

Oblivious Network RAM and Leveraging Parallelism to Achieve Obliviousness

Dana Dachman-Soled^{1 *}, Chang Liu^{1 **}, Charalampos Papamanthou^{1 ***}, Elaine Shi^{2 †}, and Uzi Vishkin^{1 ‡}

¹ University of Maryland
danadach@ece.umd.edu, liuchang@cs.umd.edu, cpap@umd.edu,
vishkin@umiacs.umd.edu

² Cornell University
runting@gmail.com

Abstract. Oblivious RAM (ORAM) is a cryptographic primitive that allows a trusted CPU to securely access untrusted memory, such that the access patterns reveal nothing about sensitive data. ORAM is known to have broad applications in secure processor design and secure multi-party computation for big data. Unfortunately, due to a logarithmic lower bound by Goldreich and Ostrovsky (Journal of the ACM, '96), ORAM is bound to incur a moderate cost in practice. In particular, with the latest developments in ORAM constructions, we are quickly approaching this limit, and the room for performance improvement is small. In this paper, we consider new models of computation in which the cost of obliviousness can be fundamentally reduced in comparison with the standard ORAM model. We propose the Oblivious Network RAM model of computation, where a CPU communicates with multiple memory banks, such that the adversary observes only which bank the CPU is communicating with, but not the address offset within each memory bank. In other words, obliviousness within each bank comes for free—either because the architecture prevents a malicious party from observing the address accessed within a bank, or because another solution is used to obfuscate memory accesses within each bank—and hence we only need to obfuscate communication patterns between the CPU and the memory banks. We present new constructions for obliviously simulating general or parallel programs in the Network RAM model. We describe applications of our new model in secure processor design and in distributed storage applications with a network adversary.

* Work supported in part by NSF CAREER award #CNS-1453045 and by a Ralph E. Powe Junior Faculty Enhancement Award.

** Work supported in part by NSF awards #CNS-1314857, #CNS-1453634, #CNS-1518765, #CNS-1514261, and Google Faculty Research Awards.

*** Work supported in part by NSF award #CNS-1514261, by a Google Faculty Research Award and by the National Security Agency.

† Work supported in part by NSF awards #CNS-1314857, #CNS-1453634, #CNS-1518765, #CNS-1514261, Google Faculty Research Awards, and a Sloan Fellowship. This work was done in part while a subset of the authors were visiting the Simons Institute for the Theory of Computing, supported by the Simons Foundation and by the DIMACS/Simons Collaboration in Cryptography through NSF award #CNS-1523467.

‡ Work supported in part by NSF award #CNS-1161857.

1 Introduction

Oblivious RAM (ORAM), introduced by Goldreich and Ostrovsky [18, 19], allows a *trusted* CPU (or a trusted computational node) to obliviously access *untrusted* memory (or storage) during computation, such that an adversary cannot gain any sensitive information by observing the data access patterns. Although the community initially viewed ORAM mainly from a theoretical perspective, there has recently been an upsurge in research on both new efficient algorithms (c.f. [8, 13, 22, 36, 39, 43, 46]) and practical systems [9, 11, 12, 21, 30, 35, 37, 38, 44, 48] for ORAM. Still the most efficient ORAM implementations [10, 37, 39] require a relatively large bandwidth blowup, and part of this is inevitable in the standard ORAM model. Fundamentally, a well-known lower bound by Goldreich and Ostrovsky states that any ORAM scheme with constant CPU cache must incur at least $\Omega(\log N)$ blowup, where N is the number of memory words, in terms of bandwidth and runtime. To make ORAM techniques practical in real-life applications, we wish to further reduce its performance overhead. However, since latest ORAM schemes [39, 43] have practical performance approaching the limit of the Goldreich-Ostrovsky lower bound, the room for improvement is small in the standard ORAM model. In this paper, we investigate the following question:

In what alternative, realistic models of computation can we significantly lower the cost of oblivious data accesses?

Motivated by practical applications, we propose the Network RAM (NRAM) model of computation and correspondingly, Oblivious Network RAM (O-NRAM). In this new model, one or more CPUs interact with M memory banks during execution. Therefore, each memory reference includes a *bank identifier*, and an *offset* within the specified memory bank. We assume that an *adversary cannot observe the address offset within a memory bank, but can observe which memory bank the CPU is communicating with*. In other words, obliviousness within each bank “comes for free”. Under such a threat model, an Oblivious NRAM (O-NRAM) can be informally defined as an NRAM whose observable memory traces (consisting of the bank identifiers for each memory request) do not leak information about a program’s private inputs (beyond the length of the execution). In other words, in an O-NRAM, the sequence of bank identifiers accessed during a program’s execution must be provably obfuscated.

1.1 Practical Applications

Our NRAM models are motivated by two primary application domains:

Secure processor architecture. Today, secure processor architectures [1, 12, 30, 35, 40, 41] are designed assuming that the memory system is *passive* and untrusted. In particular, an adversary can observe both memory contents and memory addresses during program execution. To secure against such an adversary, the trusted CPU must both encrypt data written to memory, and obfuscate memory access patterns.

Our new O-NRAM model provides a realistic alternative that has been mentioned in the architecture community [30, 31] and was inspired by the Module Parallel Computer (MPC) model of Melhorn and Vishkin [32]. The idea is to introduce *trusted* decryption

logic on the memory DIMMs (for decrypting memory addresses). This way, the CPU can encrypt the memory addresses before transmitting them over the insecure memory bus. In contrast with traditional passive memory, we refer to this new type of memory technology as *active memory*. In a simple model where a CPU communicates with a single active memory bank, obliviousness is automatically guaranteed, since the adversary can observe only *encrypted* memory contents and addresses. However, when there are multiple such active memory banks, we must obfuscate which memory bank the CPU is communicating with.

Distributed storage with a network adversary. Consider a scenario where a client (or a compute node) stores private, encrypted data on multiple distributed storage servers. We consider a setting where all endpoints (including the client and the storage servers) are *trusted*, but the network is an *untrusted* intermediary. In practice, trust in a storage server can be bootstrapped through means of trusted hardware such as the Trusted Platform Module (TPM) or as IBM 4758; and network communication between endpoints can be encrypted using standard SSL. Trusted storage servers have also been built in the systems community [3]. On the other hand, the untrusted network intermediary can take different forms in practice, e.g., an untrusted network router or WiFi access point, untrusted peers in a peer-to-peer network (e.g., Bitcoin, TOR), or packet sniffers in the same LAN. Achieving oblivious data access against such a network adversary is precisely captured by our O-NRAM model.

1.2 Background: The PRAM Model

Two of our main results deal with the parallel-RAM (PRAM) model, which is a synchronous generalization of the RAM model to the parallel processing setting. The PRAM model allows for an unbounded number of parallel processors with a shared memory. Each processor may access any shared memory cell and read/write conflicts are handled in various ways depending on the type of PRAM considered:

- **Exclusive Read Exclusive Write (EREW) PRAM:** A memory cell can be accessed by at most one processor in each time step.
- **Concurrent Read Exclusive Write (CREW) PRAM:** A memory cell can be read by multiple processors in a single time step, but can be written to by at most one processor in each time step.
- **Concurrent Read Concurrent Write (CRCW) PRAM:** A memory cell can be read and written to by multiple processors in a single time step. Reads are assumed to complete prior to the writes of the same time step. Concurrent writes are resolved in one of the following ways: (1) *Common*—all concurrent writes must write the same value; (2) *Arbitrary*—an arbitrary write request is successful; (3) *Priority*—processor id determines which processor is successful.

1.3 Results and Contributions

We introduce the Oblivious Network RAM model, and conduct the first *systematic* study to understand the “cost of obliviousness” in this model. We consider running both *sequential* programs and *parallel* programs in this setting. We propose novel algorithms

Setting	RAM to O-NRAM blowup	<i>c.f.</i> Best known ORAM blowup
Sequential-to-sequential compiler		
$W = \text{small}$	$\widehat{O}(\log N)$	$O(\log^2 N / \log \log N)$ [25]
$W = \Omega(\log^2 N)$	bandwidth: $\widehat{O}(1)$ runtime: $\widehat{O}(\log N)$	bandwidth: $\widehat{O}(\log N)$ [43] runtime: $O(\log^2 N / \log \log N)$ [25]
$W = \Omega(N^\epsilon)$	$\widehat{O}(1)$	$\widehat{O}(\log N)$ [43]
Parallel-to-sequential compiler		
$\omega(M \log N)$ -parallel	$O(1)$	Same as standard ORAM
Parallel-to-parallel compiler		
$M^{1+\delta}$ -parallel for any const $\delta > 0$	$O(\log^* N)$	best known: poly log N [7] lower bound: $\Omega(\log N)$

Table 1. A systematic study of “cost of obliviousness” in the Network ORAM model. W denotes the memory word size in # bits, N denotes the total number of memory words, and M denotes the number of memory banks. For simplicity, this table assumes that $M = O(\sqrt{N})$, and each bank has $O(\sqrt{N})$ words. Like implicit in existing ORAM works [19, 25], small word size assumes at least $\log N$ bits per word—enough to store a virtual address of the word.

that exploit the “free obliviousness” within each bank, such that the obliviousness cost is significantly lower in comparison with the standard Oblivious (Parallel) RAMs. We give a summary of our results below.

First, observe that if there are only $O(1)$ number of memory banks, there is a trivial solution with $O(1)$ cost: just make one memory access (real or dummy) to each bank for each step of execution. On the other hand, if there are $\Omega(N)$ memory banks each of constant size (where N denotes the total number of memory words), then the problem approaches standard ORAM [18, 19] or OPRAM [7]. The intermediate parameters are therefore the most interesting. For simplicity, in this section, we mainly state our results for the most interesting case when the number of banks $M = O(\sqrt{N})$, and each bank can store up to $O(\sqrt{N})$ words. In Sections 3, 4 and 5, our results will be stated for more general parameter choices. We now state our results (see also Table 1 for an overview).

“Sequential-to-sequential” compiler. First, we show that any RAM program can be obliviously simulated on a Network RAM, consuming only $O(1)$ words of local CPU cache, with $\widehat{O}(\log N)$ blowup in both runtime and bandwidth, where—throughout the paper—when we say the complexity of our scheme is $\widehat{O}(f(N))$, we mean that for any choice of $h(N) = \omega(f(N))$, our scheme attains complexity $g(N) = O(h(N))$. Further, when the RAM program has $\Omega(\log^2 N)$ memory word size, it can be obliviously simulated on Network RAM with only $\widehat{O}(1)$ bandwidth blowup (assuming non-uniform memory word sizes as used by Stefanov et al. in [38]). In comparison, the best known (constant CPU cache) ORAM scheme has roughly $\widehat{O}(\log N)$ bandwidth blowup for $\Omega(\log^2 N)$ memory word size [43]. For smaller memory words, the best known ORAM scheme has $O(\log^2 / \log \log N)$ blowup in both runtime and bandwidth [25].

“Parallel-to-sequential” compiler. We demonstrate that parallelism can facilitate obliviousness, by showing that programs with a “sufficient degree of parallelism” – specifically, programs whose degree of parallelism $P = \omega(M \log N)$ – can be obliviously simulated in the Network RAM model with only $O(1)$ blowup in runtime and bandwidth. Here, we consider parallelism as a property of the program, but are not in fact executing the program on a parallel machine. The overhead stated above is for the sequential setting, i.e., considering that both NRAM and O-NRAM have single processor. Our compiler works when the underlying PRAM program is in the EREW, CREW, common CRCW or arbitrary CRCW model.

Beyond the low overhead discussed above, our compiled sequential O-NRAM has the additional benefit that it allows for an extremely simple prefetching algorithm. In recent work, Yu et al. [49] proposed a dynamic prefetching algorithm for ORAM, which greatly improved the practical performance of ORAM. We note that our parallel-to-sequential compiler achieves prefetching essentially for free: Since the underlying PRAM program will make many parallel memory accesses to each bank, and since the compiler knows these memory addresses ahead of time, these memory accesses can automatically be prefetched. We note that a similar observation was made by Vishkin [42], who suggested leveraging parallelism for performance improvement by using (compile-time) prefetching in serial or parallel systems.

“Parallel-to-parallel” compiler. Finally, we consider oblivious simulation in the parallel setting. We show that for any parallel program executing in t parallel steps with $P = M^{1+\delta}$ processors, we can obliviously simulate the program on a Network PRAM with $P' = O(P/\log^* P)$ processors, running in $O(t \log^* P)$ time, thereby achieving $O(\log^* P)$ blowup in parallel time and bandwidth, and optimal work. In comparison, the best known OPRAM scheme has poly $\log N$ blowup in parallel time and bandwidth. The compiler works when the underlying program is in the EREW, CREW, common CRCW or arbitrary CRCW PRAM model. The resulting compiled program is in the arbitrary CRCW PRAM model.

1.4 Technical Highlights

Our most interesting technique is for the parallel-to-parallel compiler. We achieve this through an intermediate stepping stone where we first construct a parallel-to-sequential compiler (which may be of independent interest).

At a high level, the idea is to assign each virtual address to a pseudorandom memory bank (and this assignment stays the same during the entire execution). Suppose that a program is sufficiently parallel such that it always makes memory requests in $P = \omega(M \log N)$ -sized batches. For now, assume that all memory requests within a batch operate on *distinct* virtual addresses – if not we can leverage a hash table to suppress duplicates, using an additional “scratch” bank as the CPU’s working memory. Then, clearly each memory bank will in expectation serve P/M requests for each batch. With a simple Chernoff bound, we can conclude that each memory bank will serve $O(P/M)$ requests for each batch, except with *negligible* probability. In a sequential setting, we can easily achieve $O(1)$ bandwidth and runtime blowup: for each batch of memory requests, the CPU will sequentially access each bank $O(P/M)$ number of times, padding with dummy accesses if necessary (see Section 4).

However, additional difficulties arise when we try to execute the above algorithm in parallel. In each step, there is a batch of P memory requests, one coming from each processor. However, each processor cannot perform its own memory request, since the adversary can observe which processor is talking to which memory bank and can detect duplicates (note this problem did not exist in the sequential case since there was only one processor). Instead, we wish to

1. hash the memory requests into buckets according to their corresponding banks while suppressing duplicates; and
2. pad the number of accesses to each bank to a worst-case maximum – as mentioned earlier, if we suppressed duplicate addresses, each bank has $O(P/M)$ requests with probability $1 - \text{negl}(N)$.

At this point, we can assign processors to the memory requests in a round-robin manner, such that which processor accesses which bank is “fixed”. Now, to achieve the above two tasks in $O(\log^* P)$ parallel time, we need to employ non-trivial parallel algorithms for “colored compaction” [4] and “static hashing” [5, 17], for the arbitrary CRCW PRAM model, while using a scratch bank as working memory (see Section 5).

1.5 Related Work

Oblivious RAM (ORAM) was first proposed in a seminal work by Goldreich and Ostrovsky [18, 19] where they laid a vision of employing an ORAM-capable secure processor to protect software against piracy. In their work, Goldreich and Ostrovsky showed both a poly-logarithmic upper-bound (commonly referred to as the hierarchical ORAM framework) and a logarithmic lower-bound for ORAM—both under constant CPU cache. Goldreich and Ostrovsky’s hierarchical construction was improved in several subsequent works [6, 20, 22, 25, 33, 45–47]. Recently, Shi *et al.* proposed a new, tree-based paradigm for constructing ORAMs [36], thus leading to several new constructions that are simple and practically efficient [8, 13, 39, 43]. Notably, Circuit ORAM [43] partially resolved the tightness of the Goldreich-Ostrovsky lower bound, by showing that certain stronger interpretations of their lower bound are indeed tight.

Theoretically, the best known ORAM scheme (with constant CPU cache) for small $O(\log N)$ -sized memory words³ is a construction by Kushilevitz *et al.* [25], achieving $O(\log^2 N / \log \log N)$ bandwidth and runtime blowup. Path ORAM (variant with $O(1)$ CPU cache [44]) and Circuit ORAM can achieve better bounds for bigger memory words. For example, Circuit ORAM achieves $O(\log N)\omega(1)$ bandwidth blowup for a word size of $\Omega(\log^2 N)$ bits; and for $O(\log N)\omega(1)$ runtime blowup for a memory word size of N^ϵ bits where $0 < \epsilon < 1$ is any constant within the specified range.

ORAMs with larger CPU cache sizes (caching up to N^α words for any constant $0 < \alpha < 1$) have been suggested for cloud storage outsourcing applications [20, 38, 47]. In this setting, Goodrich and Mitzenmacher [20] first showed how to achieve $O(\log N)$ bandwidth and runtime blowup.

Other than secure processors and cloud outsourcing, ORAM is also noted as a key primitive for scaling secure multi-party computation to big data [23, 26, 43, 44]. In this

³ Every memory word must be large enough to store the logical memory address.

context, Wang *et al.* [43, 44] pointed out that the most relevant ORAM metric should be the circuit size rather than the traditionally considered bandwidth metrics. In the secure computation context, Lu and Ostrovsky [27] proposed a two-server ORAM scheme that achieves $O(\log N)$ runtime blowup. Similarly, ORAM can also be applied in other RAM-model cryptographic primitives such as (reusable) Garbled RAM [14–16, 28, 29].

Goodrich and Mitzenmacher [20] and Williams *et al.* [48] observed that computational tasks with inherent parallelism can be transformed into efficient, oblivious counterparts in the traditional ORAM setting—but our techniques apply to the NRAM model of computation. Finally, Oblivious RAM has been implemented in outsourced storage settings [37, 38, 45, 47, 48], on secure processors [9, 11, 12, 30, 31, 35], and atop secure multiparty computation [23, 43, 44].

Comparison of our parallel-to-parallel compiler with the work of [7]. Recently, Boyle, Chung and Pass [7] proposed Oblivious Parallel RAM, and presented a construction for oblivious simulation of PRAMs in the PRAM model. Our result is incomparable to their result: Our security model is weaker than theirs since we assume obliviousness within each memory bank comes for free; on the other hand, we obtain far better asymptotical and concrete performance. We next elaborate further on the differences in the results and techniques of the two works. [7] provide a compiler from the EREW, CREW and CRCW PRAM models to the EREW PRAM model. The security notion achieved by their compiler provides security against adversaries who see the entire access pattern, as in standard oblivious RAM. However, their compiled program incurs a poly log overhead in both the parallel time and total work. Our compiler is a compiler from the EREW, CREW, common CRCW and arbitrary CRCW PRAM models to the arbitrary CRCW PRAM model and the security notion we achieve is the weaker notion of oblivious network RAM, which protects against adversaries who see the bank being accessed, but not the offset within the bank. On the other hand, our compiled program incurs only a \log^* time overhead and its work is asymptotically the *same* as the underlying PRAM. Both our work and the work of [7] leverage previous results and techniques from the parallel computing literature. However, our techniques are primarily from the CRCW PRAM literature, while [7] use primarily techniques from the low-depth circuit literature, such as highly efficient sorting networks.

2 Definitions

2.1 Background: Random Access Machines (RAM)

We consider RAM programs to be interactive stateful systems $\langle \Pi, \text{state}, D \rangle$, consisting of a memory array D of N memory words, a CPU state denoted state , and a next instruction function Π which given the current CPU state and a value rdata read from memory, outputs the next instruction I and an updated CPU state denoted state' :

$$(\text{state}', I) \leftarrow \Pi(\text{state}, \text{rdata})$$

Each instruction I is of the form $I = (\text{op}, \dots)$, where op is called the op-code whose value is read, write, or stop. The initial CPU state is set to $(\text{start}, *, \text{state}_{\text{init}})$. Upon

input x , the RAM machine executes, computes output z and terminates. CPU state is reset to $(\text{start}, *, \text{state}_{\text{init}})$ when the computation on the current input terminates.

On input x , the execution of the RAM proceeds as follows. If $\text{state} = (\text{start}, *, \text{state}_{\text{init}})$, set $\text{state} := (\text{start}, x, \text{state}_{\text{init}})$, and $\text{rdata} := 0$. Now, repeat the $\text{doNext}()$ till termination, where $\text{doNext}()$ is defined as below:

$\text{doNext}()$

1. Compute $(I, \text{state}') = \Pi(\text{state}, \text{rdata})$. Set $\text{state} := \text{state}'$.
2. If $I = (\text{stop}, z)$ then terminate with output z .
3. If $I = (\text{write}, \text{vaddr}, \text{wdata})$ then set $D[\text{vaddr}] := \text{wdata}$.
4. If $I = (\text{read}, \text{vaddr}, \perp)$ then set $\text{rdata} := D[\text{vaddr}]$.

2.2 Network RAM (NRAM)

Network RAM. A Network RAM (NRAM) is the same as a regular RAM, except that memory is distributed across multiple banks, $\text{Bank}_1, \dots, \text{Bank}_M$. In an NRAM, every virtual address vaddr can be written in the format $\text{vaddr} := (m, \text{offset})$, where $m \in [M]$ is the bank identifier, and offset is the offset within the Bank_m .

Otherwise, the definition of NRAM is identical to the definition of RAM.

Probabilistic NRAM. Similar to the probabilistic RAM notion formalized by Goldreich and Ostrovsky [18, 19], we additionally define a *probabilistic NRAM*. A probabilistic NRAM is an NRAM whose CPU state is initialized with randomness ρ (that is unobservable to the adversary). If an NRAM is deterministic, we can simply assume that the CPU's initial randomness is fixed to $\rho := 0$. Therefore, a deterministic NRAM can be considered as a special case of a probabilistic NRAM.

Outcome of execution. Throughout the paper, we use the notation $\text{RAM}(x)$ or $\text{NRAM}(x)$ to denote the outcome of executing a RAM or NRAM on input x . Similarly, for a probabilistic NRAM, we use the notation $\text{NRAM}_\rho(x)$ to denote the outcome of executing on input x , when the CPU's initial randomness is ρ .

2.3 Oblivious Network RAM (O-NRAM)

Observable traces. To define Oblivious Network RAM, we need to first specify which part of the memory trace an adversary is allowed to observe during a program's execution. As mentioned earlier in the introduction, each memory bank has trusted logic for encrypting and decrypting the memory offset. The offset within a bank is transferred in encrypted format on the memory bus. Hence, *for each memory access* $\text{op} := \text{“read”}$ or $\text{op} := \text{“write”}$ *to virtual address* $\text{vaddr} := (m, \text{offset})$, *the adversary observes only the op-code* op *and the bank identifier* m , *but not the offset within the bank.*

Definition 1 (Observable traces). *For a probabilistic NRAM, we use the notation $\text{Tr}_\rho(\text{NRAM}, x)$ to denote its observable traces upon input x , and initial CPU randomness ρ :*

$$\text{Tr}_\rho(\text{NRAM}, x) := \{(\text{op}_1, m_1), (\text{op}_2, m_2), \dots, (\text{op}_T, m_T)\}$$

where T is the total execution time of the NRAM, and (op_i, m_i) is the op-code and memory bank identifier during step $i \in [T]$ of the execution.

We remark that one can consider a slight variant model where the opcodes $\{\text{op}_i\}_{i \in [T]}$ are also hidden from the adversary. Since to hide whether the operation is a read or write, one can simply perform one read and one write for each operation – the differences between these two models are insignificant for technical purposes. Therefore, in this paper, we consider the model whose observable traces are defined in Definition 1).

Oblivious Network RAM. Intuitively, an NRAM is said to be oblivious, if for any two inputs x_0 and x_1 resulting in the same execution time, their observable memory traces are computationally indistinguishable to an adversary.

For simplicity, we define obliviousness for NRAMs that run in deterministic T time regardless of the inputs and the CPU’s initial randomness. One can also think of T as the worst-case runtime, and that the program is always padded to the worst-case execution time. Oblivious NRAM can also be similarly defined when its runtime is randomized – however we omit the definition in this paper.

Definition 2 (Oblivious Network RAM). Consider an NRAM that runs in deterministic time $T = \text{poly}(\lambda)$. The NRAM is said to be computationally oblivious if no polynomial-time adversary \mathcal{A} can win the following security game with more than $\frac{1}{2} + \text{negl}(\lambda)$ probability. Similarly, the NRAM is said to be statistically oblivious if no adversary, even computationally unbounded ones, can win the following game with more than $\frac{1}{2} + \text{negl}(\lambda)$ probability.

- \mathcal{A} chooses two inputs x_0 and x_1 and submits them to a challenger.
- The challenger selects $\rho \in \{0, 1\}^\lambda$, and a random bit $b \in \{0, 1\}$. The challenger executes NRAM with initial randomness ρ and input x_b for exactly T steps, and gives the adversary $\text{Tr}_\rho(\text{NRAM}, x_b)$.
- \mathcal{A} outputs a guess b' of b , and wins the game if $b' = b$.

2.4 Notion of Simulation

Definition 3 (Simulation). We say that a deterministic RAM $:= \langle \Pi, \text{state}, D \rangle$ can be correctly simulated by another probabilistic NRAM $:= \langle \Pi', \text{state}', D' \rangle$ if for any input x for any initial CPU randomness ρ , $\text{RAM}(x) = \text{NRAM}_\rho(x)$. Moreover, if NRAM is oblivious, we say that NRAM is an oblivious simulation of RAM.

Below, we explain some subtleties regarding the model, and define the metrics for oblivious simulation.

Uniform vs. non-uniform memory word size. The O-NRAM simulation can either employ uniform memory word size or non-uniform memory word size. For example, the non-uniform word size model has been employed for recursion-based ORAMs in the literature [39, 43]. In particular, Stefanov *et al.* describe a parametrization trick where they use a smaller word size for position map levels of the recursion [39].

Metrics for simulation overhead. In the ORAM literature, several performance metrics have been considered. To avoid confusion, we now explicitly define two metrics

that we will adopt later. If an NRAM correctly simulates a RAM, we can quantify the overhead of the NRAM using the following metrics.

- **Runtime blowup.** If a RAM runs in time T , and its oblivious simulation runs in time T' , then the runtime blowup is defined to be T'/T . This notion is adopted by Goldreich and Ostrovsky in their original ORAM paper [18, 19].
- **Bandwidth blowup.** If a RAM transfers Y bits between the CPU and memory, and its oblivious simulation transfers Y' bits, then the bandwidth blowup is defined to be Y'/Y . Clearly, if the oblivious simulation is in a uniform word size model, then bandwidth blowup is equivalent to runtime blowup. However, bandwidth blowup may not be equal to runtime blowup in a non-uniform word size model.

In this paper, we consider oblivious simulation of RAMs in the NRAM model, and we focus on the case when the Oblivious NRAM has only $O(1)$ words of CPU cache.

3 Sequential Oblivious Simulation

We first consider oblivious (sequential) simulation of arbitrary RAMs in the NRAM model. The detailed proofs and algorithms for this section will appear in the full version. Most of the techniques used here (with the exception of how to obliviously store the position map in a separate bank) are inspired by the work on practical ORAM by Stefanov, Shi, and Song [38]. Here we describe how we have adjusted their techniques to fit the Network RAM model.

Let M denote the number of memory banks in our NRAM, where each bank has $O(N/M)$ capacity. For simplicity we first describe a simple Oblivious NRAM with $O(M)$ CPU private cache. In the beginning, every block $i \in [N]$ is assigned randomly to a bank $j \in [M]$. We also maintain locally (i) a position map that maps every block to each bank; (ii) a cache of M queues, which are initially empty. To read/write a block i :

- We retrieve its bank number x from the position map;
- We first look for block i in the local queue x . If it is not there, we send a dummy memory request to a random location. Otherwise we read and then remove block i from the memory bank x ;
- We pick a fresh random memory bank x' , and we push block i to the queue x' in the local cache.

To avoid the overflow of local queues, we use a background eviction technique from Stefanov, Shi, and Song [38], which ensures that the local queues do not grow too much, while still maintaining obliviousness. Although storing the position map takes $O(N \log M)$ bits of CPU cache, in the full version we describe a recursion technique [36, 38] that can reduce this storage to $O(1)$. Finally, to further reduce the space from $O(M)$ to $O(1)$, we can store the CPU cache in a separate memory bank. However, this is challenging, as indicated below.

Main challenge. Placing the cache in a special memory bank to achieve constant client storage might violate obliviousness, since different operations to the cache might have different memory traces. The key challenge is to design a special data structure to

store the cache inside the memory bank that ensures constant worst-case cost for each query—specifically, each queue in the eviction cache must support `pop`, `push`, `ReadAndRM` operations. Partly to design this special data structure, we modified the analysis of the deamortized Cuckoo hash table construction [2] to achieve negligible failure probability.

We defer details of our algorithms and techniques to the full version and next state our main theorem for our sequential-to-sequential compiler.

Theorem 1 (O-NRAM simulation of arbitrary RAM programs). *Any N -word RAM with a word size of $W = \Omega(\log^2 N)$ bits can be simulated by an Oblivious NRAM (with non-uniform word sizes) that consumes $O(W)$ bits of CPU cache, and with $O(M)$ memory banks each of $O(W \cdot (M + N/M + N^\delta))$ bits in size. Further, the oblivious NRAM simulation incurs $\widehat{O}(1)$ bandwidth blowup and $\widehat{O}(\log N)$ run-time blowup.*

4 Sequential Oblivious Simulation of Parallel Programs

We are eventually interested in parallel oblivious simulation of parallel programs (Section 5). As a stepping stone, we first consider sequential oblivious simulation of parallel programs. However, we emphasize that the results in this section can be of independent interest. In particular, one way to interpret these results is that “parallelism facilitates obliviousness”. Specifically, if a program exhibits a sufficient degree of parallelism, then this program can be made oblivious at only const overhead in the Network RAM model. The intuition for why this is so, is that instructions in each parallel time step can be executed in any order. Since subsequences of instructions can be executed in an arbitrary order during the simulation, many sequences of memory requests can be mapped to the same access pattern, and thus the request sequence is partially obfuscated.

4.1 Parallel RAM

To formally characterize what it means for a program to exhibit a sufficient degree of parallelism, we will formally define a P -parallel RAM. In this section, the reader should think of parallelism as a property of the program to be simulated – we actually characterize costs assuming both the non-oblivious and the oblivious programs are executed on a sequential machine (different from Section 5).

An P -parallel RAM machine is the same as a RAM machine, except the next instruction function outputs P instructions which can be executed in parallel.

Definition 4 (P -parallel RAM). *An P -Parallel RAM is a RAM which has a next instruction function $\Pi = \Pi_1, \dots, \Pi_P$ such that on input $(\text{state} = \text{state}_1 \parallel \dots \parallel \text{state}_P, \text{rdata} = \text{rdata}_1 \parallel \dots \parallel \text{rdata}_P)$, Π outputs P instructions (I_1, \dots, I_P) and P updated states $\text{state}'_1, \dots, \text{state}'_P$ such that for $p \in [P]$, $(I_p, \text{state}'_p) = \Pi_p(\text{state}_p, \text{rdata}_p)$. The instructions I_1, \dots, I_P satisfy one of the following:*

- All of I_1, \dots, I_P are set to (stop, z) (with the same z).
- All of I_1, \dots, I_P are either of the form. $(\text{read}, \text{vaddr}, \perp)$ or $(\text{write}, \text{vaddr}, \text{wdata})$.

Finally, the state state has size at most $O(P)$.

As a warmup exercise, we will first consider a special case where in each parallel step, the memory requests made by each processor in the underlying P -parallel RAM have distinct addresses—we refer to this model as a *restricted* PRAM. Later in Section 4.3, we will extend the result to the (arbitrary) CRCW PRAM case. Thus, our final compiler works when the underlying P -parallel RAM is in the EREW, CREW, common CRCW or arbitrary CRCW PRAM model.

Definition 5 (Restricted P -parallel RAM). For a P -parallel RAM denoted $\text{PRAM} := \langle D, \text{state}_1, \dots, \text{state}_P, \Pi_1, \dots, \Pi_P \rangle$, if every batch of instructions I_1, \dots, I_P have unique vaddr 's, we say that PRAM is a restricted P -parallel RAM.

4.2 Warmup: Restricted Parallel RAM to Oblivious NRAM

Our goal is to compile any P -parallel RAM (not necessarily restricted), into an efficient O-NRAM. As an intermediate step that facilitates presentation, we begin with a basic construction of O-NRAM from any *restricted*, parallel RAM. In the following section, we extend to a construction of O-NRAM from any parallel RAM (not necessarily restricted).

Let $\text{PRAM} := \langle D, \text{state}_1, \dots, \text{state}_P, \Pi_1, \dots, \Pi_P \rangle$ be a restricted P -Parallel RAM, for $P = \omega(M \log N)$. We now present an O-NRAM simulation of PRAM that requires $M + 1$ memory banks, each with $O(N/M + P)$ physical memory, where N is the database size.

Setup: Pseudorandomly assign memory words to banks. The setup phase takes the initial states of the PRAM, including the memory array D and the initial CPU state, and compiles them into the initial states of the Oblivious NRAM denoted ONRAM.

To do this, the setup algorithm chooses a secret key K , and sets $\text{ONRAM.state} = \text{PRAM.state} \parallel K$. Each memory bank of ONRAM will be initialized as a Cuckoo hash table. Each memory word in the PRAM's initial memory array D will be inserted into the bank numbered $(\text{PRF}_K(\text{vaddr}) \bmod M) + 1$, where vaddr is the virtual address of the word in PRAM. Note that the ONRAM's $(M + 1)$ -th memory bank is reserved as a scratch bank whose usage will become clear later.

`doNext()`: //We only consider read and write instructions here but not stop.

- 1: **For** $p := 1$ to P : $(\text{op}_p, \text{vaddr}_p, \text{wdata}_p) := \Pi_p(\text{state}_p, \text{rdata}_p)$
- 2: $(\text{rdata}_1, \text{rdata}_2, \dots, \text{rdata}_p) := \text{Access} \left(\left\{ \text{op}_p, \text{vaddr}_p, \text{wdata}_p \right\}_{p \in [P]} \right)$

Fig. 1. Oblivious simulation of each step of the restricted parallel RAM.

Simulating each step of the PRAM's execution. Each `doNext()` operation of the PRAM will be compiled into a sequence of instructions of the ONRAM. We now describe how this compilation works. Our presentation focuses on the case when the next instruction's op-codes are reads or writes. Wait or stop instructions are left unmodified during the compilation.

```

Access ( $\{\text{op}_p, \text{vaddr}_p, \text{wdata}_p\}_{p \in P}$ ):
1: for  $p = 1$  to  $P$  do
2:    $m \leftarrow (\text{PRF}_K(\text{vaddr}_p) \bmod M) + 1$ ;
3:    $\text{queue}[m] := \text{queue}[m].\text{push}(p, \text{op}_p, \text{vaddr}_p, \text{wdata}_p)$ ;
   // queue is stored in a separate scratch bank.
4: end for
5: for  $m = 1$  to  $M$  do
6:   if  $|\text{queue}[m]| > \text{max}$  then abort
7:   Pad  $\text{queue}[m]$  with dummy entries  $(\perp, \perp, \perp, \perp)$  so that its size is max;
8:   for  $i = 1$  to max do
9:      $(p, \text{op}, \text{vaddr}, \text{wdata}) := \text{queue}[m].\text{pop}()$ 
10:     $\text{rdata}_p := \text{ReadBank}(m, \text{vaddr})$ 
   // Each bank is a deamortized Cuckoo hash table.
11:    if  $\text{op} = \text{write}$  then  $\text{wdata} := \text{rdata}_p$ 
12:     $\text{WriteBank}(m, \text{vaddr}, \text{wdata})$ 
13:   end for
14: end for
15: return  $(\text{rdata}_1, \text{rdata}_2, \dots, \text{rdata}_P)$ 

```

Fig. 2. Obviously serving a batch of P memory requests with distinct virtual addresses.

As shown in Figure 1, for each doNext instruction, we first compute the batch of instructions I_1, \dots, I_P , by evaluating the P parallel next-instruction circuits Π_1, \dots, Π_P . This results in P parallel read or write memory operations. This batch of P memory operations (whose memory addresses are guaranteed to be distinct in the restricted parallel RAM model) will then be served using the subroutine Access.

We now elaborate on the Access subroutine. Each batch will have $P = \omega(M \log N)$ memory operations whose virtual addresses are distinct. Since each virtual address is randomly assigned to one of the M banks, in expectation, each bank will get $P/M = \omega(\log N)$ hits. Using a balls and bins analysis, we show that the number of hits for each bank is highly concentrated around the expectation. In fact, the probability of any constant factor, multiplicative deviation from the expectation is negligible in N . Therefore, we choose $\text{max} := 2(P/M)$ for each bank, and make precisely max number of accesses to each memory bank. Specifically, the Access algorithm first scans through the batch of $P = \omega(M \log N)$ memory operations, and assigns them to M queues, where the m -th queue stores requests assigned to the m -th memory bank. Then, the Access algorithm sequentially serves the requests to memory banks $1, 2, \dots, M$, padding the number of accesses to each bank to max. This way, the access patterns to the banks are guaranteed to be oblivious.

The description of Figure 2 makes use of M queues with a total size of $P = \omega(M \log N)$ words. It is not hard to see that these queues can be stored in an additional scratch bank of size $O(P)$, incurring only constant number of accesses to the scratch bank per queue operation. Further, in Figure 2, the time at which the queues are accessed, and the number of times they are accessed are not dependent on input data

(notice that Line 7 can be done by linearly scanning through each queue, incurring a \max cost each queue).

Cost analysis. Since $\max = 2(P/M)$, in Figure 2 (see Theorem 2), it is not hard to see each batch of $P = \omega(M \log N)$ memory operations will incur $\Theta(P)$ accesses to data banks in total, and $\Theta(P)$ accesses to the scratch bank. Therefore, the ONRAM incurs only a constant factor more total work and bandwidth than the underlying PRAM.

Theorem 2. *Let PRF be a family of pseudorandom functions, and PRAM be a restricted P -Parallel RAM for $P = \omega(M \log N)$. Let $\max := 2(P/M)$. Then, the construction described above is an oblivious simulation of PRAM using M banks each of $O(N/M + P)$ words in size. Moreover, the oblivious simulation performs total work that is constant factor larger than that of the underlying PRAM.*

Proof. Assuming the execution never aborts (Line 6 in Figure 2), then Theorem 2 follows immediately, since the access pattern is deterministic and independent of the inputs. Therefore, it suffices to show that the abort happens with negligible probability on Line 6. This is shown in the following lemma.

Lemma 1. *Let $\max := 2(P/M)$. For any PRAM and any input x , abort on Line 6 of Figure 2 occurs only with negligible probability (over choice of the PRF).*

Proof. We first replace PRF with a truly random function f . Note that if we can prove the lemma for a truly random function, then the same should hold for PRF, since otherwise we obtain an adversary breaking pseudorandomness.

We argue that the probability that abort occurs on Line 6 of Figure 2 in a particular step i of the execution is negligible. By taking a union bound over the (polynomial number of) steps of the execution, the lemma follows.

To upper bound the probability of abort in some step i , consider a thought experiment where we change the order of sampling the random variables: We run PRAM(x) to precompute all the PRAM's instructions up to and including the i -th step of the execution (independently of f), obtaining P distinct virtual addresses, and only then choose the outputs of the random function f on the fly. That is, when each virtual memory address vaddr_p in step i is serviced, we choose $m := f(\text{vaddr}_p)$ uniformly and independently at random. Thus, in step i of the execution, there are P distinct virtual addresses (i.e., balls) to be thrown into M memory banks (i.e., bins). Due to standard Chernoff bounds, for $P = \omega(M \log N)$, we have $P/M = \omega(\log N)$ and so the probability that there exists a bin whose load exceeds $2(P/M)$ is $N^{-\omega(1)}$, which is negligible in N .

We note that in order for the above argument to hold, the input x cannot be chosen adaptively, and must be fixed before the PRAM emulation begins.

4.3 Parallel RAM to Oblivious NRAM

Use a hash table to suppress duplicates. In Section 4.2, we describe how to obliviously simulate a restricted parallel-RAM in the NRAM model. We now generalize this result to support any P -parallel RAM, not necessarily restricted ones. The difference

```

Access ( $\{\text{op}_p, \text{vaddr}_p, \text{wdata}_p, p\}_{p \in P}$ ):
/* HTable, queue, and result data structures are stored in a scratch bank. For obliviousness,
operations on these data structures must be padded to the worst-case cost as we elaborate in
the text.*/
1: for  $p = 1$  to  $P$ : HTable[ $\text{op}_p, \text{vaddr}_p$ ] := ( $\text{wdata}_p, p$ ) // hash table insertions
2: for  $\{(\text{op}, \text{vaddr}), \text{wdata}, p\} \in \text{HTable}$  do // iterate through hash table
3:    $m := (\text{PRF}_K(\text{vaddr}) \bmod M) + 1$ 
4:   queue[ $m$ ] := queue[ $m$ ].push( $\text{op}, \text{vaddr}, \text{wdata}$ );
5: end for
6: for  $m = 1$  to  $M$  do
7:   if |queue[ $m$ ] > max then abort
8:   Pad queue[ $m$ ] with dummy entries ( $\perp, \perp, \perp$ ) so that its size is max;
9:   for  $i = 1$  to max do
10:    ( $\text{op}, \text{vaddr}, \text{wdata}, p$ ) := queue[ $m$ ].pop()
11:    result[ $p$ ] := ReadBank( $m, \text{vaddr}$ )
12:    if  $\text{op} = \text{write}$  then  $\text{wdata} := \text{rdata}$ 
13:    WriteBank( $m, \text{vaddr}, \text{wdata}$ )
14:   end for
15: end for
16: return (result[1], . . . , result[ $p$ ]) // hash table lookups

```

Fig. 3. Obviously serving a batch of P memory request, not necessarily with distinct virtual addresses.

is that for a generic P -parallel RAM, each batch of P memory operations generated by the next-instruction circuit need not have distinct virtual addresses. For simplicity, imagine that the entire batch of memory operations are reads. In the extreme case, if all $P = \omega(M \log N)$ operations correspond to the same virtual address residing in bank m , then the CPU should not read bank m as many as P number of times. To address this issue, we rely on an additional Cuckoo hash table [34] denoted HTable to suppress the duplicate requests (see Figure 3, and the doNext function is defined the same way as Section 4.2).

The HTable will be stored in the scratch bank. We can employ a standard Cuckoo hash table that need not be deamortized. As shown in Figure 3, we need to support hash table insertions, lookups, and moreover, we need to be able to iterate through the hash table. We now make a few remarks important for ensuring obliviousness. Line 1 of Figure 3 performs $P = \omega(M \log N)$ number of insertions into the Cuckoo hash table. Due to standard Cuckoo hash analysis, we know that these insertions will take $O(P)$ total time except with negligible probability. Therefore, to execute Line 1 obliviously, we simply need to pad with dummy insertions up to some $\text{max}' = c \cdot P$, for an appropriate constant c .

Next, we describe how to execute the loop at Line 2 obliviously. The total size of the Cuckoo hash table is $O(P)$. To iterate over the hash table, we simply make a linear scan through the hash table. Some entries will correspond to dummy elements. When iterating over these dummy elements, we simply perform dummy operations for the **for**

loop. Finally, observe that Line 16 performs lookups to the Cuckoo hash table, and each hash table lookup requires worst-case $O(1)$ accesses to the scratch bank.

Cost analysis. Since $\max = 2(P/M)$ (see Theorem 2), it is not hard to see each batch of $P = \omega(M \log N)$ memory operations will incur $O(P)$ accesses to data banks in total, and $O(P)$ accesses to the scratch bank. Note that this takes into account the fact that Line 1 and the for-loop starting at Line 2 are padded with dummy accesses. Therefore, the ONRAM incurs only a constant factor more total work and bandwidth than the underlying PRAM.

Theorem 3. *Let $\max = 2(P/M)$. Assume that PRF is a secure pseudorandom function, and PRAM is a P -Parallel RAM for $P = \omega(M \log N)$. Then, the above construction obviously simulates PRAM in the NRAM model, incurring only a constant factor blowup in total work and bandwidth consumption.*

Proof. (sketch.) Similar to the proof of Theorem 2, except that now we have the additional hash table. Note that obliviousness still holds, since, as discussed above, each batch of P memory requests requires $O(P)$ accesses to the scratch bank, and this can be padded with dummy accesses to ensure the number of scratch bank accesses remains the same in each execution.

5 Parallel Oblivious Simulation of Parallel Programs

In the previous section, we considered sequential oblivious simulation of programs that exhibit parallelism – there, we considered parallelism as being a property of the program which will actually be executed on a sequential machine. In this section we consider *parallel* and oblivious simulations of parallel programs. Here, the programs will actually be executed on a parallel machine, and we consider classical metrics such as parallel runtime and total work as in the parallel algorithms literature.

We introduce the *Network PRAM* model – informally, this is a Network RAM with parallel processing capability. Our goal in this section will be to compile a PRAM into an Oblivious Network PRAM (O-NPRAM), *a.k.a.*, the “parallel-to-parallel compiler”.

Our O-NPRAM is the Network RAM analog of the Oblivious Parallel RAM (OPRAM) model by Boyle *et al.* [7]. Goldreich and Ostrovsky’s logarithmic ORAM lower bound (in the sequential execution model) directly implies the following lower bound for standard OPRAM [7]: Let PRAM be an arbitrary PRAM with P processors running in parallel time t . Then, any P -parallel OPRAM simulating PRAM must incur $\Omega(t \log N)$ parallel time. Clearly, OPRAM would also work in our Network RAM model albeit not the most efficient, since it is not exploiting the fact that the addresses in each bank are inherently oblivious. In this section, we show how to perform oblivious parallel simulation of “sufficiently parallel” programs in the Network RAM model, incurring only $O(\log^* N)$ blowup in parallel runtime, and achieving optimal total work. Our techniques make use of fascinating results in the parallel algorithms literature [4, 5, 24].

5.1 Network PRAM (NPRAM) Definitions

Similar to our NRAM definition, an NPRAM is much the same as a standard PRAM, except that 1) memory is distributed across multiple banks, $\text{Bank}_1, \dots, \text{Bank}_M$; and 2)

every virtual address $vaddr$ can be written in the format $vaddr := (m, offset)$, where m is the bank identifier, and $offset$ is the offset within the $Bank_m$. We use the notation P -parallel NPRAM to denote an NPRAM with P parallel processors, each with $O(1)$ words of cache. If processors are initialized with secret randomness unobservable to the adversary, we call this a probabilistic NPRAM.

Observable traces. In the NPRAM model, we assume that an adversary can observe the following parts of the memory trace: 1) which processor is making the request; 2) whether this is a read or write request; and 3) which bank the request is going to. The adversary is unable to observe the offset within a memory bank.

Definition 6 (Observable traces for NPRAM). For a probabilistic P -parallel NPRAM, we use $Tr_\rho(\text{NPRAM}, x)$ to denote its observable traces upon input x , and initial CPU randomness ρ (collective randomness over all processors):

$$Tr_\rho(\text{NPRAM}, x) := [((op_1^1, m_1^1), \dots, (op_1^P, m_1^P)), \dots, ((op_T^1, m_T^1), \dots, (op_T^P, m_T^P))]$$

where T is the total parallel execution time of the NPRAM, and $\{(op_i^1, m_i^1), \dots, (op_i^P, m_i^P)\}$ is of the op -codes and memory bank identifiers for each processor during parallel step $i \in [T]$ of the execution.

Based on the above notion of observable memory trace, an Oblivious NPRAM can be defined in a similar manner as the notion of O-NRAM (Definition 2).

Metrics. We consider classical metrics adopted in the vast literature on parallel algorithms, namely, the parallel runtime and the total work. In particular, to characterize the oblivious simulation overhead, we will consider

- **Parallel runtime blowup.** The blowup of the parallel runtime comparing the O-NPRAM and the NPRAM.
- **Total work blowup.** The blowup of the total work comparing the O-NPRAM and the NPRAM. If the total work blowup is $O(1)$, we say that the O-NPRAM achieves *optimal* total work.

5.2 Construction of Oblivious Network PRAM

Preliminary: colored compaction. The colored compaction problem [4] is the following:

Given n objects of m different colors, initially placed in a single source array, move the objects to m different destination arrays, one for each color. In this paper, we assume that the *space for the m destination arrays are preallocated*. We use the notation d_i to denote the number of objects colored i for $i \in [m]$.

Lemma 2 (Log*-time parallel algorithm for colored compaction [4]). *There is a constant $\epsilon > 0$ such that for all given $n, m, \tau, d_1, \dots, d_m \in \mathbb{N}$, with $m = O(n^{1-\delta})$ for arbitrary fixed $\delta > 0$, and $\tau \geq \log^* n$, there exists a parallel algorithm (in the arbitrary CRCW PRAM model) for the colored compaction problem (assuming preallocated destination arrays) with n objects, m colors, and d_1, \dots, d_m number of objects for each color; executing in $O(\tau)$ time on $\lceil n/\tau \rceil$ processors, consuming $O(n + \sum_{i=1}^m d_i)$ space, and succeeding with probability at least $1 - 2^{-n^\epsilon}$.*

```

parAccess ( $\{\text{op}_p, \text{vaddr}_p, \text{wdata}_p\}_{p \in P}$ ):
/* All steps can be executed in  $O(\log^* P)$  time with  $P' = O(P/\log^* P)$  processors with all
but negligible probability.*/
1: Using the scratch bank as memory, run the parallel hashing algorithm on the batch
of  $P = M^{1+\delta}$  memory requests to suppress duplicate addresses. Denote the re-
sulting set as  $S$ , and pad  $S$  with dummy requests to the maximum length  $P$ .
2: In parallel, assign colors to each memory request in the array  $S$ . For each
real memory access  $\{\text{op}, \text{vaddr}, \text{wdata}\}$ , its color is defined as  $(\text{PRF}_K(\text{vaddr})
\text{ mod } M) + 1$ . Each dummy memory access is assigned a random color. It is not
hard to see that each color has no more than  $\text{max} := 2(P/M)$  requests except
with negligible probability.
3: Using the scratch bank as memory, run the parallel colored compaction algorithm
to assign the array  $S$  to  $M$  preallocated queues each of size  $\text{max}$  (residing in the
scratch bank).
4: Now, each queue  $i \in [M]$  contains  $\text{max}$  number of requests intended for bank  $i$  –
some real, some dummy. Serve all memory requests in the  $M$  queues in parallel.
Each processor  $i \in [P']$  is assigned the  $k$ -th memory request iff  $(k \text{ mod } P') = i$ .
Dummy requests incur accesses to the corresponding banks as well.
For each request coming from processor  $p$ , the result of the fetch is stored in an
array  $\text{result}[p]$  in the scratch bank.

```

Fig. 4. Obviously serving a batch of P memory requests using $P' := O(P/\log^* P)$ processors in $O(\log^* P)$ time. In Steps 1, 2, and 3, each processor will make exactly one access to the scratch bank in each parallel execution step – even if the processor is idle in this step, it makes a dummy access to the scratch bank. Steps 1 through 3 are always padded to the worst-case parallel runtime.

Preliminary: parallel static hashing. We will also rely on a parallel, static hashing algorithm [5, 24], by Bast and Hagerup. The static parallel hashing problem takes n elements (possibly with duplicates), and in parallel creates a hash table of size $O(n)$ of these elements, such that later each element can be visited in $O(1)$ time. In our setting, we rely on the parallel hashing to suppress duplicate memory requests. Bast and Hagerup show the following lemma:

Lemma 3 (Log^{*}-time parallel static hashing [5, 24]). *There is a constant $\epsilon > 0$ such that for all $\tau \geq \log^* n$, there is a parallel, static hashing algorithm (in the arbitrary CRCW PRAM model), such that hashing n elements (which need not be distinct) can be done in $O(\tau)$ parallel time, with $O(n/\tau)$ processors and $O(n)$ space, succeeding with $1 - 2^{-(\log n)^\tau / \log^* n} - 2^{-n^\epsilon}$ probability.*

Construction. We now present a construction that allows us to compile a P -parallel PRAM, where $P = M^{1+\delta}$ for any constant $\delta > 0$, into a $O(P/\log^* P)$ -parallel Oblivious NPRAM. The resulting NPRAM has $O(\log^* P)$ blowup in parallel runtime, and is optimal in total amount of work.

In the original P -parallel PRAM, each of the P processors does constant amount of work in each step. In the oblivious simulation, this can trivially be simulated in

$O(\log^* P)$ time with $O(P/\log^* P)$ processors. Therefore, clearly the key is how to obviously fetch a batch of P memory accesses in parallel with $O(P/\log^* P)$ processors, and $O(\log^* P)$ time. We describe such an algorithm in Figure 4. Using a scratch bank as working memory, we first call the parallel hashing algorithm to suppress duplicate memory requests. Next, we call the parallel colored compaction algorithm to assign memory request to their respective queues – depending on the destination memory bank. Finally, we make these memory accesses, including dummy ones, in parallel.

Theorem 4. *Let PRF be a secure pseudorandom function, let $M = N^\epsilon$ for any constant $\epsilon > 0$. Let PRAM be a P -parallel RAM for $P = M^{1+\delta}$, for constant $\delta > 0$. Then, there exists an Oblivious NPRAM simulation of PRAM with the following properties:*

- *The Oblivious NPRAM consumes M banks each of which $O(N/M + P)$ words in size.*
- *If the underlying PRAM executes in t parallel steps, then the Oblivious NPRAM executes in $O(t \log^* P)$ parallel steps utilizing $O(P/\log^* P)$ processors. We also say that the NPRAM has $O(\log^* P)$ blowup in parallel runtime.*
- *The total work of the Oblivious NPRAM is asymptotically the same as the underlying PRAM.*

Proof. We note that our underlying PRAM can be in the EREW, CREW, common CRCW or arbitrary CRCW models. Our compiled oblivious NPRAM is in the arbitrary CRCW model.

We now prove security and costs separately.

Security proof. Observe that Steps 1, 2, and 3 in Figure 4 make accesses only to the scratch bank. We make sure that each processor will make exactly one access to the scratch bank in every parallel step – even if the processor is idle in this step, it makes a dummy access. Further, Steps 1 through 3 are also padded to the worst-case running time. Therefore, the observable memory traces of Steps 1 through 3 are perfectly simulatable without knowing secret inputs.

For Step 4 of the algorithm, since each of the M queues are of fixed length \max , and each element is assigned to each processor in a round-robin manner, the bank number each processor will access is clearly independent of any secret inputs, and can be perfectly simulated (recall that dummy request incur accesses to the corresponding banks as well).

Costs. First, due to Lemma 1, each of the M queues will get at most $2(P/M)$ memory requests with probability $1 - \text{negl}(N)$. This part of the argument is the same as Section 4. Now, observe that the parallel runtime for Steps 2 and 4 are clearly $O(\log^* P)$ with $O(P/\log^* P)$ processors. Based on Lemmas 3 and 2, Steps 1 and 3 can be executed with a worst-case time of $O(\log^* P)$ on $O(P/\log^* P)$ processors as well. We note that the conditions $M = N^\epsilon$ and $P = M^{1+\delta}$ ensure $\text{negl}(N)$ failure probability.

References

- [1] Intel SGX for dummies (intel SGX design objectives). <https://software.intel.com/en-us/blogs/2013/09/26/protecting-application-secrets-with-intel-sgx>.

- [2] Yuriy Arbitman, Moni Naor, and Gil Segev. De-amortized cuckoo hashing: Provable worst-case performance and experimental results. In *Automata, Languages and Programming, 36th International Colloquium, ICALP 2009, Rhodes, Greece, July 5-12, 2009, Proceedings, Part I*, pages 107–118, 2009.
- [3] Sumeet Bajaj and Radu Sion. Trustddb: A trusted hardware-based database with privacy and data confidentiality. *IEEE Trans. Knowl. Data Eng.*, 26(3):752–765, 2014.
- [4] Hannah Bast and Torben Hagerup. Fast parallel space allocation, estimation, and integer sorting. *Inf. Comput.*, 123(1):72–110, November 1995.
- [5] Holger Bast and Torben Hagerup. Fast and reliable parallel hashing. In *SPAA*, pages 50–61, 1991.
- [6] Dan Boneh, David Mazieres, and Raluca Ada Popa. Remote oblivious storage: Making oblivious RAM practical. <http://dspace.mit.edu/bitstream/handle/1721.1/62006/MIT-CSAIL-TR-2011-018.pdf>, 2011.
- [7] Elette Boyle, Kai-Min Chung, and Rafael Pass. Oblivious parallel ram. <https://eprint.iacr.org/2014/594.pdf>.
- [8] Kai-Min Chung, Zhenming Liu, and Rafael Pass. Statistically-secure oram with $\tilde{O}(\log^2 n)$ overhead. *CoRR*, abs/1307.3699, 2013.
- [9] Christopher W. Fletcher, Marten van Dijk, and Srinivas Devadas. A secure processor architecture for encrypted computation on untrusted programs. In *STC*, 2012.
- [10] Christopher W. Fletcher, Ling Ren, Albert Kwon, Marten Van Dijk, Emil Stefanov, and Srinivas Devadas. Tiny ORAM: A low-latency, low-area hardware oram controller with integrity verification.
- [11] Christopher W. Fletcher, Ling Ren, Albert Kwon, Marten van Dijk, Emil Stefanov, and Srinivas Devadas. RAW Path ORAM: A low-latency, low-area hardware ORAM controller with integrity verification. *IACR Cryptology ePrint Archive*, 2014:431, 2014.
- [12] Christopher W. Fletcher, Ling Ren, Xiangyao Yu, Marten van Dijk, Omer Khan, and Srinivas Devadas. Suppressing the oblivious RAM timing channel while making information leakage and program efficiency trade-offs. In *HPCA*, pages 213–224, 2014.
- [13] Craig Gentry, Kenny A. Goldman, Shai Halevi, Charanjit S. Jutla, Mariana Raykova, and Daniel Wichs. Optimizing ORAM and using it efficiently for secure computation. In *Privacy Enhancing Technologies Symposium (PETS)*, 2013.
- [14] Craig Gentry, Shai Halevi, Steve Lu, Rafail Ostrovsky, Mariana Raykova, and Daniel Wichs. Garbled ram revisited. In *Advances in Cryptology - EUROCRYPT 2014*, volume 8441, pages 405–422. 2014.
- [15] Craig Gentry, Shai Halevi, Mariana Raykova, and Daniel Wichs. Garbled ram revisited, part i. *Cryptology ePrint Archive*, Report 2014/082, 2014. <http://eprint.iacr.org/>.
- [16] Craig Gentry, Shai Halevi, Mariana Raykova, and Daniel Wichs. Outsourcing private ram computation. *IACR Cryptology ePrint Archive*, 2014:148, 2014.
- [17] Joseph Gil, Yossi Matias, and Uzi Vishkin. Towards a theory of nearly constant time parallel algorithms. In *32nd Annual Symposium on Foundations of Computer*

- Science (FOCS)*, pages 698–710, 1991.
- [18] O. Goldreich. Towards a theory of software protection and simulation by oblivious RAMs. In *ACM Symposium on Theory of Computing (STOC)*, 1987.
 - [19] Oded Goldreich and Rafail Ostrovsky. Software protection and simulation on oblivious RAMs. *J. ACM*, 1996.
 - [20] Michael T. Goodrich and Michael Mitzenmacher. Privacy-preserving access of outsourced data via oblivious RAM simulation. In *ICALP*, 2011.
 - [21] Michael T. Goodrich, Michael Mitzenmacher, Olga Ohrimenko, and Roberto Tamassia. Practical oblivious storage. In *ACM Conference on Data and Application Security and Privacy (CODASPY)*, 2012.
 - [22] Michael T. Goodrich, Michael Mitzenmacher, Olga Ohrimenko, and Roberto Tamassia. Privacy-preserving group data access via stateless oblivious RAM simulation. In *SODA*, 2012.
 - [23] S. Dov Gordon, Jonathan Katz, Vladimir Kolesnikov, Fernando Krell, Tal Malkin, Mariana Raykova, and Yevgeniy Vahlis. Secure two-party computation in sublinear (amortized) time. In *ACM CCS*, 2012.
 - [24] Torben Hagerup. The log-star revolution. In *STACS 92, 9th Annual Symposium on Theoretical Aspects of Computer Science, Cachan, France, February 13-15, 1992, Proceedings*, pages 259–278, 1992.
 - [25] Eyal Kushilevitz, Steve Lu, and Rafail Ostrovsky. On the (in)security of hash-based oblivious RAM and a new balancing scheme. In *SODA*, 2012.
 - [26] Chang Liu, Yan Huang, Elaine Shi, Jonathan Katz, and Michael Hicks. Automating efficient ram-model secure computation. In *IEEE S & P*. IEEE Computer Society, 2014.
 - [27] Steve Lu and Rafail Ostrovsky. Distributed oblivious RAM for secure two-party computation. In *Theory of Cryptography Conference (TCC)*, 2013.
 - [28] Steve Lu and Rafail Ostrovsky. How to garble ram programs. In *EUROCRYPT*, pages 719–734, 2013.
 - [29] Steve Lu and Rafail Ostrovsky. Garbled ram revisited, part ii. Cryptology ePrint Archive, Report 2014/083, 2014. <http://eprint.iacr.org/>.
 - [30] Martin Maas, Eric Love, Emil Stefanov, Mohit Tiwari, Elaine Shi, Krste Asanovic, John Kubiatowicz, and Dawn Song. Phantom: Practical oblivious computation in a secure processor. In *CCS*, 2013.
 - [31] Martin Maas, Eric Love, Emil Stefanov, Mohit Tiwari, Elaine Shi, Krste Asanovic, John Kubiatowicz, and Dawn Song. A high-performance oblivious RAM controller on the convey hc-2ex heterogeneous computing platform. In *Workshop on the Intersections of Computer Architecture and Reconfigurable Logic (CARL)*, 2013.
 - [32] Kurt Mehlhorn and Uzi Vishkin. Randomized and deterministic simulations of prams by parallel machines with restricted granularity of parallel memories. *Acta Inf.*, 21:339–374, 1984.
 - [33] Rafail Ostrovsky and Victor Shoup. Private information storage (extended abstract). In *ACM Symposium on Theory of Computing (STOC)*, 1997.
 - [34] Rasmus Pagh and Flemming Friche Rodler. Cuckoo hashing. *J. Algorithms*, 51(2):122–144, May 2004.

- [35] Ling Ren, Xiangyao Yu, Christopher W. Fletcher, Marten van Dijk, and Srinivas Devadas. Design space exploration and optimization of path oblivious RAM in secure processors. In *ISCA*, pages 571–582, 2013.
- [36] Elaine Shi, T.-H. Hubert Chan, Emil Stefanov, and Mingfei Li. Oblivious RAM with $O((\log N)^3)$ worst-case cost. In *ASIACRYPT*, 2011.
- [37] Emil Stefanov and Elaine Shi. Oblivstore: High performance oblivious cloud storage. In *IEEE Symposium on Security and Privacy (S & P)*, 2013.
- [38] Emil Stefanov, Elaine Shi, and Dawn Song. Towards practical oblivious RAM. In *NDSS*, 2012.
- [39] Emil Stefanov, Marten van Dijk, Elaine Shi, T.-H. Hubert Chan, Christopher Fletcher, Ling Ren, Xiangyao Yu, and Srinivas Devadas. Path ORAM: an extremely simple oblivious ram protocol. In *ACM CCS*, 2013.
- [40] G. Edward Suh, Dwaine Clarke, Blaise Gassend, Marten van Dijk, and Srinivas Devadas. Aegis: architecture for tamper-evident and tamper-resistant processing. In *International conference on Supercomputing, ICS '03*, pages 160–171, 2003.
- [41] David Lie Chandramohan Thekkath, Mark Mitchell, Patrick Lincoln, Dan Boneh, John Mitchell, and Mark Horowitz. Architectural support for copy and tamper resistant software. *SIGOPS Oper. Syst. Rev.*, 34(5):168–177, November 2000.
- [42] Uzi Vishkin. Can parallel algorithms enhance serial implementation? *Commun. ACM*, 39(9):88–91, 1996.
- [43] Xiao Shaun Wang, T.-H. Hubert Chan, and Elaine Shi. Circuit ORAM: On tightness of the goldreich-ostrovksy lower bound. <http://eprint.iacr.org/2014/672.pdf>.
- [44] Xiao Shaun Wang, Yan Huang, T.-H. Hubert Chan, abhi shelat, and Elaine Shi. Scoram: Oblivious ram for secure computation. <http://eprint.iacr.org/2014/671.pdf>.
- [45] Peter Williams and Radu Sion. Usable PIR. In *Network and Distributed System Security Symposium (NDSS)*, 2008.
- [46] Peter Williams and Radu Sion. SR-ORAM: Single round-trip oblivious ram. In *ACM Conference on Computer and Communications Security (CCS)*, 2012.
- [47] Peter Williams, Radu Sion, and Bogdan Carbunar. Building castles out of mud: Practical access pattern privacy and correctness on untrusted storage. In *CCS*, 2008.
- [48] Peter Williams, Radu Sion, and Alin Tomescu. Privatefs: A parallel oblivious file system. In *CCS*, 2012.
- [49] Xiangyao Yu, Syed Kamran Haider, Ling Ren, Christopher W. Fletcher, Albert Kwon, Marten van Dijk, and Srinivas Devadas. Proram: dynamic prefetcher for oblivious RAM. In *Proceedings of the 42nd Annual International Symposium on Computer Architecture, Portland, OR, USA, June 13-17, 2015*, pages 616–628, 2015.