

S³ORAM: A Computation-Efficient and Constant Client Bandwidth Blowup ORAM with Shamir Secret Sharing*

Thang Hoang[†] Ceyhun D. Ozkaptan[†] Attila A. Yavuz[†] Jorge Guajardo[‡]
Tam Nguyen[†]

Abstract

Oblivious Random Access Machine (ORAM) enables a client to access her data without leaking her access patterns. Existing client-efficient ORAMs either achieve $\mathcal{O}(\log N)$ client-server communication blowup without heavy computation, or $\mathcal{O}(1)$ blowup but with expensive homomorphic encryptions. It has been shown that $\mathcal{O}(\log N)$ bandwidth blowup might not be practical for certain applications, while schemes with $\mathcal{O}(1)$ communication blowup incur even more delay due to costly homomorphic operations.

In this paper, we propose a new distributed ORAM scheme referred to as *Shamir Secret Sharing ORAM (S³ORAM)*, which achieves $\mathcal{O}(1)$ client-server bandwidth blowup and $\mathcal{O}(1)$ blocks of client storage *without* relying on costly partial homomorphic encryptions. S³ORAM harnesses Shamir Secret Sharing, tree-based ORAM structure and a secure multiparty multiplication protocol to eliminate costly homomorphic operations and, therefore, achieves $\mathcal{O}(1)$ client-server bandwidth blowup with a high computational efficiency. We conducted comprehensive experiments to assess the performance of S³ORAM and its counterparts on actual cloud environments, and showed that S³ORAM achieves *three orders of magnitude lower end-to-end delay* compared to alternatives with $\mathcal{O}(1)$ client communication blowup (Onion-ORAM), while it is *one order of magnitude faster* than Path-ORAM for a network with a moderate bandwidth quality. We have released the implementation of S³ORAM for further improvement and adaptation.

1 Introduction

Oblivious Random Access Memory (ORAM) [21] allows Alice to access her data outsourced to a cloud without leaking to the server which data blocks have been accessed. Despite recent progress, it has been shown that existing ORAM designs are costly due to their high communication and/or computation overhead [1, 29, 30, 6, 22].

The objective of this paper is to create an efficient ORAM scheme that simultaneously achieves (i) a low client communication overhead (i.e., $\mathcal{O}(1)$ client bandwidth blowup), (2) low computational overhead by avoiding costly partial homomorphic encryptions, and (iii) low client storage (i.e., $\mathcal{O}(1)$ block storage).

*Full version of the paper appeared in ACM CCS 2017. This work is supported by the NSF CAREER Award CNS-1652389 and an unrestricted gift from Robert Bosch LLC.

[†]Oregon State University, {hoangmin, ozkaptac, attila.yavuz, nguyeta4}@oregonstate.edu

[‡]Robert Bosch Research and Technology Center, Jorge.GuajardoMerchan@us.bosch.com

Table 1: Asymptotic and Experimental Performance Comparison of S³ORAM and its Counterparts.

Scheme	Bandwidth Blowup [†]		Server Computation	Client Block Storage [‡]	End-to-end Delay (s)	# of servers
	Client-server	Server-server				
Path-ORAM [40]	$\mathcal{O}(\log N)$	-	-	$\mathcal{O}(\lambda)$	20.3	1
Ring-ORAM [33]	$\mathcal{O}(\log N)$	-	XOR	$\mathcal{O}(\lambda)$	13.2	1
Onion-ORAM [12]	$\mathcal{O}(1)$	-	Additively HE [10]	$\mathcal{O}(1)$	10^4	1
S ³ ORAM	$\mathcal{O}(1)$	$\mathcal{O}(\log N)$	Secure addition and multiplication of SSS values	$\mathcal{O}(1)$	2.5	3

This table presents the performance result of selected ORAM schemes with 40GB database containing 128-KB blocks and the network setting with a download and upload throughput of 27 and 6 Mbps, respectively. We refer the reader to Section 5 for the details of our experiments. λ denotes the security parameter.

[†] *Bandwidth blowup* is defined as the ratio between the communication introduced by ORAM and the base case where the access pattern is not hidden.

[‡] *Client block storage* is defined as the number of data blocks that need to be temporarily stored at the client. It is equivalent to the stash component used in [40, 33].

1.1 Research Gap and Problem Statement

ORAM with $\mathcal{O}(\log N)$ bandwidth blowup. Stefanov et al. in [40] proposed Path-ORAM scheme that achieves the optimal lower bound of $\Omega(\log N)$ communication blowup under $\mathcal{O}(1)$ blocks of client storage [42, 7]. However, Path-ORAM has been shown to be costly for certain applications [30, 36, 6] due to the transmission cost of $\mathcal{O}(\log N)$ blocks per access request. The client communication blowup can be reduced by introducing computation at the server side [12, 3]. Ring-ORAM [33] improved the communication efficiency of Path-ORAM by approximately 2.5 times by allowing the server to perform XOR operations. However, it still requires $\mathcal{O}(\log N)$ communication blowup. Other ORAM schemes (e.g., [26, 11]) used single-server PIR (e.g., [41]) to reduce the communication overhead. However, they still require $\mathcal{O}(\log N)$ bandwidth blowup, and also incur significant computation cost due to single-server PIR techniques (e.g., [41]).

ORAM with $\mathcal{O}(1)$ bandwidth blowup. Recent ORAM schemes (e.g., Onion-ORAM [12], Bucket-ORAM [14], and [3]) rely on fully or partial Homomorphic Encryption (HE) (e.g., [31]) to achieve $\mathcal{O}(1)$ bandwidth blowup with $\mathcal{O}(1)$ blocks of client storage. However, these ORAMs introduce an extremely high end-to-end delay due to heavy computations incurred by HE operations. For instance, it has been shown in [28, 1] that HE operations take much longer execution time than the use of ORAMs with $\mathcal{O}(\log N)$ communication blowup (e.g., Path-ORAM [40]). Therefore, distributed ORAM schemes have been proposed to achieve a better computation performance trade-off.

Distributed ORAM without Costly HE Operations. Stefanov et al. proposed an ORAM scheme [37] that leverages two non-colluding computation-capable servers to achieve $\mathcal{O}(1)$ client-server bandwidth blowup with $\mathcal{O}(\log N)$ server-server communication blowup. However, it requires $\mathcal{O}(\sqrt{N})$ blocks of client storage due to its underlying ORAM primitive (i.e., Partition-ORAM [39]), which is extremely costly for memory-limited clients.

Some ORAM schemes (e.g., [27, 1]) attempted to use multi-server PIR (e.g., [9]) to decrease the communication overhead under $\mathcal{O}(1)$ blocks of client storage without using costly homomorphic encryption. Abraham et al. in [1] showed an asymptotically tight sub-logarithmic communication blowup bound of $\Omega(\log_{cD} N)$ for composing ORAM with PIR, where c, D are the numbers of blocks stored by the client and performed by PIR operations, respectively. Therefore, although the CHf-ORAM scheme in [27] claimed to achieve $\mathcal{O}(1)$ bandwidth blowup with $\mathcal{O}(1)$ blocks of client storage, it has been shown to violate this bound with two concrete attacks in [1]. To the best of our knowledge, there is no secure distributed ORAM scheme that can

achieve an $\mathcal{O}(1)$ client-server bandwidth blowup with $\mathcal{O}(1)$ blocks of client storage overhead.

1.2 Our Contribution

We developed a new distributed ORAM that we refer to as *Shamir Secret Shared ORAM* (S^3 ORAM). Below, we first present our main idea followed by desirable properties and contributions of our scheme.

Main idea. Although Onion-ORAM [12] is considered a theoretical construction due to its costly partial HE operations, it offers an elegant eviction strategy, which is useful to achieve $\mathcal{O}(1)$ client bandwidth blowup with $\mathcal{O}(1)$ blocks of client storage. Our main idea is to harness the “homomorphic” properties of Shamir Secret Sharing (SSS) along with a secure multi-party multiplication protocol to perform eviction operations in the line of Onion-ORAM, but in a significantly more efficient and practical manner. By doing this, S^3 ORAM inherits all desirable properties of Onion-ORAM *without* the costly homomorphic operations and, thus, requiring only a lightweight computation and suitability for larger block sizes. Table 1 outlines a high-level comparison of S^3 ORAM and its counterparts.

Desirable properties and contributions. We summarize the desirable properties of S^3 ORAM and our contributions as follows:

- Low client-server communication: S^3 ORAM achieves $\mathcal{O}(1)$ client bandwidth blowup, compared with $\mathcal{O}(\log N)$ of Path-ORAM [40] and Ring-ORAM [33] (with a fixed number of servers). Moreover, S^3 ORAM features smaller block sizes (i.e., $\Omega(\log N)$) than those of other ORAM schemes with $\mathcal{O}(1)$ communication blowup, which require fully or partial HE operations (e.g., $\Omega(\log^5 N)$ in Onion-ORAM [12], $\Omega(\log^6 N)$ in Bucket-ORAM [14]).
- Low server computation: In S^3 ORAM, servers only perform lightweight modular additions and multiplications, which are much more efficient than partial HE (e.g., [10]) operations. In particular, we show in Section 5 that, the server computation of S^3 ORAM is three orders of magnitude faster than that of Onion-ORAM.
- Low client computation: In S^3 ORAM, the client only performs lightweight computations for retrieval and eviction operations. Thus, it is more efficient than Onion-ORAM which requires a number of partial HE operations. For example, S^3 ORAM requires only a few milliseconds compared to minutes of Onion-ORAM to generate encrypted select queries (see Section 5). Moreover, since blocks in S^3 ORAM are single encrypted, the decryption is less costly and, therefore, faster than other ORAMs (e.g., [12, 37]) whose blocks are onion-encrypted.
- Low end-to-end delay: S^3 ORAM is approximately three orders of magnitude faster than Onion-ORAM, while it is one order of magnitude faster than Path-ORAM in networks with *moderate* bandwidth (e.g., < 240 Mbps).
- Compact client storage: S^3 ORAM features $\mathcal{O}(1)$ blocks of client storage, compared to $\mathcal{O}(\lambda)$ in Path-ORAM and Ring-ORAM, respectively, and $\mathcal{O}(\sqrt{N})$ of Stefanov et al. in [39].
- High security: S^3 ORAM relies on Shamir Secret Sharing and a secure multi-party multiplication protocol, and therefore, it offers information-theoretic security.

- *Full-fledged implementation and experiments:* We implemented S³ORAM and evaluated its performance in an actual cloud environment (i.e., Amazon EC2). The detailed experiments in Section 5 showed that S³ORAM is efficient in practice, and it can even be deployed on mobile devices with limited computation capacity and network connection. We also release the source code of S³ORAM for public use and wide adaptation.

S³ORAM does not rely on the direct composition of PIR and ORAM, and it requires servers to communicate with each other to execute a secure multi-party multiplication protocol with the communication blowup of $\mathcal{O}(\log N)$. Therefore, S³ORAM does not violate the asymptotic communication bound of Abraham et al. [1]. Note that a high bandwidth is available for inter-cloud communications via dedicated connections [23]. Hence, the inter-cloud communication of S³ORAM has a minimal impact on the end-to-end client-server delay as detailed in Section 5.

On the other hand, in many practical scenarios, it may not be possible to guarantee a reliable high-bandwidth communication link between the client and the servers. This is particularly true in the case of home networks and mobile devices with wireless network connectivity (e.g., Wi-Fi, LTE). Therefore, following recent work, one of the main goals of this work is to minimize the client communication overhead while at the same time requiring a low computational overhead at the client and the server sides. In Section 5, we demonstrate the advantages of S³ORAM over its counterparts with $\mathcal{O}(\log N)$ and $\mathcal{O}(1)$ communication blowup for such moderate bandwidth network settings. It turns out, that the advantages of S³ORAM over its counterparts with $\mathcal{O}(1)$ bandwidth blowup such as Onion-ORAM are significant due to its computational efficiency. This efficiency is obtained, however, at the cost of requiring multiple servers (at least three) in the distributed setting.

2 Preliminaries and Building Blocks

Notation. $x \stackrel{\$}{\leftarrow} \mathcal{S}$ denotes that x is randomly and uniformly selected from set \mathcal{S} . $|\mathcal{S}|$ denotes the cardinality of set \mathcal{S} . $|x|$ denotes the size of variable x . For any integer l , $(x_1, \dots, x_l) \stackrel{\$}{\leftarrow} \mathcal{S}$ denotes $(x_1 \stackrel{\$}{\leftarrow} \mathcal{S}, \dots, x_l \stackrel{\$}{\leftarrow} \mathcal{S})$. We denote a finite field as \mathbb{F}_p , where p is a prime. Given \mathbf{u} and \mathbf{v} as vectors with the same length, $\mathbf{u} \cdot \mathbf{v}$ denotes the inner product of \mathbf{u} and \mathbf{v} . Given an n -dimensional vector \mathbf{u} and a matrix \mathbf{I} of size $n \times m$, $\mathbf{v} = \mathbf{u} \cdot \mathbf{I}$ denotes the matrix product of \mathbf{u} and \mathbf{I} resulting in an m -dimensional vector \mathbf{v} . $\mathbf{I}[i, *]$ denotes accessing data of row i of matrix \mathbf{I} .

2.1 Model of Computation

Following the literature in distributed secure computation (e.g., [5, 19]), we assume a synchronous network which consists of a client and $\ell \geq 2t + 1$ semi-honest servers $\mathcal{S} = \{\mathcal{S}_1, \dots, \mathcal{S}_\ell\}$. It is also assumed that the channels between all the players are pairwise-secure, i.e., no player can tamper, read, or modify the contents of the communication channel of other players. We assume that all parties behave in an “honest-but-curious” manner in which parties always send messages as expected but try to learn as much as possible from the shared information received or observed. Notice that in this paper, we *do not* allow parties to provide malicious inputs, i.e., parties are not allowed to behave in a Byzantine manner.

A protocol is t -private [5] (see also e.g., [19] for similar definitions in the context of distributed PIR) if any set of at most t parties cannot compute after an execution of a protocol more than they could compute individually from their set of private inputs and outputs. Alternatively,

the parties have not “learned” anything. Our protocols in general, offer information-theoretic guarantees, unless something is said explicitly to the contrary. This implies that our solutions are secure against computationally unbounded adversaries. As it is standard, we require that all computations by the servers and client be polynomial time and efficient. Finally notice that in this paper, not only is the interaction between the servers and client performed in such a way that information-theoretic security is guaranteed but also the database being accessed is shared among the servers in a way that no coalition of up to t servers can find anything about the database contents (also in an information-theoretic manner).

2.2 Shamir Secret Sharing

We recall (t, ℓ) -threshold Shamir’s Secret Sharing scheme [34] which comprises two algorithms SSS.Create and SSS.Recover as presented in Algorithm 1. To share a secret $\alpha \in \mathbb{F}_p$ among ℓ parties, a dealer generates a random polynomial f where $f(0) = \alpha$ and evaluates $f(x_i)$ for party \mathcal{P}_i for $1 \leq i \leq \ell$, where $x_i \in \mathbb{F}_p \setminus \{0\}$ is a deterministic non-zero element of \mathbb{F}_p that uniquely identifies party \mathcal{P}_i and it is considered public information (SSS.Create algorithm). $f(x_i)$ is referred to as the share of party \mathcal{P}_i , and it is denoted by $[[\alpha]]_i$. To reconstruct the secret α , the shares of at least $t + 1$ parties have to be combined via Lagrange interpolation (SSS.Recover algorithm).

Algorithm 1 Shamir Secret Sharing (SSS) scheme [34]

$([[\alpha]]_1, \dots, [[\alpha]]_\ell) \leftarrow \text{SSS.Create}(\alpha, t)$: Create t -private shares of α

- 1: $(a_1, \dots, a_t) \xleftarrow{\$} \mathbb{F}_p$
 - 2: **for** $i = 1, \dots, \ell$ **do**
 - 3: $[[\alpha]]_i \leftarrow \alpha + \sum_{j=1}^t a_j \cdot x_i^j$
 - 4: **return** $([[\alpha]]_1, \dots, [[\alpha]]_\ell)$
-

$\alpha \leftarrow \text{SSS.Recover}(\mathcal{A}, t)$: Recover the value α from $\geq t + 1$ shares

- 1: Randomly pick $t + 1 \leq \ell$ shares $\{[[\alpha]]_{x_1}, \dots, [[\alpha]]_{x_{t+1}}\}$ in \mathcal{A}
 - 2: $g(x) \leftarrow \text{LagrangeInterpolation}(\{(x_i, [[\alpha]]_{x_i})\}_{i=1}^{t+1})$
 - 3: $\alpha \leftarrow g(0)$
 - 4: **return** α
-

We extend the notion of secret share for a value into the share for a vector in the natural way as follows: Given a vector $\mathbf{v} = (v_1, \dots, v_n)$, $[[\mathbf{v}]]_i = ([[v_1]]_i, \dots, [[v_n]]_i)$ indicates the share of \mathbf{v} for party \mathcal{P}_i , which is a vector whose elements are the shares of the elements in \mathbf{v} . Similarly, given a matrix \mathbf{I} , $[[\mathbf{I}]]$ denotes the share of \mathbf{I} , which is also a matrix with each cell $[[\mathbf{I}[i, j]]]$ being the share of the cell $\mathbf{I}[i, j]$. In some cases, to ease readability, we drop the subscript i , when the party is understood from the context.

Shamir [34] showed that SSS is information-theoretic secure and t -private in the sense that no set of t or less shares reveals any information about the secret. More precisely, $\forall m, m' \in \mathbb{F}_p, \forall \mathcal{I} \subseteq \{1, \dots, \ell\}$ s.t. $|\mathcal{I}| \leq t$ and for any set $\mathcal{A} = \{a_1, \dots, a_{|\mathcal{I}|}\}$ where $a_i \in \mathbb{F}_p$, the probability distributions of $\{s_{i \in \mathcal{I}} : (s_1, \dots, s_\ell) \leftarrow \text{SSS.Create}(m, t)\}$ and $\{s'_{i \in \mathcal{I}} : (s'_1, \dots, s'_\ell) \leftarrow \text{SSS.Create}(m', t)\}$ are identical and uniform:

$$\Pr(\{s_{i \in \mathcal{I}}\} = \mathcal{A}) = \Pr(\{s'_{i \in \mathcal{I}}\} = \mathcal{A}).$$

Ben-Or et al. [5] showed that SSS can be used to obtain t -private protocols. Lemma 1 summarizes the homomorphic properties of SSS and it was first described in [5].

Lemma 1 (SSS homomorphic properties [5]). *Let $[[\alpha]]_i^{(t)}$ be the Shamir share of value $\alpha \in \mathbb{F}_p$ with privacy level t for \mathcal{P}_i . SSS offers additively and multiplicatively homomorphic properties:*

- *Addition of two shares*

$$[[\alpha_1]]_i^{(t)} + [[\alpha_2]]_i^{(t)} = [[\alpha_1 + \alpha_2]]_i^{(t)}. \quad (1)$$

- *Multiplication w.r.t a scalar $c \in \mathbb{F}_p$*

$$c \cdot [[\alpha]]_i^{(t)} = [[c \cdot \alpha]]_i^{(t)}. \quad (2)$$

- *Partial share multiplication*

$$[[\alpha_1]]_i^{(t)} \cdot [[\alpha_2]]_i^{(t)} = [[\alpha_1 \cdot \alpha_2]]_i^{(2t)}. \quad (3)$$

The two-share partial multiplication (3) in Lemma 1 results in a share of $\alpha_1 \cdot \alpha_2$ which is t -private and represented by a $2t$ -degree polynomial. It was first observed in [5] that the resulting polynomial is not uniformly distributed. In order to achieve the uniform distribution and computation consistency over $[[\alpha_1 \cdot \alpha_2]]$, it is required to reduce the degree of the polynomial representation of $[[\alpha_1 \cdot \alpha_2]]$ from $2t$ to t and re-share the polynomial. This multiplication operation with degree reduction can be achieved via a secure multiplication protocol shown in the following section¹.

2.3 Secure Multi-party Multiplication

Gennaro et al. [17] presents a Secure Multiplication Protocol (SMP) for two Shamir secret-shared values among multiple parties. Given $\alpha_1, \alpha_2 \in \mathbb{F}_p$ shared by (t, ℓ) -threshold SSS as $[[\alpha_1]]_i^{(t)}$ and $[[\alpha_2]]_i^{(t)}$ for $1 \leq i \leq \ell$ respectively, $2t + 1$ parties \mathcal{P}_i among ℓ parties would like to compute the multiplication of α_1, α_2 without revealing the value of α_1 and α_2 . The protocol requires a Vandermonde matrix $\mathbf{V}_{\{x_i\}}$ of size $(2t + 1) \times (2t + 1)$ having the structure as follows:

$$\mathbf{V}_{\{x_1, \dots, x_{2t+1}\}} = \begin{bmatrix} x_1^0 & x_1^1 & \dots & x_1^{2t} \\ x_2^0 & x_2^1 & \dots & x_2^{2t} \\ \vdots & \vdots & \ddots & \vdots \\ x_{2t+1}^0 & x_{2t+1}^1 & \dots & x_{2t+1}^{2t} \end{bmatrix}, \quad (4)$$

where $x_i \in \mathbb{F}_p$ are unique identifiers of participating \mathcal{P}_i . We refer to \mathbf{V}^{-1} as the inverse of Vandermonde matrix. Each \mathcal{P}_i locally multiplies $[[\alpha_1]]_i^{(t)}$ and $[[\alpha_2]]_i^{(t)}$ resulting in $[[\alpha_1 \cdot \alpha_2]]_i^{(2t)}$, and creates shares of $[[\alpha_1 \cdot \alpha_2]]_i^{(2t)}$ by a new random polynomial of degree t for $2t + 1$ parties and distributes them to other $2t$ parties. Finally, each party locally performs the inner product between the received shares and $\mathbf{V}_{\{x_i\}}^{-1}[1, *]$ to obtain a new share of $\alpha_1 \cdot \alpha_2$ which is now represented by a polynomial of degree t as $[[\alpha_1 \cdot \alpha_2]]_i^{(t)}$. Protocol 1 presents this multiplication protocol.

¹Benor et al. [5] proposed a secure multiplication protocol, however the protocol of Gennaro et al. [17] is more efficient and thus, is the subject of Section 2.3.

Protocol 1 SMP Protocol [17]

Input: \mathcal{P}_i owns $\llbracket \alpha_1 \rrbracket_i^{(t)}$, $\llbracket \alpha_2 \rrbracket_i^{(t)}$ and wants to compute $\llbracket \alpha_1 \cdot \alpha_2 \rrbracket_i^{(t)}$

Output: Each \mathcal{P}_i obtains $\llbracket \beta \rrbracket_i^{(t)}$, where $\beta = \alpha_1 \cdot \alpha_2$

- 1: **for** each $\mathcal{P}_i \in \{\mathcal{P}_1, \dots, \mathcal{P}_{2t+1}\}$ **do**
 - 2: $\llbracket \beta \rrbracket_i^{(2t)} \leftarrow \llbracket \alpha_1 \rrbracket_i^{(t)} \cdot \llbracket \alpha_2 \rrbracket_i^{(t)}$
 - 3: $(\llbracket \beta \rrbracket_j^{(t)})_{j=1}^\ell \leftarrow \text{SSS.Create}(\llbracket \beta \rrbracket_i^{(2t)}, t)$
 - 4: Distribute $\llbracket \beta \rrbracket_j^{(t)}$ to all $\mathcal{P}_j \in \{\mathcal{P}_1, \dots, \mathcal{P}_{2t+1}\} \setminus \mathcal{P}_i$
 - 5: **for** each $\mathcal{P}_i \in \{\mathcal{P}_1, \dots, \mathcal{P}_{2t+1}\}$ **do**
 - 6: $\llbracket \beta \rrbracket_i^{(t)} \leftarrow \sum_{j=1}^{2t+1} \mathbf{V}^{-1}[1, j] \cdot \llbracket \beta \rrbracket_j^{(t)}$
-

Lemma 2 (SMP protocol Privacy [17]). *The SMP protocol in [17] (denoted as \star operator) offers homomorphic property for full multiplication between two SSS-shares whose result is t -private as:*

$$\llbracket \alpha_1 \cdot \alpha_2 \rrbracket_i^{(t)} = \llbracket \alpha_1 \rrbracket_i^{(t)} \star \llbracket \alpha_2 \rrbracket_i^{(t)} \quad (5)$$

2.4 Multi-server Private Information Retrieval

Private Information Retrieval (PIR) enables retrieval of a data item from an (unencrypted) public database without revealing which item being fetched. We follow the presentation of [19, 4] as follows.

Definition 1 (multi-server PIR [19, 4, 9]). *Let $\mathbf{b} = (b_1, \dots, b_n)$ be a database consisting of n items being stored in ℓ servers. A multi-server PIR protocol consists of three algorithms: PIR.CreateQuery, PIR.Retrieve and PIR.Reconstruct. Given an item b_i in \mathbf{b} to be retrieved, the client creates queries $(e_1, \dots, e_\ell) \leftarrow \text{PIR.CreateQuery}(i)$ and distributes e_j to server \mathcal{S}_j . Each server responds with an answer $a_j \leftarrow \text{PIR.Retrieve}(e_j, \mathbf{b})$. Upon receiving ℓ answers, the client computes the value of item b_i by invoking the reconstruction algorithm $b \leftarrow \text{PIR.Reconstruct}(a_1, \dots, a_\ell)$.*

Security of the protocol is defined in terms of *correctness* and *privacy*. A multi-server PIR protocol is *correct* if the client computes the correct value of b from any ℓ answers via PIR.Reconstruct algorithm with probability 1. The concept of t -privacy for protocols is applied naturally to the PIR setting and follows directly from the t -privacy of SSS and the fact that among the servers they only have access to t shares of the query vector [19].

2.5 Multi-server ORAM Security

We now define the security of multi-server ORAM in the semi-honest setting proposed in [1] as a straightforward extension of the definition in [1] to the multi-server setting.

Definition 2 (Multi-server ORAM with server computation). *Let $\mathbf{x} = ((\text{op}_1, \text{id}_1, \text{data}_1), \dots, (\text{op}_q, \text{id}_q, \text{data}_q))$ be a data request sequence of length q , where $\text{op}_j \in \{\text{Read}, \text{Write}\}$, id_j is the identifier to be read/written and data_j is the data identified by id_j to be read/written. Let $\text{ORAM}_j(\mathbf{x})$ represent the ORAM client's sequence of interactions with the server \mathcal{S}_i given a data request sequence \mathbf{x} .*

Correctness. *A multi-server ORAM is correct if for any access sequence \mathbf{x} , $\{\text{ORAM}_1(\mathbf{x}), \dots, \text{ORAM}_\ell(\mathbf{x})\}$ returns data consistent with \mathbf{x} except with a negligible probability.*

***t*-security.** A multi-server ORAM is *t*-secure if $\forall \mathcal{I} \subseteq \{1, \dots, \ell\}$ such that $|\mathcal{I}| \leq t$, for any two data access sequences \mathbf{x}, \mathbf{y} with $|\mathbf{x}| = |\mathbf{y}|$, their corresponding transcripts $\{ORAM_{i \in \mathcal{I}}(\mathbf{x})\}$ and $\{ORAM_{i \in \mathcal{I}}(\mathbf{y})\}$ observed by a coalition of up to *t* servers $\{\mathcal{S}_{i \in \mathcal{I}}\}$ are statistically/ computationally indistinguishable.

3 The Proposed S³ORAM Scheme

S³ORAM follows the typical procedure of tree-based ORAMs [35]. Specifically, given a block to be accessed, the client first retrieves it from the outsourced ORAM structure via a secure retrieval operation. The retrieved block is then assigned to a random path, and written back to the root bucket. Finally, an eviction operation is performed in order to percolate data blocks to lower levels in the ORAM structure. *The intuition behind S³ORAM access protocol is as follows:* (1) to integrate SSS with a multi-server PIR protocol to perform a private retrieval operation with some homomorphic properties; and (2) to leverage these homomorphic properties of SSS and a SMP protocol to perform block permutation and to preserve *t*-privacy level of ORAM structure in the eviction phase, without relying on costly partial HE operations. In Table 2, we outline the notation used in the S³ORAM scheme and throughout the rest of the paper.

3.1 S³ORAM Data Structure

The S³ORAM structure is a balanced binary tree denoted as \mathbb{T} with a height of H . \mathbb{T} can store up to $N \leq A \cdot 2^{H-1}$ data blocks b_i , where constant A is the eviction frequency. A node in \mathbb{T} is a *bucket* with Z slots, which can store up to Z real blocks. We index buckets in \mathbb{T} according to the top-to-bottom, left-to-right order. Hence, leaf buckets are indexed in $[2^H, 2^{H+1})$ as exemplified in Figure 1. $\mathbb{T}[i]$ and $\mathbb{T}[i, j]$ denote an access operation to the bucket with index i , and to the slot j ($1 \leq j \leq Z$) of the i -th bucket, respectively. S³ORAM has a position map $\text{pm} := (\text{id}, \langle \text{plD}, \text{plDx} \rangle)$ to store the position of real blocks in \mathbb{T} , where $2^H \leq \text{plD} < 2^{H+1}$ denotes the path assigned for the block id , and $1 \leq \text{plDx} \leq Z \cdot (H + 1)$ denotes the block’s index in the path plD . We present the construction of S³ORAM data structure in Algorithm 3. Given a database DB organized into N B -bit blocks and an ORAM tree \mathbb{T} as the input, the $\text{S}^3\text{ORAM.Setup}$ creates the shares of \mathbb{T} as the output for ℓ servers as follows. First, the client initializes every slot in each bucket of \mathbb{T} with a 0’s string of length B (step 2). The client organizes all data blocks into \mathbb{T} , wherein

Table 2: Notations.

Symbol	Description
$\mathbb{T}, \mathbb{T}[i]$	S ³ ORAM tree structure and the bucket indexing i .
B, b, c	Block size, block and block chunk, resp.
N, m	Number of blocks and number of chunks in a block.
H	Height of the S ³ ORAM tree \mathbb{T} .
A	Eviction frequency.
Z	Bucket size.
pm	Position map.
$(\text{plD}, \text{plDx}) \leftarrow \text{pm}[\text{id}]$	Precise location (path ID & path index) of a block id .
$\mathcal{I} \leftarrow \mathcal{P}(\text{plD})$	(Ordered) indexes of buckets residing in path plD .
$i \leftarrow \mathcal{P}(\text{plD}, h)$	Index of the bucket on path plD at level- h .
n_e, n_r	Current number of eviction and retrieval operations, resp.

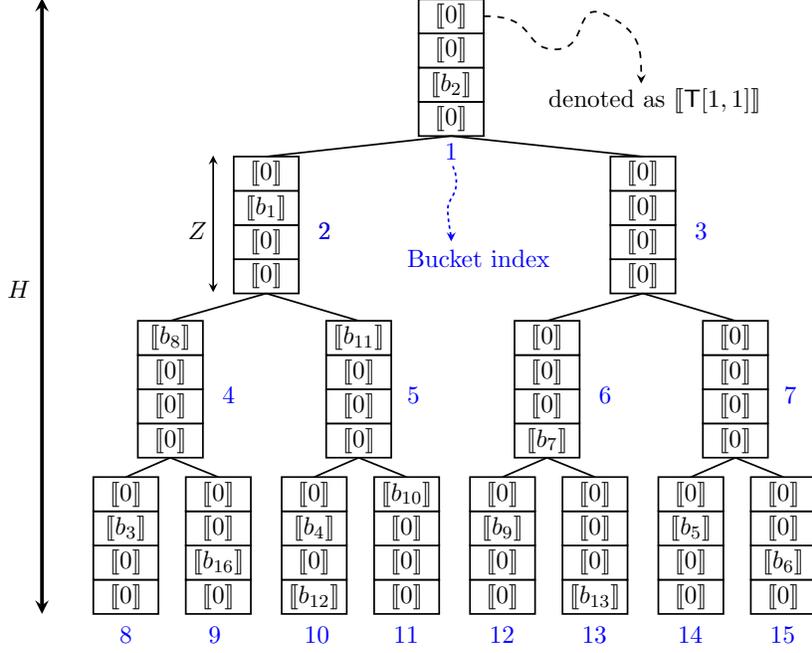


Figure 1: S^3 ORAM tree structure.

each b_i is independently assigned to a random leaf bucket of T . Notice that B can be larger than $\lceil \log_2 p \rceil$ and therefore, it might not be suitable for arithmetic computation over \mathbb{F}_p . To address this, the client splits the data in each slot of T into equal-sized chunks $c_j \in \mathbb{F}_p$ (step 11)². Finally, the client creates shares of T via SSS.Create for each chunk in each slot in T (step 13). The distributed S^3 ORAM structure consists of ℓ shares of T as $\{\llbracket T \rrbracket_1, \dots, \llbracket T \rrbracket_\ell\}$. Figure 1

²We assume implicitly that we choose an appropriate prime p such that every string c_j when interpreted as an element of \mathbb{F}_p is less than p .

Algorithm 3 $(\llbracket T \rrbracket_1, \dots, \llbracket T \rrbracket_\ell) \leftarrow S^3\text{ORAM.Setup}(\text{DB}, T)$

- 1: Organize DB into blocks (b_1, \dots, b_N) with IDs $(\text{id}_1, \dots, \text{id}_N)$
- 2: $T[i, j] \leftarrow \{0\}^B$ for $1 \leq i < 2^{H+1}, 1 \leq j \leq Z$
- 3: $n_e \leftarrow 0, n_r \leftarrow 0$
- 4: **for** $i = 1 \dots, N$ **do**
- 5: $z_i \xleftarrow{\$} [2^H, \dots, 2^{H+1})$
- 6: Select (x_i, y_i) s.t. $x_i \in \mathcal{P}(z_i)$ and $T[x_i, y_i]$ is empty
- 7: $T[x_i, y_i] \leftarrow b_i$
- 8: $\text{pm}[\text{id}_i] \leftarrow (z_i, \lceil \log_2 x_i \rceil \cdot Z + y_i)$
- 9: **for** $i = 1, \dots, 2^{H+1} - 1$ **do**
- 10: **for** $j = 1, \dots, Z$ **do**
- 11: $(c_{i,j}^{(1)}, \dots, c_{i,j}^{(m)}) \leftarrow T[i, j]$, where $c_{i,j}^{(k)} \in \mathbb{F}_p$
- 12: **for** $k = 1, \dots, m$ **do**
- 13: $(\llbracket c_{i,j}^{(k)} \rrbracket_1, \dots, \llbracket c_{i,j}^{(k)} \rrbracket_\ell) \leftarrow \text{SSS.Create}(c_{i,j}^{(k)}, t)$
- 14: $\llbracket T[i, j] \rrbracket_l \leftarrow (\llbracket c_{i,j}^{(1)} \rrbracket_l, \dots, \llbracket c_{i,j}^{(m)} \rrbracket_l)$, for all $1 \leq l \leq \ell$
- 15: **return** $(\llbracket T \rrbracket_1, \dots, \llbracket T \rrbracket_\ell)$

▷ Send $\llbracket T \rrbracket_i$ to S_i , for $1 \leq i \leq \ell$

outlines the structure of S³ORAM.

3.2 S³ORAM Access Protocol

Protocol 2 $b \leftarrow \text{S}^3\text{ORAM.Access}(\text{op}, \text{id}, b^*)$

```

1:  $b \leftarrow \text{S}^3\text{ORAM.Retrieve}(\text{id})$ 
2:  $\text{pm}[\text{id}].\text{pID} \xleftarrow{\$} [2^H, \dots, 2^{H+1})$ 
3: if  $\text{op} = \text{write}$  then
4:    $b \leftarrow b^*$ 
5:  $\text{pm}[\text{id}].\text{pIdx} \leftarrow n_r + 1$ 
6:  $(c_1, \dots, c_m) \leftarrow b$ 
7: for  $j = 1 \dots m$  do
8:    $(\llbracket c_j \rrbracket_1, \dots, \llbracket c_j \rrbracket_\ell) \leftarrow \text{SSS.Create}(c_j, t)$ 
9: for  $i = 1 \dots \ell$  do
10:  Write  $(\llbracket c_1 \rrbracket_i, \dots, \llbracket c_m \rrbracket_i)$  to slot  $\llbracket \text{T}[1, n_r + 1] \rrbracket_i$  in server  $\mathcal{S}_i$ 
11:  $n_r \leftarrow n_r + 1 \pmod A$ 
12: if  $n_r = 0$  then
13:   Execute S3ORAM.SSS-SMP-TripletEviction protocol
14:    $n_e \leftarrow n_e + 1 \pmod 2^H$ 
15: return  $b$ 

```

The S³ORAM access protocol consists of two subroutines including S³ORAM.Retrieve and S³ORAM.SSS-SMP-TripletEviction as shown in Protocol 2. We first describe a SSS-based select scheme that is used in the S³ORAM.Retrieve subroutine. We then describe our new eviction strategy based on Triplet Eviction [12]. S³ORAM eviction is performed after every A successive accesses as in [12, 33].

- *SSS-based Select Scheme:* Our objective in this select scheme is to privately retrieve a block of interest residing in the queried path from the S³ORAM structure. Recall that in single-server HE-based ORAM schemes (e.g., [12, 3]), the select query is encrypted with costly additive/fully HE. In our case, we “encrypt” S³ORAM structure with SSS that offers highly efficient additive and multiplicative homomorphic properties. We observe that multi-server PIR scheme in [19, 4] relies on SSS to create select queries and, therefore, it can serve as a suitable private retrieval tool to be used for S³ORAM structure. We describe SSS-based select scheme in Algorithm 5, and further outline it as follows:

Assume that each server \mathcal{S}_i stores a share of the database \mathbf{b} containing n items denoted as $\llbracket \mathbf{b} \rrbracket_i$, which can be interpreted as a vector with each i -th component being the share of the i -th item in \mathbf{b} . Let j be the index of the item in \mathbf{b} to be privately retrieved. The client creates an n -dimensional select vector with all zero coordinates except the j -th coordinate being set to 1 and then, secret-shares it with SSS (PIR.CreateQuery algorithm). The client then distributes these shares to their corresponding servers, each answering with the result of the inner product between the received share vector and its share of \mathbf{b} (PIR.Retrieve algorithm). Finally, the client invokes SSS.Recover function over ℓ answers to recover the desired item (PIR.Reconstruct algorithm). Note that \mathbf{b} in this context is SSS-secret shared, instead of being unencrypted as in [19, 4]. Therefore, our PIR.Reconstruct algorithm requires *at least $2t + 1$ shares instead of $t + 1$* to recover the item correctly.

Algorithm 5 SSS-based Select Scheme

$(\llbracket \mathbf{e} \rrbracket_1, \dots, \llbracket \mathbf{e} \rrbracket_\ell) \leftarrow \text{PIR.CreateQuery}(j)$: Create select queries

- 1: Let $\mathbf{e} := (e_1, \dots, e_n)$, where $e_j \leftarrow 1$, $e_i \leftarrow 0$ for $1 \leq i \neq j \leq n$
- 2: **for** $i = 1, \dots, n$ **do**
- 3: $(\llbracket e_i \rrbracket_1^{(t)}, \dots, \llbracket e_i \rrbracket_\ell^{(t)}) \leftarrow \text{SSS.Create}(e_i, t)$
- 4: $\llbracket \mathbf{e} \rrbracket_i^{(t)} := (\llbracket e_1 \rrbracket_i^{(t)}, \dots, \llbracket e_n \rrbracket_i^{(t)})$, for $1 \leq i \leq \ell$
- 5: **return** $(\llbracket \mathbf{e} \rrbracket_1^{(t)}, \dots, \llbracket \mathbf{e} \rrbracket_\ell^{(t)})$

$\llbracket b \rrbracket_i^{(2t)} \leftarrow \text{PIR.Retrieve}(\llbracket \mathbf{e} \rrbracket_i^{(t)}, \llbracket \mathbf{b} \rrbracket_i^{(t)})$: Retrieve the queried block

- 1: $\llbracket b \rrbracket_i^{(2t)} \leftarrow \llbracket \mathbf{e} \rrbracket_i^{(t)} \cdot \llbracket \mathbf{b} \rrbracket_i^{(t)}$
- 2: **return** $\llbracket b \rrbracket_i^{(2t)}$

$b \leftarrow \text{PIR.Reconstruct}(\llbracket b \rrbracket_1^{(2t)}, \dots, \llbracket b \rrbracket_\ell^{(2t)})$: Reconstruct the block

- 1: $b \leftarrow \text{SSS.Recover}(\llbracket b \rrbracket_1^{(2t)}, \dots, \llbracket b \rrbracket_\ell^{(2t)}, 2t)$
- 2: **return** b

S³ORAM Retrieval: We present S³ORAM retrieval protocol in Subroutine 1, which employs three functions of the aforementioned SSS-based select scheme. Given a block to be read, the client first determines its location in the S³ORAM structure via the position map pm and then, retrieves it using SSS-based select protocol. In this case, we interpret all slots in the retrieval path as the database input \mathbf{b} in PIR.Retrieve algorithm. Hence, the size of \mathbf{b} and the length of the query vector is $n = Z \cdot (H + 1)$. Note that there are m separate chunks in each slot, the servers need to invoke PIR.Retrieve algorithm m times with the same select query but over different \mathbf{b}_j , where each \mathbf{b}_j contains the j -th chunk of all slots in the retrieval path. Finally, the client obtains the desired block by recovering all chunks upon receiving their shares with PIR.Reconstruct algorithm (steps 9–11).

Subroutine 1 $b \leftarrow \text{S}^3\text{ORAM.Retrieve}(\text{id})$

Client:

- 1: $(s, j) \leftarrow \text{pm}[\text{id}]$
- 2: $(\llbracket \mathbf{e} \rrbracket_1^{(t)}, \dots, \llbracket \mathbf{e} \rrbracket_\ell^{(t)}) \leftarrow \text{PIR.CreateQuery}(j)$
- 3: Send $(s, \llbracket \mathbf{e} \rrbracket_i^{(t)})$ to server \mathcal{S}_i , for $1 \leq i \leq \ell$

Server: each $\mathcal{S}_i \in \{\mathcal{S}_1, \dots, \mathcal{S}_\ell\}$ receiving $(s, \llbracket \mathbf{e} \rrbracket_i^{(t)})$ **do**

- 4: $\mathcal{I} \leftarrow \mathcal{P}(s)$
- 5: **for** $j = 1, \dots, m$ **do**
- 6: Let $\llbracket \mathbf{c}_j \rrbracket_i^{(t)}$ contain j -th chunks of Z slots in $\llbracket \mathbb{T}[i'] \rrbracket_i^{(t)}, \forall i' \in \mathcal{I}$
- 7: $\llbracket \mathbf{c}_j \rrbracket_i^{(2t)} \leftarrow \text{PIR.Retrieve}(\llbracket \mathbf{e} \rrbracket_i^{(t)}, \llbracket \mathbf{c}_j \rrbracket_i^{(t)})$
- 8: Send $(\llbracket c_1 \rrbracket_i^{(2t)}, \dots, \llbracket c_m \rrbracket_i^{(2t)})$ to client

Client: On receive $(\{\llbracket c_1 \rrbracket_i^{(2t)}\}_{i=1}^\ell, \dots, \{\llbracket c_m \rrbracket_i^{(2t)}\}_{i=1}^\ell)$

- 9: **for** $j = 1, \dots, m$ **do**
 - 10: $c_j \leftarrow \text{PIR.Reconstruct}(\llbracket c_j \rrbracket_1^{(2t)}, \dots, \llbracket c_j \rrbracket_\ell^{(2t)})$
 - 11: $b \leftarrow (c_1, \dots, c_m)$
 - 12: **return** b
-

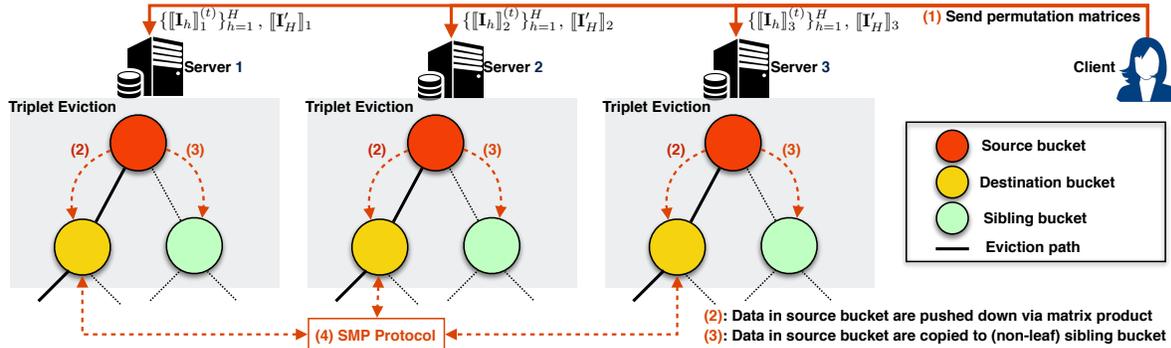


Figure 2: The SSS-SMP-based Triplet Eviction.

After the block is retrieved, the client creates its new SSS-shares (steps 7–8, Protocol 2), and then writes the share to an empty slot in the root bucket of the corresponding server (step 10). After A successive retrievals, the background eviction is performed as described below.

SSS-SMP-based Triplet Eviction: For each eviction operation, the client selects a deterministic eviction path according to the reverse lexicographical order as proposed in [18]. Given a binary ORAM tree of height H , where edges in each level are indexed by either 0 (left) or 1 (right), the collection of edges of the eviction path at the n_e -th eviction operation is determined as:

$$v = \text{DigitReverse}_2(n_e \bmod 2^H), \quad (6)$$

where $\text{DigitReverse}_2(a)$ denotes the order-reversal of base-2 digits of the integer a .

The S^3 ORAM eviction strategy is inspired on the ‘‘Triplet Eviction’’ strategy of Onion-ORAM [12]. Specifically, it percolates real data blocks in the eviction path to lower levels as much as possible, in which each real block b from each bucket T_i can be moved to one of T_i ’s children such that b still resides in its own path. This strategy incurs significant cost in Onion-ORAM due to the following reasons: (i) Onion-ORAM relies on additively HE so that after each pushdown operation, a layer of encryption is added to blocks, which increases the cost of select query creation and block decryption in the retrieval phase. (ii) It requires multiple rounds of client-server communication to (ii-1) get the precise location of the real blocks for select query creation (since this info is stored in the bucket’s metadata) and, (ii-2) bound the number of encryption layers at the leaf buckets. Notice that the latter also requires the client to perform a number of costly homomorphic encryptions and decryptions.

- *Our New Triplet Eviction Strategy:* In this paper, we propose a Triplet Eviction strategy that only requires *single-round client-server communication and lightweight client computation, and avoids accumulating multi-layer of encryption to the ORAM structure*. The main idea is to leverage SSS and SMP protocol to perform block permutation and maintain the consistency of privacy level. We present this strategy as follows:

Observe that since the precise location of real blocks is locally stored in the position map, the client is not required to interact with the server(s) to read the metadata. For each level in the eviction path, we obviously move blocks from source bucket T_i to its children T_{2i} , T_{2i+1} . We follow the terminology used in [12] to denote the buckets involved in each Triplet Eviction operation: if one child of the source bucket resides in the eviction path, it is called *destination* bucket and the other is called *sibling* bucket (see Figure 2). Eviction is performed according to the following rules:

Subroutine 3 S³ORAM.SSS-SMP-TripletEviction

Client:

```

1:  $v \leftarrow \text{DigitReverse}_2(n_e \bmod 2^H)$ 
2: Let  $\vec{u} = (u_1, \dots, u_H)$  be the indexes of src buckets on the path indicated by  $v$ , starting from the root
3: for  $h = 1, \dots, H - 1$  do
4:   Let  $\mathbf{I}_h$  be a  $2Z \times Z$  matrix, set  $\mathbf{I}_h[*,*] \leftarrow 0$ . Let  $\hat{u}_h$  be the index destination bucket of  $\mathbb{T}[u_h]$ 
5:   for each real block with id in  $\mathbb{T}[\hat{u}_h]$  do
6:      $\mathbf{I}_h[j - Z \cdot (h - 1), j - Z \cdot h] \leftarrow 1$  where  $j \leftarrow \text{pm}[\text{id}].\text{pldx}$ 
7:   for each real block with id' in  $\mathbb{T}[u_h]$  do
8:      $(s', j') \leftarrow \text{pm}[\text{id}']$ 
9:     if  $\mathcal{P}(s', h) = \hat{u}_h$  then
10:      Let  $j''$  be the index of an empty slot selected in  $\mathbb{T}[\hat{u}_h]$ 
11:       $\mathbf{I}_h[j' - Z \cdot (h - 1), j''] \leftarrow 1$ ,  $\text{pm}[\text{id}'].\text{pldx} \leftarrow Z \cdot h + j''$ 
12:     else
13:       $\text{pm}[\text{id}'].\text{pldx} \leftarrow j' + Z$ 
14: Execute steps (3)–(11) with  $h = H$  producing  $\mathbf{I}_H$  for source-to-destination permutation at leaf level
15: Execute steps (3)–(11) with  $h = H$ ,  $\hat{u}_h =$  index of sibling bucket of  $\mathbb{T}[u_h]$  producing  $\mathbf{I}'_H$  for source-to-sibling permutation at leaf level
16:  $\llbracket \mathbf{I}_h[i, j] \rrbracket_1^{(t)}, \dots, \llbracket \mathbf{I}_h[i, j] \rrbracket_\ell^{(t)} \leftarrow \text{SSS.Create}(\mathbf{I}_h[i, j], t)$  for  $1 \leq h \leq H$ ,  $1 \leq i \leq 2Z$ ,  $1 \leq j \leq Z$ 
17:  $\llbracket \mathbf{I}'_H[i, j] \rrbracket_1^{(t)}, \dots, \llbracket \mathbf{I}'_H[i, j] \rrbracket_\ell^{(t)} \leftarrow \text{SSS.Create}(\mathbf{I}'_H[i, j], t)$  for  $1 \leq i \leq 2Z$ ,  $1 \leq j \leq Z$ 
18: Send  $(v, \{\llbracket \mathbf{I}_h \rrbracket_i^{(t)}\}_{h=1}^H, \llbracket \mathbf{I}'_H \rrbracket_i^{(t)})$  to  $\mathcal{S}_i$ , for  $1 \leq i \leq \ell$ 

```

Server: each $\mathcal{S}_i \in \{\mathcal{S}_1, \dots, \mathcal{S}_\ell\}$ receiving $(v, \{\llbracket \mathbf{I}_h \rrbracket_i^{(t)}\}_{h=1}^H, \llbracket \mathbf{I}'_H \rrbracket_i^{(t)})$ **do**

```

19: Let  $\vec{u} = (u_1, \dots, u_H)$  be the indexes of src buckets on the path indicated by  $v$ , starting from the root
20:  $\llbracket \mathbb{T}[\tilde{u}_h] \rrbracket_i \leftarrow \llbracket \mathbb{T}[u_h] \rrbracket_i$  for  $1 \leq h < H$   $\triangleright$  Replace sibling bucket with source bucket at non-leaf levels
21: for  $h = 1, \dots, H$  do
22:   Let  $\hat{u}_h$  be the index of destination bucket of  $\mathbb{T}[u_h]$ 
23:   for  $j = 1, \dots, m$  do
24:     Let  $\llbracket \mathbf{c}_{h,j} \rrbracket_i^{(t)}$  contain  $j$ -th chunks from  $Z$  slots of  $\llbracket \mathbb{T}[u_h] \rrbracket_i^{(t)}$  and  $\llbracket \mathbb{T}[\hat{u}_h] \rrbracket_i^{(t)}$ 
25:      $\llbracket \mathbf{c}'_{h,j} \rrbracket_i^{(2t)} \leftarrow \llbracket \mathbf{c}_{h,j} \rrbracket_i^{(t)} \cdot \llbracket \mathbf{I}_h \rrbracket_i^{(t)}$ , and execute SMP protocol to reduce  $\llbracket \mathbf{c}'_{h,j} \rrbracket_i^{(2t)}$  to  $\llbracket \mathbf{c}'_{h,j} \rrbracket_i^{(t)}$ 
26:     Update  $j$ -th chunks in  $Z$  slots in  $\llbracket \mathbb{T}[u_h] \rrbracket_i^{(t)}$  and  $\llbracket \mathbb{T}[\hat{u}_h] \rrbracket_i^{(t)}$  with  $\llbracket \mathbf{c}'_{h,j} \rrbracket_i^{(t)}$ 
27: Execute steps (21)–(26) with  $h = H$ ,  $\llbracket \mathbf{I}_h \rrbracket_i^{(t)} = \llbracket \mathbf{I}'_H \rrbracket_i^{(t)}$ ,  $\hat{u}_h =$  index of sibling bucket of  $\mathbb{T}[u_h]$ 

```

- *Source to destination:* Let $\llbracket \mathbf{u} \rrbracket$ be a $2Z$ -dimensional share vector formed by concatenating all data in the source bucket and the destination bucket. The client creates a permutation matrix $\mathbf{I} \in \{0, 1\}^{2Z \times Z}$ such that the matrix product of $\llbracket \mathbf{u} \rrbracket$ and \mathbf{I} will result in a Z -dimensional vector $\llbracket \mathbf{v} \rrbracket$, in which data at position i in $\llbracket \mathbf{u} \rrbracket$ is moved to position j in $\llbracket \mathbf{v} \rrbracket$. That is, \mathbf{I} is a matrix, where $\mathbf{I}[i, j] \leftarrow 1$ if the block at position i in $\llbracket \mathbf{u} \rrbracket$ is expected to move to position j in $\llbracket \mathbf{v} \rrbracket$. As a result, $\mathbf{I}[i + Z, i] \leftarrow 1$ if the block currently at position i in $\llbracket \mathbf{v} \rrbracket$ stays at the same location. To hide the location information of real blocks after permutation, the client “encrypts” every single element of \mathbf{I} with SSS resulting in a share matrix $\llbracket \mathbf{I} \rrbracket \in \mathbb{F}_p^{2Z \times Z}$. Note that the matrix product between these two shares results in a share vector with each element being represented by a degree- $2t$ polynomial. To maintain the consistency of the S³ORAM structure, servers will together execute the SMP Protocol presented in Section 2 to reduce the degree of polynomial of each component in $\llbracket \mathbf{v} \rrbracket$ from $2t$ to t .

- *Source to sibling:* We can apply the same trick as in the source-to-destination operation above to push real blocks down to sibling buckets. However, since non-leaf sibling buckets are guar-

anted to be empty by previous evictions featuring a negligible bucket overflow probability (see Lemma 3), this process can be further optimized as discussed in [12] as follows. For evictions not involved with leaf buckets, the client simply requests servers to copy all data from the source buckets to sibling buckets and then, update the path index of real blocks in the position map accordingly. For leaf level, it is required to use the matrix permutation as described above since leaf buckets are not empty. This optimization can halve the bandwidth cost of client-server and server-server communication.

Generally, we can see that our eviction approach requires only one client-server communication and guarantees that all data after eviction are consistently “encrypted” by degree- t polynomials. Figure 2 illustrates this new SSS-based Triplet Eviction strategy. We present the algorithmic description of this strategy in Subroutine 3.

3.3 Asymptotic Cost Analysis

In this section, we study the cost of S³ORAM pertaining to block size and number of blocks, where the security level and other system parameters (e.g., prime field, bucket size) are fixed.

- **Communication:** The size of each select query being sent to ℓ servers in S³ORAM retrieval is $(Z \cdot (H + 1) \cdot \lceil \log_2 p \rceil)$ bits. The client sends $H + 1$ permutation matrices in the S³ORAM eviction to ℓ servers, each being of size $2Z^2 \cdot \lceil \log_2 p \rceil$ bits. Each S³ORAM access incurs one block of size B to be transferred between client and server. Since $H \geq \log N/A + 1$ and ℓ, A, Z, p are constants, the overall client-server communication complexity is $\mathcal{O}(B + \log N)$. In the eviction, each server distributes the shares of H buckets with size of $Z \cdot B$ bits to each other in H communication rounds. Hence, the inter-server communication overhead is $\mathcal{O}(B \cdot \log N)$.

Achieving $\mathcal{O}(1)$ client-server bandwidth blowup: The client bandwidth blowup is defined as the *ratio* between the number of client-server communication introduced by ORAM and the base case without ORAM being used. The communication complexity of S³ORAM shows that the size of select vector and permutation matrix is independent of block size. Note that in Onion-ORAM, the select vector size is also independent of the block size. Therefore, $\mathcal{O}(1)$ client bandwidth blowup can be achieved in S³ORAM and Onion-ORAM by selecting a suitable block size. That is, by selecting the block size B to be $\Omega(\log N)$, S³ORAM achieves $\mathcal{O}(1)$ client bandwidth blowup. In Onion-ORAM, it requires selecting a (larger) block size of $\mathcal{O}(\log^5 N)$ to absorb the size of select queries.

- **Computation:** Each server computes the inner product between the $Z \cdot (H + 1)$ -dimensional select vector and the block vector containing $Z \cdot (H + 1)$ blocks of size B . For eviction, each server computes H times the matrix product between a vector containing $2 \cdot Z$ blocks of size B and a permutation matrix of size $2Z \times Z$. After that, each server computes the share and performs degree reduction in the SMP protocol on $Z \cdot H$ blocks of size B . In total, the server computation complexity is $\mathcal{O}(B \cdot \log N)$.

The client invokes SSS.Create algorithm $Z \cdot (H + 1)$ times and $2Z^2 \cdot H$ times to create a select query and H permutation matrices, respectively. The client invokes SSS.Recover and SSS.Create algorithms to reconstruct and re-share a block of size B , respectively. Thus, the overall client computation complexity is $\mathcal{O}(B + \log N)$.

- **Storage:** S³ORAM tree structure is of height H which has $2^H \cdot Z$ slots and can store up to $N \leq A \cdot 2^{H-1}$ real blocks. Since Z and A are constants, the server storage cost is $\mathcal{O}(B \cdot N)$. Notice that the share of the value has the same size as the value (i.e., no ciphertext expansion

as in Onion-ORAM), the server storage of S^3 ORAM is constant and does not increase after a sequence of access operations.

Similar to Onion-ORAM, the block storage in S^3 ORAM is $\mathcal{O}(1)$ since the client immediately writes retrieved block back to the root bucket. The client locally stores the position map whose cost is $\mathcal{O}(N(\log N + \log \log N))$.

Achieving $\mathcal{O}(1)$ client storage via Recursion: For theoretical interest, S^3 ORAM can achieve (in total) $\mathcal{O}(1)$ client storage by storing the position map in smaller ORAMs using the recursion technique described in [39] and bucket metadata structure in [12]. Specifically, for each S^3 ORAM bucket, we create a metadata that stores the current index (pIdx) and the assigned path (pID) of blocks residing in it. For each S^3 ORAM access, the metadata of buckets along the retrieval/eviction path will be read first to get the precise location and the assigned path of blocks of interest. This information will be used to create the select query and permutation matrices. Next, we construct a series of S^3 ORAM structures $S^3\text{ORAM}_0, \dots, S^3\text{ORAM}_X$, where $X = \mathcal{O}(\log N)$, $S^3\text{ORAM}_0$ stores data blocks and $S^3\text{ORAM}_{i+1}$ stores the position map of $S^3\text{ORAM}_i$. Note that in this recursion, the position map only stores the blocks' assigned path since their precise location is already maintained in bucket metadata. We refer the reader to [39, 12] for the detailed descriptions.

In $S^3\text{ORAM}_0$, the bucket metadata is of size $\mathcal{O}(\log N)$. Each $S^3\text{ORAM}_0$ retrieval/eviction accesses the metadata from $\mathcal{O}(\log N)$ buckets resulting in $\mathcal{O}(\log^2 N)$ client-server bandwidth overhead. Therefore, to achieve $\mathcal{O}(1)$ client bandwidth blowup, the block size of $S^3\text{ORAM}_0$ needs to be $\Omega^2(\log N)$. Since the block size of $S^3\text{ORAM}_{i+1}$ is smaller than that of $S^3\text{ORAM}_i$, applying recursion technique to S^3 ORAM does not increase the bandwidth, computation or server storage overhead in the asymptotic point of view. However, it incurs $\mathcal{O}(\log N)$ communication rounds and requires $\mathcal{O}(\log N)$ factor larger block size, which might significantly increase the end-to-end delay in reality. Thus, it is recommended to maintain the position map locally, given that its size is small enough, to gain performance advantages in practice.

4 Security

The S^3 ORAM eviction follows the Triplet Eviction proposed in Onion-ORAM [12]. Therefore, they have the same failure probability. We refer the reader to [12] for the details of the proof.

Lemma 3 (Bucket Overflow Probability). *If $Z \geq A$ and $N \leq A \cdot 2^H - 1$, the probability that a bucket overflows after an eviction operation is bounded by $e^{-\frac{(2Z-A)^2}{6A}}$, where $Z = A = \Theta(\lambda)$.*

Proof. We refer the reader to [12]. □

The correctness of S^3 ORAM is shown in Theorem 1.

Theorem 1 (S^3 ORAM Correctness). *S^3 ORAM is correct by Definition 2.*

Proof. See Appendix. □

The security of S^3 ORAM is given in Theorem 2.

Theorem 2 (S^3 ORAM Security). *S^3 ORAM is unconditionally t -secure by Definition 2.*

Proof. See Appendix. □

Malicious setting: We do not consider malicious servers in this paper. However, since the S³ORAM relies on SSS as its main building block, it is possible to extend the scheme to tolerate malicious inputs. This requires more servers and additional rounds of interaction during the eviction phase due to the more involved security requirements (e.g., verifiable secret sharing [16]) in order to distribute correct shares and detect malicious behaviors. We will investigate the cost for malicious setting in the full version.

5 Experimental Evaluation

We first describe our implementation details and configuration. We then give our evaluation metrics and methodology, followed by a detailed comparison of S³ORAM and its counterparts on an actual cloud environment with various network and system settings. Note that in this evaluation, we evaluate ORAM schemes under their non-recursive form, where the position map is stored at the client since its size is practically small to be stored locally.

5.1 Implementation Details and Configuration

The implementation of S³ORAM is available at <https://github.com/thanghoang/S3ORAM>.

Software setting. We implemented S³ORAM in C++ with two external libraries: (i) Shoup’s NTL library v9.10.0³ for pseudo-random number generation and modular computations since it offers low-level (e.g., assembly) optimizations for modular multiplication and inner product; (ii) ZeroMQ library⁴ for socket programming. Our implementation supports parallelization to take full advantage of multi-core CPUs at the server. We used `libtomcrypt`⁵ with AES-CTR to implement IND-CPA encryption for S³ORAM counterparts.

Hardware setting. We conducted our experiments on two types of client devices: (i) A 2015 Macbook Pro laptop as the client, which was equipped with Intel Core i5-5287U CPU @ 2.90GHz and 16 GB RAM. (ii) A Google Nexus 6P smartphone, which ran Android 7.0 Nougat and was equipped with Qualcomm Snapdragon 810 CPU @ 2 GHz and 3 GB RAM. At the server side, we used Amazon EC2 with `c4.4xlarge` type to deploy three server instances. Each server was running Ubuntu 16.04 and equipped with 16 vCPUs Intel Xeon E5-2666 v3 @2.9 GHz, 30 GB RAM and 512 GB SSD.

Network setting. We located three servers to be geographically close to the clients as well as to each other, which results in the network latency between them to be approximately 15 ms. Servers were connected to each other via dedicated networks whose throughput for both download and upload is approximately 250 Mbps.

The laptop client was connected to a home Internet service via Wi-Fi, which offers download and upload throughputs of 29 Mbps and 5 Mbps to the servers, respectively. For mobile client, we used LTE network to communicate with server(s), which has a network latency of 25 ms, and the download and upload throughputs of approximately 27 Mbps and 9 Mbps, respectively.

³Available at <http://www.shoup.net/ntl/download.html>

⁴Available at <http://zeromq.org>

⁵Available at <https://github.com/libtom/libtomcrypt>

Database size. We evaluated the performance of all compared schemes with a randomly generated database and block size ranging from 0.5 GB to 40 GB, and from 64 KB to 768 KB, respectively.

5.2 Evaluation Metrics and Methodology

Evaluation Metrics. We compared S³ORAM with its counterparts based on: (1) Building time for ORAM structure (executed once at the beginning); (2) End-to-end delay for different database and block sizes; The cost breakdown of end-to-end delay to assess the impacts of (3) client computation overhead, (4) server computation overhead, (5) client-server communication, (6) server-server communication, (7) disk access time, (8) network bandwidth quality, (9) client and server storage overhead.

We selected Path-ORAM [40] and Onion-ORAM [12] as the main counterparts of S³ORAM, since the former achieves the optimal lower bound of $\Omega(\log N)$ communication blowup without server computation, while the latter achieves $\mathcal{O}(1)$ communication blowup with $\mathcal{O}(1)$ client storage with server computation. We also chose Ring-ORAM [33] as it is an efficient ORAM scheme with server computation. We did not consider alternatives that (i) failed to achieve $\mathcal{O}(1)$ client communication blowup but incurred more delay (e.g., [11, 26]), (ii) were shown to be insecure (e.g., [29, 27]), or (iii) incur significantly more cost than ORAMs considered in our experiments (e.g., [3]) (see Section 6 for related work). We also did not explicitly compare the performance of S³ORAM against the distributed ORAM by Stefanov et al. [37] because of the major difference in terms of client block storage between the two schemes ($\mathcal{O}(1)$ vs. $\mathcal{O}(\sqrt{N})$). Given a very large outsourced database, the storage requirement by [37] might not be suitable for resource-limited devices such as mobile phone. Moreover, if $\mathcal{O}(\sqrt{N})$ block storage is acceptable, then the lower bound in [1] might imply a better ORAM strategy than S³ORAM that leverages PIR-only technique to achieve $\mathcal{O}(1)$ client bandwidth blowup.

Evaluation Methodology. Our methodology is as follows.

- *S³ORAM*: We fully implemented S³ORAM and measured the delay of each operation (see *cost breakdown* part below). We selected the bucket size and the eviction frequency as $Z = A = 333^6$ to achieve a negligible bucket overflow probability of 2^{-80} . The cost for each S³ORAM access was measured as the retrieval delay with the write-to-root delay (step 10, Protocol 2) plus the amortized cost of eviction.
- *Path-ORAM*: We measured the delay of Path-ORAM as the time to (1) download/upload $\mathcal{O}(\log N)$ blocks, and (2) IND-CPA encryption/decryption at the client. We selected the bucket size as 5 to guarantee a negligible stash overflow probability of 2^{-80} .
- *Ring-ORAM*: We measured the delay of Ring-ORAM as the time to (1) retrieve $\mathcal{O}(1)$ block, (2) perform XOR and IND-CPA encryption/decryption at the client, (3) perform XOR operations at the server side. The amortized cost of each Ring-ORAM access is calculated by including the amortized cost of eviction and early shuffles based on the equation $(H + 1)(2Z + S)/A \cdot (1 + \text{PoisCDF}(S, A))$ given in Section 5 of [33]. We selected Ring-ORAM

⁶S³ORAM follows the eviction strategy from Onion-ORAM where the bucket overflow probability is analyzed with the assumption on the ideal PRF. For the real deployment of S³ORAM (and Onion-ORAM) in practice, we recommend to follow Lemma 3 but select Z to be larger than A , as the real-life implementation of PRF might not always guarantee the bucket load $\leq A/2$ after each triplet eviction.

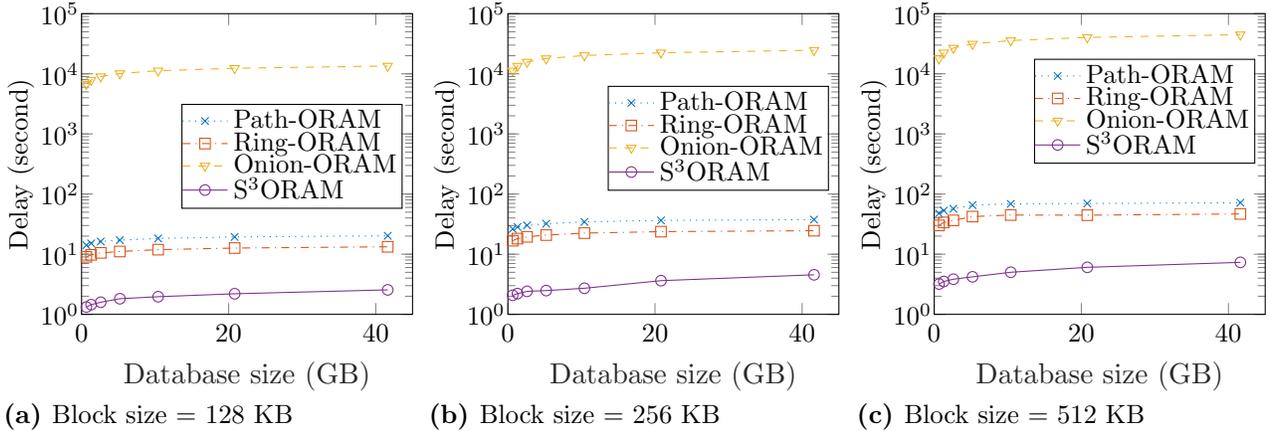


Figure 3: End-to-end delay of S³ORAM and its counterparts on a laptop with home network.

parameters $Z = 16$, $S = 25$ and $A = 20$ as stated in [33] for a negligible stash overflow probability of 2^{-80} .

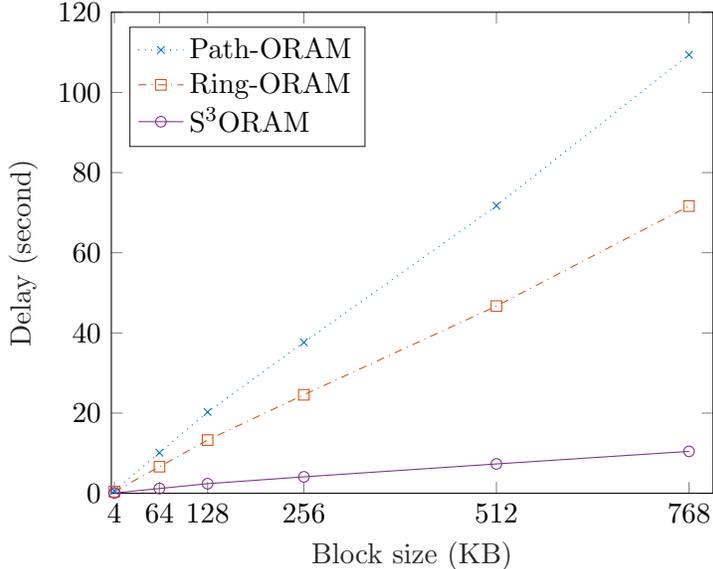
- *Onion-ORAM*: We measured the delay of Onion-ORAM as the time to (1) perform homomorphic computations at the client and server, (2) transfer $\mathcal{O}(1)$ blocks and select queries. We selected the size of RSA modulus to be 1024 bits for AHE according to [2]. Similar to S³ORAM, we selected the bucket size and the eviction frequency of Onion-ORAM as $Z = A = 333$. The cost for each Onion-ORAM access was also included with the amortized cost of eviction operation. Since Onion-ORAM is extremely computation costly, measuring its delay even on a medium database takes insurmountable amount of time. Therefore, we had to measure its delay on a very small database (i.e., 1 MB) first, and then estimate the delay for larger database sizes.

5.3 Experimental Results

- *Building Time of ORAM Structure (executed only once during offline phase)*: We first provide the total building time for constructing the distributed S³ORAM tree structures. Since S³ORAM relies on highly efficient SSS operations (e.g., basic arithmetics with modular addition/multiplication), it only took around 1 hour to create shares of a large database (i.e., 40 GB) for three servers with the laptop as the client device. This cost is comparable with IND-CPA encryption being used in traditional ORAM schemes. For instance, it took approximately 50-60 minutes to encrypt counterpart ORAM schemes with AES-CTR encryption.

• *End-to-end Delay*: Figure 3 presents the end-to-end delay of S³ORAM and its counterparts. S³ORAM outperformed its counterparts for increasing database and block sizes. As shown in Figure 3b, S³ORAM took around 4.5 seconds to access a 256-KB block, which is approximately $8.3\times$ and $5.4\times$ faster than Path-ORAM and Ring-ORAM, respectively, while being three orders of magnitude faster than Onion-ORAM (40 GB DB).

Figure 3 shows that choosing larger block sizes had a minimal impact on the delay of S³ORAM compared to its counterpart. For instance, S³ORAM took around 2.4 and 7.3 seconds to access a 128-KB block and a 512-KB block, respectively, which corresponds to a linear growth but with a small slope (40 GB DB). Figure 4 further illustrates the influence of block size on



*We excluded Onion-ORAM since its plot is far beyond the limit of y-axis.

Figure 4: End-to-end for varying block sizes for a 40GB DB.

the delay of S³ORAM and its counterparts. Although the cost of each ORAM scheme grows linearly with respect to the block size, the slope of S³ORAM is significantly smaller than that of its counterparts. Given any block size in the range from 4 KB to 768 KB, S³ORAM is always approximately 5× and 8× faster than Ring-ORAM and Path-ORAM, respectively. This gives an advantage to S³ORAM over its counterparts for applications with a large block size such as image or video storage services.

Detailed cost breakdown of S³ORAM. We now dissect the end-to-end delay of S³ORAM to investigate which factors contributed the most to the total delay. Figure 5 shows the detailed delay analysis of S³ORAM with three different block sizes on a laptop as a client device with three servers.

- *Server Computation:* The server computation only occupied a small amount (5–8%) of the total delay.

- *Client-server Communication and Disk I/O Access:* The client-server communication and server disk I/O access contributed more than 90% of the total delay. Notice that most of the client-server communication time was spent for retrieving/writing blocks over three servers. The transmission of select vector and permutation matrix only cost approximately 6–7% of the total client-server communication overhead with 512 KB block size. Due to the cache miss issue and the infrastructure of selected Amazon EC2 server instances (i.e., `c4.4xlarge` type), disk I/O access caused a significant delay, especially for large database sizes. Specifically, the maximum RAM of each server was limited to 30GB while each S³ORAM retrieval incurs a random disk I/O access of around 1-2 GB of data due to the large bucket size (i.e., $Z = 333$). As a result, after some random retrieval operations, the cache miss may appear more frequently. Moreover, S³ORAM data structure was stored in a networked storage unit called “Elastic Block Storage” (EBS), which was connected to Amazon EC2 computing unit with a maximum throughput of 160 MB per second. The disk I/O access was also limited by this throughput, and therefore, it is

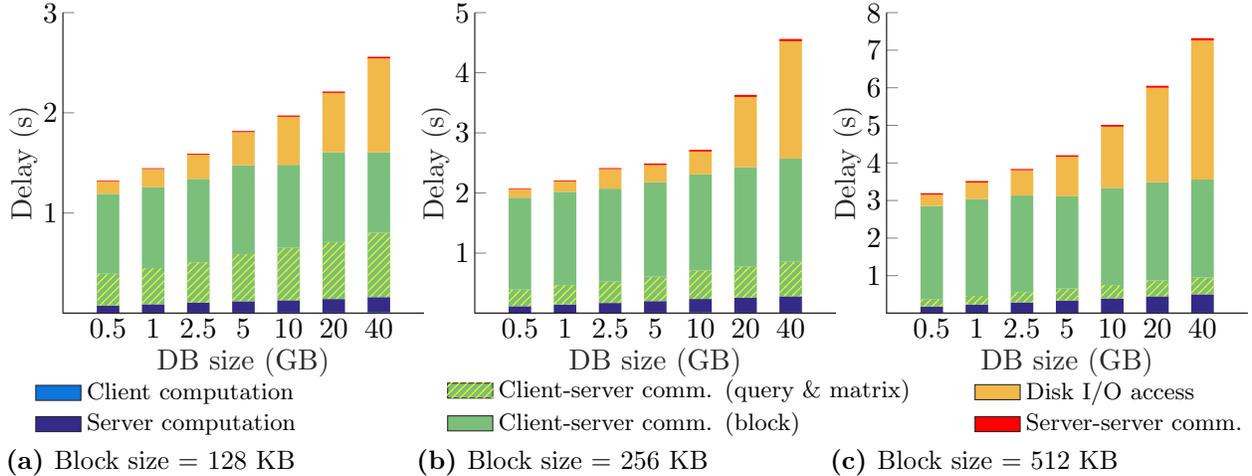


Figure 5: Detailed cost breakdown of S^3 ORAM on a laptop with home network.

much slower than a local storage setting, in which the read/write throughput of 400–500 MB/s can be achieved. Hence, we expect that the latency of S^3 ORAM can be further optimized (at least twice) by minimizing the disk I/O access via special server instances offering either local storage (e.g., internal dedicated SSD) or higher throughput (e.g., *st1* volume type).

- *Server-server Communication:* The overhead of server-server communication is minimal, since servers are connected with a high-bandwidth network (i.e., 250 Mbps). Moreover, the eviction was only performed after $A = 333$ successive retrievals and, therefore, the overhead of server-server communication was amortized.

- *Client Computation:* The client computation (e.g., block recovery, the creation of shares for blocks, the select vectors and permutation matrices) is negligible and, therefore, is difficult to observe in Figure 5.

Detailed cost breakdown of counterpart schemes.

- *Path-ORAM:* Most of the delays in Path-ORAM was due to $\mathcal{O}(\log N)$ client-server communication blowup, which accounted for 97% of overall delay. The client computation was negligible (cost less than 2%) due to IND-CPA encryption/decryption. Since the Path-ORAM bucket size is also much smaller than S^3 ORAM and Onion-ORAM (i.e., 5 vs. 333) which incurs less data to be read and to be written. Therefore, the disk I/O access time of Path-ORAM only took 1% of the overall delay.

- *Ring-ORAM:* Similar to Path-ORAM, the most significant delay of Ring-ORAM was due to the amortized communication cost of early reshuffle and eviction operations, which accounted for 96% of total delay presented in Figure 3. Client and server computations were negligible due to IND-CPA encryption and XOR operations accelerated by multi-threading, respectively. The disk I/O access was also negligible since the server only read 1 block per bucket. Thus, all operations except the communication only contributed less than 4% to the total delay.

- *Onion-ORAM:* Due to AHE, the computation cost dominated all other costs in Onion-ORAM scheme. Specifically, we estimated that given a database of size 0.5 GB containing 256-KB blocks, the server computation with multi-core processing might take approximately 3 hours, which accounted for 99% of the overall delay. Meanwhile, the client took around 38 seconds to generate a select query. Although the disk I/O access time of Onion-ORAM is similar

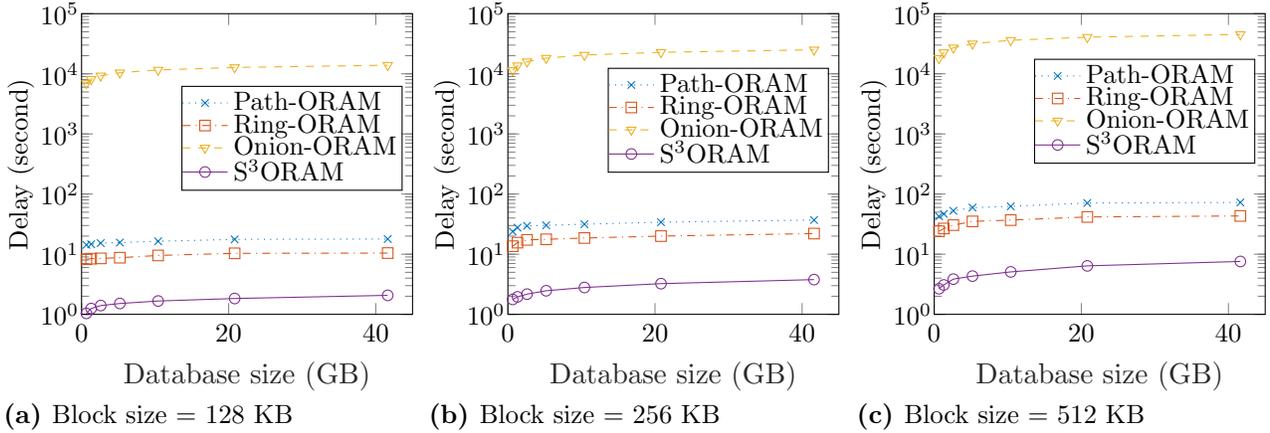


Figure 6: End-to-end delay of S³ORAM and its counterparts on a mobile client with LTE network.

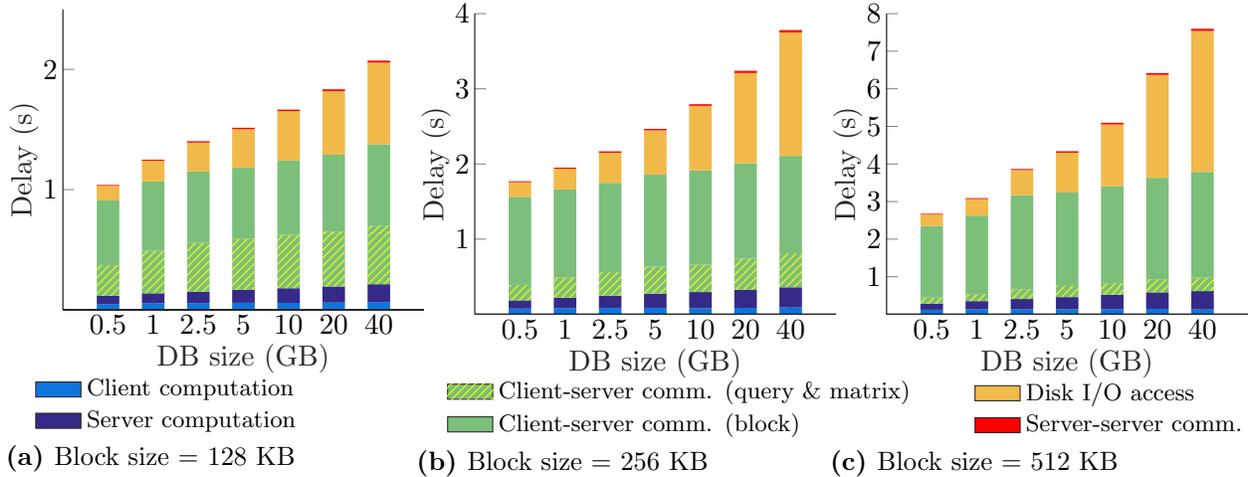
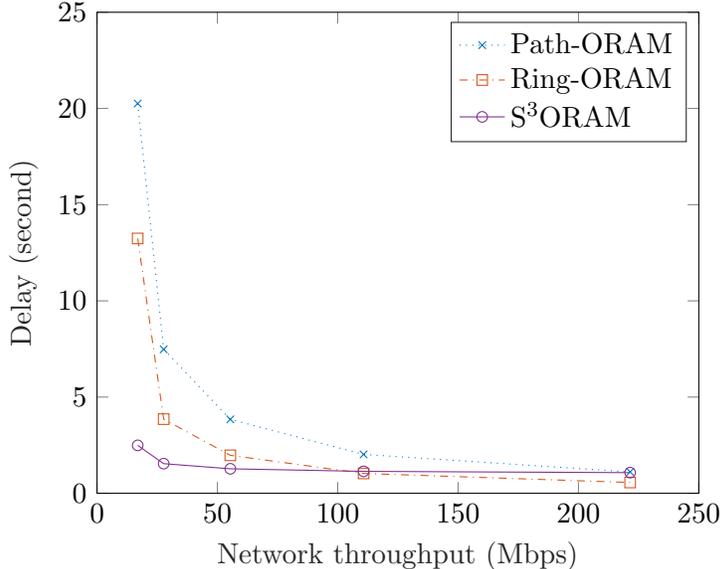


Figure 7: Detailed cost breakdown of S³ORAM on a mobile client with LTE network.

to that of S³ORAM (due to the same bucket size setting), it contributed the least. Transmission also took a small amount of time since Onion-ORAM offers $\mathcal{O}(1)$ bandwidth blowup.

Experiments with a Mobile Client Device. Figure 6 presents end-to-end delays of S³ORAM and its counterparts on a mobile client device with LTE network. The server computation of Onion-ORAM still dominated all others as in previous settings. In addition, Onion-ORAM also requires the client to perform costly computations to generate encrypted select queries, which took a few minutes with a mobile client device. Note that the client computations just took around 50 milliseconds in S³ORAM as shown in Figure 7. The performance of Path-ORAM and Ring-ORAM in the LTE network was relatively better than that of the home network since the LTE network offered a slightly higher upload throughput. S³ORAM was also slightly better but not affected much by the limited client computation. That is, although the client computation on the mobile device contributed a slightly larger delay than that on the laptop, this portion still occupied less than 5% of the total cost.

Comparison of S³ORAM with Path-ORAM and Ring-ORAM for varying network bandwidths. S³ORAM outperformed its counterparts in both home and mobile network set-



*We exclude the Onion-ORAM since its plot is far beyond the limit of the y-axis.

Figure 8: Delay for varying network throughput.

tings. However, assume that the user has high bandwidth network connection. In this case, S³ORAM might not be the best choice in terms of end-to-end delay. Hence, we ran another experiment to show that, $\mathcal{O}(\log N)$ -communication ORAMs are better than computational ORAMs after a certain threshold of network bandwidth. Specifically, we executed S³ORAM and its counterparts schemes several thousand times for increasing network bandwidth values. Figure 8 presents the performance of ORAM schemes with a database size of 40 GB and a block size of 128 KB with varying network throughputs. Observe that Path-ORAM and Ring-ORAM surpassed S³ORAM for a network throughput of approximately 240 Mbps and 110 Mbps, respectively. This is because Path-ORAM and Ring-ORAM are $\mathcal{O}(\log N)$ bandwidth blowup ORAMs and they receive a high benefit from increasing network speeds. However, S³ORAM is $\mathcal{O}(1)$ bandwidth blowup ORAM and receives a less benefit from a fast network.

Storage Overhead. Since S³ORAM and Onion-ORAM feature $\mathcal{O}(1)$ block storage, their client storage cost is lower than that of their counterparts. Given a database with 512 KB blocks, Path-ORAM and Ring-ORAM require around 32-33 MB for the stash, while S³ORAM and Onion-ORAM do not require the stash. The storage cost for a position map in non-recursive S³ORAM is slightly higher than its non-recursive counterparts. For instance, with a 16 TB database of 512-KB blocks ($N = 33554432$), S³ORAM costs 119 MB while the others (e.g., Onion-ORAM, Ring-ORAM, Path-ORAM) cost 100 MB.

In S³ORAM, for a database with N blocks of size B bits, the server storage overhead is $4N \cdot B$ bits for *each* server (recall that S³ORAM needs at least three servers). The server storage for Path-ORAM and Ring-ORAM is $10N \cdot B$ bits and $6N \cdot B$ bits, respectively. The server storage for Onion-ORAM is similar to S³ORAM for one server but will increase after a sequence of access operation due to the ciphertext expansion of AHE.

6 Related Work

Single-server ORAM without computation. The first ORAM proposed by Goldreich et al. [20] was in the context of software protection and followed by refinements (e.g., [21]). The recent ORAM schemes mainly have been considered in the client-server model to hide data access pattern over a remote server (e.g., [32]). Preliminary ORAMs were costly in terms of both communication and storage overhead, but recent ORAMs (e.g., [39, 40, 35, 42]) showed significant improvements. Path-ORAM [40], which follows the tree structure of [35], achieves $\mathcal{O}(\log N)$ communication blowup. Various ORAMs relying on Path-ORAM have been proposed for specific applications such as oblivious data structure (i.e., [44]), secure computation (e.g., ([42], [43]), Parallel ORAM [8]) and secure processor [25]. However, Path-ORAM based schemes inherit its logarithmic communication blowup [6, 30].

Single-server ORAM with computation. Ring-ORAM [33] reduced the communication cost of Path-ORAM by 2.5x given that the server performs XORs. Some other alternatives (e.g., [12, 14, 26, 3, 29]) leveraged single-server PIR or fully/partial HE to further reduce the communication cost. For instance, Onion-ORAM [12] achieves $\mathcal{O}(1)$ bandwidth blowup, where the client and server interactively run partial HE operations. Path-PIR scheme in [26] used PIR scheme in [41] with Additively HE (AHE) (i.e., [31]) on top of tree ORAM structure [35]. Bucket-ORAM in [14] used AHE on top of the underlying ORAM structure composed of tree ORAM and hierarchical ORAM. The scheme in [11] used PIR scheme in [41] on top of ObliviStore [38], which is based on Partition-ORAM in [39]. The TWORAM scheme in [15] constructed a garbled circuit [45] over the tree ORAM structure, which allows the client and server to perform secure computation to access the block.

Multi-server (Distributed) ORAM. Distributed ORAM schemes were proposed to eliminate highly costly fully/partial HE operations. CHf-ORAM [27] attempted to use four non-colluding servers to achieve $\mathcal{O}(1)$ bandwidth blowup under $\mathcal{O}(1)$ blocks of client storage. However, CHf-ORAM [27] (as well as its predecessor in [29]) was broken by Abraham et al. in [1] which also showed an asymptotically tight sub-logarithmic communication bound for composing ORAM with PIR. Abraham et al. in [1] also presented a scheme using two non-colluding servers to perform XORs for block retrieval over a k -ary ORAM tree structure. Stefanov et al. in [37] proposed an ORAM that uses two non-colluding computational-capable servers to reduce the client-server bandwidth of Partition ORAM [39]. In a different line of research, distributed ORAM schemes were proposed for secure multi-party computation (e.g., [24, 13]). In these works, the access patterns are hidden from all parties so that such ORAM schemes are integrated with some secure computation protocol (e.g., Yao’s garbled circuit [45]) and, therefore, their cost is higher than classical client-server ORAM model.

7 Conclusion

We developed a new distributed ORAM scheme that we named S³ORAM, which achieves $\mathcal{O}(1)$ client-server bandwidth blowup under $\mathcal{O}(1)$ client block storage and a low end-to-end delay by avoiding costly HE operations. S³ORAM harnesses Shamir Secret Sharing, tree-based ORAM structure, a new triplet eviction strategy, and a secure multi-party multiplication protocol in an effective manner to achieve these objectives. We performed detailed experiments in an actual cloud environment with a resource-limited mobile client to assess the effectiveness of

S³ORAM for various networking settings such as high bandwidth, home and wireless (i.e., Wi-Fi, LTE) networks. Our experiments showed that S³ORAM is three orders of magnitude faster than the existing single-server ORAM with $\mathcal{O}(1)$ client communication/storage blowup complexity (Onion-ORAM). S³ORAM is also one order of magnitude faster than Path-ORAM on a moderate network bandwidth quality, which is typical for various real-life settings (e.g., home, wireless networks and inter-state cloud deployments).

References

- [1] Ittai Abraham, Christopher W Fletcher, Kartik Nayak, Benny Pinkas, and Ling Ren. Asymptotically tight bounds for composing oram with pir. In *IACR International Workshop on Public Key Cryptography*, pages 91–120. Springer, 2017.
- [2] Anastasov Anton. Implementing onion oram: A constant bandwidth oram using ahe. <https://github.com/aanastasov/onion-oram/blob/master/doc/report.pdf>, 2016.
- [3] Daniel Apon, Jonathan Katz, Elaine Shi, and Aishwarya Thiruvengadam. Verifiable oblivious storage. In *International Workshop on Public Key Cryptography*, pages 131–148. Springer, 2014.
- [4] Amos Beimel and Yoav Stahl. Robust information-theoretic private information retrieval. In *International Conference on Security in Communication Networks*, pages 326–341. Springer, 2002.
- [5] M. Ben-Or, S. Goldwasser, and A. Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation (extended abstract). In Janos Simon, editor, *Proceedings of the 20th Annual ACM Symposium on Theory of Computing*, pages 1–10. ACM, May 2-4, 1988.
- [6] Vincent Bindschaedler, Muhammad Naveed, Xiaorui Pan, XiaoFeng Wang, and Yan Huang. Practicing oblivious access on cloud storage: the gap, the fallacy, and the new way forward. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, pages 837–849. ACM, 2015.
- [7] Elette Boyle and Moni Naor. Is there an oblivious ram lower bound? In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*, pages 357–368. ACM, 2016.
- [8] Binyi Chen, Huijia Lin, and Stefano Tessaro. Oblivious parallel ram: Improved efficiency and generic constructions. In *Theory of Cryptography Conference*, pages 205–234. Springer, 2016.
- [9] Benny Chor, Eyal Kushilevitz, Oded Goldreich, and Madhu Sudan. Private information retrieval. *Journal of the ACM (JACM)*, 45(6):965–981, 1998.
- [10] Ivan Damgård and Mads Jurik. A generalisation, a simplification and some applications of paillier’s probabilistic public-key system. In *International Workshop on Public Key Cryptography*, pages 119–136. Springer, 2001.
- [11] Jonathan Dautrich and Chinya Ravishankar. Combining oram with pir to minimize bandwidth costs. In *Proceedings of the 5th ACM Conference on Data and Application Security and Privacy*, pages 289–296. ACM, 2015.
- [12] Srinivas Devadas, Marten van Dijk, Christopher W Fletcher, Ling Ren, Elaine Shi, and Daniel Wichs. Onion oram: A constant bandwidth blowup oblivious ram. In *Theory of Cryptography Conference*, pages 145–174. Springer, 2016.
- [13] Sky Faber, Stanislaw Jarecki, Sotirios Kentros, and Boyang Wei. Three-party oram for secure computation. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 360–385. Springer, 2015.

- [14] Christopher Fletcher, Muhammad Naveed, Ling Ren, Elaine Shi, and Emil Stefanov. Bucket oram: single online roundtrip, constant bandwidth oblivious ram. Technical report, IACR Cryptology ePrint Archive, Report 2015, 1065, 2015.
- [15] Sanjam Garg, Payman Mohassel, and Charalampos Papamanthou. Tworam: round-optimal oblivious ram with applications to searchable encryption. Technical report, IACR Cryptology ePrint Archive, 2015: 1010, 2015.
- [16] Rosario Gennaro, Yuval Ishai, Eyal Kushilevitz, and Tal Rabin. The round complexity of verifiable secret sharing and secure multicast. In *Proceedings of the thirty-third annual ACM symposium on Theory of computing*, pages 580–589. ACM, 2001.
- [17] Rosario Gennaro, Michael O Rabin, and Tal Rabin. Simplified vss and fast-track multiparty computations with applications to threshold cryptography. In *Proceedings of the seventeenth annual ACM symposium on Principles of distributed computing*, pages 101–111. ACM, 1998.
- [18] Craig Gentry, Kenny A Goldman, Shai Halevi, Charanjit Julta, Mariana Raykova, and Daniel Wichs. Optimizing oram and using it efficiently for secure computation. In *International Symposium on Privacy Enhancing Technologies Symposium*, pages 1–18. Springer, 2013.
- [19] Ian Goldberg. Improving the robustness of private information retrieval. In *2007 IEEE Symposium on Security and Privacy (SP'07)*, pages 131–148. IEEE, 2007.
- [20] Oded Goldreich. Towards a theory of software protection and simulation by oblivious rams. In *Proceedings of the nineteenth annual ACM symposium on Theory of computing*, pages 182–194. ACM, 1987.
- [21] Oded Goldreich and Rafail Ostrovsky. Software protection and simulation on oblivious rams. *Journal of the ACM (JACM)*, 43(3):431–473, 1996.
- [22] Thang Hoang, Attila Altay Yavuz, and Jorge Guajardo. Practical and secure dynamic searchable encryption via oblivious access on distributed data structure. In *Proceedings of the 32nd Annual Conference on Computer Security Applications*, pages 302–313. ACM, 2016.
- [23] Ang Li, Xiaowei Yang, Srikanth Kandula, and Ming Zhang. Cloudcmp: comparing public cloud providers. In *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement*, pages 1–14. ACM, 2010.
- [24] Steve Lu and Rafail Ostrovsky. Distributed oblivious ram for secure two-party computation. In *Theory of Cryptography*, pages 377–396. Springer, 2013.
- [25] 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 *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*, pages 311–324. ACM, 2013.
- [26] Travis Mayberry, Erik-Oliver Blass, and Agnes Hui Chan. Efficient private file retrieval by combining oram and pir. In *NDSS*. Citeseer, 2014.
- [27] Tarik Moataz, Erik-Oliver Blass, and Travis Mayberry. Chf-oram: A constant communication oram without homomorphic encryption.
- [28] Tarik Moataz, Erik-Oliver Blass, and Travis Mayberry. Constant communication oram without encryption. Technical report, IACR Cryptology ePrint Archive, Report 2015/1116, 2015.
- [29] Tarik Moataz, Travis Mayberry, and Erik-Oliver Blass. Constant communication oram with small blocksize. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, pages 862–873. ACM, 2015.
- [30] Muhammad Naveed. The fallacy of composition of oblivious ram and searchable encryption. *IACR Cryptology ePrint Archive*, 2015:668, 2015.

- [31] Pascal Paillier. Public-key cryptosystems based on composite degree residuosity classes. In *International Conference on the Theory and Applications of Cryptographic Techniques*, pages 223–238. Springer, 1999.
- [32] Benny Pinkas and Tzachy Reinman. Oblivious ram revisited. In *Advances in Cryptology–CRYPTO 2010*, pages 502–519. Springer, 2010.
- [33] Ling Ren, Christopher W Fletcher, Albert Kwon, Emil Stefanov, Elaine Shi, Marten van Dijk, and Srinivas Devadas. Ring oram: Closing the gap between small and large client storage oblivious ram. *IACR Cryptology ePrint Archive*, 2014:997, 2014.
- [34] Adi Shamir. How to share a secret. *Communications of the ACM*, 22(11):612–613, 1979.
- [35] Elaine Shi, T-H Hubert Chan, Emil Stefanov, and Mingfei Li. Oblivious ram with $o((\log n)^3)$ worst-case cost. In *Advances in Cryptology–ASIACRYPT 2011*, pages 197–214. Springer, 2011.
- [36] Emil Stefanov, Charalampos Papamanthou, and Elaine Shi. Practical dynamic searchable encryption with small leakage. In *NDSS*, volume 71, pages 72–75, 2014.
- [37] Emil Stefanov and Elaine Shi. Multi-cloud oblivious storage. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*, pages 247–258. ACM, 2013.
- [38] Emil Stefanov and Elaine Shi. Oblivstore: High performance oblivious cloud storage. In *Security and Privacy (SP), 2013 IEEE Symposium on*, pages 253–267. IEEE, 2013.
- [39] Emil Stefanov, Elaine Shi, and Dawn Song. Towards practical oblivious ram. *arXiv preprint arXiv:1106.3652*, 2011.
- [40] Emil Stefanov, Marten Van Dijk, Elaine Shi, Christopher Fletcher, Ling Ren, Xiangyao Yu, and Srinivas Devadas. Path oram: an extremely simple oblivious ram protocol. In *Proceedings of the 2013 ACM SIGSAC conference on Computer and Communications security*, pages 299–310. ACM, 2013.
- [41] Jonathan Trostle and Andy Parrish. Efficient computationally private information retrieval from anonymity or trapdoor groups. In *International Conference on Information Security*, pages 114–128. Springer, 2010.
- [42] Xiao Wang, Hubert Chan, and Elaine Shi. Circuit oram: On tightness of the goldreich-ostrovsky lower bound. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, pages 850–861. ACM, 2015.
- [43] Xiao Shaun Wang, Yan Huang, TH Hubert Chan, Abhi Shelat, and Elaine Shi. Scoram: oblivious ram for secure computation. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, pages 191–202. ACM, 2014.
- [44] Xiao Shaun Wang, Kartik Nayak, Chang Liu, TH Chan, Elaine Shi, Emil Stefanov, and Yan Huang. Oblivious data structures. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, pages 215–226. ACM, 2014.
- [45] Andrew C Yao. Protocols for secure computations. In *23rd Annual Symposium on Foundations of Computer Science, 1982.*, pages 160–164. IEEE, 1982.

Appendix

We present the proof of theorems presented in Section 4.

Proof of Theorem 1. We see that $S^3\text{ORAM}$ is correct iff (i) the $S^3\text{ORAM.Retrieve}$ subroutine returns the correct value of the retrieved block, (ii) the write-to-root operation (step 10, Protocol 2) is consistent, and (iii) the $S^3\text{ORAM.SSS-SMP-TripletEviction}$ subroutine is consistent.

- *Correctness of $S^3\text{ORAM.Retrieve}$.* For each data request x , let b be the block to be retrieved and j be the location of b in its path (i.e., $j := \text{pm}[\text{id}].\text{pldx}$ where id is the identifier of b). So, the share of select query for server \mathcal{S}_i is of form: $\llbracket \mathbf{e} \rrbracket_i^{(t)} = (\llbracket e \rrbracket_1^{(t)}, \dots, \llbracket e \rrbracket_n^{(t)})$, where $n = Z \cdot (H + 1)$ and $e_i = 0$ for $1 \leq i \neq j \leq n$, $e_j = 1$. Let $\llbracket \mathbf{c}_u \rrbracket = (\llbracket c_{u1} \rrbracket, \dots, \llbracket c_{un} \rrbracket)$ be the vector consisting of the share of u -th chunks taken from Z slots in every bucket residing in the read path. For $1 \leq u \leq m$, the answer of each server \mathcal{S}_i is of form:

$$\begin{aligned} \llbracket \mathbf{e} \rrbracket_i^{(t)} \cdot \llbracket \mathbf{c}_u \rrbracket_i^{(t)} &= \sum_{k=1}^n \left(\llbracket e_k \rrbracket^{(t)} \cdot \llbracket c_{u,k} \rrbracket^{(t)} \right) \\ &= \sum_{k=1}^n \llbracket e_k \cdot c_{u,k} \rrbracket^{(2t)} && \text{by Eq. (3)} \\ &= \llbracket c_{u,j} \rrbracket^{(2t)} && \text{by Eq. (1)} \end{aligned}$$

By SSS scheme, at least $2t + 1$ shares are required to recover the secret encrypted by a random $2t$ -degree polynomial. Our system model presented in Section 2 follows this and, therefore, the client always computes the correct value of chunk c_t by $c_t \leftarrow \text{SSS.Recover}(\llbracket c_t \rrbracket_1^{(2t)}, \dots, \llbracket c_t \rrbracket_\ell^{(2t)}, 2t)$. Since all chunks of b are correctly computed, b is properly retrieved with probability 1.

- *Consistency of write-to-root (step 10, Protocol 2):* Lemma 3 implies that the root bucket is empty after eviction. The client writes the retrieved block to an empty slot in the root bucket according to the subsequent order. Since $Z = A$, this ensures that slots containing retrieved blocks are not overwritten before the eviction happens.

- *Consistency of $S^3\text{ORAM.SSS-SMP-TripletEviction}$:* Lemma 3 implies that sibling buckets are empty due to previous evictions and, therefore, they can hold all data moved from source buckets without creating inconsistency. Moving data from source buckets to destination buckets is achieved via matrix products. These computations are correct due to homomorphic properties of two-share addition and multiplication offered by SSS and the SMP protocol which was proven to be correct in [17], respectively. \square

Proof of Theorem 2. Given a request sequence \mathbf{x} of length q , where $x_j = (\text{op}_j, \text{id}_j, \text{data}_j)$ as in Definition 2, let $S^3\text{ORAM}_i(\mathbf{x})$ be the $S^3\text{ORAM}$ client's sequence of interactions with server \mathcal{S}_i including a sequence of retrieval, write-to-root and eviction operations. We have that write-to-root operation is deterministic and is performed after retrieval where the previously retrieved block is written to a publicly known slot in the root bucket (step 10, Protocol 2). The eviction is also deterministic which is performed after every A successive accesses regardless of any data being requested (step 12). Due to the independence between retrieval, write-to-bucket and eviction operations, we consider $S^3\text{ORAM}_i(\mathbf{x})$ to contain separate sequences of these operations observed by \mathcal{S}_i as:

$$S^3\text{ORAM}_i(\mathbf{x}) = \begin{cases} \vec{R}_i(\mathbf{x}) &= (R_i^{(x_1)}, \dots, R_i^{(x_q)}) \\ \vec{W}_i(\vec{\mathbf{x}}) &= (W_i^{(\tilde{x}_1)}, \dots, W_i^{(\tilde{x}_q)}), \\ \vec{E}_i(\vec{\mathbf{x}}) &= (E_i^{(\tilde{x}_1)}, \dots, E_i^{(\tilde{x}_{q/A})}) \end{cases}, \quad (7)$$

where $\vec{W}_i(\tilde{\mathbf{x}})$ and $\vec{E}_i(\tilde{\mathbf{x}})$ denote the deterministic write-to-bucket and eviction sequences, given data access sequence \mathbf{x} , respectively.

Assume that there is a coalition of up to t servers $\{\mathcal{S}_{i \in \mathcal{I}}\}$ sharing their own transcripts with each other. Let $\mathcal{I} \subseteq \{1, \dots, \ell\}$ such that $|\mathcal{I}| \leq t$. The view of $\{\mathcal{S}_{i \in \mathcal{I}}\}$ can be derived from Eq. (7) as:

$$\{\text{S}^3\text{ORAM}_{i \in \mathcal{I}}(\mathbf{x})\} = \begin{cases} \{\vec{R}_{i \in \mathcal{I}}(\mathbf{x})\} &= (\{R_{i \in \mathcal{I}}^{(x_1)}\}, \dots, \{R_{i \in \mathcal{I}}^{(x_q)}\}) \\ \{\vec{W}_{i \in \mathcal{I}}(\tilde{\mathbf{x}})\} &= (\{W_{i \in \mathcal{I}}^{(x_1)}\}, \dots, \{W_{i \in \mathcal{I}}^{(x_q)}\}), \\ \{\vec{E}_{i \in \mathcal{I}}(\tilde{\mathbf{x}})\} &= (\{E_{i \in \mathcal{I}}^{(1)}\}, \dots, \{E_{i \in \mathcal{I}}^{(q/A)}\}) \end{cases}$$

We show that, for any two access sequences \mathbf{x} and \mathbf{x}' of the same length ($|\mathbf{x}| = |\mathbf{x}'|$), the pairs $\langle \{\vec{R}_{i \in \mathcal{I}}(\mathbf{x})\}, \{\vec{W}_{i \in \mathcal{I}}(\tilde{\mathbf{x}})\}, \{\vec{E}_{i \in \mathcal{I}}(\tilde{\mathbf{x}})\} \rangle$ and $\langle \{\vec{R}_{i \in \mathcal{I}}(\mathbf{x}')\}, \{\vec{W}_{i \in \mathcal{I}}(\tilde{\mathbf{x}}')\}, \{\vec{E}_{i \in \mathcal{I}}(\tilde{\mathbf{x}}')\} \rangle$ are identically distributed.

- Retrieval transcripts: For each access request $x_j \in \mathbf{x}$, $\{\mathcal{S}_{i \in \mathcal{I}}\}$ observes a transcript $\{R_{i \in \mathcal{I}}^{(x_j)}\}$ consisting of a retrieval path \mathcal{P}_{x_j} (access pattern) which is identical for all servers (step 4, Subroutine 1) and data generated in SSS-based select scheme (steps 5–8).

The access pattern of S^3ORAM is identical for all other tree-based ORAM schemes. Specifically, each block in S^3ORAM is assigned to a leaf bucket selected randomly and independently from each other. Once a block is accessed, its position is assigned to a new bucket leaf selected randomly and independently. Therefore, access patterns generated by any data request sequences of the same length are statistically indistinguishable.

We next analyze the probability distribution of data observed at the server side in each S^3ORAM retrieval as follows. For each retrieval, the client sends to servers select queries generated by PIR.CreateQuery algorithm. Such queries are SSS shares and, therefore, achieve t -privacy. The inner product is also t -private due to Lemma 1 with addition and partial multiplication homomorphic properties (1) and (3). So, any data generated in S^3ORAM retrievals are identically distributed in the presence of t colluding servers.

By these properties, for any data request sequence \mathbf{x} , the corresponding transcripts (including access patterns) generated in the S^3ORAM retrieval phase are information-theoretically (statistically) indistinguishable from random access sequence in the presence of up to t colluding servers.

- Write-to-root transcripts: Data are written to slots in the root bucket according to subsequent order and, therefore, the access pattern is deterministic. Such written data are SSS-shared with new random polynomials so that they are t -private. Therefore, any data request sequence generates write-to-root transcripts which are identically distributed.

- Eviction transcripts: Since eviction is deterministic which follows publicly-known reverse lexicographical order like in Onion-ORAM (e.g., [12]), the access patterns of $\{E_{i \in \mathcal{I}}^{(j)}\}$ and $\{E_{i \in \mathcal{I}}^{(j')}\}$ are independent to each other for any $(j, j') \in \{0, \dots, 2^H\}$. We next show that data generated in independent evictions are identically distributed.

For each eviction, the client sends H permutation matrices which are SSS-share and, therefore, they are all t -private and uniformly distributed. Data in sibling buckets are t -private and uniformly distributed since they are merely copied from source buckets deterministically (step 20, Subroutine 3). The matrix product computations (step 25) are also t -private according to Lemma 1 with properties (1) and (3). Finally, the SMP protocol ensures that data in destination buckets are t -private and uniformly distributed (step 25) as shown in Lemma 2.

Given two request sequences \mathbf{x}, \mathbf{y} with $|\mathbf{x}| = |\mathbf{y}|$, the corresponding deterministic eviction

sequences observed by $\{\mathcal{S}_{i \in \mathcal{I}}\}$ are:

$$\{\vec{E}_{i \in \mathcal{I}}(\vec{\mathbf{x}})\} = (\{E_{i \in \mathcal{I}}^{(\bar{x}_1)}\}, \dots, \{E_{i \in \mathcal{I}}^{(\bar{x}_{q/A})}\})$$

$$\{\vec{E}_{i \in \mathcal{I}}(\vec{\mathbf{y}})\} = (\{E_{i \in \mathcal{I}}^{(\bar{y}_1)}\}, \dots, \{E_{i \in \mathcal{I}}^{(\bar{y}_{q/A})}\})$$

where $(\bar{x}_j, \bar{y}_j) \in \{0, \dots, H\}$ for $1 \leq j \leq q/A$. Since data yielded in $\{E_{i \in \mathcal{I}}^{(\bar{j}_j)}\}$ and $\{E_{i \in \mathcal{I}}^{(\bar{x}_{j'})}\}$ are identically distributed for all $(j, j') \in \{\bar{x}_1, \dots, \bar{x}_{q/A}\} \cup \{\bar{y}_1, \dots, \bar{y}_{q/A}\}$ as shown above, $\{\vec{E}_{i \in \mathcal{I}}(\vec{\mathbf{x}})\}$ and $\{\vec{E}_{i \in \mathcal{I}}(\vec{\mathbf{y}})\}$ are identically distributed.

• Final indistinguishability argument: Given any data request sequences of the same length, S³ORAM generates (i) access patterns statistically indistinguishable from random request sequence, and (ii) identically (uniform) distributed data in the presence of up to t colluding servers. This indicates that S³ORAM scheme achieves (information-theoretic) t -security according to Definition 2.

□