

Optimal Broadcast Encryption from LWE and Pairings in the Standard Model

Shweta Agrawal¹, Daniel Wichs², and Shota Yamada³

¹ IIT Madras,

`shweta.a@cse.iitm.ac.in`

² Northeastern University and NTT Research Inc,

`wichs@ccs.neu.edu`

³ National Institute of Advanced Industrial Science and Technology (AIST),

`yamada-shota@aist.go.jp`

Abstract. Broadcast Encryption with optimal parameters was a long-standing problem, whose first solution was provided in an elegant work by Boneh, Waters and Zhandry [?]. However, this work relied on multilinear maps of logarithmic degree, which is not considered a standard assumption. Recently, Agrawal and Yamada [?] improved this state of affairs by providing the first construction of optimal broadcast encryption from Bilinear Maps and Learning With Errors (LWE). However, their proof of security was in the generic bilinear group model. In this work, we improve upon their result by providing a new construction and proof in the standard model. In more detail, we rely on the Learning With Errors (LWE) assumption and the Knowledge of OrthogonALity Assumption (KOALA) [?] on bilinear groups.

Our construction combines three building blocks: a (computational) nearly linear secret sharing scheme with compact shares which we construct from LWE, an inner-product functional encryption scheme with special properties which is constructed from the bilinear Matrix Decision Diffie Hellman (MDDH) assumption, and a certain form of hyperplane obfuscation, which is constructed using the KOALA assumption. While similar to that of Agrawal and Yamada, our construction provides a new understanding of how to decompose the construction into simpler, modular building blocks with concrete and easy-to-understand security requirements for each one. We believe this sheds new light on the requirements for optimal broadcast encryption, which may lead to new constructions in the future.

1 Introduction

Broadcast encryption [?] (BE) is a novel form of encryption that enables a sender to transmit a single ciphertext over a broadcast channel so that only an authorized subset S of total N users can decrypt and recover the message. Security requires that no collusion of unauthorized users can learn anything about the encrypted message with non-negligible advantage. Evidently, broadcast

encryption is implied by public key encryption if no restriction is placed on the size of the ciphertext. However, the size of the ciphertext in broadcast encryption is of paramount importance, and is quantified in terms of ciphertext *overhead*, namely, the size of the ciphertext not counting the description of the recipient set S . Thus, in an optimal solution, the ciphertext overhead would be of size proportional to a symmetric encryption of the plaintext message (upto constant factors), aside from the description of S which is provided in the clear.

In a celebrated work, Boneh, Gentry and Waters [?] provided the first construction of broadcast encryption which achieved both optimal (constant) ciphertext overhead and short secret keys, but suffered from large public parameters, namely, linear in the number of users N . A series of elegant works provided improvements to this scheme [?, ?, ?, ?, ?, ?, ?] achieving many interesting new features such as anonymity, adaptive security and such others, but failed to improve the size of the public parameters. In 2014, Boneh, Waters and Zhandry [?] provided the first solution to the long standing problem of BE with optimal parameters, but their construction relied on the existence of multilinear maps of degree $\log N$, which is not considered a standard assumption. Recently, Agrawal and Yamada [?] improved the state of affairs by achieving the same parameters from the learning with errors assumption (LWE) along with assumptions on bilinear maps. However, this construction [?] could only be proven secure in the generic bilinear group model. Independently, Brakerski and Vaikuntanathan [?] also provided a construction of BE with optimal parameters from new assumptions on lattices, but they were unable to provide a proof of security for their scheme.

While encouraging, this state of affairs nevertheless leaves much to be desired. It is evident that for a primitive as important as broadcast encryption, we would like to have a proof from well-studied standard assumptions, and in the standard model. However, so far such a construction has been elusive.

Our Results. In this work, we make further progress towards this goal and provide the first construction for broadcast encryption with optimal parameters, from Learning with Errors (LWE) [?] and the Knowledge of OrthogonALity Assumption (KOALA) [?] in the standard model. While similar to that of Agrawal and Yamada, our construction provides a new understanding of how to decompose the construction into simpler, modular building blocks with concrete and easy-to-understand security requirements for each one. We believe this sheds new light on the requirements for optimal broadcast encryption, which may lead to new constructions in the future.

In more detail, as in [?], we provide a construction for ciphertext-policy attribute based encryption (cpABE) for NC_1 circuits, such that its ciphertext size, secret key size, and public key size are all independent of the *size* of the circuits supported by the scheme, and depend only on their input length and depth. Recall that in a cpABE scheme, a ciphertext for a message m is associated with a function (policy) f , and secret keys are associated with public attributes \mathbf{x} from the domain of f . Decryption succeeds to yield the hidden message m if

and only if the attribute satisfies the policy, namely $f(\mathbf{x}) = 1$. To see BE as a special case of cpABE, note that the circuit embedded in the ciphertext (F_S , say) can check for membership of a given user index in a set of authorised recipients S , and the attributes \mathbf{x} may encode a user’s index in the set N . Thus, a user i holds a secret key for attributes i and can decrypt a ciphertext associated with S if and only if i is a member of S . As observed in [?], the depth and input length of the circuit F_S are logarithmic in N , so it suffices to construct cpABE with parameters independent of the *width* of F_S , which is linear in N .

Building upon the construction of [?], we provide a new cpABE for NC₁ from the Learning with Errors (LWE) [?] and the Knowledge of OrthogonALity Assumption (KOALA) [?]. The LWE assumption introduced in the seminal work of Regev [?] enjoys worst case to average case hardness guarantees and is widely considered a standard assumption in the literature. The KOALA assumption introduced by Beullens and Wee [?] (also implicitly present in prior work such as [?]) may be viewed as a decisional analogue of the algebraic group model [?], which posits that the only way an adversary can compute a new group element is to take a linear combination of group elements already provided. More specifically, the KOALA assumption asserts that any adversary that can distinguish $g^{\mathbf{M}\mathbf{r}}$ from $g^{\mathbf{v}}$ for some matrix \mathbf{M} and random vectors \mathbf{r}, \mathbf{v} , must know some nontrivial vector $\mathbf{z} \neq 0$ such that $\mathbf{z} \mathbf{M} = 0$. Beullens and Wee provided a proof of the KOALA assumption in the generic group model. While KOALA is a “knowledge assumption” and therefore not considered a standard assumption, we believe it is a significant improvement over [?] to rely on the hardness of a specific assumption in the standard model, than to rely on the the generic group model for the security of the entire scheme.

Technical Overview. We proceed to outline the main ideas of our construction. As discussed above, we construct a *ciphertext-policy attribute-based encryption* (cpABE) for NC₁ circuits. The cpABE is *compact*, meaning that the size of the ciphertexts and keys are all small, proportional only to the length of the inputs and the depth of the supported circuits, but independent of the circuit size. Our construction combines three building blocks: a (computational) *nearly linear secret sharing* scheme with compact shares, which we construct from Learning With Errors (LWE), a certain form of *inner-product functional encryption* (IPFE) [?, ?, ?, ?], constructed from the bilinear Matrix Decision Diffie Hellman (MDDH) assumption, and a certain form of *hyperplane obfuscation* [?], constructed using the KOALA assumption. Next, we describe each of these primitives individually and outline how they are combined to construct our cpABE.

Nearly Linear Secret Sharing. Our main building block is a new type of secret sharing scheme. Given a message $\mu \in \{0, 1\}$ and a circuit C with ℓ -bit input, the scheme outputs 2ℓ shares $\{\text{share}_{i,b}\}_{i \in [\ell], b \in \{0,1\}}$. Each $\text{share}_{i,b}$ is a vector over \mathbb{Z}_p . For any $x \in \{0, 1\}^\ell$, let $\text{share}_x = \{\text{share}_{i,x_i}\}$, which we think of as a long vector produced by concatenating of all the component shares. If $C(x) = 0$ then share_x computationally hides the message μ , and moreover, share_x is even indistinguishable from a uniformly random vector. On the other hand, if

$C(x) = 1$ then there is an efficient method to reconstruct the message μ from share_x . Moreover, this reconstruction procedure is “nearly linear” in the sense that given C, x , one can efficiently determine some linear function f such that $f(\text{share}_x) = \mu \cdot \lceil p/2 \rceil + e$, where $|e| \ll p/2$ is some small polynomially bounded error.

We use LWE to construct this type of nearly linear secret sharing for all NC_1 circuits, where the size of the shares only depends on the security parameter and the depth of the circuit, but is independent of the circuit size. The construction closely follows the ideas behind the ABE scheme of [?] and the laconic function evaluation of [?]. As in [?], we are restricted to NC_1 because we require the magnitude of the error e to be polynomially bounded. The construction additionally relies on some uniformly random public parameters pp , which we will ignore throughout the introduction.

Towards cpABE from Secret Sharing. In our cpABE construction, to encrypt a message μ under a given policy specified by an NC_1 circuit C , the encryptor creates a “nearly linear secret sharing” of the message resulting in shares $\{\text{share}_{i,b}\}_{i \in [\ell], b \in \{0,1\}}$. At a high level, the encryptor then encrypts these shares using some form of *functional encryption* (FE) and outputs the FE ciphertext. Let us examine what kind of functional encryption would be helpful in this setting.

As a starting point, assume the shares are encrypted via an FE scheme such that a decryptor with a secret key for x only learns the subset $\text{share}_x = \{\text{share}_{i,x_i}\}$. Such an FE scheme is easy to construct by encrypting each of the 2ℓ shares under a different public key of a standard public-key encryption scheme and giving the decryptor the ℓ secret keys corresponding to the choice of x [?]. This would already provide security in the non-colluding setting – if the adversary has a secret key for a single value x such that $C(x) = 0$, then she cannot learn anything about the message by getting share_x . However, if the adversary has secret keys for even just two different values x_0, x_1 , such that $C(x_0) = C(x_1) = 0$, all bets are off; indeed, with our scheme, she could easily recover the message.

To fix the above problem, we rely on a more restricted form of FE where the decryptor with a secret key for a value x would not learn share_x in full, but rather only a *hyperplane obfuscation* of the vector share_x . A hyperplane obfuscation [?] of a vector \vec{v} allows one to test whether various affine functions h evaluate to $h(\vec{v}) = 0$, but should not reveal anything else about the obfuscated vector beyond having black-box access to such tests. When $C(x) = 1$, a hyperplane obfuscation of share_x is sufficient to decrypt the message μ , since we have a linear function f such that $f(\text{share}_x) = \mu \cdot \lceil p/2 \rceil + e$ and therefore, by testing whether $f(\text{share}_x) - e' = 0$ for all values e' in the polynomial range that e comes from, we can determine whether $\mu = 0$ or $\mu = 1$. For security, consider an adversary has secret keys for some q inputs $x^{(1)}, \dots, x^{(q)}$ such that $C(x^{(i)}) = 0$ and learns the corresponding hyperplane obfuscations of the vectors $\text{share}_{x^{(i)}}$. We know that each of the vectors $\text{share}_{x^{(i)}}$ is individually computationally indistinguishable from uniform, but mutually the vectors have non-trivial correlations and can be used to recover μ . We wish to conclude that the hyperplane obfuscations of $\text{share}_{x^{(i)}}$

are mutually indistinguishable from obfuscations of random and independent vectors. Indeed, this follows if we had a *composable* [?, ?] virtual-black-box (VBB) hyperplane obfuscator since, given black box access to each of these vectors, an adversary will never be able to find an affine function evaluates to 0 on any of them. To make this approach work, we therefore need to instantiate an appropriate hyperplane obfuscator together with a matching FE scheme that outputs a hyperplane obfuscations of share_x .

Hyperplane Obfuscation. We rely on the extremely simple hyperplane obfuscator of [?]. Let G be a cyclic group of order p with a generator g . To obfuscate a vector $\vec{v} \in \mathbb{Z}_p^n$, we choose a random $\gamma \leftarrow \mathbb{Z}_p$ and output $g^\gamma, g^{\gamma \cdot \vec{v}}$. This allows one to test if an affine function h evaluates to $h(\vec{v}) = 0$ by computing $\gamma \cdot h(\vec{v})$ in the exponent. The work of [?] shows that this is a VBB hyperplane obfuscator under a new assumption that they proposed, which can in retrospect be seen as a variant of the KOALA assumption restricted to spaces of dimension 1. However, they did not prove that the obfuscator is composable. In our work, we do not directly prove that this obfuscator satisfies composable VBB security, but rather prove that it satisfies a specialized property that suffices for us. Namely, we show that under the KOALA assumption the following holds: for any set of vectors that are individually indistinguishable from uniform but can be mutually correlated, one cannot distinguish between being given the hyperplane obfuscations of all the vectors in the set versus hyperplane obfuscations of uniformly random and independent vectors.

Functional Encryption for Inner Products. As the last step, we need to provide an appropriate (public-key) functional encryption (FE) scheme. Such an FE should allow us to encrypt a set of shares $\{\text{share}_{i,b}\}_{i \in [\ell], b \in \{0,1\}}$, and give out secret keys for value $x \in \{0,1\}^\ell$, so that such a ciphertext/key pair (only) reveals a fresh hyperplane obfuscation of share_x and nothing else. We can simplify this problem by relying on a simpler “component” FE scheme and then combining the component FEs to get what we need. The component FE should allow us to encrypt a scalar $s \in \mathbb{Z}_p$ and give out secret keys for values g^γ , so that such a ciphertext/key pair only reveals $g^{\gamma \cdot s}$ and nothing else. We want to component FE to satisfy unbounded-collusion simulation-based security. Given such a component FE scheme, we can instantiate a separate copy of it for each $i \in [\ell], b \in \{0,1\}$ and each position in the share vector. The encryptor then encrypts each position of each share vector $\text{share}_{i,b}$ under the appropriate copy of the component scheme. To create a secret key for $x \in \{0,1\}^\ell$ we choose a fresh random $\gamma \leftarrow \mathbb{Z}_p$ and give out a secret key for g^γ for each of the component schemes in locations (i, x_i) . This would ensure that, given an encryption of $\{\text{share}_{i,b}\}_{i \in [\ell], b \in \{0,1\}}$, a secret key for a value $x \in \{0,1\}^\ell$ only allows one to recover the hyperplane obfuscation of share_x given by $(g^\gamma, g^{\gamma \cdot \text{share}_x})$.

The above almost works, up to one subtlety. When we instantiate the component FE scheme, we do so using bilinear groups $(\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T)$ of order p with corresponding generators (g_1, g_2, g_T) and a bilinear map $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$. We can create an encryption of a scalar $s \in \mathbb{Z}_p$ and give out a secret key for

$g_2^\gamma \in \mathbb{G}_2$ so that the decryption of the ciphertext with the secret key reveals $g_T^{\gamma \cdot s}$. However, we can only guarantee simulation based security when the simulator is given $g_2^{\gamma \cdot s}$. In other words, there is a discrepancy between correctness (where the honest users decrypt the product in the exponent of g_T) and security (where the simulator needs to know the product in the exponent of g_2). It turns out that this suffices for us. For correctness, the decryptor gets a hyperplane obfuscation over \mathbb{G}_T , which suffices to recover the message. For security, we need to rely on the hyperplane obfuscation being secure even when given over \mathbb{G}_2 , which just requires us to assume that KOALA holds over \mathbb{G}_2 . We instantiate the above type of component FE in a black-box way using the recent primitive of “Slotted Inner Product Functional Encryption” [?, ?, ?].⁴

2 Preliminaries

In this section, we define some notation and preliminaries that we require.

Notation. We use bold letters to denote vectors. We treat a vector as a row vector by default. The notation $[a, b]$ denotes the set of integers $\{k \in \mathbb{N} \mid a \leq k \leq b\}$. We use $[n]$ to denote the set $[1, n]$. Throughout the paper, we use λ to denote the security parameter. We say a function $f(\lambda)$ is *negligible* if it is $O(\lambda^{-c})$ for all $c > 0$, and we use $\text{negl}(\lambda)$ to denote a negligible function of λ . We say $f(\lambda)$ is *polynomial* if it is $O(\lambda^c)$ for some constant $c > 0$, and we use $\text{poly}(\lambda)$ to denote a polynomial function of λ . Throughout the paper, we consider non-uniform adversaries that are modeled as polynomial-size circuits $\mathcal{A} = \{\mathcal{A}_\lambda\}_\lambda$ indexed by the security parameter. We often drop the subscript when it is clear from the context.

2.1 Bilinear Map Preliminaries

Here, we introduce our notation for bilinear maps and the bilinear generic group model following Baltico et al. [?], who specializes the framework by Barthe [?] for defining generic k -linear groups to the bilinear group settings. The definition closely follows that of Maurer [?], which is equivalent to the alternative formulation by Shoup [?].

Notation on Bilinear Maps. A bilinear group generator `GroupGen` takes as input 1^λ and outputs a group description $\mathbb{G} = (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, g_1, g_2)$, where p is a prime of $\Theta(\lambda)$ bits, $\mathbb{G}_1, \mathbb{G}_2$, and \mathbb{G}_T are cyclic groups of order p , $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$ is a non-degenerate bilinear map, and g_1 and g_2 are generators of \mathbb{G}_1 and \mathbb{G}_2 , respectively. We require that the group operations in $\mathbb{G}_1, \mathbb{G}_2$, and \mathbb{G}_T as well as the bilinear map e can be efficiently computed. We employ the implicit

⁴ In the technical sections we use IPFE directly rather than first showing that it provides an FE scheme with the above discrepancy between correctness and security and then relying on such an FE. This is purely to avoid proliferation of additional definitions/abstractions.

representation of group elements: for a matrix \mathbf{A} over \mathbb{Z}_p , we define $[\mathbf{A}]_1 := g_1^{\mathbf{A}}$, $[\mathbf{A}]_2 := g_2^{\mathbf{A}}$, $[\mathbf{A}]_T := g_T^{\mathbf{A}}$, where exponentiation is carried out component-wise. We will use similar notation for vectors.

Generic Bilinear Group Model. Let $\mathbb{G} = (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, g_1, g_2)$ be a bilinear group setting, L_1 , L_2 , and L_T be lists of group elements in \mathbb{G}_1 , \mathbb{G}_2 , and \mathbb{G}_T respectively, and let \mathcal{D} be a distribution over L_1 , L_2 , and L_T . The generic group model for a bilinear group setting \mathbb{G} and a distribution \mathcal{D} is described in Fig. 1. In this model, the challenger first initializes the lists L_1 , L_2 , and L_T by sampling the group elements according to \mathcal{D} , and the adversary receives handles for the elements in the lists. For $s \in \{1, 2, T\}$, $L_s[h]$ denotes the h -th element in the list L_s . The handle to this element is simply the pair (s, h) . An adversary running in the generic bilinear group model can apply group operations and bilinear maps to the elements in the lists. To do this, the adversary has to call the appropriate oracle specifying handles for the input elements. The challenger computes the result of a query, stores it in the corresponding list, and returns to the adversary its (newly created) handle. Handles are not unique (i.e., the same group element may appear more than once in a list under different handles).

We remark that we slightly simplify the generic group model of Baltico et. al [?]. Whereas they allow the adversary to access the equality test oracle, which is given two handles (s, h_1) and (s, h_2) and returns 1 if $L_s[h_1] = L_s[h_2]$ and 0 otherwise for all $s \in \{1, 2, T\}$, we replace this oracle with the zero-test oracle, which is given a handle (s, h) and returns 1 if $L_s[h] = 0$ and 0 otherwise only for the case of $s = T$. We claim that even with this modification, the model is equivalent to the original one. This is because we can perform the equality test for (s, h_1) and (s, h_2) using our restricted oracles as follows. Let us first consider the case of $s = T$. In this case, we can get the handle (T, h') corresponding to $L_T[h_1] - L_T[h_2]$ by calling neg_T (see Figure 1). We then make a zero-test query for (T, h') . Clearly, we get 1 if $L_s[h_1] = L_s[h_2]$ and 0 otherwise. We next consider the case of $s \in \{1, 2\}$. This case can be reduced to the case of $s = T$ by lifting the group elements corresponding to h_1 and h_2 to the group elements in \mathbb{G}_T by taking bilinear maps with an arbitrary non-unit group element in \mathbb{G}_{3-s} , which is possible by calling map_e .

Symbolic Group Model. The symbolic group model for a bilinear group setting \mathbb{G} and a distribution \mathcal{D}_P gives to the adversary the same interface as the corresponding generic group model, except that internally the challenger stores lists of elements in the field $\mathbb{Z}_p[X_1, \dots, X_n]$ instead of lists of group elements. The oracles add_s , neg_s , map , and zt computes addition, negation, multiplication, and equality in the field.

2.2 Slotted Inner Product Functional Encryption

We need slotted Inner Product Functional Encryption (IPFE) due to Lin and Vaikuntanathan [?, ?, ?]. Slotted IPFE is a hybrid between a secret-key function-hiding IPFE and a public-key IPFE. In this scheme, a vector $\mathbf{u} \in \mathbb{Z}_p^n$ is divided

State: Lists L_1, L_2, L_T over $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T$ respectively.

Initializations: Lists L_1, L_2, L_T sampled according to distribution \mathcal{D} .

Oracles: The oracles provide black-box access to the group operations, the bilinear map, and equalities.

- For all $s \in \{1, 2, T\}$: $\text{add}_s(h_1, h_2)$ appends $L_s[h_1] + L_s[h_2]$ to L_s and returns its handle $(s, |L_s|)$.
- For all $s \in \{1, 2, T\}$: $\text{neg}_s(h_1, h_2)$ appends $L_s[h_1] - L_s[h_2]$ to L_s and returns its handle $(s, |L_s|)$.
- $\text{map}_e(h_1, h_2)$ appends $e(L_1[h_1], L_2[h_2])$ to L_T and returns its handle $(T, |L_T|)$.
- $\text{zt}_T(h)$ returns 1 if $L_T[h] = 0$ and 0 otherwise.

All oracles return \perp when given invalid indices.

Fig. 1. Generic group model for bilinear group setting $\mathbb{G} = (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, g_1, g_2)$ and distribution \mathcal{D} .

into a public and private part respectively $\mathbf{u} = (\mathbf{u}_{\text{pub}}, \mathbf{u}_{\text{priv}})$ such that given the master secret key, the encryption algorithm can encrypt any vector \mathbf{u} of its choice, but given only the public key, it can encrypt only to the public slot, i.e. $\mathbf{u}_{\text{priv}} = 0$. Slotted IPFE can guarantee function hiding only with respect to the private slot. We provide the definitions from [?].

Let **GroupGen** be a group generator that outputs bilinear group $\mathbb{G} = (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, [1]_1, [1]_2)$. A slotted inner-product functional encryption (IPFE) scheme based on \mathbb{G} consists of 5 efficient algorithms:

Setup $(1^\lambda, \mathfrak{s}_{\text{pub}}, \mathfrak{s}_{\text{pri}}) \rightarrow (\text{mpk}, \text{msk})$: The setup algorithm takes as input two disjoint index sets, the public slot $\mathfrak{s}_{\text{pub}}$ and the private slot $\mathfrak{s}_{\text{pri}}$, and outputs a pair of master public key and master secret key (mpk, msk) . The whole index set \mathfrak{s} is $\mathfrak{s}_{\text{pub}} \cup \mathfrak{s}_{\text{pri}}$.

KeyGen $(\text{msk}, [\mathbf{v}]_2) \rightarrow \text{sk}_{\mathbf{v}}$: The key generation algorithm takes as input the master secret key and an encoding of a function vector $[\mathbf{v}]_2$, and outputs a secret key $\text{sk}_{\mathbf{v}}$ for $\mathbf{v} \in \mathbb{Z}_p^{\mathfrak{s}}$.

Enc $(\text{msk}, [\mathbf{u}]_1) \rightarrow \text{ct}_{\mathbf{u}}$: The encrypt algorithm takes input the master secret key and an encoding of a message vector $[\mathbf{u}]_1$ and outputs a ciphertext $\text{ct}_{\mathbf{u}}$ for $\mathbf{u} \in \mathbb{Z}_p^{\mathfrak{s}}$.

Dec $(\text{sk}_{\mathbf{v}}, \text{ct}_{\mathbf{u}}) \rightarrow T \vee \perp$: The decrypt algorithm takes as input a secret key $\text{sk}_{\mathbf{v}}$ and a ciphertext $\text{ct}_{\mathbf{u}}$, and outputs an element $T \in \mathbb{G}_T$ or \perp .

SlotEnc $(\text{mpk}, [\mathbf{u}_{\text{pub}}]_1) \rightarrow \text{ct}_{\mathbf{u}}$: The slot encryption algorithm takes as input the master public key and a vector $\mathbf{u}_{\text{pub}} \in \mathbb{Z}_p^{\mathfrak{s}_{\text{pub}}}$, sets $\mathbf{u} = (\mathbf{u}_{\text{pub}}, \mathbf{0}) \in \mathbb{Z}_p^{\mathfrak{s}}$ and outputs a ciphertext $\text{ct}_{\mathbf{u}}$.

Correctness. We say the slotted inner-product functional encryption scheme satisfies decryption correctness if for all $\lambda \in \mathbb{N}$, all index sets \mathfrak{s} and all vectors $\mathbf{u}, \mathbf{v} \in \mathbb{Z}_p^{\mathfrak{s}}$,

$$\Pr \left[\text{Dec}(\text{sk}_{\mathbf{v}}, \text{ct}_{\mathbf{u}}) = \langle \mathbf{u}, \mathbf{v} \rangle \left| \begin{array}{l} \text{msk} \leftarrow \text{Setup}(\mathfrak{s}_{\text{pub}}, \mathfrak{s}_{\text{pri}}) \\ \text{sk}_{\mathbf{v}} \leftarrow \text{KeyGen}(\text{msk}, [\mathbf{v}]_2) \\ \text{ct}_{\mathbf{u}} \leftarrow \text{Enc}(\text{msk}, [\mathbf{u}]_1) \end{array} \right. \right] = 1 .$$

We say the slotted inner-product functional encryption scheme satisfies slot-mode correctness if for all $\lambda \in \mathbb{N}$, all disjoint index sets $\mathfrak{s}_{\text{pub}}, \mathfrak{s}_{\text{pri}}$ and all vectors $\mathbf{u} \in \mathbb{Z}_p^{\mathfrak{s}_{\text{pub}}}$, the following two distributions should be identical:

$$\left\{ (\text{mpk}, \text{msk}, \text{ct}) \left| \begin{array}{l} (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, \mathfrak{s}_{\text{pub}}, \mathfrak{s}_{\text{pri}}) \\ \text{ct}_{\mathbf{u}} \leftarrow \text{Enc}(\text{msk}, [\mathbf{u} \parallel \mathbf{0}]_1) \end{array} \right. \right\}$$

and

$$\left\{ (\text{mpk}, \text{msk}, \text{ct}) \left| \begin{array}{l} (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, \mathfrak{s}_{\text{pub}}, \mathfrak{s}_{\text{pri}}) \\ \text{ct}_{\mathbf{u}} \leftarrow \text{SlotEnc}(\text{msk}, [\mathbf{u}]_1) \end{array} \right. \right\}$$

Slotted IPFE generalizes both secret-key and public-key IPFEs: we may obtain the former by setting $\mathfrak{s} = \mathfrak{s}_{\text{pri}}$ and the latter by setting $\mathfrak{s} = \mathfrak{s}_{\text{pub}}$.

Next, we define the adaptive function hiding property.

Definition 2.1 (Function Hiding Slotted IPFE).

Let $(\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec}, \text{SlotEnc})$ be a slotted IPFE scheme as defined above. The scheme is function hiding if Exp_{FH}^0 is indistinguishable from Exp_{FH}^1 for all efficient adversary $\mathcal{A} = \{\mathcal{A}_\lambda\}_\lambda$ where Exp_{FH}^b for $b \in \{0, 1\}$ is defined as follows:

1. **Setup:** Run the adversary \mathcal{A}_λ and obtain the disjoint index sets $\mathfrak{s}_{\text{pub}}, \mathfrak{s}_{\text{pri}}$ from \mathcal{A}_λ . Let $\mathfrak{s} = \mathfrak{s}_{\text{pub}} \cup \mathfrak{s}_{\text{pri}}$. Let $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, \mathfrak{s}_{\text{pub}}, \mathfrak{s}_{\text{pri}})$ and return mpk to \mathcal{A}_λ .
2. **Challenge:** Repeat the following for arbitrarily many rounds determined by \mathcal{A}_λ : In each round, \mathcal{A}_λ has 2 options:
 - \mathcal{A}_λ chooses $\mathbf{v}_j^0, \mathbf{v}_j^1 \in \mathbb{Z}_p^{\mathfrak{s}}$ and submits $[\mathbf{v}_j^0]_2, [\mathbf{v}_j^1]_2$ for a secret key. Upon receiving this, compute $\text{sk}_j \leftarrow \text{KeyGen}(\text{msk}, [\mathbf{v}_j^b]_2)$ and return this to \mathcal{A}_λ .
 - \mathcal{A}_λ chooses $\mathbf{u}_i^0, \mathbf{u}_i^1 \in \mathbb{Z}_p^{\mathfrak{s}}$ and submits $[\mathbf{u}_i^0]_1, [\mathbf{u}_i^1]_1$ for a ciphertext. Upon receiving this, compute $\text{ct}_i \leftarrow \text{Enc}(\text{msk}, [\mathbf{u}_i^b]_1)$ and return this to \mathcal{A}_λ .
3. **Guess:** \mathcal{A}_λ outputs its guess b' .

The outcome of the experiment is defined as b' if all the public components of the key queries are equal, i.e. $\mathbf{v}_j^0|_{\mathfrak{s}_{\text{pub}}} = \mathbf{v}_j^1|_{\mathfrak{s}_{\text{pub}}}$ for all j and $\langle \mathbf{u}_i^0, \mathbf{v}_j^0 \rangle = \langle \mathbf{u}_i^1, \mathbf{v}_j^1 \rangle$ for all i, j .

We will also require the following lemma by [?, ?, ?, ?, ?]:

Lemma 2.2. *Let GroupGen be a group generator that outputs bilinear group $\mathbb{G} = (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, [1]_1, [1]_2)$ and $k \geq 1$ an integer constant. If MDDH_k holds in both \mathbb{G}_1 and \mathbb{G}_2 , then there is an (adaptively) function-hiding slotted IPFE scheme on GroupGen.*

Note that the MDDH_k assumption on \mathbb{G}_s ($s \in \{1, 2\}$) says that a random group element $[\mathbf{r}]_s$ is indistinguishable from $[\mathbf{sA}]_s$ given $[\mathbf{A}]$, where $\mathbf{A} \leftarrow \mathbb{Z}_p^{k \times (k+1)}$, $\mathbf{r} \in \mathbb{Z}_p^{k+1}$, and $\mathbf{s} \in \mathbb{Z}_p^k$. The assumption is implied by the standard k -LIN assumption, which becomes progressively weaker as k becomes larger.

2.3 Attribute Based Encryption

Let $R = \{R_\lambda : A_\lambda \times B_\lambda \rightarrow \{0, 1\}\}_\lambda$ be a relation where A_λ and B_λ denote “ciphertext attribute” and “key attribute” spaces. An attribute-based encryption (ABE) scheme for R is defined by the following PPT algorithms:

- Setup**(1^λ) \rightarrow (**mpk**, **msk**): The setup algorithm takes as input the unary representation of the security parameter λ and outputs a master public key **mpk** and a master secret key **msk**.
- Enc**(**mpk**, X , μ) \rightarrow **ct**: The encryption algorithm takes as input a master public key **mpk**, a ciphertext attribute $X \in A_\lambda$, and a message bit μ . It outputs a ciphertext **ct**.
- KeyGen**(**mpk**, **msk**, Y) \rightarrow sk_Y : The key generation algorithm takes as input the master public key **mpk**, the master secret key **msk**, and a key attribute $Y \in B_\lambda$. It outputs a private key sk_Y .
- Dec**(**mpk**, **ct**, X , sk_Y , Y) \rightarrow μ or \perp : The decryption algorithm takes as input the master public key **mpk**, a ciphertext **ct**, ciphertext attribute $X \in A_\lambda$, a private key sk_Y , and private key attribute $Y \in B_\lambda$. It outputs the message μ or \perp which represents that the ciphertext is not in a valid form.

Definition 2.3 (Correctness).

An ABE scheme for relation family R is correct if for all $\lambda \in \mathbb{N}$, $X \in A_\lambda$, $Y \in B_\lambda$ such that $R(X, Y) = 1$, and for all messages $\mu \in \mathcal{M}$,

$$\Pr \left[\begin{array}{l} (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda), \text{sk}_Y \leftarrow \text{KeyGen}(\text{mpk}, \text{msk}, Y), \\ \text{ct} \leftarrow \text{Enc}(\text{mpk}, X, \mu) : \text{Dec}(\text{mpk}, \text{sk}_Y, Y, \text{ct}, X) \neq \mu \end{array} \right] = \text{negl}(\lambda)$$

where the probability is taken over the coins of **Setup**, **KeyGen**, and **Enc**.

Definition 2.4 (Sel-IND security for ABE). For an ABE scheme $\text{ABE} = \{\text{Setup}, \text{Enc}, \text{KeyGen}, \text{Dec}\}$ for a relation family $R = \{R_\lambda : A_\lambda \times B_\lambda \rightarrow \{0, 1\}\}_\lambda$ and a message space $\{\mathcal{M}_\lambda\}_{\lambda \in \mathbb{N}}$ and an efficient adversary $\mathcal{A} = \{\mathcal{A}_\lambda\}_\lambda$, let us define Sel-IND security game as follows.

1. **Choosing the Target:** At the beginning of the game, \mathcal{A}_λ chooses its target $X^* \in A_\lambda$ and sends to the challenger.
2. **Setup phase:** On input 1^λ , the challenger samples $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda)$ and gives **mpk** to \mathcal{A}_λ .
3. **Query phase:** During the game, \mathcal{A}_λ adaptively makes the following queries, in an arbitrary order. \mathcal{A}_λ can make unbounded many key queries, but can make only single challenge query.

- (a) **Key Queries:** \mathcal{A}_λ chooses an input $Y \in B_\lambda$. For each such query, the challenger replies with $\text{sk}_Y \leftarrow \text{KeyGen}(\text{mpk}, \text{msk}, Y)$.
 - (b) **Challenge Query:** At some point, \mathcal{A}_λ submits a pair of equal length messages $(\mu_0, \mu_1) \in (\mathcal{M})^2$ to the challenger. The challenger samples a random bit $\beta \leftarrow \{0, 1\}$ and replies to \mathcal{A}_λ with $\text{ct} \leftarrow \text{Enc}(\text{mpk}, X^*, \mu_\beta)$. We require that $R(X^*, Y) = 0$ holds for any Y such that \mathcal{A}_λ makes a key query for Y in order to avoid trivial attacks.
4. **Output phase:** \mathcal{A}_λ outputs a guess bit β' as the output of the experiment.

We define the advantage $\text{Adv}_{\text{ABE}, \mathcal{A}}^{\text{Sel-IND}}(1^\lambda)$ of \mathcal{A} in the above game as

$$\text{Adv}_{\text{ABE}, \mathcal{A}}^{\text{Sel-IND}}(1^\lambda) := |\Pr[\mathcal{A} \text{ outputs } 1 | \beta = 0] - \Pr[\mathcal{A} \text{ outputs } 1 | \beta = 1]|.$$

The ABE scheme ABE is said to satisfy Sel-IND security (or simply selective security) if for any efficient and stateful adversary $\mathcal{A} = \{\mathcal{A}_\lambda\}_\lambda$, there exists a negligible function $\text{negl}(\cdot)$ such that $\text{Adv}_{\text{ABE}, \mathcal{A}}^{\text{Sel-IND}}(1^\lambda) \neq \text{negl}(\lambda)$.

We can consider the following stronger version of the security where we require the ciphertext to be pseudorandom.

Definition 2.5 (Sel-INDr security for ABE). We define Sel-INDr security game similarly to Sel-IND security game except that the adversary \mathcal{A} chooses single message μ instead of (μ_0, μ_1) at the challenge phase and the challenger returns $\text{ct} \leftarrow \text{Enc}(\text{mpk}, X^*, \mu)$ if $\beta = 0$ and a random ciphertext $\text{ct} \leftarrow \mathcal{CT}$ from a ciphertext space \mathcal{CT} if $\beta = 1$. We define the advantage $\text{Adv}_{\text{ABE}, \mathcal{A}}^{\text{Sel-INDr}}(1^\lambda)$ of the adversary \mathcal{A} accordingly and say that the scheme satisfies Sel-INDr security if the quantity is negligible.

We also consider (weaker) version of the above notions, where \mathcal{A} specifies the set \mathcal{Y} of attributes for which it makes key queries along with X^* at the beginning of the game.

Definition 2.6 (VerSel-IND security for ABE). We define VerSel-IND security game as Sel-IND security game with the exception that the adversary \mathcal{A} has to choose the set $\mathcal{Y} \subseteq B_\lambda$ for which it makes key queries along with the challenge ciphertext attribute X^* before the setup phase but the choice of (μ_0, μ_1) can still be adaptive. After that, \mathcal{A}_λ can make key queries for Y_1, Y_2, \dots adaptively, but we need $Y_i \in \mathcal{Y}$ for all queries. We define the advantage $\text{Adv}_{\text{ABE}, \mathcal{A}}^{\text{VerSel-IND}}(1^\lambda)$ of the adversary \mathcal{A} accordingly and say that the scheme satisfies VerSel-IND security (or simply very selective security) if the quantity is negligible.

In the following, we define standard notions of ciphertext-policy attribute-based encryption (CP-ABE) and broadcast encryption (BE) by specifying the relation R .

CP-ABE for circuits. We define CP-ABE for circuit class $\{\mathcal{C}_\lambda\}_\lambda$ by specifying the relation. Here, \mathcal{C}_λ is a set of circuits with input length $\ell(\lambda)$ and binary output. We define $A_\lambda^{\text{CP}} = \mathcal{C}_\lambda$ and $B_\lambda^{\text{CP}} = \{0, 1\}^\ell$. Furthermore, we define the relation R_λ^{CP} as $R_\lambda^{\text{CP}}(C, \mathbf{x}) = \neg C(\mathbf{x})$.⁵

⁵ Here, we follow the standard convention in lattice-based cryptography where the decryption succeeds when $C(\mathbf{x}) = 0$ rather than $C(\mathbf{x}) = 1$.

BE. To define BE, we define $A_\lambda^{\text{BE}} = 2^{[N(\lambda)]}$ and $B_\lambda^{\text{BE}} = [N(\lambda)]$, where $N(\lambda) = \text{poly}(\lambda)$ is the number of users in the system and $2^{[N(\lambda)]}$ denotes all subsets of $[N]$. We also define $R_\lambda^{\text{BE}} : A_\lambda^{\text{BE}} \times B_\lambda^{\text{BE}} \rightarrow \{0, 1\}$ as $R_\lambda^{\text{BE}}(S, i) = 1$ when $i \in S$ and $R_\lambda^{\text{BE}}(S, i) = 0$ otherwise. For BE, we typically require that the ciphertext size should be $o(N) \cdot \text{poly}(\lambda)$, since otherwise we have a trivial construction from plain public key encryption.

Remark 2.7. We note that very selective security and selective security are in fact equivalent in the case of BE since one can convert selective adversary into very selective adversary as follows. Namely, if the selective adversary chooses its target $S \subseteq [N]$, the very selective adversary chooses the same target S and specifies the set of user indices for which it makes key queries as $[N] \setminus S$. Then, the very selective adversary can simulate the game for the selective adversary using the secret keys given by the challenger.

3 Computational Secret Sharing Scheme with Short Shares

We will use secret sharing scheme with special properties that we are going to define here. Let $\mathcal{C} = \{\mathcal{C}_\lambda\}_\lambda$ be a circuit class. A secret sharing scheme for the circuit class \mathcal{C} is defined by the following PPT algorithms:

- SS.Setup($1^\lambda, p$) \rightarrow pp: The setup algorithm takes as input the unary representation of the security parameter λ and the modulus p and outputs public parameter pp.
- SS.Share(pp, C, μ): The sharing algorithm takes as input the public parameter pp, a circuit $C \in \mathcal{C}_\lambda$ that specifies access policy, and a message $\mu \in \{0, 1\}$ to be shared and outputs a set of shares $\{\text{share}_{i,b} \in \mathbb{Z}_p^m\}_{i \in [\ell], b \in \{0,1\}}$, where ℓ is the input length of C and m is a parameter specified by λ and p .
- SS.Recon(pp, $C, x, \{\text{share}_{i,x_i}\}_{i \in [\ell]}$) \rightarrow μ or \perp : The reconstruction algorithm takes as input the public parameter pp, a circuit C , an input $x \in \{0, 1\}^\ell$ to the circuit, and shares $\{\text{share}_{i,x_i}\}_{i \in [\ell]}$ and outputs message μ or \perp .

We require correctness and security for the secret sharing scheme as defined in the following.

Definition 3.1 (Correctness). *We say that a secret sharing scheme $\text{SS} = (\text{SS.Setup}, \text{SS.Share}, \text{SS.Recon})$ for circuit class \mathcal{C} has correctness if there exists a function $p_0(\lambda)$ specified by the circuit class \mathcal{C} such that for any $p > p_0(\lambda)$, $C \in \mathcal{C}$ with input length ℓ , $x \in \{0, 1\}^\ell$ satisfying $C(x) = 1$, and $\mu \in \{0, 1\}$, we have*

$$\Pr \left[\begin{array}{l} \text{pp} \leftarrow \text{SS.Setup}(1^\lambda, p), \\ \{\text{share}_{i,b}\}_{i \in [\ell], b \in \{0,1\}} \leftarrow \text{SS.Share}(\text{pp}, C, \mu), \\ \text{SS.Recon}(\text{pp}, \{\text{share}_{i,x_i}\}_{i \in [\ell]}) = \mu \end{array} \right] = 1$$

Definition 3.2 (Security). We say that a secret sharing scheme $\text{SS} = (\text{SS.Setup}, \text{SS.Share}, \text{SS.Recon})$ for circuit class $\mathcal{C} = \{\mathcal{C}_\lambda\}_\lambda$ is secure if for any $C \in \mathcal{C}_\lambda$ with input length $\ell = \ell(\lambda)$, $x \in \{0, 1\}^\ell$ satisfying $C(x) = 0$, $\mu \in \{0, 1\}$, $p \in \mathbb{N}$, and for any efficient adversary $\mathcal{A} = \{\mathcal{A}_\lambda\}$, we have

$$\left| \Pr \left[\mathcal{A}_\lambda \left(\begin{array}{c} \text{pp}, C, x, \\ \{\text{share}_{i,x_i}\}_{i \in [\ell]} \end{array} \right) \rightarrow 1 \right] - \Pr \left[\mathcal{A}_\lambda \left(\begin{array}{c} \text{pp}, C, x, \\ \{\mathbf{v}_i\}_{i \in [\ell]} \end{array} \right) \rightarrow 1 \right] \right| = \text{negl}(\lambda)$$

where the probability is taken over the choice of $\text{pp} \leftarrow \text{SS.Setup}(1^\lambda, p)$, $\{\text{share}_{i,b}\}_{i,b} \leftarrow \text{SS.Share}(\text{pp}, C, \mu)$, $\mathbf{v}_i \leftarrow \mathbb{Z}_p^m$ for $i \in [\ell]$, and the internal coin of \mathcal{A}_λ .

We note that in the above, we not only require that the shares do not reveal μ if $C(x) = 0$, but also require that they look random.

We furthermore require following structural properties for the construction. First, we require that pp is a random string.

Definition 3.3 (Random Public Parameters). We require that pp output by $\text{SS.Setup}(1^\lambda, p)$ is statistically close to uniformly random, where the length of the string is deterministically determined by p .

Looking ahead, the above property is crucial when we prove the security of our ABE scheme. If the public parameter of the secret sharing scheme was chosen from a structured distribution, we would have to rely on the bilinear KOALA assumption (Definition 4.1) with auxiliary input chosen from the same distribution. However, we cannot hope the assumption to hold for auxiliary input with general distribution as we will discuss in Remark 4.3.

We also require that the reconstruction algorithm is structured as two steps: a function evaluation step that computes the circuit on the shares to yield the message along with noise, followed by a rounding step that removes the noise. We require that the first step is *linear*. We refer to such a reconstruction algorithm as being “almost linear”, and define it formally next.

Definition 3.4 (Almost Linear Reconstruction). We say that a secret sharing scheme $\text{SS} = (\text{SS.Setup}, \text{SS.Share}, \text{SS.Recon})$ has almost linear reconstruction if the reconstruction algorithm is divided into two steps:

- Step 1 takes as input the public parameter pp , the circuit C , and the input x . It outputs a set of coefficients $\{a_{i,j} \in \mathbb{Z}_p\}_{i \in [\ell], j \in [m]}$. We denote this step as an algorithm $\text{SS.FindCoef}(\text{pp}, C, x)$.
- Step 2 takes as input the set of shares $\{\text{share}_{i,x_i}\}_{i \in [\ell]}$ that corresponds to x and a set of linear coefficients $\{a_{i,j}\}_{i \in [\ell], j \in [m]}$ and computes

$$d := \sum_{i \in [\ell], j \in [m]} a_{i,j} \text{share}_{i,x_i,j} \pmod{p}$$

where $\text{share}_{i,x_i,j} \in \mathbb{Z}_p$ is the j -th entry of the vector share_{i,x_i} . It then outputs 1 if d is closer to $p/2$ and 0 otherwise.

We require the following property, which implies the correctness: For any $x \in \{0, 1\}^\ell$ and $C \in \mathcal{C}$ satisfying $C(x) = 1$ and $p > p_0$, if we have:

$$\begin{aligned} \text{pp} &\leftarrow \text{SS.Setup}(1^\lambda, p), \{\text{share}_{i,b}\}_{i \in [\ell]} \leftarrow \text{SS.Share}(\text{pp}, C, \mu), \\ \{a_{i,j}\}_{i,j} &\leftarrow \text{SS.FindCoef}(\text{pp}, C, x), \end{aligned}$$

where $i \in [\ell], j \in [m]$, then there exists $e \in [-B, B]$ such that

$$\sum_{i \in [\ell], j \in [m]} a_{i,j} \text{share}_{i,x_i,j} = \mu \cdot \lceil p/2 \rceil + e \pmod p$$

where $B(\lambda)$ is an integer specified by \mathcal{C} .

The following theorem asserts that we can construct a secret sharing scheme with the desired properties under the LWE assumption.

Theorem 3.5. For circuit class $\mathcal{C}_{\ell,d} = \{\mathcal{C}_{\lambda,\ell(\lambda),d(\lambda)}\}_{\lambda \in \mathbb{N}}$ consisting of circuits whose input length is $\ell(\lambda) = \text{poly}(\lambda)$ and depth $d(\lambda) = O(\log \lambda)$, we have secret sharing scheme that satisfies almost linear reconstruction (Definition 3.4) for $p_0 = \text{poly}(\lambda, 2^d, \ell)$ and has random public parameters (Definition 3.3). We can prove the security of the scheme (Definition 3.2) under the LWE assumption with approximation factor $p^\epsilon \cdot \text{poly}(\lambda)$ for some constant $\epsilon < 1$. Furthermore, the size of the parameters in the construction is as follows:

$$|\text{pp}|, |\text{share}_{i,b}| \leq \text{poly}(\lambda, d, \ell), \quad B(\lambda) \leq \text{poly}(\lambda, 2^d). \quad (3.1)$$

In particular, $B(\lambda)$ is bounded by a polynomial in λ since $d = O(\log \lambda)$.

Proof. The construction is based on the ABE scheme for circuit class $\{\mathcal{C}_{\ell,d}\}$ by [?]. Here, we first show a construction that almost works but has a problem. We then fix the problem by slightly modifying the construction. In our first construction, we put the master public key and a secret key for circuit C of the ABE scheme into pp , where the former consists of set of random matrices \mathbf{A} and $\{\mathbf{B}_{i,b}\}_{i,b}$ along with a random vector \mathbf{u} . To generate $\{\text{share}_{i,b}\}_{i,b}$, we generate LWE samples using corresponding matrices in $\{\mathbf{B}_{i,b}\}_{i,b}$ with respect to the same secret, so that $\{\text{share}_{i,x_i}\}_i$ constitutes a valid ABE ciphertext for attribute x and message μ to be shared.⁶ Most of the properties we require for the secret sharing scheme are directly implied by the corresponding properties of the ABE scheme. The correctness (Definition 3.1) and the security (Definition 3.2) of the secret sharing scheme are implied by the corresponding properties of the ABE scheme, where for the latter we use Sel-INDr security of the ABE scheme. The size requirements for the parameters (Eq. (3.1)) are satisfied by the efficiency of the ABE scheme. The almost linear reconstruction property (Definition 3.4) is also satisfied by the structure of the decryption algorithm of the ABE scheme.

⁶ In fact, the ABE ciphertext also has to include the LWE samples with respect to the matrix \mathbf{A} and the vector \mathbf{u} , where the latter will be used to mask the message. These LWE samples are put into both of $\text{share}_{1,0}$ and $\text{share}_{1,1}$.

However, the above construction does not have the property of random public parameters (Definition 3.3) and we have to change it slightly. In particular, in the above, \mathbf{pp} is chosen as follows: We first sample random matrices \mathbf{A} and $\{\mathbf{B}_{i,b}\}_{i,b}$ and a random vector \mathbf{u} , where \mathbf{A} is chosen along with a trapdoor. We then compute a matrix \mathbf{B}_C corresponding to the circuit C from the matrices $\{\mathbf{B}_{i,b}\}_{i,b}$ and then generate a vector \mathbf{r} from a Gaussian distribution over the integer lattice with the restriction

$$[\mathbf{A}|\mathbf{B}_C]\mathbf{r} = \mathbf{u} \pmod{q}$$

using the trapdoor. We then set $\mathbf{pp} = (\mathbf{A}, \mathbf{B}, \mathbf{u}, \mathbf{r})$. Because of the above relation between \mathbf{r} and \mathbf{u} , \mathbf{pp} is not random.

To address this issue, we first remove \mathbf{u} from \mathbf{pp} . We now have that the distribution of \mathbf{r} is statistically close to the Gaussian distribution over the integer lattice (without the restriction), since for a vector \mathbf{r} chosen from (sufficiently wide) Gaussian distribution over the integer lattice, \mathbf{u} defined as $\mathbf{u} = [\mathbf{A}|\mathbf{B}_C]\mathbf{r} \pmod{q}$ is statistically close to uniform. Now, such \mathbf{r} can be chosen from randomness of fixed polynomial length, in particular, without a trapdoor. We then put the randomness R used for sampling \mathbf{r} into \mathbf{pp} instead of \mathbf{r} itself. We now have that \mathbf{pp} is statistically close to random as desired.

It is easy to see that this change does not affect the properties that we want from the secret sharing scheme. In particular, the correctness is not lost since \mathbf{u} can be recovered from \mathbf{r} . We note that for the security to be preserved, we need an efficient reverse sampling algorithm that is given \mathbf{r} and samples randomness R conditioned that the Gaussian sampler outputs \mathbf{r} on input R . The reason why we need this property is that the reduction algorithm that breaks the ABE scheme using the adversary against the secret sharing scheme should simulate the randomness R for sampling \mathbf{r} only given \mathbf{r} , which is the secret key of the ABE scheme.

This property is satisfied by efficient Gaussian samplers such as [?]. To see this, let us recall the procedure of sampling Gaussian on integer lattice in [?] (See Lemma 4.3 in the paper). Without loss of generality, we can consider one-dimensional case since multi-dimensional case can be handled by running the algorithm for one-dimensional case in parallel. The sampling algorithm by [?] was based on rejection sampling. The algorithm first samples a candidate for the output uniformly at random and outputs it with certain probability. This step is repeated until it outputs something or the number of times it repeats the procedure exceeds predetermined number. The idea for the reverse sampling is to first run the algorithm until it outputs something and then replace the randomness that was used for the output with that which leads to the intended output. The former step can be done straightforwardly since it is exactly the same as the original sampling algorithm. The purpose of performing this step is to simulate the failure. The latter step can be performed efficiently as well, since the randomness that leads to the output consists of the output value itself along with the randomness that allows the sampler to accept and output the value (rather than to reject). It is easy to see that this algorithm indeed works. This completes the proof of Theorem 3.5.

The complete description of the secret sharing scheme will appear in the full version.

4 Our Security Assumptions

In this section, we will introduce the bilinear knowledge of orthogonality (KOALA) assumption, which is an analogue of the KOALA assumption introduced in [?]. Looking ahead, the assumption will be used to prove the security of our ABE scheme in Sec. 5.1. We then introduce weak KOALA assumption (wKOALA) and show that it is implied by the bilinear KOALA assumption. While the former assumption is implied by the latter, the former is handier to use and would be of independent interest.

4.1 Bilinear KOALA Assumption

Here, we introduce bilinear KOALA assumption, which is an analogue of the KOALA assumption introduced in [?].

Definition 4.1 (Bilinear KOALA Assumption). *Let $\text{Samp} = \{\text{Samp}_\lambda\}_\lambda$ be an efficient sampling algorithm that takes as input an integer p and a string aux and outputs a matrix $\mathbf{V} \in \mathbb{Z}_p^{\ell_1 \times \ell_2}$ with $\ell_1 < \ell_2$. For an efficient adversary $\mathcal{A} = \{\mathcal{A}_\lambda\}$, let us define*

$$\text{Adv}_{\mathcal{A}, \mathbb{G}, \text{Samp}}^{\text{BKOALA, dist}}(\lambda) := |\Pr[\mathcal{A}_\lambda(\mathbb{G}, \text{aux}, [\mathbf{sV}]_2) \rightarrow 1] - \Pr[\mathcal{A}_\lambda(\mathbb{G}, \text{aux}, [\mathbf{r}]_2) \rightarrow 1]|.$$

where the probabilities are taken over the choice of uniformly random aux , $\mathbb{G} = (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, [1]_1, [1]_2) \leftarrow \text{GroupGen}(1^\lambda)$, $\mathbf{V} \leftarrow \text{Samp}(p, \text{aux})$, $\mathbf{s} \leftarrow \mathbb{Z}_p^{\ell_1}$, $\mathbf{r} \leftarrow \mathbb{Z}_p^{\ell_2}$, and the coin of \mathcal{A}_λ .

Furthermore, for an efficient adversary $\mathcal{B} = \{\mathcal{B}_\lambda\}_\lambda$, we also define

$$\text{Adv}_{\mathcal{B}, \mathbb{G}, \text{Samp}}^{\text{BKOALA, find}}(\lambda) := \Pr[\mathcal{B}_\lambda(\mathbb{G}, \text{aux}) \rightarrow \mathbf{x} \wedge \mathbf{xV}^\top = \mathbf{0} \wedge \mathbf{x} \neq \mathbf{0}]$$

where the probability is taken over the choice of uniformly random aux , $\mathbb{G} = (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, [1]_1, [1]_2) \leftarrow \text{GroupGen}(1^\lambda)$, $\mathbf{V} \leftarrow \text{Samp}_\lambda(p, \text{aux})$, and the coin of \mathcal{B}_λ .

We say that the bilinear KOALA assumption holds with respect to GroupGen if for any efficient adversary \mathcal{A} and efficient sampler Samp , there exists another efficient adversary \mathcal{B} and a polynomial function $Q(\lambda)$ such that

$$\text{Adv}_{\mathcal{B}, \mathbb{G}, \text{Samp}}^{\text{BKOALA, find}}(\lambda) \geq \text{Adv}_{\mathcal{A}, \mathbb{G}, \text{Samp}}^{\text{BKOALA, dist}}(\lambda)/Q(\lambda) - \text{negl}(\lambda).$$

Remark 4.2. Our definition of the bilinear KOALA assumption differs from the original KOALA assumption defined by Beullens and Wee [?] in several points. First, we consider the assumption in groups equipped with bilinear maps whereas they consider the assumption in groups without bilinear maps.

Second, in our assumption, the adversary is given an auxiliary input \mathbf{aux} and \mathbf{V} is chosen from a distribution specified by \mathbf{Samp} whereas there is no auxiliary information and \mathbf{V} is fixed in [?]. This change is necessary to prove the security of our ABE scheme in Section 5, since we will use an adversary \mathcal{B} that is obtained from \mathcal{A} to break a computational assumption, where we put the problem instance of the assumption into \mathbf{aux} .

Remark 4.3. One could consider simpler and stronger variant of the above assumption where \mathbf{Samp} chooses \mathbf{aux} along with \mathbf{V} instead of letting \mathbf{aux} to be a random string that is not controlled by \mathbf{Samp} . However, we cannot hope this variant of the assumption to hold for all efficient samplers. For example, let us consider a sampler that outputs random \mathbf{V} along with auxiliary information $\mathbf{aux} = \mathcal{O}(C_{\mathbf{V}})$, which is an obfuscation of circuit $C_{\mathbf{V}}$ that takes as input group description \mathbb{G} and elements $[\mathbf{v}]_2$ and returns whether \mathbf{v} is in the space spanned by the rows of \mathbf{V} or not. Using $\mathcal{O}(C_{\mathbf{V}})$, one can easily distinguish $[\mathbf{sV}]_2$ from $[\mathbf{r}]_2$ with high probability. However, an efficient adversary may not be able to find a vector $\mathbf{x} \neq \mathbf{0}$ that satisfies $\mathbf{xV} = \mathbf{0}$ even given $\mathcal{O}(C_{\mathbf{V}})$, if we use sufficiently strong obfuscator to obfuscate the circuit $C_{\mathbf{V}}$. Our assumption above excludes this kind of attack by making \mathbf{aux} to be public randomness that is not touched by the sampler. Our definition is inspired by that of public coin differing input obfuscation [?], where the authors exclude similar kind of attacks [?] in the context of differing input obfuscations by restricting the distribution of auxiliary input to be random.

The following theorem justifies the bilinear KOALA assumption on the bilinear generic group model. The proof is almost the same as that for the KOALA assumption in [?], but we have to adjust it into the setting where the groups are equipped with bilinear maps and the adversary is given auxiliary input.

Theorem 4.4. *The bilinear KOALA assumption holds under the bilinear generic group model, where \mathcal{A} has access to the generic group oracles but \mathbf{Samp} does not.*

Proof. Let us fix PPT sampler \mathbf{Samp} and an adversary \mathcal{A} . We also let $Q_{\text{zt}}(\lambda)$ be the upper bound on the number of zero test queries that \mathcal{A} makes. To prove the theorem, we consider following sequence of games. Let us denote the event that \mathcal{A} outputs 1 at the end of \mathbf{Game}_x as E_x .

Game₁: In this game, \mathbf{Samp} takes as input the order of groups p and a random string \mathbf{aux} and outputs $\mathbf{V} \in \mathbb{Z}_p^{\ell_1 \times \ell_2}$. By assumption, \mathbf{Samp} does not have access to the generic group oracles. Then, the adversary \mathcal{A} is given \mathbf{aux} , handles corresponding to the group elements $[\mathbf{r}]_2$, where $\mathbf{r} = (r_1, \dots, r_{\ell_2}) \leftarrow \mathbb{Z}_p^{\ell_2}$, and access to the oracles in generic group model and outputs a bit at the end of the game.

Game₂: In this game, we switch to symbolic group model and replace $r_1, \dots, r_{\ell_2} \in \mathbb{Z}_p$ with formal variables R_1, \dots, R_{ℓ_2} . Note that all handles given to \mathcal{A} during the game refer to a group element that is represented as

$$x_0 + \sum_{j \in [\ell_2]} x_j R_j \in \mathbb{Z}_p[R_1, \dots, R_{\ell_2}]$$

where the challenger computes coefficients $\{x_j \in \mathbb{Z}_p\}_{j \in [0, \ell_2]}$ by keeping track of the group operations performed by \mathcal{A} .

We observe that this game differs from the previous game only when \mathcal{A} makes a zero test query for $x_0 + \sum_{j \in [\ell_2]} x_j R_j$ such that $x_0 + \sum_{j \in [\ell_2]} x_j R_j \neq 0$ but $x_0 + \sum_{j \in [\ell_2]} x_j r_j$. However, this occurs with probability at most $1/p$ since \mathbf{r} is chosen uniformly at random independently from anything else. Therefore, we have

$$|\Pr[\mathbf{E}_2] - \Pr[\mathbf{E}_1]| = \text{negl}(\lambda).$$

Game₃: In this game, we replace the formal variable R_j with

$$\sum_{i \in [\ell_1]} v_{i,j} S_i$$

for all $j \in [\ell_2]$, where $v_{i,j}$ is the (i, j) -th entry of \mathbf{V} and S_1, \dots, S_{ℓ_1} are set of formal variables.

Game₄: In this game, we switch back to the generic group model (rather than the symbolic group model) and provide the adversary with handles for $[\mathbf{sV}]_2$ as input.

By the same reason as the game hop from **Game₁** to **Game₂**, we have

$$|\Pr[\mathbf{E}_4] - \Pr[\mathbf{E}_3]| = \text{negl}(\lambda).$$

By the definition of the games, we have $|\Pr[\mathbf{E}_4] - \Pr[\mathbf{E}_1]| = \text{Adv}_{\mathcal{A}, \mathbb{G}, \text{Samp}}^{\text{BKOALA, dist}}(\lambda)$. Let us define $\epsilon := |\Pr[\mathbf{E}_2] - \Pr[\mathbf{E}_3]|$. By triangular inequality, we have

$$\begin{aligned} \epsilon &\geq |\Pr[\mathbf{E}_4] - \Pr[\mathbf{E}_1]| - |\Pr[\mathbf{E}_1] - \Pr[\mathbf{E}_2]| - |\Pr[\mathbf{E}_3] - \Pr[\mathbf{E}_4]| \\ &\geq \text{Adv}_{\mathcal{A}, \mathbb{G}, \text{Samp}}^{\text{BKOALA, dist}}(\lambda) - \text{negl}(\lambda). \end{aligned}$$

Therefore, it suffices to prove the following lemma to finish the proof of Theorem 4.4.

Lemma 4.5. *There exists an efficient adversary \mathcal{B} that has access to the bilinear generic group oracles and $\text{Adv}_{\mathcal{B}, \mathbb{G}, \text{Samp}}^{\text{BKOALA, find}}(\lambda) \geq \epsilon/Q_{\text{zt}}$.*

Proof. We first observe that the oracle response to \mathcal{A} in **Game₂** and **Game₃** differs only when \mathcal{A} makes a zero-test query for a handle that corresponds to $x_0 + \sum_{j \in [\ell_2]} x_j R_j$ such that $x_0 + \sum_{j \in [\ell_2]} x_j R_j \neq 0$ over $\mathbb{Z}_p[R_1, \dots, R_{\ell_2}]$ and

$$x_0 + \sum_{j \in [\ell_2]} x_j \left(\sum_{i \in [\ell_1]} v_{i,j} S_i \right) = x_0 + \sum_{i \in [\ell_1]} \left(\sum_{j \in [\ell_2]} x_j v_{i,j} \right) S_i = 0$$

over $\mathbb{Z}_p[S_1, \dots, S_{\ell_1}]$. We call such a query *bad* query. We can see that \mathcal{A} makes a bad query with probability at least ϵ in **Game₂**. We observe that for a bad query, we have $x_0 = 0$, $\mathbf{xV}^T = \mathbf{0}$, and $\mathbf{x} \neq \mathbf{0}$ for $\mathbf{x} = (x_1, \dots, x_{\ell_2})$.

To prove the theorem, we further consider the following sequence of games. In the following, let us denote \mathbf{F}_x the event that \mathcal{A} makes a bad query and the challenger does not output \perp in **Game_{2,x}**.

Game_{2,1}: This is the same as **Game₂**. Without loss of generality, we assume that the challenger simulates the generic group oracles for \mathcal{A} . By definition, we have

$$\Pr[\mathbf{F}_1] \geq \epsilon.$$

Game_{2,2}: In this game, we change the previous game so that the challenger picks a random guess k^* for the first bad query as $k^* \leftarrow [Q_{\text{zt}}]$ at the beginning of the game. Furthermore, we change the game so that the challenger outputs \perp at the end of the game if the k^* -th zero-test query is not the first bad query. Since k^* is chosen uniformly at random and independent from the view of \mathcal{A} , the guess is correct with probability $1/Q_{\text{zt}}$ conditioned on \mathbf{F}_1 . Therefore, we have

$$\Pr[\mathbf{F}_2] \geq \Pr[\mathbf{F}_1]/Q_{\text{zt}}.$$

Game_{2,3}: This game is the same as the previous game except that the challenger aborts the game and outputs \perp immediately after \mathcal{A} makes the k^* -th zero-test query. Since whether \mathbf{F}_2 occurs or not is irrelevant to how the game proceeds after the k^* -th zero-test query is made by \mathcal{A} , we clearly have

$$\Pr[\mathbf{F}_3] = \Pr[\mathbf{F}_2].$$

We then construct \mathcal{B} , which acts as the challenger in **Game_{2,3}** for \mathcal{A} as follows.

\mathcal{B} takes aux as input and then chooses random $k^* \leftarrow [Q_{\text{zt}}]$. It then runs \mathcal{A} on input aux and handles for symbols R_1, \dots, R_{ℓ_2} . \mathcal{B} then answers generic oracle queries made by \mathcal{A} honestly until the k^* -th zero test query. When \mathcal{A} makes the k^* -th zero test query, \mathcal{B} extracts $(x_0, x_1, \dots, x_{\ell_2})$ such that the query corresponds to the handle of $x_0 + \sum_{j \in [\ell_2]} x_j R_j$. This is possible by keeping track of \mathcal{A} 's group operations while simulating the generic group oracles. If $x_0 \neq 0$, \mathcal{B} aborts and outputs \perp . Otherwise, it outputs the vector $\mathbf{x} = (x_1, \dots, x_{\ell_2})$.

Since \mathcal{B} perfectly simulates **Game_{2,3}** and thus the probability that \mathcal{B} outputs \mathbf{x} such that $\mathbf{x}\mathbf{V}^\top = 0$ and $\mathbf{x} \neq \mathbf{0}$ is $\Pr[\mathbf{F}_3] = \epsilon/Q_{\text{zt}}$. This completes the proof of Lemma 4.5.

This completes the proof of Theorem 4.4.

4.2 Our New Assumption wKOALA

Here, we introduce our new assumption that we call wKOALA (for “weak” KOALA) that will be used to prove the security of our ABE scheme in Section 5. The assumption essentially says that for a sampler that outputs a set of vectors such that the vectors are individually pseudorandom but mutually correlated, it holds that the vectors appear mutually pseudorandom when they are lifted to the exponent and randomized by vector-wise randomness. We require that this hold even in the presence of random auxiliary input as is assumed for the case of the bilinear KOALA assumption.

Definition 4.6. Let $\text{Samp} = \{\text{Samp}_\lambda\}_\lambda$ be an efficient sampling algorithm that takes as input an integer p and a string aux and outputs a set of vectors $\{\mathbf{u}^{(j)} \in \mathbb{Z}_p^m\}_{j \in [t]}$. For an efficient adversary $\mathcal{A} = \{\mathcal{A}_\lambda\}_\lambda$ and $i := i(\lambda) \in \mathbb{N}$, let us define

$$\text{Adv}_{\mathcal{A}, \mathbb{G}, \text{Samp}, i}^{\text{wKOALA, single}}(\lambda) := |\Pr[\mathcal{A}_\lambda(\mathbb{G}, \text{aux}, \mathbf{u}^{(i)}) \rightarrow 1] - \Pr[\mathcal{A}_\lambda(\mathbb{G}, \text{aux}, \mathbf{v}) \rightarrow 1]|, \quad (4.1)$$

where the probabilities are taken over the choice of uniformly random aux , $\mathbb{G} = (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, [1]_1, [1]_2) \leftarrow \text{GroupGen}(1^\lambda)$, $\{\mathbf{u}^{(j)}\}_{j \in [t]} \leftarrow \text{Samp}_\lambda(p, \text{aux})$, $\mathbf{v} \leftarrow \mathbb{Z}_p^m$, and the coin of \mathcal{A}_λ . In the above, we set $\mathbf{u}^{(i)} := \mathbf{v}$ if $i > t$. Furthermore, for an efficient adversary $\mathcal{B} = \{\mathcal{B}_\lambda\}$, we define

$$\text{Adv}_{\mathcal{B}, \mathbb{G}, \text{Samp}}^{\text{wKOALA, multi}}(\lambda) := \left| \Pr \left[\mathcal{B}_\lambda \left(\begin{array}{c} \mathbb{G}, \text{aux}, \\ \{[\gamma^{(j)}]_2, [\gamma^{(j)} \mathbf{u}^{(j)}]_2\}_{j \in [t]} \end{array} \right) \rightarrow 1 \right] - \Pr \left[\mathcal{B}_\lambda \left(\begin{array}{c} \mathbb{G}, \text{aux}, \\ \{[\gamma^{(j)}]_2, [\mathbf{v}^{(j)}]_2\}_{j \in [t]} \end{array} \right) \rightarrow 1 \right] \right|, \quad (4.2)$$

where the probabilities are taken over the choice of uniformly random aux , \mathbb{G} , $\{\mathbf{u}^{(j)}\}_{j \in [t]} \leftarrow \text{Samp}_\lambda(p, \text{aux})$, $\gamma^{(j)} \leftarrow \mathbb{Z}_p$, $\mathbf{v}^{(j)} \leftarrow \mathbb{Z}_p^m$ for $j \in [t]$, and the coin of \mathcal{B}_λ . We say that wKOALA holds with respect to GroupGen if for any efficient sampler Samp such that $\text{Adv}_{\mathcal{A}, \mathbb{G}, \text{Samp}, i}^{\text{wKOALA, single}}(\lambda)$ is negligible for any efficient adversary \mathcal{A} and $i(\lambda)$, $\text{Adv}_{\mathcal{B}, \mathbb{G}, \text{Samp}}^{\text{wKOALA, multi}}(\lambda)$ is also negligible for any efficient adversary \mathcal{B} .

The following theorem shows that wKOALA is in fact implied by the bilinear KOALA assumption.

Theorem 4.7. If the bilinear KOALA assumption holds with respect to GroupGen , so does wKOALA.

Proof. For the sake of contradiction, let us assume that wKOALA does not hold with respect to GroupGen , but the bilinear KOALA assumption holds with respect to GroupGen . The former assumption implies that there exists an efficient sampler Samp such that $\text{Adv}_{\mathcal{A}, \mathbb{G}, \text{Samp}, i}^{\text{wKOALA, single}}(\lambda)$ is negligible for any efficient adversary \mathcal{A} and $i = i(\lambda)$, but there exists an efficient adversary \mathcal{B} such that $\epsilon(\lambda) := \text{Adv}_{\mathcal{B}, \mathbb{G}, \text{Samp}}^{\text{wKOALA, multi}}(\lambda)$ is non-negligible.

We then consider another sampler Samp' that takes as input p and aux and outputs matrix \mathbf{V} defined as

$$\mathbf{V} = \begin{bmatrix} 1 & \mathbf{u}^{(1)} & & & \\ & & 1 & \mathbf{u}^{(2)} & \\ & & & & \ddots \\ & & & & & 1 & \mathbf{u}^{(t)} \end{bmatrix} \in \mathbb{Z}_p^{t \times (1+m)t},$$

where $\{\mathbf{u}^{(j)}\}_{j \in [t]} \leftarrow \text{Samp}(p, \text{aux})$. For this sampler Samp' , we have

$$\text{Adv}_{\mathcal{B}, \mathbb{G}, \text{Samp}'}^{\text{BKOALA, dist}}(\lambda) = \text{Adv}_{\mathcal{B}, \mathbb{G}, \text{Samp}}^{\text{wKOALA, multi}}(\lambda) = \epsilon(\lambda),$$

which follows from the definition of $\text{Adv}_{\mathcal{B}, \mathbb{G}, \text{Samp}'}^{\text{BKOALA, dist}}(\lambda)$ (See Definition 4.1). This further implies that there exists another adversary \mathcal{B}' and polynomial function $Q(\lambda)$ such that

$$\text{Adv}_{\mathcal{B}', \mathbb{G}, \text{Samp}'}^{\text{BKOALA, find}}(\lambda) \geq \epsilon(\lambda)/Q(\lambda) - \text{negl}(\lambda)$$

from the bilinear KOALA assumption. By the definition, \mathcal{B}' takes as input aux and \mathbb{G} and outputs a vector $\mathbf{x} \in \mathbb{Z}_p^{(1+m)t}$ such that $\mathbf{x}\mathbf{V}^\top = \mathbf{0}$ and $\mathbf{x} \neq \mathbf{0}$ with probability $\epsilon'(\lambda); = \epsilon(\lambda)/Q(\lambda) - \text{negl}(\lambda)$. Let us denote

$$\mathbf{x} = (x^{(1)}, \mathbf{x}^{(1)}, x^{(2)}, \mathbf{x}^{(2)}, \dots, x^{(t)}, \mathbf{x}^{(t)})$$

where $x^{(i)} \in \mathbb{Z}_p$ and $\mathbf{x}^{(i)} \in \mathbb{Z}_p^m$ for $i \in [t]$. Then, for such vector \mathbf{x} , we have

$$x^{(i)} + \langle \mathbf{x}^{(i)}, \mathbf{u}^{(i)} \rangle = 0 \pmod p \quad \text{and} \quad (x^{(i)}, \mathbf{x}^{(i)}) \neq \mathbf{0} \pmod p$$

for some $i \in [t]$ by the structure of \mathbf{V} . This further implies that there exists fixed $i^* = i^*(\lambda)$ such that $x^{(i^*)} + \langle \mathbf{x}^{(i^*)}, \mathbf{u}^{(i^*)} \rangle = 0 \pmod p$ and $(x^{(i^*)}, \mathbf{x}^{(i^*)}) \neq \mathbf{0} \pmod p$ hold with probability at least ϵ'/t .

We then use \mathcal{B}' to construct an adversary \mathcal{A} such that $\text{Adv}_{\mathcal{A}, \mathbb{G}, \text{Samp}, i^*}^{\text{wKOALA, single}}(\lambda)$ is non-negligible, which contradicts our assumption. \mathcal{A} takes as input the group description \mathbb{G} , auxiliary information aux and a vector \mathbf{v} , which is either $\mathbf{v} \leftarrow \mathbb{Z}_p^m$ or $\mathbf{v} = \mathbf{u}^{(i^*)}$ with $\{\mathbf{u}^{(j)}\}_{j \in [t]} \leftarrow \text{Samp}(p, \text{aux})$. Then, \mathcal{A} runs \mathcal{B}' on input aux and \mathbb{G} . If \mathcal{B}' outputs something outside of \mathbb{Z}_p^m , \mathcal{A} outputs 0. Otherwise, let $\mathbf{x} \in \mathbb{Z}_p^m$ be the output by \mathcal{B}' . If $(x^{(i^*)}, \mathbf{x}^{(i^*)}) = \mathbf{0}$, \mathcal{A} outputs 0. Otherwise, \mathcal{A} checks whether

$$x^{(i^*)} + \langle \mathbf{x}^{(i^*)}, \mathbf{v} \rangle \stackrel{?}{=} 0. \tag{4.3}$$

It outputs 1 if it holds and 0 otherwise.

We evaluate the probability that \mathcal{A} outputs 1. There are two cases to consider.

- If $\mathbf{v} = \mathbf{u}^{(i^*)}$, \mathcal{B}' outputs non-zero vector $(x^{(i^*)}, \mathbf{x}^{(i^*)})$ satisfying Eq. (4.3) with probability at least ϵ'/t . \mathcal{A} outputs 1 with the same probability.
- If \mathbf{v} is chosen uniformly at random from \mathbb{Z}_p^m , Eq. (4.3) holds for $(x^{(i^*)}, \mathbf{x}^{(i^*)})$ output by \mathcal{B}' with probability at most $1/p$ unless $(x^{(i^*)}, \mathbf{x}^{(i^*)}) = \mathbf{0} \pmod p$, since \mathbf{v} is information theoretically hidden from \mathcal{B}' . Since \mathcal{A} outputs 1 only when $(x^{(i^*)}, \mathbf{x}^{(i^*)}) \neq \mathbf{0}$ and Eq. (4.3) holds, the probability that \mathcal{A} outputs 1 is at most $1/p$.

We finally observe that

$$\begin{aligned} \text{Adv}_{\mathcal{A}, \mathbb{G}, \text{Samp}, i^*}^{\text{wKOALA, single}}(\lambda) &= |\Pr[\mathcal{A}(\mathbb{G}, \text{aux}, \mathbf{u}^{(i^*)}) \rightarrow 1] - \Pr[\mathcal{A}(\mathbb{G}, \text{aux}, \mathbf{v}) \rightarrow 1]| \\ &\geq \epsilon'/t - 1/p \\ &\geq \epsilon/tQ - \text{negl}, \end{aligned}$$

where the probabilities are taken over the choice of $\mathbb{G} \leftarrow \text{GroupGen}(1^\lambda)$, $\mathbf{v} \leftarrow \mathbb{Z}_p^m$, random aux , $\{\mathbf{u}^{(j)}\}_{j \in [t]} \leftarrow \text{Samp}(p, \text{aux})$ and the internal coin of \mathcal{A} . Since $\epsilon/tQ - \text{negl}$ is non-negligible, this contradicts our initial assumption. This completes the proof of Theorem 4.2.

5 Our CP-ABE Scheme

In this section, we provide our construction of CP-ABE scheme for NC_1 whose sizes of the parameters are independent from the size of the circuits supported by the scheme and only dependent on the input length and depth of the circuits. This efficiency property is not satisfied by most of the existing schemes except for [?, ?]. Unlike [?, ?], we provide the security proof in the standard model. We then show that the CP-ABE scheme can be used to construct BE with optimal efficiency. This provides the first optimal BE scheme whose security is proven in the standard model.

5.1 Construction

Here, we provide our construction of CP-ABE scheme that supports the circuit class $\mathcal{C}_{\ell,d} = \{\mathcal{C}_{\lambda,\ell(\lambda),d(\lambda)}\}_\lambda$, which is a set of all circuits with input length $\ell(\lambda)$ and depth at most $d(\lambda)$ with arbitrary $\ell(\lambda) = \text{poly}(\lambda)$ and $d(\lambda) = O(\log \lambda)$. For our construction, we will use public key slotted IPFE scheme $\text{IPFE} = (\text{IPFE.Setup}, \text{IPFE.KeyGen}, \text{IPFE.Enc}, \text{IPFE.SlotEnc}, \text{IPFE.Dec})$, which is proposed by Lin and Luo [?], which is secure under the MDDH assumption, and a secret sharing scheme $\text{SS} = (\text{SS.Setup}, \text{SS.Share}, \text{SS.Recon})$ for $\mathcal{C}_{\ell(\lambda),d(\lambda)}$ that is provided in Section 3.

$\text{ABE.Setup}(1^\lambda)$: On input 1^λ , the setup algorithm proceeds as follows.

1. Run $\mathbb{G} = (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, [1]_1, [1]_2) \leftarrow \text{GroupGen}(1^\lambda)$. Note that p and λ specify the parameter $m := m(\lambda)$ (See syntax of secret sharing scheme in Section 3).
2. Run $\text{pp} \leftarrow \text{SS.Setup}(1^\lambda, p)$.
3. Run $(\text{IPFE.mpk}_{i,b,j}, \text{IPFE.msk}_{i,b,j}) \leftarrow \text{IPFE.Setup}(1^\lambda, \{1\}, \{2\})$ for $i \in [\ell]$, $b \in \{0, 1\}$, and $j \in [m]$. Note that here we generate slotted IPFE instances whose the first entry is for public slot and the second entry is for private slot.
4. Output $\text{ABE.mpk} = (\mathbb{G}, \text{pp}, \{\text{IPFE.mpk}_{i,b,j}\}_{i \in [\ell], b \in \{0,1\}, j \in [m]})$ and $\text{ABE.msk} = \{\text{IPFE.msk}_{i,b,j}\}_{i \in [\ell], b \in \{0,1\}, j \in [m]}$.

$\text{ABE.KeyGen}(\text{ABE.mpk}, \text{ABE.msk}, x)$: The key generation algorithm takes as input the master public key ABE.mpk , the master secret key ABE.msk , and an attribute $x \in \{0, 1\}^\ell$ and proceeds as follows.

1. Let $x_1 \cdots x_\ell \in \{0, 1\}^\ell$ be the binary representation of $x \in \{0, 1\}^\ell$.
2. Pick $\gamma \leftarrow \mathbb{Z}_p$ and compute $[\gamma]_T$.
3. Sample $\text{IPFE.sk}_{i,x_i,j} \leftarrow \text{IPFE.KeyGen}(\text{IPFE.msk}_{i,x_i,j}, [(\gamma, 0)]_2)$ for all $i \in [\ell]$ and $j \in [m]$.
4. Output $\text{ABE.sk} = ([\gamma]_T, \{\text{IPFE.sk}_{i,x_i,j}\}_{i \in [\ell], j \in [m]})$.

$\text{ABE.Enc}(\text{ABE.mpk}, C, \mu)$: The encryption algorithm takes as input the master public key ABE.mpk , a circuit C , and the message μ and proceeds as follows.

1. Run $\{\text{share}_{i,b}\}_{i \in [\ell], b \in \{0,1\}} \leftarrow \text{SS.Share}(\text{pp}, C, \mu)$.
2. Parse each $\text{share}_{i,b}$ as $\text{share}_{i,b} = \{\text{share}_{i,b,j} \in \mathbb{Z}_p\}_{j \in [m]}$.
3. Run $\text{IPFE.ct}_{i,b,j} \leftarrow \text{IPFE.SlotEnc}(\text{IPFE.mpk}_{i,b,j}, [\text{share}_{i,b,j}]_1)$ for $i \in [\ell]$, $b \in \{0,1\}$, $j \in [m]$.
4. Output $\text{ABE.ct} = \{\text{IPFE.ct}_{i,b,j}\}_{i \in [\ell], b \in \{0,1\}, j \in [m]}$.

$\text{ABE.Dec}(\text{ABE.mpk}, \text{ABE.sk}, x, \text{ABE.ct}, C)$: The decryption algorithm takes as input the master public key ABE.mpk , the secret key ABE.sk along with x , the ciphertext ABE.ct along with C and does the following:

1. Parse the ciphertext as $\text{ABE.ct} \rightarrow \{\text{IPFE.ct}_{i,b,j}\}_{i \in [\ell], b \in \{0,1\}, j \in [m]}$ and the secret key as $\text{ABE.sk} \rightarrow ([\gamma]_T, \{\text{IPFE.sk}_{i,x_i,j}\}_{i \in [\ell], j \in [m]})$.
2. Run $\text{IPFE.Dec}(\text{IPFE.sk}_{i,x_i,j}, \text{IPFE.ct}_{i,x_i,j}) \rightarrow [d_{i,j}]_T$ for $i \in [\ell]$ and $j \in [m]$.
3. Run $\text{SS.FindCoef}(\text{pp}, C, x) \rightarrow \{a_{i,j} \in \mathbb{Z}_p\}_{i \in [\ell], j \in [m]}$.
4. Compute $[d']_T = [\sum_{i \in [\ell], j \in [m]} a_{i,j} d_{i,j}]_T$ from $\{[d_{i,j}]_T\}_{i,j}$ and $\{a_{i,j}\}_{i,j}$.
5. Find $e \in [-B, B]$ and $\mu \in \{0,1\}$ such that $[d']_T = [\gamma(\mu[p/2] + e)]_T$ by Brute-force search using $[\gamma]_T$. If such a pair does not exist, output \perp . Otherwise, output μ .

Correctness. To show correctness of the scheme, we first observe that $d_{i,j} = \gamma \cdot \text{share}_{i,x_i,j}$ by the correctness of IPFE. We then observe that $d' = \sum_{i \in [\ell], j \in [m]} \gamma a_{i,j} \text{share}_{i,x_i,j} = \gamma(\mu \cdot [p/2] + e)$ for some $e \in [-B, B]$ by the almost linear reconstruction property (Definition 3.4) of ABE. Since B is polynomially bounded by Theorem 3.5, the last step in the decryption algorithm works in polynomial time and recovers the message μ .

Efficiency. The master public key of the ABE consists of $O(\ell m)$ master public keys of the IPFE and the public parameter pp of the secret sharing scheme. The ciphertext and secret key of the ABE contain $O(\ell m)$ ciphertexts and secret keys of the IPFE, respectively. By the efficiency of the secret sharing scheme, m and the size of pp are bounded by $\text{poly}(\lambda)$. Furthermore, each instance of IPFE only deals with constant dimension of vectors, which means that sizes of all the parameters of each instance of the IPFE scheme are bounded by $\text{poly}(\lambda)$. Therefore, we can see that sizes of all the parameters in the ABE scheme are bounded by $\text{poly}(\lambda, \ell)$.

5.2 Security Proof

Here, we prove the security of our ABE scheme in Sec. 5.1. Before doing so, we prove the following lemma.

Lemma 5.1. *Let $C \in \mathcal{C}_{\lambda, \ell(\lambda), d(\lambda)}$ be a circuit and $X \subseteq \{0,1\}^{\ell(\lambda)}$ be a set of strings that satisfies $C(x) = 0$ for all $x \in X$, and $\text{SS} = (\text{SS.Setup}, \text{SS.Share}, \text{SS.Recon})$ be a secure secret sharing scheme for this circuit class as per Definition 3.2. Then,*

for any efficient adversary $\mathcal{A} = \{\mathcal{A}_\lambda\}$, we have

$$\begin{aligned} & \Pr \left[\mathcal{A}_\lambda \left(\left\{ [\gamma^{(x)}]_2, \{\{\gamma^{(x)} \text{share}_{i,x_i}\}_2\}_{i \in [\ell]}\}_{x \in X} \right\}^{\mathbb{G}, \text{pp}} \right) \rightarrow 1 \right] \\ & - \Pr \left[\mathcal{A}_\lambda \left(\left\{ [\gamma^{(x)}]_2, \{\{\gamma^{(x)} \mathbf{w}_i^{(x)}\}_2\}_{i \in [\ell]}\}_{x \in X} \right\}^{\mathbb{G}, \text{pp}} \right) \rightarrow 1 \right] = \text{negl}(\lambda), \end{aligned} \quad (5.1)$$

under the bilinear KOALA assumption, where the probabilities are taken over the choice of $\mathbb{G} = (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, [1]_1, [1]_2) \leftarrow \text{GroupGen}(1^\lambda)$, $\gamma^{(x)} \leftarrow \mathbb{Z}_p$ and $\mathbf{w}_i^{(x)} \leftarrow \mathbb{Z}_p^m$ for $x \in X$, $\text{pp} \leftarrow \text{SS.Setup}(1^\lambda, p)$, and $\{\text{share}_{i,b}\}_{i \in [\ell], b \in \{0,1\}} \leftarrow \text{SS.Share}(\text{pp}, C, \mu)$.

Proof. To prove the theorem, we set **Samp** to be an algorithm that takes as input an integer p and a random string aux , sets $\text{pp} := \text{aux}$, runs $\{\text{share}_{i,b}\}_{i \in [\ell], b \in \{0,1\}} \leftarrow \text{SS.Share}(\text{pp}, C, \mu)$, and outputs

$$\left\{ \{\text{share}_{i,x_i}\}_{i \in [\ell]} \right\}_{x \in X}.$$

Then, we argue that $\text{Adv}_{\mathcal{A}, \mathbb{G}, \text{Samp}, i}^{\text{wKOALA, single}}(\lambda)$ defined in Eq. (4.1) is negligible for any i and any efficient adversary \mathcal{A} . To show this, we first fix i and \mathcal{A} and observe that the quantity is negligible if we replace aux that is input to \mathcal{A} with pp output by $\text{SS.Setup}(1^\lambda, p)$, which follows from the security of **SS** (Definition 3.2) and from the fact that $C(x) = 0$ holds for all $x \in X$. We then observe that the adversary will not notice even if we replace pp with aux since the distributions of them are statistically close by the random public parameter property (Definition 3.3) of **SS**. We therefore have $\text{Adv}_{\mathcal{A}, \mathbb{G}, \text{Samp}, i}^{\text{wKOALA, single}}(\lambda) = \text{negl}(\lambda)$. This implies that $\text{Adv}_{\mathcal{B}, \mathbb{G}, \text{Samp}}^{\text{wKOALA, multi}}(\lambda)$ defined in Eq. (4.2) is negligible for any efficient adversary \mathcal{B} by wKOALA , which is implied by the bilinear KOALA assumption. Finally, by replacing aux with pp output by $\text{SS.Setup}(1^\lambda, p)$ in Eq. (4.2), we have that Eq. (5.1) is negligible for any efficient adversary as desired.

The next theorem establishes the security of our ABE scheme.

Theorem 5.2. *Our ABE scheme satisfies very selective security under the MDDH assumption, the bilinear KOALA assumption, and the LWE assumption.*

Proof. To prove the theorem, we fix a PPT adversary $\mathcal{A} = \{\mathcal{A}_\lambda\}$. Without loss of generality, we make some simplifying assumptions on \mathcal{A} . First, we assume that \mathcal{A} always chooses $(\mu_0, \mu_1) = (0, 1)$ as its target message at the challenge phase. This can be assumed without loss of generality since our scheme is a single-bit scheme. Second, we assume that the adversary does not make key queries for the same attribute x twice. The adversary that makes key queries for the same attribute more than once can be dealt with by making the key generation algorithm deterministic by changing the scheme so that it derives randomness using a PRF, which can be instantiated from any one-way functions. Third, we assume that the adversary chooses fixed challenge attribute $C = \{C_\lambda\}$

and key queries $X = \{X_\lambda\}$. This can be assumed without loss of generality because they are chosen by the adversary at the beginning of the game only depending on the security parameter and we can derandomize \mathcal{A} by choosing the best randomness that maximizes the advantage of \mathcal{A} .

In order to prove the security, we consider following sequence of games. Let us denote the event that \mathcal{A} outputs correct guess for b at the end of \mathbf{Game}_x as E_x .

Game₁: This is the real very selective security game. To fix the notation and for the sake of concreteness, we briefly describe the game here. At the beginning of the game, the adversary chooses C and $X \subseteq \{0, 1\}^\ell$. Then, the challenger first chooses the master public key ABE.mpk and the master secret key ABE.msk of the ABE scheme. It then generates the challenge ciphertext as follows. It first chooses the message $\beta \leftarrow \{0, 1\}$ and runs $\{\text{share}_{i,b}\}_{i \in [\ell], b \in \{0,1\}} \leftarrow \text{SS.Share}(\text{pp}, C, \beta)$. It then runs

$$\text{IPFE.ct}_{i,b,j} \leftarrow \text{IPFE.SlotEnc}(\text{IPFE.mpk}_{i,b,j}, [\text{share}_{i,b,j}]_1)$$

for $i \in [\ell]$, $b \in \{0, 1\}$, $j \in [m]$ and sets the challenge ciphertext to be $\text{ABE.ct} = \{\text{IPFE.ct}_{i,b,j}\}_{i \in [\ell], b \in \{0,1\}, j \in [m]}$. It also generates secret key ABE.sk_x for all $x \in X$ as follows. It first generates $\gamma^{(x)} \leftarrow \mathbb{Z}_p$ and

$$\text{IPFE.sk}_{i,x_i,j} \leftarrow \text{IPFE.KeyGen}(\text{IPFE.msk}_{i,x_i,j}, [(\gamma^{(x)}, 0)]_2)$$

for all $i \in [\ell]$ and $j \in [m]$. It then sets $\text{ABE.sk}_x = ([\gamma]_T, \{\text{IPFE.sk}_{i,x_i,j}\}_{i \in [\ell], j \in [m]})$. Finally, the challenger returns ABE.mpk , ABE.ct , and $\{\text{ABE.sk}_x\}_{x \in X}$ to \mathcal{A} , which then outputs a bit $\hat{\beta}$ as a guess for β . By definition, the advantage of \mathcal{A} against the scheme is

$$\left| \Pr[E_1] - \frac{1}{2} \right|.$$

Game₂: In this game, we change the game so that the challenger generates the challenge ciphertext as follows. It first chooses the message $\beta \leftarrow \{0, 1\}$ and runs $\{\text{share}_{i,b}\}_{i \in [\ell], b \in \{0,1\}} \leftarrow \text{SS.Share}(\text{pp}, C, \beta)$. It then generates

$$\text{IPFE.ct}_{i,b,j} \leftarrow \text{IPFE.Enc}(\text{IPFE.msk}_{i,b,j}, [(\text{share}_{i,b,j}, 0)]_1)$$

for $i \in [\ell]$, $b \in \{0, 1\}$, $j \in [m]$ and sets the challenge ciphertext to be $\text{ABE.ct} = \{\text{IPFE.ct}_{i,b,j}\}_{i \in [\ell], b \in \{0,1\}, j \in [m]}$. By the slot-mode correctness of the IPFE, we have

$$\Pr[E_1] = \Pr[E_2].$$

Game₃: In this game, we change the way the challenge ciphertext and the secret keys are generated. In this game, the challenger generates the challenge ciphertext as follows. It first chooses the message $\beta \leftarrow \{0, 1\}$ and runs $\{\text{share}_{i,b}\}_{i \in [\ell], b \in \{0,1\}} \leftarrow \text{SS.Share}(\text{pp}, C, \beta)$. However, it ignores these values and generates

$$\text{IPFE.ct}_{i,b,j} \leftarrow \text{IPFE.Enc}(\text{IPFE.msk}_{i,b,j}, [(0, 1)]_1)$$

for $i \in [\ell]$, $b \in \{0,1\}$, $j \in [m]$ and sets the challenge ciphertext to be $\text{ABE.ct} = \{\text{IPFE.ct}_{i,b,j}\}_{i \in [\ell], b \in \{0,1\}, j \in [m]}$. We also change the way the challenger generates the secret keys as follows. For $x \in X$, it first generates $\gamma^{(x)} \leftarrow \mathbb{Z}_p$ and

$$\text{IPFE.sk}_{i,x_i,j} \leftarrow \text{IPFE.KeyGen}(\text{IPFE.msk}_{i,x_i,j}, [(\gamma^{(x)}, \gamma^{(x)} \text{share}_{i,x_i,j})]_2)$$

for all $i \in [\ell]$ and $j \in [m]$, where we use $\text{share}_{i,x_i,j}$ that is generated when creating the challenge ciphertext. It then sets $\text{ABE.sk}_x = ([\gamma^{(x)}]_T, \{\text{IPFE.sk}_{i,x_i,j}\}_{i \in [\ell], j \in [m]})$.

We can observe that for each instance of IPFE, the inner products between the vector that is encoded in the ciphertext and the vectors encoded in the secret keys are unchanged. Furthermore, the values in the public slots of the vectors encoded in the secret keys are unchanged. Therefore, this game is indistinguishable from the previous game for \mathcal{A} by the security of the IPFE, which follows from the MDDH assumption. We therefore have

$$|\Pr[\mathbf{E}_2] - \Pr[\mathbf{E}_3]| = \text{negl}(\lambda).$$

Game₄: In this game, we further change the way the secret keys are generated as follows. The challenger first generates $\gamma^{(x)} \leftarrow \mathbb{Z}_p$ and chooses $\mathbf{w}_i^{(x)} \leftarrow \mathbb{Z}_p^m$ for $i \in [\ell]$. It then generates

$$\text{IPFE.sk}_{i,x_i,j} \leftarrow \text{IPFE.KeyGen}(\text{IPFE.msk}_{i,x_i,j}, [(\gamma^{(x)}, w_{i,j}^{(x)})]_2)$$

for all $i \in [\ell]$ and $j \in [m]$, where $w_{i,j}^{(x)}$ is the j -th entry of the vector $\mathbf{w}_i^{(x)}$, and sets $\text{ABE.sk}_x = ([\gamma]_T, \{\text{IPFE.sk}_{i,x_i,j}\}_{i \in [\ell], j \in [m]})$.

We claim that this game is indistinguishable from the above game. To see this, let us assume that \mathcal{A} distinguishes the games with non-negligible advantage for the sake of contradiction. Then, for $C = \{C_\lambda\}$ and $X = \{X_\lambda\}$ chosen by \mathcal{A} , we can construct another adversary $\mathcal{B} = \{\mathcal{B}_\lambda\}$ that distinguishes two distributions in Eq. (5.1) with non-negligible advantage as follows.

\mathcal{B} takes as input \mathbb{G} , pp , $\left\{ [\gamma^{(x)}]_2, \{[\mathbf{w}_i^{(x)}]_2\}_{i \in [\ell]} \right\}_{x \in X}$, where $\mathbf{w}_i^{(x)}$ is either random or $\mathbf{w}_i^{(x)} = \gamma^{(x)} \text{share}_{i,x_i}$. It chooses $(\text{IPFE.mpk}_{i,b,j}, \text{IPFE.msk}_{i,b,j}) \leftarrow \text{IPFE.Setup}(1^\lambda, \{1\}, \{2\})$ for all i, b , and j and sets the ABE.mpk and ABE.msk accordingly. It generates the challenge ciphertext using ABE.msk . It also generates a secret key for x using $[\gamma^{(x)}]_2, \{[\mathbf{w}_i^{(x)}]_2\}_{i \in [\ell]}$ and ABE.msk . In particular, the syntax of the slotted IPFE allows us to generate the secret key for the vector $[(\gamma^{(x)}, w_{i,j}^{(x)})]_2$ without knowing the corresponding discrete logarithm, which \mathcal{B} does not know. \mathcal{B} then inputs ABE.mpk , ABE.ct , and $\{\text{ABE.sk}_x\}$ to \mathcal{A} and outputs what \mathcal{A} outputs.

Clearly, \mathcal{B} simulates **Game₃** for \mathcal{A} if $\mathbf{w}_i^{(x)} = \gamma^{(x)} \text{share}_{i,x_i}$ and **Game₄** if $\mathbf{w}_i^{(x)}$ is random. Thus, \mathcal{B} can distinguish the two distributions with the same

advantage as \mathcal{A} . This contradicts the bilinear KOALA assumption and the security of the secret sharing scheme by Lemma 5.1, where the latter follows from the LWE assumption by Theorem 3.5. Therefore, we have

$$|\Pr[\mathbf{E}_3] - \Pr[\mathbf{E}_4]| = \text{negl}(\lambda).$$

We can easily observe that the view of \mathcal{A} in **Game**₄ is independent from β and $\Pr[\mathbf{E}_4] = 1/2$. Therefore, we have

$$\left| \Pr[\mathbf{E}_1] - \frac{1}{2} \right| \leq \sum_{i \in [3]} |\Pr[\mathbf{E}_i] - \Pr[\mathbf{E}_{i+1}]| + \left| \Pr[\mathbf{E}_4] - \frac{1}{2} \right| \leq \text{negl}(\lambda)$$

as desired. This completes the proof of Theorem 5.2.

5.3 Implication to Broadcast Encryption

Here, we show that our CP-ABE scheme implies BE by restricting the circuit class of the scheme to be some specific one as was observed in [?, ?]. Let us consider the following circuit class \mathcal{F}_{BE} :

$$\mathcal{F}_{\text{BE}} = \left\{ F_S : \{0, 1\}^{\lceil \log N \rceil} \rightarrow \{0, 1\} \right\}_{S \subseteq [N]} \quad \text{where} \quad F_S(i) = \begin{cases} 1 & \text{if } i \in S \\ 0 & \text{if } i \notin S \end{cases}.$$

Here, we identify a user index $i \in [N]$ and elements in S with binary strings in $\{0, 1\}^{\lceil \log N \rceil}$ by a natural bijection map between $\{0, 1\}^{\lceil \log N \rceil}$ and $[2^{\lceil \log N \rceil}] \supseteq [N]$. Since the depth of F_S affects the efficiency of the DBE scheme, we want F_S to be as shallow as possible. For this purpose, we compute F_S by first computing $b_j := (i \stackrel{?}{=} j)$ for all $j \in S$ in parallel and then computing $\bigvee_{j \in S} b_j$. The first step can be implemented with depth $O(\log \log N)$ and the second step with $O(\log N)$. This allows us to implement F_S with depth $O(\log |S|) \leq O(\log N)$. Therefore, our CP-ABE scheme indeed supports this circuit class. Furthermore, by the definition of F_S , one can see that this CP-ABE scheme implements the functionality of BE. The obtained BE scheme has optimal efficiency in the sense that the size of the master public key, secret key, and the ciphertext is bounded by $\text{poly}(\lambda, \ell, d) = \text{poly}(\log N, \lambda) = \text{poly}(\lambda)$, which is independent of the number of users N in the system. The scheme satisfies very selective security since so is the underlying CP-ABE. We note that our scheme indeed satisfies selective security since very selective security is equivalent to selective security in the setting of BE (See Remark 2.7).

Acknowledgement. We thank the anonymous reviewers of TCC 2020 for their helpful comments. Dr. Agrawal is supported by the DST ‘‘Swarnajayanti’’ fellowship, an Indo-French CEFIPRA project and an ‘‘Indo-Israel’’ ISF-UGC project. Daniel Wichs is supported by NSF grants CNS-1314722, CNS-1413964, CNS-1750795, and the Alfred P. Sloan Research Fellowship. Shota Yamada is supported by JST CREST Grant Number JPMJCR19F6 and JSPS KAKENHI Grant Number 19H01109. This work was done in part while the first and the third authors were visiting the Simons Institute for the Theory of Computing.