Write Once Read Many Oblivious RAM

Bogdan Carbunar, Radu Sion

Abstract—We introduce WORM-ORAM, a first mechanism that combines Oblivious RAM (ORAM) access privacy and data confidentiality with Write Once Read Many (WORM) regulatory data retention guarantees. Clients can outsource their database to a server with full confidentiality and data access privacy, and, for data retention, the server ensures client access WORM semantics. In general *simple confidentiality* and *WORM* assurances are easily achievable e.g., via an encrypted outsourced data repository with server-enforced read-only access to existing records (albeit encrypted). However, this becomes hard when also *access privacy* is to be ensured – when client access patterns are necessarily hidden and the server cannot enforce access control directly. WORM-ORAM overcomes this by deploying a set of zero-knowledge proofs to convince the server that all stages of the protocol are WORM-compliant.

❖

1 INTRODUCTION

Regulatory frameworks impose a wide range of policies in finance, life sciences, health-care and the government. Examples include the Gramm-Leach-Bliley Act [1], the Health Insurance Portability and Accountability Act [2] (HIPAA), the Federal Information Security Management Act [3], the Sarbanes-Oxley Act [4], the Securities and Exchange Commission rule 17a-4 [5], the DOD Records Management Program under directive 5015.2 [6], the Food and Drug Administration 21 CFR Part 11 [7], and the Family Educational Rights and Privacy Act [8]. Over 10,000 regulations are believed to govern the management of information in the US alone [9].

A recurrent theme to be found throughout a large part of this regulatory body is the need for assured lifecycle storage of records. A main goal there is to support WORM semantics: once written, data cannot be undetectably altered or deleted before the end of its regulation-mandated life span. This naturally stems from the perception that the primary adversaries are powerful insiders with superuser powers coupled with full access to the storage system. Indeed much recent corporate malfeasance has been at the behest of CEOs and CFOs, who also have the power to order the destruction or alteration of incriminating records [10].

Major storage vendors have responded by offering compliance storage and WORM products, for on-site deployment, including IBM [11], HP [12], EMC [13]. Hitachi Data Systems [14], Zantaz [15], StorageTek [16], Sun Microsystem [17] [18], Network Appliance [19]. and Quantum Inc. [20].

However, as data management is increasingly outsourced to third party "clouds" providers such as Google, Amazon and Microsoft, existing systems simply do not work. When outsourced data lies under the incidence of both mandatory *data retention* regulation and privacy/confidentiality concerns – as it often does

in outsourced contexts – new enforcement mechanisms are to be designed.

This task is non-trivial and immediately faces an apparent contradiction. On the one hand, data retention regulation stipulates that, once generated, data records cannot be erased until their "mandated expiration time", *even by their rightful creator* – history cannot be rewritten. On the other hand, access privacy and confidentiality in outsourced scenarios mandate non-disclosure of data and patterns of access thereto to the providers' servers, and can be achieved through "Oblivious RAM" (ORAM) based client-server mechanisms [21], [22]. Yet, by their very nature, existing ORAM mechanisms allow clients unfettered read/write access to the data, including the full ability to alter or remove previously written data records – thus directly contradicting data retention requirements.

Basic *confidentiality* and *WORM* assurances are achievable e.g., via traditional systems that could encrypt outsourced data and deploy server-enforced read-only access to data records once written. Yet, when also *access privacy* is to be ensured, client access patterns become necessarily hidden and the server cannot enforce WORM semantics directly.

In this paper we introduce WORM-ORAM, a first mechanism that combines the access privacy and data confidentiality assurances of traditional ORAM with Write Once Read Many (WORM) regulatory data retention guarantees. Clients can outsource their database to a server with full confidentiality and data access privacy, and, for data retention, the server ensures client access WORM semantics, i.e., specifically that client access is append-only: – once a data record has been written it cannot be removed or altered even by its writer.

WORM-ORAM is built around a set of novel efficient zero knowledge (ZK) proofs. The main insight is to allow the client unfettered ORAM access with full privacy to the server-hosted encrypted data set while simultaneously proving to the server *in zero-knowledge* – at all stages of the ORAM access protocol – that no existing records are overwritten and WORM semantics are preserved.

- *Bogdan Carbunar is with the Applied Research Center in Motorola, Schaumburg, IL 60195.*
- *Radu Sion is with the Computer Science Department in Stony Brook University, Stony Brook, NY.*

Specifically, clients can add encrypted data records to the database ("the ORAM") hosted by a service provider. Each record will be associated with a regulatory mandated expiration time. Once stored, the client can read all data obliviously, (and add new records) – only leaking that access took place and nothing else. No access patterns or data records or any other information is leaked. The server, without having plaintext access to the data or the client access patterns then ensures – in a client-server interaction – that any client access is WORM compliant: it is either a read of an existing record, or an addition of a new record (with a new index – no overwriting permitted).

To achieve the above, at an overview level, the solution outlines as follows. The server hosts two ORAMs, one storing the actual data items (the W-ORAM) and one allowing the private retrieval of items expiring at any given time (the E-ORAM). The E-ORAM is effectively a helper data structure allowing the client to determine which items to expire at given time intervals. Client access to the E-ORAM needs to be private, but does not need to be proved correct.

The server exports an access API to the W-ORAM to the client consisting of four types of operations: write, read, expire and compliance verification. For access pattern privacy purposes as in the traditional ORAM protocols, the data set stored at the server contains both "real" and "fake" items – this is discussed later. Moreover, for the same reasons, any data access, is to look identical to the server and be indistinguishable from any future access to the same record – this in fact means all ORAM accesses will need to consist of both a read and write component – keeping the ORAM database in a constant transformation process.

Then, during any legitimate access to the W-ORAM the client will prove in ZK to the server that the item written is real, "well formed" and can be decrypted later in the case of an audit. Moreover she also proves that the access does not in fact overwrite any existing database item. Just like for the classic ORAM, the server cannot distinguish between our read and write accesses.

In the expiration operation, the client proves in ZK that the element to be removed from the W-ORAM has indeed expired. Finally, at audit time, the data has to be accessible to an authorized auditor, even in the case of a non-cooperating client (e.g., that could refuse to reveal encryption keys).

We show that our solution does not change the computational complexity of existing ORAM implementations. However, we warn that the constants involved are non-negligible and render this result of theoretical interest only for now. Future work focuses on reducing these overheads towards true practical efficiency.

2 RELATED WORK

2.1 Oblivious RAM

This paper extends the work of [23] with novel constructions that provide indistinguishability for the read and write accesses, detailed descriptions of essential components such as element expiration and with proofs for the assurances provided by the solution including the zero knowledge properties.

Oblivious RAM [21] provides access pattern privacy to clients (or software processes) accessing a remote database (or RAM), requiring only logarithmic storage at the client. The amortized communication and computational complexities are $O(log^3 n)$. Due to a large hidden constant factor, the ORAM authors offer an alternate solution with computational complexity of $O(log^4 n)$, that is more efficient for all currently plausible database sizes.

In ORAM, the database is considered a set of n encrypted blocks and supported operations are read(id), and write(id, $newvalue$). The data is organized into $log_4(n)$ levels, as a pyramid. Level i consists of up to 4^i blocks; each block is assigned to one of the 4^i buckets at this level as determined by a hash function. Due to hash collisions each bucket may contain from 0 to $log n$ blocks.

ORAM Reads. To obtain the value of block id, the client must perform a read query in a manner that maintains two invariants: (i) it never reveals which level the desired block is at, and (ii) it never looks twice in the same spot for the same block. To maintain (i), the client always scans a single bucket in every level, starting at the top (Level 0, 1 bucket) and working down. The hash function informs the client of the candidate bucket at each level, which the client then scans. *Once the client has found the desired block, the client still proceeds to each lower level, scanning random buckets instead of those indicated by their hash function.* For (ii), once all levels have been queried, the client re-encrypts the query result with a different nonce and places it in the *top* level. This ensures that when it repeats a search for this block, it will locate the block immediately (in a different location), and the rest of the search pattern will be randomized. The top level quickly fills up; how to dump the top level into the one below is described later.

ORAM Writes. Writes are performed identically to reads in terms of the data traversal pattern, with the exception that the new value is inserted into the top level at the end. Inserts are performed identically to writes, since no old value will be discovered in the query phase. Note that semantic security properties of the re-encryption function ensure the server is unable to distinguish between reads, writes, and inserts, since the access patterns are indistinguishable.

Level Overflow. Once a level is full, it is emptied into the level below. This second level is then re-encrypted and re-ordered, according to a new hash function. Thus, accesses to this new generation of the second level will hence-forth be completely independent of any previous accesses. Each level overflows once the level above it has been emptied 4 times. Any re-ordering must be performed obliviously: once complete, the adversary must be unable to make any correlation between the old block locations and the new locations. A sorting network

is used to re-order the blocks.

To enforce invariant (i), note also that all buckets must contain the same number of blocks. For example, if the bucket scanned at a particular level has no blocks in it, then the adversary would be able to determine that the desired block was *not* at that level. Therefore, each re-order process fills all partially empty buckets to the top with *fake* blocks. Recall that since every block is encrypted with a semantically secure encryption function, the adversary cannot distinguish between fake and real blocks.

Pinkas and Reinman introduce in [24] a mechanism to provide $O(\log^2 n)$ oblivious access with only logarithmic client storage. The collision-free Cuckoo hash from [25] is employed to remove a $\log n$ factor of server storage (and eliminate the corresponding query overhead). Data is stored ORAM-like, in a pyramid-shaped set of levels, with queries proceeding downward interactively.

Unfortunately, the premise of this idea is flawed. The construction of the Cuckoo hash function at a given level considers only the data to be stored at a given level, not the data already stored at lower levels. This property is necessary to achieve the desired time complexity. Thus, queries for the lower-level data have a significant chance of sharing locations with data at the current level. Because the Cuckoo hash otherwise avoids collision, any time such an occurrence is observed indicates leaks access information. That is, when the adversary sees two queries access the same pair of hash table locations, it learns that at least one of those queries was, in fact, for lower-level (less recently accessed) data. This immediately violates access privacy. The authors acknowledged this problem.

2.2 Private Information Retrieval

Another set of existing mechanisms handle access pattern privacy (but *not data confidentiality*) in the presence of *multiple clients*. Private Information Retrieval (PIR) [26] protocols aim to allow (arbitrary, multiple) clients to retrieve information from public or private databases, without revealing to the database servers which records are retrieved.

In initial results, Chor et al. [26] proved that in an information theoretic setting, any single-server solution requires $\Omega(n)$ bits of communication. PIR schemes with only sub-linear communication overheads, such as [26], require multiple non-communicating servers to hold replicated copies of the data. When the information theoretic guarantee is relaxed single-server solutions with better complexities exist; an excellent survey of PIR can be found online [27], [28].

Recently, we have shown [29] that due to computation costs, use of existing non-trivial single-server PIR protocols on current hardware is still orders of magnitude more time-consuming than trivially transferring the entire database.

2.3 Oblivious Transfer with Access Control

Camenish et al. [30] study the problem of performing k sequential oblivious transfers (OT) between a client and a server storing N values. The work makes the case that previous solutions tolerate selective failures. A selective failure occurs when the server may force the following behavior in the ith round (for any i=1..k): the round should fail if the client requests item j (of the N items) and succeed otherwise. The paper introduces security definitions to include the selective failure problem and then propose two protocols to solve the problem under the new definitions.

Coull et al. [31] propose an access control oblivious transfer problem. Specifically, the server wants to enforce access control policies on oblivious transfers performed on the data stored: The client should only access fields for which it has the credentials. However, the server should not learn which credentials the client has used and which items it accesses.

Note that the above oblivious transfer flavors do not consider by definition the problem of obliviously enforcing WORM semantics as well as writing to the data. Our regulatory compliant problem is complicated by the fact that we also allow clients to add to the database while proving that operations performed on the data do not overwrite old records. One can trivially extend OT with an add call, by imposing $O(N)$ communication and computation overheads. However, by building our solution on ORAM we can perform both read and add operations with only poly-logarithmic complexity and traffic overheads.

3 MODEL AND PRELIMINARIES

3.1 Deployment and Threat Model

In the deployment model for networked compliance storage, a legitimate client creates and stores records with a (potentially untrusted) remote WORM storage service. These records are to be available later to both the client for read as well as to auditors for audits. Network layer confidentiality is assured by mechanisms such as SSL/IPSec. Without sacrificing generality, we will assume that the data is composed of equal-sized blocks (e.g., disk blocks, or database rows).

At a later time, a previously stored record's existence is regretted and the client will do everything in her power – e.g., attempt to convince the server to remove the record – to prevent auditors from discovering the record. The main purpose of a traditional WORM storage service is to defend against such an adversary.

Moreover, numerous data regulations feature requirements of "secure deletion" of records at the end of their mandated retention periods. Then, in the WORM adversarial model the focus is mainly on preventing clients from "rewriting" history, rather than "remembering" it. Additionally, we prevent the rushed removal

of records before their retention periods. Thus, the traditional Write-Once Read-Many (WORM) systems have the following properties:

- Data records may be written by clients to the server once, read many times and not altered for the duration of their life-cycle.
- Records have associated mandatory expiration times. After expiration, they should not be accessible for either audit or read purposes.
- In the case of audits, stored data should be accessible to auditors even in the presence of a non-cooperating client refusing to reveal encryption keys. Compliant record expiration of inaccessible records should be easily proved to auditors.

Additionally, when records and their associated access patterns are sensitive they need to be concealed from a curious server. The main purpose of WORM-ORAM is to enable WORM semantics while preserving data confidentiality and access pattern privacy. This inability of the server to "see" data and associated access patterns prevents the deployment of conventional file/storage system access control mechanisms or data outsourcing techniques. Thus, we have the following additional requirements:

- Data records are encrypted from the server (*confidentiality*).
- The server cannot distinguish between different read operations targeting the same or different data records (*access privacy*).
- During a read, in the ORAM protocol, to enforce *WORM semantics*, clients will need to prove to the server that any access did not remove data records. Specifically, when re-inserting one of the read elements back into the root of the ORAM pyramid, the client needs to prove to the server in ZK that the inserted element is a correct re-encryption of the previously removed "real" element (see Section 2.1 for details).
- During a re-shuffle, in the ORAM protocol, clients can prove that no "real" elements were converted into "fake" ones.

Additionally, in the WORM-ORAM scenario, we assume the following:

- The server is allowed to distinguish between record expiration, read and write operations.
- Clients participate correctly in any record expiration protocol. This is reasonable to assume because the regulatory compliance scenario allows clients always to by-pass the server-enforced storage service and store select records elsewhere.

Several participants are of concern. First, clients have incentives to *rewrite history* and alter or completely remove *previously written* records. We note that in the regulatory scenario, there exists an apparent imbalance – clients are assume to correctly store records at the time of their creation – only later does regretting the past becoming a concern. Thus the main focus of WORM

assurances is not to prevent history but rather just its rewriting. In reality, the "regret" time interval between the creation/storage and regretting of a record is non-zero, application-specific, and often quite large. To remove any application dependency, here we consider the strongest WORM guarantees, in which records are not to be altered as soon as they are written.

Second, the storage provider (server) is *curious* and has incentives to illicitly gain information about the stored data and access patterns thereto. As the regulatory storage provider, the server is the main enforcer of WORM semantics and record expiration. Naturally, the server is assumed to not collude with clients illicitly desiring to alter their data. In summary, the server is trusted to run protocols correctly yet it may try to use information obtained from correct runs to obtain undesired information. This assumption is natural and practical as otherwise one can easily imagine a server simply deleting stored records in a denial of service attack. Basic denial of service on the client or server side is not of interest here.

We consider a server S with $O(N)$ storage and a client C with $O(\sqrt{N} \log N)$ local storage. The client stores $O(N)$ items on the server. We denote the regulatory compliance auditor by $\mathcal{A}$.

3.2 Cryptography

We require several cryptographic primitives with all the associated semantic security [32] properties including: a secure, collision-free hash function which builds a distribution from its input that is indistinguishable from a uniform random distribution, a semantically secure cryptosystem (Gen, Enc_k, Dec_k), where the encryption function Enc generates unique ciphertexts over multiple encryptions of the same item, such that a computationally bounded adversary has no non-negligible advantage at determining whether a pair of encrypted items of the same length represent the same or unique items, and a pseudo random number generator whose output is indistinguishable from a uniform random distribution over the output space.

The Decisional Diffie-Hellman (DDH) assumption over a cyclic group G of order q and a generator g states that no efficient algorithm can distinguish between two distributions (g^a, g^b, g^{ab}) and (g^a, g^b, g^c), where a, b and c are randomly chosen from $\mathbb{Z}_q$.

An integer v is said to be a *quadratic residue* modulo an integer n if there exists an integer x such that $x^2 = v \bmod n$. Let QRA be the quadratic residuosity predicate modulo n. That is, $QR(v, n) = 1$ if v is a residue mod n and $QR(v, n) = 0$ if v is a quadratic non-residue.

Given an odd integer $n = pq$, where p and q are odd primes, the quadratic residuosity (QR) assumption states that given n but not its factorization and an integer v whose Jacobi symbol $(v|n) = 1$ it is difficult to determine whether $QR(v, n)$ is 1 or 0.

The Goldwasser, Micali and Rackoff [33] *zero knowledge proof of quadratic non-residuosity* proceeds roughly as

follows. Given two parties A and B, A claims knowledge of $QR(v, n) = 0$, for $(v|n) = 1$. A proves this in zero knowledge to B, that is, without revealing n's factorization. To achieve this, B selects m random values $r_1, .., r_m$ and flips m coins. For each coin c_i, if $c_i = 0$ B computes $x_i = r_i^2 \ mod \ n$ to A, otherwise it computes $x_i = vr_i^2 \ mod \ n$. B sends all computed x_i values to A. A needs to send back the square roots of the quadratic residues it detects in the list $x_1, ..x_m$. If $QR(v, n) = 0$, then A correctly detects the residues. If $QR(v, n) = 1$, all the values received by A will be quadratic residues. A can then cheat only with probability $1/2^m$.

Notations:: Let $n = pq$ be a large composite, where p and q are primes. Let $\phi(n)$ denote the Euler totient of n. We will use $x \in_R A$ to denote the random uniform choice of x from the set A. Given a value m, let $\mathcal{P}(m)$ denote the group of permutations over the set $\{0, 1\}^m$. Let $\mathbf{k} < |n|$ be a security parameter. Let N denote the set of elements stored in the ORAM. Let $\mathbb{W}_m$ be the universe of all sets of m quadratic residues.

4 SOLUTION OVERVIEW

A WORM-ORAM system, consists of two ORAMs (W-ORAM,E-ORAM) and a set of operations (Gen, Enc, Dec, RE, Write, Read, Expire, Shuffle, Audit) that can be used to access the ORAMs. The client needs to store elements at the server while preserving the privacy of its accesses and allowing the server to preserve the data's WORM semantics. W-ORAM serves this purpose: it is used by the client to store $(label, element)$ pairs.

We organize time into epochs: each element stored at the server expires in an integer number of epochs, as determined by the client. The client needs to remember the expiration time of each element stored in the W-ORAM. The client uses the E-ORAM to achieve this, to store expiration times of labels used to index elements stored in W-ORAM. When queried with a time epoch, E-ORAM provides a list of labels expiring in that epoch. The labels are then used to retrieve the expiring elements from the W-ORAM.

The E-ORAM is stored and accessed as a regular ORAM [34]. It is used as an auxiliary storage structure by the client and it needs not be WORM compliant. The W-ORAM on the other hand stores actual elements and needs to be made WORM compliant. The W-ORAM stores two types of elements: "reals" and "fakes". A real element has a quadratic non-residue component, whereas a fake has a quadratic residue. Each time the ORAM is accessed, elements are re-encrypted to ensure access privacy. The client has then to prove in ZK that (i) an element is real or fake and (ii) a re-encrypted element decrypts to the same cleartext as the original element. We now provide a brief overview of each operation described above and follow with a detailed description in the next section.

Gen: Operation executed initially, to generate system parameters for each participant: client, server, auditor.

Enc, Dec, RE: Enc and Dec provide the basic encryption and decryption operations for elements to be stored in the W-ORAM. RE is the W-ORAM element re-encryption operation. RE is needed to ensure that the server cannot distinguish the same W-ORAM element accessed multiple times, while allowing the server to prove in zero knowledge the element's correctness.

Write: Operation used by the client to store an element on the server. The client needs to label the element and determine its expiration time. The client stores the element indexed by the label on the W-ORAM and the label indexed by the element's expiration time in the E-ORAM.

Read: Allows the client to retrieve from the W-ORAM an element indexed by an input label. The operation is based on existing ORAM reading techniques. In addition, it obliviously ensures that the client cannot remove or alter any real element from the W-ORAM.

Shuffle: Re-shuffles a level (provided as input) in the W-ORAM. Based on existing ORAM shuffling techniques, it needs to ensure that the client cannot remove or alter existing W-ORAM elements.

Expire: This operation makes use of both the E-ORAM and W-ORAM to remove all elements from the W-ORAM whose expiration time equals an input expiration time epoch. The operation needs to obliviously convince the server that only expiring elements are removed and no other W-ORAM elements are altered.

Audit: Enables an auditor to access the entire W-ORAM and search for keywords of interest.

In the following, we employ the classic ORAM operations described in Section 2.1 as the APIs for building our solution. Specifically, we use $Read_{ORAM}$ to denote the standard ORAM read operation, taking as input an ORAM and a label and returning an element stored under that label along with the list of all elements removed from the ORAM (including the one of interest). $Write_{ORAM}$ is the standard ORAM write operation, which takes as input a label and an element and stores the element indexed under the label. Note that in the standard ORAM implementation, both operations are performed in the same manner. Their operation is only different for the client. Finally, let OS be the standard ORAM re-shuffle operation (see Section 2), which takes as input a level id and generates a pseudo-random permutation of the re-encrypted elements at that level.

5 W-ORAM ELEMENT ENCRYPTION

We now define the operations for encrypting the elements to be stored in the WORM compliant ORAM.

Gen(k): Generate $p = 2p' + 1$, $q = 2q' + 1$ such that p, p', q, q' are primes. Let $n = pq$. Let G be the cyclic subgroup of order $(p - 1)(q - 1)$. DDH is believed to be intractable in G [35]. Let g be a generator of G. Let a be a random value and let $d = a^{-1} \ mod \ \phi(n)$. Let k be a random key in a semantically secure symmetric cryptosystem. Gen gives k, n, g, $h = g^a \in G$, p, q, a and

d to the client and n, g, h to the server. *Gen* also gives k, p, q, a and d to the auditor.

Enc((x,T_{exp}),k,g,h,G,f)t: Encrypt an element of value x with expiration time T_{exp}, using the client's view of *Gen*'s output as input parameters. The output of the operation is a tuple $(A, B) \in G \times G$ that can be stored on the W-ORAM. If $f = 0$, *Enc* generates a "real" W-ORAM element: the first field of such elements is a quadratic non residue, $QR(A, n) = 0$. The tuple is computed as follows. First, generate a random $r \in \{0,1\}^{\mathbf{k}}$ and use it to compute $M(x) = \{E_k(x), T_{exp}, \text{"real"}, r\}$ where $E_k(x)$ denotes the semantically secure encryption of item x with symmetric key k and "real" is a pre-defined string. The random r is chosen (using trial and failure) such that $QR(M(x), n) = 0$ (quadratic non residue mod n) whose Jacobi symbol is 1. Second, generate a random odd value $b \in_R \{0,1\}^{\mathbf{k}}$ and output the tuple $S(x) \in G \times G$ as

$$S(x) = (A, B) = (M(x)g^{2b}, h^{2b})).$$

$S(x)$ is said to be an "W-ORAM element", whose first field is the "encrypted element" and second field is called the "recovery key". Notice that since $M(x)$ is a QNR, $QR(M(x)g^{2b}, n) = 0$ with (Jacobi symbol) $(M(x)g^{2b}|n) = 1$.

If $f = 1$, *Enc* generates a "fake" W-ORAM element: the first field of fake elements is a quadratic residue, $QR(A, n) = 1$. To compute a fake element, *Enc* generates random $s, k \in_R \{0,1\}^{\mathbf{k}}$ and outputs the tuple $(s^2 \bmod n, k)$. That is, the first field in the pair is a quadratic residue, however, the "recovery key" is useless – does not recover a meaningfull message.

Dec((A,B),d,k): Decrypt a real W-ORAM element, given the secret key $d = a^{-1}$. Compute $M = AB^{-d}$. M has format $\{E, T_{exp}, \text{"real"}, r\}$. The operation outputs the tuple $Dec_k(E), T_{exp}$.

Note that the reason for not using the standard El-Gamal encryption $S(x) = (M(x)g^{ab}, g^b)$ is that the client cannot be forced to publish b or $b^{-1} \bmod \phi(n)$, since b is a per W-ORAM element random value. Without this information the verifier cannot recover $M(x)$. In the encryption used above, a single piece of information, a, can be used to decrypt the entire ORAM. For El-Gamal, each element needs another key for decryption.

RE(A,B): Re-encrypt element (A, B). Choose $u \in_R \{0,1\}^{\mathbf{k}}$, called re-encryption factor. Output pair $(A', B') = (Ag^{2u}, Bh^{2u})$. Note that knowledge of the message M encoded in (A, B) is not required. Alternatively, if M is known such that $A = Mg^{2b}$ and $B = h^{2b}$, then output $(A', B') = (Mg^{2u}, h^{2u})$. Note that u may also be used as an input parameter by $RE((A, B), u)$.

RE(L): Generalization of $RE((A, B))$, where $L = \{(A_1, B_1), .., (A_m, B_m)\}$ is a list of W-ORAM elements. Choose $\bar{u} = \{u_1, .., u_m\}$, such that $u_i \in_R \{0,1\}^{\mathbf{k}}$. $\bar{u}$ is called the re-encryption vector. Output $L' = \{RE((A_i, B_i), u_i)\}_{i=1..m}$. We also use the notation $L' = L\bar{u}$ and call L' a "correct re-shuffle" of L.

We now prove the semantic security of *Enc*.

Theorem 1: Enc is IND-CPA secure.

Proof: Let Q be an adversary that can break the semantic security of *Enc* with advantage ϵ. We then build an adversary Q^* that can break the DDH assumption in G without knowing n's factors, with probability ϵ. Let CH be a challenger. CH interacts with Q^* by sending the triple $(A = g^a, B = g^b, C = g^c)$. Q^* needs to decide whether $c = ab$ or is randomly distributed.

Q^* sends A to Q as the public key (h in our protocol). A then sends to Q^* two messages M_0 and M_1. Q^* picks a bit $\alpha \in_R \{0,1\}$ randomly and sends back to Q the tuple $(M_\alpha B^2, C^2)$. Q sends back its guess for α. If Q guesses correctly, Q^* sends to CH the value 1 ($c = ab$) or 0 (c is random).

When $c = ab$, the tuple $(M_\alpha B^2, C^2)$ is a correct ciphertext of M_α. Then, the interaction between Q^* and Q is correct and the probability of Q^* to output 1 is $1/2 + \epsilon$. When c is distributed randomly, the tuple $(M_\alpha B^2, C^2)$ is independent of α. The probability of Q^* outputting 1 is then $1/2$. Thus, Q^* has advantage ϵ in the DDH game. $\square$

Note that we can similarly prove that RE is IND-CPA. That is, given two encryptions (A_0, B_0) and (A_1, B_1) of any two messages and a re-encryption $RE(A_b, B_b)$, $b \in_R \{0,1\}$ of one of the two encryptions, an attacker cannot guess b with non-negligible probability over $1/2$.

6 THE E-ORAM

The E-ORAM is a standard ORAM, storing labels indexed under expiration time epochs. The E-ORAM needs to provide C with the means to determine how many and which labels expire at a given time epoch and also to insert a new $(epoch, label)$ pair. This is achieved in the following manner. For each T_{exp} value used to index labels in E-ORAM, a *head* value is used to store the number of labels expiring at T_{exp}: $(T_{exp}, (label, counter))$. *label* is the first label that was indexed under T_{exp}. Each of the remaining $c - 1$ labels is stored under a unique index: The ith label's index is (T_{exp}, i), that is, the label's expiration time concatenated with the label's counter at its insertion time.

We now present the most important operations for accessing the E-ORAM, Write and Enumerate.

Write(E-ORAM,T_{exp},label): Record the fact that *label* expires at time T_{exp} (see Algorithm 1, lines 1-11). Read the element currently stored under T_{exp} (line 2). If no such element exists (line 3), generate an element encoding the fact that this label is the first to be stored under T_{exp} (line 4) and store it on the E-ORAM (line 5). Moreover, run a fake E-ORAM access (line 6), whose purpose will become clear in a few lines. If a label is already stored under T_{exp} (line 7), retrieve that label (l) along with the counter c that specifies how many labels are already expiring (stored in E-ORAM) at T_{exp} (line 8). Note that the read operation performed on line 2 removes this element from the E-ORAM. Since now $c+1$ labels expire at T_{exp}, store label l and the incremented

---

Algorithm 1 E-ORAM: Write new label under expiration time. Enumerate all labels expiring at a given time. V is the list of elements returned by a Read.

---

1. **Write**$(\mathtt{E-ORAM:ORAM, T_{exp}:int, lbl:id})$
2. $(\mathtt{e,V}) := \mathtt{Read_{ORAM}(E-ORAM, T_{exp})}$;
3.   **if** $(\mathtt{e = null})$ **then**
4. $\mathtt{e' := E_k(lbl, 1)}$;
5. $\mathtt{Write_{ORAM}(T_{exp}, e')}$;
6. $\mathtt{Write_{ORAM}(null, null)}$;
7.   **else**
8. $\mathtt{(l, c) := D_k(e)}$;
9. $\mathtt{Write_{ORAM}(T_{exp}, E_k(l, c+1))}$;
10. $\mathtt{Write_{ORAM}((T_{exp}, c+1), E_k(lbl))}$;
11. **fi**
12. **end**

---

12. **Enumerate**$(\mathtt{E-ORAM:ORAM, T_{exp}:int})$
13. $\mathtt{L:id[]}$; #store result labels
14. $\mathtt{L := \emptyset}$;
15. $\mathtt{(e, A) := Read_{ORAM}(E-ORAM, T_{exp})}$;
16. **if** $(\mathtt{e! = null})$ **then**
17. $\mathtt{(l, c) := D_k(e)}$;
18. $\mathtt{L := L \cup l}$;
19.   **for** $(\mathtt{i := 2; i \leq c; i++})$ **do**
20. $\mathtt{(e, A) := Read_{ORAM}(E-ORAM, (T_{exp}, i))}$;
21. $\mathtt{l := D_k(e)}$;
22. $\mathtt{L := L \cup l}$;
23.   **od**
24. **fi**
25. **return** L;
26. **end**

---

counter in the E-ORAM under T_{exp} (line 9). Finally, store the input *label* under an index consisting of a unique value: T_{exp} concatenated with $c + 1$. This will allow the client to later enumerate all labels expiring at T_{exp} (see next). The reason for the fake E-ORAM write performed in line 6 is to make the two cases indistinguishable to the server: the E-ORAM is always accessed twice, independent of how many elements expire at T_{exp}.

Enumerate(E-ORAM,T_{exp}): Retrieve all the labels in E-ORAM that expire at T_{exp} (see Algorithm 1, lines 12-26). First, initialize the result label list (line 13). Then, read the head label stored under T_{exp} along with the counter of labels expiring at T_{exp} (lines 14,16). If such an element exists (line 15), record the head label (line 17). Then, for each of the $c - 1$ ($i = 2, .., c$) remaining labels, retrieve their actual value by reading from E-ORAM the element stored under a unique index consisting of T_{exp} concatenated with i. Note that *Enumerate* removes all labels expiring at T_{exp} from E-ORAM ($Read_{ORAM}$ removes accessed elements).

7 W-ORAM ACCESS OPERATIONS

7.1 Generating Labels

Elements in the standard ORAM model are stored as a pair $(label, value)$, where *label* may denote a memory location or the subject of an e-mail. In our case to prevent the server from launching a dictionary attack, we use the a $Label(label, lkey)$ operation to generate labels. Besides the input *label*, *Label* also uses a (random) labeling key, which is used to define a pseudo-random function F_{lkey}. The output of *Label* coincides then with the output of $F_{lkey}(label)$. We now describe the main W-ORAM accessing operations.

7.2 Writing on the Server

Write((W-ORAM,E-ORAM,v,l,T_{exp},params): Store on the server a value v under a label l, with expiration time T_{exp}, using as input also the client's view of Gen's output, $params = k, g, h, G$ (see Algorithm 2 for the pseudo-code of this operation). Generate a new *label* as described above (line 2) and call Enc to produce a W-ORAM tuple (A_u, B_u) (line 3). Generate a non-interactive zero knowledge proof of $QR(A_u, n) = 0$ (A_u's quadratic non-residuosity). If the proof verifies (line 5) the server inserts the tuple (A_u, B_u) in the top level of the W-ORAM (line 6) and stores *label* under the tuple's expiration time T_{exp} in E-ORAM (see Section 6). Otherwise, the server aborts the protocol (line 10).

7.3 Reading from the Server

Read(W-ORAM,label): Using as input the W-ORAM and a *label*, return an element of format $(label, x, T_{exp})$ (see Algorithm 3). Perform on W-ORAM a standard ORAM read on the desired *label* (line 2), returning both the W-ORAM element R of interest and the list L of elements (containing R) removed from the W-ORAM. If *label* is stored in the W-ORAM (line 3), the client computes $U = (A_u, B_u)$, a re-encryption of R (line 3) and calls ZK-POR to prove in zero knowledge that U is a re-encryption of the only real element in L (line 4). ZK-POR is described in detail in Section 7.3.1. The server verifies in ZK that $QR(A_u, n) = 0$ and also the validity of the ZK-POR proof. If the proofs are valid (line 5), the server inserts U in the first level of the W-ORAM (lines 6-7). The client decrypts the desired element R and returns the result (line 8). If any proof fails (line 9) the server restores the W-ORAM to the state before the start of Read and returns error (lines 10-11).

7.3.1 Zero Knowledge Proof of ORAM Read.

We now present ZK-POR, the zero-knowledge proof of WORM compliance of the read operation performed on the W-ORAM. ZK-POR takes as argument the list L of elements removed from W-ORAM in line 2 of Algorithm 3 and U, the re-encryption of the real element from L. For simplicity of exposition, let us assume that L also contains the elements (scanned but not removed) from the first level of W-ORAM. Let m denote the number of elements in L, $m = O(\log N)$.

Let $L = \{(s_1^2, k_1), ..., (s_{r-1}^2, k_{r-1}), S(x_r), (s_{r+1}^2, k_{r+1}), .., (s_m^2, k_m)\}$ where the elements are listed in the order in which they were removed from the W-ORAM. The

---

Algorithm 2 W-ORAM: Write value v expiring at T_{exp}.

---

1. **Write**$(W - ORAM : ORAM, E - ORAM : ORAM,$
         $v : string, l : id, T_{exp} : int)$
2. $label := newLabel(l, lkey);$
3. $(A_u, B_u) := Enc(label, v, T_{exp}, params);$
4. $ZKP := getQNRProof(A_u, n);$
5. **if** $(verify(ZKP, A_u) = 1)$ **then**
6. $T_0 := getLevel(W - ORAM, 1);$
7. $insert(T_0, (A_u, B_u));$
8. $Write(E - ORAM, T_{exp}, label);$
9. **else**
10.   return error;
11. **fi**
12. **end**

Algorithm 3 W-ORAM: Read $label$.

---

1. **Read**$(W - ORAM : ORAM, label : id)$
2. $(R, L) := Read_{ORAM}(W - ORAM, label);$
3. $U := (A_u, B_u) := RE(R);$
4. $Proof := ZK - POR(L, U);$
5. **if** $(verifyQNR(A_u, n)$
   $\&\ verify(Proof, L, U))$ **then**
6. $T_0 := getLevel(W - ORAM, 1);$
7. $insert(T_0, U);$
8.   return $Dec(R, d, k);$
9. **else**
10. $undo(W - ORAM);$
11.   return error;
12. **fi end**

---

client is interested in the item from the rth ORAM layer, $R = S(x_r)$. Let $S(x_r) = (M(x_r)g^{2t_r}, h^{2t_r}) = (A_r, B_r)$. Its first field is a quadratic non-residue. All other elements from L are fakes – their first field is a quadratic residue. Let $U = RE(R) = (M(x_r)g^{2u}, h^{2u}) = (A_u, B_u)$ be the re-encryption of $S(x_r)$. The following steps are executed s times between the client and the server.

Step 1: Proof Generation: The client selects a random permutation $\pi \in_R \mathcal{P}(m)$. The client generates $\bar{w} = \{w_1, .. w_m\}$, where each $w_i \in_R \{0, 1\}^m$ is odd and generates the proof list $P = \pi(L\bar{w})$. Let $P = \pi\{(s_1^2 g^{2w_1}, k_1 h^{2w_1}), .., (A_r g^{2w_r}, B_r h^{2w_r}),$ $.., (s_m^2 g^{2w_m}, k_m h^{2w_m})\}$, where, $(A_r g^{2w_r}, B_r h^{2w_r})$ is a re-encryption of $S(x_r)$. The client sends P to the server. The client locally stores $(w_i, s_i g^{w_i})$, $i = 1..m$. As assumed in the model, The client has $O(\sqrt{N} \log N)$ storage which is sufficient to store $m = O(\log N)$ values.

Step 2: Proof Validation: The server flips a coin b. If b is 0, the client reveals $w_1, .., w_m$. The server verifies that all w_i are odd and $\forall (A_i, B_i) \in L$, $(A_i g^{2w_i}, B_i h^{2w_i}) \in P$. If b is 1, the client sends to the server the values $s_i g^{w_i}$, $i = 1..m, i \neq r$ along with the value $\Gamma = (t_r + w_r - u)$. Note that given $s_i^2\ mod\ n$ and n's factorization, the client can easily recover s_i. The server verifies first that $(s_i g^{w_i})^2$, $i = 1..m, i \neq r$ occurs in the first field of exactly one tuple in P. That is, $m - 1$ of the elements from P are fakes. The server then verifies that $(A_r g^{2w_r}, B_r h^{2w_r}) = RE((A_u, B_u), \Gamma)$. If any verification fails, the server outputs "error" and stops.

7.3.2 Analysis

We now prove the following results.

Theorem 2: A correct execution of Read from W-ORAM has $O(\log N)$ complexity.

Proof: The overall client overhead for a single *Read* operation is

$$T_{Read}^c \approx T_{dsk} \log N + T_{RE} + s T_{RE} \log N,$$

where T_{dsk} is the time to access the disk, T_{RE} is the re-encryption time (two modular exponentiations) and T_e is the modular exponentiation time. The first factor corresponds to the $Read_{ORAM}$ in line 2 of Algorithm 3,

consisting of $\log N$ disk accesses. The second factor is due to the re-encryption in line 3. The third factor is from the proof generation step of ZK-POR. Note that $\log N$ denotes the number of ORAM levels and the number of elements for which the proof sets are built. Thus, the client overhead is $O(\log N)$. The server only participates during the proof validation steps and its cost is

$$T_{Read}^s \approx (s/2) T_{RE} \log N + (s/2)(T_{mul} \log N + T_{RE})$$

where $s/2$ is the number of bits b that come up 0 (and 1). Thus, the server's overhead is also $O(\log N)$. $\square$

Theorem 3: ZK-POR is a zero knowledge proof system of Read $\in$ WORM. That is, Read is WORM compliant.

Proof: We first need to prove that ZK-POR is a proof system (correct and sound) then prove that ZK-POR is zero knowledge. $\square$

Theorem 4: ZK-POR is complete.

Proof: First, note that it is straightforward to see that if the client is honest, it is able to build the answer to the server's challenges. We now prove that if L contains $m - 1$ fake elements and one real element and U is a re-encryption of the real element from L, an honest server (one following the protocol properly) will be convinced of this fact by an honest client.

If $b = 0$, the client reveals all values $w_1, .., w_m$ used to transform L into P. Let $< A_i, B_i >$ be a tuple in L and the corresponding factor w_i. If $(A_i g^{2w_i}, B_i h^{2w_i}) \in P$ exactly once for all $i = 1..m$, the server is convinced that P is a re-encryption of L. Also, since $QR(A_i, n) = QR(A_i g^{2w_i}, n)$, the server is convinced that the quadratic residuosity of the elements in L has been preserved in P. If $b = 1$, the client first proves that $m - 1$ of the values in P are fakes. That is, the client proves that the first field of $m - 1$ elements is a quadratic residue. For the remaining value in P, $(A_r g^{2w_r}, B_r h^{2w_r})$, the client reveals $\Gamma = t_r + w_r - u$. The server verifies that $A_u g^{2\Gamma} = M(x_r)g^{2u}g^{2(t_r + w_r - u)} = A_r g^{2w_r}$ and that $B_u h^{2\Gamma} = h^{2u}h^{2(t_r + w_r - u)} = B_r h^{2w_r}$ This proves that $U = (A_u, B_u)$ is a re-encryption of an element of P. Moreover, the fact that $QR(A_u, n) = 0$ and $QR(g^{2\Gamma}, n) = 1$ implies that $QR(A_u g^{2\Gamma}, n) = QR(A_r g^{2w_r}, n) = 0$. Thus, U is a re-encryption of the real element from P.

The conjunction of the two cases ($b = 0, 1$) convinces the server that the list L contains $m - 1$ fakes and the element to be inserted back in the ORAM is an obfuscation of the real element from L. $\square$

Theorem 5: ZK-POR is sound.

Proof: We need to prove that if the statement is false, no cheating client can convince the honest server that it is true, except with some small probability. The statement is false if the client (i) removes more than one real element from the ORAM, (ii) makes A_u, the first field of U a quadratic residue or (iii) makes the recovery key B_u of U reveal a different value from the $M(x_r)$, that is revealed by the recovery key B_r of $S(x_r)$.

We start with the observation that the server generates the coin b independently and the client has negligible advantage over $1/2$ in guessing the coin's outcome.

The first cheating attempt (i) would allow the client to surreptitiously remove one or more real elements from the ORAM, since it only has to write back a single element. Let us assume that the client wants to remove two real elements, that is L contains $m - 2$ fakes and 2 real elements. In order for this attack to succeed, the client can generate the proof list P such that it either (i.a) contains $m - 1$ fake elements OR (i.b) is a re-encryption of L. In case (i.a) where P contains $m-1$ fakes (quadratic residues) if the coin b comes up 0, the client needs to reveal the obfuscating factor for one QNR element in L that is now a QR in P. The only way this can occur is if the obfuscating factor is a QNR. However, each obfuscating factor is squared by the server before verification. In case (i.b), where P is built to be a re-encryption of L, if the coin b comes up 1, the client will fail to prove that one element in P (corresponding to one of the real elements in L) is a quadratic residue. Thus, the client can only cheat with probability $1/2$ per each independent proof list P.

In case (ii), the client attempts to change U into a fake element, allowing it later to remove it by proving it is a quadratic residue. Notice however that in Step 0, the client proves in ZK the fact that $QR(U, n) = 0$. In case (iii) the client attempts to destroy the real element before inserting it back in the ORAM - it will not be recoverable later. Let us assume that the element the client wants to write back is $U' = (M'(x'_r)g^{2u}, h^{2u'}) = (A'_u, B'_u)$, where $M'(x'_r) \neq M(x_r)$ and $u' \neq u$. Then, the client can build P either to (iii.a) contain an element that maps to U' or such that (iii.b) it is a re-encryption of L.

Let us first consider case (iii.a), where P contains an element $R = (M'(x'_r)g^{2w'_r}, h^{2w'_r})$, such that the element R is a re-encryption of U'. That is, there exists Γ such that $A'_u g^{2\Gamma} = M'(x'_r)g^{2w'_r}$ and $B'_u h^{2\Gamma} = h^{2w'_r}$. With straightforward math this leads to $w'_r - u = w''_r - u'$. Now, if the coin b comes up 0, the client needs to produce a single value w_r proving that $S(x_r)$ from L maps into element R from P. If the client can produce such a w_r, we have that $M(x_r)g^{2t_r}g^{2w_r} = M'(x'_r)g^{2w'_r}$ and $h^{2t_r}h^{2w_r} = h^{2w'_r}$. From these equations we obtain $M(x_r)g^{2w''_r} = M'(x'_r)g^{2w'_r}$. Since $w'_r - u = w''_r - u'$, we

have that $M'(x'_r) = M(x_r)g^{2w''_r - 2w'_r} = M(x_r)g^{2u' - 2u}$. Then, $U' = (M(x_r)g^{2u'}, h^{2u'})$, which is a correct re-encryption of $S(x_r)$.

Let us now consider case (iii.b), where the client builds the proof list P to be a re-encryption of L. This implies that the proof P contains an element $R = (A_r g^{2w_r}, B_r h^{2w_r}) = (M(x_r)g^{2(t_r + w_r)}, h^{2(t_r + w_r)})$ which is a re-encryption of $S(x_r)$. In this case, if b comes up 1, the client needs to provide a value Γ to prove that the element R maps into the element to be written back, U'. That is, Γ needs to satisfy the following two equations, $M'(x'_r)g^{2u}g^{2\Gamma} = M(x_r)g^{2(t_r + w_r)}$ and $h^{2u'}h^{2\Gamma} = h^{2(t_r + w_r)}$. Using straightforward math, this leads to $M'(x'_r) = M(x_r)g^{2(u' - u)}$, which implies that $U' = (M(x_r)g^{2u'}, h^{2u'})$. Thus, if the client is able to solve the $b = 1$ challenge, U' is a correct re-encryption of $S(x_r)$. $\square$

Theorem 6: ZK-POR is zero-knowledge.

Proof: It is straightforward to see that ZK-POR conveys no knowledge to an honest verifier. Even if the server has access to the plaintext message read (and written back) from the ORAM, due to the semantic security of the encryption method Enc of $WOES$, the server can associate the element U with the corresponding element from L or any proof list P with probability only negligibly larger than $1/m$.

We need to show that ZK-POR conveys no knowledge to any verifier, even one that deviates arbitrarily from the protocol. We prove this by following the approach from [33], [36]. Specifically, let S^* be an arbitrary, fixed, expected polynomial time ITM. We generate an expected polynomial time machine M^* that, without being given access to the client, produces an output whose probability distribution is identical to the probability distribution of the output of $< C, S^* >$.

We now build M^* that uses S^* as a black box many times. Whenever M^* invokes S^*, it places input $x = (L_0, L_1)$ on its input tape IT_S and a fixed sequence of random bits on its random tape, RT_S. The input x consists of $L_0 = L = \{(s_1^2, k_1), .., S(x_r), .., (s_m^2, k_m)\}$ which is the list of elements read and removed from the ORAM and $L_1 = \{(v_1^2, k'_1), .., (v_{m-1}^2, k'_{m-1}), U\}$. U is the element to be written back on the ORAM and $v_1, .., v_{m-1}, k'_1, .., k'_{m-1}$ are random numbers chosen by M^*. Since S^* is only expected polynomial time, the random bits for RT_S are selected as specified in [36]. The content of the input communication tape for S^*, CT_S will consist of pairs (P, π), where P is a set and $\pi \in \mathcal{P}(m)$.

The output of M^* consists of two tapes: the random-record tape RT_M and the communication-record tape CT_M. RT_M contains the prefix of the random bit string r read by S^*. The machine M^* works as follows (round i):

Step 1: M^* chooses a random bit $a \in_R \{0, 1\}$ and m random values $w_1, .., w_m$. Let $\bar{w} = \{w_1, .., w_m\}$. M^* also chooses a random permutation $\pi \in_R \mathcal{P}(m)$. If $a = 0$, M^* computes $P = \pi(L_0\bar{w})$. Notice that M^* can compute P without knowing $s_1, .., s_m$ or x_r and t_r. If $a = 1$, M^* com-

putes $P = \pi\{(w_1^2, k_1''), .., (w_{m-1}^2, k_{m-1}''), RE(U, w_m)\}$, for randomly chosen $k_1'', .., k_{m-1}''$. Similarly, M^* can compute P without knowing the x and t values of $U = S(x, t)$, but only U's Y_u and Θ_u fields.

Step 2: M^* sets $b = S^*\{x, r; P_1, \pi_1, .., P_{i-1}, \pi_{i-1}, P\}$. That is, b is the output of S^* on input x and random string r after receiving $i-1$ pairs P_j, π_j), $j = 1..i-1$ and proof P on its communication tape CT_S. We have the following three cases.

(Case 1). $a = b = 0$. M^* can produce $w_1, .., w_m$ and π to prove that $P = \pi(L_0 \bar{w})$. M^* sets P_i to P, π_i to π and b_i to b, appends the triple (P_i, π_i, b_i) to CT_M and proceeds to the next round (i+1).

(Case 2). $a = b = 1$. M^* can produce $w_1, .., w_m$ and π_i that prove that $m-1$ elements of P are fakes ($w_1^2, .. w_{m-1}^2$ occur in the first field of $m-1$ elements of P) and that the mth element of P is a re-encryption of U. M^* sets P_i to P, π_i to π and b_i to b, appends the triple (P_i, π_i, b_i) to CT_M and proceeds to the next round (i+1).

(Case 3). $a \neq b$. M^* discards all the values of the current iteration and repeats the current round (Step 1 and 2).

If all rounds are completed, M^* halts and outputs (x, r', CT_M), where r' is the prefix of the random bits r scanned by S^* on input x. We first prove that M^* terminates in expected polynomial time and then that the output distribution of M^* is the same as the output distribution of S_* when interacting with the client, on input (L_0, L_1).

Lemma 1: M^* terminates in expected polynomial time.

Proof: Given L and U, during the ith round P is either built from L or from U. During each run of round i, the bit a is chosen independently. Then P is also chosen independently (built from L or from U). This implies that the probability that $a = b$ is $1/2$ and the expected number of repetitions of round i is 2. S^* is expected polynomial time, which implies that M^* is also polynomial time. □

Lemma 2: The probability distribution of $< C, S^* > (L_0, L-1) >$ and of $M^*(L_0, L_1)$ are identical.

Proof: The output of $< C, S^* > (L_0, L_1) >$ and of $M^*(L_0, L_1)$ consists of a sequence of t triples of format (P, π, b). Let $\Pi_{M^*}^{(x,r,i)}$ and $\Pi_{CS^*}^{(x,r,i)}$ be the probability distributions of the first i triples output by M^* and $< C, S^* >$. We need to show that for any fixed random input r, $\Pi_{M^*}^{(x,r,t)} = \Pi_{CS^*}^{(x,r,t)}$. We prove this by induction. The base case, where $i = 0$, holds immediately. In the induction step we assume that $\Pi_{M^*}^{(x,r,i)} = \Pi_{CS^*}^{(x,r,i)} = T^{(i)}$. We need to prove that the $i + 1$st triples in $\Pi_{M^*}^{(x,r,i+1)}$, denoted by $\Pi_{M^*}^{(i+1)}$ and in $\Pi_{CS^*}^{(x,r,i+1)}$, denoted by $\Pi_{CS^*}^{(i+1)}$ have the same distribution.

For an ORAM element U, we define $\mathbb{L}_m(U)$ to be the universe of sets of format $\pi\{(v_1^2, k_1), .., (v_{m-1}^2, k_{m-1}), RE(U)\}$, for all $\pi \in \mathcal{P}(m)$, all $v_1, k_1, .., v_{m-1}, k_{m-1} \in_R \mathbb{Z}_n^*$ and all re-encryptions of element U. Then, we show that $\Pi_{M^*}^{(i+1)}$ and $\Pi_{CS^*}^{(i+1)}$ are uniform over the set

$$V = \{(P, \pi, b) | b = S^*(x, r, T^{(i)} || P) \wedge ((P = \pi(L_0 \bar{w}),$$

$$\bar{w} \in_R \mathbb{W}_m, if \ b = 0) \vee (P \in_R \mathbb{L}_m(U), if \ b = 1)\}$$

For $\Pi_{CS^*}^{(i+1)}$, this is the case, since when (i) $b = 0$, P is built as a re-encryption of L_0 and when (ii) $b = 1$, P contains a re-encryption of U and $m - 1$ fakes (elements whose first field is a quadratic residue). $\Pi_{M^*}^{(i+1)}$ is also uniformly distributed in V since M^* chooses a random permutation π, a random vector $\bar{w}$ and a random b, builds P according to b and outputs (P, π, b) only if $b = S^*(x, r, T^{(i)}, P)$. In case there is output, it is uniformly distributed in V. □

Given the above two lemmas, we have that M^* terminates in expected polynomial time and its output has the same distribution as the output of the interaction between S^* and a client. Thus, the result of the theorem follows. □

Note that the soundness property of ZK-POR ensures that a cheating client can remove an element from the ORAM during the Read operation without being detected with probability at most $1/2^s$.

8 ACCESS INDISTINGUISHABILITY

During the W-ORAM Read operation described in the previous section, the client accesses and removes $O(\log n)$ elements from the W-ORAM and writes back one element. During a Write the client only adds one element. This, in effect, breaks the access indistinguishability property of the classic ORAM solution: the server can distinguish between a read and a write.

In this section we provide a solution that combines the observable pattern of the previous Read and Write operations to create a single access operation that can be used to both read and write on the W-ORAM. As an overview, the Access operation first accesses the E-ORAM (as done in Write). Then, it accesses all the elements in the W-ORAM's first level and accesses and removes one element from all its subsequent levels (as done by a Read). Note that for a Write, all accessed elements need to be fakes. Access then writes two elements back on the W-ORAM, R and N. For a Read access, R is the re-encryption of the real element accessed in the W-ORAM and N is a new fake element. For a Write access, R is a re-encryption of one of the fakes accessed and removed previously and N is Write's input, the new element to be appended on the W-ORAM. Finally, Access proves (in zero knowledge) that R is a re-encryption of one of the elements removed from the W-ORAM and that all the other elements (removed from the W-ORAM) are fakes.

Note that we do not need to prove to the server anything about N. If N is real, the ZK proof will convince the server that the client cannot remove it later surreptitiously. If it is a fake, it stores no useful information. Moreover, the client does not need to prove that R is the re-encryption of a real element. As long as all the other elements removed from the W-ORAM are fakes, R can be the re-encryption of a fake (as will certainly be the case for a Write) and no useful information is removed from the W-ORAM.

Access((W-ORAM,E-ORAM,v,l,T_{exp},params):

If Access=Write, use l to generate a new *label* (as described in Section 7.1) and insert *label* under T_{exp} in the E-ORAM (using the Write operation described in Section 6). If Access=Read, perform a fake Write on the E-ORAM, consisting of three random accesses to the E-ORAM (one for a read and two for writes, see Section 6). Then, access all the elements in the top level of the W-ORAM and access and remove one element from each subsequent level. If Access=Write, all removed elements have to be fakes. If Access=Read, one of them is real (unless the read element was found in the top level). Let $L = \{(s_1^2, k_1),...,(s_{r-1}^2, k_{r-1}), S(x_r),$ $(s_{r+1}^2, k_{r+1}),...,(s_m^2, k_m)\}$ be the list of elements accessed in the W-ORAM, where $S(x_r)$ may be the real element accessed by a Read or a fake if accessed by a Write. Then, generate two elements R and N and send them to the server. If Access=Write, R=$RE(S(x_r))$ and N=$Enc((v, T_{exp}), k, g, h, G, 0)$ is the element to be written. If Access=Read, $R = RE(S(x_r))$ and N=$Enc((null, null), k, g, h, G, 1)$ is a fresh fake (see Section 5). The zero knowledge proof then proceeds exactly as ZK-POR.

8.1 Analysis

It is straightforward to see that Access correctly implements both W-ORAM Read and Write operations. We need to prove now that reads and writes performed using Access are indistinguishable.

Theorem 7: The server cannot decide whether an Access operation is a Read or a Write with probability significantly larger than $1/2$.

Proof: Let L_{Rd} and L_{Wr} be the lists of elements accessed by a read and a write respectively. Let R_{Rd} and N_{Rd} be the elements written back by a Read and R_{Wr} and N_{Wr} the elements written back by a Write access. We enumerate the ways for the server to distinguish a Read from a Write: (i) distinguish L_{Rd} from L_{Wr}, (ii) distinguish R_{Rd} and N_{Rd} from R_{Wr} and N_{Wr} and (iii) determine which element in $L_{Rd/Wr}$ is a re-encryption of $R_{Rd/Wr}$.

For the case (i), note that L_{Wr} is a random access sequence. If the server can distinguish L_{Rd} from L_{Wr} with probability significantly larger than $1/2$, then it can distinguish L_{Rd} from a random sequence. This would allow us to build an attacker that has a significant advantage in guessing which element is of interest in L_{Rd}, thus contradicting the Oblivious RAM property of ORAMs [21]. For the case (ii), consider the fact that R_{Rd} is a real element and N_{Rd} is a fake, whereas R_{Wr} is a fake and N_{Wr} is a real element. If the server has an advantage in distinguishing between a real and a fake element, we can immediately build an attacker that breaks the Quadratic Residuosity Assumption with the same advantage. Finally, note that the use of ZK-POR prevents the server from learning which element in $L_{Rd/Wr}$ is a re-encryption of $R_{Rd/Wr}$. $\square$

9 SHUFFLING THE W-ORAM

When the $l-1$th level of W-ORAM stores more than 4^{l-1} elements, due to element insertions occurring during Read operation, the level needs to be spilled over into level l. Let $T[l]$ denote the list of elements stored in the W-ORAM at the l-th level. The l-th level then needs to be filled with fakes. The fakes are needed to ensure that subsequent Read accesses will not run out of fakes (see [34] for more details). The l-th level then needs to be obliviously permuted, using only $O(\sqrt{N} \log N)$ client space. Let $T^{new}[l]$ denote the re-shuffled l-th level elements. Due to the WORM semantics, the client also needs to prove that the reshuffle is correct: (i) $T^{new}[l]$ is a re-encryption of the old $T[l]$ and (ii) $|T^{new}[l]| - |T[l]| - |T[l-1]|$ elements from $T^{new}[l]$ are fakes. Shuffle performs this operation.

Shuffle(W-ORAM,l): Uses as input the W-ORAM and the index of a level to reshuffle the corresponding level (see Algorithm 4). First, spill the content of level $l-1$ into level l (lines 3-6) and compute an oblivious permutation of the new level l. Then, build its ZK proof of correctness, ZK-PRS, detailed in the following (see Algorithm 4, lines 7-38 for pseudo-code).

9.1 Zero Knowledge Proof of Re-Shuffle.

Similar to ZK-POR (see Section 7.3.1), ZK-PRS consists of s rounds executed by the client and the server. During each round, a proof list P_j is built by the client (line 14 of Algorithm 4). P_j has the same number of elements as $T^{new}[l]$, $O(N)$. The client builds the list $T^{new}[l]$ and each of the s proofs P in the following steps. Initially, $T^{new}[l]$ and each proof list P_j is stored as an empty list at the server. The client generates a symmetric key k for the (G, E, D) cryptosystem.

Step 1: Element Re-Encryption: First, the client takes each element from $T[l]$ and stores a re-encrypted version in $T^{new}[l]$ and in each proof P_j (lines 7-13). That is, for each element $S_i = (A_i, B_i) \in T[l]$ (stored at the server), the client generates fresh random odd values $u_i, w_i \in \{0, 1\}^{\mathbf{k}}$ (lines 9 and 12) and produces one element S_i' to be inserted in $T^{new}[l]$ (line 10) $S_i' = E_k(A_i g^{2u_i}, B_i h^{2u_i})$ and one element P to be inserted in P_j (line 13) $P = E_k(A_i g^{2w_i}, B_i h^{2w_i}, "mv", \Gamma_1[i], \Gamma_2[i])$ where $\Gamma_1[i] = -w_i$ and $\Gamma_2[i] = (u_i - w_i)$. The string "mv" denotes that this proof element corresponds to an element from $T[l]$ moved to $T^{new}[l]$.

Step 2: Fake Insertion: The client adds f fake elements (lines 14-22). For each of the f fakes to be inserted in $T^{new}[l]$, the client generates two random values $s_i, k_i \in_R \{0, 1\}^{\mathbf{k}}$ (line 16), $i = 1..f$, where w_i is odd. The client then adds an element $E_k(s_i^2, k_i)$ in $T^{new}[l]$ (lines 17-18). It then generates a random value $w_i \in_R \{0, 1\}^{\mathbf{k}}$ (line 20) for each proof list P_j and appends an element $E_k(s_i^2 g^{2w_i}, k_i h^{2w_i}, "add", \Gamma_1[i], \Gamma_2[i])$ to P_j (lines 21-22). $\Gamma_1[i] = s_i g^{w_i}$, $\Gamma_2[i] = (u_i - w_i) \bmod \phi(N)$ and the string "add" denotes that this proof element is a fake added to level l.

Algorithm 4 Shuffle of level l.

1.**Shuffle**($W - ORAM : ORAM, l : int$)
2. $T^{new}[l] : string[] \#new\ level\ l\ array$

$\#spill\ T[l-1]\ into\ T[l]$
3. $T[l-1] := getLevel(W - ORAM, l-1);$
4. $T[l] := getLevel(W - ORAM, l);$
5. $T[l] := T[l-1] \cup T[l];$
6. $T[l-1] := \emptyset;$

$\#re-encrypt\ elements\ from\ T[l]$
7. **for** $(i := 1; i \leq |T[l]|; i++)$ **do**
8. $e := T[l][i];$
9. $u[i] := genRandom();$
10. $T^{new}[l][i] := E_k(RE(e, u[i]));$
11.   **for** $(j := 1; j \leq s; j++)$ **do**
12. $w[i] := genRandom();$
13. $P_j[i] := E_k(RE(e, w[i]),$
        $"mv", u[i], u[i] - w[i]);$
$\#add\ fakes$
14.$f := fakeCount(T[l]);$
15.**for** $(i := 1; i \leq f; i++)$ **do**
16. $(s[i], k[i]) := genRandom();$
17. $e := (s[i]^2, k[i]);$
18. $append(T^{new}[l], E_k(e));$

19.   **for** $(j := 1; j \leq s; j++)$ **do**
20. $w[i] := genRandom();$
21. $re := RE(e, w[i]),$
      $"add", s[i]g^{w[i]}, u[i] - w[i]);$
22. $append(P_j[i], E_k(re));$

$\#Shuffle\ T^{new}[l]\ and\ proofs$
23.$T^{new}[l] := OS(T^{new}[l]);$
24.**for** $(j := 1; j \leq s; j++)$ **do**
25. $P_j := OS(P_j);$
$\#decrypt\ shuffled\ elements$
26.**for** $(i := 1; i \leq |T^{new}[l]|; i++)$ **do**
27. $e := T^{new}[l][i];$
28. $T^{new}[l][i] := D_k(e);$
29.   **for** $(j := 1; j \leq s; j++)$ **do**
30. $e := P_j[i];$
31. $(A, B, str, C, D) := D_k(e);$
32. $P_j[i] := (A, B, E_k(str, C, D));$
$\#proof\ verification\ step$
34.**for** $(j := 1; i \leq s; i++)$ **do**
35.   **if** $(!verify(T[l], T^{new}[l], P_j))$ **then**
36. $undo(W - ORAM, l-1, l);$
37. $return\ error;$

$\#commit\ new\ level$
38.$T[l] := T^{new}[l];$

Note that $\Gamma_1[i]$ and $\Gamma_2[i]$ are used to keep track of the correspondence between the ith element of each P_j and its re-encryptions in $T[l]$ and $T^{new}[l]$ after the list reshuffle step (see next).

Step 3: List Reshuffle: At the end of the set generation step, the client and the server have a one-to-one correspondence between each element in $T^{new}[l]$, each element in each P_j and each element in $T[l]$. The client then calls the oblivious scramble, OS, procedure using $T^{new}[l]$ and each P_j as inputs (lines 23-25). During the OS call, elements read from $T^{new}[l]$ and P are decrypted (using k) and re-encrypted before being written back. Due to the semantic security properties of the encryption scheme employed, at the end of the OS, the server can no longer map elements from $T[l]$ to elements in the reshuffled $T^{new}[l]$ and P_j sets.

Step 4 - Decryption: The client reads each element from the reshuffled $T^{new}[l]$ list, decrypts the element and writes it back in-place (lines 26-28). The client reads each element from each proof list P_j, decrypts it and writes back $(A_ig^{2w_i}, B_ih^{2w_i}, E_k(str, \Gamma_1[i], \Gamma_2[i]))$, where str is either "mv" or "add" – moved or added fake (lines 29-32).

Step 5 - Proof Verification: The server verifies each proof list P_j (lines 34-37). If any verification fails, restore the W-ORAM to the state at the beginning of the operation and return error (lines 36-37). Each verification, for a proof list P, works as follows.

The server flips a coin b. If $b = 0$, the server asks the client to prove that P is a valid reshuffle of $T[l]$ and all the remaining elements in P are fakes. For this, the client reads each element of P, $(A_ig^{2w_i}, B_ih^{2w_i}, E_k(str, \Gamma_1[i], \Gamma_2[i]))$, retrieves $\Gamma_1[i]$ and

sends to the server, $A_ig^{2w_i}$, $B_ih^{2w_i}$, str and $\Gamma_1[i]$. If $str =$ "mv", the server first verifies that indeed $\Gamma_1[i]$ is an odd number, then verifies that $RE((A_ig^{2w_i}, B_ih^{2w_i}), \Gamma_1[i])$ appears in $T[l]$ exactly once. If $str =$ "add", the server verifies that $\Gamma_1[i]^2$ is the first field of exactly one tuple in $T^{new}[l]$. If at the end of this step the client has proved that $|T[l]|$ elements from $T^{new}[l]$ are re-encryptions of the elements from $T[l]$ and that f elements from $T^{new}[l]$ are fakes, the server continues. Otherwise it outputs "error" and stops.

If $b = 1$, the client needs to prove that P is a valid reshuffle of $T^{new}[l]$. For this, the client reads each element from P, recovers $\Gamma_2[i]$ and sends to the server the values $A_ig^{2w_i}$, $B_ih^{w_i}$ and $\Gamma_2[i]$. The server verifies that $RE((A_ig^{2w_i}, B_ih^{2w_i}), \Gamma_2[i])$ occurs in $T^{new}[l]$ exactly once.

9.2 Analysis

The following results hold.

Theorem 8: A correct execution of ZK-PRS has $O(\log N \log \log N)$ amortized complexity.

Proof: The amortized cost of the $Shuffle$ operation has four components, one for each main step: $T_{Sh} = s(T_{RE} + T_{LR} + T_{Dec} + T_{PV})$. The list reshuffle step has an amortized cost $T_{LR} = O(\log N \log \log N)$ (see [34] for details). Each of the other three steps performs a constant number of operations per element of the level to be reshuffled. For instance, the re-encryption step (including the fake insertion step) performs one re-encryption and one symmetric key encryption per moved element and one re-encryption, one modular multiplication and one symmetric key encryptions per

Algorithm 5 Operation that removes all W-ORAM elements that expire at time T.

```
1.Expire(E − ORAM,W − ORAM : ORAM,T : int
2. L : id[]; #expiring labels
3. E : string[]; #removed from W − ORAM
4. L := Enumerate(E − ORAM,T);
5. for each label in L do
6.    (R,E) := Read_ORAM(W − ORAM,label);
7.    Proof := ZK − PEE(R,E);
8.    if (verify(Proof,E) = 0) then
9.       undo(W − ORAM);
10.      return error;
11.   fi od
12.end
```

added fake element. The decryption step performs two symmetric key encryptions per level element. The proof validation step performs at most one symmetric key encryption and one re-encryption per element. Thus, for level l, containing $O(4^l)$ elements, the sum of the T_{SG}, T_{Dec} and T_{PV} components is $O(4^l)$. However, the $Shuffle$ operation for the lth level is called once every 4^{l-1} ORAM operations. Thus, the amortized cost for these three steps is

$$\sum_{l=1}^{\log N} \frac{O(4^l)}{4^{l-1}} = O(\log N)$$

Since s is a constant, we conclude that the amortized cost of the reshuffle is $O(\log N \log \log N)$. $\square$

Theorem 9: ZK-PRS is complete.

Proof: It is straightforward to see that if $T^{new}[l]$ is a permutation of $T[l]$, the client can build proof sets P such that later it is able to build the answer to the server's challenges. If $b = 0$ the client proves that $T[l] = \pi(P\bar{\Gamma_1})$ and the remaining elements of P are fakes. If $b = 1$, the client proves that $T^{new}[l] = \pi(P\bar{\Gamma_2})$. $\bar{\Gamma_1}$ and $\bar{\Gamma_2}$ are the vectors containing all $\Gamma_1[i]$ and $\Gamma_2[i]$ values. Thus, the server is convinced that $T^{new}[l]$ is a permutation of $T[l]$ and has $|T^{new}[l]| - |T[l]|$ new fakes. $\square$

Theorem 10: ZK-PRS is sound.

Proof: If $T^{new}[l]$ is not a permutation of $T[l]$, a cheating client cannot convince the server of the opposite, except with a small probability. The client can cheat by building P to be either (i) a permutation of $T[l]$ or (ii) a permutation of $T^{new}[l]$. We only study case (i), since case (ii) is similar. Let an element from P, $P_i = (A_i g^{2w_i}, B_i h^{2w_i})$ be a re-encryption of the ith element of $T[l]$ and let $T^{new}_i = (A'_i g^{2u'_i}, B'_i h^{2u''_i})$ be the corresponding element (not a re-encryption) from $T^{new}[l]$. If $b = 1$, the client has to provide a value Γ_2 such that $T^{new}_i = RE(P_i, \Gamma_2)$. If the client is able to provide such a value, it is straightforward to see that $T^{new}_i = RE((A_i, B_i), (\Gamma_2 + w_i))$. $\square$

Theorem 11: ZK-PRS is zero knowledge.

Proof: The proof follows the same pattern as the proof of Theorem 6 and [36] and we omit it due to its redundancy. $\square$

10 ELEMENT EXPIRATION

Expire(T): Use as input a time epoch T and remove all the elements from the W-ORAM that expire in that epoch (see Algorithm 5). Use the E-ORAM to enumerate all the labels that expire at T (line 4). For each such *label* (line 5) read (and remove) from the W-ORAM the corresponding element (line 6). Note that the $Read_{ORAM}$ operation also returns the entire list E of elements removed from the W-ORAM – containing $\log N$ elements. Then, build a zero knowledge proof of correctness, ZK-PEE (line 7). ZK-PEE proves that E contains one real element that expires at T and the rest ($\log N - 1$ elements) are fakes. If the proof verifies, the server accepts the expiration, otherwise restores the W-ORAM to the state before the Read of line 6 (line 9) and returns error (line 10). We now describe ZK-PEE.

10.0.1 Zero Knowledge Proof of Element Expiration.

ZK-PEE takes as input the element to be expired, R and the list E of all elements that were removed from W-ORAM when R was read (line 6). Note that $R \in E$. Let m be the number of elements in E and let $E = \{(s_1^2, k_1), .., R, .., (s_m^2, k_m)\}$. Let r be R's index in E. ZK-PEE consists of s rounds. During each round the following steps are executed by the client and the server.

Step 1: Proof Generation: The client generates a random permutation $\pi \in_R \mathcal{P}_m$ and a random vector $\bar{w} = \{w_1, .., w_m\}$, where $w_i \in_R \{0,1\}^{\mathbf{k}}$ are odd. The client computes the proof list $P = \pi(E\bar{w})$ and sends it to the server.

Step 2: Proof Verification: The server flips a bit b. If $b = 0$, the client reveals $\bar{w}$. The server verifies that all $w_i \in \bar{w}$ are odd and that $P = \pi(E\bar{w})$. If $b = 1$, the client reveals $Dec(R, d, k) = M(x) = (E_k(x), T_{exp}, "real", rnd)$ to the server along with the encryption factor uw_r and the square roots of the remaining m-1 (fake) elements in P, $s_1 g^{w_1}, .., s_m g^{w_m}$. The server verifies the revealed element: (i) its format, that is, $T_{exp} = T$ and the third field is "real" and (ii) its correctness, $(M(x)g^{2uw_r}, h^{2uw_r}) \in P$. The server also verifies that the remaining m-1 elements in P are fakes, by checking that $(s_i g^{w_i})^2$ occurs in the first field of exactly one element in P.

10.0.2 Analysis

Let e be the number of elements that expire at the same time. Then, the following result holds.

Theorem 12: A correct execution of Expire has $O(e \log N)$ complexity.

Proof: The overall client overhead for a single *Read* operation is $T^c_{Exp} \approx$
$2eT_{dsk} \log N + eT_{RE} s \log N + es/2T_e \log N$, where T_{dsk} is the time to access the disk, T_{RE} is the re-encryption time (two modular exponentiations) and T_e is the modular exponentiation time. The first factor corresponds to the time to discover the expiring labels (in the E-ORAM) and to read them from the W-ORAM. The second factor is from the proof generation step of ZK-PEE and the

last factor is from the proof verification step. The server only participates during the proof verification steps and its cost is $T^s_{Exp} \approx es/2 T_{RE} \log N + es/2(T_{mul} \log N + T_{RE})$ where $s/2$ denotes the number of bits b that come up 0 and 1. Thus, the total (client+server) overhead is $O(e \log N)$. □

We now prove the following result.

Theorem 13: ZK-PEE is a zero knowledge proof system of Expire $\in$ WORM. That is, Expire is WORM compliant.

Proof: First, note that ZK-PEE is complete: if the client is honest, it is able to build the answer to the server's challenges. If $b = 0$, the client reveals all values $w_1, .., w_m$ used to transform E into P. If $b = 1$, the client proves that $m-1$ of the values in P are fakes, that is, their first field is a quadratic residue. then, for the remaining value, the client provides the encryption factor and the decrypted, expiring element from E.

To see that ZK-PEE is sound, note that the client can only cheat by (i) removing more than 1 real element from the W-ORAM or by (ii) removing an element that does not expire. Then the client can cheat by building P to either be a permutation of E to contain m-1 fakes and the encryption of a real, expiring element. However, we follow the same reasoning used for ZK-POR and ZK-PRS. If the client can then convince the server of the validity of P for either challenge, we can show that the list E must contain m-1 fakes and one real, expiring element, which contradicts the hypothesis.

Finally, to see that ZK-PEE is zero knowledge, we follow the same approach used for ZK-POR. Due to the length and redundancy of the proof, we omit it. □

11 AUDIT

Audit(d,k).: Take as input the decryption keys d and k to search for desired elements in W-ORAM. Call $Dec((A, B), d, k)$, for all elements (A, B) in the W-ORAM. Once all the elements are recovered, they can be searched for desired keywords.

12 KEY MANAGEMENT

For the sake of presentation clarity, we have presented a simplified element encoding operation (Enc). Specifically, an element x is stored as the pair $(M(x)g^{2b}, h^{2b})$, consisting of an encrypted part and a recovery key. However, during element expiration (see Section 10) the client needs to prove to the server (in zero knowledge) the fact that one element in the list of accessed elements expires. For this, the client needs to provide the server not only with the decrypted element but also with the obfuscating exponent (b in the above example). Since an element may have been accessed and re-encrypted many times during read and reshuffle operations, the client needs to keep track of the changes in the obfuscation exponent.

We address this by storing a third field for any element: the encrypted exponent, e.g., $E(b)$ in the above case, where E is any semantically secure symmetric

Resource	Spec
Processor	2.4GHz
RAM	256KB
Disk bandwidth	120MB/s
Link bandwidth	10MB/s
Block size	1024b
T_{RE}	125 ops/s
T_{sym} on 1024b	272355 ops/s

TABLE 1
Client and server configurations.

key encryption method, whose key is private to the client. Whenever the element is re-encrypted (during read and re-shuffle operations), the new exponent is stored encrypted, replacing the existing one. The use of a semantically secure encryption method prevents the server from using this third field to correlate reshuffled elements. For fake elements the third field is random and changes whenever a fake is being "re-encrypted".

13 EXPERIMENTAL EVALUATION

We have implemented our solution using OpenSSL and we have tested it on the configuration depicted in Table 1. We used the same PC configuration (single core 2.4GHz with 256MB of RAM and 120MB/s sustained read/write rates) for both server and client platforms. As such, the server and client can perform 250 modular exponentiations per second, leading to 125 record re-encryptions per second and 272K AES encryptions on 1024 blocks. The link between client and server was a duplex 10MB/s. The outsourced dataset consists of 1024 bit records.

In the following, we look at the overheads of read and shuffle as they are the most expensive operations. The element expiration operation follows the same steps as a read and thus its cost is similar: the client needs to compute proof sets of size $\log N$ and the server either verifies their re-encryption or the quadratic residuosity of $\log N - 1$ of their elements

13.1 Read Overheads

We now focus on the overheads of the zero knowledge proofs used during a read operation. Figure 1 shows the overhead of the ZK-POR process as a function of the number of records, N. The number of proof sets employed is 40, for a client cheating probability of 2^{-40}. The x-axis shows the number of records in logarithmic scale. We have experimented with datasets ranging from 1Mb to 1Tb. Figure 1 shows the cost of a read operation, in terms of client and server computation costs and total transfer costs. For a 1Tb dataset (2^{30} records of size 1024), the client cost is under 7s and the server cost is under 4s. The transfer cost of the 41 sets of records including also the disk read/write times is only a fraction of a second. Thus, the total overhead of a read is around 10s.

Fig. 1. ZK-POR client and server overheads. Read overhead (shown in seconds) as a function of the dataset size, on the dataset size.


Fig. 2. Client, server and communication components of the amortized cost of ZK-PRS. Shuffle overhead (shown in seconds) as a function of the dataset size, $\log_2 N$.

The server overhead is roughly half the overhead of the client. This is because the server has to verify re-encryptions only on half the sets generated by the client. The other half is dominated by modular multiplications, which are three orders of magnitude faster than re-encryptions.

13.2 Shuffle Overheads

A level l re-shuffle needs to be performed every time level l becomes full (every 4^l accesses). We have measured the amortized impact of shuffles on the operation of the WORM ORAM structure and Figure 2 shows our results. The amortized cost includes the cost of re-shuffle incurred at *all* the levels in the ORAM, over all the ORAM accesses. Figure 2 shows the dependence of the cost of ZK-PRS on the number of records stored at the server, shown on the x-axis in logarithmic scale. The cost is again divided into client, server and communication components. The number of proof sets is set to 40, for a client cheating probability of 2^{-40}.

Similar to the read overheads, the client and server computation components of ZK-PRS show a logarithmic dependence on the size of the dataset. While our

theoretical result of Section 9.2 shows this dependence to be $O(\log N \log \log N)$, the costs are dominated by the generation and verification of re-encrypted permuted sets. The overheads of the actual oblivious set re-shuffles, including symmetric key operations, are quite small: more than 272K symmetric key operations can be performed per second, but only around 125 re-encryptions per second are possible.

The client and server overheads are similar, up to 20 seconds for 2^{40} datasets. The reason for this similarity is that the server has to verify set re-encryptions irrespective of the outcome of the coin flip process, leading to similar numbers of server verifications and client generations.

The communication overhead during the shuffle operation, including the time to read/write and transfer proof sets takes around 7 seconds for 1Tb datasets. This is because the client needs to shuffle not one but 41 proof sets. Thus, the total, amortized overhead of a shuffle operation is around 47s.

In ORAM, the network transfer time alone for reshuffling level i consists of about 10 sorts of $4^i logn$ data, each sort requiring $4^i log(n)log^2(4i^i logn)$ block transfers, for a total of $104^i log(n)log^2(4^i logn)2^{10}/10MB/se$. Summing over the $\log_4 n$ levels, and amortizing each level over 4^{i1} queries, ORAM has an amortized network traffic cost per query of 3.680Gb. Over the sample 10MB/s link this is a 48 sec/query amortized transfer time. Thus, by using an improved oblivious scramble protocol, we are able to support regulatory compliance *and* maintain the cost imposed by the original ORAM.

Amortizing the Shuffle Cost: In order to avoid incurring shuffle costs at once and to amortize them over client operations, we use the following techniques. The bulk of the client cost for a shuffle is generating re-encryptions, consisting of 2 expensive modular exponentiations. During each client access to the ORAM (write, read, expiration), the client generates $4(s+1)\log N$ random factors w, generates g^w and h^w and stores the encrypted tuple (w, g^w, h^w) at the server. The number of such tuples, $4(s+1)\log N$, where s is the number of proof sets and N is the dataset size, is the amortized number of re-encryptions. During an actual re-shuffle, the client reads such tuples from the server, decrypts them and uses them for re-encryptions. The remaining overhead, of modular multiplications, symmetric key cryptographic operations and oblivious set-reshuffles, is small, as shown also in the results of Figure 2.

The bulk of the server ZK-PRS cost consists in verifying set re-encryptions. The server batches shuffle operations. For each batched shuffle, the server stores partial ORAM state (for rollback). Following each client operation (read,write,expiration), the server performs an additional $4s\log N$ element re-encryption verifications from the batch queue. If any verification fails, the server rolls back to the previous state of the ORAM, associated to the currently verified shuffle. This approach introduces a computation to storage tradeoff: the server needs

to store $O(N)$ more data, including partial re-encryptions and rollback state.

14 CONCLUSIONS

In this paper we introduce WORM-ORAM, a solution that provides WORM compliant Oblivious RAMs. Our solution is based on a set of zero knowledge proofs that ensure that all ORAM operations are WORM compliant. The protocol features the same asymptotic computational complexity as ORAM.

15 ACKNOWLEDGMENTS

REFERENCES

[1] National Association of Insurance Commissioners. Graham-Leach-Bliley Act, 1999. www.naic.org/GLBA.

[2] U.S. Dept. of Health & Human Services. The Health Insurance Portability and Accountability Act (HIPAA), 1996. www.cms.gov/hipaa.

[3] U.S. Public Law 107-347. The E-Government Act, 2002.

[4] U.S. Public Law No. 107-204, 116 Stat. 745. Public Company Accounting Reform and Investor Protection Act, 2002.

[5] The U.S. Securities and Exchange Commission. Rule 17a-3&4, 17 CFR Part 240: Electronic Storage of Broker-Dealer Records. Online at http://edocket.access.gpo.gov/, 2003.

[6] The U.S. Department of Defense. Directive 5015.2: DOD Records Management Program. Online at http://www.dtic.mil/whs/directives/corres/pdf/50152std_061902/p50152s.pdf, 2002.

[7] The U.S. Department of Health and Human Services Food and Drug Administration. 21 CFR Part 11: Electronic Records and Signature Regulations. Online at http://www.fda.gov/ora/compliance_ref/part11/FRs/background/pt11finr.pdf, 1997.

[8] The U.S. Department of Education. 20 U.S.C. 1232g; 34 CFR Part 99:Family Educational Rights and Privacy Act (FERPA). Online at http://www.ed.gov/policy/gen/guid/fpco/ferpa, 1974.

[9] The Enterprise Storage Group. Compliance: The effect on information management and the storage industry. Online at http://www.enterprisestoragegroup.com/, 2003.

[10] Enron email dataset. http://www.cs.cmu.edu/enron/.

[11] IBM Corp. IBM TotalStorage Enterprise. Online at http://www-03.ibm.com/servers/storage/, 2007.

[12] HP. WORM Data Protection Solutions. Online at http://h18006.www1.hp.com/products/storageworks/wormdps/index.html, 2007.

[13] EMC. Centera Compliance Edition Plus. Online at http://www.emc.com/centera/ and http://www.mosaictech.com/pdf_docs/emc/centera.pdf, 2007.

[14] Hitachi Data Systems. The Message Archive for Compliance Solution, Data Retention Software Utility. Online at http://www.hds.com/solutions/data_life_cycle_archiving/achievingregcompliance.html, 2007.

[15] Zantaz Inc. The ZANTAZ Digital Safe Product Family. Online at http://www.zantaz.com/, 2007.

[16] StorageTek Inc. VolSafe secure tape-based write once read many (WORM) storage solution. Online at http://www.storagetek.com/, 2007.

[17] Sun Microsystems. Sun StorageTek Compliance Archiving system and the Vignette Enterprise Content Management Suite (White Paper). Online at http://www.sun.com/storagetek/white-papers/Healthcare_Sun_NAS_Vignette_EHR_080806_Final.pdf, 2007.

[18] Sun Microsystems. Sun StorageTek Compliance Archiving Software. Online at http://www.sun.com/storagetek/management_software/data_protection/compliance_archiving/, 2007.

[19] Network Appliance Inc. SnapLock Compliance and SnapLock Enterprise Software. Online at http://www.netapp.com/products/software/snaplock.html, 2007.

[20] Quantum Inc. DLTSage Write Once Read Many Solution. Online at http://www.quantum.com/Products/TapeDrives/DLT/SDLT600/DLTIce/Index.aspx and http://www.quantum.com/pdf/DS00232.pdf, 2007.

[21] Oded Goldreich and Rafail Ostrovsky. Software protection and simulation on oblivious ram. *Journal of the ACM*, 45:431–473, May 1996.

[22] Peter Williams, Radu Sion, and Bogdan Carbunar. Building Castles out of Mud: Practical Access Pattern Privacy and Correctness on Untrusted Storage. In *ACM Conference on Computer and Communication Security CCS*, 2008.

[23] Bogdan Carbunar and Radu Sion. Regulatory compliant oblivious ram. In *ACNS*, pages 456–474, 2010.

[24] Benny Pinkas and Tzachy Reinman. Oblivious ram revisited. In *Advances in Cryptology - CRYPTO 2010*, pages 502–519, 2010.

[25] Rasmus Pagh and Flemming Friche Rodler. Cuckoo hashing. *J. Algorithms*, 51:122–144, May 2004.

[26] B. Chor, O. Goldreich, E. Kushilevitz, and M. Sudan. Private information retrieval. In *IEEE Symposium on Foundations of Computer Science*, pages 41–50, 1995.

[27] W. Gasarch. A WebPage on Private Information Retrieval. Online at http://www.cs.umd.edu/~gasarch/pir/pir.html.

[28] W. Gasarch. A survey on private information retrieval. Online at http://citeseer.ifi.unizh.ch/gasarch04survey.html.

[29] Radu Sion and Bogdan Carbunar. On the Computational Practicality of Private Information Retrieval. In *Proceedings of the Network and Distributed Systems Security Symposium*, 2007. Stony Brook Network Security and Applied Cryptography Lab Tech Report 2006-06.

[30] Jan Camenisch, Gregory Neven, and Abhi Shelat. Simulatable adaptive oblivious transfer. In *EUROCRYPT '07: Proceedings of the 26th annual international conference on Advances in Cryptology*, 2007.

[31] S. Coull, M. Green, and S. Hohenberger. Controlling access to an oblivious database using stateful anonymous credentials. In *International Conference on Practice and Theory in Public Key Cryptography (PKC)*, 2009.

[32] O. Goldreich. *Foundations of Cryptography I*. Cambridge University Press, 2001.

[33] S. Goldwasser, S. Micali, and C. Rackoff. The knowledge complexity of interactive proof systems. *SIAM J. Comput.*, 18(1), 1989.

[34] Peter Williams, Radu Sion, and Bogdan Carbunar. Building castles out of mud: practical access pattern privacy and correctness on untrusted storage. In *CCS '08: Proceedings of the 15th ACM conference on Computer and communications security*, 2008.

[35] Dan Boneh. The decision diffie-hellman problem. In *ANTS-III: Proceedings of the Third International Symposium on Algorithmic Number Theory*, 1998.

[36] Oded Goldreich, Silvio Micali, and Avi Wigderson. Proofs that yield nothing but their validity or all languages in np have zero-knowledge proof systems. *J. ACM*, 38(3), 1991.