

# Verifiable Oblivious Storage

Daniel Apon, Jonathan Katz, Elaine Shi, and Aishwarya Thiruvengadam

University of Maryland, College Park, MD 20742, USA  
{dapon,jkatz,elaine,aish}@cs.umd.edu

**Abstract.** We formalize the notion of *Verifiable Oblivious Storage* (VOS), where a client outsources the storage of data to a server while ensuring data confidentiality, access pattern privacy, and integrity and freshness of data accesses. VOS generalizes the notion of Oblivious RAM (ORAM) in that it allows the server to perform computation, and also explicitly considers data integrity and freshness.

We show that allowing server-side computation enables us to construct asymptotically more efficient VOS schemes whose bandwidth overhead cannot be matched by any ORAM scheme, due to a known lower bound by Goldreich and Ostrovsky. Specifically, for large block sizes we can construct a VOS scheme with constant bandwidth per query; further, answering queries requires only poly-logarithmic server computation. We describe applications of VOS to Dynamic Proofs of Retrievability, and RAM-model secure multi-party computation.

## 1 Introduction

Oblivious RAM (ORAM) is a notion first proposed by Goldreich and Ostrovsky [20] in the context of protecting software from piracy. They consider an application in which a trusted CPU wishes to hide its memory-access patterns from an attacker who can view (and possibly modify) the entire contents of memory. Recently, as cloud computing has gained in popularity, ORAM has been recast as a means to securely outsource storage to an untrusted server, while hiding access patterns from the server.

In this paper, we propose Verifiable Oblivious Storage (VOS), which generalizes the notion of ORAM by allowing the storage medium to perform computation. In addition, it also explicitly incorporates notions of integrity and freshness. We will refer to integrity and freshness as verifiability in this paper.

**Formally Defining VOS.** Our first contribution is to formally define VOS, and to differentiate the notion of VOS from ORAM. While we are the first to formalize the VOS notion, VOS has implicitly been used by other researchers earlier, often being referred to as ORAM. For example, Williams and Sion [35] recently proposed a scheme that improves round-complexity to  $O(1)$  — since their scheme leverages server-side computation, it is implicitly a VOS scheme.

An important difference between VOS and ORAM schemes is that VOS schemes can be constructed to achieve asymptotically better bandwidth overhead than what can be achieved by any ORAM scheme. This is because all

ORAM schemes are subject to a well-known lower bound result by Goldreich and Ostrovsky [20]. This result, however, does not apply to VOS.

Several applications where ORAM was previously employed can immediately achieve asymptotic bandwidth savings if we simply replace the ORAM with a VOS construction. For example, we know from prior work that RAM-model secure multi-party computation [24] and Dynamic Proofs of Retrievability [9] can be built using ORAM as a building block. In both these applications, the party storing data (or a share of the data) can perform computation. By replacing the underlying ORAM with a VOS in these constructions [9,24], we can immediately obtain asymptotic bandwidth savings as illustrated in Section 5.

**Asymptotically Efficient VOS Construction.** We show that, by allowing server-side computation, VOS schemes can be constructed that beat the known logarithmic lower bound on the bandwidth cost for any ORAM scheme [20]. Specifically, we show that there exists a VOS scheme with block size  $\beta = \tilde{\Omega}(\lambda)$  (where  $\lambda$  is the security parameter) having  $O(\beta)$  bandwidth cost for reading or writing a block; this scheme has  $O(\beta)$  client-side storage, and uses only  $O(1)$  roundtrips and requires only  $O(\beta \cdot \text{poly log } n) \cdot \text{poly}(\lambda)$  server-side computation per data access. This is asymptotically better than what any ORAM scheme can hope to achieve since, due to the lower bound by Goldreich and Ostrovsky, any ORAM scheme must have bandwidth cost  $\Omega(\beta \log n)$  to read or write a block of  $\beta$  bits. *Note that this lower bound holds regardless of the block size  $\beta$ .*

## 1.1 Technical Highlight

**Generic ORAM-to-VOS Compiler in the Semi-honest Model.** To construct efficient VOS schemes, we rely on fully homomorphic encryption (FHE) to encrypt and outsource the entire ORAM memory, as well as the ORAM client’s secret state. The server can now perform computation on behalf of the client, without learning any secrets. The only occasion when the server needs to contact the client is to seek the client’s help to decrypt the next physical address or sequence of physical addresses to read or write. The use of FHE or PIR to outsource the ORAM client’s computation has been mentioned in earlier works [16,26]. The main challenge, however, is *how to ensure security when the server is malicious, and may arbitrarily deviate from the prescribed computation.*

**Generic ORAM-to-VOS Compiler in the Malicious Model.** Achieving security against a malicious server is much trickier in the VOS setting than in ORAM. While ORAM achieves integrity and freshness in a straightforward way by employing standard storage integrity techniques such as message authentication codes and Merkle hash trees, in VOS, we need to worry about a server that can arbitrarily deviate from the prescribed computation.

Naive applications of well-known techniques such as SNARKS [3,4,15] result in server computation that is linear in the size of the dataset. Instead, we leverage efficient Verifiable RAM computation (VC-RAM) to enforce honest server behavior. This allows us to achieve sublinear server computation. In VC-RAM,

a client outsources a large memory array to a server in a preprocessing step. Afterwards in an online stage, the client specifies a sequence of inputs, and asks the server to compute a RAM program over the outsourced memory array and the inputs. Each query made by the client can result in updates to the server’s memory array. VC-RAM allows a client to verify the result of these RAM computations, and meanwhile, the server’s computation overhead is sublinear (in the data size) for sublinear-time queries.

Although VC-RAM has been informally mentioned in earlier works [2, 4, 7], we make the contribution of explicitly formalizing *stateful* VC-RAM (for repeated queries) and its security. We also present an efficient VC-RAM scheme with constant proof size and prover computation that is comparable to the run-time of the RAM program (as opposed to the dataset size).

**Non-generic Optimizations for Specific Schemes.** We then apply these techniques to two existing ORAM schemes, the Path ORAM [33] and the Hierarchical ORAM by Goodrich and Mitzenmacher [21] that was later improved by Kushilevitz et al. [25]. The resulting VOS schemes are referred to as Path VOS and Hierarchical VOS respectively.

Applying Verifiable RAM Computation (VC-RAM) straight out-of-the-box is not sufficient to achieve the claimed asymptotic bounds for the Path VOS. We show how to tailor our VC-RAM techniques for the Path VOS to shave a  $O(\log n)$  factor off the server computation. Similarly, for the hierarchical VOS, we propose rebalancing techniques that can shave poly  $\log \log n$  to  $\log n$  factors from the bandwidth cost, at increased (but still sublinear) server computation.

While the Path VOS is asymptotically better than the hierarchical VOS, the hierarchical VOS is necessary for our dynamic PoR application, since the Path VOS does not satisfy the next-read pattern hiding property [9].

## 1.2 Related Work

Oblivious RAM (ORAM) was first proposed by Goldreich and Ostrovsky [20], and later improved in a series of works [13, 14, 16, 19, 21–23, 25, 27–30, 32–36]. Recently, ORAM has been used in outsourcing storage [21, 34, 36], and in secure two-party computation to achieve sublinear amortized cost [16, 24].

ORAM with implicit server computation has appeared in several works [16, 35], while still being referred to as ORAM. Williams and Sion rely on server-side computation to achieve a single-round ORAM scheme [35]. Their scheme ensures privacy against a malicious server but not integrity and freshness and has an asymptotic bandwidth cost of  $\tilde{O}(\beta \log^2 n)$ . In comparison, our VOS scheme is asymptotically more efficient, and ensures both privacy and integrity/freshness against a malicious server. Gentry et al. [16] proposed using homomorphic encryption to improve ORAM bandwidth cost. However, their scheme is only secure in the semi-honest model, and is also asymptotically more expensive in bandwidth than our construction. Mayberry et al. also proposed to leverage PIR techniques in combination with ORAM [26]. They too are implicitly using VOS; their scheme is not secure in the malicious model, and is asymptotically less efficient than our construction.

Private Information Retrieval (PIR) [6, 11, 12, 18] allows a client to access a dataset on the server obliviously. Single-server PIR techniques can achieve  $O(\beta)$  bandwidth cost per query using FHE techniques [6] for large enough block sizes  $\beta$ . However, single-server PIR requires server computation that is linear in the size of the dataset. Also, PIR works for public datasets; in VOS, we consider a private dataset owned by the client, which is not exposed to the server.

## 2 Definitions of Verifiable Oblivious Storage

We use  $((c\_out, c\_state), (s\_out, s\_state)) \leftarrow \text{protocol}((c\_in, c\_state), (s\_in, s\_state))$  to denote a (stateful) protocol between a client and server, where  $c\_in$  and  $c\_out$  are the client’s input and output;  $s\_in$  and  $s\_out$  are the server’s input and output; and  $c\_state$  and  $s\_state$  are the client and server’s states before and after the protocol.

We define the notion of *Verifiable Oblivious Storage* (VOS), in which a client outsources the storage of data to a server while ensuring privacy of the data and verifiability and obliviousness of access to that data.

**Definition 1 (Verifiable Oblivious Storage).** *A Verifiable Oblivious Storage (VOS) scheme consists of the following interactive protocols between a client and a server.*

$((\perp, z), (\perp, Z)) \leftarrow \text{Setup}(1^\lambda, (D, \perp), (\perp, \perp))$ : An interactive protocol where the client’s input is a memory array  $D[1..n]$  where each memory *block* has bit-length  $\beta$ ; and the server’s input is  $\perp$ . At the end of the **Setup** protocol, the client has secret state  $z$ , and server’s state is  $Z$  (which typically encodes the memory array  $D$ ).

$((\text{data}, z'), (\perp, Z')) \leftarrow \text{Access}((\text{op}, z), (\perp, Z))$ : To access data, the client starts in state  $z$ , with an input  $\text{op}$  where  $\text{op} := (\text{read}, \text{ind})$  or  $\text{op} := (\text{write}, \text{ind}, \text{data})$ ; the server starts in state  $Z$ , and has no input. In a correct execution of the protocol, the client’s output  $\text{data}$  is the current value of the memory  $D$  at location  $\text{ind}$  (for writes, the output is the old value of  $D[\text{ind}]$  before the write takes place). The client and server also update their states to  $z'$  and  $Z'$  respectively. The client outputs  $\text{data} := \perp$  if the protocol execution aborted.

We say that a VOS scheme is correct, if for any initial memory  $D \in \{0, 1\}^{\beta n}$ , for any operation sequence  $\text{op}_1, \text{op}_2, \dots, \text{op}_m$  where  $m = \text{poly}(\lambda)$ , an  $\text{op} := (\text{read}, \text{ind})$  operation would always return the last value written to the logical location  $\text{ind}$  (except with negligible probability).

### 2.1 Security Definition

We adopt a standard simulation-based definition of secure computation [8], requiring that a real-world execution “simulate” an ideal-world (reactive) functionality  $\mathcal{F}$ . At an intuitive level, our definition captures the privacy and verifiability requirements for an honest client, in the presence of a malicious server.

**Ideal World.** We define an ideal functionality  $\mathcal{F}$  that maintains an up-to-date version of the data  $D$  on behalf of the client, and answers the client’s access queries.

- *Setup.* An environment  $\mathcal{Z}$  gives an initial database  $D$  to the client. The client sends  $D$  to an ideal functionality  $\mathcal{F}$ .  $\mathcal{F}$  notifies the ideal-world adversary  $\mathcal{S}$  (of the setup operation, but not of the data contents  $D$ ). The ideal-world adversary  $\mathcal{S}$  says `ok` or `abort` to  $\mathcal{F}$ .  $\mathcal{F}$  then says `ok` or  $\perp$  to the client accordingly.
- *Access.* In each time step, the environment  $\mathcal{Z}$  specifies an operation  $\text{op} := (\text{read}, \text{ind})$  or  $\text{op} := (\text{write}, \text{ind}, \text{data})$  as the client’s input. The client sends  $\text{op}$  to  $\mathcal{F}$ .  $\mathcal{F}$  notifies the ideal-world adversary  $\mathcal{S}$  (without revealing to  $\mathcal{S}$  the operation  $\text{op}$ ). If  $\mathcal{S}$  says `ok` to  $\mathcal{F}$ ,  $\mathcal{F}$  sends  $D[\text{ind}]$  to the client, and updates  $D[\text{ind}] := \text{data}$  accordingly if this is a write operation. The client then forwards  $D[\text{ind}]$  to the environment  $\mathcal{Z}$ . If  $\mathcal{S}$  says `abort` to  $\mathcal{F}$ ,  $\mathcal{F}$  sends  $\perp$  to the client.

**Real World.** In the real world, an environment  $\mathcal{Z}$  gives an honest client a database  $D$ . The honest client runs the `Setup` protocol with the server  $\mathcal{A}$ . Then at each time step,  $\mathcal{Z}$  specifies an input  $\text{op} := (\text{read}, \text{ind})$  or  $\text{op} := (\text{write}, \text{ind}, \text{data})$  to the client. The client then runs the `Access` protocol with the server. The environment  $\mathcal{Z}$  gets the view of the adversary  $\mathcal{A}$  after every operation. The client outputs to the environment the data fetched or  $\perp$  (indicating abort).

**Definition 2 (Simulation-based security: privacy + verifiability).** *We say that a protocol  $\Pi_{\mathcal{F}}$  securely computes the ideal functionality  $\mathcal{F}$  if for any probabilistic polynomial-time real-world adversary (i.e., server)  $\mathcal{A}$ , there exists an ideal-world adversary  $\mathcal{S}$ , such that for all non-uniform, polynomial-time environment  $\mathcal{Z}$ , there exists a negligible function  $\text{negl}$  such that*

$$|\Pr[\text{REAL}_{\Pi_{\mathcal{F}}, \mathcal{A}, \mathcal{Z}}(\lambda) = 1] - \Pr[\text{IDEAL}_{\mathcal{F}, \mathcal{S}, \mathcal{Z}}(\lambda) = 1]| \leq \text{negl}(\lambda)$$

This definition is simulation-based [8] where the client is honest, and the server is corrupted. (The client is never malicious in our setting.) The definition also simultaneously captures *privacy* and *verifiability*. Intuitively, privacy ensures that the server cannot observe the data contents or the access pattern. Verifiability ensures that the client is guaranteed to read the correct data from the server — if the server happens to cheat, the client can detect it and abort the protocol.

### 3 ORAM to VOS: Generic Compilation Techniques

In this section, we describe how to generically transform any given ORAM scheme to an efficient VOS scheme. In Section 4, we give two specific VOS schemes - Path VOS and Hierarchical VOS. These are derived from the two classes of ORAM schemes, the hierarchical construction [20] and its variants [21, 23, 25, 29, 34–36], and the binary-tree scheme [30] and its variants [13, 16, 33].

### 3.1 Preliminary: Oblivious RAM

In this paper, we use a slightly different formalization of ORAM from that of Goldreich-Ostrovsky [20] to make notation simpler for our generic compiler.

An ORAM can be defined by a pair of algorithms  $\text{ORAM} := (\text{Init}, \text{Next})$ :

- $(D_o, st) \leftarrow \text{Init}(1^\lambda, D)$ : Takes in storage array  $D$  containing  $n$  blocks each of bit length  $\beta$ , produces storage array  $D_o$ , and initial ORAM client state  $st$ .
- $(\text{out}, \{\text{raddr}\}, \{\text{waddr}\}, \{\text{data}\}, st) \leftarrow \text{Next}(\text{op}, st, \{\text{fetched}\})$ : Each ORAM operation  $\text{op} := (\text{read}, \text{ind})$  or  $\text{op} := (\text{write}, \text{ind}, \text{data})$  will proceed in *multiple rounds*. Each round will invoke the  $\text{ORAM.Next}$  algorithm with the following inputs: 1) current read/write operation  $\text{op}$ ; 2) the (secret) ORAM client state  $st$ ; and 3) a set of blocks  $\{\text{fetched}\}$  fetched from the last round. If this is the first round for an operation  $\text{op}$ , this fetched set is empty by convention. The  $\text{ORAM.Next}$  function in turn outputs a set of addresses to read in the next round denoted  $\{\text{raddr}\}$ ; a set of addresses  $\{\text{waddr}\}$  and data  $\{\text{data}\}$  to write in the next round; updates the client state  $st$ ; and if this is the last round,  $\text{ORAM.Next}$  also outputs the block read out.

The Next algorithm performs one round of the ORAM client computation.

Our notation is explained in the table below:

$st$	secret ORAM client state	$\{\text{raddr}\}$	physical addr to read from
$ind$	logical index of a block	$\{\text{waddr}\}$	physical addr to write to
$\text{op} := (\text{read}, \text{ind})$ or $\text{op} := (\text{write}, \text{ind}, \text{data})$	a read/write operation	$\{\text{fetched}\}$	data blocks fetched from storage
$\text{out}$	the last logical block read	$\{\text{data}\}$	data blocks to be written

Security is defined in terms of the inability of any PPT adversary to distinguish the access patterns generated by an honest execution of the ORAM client, from those output by a simulator that does not see the sequence of logical operations.

**Definition 3 (ORAM security).** *We say that an ORAM scheme is secure, if there exists a stateful simulator  $\text{Sim}$ , such that for any PPT adversary  $\mathcal{A}$ ,*

$$\left| \Pr \left[ \mathcal{A}^{O[st](\cdot)}(1^\lambda) = 1 \right] - \Pr \left[ \mathcal{A}^{\text{Sim}(1^\lambda, m)}(1^\lambda) = 1 \right] \right| \leq \text{negl}(\lambda) \quad (1)$$

where  $m$  is the number of oracle queries made by the adversary  $\mathcal{A}$ ; and  $O[st](\cdot)$  denotes a stateful oracle  $O$ , with secret state  $st$ . The oracle  $O$  takes in an operation  $\text{op}$  and outputs a sequence of read and write physical addresses. Formally,

**Oracle O:**

**Initialization.** On input  $D$  containing  $n$  blocks each of  $\beta$  bits, initialize a storage array  $D_o$  containing  $n_o$  blocks each of size  $\beta_o$ . Initialize the set  $\{\text{fetched}\}$  to be an empty set. Run  $st := \text{ORAM.Init}(1^\lambda, n, \beta)$ .

**Data Access.** On the  $j$ -th input  $\text{op}_j$ ,  $j \in \mathbb{N}$ , perform the following:

- First, initialize the output array  $\Gamma := \emptyset$ .
- For  $\text{rnd} = 1$  to  $R_j$  where  $R_j$  is the total number of rounds for the  $j$ -th operation<sup>a</sup>:
  - Run  $(\text{out}, \{\text{raddr}\}, \{\text{waddr}\}, \{\text{data}\}, st) \leftarrow \text{ORAM.Next}(\text{op}, st, \{\text{fetched}\})$
  - Let  $D_o[\{\text{waddr}\}] := \{\text{data}\}$ , and let  $\{\text{fetched}\} := D_o[\{\text{raddr}\}]$ .
  - Append  $\{\text{raddr}\}$  and  $\{\text{waddr}\}$  to the output set  $\Gamma$ .
- Finally, output  $\Gamma$ .

<sup>a</sup> In all known ORAM constructions, due to the obliviousness requirement,  $R_j$  is a public value determined by the ORAM scheme description itself, and does not depend on the input sequence.

We use  $D_o[\{\text{waddr}\}] := \{\text{data}\}$  and  $\{\text{fetched}\} := D_o[\{\text{raddr}\}]$  to denote writing  $\{\text{data}\}$  to a set of write addresses  $\{\text{waddr}\}$ , and reading from a set of read addresses  $\{\text{raddr}\}$  respectively. We assume that  $\{\text{waddr}\}$  and  $\{\text{data}\}$  are ordered sets, and we simply write each block  $\text{data}$  into each  $\text{waddr}$  in the specified order.

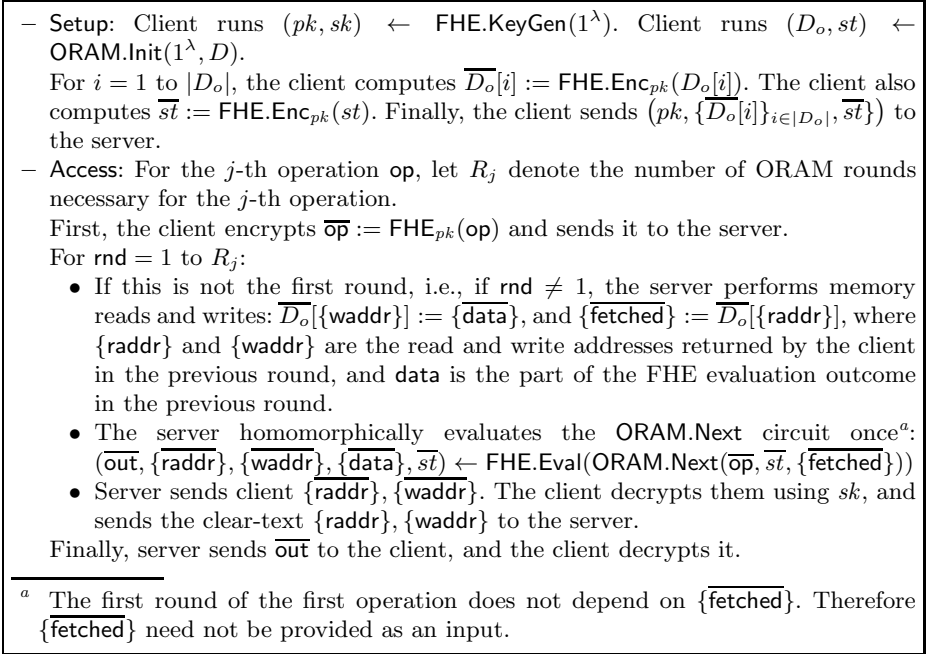
**Deterministic vs. Randomized ORAM.** In general, the ORAM client algorithms  $\text{Init}$  and  $\text{Next}$  can be randomized. However, the  $\text{Next}$  algorithm can be made deterministic by choosing a PRF key  $k$  at random and including it in the client state  $st$ . Whenever  $\text{Next}$  requires random bits, this can be generated pseudorandomly from key  $k$ . If the randomized ORAM is secure, then the resulting ORAM with a deterministic  $\text{Next}$  algorithm is also secure due to the security of the PRF. Therefore, without loss of generality, in our generic ORAM-to-VOS compiler, we will assume an ORAM scheme with a deterministic  $\text{Next}$  algorithm.

### 3.2 Compilation in the Semi-honest Model

**Intuition.** The intuition is to have the client outsource the ORAM memory encrypted under an FHE scheme to the server. The client can then outsource all its computation to the server as well, since the server can homomorphically operate over the encrypted data. In this manner, the server only contacts the client whenever it is necessary for interaction during the computation.

**ORAM-to-VOS Compiler in the Semi-honest Model.** Figure 1 describes how to transform an ORAM scheme to a VOS scheme that is secure under a semi-honest server.

**Theorem 1.** *Let  $\text{FHE} = (\text{KeyGen}, \text{Enc}, \text{Dec}, \text{Eval})$  be a semantically secure FHE scheme and let  $\text{ORAM} = (\text{Init}, \text{Next})$  be a secure ORAM scheme. Then, the generic compiler in Figure 1 gives a Verifiable Oblivious Storage (VOS) construction secure under a semi-honest server.*



**Fig. 1.** ORAM-to-VOS generic compiler: semi-honest model

The proof of Theorem 1 reduces to the security of the encryption scheme and the ORAM scheme in a straightforward manner. We refer the reader to our online technical report [1] for a detailed proof.

**Optimization: Handling Addresses Independent of Secret Information.** In the construction above, the server performs as much computation as possible and only seeks the client’s help when it needs to decrypt the next set of physical addresses to read from or write to. In many ORAM schemes, there are read/write operations whose physical addresses do not depend on secret client state, memory contents, or the logical addresses accessed. Examples are the reshuffling operations of the hierarchical ORAM scheme [20] and its variants [21–23, 25, 34, 36] and the eviction operations of the binary-tree based ORAM [30] and its variants [33]. To achieve better efficiency, such reshuffling and eviction operations, can be performed by the server (on its own) homomorphically, without seeking the client’s help to decrypt the physical addresses.

### 3.3 Handling Malicious Servers

One way to handle a malicious server is to rely on a Succinct Non-Interactive Argument of Knowledge (SNARK). However, if done naively, the circuit for the SNARK will have size that is at least linear in  $n$ , i.e., the size of the outsourced memory  $D$ . This requires the server to perform a linear amount of computation to produce a proof of correctness.



Instead, we rely on efficient Verifiable RAM Computation (VC-RAM) to enforce honest server behavior. Verifiable RAM computation has been informally introduced in the literature by Ben-Sasson et al. [2] and Bitansky et al. [4]. We, however, need a stateful version of verifiable RAM computation. Braun et al. also informally proposed and implemented verifiable RAM computation [7].

We define a *stateful* version of verifiable RAM computation, where each query can result in updates to the outsourced dataset. Below, we explicitly formalize this notion of stateful, multi-query VC-RAM. Relying on the same ideas as Ben-Sasson et al. [2] and Bitansky et al. [4], we show that verifying RAM computation can be done efficiently, resulting in server computation that is comparable to the run-time of the RAM program (as opposed to the size of the memory); succinct proofs of size  $O(\lambda)$ ; and efficient client verification time that is not too much worse than simply reading the input and output.

**Verifiable RAM Computation.** Consider a scenario where a client outsources a memory array  $D$  to a server. Let  $f$  denote a RAM program agreed upon by the client and the server. At each time step  $t$ , the client supplies a small input  $x_i$ , and the server computes the RAM program  $f$  over  $x_i$  and the current state of memory  $D$ . The RAM program produces an answer which is sent to the client. It may also update the memory contents outsourced to the server – hence our notion of VC-RAM is stateful. Verifiability requires that the client be able to check that the RAM computation results returned by the server are correct.

**Definition 4 (Verifiable RAM Computation).** *A (non-interactive) Verifiable RAM Computation (VC-RAM) scheme consists of the following algorithms:*

- $(z, Z) \leftarrow \text{Setup}(1^\lambda, D, f)$ : Given an initial memory array  $D[1..n]$  where each memory word has bit-length  $\ell$ , a RAM program description  $f$ , output initial server state  $Z$  (which typically encodes  $D$ ), and the initial client state  $z$ .
- $(\bar{y}, Z') \leftarrow \text{Compute}(x, Z)$ : Given a small input  $x$  to the RAM program  $f$ , the server's current state  $Z$ , output an encoded answer  $\bar{y}$ , and updated server state  $Z'$ .<sup>1</sup>
- $(y, b, z') \leftarrow \text{Verify}(x, \bar{y}, z)$ : Given the input  $x$ , the client's current state  $z$ , an encoded answer  $\bar{y}$ , output a decoded answer  $y$ , a bit  $b$  indicating whether to accept this answer, and updated client state  $z'$ .

Correctness is defined as usual. We require that for any parameters  $n$  and  $\ell$ , for any initial memory array  $D \in \{0, 1\}^{\ell n}$ , for any polynomial-sized RAM program  $f$  which terminates in polynomial time, for any query sequence  $x_1, x_2, \dots, x_m$  where  $m = \text{poly}(\lambda)$ ,

$$\Pr \left[ \exists i : \begin{array}{l} (y_i \neq f(D, x_1, x_2, \dots, x_i)) \\ \vee (b_i = 0) \end{array} \mid \begin{array}{l} (z, Z_0) \leftarrow \text{Setup}(1^\lambda, D, f) \\ \forall i \in \{1, 2, \dots, m\} : \\ (\bar{y}_i, Z_i) \leftarrow \text{Compute}(x_i, Z_{i-1}) \\ (y_i, b_i, z) \leftarrow \text{Verify}(x_i, \bar{y}_i, z) \end{array} \right] \leq \text{negl}(\lambda)$$

<sup>1</sup> In the specific VC-RAM construction we describe, the encoded answer  $\bar{y}$  includes the answer itself  $y$ , a proof vouching for its correctness, and an updated digest of the outsourced memory.

In particular, we use the notation  $y_i := f(D, x_1, x_2, \dots, x_i)$  to denote the outcome of the  $i$ -th query, starting with an initial memory array of  $D$ , and after computing the RAM program  $f$  on queries  $x_1, x_2, \dots, x_i$ . Note that each query is stateful, i.e., may result in updates to the memory array  $D$ .

**Definition 5 (Verifiability of VC-RAM).** *We say that a VC-RAM scheme is verifiable, if for any polynomial time (stateful) adversary  $\mathcal{A}$  the following holds.*

$$\Pr \left[ \begin{array}{l} \exists i : (b_i = 1) \wedge \\ (y_i \neq f(D, x_1, x_2, \dots, x_i)) \end{array} \middle| \begin{array}{l} (D, f) \leftarrow \mathcal{A}(1^\lambda) \\ (z, Z) \leftarrow \text{Setup}(1^\lambda, D, f) \\ (x_1, \bar{y}_1) \leftarrow \mathcal{A}(Z) \\ \forall i \in \{1, 2, \dots, m\} : \\ \quad (y_i, b_i, z) \leftarrow \text{Verify}(x_i, \bar{y}_i, z) \\ \quad (x_{i+1}, \bar{y}_{i+1}) \leftarrow \mathcal{A}(y_i, b_i) \end{array} \right] \leq \text{negl}(\lambda)$$

Note again that the adversary  $\mathcal{A}$  is stateful, and we do not write its state explicitly for simplicity.

**Theorem 2.** *There exists a non-interactive VC-RAM scheme such that for each query: the server runs in time  $\tilde{O}(\tau \log n) \text{poly}(\lambda)$  where  $\tau$  is the run-time of the RAM program in the unauthenticated setting; the verifier runs in time  $O((|x| + |y|)\lambda)$ ; and the client-server bandwidth cost is  $|x| + |y| + O(\lambda)$ .*

Note that the client-server bandwidth cost has to be at least  $|x| + |y|$ , i.e., the number of bits necessary to transmit the query  $x$  and the answer  $y$ . Therefore, the only additional cost is  $O(\lambda)$  for transmitting an updated digest of the outsourced memory and a proof vouching for the correctness of the result.

We explain the intuition for the VC-RAM construction. The full construction can be found in our online technical report [1]. The high level idea is to build a Merkle tree over all outsourced memory, such that the client keeps the up-to-date root digest. To verify a RAM computation, we build a “verifier circuit” which takes in a trace of the computation, including 1) the CPU states before and after every computation step; 2) the memory contents fetched in every computation step; and 3) the Merkle-tree digest before and after each computation step. This verifier circuit checks the trace of the computation: 1) it checks that every memory read and write is correct using memory checking; and 2) it checks that every CPU computation step is correct. The server then constructs a SNARK for this “verifier circuit”. Since this verifier circuit has size that is roughly the time of the RAM computation, we can achieve prover efficiency, i.e., the prover time is roughly the time of the RAM computation rather than the size of the entire dataset.

**Relying on VC-RAM to Enforce Honest Server Behavior.** In our semi-honest VOS construction described in Section 3.2, the client essentially outsources all of its ORAM memory (encrypted under FHE) to the server, as well as the ORAM’s secret client state (also encrypted under FHE).

During each data access operation, in every round of interaction, the server performs some RAM computation on behalf of the client, and sends a message to the client to seek its help decrypting certain physical addresses. Using VC-RAM, the server can attach a succinct proof along with every message sent to the client, vouching for the correctness of the message. If the message sent to the client deviates from correct message, the client will surely detect it (except with negligible probability).

Due to space limits, we state the theorem below, and give the formal presentation of the malicious-model ORAM-to-VOS compiler in our online technical report [1].

**Theorem 3.** *Assuming existence of SNARKs, collision resistant hash functions, and a semantically secure FHE scheme, the aforementioned VOS construction (described in detail in our online technical report [1]) is secure against a malicious server.*

*Proof.* (sketch.) Due to the proof of the semi-honest model compiler (Theorem 1), it suffices to show that a malicious server cannot deviate from the protocol without being detected — this is ensured by the security of the VC-RAM scheme.

## 4 Optimizations for Specific ORAM Schemes

### 4.1 Background on Path ORAM

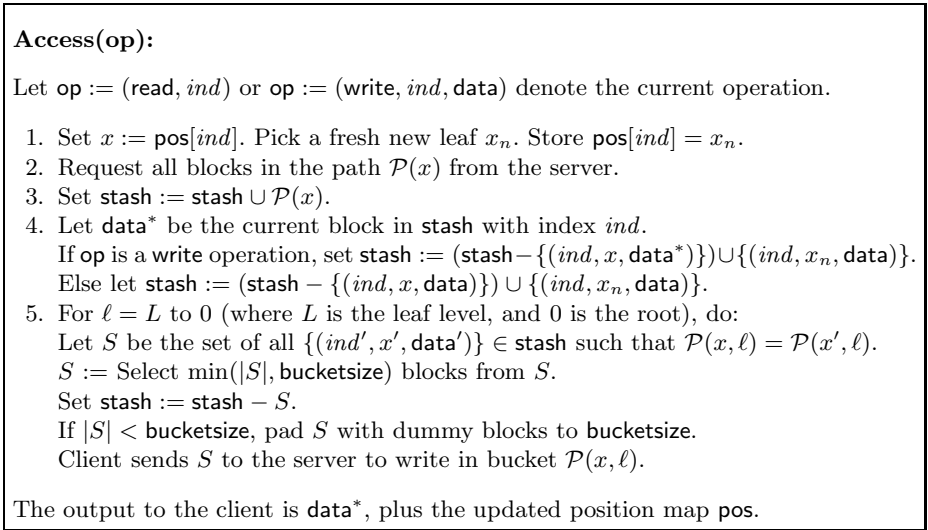
Stefanov et al. recently proposed the Path ORAM [33]. They formally prove that to achieve  $n^{-\alpha(n)}$  failure probability, the (recursive) Path ORAM construction achieves  $O(\alpha(n)\beta \log^2 n / \log \chi)$  client-side storage, and  $O(\beta \log^2 n / \log \chi)$  bandwidth cost —  $\chi$  is a term related to the block size where the block size  $\beta = \chi \log n$  bits. Specifically, to make the failure probability negligible, we can use any  $\alpha(n) := \omega(1)$ .

Of particular interest is the case when the block size is  $\Omega(\lambda)$  — in practical storage outsourcing applications, this is typically the case. Since  $n = \text{poly}(\lambda)$ , the number of recursions would be  $O(1)$ .

**Lemma 1 (Path ORAM [33]).** *For reasonably large block sizes  $\beta = \Omega(\lambda)$ , Path ORAM achieves bandwidth cost of  $O(\beta \log n)$ , a client-side storage of  $O(\alpha(n)\beta \log n)$ , and  $O(1)$  rounds, with a failure probability of  $n^{-\alpha(n)}$ . Specifically, to achieve negligible failure probability, it suffices to use any  $\alpha(n) := \omega(1)$ .*

We briefly introduce the Path ORAM algorithm below.

**Server Data Layout.** The blocks on the server are organized into a binary tree of height roughly  $\log n$ . Each node in the tree is a bucket of  $O(1)$  capacity. We use the notation  $\mathcal{P}(x)$  to denote the path from the leaf node  $x$  to the root node, containing all buckets on the path. Additionally,  $\mathcal{P}(x, \ell)$  denotes the bucket in  $\mathcal{P}(x)$  at level  $\ell$  in the tree.



**Fig. 2.** Access protocol for Path ORAM (non-recursive)

**Client Data Layout.** The client holds a position map where  $\text{pos}[\text{ind}]$  records the up-to-date designated leaf node for block  $\text{ind}$ . A block  $\text{ind}$ 's designated leaf node is  $x$  implies that the block resides somewhere along the path  $\mathcal{P}(x)$ .

The client also holds a small  $\text{stash}$  of size  $O(\alpha(n) \log n)$  for overflowing blocks, where any  $\alpha(n) := \omega(1)$  allows us to achieve negligible failure probability.

**Data Access.** To perform any data access operation  $\text{op}$ , where  $\text{op} := (\text{read}, \text{ind})$  or  $\text{op} := (\text{write}, \text{ind}, \text{data})$ , the client runs the **Access** protocol described in Figure 2. At Step 1, the block being read or written to is randomly remapped to a new leaf. At Step 2, the client requests a path of data blocks from the server. At Step 3, the local is merged with the data received from the server. At Step 4, the read/write operation is performed. At Step 5, the  $\text{stash}$  is written back into the tree, greedily pushing data blocks as close to the leaves as possible.

**Recursive Path ORAM.** The Path ORAM construction above requires the client to store a position map of  $O(n \log n)$  bits. However, the client can store the position map on the server in a smaller ORAM. This is called the recursive Path ORAM. Particularly, if the block size  $\beta := \Omega(\lambda)$ , and  $n = \text{poly}(\lambda)$ , then the depth of the recursion is constant.

## 4.2 Path VOS

We can use the generic compilation techniques described in Section 3 to compile Path ORAM to a VOS scheme — henceforth referred to as the Path VOS algorithm.

**Path VOS (non-recursive, semi-honest model)**

**Setup.** Given a memory array  $D$ , client lays out  $D$  into an initial ORAM-tree as in the Path ORAM algorithm, and creates an initial position map accordingly. The client encrypts the initial ORAM-tree under FHE, and an empty stash, and outsources both the FHE-encrypted ORAM-tree and stash to the server. The client keeps the position map locally.

**Access.** Let  $\text{op} := (\text{read}, \text{ind})$  or  $\text{op} := (\text{write}, \text{ind}, \text{data})$  denote the current operation.

- *Client:* Looks up its local position map  $x := \text{pos}[\text{ind}]$ . Pick a fresh random new leaf  $x'$ . Compute  $\overline{\text{op}} := \text{FHE.Enc}(\text{op})$ ,  $\overline{x'} := \text{FHE.Enc}(x')$ . Send  $(\overline{\text{op}}, x, \overline{x'})$  to the server.
- *Server:* Let  $\text{WritePath}(\mathcal{P}, \text{stash}, \text{op}, x')$  denote the circuit (Steps 3 to 5 in Figure 2) that on inputting a path  $\mathcal{P}$ , a stash  $\text{stash}$ , and the current operation  $\text{op}$ , returns the current value of the requested block  $\text{ind}$ , overwrites the block  $\text{ind}$ 's designated leaf tag to  $x'$ , overwrites the block  $\text{ind}$  with new  $\text{data}$  if this is a write operation, and write back blocks in  $\mathcal{P} \cup \text{stash}$  to the path  $\mathcal{P}$ , greedily packing them as close to the leaf as possible. The server homomorphically computes  $(\overline{\text{out}}, \overline{\mathcal{P}(x)}, \overline{\text{stash}}) \leftarrow \text{FHE.Eval}(\text{WritePath}(\overline{\mathcal{P}(x)}, \overline{\text{stash}}, \overline{\text{op}}, \overline{x'}))$ . The server sends to the client the FHE-encrypted result of the read  $\overline{\text{out}}$ .

**Fig. 3.** Path VOS (non-recursive, semi-honest model)

The semi-honest version of the Path VOS protocol is described in Figure 3. We can use the VC-RAM techniques described in Section 3.3 to compile the semi-honest protocol to one that is secure against a malicious server.

**Recursive Path VOS.** In the above (non-recursive) Path VOS protocol, the client needs to store a position map of size  $O(n \log n)$  bits. This client-side storage may be avoided by recursively outsourcing the position map to the server in a smaller VOS scheme. When the block size is  $\beta = \Omega(\lambda)$ , the depth of recursion is  $O(1)$ . The resulting recursive Path VOS scheme will therefore have  $O(1)$  roundtrips for each data access.

**Tailored VC-RAM Techniques for Path ORAM.** Based on the semi-honest protocol described above, and the VC-RAM techniques described in Section 3.3, we immediately obtain a Path VOS protocol with  $\tilde{O}(\beta \log^2 n) \text{poly}(\lambda)$  server computation per data access<sup>2</sup>, for block sizes  $\beta = \tilde{\Omega}(\lambda)$ . (The small increase in block size is due to FHE.)

We observe that by overlaying the Path ORAM tree structure on top of the Merkle tree, we can shave a logarithmic factor off the server computation. Recall that in our VC-RAM construction, the client maintains a Merkle-hash tree digest of the ORAM-memory outsourced to the server. To prove that any RAM computation is correct, the server computes a SNARK for a “verifier circuit” which verifies 1) that every memory access is correct (through the Merkle tree); and 2) every step of CPU computation is correct. In particular, the extra  $\log n$  factor comes from the cost.

<sup>2</sup> Throughout this paper, the notation  $\tilde{O}(f(n))$  hides  $\log(f(n))$  factors.

In the case of Path ORAM, since Path ORAM itself is a tree structure, we can overlay the Merkle tree on top of Path ORAM’s tree structure. In this way, when Path ORAM accesses a path from the root to a leaf, the underlying memory checking scheme can vouch for the correctness of the entire path with  $O(\log n)$  hashes. This can allow us to shave a  $O(\log n)$  factor off the server computation for Path VOS.

**Theorem 4 (Path VOS).** *Assume collision resistant hash functions, the ring LWE assumption with suitable parametrization [5, 17], and the  $q$ -PDH and  $q$ -PKE assumptions [15]. Let  $\alpha(n)$  denote any function such that  $\alpha(n) := \omega(1)$ .*

*There exists a secure VOS scheme for reasonably large block size  $\beta = \Omega(\lambda)$ , with  $O(\beta)$  bandwidth cost,  $\tilde{O}(\alpha(n) \cdot \beta \log n) \text{poly}(\lambda)$  server computation per data access,  $O(\beta n)$  server-side storage,  $O(\beta)$  client-side storage, and  $O(\beta + \lambda^2)$  client computation per data access. Furthermore, the failure probability is  $n^{-\alpha(n)}$ , i.e., negligible in  $n$  for any  $\alpha(n) := \omega(1)$ .*

Proof of security follows in a similar manner as the security proof for the generic compiler in the malicious model (Theorem 3).

### 4.3 The Hierarchical VOS

We propose a hierarchical VOS construction based on the Goodrich-Mitzenmacher ORAM (GM-ORAM) scheme [21] and its variants [25]. Although this hierarchical VOS construction achieves worse asymptotics than the Path VOS mentioned in the previous section, it is necessary later for our dynamic proofs of retrievability scheme — since the Path VOS scheme does not satisfy the next-read pattern hiding property (NRPH) proposed by Cash et al. [9]. (All of our VOS compilers are NRPH-preserving since they do not alter the sequence of accesses as dictated by the underlying ORAM.)

Although the basic idea is similar as before, to use FHE to outsource computation to the server, and use SNARK to enforce honest server behavior, we propose a “read/write (un)balancing” trick that allows us to reduce the bandwidth cost. The idea is that if we apply the generic ORAM-to-VOS compiler on the GM-ORAM scheme, reads will require more bandwidth than writes, since write is basically a homomorphic shuffling operation which the server can perform all on its own without interacting with the client. Therefore, we adjust the scheme to penalize writes while reducing the cost of reads. Note that in the traditional ORAM setting, writes cost more bandwidth, and that is why Kushilevitz et al. [25] propose a read/write balancing trick where they penalize reads to save on writes — our trick is the opposite of theirs since the read/write cost comparison is reverse in the VOS setting. Due to space constraints, we only give our main theorem for the Hierarchical VOS below, and defer the detailed construction to our online technical report [1].

**Theorem 5.** *Let  $g(n)$  denote some function on  $n$ . Assume collision resistant hash functions, the ring LWE assumption with suitable parametrization [5, 17], and the  $q$ -PDH and  $q$ -PKE assumptions [15]. Then, there exists a VOS scheme*

for a reasonably large block size  $\beta = \tilde{\Omega}(\lambda)$ , with  $O(\beta \log n / \log g(n))$  bandwidth cost, and  $\tilde{O}(\beta g(n) \log^3 n / \log g(n)) \text{poly}(\lambda)$  server computation (per data access), where  $n$  is the total number of blocks and  $\lambda$  is the security parameter.

The following table shows some interesting special cases of Theorem 5.

$g(n)$	server computation	bandwidth overhead
$n^\epsilon$ for constant $\epsilon < 1$	$\tilde{O}(\beta n^\epsilon \log^2 n) \text{poly}(\lambda)$	$O(\beta)$
$\log n$	$\tilde{O}(\beta \log^4 n / \log \log n) \text{poly}(\lambda)$	$O(\beta \log n / \log \log n)$
constant $c > 1$	$\tilde{O}(\beta \log^3 n) \text{poly}(\lambda)$	$O(\beta \log n)$

## 5 Applications: Efficient Dynamic Proofs of Retrievability

For applications such as Dynamic Proofs of Retrievability, and RAM-model secure multi-party computation where the party storing the data (or a share of the data) can perform computation, often, just directly replacing the ORAM scheme with a VOS scheme can reduce the asymptotic communication overhead.

We show how VOS can be useful in Dynamic Proofs of Retrievability, based on the results of Cash et al. [9]. We note that two recent results have yielded more practical dynamic PoR schemes [10, 31]. Our dynamic PoR description helps demonstrate why distinguishing between VOS and ORAM can aid theoretical understanding. For a practical implementation, the recent schemes by Shi et al. [31] and Chandran et al. [10] are recommended.

Recently Cash et al. [9] show how to leverage a blackbox ORAM scheme to construct a dynamic proof of retrievability (PoR) scheme with  $O(\beta \lambda \log^2 n)$  cost (both in terms of bandwidth and server computation) per data access. They require the underlying ORAM to have a special property which they call “*next-read pattern hiding*” (NRPH).

In the dynamic PoR scheme by Cash et al., they assume a passive server which does not perform any active computation. We observe that if we replaced the ORAM scheme in their construction with a VOS scheme (which also needs to satisfy the NRPH property), we would be able to obtain a dynamic PoR scheme (with server computation), which achieves smaller asymptotic bandwidth cost than Cash et al. [9].

The Path ORAM algorithm (and hence Path VOS too), however, does not satisfy the NRPH property, as pointed out by Cash et al. [9]. However, they showed that the GM-ORAM scheme and its variants indeed satisfy the NRPH property. Therefore, we rely on the hierarchical VOS described in Section 4.3 to build our dynamic PoR scheme.

**Theorem 6.** *Let  $g(n)$  denote some function on  $n$ . Assume collision resistant hash functions, the ring LWE assumption with suitable parametrization [5, 17], and the  $q$ -PDH and  $q$ -PKE assumptions [15]. Then, there exists a dynamic proof of retrievability scheme for reasonably large block size  $\beta = \tilde{\Omega}(\lambda)$ , with*

$O(\beta \log n / \log g(n))$  bandwidth cost and  $O(\beta g(n) \log^3 n / \log g(n)) \text{poly}(\lambda)$  server computation for each read operation;  $O(\beta \lambda \log n / \log g(n))$  bandwidth cost and  $O(\beta \lambda g(n) \log^3 n / \log g(n)) \text{poly}(\lambda)$  server computation for each write or audit operation; with  $O(\beta)$  client-storage and  $O(\beta n)$  server storage. In the above,  $n$  is the total number of blocks and  $\lambda$  is the security parameter.

Below are some interesting special cases of the above theorem. “R:” stands for read cost, and “W/A:” stands for write/audit cost.

$g(n)$	server computation	bandwidth overhead
$n^\epsilon$ for constant $\epsilon < 1$	R: $\tilde{O}(\beta n^\epsilon \log^2 n) \text{poly}(\lambda)$	R: $O(\beta)$
	W/A: $\tilde{O}(\beta \lambda n^\epsilon \log^2 n) \text{poly}(\lambda)$	W/A: $O(\beta \lambda)$
$\log n$	R: $\tilde{O}(\beta \log^4 n / \log \log n) \text{poly}(\lambda)$	R: $O(\beta \log n / \log \log n)$
	W/A: $\tilde{O}(\beta \lambda \log^4 n / \log \log n) \text{poly}(\lambda)$	W/A: $O(\beta \lambda \log n / \log \log n)$
constant $c > 1$	R: $\tilde{O}(\beta \log^3 n) \text{poly}(\lambda)$	R: $O(\beta \log n)$
	W/A: $\tilde{O}(\beta \lambda \log^3 n) \text{poly}(\lambda)$	W/A: $O(\beta \lambda \log n)$

In comparison to Cash et al. [9], using Verifiable Oblivious Storage (VOS), we can reduce the bandwidth cost to  $O(\beta \log n / \text{poly} \log \log n)$  for reads, and  $O(\beta \lambda \log n / \text{poly} \log \log n)$  for writes, with poly-logarithmic server computation. Furthermore, we can reduce the bandwidth cost to  $O(\beta)$  for reads, and  $O(\beta \lambda)$  for writes, with  $O(\beta n^\epsilon) \text{poly}(\lambda)$  amount of server computation for a constant  $\epsilon < 1$ .

**Other Applications.** Gordon et al. recently proposed to use ORAM to achieve amortized sublinear-time secure two-party computation [24]. In their setting, Alice’s input is a large database, and Bob repeatedly makes queries over the database. Alice wishes to protect the privacy of her database, while Bob wishes to protect the privacy of his query. Using ORAM, Gordon et al. show that the cost of securely querying the database can be sublinear when amortizing the ORAM setup cost over all future queries. Since both parties (each storing a share of the data) perform computation in this setting, we can simply replace the ORAM with VOS, and asymptotically, this gives savings in terms of bandwidth overhead.

## 6 Conclusion and Open Problems

This paper separates VOS from ORAM, and shows that VOS need not be subject to ORAM’s lower bounds, since it is a different model where server computation is allowed. The constructions proposed in this paper use general primitives such as FHE and SNARKs. An interesting open question is to see how to construct a practically efficient VOS scheme (potentially without FHE or SNARKs) that outperforms the best known ORAM in terms of bandwidth overhead. It would also be interesting to consider how to construct VOS schemes that are asymptotically more bandwidth efficient than ORAM from weaker assumptions, e.g., without SNARKs or non-falsifiable assumptions.



**Acknowledgments.** This research was funded by NSF under grant number CNS-1314857, by a Google Faculty Research Award, and by the US Army Research Laboratory and the UK Ministry of Defence under Agreement Number W911NF-06-3-0001. The views and conclusions contained herein are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the US Army Research Laboratory, the U.S. Government, the UK Ministry of Defense, or the UK Government. The US and UK Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

We thank Hubert Chan, Charalampos Papamanthou, Emil Stefanov, and Hong-Sheng Zhou for helpful discussions, and the anonymous reviewers for their insightful comments.

## References

1. Apon, D., Katz, J., Shi, E., Thiruvengadam, A.: Verifiable oblivious storage. Online technical report version of this paper (2013), <http://www.cs.umd.edu/~elaine/docs/vos.pdf>
2. Ben-Sasson, E., Chiesa, A., Genkin, D., Tromer, E.: Fast reductions from RAMs to delegatable succinct constraint satisfaction problems: extended abstract. In: ITCS (2013)
3. Bitansky, N., Canetti, R., Chiesa, A., Tromer, E.: From extractable collision resistance to succinct non-interactive arguments of knowledge, and back again. In: ITCS (2012)
4. Bitansky, N., Canetti, R., Chiesa, A., Tromer, E.: Recursive composition and bootstrapping for snarks and proof-carrying data. In: STOC (2013)
5. Brakerski, Z., Gentry, C., Vaikuntanathan, V.: (Leveled) fully homomorphic encryption without bootstrapping. In: ITCS (2012)
6. Brakerski, Z., Vaikuntanathan, V.: Efficient fully homomorphic encryption from (standard) LWE. In: FOCS (2011)
7. Braun, B., Feldman, A.J., Ren, Z., Setty, S., Blumberg, A.J., Walfish, M.: Verifying computations with state. In: SOSP (2013)
8. Canetti, R.: Security and composition of multiparty cryptographic protocols. *Journal of Cryptology* 13(1), 143–202 (2000)
9. Cash, D., K upc u, A., Wichs, D.: Dynamic proofs of retrievability via oblivious RAM. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 279–295. Springer, Heidelberg (2013)
10. Chandran, N., Kanukurthi, B., Ostrovsky, R.: Locally updatable and locally decodable codes. In: Lindell, Y. (ed.) TCC 2014. LNCS, vol. 8349, pp. 489–514. Springer, Heidelberg (2014)
11. Chor, B., Gilboa, N.: Computationally private information retrieval (extended abstract). In: STOC (1997)
12. Chor, B., Goldreich, O., Kushilevitz, E., Sudan, M.: Private information retrieval. In: IEEE Symposium on Foundations of Computer Science (FOCS), pp. 41–50 (1995)
13. Chung, K.-M., Pass, R.: A simple oram (2013), <https://eprint.iacr.org/2013/243.pdf>
14. Damg ard, I., Meldgaard, S., Nielsen, J.B.: Perfectly secure oblivious RAM without random oracles. In: Ishai, Y. (ed.) TCC 2011. LNCS, vol. 6597, pp. 144–163. Springer, Heidelberg (2011)

15. Gennaro, R., Gentry, C., Parno, B., Raykova, M.: Quadratic span programs and succinct NIZKs without PCPs. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 626–645. Springer, Heidelberg (2013)
16. Gentry, C., Goldman, K.A., Halevi, S., Julta, C., Raykova, M., Wichs, D.: Optimizing ORAM and using it efficiently for secure computation. In: De Cristofaro, E., Wright, M. (eds.) PETS 2013. LNCS, vol. 7981, pp. 1–18. Springer, Heidelberg (2013)
17. Gentry, C., Halevi, S., Smart, N.P.: Fully homomorphic encryption with polylog overhead. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 465–482. Springer, Heidelberg (2012)
18. Gentry, C., Ramzan, Z.: Single-database private information retrieval with constant communication rate. In: Caires, L., Italiano, G.F., Monteiro, L., Palamidessi, C., Yung, M. (eds.) ICALP 2005. LNCS, vol. 3580, pp. 803–815. Springer, Heidelberg (2005)
19. Goldreich, O.: Towards a theory of software protection and simulation by oblivious RAMs. In: STOC (1987)
20. Goldreich, O., Ostrovsky, R.: Software protection and simulation on oblivious RAMs. *J. ACM* (1996)
21. Goodrich, M.T., Mitzenmacher, M.: Privacy-preserving access of outsourced data via oblivious RAM simulation. In: Aceto, L., Henzinger, M., Sgall, J. (eds.) ICALP 2011, Part II. LNCS, vol. 6756, pp. 576–587. Springer, Heidelberg (2011)
22. Goodrich, M.T., Mitzenmacher, M., Ohrimenko, O., Tamassia, R.: Oblivious RAM simulation with efficient worst-case access overhead. In: CCSW (2011)
23. Goodrich, M.T., Mitzenmacher, M., Ohrimenko, O., Tamassia, R.: Privacy-preserving group data access via stateless oblivious RAM simulation. In: SODA (2012)
24. Gordon, S.D., Katz, J., Kolesnikov, V., Krell, F., Malkin, T., Raykova, M., Vahlis, Y.: Secure two-party computation in sublinear (amortized) time. In: ACM CCS (2012)
25. Kushilevitz, E., Lu, S., Ostrovsky, R.: On the (in)security of hash-based oblivious RAM and a new balancing scheme. In: SODA (2012)
26. Mayberry, T., Blass, E.-O., Chan, A.: Efficient private file retrieval by combining oram and pir (2013), <http://eprint.iacr.org/2013/086>
27. Ostrovsky, R.: Efficient computation on oblivious RAMs. In: STOC (1990)
28. Ostrovsky, R., Shoup, V.: Private information storage (extended abstract). In: STOC (1997)
29. Pinkas, B., Reinman, T.: Oblivious RAM revisited. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 502–519. Springer, Heidelberg (2010)
30. Shi, E., Chan, T.-H.H., Stefanov, E., Li, M.: Oblivious RAM with  $O((\log N)^3)$  worst-case cost. In: Lee, D.H., Wang, X. (eds.) ASIACRYPT 2011. LNCS, vol. 7073, pp. 197–214. Springer, Heidelberg (2011)
31. Shi, E., Stefanov, E., Papamanthou, C.: Practical dynamic proofs of retrievability. In: ACM CCS (2013)
32. Stefanov, E., Shi, E., Song, D.: Towards practical oblivious RAM. In: NDSS (2012)
33. Stefanov, E., van Dijk, M., Shi, E., Fletcher, C., Ren, L., Yu, X., Devadas, S.: Path oram: An extremely simple oblivious ram protocol. In: ACM CCS (2013)
34. Williams, P., Sion, R.: Usable PIR. In: NDSS (2008)
35. Williams, P., Sion, R.: Single round access privacy on outsourced storage. In: CCS (2012)
36. Williams, P., Sion, R., Carbunar, B.: Building castles out of mud: practical access pattern privacy and correctness on untrusted storage. In: ACM CCS (2008)