GORAM – Group ORAM for Privacy and Access Control in Outsourced Personal Records[1]

Matteo Maffei   Giulio Malavolta   Manuel Reinert   Dominique Schröder

Saarland University, CISPA
{maffei,malavolta,reinert,schroeder}@cs.uni-saarland.de

August 25, 2015

---

[1]An extended abstract of this paper has been published at IEEE S&P 2015 [1].

Abstract

Cloud storage has rapidly become a cornerstone of many IT infrastructures, constituting a seamless solution for the backup, synchronization, and sharing of large amounts of data. Putting user data in the direct control of cloud service providers, however, raises security and privacy concerns related to the integrity of outsourced data, the accidental or intentional leakage of sensitive information, the profiling of user activities and so on. Furthermore, even if the cloud provider is trusted, users having access to outsourced files might be malicious and misbehave. These concerns are particularly serious in sensitive applications like personal health records and credit score systems.

To tackle this problem, we present GORAM, a cryptographic system that protects the secrecy and integrity of outsourced data with respect to both an untrusted server and malicious clients, guarantees the anonymity and unlinkability of accesses to such data, and allows the data owner to share outsourced data with other clients, selectively granting them read and write permissions. GORAM is the first system to achieve such a wide range of security and privacy properties for outsourced storage. In the process of designing an efficient construction, we developed two new, generally applicable cryptographic schemes, namely, batched zero-knowledge proofs of shuffle and an accountability technique based on chameleon signatures, which we consider of independent interest. We implemented GORAM in Amazon Elastic Compute Cloud (EC2) and ran a performance evaluation demonstrating the scalability and efficiency of our construction.

Contents

1 Introduction

Cloud storage has rapidly gained a central role in the digital society, serving as a building block of consumer-oriented applications (e.g, Dropbox, Microsoft SkyDrive, and Google Drive) as well as particularly sensitive IT infrastructures, such as personal record management systems. For instance, credit score systems rely on credit bureaus (e.g., Experian, Equifax, and TransUnion in US) collecting and storing information about the financial status of users, which is then made available upon request. As a further example, personal health records (PHRs) are more and more managed and accessed through web services (e.g., private products like Microsoft HealthVault and PatientsLikeMe in US and national services like ELGA in Austria), since this makes PHRs readily accessible in case of emergency even without the physical presence of the e-health card and eases their synchronization across different hospitals.

Despite its convenience and popularity, cloud storage poses a number of security and privacy issues. The first problem is related to the *secrecy* of user data, which are often sensitive (e.g., PHRs give a complete picture of the health status of citizens) and, thus, should be concealed from the server. A crucial point to stress is that preventing the server from reading user data (e.g., through encryption) is necessary but not sufficient to protect the privacy of user data. Indeed, as shown in the literature [2, 3], the capability to link consecutive accesses to the same file can be exploited by the server to learn sensitive information: for instance, it has been shown that the access patterns to a DNA sequence allow for determining the patient's disease. Hence the *obliviousness* of data accesses is another fundamental property for sensitive IT infrastructures: the server should not be able to tell whether two consecutive accesses concern the same data or not, nor to determine the nature of such accesses (read or write). Furthermore, the server has in principle the possibility to modify client's data, which can be harmful for several reasons: for instance, it could drop data to save storage space or modify data to influence the statistics about the dataset (e.g., in order to justify higher insurance fees or taxes). Therefore another property that should be guaranteed is the *integrity* of user data.

Finally, it is often necessary to share outsourced documents with other clients, yet in a controlled manner, i.e., selectively granting them read and write permissions: for instance, PHRs are selectively shared with the doctor before a medical treatment and a prescription is shared with the pharmacy in order to buy a medicine. *Data sharing* complicates the enforcement of secrecy and integrity properties, which have to be guaranteed not only against a malicious server but also against malicious clients. Notice that the simultaneous enforcement of these properties is particularly challenging, since some of them are in seeming contradiction. For instance, *access control* seems to be incompatible with the obliviousness property: if the server is not supposed to learn which file the client is accessing, how can he check that the client has the rights to do so?

1.1 Our Contributions

In this work, we present GORAM, a novel framework for privacy-preserving cloud-storage. Users can share outsourced data with other clients, selectively granting them read and write permissions, and verify the integrity of such data. These are hidden from the server and access patterns are oblivious. GORAM is the first system to achieve such a wide range of security and privacy properties for storage outsourcing. More specifically, the contributions of this work are the following:

- We formalize the problem statement by introducing the notion of Group Oblivious RAM (GORAM). GORAM extends the concept of Oblivious RAM [4] (ORAM) [1] by considering multiple, possibly malicious clients, with read and/or write access to outsourced data, as opposed to a single client. We propose a formal security model that covers a variety of security and privacy properties, such as data integrity, data secrecy, obliviousness of access patterns, and anonymity.

- We first introduce a cryptographic instantiation based on a novel combination of ORAM [5], predicate encryption [6], and zero-knowledge (ZK) proofs (of shuffle) [7, 8]. This construction is secure, but building on off-the-shelf cryptographic primitives is not practical. In particular, clients prove to the server that the operations performed on the database are correct through ZK proofs of shuffle, which are expensive when the entries to be shuffled are tuples of data, as opposed to single entries.

- As a first step towards a practical instantiation, we maintain the general design, but we replace the expensive ZK proofs of shuffle with a new proof technique called *batched* ZK proofs of shuffle. A batched ZK proof of shuffle significantly reduces the number of ZK proofs by "batching" several instances and verifying them together. Since this technique is generically applicable in any setting where one is interested to perform a zero-knowledge proof of shuffle over a list of entries, each of them consisting of a tuple of encrypted blocks, we believe that it is of independent interest. This second realization greatly outperforms the first solution and is suitable for databases with relatively small entries, accessed by a few users, but it does not scale to large entries and many users.

- To obtain a scalable solution, we explore some trade-offs between security and efficiency. First, we present a new accountability technique based on chameleon signatures. The idea is to let clients perform arbitrary operations on the database, letting them verify each other's operation a-posteriori and giving them the possibility to blame misbehaving parties. Secondly, we replace the relatively expensive predicate encryption, which enables sophisticated role-based and attribute-based access control policies, with the more efficient broadcast encryption, which suffices to enforce per-user read/write permissions, as required in the personal record management systems we consider. This approach leads to a very efficient solution that scales to large files and thousands of users, with a combined communication-computation overhead of only 7% (resp. 8%) with respect to state-of-the-art, single-client ORAM constructions for reading (resp. writing) on a 1GB storage with 1MB block size (for larger datasets or block sizes, the overhead is even lower).

We have implemented GORAM in Amazon Elastic Compute Cloud (EC2) and conducted a performance evaluation demonstrating the scalability and efficiency of our construction. Although GORAM is generically applicable, the large spectrum of security and privacy properties, as well as the efficiency and scalability of the system, make GORAM particularly suitable for the management of large amounts of sensitive data, such as personal records.

1.2 Outline

Section 2 introduces the notion of Group ORAM. We discuss the general cryptographic instantiation in Section 3, the accountability-based construction in Section 4, and the efficient scheme with broadcast encryption in Section 5. Section 6 formalizes the security properties of Group ORAM and Section 7 states the security and privacy results. We implemented our system and

---

[1] ORAM is a technique originally devised to protect the access pattern of software on the local memory and then used to hide the data and the user's access pattern in storage outsourcing services.

conducted an experimental evaluation, as discussed in Section 8. Section 9 presents a case study on PHRs. The related work is discussed in Section 10. Section 11 concludes and outlines future research directions.

2 System Settings

We detail the problem statement by formalizing the concept of Group ORAM (Section 2.1), presenting the relevant security and privacy properties (Section 2.2), and introducing the attacker model (Section 2.3).

2.1 Group ORAM

We consider a data owner $\mathcal{O}$ outsourcing her database $\mathsf{DB} = d_1, \ldots, d_m$ to the server $\mathcal{S}$. A set of clients $\mathcal{C}_1, \ldots, \mathcal{C}_n$ can accesses parts of the database, as specified by the access control policy set by $\mathcal{O}$. This is formalized as an n-by-m matrix $\mathbf{AC}$, defining the permissions of the clients on the files in the database: $\mathbf{AC}(i, j)$ (i.e., the j-th entry of the i-th row) denotes the access mode for client i on data d_j. Each entry in the matrix is an element of the set $\{\bot, r, rw\}$ of access modes, denoting no access, read access, and write access, respectively.

At registration time, each client $\mathcal{C}_i$ receives a capability cap_i, which gives $\mathcal{C}_i$ access to DB as specified in the corresponding row of $\mathbf{AC}$. Furthermore, we assume the existence of a capability $cap_{\mathcal{O}}$, which grants permissions for all of the operations that can be executed by the data owner only.

In the following we formally characterize the notion of Group ORAM. Intuitively, a Group ORAM is a collection of two algorithms and four interactive protocols, used to setup the database, add clients, add an entry to the database, change the access permissions to an entry, read an entry, and overwrite an entry. In the sequel, we let $\langle A, B \rangle$ denote a protocol between the PPT machines A and B, $|\mathbf{a}|$ the length of the vector $\mathbf{a}$ of access modes, and $\mathbf{a}(i)$ the element at position i in $\mathbf{a}$. In all our protocols $|\mathsf{DB}|$ is equal to the number of columns of $\mathbf{AC}$.

Definition 1 (Group ORAM). *A Group ORAM scheme is a tuple of (interactive) PPT algorithms* $\mathsf{GORAM} = (\mathsf{gen}, \mathsf{addCl}, \mathsf{addE}, \mathsf{chMode}, \mathsf{read}, \mathsf{write})$, *such that:*

$(cap_{\mathcal{O}}, \mathsf{DB}) \leftarrow \mathsf{gen}(1^{\lambda}, n)$: *The* gen *algorithm initializes the database* $\mathsf{DB} := [\ ]$ *and the access control matrix* $\mathbf{AC} := [\ ]$, *and generates the access capability* $cap_{\mathcal{O}}$ *of the data owner. The parameter n determines the maximum number of clients. This algorithm returns* $(cap_{\mathcal{O}}, \mathsf{DB})$, *while* $\mathbf{AC}$ *is a global variable that maintains a state across the subsequent algorithm and protocol executions.*

$\{cap_i, \mathsf{deny}\} \leftarrow \mathsf{addCl}(cap_{\mathcal{O}}, \mathbf{a})$: *The* addCl *algorithm is run by the data owner, who possesses* $cap_{\mathcal{O}}$, *to register a new client, giving her access to the database as specified by the vector* $\mathbf{a}$. *If $|\mathbf{a}|$ is equal to the number of columns of* $\mathbf{AC}$, $\mathbf{a}$ *is appended to* $\mathbf{AC}$ *as the last row and the algorithm outputs a fresh capability cap_i that is assigned to that row. Otherwise, it outputs* deny.

$\{\mathsf{DB}', \mathsf{deny}\} \leftarrow \langle \mathcal{C}_{\mathsf{addE}}(cap_{\mathcal{O}}, \mathbf{a}, d), \mathcal{S}_{\mathsf{addE}}(\mathsf{DB}) \rangle$: *This protocol is run by the data owner, who possesses* $cap_{\mathcal{O}}$, *to append an element d to* DB, *assigning the vector $\mathbf{a}$ of access modes. If $|\mathbf{a}|$ is equal to the number of rows of* $\mathbf{AC}$ *then d is appended to* DB, $\mathbf{a}$ *is appended to* $\mathbf{AC}$ *as the last column, and the protocol outputs the new database* DB'; *otherwise it outputs* deny.

$\langle \mathcal{C}_{\mathsf{chMode}}(cap_{\mathcal{O}}, \mathbf{a}, j), \mathcal{S}_{\mathsf{chMode}}(\mathsf{DB}) \rangle$: *This protocol is used by the data owner, who possesses* $cap_{\mathcal{O}}$, *to change the access permissions for the j-th entry as specified by the vector $\mathbf{a}$ of access modes. If $j \leq |\mathsf{DB}|$ and $|\mathbf{a}|$ is equal to the number of rows of* $\mathbf{AC}$, *then the j-th column of* $\mathbf{AC}$ *is replaced by $\mathbf{a}$.*

5

$\{d, \mathsf{deny}\} \leftarrow \langle \mathcal{C}_{\mathsf{read}}(cap_i, j), \mathcal{S}_{\mathsf{read}}(\mathsf{DB}) \rangle :$ *The interactive* read *protocol is used by the owner of* cap_i *to read the* j-*th entry of* DB. *This protocol returns either* $d := \mathsf{DB}(j)$ *or* deny *if* $|\mathsf{DB}| < j$ *or* $\mathbf{AC}(i, j) = \bot$.

$\{\mathsf{DB}', \mathsf{deny}\} \leftarrow \langle \mathcal{C}_{\mathsf{write}}(cap_i, j, d), \mathcal{S}_{\mathsf{write}}(\mathsf{DB}) \rangle :$ *The interactive* write *protocol is used by the owner of the capability* cap_i *to overwrite the* j-*th entry of* DB *with* d. *This protocol succeeds and outputs* DB' *if and only if* $\mathbf{AC}(i, j) = rw$, *otherwise it outputs* deny.

2.2 Security and Privacy Properties

Here we briefly outline the fundamental security and privacy properties achieved by a Group ORAM. We refer to Section 6 for a precise formalization of these properties based on cryptographic games.

Secrecy: clients can only read entries they hold *read* permissions on.

Integrity: clients can only write entries they hold *write* permissions on.

Tamper-resistance: clients, eventually colluding with the server, cannot modify an entry they do not hold *write* permission on without being detected by the data owner.

Obliviousness: the server cannot determine the access pattern on the data given a clients' sequence of operation.

Anonymity: the server and the data owner cannot determine who performed a given operation, among the set of clients that are allowed to perform it.

Accountable Integrity: clients cannot write entries they do not hold *write* permission on without being detected.

2.3 The Attacker Model

We consider an adversarial model in which the data owner $\mathcal{O}$ is honest, the clients $\mathcal{C}_1, \ldots, \mathcal{C}_n$ may be malicious, and the server $\mathcal{S}$ is assumed to be honest-but-curious (HbC)[2] and not to collude with clients. These assumptions are common in the literature (see, e.g., [9, 10]) and are well justified in a cloud setting, since it is of paramount importance for service providers to keep a good reputation, which discourages them from visibly misbehaving, while they may have an incentive in passively gathering sensitive information given the commercial interest of personal data.

Although we could limit ourselves to reason about all security and privacy properties in this attacker model, we find it interesting to state and prove some of them even in a stronger attacker model, where the server can arbitrarily misbehave. This allows us to characterize which properties unconditionally hold true in our system, i.e., even if the server gets compromised (cf. the discussion in Section 6).

3 Our Construction (GORAM)

In this section, we first show how to realize a Group ORAM using a novel combination of ORAM, predicate encryption, and zero-knowledge proofs (Section 3.1 and Section 3.2). Since even the usage of the most efficient zero-knowledge proof system still yields an inefficient construction, we

---

[2]I.e., the server is regarded as a passive adversary, following the protocol but seeking to gather additional information

introduce a new proof technique called batched ZK proofs of shuffle (Section 3.3) and instantiate our general framework with this primitive.

3.1 Prerequisites

In the following, we describe the database layout, the basic cryptographic primitives, and the system assumptions.

Layout of the database. The layout of the database DB follows the one proposed by Stefanov *et al.* [5]. To store N data entries, we use a binary tree T of depth $D = O(\log N)$, where each node stores a bucket of entries, say b entries per bucket. We denote a node at depth d and row index i by $T_{d,i}$. The depth at the root ρ is 0 and increases from top to bottom; the row index increases from left to right, starting at 0. We often refer to the root of the tree as ρ instead of $T_{0,0}$. Moreover, Path-ORAM [5] uses a so-called *stash* as local storage to save entries that would overflow the root bucket. We assume the stash to be stored and shared on the server like every other node, but we leave it out for the algorithmic description. The stash can also be incorporated in the root node, which does not carry b but $b + s$ entries where s is the size of the stash. The extension of the algorithms is straight-forward (only the number of downloaded entries changes) and does not affect their computational complexity. In addition to the database, there is an index structure $\mathcal{LM}$ that maps entry indices i to leaf indices l_i. If an entry index i is mapped in $\mathcal{LM}$ to l_i then the entry with index i can be found in some node on the path from the leaf l_i to the root ρ of the tree. Finally, to initialize the database we fill it with dummy elements.

Cryptographic preliminaries. We informally review the cryptographic building blocks and introduce a few useful notations. (For formal definitions, we refer to Appendix A.)

We denote by $\Pi_{\mathsf{SE}} = (\mathsf{Gen}_{\mathsf{SE}}, \mathcal{E}, \mathcal{D})$ a private-key encryption scheme, where $\mathsf{Gen}_{\mathsf{SE}}$ is the key-generation algorithm and $\mathcal{E}$ (resp. $\mathcal{D}$) is the encryption (resp. decryption) algorithm. Analogously, we denote by $\Pi_{\mathsf{PKE}} = (\mathsf{Gen}_{\mathsf{PKE}}, \mathsf{Enc}, \mathsf{Dec})$ a public-key encryption scheme. We also require a publicly available function Rerand to rerandomize public-key ciphertexts. We require that both encryption schemes fulfill the IND-CPA-security property [11].

A predicate encryption scheme [6] $\Pi_{\mathsf{PE}} = (\mathsf{PrGen}, \mathsf{PrKGen}, \mathsf{PrEnc}, \mathsf{PrDec})$ consists of a setup algorithm PrGen, a key-generation algorithm PrKGen, and an encryption (resp. decryption) algorithm PrEnc (resp. PrDec). In a predicate encryption scheme, one can encrypt a message m under an attribute x. The resulting ciphertext can only be decrypted with a secret key that encodes a predicate f such that $f(x) = 1$. The choice of predicates determines who can decrypt which ciphertext, which makes predicate encryption a flexible cryptographic tool to enforce access control policies. We further use a predicate-only encryption scheme $\Pi_{\mathsf{PO}} = (\mathsf{PoGen}, \mathsf{PoKGen}, \mathsf{PoEnc}, \mathsf{PoDec})$. The difference from Π_{PE} is that the attribute x is encrypted only. As for public-key encryption, we require rerandomization functions PrRR and PoRR for Π_{PE} and Π_{PO}. We require that both Π_{PE} and Π_{PO} are (selectively) *attribute-hiding* [6]. This security notion says that the adversary learns nothing about the message *and* the associated attribute (except the information that is trivially leaked by the keys that the adversary has).

Intuitively, a zero-knowledge (ZK) proof system $\mathcal{ZKP}$ is a proof system that combines two fundamental properties. The first property, soundness, says that it is (computationally) infeasible to produce a ZK proof of a wrong statement. The second property, zero-knowledge, means that no information besides the validity of the proven statement is leaked. A non-interactive zero-knowledge proof is a zero-knowledge protocol consisting of one message sent by the prover to the verifier. A zero-knowledge proof of knowledge additionally ensures that the prover knows the witnesses to the given statement. We denote by $PK\{(\vec{x}) : F\}$ a zero-knowledge proof of knowledge of the variables in $\vec{x}$ such that the statement F holds. Here, $\vec{x}$
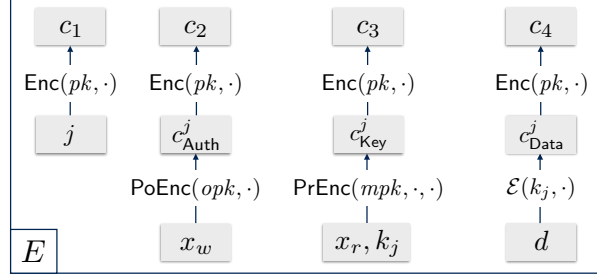
Figure 1: The structure of an entry with index j, payload d, write access regulated by the attribute x_w, and read access regulated by the attribute x_r.

is the set of witnesses, existentially quantified in the statement, and the proof does not reveal any of them. For instance, the proof $PK \{(r): c = \mathsf{Rerand}(pk, d, r)\}$ shows that two public-key ciphertexts c and d, encrypted with the same public key pk, encrypt the same plaintext, i.e., c is obtained by rerandomizing d with the secret randomness r.

Structure of an entry and access control modes. Abstractly, database entries are tuples of the form $E = (c_1, c_2, c_3, c_4)$ where $c_1, \ldots, c_4$ are ciphertexts obtained using a public-key encryption scheme (see Figure 1). In particular, c_1 is the encryption of an index j identifying the j-th entry of the database; c_2 is the encryption of a predicate-only ciphertext c^j_{Auth}, which regulates the write access to the payload stored at j using the attribute x_w; c_3 is the encryption of a ciphertext c^j_{Key}, which is in turn the predicate encryption of a private key k_j with attribute x_r, regulating the read access; c_4 is the encryption of the ciphertext c^j_{Data}, which is the encryption with the private key k_j of the data d stored at position j in the database. We use the convention that an index $j > |\mathsf{DB}|$ indicates a dummy entry and we maintain the invariant that every client may write each dummy entry.

Intuitively, in order to implement the access control modes $\bot$, r, and rw on a data index j, each client $\mathcal{C}_i$ is provided with a capability cap_i that is composed of three keys, the secret key corresponding to the top level public-key encryption scheme, a secret key for the predicate-only encryption scheme, and a secret key for the predicate encryption scheme. More specifically, if $\mathcal{C}_i$'s mode for j is $\bot$, then cap_i allows for decrypting neither c^j_{Auth} nor c^j_{Key}. If $\mathcal{C}_i$'s mode for j is r, then cap_i allows for decrypting c^j_{Key} but not c^j_{Auth}. Finally, if $\mathcal{C}_i$'s mode for j is rw, then cap_i allows for decrypting both c^j_{Auth} and c^j_{Key}. Intuitively, in order to replace an entry, a client has to successfully prove that she can decrypt the ciphertext c^j_{Auth}.

System assumptions. We assume that each client has a local storage of $O(\log N)$. Notice that the leaf index mapping has size $O(N)$, but the local client storage can be decreased to $O(\log N)$ by applying a standard ORAM construction recursively to it, as proposed by Shi *et al.* [12]. Additionally, the data owner stores a second database $\mathcal{ADB}$ that contains the attributes x_w and x_r associated to every entry in DB as well as predicates f_i associated to the client identities $\mathcal{C}_i$. Intuitively, $\mathcal{ADB}$ implements the access control matrix **AC** used in Definition 1. Since also $\mathcal{ADB}$ has size $O(N)$, we use the same technique as the one employed for the index structure. We further assume that clients establish authenticated channels with the server. These channels may be anonymous (e.g., by using anonymity networks [13] and anonymous credentials for the login [14]–[17]), but not necessarily.

Algorithm 1 $(cap_{\mathcal{O}}, \mathsf{DB}) \leftarrow \mathsf{gen}(1^{\lambda}, n)$.

Input: security parameter 1^{λ}, number of clients n

Output: the capability of the data owner $cap_{\mathcal{O}}$

1: $(pk, sk) \leftarrow \mathsf{Gen}_{\mathsf{PKE}}(1^{\lambda})$
2: $(opk, osk) \leftarrow \mathsf{PoGen}(1^{\lambda}, n)$
3: $(mpk, psk) \leftarrow \mathsf{PrGen}(1^{\lambda}, n)$
4: give pk to the server $\mathcal{S}$
5: initialize DB on $\mathcal{S}$, $\mathcal{ADB} := \{\}$, $cnt_{\mathcal{C}} := 0$, $cnt_{E} := 0$
6: **return** $cap_{\mathcal{O}} := (cnt_{\mathcal{C}}, cnt_{E}, sk, osk, psk)$

Algorithm 2 $\{cap_i, \mathsf{deny}\} \leftarrow \mathsf{addCl}(cap_{\mathcal{O}}, \mathbf{a})$.

Input: the capability of $\mathcal{O}$ $cap_{\mathcal{O}}$ and an access control list $\mathbf{a}$ for the client to be added

Output: a capability cap_i for client $\mathcal{C}_i$ in case of success, deny otherwise

1: parse $cap_{\mathcal{O}}$ as $(cnt_{\mathcal{C}}, cnt_{E}, sk, osk, psk)$
2: **if** $|\mathbf{a}| \neq cnt_{E}$ **then return** deny
3: **end if**
4: $cnt_{\mathcal{C}} := cnt_{\mathcal{C}} + 1$
5: compute f_i s.t. the following holds for $1 \leq j \leq |\mathbf{a}|$ and all $(x_{w,j}, x_{r,j}) := \mathcal{ADB}(j)$
      if $\mathbf{a}(j) = \bot$ then $f_i(x_{w,j}) = f_i(x_{r,j}) = 0$
      if $\mathbf{a}(j) = r$ then $f_i(x_{w,j}) = 0$ and $f_i(x_{r,j}) = 1$
      if $\mathbf{a}(j) = rw$ then $f_i(x_{w,j}) = f_i(x_{r,j}) = 1$
6: $\mathcal{ADB} := \mathcal{ADB}[\mathcal{C}_i \mapsto f_i]$
7: $osk_{f_i} \leftarrow \mathsf{PoKGen}(osk, f_i)$, $sk_{f_i} \leftarrow \mathsf{PrKGen}(psk, f_i)$
8: **return** $cap_i := (sk, osk_{f_i}, sk_{f_i})$

3.2 Description of the Algorithms

Implementation of $(cap_{\mathcal{O}}, \mathsf{DB}) \leftarrow \mathsf{gen}(1^{\lambda}, n)$ **(Algorithm 1).** Intuitively, the data owner initializes the cryptographic schemes (lines 1.1–1.3) as well as the rest of the infrastructure (lines 1.4–1.5), and finally outputs $\mathcal{O}$'s capability (line 1.6).[3] Notice that this algorithm takes as input the maximum number n of clients in the system, since this determines the size of the predicates ruling access control, which the predicate(-only) encryption schemes are parameterized by.

Implementation of $\{cap_i, \mathsf{deny}\} \leftarrow \mathsf{addCl}(cap_{\mathcal{O}}, \mathbf{a})$ **(Algorithm 2).** This algorithm allows $\mathcal{O}$ to register a new client in the system. Specifically, $\mathcal{O}$ creates a new capability for the new client $\mathcal{C}_i$ according to the given access permission list $\mathbf{a}$ (lines 2.5–2.8). If $\mathcal{O}$ wants to add more clients than n, the maximum number she initially decided, she can do so at the price of re-initializing the database. In particular, she has to setup new predicate-and predicate-only encryption schemes, since these depend on n. Secondly, she has to distribute new capabilities to all clients. Finally, for each entry in the database, she has to re-encrypt the ciphertexts c_{Auth} and c_{Key} with the new keys.

Implementation of $\{\mathsf{DB}', \mathsf{deny}\} \leftarrow \langle \mathcal{C}_{\mathsf{addE}}(cap_{\mathcal{O}}, \mathbf{a}, d), \mathcal{S}_{\mathsf{addE}}(\mathsf{DB}) \rangle$ **(Algorithm 3).** In this algorithm, $\mathcal{O}$ adds a new entry that contains the payload d to the database. Furthermore, the new entry is protected according to the given access permission list $\mathbf{a}$. Intuitively, $\mathcal{O}$ assigns the new entry to a random leaf and downloads the corresponding path in the database (lines 3.4–3.5). It then creates the new entry and substitutes it for a dummy entry (lines 3.6–3.9). Finally, $\mathcal{O}$

---

[3]For simplifying the notation, we assume for each encryption scheme that the public key is part of the secret key.

Algorithm 3 $\{\mathsf{DB}', \mathsf{deny}\} \leftarrow \langle \mathcal{C}_{\mathsf{addE}}(cap_{\mathcal{O}}, \mathbf{a}, d), \mathcal{S}_{\mathsf{addE}}(\mathsf{DB})\rangle$.

Input: the capability of $\mathcal{O}$ $cap_{\mathcal{O}}$, an access control list $\mathbf{a}$ and the data d for the entry to be added

Output: a changed database DB' on $\mathcal{S}$ in case of success, deny otherwise

1: parse $cap_{\mathcal{O}}$ as $(cnt_{\mathcal{C}}, cnt_E, sk, osk, psk)$
2: **if** $|\mathbf{a}| \neq cnt_{\mathcal{C}}$ **then return** deny
3: **end if**
4: $cnt_E := cnt_E + 1$, $j := cnt_E$, $l_j \leftarrow \{0,1\}^D$, $\mathcal{LM} := \mathcal{LM}[j \mapsto l_j]$
5: let $E_1, \ldots, E_{b(D+1)}$ be the path from ρ to T_{D,l_j} downloaded from $\mathcal{S}$ ($E_i = (c_{1,i}, c_{2,i}, c_{3,i}, c_{4,i})$)
6: let k be such that $\mathsf{Dec}(sk, c_{1,k}) > |\mathsf{DB}|$
7: compute $(x_{w,j}, x_{r,j})$ s.t. the following holds for $1 \leq i \leq |\mathbf{a}|$ and all $f_i := \mathcal{ADB}(\mathcal{C}_i)$
      if $\mathbf{a}(i) = \bot$ then $f_i(x_{w,j}) = f_i(x_{r,j}) = 0$
      if $\mathbf{a}(i) = r$ then $f_i(x_{w,j}) = 0$, $f_i(x_{r,j}) = 1$
      if $\mathbf{a}(i) = rw$ then $f_i(x_{w,j}) = f_i(x_{r,j}) = 1$
8: $\mathcal{ADB} := \mathcal{ADB}[j \mapsto (x_{w,j}, x_{r,j})]$
9: $E_k := (c_{1,k}, c_{2,k}, c_{3,k}, c_{4,k})$ where

$$
\begin{aligned}
k_j &\leftarrow \mathsf{Gen}_{\mathsf{SE}}(1^\lambda) & c_{1,k} &\leftarrow \mathsf{Enc}(pk, j) \\
c^j_{\mathsf{Auth}} &\leftarrow \mathsf{PoEnc}(opk, x_{w,j}) & c_{2,k} &\leftarrow \mathsf{Enc}(pk, c^j_{\mathsf{Auth}}) \\
c^j_{\mathsf{Key}} &\leftarrow \mathsf{PrEnc}(mpk, x_{r,j}, k_j) & c_{3,k} &\leftarrow \mathsf{Enc}(pk, c^j_{\mathsf{Key}}) \\
c^j_{\mathsf{Data}} &\leftarrow \mathcal{E}(k_j, d) & c_{4,k} &\leftarrow \mathsf{Enc}(pk, c^j_{\mathsf{Data}})
\end{aligned}
$$

10: **for all** $1 \leq \ell \leq b(D+1)$, $\ell \neq k$ **do**
11:     select r_ℓ uniformly at random
12:     $E'_\ell \leftarrow \mathsf{Rerand}(pk, E_\ell, r_\ell)$
13: **end for**
14: upload $E'_1, \ldots, E'_{k-1}, E_k, E'_{k+1}, \ldots, E'_{b(D+1)}$ to $\mathcal{S}$

---

rerandomizes the entries so as to hide from $\mathcal{S}$ which entry changes, and finally uploads the modified path to $\mathcal{S}$ (lines 3.10–3.14).

Eviction. In all ORAM constructions, the client has to rearrange the entries in the database in order to make subsequent accesses unlinkable to each other. In the tree construction we use [5], this is achieved by first assigning a new, randomly picked, leaf index to the read or written entry. After that, the entry might no longer reside on the path from the root to its designated leaf index and, thus, has to be moved. This procedure is called *eviction* (Algorithm 4).

This algorithm assigns the entry to be evicted to a new leaf index (line 4.1). It then locally shuffles and rerandomizes the given path according to a permutation π (lines 4.2–4.4). After replacing the old path with a new one, the evicted entry is supposed to be stored in a node along the path from the root to the assigned leaf, which always exists since the root is part of the permuted nodes. A peculiarity of our setting is that clients are not trusted and, in particular, they might store a sequence of ciphertexts in the database that is not a permutation of the original path (e.g., they could store a path of dummy entries, thereby cancelling the original data).

Integrity proofs. To tackle this problem, a first technical novelty in our construction is, in the read and write protocols, to let the client output the modified path along with a proof of shuffle correctness [18, 8], which has to be verified by the server ($s = 1$, lines 4.6–4.7). As the data owner is assumed to be honest, she does not have to send a proof in the chMode protocol ($s = 0$, line 4.9).

Algorithm 4 $(E_1'', \ldots, E_{b(D+1)}'', \pi, [P]) \leftarrow \mathsf{Evict}(E_1, \ldots, E_{b(D+1)}, s, j, k)$.

Input: a list of entries $E_1, \ldots, E_{b(D+1)}$, a bit s, an index j, and a position k in the list

Output: a permuted and rerandomized list of entries $E_1'', \ldots, E_{b(D+1)}''$, a permutation π, and a proof of shuffle correctness (if $s = 1$)

1: $l_j \leftarrow \{0,1\}^D$, $\mathcal{LM} := \mathcal{LM}[j \mapsto l_j]$
2: compute a permutation π s.t. $\pi(k) = 1$ and for all other $\ell \neq k$, π pushes ℓ down on the path from ρ ($= E_1, \ldots, E_b$) to the current leaf node ($= E_{bD+1}, \ldots, E_{b(D+1)}$) as long as the index of the ℓ-th entry still lies on the path from ρ to its designated leaf node.
3: $E_1', \ldots, E_{b(D+1)}' := E_{\pi^{-1}(1)}, \ldots, E_{\pi^{-1}(b(D+1))}$
4: let $E_1'', \ldots, E_{b(D+1)}''$ be the rerandomization of $E_1', \ldots, E_{b(D+1)}'$ as described in 3.10–3.13 (including k)
5: **if** $s = 1$ **then**
6: $\quad P := PK \left\{ \begin{array}{l} (\pi, r_1, \ldots, r_{b(D+1)}) : \\ \quad \forall \ell. \ E_\ell = \mathsf{Rerand}(pk, E_{\pi^{-1}(\ell)}, r_\ell) \end{array} \right\}$
7: $\quad$ **return** $E_1'', \ldots, E_{b(D+1)}'', \pi, P$
8: **else**
9: $\quad$ **return** $E_1'', \ldots, E_{b(D+1)}'', \pi$
10: **end if**

Implementation of $\langle \mathcal{C}_{\mathsf{chMode}}(cap_{\mathcal{O}}, \mathbf{a}, j), \mathcal{S}_{\mathsf{chMode}}(\mathsf{DB}) \rangle$ (**Algorithm 5**). In this protocol, $\mathcal{O}$ changes the access mode of the j-th entry in DB according to the new access permission list $\mathbf{a}$. Intuitively, she does so by downloading the path where the entry resides on (lines 5.4–5.5), changing the entry accordingly (lines 5.6–5.11), and uploading a modified and evicted path to the server (lines 5.12–5.13). Naïvely, $\mathcal{O}$ could simply re-encrypt the old key with the new attributes. However, if a client keeps a key for index j locally and his access on j is revoked, then he can still access the payload. Hence, $\mathcal{O}$ also picks a new key and re-encrypts the payload.

Implementation of $\{d, \mathsf{deny}\} \leftarrow \langle \mathcal{C}_{\mathsf{read}}(cap_i, j), \mathcal{S}_{\mathsf{read}}(\mathsf{DB}) \rangle$ (**Algorithm 6**). Intuitively, the client downloads the path which index j is assigned to and searches for the corresponding entry (lines 6.4–6.6). She then evicts the downloaded path, subject to the restriction that some dummy entry afterwards resides in the top position of the root node (lines 6.7–6.8). $\mathcal{C}$ uploads the evicted path together with a proof of shuffle correctness to $\mathcal{S}$ who verifies the proof and replaces the old with the new path in case of successful verification (line 6.9).

Obliviousness in presence of integrity proofs. $\mathcal{C}$ could in principle stop here since she has read the desired entry. However, in order to fulfill the notion of obliviousness (Definition 19), the read and write operations must be indistinguishable. In single-client ORAM constructions, $\mathcal{C}$ can make write indistinguishable from read by simply modifying the content of the desired entry before uploading the shuffled path to the server. This approach does not work in our setting, due to the presence of integrity proofs. Intuitively, in read, it would suffice to produce a proof of shuffle correctness, but this proof would not be the same as the one used in write, where one element in the path changes. Hence another technical novelty in our construction is the last part of the read protocol (lines 6.10–6.14), which "simulates" the write protocol despite the presence of integrity proofs. This is explained below, in the context of the write protocol.

Implementation of $\{\mathsf{DB}', \mathsf{deny}\} \leftarrow \langle \mathcal{C}_{\mathsf{write}}(cap_i, j, d), \mathcal{S}_{\mathsf{write}}(\mathsf{DB}) \rangle$ (**Algorithm 7**). Firstly, $\mathcal{C}$ reads the element that she wishes to change (line 7.1). Secondly, $\mathcal{C}$ evicts the path with the difference that here the first entry in the root node is the element that $\mathcal{C}$ wants to change, as opposed to a dummy entry like in read (line 7.8). It is important to observe that the shuffle proof sent to the server (line 4.6) is indistinguishable in read and write since it hides both the

---

Algorithm 5 $\langle \mathcal{C}_{\mathsf{chMode}}(cap_{\mathcal{O}}, \mathbf{a}, j), \mathcal{S}_{\mathsf{chMode}}(\mathsf{DB}) \rangle$.

Input: the capability of $\mathcal{O}$ $cap_{\mathcal{O}}$, an access control list $\mathbf{a}$, and an index j

Output: deny if the algorithm fails

1: parse $cap_{\mathcal{O}}$ as $(cnt_{\mathcal{C}}, cnt_E, sk, osk, psk)$
2: **if** $|\mathbf{a}| \neq cnt_{\mathcal{C}}$ or $j > cnt_E$ **then return** deny
3: **end if**
4: $l_j := \mathcal{LM}(j)$
5: let $E_1, \ldots, E_{b(D+1)}$ be the path from ρ to T_{D,l_j} downloaded from $\mathcal{S}$ ($E_i = (c_{1,i}, c_{2,i}, c_{3,i}, c_{4,i})$)

6: let k be s.t. $\mathsf{Dec}(sk, c_{1,k}) = j$
7: compute $(x_{w,j}, x_{r,j})$ according to 3.7 and subject to all f_i in $\mathcal{ADB}$, also add them to $\mathcal{ADB}$ (3.8)
8: $(x'_{w,j}, x'_{r,j}) := \mathcal{ADB}(j)$
9: let f be s.t. $f(x'_{w,j}) = f(x'_{r,j}) = 1$
   $\quad sk_f \leftarrow \mathsf{PrKGen}(psk, f) \quad c^j_{\mathsf{Key}} \leftarrow \mathsf{Dec}(sk, c_{3,k})$
10: $k'_j \leftarrow \mathsf{PrDec}(sk_f, c^j_{\mathsf{Key}}) \quad c^j_{\mathsf{Data}} \leftarrow \mathsf{Dec}(sk, c_{4,k})$
   $\quad d \leftarrow \mathcal{D}(k'_j, c^j_{\mathsf{Data}})$
11: compute E'_k as in 3.9
12: $(E''_1, \ldots, E''_{b(D+1)}, \pi) := \mathsf{Evict}(E_1, \ldots, E_{k-1}, E'_k, E_{k+1}, \ldots, E_{b(D+1)}, 0, j, k)$
13: upload $E''_1, \ldots, E''_{b(D+1)}$ to $\mathcal{S}$

---

permutation and the randomness used to rerandomize the entries. So far, we have shown how $\mathcal{C}$ can upload a shuffled and rerandomized path to the server without modifying the content of any entry.

In write, $\mathcal{C}$ can now replace the first entry in the root node with the entry containing the new payload (lines 7.12–7.13). In read, this step is simulated by rerandomizing the first entry of the root node, which is a dummy entry (line 6.12).

The integrity proofs P_{Auth} and P_{Ind} produced in read and write are indistinguishable (lines 6.11 and 6.13 for both): in both cases, they prove that $\mathcal{C}$ has the permission to write on the first entry of the root node and that the index has not changed. Notice that this proof can be produced also in read, since all clients have write access to dummy entries.

Permanent Entries. Some application scenarios of GORAM might require determined entries of the database not to be modifiable nor deletable, not even by the data owner herself (for instance, in the case of PHRs, the user should not be able to cancel diagnostic results in order to pay lower insurance fees). Even though we did not explicitly describe the construction, we mention that such a property can be achieved by assigning a binary attribute (*modifiable* or *permanent*) to each entry and storing a commitment to this in the database. Every party that tries to modify a given entry, including the data owner, has to provide a proof that the respective attribute is set to *modifiable*. Due to space constraints we omit the algorithm, but it can be efficiently instantiated using El Gamal encryption and Σ-protocols.

3.3 Batched Zero-Knowledge Proofs of Shuffle

A zero-knowledge proof of a shuffle of a set of ciphertexts proves in zero-knowledge that a new set of ciphertexts contains the same plaintexts in permuted order. In our system the encryption of an entry, for reasonable block sizes, yields in practice hundreds of ciphertexts, which means that we have to perform hundreds of shuffle proofs. These are computable in polynomial-time

---

Algorithm 6 $\{d, \mathsf{deny}\} \leftarrow \langle \mathcal{C}_{\mathsf{read}}(cap_i, j), \mathcal{S}_{\mathsf{read}}(\mathsf{DB}) \rangle$.

Input: the capability of the client executing the protocol cap_i and the index j to be read
Output: the data payload d in case of success, deny otherwise

1: parse cap_i as (sk, osk_f, sk_f)
2: **if** $j > |\mathsf{DB}|$ **then return** deny
3: **end if**
4: $l_j := \mathcal{LM}(j)$
5: let $E_1, \ldots, E_{b(D+1)}$ and k be as in lines 5.5–5.6
6: extract d from E_k as in line 5.10
7: let ℓ be s.t. $\mathsf{Dec}(sk, c_{1,\ell}) > |\mathsf{DB}|$
8: $(E_1'', \ldots, E_{b(D+1)}'', \pi, P) := \mathsf{Evict}(E_1, \ldots, E_{b(D+1)}, 1, j, \ell)$
9: upload $E_1'', \ldots, E_{b(D+1)}''$ and P to $\mathcal{S}$
10: $c_{\mathsf{Auth}}^j \leftarrow \mathsf{Dec}(sk, c_{2,1}'')$
11: $P_{\mathsf{Auth}} := PK\left\{(osk_f) : \mathsf{PoDec}(osk_f, c_{\mathsf{Auth}}^j) = 1\right\}$
12: $E_1''' := (c_{1,1}''', c_{2,1}''', c_{3,1}''', c_{4,1}''')$ where
$\quad r_1, \ldots, r_4$ are selected uniformly at random
$\quad c_{l,1}''' \leftarrow \mathsf{Rerand}(pk, c_{l,1}'', r_l)$ for $l \in \{1, 3, 4\}$
$\quad c_{2,1}''' \leftarrow \mathsf{Enc}(pk, \mathsf{PoRR}(opk, c_{\mathsf{Auth}}^j, r_2))$
13: $P_{\mathsf{Ind}} := PK\left\{(r_1) : c_{1,1}'' = \mathsf{Rerand}(pk, c_{1,1}''', r_1)\right\}$
14: upload E_1''', P_{Auth}, P_{Ind}, and the necessary information to access c_{Auth}^j to $\mathcal{S}$

---

but, even using the most efficient known solutions (e.g., [8, 19]), not fast enough for practical purposes. This problem has been addressed in the literature but the known solutions typically reveal part of the permutation (e.g., [20]), which would break obliviousness and, thus, are not applicable in our setting.

To solve this problem we introduce a new proof technique that we call *batched zero-knowledge proofs of shuffle*, based on the idea of "batching" several instances and verifying them together. Our interactive protocol takes advantage of the homomorphic property of the top layer public-key encryption scheme in order to batch the instances. On a high level, we represent the path, which the client proves the shuffle of, as an n-by-m matrix where n is the number of entries (i.e., the path length) and m is the number of blocks of ciphertexts per entry. The common inputs of the prover $\mathcal{P}$ and the verifier $\mathcal{V}$ are the two matrices $\mathbf{A}$ and $\mathbf{A}'$ characterizing the path stored on the database and the path shuffled and re-randomized by the client, respectively. $\mathcal{P}$ additionally knows the permutation π and the randomnesses $\mathbf{R}$ used for rerandomizing the entries.

Intuitively, the batching algorithm randomly selects a subset of columns (i.e., block indices) and computes the row-wise product of the corresponding blocks for each row. It then computes the proof of shuffle correctness on the resulting single-block ciphertexts. The property we would like to achieve is that modifying even a single block in a row should lead to a different product and, thus, be detected. Notice that naïvely multiplying all blocks together does not achieve the intended property, as illustrated by the following counterexample:

$$\begin{pmatrix} \mathsf{Enc}(pk, 3) & \mathsf{Enc}(pk, 4) \\ \mathsf{Enc}(pk, 5) & \mathsf{Enc}(pk, 2) \end{pmatrix} \qquad \begin{pmatrix} \mathsf{Enc}(pk, 2) & \mathsf{Enc}(pk, 6) \\ \mathsf{Enc}(pk, 5) & \mathsf{Enc}(pk, 2) \end{pmatrix}$$

In the above matrices, the rows have not been permuted but rather changed. Still, the row-wise product is preserved, i.e., 12 in the first and 10 in the second. Hence, we cannot compute the product over all columns. Instead, as proved in the long version, the intended property can be achieved with probability at least $\frac{1}{2}$ if each column is included in the product with probability $\frac{1}{2}$.

Algorithm 7 $\{\mathsf{DB}', \mathsf{deny}\} \leftarrow \langle \mathcal{C}_{\mathsf{write}}(cap_i, j, d), \mathcal{S}_{\mathsf{write}}(\mathsf{DB}) \rangle$.

Input: the capability of the client executing the protocol cap_i, the index j to be written, and the payload d

Output: deny if the algorithm fails

1: execute 6.1–6.6
8: $(E''_1, \ldots, E''_{b(D+1)}, \pi, P) := \mathsf{Evict}(E_1, \ldots, E_{b(D+1)}, 1, j, k)$
9: execute 6.9–6.11
12: $E'''_1 := (c'''_{1,1}, c'''_{2,1}, c'''_{3,1}, c'''_{4,1})$ where
$\quad\quad r_1, r_2, r_3$ are selected uniformly at random
$\quad\quad c'''_{1,1} \leftarrow \mathsf{Rerand}(pk, c''_{1,1}, r_1)$
$\quad\quad c'''_{2,1} \leftarrow \mathsf{Enc}(pk, \mathsf{PoRR}(opk, c^j_{\mathsf{Auth}}, r_2))$
$\quad\quad c'''_{3,1} \leftarrow \mathsf{Enc}(pk, \mathsf{PrRR}(mpk, c^j_{\mathsf{Key}}, r_3))$
$\quad\quad c^j_{\mathsf{Data}} \leftarrow \mathcal{E}(k_j, d)$
$\quad\quad c'''_{4,1} \leftarrow \mathsf{Enc}(pk, c^j_{\mathsf{Data}})$
13: execute 6.13–6.14

Algorithm 8 Batched ZK Proofs of Shuffle.

Input of $\mathcal{P}$: $\mathbf{A}$, $\mathbf{A}'$, π, $\mathbf{R}$

Input of $\mathcal{V}$: $\mathbf{A}$, $\mathbf{A}'$

1: $\mathcal{V}$ randomly selects $\vec{a} \leftarrow \{0, 1\}^m$ and sends it to $\mathcal{P}$.
2: $\mathcal{P}$ computes for all $1 \leq i \leq n$ the partial ciphertext products
$\quad\quad \theta_i = \prod_{j=1}^m a_j \mathbf{A}_{i,j}$ and $\theta'_i = \prod_{j=1}^m a_j \mathbf{A}'_{i,j}$
and the corresponding partial randomness sum
$\quad\quad r_i = \sum_{j=1}^m a_j \mathbf{R}_{i,j}$
where a_j is the j-th bit of $\vec{a}$. $\mathcal{V}$ also computes $\vec{\theta}$ and $\vec{\theta}'$.
3: $\mathcal{V}$ and $\mathcal{P}$ run the protocol for the proof of shuffle correctness [8] on $\vec{\theta}$, $\vec{\theta}'$, π, and $\vec{r}$.

Although a probability of $\frac{1}{2}$ is not sufficient in practice, repeating the protocol k times increases the probability to $(1 - \frac{1}{2^k})$.

The detailed construction is depicted in Algorithm 8. In line 8.1, $\mathcal{V}$ picks a challenge, which indicates which column to include in the homomorphic product. Upon receiving the challenge, in line 8.2, $\mathcal{P}$ and $\mathcal{V}$ compute the row-wise multiplication of the columns indicated by the challenge. Finally, $\mathcal{V}$ and $\mathcal{P}$ run an off-the-shelf shuffle proof on the resulting ciphertext lists (line 8.3).

It follows from the protocol design that our approach does not affect the completeness of the underlying proof of shuffle correctness, the same holds true for the zero-knowledge and proof of knowledge properties. Furthermore, any malicious prover who does not apply a correct permutation is detected by the verifier with probability at least $1/2$.

Finally, the protocol can be made non-interactive by using the Fiat-Shamir heuristic [21].

To summarize, our new approach preserves all of the properties of the underlying shuffle proof while being significantly more efficient. Our proof system eliminates the dependency of the number of proofs with respect to the block size, making it dependent only on k and on the complexity of the proof itself.

4 Accountable Integrity (A-GORAM)

In this section we relax the integrity property by introducing the concept of accountability. In particular, instead of letting the server check the correctness of client operations, we develop a technique that allows clients to detect *a posteriori* non-authorized changes on the database and blame the misbehaving party. Intuitively, each entry is accompanied by a tag (technically, a chameleon hash along with the randomness corresponding to that entry), which can only be produced by clients having write access. All clients can verify the validity of such tags and, eventually, determine which client inserted an entry with an invalid tag. This makes the construction more efficient and scalable, significantly reducing the computational complexity both on the client and on the server side, since zero-knowledge proofs are no longer necessary and, consequently, the outermost encryption can be implemented using symmetric, as opposed to asymmetric, cryptography. Such a mechanism is supposed to be paired with a data versioning protocol in order to avoid data losses: as soon as one of the clients detects an invalid entry, the misbehaving party is punished and the database is reverted to the last safe state (i.e., a state where all entries are associated with a valid tag).

4.1 Prerequisites

In the following, we review some additional cryptographic primitives and explain the structure of the log file.

Cryptographic preliminaries. Intuitively, a chameleon hash function is a randomized collision-resistant hash function that provides a trapdoor. Given the trapdoor it is possible to efficiently compute collisions. A chameleon hash function is a tuple of PPT algorithms $\Pi_{\mathsf{CHF}} = (\mathsf{Gen}_{\mathsf{CHF}}, \mathsf{CH}, \mathsf{Col})$. The setup algorithm $\mathsf{Gen}_{\mathsf{CHF}}$ takes as input a security parameter 1^λ and outputs a key pair (cpk, csk), where cpk is the public key and csk is the secret key. The chameleon hash function CH takes as input the public key cpk, a message m, and randomness r; it outputs a hash tag t. The collision function Col takes as input the secret key csk, a message m, randomness r, and another message m'; it outputs a new randomness r' such that $\mathsf{CH}(cpk, m, r) = \mathsf{CH}(cpk, m', r')$. For our construction, we use chameleon hash functions providing key-exposure freeness [22]. Intuitively, this property states that no adversary is able to find a fresh collision, without knowing the secret key csk, even after seeing polynomially many collisions.

We denote by $\Pi_{\mathsf{DS}} = (\mathsf{Gen}_{\mathsf{DS}}, \mathsf{sign}, \mathsf{verify})$ a digital signature scheme. We require a signature scheme that is existentially unforgeable. Intuitively, this notion ensures that it is infeasible for any adversary to output a forgery (i.e., a fresh signature σ on a message m without knowing the signing key) even after seeing polynomially many valid (σ, m) pairs.

Structure of the log file. We use a log file Log so as to detect who has to be held accountable in case of misbehavior. Log is append-only and consists of the list of paths uploaded to the server, each of them signed by the respective client.

4.2 Construction

Structure of entries. The structure of an entry in the database is depicted in Figure 2. An entry E is protected by a top-level private-key encryption scheme with a key $\mathcal{K}$ that is shared by the data owner $\mathcal{O}$ and all clients $\mathcal{C}_1, \ldots, \mathcal{C}_n$. Under the encryption, E contains several elements, which we explain below:
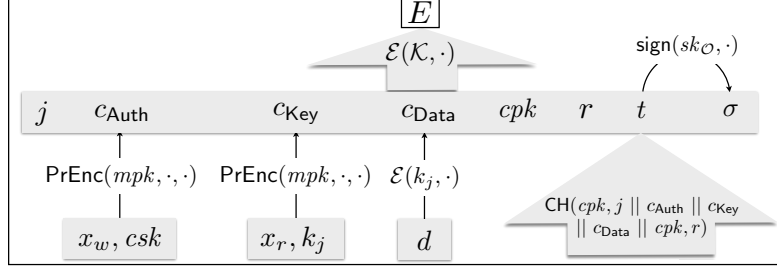
- j is the index of the entry;

Figure 2: The structure of an entry in the database.

- c_{Auth} is a predicate encryption ciphertext that encrypts the private key csk of a chameleon hash function under an attribute x_w, which regulates the write access;

- c_{Key} and c_{Data} are unchanged;

- cpk is the public key of a chameleon hash function, i.e., the counterpart of csk encrypted in c_{Auth};

- r is some randomness used in the computation of t;

- t is a chameleon hash tag, produced by hashing the concatenation of j, c_{Auth}, c_{Key}, c_{Data}, and cpk under randomness r using the public key cpk;

- σ is a signature on the chameleon hash tag t, signed by the data owner $\mathcal{O}$.

Intuitively, only clients with write access are able to decrypt c_{Auth}, and thus to retrieve the key csk required to compute a collision for the new entry d' (i.e., to find a randomness r' such that the chameleon hash t for the old entry d and randomness r is the same as the one for d' and r'). The fundamental observation is that the modification of an entry is performed without changing the respective tag. Consequently, the signature σ is the same for the old and for the new entry. Computing a collision is the only way to make the tag t, originally signed by the data owner, a valid tag also for the new entry d'. Therefore verifying the signature and the chameleon hash suffices to make sure that the entry has been only modified by authorized clients.

Basic Algorithms. The basic algorithms follow the ones defined in Section 3.2, except for natural adaptions to the new entry structure. Furthermore, the zero-knowledge proofs are no longer computed and the rerandomization steps are substituted by re-encryptions. Finally, clients upload on the server signed paths, which are stored in the Log.

Entry Verification. We introduce an auxiliary verification function that clients run in order to verify the integrity of an entry. During the execution of any protocol below we maintain the invariant that, whenever a client i (or the data owner himself) parses an entry j that he downloaded from the server, he executes Algorithm 9. If the result is $\bot$, then the client runs $\mathsf{blame}(cap_i, \mathsf{Log}, j)$.

Blame. In order to execute the function $\mathsf{blame}(cap_i, \mathsf{Log}, j)$, the client must first retrieve Log from the server. Afterwards, she parses backwards the history of modifications by decrypting the paths present in the Log. The client stops only when she finds the desired entry indexed by j in a consistent state, i.e., the data hashes to the associated tag t and the signature is valid. At this point the client moves forwards on the Log until she finds an uploaded path where the entry j is supposed to lay on (the entry might be associated with an invalid tag or

Algorithm 9 The pseudo-code for the verification of an entry in the database which is already decrypted.

Input: An entry $(j, c_{\mathsf{Auth}}, c_{\mathsf{Key}}, c_{\mathsf{Data}}, r, cpk, t, \sigma)$ and the verification key $vk_{\mathcal{O}}$ of $\mathcal{O}$.
Output: $\top$ if verification succeeds, $\bot$ otherwise.
1: **if** $t = \mathsf{CH}(cpk, j \mathbin{||} c_{\mathsf{Auth}} \mathbin{||} c_{\mathsf{Key}} \mathbin{||} c_{\mathsf{Data}} \mathbin{||} cpk, r)$ and $\top = \mathsf{verify}(\sigma, vk_{\mathcal{O}}, t)$ **then**
2:     **return** $\top$
3: **else**
4:     **return** $\bot$
5: **end if**

---

missing). The signature on the path uniquely identifies the client, whose identity is added to a list L of misbehaving clients. Finally, all of the other clients that acknowledged the changes of the inconsistent entry are also added to L, since they did not correctly verify its chameleon signature.

Discussion. As explained above, the accountability mechanism allows for the identification of misbehaving clients with a minimal computational overhead in the regular clients' operation. However, it requires the server to store a log that is linear in the number of modifications to the database and logarithmic in the number of entries. This is required to revert the database to a safe state in case of misbehaviour. Consequently, the blame algorithm results expensive in terms of computation and communication with the server, in particular for the entries that are not regularly accessed. Nonetheless, blame is supposed to be only occasionally executed, therefore we believe this design is acceptable in terms of service usability. Furthermore, we can require all the parties accessing the database to synchronize on a regular basis so as to verify the content of the whole database and to reset the Log, in order to reduce the storage on the server side and, thus, the amount of data to transfer in the blame algorithm. Such an approach could be complemented by an efficient versioning algorithm on encrypted data, which is however beyond the scope of this work and left as a future work. Finally, we also point out that the accountable-integrity property, as presented in this section, sacrifices the anonymity property, since users have to sign the paths they upload to the server. This issue can be easily overcome by using any anonymous credential system that supports revocation [23].

5   Scalable Solution (S-GORAM)

Even though the personal record management systems we consider rely on simple client-based read and write permissions, the predicate encryption scheme used in GORAM and A-GORAM support in principle a much richer class of access control policies, such as role-based access control (RBAC) or attribute-based access control (ABAC) [6]. If we stick to client-based read and write permissions, however, we can achieve a more efficient construction that scales to thousands of clients. To this end, we replace the predicate encryption scheme with a broadcast encryption scheme [24], which guarantees that a specific subset of clients is able to decrypt a given ciphertext. This choice affects the entry structure as follows (cf. Figure 2):

- c_{Key} is the broadcast encryption of k_j;

- c_{Auth} is the broadcast encryption of csk.

The subset of clients that can decrypt c_{Key} (resp. c_{Auth}) is then set to be the same subset that holds *read* (resp. *write*) permissions on the given entry. By applying the aforementioned modifications on top of A-GORAM, we obtain a much more efficient and scalable instantiation,
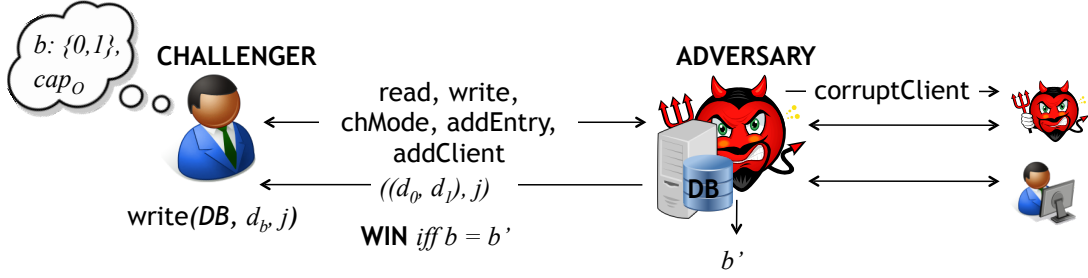
Figure 3: Game for Secrecy.

called S-GORAM, that achieves a smaller constant in the computational complexity (linear in the number of clients). For more details on the performance evaluation and a comparison with A-GORAM, we refer to Section 8.

6 Security and Privacy for Group ORAM

In order to prove the security of our constructions, we formalized the security and privacy properties of a Group ORAM by cryptographic games, which are intuitively introduced below. The formal definitions can be found in the appendix.

6.1 Security and Privacy of Group ORAM

Secrecy. Intuitively, a Group ORAM preserves the secrecy of outsourced data if no party is able to deduce any information about the content of any entry she does not have access to. We formalize this intuition through a cryptographic game, which is illustrated in Figure 3. Intuitively, the challenger initializes an empty database locally and it hands it over to the adversary so as to give him the possibility to adaptively and arbitrarily fill the content of the database. Additionally, the adversary is given the possibility of spawning and corrupting a polynomial number of clients, allowing him to perform operations on the database on their behalf. Hence, this property is proven in a strong adversarial model, without placing any assumption on the server's behavior. At some point of the game the adversary outputs two data and a database index, the challenger flips a coin and it randomly inserts either one of the two payloads in the desired database entry. In order to make the game not trivial, it must be the case that the adversary should not have corrupted any client that holds read permission on such index. We define the adversary to win the game if he correctly guesses which of the two entries has been written. Since the adversary can always randomly guess, we define the system to be *secrecy-preserving* if the adversary cannot win the game with probability non-negligibly greater than $\frac{1}{2}$.

Integrity. A Group ORAM preserves the integrity of its entries if none of the clients can modify an entry to which she does not have write permissions. The respective cryptographic game is depicted in Figure 4. Intuitively, the challenger initializes an empty database DB and a copy DB', providing the adversary with the necessary interfaces to fill the content of DB and to generate and corrupt clients. Every time the adversary queries an interface, the challenger interacts with the respective client playing the server's role and additionally executes locally the same operation on DB' in an honest manner. Note that here the adversary cannot directly operate on the database but he can only operate through the clients: this constraint reflects the honesty assumption of the server. At some point of the execution, the adversary outputs
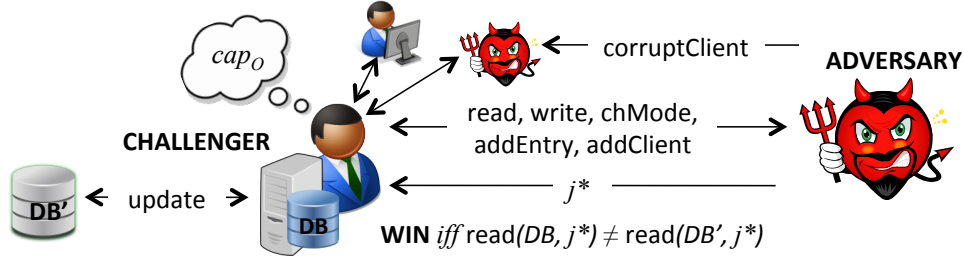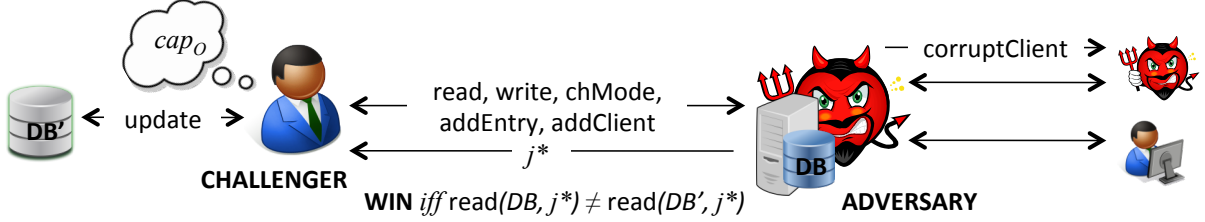
18

Figure 4: Game for Integrity.

Figure 5: Game for Tamper-resistance.

an index of the database (that none of his corrupted clients can write on) and the challenger compares the two entries stored in DB and DB', if they are not the same we say that the adversary wins the game. Since that would imply that a client could potentially be able to modify the content of an entry she does not have access to, we say that the system is *integrity-preserving* if any possible adversary cannot win the game with non-negligible probability.

Tamper-resistance. Intuitively, a Group ORAM is tamper-resistant if the server, even colluding with a subset of malicious clients, is not able to convince an honest client about the integrity of some maliciously modified data. Notice that this property refers to a strong adversarial model, where the adversary may arbitrarily misbehave and collude with clients. Naturally, tamper-resistance holds true only for entries which none of the corrupted clients had ever access to. The respective cryptographic game is depicted in Figure 5. The game is described exactly as in the previous definition except for the fact that the database is this time handed over to the adversary at the very beginning of the experiment so as to allow him to operate directly on it. The challenger maintains a local copy of the database where it performs the same operations that are triggered by the adversary but in an honest way. The winning conditions for the adversary are the same as stated above and we say that the system is *tamper-resistant* if no adversary can win this game with probability greater than a negligible value. Note that there exists a class of attacks where the adversary wins the game by simply providing an old version of the database, which are inherent to the cloud storage setting. We advocate the usage of standard techniques to deal with this kind of attacks (e.g., a gossip protocol among the clients for versioning of the entries [25]) and hence, we rule them out in our formal analysis by implicitly assuming that the information provided by the adversary are relative to the most up to date version of the database that he possesses locally.

Obliviousness. Intuitively, a Group ORAM is oblivious if the server cannot distinguish between two arbitrary query sequences which contain read and write operations. In the cryptographic game depicted in Figure 6, the adversary holds the database on his side and he gets access to the interfaces needed to adaptively and arbitrarily insert content in the database. Thus the server may arbitrarily misbehave but it is not allowed to collude with clients: the adversary can only spawn a polynomial number of them, but he cannot corrupt them. In this game the
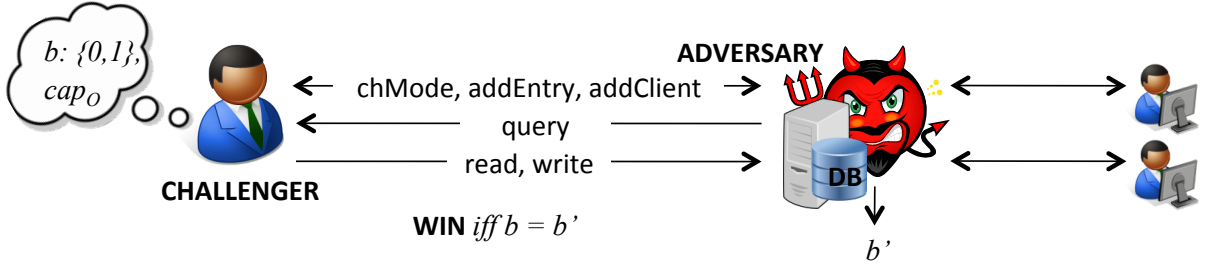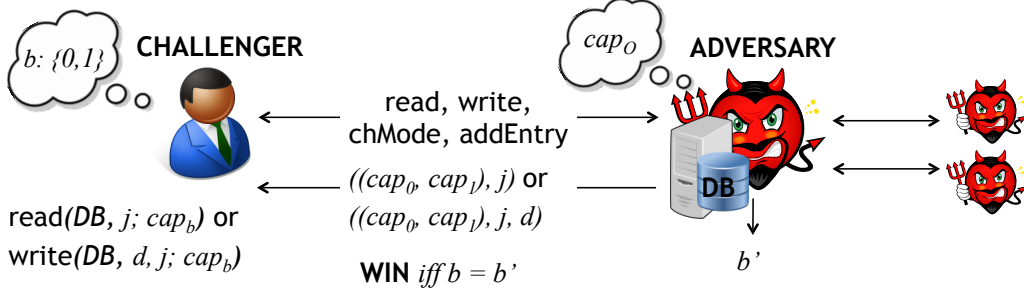
Figure 6: Game for Obliviousness.

Figure 7: Game for Anonymity.

challenger offers an additional interface where the adversary can input two arbitrary queries, i.e., on behalf of arbitrary clients and on arbitrary indices of the database. This interface can be used by the adversary polynomially many times, thus creating the two query sequences at the core of the obliviousness definition. In the beginning of the game the challenger flips a coin and then it always executes either one of the two queries, depending on the outcome of the initial random coin. In order to win, the adversary has to tell the value of the random coin of the challenger, thus distinguishing which query sequence has been executed. This would then mean that the adversary has been able to link some access to a specific memory location, hence we say that a system is *oblivious* if the adversary does not win the game with probability non-negligibly greater than $\frac{1}{2}$.

Anonymity. A Group ORAM is anonymity-preserving if the data owner cannot efficiently link a given operation to a client, among the set of clients having access to the queried index. In the cryptographic game depicted in Figure 7, the setting is equivalent to the secrecy definition except that the challenger also hands over the capability of the data owner to the adversary. Clearly the adversary does not need to corrupt clients since he can spawn them by himself. Additionally, the challenger provides the adversary with an interface that he can query with an operation associated with two arbitrary capabilities. To make the game not trivial, it must hold that both of the capabilities hold the same read and write permissions on the entry selected by the adversary. Based on some initial randomness, the challenger always executes the desired command with either one of the two capabilities and the adversary wins the game if and only if he can correctly determine which capability has been selected. Since this would imply a de-anonymization of the clients, we say that the system is *anonymity-preserving* if the adversary cannot win the game with probability non-negligibly greater than $\frac{1}{2}$.

Accountable Integrity. The server maintains an audit log Log, which holds the evidence of client operations on the database DB. Specifically, each path uploaded to the server as a result of an eviction procedure is signed by the client and appended by the server to the log. After detecting an invalid or missing entry with index j, the client retrieves Log from the server and
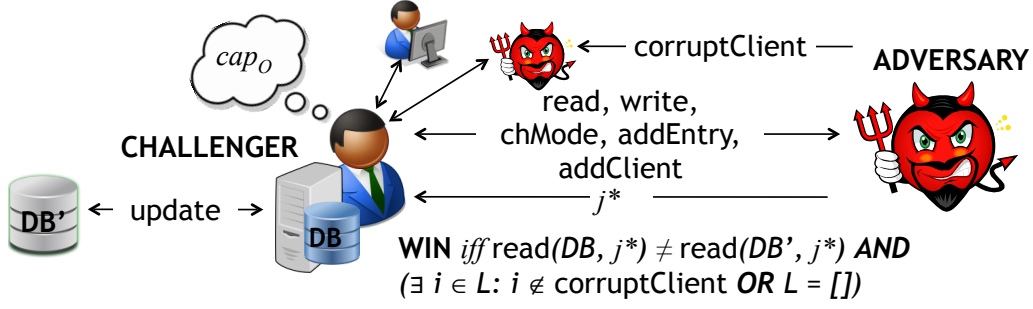
Figure 8: Game for Accountable Integrity.

Property	Server	Collusion
Secrecy	malicious	✓
(Accountable) Integrity	HbC	✗
Tamper-resistance	malicious	✓
Obliviousness	malicious	✗
Anonymity	malicious	✓

Table 1: Security and privacy properties together with their minimal assumptions.

performs the algorithm $\mathsf{blame}(cap_i, \mathsf{Log}, j)$. The output is a list of identities, which correspond to misbehaving parties.

We define the accountable integrity property through a cryptographic game, illustrated in Figure 8. The game is the same as the one for integrity, except for the winning condition, which is adjusted according to the accountability requirements. Intuitively, the adversary wins the game if he manages to modify the entry in the index he provided and the challenger is not able to identify at least one of the corrupted clients that contributed to modify that entry or it erroneously blames some honest party. This means that the blaming procedure always returns at least one of the misbehaving parties and never an honest one. The literature defines this notion of accountability as *fairness* (never blame honest parties) and *completeness* (blame at least one dishonest party) [26]. We say that a system preserves *accountable integrity* if the adversary cannot win the game with more than negligible probability.

Discussion. Table 1 summarizes the security and privacy properties presented in this section, along with the corresponding assumptions. The HbC assumption is in fact only needed for integrity, since the correctness of client operations is checked by the server, thus avoiding costly operations on the client side. We will see in Section 4 that the HbC assumption is still needed for the accountable integrity property, since the server maintains a log of accesses, which allows for blaming misbehaving parties. All other properties hold true even if the server is malicious as long as it does not collude with clients. Furthermore, secrecy, tamper-resistance, and anonymity hold true even if the server is malicious and colludes with clients. The non-collusion assumption is due to the obliviousness property, which is meant to protect the access patterns from the server. Extending this property to non-authorized clients and devising corresponding enforcement mechanisms is beyond the scope of this paper and left as an interesting future work.

Property	G.	A-G.	S-G.
Secrecy	✓	✓	✓
Integrity	✓	Accountable	Accountable
Tamper-resistance	✓	✗	✗
Obliviousness	✓	✓	✓
Anonymity	✓	✗	✗
Access control	ABAC	ABAC	R/W

Table 2: Security and privacy properties achieved by each construction where G. stands for GORAM.

7 Security and Privacy Results

In this section, we show that the Group ORAM instantiations presented in Section 3, in Section 4, and in Section 5 achieve the security and privacy properties stated in Section 6.1. The proofs are reported in Appendix D. A brief overview of the properties guaranteed by each construction is shown in Table 2. As previously discussed, dropping the computationally expensive integrity checks in favor of an accountability mechanism is crucial to achieve efficiency. It follows that A-GORAM and S-GORAM provide accountable integrity as opposed to integrity and tamper resistance. Having an accountable system trivially implies the loss of anonymity, as defined in Section 6.1, although it is still possible to achieve pseudonym-based anonymity by employing anonymous credentials. The other privacy properties of our system, namely secrecy and obliviousness, are fulfilled by all of our instantiations. Moreover, by replacing predicate encryption with broadcast encryption (S-GORAM), we sacrifice the possibility to enforce ABAC policies, although we can still handle client-based read/write permissions.

Before we present the security and privacy results, we start with a soundness result for the batched ZK proof of shuffle, which we prove in Appendix F.

Theorem 1 (Soundness)**.** *Let $\mathcal{ZKP}$ be a zero-knowledge proof system for a proof of shuffle correctness. Then the batched ZK proof of shuffle defined in Algorithm 8 is sound with probability at least $1/2$.*

The following theorems characterize the security and privacy properties achieved by each cryptographic instantiation presented in this paper.

Theorem 2 (GORAM)**.** *Let Π_{PE} and Π_{PO} be an attribute-hiding predicate and predicate-only encryption scheme, Π_{PKE} (resp. Π_{SE}) be a CPA-secure public-key (resp. private-key) encryption scheme, and $\mathcal{ZKP}$ be a zero-knowledge proof system. Then GORAM achieves secrecy, integrity, tamper-resistance, obliviousness, and anonymity.*

Theorem 3 (A-GORAM)**.** *Let Π_{PE} be an attribute-hiding predicate encryption scheme, Π_{SE} be a CPA-secure private-key encryption scheme, Π_{DS} be an existentially unforgeable digital signature scheme, and Π_{CHF} be a collision-resistant, key-exposure free chameleon hash function. Then A-GORAM achieves secrecy, accountable integrity, and obliviousness.*

Theorem 4 (S-GORAM)**.** *Let Π_{BE} be an adaptively secure broadcast encryption scheme, Π_{SE} be a CPA-secure private-key encryption scheme, Π_{DS} be an existentially unforgeable digital signature scheme, and Π_{CHF} be a collision-resistant, key-exposure free chameleon hash function. Then S-GORAM achieves secrecy, accountable integrity, and obliviousness.*

8 Implementation and Experiments

In this section, we present the concrete instantiations of the cryptographic primitives that we previously described (Section 8.1), we study their asymptotic complexity (Section 8.2), describe our implementation (Section 8.3), and discuss the experimental evaluation (Section 8.4).

8.1 Cryptographic Instantiations

Private-key and public-key encryption. We use AES [27] as private-key encryption scheme with an appropriate message padding in order to achieve the elusive-range property [28].[4] Furthermore, we employ the El Gamal encryption scheme [29] for public-key encryption as it fulfills all properties that we require for GORAM, i.e., it is rerandomizable and supports zero-knowledge proofs. We review the scheme below.

$\mathsf{Gen}_{\mathsf{PKE}}(1^\lambda)$: Let p be a prime of length λ such that $p = 2q + 1$ for a prime q. Furthermore, let $\mathbb{G}$ be the subgroup of $\mathbb{Z}_p$ of order q and g be a generator of $\mathbb{G}$. Then, draw a random $x \in \mathbb{Z}_q^*$ and compute $h = g^x$. Output the pair $(sk, pk) = ((p, q, g, x), (p, q, g, h))$.

$\mathsf{Enc}(pk, m)$: In order to encrypt a message $m \in \mathbb{G}$ using the public key $(p, q, g, h) \leftarrow pk$, draw a random $r \in \mathbb{Z}_q^*$ and output the ciphertext $c = (g^r, mh^r)$.

$\mathsf{Dec}(sk, c)$: In order to decrypt a ciphertext $c = (c_1, c_2)$ using the secret key $(p, q, g, x) \leftarrow sk$, compute $m = c_2 \cdot c_1^{-x}$.

$\mathsf{Rerand}(pk, c, r)$: In order to rerandomize a ciphertext $c = (c_1, c_2)$ using the public key $(p, q, g, h) \leftarrow pk$ and randomness $r \in \mathbb{Z}_q^*$, output $c' = (c_1 \cdot g^r, c_2 \cdot h^r)$.

It is also possible to produce information with which one can decrypt a ciphertext $c = (c_1, c_2)$ without knowing the secret key by sending c_1^{-x}.

Encryption schemes for access control. We utilize the predicate encryption scheme introduced by Katz *et al.* [6]. Its ciphertexts are rerandomizable and we also show them to be compatible with the Groth-Sahai proof system [7]. For the details, we refer to Appendix B. Concerning the implementation, the predicate encryption scheme by Katz *et al.* [6] is not efficient enough since it relies on elliptic curves on composite-order groups. In order to reach a high security parameter, the composite-order setting requires us to use much larger group sizes than in the prime-order setting, rendering the advantages of elliptic curves practically useless. Therefore, we use a scheme transformation proposed by David Freeman [30], which works in prime-order groups and is more efficient.

For implementing S-GORAM we use an adaptively secure broadcast encryption scheme by Gentry and Waters [24].

Zero-knowledge proofs. We deploy several non-interactive zero-knowledge proofs. For proving that a predicate-only ciphertext validly decrypts to 1 without revealing the key, we use Groth-Sahai non-interactive zero-knowledge proofs[5] [7]. More precisely, we apply them in the proofs created in line 6.11 (read and write, see Algorithm 6 and Algorithm 7). We employ plaintext-equivalence proofs (PEPs) [19, 31] for the proofs in line 6.13. Furthermore, we use a proof of shuffle correctness [8] and batched shuffle proofs in lines 6.8 and 7.8.

---

[4]It is clear the we cannot provably achieve elusive-range with AES, however, we still use it for practical considerations.

[5]Groth-Sahai proofs are generally not zero-knowledge. However, in our case the witnesses fulfill a special equation for which they are zero-knowledge.

Chameleon hashes and digital signatures. We use a chameleon hash function by Nyberg and Rueppel [22], which has the key-exposure freeness property. We combine the chameleon hash tags with RSA signatures [32].

Implementing permanent entries. We briefly outline how permanent entries can be implemented using El Gamal encryption and equality of discrete logarithm proofs [33]. Let $c_p = \mathsf{Enc}(pk, \mathsf{permanent}) = (g^r, g^{\mathsf{permanent}} \cdot h^r)$ be the ciphertext associated to the entry that is subject to change and $pk = (g, h)$ be the public key of the El Gamal scheme. If $\mathsf{permanent} \neq 1$ then the entry may not be removed from the database completely. Hence, if $\mathcal{O}$ attempts to remove an entry from the tree, she has to prove to $\mathcal{S}$ that $\mathsf{permanent} = 1$. The following zero-knowledge proof serves this purpose, given that $\mathsf{permanent}$ is encoded in the exponent of the message:

$$PK \left\{ \begin{array}{l} (sk): \\ \quad \log_{g^r} (g^{\mathsf{permanent}} \cdot h^r \cdot g^{-1}) = sk = \log_g h \end{array} \right\}.$$

Naturally, the re-randomization step as well as the shuffle proof step also apply to this ciphertext.

8.2 Computational Complexity

The computational and communication complexity of our constructions, for both the server and the client, is $O((B + G) \log N)$ where N is the number of the entries in the database, B is the block size of the entries in the database, and G is the number of clients that have access to the database. $O(B \log N)$ originates from the ORAM construction and we add $O(G \log N)$ for the access structure. Hence, our solution only adds a small overhead to the standard ORAM complexity. The client-side storage is $O(B \log N)$, while the server has to store $O(BN)$ many data.

8.3 Java Implementation

We implemented the four different versions of GORAM in Java (GORAM with off-the-shelf shuffle proofs and batched shuffle proofs, A-GORAM, and S-GORAM). Furthermore, we also implemented A-GORAM and S-GORAM on Amazon EC2. For zero-knowledge proofs, we build on a library [34] that implements Groth-Sahai proofs [7], which internally relies on jP-BC/PBC [35, 36].

Cryptographic setup. We use MNT curves [37] based on prime-order groups for primes of length 224 bits. This results in 112 bits of security according to different organizations [38]. We deploy AES with 128 bit keys and we instantiate the El Gamal encryption scheme, the RSA signature scheme, and the chameleon hash function with a security parameter of 2048 bits. According to NIST [38], this setup is secure until 2030.

8.4 Experiments

We evaluated the four different implementations. As a first experiment, we measured the computation times on client and server for the read and write operation for the constructions without accountable integrity. We performed these experiments on an Intel Xeon with 8 cores and 2.60GHz in order to show the efficiency gained by using batched shuffle proofs instead of off-the-shelf zero-knowledge proofs of shuffle correctness. We vary different parameters: the database size from 1GB to 1TB, the block size from 4KB to 1MB, the number of clients from 1 to 10, the number of cores from 1 to 8, and for batched shuffle proofs also the number of iterations k from 1 to 128. We fix a bucket size of 4 since Stefanov *et al.* [5] showed that this value is sufficient to prevent buckets from overflowing.

(a) GORAM.
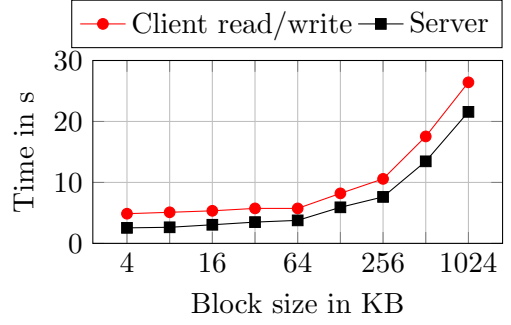
(b) GORAM with batched shuffle proofs and $k=3$.
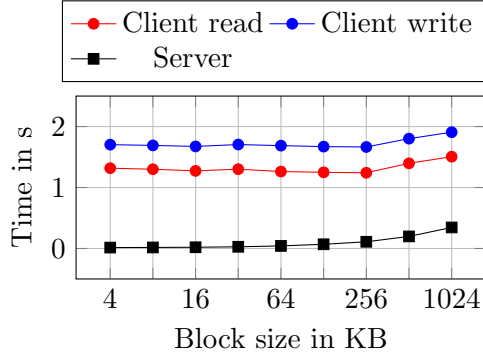
(c) A-GORAM.

(d) S-GORAM.

Figure 9: The average execution time for the read and write protocol on client and server for varying B where $BN = 1\text{GB}$ and $G = 4$.

The second experiment focuses on the solution with accountability. Here we measure also the overhead introduced by our realization with respect to a state-of-the-art ORAM construction, i.e., the price we pay to achieve a wide range of security and privacy properties in a multi-client setting. Another difference from the first experime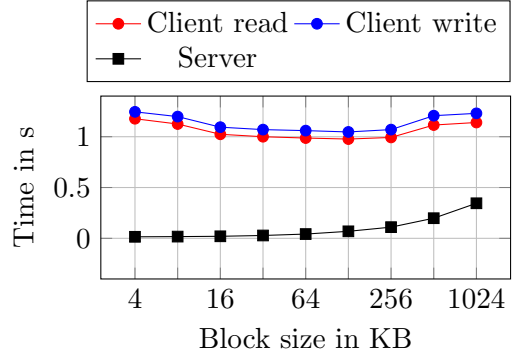nt is the hardware setup. We run the server side of the protocol in Amazon EC2 and the client side on a MacBook Pro with an Intel i7 and 2.90GHz. We vary the parameters as in the previous experiment, except for the number of clients which we vary from 1 to 100 for A-GORAM and from 1 to 10000 for S-GORAM, and the number of cores which are limited to 4. In the experiments where the number of cores is not explicitly varied, we use the maximum number of cores available.

Discussion. The results of the experiments are reported in Figure 9–14 and Table 3. As shown in Figure 9a, varying the block size has a linear effect in the construction without batched shuffle proofs. As expected, the batched shuffle proofs improve the computation time significantly (Figure 9b). The new scheme even seems to be independent of the block size, at least for block sizes less than 64KB. This effect is caused by the parallelization. Still, the homomorphic multiplication of the public-key ciphertexts before the batched shuffle proof computation depends on the block size (line 8.2). Figure 9c and Figure 9d show the results for A-GORAM and S-GORAM. Since the computation time is in practice almost independent of the block size, we can choose larger block sizes in the case of databases with large files, thereby allowing the client to read (resp. write) a file in one shot, as opposed to running multiple read (resp. write) operations. We identify a minimum computation time for 128KB as this is the optimal trade-off between the index map size and the path size. The server computation time is low and varies between 15ms and 345ms, while client operations take less than 2 seconds for A-GORAM and less than
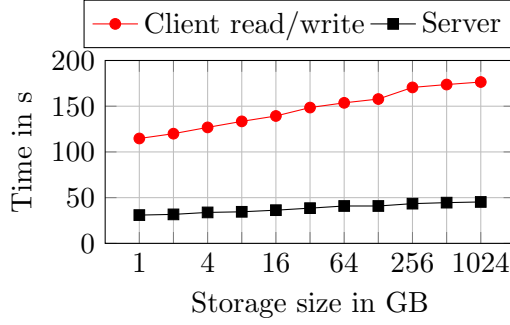
(a) GORAM with $B = 4$KB.
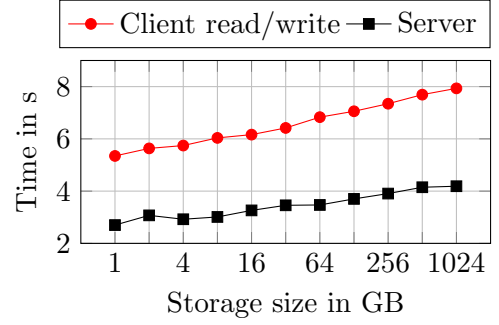
(b) GORAM with batched shuffle proofs, $B = 4$KB, and $k = 4$.
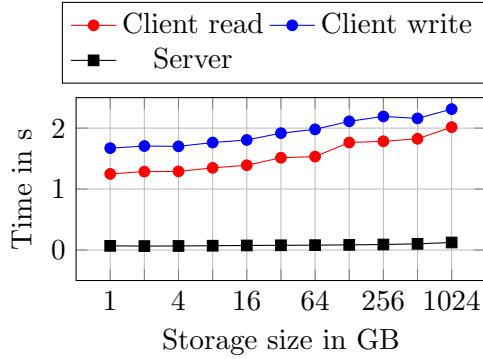
(c) A-GORAM with $B = 128$KB.

(d) S-GORAM with $B = 128$KB.

Figure 10: The average execution time for the read and write protocol on client and server for varying BN where $G = 4$.

1.3 seconds for S-GORAM. As we obtained the best results for 4KB in the experiments for GORAM and 128KB for the others, we use these block sizes in the sequel.

The results obtained by varying the storage size (Figure 10) and the number of clients (Figure 11) prove what the computational complexity suggests. Nevertheless, it is interesting to see the tremendous improvement in computation time between GORAM with and without batched shuffle proofs. The results obtained by varying the iteration time of the batched shuffle proof protocol are depicted in Figure 13 and we verify the expected linear dependency. Smaller values of k are more efficient but higher values give a better soundness probability. If we compare A-GORAM and S-GORAM in Figure 11c and Figure 11d we can see that S-GORAM scales well to a large amount of users as opposed to A-GORAM. The good scaling behavior is due to the used broadcast encryption scheme: it only computes a constant number of pairings independent of the number of users for decryption while the opposite holds for predicate encryption. Nevertheless, we identify a linear growth in the times for S-GORAM, which arises from the linear number of exponentiations that are computed. For instance, in order to write 128KB in a 1GB storage that is used by 100 users, A-GORAM needs about 20 seconds while S-GORAM only needs about 1 second. Even when increasing the number of users to 10000, S-GORAM requires only about 4 seconds, a time that A-GORAM needs for slightly more than 10 users.

Figure 12 shows the results obtained by varying the number of cores. In GORAM most of the computation, especially the zero-knowledge proof computation, can be easily parallelized. We observe this fact in both results (Figure 12a and Figure 12b). In the efficient construction
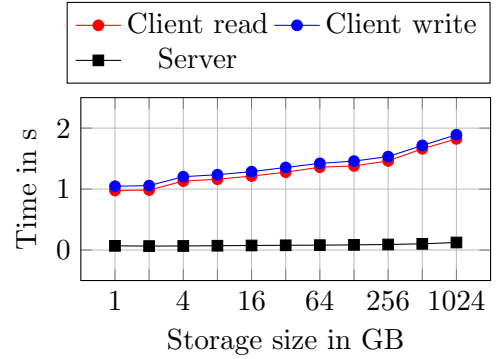
26

(a) GORAM with $B = 4$KB.

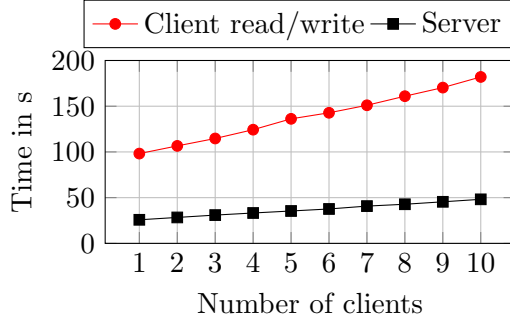(b) GORAM with batched shuffle proofs, $B = 4$KB, and $k = 4$.
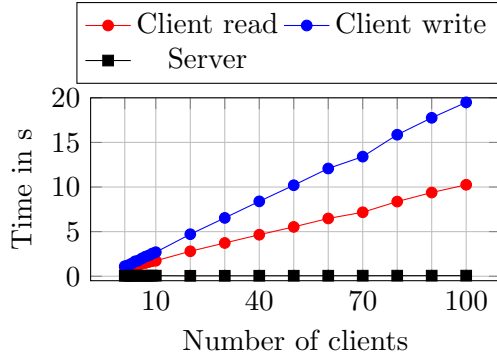
(c) A-GORAM with $B = 128$KB.
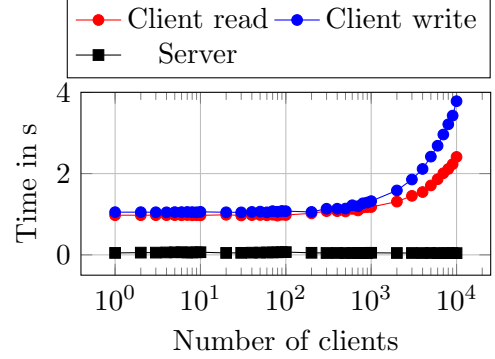
(d) S-GORAM with $B = 128$KB.

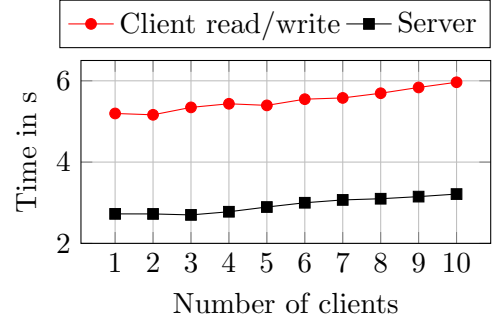Figure 11: The average execution time for the read and write protocol on client and server for varying G where $BN = 1$GB.

we can parallelize the top-level encryption and decryption, the verification of the entries, and the predicate ciphertext decryption. Also in this case parallelization significantly improves the performance (Figure 12c and Figure 12d). Notice that we run the experiments in this case for 20 clients, as opposed to 4 as done for the other constructions, because the predicate ciphertext decryption takes the majority of the computation time and, hence, longer ciphertexts take longer to decrypt and the parallelization effect can be better visualized.

Finally, Table 3 compares S-GORAM with the underlying Path-ORAM protocol. Naturally, since Path-ORAM only uses symmetric encryption, no broadcast encryption, and no verification with chameleon signatures, the computation time is much lower. However, the bottleneck of both constructions is actually the amount of data that has to be downloaded and uploaded by the client (Figure 14). The time required to upload and download data may take much more time than the computation time, given today's bandwidths. Here the overhead is only between 1.02% and 1.05%. For instance, assuming a mobile client using LTE (100Mbit/s downlink and 50Mbit/s uplink in peak) transferring 2 and 50 MB takes 480ms and 12s, respectively. Under these assumptions, considering a block size of 1MB, we get a combined computation and communication overhead of 8% for write and 7% for read, which we consider a relatively low price to pay to get a wide range of security and privacy properties in a multi-client setting.
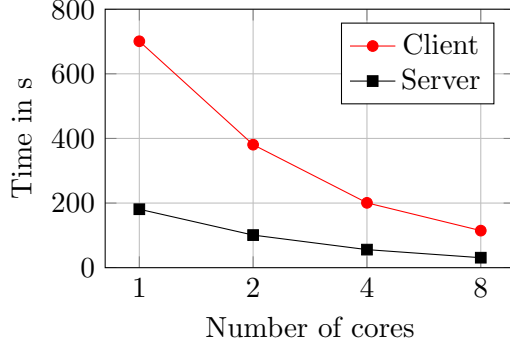
(a) GORAM with $B = 4$KB, and $G = 4$.
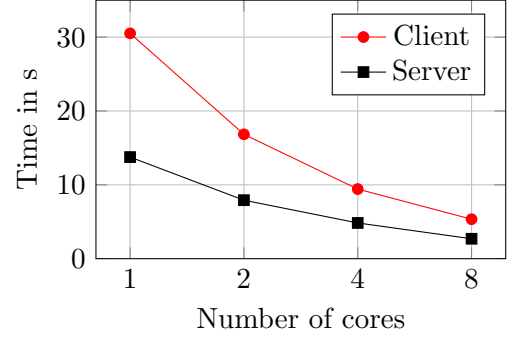
(b) GORAM with batched shuffle proofs, $B = 4$KB, $G = 4$, and $k = 4$.
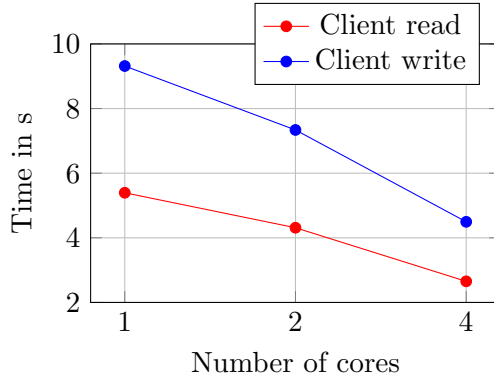
(c) A-GORAM with $BN = 1$GB, $B = 128$KB, and $G = 20$.
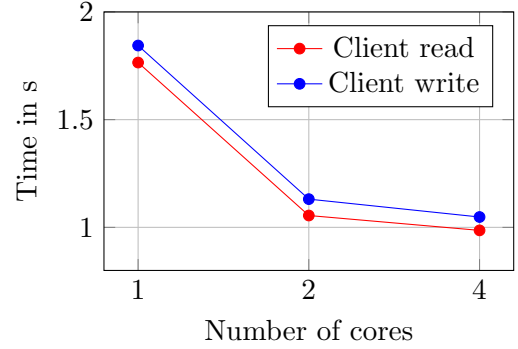
(d) S-GORAM with $BN = 1$GB, $B = 128$KB, and $G = 20$.

Figure 12: The average execution time for the read and write protocol on client and server for a varying number of cores where $BN = 1$GB.

9 Case Study: Personal Health Records

We briefly discuss a potential application of GORAM, namely, a privacy-preserving personal health record (PHR) management system. As the patient should have the control of her own record, the patient is the data owner. The server is some cloud storage provider, which may be chosen by the patient or directly by the state for all citizens (e.g., ELGA in Austria). The healthcare personal (doctors, nurses, pharmacies, and so on) constitutes the clients.

We discuss now possible real-world attacks on PHRs and how the usage of GORAM prevents them. One typical threat is the cloud provider trying to learn customer information (e.g., to sell it or to use it for targeted advertising). For instance, as previously discussed, monitoring the accesses to DNA sequences would allow the service provider to learn the patient's disease: these kinds of attacks are not possible because of obliviousness and data secrecy. Another possible attack could be a pharmacy that tries to increase its profit by changing a prescription

Scheme	Client Read	Client Write	Server
S-GORAM	0.981s	1.075s	0.068s
Path-ORAM	0.042s	0.042s	0.002s

Table 3: Comparison of the computation times between Path-ORAM [5] (single-client!) and S-GORAM on 1GB storage size, 128KB block size and 100 clients.
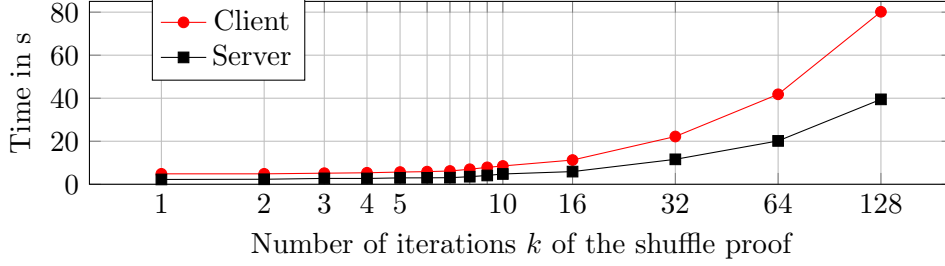
Figure 13: The average execution time for the read and write protocol on client and server for GORAM with batched shuffle proofs and varying k where $BN = 1$GB, $B = 8$KB, and $G = 4$.
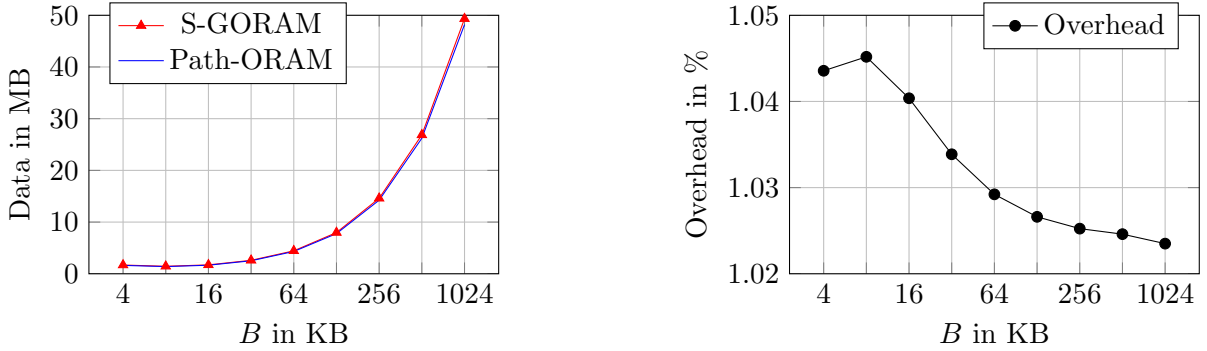
Figure 14: The up-/download amount of data compared between Path-ORAM [5] and S-GORAM for varying B while $BN = 1$GB and $G = 4$.

for a cheap medicine into one that prescribes an expensive medicine. However, in GORAM pharmacies would not have write access to prescriptions, and hence, these cannot be changed or, in A-GORAM, the misbehaving pharmacy can be blamed by the data owner. A common procedure in order to sign a contract with a health insurance is the health check. The patient might want to hide health information from the insurance in order to get a lower fee. To this end, the patient could simply try to drop this information. Dropping of entries in the database is, however, either prevented by making such documents permanent or, in A-GORAM, by letting the insurance, who sees that some documents are missing, blame the patient. Using the backup strategy, the missing documents can be restored.

Finally, we think that GORAM with batched shuffle proofs (even more so A-GORAM and S-GORAM) is a practical solution for the management of today's PHRs, since they are of rather small size. For instance, the data today stored in e-health cards is at most 128KB. The current trend is to store the remaining medical information (e.g., DNA information) on an external server, which can be accessed by using the card. This is exactly our setting, except that we allow for accessing PHRs even without the card, which is crucial in emergency situations. DNA information takes approximately 125MB[6] [39] and all our constructions offer an adequate performance for databases of a few gigabytes, with A-GORAM and S-GORAM performing better for the retrieval of large amounts of data, thanks to the possibility of using larger block sizes.

---

[6]The actual DNA sequence takes about 200GB but one usually shares only the mutations, i.e., the differences of the considered genome to the average human genome. These mutations are only 0.1% of the overall sequence.

10 Related Work

Privacy-preserving outsourced storage. Oblivious RAM (ORAM) [4] is a technique originally devised to protect the access pattern of software on the local memory and thus to prevent the reverse engineering of that software. The observation is that encryption by itself prevents an attacker from learning the content of any memory cell but monitoring how memory is accessed and modified may still leak a great amount of sensitive information. Recent advances in ORAM show that it is efficient enough to hide the data and the user's access pattern in storage outsourcing services [3, 12, 40]–[47].

While a few ORAM constructions guarantee the integrity of user data [48, 49], none of them is suitable to share data with potentially distrustful clients. Goodrich *et al.* [50] studied the problem of multi-client ORAM, but their attacker model does not include malicious, and potentially colluding, clients. Furthermore, their construction does not provide fine-grained access control mechanisms, i.e., either all members of a group have access to a certain data, or none has. Finally, this scheme does not allow the clients to verify the data integrity.

The fundamental problem in existing ORAM constructions is that all clients must have access to the ORAM key, which allows them to read and potentially disrupt the entire database.

A few recent works have started to tackle this problem. Franz *et al.* have introduced the concept of delegated ORAM [51]. The idea is to encrypt and sign each entry with a unique set of keys, initially only known to the data owner: giving a client the decryption key (resp. the decryption and signing keys) suffices to grant read (resp. write) access to that entry. This solution, however, has two drawbacks that undermine its deployment in practice. First, the data owner has to periodically visit the server for checking the validity of the signatures accompanying the data to be inserted in the database (thus tracking the individual client accesses) and reshuffling the ORAM according to the access history in order to enable further unlinkable ORAM accesses. Furthermore, revoking access for a single client requires the data owner to change (and distribute) the capabilities of all other users that have access to that file.

Huang and Goldberg have recently presented a protocol for outsourced private information retrieval [52], which is obtained by layering a private information retrieval (PIR) scheme on top of an ORAM data layout. This solution is efficient and conceals client accesses from the data owner, but it does not give clients the possibility to update data. Moreover, it assumes ℓ non-colluding servers, which is due to the usage of information theoretic multi-server PIR.

De Capitani di Vimercati *et al.* [53] proposed a storage service that uses selective encryption as a means for providing fine-grained access control. The focus of their work is to study how indexing data in the storage can leak information to clients that are not allowed to access these data, although they are allowed to know the indices. The authors do, however, neither consider verifiability nor obliviousness, which distinguishes their storage service from ours.

Finally, there have been a number of works leveraging trusted hardware to realize ORAM schemes [54, 55] including some in the multi-client setting [56, 57]. We, however, intentionally tried to strive for a solution without trusted hardware, only making use of cryptographic primitives.

Verifiable outsourced storage. Verifying the integrity of data outsourced to an untrusted server is a research problem that has recently received increasing attention in the literature. Schröder and Schröder introduced the concept of verifiable data streaming (VDS) and an efficient cryptographic realization thereof [58, 59]. In a verifiable data streaming protocol, a computationally limited client streams a long string to the server, who stores the string in its database in a publicly verifiable manner. The client has also the ability to retrieve and update any element in the database. Papamathou *et al.* [60] proposed a technique, called streaming authenticated data structures, that allows the client to delegate certain computations over

streamed data to an untrusted server and to verify their correctness. Other related approaches are proofs-of-retrievability [61]–[64], which allow the server to prove to the client that it is actually storing all of the client's data, verifiable databases [65], which differ from the previous ones in that the size of the database is fixed during the setup phase, and dynamic provable data possession [66]. All the above do not consider the privacy of outsourced data. While some of the latest work has focused on guaranteeing the confidentiality of the data [67], to the best of our knowledge no existing paper in this line of research takes into account obliviousness.

Personal Health Records. Security and privacy concerns seem to be one of the major obstacles towards the adoption of cloud-based PHRs [68, 69, 70]. Different cloud architectures have been proposed [71], as well as database constructions [72, 73], in order to overcome such concerns. However, none of these works takes into account the threat of a curious storage provider and, in particular, none of them enforces the obliviousness of data accesses.

11 Conclusion and Future Work

This paper introduces the concept of Group ORAM, which captures an unprecedented range of security and privacy properties in the cloud storage setting. The fundamental idea underlying our instantiation is to extend a state-of-the-art ORAM scheme [5] with access control mechanisms and integrity proofs while preserving obliviousness. To tackle the challenge of devising an efficient and scalable construction, we devised a novel zero-knowledge proof technique for shuffle correctness as well as a new accountability technique based on chameleon signatures, both of which are generically applicable and thus of independent interest. We showed how GORAM is an ideal solution for personal record management systems.

As a future work, we intend to relax the assumptions on the server behavior, under which some of the security and privacy properties are proven, developing suitable cryptographic techniques. A further research goal is the design of cryptographic solutions allowing clients to learn only limited information (e.g., statistics) about the dataset.

Acknowledgments

References

[1] M. Maffei, G. Malavolta, M. Reinert, and D. Schröder, "Privacy and Access Control for Outsourced Personal Records," in *Proc. IEEE Symposium on Security & Privacy (S&P'15)*. IEEE Press, 2015.

[2] M. Islam, M. Kuzu, and M. Kantarcioglu, "Access Pattern Disclosure on Searchable Encryption: Ramification, Attack and Mitigation," in *Proc. Annual Network & Distributed System Security Symposium (NDSS'12)*. Internet Society, 2012.

[3] B. Pinkas and T. Reinman, "Oblivious RAM Revisited," in *Proc. Advances in Cryptology (CRYPTO'10)*, ser. LNCS. Springer Verlag, 2010, pp. 502–519.

[4] O. Goldreich and R. Ostrovsky, "Software Protection and Simulation on Oblivious RAMs," *Journal of the ACM*, vol. 43, no. 3, pp. 431–473, May 1996.

[5] E. Stefanov, M. van Dijk, E. Shi, C. Fletcher, L. Ren, X. Yu, and S. Devadas, "Path ORAM: An Extremely Simple Oblivious RAM Protocol," in *Proc. Conference on Computer and Communications Security (CCS'13)*. ACM, 2013.

[6] J. Katz, A. Sahai, and B. Waters, "Predicate Encryption Supporting Disjunctions, Polynomial Equations, and Inner Products," in *Proc. Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT'08)*. Springer Verlag, 2008, pp. 146–162.

[7] J. Groth and A. Sahai, "Efficient Noninteractive Proof Systems for Bilinear Groups," *SIAM Journal on Computing*, vol. 41, no. 5, pp. 1193–1232, 2012.

[8] S. Bayer and J. Groth, "Efficient Zero-Knowledge Argument for Correctness of a Shuffle," in *Proc. Conference on the Theory and Applications of Cryptographic Techniques (EURO-CRYPT'12)*, ser. LNCS. Springer Verlag, 2012, pp. 263–280.

[9] I. E. Akkus, R. Chen, M. Hardt, P. Francis, and J. Gehrke, "Non-tracking Web Analytics," in *Proc. Conference on Computer and Communications Security (CCS'12)*. ACM, 2012, pp. 687–698.

[10] R. Chen, I. E. Akkus, and P. Francis, "SplitX: High-Performance Private Analytics," in *Proc. of the ACM SIGCOMM 2013*. ACM, 2013, pp. 315–326.

[11] S. Goldwasser and S. Micali, "Probabilistic Encryption & How To Play Mental Poker Keeping Secret All Partial Information," in *Proc. ACM Symposium on Theory of Computing (STOC'82)*. ACM, 1982, pp. 365–377.

[12] E. Shi, T.-H. H. Chan, E. Stefanov, and M. Li, "Oblivious RAM With $O((\log n)^3)$ Worst-Case Cost," in *Proc. International Conference on the Theory and Application of Cryptology and Information Security (ASIACRYPT'11)*, ser. LNCS. Springer Verlag, 2011, pp. 197–214.

[13] R. Dingledine, N. Mathewson, and P. Syverson, "Tor: The Second-Generation Onion Router," in *Proc. USENIX Security Symposium (USENIX'04)*. USENIX Association, 2004, pp. 303–320.

[14] M. Maffei, K. Pecina, and M. Reinert, "Security and Privacy by Declarative Design," in *Proc. Symposium on Computer Security Foundations (CSF'13)*. IEEE Press, 2013, pp. 81–96.

[15] M. Backes, S. Lorenz, M. Maffei, and K. Pecina, "Anonymous Webs of Trust," in *Proc. Privacy Enhancing Technologies Symposium (PETS'10)*, ser. LNCS. Springer Verlag, 2010, pp. 130–148.

[16] M. Backes, M. Maffei, and K. Pecina, "Automated Synthesis of Privacy-Preserving Distributed Applications," in *Proc. Annual Network & Distributed System Security Symposium (NDSS'12)*. Internet Society, 2012.

[17] F. Baldimtsi and A. Lysyanskaya, "Anonymous Credentials Light," in *Proc. Conference on Computer and Communications Security (CCS'13)*. ACM, 2013, pp. 1087–1098.

[18] D. L. Chaum, "Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms," *Communications of the ACM*, vol. 24, no. 2, pp. 84–90, 1981.

[19] M. Jakobsson and A. Juels, "Millimix: Mixing in Small Batches," DIMACS, Tech. Rep. 99-33, 1999.

[20] M. Jakobsson, A. Juels, and R. L. Rivest, "Making Mix Nets Robust for Electronic Voting by Randomized Partial Checking," in *Proc. USENIX Security Symposium (USENIX'02)*. USENIX Association, 2002, pp. 339–353.

[21] A. Fiat and A. Shamir, "How to Prove Yourself: Practical Solutions to Identification and Signature Problems," in *Proc. Advances in Cryptology (CRYPTO'86)*. Springer Verlag, 1987, pp. 186–194.

[22] G. Ateniese and B. de Medeiros, "On the Key Exposure Problem in Chameleon Hashes," in *Proc. International Conference on Security in Communication Networks (SCN'04)*, ser. LNCS. Springer Verlag, 2004, pp. 165–179.

[23] J. Camenisch, M. Kohlweiss, and C. Soriente, "An Accumulator Based on Bilinear Maps and Efficient Revocation for Anonymous Credentials," in *Proc. Practice and Theory in Public Key Cryptography (PKC'09)*, ser. LNCS. Springer Verlag, 2009, pp. 481–500.

[24] C. Gentry and B. Waters, "Adaptive Security in Broadcast Encryption Systems (with Short Ciphertexts)," in *Proc. Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT'09)*, ser. LNCS. Springer Verlag, 2009, pp. 171–188.

[25] A. Demers, D. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. Sturgis, D. Swinehart, and D. Terry, "Epidemic Algorithms for Replicated Database Maintenance," in *Proc. Symposium on Principles of Distributed Computing (PODC'87)*. ACM, 1987, pp. 1–12.

[26] R. Küsters, T. Truderung, and A. Vogt, "Accountability: Definition and Relationship to Verifiability," in *Proc. Conference on Computer and Communications Security (CCS'10)*. ACM, 2010, pp. 526–535.

[27] J. Daemen and V. Rijmen, *The Design of Rijndael, AES - The Advanced Encryption Standard*. Springer Verlag, 2002.

[28] Y. Lindell and B. Pinkas, "A Proof of Security of Yao's Protocol for Two-Party Computation," *Journal of Cryptology*, vol. 22, no. 2, pp. 161–188, Apr. 2009.

[29] T. El Gamal, "A Public Key Cryptosystem and a Signature Scheme Based on Discrete Logarithms," in *Proc. Advances in Cryptology (CRYPTO'84)*, ser. LNCS. Springer Verlag, 1985, pp. 10–18.

[30] D. M. Freeman, "Converting Pairing-Based Cryptosystems from Composite-Order Groups to Prime-Order Groups," in *Proc. Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT'10)*, ser. LNCS. Springer Verlag, 2010, pp. 44–61.

[31] C. P. Schnorr, "Efficient Identification and Signatures for Smart Cards," in *Proc. Advances in Cryptology (CRYPTO'89)*, ser. LNCS. Springer Verlag, 1989, pp. 239–252.

[32] R. L. Rivest, A. Shamir, and L. Adleman, "A Method for Obtaining Digital Signatures and Public-key Cryptosystems," *Communications of the ACM*, vol. 21, no. 2, pp. 120–126, Feb. 1978.

[33] R. Cramer, R. Gennaro, and B. Schoenmakers, "A Secure and Optimally Efficient Multi-authority Election Scheme," in *Proc. Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT'97)*, ser. LNCS. Springer Verlag, 1997, pp. 103–118.

[34] J. Backes, S. Lorenz, and K. Pecina, "Zero-knowledge Library," online at github.com/peloba/zk-library.

[35] A. D. Caro, "jPBC - Java Library for Pairing Based Cryptography," online at http://gas.dia.unisa.it/projects/jpbc/.

[36] B. Lynn, "PBC - C Library for Pairing Based Cryptography," online at http://crypto.stanford.edu/pbc/.

[37] A. Miyaji, M. Nakabayashi, and S. Takano, "Characterization of Elliptic Curve Traces under FR-Reduction," in *Proc. International Conference on Information Security and Cryptology (ICISC'00)*, ser. LNCS, vol. 2015. Springer Verlag, 2001, pp. 90–108.

[38] BlueKrypt, "Cryptograhpic Key Length Recommendation," online at www.keylength.com.

[39] R. J. Robinson, "How big is the human genome?" Online at https://medium.com/precision-medicine/how-big-is-the-human-genome-e90caa3409b0.

[40] B. Carbunar and R. Sion, "Regulatory Compliant Oblivious RAM," in *Proc. Applied Cryptography and Network Security (ACNS'10)*, ser. LNCS. Springer Verlag, 2010, pp. 456–474.

[41] M. Ajtai, "Oblivious RAMs Without Cryptographic Assumptions," in *Proc. ACM Symposium on Theory of Computing (STOC'10)*. ACM, 2010, pp. 181–190.

[42] M. T. Goodrich and M. Mitzenmacher, "Privacy-Preserving Access of Outsourced Data via Oblivious RAM Simulation," in *Proc. International Conference on Automata, Languages and Programming (ICALP'11)*, ser. LNCS. Springer Verlag, 2011, pp. 576–587.

[43] I. Damgård, S. Meldgaard, and J. B. Nielsen, "Perfectly Secure Oblivious RAM Without Random Oracles," in *Proc. Theory of Cryptography (TCC'11)*, ser. LNCS. Springer Verlag, 2011, pp. 144–163.

[44] E. Stefanov and E. Shi, "Multi-Cloud Oblivious Storage," in *Proc. Conference on Computer and Communications Security (CCS'13)*. ACM, 2013, pp. 247–258.

[45] M. Maas, E. Love, E. Stefanov, M. Tiwari, E. Shi, K. Asanovic, J. Kubiatowicz, and D. Song, "PHANTOM: Practical Oblivious Computation in a Secure Processor," in *Proc. Conference on Computer and Communications Security (CCS'13)*. ACM, 2013, pp. 311–324.

[46] E. Stefanov and E. Shi, "ObliviStore: High Performance Oblivious Cloud Storage," in *Proc. IEEE Symposium on Security & Privacy (S&P'13)*. IEEE Press, 2013, pp. 253–267.

[47] D. Apon, J. Katz, E. Shi, and A. Thiruvengadam, "Verifiable Oblivious Storage," in *Proc. Practice and Theory in Public Key Cryptography (PKC'14)*, ser. LNCS. Springer Verlag, 2014, pp. 131–148.

[48] P. Williams, R. Sion, and B. Carbunar, "Building Castles out of Mud: Practical Access Pattern Privacy and Correctness on Untrusted Storage," in *Proc. Conference on Computer and Communications Security (CCS'08)*. ACM, 2008, pp. 139–148.

[49] E. Stefanov, E. Shi, and D. Song, "Towards Practical Oblivious RAM," in *Proc. Annual Network & Distributed System Security Symposium (NDSS'12)*. Internet Society, 2012.

[50] M. T. Goodrich, M. Mitzenmacher, O. Ohrimenko, and R. Tamassia, "Privacy-Preserving Group Data Access via Stateless Oblivious RAM Simulation," in *Proc. ACM-SIAM Symposium on Discrete Algorithms (SODA'12)*. Society for Industrial and Applied Mathematics, 2012, pp. 157–167.

[51] M. Franz, C. Carbunar, R. Sion, S. Katzenbeisser, M. Sotakova, P. Williams, and A. Peter, "Oblivious Outsourced Storage with Delegation," in *Proc. Financial Cryptography and Data Security (FC'11)*. Springer Verlag, 2011, pp. 127–140.

[52] Y. Huang and I. Goldberg, "Outsourced Private Information Retrieval with Pricing and Access Control," in *Proc. Annual ACM Workshop on Privacy in the Electronic Society (WPES'13)*. ACM, 2013.

[53] S. De Capitani di Vimercati, S. Foresti, S. Jajodia, S. Paraboschi, and P. Samarati, "Private Data Indexes for Selective Access to Outsourced Data," in *Proc. Annual ACM Workshop on Privacy in the Electronic Society (WPES'11)*. ACM, 2011, pp. 69–80.

[54] M. Backes, A. Kate, M. Maffei, and K. Pecina, "ObliviAd: Provably Secure and Practical Online Behavioral Advertising," in *Proc. IEEE Symposium on Security & Privacy (S&P'12)*. IEEE Press, 2012, pp. 257–271.

[55] A. Kate, M. Maffei, P. Moreno-Sanchez, and K. Pecina, "Privacy Preserving Payments in Credit Networks," in *Proc. Annual Network & Distributed System Security Symposium (NDSS'15)*. Internet Society, 2015.

[56] A. Iliev and S. W. Smith, "Protecting Client Privacy with Trusted Computing at the Server," *IEEE Security and Privacy*, vol. 3, no. 2, pp. 20–28, Mar. 2005.

[57] J. R. Lorch, B. Parno, J. Mickens, M. Raykova, and J. Schiffman, "Shroud: Ensuring Private Access to Large-scale Data in the Data Center," in *Proc. USENIX Conference on File and Storage Technologies (FAST'13)*. USENIX Association, 2013, pp. 199–214.

[58] D. Schröder and H. Schröder, "Verifiable Data Streaming," in *Proc. Conference on Computer and Communications Security (CCS'12)*. ACM, 2012, pp. 953–964.

[59] D. Schröder and M. Simkin, "VeriStream - A Framework for Verifiable Data Streaming," in *Proc. Financial Cryptography and Data Security (FC'15)*. Springer Verlag, 2015.

[60] C. Papamanthou, E. Shi, R. Tamassia, and K. Yi, "Streaming Authenticated Data Structures ," in *Proc. Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT'13)*, 2013.

[61] H. Shacham and B. Waters, "Compact Proofs of Retrievability," in *Proc. International Conference on the Theory and Application of Cryptology and Information Security (ASIACRYPT'08)*, ser. LNCS. Springer Verlag, 2008, pp. 90–107.

[62] D. L. G. Filho and P. S. L. M. Barreto, "Demonstrating Data Possession and Uncheatable Data Transfer," Cryptology ePrint Archive, Report 2006/150, 2006, http://eprint.iacr.org/.

[63] T. Schwarz and E. L. Miller, "Store, Forget, and Check: Using Algebraic Signatures to Check Remotely Administered Storage," 2006.

[64] E. Stefanov, M. van Dijk, A. Oprea, and A. Juels, "Iris: A Scalable Cloud File System with Efficient Integrity Checks," Cryptology ePrint Archive, Report 2011/585, 2011, http://eprint.iacr.org/.

[65] S. Benabbas, R. Gennaro, and Y. Vahlis, "Verifiable Delegation of Computation Over Large Datasets," in *Proc. Advances in Cryptology (CRYPTO'11)*, ser. LNCS. Springer Verlag, 2011, pp. 111–131.

[66] C. Erway, A. Küpçü, C. Papamanthou, and R. Tamassia, "Dynamic Provable Data Possession," in *Proc. Conference on Computer and Communications Security (CCS'09)*. ACM, 2009, pp. 213–222.

[67] M. van Dijk, A. Juels, A. Oprea, R. L. Rivest, E. Stefanov, and N. Triandopoulos, "Hourglass Schemes: How to Prove That Cloud Files Are Encrypted," in *Proc. Conference on Computer and Communications Security (CCS'12)*. ACM, 2012, pp. 265–280.

[68] I. Carrión Señor, L. J. Fernández-Alemán, and A. Toval, "Are Personal Health Records Safe? A Review of Free Web-Accessible Personal Health Record Privacy Policies," *Journal of Medical Internet Research*, vol. 14, no. 4, 2012.

[69] D. Daglish and N. Archer, "Electronic Personal Health Record Systems: A Brief Review of Privacy, Security, and Architectural Issues," *World Congress on Privacy, Security, Trust and the Management of e-Business*, pp. 110–120, 2009.

[70] K. T. Win, W. Susilo, and Y. Mu, "Personal Health Record Systems and Their Security Protection," *Journal of Medical Systems*, vol. 30, no. 4, pp. 309–315, 2006.

[71] H. Löhr, A.-R. Sadeghi, and M. Winandy, "Securing the e-Health Cloud," in *Proc. ACM International Health Informatics Symposium (IHI'10)*. ACM, 2010, pp. 220–229.

[72] M. Li, S. Yu, K. Ren, and W. Lou, "Securing Personal Health Records in Cloud Computing: Patient-Centric and Fine-Grained Data Access Control in Multi-owner Settings," in *SECURECOMM'10*, 2010.

[73] P. Korde, V. Panwar, and S. Kalse, "Securing Personal Health Records in Cloud using Attribute Based Encryption," *International Journal of Engineering and Advanced Technology*, 2013.

[74] S. Goldwasser, S. Micali, and R. L. Rivest, "A Digital Signature Scheme Secure Against Adaptive Chosen-message Attacks," *SIAM Journal on Computing*, vol. 17, no. 2, pp. 281–308, Apr. 1988.

A   Cryptographic Building Blocks

Private key encryption. We recall the basic definition of private-key encryption and the respective IND-CPA security definition.

Definition 2 (Private-Key Encryption). *A private key encryption scheme is a tuple of* PPT *algorithms* $\Pi_{\mathsf{SE}} = (\mathsf{Gen}_{\mathsf{SE}}, \mathcal{E}, \mathcal{D})$, *where the generation algorithm* $\mathsf{Gen}_{\mathsf{SE}}(1^\lambda)$ *outputs a private key* k; *the encryption algorithm* $\mathcal{E}(k, m)$ *takes as input a key* k *and a message* $m \in \mathcal{M}$ *and outputs a ciphertext* c; *the decryption algorithm* $\mathcal{D}(k, c)$ *takes as input a key* k *and a ciphertext* c *and outputs a message* m.

A private key encryption scheme is *correct* if and only if, for all $k \leftarrow \mathsf{Gen}_{\mathsf{SE}}(1^\lambda)$ and all messages $m \in \mathcal{M}$ we have $\mathcal{D}(k, \mathcal{E}(k, m)) = m$.

Next, we define IND-CPA security for private key encryption schemes, where $\mathcal{O}_k(\cdot)$ is an encryption oracle that returns $\mathcal{E}(k, m)$ when queried on a message m.

Definition 3 (IND-CPA Security). *Let $\Pi_{\mathsf{SE}} = (\mathsf{Gen}_{\mathsf{SE}}, \mathcal{E}, \mathcal{D})$ be a private key encryption scheme. Π_{SE} has indistinguishable ciphertexts against chosen-plaintext attacks if for all* PPT *adversaries $\mathcal{A}$ the following probability is negligible (as function in λ):*

$$\left| \Pr[\mathsf{Exp}^{\mathsf{cpa}}_{\mathcal{A},\Pi_{\mathsf{SE}}}(\lambda, 1) = 1] - \Pr[\mathsf{Exp}^{\mathsf{cpa}}_{\mathcal{A},\Pi_{\mathsf{SE}}}(\lambda, 0) = 1] \right|$$

where $\mathsf{Exp}^{\mathsf{cpa}}_{\mathcal{A},\Pi_{\mathsf{SE}}}(\lambda, b)$ is the following experiment:

Experiment $\mathsf{Exp}^{\mathsf{cpa}}_{\mathcal{A},\Pi_{\mathsf{SE}}}(\lambda, b)$
    $k \leftarrow \mathsf{Gen}_{\mathsf{SE}}(1^\lambda)$
    $(m_0, m_1) \leftarrow \mathcal{A}^{\mathcal{O}_k(\cdot)}$
    $c_b \leftarrow \mathcal{E}(k, m_b)$
    $b' \leftarrow \mathcal{A}(c_b)$
    Output 1 if and only if, $|m_0| = |m_1|$ and $b' = b$.

Finally, we introduce the notion of elusive-range scheme, we denote the range of a key k by $\mathsf{Range}_\lambda(k) := \{\mathcal{E}(k, m)\}_{m \in \{0,1\}^\lambda}$.

Definition 4 (Elusive Range [28]). *Let $\Pi_{\mathsf{SE}} = (\mathsf{Gen}_{\mathsf{SE}}, \mathcal{E}, \mathcal{D})$ be a private key encryption scheme. Π_{SE} has elusive-range if for all* PPT *adversaries $\mathcal{A}$ the following probability is negligible in λ:*

$$\Pr_{k \leftarrow \mathsf{Gen}_{\mathsf{SE}}(1^\lambda)}[\mathcal{A}(1^\lambda) \in \mathsf{Range}_\lambda(k)]$$

Public key encryption. We recall the basic definition of public-key encryption and the corresponding IND-CPA security definition.

Definition 5 (Public Key Encryption). *A public key encryption scheme is a tuple of* PPT *algorithms $\Pi_{\mathsf{PKE}} = (\mathsf{Gen}_{\mathsf{PKE}}, \mathsf{Enc}, \mathsf{Dec})$, where the generation algorithm $\mathsf{Gen}_{\mathsf{PKE}}(1^\lambda)$ outputs a public key pk and a private key sk; the encryption algorithm $\mathsf{Enc}(pk, m)$ takes as input the public key pk and a message $m \in \mathcal{M}$ and outputs a ciphertext c; the decryption algorithm $\mathsf{Dec}(sk, c)$ takes as input the secret key sk and a ciphertext c and outputs a message m or $\perp$.*

A public key encryption scheme is *correct* if and only if, for all $(pk, sk) \leftarrow \mathsf{Gen}_{\mathsf{PKE}}(1^\lambda)$ and all messages $m \in \mathcal{M}$ we have $\mathsf{Dec}(sk, \mathsf{Enc}(pk, m)) = m$.

Next, we define IND-CPA security for public-key encryption schemes.

Definition 6 (CPA Security). *Let $\Pi_{\mathsf{PKE}} = (\mathsf{Gen}_{\mathsf{PKE}}, \mathsf{Enc}, \mathsf{Dec})$ be a public key encryption scheme. Π_{PKE} has indistinguishable ciphertexts against chosen-plaintext attacks (CPA) if for all* PPT *adversaries $\mathcal{A}$ the following probability is negligible (in the security parameter λ):*

$$|\Pr[\mathsf{ExpPKE}(\lambda, 1) = 1] - \Pr[\mathsf{ExpPKE}(\lambda, 0) = 1]|$$

where $\mathsf{ExpPKE}(\lambda, b)$ is the following experiment:

Experiment $\mathsf{ExpPKE}(\lambda, b)$
    $(pk, sk) \leftarrow \mathsf{Gen}_{\mathsf{PKE}}(1^\lambda)$
    $(m_0, m_1) \leftarrow \mathcal{A}(pk)$
    $c_b \leftarrow \mathsf{Enc}(pk, m_b)$
    $b' \leftarrow \mathcal{A}(c_b)$
    Output 1 if and only if, $|m_0| = |m_1|$ and $b' = b$.

For our purposes, we need an IND-CPA-secure public-key encryption scheme that is rerandomizable. Hence, we assume that there exists a function $\mathsf{Rerand}(pk, c, r)$ that takes as input a ciphertext c and randomness r and returns a rerandomized ciphertext c' encrypting the same content where $c \neq c'$ and both c and c' have the same distribution in the ciphertext space.

Predicate encryption. We recall the notion of predicate encryption [6]. In a predicate encryption scheme one can encrypt a message m under a certain attribute $I \in \Sigma$ using a master public key mpk where Σ is the universe of all possible attributes. Furthermore, one can decrypt the resulting ciphertext using a secret key sk_f associated with a predicate $f \in \mathcal{F}$ if and only if I fulfills f, i.e., $f(I) = 1$, where $\mathcal{F}$ is the universe of all predicates.

Definition 7 (Predicate Encryption). *A predicate encryption scheme for the universe of predicates and attributes $\mathcal{F}$ and Σ, respectively, is a tuple of PPT algorithms $\Pi_{\mathsf{PE}} = (\mathsf{PrGen}, \mathsf{PrKGen}, \mathsf{PrEnc}, \mathsf{PrDec})$, where the generation algorithm PrGen takes as input a security parameter 1^λ and returns a master public and a master secret key pair (mpk, psk); the key generation algorithm PrKGen takes as input the master secret key psk and a predicate description $f \in \mathcal{F}$ and returns a secret key sk_f associated with f; the encryption algorithm PrEnc takes as input the master public key mpk, an attribute $I \in \Sigma$, and a message m and it returns a ciphertext c; and the decryption algorithm PrDec takes as input a secret key sk_f associated with a predicate f and a ciphertext c and outputs either a message m or $\perp$.*

A predicate encryption scheme Π_{PE} is *correct* if and only if, for all λ, all key pairs $(mpk, psk) \leftarrow \mathsf{PrGen}(1^\lambda)$, all predicates $f \in \mathcal{F}$, all secret keys $sk_f \leftarrow \mathsf{PrKGen}(psk, f)$, and all attributes $I \in \Sigma$ we have that (i) if $f(I) = 1$ then $\mathsf{PrDec}(sk_f, \mathsf{PrEnc}(mpk, I, m)) = m$ and (ii) if $f(I) = 0$ then $\mathsf{PrDec}(sk_f, \mathsf{PrEnc}(mpk, I, m)) = \perp$ except with negligible probability in λ.

Next, we recall the security notion *attribute-hiding* that we require the predicate encryption scheme to hold. Suppose that there are professors, students, and employees at a university with corresponding attributes *Prof*, *Emp*, and *Stud*. Naturally, every member of a group will be equipped with a secret key sk_f such that f is either the predicate $\mathsf{mayAccProf}$, $\mathsf{mayAccEmp}$, or $\mathsf{mayAccStud}$. We use the toy policy that professors may read everything and employees and students may read only encryptions created using *Emp* and *Stud*, respectively. Now, attribute-hiding means the following: let *file* be a file encrypted using the attribute *Prof*. On the one hand, a student equipped with $sk_{\mathsf{mayAccStud}}$ can neither decrypt the file nor tell with which attribute it is encrypted except for that it was not *Stud*. On the other hand, even a professor does not learn under which attribute *file* was encrypted, she only learns the content of the file and nothing more. The following definition formalizes the intuition given above.

Definition 8 (Attribute Hiding). *Let Π_{PE} be a predicate encryption scheme with respect to $\mathcal{F}$ and Σ. Π_{PE} is attribute hiding if for all PPT adversaries $\mathcal{A}$ the following probability is negligible:*

$$|\mathsf{Pr}[\mathsf{ExpPE}(\lambda, 1) = 1] - \mathsf{Pr}[\mathsf{ExpPE}(\lambda, 0) = 1]|$$

where $\mathsf{ExpPE}(\lambda, b)$ is the following experiment:

Experiment $\mathsf{ExpPE}(\lambda, b)$
$\quad \Sigma^2 \ni (I_0, I_1) \leftarrow \mathcal{A}(1^\lambda)$
$\quad (mpk, psk) \leftarrow \mathsf{PrGen}(1^\lambda)$ *and give mpk to $\mathcal{A}$*
$\quad \mathcal{A}$ *may adaptively query keys psk_{f_i} for predicates*
$\quad\quad f_1, \ldots, f_\ell \in \mathcal{F}$ *where $f_i(I_0) = f_i(I_1)$*
$\quad (m_0, m_1) \leftarrow \mathcal{A}$ *such that $|m_0| = |m_1|$ and*
$\quad\quad$ *if there is an i such that $f_i(I_0) = f_i(I_1) = 1$,*
$\quad\quad$ *then $m_0 = m_1$ is required*

$b' \leftarrow \mathcal{A}(\mathsf{PrEnc}(mpk, I_b, m_b))$ *while $\mathcal{A}$ may continue*
   requesting keys for additional predicates
   with the same restrictions as before
output 1 if and only if, $b' = b$.

We use a variant of predicate encryption for our construction. This variant is called *predicate-only encryption*. Predicate-only encryption can be seen as a special variant of functional encryption with the universe of functions only mapping to boolean values.

Definition 9 (Predicate-only Encryption). *A predicate-only encryption scheme for the universe of predicates and attributes $\mathcal{F}$ and Σ, respectively, is a tuple of PPT algorithms $\Pi_{\mathsf{PO}} = (\mathsf{PoGen}, \mathsf{PoKGen}, \mathsf{PoEnc}, \mathsf{PoDec})$, where the generation algorithm PoGen takes as input a security parameter 1^λ and returns a master public and a master secret key pair (opk, osk); the key generation algorithm PoKGen takes as input the master secret key osk and a predicate description $f \in \mathcal{F}$ and returns a secret key osk_f associated with f; the encryption algorithm PoEnc takes as input the master public key opk and an attribute $I \in \Sigma$ and it returns a ciphertext c; and the decryption algorithm PoDec takes as input a secret key osk_f associated with a predicate f and a ciphertext c and outputs either 1 or 0.*

A predicate-only encryption scheme Π_{PO} is correct if and only if, for all λ, all key pairs $(opk, osk) \leftarrow \mathsf{PoGen}(1^\lambda)$, all predicates $f \in \mathcal{F}$, all secret keys $osk_f \leftarrow \mathsf{PoKGen}(osk, f)$, and all attributes $I \in \Sigma$ we have that $\mathsf{PoDec}(osk_f, \mathsf{PoEnc}(opk, I)) = f(I)$ except with negligible probability.

Security is defined similar to predicate encryption, but we refer the interested reader to [6] for a formal security definition.

As for public-key encryption, we require rerandomization operations $\mathsf{PrRR}(mpk, c, r)$ and $\mathsf{PoRR}(opk, c, r)$.

We briefly describe below how to encode the access control matrix through predicates and attributes. We use $\Sigma = \mathcal{F} = \mathbb{Z}_q^{n+1}$ where n is the maximum number of clients that are registered with the database owner. Let $f, I \in \mathbb{Z}_q^{n+1}$ such that $f(I) = 1$ if and only if $\langle f, I \rangle = 0$, i.e., the two vectors f and I are orthogonal. Let $(f_1, \ldots, f_n) \in \mathbb{Z}_q^{(n+1) \times n}$ be the matrix formed of all column-vectors representing the n clients. Let us furthermore assume that all the n columns are pairwise linearly independent. Now, in order to find an attribute that implements the read or write access modes of a data entry at index i for all clients, one computes a vector $I \in \mathbb{Z}_q^{n+1}$ that is orthogonal to the $k \le n$ vectors corresponding to the clients that have access to i and that is not orthogonal to the other $n - k$. Since there are at most n vectors to which I has to be orthogonal, there always exists a solution to this system of equations.

Broadcast encryption. We recall the definition of broadcast encryption and an adaptive security notion [24].

Definition 10 (Broadcast encryption). *A broadcast encryption scheme is a tuple of PPT algorithms $\Pi_{\mathsf{BE}} = (\mathsf{Setup}_{\mathsf{BE}}, \mathsf{BrKeyGen}, \mathsf{BrEnc}, \mathsf{BrDec})$:*

- *the generation algorithm $\mathsf{Setup}_{\mathsf{BE}}(1^\lambda, n)$ takes as input a security parameter λ and a maximum number of users n and outputs a key pair (bsk, bpk);*

- *the key generation algorithm $\mathsf{BrKeyGen}(i, bsk)$ takes as input an index $i \in \{1, \ldots, n\}$ and the secret key bsk and outputs a private key d_i;*

- *the encryption function $\mathsf{BrEnc}(S, bpk)$ takes as input a set $S \subseteq \{1, \ldots, n\}$ and the public key bpk and outputs a pair $\langle Hdr, K \rangle$ where Hdr is called the header and K is called the*

message encryption key. Let Π_{SE} be a symmetric encryption scheme and let $c \leftarrow \mathcal{E}(K, M)$ be the encryption of message M that is supposed to be broadcast to users in S. Then the broadcast message consists of (S, Hdr, c);

- *finally, the decryption function $\mathsf{BrDec}(S, i, d_i, Hdr, bpk)$ takes as input the set of users S, an index i with corresponding private key d_i, the header Hdr, and the public key bpk. If $i \in S$ and d_i belongs to i, then it outputs a message encryption key K that can be used to decrypt the symmetric-key ciphertext c produced during encryption.*

Definition 11 (Adaptive security). *Let Π_{BE} be a broadcast encryption scheme. Π_{BE} is adaptively secure if for all PPT adversaries $\mathcal{A}$ the following probability is negligible:*

$$|\Pr[\mathsf{ExpBE}(\lambda, 1) = 1] - \Pr[\mathsf{ExpBE}(\lambda, 0) = 1]|$$

where $\mathsf{ExpBE}(\lambda, b)$ is the following experiment:

***Experiment** $\mathsf{ExpBE}(\lambda, b)$*
  $(bsk, bpk) \leftarrow \mathsf{Setup}_{\mathsf{BE}}(1^\lambda, n)$ and give bpk to $\mathcal{A}$
  $\mathcal{A}$ may adaptively query private keys d_i for $1 \leq i \leq n$
  $S^* \leftarrow \mathcal{A}$ such that $i \notin S^*$ for all queried i's
  $\langle Hdr^*, K_0 \rangle \leftarrow \mathsf{BrEnc}(S^*, bpk)$
  K_1 from the key space
  $b' \leftarrow \mathcal{A}(Hdr^*, K_b)$
  output 1 if and only if, $b' = b$.

Chameleon hash functions. We recall the definition of chameleon hash functions as well as the notion of key-exposure freeness [22].

Definition 12 (Chameleon hash function). *A chameleon hash function is a tuple of PPT algorithms $\Pi_{\mathsf{CHF}} = (\mathsf{Gen}_{\mathsf{CHF}}, \mathsf{CH}, \mathsf{Col})$, where the generation algorithm $\mathsf{Gen}_{\mathsf{CHF}}(1^\lambda)$ outputs a key pair (cpk, csk); the chameleon hash function $\mathsf{CH}(cpk, m, r)$ takes as input the public key cpk, a message m, and randomness r and it outputs a tag t; the collision function $\mathsf{Col}(csk, m, r, m')$ takes as input the private key csk, a message m, randomness r, and a new message m' and it outputs a new randomness r' such that $\mathsf{CH}(cpk, m, r) = t = \mathsf{CH}(cpk, m', r')$.*

We define next the key-exposure freeness property for chameleon hash functions [22].

Definition 13 (Key-exposure freeness). *Let $\Pi_{\mathsf{CHF}} = (\mathsf{Gen}_{\mathsf{CHF}}, \mathsf{CH}, \mathsf{Col})$ be a chameleon hash function. Π_{CHF} is key-exposure free if for all key pairs (cpk, csk), all messages m, all randomnesses r, and all chameleon hash tags t such that $t = \mathsf{CH}(cpk, m, r)$, no PPT adversary can output a fresh message-randomness pair (m^*, r^*) such that $t = \mathsf{CH}(cpk, m^*, r^*)$ without knowing csk and even after seeing polynomially many collisions (m_i, r_i) for m and r where $m_i \neq m^*$ and $r_i \neq r^*$.*

Digital signatures. We recall the definition of digital signatures as well as the one of existential unforgeability [74].

Definition 14 (Digital signature). *A digital signature scheme is a tuple of PPT algorithms $\Pi_{\mathsf{DS}} = (\mathsf{Gen}_{\mathsf{DS}}, \mathsf{sign}, \mathsf{verify})$ where the generation algorithm $\mathsf{Gen}_{\mathsf{DS}}(1^\lambda)$ outputs a key pair (vk, sk); the signing function $\mathsf{sign}(sk, m)$ takes as input the signing key sk and a message m and it outputs a signature σ; the verification function $\mathsf{verify}(vk, \sigma, m)$ takes as input the verification key vk, a signature σ, and a message m, and it outputs either $\top$ if σ is a valid signature for m or $\top$ otherwise.*

We next define existential unforgeability [74].

Definition 15 (Existential unforgeability)**.** *Let* $\Pi_{\mathsf{DS}} = (\mathsf{Gen}_{\mathsf{DS}}, \mathsf{sign}, \mathsf{verify})$ *be a digital signature scheme.* Π_{DS} *is* existentially unforgeable against chosen message attacks *if for all* PPT *adversaries* $\mathcal{A}$ *the following probability is negligible (as function of λ):*

$$\Pr[\mathsf{Exp}^{\mathsf{eu}}_{\mathcal{A},\Pi_{\mathsf{DS}}}(\lambda) = 1]$$

where $\mathsf{Exp}^{\mathsf{eu}}_{\mathcal{A},\Pi_{\mathsf{DS}}}(\lambda)$ *is the following experiment:*

Experiment $\mathsf{Exp}^{\mathsf{eu}}_{\mathcal{A},\Pi_{\mathsf{DS}}}(\lambda)$
> $(vk, sk) \leftarrow \mathsf{Gen}_{\mathsf{DS}}(\lambda)$
> *give vk to* $\mathcal{A}$
> $\mathcal{A}$ *may adaptively query signatures σ_i for messages m_i*
> $(\sigma^*, m^*) \leftarrow \mathcal{A}$
> *output 1 if and only if,* $\mathsf{verify}(vk, \sigma^*, m^*) = \top$
> > *and for all i we have $m^* \neq m_i$.*

Proofs of shuffle correctness. Zero-knowledge proofs of shuffle correctness (also sometimes known as mix proofs) were first introduced by Chaum [18] in the context of mix networks. More formally, let $C_1, \ldots, C_n$ be a sequence of ciphertexts and $C'_1, \ldots, C'_n$ be a permuted and rerandomized version of thereof. Let furthermore π be the used permutation and $r_1, \ldots, r_n$ be the randomnesses used in the rerandomization. A zero-knowledge proof of shuffle correctness can be expressed as follows:

$$PK \left\{ \begin{array}{l} (\pi, r_1, \ldots, r_n): \\ \quad \forall i.\ C'_i = \mathsf{Rerand}(pk, C_{\pi^{-1}(i)}, r_i) \end{array} \right\}.$$

Notice that this proof reveals the old and the new ciphertext but it hides π and $r_1, \ldots, r_n$.

B Predicate Encryption and Rerandomization

We present the predicate encryption scheme of Katz *et al.* [6]. Moreover, we show how one can rerandomize ciphertexts as a public operation. This might be of independent interest. We show the rerandomization for the original scheme, however, it is straightforward to adapt the transformation to prime order groups [30] – the one we use in our implementation for efficiency reasons.

B.1 The KSW Predicate Encryption Scheme

The scheme is based on composite order groups with a bilinear map. More precisely, let $N = pqr$ be a composite number where p, q, and r are large prime numbers. Let $\mathbb{G}$ be an order-N cyclic group and $e : \mathbb{G} \times \mathbb{G} \to \mathbb{G}_T$ be a bilinear map. Recall that e is *bilinear*, i.e., $e(g^a, g^b) = e(g, g)^{ab}$, and *non-degenerate*, i.e., if $\langle g \rangle = \mathbb{G}$ then $e(g, g) \neq 1$. Then, by the chinese remainder theorem, $\mathbb{G} = \mathbb{G}_p \times \mathbb{G}_q \times \mathbb{G}_r$ where $\mathbb{G}_s$ with $s \in \{p, q, r\}$ are the s-order subgroups of $\mathbb{G}$. Moreover, given a generator g for $\mathbb{G}$, $\langle g^{pq} \rangle = \mathbb{G}_r$, $\langle g^{pr} \rangle = \mathbb{G}_q$, and $\langle g^{qr} \rangle = \mathbb{G}_p$. Another insight is the following, given for instance $a \in \mathbb{G}_p$ and $b \in \mathbb{G}_q$, we have $e(a, b) = e((g^{qr})^c, (g^{pr})^d) = e(g^{rc}, g^d)^{pqr} = 1$, i.e., a pairing of elements from different subgroups cancels out. Finally, let $\mathcal{G}$ be an algorithm that takes as input a security parameter 1^λ and outputs a description $(p, q, r, \mathbb{G}, \mathbb{G}_T, e)$. We describe the algorithms PrGen, PrKGen, PrEnc, and PrDec in the sequel.

Algorithm PoGen$(1^\lambda, n)$ and PrGen$(1^\lambda, n)$. First, the algorithm runs $\mathcal{G}(1^\lambda)$ to obtain $(p, q, r, \mathbb{G}, \mathbb{G}_T, e)$ with $\mathbb{G} = \mathbb{G}_p \times \mathbb{G}_q \times \mathbb{G}_r$. Then, it computes g_p, g_q, and g_r as generators of $\mathbb{G}_p$, $\mathbb{G}_q$, and $\mathbb{G}_r$, respectively. The algorithm selects $R_0 \in \mathbb{G}_r$, $R_{1,i}, R_{2,i} \in \mathbb{G}_r$ and $h_{1,i}, h_{2,i} \in \mathbb{G}_p$ uniformly at random for $1 \leq i \leq n$. $(N = pqr, \mathbb{G}, \mathbb{G}_T, e)$ constitutes the public parameters. The public key for the predicate-only encryption scheme is

$$opk = (g_p, g_r, Q = g_q \cdot R_0,$$
$$\{H_{1,i} = h_{1,i} \cdot R_{1,i}, H_{2,i} = h_{2,i} \cdot R_{2,i}\}_{i=1}^n)$$

and the master secret key is

$$osk = (p, q, r, g_q, \{h_{1,i}, h_{2,i}\}_{i=1}^n).$$

For the predicate encryption with messages, the algorithm additionally chooses $\gamma \in \mathbb{Z}_N$ and $h \in \mathbb{G}_p$ at random. The public key is

$$mpk = (g_p, g_r, Q = g_q \cdot R_0, P = e(g_p, h)^\gamma,$$
$$\{H_{1,i} = h_{1,i} \cdot R_{1,i}, H_{2,i} = h_{2,i} \cdot R_{2,i}\}_{i=1}^n)$$

and the master secret key is

$$psk = (p, q, r, g_q, h^{-\gamma}, \{h_{1,i}, h_{2,i}\}_{i=1}^n).$$

Algorithm PoKGen$(osk, \vec{v})$ and PrKGen$(psk, \vec{v})$. Parse $\vec{v}$ as $(v_1, \ldots, v_n)$ where $v_i \in \mathbb{Z}_N$. The algorithm picks random $r_{1,i}, r_{2,i} \in \mathbb{Z}_p$ for $1 \leq i \leq n$, random $R_5 \in \mathbb{G}_r$, random $f_1, f_2 \in \mathbb{Z}_q$, and random $Q_6 \in \mathbb{G}_q$. For the predicate-only encryption scheme it outputs a secret key

$$osk_{\vec{v}} = \begin{pmatrix} K_0 = R_5 \cdot Q_6 \cdot \prod_{i=1}^n h_{1,i}^{-r_{1,i}} \cdot h_{2,i}^{-r_{2,i}}, \\ \{K_{1,i} = g_p^{r_{1,i}} \cdot g_q^{f_1 \cdot v_i}, \\ K_{2,i} = g_p^{r_{2,i}} \cdot g_q^{f_2 \cdot v_i}\}_{i=1}^n \end{pmatrix}.$$

For the predicate encryption scheme with messages, the secret key $sk_{\vec{v}}$ is the same as $osk_{\vec{v}}$ except for

$$K_0 = R_5 \cdot Q_6 \cdot h^{-\gamma} \cdot \prod_{i=1}^n h_{1,i}^{-r_{1,i}} \cdot h_{2,i}^{-r_{2,i}}.$$

Algorithm PoEnc$(opk, \vec{x})$ and PrEnc$(mpk, \vec{x}, m)$. Parse $\vec{x}$ as $(x_1, \ldots, x_n)$ where $x_i \in \mathbb{Z}_N$. The algorithm picks random $s, \alpha, \beta \in \mathbb{Z}_N$ and random $R_{3,i}, R_{4,i} \in \mathbb{G}_r$ for $1 \leq i \leq n$. For the predicate-only encryption scheme it outputs the ciphertext

$$C = \begin{pmatrix} C_0 = g_p^s, \{C_{1,i} = H_{1,i}^s \cdot Q^{\alpha \cdot x_i} \cdot R_{3,i}, \\ C_{2,i} = H_{2,i}^s \cdot Q^{\beta \cdot x_i} \cdot R_{4,i}\}_{i=0}^n \end{pmatrix}.$$

For the predicate encryption scheme with messages notice that $m \in \mathbb{G}_T$. The ciphertext is

$$C = \begin{pmatrix} C' = m \cdot P^s, C_0 = g_p^s, \\ \{C_{1,i} = H_{1,i}^s \cdot Q^{\alpha \cdot x_i} \cdot R_{3,i}, \\ C_{2,i} = H_{2,i}^s \cdot Q^{\beta \cdot x_i} \cdot R_{4,i}\}_{i=0}^n \end{pmatrix}.$$

Algorithm $\mathsf{PoDec}(osk_{\vec{v}}, C)$ **and** $\mathsf{PrDec}(sk_{\vec{v}}, C)$**.** The predicate-only encryption outputs whether the following equation is equal to 1

$$e(C_0, K_0) \cdot \prod_{i=1}^{n} e(C_{1,i}, K_{1,i}) \cdot e(C_{2,i}, K_{2,i}).$$

The predicate encryption scheme with messages outputs the result of the following equation

$$C' \cdot e(C_0, K_0) \cdot \prod_{i=1}^{n} e(C_{1,i}, K_{1,i}) \cdot e(C_{2,i}, K_{2,i}).$$

Correctness. We start with the predicate-only encryption. Assume a secret key $osk_{\vec{v}} = (K_0, \{K_{1,i}, K_{2,i}\}_{i=0}^{n})$ and a ciphertext $C = (C_0, \{C_{1,i}, C_{2,i}\}_{i=0}^{n})$. Then

$$e(C_0, K_0) \cdot \prod_{i=0}^{n} e(C_{1,i}, K_{1,i}) \cdot e(C_{2,i}, K_{2,i})$$

$$= \; e(g_p^s, R_5 \cdot Q_6 \cdot \prod_{i=1}^{n} h_{1,i}^{-r_{1,i}} \cdot h_{2,i}^{-r_{2,i}})$$

$$\cdot \prod_{i=1}^{n} e(H_{1,i}^s \cdot Q^{\alpha \cdot x_i} \cdot R_{3,i}, g_p^{r_{1,i}} \cdot g_q^{f_1 \cdot v_i})$$

$$\cdot e(H_{2,i}^s \cdot Q^{\beta \cdot x_i} \cdot R_{4,i}, g_p^{r_{2,i}} \cdot g_q^{f_2 \cdot v_i})$$

$$= \; \prod_{i=1}^{n} e(g_p, h_{1,i})^{-s \cdot r_{1,i}} \cdot e(g_p, h_{2,i})^{-s \cdot r_{2,i}}$$

$$\cdot \prod_{i=1}^{n} e(h_{1,i}, g_p)^{s \cdot r_{1,i}} \cdot e(g_q, g_q)^{\alpha \cdot x_i \cdot f_1 \cdot v_i}$$

$$\cdot e(h_{2,i}, g_p)^{s \cdot r_{2,i}} \cdot e(g_q, g_q)^{\beta \cdot x_i \cdot f_2 \cdot v_i}$$

$$= \; \prod_{i=1}^{n} e(g_q, g_q)^{(\alpha f_1 + \beta f_2) \cdot x_i \cdot v_i}$$

$$= \; e(g_q, g_q)^{(\alpha f_1 + \beta f_2) \cdot \langle \vec{x}, \vec{v} \rangle}$$

The last equation is 1 if and only if $\langle \vec{x}, \vec{v} \rangle = 0$, as expected.

For the predicate encryption scheme with messages, we have $sk_{\vec{v}} = (K_0, \{K_{1,i}, K_{2,i}\}_{i=0}^{n})$ and a ciphertext $C = (C', C_0, \{C_{1,i}, C_{2,i}\}_{i=0}^{n})$. Then

$$C' \cdot e(C_0, K_0) \cdot \prod_{i=0}^{n} e(C_{1,i}, K_{1,i}) \cdot e(C_{2,i}, K_{2,i})$$

$$= \; m \cdot P^s \cdot e(g_p^s, R_5 \cdot Q_6 \cdot h^{-\gamma} \cdot \prod_{i=1}^{n} h_{1,i}^{-r_{1,i}} \cdot h_{2,i}^{-r_{2,i}})$$

$$\cdot \prod_{i=1}^{n} e(H_{1,i}^s \cdot Q^{\alpha \cdot x_i} \cdot R_{3,i}, g_p^{r_{1,i}} \cdot g_q^{f_1 \cdot v_i})$$

$$\cdot e(H_{2,i}^s \cdot Q^{\beta \cdot x_i} \cdot R_{4,i}, g_p^{r_{2,i}} \cdot g_q^{f_2 \cdot v_i})$$

$$= \; m \cdot e(g_p, h)^{s \cdot \gamma} \cdot e(g_p, h)^{-s \cdot \gamma}$$

43

$$\cdot \prod_{i=1}^{n} e(g_p, h_{1,i})^{-s \cdot r_{1,i}} \cdot e(g_p, h_{2,i})^{-s \cdot r_{2,i}}$$

$$\cdot \prod_{i=1}^{n} e(h_{1,i}, g_p)^{s \cdot r_{1,i}} \cdot e(g_q, g_q)^{\alpha \cdot x_i \cdot f_1 \cdot v_i}$$

$$\cdot e(h_{2,i}, g_p)^{s \cdot r_{2,i}} \cdot e(g_q, g_q)^{\beta \cdot x_i \cdot f_2 \cdot v_i}$$

$$= \quad m \cdot \prod_{i=1}^{n} e(g_q, g_q)^{(\alpha f_1 + \beta f_2) \cdot x_i \cdot v_i}$$

$$= \quad m \cdot e(g_q, g_q)^{(\alpha f_1 + \beta f_2) \cdot \langle \vec{x}, \vec{v} \rangle}$$

The second factor in the last line cancels out only if $\langle \vec{x}, \vec{v} \rangle = 0$, as expected.

B.2 Rerandomizing KSW Ciphertexts

The correctness validation in the previous section already suggests that rerandomization of Π_{PO} and Π_{PE} ciphertexts is possible since all terms that involve randomness s cancel out in the end. As terms including s only occur in the ciphertext we can easily have public rerandomization functions PoRR and PrRR as follows.

Algorithms PoRR(C) and PrRR(C). The algorithm picks fresh randomness $s' \in \mathbb{Z}_N$ and computes C_R in the predicate-only encryption scheme as

$$C_R = (C_0 \cdot g_p^{s'}, \{C_{1,i} \cdot H_{1,i}^{s'}, C_{2,i} \cdot H_{2,i}^{s'}\}_{i=1}^{n}).$$

In the predicate encryption scheme with messages it returns

$$C_R = (C' \cdot P^{s'}, C_0 \cdot g_p^{s'}, \{C_{1,i} \cdot H_{1,i}^{s'}, C_{2,i} \cdot H_{2,i}^{s'}\}_{i=1}^{n}).$$

This transformation constitutes an additive randomization in the sense that in every exponent where s occurs, it now contains exponent $s + s'$. Therefore, also the correctness is preserved.

B.3 Proving Knowledge of Secret Keys in Groth-Sahai

In our construction, the client has to prove to the server that she is eligible to write an entry whenever she wants to replace an entry in the database. The proof (see line 6.11) has the general form

$$PK \{(osk_f) : \; \mathsf{PoDec}(osk_f, c_{\mathsf{Auth}}) = 1\}.$$

In our instantiation where we use KSW predicate-only encryption for regulating the write access, c_{Auth} is of the form

$$(C_0, \{C_{1,i}\}_{1 \leq i \leq n}, \{C_{2,i}\}_{1 \leq i \leq n})$$

and secret keys are of the form

$$(K_0, \{K_{1,i}\}_{1 \leq i \leq n}, \{K_{2,i}\}_{1 \leq i \leq n}).$$

This means that the concrete proof is of the form

$$PK \left\{ \begin{array}{l} (K_0, \{K_{1,i}\}_{1 \leq i \leq n}, \{K_{2,i}\}_{1 \leq i \leq n}) : \\ \quad e(C_0, K_0) \cdot \\ \quad \prod_{i=1}^{n} e(C_{1,i}, K_{1,i}) e(C_{2,i}, K_{2,i}) = 1 \end{array} \right\}.$$

The Groth-Sahai proof system [7] allows for proving relations of the above form. More precisely, given vectors of witnesses $\vec{X} \in \mathbb{G}_1^m$, $\vec{Y} \in \mathbb{G}_2^n$, we can prove the following equality while disclosing neither $\vec{X}$ nor $\vec{Y}$:

$$\prod_{i=1}^{n} e(A_i, Y_i) \cdot \prod_{i=1}^{m} e(X_i, B_i) \cdot \prod_{i=1}^{n} \prod_{j=1}^{m} e(X_i, Y_j)^{\gamma_{ij}} = t_T$$

where $\vec{A} \in \mathbb{G}_1^n$, $\vec{B} \in \mathbb{G}_2^m$, $\Gamma \in \mathbb{Z}_q^{n \times m}$, and $t_T \in \mathbb{G}_T$ are the public components of the proof. In our special case, it is sufficient to consider the following special form of this equation since we only want to keep the secret key hidden, which is in $\mathbb{G}_2$:

$$\prod_{i=1}^{n} e(A_i, Y_i) = t_T.$$

Furthermore, $t_T = 1$ where 1 stands for the neutral element of the group operation. We construct the vectors $\vec{A}$ and $\vec{Y}$ as

$$\vec{A} = (C_0, C_{1,i}, \ldots, C_{1,n}, C_{2,1}, \ldots, C_{2,n})^\top \qquad \vec{Y} = (K_0, K_{1,i}, \ldots, K_{1,n}, K_{2,1}, \ldots, K_{2,n})^\top$$

We do not review the proof construction here but refer the interested reader to [7] for a concise explanation. We observe that since $t_T = 1$, the proofs for our scenario are indeed zero-knowledge.

C Formal Definitions

C.1 Secrecy

Definition 16 (Secrecy). *A Group ORAM* GORAM $=$ (gen, addCl, addE, chMode, read, write) *preserves secrecy, if for every* PPT *adversary* $\mathcal{A}$ *the following probability is negligible in the security parameter* λ:

$$\left| \Pr[\mathsf{ExpSec}_{\mathsf{GORAM}}^{\mathcal{A}}(\lambda, 1) = 1] - \Pr[\mathsf{ExpSec}_{\mathsf{GORAM}}^{\mathcal{A}}(\lambda, 0) = 1] \right|$$

where $\mathsf{ExpSec}_{\mathsf{GORAM}}^{\mathcal{A}}(\lambda, b)$ *is the following game:*
Setup: *The challenger runs* $(cap_{\mathcal{O}}, \mathsf{DB}) \leftarrow \mathsf{gen}(1^\lambda)$, *sets* $\mathbf{AC} := []$, *and runs a black-box simulation of* $\mathcal{A}$ *to which it hands over* DB.
Queries: *The challenger provides* $\mathcal{A}$ *with interactive interfaces* addCl, addE, chMode, read, write, *and* corCl *that* $\mathcal{A}$ *may query adaptively and in any order. Each round* $\mathcal{A}$ *can query exactly one interface. These interfaces are described below:*

(1) On input addCl(**a**) *by* $\mathcal{A}$, *the challenger executes* addCl($cap_{\mathcal{O}}$, **a**) *locally and stores the capability* cap_i *returned by the algorithm.*

(2) On input addE(**a**, d) *by* $\mathcal{A}$, *the challenger executes* $\langle \mathcal{C}_{\mathsf{addE}}(cap_{\mathcal{O}}, \mathbf{a}, d), \mathcal{S}_{\mathsf{addE}}(\mathsf{DB}) \rangle$ *in interaction with* $\mathcal{A}$, *where the former plays the role of the client while the latter plays the role of the server.*

(3) On input chMode(**a**, j) *by* $\mathcal{A}$, *the challenger executes* $\langle \mathcal{C}_{\mathsf{chMode}}(cap_{\mathcal{O}}, \mathbf{a}, j), \mathcal{S}_{\mathsf{chMode}}(\mathsf{DB}) \rangle$ *in interaction with* $\mathcal{A}$.

(4) On input corCl(i) *by* $\mathcal{A}$, *the challenger hands over the capability* cap_i *related to the i-th client in the access control matrix* $\mathbf{AC}$.

(5) *On input* read(i, j) *by* $\mathcal{A}$, *the challenger executes* $\langle \mathcal{C}_{\mathsf{read}}(cap_i, j), \mathcal{S}_{\mathsf{read}}(\mathsf{DB}) \rangle$ *in interaction with* $\mathcal{A}$.

(6) *On input* write(i, j, d) *by* $\mathcal{A}$, *the challenger executes* $\langle \mathcal{C}_{\mathsf{write}}(cap_i, j, d), \mathcal{S}_{\mathsf{write}}(\mathsf{DB}) \rangle$ *in interaction with* $\mathcal{A}$.

Challenge: *Finally,* $\mathcal{A}$ *outputs* $(j, (d_0, d_1))$, *where* j *is an index denoting the database entry on which* $\mathcal{A}$ *wants to be challenged and* (d_0, d_1) *is a pair of entries such that* $|d_0| = |d_1|$. *The challenger accepts the request only if* $\mathbf{AC}(i, j) = \bot$, *for every* i *corrupted by* $\mathcal{A}$ *in the query phase. Afterwards it invokes* $\langle \mathcal{C}_{\mathsf{write}}(cap_{\mathcal{O}}, j, d_b), \mathcal{S}_{\mathsf{write}}(\mathsf{DB}) \rangle$ *in interaction with* $\mathcal{A}$.

Output: *In the output phase* $\mathcal{A}$ *still has access to the interfaces except for* addCl *on input* $\mathbf{a}$ *such that* $\mathbf{a}(j) \neq \bot$; corCl *on input* i *such that* $\mathbf{AC}(i, j) \neq \bot$; *and* chMode *on input* $\mathbf{a}, i$ *with* $\mathbf{a}(i) \neq \bot$ *for some previously corrupted client* i. *Eventually,* $\mathcal{A}$ *stops, outputting a bit* b'. *The challenger outputs 1 if and only if* $b = b'$.

C.2 Integrity

Definition 17 (Integrity). *A Group ORAM* GORAM $=$ (gen, addCl, addE, chMode, read, write) *preserves integrity, if for every* PPT *adversary* $\mathcal{A}$ *the following probability is negligible in the security parameter:*

$$\Pr[\mathsf{ExpInt}_{\mathsf{GORAM}}^{\mathcal{A}}(\lambda) = 1]$$

where $\mathsf{ExpInt}_{\mathsf{GORAM}}^{\mathcal{A}}(\lambda)$ *is the following game:*

Setup: *The challenger runs* $(cap_{\mathcal{O}}, \mathsf{DB}) \leftarrow \mathsf{gen}(1^{\lambda})$, *sets* $\mathbf{AC} := []$, *and runs a black-box simulation of* $\mathcal{A}$. *Furthermore, the challenger initializes a second database* $\mathsf{DB}' := []$ *which is managed locally.*

Queries: *The challenger provides* $\mathcal{A}$ *with the same interfaces as in* Definition 16, *which* $\mathcal{A}$ *may query adaptively and in any order. Since* DB *is maintained on the challenger's side, the queries to* addE, chMode, read *and* write *are locally executed by the challenger. Furthermore, the challenger updates* DB' *locally for all affecting interface calls.*

Challenge: *Finally, the adversary outputs an index* j^* *which he wants to be challenged on. If there exists a capability* cap_i *provided to* $\mathcal{A}$ *with* $\mathbf{AC}(i, j^*) = rw$, *the challenger aborts. Otherwise it runs* $d^* \leftarrow \langle \mathcal{C}_{\mathsf{read}}(cap_{\mathcal{O}}, j^*), \mathcal{S}_{\mathsf{read}}(\mathsf{DB}) \rangle$ *locally.*

Output: *It outputs 1 if and only if* $d^* \neq \mathsf{DB}'(j^*)$.

C.3 Tamper Resistance

Definition 18 (Tamper resistance). *A Group ORAM* GORAM $=$ (gen, addCl, addE, chMode, read, write) *is* tamper-resistant, *if for every* PPT *adversary* $\mathcal{A}$ *the following probability is negligible in the security parameter:*

$$\Pr[\mathsf{ExpTamRes}_{\mathsf{GORAM}}^{\mathcal{A}}(\lambda) = 1]$$

where $\mathsf{ExpTamRes}_{\mathsf{GORAM}}^{\mathcal{A}}(\lambda)$ *is the following game:*

Setup: *The challenger runs the Setup phase as in* Definition 16. *Furthermore, it forwards* DB *to* $\mathcal{A}$ *and initializes a second database* DB' *which is managed locally.*

Queries: *The challenger provides* $\mathcal{A}$ *with the same interfaces as in* Definition 16, *which* $\mathcal{A}$ *may query adaptively and in any order. Furthermore, the challenger updates* DB' *locally for all affecting interface calls.*

Challenge: *Finally, the adversary outputs an index* j^* *which he wants to be challenged on. If there exists a capability* cap_i *that has ever been provided to* $\mathcal{A}$ *such that* $\mathbf{AC}(i, j^*) = rw$, *then*

the challenger aborts. The challenger runs $d^* \leftarrow \langle \mathcal{C}_{\mathsf{read}}(cap_{\mathcal{O}}, j^*), \mathcal{S}_{\mathsf{read}}(\mathsf{DB}) \rangle$ in interaction with $\mathcal{A}$.

Output: It outputs 1 if and only if $d^* \neq \mathsf{DB}'(j^*)$.

C.4 Obliviousness

Definition 19 (Obliviousness)**.** *A Group ORAM* $\mathsf{GORAM} = (\mathsf{gen}, \mathsf{addCl}, \mathsf{addE}, \mathsf{chMode}, \mathsf{read}, \mathsf{write})$ *is* obliviousness, *if for every* PPT *adversary* $\mathcal{A}$ *the following probability is negligible in the security parameter:*

$$\big| \Pr[\mathsf{ExpObv}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda, 1) = 1] - \\ \Pr[\mathsf{ExpObv}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda, 0) = 1] \big|$$

where $\mathsf{ExpObv}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda, b)$ *is the following game:*
Setup: *The challenger runs* $(cap_{\mathcal{O}}, \mathsf{DB}) \leftarrow \mathsf{gen}(1^{\lambda})$ *as in Definition 16 and it forwards* DB *to* $\mathcal{A}$.
Queries: *The challenger provides* $\mathcal{A}$ *with the same interfaces as in Definition 16 except* corCl, *which* $\mathcal{A}$ *may query adaptively and in any order. Furthermore,* $\mathcal{A}$ *is provided with the following additional interface:*

(1) *On input* $\mathsf{query}(\{(i_0, j_0), (i_0, j_0, d_0)\}, \{(i_1, j_1), (i_1, j_1, d_1)\})$ *by* $\mathcal{A}$, *the challenger checks whether* $j_0 \leq |\mathsf{DB}|$, $j_1 \leq |\mathsf{DB}|$, *and* i_0, i_1 *are valid clients. Furthermore, it checks that the operations requested by* $\mathcal{A}$ *are allowed by* **AC**. *If not it aborts. Otherwise it executes* $\langle \mathcal{C}_{\mathsf{read}}(cap_{i_b}, j_b), \mathcal{S}_{\mathsf{read}}(\mathsf{DB}) \rangle$ *or* $\langle \mathcal{C}_{\mathsf{write}}(cap_{i_b}, j_b, d_b), \mathcal{S}_{\mathsf{write}}(\mathsf{DB}) \rangle$ *depending on the input, in interaction with* $\mathcal{A}$. *Here the challenger plays the role of the client and* $\mathcal{A}$ *plays the role of the server.*

Output: *Finally,* $\mathcal{A}$ *outputs a bit* b'. *The challenger outputs 1 if and only if* $b = b'$.

C.5 Anonymity

Definition 20 (Anonymity)**.** *A Group ORAM* $\mathsf{GORAM} = (\mathsf{gen}, \mathsf{addCl}, \mathsf{addE}, \mathsf{chMode}, \mathsf{read}, \mathsf{write})$ *is* anonymity-preserving, *if for every* PPT *adversary* $\mathcal{A}$ *the following probability is negligible in the security parameter:*

$$\big| \Pr[\mathsf{ExpAnon}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda, 1) = 1] - \\ \Pr[\mathsf{ExpAnon}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda, 0) = 1] \big|$$

where $\mathsf{ExpAnon}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda, b)$ *is the following game:*
Setup: *The challenger runs* $(cap_{\mathcal{O}}, \mathsf{DB}) \leftarrow \mathsf{gen}(1^{\lambda})$ *and it forwards* $cap_{\mathcal{O}}$ *and* DB *to* $\mathcal{A}$.
Queries: *The challenger provides* $\mathcal{A}$ *with* read *and a* write *interactive interfaces that* $\mathcal{A}$ *may query adaptively and in any order. Each round* $\mathcal{A}$ *can query exactly one interface. The interfaces are described below:*

(3) *On input* $\mathsf{read}(cap_i, j)$ *by* $\mathcal{A}$, *the challenger executes* $\langle \mathcal{C}_{\mathsf{read}}(cap_i, j), \mathcal{S}_{\mathsf{read}}(\mathsf{DB}) \rangle$ *in interaction with* $\mathcal{A}$, *where the former plays the role of the server and the latter plays the role of the client.*

(4) *On input* $\mathsf{write}(cap_i, j, d)$ *by* $\mathcal{A}$, *the challenger executes* $\langle \mathcal{C}_{\mathsf{write}}(cap_i, j, d), \mathcal{S}_{\mathsf{write}}(\mathsf{DB}) \rangle$ *in interaction with* $\mathcal{A}$, *where the former plays the role of the server and the latter plays the role of the client.*

Challenge: $\mathcal{A}$ outputs $((cap_{i_0}, cap_{i_1}), \{j, (j, d)\})$, where (cap_{i_0}, cap_{i_1}) is a pair of capabilities, j is an index denoting the database entry on which $\mathcal{A}$ wishes to be challenged, and d is some data. The challenger checks whether $\mathbf{AC}(i_0, j) = \mathbf{AC}(i_1, j)$: if not, then it aborts, otherwise it executes $\langle \mathcal{C}_{\mathsf{read}}(cap_{i_b}, j), \mathcal{S}_{\mathsf{read}}(\mathsf{DB})\rangle$ or $\langle \mathcal{C}_{\mathsf{write}}(cap_{i_b}, j, d), \mathcal{S}_{\mathsf{write}}(\mathsf{DB})\rangle$ in interaction with $\mathcal{A}$.
Output: Finally, $\mathcal{A}$ outputs a bit b'. The challenger outputs 1 if and only if $b = b'$.

C.6 Accountability

Definition 21 (Accountability). *A Group ORAM* $\mathsf{GORAM} = (\mathsf{gen}, \mathsf{addCl}, \mathsf{addE}, \mathsf{chMode}, \mathsf{read}, \mathsf{write}, \mathsf{blame})$ *is accountable, if for every* PPT *adversary* $\mathcal{A}$ *the following probability is negligible in the security parameter:*
$$\Pr[\mathsf{ExpAcc}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda) = 1]$$

where $\mathsf{ExpAcc}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda)$ *is the following game:*
Setup: *The challenger runs the Setup phase as in Definition 17.*
Queries: *The challenger runs the Query phase as in Definition 17.*
Challenge: *Finally, the adversary outputs an index* j^* *which he wants to be challenged on. If there exists a capability* cap_i *provided to* $\mathcal{A}$ *such that* $\mathbf{AC}(i, j^*) = rw$, *then the challenger aborts. The challenger runs* $d^* \leftarrow \langle \mathcal{C}_{\mathsf{read}}(cap_{\mathcal{O}}, j^*), \mathcal{S}_{\mathsf{read}}(\mathsf{DB})\rangle$ *and* $L \leftarrow \mathsf{blame}(cap_{\mathcal{O}}, \mathsf{Log}, j^*)$ *locally.*
Output: *It outputs 1 if and only if* $d^* \neq \mathsf{DB}'(j^*)$ *and* $\exists\, i \in L$ *that has not been queried by* $\mathcal{A}$ *to the interface* $\mathsf{corCl}(\cdot)$ *or* $L = []$.

D Full Cryptographic Proofs

In this section we formally define and sketch the proof for the correctness of our scheme (Section D.1) and we prove the theorems presented in Section 7 (Section D.2).

D.1 Correctness

In the following we state the conditions that determine the correctness of a multi-client ORAM. Intuitively, the primitive is correct if, given a successful execution of any algorithm, its outcome satisfies some specific constraints which guarantee the reliability of the whole scheme. We formalize the notion of correctness up to each operation defined within the multi-client ORAM.

Definition 22 (Correctness). *A Group ORAM* $(\mathsf{gen}, \mathsf{addCl}, \mathsf{addE}, \mathsf{chMode}, \mathsf{read}, \mathsf{write})$ *is correct, if the following statements are true except with negligible probability in the security parameter: let* D *be the payload domain and* $cnt_{\mathcal{C}}$ *be the number of registered users.*
$\mathsf{addCl}.$

$$\forall d \in D, \forall \mathbf{a} \in \{\bot, r, rw\}^{|\mathsf{DB}|}, \forall j \in [1, |\mathsf{DB}|]:$$
$$cap_i \leftarrow \mathsf{addCl}(cap_{\mathcal{O}}, \mathbf{a}) \implies$$
$$(d \leftarrow \langle \mathcal{C}_{\mathsf{read}}(cap_i, j), \mathcal{S}_{\mathsf{read}}(\mathsf{DB})\rangle \iff$$
$$d = \mathsf{DB}(j) \,\wedge\, (\mathbf{a}(j) = r \,\vee\, \mathbf{a}(j) = rw))$$
$$\wedge\, (\mathsf{DB}' \leftarrow \langle \mathcal{C}_{\mathsf{write}}(cap_i, j, d), \mathcal{S}_{\mathsf{write}}(\mathsf{DB})\rangle \iff$$
$$\mathsf{DB}' = \mathsf{DB}[j \mapsto d] \,\wedge\, \mathbf{a}(j) = rw)$$

addE.

$$\forall d \in D, \forall \mathbf{a} \in \{\bot, r, rw\}^{|cnt_C|}, \forall i \in [1, cnt_C], \exists j :$$
$$\mathsf{DB}' \leftarrow \langle \mathcal{C}_{\mathsf{addE}}(cap_{\mathcal{O}}, \mathbf{a}, d), \mathcal{S}_{\mathsf{addE}}(\mathsf{DB}) \rangle \implies$$
$$(|\mathsf{DB}'| = j) \wedge (|\mathsf{DB}| = j - 1)$$
$$\wedge (\mathsf{DB}' = \mathsf{DB}[j \mapsto d])$$
$$\wedge (d \leftarrow \langle \mathcal{C}_{\mathsf{read}}(cap_i, j), \mathcal{S}_{\mathsf{read}}(\mathsf{DB}') \rangle \iff$$
$$d = \mathsf{DB}'(j) \ \wedge \ (\mathbf{a}(j) = r \ \vee \ \mathbf{a}(j) = rw))$$
$$\wedge \forall d'. (\mathsf{DB}'' \leftarrow \langle \mathcal{C}_{\mathsf{write}}(cap_i, j, d'), \mathcal{S}_{\mathsf{write}}(\mathsf{DB}') \rangle$$
$$\iff \mathsf{DB}'' = \mathsf{DB}'[j \mapsto d'] \ \wedge \ \mathbf{a}(j) = rw)$$

chMode.

$$\forall d \in D, \forall \mathbf{a} \in \{\bot, r, rw\}^{cnt_C},$$
$$\forall j \in [1, |\mathsf{DB}|], \forall i \in [1, cnt_C] :$$
$$\langle \mathcal{C}_{\mathsf{chMode}}(cap_{\mathcal{O}}, \mathbf{a}, j), \mathcal{S}_{\mathsf{chMode}}(\mathsf{DB}) \rangle \implies$$
$$(d \leftarrow \langle \mathcal{C}_{\mathsf{read}}(cap_i, j), \mathcal{S}_{\mathsf{read}}(\mathsf{DB}) \rangle \iff$$
$$d = \mathsf{DB}(j) \ \wedge \ (\mathbf{a}(j) = r \ \vee \ \mathbf{a}(j) = rw))$$
$$\wedge (\mathsf{DB}' \leftarrow \langle \mathcal{C}_{\mathsf{write}}(cap_i, j, d), \mathcal{S}_{\mathsf{write}}(\mathsf{DB}) \rangle \iff$$
$$\mathsf{DB}' = \mathsf{DB}[j \mapsto d] \ \wedge \ \mathbf{a}(j) = rw)$$

read.

$$\forall d \in D, \forall j \in [1, |\mathsf{DB}|], \forall i \in [1, cnt_C] :$$
$$d \leftarrow \langle \mathcal{C}_{\mathsf{read}}(cap_i, j), \mathcal{S}_{\mathsf{read}}(\mathsf{DB}) \rangle \implies$$
$$d = \mathsf{DB}(j) \ \wedge \ (\mathbf{AC}(i, j) = r \ \vee \ \mathbf{AC}(i, j) = rw)$$

write.

$$\forall d \in D, \forall j \in [1, |\mathsf{DB}|], \forall i \in [1, cnt_C] :$$
$$\mathsf{DB}' \leftarrow \langle \mathcal{C}_{\mathsf{write}}(cap_i, j, d), \mathcal{S}_{\mathsf{write}}(\mathsf{DB}) \rangle \implies$$
$$\mathsf{DB}' = \mathsf{DB}[j \mapsto d] \ \wedge \ \mathbf{AC}(i, j) = rw$$

Theorem 5 (Correctness). *Let* Π_{PE} *and* Π_{PO} *be a predicate (resp. predicate-only) encryption scheme,* Π_{PKE} *and* Π_{SE} *be a public-key (resp. private-key) encryption scheme, and* $\mathcal{ZKP}$ *be a zero-knowledge proof system such that they all fulfill completeness. Then* GORAM *constructed in Section 3.2 satisfies the definition of correctness (Definition 22).*

Proof sketch. The proof is conducted by protocol inspection on the implementation of each algorithm. Under the assumption that all of the encryption schemes, as well as the zero-knowledge proof system, are complete except with negligible probability, it directly follows from the analysis of the protocols that all of our instantiations fulfill the definition of correctness. $\square$

D.2 Security Proofs

Proof of Theorem 2. In the following we separately prove the security of GORAM for each specified property.

Lemma 1. *Let* Π_{PE} *be an attribute-hiding predicate encryption scheme and* Π_{SE} *be a CPA-secure private-key encryption scheme. Then* GORAM *achieves* secrecy.

Proof of Lemma 1. The proof is constructed by fixing the choice of the challenger over the sampling of the random coin and define intermediate hybrid games, where the two extremes are the $\mathsf{ExpSec}_{\mathsf{GORAM}}^{\mathcal{A}}(\lambda, b)$ experiment over the two values of b. We start by defining a new experiment $\mathsf{ExpSec}'^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda, b)$ that slightly differs from the original one. For sake of readability we introduce the following notation:

- GAME 1 := $\mathsf{ExpSec}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda, 0)$

- GAME 2 := $\mathsf{ExpSec}'^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda, 0)$

- GAME 3 := $\mathsf{ExpSec}'^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda, 1)$

- GAME 4 := $\mathsf{ExpSec}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda, 1)$

Then we show that the difference among any two neighboring games is bounded by a negligible value in the security parameter, therefore the advantage of the adversary in $\mathsf{ExpSec}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda, b)$ turns to be a sum of negligible values, which is still negligible. In particular we demonstrate the following:

$$\text{GAME } 1 \approx \text{GAME } 2 \approx \text{GAME } 3 \approx \text{GAME } 4$$

New Experiment. We define $\mathsf{ExpSec}'^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda, b)$ as the following game:

Setup. The challenger runs $(cap_{\mathcal{O}}, \mathsf{DB}) \leftarrow \mathsf{gen}(1^\lambda)$, sets $\mathbf{AC} := []$, and runs a black-box simulation of $\mathcal{A}$ to which it hands over DB.

Queries. The challenger provides $\mathcal{A}$ with an addCl, an addE, a chMode, a read, a write, and a corCl interactive interface that $\mathcal{A}$ may query adaptively and in any order. Each round $\mathcal{A}$ can query exactly one interface. These interfaces are described below:

(1) On input $\mathsf{addCl}(\mathbf{a})$ by $\mathcal{A}$, the challenger executes $\mathsf{addCl}(cap_{\mathcal{O}}, \mathbf{a})$ locally and stores the capability cap_i returned by the algorithm.

(2) On input $\mathsf{addE}(\mathbf{a}, d)$ by $\mathcal{A}$, the challenger executes $\langle \mathcal{C}_{\mathsf{addE}}(cap_{\mathcal{O}}, \mathbf{a}, d), \mathcal{S}_{\mathsf{addE}}(\mathsf{DB}) \rangle$ in interaction with $\mathcal{A}$, where the former plays the role of the client while the latter plays the role of the server.

(3) On input $\mathsf{chMode}(\mathbf{a}, j)$ by $\mathcal{A}$, the challenger executes $\langle \mathcal{C}_{\mathsf{chMode}}(cap_{\mathcal{O}}, \mathbf{a}, j), \mathcal{S}_{\mathsf{chMode}}(\mathsf{DB}) \rangle$ in interaction with $\mathcal{A}$.

(4) On input $\mathsf{corCl}(i)$ by $\mathcal{A}$, the challenger hands over the capability cap_i related to the i-th client in the access control matrix $\mathbf{AC}$.

(5) On input $\mathsf{read}(i, j)$ by $\mathcal{A}$, the challenger executes $\langle \mathcal{C}_{\mathsf{read}}(cap_i, j), \mathcal{S}_{\mathsf{read}}(\mathsf{DB}) \rangle$ in interaction with $\mathcal{A}$.

(6) On input $\mathsf{write}(i, j, d)$ by $\mathcal{A}$, the challenger executes $\langle \mathcal{C}_{\mathsf{write}}(cap_i, j, d), \mathcal{S}_{\mathsf{write}}(\mathsf{DB}) \rangle$ in interaction with $\mathcal{A}$.

Challenge. Finally, $\mathcal{A}$ outputs $(j, (d_0, d_1))$, where j is an index denoting the database entry on which $\mathcal{A}$ wishes to be challenged and (d_0, d_1) is a pair of entries such that $|d_0| = |d_1|$. The challenger accepts the request only if $\mathbf{AC}(i, j) = \bot$, for every i corrupted by $\mathcal{A}$ in the query phase. Afterwards, the challenger invokes $\langle \mathcal{C}_{\mathsf{write}}(cap_{\mathcal{O}}, j, d_b), \mathcal{S}_{\mathsf{write}}(\mathsf{DB}) \rangle$ in interaction with $\mathcal{A}$, as explained in Section 3, with the difference that the new entry is computed as follows:

$$E'_j = \begin{pmatrix} c'_{1,j} \leftarrow \mathsf{Rerand}(pk, c_{1,j}, r_1) \\ c'_{2,j} \leftarrow \mathsf{Enc}(pk, \mathsf{PoRR}(opk, c_{\mathsf{Auth}}, r_2)) \\ c'_{3,j} \leftarrow \mathsf{Enc}(pk, \mathsf{PrEnc}(mpk, I, r)) \\ c'_{4,j} \leftarrow \mathsf{Enc}(pk, c'_{\mathsf{Data}}) \end{pmatrix}$$

where I is the attribute previously associated to the entry j and r is a random string such that $|r| = |k(j)|$.

Output. In the output phase $\mathcal{A}$ still has access to the interfaces except for addCl on input $\mathbf{a}$ such that $\mathbf{a}(j) \neq \bot$; corCl on input i such that $\mathbf{AC}(i,j) \neq \bot$; and chMode on input $\mathbf{a}$ and j with $\mathbf{a}(i) \neq \bot$ for some previously corrupted client i. Eventually, $\mathcal{A}$ stops, outputting a bit b'. The challenger outputs 1 if and only if $b = b'$.

GAME 1 $\approx$ GAME 2. We assume toward contradiction that there exists a PPT adversary $\mathcal{A}$ that is able to distinguish among GAME 1 and GAME 2 with non-negligible probability, namely:

$$|\mathrm{Pr}[\text{GAME 1} = 1] - \mathrm{Pr}[\text{GAME 2} = 1]| \geq \epsilon(\lambda)$$

for some non-negligible $\epsilon(\lambda)$. Then we show that we can use such an adversary to build the following reduction $\mathcal{B}$ against the attribute-hiding property of the predicate-encryption scheme Π_{PE} defined in Definition 8. The simulation is elaborated below.

Setup. $\mathcal{B}$ receives as input the security parameter 1^λ from the challenger and it forwards it to $\mathcal{A}$. $\mathcal{B}$ initializes uniformly at random an attribute I and it sends to the challenger the tuple (I, I), who replies with the public key of the predicate-encryption scheme mpk^*. Finally $\mathcal{B}$ runs $(cap_{\mathcal{O}}, \mathsf{DB}) \leftarrow \mathsf{gen}(1^\lambda)$ as described in Section 3 without $\mathsf{PrGen}(1^\lambda)$, setting $mpk = mpk^*$ instead. Subsequently $\mathcal{B}$ gives pk_O and DB to $\mathcal{A}$.

Queries. $\mathcal{B}$ provides then $\mathcal{A}$ with the following interfaces:

(1) On input addCl($\mathbf{a}$) by $\mathcal{A}$, $\mathcal{B}$ initializes a predicate f such that $f(I) = \bot$ and $\forall j \in \mathsf{DB}$ it holds that $f(I_j) = 0$ whenever $\mathbf{a}(j) = \bot$ and $f(I_j) = 1$ otherwise.

(2) On input addE($\mathbf{a}, d$) by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{addE}}(cap_{\mathcal{O}}, \mathbf{a}, d), \mathcal{S}_{\mathsf{addE}}(\mathsf{DB}) \rangle$ in interaction with $\mathcal{A}$.

(3) On input chMode($\mathbf{a}, j$) by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{chMode}}(cap_{\mathcal{O}}, \mathbf{a}, j), \mathcal{S}_{\mathsf{chMode}}(\mathsf{DB}) \rangle$ in interaction with $\mathcal{A}$.

(4) On input corCl(i) by $\mathcal{A}$, $\mathcal{B}$ queries the oracle provided by the challenger on f_i so to retrieve the corresponding key sk_{f_i}. $\mathcal{B}$ constructs cap_i using such a key, which is handed over to $\mathcal{A}$.

(5) On input read(i, j) by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{read}}(cap_i, j), \mathcal{S}_{\mathsf{read}}(\mathsf{DB}) \rangle$ in interaction with $\mathcal{A}$.

(6) On input write(i, j, d) by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{write}}(cap_i, j, d), \mathcal{S}_{\mathsf{write}}(\mathsf{DB}) \rangle$ in interaction with $\mathcal{A}$.

Challenge. Finally, $\mathcal{A}$ outputs $(j, (d_0, d_1))$, where j is an index denoting the database entry on which $\mathcal{A}$ wishes to be challenged and (d_0, d_1) is a pair of entries such that $|d_0| = |d_1|$. $\mathcal{B}$ accepts the tuple only if $\mathbf{AC}(i,j) = \bot$, for every i corrupted by $\mathcal{A}$ in the query phase. $\mathcal{B}$ parses then the j-th entry as $E_j = (c_{1,j}, c_{2,j}, c_{3,j}, c_{4,j})$, it fetches $c_{\mathsf{Key}} \leftarrow \mathsf{Dec}(sk, c_{3,j})$, and it finally gets $k(j) \leftarrow \mathsf{PrDec}(sk_O^f, c_{\mathsf{Key}})$, for some suitable sk_O^f. Afterwards, $\mathcal{B}$ sends the tuple $(m_0, m_1) = (r, k(j))$ to the challenger, where r is a random string such that $|r| = |k(j)|$. The challenger answers back with the challenge ciphertext $c^* \leftarrow \mathsf{PrEnc}(mpk, I, m_b)$ that $\mathcal{B}$ uses to execute $\langle \mathcal{C}_{\mathsf{write}}(cap_{\mathcal{O}}, j, d_0), \mathcal{S}_{\mathsf{write}}(\mathsf{DB}) \rangle$, computing the new entry in the following manner:

$$E'_j = \left(\begin{array}{l} c'_{1,j} \leftarrow \mathsf{Rerand}(pk, c_{1,j}, r_1) \\ c'_{2,j} \leftarrow \mathsf{Enc}(pk, \mathsf{PoRR}(opk, c_{\mathsf{Auth}}, r_2)) \\ c'_{3,j} \leftarrow \mathsf{Enc}(pk, c^*) \\ c'_{4,j} \leftarrow \mathsf{Enc}(pk, c'_{\mathsf{Data}}) \end{array} \right).$$

Output. In the output phase $\mathcal{A}$ still has access to the interfaces except for addCl on input $\mathbf{a}$ such that $\mathbf{a}(j) \neq \bot$; corCl on input i such that $\mathbf{AC}(i,j) \neq \bot$; and chMode on input $\mathbf{a}, j$ with $\mathbf{a}(i) \neq \bot$ for some previously corrupted client i. Note that in case there exists some non-corrupted i such that $\mathbf{a}(i) \neq \bot$, $\mathcal{B}$ just simulates the $\langle \mathcal{C}_{\mathsf{chMode}}(cap_{\mathcal{O}}, \mathbf{a}, j), \mathcal{S}_{\mathsf{chMode}}(\mathsf{DB}) \rangle$ protocol

by rerandomizing the challenge ciphertext rather than re-encrypting it. Eventually, $\mathcal{A}$ stops, outputting a bit b'. $\mathcal{B}$ outputs 1 if and only if $b' = 0$.

The simulation is clearly efficient, also it is easy to see that whenever the challenger samples $b = 1$, the simulation perfectly reproduces the inputs that $\mathcal{A}$ is expecting in GAME 1. The only difference is indeed that in the challenge phase c'_{Data} is re-encrypted in the simulation, while in the real experiment it is just rerandomized; also in the output phase, the interface chMode on j is simulated with a rerandomization, rather than a re-encryption on c_{Key}. By definition of rerandomization, however, these two operations are indistinguishable to $\mathcal{A}$. Thus, we can state the following:

$$\Pr[\mathcal{B} \mapsto 1 | b = 1] \approx \Pr[\mathsf{GAME}\ 1 = 1].$$

On the other hand, in case the challenger initializes $b = 0$, then $\mathcal{B}$ perfectly simulates the environment that $\mathcal{A}$ is expecting in GAME 2. Therefore we can assert that:

$$\Pr[\mathcal{B} \mapsto 1 | b = 0] \approx \Pr[\mathsf{GAME}\ 2 = 1].$$

However, it follows from the initial assumption that

$$|\Pr[\mathcal{B} \mapsto 1 | b = 1] - \Pr[\mathcal{B} \mapsto 1 | b = 0]| \geq \epsilon(\lambda),$$

which is clearly a contradiction with respect to the attribute-hiding property of the predicate encryption scheme Π_{PE}, and it proves the initial lemma.

GAME 2 $\approx$ GAME 3. We assume toward contradiction that there exists a PPT adversary $\mathcal{A}$ that is able to distinguish among GAME 2 and GAME 3 with non-negligible probability, namely:

$$|\Pr[\mathsf{GAME}\ 2 = 1] - \Pr[\mathsf{GAME}\ 3 = 1]| \geq \epsilon(\lambda)$$

For some non-negligible $\epsilon(\lambda)$. Then we show that we can we can use such an adversary to build the following reduction $\mathcal{B}$ against the CPA-security property of the private-key encryption scheme Π_{SE}. The simulation is elaborated below.

Setup. $\mathcal{B}$ receives as input the security parameter 1^λ from the challenger and it forwards it to $\mathcal{A}$. $\mathcal{B}$ then runs $(cap_{\mathcal{O}}, \mathsf{DB}) \leftarrow \mathsf{gen}(1^\lambda)$ as described in Section 3 and it gives pk_O and DB to $\mathcal{A}$.

Queries. $\mathcal{B}$ provides then $\mathcal{A}$ with the following interfaces:

(1) On input $\mathsf{addCl}(\mathbf{a})$ by $\mathcal{A}$, $\mathcal{B}$ executes $\mathsf{addCl}(cap_{\mathcal{O}}, \mathbf{a})$ locally and it stores the capability cap_i returned by the algorithm.

(2) On input $\mathsf{addE}(\mathbf{a}, d)$ by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{addE}}(cap_{\mathcal{O}}, \mathbf{a}, d), \mathcal{S}_{\mathsf{addE}}(\mathsf{DB}) \rangle$ in interaction with $\mathcal{A}$.

(3) On input $\mathsf{chMode}(\mathbf{a}, j)$ by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{chMode}}(cap_{\mathcal{O}}, \mathbf{a}, j), \mathcal{S}_{\mathsf{chMode}}(\mathsf{DB}) \rangle$ in interaction with $\mathcal{A}$.

(4) On input $\mathsf{corCl}(i)$ by $\mathcal{A}$, $\mathcal{B}$ hands over the capability cap_i related to the i-th client in the access control matrix $\mathbf{AC}$.

(5) On input $\mathsf{read}(i, j)$ by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{read}}(cap_i, j), \mathcal{S}_{\mathsf{read}}(\mathsf{DB}) \rangle$ in interaction with $\mathcal{A}$.

(6) On input $\mathsf{write}(i, j, d)$ by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{write}}(cap_i, j, d), \mathcal{S}_{\mathsf{write}}(\mathsf{DB}) \rangle$ in interaction with $\mathcal{A}$.

Challenge. Finally, $\mathcal{A}$ outputs $(j, (d_0, d_1))$, where j is an index denoting the database entry on which $\mathcal{A}$ wishes to be challenged and (d_0, d_1) is a pair of entries such that $|d_0| = |d_1|$. $\mathcal{B}$ accepts the tuple only if $\mathbf{AC}(i, j) = \perp$, for every i corrupted by $\mathcal{A}$ in the query phase. $\mathcal{B}$ sends the tuple $(m_0, m_1) = (d_0, d_1)$ to the challenger, who answers back with the challenge ciphertext $c^* \leftarrow$

$\mathcal{E}(k, d_b)$ that $\mathcal{B}$ uses to perform a local execution of $\langle \mathcal{C}_{\mathsf{write}}(cap_{\mathcal{O}}, j, d_b), \mathcal{S}_{\mathsf{write}}(\mathsf{DB}) \rangle$, computing the new entry in the following manner:

$$E'_j = \begin{pmatrix} c'_{1,j} \leftarrow \mathsf{Rerand}(pk, c_{1,j}, r_1) \\ c'_{2,j} \leftarrow \mathsf{Enc}(pk, \mathsf{PoRR}(opk, c_{\mathsf{Auth}}, r_2)) \\ c'_{3,j} \leftarrow \mathsf{Enc}(pk, \mathsf{PrEnc}(mpk, I, r)) \\ c'_{4,j} \leftarrow \mathsf{Enc}(pk, c^*) \end{pmatrix}$$

where I is the attribute previously associated to the entry j and r is a random string such that $|r| = |k|$.

Output. In the output phase $\mathcal{A}$ still has access to the interfaces except for addCl on input $\mathbf{a}$ such that $\mathbf{a}(j) \neq \perp$; corCl on input i such that $\mathbf{AC}(i, j) \neq \perp$; and chMode on input $\mathbf{a}, j$ with $\mathbf{a}(i) \neq \perp$ for some previously corrupted client i. Eventually, $\mathcal{A}$ stops, outputting a bit b'. $\mathcal{B}$ outputs b' as well.

The simulation is obviously efficient, also it is easy to see that whenever the challenger samples $b = 0$, the simulation perfectly reproduces the inputs that $\mathcal{A}$ is expecting in GAME 2. Thus, we can state the following:

$$\Pr[\mathcal{B} \mapsto 0 | b = 0] \approx \Pr[\mathsf{GAME\ 2} = 1].$$

On the other hand, in case the challenger initializes $b = 1$, then $\mathcal{B}$ perfectly simulates the environment that $\mathcal{A}$ is expecting in GAME 3. Therefore we can assert that:

$$\Pr[\mathcal{B} \mapsto 1 | b = 1] \approx \Pr[\mathsf{GAME\ 3} = 1].$$

However, it follows from the initial assumption that:

$$|\Pr[\mathcal{B} \mapsto 1 | b = 1] - \Pr[\mathcal{B} \mapsto 0 | b = 0]| \geq \epsilon(\lambda),$$

$$\left| \Pr[\mathsf{Exp}^{\mathsf{cpa}}_{\mathcal{A}, \Pi_{\mathsf{SE}}}(\lambda, 1) = 1] - \Pr[\mathsf{Exp}^{\mathsf{cpa}}_{\mathcal{A}, \Pi_{\mathsf{SE}}}(\lambda, 0) = 1] \right| \\ \geq \epsilon(\lambda),$$

which implies a non-negligible difference in the success probability of $\mathcal{B}$ with respect to the random choice of b and it clearly represents a contradiction with respect to the CPA-security property of the private-key encryption scheme Π_{SE}. This proves the initial lemma.

GAME 3 $\approx$ GAME 4. The proof is conducted with the same pipeline of the indistinguishability between GAME 1 and GAME 2; the simulation also works correspondingly, except that in this case the reduction $\mathcal{B}$ encrypts the message d_1 rather than d_0. However, the analogous argument applies.

GAME 1 $\approx$ GAME 4. By the previous lemmas it directly follows that the difference among each couple of neighboring games is bounded by a negligible value, thus the difference between GAME 1 and GAME 4 is a sum of negligible terms, which is, again, negligible. In particular

$$\mathsf{GAME\ 1} \approx \mathsf{GAME\ 4}$$

directly implies that:

$$\mathsf{ExpSec}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda, 0) \approx \mathsf{ExpSec}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda, 1)$$

thus, $\forall$ PPT adversary the two experiments look indistinguishable. This concludes our proof. $\quad\square$

Lemma 2. *Let $\mathcal{ZKP}$ be a zero-knowledge proof system. Then* GORAM *achieves integrity.*

Proof of Lemma 2. The proof is conducted by contradiction. Assume that there exists a PPT $\mathcal{A}$ that wins the experiment defined in Definition 17 with non-negligible probability. Then we construct an adversary $\mathcal{B}$ that breaks the soundness of the zero-knowledge proof system $\mathcal{ZKP}$. The simulation works exactly as specified in Definition 17. Since the database is held on the side of $\mathcal{B}$, the adversary can only modify data by triggering interfaces and all of the resulting operations are executed locally by $\mathcal{B}$ as described in Section 3.2 except for the read and write protocols. Whenever an operation affects the database DB, $\mathcal{B}$ reflects these changes on DB', especially in the protocols addE, chMode, and write. This implies that the opportunity for the adversary to inject some data such that $\mathcal{B}$ does not update the local database DB' accordingly, is restricted to the read and write algorithms. We will briefly recall how the operations are performed from Section 3.2.

Read and write. When $\mathcal{A}$ triggers the read or write on a certain index j, $\mathcal{B}$ releases the path for the leaf l_j associated to such an entry $E = (E_1, \ldots, E_{b(D+1)})$, from ρ down to T_{D,l_j}. The database DB is kept blocked by $\mathcal{B}$ until $\mathcal{A}$ frees it again by submitting an updated path for l_j along with a valid shuffle proof. The proof looks as follows:

$$P = PK \left\{ \begin{array}{l} (\pi, r_1, \ldots, r_{b(D+1)}) : \\ \quad \forall \ell. \ E_\ell = \mathsf{Rerand}(pk, E_{\pi^{-1}(\ell)}, r_\ell) \end{array} \right\}. \tag{1}$$

It is easy to see that, by construction, the proof guarantees that the new path is just a shuffling and a re-randomization with respect to the old one. $\mathcal{B}$ verifies the proof against the new and the old path, i.e., it checks whether the proof of shuffle correctness verifies cryptographically, whether the inputs correspond to the components of the E_j, and whether the outputs correspond to the components of the E'_j. If all these checks succeed, $\mathcal{B}$ replaces the old with the new path. Subsequently $\mathcal{A}$ sends an updated top entry $E' = (c'_1, c'_2, c'_3, c'_4)$ for an old top entry $E = (c_1, c_2, c_3, c_4)$ along with a proof of the decryption for the top entry of the updated path:

$$P_{\mathsf{Auth}} = PK \{(osk_f) : \ \mathsf{PoDec}(osk_f, c_{\mathsf{Auth}}) = 1\} \tag{2}$$

together with information to access c_{Auth} (see Section 8.1 for how this information looks like in our concrete instantiation), and a proof that he did not change the index

$$P = PK \left\{ [\![c'_1]\!] = [\![c_1]\!] \right\}. \tag{3}$$

$\mathcal{B}$ checks then P_{Auth} against the content of c_2 of E and P against c_1 and c'_1. If all of the checks verify, $\mathcal{B}$ replaces the old by the new entry in the first position of the tree.

Analysis. We assume toward contradiction that there exists an $\mathcal{A}$ such that the experiment $\mathsf{ExpInt}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda)$ outputs 1 with non-negligible probability, namely

$$\mathsf{Pr}[\mathsf{ExpInt}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda) = 1] \geq \epsilon(\lambda).$$

As we argued before, by construction of the experiment we know that $\mathcal{A}$ can only inject entries by cheating in the read and write protocols. Thus we can restrict the success probability of $\mathcal{A}$ to the interaction within the read and write interfaces. Furthermore we split the success probability over his ability of breaking the zero-knowledge proof system $\mathcal{ZKP}$. We define BREAK to be the event in which $\mathcal{A}$ computes one zero-knowledge proof that verifies but without knowing the witness, i.e., he convinces the verifier of a false statement. In particular it follows from the initial assumption that

$$\begin{aligned} &\mathsf{Pr}[\mathsf{ExpInt}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda) = 1] = \\ &\quad \mathsf{Pr}[\mathcal{A} \text{ wins} \mid \mathsf{BREAK}] \cdot \mathsf{Pr}[\mathsf{BREAK}] + \\ &\quad \mathsf{Pr}[\mathcal{A} \text{ wins} \mid \neg \ \mathsf{BREAK}] \cdot \mathsf{Pr}[\neg \ \mathsf{BREAK}] \geq \epsilon(\lambda). \end{aligned}$$

It is easy to see that, given that $\mathcal{A}$ can compute a false statement that convinces the verifier in the zero-knowledge proof system $\mathcal{ZKP}$, he can win the game with probability 1, e.g., writing on an entry which he does not have writing access to. Therefore

$$\Pr[\mathcal{A} \text{ wins} \mid \mathsf{BREAK}] = 1.$$

On the other hand, given that $\mathcal{A}$ does not break the soundness of the zero-knowledge proof system $\mathcal{ZKP}$, the probability that $\mathcal{A}$ modifies an entry without $\mathcal{B}$ noticing is negligible in both read and write algorithms. This follows directly from the notion of correctness of the primitive Definition 22. Thus we can rewrite

$$\Pr[\mathcal{A} \text{ wins} \mid \neg\, \mathsf{BREAK}] \leq negl(\lambda).$$

It follows that we can bind the success probability of $\mathcal{A}$ as

$$\begin{aligned}
\Pr[\mathsf{ExpInt}_{\mathsf{GORAM}}^{\mathcal{A}}(\lambda) = 1] &\approx \\
1 \cdot \Pr[\mathsf{BREAK}] + negl(\lambda) \cdot \Pr[\neg\, \mathsf{BREAK}] &\approx \\
\Pr[\mathsf{BREAK}] &\geq \epsilon(\lambda).
\end{aligned}$$

Since $\mathcal{A}$ can query the interfaces only a polynomial number of times, say q, $\mathcal{B}$ can simply store all of the transcripts of the queries to the read and write interfaces and output one zero-knowledge proof chosen uniformly at random. There are at most $3 \cdot q$ zero-knowledge proof transcripts by construction of the protocol, therefore the probability that $\mathcal{B}$ outputs a zero-knowledge proof which verifies a false statement is lower bounded by $\epsilon(\lambda) \cdot \frac{1}{3 \cdot q}$ which is still a non-negligible value. This is clearly a contradiction with respect to the soundness property of the zero-knowledge proof system and it concludes our proof. $\qquad\square$

Lemma 3. $\mathcal{ZKP}$ *be a zero-knowledge proof system and* Π_{PE} *be an attribute-hiding predicate encryption scheme. Then* GORAM *achieves* tamper-resistance.

Proof of Lemma 3. The proof is elaborated by contradiction, assuming we have a PPT adversary $\mathcal{A}$ such that he is able to break the security of the game with non-negligible probability, we build a distinguisher $\mathcal{B}$ such that it runs the aforementioned $\mathcal{A}$ as a black-box to break the attribute-hiding property of the underlying predicate encryption scheme Π_{PE}.

In particular, we assume toward contradiction that the following inequality holds:

$$\Pr[\mathsf{ExpTamRes}_{\mathsf{GORAM}}^{\mathcal{A}}(\lambda) = 1] \geq \epsilon(\lambda).$$

Note that, by construction of the experiment, the challenge entry outputted by $\mathcal{A}$ is accepted by the challenger only if $\mathcal{A}$ never had read or write access to it. We now define a new event GUESS as the event in which $\mathcal{A}$ guesses the attribute encrypted in the c_{Key} cipher. It is easy to see that, with such a knowledge, $\mathcal{A}$ can always win the game by simply re-encrypting another symmetric key $k(j^*)'$ within the c_{Key} and using the same $k(j^*)'$ for encrypting another arbitrary plaintext into c_{Data}. It follows that when the challenger attempts to read the challenged index j^*, she will always succeed and she outputs 1 with overwhelming probability. Therefore we split the success probability of $\mathcal{A}$ over the probability that GUESS occurs:

$$\begin{aligned}
\Pr[\mathsf{ExpTamRes}_{\mathsf{GORAM}}^{\mathcal{A}}(\lambda) = 1] &= \\
\Pr[\mathsf{ExpTamRes}_{\mathsf{GORAM}}^{\mathcal{A}}(\lambda) = 1 \mid \mathsf{GUESS}] \cdot \Pr[\mathsf{GUESS}] &+ \\
\Pr[\mathsf{ExpTamRes}_{\mathsf{GORAM}}^{\mathcal{A}}(\lambda) = 1 \mid \neg\, \mathsf{GUESS}] \cdot \Pr[\neg\, \mathsf{GUESS}].
\end{aligned}$$

As reasoned above, the success probability of $\mathcal{A}$ given his knowledge of the attribute encrypted within c_{Key} is overwhelming, thus we can rewrite:

$$\Pr[\mathsf{ExpTamRes}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda) = 1] =$$
$$\Pr[\mathsf{GUESS}] +$$
$$\Pr[\mathsf{ExpTamRes}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda) = 1 | \neg\, \mathsf{GUESS}] \cdot \Pr[\neg\, \mathsf{GUESS}].$$

We now consider the success probability of $\mathcal{A}$ given that he does not know the attribute of the Π_{PE} cipher relative to the challenge entry. It follows from the proof of Lemma 2 that $\mathcal{A}$ cannot modify the cipher via the interfaces provided by the challenger without being detected. However, since $\mathcal{A}$ stores the database locally, he can still perform local modifications. Nevertheless, in order to win the experiment, it must be the case that the challenger succeeds in reading j^*, that implies the challenge entry to be encrypted such that the capability held by the challenger allows her to access such an entry. In practice, this means that the attribute I of the cipher c_{Key} in the challenge entry is orthogonal with respect to the predicate f held by the challenger, as defined in Definition 7. By assumption $\mathcal{A}$ knows neither the attribute I in c_{Key} nor the predicate f, therefore the probability that he computes an attribute I' such that it is orthogonal to f is negligible by construction. This implies that $\mathcal{A}$ can only win by locally modifying the cipher c_{Data} such that it successfully decrypts to another arbitrary message, given that $\mathcal{A}$ has neither knowledge nor the possibility to modify the key $k(j^*)$. This, however, is a contradiction to the elusive-range property of the symmetric encryption scheme Π_{SE} (see Definition 4). Hence, it follows that the probability that $\mathcal{A}$ wins the experiment, given that he does not guess the attribute, is upper bounded by a negligible value. Then we have:

$$\Pr[\mathsf{ExpTamRes}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda) = 1] =$$
$$\Pr[\mathsf{GUESS}] + negl(\lambda) \cdot \Pr[\neg\, \mathsf{GUESS}],$$

which by assumption is lower bounded by:

$$\Pr[\mathsf{ExpTamRes}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda) = 1] \approx \Pr[\mathsf{GUESS}] \geq \epsilon(\lambda).$$

New Experiment. We now define the experiment $\mathsf{ExpSInt'}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda)$ analogously to the game $\mathsf{ExpTamRes}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda)$, except that in the challenge phase the adversary outputs the attribute I that is encrypted in the c_{Key} cipher of such an entry, along with the challenge index j^*. By the argument above it follows that $\mathcal{A}$ has the same success probability in both games. The experiment is defined as follows:

Setup. The challenger runs $(cap_{\mathcal{O}}, \mathsf{DB}) \leftarrow \mathsf{gen}(1^{\lambda})$ as in Definition 16. Furthermore, it forwards DB to $\mathcal{A}$ and it initializes a second database DB' which is managed locally.

Queries. The challenger provides $\mathcal{A}$ with the same interfaces as in Definition 16, which $\mathcal{A}$ may query adaptively and in any order. In all interactive protocols, $\mathcal{A}$ plays the role of the server and the challenger the one of the client. Furthermore, the challenger updates DB' whenever the protocols triggered by addE (2), chMode (3), and write (6) complete successfully.

Challenge. Finally, the adversary outputs an index j^* on which he wants to be challenged, along with the attribute I^*. If there exists a capability cap_i that has ever been provided to $\mathcal{A}$ such that $\mathbf{AC}(i, j) \neq \bot$, then the challenger aborts. The challenger runs $d^* \leftarrow \langle \mathcal{C}_{\mathsf{read}}(cap_{\mathcal{O}}, j^*), \mathcal{S}_{\mathsf{read}}(\mathsf{DB}) \rangle$ in interaction with $\mathcal{A}$.

Output. It outputs 1 if and only if $d^* \neq \mathsf{DB}'(j^*)$ and I^* is the attribute associated with the entry j^*.

Reduction. Note that the experiment $\mathsf{ExpSInt'}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda)$ reproduces the experiment $\mathsf{ExpTamRes}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda)$ except that $\mathcal{A}$ must hold the knowledge of the attribute I associated

to the challenge entry in order to win the former. However, since we argued that the success probability of $\mathcal{A}$ in the latter game implies such a knowledge, we can conclude that

$$\Pr[\mathsf{ExpTamRes}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda) = 1] \approx \Pr[\mathsf{GUESS}]$$
$$\approx \Pr[\mathsf{ExpSInt'}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda) = 1]$$
$$\geq \epsilon(\lambda).$$

Under this assumption, we can build the following reduction $\mathcal{B}$ against the attribute-hiding for multiple messages property of the predicate encryption scheme Π_{PE}. Note that, even though we did not explicitly state such a property so far, it is implied by the attribute-hiding notion and we will subsequently show why this is the case. We define the experiment $\mathsf{ExpPEMulti}(\lambda, b)$ as follows:

Experiment $\mathsf{ExpPEMulti}(\lambda, b)$
    $\Sigma^2 \ni (I_0, I_1) \leftarrow \mathcal{A}(1^\lambda)$
    $(mpk, psk) \leftarrow \mathsf{PrGen}(1^\lambda)$ and give mpk to $\mathcal{A}$
    $\mathcal{A}$ may adaptively query psk_{f_i} for predicates
        $f_1, \ldots, f_\ell \in \mathcal{F}$ where $f_i(I_0) = f_i(I_1)$
    $\mathbf{M} \leftarrow \mathcal{A}$ such that
        $\mathbf{M} = ((m_{0,0}, m_{1,0}), ..., (m_{0,t}, m_{1,t}))$
        and $|m_{0,j}| = |m_{1,j}|$ and
        if there is an i such that $f_i(I_0) = f_i(I_1) = 1$,
        then $m_{0,j} = m_{1,j}$ is required for all j
    $b' \leftarrow \mathcal{A}(\mathbf{C_b})$ where
        $\mathbf{C_b} = (\mathsf{PrEnc}(mpk, I_b, m_{b,0}), \ldots,$
        $\mathsf{PrEnc}(mpk, I_b, m_{b,t}))$
        while $\mathcal{A}$ may continue requesting keys for
        additional predicates with the same
        restrictions as before
    output 1 if and only if, $b' = b$.

The simulation works as follows:

Setup. $\mathcal{B}$ receives as input the security parameter 1^λ from the challenger and it forwards it to $\mathcal{A}$. $\mathcal{B}$ initializes uniformly at random an attribute pair (I_0, I_1) and it sends it to the challenger, who replies with the public key of the predicate-encryption scheme mpk^*. Finally $\mathcal{B}$ runs $(cap_{\mathcal{O}}, \mathsf{DB}) \leftarrow \mathsf{gen}(1^\lambda)$ as described in Section 3 without $\mathsf{PrGen}(1^\lambda)$, setting $mpk = mpk^*$ instead. Subsequently $\mathcal{B}$ gives $pk_{\mathcal{O}}$ to $\mathcal{A}$, and it initializes a second database DB' which is managed locally.

Queries. $\mathcal{B}$ provides then $\mathcal{A}$ with the following interfaces:

(1) On input $\mathsf{addCl}(\mathbf{a})$ by $\mathcal{A}$, $\mathcal{B}$ initializes a predicate f such that $\forall j \in \mathsf{DB}$ it holds that $f(I_j) = \mathbf{a}(j)$. Note that some entry j may be encrypted under either the I_0 or I_1 attribute, in this case f is chosen such that $f(I_0) = f(I_1) = \mathbf{a}(j)$. Then it queries the oracle provided by the challenger on f so to retrieve the corresponding key $sk_{\mathcal{O}}^f$. $\mathcal{B}$ constructs cap_i using such a key, which is stored locally.

(2) On input $\mathsf{addE}(\mathbf{a}, d)$ by $\mathcal{A}$, $\mathcal{B}$ checks whether there exists some corrupted i such that $\mathbf{AC}(i, j) \neq \bot$, if this is the case it executes $\langle \mathcal{C}_{\mathsf{addE}}(cap_{\mathcal{O}}, \mathbf{a}, d), \mathcal{S}_{\mathsf{addE}}(\mathsf{DB}) \rangle$ in interaction with $\mathcal{A}$ who holds DB. Otherwise, $\mathcal{B}$ initializes a new symmetric key $k(j)$ and it sends the tuple $(m_0, m_1) = (k(j), k(j))$ to the challenger, who answers back with the challenge ciphertext $c^* \leftarrow \mathsf{PrEnc}(mpk, I_b, m_b)$ that $\mathcal{B}$ uses to perform an execution of

$\langle \mathcal{C}_{\mathsf{addE}}(cap_\mathcal{O}, \mathbf{a}, d), \mathcal{S}_{\mathsf{addE}}(\mathsf{DB})\rangle$ in interaction with $\mathcal{A}$, setting $c_{\mathsf{Key}} \leftarrow c^*$. In both cases, $\mathcal{B}$ updates DB' with the output of the algorithm.

(3) On input $\mathsf{chMode}(\mathbf{a}, j)$ by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{chMode}}(cap_\mathcal{O}, \mathbf{a}, j), \mathcal{S}_{\mathsf{chMode}}(\mathsf{DB})\rangle$ in interaction with $\mathcal{A}$, who holds DB. Note that, if it is still the case that no corrupted i has access to the entry j, the c_{Key} of such an entry is rerandomized during the execution of the protocol, rather than re-encrypted.

(4) On input $\mathsf{corCl}(i)$ by $\mathcal{A}$, $\mathcal{B}$ recomputes the predicate related to the i-th client in the access control matrix $\mathbf{AC}$ such that it fulfills the access control policy described by $\mathbf{AC}$. Note that some entry j may be encrypted under either the I_0 or I_1 attribute, in this case f is chosen such that $f(I_0) = f(I_1) = \mathbf{AC}(i, j)$. $\mathcal{B}$ then queries the oracle provided by the challenger on such a predicate and it receives back a secret key sk_f which it sends to $\mathcal{A}$ together with the rest of the keys that form the capability cap_i.

(5) On input $\mathsf{read}(i, j)$ by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{read}}(cap_i, j), \mathcal{S}_{\mathsf{read}}(\mathsf{DB})\rangle$ in interaction with $\mathcal{A}$, where the former plays the role of the client and the latter plays the role of the server.

(6) On input $\mathsf{write}(i, j, d)$ by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{write}}(cap_i, j, d), \mathcal{S}_{\mathsf{write}}(\mathsf{DB})\rangle$ in interaction with $\mathcal{A}$, where the former plays the role of the client and the latter plays the role of the server. Eventually $\mathcal{B}$ updates DB' with the output of the algorithm.

Challenge. Finally, $\mathcal{A}$ outputs (j^*, I^*). $\mathcal{B}$ then checks whether $I^* = I_1$, if this is the case it outputs 1, otherwise it outputs 0.

It is easy to see that the reduction $\mathcal{B}$ is efficient and it perfectly simulates the inputs that $\mathcal{A}$ is expecting in the experiment $\mathsf{ExpSInt}'^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda)$, therefore we can bind the success probability of $\mathcal{B}$ over the random choice of b in $\mathsf{ExpPEMulti}^{\mathcal{B}}_{\mathsf{GORAM}}$, to the success probability of $\mathcal{A}$. It follows from the initial assumption that

$$\mathsf{Pr}[\mathsf{ExpPEMulti}^{\mathcal{B}}_{\mathsf{GORAM}}(\lambda, b) = 1] \approx$$
$$\mathsf{Pr}[\mathsf{ExpSInt}'^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda) = 1] \geq \epsilon(\lambda)$$

which is clearly a contradiction with respect the attribute-hiding for multiple messages property of the predicate encryption scheme Π_{PE}.

Attribute-Hiding for Multiple Messages. What is left to show, is that the security notion of attribute-hiding of a predicate encryption scheme implies the attribute-hiding for multiple messages. The demonstration consists of a standard hybrid argument over the vector of ciphertexts. We define a hybrid distribution H_i where the first i ciphertexts contain the encryption of the message m_0 and the remaining $n - i$ are the encryption of m_1. Observe that $H_0 = C_0$ and $H_1 = C_1$. Assuming toward contradiction that there exists a distinguisher $\mathcal{A}$ such that:

$$|\mathsf{Pr}[\mathcal{A}(H_0) = 1] - \mathsf{Pr}[\mathcal{A}(H_1) = 1]| \geq \epsilon(\lambda),$$

then it must be the case that there exists some i such that

$$|\mathsf{Pr}[\mathcal{A}(H_i) = 1] - \mathsf{Pr}[\mathcal{A}(H_{i+1}) = 1]| \geq \epsilon(\lambda).$$

Notice that the only difference between H_i and H_{i+1} is that the $(i+1)$-th cipher is the encryption of 1 in H_i, while it is the encryption of 0 in H_{i+1}. Now, given a cipher c_b, it is easy to construct a distribution $H = (c_{0,0}, ..., c_{0,i}, c_b, c_{1,i+2}, ..., c_{1,n})$ which is equal to H_i if and only if c_b is the encryption of 1 and it is equal to H_{i+1} otherwise. It follows that it is possible to distinguish

the encryption of either m_0 or m_1 with the same probability with which $\mathcal{A}$ distinguishes H_i and H_{i+1}. This, however, contradicts our initial assumption and it proves the implication.

Versioning. Note that by the previous argument we only ruled out the cases where the server maliciously modifies an entry of the database without the client noticing it. However it is still possible for an adversary to win the game just by querying the write interface for an allowed modification on a given entry and then provide an old version of DB in the challenge phase. Note that this is a class of attacks that is inherent to the cloud storage design and can be prevented via standard techniques (e.g., by using gossip protocols [25] among the clients). For the sake of simplicity we do not consider such attacks in our proof and we implicitly assume that the adversary always runs the read algorithm on top of the most recent version of DB in the challenge phase of the cryptographic game. □

Lemma 4. *Let $\mathcal{ZKP}$ be a zero-knowledge proof system and Π_{PKE} be a CPA-secure public-key encryption scheme. Then GORAM achieves obliviousness.*

Proof of Lemma 4. The proof consists of the analysis of the distribution of the read and write operations over the access pattern of the paths in the binary tree. Combining the uniform distribution over the retrieved path with the indistinguishability among the read and write algorithms, it directly follows that any two sequences of queries are indistinguishable, i.e., it cannot exist any adversary who wins the experiment $\mathsf{ExpObv}_{\mathsf{GORAM}}^{\mathcal{A}}(\lambda, b)$ with non-negligible probability.

It follows from the design of the primitive that each bucket independently represents a trivial ORAM and therefore preserves oblivious access. Indeed, every bucket is always retrieved as a whole and it is re-randomized upon every access. By the CPA-security of the top layer public key encryption scheme Π_{PKE}, we can state that the rerandomization operation completely hides any modification of the data. Thus, we can restrict the information leaked by each operation to the path of the binary tree that gets retrieved and, in order to prove obliviousness, it is sufficient to demonstrate that each client's access leads to the same distribution over the choice of such a path.

Read. In the read algorithm the path associated with the target entry is initially retrieved by the client, note that the association leaf-entry was sampled uniformly at random. On the client-side, a new leaf-entry association is uniformly generated for the target entry and the path is arranged such that each entry is pushed as far as possible toward its designated leaf in the tree. In this process the client must make sure that a dummy entry is placed on top of the path, however the server does not gather any additional information from this procedure because of the CPA-security of Π_{PKE}. The path is then rerandomized and uploaded on the hosting server. It is easy to see that any further access of any entry (including the most recently accessed one) will always determine a path only depending on the association between leaf and entry, which is uniformly distributed. It follows that, for all accesses, the distribution over the choice of the path is uniform and independent with respect to the accessed entry. After the shuffling the client overwrites the dummy entry placed on top of the root node, however this does not affect the obliviousness of the algorithm since the memory location accessed is fixed.

Write. The write algorithm works analogously to the read, so the argument above still applies. The only difference in this scenario is that a target entry is selected to be placed on top of the root, instead of a dummy one. This, however, does not introduce any difference in the server's view since the choice of the entry to set on top of the path is again hidden by the CPA-security of Π_{PKE}. Thus, we achieve uniform distribution of the memory accesses in the write operation and consequently indistinguishability among the read and write protocols.

Zero-Knowledge. In both read and write cases, the client attaches along with the data sent to the server three non-interactive zero-knowledge proves that the server must verify in order to protect the integrity of the data. Specifically this set of proofs is composed of a proof of a shuffle correctness (line 6.8 and line 7.8), a proof of writing eligibility (line 6.11), and a proof of plaintext-equivalence for the index (line 6.13). However, due to the zero-knowledge property of $\mathcal{ZKP}$ the proved statements do not reveal any information so as to give $\mathcal{A}$ additional information that would allow him to break obliviousness. □

Lemma 5. *Let $\mathcal{ZKP}$ be a zero-knowledge proof system. Then* GORAM *achieves* anonymity.

Proof of Lemma 5. The proof proceeds by contradiction. We show that, given an adversary who is able to win the experiment $\mathsf{ExpAnon}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda, b)$ with non-negligible probability, we can construct an efficient algorithm that breaks the zero-knowledge property of $\mathcal{ZKP}$.

By construction of the game, the only information available to the adversary in order to distinguish between the two capabilities is the execution of the read or write protocol, depending on the input he provides to the challenger. In both cases, it follows from the inspection of the protocol (see Section 3) that the only step which is not independent from the capability held by the client is the formulation of the authorization proof. We recall that such a statement proves the knowledge of a key osk_f which can be used to successfully decrypt the ciphertext c_{Auth}. The proof looks as follows:

$$P_{\mathsf{Auth}} = PK\left\{(osk_f): \ \mathsf{PoDec}(osk_f, c_{\mathsf{Auth}}) = 1\right\},$$

note, indeed, that in the instantiation of our protocol, the capability of the client is implemented as the client's secret key osk_f. Thus, all of the read or write transcripts, except for such a zero-knowledge proof, are trivially indistinguishable over the choice of the capability. It follows that any information that the adversary gathers, can only be derived from the transcript of the authentication proof. Assuming toward contradiction that

$$\Pr[\mathsf{ExpAnon}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda, b) = 1] \geq \frac{1}{2} + \epsilon(\lambda),$$

then it must be the case that $\mathcal{A}$ is able to correctly guess the capability chosen by the challenger to formulate the proof P_{Auth} with advantage at least $\epsilon(\lambda)$ over the random choice of b. We define EXTRACT as the event in which $\mathcal{A}$ extracts some additional information from P_{Auth} about the capability used to construct the proof. It follows from the initial assumption that

$$\begin{aligned}
\Pr[\mathsf{ExpAnon}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda, b) = 1] = \\
\Pr[\mathcal{A} \text{ wins} \mid \mathsf{EXTRACT}] \cdot \Pr[\mathsf{EXTRACT}] + \\
\Pr[\mathcal{A} \text{ wins} \mid \neg\,\mathsf{EXTRACT}] \cdot \Pr[\neg\,\mathsf{EXTRACT}] \\
\geq \tfrac{1}{2} + \epsilon(\lambda).
\end{aligned}$$

It is easy to see that, given that $\mathcal{A}$ can deduce some information about the capability from P_{Auth}, i.e., that EXTRACT happens, he can win the game with probability 1. Therefore

$$\Pr[\mathcal{A} \text{ wins} \mid \mathsf{EXTRACT}] = 1.$$

On the other hand, given that $\mathcal{A}$ does not break the zero-knowledge of P_{Auth}, the probability that $\mathcal{A}$ correctly guesses the capability sampled by the challenge is negligibly bigger than $1/2$ in both read and write algorithms, as argued above. Thus we can rewrite

$$\Pr[\mathcal{A} \text{ wins} \mid \neg\,\mathsf{EXTRACT}] \approx \frac{1}{2}.$$

It follows that we can bind the success probability of $\mathcal{A}$ as

$$\Pr[\mathsf{ExpAnon}_{\mathsf{GORAM}}^{\mathcal{A}}(\lambda, b) = 1] \approx$$
$$1 \cdot \Pr[\mathsf{EXTRACT}] + \tfrac{1}{2} \cdot \Pr[\neg \ \mathsf{EXTRACT}] \geq \tfrac{1}{2} + \epsilon(\lambda),$$

then we have

$$\Pr[\mathsf{EXTRACT}] + \frac{1}{2} \cdot (1 - \Pr[\mathsf{EXTRACT}]) \geq \frac{1}{2} + \epsilon(\lambda)$$
$$\frac{1}{2} \cdot \Pr[\mathsf{EXTRACT}] \geq \epsilon(\lambda)$$
$$\Pr[\mathsf{EXTRACT}] \geq 2 \cdot \epsilon(\lambda),$$

which is still a non-negligible value. This implies that $\mathcal{A}$ must be able to extract additional information from the proof without knowing the witnesses, which is clearly a contradiction with respect to the zero-knowledge property of $\mathcal{ZKP}$ and it concludes our proof. $\qquad \square$

$\hfill \square$

Proof of Theorem 3. In the following we separately prove the security of A-GORAM for each specified property.

Lemma 6. *Let* Π_{PE} *be an attribute-hiding predicate encryption scheme and* Π_{SE} *be a CPA-secure private-key encryption scheme. Then* A-GORAM *achieves* secrecy.

Proof of Lemma 6. The proof works analogously to Lemma 1. $\qquad \square$

Lemma 7. *Let* Π_{CHF} *be a collision-resistant, key-exposure free chameleon hash function and* Π_{DS} *be an existentially unforgeable digital signature scheme. Then* A-GORAM *achieves* accountable-integrity.

Proof of Lemma 7. The proof is conducted by splitting the success probability of the adversary and showing that such probability is upper bounded by a sum of negligible values. We first define the event COLL as the event in which the adversary is able to find a collision on the chameleon hash function Π_{CHF} without knowing the secret key, for the challenge entry of the experiment. We can express the probability that the adversary wins the experiment defined in Definition 21 as follows:

$$\Pr[\mathsf{ExpAcc}_{\mathsf{GORAM}}^{\mathcal{A}}(\lambda) = 1] =$$
$$\Pr[\mathcal{A} \ \mathsf{wins} \mid \mathsf{COLL}] \cdot \Pr[\mathsf{COLL}] + $$
$$\Pr[\mathcal{A} \ \mathsf{wins} \mid \neg \ \mathsf{COLL}] \cdot \Pr[\neg \ \mathsf{COLL}]$$

It is easy to see that whenever the event COLL happens the adversary can easily win the game by arbitrarily modifying the entry and computing a collision for the chameleon hash function. In the history of changes all of the versions of that entry will correctly verify and the challenger will not be able to blame any client in particular, thus making the adversary succeed with probability 1. By the above reasoning we can rewrite:

$$\Pr[\mathsf{ExpAcc}_{\mathsf{GORAM}}^{\mathcal{A}}(\lambda) = 1] =$$
$$\Pr[\mathsf{COLL}] + \Pr[\mathcal{A} \ \mathsf{wins} \mid \neg \ \mathsf{COLL}] \cdot \Pr[\neg \ \mathsf{COLL}]$$

We will now show that the probability that COLL happens is upper bounded by a negligible function. In order to do so we first define an intermediate game $\mathsf{ExpAcc`}_{\mathsf{GORAM}}^{\mathcal{A}}(\lambda)$, we then show that such a game is indistinguishable from the original and finally we prove that in this

latter experiment the probability of COLL is negligible. It directly follows that the probability of COLL is also negligible in the original experiment $\mathsf{ExpAcc}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda)$ since the two games are indistinguishable to the view of the adversary.

New Experiment. We define the intermediate game $\mathsf{ExpAcc'}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda)$ as follows:

Setup:. The challenger runs the Setup phase as in Definition 21. Additionally it sets a polynomial upper bound p on the number of queries to the interfaces addE and chMode and it picks a $q \in \{1...p\}$ uniformly at random.

Queries:. The challenger runs the Query phase as in Definition 21, except for the following interfaces:

(1) On input $\mathsf{addE}(\mathbf{a}, d)$ by $\mathcal{A}$, the challenger executes $\langle \mathcal{C}_{\mathsf{addE}}(cap_{\mathcal{O}}, \mathbf{a}, d), \mathcal{S}_{\mathsf{addE}}(\mathsf{DB}) \rangle$ locally. If it holds that for all corrupted i and for the new entry index j, $\mathbf{AC}(i, j) \neq rw$ and the query is the q-th query, then the new entry is computed with $c_{\mathsf{Auth}} \leftarrow \mathsf{PrEnc}(mpk, x_w, s)$ where s is a random string such that $|s| = |csk|$.

(2) On input $\mathsf{chMode}(\mathbf{a}, j)$ by $\mathcal{A}$, the challenger executes $\langle \mathcal{C}_{\mathsf{chMode}}(cap_{\mathcal{O}}, \mathbf{a}, j), \mathcal{S}_{\mathsf{chMode}}(\mathsf{DB}) \rangle$ locally. If it holds that for all corrupted i and for the new entry index j, $\mathbf{AC}(i, j) \neq rw$ and the query is the q-th query, then the new entry is computed with $c_{\mathsf{Auth}} \leftarrow \mathsf{PrEnc}(mpk, x_w, s)$ where s is a random string such that $|s| = |csk|$. On the other hand, if there exists a corrupted i such that $\mathbf{AC}(i, j) = rw$, c_{Auth} is reverted to be consistent with the entry structure.

Challenge. The challenger runs the Challenge phase as in Definition 21.

Output. The challenger runs the Output phase as in Definition 21.

$\mathsf{ExpAcc}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda) \approx \mathsf{ExpAcc'}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda).$ We prove the claim with a reduction against the attribute-hiding property of the predicate-encryption scheme Π_{PE}. That is, given an adversary $\mathcal{A}$ that can efficiently distinguish the two games, we create a simulation $\mathcal{B}$ that breaks the attribute-hiding property with the same probability, thus such an adversary cannot exist. The reduction is depicted below.

Setup. $\mathcal{B}$ receives as input the security parameter 1^{λ} from the challenger and it forwards it to $\mathcal{A}$. $\mathcal{B}$ initializes uniformly at random an attribute I and it sends to the challenger the tuple (I, I), who replies with the public key of the predicate-encryption scheme mpk^*. $\mathcal{B}$ then runs $(cap_{\mathcal{O}}, \mathsf{DB}) \leftarrow \mathsf{gen}(1^{\lambda})$ without $\mathsf{PrGen}(1^{\lambda})$, setting $mpk = mpk^*$ instead. Subsequently $\mathcal{B}$ gives pk_O to $\mathcal{A}$. Finally, it sets a polynomial upper bound p on the number of queries to the interfaces addE and chMode and it picks a $q \in \{1...p\}$ uniformly at random.

Queries. $\mathcal{B}$ provides then $\mathcal{A}$ with the following interfaces:

(1) On input $\mathsf{addCl}(\mathbf{a})$ by $\mathcal{A}$, $\mathcal{B}$ initializes a predicate f such that $f(I) = \perp$ and $\forall j \in \mathsf{DB}$ it holds that $f(I_j) = 0$ whenever $\mathbf{a}(j) = \perp$ and $f(I_j) = 1$ otherwise.

(2) On input $\mathsf{addE}(\mathbf{a}, d)$ by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{addE}}(cap_{\mathcal{O}}, \mathbf{a}, d), \mathcal{S}_{\mathsf{addE}}(\mathsf{DB}) \rangle$ locally. If it holds that for all corrupted i and for the new entry index j, $\mathbf{AC}(i, j) \neq rw$ and the query is the q-th query, then $\mathcal{B}$ sends to the challenger the pair $(csk, s) = (m_0, m_1)$ where csk is the chameleon secret key relative to that entry and s is a random string such that $|s| = |csk|$. The challenger replies with $c^* \leftarrow \mathsf{PrEnc}(mpk, I, m_b)$ and $\mathcal{B}$ sets $c_{\mathsf{Auth}} = c^*$ for the target entry.

(3) On input $\mathsf{chMode}(\mathbf{a}, j)$ by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{chMode}}(cap_{\mathcal{O}}, \mathbf{a}, j), \mathcal{S}_{\mathsf{chMode}}(\mathsf{DB}) \rangle$ locally. If it holds that for all corrupted i and for the new entry index j, $\mathbf{AC}(i, j) \neq rw$ and the query is the q-th query, then $\mathcal{B}$ sends to the challenger the pair $(csk, s) = (m_0, m_1)$ where csk is the chameleon secret key relative to that entry and s is a random string such that $|s| = |csk|$.

The challenger replies with $c^* \leftarrow \mathsf{PrEnc}(mpk, I, m_b)$ and $\mathcal{B}$ sets $c_{\mathsf{Auth}} = c^*$ for the target entry. On the other hand, if there exists a corrupted i such that $\mathbf{AC}(i, j) = rw$, c_{Auth} is reverted to be consistent with the entry structure.

(4) On input $\mathsf{corCl}(i)$ by $\mathcal{A}$, $\mathcal{B}$ queries the oracle provided by the challenger on the relative predicate f_i so to retrieve the corresponding key osk_{f_i}. $\mathcal{B}$ constructs cap_i using such a key, which is then handed over to $\mathcal{A}$.

(5) On input $\mathsf{read}(i, j)$ by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{read}}(cap_i, j), \mathcal{S}_{\mathsf{read}}(\mathsf{DB}) \rangle$ locally or in interaction with $\mathcal{A}$, depending whether i is corrupted.

(6) On input $\mathsf{write}(i, j, d)$ by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{write}}(cap_i, j, d), \mathcal{S}_{\mathsf{write}}(\mathsf{DB}) \rangle$ locally or in interaction with $\mathcal{A}$, depending whether i is corrupted.

Challenge. Finally, $\mathcal{A}$ outputs an index j^* which he wants to be challenged on. If there exists a capability cap_i provided to $\mathcal{A}$ such that $\mathbf{AC}(i, j^*) = rw$, then $\mathcal{B}$ aborts. $\mathcal{B}$ runs $d^* \leftarrow \langle \mathcal{C}_{\mathsf{read}}(cap_{\mathcal{O}}, j^*), \mathcal{S}_{\mathsf{read}}(\mathsf{DB}) \rangle$ and $L \leftarrow \langle \mathsf{blame}(cap_{\mathcal{O}}, \mathsf{Log}, j^*) \rangle$ locally.

Output. $\mathcal{B}$ sends to $\mathcal{A}$ 1 if and only if $d^* \neq \mathsf{DB}'(j^*)$ and $\exists\, i \in L$ that has not been queried by $\mathcal{A}$ to the interface $\mathsf{corCl}(\cdot)$ or $L = []$. At any point of the execution $\mathcal{A}$ can output 0 or 1 depending on his guess about which game he is facing, $\mathcal{B}$ simply forwards such a bit to the challenger and it stops the simulation.

The simulation above it is clearly efficient, also it is easy to see that whenever the challenger samples its internal coin $b = 0$, the simulation of $\mathcal{B}$ perfectly reproduces $\mathsf{ExpAcc}_{\mathsf{GORAM}}^{\mathcal{A}}(\lambda)$, thus:

$$\Pr[\mathcal{B} \mapsto 1 | b = 0] \approx \Pr[\mathcal{A} \mapsto 1 | b = 0].$$

Instead, whenever $b = 1$ the protocol executed by $\mathcal{B}$ perfectly simulates $\mathsf{ExpAcc}`_{\mathsf{GORAM}}^{\mathcal{A}}(\lambda)$,

$$\Pr[\mathcal{B} \mapsto 1 | b = 1] \approx \Pr[\mathcal{A} \mapsto 1 | b = 1].$$

By our initial assumption $\mathcal{A}$ was able to distinguish between the two games with non-negligible probability, therefore the probability carries over

$$|\Pr[\mathcal{B} \mapsto 1 | b = 0] - \Pr[\mathcal{B} \mapsto 1 | b = 1]| \geq \epsilon(\lambda),$$

which is clearly a contradiction to the attribute-hiding property of Π_{PE} and it proves our lemma.

$\Pr[\mathsf{COLL}]$ *in* $\mathsf{ExpAcc}`_{\mathsf{GORAM}}^{\mathcal{A}}(\lambda) \leq negl(\lambda)$. We demonstrate the claim via a reduction against the property of collision-resistance with key-exposure freeness of the chameleon hash function. Assume towards contradiction that there exists an adversary $\mathcal{A}$ such that the event COLL happens in $\mathsf{ExpAcc}`_{\mathsf{GORAM}}^{\mathcal{A}}(\lambda)$ with non-negligible probability, we build the following algorithm $\mathcal{B}$ to efficiently break the collision-resistance with key-exposure freeness property:

Setup. $\mathcal{B}$ sets a polynomial upper bound p on the number of queries to the interfaces addE and chMode and it picks a $q \in \{1...p\}$ uniformly at random. Then it runs $(cap_{\mathcal{O}}, \mathsf{DB}) \leftarrow \mathsf{gen}(1^{\lambda})$ and it hands over pk_O to $\mathcal{A}$.

Queries. $\mathcal{B}$ provides then $\mathcal{A}$ with the following interfaces:

(1) On input $\mathsf{addCl}(\mathbf{a})$ by $\mathcal{A}$, $\mathcal{B}$ executes $\mathsf{addCl}(cap_{\mathcal{O}}, \mathbf{a})$ locally and stores the capability cap_i returned by the algorithm.

(2) On input $\mathsf{addE}(\mathbf{a}, d)$ by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{addE}}(cap_{\mathcal{O}}, \mathbf{a}, d), \mathcal{S}_{\mathsf{addE}}(\mathsf{DB}) \rangle$ locally. If it holds that for all corrupted i and for the new entry index j, $\mathbf{AC}(i, j) \neq rw$ and the query is the q-th query, then the new entry is computed with $c_{\mathsf{Auth}} \leftarrow \mathsf{PrEnc}(mpk, x_w, s)$ where s is a random string such that $|s| = |csk|$.

(3) On input $\mathsf{chMode}(\mathbf{a}, j)$ by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{chMode}}(cap_{\mathcal{O}}, \mathbf{a}, j), \mathcal{S}_{\mathsf{chMode}}(\mathsf{DB}) \rangle$ locally. If it holds that for all corrupted i and for the new entry index j, $\mathbf{AC}(i, j) \neq rw$ and the query is the q-th query, then the new entry is computed with $c_{\mathsf{Auth}} \leftarrow \mathsf{PrEnc}(mpk, x_w, s)$ where s is a random string such that $|s| = |csk|$. On the other hand, if there exists a corrupted i such that $\mathbf{AC}(i, j) = rw$, c_{Auth} is reverted to be consistent with the entry structure.

(4) On input $\mathsf{corCl}(i)$ by $\mathcal{A}$, $\mathcal{B}$ hands over the capability cap_i related to the i-th client in the access control matrix $\mathbf{AC}$.

(5) On input $\mathsf{read}(i, j)$ by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{read}}(cap_i, j), \mathcal{S}_{\mathsf{read}}(\mathsf{DB}) \rangle$ locally or in interaction with $\mathcal{A}$, depending whether i is corrupted.

(6) On input $\mathsf{write}(i, j, d)$ by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{write}}(cap_i, j, d), \mathcal{S}_{\mathsf{write}}(\mathsf{DB}) \rangle$ locally or in interaction with $\mathcal{A}$, depending whether i is corrupted.

Challenge. Finally, $\mathcal{A}$ outputs an index j^* which he wants to be challenged on. If there exists a capability cap_i provided to $\mathcal{A}$ such that $\mathbf{AC}(i, j^*) = rw$, then the $\mathcal{B}$ aborts. It then runs $d^* \leftarrow \langle \mathcal{C}_{\mathsf{read}}(cap_{\mathcal{O}}, j^*), \mathcal{S}_{\mathsf{read}}(\mathsf{DB}) \rangle$ and $L \leftarrow \langle \mathsf{blame}(cap_{\mathcal{O}}, \mathsf{Log}, j^*) \rangle$ locally.

Output. $\mathcal{B}$ outputs 1 if and only if $d^* \neq \mathsf{DB}'(j^*)$ and $\exists i \in L$ that has not been queried by $\mathcal{A}$ to the interface $\mathsf{corCl}(\cdot)$ or $L = []$.

The simulation is efficient and it perfectly reproduces the game that $\mathcal{A}$ is expecting. Note that by assumption we have that COLL happens with probability $\epsilon(\lambda)$, thus it must be the case that the adversary was able to compute a collision of the chameleon hash function in the challenge entry with non-negligible probability. Note that $\mathcal{B}$ selects the challenge entry for storing s in c_{Auth} with probability at least $\frac{1}{p}$. Thus, with probability at least $\frac{1}{p} \cdot \epsilon(\lambda)$ $\mathcal{A}$ was able to compute a collision without having any information on the secret key csk. The probability is still non-negligible, therefore this constitutes a contradiction to the collision resistance with key-exposure freeness of Π_{CHF}. This proves our lemma.

We have demonstrated that the event COLL does not occur with more than negligible probability, therefore we can rewrite the total success probability of the adversary as follows:

$$\Pr[\mathsf{ExpAcc}_{\mathsf{GORAM}}^{\mathcal{A}}(\lambda) = 1] =$$
$$negl(\lambda) + \Pr[\mathcal{A} \text{ wins} \mid \neg \mathsf{COLL}] \cdot (1 - negl(\lambda)) \approx .$$
$$\Pr[\mathcal{A} \text{ wins} \mid \neg \mathsf{COLL}]$$

Thus, what is left to show is that the success probability of the adversary given that he is not able to compute a collision for Π_{CHF}, is at most a negligible value in the security parameter. This is demonstrated through a reduction against the existential unforgeability of the digital signature scheme Π_{DS}. Assuming towards contradiction that there exists an adversary $\mathcal{A}$ such that $\Pr[\mathcal{A} \text{ wins} \mid \neg \mathsf{COLL}] \geq \epsilon(\lambda)$ we can build a reduction $\mathcal{B}$ against the existential unforgeability $\mathsf{Exp}_{\mathcal{A}, \Pi_{\mathsf{DS}}}^{\mathsf{eu}}$ of Π_{DS} as follows:

Setup. $\mathcal{B}$ receives as input the security parameter 1^λ and the verification key vk^*. It runs $(cap_{\mathcal{O}}, \mathsf{DB}) \leftarrow \mathsf{gen}(1^\lambda)$ setting $vk = vk^*$ and it hands over $pk_{\mathcal{O}}$ to $\mathcal{A}$.

Queries. $\mathcal{B}$ provides then $\mathcal{A}$ with the following interfaces:

(1) On input $\mathsf{addCl}(\mathbf{a})$ by $\mathcal{A}$, $\mathcal{B}$ executes $\mathsf{addCl}(cap_{\mathcal{O}}, \mathbf{a})$ locally and stores the capability cap_i returned by the algorithm.

(2) On input $\mathsf{addE}(\mathbf{a}, d)$ by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{addE}}(cap_{\mathcal{O}}, \mathbf{a}, d), \mathcal{S}_{\mathsf{addE}}(\mathsf{DB}) \rangle$ locally. In order to compute the correct signature on the respective chameleon hash tag t, $\mathcal{B}$ queries the signing oracle provided by the challenger and retrieves the signature tag σ.

(3) On input $\mathsf{chMode}(\mathbf{a}, j)$ by $\mathcal{A}$, the challenger executes $\langle \mathcal{C}_{\mathsf{chMode}}(cap_{\mathcal{O}}, \mathbf{a}, j), \mathcal{S}_{\mathsf{chMode}}(\mathsf{DB}) \rangle$ locally.

(4) On input $\mathsf{corCl}(i)$ by $\mathcal{A}$, $\mathcal{B}$ hands over the capability cap_i related to the i-th client in the access control matrix $\mathbf{AC}$.

(5) On input $\mathsf{read}(i, j)$ by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{read}}(cap_i, j), \mathcal{S}_{\mathsf{read}}(\mathsf{DB}) \rangle$ locally or in interaction with $\mathcal{A}$, depending whether i is corrupted.

(6) On input $\mathsf{write}(i, j, d)$ by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{write}}(cap_i, j, d), \mathcal{S}_{\mathsf{write}}(\mathsf{DB}) \rangle$ locally or in interaction with $\mathcal{A}$, depending whether i is corrupted.

Challenge. Finally, $\mathcal{A}$ outputs an index j^* which he wants to be challenged on. $\mathcal{B}$ parses the Log to search one version of that entry that contains a pair (t', σ') that has not been queried to the signing oracle and such that $\mathsf{verify}(vk, t', \sigma') = 1$.

Output. $\mathcal{B}$ outputs such a pair (t', σ') and it interrupts the simulation.

The simulation is clearly efficient. It is easy to see that, in order to win the game, $\mathcal{A}$ must be able to change an entry without writing permission and by leaving it in a consistent state. This can be done by either computing a collision in the chameleon hash function or by forging a valid signature on the tag t. By assumption we ruled out the first hypothesis, thus the winning condition of the adversary implies the forge of a verifying message-signature pair. Note that the $\mathcal{A}$ could also just roll back to some previous version of the entry but this can be easily prevented by including some timestamp in the computation of the chameleon hash. The winning probability of the reduction then carries over:

$$\Pr[\mathcal{B} \text{ wins}] \approx \Pr[\mathcal{A} \text{ wins} \mid \neg \text{ COLL}] \geq \epsilon(\lambda).$$

In this way we built an efficient adversary $\mathcal{B}$ that breaks the existential unforgeability of Π_{DS} with non negligible probability, which is clearly a contradiction. Therefore it must hold that $\Pr[\mathcal{A} \text{ wins} \mid \neg \text{ COLL}]$ is a negligible function in the security parameter. Finally we have that:

$$\Pr[\mathsf{ExpAcc}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda) = 1] \approx$$
$$\Pr[\mathcal{A} \text{ wins} \mid \neg \text{ COLL}] \leq negl(\lambda) \quad,$$

which concludes our proof. $\qquad\square$

Lemma 8. *Let $\mathcal{ZKP}$ be a zero-knowledge proof system and Π_{SE} be a CPA-secure private-key encryption scheme. Then* A-GORAM *achieves obliviousness.*

Proof of Lemma 8. The proof works analogously to Lemma 4. $\qquad\square$

$\qquad\square$

Proof of Theorem 4. In the following we separately prove the security of S-GORAM for each specified property.

Lemma 9. *Let Π_{BE} be an be an adaptively secure broadcast encryption scheme and Π_{SE} be a CPA-secure private-key encryption scheme. Then* S-GORAM *achieves secrecy.*

Proof of Lemma 9. The proof is constructed by fixing the choice of the challenger over the sampling of the random coin and define intermediate hybrid games, where the two extremes are the $\mathsf{ExpSec}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda, b)$ experiment over the two values of b. We start by defining a new experiment $\mathsf{ExpSec}'^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda, b)$ that slightly differs from the original one. For sake of readability we introduce the following notation:

- GAME 1 := $\mathsf{ExpSec}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda, 0)$

- GAME 2 := $\mathsf{ExpSec'}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda, 0)$

- GAME 3 := $\mathsf{ExpSec'}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda, 1)$

- GAME 4 := $\mathsf{ExpSec}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda, 1)$

Then we show that the difference among any two neighboring games is bounded by a negligible value in the security parameter, therefore the advantage of the adversary in $\mathsf{ExpSec}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda, b)$ turns to be a sum of negligible values, which is still negligible. In particular we demonstrate the following:

$$\text{GAME } 1 \approx \text{GAME } 2 \approx \text{GAME } 3 \approx \text{GAME } 4$$

New Experiment. We define $\mathsf{ExpSec'}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda, b)$ as the following game:

Setup. The challenger runs $(cap_{\mathcal{O}}, \mathsf{DB}) \leftarrow \mathsf{gen}(1^{\lambda})$, sets $\mathbf{AC} := []$, and runs a black-box simulation of $\mathcal{A}$ to which it hands over DB.

Queries. The challenger provides $\mathcal{A}$ with an addCl, an addE, a chMode, a read, a write, and a corCl interactive interface that $\mathcal{A}$ may query adaptively and in any order. Each round $\mathcal{A}$ can query exactly one interface. These interfaces are described below:

(1) On input $\mathsf{addCl}(\mathbf{a})$ by $\mathcal{A}$, the challenger executes $\mathsf{addCl}(cap_{\mathcal{O}}, \mathbf{a})$ locally and stores the capability cap_i returned by the algorithm.

(2) On input $\mathsf{addE}(\mathbf{a}, d)$ by $\mathcal{A}$, the challenger executes $\langle \mathcal{C}_{\mathsf{addE}}(cap_{\mathcal{O}}, \mathbf{a}, d), \mathcal{S}_{\mathsf{addE}}(\mathsf{DB}) \rangle$ in interaction with $\mathcal{A}$, where the former plays the role of the client while the latter plays the role of the server.

(3) On input $\mathsf{chMode}(\mathbf{a}, j)$ by $\mathcal{A}$, the challenger executes $\langle \mathcal{C}_{\mathsf{chMode}}(cap_{\mathcal{O}}, \mathbf{a}, j), \mathcal{S}_{\mathsf{chMode}}(\mathsf{DB}) \rangle$ in interaction with $\mathcal{A}$.

(4) On input $\mathsf{corCl}(i)$ by $\mathcal{A}$, the challenger hands over the capability cap_i related to the i-th client in the access control matrix $\mathbf{AC}$.

(5) On input $\mathsf{read}(i, j)$ by $\mathcal{A}$, the challenger executes $\langle \mathcal{C}_{\mathsf{read}}(cap_i, j), \mathcal{S}_{\mathsf{read}}(\mathsf{DB}) \rangle$ in interaction with $\mathcal{A}$.

(6) On input $\mathsf{write}(i, j, d)$ by $\mathcal{A}$, the challenger executes $\langle \mathcal{C}_{\mathsf{write}}(cap_i, j, d), \mathcal{S}_{\mathsf{write}}(\mathsf{DB}) \rangle$ in interaction with $\mathcal{A}$.

Challenge. Finally, $\mathcal{A}$ outputs $(j, (d_0, d_1))$, where j is an index denoting the database entry on which $\mathcal{A}$ wishes to be challenged and (d_0, d_1) is a pair of entries such that $|d_0| = |d_1|$. The challenger accepts the request only if $\mathbf{AC}(i, j) = \bot$, for every i corrupted by $\mathcal{A}$ in the query phase. Afterwards, the challenger invokes $\langle \mathcal{C}_{\mathsf{write}}(cap_{\mathcal{O}}, j, d_b), \mathcal{S}_{\mathsf{write}}(\mathsf{DB}) \rangle$ in interaction with $\mathcal{A}$, as explained in Section 5, with the difference that the new entry is computed as follows:

$$E'_j = \begin{pmatrix} c'_{1,j} \leftarrow \mathcal{E}(\mathcal{K}, j) \\ c'_{2,j} \leftarrow \mathcal{E}(\mathcal{K}, \mathsf{BrEnc}(bpk_w, W^*, csk)) \\ c'_{3,j} \leftarrow \mathcal{E}(\mathcal{K}, \mathsf{BrEnc}(bpk_r, R^*, K^*)) \\ c'_{4,j} \leftarrow \mathcal{E}(\mathcal{K}, c'_{\mathsf{Data}}) \end{pmatrix}$$

where R^* is the subset of users having read access on the entry j, W^* is the subset of users having write access on it and K^* is a random string such that $|K^*| = |k_j|$ and in particular it is not necessarily the same key used to encrypt c'_{Data}.

Output. In the output phase $\mathcal{A}$ still has access to the interfaces except for addCl on input $\mathbf{a}$ such that $\mathbf{a}(j) \neq \bot$; corCl on input i such that $\mathbf{AC}(i,j) \neq \bot$; and chMode on input $\mathbf{a}$ and j with $\mathbf{a}(i) \neq \bot$ for some previously corrupted client i. Eventually, $\mathcal{A}$ stops, outputting a bit b'. The challenger outputs 1 if and only if $b = b'$.

GAME 1 $\approx$ GAME 2. We assume toward contradiction that there exists a PPT adversary $\mathcal{A}$ that is able to distinguish among GAME 1 and GAME 2 with non-negligible probability, namely:

$$|\Pr[\mathsf{GAME\ 1} = 1] - \Pr[\mathsf{GAME\ 2} = 1]| \geq \epsilon(\lambda)$$

for some non-negligible $\epsilon(\lambda)$. Then we show that we can use such an adversary to build the following reduction $\mathcal{B}$ against the adaptive-security property of the broadcast encryption scheme Π_{BE} defined in Definition 11. The simulation is elaborated below.

Setup. $\mathcal{B}$ receives as input the security parameter 1^λ and the public key bpk^* from the challenger and it forwards 1^λ to $\mathcal{A}$. $\mathcal{B}$ runs $(cap_\mathcal{O}, \mathsf{DB}) \leftarrow \mathsf{gen}(1^\lambda)$ as described in Section 5 without $\mathsf{Setup}_{\mathsf{BE}}(1^\lambda)$, setting $bpk_r = bpk^*$ instead. Additionally, $\mathcal{B}$ initializes an empty set S of clients. Finally $\mathcal{B}$ gives pk_O and DB to $\mathcal{A}$.

Queries. $\mathcal{B}$ provides then $\mathcal{A}$ with the following interfaces:

(1) On input $\mathsf{addCl}(\mathbf{a})$ by $\mathcal{A}$, $\mathcal{B}$ adds one client to the set S of clients.

(2) On input $\mathsf{addE}(\mathbf{a}, d)$ by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{addE}}(cap_\mathcal{O}, \mathbf{a}, d), \mathcal{S}_{\mathsf{addE}}(\mathsf{DB}) \rangle$ in interaction with $\mathcal{A}$.

(3) On input $\mathsf{chMode}(\mathbf{a}, j)$ by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{chMode}}(cap_\mathcal{O}, \mathbf{a}, j), \mathcal{S}_{\mathsf{chMode}}(\mathsf{DB}) \rangle$ in interaction with $\mathcal{A}$.

(4) On input $\mathsf{corCl}(i)$ by $\mathcal{A}$, $\mathcal{B}$ queries the oracle provided by the challenger on i so to retrieve the corresponding key bsk_i. $\mathcal{B}$ constructs cap_i using such a key, which is handed over to $\mathcal{A}$.

(5) On input $\mathsf{read}(i, j)$ by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{read}}(cap_i, j), \mathcal{S}_{\mathsf{read}}(\mathsf{DB}) \rangle$ in interaction with $\mathcal{A}$.

(6) On input $\mathsf{write}(i, j, d)$ by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{write}}(cap_i, j, d), \mathcal{S}_{\mathsf{write}}(\mathsf{DB}) \rangle$ in interaction with $\mathcal{A}$.

Challenge. Finally, $\mathcal{A}$ outputs $(j, (d_0, d_1))$, where j is an index denoting the database entry on which $\mathcal{A}$ wishes to be challenged and (d_0, d_1) is a pair of entries such that $|d_0| = |d_1|$. $\mathcal{B}$ accepts the tuple only if $\mathbf{AC}(i,j) = \bot$, for every i corrupted by $\mathcal{A}$ in the query phase. $\mathcal{B}$ sets S^* to be the set of clients that have access to the j-th entry and it sends it to the challenger, who replies with the tuple (Hdr^*, K^*). $\mathcal{B}$ then executes $\langle \mathcal{C}_{\mathsf{write}}(cap_\mathcal{O}, j, d_0), \mathcal{S}_{\mathsf{write}}(\mathsf{DB}) \rangle$, computing the new entry in the following manner:

$$E'_j = \begin{pmatrix} c'_{1,j} \leftarrow \mathcal{E}(\mathcal{K}, j) \\ c'_{2,j} \leftarrow \mathcal{E}(\mathcal{K}, \mathsf{BrEnc}(bpk_w, W^*, csk)) \\ c'_{3,j} \leftarrow \mathcal{E}(\mathcal{K}, Hdr^*) \\ c'_{4,j} \leftarrow \mathcal{E}(\mathcal{K}, \mathcal{E}(K^*, d_0)) \end{pmatrix}$$

Output. In the output phase $\mathcal{A}$ still has access to the interfaces except for addCl on input $\mathbf{a}$ such that $\mathbf{a}(j) \neq \bot$; corCl on input i such that $\mathbf{AC}(i,j) \neq \bot$; and chMode on input $\mathbf{a}, j$ with $\mathbf{a}(i) \neq \bot$ for some previously corrupted client i. Note that in case there exists some non-corrupted i such that $\mathbf{a}(i) \neq \bot$, $\mathcal{B}$ just simulates the $\langle \mathcal{C}_{\mathsf{chMode}}(cap_\mathcal{O}, \mathbf{a}, j), \mathcal{S}_{\mathsf{chMode}}(\mathsf{DB}) \rangle$ protocol by rerandomizing the challenge ciphertext rather than re-encrypting it. Eventually, $\mathcal{A}$ stops, outputting a bit b'. $\mathcal{B}$ outputs 1 if and only if $b' = 0$.

The simulation is clearly efficient, also it is easy to see that whenever the challenger samples $b = 0$, the simulation perfectly reproduces the inputs that $\mathcal{A}$ is expecting in GAME 1. The only

difference is indeed that in the challenge phase c'_{Data} is re-encrypted in the simulation, while in the real experiment it is just rerandomized; also in the output phase, the interface chMode on j is simulated with a rerandomization on Hdr^*, rather than a re-encryption on c_{Key}. By definition of rerandomization, however, these two operations are indistinguishable to $\mathcal{A}$. Note that we assume for simplicity the broadcast encryption scheme to be rerandomizable, however this feature is not strictly necessary since we could simulate the rerandomization by asking the challenger another challenge cipher. This does not affect the security of the scheme by standard hybrid argument. Thus, we can state the following:

$$\Pr[\mathcal{B} \mapsto 1 | b = 0] \approx \Pr[\mathsf{GAME\ 1} = 1].$$

On the other hand, in case the challenger initializes $b = 1$, then $\mathcal{B}$ perfectly simulates the environment that $\mathcal{A}$ is expecting in $\mathsf{GAME\ 2}$. Therefore we can assert that:

$$\Pr[\mathcal{B} \mapsto 1 | b = 1] \approx \Pr[\mathsf{GAME\ 2} = 1].$$

However, it follows from the initial assumption that

$$|\Pr[\mathcal{B} \mapsto 1 | b = 1] - \Pr[\mathcal{B} \mapsto 1 | b = 0]| \geq \epsilon(\lambda),$$

which is clearly a contradiction with respect to the adaptive-security property of the broadcast encryption scheme Π_{BE}, and it proves the initial lemma.

$\mathsf{GAME\ 2} \approx \mathsf{GAME\ 3}$. We assume toward contradiction that there exists a PPT adversary $\mathcal{A}$ that is able to distinguish among $\mathsf{GAME\ 2}$ and $\mathsf{GAME\ 3}$ with non-negligible probability, namely:

$$|\Pr[\mathsf{GAME\ 2} = 1] - \Pr[\mathsf{GAME\ 3} = 1]| \geq \epsilon(\lambda)$$

For some non-negligible $\epsilon(\lambda)$. Then we show that we can we can use such an adversary to build the following reduction $\mathcal{B}$ against the CPA-security property of the private-key encryption scheme Π_{SE}. The simulation is elaborated below.

Setup. $\mathcal{B}$ receives as input the security parameter 1^λ from the challenger and it forwards it to $\mathcal{A}$. $\mathcal{B}$ then runs $(cap_{\mathcal{O}}, \mathsf{DB}) \leftarrow \mathsf{gen}(1^\lambda)$ as described in Section 5 and it gives pk_O and DB to $\mathcal{A}$.

Queries. $\mathcal{B}$ provides then $\mathcal{A}$ with the following interfaces:

(1) On input $\mathsf{addCl}(\mathbf{a})$ by $\mathcal{A}$, $\mathcal{B}$ executes $\mathsf{addCl}(cap_{\mathcal{O}}, \mathbf{a})$ locally and it stores the capability cap_i returned by the algorithm.

(2) On input $\mathsf{addE}(\mathbf{a}, d)$ by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{addE}}(cap_{\mathcal{O}}, \mathbf{a}, d), \mathcal{S}_{\mathsf{addE}}(\mathsf{DB}) \rangle$ in interaction with $\mathcal{A}$.

(3) On input $\mathsf{chMode}(\mathbf{a}, j)$ by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{chMode}}(cap_{\mathcal{O}}, \mathbf{a}, j), \mathcal{S}_{\mathsf{chMode}}(\mathsf{DB}) \rangle$ in interaction with $\mathcal{A}$.

(4) On input $\mathsf{corCl}(i)$ by $\mathcal{A}$, $\mathcal{B}$ hands over the capability cap_i related to the i-th client in the access control matrix $\mathbf{AC}$.

(5) On input $\mathsf{read}(i, j)$ by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{read}}(cap_i, j), \mathcal{S}_{\mathsf{read}}(\mathsf{DB}) \rangle$ in interaction with $\mathcal{A}$.

(6) On input $\mathsf{write}(i, j, d)$ by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{write}}(cap_i, j, d), \mathcal{S}_{\mathsf{write}}(\mathsf{DB}) \rangle$ in interaction with $\mathcal{A}$.

Challenge. Finally, $\mathcal{A}$ outputs $(j, (d_0, d_1))$, where j is an index denoting the database entry on which $\mathcal{A}$ wishes to be challenged and (d_0, d_1) is a pair of entries such that $|d_0| = |d_1|$. $\mathcal{B}$ accepts the tuple only if $\mathbf{AC}(i, j) = \bot$, for every i corrupted by $\mathcal{A}$ in the query phase. $\mathcal{B}$ sends the tuple $(m_0, m_1) = (d_0, d_1)$ to the challenger, who answers back with the challenge ciphertext $c^* \leftarrow$

$\mathcal{E}(k, d_b)$ that $\mathcal{B}$ uses to perform a local execution of $\langle \mathcal{C}_{\text{write}}(cap_{\mathcal{O}}, j, d_b), \mathcal{S}_{\text{write}}(\text{DB})\rangle$, computing the new entry in the following manner:

$$E'_j = \begin{pmatrix} c'_{1,j} \leftarrow \mathcal{E}(\mathcal{K}, j) \\ c'_{2,j} \leftarrow \mathcal{E}(\mathcal{K}, \text{BrEnc}(bpk_w, W^*, csk)) \\ c'_{3,j} \leftarrow \mathcal{E}(\mathcal{K}, \text{BrEnc}(bpk_r, R^*, K^*)) \\ c'_{4,j} \leftarrow \mathcal{E}(\mathcal{K}, c^*) \end{pmatrix}$$

where K^* is a random string such that $|K^*| = |k|$.

Output. In the output phase $\mathcal{A}$ still has access to the interfaces except for addCl on input $\mathbf{a}$ such that $\mathbf{a}(j) \neq \perp$; corCl on input i such that $\mathbf{AC}(i, j) \neq \perp$; and chMode on input $\mathbf{a}, j$ with $\mathbf{a}(i) \neq \perp$ for some previously corrupted client i. Eventually, $\mathcal{A}$ stops, outputting a bit b'. $\mathcal{B}$ outputs b' as well.

The simulation is obviously efficient, also it is easy to see that whenever the challenger samples $b = 0$, the simulation perfectly reproduces the inputs that $\mathcal{A}$ is expecting in GAME 2. Thus, we can state the following:

$$\Pr[\mathcal{B} \mapsto 0 | b = 0] \approx \Pr[\text{GAME 2} = 1].$$

On the other hand, in case the challenger initializes $b = 1$, then $\mathcal{B}$ perfectly simulates the environment that $\mathcal{A}$ is expecting in GAME 3. Therefore we can assert that:

$$\Pr[\mathcal{B} \mapsto 1 | b = 1] \approx \Pr[\text{GAME 3} = 1].$$

However, it follows from the initial assumption that:

$$|\Pr[\mathcal{B} \mapsto 1 | b = 1] - \Pr[\mathcal{B} \mapsto 0 | b = 0]| \geq \epsilon(\lambda),$$

$$\left| \Pr[\text{Exp}^{\text{cpa}}_{\mathcal{A}, \Pi_{\text{SE}}}(\lambda, 1) = 1] - \Pr[\text{Exp}^{\text{cpa}}_{\mathcal{A}, \Pi_{\text{SE}}}(\lambda, 0) = 1] \right|$$
$$\geq \epsilon(\lambda),$$

which implies a non-negligible difference in the success probability of $\mathcal{B}$ with respect to the random choice of b and it clearly represents a contradiction with respect to the CPA-security property of the private-key encryption scheme Π_{SE}. This proves the initial lemma.

GAME 3 $\approx$ GAME 4. The proof is conducted with the same pipeline of the indistinguishability between GAME 1 and GAME 2; the simulation also works correspondingly, except that in this case the reduction $\mathcal{B}$ encrypts the message d_1 rather than d_0. However, the analogous argument applies.

GAME 1 $\approx$ GAME 4. By the previous lemmas it directly follows that the difference among each couple of neighboring games is bounded by a negligible value, thus the difference between GAME 1 and GAME 4 is a sum of negligible terms, which is, again, negligible. In particular

$$\text{GAME 1} \approx \text{GAME 4}$$

directly implies that:

$$\text{ExpSec}^{\mathcal{A}}_{\text{GORAM}}(\lambda, 0) \approx \text{ExpSec}^{\mathcal{A}}_{\text{GORAM}}(\lambda, 1)$$

thus, $\forall$ PPT adversary the two experiments look indistinguishable. This concludes our proof. $\square$

Lemma 10. *Let Π_{CHF} be a collision-resistant, key-exposure free chameleon hash function and Π_{DS} be an existentially unforgeable digital signature scheme. Then S-GORAM achieves accountable-integrity.*

Proof of Lemma 10. The proof is conducted by splitting the success probability of the adversary and showing that such probability is upper bounded by a sum of negligible values. We first define the event COLL as the event in which the adversary is able to find a collision on the chameleon hash function Π_{CHF} without knowing the secret key, for the challenge entry of the experiment. We can express the probability that the adversary wins the experiment defined in Definition 21 as follows:

$$\Pr[\mathsf{ExpAcc}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda) = 1] =$$
$$\Pr[\mathcal{A} \text{ wins} \mid \mathsf{COLL}] \cdot \Pr[\mathsf{COLL}]+$$
$$\Pr[\mathcal{A} \text{ wins} \mid \neg \mathsf{COLL}] \cdot \Pr[\neg \mathsf{COLL}]$$

It is easy to see that whenever the event COLL happens the adversary can easily win the game by arbitrarily modifying the entry and computing a collision for the chameleon hash function. In the history of changes all of the versions of that entry will correctly verify and the challenger will not be able to blame any client in particular, thus making the adversary succeed with probability 1. By the above reasoning we can rewrite:

$$\Pr[\mathsf{ExpAcc}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda) = 1] =$$
$$\Pr[\mathsf{COLL}] + \Pr[\mathcal{A} \text{ wins} \mid \neg \mathsf{COLL}] \cdot \Pr[\neg \mathsf{COLL}]$$

We will now show that the probability that COLL happens is upper bounded by a negligible function. In order to do so we first define an intermediate game $\mathsf{ExpAcc}^{\prime \mathcal{A}}_{\mathsf{GORAM}}(\lambda)$, we then show that such a game is indistinguishable from the original and finally we prove that in this latter experiment the probability of COLL is negligible. It directly follows that the probability of COLL is also negligible in the original experiment $\mathsf{ExpAcc}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda)$ since the two games are indistinguishable to the view of the adversary.

New Experiment. We define the intermediate game $\mathsf{ExpAcc}^{\prime \mathcal{A}}_{\mathsf{GORAM}}(\lambda)$ as follows:
Setup:. The challenger runs the Setup phase as in Definition 21. Additionally it sets a polynomial upper bound p on the number of queries to the interfaces addE and chMode and it picks a $q \in \{1...p\}$ uniformly at random.
Queries:. The challenger runs the Query phase as in Definition 21, except for the following interfaces:

(1) On input $\mathsf{addE}(\mathbf{a}, d)$ by $\mathcal{A}$, the challenger executes $\langle \mathcal{C}_{\mathsf{addE}}(cap_{\mathcal{O}}, \mathbf{a}, d), \mathcal{S}_{\mathsf{addE}}(\mathsf{DB}) \rangle$ locally. If it holds that for all corrupted i and for the new entry index j, $\mathbf{AC}(i, j) \neq rw$ and the query is the q-th query, then the new entry is computed with $c_{\mathsf{Auth}} \leftarrow \mathsf{BrEnc}(bpk_w, W^*, K^*)$ where K^* is a random string such that $|K^*| = |csk_j|$.

(2) On input $\mathsf{chMode}(\mathbf{a}, j)$ by $\mathcal{A}$, the challenger executes $\langle \mathcal{C}_{\mathsf{chMode}}(cap_{\mathcal{O}}, \mathbf{a}, j), \mathcal{S}_{\mathsf{chMode}}(\mathsf{DB}) \rangle$ locally. If it holds that for all corrupted i and for the new entry index j, $\mathbf{AC}(i, j) \neq rw$ and the query is the q-th query, then the new entry is computed with $c_{\mathsf{Auth}} \leftarrow \mathsf{BrEnc}(bpk_w, W^*, K^*)$ where K^* is a random string such that $|K^*| = |csk_j|$. On the other hand, if there exists a corrupted i such that $\mathbf{AC}(i, j) = rw$, c_{Auth} is reverted to be consistent with the entry structure.

Challenge. The challenger runs the Challenge phase as in Definition 21.
Output. The challenger runs the Output phase as in Definition 21.

$\mathsf{ExpAcc}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda) \approx \mathsf{ExpAcc}^{\prime \mathcal{A}}_{\mathsf{GORAM}}(\lambda)$. We prove the claim with a reduction against the adaptive-security property of the broadcast encryption scheme Π_{BE}. That is, given an adversary $\mathcal{A}$ that can efficiently distinguish the two games, we create a simulation $\mathcal{B}$ that breaks

the adaptive-security property with the same probability, thus such an adversary cannot exist. The reduction is depicted below.

Setup. $\mathcal{B}$ receives as input the security parameter 1^λ and the public key bpk^* from the challenger and it forwards 1^λ to $\mathcal{A}$. $\mathcal{B}$ then runs $(cap_\mathcal{O}, \mathsf{DB}) \leftarrow \mathsf{gen}(1^\lambda)$ without $\mathsf{Setup}_{\mathsf{BE}}(1^\lambda)$, setting $bpk_w = bpk^*$ instead. Subsequently $\mathcal{B}$ initializes an empty set S of clients and it gives pk_O to $\mathcal{A}$. Finally, it sets a polynomial upper bound p on the number of queries to the interfaces addE and chMode and it picks a $q \in \{1...p\}$ uniformly at random.

Queries. $\mathcal{B}$ provides then $\mathcal{A}$ with the following interfaces:

(1) On input $\mathsf{addCl}(\mathbf{a})$ by $\mathcal{A}$, $\mathcal{B}$ adds one client to the set S of clients.

(2) On input $\mathsf{addE}(\mathbf{a}, d)$ by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{addE}}(cap_\mathcal{O}, \mathbf{a}, d), \mathcal{S}_{\mathsf{addE}}(\mathsf{DB}) \rangle$ locally. If it holds that for all corrupted i and for the new entry index j, $\mathbf{AC}(i, j) \neq rw$ and the query is the q-th query, then $\mathcal{B}$ sends to the challenger the set S^* of clients having write access to the j-th entry. The challenger replies with the tuple (Hdr^*, K^*) and $\mathcal{B}$ sets $c_{\mathsf{Auth}} \leftarrow Hdr^*$ and $csk_j \leftarrow K^*$ (cpk_j is also changed accordingly) for the target j-th entry.

(3) On input $\mathsf{chMode}(\mathbf{a}, j)$ by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{chMode}}(cap_\mathcal{O}, \mathbf{a}, j), \mathcal{S}_{\mathsf{chMode}}(\mathsf{DB}) \rangle$ locally. If it holds that for all corrupted i and for the new entry index j, $\mathbf{AC}(i, j) \neq rw$ and the query is the q-th query, then $\mathcal{B}$ sends to the challenger the set S^* of clients having write access to the j-th entry. The challenger replies with the tuple (Hdr^*, K^*) and $\mathcal{B}$ sets $c_{\mathsf{Auth}} \leftarrow Hdr^*$ and $csk_j \leftarrow K^*$ (cpk_j is also changed accordingly) for the target j-th entry. On the other hand, if there exists a corrupted i such that $\mathbf{AC}(i, j) = rw$, c_{Auth} is reverted to be consistent with the entry structure.

(4) On input $\mathsf{corCl}(i)$ by $\mathcal{A}$, $\mathcal{B}$ queries the oracle provided by the challenger on i so to retrieve the corresponding key bsk_i. $\mathcal{B}$ constructs cap_i using such a key, which is then handed over to $\mathcal{A}$.

(5) On input $\mathsf{read}(i, j)$ by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{read}}(cap_i, j), \mathcal{S}_{\mathsf{read}}(\mathsf{DB}) \rangle$ locally or in interaction with $\mathcal{A}$, depending whether i is corrupted.

(6) On input $\mathsf{write}(i, j, d)$ by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{write}}(cap_i, j, d), \mathcal{S}_{\mathsf{write}}(\mathsf{DB}) \rangle$ locally or in interaction with $\mathcal{A}$, depending whether i is corrupted.

Challenge. Finally, $\mathcal{A}$ outputs an index j^* which he wants to be challenged on. If there exists a capability cap_i provided to $\mathcal{A}$ such that $\mathbf{AC}(i, j^*) = rw$, then $\mathcal{B}$ aborts. $\mathcal{B}$ runs $d^* \leftarrow \langle \mathcal{C}_{\mathsf{read}}(cap_\mathcal{O}, j^*), \mathcal{S}_{\mathsf{read}}(\mathsf{DB}) \rangle$ and $L \leftarrow \langle \mathsf{blame}(cap_\mathcal{O}, \mathsf{Log}, j^*) \rangle$ locally.

Output. $\mathcal{B}$ sends to $\mathcal{A}$ 1 if and only if $d^* \neq \mathsf{DB}'(j^*)$ and $\exists i \in L$ that has not been queried by $\mathcal{A}$ to the interface $\mathsf{corCl}(\cdot)$ or $L = []$. At any point of the execution $\mathcal{A}$ can output 0 or 1 depending on his guess about which game he is facing, $\mathcal{B}$ simply forwards such a bit to the challenger and it stops the simulation.

The simulation above it is clearly efficient, also it is easy to see that whenever the challenger samples its internal coin $b = 0$, the simulation of $\mathcal{B}$ perfectly reproduces $\mathsf{ExpAcc}^\mathcal{A}_{\mathsf{GORAM}}(\lambda)$, thus:

$$\Pr[\mathcal{B} \mapsto 1 | b = 0] \approx \Pr[\mathcal{A} \mapsto 1 | b = 0].$$

Instead, whenever $b = 1$ the protocol executed by $\mathcal{B}$ perfectly simulates $\mathsf{ExpAcc}'^\mathcal{A}_{\mathsf{GORAM}}(\lambda)$,

$$\Pr[\mathcal{B} \mapsto 1 | b = 1] \approx \Pr[\mathcal{A} \mapsto 1 | b = 1].$$

By our initial assumption $\mathcal{A}$ was able to distinguish between the two games with non-negligible probability, therefore the probability carries over

$$|\Pr[\mathcal{B} \mapsto 1 | b = 0] - \Pr[\mathcal{B} \mapsto 1 | b = 1]| \geq \epsilon(\lambda),$$

which is clearly a contradiction to the adaptive-security property of Π_{BE} and it proves our lemma.

$\Pr[\mathsf{COLL}]$ *in* $\mathsf{ExpAcc`}_{\mathsf{GORAM}}^{\mathcal{A}}(\lambda) \leq negl(\lambda)$. We demonstrate the claim via a reduction against the property of collision-resistance with key-exposure freeness of the chameleon hash function. Assume towards contradiction that there exists an adversary $\mathcal{A}$ such that the event COLL happens in $\mathsf{ExpAcc`}_{\mathsf{GORAM}}^{\mathcal{A}}(\lambda)$ with non-negligible probability, we build the following algorithm $\mathcal{B}$ to efficiently break the collision-resistance with key-exposure freeness property:

Setup. $\mathcal{B}$ sets a polynomial upper bound p on the number of queries to the interfaces addE and chMode and it picks a $q \in \{1...p\}$ uniformly at random. Then it runs $(cap_{\mathcal{O}}, \mathsf{DB}) \leftarrow \mathsf{gen}(1^{\lambda})$ and it hands over pk_O to $\mathcal{A}$.

Queries. $\mathcal{B}$ provides then $\mathcal{A}$ with the following interfaces:

(1) On input $\mathsf{addCl}(\mathbf{a})$ by $\mathcal{A}$, $\mathcal{B}$ executes $\mathsf{addCl}(cap_{\mathcal{O}}, \mathbf{a})$ locally and stores the capability cap_i returned by the algorithm.

(2) On input $\mathsf{addE}(\mathbf{a}, d)$ by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{addE}}(cap_{\mathcal{O}}, \mathbf{a}, d), \mathcal{S}_{\mathsf{addE}}(\mathsf{DB}) \rangle$ locally. If it holds that for all corrupted i and for the new entry index j, $\mathbf{AC}(i, j) \neq rw$ and the query is the q-th query, then the new entry is computed with $c_{\mathsf{Auth}} \leftarrow \mathsf{BrEnc}(bpk_w, W^*, K^*)$ where K^* is a random string such that $|K^*| = |csk_j|$.

(3) On input $\mathsf{chMode}(\mathbf{a}, j)$ by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{chMode}}(cap_{\mathcal{O}}, \mathbf{a}, j), \mathcal{S}_{\mathsf{chMode}}(\mathsf{DB}) \rangle$ locally. If it holds that for all corrupted i and for the new entry index j, $\mathbf{AC}(i, j) \neq rw$ and the query is the q-th query, then the new entry is computed with $c_{\mathsf{Auth}} \leftarrow \mathsf{BrEnc}(bpk_w, W^*, K^*)$ where K^* is a random string such that $|K^*| = |csk_j|$. On the other hand, if there exists a corrupted i such that $\mathbf{AC}(i, j) = rw$, c_{Auth} is reverted to be consistent with the entry structure.

(4) On input $\mathsf{corCl}(i)$ by $\mathcal{A}$, $\mathcal{B}$ hands over the capability cap_i related to the i-th client in the access control matrix $\mathbf{AC}$.

(5) On input $\mathsf{read}(i, j)$ by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{read}}(cap_i, j), \mathcal{S}_{\mathsf{read}}(\mathsf{DB}) \rangle$ locally or in interaction with $\mathcal{A}$, depending whether i is corrupted.

(6) On input $\mathsf{write}(i, j, d)$ by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{write}}(cap_i, j, d), \mathcal{S}_{\mathsf{write}}(\mathsf{DB}) \rangle$ locally or in interaction with $\mathcal{A}$, depending whether i is corrupted.

Challenge. Finally, $\mathcal{A}$ outputs an index j^* which he wants to be challenged on. If there exists a capability cap_i provided to $\mathcal{A}$ such that $\mathbf{AC}(i, j^*) = rw$, then the $\mathcal{B}$ aborts. It then runs $d^* \leftarrow \langle \mathcal{C}_{\mathsf{read}}(cap_{\mathcal{O}}, j^*), \mathcal{S}_{\mathsf{read}}(\mathsf{DB}) \rangle$ and $L \leftarrow \langle \mathsf{blame}(cap_{\mathcal{O}}, \mathsf{Log}, j^*)$ locally.

Output. $\mathcal{B}$ outputs 1 if and only if $d^* \neq \mathsf{DB}'(j^*)$ and $\exists i \in L$ that has not been queried by $\mathcal{A}$ to the interface $\mathsf{corCl}(\cdot)$ or $L = []$.

The simulation is efficient and it perfectly reproduces the game that $\mathcal{A}$ is expecting. Note that by assumption we have that COLL happens with probability $\epsilon(\lambda)$, thus it must be the case that the adversary was able to compute a collision of the chameleon hash function in the challenge entry with non-negligible probability. Note that $\mathcal{B}$ selects the challenge entry for storing K^* in c_{Auth} with probability at least $\frac{1}{p}$. Thus, with probability at least $\frac{1}{p} \cdot \epsilon(\lambda)$ $\mathcal{A}$ was able to compute a collision without having any information on the secret key csk. The probability is still non-negligible, therefore this constitutes a contradiction to the collision resistance with key-exposure freeness of Π_{CHF}. This proves our lemma.

We have demonstrated that the event COLL does not occur with more than negligible probability, therefore we can rewrite the total success probability of the adversary as follows:

$$\Pr[\mathsf{ExpAcc}_{\mathsf{GORAM}}^{\mathcal{A}}(\lambda) = 1] =$$
$$negl(\lambda) + \Pr[\mathcal{A} \text{ wins} \mid \neg \mathsf{COLL}] \cdot (1 - negl(\lambda)) \approx .$$
$$\Pr[\mathcal{A} \text{ wins} \mid \neg \mathsf{COLL}]$$

Thus, what is left to show is that the success probability of the adversary given that he is not able to compute a collision for Π_{CHF}, is at most a negligible value in the security parameter. This is demonstrated through a reduction against the existential unforgeability of the digital signature scheme Π_{DS}. Assuming towards contradiction that there exists an adversary $\mathcal{A}$ such that $\Pr[\mathcal{A} \text{ wins} \mid \neg \mathsf{COLL}] \geq \epsilon(\lambda)$ we can build a reduction $\mathcal{B}$ against the existential unforgeability $\mathsf{Exp}_{\mathcal{A},\Pi_{\mathsf{DS}}}^{\mathsf{eu}}$ of Π_{DS} as follows:

Setup. $\mathcal{B}$ receives as input the security parameter 1^λ and the verification key vk^*. It runs $(cap_{\mathcal{O}}, \mathsf{DB}) \leftarrow \mathsf{gen}(1^\lambda)$ setting $vk = vk^*$ and it hands over pk_O to $\mathcal{A}$.

Queries. $\mathcal{B}$ provides then $\mathcal{A}$ with the following interfaces:

(1) On input $\mathsf{addCl}(\mathbf{a})$ by $\mathcal{A}$, $\mathcal{B}$ executes $\mathsf{addCl}(cap_{\mathcal{O}}, \mathbf{a})$ locally and stores the capability cap_i returned by the algorithm.

(2) On input $\mathsf{addE}(\mathbf{a}, d)$ by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{addE}}(cap_{\mathcal{O}}, \mathbf{a}, d), \mathcal{S}_{\mathsf{addE}}(\mathsf{DB}) \rangle$ locally. In order to compute the correct signature on the respective chameleon hash tag t, $\mathcal{B}$ queries the signing oracle provided by the challenger and retrieves the signature tag σ.

(3) On input $\mathsf{chMode}(\mathbf{a}, j)$ by $\mathcal{A}$, the challenger executes $\langle \mathcal{C}_{\mathsf{chMode}}(cap_{\mathcal{O}}, \mathbf{a}, j), \mathcal{S}_{\mathsf{chMode}}(\mathsf{DB}) \rangle$ locally.

(4) On input $\mathsf{corCl}(i)$ by $\mathcal{A}$, $\mathcal{B}$ hands over the capability cap_i related to the i-th client in the access control matrix $\mathbf{AC}$.

(5) On input $\mathsf{read}(i, j)$ by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{read}}(cap_i, j), \mathcal{S}_{\mathsf{read}}(\mathsf{DB}) \rangle$ locally or in interaction with $\mathcal{A}$, depending whether i is corrupted.

(6) On input $\mathsf{write}(i, j, d)$ by $\mathcal{A}$, $\mathcal{B}$ executes $\langle \mathcal{C}_{\mathsf{write}}(cap_i, j, d), \mathcal{S}_{\mathsf{write}}(\mathsf{DB}) \rangle$ locally or in interaction with $\mathcal{A}$, depending whether i is corrupted.

Challenge. Finally, $\mathcal{A}$ outputs an index j^* which he wants to be challenged on. $\mathcal{B}$ parses the Log to search one version of that entry that contains a pair (t', σ') that has not been queried to the signing oracle and such that $\mathsf{verify}(vk, t', \sigma') = 1$.

Output. $\mathcal{B}$ outputs such a pair (t', σ') and it interrupts the simulation.

The simulation is clearly efficient. It is easy to see that, in order to win the game, $\mathcal{A}$ must be able to change an entry without writing permission and by leaving it in a consistent state. This can be done by either computing a collision in the chameleon hash function or by forging a valid signature on the tag t. By assumption we ruled out the first hypothesis, thus the winning condition of the adversary implies the forge of a verifying message-signature pair. Note that the $\mathcal{A}$ could also just roll back to some previous version of the entry but this can be easily prevented by including some timestamp in the computation of the chameleon hash. The winning probability of the reduction then carries over:

$$\Pr[\mathcal{B} \text{ wins}] \approx \Pr[\mathcal{A} \text{ wins} \mid \neg \mathsf{COLL}] \geq \epsilon(\lambda).$$

In this way we built an efficient adversary $\mathcal{B}$ that breaks the existential unforgeability of Π_{DS} with non negligible probability, which is clearly a contradiction. Therefore it must hold that $\Pr[\mathcal{A} \text{ wins} \mid \neg \mathsf{COLL}]$ is a negligible function in the security parameter. Finally we have that:

$$\Pr[\mathsf{ExpAcc}^{\mathcal{A}}_{\mathsf{GORAM}}(\lambda) = 1] \approx$$
$$\Pr[\mathcal{A} \text{ wins} \mid \neg \mathsf{COLL}] \leq negl(\lambda) \quad,$$

which concludes our proof. $\qquad\square$

Lemma 11. *Let $\mathcal{ZKP}$ be a zero-knowledge proof system and Π_{SE} be a CPA-secure private-key encryption scheme. Then* S-GORAM *achieves obliviousness.*

Proof of Lemma 11. The proof works analogously to Lemma 4. $\qquad\square$

$\qquad\square$

E   Algorithms for GORAM with Accountable Integrity

Implementation of $(cap_{\mathcal{O}}, \mathsf{DB}) \leftarrow \mathsf{gen}(1^\lambda, n)$**.** Additionally to Algorithm 1 we include in the capability of the data owner $cap_{\mathcal{O}}$ the key pair $(vk_{\mathcal{O}}, sk_{\mathcal{O}}) \leftarrow \mathsf{Gen}_{\mathsf{DS}}(1^\lambda)$, used for signing the chameleon hash tag t. We also add another predicate-encryption key pair (mpk, psk) instead of the predicate-only (opk, osk) (line 1.2). We refer to the two key pairs in the following as $(mpk_{\mathsf{Auth}}, psk_{\mathsf{Key}})$ and $(mpk_{\mathsf{Key}}, psk_{\mathsf{Key}})$.

Implementation of $\{cap_i, \mathsf{deny}\} \leftarrow \mathsf{addCl}(cap_{\mathcal{O}}, \mathbf{a})$**.** The difference from Algorithm 2 is that we instrument the capability cap_i with another predicate-encryption key sk_{f_i} instead of the predicate-only osk_{f_i} (line 2.7). The two predicate encryption secret keys are denoted by sk_{Auth} and sk_{Key} in the following.

Implementation of $\{\mathsf{DB}', \mathsf{deny}\} \leftarrow \langle \mathcal{C}_{\mathsf{addE}}(cap_{\mathcal{O}}, \mathbf{a}, d), \mathcal{S}_{\mathsf{addE}}(\mathsf{DB}) \rangle$**.** The difference from Algorithm 3 is line 3.9, where the new entry is instead composed of $E_k = (j, c^j_{\mathsf{Auth}}, c^j_{\mathsf{Key}}, c^j_{\mathsf{Data}}, r^j, cpk, t^j, \sigma^j)$ where

$$
\begin{aligned}
&(cpk, csk) \leftarrow \mathsf{Gen}_{\mathsf{CHF}}(1^\lambda) && c^j_{\mathsf{Auth}} \leftarrow \mathsf{PrEnc}(mpk_{\mathsf{Auth}}, x_{w,j}, csk)\\
&r^j \leftarrow \{0,1\}^\lambda && c^j_{\mathsf{Key}} \leftarrow \mathsf{PrEnc}(mpk_{\mathsf{Key}}, x_{r,j}, k_j)\\
&k_j \leftarrow \mathsf{Gen}_{\mathsf{SE}}(1^\lambda) && c^j_{\mathsf{Data}} \leftarrow \mathcal{E}(k_j, d)\\
&t^j \leftarrow \mathsf{CH}(cpk, j \parallel c^j_{\mathsf{Auth}} \parallel c^j_{\mathsf{Key}} \parallel c^j_{\mathsf{Data}} \parallel cpk, r^j)\\
&\sigma^j \leftarrow \mathsf{sign}(sk_{\mathcal{O}}, t^j).
\end{aligned}
$$

Furthermore, the rerandomization in 3.12 is substituted by the re-encryption under the same symmetric key.

Eviction. The eviction algorithm is implemented as in Algorithm 4 with the difference that the proof P is never computed. Additionally, the rerandomization step (line 4.4) is substituted with a re-encryption of the top-level encryption layer.

Implementation of $\langle \mathcal{C}_{\mathsf{chMode}}(cap_{\mathcal{O}}, \mathbf{a}, j), \mathcal{S}_{\mathsf{chMode}}(\mathsf{DB}) \rangle$**.** The algorithm is defined as in Algorithm 5, except for line 5.11 where the new entry is initialized as defined in the above addE.

Implementation of $\{d, \mathsf{deny}\} \leftarrow \langle \mathcal{C}_{\mathsf{read}}(cap_i, j), \mathcal{S}_{\mathsf{read}}(\mathsf{DB}) \rangle$**.** The read algorithm follows Algorithm 6 until line 6.9 and stops there. Furthermore, the proof P is no longer computed.

Implementation of $\{\mathsf{DB}', \mathsf{deny}\} \leftarrow \langle \mathcal{C}_{\mathsf{write}}(cap_i, j, d), \mathcal{S}_{\mathsf{write}}(\mathsf{DB}) \rangle$**.** The write algorithm is equivalent to the aforementioned read algorithm up to the point where we upload the entry E_k stored

at the index j, which is modified as follows:

$$csk^j := \mathsf{PrDec}(sk_{\mathsf{Auth}}, c^j_{\mathsf{Auth}}) \quad k_j := \mathsf{PrDec}(sk_{\mathsf{Key}}, c^j_{\mathsf{Key}})$$
$$\hat{c}^j_{\mathsf{Data}} \leftarrow \mathcal{E}(k_j, d)$$
$$\hat{r}^j \leftarrow \mathsf{Col}(csk^j, (j \ || \ c^j_{\mathsf{Auth}} \ || \ c^j_{\mathsf{Key}} \ || \ c^j_{\mathsf{Data}} \ || \ cpk^j), r^j,$$
$$\qquad (j \ || \ c^j_{\mathsf{Auth}} \ || \ c^j_{\mathsf{Key}} \ || \ \hat{c}^j_{\mathsf{Data}} \ || \ cpk^j))$$
$$E_k := (j, c^j_{\mathsf{Auth}}, c^j_{\mathsf{Key}}, \hat{c}^j_{\mathsf{Data}}, \hat{r}^j, cpk^j, t^j, \sigma^j)$$

F   Proof of Soundness for the Batched Zero-Knowledge Proof of Shuffle

In the following we argue why our novel batched ZK proof of shuffle (BZKPS) approach preserves all of the properties of the underlying zero-knowledge protocol for proving the correctness of the shuffle. In this section we abstract away the instantiation of the concrete protocol and we simply assume that it fulfills the properties of a zero-knowledge proof of knowledge system $\mathcal{ZKP}$ specified in Section 3.1. We do, however, assume that the common input to the prover and the verifier in such a proof are the two lists of ciphertexts $\vec{\theta}$ and $\vec{\theta'}$ and the additional input for the prover $\mathcal{P}$ is the permutation π and the randomnesses $\vec{r}$.

Correctness. For our BZKPS, it is easy to see that whenever the two input matrices of ciphers $\mathbf{A}$ and $\mathbf{A}'$ are honestly computed and the two parties do not deviate from the protocol specifications, the probability that the verifier is convinced of the statement does not vary with respect to the underlying $\mathcal{ZKP}$. This follows directly from the fact that, as long as $\mathbf{A}$ differs from $\mathbf{A}'$ only in the order of the rows, the resulting vectors $\vec{\theta}$ and $\vec{\theta'}$ will contain the same entries, in the same permuted order. Thus correctness holds true.

Zero-Knowledge and proof of knowledge. The zero-knowledge and the proof of knowledge properties can be also derived by the $\mathcal{ZKP}$ since the only procedure performed in BZKPS, outside of the original protocol, is a public operation, namely, the multiplication of ciphertexts which causes homomorphically a multiplication of plaintexts.

Soundness. Finally, in order to show that our approach preserves the soundness of the $\mathcal{ZKP}$, we have to prove that any malicious prover trying to fool the protocol succeeds with probability at most $1/2$. It is clear that this fact suffices by itself since with k protocol runs the success probability of the adversary drops to $1/2^k$, which is exponentially low in k. The intuition behind the proof is that the adversary cannot modify a single block in a row of $\mathbf{A}'$ without the verifier noticing at least half of the times, thus he needs to modify at least one more spot so to cancel out the multiplied values in the total product. But, again, this block will be selected by the verifier to contribute to the product just half of the times in expectation, so the success probability does not increase. This holds true no matter how many blocks in the same row the adversary modifies.

Proof of Theorem 1. In the soundness scenario a malicious prover $\mathcal{P}^*$ wants to convince the verifier $\mathcal{V}$ that the matrices of ciphers $\mathbf{A}$ and $\mathbf{A}'$ contain the same plaintexts up to the order of the rows, i.e., they are permuted with respect to π, and the ciphertexts are rerandomized with respect to $\mathbf{R}$. Let $|\mathbf{A}| = |\mathbf{A}'| = n \times m$. Note that $\mathbf{A}'$ is generated by $\mathcal{P}^*$ and then sent to $\mathcal{V}$ in an early stage of the protocol. Since the underlying $\mathcal{ZKP}$ is assumed to be sound, it follows that, any time an element θ_i in the resulting vector $\vec{\theta}$ will differ with respect to the correspondent permuted element $\theta'_{\pi(i)}$ in the vector $\vec{\theta'}$, the verifier will notice it with overwhelming probability. Thus, in order for $\mathcal{P}^*$ to succeed, it must be the case that for all i, $[\![\theta_i]\!] = [\![\theta'_{\pi(i)}]\!]$ where $[\![c]\!]$ denotes the plaintext encrypted in c.

Let Δ_i be the quotient between θ_i and $\theta'_{\pi(i)}$ and $\delta_i = [\![\Delta_i]\!]$. It follows from the argument above and from the homomorphic property of the encryption scheme that, in order for the protocol to successfully terminate, for all i, δ_i must be equal to 1. Since $\mathcal{P}^*$ has no control over the input $\mathbf{A}$, and therefore over $\vec{\theta}$, the value of δ_i is directly derived from the modifications that $\mathcal{P}^*$ introduced in the i-th row of the matrix $\mathbf{A}'$. We next prove the following: if $\mathcal{P}^*$ performed at least one modification on a given row i of $\mathbf{A}'$ (i.e., he multiplied some value different from one to the plaintext of $\mathbf{A}'_{i,j}$ for some $1 \leq j \leq m$), then δ_i has any possible fixed value with probability at most $1/2$. Intuitively, this statement is true since with probability $1/2$, the verifier does not pick a column that contributes to the difference of the plaintext product. Since any δ_i must be one for $\mathcal{P}^*$ to succeed, this directly proves our theorem.

The proof is conducted by induction on ℓ, the number of modified blocks in the given row:

$\ell = 1$: Assume without loss of generality that $\mathcal{P}^*$ introduces an arbitrary modification $z \neq 1$ on a given block (i.e., on the cipher identified by some column index j), then such a block is selected to contribute to the product θ'_i with probability exactly $1/2$. Thus it holds that:

$$\Pr[\delta_i^1 = 0] = \frac{1}{2} \qquad\qquad \Pr[\delta_i^1 = z] = \frac{1}{2}$$

$\ell \to \ell + 1$: Assume without loss of generality that $\mathcal{P}^*$ introduces some arbitrary modifications different than one on ℓ blocks of the given row and that he multiplied the value z to the $(\ell+1)$-th spot. We denote by $\mathsf{CHOOSE}(\ell)$ the event that the ℓ-th block is picked to contribute to the product θ'_i. Then, for all values y it holds that

$$\begin{aligned}\Pr[\delta_i^{\ell+1} = y] = {}& \Pr[\delta_i^\ell = y/z \mid \mathsf{CHOOSE}(\ell+1)] \\ & \cdot \Pr[\mathsf{CHOOSE}(\ell+1)] + \\ & \Pr[\delta_i^\ell = y \mid \neg\, \mathsf{CHOOSE}(\ell+1)] \\ & \cdot \Pr[\neg\, \mathsf{CHOOSE}(\ell+1)].\end{aligned}$$

Since $z \neq 1$ it follows that $y \neq y/z$, additionally the $(\ell+1)$-th spot is chosen to contribute to the product θ'_i with probability $1/2$, so we have

$$\begin{aligned}\Pr[\delta_i^{\ell+1} = y] = {}& \\ & \frac{1}{2} \cdot \Pr[\delta_i^\ell = y/z \mid \mathsf{CHOOSE}(\ell+1)] + \\ & \frac{1}{2} \cdot \Pr[\delta_i^\ell = y \mid \neg\, \mathsf{CHOOSE}(\ell+1)].\end{aligned}$$

Furthermore the two events are independent, thus we can rewrite

$$\begin{aligned}\Pr[\delta_i^{\ell+1} = y] = {}& \\ & \frac{1}{2} \cdot \Pr[\delta_i^\ell = y/z] + \frac{1}{2} \cdot \Pr[\delta_i^\ell = y].\end{aligned}$$

By induction hypothesis it holds that for all i, δ_i has any fixed value with probability at most $\frac{1}{2}$, then

$$\Pr[\delta_i^{\ell+1} = y] = \frac{1}{2} \cdot \frac{1}{2} + \frac{1}{2} \cdot \frac{1}{2} = \frac{1}{2}.$$

The induction above proves that for all i, $\delta_i = 1$ with probability at most $1/2$, therefore the verifier notifies the cheating prover with probability at least $1/2$. This concludes our proof. $\quad\square$