Unbounded Distributed Broadcast Encryption and Registered ABE from Succinct LWE

Hoeteck Wee	David J. Wu
NTT Research	UT Austin
wee@di.ens.fr	dwu4@cs.utexas.edu

Abstract

We construct distributed broadcast encryption and registered attribute-based encryption (ABE) that support an arbitrary polynomial of users from the succinct LWE assumption. Specifically, if we take λ to be the security parameter and N to be the number of users, we obtain the following:

- We obtain a distributed broadcast encryption scheme where the size of the public parameters, user public/secret keys, and ciphertexts are optimal (i.e., have size poly(λ, log N)). Security relies on the poly(λ, log N)-succinct LWE assumption. Previously, this was only known from indistinguishability obfuscation or witness encryption. All constructions that did not rely on these general tools could only support an a priori bounded number of users.
- We obtain a key-policy registered ABE scheme that supports arbitrary bounded-depth Boolean circuit policies from the poly(λ , d, log N)-succinct LWE assumption in the random oracle model, where d is the depth of the circuit computing the policy. The public parameters, user public/secret keys, and ciphertexts have size poly(λ , d, log N), which are optimal up to the poly(d) factor. This is the first registered ABE scheme with nearly-optimal parameters. All previous schemes (including constructions based on indistinguishability obfuscation, witness encryption, or evasive LWE) either have ciphertexts that scale with the policy size and attribute length, or can only support a bounded number of users (with long public parameters and public keys that scale with the number of users).

1 Introduction

In recent years, registration-based cryptography [GHMR18] has emerged as a popular paradigm for realizing advanced cryptographic primitives without a trusted authority. In this work, we study two trustless cryptographic primitives: distributed broadcast encryption and registered attribute-based encryption (registered ABE):

- Distributed broadcast encryption: In broadcast encryption [FN93], a sender can encrypt a message to
 an arbitrary set of recipients with a ciphertext whose size scales sublinearly with the number of recipients.
 Traditionally, broadcast encryption requires a central authority who issues decryption keys to each recipient
 in the system. Distributed broadcast encryption [WQZD10, BZ14] removes the central authority. Instead,
 each receiver independently generates their own public/secret key and publishes it in a public-key directory.
 Thereafter, a sender can encrypt a message to an arbitrary set of public keys with a ciphertext whose size scales
 sublinearly with the number of recipients.
- **Registered ABE:** Attribute-based encryption [SW05, GPSW06] is a generalization of public-key encryption that provides fine-grained access control to encrypted data. In key-policy ABE, decryption keys are associated with a policy f while ciphertexts are associated with an attribute x. Decryption recovers the associated message m only if the decryption policy is satisfied (e.g., if f(x) = 0). Like the case with broadcast encryption, the standard notion of an ABE scheme assumes the existence of a central trusted authority that issues keys to users. Registered ABE [HLWW23] removes this trusted authority. In this model, each user independently generates their own public/secret key and publishes it to a public-key directory. Then, there is a process that aggregates all of the individual public keys (with their associated decryption policies) into a succinct master public key

mpk. The aggregated master public key functions as the public key for an ABE scheme (i.e., one can encrypt a message with respect to any attribute x with the guarantee that only registered users associated with an accepting policy f can decrypt).

In both settings, the goal is to support the capabilities provided by broadcast encryption or attribute-based encryption without needing to assume the existence of a trusted authority. In the traditional centralized versions of these primitives, there is a central point of failure. If the adversary compromises the central authority, they immediately compromise the security of every user in the system. Registration-based cryptography ensures that individual users retain control of their secret keys and that trust is not concentrated in any single entity.

Bounded vs. unbounded. In the last few years, there has been a flurry of activity in constructing distributed broadcast encryption [WQZD10, BZ14, KMW23, FWW23, GKPW24, CW24, CHW25] and registered ABE [HLWW23, FWW23, ZZGQ23, AT24, GLWW24, CHW25, ZZC⁺25] from different cryptographic assumptions. Thus far, these works can be categorized into two main categories:

- Unbounded constructions: The first class of constructions are those that support an arbitrary polynomial of users. Namely, these are schemes where any number of users can join the system by posting their public key to the public-key directory. This is the most natural formulation of these primitives (and an implicit requirement in the original work introducing registration-based cryptography [GHMR18]). Constructing schemes that support an arbitrary number of users has thus far relied on strong tools such as indistinguishability obfuscation [BZ14, HLWW23], witness encryption for NP [FWW23], or strong non-falsifiable assumptions [ZZC⁺25].
- **Bounded constructions:** A second line of work has focused on direct algebraic constructions of trustless cryptographic primitives that do not need general-purpose tools like indistinguishability obfuscation or witness encryption. These constructions often have the advantage of being simpler, concretely-efficient, and in the case of lattice-based schemes, plausibly post-quantum secure. However, with the exception of [ZZC⁺25], which is based on private-coin evasive LWE, all of the constructions in this family [WQZD10, HLWW23, KMW23, ZZGQ23, AT24, GLWW24, CW24, GKPW24, CHW25] make a significant compromise by assuming there is an a priori *maximum* number of users that the scheme can support. Moreover, in these schemes, the parameters as well as the size of each user's public key grows linearly (or worse) with the maximum number of users. For schemes with large number of users, the size of these parameters introduces significant overhead.

A major challenge in the study of trustless cryptography is developing new techniques that support an arbitrary number of users from simple (and falsifiable) assumptions (and without relying on heavy machinery such as program obfuscation or witness encryption).

1.1 Our Results

In this work, we develop new techniques for constructing distributed broadcast encryption and (key-policy) registered ABE for general policies that support an *a priori* unbounded number of users. Security of both constructions relies on the succinct LWE assumption [Wee24], and our registered ABE scheme additionally relies on the random oracle model. The succinct LWE assumption is a simple and falsifiable lattice assumption (which is also implied by public-coin evasive LWE). We now provide a brief comparison of our results with those from prior work (see also Tables 1 and 2).

Distributed broadcast encryption. Our first result is a distributed broadcast encryption scheme that supports an arbitrary polynomial number of users from the poly(λ , log N)-succinct LWE assumption, where λ is the security parameter and N is the total number of users. The size of the public parameters, the user public/secret keys, as well as the ciphertext in our scheme are all poly(λ , log N). Previously, this was only known from indistinguishability obfuscation [BZ14] or from plain witness encryption [FWW23]. Compared to these approaches, our algebraic approach is conceptually simpler. For instance, the [FWW23] construction relies on witness encryption together with a function-binding hash function; this latter primitive in turn relies on fully homomorphic encryption. In contrast, our approach for distributed broadcast encryption does not need any kind of homomorphic computation machinery.

Scheme	Assumption	pp	pk	sk	ct	ТР	PQ	AD
[BZ14] [FWW23]	iO + one-way function witness encryption + LWE	- 1	1 1	1 1	1 1	✓ ✓	× √	✓* ✓*
[WQZD10] [KMW23] [GKPW24]	bilinear Diffie-Hellman exponent bilinear Diffie-Hellman exponent generic bilinear group	N N N	$egin{array}{c} N^2 \ N \ N \ N \end{array}$	N 1 1	1 1 1	✓ × ×	X X X	× ✓* ×
[CW24] [CHW25]	poly (λ, N) -succinct LWE poly (λ, N) -succinct LWE †	$N^2 onumber N^2$	N N	1 1	1 1	x x	√ √	× ✓
This work	$poly(\lambda, \log N)$ -succinct LWE	1	1	1	1	X^{\ddagger}	1	×

*These schemes were originally shown to satisfy *semi-static* security. Adaptive security can then be obtained by using the Gentry-Waters compiler [GW09] in the random oracle model or the more recent Hsieh-Waters-Wu compiler [HWW25], which gives adaptive security in the plain model.

[†]Security of this scheme additionally relies on the random oracle model.

 ‡ We can obtain a variant of this construction with transparent setup by using the recently-introduced decomposed LWE assumption [AMR25]. We refer to Remark 4.13 for more details.

Table 1: Comparison with previous distributed broadcast encryption schemes. For each scheme, we report the size of the public parameters pp, the user public key pk, the user secret key sk, and the ciphertext ct as a function of the number of users *N*. For ease of comparison, we suppress $poly(\lambda, \log N)$ factors, where λ is the security parameter. For each scheme, we also indicate whether the public parameters pp (if present) can be sampled with a *transparent* setup (TP), whether it is plausibly post-quantum secure (PQ), and whether it is proven to be adaptively secure (AD). We write *iO* to denote indistinguishability obfuscation [BGI⁺01].

Compared to direct algebraic constructions based on bilinear maps [WQZD10, KMW23, GKPW24] or succinct LWE [CW24, CHW25], our scheme is the only one that supports an unbounded number of users. In previous schemes, both the size of the public parameters and the size of an individual user's public key scale with the number of users *N*. As we discuss more in Section 2.1, all of these schemes were bounded because each user's public key needed to include a "cross-term" for every other user in the system. A key technical contribution of this work is a new technique for efficiently deriving the cross-terms from a succinct commitment. This enables algebraic schemes that support an unbounded number of users. We provide a more detailed comparison in Table 1.

Key-policy registered ABE. Our second construction is a key-policy registered ABE scheme for general policies (modeled as bounded-depth Boolean circuits) from the succinct LWE assumption in the random oracle model. Our scheme supports an arbitrary number of users and moreover, has succinct ciphertexts (of size $poly(\lambda, d, \log N)$, where *d* is the depth of the circuit and *N* is the number of users). Notably, the size of the ciphertext does not scale with the size of the policy circuit or with the length of the attribute. Prior to this work, the only scheme with this level of ciphertext succinctness is the scheme from [CHW25] from succinct LWE (in the random oracle model), but that scheme only supports a bounded number of users (and requires long public parameters and user public keys). The other constructions from obfuscation [HLWW23], witness encryption [FWW23], or private-coin evasive LWE [ZZC⁺25] all have ciphertexts whose size scales linearly with the attribute length. Thus, we obtain the first unbounded registered ABE scheme with nearly-optimal ciphertext size (i.e., optimal up to the poly(*d*) factor). We provide a more detailed comparison in Table 2.

Concurrent work. In a concurrent and independent work, Abram, Malavolta, and Roy [AMR25] introduced the decomposed LWE assumption, a weaker variant of the succinct LWE assumption, and show (among other things) how to obtain a registered ABE scheme that supports general circuit policies in the plain model. We provide a brief comparison of our two schemes:

• **Bounded vs. unbounded:** The [AMR25] registered ABE scheme relies on a structured reference string that contains *N* secret keys for a (centralized) ABE scheme, where *N* is a bound on the number of parties. As a result,

Scheme	Assumption	pp	pk	sk	ct	ТР	PQ	AD
[HLWW23]	iO + SSB hash function	1	1	1	$poly(C , \mathbf{x})$	\	×	√
[FWW23]	witness encryption + LWE	1	1	1	$poly(C , \mathbf{x})$	\	√	×
[ZZC ⁺ 25]	private-coin evasive LWE poly(λ, d, N)-succinct LWE*	poly(d)	poly(d)	poly(<i>d</i>)	$poly(d, \mathbf{x})$	√	\$	x
[CHW25]		$N^2 \cdot poly(d)$	$N \cdot poly(d)$	poly(<i>d</i>)	poly(d)	×	\$	x
This work	$poly(\lambda, d, \log N)$ -succinct LWE*	poly(d)	poly(d)	poly(d)	poly(d)	X‡	1	×

*Security of these schemes are in the random oracle model.

 ‡ We can obtain a variant of this construction with transparent setup by using the recently-introduced decomposed LWE assumption [AMR25]. We refer to Remark 5.31 for more details.

Table 2: Comparison with previous (key-policy) registered ABE schemes that support general circuit policies. For each scheme, we report the size of the public parameters pp, the user public key pk, the user secret key sk, and the ciphertext ct as a function of the number of users *N*, the attribute **x**, and the policy circuit *C*. We write *d* to denote the depth of the circuit *C*. We assume the decryption algorithm is provided the description of the policy circuit *C* as well as the attribute $|\mathbf{x}|$ as input, so these do not necessarily have to be encoded as part of the secret key or the ciphertext. For ease of comparison, we suppress poly(λ , log *N*) factors, where λ is the security parameter. For each scheme, we also indicate whether the public parameters pp can be sampled with a *transparent* setup (TP), whether it is plausibly post-quantum secure (PQ), and whether it is proven to be adaptively secure (AD). We write *iO* to denote indistinguishability obfuscation [BGI⁺01] and "SSB hash function" to refer to a somewhere-statistically-binding hash function [HW15].

their scheme only supports an a priori bounded number of parties. Our scheme requires a structured reference string whose size scales *logarithmically* with *N*, and thus, can support an arbitrary polynomial number of parties. Supporting an unbounded polynomial number of parties is the central goal of this work. As noted in Remark 5.31, we can also obtain a version of our scheme with a transparent setup (where the public parameters now consist of a uniform random string) by using the decomposed LWE assumption from [AMR25].

- Security definition: The work of [AMR25] analyze security of their scheme under a "very selective" model where the adversary has to choose its public keys before seeing the CRS and the honest parties' keys. In our setting, we consider the standard security definition for registered ABE where the adversary can register malicious keys of its choosing (*after* seeing the CRS and the honest parties' keys). The standard security model for registered ABE captures adversarial strategies such as rogue-key attacks (see [RY07, BDN18] for examples of rogue-key attacks in the context of aggregate signatures). Since the broader goal of registration-based cryptography is to allow users to choose their own keys, it seems unreasonable to remove this capability from the adversary. As discussed in more detail in [CHW25, §1.2], we are not aware of any generic techniques that lifts a registered ABE scheme secure in the very selective model to the standard security definition for registered ABE.
- Hardness assumption: Security of the [AMR25] construction relies on the decomposed LWE assumption in the plain model whereas our registered ABE scheme relies on the succinct LWE assumption in the random oracle model. The work [AMR25] show that the decomposed LWE assumption is a weaker assumption than succinct LWE. As we note in Remarks 4.13 and 5.31, we can also modify our scheme to obtain security based on the decomposed LWE assumption (and in fact, the resulting scheme would have the added benefit of having a transparent setup). In either case, security of our registered ABE scheme would still rely on the random oracle. Like [CHW25], we rely on the random oracle to handle the ability of the adversary to register malicious keys.

2 Technical Overview

In the following, let n, m, q be lattice parameters where $m = O(n \log q)$. We use curly underlines to suppress low-norm errors. Namely, we write $\mathbf{s}_{q}^{\mathsf{T}}\mathbf{A}$ to denote $\mathbf{s}^{\mathsf{T}}\mathbf{A} + \mathbf{e}^{\mathsf{T}}$, where \mathbf{e} is a low-norm error vector. For a matrix $\mathbf{B} \in \mathbb{Z}_{q}^{n \times m}$ and

a target vector $\mathbf{z} \in \mathbb{Z}_q^n$, we write $\mathbf{y} \leftarrow \mathbf{B}^{-1}(\mathbf{z})$ to denote sampling $\mathbf{y} \in \mathbb{Z}_q^m$ from a discrete Gaussian distribution conditioned on $\mathbf{B}\mathbf{y} = \mathbf{z}$. We write G to denote the standard gadget matrix [MP12].

 ℓ -succinct LWE. Our constructions rely on the succinct LWE assumption introduced by Wee [Wee24]. For a parameter ℓ , the ℓ -succinct LWE assumption asserts that the following distributions are computationally indistinguishable:

$$(\mathbf{B}, \mathbf{s}^{\mathsf{T}}\mathbf{B}, \mathbf{W}, \mathbf{T})$$
 and $(\mathbf{B}, \mathbf{u}^{\mathsf{T}}, \mathbf{W}, \mathbf{T})$

where $\mathbf{B} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^{n \times m}$, $\mathbf{s} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^n$, $\mathbf{W} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^{\ell n \times m}$, \mathbf{T} is a random Gaussian matrix where $[\mathbf{I}_\ell \otimes \mathbf{B} \mid \mathbf{W}]\mathbf{T} = \mathbf{I}_\ell \otimes \mathbf{G}$, and \mathbf{I}_ℓ is the identity matrix of dimension ℓ . Equivalently, ℓ -succinct LWE asserts that LWE is hard with respect to the (random) matrix \mathbf{B} given a trapdoor for the related matrix $[\mathbf{I}_\ell \otimes \mathbf{B} \mid \mathbf{W}]$. In the following, we will often refer to $(\mathbf{B}, \mathbf{W}, \mathbf{T})$ as the public parameters for a succinct LWE instance of dimension ℓ .

Matrix commitment scheme. The key ingredient we use underlying our distributed broadcast encryption and registered ABE schemes is the recent matrix commitment scheme by Wee [Wee25]. Specifically, the work of [Wee25] shows that given the public parameters $pp_{com} = (B, W, T)$ for a succinct LWE instance of dimension $2m^2$ (i.e., $W \in \mathbb{Z}_q^{2m^2n \times m}$ and $[I_{2m^2} \otimes B \mid W]T = I_{2m^2} \otimes G$) and any matrix $M \in \mathbb{Z}_q^{n \times N}$, there is an efficient and deterministic algorithm to compute a commitment $C \in \mathbb{Z}_q^{n \times m}$ and a low-norm opening $Z^{m \times N}$ such that

$$\mathbf{C} \cdot \mathbf{V}_N = \mathbf{M} - \mathbf{B} \cdot \mathbf{Z} \in \mathbb{Z}_a^{n \times N},\tag{2.1}$$

where $\mathbf{V}_N \in \mathbb{Z}_q^{m \times N}$ is a fixed low-norm verification matrix that is publicly derived from pp_{com} and the width *N*. Wee used the matrix commitment scheme to construct key-policy and ciphertext-policy ABE schemes for bounded-depth Boolean circuits from the poly(λ , *d*)-succinct LWE assumption, where *d* denotes the depth of the Boolean circuit.

2.1 Distributed Broadcast Encryption

We now show how to use matrix commitments to construct a distributed broadcast encryption scheme. In distributed broadcast encryption, each public key is associated with an index $i \in \mathbb{N}$, and one can encrypt to any set of public keys, provided that each key has a different index. When $N = 2^{\lambda}$, we can interpret each index as an identity (or a hash of an identity) and we say the scheme supports an arbitrary or unbounded polynomial number of users. The main technical challenges for constructing a distributed broadcast encryption that support an unbounded number of users are twofold:

- First, the size of all of the scheme components (i.e., the public parameters pp, the user public key pk, the user secret key sk, and the ciphertext ct) must all be bounded by poly(λ, log N),
- Second, the running time of key-generation needs to be bounded by $poly(\lambda, \log N)$, and that of encryption and decryption must be bounded by $|S| \cdot poly(\lambda, \log N)$, where *S* is the broadcast set.

As a warm-up, we first describe a scheme with short parameters (i.e., a scheme that addresses the first challenge), but where the running time of key-generation time, encryption, and decryption is *slow* (i.e., scaling with $poly(\lambda, N)$).

- **Public parameters:** The public parameters for the distributed broadcast encryption scheme pp = (pp_{com}, A, p) consists of the public parameters $pp_{com} = (B, W, T)$ for a succinct LWE instance of dimension $2m^2$, a matrix $A \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^{n \times m}$, and a vector $\mathbf{p} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^n$. The vector \mathbf{p} can be viewed as a public key for a dual Regev encryption scheme and the matrix A is used to program the challenge set into the public parameters in the security analysis. Let $\mathbf{V}_N \in \mathbb{Z}_q^{m \times N}$ be the low-norm verification matrix associated with pp_{com} and matrices with width N. Let $\mathbf{V}_N = [\mathbf{v}_1 | \cdots | \mathbf{v}_N]$ where $\mathbf{v}_i \in \mathbb{Z}_q^m$ is the *i*th column of \mathbf{V}_N .
- User key-generation: To generate a public key for a slot $i \in [N]$, the user samples $\mathbf{r}_i \xleftarrow{\mathbb{R}} \{0, 1\}^m$. The public key is $\mathbf{t}_i = \mathbf{Br}_i + \mathbf{p} \mathbf{Av}_i$ and the secret key is \mathbf{r}_i .
- Encryption: Let $S \subseteq [N]$ be a set of indices and let $\{(j, \mathbf{t}_j)\}_{j \in S}$ be a set of public keys. To encrypt a message $\mu \in \{0, 1\}$ to this set of public keys, the encrypter proceeds as follows:

− For each $j \in S$, let $C_j \in \mathbb{Z}_q^{n \times m}$ be the commitment to $\mathbf{u}_j^{\mathsf{T}} \otimes \mathbf{t}_j$ where $\mathbf{u}_j \in \{0, 1\}^N$ is the *j*th standard basis vector. Let $\mathbf{Z}_j \in \mathbb{Z}_q^{m \times N}$ be the respective opening. By Eq. (2.1), this means

$$\mathbf{C}_{i}\mathbf{V}_{N} = (\mathbf{u}_{i}^{\mathsf{T}} \otimes \mathbf{t}_{i}) - \mathbf{B}\mathbf{Z}_{i}.$$

If we parse $\mathbf{Z}_j = [\mathbf{z}_{j,1} | \cdots | \mathbf{z}_{j,N}]$, then this means

$$\mathbf{C}_{j}\mathbf{v}_{j} = \mathbf{t}_{j} - \mathbf{B}\mathbf{z}_{j,j}$$
 and $\forall i \neq j : \mathbf{C}_{j}\mathbf{v}_{i} = -\mathbf{B}\mathbf{z}_{j,i}.$ (2.2)

We often refer to $\mathbf{z}_{j,i}$ for $j \neq i$ as a cross-term since it recodes from **B** to the product $\mathbf{C}_j \mathbf{v}_i$ of user *i*'s public key \mathbf{C}_i with user *j*'s decryption component \mathbf{v}_j .

- The encrypter now samples an LWE secret $\mathbf{s} \leftarrow \mathbb{Z}_{q}^{n}$ and outputs the ciphertext

$$ct = \left(\mathbf{s}^{\mathsf{T}}\mathbf{B}, \mathbf{s}^{\mathsf{T}}(\mathbf{A} + \sum_{j \in S} \mathbf{C}_{j}), \mathbf{s}^{\mathsf{T}}\mathbf{p} + \mu \cdot \lfloor q/2 \rfloor\right)$$

• **Decryption:** Take any user $i \in S$. By Eq. (2.2), we have

$$\mathbf{s}^{\mathsf{T}} \left(\mathbf{A} + \sum_{j \in S} \mathbf{C}_{j} \right) \cdot \mathbf{v}_{i} \approx \mathbf{s}^{\mathsf{T}} \mathbf{A} \mathbf{v}_{i} + \sum_{j \in S} \mathbf{s}^{\mathsf{T}} \mathbf{C}_{j} \mathbf{v}_{i} = \mathbf{s}^{\mathsf{T}} \mathbf{A} \mathbf{v}_{i} + \mathbf{s}^{\mathsf{T}} \mathbf{t}_{i} - \sum_{j \in S} \mathbf{s}^{\mathsf{T}} \mathbf{B} \mathbf{z}_{j,i}$$

Observe now that given the users' public keys $\{(j, \mathbf{t}_j)\}_{j \in S}$, the decrypter can compute $\mathbf{z}_{j,i}$ itself $(\mathbf{z}_{j,i}$ is the *i*th column of the opening \mathbf{Z}_j to the matrix $\mathbf{u}_j^{\mathsf{T}} \otimes \mathbf{t}_j$). Next, using the fact that $\mathbf{t}_i = \mathbf{Br}_i + \mathbf{p} - \mathbf{Av}_i$, and knowledge of the secret key \mathbf{r}_i , the decrypter computes

$$\mathbf{s}^{\mathsf{T}} \Big(\mathbf{A} + \sum_{j \in S} \mathbf{C}_j \Big) \cdot \mathbf{v}_i - \mathbf{s}^{\mathsf{T}} \mathbf{B} \Big(\mathbf{r}_i - \sum_{j \in S} \mathbf{z}_{j,i} \Big) \approx \mathbf{s}^{\mathsf{T}} \mathbf{A} \mathbf{v}_i + \mathbf{s}^{\mathsf{T}} \big(\mathbf{B} \mathbf{r}_i + \mathbf{p} - \mathbf{A} \mathbf{v}_i \big) - \mathbf{s}^{\mathsf{T}} \mathbf{B} \mathbf{r}_i = \mathbf{s}^{\mathsf{T}} \mathbf{p}.$$

Subtracting this from $\mathbf{s}^{\mathsf{T}}\mathbf{p} + b \cdot \lfloor q/2 \rfloor$ and rounding recovers the message μ .

By construction, the size of the public parameters pp, the users' public/secret keys (pk, sk), and the size of the ciphertext ct are poly(λ , log N). However, the running time of key generation, encryption, and decryption is poly(λ , N) because they compute commitments and/or openings to matrices of width N. When N is super-polynomial, this means key-generation, encryption, and decryption are no longer efficient algorithms. In order to support an arbitrary number of users (or alternatively, an identity-based distributed broadcast encryption scheme) with $N = 2^{\lambda}$, we need to reduce the running times of these algorithms to poly(λ , log N). We discuss this below (after sketching the security analysis).

Arguing security. We consider selective security where the adversary declares the set $S \subseteq [N]$ of challenge indices at the beginning of the security game. To prove security from succinct LWE, we need to show how to simulate the challenge ciphertext given $(\mathbf{B}, \mathbf{s}_{\mathcal{D}}^{\mathsf{T}}\mathbf{B}, \mathbf{W}, \mathbf{T})$. We do so by programming the matrix \mathbf{A} and the public keys pk_i for the users $i \in S$. Here, we crucially exploit the fact that the challenge ciphertext does not depend on the public keys outside S (which are chosen adversarially). As an aside, we note that unlike standard broadcast encryption, we do not need to explicitly simulate secret keys for users outside S, since malicious users outside the set S can register any (possibly malformed) public key. The reduction has the following high-level procedure:

- For all $j \in S$, the challenger samples the public key as $\mathbf{t}_i \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^n$. By the leftover hash lemma, the honestlygenerated public keys $\mathbf{t}_j = \mathbf{Br}_j + \mathbf{p} - \mathbf{Av}_j$ are statistically indistinguishable from a uniform random vector $\mathbf{t}_j \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^n$. Importantly, the public keys are now independent of **A**, and thus, the challenger can now sample the public keys \mathbf{t}_j *before* it samples the matrix **A**.
- When setting the public parameters, the challenger embeds the challenge set into the public parameters. Namely, after sampling the public keys $\mathbf{t}_j \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^n$ for the honest users $j \in S$, the challenger now sets $\mathbf{A} = \mathbf{BR}_{\mathbf{A}} \sum_{j \in S} \mathbf{C}_j$, where $\mathbf{R}_{\mathbf{A}} \stackrel{\mathbb{R}}{\leftarrow} \{0, 1\}^{m \times m}$ and \mathbf{C}_j is a commitment to the vector $\mathbf{u}_j^{\mathsf{T}} \otimes \mathbf{t}_j$. It also sets $\mathbf{p} = \mathbf{Br}_{\mathbf{p}}$ where $\mathbf{r}_{\mathbf{p}} \stackrel{\mathbb{R}}{\leftarrow} \{0, 1\}^{m}$. Again by the leftover hash lemma, the distribution of \mathbf{A} and \mathbf{p} are statistically close to uniform, which coincides with their distribution in the real scheme.

• With the above modifications, the challenge ciphertext can now be written as

$$\mathsf{ct} = \left(\mathbf{s}^{\mathsf{T}} \mathbf{B} , \ \mathbf{s}^{\mathsf{T}} \left(\mathbf{A} + \sum_{j \in S} \mathbf{C}_{j} \right) , \ \mathbf{s}^{\mathsf{T}} \mathbf{p} + \mu \cdot \lfloor q/2 \rfloor \right) \approx \left(\mathbf{s}^{\mathsf{T}} \mathbf{B} , \ \mathbf{s}^{\mathsf{T}} \mathbf{B} \cdot \mathbf{R}_{\mathbf{A}} , \ \mathbf{s}^{\mathsf{T}} \mathbf{B} \cdot \mathbf{r}_{\mathbf{p}} + \mu \cdot \lfloor q/2 \rfloor \right),$$

which can be simulated from $\mathbf{s}^{\mathsf{T}}\mathbf{B}$. By the $2m^2$ -succinct LWE assumption, $\mathbf{s}^{\mathsf{T}}\mathbf{B}$ is pseudorandom given $pp_{com} = (\mathbf{B}, \mathbf{W}, \mathbf{T})$. This implies the ciphertext is pseudorandom and security holds.

We refer to Section 4 for the formal description of the construction and security analysis.

Local openings. To support $N = 2^{\lambda}$ (i.e., an arbitrary polynomial number of users), we show how to implement the key-generation, encryption, and decryption algorithms in poly(λ , log N) time. The modification is purely algorithmic and the modifications have no effect on the correctness or security analysis. Our approach relies on two key observations:

- The above scheme only commits to *sparse* matrices (of the form u^T_i ⊗ t_i). While this matrix has N columns, only a single column is non-zero.
- The key-generation, encryption, and decryption algorithms only needs local access to the verification matrix V_N and the openings Z_j. In fact, each algorithm only reads a *single* column of the verification matrix V_L or the opening matrix Z_j.

In Appendix A, we show that the [Wee25] matrix commitment scheme satisfies the following two properties:

- The commitment $\mathbf{C} \in \mathbb{Z}_q^{n \times m}$ to a sparse matrix $\mathbf{M} \in \mathbb{Z}_q^{n \times N}$ with *K* non-zero columns can be computed in time poly(*m*, log *q*, log *N*, *K*).
- There is an algorithm running in time $poly(m, \log q, \log N)$ for computing the *i*th column of the verification matrix **V**_N and an algorithm running in time $poly(m, \log q, \log N, K)$ for computing the *i*th column of the opening **Z**_{*i*} to a matrix **M** with *K* non-zero columns.

In some sense, the matrix commitment scheme from [Wee25] has a Merkle-tree-like structure where a commitment to a matrix $\mathbf{M} = [\mathbf{M}_L \mid \mathbf{M}_R]$ is derived by first committing to its left half \mathbf{M}_L and its right half \mathbf{M}_R , and finally committing to the resulting commitments $[\mathbf{C}_L \mid \mathbf{C}_R]$. We can efficiently commit and provide local openings for (exponentially-long) sparse vectors with a Merkle tree, and the same is true for the matrix commitment scheme of [Wee25]. This yields a distributed broadcast encryption scheme that supports $N = 2^{\lambda}$ users.

Comparison with [CW24]. The structure of our distributed broadcast encryption scheme shares some similarities with that from [CW24]. In our notation, the public keys in [CW24] consists of a (random) matrix C_i together with a collection of low-norm cross-terms $z_{i,j}$ where $Bz_{i,j} = -C_i v_j$. The secret key is a low-norm vector $z_{i,i}$ where $Bz_{i,i} = p + Av_i - C_i v_i$. The public parameters include the matrices A, B, the target vector p, the vectors v_1, \ldots, v_N , along with a (sufficiently-large) succinct LWE trapdoor that is used to sample public keys C_i , cross-terms $z_{i,j}$, and secret keys $z_{i,i}$. From a structural perspective, the public/secret keys between our scheme and the [CW24] scheme are very similar (and likewise for the ciphertexts). The key difference is the following:

- In [CW24], the key-generation algorithm jointly *samples* the public-key matrix C_i together with the cross-terms $z_{i,j}$. This leads to a scheme with O(N)-size public keys. Notably, there is no compact description of the cross-terms needed for decryption (i.e., the terms $z_{j,i}$ that recode from B to $C_j v_i$). Note that publishing the randomness used to sample C_i and $z_{i,j}$ for $j \neq i$ is not sufficient since the randomness would also leak the secret key $z_{i,i}$.
- In our scheme, the public key is simply a vector \mathbf{t}_i , and the associated public-key matrix \mathbf{C}_i and the cross-terms $\mathbf{z}_{i,j}$ are all *deterministically* derived from \mathbf{t}_i using the matrix commitment scheme. The secret key is the randomness used to sample \mathbf{t}_i . Because the matrix \mathbf{C}_i and the cross-terms have a *compact* description, our scheme supports an unbounded number of users. A key technical contribution of this work is showing that the structure of the [Wee25] matrix commitments enables us to compress cross-terms (or alternatively, derive cross-terms from a public procedure).

We note here that the need to include a cross-term for every other user is a standard feature in nearly all constructions of distributed broadcast encryption and registered ABE that do not go through obfuscation or witness encryptions [WQZD10, HLWW23, KMW23, ZZGQ23, AT24, GLWW24, CW24, GKPW24, CHW25]. This is also the main reason these schemes cannot handle an arbitrary number of users. While we still rely on the same type of cross-term cancellation in this work, our use of matrix commitments enables a new compact description for these cross terms. This enables schemes that support an arbitrary number of users.

Comparison with [FWW23]. Our distributed broadcast encryption construction also shares some high-level similarities with the [FWW23] construction based on witness encryption. Both constructions embed a succinct commitment (alternatively, a hash) of the public keys of the users in the broadcast set *S* in the ciphertext, and moreover, in the security analysis, both reductions modify the distribution of the public keys for the users in *S*. In a bit more detail:

- In [FWW23], the commitment to the public keys is the instance used in the witness encryption ciphertext. In their setting, the commitment is a (function-binding) hash of the users' public keys. In our scheme, the commitment to the public keys is Σ_{j∈S} C_j and we embed it as an LWE sample s^T(A + Σ_{j∈S} C_j) in the ciphertext. In both schemes, the commitment has a tree-like structure, and decryption relies on a local opening to the user's public key, which can be derived given just the public keys for the set S (along with the public parameters).
- In order to invoke semantic security of the witness encryption scheme in the [FWW23] security proof, they not only modify the distribution of the public keys in *S*, and also rely on the commitment satisfying a "function binding" property (which can be based on LWE). In contrast, our security proof is simpler and follows by directly programming the set *S* into the public parameters (e.g., setting $\mathbf{A} = -\sum_{j \in S} C_j$, after modifying \mathbf{t}_j , and thus C_j , to be independent of A). This partitioning strategy (for arguing selective security) is a standard approach for analyzing the security of (distributed) broadcast encryption schemes (c.f., [BGW05, KMW23, CW24]). On the flip side, we note that the proof strategy in [FWW23] based on function-binding hash functions shows their scheme to satisfy *semi-static* security; this can in turn be lifted to full adaptive security via the [GW09, HWW25] transformations.

2.2 Key-Policy Registered ABE

We can combine our techniques for distributed broadcast encryption with ideas from the recent work of [CHW25] to obtain a key-policy registered ABE scheme with succinct ciphertexts in the random oracle model. In some sense, the work of [CHW25] starts with the distributed broadcast encryption scheme from [CW24] and shows how to extend it to a registered ABE scheme. In this work, we start from our new distributed broadcast encryption scheme (from Section 2.1) and show how to apply the [CHW25] techniques to lift the scheme to a registered ABE scheme. Since our underlying distributed broadcast encryption scheme supports an unbounded number of users, our registered ABE scheme ABE scheme ABE scheme that can simultaneously support a bounded number of users. Our techniques yield the first registered ABE scheme that can simultaneously support an arbitrary polynomial number of users and which has succinct ciphertexts (see Table 2). Such a scheme was not previously known even from witness encryption or indistinguishability obfuscation. Here, we provide a brief overview of how we augment our distributed broadcast encryption scheme to obtain a (key-policy) registered ABE scheme.

Lattice-based homomorphic evaluation. Our construction relies on the classic homomorphic evaluation machinery from [GSW13, BGG⁺14]. Specifically, given a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times \ell m}$, a Boolean function $f: \{0, 1\}^{\ell} \to \{0, 1\}$, and an input $\mathbf{x} \in \{0, 1\}^{\ell}$, there exists a low-norm matrix $\mathbf{H}_{\mathbf{A}, f, \mathbf{x}}$ such that

$$(\mathbf{A} - \mathbf{x}^{\mathsf{T}} \otimes \mathbf{G}) \cdot \mathbf{H}_{\mathbf{A}, f, \mathbf{x}} = \mathbf{A}_{f} - f(\mathbf{x}) \cdot \mathbf{G},$$
(2.3)

where A_f is a matrix that only depends on A and f.

Registered key-policy ABE. We now describe the general structure of our key-policy registered ABE scheme. Specifically, we describe a "slotted" registered ABE scheme where each key is associated with an index $i \in [N]$, and instead of users joining the system dynamically, there is instead an aggregation algorithm that takes as input a collection of *N* public keys pk_1, \ldots, pk_N along with their associated policies f_1, \ldots, f_N and aggregates them together into a master public key mpk and a set of helper decryption keys hsk_1, \ldots, hsk_N for the *N* users. To decrypt, the user combines their secret key and their helper decryption key with the ciphertext. The work of [HLWW23] show that the slotted primitive generically implies standard the usual notion of registered ABE (that supports dynamic registrations) with only $poly(\lambda, \log N)$ overhead. In this work, we focus exclusively on the simpler slotted primitive.

For ease of exposition, we describe our construction with long ciphertexts (that scale with the attribute length). We can then apply the ciphertext compression approach from [Wee24, Wee25] to obtain a registered ABE scheme with succinct ciphertexts. Essentially, instead of encrypting with respect to $\mathbf{x}^T \otimes \mathbf{G} \in \mathbb{Z}_q^{n \times tm}$ (as in the classic [BGG⁺14] ABE scheme), we instead encrypt to a [Wee25] commitment $C_{\mathbf{x}} \in \mathbb{Z}_q^{n \times tm}$ to $\mathbf{x}^T \otimes \mathbf{G}$. We start with a basic version of the scheme:

- **Public parameters:** The public parameters $pp = (pp_{com}, A, p, \{d_i\}_{i \in [N]})$ have essentially the same structure as that for our distributed broadcast encryption scheme. Here $pp_{com} = (B, W, T)$ are the public parameters for a succinct LWE instance of dimension $2m^2$, $A \notin \mathbb{Z}_q^{n \times \ell m}$ is the matrix used to embed the attribute, $\mathbf{p} \notin \mathbb{Z}_q^m$ is a dual Regev public key (for encoding the message), and $\mathbf{d}_1, \ldots, \mathbf{d}_N \notin \mathbb{Z}_q^n$ are vectors used for noise smudging (in the security analysis). As described, the size of the CRS scales linearly with *N*, but since the \mathbf{d}_i vectors are uniformly random, we can compress them by working in the random oracle model and setting $\mathbf{d}_i = H_1(i)$, where H_1 is modeled as a random oracle.
- User key-generation: To generate a public key for a slot $i \in [N]$ and a function f, the user samples $\mathbf{r}_i \in \{0, 1\}^m$ and sets the secret key to be $\mathbf{t}_i = \mathbf{Br}_i + \mathbf{p} + \mathbf{A}_f \mathbf{G}^{-1}(\mathbf{d}_i) \in \mathbb{Z}_q^n$. The secret key is the randomness \mathbf{r}_i .
- Key aggregation: Given a collection of public keys $\mathbf{t}_1, \ldots, \mathbf{t}_N$ for functions f_1, \ldots, f_N , the aggregation algorithm computes commitments $\mathbf{C}_i \in \mathbb{Z}_q^{n \times m}$ and openings $\mathbf{Z}_i \in \mathbb{Z}_q^{m \times L}$ to $\mathbf{u}_i^{\mathsf{T}} \otimes \mathbf{t}_i$ where $\mathbf{u}_i \in \{0, 1\}^N$ is the *i*th unit vector. It parses $\mathbf{Z}_i = [\mathbf{z}_{i,1} | \cdots | \mathbf{z}_{i,N}]$. The master public key mpk and helper decryption key hsk_i for user *i* are defined to be

$$mpk = \widehat{\mathbf{C}} = \sum_{j \in [N]} \mathbf{C}_j \quad and \quad hsk_i = \widehat{\mathbf{z}}_i = \sum_{j \in [N]} \mathbf{z}_{j,i}.$$
(2.4)

• Encryption: To encrypt a message $\mu \in \{0, 1\}$ with attribute $\mathbf{x} \in \{0, 1\}^{\ell}$, the encryption algorithm samples $\mathbf{s} \leftarrow \mathbb{Z}_{q}^{n}$ and outputs the ciphertext

$$ct = \left(\underbrace{\mathbf{s}^{\mathsf{T}}\mathbf{B}}_{\sim}, \ \underbrace{\mathbf{s}^{\mathsf{T}}\widehat{\mathbf{C}}}_{\sim}, \ \underbrace{\mathbf{s}^{\mathsf{T}}(\mathbf{A} - \mathbf{x}^{\mathsf{T}} \otimes \mathbf{G})}_{\sim}, \ \underbrace{\mathbf{s}^{\mathsf{T}}\mathbf{p} + \lfloor q/2 \rfloor \cdot \mu}_{\sim}\right).$$
(2.5)

• **Decryption:** To decrypt using a secret key \mathbf{r}_i for slot *i* and an associated function f_i where $f_i(\mathbf{x}) = 0$, the decrypter first computes

$$\mathbf{s}^{\mathsf{T}}(\mathbf{A} - \mathbf{x}^{\mathsf{T}} \otimes \mathbf{G}) \cdot \mathbf{H}_{\mathbf{A}, f, \mathbf{x}} \approx \mathbf{s}^{\mathsf{T}}(\mathbf{A}_{f} - f(\mathbf{x}) \cdot \mathbf{G}) = \mathbf{s}^{\mathsf{T}} \mathbf{A}_{f}$$

by Eq. (2.3). Let $\mathbf{V}_N \in \mathbb{Z}_q^{m \times N}$ be the low-norm verification matrix associated with pp_{com} and matrices of width N. Write $\mathbf{V} = [\mathbf{v}_1 | \cdots | \mathbf{v}_N]$. Then, using the fact that \mathbf{C}_i is a commitment to $\mathbf{u}_i^{\mathsf{T}} \otimes \mathbf{t}_i$ and \mathbf{Z}_i is the associated opening, we appeal to Eqs. (2.1) and (2.4) to write

$$\mathbf{s}_{\sim}^{\mathsf{T}} \widehat{\mathbf{C}} \cdot \mathbf{v}_{i} \approx \mathbf{s}^{\mathsf{T}} \sum_{j \in [N]} \mathbf{C}_{j} \mathbf{v}_{i} = \mathbf{s}^{\mathsf{T}} \mathbf{t}_{i} - \mathbf{s}^{\mathsf{T}} \sum_{j \in [N]} \mathbf{B} \mathbf{z}_{j,i} = \mathbf{s}^{\mathsf{T}} \mathbf{t}_{i} - \mathbf{s}^{\mathsf{T}} \mathbf{B} \widehat{\mathbf{z}}.$$

Finally, using the fact that $\mathbf{t}_i = \mathbf{Br}_i + \mathbf{p} + \mathbf{A}_f \mathbf{G}^{-1}(\mathbf{d}_i)$, the decrypter can use its secret key \mathbf{r}_i to compute

$$\begin{split} \mathbf{s}_{\sim}^{\mathsf{T}} & \widehat{\mathbf{C}} \cdot \mathbf{v}_{i} - \mathbf{s}_{i}^{\mathsf{T}} (\mathbf{A} - \mathbf{x}^{\mathsf{T}} \otimes \mathbf{G}) \cdot \mathbf{H}_{\mathbf{A}, f, \mathbf{x}} \cdot \mathbf{G}^{-1}(\mathbf{d}_{i}) - \mathbf{s}_{\sim}^{\mathsf{T}} & \mathbf{B} \cdot (\mathbf{r}_{i} - \hat{\mathbf{z}}) \\ & \approx \mathbf{s}^{\mathsf{T}} (\mathbf{B} \mathbf{r}_{i} + \mathbf{p} + \mathbf{A}_{f} \mathbf{G}^{-1}(\mathbf{d}_{i})) - \mathbf{s}^{\mathsf{T}} \mathbf{B} \hat{\mathbf{z}} - \mathbf{s}^{\mathsf{T}} \mathbf{A}_{f} \mathbf{G}^{-1}(\mathbf{d}_{i}) - \mathbf{s}^{\mathsf{T}} \mathbf{B} \mathbf{r}_{i} + \mathbf{s}^{\mathsf{T}} \mathbf{B} \hat{\mathbf{z}} \\ & = \mathbf{s}^{\mathsf{T}} \mathbf{p}. \end{split}$$

Taking the difference with $\mathbf{s}^{\mathsf{T}}\mathbf{p} + \lfloor q/2 \rfloor \cdot \mu$ and rounding now recovers the message μ .

Arguing security. We consider attribute-selective security where the adversary commits to the challenge attribute **x** at the beginning of the security game. To prove security from the succinct LWE assumption, we again need to show how to simulate the challenge ciphertext (Eq. (2.5)) from $(\mathbf{B}, \mathbf{s}^{\mathsf{T}}\mathbf{B}, \mathbf{W}, \mathbf{T})$. Since the adversary commits to the challenge attribute **x**, we can use the strategy from $[\mathbf{BGG}^+14]$ and program **x** into the matrix **A** (e.g., by setting $\mathbf{A} = \mathbf{BR}_{\mathbf{A}} + \mathbf{x}^{\mathsf{T}} \otimes \mathbf{G}$). Now the reduction can simulate the attribute-embedding component in the challenge ciphertext as

$$\underbrace{s^{\mathsf{T}}(A - x^{\mathsf{T}} \otimes G)}_{(A - x^{\mathsf{T}} \otimes G)} \approx s^{\mathsf{T}} B R_A \approx \underbrace{s^{\mathsf{T}}_{(A - x^{\mathsf{T}} \otimes G)}}_{(A - x^{\mathsf{T}} \otimes G)} \approx s^{\mathsf{T}} B R_A$$

Similarly, the challenger sets $\mathbf{p} = \mathbf{Ar}_{\mathbf{p}}$ for a low-norm \mathbf{p} . Then, it can simulate the message-embedding component of the challenger ciphertext as

$$\mathbf{s}^{\mathsf{T}}\mathbf{p} + \lfloor q/2 \rfloor \cdot \mu \approx \mathbf{s}^{\mathsf{T}}\mathbf{B}\mathbf{r}_{\mathbf{p}} + \lfloor q/2 \rfloor \cdot \mu \approx \mathbf{s}^{\mathsf{T}}\mathbf{B} \cdot \mathbf{r}_{\mathbf{p}} + \lfloor q/2 \rfloor \cdot \mu.$$

The difficult term is simulating $s_{i}^{\mathsf{T}} \widehat{\mathbf{C}}$. Here, $\widehat{\mathbf{C}} = \sum_{i \in [N]} \mathbf{C}_i$, where \mathbf{C}_i is the commitment to the public key of user *i*. In registered ABE, these keys can be chosen *maliciously*. This is a major distinction between registered ABE and distributed broadcast encryption. In distributed broadcast encryption, the keys that influence the challenge ciphertext are honestly generated, so the reduction algorithm can program them so as to be able to simulate the challenge ciphertext. This is not the case in registered ABE. We solve this problem by using the ciphertext re-randomization technique from [CHW25]. The idea is to start with the following *randomized* aggregation algorithm:

- The aggregation algorithm samples $C_0 \leftarrow \mathbb{Z}_q^{n \times m}$ together with low-norm vectors $\mathbf{z}_{0,i} \in \mathbb{Z}_q^m$ where $C_0 \mathbf{v}_i = -\mathbf{B}\mathbf{z}_{0,i}$ for all $i \in [N]$. We can view C as a *random* [Wee25] commitment to the all-zeroes matrix $\mathbf{0}^{n \times N}$ with opening $Z_0 = [\mathbf{z}_{0,1} | \cdots | \mathbf{z}_{0,N}]$. We show that there is an efficient algorithm to sample C_0 together with Z_0 using the succinct LWE trapdoor (see Section 5.1 and Theorem 5.9).
- The aggregated master public key is then mpk = $\widehat{\mathbf{C}} = \mathbf{C}_0 + \sum_{j \in [N]} \mathbf{C}_j$ and the corresponding helper decryption key for user *i* is $hsk_i = \hat{\mathbf{z}}_i = \mathbf{z}_{0,i} + \sum_{j \in [N]} \mathbf{z}_{j,i}$.

It is not hard to see that this modification still preserves correctness of the scheme. Now, in the security proof, the reduction algorithm will set $C_0 = BR_{\widehat{C}} - \sum_{j \in [N]} C_j$, where $R_{\widehat{C}}$ is a low-norm matrix sampled by the reduction. This allows the reduction to simulate the challenge ciphertext component as

$$\mathbf{s}_{\widetilde{C}}^{\mathsf{T}}\widehat{\mathbf{C}} \approx \mathbf{s}^{\mathsf{T}}\left(\mathbf{C}_{0} + \sum_{j \in [N]} \mathbf{C}_{j}\right) = \mathbf{s}^{\mathsf{T}}\mathbf{B}\mathbf{R}_{\widehat{\mathbf{C}}} \approx \mathbf{s}_{\widetilde{C}}^{\mathsf{T}}\mathbf{B} \cdot \mathbf{R}_{\widehat{\mathbf{C}}}.$$

Of course, there is still the caveat that the reduction algorithm now needs to simulate the low-norm vectors $\mathbf{z}_{0,i}$ where $\mathbf{C}_0 \mathbf{v}_i = -\mathbf{B} \mathbf{z}_{0,i}$. To do so, the reduction algorithm will first set $\mathbf{d}_i = \mathbf{Br}_{\mathbf{d}_i}$ for a low-norm $\mathbf{r}_{\mathbf{d}_i}$ known to the reduction. Now, we consider two possibilities depending on whether a key for a slot *i* was honestly-generated (by the reduction algorithm) or chosen maliciously by the adversary:

• Suppose a slot *i* is associated with an honestly-generated key. Using again the fact that for all $j \neq i$, $C_j v_i = -Bz_{j,i}$ and $C_i v_i = t_i - Bz_{i,i}$, we have

$$\mathbf{BR}_{\widehat{\mathbf{C}}}\mathbf{v}_i + \sum_{i \in [N]} \mathbf{B}\mathbf{z}_{j,i} = (\mathbf{C}_0 + \sum_{j \in [N]} \mathbf{C}_j)\mathbf{v}_i + \sum_{j \in [N]} \mathbf{B}\mathbf{z}_{j,i} = \mathbf{C}_0\mathbf{v}_i + \mathbf{t}_i.$$
(2.6)

When generating the keys for the honest users, the reduction algorithm programs the public key to be $\mathbf{t}_i = \mathbf{Br}_i + \mathbf{d}_i = \mathbf{Br}_i + \mathbf{Br}_{\mathbf{d}_i}$. Combined with Eq. (2.6), this means

$$\mathbf{B} \cdot \underbrace{\left(\mathbf{r}_{\mathbf{d}_{i}} + \mathbf{r}_{i} - \mathbf{R}_{\widehat{\mathbf{C}}}\mathbf{v}_{i} - \sum_{j \in [N]} \mathbf{z}_{j,i}\right)}_{\mathbf{z}_{0,i}} = \mathbf{t}_{i} - \mathbf{C}_{0}\mathbf{v}_{i} - \mathbf{t}_{i} = -\mathbf{C}_{0}\mathbf{v}_{i}.$$

• Suppose a slot *i* is associated with an adversarially-chosen key t_i . In this case, the associated policy f_i is not satisfied by the challenge attribute \mathbf{x} (i.e., $f(\mathbf{x}) = 1$). To construct $\mathbf{z}_{0,i}$ in this case, the reduction algorithm will

need to know a short \mathbf{r}_i such that $\mathbf{t}_i = \mathbf{Br}_i + \mathbf{p} + \mathbf{A}_{f_i}\mathbf{G}^{-1}(\mathbf{d}_i)$. As in [CHW25], we facilitate this by requiring each public key include a non-interactive zero-knowledge (NIZK) proof of knowledge of the associated secret key \mathbf{r}_i . The reduction algorithm in turn extracts the associated secret key from each public key chosen by the adversary. In this case, since $f(\mathbf{x}) = 1$,

$$\mathbf{BR}_{\mathbf{A}} \cdot \mathbf{H}_{\mathbf{A}, f_i, \mathbf{x}} \mathbf{G}^{-1}(\mathbf{d}_i) = (\mathbf{A} - \mathbf{x}^{\mathsf{T}} \otimes \mathbf{G}) \cdot \mathbf{H}_{\mathbf{A}, f_i, \mathbf{x}} \mathbf{G}^{-1}(\mathbf{d}_i) = (\mathbf{A}_{f_i} - \mathbf{G}) \cdot \mathbf{G}^{-1}(\mathbf{d}_j) = \mathbf{A}_{f_i} \mathbf{G}^{-1}(\mathbf{d}_i) - \mathbf{Br}_{\mathbf{d}_i}.$$

This means

$$\mathbf{t}_i = \mathbf{B}\mathbf{r}_i + \mathbf{p} + \mathbf{A}_{f_i}\mathbf{G}^{-1}(\mathbf{d}_i) = \mathbf{B}\mathbf{r}_i + \mathbf{B}\mathbf{r}_{\mathbf{p}} + \mathbf{B}\mathbf{R}_{\mathbf{A}}\mathbf{H}_{\mathbf{A},f_i,\mathbf{x}}\mathbf{G}^{-1}(\mathbf{d}_i) + \mathbf{B}\mathbf{r}_{\mathbf{d}_i}.$$

By Eq. (2.6), we have

$$\mathbf{B} \cdot \underbrace{\left(\mathbf{r}_{\mathbf{d}_{i}} + \mathbf{r}_{i} + \mathbf{r}_{p} + \mathbf{R}_{A}\mathbf{H}_{A,f_{i},\mathbf{x}}\mathbf{G}^{-1}(\mathbf{d}_{i}) - \mathbf{R}_{\widehat{\mathbf{C}}}\mathbf{v}_{i} - \sum_{j \in [N]} \mathbf{z}_{j,i}\right)}_{\mathbf{z}_{0,i}} = \mathbf{t}_{i} - \mathbf{C}_{0}\mathbf{v}_{i} - \mathbf{t}_{i} = -\mathbf{C}_{0}\mathbf{v}_{i}$$

In both cases, the reduction algorithm is able to construct a low-norm $\mathbf{z}_{0,i}$ such that $\mathbf{B}\mathbf{z}_{0,i} = -\mathbf{C}_0\mathbf{v}_i$. Thus, using this procedure, the reduction algorithm can program $\mathbf{C}_0 = \mathbf{B}\mathbf{R}_{\widehat{\mathbf{C}}} - \sum_{j \in [N]} \mathbf{C}_j$ and simulate the challenge ciphertext. There are two remaining issues we need to address:

- First, in registered ABE, we require aggregation to be deterministic. Following [CHW25], we can derandomize the aggregation algorithm by having the real scheme sample ($C_0, z_{0,1}, \ldots, z_{0,N}$) using the random oracle (by hashing the input to the aggregation algorithm). To implement this strategy, we require an "explainable" sampling procedure for sampling ($C_0, z_{0,1}, \ldots, z_{0,N}$). Namely, given any (correctly-distributed) tuple ($C_0, z_{0,1}, \ldots, z_{0,N}$), there is an efficient algorithm that outputs a set of random coins to the sampling algorithm that would produce the tuple. In Section 5.1 (Theorem 5.9), we show that the explainable discrete Gaussian preimage sampler from [CHW25] can be used in conjunction with the succinct LWE trapdoor to build such a sampler. To argue that the reduction's strategy for constructing $z_{0,i}$ follows the *same* distribution as that output by the sampler, we use a similar noise smudging argument as in [CHW25]. Specifically, the r_{d_i} component in each $z_{0,i}$ is sampled from a sufficiently-wide (but still low-norm) Gaussian distribution so as to drown out the remaining low-order terms. We give the details of the explainable sampler in Section 5.2 (see the proof of Theorem 5.14).
- Second, to achieve succinct ciphertexts that are independent of the attribute length, we use the compression technique from [Wee24, Wee25]. Specifically, let C_x be a commitment to $x^T \otimes G$, and let Z_x be the associated opening. Let $V_{\ell m}$ be the verification matrix associated with pp_{com} and matrices of length ℓm . Then Eq. (2.1) says

$$C_x V_{\ell m} = x^T \otimes G - BZ_x$$

Sample $\mathbf{B}_0 \leftarrow \mathbb{Z}_q^{n \times m}$ and define $\mathbf{A} = -\mathbf{B}_0 \mathbf{V}_{\ell m}$. The attribute-embedding component $\mathbf{s}^{\mathsf{T}}(\mathbf{A} - \mathbf{x}^{\mathsf{T}} \otimes \mathbf{G})$ in the ciphertext is now replaced by $\mathbf{s}^{\mathsf{T}}(\mathbf{B}_0 + \mathbf{C}_{\mathbf{x}})$. The observation now is that

$$\begin{bmatrix} \mathbf{s}^{\mathsf{T}}\mathbf{B} \mid \mathbf{s}^{\mathsf{T}}(\mathbf{B}_{0} + \mathbf{C}_{\mathbf{x}}) \end{bmatrix} \cdot \begin{bmatrix} -\mathbf{Z}_{\mathbf{x}} \\ -\mathbf{V}_{\ell m} \end{bmatrix} \approx -\mathbf{s}^{\mathsf{T}}\mathbf{B}\mathbf{Z}_{\mathbf{x}} - \mathbf{s}^{\mathsf{T}}\mathbf{B}_{0}\mathbf{V}_{\ell m} - \mathbf{s}^{\mathsf{T}}\mathbf{C}_{\mathbf{x}}\mathbf{V}_{\ell m} = \mathbf{s}^{\mathsf{T}}(\mathbf{A} - \mathbf{x}^{\mathsf{T}}\otimes\mathbf{G}).$$

Thus $\mathbf{s}^{\mathsf{T}}(\mathbf{B}_0 + \mathbf{C}_{\mathbf{x}}) \in \mathbb{Z}_q^m$ can be viewed as a *compressed* representation of $\mathbf{s}^{\mathsf{T}}(\mathbf{A} - \mathbf{x}^{\mathsf{T}} \otimes \mathbf{G})$.

We give the full details in Construction 5.11. Taken together, we obtain a key-policy registered ABE scheme with succinct ciphertexts and which supports an unbounded number of users from the $2m^2$ -succinct LWE assumption.

3 Preliminaries

Throughout this work, we write $\lambda \in \mathbb{N}$ to denote the security parameter. For a positive integer $n \in \mathbb{N}$, we write $[n] := \{1, ..., n\}$. We write poly(λ) to denote a function that is bounded by a fixed polynomial in λ and negl(λ) to denote a function that is negligible in λ (i.e., $f(\lambda) = \text{negl}(\lambda)$ if $f = o(\lambda^{-c})$ for all constants $c \in \mathbb{N}$). For a finite set

S, we write $x \stackrel{\mathbb{R}}{\leftarrow} S$ to denote a uniform random draw from S. For a distribution \mathcal{D} , we write $x \leftarrow \mathcal{D}$ to denote a sample from \mathcal{D} . When \mathcal{A} is a deterministic algorithm, we write $x = \mathcal{A}(\cdot)$ to denote assigning x to the output of \mathcal{A} . We say an algorithm is efficient if it runs in probabilistic polynomial time in the length of its input. We say that two distribution ensembles $\mathcal{D}_0 = {\mathcal{D}_{0,\lambda}}_{\lambda \in \mathbb{N}}$ and $\mathcal{D}_1 = {\mathcal{D}_{1,\lambda}}_{\lambda \in \mathbb{N}}$ are computationally indistinguishable if for all efficient adversaries \mathcal{A} , there exists a negligible function negl(\cdot) such that for all $\lambda \in \mathbb{N}$,

$$|\Pr[\mathcal{A}(1^{\lambda}, x) : x \leftarrow \mathcal{D}_{0,\lambda}] - \Pr[\mathcal{A}(1^{\lambda}, x) : x \leftarrow \mathcal{D}_{1,\lambda}]| = \operatorname{negl}(\lambda).$$

We say that \mathcal{D}_0 and \mathcal{D}_1 are statistically indistinguishable if the statistical distance between $\mathcal{D}_{0,\lambda}$ and $\mathcal{D}_{1,\lambda}$ is bounded by a negligible function negl(λ), and that they are identical if the statistical distance is identically 0.

3.1 Lattice Preliminaries

We use bold uppercase letters (e.g., **B**, **B**) for matrices and bold lowercase letters (e.g., **u**, **v**) for vectors. We use non-boldface letters for their components (e.g., $\mathbf{v} = [v_1, \ldots, v_n]$). For a matrix $\mathbf{V} \in \mathbb{Z}^{n \times n'}$ over the integers, we write $\|\mathbf{V}\|$ to denote the maximum absolute value of its entries. When $\mathbf{V} \in \mathbb{Z}_q^{n \times n'}$, we write $\|\mathbf{V}\| := \|\mathbf{V}_{\mathbb{Z}}\|$, where $\mathbf{V}_{\mathbb{Z}} \in \mathbb{Z}^{n \times n'}$ is the matrix obtained by replacing each component of \mathbf{V} with its integer representative in the interval (-q/2, q/2). We write $\mathbf{A} \otimes \mathbf{B}$ to denote the Kronecker product between matrices \mathbf{A} and \mathbf{B} . For matrices \mathbf{A} , \mathbf{B} , \mathbf{C} , \mathbf{D} with compatible dimensions, we have

$$(\mathbf{A} \otimes \mathbf{C})(\mathbf{B} \otimes \mathbf{D}) = \mathbf{A}\mathbf{B} \otimes \mathbf{C}\mathbf{D}.$$
(3.1)

For a matrix A, we write vec(A) to denote the vector obtained by concatenating the columns of A. We use the following identity:

$$\operatorname{vec}(\mathbf{ABC}) = (\mathbf{C}^{\mathsf{v}} \otimes \mathbf{A}) \cdot \operatorname{vec}(\mathbf{B}). \tag{3.2}$$

Next, we recall the (generalized) leftover hash lemma from [DORS08]. We state the specific formulation from [ABB10]:

Lemma 3.1 (Generalized Leftover Hash Lemma [ABB10, Lemma 13, adapted]). Let n, m, q be integers such that $m \ge 2n \log q$ and q > 2 is prime. Then, for all fixed vectors $\mathbf{e} \in \mathbb{Z}_q^m$ and all k = poly(n), the statistical distance between the following distributions is negl(n):

- (A, AR, $\mathbf{e}^{\mathsf{T}}\mathbf{R}$) where $\mathbf{A} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^{n \times m}, \mathbf{R} \stackrel{\mathbb{R}}{\leftarrow} \{0, 1\}^{m \times k}$.
- $(\mathbf{A}, \mathbf{U}, \mathbf{e}^{\mathsf{T}}\mathbf{R})$ where $\mathbf{A} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^{n \times m}, \mathbf{U} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^{n \times k}, \mathbf{R} \stackrel{\mathbb{R}}{\leftarrow} \{0, 1\}^{m \times k}$.

Discrete Gaussians. We write $D_{\mathbb{Z},\sigma}$ to denote the discrete Gaussian distribution over \mathbb{Z} with width parameter $\sigma > 0$. For a matrix $\mathbf{B} \in \mathbb{Z}_q^{n \times m}$ and a target $\mathbf{Z} \in \mathbb{Z}_q^{n \times n'}$ in the image of \mathbf{B} , we write $\mathbf{B}_{\sigma}^{-1}(\mathbf{Z})$ to denote the random variable $\mathbf{Y} \leftarrow D_{\mathbb{Z},\sigma}^{m \times n'}$ conditioned on $\mathbf{B}\mathbf{Y} = \mathbf{Z} \mod q$. For positive integers $n, q \in \mathbb{N}$ and $m \ge n \lceil \log q \rceil$, we write $\mathbf{G}_n := \mathbf{I}_n \otimes \mathbf{g}^{\mathsf{T}} \in \mathbb{Z}_q^{n \times m}$ to denote the gadget matrix [MP12], where \mathbf{I}_n is the identity matrix of dimension n, $\mathbf{g}^{\mathsf{T}} = [1, 2, \dots, 2^{\lceil \log q \rceil - 1}, 0, \dots, 0] \in \mathbb{Z}^m$. We write $\mathbf{G}_n^{-1}(\cdot) : \mathbb{Z}_q^n \to \mathbb{Z}_q^m$ to denote the standard deterministic entry-wise bit decomposition (and padding with 0s if $m \ge n \cdot \lceil \log q \rceil$). When the context is clear, we omit the subscript n and simply write \mathbf{G} and $\mathbf{G}^{-1}(\cdot)$. We now recall some useful properties on the discrete Gaussian distribution that we will use:

Lemma 3.2 (Gaussian Tail Bound [MP12, Lemma 2.6, adapted]). Let n, m, q be lattice parameters where $m \ge 2n \log q$. For all but a negl(n)-fraction of matrices $\mathbf{B} \in \mathbb{Z}_q^{n \times m}$, all $\sigma \ge \log m$, and all vectors $\mathbf{y} \in \mathbb{Z}_q^n$ in the column space of \mathbf{B} ,

$$\Pr[\|\mathbf{u}\| \ge \sqrt{m}\sigma : \mathbf{u} \leftarrow \mathbf{B}_{\sigma}^{-1}(\mathbf{y})] \le O(2^{-m}).$$

In addition, for all $\lambda \in \mathbb{N}$,

$$\Pr[|x| \ge \sqrt{\lambda\sigma} : x \leftarrow D_{\mathbb{Z},\sigma}] \le 2^{-\lambda}.$$

Lemma 3.3 (Gaussian Preimages [GPV08]). Let n, m, q be lattice parameters where $m \ge 2n \log q$ and q is prime. Let $\sigma \ge \log m$. There exist a negligible function $\operatorname{negl}(\cdot)$ such that for all but a $\operatorname{negl}(n)$ -fraction of matrices $\mathbf{B} \in \mathbb{Z}_q^{n \times m}$, the statistical distance between the following distributions is at most $\operatorname{negl}(n)$:

$$\{(\mathbf{y},\mathbf{B}\mathbf{y}):\mathbf{y}\leftarrow D_{\mathbb{Z}\sigma}^m\}\quad and\quad \{(\mathbf{y},\mathbf{z}):\mathbf{z}\leftarrow^{\mathbb{R}}\mathbb{Z}_q^n,\mathbf{y}\leftarrow\mathbf{B}_{\sigma}^{-1}(\mathbf{z})\}.$$

Moreover, this property holds when $\sigma \ge \log m$ and $\mathbf{B} = \mathbf{G}_n$.

Lemma 3.4 (Marginal of Gaussian Preimages [WW23, Corollary 2.11, adapted]). Let n, m, q be lattice parameters where $m \ge 2n \log q$ and q is prime. Let $\ell, k = \text{poly}(n, \log q)$. There exist a negligible function $\text{negl}(\cdot)$ such that for all but $a q^{-n}$ -fraction of matrices $\mathbf{B} \in \mathbb{Z}_q^{n \times m}$, all matrices $\mathbf{W} \in \mathbb{Z}_q^{n\ell \times k}$ and matrices $\mathbf{C} = [\mathbf{I}_\ell \otimes \mathbf{B} \mid \mathbf{W}]$, all target vectors $\mathbf{y} \in \mathbb{Z}_q^{n\ell}$, and all width parameters $\sigma \ge 4 \log(\ell m)$, the statistical distance between the following distributions is at most negl(n):

$$\left\{\mathbf{v}:\mathbf{v}\leftarrow\mathbf{C}_{\sigma}^{-1}(\mathbf{y})\right\}\quad and\quad \left\{\begin{bmatrix}\mathbf{v}_{1}\\\mathbf{v}_{2}\end{bmatrix}:\quad \frac{\mathbf{v}_{2}\leftarrow D_{\mathbb{Z},\sigma}^{k}}{\mathbf{v}_{1}\leftarrow(\mathbf{I}_{\ell}\otimes\mathbf{B})_{\sigma}^{-1}(\mathbf{y}-\mathbf{B}\mathbf{v}_{2})}\right\}.$$

Gadget trapdoors. We now recall the notion of a gadget trapdoor:

Lemma 3.5 (Gadget Trapdoor [Ajt96, GPV08, MP12]). Let n, m, q be lattice parameters with $m \ge 3n \log q$. There exists efficient algorithms (TrapGen, SamplePre) with the following syntax:

- TrapGen $(1^n, 1^m, q) \rightarrow (\mathbf{B}, \mathbf{T})$: On input the lattice dimension n, the width m, and the modulus q, the trapdoorgeneration algorithm outputs a matrix $\mathbf{B} \in \mathbb{Z}_q^{n \times m}$ together with a trapdoor $\mathbf{T} \in \mathbb{Z}_q^{m \times m}$.
- SamplePre(B, T, Z, σ) \rightarrow Y: On input a matrix $\mathbf{B} \in \mathbb{Z}_q^{n \times m}$, a trapdoor $\mathbf{T} \in \mathbb{Z}_q^{m \times m}$, a target matrix $\mathbf{Z} \in \mathbb{Z}_q^{n \times n'}$, and a width parameter $\sigma > 0$, the preimage-sampling algorithm outputs $\mathbf{Y} \in \mathbb{Z}_q^{m \times n'}$.

Moreover, the above algorithms satisfy the following properties:

- **Trapdoor distribution:** $If(B, T) \leftarrow TrapGen(1^n, 1^m, q)$, then the distribution of **B** is negl(n)-close to the uniform distribution over $\mathbb{Z}_a^{n \times m}$. Moreover, $\mathbf{BT} = \mathbf{G} \in \mathbb{Z}_a^{n \times m}$ and $\|\mathbf{T}\| = 1$.
- **Preimage sampling:** For all matrices $\mathbf{B} \in \mathbb{Z}_q^{n \times m}$, $\mathbf{T} \in \mathbb{Z}_q^{m \times m}$, width parameters $\sigma > 0$, and all $\mathbf{Z} \in \mathbb{Z}_q^{n \times n'}$ in the image of \mathbf{B} , the output $\mathbf{Y} \leftarrow \text{SamplePre}(\mathbf{B}, \mathbf{T}, \mathbf{Z}, \sigma)$ satisfies $\mathbf{B}\mathbf{Y} = \mathbf{Z}$.
- **Preimage distribution:** There exists a negligible function negl(·) such that for all $\mathbf{B} \in \mathbb{Z}_q^{n \times m}$ and $\mathbf{T} \in \mathbb{Z}_q^{m \times m}$ where $\mathbf{BT} = \mathbf{G}$, all $\sigma \ge m \|\mathbf{T}\| \log n$, and all targets $\mathbf{Z} \in \mathbb{Z}_q^{n \times n'}$, the output of SamplePre($\mathbf{B}, \mathbf{T}, \mathbf{Z}, \sigma$) is negl(*n*)-close to the distribution $\mathbf{B}_{\sigma}^{-1}(\mathbf{Z})$.

Homomorphic computation. Our construction of registered ABE will rely on the lattice homomorphic evaluation machinery from [GSW13, BGG⁺14].

Theorem 3.6 (Homomorphic Encodings [GSW13, BGG⁺14]). Let λ be a security parameter and n, m, q be lattice parameters where $m \ge 2n \log q$. Let $\mathcal{F}_{\ell,d}$ be a family of functions $f : \{0,1\}^{\ell} \to \{0,1\}$ that can be computed by a Boolean circuit of depth d. There exists a pair of efficient algorithms (EvalF, EvalFX) with the following properties:

- EvalF(A, f) \rightarrow A_f: On input a matrix A $\in \mathbb{Z}_q^{n \times \ell m}$ and a function $f \in \mathcal{F}_{\ell,d}$, output a matrix A_f $\in \mathbb{Z}_q^{n \times m}$.
- EvalFX(A, f, \mathbf{x}) \rightarrow H_{A, f,\mathbf{x}}: On input a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times \ell m}$, a function $f \in \mathcal{F}_{\ell,d}$, and an input $\mathbf{x} \in \{0, 1\}^{\ell}$, output a matrix $\mathbf{H}_{\mathrm{A},f,\mathbf{x}} \in \mathbb{Z}_q^{\ell m \times m}$.

For all matrices $\mathbf{A} \in \mathbb{Z}_q^{n \times \ell m}$, functions $f \in \mathcal{F}_{\ell,d}$, and inputs $\mathbf{x} \in \{0,1\}^{\ell}$, the matrices $\mathbf{A}_f \leftarrow \text{EvalF}(\mathbf{A}, f)$ and $\mathbf{H}_{\mathbf{A}, f, \mathbf{x}} \leftarrow \text{EvalFX}(\mathbf{A}, f, \mathbf{x})$ satisfy the following properties:

- $(\mathbf{A} \mathbf{x}^{\mathsf{T}} \otimes \mathbf{G}) \cdot \mathbf{H}_{\mathbf{A}, f, \mathbf{x}} = \mathbf{A}_{f} f(\mathbf{x}) \cdot \mathbf{G}.$
- $\|\mathbf{H}_{\mathbf{A},f,\mathbf{x}}\| \leq m^{O(d)}$.

Learning with errors. The learning with errors (LWE) assumption [Reg05] with parameters (n, m, q, χ) states that

$$(\mathbf{B}, \mathbf{s}^{\mathsf{T}}\mathbf{B} + \mathbf{e}^{\mathsf{T}}) \stackrel{c}{\approx} (\mathbf{B}, \mathbf{c}^{\mathsf{T}}),$$

where $\mathbf{B} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^{n \times m}$, $\mathbf{s} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^n$, $\mathbf{e} \leftarrow D_{\mathbb{Z}, \chi}^m$, and $\mathbf{c} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^m$.

 ℓ -succinct LWE. The ℓ -succinct LWE assumption [Wee24] states that $(\mathbf{B}, \mathbf{s}^{\mathsf{T}}\mathbf{B} + \mathbf{e}^{\mathsf{T}})$ is pseudorandom even given a trapdoor for the matrix $[\mathbf{I}_{\ell} \otimes \mathbf{B} \mid \mathbf{W}]$, where $\mathbf{W} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_{q}^{\ell n \times m}$. We give the formal statement below:

Assumption 3.7 (ℓ -Succinct LWE [Wee24]). Let λ be a security parameter and (n, m, q, χ) be LWE parameters. The (ℓ, σ) -succinct LWE assumption states that

$$(\mathbf{B}, \mathbf{s}^{\mathsf{T}}\mathbf{B} + \mathbf{e}^{\mathsf{T}}, \mathbf{W}, \mathbf{T}) \stackrel{c}{\approx} (\mathbf{B}, \mathbf{c}^{\mathsf{T}}, \mathbf{W}, \mathbf{T}),$$

where $\mathbf{B} \leftarrow \mathbb{Z}_q^{n \times m}$, $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, $\mathbf{e} \leftarrow D_{\mathbb{Z},\chi}^m$, $\mathbf{c} \leftarrow \mathbb{Z}_q^m$, $\mathbf{W} \leftarrow \mathbb{Z}_q^{\ell n \times m}$, and $\mathbf{T} \leftarrow [\mathbf{I}_\ell \otimes \mathbf{B} \mid \mathbf{W}]_{\sigma}^{-1}(\mathbf{I}_\ell \otimes \mathbf{G})$. We abbreviate the assumption to ℓ -succinct LWE when $\sigma = \text{poly}(\lambda, \ell, m)$.

3.2 Matrix Commitments

Our constructions rely on the matrix commitments recently introduced in the work of Wee [Wee25]. Specifically, a commitment to a matrix $\mathbf{M} \in \mathbb{Z}_q^{n \times L}$ with respect to $\mathbf{B} \in \mathbb{Z}_q^{n \times m}$ is a matrix $\mathbf{C} \in \mathbb{Z}_q^{n \times m}$ such that

$$\mathbf{C} \cdot \mathbf{V}_L = \mathbf{M} - \mathbf{B} \cdot \mathbf{Z},\tag{3.3}$$

where $\mathbf{V}_L \in \mathbb{Z}_q^{m \times L}$ is a *fixed* verification matrix that only depends on the dimension L and $\mathbf{Z} \in \mathbb{Z}_q^{m \times L}$ is a short opening. The work of [Wee25] shows how to sample a commitment **C** to an arbitrary matrix $\mathbf{M} \in \mathbb{Z}_q^{n \times L}$ (for *any* L) given pp = (**B**, **W**, **T**) where $\mathbf{B} \in \mathbb{Z}_q^{n \times m}$, $\mathbf{W} \in \mathbb{Z}_q^{2m^2n \times m}$, $\mathbf{T} \in \mathbb{Z}_q^{(2m^2+1)m \times 2m^3}$, and $[\mathbf{I}_{2m^2} \otimes \mathbf{B} | \mathbf{W}] \cdot \mathbf{T} = \mathbf{I}_{2m^2} \otimes \mathbf{G}$. In this context, **T** is a trapdoor for a succinct LWE instance with dimension $2m^2$. We now give the formal statement, and for completeness, include the description of the algorithms from [Wee25, §3.2] in Appendix A:

Lemma 3.8 (Matrix Commitment [Wee25, adapted]). Let n, m, q be lattice parameters with $m \ge 2n \log q$. There exist a triple of efficient and deterministic algorithms (Com^{mx}, Ver^{mx}, Open^{mx}) with the following syntax:

- Com^{mx}(pp, M) \rightarrow C: On input public parameters pp and a matrix M $\in \mathbb{Z}_q^{n \times L}$, output a matrix C $\in \mathbb{Z}_q^{n \times m}$.
- Ver^{mx}(pp, 1^L) \rightarrow V_L: On input public parameters pp and the length parameter L, output a matrix V_L $\in \mathbb{Z}_q^{m \times L \lceil \log q \rceil}$.
- Open^{mx}(pp, M) \rightarrow Z: On input public parameters pp and a matrix M $\in \mathbb{Z}_q^{n \times L}$, output a matrix Z $\in \mathbb{Z}_q^{m \times L \lceil \log q \rceil}$.

Moreover, for all pp = (B, W, T) where $B \in \mathbb{Z}_q^{n \times m}$, $W \in \mathbb{Z}_q^{2m^2n \times m}$, $T \in \mathbb{Z}_q^{(2m^2+1)m \times 2m^3}$, $[I_{2m^2} \otimes B | W] \cdot T = I_{2m^2} \otimes G$, all parameters $L \in \mathbb{N}$, all matrices $M \in \mathbb{Z}_q^{n \times L}$, and setting $C = \text{Com}^{mx}(\text{pp}, \mathbf{x})$, $V_L = \text{Ver}^{mx}(\text{pp}, 1^L)$, and $Z = \text{Open}^{mx}(\text{pp}, \mathbf{x})$,

$$\mathbf{C} \cdot \mathbf{V}_L = \mathbf{M} \cdot \mathbf{G}_L - \mathbf{B} \cdot \mathbf{Z} \quad and \quad \|\mathbf{V}_L\| \le O(\|\mathbf{T}\| \cdot m^4 \log q) \quad and \quad \|\mathbf{Z}\| \le O(\|\mathbf{T}\| \cdot m^7 \log q \log L).$$

Remark 3.9 (Committing to M). The commitment relation $(C \cdot V_L = M \cdot G_L - B \cdot Z)$ in Lemma 3.8 does not completely match Eq. (3.3), but this is easy to fix by simply multiplying by $G_L^{-1}(I_L)$. In this case, we have

$$\mathbf{C} \cdot \underbrace{\mathbf{V}_L \cdot \mathbf{G}_L^{-1}(\mathbf{I}_L)}_{\tilde{\mathbf{V}}_L} = \mathbf{M} \cdot \mathbf{G}_L \cdot \mathbf{G}_L^{-1}(\mathbf{I}_L) - \mathbf{B} \cdot \mathbf{Z} \cdot \mathbf{G}^{-1}(\mathbf{I}_L) = \mathbf{M} - \mathbf{B} \cdot \underbrace{\mathbf{Z} \cdot \mathbf{G}_L^{-1}(\mathbf{I}_L)}_{\tilde{\mathbf{Z}}_L}.$$

Since the columns of $G_L^{-1}(I_L)$ have Hamming weight 1, the matrices \tilde{V} and \tilde{Z} are submatrices of V and Z, respectively. This means $\|\tilde{V}\| \leq \|V\|$ and $\|\tilde{Z}\| \leq \|Z\|$. Since Eq. (3.3) is more convenient to work with in our constructions, we define algorithms (Com^{mat}, Ver^{mat}, Open^{mat}) as follows:

- Com^{mat}(pp, M): Output C = Com^{mx}(pp, M) $\in \mathbb{Z}_q^{n \times m}$.
- Ver^{mat}(pp, 1^{*L*}): Compute $\mathbf{V}'_L = \operatorname{Ver}^{mx}(pp, 1^L)$ and output $\mathbf{V}_L = \mathbf{V}'_L \cdot \mathbf{G}_L^{-1}(\mathbf{I}_L) \in \mathbb{Z}_q^{m \times L}$.
- Open^{mat}(pp, M): Compute $\mathbf{Z}'_{L} = \text{Open}^{mx}(pp, M)$ and output $\mathbf{Z} = \mathbf{Z}'_{L} \cdot \mathbf{G}_{L}^{-1}(\mathbf{I}_{L}) \in \mathbb{Z}_{q}^{m \times L}$

Moreover, for all pp = (B, W, T) where $\mathbf{B} \in \mathbb{Z}_q^{n \times m}$, $\mathbf{W} \in \mathbb{Z}_q^{2m^2n \times m}$, $\mathbf{T} \in \mathbb{Z}_q^{(2m^2+1)m \times 2m^3}$, $[\mathbf{I}_{2m^2} \otimes \mathbf{B} \mid \mathbf{W}] \cdot \mathbf{T} = \mathbf{I}_{2m^2} \otimes \mathbf{G}$, all parameters $L \in \mathbb{N}$, all matrices $\mathbf{M} \in \mathbb{Z}_q^{n \times L}$, and setting $\mathbf{C} = \operatorname{Com}^{mx}(\operatorname{pp}, \mathbf{x})$, $\mathbf{V}_L = \operatorname{Ver}^{mx}(\operatorname{pp}, \mathbf{1}^L)$, and $\mathbf{Z} = \operatorname{Open}^{mx}(\operatorname{pp}, \mathbf{x})$,

$$\mathbf{C} \cdot \mathbf{V}_L = \mathbf{M} - \mathbf{B} \cdot \mathbf{Z}$$
 and $\|\mathbf{V}_L\| \le O(\|\mathbf{T}\| \cdot m^4 \log q)$ and $\|\mathbf{Z}\| \le O(\|\mathbf{T}\| \cdot m^7 \log q \log L)$.

Committing to sparse matrices. Our applications to distributed broadcast encryption and registered ABE relies on the ability for users to commit to *sparse* matrices $\mathbf{M} \in \mathbb{Z}_q^{n \times L}$ where *L* is exponential (e.g., $L = 2^{\lambda}$), but where **M** only contains a polynomial number of non-zero columns. This property allows us to construct schemes that support an a priori *unbounded* number of users. When committing to exponentially-wide matrices, the verification matrix \mathbf{V}_L and the opening **Z** are also exponentially wide. Our applications also require an algorithm that provides *local* access to the columns of \mathbf{V}_L and **Z**. Namely, there is an efficient algorithm that runs in time poly(m, log q, log L) which can compute the i^{th} column of the verification matrix \mathbf{V}_L . Similarly, there is an algorithm that runs in time poly(K, m, log q, log L) that computes the i^{th} column of the opening matrix **Z**, where K is the number of non-zero columns in the matrix **M**. It is straightforward to adapt the construction from [Wee25, §3.2] to support these properties. We summarize the properties we need in the following lemma, and provide a formal proof of it in Appendix A.1.

Lemma 3.10 (Committing to Sparse Matrices and Computing Local Openings). There exists a tuple of efficient algorithms (ComSparse^{mat}, VerLocal^{mat}, OpenLocal^{mat}) with the following syntax and properties:

- ComSparse^{mat}(pp, M) \rightarrow C: On input the public parameters pp (with lattice parameters n, m, q) and a matrix $\mathbf{M} \in \mathbb{Z}_q^{n \times L}$, ComSparse^{mat}(pp, M) outputs $\mathbf{C} = \text{Com}^{\text{mat}}(\text{pp}, \mathbf{M})$ in time poly(K, m, log q, log L), where K is the number of non-zero columns in M.
- VerLocal^{mat}(pp, L, i) $\rightarrow \mathbf{v}_{L,i}$: On input the public parameters pp (with lattice parameters n, m, q), the length parameter L (in binary) and a column index $i \in [L]$, VerLocal^{mat}(pp, L, i) outputs the ith column $\mathbf{v}_{L,i} \in \mathbb{Z}_q^m$ of the matrix $\mathbf{V}_L = \text{Ver}^{\text{mat}}(\text{pp}, 1^L)$ in time poly(m, log q, log L).
- OpenLocal^{mat}(pp, M, i) $\rightarrow z_i$: On input the public parameters pp (with lattice parameters n, m, q), a sparse matrix $\mathbf{M} \in \mathbb{Z}_q^{n \times L}$, and an index $i \in [L]$, OpenLocal^{mat}(pp, M, i) outputs the *i*th column $\mathbf{z}_i \in \mathbb{Z}_q^m$ of $\mathbf{Z}_i = \text{Open}^{\text{mat}}(\text{pp}, \mathbf{M})$ in time poly(K, m, log q, log L), where K is the number of non-zero columns in \mathbf{M} .

4 Distributed Broadcast Encryption

In this section, we show how to construct an distributed broadcast encryption [WQZD10, BZ14] scheme from the succinct LWE assumption. In distributed broadcast encryption, users choose their own public and secret keys and post their public keys to a public-key directory. One can then encrypt a message to an arbitrary subset of public keys with a ciphertext whose size scales polylogarithmically with the number of users in the broadcast set. Previous distributed broadcast encryption schemes based on bilinear maps [KMW23, FWW23] or succinct LWE [CW24] could only support a bounded number of users, where the size of the public parameters and individual public keys scale linearly (or worse) with the bound on the number of users. We give the first distributed broadcast encryption scheme from succinct LWE that supports an arbitrary polynomial number of users. Previous schemes that could support an unbounded number of users relied either on indistinguishability obfuscation [BZ14] or witness encryption [FWW23].

Distributed broadcast encryption. We begin by recalling the notion of distributed broadcast encryption. Our definition is adapted from the works of [BZ14, KMW23].¹

Definition 4.1 (Distributed Broadcast Encryption [BZ14, KMW23]). Let λ be a security parameter. A distributed broadcast encryption scheme Π_{DBE} is a tuple of efficient algorithms Π_{DBE} = (Setup, KeyGen, Encrypt, Decrypt) with the following syntax:

- Setup $(1^{\lambda}, N) \rightarrow pp$: On input the security parameter λ and the number of users N (in *binary*), the setup algorithm outputs the public parameters pp. We assume that pp contains 1^{λ} and N.
- KeyGen(pp, i) \rightarrow (pk_i, sk_i): On input the public parameters pp and an index $i \in [N]$, the key-generation algorithm outputs a public key pk_i and secret key sk_i.

¹Some previous lattice-based schemes [CW24, CHW25] also required an IsValid algorithm that is used to (publicly) decide whether a user public key is well-formed or not, and only required correctness/security to hold when encrypting to well-formed keys. Our distributed broadcast encryption construction (Construction 4.2) will not require this property, so for ease of exposition, we omit this algorithm from the formal syntax. The syntax we use matches that from [BZ14, KMW23].

- Encrypt(pp, $\{(i, pk_i)\}_{i \in S}, \mu$) \rightarrow ct: On input the public parameters pp, a collection of public keys pk_i and a message $\mu \in \{0, 1\}$, the encryption algorithm outputs a ciphertext ct.
- Decrypt(pp, $\{(i, pk_i)\}_{i \in S}$, ct, (i^*, sk_{i^*})) $\rightarrow \mu$: On input the public parameters pp, a collection of public keys pk_i , a ciphertext ct, and a secret key sk_{i^*} for an index i^* , the decryption algorithm outputs a message $\mu \in \{0, 1\}$.

We require that Π_{DBE} satisfy the following properties:

• **Correctness:** For all parameters $\lambda, N \in \mathbb{N}$, all pp in the support of Setup $(1^{\lambda}, N)$, all sets $S \subseteq [N]$, all indices $i^* \in S$, all public keys pk_i for $i \in S \setminus \{i^*\}$, and all messages $\mu \in \{0, 1\}$,

$$\Pr\left| \operatorname{Decrypt}(\operatorname{pp}, \{(i, \operatorname{pk}_i)\}_{i \in S}, \operatorname{ct}, (i^*, \operatorname{sk}_{i^*})) = \mu : \begin{array}{c} (\operatorname{pk}_{i^*}, \operatorname{sk}_{i^*}) \leftarrow \operatorname{KeyGen}(\operatorname{pp}, i^*) \\ \operatorname{ct} \leftarrow \operatorname{Encrypt}(\operatorname{pp}, \{(i, \operatorname{pk}_i)\}_{i \in S}, \mu) \end{array} \right| = 1$$

- Selective security: For a security parameter λ , a bound N, and a bit $b \in \{0, 1\}$, we define the selective security game between an adversary \mathcal{A} and a challenger as follows:
 - On input the security parameter 1^{λ} and the bound N, the adversary outputs the challenge set $S^* \subseteq [N]$.
 - The challenger samples $pp \leftarrow \text{Setup}(1^{\lambda}, N)$ and $(pk_i, sk_i) \leftarrow \text{KeyGen}(pp, i)$ for $i \in S^*$. Finally, it computes $ct_b \leftarrow \text{Encrypt}(pp, \{pk_i\}_{i \in S^*}, b, S^*)$ and sends $(pp, \{pk_i\}_{i \in S^*}, ct_b)$ to \mathcal{A} .
 - At the end of the game, algorithm \mathcal{A} outputs a bit $b' \in \{0, 1\}$, which is the output of the experiment.

We say the distributed broadcast encryption scheme is selectively secure if for all efficient adversaries \mathcal{A} and all bounds $N \leq 2^{\lambda}$, there exists a negligible function negl(·) such that for all $\lambda \in \mathbb{N}$,

$$|\Pr[b' = 1 | b = 1] - \Pr[b' = 1 | b = 0]| = \operatorname{negl}(\lambda)$$

in the selective security game.

• Succinctness: There exists a fixed polynomial $poly(\cdot, \cdot)$ such that for all $\lambda, N \in \mathbb{N}$, all subsets $S \subseteq [N]$, all public parameters pp in the support of $Setup(1^{\lambda}, N)$, all public keys pk_i , all messages $\mu \in \{0, 1\}$, and all ciphertexts ct in the support of $Encrypt(pp, \{pk_i\}_{i \in S}, \mu, S)$, it holds that $|ct| \leq poly(\lambda, \log N)$.

Construction 4.2 (Distributed Broadcast Encryption). Let λ be a security parameter and (n, m, q, χ) be LWE parameters (that can be functions of λ and N). Let σ be a Gaussian width parameter. We construct a distributed broadcast encryption scheme as follows:

• Setup $(1^{\lambda}, N)$: On input the security parameter λ and the number of users N, sample

$$(\mathbf{B}, \mathbf{T}_{\mathbf{B}}) \leftarrow \mathsf{TrapGen}(1^{n}, 1^{m}, q), \mathbf{W} \overset{\mathbb{R}}{\leftarrow} \mathbb{Z}_{q}^{2m^{2}n \times m}$$
$$\mathbf{T} \leftarrow \mathsf{SamplePre}([\mathbf{I}_{2m^{2}} \otimes \mathbf{B} \mid \mathbf{W}], \begin{bmatrix} \mathbf{I}_{2m^{2}} \otimes \mathbf{T}_{\mathbf{B}} \\ \mathbf{0} \end{bmatrix}, \mathbf{I}_{2m^{2}} \otimes \mathbf{G}, \sigma)$$
$$\mathbf{A} \overset{\mathbb{R}}{\leftarrow} \mathbb{Z}_{q}^{n \times m}, \mathbf{p} \overset{\mathbb{R}}{\leftarrow} \mathbb{Z}_{q}^{n}.$$

If $||\mathbf{T}|| > \sqrt{m\sigma}$, then set $\mathbf{T} = \begin{bmatrix} \mathbf{I}_{2m_0^2 \otimes \mathbf{T}_B} \\ \mathbf{0} \end{bmatrix}$. Let $pp_{com} = (\mathbf{B}, \mathbf{W}, \mathbf{T})$ and output the public parameters $pp = (N, pp_{com}, \mathbf{A}, \mathbf{p})$.

• KeyGen(pp, *i*): On input the public parameters $pp = (N, pp_{com}, A, p)$ and an index $i \in [N]$, sample

$$\mathbf{r}_i \leftarrow \{0, 1\}^m$$

 $\mathbf{v}_i = \text{VerLocal}^{\text{mat}}(\text{pp}_{\text{com}}, N, i)$

Output the public key $\mathbf{t}_i = \mathbf{Br}_i + \mathbf{p} - \mathbf{Av}_i \in \mathbb{Z}_q^n$ and the secret key sk = \mathbf{r}_i .

• Encrypt(pp, $\{(i, pk_i)\}_{i \in S}, \mu$): On input the public parameters $pp = (N, pp_{com}, A, p)$ where $pp_{com} = (B, W, T)$, the public keys $pk_i = t_i \in \mathbb{Z}_q^{n \times m}$ and a message $\mu \in \{0, 1\}$, sample the following:

$$\mathbf{s} \stackrel{\mathsf{R}}{\leftarrow} \mathbb{Z}_q^n$$
, $\mathbf{e} \leftarrow D_{\mathbb{Z},\chi}^m$, $\mathbf{D}_1 \stackrel{\mathsf{R}}{\leftarrow} \{0,1\}^{m \times m}$, $\mathbf{d}_2 \stackrel{\mathsf{R}}{\leftarrow} \{0,1\}^m$.

If $\|\mathbf{e}\| > \sqrt{m\chi}$, then set $\mathbf{e} = \mathbf{0}^m$. Then, for all $i \in S$, compute $C_i = \text{ComSparse}^{\text{mat}}(\text{pp}_{\text{com}}, \mathbf{u}_i^{\mathsf{T}} \otimes \mathbf{t}_i)$, where $\mathbf{u}_i \in \{0, 1\}^N$ is the *i*th unit vector. Output the ciphertext

$$ct = \left(\mathbf{s}^{\mathsf{T}}\mathbf{B} + \mathbf{e}^{\mathsf{T}}, \ \mathbf{s}^{\mathsf{T}}\left(\mathbf{A} + \sum_{i \in S} \mathbf{C}_{i}\right) + \mathbf{e}^{\mathsf{T}}\mathbf{D}_{1}, \ \mathbf{s}^{\mathsf{T}}\mathbf{p} + \mathbf{e}^{\mathsf{T}}\mathbf{d}_{2} + \lfloor q/2 \rfloor \cdot \mu\right)$$

• Decrypt(pp, { (i, pk_i) }_{$i \in S$}, ct, (i^*, sk_{i^*})): On input the public parameters pp = (N, pp_{com}, A, p) where pp_{com} = (B, W, T), the public keys $pk_i = t_i \in \mathbb{Z}_q^{n \times m}$, the ciphertext ct = (c_1, c_2, c_3) , and a secret key $sk_{i^*} = r_{i^*}$ for an index $i^* \in S$, the decryption algorithm first computes

$$\mathbf{v}_{i^*} = \mathsf{VerLocal}^{\mathsf{mat}}(\mathsf{pp}_{\mathsf{com}}, N, i^*)$$

$$\forall i \in S : \mathbf{z}_{i,i^*} = \mathsf{OpenLocal}^{\mathsf{mat}}(\mathsf{pp}_{\mathsf{com}}, \mathbf{u}_i^{\mathsf{T}} \otimes \mathbf{t}_i, i^*)$$

Finally, compute

$$\tilde{\mu} = c_3 + \mathbf{c}_1^{\mathsf{T}} \mathbf{r}_{i^*} - \mathbf{c}_2^{\mathsf{T}} \mathbf{v}_{i^*} - \sum_{i \in S} \mathbf{c}_1^{\mathsf{T}} \mathbf{z}_{i,i^*}$$

and output 0 if $-q/4 < \xi < q/4$ and 1 otherwise.

Theorem 4.3 (Correctness). Suppose $q > N \cdot O(m^9 \chi \sigma \log q \log N)$. Then, Construction 4.2 is correct.

Proof. Take any $\lambda, N \in \mathbb{N}$, any pp in the support of Setup $(1^{\lambda}, N)$, any set $S \subseteq [N]$, any index $i^* \in S$, any collection of public keys $\{pk_i\}_{i \in S \setminus \{i^*\}}$, and any message $\mu \in \{0, 1\}$. Let $(pk_{i^*}, sk_{i^*}) \leftarrow \text{KeyGen}(pp, i^*)$ and ct $\leftarrow \text{Encrypt}(pp, \{(i, pk_i)\}_{i \in S}, \mu)$. Consider Decrypt $(pp, \{(i, pk_i)\}_{i \in S}, ct, (i^*, sk_{i^*}))$:

- Let $pp = (N, pp_{com}, A, p)$ where $pp_{com} = (B, W, T)$. By construction of Setup, $[I_{2m^2} \otimes B | W]T = I_{2m^2} \otimes G$ and $||T|| \le \sqrt{m\sigma}$.
- By definition, $\mathsf{pk}_{i^*} = \mathbf{t}_{i^*} = \mathbf{Br}_{i^*} + \mathbf{p} \mathbf{Av}_{i^*}$ where $\mathbf{r}_{i^*} \in \{0, 1\}^m$ and $\mathbf{v}_{i^*} = \mathsf{VerLocal}^{\mathsf{mat}}(\mathsf{pp}_{\mathsf{com}}, N, i^*)$.
- Let s, e, \mathbf{D}_1 , \mathbf{d}_2 be the components sampled by Encrypt. By definition, $\|\mathbf{e}\| \le \sqrt{m\chi}$ and $\|\mathbf{D}_1\|$, $\|\mathbf{d}_2\| \le 1$. The ciphertext can be written as

$$\mathsf{ct} = (\mathbf{c}_1^\mathsf{T}, \mathbf{c}_2^\mathsf{T}, \mathbf{c}_3) = \left(\mathbf{s}^\mathsf{T}\mathbf{B} + \mathbf{e}^\mathsf{T}, \ \mathbf{s}^\mathsf{T}\left(\mathbf{A} + \sum_{i \in S} \mathbf{C}_i\right) + \mathbf{e}^\mathsf{T}\mathbf{D}_1, \ \mathbf{s}^\mathsf{T}\mathbf{p} + \mathbf{e}^\mathsf{T}\mathbf{d}_2 + \lfloor q/2 \rfloor \cdot \mu\right)$$

• Let \mathbf{v}_{i^*} and \mathbf{z}_{i,i^*} be the vectors computed by Decrypt. For each $i \in S$, let $\mathbf{C}_i = \text{ComSparse}^{\text{mat}}(\text{pp}_{\text{com}}, \mathbf{u}_i^{\mathsf{T}} \otimes \mathbf{t}_i)$. By Remark 3.9 and Lemma 3.10, we have for all $i \in S$

$$\mathbf{C}_{i}\mathbf{v}_{i^{*}} = \begin{cases} \mathbf{t}_{i^{*}} - \mathbf{B}\mathbf{z}_{i^{*},i^{*}} & i = i^{*} \\ -\mathbf{B}\mathbf{z}_{i,i^{*}} & i \neq i^{*}. \end{cases}$$

Thus, we can now write

$$\mathbf{c}_{2}^{\mathsf{T}}\mathbf{v}_{i^{*}} + \sum_{i \in S} \mathbf{c}_{1}^{\mathsf{T}}\mathbf{z}_{i,i^{*}} = \mathbf{s}^{\mathsf{T}}\mathbf{A}\mathbf{v}_{i^{*}} + \sum_{i \in S} \mathbf{s}^{\mathsf{T}}\mathbf{C}_{i}\mathbf{v}_{i^{*}} + \sum_{i \in S} \mathbf{s}^{\mathsf{T}}\mathbf{B}\mathbf{z}_{i,i^{*}} + \mathbf{e}^{\mathsf{T}}\mathbf{D}_{1}\mathbf{v}_{i^{*}} + \sum_{i \in S} \mathbf{e}^{\mathsf{T}}\mathbf{z}_{i,i^{*}}$$
$$= \mathbf{s}^{\mathsf{T}}(\mathbf{A}\mathbf{v}_{i^{*}} + \mathbf{t}_{i^{*}}) + \tilde{e}_{1}$$
$$= \mathbf{s}^{\mathsf{T}}(\mathbf{p} + \mathbf{B}\mathbf{r}_{i^{*}}) + \tilde{e}_{1}.$$

where $\tilde{e}_1 = \mathbf{e}^{\mathsf{T}} \mathbf{D}_1 \mathbf{v}_{i^*} + \sum_{i \in S} \mathbf{e}^{\mathsf{T}} \mathbf{z}_{i,i^*}$. Then

$$\begin{split} \tilde{\mu} &= c_3 + \mathbf{c}_1^{\mathsf{T}} \mathbf{r}_{i^*} - \mathbf{c}_2^{\mathsf{T}} \mathbf{v}_{i^*} - \sum_{i \in S} \mathbf{c}_1^{\mathsf{T}} \mathbf{z}_{i,i^*} \\ &= \mathbf{s}^{\mathsf{T}} \mathbf{p} + \mathbf{e}^{\mathsf{T}} \mathbf{d}_2 + \lfloor q/2 \rfloor \cdot \mu + \mathbf{s}^{\mathsf{T}} \mathbf{B} \mathbf{r}_{i^*} + \mathbf{e}^{\mathsf{T}} \mathbf{r}_{i^*} - \mathbf{s}^{\mathsf{T}} (\mathbf{p} + \mathbf{B} \mathbf{r}_{i^*}) - \tilde{e}_1 \\ &= \lfloor q/2 \rfloor \cdot \mu + \mathbf{e}^{\mathsf{T}} \mathbf{d}_2 + \mathbf{e}^{\mathsf{T}} \mathbf{r}_{i^*} - \tilde{e}_1, \end{split}$$

Let $\tilde{e} = \mathbf{e}^{\mathsf{T}} \mathbf{d}_2 + \mathbf{e}^{\mathsf{T}} \mathbf{r}_{i^*} - \tilde{e}_1$. As long as $|\tilde{e}| < q/4$, correctness holds.

• We now bound $|\tilde{e}|$. First, for all $i \in S$, by Remark 3.9 and Lemma 3.10,

$$\begin{aligned} \|\mathbf{v}_{i^*}\| &\leq O(\|\mathbf{T}\| \cdot m^4 \log q) \leq O(m^{9/2} \sigma \log q) \\ \|\mathbf{z}_{i,i^*}\| &\leq O(\|\mathbf{T}\| \cdot m^7 \log q \log N) \leq O(m^{15/2} \sigma \log q \log N). \end{aligned}$$

First, we bound $|\tilde{e}_1|$. Since $||\mathbf{e}|| \le \sqrt{m}\chi$, we have

$$\begin{split} |\tilde{e}_1| &\leq \left\| \mathbf{e}^{\mathsf{T}} \mathbf{D}_1 \mathbf{v}_{i^*} \right\| + \sum_{i \in S} \left\| \mathbf{e}^{\mathsf{T}} \mathbf{z}_{i,i^*} \right\| \leq O(m^7 \chi \sigma \log q) + N \cdot O(m^9 \chi \sigma \log q \log N) \\ &= N \cdot O(m^9 \chi \sigma \log q \log N). \end{split}$$

Next,

$$|\tilde{e}| \leq |\mathbf{e}^{\mathsf{T}}\mathbf{d}_2| + |\mathbf{e}^{\mathsf{T}}\mathbf{r}_{i^*}| + |\tilde{e}_1| \leq N \cdot O(m^9 \chi \sigma \log q \log N).$$

Setting $q > N \cdot O(m^9 \chi \sigma \log q \log N)$ suffices for correctness.

Theorem 4.4 (Selective Security). Suppose $n \ge \lambda$, $m \ge 3n \log q$, and $\sigma \ge O(m^3 \log m)$. Then, under the $(2m^2, \sigma)$ -succinct LWE assumption with parameters (n, m, q, χ) , Construction 4.2 is selectively secure.

Proof. Take any $N \leq 2^{\lambda}$ and any efficient adversary \mathcal{A} for the selective security game. We now define a sequence of hybrid experiments:

- $Hyb_0^{(b)}$: This is the selective security experiment where the challenger encrypts the bit $b \in \{0, 1\}$. Specifically, the game proceeds as follows:
 - On input the security parameter 1^{λ} and the bound *N*, algorithm \mathcal{A} outputs the challenge set $S^* \subseteq [N]$.
 - The challenger then samples $pp \leftarrow \text{Setup}(1^{\lambda}, N)$ by computing

$$(\mathbf{B}, \mathbf{T}_{\mathbf{B}}) \leftarrow \mathsf{TrapGen}(1^{n}, 1^{m}, q), \mathbf{W} \leftarrow \mathbb{Z}_{q}^{2m^{2}n \times m}$$
$$\mathbf{T} \leftarrow \mathsf{SamplePre}([\mathbf{I}_{2m^{2}} \otimes \mathbf{B} \mid \mathbf{W}], \begin{bmatrix} \mathbf{I}_{2m^{2}} \otimes \mathbf{T}_{\mathbf{B}} \\ \mathbf{0} \end{bmatrix}, \mathbf{I}_{2m^{2}} \otimes \mathbf{G}, \sigma)$$
$$\mathbf{A} \leftarrow \mathbb{Z}_{q}^{n \times m}, \mathbf{p} \leftarrow \mathbb{Z}_{q}^{n}.$$

If $||T|| > \sqrt{m\sigma}$, then it sets $T = \begin{bmatrix} I_{2m^2 \otimes T_B} \\ 0 \end{bmatrix}$. The challenger sets $pp_{com} = (B, W, T)$ and $pp = (N, pp_{com}, A, p)$.

- Next, for each $i \in S^*$, the challenger samples a public key by computing $pk_i \leftarrow KeyGen(pp, i)$. Specifically, for each $i \in S^*$, the challenger samples

$$\mathbf{r}_{i} \stackrel{\text{\tiny R}}{\leftarrow} \{0, 1\}^{m}$$

$$\mathbf{v}_{i} = \text{VerLocal}^{\text{mat}}(\text{pp}_{\text{com}}, N, i)$$

and sets $\mathsf{pk}_i = \mathbf{t}_i = \mathbf{Br}_i + \mathbf{p} - \mathbf{Av}_i \in \mathbb{Z}_q^{n \times m}$.

- Finally, the challenger constructs the challenge ciphertext $ct^* \leftarrow Encrypt(pp, \{(i, pk_i)\}_{i \in [N]}, b)$. Concretely, the challenger starts by sampling

$$\mathbf{s} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^n$$
, $\mathbf{e} \leftarrow D_{\mathbb{Z},\chi}^m$, $\mathbf{D}_1 \stackrel{\mathbb{R}}{\leftarrow} \{0,1\}^{m \times m}$, $\mathbf{d}_2 \stackrel{\mathbb{R}}{\leftarrow} \{0,1\}^m$.

If $\|\mathbf{e}\| > \sqrt{m\chi}$, then the challenger sets $\mathbf{e} = \mathbf{0}^m$. Then, for all $i \in S$, the challenger computes $\mathbf{C}_i = \text{ComSparse}^{\text{mat}}(\text{pp}_{\text{com}}, \mathbf{u}_i^{\mathsf{T}} \otimes \mathbf{t}_i)$. It defines the challenge ciphertext ct* to be

$$\mathbf{ct}^* = \left(\mathbf{s}^\mathsf{T}\mathbf{B} + \mathbf{e}^\mathsf{T}, \ \mathbf{s}^\mathsf{T}\left(\mathbf{A} + \sum_{i \in S^*} \mathbf{C}_i\right) + \mathbf{e}^\mathsf{T}\mathbf{D}_1, \ \mathbf{s}^\mathsf{T}\mathbf{p} + \mathbf{e}^\mathsf{T}\mathbf{d}_2 + \lfloor q/2 \rfloor \cdot b\right).$$

- The challenger gives $(pp, \{(i, pk_i)\}_{i \in S^*}, ct^*)$ to \mathcal{A} .

- At the end of the game, algorithm \mathcal{A} outputs a bit $b' \in \{0, 1\}$, which is the output of the experiment.

• $Hyb_1^{(b)}$: Same as $Hyb_0^{(b)}$, except when sampling the public parameters, the challenger samples

 $\mathbf{T} \xleftarrow{\mathbb{R}} [\mathbf{I}_{2m^2} \otimes \mathbf{B} \mid \mathbf{W}]_{\sigma}^{-1}(\mathbf{I}_{2m^2} \otimes \mathbf{G}) \quad \text{and} \quad \mathbf{B} \xleftarrow{\mathbb{R}} \mathbb{Z}_{q}^{n \times m}.$

- $Hyb_2^{(b)}$: Same as $Hyb_1^{(b)}$, except when sampling the public parameters, the challenger no longer checks the condition $||\mathbf{T}|| > \sqrt{m\sigma}$. Similarly, when constructing the challenge ciphertext ct^{*}, the challenger no longer checks if $||\mathbf{e}|| > \sqrt{m\chi}$.
- Hyb₃^(b): Same as Hyb₂^(b), except the challenger samples $\mathbf{t}_i \leftarrow \mathbb{Z}_q^{n \times m}$ for all $i \in S^*$.
- $Hyb_4^{(b)}$: Same as $Hyb_3^{(b)}$, except the challenger now sets $\mathbf{A} = \mathbf{BD}_1 \sum_{i \in S^*} \mathbf{C}_i$ and $\mathbf{p} = \mathbf{Bd}_2$. Specifically, this experiment operates as follows:
 - On input the security parameter 1^{λ} , algorithm \mathcal{A} outputs the bound N and a challenge set $S^* \subseteq [N]$.
 - The challenger samples

$$\mathbf{B} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^{n \times m}, \mathbf{W} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^{2m^2n \times m}$$
$$\mathbf{T} \leftarrow [\mathbf{I}_{2m^2} \otimes \mathbf{B} \mid \mathbf{W}]_{\sigma}^{-1}(\mathbf{I}_{2m^2} \otimes \mathbf{G})$$
$$\mathbf{t}_i \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^n \text{ for all } i \in S^*$$
$$\mathbf{D}_1 \stackrel{\mathbb{R}}{\leftarrow} \{0, 1\}^{m \times m}, \mathbf{d}_2 \stackrel{\mathbb{R}}{\leftarrow} \{0, 1\}^m$$

The challenger sets $pp_{com} = (B, W, T)$. For each $i \in S^*$, the challenger sets $pk_i = t_i$ and also computes $C_i = ComSparse^{mat}(pp_{com}, \mathbf{u}_i^T \otimes t_i)$. The challenger sets $\mathbf{A} = \mathbf{BD}_1 - \sum_{i \in S^*} C_i$, $\mathbf{p} = \mathbf{Bd}_2$, and $pp = (N, pp_{com}, \mathbf{A}, \mathbf{p})$.

- To construct the challenge ciphertext, the challenger samples $\mathbf{s} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^n$ and $\mathbf{e} \leftarrow D_{\mathbb{Z},\chi}^m$. It then defines the challenge ciphertext ct^* to be

$$ct^* = \left(\mathbf{s}^{\mathsf{T}}\mathbf{B} + \mathbf{e}^{\mathsf{T}}, \ \mathbf{s}^{\mathsf{T}}\left(\mathbf{A} + \sum_{i \in S^*} \mathbf{C}_i\right) + \mathbf{e}^{\mathsf{T}}\mathbf{D}_1, \ \mathbf{s}^{\mathsf{T}}\mathbf{p} + \mathbf{e}^{\mathsf{T}}\mathbf{d}_2 + \lfloor q/2 \rfloor \cdot b\right)$$
$$= \left(\mathbf{s}^{\mathsf{T}}\mathbf{B} + \mathbf{e}^{\mathsf{T}}, \ (\mathbf{s}^{\mathsf{T}}\mathbf{B} + \mathbf{e}^{\mathsf{T}})\mathbf{D}_1, \ (\mathbf{s}^{\mathsf{T}}\mathbf{B} + \mathbf{e}^{\mathsf{T}})\mathbf{d}_2 + \lfloor q/2 \rfloor \cdot b\right).$$

- The challenger gives (pp, $\{(i, pk_i)\}_{i \in S^*}$, ct^{*}) to \mathcal{A} .
- At the end of the game, algorithm \mathcal{A} outputs a bit $b' \in \{0, 1\}$, which is the output of the experiment.
- $Hyb_5^{(b)}$: Same as $Hyb_4^{(b)}$, except the challenger samples $\mathbf{c}_1 \leftarrow \mathbb{Z}_q^m$ and defines the challenge ciphertext to be

$$\mathbf{ct}^* = (\mathbf{c}_1^\mathsf{T}, \mathbf{c}_1^\mathsf{T}\mathbf{D}_1, \mathbf{c}_1^\mathsf{T}\mathbf{d}_2 + \lfloor q/2 \rfloor \cdot b).$$

• $Hyb_6^{(b)}$: Same as $Hyb_5^{(b)}$, except the challenger samples $\mathbf{c}_2 \leftarrow \mathbb{Z}_q^m$ and $c_3 \leftarrow \mathbb{Z}_q$. The challenger defines the challenge ciphertext to be $ct^* = (\mathbf{c}_1^T, \mathbf{c}_2^T, c_3)$. Notably, in this experiment, the adversary's view is independent of the bit *b*.

We write $\mathsf{Hyb}_i^{(b)}(\mathcal{A})$ to denote the distribution of the output of $\mathsf{Hyb}_i^{(b)}$ with adversary \mathcal{A} . We now analyze each pair of adjacent distributions.

Lemma 4.5. If $n \ge \lambda$, $m \ge 3n \log q$, and $\sigma \ge O(m^3 \log m)$, then for all $b \in \{0, 1\}$, $\mathsf{Hyb}_0^{(b)}$ and $\mathsf{Hyb}_1^{(b)}$ are statistically indistinguishable.

Proof. By Lemma 3.5, $||\mathbf{T}_{\mathbf{B}}|| = 1$. As long as $\sigma \ge (2m^2 + 1)m ||\mathbf{T}_{\mathbf{B}}|| \log(2m^2n) = O(m^3 \log m)$, these two distributions are statistically indistinguishable by Lemma 3.5.

Lemma 4.6. If $n \ge \lambda$, $m \ge 2n \log q$, q is prime, and $\sigma \ge O(\log m)$, then for all $b \in \{0, 1\}$, $\mathsf{Hyb}_1^{(b)}$ and $\mathsf{Hyb}_2^{(b)}$ are statistically indistinguishable.

Proof. By Lemmas 3.2 and 3.4, with overwhelming probability over the choice of $\mathbf{B} \notin \mathbb{Z}_q^{n \times m}$, we have that $\|\mathbf{T}\| \leq \sqrt{m\sigma}$. By Lemma 3.4, we have that $\|\mathbf{e}\| \leq \sqrt{m\chi}$ with probability at least $1 - m \cdot O(2^{-m})$. Thus, the conditions the challenger checks in $\mathsf{Hyb}_1^{(b)}$ hold with negligible probability so the two experiments are statistically indistinguishable.

Lemma 4.7. If $n \ge \lambda$, $m \ge 2n \log q$, and q > 2 is prime, then for all $b \in \{0, 1\}$, $\mathsf{Hyb}_2^{(b)}$ and $\mathsf{Hyb}_3^{(b)}$ are statistically indistinguishable by

Proof. Specifically, by Lemma 3.1, the distributions (**B**, **Br**) and (**B**, **t**) where $\mathbf{B} \leftarrow \mathbb{Z}_q^{n \times m}$, $\mathbf{r} \leftarrow \{0, 1\}^{m \times m}$, and $\mathbf{t} \leftarrow \mathbb{Z}_q^{n \times m}$ are statistically indistinguishable. Since $|S^*| = \text{poly}(\lambda)$, the claim now follows by a standard hybrid argument.

Lemma 4.8. If $n \ge \lambda$, $m \ge 2n \log q$, and q > 2 is prime, then for all $b \in \{0, 1\}$, $\mathsf{Hyb}_3^{(b)}$ and $\mathsf{Hyb}_4^{(b)}$ are statistically indistinguishable.

Proof. Specifically, invoking Lemma 3.1 with the error vector $\mathbf{e} \in \mathbb{Z}_q^m$ the challenger samples in $\text{Hyb}_3^{(b)}$ and $\text{Hyb}_4^{(b)}$, we can conclude that the distributions $(\mathbf{B}, \mathbf{BD}_1, \mathbf{e}^{\mathsf{T}}\mathbf{D}_1)$ and $(\mathbf{B}, \mathbf{A}^*, \mathbf{e}_1^{\mathsf{T}}\mathbf{D}_1)$ are statistically indistinguishable when $\mathbf{B} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^{n \times m}$, $\mathbf{D}_1 \stackrel{\mathbb{R}}{\leftarrow} \{0, 1\}^{m \times m}$, $\mathbf{A}^* \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^{n \times m}$. In $\text{Hyb}_3^{(b)}$ and $\text{Hyb}_4^{(b)}$, write $\mathbf{A} = \mathbf{A}^* - \sum_{i \in S^*} \mathbf{C}_i$. When $\mathbf{A}^* = \mathbf{BD}_1$, we obtain the distribution of \mathbf{A} in $\text{Hyb}_3^{(b)}$ and when $\mathbf{A}^* \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^{n \times m}$, we obtain the distribution of \mathbf{A} in $\text{Hyb}_4^{(b)}$ (since \mathbf{A}^* is sampled independently of all other quantities in the experiment). Similarly, the distributions $(\mathbf{B}, \mathbf{Bd}_2, \mathbf{e}^{\mathsf{T}}\mathbf{d}_2)$ and $(\mathbf{B}, \mathbf{p}, \mathbf{e}^{\mathsf{T}}\mathbf{d}_2)$ are statistically indistinguishable when $\mathbf{B} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^{n \times m}$, $\mathbf{d}_2 \stackrel{\mathbb{R}}{\leftarrow} \{0, 1\}^m$, $\mathbf{p} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^n$ and the claim holds.

Lemma 4.9. Under the $(2m^2, \sigma)$ -succinct LWE assumption with lattice parameters (n, m, q, χ) , for all $b \in \{0, 1\}$, $\mathsf{Hyb}_4^{(b)}$ and $\mathsf{Hyb}_5^{(b)}$ are computationally indistinguishable.

Proof. Suppose $|\Pr[Hyb_4^{(b)}(\mathcal{A}) = 1] - \Pr[Hyb_5^{(b)}(\mathcal{A}) = 1]| \ge \varepsilon$ for some non-negligible ε . We use \mathcal{A} to construct an adversary \mathcal{B} for the $(2m^2)$ -succinct LWE assumption:

- 1. On input the succinct LWE challenge (**B**, **c**₁, **W**, **T**), algorithm \mathcal{B} starts running algorithm \mathcal{A} which outputs the bound $N = N(\lambda)$ and the challenge set $S^* \subseteq [N]$.
- 2. Algorithm \mathcal{B} sets $pp_{com} = (\mathbf{B}, \mathbf{W}, \mathbf{T})$ and samples $\mathbf{D}_1 \stackrel{\mathbb{R}}{\leftarrow} \{0, 1\}^{m \times m}$ and $\mathbf{d}_2 \stackrel{\mathbb{R}}{\leftarrow} \{0, 1\}^m$. Next, for each $i \in S^*$, algorithm \mathcal{B} samples $pk_i = \mathbf{P}_i \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^{n \times m}$ and computes $\mathbf{C}_i = \text{ComSparse}^{\text{mat}}(pp_{com}, \mathbf{u}_i^{\mathsf{T}} \otimes \mathbf{P}_i)$. Finally, algorithm \mathcal{B} sets $\mathbf{A} = \mathbf{B}\mathbf{D}_1 \sum_{i \in S^*} \mathbf{C}_i$, $\mathbf{p} = \mathbf{B}\mathbf{d}_2$, and $pp = (N, pp_{com}, \mathbf{A}, \mathbf{p})$.
- 3. Finally, algorithm \mathcal{B} constructs the challenger ciphertext as $ct^* = (c_1^T, c_1^T D_1, ct_1^T d_2 + \lfloor q/2 \rfloor \cdot b)$.
- 4. Algorithm \mathcal{B} gives (pp, $\{(i, pk_i)\}_{i \in S^*}$, ct^{*}) to \mathcal{A} and outputs whatever \mathcal{A} outputs.

By definition, the challenger samples $\mathbf{B} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^{n \times m}$, $\mathbf{W} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^{2m^2n \times m}$, and $\mathbf{T} \leftarrow [\mathbf{I}_{2m^2} \otimes \mathbf{B} \mid \mathbf{W}]_{\sigma}^{-1}(\mathbf{I}_{2m^2} \otimes \mathbf{G})$. Thus, algorithm \mathcal{B} perfectly simulates pp according to the distribution of $\mathsf{Hyb}_4^{(b)}$ and $\mathsf{Hyb}_5^{(b)}$. The remaining components in the public parameters and the public keys are sampled exactly as in $\mathsf{Hyb}_4^{(b)}$ and $\mathsf{Hyb}_5^{(b)}$. It suffices to consider the challenger ciphertext:

• If the challenger sets $\mathbf{c}_1^{\mathsf{T}} = \mathbf{s}^{\mathsf{T}} \mathbf{B} + \mathbf{e}^{\mathsf{T}}$ where $\mathbf{s} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^n$ and $\mathbf{e} \stackrel{\mathbb{R}}{\leftarrow} D_{\mathbb{Z},\chi}^m$, then algorithm \mathcal{B} perfectly simulates the distribution of $\mathsf{Hyb}_4^{(b)}$.

• If the challenger samples $\mathbf{c}_1 \xleftarrow{\mathbb{R}} \mathbb{Z}_q^m$, then algorithm \mathcal{B} perfectly simulates the distribution of $\mathsf{Hyb}_5^{(b)}$.

Thus, algorithm $\mathcal B$ breaks succinct LWE security with the same advantage ε .

Lemma 4.10. If $n \ge \lambda$, $m \ge 2n \log q$, and q > 2 is prime, then for all $b \in \{0, 1\}$, $\mathsf{Hyb}_5^{(b)}$ and $\mathsf{Hyb}_6^{(b)}$ are statistically indistinguishable.

Proof. Applying Lemma 3.1 with respect to the matrix $\begin{bmatrix} \mathbf{B} \\ \mathbf{c}_1^T \end{bmatrix}$, the following distributions are statistically indistinguishable:

$$(\mathbf{B}, \mathbf{c}_1, \mathbf{B}\mathbf{D}_1, \mathbf{B}\mathbf{d}_2, \mathbf{c}_1^{\mathsf{T}}\mathbf{D}_1, \mathbf{c}_1^{\mathsf{T}}\mathbf{d}_2)$$
 and $(\mathbf{B}, \mathbf{c}_1, \mathbf{A}^*, \mathbf{p}, \mathbf{c}_2, \mathbf{c}_3)$,

when $\mathbf{B} \leftarrow \mathbb{Z}_q^{n \times m}$, $\mathbf{c}_1 \leftarrow \mathbb{Z}_q^m$, $\mathbf{D}_1 \leftarrow \{0, 1\}^{m \times m}$, $\mathbf{d}_2 \leftarrow \{0, 1\}^m$, $\mathbf{A}^* \leftarrow \mathbb{Z}_q^{n \times m}$, $\mathbf{p} \leftarrow \mathbb{Z}_q^n$, $\mathbf{c}_2 \leftarrow \mathbb{Z}_q^m$, and $\mathbf{c}_3 \leftarrow \mathbb{Z}_q$. Applying Lemma 3.1 again with respect to the matrix \mathbf{B} , we conclude that the following two distributions are also statistically indistinguishable:

 $(B, c_1, A^*, p, c_2, c_3)$ and $(B, c_1, BD_1, Bd_2, c_2, c_3)$,

By a hybrid argument, we conclude that

$$(\mathbf{B}, \mathbf{c}_1, \mathbf{B}\mathbf{D}_1, \mathbf{B}\mathbf{d}_2, \mathbf{c}_1^{\mathsf{T}}\mathbf{D}_1, \mathbf{c}_1^{\mathsf{T}}\mathbf{d}_2)$$
 and $(\mathbf{B}, \mathbf{c}_1, \mathbf{B}\mathbf{D}_1, \mathbf{B}\mathbf{d}_2, \mathbf{c}_2, \mathbf{c}_3)$

are statistically indistinguishable. The left distribution maps to $Hyb_5^{(b)}$ while the right distribution maps to $Hyb_6^{(b)}$.

By construction, the challenger's behavior in $Hyb_6^{(b)}$ is independent of the bit $b \in \{0,1\}$. Thus, $Hyb_6^{(0)}(\mathcal{A}) \equiv Hyb_6^{(1)}(\mathcal{A})$. The claim now follows by combining Lemmas 4.5 to 4.10.

Theorem 4.11 (Succinctness). Suppose $m \log q = poly(n, \log N)$. Then Construction 4.2 is succinct.

Proof. A ciphertext in Construction 4.2 can be written as $(\mathbf{c}_1, \mathbf{c}_2, c_3)$ where $\mathbf{c}_1, \mathbf{c}_2 \in \mathbb{Z}_q^m$ and $c_3 \in \mathbb{Z}_q$. Thus, the total size of a ciphertext is $O(m \log q) = \text{poly}(n, \log N)$, as required.

Parameter instantiation. We now describe one instantiation of the parameters in Construction 4.2 to satisfy Theorems 4.3, 4.4, and 4.11. Let λ be the security parameter and N be the maximum number of users. Take any constant $0 < \varepsilon < 1$. We instantiate the LWE parameters (n, m, q, χ) and the width parameter σ where

$$n = \lambda \log^{1/\epsilon} N$$

$$m = n \cdot \operatorname{poly}(\lambda, \log N)$$

$$\chi = \operatorname{poly}(\lambda, \log N)$$

$$\sigma = O(m^3 \log m)$$

$$q = N \cdot \operatorname{poly}(\lambda, \log N)$$

and the $(2m^2, \sigma)$ -succinct LWE assumption holds with LWE parameters (n, m, q, χ) . This corresponds to a succinct LWE instance where the modulus-to-noise ratio is $N \cdot \text{poly}(\lambda, \log N) = 2^{\tilde{O}(n^{\varepsilon})}$. For this choice of parameters, we obtain a distributed broadcast encryption scheme with the following properties:

$$|pp| = poly(\lambda, \log N)$$
, $|pk| = poly(\lambda, \log N)$, $|sk| = poly(\lambda, \log N)$, $|ct| = poly(\lambda, \log N)$

In particular, the size of the public parameters, the size of individual public/secret keys, and the size of the ciphertext are optimal. Moreover, in the case where $N = \text{poly}(\lambda)$, then $q = \text{poly}(\lambda)$. In this case, we can rely on succinct LWE with a *polynomial* modulus-to-noise-ratio. We summarize our instantiation with the following corollary:

Corollary 4.12 (Distributed Broadcast Encryption). Let λ be a security parameter and N be the number of users. Under the poly(λ , log N)-succinct LWE assumption with a sub-exponential modulus-to-noise ratio, there exists a distributed broadcast encryption scheme supporting up to N users where the size of the public parameters, the size of individual public/secret keys, and the size of the ciphertext is poly(λ , log N). Moreover, when N = poly(λ), then security can be based on the poly(λ , log N)-succinct LWE assumption with a polynomial modulus-to-noise ratio.

Remark 4.13 (Transparent Setup via Decomposed LWE). Our distributed broadcast encryption scheme (Construction 4.2) makes black-box use of the [Wee25] matrix commitment scheme. Moreover, the cryptographic assumption in the security proof (Theorem 4.4) can equivalently be restated as asking that

$$(pp_{com}, \mathbf{s}^{\mathsf{T}}\mathbf{B} + \mathbf{e}^{\mathsf{T}}) \stackrel{\sim}{\approx} (pp_{com}, \mathbf{c}^{\mathsf{T}}),$$

where $pp_{com} = (\mathbf{B}, \mathbf{W}, \mathbf{T})$ is the public parameters for the matrix commitment scheme, $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, $\mathbf{e} \leftarrow \mathbb{D}_{\mathbb{Z},\chi}^m$, and $\mathbf{c} \leftarrow \mathbb{Z}_q^m$. In [Wee25, Appendix C], Wee shows an alternative way to realize a matrix commitment scheme from the decomposed LWE assumption introduced in [AMR25]. Specifically, the decomposed LWE assumption (with dimension ℓ and width parameter σ) along with LWE parameters (n, m, q, χ) asserts that

$$(\mathbf{B}, \mathbf{s}^{\mathsf{T}}(\mathbf{W}(\mathbf{I}_{\ell} \otimes \mathbf{R}) + \operatorname{vec}(\mathbf{I}_{\ell})^{\mathsf{T}} \otimes \mathbf{G}) + \mathbf{e}^{\mathsf{T}}, \mathbf{W}, \mathbf{R}) \stackrel{c}{\approx} (\mathbf{B}, \mathbf{c}^{\mathsf{T}}, \mathbf{W}, \mathbf{R}),$$

where $\mathbf{B} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^{n \times m}$, $\mathbf{s} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^n$, $\mathbf{W} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^{n \times \ell m}$, $\mathbf{R} \leftarrow D_{\mathbb{Z},\sigma}^{m \times \ell m}$, $\mathbf{e} \leftarrow D_{\mathbb{Z},\chi}^{\ell^2 m}$, and $\mathbf{c} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^{\ell^2 m}$. We then shows that an analogous matrix commitment scheme (satisfying the properties in Lemma 3.8) where the public parameters are given by $pp_{com} = (\hat{\mathbf{B}}, \mathbf{W}, \mathbf{R})$, where $\mathbf{W} \in \mathbb{Z}_q^{n \times 2m^3}$, $\mathbf{R} \in D_{\mathbb{Z},\sigma}^{m \times 2m^3}$, and

$$\hat{\mathbf{B}} := \mathbf{W}(\mathbf{I}_{2m^2 \otimes \mathbf{R}} + \operatorname{vec}(\mathbf{I}_{2m^2})^{\mathsf{T}} \otimes \mathbf{G}) \in \mathbb{Z}_q^{n \times 4m^5}.$$
(4.1)

Then, the decomposed LWE assumption (with dimension $2m^2$ and width parameter σ) asserts that

$$(pp_{com}, \mathbf{s}^{T}\hat{\mathbf{B}} + \mathbf{e}^{T}) \stackrel{\sim}{\approx} (pp_{com}, \mathbf{c}^{T}),$$

where $pp_{com} = (\hat{\mathbf{B}}, \mathbf{W}, \mathbf{R}), \mathbf{W} \leftarrow \mathbb{Z}_q^{n \times 2m^3}, \mathbf{R} \leftarrow D_{\mathbb{Z},\sigma}^{m \times 2m^3}, \hat{\mathbf{B}} \in \mathbb{Z}_q^{n \times 4m^5}$ as in Eq. (4.1), $\mathbf{s} \leftarrow \mathbb{Z}_q^n, \mathbf{e} \leftarrow D_{\mathbb{Z},\chi}^{4m^5}$, and $\mathbf{c} \leftarrow \mathbb{Z}_q^{4m^5}$. As such, we can substitute this alternative instantiation for the matrix commitment in Construction 4.2 to obtain a distributed broadcast encryption scheme from the decomposed LWE assumption with dimension $2m^2$ and width parameter σ . Compared to the basic instantiation from Construction 4.2, this has the following advantages and disadvantages:

- **Transparent setup:** The public parameters p_{com} for the matrix commitment scheme in the above instantiation can be described by $\mathbf{W} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^{n \times 2m^3}$ and $\mathbf{R} \leftarrow D_{\mathbb{Z},\sigma}^{m \times 2m^3}$. These can be derived from a uniform random string. The remaining components of the public key in Construction 4.2 is a uniform random matrix **A** and a uniform random vector **p**. Thus, this yields a construction with a *transparent* setup process. The construction based on succinct LWE requires including the succinct LWE trapdoor **T** as part of the commitment parameters, which results in the need for a structured string.
- Larger parameters: A downside of using decomposed LWE instead of succinct LWE is the matrix $\hat{\mathbf{B}} \in \mathbb{Z}_q^{n \times 4m^5}$ now has width $4m^5$ (compared to $\mathbf{B} \in \mathbb{Z}_q^{n \times m}$ in Construction 4.2). Thus, the ciphertexts are longer in the modified construction. Note though that this is a polynomial blowup in the parameter size (since $m = \text{poly}(\lambda, \log N)$).

5 Key-Policy Registered Attribute-Based Encryption

In this section, we show how to construct a registered key-policy ABE for general circuits from the succinct LWE assumption in the random oracle model. Our construction combines ideas from the recent registered ABE construction from [CHW25] with our distributed broadcast encryption scheme from Section 4.

Slotted registered ABE. Similar to prior works on registered ABE [HLWW23, FWW23, ZZGQ23, GLWW24, AT24, CHW25], we focus on the simpler "slotted" primitive introduced in [HLWW23]. In a slotted registered ABE scheme, there is a *fixed* number of users *N* and each user is associated with a specific slot index. Instead of users joining the system at arbitrary times, there is instead an aggregation algorithm that takes all *N* keys (together with their associated policies) and aggregates them into a single succinct master public key. The work of [HLWW23] show that a slotted registered ABE scheme implies the normal notion that supports dynamic registrations with only polylogarithmic overhead (using a standard powers-of-two transformation implicit in earlier works such as [GHMR18]). In this work, we focus exclusively on the slotted primitive because it is simpler to describe and construct.

Definition 5.1 (Slotted Key-Policy Registered ABE [HLWW23, adapted]). Let λ be a security parameter and τ be a policy-family parameter. Let $\mathcal{X} = {\mathcal{X}_{\tau}}_{\tau \in \mathbb{N}}$ be a family of attributes and $\mathcal{F} = {\mathcal{F}_{\tau}}_{\tau \in \mathbb{N}}$ be a set of policies (where each $f \in \mathcal{F}_{\tau}$ is a function $f : \mathcal{F}_{\tau} \to {0, 1}$). Throughout this work, we adopt the convention that an attribute $x \in \mathcal{X}_{\tau}$ satisfies a policy $f \in \mathcal{F}_{\tau}$ if f(x) = 0 (and does not satisfy the policy if f(x) = 1). A slotted key-policy registered ABE scheme with attribute space \mathcal{X} and policy space \mathcal{F} is a tuple of efficient algorithms $\Pi_{sRABE} = (Setup, KeyGen, IsValid, Aggregate, Encrypt, Decrypt)$ with the following syntax:

- Setup $(1^{\lambda}, 1^{\tau}, N) \rightarrow pp$: On input the security parameter λ , the policy-family parameter τ , and the number of slots *N* (in *binary*), the setup algorithm outputs the public parameters pp. We assume that pp implicitly defines $1^{\lambda}, 1^{\tau}$, and *N*.
- KeyGen(pp, $i, f) \rightarrow (pk, sk)$: On input the public parameters pp, a slot index $i \in [N]$ and a function $f \in \mathcal{F}_{\tau}$, the key-generation algorithm outputs a public key pk and a secret key sk.
- IsValid(pp, i, f, pk) $\rightarrow b$: On input the public parameters pp, a slot index $i \in [N]$, a function $f \in \mathcal{F}_{\tau}$, and a public key pk, the validity-checking algorithm outputs a bit $b \in \{0, 1\}$.
- Aggregate(pp, $(pk_1, f_1), \ldots, (pk_N, f_N)$) \rightarrow (mpk, hsk₁, ..., hsk_N): On input the common reference string pp, a collection of N public keys pk_1, \ldots, pk_N , and their respective functions $f_1, \ldots, f_N \in \mathcal{F}_{\tau}$, the aggregation algorithm outputs a master public key mpk and a collection of helper decryption keys hsk₁, ..., hsk_N. This algorithm is deterministic.
- Encrypt(mpk, x, μ) \rightarrow ct: On input the master public key mpk, an attribute $x \in X_{\tau}$, and a message $\mu \in \{0, 1\}$, the encryption algorithm outputs a ciphertext ct.
- Decrypt(sk, hsk, f, x, ct) $\rightarrow \mu$: On input the secret key sk, a helper decryption key hsk, a policy $f \in \mathcal{F}_{\tau}$, an attribute $x \in X_{\tau}$, and a ciphertext ct, the decryption algorithm outputs a message $\mu \in \{0, 1\}$.

Moreover, Π_{sRABE} should satisfy the following properties:

• **Completeness:** For all $\lambda, \tau \in \mathbb{N}$, all $N \leq 2^{\lambda}$, all indices $i \in [N]$, and all policies $f \in \mathcal{F}_{\tau}$,

$$\Pr\left[\text{IsValid}(\text{pp}, i, f, \text{pk}) = 1: \begin{array}{c} \text{pp} \leftarrow \text{Setup}(1^{\lambda}, 1^{\tau}, N) \\ (\text{pk}, \text{sk}) \leftarrow \text{KeyGen}(\text{pp}, i, f) \end{array}\right] = 1$$

• **Correctness:** We say Π_{sRABE} is correct if for all λ, τ , all $N \leq 2^{\lambda}$, all indices $i \in [N]$, all policies $f_i \in \mathcal{F}_{\tau}$, all attributes $x \in X_{\tau}$ where $f_i(x) = 0$, all pp in the support of Setup $(1^{\lambda}, 1^{\tau}, N)$, all $\{(j, f_j, pk_j)\}_{j \neq i}$ where IsValid(pp, $j, f_j, pk_j) = 1$, and all messages $\mu \in \{0, 1\}$,

 $\Pr\left[\mathsf{Decrypt}(\mathsf{sk}_i,\mathsf{hsk}_i,f,x,\mathsf{ct}) = \mu: \ \mathsf{ct} \leftarrow \mathsf{Encrypt}(\mathsf{mpk},x,\mu)\right] = 1,$

where $(mpk, hsk_1, ..., hsk_N) = Aggregate(pp, pk_1, ..., pk_N)$.

• **Compactness:** There exists a polynomial poly (\cdot, \cdot) such that for all $\lambda, \tau \in \mathbb{N}$, all $N \leq 2^{\lambda}$, all pp in the support of Setup $(1^{\lambda}, 1^{\tau}, N)$, all $(\mathsf{pk}_1, f_1), \ldots, (\mathsf{pk}_N, f_N)$ where $\mathsf{IsValid}(\mathsf{pp}, i, f_i, \mathsf{pk}_i) = 1$ for all $i \in [N]$, and all $(\mathsf{mpk}, \mathsf{hsk}_1, \ldots, \mathsf{hsk}_N)$ in the support of Aggregate $(\mathsf{pp}, (\mathsf{pk}_1, f_1), \ldots, (\mathsf{pk}_N, f_N))$, it holds that

$$|\mathsf{mpk}| \le \mathsf{poly}(\lambda, \tau)$$
 and $\forall i \in [N] : |\mathsf{hsk}_i| \le \mathsf{poly}(\lambda, \tau)$

- Security: Let $b \in \{0, 1\}$ be a bit. For a policy-family parameter $\tau \in \mathbb{N}$, a slot number N, and an adversary \mathcal{A} , we define the following security game between \mathcal{A} and a challenger:
 - Setup phase: The challenger begins by sampling $pp \leftarrow \text{Setup}(1^{\lambda}, 1^{\tau}, N)$ and gives pp to \mathcal{A} . The challenger also initializes a counter ctr = 0 and an (empty) dictionary D to keep track of key-generation queries.
 - Query phase: Adversary \mathcal{A} can now issue the following queries:
 - * **Key-generation query:** In a key-generation query, the adversary specifies a slot index $i \in [N]$ and a function $f \in \mathcal{F}_{\tau}$. The challenger increments the counter ctr = ctr + 1 and samples (pk_{ctr}, sk_{ctr}) \leftarrow KeyGen(pp, *i*, *f*). The challenger replies to \mathcal{A} with (ctr, pk_{ctr}) to \mathcal{A} and then adds the mapping ctr \mapsto (*i*, *f*, pk_{ctr}, sk_{ctr}) to D.
 - * **Corruption query:** In a corruption query, the adversary specifies an index $1 \le c \le \text{ctr.}$ In response, the challenger looks up the tuple (i', f', pk', sk') = D[c] and replies to \mathcal{A} with sk'.
 - **Challenge phase:** Algorithm \mathcal{A} now outputs a list of tuples $(c_1, f_1, pk_1), \ldots, (c_N, f_N, pk_N)$ where either $c_i \in \{1, \ldots, ctr\}$ to reference a challenger-generated key or $c_i = \bot$ to reference a key outside this set. The adversary also specifies a challenge attribute $x^* \in X_\tau$. The challenger then checks the following:
 - * If $c_i \in \{1, ..., ctr\}$, then the challenger looks up the entry $D[c_i] = (i', f', pk', sk')$ and checks that i = i'. If not, the challenger outputs 0. Otherwise, the challenger sets $pk_i^* = pk'$ and $f_i^* = f'$. In addition, if the adversary issued a "corruption" query on index c_i , then the challenger additionally checks that $f'(x^*) = 1$ and outputs 0 if not.
 - * If $c_i = \bot$, then the challenger checks that $IsValid(pp, i, f_i, pk_i) = 1$ and $f_i(x^*) = 1$. If so, then the challenger sets $pk_i^* = pk_i$ and $f_i^* = f_i$. Otherwise, the challenger outputs 0.

The challenger computes $(mpk, hsk_1, ..., hsk_N) = \text{Aggregate}(pp, (pk_1^*, f_1^*) ..., (pk_N^*, f_N^*))$ and replies with the challenge ciphertext ct^{*} \leftarrow Encrypt (mpk, x^*, b) .

- **Output phase:** At the end of the experiment, algorithm \mathcal{A} outputs a bit $b' \in \{0, 1\}$, which is the output of the experiment. If \mathcal{A} aborts before this point, then the output of the experiment is 0.

We say that a slotted registered ABE scheme is secure if for all polynomials $\tau = \tau(\lambda)$ and $N = N(\lambda)$, and all efficient adversaries \mathcal{A} , there exists a negligible function negl(·) such that for all $\lambda \in \mathbb{N}$,

$$|\Pr[b' = 1 : b = 0] - \Pr[b' = 1 : b = 1]| = \operatorname{negl}(\lambda)$$

in the above security game.

Relaxed security notions. We now describe two relaxations of the security definition in Definition 5.1.

Definition 5.2 (Attribute-Selective Security). We say a slotted registered ABE scheme Π_{sRABE} satisfies attributeselective security if the adversary must declare its challenger attribute $x^* \in X_{\tau}$ at the beginning of the setup phase (before seeing pp).

Definition 5.3 (Security without Corruptions). We say a slotted registered ABE scheme Π_{sRABE} satisfies security without corruptions if the adversary in the security game of Definition 5.1 is not allowed to make corruption queries.

Remark 5.4 (Relationship between Definitions). While Definitions 5.2 and 5.3 relax the main security notion in Definition 5.1, there are generic ways to transform a scheme satisfying the relaxed notions of security into one that satisfies full security:

- First, if Π_{sRABE} satisfies attribute-selective security, we can obtain an adaptively-secure scheme (where the attribute is chosen in the challenge phase) via standard complexity leveraging [BB04]. Using complexity leveraging will require assuming sub-exponential hardness, and moreover, the size of the scheme parameters will scale with the attribute length.
- If Π_{sRABE} satisfies security without corruptions, we can apply the generic transformation from [FWW23] to obtain a scheme with security against adversaries that can make corruption queries in the random oracle model. Since our base registered ABE scheme is already in the random oracle model, this transformation is essentially free (the transformation itself only incurs constant overhead over the base scheme).

5.1 Additional Building Blocks

Similar to the registered ABE scheme from [CHW25], our construction also relies on simulation-sound extractable NIZK arguments, an explainable discrete Gaussian sampler, as well as a Gaussian preimage smudging lemma. We review these notions below:

Simulation-sound extractable NIZKs. The first primitive we require is a simulation-sound extractable noninteractive zero-knowledge (NIZK) argument for NP [BFM88, FLS90, Sah99, DDO⁺01]. Similar to [CHW25], we will use the NIZK argument of knowledge to extract secret keys in the security analysis of the registered ABE scheme. We give the definition below (taken mostly verbatim from [CHW25]):

Definition 5.5 (Simulation-Sound Extractable NIZK). A simulation-sound extractable NIZK argument Π_{NIZK} for NP is a tuple of efficient algorithms Π_{NIZK} = (Setup, TrapSetup, Prove, Verify, Sim, Extract) with the following syntax:

- Setup $(1^{\lambda}) \rightarrow \text{crs:}$ On input the security parameter λ , the setup algorithm outputs a common reference string crs.
- TrapSetup $(1^{\lambda}) \rightarrow (crs, td)$: On input the security parameter λ , the trapdoor setup algorithm outputs a common reference string crs and a trapdoor td.
- Prove(crs, C, x, w) $\rightarrow \pi$: On input the common reference string crs, a Boolean circuit $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$, a statement $x \in \{0, 1\}^n$, and a witness $w \in \{0, 1\}^h$, the prove algorithm outputs a proof π .
- Verify(crs, C, x, π) $\rightarrow b$: On input the common reference string crs, a Boolean circuit $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$, a statement $x \in \{0, 1\}^n$, and a proof π , the verification algorithm outputs a bit $b \in \{0, 1\}$.
- Sim(td, $C, x) \to \pi$: On input the trapdoor td, a Boolean circuit $C: \{0, 1\}^n \times \{0, 1\}^h \to \{0, 1\}$, and a statement $x \in \{0, 1\}^n$, the simulation algorithm outputs a proof π .
- Extract(td, C, x, π) \rightarrow w: On input the trapdoor td, a Boolean circuit $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$, a statement $x \in \{0, 1\}^n$, the extraction algorithm outputs a witness $w \in \{0, 1\}^h$ (or a special symbol \perp).

We require that Π_{NIZK} satisfy the following properties:

• **Completeness:** For all $\lambda \in \mathbb{N}$, all Boolean circuits $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$, all statements $x \in \{0, 1\}^n$ and witnesses $w \in \{0, 1\}^h$ where C(x, w) = 1,

$$\Pr\left[\operatorname{Verify}(\operatorname{crs}, C, x, \pi) = 1: \begin{array}{c} \operatorname{crs} \leftarrow \operatorname{Setup}(1^{\lambda}) \\ \pi \leftarrow \operatorname{Prove}(\operatorname{crs}, C, x, w) \end{array}\right] = 1.$$

- **Zero-knowledge:** For a security parameter λ , an adversary \mathcal{A} , and a bit $b \in \{0, 1\}$, we define the zero-knowledge security game as follows:
 - If b = 0, the challenger samples $crs \leftarrow Setup(1^{\lambda})$ and if b = 1, the challenger samples $(crs, td) \leftarrow TrapSetup(1^{\lambda})$. The challenger gives crs to \mathcal{A} .
 - Algorithm \mathcal{A} can now make adaptive queries of the form (C, x, w), where $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$ is a Boolean circuit, $x \in \{0, 1\}^n$ is a statement, and $w \in \{0, 1\}^h$ is a witness.
 - * The challenger first checks if C(x, w) = 1. If not, the challenger responds with \perp .
 - * Otherwise, if b = 0, the challenger replies with $\pi \leftarrow \text{Prove}(\text{crs}, C, x, w)$. If b = 1, the challenger replies with $\pi \leftarrow \text{Sim}(\text{td}, C, x)$.
 - After \mathcal{A} is finished making queries, it outputs a bit $b' \in \{0, 1\}$, which is the output of the experiment.

We say that Π_{NIZK} satisfies computational zero-knowledge if for all efficient adversaries \mathcal{A} , there exists a negligible function negl(·) such that for all $\lambda \in \mathbb{N}$, $|\Pr[b' = 1 | b = 0] - \Pr[b' = 1 | b = 1]| = \operatorname{negl}(\lambda)$ in the zero-knowledge security game.

- Simulation extractability: For a security parameter λ, and an adversary A, we define the simulation extractability games as follows:
 - The challenger starts by sampling (crs, td) \leftarrow TrapSetup(1^{λ}) and gives crs to \mathcal{A} . The challenger also initializes an (empty) list Q.
 - Algorithm \mathcal{A} can now make adaptive queries (C, x) where $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$ is a Boolean circuit and $x \in \{0, 1\}^n$ is a statement. The challenger replies with $\pi \leftarrow \text{Sim}(\text{td}, C, x)$ and adds (C, x, π) to Q.
 - After \mathcal{A} is finished making queries, it outputs a Boolean circuit $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$, a statement $x \in \{0, 1\}^n$, and a proof π .
 - The challenger computes $w = \text{Extract}(\text{td}, C, x, \pi)$ and outputs b' = 1 if $\text{Verify}(\text{crs}, C, x, \pi) = 1$, $(C, x, \pi) \notin Q$ and C(x, w) = 0. Otherwise, the challenger outputs b' = 0.

We say that Π_{NIZK} satisfies simulation extractability if for all efficient adversaries \mathcal{A} , there exists a negligible function negl(·) such that for all $\lambda \in \mathbb{N}$, $\Pr[b' = 1] = \operatorname{negl}(\lambda)$ in the simulation extractability game.

The work of [DDO⁺01] show how to construct a simulation-sound extractable NIZK for NP from any NIZK for NP together with a public-key encryption scheme (and a one-time signature scheme, which is implied by public-key encryption). Both NIZKs for NP [PS19, Wat24, WWW25, BCD⁺25] and public-key encryption [Reg05] are known from the plain LWE assumption.

Remark 5.6 (On Semi-Malicious Security). In our registered ABE construction (Construction 5.11), each user's public key includes a NIZK proof of knowledge of the randomness used to sample the key. Essentially, the NIZK serves to provide a proof of well-formedness for public keys. A natural question is whether we can provide a more modular description where we first show that the registered ABE scheme without NIZK proofs of well-formedness is secure against *semi-malicious* adversaries. In this model, when the adversary specifies a public key in the challenge phase, it must also reveal the key-generation randomness used to sample the key. We can then upgrade a scheme with semi-malicious security to one with full security by requiring public keys to include a NIZK proof of knowledge of the key-generation randomness; indeed, such a transformation was recently described in [LWW25] in the context of multi-authority registered ABE. Unfortunately, this modular approach does not apply in our setting (for the same reason it did not apply in [CHW25]). This is because in the security proof, the reduction algorithm needs to know the key-generation randomness associated with adversarially-chosen public keys when it is simulating *random oracle queries* (before the challenge phase). It is not meaningful to restrict the adversary to only query the random oracle on inputs that are well-formed. As a result, our current proof strategy critically relies on the ability to extract the key-generation randomness from *any* well-formed public key (as opposed to only the subset of public keys chosen during the challenge phase). We refer to the proof of Theorem 5.14 for the full details.

Explainable Gaussian sampling. Our construction relies on the re-randomization technique introduced in [CHW25]. Namely, the key-aggregation algorithm will sample and register a "dummy key," and the randomness for sampling the dummy key will be derived from the random oracle. In the security analysis, the reduction algorithm programs the aggregation randomness in order to correctly simulate the challenge ciphertext. To implement this, the work of [CHW25] relies on an *explainable* discrete Gaussian sampler for sampling from the distribution $A_{\sigma}^{-1}(z)$. Specifically, there is an Explain algorithm that takes a preimage $\mathbf{y} \leftarrow A_{\sigma}^{-1}(\mathbf{z})$ and outputs the randomness (for the sampler algorithm) that would produce \mathbf{x} . The work of [CHW25] observed that the Gentry-Peikert-Vaikuntanathan preimage sampling algorithm [GPV08] is explainable (when instantiated with an explainable discrete Gaussian sampler over the integers [LW22]). We give the formal definition and the required properties from [CHW25] below:

Definition 5.7 (Explainable Discrete Gaussian Preimage Sampler [CHW25, Definition 4.1]). Let λ be a security parameter, and *n*, *m*, *q* be lattice parameters. A (ρ , σ_{loss})-explainable discrete Gaussian preimage sampler Π_{DGS} with randomness length ρ and width parameter σ_{loss} is a pair of efficient algorithms $\Pi_{DGS} = (SamplePre, Explain)$ with the following syntax:

• SamplePre(1^{λ}, **A**, **T**, **z**, σ ; r): On input the security parameter λ , a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, a trapdoor $\mathbf{T} \in \mathbb{Z}_q^{m \times m}$, a target vector $\mathbf{z} \in \mathbb{Z}_q^n$, a width parameter $\sigma > 0$, and randomness $r \in \{0, 1\}^{\rho}$, the preimage sampling algorithm outputs a vector $\mathbf{y} \in \mathbb{Z}_q^m$.

• Explain $(1^{\lambda}, 1^{\kappa}, \mathbf{A}, \mathbf{T}, \mathbf{z}, \mathbf{y}, \sigma) \rightarrow r$: On input the security parameter λ , the precision parameter κ , a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, a trapdoor $\mathbf{T} \in \mathbb{Z}_q^{m \times m}$, a target vector $\mathbf{z} \in \mathbb{Z}_q^n$, a preimage $\mathbf{y} \in \mathbb{Z}_q^m$, and a width parameter $\sigma > 0$, the explain algorithm outputs a string $r \in \{0, 1\}^{\rho}$.

Moreover, there exists a negligible function $negl(\cdot)$ such that for all $\lambda \in \mathbb{N}$, all matrices $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, trapdoors $\mathbf{T} \in \mathbb{Z}_q^{m \times m}$ where $\mathbf{A}\mathbf{T} = \mathbf{G}$, all targets $\mathbf{z} \in \mathbb{Z}_q^n$ where $\|\mathbf{z}\| \le 2^{\lambda}$, and all width parameters σ where $\|\mathbf{T}\| \cdot \sigma_{\text{loss}} \le \sigma \le 2^{\lambda}$, the following two properties hold:

• Correctness: The following distributions are statistically indistinguishable:

 $\{\mathbf{x} \leftarrow \text{SamplePre}(1^{\lambda}, \mathbf{A}, \mathbf{T}, \mathbf{y}, \sigma)\} \text{ and } \{\mathbf{x} \leftarrow \mathbf{A}_{\sigma}^{-1}(\mathbf{y})\}.$

Moreover, for all **x** in the support of SamplePre(1^{λ} , **A**, **T**, **y**, σ), we have **Ax** = **y**.

- **Explainability:** For all $\kappa \in \mathbb{N}$, the statistical distance between the following distributions is bounded by $1/\kappa + \operatorname{negl}(\lambda)$:
 - $\mathcal{D}_{\text{SamplePre}}$: Sample $r \stackrel{\mathbb{R}}{\leftarrow} \{0, 1\}^{\rho}$ and $\mathbf{y} \leftarrow \text{SamplePre}(1^{\lambda}, \mathbf{A}, \mathbf{T}, \mathbf{z}, \sigma; r)$. Output (\mathbf{y}, r) .
 - $\mathcal{D}_{\text{Explain}}$: Sample $r' \leftarrow \{0, 1\}^{\rho}$ and $\mathbf{y} \leftarrow \text{SamplePre}(1^{\lambda}, \mathbf{A}, \mathbf{T}, \mathbf{z}, \sigma; r')$ and $r \leftarrow \text{Explain}(1^{\lambda}, 1^{\kappa}, \mathbf{A}, \mathbf{T}, \mathbf{z}, \mathbf{y}, \sigma)$. Output (\mathbf{y}, r) .

Theorem 5.8 (Explainable Discrete Gaussian Preimage Sampler [CHW25, Theorem 4.2]). Let λ be a security parameter, and n, m, q be lattice parameters. There exists an explicit (ρ, σ_{loss}) -explainable discrete Gaussian preimage sampler where $\rho = \text{poly}(\lambda, n, m, \log q)$ and $\sigma_{loss} = 18m^{3/2} \log(m\lambda) \log \log q$.

Explainable re-randomizer. Let pp = (B, W, T) be the public parameters for a matrix commitment scheme. To apply the re-randomization technique from [CHW25] to prove security of our registered ABE scheme, we require an explainable sampler for sampling from the following distribution:

$$\left\{ (\mathbf{C}, \mathbf{y}_1, \dots, \mathbf{y}_N) : \begin{array}{c} \mathbf{C} \xleftarrow{\mathbb{R}} \mathbb{Z}_q^{n \times m} \\ \mathbf{y}_i \leftarrow \mathbf{B}_{\sigma}^{-1}(-\mathbf{C}\mathbf{v}_i) \end{array} \right\}$$

where $\mathbf{V} = [\mathbf{v}_1 | \cdots \mathbf{v}_N] = \text{Ver}^{\text{mat}}(\text{pp}, 1^N)$. To simplify the exposition of our main construction, we abstract out the properties we require and then show how to build an explainable sampler for this distribution below:

Theorem 5.9 (Explainable Re-randomizer). Let λ be a security parameter, n, m, q be lattice parameters, and ρ be a randomness parameter. There exists a pair of efficient algorithms (Sample, Explain) with the following syntax:

- Sample(pp, $1^{\lambda}, 1^{N}, \sigma; r) \rightarrow (\mathbf{C}, \mathbf{y}_{1}, \dots, \mathbf{y}_{N})$: On input the public parameters pp, the security parameter λ , the dimension N, a width parameter $\sigma > 0$, and randomness $r \in \{0, 1\}^{\rho}$, output $\mathbf{C} \in \mathbb{Z}_{q}^{n \times m}$ and $\mathbf{y}_{1}, \dots, \mathbf{y}_{N} \in \mathbb{Z}_{q}^{m}$.
- Explain(pp, 1^{λ} , 1^{κ} , $(\mathbf{C}, \mathbf{y}_1, \dots, \mathbf{y}_N), \sigma$) $\rightarrow r$: On input the public parameters pp, a security parameter λ , a precision parameter κ , a matrix $\mathbf{C} \in \mathbb{Z}_q^{n \times m}$, vectors $\mathbf{y}_1, \dots, \mathbf{y}_N \in \mathbb{Z}_q^m$, and a width parameter $\sigma > 0$, output randomness $r \in \{0, 1\}^{\rho}$.

There exists a polynomial $\rho = \text{poly}(\lambda, N, n, m, \log q)$ such that for all $n \ge \lambda$ and $m \ge 2n \log q$, for all but a negligible fractions of matrices $\mathbf{B} \in \mathbb{Z}_q^{n \times m}$, and all $\mathbf{W} \in \mathbb{Z}_q^{2m^2n \times m}$, $\mathbf{T} \in \mathbb{Z}_q^{(2m^2+1)m \times 2m^3}$ where $[\mathbf{I}_{2m^2} \otimes \mathbf{B} \mid \mathbf{W}] \cdot \mathbf{T} = \mathbf{I}_{2m^2} \otimes \mathbf{G}$, all parameters $N \in \mathbb{N}$, and setting pp = ($\mathbf{B}, \mathbf{W}, \mathbf{T}$), the following properties hold:

- Sampling distribution: For all width parameters $\|\mathbf{T}\| \cdot O(m^{12}N^3) < \sigma < 2^{\lambda}$, the statistical distance between the following distributions is bounded by $\operatorname{negl}(\lambda)$:
 - $(\mathbf{C}, \mathbf{y}_1, \dots, \mathbf{y}_N) \leftarrow \text{Sample}(\text{pp}, 1^{\lambda}, 1^N, \sigma; r).$ - $(\mathbf{C}, \mathbf{y}_1, \dots, \mathbf{y}_N)$ where $\mathbf{C} \leftarrow \mathbb{R}_q^{\mathbb{R} \times m}$ and $\mathbf{y}_i \leftarrow \mathbf{B}_{\sigma}^{-1}(-\mathbf{C}\mathbf{v}_i)$, and $[\mathbf{v}_1 | \cdots | \mathbf{v}_N] = \text{Ver}^{\text{mat}}(\text{pp}, 1^N).$

Moreover, for all (C, y_1, \ldots, y_N) in the support of Sample $(pp, 1^{\lambda}, 1^N, \sigma)$, it holds that $By_i = -Cv_i$ and $||y_i|| \le \sqrt{m\sigma}$.

- Explainability: There exists a negligible function negl(·) such that for all precision parameters κ , dimensions $N \in \mathbb{N}$, and width parameters $||\mathbf{T}|| \cdot O(m^{12}N^3) < \sigma < 2^{\lambda}$, the statistical distance between the following distributions is bounded by $1/\kappa + \text{negl}(\lambda)$:
 - $\mathcal{D}_{\text{Sample}}$: Sample $r \leftarrow \{0, 1\}^{\rho}$ and $(C, y_1, \dots, y_N) \leftarrow \text{Sample}(\text{pp}, 1^{\lambda}, 1^N, \sigma; r)$. Output (C, y_1, \dots, y_N, r) .
 - $\mathcal{D}_{\text{Explain}}$: Sample $r' \stackrel{\mathbb{R}}{\leftarrow} \{0,1\}^{\rho}$ and $(\mathbf{C},\mathbf{y}_1,\ldots,\mathbf{y}_N) \leftarrow \text{Sample}(\text{pp},1^{\lambda},1^N,\sigma;r')$. Then, sample the randomness $r \leftarrow \text{Explain}(\text{pp},1^{\lambda},1^{\kappa},(\mathbf{C},\mathbf{y}_1,\ldots,\mathbf{y}_N),\sigma)$. Output $(\mathbf{C},\mathbf{y}_1,\ldots,\mathbf{y}_N,r)$.

Proof. We define the algorithms as follows:

- Sample(pp, 1^λ, 1^N, σ; r): On input the public parameters pp = (B, W, T), the security parameter λ, the dimension N ∈ N, a width parameter σ > 0, and randomness r ∈ {0, 1}^ρ, the sampling algorithm proceeds as follows:
 - 1. Compute $\mathbf{V}_N = \operatorname{Ver}^{\operatorname{mat}}(\operatorname{pp}, 1^N)$. Parse $\mathbf{V}_N = [\mathbf{v}_1 | \cdots | \mathbf{v}_N]$ where each $\mathbf{v}_i \in \mathbb{Z}_q^m$.

п

- 2. Let $\mathbf{G}_N = \mathbf{I}_N \otimes \mathbf{G} \in \mathbb{Z}_q^{nN \times mN}$, and let $\mathbf{m}_1, \ldots, \mathbf{m}_{mN} \in \mathbb{Z}_q^{nN}$ be the columns of \mathbf{G}_N . Let $\mathbf{M}_i \in \mathbb{Z}_q^{n \times N}$ be the matrix where $\operatorname{vec}(\mathbf{M}_i) = \mathbf{m}_i$. Then, for each $i \in [mN]$, compute the commitment $\mathbf{C}_i = \operatorname{Com}^{\max}(\operatorname{pp}, \mathbf{M}_i) \in \mathbb{Z}_q^{n \times m}$ and the opening $\mathbf{Z}_i = \operatorname{Open}^{\max}(\operatorname{pp}, \mathbf{M}_i) \in \mathbb{Z}_q^{m \times N}$.
- 3. Define the vector $\mathbf{r}_i \in \mathbb{Z}_q^{mN+k}$

$$\mathbf{r}_i = \begin{bmatrix} \operatorname{vec}(\mathbf{Z}_i) \\ \mathbf{G}_{nm}^{-1}(\operatorname{vec}(\mathbf{C}_i)) \end{bmatrix} \in \mathbb{Z}_q^{mN+k}.$$

Let $\mathbf{R} = [\mathbf{r}_1 \mid \cdots \mid \mathbf{r}_{mN}] \in \mathbb{Z}_q^{(mN+k) \times mN}$.

4. Let $k = nm \lceil \log q \rceil$ and let $\mathbf{D} = [\mathbf{D}_1 \mid \cdots \mid \mathbf{D}_k] \in \mathbb{Z}_q^{n \times mk}$ be the matrix where $vec(\mathbf{D}_i)$ is the *i*th column of \mathbf{G}_{nm} . Now define the matrix \mathbf{A} as follows:

$$\mathbf{A} = \begin{bmatrix} \mathbf{B} & & \mathbf{D}(\mathbf{I}_k \otimes \mathbf{v}_1) \\ & \ddots & & \vdots \\ & & \mathbf{B} & \mathbf{D}(\mathbf{I}_k \otimes \mathbf{v}_N) \end{bmatrix} \in \mathbb{Z}_q^{nN \times (mN+k)}.$$
(5.1)

5. Sample

$$\begin{vmatrix} \mathbf{y}_1 \\ \vdots \\ \mathbf{y}_N \\ \mathbf{y}_0 \end{vmatrix} \leftarrow \mathsf{DGS.Sample}(1^\lambda, \mathbf{A}, \mathbf{R}, \mathbf{0}^{nN}, \sigma; r),$$

where $\mathbf{y}_1, \ldots, \mathbf{y}_N \in \mathbb{Z}_q^m$ and $\mathbf{y}_0 \in \mathbb{Z}_q^k$. If $\|\mathbf{y}_i\| > \sqrt{m\sigma}$ for any $i \in [N]$, then output $\mathbf{C} = \mathbf{0}^{n \times m}$ and $\mathbf{y}_1, \ldots, \mathbf{y}_N = \mathbf{0}^m$. Otherwise, output the matrix $\mathbf{C} = \mathbf{D}(\mathbf{y}_0 \otimes \mathbf{I}_m)$ and the vectors $\mathbf{y}_1, \ldots, \mathbf{y}_N$.

- Explain(pp, 1^{λ} , 1^{κ} , (C, y_1 , ..., y_N), σ): On input the public parameters pp = (B, W, T), the precision parameter κ , a target (C, y_1 , ..., y_N), and a width parameter σ , the explain algorithm proceeds as follows:
 - Compute $y_0 \leftarrow \text{SamplePre}(\mathbf{G}_{nm}, \mathbf{I}_k, \text{vec}(\mathbf{C}), \sigma)$.
 - Let $\mathbf{y}^{\mathsf{T}} = [\mathbf{y}_1^{\mathsf{T}} | \cdots | \mathbf{y}_N^{\mathsf{T}} | \mathbf{y}_0^{\mathsf{T}}]^{\mathsf{T}}$. Compute \mathbf{A}, \mathbf{R} as in Sample(pp, $1^{\lambda}, 1^N, \sigma$). By construction, \mathbf{A}, \mathbf{R} is a deterministic function of pp. Output $r \leftarrow \mathsf{DGS}.\mathsf{Explain}(1^{\lambda}, 1^{\kappa}, \mathbf{A}, \mathbf{R}, \mathbf{0}^{nN}, \mathbf{y}, \sigma)$.

We now show that (Sample, Explain) satisfy the requirements in Theorem 5.9.

Sampling distribution. First, we argue that $AR = I_N \otimes G$. By Remark 3.9, for all $i \in [mN]$, $C_i V_N = M_i - BZ_i$. By Eq. (3.2), this means

$$\mathbf{m}_{i} = \operatorname{vec}(\mathbf{M}_{i}) = (\mathbf{I}_{N} \otimes \mathbf{B}) \cdot \operatorname{vec}(\mathbf{Z}_{i}) + (\mathbf{I}_{N} \otimes \mathbf{C}_{i}) \cdot \operatorname{vec}(\mathbf{V}_{N}).$$
(5.2)

Next, recall that $\mathbf{D} = [\mathbf{D}_1 | \cdots | \mathbf{D}_k]$ where $\mathbf{D}_i \in \mathbb{Z}_q^{n \times m}$. For any vector $\mathbf{x} \in \mathbb{Z}_q^k$, we have

$$\mathbf{D}(\mathbf{x} \otimes \mathbf{I}_m) = [\mathbf{D}_1 | \cdots | \mathbf{D}_k](\mathbf{x} \otimes \mathbf{I}_m) = \sum_{i \in [k]} x_i \mathbf{D}_i.$$

Then,

$$\operatorname{vec}(\mathbf{D}(\mathbf{x}\otimes\mathbf{I}_m)) = \operatorname{vec}\left(\sum_{i\in[k]}x_i\mathbf{D}_i\right) = \sum_{i\in[k]}x_i\cdot\operatorname{vec}(\mathbf{D}_i) = [\operatorname{vec}(\mathbf{D}_1) \mid \cdots \mid \operatorname{vec}(\mathbf{D}_k)] \cdot \mathbf{x} = \mathbf{G}_{nm} \cdot \mathbf{x}.$$
 (5.3)

This means

$$\operatorname{vec}(\mathbf{C}_i) = \mathbf{G}_{nm} \cdot \mathbf{G}_{nm}^{-1}(\operatorname{vec}(\mathbf{C}_i)) = \operatorname{vec}(\mathbf{D}(\mathbf{G}_{nm}^{-1}(\operatorname{vec}(\mathbf{C}_i)) \otimes \mathbf{I}_m)).$$

Since C_i and $D(G_{nm}^{-1}(vec(C_i)) \otimes I_m)$ are matrices with the same dimension, we conclude that

$$\mathbf{C}_i = \mathbf{D}(\mathbf{G}_{nm}^{-1}(\operatorname{vec}(\mathbf{C}_i)) \otimes \mathbf{I}_m).$$
(5.4)

Now,

$$\begin{aligned} \mathbf{Ar}_{i} &= (\mathbf{I}_{N} \otimes \mathbf{B}) \cdot \operatorname{vec}(\mathbf{Z}_{i}) + \begin{bmatrix} \mathbf{D}(\mathbf{I}_{k} \otimes \mathbf{v}_{1}) \\ \vdots \\ \mathbf{D}(\mathbf{I}_{k} \otimes \mathbf{v}_{N}) \end{bmatrix} \mathbf{G}_{nm}^{-1}(\operatorname{vec}(\mathbf{C}_{i})) & \text{by Eq. (5.1)} \end{aligned}$$

$$= (\mathbf{I}_{N} \otimes \mathbf{B}) \cdot \operatorname{vec}(\mathbf{Z}_{i}) + \begin{bmatrix} \mathbf{D}(\mathbf{G}_{nm}^{-1}(\operatorname{vec}(\mathbf{C}_{i})) \otimes \mathbf{I}_{m})\mathbf{v}_{1} \\ \vdots \\ \mathbf{D}(\mathbf{G}_{nm}^{-1}(\operatorname{vec}(\mathbf{C}_{i})) \otimes \mathbf{I}_{m})\mathbf{v}_{N} \end{bmatrix}$$

$$= (\mathbf{I}_{N} \otimes \mathbf{B}) \cdot \operatorname{vec}(\mathbf{Z}_{i}) + (\mathbf{I}_{N} \otimes \mathbf{C}_{i})\operatorname{vec}(\mathbf{V}_{N}) & \text{by Eq. (5.4)} \end{aligned}$$

$$= \mathbf{m}_{i}$$

$$\begin{aligned} \mathbf{M}_{i} = \mathbf{M}_{i} & \text{by Eq. (5.2).} \end{aligned}$$

We conclude that $AR = [m_1 | \cdots | m_{mN}] = G_N$, as required. By Remark 3.9, for all $i \in [mN]$,

 $\|\mathbf{Z}_i\| \le O(\|\mathbf{T}\| \cdot m^7 \log q \log N).$

Thus, $\|\mathbf{R}\| \leq O(\|\mathbf{T}\| \cdot m^7 \log q \log N)$. Let $\sigma'_{\text{loss}} = 18(mN+k)^{3/2} \log((mN+k)\lambda) \log \log q \leq O(m^4N^2)$. Then,

$$\|\mathbf{R}\| \cdot \sigma_{\text{loss}}' \le \|\mathbf{T}\| \cdot O(m^{12}N^3).$$

If $\|\mathbf{R}\| \cdot \sigma'_{\text{loss}} \leq \|\mathbf{T}\| \cdot O(m^{12}N^3) \leq \sigma \leq 2^{\lambda}$, the distribution of $(\mathbf{y}_0, \mathbf{y}_1, \dots, \mathbf{y}_N)$ output by DGS.Sample $(1^{\lambda}, \mathbf{A}, \mathbf{R}, \mathbf{0}^{nN}, \sigma; r)$ is statistically close to sampling

$$\begin{array}{c} \mathbf{y}_{1} \\ \vdots \\ \mathbf{y}_{N} \\ \mathbf{y}_{0} \end{array} \leftarrow \mathbf{A}_{\sigma}^{-1}(\mathbf{0}^{nN}).$$
 (5.5)

We now characterize the distribution of $(\mathbf{y}_0, \mathbf{y}_1, \dots, \mathbf{y}_N)$. We proceed by a hybrid argument:

• \mathcal{D}_0 : Sample $(\mathbf{y}_0, \mathbf{y}_1, \dots, \mathbf{y}_N)$ according to Eq. (5.5). Set $\mathbf{C} = \mathbf{D}(\mathbf{y}_0 \otimes \mathbf{I}_m)$ and output $(\mathbf{C}, \mathbf{y}_1, \dots, \mathbf{y}_N)$ if for all $i \in [N], \|\mathbf{y}_i\| \le \sqrt{m\sigma}$. Otherwise, if $\|\mathbf{y}_i\| > \sqrt{m\sigma}$ for some $i \in [N]$, output $(\mathbf{0}^{n \times m}, \mathbf{0}^m, \dots, \mathbf{0}^m)$.

- \mathcal{D}_1 : Sample $\mathbf{y}_0 \leftarrow D_{\mathbb{Z},\sigma}^k$ and set $\mathbf{C} = \mathbf{D}(\mathbf{y}_0 \otimes \mathbf{I}_m)$. Then, for $i \in [N]$, sample $\mathbf{y}_i \leftarrow \mathbf{B}_{\sigma}^{-1}(-\mathbf{C}\mathbf{v}_i)$. Output $(\mathbf{C}, \mathbf{y}_1, \dots, \mathbf{y}_N)$ if for all $i \in [N]$, $\|\mathbf{y}_i\| \leq \sqrt{m\sigma}$. Otherwise, if $\|\mathbf{y}_i\| > \sqrt{m\sigma}$ for some $i \in [N]$, output $(\mathbf{0}^{n \times m}, \mathbf{0}^m, \dots, \mathbf{0}^m)$.
- \mathcal{D}_2 : Sample $C \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^{n \times m}$ and for each $i \in [N]$, sample $\mathbf{y}_i \leftarrow \mathbf{B}_{\sigma}^{-1}(-C\mathbf{v}_i)$. Output $(C, \mathbf{y}_1, \dots, \mathbf{y}_N)$ if for all $i \in [N]$, $\|\mathbf{y}_i\| \le \sqrt{m\sigma}$. Otherwise, if $\|\mathbf{y}_i\| > \sqrt{m\sigma}$ for some $i \in [N]$, output $(\mathbf{0}^{n \times m}, \mathbf{0}^m, \dots, \mathbf{0}^m)$.
- \mathcal{D}_3 : Sample $\mathbf{C} \leftarrow \mathbb{Z}_q^{n \times m}$ and for each $i \in [N]$, sample $\mathbf{y}_i \leftarrow \mathbf{B}_{\sigma}^{-1}(-\mathbf{C}\mathbf{v}_i)$. Output $(\mathbf{C}, \mathbf{y}_1, \dots, \mathbf{y}_N)$. In particular, there is no additional check on the norms of $\mathbf{y}_1, \dots, \mathbf{y}_N$.

As argued above, when $\|\mathbf{T}\| \cdot \sigma_{\text{loss}} \leq \sigma \leq 2^{\lambda}$, the output of Sample(pp, $1^{\lambda}, 1^{N}, \sigma; r$) is statistically close to \mathcal{D}_{0} . We now show that \mathcal{D}_{0} is statistically close to \mathcal{D}_{3} :

• By Lemma 3.4, for all but a negligible fraction of matrices $\mathbf{B} \in \mathbb{Z}_q^{n \times m}$, the distribution of $(\mathbf{y}_0, \mathbf{y}_1, \dots, \mathbf{y}_N)$ is statistically close to sampling $\mathbf{y}_0 \leftarrow D_{\mathbb{Z},\sigma^k}$ and $\mathbf{y}_i \leftarrow \mathbf{B}_{\sigma}^{-1}(-\mathbf{D}(\mathbf{I}_k \otimes \mathbf{v}_i)\mathbf{y}_0)$. By Eq. (3.1), we can write

$$\mathbf{D}(\mathbf{I}_k \otimes \mathbf{v}_i)\mathbf{y}_0 = \mathbf{D}(\mathbf{y}_0 \otimes \mathbf{I}_m)\mathbf{v}_i = \mathbf{C}\mathbf{v}_i.$$

This is the distribution in \mathcal{D}_1 .

- By Eq. (5.3), we can write $vec(C) = vec(D(\mathbf{y}_0 \otimes \mathbf{I}_m)) = \mathbf{G}_{nm} \cdot \mathbf{y}_0$. In $\mathcal{D}_1, \mathbf{y}_0 \leftarrow D_{\mathbb{Z},\sigma}^k$ so by Lemma 3.3, $\mathbf{G}_{nm} \cdot \mathbf{y}_0$ is statistically close to uniform over \mathbb{Z}_q^{nm} . Correspondingly, this means that the marginal distribution of C is statistically close to uniform over $\mathbb{Z}_q^{n\times m}$, as required.
- By Lemma 3.2, for all $i \in [N]$, we have $\Pr[||\mathbf{y}_i|| \ge \sqrt{m\sigma}] \le O(2^{-m})$. Since $m \ge 2\lambda$ and $N \le 2^{\lambda}$, the probability that there exists $i \in [N]$ where $||\mathbf{y}_i|| \ge \sqrt{m\sigma}$ is at most $2^{\lambda}/2^m \le 2^{-\lambda}$. Thus, with overwhelming probability, the vectors $||\mathbf{y}_i||$ sampled in \mathcal{D}_2 all satisfy $||\mathbf{y}_i|| \le \sqrt{m\sigma}$, so \mathcal{D}_2 and \mathcal{D}_3 are statistically indistinguishable.

The claim now follows by a hybrid argument. Finally, the worst-case guarantee that $\mathbf{B}\mathbf{y}_i = -\mathbf{C}\mathbf{v}_i$ follows immediately by the worst-case guarantee of Π_{DGS} . Namely, the $(\mathbf{y}_0, \mathbf{y}_1, \dots, \mathbf{y}_N)$ output by DGS.Sample always satisfies $\mathbf{y}_i = \mathbf{B}(-\mathbf{D}(\mathbf{I}_k \otimes \mathbf{v}_i)\mathbf{y}_0) = \mathbf{B}(-\mathbf{C}\mathbf{v}_i)$. The norm constraint is ensured by construction.

Explainability. We proceed via a hybrid argument. We start by defining a sequence of distributions:

• \mathcal{D}_0 : This is distribution $\mathcal{D}_{\text{Sample}}$. Namely, sample $r \stackrel{\mathbb{R}}{\leftarrow} \{0, 1\}^{\rho}$, and $(C, y_1, \dots, y_N) \leftarrow \text{Sample}(\text{pp}, 1^{\lambda}, 1^N, \sigma; r)$. Concretely, compute A, **R** as in Sample(pp, $1^{\lambda}, 1^N, \sigma$). Then sample

$$\mathbf{y} = \begin{bmatrix} \mathbf{y}_1 \\ \vdots \\ \mathbf{y}_N \\ \mathbf{y}_0 \end{bmatrix} \leftarrow \mathsf{DGS}.\mathsf{Sample}(1^\lambda, \mathbf{A}, \mathbf{R}, \mathbf{0}^{nN}, \sigma; r),$$

where $\mathbf{y}_1, \ldots, \mathbf{y}_N \in \mathbb{Z}_q^m$ and $\mathbf{y}_0 \in \mathbb{Z}_q^k$. Let $\mathbf{C} = \mathbf{D}(\mathbf{y}_0 \otimes \mathbf{I}_m)$. If there exists $i \in [N]$ where $\|\mathbf{y}_i\| > \sqrt{m}\sigma$, then set $\mathbf{C} = \mathbf{0}^{n \times m}$ and $\mathbf{y}_1, \ldots, \mathbf{y}_N = \mathbf{0}^m$. Output $(\mathbf{C}, \mathbf{y}_1, \ldots, \mathbf{y}_N, r)$.

- \mathcal{D}_1 : Same as \mathcal{D}_0 except the distribution no longer checks the norm constraints $||\mathbf{y}_i|| > \sqrt{m\sigma}$.
- \mathcal{D}_2 : Same as \mathcal{D}_1 except after sampling y and C, sample $r' \leftarrow \mathsf{DGS}.\mathsf{Explain}(1^\lambda, 1^\kappa, \mathbf{A}, \mathbf{R}, \mathbf{0}^{nN}, \mathbf{y}, \sigma)$ and output $(\mathbf{C}, \mathbf{y}_1, \dots, \mathbf{y}_N, r')$.
- \mathcal{D}_3 : Same as \mathcal{D}_2 , except sample $\mathbf{y} \leftarrow \mathbf{A}_{\sigma}^{-1}(\mathbf{0}^{nN})$.
- \mathcal{D}_4 : Same as \mathcal{D}_3 , except instead sample $\mathbf{y}_0 \leftarrow D_{\mathbb{Z},\sigma}^k$. Then set $\mathbf{C} = \mathbf{D}(\mathbf{y}_0 \otimes \mathbf{I}_m)$ and sample $\mathbf{y}_i \leftarrow \mathbf{B}_{\sigma}^{-1}(-\mathbf{C}\mathbf{v}_i)$ for each $i \in [N]$. In this experiment, the vector \mathbf{y} is still defined as

$$\mathbf{y} = \begin{bmatrix} \mathbf{y}_1 \\ \vdots \\ \mathbf{y}_N \\ \mathbf{y}_0 \end{bmatrix}.$$

- \mathcal{D}_5 : Same as \mathcal{D}_4 , except sample $C \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^{n \times m}$ and $\mathbf{y}_0 \leftarrow (\mathbf{G}_{nm})_{\sigma}^{-1}(\operatorname{vec}(\mathbf{C}))$.
- \mathcal{D}_6 : Same as \mathcal{D}_5 except sample $\mathbf{y}'_0 \leftarrow D^k_{\mathbb{Z},\sigma}$ and set $\mathbf{C} = \mathbf{D}(\mathbf{y}'_0 \otimes \mathbf{I}_m)$. Next, sample $\mathbf{y}_0 \leftarrow (\mathbf{G}_{nm})^{-1}_{\sigma}(\operatorname{vec}(\mathbf{C}))$ and for each $i \in [N]$, sample $\mathbf{y}_i \leftarrow \mathbf{B}^{-1}_{\sigma}(-\mathbf{C}\mathbf{v}_i)$.
- \mathcal{D}_7 : Same as \mathcal{D}_6 except sample

$$\begin{array}{c|c} \mathbf{y}_1 \\ \vdots \\ \mathbf{y}_N \\ \mathbf{y}'_0 \\ \end{array} \leftarrow \mathbf{A}_{\sigma}^{-1}(\mathbf{0}^{nN})$$

Then set $C = D(\mathbf{y}_0' \otimes \mathbf{I}_m)$ and compute $\mathbf{y}_0 \leftarrow (\mathbf{G}_{nm})_{\sigma}^{-1}(\text{vec}(C))$. Let $\mathbf{y} = [\mathbf{y}_1^{\mathsf{T}} | \cdots | \mathbf{y}_N^{\mathsf{T}} | \mathbf{y}_0^{\mathsf{T}}]^{\mathsf{T}}$ and $r' \leftarrow \text{DGS.Explain}(\mathbf{1}^{\lambda}, \mathbf{1}^{\kappa}, \mathbf{A}, \mathbf{R}, \mathbf{0}^{nN}, \mathbf{y}, \sigma)$ as before.

• \mathcal{D}_8 : Same as \mathcal{D}_7 except sample

$$\begin{array}{c|c} \mathbf{y}_1 \\ \vdots \\ \mathbf{y}_N \\ \mathbf{y}'_0 \end{array} \leftarrow \mathsf{DGS.Sample}(1^\lambda, \mathbf{A}, \mathbf{R}, \mathbf{0}^{nN}, \sigma). \end{array}$$

Set $C = D(y'_0 \otimes I_m)$ and compute $y_0 \leftarrow (G_{nm})^{-1}_{\sigma}(vec(C))$. Let $y = [y_1^T | \cdots | y_N^T | y_0^T]^T$ and $r' \leftarrow DGS.Explain(1^{\lambda}, 1^{\kappa}, A, R, 0^{nN}, y, \sigma)$ as before.

• \mathcal{D}_9 : Same as \mathcal{D}_8 except sample $\mathbf{y}_0 \leftarrow$ SamplePre($\mathbf{G}_{nm}, \mathbf{I}_k, \text{vec}(\mathbf{C}), \sigma$). This is the distribution $\mathcal{D}_{\text{Explain}}$.

We now analyze each adjacent pair of distributions. As in the analysis of the sampling distribution, we assume that $\|\mathbf{R}\| \cdot \sigma'_{\text{loss}} \leq \|\mathbf{T}\| \cdot O(m^{12}N^3) \leq \sigma \leq 2^{\lambda}$, where $\sigma'_{\text{loss}} = 18(mN+k)^{3/2} \log((mN+k)\lambda) \log \log q \leq O(m^4N^2)$.

- Distributions \mathcal{D}_0 and \mathcal{D}_1 are statistically indistinguishable by the above analysis on the sampling distribution.
- Distributions \mathcal{D}_1 and \mathcal{D}_2 have statistical distance $1/\kappa + \text{negl}(\lambda)$ by the explainability property of Π_{DCS} .
- Distributions \mathcal{D}_2 and \mathcal{D}_3 are statistically indistinguishable by the correctness property of Π_{DGS} .
- Distributed \mathcal{D}_3 and \mathcal{D}_4 are statistically indistinguishable by Lemma 3.4. Specifically, for all but a negligible fraction of matrices $\mathbf{B} \in \mathbb{Z}_q^{n \times m}$, the distribution of $\mathbf{y} \leftarrow \mathbf{A}_{\sigma}^{-1}(\mathbf{0}^{nN})$ is statistically close to sampling $\mathbf{y}_0 \leftarrow D_{\mathbb{Z},\sigma^k}$ and $\mathbf{y}_i \leftarrow \mathbf{B}_{\sigma}^{-1}(-\mathbf{C}\mathbf{v}_i)$ where $\mathbf{C} = \mathbf{D}(\mathbf{y}_0 \otimes \mathbf{I}_m)$. The former sampling procedure corresponds to \mathcal{D}_2 while the latter corresponds to \mathcal{D}_3 .
- Distributions \mathcal{D}_4 and \mathcal{D}_5 are statistically indistinguishable by Lemma 3.3. Specifically, by Eq. (5.3), we can write $\operatorname{vec}(\mathbf{C}) = \operatorname{vec}(\mathbf{D}(\mathbf{y}_0 \otimes \mathbf{I}_m)) = \mathbf{G}_{nm} \cdot \mathbf{y}_0$. By Lemma 3.3, the following distributions are statistically indistinguishable:

- (C,
$$\mathbf{y}_0$$
) where $\mathbf{y}_0 \leftarrow D_{\mathbb{Z},\sigma}^k$ and $\mathbf{C} = \mathbf{D}(\mathbf{y}_0 \otimes \mathbf{I}_m)$.

- $(\mathbf{C}, \mathbf{y}_0)$ where $\mathbf{C} \leftarrow \mathbb{Z}_{q,\sigma}^{n \times m}$, $\mathbf{y}_0 \leftarrow (\mathbf{G}_{nm})_{\sigma}^{-1}(\operatorname{vec}(\mathbf{C}))$.

The left distribution corresponds to \mathcal{D}_4 while the right distribution corresponds to \mathcal{D}_5 .

- Distributions \mathcal{D}_5 and \mathcal{D}_6 are also statistically indistinguishable by Lemma 3.3 (via the same argument as in the previous case).
- Distributions \mathcal{D}_6 and \mathcal{D}_7 are statistically indistinguishable by Lemma 3.4 (via the same argument as used to argue indistinguishability of \mathcal{D}_2 and \mathcal{D}_3).
- Distributions D_7 and D_8 are statistically indistinguishable by the correctness property of Π_{DGS} .

• Distributions \mathcal{D}_8 and \mathcal{D}_9 are statistically indistinguishable by Lemma 3.5.

By a hybrid argument, the statistical distance between $\mathcal{D}_0 \equiv \mathcal{D}_{\text{Sample}}$ and $\mathcal{D}_9 \equiv \mathcal{D}_{\text{Explain}}$ can be bounded by $1/\kappa + \text{negl}(\lambda)$, as required.

Gaussian preimage smudging. We will also need the following noise smudging lemma from [CHW25]:

Lemma 5.10 (Gaussian Preimage Smudging [CHW25, Theorem 4.3, adapted]). Let n, m, q be lattice parameters such that $m \ge 2n \log q$ and q is prime. Then, for all but a q^{-n} -fraction of matrices $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, for all vectors $\mathbf{y} \in \mathbb{Z}_q^n$ in the column-span of \mathbf{A} , all $\mathbf{z} \in \mathbb{Z}_q^m$, and all width parameters $\sigma > \lambda^{\omega(1)} \sqrt{m} ||\mathbf{z}||$, then the following distributions are statistically indistinguishable:

$$\{\mathbf{A}_{\sigma}^{-1}(\mathbf{y}+\mathbf{A}\mathbf{z})\}\$$
 and $\{\mathbf{A}_{\sigma}^{-1}(\mathbf{y})+\mathbf{z}\}\$

5.2 Key-Policy Registered ABE Construction

In this section, we give our construction of a key-policy registered ABE scheme that supports arbitrary (boundeddepth) Boolean circuits. Specifically, the scheme has short parameters and thus, can be used to obtain a registered ABE scheme that supports an arbitrary polynomial number of users. All previous registered ABE schemes [HLWW23, ZZGQ23, GLWW24, AT24, CHW25] that does not rely on indistinguishability obfuscation or witness encryption could only support an *a priori* bounded number of users.

Construction 5.11 (Key-Policy Slotted Registered ABE). Let λ be a security parameter and τ be a policy-family parameter. We define the following quantities:

- Let (n, m, q, χ) be LWE parameters (which may be functions of λ, τ). Let $m' = n \lceil \log q \rceil$. Let $\sigma_{td}, \sigma_{agg} > 0$ be width parameters.
- Let $\ell = \ell(\tau)$ be the attribute length. Let $X = \{X_{\tau}\}_{\tau \in \mathbb{N}}$ where $X_{\tau} = \{0, 1\}^{\ell(\tau)}$ be the attribute space.
- Let \mathcal{F}_{τ} be the family of functions *f* that can be computed by a Boolean circuit of depth at most $d = d(\tau)$.
- Let λ_{DGS} be the security parameter for the explainable re-randomizer (Sample, Explain) from Theorem 5.9. Let $\rho = \text{poly}(\lambda_{\text{DGS}}, N, n, m, \log q)$ be the randomness length from Theorem 5.9.
- Let $\Pi_{NIZK} = (NIZK.Setup, NIZK.TrapSetup, NIZK.Prove, NIZK.Verify, NIZK.Sim, NIZK.Extract) be a simulation$ $sound extractable NIZK argument for NP. Let <math>C_{valid}$ be the Boolean circuit that takes a statement (*i*, A, d, p, B, t), a witness r, and outputs 1 if t = Br + AG⁻¹(d) + p.²
- Let $H_1: \mathbb{N} \to \mathbb{Z}_q^m$ and $H_2: \{0, 1\}^* \to \{0, 1\}^{\rho}$ be hash functions (which we model as random oracles in the security analysis).

We construct a key-policy slotted registered ABE scheme Π_{sRABE} = (Setup, KeyGen, IsValid, Aggregate, Encrypt, Decrypt) as follows:

• Setup $(1^{\lambda}, 1^{\tau})$: On input the security parameter λ and the policy family parameter τ , sample

$$\begin{aligned} \operatorname{crs}_{\mathsf{NIZK}} &\leftarrow \mathsf{NIZK}.\mathsf{Setup}(1^{\lambda}) \\ (\mathbf{B},\mathbf{T}_{\mathbf{B}}) &\leftarrow \mathsf{TrapGen}(1^{n},1^{m},q), \mathbf{W} \xleftarrow{\mathbb{R}} \mathbb{Z}_{q}^{2m^{2}n \times m} \\ \mathbf{T} &\leftarrow \mathsf{SamplePre}([\mathbf{I}_{2m^{2}} \otimes \mathbf{B} \mid \mathbf{W}], \begin{bmatrix} \mathbf{I}_{2m^{2}} \otimes \mathbf{T}_{\mathbf{B}} \\ \mathbf{0} \end{bmatrix}, \mathbf{I}_{2m^{2}} \otimes \mathbf{G}, \sigma_{\mathsf{td}}) \\ \mathbf{B}_{0} \xleftarrow{\mathbb{R}} \mathbb{Z}_{q}^{n \times m}, \mathbf{p} \xleftarrow{\mathbb{R}} \mathbb{Z}_{q}^{n}. \end{aligned}$$

If $||\mathbf{T}|| > \sqrt{m}\sigma_{td}$, then set $\mathbf{T} = \begin{bmatrix} I_{2m^2 \otimes T_B} \\ \mathbf{0} \end{bmatrix}$. Let $pp_{com} = (\mathbf{B}, \mathbf{W}, \mathbf{T})$ and output the public parameters $pp = (crs_{NIZK}, pp_{com}, \mathbf{B}_0, \mathbf{p})$.

²Note that the circuit ignores the index *i*, but including it as part of the statement allows us to bind a proof to the specific index *i* (when appealing to simulation-sound extractability).

• KeyGen(pp, *i*, *f*): On input the public parameters pp = (crs_{NIZK}, pp_{com}, **B**₀, **p**), where pp_{com} = (**B**, **W**, **T**), compute

$$\mathbf{d}_{i} = H_{1}(i) \in \mathbb{Z}_{q}^{n}$$

$$\mathbf{V}_{\ell m} = \operatorname{Ver}^{\operatorname{mat}}(\operatorname{pp}_{\operatorname{com}}, 1^{\ell m}) \in \mathbb{Z}_{q}^{m \times \ell m}$$

$$\mathbf{A} = -\mathbf{B}_{0}\mathbf{V}_{\ell m} \in \mathbb{Z}_{q}^{n \times \ell m}$$

$$\mathbf{A}_{f} = \operatorname{EvalF}(\mathbf{A}, f) \in \mathbb{Z}_{q}^{n \times m}.$$
(5.6)

Sample $\mathbf{r} \leftarrow^{\mathbb{R}} \{0, 1\}^m$ and compute

$$\mathbf{t} = \mathbf{B}\mathbf{r} + \mathbf{A}_f \mathbf{G}^{-1}(\mathbf{d}_i) + \mathbf{p} \in \mathbb{Z}_q^n.$$

Compute a NIZK proof $\pi \leftarrow \text{NIZK}$.Prove $(C_{\text{valid}}, (i, \mathbf{A}_f, \mathbf{d}_i, \mathbf{p}, \mathbf{B}, \mathbf{t}), \mathbf{r})$. Output the public key $p\mathbf{k} = (\mathbf{t}, \pi)$ and the secret key $s\mathbf{k} = (\mathbf{t}, \mathbf{r})$.

• IsValid(pp, *i*, *f*, pk): On input the public parameters pp = (crs_{NIZK}, pp_{com}, **B**₀, **p**) where pp_{com} = (**B**, **W**, **T**) and a public key pk = (t, π), compute **d**_{*i*}, **A**_{*f*} according to Eq. (5.6) and output 1 if

NIZK.Verify(crs_{NIZK},
$$(i, \mathbf{A}_f, \mathbf{d}_i, \mathbf{p}, \mathbf{B}, \mathbf{t}), \pi$$
) = 1.

Otherwise, output 0.

• Aggregate(pp, $(pk_1, f_1), \ldots, (pk_N, f_N)$): On input the public parameters $pp = (crs_{NIZK}, pp_{com}, B_0, p)$ where $pp_{com} = (B, W, T)$, a list of public keys $pk_1 = (t_1, \pi_1), \ldots, pk_N = (t_N, \pi_N)$ and their associated policies f_1, \ldots, f_N , the aggregation algorithm computes for each $i \in [N]$,

$$\begin{aligned} \mathbf{C}_{i} &= \operatorname{Com}^{\mathrm{mat}}(\mathrm{pp}_{\mathrm{com}}, \mathbf{u}_{i}^{\mathsf{T}} \otimes \mathbf{t}_{i}) \in \mathbb{Z}_{q}^{n \times m} \\ \mathbf{Z}_{i} &= \operatorname{Open}^{\mathrm{mat}}(\mathrm{pp}_{\mathrm{com}}, \mathbf{u}_{i}^{\mathsf{T}} \otimes \mathbf{t}_{i}) \in \mathbb{Z}_{q}^{n \times N}, \end{aligned}$$

where $\mathbf{u}_i \in \mathbb{Z}_q^N$ is the *i*th unit vector. The aggregation algorithm parses $\mathbf{Z}_i = [\mathbf{z}_{i,1} | \cdots | \mathbf{z}_{i,N}]$. Next, it computes the re-randomization terms

$$\xi = (pp, (pk_1, f_1), \dots, (pk_N, f_N))$$
$$(C_0, \mathbf{z}_{0,1}, \dots, \mathbf{z}_{0,N}) = Sample(pp_{com}, 1^{\lambda_{DCS}}, 1^N, \sigma_{agg}; H_2(\xi)).$$

Output mpk = $C_0 + \sum_{i \in [N]} C_i$ and hsk_i = $(N, f_i, \mathbf{z}_{0,i} + \sum_{j \in [N]} \mathbf{z}_{j,i})$ for all $i \in [N]$.

• Encrypt(mpk, \mathbf{x}, μ): On input the master public key mpk = $\widehat{\mathbf{C}}$, an attribute $\mathbf{x} \in \{0, 1\}^{\ell}$, and a message $\mu \in \{0, 1\}$, sample the following:

$$\mathbf{s} \leftarrow^{\mathbb{R}} \mathbb{Z}_{q}^{n}$$
, $\mathbf{e} \leftarrow D_{\mathbb{Z},\chi}^{m}$, $\mathbf{R}_{\widehat{\mathbf{C}}}$, $\mathbf{R}_{\mathbf{B}_{0}} \leftarrow^{\mathbb{R}} \{0,1\}^{m \times m}$, $\mathbf{r}_{\mathbf{p}} \leftarrow^{\mathbb{R}} \{0,1\}^{m}$

If $\|\mathbf{e}\| > \sqrt{m\chi}$, then set $\mathbf{e} = \mathbf{0}^m$. Next, compute $C_{\mathbf{x}} = \operatorname{Com}^{mat}(\operatorname{pp}_{com}, \mathbf{x}^{\mathsf{T}} \otimes \mathbf{G}) \in \mathbb{Z}_q^{n \times m}$ and output the ciphertext

$$\mathsf{ct} = \left(\mathbf{s}^{\mathsf{T}} \mathbf{B} + \mathbf{e}^{\mathsf{T}} , \ \mathbf{s}^{\mathsf{T}} \widehat{\mathbf{C}} + \mathbf{e}^{\mathsf{T}} \mathbf{R}_{\widehat{\mathbf{C}}} , \ \mathbf{s}^{\mathsf{T}} (\mathbf{B}_0 + \mathbf{C}_{\mathbf{x}}) + \mathbf{e}^{\mathsf{T}} \mathbf{R}_{\mathbf{B}_0} , \ \mathbf{s}^{\mathsf{T}} \mathbf{p} + \mathbf{e}^{\mathsf{T}} \mathbf{r}_{\mathbf{p}} + \lfloor q/2 \rfloor \cdot \mu \right).$$

Decrypt(sk, hsk, f, x, ct): On input the secret key sk = (t, r), the helper decryption key hsk = (N, f, ẑ), the function f, the attribute x ∈ {0, 1}^ℓ, and the ciphertext ct = (c^T₁, c^T₂, c^T₃, c₄), the decryption algorithm first computes

$$\begin{aligned} \mathbf{V}_{\ell m} &= \mathsf{Ver}^{\mathsf{mat}}(\mathsf{pp}_{\mathsf{com}}, \mathbf{1}^{\ell m}) \in \mathbb{Z}_q^{m \times \ell m} \\ \mathbf{Z}_{\mathbf{x}} &= \mathsf{Open}^{\mathsf{mat}}(\mathsf{pp}_{\mathsf{com}}, \mathbf{x}^{\mathsf{T}} \otimes \mathbf{G}) \in \mathbb{Z}_q^{m \times \ell m} \end{aligned}$$

Next, it computes the homomorphic evaluation matrices

$$\mathbf{A} = -\mathbf{B}_0 \mathbf{V}_{\ell m} \in \mathbb{Z}_q^{n \times \ell m}$$
$$\mathbf{A}_f = \mathsf{EvalF}(\mathbf{A}, f) \in \mathbb{Z}_q^{n \times m}$$
$$\mathbf{H}_{\mathbf{A}, f, \mathbf{x}} = \mathsf{EvalFX}(\mathbf{A}, f, \mathbf{x}) \in \mathbb{Z}_q^{\ell m \times m}$$

Finally, it computes $\mathbf{d} = H_1(i) \in \mathbb{Z}_q^n$ and

$$\tilde{c}_{3}^{\mathsf{T}} = [\mathbf{c}_{1}^{\mathsf{T}} \mid \mathbf{c}_{3}^{\mathsf{T}}] \cdot \begin{bmatrix} -\mathbf{Z}_{\mathbf{x}} \\ -\mathbf{V}_{\ell m} \end{bmatrix} \cdot \mathbf{H}_{\mathbf{A}, f_{i}, \mathbf{x}} \mathbf{G}^{-1}(\mathbf{d}).$$

Next, it computes $\mathbf{v} = \text{VerLocal}^{\text{mat}}(\text{pp}_{\text{com}}, Nm, i)$ and the (noisy) encoding of the message

$$\tilde{\mu} = c_4 + \mathbf{c}_1^{\mathsf{T}}(\mathbf{r} - \hat{\mathbf{z}}) - \mathbf{c}_2^{\mathsf{T}}\mathbf{v} + \tilde{c}_3.$$

It outputs 0 if $-q/4 < \tilde{\mu} < q/4$ and 1 otherwise.

Theorem 5.12 (Completeness). If Π_{NIZK} is complete, then Construction 5.11 is complete.

Proof. Take any $\lambda, \tau \in \mathbb{N}$, any $N \leq 2^{\lambda}$, any $i \in [N]$, and policy f. Take any pp = $(\operatorname{crs}_{NIZK}, \operatorname{pp}_{com}, \mathbf{B}_0, \mathbf{p})$ in the support of Setup $(1^{\lambda}, 1^{\tau}, N)$ and any (pk, sk) in the support of KeyGen(pp, i, f). By construction, we can parse pp_{com} = (**B**, **W**, **T**) and pk = (**t**, π) where

$$\begin{aligned} \mathbf{d}_{i} &= H_{1}(i) \in \mathbb{Z}_{q}^{n} \\ \mathbf{V}_{\ell m} &= \operatorname{Ver}^{\mathrm{mat}}(\mathrm{pp}_{\mathrm{com}}, 1^{\ell m}) \in \mathbb{Z}_{q}^{m \times \ell m} \\ \mathbf{A} &= -\mathbf{B}_{0}\mathbf{V}_{\ell m} \in \mathbb{Z}_{q}^{n \times \ell m} \\ \mathbf{A}_{f} &= \operatorname{EvalF}(\mathbf{A}, f) \in \mathbb{Z}_{q}^{n \times m} \\ \mathbf{t} &= \mathbf{Br} + \mathbf{A}_{f}\mathbf{G}^{-1}(\mathbf{d}_{i}) + \mathbf{p} \in \mathbb{Z}_{q}^{n} \\ \pi \leftarrow \operatorname{NIZK}.\operatorname{Prove}(C_{\mathrm{valid}}, (i, \mathbf{A}_{f}, \mathbf{d}_{i}, \mathbf{p}, \mathbf{B}, \mathbf{t}), \mathbf{r}), \end{aligned}$$

and $\mathbf{r} \in \{0, 1\}^m$. By design, $C_{\text{valid}}((i, \mathbf{A}_f, \mathbf{d}_i, \mathbf{p}, \mathbf{B}, \mathbf{t}), \mathbf{r}) = 1$ so by completeness of Π_{NIZK} ,

NIZK.Verify
$$(C_{\text{valid}}, (i, \mathbf{A}_f, \mathbf{d}_i, \mathbf{p}, \mathbf{B}, \mathbf{t}), \pi) = 1.$$

Correspondingly, IsValid(pp, *i*, *f*, pk) outputs 1 and the claim holds.

Theorem 5.13 (Correctness). Suppose $m \ge n \log q$ and suppose $q > m^{O(d)} \cdot \chi \cdot (\sigma_{td} \cdot \ell \log \ell + \sigma_{td} \cdot N \log N + \sigma_{agg})$.

Proof. Take $\lambda, \tau \in \mathbb{N}, N \leq 2^{\lambda}$, any index $i \in [N]$, any policy $f_i \in \mathcal{F}_{\tau}$, any attribute $\mathbf{x} \in \{0, 1\}^{\ell}$ where $f_i(\mathbf{x}) = 0$, any pp in the support of Setup $(1^{\lambda}, 1^{\tau}, N)$, any $(\mathsf{pk}_i, \mathsf{sk}_i)$ in the support of KeyGen (pp, i, f_i) , any set $\{(j, f_j, \mathsf{pk}_j)\}_{j \neq i}$ where IsValid $(\mathsf{pp}, j, f_j, \mathsf{pk}_j) = 1$ for all $j \neq i$, and any message $\mu \in \{0, 1\}$. Let

$$(\mathsf{mpk}, \mathsf{hsk}_1, \dots, \mathsf{hsk}_N) = \mathsf{Aggregate}(\mathsf{pp}, \mathsf{pk}_1, \dots, \mathsf{pk}_N)$$
$$\mathsf{ct} \leftarrow \mathsf{Encrypt}(\mathsf{mpk}, \mathbf{x}, \mu).$$

Then the following holds:

- First, pp = (crs_{NIZK}, pp_{com}, **B**₀, **p**) where pp_{com} = (**B**, **W**, **T**). By construction, $[I_{2m^2} \otimes \mathbf{B} | \mathbf{W}]\mathbf{T} = I_{2m^2} \otimes \mathbf{G}$ and $||\mathbf{T}|| \le \sqrt{m}\sigma_{td}$.
- Let $\mathbf{V}_{\ell m} = \operatorname{Ver}^{\operatorname{mat}}(\operatorname{pp}_{\operatorname{com}}, 1^{\ell m}), \mathbf{A} = -\mathbf{B}_0 \mathbf{V}_{\ell m}$. For each $j \in [N]$, let $\mathbf{A}_{f_i} = \operatorname{EvalF}(\mathbf{A}, f_j)$ and $\mathbf{d}_j = H_1(j)$.
- Since (pk_i, sk_i) is in the support of KeyGen (pp, i, f_i) , we can write $pk_i = (t_i, \pi_i)$ where

$$\mathbf{t}_i = \mathbf{B}\mathbf{r}_i + \mathbf{A}_{f_i}\mathbf{G}^{-1}(\mathbf{d}_i) + \mathbf{p}$$

and $\mathbf{sk}_i = (\mathbf{t}_i, \mathbf{r}_i)$ for some $\mathbf{r}_i \in \{0, 1\}^m$.

• Consider the quantities constructed by the aggregation algorithm. For each $j \in [N]$, the aggregation algorithm computes

$$C_j = Commat(ppcom, \mathbf{u}_j^{\mathsf{T}} \otimes \mathbf{t}_j)$$

$$Z_j = Openmat(ppcom, \mathbf{u}_j^{\mathsf{T}} \otimes \mathbf{t}_j).$$

The aggregation algorithm parses $Z_j = [z_{j,1} | \cdots | z_{j,N}]$. Then it computes

$$\xi = (pp, (pk_1, f_1), \dots, (pk_N, f_N))$$
$$(C_0, \mathbf{z}_{0,1}, \dots, \mathbf{z}_{0,N}) = Sample(pp_{com}, 1^{\lambda_{DGS}}, 1^N, \sigma_{agg}; H_2(\xi)).$$

The master public key mpk and helper decryption key hsk_i is then defined to be

$$\mathsf{mpk} = \widehat{\mathsf{C}} = \mathsf{C}_0 + \sum_{j \in [N]} \mathsf{C}_j \quad \text{and} \quad \mathsf{hsk}_i = \widehat{\mathsf{z}}_i = \mathsf{z}_{0,i} + \sum_{j \in [N]} \mathsf{z}_{j,i}$$

• Next, consider the ciphertext ct. By definition,

$$\mathsf{ct} = (\mathbf{c}_1^\mathsf{T}, \mathbf{c}_2^\mathsf{T}, \mathbf{c}_3^\mathsf{T}, \mathbf{c}_4) = \left(\mathbf{s}^\mathsf{T}\mathbf{B} + \mathbf{e}^\mathsf{T}, \ \mathbf{s}^\mathsf{T}\widehat{\mathbf{C}} + \mathbf{e}^\mathsf{T}\mathbf{R}_{\widehat{\mathbf{C}}}, \ \mathbf{s}^\mathsf{T}(\mathbf{B}_0 + \mathbf{C}_{\mathbf{x}}) + \mathbf{e}^\mathsf{T}\mathbf{R}_{\mathbf{B}_0}, \ \mathbf{s}^\mathsf{T}\mathbf{p} + \mathbf{e}^\mathsf{T}\mathbf{r}_{\mathbf{p}} + \lfloor q/2 \rfloor \cdot \mu.\right),$$

where $\mathbf{C}_{\mathbf{x}} = \operatorname{Com}^{\mathrm{mat}}(\operatorname{pp}_{\mathrm{com}}, \mathbf{x}^{\mathsf{T}} \otimes \mathbf{G}), \mathbf{s} \in \mathbb{Z}_{q}^{n}, \mathbf{e} \leftarrow D_{\mathbb{Z},\chi}^{m}, \mathbf{R}_{\widehat{\mathbf{C}}}, \mathbf{R}_{\mathbf{B}_{0}} \in \{0,1\}^{m \times m} \text{ and } \mathbf{r}_{\mathbf{p}} \in \{0,1\}^{m}.$

Consider now Decrypt(sk, hsk_i, f_i, **x**, ct):

• Let $\mathbf{V}_{\ell m} = \operatorname{Ver}^{\operatorname{mat}}(\operatorname{pp}_{\operatorname{com}}, 1^{\ell m})$ and $\mathbf{Z}_{\mathbf{x}} = \operatorname{Open}^{\operatorname{mat}}(\operatorname{pp}_{\operatorname{com}}, \mathbf{x}^{\mathsf{T}} \otimes \mathbf{G})$. By Remark 3.9, we have

 $\mathbf{C}_{\mathbf{x}} \cdot \mathbf{V}_{\ell m} = \mathbf{x}^{\mathsf{T}} \otimes \mathbf{G} - \mathbf{B} \cdot \mathbf{Z}_{\mathbf{x}}.$

Recalling that $\mathbf{A} = -\mathbf{B}_0 \mathbf{V}_{\ell m}$, this means

$$\begin{bmatrix} \mathbf{B} \mid \mathbf{B}_0 + \mathbf{C}_{\mathbf{x}} \end{bmatrix} \cdot \begin{bmatrix} -\mathbf{Z}_{\mathbf{x}} \\ -\mathbf{V}_{\ell m} \end{bmatrix} = -\mathbf{B} \cdot \mathbf{Z}_{\mathbf{x}} - \mathbf{B}_0 \cdot \mathbf{V}_{\ell m} - \mathbf{C}_{\mathbf{x}} \cdot \mathbf{V}_{\ell m} = \mathbf{A} - \mathbf{x}^{\mathsf{T}} \otimes \mathbf{G}_{\mathbf{x}}$$

Let $H_{A,f_i,x}$ = EvalFX(A, f_i, x). Since $f_i(x) = 0$, we appeal to Theorem 3.6 to conclude that

$$(\mathbf{A} - \mathbf{x}^{\mathsf{T}} \otimes \mathbf{G}) \cdot \mathbf{H}_{\mathbf{A}, f_i, \mathbf{x}} = \mathbf{A}_{f_i} - f_i(\mathbf{x}) \cdot \mathbf{G} = \mathbf{A}_{f_i}.$$

Thus,

$$\tilde{c}_{3}^{\mathsf{T}} = [\mathbf{c}_{1}^{\mathsf{T}} \mid \mathbf{c}_{3}^{\mathsf{T}}] \begin{bmatrix} -\mathbf{Z}_{\mathbf{x}} \\ -\mathbf{V}_{\ell m} \end{bmatrix} \cdot \mathbf{H}_{\mathbf{A}, f_{i}, \mathbf{x}} \mathbf{G}^{-1}(\mathbf{d}_{i})$$

$$= \mathbf{s}^{\mathsf{T}} [\mathbf{B} \mid \mathbf{B}_{0} + \mathbf{C}_{\mathbf{x}}] \begin{bmatrix} -\mathbf{Z}_{\mathbf{x}} \\ -\mathbf{V}_{\ell m} \end{bmatrix} \cdot \mathbf{H}_{\mathbf{A}, f_{i}, \mathbf{x}} \mathbf{G}^{-1}(\mathbf{d}_{i}) + \underbrace{\mathbf{e}^{\mathsf{T}} [\mathbf{I}_{m} \mid \mathbf{R}_{\mathbf{B}_{0}}] \begin{bmatrix} -\mathbf{Z}_{\mathbf{x}} \\ -\mathbf{V}_{\ell m} \end{bmatrix} \cdot \mathbf{H}_{\mathbf{A}, f_{i}, \mathbf{x}} \mathbf{G}^{-1}(\mathbf{d}_{i})}_{\tilde{e}_{1}}$$

$$= \mathbf{s}^{\mathsf{T}} \mathbf{A}_{f_{i}} \mathbf{G}^{-1}(\mathbf{d}_{i}) + \tilde{e}_{1}.$$

• Let $\mathbf{V}_N = [\mathbf{v}_1 | \cdots | \mathbf{v}_N] = \text{Ver}^{\text{mat}}(\text{pp}_{\text{com}}, 1^N)$. For all $j \in [N]$, \mathbf{C}_j is a commitment to $\mathbf{u}_j^{\mathsf{T}} \otimes \mathbf{t}_j$ and \mathbf{Z}_j is an opening to $\mathbf{u}_j^{\mathsf{T}} \otimes \mathbf{t}_j$. By Remark 3.9, $\mathbf{C}_j \mathbf{V}_N = (\mathbf{u}_j^{\mathsf{T}} \otimes \mathbf{t}_j) - \mathbf{B}\mathbf{Z}_j$, and in particular,

$$\begin{aligned} \forall j \neq i : \mathbf{C}_{j} \mathbf{v}_{i} &= -\mathbf{B} \mathbf{z}_{j,i} \\ \mathbf{C}_{i} \mathbf{v}_{i} &= \mathbf{t}_{i} - \mathbf{B} \mathbf{z}_{i,i} = \mathbf{B} \mathbf{R}_{i} + \mathbf{A}_{f_{i}} \mathbf{G}^{-1}(\mathbf{d}_{i}) + \mathbf{p} - \mathbf{B} \mathbf{z}_{i,i}. \end{aligned}$$

By Theorem 5.9, we have that $\mathbf{Bz}_{0,j} = -\mathbf{C}_0 \mathbf{v}_j$ for all $j \in [N]$. Combining these two relations, we conclude

$$\widehat{\mathbf{C}}\mathbf{v}_i + \mathbf{B}\widehat{\mathbf{z}}_i = \mathbf{C}_0\mathbf{v}_i + \sum_{j \in [N]} \mathbf{C}_j\mathbf{v}_i + \mathbf{B}\mathbf{z}_{0,i} + \sum_{j \in [N]} \mathbf{B}\mathbf{z}_{j,i} = \mathbf{C}_i\mathbf{v}_i = \mathbf{B}\mathbf{r}_i + \mathbf{A}_{f_i}\mathbf{G}^{-1}(\mathbf{d}_i) + \mathbf{p}.$$

Thus,

$$\mathbf{c}_1^{\mathsf{T}}(\mathbf{r}_i - \hat{\mathbf{z}}_i) - \mathbf{c}_2^{\mathsf{T}} \mathbf{v}_i + \tilde{c}_3 = \mathbf{s}^{\mathsf{T}} (\mathbf{B} \mathbf{r}_i - \mathbf{B} \hat{\mathbf{z}}_i - \widehat{\mathbf{C}} \mathbf{v}_i + \mathbf{A}_{f_i} \mathbf{G}^{-1}(\mathbf{d}_i)) + \tilde{e}_2$$
$$= -\mathbf{s}^{\mathsf{T}} \mathbf{p} + \tilde{e}_2$$

where

$$\tilde{e}_2 = \mathbf{e}^{\mathsf{T}} \mathbf{r}_i - \mathbf{e}^{\mathsf{T}} \hat{\mathbf{z}}_i - \mathbf{e}^{\mathsf{T}} \mathbf{R}_{\widehat{\mathbf{C}}} \mathbf{v}_i + \tilde{e}_1$$

Thus, we have

$$\tilde{\mu} = c_4 + \mathbf{c}_1^{\mathsf{T}}(\mathbf{r}_i - \hat{\mathbf{z}}_i) - \mathbf{c}_2^{\mathsf{T}}\mathbf{v}_i + \tilde{c}_3$$

= $\mathbf{s}^{\mathsf{T}}\mathbf{p} + \mathbf{e}^{\mathsf{T}}\mathbf{r}_{\mathbf{p}} + \lfloor q/2 \rfloor \cdot \mu - \mathbf{s}^{\mathsf{T}}\mathbf{p} + \tilde{e}_2$
= $\lfloor q/2 \rfloor \cdot \mu + \mathbf{e}^{\mathsf{T}}\mathbf{r}_{\mathbf{p}} + \tilde{e}_2.$

As long as $|\mathbf{e}^{\mathsf{T}}\mathbf{r}_{\mathbf{p}} + \tilde{e}_2| < q/4$, correctness holds.

To complete, the proof, it suffices to bound the error $|\mathbf{e}^{\mathsf{T}}\mathbf{r}_{\mathbf{p}} - \tilde{e}_2|$.

- By definition, we have $\|\mathbf{T}\| \leq \sqrt{m}\sigma_{\mathsf{td}}$ and $\|\mathbf{e}\| \leq \sqrt{m}\chi$. In addition, $\mathbf{R}_{\widehat{\mathsf{C}}}, \mathbf{R}_{\mathbf{B}_0} \in \{0, 1\}^{m \times m}$ and $\mathbf{r}_i, \mathbf{r}_p \in \{0, 1\}^m$.
- By Remark 3.9, for all $j \in [N]$, we have

$$\begin{aligned} \|\mathbf{V}_{\ell m}\|, \|\mathbf{V}_{N}\| &\leq O(\|\mathbf{T}\| \cdot m^{4} \log q) = O(\sigma_{td} m^{9/2} \log q) \\ \|\mathbf{Z}_{\mathbf{x}}\| &\leq O(\|\mathbf{T}\| \cdot m^{7} \log q \log(\ell m)) = O(\sigma_{td} m^{15/2} \log q \log(\ell m)) \\ \|\mathbf{Z}_{\mathbf{j}}\| &\leq O(\|\mathbf{T}\| \cdot m^{7} \log q \log N) = O(\sigma_{td} m^{15/2} \log q \log N). \end{aligned}$$

From Theorem 5.9, we have that $\|\mathbf{z}_{0,i}\| \leq \sqrt{m}\sigma_{\text{agg}}$. Since $\hat{\mathbf{z}}_i = \mathbf{z}_{0,i} + \sum_{j \in [N]} \mathbf{z}_{j,i}$, this means

$$\|\hat{\mathbf{z}}_{i}\| \leq \|\mathbf{z}_{0,i}\| + \sum_{j \in [N]} \|\mathbf{z}_{j,i}\| \leq \sqrt{m}\sigma_{\text{agg}} + N \cdot O(\sigma_{\text{td}}m^{15/2}\log q\log N).$$

- By Lemma 3.5, we have $\|\mathbf{H}_{\mathbf{A},f_i,\mathbf{x}}\| \le m^{O(d)}$.
- By definition, $\tilde{e}_1 = \mathbf{e}^{\mathsf{T}}[\mathbf{I}_m \mid \mathbf{R}_{\mathbf{B}_0}] \begin{bmatrix} -\mathbf{Z}_{\mathbf{x}} \\ -\mathbf{V}_{\ell m} \end{bmatrix} \cdot \mathbf{H}_{\mathbf{A}, f_i, \mathbf{x}} \mathbf{G}^{-1}(\mathbf{d}_i)$. This means

$$\begin{split} |\tilde{e}_1| &\leq (\sqrt{m}\chi) \cdot m \cdot O(\sigma_{\rm td}m^{15/2}\log q\log(\ell m)) \cdot (2m) \cdot m^{O(d)} \cdot \ell m \cdot m \\ &\leq m^{O(d)} \cdot \sigma_{\rm td}\ell\chi \log q\log\ell. \end{split}$$

• Next $\tilde{e}_2 = \mathbf{e}^{\mathsf{T}} \mathbf{r}_i - \mathbf{e}^{\mathsf{T}} \hat{\mathbf{z}}_i - \mathbf{e}^{\mathsf{T}} \mathbf{R}_{\widehat{\mathbf{C}}} \mathbf{v}_i + \tilde{e}_1$. Then,

$$|\tilde{e}_2| \le m^{O(d)} \cdot \sigma_{\mathsf{td}} \ell \chi \log q \log \ell + m^2 \chi \sigma_{\mathsf{agg}} + N \cdot O(\sigma_{\mathsf{td}} \chi m^9 \log q \log N).$$

Finally, $|\mathbf{e}^{\mathsf{T}}\mathbf{r}_{\mathbf{p}}| \leq m\sqrt{m}\chi$.

If $m \ge n \log q$, setting

 $q > m^{O(d)} \cdot \chi \cdot (\sigma_{td} \cdot \ell \log \ell + \sigma_{td} \cdot N \log N + \sigma_{agg})$

suffices for correctness.

Theorem 5.14 (Attribute-Selective Security without Corruptions). Suppose Π_{NIZK} is complete, zero-knowledge, and simulation-sound extractable, and that the $(2m^2, \sigma_{td})$ -succinct LWE assumption holds with parameters (n, m, q, χ) . Suppose also $n \ge \lambda$, $m \ge O(n \log q)$, q > 2 is prime, $\chi \ge \log m$, $\sigma_{td} > O(m^3 \log m)$, $\sigma_{agg} > \lambda^{\omega(1)} \cdot \sigma_{td} \cdot m^{O(d)} \cdot N^3 \log \ell$, and $\lambda_{DGS} > \log \sigma_{agg}$. Then, Construction 5.11 satisfies attribute-selective security without corruptions.

Proof. Our proof follows a similar structure as the corresponding proof from [CHW25]:

- First, we replace the NIZK proofs in the keys sampled by the challenger with simulated proofs. For the adversarially-generated keys, the reduction extracts an associated secret key.
- Next, we change the distribution of the honest keys. Instead of sampling them as $\mathbf{t} = \mathbf{Br} + \mathbf{A}_f \mathbf{G}^{-1}(\mathbf{d}_i) + \mathbf{p}$, they are replaced with $\mathbf{t} = \mathbf{Br} + \mathbf{d}_i$. In particular, the reduction algorithm no longer has a secret key for the honestly-generated public keys.
- Next, we program the challenge attribute x into the public parameters (i.e., we set $B_0 BR_{B_0} C_x$, where C_x is a commitment to $x^T \otimes G$ and x is the challenge attribute).
- Then we use the explainable re-randomizer to program the challenge public keys into C_0 . Specifically, we choose C_0 such that the aggregated key \widehat{C} can be written as $\widehat{C} = C_0 + \sum_{i \in [N]} C_i^* = BR_{C_0}$, where C_i^* is the commitment to the (adversarially-chosen) public key for slot *i*. Just as in [CHW25], the re-randomization term C_0 is chosen to "cancel" out the adversarially-chosen public keys. This is the critical step that enables the reduction algorithm (to succinct LWE) to simulate the challenge ciphertext.
- Finally, we rely on succinct LWE to argue that the challenge ciphertext is pseudorandom.

We now give the formal argument. Take polynomials $\tau = \tau(\lambda)$ and $N = N(\lambda)$ and let \mathcal{A} be an efficient adversary for the attribute-selective security game (without corruptions). In the security analysis, we model the hash functions H_1 and H_2 as random oracles. For ease of exposition, we assume that \mathcal{A} has the following properties:

- Algorithm \mathcal{A} does not query H_1 or H_2 on the same input more than once.
- Algorithm \mathcal{A} always queries H_2 on the tuple $\xi^* = (pp, (pk_1, f_1), \dots, (pk_N, f_N))$ associated with the challenge query before entering the challenge phase.

Observe that both properties hold without loss of generality. Namely, any efficient algorithm \mathcal{A} that does not satisfy these properties can be generically compiled into an algorithm that does. Finally, let Q_{ro} be a bound on the number of random oracle queries algorithm \mathcal{A} makes (to either H_1 or H_2).

Hybrid experiments. We now define the sequence of hybrid experiments for our security analysis. Each experiment is parameterized by a bit $b \in \{0, 1\}$ and a precision parameter $\kappa = \kappa(\lambda)$ (i.e., the input to the explainable sampling procedure). For simplicity of notation, we omit the index κ when the behavior of the experiment is independent of the choice of κ .

- $Hyb_0^{(b)}$: This is the semi-malicious attribute-selective security experiment with challenge bit b:
 - Setup phase: On input the security parameter 1^{λ} , the policy-family parameter 1^{τ} , and the number of slots *N*, algorithm \mathcal{A} outputs the challenge attribute $\mathbf{x} \in \{0, 1\}^{\ell(\tau)}$. The challenger samples

$$\begin{aligned} \operatorname{crs}_{\mathsf{NIZK}} &\leftarrow \mathsf{NIZK}.\mathsf{Setup}(1^{\lambda}) \\ (\mathbf{B},\mathbf{T}_{\mathbf{B}}) &\leftarrow \mathsf{TrapGen}(1^{n},1^{m},q), \mathbf{W} \xleftarrow{\mathbb{R}} \mathbb{Z}_{q}^{2m^{2}n \times m} \\ \mathbf{T} &\leftarrow \mathsf{SamplePre}([\mathbf{I}_{2m^{2}} \otimes \mathbf{B} \mid \mathbf{W}], \begin{bmatrix} \mathbf{I}_{2m^{2}} \otimes^{\mathsf{T}_{\mathbf{B}}} \end{bmatrix}, \mathbf{I}_{2m^{2}} \otimes \mathbf{G}, \sigma_{\mathsf{td}}) \\ \mathbf{B}_{0} \xleftarrow{\mathbb{R}} \mathbb{Z}_{a}^{n \times m}, \mathbf{p} \xleftarrow{\mathbb{R}} \mathbb{Z}_{a}^{n}. \end{aligned}$$

If $||\mathbf{T}|| > \sqrt{m}\sigma_{td}$, then the challenger sets $\mathbf{T} = \begin{bmatrix} I_{2m^2 \otimes T_B} \\ \mathbf{0} \end{bmatrix}$. Let $pp_{com} = (\mathbf{B}, \mathbf{W}, \mathbf{T})$. The challenger gives the public parameters $pp = (crs_{NIZK}, pp_{com}, \mathbf{B}_0, \mathbf{p})$ to \mathcal{A} . The challenger also initializes a counter ctr = 0 and

a dictionary D. In addition, whenever \mathcal{A} queries H_1 on an index $i \in [N]$, the challenger responds with $\mathbf{d}_i \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^n$. Whenever algorithm \mathcal{A} queries H_2 on a string $\xi \in \{0, 1\}^*$, the challenger responds with a string $\gamma \stackrel{\mathbb{R}}{\leftarrow} \{0, 1\}^{\rho}$.

- Query phase: When \mathcal{A} makes a key-generation query on an index $i \in [N]$ and a function f, the challenger increments the counter ctr = ctr + 1. Then it computes

$$\mathbf{d}_{i} = H_{1}(i) \in \mathbb{Z}_{q}^{n}$$

$$\mathbf{V}_{\ell m} = \operatorname{Ver}^{\operatorname{mat}}(\operatorname{pp}_{\operatorname{com}}, 1^{\ell m}) \in \mathbb{Z}_{q}^{m \times \ell m}$$

$$\mathbf{A} = -\mathbf{B}_{0}\mathbf{V}_{\ell m} \in \mathbb{Z}_{q}^{n \times \ell m}$$

$$\mathbf{A}_{f} = \operatorname{EvalF}(\mathbf{A}, f) \in \mathbb{Z}_{q}^{n \times m}.$$
(5.7)

Then, it samples $\mathbf{r} \leftarrow \{0, 1\}^m$ and computes

$$\mathbf{t} = \mathbf{B}\mathbf{r} + \mathbf{A}_f \mathbf{G}^{-1}(\mathbf{d}_i) + \mathbf{p} \in \mathbb{Z}_a^n.$$

The challenger then computes $\pi \leftarrow \mathsf{NIZK}.\mathsf{Prove}(C_{\mathsf{valid}}, (i, \mathbf{A}_f, \mathbf{d}_i, \mathbf{p}, \mathbf{B}, \mathbf{t}), \mathbf{r})$ and responds with the public key $\mathsf{pk} = (\mathbf{t}, \pi)$. The challenger adds the mapping $\mathsf{ctr} \mapsto (i, f, \mathbf{t})$ to D.

- Challenge phase: Let $((c_1, f_1, pk_i), ..., (c_N, f_N, pk_N))$ be algorithm \mathcal{A} 's challenge query. For each $i \in [N]$, the challenger proceeds as follows:
 - * If $c_i \in \{1, ..., ctr\}$, then the challenger looks up $(i', f', pk') = D[c_i]$ and checks that i = i'. If not, the challenger outputs 0. Otherwise, the challenger sets $pk_i^* = pk' = (t_i, \pi_i)$ and $f_i^* = f'$.
 - * If $c_i = \bot$, then the challenger checks that $f_i(\mathbf{x}) = 1$. If so, it parses $pk_i = (\mathbf{t}_i, \pi_i)$. Then, it checks that IsValid(pp, *i*, *f_i*, pk_{*i*}) = 1. If so, the challenger sets $pk_i^* = pk_i$ and $f_i^* = f_i$. Otherwise, the challenger outputs 0.

Next, for each $i \in [N]$, the challenger computes $\mathbf{C}_i^* = \operatorname{Com}^{\operatorname{mat}}(\operatorname{pp}_{\operatorname{com}}, \mathbf{u}_i^{\mathsf{T}} \otimes \mathbf{t}_i)$ and the re-randomization matrix \mathbf{C}_0 as

$$\xi^{*} = (pp, (pk_{1}^{*}, f_{1}^{*}), \dots, (pk_{N}^{*}, f_{N}^{*}))$$

$$\gamma^{*} = H_{2}(\xi^{*})$$
(C₀, z_{0,1}, ..., z_{0,N}) = Sample(pp_{com}, 1^{\lambda}_{DGS}, 1^N, \sigma_{agg}; \gamma^{*}).
(5.8)

Finally, the challenger sets

$$\widehat{\mathbf{C}} = \mathbf{C}_0 + \sum_{i \in [N]} \mathbf{C}_i^*$$

Next, to construct the challenge ciphertext, the challenger samples the following:

$$\mathbf{s} \leftarrow^{\mathbb{R}} \mathbb{Z}_q^n$$
, $\mathbf{e} \leftarrow D_{\mathbb{Z},\chi}^m$, $\mathbf{R}_{\widehat{\mathbf{C}}}, \mathbf{R}_{\mathbf{B}_0} \leftarrow^{\mathbb{R}} \{0,1\}^{m \times m}$, $\mathbf{r}_p \leftarrow^{\mathbb{R}} \{0,1\}^m$.

If $\|\mathbf{e}\| > \sqrt{m\chi}$, it sets $\mathbf{e} = \mathbf{0}^m$. Next, compute $C_x = \text{Com}^{\text{mat}}(\text{pp}_{\text{com}}, \mathbf{x}^{\mathsf{T}} \otimes \mathbf{G}) \in \mathbb{Z}_q^{n \times m}$ and the ciphertext

$$\mathsf{ct} = \left(\mathbf{s}^{\mathsf{T}}\mathbf{B} + \mathbf{e}^{\mathsf{T}}, \ \mathbf{s}^{\mathsf{T}}\widehat{\mathbf{C}} + \mathbf{e}^{\mathsf{T}}\mathbf{R}_{\widehat{\mathbf{C}}}, \ \mathbf{s}^{\mathsf{T}}(\mathbf{B}_{0} + \mathbf{C}_{\mathbf{x}}) + \mathbf{e}^{\mathsf{T}}\mathbf{R}_{\mathbf{B}_{0}}, \ \mathbf{s}^{\mathsf{T}}\mathbf{p} + \mathbf{e}^{\mathsf{T}}\mathbf{r}_{\mathbf{p}} + \lfloor q/2 \rfloor \cdot b\right).$$

The challenger gives ct to \mathcal{A} .

- **Output phase:** At the end of the game, algorithm \mathcal{A} outputs a bit $b' \in \{0, 1\}$, which is the output of the experiment.
- $Hyb_1^{(b)}$: Same as $Hyb_0^{(b)}$ except at the beginning of the experiment, the challenger samples an index ind $\stackrel{\mathbb{R}}{\leftarrow} [Q_{ro}]$. In the challenge phase, after the challenger computes ξ^* according to Eq. (5.8), the challenger additionally checks that algorithm \mathcal{A} has made at least ind queries to H_2 , and if so, that the indth query $\xi_{ind} \in \{0, 1\}^*$ algorithm \mathcal{A} made to H_2 satisfies $\xi_{ind} = \xi^*$. If this is not the case, then the challenger halts with output 0.

- $Hyb_2^{(b)}$: Same as $Hyb_1^{(b)}$ except when responding to key-generation queries (on an index *i* and function *f*), the challenger outputs 0 if IsValid(pp, *i*, *f*, pk) = 0.
- Hyb₃^(b): Same as Hyb₂^(b) except the challenger replaces the NIZK proofs in the key-generation queries with simulated NIZK proofs:
 - In the setup phase, the challenger samples (crs_{NIZK}, td_{NIZK}) \leftarrow NIZK.TrapSetup(1^{λ}).
 - On each key-generation query, the challenger computes $\pi \leftarrow \mathsf{NIZK}.\mathsf{Sim}(\mathsf{td}_{\mathsf{NIZK}}, C_{\mathsf{valid}}, (i, \mathbf{A}_f, \mathbf{d}_i, \mathbf{p}, \mathbf{B}, \mathbf{t})).$
- $Hyb_4^{(b)}$: Same as $Hyb_3^{(b)}$ except the challenger extracts secret keys for the public keys associated with the indth query ξ_{ind} to H_2 . Specifically, the experiment proceeds as follows:
 - During the setup phase, the challenger initializes an empty dictionary D_{sk}.
 - Whenever \mathcal{A} makes a key-generation query on an index *i* and function *f*, after the challenger samples $\mathbf{r} \stackrel{\mathbb{R}}{\leftarrow} \{0, 1\}^m$ and computes $\mathbf{t} = \mathbf{Br} \mathbf{A}_f \mathbf{G}^{-1}(\mathbf{d}_i) \mathbf{p}$, the challenger adds the mapping $(i, \mathbf{A}_f, \mathbf{t}) \mapsto (0, \mathbf{r})$ to D_{sk} if $(i, \mathbf{A}_f, \mathbf{t})$ is not already contained in D_{sk} .

When responding to the indth query ξ_{ind} to H_2 , the challenger proceeds as follows:

- Parse $\xi_{ind} = (pp^*, (pk_1^*, f_1^*), ..., (pk_N^*, f_N^*))$, where $pk_i^* = (t_i^*, \pi_i^*)$. If ξ_{ind} does not have this form or $pp^* \neq pp$, then the challenger outputs 0. If IsValid(pp, *i*, $f_i^*, pk_i^*) = 0$ for any *i* ∈ [*N*], the challenger outputs 0.
- − For each $i \in [N]$, the challenger computes $A_{f_i^*} = \text{EvalF}(A, f_i^*)$. If $(i, A_{f_i^*}, t_i^*)$ is not contained in D_{sk} , the challenger first checks that $f_i^*(\mathbf{x}) = 1$. If not, the challenger outputs 0. Otherwise, the challenger computes

 $\mathbf{r}_{i}^{*} = \mathsf{NIZK}.\mathsf{Extract}(\mathsf{td}_{\mathsf{NIZK}}, C_{\mathsf{valid}}, (i, \mathbf{A}_{f_{i}^{*}}, \mathbf{d}_{i}, \mathbf{p}, \mathbf{B}, \mathbf{t}_{i}^{*}), \pi_{i}^{*}).$

If $\mathbf{r}_i^* \notin \{0,1\}^m$ or $\mathbf{Br}_i^* \neq \mathbf{t}_i^* + \mathbf{A}_{f_i^*}\mathbf{G}^{-1}(\mathbf{d}_i) + \mathbf{p}$, then the challenger outputs 0. Otherwise, the challenger adds the mapping $(i, \mathbf{A}_{f_i^*}, \mathbf{t}_i^*) \mapsto (1, \mathbf{r}_i^*)$ to D_{sk} .

If all of the checks pass, the challenger samples $\gamma^* \stackrel{\mathbb{R}}{\leftarrow} \{0, 1\}^{\rho}$ and responds with $H_2(\xi_{ind}) \coloneqq \gamma^*$. In this experiment, for every $pk_i^* = (\mathbf{t}_i^*, \pi_i^*)$ and accompanying policy f_i^* in ξ_{ind} , there is a mapping $(i, \mathbf{A}_{f_i^*}, \mathbf{t}_i^*) \mapsto (b_i^*, \mathbf{r}_i^*)$ in D_{sk} . The bit b_i^* indicates whether \mathbf{r}_i^* was sampled by the challenger (in response to a key-generation query) or if \mathbf{r}_i^* was chosen by the adversary.

• $Hyb_5^{(b)}$: Same as $Hyb_4^{(b)}$, except when sampling the public parameters, the challenger samples

 $\mathbf{B} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^{n \times m}$ and $\mathbf{T} \leftarrow [\mathbf{I}_{2m^2} \otimes \mathbf{B} \mid \mathbf{W}]_{\sigma_{\mathrm{td}}}^{-1}(\mathbf{I}_{2m^2} \otimes \mathbf{G}).$

- $Hyb_6^{(b)}$: Same as $Hyb_5^{(b)}$, except the challenger samples the matrices $\mathbf{R}_{\widehat{\mathbf{C}}}$, $\mathbf{R}_{\mathbf{B}_0} \stackrel{\mathbb{R}}{\leftarrow} \{0, 1\}^m \leftarrow \{0, 1\}^m$ during the setup phase instead. Then, the challenger computes $\mathbf{C}_{\mathbf{x}} = \operatorname{Com}^{mat}(\operatorname{pp}_{com}, \mathbf{x}^T \otimes \mathbf{G})$. It sets $\mathbf{B}_0 = \mathbf{B}\mathbf{R}_{\mathbf{B}_0} \mathbf{C}_{\mathbf{x}}$ and $\mathbf{p} = \mathbf{B}\mathbf{r}_p$. When responding to a key-generation query (on an index *i* and function *f*), the challenger instead sets $\mathbf{t} = \mathbf{Br} + \mathbf{d}_i$. The challenger still adds the mapping $(i, \mathbf{A}_f, \mathbf{t}) \mapsto (0, \mathbf{r})$ to D_{sk} as before.
- $Hyb_{7,\kappa}^{(b)}$: Same as $Hyb_6^{(b)}$, except when the adversary makes the indth query ξ_{ind} to H_2 , the challenger now samples

$$\begin{split} \gamma &\leftarrow^{\mathbb{R}} \{0,1\}^{\rho} \\ (C_0, \mathbf{z}_{0,1}, \dots, \mathbf{z}_{0,N}) &= \mathsf{Sample}(\mathsf{pp}_{\mathsf{com}}, 1^{\lambda_{\mathsf{DGS}}}, 1^N, \sigma_{\mathsf{agg}}; \gamma) \\ \gamma^* &\leftarrow \mathsf{Explain}(\mathsf{pp}_{\mathsf{com}}, 1^{\lambda_{\mathsf{DGS}}}, 1^\kappa, (C_0, \mathbf{z}_{0,1}, \dots, \mathbf{z}_{0,N}), \sigma_{\mathsf{agg}}) \end{split}$$

The challenger replies to \mathcal{A} with γ^* (i.e., implicitly setting $H_2(\xi_{ind}) \coloneqq \gamma^*$).

• $Hyb_{8,\kappa}^{(b)}$: Same as $Hyb_{7,\kappa}^{(b)}$, except when responding to the indth query ξ_{ind} to H_2 , the challenger now samples

$$\mathbf{C}_0 \xleftarrow{\mathbb{R}} \mathbb{Z}_q^{n \times m}$$
 and $\forall i \in [N] : \mathbf{z}_{0,i} \leftarrow \mathbf{B}_{\sigma_{a \sigma \sigma}}^{-1}(-\mathbf{C}_0 \mathbf{v}_i)$

where $\mathbf{V}_N = [\mathbf{v}_1 | \cdots | \mathbf{v}_N] = \text{Ver}^{\text{mat}}(\text{pp}_{\text{com}}, 1^N).$

• $Hyb_{9,\kappa}^{(b)}$: Same as $Hyb_{8,\kappa}^{(b)}$ except the challenger changes how it samples C_0 and $z_{0,i}$ when responding to the indth query ξ_{ind} to H_2 . Specifically, let $\xi_{ind} = (pp, (pk_1^*, f_1^*), \dots, (pk_N^*, f_N^*))$ where $pk_i^* = (t_i^*, \pi_i^*)$. If ξ_{ind} does not have this form, then the challenger outputs 0 (as in $Hyb_4^{(b)}$ and all subsequent hybrids). The challenger populates the dictionary D_{sk} usually the same procedure described in $Hyb_4^{(b)}$. Then, for each $i \in [N]$, the challenger computes

$$\mathbf{C}_{i}^{*} = \operatorname{Com}^{\mathrm{mat}}(\operatorname{pp}_{\mathrm{com}}, \mathbf{u}_{i}^{\mathsf{T}} \otimes \mathbf{t}_{i}^{*})$$
$$\mathbf{Z}_{i}^{*} = \operatorname{Open}^{\mathrm{mat}}(\operatorname{pp}_{\mathrm{com}}, \mathbf{u}_{i}^{\mathsf{T}} \otimes \mathbf{t}_{i}^{*})$$

The challenger parses $\mathbf{Z}_i^* = [\mathbf{z}_{i,1}^* | \cdots | \mathbf{z}_{i,N}^*]$. Next, the challenger sets $\mathbf{C}_0 = \mathbf{BR}_{\widehat{\mathbf{C}}} - \sum_{j \in [N]} \mathbf{C}_j^*$. By construction, if the challenger has not halted, we have for every $i \in [N]$, there exists a mapping $(i, \mathbf{A}_{f_i^*}, \mathbf{t}_i^*) \mapsto (b_i^*, \mathbf{r}_i^*)$ in D_{sk} . Then, for each $i \in [N]$, it defines $\mathbf{z}_{0,i}$ as follows:

- Suppose $(i, \mathbf{A}_{f_i^*}, \mathbf{t}_i^*) \mapsto (0, \mathbf{r}_i^*)$ in D_{sk} . Then, the challenger sets

$$\mathbf{z}_{0,i} \leftarrow \mathbf{B}_{\sigma_{\mathrm{agg}}}^{-1} \left(\mathbf{d}_i - \mathbf{B} \big(\mathbf{R}_{\widehat{\mathbf{C}}} \mathbf{v}_i - \mathbf{r}_i^* + \sum_{j \in [N]} \mathbf{z}_{j,i}^* \big) \right).$$

- Suppose $(i, \mathbf{A}_{f_i^*}, \mathbf{t}_i^*) \mapsto (1, \mathbf{r}_i^*)$ in D_{sk} . Then, the challenger sets

$$\mathbf{z}_{0,i} \leftarrow \mathbf{B}_{\sigma_{\text{agg}}}^{-1} \left(\mathbf{d}_i - \mathbf{B} \left(\mathbf{R}_{\widehat{\mathbf{C}}} \mathbf{v}_i - \mathbf{r}_i^* + (\mathbf{Z}_{\mathbf{x}} + \mathbf{R}_{\mathbf{B}_0} \mathbf{V}_{\ell m}) \mathbf{H}_{\mathbf{A}, f_i^*, \mathbf{x}} \mathbf{G}^{-1}(\mathbf{d}_i) - \mathbf{r}_{\mathbf{p}} + \sum_{j \in [N]} \mathbf{z}_{j,i}^* \right) \right).$$

The rest of the experiment proceeds as in Hyb $_{8\kappa}^{(b)}$.

• $\mathsf{Hyb}_{10,\kappa}^{(b)}$: Same as $\mathsf{Hyb}_{9,\kappa}^{(b)}$ except the challenger changes how it samples $z_{0,i}$ when responding to the indth query ξ_{ind} to H_2 . As usual, let $\xi_{ind} = (\mathsf{pp}, (\mathsf{pk}_1^*, f_1^*), \dots, (\mathsf{pk}_N^*, f_N^*))$ where $\mathsf{pk}_i^* = (\mathsf{t}_i^*, \pi_i^*)$.

- Suppose $(i, \mathbf{A}_{f_i^*}, \mathbf{t}_i^*) \mapsto (0, \mathbf{r}_i^*)$ in D_{sk} . Then, the challenger samples $\tilde{\mathbf{z}}_{0,i} \leftarrow \mathbf{B}_{\sigma_{\mathsf{age}}}^{-1}(\mathbf{d}_i)$ and sets

$$\mathbf{z}_{0,i} = \tilde{\mathbf{z}}_{0,i} - \mathbf{R}_{\widehat{\mathbf{C}}} \mathbf{v}_i + \mathbf{r}_i^* - \sum_{j \in [N]} \mathbf{z}_{j,i}^*$$

- Suppose $(i, \mathbf{A}_{f_i^*}, \mathbf{t}_i^*) \mapsto (1, \mathbf{r}_i^*)$ in D_{sk} . Then, the challenger samples $\tilde{\mathsf{z}}_{0,i} \leftarrow \mathbf{B}_{\sigma_{\mathsf{arg}}}^{-1}(\mathbf{d}_i)$ and sets

$$z_{0,i} = \tilde{z}_{0,i} - R_{\widehat{C}} v_i + r_i^* - (Z_x + R_{B_0} V_{\ell m}) H_{A, f_i^*, x} G^{-1}(d_i) + r_p - \sum_{j \in [N]} z_{j,i}^*.$$

- $\operatorname{Hyb}_{11,\kappa}^{(b)}$: Same as $\operatorname{Hyb}_{10,\kappa}^{(b)}$, except when simulating $\mathbf{d}_i = H_1(i)$, instead of sampling $\mathbf{d}_i \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^n$, the challenger instead samples $\tilde{\mathbf{z}}_{0,i} \leftarrow D_{\mathbb{Z},\sigma_{\operatorname{agg}}}^m$ and sets $\mathbf{d}_i = \mathbf{B}\tilde{\mathbf{z}}_{0,i}$. When answering the indth query to H_2 , the challenger uses these values of $\tilde{\mathbf{z}}_{0,i}$ for all $i \in [N]$ instead of sampling them.
- $Hyb_{12,\kappa}^{(b)}$: Same as $Hyb_{11,\kappa}^{(b)}$, except when sampling the public parameters, the challenger no longer checks the condition $||\mathbf{T}|| > \sqrt{m}\sigma_{td}$. Similarly, when constructing the challenge ciphertext ct^{*}, the challenger no longer checks if $||\mathbf{e}|| > \sqrt{m}\chi$.
- $\mathsf{Hyb}_{13,\kappa}^{(b)}$: Same as $\mathsf{Hyb}_{12,\kappa}^{(b)}$ except when simulating the challenge ciphertext, the challenger now samples $\tilde{\mathbf{c}} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^m$ and defines

$$ct = (\tilde{c}^{\dagger}, \tilde{c}^{\dagger}R_{\widehat{C}}, \tilde{c}^{\dagger}R_{B_0}, \tilde{c}^{\dagger}r_p + \lfloor q/2 \rfloor \cdot b).$$

- Hyb^(b)_{14, κ}: Same as Hyb^(b)_{13, κ} except when responding to the indth query ξ_{ind} to H_2 , the challenger samples $z_{0,i}$ as follows:
 - Suppose $(i, \mathbf{A}_{f_i^*}, \mathbf{t}_i^*) \mapsto (0, \mathbf{r}_i^*)$ in D_{sk} . Then, the challenger sets

$$\mathbf{z}_{0,i} \leftarrow \mathbf{B}_{\sigma_{\mathrm{agg}}}^{-1} \left(\mathbf{d}_i - \mathbf{B} \big(\mathbf{R}_{\widehat{\mathbf{C}}} \mathbf{v}_i - \mathbf{r}_i^* + \sum_{j \in [N]} \mathbf{z}_{j,i}^* \big) \right).$$

- Suppose $(i, \mathbf{A}_{f_i^*}, \mathbf{t}_i^*) \mapsto (1, \mathbf{r}_i^*)$ in D_{sk} . Then, the challenger sets

$$\mathbf{z}_{0,i} \leftarrow \mathbf{B}_{\sigma_{\text{agg}}}^{-1} \left(\mathbf{d}_i - \mathbf{B} \left(\mathbf{R}_{\widehat{\mathbf{C}}} \mathbf{v}_i - \mathbf{r}_i^* + (\mathbf{Z}_{\mathbf{x}} + \mathbf{R}_{B_0} \mathbf{V}_{\ell m}) \mathbf{H}_{\mathbf{A}, f_i^*, \mathbf{x}} \mathbf{G}^{-1}(\mathbf{d}_i) - \mathbf{r}_{\mathbf{p}} + \sum_{j \in [N]} \mathbf{z}_{j,i}^* \right) \right).$$

Note that this is equivalent to sampling

$$\mathbf{z}_{0,i} \leftarrow \mathbf{B}_{\sigma_{\text{agg}}}^{-1} \left(\mathbf{d}_i - \mathbf{B} \left(\mathbf{R}_{\widehat{\mathbf{C}}} \mathbf{v}_i - \mathbf{r}_i^* + (\mathbf{Z}_{\mathbf{x}} + \mathbf{R}_{\mathbf{B}_0} \mathbf{V}_{\ell m}) \mathbf{H}_{\mathbf{A}, f_i^*, \mathbf{x}} \mathbf{G}^{-1}(\mathbf{d}_i) + \sum_{j \in [N]} \mathbf{z}_{j,i}^* \right) + \mathbf{p} \right).$$

This latter expression only depends on p and not r_p .

• $Hyb_{15,\kappa}^{(b)}$: Same as $Hyb_{14,\kappa}^{(b)}$ except when simulating the challenger ciphertext, the challenger samples $c_4 \leftarrow \mathbb{Z}_q$ and outputs

$$\mathsf{ct} = \left(\widetilde{\mathbf{c}}^{\mathsf{T}} , \ \widetilde{\mathbf{c}}^{\mathsf{T}} \mathbf{R}_{\widehat{\mathbf{C}}} , \ \widetilde{\mathbf{c}}^{\mathsf{T}} \mathbf{R}_{\mathbf{B}_0} , \ \boldsymbol{c}_4
ight).$$

Notably, the challenger's behavior in this experiment is *independent* of the bit $b \in \{0, 1\}$.

We write $Hyb_i^{(b)}(\mathcal{A})$ to denote the distribution of the output of $Hyb_i^{(b)}$ with adversary \mathcal{A} . We now analyze each pair of adjacent distributions.

Lemma 5.15. For all $b \in \{0, 1\}$, $\Pr[Hyb_1^{(b)}(\mathcal{A}) = 1] = \frac{1}{Q_{ro}} \Pr[Hyb_0^{(b)}(\mathcal{A}) = 1]$.

Proof. By construction, the view of adversary \mathcal{A} is identical in $Hyb_0^{(b)}$ and $Hyb_1^{(b)}$. By assumption, algorithm \mathcal{A} is guaranteed to query H_2 on ξ^* prior to the challenge phase. Algorithm \mathcal{A} makes at most Q_{ro} queries to H_2 so let ind^{*} $\in [Q_{ro}]$ be the index of query ξ^* . Since all the queries \mathcal{A} makes to H_2 are distinct, with probability $1/Q_{ro}$ over the choice of ind $\stackrel{\mathbb{R}}{\leftarrow} [Q_{ro}]$, it will be the case that ind = ind^{*}. The output in experiment $Hyb_1^{(b)}$ is 1 if and only if ind = ind^{*} and the output in experiment $Hyb_0^{(b)}$ is 1. We conclude

$$\Pr[\mathsf{Hyb}_1^{(b)}(\mathcal{A}) = 1] = \Pr[\mathsf{Hyb}_0^{(b)}(\mathcal{A}) = 1 \land \mathsf{ind} = \mathsf{ind}^*] = \frac{1}{Q_{\mathsf{ro}}} \Pr[\mathsf{Hyb}_0^{(b)}(\mathcal{A}) = 1].$$

Lemma 5.16. Suppose Π_{NIZK} is complete. Then, for all $b \in \{0, 1\}$, $\text{Hyb}_1^{(b)}(\mathcal{A})$ and $\text{Hyb}_2^{(b)}(\mathcal{A})$ are identically distributed.

Proof. By Theorem 5.12, Construction 5.11 is complete so IsValid(pp, i, f, pk) = 1 for all pk in the support of KeyGen(pp, *i*, *f*). Thus, the additional abort condition never triggers and the two experiments are identical.

Lemma 5.17. Suppose Π_{NIZK} satisfies zero-knowledge. Then, for all $b \in \{0,1\}$, $Hyb_2^{(b)}(\mathcal{A})$ and $Hyb_3^{(b)}(\mathcal{A})$ are computationally indistinguishable.

Proof. Suppose $|\Pr[Hyb_2^{(b)}(\mathcal{A}) = 1] - \Pr[Hyb_3^{(b)}(\mathcal{A}) = 1]| = \varepsilon$ for some non-negligible ε . We use \mathcal{A} to construct an efficient adversary \mathcal{B} that breaks zero-knowledge of Π_{NIZK} :

• Setup phase: At the beginning of the game, algorithm \mathcal{B} receives a common reference string crs_{NIZK} from the zero-knowledge challenger. Algorithm \mathcal{B} samples an index ind $\stackrel{\mathbb{R}}{\leftarrow} [Q_{ro}]$ and starts running \mathcal{A} on input

the security parameter 1^{λ} , the policy-family parameter 1^{τ} , and the number of slots *N*. Algorithm \mathcal{A} outputs the challenge attribute $\mathbf{x} \in \{0, 1\}^{\ell(\tau)}$. Algorithm \mathcal{B} now samples

$$(\mathbf{B}, \mathbf{T}_{\mathbf{B}}) \leftarrow \mathsf{TrapGen}(1^{n}, 1^{m}, q), \mathbf{W} \xleftarrow{\mathbb{R}} \mathbb{Z}_{q}^{2m^{2}n \times m}$$
$$\mathbf{T} \leftarrow \mathsf{SamplePre}([\mathbf{I}_{2m^{2}} \otimes \mathbf{B} \mid \mathbf{W}], \begin{bmatrix} \mathbf{I}_{2m^{2}} \otimes \mathbf{T}_{\mathbf{B}} \\ \mathbf{0} \end{bmatrix}, \mathbf{I}_{2m^{2}} \otimes \mathbf{G}, \sigma_{\mathsf{td}})$$
$$\mathbf{B}_{0} \xleftarrow{\mathbb{R}} \mathbb{Z}_{a}^{n \times m}, \mathbf{p} \xleftarrow{\mathbb{R}} \mathbb{Z}_{a}^{n}.$$

If $||\mathbf{T}|| > \sqrt{m}\sigma_{td}$, then algorithm \mathcal{B} sets $\mathbf{T} = \begin{bmatrix} \mathbf{I}_{2m^2 \otimes \mathbf{T}_B} \\ \mathbf{0} \end{bmatrix}$. Let $pp_{com} = (\mathbf{B}, \mathbf{W}, \mathbf{T})$. Algorithm \mathcal{B} gives the public parameters $pp = (\operatorname{crs}_{\mathsf{NIZK}}, pp_{com}, \mathbf{B}_0, \mathbf{p})$ to \mathcal{A} . Algorithm \mathcal{B} also initializes a counter $\operatorname{ctr} = 0$ and a dictionary D. In addition, whenever \mathcal{A} queries H_1 on an index $i \in [N]$, algorithm \mathcal{B} responds with $\mathbf{d}_i \stackrel{\mathcal{R}}{\leftarrow} \mathbb{Z}_q^n$. Whenever algorithm \mathcal{A} queries H_2 on a string $\xi \in \{0, 1\}^*$, algorithm \mathcal{B} responds with a string $\gamma \stackrel{\mathcal{R}}{\leftarrow} \{0, 1\}^{\rho}$.

- Query phase: When \mathcal{A} makes a key-generation query on an index $i \in [N]$ and a function f, algorithm \mathcal{B} increments the counter ctr = ctr + 1. Then it computes $\mathbf{d}_i, \mathbf{V}_\ell, \mathbf{A}, \mathbf{A}_f$ according to Eq. (5.7) and samples $\mathbf{r} \stackrel{\mathbb{R}}{\leftarrow} \{0, 1\}^m$. Algorithm \mathcal{B} computes $\mathbf{t} = \mathbf{Br} + \mathbf{A}_f \mathbf{G}^{-1}(\mathbf{d}_i) + \mathbf{p}$ and submits $(C_{\text{valid}}, (i, \mathbf{A}_f, \mathbf{d}_i, \mathbf{p}, \mathbf{B}, \mathbf{t}), \mathbf{r})$ to the zero-knowledge challenger. The zero-knowledge challenger replies with a proof π . Algorithm \mathcal{B} responds to \mathcal{A} with the public key $p\mathbf{k} = (\mathbf{t}, \pi)$ and also adds the mapping ctr $\mapsto (i, f, \mathbf{t})$ to D. Finally, algorithm \mathcal{B} checks if IsValid(pp, $i, f, p\mathbf{k}) = 0$ and outputs 0 if so.
- Challenge phase: Let $((c_1, f_1, pk_i), ..., (c_N, f_N, pk_N))$ be algorithm \mathcal{A} 's challenge query. For each $i \in [N]$, algorithm \mathcal{B} proceeds as follows:
 - If $c_i \in \{1, ..., ctr\}$, then algorithm \mathcal{B} looks up $(i', f', pk') = D[c_i]$ and checks that i = i'. If not, algorithm \mathcal{B} outputs 0. Otherwise, algorithm \mathcal{B} sets $pk_i^* = pk' = (t_i, \pi_i)$ and $f_i^* = f_i$.
 - If $c_i = \bot$, then algorithm \mathcal{B} checks that $f_i(\mathbf{x}) = 1$. If so, it parses $\mathsf{pk}_i = (\mathbf{t}_i, \pi_i)$. Then, it checks that NIZK.Verify($\mathsf{crs}_{\mathsf{NIZK}}, (i, \mathbf{A}_{f_i}, \mathbf{d}_i, \mathbf{p}, \mathbf{B}, \mathbf{t}_i), \pi_i$) = 1 where $\mathbf{A}_{f_i} = \mathsf{EvalF}(\mathbf{A}, f_i)$, and $\mathbf{d}_i = H_1(i)$). If the checks pass, algorithm \mathcal{B} sets $\mathsf{pk}_i^* = \mathsf{pk}_i$ and $f_i^* = f_i$. Otherwise, algorithm \mathcal{B} outputs 0.

Next, for each $i \in [N]$, algorithm \mathcal{B} computes $\mathbf{C}_i^* = \operatorname{Com}^{\max}(\operatorname{pp}_{\operatorname{com}}, \mathbf{u}_i^{\mathsf{T}} \otimes \mathbf{t}_i)$ and the re-randomization matrix \mathbf{C}_0 as

$$\xi^* = (pp, (pk_1^*, f_1^*), \dots, (pk_N^*, f_N^*))$$
$$\gamma^* = H_2(\xi^*)$$
$$(C_0, z_{0,1}, \dots, z_{0,N}) = \text{Sample}(pp_{\text{com}}, 1^{\lambda_{\text{DGS}}}, 1^N, \sigma_{\text{agg}}; \gamma^*).$$

If adversary \mathcal{A} has not yet made at least ind queries to H_2 , or if its indth query ξ_{ind} to H_2 satisfies $\xi_{ind} \neq \xi^*$, then algorithm \mathcal{B} outputs 0. Otherwise, algorithm \mathcal{B} computes

$$\widehat{\mathbf{C}} = \mathbf{C}_0 + \sum_{i \in [N]} \mathbf{C}_i^* , \ \mathbf{s} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^n , \ \mathbf{e} \leftarrow D_{\mathbb{Z},\chi}^m , \ \mathbf{R}_{\widehat{\mathbf{C}}}, \mathbf{R}_{\mathbf{B}_0} \stackrel{\mathbb{R}}{\leftarrow} \{0,1\}^{m \times m} , \ \mathbf{r}_p \stackrel{\mathbb{R}}{\leftarrow} \{0,1\}^m .$$

If $\|\mathbf{e}\| > \sqrt{m\chi}$, algorithm \mathcal{B} sets $\mathbf{e} = \mathbf{0}^m$. Next, algorithm \mathcal{B} computes $C_{\mathbf{x}} = \text{Com}^{\text{mat}}(\text{pp}_{\text{com}}, \mathbf{x}^{\mathsf{T}} \otimes \mathbf{G}) \in \mathbb{Z}_q^{n \times m}$ and the ciphertext

$$\mathsf{ct} = \left(\mathbf{s}^{\mathsf{T}}\mathbf{B} + \mathbf{e}^{\mathsf{T}}, \ \mathbf{s}^{\mathsf{T}}\widehat{\mathbf{C}} + \mathbf{e}^{\mathsf{T}}\mathbf{R}_{\widehat{\mathbf{C}}}, \ \mathbf{s}^{\mathsf{T}}(\mathbf{B}_{0} + \mathbf{C}_{\mathbf{x}}) + \mathbf{e}^{\mathsf{T}}\mathbf{R}_{\mathbf{B}_{0}}, \ \mathbf{s}^{\mathsf{T}}\mathbf{p} + \mathbf{e}^{\mathsf{T}}\mathbf{r}_{\mathbf{p}} + \lfloor q/2 \rfloor \cdot b\right).$$

Algorithm \mathcal{B} gives ct to \mathcal{A} .

• **Output phase:** At the end of the game, algorithm \mathcal{A} outputs a bit $b' \in \{0, 1\}$, which \mathcal{B} also outputs.

By construction, on every key-generation query, we have that

$$\mathbf{t} = \mathbf{B}\mathbf{r} + \mathbf{A}_f \mathbf{G}^{-1}(\mathbf{d}_i) + \mathbf{p},$$

so $C_{\text{valid}}((i, \mathbf{A}_f, \mathbf{d}_i, \mathbf{p}, \mathbf{B}, \mathbf{t}), \mathbf{r}) = 1$. We now consider two possibilities:

- Suppose the zero-knowledge challenger sampled $\operatorname{crs}_{NIZK} \leftarrow \operatorname{NIZK}.\operatorname{Setup}(1^{\lambda})$ and constructed the proofs π by setting $\pi \leftarrow \operatorname{NIZK}.\operatorname{Prove}(\operatorname{crs}_{NIZK}, C_{\operatorname{valid}}, (i, \mathbf{A}_f, \mathbf{d}_i, \mathbf{p}, \mathbf{B}, \mathbf{t}), \mathbf{r})$. Then, algorithm \mathcal{B} perfectly simulates an execution of $\operatorname{Hyb}_2^{(b)}$ and outputs 1 with probability $\operatorname{Pr}[\operatorname{Hyb}_2^{(b)}(\mathcal{A}) = 1]$.
- Suppose the zero-knowledge challenger sampled $(\operatorname{crs}_{NIZK}, \operatorname{td}_{NIZK}) \leftarrow \operatorname{NIZK}.\operatorname{TrapSetup}(1^{\lambda})$ and constructed the proof π by setting $\pi \leftarrow \operatorname{NIZK}.\operatorname{Sim}(\operatorname{td}_{NIZK}, C_{\operatorname{valid}}, (i, \mathbf{A}_f, \mathbf{d}_i, \mathbf{p}, \mathbf{B}, \mathbf{t}))$. Then, algorithm \mathcal{B} perfectly simulates an execution of $\operatorname{Hyb}_3^{(b)}$ and outputs 1 with probability $\Pr[\operatorname{Hyb}_3^{(b)}(\mathcal{A}) = 1]$.

We conclude that algorithm $\mathcal B$ breaks zero-knowledge with the same advantage ε .

Lemma 5.18. Suppose $n \ge \lambda$, $m \ge 2n \log q$, and q is prime. If Π_{NIZK} satisfies simulation-sound extractability, then for all $b \in \{0, 1\}$, $Hyb_3^{(b)}(\mathcal{A})$ and $Hyb_4^{(b)}(\mathcal{A})$ are computationally indistinguishable.

Proof. The only difference between $Hyb_3^{(b)}$ and $Hyb_4^{(b)}$ is the extra check that the challenger performs when responding to the indth query ξ_{ind} to H_2 . Specifically, the two experiments can only differ when the following occurs:

• Let $((c_1, f_1, pk_i), \dots, (c_N, f_N, pk_N))$ be adversary \mathcal{A} 's challenge query and ξ^* be the challenger's input to H_2 from Eq. (5.8). Then,

$$\xi_{\text{ind}} = \xi^* = (\text{pp}, (\text{pk}_1^*, f_1^*), \dots, (\text{pk}_N^*, f_N^*)),$$

where $pk_i^* = (t_i^*, \pi_i^*)$. Otherwise, the output in both experiments output 0. Moreover, it is also the case that $IsValid(pp, i, f_i^*, pk_i^*) = 1$ for all $i \in [N]$. To see this, we consider two cases:

- Suppose $c_i \in \{1, ..., ctr\}$. In this case, pk_i^* was obtained as the result of an honest key-generation query (on index *i* and function f_i^*). In both experiments, the challenger outputs 0 if IsValid(pp, *i*, f_i^* , pk_i^*) = 0.
- Suppose $c_i = \bot$. Then, the challenger in both experiments affirms that IsValid(pp, *i*, *f_i*, pk^{*}_i) = 1. In this case, the challenger in both experiments sets $f_i^* = f_i$. We conclude that IsValid(pp, *i*, f_i^* , pk^{*}_i) = 1.
- For each $i \in [N]$, let $\mathbf{A}_{f_i^*} = \text{EvalF}(\mathbf{A}, f_i^*)$. It must be the case that there exists some index $i \in [N]$ where $(i, \mathbf{A}_{f_i^*}, \mathbf{t}_i^*)$ is not contained in D_{sk} at the time algorithm \mathcal{A} queried H_2 on ξ_{ind} . If not, then the challenger's behavior in the two experiments is identical. Moreover, one of the following conditions must occur for one such index $i \in [N]$ where $(i, \mathbf{A}_{f_i^*}, \mathbf{t}_i^*)$ is not contained in D_{sk} :
 - $f_i^*(\mathbf{x}) \neq 1$; or
 - The extracted vector $\mathbf{r}_i^* = \text{NIZK}.\text{Extract}(\text{td}_{\text{NIZK}}, C_{\text{valid}}, (i, \mathbf{A}_{f_i^*}, \mathbf{d}_i, \mathbf{p}, \mathbf{B}, \mathbf{t}_i^*), \pi_i^*)$ satisfies either $\mathbf{r}_i^* \notin \{0, 1\}^m$ or $\mathbf{Br}_i^* \neq \mathbf{t}_i^* + \mathbf{A}_{f_i^*}\mathbf{G}^{-1}(\mathbf{d}_i) + \mathbf{p}$.

Suppose $|\Pr[Hyb_3^{(b)}(\mathcal{A}) = 1] - \Pr[Hyb_4^{(b)}(\mathcal{A}) = 1]| = \varepsilon$ for some non-negligible ε . Then, it must be the case that in an execution of $Hyb_3^{(b)}$, the above conditions hold with probability ε . We consider the two possibilities:

- Suppose there exists some index $i \in [N]$ where $(i, \mathbf{A}_{f_i^*}, \mathbf{t}_i^*)$ is not contained in D_{sk} and $f_i^*(\mathbf{x}) \neq 1$. In this case, the challenger in $\mathsf{Hyb}_4^{(b)}$ always outputs 0. We claim that this is the case with overwhelming probability in $\mathsf{Hyb}_3^{(b)}$. Recall also that $\xi_{\mathsf{ind}} = \xi^*$. By construction of ξ^* , we now have the following:
 - Suppose $c_i \in \{1, \ldots, \text{ctr}\}$. In this case, the challenger in $\text{Hyb}_3^{(b)}$ sets $\text{pk}_i^* = \text{pk'}$ where $(i', f', \text{pk'}) = D[c_i]$. By construction, $\text{pk}_i^* = (\mathbf{t}_i^*, \pi_i^*)$ is the public key the challenger generated when responding to a keygeneration query on index *i* and function $f' = f_i^*$. Also, i = i' as otherwise, the output in $\text{Hyb}_3^{(b)}$ is also 0. In $\text{Hyb}_3^{(b)}$, the challenger would then add the mapping $(i, \mathbf{A}_{f_i^*}, \mathbf{t}_i^*)$ to D_{sk} . However, since $(i, \mathbf{A}_{f_i^*}, \mathbf{t}_i^*)$ is not contained in D_{sk} at the time \mathcal{A} made its indth query to H_2 , this case can only happen if the challenger inserted the entry $(i, \mathbf{A}_{f_i^*}, \mathbf{t}_i^*)$ into D_{sk} after algorithm \mathcal{A} queried H_2 on ξ_{ind} . This can only happen if the challenger sampled \mathbf{t}_i^* as the public key in one of the subsequent key-generation queries. However, in a key-generation query, the challenger samples $\mathbf{r} \stackrel{\text{\tiny eff}}{\leftarrow} \{0, 1\}^m$ and then sets $\mathbf{t} = \mathbf{Br} + \mathbf{A}_f \mathbf{G}^{-1}(\mathbf{d}_i) + \mathbf{p}$. Since $n \ge \lambda$, $m \ge 2n \log q$, and q > 2 is prime, the marginal distribution of \mathbf{Br} is statistically close to

uniform over \mathbb{Z}_q^n . Thus, with overwhelming probability, it will be the case that $\mathbf{t} \neq \mathbf{t}_i^*$. Since the adversary can make at most a polynomial number of key-generation queries, we conclude by a union bound that the probability that the challenger samples \mathbf{t}^* in a key-generation query (after \mathcal{A} queries H_2 on ξ_{ind}) is negligible so this case occurs with negligible probability.

- Suppose $c_i = \bot$. In this case, the challenger in Hyb₃^(b) outputs 0 if $f_i^*(\mathbf{x}) = 1$.

We conclude that in this case, the challenger in $Hyb_3^{(b)}$ also outputs 0 with overwhelming probability.

• Since the challenger's behavior in $Hyb_3^{(b)}$ and $Hyb_4^{(b)}$ is identical with overwhelming probability when $f_i^*(\mathbf{x}) = 0$, it must be the case that with probability $\varepsilon - negl(\lambda)$, $f_i^*(\mathbf{x}) = 1$ and the extracted vector \mathbf{r}_i^* satisfies either $\mathbf{r}_i^* \notin \{0, 1\}^m$ or $\mathbf{t}_i^* \neq \mathbf{Br}_i^* + \mathbf{A}_{f_i^*}\mathbf{G}^{-1}(\mathbf{d}_i) + \mathbf{p}$. In particular, this means that $C_{valid}((i, \mathbf{A}_{f_i^*}, \mathbf{d}_i, \mathbf{p}, \mathbf{B}, \mathbf{t}_i^*), \mathbf{r}_i^*) = 0$. We show below that this implies an adversary \mathcal{B} that can break simulation-extractability of Π_{NIZK} with the same advantage $\varepsilon - negl(\lambda)$.

We now use \mathcal{A} to construct an adversary \mathcal{B} for the simulation-sound extractability game:

- At the beginning of the game, algorithm \mathcal{B} receives a common reference string $\operatorname{crs}_{NIZK}$ from the simulationsound extractability challenger. Algorithm \mathcal{B} samples an index ind $\stackrel{\mathbb{R}}{\leftarrow} [Q_{ro}]$ and starts running \mathcal{A} on input the security parameter 1^{λ} , the policy-family parameter 1^{τ} , and the number of slots N.
- Algorithm \mathcal{A} outputs the challenge attribute $\mathbf{x} \in \{0, 1\}^{\ell(\tau)}$. Algorithm \mathcal{B} now samples

$$(\mathbf{B}, \mathbf{T}_{\mathbf{B}}) \leftarrow \mathsf{TrapGen}(1^{n}, 1^{m}, q), \mathbf{W} \xleftarrow{\mathbb{R}} \mathbb{Z}_{q}^{2m^{2}n \times m}$$
$$\mathbf{T} \leftarrow \mathsf{SamplePre}([\mathbf{I}_{2m^{2}} \otimes \mathbf{B} \mid \mathbf{W}], \begin{bmatrix} \mathbf{I}_{2m^{2}} \otimes \mathbf{T}_{\mathbf{B}} \\ \mathbf{0} \end{bmatrix}, \mathbf{I}_{2m^{2}} \otimes \mathbf{G}, \sigma_{\mathsf{td}})$$
$$\mathbf{B}_{0} \xleftarrow{\mathbb{R}} \mathbb{Z}_{q}^{n \times m}, \mathbf{p} \xleftarrow{\mathbb{R}} \mathbb{Z}_{q}^{n}.$$

If $||\mathbf{T}|| > \sqrt{m}\sigma_{td}$, then algorithm \mathcal{B} sets $\mathbf{T} = \begin{bmatrix} I_{2m^2 \otimes T_B} \\ 0 \end{bmatrix}$. Let $pp_{com} = (\mathbf{B}, \mathbf{W}, \mathbf{T})$. Algorithm \mathcal{B} gives the public parameters $pp = (crs_{NIZK}, pp_{com}, \mathbf{B}_0, \mathbf{p})$ to \mathcal{A} . In addition, whenever \mathcal{A} queries H_1 on an index $i \in [N]$, algorithm \mathcal{B} responds with $\mathbf{d}_i \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^n$. Whenever algorithm \mathcal{A} queries H_2 on a string $\xi \in \{0, 1\}^*$, algorithm \mathcal{B} responds with a string $\gamma \stackrel{\mathbb{R}}{\leftarrow} \{0, 1\}^{\rho}$. Finally, algorithm \mathcal{B} also initializes an empty dictionary D_{sk} .

- Whenever \mathcal{A} makes a key-generation query on an index $i \in [N]$ and a function f, algorithm \mathcal{B} increments the counter ctr = ctr + 1. Then it computes $\mathbf{d}_i, \mathbf{V}_\ell, \mathbf{A}, \mathbf{A}_f$ according to Eq. (5.7) and samples $\mathbf{r} \leftarrow \{0, 1\}^m$. Algorithm \mathcal{B} computes $\mathbf{t} = \mathbf{Br} + \mathbf{A}_f \mathbf{G}^{-1}(\mathbf{d}_i) + \mathbf{p}$ and submits $(C_{\text{valid}}, (i, \mathbf{A}_f, \mathbf{d}_i, \mathbf{p}, \mathbf{B}, \mathbf{t}))$ to the simulation-sound extractability challenger and receives a proof π . Algorithm \mathcal{B} responds to \mathcal{A} with the public key $\mathbf{pk} = (\mathbf{t}, \pi)$ and adds the mapping $(i, \mathbf{A}_f, \mathbf{t}) \mapsto (0, \mathbf{r})$ to \mathbf{D}_{sk} if such a mapping does not already exist. Finally, algorithm \mathcal{B} also checks if $\mathbf{IsValid}(\mathbf{pp}, i, f, \mathbf{pk}) = 0$ and outputs 0 if so.
- When \mathcal{A} makes its indth query to H_2 , algorithm \mathcal{B} parses $\xi_{ind} = (pp^*, (pk_1^*, f_1^*), \dots, (pk_N^*, f_N^*))$, where $pk_i^* = (\mathbf{t}_i^*, \pi_i^*)$. If ξ_{ind} does not have this form or $pp^* \neq pp$, then algorithm \mathcal{B} outputs \perp . Otherwise, algorithm \mathcal{B} samples a random index $i^* \notin [N]$ and and outputs $(C_{valid}, (i, \mathbf{A}_{f_i^*}, \mathbf{d}_i, \mathbf{p}, \mathbf{B}, \mathbf{t}_{i^*}^*), \pi_{i^*}^*)$.
- If \mathcal{A} does not make indth queries to H_2 prior to the challenge phase, algorithm \mathcal{B} outputs \perp .

Since the challenger samples $(\operatorname{crs}_{NIZK}, \operatorname{td}_{NIZK}) \leftarrow NIZK$. TrapSetup (1^{λ}) and constructs the (simulated) proofs π as $\pi \leftarrow \operatorname{NIZK}.Sim(\operatorname{td}_{NIZK}, C_{\operatorname{valid}}, (i, \mathbf{A}_f, \mathbf{d}_i, \mathbf{p}, \mathbf{B}, \mathbf{t}))$, algorithm \mathcal{B} perfectly simulates an execution of $\operatorname{Hyb}_3^{(b)}$ and $\operatorname{Hyb}_4^{(b)}$ for \mathcal{A} . By assumption, with probability ε – negl (λ) , the query $\xi_{\operatorname{ind}} = (\operatorname{pp}^*, (\operatorname{pk}_1^*, f_1^*), \dots, (\operatorname{pk}_N^*, f_N^*))$ has the property that there exists an index $i \in [N]$ where

- $(i, \mathbf{A}_{f_i^*}, \mathbf{t}_i^*) \notin \mathbf{D}_{sk}$. This means algorithm \mathcal{B} did not query for a proof on $(i, \mathbf{A}_{f_i^*}, \mathbf{d}_i, \mathbf{p}, \mathbf{B}, \mathbf{t}_i^*)$.
- IsValid(pp^*, i, f_i^*, pk_i^*) = 1 which means NIZK.Verify($crs_{NIZK}, (i, A_{f_i^*}, d_i, p, B, t_i^*), \pi^*$) = 1.

• The extracted vector $\mathbf{r}_i^* = \text{NIZK}.\text{Extract}(\text{td}_{\text{NIZK}}, C_{\text{valid}}, (i, \mathbf{A}_{f_i^*}, \mathbf{d}_i, \mathbf{p}, \mathbf{B}, \mathbf{t}_i^*), \pi_i^*)$ satisfies either $\mathbf{r}_i^* \notin \{0, 1\}^m$ or $\mathbf{t}_i^* \neq \mathbf{Br}_i^* + \mathbf{A}_{f_i^*}\mathbf{G}^{-1}(\mathbf{d}_i) + \mathbf{p}$. In particular, this means $C_{\text{valid}}((i, \mathbf{A}_{f_i^*}, \mathbf{d}_i, \mathbf{p}, \mathbf{B}, \mathbf{t}_i^*), \mathbf{r}^*) = 0$.

Observe that these precisely coincide with the winning conditions in the simulation-extractability game. Thus, whenever $i^* = i$, then algorithm \mathcal{B} successfully breaks simulation-extractability. Since algorithm \mathcal{B} samples $i^* \stackrel{\mathbb{R}}{\leftarrow} [N]$, algorithm \mathcal{B} breaks simulation-extractability with advantage at least $(\varepsilon - \operatorname{negl}(\lambda))/N$, which is non-negligible since $N = \operatorname{poly}(\lambda)$.

Lemma 5.19. Suppose $n \ge \lambda$, $m \ge 3n \log q$, and $\sigma_{td} \ge O(m^3 \log m)$. Then, for all $b \in \{0,1\}$, $Hyb_4^{(b)}(\mathcal{A})$ and $Hyb_5^{(b)}(\mathcal{A})$ are statistically indistinguishable.

Proof. The lemma follows from Lemma 3.5. Specifically, the claim follows from Lemma 3.5 as long as

$$\sigma_{\rm td} \ge (2m^3 + m)\log(2m^2n) = O(m^3\log m).$$

Lemma 5.20. Suppose $n \ge \lambda$, $m \ge 2n \log q$, and q > 2 is prime. Then, for all $b \in \{0, 1\}$, $\mathsf{Hyb}_5^{(b)}(\mathcal{A})$ and $\mathsf{Hyb}_6^{(b)}(\mathcal{A})$ are statistically indistinguishable.

Proof. This follows by the leftover hash lemma (Lemma 3.1). First, consider the distribution of \mathbf{B}_0 in the two experiments. In $\mathsf{Hyb}_5^{(b)}$, the challenger samples $\mathbf{B}_0 \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^{n \times m}$ whereas in $\mathsf{Hyb}_6^{(b)}$, the challenger samples $\mathbf{R}_0 \stackrel{\mathbb{R}}{\leftarrow} \{0, 1\}^{m \times m}$ and sets $\mathbf{B}_0 = \mathbf{BR}_{\mathbf{B}_0} - \mathbf{C}_{\mathbf{x}}$. We claim that these two distributions are statistically close:

- First, suppose we sample $\mathbf{B}_0 = \mathbf{B}_0^* \mathbf{C}_{\mathbf{x}}$ where $\mathbf{B}_0^* \leftarrow^{\mathbb{R}} \mathbb{Z}_q^{n \times m}$. Since \mathbf{B}_0^* is sampled independently of $\mathbf{C}_{\mathbf{x}}$, the distribution of \mathbf{B}_0 remains uniform over $\mathbb{Z}_q^{n \times m}$.
- By Lemma 3.1, the distributions

 $(\mathbf{B}, \mathbf{B}\mathbf{R}_{\mathbf{B}_0}, \mathbf{e}^{\mathsf{T}}\mathbf{R}_{\mathbf{B}_0})$ and $(\mathbf{B}, \mathbf{B}_0^*, \mathbf{e}^{\mathsf{T}}\mathbf{R}_{\mathbf{B}_0})$

are statistically indistinguishable when $\mathbf{B} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^{n \times m}$, $\mathbf{R}_{\mathbf{B}_0} \stackrel{\mathbb{R}}{\leftarrow} \{0,1\}^{m \times m}$, and $\mathbf{B}_0^* \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^{n \times m}$.

Combining the above statements, we conclude that the distribution of B_0 in the two experiments are statistically indistinguishable. Next, consider the distribution of **p**. By Lemma 3.1, the distributions

 $(\mathbf{B}, \mathbf{Br}_{\mathbf{p}}, \mathbf{e}^{\mathsf{T}}\mathbf{r}_{\mathbf{p}})$ and $(\mathbf{B}, \mathbf{p}, \mathbf{e}^{\mathsf{T}}\mathbf{r}_{\mathbf{p}})$

are statistically indistinguishable when $\mathbf{B} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^{n \times m}$, $\mathbf{r_p} \stackrel{\mathbb{R}}{\leftarrow} \{0, 1\}^m$, and $\mathbf{p} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^n$. The left distribution maps to $\mathsf{Hyb}_6^{(b)}$. Finally, consider the key-generation queries:

- In Hyb₅^(b), the challenger sets $\mathbf{t} = \mathbf{Br} + \mathbf{A}_f \mathbf{G}^{-1}(\mathbf{d}_i) + \mathbf{p}$ where $\mathbf{r} \leftarrow \{0, 1\}^m$.
- In Hyb₆^(b), the challenger sets $\mathbf{t} = \mathbf{Br} + \mathbf{d}_i$ where $\mathbf{r} \leftarrow \{0, 1\}^m$.

By the Lemma 3.1, the distribution of $(\mathbf{B}, \mathbf{Br})$ is statistically indistinguishable from $(\mathbf{B}, \mathbf{t}^*)$ where $\mathbf{t}^* \leftarrow \mathbb{Z}_q^n$. The claim now follows via a similar argument as used to analyze \mathbf{B}_0 .

Lemma 5.21. Suppose $n \ge \lambda$, $m \ge 2n \log q$, and $\sigma_{td} \cdot O(m^{25/2}N^3) < \sigma_{agg} < 2^{\lambda_{DGS}}$. Then, for all polynomials $\kappa = \kappa(\lambda)$ and all $b \in \{0, 1\}$, there exists a negligible function $negl(\cdot)$ such that the statistical distance between $Hyb_6^{(b)}(\mathcal{A})$ and $Hyb_{7,\kappa}^{(b)}(\mathcal{A})$ is at most $1/\kappa + negl(\lambda)$.

Proof. By construction, in $\text{Hyb}_{7,\kappa}^{(b)}(\mathcal{A})$ and $\text{Hyb}_{8,\kappa}^{(b)}(\mathcal{A})$, the challenger samples $\mathbf{B} \leftarrow \mathbb{Z}_q^{n \times m}$ and sets T such that $\|\mathbf{T}\| \leq \sqrt{m}\sigma_{\text{td}}$. The claim now follows by Theorem 5.9 (specifically, the explainability property).

Lemma 5.22. Suppose $n \ge \lambda$, $m \ge 2n \log q$, and $\sigma_{td} \cdot O(m^{25/2}N^3) < \sigma_{agg} < 2^{\lambda_{DGS}}$. Then, for all polynomials $\kappa = \kappa(\lambda)$ and all $b \in \{0, 1\}$, $Hyb_{7,\kappa}^{(b)}(\mathcal{A})$ and $Hyb_{8\kappa}^{(b)}(\mathcal{A})$ are statistically indistinguishable.

Proof. By construction, in $\text{Hyb}_{7,\kappa}^{(b)}(\mathcal{A})$ and $\text{Hyb}_{8,\kappa}^{(b)}(\mathcal{A})$, the challenger samples $\mathbf{B} \leftarrow \mathbb{Z}_q^{n \times m}$ and sets T such that $\|\mathbf{T}\| \leq \sqrt{m}\sigma_{\text{td}}$. The claim now follows by Theorem 5.9 (specifically, the sampling distribution property).

Lemma 5.23. Suppose $n \ge \lambda$, $m \ge 2n \log q$, and q is prime. Then, for all polynomials $\kappa = \kappa(\lambda)$ and all $b \in \{0, 1\}$, $Hyb_{8\kappa}^{(b)}(\mathcal{A})$ and $Hyb_{9\kappa}^{(b)}(\mathcal{A})$ are statistically indistinguishable.

Proof. By construction, the only difference between these two distributions is the distribution of C_0 and $z_{0,i}$ for $i \in [N]$. We show that these components are statistically indistinguishable:

- Since $\mathbf{B} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^{n \times m}$ and $\mathbf{R}_{\widehat{\mathbf{C}}} \stackrel{\mathbb{R}}{\leftarrow} \{0, 1\}^{m \times m}$, the distribution of $\mathbf{BR}_{\widehat{\mathbf{C}}}$ is statistically close to uniform over $\mathbb{Z}_q^{n \times m}$ by the leftover hash lemma (Lemma 3.1). Thus, the distribution of $\mathbf{C}_0 = \mathbf{BR}_{\widehat{\mathbf{C}}} \sum_{j \in [N]} \mathbf{C}_j^*$ in $\mathsf{Hyb}_{9,\kappa}^{(b)}$ is statistically close to uniform over $\mathbb{Z}_q^{n \times m}$, which coincides with the distribution of \mathbf{C}_0 in $\mathsf{Hyb}_{8,\kappa}^{(b)}$
- Consider $\mathbf{z}_{0,i}$ for an index $i \in [N]$ where $(i, \mathbf{A}_{f_i^*}, \mathbf{t}_i^*) \mapsto (0, \mathbf{r}_i^*)$ in D_{sk} . By construction, this means $\mathbf{t}_i^* = \mathbf{d}_i + \mathbf{Br}_i^*$ where $\mathbf{r}_i^* \in \{0, 1\}^m$. Now,

$$\mathbf{C}_0 \mathbf{v}_i = \mathbf{B} \mathbf{R}_{\widehat{\mathbf{C}}} \mathbf{v}_i - \sum_{j \in [N]} \mathbf{C}_j^* \mathbf{v}_i.$$

Since C_j^* is a matrix commitment to $\mathbf{u}_j^{\mathsf{T}} \otimes \mathbf{t}_j^*$ and \mathbf{Z}_j^* is the associated opening, we appeal to Lemma 3.8 to conclude that

$$\begin{aligned} \forall j \neq i : \mathbf{C}_{j}^{*} \mathbf{v}_{i} &= -\mathbf{B} \mathbf{z}_{j,i}^{*} \\ \mathbf{C}_{i}^{*} \mathbf{v}_{i} &= \mathbf{t}_{i}^{*} - \mathbf{B} \mathbf{z}_{i,i}^{*} \\ &= \mathbf{d}_{i} + \mathbf{B} \mathbf{r}_{i}^{*} - \mathbf{B} \mathbf{z}_{i,i}^{*} \end{aligned}$$

Thus, in $Hyb_{9,\kappa}^{(b)}$, we have

$$C_0 \mathbf{v}_i = \mathbf{B} \mathbf{R}_{\widehat{\mathbf{C}}} \mathbf{v}_i - \mathbf{d}_i - \mathbf{B} \mathbf{r}_i^* + \sum_{j \in [N]} \mathbf{B} \mathbf{z}_{j,i}^*$$
$$= -\mathbf{d}_i + \mathbf{B} \left(\mathbf{R}_{\widehat{\mathbf{C}}} \mathbf{v}_i - \mathbf{r}_i^* + \sum_{j \in [N]} \mathbf{z}_{j,i}^* \right)$$

In this case, the distribution of $\mathbf{z}_{0,i}$ in $\mathsf{Hyb}_{9,\kappa}^{(b)}$ is precisely $\mathbf{B}_{\sigma_{\mathrm{agg}}}^{-1}(-\mathbf{C}_0\mathbf{v}_i)$, which matches the distribution in $\mathsf{Hyb}_{8,\kappa}^{(b)}$

• Consider $\mathbf{z}_{0,i}$ for an index $i \in [N]$ where $(i, \mathbf{A}_{f_i^*}, \mathbf{t}_i^*) \mapsto (1, \mathbf{r}_i^*)$ in D_{sk} . By construction, this means $\mathbf{Br}_i^* = \mathbf{t}_i^* + \mathbf{A}_{f_i^*} \mathbf{G}^{-1}(\mathbf{d}_i) + \mathbf{p}$ and moreover $f_i^*(\mathbf{x}) = 1$. Let $\mathbf{Z}_{\mathbf{x}} = \mathsf{Open}^{\mathsf{mat}}(\mathsf{pp}_{\mathsf{com}}, \mathbf{x}^\mathsf{T} \otimes \mathbf{G})$. Using the fact that $\mathbf{A} = -\mathbf{B}_0 \mathbf{V}_{\ell m}$, we have

$$\begin{bmatrix} \mathbf{B} \mid \mathbf{B}_0 + \mathbf{C}_{\mathbf{x}} \end{bmatrix} \cdot \begin{bmatrix} -\mathbf{Z}_{\mathbf{x}} \\ -\mathbf{V}_{\ell m} \end{bmatrix} = -\mathbf{B}\mathbf{Z}_{\mathbf{x}} - \mathbf{B}_0\mathbf{V}_{\ell m} - \mathbf{C}_{\mathbf{x}}\mathbf{V}_{\ell m} = \mathbf{A} - \mathbf{x}^{\mathsf{T}} \otimes \mathbf{G}$$

Let $H_{A,f_i^*,x} = EvalFX(A, f_i^*, x)$. Using the fact that $B_0 = BR_{B_0} - C_x$, we now have

$$\mathbf{B}[\mathbf{I}_m \mid \mathbf{R}_{\mathbf{B}_0}] \cdot \begin{bmatrix} -\mathbf{Z}_{\mathbf{x}} \\ -\mathbf{V}_{\ell m} \end{bmatrix} \cdot \mathbf{H}_{\mathbf{A}, f_i^*, \mathbf{x}} = (\mathbf{A} - \mathbf{x}^{\mathsf{T}} \otimes \mathbf{G}) \cdot \mathbf{H}_{\mathbf{A}, f_i^*, \mathbf{x}} = \mathbf{A}_{f_i^*} - f_i^*(\mathbf{x}) \cdot \mathbf{G} = \mathbf{A}_{f_i^*} - \mathbf{G}_{\mathbf{A}_i^*} - \mathbf{G}_{\mathbf{A$$

This means we can write

$$\mathbf{A}_{f_i^*} \mathbf{G}^{-1}(\mathbf{d}_i) = \mathbf{d}_i + \mathbf{B}(-\mathbf{Z}_{\mathbf{x}} - \mathbf{R}_{\mathbf{B}_0} \mathbf{V}_{\ell m}) \mathbf{H}_{\mathbf{A}, f_i^*, \mathbf{x}} \mathbf{G}^{-1}(\mathbf{d}_i).$$
(5.9)

Using the fact that $\mathbf{p} = \mathbf{Br}_{\mathbf{p}}$, we now have

$$\begin{aligned} \forall j \neq i : \mathbf{C}_{j}^{*} \mathbf{v}_{i} &= -\mathbf{B}\mathbf{z}_{j,i}^{*} \\ \mathbf{C}_{i}^{*} \mathbf{v}_{i} &= \mathbf{t}_{i}^{*} - \mathbf{B}\mathbf{z}_{i,i}^{*} \\ &= \mathbf{B}\mathbf{r}_{i}^{*} + \mathbf{A}_{f_{i}^{*}}\mathbf{G}^{-1}(\mathbf{d}_{i}) + \mathbf{p} - \mathbf{B}\mathbf{z}_{i,i}^{*} \\ &= \mathbf{B}\mathbf{r}_{i}^{*} + \mathbf{d}_{i} + \mathbf{B}(-\mathbf{Z}_{x} - \mathbf{R}_{B_{0}}\mathbf{V}_{\ell m})\mathbf{H}_{\mathbf{A},f_{i}^{*},x}\mathbf{G}^{-1}(\mathbf{d}_{i}) + \mathbf{B}\mathbf{r}_{p} - \mathbf{B}\mathbf{z}_{i,i}^{*} \\ &= \mathbf{d}_{i} + \mathbf{B}(\mathbf{r}_{i}^{*} - (\mathbf{Z}_{x} + \mathbf{R}_{B_{0}}\mathbf{V}_{\ell m})\mathbf{H}_{\mathbf{A},f_{i}^{*},x}\mathbf{G}^{-1}(\mathbf{d}_{i}) + \mathbf{r}_{p} - \mathbf{z}_{i,i}^{*}). \end{aligned}$$

In this case,

$$\mathbf{C}_{0}\mathbf{v}_{i} = \mathbf{B}\mathbf{R}_{\widehat{\mathbf{C}}}\mathbf{v}_{i} - \sum_{j \in [N]} \mathbf{C}_{j}^{*}\mathbf{v}_{i} = -\mathbf{d}_{i} + \mathbf{B}\left(\mathbf{R}_{\widehat{\mathbf{C}}}\mathbf{v}_{i} - \mathbf{r}_{i}^{*} + (\mathbf{Z}_{\mathbf{x}} + \mathbf{R}_{\mathbf{B}_{0}}\mathbf{V}_{\ell m})\mathbf{H}_{\mathbf{A},f_{i}^{*},\mathbf{x}}\mathbf{G}^{-1}(\mathbf{d}_{i}) - \mathbf{r}_{\mathbf{p}} + \sum_{j \in [N]} \mathbf{z}_{j,i}^{*}\right).$$

Thus the distribution of $\mathbf{z}_{0,i}$ in this case in $\mathsf{Hyb}_{9,\kappa}^{(b)}$ is precisely $\mathbf{B}_{\sigma_{\text{agg}}}^{-1}(-\mathbf{C}_{0}\mathbf{v}_{i})$, which is the distribution in $\mathsf{Hyb}_{8,\kappa}^{(b)}$.

Since the distribution of C_0 is statistically indistinguishable between $Hyb_{8,\kappa}^{(b)}$ and $Hyb_{9,\kappa}^{(b)}$ and $z_{0,i}$ is constructed using identical procedures, the two distributions are statistically indistinguishable.

Lemma 5.24. Suppose $\sigma_{\text{agg}} > \lambda^{\omega(1)} \cdot \sigma_{\text{td}} \cdot m^{O(d)} \log q \cdot (N \log N + \log(\ell m))$. Then, for all polynomials $\kappa = \kappa(\lambda)$ and all $b \in \{0, 1\}$, $\text{Hyb}_{9\kappa}^{(b)}(\mathcal{A})$ and $\text{Hyb}_{10\kappa}^{(b)}(\mathcal{A})$ are statistically indistinguishable.

Proof. This follows by the Gaussian preimage smudging lemma (Lemma 5.10), as long as σ_{agg} is sufficiently large. In the following analysis, it suffices to consider the setting where $||\mathbf{T}|| \leq \sqrt{m}\sigma_{td}$. Otherwise, the output in both experiments is 0. In $\text{Hyb}_{10,\kappa}^{(b)}$, each $\mathbf{z}_{0,i}$ can be written as $\mathbf{z}_{0,i} = \tilde{\mathbf{z}}_{0,i} + \mathbf{z}'_{0,i}$. By Lemma 5.10, if $\sigma_{agg} > \lambda^{\omega(1)} \cdot \sqrt{m} ||\mathbf{z}'_{0,i}||$, then the distributions of $\mathbf{z}_{0,i}$ in $\text{Hyb}_{9,\kappa}^{(b)}(\mathcal{A})$ and $\text{Hyb}_{10,\kappa}^{(b)}(\mathcal{A})$ are statistically indistinguishable. The claim then holds by a hybrid argument over all $N = \text{poly}(\lambda)$ indices *i*. It suffices to analyze $||\mathbf{z}'_{0,i}||$ for each $i \in [N]$:

• Suppose $(i, \mathbf{A}_{f_i^*}, \mathbf{t}_i^*) \mapsto (0, \mathbf{r}_i^*)$ in D_{sk} . In this case, $\mathbf{z}'_{0,i} = -\mathbf{R}_{\widehat{\mathsf{C}}}\mathbf{v}_i + \mathbf{r}_i^* - \sum_{j \in [N]} \mathbf{z}_{j,i}^*$. By Lemma 3.8,

$$\|\mathbf{v}_{i}\| \leq O(\|\mathbf{T}\| \cdot m^{4} \log q) = O(\sigma_{\rm td} \cdot m^{9/2} \log q) \|\mathbf{z}_{i,i}^{*}\| \leq O(\|\mathbf{T}\| \cdot m^{7} \log q \log N) \leq O(\sigma_{\rm td} \cdot m^{15/2} \log q \log N).$$
(5.10)

Since $\mathbf{R}_{\widehat{\mathbf{C}}} \in \{0,1\}^{m \times m}, \mathbf{r}_i^* \in \{0,1\}^m$, we conclude that

$$\|\mathbf{z}_{0,i}'\| = \left\| -\mathbf{R}_{\widehat{\mathbf{C}}}\mathbf{v}_i + \mathbf{r}_i^* - \sum_{j \in [N]} \mathbf{z}_{j,i}^* \right\| \le O(N\sigma_{\mathsf{td}} \cdot m^{15/2} \log q \log N)$$

• Suppose $(i, \mathbf{A}_{f_i^*}, \mathbf{t}_i^*) \mapsto (1, \mathbf{r}_i^*)$ in D_{sk} . In this case,

$$\mathbf{z}_{0,i}' = -\mathbf{R}_{\widehat{\mathbf{C}}} \mathbf{v}_{i} + \mathbf{r}_{i}^{*} - (\mathbf{Z}_{\mathbf{x}} + \mathbf{R}_{\mathbf{B}_{0}} \mathbf{V}_{\ell m}) \mathbf{H}_{\mathbf{A}, f_{i}^{*}, \mathbf{x}} \mathbf{G}^{-1}(\mathbf{d}_{i}) + \mathbf{r}_{\mathbf{p}} - \sum_{j \in [N]} \mathbf{z}_{j,i}^{*}$$

By Theorem 3.6, $\|\mathbf{H}_{\mathbf{A}, f_i^*, \mathbf{x}}\| \le m^{O(d)}$. By Lemma 3.8,

$$\begin{aligned} \|\mathbf{V}_{\ell m}\| &\le O(\|\mathbf{T}\| \cdot m^4 \log q) = O(\sigma_{\rm td} \cdot m^{9/2} \log q) \\ \|\mathbf{Z}_{\mathbf{x}}\| &\le O(\|\mathbf{T}\| \cdot m^7 \log q \log(\ell m)) = O(\sigma_{\rm td} \cdot m^{15/2} \log q \log(\ell m)). \end{aligned}$$

Next, $\mathbf{R}_{\mathbf{B}_0} \in \{0, 1\}^{m \times m}$. Combined with Eq. (5.10), we can bound

$$\left\|\mathbf{z}_{0,i}^{\prime}\right\| \leq \sigma_{\mathsf{td}} \cdot m^{O(d)} \log q \cdot (N \log N + \log(\ell m)).$$

If $\sigma_{\text{agg}} > \lambda^{\omega(1)} \cdot \sigma_{\text{td}} \cdot m^{O(d)} \log q \cdot (N \log N + \log(\ell m))$, then $\sigma_{\text{agg}} > \lambda^{\omega(1)} \cdot \sqrt{m} \|\mathbf{z}_{0,i'}\|$ for all $i \in [N]$ and the claim holds.

Lemma 5.25. Suppose $n \ge \lambda$, $m \ge 2n \log q$, q is prime, and $\sigma_{\text{agg}} > \log m$. Then, for all polynomials $\kappa = \kappa(\lambda)$ and all $b \in \{0, 1\}$, $\text{Hyb}_{10,\kappa}^{(b)}(\mathcal{A})$ and $\text{Hyb}_{11,\kappa}^{(b)}(\mathcal{A})$ are statistically indistinguishable.

Proof. The only difference between these two experiments is the distribution of $(\tilde{\mathbf{z}}_{0,i}, \mathbf{d}_i)$ for each $i \in [N]$. In Hyb^(b)_{10, κ}, the challenger samples $\mathbf{d}_i \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}^n_q$ and $\tilde{\mathbf{z}}_{0,i} \leftarrow \mathbf{B}_{\sigma_{\text{agg}}}^{-1}(\mathbf{d}_i)$ whereas in Hyb^(b)_{11, κ}, the challenger samples $\tilde{\mathbf{z}}_{0,i} \leftarrow D^m_{\mathbb{Z},\sigma_{\text{agg}}}$ and sets $\mathbf{d}_i = \mathbf{B}\tilde{\mathbf{z}}_{0,i}$. Under the given conditions, these two distributions are statistically indistinguishable by Lemma 3.3. \Box

Lemma 5.26. Suppose $n \ge \lambda$, $m \ge 2n \log q$, q is prime, $\chi \ge \log m$, and $\sigma_{td} \ge O(\log m)$. Then, for all polynomials $\kappa = \kappa(\lambda)$ and all $b \in \{0, 1\}$, $\mathsf{Hyb}_{11,\kappa}^{(b)}(\mathcal{A})$ and $\mathsf{Hyb}_{12,\kappa}^{(b)}(\mathcal{A})$ are statistically indistinguishable.

Proof. Follows immediately from Lemmas 3.2 and 3.4.

Lemma 5.27. Suppose the $(2m^2, \sigma_{td})$ -succinct LWE assumption with lattice parameters (n, m, q, χ) holds. Then, for all polynomials $\kappa = \kappa(\lambda)$ and all $b \in \{0, 1\}$, $Hyb_{12\kappa}^{(b)}(\mathcal{A})$ and $Hyb_{13\kappa}^{(b)}(\mathcal{A})$ is computationally indistinguishable.

Proof. Suppose $|\Pr[Hyb_{12,\kappa}^{(b)}(\mathcal{A}) = 1] - \Pr[Hyb_{13,\kappa}^{(b)}(\mathcal{A}) = 1]| = \varepsilon$ for some non-negligible ε . We use \mathcal{A} to construct an efficient adversary \mathcal{B} for the $(2m^2, \sigma_{td})$ -succinct LWE assumption:

• Setup phase: At the beginning of the game, algorithm \mathcal{B} receives the challenge (B, \tilde{c} , W, T). Algorithm \mathcal{B} starts running algorithm \mathcal{A} on input 1^{λ} , 1^{τ} , and *N*. Algorithm \mathcal{A} begins by outputting the attribute $\mathbf{x} \in \{0, 1\}^{\ell}$. Algorithm \mathcal{B} sets pp_{com} = (B, W, T) and then samples the following:

$$\begin{split} (\mathsf{crs}_{\mathsf{NIZK}},\mathsf{td}_{\mathsf{NIZK}}) & \leftarrow \mathsf{NIZK}.\mathsf{TrapSetup}(1^{\texttt{A}}) \\ \mathbf{C}_{\mathbf{x}} &= \mathsf{Com}^{\mathsf{mat}}(\mathsf{pp}_{\mathsf{com}},\mathbf{x}^{\mathsf{T}}\otimes\mathbf{G}) \\ \mathbf{R}_{\widehat{C}},\mathbf{R}_{B_{0}} \overset{\mathbb{R}}{\leftarrow} \{0,1\}^{m\times m}, \ \mathbf{r}_{p} \overset{\mathbb{R}}{\leftarrow} \{0,1\}^{m} \\ \mathbf{B}_{0} &= \mathbf{BR}_{B_{0}} - \mathbf{C}_{\mathbf{x}}, \ \mathbf{p} = \mathbf{Br}_{p}. \end{split}$$

Algorithm \mathcal{B} sets pp = (crs_{NIZK}, pp_{com}, **B**₀, **p**) and gives pp to \mathcal{A} . Algorithm \mathcal{B} also initializes a counter ctr = 0, dictionaries D, D_{sk}, and samples an index ind $\stackrel{R}{\leftarrow} [Q_{ro}]$.

- Queries to H_1 : Whenever \mathcal{A} queries H_1 on an index $i \in \mathbb{N}$, algorithm \mathcal{B} samples $\tilde{z}_{0,i} \leftarrow D^m_{\mathbb{Z},\sigma_{agg}}$ and responds with $H_1(i) = \mathbf{d}_i = \mathbf{B}\tilde{z}_{0,i}$.
- Queries to H_2 : Whenever \mathcal{A} queries H_2 on an input $\xi \in \{0, 1\}^*$, if this is not the indth query to H_2 , algorithm \mathcal{B} responds with $\gamma \xleftarrow{\mathbb{R}} \{0, 1\}^{\rho}$. If it is the indth query, then algorithm \mathcal{B} proceeds as follows:
 - Algorithm \mathcal{B} parses $\xi_{ind} = (pp^*, (pk_1^*, f_1^*), ..., (pk_N^*, f_N^*))$, where $pk_i^* = (t_i^*, \pi_i^*)$. If ξ_{ind} does not have this form or $pp^* \neq pp$, then \mathcal{B} outputs 0. If IsValid(pp, *i*, $f_i^*, pk_i^*) = 0$ for any *i* ∈ [*N*], algorithm \mathcal{B} outputs 0.
 - Algorithm \mathcal{B} computes $V_{\ell m} = \text{Ver}^{\text{mat}}(\text{pp}_{\text{com}}, 1^{\ell m})$ and $A = -B_0 V_{\ell m}$.
 - For each $i \in [N]$, algorithm \mathcal{B} computes $\mathbf{A}_{f_i^*} = \text{EvalF}(\mathbf{A}, f_i^*)$. If $(i, \mathbf{A}_{f_i^*}, \mathbf{t}_i^*)$ is not contained in D_{sk} , algorithm \mathcal{B} checks that $f_i^*(\mathbf{x}) = 1$. If not, it outputs 0. Otherwise, algorithm \mathcal{B} computes

 $\mathbf{r}_{i}^{*} = \mathsf{NIZK}.\mathsf{Extract}(\mathsf{td}_{\mathsf{NIZK}}, C_{\mathsf{valid}}, (i, \mathbf{A}_{f_{i}^{*}}, \mathbf{d}_{i}, \mathbf{p}, \mathbf{B}, \mathbf{t}_{i}^{*}), \pi_{i}^{*}).$

If $\mathbf{r}_i^* \notin \{0,1\}^m$ or $\mathbf{Br}_i^* \neq \mathbf{t}_i^* + \mathbf{A}_{f_i^*}\mathbf{G}^{-1}(\mathbf{d}_i) + \mathbf{p}$, then algorithm \mathcal{B} outputs 0. Otherwise, the challenger adds the mapping $(i, \mathbf{A}_{f_i^*}, \mathbf{t}_i^*) \mapsto (1, \mathbf{r}_i^*)$ to D_{sk} .

- Algorithm \mathcal{B} computes $\mathbf{V}_N = \text{Ver}^{\text{mat}}(\text{pp}_{\text{com}}, 1^N)$ and parses $\mathbf{V}_N = [\mathbf{v}_1 | \cdots | \mathbf{v}_N]$. Now, for each $i \in [N]$, algorithm \mathcal{B} computes

$$C_i^* = \text{Com}^{\text{mat}}(\text{pp}_{\text{com}}, \mathbf{u}_i^{\mathsf{T}} \otimes \mathbf{t}_i^*)$$
$$Z_i^* = \text{Open}^{\text{mat}}(\text{pp}_{\text{com}}, \mathbf{u}_i^{\mathsf{T}} \otimes \mathbf{t}_i^*)$$

It parses $\mathbf{Z}_i^* = [\mathbf{z}_{i,1}^* | \cdots | \mathbf{z}_{i,N}^*]$ and sets $\mathbf{C}_0 = \mathbf{BR}_{\widehat{\mathbf{C}}} - \sum_{j \in [N]} \mathbf{C}_j^*$. For each $i \in [N]$, algorithm \mathcal{B} defines $\mathbf{z}_{0,i}$ as follows:

* Suppose $(i, \mathbf{A}_{f_i^*}, \mathbf{t}_i^*) \mapsto (0, \mathbf{r}_i^*)$ in D_{sk} . Then, it sets

$$\mathbf{z}_{0,i} = \tilde{\mathbf{z}}_{0,i} - \mathbf{R}_{\widehat{\mathbf{C}}}\mathbf{v}_i + \mathbf{r}_i^* - \sum_{j \in [N]} \mathbf{z}_{j,i}^*.$$

* Suppose $(i, \mathbf{A}_{f_i^*}, \mathbf{t}_i^*) \mapsto (1, \mathbf{r}_i^*)$ in D_{sk} . Then, it computes $\mathbf{Z}_{\mathbf{x}} = \mathsf{Open}^{\mathsf{mat}}(\mathsf{pp}_{\mathsf{com}}, \mathbf{x}^{\mathsf{T}} \otimes \mathbf{G})$ and $\mathsf{H}_{\mathsf{A}, f_i^*, \mathbf{x}} = \mathsf{EvalFX}(\mathsf{A}, f_i^*, \mathbf{x})$ and sets

$$z_{0,i} = \tilde{z}_{0,i} - R_{\widehat{C}} v_i + r_i^* - (Z_x + R_{B_0} V_{\ell m}) H_{A, f_i^*, x} G^{-1}(d_i) + r_p - \sum_{j \in [N]} z_{j,i}^*$$

Here, $\tilde{\mathbf{z}}_{0,i} \in \mathbb{Z}_q^m$ is the value algorithm \mathcal{B} sampled when responding to a query on $H_1(i)$. If algorithm \mathcal{A} has not yet made a query to H_1 on some index $i \in [N]$, algorithm \mathcal{B} samples $\tilde{\mathbf{z}}_{0,i} \leftarrow D_{\mathbb{Z},\sigma_{\text{agg}}}^m$ and programs $H_1(i) = \mathbf{d}_i = \mathbf{B}\tilde{\mathbf{z}}_{0,i}$.

– Finally, algorithm \mathcal{B} computes

$$\gamma^* \leftarrow \text{Explain}(\text{pp}_{\text{com}}, 1^{\lambda_{\text{DGS}}}, 1^{\kappa}, (C_0, \mathbf{z}_{0,1}, \dots, \mathbf{z}_{0,N}), \sigma_{\text{agg}}).$$

It responds to \mathcal{A} with $H_2(\xi_{ind}) = \gamma^*$.

- **Key-generation queries:** Whenever \mathcal{A} makes a key-generation query on an index $i \in [N]$ and a function f, algorithm \mathcal{B} samples $\mathbf{r} \leftarrow \{0, 1\}^m$ and sets $\mathbf{t} = \mathbf{Br} + \mathbf{d}_i$. It also computes $\mathbf{A}_f = \text{EvalF}(\mathbf{A}, f)$ and $\mathbf{d}_i = H_1(i)$. Then, it generates a (simulated) proof $\pi \leftarrow \text{NIZK.Sim}(\text{td}_{\text{NIZK}}, C_{\text{valid}}, (i, \mathbf{A}_f, \mathbf{d}_i, \mathbf{p}, \mathbf{B}, \mathbf{t}))$ and responds to \mathcal{A} with the public key $\mathbf{pk} = (\mathbf{t}, \pi)$. Algorithm \mathcal{B} adds the mapping ctr $\mapsto (i, f, \mathbf{t})$ to D and $(i, \mathbf{A}_f, \mathbf{t}) \mapsto (0, \mathbf{r})$ to D_{sk} if such a mapping does not already exist. In addition, algorithm \mathcal{B} checks that IsValid(pp, $i, f, \mathbf{pk}) = 1$ and outputs 0 if not.
- Challenge query: Let $((c_1, f_1, pk_i), ..., (c_N, f_N, pk_N))$ be algorithm \mathcal{A} 's challenge query. For each $i \in [N]$, algorithm \mathcal{B} proceeds as follows:
 - If $c_i \in \{1, ..., ctr\}$, then algorithm \mathcal{B} looks up $(i', f', pk') = D[c_i]$ and checks that i = i'. If not, algorithm \mathcal{B} outputs 0. Otherwise, it sets $pk_i^* = pk' = (t_i, \pi_i)$ and $f_i^* = f'$.
 - If $c_i = \bot$, then algorithm \mathcal{B} checks that $f_i(\mathbf{x}) = 1$. If so, it parses $pk_i = (t_i, \pi_i)$ and checks that $IsValid(pp, i, f_i, pk_i) = 1$. If so, algorithm \mathcal{B} sets $pk_i^* = pk_i$ and $f_i^* = f_i$. Otherwise, it outputs 0.

Let $\xi^* = (pp, (pk_1^*, f_1^*), \dots, (pk_N^*, f_N^*))$. If algorithm \mathcal{A} has not made at least ind queries to H_2 , then algorithm \mathcal{B} outputs 0. Otherwise, let ξ_{ind} be the indth query algorithm \mathcal{A} made to H_2 . If $\xi^* \neq \xi_{ind}$, then algorithm \mathcal{B} outputs 0. Otherwise, algorithm \mathcal{B} constructs the challenge ciphertext as follows:

$$ct = (\tilde{\mathbf{c}}^{\mathsf{T}}, \ \tilde{\mathbf{c}}^{\mathsf{T}}\mathbf{R}_{\widehat{\mathbf{C}}}, \ \tilde{\mathbf{c}}^{\mathsf{T}}\mathbf{R}_{\mathbf{B}_{0}}, \ \tilde{\mathbf{c}}^{\mathsf{T}}\mathbf{r}_{\mathbf{p}} + \lfloor q/2 \rfloor \cdot b).$$

The challenger gives ct to \mathcal{A} .

• **Output:** At the end of the game, algorithm \mathcal{A} outputs a bit $b' \in \{0, 1\}$, which algorithm \mathcal{B} also outputs.

If \mathcal{A} is efficient, then algorithm \mathcal{B} is also efficient by construction. By definition, the succinct LWE challenger samples $\mathbf{B} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^{n \times m}$, $\mathbf{W} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^{2m^2n \times m}$, and $\mathbf{T} \leftarrow [\mathbf{I}_{2m^2} \otimes \mathbf{B} \mid \mathbf{W}]_{\sigma}^{-1}(\mathbf{I}_{2m^2} \otimes \mathbf{G})$. Thus, algorithm \mathcal{B} perfectly simulates the public parameters pp according to the specification of $\mathsf{Hyb}_{12,\kappa}^{(b)}$ and $\mathsf{Hyb}_{13,\kappa}^{(b)}$. The random oracle queries and key-generation queries are also simulated exactly according to the specification of $\mathsf{Hyb}_{12,\kappa}^{(b)}$ and $\mathsf{Hyb}_{13,\kappa}^{(b)}$. Consider now the distribution of the challenge ciphertext:

• Suppose $\tilde{\mathbf{c}}^{\mathsf{T}} = \mathbf{s}^{\mathsf{T}} \mathbf{B} + \mathbf{e}^{\mathsf{T}}$ where $\mathbf{s} \leftarrow \mathbb{Z}_q^n$ and $\mathbf{e} \leftarrow D_{\mathbb{Z},Y}^m$. Observe that when $\xi_{\mathsf{ind}} = \xi^*$, we have that

$$\mathbf{BR}_{\widehat{\mathbf{C}}} = \mathbf{C}_0 + \sum_{i \in [N]} \mathbf{C}_i^* = \widehat{\mathbf{C}}$$

Similarly, by definition of \mathbf{B}_0 and \mathbf{p} , we further have

$$BR_{B_0} = B_0 + C_x$$
$$Br_p = p.$$

Then, we can write

$$\begin{aligned} \mathbf{c} &= \left(\mathbf{s}^{\mathsf{T}}\mathbf{B} + \mathbf{e}^{\mathsf{T}} , \ \mathbf{s}^{\mathsf{T}}\mathbf{B}\mathbf{R}_{\widehat{\mathbf{C}}} + \mathbf{e}^{\mathsf{T}}\mathbf{R}_{\widehat{\mathbf{C}}} , \ \mathbf{s}^{\mathsf{T}}\mathbf{B}\mathbf{R}_{B_{0}} + \mathbf{e}^{\mathsf{T}}\mathbf{R}_{B_{0}} , \ \mathbf{s}^{\mathsf{T}}\mathbf{B}\mathbf{r}_{\mathbf{p}} + \mathbf{e}^{\mathsf{T}}\mathbf{r}_{\mathbf{p}} + \lfloor q/2 \rfloor \cdot b \right) \\ &= \left(\mathbf{s}^{\mathsf{T}}\mathbf{B} + \mathbf{e}^{\mathsf{T}} , \ \mathbf{s}^{\mathsf{T}}\widehat{\mathbf{C}} + \mathbf{e}^{\mathsf{T}}\mathbf{R}_{\widehat{\mathbf{C}}} , \ \mathbf{s}^{\mathsf{T}}(\mathbf{B}_{0} + \mathbf{C}_{\mathbf{x}}) + \mathbf{e}^{\mathsf{T}}\mathbf{R}_{B_{0}} , \ \mathbf{s}^{\mathsf{T}}\mathbf{p} + \mathbf{e}^{\mathsf{T}}\mathbf{r}_{\mathbf{p}} + \lfloor q/2 \rfloor \cdot b \right). \end{aligned}$$

This is the ciphertext distribution in $Hyb_{12\kappa}^{(b)}$

• Suppose $\tilde{\mathbf{c}} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_p^n$. Then, the challenge ciphertext is distributed exactly as in Hyb^(b)_{13 k}.

We conclude that algorithm $\mathcal B$ breaks the succinct LWE assumption with the same advantage ϵ .

Lemma 5.28. Suppose $\sigma_{\text{agg}} > \lambda^{\omega(1)} \cdot \sigma_{\text{td}} \cdot m^{O(d)} \log q \cdot (N \log N + \log(\ell m))$. Then, for all polynomials $\kappa = \kappa(\lambda)$ and all $b \in \{0, 1\}$, $\text{Hyb}_{13\kappa}^{(b)}(\mathcal{A})$ and $\text{Hyb}_{14\kappa}^{(b)}(\mathcal{A})$ are statistically indistinguishable.

Proof. Follows by the same argument as the proof of Lemma 5.25 (via the Gaussian preimage smudging lemma). \Box

Lemma 5.29. Suppose $n \ge \lambda$, $m \ge 2n \log q$, and q > 2 is prime. Then, for all polynomials $\kappa = \kappa(\lambda)$ and all $b \in \{0, 1\}$, $Hyb_{14\kappa}^{(b)}(\mathcal{A})$ and $Hyb_{15\kappa}^{(b)}(\mathcal{A})$ are statistically indistinguishable.

Proof. By the leftover hash lemma (Lemma 3.1), the distributions ($\mathbf{B}, \tilde{\mathbf{c}}, \mathbf{Ar_p}, \tilde{\mathbf{c}}^{\mathsf{T}}\mathbf{r}_p + b \cdot \lfloor q/2 \rfloor$) and ($\mathbf{B}, \tilde{\mathbf{c}}, \mathbf{p}, c_4 + b \cdot \lfloor q/2 \rfloor$) where $\mathbf{B} \leftarrow \mathbb{Z}_q^{n \times m}, \tilde{\mathbf{c}} \leftarrow \mathbb{Z}_q^n, \mathbf{r}_p \leftarrow \{0, 1\}^m, \mathbf{p} \leftarrow \mathbb{Z}_q^n$, and $c_4 \leftarrow \mathbb{Z}_q$ are statistically indistinguishable. The first distribution corresponds to $\mathsf{Hyb}_{14,\kappa}^{(b)}$ while the second corresponds to $\mathsf{Hyb}_{15,\kappa}^{(b)}$.

Completing the proof. To complete the proof of Theorem 5.14, suppose

$$|\Pr[\mathsf{Hyb}_0^{(0)}(\mathcal{A}) = 1] - \Pr[\mathsf{Hyb}_0^{(1)}(\mathcal{A}) = 1]| = \varepsilon(\lambda)$$

for some non-negligible function. By Lemma 5.15, this means

$$|\Pr[\mathsf{Hyb}_{1}^{(1)}(\mathcal{A}) = 1] - \Pr[\mathsf{Hyb}_{1}^{(1)}(\mathcal{A}) = 1]| = \frac{\varepsilon(\lambda)}{Q_{\mathsf{ro}}}.$$
(5.11)

Since $Q_{ro} = \text{poly}(\lambda)$, the quantity $\varepsilon(\lambda)/Q_{ro}$ is also non-negligible. Thus, there exists a polynomial $\kappa'(\lambda)$ such that for infinitely-many $\lambda \in \mathbb{N}$, $\varepsilon(\lambda)/Q_{ro} > 1/\kappa'(\lambda)$. Let $\kappa(\lambda) = 3\kappa'(\lambda)$. For $i \le 6$, we write $\text{Hyb}_{i,\kappa}^{(b)}(\mathcal{A})$ to denote $\text{Hyb}_{i}^{(b)}(\mathcal{A})$. Moreover, since the challenger's behavior in $\text{Hyb}_{15}^{(b)}$ is independent of the bit $b \in \{0, 1\}$,

$$\Pr[\mathsf{Hyb}_{15}^{(0)}(\mathcal{A}) = 1] = \Pr[\mathsf{Hyb}_{15}^{(1)}(\mathcal{A}) = 1]$$

By Lemmas 5.16 to 5.29, we now have for all $\lambda \in \mathbb{N}$,

$$\begin{aligned} |\Pr[\mathsf{Hyb}_{0}^{(1)}(\mathcal{A}) = 1] - \Pr[\mathsf{Hyb}_{1}^{(1)}(\mathcal{A}) = 1]| &\leq \sum_{i=1}^{15} |\Pr[\mathsf{Hyb}_{i,\kappa}^{(0)}(\mathcal{A}) = 1] - \Pr[\mathsf{Hyb}_{i+1,\kappa}^{(0)}(\mathcal{A}) = 1]| \\ &+ |\Pr[\mathsf{Hyb}_{15}^{(0)}(\mathcal{A}) = 1] - \Pr[\mathsf{Hyb}_{15}^{(1)}(\mathcal{A}) = 1]| \\ &+ \sum_{i=1}^{15} |\Pr[\mathsf{Hyb}_{i+1,\kappa}^{(1)}(\mathcal{A}) = 1] - \Pr[\mathsf{Hyb}_{i,\kappa}^{(1)}(\mathcal{A}) = 1]| \\ &\leq 2/\kappa(\lambda) + \operatorname{negl}(\lambda). \end{aligned}$$

However, this contradicts Eq. (5.11) which asserts that there are infinitely-many $\lambda \in \mathbb{N}$ where

$$|\Pr[\mathsf{Hyb}_1^{(1)}(\mathcal{A}) = 1] - \Pr[\mathsf{Hyb}_1^{(1)}(\mathcal{A}) = 1]| = \frac{\varepsilon(\lambda)}{Q_{\mathsf{ro}}} > \frac{1}{\kappa'(\lambda)} = \frac{3}{\kappa(\lambda)}$$

Hence, we conclude that it must be the case that $\varepsilon(\lambda)$ is negligible, and attribute-selective security holds.

Parameter instantiation. We now provide one example instantiation of the lattice parameters for Construction 5.11 so as to satisfy Theorems 5.13 and 5.14. Let λ be a security parameter and τ be the policy parameter. Let $d = d(\tau)$ and $\ell = \ell(\tau)$ be the bound on the policy depth and the attribute length, respectively. Let $N \leq 2^{\lambda}$ be the number of slots. Let $\varepsilon \in (0, 1)$ be a constant. Then, we set

$$\begin{split} n &= (\lambda d)^{1/\varepsilon} \cdot \operatorname{poly}(\log \lambda, \log d, \log \ell, \log N) \\ m &= nd \cdot \operatorname{poly}(\log \lambda, \log d, \log \ell, \log N) \\ \chi &= \operatorname{poly}(n, \log d, \log \ell) \\ \sigma_{td} &= O(m^3 \log m) \\ \sigma_{agg} &= \lambda^{\omega(1)} \cdot m^{O(d)} \cdot N^3 \log \ell \\ q &= \lambda^{\omega(1)} \cdot m^{O(d)} \cdot \operatorname{poly}(n, N, \ell, \log d) \\ \lambda_{DGS} &= (\lambda + d) \cdot \operatorname{poly}(\log \lambda, \log d, \log \ell) \end{split}$$

For this choice of parameters, $q < 2^{n^{e}}$. In this case, security relies on the $(2m^{2}, \sigma_{td})$ -succinct LWE assumption with LWE parameters (n, m, q, χ) ; this corresponds to a sub-exponential modulus-to-noise ratio. We summarize our instantiation in the following corollary:

Corollary 5.30 (Slotted Key-Policy Registered ABE). Let λ be a security parameter and $N \leq 2^{\lambda}$ be the number of slots. Let $d \leq 2^{\lambda}$ be a depth bound and $\ell \leq 2^{\lambda}$ be an attribute length. Then, under the poly (λ, d) -succinct LWE assumption with a sub-exponential modulus-to-noise ratio, there exists a slotted key-policy registered ABE scheme with N slots where the size of the public parameters, the size of the individual public/secret keys, and the size of the ciphertext is $poly(\lambda, d, \log N, \log \ell)$.

Remark 5.31 (Transparent Setup via Decomposed LWE). Similar to Remark 4.13, we can replace the matrix commitment scheme in Construction 5.11 with the alternative instantiation from [Wee25, Appendix C] based on decomposed LWE. This yields a key-policy registered ABE scheme that supports an unbounded number of users with a transparent setup.

Acknowledgments

This work was initiated at Shonan Meeting No.195 "Workshop on Encrypted Computation / Enhancing Functionality in Cryptography." David J. Wu is supported by NSF CNS-2140975, CNS-2318701, a Microsoft Research Faculty Fellowship, and a Google Research Scholar award.

References

- [ABB10] Shweta Agrawal, Dan Boneh, and Xavier Boyen. Efficient lattice (H)IBE in the standard model. In *EUROCRYPT*, 2010.
- [Ajt96] Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In STOC, 1996.
- [AMR25] Damiano Abram, Giulio Malavolta, and Lawrence Roy. Key-homomorphic computations for RAM: fully succinct randomised encodings and more. In *CRYPTO*, 2025.
- [AT24] Nuttapong Attrapadung and Junichi Tomida. A modular approach to registered abe for unbounded predicates. In *CRYPTO*, 2024.
- [BB04] Dan Boneh and Xavier Boyen. Efficient selective-id secure identity-based encryption without random oracles. In *EUROCRYPT*, 2004.
- [BCD⁺25] Pedro Branco, Arka Rai Choudhuri, Nico Döttling, Abhishek Jain, Giulio Malavolta, and Akshayaram Srinivasan. Black-box non-interactive zero knowledge from vector trapdoor hash. In *EUROCRYPT*, 2025.
- [BDN18] Dan Boneh, Manu Drijvers, and Gregory Neven. Compact multi-signatures for smaller blockchains. In *ASIACRYPT*, 2018.
- [BFM88] Manuel Blum, Paul Feldman, and Silvio Micali. Non-interactive zero-knowledge and its applications (extended abstract). In *STOC*, 1988.
- [BGG⁺14] Dan Boneh, Craig Gentry, Sergey Gorbunov, Shai Halevi, Valeria Nikolaenko, Gil Segev, Vinod Vaikuntanathan, and Dhinakaran Vinayagamurthy. Fully key-homomorphic encryption, arithmetic circuit ABE and compact garbled circuits. In *EUROCRYPT*, 2014.
- [BGI⁺01] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. In *CRYPTO*, 2001.
- [BGW05] Dan Boneh, Craig Gentry, and Brent Waters. Collusion resistant broadcast encryption with short ciphertexts and private keys. In *CRYPTO*, 2005.
- [BZ14] Dan Boneh and Mark Zhandry. Multiparty key exchange, efficient traitor tracing, and more from indistinguishability obfuscation. In *CRYPTO*, 2014.
- [CHW25] Jeffrey Champion, Yao-Ching Hsieh, and David J. Wu. Registered ABE and adaptively-secure broadcast encryption from succinct LWE. In *CRYPTO*, 2025.
- [CW24] Jeffrey Champion and David J. Wu. Distributed broadcast encryption from lattices. In TCC, 2024.
- [DDO⁺01] Alfredo De Santis, Giovanni Di Crescenzo, Rafail Ostrovsky, Giuseppe Persiano, and Amit Sahai. Robust non-interactive zero knowledge. In *CRYPTO*, 2001.
- [DORS08] Yevgeniy Dodis, Rafail Ostrovsky, Leonid Reyzin, and Adam D. Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. *SIAM J. Comput.*, 38(1), 2008.
- [FLS90] Uriel Feige, Dror Lapidot, and Adi Shamir. Multiple non-interactive zero knowledge proofs based on a single random string (extended abstract). In *FOCS*, 1990.
- [FN93] Amos Fiat and Moni Naor. Broadcast encryption. In *CRYPTO*, 1993.
- [FWW23] Cody Freitag, Brent Waters, and David J. Wu. How to use (plain) witness encryption: Registered ABE, flexible broadcast, and more. In *CRYPTO*, 2023.
- [GHMR18] Sanjam Garg, Mohammad Hajiabadi, Mohammad Mahmoody, and Ahmadreza Rahimi. Registrationbased encryption: Removing private-key generator from IBE. In *TCC*, 2018.

- [GKPW24] Sanjam Garg, Dimitris Kolonelos, Guru-Vamsi Policharla, and Mingyuan Wang. Threshold encryption with silent setup. In *CRYPTO*, 2024.
- [GLWW24] Rachit Garg, George Lu, Brent Waters, and David J. Wu. Reducing the CRS size in registered ABE systems. In *CRYPTO*, 2024.
- [GPSW06] Vipul Goyal, Omkant Pandey, Amit Sahai, and Brent Waters. Attribute-based encryption for fine-grained access control of encrypted data. In *ACM CCS*, 2006.
- [GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, 2008.
- [GSW13] Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In *CRYPTO*, 2013.
- [GW09] Craig Gentry and Brent Waters. Adaptive security in broadcast encryption systems (with short ciphertexts). In *EUROCRYPT*, 2009.
- [HLWW23] Susan Hohenberger, George Lu, Brent Waters, and David J. Wu. Registered attribute-based encryption. In *EUROCRYPT*, 2023.
- [HW15] Pavel Hubácek and Daniel Wichs. On the communication complexity of secure function evaluation with long output. In *ITCS*, 2015.
- [HWW25] Yao-Ching Hsieh, Brent Waters, and David J. Wu. A generic approach to adaptively-secure broadcast encryption in the plain model. In *EUROCRYPT*, 2025.
- [KMW23] Dimitris Kolonelos, Giulio Malavolta, and Hoeteck Wee. Distributed broadcast encryption from bilinear groups. In *ASIACRYPT*, 2023.
- [LW22] George Lu and Brent Waters. How to sample a discrete gaussian (and more) from a random oracle. In *TCC*, 2022.
- [LWW25] George Lu, Brent Waters, and David J. Wu. Multi-authority registered attribute-based encryption. In *EUROCRYPT*, 2025.
- [MP12] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *EUROCRYPT*, 2012.
- [PS19] Chris Peikert and Sina Shiehian. Noninteractive zero knowledge for NP from (plain) learning with errors. In *CRYPTO*, 2019.
- [Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In STOC, 2005.
- [RY07] Thomas Ristenpart and Scott Yilek. The power of proofs-of-possession: Securing multiparty signatures against rogue-key attacks. In *EUROCRYPT*, 2007.
- [Sah99] Amit Sahai. Non-malleable non-interactive zero knowledge and adaptive chosen-ciphertext security. In *FOCS*, 1999.
- [SW05] Amit Sahai and Brent Waters. Fuzzy identity-based encryption. In *EUROCRYPT*, 2005.
- [Wat24] Brent Waters. A new approach for non-interactive zero-knowledge from learning with errors. In *STOC*, 2024.
- [Wee24] Hoeteck Wee. Circuit ABE with poly(depth, λ)-sized ciphertexts and keys from lattices. In *CRYPTO*, 2024.
- [Wee25] Hoeteck Wee. Almost optimal KP and CP-ABE for circuits from succinct LWE. In *EUROCRYPT*, 2025.

- [WQZD10] Qianhong Wu, Bo Qin, Lei Zhang, and Josep Domingo-Ferrer. Ad hoc broadcast encryption. In ACM CCS, 2010.
- [WW23] Hoeteck Wee and David J. Wu. Succinct vector, polynomial, and functional commitments from lattices. In *EUROCRYPT*, 2023.
- [WWW25] Brent Waters, Hoeteck Wee, and David J. Wu. New techniques for preimage sampling: Improved NIZKs and more from LWE. In *EUROCRYPT*, 2025.
- [ZZC⁺25] Ziqi Zhu, Kai Zhang, Zhili Chen, Junqing Gong, and Haifeng Qian. Black-box registered ABE from lattices. IACR Cryptol. ePrint Arch., 2025. Available at https://eprint.iacr.org/2025/051.pdf.
- [ZZGQ23] Ziqi Zhu, Kai Zhang, Junqing Gong, and Haifeng Qian. Registered ABE via predicate encodings. In ASIACRYPT, 2023.

A The [Wee25] Matrix Commitment Scheme

In this section, we recall the matrix commitment scheme from [Wee25, §3.2]. Let *n*, *m*, *q* be lattice parameters where $m \ge 2n \log q$. In the following, we write

$$pp = (B, W, T)$$
 where $[I_{2m^2} \otimes B | W] \cdot T = I_{2m^2} \otimes G$,

where $\mathbf{B} \in \mathbb{Z}_q^{n \times m}$, $\mathbf{W} \in \mathbb{Z}_q^{2m^2n \times m}$, and $\mathbf{T} \in \mathbb{Z}_q^{(2m^2+1)m \times 2m^3}$. Next, we define the Split function that takes as input (pp, *L*) where $L \leq 2m$ and outputs submatrices \mathbf{W}_L , \mathbf{T}_L of \mathbf{W} , \mathbf{T} , respectively, such that

$$[\mathbf{I}_{Lm} \otimes \mathbf{B} \mid \mathbf{W}_L] \cdot \mathbf{T}_L = \mathbf{I}_{Lm} \otimes \mathbf{G}.$$

Specifically, the Split function operates as follows:

• Split(pp, *L*): On input pp = (B, W, T) and $L \le 2m$, parse

$$\mathbf{W} = \begin{bmatrix} \mathbf{W}^{(1)} \\ \vdots \\ \mathbf{W}^{(2m^2)} \end{bmatrix} \quad \text{and} \quad \mathbf{T} = \begin{bmatrix} \mathbf{T}_{\mathrm{LT}}^{(1)} \\ \vdots \\ \mathbf{T}_{\mathrm{LT}}^{(2m^2)} \\ \underline{\mathbf{T}}_{\mathrm{LT}} \end{bmatrix},$$

where $\mathbf{W}^{(i)} \in \mathbb{Z}_q^{n \times m}$ and $\mathbf{T}_{\text{LT}}^{(i)}, \mathbf{\underline{T}}_{\text{LT}} \in \mathbb{Z}_q^{m \times Lm^2}$, and $\mathbf{T}_{\text{RT}} \in \mathbb{Z}_q^{(2m^2+1)m \times (2m^2-Lm)m}$. The split function outputs

$$\mathbf{W}_{L} \coloneqq \begin{bmatrix} \mathbf{W}^{(1)} \\ \vdots \\ \mathbf{W}^{(Lm)} \end{bmatrix} \in \mathbb{Z}_{q}^{Lmn \times m} \quad \text{and} \quad \mathbf{T}_{L} \coloneqq \begin{bmatrix} \mathbf{T}_{\mathrm{LT}}^{(1)} \\ \vdots \\ \mathbf{T}_{\mathrm{LT}}^{(Lm)} \\ \mathbf{T}_{\mathrm{T}} \end{bmatrix} \in \mathbb{Z}_{q}^{(Lm+1)m \times Lm^{2}}$$

In the following, we parse $\mathbf{T}_L \in \mathbb{Z}_q^{(Lm+1)m \times Lm^2}$ as $\mathbf{T}_L \coloneqq \begin{bmatrix} \overline{\mathbf{T}}_L \\ \underline{\mathbf{T}}_L \end{bmatrix}$ where $\overline{\mathbf{T}}_L \in \mathbb{Z}_q^{Lm^2 \times Lm^2}$ and $\underline{\mathbf{T}}_L \in \mathbb{Z}_q^{m \times Lm^2}$. In addition, define

$$bits(\mathbf{M}) \coloneqq vec(\mathbf{G}^{-1}(\mathbf{M})) \in \mathbb{Z}_q^{Lm}.$$

Let $\mathbf{J}_L \in \{0, 1\}^{Lm^2 \times L\lceil \log q \rceil}$ be the fixed matrix from [Wee25, Lemma 4] where for all $\mathbf{M} \in \mathbb{Z}_q^{n \times L}$,

$$(bits(\mathbf{M}) \otimes \mathbf{G}) \cdot \mathbf{J}_L = \mathbf{M} \cdot \mathbf{G}_L$$

The matrix commitment from [Wee25] is recursive, where the base case corresponds to committing to a matrix $\mathbf{M} \in \mathbb{Z}_q^{n \times L}$ where $L \leq 2m$. Without loss of generality, when L > 2m, we always pad the width of \mathbf{M} (with zeroes) to a value of L where $L = 2^k \cdot \ell$, for some $k \in \mathbb{N}$ and $\ell \in [2m]$. The associated verification matrix and openings can then be derived by truncating the corresponding matrices for the padded matrix. Specifically, if $C \cdot [\mathbf{V}_{LT} \mid \mathbf{V}_{RT}] = [\mathbf{M}_{LT} \mid \mathbf{M}_{RT}] - \mathbf{B} \cdot [\mathbf{Z}_{LT} \mid \mathbf{Z}_{RT}]$, then $C \cdot \mathbf{V}_{LT} = \mathbf{M}_{LT} - \mathbf{B} \cdot \mathbf{Z}_{LT}$.

- Com^{mx}(pp, M): On input pp = (B, W, T) and $M \in \mathbb{Z}_q^{n \times L}$:
 - If $L \leq 2m$, output $\mathbf{C} = (\text{bits}(\mathbf{M})^{\mathsf{T}} \otimes \mathbf{I}_n) \mathbf{W}_L \in \mathbb{Z}_a^{n \times m}$ where $(\mathbf{W}_L, \mathbf{T}_L) \leftarrow \text{Split}(\text{pp}, L)$.
 - If $L = 2^k \cdot \ell$ for $k \in \mathbb{N}$ and $\ell \in [2m]$, parse $\mathbf{M} = [\mathbf{M}_0 \mid \mathbf{M}_1]$ where $\mathbf{M}_{\beta} \in \mathbb{Z}_q^{n \times L/2}$. For $\beta \in \{0, 1\}$, compute $C_{\beta} = \operatorname{Com}^{m_x}(\operatorname{pp}, \mathbf{M}_{\beta}) \in \mathbb{Z}_q^{n \times m}$. Output $\mathbf{C} = \operatorname{Com}^{m_x}(\operatorname{pp}, [\mathbf{C}_0, | \mathbf{C}_1]) \in \mathbb{Z}_q^{n \times m}$.
- Ver^{mx}(pp, 1^{*L*}): On input pp = (**B**, **W**, **T**) and $L \in \mathbb{N}$:
 - If $L \leq 2m$, output $\mathbf{V}_L = \underline{\mathbf{T}}_L \cdot \mathbf{J}_L \in \mathbb{Z}_q^{m \times L \lceil \log q \rceil}$ where $(\mathbf{W}_L, \mathbf{T}_L) \leftarrow \text{Split}(\text{pp}, L)$.
 - If $L = 2^k \cdot \ell$ for $k \in \mathbb{N}$ and $\ell \in [2m]$, compute $\mathbf{V}_{L/2} = \operatorname{Ver}^{\mathsf{mx}}(\mathsf{pp}, 1^{L/2})$ and $\mathbf{V}_{2m} = \operatorname{Ver}^{\mathsf{mx}}(\mathsf{pp}, 1^{2m})$. Output $\mathbf{V}_L = \mathbf{V}_{2m}(\mathbf{I}_2 \otimes \mathbf{G}_m^{-1}(\mathbf{V}_{L/2})) \in \mathbb{Z}_q^{m \times L \lceil \log q \rceil}$.
- Open^{mx}(pp, M): On input pp = (B, W, T) and $M \in \mathbb{Z}_q^{n \times L}$:
 - If $L \leq 2m$, output $\mathbf{Z} = (\text{bits}(\mathbf{M})^{\mathsf{T}} \otimes \mathbf{I}_m) \overline{\mathbf{T}}_L \cdot \mathbf{J}_L \in \mathbb{Z}_q^{m \times L \lceil \log q \rceil}$ where $(\mathbf{W}_L, \mathbf{T}_L) \leftarrow \text{Split}(\text{pp}, L)$.
 - If $L = 2^k \cdot \ell$ for $k \in \mathbb{N}$ and $\ell \in [2m]$, parse $\mathbf{M} = [\mathbf{M}_0 \mid \mathbf{M}_1]$ where $\mathbf{M}_{\beta} \in \mathbb{Z}_q^{n \times L/2}$. For $\beta \in \{0, 1\}$, compute $\mathbf{Z}_{\beta} = \operatorname{Open}^{\mathsf{mx}}(\mathsf{pp}, \mathbf{M}_{\beta})$ and $\mathbf{C}_{\beta} = \operatorname{Com}^{\mathsf{mx}}(\mathsf{pp}, \mathbf{M}_{\beta})$. In addition, compute $\mathbf{V}_{L/2} = \operatorname{Ver}^{\mathsf{mx}}(\mathsf{pp}, \mathbf{1}^{L/2})$. Finally, compute $\mathbf{Z}' = \operatorname{Open}^{\mathsf{mx}}(\mathsf{pp}, [\mathbf{C}_0 \mid \mathbf{C}_1])$ and output $\mathbf{Z} = \mathbf{Z}' \cdot (\mathbf{I}_2 \otimes \mathbf{G}_m^{-1}(\mathbf{V}_{L/2})) + [\mathbf{Z}_0 \mid \mathbf{Z}_1] \in \mathbb{Z}_q^{\mathsf{m} \times L\lceil \log q \rceil}$.

A.1 Committing to Sparse Matrices and Supporting Local Openings

In this section, we give a proof of Lemma 3.10 and show that the [Wee25] matrix commitment scheme supports efficient commitment to sparse matrices, and moreover, there is an efficient algorithm to locally compute the columns of the verification matrix and the opening matrix.

Committing to sparse matrices. To support commitment to sparse matrices, we first observe that for all pp = (B, W, T) and all $L \in \mathbb{N}$, $Com^{mx}(pp, 0^{n \times L}) = 0^{n \times m}$. Thus, we can define the ComSparse^{mat}(pp, M) algorithm as follows:

- ComSparse^{mat}(pp, M): On input pp = (B, W, T) and a sparse matrix $M \in \mathbb{Z}_q^{n \times L}$:
 - If $\mathbf{M} = \mathbf{0}^{n \times L}$, return $\mathbf{0}^{n \times m}$.
 - If $L \leq 2m$, output $C = Com^{mx}(pp, M)$.
 - If $L = 2^k \cdot \ell$ for $k \in \mathbb{N}$ and $\ell \in [2m]$, parse $\mathbf{M} = [\mathbf{M}_0 \mid \mathbf{M}_1]$ where $\mathbf{M}_\beta \in \mathbb{Z}_q^{n \times L/2}$. For $\beta \in \{0, 1\}$, compute $C_\beta = \text{ComSparse}^{\text{mat}}(\text{pp}, \mathbf{M}_\beta) \in \mathbb{Z}_q^{n \times m}$. Output $\mathbf{C} = \text{Com}^{\text{mx}}(\text{pp}, [\mathbf{C}_0, \mid \mathbf{C}_1]) \in \mathbb{Z}_q^{n \times m}$.

Since $\text{Com}^{\text{mx}}(\text{pp}, \mathbf{0}^{n \times L}) = \mathbf{0}^{n \times m}$, we have that $\text{Com}\text{Sparse}^{\text{mat}}(\text{pp}, \mathbf{M}) = \text{Com}^{\text{mx}}(\text{pp}, \mathbf{M}) = \text{Com}^{\text{mat}}(\text{pp}, \mathbf{M})$. It suffices to analyze the running time of $\text{Com}\text{Sparse}^{\text{mat}}(\text{pp}, \mathbf{M})$. Take any matrix $\mathbf{M} \in \mathbb{Z}_q^{n \times L}$ where \mathbf{M} contains at most K non-zero columns. Throughout, we assume that the matrix \mathbf{M} (and its submatrices) are encoded in a sparse format (e.g., the description length of \mathbf{M} is poly($K, m, \log q, \log L$)). We show that $\text{Com}\text{Sparse}^{\text{mat}}(\text{pp}, \mathbf{M})$ runs in time poly($K, m, \log q, \log L$). It suffices to consider the case where L > 2m. Write $L = 2^k \cdot \ell$ where $k \in \mathbb{N}$ and $\ell \in [2m]$. Without loss of generality, suppose that $2\ell > 2m$ (if not, we can alternatively write $L = 2^{k-1} \cdot (2\ell)$). We now define a complete binary tree \mathcal{T} with height k as follows:

- First, parse the matrix $\mathbf{M} = [\mathbf{M}_0 | \cdots | \mathbf{M}_{2^{k}-1}]$, where $\mathbf{M}_i \in \mathbb{Z}_q^{n \times \ell}$. We associate the *i*th leaf node (from left to right) of \mathcal{T} with the sub-matrix \mathbf{M}_i .
- With each intermediate node σ in \mathcal{T} , let $\mathbf{M}_i, \mathbf{M}_{i+1}, \dots, \mathbf{M}_j$ be the matrices associated with the leaf nodes in the subtree rooted at σ (from left to right). We associate the matrix $\mathbf{M}_{\sigma} = [\mathbf{M}_i | \cdots | \mathbf{M}_j]$ with σ .

We now compute the running time of ComSparse^{mat}. To do so, we associate a value t_{σ} with each node $\sigma \in \mathcal{T}$, where t_{σ} is the running time of ComSparse^{mat}(pp, \mathbf{M}_{σ}) and \mathbf{M}_{σ} is the (sparse) matrix associated with σ . First, let $T = \text{poly}(K, m, \log q, \log L)$ be the running time needed to (1) check if $\mathbf{M} = \mathbf{0}$; (2) partition $\mathbf{M} = [\mathbf{M}_0 | \mathbf{M}_1]$; and (3) invoke $\text{Com}^{\text{mx}}(\text{pp}, \cdot)$ on a matrix with width 2m. This corresponds to the non-recursive components of the recursive step. Note that T also bounds the cost of a base case of the recursion. Then, we can define the running times t_{σ} as follows:

- If σ is associated with the zero matrix or σ is a leaf node, then $t_{\sigma} \leq T$.
- If σ is an internal node associated with a non-zero matrix, let σ_{LT} and σ_{RT} be its left and right children. Then, $t_{\sigma} \leq t_{\sigma_{LT}} + t_{\sigma_{RT}} + T$.

We now show the following inductive invariant. For a node $\sigma \in \mathcal{T}$, let K_{σ} denote the number of leaf nodes in the sub-tree rooted at σ that are associated with non-zero matrices, and let h_{σ} denote the height of the sub-tree rooted at σ (i.e., the length of the longest path in the sub-tree). Then, for all $\sigma \in \mathcal{T}$, we show that $t_{\sigma} \leq (1 + 2K_{\sigma}h_{\sigma})T$. We show this inductively:

- Consider a leaf node $\sigma \in \mathcal{T}$. In this case, $h_{\sigma} = 0$ and $t_{\sigma} \leq T = (1 + 2K_{\sigma}h_{\sigma})T$.
- Consider an internal node $\sigma \in \mathcal{T}$ that is associated with the zero matrix. Then $K_{\sigma} = 0$ and $t_{\sigma} \leq T = (1+2K_{\sigma}h_{\sigma})T$.
- Consider an internal node $\sigma \in \mathcal{T}$ that is associated with a nonzero matrix. This means $K_{\sigma} \geq 1$ Then, $t_{\sigma} \leq t_{\sigma_{LT}} + t_{\sigma_{RT}} + T$. By definition, $h_{\sigma_{LT}} = h_{\sigma} - 1$ and $K_{\sigma_{LT}} + K_{\sigma_{RT}} = K_{\sigma}$. Now, by the inductive hypothesis,

$$\begin{split} t_{\sigma} &\leq t_{\sigma_{\mathrm{LT}}} + t_{\sigma_{\mathrm{RT}}} + T \\ &\leq (1 + 2K_{\sigma_{\mathrm{LT}}}(h_{\sigma} - 1))T + (1 + 2K_{\sigma_{\mathrm{RT}}}(h_{\sigma} - 1))T + T \\ &= (1 + 2K_{\sigma}h_{\sigma})T + 2(1 - K_{\sigma})T \\ &\leq (1 + 2K_{\sigma}h_{\sigma})T, \end{split}$$

since $1 - K_{\sigma} \leq 0$.

By induction on \mathcal{T} , we conclude for all $\sigma \in \mathcal{T}$, $t_{\sigma} \leq (1 + 2K_{\sigma}h_{\sigma})T$. The running time of ComSparse^{mat}(pp, **M**) is then bounded by $(1 + 2Kk)T = \text{poly}(K, m, \log q, \log L)$, as required.

Local access to the verification matrix. Next, we show that given pp and *L*, we can compute any single column of the verification matrix \mathbf{V}_L in time poly(*m*, log *q*, log *L*). We first define a local version of Ver^{mx}:

- VerLocal^{mx}(pp, *L*, *i*): On input the public parameters pp, the length $L \in \mathbb{N}$ in *binary*, and an index $i \leq L \lceil \log q \rceil$:
 - If $L \leq 2m$, compute $\mathbf{V}_L = \operatorname{Ver}^{mx}(pp, 1^L)$ and output the *i*th column $\mathbf{v}_{L,i}$ of \mathbf{V}_L .
 - If $L = 2^k \cdot \ell$ for $k \in \mathbb{N}$ and $\ell \in [2m]$, first compute $\mathbf{V}_{2m} = \operatorname{Ver}^{\mathrm{mx}}(\mathrm{pp}, 1^{2m}) \in \mathbb{Z}_q^{m \times 2m \lceil \log q \rceil}$. Parse $\mathbf{V}_{2m} = [\mathbf{V}_{2m,\mathrm{LT}} \mid \mathbf{V}_{2m,\mathrm{RT}}]$ where $\mathbf{V}_{2m,\mathrm{LT}}, \mathbf{V}_{2m,\mathrm{RT}} \in \mathbb{Z}_q^{m \times m \lceil \log q \rceil}$. Then output

$$\mathbf{v}_{L,i} = \begin{cases} \mathbf{V}_{2m,\mathrm{LT}} \cdot \mathbf{G}_m^{-1}(\mathbf{v}_{L/2,i}) & i \le L \left\lceil \log q \right\rceil / 2 \\ \mathbf{V}_{2m,\mathrm{RT}} \cdot \mathbf{G}_m^{-1}(\mathbf{v}_{L/2,i-L \left\lceil \log q \right\rceil / 2}) & i > L \left\lceil \log q \right\rceil / 2, \end{cases}$$

where $\mathbf{v}_{L/2,i} = \text{VerLocal}^{\text{mx}}(\text{pp}, L/2, i)$ and $\mathbf{v}_{L/2,i-L\lceil \log q \rceil/2} = \text{VerLocal}^{\text{mx}}(\text{pp}, L/2, i-L\lceil \log q \rceil/2)$.

First, we analyze the running time of VerLocal^{mx}. Let T(n, m, q, L) be the running time of VerLocal^{mx}(pp, *L*, *i*) on any set of public parameters pp with lattice parameters (n, m, q), length *L*, and any index $i \le L \lceil \log q \rceil$.

• If $L \leq 2m$, then $T(n, m, q, L) = p_1(m, \log q)$ for a fixed polynomial p_1 .³

³Recall that we are only considering public parameters pp where $m \ge 2n \log q$. Thus, any polynomial dependence in the running time on n can be absorbed by a poly(m) factor.

• If L > 2m, then the running time is $p_2(m, \log q, \log L) + T(n, m, q, L/2)$ for some fixed polynomial p_2 .

Thus,

$$T(n, m, q, L) \leq \log L \cdot p_2(m, \log q, \log L) + p_1(m, \log q) = \mathsf{poly}(m, \log q, \log L),$$

as required. It suffices to show that $VerLocal^{mx}(pp, L, i)$ is correct. We proceed by (strong) induction on L.

- If $L \leq 2m$, correctness holds by construction.
- Suppose L > 2m. By definition of Ver^{mx}, this means

$$\mathbf{V}_{L} = \mathbf{V}_{2m}(\mathbf{I}_{2} \otimes \mathbf{G}_{m}^{-1}(\mathbf{V}_{L/2})) = \begin{bmatrix} \mathbf{V}_{2m,\text{LT}} \mid \mathbf{V}_{2m,\text{RT}} \end{bmatrix} \cdot \begin{bmatrix} \mathbf{G}_{m}^{-1}(\mathbf{V}_{L/2}) & \mathbf{0} \\ \mathbf{0} & \mathbf{G}_{m}^{-1}(\mathbf{V}_{L/2}) \end{bmatrix}.$$

In particular, the i^{th} column $\mathbf{v}_{L,i}$ of \mathbf{V}_L is then

$$\mathbf{v}_{L,i} = \begin{cases} \mathbf{V}_{2m,\text{LT}} \cdot \mathbf{G}_m^{-1}(\mathbf{v}_{L/2,i}) & i \le L \lceil \log q \rceil / 2 \\ \mathbf{V}_{2m,\text{RT}} \cdot \mathbf{G}_m^{-1}(\mathbf{v}_{L/2,i-L \lceil \log q \rceil / 2}) & i > L \lceil \log q \rceil / 2, \end{cases}$$

where $\mathbf{v}_{L/2,i}$ is the *i*th row of $\mathbf{V}_{L/2} = \text{Ver}^{\text{mx}}(\text{pp}, L/2)$ and $\mathbf{v}_{L/2,i-L\lceil \log q \rceil/2}$ is the $(i - L \lceil \log q \rceil/2)^{\text{th}}$ row of $\mathbf{V}_{L/2}$. By the inductive hypothesis,

$$\mathbf{v}_{L/2,i} = \text{VerLocal}^{\text{mx}}(\text{pp}, L/2, i)$$
$$\mathbf{v}_{L/2,i-L\lceil \log q \rceil/2} = \text{VerLocal}^{\text{mx}}(\text{pp}, L/2, i - L\lceil \log q \rceil/2)$$

The claim now follows by (strong) induction on *L*.

Given VerLocal^{mx}(pp, L, i), it is straightforward to construct the local version of Ver^{mat}:

• VerLocal^{mat}(pp, L, i): Output VerLocal^{mx}(pp, L, $(i - 1) \lceil \log q \rceil + 1$).

Correctness follows by construction. Namely, $\operatorname{Ver}^{\operatorname{mat}}(\operatorname{pp}, L)$ outputs $\mathbf{V}_L \cdot \mathbf{G}_L^{-1}(\mathbf{I}_L)$, where $\mathbf{V}_L = \operatorname{Ver}^{\operatorname{mx}}(\operatorname{pp}, L)$. By construction, the *i*th column of $\mathbf{G}_L^{-1}(\mathbf{I}_L)$ is the unit vector $\mathbf{u}_{(i-1)} \cdot \lceil \log q \rceil + 1 \in \mathbb{Z}_q^{L \lceil \log q \rceil \times L}$. Correspondingly, the *i*th column of $\mathbf{V}_L \cdot \mathbf{G}_L^{-1}(\mathbf{I}_L)$ is then the $((i-1) \cdot \lceil \log q \rceil + 1)^{\operatorname{th}}$ -column of \mathbf{V}_L .

Computing local openings. Finally, combining the above two procedures, we obtain an analogous algorithm for computing any single column of the opening $\mathbf{Z} = \operatorname{Open}^{\mathsf{mx}}(\mathsf{pp}, \mathbf{M})$ in time $\mathsf{poly}(K, m, \log q, \log L)$, where *K* is the number of non-zero columns of \mathbf{M} . We start by defining the algorithm $\operatorname{OpenLocal}^{\mathsf{mx}}$:

- OpenLocal^{mx}(pp, M, *i*): On input pp = (B, W, T), a sparse matrix $M \in \mathbb{Z}_q^{n \times L}$, and an index $i \leq L \lceil \log q \rceil$:
 - If $L \leq 2m$, compute **Z** = Open^{mx}(pp, **M**) and output the *i*th column **z**_{*i*} of **Z**.
 - If L = 2^k · ℓ for k ∈ ℕ and ℓ ∈ [2m], parse M = [M₀ | M₁] where M_β ∈ Z^{n×L/2}_q. Then compute the following:
 * For β ∈ {0, 1}, let C_β = ComSparse^{mat}(pp, M_β). Then compute Z' = Open^{mx}(pp, [C₀ | C₁]). Parse Z' = [Z'_{LT} | Z'_{RT}] where Z'_{LT}, Z'_{RT} ∈ Z^{m×m[log q]}.
 - * Compute the vector $\hat{\mathbf{v}} \in \mathbb{Z}_q^m$ as

$$\hat{\mathbf{v}} = \begin{cases} \operatorname{VerLocal^{mx}}(\operatorname{pp}, L/2, i) & i \le L \left\lceil \log q \right\rceil / 2 \\ \operatorname{VerLocal^{mx}}(\operatorname{pp}, L/2, i - L \left\lceil \log q \right\rceil / 2) & i > L \left\lceil \log q \right\rceil / 2. \end{cases}$$

* Compute the vector $\hat{\mathbf{z}} \in \mathbb{Z}_q^m$ as

$$\hat{\mathbf{z}} = \begin{cases} \text{OpenLocal}^{\text{mx}}(\text{pp}, \mathbf{M}_0, i) & i \le L \lceil \log q \rceil / 2 \\ \text{OpenLocal}^{\text{mx}}(\text{pp}, \mathbf{M}_1, i - L \lceil \log q \rceil / 2) & i > L \lceil \log q \rceil / 2. \end{cases}$$

Output the vector

$$\mathbf{z}_{i} = \begin{cases} \mathbf{Z}_{\text{LT}}' \cdot \mathbf{G}_{m}^{-1}(\hat{\mathbf{v}}) + \hat{\mathbf{z}} & i \leq L \lceil \log q \rceil / 2 \\ \mathbf{Z}_{\text{RT}}' \cdot \mathbf{G}_{m}^{-1}(\hat{\mathbf{v}}) + \hat{\mathbf{z}} & i > L \lceil \log q \rceil / 2. \end{cases}$$

First, we analyze the running time of OpenLocal^{mx}. Let T(n, m, q, L, K) be the running time of OpenLocal^{mx}(pp, M, *i*) on any set of public parameters pp with lattice parameters (n, m, q), a matrix $\mathbf{M} \in \mathbb{Z}_q^{n \times L}$ with at most *K* non-zero columns, and any index $i \leq L \lceil \log q \rceil$.

- If $L \le 2m$, then $T(n, m, q, L) = p_1(m, \log q)$ for a fixed polynomial p_1 .
- If L > 2m, then OpenLocal^{mx}(pp, M, *i*) computes the following quantities:
 - Computing ComSparse^{mat}(pp, \mathbf{M}_{β}) for $\beta \in \{0, 1\}$, requires time poly(*K*, *m*, log *q*, log *L*). Computing **Z**' from C₀ and C₁ requires time poly(*m*, log *q*).
 - Computing the vector $\hat{\mathbf{v}}$ requires time poly(m, log q, log L).
 - Computing the vector $\hat{\mathbf{z}}$ requires time T(n, m, q, L/2, K).
 - Computing the final output from \mathbf{Z}' , $\hat{\mathbf{v}}$, and $\hat{\mathbf{z}}$ requires time poly(m, log q).

The total running time in this case is then $T(n, m, q, L/2, K) = p_2(K, m, \log q, \log L) + T(n, m, q, L/2, K)$, where p_2 is a fixed polynomial.

This means

$$T(n, m, q, L, K) = \log L \cdot p_2(K, m, \log q, \log L) + p_1(m, \log q) = \text{poly}(K, m, \log q, \log L).$$

To complete the proof, we show that $OpenLocal^{mx}(pp, M, i)$ is correct. Similar to the previous case, we proceed by (strong) induction on *L*:

- If $L \leq 2m$, correctness holds by construction.
- Suppose L > 2m. Then, the opening algorithm $Open^{mx}(pp, \mathbf{M})$ computes $\tilde{C}_{\beta} = Com^{mx}(pp, \mathbf{M}_{\beta})$ for $\beta \in \{0, 1\}$. and $\tilde{Z}' = Open^{mx}(pp, [\tilde{C}_0 | \tilde{C}_1])$. It also computes $\tilde{Z}_{\beta} = Open^{mx}(pp, \mathbf{M}_{\beta})$ for $\beta \in \{0, 1\}$, $\tilde{V}_{L/2} = Ver^{mx}(pp, 1^{L/2})$, and finally, sets $\tilde{Z} = \tilde{Z}' \cdot (\mathbf{I}_2 \otimes \tilde{G}^{-1}_m(\mathbf{V}_{L/2})) + [\tilde{Z}_0 | \tilde{Z}_1]$. Consider the *i*th column \tilde{z}_i of \tilde{Z} . We can write it as

$$\tilde{\mathbf{z}}_{i} = \begin{cases} \tilde{\mathbf{Z}}'_{\text{LT}} \cdot \mathbf{G}_{m}^{-1}(\tilde{\mathbf{v}}_{L/2,i}) + \tilde{\mathbf{z}}_{0,i} & i \leq L \left\lceil \log q \right\rceil / 2 \\ \tilde{\mathbf{Z}}'_{\text{KT}} \cdot \mathbf{G}_{m}^{-1}(\tilde{\mathbf{v}}_{L/2,i-L \left\lceil \log q \right\rceil / 2}) + \tilde{\mathbf{z}}_{1,i-L \left\lceil \log q \right\rceil / 2} & i > L \left\lceil \log q \right\rceil / 2, \end{cases}$$

where $\tilde{\mathbf{Z}} = [\mathbf{Z}'_{LT} \mid \tilde{\mathbf{Z}}'_{RT}]$, $\tilde{\mathbf{v}}_{L/2,i}$ denotes the *i*th column of $\tilde{\mathbf{V}}_{L/2}$ and $\tilde{\mathbf{z}}_{0,i}, \tilde{\mathbf{z}}_{1,i}$ denote the *i*th column of $\tilde{\mathbf{Z}}_0$ and $\tilde{\mathbf{Z}}_1$, respectively. Consider now OpenLocal^{mx}(pp, M, *i*). By correctness of ComSparse^{mat}, we have that

$$C_{\beta} = \text{ComSparse}^{\text{mat}}(\text{pp}, \mathbf{M}_{\beta}) = \text{Com}^{\text{mx}}(\text{pp}, \mathbf{M}_{\beta}) = C_{\beta}$$

for $\beta \in \{0, 1\}$. This means

$$\mathbf{Z}' = \operatorname{Open}^{\mathsf{mx}}(\mathsf{pp}, [\mathbf{C}_0 \mid \mathbf{C}_1]) = \operatorname{Open}^{\mathsf{mx}}(\mathsf{pp}, [\tilde{\mathbf{C}}_0 \mid \tilde{\mathbf{C}}_1]) = \tilde{\mathbf{Z}}'.$$

By correctness of VerLocal^{mx}, we have

 $\hat{\mathbf{v}} = \text{VerLocal}^{\text{mx}}(\text{pp}, L/2, i) = \tilde{\mathbf{v}}_{L/2, i}$ if $i \le L \lceil \log q \rceil / 2$

$$\hat{\mathbf{v}} = \text{VerLocal}^{\text{mx}}(\text{pp}, L/2, i - L \lceil \log q \rceil / 2) = \tilde{\mathbf{v}}_{L/2, i - L \lceil \log q \rceil / 2} \quad \text{if } i > L \lceil \log q \rceil / 2$$

By the inductive hypothesis, we have

$\hat{\mathbf{z}} = \text{OpenLocal}^{\text{mx}}(\text{pp}, \mathbf{M}_0, i) = \tilde{\mathbf{z}}_{0,i}$	if $i \leq L \lceil \log q \rceil / 2$
$\hat{\mathbf{z}} = \text{OpenLocal}^{\text{mx}}(\text{pp}, \mathbf{M}_1, i - L \lceil \log q \rceil / 2) = \tilde{\mathbf{z}}_{1, i - L \lceil \log q \rceil / 2}$	if $i > L \lceil \log q \rceil / 2$.

Putting the pieces together, we now have

$$\begin{aligned} \mathbf{z}_{i} &= \mathbf{Z}_{\mathrm{LT}}' \cdot \mathbf{G}_{m}^{-1}(\hat{\mathbf{v}}) + \hat{\mathbf{z}} = \tilde{\mathbf{Z}}_{\mathrm{LT}}' \cdot \mathbf{G}_{m}^{-1}(\tilde{\mathbf{v}}_{L/2,i}) + \tilde{\mathbf{z}}_{0,i} = \tilde{\mathbf{z}}_{i} & \text{if } i \leq L \left\lceil \log q \right\rceil / 2 \\ \mathbf{z}_{i} &= \mathbf{Z}_{\mathrm{RT}}' \cdot \mathbf{G}_{m}^{-1}(\hat{\mathbf{v}}) + \hat{\mathbf{z}} = \tilde{\mathbf{Z}}_{\mathrm{RT}}' \cdot \mathbf{G}_{m}^{-1}(\tilde{\mathbf{v}}_{L/2,i-L \left\lceil \log q \right\rceil / 2}) + \tilde{\mathbf{z}}_{1,i-L \left\lceil \log q \right\rceil / 2} = \tilde{\mathbf{z}}_{i} & \text{if } i > L \left\lceil \log q \right\rceil / 2. \end{aligned}$$

$$\mathbf{z}_{i} = \mathbf{Z}_{RT} \cdot \mathbf{G}_{m} (\mathbf{v}) + \mathbf{z} = \mathbf{Z}_{RT} \cdot \mathbf{G}_{m} (\mathbf{v}_{L/2, i-L[\log q]/2}) + \mathbf{z}_{1, i-L[\log q]/2} = \mathbf{z}_{i} \qquad \text{if } t > L |\log q|$$

Thus for all $i \leq L \lceil \log q \rceil$, we have that $\mathbf{z}_i = \tilde{\mathbf{z}}_i$, and correctness holds.

Given OpenLocal^{mx}(pp, L, i), it is straightforward to construct the local version of Open^{mat}:

• OpenLocal^{mat}(pp, M, *i*): Output OpenLocal^{mx}(pp, M, $(i - 1) \lceil \log q \rceil + 1$).

Correctness via the same analysis as for VerLocal^{mat} above.