

Multi-Authority ABE from Lattices without Random Oracles

Brent Waters^{1,2}, Hoeteck Wee^{2,3}, and David J. Wu¹

¹ University of Texas at Austin, Austin, TX, USA

² NTT Research, Sunnyvale, CA, USA

³ CNRS, ENS, PSL, Paris, France

Abstract. Attribute-based encryption (ABE) extends public-key encryption to enable fine-grained control to encrypted data. However, this comes at the cost of needing a central *trusted* authority to issue decryption keys. A multi-authority ABE (MA-ABE) scheme decentralizes ABE and allows anyone to serve as an authority. Existing constructions of MA-ABE only achieve security in the random oracle model.

In this work, we develop new techniques for constructing MA-ABE for the class of subset policies (which captures policies such as conjunctions and DNF formulas) whose security can be based in the plain model *without* random oracles. We achieve this by relying on the recently-proposed “evasive” learning with errors (LWE) assumption by Wee (EUROCRYPT 2022) and Tsabury (CRYPTO 2022).

Along the way, we also provide a modular view of the MA-ABE scheme for DNF formulas by Datta et al. (EUROCRYPT 2021) in the random oracle model. We formalize this via a general version of a related-trapdoor LWE assumption by Brakerski and Vaikuntanathan (ITCS 2022), which can in turn be reduced to the plain LWE assumption. As a corollary, we also obtain an MA-ABE scheme for subset policies from *plain* LWE with a *polynomial* modulus-to-noise ratio in the random oracle model. This improves upon the Datta et al. construction which relied on LWE with a *sub-exponential* modulus-to-noise ratio. Moreover, we are optimistic that the generalized related-trapdoor LWE assumption will also be useful for analyzing the security of other lattice-based constructions.

1 Introduction

Attribute-based encryption (ABE) [SW05, GPSW06] extends classic public-key encryption to support fine-grained access control on encrypted data. For instance, in a ciphertext-policy ABE (CP-ABE) scheme, each ciphertext ct is associated with a policy f together with a message μ while decryption keys sk are associated with an attribute x . Decryption successfully recovers the message μ when x satisfies f . Security requires that an adversary who only possesses secret keys for a collection of attributes x_1, \dots, x_n that do not satisfy f does not learn anything about the message. In this work, we are interested in systems that are secure against *unbounded* collusions: that is, security holds against an adversary that has any arbitrary (polynomial) number of non-satisfying attributes.

Multi-authority ABE. In a traditional ABE scheme, there exists a central trusted authority that generates and issues decryption keys. The central authority has the ability to decrypt *all* ciphertexts encrypted using the system. To mitigate the reliance on a single central trusted authority, a line of works [Cha07, LCLS08, MKE08, CC09] have introduced and studied the notion of a “multi-authority” ABE (MA-ABE) scheme where *anyone* can become an authority. In an MA-ABE scheme, each authority controls different attributes and can *independently* issue secret keys corresponding to the set of attributes under their control. Policies in an MA-ABE system are formulated with respect to the attributes of one or more authorities. To decrypt, a user combines the secret keys for attributes from a set of authorities that satisfy the policy. Security is still required to hold against users who possess an arbitrary number of unauthorized secret keys, with an additional challenge that some subset of the authorities (associated with the ciphertext policy) could now be corrupted and colluding with the adversary.

Earlier constructions of MA-ABE had various limitations in terms of functionality or security (or both). The first construction that achieved the first fully decentralized MA-ABE scheme was by Lewko and Waters [LW11]. Unlike previous schemes, the Lewko-Waters scheme allows any user to become an authority, and moreover, the only coordination needed among users and authorities is a one-time sampling of a set of global parameters. The Lewko-Waters construction supports any access policy computable by an NC^1 circuit (i.e., a Boolean formula) and security relies on assumptions on groups with bilinear maps and in the random oracle model. Subsequently, a number of works have realized new constructions for NC^1 policies based on bilinear maps [RW15, DKW21b], and recently, Datta et al. [DKW21a] showed how to construct an MA-ABE scheme for access policies computable by DNF formulas (of *a priori* bounded size) from the learning with errors (LWE) assumption [DKW21a]. All of these constructions rely on the random oracle model. This motivates the following question:

Can we construct a multi-authority ABE scheme without random oracles?

1.1 Our Contributions

In this work, we show how to leverage the recently-introduced evasive LWE assumption [Wee22, Tsa22] to obtain an MA-ABE scheme for subset policies *without* random oracles. Subset policies capture DNF policies as in [DKW21a].⁴ Moreover, our MA-ABE construction supports subset policies and DNFs of arbitrary polynomial size which improves upon the previous lattice-based construction in the random oracle model [DKW21a]. We summarize this result in the following informal theorem and provide the full details in Section 6:

⁴As noted in [DKW21a, Remark 6.1], the MA-ABE scheme therein requires a *monotone* secret-sharing scheme where reconstruction has small coefficients and the joint distribution of the unauthorized shares are uniformly random; such a scheme is only known for subset policies and DNFs.

Theorem 1.1 (Informal). *Assuming polynomial hardness of LWE and of evasive LWE (both with a sub-exponential modulus-to-noise ratio), there exists a statically-secure multi-authority ABE for subset policies (of arbitrary polynomial size).*

Understanding the evasive LWE assumption. While the evasive LWE assumption is much less well-understood compared to the plain LWE assumption, our construction provides a new avenue towards realizing MA-ABE *without* random oracles. In particular, putting assumptions aside, our construction constitutes the first heuristic MA-ABE without random oracles. In all previous constructions of multi-authority ABE, the random oracle was used to hash a global user identifier (denoted gid) to obtain common randomness that is used to bind different keys to a single user. For the particular case of [DKW21a], the random oracle was used to hash an identifier to obtain a discrete Gaussian sample. Our candidate replaces the random oracle with a subset product of public low-norm matrices. To prove security of the resulting scheme, we rely on the fact that under LWE, multiplying a secret key by a subset product of (public) low-norm matrices yields a pseudorandom function [BLMR13] in addition to the evasive LWE assumption.

A modular approach in the random oracle model. The starting point of our construction is the MA-ABE construction for (bounded-size) DNF policies by Datta et al. [DKW21a]. Along the way to our construction without random oracles (Theorem 1.1), we provide a more modular description of the Datta et al. scheme. Specifically, we extract a new trapdoor sampling lemma that is implicitly used in their construction. This lemma can be viewed as a generalization of the related trapdoor LWE lemma from the recent work of Brakerski and Vaikuntanathan [BV22], and may prove useful for constructing other primitives from the *standard* LWE assumption. We provide an overview of our generalized related-trapdoor lemma in Section 2 and provide the full details in Section 4.

Using our generalized related-trapdoor LWE lemma, we in turn provide a more modular description of the MA-ABE scheme of Datta et al. [DKW21a], and moreover, base hardness on the plain LWE assumption with a *polynomial* modulus-to-noise ratio in the random oracle model. Previously, Datta et al. relied on noise smudging for trapdoor sampling in their security analysis⁵, and consequently, could only reduce security to LWE with a *sub-exponential* modulus-to-noise ratio. We summarize these results in the following (informal) theorem and provide the full details in Section 5:

Theorem 1.2 (Informal). *Let λ be a security parameter. Assuming polynomial hardness of LWE with a polynomial modulus-to-noise ratio, there exists a statically-secure multi-authority ABE scheme for subset policies of a priori bounded length $L = L(\lambda)$ in the random oracle model. The size of the ciphertext is quasi-linear in the bound L .*

⁵See the descriptions of Hybrid 5 and the analysis of Lemmas 5.5 and 6.5 in [DKW21a], where noise smudging is used for simulating secret keys.

Like previous lattice-based MA-ABE constructions in the random oracle model [DKW21a], the global public parameters in Theorem 1.2 imposes an *a priori* bound L on the size of the policies that can be associated with ciphertexts, and moreover, the ciphertext size increases as a function of L . We note that our construction based on the stronger evasive LWE assumption (Theorem 1.1) supports policies of arbitrary polynomial size in the plain model.

1.2 Additional Related Work

Kim [Kim19] and Wang et al. [WFL19] also studied constructions of multi-authority ABE (for bounded-depth circuits and Boolean formulas, respectively) from lattice-based assumptions. However, both schemes operate in a model where there is a single central authority that generates the public keys and secret keys for each of the authorities in the system. Relying on a central trusted party runs against the original goal of *decentralizing* trust. Moreover, these constructions only ensure security against bounded collusions. In this work, we focus exclusively on the fully decentralized setting introduced by Lewko and Waters [LW11] that neither requires a centralized setup nor assumes an *a priori* bound on the number of authorities or corruptions.

Recently, Tsabury [Tsa22] and Vaikuntanathan et al. [VWW22] showed how to build witness encryption from a *stronger* variant of the evasive LWE assumption with private-coin auxiliary input and sub-exponential hardness. In contrast, our multi-authority ABE construction in the standard model relies on evasive LWE with *public-coin* auxiliary input and polynomial hardness with a sub-exponential modulus-to-noise ratio; this was also the case for the optimal broadcast encryption scheme by Wee [Wee22]. While vanilla witness encryption implies single-authority ABE [GGSW13], we currently do not know how to construct multi-authority ABE from *vanilla* witness encryption.

2 Technical Overview

In this section, we provide a technical overview of our lattice-based MA-ABE constructions. Throughout this work, we focus exclusively on *subset policies* (which suffices for supporting DNF formulas). In an ABE scheme for subset policies, ciphertexts are associated with a set A and secret keys are associated with a set B . Decryption succeeds if $A \subseteq B$.

Lattice preliminaries. The learning with errors (LWE) assumption [Reg05] says that the distribution $(\mathbf{A}, \mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top)$ is computationally indistinguishable from $(\mathbf{A}, \mathbf{u}^\top)$ where $\mathbf{A} \xleftarrow{\mathcal{R}} \mathbb{Z}_q^{n \times m}$, $\mathbf{s} \xleftarrow{\mathcal{R}} \mathbb{Z}_q^n$, $\mathbf{e} \leftarrow D_{\mathbb{Z}, \chi}^m$, and $\mathbf{u} \xleftarrow{\mathcal{R}} \mathbb{Z}_q^m$, where n, m, q, χ are lattice parameters and $D_{\mathbb{Z}, \chi}$ is the discrete Gaussian distribution with parameter χ . To simplify the presentation in the technical overview, we will use *curly underlines* in place of (small) noise terms. Namely, instead of writing $\mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top$, we simply write $\underline{\mathbf{s}^\top \mathbf{A}}$.

For a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and a target vector $\mathbf{y} \in \mathbb{Z}_q^n$, we write $\mathbf{A}_\chi^{-1}(\mathbf{y})$ to denote a random variable $\mathbf{x} \in \mathbb{Z}_q^m$ whose distribution is a discrete Gaussian distribution $D_{\mathbb{Z}, \chi}^m$ conditioned on $\mathbf{A}\mathbf{x} = \mathbf{y}$. For ease of notation, we will drop the subscript χ in this technical overview. A sequence of works [Ajt96, GPV08, ABB10b, ABB10a, CHKP10, MP12] (see also Section 3.2) have shown how to sample a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ together with a trapdoor $\text{td}_{\mathbf{A}}$ to enable efficient sampling from the distribution $\mathbf{A}^{-1}(\mathbf{y})$ for any target $\mathbf{y} \in \mathbb{Z}_q^n$.

In the following description, we write $\mathbf{I}_n \in \mathbb{Z}_q^{n \times n}$ to denote the n -by- n identity matrix and $\mathbf{G} = \mathbf{I}_n \otimes \mathbf{g}^\top \in \mathbb{Z}_q^{n \times m}$, where $\mathbf{g}^\top = [1 \mid 2 \mid \dots \mid 2^{\lfloor \log q \rfloor}]$, to denote the standard gadget matrix [MP12].

2.1 Starting Point: Single-Authority CP-ABE for Subset Policies

We start by describing a simple CP-ABE for subset policies that lies at the core of our MA-ABE scheme. In the following, let $[L]$ be the universe of attributes. Each ciphertext is associated with a subset $A \subseteq [L]$ and each secret key is associated with a subset $B \subseteq [L]$; decryption succeeds as long as $A \subseteq B$.

- The master public key consists of $(\mathbf{A}_1, \mathbf{B}_1, \mathbf{p}_1), \dots, (\mathbf{A}_L, \mathbf{B}_L, \mathbf{p}_L) \xleftarrow{\mathbb{R}} \mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^{n \times m(2L-1)} \times \mathbb{Z}_q^n$.
- The master secret key consist of the trapdoors $\text{td}_{\mathbf{A}_1}, \dots, \text{td}_{\mathbf{A}_L}$ for $\mathbf{A}_1, \dots, \mathbf{A}_L$, respectively.
- An encryption of a message bit $\mu \in \{0, 1\}$ with respect to a set $X \subseteq [L]$ is a tuple

$$\text{ct} = \left(\{\underbrace{\mathbf{s}^\top \mathbf{A}_i}_{i \in X}, \underbrace{\mathbf{s}^\top \sum_{i \in X} \mathbf{B}_i}, \underbrace{\mathbf{s}^\top \sum_{i \in X} \mathbf{p}_i + \mu \cdot \lfloor q/2 \rfloor} \right),$$

where $\mathbf{s} \xleftarrow{\mathbb{R}} \mathbb{Z}_q^n$.

- A secret key for a set $Y \subseteq [L]$ consists of a tuple

$$\text{sk} = \left(\{\mathbf{A}_i^{-1}(\mathbf{p}_i + \mathbf{B}_i \mathbf{r})\}_{i \in Y}, \mathbf{r} \right),$$

where $\mathbf{r} \leftarrow D_{\mathbb{Z}, \chi}^{m(2L-1)}$ is sampled from a discrete Gaussian distribution.

Decryption uses the fact that

$$-\left(\underbrace{\mathbf{s}^\top \sum_{i \in X} \mathbf{B}_i}_{i \in X}\right) \cdot \mathbf{r} + \sum_{i \in X} \underbrace{\mathbf{s}^\top \mathbf{A}_i}_{i \in X} \cdot \mathbf{A}_i^{-1}(\mathbf{p}_i + \mathbf{B}_i \mathbf{r}) \approx -\mathbf{s}^\top \sum_{i \in X} \mathbf{B}_i \mathbf{r} + \mathbf{s}^\top \sum_{i \in X} (\mathbf{p}_i + \mathbf{B}_i \mathbf{r}) = \mathbf{s}^\top \sum_{i \in X} \mathbf{p}_i,$$

since \mathbf{r} and $\mathbf{A}^{-1}(\cdot)$ are small. Looking ahead to our multi-authority construction, observe that key generation can be carried out in a decentralized manner: given a “public” Gaussian vector \mathbf{r} , computing the secret-key components $\mathbf{A}_i^{-1}(\mathbf{p}_i + \mathbf{B}_i \mathbf{r})$ associated with index i only requires knowledge of $\mathbf{B}_i, \mathbf{p}_i$ and the trapdoor for \mathbf{A}_i , which are all *specific* to attribute i (and could be independently generated by the i^{th} authority).

Selective security. To argue that this CP-ABE scheme is selectively secure⁶, we proceed as follows:

1. First, we show how to sample a secret key for a set $Y \subseteq [L]$ given a trapdoor for \mathbf{B}_Y , where $\mathbf{B}_Y \in \mathbb{Z}_q^{n|Y| \times m(2L-1)}$ is the matrix formed by vertically concatenating \mathbf{B}_i for all $i \in Y$.
2. Next, we show that under the LWE assumption, $\mathbf{s}^\top \sum_{i \in X} \mathbf{B}_i$ is pseudorandom even given an oracle for $\mathbf{B}_Y^{-1}(\cdot)$ for arbitrary $Y \subseteq [L]$ of the adversary's choosing, provided that for each Y , it is the case that $X \not\subseteq Y$. Here, $X \subseteq [L]$ is the set associated with the challenge ciphertext. Technically, we additionally require that $\mathbf{s}^\top \mathbf{A}_i$ and $\mathbf{s}^\top \sum_{i \in X} \mathbf{p}_i$ are also pseudorandom, but these components are easily handled by the standard LWE assumption. For ease of exposition, we do not focus on these additional components in this overview and refer instead to [Sections 4 and 5](#) for the full description.

For the second step, we prove a more general statement which generalizes the related-trapdoor LWE lemma previously introduced by Brakerski and Vaikuntanathan [\[BV22\]](#) in the context of constructing compact CP-ABE for circuits.

Generalized related-trapdoor LWE. Our generalized related-trapdoor LWE assumption asserts that for any non-zero vector $\mathbf{u} \in \{0, 1\}^L$, the vector $\mathbf{s}^\top (\mathbf{u}^\top \otimes \mathbf{I}_n) \mathbf{B}$ is pseudorandom given an oracle for the function $(\mathbf{M}, \mathbf{t}) \mapsto ((\mathbf{M} \otimes \mathbf{I}_n) \mathbf{B})^{-1}(\mathbf{t})$, as long as the matrix $\bar{\mathbf{M}} = \begin{bmatrix} \mathbf{M} \\ \mathbf{u}^\top \end{bmatrix} \in \mathbb{Z}_q^{(k+1) \times L}$ is full rank (and $k < L$).⁷ To show that the standard LWE assumption implies the generalized related-trapdoor LWE assumption, we take an LWE matrix $\hat{\mathbf{A}}$ and the vector $\mathbf{u} \in \{0, 1\}^L$, and we set the matrix \mathbf{B} to be

$$\mathbf{B} = \begin{bmatrix} \hat{\mathbf{A}} & \hat{\mathbf{A}}\mathbf{R} + \mathbf{U}^\perp \otimes \mathbf{G} \end{bmatrix}$$

where \mathbf{R} is a (random) low-norm matrix and $\mathbf{U}^\perp \in \{0, 1\}^{L \times (L-1)}$ is a full-rank basis for the kernel of \mathbf{u}^\top . By design, $(\mathbf{u} \otimes \mathbf{I}_n) \mathbf{B} = [(\mathbf{u} \otimes \mathbf{I}_n) \hat{\mathbf{A}} \mid (\mathbf{u} \otimes \mathbf{I}_n) \hat{\mathbf{A}} \mathbf{R}]$ which means we do *not* know a trapdoor for $(\mathbf{u} \otimes \mathbf{I}_n) \mathbf{B}$. On the other hand,

$$(\mathbf{M} \otimes \mathbf{I}_n) \mathbf{B} \underbrace{\begin{bmatrix} -\mathbf{R} \\ \mathbf{I}_{m(L-1)} \end{bmatrix}}_{\tilde{\mathbf{R}}} = (\mathbf{M} \otimes \mathbf{I}_n) (\mathbf{U}^\perp \otimes \mathbf{G}) = \mathbf{M} \mathbf{U}^\perp \otimes \mathbf{G}.$$

When $\bar{\mathbf{M}} = \begin{bmatrix} \mathbf{M} \\ \mathbf{u}^\top \end{bmatrix}$ is full rank, then $\mathbf{M} \mathbf{U}^\perp$ is also full rank. Since $\tilde{\mathbf{R}}$ is low-norm, it is a trapdoor for $(\mathbf{M} \otimes \mathbf{I}_n) \mathbf{B}$ (see [\[MP12\]](#) and [Corollary 3.12](#)).

⁶In the selective security game, the adversary starts by committing to the set X associated with the challenge ciphertext. The reduction algorithm is then allowed to program X into the public parameters of the scheme.

⁷Some restriction on \mathbf{M} is also necessary. For instance, it is easy to distinguish $\mathbf{s}^\top (\mathbf{u} \otimes \mathbf{I}_n) \mathbf{B}$ if $\mathbf{M} = \mathbf{u}^\top$, or more generally, if $\mathbf{u}_0^\top \mathbf{M} = \mathbf{u}$ for some $\mathbf{u}_0 \in \{0, 1\}^k$.

Returning to the proof of selective security for the above CP-ABE scheme, observe that showing $\mathbf{s}^\top \sum_{i \in X} \mathbf{B}_i$ given an oracle for $\mathbf{B}_Y^{-1}(\cdot)$ directly maps to an instance of the related-trapdoor LWE assumption:

- Let $\mathbf{B} \in \mathbb{Z}_q^{nL \times m(2L-1)}$ be the matrix obtained by vertically stacking $\mathbf{B}_1, \dots, \mathbf{B}_L \in \mathbb{Z}_q^{n \times m(2L-1)}$.
- The vector $\mathbf{u} \in \{0, 1\}^L$ is the indicator vector for the challenge set X . Namely, $u_i = 1$ if $i \in X$ and 0 otherwise. Then, $(\mathbf{u} \otimes \mathbf{I}_n)\mathbf{B} = \sum_{i \in X} \mathbf{B}_i$.
- The oracle $\mathbf{B}_Y^{-1}(\cdot)$ can be simulated by querying the related-trapdoor oracle on matrix $\mathbf{M}_Y \in \mathbb{Z}_q^{|Y| \times L}$ formed by taking the rows of \mathbf{I}_L corresponding to the indices in Y . In this case $(\mathbf{M} \otimes \mathbf{I}_n)\mathbf{B} = \mathbf{B}_Y$ defined previously. Moreover, by construction of \mathbf{M}_Y , whenever $X \not\subseteq Y$, we have that \mathbf{u}^\top is *not* in the row-span of \mathbf{M}_Y .

Finally, we remark here that the original version of the related-trapdoor LWE assumption formulated by Brakerski and Vaikuntanathan [BV22] considered the special case where the matrix \mathbf{M} is a row vector with a specific structure.⁸ Our formulation considers a general matrix \mathbf{M} which is useful for constructing an ABE scheme with a *distributed* setup. We also note that this type of trapdoor sampling was also implicit in the CP-ABE construction of Datta et al. [DKW21a]; however, they critically relied on noise flooding to simulate the analog of the $((\mathbf{M} \otimes \mathbf{I}_n)\mathbf{B})^{-1}(\cdot)$ oracle. As a result, the security of their scheme relied on LWE with a super-polynomial modulus-to-noise ratio in the random oracle model. In this work, we both provide a modular description of the core trapdoor sampling lemma (Section 4) and then show how to leverage it to obtain a multi-authority ABE for subset policies using LWE with a polynomial modulus-to-noise ratio in the random oracle model (Section 5). We are optimistic that our generalized version of the related trapdoor LWE assumption will also be useful for analyzing the security of other lattice-based constructions.

2.2 MA-ABE for Subset Policies in the Random Oracle Model

First, we observe that our core CP-ABE scheme naturally extends to yield a MA-ABE scheme for subset policies in the random oracle model. We make the following modifications to the base scheme:

- The authority associated with attribute i samples $\mathbf{A}_i, \mathbf{B}_i, \mathbf{p}_i$ along with a trapdoor $\text{td}_{\mathbf{A}_i}$ for \mathbf{A}_i .
- To generate a key for a user with identifier gid , we derive \mathbf{r} *deterministically* from $H(\text{gid})$ and output $\mathbf{A}^{-1}(\mathbf{p}_i + \mathbf{B}_i \mathbf{r})$.

Security of the core CP-ABE implies that the ensuing MA-ABE scheme remains secure as long as no authority is corrupted. On the other hand, it is easy to see that the scheme is insecure if we allow authority corruptions, since we can use an authority's trapdoor to recover the LWE secret \mathbf{s} from $\mathbf{s}^\top \mathbf{A}_i$.

⁸Concretely, $\mathbf{u}^\top = [1 \mid \mathbf{x}^\top]$ and $\mathbf{M} = [1 \mid \mathbf{y}^\top]$ for some $\mathbf{x}, \mathbf{y} \in \{0, 1\}^{L-1}$. The adversary is restricted to queries $\mathbf{y} \neq \mathbf{x}$, which is implied by our requirement that \mathbf{M} has full rank.

Security with authority corruptions. To defend against corrupted authorities, we modify the ciphertext structure. Instead of having a single LWE secret \mathbf{s} that is shared across *authorities*, we instead sample a fresh \mathbf{s}_i for each attribute $i \in X$. That is, the ciphertext is now given by:

$$\text{ct} = \left(\underbrace{\{\mathbf{s}_i^\top \mathbf{A}_i\}}_{i \in X}, \underbrace{\sum_{i \in X} \mathbf{s}_i^\top \mathbf{B}_i}, \underbrace{\sum_{i \in X} \mathbf{s}_i^\top \mathbf{p}_i + \mu \cdot \lfloor q/2 \rfloor} \right)$$

Key generation proceeds as before. Decryption still follows from a similar relation as before:

$$-\left(\underbrace{\sum_{i \in X} \mathbf{s}_i^\top \mathbf{B}_i}\right) \cdot \mathbf{r} + \sum_{i \in X} \underbrace{\mathbf{s}_i^\top \mathbf{A}_i} \cdot \mathbf{A}_i^{-1}(\mathbf{p}_i + \mathbf{B}_i \mathbf{r}) = \sum_{i \in X} \mathbf{s}_i^\top \mathbf{p}_i.$$

Static security with authority corruptions. We now argue that the resulting MA-ABE scheme is statically secure.⁹ Let \mathcal{C} denote the set of authorities that are corrupted. The adversary gets to choose the public keys and secret keys for authorities in \mathcal{C} . In the multi-authority setting, a secret-key query consists of a pair (Y, gid) where Y is a set of *honest* authorities (i.e., $Y \cap \mathcal{C} = \emptyset$) and gid is the user identifier. Let X be the set of authorities associated with the challenge ciphertext. The admissibility criterion is that $X \not\subseteq Y \cup \mathcal{C}$.

The proof of security proceeds similarly to that of our core CP-ABE, except we replace the challenge set X with the set $X \setminus \mathcal{C}$. Since $Y \cap \mathcal{C} = \emptyset$, the MA-ABE admissibility criterion $X \not\subseteq Y \cup \mathcal{C}$ is equivalent to $X \setminus \mathcal{C} \not\subseteq Y$, which coincides with the criterion from our CP-ABE analysis. In particular, the security reduction can basically *ignore* the ciphertext components associated with corrupted authorities (since the ciphertext component of each authority is associated with *independent* LWE secrets \mathbf{s}_i) and just focus on the attributes controlled by the honest authorities. The general argument again relies on our (generalized) related-trapdoor LWE assumption:

1. First, we show how to sample a secret key for Y given a trapdoor for \mathbf{B}_Y (where $\mathbf{B}_Y \in \mathbb{Z}_q^{n|Y| \times m(2L-1)}$ is again the matrix formed by vertically stacking the matrices \mathbf{B}_i associated with the authorities $i \in Y$).
2. As in the analysis of the CP-ABE scheme, we use the oracle in the related-trapdoor LWE assumption to compute $\mathbf{B}_Y^{-1}(\cdot)$ in the proof. Arguing the correctness of this step additionally requires the ability to “program” the random oracle. This is because in the real scheme, the secret keys are sampled by computing $\mathbf{r} \leftarrow \mathbf{H}(\text{gid})$ and then sampling $\mathbf{u}_i \leftarrow \mathbf{A}_i(\mathbf{p}_i + \mathbf{B}_i \mathbf{r})$ for each $i \in X$. The reduction algorithm will instead sample $\mathbf{u}_i \leftarrow D_{\mathbb{Z}, X}^m$ itself and then

⁹In the static security model [RW15], we require the adversary to commit to the set of corrupted authorities, the secret-key queries, and the challenge ciphertext query at the beginning of the security game. Previous lattice-based MA-ABE constructions were also analyzed in the static security model [DKW21a].

- obtain $\mathbf{r} \in \mathbb{Z}_q^{m(2L-1)}$ using its oracle $\mathbf{B}_Y^{-1}(\cdot)$. In the random oracle model, the reduction then programs $\mathbf{H}(\text{gid})$ to \mathbf{r} . We refer to [Section 5](#) for more details.
3. Finally, to simulate the challenge ciphertext, the reduction algorithm samples a random $\mathbf{s}_i \xleftarrow{\mathbb{R}} \mathbb{Z}_q^n$ for each corrupted authority $i \in \mathcal{C}$. For the honest authorities $i \in X \setminus \mathcal{C}$, the reduction sets the secret key to be $\hat{\mathbf{s}}_i$ and programs $\mathbf{s}_i := \mathbf{s} + \hat{\mathbf{s}}_i$, where \mathbf{s} is the secret in the related-trapdoor

We provide the formal analysis in [Section 5](#). This construction yields a MA-ABE scheme for subset policies from the related-trapdoor LWE assumption in the random oracle model. The related-trapdoor LWE assumption we rely on here reduces to the standard LWE assumption with a polynomial modulus-to-noise ratio. This yields [Theorem 1.2](#).

2.3 Removing Random Oracles via Evasive LWE

To obtain an MA-ABE construction *without* random oracles, we describe a way to *concretely* implement the hash function \mathbf{H} in our basic construction above. Our specific instantiation relies on computing a subset product of low-norm matrices. Specifically, let $\mathbf{D}_0, \mathbf{D}_1 \in \mathbb{Z}_q^{m \times m}$ be low-norm matrices. These are fixed public matrices that will be included as part of the global parameters. For an input $x \in \{0, 1\}^\ell$, we define $\mathbf{H}(x) := (\prod_{i \in [\ell]} \mathbf{D}_{x_i}) \boldsymbol{\eta} \in \mathbb{Z}_q^m$, where $\boldsymbol{\eta} \in \mathbb{Z}_q^m$ is the first canonical basis vector. Previously, Boneh et al. [\[BLMR13\]](#) showed that for any sequence of $x_1, \dots, x_k \in \{0, 1\}^\ell$ the values $\{\mathbf{s}^\top \mathbf{H}(x_i)\}_{i \in [k]}$ are pseudorandom.

While we do not know how to prove security of the MA-ABE construction instantiated with this subset-product hash function using the plain learning with errors assumption, we show how to do so using the recently-introduced evasive LWE assumption by Wee [\[Wee22\]](#) and Tsabury [\[Tsa22\]](#).

Evasive LWE. We start by describing a variant of the evasive LWE assumption introduced by Wee [\[Wee22\]](#) and refer to [Section 3.2](#) for the formal description. Let $\mathbf{P}_1, \dots, \mathbf{P}_\ell$ be drawn from some efficiently-sampleable distribution of matrices. Roughly speaking, the evasive LWE assumption says that if the distribution $\{\mathbf{A}_i, \mathbf{s}^\top \mathbf{P}_i\}_{i \in [\ell]}$ is pseudorandom, then the distributions

$$\{\mathbf{A}_i, \mathbf{s}^\top \mathbf{A}_i, \mathbf{A}_i^{-1}(\mathbf{P}_i)\}_{i \in [\ell]} \quad \text{and} \quad \{\mathbf{A}_i, \mathbf{u}_i^\top, \mathbf{A}_i^{-1}(\mathbf{P}_i)\}_{i \in [\ell]}$$

are computationally indistinguishable. Intuitively, the evasive LWE assumption says that the presence of $\mathbf{A}_i^{-1}(\mathbf{P}_i)$ does *not* help break LWE so long as $\mathbf{s}^\top \mathbf{P}_i$ is pseudorandom. Indeed, if the distinguisher multiplied $\mathbf{s}^\top \mathbf{A}$ with $\mathbf{A}^{-1}(\mathbf{P})$, then it roughly obtains $\mathbf{s}^\top \mathbf{P}$, which is pseudorandom by assumption.

In the context of our MA-ABE scheme, the matrices $\mathbf{A}_1, \dots, \mathbf{A}_\ell$ will be associated with the public keys for the honest authorities, and the columns of \mathbf{P}_i will consist of $\mathbf{p}_i + \mathbf{B}_i \mathbf{r}_{\text{gid}}$ for the user identifiers gid that appear in the adversary's secret-key queries. By setting \mathbf{P}_i properly (see [Section 6](#)), the reduction algorithm

can in turn answer the secret-key queries *without* switching to using a trapdoor for \mathbf{B}_Y to answer key queries. We highlight the key differences in reduction strategies here:

- Previously (Section 2.2), the reduction sampled \mathbf{u}_i itself and used the trapdoor for \mathbf{B}_Y to sample $\mathbf{r} = \mathbf{H}(\text{gid})$. This was necessary because the reduction did not (and cannot) possess a trapdoor for each \mathbf{A}_i to sample \mathbf{u}_i as in the real scheme. If the reduction did possess such a trapdoor for every i that appears in the challenge ciphertext, then it could trivially break security itself. Then, to ensure consistency of the sampled key with respect to the outputs of \mathbf{H} , this requires the reduction to *program* the outputs of \mathbf{H} . Hence, we model \mathbf{H} as a random oracle in this case.
- In contrast, when we use evasive LWE, the reduction computes $\mathbf{r} = \mathbf{H}(\text{gid})$ normally and then directly constructs \mathbf{u}_i using the terms provided in the evasive LWE challenge. These terms can be simulated *without* knowledge of a trapdoor for \mathbf{A}_i . Observe that this strategy only relies on the ability to compute $\mathbf{H}(\cdot)$, *not* the ability to program its outputs. In general, the evasive LWE assumptions allows us to reduce the task of proving security to that of reasoning about the pseudorandomness of LWE samples with respect to correlated public matrices. In the latter distribution, there are no Gaussian samples, and no need to implement any kind of trapdoor sampling.

When we use evasive LWE, the computation of $\tilde{\mathbf{s}}^\top \tilde{\mathbf{P}}$ essentially translates to computing $\tilde{\mathbf{s}}^\top \mathbf{H}(\text{gid})$, which is pseudorandom by the Boneh et al. [BLMR13] analysis. We refer to Section 6 for the formal description.

While the evasive LWE assumption is much less well understood compared to the classic LWE assumption, proving security under evasive LWE at the minimum indicates that replacing the random oracle with a subset-product hash function is a sound *heuristic* for constructing an MA-ABE scheme in the plain model. It is an interesting challenge to try and prove the security of our construction from the plain LWE assumption; such a proof would provide the first construction of MA-ABE from standard assumptions in the plain model. Alternatively, it is also interesting to further cryptanalyze the evasive LWE assumption.

3 Preliminaries

We write λ to denote the security parameter. For a positive integer $n \in \mathbb{N}$, we write $[n]$ to denote the set $\{1, \dots, n\}$. For a positive integer $q \in \mathbb{N}$, we write \mathbb{Z}_q to denote the integers modulo q . We use bold uppercase letters to denote matrices (e.g., \mathbf{A}, \mathbf{B}) and bold lowercase letters to denote vectors (e.g., \mathbf{u}, \mathbf{v}). We use non-boldface letters to refer to their components: $\mathbf{v} = (v_1, \dots, v_n)$. For matrices $\mathbf{A}_1, \dots, \mathbf{A}_\ell \in \mathbb{Z}_q^{n \times m}$, we write $\text{diag}(\mathbf{A}_1, \dots, \mathbf{A}_\ell) \in \mathbb{Z}_q^{n\ell \times m\ell}$ to denote the block diagonal matrix with blocks $\mathbf{A}_1, \dots, \mathbf{A}_\ell$ along the main diagonal (and 0s elsewhere).

We write $\text{poly}(\lambda)$ to denote a function that is $O(\lambda^c)$ for some $c \in \mathbb{N}$ and $\text{negl}(\lambda)$ to denote a function that is $o(\lambda^{-c})$ for all $c \in \mathbb{N}$. An algorithm is

efficient if it runs in probabilistic polynomial time in its input length. We say that two families of distributions $\mathcal{D}_1 = \{\mathcal{D}_{1,\lambda}\}_{\lambda \in \mathbb{N}}$ and $\mathcal{D}_2 = \{\mathcal{D}_{2,\lambda}\}_{\lambda \in \mathbb{N}}$ are computationally indistinguishable if no efficient algorithm can distinguish them with non-negligible probability. We denote this by writing $\mathcal{D}_1 \stackrel{c}{\approx} \mathcal{D}_2$. We say they are statistically indistinguishable if the statistical distance $\Delta(\mathcal{D}_1, \mathcal{D}_2)$ is bounded by a negligible function in λ and denote this by writing $\mathcal{D}_1 \stackrel{s}{\approx} \mathcal{D}_2$. We say a distribution \mathcal{D} is B -bounded if $\Pr[|x| \leq B : x \leftarrow \mathcal{D}] = 1$.

3.1 Multi-Authority Attribute-Based Encryption

In this section, we introduce the syntax of a multi-authority ABE scheme [LW11]. We start with the definition of a monotone access structure [Bei96].

Definition 3.1 (Access Structure [Bei96]). Let S be a set and let 2^S denote the power set of S (i.e., the set of all subsets of S). An access structure on S is a set $\mathbb{A} \subseteq 2^S \setminus \emptyset$ of non-empty subsets of S . We refer to the elements of \mathbb{A} as the authorized sets and those not in \mathbb{A} as the unauthorized sets. We say an access structure is monotone if for all sets $B, C \in 2^S$, if $B \in \mathbb{A}$ and $B \subseteq C$, then $C \in \mathbb{A}$.

Definition 3.2 (Multi-Authority ABE [LW11, RW15, adapted]). Let λ be a security parameter, \mathcal{M} be a message space, $\mathcal{AU} = \{\mathcal{AU}_\lambda\}_{\lambda \in \mathbb{N}}$ be the universe of authority identifiers, and $\mathcal{GID} = \{\mathcal{GID}_\lambda\}_{\lambda \in \mathbb{N}}$ be the universe of global identifiers for users. To simplify the exposition, we follow the convention in [RW15, DKW21a] and assume that each authority controls a single attribute; this definition generalizes naturally to the setting where each authority controls an arbitrary polynomial number of attributes (see [RW15]). A multi-authority attribute-based encryption scheme for a class of policies $\mathcal{P} = \{\mathcal{P}_\lambda\}_{\lambda \in \mathbb{N}}$ (each described by a monotone access structure on a subset of \mathcal{AU}) consists of a tuple of efficient algorithms $\Pi_{\text{MA-ABE}} = (\text{GlobalSetup}, \text{AuthSetup}, \text{KeyGen}, \text{Encrypt}, \text{Decrypt})$ with the following properties:

- $\text{GlobalSetup}(1^\lambda) \rightarrow \text{gp}$: On input the security parameter λ , the global setup algorithm outputs the global parameters gp .
- $\text{AuthSetup}(\text{gp}, \text{aid}) \rightarrow (\text{pk}_{\text{aid}}, \text{msk}_{\text{aid}})$: On input the global parameters gp and an authority identifier $\text{aid} \in \mathcal{AU}$, the authority setup algorithm outputs a public key pk_{aid} and a master secret key msk_{aid} .
- $\text{KeyGen}(\text{gp}, \text{msk}, \text{gid}) \rightarrow \text{sk}$: On input the global parameters gp , the authority's master secret key msk , and the user identifier $\text{gid} \in \mathcal{GID}$, the key-generation algorithm outputs a decryption key sk .
- $\text{Encrypt}(\text{gp}, \mathbb{A}, \{\text{pk}_{\text{aid}}\}_{\text{aid} \in A}, \mu) \rightarrow \text{ct}$: On input the global parameters gp , an access structure $\mathbb{A} \in \mathcal{P}$ on a set of authorities $A \subseteq \mathcal{AU}$, the set of public keys pk_{aid} associated with each authority $\text{aid} \in A$, and a message $\mu \in \mathcal{M}$, the encryption algorithm outputs a ciphertext ct .
- $\text{Decrypt}(\text{gp}, \{\text{sk}_{\text{aid}}\}_{\text{aid} \in A}, \text{ct}) \rightarrow \mu$: On input the global parameters gp , a collection of secret keys sk_{aid} issued by a set of authorities $\text{aid} \in A$, and a ciphertext ct , the decryption algorithm outputs a message $\mu \in \mathcal{M} \cup \{\perp\}$.

Moreover, $\Pi_{\text{MA-ABE}}$ should satisfy the following properties:

- **Correctness:** There exists a negligible function $\text{negl}(\cdot)$ such that for every $\lambda \in \mathbb{N}$, every message $\mu \in \mathcal{M}$, every identifier $\text{gid} \in \mathcal{GID}_\lambda$, every set of authorities $A \subseteq \mathcal{AU}_\lambda$, every access structure $\mathbb{A} \in \mathcal{P}_\lambda$ on A , and every subset of authorized authorities $B \in \mathbb{A}$,

$$\Pr \left[\begin{array}{l} \text{gp} \leftarrow \text{GlobalSetup}(1^\lambda); \\ \forall \text{aid} \in A : (\text{pk}_{\text{aid}}, \text{msk}_{\text{aid}}) \leftarrow \text{AuthSetup}(\text{gp}, \text{aid}); \\ \mu' = \mu : \forall \text{aid} \in B : \text{sk}_{\text{gid}, \text{aid}} \leftarrow \text{KeyGen}(\text{gp}, \text{msk}_{\text{aid}}, \text{gid}); \\ \text{ct} \leftarrow \text{Encrypt}(\text{gp}, \mathbb{A}, \{\text{pk}_{\text{aid}}\}_{\text{aid} \in A}, \mu); \\ \mu' \leftarrow \text{Decrypt}(\text{gp}, \{\text{sk}_{\text{gid}, \text{aid}}\}_{\text{aid} \in B}, \text{ct}) \end{array} \right] = 1 - \text{negl}(\lambda).$$

- **Static security:** For a security parameter $\lambda \in \mathbb{N}$, an adversary \mathcal{A} , and a bit $b \in \{0, 1\}$, we define the static security game for an multi-authority ABE scheme as follows:
 - **Setup:** The challenger starts by sampling $\text{gp} \leftarrow \text{GlobalSetup}(1^\lambda)$ and gives gp to \mathcal{A} .
 - **Attacker queries.** The adversary \mathcal{A} now specifies the following:
 - * A set $\mathcal{C} \subseteq \mathcal{AU}_\lambda$ of corrupt authorities together with a public key pk_{aid} for each corrupt authority $\text{aid} \in \mathcal{C}$.
 - * A set $\mathcal{N} \subseteq \mathcal{AU}_\lambda$ of non-corrupt authorities, where $\mathcal{N} \cap \mathcal{C} = \emptyset$.
 - * A set $\mathcal{Q} = \{(\text{gid}, A)\}$ of secret key queries where each query consists of a global identifier $\text{gid} \in \mathcal{GID}_\lambda$ and a subset of non-corrupt authorities $A \subseteq \mathcal{N}$.
 - * A pair of challenge messages $\mu_0, \mu_1 \in \mathcal{M}$, a set of authorities $A^* \subseteq \mathcal{C} \cup \mathcal{N}$, and an access structure $\mathbb{A} \in \mathcal{P}_\lambda$ on A^* .
 - **Challenge.** The challenger then samples $(\text{pk}_{\text{aid}}, \text{msk}_{\text{aid}}) \leftarrow \text{AuthSetup}(\text{gp}, \text{aid})$ for each authority $\text{aid} \in \mathcal{N}$. It responds to the adversary with the following:
 - * The public keys pk_{aid} for the non-corrupted authority $\text{aid} \in \mathcal{N}$.
 - * For each secret-key query (gid, A) , the secret keys $\text{sk}_{\text{gid}, \text{aid}} \leftarrow \text{KeyGen}(\text{gp}, \text{msk}_{\text{aid}}, \text{gid})$ for each $\text{aid} \in A$.
 - * The challenge ciphertext $\text{ct}_b \leftarrow \text{Encrypt}(\text{gp}, \mathbb{A}, \{\text{pk}_{\text{aid}}\}_{\text{aid} \in A^*}, \mu_b)$.
 - **Output phase:** Finally, algorithm \mathcal{A} outputs a bit $b' \in \{0, 1\}$, which is the output of the experiment.

We say an adversary \mathcal{A} is admissible for the above security game if $A^* \cap \mathcal{C} \notin \mathbb{A}$ and moreover, for every secret key query (gid, A) , it holds that $(A \cup \mathcal{C}) \cap A^* \notin \mathbb{A}$. Finally, we say $\Pi_{\text{MA-ABE}}$ satisfies static security if for all efficient and admissible adversaries \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $|\Pr[b' = 1 | b = 0] - \Pr[b' = 1 | b = 1]| = \text{negl}(\lambda)$ in the above security game.

Remark 3.3 (Static Security in the Random Oracle Model). Following [RW15, DKW21a], we also extend Definition 3.2 to the random oracle model [BR93]. In this setting, we assume that a global hash function \mathbf{H} (modeled as a random oracle) is published as part of the global public parameters and accessible to all of the parties in the system. When extending static security to the random oracle

model, we require that the adversary submits its random oracle queries as part of its initial query in the static security game. The challenger then includes the responses to the random oracle queries as part of the challenge. We also allow the adversary to further query the random oracle during the challenge phase of the game.

Remark 3.4 (Security Notions). The static security requirement in Definition 3.2 requires that the adversary commits to *all* of its queries upfront. A stronger notion of security is adaptive security under static corruptions [LW11] which requires the adversary pre-commit to the set of corrupted authorities, but thereafter, the adversary can adaptively make secret-key queries both before and after making its challenge ciphertext query. We can also consider intermediate notions where the adversary needs to commit to the policy associated with the challenge ciphertext, but can then issue secret key queries adaptively (i.e., the analog of “selective security” in single-authority ABE). Achieving stronger notions of security (beyond static security) for multi-authority ABE from lattice-based assumptions is an interesting open problem.

Multi-authority ABE for subset policies. Our focus in this work is on constructing multi-authority ABE for the class of subset policies. Here, the ciphertext is associated with a set of authorities A and decryption succeeds whenever a user possesses keys from a set of authorities B where $A \subseteq B$. We define this more formally below.

Definition 3.5 (Multi-Authority ABE for Subset Policies). *Let λ be a security parameter and $\mathcal{AU} = \{\mathcal{AU}_\lambda\}_{\lambda \in \mathbb{N}}$ be the universe of authority identifiers. We define the class of subset policies $\mathcal{P} = \{\mathcal{P}_\lambda\}_{\lambda \in \mathbb{N}}$ to be the set*

$$\mathcal{P}_\lambda = \{\mathbb{A} : \mathbb{A} = \{B : A \subseteq B\} \text{ where } A \subseteq \mathcal{AU}_\lambda\}.$$

Notably, an access structure \mathbb{A} for a subset policy is fully determined by the set $A \subseteq \mathcal{AU}_\lambda$. Thus, when describing an MA-ABE scheme $\Pi_{\text{MA-ABE}} = (\text{GlobalSetup}, \text{AuthSetup}, \text{KeyGen}, \text{Encrypt}, \text{Decrypt})$ for the class of subset policies, we omit the specification of \mathbb{A} in the encryption algorithm and have the encryption algorithm only take as input the public keys associated with the authorities in A . More precisely, we modify the syntax of the encryption algorithm as follows:

- $\text{Encrypt}(\text{gp}, \{\text{pk}_{\text{aid}}\}_{\text{aid} \in A}, \mu) \rightarrow \text{ct}$: *On input the global parameters gp , the set of public keys pk_{aid} associated with each authority $\text{aid} \in A$, and a message $\mu \in \mathcal{M}$, the encryption algorithm outputs a ciphertext ct .*

Remark 3.6 (Multi-Authority ABE for DNFs). A multi-authority ABE scheme for subset policies directly implies a multi-authority ABE scheme for access structures that can be decided by a polynomial-size conjunction or a DNF formula. First, we define the notion of an access structure decidable by a Boolean formula. Let \mathbb{A} be an access structure on a set $A = \{a_1, \dots, a_n\}$. For a subset $B \subseteq A$, we define indicator bits b_1, \dots, b_n where $b_i = 1$ if $a_i \in B$ and 0 otherwise.

We say that \mathbb{A} can be computed by a Boolean formula φ if there exists a Boolean formula $\varphi: \{0, 1\}^n \rightarrow \{0, 1\}$ such that $B \in \mathbb{A}$ if and only if $\varphi(b_1, \dots, b_n) = 1$. It is straightforward to use an MA-ABE scheme for subset policies to construct MA-ABE schemes for policies computable by either a conjunction or a DNF:

- **Conjunction:** Let \mathbb{A} be an access structure on A that is computable by a conjunction on variables b_{i_1}, \dots, b_{i_d} . This is equivalent to a subset policy for the set $\{a_{i_1}, \dots, a_{i_d}\}$.
- **DNF formulas:** Let \mathbb{A} be an access structure on A that is computable by a DNF $\varphi: \{0, 1\}^n \rightarrow \{0, 1\}$. By construction, we can write $\varphi(x_1, \dots, x_n) = \bigvee_{i \in [t]} \varphi_i(x_1, \dots, x_n)$, where each φ_i is a conjunction. In this case, decryption succeeds as long as at least one of the φ_i is satisfied. In this case, we simply concatenate t ciphertexts together, where the i^{th} ciphertext is an encryption to the i^{th} conjunction φ_i . Correctness follows by construction while security follows by a standard hybrid argument.

Remark 3.7 (Multi-Authority ABE for k -CNFs). In the single-authority setting, ABE for subset policies implies an ABE scheme for k -CNF formulas for constant $k \in \mathbb{N}$ [Tsa19, GLW21]. However, this generic approach does not easily translate to the multi-authority setting. Here, a k -CNF formula $\varphi: \{0, 1\}^n \rightarrow \{0, 1\}$ can be written as $\varphi(x_1, \dots, x_n) = \bigwedge_{i \in [t]} \varphi_i(x_1, \dots, x_n)$, where each clause $\varphi_i(x_1, \dots, x_n)$ is a disjunction on up to k variables. To support k -CNF formulas $\varphi: \{0, 1\}^n \rightarrow \{0, 1\}$ on a set $A = \{a_1, \dots, a_n\}$, the approach is to first define a universe U of size $|U| = O(kn^k)$, where each element $u \in U$ is associated with a distinct subset of $S_u \subseteq A$ of size $|S_u| \leq k$. A secret key for a_i consists of secret keys for all $u \in U$ where $a_i \in S_u$. A k -CNF policy $\varphi(x_1, \dots, x_n) = \bigwedge_{i \in [t]} \varphi_i(x_1, \dots, x_n)$ where each clause φ_i depends on a set $T_i \subseteq A$ of at most k variables corresponds to a subset policy for the set $\{u_{T_1}, \dots, u_{T_t}\}$.

In the multi-authority setting, different authorities own the different attributes a_1, \dots, a_n . To implement k -CNF policies as subset policies via the above transformation, we require a multi-authority ABE scheme that supports subset policies where the basic attributes are *combinations* of attributes from *different* authorities. This conflicts with the requirement that authorities be independent in the multi-authority setting. It is an interesting question to construct a multi-authority ABE scheme capable of supporting k -CNF formulas from one that supports subset policies.

3.2 Lattice Preliminaries

Throughout this work, we always use the ℓ_∞ norm for vectors and matrices. Specifically, for a vector \mathbf{u} , we write $\|\mathbf{u}\| := \max_i |x_i|$, and for a matrix \mathbf{A} , we write $\|\mathbf{A}\| = \max_{i,j} |A_{i,j}|$. For a dimension $k \in \mathbb{N}$, we write $\mathbf{I}_k \in \mathbb{Z}_q^{k \times k}$ to denote the k -by- k identity matrix.

Discrete Gaussians. We write $D_{\mathbb{Z}, \chi}$ to denote the (centered) discrete Gaussian distribution over \mathbb{Z} with parameter $\chi \in \mathbb{R}^+$. For a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times t}$, and a vector

$\mathbf{v} \in \mathbb{Z}_q^n$, we write $\mathbf{A}_\chi^{-1}(\mathbf{v})$ to denote a random variable $\mathbf{x} \leftarrow D_{\mathbb{Z}, \chi}^n$ conditioned on $\mathbf{A}\mathbf{x} = \mathbf{v} \bmod q$. We extend \mathbf{A}_s^{-1} to matrices by applying \mathbf{A}_s^{-1} to each column of the input. Throughout this work, we will use the following standard tail bound on Gaussian distributions:

Fact 3.8 (Gaussian Tail Bound). Let λ be a security parameter and $s = s(\lambda)$ be a Gaussian width parameter. Then, for all polynomials $n = n(\lambda)$, there exists a negligible function $\text{negl}(\lambda)$ such that for all $\lambda \in \mathbb{N}$,

$$\Pr [\|\mathbf{v}\| > \sqrt{\lambda} s : \mathbf{v} \leftarrow D_{\mathbb{Z}, s}^n] = \text{negl}(\lambda).$$

Assumption 3.9 (Learning with Errors [Reg05]). Let λ be a security parameter and let $n = n(\lambda)$, $m = m(\lambda)$, $q = q(\lambda)$, $\chi = \chi(\lambda)$ be integers. Then, the decisional learning with errors assumption $\text{LWE}_{n, m, q, \chi}$ states that for $\mathbf{A} \xleftarrow{\mathcal{R}} \mathbb{Z}_q^{n \times m}$, $\mathbf{s} \xleftarrow{\mathcal{R}} \mathbb{Z}_q^n$, $\mathbf{e} \leftarrow D_{\mathbb{Z}, \chi}^m$, and $\mathbf{u} \xleftarrow{\mathcal{R}} \mathbb{Z}_q^m$,

$$(\mathbf{A}, \mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top) \stackrel{c}{\approx} (\mathbf{A}, \mathbf{u}).$$

The gadget matrix. We recall the definition of the gadget matrix [MP12]. For positive integers $n, q \in \mathbb{N}$, let $\mathbf{G}_n = \mathbf{I}_n \otimes \mathbf{g}^\top \in \mathbb{Z}_q^{n \times m}$ be the gadget matrix where $\mathbf{g}^\top = [1, 2, \dots, 2^{\log q - 1}]$ and $m = n \lceil \log q \rceil$. The inverse function $\mathbf{G}_n^{-1}: \mathbb{Z}_q^{n \times t} \rightarrow \mathbb{Z}_q^{m \times t}$ expands each entry $x \in \mathbb{Z}_q$ into a column of size $\lceil \log q \rceil$ consisting of the bits in the binary representation of x . By construction, for every matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times t}$, it follows that $\mathbf{G}_n \cdot \mathbf{G}_n^{-1}(\mathbf{A}) = \mathbf{A} \bmod q$. When the lattice dimension n is clear, we will omit the subscript and simply write \mathbf{G} and $\mathbf{G}^{-1}(\cdot)$ to denote \mathbf{G}_n and $\mathbf{G}_n^{-1}(\cdot)$.

Lattice trapdoors. In this work, we use the gadget trapdoors introduced by Micciancio and Peikert [MP12]. Our description below follows many of the notational conventions from [BTVW17].

Theorem 3.10 (Lattice Trapdoors [Ajt96, GPV08, ABB10b, ABB10a, CHKP10, MP12]). Let n, m, q be lattice parameters. Then there exist efficient algorithms ($\text{TrapGen}, \text{SamplePre}$) with the following syntax:

- $\text{TrapGen}(1^n, q, m) \rightarrow (\mathbf{A}, \text{td}_{\mathbf{A}})$: On input the lattice dimension n , the modulus q , the number of samples m , the trapdoor-generation algorithm outputs a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ together with a trapdoor $\text{td}_{\mathbf{A}}$.
- $\text{SamplePre}(\mathbf{A}, \text{td}_{\mathbf{A}}, \mathbf{v}, s) \rightarrow \mathbf{u}$: On input a matrix \mathbf{A} , a trapdoor $\text{td}_{\mathbf{A}}$, a target vector \mathbf{v} , and a Gaussian width parameter s , the preimage-sampling algorithm outputs a vector \mathbf{u} .

Moreover, there exists a polynomial $m_0 = m_0(n, q) = O(n \log q)$ such that for all $m \geq m_0$, the above algorithms satisfy the following properties:

- **Trapdoor distribution:** The matrix \mathbf{A} output by $\text{TrapGen}(1^n, q, m)$ is statistically close to uniform. Specifically, if $(\mathbf{A}, \text{td}_{\mathbf{A}}) \leftarrow \text{TrapGen}(1^n, q, m)$ and $\mathbf{A}' \xleftarrow{\mathcal{R}} \mathbb{Z}_q^{n \times m}$, then $\Delta(\mathbf{A}, \mathbf{A}') \leq 2^{-n}$.

- **Trapdoor quality:** The trapdoor $\text{td}_{\mathbf{A}}$ output by $\text{TrapGen}(1^n, q, m)$ is a τ -trapdoor where $\tau = O(\sqrt{n \log q \log n})$. We refer to the parameter τ as the quality of the trapdoor.
- **Preimage sampling:** Suppose $\text{td}_{\mathbf{A}}$ is a τ -trapdoor for \mathbf{A} . Then, for all $s \geq \tau \cdot \omega(\sqrt{\log n})$ and all target vectors $\mathbf{v} \in \mathbb{Z}_q^n$, the statistical distance between the following distributions is at most 2^{-n} :

$$\{\mathbf{u} \leftarrow \text{SamplePre}(\mathbf{A}, \text{td}_{\mathbf{A}}, \mathbf{v}, s)\} \quad \text{and} \quad \{\mathbf{u} \leftarrow \mathbf{A}_s^{-1}(\mathbf{v})\}.$$

Gadget trapdoors. In this work, we will work with the gadget trapdoors introduced by Micciancio and Peikert [MP12]. We recall the key properties of gadget trapdoors from [MP12] and then state a direct corollary that we will use in this work (Corollary 3.12).

Theorem 3.11 (Gadget Trapdoors [MP12]). *The gadget matrix $\mathbf{G} \in \mathbb{Z}_q^{n \times m}$ has a public τ -trapdoor $\text{td}_{\mathbf{G}}$ where $\tau = O(1)$. In addition, if $\mathbf{AR} = \mathbf{HG}$ where $\mathbf{A} \in \mathbb{Z}_q^{n \times m'}$, $\mathbf{R} \in \mathbb{Z}_q^{m' \times m}$, $m = n \lceil \log q \rceil$, and $\mathbf{H} \in \mathbb{Z}_q^{n \times n}$ is invertible, then $\text{td}_{\mathbf{A}} = (\mathbf{R}, \mathbf{H})$ can be used as a τ -trapdoor (by extending SamplePre from Theorem 3.10 accordingly) for \mathbf{A} where $\tau = s_1(\mathbf{R})$ and $s_1(\mathbf{R}) \leq \sqrt{mm'} \|\mathbf{R}\|$ denotes the largest singular value of \mathbf{R} .*

Corollary 3.12 (Gadget Trapdoors). *Let $\mathbf{H} \in \mathbb{Z}_q^{k \times t}$ be a full rank matrix where $k \leq t$ (i.e., \mathbf{H} has full row rank). Suppose $\mathbf{AR} = \mathbf{H} \otimes \mathbf{G}$. Let $\mathbf{A} \in \mathbb{Z}_q^{kn \times m'}$ and $\mathbf{R} \in \mathbb{Z}_q^{m' \times mt}$ with $m = n \lceil \log q \rceil$. Then, $\text{td}_{\mathbf{A}} = (\mathbf{R}, \mathbf{H})$ can be used as a τ -trapdoor for \mathbf{A} where $\tau \leq \sqrt{kmm'} \cdot mt \|\mathbf{R}\|$.*

Proof. We can write $\mathbf{H} \otimes \mathbf{G} = (\mathbf{H} \otimes \mathbf{I}_n)(\mathbf{I}_t \otimes \mathbf{G}) = (\mathbf{H} \otimes \mathbf{I}_n)\mathbf{G}_{nt}$. Since \mathbf{H} is full rank (with $k \leq t$), there exists a matrix $\mathbf{H}^* \in \mathbb{Z}_q^{t \times k}$ such that $\mathbf{HH}^* = \mathbf{I}_k$. Correspondingly, $(\mathbf{H} \otimes \mathbf{I}_n)(\mathbf{H}^* \otimes \mathbf{I}_n) = \mathbf{I}_{kn}$. Let $\bar{\mathbf{R}} = \mathbf{R}\mathbf{G}_{nt}^{-1}((\mathbf{H}^* \otimes \mathbf{I}_n)\mathbf{G}_{kn}) \in \mathbb{Z}_q^{m' \times km}$. Now, we can write

$$\mathbf{A}\bar{\mathbf{R}} = \mathbf{A}\mathbf{R}\mathbf{G}_{nt}^{-1}((\mathbf{H}^* \otimes \mathbf{I}_n)\mathbf{G}_{kn}) = (\mathbf{H} \otimes \mathbf{I}_n)\mathbf{G}_{nt}\mathbf{G}_{nt}^{-1}((\mathbf{H}^* \otimes \mathbf{I}_n)\mathbf{G}_{kn}) = \mathbf{G}_{kn},$$

and so $\bar{\mathbf{R}}$ is a trapdoor for \mathbf{A} (Theorem 3.11). Moreover, $\|\bar{\mathbf{R}}\| \leq mt \|\mathbf{R}\|$, and the claim follows.

Preimage sampling. We will also use the following property of discrete Gaussian distributions which follows from [GPV08]:

Lemma 3.13 (Preimage Sampling [GPV08, adapted]). *Let n, m, q be lattice parameters. There exists polynomials $m_0(n, q) = O(n \log q)$ and $\chi_0(n, q) = \sqrt{n \log q} \cdot \omega(\sqrt{\log n})$ such that for all $m \geq m_0(n, q)$ and $\chi \geq \chi_0(n, q)$, the statistical distance between the following distributions is $\text{negl}(n)$:*

$$\{(\mathbf{A}, \mathbf{x}, \mathbf{Ax}) : \mathbf{A} \xleftarrow{\mathcal{R}} \mathbb{Z}_q^{n \times m}, \mathbf{x} \leftarrow D_{\mathbb{Z}, \chi}^m\} \quad \text{and} \quad \{(\mathbf{A}, \mathbf{x}, \mathbf{y}) : \mathbf{A} \xleftarrow{\mathcal{R}} \mathbb{Z}_q^{n \times m}, \mathbf{y} \xleftarrow{\mathcal{R}} \mathbb{Z}_q^n, \mathbf{x} \leftarrow \mathbf{A}_{\chi}^{-1}(\mathbf{y})\}.$$

Lemma 3.14 (Leftover Hash Lemma [ABB10a]). *Let n, m, q be lattice parameters where $q > 2$ is prime. There exists a polynomial $m_0(n, q) = O(n \log q)$ such that for all $m \geq m_0(n, q)$, all vectors $\mathbf{e} \in \mathbb{Z}_q^m$, and all polynomials $k = k(n)$, the statistical distance between the following distributions is $\text{negl}(n)$:*

$$\begin{aligned} & \{(\mathbf{A}, \mathbf{AR}, \mathbf{e}^\top \mathbf{R}) : \mathbf{A} \xleftarrow{\mathbb{R}} \mathbb{Z}_q^{n \times m}, \mathbf{R} \xleftarrow{\mathbb{R}} \{-1, 1\}^{m \times k}\} \\ & \text{and} \\ & \{(\mathbf{A}, \mathbf{B}, \mathbf{e}^\top \mathbf{R}) : \mathbf{A} \xleftarrow{\mathbb{R}} \mathbb{Z}_q^{n \times m}, \mathbf{B} \xleftarrow{\mathbb{R}} \mathbb{Z}_q^{n \times k}, \mathbf{R} \xleftarrow{\mathbb{R}} \{-1, 1\}^{m \times k}\}. \end{aligned} \quad (3.1)$$

Smudging lemma. We will also use the following standard smudging lemma (see [BDE⁺18] for a proof):

Lemma 3.15 (Smudging Lemma). *Let λ be a security parameter. Take any $e \in \mathbb{Z}$ where $|e| \leq B$. Suppose $\chi \geq B \cdot \lambda^{\omega(1)}$. Then, the statistical distance between the distributions $\{z : z \leftarrow D_{\mathbb{Z}, \chi}\}$ and $\{z + e : z \leftarrow D_{\mathbb{Z}, \chi}\}$ is $\text{negl}(\lambda)$.*

The evasive LWE assumption. We now state a variant of the evasive LWE assumption introduced by Wee [Wee22] and Tsabury [Tsa22]. We compare our formulation with the original version by Wee in Remark 3.18.

Assumption 3.16 (Evasive LWE). Let λ be a security parameter, and let $n = n(\lambda), m = m(\lambda), q = q(\lambda), \chi = \chi(\lambda), s = s(\lambda)$ with $s \geq O(\sqrt{m \log q})$. Let **Samp** be an algorithm that takes the security parameter 1^λ as input and outputs a matrix $\mathbf{B} \in \mathbb{Z}_q^{n \times m'}$, a set of ℓ target matrices $\mathbf{P}_1 \in \mathbb{Z}_q^{n \times N_1}, \dots, \mathbf{P}_\ell \in \mathbb{Z}_q^{n \times N_\ell}$, and auxiliary information $\text{aux} \in \{0, 1\}^*$. Then, for adversaries \mathcal{A}_0 and \mathcal{A}_1 , we define advantage functions

$$\begin{aligned} \text{Adv}_{\mathcal{A}_0}^{(\text{PRE})}(\lambda) &:= \left| \Pr [\mathcal{A}_0(\{(\mathbf{A}_i, \mathbf{s}_i^\top \mathbf{A}_i + \mathbf{e}_{1,i}^\top)\}_{i \in [\ell]}, \mathbf{B}, \mathbf{s}^\top \mathbf{B} + \mathbf{e}_2^\top, \{\mathbf{s}_i^\top \mathbf{P}_i + \mathbf{e}_{3,i}^\top\}_{i \in [\ell]}, \text{aux}) = 1] \right. \\ &\quad \left. - \Pr [\mathcal{A}_0(\{(\mathbf{A}_i, \mathbf{u}_{1,i}^\top)\}_{i \in [\ell]}, \mathbf{B}, \mathbf{u}_2^\top, \{\mathbf{u}_{3,i}^\top\}_{i \in [\ell]}, \text{aux}) = 1] \right| \\ \text{Adv}_{\mathcal{A}_1}^{(\text{POST})}(\lambda) &:= \left| \Pr [\mathcal{A}_1(\{(\mathbf{A}_i, \mathbf{s}_i^\top \mathbf{A}_i + \mathbf{e}_{1,i}^\top)\}_{i \in [\ell]}, \mathbf{B}, \mathbf{s}^\top \mathbf{B} + \mathbf{e}_2^\top, \{\mathbf{K}_i\}_{i \in [\ell]}, \text{aux}) = 1] \right. \\ &\quad \left. - \Pr [\mathcal{A}_1(\{(\mathbf{A}_i, \mathbf{u}_{1,i}^\top)\}_{i \in [\ell]}, \mathbf{B}, \mathbf{u}_2^\top, \{\mathbf{K}_i\}_{i \in [\ell]}, \text{aux}) = 1] \right|, \end{aligned}$$

where

$$\begin{aligned} & (\mathbf{B}, \mathbf{P}_1, \dots, \mathbf{P}_\ell, \text{aux}) \leftarrow \text{Samp}(1^\lambda), \\ & \mathbf{A}_1, \dots, \mathbf{A}_\ell \xleftarrow{\mathbb{R}} \mathbb{Z}_q^{n \times m}, \\ & \mathbf{s}_1, \dots, \mathbf{s}_\ell \xleftarrow{\mathbb{R}} \mathbb{Z}_q^n, \mathbf{s}^\top \leftarrow [\mathbf{s}_1^\top \mid \dots \mid \mathbf{s}_\ell^\top] \in \mathbb{Z}_q^{n\ell}, \\ & \mathbf{u}_{1,i} \xleftarrow{\mathbb{R}} \mathbb{Z}_q^m, \mathbf{e}_{1,i} \leftarrow D_{\mathbb{Z}, \chi}^m \quad \forall i \in [\ell], \\ & \mathbf{u}_2 \xleftarrow{\mathbb{R}} \mathbb{Z}_q^{m'}, \mathbf{e}_2 \leftarrow D_{\mathbb{Z}, \chi}^{m'}, \\ & \mathbf{u}_{3,i} \xleftarrow{\mathbb{R}} \mathbb{Z}_q^{N_i}, \mathbf{e}_{3,i} \leftarrow D_{\mathbb{Z}, \chi}^{N_i} \quad \forall i \in [\ell], \\ & \mathbf{K}_i \leftarrow (\mathbf{A}_i)_s^{-1}(\mathbf{P}_i) \quad \forall i \in [\ell]. \end{aligned}$$

We say that the evasive LWE assumption holds if for every efficient sampler **Samp** and every efficient adversary \mathcal{A}_1 , there exists an efficient algorithm \mathcal{A}_0 , polynomial $\text{poly}(\cdot)$, and negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$,

$$\text{Adv}_{\mathcal{A}_0}^{(\text{PRE})}(\lambda) \geq \text{Adv}_{\mathcal{A}_1}^{(\text{POST})}(\lambda) / \text{poly}(\lambda) - \text{negl}(\lambda).$$

Remark 3.17 (Auxiliary Input Distribution). As in [Wee22], we only require that the assumption holds for samplers where \mathbf{aux} additionally contains all of the coin tosses used by \mathbf{Samp} (i.e., *public-coin* samplers). This avoids obfuscation-based counter-examples where \mathbf{aux} contains an obfuscation of a program related to a trapdoor for matrix \mathbf{B} or \mathbf{P}_i . This is a *weaker* assumption compared to the evasive LWE assumptions needed to realize witness encryption (which rely on security of evasive LWE to hold for private-coin samplers) [Tsa22, VWW22].

Remark 3.18 (Comparison with [Wee22]). The original formulation of the evasive LWE assumption by Wee [Wee22] corresponds to the special case where $\ell = 1$ (i.e., there is just a single matrix \mathbf{A}_1 and single target \mathbf{P}_1). When constructing multi-authority ABE, we rely on multiple *independent* matrices $\mathbf{A}_1, \dots, \mathbf{A}_\ell$ (one associated with each authority). It is an interesting question to reduce Assumption 3.16 to the simpler setting of $\ell = 1$. We note that the justification given in [Wee22] for evasive LWE are equally applicable to this setting.

4 Generalized Related-Trapdoor LWE Assumption

In this section, we introduce a generalized variant of the related-trapdoor robust LWE assumption of Brakerski and Vaikuntanathan [BV22] and then show that its hardness can be based on the standard LWE assumption (Theorem 4.2). As described in Section 2, the generalized related-trapdoor LWE assumption essentially asserts that given a vector $\mathbf{u} \in \{0, 1\}^L$, an LWE sample with respect to $(\mathbf{u} \otimes \mathbf{I}_n)\mathbf{B}$ is pseudorandom (where $\mathbf{B} \in \mathbb{Z}_q^{n \times mL}$) given an oracle that takes as input (\mathbf{M}, \mathbf{t}) and outputs $(\mathbf{M} \otimes \mathbf{I}_n)\mathbf{B}^{-1}(\mathbf{t})$ whenever $\bar{\mathbf{M}} = \begin{bmatrix} \mathbf{M} \\ \mathbf{u}^\top \end{bmatrix} \in \mathbb{Z}_q^{(k+1) \times L}$ is full rank. The original formulation of the related trapdoor assumption in [BV22] (for the setting of *single-authority* ciphertext-policy ABE) considered the special case where the matrix $\mathbf{M} \in \mathbb{Z}_q^{1 \times L}$ is a row vector. Here, we consider the case where \mathbf{M} can be an arbitrary matrix. This generalization will be useful for distributing the setup in an ABE scheme to obtain a multi-authority ABE (see Section 5).

A similar approach is also implicit in the ciphertext-policy ABE scheme by Datta et al. [DKW21a]. Their approach relied on noise smudging to simulate the preimage-sampling oracle, and as such, security relied on a super-polynomial modulus. In this work, we abstract out the core technique through the related-trapdoor LWE assumption and then show a direct reduction to LWE without relying on noise smudging. This allows us to base security on LWE with a *polynomial* modulus.

Assumption 4.1 (Generalized Related-Trapdoor LWE). Let $\lambda \in \mathbb{N}$ be a security parameter, and $n = n(\lambda)$, $m = m(\lambda)$, $\hat{m} = \hat{m}(\lambda)$, and $\chi = \chi(\lambda)$ be lattice parameters. Let $q = q(\lambda)$ be a prime modulus. Let $L = L(\lambda)$ be a length parameter. For a bit $b \in \{0, 1\}$, we define the related-trapdoor LWE game between a challenger and an adversary \mathcal{A} :

1. The adversary \mathcal{A} starts by choosing a non-zero vector $\mathbf{u} \in \{0, 1\}^L$.

2. The challenger samples matrices $\mathbf{A} \xleftarrow{\mathbb{R}} \mathbb{Z}_q^{n \times m}$ and $\mathbf{B} \xleftarrow{\mathbb{R}} \mathbb{Z}_q^{nL \times \hat{m}(2L-1)}$ and constructs the challenge as follows:
 - If $b = 0$, the challenger samples $\mathbf{s} \xleftarrow{\mathbb{R}} \mathbb{Z}_q^n$, $\mathbf{R} \xleftarrow{\mathbb{R}} \{-1, 1\}^{\hat{m}L \times \hat{m}(L-1)}$, $\mathbf{e} \leftarrow D_{\mathbb{Z}, \chi}^m$, $\hat{\mathbf{e}}_0 \leftarrow D_{\mathbb{Z}, \chi}^{\hat{m}L}$, $\hat{\mathbf{e}}^\top \leftarrow \hat{\mathbf{e}}_0^\top [\mathbf{I}_{\hat{m}L} \mid \mathbf{R}] \in \mathbb{Z}_q^{\hat{m}(2L-1)}$, and gives $(\mathbf{A}, \mathbf{B}, \mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top, \mathbf{s}^\top (\mathbf{u}^\top \otimes \mathbf{I}_n) \mathbf{B} + \hat{\mathbf{e}}^\top)$ to \mathcal{A} .
 - If $b = 1$, the challenger samples $\mathbf{v} \xleftarrow{\mathbb{R}} \mathbb{Z}_q^m$, $\hat{\mathbf{v}} \xleftarrow{\mathbb{R}} \mathbb{Z}_q^{\hat{m}(2L-1)}$ and gives $(\mathbf{A}, \mathbf{B}, \mathbf{v}^\top, \hat{\mathbf{v}}^\top)$ to \mathcal{A} .
3. Adversary \mathcal{A} can now make queries of the form (\mathbf{M}, \mathbf{t}) where $\mathbf{M} \in \mathbb{Z}_q^{k \times L}$ where $k < L$ and $\mathbf{t} \in \mathbb{Z}_q^{kn}$.
 - Define the matrix $\bar{\mathbf{M}} = \begin{bmatrix} \mathbf{M} \\ \mathbf{u}^\top \end{bmatrix}$. If $\bar{\mathbf{M}}$ is not full rank (over \mathbb{Z}_q), the challenger replies with \perp .
 - If \mathbf{t} is not in the column span of $(\mathbf{M} \otimes \mathbf{I}_n) \mathbf{B}$, then the challenger also replies with \perp .
 - Otherwise, it samples and replies with $\mathbf{y} \leftarrow ((\mathbf{M} \otimes \mathbf{I}_n) \mathbf{B})_\chi^{-1}(\mathbf{t})$. Namely, $\mathbf{y} \in \mathbb{Z}_q^{m(2L-1)}$ is sampled from the distribution $D_{\mathbb{Z}, \chi}^{m(2L-1)}$ conditioned on $(\mathbf{M} \otimes \mathbf{I}_n) \mathbf{B} \mathbf{y} = \mathbf{t}$.
4. At the end of the game, algorithm \mathcal{A} outputs a bit $b' \in \{0, 1\}$, which is also the output of the experiment.

We say that the $\text{RTLWE}_{n, m, \hat{m}, q, \chi, L}$ assumption holds if for all efficient adversaries \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $|\Pr[b' = 1 \mid b = 0] - \Pr[b' = 1 \mid b = 1]| = \text{negl}(\lambda)$ in the above security game.

The generalized related-trapdoor LWE assumption reduces to the vanilla LWE assumption. We state the formal theorem below and refer to the full version of this paper [WWW22] for a proof.

Theorem 4.2 (Generalized Related-Trapdoor LWE). *Let λ be a security parameter, and let $n = n(\lambda)$, $q = q(\lambda)$, $m = m(\lambda)$, $\hat{m} = \hat{m}(\lambda)$, and $\chi = \chi(\lambda)$ be lattice parameters. Suppose that $q > 2$ is a prime and $\chi > 2\hat{m}^2 L^2 \cdot \omega(\sqrt{\log n})$. Then, there exists a fixed polynomial $m_0(n, q) = O(n \log q)$ such that for all $\hat{m} > m_0(n, q)$ and under the $\text{LWE}_{n, m + \hat{m}L, q, \chi}$ assumption, the $\text{RTLWE}_{n, m, \hat{m}, q, \chi, L}$ assumption holds.*

5 Multi-Authority ABE from LWE in the Random Oracle Model

In this section, we describe our construction of multi-authority ABE for the family of subset policies in the random oracle model. Our construction follows a similar structure as the multi-authority ABE scheme of Datta et al. [DKW21a] except we provide a direct reduction to the (generalized) related trapdoor LWE problem (Section 4). Notably, this allows us to base security on polynomial hardness of the plain LWE assumption with a *polynomial* modulus. The previous construction of Datta et al. relied on LWE with a *super-polynomial* modulus-to-noise ratio.

Construction 5.1 (Multi-Authority ABE in the Random Oracle Model).

Let λ be a security parameter, and $n = n(\lambda)$, $m = m(\lambda)$, $q = q(\lambda)$, and $\chi = \chi(\lambda)$ be lattice parameters. Let $L = L(\lambda)$ be a bound on the number of attributes associated with a ciphertext. Let $\mathcal{GID} = \{0, 1\}^\lambda$ be the set of user identifiers and $\mathcal{AU} = \{0, 1\}^\lambda$ be the set of authority identifiers. The construction will rely on a hash function $H: \mathcal{GID} \rightarrow \mathbb{Z}_q^{m(2L-1)}$, which will be modeled as a random oracle as follows:

- For ease of exposition in the following description, we will start by assuming that the outputs of the random oracle H are distributed according to a *discrete Gaussian distribution*. Specifically, on every input $\text{gid} \in \mathcal{GID}$, the output $H(\text{gid})$ is a sample from the distribution $D_{\mathbb{Z}, \chi}^{m(2L-1)}$. In [Section 5.1](#) and [Remark 5.7](#), we show that using inversion sampling, we can implement H using a *standard* random oracle $H': \mathcal{GID} \rightarrow \{0, 1\}^{\lambda m(2L-1)}$, where the outputs of $H'(\text{gid})$ are distributed uniformly over $\{0, 1\}^{\lambda m(2L-1)}$ as usual.

We construct a multi-authority ABE scheme for subset policies with message space $\mathcal{M} = \{0, 1\}$ as follows:

- **GlobalSetup**(1^λ): Output the global parameters $\text{gp} = (\lambda, n, m, q, \chi, L, H)$.
- **AuthSetup**(gp, aid): On input the global parameters gp and an authority identifier $\text{aid} \in \mathcal{AU}$, sample $(\mathbf{A}_{\text{aid}}, \text{td}_{\text{aid}}) \leftarrow \text{TrapGen}(1^n, q, m)$, $\mathbf{p}_{\text{aid}} \xleftarrow{\mathbb{R}} \mathbb{Z}_q^n$, and $\mathbf{B}_{\text{aid}} \xleftarrow{\mathbb{R}} \mathbb{Z}_q^{n \times m(2L-1)}$. Output the authority public key $\text{pk}_{\text{aid}} \leftarrow (\mathbf{A}_{\text{aid}}, \mathbf{B}_{\text{aid}}, \mathbf{p}_{\text{aid}})$ and the authority secret key $\text{msk}_{\text{aid}} = \text{td}_{\text{aid}}$.
- **KeyGen**($\text{gp}, \text{msk}, \text{pk}, \text{gid}$): On input the global parameters $\text{gp} = (\lambda, n, m, q, \chi, L, H)$, the master secret key $\text{msk} = \text{td}$, the public key $\text{pk} = (\mathbf{A}, \mathbf{B}, \mathbf{p})$, and the user identifier gid , the key-generation algorithm computes $\mathbf{r} \leftarrow H(\text{gid}) \in \mathbb{Z}_q^{m(2L-1)}$ and uses td to sample $\mathbf{u} \leftarrow \mathbf{A}_\chi^{-1}(\mathbf{p} + \mathbf{B}\mathbf{r})$. It outputs $\text{sk}_{\text{aid}, \text{gid}} = \mathbf{u}$.
- **Encrypt**($\text{gp}, \{\text{pk}_{\text{aid}}\}_{\text{aid} \in A}, \mu$): On input the global parameters $\text{gp} = (\lambda, n, m, q, \chi, L, H)$, a set of public keys $\text{pk}_{\text{aid}} = (\mathbf{A}_{\text{aid}}, \mathbf{B}_{\text{aid}}, \mathbf{p}_{\text{aid}})$ associated with a set of authorities A , and the message $\mu \in \{0, 1\}$, the encryption algorithm samples $\mathbf{s}_{\text{aid}} \xleftarrow{\mathbb{R}} \mathbb{Z}_q^n$, $\mathbf{e}_{1, \text{aid}} \leftarrow D_{\mathbb{Z}, \chi}^n$, $\mathbf{R} \xleftarrow{\mathbb{R}} \{0, 1\}^{mL \times m(L-1)}$, $\hat{\mathbf{e}}_2 \leftarrow D_{\mathbb{Z}, \chi}^{mL}$, and $\mathbf{e}_2^\top \leftarrow \hat{\mathbf{e}}_2^\top [\mathbf{I}_{mL} \mid \mathbf{R}]$, and $e_3 \leftarrow D_{\mathbb{Z}, \chi}$ for each $\text{aid} \in A$. It outputs the ciphertext

$$\text{ct} = \left(\left\{ \mathbf{s}_{\text{aid}}^\top \mathbf{A}_{\text{aid}} + \mathbf{e}_{1, \text{aid}}^\top \right\}_{\text{aid} \in A}, \sum_{\text{aid} \in A} \mathbf{s}_{\text{aid}}^\top \mathbf{B}_{\text{aid}} + \mathbf{e}_2^\top, \sum_{\text{aid} \in A} \mathbf{s}_{\text{aid}}^\top \mathbf{p}_{\text{aid}} + e_3 + \mu \cdot \lfloor q/2 \rfloor \right).$$

- **Decrypt**($\text{gp}, \{\text{sk}_{\text{aid}, \text{gid}}\}_{\text{aid} \in A}, \text{ct}, \text{gid}$): On input the global parameters $\text{gp} = (\lambda, n, m, q, \chi, L, H)$, a set of secret keys $\text{sk}_{\text{aid}, \text{gid}} = \mathbf{u}_{\text{aid}, \text{gid}}$ associated with authorities $\text{aid} \in A$ and user identifier gid , and a ciphertext $\text{ct} = (\{\mathbf{c}_{1, \text{aid}}^\top\}_{\text{aid} \in A}, \mathbf{c}_2^\top, c_3)$, the decryption algorithm computes $\mathbf{r} \leftarrow H(\text{gid})$ and outputs

$$\left\lfloor \frac{2}{q} \cdot \left(c_3 + \mathbf{c}_2^\top \mathbf{r} - \sum_{\text{aid} \in A} \mathbf{c}_{1, \text{aid}}^\top \mathbf{u}_{\text{aid}, \text{gid}} \bmod q \right) \right\rfloor.$$

We now state the correctness and security theorems for [Construction 5.1](#), but defer their formal proofs to the full version of this paper [\[WWW22\]](#).

Theorem 5.2 (Correctness). *Suppose the conditions of [Theorem 3.10](#) and [Lemma 3.13](#) hold (i.e., $m \geq m_0(n, q) = O(n \log q)$ and $\chi > \chi_0(n, q) = \sqrt{n \log q} \cdot \omega(\sqrt{\log n})$). Then, there exists a polynomial $q_0 = O(\lambda \chi^2 m^2 L^2)$ such that for all $q > q_0$, [Construction 5.1](#) is correct.*

Theorem 5.3 (Static Security). *Suppose the conditions of [Theorem 3.10](#) and [Lemma 3.13](#) hold (i.e., $m \geq m_0(n, q) = O(n \log q)$ and $\chi > \chi_0(n, q) = \sqrt{n \log q} \cdot \omega(\sqrt{\log n})$). Then, under the $\text{RTLWE}_{n, Lm+1, m, q, \chi, L}$ assumption and modeling $\mathbf{H}: \mathcal{GID} \rightarrow \mathbb{Z}_q^{m(2L-1)}$ as a random oracle (with outputs distributed according to $D_{\mathbb{Z}, \chi}^{m(2L-1)}$), [Construction 5.1](#) is statically secure.*

Parameter setting. Let λ be a security parameter. We can now instantiate [Construction 5.1](#) as follows:

- We set the lattice dimension $n = \lambda$.
- To rely on [Theorem 5.3](#), we rely on the $\text{RTLWE}_{n, Lm+1, m, q, \chi, L}$ assumption. By [Theorem 4.2](#), this reduces to $\text{LWE}_{n, 2Lm+1, q, \chi}$ if we set $m = O(n \log q)$, $q > 2$ to a prime, and $\chi = O(m^2 L^2 \log n)$.
- For correctness ([Theorem 5.2](#)), we additionally require $q = O(\lambda \chi^2 m^2 L^2)$.

In particular, this means we can choose m, q, χ to be polynomials in λ , and thus, base hardness on LWE with a *polynomial* modulus-to-noise ratio. We summarize the instantiation below:

Corollary 5.4 (Multi-Authority ABE for Subset Policies in the Random Oracle Model). *Let λ be a security parameter. Assuming polynomial hardness of LWE with a polynomial modulus-to-noise ratio, there exists a statically-secure multi-authority ABE scheme for subset policies of a priori bounded length $L = L(\lambda)$ in the random oracle model. The size of the ciphertext scales quasi-linearly with the bound L .*

5.1 Instantiating using a Random Oracle with Uniform Outputs

As described, [Construction 5.1](#) and [Corollary 5.4](#) relies on a random oracle $\mathbf{H}: \mathcal{GID} \rightarrow \mathbb{Z}_q^{m(2L-1)}$ whose output distribution is the discrete Gaussian distribution $D_{\mathbb{Z}, \chi}^{m(2L-1)}$. Since $\chi = \text{poly}(\lambda)$ in our setting, we describe a simple way to instantiate \mathbf{H} using a random oracle $\mathbf{H}': \mathcal{GID} \rightarrow \{0, 1\}^{\lambda m(2L-1)}$ whose output distribution is the uniform distribution via inversion sampling. The function \mathbf{H}' coincides with the usual way we model the output distribution of a random oracle [[BR93](#)].

Previously, Brakerski et al. [[BCTW16](#)] sketched an alternative approach for instantiating a random oracle outputting samples from a discrete Gaussian distribution by adapting the rejection sampler of Lyubashevsky and Wichs [[LW15](#)]. Datta et al. [[DKW21a](#)] rely on noise smudging in their setting (which would in turn necessitate using a super-polynomial modulus-to-noise ratio). In our setting where we have a distribution with polynomial-size support, we describe a simple alternative based on inversion sampling. This is a simple approach used in concrete implementations of lattice-based cryptography [[BCD⁺16](#)].

Lemma 5.5 (Inversion Sampling). *Let λ be a security parameter, $t = t(\lambda)$ be an input length, and D be a discrete B -bounded distribution with an efficiently-computable cumulative distribution function. Then, there exists a pair of efficient algorithms (Project, SampleR) with the following properties:*

- **Project**(x) $\rightarrow y$: On input an input $x \in \{0, 1\}^t$, the projection algorithm outputs a sample $y \in [-B, B]$. The projection algorithm is deterministic.
- **SampleR**(y) $\rightarrow x$: On input a value $y \in [-B, B]$, the reverse sampling algorithm outputs an $x \in \{0, 1\}^t$.

In addition, the following properties hold:

- **Correctness:** For all $y \in [-B, B]$, $\Pr[\text{Project}(\text{SampleR}(y)) = y] = 1$.
- **Reverse-sampleability:** For all $t > \log B + \omega(\log \lambda)$, the following two distributions are statistically indistinguishable:

$$\{(x, \text{Project}(x)) : x \xleftarrow{R} \{0, 1\}^t\} \quad \text{and} \quad \{(\text{SampleR}(y), y) : y \leftarrow D\}.$$

Proof. We take (Project, SampleR) to be the standard inversion sampling algorithm. Let $f : [-B - 1, B] \rightarrow [0, 1]$ be the cumulative distribution function for D , and let $T = 2^t - 1$. We construct the two algorithms as follows:

- **Project**(x): On input $x \in \{0, 1\}^t$, let $X \in [0, T]$ be the integer whose binary representation is x . Output $y \in [-B, B]$ where $T \cdot f(y - 1) < X \leq T \cdot f(y)$.
- **SampleR**(y): On input $y \in [-B, B]$, let $x_0 \leftarrow T \cdot f(y - 1)$ and $x_1 \leftarrow T \cdot f(y)$. Output the binary representation of the element $x \xleftarrow{R} (x_0, x_1] \cap \mathbb{Z}$.

Since the cumulative distribution function f is efficiently-computable and the Project algorithm can be computed with $\text{polylog}(B)$ calls to f (e.g., using binary search), the Project algorithm is efficiently-computable. The SampleR algorithm only requires making two calls to f and is likewise efficient. Next, correctness of the algorithm follows by construction. Finally, for the reverse-sampleability property, take any $Y \in [-B, B]$. Then,

$$\begin{aligned} \Pr[\text{Project}(x) = Y : x \xleftarrow{R} \{0, 1\}^t] &= \frac{\lfloor T \cdot f(Y) \rfloor - \lfloor T \cdot f(Y - 1) \rfloor}{T} = f(Y) - f(Y - 1) + e \\ &= \Pr[y = Y : y \leftarrow D] + e \end{aligned}$$

where $|e| \leq 2/T$. Thus, the statistical distance between $\{\text{Project}(x) : x \xleftarrow{R} \{0, 1\}^t\}$ and D is at most $2(2B+1)/T = \text{negl}(\lambda)$. Finally, on input $y \in [-B, B]$, SampleR(y) outputs a uniform $x \xleftarrow{R} \{0, 1\}^t$ conditioned on $\text{Project}(x) = y$.

Remark 5.6 (Extending to Product Distributions). We can extend (Project, SampleR) to sample from a product distribution D^n in the natural way. The projection algorithm takes as input a vector of bit-strings $\mathbf{x} \in (\{0, 1\}^t)^n$ and applies the projection operator component-wise. The reverse sampling algorithm is defined analogously. Correctness and reverse-sampleability then follow via a standard hybrid argument.

Remark 5.7 (Implementing the Random Oracle in Corollary 5.4). We can now implement the random oracle $H: \mathcal{GID} \rightarrow \mathbb{Z}_q^{m(2L-1)}$ in Corollary 5.4 (whose outputs are distributed according to $D_{\mathbb{Z},\chi}^{m(2L-1)}$) with a random oracle $H': \mathcal{GID} \rightarrow \{0,1\}^{\lambda m(2L-1)}$ whose outputs are uniform as follows:

- Let $\tilde{D}_{\mathbb{Z},\chi}$ be the discrete Gaussian distribution $D_{\mathbb{Z},\chi}$ truncated to the interval $[-\sqrt{\lambda}\chi, \sqrt{\lambda}\chi]$. Namely, to sample $\tilde{x} \leftarrow \tilde{D}_{\mathbb{Z},\chi}$, we first sample $x \leftarrow D_{\mathbb{Z},\chi}$ and output x if $x \in [-\sqrt{\lambda}\chi, \sqrt{\lambda}\chi]$ and output 0 otherwise. By Fact 3.8, $\tilde{D}_{\mathbb{Z},\chi}$ is statistically indistinguishable from $D_{\mathbb{Z},\chi}$. In addition, $\tilde{D}_{\mathbb{Z},\chi}$ is B -bounded for $B = \sqrt{\lambda}\chi$.
- Let (Project, SampleR) be the inversion sampling algorithm from Lemma 5.5 and Remark 5.6 for the product distribution $\tilde{D}_{\mathbb{Z},\chi}^{m(2L-1)}$. We now define

$$H(\text{gid}) := \text{Project}(H'(\text{gid})).$$

Since $\chi = \chi(\lambda)$ is polynomially-bounded, the cumulative distribution function of $\tilde{D}_{\mathbb{Z},\chi}$ is efficiently-computable. Then, by Lemma 5.5 and Remark 5.6, for all polynomial-size collections of distinct inputs $\text{gid}_1, \dots, \text{gid}_\ell \in \mathcal{GID}$, the joint distributions of

$$\{H(\text{gid}_i)\}_{i \in [\ell]} \quad \text{and} \quad \left\{ \mathbf{r}_i \leftarrow D_{\mathbb{Z},\chi}^{m(2L-1)} \right\}_{i \in [\ell]}$$

are statistically indistinguishable.

- Finally, the proof of Theorem 5.3 critically relies on the ability to *program* the outputs of the random oracle in the reduction. Here, we rely on the SampleR algorithm. Namely, to program $H(\text{gid})$ to a vector $\mathbf{r}_{\text{gid}} \leftarrow D_{\mathbb{Z},\chi}^{m(2L-1)}$, the reduction algorithm would sample $x_{\text{gid}} \leftarrow \text{SampleR}(\mathbf{r}_{\text{gid}})$ and program $H'(\text{gid})$ to x_{gid} . This induces the correct distribution by Lemma 5.5 and Remark 5.6.

6 Multi-Authority ABE without Random Oracles

We now give our construction of a multi-authority ABE scheme without random oracles. Specifically, we instantiate the hash function from Construction 5.1 with a subset-product construction (i.e., the lattice-based PRF from Theorem 6.1) and then prove security under the evasive LWE assumption (Assumption 3.16) and lattice-based PRFs [BPR12, BLMR13].

Lattice-based PRFs. Our analysis will rely on an *unrounded* lattice-based PRF. We state the theorem and provide a proof sketch below, and refer readers to [BPR12, Theorem 5.2] for a more formal exposition. Our presentation here is adapted from the work of Chen et al. [CVW18, Lemma 7.4] who use a similar theorem for analyzing the security of their private constrained PRF construction. We state the theorem here and refer to [BPR12, BLMR13, CVW18] for the proof.

Theorem 6.1 (Lattice-Based PRFs [BPR12, BLMR13]). *Let λ be a security parameter and let $n = n(\lambda)$, $q = q(\lambda)$, $\chi = \chi(\lambda)$, $k = k(\lambda)$ be integers. Let $\chi_{\text{smudge}} = \chi_{\text{smudge}}(\lambda)$ be a noise parameter that will be used for noise smudging. Let $\boldsymbol{\eta} \in \mathbb{Z}_q^k$ be the first elementary basis vector (i.e., $\eta_1 = 1$ and $\eta_i = 0$ for all $i \neq 1$). For a bit $b \in \{0, 1\}$, an input length $\tau = \tau(\lambda)$, and an adversary \mathcal{A} , define the following pseudorandomness game between a challenger and \mathcal{A} :*

1. *The challenger begins by sampling $(\mathbf{D}_0, \mathbf{D}_1) \xleftarrow{\mathbb{R}} D_{\mathbb{Z}, \chi}^{k \times k}$ and a secret key $\mathbf{s} \xleftarrow{\mathbb{R}} \mathbb{Z}_q^k$. It gives \mathbf{D}_0 and \mathbf{D}_1 to \mathcal{A} .*
2. *Algorithm \mathcal{A} can now adaptively submit queries $x \in \{0, 1\}^\tau$ to the challenger. If $b = 0$, the challenger samples $e_x \xleftarrow{\mathbb{R}} D_{\mathbb{Z}, \chi_{\text{smudge}}}$ and outputs*

$$f_{\mathbf{D}_0, \mathbf{D}_1, \mathbf{s}}(x) := \mathbf{s}^\top \left(\prod_{i \in [\tau]} \mathbf{D}_{x_i} \right) \boldsymbol{\eta} + e_x \in \mathbb{Z}_q. \quad (6.1)$$

Otherwise, if $b = 1$, the challenger replies with $y \xleftarrow{\mathbb{R}} \mathbb{Z}_q$.

3. *After \mathcal{A} is done making queries, it outputs a bit $b' \in \{0, 1\}$, which is the output of the experiment.*

An adversary \mathcal{A} is admissible if all of the queries it submits are distinct. Then, for all polynomials $\tau = \tau(\lambda)$, $q = q(\lambda)$, parameters $k \geq 6n \log q$, $\chi = \Omega(\sqrt{n} \log q)$, $\chi_{\text{smudge}} > \lambda^{\tau + \omega(1)} \cdot (k\chi)^\tau$, and assuming the $\text{LWE}_{n, m, q, \chi}$ assumption for some $m = \text{poly}(k, \tau, Q)$, for all efficient and admissible adversaries \mathcal{A} making up to Q queries, there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $|\Pr[b' = 1 : b = 0] - \Pr[b' = 1 : b = 1]| = \text{negl}(\lambda)$.

MA-ABE for subset policies without random oracles. We now give the full construction of our MA-ABE scheme without random oracles. As described in Section 2, our construction essentially instantiates the random oracle in Construction 5.1 with a subset-product of low-norm matrices (which can be used as the basis for constructing a PRF according to Theorem 6.1). Arguing security in turn relies on the evasive LWE assumption (Assumption 3.16). Using the evasive LWE assumption to argue security has the extra benefit of allowing support for policies of arbitrary (polynomial) length (recall that Construction 5.1 as well as the previous lattice-based construction of Datta et al. [DKW21a] required imposing an *a priori* bound on the policy length, and the size of the ciphertext in turn grew with the maximum length).

Construction 6.2 (Multi-Authority ABE without Random Oracles).

Let λ be a security parameter, and $n = n(\lambda)$, $m = m(\lambda)$, $q = q(\lambda)$, and $\chi = \chi(\lambda)$ be lattice parameters. Let $\chi_{\text{PRF}} = \chi_{\text{PRF}}(\lambda)$ be a Gaussian width parameter used to define the hash function. Let $\tau = \tau(\lambda)$ be the bit-length of identities and let $\mathcal{GID} = \{0, 1\}^\tau$ be the set of user identifiers. Let $\mathcal{AU} = \{0, 1\}^\lambda$ be the set of authorities. We construct an MA-ABE scheme for subset policies (Definition 3.5) with message space $\mathcal{M} = \{0, 1\}$ as follows:

- **GlobalSetup**(1^λ): Sample $\mathbf{D}_0, \mathbf{D}_1 \leftarrow D_{\mathbb{Z}, \chi_{\text{PRF}}}^{m \times m}$. Define the hash function $H: \{0, 1\}^\tau \rightarrow \mathbb{Z}_q^m$ by the function $H(x) := \left(\prod_{i \in [\tau]} \mathbf{D}_{x_i} \right) \boldsymbol{\eta}$ where $\boldsymbol{\eta} \in \mathbb{Z}_q^m$ is the first canonical basis vector (i.e., $\eta_1 = 1$ and $\eta_i = 0$ for all $i \neq 1$). Output

$$\mathbf{gp} = (\lambda, n, m, q, \chi, \chi_{\text{PRF}}, \tau, \mathbf{D}_0, \mathbf{D}_1).$$

For ease of exposition, whenever we write $H(\cdot)$ in the following, we refer to the hash function defined by the matrices $\mathbf{D}_0, \mathbf{D}_1$ in the global parameters.

- **AuthSetup**(\mathbf{gp}, aid): On input the global parameters $\mathbf{gp} = (\lambda, n, m, q, \chi, \chi_{\text{PRF}}, \tau, \mathbf{D}_0, \mathbf{D}_1)$ and an authority identifier $\text{aid} \in \mathcal{AU}$, sample $(\mathbf{A}_{\text{aid}}, \text{td}_{\text{aid}}) \leftarrow \text{TrapGen}(1^n, q, m)$, $\mathbf{p}_{\text{aid}} \xleftarrow{\mathbb{R}} \mathbb{Z}_q^n$, and $\mathbf{B}_{\text{aid}} \xleftarrow{\mathbb{R}} \mathbb{Z}_q^{n \times m}$. Output the authority public key $\mathbf{pk}_{\text{aid}} \leftarrow (\mathbf{A}_{\text{aid}}, \mathbf{B}_{\text{aid}}, \mathbf{p}_{\text{aid}})$ and the authority secret key $\mathbf{msk}_{\text{aid}} = \text{td}_{\text{aid}}$.
- **KeyGen**($\mathbf{gp}, \mathbf{msk}, \mathbf{pk}, \text{gid}$): On input the global parameters $\mathbf{gp} = (\lambda, n, m, q, \chi, \chi_{\text{PRF}}, \tau, \mathbf{D}_0, \mathbf{D}_1)$, the master secret key $\mathbf{msk} = \text{td}$, the public key $\mathbf{pk} = (\mathbf{A}, \mathbf{B}, \mathbf{p})$, the user identifier $\text{gid} \in \{0, 1\}^\tau$, the key-generation algorithm computes $\mathbf{r} \leftarrow H(\text{gid}) \in \mathbb{Z}_q^m$ and uses td to sample $\mathbf{u} \leftarrow \mathbf{A}_{\chi}^{-1}(\mathbf{p} + \mathbf{B}\mathbf{r})$. It outputs $\mathbf{sk}_{\text{aid}, \text{gid}} = \mathbf{u}$.
- **Encrypt**($\mathbf{gp}, \{\mathbf{pk}_{\text{aid}}\}_{\text{aid} \in A}, \mu$): On input the global parameters $\mathbf{gp} = (\lambda, n, m, q, \chi, \chi_{\text{PRF}}, \tau, \mathbf{D}_0, \mathbf{D}_1)$, a set of public keys $\mathbf{pk}_{\text{aid}} = (\mathbf{A}_{\text{aid}}, \mathbf{B}_{\text{aid}}, \mathbf{p}_{\text{aid}})$ associated with a set of authorities A , and the message $\mu \in \{0, 1\}$, the encryption algorithm samples $\mathbf{s}_{\text{aid}} \xleftarrow{\mathbb{R}} \mathbb{Z}_q^n$, $\mathbf{e}_{1, \text{aid}} \leftarrow D_{\mathbb{Z}, \chi}^m$, $\mathbf{e}_2 \leftarrow D_{\mathbb{Z}, \chi}^m$, and $e_3 \leftarrow D_{\mathbb{Z}, \chi}$ for each $\text{aid} \in A$. It outputs the ciphertext

$$\text{ct} = \left(\left\{ \mathbf{s}_{\text{aid}}^\top \mathbf{A}_{\text{aid}} + \mathbf{e}_{1, \text{aid}}^\top \right\}_{\text{aid} \in A}, \sum_{\text{aid} \in A} \mathbf{s}_{\text{aid}}^\top \mathbf{B}_{\text{aid}} + \mathbf{e}_2^\top, \sum_{\text{aid} \in A} \mathbf{s}_{\text{aid}}^\top \mathbf{p}_{\text{aid}} + e_3 + \mu \cdot \lfloor q/2 \rfloor \right).$$

- **Decrypt**($\mathbf{gp}, \{\mathbf{sk}_{\text{aid}, \text{gid}}\}_{\text{aid} \in A}, \text{ct}, \text{gid}$): On input the global parameters $\mathbf{gp} = (\lambda, n, m, q, \chi, \chi_{\text{PRF}}, \tau, \mathbf{D}_0, \mathbf{D}_1)$, a set of secret keys $\mathbf{sk}_{\text{aid}, \text{gid}} = \mathbf{u}_{\text{aid}, \text{gid}}$ associated with authorities $\text{aid} \in A$ and user identifier gid , a ciphertext $\text{ct} = (\{\mathbf{c}_{1, \text{aid}}^\top\}_{\text{aid} \in A}, \mathbf{c}_2, \mathbf{c}_3)$, the decryption algorithm computes $\mathbf{r} \leftarrow H(\text{gid})$ and outputs

$$\left\lfloor \frac{2}{q} \cdot \left(c_3 + \mathbf{c}_2^\top \mathbf{r} - \sum_{\text{aid} \in A} \mathbf{c}_{1, \text{aid}}^\top \mathbf{u}_{\text{aid}, \text{gid}} \bmod q \right) \right\rfloor.$$

Correctness and security analysis. We now state the correctness and security properties of [Construction 6.2](#), but due to space limitations, defer their proofs to the full version of this paper [\[WWW22\]](#).

Theorem 6.3 (Correctness). *Let $L = L(\lambda)$ be a bound on the number of attributes associated with a ciphertext. Suppose the conditions of [Theorem 3.10](#) and [Lemma 3.13](#) hold (i.e., $m \geq m_0(n, q) = O(n \log q)$ and $\chi > \chi_0(n, q) = \sqrt{n \log q} \cdot \omega(\sqrt{\log n})$). Then, there exists $q_0 = O(Lm\lambda\chi^2 + (\sqrt{\lambda}m\chi_{\text{PRF}})^{\tau+1}\chi)$ such that for all $m > m_0$, $q > q_0$, and $\chi > \chi_0$, [Construction 6.2](#) is correct.*

Theorem 6.4 (Static Security). *There exists a polynomial $m_0(n, q) = O(n \log q)$ such that under the following conditions and assumptions, [Construction 6.2](#) is statically secure:*

- The number of samples m satisfies $m \geq m_0$.
- Let $\chi_{\text{smudge}} = \chi_{\text{smudge}}(\lambda)$ be a smudging parameter where $\chi_{\text{smudge}} > \lambda^{\tau+\omega(1)}(m\chi_{\text{PRF}})^{\tau+1}$.
- The noise parameter χ satisfies $\chi > \lambda^{\omega(1)}\ell\chi_{\text{smudge}}$.
- The $\text{LWE}_{n,m',q,\chi_{\text{PRF}}}$ assumption holds where $m' = \text{poly}(m, \tau, Q)$ and Q is a bound on the number of secret-key queries the adversary makes.
- The evasive LWE assumption with parameters $n, m, q, \chi, s = \chi$ holds (in particular, the preimages $\mathbf{K} \leftarrow \mathbf{A}^{-1}(\mathbf{P})$ are distributed according to a discrete Gaussian with parameter $s = \chi$).

Parameter setting. Let λ be a security parameter. We now instantiate [Construction 6.2](#) as follows:

- Let the lattice dimension be $n = \lambda^{1/\varepsilon}$ for some constant $\varepsilon > 0$.
- We can set the length of the identities gid to be $\tau = \lambda$.
- For security ([Theorem 6.4](#)), we require that $\chi_{\text{smudge}} > \lambda^{\lambda+\omega(1)}(m\chi_{\text{PRF}})^{\lambda+1}$ and $\chi > \lambda^{\omega(1)}\ell\chi_{\text{smudge}}$. Each of $\ell = \ell(\lambda), m = m(\lambda), \chi_{\text{PRF}} = \chi_{\text{PRF}}(\lambda)$ are polynomially-bounded. Thus, we can set $\chi = 2^{\tilde{O}(n^\varepsilon)}$ to satisfy these requirements, where $\tilde{O}(\cdot)$ suppresses constant and logarithmic factors.
- To support arbitrary polynomial-size ciphertext policies, we set the bound $L = 2^\lambda$ in [Theorem 6.3](#). To ensure correctness, we can set $m = O(n \log q)$ and $q = O(2^\lambda m \lambda \chi^2 + (\lambda m \chi_{\text{PRF}})^{\lambda+1} \chi)$. Setting $q = 2^{\tilde{O}(n^\varepsilon)}$ suffices to satisfy these requirements.

This yields the following corollary:

Corollary 6.5 (Multi-Authority ABE for Subset Policies from Evasive LWE). *Assuming polynomial hardness of LWE and of the evasive LWE assumption (both with a sub-exponential modulus-to-noise ratio), there exists a statically-secure multi-authority ABE for subset policies (of arbitrary polynomial size).*

Acknowledgments

We thanks the TCC reviewers for helpful suggestions. B. Waters is supported by NSF CNS-1908611, a Simons Investigator award, and the Packard Foundation Fellowship. D. J. Wu is supported by NSF CNS-2151131, CNS-2140975, a Microsoft Research Faculty Fellowship, and a Google Research Scholar award.

References

- ABB10a. Shweta Agrawal, Dan Boneh, and Xavier Boyen. Efficient lattice (H)IBE in the standard model. In *EUROCRYPT*, pages 553–572, 2010.
- ABB10b. Shweta Agrawal, Dan Boneh, and Xavier Boyen. Lattice basis delegation in fixed dimension and shorter-ciphertext hierarchical IBE. In *CRYPTO*, pages 98–115, 2010.

- Ajt96. Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In *STOC*, pages 99–108, 1996.
- BCD⁺16. Joppe W. Bos, Craig Costello, Léo Ducas, Ilya Mironov, Michael Naehrig, Valeria Nikolaenko, Ananth Raghunathan, and Douglas Stebila. Frodo: Take off the ring! practical, quantum-secure key exchange from LWE. In *ACM CCS*, pages 1006–1018, 2016.
- BCTW16. Zvika Brakerski, David Cash, Rotem Tsabary, and Hoeteck Wee. Targeted homomorphic attribute-based encryption. In *TCC*, pages 330–360, 2016.
- BDE⁺18. Jonathan Bootle, Claire Delaplace, Thomas Espitau, Pierre-Alain Fouque, and Mehdi Tibouchi. LWE without modular reduction and improved side-channel attacks against BLISS. In *ASIACRYPT*, pages 494–524, 2018.
- Bei96. Amos Beimel. *Secure Schemes for Secret Sharing and Key Distribution*. PhD thesis, Technion, 1996.
- BLMR13. Dan Boneh, Kevin Lewi, Hart William Montgomery, and Ananth Raghunathan. Key homomorphic prfs and their applications. In *CRYPTO*, pages 410–428, 2013.
- BPR12. Abhishek Banerjee, Chris Peikert, and Alon Rosen. Pseudorandom functions and lattices. In *EUROCRYPT*, pages 719–737, 2012.
- BR93. Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *ACM CCS*, pages 62–73, 1993.
- BTWV17. Zvika Brakerski, Rotem Tsabary, Vinod Vaikuntanathan, and Hoeteck Wee. Private constrained prfs (and more) from LWE. In *TCC*, pages 264–302, 2017.
- BV22. Zvika Brakerski and Vinod Vaikuntanathan. Lattice-inspired broadcast encryption and succinct ciphertext-policy ABE. In *ITCS*, pages 28:1–28:20, 2022.
- CC09. Melissa Chase and Sherman S. M. Chow. Improving privacy and security in multi-authority attribute-based encryption. In *ACM CCS*, pages 121–130, 2009.
- Cha07. Melissa Chase. Multi-authority attribute based encryption. In *TCC*, pages 515–534, 2007.
- CHKP10. David Cash, Dennis Hofheinz, Eike Kiltz, and Chris Peikert. Bonsai trees, or how to delegate a lattice basis. In *EUROCRYPT*, pages 523–552, 2010.
- CVW18. Yilei Chen, Vinod Vaikuntanathan, and Hoeteck Wee. GGH15 beyond permutation branching programs: Proofs, attacks, and candidates. In *CRYPTO*, pages 577–607, 2018.
- DKW21a. Pratish Datta, Ilan Komargodski, and Brent Waters. Decentralized multi-authority ABE for DNFs from LWE. In *EUROCRYPT*, pages 177–209, 2021.
- DKW21b. Pratish Datta, Ilan Komargodski, and Brent Waters. Decentralized multi-authority ABE for nc^1 from computational-bdh. *IACR Cryptol. ePrint Arch.*, page 1325, 2021.
- GGSW13. Sanjam Garg, Craig Gentry, Amit Sahai, and Brent Waters. Witness encryption and its applications. In *STOC*, pages 467–476, 2013.
- GLW21. Rishab Goyal, Jiahui Liu, and Brent Waters. Adaptive security via deletion in attribute-based encryption: Solutions from search assumptions in bilinear groups. In *ASIACRYPT*, pages 311–341, 2021.
- GPSW06. Vipul Goyal, Omkant Pandey, Amit Sahai, and Brent Waters. Attribute-based encryption for fine-grained access control of encrypted data. In *ACM CCS*, pages 89–98, 2006.

- GPV08. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pages 197–206, 2008.
- Kim19. Sam Kim. Multi-authority attribute-based encryption from LWE in the OT model. *IACR Cryptol. ePrint Arch.*, page 280, 2019.
- LCLS08. Huang Lin, Zhenfu Cao, Xiaohui Liang, and Jun Shao. Secure threshold multi authority attribute based encryption without a central authority. In *INDOCRYPT*, pages 426–436, 2008.
- LW11. Allison B. Lewko and Brent Waters. Decentralizing attribute-based encryption. In *EUROCRYPT*, pages 568–588, 2011.
- LW15. Vadim Lyubashevsky and Daniel Wichs. Simple lattice trapdoor sampling from a broad class of distributions. In *PKC*, pages 716–730, 2015.
- MKE08. Sascha Müller, Stefan Katzenbeisser, and Claudia Eckert. Distributed attribute-based encryption. In *ICISC*, pages 20–36, 2008.
- MP12. Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *EUROCRYPT*, pages 700–718, 2012.
- Reg05. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In *STOC*, pages 84–93, 2005.
- RW15. Yannis Rouselakis and Brent Waters. Efficient statically-secure large-universe multi-authority attribute-based encryption. In *Financial Cryptography and Data Security*, pages 315–332, 2015.
- SW05. Amit Sahai and Brent Waters. Fuzzy identity-based encryption. In *EUROCRYPT*, pages 457–473, 2005.
- Tsa19. Rotem Tsabary. Fully secure attribute-based encryption for t-CNF from LWE. In *CRYPTO*, pages 62–85, 2019.
- Tsa22. Rotem Tsabary. Candidate witness encryption from lattice techniques. In *CRYPTO*, 2022.
- VWW22. Vinod Vaikuntanathan, Hoeteck Wee, and Daniel Wichs. Witness encryption and null-IO from evasive LWE. In *ASIACRYPT*, 2022.
- Wee22. Hoeteck Wee. Optimal broadcast encryption and CP-ABE from evasive lattice assumptions. In *EUROCRYPT*, 2022.
- WFL19. Zhedong Wang, Xiong Fan, and Feng-Hao Liu. FE for inner products and its application to decentralized ABE. In *PKC*, pages 97–127, 2019.
- WWW22. Brent Waters, Hoeteck Wee, and David J. Wu. Multi-authority ABE from lattices without random oracles. *IACR Cryptol. ePrint Arch.*, 2022.