

Adaptively Secure Identity-Based Encryption from Lattices with Asymptotically Shorter Public Parameters

Shota Yamada^{*1}

¹National Institute of Advanced Industrial Science and Technology (AIST).

February 16, 2016

Abstract

In this paper, we present two new adaptively secure identity-based encryption (IBE) schemes from lattices. The size of the public parameters, ciphertexts, and private keys are $\tilde{O}(n^2\kappa^{1/d})$, $\tilde{O}(n)$, and $\tilde{O}(n)$ respectively. Here, n is the security parameter, κ is the length of the identity, and $d \in \mathbb{N}$ is a flexible constant that can be set arbitrary (but will affect the reduction cost). Ignoring the poly-logarithmic factors hidden in the asymptotic notation, our schemes achieve the best efficiency among existing adaptively secure IBE schemes from lattices. In more detail, our first scheme is anonymous, but proven secure under the LWE assumption with approximation factor $n^{\omega(1)}$. Our second scheme is not anonymous, but proven adaptively secure assuming the LWE assumption for all polynomial approximation factors.

As a side result, based on a similar idea, we construct an attribute-based encryption scheme for branching programs that simultaneously satisfies the following properties for the first time: Our scheme achieves compact secret keys, the security is proven under the LWE assumption with polynomial approximation factors, and the scheme can deal with unbounded length branching programs.

1 Introduction

Background. Identity-based encryption (IBE) is an advanced form of public key encryption (PKE) where any string such as an email address can be used as a public key. The notion of IBE was proposed by Shamir in 1984 [42]. Since then, it took nearly 20 years for the first realizations of IBE [41, 10, 18] to appear. Boneh and Franklin [10] and Sakai, Ohgishi, and Kasahara [41] used groups equipped with efficiently computable bilinear maps to construct the first IBE. On the other hand, Cocks [18] used quadratic residue for a composite modulus. These constructions are only proven secure in the random oracle model. In subsequent works, pairing-based schemes in the standard model appeared [15, 8, 9, 47, 48]. While earlier works [15, 8] focus on the constructions that are only selectively secure, later works [9, 47, 48] focus on a much more realistic security, i.e., adaptive security.

Another important line of research is construction of IBE from lattices. The first lattice-based IBE was proposed in the seminal work by Gentry, Peikert, and Vaikuntanathan [25] in the random oracle model. Later, constructions in the standard model were proposed [1, 16, 12]. To achieve

^{*}yamada-shota@aist.go.jp

adaptive security in the lattice-based settings, we have to either rely on an analogue of Waters’ hash [47] or an admissible hash [9, 16]. In any case, we require $O(\kappa)$ number of basic matrices in the public parameters (master public key), where κ is the bit length of the identities. This results in very large public parameters with size $\tilde{O}(n^2\kappa)$. Here, n is the security parameter (dimension of the lattices). On the other hand, in the selectively secure variant of lattice IBE in [1], we only require small constant number of basic matrices in the public parameters. This stands in sharp contrast to pairing-based settings, in which we have adaptively secure IBE schemes [17, 31] that are as efficient as selectively secure ones [8], up to only small constant factors. A natural important question is:

Can we construct adaptively secure IBE schemes from lattices, which is as efficient as selectively secure ones? In particular, can we reduce the size of the public parameters?

Difficulties. A natural approach to achieve short public parameters in lattice based IBE schemes would be to mimic the technique for pairing based IBE schemes. However, all IBE schemes with short public parameters based on pairings are constructed using dual system encryption methodology [48], for which there is still no lattice analogue. The realization of the dual system encryption methodology in the lattice settings is an important open problem [38]. Another possible approach would be to use a technique from Naccache’s IBE scheme [36], as is done in [44]. Using this approach, we can obtain a scheme with the public parameters shorter by a factor of u , at the cost of 2^u -loss in security. Therefore, using this approach, we are only allowed to reduce the size of public parameters up to logarithmic factor.

Our Contribution. Instead of taking the above approaches, we use a technique unique to the lattice setting. Namely, we use the fully homomorphic computation of trapdoors, which is recently devised in [11] to reduce the size of the public parameters. We obtain the following two different IBE schemes with trade-off between the security, efficiency, and underlying hardness assumptions. See Table 1 in Section 6 for the overview.

- We propose an adaptively secure and anonymous IBE with asymptotically short parameters. In particular, the size of the public parameters, ciphertexts, and private keys are $\tilde{O}(n^2\kappa^{1/d})$, $\tilde{O}(n)$, and $\tilde{O}(n)$ respectively. Here, $d \in \mathbb{N}$ is a flexible constant which can be set arbitrary. Ignoring poly-logarithmic factors hidden in the asymptotic notation, our scheme achieves the best efficiency among all previous adaptively secure IBE schemes from lattices. The security of the scheme is proven under the LWE assumption with super-polynomial approximation factors.
- We propose an adaptively secure IBE (without anonymity) that achieves asymptotically the same efficiency as the above scheme. The difference from the above scheme is that our scheme can be proven secure assuming the LWE assumption with all polynomial approximation factor. The assumption is weaker than the one used in the above scheme, but the sizes of the public parameters, ciphertexts, and private keys are larger than the above scheme by a super-constant factor.

In the second construction, different from lattice IBE schemes in the literature [1, 16, 2, 12], we have to rely on the LWE assumption for *all* polynomial approximation factors, rather than some *fixed* polynomial approximation factor (e.g., $O(n^3)$). The interesting feature of the reduction is that the problem we reduce the security to varies according to the power of the adversary. More specifically, as the number of key extraction queries grows or as the advantage of the adversary drops, we would need the LWE assumption with larger approximation factor. This is somewhat

similar to the security proof based on the q -type assumptions (e.g., [24]), in which the problem that the reduction algorithm solves depends on the number of key extraction queries made by the adversary. However, unlike the q -type assumptions, our assumptions enjoy reduction to the worst case lattice problems [40, 37, 13].

To present our schemes in a unified manner, we define the new notion of parametrized IBE (PIBE). The syntax of PIBE is the same as that of ordinary IBE except that it is parametrized by a variable c . As for the security, roughly speaking, we require the advantage of any adversary to be at most $1/n^c$ if the number of key extraction queries is bounded by n^c . In the case of c is a super-constant function, the notion of PIBE corresponds to that of (ordinary) IBE. We then construct a specific PIBE scheme from the LWE assumption. By setting c to be a super-constant function, we obtain our first IBE scheme. Our second IBE scheme is obtained by running several instances of the PIBE scheme in parallel with different values of c . This is captured as a generic conversion from PIBE to (ordinary) IBE.

We note that our IBE schemes might not be as efficient as previous adaptively secure lattice IBE schemes [1, 12] for a practical choice of parameters, due to the super-constant factors hidden in the asymptotic notation. However, we believe that our technique would be of theoretical interest. In particular, the security proof of our PIBE scheme is based on the traditional partitioning technique [47] with some novel ideas. In addition, our technique used in the generic construction of IBE from PIBE, inspired by [7], would be useful for other settings.

Other Application of Our Technique. As a side result, we show an application of our technique to attribute-based encryption (ABE). In particular, we obtain the first ABE scheme that simultaneously satisfies the following properties: an unbounded length branching program is usable as an attribute, the sizes of the private keys are compact, the security is proven under the LWE problem for all polynomial approximation factors. We obtain such a scheme by applying a simple conversion to the recent ABE scheme for branching programs by Gorbunov and Vinayagamurthy [28]. The idea for the conversion is similar in spirit to our PIBE-to-IBE conversion. We note that the original ABE scheme of [28] is either based on the super-polynomial LWE while dealing with unbounded length branching programs or based on the polynomial LWE while only dealing with bounded length branching programs. See Appendix B for the details.

Related Works. We can obtain efficient PKE as well as IBE schemes over ideal lattices [45, 22]. By switching to the ring setting, we can generally reduce the size of the public parameters by a factor of $O(n)$. However, we have to rely on the ring LWE (RLWE) assumption [33, 34], which is a stronger assumption than the LWE assumption.

The techniques for constructing IBE and signatures are somewhat similar and related. Indeed, we can obtain secure signature from (adaptively) secure IBE, via the Naor transformation [10]. A construction of short signature with short public parameters from weak assumptions has been an important research topic. This problem has been addressed by several previous works [32, 30, 7, 23, 4]. However, their techniques heavily depend on the fact that we can convert a non-adaptively secure signature scheme into adaptively secure (or equivalently, EUF-CMA secure) one by using chameleon hash functions [43]. There is no known analogue of the conversion in the setting of IBE. We also note that our technique of converting PIBE into IBE is similar to the “on the fly adaptation technique” in [21], which was used to improve the efficiency and the reduction cost of the Naor-Reingold PRF.

2 Overview of Our Technique

2.1 Overview of the Construction

We follow the general framework for constructing lattice-based IBE schemes, which is an abstraction of many existing schemes [16, 1, 2]. In the template, we associate each identity ID with the following matrix:

$$(\mathbf{A}|\mathbf{H}(\text{ID})) \in \mathbb{Z}_q^{n \times (m+m')}$$

where $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and $\mathbf{H}(\cdot)$ is a function that maps an identity to a matrix in $\mathbb{Z}_q^{n \times m'}$ for some $n, m, m' \in \mathbb{N}$ and some prime number q . A ciphertext for an identity ID includes a vector of the following form:

$$\mathbf{s}^\top (\mathbf{A}|\mathbf{H}(\text{ID})) + (\mathbf{x}_1^\top | \mathbf{x}_2^\top)$$

where \mathbf{s} is a random vector in \mathbb{Z}_q^n and $\mathbf{x}_1 \in \mathbb{Z}_q^m$ and $\mathbf{x}_2 \in \mathbb{Z}_q^{m'}$ are small error terms. A private key is a short vector $\mathbf{e} \in \mathbb{Z}^{m+m'}$ that satisfies

$$(\mathbf{A}|\mathbf{H}(\text{ID}))\mathbf{e} = \mathbf{u} \pmod{q}$$

for some fixed $\mathbf{u} \in \mathbb{Z}_q^n$. In the adaptively secure variant of the IBE scheme in [1], the function $\mathbf{H}(\text{ID})$ is defined as

$$\mathbf{H}(\text{ID}) = \mathbf{B}_0 + \sum_{\{i \in [1, \kappa] \mid \text{ID}_i = 1\}} \mathbf{B}_i$$

where $\mathbf{B}_0, \mathbf{B}_1, \dots, \mathbf{B}_\kappa \in \mathbb{Z}_q^{n \times m}$ are matrices that are included in the public parameters and ID_i is the i -th bit of the bit string $\text{ID} \in \{0, 1\}^\kappa$. We typically set $\kappa = O(n)$ and require rather long public parameters $\mathbf{B}_0, \mathbf{B}_1, \dots, \mathbf{B}_\kappa$.

Our first idea is to use the technique called fully homomorphic trapdoor computation, which is introduced in [11], to reduce the size of the public parameters. Namely, we set $\ell = \lceil \sqrt{\kappa} \rceil$ and the public parameters as matrices $\mathbf{B}_{1,1}, \dots, \mathbf{B}_{1,\ell}, \mathbf{B}_{2,1}, \dots, \mathbf{B}_{2,\ell} \in \mathbb{Z}_q^{n \times m}$. We also introduce an injective map $S : \{0, 1\}^\kappa \rightarrow 2^{[\ell] \times [\ell]}$ that maps an identity to a subset of the set $[\ell] \times [\ell]$. Then, we change the definition of the function as

$$\mathbf{H}(\text{ID}) = \mathbf{B}_0 + \sum_{(i,j) \in S(\text{ID})} \mathbf{B}_{1,i} \cdot \mathbf{G}^{-1}(\mathbf{B}_{2,j}),$$

where \mathbf{G} is a gadget matrix whose trapdoor is publicly known [35] and \mathbf{G}^{-1} is a deterministic function* that maps a matrix in $\mathbf{U} = \mathbb{Z}_q^{n \times m}$ to a matrix in $\mathbf{V} = \{0, 1\}^{m \times m}$ such that $\mathbf{G}\mathbf{V} = \mathbf{U}$. By this change, we are able to reduce the number of basic matrices from $O(\kappa)$ to $O(\sqrt{\kappa})$.[†]

2.2 Overview of the Security Proof

We prove the security of the scheme under the LWE assumption. Let the input to the reduction algorithm be $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and $\mathbf{v} \in \mathbb{Z}_q^m$. The task of the algorithm is to distinguish whether

*Note that we are abusing the notation here. \mathbf{G}^{-1} is not an inverse matrix of \mathbf{G} , but a function.

[†]For the sake of simplicity, we present a scheme that is a special case of our scheme in Section 5. More generally, we can further reduce the number of basic matrices from $O(\sqrt{\kappa})$ to be $O(\kappa^{1/d})$ for any constant $d \in \mathbb{N}$.

$\mathbf{v}^\top = \mathbf{s}^\top \mathbf{A} + \mathbf{x}^\top \pmod q$ for some $\mathbf{s} \in \mathbb{Z}_q^n$ and small $\mathbf{x} \in \mathbb{Z}^m$, or, \mathbf{v} is a random vector. In the security proof, we pick random $y_0, y_{1,1}, \dots, y_{1,\ell}, y_{2,1}, \dots, y_{2,\ell} \in \mathbb{Z}_q$ from certain domains, whose sizes grow proportion to the number of key extraction queries Q that the adversary makes (similarly to in [47]). Since we assume that Q is much smaller than q , these random values are bounded by some “small” polynomial. Then, the reduction algorithm picks $\mathbf{R}_0, \mathbf{R}_{i,j} \stackrel{\$}{\leftarrow} \{-1, 1\}^{m \times m}$ and embeds these values into the public parameters as

$$\mathbf{B}_0 = \mathbf{A}\mathbf{R}_0 + y_0\mathbf{G}, \quad \mathbf{B}_{i,j} = \mathbf{A}\mathbf{R}_{i,j} + y_{i,j}\mathbf{G}$$

for $(i, j) \in \{1, 2\} \times [1, \ell]$. Then, we have

$$\begin{aligned} \mathbf{H}(\text{ID}) &= (\mathbf{A}\mathbf{R}_0 + y_0\mathbf{G}) + \sum_{(i,j) \in \mathcal{S}(\text{ID})} (\mathbf{A}\mathbf{R}_{1,i} + y_{1,i}\mathbf{G}) \cdot \mathbf{G}^{-1}(\mathbf{B}_{2,j}) \\ &= (\mathbf{A}\mathbf{R}_0 + y_0\mathbf{G}) + \sum_{(i,j) \in \mathcal{S}(\text{ID})} (\mathbf{A}\mathbf{R}_{1,i}\mathbf{G}^{-1}(\mathbf{B}_{2,j}) + y_{1,i}\mathbf{B}_{2,j}) \\ &= \mathbf{A} \left(\underbrace{\mathbf{R}_0 + \sum_{(i,j) \in \mathcal{S}(\text{ID})} (\mathbf{R}_{1,i}\mathbf{G}^{-1}(\mathbf{B}_{2,j}) + y_{1,i}\mathbf{R}_{2,j})}_{:=\mathbf{R}_{\text{ID}}, \text{ which is “small”}} \right) + \left(\underbrace{y_0 + \sum_{(i,j) \in \mathcal{S}(\text{ID})} y_{1,i}y_{2,j}}_{:=\mathbf{F}_{\mathbf{y}}(\text{ID})} \right) \cdot \mathbf{G} \\ &= \mathbf{A}\mathbf{R}_{\text{ID}} + \mathbf{F}_{\mathbf{y}}(\text{ID})\mathbf{G}. \end{aligned}$$

The reduction algorithm has a trapdoor for the matrix $(\mathbf{A} \parallel \mathbf{H}(\text{ID}))$ if $\mathbf{F}_{\mathbf{y}}(\text{ID}) \neq 0 \pmod q$ and thus can simulate a private key for such an identity ID . (\mathbf{R}_{ID} corresponds to the \mathbf{G} -trapdoor [35] of $(\mathbf{A} \parallel \mathbf{H}(\text{ID}))$.) On the other hand, the reduction algorithm expects the challenge identity ID^* to satisfy $\mathbf{F}_{\mathbf{y}}(\text{ID}^*) = 0$, for which it does not know the trapdoor. If these conditions are not satisfied, the reduction fails. We have to estimate the probability that it does not abort. In particular, we have to show that

$$\Pr[\mathbf{F}_{\mathbf{y}}(\text{ID}^*) = 0 \wedge \mathbf{F}_{\mathbf{y}}(\text{ID}_1) \neq 0 \dots \wedge \mathbf{F}_{\mathbf{y}}(\text{ID}_Q) \neq 0] \quad (1)$$

is noticeable. Here, $\text{ID}_1, \dots, \text{ID}_Q$ are identities for which key extraction queries are made. By a similar analysis to [47, 6], to show a lower bound for the probability of (1), it suffices to show an upper bound for the following probability

$$\Pr[\mathbf{F}_{\mathbf{y}}(\text{ID}^*) = 0 \wedge \mathbf{F}_{\mathbf{y}}(\text{ID}_i) = 0] \quad (2)$$

for identities ID^* and ID_i where $\text{ID}^* \neq \text{ID}_i$. To show an upper bound for (2), we first observe that

$$\begin{aligned} &\mathbf{F}_{\mathbf{y}}(\text{ID}^*) = 0 \wedge \mathbf{F}_{\mathbf{y}}(\text{ID}_i) = 0 \\ \Leftrightarrow &\mathbf{F}_{\mathbf{y}}(\text{ID}^*) = 0 \wedge \mathbf{F}_{\mathbf{y}}(\text{ID}_i) - \mathbf{F}_{\mathbf{y}}(\text{ID}^*) = 0 \\ \Leftrightarrow &\underbrace{\left(y_0 + \sum_{(j,k) \in \mathcal{S}(\text{ID}^*)} y_{1,j}y_{2,k} = 0 \right)}_{\text{Event (A)}} \wedge \underbrace{\left(\sum_{(j,k) \in \mathcal{S}(\text{ID}_i)} y_{1,j}y_{2,k} - \sum_{(j,k) \in \mathcal{S}(\text{ID}^*)} y_{1,j}y_{2,k} = 0 \right)}_{\text{Event (B)}}. \end{aligned}$$

The value of y_0 is clearly independent of the Event (B). Therefore, we can easily estimate the probability of Event (A) occurring, conditioned on that Event (B) occurs. Thus, it suffices to

show an upper bound on the probability of Event (B) occurring. This can be accomplished by using the Schwartz-Zippel lemma.

Proof Continued. Based on the idea we have explained above, we can simulate key extraction queries with sufficiently high success probability. However, two problems remain in order to complete the security proof.

(C) In the above discussion, we assumed that q is much larger than Q . Therefore, if q is bounded by some polynomial, so is Q . In such a setting, we can only prove “bounded” security, where the number of key extraction queries is bounded by a predetermined polynomial.

(D) Furthermore, we are not able to generate a properly distributed challenge ciphertext, as we explain below.

Let us explain the problem (D). Assume that for the challenge identity ID^* , we have $F_{\mathbf{y}}(ID^*) = 0$ and thus $H(ID^*) = \mathbf{A}\mathbf{R}_{ID^*}$. To prove security, we have to embed the LWE problem instance \mathbf{A} and \mathbf{v} into the challenge ciphertext, where $\mathbf{v}^\top = \mathbf{s}^\top \mathbf{A} + \mathbf{x}^\top$ or \mathbf{v} a random vector. A natural way to do this is to implicitly set

$$\mathbf{x}_1 = \mathbf{x}, \quad \mathbf{x}_2 = \mathbf{R}_{ID^*}^\top \mathbf{x}$$

and compute the challenge ciphertext as

$$\mathbf{s}^\top (\mathbf{A}|H(ID)) + (\mathbf{x}_1|\mathbf{x}_2) = (\mathbf{v}^\top | \mathbf{v}^\top \mathbf{R}_{ID^*}).$$

The problem with this approach is that the vector \mathbf{x}_2 is highly correlated to the value of \mathbf{R}_{ID^*} , which includes the information of $\mathbf{y} = (y_0, \{y_{i,j}\}_{(i,j) \in [1,2] \times [1,\ell]})$ and additionally $\mathbf{R}_0, \mathbf{R}_{1,1} \dots, \mathbf{R}_{1,\ell}, \mathbf{R}_{2,1} \dots, \mathbf{R}_{2,\ell}$. While a similar (but simpler) problem is resolved in a previous work [1] using a generalized form of the leftover hash lemma [20], we are not able to do the same argument due to the additional correlation to \mathbf{y} .

We can resolve the problem by a standard technique. Namely, we “smudge out” or “eat” the problematic term $\mathbf{R}_{ID^*}^\top \mathbf{x}$ by adding a large enough term $\mathbf{x}' \in \mathbb{Z}_q^m$ to it. This makes the error terms essentially statistically independent from \mathbf{R}_{ID^*} . The size of the term \mathbf{x}' should be super-polynomially larger than the size of $\mathbf{R}_{ID^*}^\top \mathbf{x}$, but it should be polynomially smaller than q . Therefore, the size of q should be super-polynomially large, which also resolves the problem (C) at the same time. Appropriately setting the parameters, we obtain our new adaptively secure and anonymous IBE scheme.

2.3 An Additional Idea

However, making q super-polynomially large is not quite desirable because of the following two reasons. Firstly, this would negatively impact the performance of the system. Secondly, since the error term (in our case \mathbf{x}) is super-polynomially smaller compared to q , the corresponding LWE problem becomes easier. While we are not able to resolve the first problem, we present an idea to avoid the second problem.

Our first observation is that for any constant $c \in \mathbb{N}$, by making q and \mathbf{x}' sufficiently large (but polynomial size), we can show that any PPT adversary whose number of key extraction queries is bounded by n^c cannot break the security of IBE with advantage non-negligibly larger than $1/n^c$. Of course, this is not sufficient because we need the adversary to have only negligible (rather than inverse of polynomial) advantage, even if the number of key extraction queries is unbounded.

In order to accomplish this, we prepare several instances of IBE scheme with different size of q . We call each instance of the IBE scheme as a sub-scheme. The number of sub-schemes is *super-constant* (rather than super-polynomial) and therefore the resulting scheme is still efficient. The size of q varies from very small polynomial to super-polynomial. Furthermore, we “glue” them so that an adversary must break the security of all of the sub-schemes, in order to break the resulting IBE scheme. This can easily be accomplished by splitting the message by k -out-of- k secret sharing scheme, and then encrypt them by each of the sub-schemes.

In the security proof, we assume an PPT adversary \mathcal{A} that breaks the resulting IBE scheme. Since \mathcal{A} is polynomial time and has non-negligible advantage, there exists some constant $c \in \mathbb{N}$ such that the number of the key extraction queries that \mathcal{A} makes is smaller than n^c and \mathcal{A} 's advantage is non-negligibly larger than $1/n^c$. Thus, there exists at least one sub-scheme whose size of q fits for \mathcal{A} , and q is polynomial size. We transform the adversary \mathcal{A} into another adversary \mathcal{B} that breaks the sub-scheme. Since q is polynomial size, we can reduce the security to the LWE assumption with *polynomial* approximation factor. Note that similar technique is used in [21] to improve the efficiency and the reduction cost of the Naor-Reingold PRF. There, the reduction algorithm chooses the target sub-scheme based on the number of queries that the adversary makes. In our reduction, we choose the target depending on the advantage of the adversary in addition to the number of key extraction queries.

To present our results in a unified and modular manner, we introduce the notion of PIBE. Roughly speaking, PIBE is an IBE scheme that is parametrized by a variable c . Our technique to avoid super-polynomial factor we discussed above can be generalized to be a generic conversion from PIBE to IBE. Furthermore, our scheme we discussed in the previous subsection also can be captured as a special case of PIBE, in that c is set to be a super-constant.

3 Preliminaries

Notation. We denote by $[n]$ a set $\{1, 2, \dots, n\}$ for any integer $n \in \mathbb{N}$. We treat a vector as a column vector. If \mathbf{A}_1 is $n \times m$ and \mathbf{A}_2 is $n \times m'$ matrix, then $(\mathbf{A}_1 | \mathbf{A}_2)$ denotes the $n \times (m+m')$ matrix formed by concatenating \mathbf{A}_1 and \mathbf{A}_2 . We use similar notation for vectors. A function $f : \mathbb{N} \rightarrow \mathbb{R}_{\geq 0}$ is said to be negligible, if for all c , there exists N such that $f(n) < 1/n^c$ for all $n > N$. We denote by $\text{negl}(n)$ a negligible function. We denote by $x \stackrel{\$}{\leftarrow} X$ the process of sampling a value x according to the distribution X . Similarly, for a finite set S , we denote by $x \stackrel{\$}{\leftarrow} S$ the process of sampling a value x according to the uniform distribution over S . Statistical distance between two random variables X and Y with support Ω is defined as $\Delta(X; Y) = \frac{1}{2} \sum_{s \in \Omega} |\Pr[X = s] - \Pr[Y = s]|$. For ensembles of random variable $\{X(n)\}_{n \in \mathbb{N}}$ and $\{Y(n)\}_{n \in \mathbb{N}}$, we say that they are $\text{negl}(n)$ -close if $\Delta(X(n); Y(n)) = \text{negl}(n)$.

3.1 Identity-Based Encryption

Syntax. Let \mathcal{ID} be the ID space of the scheme. If a collision resistant hash function $CRH : \{0, 1\}^* \rightarrow \mathcal{ID}$ is available, one can use an arbitrary string as an identity. An IBE scheme is defined by the following four algorithms.

Setup(1^n) \rightarrow (mpk, msk): The setup algorithm takes as input a security parameter 1^n and outputs a master public key mpk and a master secret key msk.

KeyGen(mpk, msk, ID) \rightarrow sk_{ID} : The key generation algorithm takes as input the master public key mpk, the master secret key msk, and an identity $\text{ID} \in \mathcal{ID}$. It outputs a private key sk_{ID} . We assume that ID is implicitly included in sk_{ID} .

$\text{Encrypt}(\text{mpk}, \text{ID}, \text{M}) \rightarrow C$: The encryption algorithm takes as input a master public key mpk , an identity $\text{ID} \in \mathcal{ID}$, and a message M , It outputs a ciphertext C .

$\text{Decrypt}(\text{mpk}, \text{sk}_{\text{ID}}, C) \rightarrow \text{M}$ or \perp : The decryption algorithm takes as input the master public key mpk , a private key sk_{ID} , and a ciphertext C . It outputs the message M or \perp , which means that the ciphertext is not in a valid form.

Correctness. We require correctness of decryption: that is, for all n , all $\text{ID} \in \mathcal{ID}$, and all M in the specified message space,

$$\Pr[\text{Decrypt}(\text{mpk}, \text{sk}_{\text{ID}}, \text{Encrypt}(\text{mpk}, \text{ID}, \text{M})) = \text{M}] = 1 - \text{negl}(n)$$

holds, where the probability is taken over the randomness used in $(\text{mpk}, \text{msk}) \xleftarrow{\$} \text{Setup}(1^n)$, $\text{sk}_{\text{ID}} \xleftarrow{\$} \text{KeyGen}(\text{mpk}, \text{msk}, \text{ID})$, and $\text{Encrypt}(\text{mpk}, \text{ID}, \text{M})$.

Security. We now define the security for an IBE scheme Π . This security notion is defined by the following game between a challenger and an adversary \mathcal{A} .

- **Setup.** At the outset of the game, the challenger runs $\text{Setup}(1^n) \rightarrow (\text{mpk}, \text{msk})$ and gives mpk to \mathcal{A} .

- **Phase 1.** \mathcal{A} may adaptively make key-extraction queries. If \mathcal{A} submits $\text{ID} \in \mathcal{ID}$ to the challenger, the challenger returns $\text{sk}_{\text{ID}} \leftarrow \text{KeyGen}(\text{mpk}, \text{msk}, \text{ID})$.

- **Challenge Phase.** At some point, \mathcal{A} outputs a message M and an identity $\text{ID}^* \in \mathcal{ID}$, on which it wishes to be challenged. Then, the challenger picks a random coin $\text{coin} \xleftarrow{\$} \{0, 1\}$ and a random ciphertext $C \xleftarrow{\$} \mathcal{C}$ from the ciphertext space. If $\text{coin} = 0$, it runs $\text{Encrypt}(\text{mpk}, \text{ID}^*, \text{M}) \rightarrow C^*$ and gives the challenge ciphertext C^* to \mathcal{A} . If $\text{coin} = 1$, it sets the challenge ciphertext as $C^* = C$ and gives it to \mathcal{A} .

- **Phase 2.** After the challenge query, \mathcal{A} may continue to make key-extraction queries, with the added restriction that $\text{ID} \neq \text{ID}^*$.

- **Guess.** Finally, \mathcal{A} outputs guess a $\widehat{\text{coin}}$ for coin . The advantage of \mathcal{A} is defined as

$$\text{Adv}_{\mathcal{A}, \Pi}^{\text{IBE}} = \left| \Pr[\widehat{\text{coin}} = \text{coin}] - \frac{1}{2} \right|.$$

We say that Π is adaptively anonymous, if the advantage of any PPT \mathcal{A} is negligible.

We also define adaptive security (without anonymity) for Π via a similar game to the above. To define adaptive security, we change the challenge phase as follows.

- **Challenge Phase.** \mathcal{A} outputs two messages M_0, M_1 and an identity $\text{ID}^* \in \mathcal{ID}$, on which it wishes to be challenged. Then, the challenger picks a random coin $\text{coin} \xleftarrow{\$} \{0, 1\}$, runs $\text{Encrypt}(\text{mpk}, \text{ID}^*, \text{M}_{\text{coin}}) \rightarrow C^*$, and gives the challenge ciphertext C^* to \mathcal{A} .

We also say that Π is adaptively secure, if the advantage of any PPT \mathcal{A} is negligible. We note that the adaptive anonymity implies the adaptive security. Namely, the former is a stronger security notion.

3.2 Lattice Preliminaries

For positive integers q, m, n , a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, and a vector $\mathbf{u} \in \mathbb{Z}_q^m$, the m -dimensional integer lattices $\Lambda_q^\perp(\mathbf{A})$ and $\Lambda_q^{\mathbf{u}}(\mathbf{A})$ are defined as

$$\Lambda_q^\perp(\mathbf{A}) = \{\mathbf{e} \in \mathbb{Z}^m : \mathbf{A}\mathbf{e} = \mathbf{0} \pmod{q}\}$$

$$\Lambda_q^{\mathbf{u}}(\mathbf{A}) = \{\mathbf{e} \in \mathbb{Z}^m : \mathbf{A}\mathbf{e} = \mathbf{u} \pmod{q}\}.$$

Let $D_{\Lambda, \mathbf{c}, \sigma}$ denote the discrete Gaussian distribution over Λ with center \mathbf{c} and parameter γ . When \mathbf{c} is omitted, we set $\mathbf{c} = \mathbf{0}$.

Matrix Norms. For a vector \mathbf{u} , we let $\|\mathbf{u}\|$ and $\|\mathbf{u}\|_{\infty}$ denote its ℓ_2 and ℓ_{∞} norm respectively. For a matrix $\mathbf{R} \leq \mathbb{Z}^{k \times m}$ we denote three matrix norms:

$\|\mathbf{R}\|$ denotes the ℓ_2 length of the longest column of \mathbf{R} .

$\|\mathbf{R}\|_{\text{GS}}$ denotes $\|\tilde{\mathbf{R}}\|$ where $\tilde{\mathbf{R}}$ is the result of applying Gram-Schmidt to the columns of \mathbf{R} .

$\|\mathbf{R}\|_2$ is the operator norm of \mathbf{R} defined as $\|\mathbf{R}\|_2 = \sup_{\|\mathbf{x}\|=1} \|\mathbf{R}\mathbf{x}\|$.

We have that the following lemma holds [1].

Lemma 1. *Let m, n, q be positive integers with $m > n$, $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ be a matrix, $\mathbf{u} \in \mathbb{Z}_q^n$ be a vector, $\mathbf{T}_{\mathbf{A}}$ be a basis for $\Lambda_q^{\perp}(\mathbf{A})$, and $\sigma > \|\mathbf{T}_{\mathbf{A}}\|_{\text{GS}} \cdot \omega(\sqrt{\log m})$. Then we have $\Pr[\mathbf{x} \stackrel{\$}{\leftarrow} D_{\Lambda_q^{\mathbf{u}}(\mathbf{A}), \sigma} : \|\mathbf{x}\| > \sqrt{m}\sigma] < \text{negl}(n)$.*

Trapdoor Generators and Related Operations.

Lemma 2. *Let $n, m, q > 0$ be integers with q prime. There are polynomial time algorithms such that*

1. ([3, 5]): $\text{TrapGen}(1^n, 1^m, q) \rightarrow (\mathbf{A}, \mathbf{T}_{\mathbf{A}})$
a randomized algorithm that, when $m \geq 6n \lceil \log q \rceil$, outputs a full rank matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and a basis $\mathbf{T}_{\mathbf{A}} \in \mathbb{Z}^{m \times m}$ for $\Lambda_q^{\perp}(\mathbf{A})$ such that \mathbf{A} is $\text{negl}(n)$ -close to uniform and $\|\mathbf{T}_{\mathbf{A}}\|_{\text{GS}} = O(\sqrt{n \log q})$ with all but negligible probability in n .
2. ([16]): $\text{SampleLeft}(\mathbf{A}, \mathbf{F}, \mathbf{u}, \mathbf{T}_{\mathbf{A}}, \sigma) \rightarrow \mathbf{e}$
a randomized algorithm that, given a full rank matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, a matrix $\mathbf{F} \in \mathbb{Z}_q^{n \times m}$, a vector $\mathbf{u} \in \mathbb{Z}_q^n$, a basis $\mathbf{T}_{\mathbf{A}}$ for $\Lambda_q^{\perp}(\mathbf{A})$, and a Gaussian parameter $\sigma > \|\mathbf{T}_{\mathbf{A}}\|_{\text{GS}} \cdot \omega(\sqrt{\log m})$, outputs a vector $\mathbf{e} \in \mathbb{Z}^{2m}$ sampled from a distribution which is $\text{negl}(n)$ -close to $D_{\Lambda_q^{\mathbf{u}}(\mathbf{A}|\mathbf{F}), \sigma}$.
3. ([1]): $\text{SampleRight}(\mathbf{A}, \mathbf{G}, \mathbf{R}, y, \mathbf{u}, \mathbf{T}_{\mathbf{G}}, \sigma) \rightarrow \mathbf{e}$ where $\mathbf{F} = \mathbf{A}\mathbf{R} + y\mathbf{G}$
a randomized algorithm that, given a full rank matrix $\mathbf{A}, \mathbf{G} \in \mathbb{Z}_q^{n \times m}$, $y \in \mathbb{Z}_q \setminus \{0\}$, a matrix $\mathbf{R} \in \mathbb{Z}^{m \times m}$, a vector $\mathbf{u} \in \mathbb{Z}_q^n$, a basis $\mathbf{T}_{\mathbf{G}}$ for $\Lambda_q^{\perp}(\mathbf{G})$, and a Gaussian parameter $\sigma > \|\mathbf{T}_{\mathbf{G}}\|_{\text{GS}} \cdot \|\mathbf{R}\|_2 \cdot \omega(\sqrt{\log m})$ outputs a vector $\mathbf{e} \in \mathbb{Z}^{2m}$ sampled from a distribution which is $\text{negl}(n)$ -close to $D_{\Lambda_q^{\mathbf{u}}(\mathbf{A}|\mathbf{F}), \sigma}$.
4. ([35]): *Let $m > n \lceil \log q \rceil$. Then there is a fixed full-rank matrix $\mathbf{G} \in \mathbb{Z}_q^{n \times m}$ such that the lattice $\Lambda_q^{\perp}(\mathbf{G})$ has a publicly known basis $\mathbf{T}_{\mathbf{G}} \in \mathbb{Z}^{m \times m}$ with $\|\mathbf{T}_{\mathbf{G}}\|_{\text{GS}} \leq \sqrt{5}$. Furthermore, there exists a deterministic polynomial-time algorithm \mathbf{G}^{-1} which takes the input $\mathbf{U} \in \mathbb{Z}_q^{n \times m}$ and outputs $\mathbf{R} = \mathbf{G}^{-1}(\mathbf{U})$ such that $\mathbf{R} \in \{0, 1\}^{m \times m}$ and $\mathbf{G}\mathbf{R} = \mathbf{U}$.*

Note that in the above, we are abusing notation and \mathbf{G}^{-1} is not a matrix but rather a function. Namely, for any \mathbf{U} there are many choices of \mathbf{R} such that $\mathbf{G}\mathbf{R} = \mathbf{U}$, and $\mathbf{G}^{-1}(\mathbf{U})$ deterministically outputs a particular short matrix from this set. Since we have $\|\mathbf{R}\|_2 \leq m$ for any $\mathbf{R} \in \{-1, 0, 1\}^{m \times m}$, $\|\mathbf{G}^{-1}(\mathbf{U})\|_2 \leq m$ holds for any $\mathbf{U} \in \mathbb{Z}_q^{n \times m}$.

Learning with Errors. The learning with errors (LWE) problem was introduced by Regev who showed that solving it on the average is as hard as (quantumly) solving several standard lattice problems in the worst case.

Definition 1 (LWE). For an integers n , $m = m(n)$, a prime integer $q = q(n) > 2$, an error distribution $\chi = \chi(n)$ over \mathbb{Z}_q , and an PPT algorithm \mathcal{A} , an advantage for the learning with errors problem $\text{dLWE}_{n,m,q,\chi}$ of \mathcal{A} is defined as follows:

$$\text{Adv}_{\mathcal{A}}^{\text{dLWE}_{n,m,q,\chi}} = |\Pr[\mathcal{A}(\mathbf{A}, \mathbf{s}^\top \mathbf{A} + \mathbf{x}^\top) \rightarrow 1] - \Pr[\mathcal{A}(\mathbf{A}, \mathbf{v}^\top) \rightarrow 1]|$$

where $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n \times m}$, $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n$, $\mathbf{x} \xleftarrow{\$} \chi^m$, $\mathbf{v} \xleftarrow{\$} \mathbb{Z}_q^m$. We say that $\text{dLWE}_{n,m,q,\chi}$ assumption holds if $\text{Adv}_{\mathcal{A}}^{\text{dLWE}_{n,m,q,\chi}}$ is negligible for all PPT \mathcal{A} .

Let $B = B(n) \in \mathbb{N}$. A family of distributions $\chi = \{\chi_n\}$ is called B -bounded if $\Pr[\chi \in [-B, B]] = 1$. For any constant $d > 0$ and sufficiently large q , Regev [40] through a quantum reduction showed that taking χ as a q/n^d -bounded (truncated) discretized Gaussian distribution, the $\text{dLWE}_{n,m,q,\chi}$ problem is as hard as approximating the worst-case GapSVP to $n^{O(d)}$ factors, which is believed to be hard. In subsequent works, (partial) dequantization of the Regev's reduction were achieved [37, 13]. More generally, let $\chi_{\max} < q$ be the bound on the noise distribution. The difficulty of the problem is measured by the ratio q/χ_{\max} . This ratio is always bigger than 1 and the smaller it is the harder the problem. The problem appears to remain hard even when $q/\chi_{\max} < 2^{n^\epsilon}$ for some fixed ϵ that is $0 < \epsilon < 1/2$.

3.3 Basic Facts

Injective map. Let d and κ be some integers. Furthermore, let ℓ be $\ell = \lceil \kappa^{1/d} \rceil$. Then, an element of $[1, \kappa]$ can be written as an element of $[1, \ell]^d$ using some canonical map. Furthermore, it is also possible to write a subset of $[1, \kappa]$ as a subset of $[1, \ell]^d$, by naturally extending the canonical map. By identifying a bit string in $\{0, 1\}^\kappa$ with a subset of $[1, \kappa]$ (for example, by regarding the former as the indicator vector of a subset of $[1, \kappa]$), we can define an efficiently computable injective map S that maps a bit string $\text{ID} \in \{0, 1\}^\kappa$ to a subset $S(\text{ID})$ of $[1, \ell]^d$.

The following lemma can be shown by a simple calculation.

Lemma 3. (Smudging out Lemma.) Let $\mathbf{x}_0 \in \mathbb{Z}^m$ be a (fixed) vector such that $\|\mathbf{x}_0\|_\infty \leq \delta$ and let $\mathbf{x} \in \mathbb{Z}^m$ be a random vector that is chosen as $\mathbf{x} \xleftarrow{\$} [-B', B']^m$. Then, the following distributions are within statistical distance $m\delta/B'$:

$$\mathbf{x}_0 + \mathbf{x} \approx \mathbf{x}.$$

As observed in [40, 1], the following lemma is obtained as a corollary to the (general) leftover hash lemma.

Lemma 4. (Leftover Hash Lemma.) Let $q \in \mathbb{N}$ be an odd prime and let $m > (n + 1) \log q + \omega(\log n)$. Let $\mathbf{R} \xleftarrow{\$} \{-1, 1\}^{m \times m}$ and $\mathbf{A}, \mathbf{A}' \xleftarrow{\$} \mathbb{Z}_q^{n \times m}$ be uniformly random matrices. Then the distribution of $(\mathbf{A}, \mathbf{A}\mathbf{R})$ is $\text{negl}(n)$ -close to the distribution of $(\mathbf{A}, \mathbf{A}')$.

The following lemma is implicitly shown in [6]. The proof can be found in Appendix A.1.

Lemma 5. Let $a_1, \dots, a_n \in \mathbb{R}$ be real numbers such that $|\sum_{i=1}^n a_i| = \epsilon$ and $\sum_{i=1}^n |a_i| \leq 1/2$. Furthermore, let $\gamma_1, \dots, \gamma_n \in \mathbb{R}$ be real numbers such that $0 < \gamma_{\min} \leq \gamma_i \leq \gamma_{\max}$ for $i \in [n]$. Then, we have $|\sum_{i=1}^n \gamma_i a_i| \geq \gamma_{\min} \epsilon - (\gamma_{\max} - \gamma_{\min})/2$.

4 Parametrized IBE

In this section, we introduce the notion of parametrized IBE (PIBE), which is a slight extension of the ordinary notion of IBE. The syntax and the security notion for PIBE is almost the same, except that it is parametrized by an integer c . Roughly speaking, the larger c becomes, the more secure PIBE becomes. In particular, when c is super-constant in n , the security notion for PIBE corresponds to that for ordinary IBE. However, in our construction of PIBE in Section 5, in order to prove the security of the scheme for super-constant c , we need to assume super-polynomial LWE, which is a stronger assumption than the assumption that is needed for constant c . In this section, to base the scheme on a weaker assumption, we provide generic construction of adaptively secure IBE scheme from PIBE scheme that is secure *only for constant c* .

4.1 Definition of Parametrized IBE

Here, we define PIBE. The syntax of PIBE is the same as ordinary IBE except that the `Setup` algorithm is parametrized by an integer $c = c(n)$. Namely, `Setup` takes as inputs 1^n and 1^c and outputs a master public key `mpk` and a master secret key `msk`. Other algorithms, `KeyGen`, `Encrypt`, and `Decrypt` are defined as in ordinary IBE. We require that these algorithms work within a time that is polynomial in n and c .

As for the security, we define advantage $\text{Adv}_{\mathcal{A}, \Pi}^{\text{PIBE}}$ of an adversary \mathcal{A} for a PIBE scheme Π via a game that is almost the same as that of an ordinary IBE scheme. The only difference is that `mpk` and `msk` are generated by `Setup`($1^n, 1^c$) at the beginning of the game. The rest of the game is the same. We say that the scheme is c -adaptively anonymous, if for any PPT adversary \mathcal{A} such that $Q(n) \leq n^c/2 - 1$,

$$\frac{\text{Adv}_{\mathcal{A}, \Pi}^{\text{PIBE}}}{Q + 1} < \frac{1}{n^c} + \text{negl}(n) \quad (3)$$

holds for some negligible function $\text{negl}(n)$. Here $Q = Q(n)$ is the upper bound for the number of key extraction queries made by \mathcal{A} during the game.

When $c(n)$ is a constant, the c -adaptive anonymity is a weaker security notion than the adaptive anonymity for IBE, since it allows an adversary to have non-negligible advantage. Furthermore, there is a bound on the number of key extraction queries. On the other hand, when $c(n)$ is super-constant, the security definition of c -adaptive anonymity corresponds to that of adaptive anonymity for (ordinary) IBE. More precisely, we have the following theorem.

Theorem 1. *If $\Pi = (\text{Setup}, \text{KeyGen}, \text{Encrypt}, \text{Decrypt})$ is c' -adaptively anonymous for some super-constant function $c'(n) = \omega(1)$ such that $c'(n) < \text{poly}(n)$, $\Pi' = (\text{Setup}', \text{KeyGen}, \text{Encrypt}, \text{Decrypt})$ is adaptively anonymous (as an ordinary IBE) if we set*

$$\text{Setup}'(1^n) = \text{Setup}(1^n, 1^{c'(n)}).$$

Proof. Since $c'(n) < \text{poly}(n)$, `Setup'`, `KeyGen`, `Encrypt`, and `Decrypt` run in polynomial time. In addition, since $c'(n) = \omega(1)$ and thus $n^{c'}$ is super-polynomial, there is no bound on the number of key extraction queries for the adversary in the c' -adaptive anonymity game. Furthermore, since $1/n^{c'}$ is a negligible function, by Equation (3), we have

$$\text{Adv}_{\mathcal{A}, \Pi}^{\text{PIBE}} < (Q + 1) \left(\frac{1}{n^{c'}} + \text{negl}(n) \right) = \text{negl}(n)$$

for any adversary \mathcal{A} . Thus, Π' defined as above is adaptively anonymous. \square

Comparison with Bounded Collusion IBE. Our notion of PIBE is similar to the notion of bounded collusion IBE [19] (also called k -resilient IBE [29]), in that adversaries only learn private keys of an a-priori bounded number of identities. The security requirement for the former is weaker than that for the latter, because we allow adversaries to have non-negligible advantages (in the case of c is a constant). On the other hand, we pose more severe requirement on the efficiency for the former. We require the algorithms of PIBE to work in polynomial time in c , rather than in n^c . Because of this, existing bounded collusion IBE schemes [19, 29, 49, 26, 46] do not satisfy the requirement of PIBE.

4.2 IBE from PIBE

In this section, we show a conversion from a PIBE scheme $\Pi = (\text{PIBE.Setup}, \text{PIBE.KeyGen}, \text{PIBE.Encrypt}, \text{PIBE.Decrypt})$ to an (ordinary) IBE scheme $\Pi' = (\text{IBE.Setup}, \text{IBE.KeyGen}, \text{IBE.Encrypt}, \text{IBE.Decrypt})$. In the following, let $\eta(n)$ be any function such that $\eta(n) = \omega(1)$ (e.g., $\eta(n) = \log \log(n)$). We also let the message space of Π and Π' be $\{0, 1\}^{\ell_M}$ for some $\ell_M \in \mathbb{N}$.

$\text{IBE.Setup}(1^n)$: It runs $\text{PIBE.Setup}(1^n, 1^i) \rightarrow (\text{mpk}^{(i)}, \text{msk}^{(i)})$ for $i = 1, \dots, \eta$. It outputs

$$\text{mpk} = (\text{mpk}^{(1)}, \text{mpk}^{(2)}, \dots, \text{mpk}^{(\eta)}) \quad \text{and} \quad \text{msk} = (\text{msk}^{(1)}, \text{msk}^{(2)}, \dots, \text{msk}^{(\eta)}).$$

$\text{IBE.KeyGen}(\text{mpk}, \text{msk}, \text{ID})$: It runs $\text{PIBE.KeyGen}(\text{mpk}^{(i)}, \text{msk}^{(i)}, \text{ID}) \rightarrow \text{sk}_{\text{ID}}^{(i)}$ for $i = 1, \dots, \eta$. It outputs

$$\text{sk}_{\text{ID}} = (\text{sk}_{\text{ID}}^{(1)}, \text{sk}_{\text{ID}}^{(2)}, \dots, \text{sk}_{\text{ID}}^{(\eta)}).$$

$\text{Encrypt}(\text{mpk}, \text{ID}, M)$: To encrypt $M = \{0, 1\}^{\ell_M}$, it picks random $M^{(i)} \in \{0, 1\}^{\ell_M}$ for $i \in [\eta]$ subject to constraint that $M = \bigoplus_{i=1}^{\eta} M^{(i)}$, where \bigoplus denotes bitwise exclusive or. Then it runs

$$\text{PIBE.Encrypt}(\text{mpk}^{(i)}, \text{ID}, M^{(i)}) \rightarrow C^{(i)} \quad \text{for } i = 1, \dots, \eta.$$

Finally, it outputs the ciphertext $C = (C^{(1)}, \dots, C^{(\eta)})$.

$\text{Decrypt}(\text{mpk}, \text{sk}_{\text{ID}}, C)$: It first parses the ciphertext and the private key as $C \rightarrow (C^{(1)}, \dots, C^{(\eta)})$ and $\text{sk}_{\text{ID}} \rightarrow (\text{sk}_{\text{ID}}^{(1)}, \dots, \text{sk}_{\text{ID}}^{(\eta)})$. Then, it runs

$$\text{PIBE.Decrypt}(\text{mpk}^{(i)}, \text{sk}_{\text{ID}}^{(i)}, C^{(i)}) \rightarrow M^{(i)} \quad \text{for } i = 1, \dots, \eta.$$

Finally, it outputs $M = \bigoplus_{i=1}^{\eta} M^{(i)}$.

Correctness of the scheme can be shown very easily. The following theorem addresses the security of the scheme. Note that the resulting IBE scheme is not anonymous even if the original PIBE scheme is anonymous.

Theorem 2. *Assume that PIBE Π is secure for all (constant) $c \in \mathbb{N}$. Then, Π' is adaptively secure as an (ordinary, not parametrized) IBE scheme.*

Proof. Assume an adversary \mathcal{A} that breaks Π' with non-negligible probability. Since \mathcal{A} is a PPT algorithm, there exist constants $c' \in \mathbb{N}$ and $c'' \in \mathbb{N}$ such that

- The advantage $\epsilon(n)$ of \mathcal{A} is greater than $2/n^{c'}$ for infinitely many n .

- The number $Q(n)$ of key extraction queries that \mathcal{A} makes is bounded by $n^{c''}/2 - 1$.

Let i^* be $i^* = c' + c''$. Then, we have

$$\frac{\epsilon(n)}{2(Q(n) + 1)} - \frac{1}{n^{i^*}} \geq \frac{2}{n^{c'+c''}} - \frac{1}{n^{i^*}} = \frac{1}{n^{i^*}} \quad (4)$$

for infinitely many n . In particular, $\epsilon/2(Q + 1) - 1/n^{i^*}$ cannot be bounded by any negligible function. To show the theorem, we construct an adversary \mathcal{B} against i^* -adaptive anonymity of PIBE Π from \mathcal{A} . In the following, we assume $\eta \geq i^*$. Since $\eta(n) = \omega(1)$, this holds for sufficiently large n .

Setup. First, $\text{PIBE.Setup}(1^n, 1^{i^*}) \rightarrow (\text{mpk}^{(i^*)}, \text{msk}^{(i^*)})$ is run and $\text{mpk}^{(i^*)}$ is given to \mathcal{B} . Then, \mathcal{A} runs $\text{PIBE.Setup}(1^n, 1^i) \rightarrow (\text{mpk}^{(i)}, \text{msk}^{(i)})$ for $i = [1, \eta] \setminus \{i^*\}$ and sets $\text{mpk} = (\text{mpk}^{(1)}, \text{mpk}^{(2)}, \dots, \text{mpk}^{(\eta)})$. \mathcal{B} keeps $\text{msk}^{(i)}$ for $i \in [1, \eta] \setminus \{i^*\}$ secret, and returns mpk to \mathcal{A} .

Phase 1 and 2. When \mathcal{A} makes a key extraction query for an identity ID , \mathcal{B} queries a private key for the same ID to its challenger. Then, $\text{PIBE.KeyGen}(\text{mpk}^{(i^*)}, \text{msk}^{(i^*)}, \text{ID}) \rightarrow \text{sk}_{\text{ID}}^{(i^*)}$ is run and $\text{sk}_{\text{ID}}^{(i^*)}$ is given to \mathcal{B} . Then \mathcal{B} runs $\text{PIBE.KeyGen}(\text{mpk}^{(i)}, \text{msk}^{(i^*)}, \text{ID}) \rightarrow \text{sk}_{\text{ID}}^{(i)}$ for $i \in [1, \eta] \setminus \{i^*\}$ and returns $\text{sk}_{\text{ID}} = (\text{sk}_{\text{ID}}^{(1)}, \dots, \text{sk}_{\text{ID}}^{(\eta)})$ to \mathcal{A} .

Challenge. When \mathcal{A} makes a challenge query for (ID^*, M_0, M_1) , \mathcal{B} first picks random $M^{(i)} \xleftarrow{\$} \{0, 1\}^{\ell_M}$ for $i \in [1, \eta] \setminus \{i^*\}$. Then, it sets

$$M_b^{(i^*)} = M_b \oplus \left(\bigoplus_{i \in [1, \eta] \setminus \{i^*\}} M^{(i)} \right) \quad \text{for } b \in \{0, 1\}$$

and runs $\text{PIBE.Encrypt}(\text{mpk}^{(i)}, \text{ID}, M^{(i)}) \rightarrow C^{(i)}$ for $i \in [1, \eta] \setminus \{i^*\}$. Then, it picks random coin $\text{coin}' \xleftarrow{\$} \{0, 1\}$ and makes the challenge query for $(\text{ID}^*, M_{\text{coin}'})^{(i^*)}$ to its challenger. Then, the challenger picks a coin $\text{coin} \xleftarrow{\$} \{0, 1\}$ and returns C^* to \mathcal{B} . If $\text{coin} = 0$, we have $\text{PIBE.Encrypt}(\text{mpk}^{(i^*)}, \text{ID}^*, M_{\text{coin}'}^{(i^*)}) \rightarrow C^*$. Otherwise, C^* is a random element of the ciphertext space. Given C^* , \mathcal{B} returns the challenge ciphertext

$$(C^{(1)}, \dots, C^{(i^*-1)}, C^*, C^{(i^*+1)}, \dots, C^{(\eta)})$$

to \mathcal{A} .

Guess. Finally, \mathcal{A} outputs a guess $\widehat{\text{coin}}$ for coin' . If $\widehat{\text{coin}} = \text{coin}'$, \mathcal{B} outputs 0 as its guess for coin and outputs 1 otherwise.

Analysis. We can see that \mathcal{B} is a valid adversary for the parametrized IBE Π since \mathcal{A} does not make a key extraction query for ID^* . Furthermore, \mathcal{B} makes the same number of key extraction queries as \mathcal{A} and in particular, we have $Q(n) < n^{i^*}/2 - 1$. It is easy to see that the view of the adversary \mathcal{A} corresponds to that in adaptive security game for IBE Π' when $\text{coin} = 0$. It can also be seen that the view of the adversary is independent of coin' when $\text{coin} = 1$. Therefore, we have

$$\begin{aligned} \text{Adv}_{\mathcal{B}, \Pi}^{\text{PIBE}} &= \left| \frac{1}{2} \Pr[\widehat{\text{coin}} = \text{coin}' | \text{coin} = 0] + \frac{1}{2} \Pr[\widehat{\text{coin}} \neq \text{coin}' | \text{coin} = 1] - \frac{1}{2} \right| \\ &= \frac{1}{2} \left| \Pr[\widehat{\text{coin}} = \text{coin}' | \text{coin} = 0] - \frac{1}{2} \right| = \frac{1}{2} \epsilon(n). \end{aligned}$$

Thus, by Equation (4), \mathcal{B} is a successful attacker against the i^* -adaptive anonymity of Π . \square

More Efficient Conversion. In the above conversion, we run η instances of PIPE scheme in parallel. The number of instances can be reduced to $O(\log \eta)$. We briefly sketch the construction and the security proof for it. Let us assume that η is a power of 2. In the setup algorithm of the variant, we run $\text{PIPE.Setup}(1^n, 1^i) \rightarrow (\text{mpk}^{(i)}, \text{msk}^{(i)})$ for $i = 1, 2, 4, \dots, 2^i, \dots, 2^{\log \eta} (= \eta)$, instead of $i = 1, 2, \dots, \eta$. Other algorithms are defined similarly to the above. In the security proof, the target of the reduction algorithm is set to be i^* such that $2^{i^*-1} \leq c' + c'' < 2^{i^*}$.

5 Our Construction of PIPE from Lattices

Here, we show our constructions of PIPE from lattices. By setting the parameter c super-constant or applying the conversions in Section 4.2, we obtain IBE schemes that provide trade-off between the efficiency, security, and the underlying assumptions. (See Section 6 for the overview). In this section, we first introduce some functions that will be needed to describe our construction. Then, we show our construction of PIPE scheme for single-bit message space. We then prove the security of the scheme. Finally, we discuss extension of the scheme to the multi-bit variant.

5.1 Homomorphic Computation

Let d be a natural number. We introduce a function $\text{PubEval}_d : (\mathbb{Z}_q^{n \times m})^d \rightarrow \mathbb{Z}_q^{n \times m}$ which takes a set of matrices $\mathbf{B}_1, \mathbf{B}_2, \dots, \mathbf{B}_d \in \mathbb{Z}_q^{n \times m}$ as inputs and outputs a matrix in $\mathbb{Z}_q^{n \times m}$. The function is defined recursively as follows:

$$\text{PubEval}_d(\mathbf{B}_1, \dots, \mathbf{B}_d) = \begin{cases} \mathbf{B}_1 & \text{if } d = 1 \\ \mathbf{B}_1 \cdot \mathbf{G}^{-1}(\text{PubEval}_{d-1}(\mathbf{B}_2, \dots, \mathbf{B}_d)) & \text{if } d \geq 2. \end{cases}$$

We have the following lemma. The proof of the lemma can be found in Appendix A.2.

Lemma 6. *Let $\mathbf{A}, \mathbf{B}_1, \dots, \mathbf{B}_d$ be matrices in $\mathbb{Z}_q^{n \times m}$ and $\mathbf{R}_1, \dots, \mathbf{R}_d$ be matrices in $\mathbb{Z}^{m \times m}$ such that $\mathbf{B}_i = \mathbf{A}\mathbf{R}_i + y_i\mathbf{G}$ for $i \in [d]$. Furthermore, we assume that $\|\mathbf{R}_i\|_2 \leq m$, $|y_i| \leq \delta$ for $i \in [d]$, and $\delta > m$. Then, there exists an efficient algorithm TrapEval_d that takes $\mathbf{R}_1, \dots, \mathbf{R}_d, y_1, \dots, y_d$ as inputs and outputs \mathbf{R}' such that*

$$\text{PubEval}_d(\mathbf{B}_1, \dots, \mathbf{B}_d) = \mathbf{A}\mathbf{R}' + y_1 \cdots y_d \cdot \mathbf{G} \quad (5)$$

and $\|\mathbf{R}'\|_2 \leq md\delta^{d-1}$.

5.2 Our Construction

In the following, we present our PIPE scheme. Let d be a (flexible) constant. In addition, let the identity space of the scheme be $\mathcal{ID} = \{0, 1\}^\kappa$ for some $\kappa \in \mathbb{N}$ and the message space be $\{0, 1\}$. For our construction, we consider an efficiently computable injective map S that maps an identity $\text{ID} \in \{0, 1\}^\kappa$ to a subset $S(\text{ID})$ of $[1, \ell]^d$, where $\ell = \lceil \kappa^{1/d} \rceil$. Such a map can be constructed easily as we explained in Section 3.3. We would typically set $\kappa = O(n)$, and thus $\ell = O(n^{1/d})$ in such a case.

Setup($1^n, 1^c$): On input 1^n and 1^c , it sets the parameters q, m, σ, B, B' , and a distribution χ as specified in Section 5.3, where q is a prime number. Then, it picks random matrices $\mathbf{B}_0 \xleftarrow{\$} \mathbb{Z}_q^{n \times m}$, $\mathbf{B}_{i,j} \xleftarrow{\$} \mathbb{Z}_q^{n \times m}$ for $(i, j) \in [d, \ell]$ and a vector $\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^n$. It also picks

$\text{TrapGen}(1^n, 1^m, q) \rightarrow (\mathbf{A}, \mathbf{T}_\mathbf{A}) \in \mathbb{Z}_q^{n \times m} \times \mathbb{Z}^{m \times m}$ such that $\|\mathbf{T}_\mathbf{A}\|_{\text{GS}} = O(\sqrt{n \log q})$. It finally outputs

$$\text{mpk} = (\mathbf{A}, \mathbf{B}_0, \{\mathbf{B}_{i,j}\}_{(i,j) \in [d,\ell]}, \mathbf{u}) \quad \text{and} \quad \text{msk} = \mathbf{T}_\mathbf{A}.$$

In the following, we use a deterministic function $\text{H} : \mathcal{ID} \rightarrow \mathbb{Z}_q^{n \times m}$ that is defined as follows.

$$\text{H}(\text{ID}) = \mathbf{B}_0 + \sum_{(j_1, \dots, j_d) \in \mathcal{S}(\text{ID})} \text{PubEval}_d(\mathbf{B}_{1,j_1}, \mathbf{B}_{2,j_2}, \dots, \mathbf{B}_{d,j_d}) \in \mathbb{Z}_q^{n \times m}.$$

$\text{KeyGen}(\text{mpk}, \text{msk}, \text{ID})$: It first computes $\text{H}(\text{ID})$ and picks $\mathbf{e} \in \mathbb{Z}^{2m}$ such that

$$(\mathbf{A} | \text{H}(\text{ID})) \cdot \mathbf{e} = \mathbf{u}$$

by running $\text{SampleLeft}(\mathbf{A}, \text{H}(\text{ID}), \mathbf{u}, \mathbf{T}_\mathbf{A}, \sigma) \rightarrow \mathbf{e}$. It returns $\text{sk}_{\text{ID}} = \mathbf{e}$.

$\text{Encrypt}(\text{mpk}, \text{ID}, b)$: To encrypt a message $b \in \{0, 1\}$, it picks $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n$, $x_0 \xleftarrow{\$} \chi$, $\mathbf{x}_1 \xleftarrow{\$} \chi^m$, $\mathbf{x}_2 \xleftarrow{\$} [-B', B']^m$ and computes

$$c_0 = \mathbf{s}^\top \mathbf{u} + x_0 + b \cdot \lceil q/2 \rceil, \quad \mathbf{c}_1^\top = \mathbf{s}^\top (\mathbf{A} | \text{H}(\text{ID})) + (\mathbf{x}_1^\top | \mathbf{x}_2^\top).$$

Finally, it returns the ciphertext $C = (c_0, \mathbf{c}_1)$.

$\text{Decrypt}(\text{mpk}, \text{sk}_{\text{ID}}, C)$: To decrypt a ciphertext $C = (c_0, \mathbf{c}_1)$ using a private key $\text{sk}_{\text{ID}} := \mathbf{e}$, it first computes

$$w = c_0 - \mathbf{c}_1^\top \cdot \mathbf{e} \in \mathbb{Z}_q.$$

Then it returns 1 if $|w - \lceil q/2 \rceil| < \lceil q/4 \rceil$ and 0 otherwise.

5.3 Correctness and Parameter Selection

When the cryptosystem is operated as specified, we have during decryption,

$$w = c_0 - \mathbf{c}_1^\top \cdot \mathbf{e} = b \cdot \lceil q/2 \rceil + \underbrace{x_0 - (\mathbf{x}_1^\top | \mathbf{x}_2^\top) \cdot \mathbf{e}}_{\text{error term}}.$$

Lemma 7. *Assuming $B' > B$, the error term is bounded by $O(B' \sigma m)$ with overwhelming probability.*

Proof. Since χ is B -bounded distribution, with overwhelming probability, we have

$$\begin{aligned} |x_0 - (\mathbf{x}_1^\top | \mathbf{x}_2^\top) \cdot \mathbf{e}| &\leq |x_0| + |(\mathbf{x}_1^\top | \mathbf{x}_2^\top) \cdot \mathbf{e}| \\ &\leq |x_0| + \|(\mathbf{x}_1^\top | \mathbf{x}_2^\top)\| \cdot \|\mathbf{e}\| \\ &\leq B + \max\{B, B'\} \cdot \sqrt{2m} \cdot \sigma \sqrt{2m} \\ &= O(B' \sigma m). \end{aligned}$$

The second inequality above follows from Cauchy-Schwartz and the third inequality follows from Lemma 1. \square

Parameter selection. Now, to satisfy the correctness requirement and make the security proof work, we need that

- the error term is less than $q/5$ with overwhelming probability (i.e., $\Omega(B'\sigma m) < q$),
- that q is sufficiently large so that the simulation works (i.e., $q > \Theta(\kappa(dn^c)^d)$),
- that TrapGen can operate (i.e., $m \geq 6n \lceil \log q \rceil$),
- that the leftover hash lemma (Lemma 4) can be applied in the security proof (i.e., $m = (n+1) \log q + \omega(\log n)$),
- that σ is sufficiently large so that SampleLeft and SampleRight work, (i.e., $\sigma > O(\sqrt{n \log q}) \cdot \omega(\sqrt{\log m})$ and $\sigma > m(1 + \kappa d^d n^{c(d-1)}) \cdot \omega(\sqrt{\log m})$, where the latter condition turns out to be more restrictive),
- that the “noise smudging step” in the security proof works (i.e., $m^{5/2}(1 + \kappa d^d n^{c(d-1)})B/B' \leq d/(\kappa+1)(dn^c)^{d+1}$. See Equation (11).)

To satisfy the above requirements, we set the parameters as follows:

$$\begin{aligned} m &= O(n \log q), & q &= O(n^{3c(d-1)+3c'+6}), & \chi &= D_{\mathbb{Z}, \sqrt{n}}, \\ \sigma &= m\kappa n^{c(d-1)} \cdot \omega(\sqrt{\log m}), & B &= O(n), & B' &= O(m^{5/2} \kappa^2 n^{2cd+1}), \end{aligned}$$

where c' is a constant such that $\kappa = O(n^{c'})$. Typically, we would set $c' = 1$.

5.4 Security Proof

The following theorem addresses the security of the scheme. The proof is based on the partitioning technique, similarly to [47, 6, 1, 12]. For simplicity, we opt to use the framework of [6] in our analysis, which does not require the artificial abort step [47]. The analysis with the artificial abort step is also possible, and it might lead to a scheme with slightly better efficiency (up to constant factors).

Theorem 3. *The above scheme is c -adaptive anonymous assuming $\text{dLWE}_{n,m+1,q,\chi}$ is hard, where the ciphertext space is $\mathcal{C} = \mathbb{Z}_q \times \mathbb{Z}_q^{2m}$.*

Proof. Let \mathcal{A} be a PPT adversary that breaks c -adaptive anonymity of the scheme. In addition, let $\epsilon = \epsilon(n)$ and $Q = Q(n)$ be its advantage and the upper bound of the number of key extraction queries, respectively. Without loss of generality, we assume that \mathcal{A} always makes exactly Q key extraction queries. Let us define \tilde{c} as a constant that satisfies

$$Q \leq \frac{n^{\tilde{c}}}{2} - 1 \quad \text{and} \quad \frac{\epsilon}{Q+1} - \frac{1}{n^{\tilde{c}}} = \text{nonneg}(n) \quad (6)$$

where $\text{nonneg}(n)$ is some non-negligible function. We explain such \tilde{c} always exist. In the case of $c = c(n)$ is a constant, we simply let $\tilde{c} = c$. Let us consider the case of $c(n) = \omega(1)$. Since \mathcal{A} is a PPT algorithm, there exists a constant c' such that $Q(n) \leq n^{c'}/2 - 1$. Furthermore, since \mathcal{A} breaks c -adaptive anonymity of the scheme and $1/n^c$ is negligible, $\epsilon/(Q+1)$ is non-negligible. Therefore, there exists a constant c'' such that $\epsilon/(Q+1) > 2/n^{c''}$ holds for infinitely many n . By setting $\tilde{c} = \max\{c', c''\}$, we are done. We note that in any case, $\tilde{c}(n) \leq c(n)$ holds for sufficiently large n .

We show the security of the scheme via the following games. In each game, a value $\text{coin}' \in \{0, 1\}$ is defined. While it is set $\text{coin}' = \widehat{\text{coin}}$ in the first game, these values might be different in the later games. In the following, we define X_i be the event that $\text{coin}' = \text{coin}$.

Game₀ : This is the real security game. Recall that since the ciphertext space is $\mathcal{C} = \mathbb{Z}_q \times \mathbb{Z}_q^{2m}$, in the challenge phase, the challenge ciphertext is set as $C^* = (c_0, \mathbf{c}_1) \xleftarrow{\$} \mathbb{Z}_q \times \mathbb{Z}_q^{2m}$ if $\text{coin} = 1$. At the end of the game, \mathcal{A} outputs a guess $\widehat{\text{coin}}$ for coin . Finally, the challenger sets $\text{coin}' = \widehat{\text{coin}}$. By the definition, we have

$$\left| \Pr[X_0] - \frac{1}{2} \right| = \left| \Pr[\text{coin}' = \text{coin}] - \frac{1}{2} \right| = \left| \Pr[\widehat{\text{coin}} = \text{coin}] - \frac{1}{2} \right| = \epsilon.$$

Game₁ : In this game, we change **Game₀** so that the challenger performs the following additional step at the end of the game. First, the challenger picks $\mathbf{y} = (y_0, \{y_{i,j}\}_{(i,j) \in [d,\ell]})$ as

$$y_0 \xleftarrow{\$} [-(\kappa + 1)(dn^{\tilde{c}})^d + 1, 0] \quad \text{and} \quad y_{i,j} \xleftarrow{\$} [1, dn^{\tilde{c}}] \quad \text{for} \quad (i, j) \in [d] \times [\ell].$$

We define a function $F_{\mathbf{y}} : \mathcal{ID} \rightarrow \mathbb{Z}_q$ as follows:

$$F_{\mathbf{y}}(\text{ID}) = y_0 + \sum_{(j_1, \dots, j_d) \in S(\text{ID})} y_{1,j_1} \cdots y_{d,j_d}.$$

Then the challenger checks whether the following condition holds:

$$F_{\mathbf{y}}(\text{ID}^*) = 0 \wedge F_{\mathbf{y}}(\text{ID}_1) \neq 0 \wedge F_{\mathbf{y}}(\text{ID}_2) \neq 0 \wedge \cdots \wedge F_{\mathbf{y}}(\text{ID}_Q) \neq 0 \quad (7)$$

where ID^* is the challenge identity, and $\text{ID}_1, \dots, \text{ID}_Q$ are identities for which \mathcal{A} has made key extraction queries. If it does not hold, the challenger ignores the output $\widehat{\text{coin}}$ of \mathcal{A} , and sets $\text{coin}' \xleftarrow{\$} \{0, 1\}$. In this case, we say that the challenger aborts. If condition (7) holds, the challenger sets $\text{coin}' = \widehat{\text{coin}}$. As we will show in Lemma 8, we have

$$\left| \Pr[X_1] - \frac{1}{2} \right| \geq \frac{1}{\kappa + 1} \cdot \left(\frac{1}{dn^{\tilde{c}}} \right)^d \cdot \left(\epsilon - \frac{Q}{n^{\tilde{c}}} \right).$$

So as not to interrupt the proof of Theorem 3, we intentionally skip the proof for the time being.

Game₂ : In this game, we change the way \mathbf{B}_0 and $\mathbf{B}_{i,j}$ are chosen. At the beginning of the game, the challenger picks $\mathbf{R}_0, \mathbf{R}_{i,j} \xleftarrow{\$} \{-1, 1\}^{m \times m}$ for $(i, j) \in [d] \times [\ell]$. It also picks \mathbf{y} as in **Game₁**. Then, \mathbf{A} , \mathbf{B}_0 , and $\mathbf{B}_{i,j}$ are defined as

$$\mathbf{B}_0 = \mathbf{A}\mathbf{R}_0 + y_0\mathbf{G}, \quad \mathbf{B}_{i,j} = \mathbf{A}\mathbf{R}_{i,j} + y_{i,j}\mathbf{G} \quad (8)$$

for $(i, j) \in [d] \times [\ell]$. The rest of the game is the same as in **Game₁**.

Then, we bound $|\Pr[X_2] - \Pr[X_1]|$. By Lemma 4, the distributions

$$\left(\mathbf{A}, \mathbf{A}\mathbf{R}_0 + y_0\mathbf{G}, \{\mathbf{A}\mathbf{R}_{i,j} + y_{i,j}\mathbf{G}\}_{(i,j) \in [d] \times [\ell]} \right) \quad \text{and} \quad \left(\mathbf{A}, \mathbf{B}_0, \{\mathbf{B}_{i,j}\}_{(i,j) \in [d] \times [\ell]} \right)$$

are $\text{negl}(n)$ -close, where $\mathbf{B}_0, \mathbf{B}_{i,j} \xleftarrow{\$} \mathbb{Z}_q^{n \times m}$. Therefore, we have

$$|\Pr[X_1] - \Pr[X_2]| = \text{negl}(n).$$

Before describing the next game, we define \mathbf{R}_{ID} for an identity $\text{ID} \in \mathcal{ID}$ as

$$\mathbf{R}_{\text{ID}} = \mathbf{R}_0 + \sum_{(j_1, \dots, j_d) \in S(\text{ID})} \text{TrapEval}(\mathbf{R}_{1, j_1}, \dots, \mathbf{R}_{d, j_d}, y_{1, j_1}, \dots, y_{d, j_d}). \quad (9)$$

Note that by Lemma 6, we have

$$\begin{aligned} \|\mathbf{R}_{\text{ID}}^\top\|_2 &= \|\mathbf{R}_{\text{ID}}\|_2 \\ &\leq \|\mathbf{R}_0\|_2 + \sum_{(j_1, \dots, j_d) \in S(\text{ID})} \|\text{TrapEval}(\mathbf{R}_{1, j_1}, \dots, \mathbf{R}_{d, j_d}, y_{1, j_1}, \dots, y_{d, j_d})\|_2 \\ &\leq (m + \kappa(md \cdot (dn^{\tilde{c}})^{d-1})) \\ &\leq m(1 + \kappa d^d n^{c(d-1)}) \end{aligned} \quad (10)$$

for any $\text{ID} \in \mathcal{ID}$. The last inequality above follows from $\tilde{c} \leq c$.

Game₃ : In this game, we change the way the challenge ciphertext is created when $\text{coin} = 0$. If $\text{coin} = 0$, to create the challenge ciphertext **Game₃** challenger first picks $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n$, $x_0 \xleftarrow{\$} \chi$, $\mathbf{x}_1 \xleftarrow{\$} \chi^m$, $\mathbf{x}_2 \xleftarrow{\$} [-B', B']^m$ and computes \mathbf{R}_{ID^*} . Then, the challenge ciphertext $C^* = (c_0, \mathbf{c}_1)$ is computed as

$$c_0 = \mathbf{s}^\top \mathbf{u} + x_0 + b \cdot \lceil q/2 \rceil, \quad \mathbf{c}_1^\top = \mathbf{s}^\top (\mathbf{A} | \mathbf{H}(\text{ID}^*)) + (\mathbf{x}_1^\top | \mathbf{x}_1^\top \mathbf{R}_{\text{ID}^*} + \mathbf{x}_2^\top)$$

where $b \in \{0, 1\}$ is the message chosen by \mathcal{A} .

We then proceed to bound $|\Pr[X_3] - \Pr[X_2]|$. Since \mathbf{x}_1 is chosen from a B -bounded distribution, we have

$$\|\mathbf{R}_{\text{ID}^*}^\top \mathbf{x}_1\|_\infty \leq \|\mathbf{R}_{\text{ID}^*}^\top \mathbf{x}_1\|_2 \leq \|\mathbf{R}_{\text{ID}^*}^\top\|_2 \cdot \|\mathbf{x}_1\| \leq m^{3/2} (1 + \kappa d^d n^{c(d-1)}) B.$$

When all randomness other than \mathbf{x}_2 in this game is fixed, the distributions \mathbf{x}_2 and $\mathbf{R}_{\text{ID}^*}^\top \cdot \mathbf{x}_1 + \mathbf{x}_2$ are within statistical distance

$$m \cdot \|\mathbf{R}_{\text{ID}^*}^\top \mathbf{x}_1\|_\infty / B' = m^{5/2} (1 + \kappa d^d n^{c(d-1)}) B / B' \leq \frac{d}{\kappa + 1} \cdot \left(\frac{1}{dn^c} \right)^{d+1} \quad (11)$$

by Lemma 3. Averaging over all other randomness, we have that the distribution of the challenge ciphertext is within statistical distance $d/(\kappa + 1)(dn^c)^{d+1}$ from the previous game, when $\text{coin} = 0$. In the case of $\text{coin} = 1$, the view of \mathcal{A} is unchanged. Therefore, we conclude that the view of \mathcal{A} in this game is within statistical distance $d/(\kappa + 1)(dn^c)^{d+1}$ from the previous game. Thus, we have

$$|\Pr[X_2] - \Pr[X_3]| \leq \frac{d}{\kappa + 1} \cdot \left(\frac{1}{dn^c} \right)^{d+1}.$$

Game₄ Recall that in the previous game, the challenger aborts at the end of the game, if the condition (7) is not satisfied. In this game, we change the game so that the challenger aborts as soon as the abort condition becomes true. Since this is only a conceptual change, we have

$$\Pr[X_3] = \Pr[X_4].$$

Game₅ In this game, we change the way the matrix \mathbf{A} is sampled. Namely, **Game₅** challenger picks $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n \times m}$ instead of generating it with a trapdoor. By Lemma 2, this makes only negligible difference. Furthermore, we also change the way the key extraction queries are answered. When \mathcal{A} makes a key extraction query for an identity ID , the challenger first computes \mathbf{R}_{ID} as in Equation (9). By the definition of \mathbf{R}_{ID} , it holds that

$$\mathbf{H}(\text{ID}) = \mathbf{A} \cdot (\mathbf{R}_{\text{ID}} + \mathbf{F}_y(\text{ID})\mathbf{G}).$$

If $\mathbf{F}_y(\text{ID}) = 0$, it aborts, as the previous game. Otherwise, it runs

$$\text{SampleRight}(\mathbf{A}, \mathbf{G}, \mathbf{R}_{\text{ID}}, \mathbf{F}_y(\text{ID}), \mathbf{u}, \mathbf{T}_{\mathbf{G}}, \sigma) \rightarrow \mathbf{e},$$

and returns \mathbf{e} to \mathcal{A} . Note that the private key was sampled as

$$\text{SampleLeft}(\mathbf{A}, \mathbf{H}(\text{ID}), \mathbf{u}, \mathbf{T}_{\mathbf{A}}, \sigma) \rightarrow \mathbf{e}$$

in the previous game. By Equation (10) and the choice of σ , the output distribution of **SampleRight** is $\text{negl}(n)$ -close to $D_{\Lambda_q^u(\mathbf{A}|\mathbf{H}(\text{ID})), \sigma}$. Similarly, by the choice of σ , the output distribution of **SampleLeft** is also $\text{negl}(n)$ -close to $D_{\Lambda_q^u(\mathbf{A}|\mathbf{H}(\text{ID})), \sigma}$. Therefore, the above change alters the view of the adversary only negligibly. Thus, we have

$$|\Pr[X_4] - \Pr[X_5]| = \text{negl}(n).$$

Game₆ In this game, we change the way the challenge ciphertext is created when $\text{coin} = 0$. If $\text{coin} = 0$, to create the challenge ciphertext for the identity ID^* and the message b , **Game₆** challenger first picks $v_0 \xleftarrow{\$} \mathbb{Z}_q$, $\mathbf{v}_1 \xleftarrow{\$} \mathbb{Z}_q^m$, $\mathbf{x}_2 \xleftarrow{\$} [-B', B']^m$ and computes \mathbf{R}_{ID^*} . Then, it sets the challenge ciphertext $C^* = (c_0, \mathbf{c}_1)$ as

$$c_0 = v_0 + b \cdot \lceil q/2 \rceil, \quad \mathbf{c}_1^\top = (\mathbf{v}_1^\top | \mathbf{v}_1^\top \mathbf{R}_{\text{ID}^*}) + (\mathbf{0}_m^\top | \mathbf{x}_2^\top).$$

As we will show in Lemma 9, assuming $\text{dLWE}_{n, m+1, q, \chi}$ is hard, we have

$$|\Pr[X_5] - \Pr[X_6]| = \text{negl}(n).$$

Game₇ In this game, we change the challenge ciphertext to be a random vector, regardless of whether $\text{coin} = 0$ or $\text{coin} = 1$. Namely, **Game₇** challenger generates the challenge ciphertext (c_0, \mathbf{c}_1) as

$$c_0 \xleftarrow{\$} \mathbb{Z}_q, \quad \text{and} \quad \mathbf{c}_1 \xleftarrow{\$} \mathbb{Z}_q^m.$$

We now proceed to bound $|\Pr[X_7] - \Pr[X_6]|$. Since **Game₆** and **Game₇** differ only in the creation of the challenge ciphertext when $\text{coin} = 0$, we focus on this case. First, it is easy to see that c_0 is uniformly random over \mathbb{Z}_q in both of **Game₆** and **Game₇**. We also have to show that the distribution of \mathbf{c}_1 is $\text{negl}(n)$ -close to the uniform distribution over \mathbb{Z}_q^{2m} . To see this, it suffices to show that $(\mathbf{v}_1^\top | \mathbf{v}_1^\top \mathbf{R}_{\text{ID}^*})$ is distributed statistically close to uniform distribution over \mathbb{Z}_q^{2m} . Observe that the following distributions are $\text{negl}(n)$ -close:

$$(\mathbf{A}, \mathbf{A}\mathbf{R}_0, \mathbf{v}_1^\top, \mathbf{v}_1^\top \mathbf{R}_0) \approx (\mathbf{A}, \mathbf{A}', \mathbf{v}_1^\top, \mathbf{v}_1'^\top) \approx (\mathbf{A}, \mathbf{A}\mathbf{R}_0, \mathbf{v}_1^\top, \mathbf{v}_1'^\top), \quad (12)$$

where $\mathbf{A}, \mathbf{A}' \xleftarrow{\$} \mathbb{Z}_q^{n \times m}$, $\mathbf{R}_0 \xleftarrow{\$} \{-1, 1\}^{m \times m}$, $\mathbf{v}_1, \mathbf{v}'_1 \xleftarrow{\$} \mathbb{Z}_q^m$. It can be seen that the first and the second distributions are $\text{negl}(n)$ -close, by applying Lemma 4 for $(\mathbf{A}^\top | \mathbf{v})^\top \in \mathbb{Z}_{(n+1) \times m}$ and \mathbf{R}_0 . It can also be seen that the second and the third distributions are $\text{negl}(n)$ -close, by applying the same lemma for \mathbf{A} and \mathbf{R}_0 . From the above, we have that the following distributions are statistically close:

$$\begin{aligned}
& (\mathbf{A}, \mathbf{A}\mathbf{R}_0, \mathbf{v}_1^\top, \mathbf{v}'_1^\top \mathbf{R}_{\text{ID}}^\star) \\
&= \left(\mathbf{A}, \mathbf{A}\mathbf{R}_0, \mathbf{v}_1^\top, \mathbf{v}'_1^\top \left(\mathbf{R}_0 + \sum_{(j_1, \dots, j_d) \in S(\text{ID})} \text{TrapEval}(\mathbf{R}_{1, j_1}, \dots, \mathbf{R}_{d, j_d}, y_{1, j_1}, \dots, y_{d, j_d}) \right) \right) \\
&\approx \left(\mathbf{A}, \mathbf{A}\mathbf{R}_0, \mathbf{v}_1^\top, \mathbf{v}'_1^\top + \mathbf{v}_1^\top \left(\sum_{(j_1, \dots, j_d) \in S(\text{ID})} \text{TrapEval}(\mathbf{R}_{1, j_1}, \dots, \mathbf{R}_{d, j_d}, y_{1, j_1}, \dots, y_{d, j_d}) \right) \right) \\
&\approx (\mathbf{A}, \mathbf{A}\mathbf{R}_0, \mathbf{v}_1^\top, \mathbf{v}'_1^\top)
\end{aligned}$$

where $\mathbf{A}, \mathbf{A}' \xleftarrow{\$} \mathbb{Z}_q^{n \times m}$, $\mathbf{R}_0 \xleftarrow{\$} \{-1, 1\}^{m \times m}$, $\mathbf{v}_1, \mathbf{v}'_1 \xleftarrow{\$} \mathbb{Z}_q^m$. The second and the third distributions above are $\text{negl}(n)$ -close by Equation (12). Therefore, we may conclude that

$$|\Pr[X_6] - \Pr[X_7]| = \text{negl}(n).$$

Analysis. From the above, we have

$$\begin{aligned}
\left| \Pr[X_7] - \frac{1}{2} \right| &= \left| \Pr[X_1] - \frac{1}{2} + \sum_{i=1}^6 \Pr[X_{i+1}] - \Pr[X_i] \right| \\
&\geq \left| \Pr[X_1] - \frac{1}{2} \right| - \sum_{i=1}^6 |\Pr[X_{i+1}] - \Pr[X_i]| \\
&\geq \frac{1}{\kappa + 1} \cdot \left(\frac{1}{dn^{\tilde{c}}} \right)^d \cdot \left(\epsilon - \frac{Q}{n^{\tilde{c}}} \right) - \frac{d}{\kappa + 1} \cdot \left(\frac{1}{dn^c} \right)^{d+1} - \text{negl}(n) \\
&\geq \frac{1}{\kappa + 1} \cdot \left(\frac{1}{dn^{\tilde{c}}} \right)^d \cdot \left(\epsilon - \frac{Q}{n^{\tilde{c}}} \right) - \frac{d}{\kappa + 1} \cdot \left(\frac{1}{dn^{\tilde{c}}} \right)^{d+1} - \text{negl}(n) \\
&= \frac{1}{\kappa + 1} \cdot \left(\frac{1}{dn^{\tilde{c}}} \right)^d \cdot (Q + 1) \cdot \left(\frac{\epsilon}{Q + 1} - \frac{1}{n^{\tilde{c}}} \right) - \text{negl}(n) \\
&= \frac{1}{\text{poly}(n)} \cdot \left(\frac{\epsilon}{Q + 1} - \frac{1}{n^{\tilde{c}}} \right) - \text{negl}(n). \tag{13}
\end{aligned}$$

The third inequality above follows from $c \geq \tilde{c}$. Since the challenge ciphertext is independent from the value of coin in Game_7 , we have $\Pr[X_7] = 1/2$ and thus $|\Pr[X_7] - 1/2| = 0$. Therefore, from inequality (13),

$$\frac{\epsilon}{Q + 1} < \frac{1}{n^{\tilde{c}}} + \text{negl}(n)$$

follows. However, this contradicts to Equation (6). \square

To complete the proof of Theorem 3, it remains to show Lemma 8 and 9.

Lemma 8. For any PPT adversary \mathcal{A} , we have

$$\left| \Pr[X_1] - \frac{1}{2} \right| \geq \frac{1}{\kappa + 1} \cdot \left(\frac{1}{dn^{\tilde{c}}} \right)^d \cdot \left(\epsilon - \frac{Q}{n^{\tilde{c}}} \right).$$

Proof. For a sequence of identities $\mathbb{ID} = (\text{ID}^*, \text{ID}_1, \dots, \text{ID}_Q) \in \mathcal{ID}^{Q+1}$, we define $\gamma(\mathbb{ID})$ as

$$\gamma(\mathbb{ID}) = \Pr_{\mathbf{y}}[\mathbf{F}_{\mathbf{y}}(\text{ID}^*) = 0 \wedge \mathbf{F}_{\mathbf{y}}(\text{ID}_1) \neq 0 \wedge \mathbf{F}_{\mathbf{y}}(\text{ID}_2) \neq 0 \wedge \dots \wedge \mathbf{F}_{\mathbf{y}}(\text{ID}_Q) \neq 0]$$

where the probability is taken over $\mathbf{y} = (y_0, \{y_{i,j}\}_{(i,j) \in [d, \ell]})$, which is chosen as specified in Game_1 . To show the lemma, we first show the following claim, which gives an upper and lower bounds for $\gamma(\mathbb{ID})$.

Claim 1. For any $\mathbb{ID} = (\text{ID}^*, \text{ID}_1, \dots, \text{ID}_Q)$ such that $\text{ID}^* \neq \text{ID}_i$ for all $i \in [Q]$,

$$\frac{1}{\kappa + 1} \cdot \left(\frac{1}{dn^{\tilde{c}}} \right)^d \cdot \left(1 - \frac{Q}{n^{\tilde{c}}} \right) \leq \gamma(\mathbb{ID}) \leq \frac{1}{\kappa + 1} \cdot \left(\frac{1}{dn^{\tilde{c}}} \right)^d.$$

Proof. Showing the upper bound of the probability is very easy. For any $\{y_{i,j}\}_{(i,j) \in [d] \times [\ell]}$, there exists exactly one $y_0 \in [-(\kappa + 1)(dn^{\tilde{c}})^d + 1, 0]$ such that $\mathbf{F}_{\mathbf{y}}(\text{ID}^*) = 0$, since for any $\{y_{i,j}\}_{(i,j) \in [d] \times [\ell]}$ and ID , we have

$$0 \leq \sum_{(j_1, \dots, j_d) \in S(\text{ID})} y_{1, j_1} \cdots y_{d, j_d} \leq \sum_{(j_1, \dots, j_d) \in S(\text{ID})} (dn^{\tilde{c}})^d < (\kappa + 1)(dn^{\tilde{c}})^d$$

Therefore, we have

$$\gamma(\mathbb{ID}) \leq \Pr_{\mathbf{y}}[\mathbf{F}_{\mathbf{y}}(\text{ID}^*) = 0] = \frac{1}{\kappa + 1} \cdot \left(\frac{1}{dn^{\tilde{c}}} \right)^d.$$

We then proceed to show the lower bound.

$$\begin{aligned} \gamma(\mathbb{ID}) &= \Pr_{\mathbf{y}}[\mathbf{F}_{\mathbf{y}}(\text{ID}^*) = 0 \wedge \mathbf{F}_{\mathbf{y}}(\text{ID}_1) \neq 0 \wedge \mathbf{F}_{\mathbf{y}}(\text{ID}_2) \neq 0 \wedge \dots \wedge \mathbf{F}_{\mathbf{y}}(\text{ID}_Q) \neq 0] \\ &\geq \Pr_{\mathbf{y}}[\mathbf{F}_{\mathbf{y}}(\text{ID}^*) = 0] - \sum_{i \in [Q]} \Pr_{\mathbf{y}}[\mathbf{F}_{\mathbf{y}}(\text{ID}^*) = 0 \wedge \mathbf{F}_{\mathbf{y}}(\text{ID}_i) = 0] \\ &= \frac{1}{\kappa + 1} \cdot \left(\frac{1}{dn^{\tilde{c}}} \right)^d - \sum_{i \in [Q]} \Pr_{\mathbf{y}}[\mathbf{F}_{\mathbf{y}}(\text{ID}^*) = 0 \wedge \mathbf{F}_{\mathbf{y}}(\text{ID}_i) = 0]. \end{aligned} \quad (14)$$

It suffices to show an upper bound for $\Pr[\mathbf{F}_{\mathbf{y}}(\text{ID}^*) = 0 \wedge \mathbf{F}_{\mathbf{y}}(\text{ID}_i) = 0]$. For $i \in [Q]$, we have

$$\begin{aligned} &\Pr_{\mathbf{y}}[\mathbf{F}_{\mathbf{y}}(\text{ID}^*) = 0 \wedge \mathbf{F}_{\mathbf{y}}(\text{ID}_i) = 0] \\ &= \Pr_{\mathbf{y}}[\mathbf{F}_{\mathbf{y}}(\text{ID}^*) = 0 \wedge \mathbf{F}_{\mathbf{y}}(\text{ID}^*) - \mathbf{F}_{\mathbf{y}}(\text{ID}_i) = 0] \\ &= \Pr_{\mathbf{y}}[\mathbf{F}_{\mathbf{y}}(\text{ID}^*) = 0 \mid \mathbf{F}'_{\mathbf{y}}(\text{ID}^*, \text{ID}_i) = 0] \cdot \Pr_{\mathbf{y}}[\mathbf{F}'_{\mathbf{y}}(\text{ID}^*, \text{ID}_i) = 0] \\ &= \Pr_{\mathbf{y}} \left[y_0 = - \sum_{(j_1, \dots, j_d) \in S(\text{ID}^*)} y_{1, j_1} \cdots y_{d, j_d} \mid \mathbf{F}'_{\mathbf{y}}(\text{ID}^*, \text{ID}_i) = 0 \right] \cdot \Pr_{\mathbf{y}}[\mathbf{F}'_{\mathbf{y}}(\text{ID}^*, \text{ID}_i) = 0] \end{aligned}$$

$$= \frac{1}{\kappa + 1} \cdot \left(\frac{1}{dn^{\tilde{c}}} \right)^d \cdot \Pr_{\mathbf{y}}[F'_{\mathbf{y}}(\text{ID}^*, \text{ID}_i) = 0]. \quad (15)$$

In the above, we defined $F'_{\mathbf{y}}(\text{ID}^*, \text{ID}_i)$ as

$$\begin{aligned} F'_{\mathbf{y}}(\text{ID}^*, \text{ID}_i) &:= F_{\mathbf{y}}(\text{ID}^*) - F_{\mathbf{y}}(\text{ID}_i) \\ &= \sum_{(j_1, \dots, j_d) \in S(\text{ID}^*)} y_{1,j_1} \cdots y_{d,j_d} - \sum_{(j_1, \dots, j_d) \in S(\text{ID}_i)} y_{1,j_1} \cdots y_{d,j_d}. \end{aligned}$$

The last equation in Equation (15) follows since y_0 is independent from $F'_{\mathbf{y}}(\text{ID}^*, \text{ID}_i)$. (Observe that y_0 does not appear in the definition of $F'_{\mathbf{y}}(\text{ID}^*, \text{ID}_i)$.)

We then finally bound $\Pr_{\mathbf{y}}[F'_{\mathbf{y}}(\text{ID}^*, \text{ID}_i) = 0]$. Since $\text{ID}^* \neq \text{ID}_i$ and S is an injective map, we have $S(\text{ID}^*) \neq S(\text{ID}_i)$. Therefore, there exists $(j_1^*, \dots, j_d^*) \in [\ell]^d$ such that $(j_1^*, \dots, j_d^*) \in S(\text{ID}^*) \triangle S(\text{ID}_i)$, where $S(\text{ID}^*) \triangle S(\text{ID}_i)$ denotes the symmetric difference of $S(\text{ID}^*)$ and $S(\text{ID}_i)$. Thus, $F'_{\mathbf{y}}(\text{ID}^*, \text{ID}_i)$ is not a zero-polynomial when we regard it as a polynomial in indeterminates $\{y_{j,k}\}_{(j,k) \in [d] \times [\ell]}$. Since each $y_{j,k}$ is uniformly random over $[1, dn^{\tilde{c}}]$ and $F'_{\mathbf{y}}(\text{ID}^*, \text{ID}_i)$ is a polynomial with degree d , by the Schwartz-Zippel lemma, it follows that

$$\Pr_{\mathbf{y}}[F'_{\mathbf{y}}(\text{ID}^*, \text{ID}_i) = 0] \leq \frac{d}{dn^{\tilde{c}}} \leq \frac{1}{n^{\tilde{c}}}.$$

By combining this with Equation (14) and (15), the claim follows. \square

We then proceed to show a lower bound for $|\Pr[X_1] - 1/2|$. For $\mathbb{ID} = (\text{ID}^*, \text{ID}_1, \dots, \text{ID}_Q)$ such that $\text{ID}^* \neq \text{ID}_i$ for all $i \in [Q]$, we define γ_{\max} and γ_{\min} as the largest and the smallest value of $\gamma(\mathbb{ID})$ taken over all such \mathbb{ID} , respectively. We define $\mathbf{Q}(\mathbb{ID})$ as the event that \mathcal{A} chooses ID^* as its challenge identity and it makes key extraction queries for $\text{ID}_1, \dots, \text{ID}_Q$. We also define Abort as the event that the challenger aborts. Then, we have

$$\begin{aligned} & \left| \Pr[X_1] - \frac{1}{2} \right| \\ &= \left| \Pr[\text{coin}' = \text{coin}] - \frac{1}{2} \right| \\ &= \left| \sum_{\mathbb{ID}} \Pr[\mathbf{Q}(\mathbb{ID})] \cdot \Pr[\text{coin}' = \text{coin} | \mathbf{Q}(\mathbb{ID})] - \frac{1}{2} \right| \\ &= \left| \sum_{\mathbb{ID}} \Pr[\mathbf{Q}(\mathbb{ID})] \cdot \left(\Pr[\text{coin}' = \text{coin} \wedge \neg \text{Abort} | \mathbf{Q}(\mathbb{ID})] + \Pr[\text{coin}' = \text{coin} \wedge \text{Abort} | \mathbf{Q}(\mathbb{ID})] - \frac{1}{2} \right) \right| \\ &= \left| \sum_{\mathbb{ID}} \Pr[\mathbf{Q}(\mathbb{ID})] \cdot \left(\Pr[\widehat{\text{coin}} = \text{coin} | \mathbf{Q}(\mathbb{ID})] \cdot \gamma(\mathbb{ID}) + \frac{1}{2} \cdot (1 - \gamma(\mathbb{ID})) - \frac{1}{2} \right) \right| \\ &= \left| \sum_{\mathbb{ID}} \gamma(\mathbb{ID}) \cdot \Pr[\mathbf{Q}(\mathbb{ID})] \cdot \left(\Pr[\widehat{\text{coin}} = \text{coin} | \mathbf{Q}(\mathbb{ID})] - \frac{1}{2} \right) \right| \\ &\geq \gamma_{\min} \cdot \epsilon - \frac{\gamma_{\max} - \gamma_{\min}}{2}. \end{aligned}$$

In the third equation above, we used the fact $\sum_{\mathbb{ID}} \Pr[\mathbf{Q}(\mathbb{ID})] = 1$. The fourth equation above follows from the fact that the probability of the abort is $\gamma(\mathbb{ID})$, when conditioned on $\mathbf{Q}(\mathbb{ID})$

(regardless of the value of $\widehat{\text{coin}}$). The last inequality above follows by Lemma 5, since we have

$$\begin{aligned} \left| \sum_{\mathbb{ID}} \Pr[\mathbf{Q}(\mathbb{ID})] \cdot \left(\Pr[\widehat{\text{coin}} = \text{coin} | \mathbf{Q}(\mathbb{ID})] - \frac{1}{2} \right) \right| &= \left| \sum_{\mathbb{ID}} \Pr[\widehat{\text{coin}} = \text{coin} \wedge \mathbf{Q}(\mathbb{ID})] - \frac{1}{2} \right| \\ &= \left| \Pr[\widehat{\text{coin}} = \text{coin}] - \frac{1}{2} \right| = \epsilon \end{aligned}$$

and

$$\sum_{\mathbb{ID}} \left| \Pr[\mathbf{Q}(\mathbb{ID})] \cdot \left(\Pr[\widehat{\text{coin}} = \text{coin} | \mathbf{Q}(\mathbb{ID})] - \frac{1}{2} \right) \right| \leq \sum_{\mathbb{ID}} \Pr[\mathbf{Q}(\mathbb{ID})] \cdot \frac{1}{2} = \frac{1}{2}.$$

We complete the proof of Lemma 8 by observing

$$\begin{aligned} & \gamma_{\min} \cdot \epsilon - \frac{\gamma_{\max} - \gamma_{\min}}{2} \\ & \geq \frac{1}{\kappa + 1} \cdot \left(\frac{1}{dn^{\bar{c}}} \right)^d \cdot \left(1 - \frac{Q}{n^{\bar{c}}} \right) \cdot \epsilon - \frac{1}{2(\kappa + 1)} \cdot \left(\frac{1}{dn^{\bar{c}}} \right)^d \cdot \left(1 - \left(1 - \frac{Q}{n^{\bar{c}}} \right) \right) \\ & \geq \frac{1}{\kappa + 1} \cdot \left(\frac{1}{dn^{\bar{c}}} \right)^d \cdot \left(\epsilon - \frac{Q}{n^{\bar{c}}} \right). \end{aligned}$$

The last inequality follows from $\epsilon \leq 1/2$. \square

Lemma 9. *For any PPT adversary \mathcal{A} , there exists another PPT adversary \mathcal{B} such that*

$$|\Pr[X_5] - \Pr[X_6]| \leq \text{Adv}_{\mathcal{B}}^{\text{dLWE}_{n,m+1,q,\chi}}.$$

In particular, under the $\text{dLWE}_{n,m+1,q,\chi}$ assumption, we have $|\Pr[X_5] - \Pr[X_6]| = \text{negl}(n)$.

Proof. Suppose an adversary \mathcal{A} that has non-negligible advantage in distinguishing Game_5 and Game_6 . We use \mathcal{A} to construct an LWE algorithm denoted \mathcal{B} , which proceeds as follows.

Instance. \mathcal{B} is given the problem instance of LWE $(\mathbf{A}', \mathbf{v}') \in \mathbb{Z}_q^{n \times (m+1)} \times \mathbb{Z}_q^{m+1}$. Let the first column of \mathbf{A}' be $\mathbf{u} \in \mathbb{Z}_q^n$ and the last m column be $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$. It also sets the first coefficient of \mathbf{v}' be v_0 and the last m coefficients be \mathbf{v}_1 .

Setup. To construct master public key mpk , \mathcal{B} first picks \mathbf{y} as in Game_1 . It also picks $\mathbf{R}_0, \mathbf{R}_{i,j} \xleftarrow{\$} \{-1, 1\}^{m \times m}$ and sets \mathbf{B}_0 and $\mathbf{B}_{i,j}$ as Equation (8). Finally, it returns $\text{mpk} = (\mathbf{A}, \mathbf{B}_0, \{\mathbf{B}_{i,j}\}_{(i,j) \in [d,\ell]}, \mathbf{u})$ to \mathcal{A} . \mathcal{B} also picks a random bit $\text{coin} \xleftarrow{\$} \{0, 1\}$ and keeps it secret.

Phase 1 and Phase 2. When \mathcal{A} makes a key extraction query for ID , \mathcal{B} first computes $F_{\mathbf{y}}(\text{ID})$. It aborts and sets $\text{coin}' \xleftarrow{\$} \{0, 1\}$ if $F_{\mathbf{y}}(\text{ID}) = 0$. Otherwise, \mathcal{B} generates the private key as in Game_5 .

Challenge Query. When \mathcal{A} makes the challenge query for the challenge identity ID^* and the message b , \mathcal{B} first computes $F_{\mathbf{y}}(\text{ID}^*)$. Then, it aborts and sets $\text{coin}' \xleftarrow{\$} \{0, 1\}$ if $F_{\mathbf{y}}(\text{ID}^*) \neq 0$. Otherwise, it proceeds as follows. If $\text{coin} = 0$, it computes \mathbf{R}_{ID^*} and picks $\mathbf{x}_2 \xleftarrow{\$} [-B', B']^m$. Then, it sets the challenge ciphertext as

$$c_0 = v_0 + b \cdot \lceil q/2 \rceil, \quad \mathbf{c}_1^\top = (\mathbf{v}_1^\top | \mathbf{v}_1^\top \mathbf{R}_{\text{ID}^*}) + (\mathbf{0}_m^\top | \mathbf{x}_2^\top)$$

and returns $C^* = (c_0, \mathbf{c}_1)$ to \mathcal{A} . In the case of $\text{coin} = 1$, \mathcal{B} picks $c_0 \xleftarrow{\$} \mathbb{Z}_q$, $\mathbf{c}_1 \xleftarrow{\$} \mathbb{Z}_q^{2m}$ and returns the challenge ciphertext $C^* = (c_0, \mathbf{c}_1)$ to \mathcal{A} .

Guess. At last, \mathcal{A} outputs its guess $\widehat{\text{coin}}$ (if the abort condition has not been satisfied). Then, \mathcal{B} sets $\text{coin}' = \widehat{\text{coin}}$. Finally, \mathcal{B} outputs 1 if $\text{coin}' = \text{coin}$ and 0 otherwise.

Analysis. We now show that \mathcal{B} perfectly simulates the view of \mathcal{A} in Game_5 if $(\mathbf{A}', \mathbf{v}')$ is a valid LWE sample (i.e., $\mathbf{v}'^\top = \mathbf{s}^\top \mathbf{A}' + \mathbf{x}^\top$ for $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n$ and $\mathbf{x} \xleftarrow{\$} \chi^{m+1}$), and Game_6 if $\mathbf{v}' \xleftarrow{\$} \mathbb{Z}_q^{m+1}$. Note that these games differ only in the generation of the challenge ciphertext in the case of $\text{coin} = 0$. Furthermore, it is easy to see that the simulation of the master public key, **Phase 1**, **Phase 2**, and the challenge ciphertext for the case of $\text{coin} = 1$ are perfect. Therefore, in the following, we focus on the generation of the challenge ciphertext in the case of $\text{coin} = 0$.

We first show that if $(\mathbf{A}', \mathbf{v}')$ is a valid LWE sample, i.e., $\mathbf{v}'^\top = \mathbf{s}^\top \mathbf{A}' + \mathbf{x}^\top$ for $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n$ and $\mathbf{x} \xleftarrow{\$} \chi^{m+1}$, the distribution of the challenge ciphertext corresponds to that of Game_5 . Let us denote $\mathbf{x}^\top = (x_0, \mathbf{x}_1^\top)$ and assume that $\mathbf{F}_y(\text{ID}^*) = 0$ holds. Then, we have

$$\begin{aligned} c_0 &= v_0 + b \cdot \lceil q/2 \rceil = (\mathbf{u}^\top \mathbf{s} + x_0) + b \cdot \lceil q/2 \rceil \quad \text{and} \\ \mathbf{c}_1 &= (\mathbf{v}_1^\top | \mathbf{v}_1^\top \mathbf{R}_{\text{ID}^*}) + (\mathbf{0}_m^\top | \mathbf{x}_2^\top) \\ &= (\mathbf{s}^\top \mathbf{A} + \mathbf{x}_1^\top | (\mathbf{s}^\top \mathbf{A} + \mathbf{x}_1^\top) \mathbf{R}_{\text{ID}^*}) + (\mathbf{0}_m^\top | \mathbf{x}_2^\top) \\ &= \mathbf{s}^\top (\mathbf{A} | \mathbf{A} \mathbf{R}_{\text{ID}^*}) + (\mathbf{x}_1^\top | \mathbf{x}_1^\top \mathbf{R}_{\text{ID}^*} + \mathbf{x}_2^\top) \\ &= \mathbf{s}^\top (\mathbf{A} | \mathbf{H}(\text{ID}^*)) + (\mathbf{x}_1^\top | \mathbf{x}_1^\top \mathbf{R}_{\text{ID}^*} + \mathbf{x}_2^\top). \end{aligned}$$

The last equation follows because $\mathbf{F}_y(\text{ID}^*) = 0$. Therefore, the challenge ciphertext is distributed as in Game_5 in this case. It is easy to see that the challenge ciphertext is distributed as in Game_6 , if $\mathbf{v}' \xleftarrow{\$} \mathbb{Z}_q^{m+1}$.

Therefore, we have $\text{Adv}_{\mathcal{B}}^{\text{dLWE}_{n,m+1,q,\chi}} = |\Pr[X_5] - \Pr[X_6]|$ as desired. \square

5.5 Multi-bit Encryption

Here, we explain that our scheme can be extended to deal with multi-bit messages without much increasing the sizes of public parameters and ciphertexts, similarly to [39, 1]. To modify the scheme so that it can encrypt messages with N -bit, we replace $\mathbf{u} \in \mathbb{Z}_q^n$ in mpk with $\mathbf{u}_1, \dots, \mathbf{u}_N \in \mathbb{Z}_q^n$. The component $c_0 = \langle \mathbf{u}, \mathbf{s} \rangle + x_0 + b \lceil \frac{q}{2} \rceil$ in the ciphertext is replaced with $\mathbf{c}_0 = \{ \langle \mathbf{u}_i, \mathbf{s} \rangle + x_{0,i} + b_i \lceil \frac{q}{2} \rceil \}_{i=1}^N$ where $x_{0,i} \xleftarrow{\$} \chi$ and $b_i \in \{0, 1\}$ is the i -th bit of the message. Furthermore, the private key is changed to be short vectors $\mathbf{e}_1, \dots, \mathbf{e}_N \in \mathbb{Z}^m$ such that $(\mathbf{A} | \mathbf{H}(\text{ID})) \mathbf{e}_i = \mathbf{u}_i$ for $i = 1, \dots, N$. We can prove the security for the variant from $\text{dLWE}_{n,m+N,q,\chi}$ by naturally extending the proof of Theorem 3.

As for the efficiency, the size of the master public key and the ciphertexts become $O((\ell m + N)n \log q)$ and $O((m + N) \log q)$ respectively, and these are asymptotically the same as the case of single-bit encryption when $N < O(m)$. The case of $N > O(m)$ can also be handled without increasing the size of parameters, by employing the KEM-DEM approach. Namely, we encrypt a random ephemeral key of sufficient length (e.g., $O(n)$) by IBE and then encrypt the message by the ephemeral key using a symmetric cipher.

6 Comparisons and Discussions

From the PIBE scheme in Section 5, we can obtain the following new IBE schemes:

- By setting $c = \omega(1)$, we obtain adaptively anonymous IBE by Theorem 1. However, we have to rely on super-polynomial LWE assumption, namely, $\text{dLWE}_{n,m,q,\chi}$ with $q/\chi_{\max} = n^{\omega(1)}$.

- By applying PIBE-to-IBE conversion in Section 4.2 to our PIBE in Section 5, we obtain (non-anonymous) adaptively secure IBE from polynomial LWE. More precisely, the security of the scheme can be proven under the assumption that $\text{dLWE}_{n,m,q,\chi}$ is hard for all $q/\chi_{\max} = \text{poly}(n)$.

For concreteness, we would set $c(n) = O(\log \log n)$ in the first construction, and $c(n) = \log \log n$ and $\eta(n) = \log \log n$ for the second construction. Ignoring poly-logarithmic factors hidden in the asymptotic notation $\tilde{O}(\cdot)$, both of our schemes achieve the best efficiency among existing adaptively secure IBE schemes. See Table 1 for the comparison. Comparing in more details, ciphertexts and private keys of both of our schemes are longer than [1, 12] by a super-constant factor. This is because we need to use super polynomially large q . On the other hand, in both of our schemes, the sizes of master public keys are asymptotically smaller than [1, 12], even though we have to use larger q . This is because we require smaller number of basic matrices in the master public keys. Our first scheme is more efficient than our second scheme by super-constant factors, because the conversion in Section 4.2 incurs super-constant efficiency loss. We also note that our security reduction is very loose even compared to non-tight reduction of [1, 12]. The security degrades exponentially as d grows. Therefore, in order to have polynomial reduction, we have to set d to be a (possibly small) constant.

Table 1: Comparison of IBE from the LWE assumption in the Standard Model.

Schemes	$ \text{mpk} $	$ C $	$ \text{sk}_{\text{ID}} $	Anonymous?	Selective or Adaptive	q/χ_{\max} for LWE Assumption
[1]	$\tilde{O}(n^2)$	$\tilde{O}(n)$	$\tilde{O}(n)$	Yes	Selective	Fixed $\text{poly}(n)$
[16]	$\tilde{O}(n^2\kappa)$	$\tilde{O}(n\kappa)$	$\tilde{O}(n^2)$	Yes	Adaptive	Fixed $\text{poly}(n)$
[1]+[12]*	$\tilde{O}(n^2\kappa)$	$\tilde{O}(n)$	$\tilde{O}(n)$	Yes	Adaptive	Fixed $\text{poly}(n)$
Ours: Sec. 5 + Th. 1.	$\tilde{O}(n^2\kappa^{1/d})$	$\tilde{O}(n)$	$\tilde{O}(n)$	Yes	Adaptive	$n^{\omega(1)}$
Ours: Sec. 5 + Th. 2.	$\tilde{O}(n^2\kappa^{1/d})$	$\tilde{O}(n)$	$\tilde{O}(n)$	No	Adaptive	All $\text{poly}(n)$

We compare IBE schemes from the LWE assumption in the standard model. $|\text{mpk}|$, $|C|$, and $|\text{sk}_{\text{ID}}|$ show the size of the master public keys, ciphertexts, and private keys, respectively. In the table, κ denotes the length of the identity (which corresponds to the output length of the collision resistant hash if we first hash the bit string representing identity in the scheme). $d \in \mathbb{N}$ is a flexible constant, which can be set to be any value. “Anonymous?” shows whether the scheme is anonymous. “Selective/Adaptive” shows whether the scheme is selectively secure or adaptively secure. “ q/χ_{\max} ” for LWE assumption refers to the ratio of the modulus to the error size of the underlying LWE assumption used in the security reduction. “Fixed $\text{poly}(n)$ ” means that the corresponding scheme is proven secure under the LWE assumption with q/χ_{\max} being some fixed polynomial (e.g., n^3). “All $\text{poly}(n)$ ” mean that we have to assume the LWE assumption for all polynomial q/χ_{\max} .

* In the security proof for the adaptively secure variant of IBE in [1], we have a restriction that $q > Q$. Namely, only bounded form of the security is proven. This restriction is removed in the refined analysis due to Boyen [12].

Acknowledgement. The author would like to thank all members of the study group “Shin-Akarui-Angou-Benkyou-Kai” for fruitful discussion. In particular, the author thanks Shuichi Katsumata for his comments on improving the presentation, Goichiro Hanaoka and Jacob. C. N. Schuldt for their helpful advice in the rebuttal phase. The author also thanks the anonymous reviewers of Eurocrypt 2016 for their insightful comments.

References

- [1] S. Agrawal, D. Boneh, and X. Boyen. Efficient Lattice (H)IBE in the Standard Model. In *EUROCRYPT*, pp. 553–572, 2010.
- [2] S. Agrawal, D. Boneh, and X. Boyen. Lattice Basis Delegation in Fixed Dimension and Shorter-Ciphertext Hierarchical IBE. In *CRYPTO*, pp. 98–115, 2010.
- [3] M. Ajtai. Generating Hard Instances of the Short Basis Problem. *ICALP*, pp. 1–9, 1999.
- [4] J. Alperin-Sheriff. Short Signatures with Short Public Keys from Homomorphic Trapdoor Functions. In *PKC*, pp. 236–255, 2015.
- [5] J. Alwen and C. Peikert. Generating Shorter Bases for Hard Random Lattices. In *STACS*, pp. 75–86, 2009.
- [6] M. Bellare and T. Ristenpart. Simulation without the artificial abort: Simplified proof and improved concrete security for waters’ IBE scheme. In *EUROCRYPT*, pp. 407–424, 2009.
- [7] F. Böhl, D. Hofheinz, T. Jager, J. Koch, J. Seo, and C. Striecks. Practical Signatures from Standard Assumptions. In *EUROCRYPT*, pp. 461–485, 2013.
- [8] D. Boneh and X. Boyen. Efficient selective-id secure identity-based encryption without random oracles. In *EUROCRYPT*, pp. 223–238, 2004.
- [9] D. Boneh and X. Boyen. Secure identity based encryption without random oracles. In *CRYPTO*, pp. 443–459, 2004.
- [10] D. Boneh and M. Franklin. Identity-based encryption from the weil pairing. In *CRYPTO*, pp. 213–229, 2001.
- [11] D. Boneh, C. Gentry, S. Gorbunov, S. Halevi, V. Nikolaenko, G. Segev, V. Vaikuntanathan, and D. Vinayagamurthy. Fully Key-Homomorphic Encryption, Arithmetic Circuit ABE and Compact Garbled Circuits. In *EUROCRYPT*, pp. 533–556, 2014.
- [12] X. Boyen. Lattice Mixing and Vanishing Trapdoors: A Framework for Fully Secure Short Signatures and More. In *PKC*, pp. 499–517, 2010.
- [13] Z. Brakerski, A. Langlois, C. Peikert, O. Regev, and D. Stehlé. Classical hardness of learning with errors. In *STOC*, pp. 575–584, 2013.
- [14] Z. Brakerski and V. Vaikuntanathan. Lattice-based FHE as secure as PKE. In *ITCS*, pp. 1–12, 2014.
- [15] R. Canetti, S. Halevi, and J. Katz. A forward-secure public-key encryption scheme. In *EUROCRYPT*, pp. 255–271, 2003.
- [16] D. Cash, D. Hofheinz, E. Kiltz, and C. Peikert. Bonsai Trees, or How to Delegate a Lattice Basis. In *EUROCRYPT*, pp. 523–552, 2010.
- [17] J. Chen and H. Wee. Fully, (Almost) Tightly Secure IBE and Dual System Groups. In *CRYPTO*, pp. 435–460, 2013.
- [18] C. Cocks. An identity based encryption scheme based on quadratic residues. In *IMA Int. Conf.*, pp. 360–363, 2001.
- [19] Y. Dodis, J. Katz, S. Xu, and M. Yung. Key-Insulated Public Key Cryptosystems. In *EUROCRYPT*, pp. 65–82, 2002.
- [20] Y. Dodis, R. Ostrovsky, L. Reyzin, and A. Smith. Fuzzy Extractors: How to Generate Strong Keys from Biometrics and Other Noisy Data. In *SIAM J. Comput.* 38(1), pp. 97–139, 2008.
- [21] N. Döttling and D. Schröder. Efficient Pseudorandom Functions via On-the-Fly Adaptation. In *CRYPTO*, pp. 329–350, 2015.
- [22] L. Ducas, V. Lyubashevsky, and T. Prest. Efficient Identity-Based Encryption over NTRU Lattices. In *ASIACRYPT* (2), pp. 22–41, 2014.
- [23] L. Ducas and D. Micciancio. Improved Short Lattice Signatures in the Standard Model. In *CRYPTO*, pp. 335–352, 2014.

- [24] C. Gentry. Practical identity-based encryption without random oracles. In *EUROCRYPT*, pp. 445–464, 2006.
- [25] C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pp. 197–206, 2008.
- [26] S. Goldwasser, A. B. Lewko, and D. A. Wilson. Bounded-Collusion IBE from Key Homomorphism. In *TCC*, pp. 564–581, 2012.
- [27] S. Gorbunov, V. Vaikuntanathan, and H. Wee. Attribute-based encryption for circuits. In *STOC*, pp. 545–554, 2013.
- [28] S. Gorbunov and D. Vinayagamurthy. Riding on Asymmetry: Efficient ABE for Branching Programs. In *ASIACRYPT*, pp. 550–574, 2015.
- [29] S. Heng and K. Kurosawa. k -Resilient Identity-Based Encryption in the Standard Model. In *CT-RSA*, pp. 67–80, 2004.
- [30] S. Hohenberger and B. Waters. Short and Stateless Signatures from the RSA Assumption. In *CRYPTO*, pp. 654–670, 2009.
- [31] C. S. Jutla, A. Roy: Shorter Quasi-Adaptive NIZK Proofs for Linear Subspaces. In *ASIACRYPT (1)* pp. 1–20, 2013.
- [32] V. Lyubashevsky and D. Micciancio. Asymptotically Efficient Lattice-Based Digital Signatures. In *TCC*, pp. 37–54, 2008.
- [33] V. Lyubashevsky, C. Peikert, and O. Regev. On Ideal Lattices and Learning with Errors over Rings. In *EUROCRYPT*, pp. 1–23, 2010.
- [34] V. Lyubashevsky, C. Peikert, and O. Regev. A Toolkit for Ring-LWE Cryptography. In *EUROCRYPT*, pp. 35–54, 2013.
- [35] D. Micciancio and C. Peikert. Trapdoors for Lattices: Simpler, Tighter, Faster, Smaller. In *EUROCRYPT*, pp. 700–718, 2012.
- [36] D. Naccache. Secure and *practical* identity-based encryption. In *IET Information Security*, volume 1(2): pp. 59–64, 2007.
- [37] C. Peikert. Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In *STOC*, pp. 333–342, 2009.
- [38] C. Peikert. A decade of lattice cryptography. *IACR Cryptology ePrint Archive*, Report 2015/939.
- [39] C. Peikert, V. Vaikuntanathan, and B. Waters. A Framework for Efficient and Composable Oblivious Transfer. In *CRYPTO*, pp. 554–571, 2008.
- [40] O. Regev. On Lattices, Learning with Errors, Random Linear Codes, and Cryptography. In *STOC*, pp. 843–873, 2005.
- [41] R. Sakai, K. Ohgishi, and M. Kasahara. Cryptosystems based on pairing over elliptic curve. In *The 2000 Symposium on Cryptography and Information Security*, 2000. (in Japanese).
- [42] A. Shamir. Identity-based cryptosystems and signature schemes. In *CRYPTO*, pp. 47–53, 1984.
- [43] A. Shamir and Y. Tauman. Improved Online/Offline Signature Schemes. In *CRYPTO*, pp. 355–367, 2001.
- [44] K. Singh, C. Pandu Rangan, and A. K. Banerjee. Adaptively Secure Efficient Lattice (H)IBE in Standard Model with Short Public Parameters. In *SPACE*, pp. 153–172, 2012.
- [45] D. Stehlé, R. Steinfeld, K. Tanaka, and K. Xagawa. Efficient Public Key Encryption Based on Ideal Lattices. In *ASIACRYPT*, pp. 617–635, 2009.
- [46] S. Tessaro and D. A. Wilson. Bounded-Collusion Identity-Based Encryption from Semantically-Secure Public-Key Encryption: Generic Constructions with Short Ciphertexts. In *PKC*, pp. 257–274, 2014.
- [47] B. Waters. Efficient identity-based encryption without random oracles. In *EUROCRYPT*, pp. 114–127, 2005.
- [48] B. Waters. Dual system encryption: Realizing fully secure IBE and HIBE under simple assumptions. In *CRYPTO*, pp. 619–636, 2009.
- [49] S. Yamada, G. Hanaoka, N. Kunihiro. Two-Dimensional Representation of Cover Free Families and Its Applications: Short Signatures and More. In *CT-RSA*, pp. 260–277, 2012.

A Omitted Proofs

A.1 The Proof of Lemma 5

Proof. We have

$$\begin{aligned}
\left| \sum_{i=1}^n \gamma_i a_i \right| &= \left| \sum_{i=1}^n \gamma_{\min} a_i + \sum_{i=1}^n (\gamma_i - \gamma_{\min}) a_i \right| \\
&\geq \left| \sum_{i=1}^n \gamma_{\min} a_i \right| - \sum_{i=1}^n |(\gamma_i - \gamma_{\min}) a_i| \\
&\geq \gamma_{\min} \cdot \epsilon - (\gamma_{\max} - \gamma_{\min}) \sum_{i=1}^n |a_i| \\
&\geq \gamma_{\min} \cdot \epsilon - (\gamma_{\max} - \gamma_{\min})/2.
\end{aligned}$$

□

A.2 The Proof of Lemma 6

Proof. We prove it by induction. The base case (the case of $d = 1$) is trivial. Therefore, let us assume the hypothesis for $d - 1$ where $d \geq 2$. Then, we have

$$\|\mathbf{R}''\|_2 \leq m(d-1)\delta^{d-2} \quad \text{and} \quad \text{PubEval}_{d-1}(\mathbf{B}_2, \dots, \mathbf{B}_d) = \mathbf{A}\mathbf{R}'' + y_2 \cdots y_d \mathbf{G}$$

for efficiently computable \mathbf{R}'' . Therefore, by the definition of PubEval_d , we have

$$\begin{aligned}
&\text{PubEval}_d(\mathbf{B}_1, \dots, \mathbf{B}_d) \\
&= (\mathbf{A}\mathbf{R}_1 + y_1 \mathbf{G}) \cdot \mathbf{G}^{-1}(\text{PubEval}_{d-1}(\mathbf{B}_2, \dots, \mathbf{B}_d)) \\
&= \mathbf{A}\mathbf{R}_1 \cdot \mathbf{G}^{-1}(\text{PubEval}_{d-1}(\mathbf{B}_2, \dots, \mathbf{B}_d)) + y_1 \cdot \text{PubEval}_{d-1}(\mathbf{B}_2, \dots, \mathbf{B}_d) \\
&= \mathbf{A}\mathbf{R}_1 \cdot \mathbf{G}^{-1}(\text{PubEval}_{d-1}(\mathbf{B}_2, \dots, \mathbf{B}_d)) + y_1(\mathbf{A}\mathbf{R}'' + y_2 \cdots y_d \mathbf{G}) \\
&= \mathbf{A}(\mathbf{R}_1 \cdot \mathbf{G}^{-1}(\text{PubEval}_{d-1}(\mathbf{B}_2, \dots, \mathbf{B}_d)) + y_1 \mathbf{R}'') + y_1 y_2 \cdots y_d \mathbf{G}.
\end{aligned}$$

It can be seen that Equation (5) holds by setting

$$\mathbf{R}' = \mathbf{R}_1 \cdot \mathbf{G}^{-1}(\text{PubEval}_{d-1}(\mathbf{B}_2, \dots, \mathbf{B}_d)) + y_1 \mathbf{R}''.$$

It is clear that it can be efficiently computable. Furthermore, we have

$$\begin{aligned}
\|\mathbf{R}'\|_2 &\leq \|\mathbf{R}_1\|_2 \cdot \|\mathbf{G}^{-1}(\text{PubEval}_{d-1}(\mathbf{B}_2, \dots, \mathbf{B}_d))\|_2 + |y_1| \cdot \|\mathbf{R}''\|_2 \\
&\leq m \cdot m + \delta \cdot m(d-1)\delta^{d-2} \leq m \cdot \delta^{d-1} + m(d-1)\delta^{d-1} \leq md\delta^{d-1}.
\end{aligned}$$

□

B Application of Our Technique to Attribute-Based Encryption

Here, we discuss another application of our technique to the context of the attribute-based encryption.

Recently, Gorbunov and Vinayagamurthy [28] proposed a new attribute-based encryption for branching programs from the LWE assumption, by cleverly adapting a technique from the fully homomorphic encryption [14]. In particular, they showed the following:

Theorem 4. *There exists selectively-secure attribute-based encryption for branching programs assuming $\text{dLWE}_{n,m,q,\chi}$ is hard for $\chi = D_{\mathbb{Z},\sqrt{n}}$, $m = O(n \log q)$, and $q = \tilde{O}(n^7 \cdot L^2)$, where ℓ and L denote the length of the input to branching programs and the upper bound for the length of the branching programs, respectively. Furthermore, the running times of the setup algorithm, key generation algorithm, encryption algorithm, and decryption algorithm are bounded by $\text{poly}(n, \ell, \log L)$. In particular, the size of the private keys is compact: It is bounded by $\text{poly}(n, \log L)$ [‡], which depends on L only logarithmically.*

As corollaries, we have the following, which provide trade-off between the functionality and the underlying assumptions:

- By setting $L = n^{\omega(1)}$, we obtain ABE for branching programs with small private keys, which can deal with essentially unbounded length branching programs. However, we have to rely on the LWE assumption with $q/\chi_{\max} = n^{\omega(1)}$ to prove the security.
- By setting $L = O(n^c)$ for some constant c , we obtain ABE for branching programs with small private keys from the LWE assumption with $q/\chi_{\max} = \tilde{O}(n^{2c+6.5})$, which is a weaker assumption than the above. However, the length of the branching programs is bounded by L .

A natural question is whether we can obtain the best of both worlds. In this section, we show that it is possible to obtain an ABE scheme for branching programs with compact private keys that can deal with *unbounded length* branching programs only assuming the LWE assumption with $q/\chi_{\max} = n^c$ for *all constant* $c \in \mathbb{N}$. The assumption is weaker than the LWE assumption with $q/\chi_{\max} = n^{\omega(1)}$. We obtain such a scheme by applying simple conversion that is similar to our PIBE-to-IBE conversion in Section 4.2 to the scheme in [28]. Namely, we run several instances of the scheme of [28] with different sizes of L . The size of L varies from a very small polynomial to a super-polynomial. The conversion is slightly more complex compared to our PIBE-to-IBE conversion, reflecting the fact that we are dealing with ABE, which is more complex object than IBE.

In the following, we first introduce necessary definitions. Then, we present our conversion. Finally, we explain the property of our new scheme obtained by applying the conversion to the ABE scheme in [28].

B.1 Attribute-Based Encryption

Here, we define attribute-based encryption.

Syntax. Let $R : \mathcal{X} \times \mathcal{Y} \rightarrow \{0, 1\}$ be a relation family where \mathcal{X} and \mathcal{Y} denote “key attribute” and “ciphertext attribute” spaces. These spaces are equipped with a function $|\cdot| : \mathcal{X} \cup \mathcal{Y} \rightarrow \mathbb{N}$, which represents the size of the attribute. An ABE scheme for the relation R is defined by the following algorithms:

Setup($1^n, 1^\ell, L$) \rightarrow (mpk, msk): The setup algorithm takes as input a security parameter 1^n , the length ℓ (in a unary form) of the size of ciphertext attributes, and the bound L (in a binary form) for the size of key attributes and outputs a master public key mpk and a master secret key msk.

[‡] When measuring the size of private keys, we ignore the size of the representation for the branching programs that is associated with the key. If we take it into account, the size is $L + \text{poly}(n, \log L)$.

$\text{KeyGen}(\text{mpk}, \text{msk}, x) \rightarrow \text{sk}_x$: The key generation algorithm takes as input the master public key mpk , the master secret key msk , and an key attribute $x \in \mathcal{X}$ such that $|x| \leq L$. It outputs a private key sk_x . We assume that x is implicitly included in sk_x .

$\text{Encrypt}(\text{mpk}, y, M) \rightarrow C$: The encryption algorithm takes as input the master public key mpk , a ciphertext attribute $y \in \mathcal{Y}$ such that $|y| = \ell$, and a message M . It outputs a ciphertext C . We assume that y is implicitly included in C . (Namely, we do not consider anonymity here.)

$\text{Decrypt}(\text{mpk}, \text{sk}_x, C) \rightarrow M$ or \perp : The decryption algorithm takes as input the master public key mpk , a private key sk_x , and a ciphertext C . It outputs the message M or \perp which means that the ciphertext is not in a valid form.

The reason why the setup algorithm takes L in a *binary* form (rather than in a *unary* form) as input is that we will treat a scheme with super-polynomial L , which leads to a scheme that can deal with essentially unbounded size of key attributes. On the other hand, ℓ is bounded by some polynomial.

Correctness. We require (standard) correctness of decryption: that is, for all $n, \ell, L \in \mathbb{N}$, all $x \in \mathcal{X}$ with $|x| \leq L$, $y \in \mathcal{Y}$ with $|y| = \ell$, and all M in the specified message space,

$$\Pr[\text{Decrypt}(\text{mpk}, \text{sk}_x, \text{Encrypt}(\text{mpk}, M, y)) = M] = 1 - \text{negl}(n)$$

holds if $R(x, y) = 1$. Here, the probability is taken over the randomness used in $(\text{mpk}, \text{msk}) \xleftarrow{\$} \text{Setup}(1^n, 1^\ell, L)$, $\text{sk}_x \xleftarrow{\$} \text{KeyGen}(\text{msk}, \text{mpk}, x)$, and $\text{Encrypt}(\text{mpk}, M, y)$.

Relation for Branching Programs. Since the precise definition of branching programs is not necessary for our purpose, we omit this and refer to [28]. In the scheme in [28], the key attribute space \mathcal{X} is a set of all branching programs and the ciphertext attribute space \mathcal{Y} is $\{0, 1\}^*$, which corresponding to the input to the branching programs. For $x \in \mathcal{X}$, $|x|$ is the length of the branching program x and for $y \in \mathcal{Y}$, $|y|$ is the length of the bit string y . We note that the relation family R for branching programs satisfies the following (natural) requirement:

There exists a special element $\top \in \mathcal{X}$ such that $|\top| = 1$ and $R(\top, y) = 1$ for all $y \in \mathcal{Y}$.

We will use this property in Section B.2.

Security. We now define the selective security for an ABE scheme Π . This security notion is defined by the following game between a challenger and an adversary \mathcal{A} .

- **Setup.** At the outset of the game, given 1^n and 1^ℓ as an input, the adversary \mathcal{A} chooses the challenge attribute y^* , which satisfies $|y^*| = \ell$. Then, the challenger runs $\text{Setup}(1^n, 1^\ell, L) \rightarrow (\text{mpk}, \text{msk})$ and gives mpk to \mathcal{A} .

- **Phase 1.** \mathcal{A} may adaptively make key-extraction queries. If \mathcal{A} submits $x \in \mathcal{X}$ such that $|x| \leq L$ and $R(x, y^*) = 0$ to the challenger, the challenger returns $\text{sk}_x \leftarrow \text{KeyGen}(\text{mpk}, \text{msk}, x)$.

- **Challenge Phase.** \mathcal{A} outputs two messages M_0, M_1 on which it wishes to be challenged. Then, the challenger picks a random coin $\text{coin} \xleftarrow{\$} \{0, 1\}$, runs $\text{Encrypt}(\text{mpk}, y^*, M_{\text{coin}}) \rightarrow C^*$, and gives the challenge ciphertext C^* to \mathcal{A} .

- **Phase 2.** After the challenge query, \mathcal{A} may continue to make key-extraction queries, with the same restriction as **Phase 1**.

- **Guess.** Finally, \mathcal{A} outputs a guess $\widehat{\text{coin}}$ for coin . The advantage of \mathcal{A} is defined as

$$\text{Adv}_{\mathcal{A}, \Pi}^{\text{IBE}} = \left| \Pr[\widehat{\text{coin}} = \text{coin}] - \frac{1}{2} \right|.$$

We say that Π is selectively-secure, if the advantage of any PPT \mathcal{A} is negligible.

B.2 Our Transformation

Here, we show a conversion from an ABE scheme $\Pi = (\text{ABE.Setup}, \text{ABE.KeyGen}, \text{ABE.Encrypt}, \text{ABE.Decrypt})$ for a relation family R from another ABE scheme $\Pi' = (\text{ABE.Setup}', \text{ABE.KeyGen}', \text{ABE.Encrypt}', \text{ABE.Decrypt}')$ for the same relation. A nice property of the resulting scheme Π' is that it can deal with unbounded size of key attributes, and at the same time, its security can be reduced to that of Π with polynomial L (rather than super-polynomial L). By applying this transformation to the ABE scheme for branching programs in [28], we obtain an ABE with a new property, as we will discuss in Section B.3. In the following, let $\eta(n)$ be any function such that $\eta(n) = \omega(1)$ (e.g., $\eta(n) = \log \log(n)$). We also let the message space of Π and Π' be $\{0, 1\}^{\ell_M}$ for some $\ell_M \in \mathbb{N}$.

$\text{ABE.Setup}'(1^n, 1^\ell, L := n^\eta)$: It runs $\text{ABE.Setup}(1^n, 1^\ell, n^i) \rightarrow (\text{mpk}^{(i)}, \text{msk}^{(i)})$ for $i = 1, \dots, \eta$. It outputs

$$\text{mpk} = (\text{mpk}^{(1)}, \text{mpk}^{(2)}, \dots, \text{mpk}^{(\eta)}) \quad \text{and} \quad \text{msk} = (\text{msk}^{(1)}, \text{msk}^{(2)}, \dots, \text{msk}^{(\eta)}).$$

$\text{ABE.KeyGen}'(\text{mpk}, \text{msk}, \mathbf{x})$: Let us assume that $|\mathbf{x}| < n^\eta$ holds. Then, there exists unique $i^* \in [\eta]$ such that

$$n^{i^*-1} \leq |\mathbf{x}| < n^{i^*}.$$

The key generation algorithm runs $\text{ABE.KeyGen}(\text{mpk}^{(i)}, \text{msk}^{(i)}, \top) \rightarrow \text{sk}_\top^{(i)}$ for $i = 1, 2, \dots, i^* - 1$ and $\text{ABE.KeyGen}(\text{mpk}^{(i)}, \text{msk}^{(i)}, \mathbf{x}) \rightarrow \text{sk}_\mathbf{x}^{(i)}$ for $i = i^*, i^* + 1, \dots, \eta$. Then it outputs

$$\text{sk}_\mathbf{x} = (\text{sk}_\top^{(1)}, \text{sk}_\top^{(2)}, \dots, \text{sk}_\top^{(i^*-1)}, \text{sk}_\mathbf{x}^{(i^*)}, \text{sk}_\mathbf{x}^{(i^*+1)}, \dots, \text{sk}_\mathbf{x}^{(\eta)}).$$

Note that we can use any polynomial-size \mathbf{x} as an key attribute (in asymptotic sense), since n^η is super-polynomial. In other words, the scheme can deal with *unbounded* size of \mathbf{x} .

$\text{ABE.Encrypt}(\text{mpk}, \mathbf{y}, \mathbf{M})$: To encrypt $\mathbf{M} = \{0, 1\}^{\ell_M}$, it picks random $\mathbf{M}^{(i)} \in \{0, 1\}^{\ell_M}$ for $i \in [\eta]$ subject to constraint that $\mathbf{M} = \bigoplus_{i=1}^\eta \mathbf{M}^{(i)}$. Then it runs

$$\text{ABE.Encrypt}(\text{mpk}^{(i)}, \mathbf{y}, \mathbf{M}^{(i)}) \rightarrow \mathbf{C}^{(i)} \quad \text{for } i = 1, \dots, \eta.$$

Finally, it outputs the ciphertext $\mathbf{C} = (\mathbf{C}^{(1)}, \dots, \mathbf{C}^{(\eta)})$.

$\text{Decrypt}(\text{mpk}, \text{sk}_\mathbf{x}, \mathbf{C})$: It first parses the ciphertext and the private key as $\mathbf{C} \rightarrow (\mathbf{C}^{(1)}, \dots, \mathbf{C}^{(\eta)})$ and $\text{sk}_\mathbf{x} \rightarrow (\text{sk}_\top^{(1)}, \dots, \text{sk}_\top^{(i^*-1)}, \text{sk}_\mathbf{x}^{(i^*)}, \dots, \text{sk}_\mathbf{x}^{(\eta)})$. Then, it runs

$$\begin{aligned} & \text{ABE.Decrypt}(\text{mpk}^{(i)}, \text{sk}_\top^{(i)}, \mathbf{C}^{(i)}) \rightarrow \mathbf{M}^{(i)} \quad \text{for } i = 1, \dots, i^* - 1 \\ \text{and} \quad & \text{ABE.Decrypt}(\text{mpk}^{(i)}, \text{sk}_\mathbf{x}^{(i)}, \mathbf{C}^{(i)}) \rightarrow \mathbf{M}^{(i)} \quad \text{for } i = i^*, \dots, \eta. \end{aligned}$$

Finally, it outputs $\mathbf{M} = \bigoplus_{i=1}^\eta \mathbf{M}^{(i)}$.

Correctness. Let us assume $R(\mathbf{x}, \mathbf{y}) = 1$. Since $R(\top, \mathbf{y}) = 1$ for any \mathbf{y} and $n^i \geq |\mathbf{y}| = 1$ for all $i \in \mathbb{N}$, $\text{ABE.Decrypt}(\text{mpk}^{(i)}, \text{sk}_\top^{(i)}, \mathbf{C}^{(i)})$ can recover $\mathbf{M}^{(i)}$ for $i = 1, \dots, i^* - 1$. Furthermore, since $n^i > |\mathbf{x}|$ for all $i \geq i^*$, $\text{ABE.Decrypt}(\text{mpk}^{(i)}, \text{sk}_\mathbf{x}^{(i)}, \mathbf{C}^{(i)})$ can recover $\mathbf{M}^{(i)}$ also for $i = i^*, \dots, \eta$. Therefore, the decryption algorithm can correctly recover the message \mathbf{M} .

Theorem 5. Assuming Π is selective secure for all polynomial $L = L(n)$, so is Π' .

Proof. Assume an adversary \mathcal{A} that breaks Π' with non-negligible probability. Let ℓ_{\max} be the maximum size of the key attributes for which \mathcal{A} makes key extraction queries. Since \mathcal{A} is a PPT algorithm, ℓ_{\max} is bounded by some polynomial in n . In particular, there exists a constant $i^* \in \mathbb{N}$ such that

$$\ell_{\max} < n^{i^*} \quad \text{and} \quad i^* \leq \eta(n)$$

hold for all sufficiently large n . We will construct an adversary \mathcal{B} against Π with $L = n^{i^*}$ from \mathcal{A} in the following.

Setup. First, given 1^n and 1^ℓ , \mathcal{B} passes the same values to \mathcal{A} . Then, \mathcal{A} outputs y^* satisfying $|y^*| = \ell$. \mathcal{B} outputs the same y^* as its target. Then, $\text{ABE.Setup}(1^n, 1^\ell, n^{i^*}) \rightarrow (\text{mpk}^{(i^*)}, \text{msk}^{(i^*)})$ is run and $\text{mpk}^{(i^*)}$ is given to \mathcal{B} . \mathcal{B} then runs $\text{ABE.Setup}(1^n, 1^\ell, n^i) \rightarrow (\text{mpk}^{(i)}, \text{msk}^{(i)})$ for $i = [1, \eta] \setminus \{i^*\}$ and sets $\text{mpk} = (\text{mpk}^{(1)}, \text{mpk}^{(2)}, \dots, \text{mpk}^{(\eta)})$. \mathcal{B} keeps $\text{msk}^{(i)}$ for $i \in [1, \eta] \setminus \{i^*\}$ secret, and returns mpk to \mathcal{A} .

Phase 1 and 2. When \mathcal{A} makes a key extraction query for x , \mathcal{B} proceeds as follows. \mathcal{B} first finds i^* such that $n^{i^*-1} \leq |x| < n^{i^*}$. Since we have $|x| \leq \ell_{\max}$ and $\ell_{\max} < n^{i^*}$, it holds that $i^* \leq i^*$. \mathcal{B} then queries a private key for x to its challenger. Note that since $|x| \leq \ell_{\max} < n^{i^*}$, it is a valid query. Then, $\text{ABE.KeyGen}(\text{mpk}^{(i^*)}, \text{msk}^{(i^*)}, x) \rightarrow \text{sk}_x^{(i^*)}$ is run and $\text{sk}_x^{(i^*)}$ is given to \mathcal{B} . \mathcal{B} also runs $\text{ABE.KeyGen}(\text{mpk}^{(i)}, \text{msk}^{(i)}, \top) \rightarrow \text{sk}_\top^{(i)}$ for $i = 1, \dots, i^* - 1$. Finally, \mathcal{B} runs $\text{ABE.KeyGen}(\text{mpk}^{(i)}, \text{msk}^{(i)}, x) \rightarrow \text{sk}_x^{(i)}$ for $i \in [i^*, \eta] \setminus \{i^*\}$ and returns

$$\text{sk}_x = (\text{sk}_\top^{(1)}, \text{sk}_\top^{(2)}, \dots, \text{sk}_\top^{(i^*-1)}, \text{sk}_x^{(i^*)}, \text{sk}_x^{(i^*+1)}, \dots, \text{sk}_{\text{ID}}^{(\eta)}).$$

to \mathcal{A} . Note that the simulation is possible since $i^* \leq i^*$. Otherwise, if $i^* > i^*$, \mathcal{B} has to give $\text{sk}_\top^{(i^*)}$ to \mathcal{A} and thus the simulation is impossible. (\mathcal{B} is not allowed to make a key extraction query for \top , since it allows \mathcal{B} to decrypt the challenge ciphertext and trivially win the game).

Challenge. In the challenge phase, \mathcal{A} outputs two messages (M_0, M_1) . Then, \mathcal{B} first picks random $M^{(i)} \xleftarrow{\$} \{0, 1\}^{\ell_M}$ for $i \in [1, \eta] \setminus \{i^*\}$ and sets

$$M_b^{(i^*)} = M_b \oplus \left(\bigoplus_{i \in [1, \eta] \setminus \{i^*\}} M^{(i)} \right) \quad \text{for} \quad b \in \{0, 1\}.$$

Then, \mathcal{B} runs $\text{ABE.Encrypt}(\text{mpk}^{(i)}, y^*, M^{(i)}) \rightarrow C^{(i)}$ for $i \in [1, \eta] \setminus \{i^*\}$ and makes the challenge query for messages $(M_0^{(i^*)}, M_1^{(i^*)})$ to its challenger. The challenger picks a coin $\text{coin} \xleftarrow{\$} \{0, 1\}$, runs $\text{ABE.Encrypt}(\text{mpk}^{(i^*)}, y^*, M_{\text{coin}}^{(i^*)}) \rightarrow C^*$, and gives C^* to \mathcal{B} . Finally, \mathcal{B} gives the challenge ciphertext

$$(C^{(1)}, \dots, C^{(i^*-1)}, C^*, C^{(i^*+1)}, \dots, C^{(\eta)})$$

to \mathcal{A} .

Guess. Finally, \mathcal{A} outputs a guess $\widehat{\text{coin}}$ for coin . \mathcal{B} outputs the same $\widehat{\text{coin}}$ as its guess.

Analysis. We show that \mathcal{B} is a valid adversary for Π . First, since \mathcal{A} does not make a key extraction query for x such that $R(x, y^*) = 1$, \mathcal{B} does not make such a query neither. It can also be seen that the challenge ciphertext is distributed correctly. Since the simulation by \mathcal{B} is perfect, \mathcal{B} has the same advantage as \mathcal{A} , as desired. \square

More Efficient Conversion. Similarly to our conversion in Section 4.2, the number of instances of the sub-scheme Π can be reduced to $O(\log \eta)$ from $O(\eta)$. The idea is to run $\text{ABE.Setup}(1^n, 1^\ell, n^i)$ for $i = 1, 2, 4, \dots, 2^i, \dots, 2^{\lceil \log \eta \rceil}$. We note that this leads to a security reduction with looser approximation factors.

B.3 Implication and Comparison

By applying our transformation in Section B.2 to the ABE scheme in [28], we obtain a new ABE scheme. The security of the scheme is proven from that of the scheme in [28] for polynomial size L , which can be reduced to $\text{dLWE}_{q,m,q,\chi}$ with polynomial q/χ_{\max} . The latter can be reduced to the worst case lattice problems with polynomial approximation factor [40]. Furthermore, the size of the private key is $\eta(n) \cdot \text{poly}(n, \log(n^n)) = \text{poly}(n)$, which is independent of L . Therefore, our scheme

- is proven secure from the worst case lattice problems with polynomial approximation factors,
- achieves compact private keys, which means that the size of private keys is independent of the length of the branching programs, and
- can deal with unbounded length branching programs for key attribute.

Our scheme is the first scheme that satisfies all of the above properties simultaneously. The scheme in [27] can deal with unbounded length branching programs, but the private keys are not compact (linear in the length of branching programs). We can obtain ABE for branching programs with short private keys from the ABE for general circuit in [11]. However, the security of the scheme depends on $\text{dLWE}_{q,m,n,\chi}$ with $q/\chi_{\max} = n^{\log(n)}$. In the original scheme in [28], either the length of branching programs are bounded or the security is based on the super-polynomial LWE assumption.

Contents

1	Introduction	1
2	Overview of Our Technique	4
2.1	Overview of the Construction	4
2.2	Overview of the Security Proof	4
2.3	An Additional Idea	6
3	Preliminaries	7
3.1	Identity-Based Encryption	7
3.2	Lattice Preliminaries	8
3.3	Basic Facts	10
4	Parametrized IBE	11
4.1	Definition of Parametrized IBE	11
4.2	IBE from PIBE	12
5	Our Construction of PIBE from Lattices	14
5.1	Homomorphic Computation	14
5.2	Our Construction	14
5.3	Correctness and Parameter Selection	15
5.4	Security Proof	16
5.5	Multi-bit Encryption	24
6	Comparisons and Discussions	24
	References	26
A	Omitted Proofs	28
A.1	The Proof of Lemma 5	28
A.2	The Proof of Lemma 6	28
B	Application of Our Technique to Attribute-Based Encryption	28
B.1	Attribute-Based Encryption	29
B.2	Our Transformation	31
B.3	Implication and Comparison	33