problems

In this paper, we introduced a new multi-server model for oblivious RAM and constructed a two-server scheme in this model. The scheme is secure against honest-but-curious servers assuming one-way functions exist. The parameters of the scheme – $O(1)$

client memory, $O(n)$ server memory, and $O(\log n)$ overhead – match the lower bound of single-server oblivious RAM. The natural open problem to ask is whether or not the same lower bound holds, or if a better scheme can be constructed in this new model.

Our scheme was constructed under the assumption of the existence of one-way functions. We ask the open question of whether or not information-theoretic multi-server oblivious RAM can be constructed with similar parameters. One naive way of doing so would be to duplicate each server and use information-theoretic secret sharing between each server and its duplicate in order to replace encryption. The interesting question is to ask whether one can do so with fewer servers or perhaps better performance.

In the follow-up paper we have shown how to make garble RAM programs non-interactive with poly-logarithmic communication overhead [28]. In this paper, we showed how to make the overhead logarithmic. This improves existing constructions [33, 29, 19, 28]. Notice, however, that unlike the non-interactive solution of [28], the solution presented in this paper is highly interactive. The task of achieving logarithmic overhead for non-interactive secure execution of RAM programs remains an interesting open question.

References

- [1] Miklós Ajtai. Oblivious RAMs without cryptographic assumptions. In *STOC*, pages 181–190, 2010.
- [2] Miklós Ajtai, János Komlós, and Endre Szemerédi. An $O(n \log n)$ sorting network. In *STOC*, pages 1–9, 1983.
- [3] Kenneth E. Batchier. Sorting networks and their applications. In *AFIPS Spring Joint Computing Conference*, pages 307–314, 1968.
- [4] Michael Ben-Or, Shafi Goldwasser, Joe Kilian, and Avi Wigderson. Multi-prover interactive proofs: How to remove intractability assumptions. In *STOC*, pages 113–131, 1988.
- [5] Burton H. Bloom. Space/time trade-offs in hash coding with allowable errors. *Commun. ACM*, 13(7):422–426, 1970.
- [6] Dan Boneh, Eyal Kushilevitz, Rafail Ostrovsky, and William E. Skeith III. Public key encryption that allows PIR queries. In *CRYPTO*, pages 50–67, 2007.
- [7] Dan Boneh, David Mazieres, and Raluca Ada Popa. Remote oblivious storage: Making oblivious RAM practical. CSAIL Technical Report, MIT-CSAIL-TR-2011-018, 2011.
- [8] Nishanth Chandran, Rafail Ostrovsky, and William E. Skeith III. Public-key encryption with efficient amortized updates. In *SCN*, pages 17–35, 2010.
- [9] Benny Chor, Oded Goldreich, Eyal Kushilevitz, and Madhu Sudan. Private information retrieval. In *FOCS*, pages 41–50, 1995.
- [10] Ivan Damgård, Sigurd Meldgaard, and Jesper Buus Nielsen. Perfectly secure oblivious RAM without random oracles. In *TCC*, pages 144–163, 2011.
- [11] Oded Goldreich. Towards a theory of software protection and simulation by oblivious RAMs. In *STOC*, pages 182–194, 1987.
- [12] Oded Goldreich. *Foundations of Cryptography: Basic Tools*. Cambridge University Press, Cambridge, UK, 2001.
- [13] Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game or a completeness theorem for protocols with honest majority. In *STOC*, pages 218–229, 1987.
- [14] Oded Goldreich and Rafail Ostrovsky. Software protection and simulation on oblivious RAMs. *J. ACM*, 43(3):431–473, 1996.
- [15] Michael T. Goodrich. Randomized shellsort: A simple oblivious sorting algorithm. In *SODA*, pages 1262–1277, 2010.
- [16] Michael T. Goodrich and Michael Mitzenmacher. Privacy-preserving access of outsourced data via oblivious RAM simulation. In *ICALP*, pages 576–587, 2011.

- [17] Michael T. Goodrich, Michael Mitzenmacher, Olga Ohrimenko, and Roberto Tamassia. Oblivious RAM simulation with efficient worst-case access overhead. In *CCSW*, pages 95–100, 2011.
- [18] Michael T. Goodrich, Michael Mitzenmacher, Olga Ohrimenko, and Roberto Tamassia. Privacy-preserving group data access via stateless oblivious RAM simulation. In *SODA*, pages 157–167, 2012.
- [19] Dov Gordon, Jonathan Katz, Vladimir Kolesnikov, Tal Malkin, Mariana Raykova, and Yevgeniy Vahlis. Secure computation with sublinear amortized work. Cryptology ePrint Archive, Report 2011/482, 2011.
- [20] Yan Huang, David Evans, Jonathan Katz, and Lior Malka. Faster secure two-party computation using garbled circuits. In *USENIX Security Symposium*, 2011.
- [21] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Cryptography with constant computational overhead. In *STOC*, pages 433–442, 2008.
- [22] Yuval Ishai, Manoj Prabhakaran, and Amit Sahai. Founding cryptography on oblivious transfer - efficiently. In *CRYPTO*, pages 572–591, 2008.
- [23] Adam Kirsch, Michael Mitzenmacher, and Udi Wieder. More robust hashing: Cuckoo hashing with a stash. In *ESA*, pages 611–622, 2008.
- [24] Benjamin Kreuter, abhi shelat, and Chih-hao Shen. Towards billion-gate secure computation with malicious adversaries. Cryptology ePrint Archive, Report 2012/179, 2012.
- [25] Eyal Kushilevitz, Steve Lu, and Rafail Ostrovsky. On the (in)security of hash-based oblivious RAM and a new balancing scheme. Cryptology ePrint Archive, Report 2011/327, 2011.
- [26] Eyal Kushilevitz, Steve Lu, and Rafail Ostrovsky. On the (in)security of hash-based oblivious RAM and a new balancing scheme. In *SODA*, pages 143–156, 2012.
- [27] Steve Lu and Rafail Ostrovsky. Multi-server oblivious RAM. *IACR Cryptology ePrint Archive*, 2011:384, 2011.
- [28] Steve Lu and Rafail Ostrovsky. How to garble RAM programs. *IACR Cryptology ePrint Archive*, 2012:601, 2012.
- [29] Moni Naor and Kobbi Nissim. Communication preserving protocols for secure function evaluation. In *STOC*, pages 590–599, 2001.
- [30] Rafail Ostrovsky. Apparatus system and method to efficiently search and modify information stored on remote servers, while hiding the access pattern. U.S. Patent Application No. 12,768,617. April 27, 2010.
- [31] Rafail Ostrovsky. *Software Protection and Simulation On Oblivious RAMs*. PhD thesis, Massachusetts Institute of Technology, 1992.
- [32] Rafail Ostrovsky. Efficient computation on oblivious RAMs. In *STOC*, pages 514–523, 1990.
- [33] Rafail Ostrovsky and Victor Shoup. Private information storage (extended abstract). In *STOC*, pages 294–303, 1997.
- [34] Dag Arne Osvik, Adi Shamir, and Eran Tromer. Cache attacks and countermeasures: The case of AES. In *CT-RSA*, pages 1–20, 2006.
- [35] Rasmus Pagh and Flemming Friche Rodler. Cuckoo hashing. In *ESA*, pages 121–133, 2001.
- [36] Benny Pinkas and Tzachy Reinman. Oblivious RAM revisited. In *CRYPTO*, pages 502–519, 2010.
- [37] Nicholas Pippenger and Michael J. Fischer. Relations among complexity measures. *J. ACM*, 26(2):361–381, 1979.
- [38] Elaine Shi, T.-H. Hubert Chan, Emil Stefanov, and Mingfei Li. Oblivious RAM with $O((\log N)^3)$ worst-case cost. In *ASIACRYPT*, pages 197–214, 2011.
- [39] Emil Stefanov, Elaine Shi, and Dawn Song. Towards practical oblivious RAM. In *NDSS*, 2012.
- [40] Peter Williams and Radu Sion. Usable PIR. In *NDSS*, 2008.
- [41] Peter Williams, Radu Sion, and Bogdan Carbunar. Building castles out of mud: practical access pattern privacy and correctness on untrusted storage. In *ACM Conference on Computer and Communications Security*, pages 139–148, 2008.
- [42] Andrew Chi-Chih Yao. Protocols for secure computations (extended abstract). In *FOCS*, pages 160–164, 1982.