Research Article

Marco Cianfriglia, Elia Onofri*, and Marco Pedicini

mR_{LWE}-CP-ABE: A revocable CP-ABE for post-quantum cryptography

https://doi.org/10.1515/jmc-2023-0026
received September 04, 2023; accepted October 31, 2023

Abstract: We address the problem of user fast revocation in the lattice-based Ciphertext Policy Attribute-Based Encryption (CP-ABE) by extending the scheme originally introduced by Zhang and Zhang [Zhang J, Zhang Z. *A ciphertext policy attribute-based encryption scheme without pairings*. In: International Conference on Information Security and Cryptology. Springer; 2011. p. 324–40. doi: https://doi.org/10.1007/978-3-642-34704-7_23.]. While a lot of work exists on the construction of revocable schemes for CP-ABE based on pairings, works based on lattices are not so common, and – to the best of our knowledge – we introduce the first server-aided revocation scheme in a lattice-based CP-ABE scheme, hence being embedded in a post-quantum secure environment. In particular, we rely on semi-trusted "mediators" to provide a multi-step decryption capable of handling mediation without re-encryption. We comment on the scheme and its application, and we provide performance experiments on a prototype implementation in the Attribute-Based Encryption spin-off library of Palisade to evaluate the overhead compared with the original scheme.

Keywords: mediated CP-ABE, lattice-based cryptography, learning with errors, post-quantum cryptography, cloud cryptography

MSC 2020: 94A60, 68P25, 68Q25

1 Introduction

In this work, we tackle the problem of designing a fast key-revoking system in a Ciphertext Policy Attribute-Based Encryption (CP-ABE) constructed on some presumed post-quantum resistant algebraic setting. The presented approach involves a Dual-Regev CP-ABE scheme, which combines the advantages of attribute-based encryption with the security properties of the Regev encryption scheme [1] and provides a flexible and secure mechanism for access control and data encryption.

The Regev encryption scheme is a lattice-based encryption scheme based on the hardness of the Learning with Errors problem (the so-called *learning with errors (LWE) assumption*), which is considered to be resistant to quantum attacks by the worst-case complexity of the decision version of the shortest vector problem (GapSVP) and shortest independent vector problem (SIVP) on lattices. It represents messages as vectors, and encryption is achieved by adding noise to those vectors. Decryption, conversely, can only be efficiently done by the intended recipient who possesses a secret key.

In a Dual-Regev CP-ABE scheme, ciphertexts are associated with access policies represented as pattern strings, where symbols can be 0, 1, or *. Users possess secret keys corresponding to their attributes (represented as bit strings). Decryption succeeds if the user's attribute matches the policy pattern specified in the

*** Corresponding author: Elia Onofri,** Department of Mathematics and Physics, Roma Tre University, 00146 Rome, Italy; Institute for Applied Mathematics, National Research Council, 00185 Rome, Italy, e-mail: elia.onofri@uniroma3.it
Marco Cianfriglia: Department of Mathematics and Physics, Roma Tre University, 00146 Rome, Italy, e-mail: marco.cianfriglia@uniroma3.it
Marco Pedicini: Department of Mathematics and Physics, Roma Tre University, 00146 Rome, Italy, e-mail: marco.pedicini@uniroma3.it
ORCID: Marco Cianfriglia 0000-0002-6775-7804; Elia Onofri 0000-0001-8391-2563; Marco Pedicini 0000-0002-9016-074X

ciphertext. This allows for fine-grained access control, where access to encrypted data is granted based on attribute matching. In contrast to most of the CP-ABE schemes that are based on bilinear maps, these schemes do not rely on pairings. The absence of pairings in such schemes offers advantages in terms of simplicity and efficiency.

In this context, revocation implies the capability of making the user's attributes/keys no longer valid, if needed; situations where revocation is required are typical in the business context, where, e.g., a product owner requires the ability to revoke privileges to users without needing to re-encrypt all the data. This is particularly true if the product owner has no direct control over the data, e.g., if they are hosted on third-party platforms or they actually belong to other users (like in the case of sharing platforms). Although this last example might seem unusual, it is common that cloud storage companies store data on servers that are, e.g., located in different countries and are hence subject to different laws concerning data regulation and privacy. In such a scenario, a company might want to retain complete control over data access without having to store them in first person.

To address the issue of user revocation, we propose mR_{LWE}-CP-ABE, a novel solution that builds upon the Dual Regev CP-ABE scheme introduced in the study by Zhang and Zhang [2] by enforcing a security-mediated public key encryption (PKE). In particular, mR_{LWE}-CP-ABE shares similarities with the ideas introduced in USENIX 2001 Boneh et al.'s article, [3].

The main idea presented in this article is the use of a (semi)trusted third party, called the *security mediator* (SM), to check the user whenever she wants to decrypt a ciphertext. The user requires assistance of the security mediator because the secret key is separated into two (or more) portions during key generation, with one portion given to the user while the remaining parts are given to (possibly multiple) SMs. The user requires the security mediator's help in order to enable full user secret key recovery, and decrypt or sign messages.

We implement the proposed scheme in Palisade [4]. In particular, by building such practical tests on the implementation of the Palisade attribute-based encryption (ABE) project spin-off, we implicitly show its effectiveness.

1.1 Related works

ABE, first proposed in the study by Sahai and Waters [5], is an asymmetric cryptographic primitive for one-to-many encryption that, as highlighted by high number of surveys in the last few years on it [6–10], attracted many interests along the years as it provides fine-grained access control over data. An ABE scheme allows a data owner to encrypt some data once and to share them with many along with a set of required attributes that define an access policy; the set of valid recipients is not required to be known in advance: all we need is that an authorised user must retain a set of valid attributes that satisfy the access policy. Each user is identified by the set of attributes of his/her owns. Over the years, two variants of ABE have been proposed in the literature: the CP-ABE [11] and the Key Policy Attribute-Based Encryption (KP-ABE) [12]. In CP-ABE, the access policy is applied to the ciphertext, conversely, in KP-ABE it is associated with the secret key, so usually CP-ABE is preferred as it is more flexible. Different from classical public key schemes where a user who wants to share encrypted data with many others is required to perform many encryptions, one for each of valid recipient, in ABE schemes, the encryption is performed only once for many users: for this reason, in cloud environments, ABE schemes are a common choice. However, in this context, usually, the set of users changes frequently so the ability to revoke some users is a necessary requirement for any ABE scheme.

In the literature [13], the revocation mechanism was categorised into three classes: direct, indirect, and server-aided.

The *direct revocation* follows the approach of conventional public key management systems (PKMS) where a certificate revocation list (CRL) is distributed. Once a user needs to be revoked, the key authority in the PKMS adds the user identifier to the CRL and shares the updated list. Some examples of ABE schemes that implement direct revocation were given in the studies by Liu et al. [14] and Phuong et al. [15]. The major drawback of direct revocation is, of course, related to the distribution of the updated CRL. Any data owner must update his/

her CRL before encrypting new data to exclude revoked users. Furthermore, as the revoked user set grows, so does the size of the CRL. Liu et al. [14] proposed a solution to overcome both issues by setting expiration dates on keys, by embedding the revocation list along with the ciphertext, and by removing revoked keys from the list once expired; however, in such schemes, data owner still needs to update his/her CRL to be sure not to miss any recently revoked user.

In *indirect revocation*, every time a user is revoked, the key authority generates new keys only for the remaining non-revoked users. The benefit of this approach is that the server only needs to work on the subset of still active users and does not need to periodically share the CRL. A few examples of CP-ABE indirect revocable schemes were mentioned in the studies by Sahai et al. [16] and Yu et al. [17] and in the study by Xie et al. [18]. For instance, in both studies by Sahai et al. [16] and by Yu et al. [17], the authors proposed to update both the keys for still active users and the older ciphertexts, stored on the cloud, to not letting a revoked user to decrypt them anymore. The approach proposed in in the study by Xie et al. [18] is slightly different: they update the keys but each user has two different keys, an individual and a group key, both needed for the decryption.

Server-aided revocation solutions try to avoid the need for key updates and the distribution of the CRL. They required, as the system we propose in this article, to leverage third-party cooperation to decrypt. Here, following the approach first proposed in the study by Boneh et al. [3] and then applied also in the studies by Yang et al. [19] and Cul et al. [20], we rely on key-splitting feature for the revocation. Different from the literature, to the best of our knowledge, mR$_{LWE}$-CP-ABE is the first application of such a revocation technique to a lattice-based ABE scheme; we believe this is an important step as our system, uniquely with respect to all the previous works, is embedded in a post-quantum secure environment.

1.2 Contributions

Here, in the following, we list the main contributions of this work

- Inspired by the study by Boneh et al. [3], we propose mR$_{LWE}$-CP-ABE the first, to the best of our knowledge, CP-ABE revocation scheme based on lattices, a presumed post-quantum resistant algebraic setting. We start with the CP-ABE scheme presented in the study by Zhang and Zhang [2], and we extend and modify it to support key revocation. We rely on (semi)trusted third party, called the *security mediator*, to perform fast and efficient user key revocation.
- We provide a formal description of the proposed scheme along with the analysis of its parameter and its security proof.
- We implement mR$_{LWE}$-CP-ABE scheme on Palisade, a well-known crypto library, and we experimentally evaluate the overhead introduced by the revocation mechanism in terms of performance.
- We will release the implementation of our scheme to let the community independently test and evaluate it.

1.3 Organisation of the article

The rest of this article is organised as follows. Section 2 wraps up the notation and the mathematical basics used in the rest of this article. Section 3 reviews the mathematical background used throughout this article (a confident reader can safely skip this section): in particular, the SIVP problem (Section 3.1), the discrete Gaussian (Section 3.2), the learning with error problem (Section 3.3), and some useful algorithms on lattices (Section 3.4) are recalled. Section 4 introduces the system model of CP-ABE (Section 4.1) and analyses the scheme presented in the study by Zhang and Zhang [2] by reworking its definition (Section 4.2) and analysing its parameters (Section 4.3), providing a few small changes in the notation to better prepare the ground for the mediated scheme. Section 5 holds the main contribution of the article, defining the mediated CP-ABE system model (Section 5.1), introducing the novel scheme (Section 5.2), and analysing its parameters (Section 5.3). Section 6 analyses both the threat model (Section 6.1) and the classical security (Section 6.2) of the proposed

scheme. Section 7 introduces the multi-bit variation on the original and mediated scheme, following the build from the study by Zhang and Zhang [2]. Section 8 presents some benchmarks and results on the proposed scheme. Finally, Section 9 closes this article resuming the contributions and providing some hints on future works.

2 Notation

Numeric sets of positive integers, integers, and real numbers are denoted with blackboard bold letters $\mathbb{N}$, $\mathbb{Z}$, and $\mathbb{R}$, respectively. The quotient group modulo $q, q \in \mathbb{N}$, is denoted by $\mathbb{Z}_q = \mathbb{Z}/q\mathbb{Z} = \{0, \ldots, q-1\}$. Probabilities are defined by capital letter $\mathbb{P}[\cdot]$, and distributions are denoted usually with χ, and we say that a is sampled from it by writing $a \hookleftarrow_{\$} \chi$. In particular, the uniform distribution over a set S is denoted by $U(S)$.

Matrices are usually denoted by upper-case letters ($\mathbf{A}, \mathbf{B}, \ldots$), while vectors are interpreted as single-column matrices and usually denoted by lower-case letters ($\mathbf{a}, \mathbf{b}, \ldots$). Matrices (and vectors) can be transposed ($\mathbf{A}^T$), concatenated by columns ($[\mathbf{A}||\mathbf{B}]$), or concatenated by rows ($\mathbf{A};\mathbf{B}$). The scalar product is denoted by $\langle\cdot,\cdot\rangle$, while the Euclidean and the infinity norm of a vector are denoted by $||\mathbf{a}||$ and $||\mathbf{a}||_\infty$, respectively. By the abuse of notation, we define the norm of a matrix as the infinity norm over the Euclidean norm of its columns, i.e. if $\mathbf{A} = [\mathbf{a}_1||\ldots||\mathbf{a}_n]$, then $||\mathbf{A}|| = \max_i||\mathbf{a}_i||$. Finally, if the columns of a matrix $\mathbf{A} = [\mathbf{a}_1||\ldots||\mathbf{a}_n]$ are linearly independent, we denote with $\tilde{\mathbf{A}} = [\tilde{\mathbf{a}}_1||\ldots||\tilde{\mathbf{a}}_n]$ the Gram-Schmidt orthogonalisation of vectors $a_1, \ldots, a_n$ taken in that order.

We refer to attributes with calligraphic capital letters; in particular, $\mathcal{R}$ denotes the admissible attributes, $\mathcal{S}$ denotes the user attribute specifications, and $\mathcal{W}$ denotes the ciphertext access structures. If $\mathcal{S}$ is compatible with $\mathcal{W}$, we say that it satisfies the access structure and we write $\mathcal{S} \vdash \mathcal{W}$; otherwise, we write $\mathcal{S} \nvdash \mathcal{W}$.

The security parameter throughout this article is n, and all other quantities are the implicit functions of it. We use standard notations big-O, big-$\tilde{O}$, and small-ω to denote asymptotic classes, we write poly(n) to determine functions $f(n) = O(n^c)$ for some constant c, and we write negl(n) to determine negligible functions $f(n)$, i.e., eventually upper bounded by $\frac{1}{n^c}$. Finally, we say that a probability is overwhelming if it is $1 - \text{negl}(n)$.

3 Prerequisites on lattices

An n-dimensional lattice of rank $m \leq n$ is a subset of $\mathbb{R}^n$ given by the span of m linear independent vectors $\mathbf{b}_1, \ldots, \mathbf{b}_m \in \mathbb{R}^n$. In formulas, we have

$$\mathbf{\Lambda} = \mathcal{L}(\mathbf{B}) = \{\langle \mathbf{B}, \mathbf{c}\rangle | \mathbf{c} \in \mathbb{Z}^m\},$$

where $\mathbf{B} \in \mathbb{R}^{n\times m} = [\mathbf{b}_1||\ldots||\mathbf{b}_m]$ is called the *basis* of the lattice.

The set of linear functionals that take integer values on each point of $\mathbf{\Lambda}$ is called *dual lattice*, and it is denoted by:

$$\mathbf{\Lambda}^* = \{x \in \mathbb{R}^n | \langle \mathbf{x}, \mathbf{v}\rangle \in \mathbb{Z}, \quad \text{for all } v \in \mathbf{\Lambda}\}.$$

Given a matrix $\mathbf{A} \in \mathbb{Z}^{n\times m}$, the set of vectors that nullifies $\mathbf{A}$ is an m-dimensional lattice, called *orthogonal lattice of* $\mathbf{A}$, and it is denoted by:

$$\mathbf{\Lambda}^{\perp}(A) = \{\mathbf{e} \in \mathbb{Z}^m | \langle \mathbf{A}, \mathbf{e}\rangle = \mathbf{0}\}.$$

Orthogonal lattices are particularly useful when working in modular arithmetic; given a matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$, we analogously define

$$\mathbf{\Lambda}_q^{\perp}(\mathbf{A}) = \{\mathbf{e} \in \mathbb{Z}^m | \langle \mathbf{A}, \mathbf{e}\rangle \equiv_q \mathbf{0}\}.$$

We further observe that for any square matrix $\mathbf{B} \in \mathbb{Z}_q^{n\times n}$, we have $\langle \mathbf{A}, \mathbf{x}\rangle = 0 \Leftrightarrow \langle \mathbf{B}, \langle \mathbf{A}, \mathbf{x}\rangle\rangle = 0$; hence, $\Lambda_q^{\perp}(\mathbf{A}) = \Lambda_q^{\perp}(\langle \mathbf{B}, \mathbf{A}\rangle)$.

3.1 Hard problems

Many cryptographic primitives have been constructed whose security is based on the (worst-case) hardness of SIVP or closely related lattice problems. In particular, the (worst-case) hardness of the SIVP for poly(n) approximation factors implies the existence of several fundamental cryptographic primitives. Blömer and Seifert [21] showed that the SIVP is NP-hard to approximate for any constant approximation factor γ. Their result is shown only for the Euclidean norm, and their proofs were extended to arbitrary norms by the study by Aggarwal and Chung [22].

The norm of vector $\mathbf{x}$, denoted by $||\mathbf{x}||$, is defined with respect to integer p:

$$||\mathbf{x}||_p \coloneqq \left(\sum_{i=1}^{n} |x_i|^p\right)^{1/p}.$$

We write SIVP$_p$ as a notation respective to p. Hence, SIVP$_2$ is the case considered in the study by Blömer and Seifert [21].

Hereinafter, we suppose fixed $p = 2$ and we omit from explicitly mentioning it in the norm.

A basic parameter of the lattice $\mathbf{\Lambda}$ is the length of the shortest non-zero vector in the lattice. The parameter λ_1 is also indicated as the *first successive* of $\mathbf{\Lambda}$ and denoted by λ_1. It is important to know the lower and upper bounds for λ_1, which, of course, depend on p: a lower bound is given by the length of the shortest vector in the Gram–Schmidt reduced form of the basis: $\lambda_1 \le \min_i||\tilde{\mathbf{b}}_i||$. Similarly, for $i = 1, \ldots, n$, the i-th successive minimum, denoted by $\lambda_i(\mathbf{\Lambda})$, is the smallest l such that there are i non-zero linearly independent lattice vectors that have length at most l.

The SIVP consists in finding n independent and "short" vectors: given a basis $\mathbf{B} \in \mathbb{Z}^{n\times n}$, find independent vectors $\mathbf{u}_1, \ldots, \mathbf{u}_n$ such that $||\mathbf{u}_i|| \le \lambda_n$ for $i = 1, \ldots, n$, [23].

Proposition 1. (Theorem 2 from [22]) *Under the (randomised) gap exponential time hypothesis, for any $p \ge 1$, there exists $\gamma > 1$, $\varepsilon > 0$ such that γ-SIVP$_p$ with rank n is not solvable in $2^{\varepsilon n}$ time.*

The Gap-Exponential Time Hypothesis is a fine-grained complexity-theoretic hypothesis introduced in the study by Impagliazzo and Paturi [24], and it is required to exclude sub-exponential algorithms.

3.2 Discrete Gaussians

We recall the definition of Gaussian function centred in $\mathbf{c}$ and scaled by a factor of s to be

$$\rho_{s,\mathbf{c}}(\mathbf{x}) = \exp\left(-\pi\frac{||\mathbf{x}-\mathbf{c}||^2}{s^2}\right), \quad \mathbf{x} \in \mathbb{R}^n.$$

A Gaussian function is typically used to build (continuous) probability distributions as:

$$D_{s,\mathbf{c}}(\mathbf{x}) = \frac{\rho_{s,\mathbf{c}}(\mathbf{x})}{s^N}, \quad \mathbf{x} \in \mathbb{R}^n,$$

being $s^N = \int_{\mathbf{x}\in\mathbb{R}^n}\rho_{s,\mathbf{c}}(\mathbf{x})\mathrm{d}\mathbf{x}$ the total measure associated with $\rho_{s,\mathbf{c}}$.

Given a lattice $\mathbf{\Lambda} \subset \mathbb{Z}^n$, we can discretise the distributions $D_{s,\mathbf{c}}$ on it by distributing $\mathbf{x} \in \mathbb{R}^n$ according to $D_{s,\mathbf{c}}$ and conditioning $\mathbf{x} \in \mathbf{\Lambda}$, thus obtaining

$$D_{\Lambda,s,\mathbf{c}}(\mathbf{x}) = \frac{D_{s,\mathbf{c}}(\mathbf{x})}{D_{s,\mathbf{c}}(\Lambda)} = \frac{\rho_{s,\mathbf{c}}(\mathbf{x})}{\rho_{s,\mathbf{c}}(\Lambda)},$$

with $\rho_{s,\mathbf{c}}(\Lambda)$ being the proper normalisation constant evaluated as $\rho_{s,\mathbf{c}}(\Lambda) = \sum_{\mathbf{y}\in\Lambda}\rho_{s,\mathbf{c}}(\mathbf{y})$. We call such distribution a *discrete Gaussian function* with centre $\mathbf{c}$ and parameter s, and we omit the subscripts s and $\mathbf{c}$ if equal, respectively, to 1 and to the origin $\mathbf{0}$.

Given a parameter $\varepsilon \in \mathbb{R}^+$, we further recall from the study by Micciancio and Regev [25] the definition of *smoothing parameter* η_ε as:

$$\eta_\varepsilon = \min\{s \in \mathbb{R}^+ | \rho_{\frac{1}{s}}(\Lambda^* \backslash \{\mathbf{0}\}) \leq \varepsilon\}.$$

In particular, if $s \geq \eta_\varepsilon$, we can bound the dispersion of the Gaussian as per the following.

Lemma 1. (Lemma 4.4 from [25]) *For any n-dimensional lattice* Λ, *for any centre* $\mathbf{c} \in \mathbb{R}^n$, *and for any* $\varepsilon \in (0,1)$, *we have that if* $s \geq \eta_\varepsilon(\Lambda)$, *then*

$$\mathbb{P}_{\mathbf{x}\leftarrow_\$ D_{\Lambda,s,\mathbf{c}}}[||\mathbf{x} - \mathbf{c}|| > s\sqrt{n}\,] \leq \frac{1-\varepsilon}{1+\varepsilon} \cdot 2^{-n}.$$

3.3 LWE

Originally presented in the study by Regev [1] and later extended in the study [26], LWE is a hard lattice problem founding in Fully Homomorphic Encryption. Its hardness has been proven in the study by Regev [1] via a quantum reduction to SIVP and GapSVP and in the study by Peikert [27] via a classical reduction to a variation of GapSVP.

Let $q \in \mathbb{N}$ and let χ be a probability distribution on $\mathbb{Z}_q$. For any $\mathbf{s} \in \mathbb{Z}_q^n$, LWE instances with secret $\mathbf{s}$ are defined as samples from:

$$A_{\mathbf{s},\chi} = \{(\mathbf{a}, \mathbf{y}) \in \mathbb{Z}_q^n \times \mathbb{Z}_q | \mathbf{y} = \mathbf{a}^T\mathbf{s} + x, \text{ with } \mathbf{a} \leftarrow_\$ U(\mathbb{Z}_q)^n, x \leftarrow_\$ \chi\}.$$

LWE can be either formulated as a search or a decision problem, being the first to recover $\mathbf{s}$ given multiple samples of $A_{\mathbf{s},\chi}$ and the second to distinguish between $A_{\mathbf{s},\chi}$ and $U(\mathbb{Z}_q)^n \times U(\mathbb{Z}_q)$. In particular, if $q = \text{poly}(n)$, the two problems are polynomially equivalent [26]. In the following, we denote with $\text{LWE}_{q,\chi}$ a generic instance of LWE with a specific parameter $q \in \mathbb{N}$.

Let us denote by Ψ_α, a periodisation of the normal distribution with mean 0 and variance $\frac{\beta^2}{2\pi}$ and by $\bar{\Psi}_\alpha$ its discretisation, then we have:

Proposition 2. (Theorem 1.1 from [26]) *Let* $\alpha = \alpha(n) \in (0,1)$ *and let* $q \in \mathbb{N}$ *be such that* $\alpha q > 2\sqrt{n}$ *holds. Assuming we have access to an oracle that solves* $\text{LWE}_{q,\bar{\Psi}_\alpha}$ *given a polynomial number of samples, then there exists an efficient quantum algorithm for solving the decision version of GapSVP and SIVP to within* $\tilde{O}(\frac{n}{\alpha})$ *in the worst case.*

More formally, for $r \in [0,1)$, we have

$$\Psi_\alpha(r) = \sum_{k=-\infty}^{\infty} \frac{1}{\alpha} \cdot \exp\left(-\pi\left(\frac{r-k}{\alpha}\right)^2\right) \bmod 1$$

and

$$\bar{\Psi}_\alpha(r) = \lfloor q \cdot \Psi_\alpha(r) \rceil \bmod q.$$

In particular, we can characterise the distribution $\bar{\Psi}_\alpha^m$ as follows:

Lemma 2. (Lemma 12 from [28]) *Let* $\mathbf{e} \in \mathbb{Z}^m$ *and* $\mathbf{y} \leftarrow_\$ \bar{\Psi}_\alpha^m$. *Then, the following relation in* $\mathbb{Z}_q$ *holds* (*but for negligible probability in m*):

$$|\mathbf{e}^T\mathbf{y}| \leq ||\mathbf{e}|| \cdot q\alpha \cdot \omega(\sqrt{\log m}) + ||\mathbf{e}|| \frac{\sqrt{m}}{2}.$$

In particular, for $x \hookleftarrow_{\$} \bar{\Psi}_\alpha$, it holds in $\mathbb{Z}_q$ (but for negligible probability in m):

$$|x| \leq q\alpha \cdot \omega(\sqrt{\log m}) + \frac{1}{2}.$$

3.4 Literature algorithms on lattices

In the following, we recall four algorithms from the literature that are later used both in the original CP-ABE scheme and in mR$_{LWE}$-CP-ABE.

Function 1. (`SampleGaussian`, Theorem 4.1 from [29]) Let $\mathbf{\Lambda} = \mathcal{L}(\mathbf{B}) \subset \mathbb{R}^m$ be a m-dimensional lattice with basis $\mathbf{B}$. Given a Gaussian parameter $s \in \mathbb{R}^+$ such that $s \geq ||\tilde{\mathbf{B}}|| \cdot \omega(\sqrt{\log m})$ and for any centre $\mathbf{c} \in \mathbb{R}^m$, there exists a probabilistic polynomial-time algorithm `SampleGaussian`$(\mathbf{B}, s, \mathbf{c})$ that samples a vector $\mathbf{x} \in \mathbf{\Lambda}$ with a distribution statistically close to the discrete Gaussian $D_{\mathbf{\Lambda},s,\mathbf{c}}$.

Function 2. (`TrapGen`, Algorithm 1 from [30]) Let $q \in \mathbb{N}$ be an odd prime associated with a security parameter n and let $m \in \mathbb{N}$ be a dimension such that $m \geq (5 + 3\delta_0)n \log q$, for any $\delta_0 \in \mathbb{R}^+$. There exists a probabilistic polynomial-time algorithm `TrapGen`(n, m, q) that generates a statistically $(mq^{-\delta_0 \frac{n}{2}})$-close to uniform matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and a with-overwhelming-probability-short basis $\mathbf{T_A}$ of the orthogonal lattice $\mathbf{\Lambda}_q^{\perp}(\mathbf{A})$, i.e., such that $||\mathbf{T_A}|| \leq O(n \log q)$ and $||\tilde{\mathbf{T}}_\mathbf{A}|| \leq O(\sqrt{n \log q})$.

In the following, we choose $\delta_0 = \frac{1}{3}$ so that we obtain $m \geq \lceil 6n \log q \rceil$.

Function 3. (`SamplePre`, Section 5.2 from [29]) Let $q \in \mathbb{N}$ be an odd prime associated with a security parameter n, let $m \in \mathbb{N}$ be a dimension such that $m \geq 2n \log q$, and let $s \in \mathbb{R}$ be a Gaussian parameter such that $s \geq \omega(\sqrt{\log m})$. In general, for all (but a $2q^{-n}$ fraction of) $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, the distribution of the syndrome $\mathbf{u} = \mathbf{Ae} \bmod q$ yielded by $\mathbf{e} \hookleftarrow_{\$} D_{\mathbb{Z}^m,s}$ is statistically close to $U(\mathbb{Z}_q^n)$. In particular, for such values, there exists a probabilistic polynomial time algorithm `SamplePre`$(\mathbf{A}, \mathbf{T_A}, s, \mathbf{u})$ that samples $\mathbf{e}$ given a short basis $\mathbf{T_A}$ of the orthogonal lattice $\mathbf{\Lambda}_q^{\perp}(\mathbf{A})$, conditioned on s being such that $s \geq ||\tilde{\mathbf{T}}_\mathbf{A}|| \cdot \omega(\sqrt{\log m})$.

Function 4. (`GenSamplePre`, Theorem 3.4 from [31]) Let $q \in \mathbb{N}$ be an odd prime associated with a security parameter n and let $m \in \mathbb{N}$ be a dimension such that $m \geq 2n \log q$. Assume $\mathbf{A} = [\mathbf{A}_1||\ldots||\mathbf{A}_k] \in \mathbb{Z}_q^{n \times mk}$ and consider $J = \{j_1, \ldots j_{|J|}\} \subset \{1, \ldots, k\}$ be a set of indices of the $\mathbf{A}_i$ matrices[1]. Let $\mathbf{A}_J = [\mathbf{A}_{j_1}||\ldots||\mathbf{A}_{j_{|J|}}]$ and let $\mathbf{T}_{\mathbf{A}_J}$ be a basis of the orthogonal lattice $\mathbf{\Lambda}_q^{\perp}(\mathbf{A}_J)$. There exists a probabilistic polynomial-time algorithm `GenSamplePre` $(\mathbf{A}, \mathbf{T}_{\mathbf{A}_J}, J, s, \mathbf{u})$ that samples $\mathbf{e} \hookleftarrow_{\$} D_{\mathbb{Z}^{mk},s}$ condition on $\langle \mathbf{A}, \mathbf{e} \rangle = \mathbf{u}$, with $s \geq ||\tilde{\mathbf{T}}_{\mathbf{A}_J}|| \cdot \omega(\sqrt{\log km})$ (hence independent of the choice and size of J).

In particular, to build such an algorithm, consider $\bar{J} = \{1, \ldots, k\} \backslash J$. We can retrieve $\mathbf{e}_i$ for $i \in \bar{J}$ directly from $\mathbf{e}_{\bar{J}} \hookleftarrow_{\$} D_{\mathbb{Z}^{m \cdot (k - ||J||)},s}$, while $\mathbf{e}_j$ for $j \in J$ can be retrieved from $\mathbf{e}_J = \text{SamplePre}(\mathbf{B}, \mathbf{T_B}, s, \mathbf{u} - \langle \mathbf{A}, \mathbf{e}_{\bar{J}} \rangle)$, so building $\mathbf{e}$ such that $\langle \mathbf{A}, \mathbf{e} \rangle = \mathbf{u}$.

4 CP-ABE scheme on lattices

We open this section by recalling the formal definition of a CP-ABE scheme and a possible security model for it. Then, we review the CP-ABE scheme presented in the study by Zhang and Zhang [2], we extend later in Section 5, and we discuss its parameters requirements and security.

1 More in general, at least n columns of the matrix $\mathbf{A}$ are required; however, for the sake of simplicity, we consider only blocks $\mathbf{A}_i$.

4.1 System model

A CP-ABE scheme is a framework to perform secure data sharing where recipients are not specific users – like in classic PKE schemes – but rather users with specific attributes. A trusted central authority is needed for what concerns user key creation; however, data encryption and decryption can be performed without its further collaboration; in particular, also data owners outside the accredited users can encrypt data.

More formally, we have

Scheme 1. (CP-ABE) A CP-ABE scheme consists of four algorithms:

- Setup($\boldsymbol{\sigma}, \mathcal{R}$) $\rightarrow$ (msk, pk) is the initialisation algorithm executed by a central authority to set up a pair of public key (pk) and master secret key (msk) starting from a set of security parameters $\boldsymbol{\sigma}$ and a set of admissible attributes $\mathcal{R}$. msk is used for the creation of users' keys, while pk is used for message encryption.
- KGen(msk, $\mathcal{S}$) $\rightarrow$ sk is the algorithm the authority runs to accredit a user with an attribute specification $\mathcal{S}$, hence building a private key sk capable of decrypt ciphertexts only with access structure $\mathcal{W}$ such that $\mathcal{S} \vdash \mathcal{W}$.
- Enc(pk, $\mathcal{W}, M$) $\rightarrow C$ is the encryption algorithm run by a data owner to encrypt the message M in a ciphertext C with access structure $\mathcal{W}$. Only the public key pk is needed to perform this operation.
- Dec(sk, C) $\rightarrow M'$ or $\perp$ is the decryption algorithm run by a user to retrieve the message M associated with the ciphertext C. The equality $M = M'$ is required with overwhelming probability if the attribute specification $\mathcal{S}$ of the private key pk satisfies the access structure $\mathcal{W}$ of the ciphertext C (i.e., $\mathcal{S} \vdash \mathcal{W}$). On the contrary, if $\mathcal{S} \nvdash \mathcal{W}$, the output must be $\perp$.

Following the structure of the original article [2], we propose the *selective chosen plaintext attack* (sCPA) for assessing security. In sCPA, a challenge access structure $\mathcal{W}$ is initially specified by the attacker and then an interactive game is run. In the game, the attacker submits two plaintexts, one of which is randomly chosen and encrypted by the challenger. The attacker is then required to determine which plaintext corresponds to the given ciphertext.

More formally, consider the following indistinguishability game (IND-sCPA) between a challenger C that acts as a central authority and an adversary $\mathcal{A}$ that acts as an attacker:

Init. $\mathcal{A}$ chooses a challenge access structure $\mathcal{W}$ and prompts it to C.

Setup. C performs all the setup tasks and eventually prompts the public key pk to $\mathcal{A}$.

Key generation queries. $\mathcal{A}$ is allowed us to make a polynomial number of adaptive key generation queries on any attribute specification $\mathcal{S}$ such that $\mathcal{S} \nvdash \mathcal{W}$.

Challenge. $\mathcal{A}$ submits two messages of equal length M_0 and M_1 to C, who randomly chooses $b \in \{0, 1\}$ and returns to $\mathcal{A}$ the ciphertext associated with M_b, i.e. returns Enc(pk, $\mathcal{W}, M_b$).

Guess. $\mathcal{A}$ is allowed us to perform one more round of Key generation queries and eventually outputs a bit b'.

The advantage of an adversary $\mathcal{A}$ w.r.t. the previous game is defined as:

$$\mathrm{Adv}_{\mathcal{A}}^{\text{IND-sCPA}}(\boldsymbol{\sigma}) = \left| \mathbb{P}[b = b'] - \frac{1}{2} \right|.$$

We can further define a CP-ABE scheme to be secure against sCPA if, for any polynomial time adversary $\mathcal{A}$, the advantage $\mathrm{Adv}_{\mathcal{A}}^{\text{IND-sCPA}}(\boldsymbol{\sigma})$ is a negligible function in the security parameters σ.

As shown in the recent literature (see e.g., [32]), carefulness should, however, be made if protocol security is obtained via classical proofs, since their translation to the post-quantum context (in particular when interactive proofs are considered, like in this case) is not guaranteed. In light of these recent results, in the present work, we only discuss classical security. More work is to be carried out in this direction; although the scheme here presented is based, under the LWE assumption, on a problem that is commonly thought to be hard also in a quantum computing setting, the security of the scheme is assessed as in the study by Zhang and Zhang [2] only against a classical attacker.

4.2 Scheme

The scheme we work on was originally introduced by Zhang and Zhang [2], and it is somehow inspired by Shamir Secret Sharing [33] technique, where a randomly chosen shared secret $\mathbf{s} \hookleftarrow_{\$} \mathbb{Z}_q^n$ is hidden through multiple LWE samples and it is used in a LWE-PKE [29] fashion to build a ciphertext.

The main idea is to provide a given user with a fixed attribute (say 0) and a set of variable attributes $\mathcal{R} = \{1, \ldots, |\mathcal{R}|\}$ that can either be assigned (say i^+) or not (say i^-) for a total of $r = |\mathcal{R}| + 1$ attributes. Then, access structures $\mathcal{W}$ can either specify a given attribute (both in a positive or in a negative way) or not (actually providing them both).

More formally, a user attribute specification is a 2-partition $\mathcal{S} = (S^+, S^-)$ of $\mathcal{R}$ (i.e., $S^+ \cup S^- = \mathcal{R}$ and $S^+ \cap S^- = \emptyset$), while an access structure is a 2-covering $\mathcal{W} = (W^+, W^-)$ of $\mathcal{R}$ (i.e. $W^+ \cup W^- = \mathcal{R}$, but $W^+ \cap W^- = \emptyset$ is not required), where S^+ and W^+ represent the sets of positive attributes; moreover, S^- and W^- are the sets of negative attributes. In particular, we say that user attributes $\mathcal{S}$ *satisfy* the access structure $\mathcal{W}$ if $S^+ \subseteq W^+$ and $S^- \subseteq W^-$: in such case, we write $\mathcal{S} \vdash \mathcal{W}$; otherwise, we write $\mathcal{S} \nvdash \mathcal{W}$.

The advantage of providing user-attribute specifications as 2-partition consists in always having the same number of attributes, hence being able to build a matrix $\mathbf{D} \in \mathbb{Z}_q^{n \times mr}$ to be used in `GenSamplePre`. At the same time, the fixed attribute 0 provides an excellent point to evaluate the short basis needed by `GenSamplePre` (hence assuming $J = \{0\}$): in fact, it is fixed amongst all the possible user attribute specifications and it can be pre-evaluated efficiently via `TrapGen` algorithm.

Formally, the scheme is parameterised on the modulus q, the dimension m, the security parameter n, the Gaussian parameter s, and the error distribution χ with parameter α. Requirements on these parameters are analysed later in the next section.

The definition of the four functions from Scheme 1 are provided in Algorithms from 1 to 4.

Algorithm 1: Setup($n, m, q, \mathcal{R}$) → (pk, msk)

Input: the parameters $n, m, q \in \mathbb{N}$ and the set of attributes $\mathcal{R} = \{1, \ldots, r-1\}$

Output: the public key pk and the master secret key msk

$(\mathbf{B}_0, \mathbf{T}_{\mathbf{B}_0}) \leftarrow \text{TrapGen}(n, m, q)$;

for each $i \in \mathcal{R}$ **do**

  $\mathbf{B}_i^+, \mathbf{B}_i^- \hookleftarrow_{\$} U(\mathbb{Z}_q^{n \times m})$;

$\mathbf{u} \hookleftarrow_{\$} U(\mathbb{Z}_q^n)$;

$\text{pk} \leftarrow (\mathbf{B}_0, \{\mathbf{B}_i^+, \mathbf{B}_i^-\}_{i \in \mathcal{R}}, \mathbf{u})$;

$\text{msk} \leftarrow (\text{pk}, \mathbf{T}_{\mathbf{B}_0})$

Algorithm 2: KGen(msk, $\mathcal{S}$) → sk

Input: the master secret key msk and a user attribute spec. $\mathcal{S} = (S^+, S^-)$

Output: the user secret key sk holding the attribute specification $\mathcal{S}$ and the private secret $\mathbf{e} \hookleftarrow_{\$} D_{\mathbb{Z}^{mr}, s}$

for each $i \in \mathcal{R}$

  $\mathbf{A}_i \leftarrow \begin{cases} \mathbf{B}_i^+ & \text{if } i \in S^+ \\ \mathbf{B}_i^- & \text{if } i \in S^- \end{cases}$;

$\mathbf{A} \leftarrow [\mathbf{B}_0 || \mathbf{A}_1 || \ldots || \mathbf{A}_{|\mathcal{R}|}]$;

$\mathbf{e} \leftarrow \text{GenSamplePre}(\mathbf{A}, \mathbf{T}_{\mathbf{B}_0}, \{0\}, s, \mathbf{u})$;

$\text{sk} \leftarrow (\mathcal{S}, \mathbf{e})$;

Algorithm 3: Enc(pk, W, M) $\rightarrow C$

Input: the public key pk, an access structure $\mathcal{W} = (W^+, W^-)$ and a message $M \in \{0, 1\}$

Output: the ciphertext structure C holding the LWE-PKE encrypted message $z \in \mathbb{Z}_q$ and the coefficients $\mathbf{c}_i^{\pm}$ to allow the random secret retrieval (if the access structure is satisfied)

$\mathbf{s} \leftarrow_{\$} U(\mathbb{Z}_q^n)$;

$x_z \leftarrow_{\$} \chi$;

$z \leftarrow \langle \mathbf{u}^T, \mathbf{s} \rangle + x_z + M\lfloor \frac{q}{2} \rfloor$;

$\mathbf{x} \leftarrow_{\$} \chi^m$;

$\mathbf{c}_0 \leftarrow \langle \mathbf{B}_0^T, \mathbf{s} \rangle + \mathbf{x}$;

for each $i \in W^+$

$\mathbf{x} \leftarrow_{\$} \chi^m$;

$\mathbf{c}_i^+ \leftarrow \langle \mathbf{B}_i^{+T}, \mathbf{s} \rangle + \mathbf{x}$;

for each $i \in W^-$ **do**

$\mathbf{x} \leftarrow_{\$} \chi^m$;

$\mathbf{c}_i^- \leftarrow \langle \mathbf{B}_i^{-T}, \mathbf{s} \rangle + \mathbf{x}$;

$C \leftarrow (W, z, \mathbf{c_0}, \{\mathbf{c}_i^+\}_{i \in W^+}, \{\mathbf{c}_i^-\}_{i \in W^-})$;

Algorithm 4: Dec(C, sk) $\rightarrow M$ or $\perp$

Input: a ciphertext structure C and a secret key sk

Output: the message $M' \in \{0, 1\}$ which corresponds to the original message M if $|x_z - x'| < \frac{q}{4}$ (say $\leq \frac{q}{\ell}$ for each $\ell > 4$)

if $\mathcal{S} \nvDash \mathcal{W}$

return $\perp$;

for each $i \in \mathcal{R}$ **do**

$$\mathbf{y}_i \leftarrow \begin{cases} \mathbf{c}_i^+ & \text{if } i \in S^+ \\ \mathbf{c}_i^- & \text{if } i \in S^- \end{cases};$$

$\mathbf{y} \leftarrow [\mathbf{c}_0; \mathbf{y}_1; \ldots; \mathbf{y}_{|\mathcal{R}|}]$;

$a \leftarrow \langle \mathbf{e}^T, \mathbf{y} \rangle$; //$\langle \mathbf{e}^T, \mathbf{y} \rangle = \langle \mathbf{e}^T, \langle \mathbf{A}^T, \mathbf{s} \rangle \rangle + \langle \mathbf{e}^T, \mathbf{x} \rangle = \langle \mathbf{u}^T, \mathbf{s} \rangle + x'$

$b \leftarrow z - a$; //$z - a = x_z - x' + M\lfloor \frac{q}{2} \rfloor$

$$M' \leftarrow \begin{cases} 1 & \text{if } \lfloor \frac{q}{4} \rfloor \leq b \leq \lfloor \frac{3q}{4} \rfloor \\ 0 & \text{otherwise} \end{cases};$$

4.3 Parameter requirements and security

We analyse the scheme parameters considering the requisites (i) of having a correct decryption (cf. Algorithm 4), (ii) required by Proposition 2, and (iii) required by Functions 1–4. Then, we obtain:

(i) $m \geq \lceil 6n \log q \rceil$ as required by `TrapGen` (Function 2);

(ii) $s \geq \|\tilde{\mathbf{T}}_{\mathbf{B}_0}\| \cdot \omega(\sqrt{\log(mr)})$ as required by `GenSamplePre` (Function 4) and by the security proof;

(iii) $|x_z - x'| \leq \frac{q}{\ell}$, with $\ell > 4$, for correct decryption (Algorithm 4);

(iv) $\alpha q \geq 2\sqrt{n}$ for LWE hardness (Proposition 2).

(i) suggests us to parameterise m over a value $\delta \in \mathbb{R}$ being such that $n^\delta > \lceil \log q \rceil$, thus obtaining

$$m = 6n^{1+\delta}.$$

Furthermore, we know from Function 2 that $\|\tilde{\mathbf{T}}_{\mathbf{B}_0}\| \leq O(\sqrt{n \log q})$, or, in other terms, that $\|\tilde{\mathbf{T}}_{\mathbf{B}_0}\| \leq O(\sqrt{m})$; hence, from the second condition, we obtain

$$s = \sqrt{m} \cdot \omega(\sqrt{\log(mr)}).$$

In order to tackle (iii), we recall from Lemma 2 that $|x_z| \leq q\alpha \cdot \omega(\sqrt{\log m}) + \frac{1}{2}$ and $|x'| = |\langle \mathbf{e}^T, \mathbf{x} \rangle| \leq \|\mathbf{e}\| \cdot q\alpha \cdot \omega(\sqrt{\log m}) + \|\mathbf{e}\| \frac{\sqrt{m}}{2}$ and from Lemma 1 that $\|\mathbf{e}\| \leq s\sqrt{mr}$; hence, due to triangular inequality, we obtain

$$\begin{aligned}
|x_z - x'| &\leq |x_z| + |x'| \\
&\leq q\alpha \cdot (\omega(\sqrt{\log m}) + \|\mathbf{e}\| \cdot \omega(\sqrt{\log(mr)})) + \frac{1}{2}(1 + \|\mathbf{e}\|\sqrt{m}) \\
&\leq q\alpha \cdot s\sqrt{mr} \cdot \omega(\sqrt{\log(mr)}) + \frac{1}{2}(1 + sm\sqrt{r}) \\
&\leq sq\alpha\sqrt{mr} \cdot \omega(\sqrt{\log(mr)}) + smr.
\end{aligned}$$

Plugging the inequality in (iii) and letting $\hat{\omega} = \omega(\sqrt{\log(mr)})$, we obtain

$$\begin{aligned}
\ell sq\alpha \cdot \sqrt{mr} \cdot \hat{\omega} + \ell smr \leq q \\
q \cdot (\ell s\alpha \cdot \sqrt{mr} \cdot \hat{\omega} - 1) \leq \ell smr,
\end{aligned}$$

which suggests us requiring

$$\alpha = (s \cdot \sqrt{mr} \cdot \omega(\sqrt{\log(mr)}))^{-1},$$

hence obtaining from the previous inequality that

$$q \cdot (\ell \cdot \omega(1) - 1) \leq \ell smr.$$

Furthermore, in order to satisfy (iv), we obtain

$$q > 2\sqrt{n} \cdot \alpha^{-1} = s \cdot \sqrt{4nmr} \cdot \hat{\omega}.$$

Recalling from (i) that $m > 4n$, a suitable solution is given by:

$$q = smr \cdot \omega(\sqrt{\log(mr)}),$$

solution yet still satisfying the sequence of inequalities we built for (iii).

We can resume the above-stated conditions as follows:

$$\begin{aligned}
m &= 6n^{1+\delta}, \quad \text{with } \delta \in \mathbb{R} | n^\delta > \lceil \log q \rceil, \\
s &= \sqrt{m} \cdot \omega(\sqrt{\log(mr)}), \\
q &= smr \cdot \omega(\sqrt{\log(mr)}), \\
\alpha &= (s \cdot \sqrt{mr} \cdot \omega(\sqrt{\log(mr)}))^{-1},
\end{aligned} \tag{†}$$

in order to provide the scheme security claim.

Proposition 3. (Theorem 1 from [2]) *Let $\chi = \bar{\Psi}_\alpha$ and let m, s, q, and α be as from* (†)*. Then, if* $\mathrm{LWE}_{q,\chi}$ *is hard, the CP-ABE scheme* (`Setup`, `KGen`, `Enc`, `Dec`) *defined by Algorithms 1–4 is secure against sCPA.*

In particular, if there exists an adversary $\mathcal{A}$ that breaks its sCPA security with advantage ε, then there exists an algorithm $\mathcal{B}$ solving $\mathrm{LWE}_{q,\chi}$ *with probability ε.*

5 mR_{LWE}-CP-ABE

In this section, we describe the architecture of mR_{LWE}-CP-ABE, our new CP-ABE encryption scheme, based on lattices, able to efficiently revoke a target user; in particular, we propose a general scheme that extends Scheme 1 by considering the addition of SMs. Then, we provide a full description of the mediated scheme built on top of Zhang and Zhang's scheme from Section 4.2, and we discuss the requirements changes on the scheme parameters.

5.1 The system model

We rely on a new server-aided approach to provide a fast and reliable solution in order to avoid some inefficiency intrinsically derived by direct and indirect revocation mechanisms. Our system is logically composed by four kinds of entities:

- the key generation server (KGS): a trusted server that is able to generate a public key and, for each user, the corresponding secret key.
- A set of k SM: each SM is a semi-trusted entity that has access to the *mediator keys* of a (sub)set of users.
- The data owner: someone who wants to encrypt some data for a set of, possibly unknown, users.
- A set of users that belong to the system: each user has an attribute specification that specifies his/her access rights. The attributes are associated with the secret user key generated by the KGS.

We define our scheme on top of the one introduced in Section 4.2; namely, for each user, the KGS generates a tuple of keys, $(\mathrm{sk}, \mathrm{mk}_1, \ldots, \mathrm{mk}_k)$; the sk is the user key and it is given to the user while the keys mk_j, with $1 \le j \le k$, are the mediator keys that are distributed one for each SM involved. In order for a user to successfully decrypt a ciphertext, two conditions are required: first, as usually in CP-ABE, the user must have an attribute specification $\mathcal{S}$ that satisfies the ciphertext policy; second, all the k SMs must contribute to the decryption.

Scheme 2. (Revocable CP-ABE) Our revocable CP-ABE scheme consists of five algorithms:

- $\mathrm{Setup}(\boldsymbol{\sigma}, \mathbb{R}) \to (\mathrm{msk}, \mathrm{pk})$ is the initialisation algorithm executed by the KGS. It behaves like in regular CP-ABE.
- $\mathrm{MKGen}(\mathrm{msk}, \mathcal{S}, k) \to (\mathrm{sk}, \{\mathrm{mk}_j\}_{j=1}^{k})$ is the algorithm the KGS runs to accredit a user with an attribute specification $\mathcal{S}$. In contrast to regular CP-ABE, the key is segmented in $k + 1$ parts, k of which are provided to k SMs. The specified access structure $\mathcal{W}$ is stored with the user private key sk only, making mediators unaware of users' capabilities.
- $\mathrm{Enc}(\mathrm{pk}, \mathcal{W}, M) \to C$ is the encryption algorithm the data owner runs to encrypt the message M. It behaves like in regular CP-ABE schemes. As a consequence, data owners are not influenced by any means by this scheme.
- $\mathrm{MDec}(C, \mathrm{sk}) \to M$ or $\perp$ is the decryption algorithm a user runs to retrieve the message M associated with the ciphertext C. It requires the cooperation of all the k SMs that should return the result of PDec in order to make the user able to evaluate $M' = M$ with overwhelming probability (if $\mathcal{S} \vdash \mathcal{W}$). Like in regular CP-ABE schemes, if $\mathcal{S} \nvdash \mathcal{W}$, the output must be $\perp$ (regardless of the possible collusion with SMs).
- $\mathrm{PDec}(\mathbf{y}, \mathrm{mk}) \to a$ is the algorithm run by SMs that allows them to produce a partial decryption information a from $\mathbf{y}$ and the mediator key mk. Here, $\mathbf{y}$ is derived from the ciphertext C by the user requiring the partial decryption within MDec function.

If a user is revoked, the KGS only needs to send this information to the SMs that have a mediator key for that user and they will stop to collaborate in the decryption process. In particular, it is sufficient that just a single SM refuses to cooperate to defeat the decryption process. This guarantees that, if at least one SM follows the protocol, a revoked user cannot decrypt anymore.

Please note that, differently from previous schemes in the literature, we do not require to update keys or re-encrypt ciphertexts, in order to revoke a user, we just need to notify the SMs. Furthermore, already encrypted ciphertexts that have not been decrypted before revoke occurs, are evenly secure against the revoked user. This is also different from direct revocation where the CRL, as this revocation process does, does not involve the encryption process. In order to support fast and secure revocation, our system incurs, of course, in some overhead compared to that in the study by Zhang and Zhang [2]. For instance,

- The KGS has to generate $k + 1$ keys for each user;
- The decryption process of a ciphertext C requires $k + 1$ partial decryption plus k error generation that is added by each SM to protect their mediator keys.

To experimentally evaluate the impact of revocation, we report some experiments in Section 8.

It is important to highlight that, despite SM's need to be reachable at decryption time, hence making the protocol interactive, the approach preserves the advantages of CP-ABE over classical PKE schemes. In fact, data owners still produce encrypted data offline and without suffering any overhead w.r.t. non-mediated CP-ABE schemes. Furthermore, access to data is still preserved by data owners and final users only, since SMs have blind-access to data.

It is, however, important to note that the here proposed scheme is weak against any kind of denial of service attack since being a single SM unreachable or uncooperative causes the decryption procedure to fail. This outcome is common in any (k, k) secret sharing scheme, where k participants out of k are required to collaborate to retrieve the key. Just analogously to secret sharing schemes, it is hence possible, depending on the use case requirement, to mitigate this issue by building a (t, k) threshold scheme where only t out of k SMs are required to carry out the partial decryption. Simpler solutions are also available, e.g., the possibility to give SM keys with a certain amount of redundancy, actually creating a (possibly) non-homogeneous network of SMs. Such a solution also allows us to distribute the decryption workload amongst multiple parties. Further description of these approaches is, however, out of the scope of this article and will be the object of further future analysis.

5.2 Scheme

mR$_{LWE}$-CP-ABE shares the same parameter structure with the regular CP-ABE scheme presented in Section 4. Requirements on these parameters are very similar too and are discussed in the next section. For this reason, Setup function is defined as in Algorithm 1 without any particular changes.

Concurrently, as described in Section 5.1, the encryption procedure is not modified by the mediation process; hence, Enc function is defined as in Algorithm 3.

The definition of the three remaining functions defining a revocable CP-ABE scheme follows in Algorithms 5 to 7.

Algorithm 5: MKGen(msk, $\mathcal{S}, k$) $\rightarrow$ (sk, $\{\text{mk}_j\}_{j=1}^{k}$)

Input: the master secret key msk, a user attribute specification $\mathcal{S} = (S^+, S^-)$, and the number of mediators k

Output: the user secret key sk holding the attribute specification $\mathcal{S}$ and the private secret $\mathbf{e} \leftarrow_{\$} D_{\mathbb{Z}^{mr},s}$

Output: the mediator secret key mk_j holding the private secret $\mathbf{mk}_j \leftarrow_{\$} D_{\mathbb{Z}^{mr},s}$, for each $0 < j \leq k$

for each $i \in \mathcal{R}$ **do**

$$\left| \mathbf{A}_i \leftarrow \begin{cases} \mathbf{B}_i^+ & \text{if } i \in S^+ \\ \mathbf{B}_i^- & \text{if } i \in S^- \end{cases}; \right.$$

$\mathbf{A} \leftarrow [\mathbf{B}_0 || \mathbf{A}_1 || \ldots || \mathbf{A}_{|\mathcal{R}|}]$;

for $j = 1, \ldots, k$

$\lfloor \mathbf{u}_j \hookleftarrow_{\$} U(\mathbb{Z}_q^n)$;
$\mathbf{mk}_j \leftarrow \text{GenSamplePre}(\mathbf{A}, \mathbf{T}_{\mathbf{B}_0}, \{0\}, s, \mathbf{u}_j)$;
$\mathbf{u}_0 \leftarrow \mathbf{u} - \sum_{j=1}^{k} \mathbf{u}_j$;
$\mathbf{e} \leftarrow \text{GenSamplePre}(\mathbf{A}, \mathbf{T}_{\mathbf{B}_0}, \{0\}, s, \mathbf{u}_0)$;
$\text{sk} \leftarrow (\mathcal{S}, \mathbf{e})$;
$\text{mk}_j \leftarrow (\mathbf{mk}_j)$;

Algorithm 6: $\text{PDec}(\mathbf{y}, \text{mk}_j) \rightarrow a_j$

Input: a vector $\mathbf{y}$ holding the information about the shared secret of a ciphertext and a mediator key mk_j
Output: the decryption information a_j

$x_j \hookleftarrow_{\$} \chi$;
$a_j \leftarrow \langle \mathbf{mk}_j^T, \mathbf{y} \rangle + x_j$

$\langle \mathbf{mk}_j^T, \mathbf{y} \rangle + x_j = \langle \mathbf{mk}_j^T, \langle \mathbf{A}^T, \mathbf{s} \rangle \rangle + \langle \mathbf{mk}_j^T, \mathbf{x} \rangle + x_j = \langle \mathbf{u}_j^T, \mathbf{s} \rangle + x_j' + x_j$

Algorithm 7: $\text{MDec}(C, \text{sk}) \rightarrow M$ or $\perp$

Input: a ciphertext structure C and a user secret key sk
Output: the message $M' \in \{0, 1\}$ which corresponds to the original message M if $|x_z - \sum_{j=0}^{k} x_j' - \sum_{j=1}^{k} x_j| < \frac{q}{4}$ (say $\leq \frac{q}{\hat{\ell}}$ for each $\hat{\ell} > 4$)

if $\mathcal{S} \nvDash \mathcal{W}$
$\lfloor$ **return** $\perp$
for each $i \in \mathcal{R}$ **do**
$\left\lfloor \mathbf{y}_i \leftarrow \begin{cases} \mathbf{c}_i^+ & \text{if } i \in S^+ \\ \mathbf{c}_i^- & \text{if } i \in S^- \end{cases}; \right.$
$\mathbf{y} \leftarrow [\mathbf{c}_0 || \mathbf{y}_1 || \ldots || \mathbf{y}_{|\mathcal{R}|}]$;
$a_0 \leftarrow \langle \mathbf{e}^T, \mathbf{y} \rangle$; $//\langle \mathbf{e}^T, \mathbf{y} \rangle = \langle \mathbf{e}^T, \langle \mathbf{A}^T, \mathbf{s} \rangle \rangle + \langle \mathbf{e}^T, \mathbf{x} \rangle = \langle \mathbf{u}_0^T, \mathbf{s} \rangle + x_0'$
for $j = 1, \ldots, k$ // Queries to SMs can be performed asynchronously
$\lfloor$ Request to the j-th mediator $a_j \leftarrow \text{PDec}(\mathbf{y}, \circ)$;
$a \leftarrow \sum_{j=0}^{k} a_j$; $//\sum_{j=0}^{k} a_j = \sum_{j=0}^{k} (\langle \mathbf{u}_j^T, \mathbf{s} \rangle + x_j') + \sum_{j=1}^{k} x_j = \langle \mathbf{u}^T, \mathbf{s} \rangle + \sum_{j=0}^{k} x_j' + \sum_{j=1}^{k} x_j$
$b \leftarrow z - a$; // $z - a = x_z - \sum_{j=0}^{k} x_j' - \sum_{j=1}^{k} x_j + M\lfloor \frac{q}{2} \rfloor$
$M' \leftarrow \begin{cases} 1 & \text{if } \lfloor \frac{q}{4} \rfloor \leq b \leq \lfloor \frac{3q}{4} \rfloor \\ 0 & \text{otherwise} \end{cases}$;

5.3 Parameter requirements

Requirements introduced in Section 4.3 still hold. However, error grows higher in `MDec` if compared to the lattice-based CP-ABE scheme `Dec`. In fact, the requirement for correct decryption is as follows:

$$\left| x_z - \sum_{j=0}^{k} x_j' - \sum_{j=1}^{k} x_j \right| \leq \frac{q}{\hat{\ell}}.$$

Do note that $\{x_j\}_{j=1}^{k}$ are sampled from the same distribution as x_z and $\{x'_j\}_{j=1}^{k}$ are obtained as it was for x' in the original algorithm; hence, Lemma 2 still applies. Due to triangular inequality and following the same reductions as mentioned earlier, we obtain

$$\left| x_z - \sum_{j=0}^{k} x'_j - \sum_{j=1}^{k} x_j \right| \leq |x_z| + \sum_{j=0}^{k} |x'_j| + \sum_{j=1}^{k} |x_j| \leq (k+1)sq\alpha\sqrt{mr} \cdot \hat{\omega} + (k+1)smr.$$

Plugging the inequality in the requirement for decryption, we obtain

$$(k+1)\hat{\ell}sq\alpha \cdot \sqrt{mr} \cdot \hat{\omega} + (k+1)\hat{\ell}smr \leq q$$
$$q \cdot ((k+1)\hat{\ell}s\alpha \cdot \sqrt{mr} \cdot \hat{\omega} - 1) \leq (k+1)\hat{\ell}smr,$$

whose solution is comparable to the one of the original scheme if we consider $\ell = \hat{\ell} \cdot (k+1)$ since the only requirement imposed on ℓ is $\ell > 4$, which still holds.

6 About threat and security to mR$_{LWE}$-CP-ABE

In this section, we analyse the security of our proposed scheme by first defining a suitable threat model, including the possible collusion of one or more mediators with the attacker. Then, we discuss the security of our solution against the IND-sCPA game described in Section 4.1 suitably tackled in terms of the proposed threat model.

6.1 Threat model

We now describe the threat model of our system by means of five entities: the KGS, the set of SMs, the data owner, the set of users, and an external attacker. We remember that the KGS is a trusted entity, whereas the SMs are semi-trusted. The data owner is also trusted, whereas the users and, of course, the external attacker are untrusted. We identify the following possible threats that may affect our system:

- SMs collusion: multiple SMs involved in the decryption of the same ciphertext may collude together to decrypt without the user's aid;
- Users collusion: multiple users may collude to decrypt a ciphertext they are not authorised to;
- Ineffective revocation: a revoked user is still able to decrypt;
- DOS-decryption: an attacker who compromises at least one SM can prevent legitimate users to decrypt.

We formally analyse the security of SMs and user collusion and the ineffective revocation in Theorem 1.

The DOS-decryption attack, indeed, can be mitigated by providing redundant mediated keys to SM as already described in Section 5.1.

6.2 Security analysis

The here presented mediated scheme is equivalent to the original scheme from the point of view of an external attacker. In fact, the encryption and decryption functions behave the same as in the original scheme but for the addition of more noise (the more the mediators, the higher the noise). We can further claim, analogously to Proposition 3, the security of the scheme under sCPA:

Theorem 1. (Security of $\mathrm{mR}_{\mathrm{LWE}}$-CP-ABE (external)) *Let* $\chi = \bar{\Psi}_\alpha$ *and let* m, s, q, *and* α *be as from* (†). *Then, if* $\mathrm{LWE}_{q,\chi}$ *is hard, the revocable CP-ABE scheme* (`Setup`, `MKGen`, `Enc`, `PDec`, `MDec`) *defined by Algorithms 1, 5, 3, 6, and 7 is secure against the sCPA.*

In particular, if there exists an adversary $\mathcal{A}$ *that breaks its sCPA security, then there exists an adversary* $\mathcal{B}$ *that solves the* $\mathrm{LWE}_{q,\chi}$ *decision problem.*

The proof of the theorem is analogous to the one from the study by Zhang and Zhang [2]; however, we report it for completeness.

Proof. Assume there exists a polynomial-time adversary $\mathcal{A}$ capable of breaking IND-sCPA for the mediated scheme with advantage ε using at most q key generation queries by obtaining both user and mediator keys.

Let $O(\circ)$ be an oracle that samples always either from $A_{\mathbf{s},\chi}$ or from the uniform distribution $U(\mathbb{Z}_q)$. Let $\mathcal{B}$ be an attacker who cooperates with $\mathcal{A}$ and wants to decide whether one of the two distributions $O(\circ)$ is sampling from.

The idea of the cooperation is to build a CP-ABE game – with $\mathcal{A}$ as the attacker and $\mathcal{B}$ as the challenger – that can be won with probability noticeably greater than $\frac{1}{2}$ if and only if $O(\circ)$ is sampling from $_{\mathbf{s},\chi}$. Assume that such a game exists, and then, $\mathcal{B}$ can discriminate between the two distributions.

We recall that in order to run the game, the challenger $\mathcal{B}$ is only required to be able to (i) provide a public key to $\mathcal{A}$, (ii) encrypt a message, and (iii) make q generations of valid keys (with respect to the provided public key) on attribute specifications that do not satisfy the challenged access structure; decryption is hence not required as well as being able to generate secret keys for attribute specification satisfying the challenged access structure.

Let us formally define the game:

Init. $\mathcal{A}$ chooses a challenge $\mathcal{W} = (W^+, W^-)$ and prompts it to $\mathcal{B}$.

Setup. $\mathcal{B}$ samples $(\mathbf{a}, y) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$ pairs multiple times from $O(\circ)$ in order to build the matrices $\mathbf{B}_0$ and $\mathbf{B}_i^{\pm}$ needed by the chosen access structure (out of the vectors $\mathbf{a}_i$) and to save (potentially) LWE-valid vectors $\mathbf{c}_i$ for the ciphertext creation. The total number of samples required is $(|S^+| + |S^-| + 1)\cdot m + 1$, and they are used to build the following couples:

- $(\mathbf{B}_0, \mathbf{v}_0) \in \mathbb{Z}_q^{n\times m} \times \mathbb{Z}_q^m$;
- $(\mathbf{u}, v_u) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$;
- $(\mathbf{B}_i^+, \mathbf{v}_i^+) \in \mathbb{Z}_q^{n\times m} \times \mathbb{Z}_q^m$ for each $i \in S^+$;
- $(\mathbf{B}_i^-, \mathbf{v}_i^-) \in \mathbb{Z}_q^{n\times m} \times \mathbb{Z}_q^m$ for each $i \in S^-$.

Then, in order to create the missing matrices $\mathbf{B}_i^+$ and $\mathbf{B}_i^-$ (respectively, for $i \notin S^+$ and $i \notin S^-$) and coherently being able to run `MKGen` algorithm on $\mathcal{S}$ such that $\mathcal{S} \nvdash \mathcal{W}$, the challenger $\mathcal{B}$ computes the following:

- $(\mathbf{B}_i^+, \mathbf{T}_{\mathbf{B}_i^+}) \leftarrow \mathrm{TrapGen}(n, m, q)$ for each $i \notin S^+$;
- $(\mathbf{B}_i^-, \mathbf{T}_{\mathbf{B}_i^-}) \leftarrow \mathrm{TrapGen}(n, m, q)$ for each $i \notin S^-$.

Finally, the challenger stores $(\{\mathbf{T}_{\mathbf{B}_i^+}\}_{i\in S^+}, \{\mathbf{T}_{\mathbf{B}_i^+}\}_{i\in S^+})$ for the key generation, stores $(\mathbf{v}_0, v_u, \{\mathbf{v}_i^+\}_{i\in S^+}, \{\mathbf{v}_i^-\}_{i\in S^-})$ for the ciphertext creation, and outputs the public key $\mathrm{pk} = (\mathbf{B}_0, \{\mathbf{B}_i^+, \mathbf{B}_i^-\}_{i\in\mathcal{R}}, \mathbf{u})$ to the attacker $\mathcal{A}$.

Keygen query. Upon receiving a user attribute specification $\mathcal{S}$ from $\mathcal{A}$, if $\mathcal{S} \vdash \mathcal{W}$, then $\mathcal{B}$ outputs $\perp$. Otherwise, there exists at least one attribute $i \in S^+$ such that $i \notin W^+$ or $i \in S^-$ such that $i \notin W^-$; let $\hat{\mathbf{T}}$ be the short basis generated by `TrapGen` during setup associated with such an attribute. $\mathcal{B}$ finally runs and outputs $\mathrm{MKGen}((\mathrm{pk}, \hat{\mathbf{T}}), \mathcal{S})$ to $\mathcal{A}$. Do note that the so-formed master secret key is valid for $\mathcal{S}$ (and for all the user specifications containing i as does $\mathcal{S}$) since, according with Function 4, `GenSamplePre` requires whatever short basis generated from a subset of m linearly independent vectors of $\mathbf{A}$ (the matrix in $\mathbb{Z}_q^{n\times mk}$ defined in Algorithm 5).

Challenge. The attacker $\mathcal{A}$ submits $M_0, M_1 \in \{0, 1\}$ to the challenger $\mathcal{B}$, who randomly chooses $b \in \{0, 1\}$ and returns the (possibly valid) ciphertext associated with M_b. However, since $\mathcal{B}$ wants to output a valid

ciphertext only if $O(\circ)$ is sampling from $A_{\mathbf{c},s}$, the idea is to use the stored values from the setup in order to emulate the LWE instances of the Enc function. Therefore, $\mathcal{B}$ computes and outputs $C = (W, z, \mathbf{c_0}, \{\mathbf{c}_i^+\}_{i\in W^+}, \{\mathbf{c}_i^-\}_{i\in W^-})$ with:

- $z \leftarrow v_u + M_b\lfloor\frac{q}{2}\rfloor$, where v_u emulates $\langle \mathbf{u}^T, \mathbf{s}\rangle + x_z$;
- $\mathbf{c}_0 \leftarrow \mathbf{v}_0$ to emulate $\langle \mathbf{B}_0^T, \mathbf{s}\rangle + \mathbf{x}$;
- $\mathbf{c}_i^+ \leftarrow \mathbf{v}_i^+$ to emulate $\langle \mathbf{B}_i^{+T}, \mathbf{s}\rangle + \mathbf{x}$, for each $i \in W^+$;
- $\mathbf{c}_i^- \leftarrow \mathbf{c}_i^-$ to emulate $\langle \mathbf{B}_i^{-T}, \mathbf{s}\rangle + \mathbf{x}$, for each $i \in W^-$.

$\mathcal{A}$ is allowed us to make more key generation queries after the challenge has been set. Eventually, it outputs a guess b' for b that is correct either with probability $\frac{1}{2} + \varepsilon$ if $O(\circ)$ is sampling from $A_{\mathbf{c},s}$ or with probability $\frac{1}{2}$ if it is sampling from $U(\mathbb{Z}_q^n \times \mathbb{Z}_q)$. Hence, $\mathcal{B}$ guesses $A_{\mathbf{c},s}$ if $b = b'$ (i.e. $\mathcal{A}$ is correct) or guesses $U(\mathbb{Z}_q^n \times \mathbb{Z}_q)$ if $b \neq b'$ (i.e. $\mathcal{A}$ is wrong).

Do note that if $O(\circ)$ is sampling from $A_{\mathbf{c},s}$, then $\mathcal{B}$ guesses right with the same non-negligible advantage as $\mathcal{A}$ does. So, if such $\mathcal{A}$ does exist, $\mathcal{B}$ can solve LWE problems, which yields the claim. □

However, due to the revocation requirement, security must also be ensured if the attacker is one or more users, in the sense that

(i-u) A user is not able to decrypt as far as at least an SM denies its cooperation;
(ii-u) A user cannot reject a revocation in polynomial time, even after a polynomial number of correct decryptions (meaning that he can neither break nor forge an SM key);
(iii-u) Multiple users cannot collude to retrieve an SM key or to decrypt a message;

and if the attacker is one or more SMs, in the sense that:

(i-m) One or more SM cannot collude to decrypt a message, even upon receiving many mediator keys from different users;
(ii-m) One or more SM cannot collude to retrieve a user secret $\mathbf{e}$.

We do note that mediator keys and user keys are complementary in the decryption phase and they are generated all starting from $\mathbf{u}_i$ distributed uniformly at random: $\mathbf{u}_0$ is given by the difference of vector distributed uniformly at random; hence, it is still distributed uniformly at random. Since $\mathbf{u}_i$ are generated randomly for each user, it follows that different users $\mathbf{e}$ are independent one from the other; hence, combining information from different keys ensures no further knowledge, ensuring (iii-u).

A similar outcome can be derived for SMs holding mediator keys related to different users. Moreover, SM and users receive no information about mutual keys by-design and mediator keys, if equipped with $\mathcal{S}$, are equivalent to the user ones (i.e., SM keys are weaker than user ones), we have that (i-u) implies (i-m) (i.e., if a user is not able to decrypt without even a single SM, then decryption cannot occur without the collaboration of k parties out of the $k + 1$, no matters which one is missing).

We claim the following theorem to prove (i-u):

Theorem 2. (Security of PDec algorithm (break)) *Let $\chi = \bar{\Psi}_\alpha$ and let m, s, q, and α be as from* (†). *Then, if* $\mathrm{LWE}_{q,\chi}$ *is hard, the CP-ABE scheme* (Setup, MKGen, Enc, PDec, MDec) *defined by Algorithms* 1, 5, 3, 6, *and* 7 *is secure against the sCPA carried out by a user if at least one mediator does not participate in the decryption.*

In particular, if there exists an adversary $\mathcal{A}$ that breaks its sCPA security, then there exists an adversary $\mathcal{B}$ that solves the $\mathrm{LWE}_{q,\chi}$ *decision problem.*

Proof. If at least a mediator (say $\hat{j}$) does not take part in the decryption procedure, from client's perspective, the problem resembles solving the non-mediated version of the scheme with higher noise on z. In fact, he can compute $z' \leftarrow z - \sum_{j=0, j\neq\hat{j}}^{k} a_j$ and $C' = (W, z', \mathbf{c_0}, \{\mathbf{c}_i^+\}_{i\in W^+}, \{\mathbf{c}_i^-\}_{i\in W^-})$ is a valid ciphertext (apart from the potentially higher noise) for the non-mediated scheme with public key $\mathrm{pk} = (\mathbf{B}_0, \{\mathbf{B}_i^+, \mathbf{B}_i^-\}_{i\in\mathcal{R}}, \mathbf{u}_{\hat{j}})$.

The only advantage the client has is the knowledge of its $\mathbf{e}$ that, however, reveals further information neither about the short basis $\mathbf{T}_{\mathbf{B}_0}$ nor about other keys[2] since they are obtained from GenSamplePre applied to uniformly random $\mathbf{u}$.

It follows that if there exists an adversary $\mathcal{A}$ that breaks sCPA for the mediated scheme under these assumptions, then there exists an adversary $\mathcal{A}'$ that breaks sCPA for the non-mediated scheme, and hence, by Proposition 3, there exists an adversary $\mathcal{B}$ that solves the $\mathrm{LWE}_{q,\chi}$ decision problem. □

Furthermore, by design, SMs receive no information about the decryption procedure when a request is submitted by a user; hence, they can learn anything about the user's secret neither from the query itself nor from other sources. Analogously, (ii-m) follows.

Do also consider that mediators receive no information about the attribute specification $\mathcal{S}$ as well; however, if mediators do have access to the database of ciphertexts and can guess which ciphertext was delivered by the user, they can guess the attribute specification by matching vector $\mathbf{y}$ with the allowed $\mathbf{c}_i^{\pm}$. In particular, guessing the correct match between delivered $\mathbf{y}$ and ciphertext is mainly a combinatorial matter that has not received much attention in the literature. However, we point out that, whenever it would get important to preserve the privacy between which users required which ciphertext (e.g., to prevent user profiling), a possible solution for the user would be to protect $\mathbf{y}$ by adding a noise $\mathbf{x} \leftarrow_{\$} \chi^{mr}$. It is out of the scope of this article to show the complete proof of correctness; however, do note that $\mathbf{y} + \mathbf{x}$ would still be considered as a valid LWE sample (see also later the proof of Theorem 3) with higher noise and decryption would still be correct with some correction on total noise size.

Finally, we claim the following theorem that tackles (ii-u):

Theorem 3. (Security of PDec algorithm) *Let* $\chi = \bar{\Psi}_\alpha$ *and let* m, s, q, *and* α *be as from* (†). *Then, if* $\mathrm{LWE}_{q,\chi}$ *is hard, the function* PDec *defined by Algorithm 6 is hard to*:

- *Break, in the sense that there exists no polynomial-time algorithm to retrieve* mk *from a polynomially bounded number of pair* $(\mathbf{y}, a)$, *where* $a \leftarrow \mathrm{PDec}(\mathbf{y}, \mathrm{mk})$.
- *Forge, in the sense that there exists no polynomial-time algorithm to evaluate* a^* *from an arbitrary* $\mathbf{y}^*$, *where* $a^* \leftarrow \mathrm{PDec}(\mathbf{y}^*, \mathrm{mk})$, *without knowing* mk *from a priorly obtained polynomially bounded number of pair* $(\mathbf{y}_i, a_i)$, *where* $a_i \leftarrow \mathrm{PDec}(\mathbf{y}_i, \mathrm{mk})$.

In particular, if there exists an adversary $\mathcal{A}$ *that breaks (forges)* PDec, *then there exists an adversary* $\mathcal{B}$ *that breaks the* $\mathrm{LWE}_{q,\chi}$ *search (decision) problem.*

Proof. We recall $\mathbf{y} \leftarrow_{\$} A_{\mathbf{s},\chi}$ since $\mathbf{y} = \langle \mathbf{A}, \mathbf{s} \rangle + \mathbf{x}$; hence, $\mathbf{y}$ is indistinguishable from the uniform distribution for the $\mathrm{LWE}_{q,\chi}$ hardness.

It is easy to see that $a \in A_{\mathbf{mk},s}$; in fact, $a = \langle \mathbf{mk}^T, \mathbf{y} \rangle + x$, where $\mathbf{y}$ is statistically uniform to random and $x \leftarrow_{\$} \chi$. It follows that PDec is an actual instance of $\mathrm{LWE}_{q,\chi}$, hence proving the claim. □

7 Multiple-bit encryption

The original CP-ABE scheme introduced in Section 4 was also proposed as an N-bit encryption scheme (with $N \in \mathbb{N}$), where the same shared secret $\mathbf{s}$ was used to encrypt a vector of message bits $\mathbf{M} \in \{0, 1\}^N$.

The authors introduced a public matrix $\mathbf{U} \in \mathbb{Z}_q^{n \times N}$ and a user secret matrix $\mathbf{E}$ of size $mr \times N$, where each of the N columns is generated by applying GenSamplePre to a different column of $\mathbf{U}$. Here, encryption works analogously, with the only difference within the evaluation of z, now being a vector $\mathbf{z}$:

2 If this was the case, do note also that users could forge other users keys at first glance.

$$\mathbf{z} \leftarrow \langle \mathbf{U}^T, \mathbf{s} \rangle + \mathbf{x}_z + \mathbf{M}\lfloor \frac{q}{2} \rfloor, \quad \text{with } \mathbf{x}_z \hookleftarrow_{\$} \chi^N.$$

Decryption, at the same glance, does not require any further care; in fact, once retrieved the suitable $\mathbf{y}$ vector, we can perform:

$$\mathbf{a} \leftarrow \langle \mathbf{E}^T, \mathbf{y} \rangle = \langle \mathbf{E}^T, \langle \mathbf{A}^T, \mathbf{s} \rangle \rangle + \langle \mathbf{E}^T, \mathbf{x} \rangle = \langle \mathbf{U}^T, \mathbf{s} \rangle + \mathbf{x}',$$
$$\mathbf{b} \leftarrow \mathbf{z} - \mathbf{a} = \mathbf{x}_z - \mathbf{x}' + \mathbf{M}\left\lfloor \frac{q}{2} \right\rfloor,$$

and, finally, by identifying $\mathbf{M} = (M_1, \ldots, M_N)$ and $\mathbf{b} = (b_1, \ldots, b_N)$:

$$M_i \leftarrow \begin{cases} 1, & \text{if } \left\lfloor \frac{q}{4} \right\rfloor \le b_i \le \left\lfloor \frac{3q}{4} \right\rfloor, \quad \text{for } i = 1, \ldots, N. \\ 0, & \text{otherwise} \end{cases}$$

A notable advantage of this approach is given by the size of the ciphertext, since only a single copy of $\mathbf{c}_0$, $\mathbf{c}_i^+$ (for each $i \in W^+$) and $\mathbf{c}_i^-$ (for each $i \in W^-$) is needed regardless of the number N of encrypted bits. Therefore, C is made of $N + m \cdot |W^+| \cdot |W^-| \le N + 2mr$ values in $\mathbb{Z}_q$, compared with the $2Nmr$ generated by N different single-bit encryptions.

Clearly, the mediated scheme introduced in this article can benefit from the same approach, where mediators and user both receive a matrix $\mathbf{MK}_j$ and $\mathbf{E}$ of size $mr \times N$, built upon random matrices $\mathbf{U}_j \hookleftarrow_{\$} U(\mathbb{Z}_q^{n \times N})$ and upon $\mathbf{U}_0 = \mathbf{U} - \sum_{j=1}^{k} \mathbf{U}_j$, respectively.

Furthermore, security is ensured with the same claims as per the original manuscript.

8 Experiments with mR$_{LWE}$-CP-ABE

Here, we report the results of the performance experiments we carry out to evaluate the overhead introduced by the generation of the mediator keys and by their application in the decryption phase. In order to do that, we implement mR$_{LWE}$-CP-ABE on top of the Palisade-ABE implementation[3] of the study by Zhang and Zhang [2] and we compare the execution time of `KGen` vs `MKGen` and of `Dec` vs (`PDec` + `MDec`) vs `Enc`. In particular, we compare the time for (`PDec` + `MDec`) also against the time needed for `Enc` algorithm as the procedure `PDec` spends most of the time generating the error x_j (cf. Algorithm 6 l.1) to protect the mediator key (which corresponds to actually performing an LWE scheme). It is important to note that such an error generation does not depend on the inputs of `PDec` and can be also generated offline to save computational time.

We carry out the following performance experiments on an Nvidia DGX-1 equipped with 512 GB of memory; for each experiment, we report in Table 1 the average execution time over 20 repetitions for the three different security levels, namely, HEStd_128_classic, HEStd_192_classic, and HEStd_256_classic, and for five values for the number of attributes, for instance, 6, 8, 16, 20, and 34. We conduct experiments fixing the number of encrypted bits to 10,000 and setting $k = 1$.

As mentioned earlier, by seeing Table 1, it is clear that the decryption in our system (`PDec` + `MDec`) requires almost the time of the encryption. For what concern the `KGen` vs `MKGen`, the latter requires to double the time of the former; the behaviour is what we expect as `MKGen` needs to generate both the user key and the mediator keys and in the experiments setup, as mentioned, $k = 1$ so actually it generates two keys.

When k is increased, we experience a linear growth of both the `MKGen` and the `MDec` execution time. However, it is worthwhile noticing that, despite the total execution time for `MDec` clearly scales linearly with k, the time of a single `PDec` does not. As a consequence, if in `MDec` the queries to the SMs are implemented asynchronously (cf. Algorithm 7 ll.7–8), then the resulting computational time is the same regardless of the number of SMs.

3 Last access time 02-12-2022 at https://gitlab.com/palisade/palisade-abe.

Table 1: Average execution time (ms) of KGen, MKGen, Dec, PDec, MDec, and Enc algorithms. We run experiments by varying the security level and the values for attributes. We fix the number of SMs to $k = 1$ and the size of plaintext to 10,000 bits

Parameters		Time (ms)				
Security level	**#Attributes**	KGen	MKGen	Dec	PDec + MDec	Enc
HEStd_128_classic	6	179	354	1	54	47
HEStd_128_classic	8	223	451	2	75	73
HEStd_128_classic	16	386	769	3	141	129
HEStd_128_classic	20	465	938	4	185	174
HEStd_128_classic	32	725	1,463	7	309	276
HEStd_192_classic	6	108	212	0	23	21
HEStd_192_classic	8	125	248	1	32	31
HEStd_192_classic	16	195	387	1	59	56
HEStd_192_classic	20	227	454	2	80	73
HEStd_192_classic	32	336	657	3	117	112
HEStd_256_classic	6	341	681	3	103	99
HEStd_256_classic	8	435	864	4	149	147
HEStd_256_classic	16	771	1,550	8	282	272
HEStd_256_classic	20	929	1,869	9	381	363
HEStd_256_classic	32	1,470	2,947	16	687	574

9 Conclusions and further work

In this article, we have presented – to the best of our knowledge – the first scheme for revocable CP-ABE based on the LWE problem over lattices, hence being embedded in a post-quantum secure environment. The scheme takes advantage of the lattice-based CP-ABE scheme first presented in the study by Zhang and Zhang [2] by building upon it a server-aided fine-grained revoking system (mR_{LWE}-CP-ABE). The servers involved are considered semi-trusted; hence, the security proofs are given against different threat models. Security and applications are discussed both in the single-bit and in the multi-bit approach. For the sake of completeness, the here proposed scheme mR_{LWE}-CP-ABE is implemented on the ABE spin-off of the well-established open-source library Palisade to experimentally validate and provide some early performance estimation with particular attention to the overhead with respect to the original scheme. The implementation will be released as open source to let the community independently test and evaluate it.

In future implementations, we plan to develop a similar approach on two schemes similar to the one presented in Section 4.2: the first, proposed by Zhang et al. [34], introduced support to multi-valued attributes, while the second, introduced by Chen et al. [35], extends over [34] to support Ring-LWE. Furthermore, we also plan to provide support to other libraries, including e.g., Microsoft SEAL [36] and Pyfhel, as well as to carry out a more in-depth performance analysis of the system. On the side, future work might also be focused on investigating threshold systems for allowing more complex distribution of secrets amongst the SMs, hence providing more strength and robustness to the proposed system. Finally, and arguably most importantly, work is currently in progress to provide a formal proof of the post-quantum resistance of mR_{LWE}-CP-ABE, in particular to tackle the problems highlighted by Lombardi *et al.* in the recent literature [32].

Acknowledgements: E. Onofri acknowledges the Gruppo Nazionale per il Calcolo Scientifico (GNCS) of Istituto Nazionale di Alta Matematica (INdAM). M. Pedicini acknowledges the Gruppo Nazionale per le Strutture Algebriche, Geometriche e le loro Applicazioni (GNSAGA) of Istituto Nazionale di Alta Matematica (INdAM). This work has been accepted for presentation at CIFRIS23, the Congress of the Italian association of cryptography “De Componendis Cifris.”

Funding information: The authors state no funding involved.

Conflict of interest: The authors state no conflict of interest.

References

[1] Regev O. On lattices, learning with errors, random linear codes, and cryptography. In: STOC'05: Proceedings of the 37th Annual ACM Symposium on Theory of Computing. New York: ACM; 2005. p. 84–93. doi: https://doi.org/10.1145/1060590.1060603.

[2] Zhang J, Zhang Z. A ciphertext policy attribute-based encryption scheme without pairings. In: International Conference on Information Security and Cryptology. Springer; 2011. p. 324–40. doi: https://doi.org/10.1007/978-3-642-34704-7_23.

[3] Boneh D, Ding X, Tsudik G, Wong C. A method for fast revocation of public key certificates and security capabilities. In: Wallach DS, editor. 10th USENIX Security Symposium, August 13–17, 2001, Washington, D.C., USA. USENIX; 2001. http://www.usenix.org/publications/library/proceedings/sec01/boneh.html.

[4] PALISADE Lattice Cryptography Library (release 1.11.2); 2021. https://palisade-crypto.org/.

[5] Sahai A, Waters B. Fuzzy identity-based encryption. In: Cramer R, editor. Advances in Cryptology - EUROCRYPT 2005. Berlin, Heidelberg: Springer; 2005. p. 457–73. doi: https://doi.org/10.1007/11426639_27.

[6] Al-Dahhan RR, Shi Q, Lee GM, Kifayat K. Survey on revocation in Ciphertext-policy attribute-based encryption. Sensors (Basel). 2019 Apr;19(7):1695. doi: https://doi.org/10.3390/s19071695.

[7] Mascia C, Sala M, Villa I. A survey on functional encryption. Adv Math Commun. 2023;17(5):1251–89. doi: https://doi.org/10.3934/amc.2021049.

[8] Moffat S, Hammoudeh M, Hegarty R. A survey on ciphertext-policy attribute-based encryption (CP-ABE) approaches to data security on mobile devices and its application to IoT. In: Proceedings of the International Conference on Future Networks and Distributed Systems. ICFNDS '17. New York, NY, USA: Association for Computing Machinery; 2017. doi: https://doi.org/10.1145/3102304.3102338.

[9] Rasori M, Manna ML, Perazzo P, Dini G. A survey on attribute-based encryption schemes suitable for the Internet of things. IEEE Internet Things J. 2022 June;9(11):8269–90. doi: https://doi.org/10.1109/JIOT.2022.3154039.

[10] Zhang Y, Deng RH, Xu S, Sun J, Li Q, Zheng D. Attribute-based encryption for cloud computing access control: a survey. ACM Comput Surv. 2020 Aug;53(4):1–41. doi: https://doi.org/10.1145/3398036.

[11] Bethencourt J, Sahai A, Waters B. Ciphertext-policy attribute-based encryption. In: 2007 IEEE Symposium on Security and Privacy (SP '07); 2007. p. 321–34. doi: https://doi.org/10.1109/SP.2007.11.

[12] Goyal V, Pandey O, Sahai A, Waters B. Attribute-based encryption for fine-grained access control of encrypted data. In: Proceedings of the 13th ACM Conference on Computer and Communications Security. CCS '06. New York, NY, USA: Association for Computing Machinery; 2006. p. 89–98. doi: https://doi.org/10.1145/1180405.1180418.

[13] Xu S, Yang G, Mu Y. Revocable attribute-based encryption with decryption key exposure resistance and ciphertext delegation. Inform Sci. 2019;479:116–34. doi: https://doi.org/10.1016/j.ins.2018.11.031.

[14] Liu JK, Yuen TH, Zhang P, Liang K. Time-based direct revocable ciphertext-policy attribute-based encryption with short revocation list. In: Preneel B, Vercauteren F, editors. Applied Cryptography and Network Security. Cham: Springer International Publishing; 2018. p. 516–34. doi: https://doi.org/10.1007/978-3-319-93387-0_27.

[15] Phuong TVX, Yang G, Susilo W, Chen X. Attribute based broadcast encryption with short ciphertext and decryption key. In: Pernul G, Y A Ryan P, Weippl E, editors. Computer Security - ESORICS 2015. Cham: Springer International Publishing; 2015. p. 252–69. doi: https://doi.org/10.1007/978-3-319-24177-7_13.

[16] Sahai A, Seyalioglu H, Waters B. Dynamic credentials and ciphertext delegation for attribute-based encryption. In: Safavi-Naini R, Canetti R, editors. Advances in Cryptology - CRYPTO 2012. Berlin, Heidelberg: Springer; 2012. p. 199–217. doi: https://doi.org/10.1007/978-3-642-32009-5.

[17] Yu S, Wang C, Ren K, Lou W. Attribute based data sharing with attribute revocation. In: Proceedings of the 5th ACM Symposium on Information, Computer and Communications Security. ASIACCS '10. New York, NY, USA: Association for Computing Machinery; 2010. p. 261–70. doi: https://doi.org/10.1145/1755688.1755720.

[18] Xie X, Ma H, Li J, Chen X. An efficient ciphertext-policy attribute-based access control towards revocation in cloud computing. J Universal Comput Sci. 2013;19(16):2349–67. doi: https://doi.org/10.3217/jucs-019-16-2349.

[19] Yang Y, Ding X, Lu H, Wan Z, Zhou J. Achieving revocable fine-grained cryptographic access control over cloud data. In: Desmedt Y, editor. Information security. Cham: Springer International Publishing; 2015. p. 293–308. doi: https://doi.org/10.1007/978-3-319-27659-5_21.

[20] Cui H, Deng RH, Ding X, Li Y. Attribute-based encryption with granular revocation. In: Deng R, Weng J, Ren K, Yegneswaran V, editors. Security and Privacy in Communication Networks. Cham: Springer International Publishing; 2017. p. 165–81. doi: https://doi.org/10.1007/978-3-319-59608-2_9.

[21] Blömer J, Seifert JP. On the complexity of computing short linearly independent vectors and short bases in a lattice. In: Proceedings of the Thirty-First Annual ACM Symposium on Theory of Computing. STOC '99. New York, NY, USA: Association for Computing Machinery; 1999. p. 711–20. doi: https://doi.org/10.1145/301250.301441.

[22] Aggarwal D, Chung E. A note on the concrete hardness of the shortest independent vector in lattices. Inform Process Lett. 2021;167:106065. doi: https://doi.org/10.1016/j.ipl.2020.106065.
[23] Bennett H, Golovnev A, Stephens-Davidowitz N. On the quantitative hardness of CVP. In: 2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS); 2017. p. 13–24. doi: https://doi.org/10.1109/FOCS.2017.11.
[24] Impagliazzo R, Paturi R. On the Complexity of k-SAT. J Comput Syst Sci. 2001;62(2):367–75. doi: https://doi.org/10.1006/jcss.2000.1727.
[25] Micciancio D, Regev O. Worst-case to average-case reductions based on Gaussian measures. SIAM J Comput. 2007;37(1):267–302. doi: https://doi.org/10.1137/S0097539705447360.
[26] Regev O. On lattices, learning with errors, random linear codes, and cryptography. J ACM (JACM). 2009;56(6):1–40. doi: https://doi.org/10.1145/1568318.1568324.
[27] Peikert C. Some recent progress in lattice-based cryptography. In: Theory of Cryptography. Berlin Heidelberg: Springer; 2009. p. 72–2. doi: https://doi.org/10.1007/978-3-642-00457-5_5.
[28] Agrawal S, Boneh D, Boyen X.. Efficient Lattice (H) IBE in the standard model. Eurocrypt'10 and PKC'10 joint work.2010. http://boneh.com/pubs/papers/latticebb.pdf.
[29] Gentry C, Peikert C, Vaikuntanathan V. Trapdoors for hard lattices and new cryptographic constructions. In: Proceedings of the fortieth annual ACM symposium on Theory of computing; 2008. p. 197–206. doi: https://doi.org/10.1145/1374376.1374407.
[30] Alwen J, Peikert C. Generating shorter bases for hard random lattices. In: Albers S, Marion JY, editors. 26th International Symposium on Theoretical Aspects of Computer Science STACS 2009. Proceedings of the 26th Annual Symposium on the Theoretical Aspects of Computer Science. Freiburg, Germany: IBFI Schloss Dagstuhl; 2009. p. 75–86. https://hal.inria.fr/inria-00359718.
[31] Cash D, Hofheinz D, Kiltz E. How to delegate a Lattice basis; 2009. Cryptology ePrint Archive, Paper 2009/351. https://eprint.iacr.org/2009/351.
[32] Lombardi A, Mook E, Quach W, Wichs D. Post-quantum insecurity from LWE. In: Theory of Cryptography. Springer Nature Switzerland; 2022. p. 3–32. doi: https://doi.org/10.1007/978-3-031-22318-1_1.
[33] Shamir A. How to share a secret. Commun ACM. 1979 Nov;22(11):612–3. doi: https://doi.org/10.1145/359168.359176.
[34] Zhang J, Zhang Z, Ge A. Ciphertext policy attribute-based encryption from lattices. In: Proceedings of the 7th ACM Symposium on Information, Computer and Communications Security; 2012. p. 16–7. doi: https://doi.org/10.1145/2414456.2414464.
[35] Chen Z, Zhang P, Zhang F, Huang J. Ciphertext policy attribute-based encryption supporting unbounded attribute space from R-LWE. KSII Trans Internet Inform Syst (TIIS). 2017;11(4):2292–309.
[36] Microsoft SEAL (release 4.0); 2022. Microsoft Research, Redmond, WA. https://github.com/Microsoft/SEAL.