

Optimal Broadcast Encryption from Pairings and LWE

Shweta Agrawal¹ and Shota Yamada²

¹ IIT Madras, Chennai, India

shweta.a@cse.iitm.ac.in

² National Institute of Advanced Industrial Science and Technology (AIST), Tokyo, Japan

yamada-shota@aist.go.jp

Abstract. Boneh, Waters and Zhandry (CRYPTO 2014) used multilinear maps to provide a solution to the long-standing problem of public-key broadcast encryption (BE) where all parameters in the system are small. In this work, we improve their result by providing a solution that uses only *bilinear* maps and Learning With Errors (LWE). Our scheme is fully collusion-resistant against any number of colluders, and can be generalized to an identity-based broadcast system with short parameters. Thus, we reclaim the problem of optimal broadcast encryption from the land of “Obfustopia”.

Our main technical contribution is a ciphertext policy attribute based encryption (CP-ABE) scheme which achieves special efficiency properties – its ciphertext size, secret key size, and public key size are all independent of the size of the circuits supported by the scheme. We show that this special CP-ABE scheme implies BE with optimal parameters; but it may also be of independent interest. Our constructions rely on a novel interplay of bilinear maps and LWE, and are proven secure in the generic group model.

1 Introduction

Broadcast Encryption (BE) [FN94] enables a sender to encrypt a message for a subset of users who are listening on a broadcast channel. In more detail, in a BE system, a sender can encrypt to any set S of its choice, and any user in S can decrypt the broadcast using its secret key. The system is said to be fully collusion resistant if no collection of users outside S can learn anything about the plaintext.

Introduced in a seminal work by Fiat and Naor [FN94], the primitive of broadcast encryption has received significant attention, with diverse constructions achieving different tradeoffs in the sizes of ciphertext, secret key and public parameters. Of particular importance is the size of the ciphertext overhead: namely, the size of the ciphertext beyond what is necessary for the description of the recipient set S and the symmetric encryption of the plaintext message. A BE scheme is said to have low overhead if the ciphertext overhead depends at most logarithmically on the number of users in the system (N , say). In this work, we focus on BE systems that are public key, have low ciphertext overhead and are fully collusion resistant.

The first work to satisfy the above desiderata was by Boneh, Gentry, and Waters [BGW05], and was based on hardness assumptions on bilinear maps. This construction achieved optimal (constant) ciphertext overhead and short secret keys, but suffered from

public parameter size which is linear in the number of users N . Follow-ups based on bilinear maps improved some aspects of this construction [GW09, DPP07, GW09, Del07, SF, AL10, HWL⁺16], but could not improve the public key size. Indeed, even relying on the existence of the powerful indistinguishability obfuscation [GGH⁺13], BE with short public key remained elusive (though it achieved other remarkable properties) [BZ17].

This state of affairs was improved considerably by the work of Boneh, Waters and Zhandry [BWZ14] who provided the first construction of broadcast encryption, achieving optimal parameters including short public key, by relying on multilinear maps. This marked the first solution to a long standing open problem. However, the constructions suggested by [BWZ14] also have some limitations. In more detail, the [BWZ14] provide three broadcast encryption systems that use an $O(\log N)$ way multilinear map – this necessitates the degree of the map to be *polynomial* when N is exponential. More importantly, existing candidates of multilinear maps have been subject to many attacks [CHL⁺15, CGH⁺15, HJ15, CJL, CFL⁺, MSZ16, CLLT16, ADGM16] and their security is poorly understood. Thus, the question of broadcast encryption with optimal parameter size has so far, remained squarely in the land of “Obfustopia”.

Our Results. In this work, we reclaim broadcast encryption from Obfustopia by providing a solution that uses only *bilinear* maps and Learning With Errors (LWE). Our scheme is public key, fully collusion-resistant against any number of colluders, and can be generalized to an identity-based broadcast system with short parameters. Along the way, we provide the first ciphertext policy attribute based encryption scheme whose ciphertext size, secret key size, and public key size are all independent of the size of the circuits supported by the scheme. This construction may be of independent interest. Our constructions rely on a novel interplay between bilinear maps and LWE and are proven secure in the generic group model.

1.1 Our Techniques.

Recasting BE as CP-ABE: Our starting observation is that the question of broadcast encryption can be re-stated in terms of the notion of *ciphertext policy attribute based encryption* (CP-ABE). In a CP-ABE scheme, a ciphertext for a message m is labelled with a function (policy) f , and secret keys are labelled 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 CP-ABE, note that the ciphertext may encode a circuit F_S that checks membership of a given user index in a set of recipients S , and the attributes \mathbf{x} may encode user index in the set N . Thus, a user i can use her CP-ABE secret key to test whether i is a member of the set S encoded in the ciphertext, and recover the message m if and only if this is true. Then, a natural approach to construct BE is to leverage CP-ABE schemes. However, unsurprisingly, constructions of CP-ABE achieving optimal parameters that suffice for BE, has been elusive.

From Pairings to LWE: All known constructions of BE from standard assumptions (i.e. without relying on the existence of multilinear maps or indistinguishability obfuscation) are based on various assumptions on bilinear groups. Since the question of

optimal BE from pairings has met with little progress for over a decade, it is evidently meaningful to look at assumptions on other mathematical structures to seek a way forward. The most obvious candidate that presents itself is the versatile Learning With Errors (LWE) assumption, which has led to breakthroughs in similar primitives, notably in fully homomorphic encryption [BV11, BGV12, Bra12].

Let us then examine what is known from LWE in this context. The dual notion of key-policy ABE has met with fantastic success from LWE – the works of Gorbunov et al. [GVW13] and Boneh et al. [BGG⁺14] show how to construct KP-ABE for all circuits (on the other hand, constructions based on pairings could only support the much weaker circuit class NC₁). KP-ABE is the same as CP-ABE with the roles of circuit and attributes swapped. Additionally, the KP-ABE construction of Boneh et al. [BGG⁺14], henceforth denoted as BGG⁺, manages to encode the circuit very succinctly – in more detail, the size of the public and secret keys in the BGG⁺ construction are independent of the circuit size and depend only on the *depth* of the circuit. Additionally, the size of the ciphertext is also independent of the circuit size and depends only on input length. Since the input length for the circuit F_S that checks membership in S is an encoding of a user index, it is of size $O(\log N)$. Moreover, it is easy to check that the depth of F_S is also $O(\log N)$. Therefore, if we have a CP-ABE with analogous efficiency, namely, so that the public key size, secret key size, and ciphertext size do not depend on the size but only input length and depth of the circuit, it follows that we can obtain BE with optimal parameters.

Constructing CP-ABE from LWE: Thus, it suffices to ask whether we can have a CP-ABE scheme, denoted by cpABE, with the desired efficiency. To leverage the succinctness of the circuit encoding of BGG⁺, a naive idea is to set $\text{cpABE.CT}(F_S) = \text{BGG}^+.\text{SK}(F_S)$. Two immediate problems present themselves: i) Where to embed the message m ³, and ii) Computing $\text{BGG}^+.\text{SK}(F_S)$ requires the master secret but encryption is a public key algorithm.

To address these challenges, a first idea is to exploit the *decomposability* of BGG⁺. In more detail, decomposability means that the ciphertext for attribute \mathbf{x} and message m may be decomposed into $|\mathbf{x}| + 1$ encodings, one for each bit x_i of the attribute string and message m – these are tied together using common randomness used during their generation. Let us denote the encoding corresponding to bit x_i as ψ_{i,x_i} . Then, a natural idea is to let the encryptor sample a fresh instance of the BGG⁺ scheme, generate $\text{BGG}^+.\text{SK}(F_S)$ and encrypt each $\psi_{i,b}$ using a different public key encryption scheme, say with PKE key $\text{PK}_{i,b}$. This yields a CP-ABE with the desired efficiency, inherited directly from the succinctness of the BGG⁺ key and the decomposability of the BGG⁺ ciphertext.

Constraining the Information Leaked (Or, Back to Pairings): However, this scheme is obviously not collusion resistant: a user with keys for \mathbf{x} and $\bar{\mathbf{x}}$ can decrypt every ciphertext. To make the scheme collusion resistant, we would like to replace the naive use of public key encryption above into a more sophisticated scheme, which hides all but the output of the BGG⁺ decryption algorithm. This description bears close resemblance to a functional encryption scheme for some restricted functionality, for which we turn

³ This question is surprisingly non-trivial even in the symmetric key setting

to — pairings! In particular, we isolate the $\psi_{i,b}$ by randomizing and lifting them to the exponent of a bilinear group. The hope is that we may provide a secret key for some attribute x such that it only allows the appropriate ψ_{i,x_i} to be selected and combined so that only the output of the BGG^+ decryption is revealed, and that the randomization, which will be unique to every cpABE ciphertext and secret key pair, prevents collusion attacks.

Evaluation of NC_1 Circuit in the Exponent: Several questions arise. First, we discussed above that the circuit for checking membership in set S is in NC_1 – however, pairings are only capable of supporting at most quadratic operations. How then, do we hope to compute an NC_1 circuit in the exponent of a bilinear group? The answer lies in the specific structure of the BGG^+ evaluation algorithm, which, even for a circuit in P is *linear* in the encodings and the secret key, followed by a final *rounding* step to remove the noise. Indeed, the knowledgeable reader may observe that this very linearity of the BGG^+ evaluation procedure has been the cause of attacks in other contexts [Agr17] – what is a “curse” there is a “blessing” here! However, the rounding step remains – this is in NC_1 and clearly cannot be performed in the exponent.

An approach is to perform the linear computation (which represents the circuit F_S) in the exponent, recover the output via discrete log, and then compute the rounding in the clear. Again, it is unclear this satisfies either correctness or security. For the former, note that recovering the encoded output from the exponent requires that the output be polynomially bounded. In this case, the output is the message bit plus some noise that resulted from the homomorphic evaluation. While the noise in this context may be superpolynomial in general, we can convert our NC_1 circuit into a branching program and leverage the asymmetric noise growth for BP evaluation of BGG^+ encodings to ensure that the noise is bounded by a polynomial [GV15].

The more worrisome issue is that of security. It is well known that the noise that results from homomorphic evaluation of encodings leaks the noise in the original encodings and is often a security threat – in fact, the savvy reader may have observed that this leakage is one of the main barriers in constructing iO from standard assumptions [Agr19, AJL⁺19]. However, here we are rescued by the serendipitous fact that what we are trying to build here is a kind of attribute based encryption, not functional encryption! In more detail, the leakage caused by the noise is a security threat in the context of functional encryption, as formalized in [Agr17, Agr19, AJL⁺19] – this is because a decryptor who possesses some secret keys for a functional encryption scheme must still not be able to learn anything about the encrypted message beyond what the keys reveal. On the other hand, attribute based encryption is a much simpler “all or nothing” primitive – if the adversary possesses a single key that decrypts a ciphertext, there are no more secrets the scheme withholds from her. Hence, if the adversary has a key that lets her recover the value encoded in the exponent, the additional leakage created by the noise terms do not pose a security threat.

Preventing Mix and Match Attacks: To prevent collusions, we design the decryption algorithm so that the decryptor obtains a randomized version of the ciphertext components in the exponents as $g_T^{\delta\psi_{i,x_i}}$, where g_T is the target group, ψ_{i,x_i} are BGG^+ encodings as defined above and δ is user specific randomness. Since δ is user specific,

the attacker cannot combine partial decryption results of multiple users, preventing mix and match attacks.

Hence, it suffices to restrict our attention to the single user case. Here, we must ensure that the adversary only gets components $\{g_T^{\delta\psi_{i,x_i}}\}_i$ corresponding to the particular key \mathbf{x} issued to her, instead of all the components $\{g_T^{\delta\psi_{i,b}}\}_{i,b}$ for $b \in \{0, 1\}$. Furthermore, we must ensure that the attribute vector \mathbf{x} is processed in the correct sequence – i.e., its bits are not permuted. To prevent these attacks, we bind each entry of the ciphertext and each bit of the secret key attribute \mathbf{x} to the corresponding positions. This is possible by setting the master public key to be $\{g^{w_{i,b}}\}_{i,b}$ where $w_{i,b}$ are randomly chosen for each position i and $b \in \{0, 1\}$ and setting the secret key and ciphertext as $\{g_2^{\delta/w_{i,x_i}}\}_i$ and $\{g_1^{\psi_{i,b}w_{i,b}}\}_{i,b}$ respectively. We remark that we need to use asymmetric pairings to prevent ciphertext (respectively key) components from being paired between themselves to leak information. By tying element values to their positions, we ensure that pairing of the ciphertext and secret key components corresponding to different positions result in a term which looks like $g_T^{\delta\psi_{i,b}w_{i,b}/w_{i',b'}}$ for $(i, b) \neq (i', b')$. Now, we claim that a term of the form $g^{\delta w_{i,b}/w_{i',b'}}$ is useless to the attacker – to see this, note that in the generic group model, an attacker cannot obtain any information about a value encoded in the exponent unless she finds a non-trivial linear relation that contains that term. However, since the term $\delta w_{i,b}/w_{i',b'}$ appears only when we pair the ciphertext component with position (i, b) with the secret key component with position (i', b') and cannot appear anywhere else, it follows that it cannot appear as a term in a linear combination that results in 0 (except with negligible probability). Thus, by using $\{w_{i,b}\}_{i,b}$, we enforce that the computation follows the desired path.

Combining the above ideas, we obtain our final CP-ABE scheme. By setting the circuit class appropriately, this yields BE and even Identity Based BE (or IBBE). Please see Section 3 for the CP-ABE and Section 5 for the construction of BE and IBBE.

Security in the Generic Group Model: We prove security in the generic group model, which closely follows the intuition we explained so far. Specifically, we will show that the adversary cannot find any non-trivial linear relation among the partial decryption results of the ciphertext components. The main challenge in the security proof is that the partial decryption results obtained by using different secret keys are correlated – in more detail, they can contain terms $g_1^{\delta\psi_{i,0}}$ and $g_1^{\delta'\psi_{i,1}}$ where $\psi_{i,0}$ and $\psi_{i,1}$, if learned simultaneously, lead to a complete break of security. Simulating these in the standard model using the security of BGG^+ appears difficult.

To address the issue we first observe that the adversary cannot take a linear combination among partial decryption results obtained by two different secret keys in a meaningful way, since they are randomized by the user specific randomness introduced for preventing collusions. This implies that if the adversary manages to find a non-trivial linear relation among the partial decryption results, all the terms involved should be obtained from the same secret key. We also observe that until the point when the adversary finds the *first* non-trivial linear relation, the simulator can simulate the generic group oracles without knowing the corresponding encodings. This can be done by simply pretending that there is no non-trivial linear relation among the terms.

The above observations allow us to concentrate on the security proof for the single-key case without worrying about the partial decryption results by other keys. We can then conclude by using the security of the BGG^+ scheme. In more detail, an adversary who can find a non-trivial linear relation among the partial decryption results can be used to distinguish a BGG^+ ciphertexts from random ones, since the partial decryption result by a single key essentially corresponds to a BGG^+ ciphertext in exponent and it cannot find any non-trivial linear relation among the random ciphertext components as long as the modulus size is exponential.

1.2 Related Works.

In an independent work (that predates ours), Brakerski and Vaikuntanathan [BV20] also construct broadcast encryption achieving optimal parameters. Their techniques as well as final result are very different from ours – while our work crucially uses pairings in conjunction with LWE, they rely entirely on LWE and new assumptions in the regime of lattices. Both works can be seen as following the broad approach of starting with a succinct single-key CP-ABE from LWE⁴, and adding collusion resistance using pairings (ours) or new techniques in the lattice regime (theirs).

The techniques in our work are similar in spirit to a growing line of work that uses “the best of both” of pairings and LWE [AJL⁺19, Agr19, JLS19, GQWW19], but quite different in details. Closest to our work are techniques used to construct *key policy functional encryption* [AJL⁺19, Agr19, JLS19], which use FHE (based on LWE) for encrypted evaluation and pairings for performing FHE decryption in the exponent. While a major challenge in these constructions is the leakage caused by FHE decryption noise, we sidestep this issue altogether because BE is an “all or nothing” primitive” with no secrets from a legitimate key holder. On the other hand, we need new tricks to handle the functionality and security of a *ciphertext-policy* scheme – for instance, we need to use position-wise randomness on the exponent to prevent ciphertext and secret key components from being paired in illegitimate positions to leak information.

2 Preliminaries

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

Notation. We use bold letters to denote vectors and the notation $[a, b]$ to denote the set of integers $\{k \in \mathbb{N} \mid a \leq k \leq b\}$. 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 .

2.1 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:

⁴ The single key CP-ABE with succinct CT was also discovered by Boneh and Kim [BK].

$\text{Setup}(1^\lambda) \rightarrow (\text{mpk}, \text{msk})$: 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 .
 $\text{Enc}(\text{mpk}, X, \mu) \rightarrow \text{ct}$: 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 .
 $\text{KeyGen}(\text{mpk}, \text{msk}, Y) \rightarrow \text{sk}_Y$: 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 .
 $\text{Dec}(\text{mpk}, \text{ct}, X, \text{sk}_Y, Y) \rightarrow \mu$ or \perp : We assume that the decryption algorithm is deterministic. 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 2.1 (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 2.2 (Ada-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 adversary \mathcal{A} , let us define Ada-IND security game as follows.

1. **Setup phase:** On input 1^λ , the challenger samples $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda)$ and gives mpk to \mathcal{A} .
2. **Query phase:** During the game, \mathcal{A} adaptively makes the following queries, in an arbitrary order. \mathcal{A} can make unbounded many key queries, but can make only single challenge query.
 - (a) **Key Queries:** \mathcal{A} 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} submits a pair of equal length messages $(\mu_0, \mu_1) \in (\mathcal{M})^2$ and the target $X^* \in A_\lambda$ to the challenger. The challenger samples a random bit $b \leftarrow \{0, 1\}$ and replies to \mathcal{A} with $\text{ct} \leftarrow \text{Enc}(\text{mpk}, X^*, \mu_b)$. We require that $R(X^*, Y) = 0$ holds for any Y such that \mathcal{A} makes a key query for Y in order to avoid trivial attacks.
3. **Output phase:** \mathcal{A} outputs a guess bit b' as the output of the experiment.

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

$$\text{Adv}_{\text{ABE}, \mathcal{A}}^{\text{Ada-IND}}(1^\lambda) := \left| \Pr[\text{Exp}_{\text{ABE}, \mathcal{A}}(1^\lambda) = 1 | b = 0] - \Pr[\text{Exp}_{\text{ABE}, \mathcal{A}}(1^\lambda) = 1 | b = 1] \right|.$$

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

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

Definition 2.3 (Ada-INDr security for ABE). We define Ada-INDr security game similarly to Ada-IND security game except that the adversary A chooses single message μ instead of (μ_0, μ_1) at the challenge phase and the challenger returns $ct \leftarrow \text{Enc}(\text{mpk}, X^*, \mu)$ if $b = 0$ and a random ciphertext $ct \leftarrow CT$ from a ciphertext space CT if $b = 1$. We define the advantage $\text{Adv}_{\text{ABE}, A}^{\text{Ada-INDr}}(1^\lambda)$ of the adversary A accordingly and say that the scheme satisfies Ada-INDr security if the quantity is negligible.

We also consider (weaker) selective versions of the above notions, where A specifies its target X^* at the beginning of the game.

Definition 2.4 (Sel-IND security for ABE). We define Sel-IND security game as Ada-IND security game with the exception that the adversary A has to choose the challenge ciphertext attribute X^* before the setup phase but key queries Y_1, Y_2, \dots and choice of (μ_0, μ_1) can still be adaptive. We define the advantage $\text{Adv}_{\text{ABE}, A}^{\text{Sel-IND}}(1^\lambda)$ of the adversary A accordingly and say that the scheme satisfies Sel-INDr security (or simply selective security) if the quantity is negligible.

Definition 2.5 (Sel-INDr security for ABE). We define Sel-INDr security game as Ada-INDr security game with the exception that the adversary A has to choose the challenge ciphertext attribute X^* before the setup phase but key queries Y_1, Y_2, \dots and choice of μ can still be adaptive. We define the advantage $\text{Adv}_{\text{ABE}, A}^{\text{Sel-INDr}}(1^\lambda)$ of the adversary A accordingly and say that the scheme satisfies Sel-INDr security if the quantity is negligible.

In the following, we recall definitions of various ABEs by specifying the relation. We start with the standard notions of ciphertext-policy attribute-based encryption (CP-ABE) and key-policy attribute-based encryption (KP-ABE).

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

KP-ABE for circuits. To define KP-ABE for circuits, we simply swap key and ciphertext attributes in CP-ABE for circuits. More formally, to define KP-ABE for circuits, we define $A_\lambda^{\text{KP}} = \{0, 1\}^\ell$ and $B_\lambda^{\text{KP}} = \mathcal{C}_\lambda$. We also define $R_\lambda^{\text{KP}} : A_\lambda^{\text{KP}} \times B_\lambda^{\text{KP}} \rightarrow \{0, 1\}$ as $R_\lambda^{\text{KP}}(\mathbf{x}, C) = \neg C(\mathbf{x})$.

We can also capture identity-based broadcast encryption (IBBE) and broadcast encryption (BE) as special cases of ABE by specifying the relations.

IBBE. To define IBBE, we define $A_\lambda^{\text{IBBE}} = \mathcal{ID}(\lambda)^{\leq t}$ and $B_\lambda^{\text{IBBE}} = \mathcal{ID}(\lambda)$, where $\mathcal{ID}(\lambda)$ is the identity space and $\mathcal{ID}(\lambda)^{\leq t}$ denotes all subsets of $\mathcal{ID}(\lambda)$ with size at most

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

t. We also define $R_\lambda^{\text{IBBE}} : A_\lambda^{\text{IBBE}} \times B_\lambda^{\text{IBBE}} \rightarrow \{0, 1\}$ as $R_\lambda^{\text{IBBE}}(S, \text{id}) = \begin{cases} 1 & \text{if } \text{id} \in S \\ 0 & \text{if } \text{id} \notin S \end{cases}$.

For IBBE, we typically require that the ciphertext size should be $o(t) \cdot \text{poly}(\lambda)$, since otherwise we have a trivial construction from IBE.

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.

We also define dual versions of BE and IBBE where the ciphertext and secret key attributes are swapped.

Dual IBBE (DIBBE). To define DIBBE, we define $A_\lambda^{\text{DIBBE}} = \mathcal{ID}(\lambda)$ and $B_\lambda^{\text{DIBBE}} = \mathcal{ID}(\lambda)^{\leq t}$, where $\mathcal{ID}(\lambda)$ is the identity space. We define $R_\lambda^{\text{DIBBE}} : A_\lambda^{\text{DIBBE}} \times B_\lambda^{\text{DIBBE}} \rightarrow \{0, 1\}$ as $R_\lambda^{\text{DIBBE}}(\text{id}, S) = 1$ if $\text{id} \in S$ and $R_\lambda^{\text{DIBBE}}(\text{id}, S) = 0$ otherwise.

Dual BE (DBE). To define DBE, we define $A_\lambda^{\text{DBE}} = [N(\lambda)]$ and $B_\lambda^{\text{DBE}} = 2^{[N(\lambda)]}$, where $N(\lambda) = \text{poly}(\lambda)$ is the number of users in the system. We also define $R_\lambda^{\text{DBE}} : A_\lambda^{\text{DBE}} \times B_\lambda^{\text{DBE}} \rightarrow \{0, 1\}$ as $R_\lambda^{\text{DBE}}(i, S) = 1$ when $i \in S$ and $R_\lambda^{\text{DBE}}(i, S) = 0$ otherwise.

2.2 Lattice Preliminaries

Here, we recall some facts on lattices that are needed for the exposition of our construction. Throughout this section, n , m , and q are integers such that $n = \text{poly}(\lambda)$ and $m \geq n \lceil \log q \rceil$. In the following, let $\text{SampZ}(\gamma)$ be a sampling algorithm for the truncated discrete Gaussian distribution over \mathbb{Z} with parameter $\gamma > 0$ whose support is restricted to $z \in \mathbb{Z}$ such that $|z| \leq \sqrt{n}\gamma$.

Learning with Errors. We introduce then learning with errors (LWE) problem.

Definition 2.6 (The LWE Assumption). Let $n = n(\lambda)$, $m = m(\lambda)$, and $q = q(\lambda) > 2$ be integers and $\chi = \chi(\lambda)$ be a distribution over \mathbb{Z}_q . We say that the $\text{LWE}(n, m, q, \chi)$ hardness assumption holds if for any PPT adversary \mathbf{A} we have

$$|\Pr[\mathbf{A}(\mathbf{A}, \mathbf{s}^\top \mathbf{A} + \mathbf{x}^\top) \rightarrow 1] - \Pr[\mathbf{A}(\mathbf{A}, \mathbf{v}^\top) \rightarrow 1]| \leq \text{negl}(\lambda)$$

where the probability is taken over the choice of the random coins by the adversary \mathbf{A} and $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$, $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, $\mathbf{x} \leftarrow \chi^m$, and $\mathbf{v} \leftarrow \mathbb{Z}_q^m$. We also say that $\text{LWE}(n, m, q, \chi)$ problem is subexponentially hard if the above probability is bounded by $2^{-n^\epsilon} \cdot \text{negl}(\lambda)$ for some constant $0 < \epsilon < 1$ for all PPT \mathbf{A} .

As shown by previous works [Reg09, BLP⁺13], if we set $\chi = \text{SampZ}(\gamma)$, the $\text{LWE}(n, m, q, \chi)$ problem is as hard as solving worst case lattice problems such as gapSVP and SIVP with approximation factor $\text{poly}(n) \cdot (q/\gamma)$ for some $\text{poly}(n)$. Since the best known algorithms for 2^k -approximation of gapSVP and SIVP run in time $2^{\tilde{O}(n/k)}$, it follows that the above $\text{LWE}(n, m, q, \chi)$ with noise-to-modulus ratio 2^{-n^ϵ} is likely to be (subexponentially) hard for some constant ϵ .

Trapdoors. Let us consider a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$. For all $\mathbf{V} \in \mathbb{Z}_q^{n \times m'}$, we let $\mathbf{A}_\gamma^{-1}(\mathbf{V})$ be an output distribution of $\text{SampZ}(\gamma)^{m \times m'}$ conditioned on $\mathbf{A} \cdot \mathbf{A}_\gamma^{-1}(\mathbf{V}) = \mathbf{V}$. A γ -trapdoor for \mathbf{A} is a trapdoor that enables one to sample from the distribution $\mathbf{A}_\gamma^{-1}(\mathbf{V})$ in time $\text{poly}(n, m, m', \log q)$ for any \mathbf{V} . We slightly overload notation and denote a γ -trapdoor for \mathbf{A} by \mathbf{A}_γ^{-1} . We also define the special gadget matrix $\mathbf{G} \in \mathbb{Z}_q^{n \times m}$ as the matrix obtained by padding $\mathbf{I}_n \otimes (1, 2, 4, 8, \dots, 2^{\lceil \log q \rceil})$ with zero-columns. The following properties had been established in a long sequence of works [GPV08, CHKP10, ABB10a, ABB10b, MP12, BLP⁺13].

Lemma 2.7 (Properties of Trapdoors). *Lattice trapdoors exhibit the following properties.*

1. Given \mathbf{A}_τ^{-1} , one can obtain $\mathbf{A}_{\tau'}^{-1}$ for any $\tau' \geq \tau$.
2. Given \mathbf{A}_τ^{-1} , one can obtain $[\mathbf{A} \parallel \mathbf{B}]_\tau^{-1}$ and $[\mathbf{B} \parallel \mathbf{A}]_\tau^{-1}$ for any \mathbf{B} .
3. There exists an efficient procedure $\text{TrapGen}(1^n, 1^m, q)$ that outputs $(\mathbf{A}, \mathbf{A}_{\tau_0}^{-1})$ where $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ for some $m = O(n \log q)$ and is 2^{-n} -close to uniform, where $\tau_0 = \omega(\sqrt{n \log q \log m})$.

Lattice Evaluation. The following is an abstraction of the evaluation procedure in previous LWE based FHE and ABE schemes. We follow the presentation by Tsabary [Tsa19], but with different parameters.

Lemma 2.8 (Fully Homomorphic Computation [GV15]). *There exists a pair of deterministic algorithms $(\text{EvalF}, \text{EvalFX})$ with the following properties.*

- $\text{EvalF}(\mathbf{B}, F) \rightarrow \mathbf{H}_F$. Here, $\mathbf{B} \in \mathbb{Z}_q^{n \times m^\ell}$ and $F : \{0, 1\}^\ell \rightarrow \{0, 1\}$ is a circuit.
- $\text{EvalFX}(F, \mathbf{x}, \mathbf{B}) \rightarrow \widehat{\mathbf{H}}_{F, \mathbf{x}}$. Here, $\mathbf{x} \in \{0, 1\}^\ell$ with $x_1 = 1^6$ and $F : \{0, 1\}^\ell \rightarrow \{0, 1\}$ is a circuit with depth d . We have $[\mathbf{B} - \mathbf{x} \otimes \mathbf{G}] \widehat{\mathbf{H}}_{F, \mathbf{x}} = \mathbf{B} \mathbf{H}_F - F(\mathbf{x}) \mathbf{G} \pmod q$, where we denote $[x_1 \mathbf{G} \parallel \dots \parallel x_k \mathbf{G}]$ by $\mathbf{x} \otimes \mathbf{G}$. Furthermore, we have $\|\mathbf{H}_F\|_\infty \leq m \cdot 2^{O(d)}$, $\|\widehat{\mathbf{H}}_{F, \mathbf{x}}\|_\infty \leq m \cdot 2^{O(d)}$.
- The running time of $(\text{EvalF}, \text{EvalFX})$ is bounded by $\text{poly}(n, m, \log q, 2^d)$.

The above algorithms are taken from [GV15], which is a variant of similar algorithms proposed by Boneh et al. [BGG⁺14]. The algorithms in [BGG⁺14] work for any polynomial-sized circuit F , but $\|\mathbf{H}_F\|_\infty$ and $\|\widehat{\mathbf{H}}_{F, \mathbf{x}}\|_\infty$ become super-polynomial even if the depth of the circuit is shallow (i.e., logarithmic depth). On the other hand, the above algorithms run in polynomial time only when F is of logarithmic depth, but $\|\mathbf{H}_F\|_\infty$ and $\|\widehat{\mathbf{H}}_{F, \mathbf{x}}\|_\infty$ can be polynomially bounded. The latter property is crucial for our purpose.

2.3 KP-ABE Scheme by Boneh et al. [BGG⁺14].

We will use a variant of the KP-ABE scheme proposed by Boneh et al. [BGG⁺14] as a building block of our construction of CP-ABE. We call the scheme BGG^+ and provide the description of the scheme in the following. We focus on the case where the policies associated with secret keys are limited to circuits with logarithmic depth rather than

⁶ This condition may be necessary for the lemma to hold for arbitrary F .

arbitrary polynomially bounded depth, so that we can use the evaluation algorithm due to Gorbunov and Vinayagamurthy [GV15] (see Lemma 2.8). This allows us to bound the noise growth during the decryption by a polynomial factor, which is crucial for our application.

The scheme supports the circuit class $\mathcal{C}_{\ell(\lambda), d(\lambda)}$, 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 $d(\lambda) = O(\log \lambda)$.

Setup(1^λ): On input 1^λ , the setup algorithm defines the parameters $n = n(\lambda)$, $m = m(\lambda)$, noise distribution χ over \mathbb{Z} , τ_0 , τ , and $B = B(\lambda)$ as specified later. It then proceeds as follows.

1. Sample $(\mathbf{A}, \mathbf{A}_{\tau_0}^{-1}) \leftarrow \text{TrapGen}(1^n, 1^m, q)$ such that $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$.
2. Sample random matrix $\mathbf{B} = (\mathbf{B}_1, \dots, \mathbf{B}_\ell) \leftarrow (\mathbb{Z}_q^{n \times m})^\ell$ and a random vector $\mathbf{u} \leftarrow \mathbb{Z}_q^n$.
3. Output the master public key $\text{mpk} = (\mathbf{A}, \mathbf{B}, \mathbf{u})$ and the master secret key $\text{msk} = \mathbf{A}_{\tau_0}^{-1}$.

KeyGen($\text{mpk}, \text{msk}, F$): The key generation algorithm takes as input the master public key mpk , the master secret key msk , and a circuit $F \in \mathcal{F}_\lambda$ and proceeds as follows.

1. Compute $\mathbf{H}_F = \text{EvalF}(\mathbf{B}, F)$ and $\mathbf{B}_F = \mathbf{B}\mathbf{H}_F$.
2. Compute $[\mathbf{A} \parallel \mathbf{B}_F]_{\tau}^{-1}$ from $\mathbf{A}_{\tau_0}^{-1}$ and sample $\mathbf{r} \in \mathbb{Z}^{2m}$ as $\mathbf{r} \leftarrow [\mathbf{A} \parallel \mathbf{B}_F]_{\tau}^{-1}(\mathbf{u})$.
3. Output the secret key $\text{sk}_F := \mathbf{r}$.

Enc($\text{mpk}, \mathbf{x}, \mu$): The encryption algorithm takes as input the master public key mpk , an attribute $\mathbf{x} \in \{0, 1\}^\ell$ with $x_1 = 1$,⁷ and a message $\mu \in \{0, 1\}$ and proceeds as follows.

1. Sample $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, $e_1 \leftarrow \chi$, $\mathbf{e}_2 \leftarrow \chi^m$, and $\mathbf{S}_{i,b} \leftarrow \{-1, 1\}^{m \times m}$ for $i \in [\ell]$ and $b \in \{0, 1\}$. Then, set $\mathbf{e}_{i,b} := \mathbf{S}_{i,b}^\top \mathbf{e}_2$ for $i \in [\ell]$ and $b \in \{0, 1\}$.
2. Compute

$$\begin{aligned} \psi_1 &:= \mathbf{s}^\top \mathbf{u} + e_1 + \mu \lceil q/2 \rceil \in \mathbb{Z}_q, & \psi_2^\top &:= \mathbf{s}^\top \mathbf{A} + \mathbf{e}_2^\top \in \mathbb{Z}_q^m, \\ \psi_{i,b}^\top &:= \mathbf{s}^\top (\mathbf{B} - x_i \mathbf{G}) + \mathbf{e}_{i,b}^\top \in \mathbb{Z}_q^m & \text{for all } i \in [\ell] \text{ and } b \in \{0, 1\}. \end{aligned}$$

3. Output the ciphertext $\text{ct}_{\mathbf{x}} := (\psi_1, \psi_2, \{\psi_{i,x_i}\}_{i \in [\ell]})$, where x_i is the i -th bit of \mathbf{x} .

Dec($\text{mpk}, \text{sk}_{\mathbf{x}}, \mathbf{x}, F, \text{ct}_F$): The decryption algorithm takes as input the master public key mpk , a secret key sk_F for a circuit F , and a ciphertext $\text{ct}_{\mathbf{x}}$ for an attribute \mathbf{x} and proceeds as follows.

1. Parse $\text{ct}_{\mathbf{x}} \rightarrow (\psi_1 \in \mathbb{Z}_q, \psi_2 \in \mathbb{Z}_q^m, \{\psi_{i,x_i} \in \mathbb{Z}_q^m\}_{i \in [\ell]})$, and $\text{sk}_F \in \mathbb{Z}^{2m}$. If any of the component is not in the corresponding domain or $F(\mathbf{x}) = 1$, output \perp .
2. Concatenate $\{\psi_{i,x_i}\}_{i \in [\ell]}$ to form $\psi_3^\top = (\psi_{1,x_1}^\top, \dots, \psi_{\ell,x_\ell}^\top)$.
3. Compute $\psi' := \psi_1 - [\psi_2^\top \parallel \psi_3^\top] \mathbf{r}$.
4. Output 0 if $\psi' \in [-B, B]$ and 1 if $[-B + \lceil q/2 \rceil, B + \lceil q/2 \rceil]$.

⁷ This restriction is required to apply Lemma 2.8. We can remove the condition by increasing the dimension of \mathbf{x} by 1 and considering function F that ignores the first bit.

Remark 2.9. We note that the encryption algorithm above computes redundant components $\{\psi_{i, \neg x_i}\}_{i \in [\ell]}$ in the second step, which are discarded in the third step. However, due to this redundancy, the scheme has the following special structure that will be useful for us. Namely, the first and the second steps of the encryption algorithm can be executed without knowing \mathbf{x} . Only the third step of the encryption algorithm needs the information of \mathbf{x} , where it chooses $\{\psi_{i, x_i}\}_{i \in [\ell]}$ from $\{\psi_{i, b}\}_{i \in [\ell], b \in \{0,1\}}$ depending on each bit of \mathbf{x} and then output the former terms along with ψ_1 and ψ_2 . Looking ahead, our construction of CP-ABE in Sec. 3 crucially relies on this special structure. There, the encryption algorithm, who takes as input a circuit C that specifies the policy and does not know the corresponding input \mathbf{x} , executes the first two steps of the above encryption algorithm. This is possible since these two steps do not need the knowledge of \mathbf{x} .

Parameters and Security. We choose the parameters for the scheme as follows:

$$\begin{aligned} m &= n^{1.1} \log q, & q &= 2^{\Theta(\lambda)}, & \chi &= \text{SampZ}(3\sqrt{n}), \\ \tau_0 &= n \log q \log m, & \tau &= m^{3.1} \ell \cdot 2^{O(d)} & B &= n^2 m^2 \tau \cdot 2^{O(d)}. \end{aligned}$$

The parameter n will be chosen depending on whether we need Sel-INDr security or Ada-INDr security for the scheme. If it suffices to have Sel-INDr security, we set $n = \lambda^c$ for some constant $c > 1$. If we need Ada-INDr security, we have to enlarge the parameter to be $n = (\ell\lambda)^c$ in order to compensate for the security loss caused by the complexity leveraging.

We remark that if we were to use the above ABE scheme stand-alone, we would have been able to set q polynomially bounded as in [GV15]. The reason why we set q exponentially large is that we combine the scheme with bilinear maps of order q to lift the ciphertext components to the exponent so that they are “hidden” in some sense (See Sec. 4). In order to use the security of the bilinear map, we set the group order q to be exponentially large.

The following theorem summarizes the security and efficiency properties of the construction. There are two parameter settings depending on whether we assume subexponential hardness of LWE or not.

Theorem 2.10 (Adapted from [GV15, BGG⁺14]). *Assuming hardness of $\text{LWE}(n, m, q, \chi)$ with $\chi = \text{SampZ}(3\sqrt{n})$ and $q = O(2^{n^{1/\epsilon}})$ for some constant $\epsilon > 1$, the above scheme satisfies Sel-INDr security. Assuming subexponential hardness of $\text{LWE}(n, m, q, \chi)$ with the same parameters, the above scheme satisfies Ada-INDr security with respect to the ciphertext space $\mathcal{CT} := \mathbb{Z}_q^{m(\ell+1)+1}$*

2.4 Bilinear Map Preliminaries

Here, we introduce our notation for bilinear maps and the bilinear generic group model following Baltico et.al [BCFG17], who specializes the framework by Barthe [BFF⁺14] for defining generic k -linear groups to the bilinear group settings. The definition closely follows that of Maurer [Mau05], which is equivalent to the alternative formulation by Shoup [Sho97].

Notation on Bilinear Maps. A bilinear group generator takes as input 1^λ and outputs a group description $\mathbb{G} = (q, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, g_1, g_2)$, where q is a prime of $\Theta(\lambda)$ bits, $\mathbb{G}_1, \mathbb{G}_2$, and \mathbb{G}_T are cyclic groups of order q , $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$ 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 representation of group elements: for a matrix \mathbf{A} over \mathbb{Z}_q , 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 also use the following less standard notations. For vectors $\mathbf{w} = (w_1, \dots, w_\ell)^\top \in \mathbb{Z}_q^\ell$ and $\mathbf{v} = (v_1, \dots, v_\ell)^\top \in \mathbb{Z}_q^\ell$ of the same length, $\mathbf{w} \odot \mathbf{v}$ denotes the vector that is obtained by component-wise multiplications. Namely, $\mathbf{v} \odot \mathbf{w} = (v_1 w_1, \dots, v_\ell w_\ell)^\top$. When $\mathbf{w} \in (\mathbb{Z}_q^*)^\ell$, $\mathbf{v} \oslash \mathbf{w}$ denotes the vector $\mathbf{v} \oslash \mathbf{w} = (v_1/w_1, \dots, v_\ell/w_\ell)^\top$. It is easy to verify that for vectors $\mathbf{c}, \mathbf{d} \in \mathbb{Z}_q^\ell$ and $\mathbf{w} \in (\mathbb{Z}_q^*)^\ell$, we have $(\mathbf{c} \odot \mathbf{w}) \odot (\mathbf{d} \oslash \mathbf{w}) = \mathbf{c} \odot \mathbf{d}$. For group elements $[\mathbf{v}]_1 \in \mathbb{G}_1^\ell$ and $[\mathbf{w}]_1 \in \mathbb{G}_1^\ell$, $[\mathbf{v}]_1 \odot [\mathbf{w}]_2$ denotes $([v_1 w_1]_T, \dots, [v_\ell w_\ell]_T)^\top$, which is efficiently computable from $[\mathbf{v}]_1$ and $[\mathbf{w}]_2$ using the bilinear map e .

Generic Bilinear Group Model. Let $\mathbb{G} = (q, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, g_1, g_2)$ 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 definition of the generic group model by Baltico et. al [BCFG17]. 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 and add_T . 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

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 $-L_s[h_1]$ 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 $\mathbb{G} = (q, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, g_1, g_2)$ and distribution \mathcal{D} .

generic group model, except that internally the challenger stores lists of element 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. In our work, we will use the subring $\mathbb{Z}_p[X_1, \dots, X_n, 1/X_1, \dots, 1/X_n]$ of the entire field $\mathbb{Z}_p(X_1, \dots, X_n)$. Note that any element f in $\mathbb{Z}_p[X_1, \dots, X_n, 1/X_1, \dots, 1/X_n]$ can be represented as $f(X_1, \dots, X_n) = \sum_{(c_1, \dots, c_n) \in \mathbb{Z}^n} a_{c_1, \dots, c_n} X_1^{c_1} \cdots X_n^{c_n}$ using $\{a_{c_1, \dots, c_n} \in \mathbb{Z}_p\}_{(c_1, \dots, c_n) \in \mathbb{Z}^n}$, where we have $a_{c_1, \dots, c_n} = 0$ for all but finite $(c_1, \dots, c_n) \in \mathbb{Z}^n$. Note that this expression is unique.

3 Our Construction of CP-ABE

Here, we describe our new construction of CP-ABE scheme. Our construction can deal with any circuit class $\mathcal{F} = \{\mathcal{F}_\lambda\}_\lambda$ that is subclass of $\{\mathcal{C}_{\ell(\lambda), d(\lambda)}\}_\lambda$ with arbitrary $\ell(\lambda) \leq \text{poly}(\lambda)$ and $d(\lambda) = O(\log \lambda)$, where $\mathcal{C}_{\ell(\lambda), d(\lambda)}$ is a set of circuits with input length $\ell(\lambda)$ and depth at most $d(\lambda)$. As we will see in Sec. 5, we can obtain new constructions of BE, IBBE, CP-ABE by setting the circuit class \mathcal{F} appropriately. In order to get the scheme, we use the KP-ABE scheme BGG^+ for the circuit class $\mathcal{F} = \{\mathcal{F}_\lambda\}_\lambda$ that is described in Sec. 2.3 as an ingredient. Our construction below can be seen as a conversion from an ABE scheme to another ABE scheme with dual predicate.

Setup(1^λ): On input 1^λ , the setup algorithm defines the parameters $n = n(\lambda)$, $m = m(\lambda)$, noise distribution χ over \mathbb{Z} , τ_0, τ , and $B = B(\lambda)$ as specified in Sec. 2.3.

It samples a group description $\mathbb{G} = (q, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, [1]_1, [1]_2)$. It then sets $L := (2\ell + 1)m + 2$ and proceeds as follows.

1. Sample $\mathbf{w} \leftarrow (\mathbb{Z}_q^*)^L$ and compute $[\mathbf{w}]_1$.
2. Output $\text{mpk} = ([1]_1, [1]_2, [\mathbf{w}]_1)$ and $\text{msk} = \mathbf{w}$.

KeyGen($\text{mpk}, \text{msk}, \mathbf{x}$): The key generation algorithm takes as input the master public key mpk , the master secret key msk , and an attribute $\mathbf{x} \in \{0, 1\}^\ell$ with $x_1 = 1$ and proceeds as follows.

1. Let $\mathbf{1} := (1, \dots, 1)^\top \in \mathbb{Z}_q^m$ and $\mathbf{0} := (0, \dots, 0)^\top \in \mathbb{Z}_q^m$. Set

$$\begin{aligned} \phi_0 &= \mathbf{1} \in \mathbb{Z}_q, & \phi_1 &= \mathbf{1} \in \mathbb{Z}_q, & \phi_2 &:= \mathbf{1} \in \mathbb{Z}_q^m, \\ \phi_{i,b} &:= \begin{cases} \mathbf{1} \in \mathbb{Z}_q^m & \text{if } b = x_i \\ \mathbf{0} \in \mathbb{Z}_q^m & \text{if } b \neq x_i \end{cases} & \text{for } i \in [\ell] \text{ and } b \in \{0, 1\}. \end{aligned} \quad (3.1)$$

2. Vectorize $(\phi_0, \phi_1, \phi_2, \{\phi_{i,b}\}_{i,b})$ to form a vector $\mathbf{d} \in \mathbb{Z}_q^L$ by concatenating each entry of the vectors in a predetermined order.
3. Sample $\delta \leftarrow \mathbb{Z}_q^*$.
4. Compute $[\delta \mathbf{d} \odot \mathbf{w}]_2 \in \mathbb{G}_2^L$ from $\text{msk} = \mathbf{w}$ in msk .
5. Output $\text{sk}_{\mathbf{x}} = [\delta \mathbf{d} \odot \mathbf{w}]_2$.

$\text{Enc}(\text{mpk}, F, \mu)$: The encryption algorithm takes as input the master public key mpk , the circuit F , and a message $\mu \in \{0, 1\}$ and proceeds as follows.

1. Sample fresh BGG⁺ scheme:

- (a) Sample $(\mathbf{A}, \mathbf{A}_{\tau_0}^{-1}) \leftarrow \text{TrapGen}(1^n, 1^m, q)$ such that $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$.
- (b) Sample random matrix $\mathbf{B} = (\mathbf{B}_1, \dots, \mathbf{B}_\ell) \leftarrow (\mathbb{Z}_q^{n \times m})^\ell$ and a random vector $\mathbf{u} \leftarrow \mathbb{Z}_q^n$.

2. Compute BGG⁺ function key for circuit F :

- (a) Compute $\mathbf{H}_F = \text{EvalF}(\mathbf{B}, F)$ and $\mathbf{B}_F = \mathbf{B} \mathbf{H}_F$.
- (b) Compute $[\mathbf{A} \parallel \mathbf{B}_F]_\tau^{-1}$ from $\mathbf{A}_{\tau_0}^{-1}$ and sample $\mathbf{r} \in \mathbb{Z}^{2m}$ as $\mathbf{r} \leftarrow [\mathbf{A} \parallel \mathbf{B}_F]_\tau^{-1}(\mathbf{u})$.

3. Compute BGG⁺ ciphertext for all possible inputs:

- (a) Sample $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, $e_1 \leftarrow \chi$, $e_2 \leftarrow \chi^m$, and $\mathbf{S}_{i,b} \leftarrow \{-1, 1\}^{m \times m}$ for $i \in [\ell]$ and $b \in \{0, 1\}$. Then, set $\mathbf{e}_{i,b} := \mathbf{S}_{i,b}^\top e_2$ for $i \in [\ell]$ and $b \in \{0, 1\}$.
- (b) Compute

$$\begin{aligned} \psi_0 &:= \mathbf{1} \in \mathbb{Z}_q, & \psi_1 &:= \mathbf{s}^\top \mathbf{u} + e_1 + \mu \lceil q/2 \rceil \in \mathbb{Z}_q, \\ \psi_2^\top &:= \mathbf{s}^\top \mathbf{A} + \mathbf{e}_2^\top \in \mathbb{Z}_q^m, \\ \psi_{i,b}^\top &:= \mathbf{s}^\top (\mathbf{B}_i - b \mathbf{G}) + \mathbf{e}_{i,b}^\top \in \mathbb{Z}_q^m & \text{for } i \in [\ell] \text{ and } b \in \{0, 1\}. \end{aligned} \quad (3.2)$$

4. Encode BGG⁺ ciphertexts in exponent of bilinear group:

- (a) Vectorize $(\psi_0, \psi_1, \psi_2, \{\psi_{i,b}\}_{i,b})$ to form a vector $\mathbf{c} \in \mathbb{Z}_q^L$ by concatenating each entry of the vectors in a predetermined order (that aligns with the one used in the key generation algorithm).
- (b) Sample $\gamma \leftarrow \mathbb{Z}_q^*$.
- (c) Compute $[\gamma \mathbf{c} \odot \mathbf{w}]_1 \in \mathbb{G}_2^L$ from γ , \mathbf{c} , and $[\mathbf{w}]_1$ in mpk .

5. Output $\text{ct}_F = (\text{ct}_0 = (\mathbf{A}, \mathbf{B}), \text{ct}_1 = [\gamma \mathbf{c} \odot \mathbf{w}]_1, \text{ct}_2 = \mathbf{r})$.

$\text{Dec}(\text{mpk}, \text{sk}_{\mathbf{x}}, \mathbf{x}, F, \text{ct}_F)$: The decryption algorithm takes as input the master public key mpk , the secret key $\text{sk}_{\mathbf{x}}$ for an attribute \mathbf{x} , and the ciphertext ct_F for a circuit F and proceeds as follows.

1. Parse $\text{ct}_F \rightarrow (\text{ct}_0 = (\mathbf{A} \in \mathbb{Z}_q^{n \times m}, \mathbf{B} \in \mathbb{Z}_q^{n \times m \ell}), \text{ct}_1 \in \mathbb{G}_1^L, \text{ct}_2 \in \mathbb{Z}^{2m})$ and $\text{sk}_{\mathbf{x}} \in \mathbb{G}_2^L$. If any of the component is not in the corresponding domain or $F(\mathbf{x}) = 1$, output \perp .

2. Unmask BGG⁺ ciphertexts corresponding to \mathbf{x} by using secret key:
 Compute $[\mathbf{v}]_T := \text{ct}_1 \odot \text{sk}_{\mathbf{x}}$ and de-vectorize $[\mathbf{v}]_T$ to obtain
 $[v_0]_T \in \mathbb{G}_T, [v_1]_T \in \mathbb{G}_T, [v_2]_T \in \mathbb{G}_T^m, [\mathbf{v}_{i,b}]_T \in \mathbb{G}_T^m, \text{ for } i \in [\ell], b \in \{0, 1\}.$
3. Evaluate circuit F on BGG⁺ ciphertexts in the exponent:
 Compute $\widehat{\mathbf{H}}_{F,\mathbf{x}} = \text{EvalF}(F, \mathbf{x}, \mathbf{B}).$
4. Perform BGG⁺ decryption in the exponent:
 Form $[\mathbf{v}_{\mathbf{x}}^\top]_T = [\mathbf{v}_{1,x_1}^\top, \dots, \mathbf{v}_{\ell,x_\ell}^\top]_T$ and $\text{ct}_2^\top = (\mathbf{r}_1^\top \in \mathbb{Z}_q^m, \mathbf{r}_2^\top \in \mathbb{Z}_q^m).$ Then compute

$$[v']_T := [v_1 - (\mathbf{v}_2^\top \mathbf{r}_1 + \mathbf{v}_{\mathbf{x}}^\top \widehat{\mathbf{H}}_{F,\mathbf{x}} \mathbf{r}_2)]_T$$
 from $[v_1]_T, [v_2]_T, [\mathbf{v}_{\mathbf{x}}]_T, \mathbf{r}_1, \mathbf{r}_2,$ and $\widehat{\mathbf{H}}_{F,\mathbf{x}}.$
5. Recover exponent via brute force if $F(\mathbf{x}) = 0$:
 Find $\eta \in [-B, B] \cup [-B + \lceil q/2 \rceil, B + \lceil q/2 \rceil]$ such that $[v_0]_T^\eta = [v']_T$ by brute-force search. If there is no such η , output \perp . To speed up the operation, one can employ the baby-step giant-step algorithm.
6. Output 0 if $\eta \in [-B, B]$ and 1 if $[-B + \lceil q/2 \rceil, B + \lceil q/2 \rceil]$.

Correctness. To see correctness of the scheme, we first observe that we have $\text{ct}_1 \odot \text{sk}_{\mathbf{x}} = [\gamma\delta \cdot \mathbf{c} \odot \mathbf{d}]_T$ and thus

$$v_0 = \gamma\delta, \quad v_1 = \gamma\delta (\mathbf{s}^\top \mathbf{u} + e_1 + \mu \lceil q/2 \rceil), \quad \mathbf{v}_2^\top = \gamma\delta (\mathbf{s}^\top \mathbf{A} + \mathbf{e}_2^\top),$$

$$\mathbf{v}_{i,b}^\top = \begin{cases} \gamma\delta (\mathbf{s}^\top (\mathbf{B}_i - x_i \mathbf{G}) + \mathbf{e}_{i,x_i}^\top) & \text{if } b = x_i \\ \mathbf{0} & \text{if } b = 1 - x_i \end{cases}.$$

From the above, we have $\mathbf{v}_{\mathbf{x}}^\top = \mathbf{s}^\top (\mathbf{B} - \mathbf{x} \otimes \mathbf{G}) + \mathbf{e}_{\mathbf{x}}^\top$ for $\mathbf{e}_{\mathbf{x}}^\top := (\mathbf{e}_{1,x_1}^\top, \dots, \mathbf{e}_{\ell,x_\ell}^\top).$ We then have

$$\begin{aligned} \mathbf{v}_2^\top \mathbf{r}_1 + \mathbf{v}_{\mathbf{x}}^\top \widehat{\mathbf{H}}_{F,\mathbf{x}} \mathbf{r}_2 &= \gamma\delta (\mathbf{s}^\top \mathbf{A} + \mathbf{e}_2^\top) \mathbf{r}_1 + \gamma\delta (\mathbf{s}^\top (\mathbf{B} - \mathbf{x} \otimes \mathbf{G}) + \mathbf{e}_{\mathbf{x}}^\top) \widehat{\mathbf{H}}_{F,\mathbf{x}} \mathbf{r}_2 \\ &= \gamma\delta (\mathbf{s}^\top (\mathbf{A} \mathbf{r}_1 + \mathbf{B}_F \mathbf{r}_2) + \mathbf{e}_2^\top \mathbf{r}_1 + \mathbf{e}_{\mathbf{x}}^\top \widehat{\mathbf{H}}_{F,\mathbf{x}} \mathbf{r}_2) \\ &= \gamma\delta (\mathbf{s}^\top \mathbf{u} + \mathbf{e}_2^\top \mathbf{r}_1 + \mathbf{e}_{\mathbf{x}}^\top \widehat{\mathbf{H}}_{F,\mathbf{x}} \mathbf{r}_2) \end{aligned}$$

where the second equation follows from $(\mathbf{B} - \mathbf{x} \otimes \mathbf{G}) \widehat{\mathbf{H}}_{F,\mathbf{x}} = \mathbf{B}_F$ and the third equation follows from $[\mathbf{A} \parallel \mathbf{B}_F] \mathbf{r} = \mathbf{u}$. This implies

$$v' = \gamma\delta (\mu \lceil q/2 \rceil + e_1 - \mathbf{e}_2^\top \mathbf{r}_1 - \mathbf{e}_{\mathbf{x}}^\top \widehat{\mathbf{H}}_{F,\mathbf{x}} \mathbf{r}_2).$$

Recall that we set $\chi = \text{SampZ}(3\sqrt{n})$. By the definition of SampZ , we have $\|e_1\|_\infty \leq 3n$ and $\|e_2\|_\infty \leq 3n$. Furthermore, we have $\|e_{i,b}\|_\infty = \|\mathbf{S}_{i,b}^\top e_2\|_\infty \leq 3mn$ for $i \in [\ell]$ and $b \in \{0, 1\}$, $\|\mathbf{r}\|_\infty \leq \sqrt{n}\tau$, and $\|\widehat{\mathbf{H}}_{F,\mathbf{x}}\|_\infty \leq m \cdot 2^{O(d)}$, where the last inequality follows from Lemma 2.8. Thus, we have

$$\|e_1 - \mathbf{e}_2^\top \mathbf{r}_1 - \mathbf{e}_{\mathbf{x}}^\top \widehat{\mathbf{H}}_{F,\mathbf{x}} \mathbf{r}_2\|_\infty \leq O(n^{1.5} m^2 \tau \cdot 2^{O(d)}) \leq B$$

by our choice of B . The correctness therefore follows. Note that since $B = \text{poly}(n, \ell) \cdot 2^{O(d)} = \text{poly}(\lambda)$, the decryption algorithm runs in polynomial time.

Efficiency of the Scheme. Here, we evaluate the efficiency of the above scheme. In particular, we measure the sizes of the parameters. The master public key of the scheme consists of $L + 2$ group elements. Since $L = O(m\ell)$, we have that the master public key can be represented by a binary string of length $\ell \cdot \text{poly}(\lambda)$. Next, we observe that a secret key in the scheme consists of L group elements, which can be represented by a binary string of length $\ell \cdot \text{poly}(\lambda)$. Finally, a ciphertext in the scheme consists of $O(nm)$ elements of \mathbb{Z}_q and L group elements. The former elements are represented by a binary string of length $\text{poly}(\lambda)$ if we only need Sel-INDr security for the underlying KP-ABE scheme. If we need Ada-INDr security, the length of the binary string is $\text{poly}(\ell, \lambda)$. Therefore, the length of the whole ciphertext is $\ell \cdot \text{poly}(\lambda)$ if we only need Sel-INDr security for the underlying KP-ABE scheme and $\text{poly}(\ell, \lambda)$ if we need Ada-INDr security. In any case, the sizes of all parameters in the system are independent of the size of the circuits being supported by the scheme, which is a notable feature of the scheme.

4 Security Proof for Our CP-ABE

This section is devoted to prove the following theorem that asserts the security of our CP-ABE scheme in Sec. 3.

Theorem 4.1. *Our CP-ABE scheme for function class \mathcal{F} satisfies Ada-IND security in the generic group model assuming that the KP-ABE BGG^+ for function class \mathcal{F} satisfies Ada-INDr security.*

Overview of the Proof. Before going to the formal proof, we give its overview. The proof is done by considering a sequence of games and consists of two parts. In the first part of the proof, which is captured by a series of game hops from **Game**₀ through **Game**₅ defined below, we prove that it is pointless for the adversary to take pairing products between unmatching positions of the ciphertext and secret key components and then take linear combinations among them. Therefore, the only possible strategy for the adversary is to take linear combination among “partial decryption results” obtained by taking pairing products between matching positions of the ciphertext and secret key components and infer information of the message being encrypted. In the second step of the proof, which is captured by the game hop from **Game**₅ to **Game**₆, we show that this type of attack does not work either. To do so, we further consider a sequence of subgames from **Game**_{5,0} through **Game**_{5,8}. We first prove that taking linear combinations among partial decryption results from different secret keys is useless. This is the key step that excludes the collusion attack and is captured by the game hop from **Game**_{5,3} to **Game**_{5,4}. At this point, the only strategy for the adversary is to take linear combination among partial decryption result obtained by single secret key. Finally, in the step from **Game**_{5,7} to **Game**_{5,8}, we use the security of the BGG^+ ABE to conclude that this strategy does not work either. To invoke the security of BGG^+ ABE, we use the fact that the partial decryption result obtained by secret key for x forms randomized version of BGG^+ ABE ciphertext for attribute x in the exponent.

Proof. To prove the theorem, we fix a PPT adversary A that makes at most $Q_{\text{kq}}(\lambda)$ key queries and $Q_{\text{zt}}(\lambda)$ zero-test queries during the game. Furthermore, we assume that 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. In order to prove the security, we consider following sequence of games. Let us denote the event that A outputs correct guess for b at the end of Game_x as E_x .

Game₀: This is the real game in the generic group model. To fix the notation and for the sake of concreteness, we briefly describe the game here. Without loss of generality, we assume that the challenger simulates the generic group oracle for A. At the beginning of the game, the challenger picks $\mathbf{w} \leftarrow (\mathbb{Z}_q^*)^L$ and sets the master public key $\text{mpk} = ([1]_1, [1]_2, [\mathbf{w}]_1)$ and the master secret key $\text{msk} = \mathbf{w}$. Then, it gives handles to the group elements in mpk to A. To respond to the j -th key query $\mathbf{x}^{(j)}$ made by A, the challenger samples $\delta_j \leftarrow \mathbb{Z}_q^*$, sets $\mathbf{d}^{(j)} \in \mathbb{Z}_q^L$ as specified in the key generation algorithm, and sets $\text{sk}^{(j)} = [\delta_j \mathbf{d}^{(j)} \odot \mathbf{w}]_2$. It then gives the handles corresponding to the group elements in $\text{sk}^{(j)}$ to A. To answer the challenge query for a circuit F , the challenger first picks the message $b \leftarrow \{0, 1\}$ to be encrypted, chooses $\gamma \leftarrow \mathbb{Z}_q^*$, computes $\mathbf{A}, \mathbf{B}, \mathbf{r}, \mathbf{c}$ as specified in the encryption algorithm (where b is encrypted), and forms the challenge ciphertext as $\text{ct}_F = (\text{ct}_0 = (\mathbf{A}, \mathbf{B}), \text{ct}_1 = [\gamma \mathbf{c} \odot \mathbf{w}]_1, \text{ct}_2 = \mathbf{r})$. It then returns $\text{ct}_0 = (\mathbf{A}, \mathbf{B})$, handles to $\text{ct}_1 = [\gamma \mathbf{c} \odot \mathbf{w}]_1$, and ct_2 to A. By definition, the advantage of A against the scheme is $|\Pr[E_0] - \frac{1}{2}|$.

Game₁: This game is the same as the previous game except that the challenger samples $\mathbf{w} = (w_1, \dots, w_L)^\top$, $\delta_1, \dots, \delta_{Q_{\text{kq}}}$, $\mathbf{A}, \mathbf{B}, \mathbf{u}, \gamma, b$, and $\mathbf{c} = (c_1, \dots, c_L)^\top$ at the beginning of the game. Note that \mathbf{c} is sampled from the distribution that is only dependent on the bit b being encrypted, and is *independent* of the circuit F that is specified by A later in the game. Therefore, this game is well-defined. As we prove in Lemma 4.2, we have $\Pr[E_0] = \Pr[E_1]$.

Game₂: In this game, we (partially) switch to the symbolic group model and replace $\{w_i\}_{i \in [L]}$, $\{\delta_j\}_{j \in [Q_{\text{kq}}]}$, γ , and $\{c_i\}_{i \in [L]}$ in \mathbb{Z}_q with the formal variables $\{W_i\}_{i \in [L]}$, $\{\Delta_j\}_{j \in [Q_{\text{kq}}]}$, Γ , and $\{C_i\}_{i \in [L]}$ respectively. As a result, all handles given to A refer to elements in the ring

$$\mathbb{T} := \mathbb{Z}_q[W_1, \dots, W_L, 1/W_1, \dots, 1/W_L, \Delta_1, \dots, \Delta_{Q_{\text{kq}}}, \Gamma, C_1, \dots, C_L],$$

where $\{1/W_i\}_i$ are needed to represent the components in the secret keys. However, when the challenger answers the zero-test queries, it substitutes the formal variables with corresponding elements in \mathbb{Z}_q . Namely, in this game, the challenger picks $\{w_i\}_i$, $\{\delta_j\}_j$, γ , and $\{c_i\}_i$ at the beginning of the game as specified in the previous game and when A makes a zero-test query for a handle corresponding to $f(W_1, \dots, W_L, \Delta_1, \dots, \Delta_{Q_{\text{kq}}}, \Gamma, C_1, \dots, C_L) \in \mathbb{T}$, the challenger returns 1 if

$$f(w_1, \dots, w_L, \delta_1, \dots, \delta_{Q_{\text{kq}}}, \gamma, c_1, \dots, c_L) = 0$$

holds over \mathbb{Z}_q and 0 otherwise. As we prove in Lemma 4.3, we have $\Pr[E_1] = \Pr[E_2]$.

Here, we list all the components in \mathbb{T} for which corresponding handles are given to A in **Game₂** as either handles to the group elements in mpk , the challenge ciphertext,

or secret keys:

$$S_1 := \{ 1, W_i, \{C_i \Gamma W_i\}_{i \in [L]} \}, \quad S_2 := \left\{ 1, \{d_i^{(j)} \Delta_j / W_i\}_{i \in [L], j \in [Q_{\text{kq}}]} \right\}$$

where $d_i^{(j)} \in \{0, 1\}$ is the i -th entry of $\mathbf{d}^{(j)}$. Note that S_1 and S_2 correspond to handles for elements in \mathbb{G}_1 and \mathbb{G}_2 , respectively. We then define S_T as $S_T := \{X \cdot Y : X \in S_1, Y \in S_2, X \cdot Y \neq 0\}$. If we explicitly write down S_T , we have $S_T = S_{T,1} \cup S_{T,2}$ where

$$S_{T,1} := \left\{ \begin{array}{ll} 1, & \\ W_i, C_i \Gamma W_i, & \text{for } i \in [L], \\ \Delta_j, & \text{for } j \in [Q_{\text{kq}}], \\ \Delta_j / W_i, & \text{for } i \in [L], j \in [Q_{\text{kq}}] \text{ such that } d_i^{(j)} = 1, \\ \Delta_j W_{i'} / W_i, & \text{for } i, i' \in [L], j \in [Q_{\text{kq}}] \text{ such that } i \neq i' \text{ and } d_i^{(j)} = 1 \\ C_{i'} \Gamma \Delta_j W_{i'} / W_i & \text{for } i, i' \in [L], j \in [Q_{\text{kq}}] \text{ such that } i \neq i' \text{ and } d_i^{(j)} = 1 \end{array} \right\}$$

and $S_{T,2} = \{C_i \Gamma \Delta_j \text{ for } i \in [L], j \in [Q_{\text{kq}}] \text{ such that } d_i^{(j)} = 1\}$. Here, $S_{T,2}$ consists of terms that are obtained by taking product between matching positions of the ciphertext and secret keys, whereas $S_{T,1}$ consists of terms that are obtained by taking product between unmatching positions of the ciphertext and secret keys or between master public key and the ciphertext or secret keys. Note that any handle submitted to the zero-test oracle by A during the game refers to an element f in \mathbb{T} that can be represented as

$$f(W_1, \dots, W_L, \Delta_1, \dots, \Delta_{Q_{\text{kq}}}, \Gamma, C_1, \dots, C_L) = \sum_{Z \in S_T} a_Z Z \quad (4.1)$$

where the coefficients $\{a_Z \in \mathbb{Z}_q\}_{Z \in S_T}$ can be efficiently computed. Furthermore, $\{a_Z \in \mathbb{Z}_q\}_{Z \in S_T}$ satisfying the above equation is unique since all monomials in S_T are distinct.

Game₃: In this game, we change the game so that $\{W_i\}_{i \in [L]}$, $\{\Delta_j\}_{j \in [Q_{\text{kq}}]}$, Γ are treated as formal variables rather than elements in \mathbb{Z}_q even when answering zero-test queries. Namely, the challenger no longer samples $\{w_i\}_{i \in [L]}$, $\{\delta_j\}_{j \in [Q_{\text{kq}}]}$, and γ at the beginning of the game and when A makes a zero-test query for a handle corresponding to $f(W_1, \dots, W_L, \Delta_1, \dots, \Delta_{Q_{\text{kq}}}, \Gamma, C_1, \dots, C_L) \in \mathbb{T}$, the challenger returns 1 if

$$f(W_1, \dots, W_L, \Delta_1, \dots, \Delta_{Q_{\text{kq}}}, \Gamma, c_1, \dots, c_L) = 0 \quad (4.2)$$

holds over \mathbb{T} and 0 otherwise, where $\{c_i\}_{i \in [L]}$ are sampled at the beginning of the game as specified in the previous game. As we prove in Lemma 4.4, we have $|\Pr[\mathbf{E}_2] - \Pr[\mathbf{E}_3]| \leq Q_{\text{zt}}(L+3)^2/q$.

Game₄: This game is the same as the previous game except that the challenger aborts the game and enforces the adversary to output a random bit when there exists $i \in [L]$ such that $c_i = 0$, where $\mathbf{c} = (c_1, \dots, c_L)^\dagger$ is sampled as in the previous game. As we prove in Lemma 4.5, we have $|\Pr[\mathbf{E}_3] - \Pr[\mathbf{E}_4]| \leq L/q$.

Game₅: In this game, we further change the way zero-test queries are answered. In particular, when A makes a zero-test query for a handle corresponding to $f \in \mathbb{T}$ that can be represented as

$$f(W_1, \dots, W_L, \Delta_1, \dots, \Delta_{Q_{kq}}, \Gamma, C_1, \dots, C_L) = \sum_{Z \in S_{T,1}} a_Z Z + \sum_{Z \in S_{T,2}} a_Z Z, \quad (4.3)$$

the challenger returns 0 if there exists $Z \in S_{T,1}$ such that $a_Z \neq 0$. Otherwise, the challenger answers the query as in the previous game. As we prove in Lemma 4.6, we have $\Pr[E_4] = \Pr[E_5]$.

Game₆: In this game, we change the game so that zero-test queries are performed over the ring \mathbb{T} . Namely, when A makes a zero-test query for a handle corresponding to $f \in \mathbb{T}$ the challenger returns 0 if $f \neq 0$ over \mathbb{T} . Equivalently, the challenger returns 0 if there exists $Z \in S_T$ such that $a_Z = 0$ when A makes a zero-test query for a handle corresponding to $f \in \mathbb{T}$ that is represented as Eq. (4.1). Note that (c_1, \dots, c_L) is not used in this game and the challenger does not have to sample it any more. As we prove in Lemma 4.7, there exists a PPT adversary B such that $|\Pr[E_5] - \Pr[E_6]| \leq Q_{kq} Q_{zt} \cdot (\text{Adv}_{\text{BGG}^+, \text{B}}^{\text{Ada-INDr}}(1^\lambda) + 1/q)$.

We can see that the adversary cannot obtain any information about the encrypted message b in **Game₆** since the challenge ciphertext is replaced by formal variables (C_1, \dots, C_L) that does not contain any information of b and the answers to the zero test queries do not depend on b neither. Therefore, we have $\Pr[E_6] = 1/2$. Thus, there exists a PPT adversary B against Ada-INDr security of BGG^+ such that

$$\left| \Pr[E_0] - \frac{1}{2} \right| \leq Q_{kq} Q_{zt} \cdot \left(\text{Adv}_{\text{BGG}^+, \text{B}}^{\text{Ada-INDr}}(1^\lambda) + \frac{1}{q} \right) + \frac{Q_{zt}(L+3)^2 + L}{q}.$$

In particular, assuming BGG^+ satisfies Ada-INDr security, the above quantity is negligible as desired.

To finish the proof of Theorem 4.1, it remains to prove Lemma 4.2, 4.3, 4.4, 4.5, 4.6, and 4.7 in the following.

Lemma 4.2 (Game₀ \equiv Game₁). *We have $\Pr[E_0] = \Pr[E_1]$.*

Proof. Since this is only a conceptual change, the lemma immediately follows.

Lemma 4.3 (Game₁ \equiv Game₂). *We have $\Pr[E_1] = \Pr[E_2]$.*

Proof. Since zero-test queries in **Game₂** are answered by using $\{w_i\}_i, \{\delta_j\}_j, \gamma$, and $\{c_i\}_i$ that are sampled from exactly the same distribution as that in **Game₁**, the view of A in **Game₂** is not altered from that in **Game₁**. The lemma therefore follows.

Lemma 4.4 (Game₂ \approx_s Game₃). *We have $|\Pr[E_2] - \Pr[E_3]| \leq Q_{zt}(L+3)^2/q$.*

Proof. Let us observe that **Game₂** and **Game₃** differ only when A submits a handle corresponding to a polynomial $f(W_1, \dots, W_L, \Delta_1, \dots, \Delta_{Q_{kq}}, \Gamma, C_1, \dots, C_L) \in \mathbb{T}$

satisfying $f(w_1, \dots, w_L, \delta_1, \dots, \delta_{Q_{kq}}, \gamma, c_1, \dots, c_L) = 0$ and $f(W_1, \dots, W_L, \Delta_1, \dots, \Delta_{Q_{kq}}, \Gamma, c_1, \dots, c_L) \neq 0$ to the zero-test oracle. Let F denote the event. It suffices to bound the probability of F occurring in \mathbf{Game}_2 . To do so, let us fix an element f in \mathbb{T} and c_1, \dots, c_L in \mathbb{Z}_q . We then define a polynomial $g(W_1, \dots, W_L, \Delta_1, \dots, \Delta_{Q_{kq}}, \Gamma) \in \mathbb{Z}_q[W_1, \dots, W_L, \Delta_1, \dots, \Delta_{Q_{kq}}, \Gamma]$ as

$$g(W_1, \dots, W_L, \Delta_1, \dots, \Delta_{Q_{kq}}, \Gamma) := \left(\prod_{i \in [L]} W_i \right) \cdot f(W_1, \dots, W_L, \Delta_1, \dots, \Delta_{Q_{kq}}, \Gamma, c_1, \dots, c_L).$$

Note that in the above, the term $(\prod_i W_i)$ is introduced in order to clear the denominators that possibly appear in f and to make sure that g is in the ring $\mathbb{Z}_q[W_1, \dots, W_L, \Delta_1, \dots, \Delta_{Q_{kq}}, \Gamma]$ rather than in \mathbb{T} . We observe that F occurs if and only if $g(w_1, \dots, w_L, \delta_1, \dots, \delta_{Q_{kq}}, \gamma) = 0$ and $g(W_1, \dots, W_L, \Delta_1, \dots, \Delta_{Q_{kq}}, \Gamma) \neq 0$ since we have $w_i \neq 0$ for all $i \in [L]$. We can bound this probability by $(L+3)^2/q$ using Schwartz-Zippel lemma since g is a polynomial in $\mathbb{Z}_q[W_1, \dots, W_L, \Delta_1, \dots, \Delta_{Q_{kq}}, \Gamma]$ with degree at most $L+3$. (Recall that f can be represented as a linear combination of the terms in S_T .) Since A makes at most Q_{zt} zero-test queries, the lemma follows by the union bound.

Lemma 4.5 ($\mathbf{Game}_3 \approx_s \mathbf{Game}_4$). *We have $|\Pr[E_3] - \Pr[E_4]| \leq L/q$.*

Proof. We observe that each entry of $\mathbf{c} = (c_1, \dots, c_L)$ is either fixed to be 1 or distributed uniformly at random over \mathbb{Z}_q . Therefore, by the union bound, the probability that there is $i \in [L]$ such that $c_i = 0$ can be bounded by L/q . The lemma therefore follows.

Lemma 4.6 ($\mathbf{Game}_4 \equiv \mathbf{Game}_5$). *We have $\Pr[E_4] = \Pr[E_5]$.*

Proof. We observe that \mathbf{Game}_4 and \mathbf{Game}_5 differ only when A makes a zero-test query for a handle corresponding to $f \in \mathbb{T}$ that satisfies Eq. (4.2) and there exists $Z \in S_{T,1}$ such that $a_Z \neq 0$ when we express f as Eq. (4.3). We claim that such f does not exist and two games are actually equivalent. For the sake of contradiction, assume that such f exists. Then Eq. (4.2) implies

$$\sum_{Z \in S_{T,1}} a_Z Z(c_1, \dots, c_L) + \sum_{Z \in S_{T,2}} a_Z Z(c_1, \dots, c_L) = 0,$$

where $Z(c_1, \dots, c_L)$ denotes $Z(W_1, \dots, W_L, \Delta_1, \dots, \Delta_{Q_{kq}}, \Gamma, c_1, \dots, c_L) \in \mathbb{T}$ in the above. We can see that $\sum_{Z \in S_{T,1}} a_Z Z(c_1, \dots, c_L) = 0$ holds since we have

$$\left\{ \sum_{Z \in S_{T,1}} a'_Z Z(c_1, \dots, c_L) : a'_Z \in \mathbb{Z}_q \right\} \cap \left\{ \sum_{Z \in S_{T,2}} a''_Z Z(c_1, \dots, c_L) : a''_Z \in \mathbb{Z}_q \right\} = \{0\},$$

which follows from the fact that monomials in $S_{T,1}$ and $S_{T,2}$ are distinct even if we substitute $\{C_i\}_i$ in $S_{T,1}$ and $S_{T,2}$ with $\{c_i\}_i$ and ignore the difference between the

coefficients of the monomials. Furthermore, $\sum_{Z \in S_{T,1}} a_Z Z(c_1, \dots, c_L) = 0$ implies $a_Z = 0$ for all $Z \in S_{T,1}$, which follows from $\mathbf{c} \in (\mathbb{Z}_q^*)^L$ and from the fact that all monomials in $S_{T,1}$ are distinct even if we substitute $\{C_i\}_i$ with $\{c_i\}_i$ and ignore the difference between the coefficients of the monomials. However, this contradicts the assumption that there exists $Z \in S_{T,1}$ such that $a_Z \neq 0$. This completes the proof of the lemma.

Lemma 4.7 (Game₅ \approx_c Game₆). *There exists a PPT adversary B such that $|\Pr[E_5] - \Pr[E_6]| \leq Q_{\text{kq}} Q_{\text{zt}} \cdot (\text{Adv}_{\text{BGG}^+, \text{B}}^{\text{Ada-INDr}}(1^\lambda) + 1/q)$.*

Proof. We first observe that **Game₅** and **Game₆** differ only when A makes a zero-test query for a handle corresponding to $f \in \mathbb{T}$ that can be represented as

$$f(W_1, \dots, W_L, \Delta_1, \dots, \Delta_{Q_{\text{kq}}}, \Gamma, C_1, \dots, C_L) = \sum_{Z \in S_{T,2}} a_Z Z \quad (4.4)$$

and satisfies $f \neq 0$ over \mathbb{T} and Eq. (4.2). We call such a query *bad*. In the following, we prove that the probability that A makes a bad query in **Game₅** is negligible. To do so, we consider following sequence of games. We define F_x as the event that A makes a bad query in **Game_{5,x}** and the challenger does not abort.

Game_{5,0}: This game is the same as **Game₅**. By definition, the probability that A makes a bad query in **Game₅** is $\Pr[F_0]$.

Game_{5,1}: 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 aborts 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 A , the guess is correct with probability $1/Q_{\text{zt}}$ conditioned on F_0 . Therefore, we have $\Pr[F_1] = \Pr[F_0]/Q_{\text{zt}}$.

Game_{5,2}: This game is the same as the previous game except that the challenger aborts the game immediately after A makes the k^* -th zero-test query. Since whether F_1 occurs or not is irrelevant to how the game proceeds after the k^* -th zero-test query is made by A , we clearly have $\Pr[F_2] = \Pr[F_1]$.

Game_{5,3}: In this game, we change the game so that the challenger answers the first $k^* - 1$ zero-test queries by performing zero tests over \mathbb{T} . Furthermore, we change the game so that the sampling of \mathbf{c} is deferred until the k^* -th zero-test query is made by A . We first observe that the game is well-defined since \mathbf{c} is used only for the k^* -th zero-test query. Furthermore, since the first $k^* - 1$ zero-test queries that refer to $f \in \mathbb{T}$ such that $f \neq 0$ are answered by 0 whenever F_2 happens, we have $\Pr[F_3] \geq \Pr[F_2]$.

Game_{5,4}: To define the game, we first define the set $S_{T,2,j} := \{C_i \Gamma \Delta_j\}_{i \in [L]}$ s.t. $d_i^{(j)} = 1$. By definition, we have $S_{T,2} = \cup_{j \in [Q_{\text{kq}}]} S_{T,2,j}$. Using this notation, any $f \in \mathbb{T}$ referred by a bad query can be represented as

$$f(W_1, \dots, W_L, \Delta_1, \dots, \Delta_{Q_{\text{kq}}}, \Gamma, C_1, \dots, C_L) = \sum_{j \in [Q_{\text{kq}}]} \sum_{Z \in S_{T,2,j}} a_Z Z. \quad (4.5)$$

In this game, we change the game so that the challenger aborts the game if the bad query made by A refers to f such that there *does not* exist $j \in [Q_{\text{kq}}]$ satisfying

$$\sum_{Z \in S_{T,2,j}} a_Z Z \neq 0 \quad \text{and} \quad \sum_{Z \in S_{T,2,j}} a_Z Z(c_1, \dots, c_L) = 0, \quad (4.6)$$

where $Z(c_1, \dots, c_L)$ denotes $Z(W_1, \dots, W_L, \Delta_1, \dots, \Delta_{Q_{\text{kq}}}, \Gamma, c_1, \dots, c_L) \in \mathbb{T}$ above. We claim that this actually cannot happen. To see this, we first observe that since we have $f \neq 0$ for a bad query, there exists $j \in [Q_{\text{kq}}]$ satisfying $\sum_{Z \in S_{T,2,j}} a_Z Z \neq 0$. Furthermore, we have

$$\sum_{Z \in S_{T,2,j}} a_Z Z(c_1, \dots, c_L) = - \sum_{j' \neq j} \sum_{Z \in S_{T,2,j'}} a_Z Z(c_1, \dots, c_L)$$

from Eq. (4.2). However, the above is impossible unless the left hand side equals to 0 since any monomial in $S_{T,2,j}$ never appears in $S_{T,2,j'}$ for $j' \neq j$ even if we replace $\{C_i\}_i$ with $\{c_i\}_i$ and ignore the difference between the coefficients of the monomials. Therefore, the change made in this game is only conceptual and we have $\Pr[F_4] = \Pr[F_3]$.

Game_{5,5}: In this game, we change the previous game so that the challenger picks $j^* \leftarrow [Q_{\text{kq}}]$ uniformly at random at the beginning of the game. Furthermore, we add the abort condition that the challenger aborts if Eq. (4.6) does not hold with respect to $j = j^*$ for f that is referred by the k^* -th zero-test query. Since there exists $j' \in [Q_{\text{kq}}]$ that satisfies Eq. (4.6) as long as F_4 occurs and j^* is chosen uniformly at random and independent from the view of A, we have $\Pr[F_5] \geq \Pr[F_4]/Q_{\text{kq}}$.

Game_{5,6}: In this game, we further change the game so that the challenger aborts the game if the j^* -th key query has not been made yet at the point when the k^* -th zero-test query is made. We claim that conditioned on F_5 happens, the challenger never aborts. To see this, we observe that if the j^* -th key query has not been made then terms that contain Δ_{j^*} has not been given to A and there is no way to make a zero-test query for f such that $\sum_{Z \in S_{T,2,j^*}} a_Z Z \neq 0$, since all terms in $S_{T,2,j^*}$ are multiples of Δ_{j^*} . We therefore have $\Pr[F_6] = \Pr[F_5]$.

Game_{5,7}: In this game, we further change the game so that the challenger samples c_i only for $i \in [L]$ such that $d_i^{(j^*)} = 1$, where j^* is chosen at the beginning of the game as in **Game_{5,5}**. The game is still well-defined since the only place in the game where we need the information of \mathbf{c} is when checking Eq. (4.6) and we only need $\{c_i\}_{i \in [L] \text{ s.t. } d_i^{(j^*)} = 1}$ there. (Recall that we have $S_{T,2,j} = \{C_i \Gamma \Delta_j\}_{i \in [L] \text{ s.t. } d_i^{(j)} = 1}$.) Clearly, this does not change the view of A. We therefore have $\Pr[F_7] = \Pr[F_6]$.

From Eq. (3.1) and (3.2), we can see that $\{c_i\}_{i \in [L] \text{ s.t. } d_i^{(j^*)} = 1}$ consists of the following components:

$$\begin{aligned} \psi_0 &= 1, \quad \psi_1 := \mathbf{s}^\top \mathbf{u} + e_1 + \mu[q/2], \quad \psi_2^\top := \mathbf{s}^\top \mathbf{A} + \mathbf{e}_2^\top, \\ \psi_{i,x_i^{(j^*)}}^\top &:= \mathbf{s}^\top (\mathbf{B}_i - x_i^{(j^*)} \mathbf{G}) + \mathbf{e}_{i,x_i^{(j^*)}}^\top \quad \text{for } i \in [\ell], \end{aligned}$$

where $x_i^{(j^*)}$ is the i -th entry of $\mathbf{x}^{(j^*)}$.

Game_{5,8}: In this game, we further change the game so that the challenger samples

$$\psi_0 := 1 \in \mathbb{Z}_q, \quad \psi_1 \leftarrow \mathbb{Z}_q, \quad \psi_2 \leftarrow \mathbb{Z}_q^m, \quad \psi_{i,b} \leftarrow \mathbb{Z}_q^m \quad \text{for } i \in [\ell] \text{ and } b \in \{0, 1\}$$

and sets $\{c_i\}_{i \in [L]}$ s.t. $d_i^{(j^*)} = 1$ from the above components.⁸ As we prove in Lemma 4.8, there exists a PPT adversary B such that $\text{Adv}_{\text{BGG}^+, B}^{\text{Ada-INDr}}(1^\lambda) \geq |\Pr[F_7] - \Pr[F_8]|$.

As we will prove in Lemma 4.9, we have $\Pr[F_8] \leq 1/q$. This allows us to bound $\Pr[F_0]$ as $\Pr[F_0] \leq Q_{\text{kq}} Q_{\text{zt}} \cdot (\text{Adv}_{\text{BGG}^+, B}^{\text{Ada-INDr}}(1^\lambda) + 1/q)$, where B is a PPT adversary. This completes the proof of Lemma 4.7.

It remains to prove Lemma 4.8 and 4.9 in the following.

Lemma 4.8 (Game_{5,7} \approx_c Game_{5,8}). *There exists a PPT adversary B such that $\text{Adv}_{\text{BGG}^+, B}^{\text{Ada-INDr}}(1^\lambda) \geq |\Pr[F_7] - \Pr[F_8]|$.*

Proof. We show that if A can distinguish Game_{5,7} from Game_{5,8}, we can build another adversary B against Ada-INDr security of BGG^+ . The adversary B acts as the challenger and simulates the game for A . Looking ahead, setup phase and key queries are trivial to handle since they do not need any parameter of BGG^+ . The only steps we need care are the simulation of the challenge phase and the k^* -th zero-test query, where B needs to interact with its challenger in order to handle them. We describe how B proceeds in the following.

Setup phase. At the beginning of the game, B is given 1^λ and the master public key of BGG^+ $(\mathbf{A}, \mathbf{B}, \mathbf{u})$. It then gives the handles to $1, W_1, \dots, W_L$ corresponding to \mathbb{G}_1 and the handle to 1 corresponding to \mathbb{G}_2 to A . These handles correspond to the master public key. B also samples $j^* \leftarrow [Q_{\text{kq}}], k^* \leftarrow [Q_{\text{zt}}]$, and $b \leftarrow \{0, 1\}$ and keeps them secret.

Key Queries. Given the j -th secret key query for $\mathbf{x}^{(j)}$ made by A , B proceeds as follows. B first forms $\mathbf{d}^{(j)} \in \mathbb{Z}_q^L$ as specified in the key generation algorithm and returns the handles corresponding to $(d_1^{(j)} \Delta_j / W_1, \dots, d_L^{(j)} \Delta_j / W_L)$ in \mathbb{G}_2 to A .

Challenge Query. When A makes the challenge query for a circuit F , B makes a *secret key query* for F to its challenger and is given \mathbf{r} sampled as $\mathbf{r} \leftarrow [\mathbf{A} \parallel \mathbf{B}_F]_{\tau}^{-1}(\mathbf{u})$. B then sets $\text{ct}_0 = (\mathbf{A}, \mathbf{B})$, $\text{ct}_2 := \mathbf{r}$, and ct_1 as the handles corresponding to the formal variables (C_1, \dots, C_L) and gives $\text{ct} = (\text{ct}_0, \text{ct}_1, \text{ct}_2)$ to A as the challenge ciphertext.

Generic Group Queries. B honestly handles the queries for the generic group oracle corresponding to addition, negation, and multiplication (bilinear map) made by A by keeping track of the underlying encodings in \mathbb{T} associated with the handles. For the k -th zero-test query that refers to an element f in \mathbb{T} , B returns 1 if $f = 0$ over \mathbb{T} and 0 otherwise if $k < k^*$. If $k = k^*$, B first checks whether the j^* -th key query has already been made and aborts otherwise, as specified in Game_{5,6}. It then makes the

⁸ Note that until this step, we have not changed the distribution of $\{c_i\}_{i \in [L]}$ except that we stop sampling c_i for i such that $d_i^{(j^*)} = 1$ in Game_{5,7}.

challenge query for the attribute $\mathbf{x}^{(j^*)}$, where $\mathbf{x}^{(j^*)}$ is the attribute for which A has made the j^* -th key query, and the message b to its challenger. Then B obtains its challenge ciphertext $(\psi_1, \psi_2, \{\psi_{i, x_i^{(j^*)}}\}_{i \in [\ell]})$. It then sets $\psi_0 = 1$ and forms $\{c_i\}_{i \in [L]}$ s.t. $d_i^{(j^*)} = 1$ by vectorizing the terms appropriately. Finally, it checks whether Eq. (4.6) holds or not using $\{c_i\}_{i \in [L]}$ s.t. $d_i^{(j^*)} = 1$ as specified in $\mathbf{Game}_{5.7}$ and outputs 1 if it holds and 0 otherwise.

Analysis. It is easy to see that B simulates $\mathbf{Game}_{5.7}$ if the challenge ciphertext for B is the real one and $\mathbf{Game}_{5.8}$ if it is chosen uniformly at random from the ciphertext space. Therefore, it can be seen that B outputs 1 with probability $\Pr[\mathbf{F}_7]$ if the challenge bit for B is 0 and $\Pr[\mathbf{F}_8]$ otherwise. Therefore, B's advantage against \mathbf{BGG}^+ is $|\Pr[\mathbf{F}_7] - \Pr[\mathbf{F}_8]|$. This completes the proof of the lemma.

Lemma 4.9. *We have $\Pr[\mathbf{F}_8] = 1/q$.*

Proof. We observe that \mathbf{F}_8 occurs only when A makes a zero-test query that refers to a handle $f \neq 0$ that can be represented as Eq. (4.4) and satisfies Eq. (4.6) with respect to j^* where $\{c_i\}_{i \in [L]}$ s.t. $d_i^{(j^*)} = 1$ are chosen as $\mathbf{Game}_{5.8}$. However, Eq. (4.6) can happen only with probability at most $1/q$ since f is represented as a linear combination of $\{C_i \Gamma \Delta_j\}_{i,j}$ and all entries of $\{c_i\}_{i \in [L]}$ s.t. $d_i^{(j^*)} = 1$ are chosen uniformly at random except for the entry that is fixed to be 1.

5 Implications to CP-ABE, BE, and IBBE

In this section, we show that by setting the circuit class supported by our CP-ABE scheme in Sec. 3 appropriately, we can obtain various new schemes with different security and efficiency tradeoffs. In particular, we obtain new CP-ABE, BE, and IBBE schemes from the LWE assumption in the bilinear generic group model. Our CP-ABE scheme achieves the notable efficiency property that the sizes of all the parameters in the system do not depend on the size of the circuits supported by the scheme. Similarly, our BE (resp., IBBE) schemes achieve optimal parameter size, in the sense that the sizes of all parameters in the system are bounded by a fixed polynomial that is independent from the number of users (resp., upper bound on the number of recipients). These efficiency properties have never been achieved without using indistinguishability obfuscation or multilinear maps.

5.1 New CP-ABE scheme

By setting $\mathcal{F}_{\text{CP}} := \{C_{\ell(\lambda), d(\lambda)}\}_\lambda$ in the construction in Sec. 3, we obtain a CP-ABE scheme that can deal with the set of circuits whose input length and depth are $\ell(\lambda)$ and $d(\lambda)$, respectively. In order to prove Ada-IND security for the resulting scheme, we need to be able to prove Ada-INDr security for the KP-ABE scheme \mathbf{BGG}^+ for the same circuit class as stated in Theorem 4.1. This is possible by assuming subexponential hardness of LWE as we see in Theorem 2.10. The notable feature of the resulting scheme is that the sizes of the master public key, ciphertexts, and secret keys are independent

from the size of the circuits supported by the scheme. The sizes of these parameters are only dependant on the input length and the depth of the circuits.

Summarizing the above discussion, we get the following theorem.

Theorem 5.1. *Assuming the subexponential hardness of LWE, we have a CP-ABE scheme for circuit class $\mathcal{C}_{\ell,d}$ for arbitrary $\ell = \text{poly}(\lambda)$ and $d = O(\log \lambda)$ that satisfies Ada-IND security in the bilinear generic group model. The sizes of the master public key, ciphertexts, and secret keys are bounded by $\text{poly}(\lambda, \ell, d)$.*

We note that in all previous CP-ABE scheme (e.g., [BSW07, Wat11, RW13]) for NC_1 , either the ciphertext or secret key size depends on the circuit size supported by the scheme.

5.2 New BE scheme with Optimal Parameter Size

Here, we show that we can obtain a BE scheme with optimal parameter size by setting the circuit class \mathcal{F} supported by the CP-ABE scheme in Sec. 3 appropriately.

Obtaining DBE from KP-ABE. In order to get the BE scheme, we first observe that we can implement a DBE scheme by a KP-ABE scheme for the following circuit class \mathcal{F}_{BE} defined as $\mathcal{F}_{\text{BE}} = \{F_S : \{0, 1\}^{\lceil \log N \rceil} \rightarrow \{0, 1\}\}_{S \subseteq [N]}$ where $F_S(i) = \begin{cases} 1 & \text{if } i \in S \\ 0 & \text{if } i \notin S \end{cases}$.

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 $O(\log |S|) \leq O(\log N)$. By the definition of F_S , one can see that this KP-ABE scheme implements the functionality of DBE.

Plugging the DBE into Our Construction in Sec. 3. We then instantiate the KP-ABE for the circuit class \mathcal{F}_{BE} with BGG^+ and plug this scheme into our CP-ABE construction in Sec. 3. Since the ciphertext and key attributes of the CP-ABE scheme are swapped from the underlying KP-ABE scheme, we obtain a BE scheme as a result. This instantiation is possible since the depth of the circuits is bounded by $O(\log N) \leq O(\log \lambda)$ and we can take the upper bound on the depth $d(\lambda)$ to be larger than this. The sizes of the master public key, ciphertexts, and secret keys in the resulting BE scheme are bounded by $\text{poly}(\log N, \lambda) = \text{poly}(\lambda)$, which is independent of the number of users, since the depth and input length of the circuits in \mathcal{F}_{BE} is bounded by $O(\log N)$. Note that we crucially rely on the efficiency property of our CP-ABE scheme that the sizes of all parameters in the system are independent of the size of the circuits being supported, where the latter can be as large as $O(N)$ for \mathcal{F}_{BE} .

Security of the resulting BE Scheme. In order for the resulting BE scheme to have Ada-IND security, we need the underlying KP-ABE scheme BGG^+ to have Ada-INDr security as stated in Theorem 4.1. In the general case where the input length for the circuits is of $\text{poly}(\lambda)$, we need to assume subexponential hardness of LWE to prove

Ada-INDr security for BGG^+ as we see in Theorem 2.10. However, since we restrict the circuit class for BGG^+ to be \mathcal{F}_{BE} here, we can avoid assuming subexponential hardness of LWE and base the security of our scheme on polynomial hardness of LWE. To see this, we first recall that for proving Sel-INDr security for BGG^+ , polynomial hardness of LWE is enough (Theorem 2.10). We then observe that in the special case of DBE, Sel-INDr and Ada-INDr are actually equivalent, since one can guess the target attribute $i^* \in [N]$ chosen by the adversary in the security game with only polynomial security loss.

Summarizing the above discussion, we get the following theorem.

Theorem 5.2. *Assuming the LWE assumption, we have a BE scheme that satisfies Ada-IND security in the bilinear generic group model. The sizes of the master public key, ciphertexts, and secret keys are bounded by a fixed polynomial $\text{poly}(\lambda)$ that is independent of N .*

5.3 New IBBE scheme with Optimal Parameter Size

Here, we show that we can obtain an IBBE scheme with optimal parameter size by setting the circuit class \mathcal{F} supported by the CP-ABE scheme in Sec. 3 appropriately. This can be done similarly to the case of BE and we only highlight the difference below.

Obtaining DIBBE from KP-ABE. Similarly to the case of BE, we first observe that we can implement a DIBBE scheme by a KP-ABE scheme by setting the supported circuit class \mathcal{F}_{IBBE} appropriately. In order to do so, we set \mathcal{F}_{IBBE} to be the same as \mathcal{F}_{BE} except that we set N to be $2^{O(\lambda)}$ and the identity space to be $[N]$. We can implement the circuit F_S with depth $O(\log \log N + \log t)$. (Recall that the size of S is bounded by t .) To support identities with arbitrary length, we hash an identity into an integer in $[N]$ by a collision resistant hash function.

Plugging the DIBBE into Our Construction in Sec. 3. We then instantiate the KP-ABE scheme for \mathcal{F}_{IBBE} with BGG^+ and plug the scheme into our construction of CP-ABE in Sec. 3. Note that the instantiation is possible since the depth of the circuits is bounded by $O(\log t + \log \log N) \leq (\log \lambda)$ and we can take the bound on the depth $d(\lambda)$ to be larger than this. As a result, we obtain an IBBE scheme whose sizes of the master public key, ciphertexts, and secret keys are bounded by $O(\lambda)$.

Security of the resulting IBBE Scheme. In order for the resulting IBBE scheme to have Ada-IND security, we need the underlying KP-ABE scheme BGG^+ to have Ada-INDr security. Unlike the BE case, we do not have equivalence between Sel-INDr and Ada-INDr here, since N is exponentially large. Instead, we rely on the random oracle model to transform the scheme that is Sel-INDr secure into a scheme with Ada-INDr security. This can be done similarly to the selective-to-adaptive conversion for IBE (Theorem 7.3 in [BB04]) by applying a hash function that is modeled as a random oracle to the identities before they are input to the encryption or key generation algorithms. The reduction from Ada-INDr security to Sel-INDr security is done by guessing when the hash query for the target identity id is made during the game. The guess is correct with noticeable probability since the adversary makes only polynomial number of hash queries and thus the reduction works.

Summarizing the above discussion, we get the following theorem.

Theorem 5.3. *Assuming the LWE assumption, we have an IBBE scheme that satisfies Ada-IND security in the bilinear generic group model and the random oracle model. The sizes of the master public key, ciphertexts, and secret keys are bounded by a fixed polynomial $\text{poly}(\lambda)$ that is independent of t , which is the upper bound of the number of recipients that is specified by the scheme.*

6 Conclusion

We provided the first BE and IBBE constructions with optimal parameters from bilinear maps and pairings. Our techniques are new and demonstrate a novel interplay of LWE and pairings that may be of independent interest. Several interesting questions arise from our work. An immediate question is whether we can modify the construction so as to obtain a proof in the standard model. It would also be interesting to explore if our techniques can have applications in constructing other primitives that have been only constructed from multilinear maps or indistinguishability obfuscation.

Acknowledgement. We thank anonymous reviewers for helpful comments. Shota Yamada is supported by JST CREST Grant Number JPMJCR19F6 and JSPS KAKENHI Grant Number 16K16068.

References

- ABB10a. Shweta Agrawal, Dan Boneh, and Xavier Boyen. Efficient lattice (H)IBE in the standard model. In *EUROCRYPT*, pages 553–572, 2010.
- ABB10b. Shweta Agrawal, Dan Boneh, and Xavier Boyen. Lattice basis delegation in fixed dimension and shorter-ciphertext hierarchical IBE. In *CRYPTO*, pages 98–115, 2010.
- ADGM16. Daniel Apon, Nico Döttling, Sanjam Garg, and Pratyay Mukherjee. Cryptanalysis of indistinguishability obfuscations of circuits over ggh13 . eprint 2016, 2016.
- Agr17. Shweta Agrawal. Stronger security for reusable garbled circuits, new definitions and attacks. In *Crypto*, 2017.
- Agr19. Shweta Agrawal. Indistinguishability obfuscation without multilinear maps: New techniques for bootstrapping and instantiation. In *Eurocrypt*, 2019.
- AJL⁺19. Prabhanjan Ananth, Aayush Jain, Huijia Lin, Christian Matt, and Amit Sahai. Indistinguishability obfuscation without multilinear maps: io from lwe , bilinear maps, and weak pseudorandomness. In *Crypto*, 2019.
- AL10. Nuttapon Attrapadung and Benoît Libert. Functional encryption for inner product: Achieving constant-size ciphertexts with adaptive security or support for negation. In *PKC*, pages 384–402, 2010.
- BB04. Dan Boneh and Xavier Boyen. Efficient selective-id secure identity-based encryption without random oracles. In *EUROCRYPT*, 2004.
- BCFG17. Carmen Elisabetta Zaira Baltico, Dario Catalano, Dario Fiore, and Romain Gay. Practical functional encryption for quadratic functions with applications to predicate encryption. In *CRYPTO*, 2017.
- BFF⁺14. Gilles Barthe, Edvard Fagerholm, Dario Fiore, John C. Mitchell, Andre Scedrov, and Benedikt Schmidt. Automated analysis of cryptographic assumptions in generic group models. In *CRYPTO*, 2014.

- BGG⁺14. Dan Boneh, Craig Gentry, Sergey Gorbunov, Shai Halevi, Valeria Nikolaenko, Gil Segev, Vinod Vaikuntanathan, and Dhinakaran Vinayagamurthy. Fully key-homomorphic encryption, arithmetic circuit ABE and compact garbled circuits. In *EUROCRYPT*, pages 533–556, 2014.
- BGV12. Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (Leveled) fully homomorphic encryption without bootstrapping. In *ITCS*, pages 309–325, 2012.
- BGW05. Dan Boneh, Craig Gentry, and Brent Waters. Collusion resistant broadcast encryption with short ciphertexts and private keys. In *Proceedings of the 25th Annual International Conference on Advances in Cryptology, CRYPTO’05*, 2005.
- BK. Dan Boneh and Sam Kim. Single key cp-abe. Personal Communication.
- BLP⁺13. Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of learning with errors. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing, STOC ’13*. ACM, 2013.
- Bra12. Zvika Brakerski. Fully homomorphic encryption without modulus switching from classical gapsvp. In *CRYPTO*, 2012.
- BSW07. John Bethencourt, Amit Sahai, and Brent Waters. Ciphertext-policy attribute-based encryption. In *IEEE Symposium on Security and Privacy*, pages 321–334, 2007.
- BV11. Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In *FOCS*, 2011.
- BV20. Zvika Brakerski and Vinod Vaikuntanathan. Lattice-inspired broadcast encryption and succinct ciphertext policy abe. Personal communication, 2020.
- BWZ14. Dan Boneh, Brent Waters, and Mark Zhandry. Low overhead broadcast encryption from multilinear maps. In Juan A. Garay and Rosario Gennaro, editors, *Advances in Cryptology – CRYPTO 2014*, 2014.
- BZ17. Dan Boneh and Mark Zhandry. Multiparty key exchange, efficient traitor tracing, and more from indistinguishability obfuscation. *Algorithmica*, 79(4), December 2017.
- CFL⁺. Jung Hee Cheon, Pierre-Alain Fouque, Changmin Lee, Brice Minaud, and Hansol Ryu. Cryptanalysis of the new clt multilinear map over the integers. Eprint 2016/135.
- CGH⁺15. Jean-Sébastien Coron, Craig Gentry, Shai Halevi, Tancrede Lepoint, Hemanta K Maji, Eric Miles, Mariana Raykova, Amit Sahai, and Mehdi Tibouchi. Zeroizing without low-level zeroes: New mmap attacks and their limitations. In *Advances in Cryptology–CRYPTO 2015*, pages 247–266. Springer, 2015.
- CHKP10. David Cash, Dennis Hofheinz, Eike Kiltz, and Chris Peikert. Bonsai trees, or how to delegate a lattice basis. In *EUROCRYPT*, pages 523–552, 2010.
- CHL⁺15. J.-H. Cheon, K. Han, C. Lee, H. Ryu, and D. Stehlé. Cryptanalysis of the multilinear map over the integers. In *Proc. of EUROCRYPT*, volume 9056 of *LNCS*, pages 3–12. Springer, 2015.
- CJL. Jung Hee Cheon, Jinhyuck Jeong, and Changmin Lee. An algorithm for ntru problems and cryptanalysis of the ggh multilinear map without a low level encoding of zero. Eprint 2016/139.
- CLLT16. Jean-Sébastien Coron, Moon Sung Lee, Tancrede Lepoint, and Mehdi Tibouchi. Zeroizing attacks on indistinguishability obfuscation over clt13. Eprint 2016, 2016.
- Del07. Cecile Delerablee. Identity-based broadcast encryption with constant size ciphertexts and private keys. In *ASIACRYPT*, 2007.
- DPP07. Cécile Delerablée, Pascal Paillier, and David Pointcheval. Fully collusion secure dynamic broadcast encryption with constant-size ciphertexts or decryption keys. In *Proceedings of the First International Conference on Pairing-Based Cryptography, Pairing’07*, 2007.
- FN94. Amos Fiat and Moni Naor. Broadcast encryption. In Douglas R. Stinson, editor, *Advances in Cryptology — CRYPTO’ 93*, 1994.

- GGH⁺13. Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *FOCS*, 2013. <http://eprint.iacr.org/>.
- GPV08. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pages 197–206, 2008.
- GQWW19. Rishab Goyal, Willy Quach, Brent Waters, and Daniel Wichs. Broadcast and trace with n^ϵ ciphertext size from standard assumptions. In *Crypto*, 2019. <https://eprint.iacr.org/2019/636>.
- GV15. Sergey Gorbunov and Dhinakaran Vinayagamurthy. Riding on asymmetry: Efficient ABE for branching programs. In *ASIACRYPT*, 2015.
- GVW13. Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Attribute based encryption for circuits. In *STOC*, 2013.
- GW09. Craig Gentry and Brent Waters. Adaptive security in broadcast encryption systems with short ciphertexts. In *Proceedings of the 28th Annual International Conference on Advances in Cryptology - EUROCRYPT 2009 - Volume 5479*, 2009.
- HJ15. Y. Hu and H. Jia. Cryptanalysis of GGH map. Cryptology ePrint Archive: Report 2015/301, 2015.
- HWL⁺16. Kai He, Jian Weng, Jia-Nan Liu, Joseph K. Liu, Wei Liu, and Robert H. Deng. Anonymous identity-based broadcast encryption with chosen-ciphertext security. In *Proceedings of the 11th ACM on Asia Conference on Computer and Communications Security*, ASIA CCS '16, 2016.
- JLS19. Aayush Jain, Huijia Lin, and Amit Sahai. Simplifying constructions and assumptions for $i\mathcal{O}$. Cryptology ePrint Archive, Report 2019/1252, 2019. <https://eprint.iacr.org/2019/1252>.
- Mau05. Ueli Maurer. Abstract models of computation in cryptography. In *IMA Int. Conf.*, volume 3796 of *Lecture Notes in Computer Science*, pages 1–12. Springer, 2005.
- MP12. Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *Advances in Cryptology - EUROCRYPT 2012 - 31st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cambridge, UK, April 15-19, 2012. Proceedings*, pages 700–718, 2012.
- MSZ16. Eric Miles, Amit Sahai, and Mark Zhandry. Annihilation attacks for multilinear maps: Cryptanalysis of indistinguishability obfuscation over ggh13. In *Crypto*, 2016.
- Reg09. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J.ACM*, 56(6), 2009. extended abstract in STOC'05.
- RW13. Yannis Rouselakis and Brent Waters. Practical constructions and new proof methods for large universe attribute-based encryption. In *ACM-CCS*, pages 463–474, 2013.
- SF. Ryuichi Sakai and Jun Furukawa. Identity-based broadcast encryption. IACR Cryptology ePrint Archive, 2007.
- Sho97. Victor Shoup. Lower bounds for discrete logarithms and related problems. In *Eurocrypt*, pages 256–266. Springer-Verlag, 1997.
- Tsa19. Rotem Tsabary. Fully secure attribute-based encryption for t-cnf from LWE. In *CRYPTO*, 2019.
- Wat11. Brent Waters. Ciphertext-policy attribute-based encryption: An expressive, efficient, and provably secure realization. In *PKC*, pages 53–70, 2011.