TWORAM: Efficient Oblivious RAM in Two Rounds with Applications to Searchable Encryption

Sanjam Garg[1], Payman Mohassel[2], and Charalampos Papamanthou[3](✉)

[1] University of California, Berkeley, USA
[2] Visa Research, Foster City, USA
[3] University of Maryland, College Park, USA
cpap@umd.edu

Abstract. We present TWORAM, an asymptotically efficient oblivious RAM (ORAM) protocol providing oblivious access (read and write) of a memory index y in exactly *two* rounds: The client prepares an encrypted query encapsulating y and sends it to the server. The server accesses memory M obliviously and returns encrypted information containing the desired value $\mathsf{M}[y]$. The cost of TWORAM is only a multiplicative factor of security parameter higher than the tree-based ORAM schemes such as the path ORAM scheme of Stefanov et al. [34].

TWORAM gives rise to interesting applications, and in particular to a 4-round symmetric searchable encryption scheme where search is sublinear in the worst case and the search pattern is not leaked—the access pattern can also be concealed assuming the documents are stored in the obliviously accessed memory M.

1 Introduction

Oblivious RAM (ORAM) is a cryptographic primitive for accessing a remote memory M of n entries in a way that memory accesses do not reveal anything about the accessed index $y \in \{1, \ldots, n\}$. Goldreich and Ostrovsky [16] were the first to show that ORAM can be built with $\mathsf{poly}(\log n)$ bandwidth overhead[1], and since then, there has been a fruitful line of research on substantially reducing this overhead [9,29,34,36], in part motivated by the tree ORAM framework proposed by Shi et al. [31]. However, *most* existing practical ORAM protocols are interactive, requiring the client to perform a "download-decrypt-compute-encrypt-upload" operation several times (typically $O(\log n)$ rounds are involved). This can be a bottleneck for applications where low latency is important.

In this paper, we consider the problem of building an efficient round-optimal ORAM scheme. In particular, we propose TWORAM, an ORAM scheme enabling a client to obliviously access a memory location $\mathsf{M}[y]$ in two rounds, where the client sends an encrypted message to the server that encapsulates y, the server

[1] We define *bandwidth overhead* as the number of bits transferred between the client and the server during a single memory access, including the data block.

© International Association for Cryptologic Research 2016
M. Robshaw and J. Katz (Eds.): CRYPTO 2016, Part III, LNCS 9816, pp. 563–592, 2016.
DOI: 10.1007/978-3-662-53015-3_20

performs the oblivious computation, and sends a message back to the client, from which the client can retrieve the desired value $\mathsf{M}[y]$.

TWORAM's worst-case bandwidth overhead is $O(\kappa \cdot p)$ where p is the bandwidth overhead of a tree-based ORAM scheme and κ is the security parameter. For instance, in Path-ORAM [34], it is $p = \log^3 n$ for a block of size $O(\log n)$ bits. In other words, in order to obliviously read a data block of $O(\log n)$ bits using TWORAM, one needs to exchange, *in the worst case*, a $O(\kappa \cdot \log^3 n)$ bits with the server, just in two rounds.

1.1 Existing Round-Optimal ORAM Protocols

Williams and Sion [37] devised a round-optimal ORAM scheme based on a customized garbling scheme and Bloom filters. Lu and Ostrovsky also include an optimized construction for single-round oblivious RAM in their seminal garbled RAM paper [28]. Subsequent to our work, Fletcher et al. [10] also provide single-round ORAM by generalizing the approach of [37] to use a garbling scheme for branching programs. All aforementioned approaches are symmetric-key and are built on top of the hierarchical ORAM framework as introduced by Goldreich and Ostrovsky [16]. Our approach however is based on the tree-based ORAM framework as introduced by Shi et al. [31], yielding worst-case logarithmic costs by construction, thus avoiding involved deamortization procedures. Burst ORAM [21] is also round-optimal, yet it requires linear storage at the client side.

Other less efficient approaches to construct round-optimal ORAM schemes are generic constructions based on garbled RAM [11,12,14]. However, such generic approaches are prohibitively inefficient. For instance, for the non-black-box Garbled RAM approaches [12,14], the bandwidth overhead grows with $\mathsf{poly}(\log n, \kappa, |f|)$, where $|f|$ is the size of the circuit for computing the one-way function f and κ is the security parameter. This leads to inefficient constructions, that are only of theoretical interest. Also, for the black-box Garbled RAM approach [11] the bandwidth overhead grows with $\mathsf{poly}(\log n, \kappa)$, and is independent of $|f|$. However, the construction itself is asymptotically very inefficient. Specifically in [11] the authors do not provide details on how large the involved polynomials are, which will depend on the choice of various parameters. According to our back-of-the-envelope calculation, however, the polynomial is at least $\kappa^5 \cdot \log^7 n$. A key reason for this inefficiency is that they require certain expensive ORAM operations, specifically "eviction," to be performed inside a garbled circuit. We eliminate this source of inefficiency by moving these expensive ORAM operations outside of the garbled circuits.

1.2 TWORAM's Technical Highlights

Our construction is inspired by the ideas from the recent, black-box garbled RAM work by Garg et al. [11]. We specifically use those ideas on top of the tree ORAM algorithms [31]. Our new ideas help avoid certain inefficiencies involved in the original construction of [11], yielding an asymptotically better protocol.

Our first step is to abstract away certain details of eviction-based tree ORAM algorithms, such as Path-ORAM [34], circuit ORAM [36] and Onion ORAM [9]. These algorithms work as follows: The memory M that must be accessed obliviously is stored as a sequence of L trees $T_1, T_2, \ldots, T_L$. The actual data of M is stored encrypted in the tree T_L, while the other trees store *position map* information (also encrypted). Only T_1 is stored on the client side. Roughly speaking, to access an index y in M, the client accesses T_1 and sends a path index p_2 to the server. The server then, successively accesses paths $p_2, p_3, \ldots, p_L$ in $T_2, T_3, \ldots, T_L$. However the paths are accessed adaptively: in order to learn p_i, one needs to first access p_{i-1} in T_{i-1}, and have all the information (also known as buckets) stored in its nodes decrypted. This is where existing approaches require $O(L)$ rounds of interaction: decryption can only take place at the client side, which means all the information on the paths must be communicated back to the client.

TWORAM's Core Idea. In order to avoid the roundtrips described above, we do not use standard encryption. Instead, we hardcode the content of each bucket inside a garbled circuit [38]. In other words, after the trees $T_2, T_3, \ldots, T_L$ are produced, the client generates one garbled circuit per each internal node in each tree. The function of this garbled circuit is very simple: Informally, it takes as input an index x; loops through the blocks $\mathsf{bucket}[i]$ contained in the current bucket until it finds $\mathsf{bucket}[x]$, and returns the index $\pi = \mathsf{bucket}[x]$ of the next path to be followed. Note that the index π is returned in form of a *garbled input* for the next garbled circuit, so that the execution can proceed by the server until T_L is reached, and the final desired value can be returned to the client (see Fig. 3 for a more formal description).

This simplified description ignores some technical hurdles. Firstly, security of the underlying ORAM scheme requires that the location where $\mathsf{bucket}[x]$ is found remains hidden. In particular, the garbled circuit which has the value $\mathsf{bucket}[x]$ inside should not be identifiable by the server. We resolve this issue as follows. For every bucket that the underlying ORAM needs to touch, all the corresponding garbled circuits are executed in a specific order and the value of interest is carried along the way and output only by the final evaluated circuit in that tree.

Secondly, the above approach only works well for a single memory access, since the garbled circuits can only be used once. Fortunately, as we show in the paper, only a logarithmic number of garbled circuits are touched for each access. These circuits can be downloaded by the client who decodes the hardcoded values, performs the eviction strategy locally (on plaintext data), and sends fresh garbled circuits back to the server. This step does not increase the number of rounds (from two to three), since sending the fresh garbled circuits to the server can be "piggybacked" onto the message the client prepares for the next memory access.

Finally, in order to ensure the desired efficiency, and to avoid a blowup of polynomial multiplicative factor in security parameter, we develop optimizations that help ensure that the sizes of the circuits garbled in our construction remain small and proportional to the underlying ORAM.

1.3 Application: 4-Round Searchable Encryption with No Search Pattern Leakage

An SSE scheme allows a client to outsource a database (defined as a set of document/keyword set pairs $\mathsf{DB} = (d_i, W_i)_{i=1}^N$) to a server in an encrypted format, where a search query for w returns d_i where $w \in W_i$.

Several recent works [3,20,26,39] demonstrate attacks against property-preserving encryption schemes (which also enable search on encrypted data), by taking advantage of the leakage associated with these schemes. Thought these attacks do not lead to concrete attacks against existing SSE schemes, they underline the importance of examining the feasibility of solutions that avoid leakage. A natural building block for doing so is ORAM. We use TWORAM to obtain the first constant-round, and asymptotically efficient SSE that can hide search/access patterns.

Our construction combines TWORAM and a non-recursive Path-ORAM (whose position map of the first level is not outsourced) such that searching for w requires (i) a single access on TWORAM; (ii) $|\mathsf{DB}(w)|$ parallel accesses to the non-recursive Path-ORAM (note that an access to a non-recursive ORAM requires only two rounds).

In particular, we use TWORAM to store pairs of the form $(w, (count_w, access_w))$, where w is a keyword, $count_w$ is the number of documents containing w and $access_w$ is the number of times w has been accessed so far. The keyword/document pairs $(w||i, d_i)$ (where d_i is the i-th document containing w) are then stored in the non-recursive Path-ORAM where their position in the Path-ORAM tree (namely the random path they are mapped to) is determined on the fly by using a PRF F as $F_k(w||i, access_w)$ (therefore there is no need to store the position map locally). To search for keyword w, we first access TWORAM to obtain $(count_w, access_w)$ (and increment $access_w$), and then generate all positions to look up in the Path-ORAM using the PRF F. These lookups can be parallelized and updating the paths can be piggybacked to the next search.

The above yields a construction with 4 rounds of interaction. Note that naively using ORAM for SSE would incur $|\mathsf{DB}(w)|$ ORAM accesses which imply *at least* $|\mathsf{DB}(w)|$ roundtrips (depending on the number of rounds of the underlying ORAM). As we said before, our construction *does not leak the search pattern*, by providing randomly generated tokens every time a search is performed. If we choose to store all documents in the obliviously-accessed memory, the access pattern can also be concealed.

1.4 Other Related Work

Oblivious RAM. ORAM protocols with a non-constant number of roundtrips can be categorized into *hierarchical* [17,18,24,27], motivated by the seminal work of Goldreich and Ostrovsky [16], and *tree-based* [9,29,34,36], motivated by the seminal work of Shi et al. [31]. We note however, that, by picking the data block size to be big (e.g., $\sqrt{n}$ bits), the number of rounds in tree-based ORAMs can be made constant, yet bandwidth increases beyond polylogarithmic, so such a parameter selection is not interesting.

Searchable Encryption. Song et al. [32] were the first to explore feasibility of searchable encryption. Since then, many follow-up works have designed new schemes for both static data [4,6,8] and dynamic data [5,15,22,23,33,35]. The security definitions also evolved over time and were eventually established in the work of [6,8]. Unlike our construction, *all these approaches use deterministic tokens*, an therefore leak the search patterns. The only proposed approaches that are constant-round and have randomized tokens (apart from constructing SSE through Garbled RAM) are the ones based on functional encryption [30]. However, such approaches incur a linear search overhead. We also note that one can obtain SSE with no search pattern leakage by using interactive ORAMs such as Path-ORAM [34], or other variants optimized for binary search [13].

Secure Computation for RAM Programs. A recent line of work studies efficient secure two-party computation of RAM programs based on garbled circuits [1,19]. These constructions can also be used to design SSE that hide the search pattern—yet these approaches do not lead to constant-round SSE schemes, requiring the client to perform computation proportional to the size of the search result.

2 Definitions for Garbled Circuits and Oblivious RAM

In this section, we recall definitions and describe building blocks we use in this paper. We use the notation $\langle C', S'\rangle \leftrightarrow \Pi\langle C, S\rangle$ to indicate that a protocol Π is executed between a client with input C and a server with input S. After the execution of the protocol the client receives C' and the server receives S'. For non-interactive protocols, we just use the left arrow notation ($\leftarrow$) instead.

2.1 Garbled Circuits

Garbled circuits were first constructed by Yao [38] (see Lindell and Pinkas [25] and Bellare et al. [2] for a detailed proof and further discussion). A *circuit garbling scheme* is a tuple of PPT algorithms $(\mathsf{GCircuit}, \mathsf{Eval})$, where $\mathsf{GCircuit}$ is the circuit garbling procedure and Eval the corresponding evaluation procedure. More formally:

- $(\tilde{C}, \mathsf{lab}) \leftarrow \mathsf{GCircuit}\,(1^\kappa, C)$: $\mathsf{GCircuit}$ takes as input a security parameter κ, and a Boolean circuit C. This procedure outputs a *garbled circuit* $\tilde{C}$ and *input labels* lab, which is a set of pairs of random strings. Each pair in lab corresponds to every input wire of C (and in particular each element in the pair represents either 0 or 1).
- $y \leftarrow \mathsf{Eval}(\tilde{C}, \mathsf{lab}_x)$: Given a garbled circuit $\tilde{C}$ and *garbled input* lab_x, Eval outputs $y = C(x)$.

Input Labels and Garbled Inputs. For a specific input x, we denote with lab_x the *garbled inputs*, a "projection" of x on the input labels. E.g., for a Boolean circuit of two input bits z and w, it is $\mathsf{lab} = \{(z_0, z_1), (w_0, w_1)\}$, $\mathsf{lab}_{00} = \{z_0, w_0\}$, $\mathsf{lab}_{01} = \{z_0, w_1\}$, etc.

Correctness. Let $(\mathsf{GCircuit}, \mathsf{Eval})$ be a circuit garbling scheme. For correctness, we require that for any circuit C and an input x for C, we have that that $C(x) = \mathsf{Eval}(\tilde{C}, \mathsf{lab}_x)$, where $(\tilde{C}, \mathsf{lab}) \leftarrow \mathsf{GCircuit}\,(1^\kappa, C)$.

Security. Let $(\mathsf{GCircuit}, \mathsf{Eval})$ be a circuit garbling scheme. For security, we require that for any PPT adversary A, there is a PPT simulator Sim such that the following distributions are computationally indistinguishable:

- $\mathrm{Real}_{\mathsf{A}}(\kappa)$: A chooses a circuit C. Experiment runs $(\tilde{C}, \mathsf{lab}) \leftarrow \mathsf{GCircuit}\,(1^\kappa, C)$ and sends $\tilde{C}$ to A. A then chooses an input x. The experiment uses lab and x to derive lab_x and sends lab_x to A. Then it outputs the output of the adversary.
- $\mathrm{Ideal}_{\mathsf{A},\mathsf{Sim}}(\kappa)$: A chooses a circuit C. Experiment runs $(\tilde{C}, \sigma) \leftarrow \mathsf{Sim}(1^\kappa, |C|)$ and sends $\tilde{C}$ to A. A then chooses an input x. The experiment runs $\mathsf{lab}_x \leftarrow \mathsf{Sim}(1^\kappa, \sigma)$ and sends lab_x to A. Then it outputs the output of the adversary.

The above definition guarantees adaptive security, since the adversary gets to choose input x after seeing the garbled circuit $\tilde{C}$. We only know how to instantiate garbling schemes with adaptive security in the random oracle model. In the standard model, existing garbling schemes achieve a weaker static variant of the above definition where the adversary chooses both C and input x at the same time and before receiving $\tilde{C}$.

Concerning complexity, we note that if the cleartext circuit C has $|C|$ gates, the respective garbled circuit has size $O(|C|\kappa)$. This is because every gate in the circuit is typically replaced with a table of four rows, each row storing encryptions of labels (each encryption has κ bits).

2.2 Oblivious RAM

We recall *Oblivious RAM* (ORAM), a notion introduced and first studied by Goldreich and Ostrovsky [16]. ORAM can be thought of as a compiler that encodes the memory into a special format such that accesses on the compiled memory do not reveal the underlying access patterns on the original memory. An ORAM scheme consists of protocols (Setup, ObliviousAccess).

- $\langle \sigma, \mathsf{EM} \rangle \leftrightarrow \textsc{Setup}\langle (1^\kappa, \mathsf{M}), \perp \rangle$: Setup takes as input the security parameter κ and a memory array M and outputs a secret state σ (for the client), and an encrypted memory EM (for the server).
- $\langle (\mathsf{M}[y], \sigma'), \mathsf{EM}' \rangle \leftrightarrow \textsc{ObliviousAccess}\langle (\sigma, y, v), \mathsf{EM} \rangle$: ObliviousAccess is a protocol between the client and the server, where the client's input is the secret state σ, an index y and a value v which is set to null in case the access is a read operation (and not a write). Server's input is the encrypted memory EM. Client's output is $\mathsf{M}[y]$ and an updated secret state σ' and the server's output is an updated encrypted memory EM' where $\mathsf{M}[y] = v$, if $v \neq \mathsf{null}$.

Correctness. Consider the following correctness experiment. Adversary A chooses memory M_0. Consider EM_0 generated with $\langle \sigma_0, \mathsf{EM}_0 \rangle \leftrightarrow \textsc{Setup}\langle (1^\kappa, \mathsf{M}_0), \perp \rangle)$. The adversary then adaptively chooses memory locations

to read and write. Denote the adversary's read/write queries by $(y_1, v_1), \ldots,$ (y_q, v_q) where $v_i = \mathsf{null}$ for read operations. A wins in the correctness game if $\langle (\mathsf{M}_{i-1}[y_i], \sigma_i), \mathsf{EM}_i \rangle$ are not the final outputs of the protocol ObliviousAccess$\langle (\sigma_{i-1}, y_i, v_i), \mathsf{EM}_{i-1} \rangle$ for any $1 \leq i \leq q$, where M_i, EM_i, σ_i are the memory array, the encrypted memory array and the secret state, respectively, after the i-th access operation, and ObliviousAccess is run between an honest client and server. The ORAM scheme is correct if the probability of A in winning the game is negligible in κ.

Security. An ORAM scheme is secure in the semi-honest model if for any PPT adversary A, there exists a PPT simulator Sim such that the following two distributions are computationally indistinguishable.

- $\text{Real}_{\mathsf{A}}(\kappa)$: A chooses M_0. Experiment then runs $\langle \sigma_0, \mathsf{EM}_0 \rangle \leftrightarrow$ Setup $\langle (1^\kappa, \mathsf{M}_0), \perp \rangle$. For $i = 1, \ldots, q$, A makes adaptive read/write queries (y_i, v_i) where $v_i = \mathsf{null}$ on reads, for which the experiment runs the protocol

$$\langle (\mathsf{M}_{i-1}[y_i], \sigma_i), \mathsf{EM}_i \rangle \leftrightarrow \text{ObliviousAccess}\langle (\sigma_{i-1}, y_i, v_i), \mathsf{EM}_{i-1} \rangle.$$

  Denote the full transcript of the above protocol by t_i. Eventually, the experiment outputs $(\mathsf{EM}_q, t_1, \ldots, t_q)$ where q is the total number of read/write queries.
- $\text{Ideal}_{\mathsf{A},\mathsf{Sim}}(\kappa)$: The experiment outputs $(\mathsf{EM}_q, t'_1, \ldots, t'_q) \leftrightarrow \mathsf{Sim}(q, |\mathsf{M}_0|, 1^\kappa)$.

3 TWORAM Construction

Our TWORAM construction uses an abstraction of tree-based ORAM schemes, e.g., Path-ORAM [34]. We start by describing this abstraction informally. Then we show how to turn the interactive Path-ORAM protocol (e.g., the one by Stefanov et al. [34]) into a two-round ORAM protocol, using the abstraction that we present below. We now give some necessary notation that we need for understanding our abstraction.

3.1 Notation

Let $n = 2^L$ be the size of the initial memory that we wish to access obliviously. This memory is denoted by $A_L[1], A_L[2], \ldots, A_L[n]$ where $A_L[i]$ is the i-th *block* of the memory. Given location y that we wish to access, let $y_L, y_{L-1}, \ldots, y_1$ be defined recursively as $y_L = y$ and $y_i = ceil(y_{i+1}/2)$, for all $i = L-1, L-2, \ldots, 1$. For example, for $L = 4$ and $y = 13$, we have

- $y_1 = ceil(ceil(ceil(y/2)/2)/2) = 2$.
- $y_2 = ceil(ceil(y/2)/2) = 4$.
- $y_3 = ceil(y/2) = 7$.
- $y_4 = 13$.

Also define $b_i = 1 - y_i\%2$ to be a bit (namely b_i indicates if y_i is even or not). Finally, on input a value x of $2 \cdot L$ bits, $\mathsf{select}(x, 0)$ selects the first L bits of x, while $\mathsf{select}(x, 1)$ selects the last L bits of x. We note here that both y_i and b_i are functions of y, but we do not indicate this explicitly so that not to clutter notation.

3.2 Path-ORAM Abstraction

We start by describing our abstraction of Path-ORAM construction. In Appendix A we describe formally how this abstraction can be used to implement the interactive Path-ORAM algorithm [34] (with $\log n$ rounds of interaction). We note that the details in Appendix A are provided only for helping better understanding. Our construction can be understood based on just the abstraction defined below.

Roughly speaking, Path-ORAM algorithms encode the original memory A_L in the form of L memories

$$A_L, A_{L-1}, \ldots, A_1,$$

where A_L stores the original data and the remaining memories A_i store information required for accessing data in A_L *obliviously*. Each A_i has 2^i entries, each one storing blocks of $2 \cdot L$ bits (for ease of presentation we assume the block size is $\Theta(\log n)$ but our results apply with other block parameterizations as well). Memories $A_L, A_{L-1}, \ldots, A_2$ are stored in trees $T_L, T_{L-1}, \ldots, T_2$ respectively. The smallest memory A_1 is kept locally by the client. The invariant that is maintained is that any block $A_i[x]$ will reside in *some* leaf-to-root path of tree T_i, and specifically on the path that starts from leaf x_i in T_i. The value x_i itself can be retrieved by accessing A_{i-1}, as we detail in the following.

Reading a Value $A_L[y]$. To read a value $A_L[y]$, one first reads $A_1[y_1]$ from local storage and computes $x_2 \leftarrow \mathsf{select}(A_1[y_1], b_1)$ (recall definitions of y_1 and b_1 from Sect. 3.1). Then one traverses the path starting from leaf x_2 in T_2. This path is denoted with $T_2(x_2)$. Block $A_2[y_2]$ is guaranteed to be on $T_2(x_2)$. Then one computes $x_3 \leftarrow \mathsf{select}(A_2[y_2], b_2)$, and continues in this way. In the end, one will traverse path $T_L(x_L)$ and will eventually retrieve block $A_L[y]$. See Fig. 1.

Updating the Paths. Once the above process finishes, we need to make sure that we do not access the same leaf-to-root paths in case we access $A_L[y]$ again in the future—this would violate obliviousness. Thus, for $i = 2, \ldots, L$, we perform the following tasks:

1. We remove all blocks from $T_i(x_i)$ and copy them into a data structure C_i called *stash*. In our abstraction, stash C_i is viewed as an extension of the root of tree T_i;
2. In the stash C_{i-1}, we set $\mathsf{select}(A_{i-1}[y_{i-1}], b_{i-1}) \leftarrow r_i$, where r_i is a fresh random number in $[1, 2^i]$ that replaces x_i from above. This effectively means that block $A_i[y_i]$ should be reinserted on path $T_i(r_i)$, when eviction from stash C_i takes place;

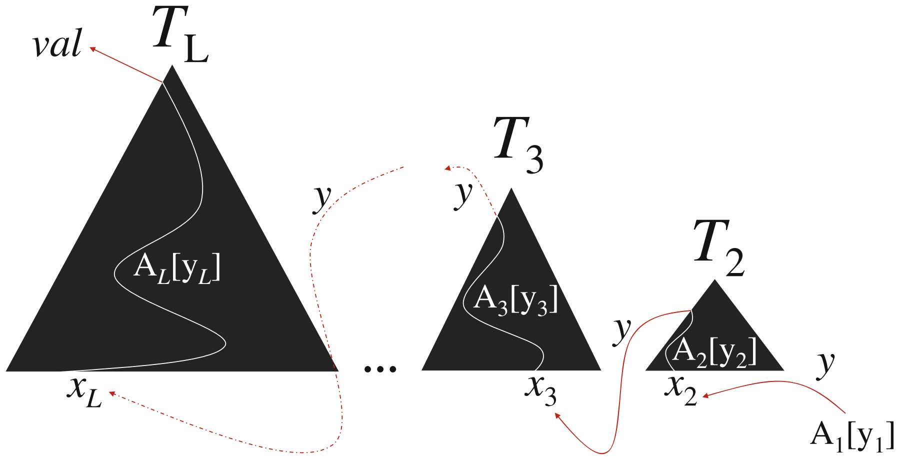

Fig. 1. Our Path-ORAM abstraction for reading a value $val = A_L[y]$. $A_1[y_1]$ is read from local storage and defines x_2. x_2 defines a path p_2 in T_2. By traversing p_2 the algorithm will retrieve $A_2[y_2]$, which will yield x_3, which defines a path p_3 in T_3. Repeating this process yields a path p_L in T_L, traversing which yields the final value $A_L[y_L] = A_L[y]$. Note that y is passed from tree T_{i-1} to tree T_i so that the index y_i (and the bit b_i) can be computed for searching for the right block on path p_i.

3. We evict blocks from stash C_i back to tree $T_i(x_i)$, respecting the new assignments made above.

Syntax. A *Path-ORAM* consists of three procedures (INITIALIZE, EXTRACT, UPDATE) with syntax:

- $\mathcal{T} \leftarrow$ INITIALIZE$(1^\kappa, A_L)$: Given a security parameter κ and memory A_L as input, SETUP outputs a set of $L-1$ trees $\mathcal{T} = \{T_2, T_3, \ldots, T_L\}$ and an array of two entries A_1. A_1 is stored locally with the client and $T_2, \ldots, T_L$ are stored with the server.
- $x_{i+1} \leftarrow$ EXTRACT$(i, y, T_i(x_i))$ for $i = 2, \ldots, L$. Given the tree number i, the final memory location of interest y and a leaf-to-root path $T_i(x_i)$ (that starts from leaf x_i) in tree T_i, EXTRACT outputs an index x_{i+1} to be read in the next tree T_{i+1}. The client can obtain x_2 from local storage as $x_2 \leftarrow \mathsf{select}(A_1[y_1], b_1)$. The obtained value x_2 is sent to the server in order for the server to continue execution. Finally, the server outputs x_{L+1}, which is the desired value $A_L[y]$.

  EXTRACTBUCKET Algorithm. In Path-ORAM [34], internal nodes of the trees store more than one block $(z, A_i[z])$, in the form of *buckets*. We note that EXTRACT can be broken to work on individual buckets along a root-to-leaf path in a tree T_i. In particular, we can define the algorithm $\pi \leftarrow$ EXTRACTBUCKET(i, y, b) where i is the tree of interest, y is the memory location that needs to be accessed, and b is a bucket corresponding to a particular

node on the leaf-to-root path. π will be found at one of the nodes on the leaf-to-root path. Note that the algorithm EXTRACT can be implemented by repeatedly calling EXTRACTBUCKET for every b on $T_i(x_i)$.

- $\{A_1, T_2(x_2), \ldots, T_L(x_L)\} \leftarrow \text{UPDATE}(y, op, val, A_1, T_2(x_2), \ldots, T_L(x_L)$. Procedure UPDATE takes as input the leaf-to-root paths (and local storage A_1) that were traversed during the access and accordingly updates these paths (and local storage A_1). Additionally, UPDATE ensures the new value *val* is written to $A_L[y]$, if operation *op* is a "write" operation.

An implementation of the above abstractions, for Path-ORAM [34], is given in Algorithms 1, 2 and 3 in Appendix A.1. Note that the description of the UPDATE procedure [34] abstracts away the details of the eviction strategy. The SETUP and OBLIVIOUSACCESS protocols of the interactive Path-ORAM *using these abstractions* are given in Figs. 6 and 7 respectively in the Appendix A.2. It is easy to see that the OBLIVIOUSACCESS protocol has $\log n$ rounds of interactions. By the proof of Stefanov et al. [34], we get the following:

Corollary 1. *The protocols* SETUP *and* OBLIVIOUSACCESS *from* Figs. 6 and 7 *respectively in* Appendix A.2 *comprise a secure ORAM scheme (as defined in* Sect. 2.2*) with* $O(\log n)$ *rounds, assuming the encryption scheme used is CPA-secure.*

We recall that the bandwidth overhead for Path-ORAM [34] is $O(\log^3 n)$ bits and the client storage is $O(\log^2 n) \cdot \omega(1)$ bits, for a block size of $2 \cdot L = 2 \cdot \log n$ bits.

3.3 From $\log n$ Rounds to Two Rounds

Existing Path-ORAM protocols implementing our abstraction require $\log n$ rounds (see OBLIVIOUSACCESS protocol in Fig. 7). The main reason for that is the following: In order for the server to determine the index of leaf x_i from which the next path traversal begins, the server needs to access $A_{i-1}[y_{i-1}]$, which is stored *encrypted* at some node on the path starting from leaf x_{i-1} in tree T_{i-1}—see Fig. 1. Therefore the server has to return all encrypted nodes on $T_{i-1}(x_{i-1})$ to the client, who performs the decryption locally, searches for $A_{i-1}[y_{i-1}]$ (via the EXTRACTBUCKET procedure) and returns the value x_i to the server (see Line 10 of the OBLIVIOUSACCESS protocol in Fig. 7).

Our Approach. To overcome this difficulty, we do not encrypt the blocks in the buckets. Instead, for each bucket stored at a tree node u, we prepare a garbled circuit that hardcodes, among other things, the blocks that are contained in the bucket. Subsequently, this garbled circuit executes the EXTRACTBUCKET algorithm on the hardcoded blocks and outputs either $\perp$ or the next leaf index π, depending on whether the search performed by EXTRACTBUCKET was successful or not. The output, whatever that is, is fed as a garbled input to either the left child bucket or the right child bucket (depending on the currently traversed path) or the next root bucket (in case u is a leaf) of u. In this way, by the time the server has executed all the garbled circuits along the currently traversed

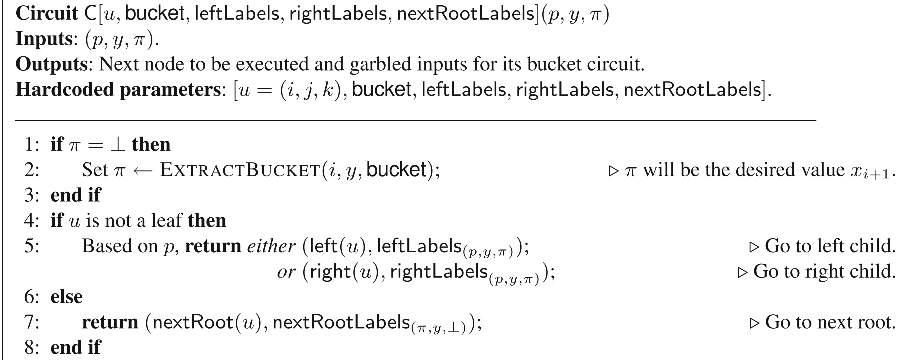

Circuit $\mathsf{C}[u, \mathsf{bucket}, \mathsf{leftLabels}, \mathsf{rightLabels}, \mathsf{nextRootLabels}](p, y, \pi)$
Inputs: (p, y, π).
Outputs: Next node to be executed and garbled inputs for its bucket circuit.
Hardcoded parameters: $[u = (i, j, k), \mathsf{bucket}, \mathsf{leftLabels}, \mathsf{rightLabels}, \mathsf{nextRootLabels}]$.

1: **if** $\pi = \perp$ **then**
2: Set $\pi \leftarrow$ ExtractBucket(i, y, bucket); ▷ π will be the desired value x_{i+1}.
3: **end if**
4: **if** u is not a leaf **then**
5: Based on p, **return** *either* $(\mathsf{left}(u), \mathsf{leftLabels}_{(p,y,\pi)})$; ▷ Go to left child.
or $(\mathsf{right}(u), \mathsf{rightLabels}_{(p,y,\pi)})$; ▷ Go to right child.
6: **else**
7: **return** $(\mathsf{nextRoot}(u), \mathsf{nextRootLabels}_{(\pi,y,\perp)})$; ▷ Go to next root.
8: **end if**

Fig. 2. Formal description of the naive bucket circuit. Notation: Given lab, the set of *input labels* for a garbled circuit, we let lab_a denote the garbled input labels (i.e., the labels taken from lab) corresponding to the input value a.

path, he will be able to pass the index π to the next tree as a garbled input, and continue the execution in the same way without having to interact with the client. Therefore the client can obliviously retrieve his value $A_L[y]$ in only two rounds of communication.

Unfortunately, once these garbled circuits have been consumed, they cannot be used again since this would violate security of garbled circuits. To avoid this problem, the client downloads all the data that was accessed before, decrypts them, runs the Update procedure locally, recomputes the garbled circuits that were consumed before, and stores the new garbled circuits locally. In the next access, these garbled circuits will be sent along with the query. Therefore the total number of communication rounds is equal to two (note that this approach requires permanent client storage—for transient storage, the client will have to send the garbled circuits immediately which would increase the rounds to three). We now continue with describing the bucket circuit that needs to be garbled for our construction.

Naive Bucket Circuit. To help the reader, in Fig. 2 we describe a naive version of our bucket circuit that leads to an inefficient construction. Then we give the full-fledged description of our bucket circuit in Fig. 3. The naive bucket circuit has *inputs*, *outputs* and *hardcoded parameters*, which we detail in the following.

Inputs. The input of the circuit is a triplet consisting of the following information:

1. The index of the leaf p from which the currently explored path begins;
2. The final location to be accessed y;
3. The output from previous bucket π (can be the actual value of the next index to be explored or $\perp$).

Outputs. The outputs of the circuit are the next node to be executed, along with its garbled inputs. For example, if the current node u is not a leaf (see Lines 4 and 5

in Fig. 2), the circuit outputs the garbled inputs of either the left or the right child, whereas if the current node is a leaf (see Lines 6–8 in Fig. 2), the circuit outputs the garbled inputs of the next root to be executed. Note that outputting the garbled inputs is easy, since the bucket circuit hardcodes the input labels of the required circuits. Finally we note that the $\textsc{ExtractBucket}(i, y, \mathsf{bucket})$ algorithm used in Fig. 2 can be found in Appendix A.1—see Algorithm 2.

Hardcoded Parameters. The circuit for node u hardcodes:

1. The node identifier u that consists of a triplet (i, j, k) where
   - $i \in \{2, \ldots, L\}$ is the tree number where node u belongs to;
   - $j \in \{0, \ldots, 2^{i-1}\}$ is the depth of node u;
   - $k \in \{0, \ldots, 2^j - 1\}$ is the oder of node u in the specific level.

   For example, the root of tree T_3 will be denoted $(3, 0, 0)$, while its right child will be $(3, 1, 1)$.
2. The bucket information bucket (i.e., blocks $(x, A_i[x], r)$ contained in node u—recall r is the path index in T_i assigned to $A_i[x]$);
3. The *input labels* leftLabels, rightLabels and nextRootLabels that are used to compute the *garbled inputs* for the next circuit to be executed. Note that leftLabels and rightLabels are used to prepare the next garbled inputs when node u is an internal node (to go either to the left or the right child), while nextRootLabels are used when node u is a leaf (to go to the next root).

Final Bucket Circuit. In the naive circuit presented before, we hardcode the input labels of the root node root of every tree T_i into all the nodes/circuits of tree T_{i-1}. Unfortunately, in every oblivious access, the garbled circuits of all roots are consumed (and therefore root's circuit as well), hence *all* the garbled circuits of tree T_{i-1} will have to be recomputed from scratch. This cost is $O(n)$, thus very inefficient. We would like to mimimize the number of circuits in T_{i-1} that need to be recomputed and ideally make this cost proportional to $O(\log n)$.

To achieve that, we observe that, *instead of hardcoding input labels* nextRootLabels *in the garbled circuit of every node of tree* T_{i-1} *, we can just pass them as garbled inputs to the garbled circuit of every node of tree* T_{i-1}. The final circuit is given in Fig. 3. Note that the only difference of the new circuit from the naive circuit is in the computation of the garbled inputs

$$\mathsf{leftNewLabels}_{(p,y,\pi,\mathsf{nextRootLabels})}$$

and

$$\mathsf{rightNewLabels}_{(p,y,\pi,\mathsf{nextRootLabels})},$$

where nextRootLabels is added in the subscript (see Line 5 of both Figs. 3 and 2), to account for the new input of the new circuit. Note also that we indicate the change in the input format by using "leftNewLabels" instead of just "leftLabels" and "rightNewLabels" instead of just "rightLabels". nextRootLabels have the same meaning in both circuits.

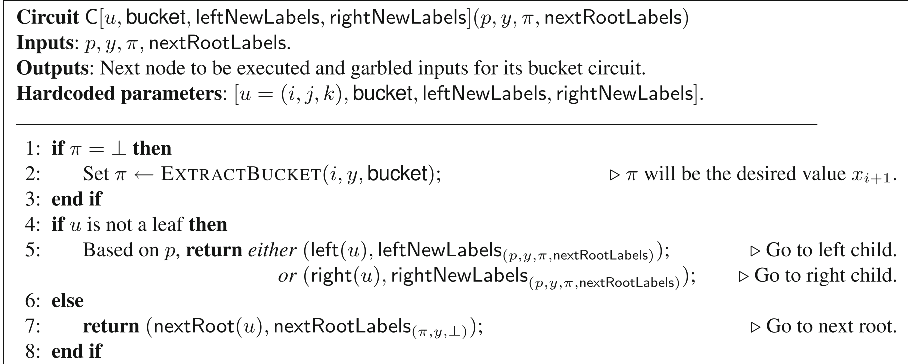

Circuit $\mathsf{C}[u, \mathsf{bucket}, \mathsf{leftNewLabels}, \mathsf{rightNewLabels}](p, y, \pi, \mathsf{nextRootLabels})$
Inputs: $p, y, \pi, \mathsf{nextRootLabels}$.
Outputs: Next node to be executed and garbled inputs for its bucket circuit.
Hardcoded parameters: $[u = (i, j, k), \mathsf{bucket}, \mathsf{leftNewLabels}, \mathsf{rightNewLabels}]$.

1: **if** $\pi = \perp$ **then**
2: Set $\pi \leftarrow$ ExtractBucket(i, y, bucket); ▷ π will be the desired value x_{i+1}.
3: **end if**
4: **if** u is not a leaf **then**
5: Based on p, **return** *either* $(\mathsf{left}(u), \mathsf{leftNewLabels}_{(p,y,\pi,\mathsf{nextRootLabels})})$; ▷ Go to left child.
or $(\mathsf{right}(u), \mathsf{rightNewLabels}_{(p,y,\pi,\mathsf{nextRootLabels})})$; ▷ Go to right child.
6: **else**
7: **return** $(\mathsf{nextRoot}(u), \mathsf{nextRootLabels}_{(\pi,y,\perp)})$; ▷ Go to next root.
8: **end if**

Fig. 3. Formal description of the final bucket circuit.

3.4 Protocols SETUP and OBLIVIOUSACCESS of our construction

We now describe in detail the Setup and ObliviousAccess protocols of TWORAM.

Setup. The Setup protocol is described in Fig. 4. Just like the setup for the interactive ORAM protocol (see Fig. 6 in Appendix A.2), in TWORAM, the client does some computation locally in the beginning (using his secret key) and then outputs some "garbled information" that is being sent to the server. In particular:

1. After producing the trees $T_2, T_3, \ldots, T_L$ using algorithm Initialize, the client prepares the garbled circuit of Fig. 3 for all the nodes $u \in T_i$, for all trees T_i. It is important this computation takes place from the leaves towards the root (that is why we write $j \in \{i-1, \ldots, 0\}$ in Line 2 of Fig. 4), since a garbled circuit of a node u hardcodes the input labels of the garbled circuits of its children—so these need to be readily available by the time u's garbled circuit is computed.
2. Apart from the garbled circuits, the client needs to prepare garbled inputs for the nextRootLabels inputs of all the roots of the trees T_i. These are essentially the β_i's computed in Line 4 of Fig. 4.

ObliviousAccess. The ObliviousAccess protocol of TWORAM is described in Fig. 5. The first step of the protocol is similar to that of the interactive scheme (see Fig. 7 in Appendix), where the client accesses local storage A_1 to compute the path index x_2 that must be traversed in T_2. However, the main difference is that, instead of sending x_2 directly, the client sends the *garbled input that corresponds to* x_2 for the root circuit of tree T_2, denoted with α in Fig. 5.

We note here that α is not enough for the first garbled circuit to start executing, and therefore the server complements this garbled input with β_2 (see Server Line 1), the other half that was sent by the client before and that represents the garbled inputs for the input labels of the next root. Subsequently, the server

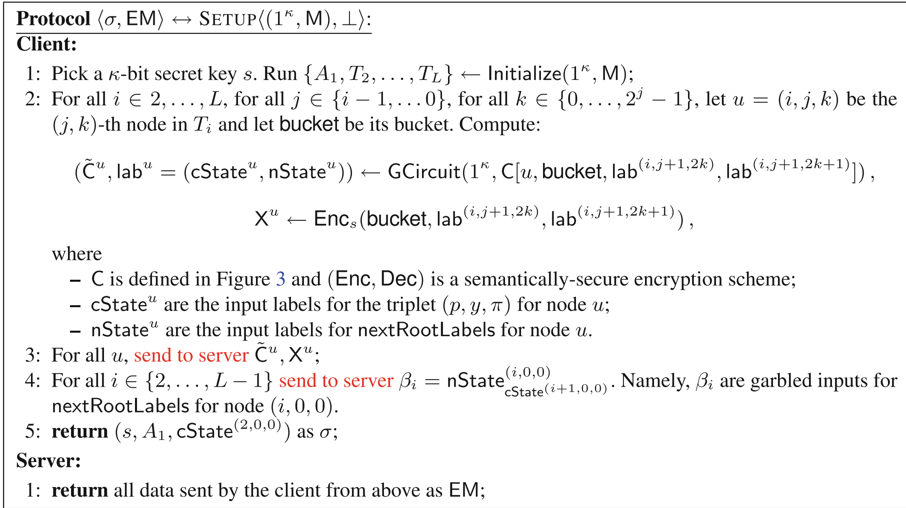

Protocol $\langle \sigma, \mathsf{EM} \rangle \leftrightarrow$ SETUP$\langle (1^\kappa, \mathsf{M}), \perp \rangle$:

Client:

1: Pick a κ-bit secret key s. Run $\{A_1, T_2, \ldots, T_L\} \leftarrow \mathsf{Initialize}(1^\kappa, \mathsf{M})$;

2: For all $i \in 2, \ldots, L$, for all $j \in \{i-1, \ldots 0\}$, for all $k \in \{0, \ldots, 2^j - 1\}$, let $u = (i, j, k)$ be the (j, k)-th node in T_i and let bucket be its bucket. Compute:

$$(\tilde{\mathsf{C}}^u, \mathsf{lab}^u = (\mathsf{cState}^u, \mathsf{nState}^u)) \leftarrow \mathsf{GCircuit}(1^\kappa, \mathsf{C}[u, \mathsf{bucket}, \mathsf{lab}^{(i,j+1,2k)}, \mathsf{lab}^{(i,j+1,2k+1)}]),$$

$$\mathsf{X}^u \leftarrow \mathsf{Enc}_s(\mathsf{bucket}, \mathsf{lab}^{(i,j+1,2k)}, \mathsf{lab}^{(i,j+1,2k+1)}),$$

where
- C is defined in Figure 3 and $(\mathsf{Enc}, \mathsf{Dec})$ is a semantically-secure encryption scheme;
- cState^u are the input labels for the triplet (p, y, π) for node u;
- nState^u are the input labels for $\mathsf{nextRootLabels}$ for node u.

3: For all u, send to server $\tilde{\mathsf{C}}^u, \mathsf{X}^u$;

4: For all $i \in \{2, \ldots, L-1\}$ send to server $\beta_i = \mathsf{nState}^{(i,0,0)}_{\mathsf{cState}^{(i+1,0,0)}}$. Namely, β_i are garbled inputs for $\mathsf{nextRootLabels}$ for node $(i, 0, 0)$.

5: **return** $(s, A_1, \mathsf{cState}^{(2,0,0)})$ as σ;

Server:

1: **return** all data sent by the client from above as EM;

Fig. 4. SETUP protocol for TWORAM.

starts executing the garbled circuits one-by-one, using the outputs of the first circuit, as garbled inputs to the second one, and so on. Eventually, the clients reads and decrypts all paths $T_i(x_i)$, retrieving the desired value (see Client Line 2). Finally, the client runs the UPDATE, re-garbles the circuits that got consumed and waits until the next query to send them back. We can now state the main result of our paper.

Theorem 1. *The protocols* SETUP *and* OBLIVIOUSACCESS *from* Figs. 4 and 5 *respectively comprise a two-round secure ORAM scheme (as defined in* Sect. 2.2*), assuming the garbling scheme used is secure (as defined in* Sect. 2.1*) and the encryption scheme used is CPA-secure.*

The proof of the above theorem can be found in Appendix A.3. Concerning complexity, it is clear that the only overhead that we are adding on Path-ORAM [34] is a garbled circuit per bucket—this adds a multiplicative security parameter factor on all the complexity measures of Path-ORAM. E.g., the bandwidth overhead of our construction is $O(\kappa \cdot \log^3 n)$ bits (for blocks of $2 \log n$ bits).

3.5 Optimizations

Recall that in the garbling procedure of a circuit C, one has the following choices: (i) either to garble C in a way that during evaluation of the garbled circuit on x the output is the cleartext value $C(x)$; (ii) or to garble C in a way that during evaluation of the garbled circuit on x the output is the garbled labels corresponding to the value $C(x)$. We now describe an optimization for a specific circuit C that we will be using in our construction that uses the above observation.

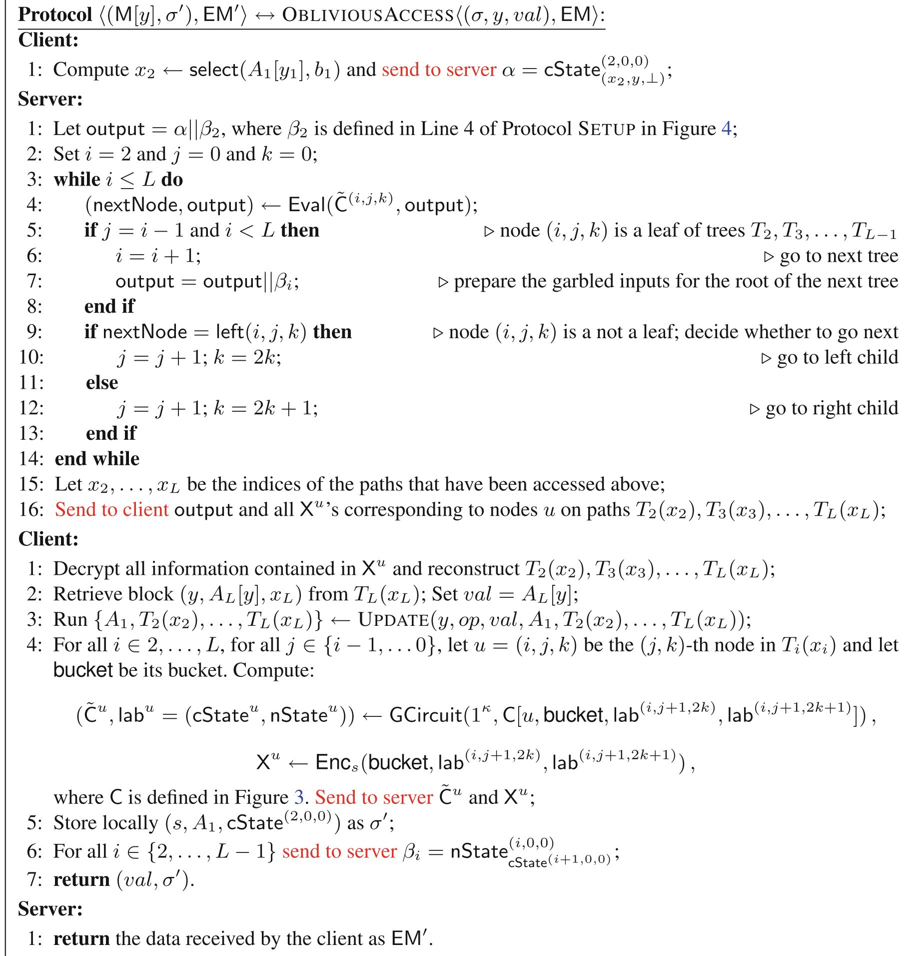

Protocol $\langle(\mathsf{M}[y], \sigma'), \mathsf{EM}'\rangle \leftrightarrow$ OBLIVIOUSACCESS$\langle(\sigma, y, val), \mathsf{EM}\rangle$:

Client:

1: Compute $x_2 \leftarrow \mathsf{select}(A_1[y_1], b_1)$ and send to server $\alpha = \mathsf{cState}^{(2,0,0)}_{(x_2,y,\perp)}$;

Server:

1: Let output $= \alpha || \beta_2$, where β_2 is defined in Line 4 of Protocol SETUP in Figure 4;
2: Set $i = 2$ and $j = 0$ and $k = 0$;
3: **while** $i \leq L$ **do**
4: (nextNode, output) $\leftarrow \mathsf{Eval}(\tilde{\mathsf{C}}^{(i,j,k)}, \mathsf{output})$;
5: **if** $j = i - 1$ and $i < L$ **then** ▷ node (i, j, k) is a leaf of trees $T_2, T_3, \ldots, T_{L-1}$
6: $i = i + 1$; ▷ go to next tree
7: output $=$ output$||\beta_i$; ▷ prepare the garbled inputs for the root of the next tree
8: **end if**
9: **if** nextNode $= \mathsf{left}(i, j, k)$ **then** ▷ node (i, j, k) is a not a leaf; decide whether to go next
10: $j = j + 1$; $k = 2k$; ▷ go to left child
11: **else**
12: $j = j + 1$; $k = 2k + 1$; ▷ go to right child
13: **end if**
14: **end while**
15: Let $x_2, \ldots, x_L$ be the indices of the paths that have been accessed above;
16: Send to client output and all X^u's corresponding to nodes u on paths $T_2(x_2), T_3(x_3), \ldots, T_L(x_L)$;

Client:

1: Decrypt all information contained in X^u and reconstruct $T_2(x_2), T_3(x_3), \ldots, T_L(x_L)$;
2: Retrieve block $(y, A_L[y], x_L)$ from $T_L(x_L)$; Set $val = A_L[y]$;
3: Run $\{A_1, T_2(x_2), \ldots, T_L(x_L)\} \leftarrow$ UPDATE$(y, op, val, A_1, T_2(x_2), \ldots, T_L(x_L))$;
4: For all $i \in 2, \ldots, L$, for all $j \in \{i - 1, \ldots 0\}$, let $u = (i, j, k)$ be the (j, k)-th node in $T_i(x_i)$ and let bucket be its bucket. Compute:

$$(\tilde{\mathsf{C}}^u, \mathsf{lab}^u = (\mathsf{cState}^u, \mathsf{nState}^u)) \leftarrow \mathsf{GCircuit}(1^\kappa, \mathsf{C}[u, \mathsf{bucket}, \mathsf{lab}^{(i,j+1,2k)}, \mathsf{lab}^{(i,j+1,2k+1)}]),$$

$$\mathsf{X}^u \leftarrow \mathsf{Enc}_s(\mathsf{bucket}, \mathsf{lab}^{(i,j+1,2k)}, \mathsf{lab}^{(i,j+1,2k+1)}),$$

where C is defined in Figure 3. Send to server $\tilde{\mathsf{C}}^u$ and X^u;
5: Store locally $(s, A_1, \mathsf{cState}^{(2,0,0)})$ as σ';
6: For all $i \in \{2, \ldots, L - 1\}$ send to server $\beta_i = \mathsf{nState}^{(i,0,0)}_{\mathsf{cState}^{(i+1,0,0)}}$;
7: **return** (val, σ').

Server:

1: **return** the data received by the client as EM'.

Fig. 5. OBLIVIOUSACCESS protocol for TWORAM.

General Optimization. Consider a circuit that performs the following task: It hardcodes two k-bit strings s_0 and s_1, takes an input a bit b and outputs s_b. This cleartext circuit has size $O(k)$, so the garbled circuit for that will have size $O(k^2)$. To improve upon that we consider a circuit C' that takes as input bit b and outputs the same bit b! This cleartext circuit has size $O(1)$. However, to make sure that the output of the garbled version of C' is always s_b, we garble C' by outputting the garbled label corresponding to b, namely s_b (i.e., using (ii) from above). In particular, during the garbling procedure we use s_0 as the garbled label output for output $b = 0$ and we use s_1 as the garbled label output for the output $b = 1$. Note that the size of the new garbled circuit has size $O(k)$, yet it has exactly the same I/O behavior with the garbling of C, which has size $O(k^2)$.

- **Improving cState —Not Hard-Coding Input Labels Inside the Bucket Circuit.** In the construction we described, we include the input labels $\mathsf{leftLabels}, \mathsf{rightLabels}$ in the circuit $\mathsf{C}[u, \mathsf{bucket}, \mathsf{leftLabels}, \mathsf{rightLabels}]$. Consequently, the size of the ungarbled version of this circuit grows with the size of $\mathsf{leftLabels}$ and $\mathsf{rightLabels}$ which is $\kappa \cdot |\mathsf{cState}|$. We can easily use the general optimization described above, for each bit of $|\mathsf{cState}|$, to make the size of the ungarbled version of our circuit only grow with $|\mathsf{cState}|$.
- **Improving nState —Input Labels Passing.** In the construction described previously, for each tree, an input value nState is passed from the root to a leaf node in the tree. However this value is used only at the leaf node. Recall that the nState value passed from the root to a leaf garbled circuits in the tree T_i is exactly the value $\mathsf{cState}^{i+1,0,0}$, the input labels of the root garbled circuit of the tree T_{i+1}. Since each ungarbled circuit gets this value as input, therefore each of one of them needs to grow with $\kappa \cdot |\mathsf{cState}|$.[2] We will now describe an optimization such that the size of the garbled version, rather than the clear version, grows linearly in $\kappa \cdot |\mathsf{cState}|$.

  Note that in our construction the value $\mathsf{cState}^{i+1,0,0}$ is not used at all in the intermediate circuits as it gets passed along the garbled circuits for tree T_i. In order to avoid this wastefulness, for all nodes $i \in \{1, \ldots, L\}, j \in [i], k \in [2^j]$ we sample a value $r^{(i,j,k)}$ of length $\kappa \cdot |\mathsf{cState}|$ and hardcode the values $r^{(i,j,k)} \oplus r^{(i,j+1,2k)}$ and $r^{(i,j,k)} \oplus r^{(i,j+1,2k+1)}$ inside the garbed circuit $\tilde{C}^{i,j,k}$ which output the first of two values if the execution goes left and the second if the execution goes right. Note that a garbled circuits grows only additively in $\kappa \cdot |\mathsf{cState}|$ because of this change. This follows by using the first optimization. Additionally, we include the value $\mathsf{cState}^{i+1,0,0} \oplus r^{(i,0,0)}$ with the root node of the tree T_i. The leaf garbled circuit $(i, i-1, k)$ in tree T_i is constructed assuming $r^{(i,i-1,k)}$ is the sequence of input labels for the root garbled circuit of the tree T_{i+1}.[3] Let $\alpha_0, \ldots \alpha_{i-1}$ be the strings output during the root to a leaf traversal in tree T_i. Now observe that $\mathsf{cState}^{i+1,0,0} \oplus r^{(i,0,0)} \oplus_{j \in [i]} \alpha_j$ is precisely $\mathsf{cState}^{i+1,0,0} \oplus r^{(i,i-1,k)}$ where k is the leaf node in the traversed path. At this point it is easy to see that given the output of the leaf grabled circuit for tree T_i one can compute the required input labels for the root of tree T_{i+1}.

  The update mechanism in our construction can be easily adapted to work with this change. Here note that we would now include the values $r^{(i,j,k)}, r^{(i,j+1,2k)}$ and $r^{(i,j+1,2k+1)}$ in the ciphertext $X^{(i,j,k)}$. Also note that we will use fresh $r^{(\cdot,\cdot,\cdot)}$ values whenever a fresh garbled circuit for a node is generated. The security argument now additionally uses the fact that the outputs generated by garbled circuits in two separate root to leaf traversals depend on completely independent $r^{(\cdot,\cdot,\cdot)}$ values.

  Note that the above modification leaks what value is passed by the executed leaf garbled circuit in tree T_i to the root garbled circuit in tree T_{i+1}. This can be deduced based on what bit values of $\mathsf{cState}^{i+1,0,0} \oplus r^{(i,0,0)}$ are revealed.

[2] This efficiency is achieved when the first optimization is used.

[3] Note that here the first optimization allows us to ensure that the size of the garbled leaf circuit, rather than the clear leaf circuit, grows with the length of $r^{(i,i-1,k)}$ as these hard-codings are performed.

This can be tackled by randomly permuting the labels in $\mathsf{cState}^{i+1,0,0}$ and passing the information on this permutations along with in the tree to leaf garbled circuits. Note that the size of this information is small.

Taken together these two optimizations reduce the size of each garbled circuit to $O(\kappa \cdot (|\mathsf{bucket}| + |\mathsf{cState}|))$. Since $|\mathsf{bucket}| > |\mathsf{cState}|$ this expression reduces to $O(\kappa \cdot |\mathsf{bucket}|)$. This implies that the overhead of our construction is just κ times the overhead of the underlying Path ORAM scheme.

4 Searchable Encryption Construction Using TWORAM

The natural way of designing an SSE scheme that does not leak the search and access patterns using an ORAM scheme is to first use a data structure for storing keyword-document pairs, setup the data structure in memory using an ORAM setup and then read/write from it using ORAM operations. Since ORAM hides the read/write access patterns, but it does not hide the number of memory accesses, one needs to ensure that the number of memory accesses for each operation is also data-independent. Fortunately, this can be achieved by not letting the key used for the hash table be the output of a pseudorandom function applied to the keyword w, and not the keyword w itself.

We start by giving some definitions and then describe constructions that can be instantiated using any ORAM scheme. We then show how to obtain a significantly more efficient instantiation using a combination of TWORAM and a non-recursive Path-ORAM scheme.

4.1 Hash Table Definition

A hash table is a data structure commonly used for mapping keys to values [7]. It often uses a hash function h that maps a key to an index (or a set of indices) in a memory array M where the value associated with the key may be found. In particular, h takes as input a keyword key and outputs a set of indices $i_1, \ldots, i_c$ where c is a parameter. The value associated with key is in one of the locations $\mathsf{M}[i_1], \ldots \mathsf{M}[i_c]$. The keyword is not in the table if it is not in one of those locations. Similarly, to write a new $(key, value)$ pair into the table, $(key, value)$ is written into the first empty location among $i_1, \ldots, i_c$. More formally, we define a hash table $H = (\mathsf{hsetup}, \mathsf{hlookup}, \mathsf{hwrite})$ using a tuple of algorithms and a parameter c denoting an upper bound on the number of locations to search.

- $(h, \mathsf{M}) \leftarrow \mathsf{hsetup}(S, size)$: hsetup takes as input an initial set S of keyword-value pairs and a maximum table size $size$ and outputs a hash function h and a memory array M.
- $value \leftarrow \mathsf{hlookup}(key)$: hlookup computes $\{i_1, \ldots, i_c\} \leftarrow h(key)$, looks for a key-value pair $(key, \cdot)$ in $\mathsf{M}[i_1], \ldots, \mathsf{M}[i_c]$. If such a pair is found it returns the second component of the pair (i.e., the value), else it returns $\perp$.
- $\mathsf{M} \leftarrow \mathsf{hwrite}(key, value)$: hwrite computes $i_1, \ldots, i_c \leftrightarrow h(key)$, if $(key, value)$ already exists in one of those indices in M it does nothing, else it stores $(key, value)$ in the first empty index.

4.2 Searchable Encryption Definition

A database D is a set of document/keyword-set pair

$$\mathsf{DB} = (d_i, W_i)_{i=1}^N.$$

Let $W = \cup_{i=1}^N W_i$ be the universe of keywords. A keyword search query for w should return all d_i where $w \in W_i$. We denote this subset of DB by $\mathsf{DB}(w)$. A searchable symmetric encryption scheme consists of protocols SSESetup, SSESearch and SSEAdd. The following formalization first appeared in [6,8].

- $\langle \sigma, \mathsf{EDB}\rangle \leftrightarrow \text{SSESetup}\langle(1^\kappa, \mathsf{DB}), \perp\rangle$: SSESetup takes as client's input database DB and outputs a secret state σ (for the client), and an encrypted database EDB which is outsourced to the server.
- $\langle(\mathsf{DB}(w), \sigma'), \mathsf{EDB}'\rangle \leftrightarrow \text{SSESearch}\langle(\sigma, w), \mathsf{EDB}\rangle$: SSESearch is a protocol between the client and the server, where client's input is the secret state σ and the keyword w he is searching for. Server's input is the encrypted database EDB. Client's output is the set of documents containing w, i.e. $\mathsf{DB}(w)$ as well an updated secret state σ' and the server obtains an updated encrypted database EDB'.
- $\langle \sigma', \mathsf{EDB}'\rangle \leftrightarrow \text{SSEAdd}\langle(\sigma, d), \mathsf{EDB}\rangle$: SSEAdd is a protocol between the client and the server, where client's input is the secret state σ and a document d to be inserted into the database. Server's input is the encrypted database EDB. Client's output is an updated secret state σ' and the server's output is an updated encrypted database EDB' which now contains the new document d.

Correctness. Consider the following correctness experiment. An adversary A chooses a database DB_0. Consider the encrypted database EDB_0 generated using SSESetup (i.e., $\langle \sigma_0, \mathsf{EDB}_0\rangle \leftrightarrow \text{SSESetup}\langle(1^\kappa, \mathsf{DB}_0), \perp\rangle$). The adversary then adaptively chooses keywords to search and documents to add to the database, and the respective protocols SSESearch and SSEAdd are run between an honest client and server, outputting the updated EDB, DB and σ. Denote the operations chosen by the adversary with $w_1, \ldots, w_q$. A wins in the correctness game if for some search query w_i it is

$$\langle(\mathsf{DB}_i(w_i), \sigma_i), \mathsf{EDB}_i\rangle \neq \text{SSESearch}\langle(\sigma_{i-1}, w_i), \mathsf{EDB}_{i-1}\rangle,$$

where $\mathsf{DB}_i, \mathsf{EDB}_i$ are the database and encrypted database, respectively, after the i-th search. The SSE scheme is correct if the probability of A winning the game is negligible in κ.

Security. We discuss security in the semi-honest model. It is parametrized by a leakage function $\mathcal{L}$, which explains what the adversary (the server) learns about the database and the search and update queries, while interacting with a secure SSE scheme. A SSE scheme is $\mathcal{L}$-secure if for any PPT adversary A, there exist a simulator Sim such that the following two distributions are computationally indistinguishable.

– $\text{Real}_A(\kappa)$: A chooses DB_0. The experiment then runs

$$\langle \sigma_0, \mathsf{EDB}_0 \rangle \leftrightarrow \text{SSESetup}\langle (1^\kappa, \mathsf{DB}_0), \perp \rangle.$$

A then adaptively makes search queries w_i, which the experiment answers by running the protocol $\langle \mathsf{DB}_{i-1}(w_i), \sigma_i \rangle \leftrightarrow \text{SSESearch}\langle (\sigma_{i-1}, w_i), \mathsf{EDB}_{i-1} \rangle$. Denote the full transcripts of the protocol by t_i and with EDB' the final encrypted database. Add queries are handled in a similar way. Eventually, the experiment outputs

$$(\mathsf{EDB}, t_1, \ldots, t_q),$$

where q is the total number of search/add queries made by A.

– $\text{Ideal}_{\mathsf{A},\mathsf{Sim},\mathcal{L}}(\kappa)$: A chooses DB_0. The experiment runs

$$(st_0, \mathsf{EDB}_0) \leftrightarrow \mathsf{Sim}(\mathcal{L}(\mathsf{DB}_0)),$$

where st_0 is the initial state of the simulator. On input any search query w_i from A, the experiment adds $(w_i, search)$ to the history H, and on an add query d_i it adds (d_i, add) to H. It then runs $(t_i, st_i) \leftrightarrow \mathsf{Sim}(st_{i-1}, \mathcal{L}(\mathsf{DB}_{i-1}, H))$. Eventually, the experiment outputs $(\mathsf{EDB}', t_1, \ldots, t_q)$ where q is the total number of search/add queries made by A.

Leakage. The level of security one obtains from a SSE scheme depends on the leakage function $\mathcal{L}$. Ideally $\mathcal{L}$ should only output the total number $\sum_{w \in W} |\mathsf{DB}(w)|$ of (w, d) pairs, the total number of unique keywords $|W|$ and $|\mathsf{DB}(w)|$ for any searched keyword w. Achieving this level of security is only possible if the SSESearch operation outputs the documents themselves to the client. If instead (as is common for applications with large document sizes), it returns document identifiers which the client then uses to retrieve the actual documents, any SSE protocol would also leak the access pattern.

4.3 SSE from any ORAM

First Approach. The common way of storing a database of documents in a hash table is to insert a key-value pair (w, d) into the table for any keyword w in a document d. Searching for a document with keyword w then reduces to looking up w in the table. If there is more than one document containing a keyword w, a natural solution is to create a bucket B_w storing all the documents containing w and storing the bucket in position pt_w of an array A. One then inserts (w, pt_w) in a hash table. Now, to search for a keyword w, we first look up (w, pt_w), and then access $\mathsf{A}[pt_w]$ to obtain the bucket B_w of all the desired documents. A subtle issue is that the distribution of bucket sizes would leak information about the database even before any keyword is searched. As a result, for this approach to be fully-secure, one needs to pad each bucket to an upperbound on the number of searchable documents per keyword.

Next we describe the SSE scheme more formally. Given a hash table $H = (\mathsf{hsetup}, \mathsf{hlookup}, \mathsf{hwrite})$, and an ORAM scheme $ORAM = (\text{Setup}, \text{ObliviousAccess})$, we construct an SSE scheme $(\text{SSESetup}, \text{SSESearch}, \text{SSEAdd})$ as follows.

1. $\langle \sigma, \mathsf{EDB} \rangle \leftrightarrow$ SSESETUP$\langle (1^\kappa, max, \mathsf{DB}), \perp \rangle$: Given an initial set of documents DB, client lets S be the set of key-value pairs (w, pt_w) where pt_w is an index to an array of buckets A such that $\mathsf{A}[pt_w]$ stores the bucket of all documents in DB containing w. Each bucket is padded to the maximum size max with dummy documents.

   Client first runs $\mathsf{hsetup}(S, size)$ to obtain (h, M). $size$ is the maximum size of hash table H. Then client and server run $\langle \sigma_1, \mathsf{EM} \rangle \leftrightarrow$ SETUP$\langle (1^\kappa, \mathsf{M}), \perp \rangle$. Cleint and server also run $\langle \sigma_2, \mathsf{EA} \rangle \leftrightarrow$ SETUP$\langle (1^\kappa, \mathsf{A}), \perp \rangle$

   Note that server's output is $\mathsf{EDB} = (\mathsf{EM}, \mathsf{EA})$ and client's output is $\sigma = (\sigma_1, h, \sigma_2)$.
2. SSESEARCH$\langle (\sigma, w), \mathsf{EDB} \rangle$: Client computes $i_1, \ldots, i_c \leftarrow h(w)$. Then, client and server run OBLIVIOUSACCESS$\langle ((\sigma_1, i_j, \mathsf{null}), \mathsf{EM} \rangle$ for $j \in \{1, \ldots, c\}$ for client to obtain $\mathsf{M}[i_j]$. If client does not find (w, pt_w) in one of the retrieved locations it lets $pt_w = 0$, corresponding to a dummy access to the index 0 in A.

   Client and server then run OBLIVIOUSACCESS$\langle (\sigma_2, pt_w, \mathsf{null}), \mathsf{EA} \rangle)$ for client to obtain the bucket B_w stored in $\mathsf{A}[pt_w]$. Client outputs all the non-dummy documents in B_w.
3. SSEADD$\langle (\sigma, d), \mathsf{EDB} \rangle$: For every w in d, client computes $i_1, \ldots, i_c \leftarrow h(w)$ and client and server run OBLIVIOUSACCESS$\langle (\sigma_1, i_j, \mathsf{null}), \mathsf{EM} \rangle$ for $j \in \{1, \ldots, c\}$ for client to obtain $\mathsf{M}[i_j]$. If (w, pt_w) is in the retrieved locations let i_j^* be the location it was found at. If not, let pt_w be the first empty location in A, and let $i*_j$ be the first empty location from the retrieved ones in M. Client and server run OBLIVIOUSACCESS$\langle (\sigma_1, i_j^*, (w, pt_w)), \mathsf{EM} \rangle$.

   Client and server run OBLIVIOUSACCESS$\langle (\sigma_2, pt_w, \mathsf{null}), \mathsf{EA} \rangle$ to retrieve $\mathsf{A}[pt_w]$. Let B_w be the retrieved bucket. Client inserts d in the first dummy entry of B_w, denoting the new bucket by B'_w. Client and server run

$$\text{OBLIVIOUSACCESS}\langle (\sigma_2, pt_w, B'_w), \mathsf{EA} \rangle.$$

The main disadvantage of the above construction is that we need to anticipate an upper bound on the bucket sizes, and pad all buckets to that size. Given that in practice there are often keywords that appear in a large number of documents, and keywords that only appear in a few, the padding will lead to inefficiency. Our next solution addresses this issue but instead has a higher round complexity.

Second Approach. Instead of storing all documents matching a keyword w in one bucket, we store each of them separately in the hash table, using a different keyword. In particular, we can store the key-value pair $(w||i, d)$ in the hash table for the ith document d containing w. This works fine except that it requires looking up $w||count$ for an incremental counter $count$ until the keyword is no longer found in the table.

To make this approach cleaner and the write operations more efficient, we maintain two hash tables, one for storing the counter representing the number of documents containing the keyword, and one storing the incremental key-value pairs as described above. To lookup a keyword w, one first looks up the counter $count$ in the first table and then makes $count$ lookup queries to the second table.

We now describe the above SSE scheme in more detail. Given a hash table $H = (\mathsf{hsetup}, \mathsf{hlookup}, \mathsf{hwrite})$ and a scheme $ORAM = (\textsc{Setup}, \textsc{ObliviousAccess})$, we construct an SSE scheme $(\text{SSE}\textsc{Setup}, \text{SSE}\textsc{Search}, \text{SSE}\textsc{Add})$ as follows:

1. $\langle \sigma, \mathsf{EDB} \rangle \leftrightarrow \text{SSE}\textsc{Setup}\langle (1^\kappa, \mathsf{DB}), \perp \rangle$: Given an initial set of documents DB. Let S_1 be the set of $(w, count_w)$ pairs and S_2 be the set of key-value pairs $(w||i, d_i)$ for $1 \leq i \leq count_w$ where $count_w$ is the number of documents containing w, and d_i denotes the ith document in DB containing w.

   Cleint runs $\mathsf{hsetup}(S_i, size_i)$ to obtain (h_i, M_i). $size_i$ is the maximum size of the hash table H_i. Then client and server run $\langle \sigma_i, \mathsf{EM}_i \rangle \leftrightarrow \textsc{Setup}\langle (1^\kappa, \mathsf{M}_i), \perp \rangle$. Note that server's output is $\mathsf{EDB} = (\mathsf{EM}_1, \mathsf{EM}_2)$ and client's output is $\sigma = (\sigma_1, \sigma_2, h_1, h_2)$.
2. $\text{SSE}\textsc{Search}\langle (\sigma, w), \mathsf{EDB} \rangle$: Client computes $i_1, \ldots, i_c \leftarrow h_1(w)$ and client and server run $\textsc{ObliviousAccess}\langle (\sigma_1, i_j, \mathsf{null}), \mathsf{EM}_1 \rangle \rangle$ for $j \in \{1, \ldots, c\}$ for client to obtain $(w, count_w)$ among the retrieved locations. If such a pair is not found, client lets $count_w = 0$.

   For $1 \leq k \leq count_w$, client computes $i_1^k, \ldots, i_c^k \leftarrow h_2(w||k)$ and client and server run $\textsc{ObliviousAccess}\langle (\sigma_2, i_j^k, \mathsf{null}), \mathsf{EM}_2 \rangle \rangle$ for $j \in \{1, \ldots, c\}$ for client to obtain $\mathsf{M}_2[i_j^k]$. Client outputs d for all d where $(w||k, d)$ is in the retrieved locations from M_2.
3. $\text{SSE}\textsc{Add}\langle (\sigma, d), \mathsf{EDB} \rangle$: For every w in d, client computes $i_1, \ldots, i_c \leftarrow h_1(w)$ and client and server run $\textsc{ObliviousAccess}\langle (\sigma_1, i_j, \mathsf{null}), \mathsf{EM}_1 \rangle$ for $j \in \{1, \ldots, c\}$ for client to obtain $\mathsf{M}_1[i_j]$. If $(w, count_w)$ is in the retrieved locations let i_j^* be the location it was found at. If not, let $count_w = 0$ and let i_j^* be the first empty location from the retrieved ones. Client and server run $\textsc{ObliviousAccess}\langle (\sigma_1, i_j^*, (w, count_w + 1)), \mathsf{EM}_1 \rangle$ to increase the counter by one.

   Client then computes $i'_1, \ldots, i'_c \leftarrow h_2(w||count_w + 1)$ and client and server run $\textsc{ObliviousAccess}\langle (\sigma_2, i'_j, \mathsf{null}), \mathsf{EM}_2 \rangle$ to retrieve $\mathsf{M}_2[i'_j]$ for $j \in \{1, \ldots, c\}$. Let i'_k be the first empty location among them. Client and server run

$$\textsc{ObliviousAccess}\langle (\sigma_2, i'_k, (w||count + 1)), \mathsf{EM}_2 \rangle.$$

The main disadvantage of our second approach is that for each search, it requires $count_w$ ORAM accesses to retrieve all matching documents. This means that the bandwidth/computation overhead of ORAM scheme is multiplied by $count_w$ which can be large for some keywords. More importantly, it would require $O(count_w)$ rounds since the ORAM accesses cannot be parallelized in our constant-round ORAM construction. In particular, note that each memory garbled circuit in the construction can only be used once and needs to be replaced before the next memory access. Finally, the constant-round ORAM needs to store a memory array that is proportional to the number of (w, d) tuples associated with the database, which is significantly larger than the number of unique keywords, increasing the storage overhead of the resulting SSE scheme.

Next, we address all these efficiency concerns, showing a construction that only requires a single ORAM access using our constant-round construction.

4.4 SSE from Path-ORAM

The idea is to not only store a per-keyword counter $count_w$ as before, but also to store a $access_w$ that represents the number of search/add queries performed on w so far. Similar to the previous approach, the tuple $(w, (count_w, access_w))$ is stored in a hash table that is implemented using our constant-round ORAM scheme TWORAM. The $count_w$ is incremented whenever a new document containing w is added and the $access_w$ is incremented after each search/add query for w.

The tuples $(w||i, d_i)$ for all d_i containing w are then stored in a one-level (non-recursive) Path-ORAM. In order to avoid storing a large client-side position map for this non-recursive Path-ORAM, we generate/update the positions pseudorandomly using a PRF $F_K(w||i||access_w)$. Since each document d_i has a different index and each search/add query for w will increment $access_w$, the pseudorandomness property of F ensures that this way of generating the position maps is indistinguishable from generating them at random. Now the client only needs to keep the secret key K. Note that since we are using a one-level Path-ORAM to store the documents, we can handle multiple parallel accesses without any problems, hence obtaining a constant-round search/add complexity. Furthermore, we only access TWORAM(which uses garbled circuits) once per keyword search to retrieve the tuple $(w, (count_w, access_w))$, so TWORAM's overhead is not multiplied by $count_w$ for each search/add query. Similarly, the storage overhead of $\mathsf{TWORAM}$is only for a memory array of size $|W|$ (number of unique keywords in documents) which is significantly smaller than the number of keyword-document pairs needed in the general approach.

We need to make a few small modifications to the syntax of the abstraction for Path-ORAM here. First, since we generate the position map on the fly using a PRF, it is convenient to modify the syntax of the UPDATE procedure to take the new random position as input, instead of internally generating it in our original syntax. Also, since we are not extracting an index y from the Path-ORAM and instead are extracting a tuple of the form $(w||i, d_i)$, we will pass $w||i$ as input in place of y in the EXTRACT and UPDATE operations.

We now describe the SSE scheme. Given a hash table $H = (\mathsf{hsetup}, \mathsf{hlookup}, \mathsf{hwrite})$, our constant-round ORAM scheme $\mathsf{TWORAM} = (\text{SETUP}, \text{OBLIVIOUSACCESS})$, a single level Path-ORAM scheme with procedures $(\text{INITIALIZE}, \text{EXTRACT}, \text{UPDATE})$, and a PRF function F, we build an SSE scheme $(\text{SSESETUP}, \text{SSESEARCH}, \text{SSEADD})$ as follows:

1. $\langle \sigma, \mathsf{EDB} \rangle \leftrightarrow \text{SSESETUP}\langle (1^\kappa, \mathsf{DB}), \perp \rangle$: Given an initial set of documents DB, let S be the set of $(w, (count_w, access_w = 0))$ where $count_w$ is the number of documents containing w, and $access_w$ denotes the number of times the keyword w has been searched/added.

   Client runs $\mathsf{hsetup}(S, size)$ to obtain (h, M). $size$ is the anticipated maximum size of the hash table H. Then client and server run $\langle \sigma_s, \mathsf{EM} \rangle \leftrightarrow \text{SETUP}\langle (1^\kappa, \mathsf{M}), \perp \rangle$.

   Let A_L be an initially empty memory array with a size that estimates an upper bound on total number of (w, d) pairs ind DB. Client runs $\mathcal{T} \leftarrow$

Initialize($1^\kappa, A_L$), and only sends the tree T_L for the last level to server, and discards the rest.

Client generates a PRF key $K \leftarrow \{0,1\}^\kappa$.

For every item $(w, (count_w, access_w))$ in S, and for $1 \leq i \leq count_w$ (in parallel):

(a) Client lets $val_{w,i} = (w||i, d_i)$ where d_i denotes the ith document in DB containing w.

(b) Client lets $x_{w,i} = F_K(w||i||access_w)$ and sends $x_{w,i}$ to server who returns the encrypted buckets on path $T_L(x_{w,i})$ which client decrypts itself.

(c) Client runs $\{T_L(x_{w,i})\} \leftarrow$ Update$(w||i, write, val_{w,i}, T_L(x_{w,i}), x'_{w,i})$, where $x'_{w,i} = F_K(w||i||access_w + 1)$, to insert $val_{w,i}$ into the path along its new path $T_L(x'_{w,i})$. Client then encrypts the updated path $T_L(x_{w,i})$ and sends it to server who updates T_L.

Note that server's output is $\mathsf{EDB} = (\mathsf{EM}, T_L)$ and client's output is $\sigma = (\sigma_s, h, K)$.

2. SSESearch$\langle(\sigma, w), \mathsf{EDB}\rangle$: Client computes $i_1, \ldots, i_c \leftarrow h(w)$ and client and server run ObliviousAccess$\langle(\sigma_s, i_j, \mathsf{null}), \mathsf{EM}\rangle)$ for $j \in \{1, \ldots, c\}$. If client finds $(w, (count_w, access_w))$ in one of the retrieved locations, let i_j^* be the location it was found at. If such a pair is not found the search ends here. Client and server run ObliviousAccess$\langle(\sigma_s, i_j^*, (w, count_w, access_w + 1)), \mathsf{EM}\rangle$ to increase the $access_w$ by one.

   For $1 \leq i \leq count_w$ (in parallel):

   (a) Client lets $x_{w,i} = F_K(w||i||access_w)$ and sends $x_{w,i}$ to server who returns $T_L(x_{w,i})$ which client decrypts.

   (b) Client runs $(w||i, d_i) \leftarrow$ Extract$(L, w||i, T_L(x_{w,i}))$, and outputs d_i. Client runs $\{T_L(x_{w,i})\} \leftarrow$ Update$(w||i, read, (w||i, d_i), T_L(x_{w,i}), x'_{w,i} = F_K(w||i||access_w + 1))$ to update the location of $(w||i, d_i)$ to $x'_{w,i}$. Client then encrypts the updated path and sends it to server to update T_L.

3. SSEAdd$\langle(\sigma, d), \mathsf{EDB}\rangle$:

   For every w in d:

   (a) Client computes $i_1, \ldots, i_c \leftarrow h(w)$ and client and server run

   $$\text{ObliviousAccess}\langle(\sigma_s, i_j, \mathsf{null}), \mathsf{EM}\rangle),$$

   for $j \in \{1, \ldots, c\}$. If client finds $(w, (count_w, access_w))$ in one of the retrieved locations, let i_j^* be the location it was found at. Else, it lets i_j^* be the first empty location among the retrieved ones.

   (b) Client and server run ObliviousAccess$\langle(\sigma_s, i_j^*, (w, (count_w + 1, access_w + 1))), \mathsf{EM}\rangle$ to increase $count_w$ and $access_w$ by one.

   (c) Client lets $x_{w,count_w} = F_K(w||count_w||acess_w)$ and sends $x_{w,count_w}$ to server who returns encrypted $T_L(x_{w,count_w})$ back. Client decrypts the path.

   (d) Client lets $x' = F_K(w||count_w + 1||access_w + 1)$ and runs $\{T_L(x_{w,count_w})\} \leftarrow$ Update$(w||i, write, (w||count_w + 1, d), T_L(x_{w,count_w}), x')$ to update the path. Client then encrypts the updated path and sends it to server to update T_L.

Before stating the security theorem for the above SSE scheme, we first need to make the leakage function associated with the scheme more precise. The leakage function $\mathcal{L}(\mathsf{DB}, H)$ for our scheme outputs the following (DB is the database and H is the search/add history): $|W|$, number unique keywords in all documents; $|\mathsf{DB}(w)|$ for every w searched; $\sum_{w \in W} |\mathsf{DB}(w)|$ i.e. the number of (w, d) pairs where w is in d. See Appendix A.4 for the proof.

Theorem 2. *The above SSE scheme is $\mathcal{L}$-secure (cf. Definition of* Sect. 4*), if* TWORAM *is secure (cf. Definition in* Sect. 2.2*), F is a PRF, and the encryption used in the one-level Path-ORAM is CPA-secure.*

Efficiency. The setup cost for our SSE scheme is the sum of the setup cost for TWORAM for a memory of size $|W|$, and the setup for a one-level Path-ORAM of size $n = \sum_{w \in W} |\mathsf{DB}(w)|$ which is $O(n \log n \operatorname{loglog} n)$.

The bandwidth cost for each search/add query w is the cost of one ORAM read in TWORAMplus $O(|\mathsf{DB}(w)| * (\log n \operatorname{loglog} n))$ for $n = \sum_{w \in W} |\mathsf{DB}(w)|$.

Acknowledgments. This work was done in part while the authors were visiting the Simons Institute for the Theory of Computing, supported by the Simons Foundation and by the DIMACS/Simons Collaboration in Cryptography through NSF grant #CNS-1523467. Sanjam Garg was supported in part from a DARPA/ARL SAFEWARE award, AFOSR Award FA9550-15-1-0274, and NSF CRII Award 1464397. Charalampos Papamanthou was supported in part by NSF grants #1514261 and #1526950, by a NIST award, by a Google Faculty Research Award and by Yahoo! Labs through the Faculty Research Engagement Program (FREP). The views expressed are those of the authors and do not reflect the official policy or position of the Department of Defense, the National Science Foundation, or the U.S. Government.

A More Details on Path ORAM

A.1 Path ORAM Abstraction Algorithms

Algorithm 1. Setting up path ORAM data structures.

1: **procedure** $\mathcal{T} \leftarrow$ INITIALIZE$(1^\kappa, A_L)$
2: Let π_L be a random permutation from $[n]$ to $[n]$;
3: Store $(x, A_L[x], \pi_L(x))$ at leaf $\pi_L(x)$ of tree T_L;
4: **for** $i = L$ down to 3 **do**
5: Set $A_{i-1}[x] = \pi_i(2x-1)||\pi_i(2x)$ for $x = 1, \ldots, 2^{i-1}$;
6: Let π_{i-1} be a random permutation from $[2^{i-1}]$ to $[2^{i-1}]$;
7: Store $(x, A_{i-1}[x], \pi_{i-1}(x))$ at leaf $\pi_{i-1}(x)$ of tree T_{i-1};
8: **end for**
9: Let A_1 be an array of 2 entries such that $A_1[x] = \pi_2(2x-1)||\pi_2(2x)$ for $x = 1, 2$;
10: **return** $\{A_1, T_2, \ldots, T_L\}$;
11: **end procedure**

Algorithm 2. Extraction algorithm for buckets.

1: **procedure** $\pi \leftarrow$ExtractBucket(i, y, b)
2: Search bucket b to retrieve block $(y_i, A_i[y_i], p)$;
3: **if** found **then**
4: **return** $\pi \leftarrow \mathsf{select}(A_i[y_i], b_i)$; ▷ π is the index of the path to be explored in T_{i+1}.
5: **else**
6: **return** $\perp$;
7: **end if**
8: **end procedure**

Algorithm 3. Update algorithm. It takes as input $L-1$ paths and local storage A_1 and outputs new paths, based on the new assignments of positions.

1: **procedure** $\{A_1, T_2(x_2), \ldots, T_L(x_L)\} \leftarrow$Update$(y, val, A_1, T_2(x_2), \ldots, T_L(x_L))$
2: $\mathsf{select}(A_1[y_1], b_1) \leftarrow r_2$; ▷ r_i is random in $[1, 2^{i+1}]$.
3: **for** $i = 2$ to L **do**
4: $T_i.\mathsf{root} \leftarrow T_i.\mathsf{root} \cup \mathsf{readPath}(T_i(x_i))$; ▷ $T_i.\mathsf{root}$ serves as the stash C_i.
5: Update block $(y_i, A_i[y_i], x_i)$ to $(y_i, A_i[y_i], r_i)$ in $T_i.\mathsf{root}$;
6: $\mathsf{select}(A_i[y_i], b_i) \leftarrow r_{i+1}$; ▷ if $i = L$ **do** if $val \neq \mathsf{null}$, $A_L[y] \leftarrow val$, else NOOP.
7: $[T_i.\mathsf{root}, T_i(x_i)] \leftarrow \mathsf{evictPath}(T_i.\mathsf{root})$;
8: **end for**
9: **return** $A_1, T_2(x_2), T_3(x_3), \ldots, T_L(x_L)$;
10: **end procedure**

Protocol $\langle \sigma, \mathsf{EM} \rangle \leftrightarrow$ Setup$\langle (1^\kappa, \mathsf{M}), \perp \rangle$:

Client:

1: Pick a κ-bit s; Run $\{A_1, T_2, \ldots, T_L\} \leftarrow$ Initialize$(1^\kappa, \mathsf{M})$;
2: For all $i > 1$, for all $u \in T_i$, set $\mathsf{B}_u \leftarrow \mathsf{Enc}_s(\mathsf{bucket}_u)$, where $\mathsf{Enc}_s(.)$ is a CPA-secure encryption;
3: For all $i > 1$, for all $u \in T_i$, send to server data B_u;
4: **return** s and A_1 as σ;

Server:

1: **return** all data B_u sent by the client from above as EM;

Fig. 6. Formal description of the Setup protocol for the interactive ORAM [34].

A.2 Path ORAM Protocols with $\log n$ Rounds of Interaction Using the Abstraction

A.3 Proof of Security for TWORAM

Now we prove TWORAM is a secure realization of an oblivious RAM scheme as described in Sect. 2.2. We start by arguing correctness. Note that the garbled circuits implement the exact same procedures as are required in our abstraction. Therefore the correctness of our scheme follows directly from the correctness of the underlying Path ORAM scheme and garbled circuits construction. Next we argue security. In other words we need to argue that for any adversary A, there

exists a simulator Sim for which the following two distributions are computationally indistinguishable.

- $\mathrm{Real}_A^{\Pi}(\kappa)$: A chooses M. The experiment then runs $\langle \sigma, \mathsf{EM} \rangle \leftrightarrow$ $\textsc{Setup}\langle (1^\kappa, \mathsf{M}), \perp \rangle$. A then provides read/write queries (y_i, v) where $v = \mathsf{null}$ on reads, for which the experiment runs the protocol

$$\langle (\mathsf{M}[y_i], \sigma_i), \mathsf{EM}_i \rangle \leftrightarrow \textsc{ObliviousAccess}\langle (\sigma_{i-1}, y_i, v), \mathsf{EM}_{i-1} \rangle.$$

  Denote the full transcript of the protocol by t_i. Eventually, the experiment outputs $(\mathsf{EM}, t_1, \ldots, t_q)$ where q is the total number of read/write queries.
- $\mathrm{Ideal}_{\mathsf{Sim}}^{\Pi}(\kappa)$: The experiment outputs $(\mathsf{EM}, t_1', \ldots, t_q') \leftarrow \mathsf{Sim}(q, |\mathsf{M}|, 1^\kappa)$.

Our Simulator. Note that the simulator needs to provide to the server, for all u $\tilde{\mathsf{C}}^u, \mathsf{X}^u$ and for all $i \in \{2, \ldots, L\}$ $\beta_i := \mathsf{nState}^{(i,0,0)}_{\mathsf{cState}^{(i+1,0,0)}}$. Furthermore replacement circuits need to be provided as read/write queries are implemented. Our simulator Sim generates these as follows:

- For each $u = (i, j, k)$, let $(\tilde{\mathsf{C}}^u, \mathsf{lab}^u \leftarrow \mathsf{GCircuit}(1^\kappa, \mathsf{P}[u, b_u, \mathsf{lab}_{\mathbf{0}}^{(i,j+1,2k+b)}])$, where b_u is random bit and P is a circuit that, if $j = i$ outputs $(\mathsf{nextRoot}, \mathsf{lab}_{\mathbf{0}}^{(i+1,0,0)})$, else if $b = 0$ then it outputs $(\mathsf{left}, \mathsf{lab}_{\mathbf{0}}^{(i,j+1,2k+b)})$ and $(\mathsf{right}, \mathsf{lab}_{\mathbf{0}}^{(i,j+1,2k+b)})$ otherwise.
- Each X^u is generated as as encryption of a zero-string, namely $\mathsf{Enc}_s(\mathbf{0})$. Similarly for all $i \in \{2, \ldots, L\}$ $\beta_i := \mathsf{nState}_{\mathbf{0}}^{(i,0,0)}$.

Note that as the provided garbled circuits are executed, replacements circuits need to be given and they are generated in the same manner as above.

Proof of Indistinguishability. The proof follows by a hybrid argument.

- H_0: This hybrid corresponds to the honest execution $\mathrm{Real}_A^{\Pi}(\kappa)$ as done honestly.
- H_1: This hybrid is same as H_0 except that we now generate all the X^u values as encryptions of zero-strings of appropriate length. Specifically, for each u we set $\mathsf{X}^u \leftarrow \mathsf{Enc}_s(\mathbf{0})$.
  The indistinguishability between H_0 and H_1 follows from the security of the encryption scheme $(\mathsf{Enc}, \mathsf{Dec})$.
- H_2: In this hybrid the simulator maintains the entire Path ORAM tree internally but does not include it in the provided garbled circuits. In other words garbled circuits are generated as follows:
  - For each $u = (i, j, k)$, let $(\tilde{\mathsf{C}}^u, \mathsf{lab}^u \leftarrow \mathsf{GCircuit}(1^\kappa, \mathsf{P}[u, b_u, \mathsf{lab}_{\mathbf{0}}^{(i,j+1,2k+b)}])$, where b_u is 0 or 1 depending on whether the execution as per ORAM would go left or right and P is a circuit that, if $j = i$ outputs $(\mathsf{nextRoot}, \mathsf{lab}_{\mathbf{0}}^{(i+1,0,0)})$, else if $b = 0$ then it outputs $(\mathsf{left}, \mathsf{lab}_{\mathbf{0}}^{(i,j+1,2k+b)})$ and $(\mathsf{right}, \mathsf{lab}_{\mathbf{0}}^{(i,j+1,2k+b)})$ otherwise.
  - Each X^u is generated as as encryption of a zero-string, namely $\mathsf{Enc}_s(\mathbf{0})$. Similarly for all $i \in \{2, \ldots, L\}$ $\beta_i := \mathsf{nState}_{\mathbf{0}}^{(i,0,0)}$.

Protocol $\langle(\mathsf{M}[y], \sigma'), \mathsf{EM}'\rangle \leftrightarrow$ OBLIVIOUSACCESS$\langle(\sigma, y, val), \mathsf{EM}\rangle$:

Client(1):

1: Compute $x_2 \leftarrow \mathsf{select}(A_1[y_1], b_1)$. Send to server index x_2; ▷ run **Server**(2)

Server(i):

1: For all $u \in T_i(x_i)$ send to client B_u; ▷ run **Client(i)**

Client(i):

1: $\pi = \perp$;
2: **for** $u \in T_i(x_i)$ **do**
3: $\mathsf{bucket}_u \leftarrow \mathsf{Dec}_s(\mathsf{B}_u)$;
4: **if** $\pi = \perp$ **then**
5: $\pi \leftarrow$ EXTRACTBUCKET$(i, y, \mathsf{bucket}_u)$;
6: $x_{i+1} = \pi$;
7: **end if**
8: **end for**
9: **if** $i < L$ **then**
10: Send to server new index x_{i+1}; ▷ run **Server($i + 1$)**
11: **else**
12: $\{A_1, T_2(x_2), \ldots, T_L(x_L)\} \leftarrow$ UPDATE$(y, val, A_1, T_2(x_2), \ldots, T_L(x_L))$;
13: For all $i > 1$, for all $u \in T_i(x_i)$, send to server $\mathsf{B}_u \leftarrow \mathsf{Enc}_s(\mathsf{bucket}_u)$; ▷ run **Server($L$)**
14: **return** x_{L+1} as $\mathsf{M}[y]$ and A_1 as σ';
15: **end if**

Server(L):

1: **return** the data received by the client as EM';

Fig. 7. Formal description of the OBLVIOUSACCESS protocol for the interactive ORAM [34].

The indistinguishability between H_1 and H_2 follows by a sequence of hybrids where each garbled circuit is replaced by a simulated garbled circuit. Here these hybrids must be performed in sequence in which garbled circuits are consumed. Note that for the unconsumed garbled circuits the input labels aren't provided (or hardcoded inside any other circuit) and hence they can also be simulated.

– H_3: Same as H_2, except that each b_u is now chosen uniformly random, independent of the Path ORAM execution. Note that this is same as the simulator. The indistinguishability between H_2 and H_3 from the security of the Path ORAM scheme.

This concludes the proof. □

A.4 Proof of Security for the SSE scheme

We prove Theorem 2 on security of the SSE scheme next, Following the definition of Sect. 4, we first describe a simulator Sim who generates the transcripts for the ideal distribution $\mathrm{Ideal}^{\Pi}_{A,\mathsf{Sim},\mathcal{L}}(\kappa)$. Sim takes as input $\mathcal{L}(\mathsf{DB}, H)$, and does the following: To generate full transcripts of the constant round ORAM scheme for the adversary A, Sim runs Sim', the simulator that exists for that scheme

due its security (see definition of Sect. 2.2). That is, he runs $(\mathsf{EM}, t_1, \ldots, t_q) \leftarrow \mathsf{Sim}'(q, |\mathsf{M}|, 1^\kappa)$, where he drives $|\mathsf{M}|$ from $|W|$. To simulate the transcripts of the path-ORAM component, it generates a one-level path ORAM tree T_L for a memory array of size $\sum_{w \in W} |\mathsf{DB}(w)|$ filled with all 0 values. For each read/add query, it replaces the PRF-genenerated paths by uniformly random paths, and generates freshly generated ciphertexts of 0 for updated paths. Sim knows the number of paths to retrieve/update for each query from the leakage function which outputs $|\mathsf{DB}(w)|$ for every query w. This completes the description of the simulator. We now need to show that $\text{Ideal}^{\Pi}_{A,\mathsf{Sim},\mathcal{L}}(\kappa)$ is indinstinguishable from $\text{Real}^{\Pi}_{A}(\kappa)$, which constitutes the first in the sequence of our Hybrids:

Proof of Indistinguishability. The proof follows by a hybrid argument.

- H_0: This hybrid corresponds to the honest execution $\text{Real}^{\Pi}_{A}(\kappa)$ for the SSE scheme which we repeat here for completeness. A chooses DB. The experiment then runs $\langle \mathsf{EDB}, \sigma \rangle \leftrightarrow \text{SSESetup}\langle (1^\kappa, \mathsf{DB}), \perp \rangle$. A then adaptively makes search queries w_i, which the experiment answers by running the protocol $\langle \mathsf{DB}_{i-1}(w_i), \sigma_i \rangle \leftrightarrow \text{SSESearch}\langle (\sigma_{i-1}, w_i), \mathsf{EDB}_{i-1} \rangle$. Denote the full transcripts of the protocol by t_i. Add queries are handled in a similar way. Eventually, the experiment outputs $(\mathsf{EDB}, t_1, \ldots, t_q)$ where q is the total number of search/add queries made by A.
- H_1:Similar to H_0, except that the portions of t_i's corresponding to the constant-round ORAM are instead generated by $\mathsf{Sim}'(q, |\mathsf{M}|, 1^\kappa)$ where Sim' is the simulator in the proof of the ORAM scheme.
  The indistinguishability of H_0 and H_1 follows from security of the ORAM scheme.
- H_2: Similar to H_1 except that all ciphertexts in the path ORAM tree are replaced by encryptions of 0, and all updated ciphertexts will be fresh encryption of 0.
  The indistinguishability of H_2 and H_1 follows from the semantic security of the encryption scheme used in the path ORAM.
- H_3: Similar to H_2 except that all PRF-generated positions are replaced by uniformly random positions. Note that H_3 is essentially $\text{Ideal}^{\Pi}_{A,\mathsf{Sim},\mathcal{L}}(\kappa)$.
  The indistinguishability of H_3 and H_2 follows from the pseudorandomness of the the PRF.

This concludes the proof. □

References

1. Afshar, A., Hu, Z., Mohassel, P., Rosulek, M.: How to efficiently evaluate RAM programs with malicious security. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015. LNCS, vol. 9056, pp. 702–729. Springer, Heidelberg (2015)
2. Bellare, M., Hoang, V.T., Rogaway, P.: Foundations of garbled circuits. In: CCS, pp. 784–796 (2012)
3. Cash, D., Grubbs, P., Perry, J., Ristenpart, T.: Leakage-abuse attacks against searchable encryption. In: CCS, pp. 668–679 (2015)

4. Cash, D., Jarecki, S., Jutla, C., Krawczyk, H., Roşu, M.-C., Steiner, M.: Highly-scalable searchable symmetric encryption with support for boolean queries. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part I. LNCS, vol. 8042, pp. 353–373. Springer, Heidelberg (2013)
5. Chang, Y.-C., Mitzenmacher, M.: Privacy preserving keyword searches on remote encrypted data. In: Ioannidis, J., Keromytis, A.D., Yung, M. (eds.) ACNS 2005. LNCS, vol. 3531, pp. 442–455. Springer, Heidelberg (2005)
6. Chase, M., Kamara, S.: Structured encryption and controlled disclosure. In: Abe, M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 577–594. Springer, Heidelberg (2010)
7. Cormen, T.H., Stein, C., Rivest, R.L., Leiserson, C.E.: Introduction to Algorithms, 2nd edn. McGraw-Hill Higher Education, New York (2001)
8. Curtmola, R., Garay, J., Kamara, S., Ostrovsky, R.: Searchable symmetric encryption: improved definitions and efficient constructions. In: CCS, pp. 79–88 (2006)
9. Devadas, S., van Dijk, M., Fletcher, C.W., Ren, L., Shi, E., Wichs, D.: Onion ORAM: a constant bandwidth blowup oblivious RAM. In: TCC, pp. 145–174 (2016)
10. Fletcher, C., Naveed, M., Ren, L., Shi, E., Stefanov, E.: Bucket ORAM: single online roundtrip, constant bandwidth oblivious RAM. Cryptology ePrint Archive, Report 2015/1065 (2015). http://eprint.iacr.org/
11. Garg, S., Lu, S., Ostrovsky, R.: Black-box garbled RAM. In: FOCS, pp. 210–229 (2015)
12. Garg, S., Lu, S., Ostrovsky, R., Scafuro, A.: Garbled RAM from one-way functions. In: STOC, pp. 449–458 (2015)
13. Gentry, C., Goldman, K.A., Halevi, S., Julta, C., Raykova, M., Wichs, D.: Optimizing ORAM and using it efficiently for secure computation. In: De Cristofaro, E., Wright, M. (eds.) PETS 2013. LNCS, vol. 7981, pp. 1–18. Springer, Heidelberg (2013)
14. Gentry, C., Halevi, S., Lu, S., Ostrovsky, R., Raykova, M., Wichs, D.: Garbled RAM revisited. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 405–422. Springer, Heidelberg (2014)
15. Goh, E.-J.: Secure indexes. Cryptology ePrint Archive, Report 2003/216 (2003). http://eprint.iacr.org/2003/216/
16. Goldreich, O., Ostrovsky, R.: Software protection and simulation on oblivious RAMs. J. ACM **43**(3), 431–473 (1996)
17. Goodrich, M.T., Mitzenmacher, M.: Privacy-preserving access of outsourced data via oblivious RAM simulation. In: Aceto, L., Henzinger, M., Sgall, J. (eds.) ICALP 2011, Part II. LNCS, vol. 6756, pp. 576–587. Springer, Heidelberg (2011)
18. Goodrich, M.T., Mitzenmacher, M., Ohrimenko, O., Tamassia, R.: Privacy-preserving group data access via stateless oblivious RAM simulation. In: SODA, pp. 157–167 (2012)
19. Gordon, S.D., Katz, J., Kolesnikov, V., Krell, F., Malkin, T., Raykova, M., Vahlis, Y.: Secure two-party computation in sublinear (amortized) time. In: CCS, pp. 513–524 (2012)
20. Islam, M.S., Kuzu, M., Kantarcioglu, M.: Access pattern disclosure on searchable encryption ramification, attack and mitigation. In: NDSS (2012)
21. Dautrich Jr., J.L., Stefanov, E., Shi, E.: Burst ORAM: minimizing ORAM response times for bursty access patterns. In: Usenix Security, pp. 749–764 (2014)
22. Kamara, S., Papamanthou, C.: Parallel and dynamic searchable symmetric encryption. In: FC, pp. 258–274 (2013)

23. Kamara, S., Papamanthou, C., Roeder, T.: Dynamic searchable symmetric encryption. In: CCS, pp. 965–976 (2012)
24. Kushilevitz, E., Lu, S., Ostrovsky, R.: On the (in)security of hash-based oblivious RAM and a new balancing scheme. In: SODA, pp. 143–156 (2012)
25. Lindell, Y., Pinkas, B.: A proof of security of Yao's protocol for two-party computation. J. Cryptol. **22**(2), 161–188 (2009)
26. Liu, C., Zhu, L., Wang, M., Tan, Y.-A.: Search pattern leakage in searchable encryption: attacks and new construction. Inf. Sci. **265**, 176–188 (2014)
27. Lu, S., Ostrovsky, R.: Distributed oblivious RAM for secure two-party computation. In: Sahai, A. (ed.) TCC 2013. LNCS, vol. 7785, pp. 377–396. Springer, Heidelberg (2013)
28. Lu, S., Ostrovsky, R.: How to garble RAM programs? In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 719–734. Springer, Heidelberg (2013)
29. Moataz, T., Mayberry, T., Blass, E.-O.: Constant communication ORAM with small blocksize. In: CCS, pp. 862–873 (2015)
30. Shen, E., Shi, E., Waters, B.: Predicate privacy in encryption systems. In: Reingold, O. (ed.) TCC 2009. LNCS, vol. 5444, pp. 457–473. Springer, Heidelberg (2009)
31. Shi, E., Chan, T.-H.H., Stefanov, E., Li, M.: Oblivious RAM with $O((\log N)^3)$ worst-case cost. In: Lee, D.H., Wang, X. (eds.) ASIACRYPT 2011. LNCS, vol. 7073, pp. 197–214. Springer, Heidelberg (2011)
32. Song, D.X., Wagner, D., Perrig, A.: Practical techniques for searches on encrypted data. In: IEEE Symposium on Security and Privacy, pp. 44–55 (2000)
33. Stefanov, E., Papamanthou, C., Shi, E.: Practical dynamic searchable encryption with small leakage. In: NDSS (2014)
34. Stefanov, E., van Dijk, M., Shi, E., Fletcher, C.W., Ren, L., Xiangyao, Y., Devadas, S.: Path ORAM: an extremely simple oblivious RAM protocol. In: CCS, pp. 299–310 (2013)
35. van Liesdonk, P., Sedghi, S., Doumen, J., Hartel, P., Jonker, W.: Computationally efficient searchable symmetric encryption. In: Jonker, W., Petković, M. (eds.) SDM 2010. LNCS, vol. 6358, pp. 87–100. Springer, Heidelberg (2010)
36. Wang, X.S., Hubert Chan, T.-H., Shi, E., Circuit, O.: On tightness of the Goldreich-Ostrovsky lower bound. In: CCS, pp. 191–202 (2015)
37. Williams, P., Sion, R.: Single round access privacy on outsourced storage. In: CCS, pp. 293–304 (2012)
38. Yao, A.C.: Protocols for secure computations (extended abstract). In: FOCS (1982)
39. Zhang, Y., Katz, J., Papamanthou, C.: All your queries are belong to us: the power of file-injection attacks on searchable encryption. In: Usenix Security (2016)