Fully Secure Lattice-based ABE from Noisy Linear Functional Encryption

Geng Wang, Ming Wan, Zhen Liu and Dawu Gu

School of Electronic Information and Electrical Engineering
Shanghai Jiao Tong University, 100072, P.R.China

Abstract. Constructing lattice-based fully secure attribute-based encryption (ABE) has always been a challenging task. Although there are many selective secure ABE schemes from the hardness of learning with errors (LWE) problem, it is hard to extend them to fully security, since the dual system technique in pairing-based cryptography cannot be applied to lattice-based constructions.
In this paper, we take a different approach: constructing fully secure ABE from another primitive called noisy linear functional encryption (NLinFE) which can be constructed from LWE problem. We give a fully secure ciphertext-policy ABE scheme for CNF formulae which security relies on the security of NLinFE and hardness of LWE. Since current constructions for NLinFE only satisfy bounded collusion security, our resulting scheme is also bounded collusion only, but it can be easily extended into unbounded security if unbounded NLinFE can be shown to exist. Also, since existing NLinFE schemes are inefficient, we give a new construction for NLinFE with better efficiency, hence our ABE construction is more efficient than other existing bounded collusion ABE/FE schemes.

Keywords: Attribute-based encryption, Noisy Linear Functional Encryption, LWE, Lattice-based cryptography

1 Introduction

Attribute-based Encryption (ABE for short) was first brought by Sahai and Waters in 2005 [37]. In an ABE scheme, the decryption is correct only if the provided attribute set satisfies a certain access policy. By using different types of access policies, ABE can handle flexible access control matters, without using complex key distribution techniques. There are mainly two types of ABE, one is called key-policy ABE (KP-ABE) [25], other is called ciphertext-policy ABE (CP-ABE) [10]. In KP-ABE, the access policy is embedded in the decryption key, while the ciphertext is related to a set of attributes; in CP-ABE, the access policy is embedded in the ciphertext, and attributes are related to the decryption key, held by the users. In [13], ABE is considered as a special case of a more generalized primitive called functional encryption (FE), which given an encrypted data $Enc(x)$, calculate the function output $f(x)$ of an encrypted message for a certain class of function class $f \in \mathcal{F}$.

Most of these early ABE schemes [25, 10, 33, 19, 41, 36] are from a weak security model, which is called selective security. In a selective security model, the adversary must first give the challenge policy (for CP-ABE) or challenge attribute set (for KP-ABE) before it was allowed to get the public key and query for secret keys. It is easy to see that the selective security model greatly restricts the ability of the adversary, and cannot handle many types of real world attacks.

Many researchers focus on removing the restriction to get full security for ABE schemes. Many different approaches have been proposed, but the most successful one among them is the dual-system encryption method, given by Waters in 2009 [40]. Although the original method is for IBE and HIBE, which are only simplified versions of ABE, it was soon used to construct fully secure ABE schemes for various access policies, as in [29, 30, 9, 42, 27, 18, 28].

The schemes above are constructed in bilinear groups, which suffer from quantum attacks. Recently, many researchers have been working on constructing ABE schemes using lattice assumptions, such as learning with error (LWE) problem [4, 3, 14]. Lattice-based ABE schemes are not only quantum secure, but also more powerful than schemes in bilinear groups, as they support much richer classes of access policies, even for arbitrary circuits [23, 12, 16].

However, the existing schemes are only selective secure, except for some recent works [38, 26] that can only support a quite weak class of access policies. Since the original dual-system method is highly related to the properties of pairing in bilinear groups, it was not known whether there exists an analogue for dual-system in lattice, which could be used to prove the full security of lattice-based ABE schemes. This question has been raised in many earlier works, and has been considered as a long time open problem in lattice-based cryptography. In this paper, we present a similar method for constructing fully secure lattice-based ABE schemes using noisy linear functional encryption (NLinFE), also give a CP-ABE scheme supporting CNF policies and prove its full security in the standard model.

The security properties of our constructed CP-ABE scheme rely on the security properties of the underlying NLinFE scheme. With bounded collusion public-key NLinFE in [7] (eprint version), and secret-key NLinFE in [1, 6], we get both a fully secure bounded collusion CP-ABE and a fully secure secret-key CP-ABE from lattice assumptions. We also give a construction for NLinFE with weaker security which only supports random key queries, but outperforms [7] in the ciphertext size, hence get a bounded collusion CP-ABE which has shorter ciphertexts. If public key unbounded NLinFE can be proven to exist, we can simply construct a fully secure unbounded CP-ABE supporting CNF policies, which beats all current results.

1.1 Related Works

There are currently a few researches working on lattice-based fully secure identity-based encryption (IBE)[2, 17, 15, 43], which can be considered as ABE which access policy is point function. These security proofs rely on various primitives, such as admissible hash or pseudorandom functions. It is not known how these

techniques can be used for other access policies. In [23], the authors claimed that using a result from [11], the selective security of the scheme can be extended to full security assuming the subexponential hardness of LWE. Despite the non-standardness of the hardness assumption, it seems that this method cannot be extended into CP-ABE schemes. In [16, 24], the authors focused on semi-adaptive security of ABE schemes. Although stronger than selective security, it is still weaker than full security.

In [38], the author gave the first fully secure ABE scheme (other than IBE) from standard LWE assumption using a new primitive called conforming cPRF, which is a huge step forward. However, the access policy is only t-CNF for a constant t, which means that each clause exactly contains t literals. A similar construction is from [26] which supports inner products. These are weaker than our access policy, which is (unrestricted) CNF. The authors claimed that the access policy is only related to the expressibility of the conforming cPRF, however, constructing conforming cPRF supporting various access policies seems to be extremely difficult. Despite the complexity in the conforming cPRF itself, the function needs to be evaluated through key-homomorphic encryption [12], which makes the scheme almost impossible for implementation, while our scheme is more simple and efficient for implementation.

In [39], a fully secure decentralized ABE is constructed from inner product encryption based on LWE assumption [5]. The idea of using NLinFE in our scheme is partly borrowed from the use of (non-approximate) IPE in their work. Their construction is also with bounded collusion. We note that our result highly overlaps with this work. (Using similar techniques, it seems that our scheme can also be made decentralized, but we will not discuss that in this paper.) However, the number of calls to IPE is related to the vector dimension of IPE, hence the number of key queries, while our number of calls to NLinFE is only related to the access policy, which leads to a more efficient scheme.

It seems that fully secure ABE with bounded collusion can also be instantiated by fully secure functional encryption with bounded collusion as in [22, 21, 5, 7]. However, since these constructions are for arbitrary polynomial circuits, their schemes have larger ciphertexts, and are more complex and hard to implement compared with ours.

The table below compares the ciphertext size between this work and other ABE or FE schemes with bounded collusion, where Q is the number of key queries. We also point out that in [8], the authors introduced a new technique on any bounded collusion schemes, such that the ciphertext size which is polynomial in Q becomes polynomial in the security parameter λ and linear in Q (e.g. the ciphertext size of our construction becomes $O(Q\lambda \log \lambda)$ instead of $O(Q \log Q)$). This does not change the fact that our ciphertext size is smaller than others.

2 Preliminaries

Notations. $x \leftarrow \chi$ for a distribution χ means that x is sampled from χ. $x \leftarrow X$ for a set X means that x is uniformly random chosen from X. For any odd

Table 1: Compare with other bounded collusion ABE and FE

	Supporting Functions	Ciphertext Size	CT Size with [AV19]
GVW12[22]	P/poly	$O(Q^4)$	$O(Q\lambda^4)$
AR17[7]	P/poly	$O(Q^2)$	$O(Q\lambda^2)$
WFL19[39]	0-1 LSSS	$O(Q^2 \log Q)$	$O(Q\lambda^2 \log \lambda)$
This work	CNF	$O(Q \log Q)$	$O(Q\lambda \log \lambda)$

modulus q, $\mathbb{Z}_q$ and the operation mod q takes value from $[-\frac{q-1}{2}, \frac{q-1}{2}]$. We say that ϵ is negligible in λ, if $\epsilon < 1/\Omega(\lambda^c)$ for any $c > 0$ for sufficently large λ. For two distributions X, Y, let $\Delta(X, Y)$ be the statistical distance between X and Y. $\|.\|$ is the 2-norm, while $\|.\|_\infty$ is the infinity norm.

2.1 Conjunctive Normal Form

Definition 2.1. *Let* $\mathbf{L}$ *be a set of literals (a literal is either* α *or* $\neg\alpha$ *for some variable* α*), and* $T_1, ..., T_k \subseteq \mathbf{L}$ *be a set of* clauses.

A conjunctive normal form (CNF) is a boolean function $f = \bigwedge_{i=1}^{k}(\bigvee T_i)$*, which inputs a set of literals* $L \subseteq \mathbf{L}$ *(for each variable* α*,* α *and* $\neg\alpha$ *not both in* L*), and outputs the value* $f(L) = \bigwedge_{i=1}^{k}(\bigvee T_i(L))$*. Here* $\bigvee T_i(L) = 1$ *if and only if* $T_i \cap L \neq \emptyset$*.*

Let $l = |\mathbf{L}|$*, and we label the literals in* $\mathbf{L}$ *by 1 to* l*.*

Note that we do not consider the relationship between α and $\neg\alpha$, and simply let them be two different elements. Such representation does not lower the expressibility of CNF policy. In fact, our definition is stronger than boolean formulas: for an attribute (literal) set L, we allow that both α and $\neg\alpha$ are not in L, which means that we "do not care" the value of α, as in [19].

2.2 Ciphertext-Policy Attribute-based Encryption

Definition 2.2. *A CP-ABE scheme for CNF formula* f *consists of four algorithms* $(\mathsf{Setup}, \mathsf{Enc}, \mathsf{KeyGen}, \mathsf{Dec})$*:*

- $\mathsf{Setup}(1^\lambda) \to (\mathsf{mpk}, \mathsf{msk})$*: The setup algorithm gets as input the security parameter* λ*, and outputs the public parameter* mpk*, and the master key* msk*.*
- $\mathsf{Enc}(\mathsf{mpk}, f, m) \to \mathsf{ct}_f$*: The encryption algorithm gets as input* mpk*, a CNF formula* f*, and a message* $m \in \mathcal{M}$*. It outputs a ciphertext* ct_f*. Note that the policy is known if we know the ciphertext.*
- $\mathsf{KeyGen}(\mathsf{msk}, L) \to \mathsf{sk}_L$*: The key generation algorithm gets as input* msk *and a set of literals* L*. It outputs a secret key* sk_L*.*
- $\mathsf{Dec}(\mathsf{sk}_L, \mathsf{ct}_f) \to m$*: The decryption algorithm gets as input a secret key and a ciphertext, and outputs either* $\perp$ *or a message* $m \in \mathcal{M}$*.*

The CP-ABE scheme is correct if and only if the decryption algorithm returns the correct message when $f(L) = 1$*, and returns* $\perp$ *when* $f(L) = 0$*.*

Definition 2.3. *A CP-ABE scheme is fully secure, if for any adversary, the advantage of winning the following CPA-CP-ABE game is negligible:*

Setup. *The challenger runs* ***Setup*** *and gives the adversary* mpk.

Phase 1. *The adversary submits a set of literals L for a* ***KeyGen*** *query. These queries can be repeated adaptively.*

Challenge. *The adversary submits two messages m_0 and m_1 of equal length, and a CNF formula f, and $f(L) = 0$ for all previously queried L. The challenger chooses a random bit $b \in \{0, 1\}$, and encrypts m_b under f. The encrypted ciphertext ct_f is returned to the adversary.*

Phase 2. *The adversary repeats* ***Phase 1*** *to get more secret keys. Each queried L must have $f(L) = 0$.*

Guess. *The adversary outputs a guess b' for b.*

The advantage of an adversary $\mathcal{A}$ in the CPA-CP-ABE game is defined by $\mathsf{Adv}_{\mathcal{A}}^{\mathrm{ABE}}(\lambda) = |Pr[b' = b] - 1/2|$.

2.3 Lattice and Smoothing Parameters

Definition 2.4. *Let $\mathbf{b}_1, ..., \mathbf{b}_n$ be a set of vectors in $\mathbb{R}^m$ for $m \geq n$. A lattice Λ is defined as $\{\sum_{i=1}^{n} c_i \mathbf{b}_i : c_1, ..., c_n \in \mathbb{Z}\}$, and $\mathbf{b}_1, ..., \mathbf{b}_n$ is called a basis of Λ.*

Definition 2.5. *[34]*

Given its center $\mathbf{c} \in \mathbb{Z}^m$, for any vector $\mathbf{x} \in \mathbb{Z}^m$, let $\rho_{s,\mathbf{c}}(\mathbf{x}) = \exp(-\pi\|\mathbf{x} - \mathbf{c}\|^2/s^2)$.

For a lattice Λ and $\mathbf{c} \in \mathbb{Z}^m$, the discrete Gaussian distribution $D_{\Lambda,s,\mathbf{c}}$ is defined as:

$$D_{\Lambda,s,\mathbf{c}}(\mathbf{x}) = \frac{\rho_{s,\mathbf{c}}(\mathbf{x})}{\sum_{\mathbf{v}\in\Lambda} \rho_s(\mathbf{v})}.$$

$D_{\Lambda,s,\mathbf{c}}$ is sometimes also written as $D_{\Lambda+\mathbf{c},s}$, $\Lambda + \mathbf{c}$ is a lattice coset. We also write $\rho_s(\Lambda) = \sum_{\mathbf{v}\in\Lambda} \rho_s(\mathbf{x})$.

There is an important property for lattice called smoothing parameter, defined as below:

Definition 2.6. *[32] For any n-dimensional lattice Λ and $\epsilon > 0$, the smoothing parameter $\eta_\epsilon(\Lambda)$ is the smallest s such that $\rho(\Lambda^* - \{0\}) \leq \epsilon$, where $\Lambda^* = \{\mathbf{x} \in \mathbb{R}^n : \forall \mathbf{v} \in \Lambda, \langle \mathbf{x}, \mathbf{v} \rangle \in \mathbb{Z}\}$.*

The following properties for lattice are related to its smoothing parameter, and will be used in our proof.

Lemma 2.1. *[20] For any n-dimensional lattice Λ, $\epsilon > 0$, and any $\omega(\sqrt{\log n})$ function, there is a negligible $\epsilon(n)$ for which $\eta_\epsilon(\Lambda) \leq \omega(\sqrt{\log n})/\lambda_1^\infty(\Lambda^*)$, $\lambda_1^\infty(\Lambda^*)$ is the length of shortest non-zero vector in Λ^*.*

Lemma 2.2. *[20] For any n-dimensional lattice Λ, $\mathbf{c} \in \mathrm{span}(\Lambda)$, real $\epsilon \in (0, 1)$, and $s \geq \eta_\epsilon(\Lambda)$, $Pr_{\mathbf{x}\leftarrow D_{\Lambda,s,\mathbf{c}}}[\|\mathbf{x} - \mathbf{c}\| > s\sqrt{n}] \leq \frac{1+\epsilon}{1-\epsilon} \cdot 2^{-n}$.*

Lemma 2.3. *[38] Let $y \in \mathbb{Z}$, the statistical difference between $D_{\mathbb{Z},\sigma}$ and $D_{\mathbb{Z},\sigma} + y$ is at most $|y|/\sigma$.*

2.4 Lattice Trapdoor and Learning with Error

The following lemma in [20, 31] shows that there exists a trapdoor and a preimage sampling algorithm for discrete Gaussian distribution.

Lemma 2.4. *[20, 31] There is an efficient randomized algorithm* $\mathsf{TrapSamp}(1^n, 1^m, q)$ *that, given* $n \geq 1$, $q \geq 2$, $m = \Omega(n \log q)$, *outputs* $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ *and a "trapdoor"* T *such that the distribution of* $\mathbf{A}$ *is* $\mathrm{negl}(n)$*-close to uniform.*

Moreover, let $\Lambda_{\mathbf{u}}^{\perp}(\mathbf{A}) = \{\mathbf{x} : \mathbf{A}\mathbf{x} = \mathbf{u}\}$ *(which is a lattice coset). Then there is an efficient randomized algorithm* $\mathsf{SamplePre}$ *that for any* $\mathbf{u} \in \mathbb{Z}_q^n$, $s = \Omega(\sqrt{n \log q})$, $\mathsf{SamplePre}(\mathbf{A}, \mathbf{T}, \mathbf{u}, s)$ *outputs a vector* $\mathbf{r} \in \mathbb{Z}^m$, *which distribution is statistically close to* $D_{\Lambda_{\mathbf{u}}^{\perp}(\mathbf{A}), s}$ *(with negligible distance).*

We sometimes omit the parameter s if there is no confusion.

The following lemma is required for our security proof:

Lemma 2.5. *Let* $(\mathbf{A}, \mathbf{T}) \leftarrow \mathsf{TrapSamp}(1^n, 1^m, q)$, *and* $(\mathbf{A}', \mathbf{T}') \leftarrow \mathsf{TrapSamp}(1^{n'}, 1^m, q)$, $n' > n$, *and we write* $\mathbf{A}' = \binom{\bar{\mathbf{A}}}{\tilde{\mathbf{A}}}$, $\bar{\mathbf{A}} \in \mathbb{Z}_q^{n \times m}$, *and* $\tilde{\mathbf{A}} \in \mathbb{Z}_q^{(n'-n) \times m}$. *Then there exists* $s > 0$ *such that the following two distribution are statistically indistinguishable:*

- $\mathbf{A}, \mathbf{x} \leftarrow \mathsf{SamplePre}(\mathbf{A}, \mathbf{T}, \mathbf{u}, s)$;
- $\bar{\mathbf{A}}, \bar{\mathbf{x}} \leftarrow \mathsf{SamplePre}(\mathbf{A}', \mathbf{T}', \binom{\mathbf{u}}{\mathbf{b}}, s)$, *where* $\mathbf{b} \leftarrow \mathbb{Z}_q^{n'-n}$.

Proof. See Appendix A. □

Now we introduce our hardness assumption: the (decisional) learning with error (LWE) problem, first introduced in [35]. It has the nice property called worst-case to average-case reduction: solving LWE on the average is as hard as (quantumly) solving GapSVP and SIVP problems in the worst case.

Definition 2.7 (LWE problem). *[35] For a vector* $\mathbf{s} \in \mathbb{Z}_q^n$ *called the secret, the LWE distribution* $A_{\mathbf{s},\chi}$ *over* $\mathbb{Z}_q^n \times \mathbb{Z}_q$ *is sampled by choosing* $\mathbf{a} \leftarrow \mathbb{Z}_q^n$ *uniformly at random, choosing* $e \leftarrow \chi$, *and outputting* $(\mathbf{a}, b = \mathbf{s}^T\mathbf{a} + e \bmod q)$.

The decisional learning with errors (LWE) problem $\mathsf{LWE}_{n,q,\chi,m}$ *is that given* m *independent samples* $(\mathbf{a}_i, b_i) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$ *where the samples are distributed according to either* $A_{\mathbf{s},\chi}$ *for a uniformly random* $\mathbf{s}$ *or the uniform distribution, distinguish which is the case with non-negligible advantage.*

For parameters, it is often required that $m = \mathrm{poly}(n)$, $q = O(2^{n^\epsilon})$ for some $\epsilon > 0$, and χ is the discrete Gaussian. We say that the distribution χ is β-bounded, if $|\chi| \leq \beta$ with overwhelming probability. We can choose appropriate parameters for χ to be β-bounded given $\beta = \mathrm{poly}(\lambda)$ such that $\mathsf{LWE}_{n,q,\chi,m}$ is hard.

We give a lemma which will be used in our proof:

Lemma 2.6. *For* $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, *let* $\{(\mathbf{a}_i, b_i)\}_{i \in [m]}$ *be sampled from* $A_{\mathbf{s},\chi}$. *Let* $M \subseteq [m]$, *and* $\{(\mathbf{a}'_i, b'_i)\}_{i \in [m]}$ *be defined as: for* $i \in M$, $(\mathbf{a}'_i, b'_i) \leftarrow A_{\mathbf{s},\chi}$, *otherwise* $(\mathbf{a}'_i, b'_i)$ *is uniformly random. Then* $\{(\mathbf{a}_i, b_i)\}_{i \in [m]}$ *and* $\{(\mathbf{a}'_i, b'_i)\}_{i \in [m]}$ *are indistinguishable assuming the hardness of* $\mathsf{LWE}_{n,q,\chi,m}$.

Proof. Let $\{(\mathbf{a}_i^*, b_i^*)\}_{i\in[m]}$ be a set of m uniformly random samples, then it is indistinguishable with $\{(\mathbf{a}_i, b_i)\}_{i\in[m]}$ from the hardness of $\mathsf{LWE}_{n,q,\chi,m}$. For those $i \in M$, we replace only $(\mathbf{a}_i^*, b_i^*)$ by LWE samples from $A_{\mathbf{s},\chi}$ to get $\{(\mathbf{a}_i', b_i')\}_{i\in[m]}$, and the two are also indistinguishable from the hardness of $\mathsf{LWE}_{n,q,\chi,m}$. □

Below we also use another assumption called mheLWE [5], which hardness can be reduced to the standard LWE assumption.

Definition 2.8. *[5] Let q, m, t be integers, σ be a real and τ be a distribution over $\mathbb{Z}^{t\times m}$, all of them functions of a parameter n. The multi-hint extended-LWE problem* $\mathsf{mheLWE}_{n,q,\sigma,m,t,\tau}$ *is to distinguish between the distributions of the tuples: $(\mathbf{A}, \mathbf{A}\cdot\mathbf{s}+\mathbf{e}, \mathbf{Z}, \mathbf{Z}\cdot\mathbf{e})$ and $(\mathbf{A}, \mathbf{u}, \mathbf{Z}, \mathbf{Z}\cdot\mathbf{e})$, where $\mathbf{A} \leftarrow \mathbb{Z}_q^{m\times n}$, $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, $\mathbf{u} \leftarrow \mathbb{Z}_q^m$, $\mathbf{e} \leftarrow D_{\mathbb{Z},\sigma}^m$, and $\mathbf{Z} \leftarrow \tau$.*

Lemma 2.7. *[5] Let $n \geq 100$, $q \geq 2$,$t < n$ and m with $m = \Omega(n\log n)$ and $m \leq n^{O(1)}$. There exists $\xi \leq O(n^4m^2\log 5/2n)$ and a distribution τ over $\mathbb{Z}^{t\times m}$ such that the following statements hold:*

- *There is a reduction from* $\mathsf{LWE}_{n-t,q,\sigma,m}$ *in dimension to* $\mathsf{mheLWE}_{n,q,\sigma\xi,m,t,\tau}$ *that reduces the advantage by at most $2^{\Omega(t-n)}$;*
- *It is possible to sample from τ in time polynomial in n;*
- *Each entry of matrix τ is an independent discrete Gaussian $\tau_{i,j} = D_{\mathbb{Z},\sigma_{i,j},c_{i,j}}$ for some $c_{i,j} \in \{0,1\}$ and $\sigma_{i,j} \geq \Omega(\sqrt{mn\log m})$;*
- *All rows from a sample from τ have norms $\leq \xi$ without a negligible probability.*

3 Noisy Linear Functional Encryption with Bounded Collusion

In this section, we construct an indistinguishability-based secure noisy linear functional encryption scheme with random key queries. Our construction is similar to the inner-product encryption scheme in [5].

Definition 3.1. *An NLinFE scheme consists of the following algorithms:*

- $\mathsf{Setup}(1^\lambda, 1^l)$*: output a pair (PK, MSK).*
- $\mathsf{KeyGen}(MSK, \mathbf{x})$*: for $\mathbf{x} \in \mathbb{Z}_q^l$, output a secret key $sk_\mathbf{x}$.*
- $\mathsf{Enc}(PK, \mathbf{y})$*: for $\mathbf{y} \in \mathbb{Z}_q^l$, output a ciphertext $ct_\mathbf{y}$.*
- $\mathsf{Dec}(ct_\mathbf{y}, sk_\mathbf{x})$*: Output an approximate inner product for $\mathbf{y}, \mathbf{x}$.*

An NLinFE scheme is γ-correct if for any $ct_\mathbf{y} \leftarrow \mathsf{Enc}(PK, \mathbf{y})$, $sk_\mathbf{x} \leftarrow \mathsf{KeyGen}(MSK, \mathbf{x})$, $|\mathsf{Dec}(ct_\mathbf{y}, sk_\mathbf{x}) - \langle\mathbf{y}, \mathbf{x}\rangle| \mod q \leq \gamma$ except for a negligible probability.

Now we give our construction for NLinFE with random keys.

- $\mathsf{Setup}(1^\lambda, 1^l)$: Let τ be a distribution over $\mathbb{Z}^{(l+1)\times m}$ as in the definition of mheLWE. Sample $\mathbf{A} \leftarrow \mathbb{Z}_q^{m\times n}$ and $\mathbf{Z} \leftarrow \tau$ (as defined in Lemma 2.7), compute $\mathbf{U} = \mathbf{Z}\cdot\mathbf{A} \in \mathbb{Z}_q^{(l+1)\times n}$. Let $PK = (\mathbf{A}, \mathbf{U})$ and $MSK = \mathbf{Z}$.

- $\mathsf{KeyGen}(MSK, \mathbf{x})$: Given a vector $\mathbf{x} \in \mathbb{Z}^l$ which is indistinguishable from $D_{\mathbb{Z}^l,\sigma}$, let $\bar{\mathbf{x}} = \binom{\mathbf{x}}{0}$, return the secret key $sk_{\mathbf{x}} = (\mathbf{x}, \mathbf{z_x} := \bar{\mathbf{x}}^T \mathbf{Z} \in \mathbb{Z}^m)$.
- $\mathsf{Enc}(PK, \mathbf{y})$: To encrypt a vector $\mathbf{y} \in \mathbb{Z}_q^l$, let $\bar{\mathbf{y}} = \binom{\mathbf{y}}{\alpha}$, $\alpha \leftarrow \mathbb{Z}_q$. Sample $\mathbf{s} \in \mathbb{Z}_q^n$, $\mathbf{e}_0 \leftarrow D_{\mathbb{Z},\sigma}^m$, $\mathbf{e}_1 \leftarrow D_{\mathbb{Z},\sigma'}^{l+1}$, and compute $ct = (\mathbf{c}_0 = \mathbf{A} \cdot \mathbf{s} + \mathbf{e}_0, \mathbf{c}_1 = \mathbf{U} \cdot \mathbf{s} + \mathbf{e}_1 + \bar{\mathbf{y}})$.
- $\mathsf{Dec}(ct, sk_{\mathbf{x}})$: Let $ct = (\mathbf{c}_0, \mathbf{c}_1)$, compute $\mu = \langle \mathbf{x}, \mathbf{c}_1 \rangle - \langle \mathbf{z_x}, \mathbf{c}_0 \rangle$.

Note that we add restrictions on $\mathbf{x}$ in each KeyGen query, such that $\mathbf{x}$ is indistinguishable from a specific distribution: say, $D_{\mathbb{Z}^l,\sigma}$. This limits the usage of our NLinFE scheme. However, it is enough to construct the required ABE scheme.

Correctness. Since $\langle \mathbf{x}, \mathbf{y} \rangle = \langle \bar{\mathbf{x}}, \bar{\mathbf{y}} \rangle$, we see that $(\langle \bar{\mathbf{x}}, \mathbf{c}_1 \rangle - \langle \mathbf{z_x}, \mathbf{c}_0 \rangle) - \langle \mathbf{x}, \mathbf{y} \rangle = \langle \bar{\mathbf{x}}, \mathbf{e}_1 \rangle - \langle \mathbf{z_x}, \mathbf{e}_0 \rangle$. By Lemma 2.2, we have that $\|\bar{\mathbf{x}}\| \le \sigma\sqrt{l}$, $\|\mathbf{z_x}\| \le \sigma \max\{\sigma_{i,j}\} l\sqrt{m}$, $\|\mathbf{e}_0\| \le \sigma\sqrt{m}$, $\|\mathbf{e}_1\| \le \sigma'\sqrt{l}$. So the scheme is γ-correct for $\gamma \ge \sigma\sigma' l + \sigma^2 \max\{\sigma_{i,j}\} lm$.

We define the fully indistinguishability-based security by the following interactive game:

Definition 3.2. *An NLinFE scheme with random keys is fully β-indistinguishability-based secure, if for any adversary, the advantage of winning the following game is negligible:*

***Setup**. The challenger runs the **Setup** algorithm and gives the adversary PK.*

***Phase 1**. The adversary submits a vector $\mathbf{x}$ for a **KeyGen** query which distribution is indistinguishable from $D_{\mathbb{Z}^l,\sigma}$. The challenger answers with a secret key $\mathsf{sk}_{\mathbf{x}}$ for $\mathbf{x}$. These queries can be repeated adaptively.*

***Challenge**. The adversary chooses two challenge messages $\mathbf{y}_0, \mathbf{y}_1$ and gives it to the challenger. The challenger first checks whether for all queried $\mathbf{x}$, there is $|\langle \mathbf{x}, \mathbf{y}_0 - \mathbf{y}_1 \rangle| \le \beta$. If this does not hold, then the challenger aborts. Otherwise, it chooses a random bit $b \in \{0, 1\}$, and returns $\mathsf{ct}_{\mathbf{y}} = \mathsf{Enc}(PK, \mathbf{y}_b)$ to the adversary.*

***Phase 2**. The adversary repeats Phase 1, under the restriction that each queried $\mathbf{x}$ satisfies that $|\langle \mathbf{x}, \mathbf{y}_0 - \mathbf{y}_1 \rangle| \le \beta$.*

***Guess**. The adversary outputs a guess b' for b, and the winning advantage is defined as $|Pr[b' = b] - 1/2|$.*

Theorem 3.1. *For some properly chosen $\sigma, \tau, \sigma', q, n, l < n$ and $m = \Theta(n \log q)$, the NLinFE scheme above is fully indistinguishability-based secure with k-bounded collusion for $k \le l/3$, assuming the hardness of $\mathsf{mheLWE}_{n,q,\sigma,m,l,\tau}$.*

Proof. The proof of this theorem is similar to the proof in [5]. We refer the readers to Appendix B.

We see that the number of maximal key queries can be bounded by $l/3$ where l is the vector dimension, hence the vector dimension must grow linearly in Q. We note that in order to successfully decrypt, the modulus q should also grow polynomially in the vector dimension, so it leads to a ciphertext size of $O(Q \log Q)$. This is better than [7] where the modulus q grows exponentially in the vector dimension, which leads to $O(Q^2)$ ciphertext size.

4 Fully Secure CP-ABE Scheme for CNF policies

4.1 Construction

Let NLinFE be γ-correct and β-indistinguishability-based secure as in Section 3. We choose a β-bounded error distribution χ, and $\beta/\gamma = O(2^{-\lambda^\epsilon})$ for some $\epsilon > 0$. The CP-ABE scheme is constructed as follows:

- $\mathsf{Setup}(1^\lambda, 1^l)$: Let l be the maximal number of literals. Sample $l+1$ uniformly random matrixes in $\mathbb{Z}_q^{n\times m}$ along with trapdoor: $(\mathbf{A}_1, \mathbf{T}_1), ..., (\mathbf{A}_l, \mathbf{T}_l), (\mathbf{A}, \mathbf{T})$, and a uniformly random vector $\mathbf{u} \in \mathbb{Z}_q^n$. Run $\mathsf{NLinFE.Setup}$ $l+1$ times to generate $(PK_1, MSK_1), ..., (PK_l, MSK_l), (PK, MSK)$. Output $mpk = (\mathbf{A}_1, ..., \mathbf{A}_l, \mathbf{A}, \mathbf{u}, PK_1, ..., PK_l, PK)$, and $msk = (\mathbf{T}_1, ..., \mathbf{T}_l, \mathbf{T}, MSK_1, ..., MSK_l, MSK)$.
- $\mathsf{KeyGen}(msk, L)$: Randomly choose $\mathbf{a} \in \mathbb{Z}_q^n$. Sample $\mathbf{x} \in \mathbb{Z}^m$ such that $\mathbf{A}\mathbf{x} = \mathbf{a}+\mathbf{u}$, and use NLinFE.KeyGen$(MSK, \mathbf{x})$ to generate an NLinFE secret key K. For each literal $i \in L$, sample $\mathbf{x}_i \in \mathbb{Z}^m$ such that $\mathbf{A}_i\mathbf{x}_i = \mathbf{a}$, and use $\mathsf{NLinFE.KeyGen}(MSK_i, \mathbf{x}_i)$ to generate an NLinFE secret key K_i. Return the secret key $K, \{K_i\}_{i\in L}$.
- $\mathsf{Enc}(mpk, f, \mu)$: Let $T_1, ..., T_k$ be clauses in f. Generate uniform $\mathbf{s}_1, ..., \mathbf{s}_k \in \mathbb{Z}_q^n$. For each $j \in T_i$, let $C_{i,j} = \mathsf{NLinFE.Enc}(PK_j, \mathbf{s}_i^T\mathbf{A}_j)$. Let $C = \mathsf{NLinFE.Enc}$ $(PK, (\sum_{i=1}^k \mathbf{s}_i)^T\mathbf{A})$, and $C' = (\sum_{i=1}^k \mathbf{s}_i)^T\mathbf{u} + \mu\lfloor q/2\rfloor + \bar{e}$, $\bar{e} \leftarrow \chi$. Return the ciphertext $(\{C_{i,j}\}_{i\in[k], j\in T_i}, C, C')$.
- $\mathsf{Dec}(ct_f, sk_L)$: First check if L satisfy the policy f. If $f(L) = 1$, then for each $i \in [k]$, there is at least one literal $l_i \in L \cap T_i$, let $d_i = \mathsf{NLinFE.Dec}(K_{l_i}, C_{i,l_i})$. Let $d = \mathsf{NLinFE.Dec}(K, C)$. Calculate $(\sum_{i=1}^k d_i) - d + C'$, if the value is close to 0, return 0; if the value is close to $q/2$, return 1.

Theorem 4.1. *Let $q > 4(l+1)\gamma + 4\beta$, and NLinFE is γ-correct. Then the CP-ABE scheme above is correct.*

Proof. First, by the correctness of NLinFE, for $j \in L \cap T_i$, $d_i = \mathbf{s}_i^T\mathbf{A}_j\mathbf{x}_j + e_i = \mathbf{s}_i^T\mathbf{a} + e_j$, $|e_j| \le \gamma$. Also, $d = (\sum_{i=1}^k \mathbf{s}_i)^T\mathbf{A}\mathbf{x} + e = (\sum_{i=1}^k \mathbf{s}_i)^T(\mathbf{a}+\mathbf{u}) + e$, $|e| \le \gamma$.

So $(\sum_{i=1}^k d_i) - d + C' = \mu\lfloor q/2\rfloor + \sum_{i=1}^k e_i - e + \bar{e}$, which is $(l+1)\gamma + \beta$-close to 0 or $\lfloor q/2\rfloor$. Since $(l+1)\gamma + \beta < q/4$, we can get the correct message. □

Now we give the security result of the scheme above.

Theorem 4.2. *The construction above is fully secure under bounded collusion, assuming the existence of an indistinguishability-based secure NLinFE scheme with bounded collusion and the hardness of LWE problem.*

We combine Theorem 4.2, Theorem 3.1 and Lemma 2.7, and immediately get the following result:

Corollary 4.3 *The construction above is fully secure under bounded collusion, assuming the the hardness of LWE problem.*

4.2 Hyper-functional Keys and Semi-functional Ciphertexts

Now we are ready to prove Theorem 4.2. But before we start the security proof, we first define hyper-functional secret keys and semi-functional ciphertexts.

Hyper-functional key. For a hyper-functional key, we not only change the key generation algorithm, but also the setup algorithm. In Setup, instead of generating $\mathbf{A}$ along with its trapdoor, we generate $\mathbf{A}' \in \mathbb{Z}_q^{(n+1)\times m}$ along with its trapdoor $\mathbf{T}'$. We write the first n rows of $\mathbf{A}'$ as $\mathbf{A}$, and the last row as $\tilde{\mathbf{a}}^T$, which means that $\mathbf{A}' = \binom{\mathbf{A}}{\tilde{\mathbf{a}}^T}$. $\mathbf{A}$ is included in the public key as normal. We also generate $\mathbf{t} \leftarrow \mathbb{Z}_q^n$.

For KeyGen queries, we first sample $e', e \leftarrow \chi$. Let $\mathbf{x} \leftarrow \mathsf{SamplePre}(\mathbf{A}', \mathbf{T}', \binom{\mathbf{a}+\mathbf{u}}{\mathbf{t}^T(\mathbf{a}+\mathbf{u})+e'+\bar{e}+e})$. Then we have $(\tilde{\mathbf{a}}^T - \mathbf{t}^T\mathbf{A})\mathbf{x} = e' + \bar{e} + e \approx 0$. Let $K \leftarrow \mathsf{NLinFE.KeyGen}(MSK, \mathbf{x})$ and other key elements generated the same as normal. We say that the secret key is hyper-functional related to $\tilde{\mathbf{a}}^T - \mathbf{t}^T\mathbf{A}$.

Note that we also say that a the secret key is "normal", if $\mathbf{x} \leftarrow \mathsf{SamplePre}(\mathbf{A}', \mathbf{T}', \binom{\mathbf{a}+\mathbf{u}}{b})$ for $b \leftarrow \mathbb{Z}_q$.

For the indistinguishability between hyper-functional and normal keys, we have the following lemma:

Lemma 4.4 *Let* $(\mathbf{A}_0, \mathbf{T}_0) \leftarrow \mathsf{TrapSamp}(1^n, 1^m, q)$. *For* $i \in [Q]$ *and* $\mathbf{a}^i \leftarrow \mathbb{Z}_q^n$, $\mathbf{x}_0^i = \mathsf{SamplePre}(\mathbf{A}_0, \mathbf{T}_0, \mathbf{a}^i)$. *Let* $(\mathbf{A}', \mathbf{T}') \leftarrow \mathsf{TrapSamp}(1^{n+1}, 1^m, q)$, $\mathbf{x}_1^i = \mathsf{SamplePre}(\mathbf{A}', \mathbf{T}', \binom{\mathbf{a}^i}{a'^i+e^i})$, *where* $\mathbf{A}' = \binom{\mathbf{A}_1}{\tilde{\mathbf{a}}}$, $e^i \leftarrow \chi$, $a'^i \in \mathbb{Z}_q$. *Then* $(\mathbf{A}_0, \{\mathbf{x}_0^i\}_{i\in[Q]})$ *is computationally indistinguishable from* $(\mathbf{A}_1, \{\mathbf{x}_1^i\}_{i\in[Q]})$ *assuming the hardness of LWE.*

Proof. We prove the lemma by showing the following distributions are pairwise indistinguishable (either statistical or computational).

- (1) Let $(\mathbf{A}'', \mathbf{T}'') \leftarrow \mathsf{TrapSamp}(1^{2n+1}, 1^m, q)$, $\mathbf{A}''^T = (\mathbf{A}_2^T | \bar{\mathbf{A}}^T | \bar{\mathbf{a}}^T)$. Let $\mathbf{x}_2^i = \mathsf{SamplePre}(\mathbf{A}'', \mathbf{T}'', (\mathbf{a}^{iT} | \bar{\mathbf{b}}^{iT} | \bar{b}^i)^T)$, where $\bar{\mathbf{b}}^i \leftarrow \mathbb{Z}_q^n$ and $\bar{b}^i \leftarrow \mathbb{Z}_q$. By Lemma 2.5, we have $(\mathbf{A}_0, \{\mathbf{x}_0^i\}_{i\in[Q]})$ is statistically indistinguishable from $(\mathbf{A}_2, \{\mathbf{x}_2^i\}_{i\in[Q]})$.
- (2) We first choose $\tilde{b}^i \leftarrow \mathbb{Z}_q^n$ and write $\bar{b}^i = \tilde{b}^i + a'^i$. This does not change the distribution.
- (3) We first choose $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, let $\bar{b}'^i = \mathbf{s}^T\bar{\mathbf{b}}^i + e^i + a'^i$, and let $\mathbf{x}_2'^i = \mathsf{SamplePre}(\mathbf{A}'', \mathbf{T}'', (\mathbf{a}^{iT} | \bar{\mathbf{b}}^{iT} | \bar{b}'^i)^T)$. By the hardness of LWE problem, any adversary cannot distinguish between $\bar{\mathbf{b}}^i, \tilde{b}^i$ and $\bar{\mathbf{b}}^i, \mathbf{s}^T\bar{\mathbf{b}}^i + e^i$, hence cannot distinguish between $\mathbf{x}_2^i$ and $\mathbf{x}_2'^i$.
- (4) Let $\tilde{\mathbf{a}} = \bar{\mathbf{a}} - \bar{\mathbf{A}}^T\mathbf{s}$, and we have $\tilde{\mathbf{a}}^T\mathbf{x}_2'^i = a'^i + e^i$.
- (5) This time we write $\mathbf{A}''^T = (\mathbf{A}_2^T | \bar{\mathbf{A}}^T | \tilde{\mathbf{a}}^T)$, and set $\mathbf{x}_2''^i = \mathsf{SamplePre}(\mathbf{A}'', \mathbf{T}'', (\mathbf{a}^{iT} | \bar{\mathbf{b}}^{iT} | a'^i + e^i)^T)$. Then $\mathbf{x}_2'^i$ and $\mathbf{x}_2''^i$ are from the same distribution.
- (6) By Lemma 2.5, $(\mathbf{A}_2, \{\mathbf{x}_2''^i\}_{i\in[Q]})$ is statistically indistinguishable from $(\mathbf{A}_1, \{\mathbf{x}_1^i\}_{i\in[Q]})$.

□

Semi-functional ciphertext.

A ciphertext is semi-functional, if the ciphertext element C is $\mathsf{NLinFE.Enc}(PK, (\sum_{i=1}^{k} \mathbf{s}_i - \mathbf{t})^T \mathbf{A} + \tilde{\mathbf{a}})$ instead of $\mathsf{NLinFE.Enc}(PK, (\sum_{i=1}^{k} \mathbf{s}_i)^T \mathbf{A})$.

It follows directly from the indistinguishable security of NLinFE that a semi-functional ciphertext element is indistinguishable from a normal one if all secret keys are hyper-functional.

Along with hyper-functional keys and semi-functional ciphertexts, we also define temporary hyper-functional keys and i-temporary semi-functional ciphertexts, which will be used in our security proof. We note that in our definition, "hyper-functional" and "temporary hyper-functional" form two independent dimensions: a temporary hyper-functional key can be either normal or hyper-functional.

Temporary hyper-functional key. Let l be the number of literals. Like the definition of hyper-functional keys, we not only change the key generation algorithm, but also the setup algorithm. In Setup, instead of generating $\mathbf{A}_j$, $j \in [l]$ along with its trapdoor, we generate $\mathbf{A}'_j \in \mathbb{Z}_q^{(n+1)\times m}$ along with its trapdoor $\mathbf{T}'_j$. We write the first n rows of $\mathbf{A}'_j$ as $\mathbf{A}_j$, and the last row as $\tilde{\mathbf{a}_j}^T$, which means that $\mathbf{A}'_j = \binom{\mathbf{A}_j}{\tilde{\mathbf{a}_j}^T}$. $\mathbf{A}_j$ is included in the public key as normal.

For KeyGen queries, let L be the queried literal set. For $j \in L$, let $\mathbf{x}_j \leftarrow \mathsf{SamplePre}(\mathbf{A}'_j, \mathbf{T}'_j, \binom{\mathbf{a}}{\mathbf{t}^T\mathbf{a}+e'+e_j})$, where $e_j \leftarrow \chi$, and if the key is normal, we sample $e' \leftarrow \chi$, if the key is hyper-functional, we use the same e' as in the generation of $\mathbf{x}$. Then we have $(\tilde{\mathbf{a}_j}^T - \mathbf{t}^T\mathbf{A}_j)\mathbf{x}_j = e' + e_j \approx 0$. Let $K_j \leftarrow \mathsf{NLinFE.KeyGen}(MSK_j, \mathbf{x}_j)$. We say that the secret key is temporary hyper-functional related to $\{\tilde{\mathbf{a}}_j^T - \mathbf{t}^T\mathbf{A}_j\}_{j\in S}$.

We can also use Lemma 4.4 to prove the indistinguishability between normal/hyper-functional keys and temporary hyper-functional keys.

i-Temporary semi-functional ciphertext. A ciphertext is i-temporary semi-functional, if each ciphertext element $C_{i,j}$ is $\mathsf{NLinFE.Enc}(PK_j, (\mathbf{s}_i - \mathbf{t})^T\mathbf{A}_j + \tilde{\mathbf{a}_j})$ instead of $\mathsf{NLinFE.Enc}(PK_j, \mathbf{s}_i^T\mathbf{A}_j)$.

It follows directly from the indistinguishability-based security of NLinFE that a temporary semi-functional ciphertext element is indistinguishable from a semi-functional one if all secret keys but those $sk_L, L \cap T_i = \emptyset$ are temporary hyper-functional.

4.3 Security Proof

We first give the outline of our proof.

- Switch all queried secret keys into hyper-functional keys.
- Switch the challenge ciphertext into semi-functional ciphertext.
- For the p-th query in Phase 1 which challenge literal set is L:
  - Switch all secret keys into temporary hyper-functional secret keys.
  - Switch the ciphertext into i-temporary semi-functional ciphertext such that $L \cap T_i = \emptyset$.
  - Switch the p-th secret key into a normal one using LWE assumption.
  - Switch the ciphertext into a non-temporary semi-functional ciphertext.

  - Switch all secret keys into non-temporary normal or hyper-functional secret keys.
- For queries in Phase 2, and $i \in [k]$, k is the maximal number of clauses:
  - Switch all secret keys into temporary hyper-functional secret keys.
  - Switch the ciphertext into i-temporary semi-functional ciphertext.
  - Switch all Phase 2 secret keys such that $L \cap T_i = \emptyset$ into a normal one using LWE assumption.
  - Switch the ciphertext into a non-temporary semi-functional ciphertext.
  - Switch all secret keys into non-temporary normal or hyper-functional secret keys.
- Now C is uniformly random, independent with any queried secret keys. We further switch C' into a uniformly random element, and thus have our result.

Now we define the game sequence.

Game 0 is the original game.

Game 1: Each queried secret key is a hyper-functional key. **Game** 0 and **Game** 1 are indistinguishable by Lemma 4.4.

Game 2: The challenge ciphertext is semi-functional. We first define **Game** 1a and **Game** 2a as follows:

- The Setup phase and Phase 1 are the same as Game 1 or Game 2.
- Let Q_2 be the maximal number of Phase 2 queries. In the Challenge phase, before the challenge ciphertext is given, for each $r \in [Q_2]$, we generate $\mathbf{a}^r \leftarrow \mathbb{Z}_q^n$, $e'^r, e^r \leftarrow \chi$, and $\mathbf{x}^r \leftarrow \mathsf{SamplePre}(\mathbf{A}', \mathbf{T}', \binom{\mathbf{a}^r+\mathbf{u}}{\mathbf{t}^T(\mathbf{a}^r+\mathbf{u})+e'^r+\bar{e}+e^r})$. Let $K^r = \mathsf{NLinFE.KeyGen}(MSK, \mathbf{x}^r)$.
- In Game 1a, the challenger returns a normal ciphertext, and in Game 2a, it returns a semi-functional one.
- In the r-th Phase 2 query, we let $e' = e'^r$, $e = e^r$, $\mathbf{a} = \mathbf{a}^r$, and the key element $K = K^r$. Other key elements are generated as before.

It is easy to see that **Game** 1 and **Game** 1a; **Game** 2 and **Game** 2a are the same from the adversary's point of view. We now show that **Game** 1a and **Game** 2a are indistinguishable.

For the challenger, instead of generating all Ks and C itself, it now runs a indistinguishable game for NLinFE, get K by the KeyGen query of NLinFE, and get C as the challenge ciphertext of NLinFE. Because $|(\tilde{\mathbf{a}}^T - \mathbf{t}^T\mathbf{A})\mathbf{x}| \leq 3\beta$ and $\beta/\gamma = O(2^{-\lambda^\epsilon})$, we have the indistinguishable result by Theorem 3.1.

Game $2(p)$, $p \in [Q_1 + 1]$, Q_1 is the number of phase 1 queries: The first $p - 1$ Phase 1 keys are normal, and the rest of the keys are hyper-functional; the challenge ciphertext is semi-functional. Then **Game** 2(1) is **Game** 2, and in **Game** $2(Q_1 + 1)$, all Phase 1 keys are normal. We prove the following result:

Lemma 4.5 ***Game*** *2(p) and* ***Game*** *2(p + 1) are indistinguishable assuming the security of NLinFE and the hardness of LWE.*

Proof. We prove this by the following game sequence:

Game 2-1(p): **Game** 2-1(p) is same as **Game** 2(p), except that we change all keys into temporary hyper-functional keys. **Game** 2-1(p) is indistinguishable from **Game** 2(p) according to Lemma 4.4.

Let L be the challenge literal set in the p-th query of Phase 1. So there must be clause T_i such that $L \cap T_i = \emptyset$. This i will be used in the following games.

Game 2-2(p, j): **Game** 2-2(p, j) is same as **Game** 2-1(p), except that for any $C_{i,j'}$ such that $j' \leq j$ and $j' \in T_i$, $C_{i,j'}$ is generated as $\mathsf{NLinFE.Enc}(PK_{j'}, (\mathbf{s}_i - \mathbf{t})^T \mathbf{A}_{j'} + \tilde{\mathbf{a}}_{j'}^T)$. So **Game** 2-2($p, 0$) is **Game** 2-1($p$), and in **Game** 2-2($p, l$), the ciphertext is i-temporary semi-functional. We now show that **Game** 2-2($p, j-1$) is indistinguishable from Game 2-2(p, j).

We define **Game** 2-2a(p, j) and **Game** 2-2b(p, j) as follows:

Game 2-2a(p, j): The game is the same as **Game** 2-2(p, j), except that:

- In the Challenge phase, before the challenge ciphertext is given, we first check whether $j + 1 \in T_i$. If $j + 1 \notin T_i$, the game proceeds as **Game** 2-2(p, j). If $j+1 \in T_i$, for each $r \in [Q_2]$, we generate $\mathbf{a}^r \leftarrow \mathbb{Z}_q^n$, $e'^r, e_{j+1}^r \leftarrow \chi$, and $\mathbf{x}_{j+1}^r \leftarrow \mathsf{SamplePre}(\mathbf{A}'_{j+1}, \mathbf{T}'_{j+1}, \binom{\mathbf{a}^r}{\mathbf{t}^T\mathbf{a}^r + e'^r + e_{j+1}^r})$. Let $K_{j+1}^r = \mathsf{NLinFE.KeyGen}(MSK_{j+1}, \mathbf{x}_{j+1}^r)$.
- In the r-th Phase 2 query, if $j + 1 \in T_i$, we let $e' = e'^r$, $e_{j+1} = e_{j+1}^r$, $\mathbf{a} = \mathbf{a}^r$, and the key element $K_{j+1} = K_{j+1}^r$. Then, generate other key elements as in **Game** 2-2(p, j).

Game 2-2b(p, j): The game is the same as **Game** 2-2(p, j), except that:

- In the Challenge phase, before the challenge ciphertext is given, we first check whether $j \in T_i$. If $j \notin T_i$, the game proceeds as **Game** 2-2(p, j). If $j \in T_i$, for each $r \in [Q_2]$, we generate $\mathbf{a}^r \leftarrow \mathbb{Z}_q^n$, $e'^r, e_j^r \leftarrow \chi$, and $\mathbf{x}_j^r \leftarrow \mathsf{SamplePre}(\mathbf{A}'_j, \mathbf{T}'_j, \binom{\mathbf{a}^r}{\mathbf{t}^T\mathbf{a}^r + e'^r + e_j^r})$. Let $K_j^r = \mathsf{NLinFE.KeyGen}(MSK_j, \mathbf{x}_j^r)$.
- In the r-th Phase 2 query, if $j \in T_i$, we let $e' = e'^r$, $e_j = e_j^r$, $\mathbf{a} = \mathbf{a}^r$, and the key element $K_j = K_j^r$. Then, generate other key elements as in **Game** 2-2(p, j).

It is easy to see that **Game** 2-2(p, j), **Game** 2-2a(p, j) and **Game** 2-2b(p, j) are the same from the adversary's point of view. We now show that **Game** 2-2a($p, j-1$) and **Game** 2-2b(p, j) are indistinguishable.

For the challenger, instead of generating all K_js and $C_{i,j}$ itself, it now runs a indistinguishable game for NLinFE, get K_j by the KeyGen query of NLinFE, and get $C_{i,j}$ as the challenge ciphertext of NLinFE. Since $|(\tilde{\mathbf{a}}_j^T - \mathbf{t}^T\mathbf{A}_j)\mathbf{x}_j| \leq 2\beta$ and $\beta/\gamma = O(2^{-\lambda^\epsilon})$ by assumption, we only need to show that the NLinFE game can proceed correctly. If $j \in T_i$, all KeyGen queries are made before the challenge ciphertext, which is legal in the NLinFE game. If $j \notin T_i$, the NLinFE challenge ciphertext is never required, so all KeyGen queries can be made correctly. Thus we have the indistinguishable result by Theorem 3.1.

Thus we have that **Game** 2-1(p) is indistinguishable from **Game** 2-2(p, l).

Game 2-3(p): The game is the same as **Game** 2-2(p, l), except that:

- In the challenge phase, we generate $\bar{\mathbf{s}} \leftarrow \mathbb{Z}_q^n$, $\{\mathbf{s}_{i'}\}_{i' \neq i} \leftarrow \mathbb{Z}_q^n$, and write $C_{i,j}$ for any $j \in T_i$ as $\mathsf{NLinFE.Enc}(PK_j, (\bar{\mathbf{s}} - \sum_{i' \neq i} \mathbf{s}_{i'})^T \mathbf{A}_j + \tilde{\mathbf{a}}_j)$.
- We also write $C = \mathsf{NLinFE.Enc}(PK, \bar{\mathbf{s}}^T \mathbf{A} + \tilde{\mathbf{a}}^T)$, and $C' = (\bar{\mathbf{s}} + \mathbf{t})^T \mathbf{u} + \mu \lfloor q/2 \rfloor + \bar{e}$.

Note that in **Game** 2-3(p), we implicitly set $\mathbf{s}_i = \bar{\mathbf{s}} + \mathbf{t} - \sum_{i' \neq i} \mathbf{s}_{i'}$, so that for the adversary, **Game** 2-3(p) is the same as **Game** 2-2(p, l). Now we see that $\mathbf{t}$ only occurs in C' and in KeyGen queries. All these occurrences of $\mathbf{t}$ take the form of LWE samples: $\mathbf{t}^T \mathbf{a} + e'$, and $\mathbf{t}^T \mathbf{u} + \bar{e}$.

Game 2-4(p): For the p-th query, we choose a uniformly random $\tilde{b} \leftarrow \mathbb{Z}_q$, and let $\mathbf{x} \leftarrow \mathsf{SamplePre}(\mathbf{A}', \mathbf{T}', \binom{\mathbf{a}}{\tilde{b} + \mathbf{t}^T \mathbf{u} + \bar{e} + e})$. For $i \in [l]$, $\mathbf{x}_i \leftarrow \mathsf{SamplePre}(\mathbf{A}'_i, \mathbf{T}'_i, \binom{\mathbf{a}}{\tilde{b} + e_i})$.

Game 2-3(p) and **Game** 2-4(p) are indistinguishable using Lemma 2.6, by the hardness of LWE problem. We also define **Game** 2-4a(p), which removes $\bar{\mathbf{s}}$, and $\mathbf{s}_i$ is uniformly sampled in the challenge phase. **Game** 2-4a(p) is the same as **Game** 2-4(p) from the adversary's point of view.

Game 2-5(p, j): **Game** 2-5(p, j) is same as **Game** 2-4a(p), except that for any $C_{i,j'}$ such that $j' \leq j$ and $j' \in T_i$, $C_{i,j'}$ is generated as $\mathsf{NLinFE.Enc}(PK_{j'}, \mathbf{s}_i^T \mathbf{A}_{j'}$. So **Game** 2-5($p, 0$) is **Game** 2-4a($p$), and in **Game** 2-5($p, l$), the ciphertext is (non-temporary) semi-functional.

The indistinguishability between **Game** 2-5($p, j-1$) and **Game** 2-5(p, j) is nearly the same as **Game** 2-2($p, j-1$) and **Game** 2-2(p, j), except that this time, for the p-th query with literal set L, $|(\tilde{\mathbf{a}_j}^T - \mathbf{t}^T \mathbf{A}_j)\mathbf{x}_j|$ may not be small. However, since $L \cap T_i = \emptyset$, for each $j \in T_i$ where it is required to generate the ciphertext element $C_{i,j}$, the corresponding key element $\mathbf{x}_j$ does not occur. So the ciphertext can be generated correctly in the reduction. Now we have that **Game** 2-4a(p) is indistinguishable from **Game** 2-5(p, l).

Game 2-6(p): **Game** 2-6(p) is same as **Game** 2-5(p, l) except that in the p-th KeyGen query, instead of generating random $\tilde{b}$, we sample $b \leftarrow \mathbb{Z}_q$, and set $\tilde{b} = b - \mathbf{t}^T \mathbf{u} - \bar{e} - e$. **Game** 2-6($p$) is same as **Game** 2-5(p, l) from the adversary's point of view. We can see that **Game** 2-6(p) is indistinguishable from **Game** 2($p+1$) from Lemma 4.4. □

Game 3(i), $i \in [k+1]$, k is the number of clauses in the challenge access policy: The Phase 1 keys are normal, and for the Phase 2 keys which challenge literal set is L, the key is normal iff there exists $i' < i$ such that $L \cap T_{i'} = \emptyset$. Game 3(1) is the same as Game 2($Q_1 + 1$). Since L must not satisfy the access policy, it is easy to see that in **Game** 3($k+1$), all keys are normal.

Lemma 4.6 ***Game*** *3(i) and* ***Game*** *3(i+1) are indistinguishable assuming the security of NLinFE and the hardness of LWE.*

Proof. The proof is essentially the same as Lemma 4.5. We omit the details here. □

Game 4: **Game** 4 is same as **Game** 3($k+1$), except that all secret keys are temporary hyper-functional keys. **Game** 4 is indistinguishable from **Game** 3($k+1$) by Lemma 4.4.

Game 5: **Game** 5 is same as **Game** 4, except that the challenge ciphertext is 1-temporary semi-functional. Using similar discussion from **Game** 2-2(p, j) in Lemma 4.5, we have that **Game** 4 and **Game** 5 are indistinguishable by Theorem 3.1.

Game 6: The game is the same as **Game** 5, except that:

- In the challenge phase, we generate $\bar{\mathbf{s}} \leftarrow \mathbb{Z}_q^n$, $\{\mathbf{s}_{i'}\}_{i' \neq 1} \leftarrow \mathbb{Z}_q^n$, and write $C_{1,j}$ for any $j \in T_1$ as $\mathsf{NLinFE.Enc}(PK_j, (\bar{\mathbf{s}} - \sum_{i' \neq 1} \mathbf{s}_{i'})^T \mathbf{A}_j + \tilde{\mathbf{a}}_j)$.
- We also write $C = \mathsf{NLinFE.Enc}(PK, \bar{\mathbf{s}}^T \mathbf{A} + \tilde{\mathbf{a}}^T)$, and $C' = (\bar{\mathbf{s}} + \mathbf{t})^T \mathbf{u} + \mu \lfloor q/2 \rfloor + \bar{e}$.

Game 6 is the same as **Game** 5 from the adversary's point of view. Note that this time, $\mathbf{t}^T \mathbf{u}$ only occurs in C'.

Game 7: The game is the same as **Game** 6, except that in the challenge phase, C' is generated by $\bar{\mathbf{s}}^T \mathbf{u} + v + \mu \lfloor q/2 \rfloor$, $v \leftarrow \mathbb{Z}_q$. **Game** 7 is indistinguishable from **Game** 6 by Lemma 2.6 from LWE assumption.

Game 8: The game is the same as **Game** 7, except that in the challenge phase, we let $v' \leftarrow \mathbb{Z}_q$, and $v = v' - \bar{\mathbf{s}}^T \mathbf{u} - \mu \lfloor q/2 \rfloor$, so $C' = v'$. **Game** 7 and **Game** 8 are the same from the adversary's point of view. Then in **Game** 8, the ciphertext contains no information on μ, so the advantage for any adversary is $1/2$. Thus we finish our proof.

5 Conclusion and Future Works

In this paper, we give a construction for lattice-based fully secure ABE schemes from noisy linear functional encryption, which can be considered as a lattice version of the widely used dual-system method from pairing-based cryptography. Our scheme supports CNF formula as its access policy, and any predetermined number of key queries. Compared with other methods for constructing bounded collusion fully secure ABE, our scheme is simpler and has smaller ciphertext size. Since dual-system encryption has shown to be useful in pairing-based cryptography, we hope that we can also extend our scheme for richer functionalities.

Although our scheme supports only bounded collusion, it is easy to see that it can be extended into unbounded case if there exists an unbounded NLinFE scheme. Although in [1], a secret key version of unbounded NLinFE has been introduced, it is currently unknown how to transform it into a public key scheme. This shall be our future work.

Acknowledgements. This work is partially supported by the National Natural Science Foundation of China (No. 62072305, No. 61672339), the National Cryptography Development Fund (No. MMJJ20170111), and the Foundation of Science and Technology on Information Assurance Laboratory (No. KJ-17-109).

References

1. Shweta Agrawal. Indistinguishability obfuscation without multilinear maps: New methods for bootstrapping and instantiation. In Yuval Ishai and Vincent Rijmen,

editors, *Advances in Cryptology - EUROCRYPT 2019 - 38th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Darmstadt, Germany, May 19-23, 2019, Proceedings, Part I*, volume 11476 of *Lecture Notes in Computer Science*, pages 191–225. Springer, 2019.

2. Shweta Agrawal, Dan Boneh, and Xavier Boyen. Lattice basis delegation in fixed dimension and shorter-ciphertext hierarchical IBE. In *Advances in Cryptology - CRYPTO 2010*, pages 98–115, 2010.
3. Shweta Agrawal, Xavier Boyen, Vinod Vaikuntanathan, Panagiotis Voulgaris, and Hoeteck Wee. Functional encryption for threshold functions (or fuzzy ibe) from lattices. *International Workshop on Public Key Cryptography*.
4. Shweta Agrawal, David Mandell Freeman, and Vinod Vaikuntanathan. Functional encryption for inner product predicates from learning with errors. *International Conference on the Theory and Application of Cryptology and Information Security*, 2011:21–40, 2011.
5. Shweta Agrawal, Benoît Libert, and Damien Stehlé. Fully secure functional encryption for inner products, from standard assumptions. In *Annual International Cryptology Conference*, pages 333–362. Springer, 2016.
6. Shweta Agrawal and Alice Pellet-Mary. Indistinguishability obfuscation without maps: Attacks and fixes for noisy linear FE. In Anne Canteaut and Yuval Ishai, editors, *Advances in Cryptology - EUROCRYPT 2020 - 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, May 10-14, 2020, Proceedings, Part I*, volume 12105 of *Lecture Notes in Computer Science*, pages 110–140. Springer, 2020.
7. Shweta Agrawal and Alon Rosen. Functional encryption for bounded collusions, revisited. In *Theory of Cryptography - 15th International Conference, TCC 2017, Baltimore, MD, USA, November 12-15, 2017, Proceedings, Part I*, pages 173–205, 2017.
8. Prabhanjan Ananth and Vinod Vaikuntanathan. Optimal bounded-collusion secure functional encryption. In Dennis Hofheinz and Alon Rosen, editors, *Theory of Cryptography - 17th International Conference, TCC 2019, Nuremberg, Germany, December 1-5, 2019, Proceedings, Part I*, volume 11891 of *Lecture Notes in Computer Science*, pages 174–198. Springer, 2019.
9. Nuttapong Attrapadung. Dual system encryption via doubly selective security: framework, fully secure functional encryption for regular languages, and more. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 557–577. Springer, 2014.
10. John Bethencourt, Amit Sahai, and Brent Waters. Ciphertext-policy attribute-based encryption. *IEEE Symposium on Security and Privacy*, pages 321–334, 2007.
11. Dan Boneh and Xavier Boyen. Efficient selective-id secure identity-based encryption without random oracles. *Annual International Conference on the Theory and Application of Cryptographic Techniques*, 2004:223–238, 2004.
12. 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. *Annual International Conference on the Theory and Application of Cryptographic Techniques*, 2014:533–556, 2014.
13. Dan Boneh, Amit Sahai, and Brent Waters. Functional encryption: Definitions and challenges. In *Theory of Cryptography Conference*, pages 253–273. Springer, 2011.
14. Xavier Boyen. Attribute-based functional encryption on lattices. *Theory of Cryptography Conference*, 2012:122–142, 2013.

15. Xavier Boyen and Qinyi Li. Towards tightly secure lattice short signature and id-based encryption. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 404–434. Springer, 2016.
16. Zvika Brakerski and Vinod Vaikuntanathan. Circuit-abe from lwe: unbounded attributes and semi-adaptive security. In *Annual International Cryptology Conference*, pages 363–384. Springer, 2016.
17. David Cash, Dennis Hofheinz, Eike Kiltz, and Chris Peikert. Bonsai trees, or how to delegate a lattice basis. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 523–552. Springer, 2010.
18. Jie Chen, Junqing Gong, Lucas Kowalczyk, and Hoeteck Wee. Unbounded abe via bilinear entropy expansion, revisited. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 503–534. Springer, 2018.
19. Ling Cheung and Calvin Newport. Provably secure ciphertext policy abe. In *Proceedings of the 14th ACM conference on Computer and communications security*, pages 456–465, 2007.
20. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. *Symposium on the Theory of Computing*, 2007:197–206, 2008.
21. Shafi Goldwasser, Yael Tauman Kalai, Raluca A. Popa, Vinod Vaikuntanathan, and Nickolai Zeldovich. Reusable garbled circuits and succinct functional encryption. In *Symposium on Theory of Computing Conference, STOC'13*, pages 555–564, 2013.
22. Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Functional encryption with bounded collusions via multi-party computation. In Reihaneh Safavi-Naini and Ran Canetti, editors, *Advances in Cryptology - CRYPTO 2012 - 32nd Annual Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2012. Proceedings*, volume 7417 of *Lecture Notes in Computer Science*, pages 162–179. Springer, 2012.
23. Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Attribute-based encryption for circuits. *Symposium on the Theory of Computing*, 2013:545–554, 2013.
24. Rishab Goyal, Venkata Koppula, and Brent Waters. Semi-adaptive security and bundling functionalities made generic and easy. In *Theory of Cryptography Conference*, pages 361–388. Springer, 2016.
25. Vipul Goyal, Omkant Pandey, Amit Sahai, and Brent Waters. Attribute-based encryption for fine-grained access control of encrypted data. *Proceedings of the 13th ACM conference on Computer and Communications Security*, 2006:89–98, 2006.
26. Shuichi Katsumata, Ryo Nishimaki, Shota Yamada, and Takashi Yamakawa. Adaptively secure inner product encryption from LWE. In Shiho Moriai and Huaxiong Wang, editors, *Advances in Cryptology - ASIACRYPT 2020 - 26th International Conference on the Theory and Application of Cryptology and Information Security, Daejeon, South Korea, December 7-11, 2020, Proceedings, Part III*, volume 12493 of *Lecture Notes in Computer Science*, pages 375–404. Springer, 2020.
27. Lucas Kowalczyk and Allison Bishop Lewko. Bilinear entropy expansion from the decisional linear assumption. In *Annual Cryptology Conference*, pages 524–541. Springer, 2015.
28. Lucas Kowalczyk and Hoeteck Wee. Compact adaptively secure abe for NC^1 from k-lin. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 3–33. Springer, 2019.

29. Allison Lewko, Tatsuaki Okamoto, Amit Sahai, Katsuyuki Takashima, and Brent Waters. Fully secure functional encryption: Attribute-based encryption and (hierarchical) inner product encryption. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 62–91. Springer, 2010.
30. Allison B. Lewko and Brent Waters. New proof methods for attribute-based encryption: Achieving full security through selective techniques. In *Advances in Cryptology - CRYPTO 2012*, pages 180–198.
31. Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. *Annual International Conference on the Theory and Application of Cryptographic Techniques*, 2011:700–718, 2012.
32. Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on gaussian measures. *SIAM Journal on Computing*, 37(1):267–302, 2007.
33. Rafail Ostrovsky, Amit Sahai, and Brent Waters. Attribute-based encryption with non-monotonic access structures. In *Proceedings of the 14th ACM conference on Computer and communications security*, pages 195–203, 2007.
34. Oded Regev. New lattice-based cryptographic constructions. *Journal of the ACM (JACM)*, 51(6):899–942, 2004.
35. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *Journal of the ACM*, 56(6):34, 2009.
36. Yannis Rouselakis and Brent Waters. Practical constructions and new proof methods for large universe attribute-based encryption. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*, pages 463–474, 2013.
37. Amit Sahai and Brent Waters. Fuzzy identity-based encryption. *Annual International Conference on the Theory and Application of Cryptographic Techniques*, 2004:457–473, 2005.
38. Rotem Tsabary. Fully secure attribute-based encryption for t-cnf from LWE. In *Advances in Cryptology - CRYPTO 2019*, pages 62–85, 2019.
39. Zhedong Wang, Xiong Fan, and Feng-Hao Liu. FE for inner products and its application to decentralized ABE. In *Public-Key Cryptography - PKC 2019*, pages 97–127, 2019.
40. Brent Waters. Dual system encryption: Realizing fully secure IBE and HIBE under simple assumptions. In *Advances in Cryptology - CRYPTO 2009*, pages 619–636, 2009.
41. Brent Waters. Ciphertext-policy attribute-based encryption: An expressive, efficient, and provably secure realization. In *International Workshop on Public Key Cryptography*, pages 53–70. Springer, 2011.
42. Hoeteck Wee. Dual system encryption via predicate encodings. In *Theory of Cryptography Conference*, pages 616–637. Springer, 2014.
43. Shota Yamada. Adaptively secure identity-based encryption from lattices with asymptotically shorter public parameters. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 32–62. Springer, 2016.

A Proof of Lemma 2.5

We first give the following lemma which is proven in [32, 20].

Lemma A.1. *[20]*

For any $\epsilon \in (0,1)$, there exists $\eta > 0$, such that for $s \geq \eta$, $\rho_s(\Lambda^{\perp}_{\mathbf{u}}(\mathbf{A})) \in [\frac{1-\epsilon}{1+\epsilon}, 1] \cdot \rho_s(\Lambda^{\perp}_{\mathbf{0}}(\mathbf{A}))$.

By Lemma 2.1, we have that the distribution of $\mathbf{x}$ is statistically close to $D_{\Lambda^{\perp}_{\mathbf{u}}(\mathbf{A}),s}$. So we only need to show that the distribution of $\mathbf{x}'$ is statistically close to $D_{\Lambda^{\perp}_{\mathbf{u}}(\mathbf{A}'),s}$.

It is easy to see that $\{\Lambda^{\perp}_{(\mathbf{u}^T|\mathbf{b}^T)^T}(\mathbf{A}')\}_{\mathbf{b}\in\mathbb{Z}_q^{n'-n}}$ forms a partition of the lattice co-set $\Lambda^{\perp}_{\mathbf{u}}(\bar{\mathbf{A}})$. So by the definition of discrete Gaussian, we have that, for any $\mathbf{c}\in\Lambda^{\perp}_{\mathbf{u}}(\mathbf{A}')$, let $\mathbf{b}=\tilde{\mathbf{A}}\mathbf{c}$, we have $Pr(\mathbf{x}=\mathbf{c})=q^{-(n'-n)}\rho_s(\mathbf{c})/\rho_s(\Lambda^{\perp}_{(\mathbf{u}^T|\mathbf{b}^T)^T}(\mathbf{A}'))$. For a negligible ϵ, we choose s satisfies Lemma A.1. Then we have that for any $\mathbf{b}'$, $\rho_s(\Lambda^{\perp}_{(\mathbf{u}^T|\mathbf{b}^T)^T}(\mathbf{A}'))/\rho_s(\Lambda^{\perp}_{(\mathbf{u}^T|\mathbf{b}'^T)^T}(\mathbf{A}'))\in[\frac{1-\epsilon}{1+\epsilon},\frac{1+\epsilon}{1-\epsilon}]$.

By definition, we have:

$$D_{\Lambda^{\perp}_{\mathbf{u}}(\mathbf{A}'),s}(\mathbf{c})=\frac{\rho_s(\mathbf{c})}{\rho_s(\Lambda^{\perp}_{\mathbf{u}}(\bar{\mathbf{A}}))}=\frac{\rho_s(\mathbf{c})}{\sum_{\mathbf{b}'^T\in\mathbb{Z}_q^{n'-n}}\rho_s(\Lambda^{\perp}_{(\mathbf{u}^T|\mathbf{b}'^T)^T}(\mathbf{A}'))}.$$

So:

$$\frac{1-\epsilon}{1+\epsilon}\cdot\frac{\rho_s(\mathbf{c})}{q^{n'-n}\rho_s(\Lambda^{\perp}_{(\mathbf{u}^T|\mathbf{b}^T)^T}(\mathbf{A}'))}\le D_{\Lambda^{\perp}_{\mathbf{u}}(\mathbf{A}'),s}(\mathbf{c})\le\frac{1+\epsilon}{1-\epsilon}\cdot\frac{\rho_s(\mathbf{c})}{q^{n'-n}\rho_s(\Lambda^{\perp}_{(\mathbf{u}^T|\mathbf{b}^T)^T}(\mathbf{A}'))}.$$

Now we have that the statistical distance between the two distributions is no more than 2ϵ, thus we have our result.

B Proof of Theorem 3.1

We prove this by a sequence of interactive games. Let Game 0 be the full security game defined above.

Game 1: Instead of $\mathbf{c}_1=\mathbf{U}\cdot\mathbf{s}+\mathbf{e}_1+\mathbf{y}_\beta$, we compute $\mathbf{c}_1=\mathbf{Z}\cdot\mathbf{c}_0-\mathbf{Z}\cdot\mathbf{e}_0+\mathbf{e}_1+\mathbf{y}_\beta$. Game 1 is the same as Game 0.

Game 2: In Game 2, $\mathbf{c}_0$ is chosen uniform randomly from $\mathbb{Z}_q^m$ instead of $\mathbf{As}+\mathbf{e}_0$. Game 2 is indistinguishable from Game 1 by the hardness of mheLWE.

Now, we remain to prove that in Game 2, the distinguishing advantage for any adversary is negligible. Let $\mathbf{x}_1,...,\mathbf{x}_\kappa$ be the largest set of independent vectors in the key query, and we write $\mathbf{X}=(\mathbf{x}_1|...|\mathbf{x}_\kappa)$, and $\kappa\le k$. We write the ciphertext $ct_\beta=(\mathbf{c}_0,\mathbf{c}_1^\beta)$. By the construction of our scheme, we only need to show that any adversary cannot distinguish between $(\mathbf{A},\mathbf{ZA},\mathbf{X},\mathbf{XZ},\mathbf{c}_0,\mathbf{c}_1^0)$ and $(\mathbf{A},\mathbf{ZA},\mathbf{X},\mathbf{XZ},\mathbf{c}_0,\mathbf{c}_1^1)$ with non-negligible probability.

Let $\mathbf{y}=\mathbf{c}_1^0-\mathbf{c}_1^1=\binom{\mathbf{y}_0-\mathbf{y}_1}{\alpha_0-\alpha_1}$ for $\alpha_0,\alpha_1\leftarrow\mathbb{Z}_q$. Since the last row of $\mathbf{X}$ is 0, so $\mathbf{y}$ is linearly independent with $\mathbf{X}$ except for a negligible probability. We find a short solution $\mathbf{t}$, such that $\mathbf{X}^T\mathbf{t}=0$, $\mathbf{y}^T\mathbf{t}\neq 0$, the coefficients of $\mathbf{t}$ is co-prime, and $\|\mathbf{t}\|=O(\mathrm{poly}(n))$. The solution exists by Siegel's Lemma. We append vectors orthogonal to $\mathbf{t},\mathbf{y}$ and linear independent with $\mathbf{X}$ to form a invertible $n\times n$ matrix (modulus q), written as $\bar{\mathbf{X}}=(\mathbf{X}|\mathbf{y}|\mathbf{X}')$.

Given the invertible matrix $\bar{\mathbf{X}}$, we have that $(\mathbf{A},\mathbf{ZA},\mathbf{X},\mathbf{XZ},\mathbf{c}_0,\mathbf{c}_1^0)$ and $(\mathbf{A},\mathbf{ZA},\mathbf{X},\mathbf{XZ},\mathbf{c}_0,\mathbf{c}_1^1)$ are indistinguishable if and only if $(\mathbf{A},\mathbf{ZA},\mathbf{X},\mathbf{XZ},\mathbf{c}_0,\bar{\mathbf{X}}^T\mathbf{c}_1^0)$ and $(\mathbf{A},\mathbf{ZA},\mathbf{X},\mathbf{XZ},\mathbf{c}_0,\bar{\mathbf{X}}^T\mathbf{c}_1^1)$ are indistinguishable.

We then write $\bar{\mathbf{X}}^T\mathbf{c}_1^\beta$ as $(\mathbf{X}^T\mathbf{c}_1^\beta, \mathbf{X}'^T\mathbf{c}_1^\beta, \mathbf{y}^T\mathbf{c}_1^\beta)$. By the choice of $\mathbf{X}'$, we have that $\mathbf{X}'^T\mathbf{c}_1^0 = \mathbf{X}'^T\mathbf{c}_1^1$.

By the definition of β-indistinguishability-based security, we have that $|\langle \mathbf{x}_i, \mathbf{y}_0\rangle - \langle \mathbf{x}_i, \mathbf{y}_0\rangle| \leq \beta$. So we have that $\mathbf{X}^T\mathbf{c}_1^0 = \mathbf{X}^T\mathbf{Z}(\mathbf{c}_0 - \mathbf{e}_0) + \mathbf{X}^T\mathbf{e}_1 + \mathbf{X}^T\mathbf{y}_0 = \mathbf{X}^T\mathbf{Z}(\mathbf{c}_0 - \mathbf{e}_0) + \mathbf{X}^T\mathbf{e}_1 + \mathbf{X}^T\mathbf{y}_1 + \mathbf{b}$ where $\|\mathbf{b}\|_\infty \leq \beta$. By the lemma below, we show that $\mathbf{X}^T\mathbf{e}_1$ is indistinguishable from $\mathbf{X}^T\mathbf{e}_1 + \mathbf{b}$.

Lemma B.1 *Given $\mathbf{A} \in \mathbb{Z}^{n\times m}$, where each row of $\mathbf{A}$ is independently sampled from $D_{\mathbb{Z}^m,\sigma}$, $\sigma = O(\mathrm{poly}(n))$, $m \geq 3n$, $\mathbf{b} \in \mathbb{Z}^n$, and $\|\mathbf{b}\|_\infty \leq \beta = O(\mathrm{poly}(n))$. Then there exists $\mathbf{x} \in \mathbb{Z}^m$ and $\|\mathbf{x}\|_\infty \leq \delta = O(\mathrm{poly}(n))$ such that $\mathbf{A}\mathbf{x} = \mathbf{b}$ except for a negligible probability.*

Proof. This proof is using standard methods in linear algebra and number theory, we only give a proof sketch due to the page limits.

The proof consists of the several steps:

- For $\mathbf{A} \in \mathbb{Z}^{n\times m}$, show that $\mathbf{A}\mathbf{x} = \mathbf{b}$ has an integer solution iff the determinants of all $n \times n$ sub-matrixes of $\mathbf{A}$ are co-prime. This is proven by constructing the elementary row/column transformations that transform $\mathbf{A}$ into $\mathbf{I}|\mathbf{0}$.
- Show that for $\mathbf{A}$ sampled as defined and each prime $p < q$, the probability that the determinants of all $n \times n$ sub-matrixes of $\mathbf{A}$ are a multiple of p is negligible, hence the probability of $\mathbf{A}\mathbf{x} = \mathbf{b}$ has no integer solution is negligible. This is proven by induction on n: as long as there is at least one $(k-1)\times(k-1)$ sub-matrix of $\mathbf{A}$ which determinant is not a multiple of p, there is at least one $k \times k$ sub-matrix which determinant is not a multiple of p except for a negligible probability.
- We write $\mathbf{A}_0$ as the first $n-1$ rows of $\mathbf{A}$, and $\mathbf{a}^T$ as the last row of $\mathbf{A}$. Using Siegel's lemma, $\mathbf{A}_0\mathbf{x} = \mathbf{0}$ has a set of linear independent solutions with norm at most $\mathrm{poly}(n)$, we write them as $\mathbf{x}_1, ..., \mathbf{x}_{m-n+1}$. Let $c_i = \mathbf{a}^T\mathbf{x}_i$, then $c_i = \mathrm{poly}(n)$ and $c_1, ..., c_{m-n+1}$ are co-prime (otherwise there is no integer solution for $\mathbf{A}\mathbf{x} = \mathbf{e}_n$, $\mathbf{e}_n = (0, ..., 0, 1)^T$). By Bezout's lemma, we construct $d_1, ..., d_{m-n+1}$ such that $d_i = \mathrm{poly}(n)$ and $c_1d_1 + ... + c_{m-n+1}d_{m-n+1} = 1$, so $d_1\mathbf{x}_1 + ... + d_{m-n+1}\mathbf{x}_{m-n+1}$ is an integer solution of $\mathbf{A}\mathbf{x} = \mathbf{e}_n$ with norm at most $\mathrm{poly}(n)$.
- Similarly, we construct integer solutions for $\mathbf{A}\mathbf{x} = \mathbf{e}_i$ for $i \in [n]$, and use them to construct a solution for $\mathbf{A}\mathbf{x} = \mathbf{b}$ with norm at most $\mathrm{poly}(n)$.

□

Now we find $\mathbf{r}$ such that $\mathbf{X}^T\mathbf{r} = \mathbf{b}$ and $\|\mathbf{r}\|_\infty \leq \delta$, and we can write $\mathbf{X}^T\mathbf{e}_1 + \mathbf{b}$ as $\mathbf{X}^T(\mathbf{e}_1 + \mathbf{r})$. So we only need to show that $\mathbf{e}_1$ and $\mathbf{e}_1 + \mathbf{r}$ are indistinguishable. By Lemma 2.3, we can choose large enough σ' such that $\mathbf{e}_1$ is statistical indistinguishable from $\mathbf{e}_1 + \mathbf{r}$.

We write $\mathbf{X}_{top} = (\mathbf{X}|\mathbf{X}')$. Now we only need to show that given $\mathbf{A}, \mathbf{Z}\mathbf{A}, \mathbf{X}, \mathbf{X}^T\mathbf{Z}, \mathbf{c}_0, \mathbf{X}_{top}^T\mathbf{c}_1^0$, $\mathbf{y}^T\mathbf{c}_1^0$ is indistinguishable from $\mathbf{y}^T\mathbf{c}_1^1$. The discussion is exactly the same as Theorem 2 in [5], except that the vector orthogonal to $\mathbf{X}_{top}$ here is $\mathbf{t}$, instead of $\mathbf{y}$. We omit the details here due to the page limits.