

Tighter Security Proofs for GPV-IBE in the Quantum Random Oracle Model

Shuichi Katsumata^{*†}, Shota Yamada[†], and Takashi Yamakawa[‡]

^{*}The University of Tokyo

`shuichi_katsumata@it.k.u-tokyo.ac.jp`

[†]National Institute of Advanced Industrial Science and Technology (AIST)

`yamada-shota@aist.go.jp`

[‡]NTT Secure Platform Laboratories

`yamakawa.takashi@lab.ntt.co.jp`

November 20, 2018

Abstract

In (STOC, 2008), Gentry, Peikert, and Vaikuntanathan proposed the first identity-based encryption (GPV-IBE) scheme based on a post-quantum assumption, namely, the learning with errors (LWE) assumption. Since their proof was only made in the random oracle model (ROM) instead of the *quantum* random oracle model (QROM), it remained unclear whether the scheme was truly post-quantum or not. In (CRYPTO, 2012), Zhandry developed new techniques to be used in the QROM and proved security of GPV-IBE in the QROM, hence answering in the affirmative that GPV-IBE is indeed post-quantum. However, since the general technique developed by Zhandry incurred a large reduction loss, there was a wide gap between the concrete efficiency and security level provided by GPV-IBE in the ROM and QROM. Furthermore, regardless of being in the ROM or QROM, GPV-IBE is not known to have a tight reduction in the multi-challenge setting. Considering that in the real-world an adversary can obtain many ciphertexts, it is desirable to have a security proof that does not degrade with the number of challenge ciphertext.

In this paper, we provide a much tighter proof for the GPV-IBE in the QROM in the single-challenge setting. In addition, we also show that a slight variant of the GPV-IBE has an almost tight reduction in the multi-challenge setting both in the ROM and QROM, where the reduction loss is independent of the number of challenge ciphertext. Our proof departs from the traditional partitioning technique and resembles the approach used in the public key encryption scheme of Cramer and Shoup (CRYPTO, 1998). Our proof strategy allows the reduction algorithm to program the random oracle the same way for all identities and naturally fits the QROM setting where an adversary may query a superposition of all identities in one random oracle query. Notably, our proofs are much simpler than the one by Zhandry and conceptually much easier to follow for cryptographers not familiar with quantum computation. Although at a high level, the techniques used for the single and multi-challenge setting are similar, the technical details are quite different. For the multi-challenge setting, we rely on the Katz-Wang technique (CCS, 2003) to overcome some obstacles regarding the leftover hash lemma.

Keywords. Identity-based encryption, quantum random oracle models, LWE assumption, tight security reduction, multi-challenge security.

Contents

1	Introduction	3
1.1	Background	3
1.2	Our Contribution	4
1.3	Technical Overview	4
1.4	Discussion.	8
1.5	Related Work	9
2	Preliminaries	10
2.1	Quantum Computation	10
2.2	Pseudorandom Function.	11
2.3	Identity-Based Encryption	12
2.4	Background on Lattices	13
3	Tightly Secure Single Challenge GPV-IBE	16
3.1	Construction	16
3.2	Correctness and Parameter Selection	17
3.3	Security Proof in ROM	18
3.4	Security Proof in QROM	21
4	(Almost) Tightly Secure Multi-Challenge IBE	24
4.1	Randomness Extraction	24
4.2	Construction	26
4.3	Correctness and Parameter Selection	27
4.4	Security Proof in ROM	28
4.5	Security Proof in QROM	31

1 Introduction

1.1 Background

Shor [Sho94] in his breakthrough result showed that if a quantum computer is realized, then almost all cryptosystems used in the real world will be broken. Since then, a significant amount of studies have been done in the area of post-quantum cryptography, whose motivation is constructing cryptosystems secure against quantum adversaries. Recently in 2016, the National Institute of Standards and Technology (NIST) initiated the Post-Quantum Cryptography Standardization, and since then post-quantum cryptography has been gathering increasingly more attention.

Random Oracles in Quantum World. In general, security proofs of practical cryptographic schemes are given in the random oracle model (ROM) [BR93], which is an idealized model where a hash function is modeled as a publicly accessible oracle that computes a random function. Boneh et al. [BDF⁺11] pointed out that the ROM as in the classical setting is not reasonable when considering security against quantum adversaries, since quantum adversaries may compute hash functions over quantum superpositions of many inputs. Considering this fact, as a reasonable model against quantum adversaries, they proposed a new model called the quantum random oracle model (QROM), where a hash function is modeled as a *quantumly accessible* random oracle. As discussed in [BDF⁺11], many commonly-used proof techniques in the ROM do not work in the QROM. Therefore even if we have a security proof in the ROM, we often require new techniques to obtain similar results in the QROM.

Identity-based Encryption in QROM. Identity-Based Encryption (IBE) is a generalization of a public key encryption scheme where the public key of a user can be any arbitrary string such as an e-mail address. The first IBE scheme based on a post-quantum assumption is the one proposed by Gentry, Peikert and Vaikuntanathan (GPV-IBE) [GPV08], which is based on the learning with errors (LWE) assumption [Reg05]. To this date, GPV-IBE is still arguably the most efficient IBE scheme that is based on a hardness assumption that resists quantum attacks. However, since their original security proof was made in the ROM instead of the QROM, it was unclear if we could say the scheme is truly post-quantum. Zhandry [Zha12b] answered this in the affirmative by proving that the GPV-IBE is indeed secure in the QROM under the LWE assumption, hence truly post-quantum, by developing new techniques in the QROM.

Tight Security of GPV-IBE. However, if we consider the tightness of the reduction, the security proof of the GPV-IBE by Zhandry [Zha12b] does not provide satisfactory security. Specifically, GPV-IBE may be efficient in the ROM, but it is no longer efficient in the QROM. In general, a cryptographic scheme is said to be tightly secure under some assumption if breaking the security of the scheme is as hard as solving the assumption. More precisely, suppose that we proved that if there exists an adversary breaking the security of the scheme with advantage ϵ and running time T , we can break the underlying assumption with advantage ϵ' and running time T' . We say that the scheme is tightly-secure if we have $\epsilon'/T' \approx \epsilon/T$. By using this notation, Zhandry gave a reduction from the security of GPV-IBE to the LWE assumption with $\epsilon' \approx \epsilon^2/(Q_H + Q_{ID})^4$ and $T' \approx T + (Q_H + Q_{ID})^2 \cdot \text{poly}(\lambda)$ where Q_H denotes the number of hash queries, Q_{ID} denotes the number of secret key queries, λ denotes the security parameter, and poly denotes some fixed polynomial. Though the reduction is theoretically interesting, the meaning of the resulting security bound in a realistic setting is unclear. For example, if we want to obtain 128-bit security for the resulting IBE, and say we had $\epsilon = 2^{-128}$, $Q_H = 2^{100}$, $Q_{ID} = 2^{20}$, then even if we ignore the blowup for the running time, we would have to start from at least a 656-bit secure LWE assumption, which incurs a significant blowup of the parameters. Indeed, Zhandry left it as an open problem to give

a tighter reduction for the GPV-IBE.

Multi-Challenge Tightness. The standard security notion of IBE considers the setting where an adversary obtains only one challenge ciphertext. This is because security against adversaries obtaining many challenge ciphertexts can be reduced to the security in the above simplified setting. However, as pointed out by Hofheinz and Jager [HJ12], tightness is not preserved in the above reduction since the security degrades by the number of ciphertexts. Therefore tightly secure IBE in the single-challenge setting does not imply tightly secure IBE in the multi-challenge setting. On the other hand, in the real world, it is natural to assume that an adversary obtains many ciphertexts, and thus tight security in the multi-challenge setting is desirable. However, there is no known security proof for the GPV-IBE or its variant that does not degrade with the number of challenge ciphertexts even in the classical setting.

1.2 Our Contribution

We provide much tighter security proofs for the GPV-IBE in the QROM in the single-challenge setting. Furthermore, we provide a multi-challenge tight variant of GPV-IBE that is secure both in the ROM and QROM. In the following, we describe the tightness of our security proofs by using the same notation as in the previous section.

- In the single-challenge setting, we give a reduction from the security of GPV-IBE to the LWE assumption with $\epsilon' \approx \epsilon$ and $T' = T + (Q_H + Q_{ID})^2 \cdot \text{poly}(\lambda)$. If we additionally assume quantumly secure pseudorandom functions (PRFs), then we further obtain a tighter reduction, which gives $\epsilon' \approx \epsilon$ and $T' = T + (Q_H + Q_{ID}) \cdot \text{poly}(\lambda)$. This is the first security proof for GPV-IBE whose security bound does not degrade with Q_H or Q_{ID} even in the classical setting. We note that the same security bound can be achieved without assuming PRFs in the classical ROM.
- We give a slight variant of GPV-IBE scheme whose multi-challenge security is reduced to the LWE assumption with $\epsilon' = \epsilon/\text{poly}(\lambda)$ and $T' \approx T + (Q_H + Q_{ID} + Q_{ch})^2 \cdot \text{poly}(\lambda)$ where Q_{ch} denotes the number of challenge queries. If we additionally assume quantumly secure PRFs, then we further obtain a tighter reduction. Namely, ϵ' is the same as the above, and $T' = T + (Q_H + Q_{ID} + Q_{ch}) \cdot \text{poly}(\lambda)$. This is the first variant of the GPV-IBE scheme whose security bound does not degrade with Q_{ch} even in the classical setting. We note that the same security bound can be achieved without assuming PRFs in the classical ROM.

Moreover, our security proofs are much simpler than the one by Zhandry [Zha12b]. In his work, he introduced new techniques regarding indistinguishability of oracles against quantum adversaries. Though his techniques are general and also useful in other settings (e.g., [Zha12a]), it involves some arguments on quantum computation, and they are hard to follow for cryptographers who are not familiar with quantum computation. On the other hand, our proofs involve a minimal amount of discussions about quantum computation, and our proofs are done almost similar to the counterparts in the classical ROM.

1.3 Technical Overview

GPV-IBE. First, we briefly describe the GPV-IBE [GPV08], which is the main target of this paper. A master public key is a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and a master secret key is its trapdoor $\mathbf{T}_A \in \mathbb{Z}^{m \times m}$, which enables one to compute a short vector $\mathbf{e} \in \mathbb{Z}_q^m$ such that $\mathbf{A}\mathbf{e} = \mathbf{u}$ given an arbitrary vector $\mathbf{u} \in \mathbb{Z}_q^n$. A private key sk_{ID} for an identity $ID \in \mathcal{ID}$ is a short vector $\mathbf{e} \in \mathbb{Z}_q^m$

such that $\mathbf{A}\mathbf{e} = \mathbf{u}_{\text{ID}}$ where $\mathbf{u}_{\text{ID}} = \text{H}(\text{ID})$ for a hash function $\text{H} : \mathcal{ID} \rightarrow \mathbb{Z}_q^n$, which is modeled as a random oracle. A ciphertext for a message $M \in \{0, 1\}$ consists of $c_0 = \mathbf{u}_{\text{ID}}^\top \mathbf{s} + x + M \lfloor q/2 \rfloor$ and $\mathbf{c}_1 = \mathbf{A}^\top \mathbf{s} + \mathbf{x}$. Here \mathbf{s} is a uniformly random vector over \mathbb{Z}_q^n and x, \mathbf{x} are small “noise” terms where each entries are sampled from some specific Gaussian distribution χ . Decryption can be done by computing $w = c_0 - \mathbf{c}_1^\top \mathbf{e}_{\text{ID}} \in \mathbb{Z}_q$ and deciding if w is closer to 0 or to $\lfloor q/2 \rfloor$ modulo q .

Security Proof in Classical ROM. The above IBE relies its security on the LWE assumption, which informally states the following: given a uniformly random matrix $[\mathbf{A}|\mathbf{u}] \leftarrow \mathbb{Z}_q^{n \times (m+1)}$ and some vector $\mathbf{b} \in \mathbb{Z}_q^{m+1}$, there is no PPT algorithm that can decide with non-negligible probability whether \mathbf{b} is of the form $[\mathbf{A}|\mathbf{u}]^\top \mathbf{s} + \mathbf{x}'$ for some $\mathbf{s} \leftarrow \mathbb{Z}_q^n$ and $\mathbf{x}' \leftarrow \chi^{m+1}$, or a uniformly random vector over \mathbb{Z}_q^{m+1} , i.e., $\mathbf{b} \leftarrow \mathbb{Z}_q^{m+1}$. Below, we briefly recall the original security proof in the classical ROM given by Gentry et al. [GPV08] and see how the random oracle is used by the reduction algorithm. The proof relies on a key lemma which states that we can set $\text{H}(\text{ID})$ and \mathbf{e} in the “reverse order” from the real scheme. That is, we can first sample \mathbf{e} from some distribution and program $\text{H}(\text{ID}) := \mathbf{A}\mathbf{e}$ so that their distributions are close to uniformly random as in the real scheme. In the security proof, a reduction algorithm guesses $i \in [Q]$ such that the adversary’s i -th hash query is the challenge identity ID^* where Q denotes the number of hash queries made by the adversary. Then for all but the i -th hash query, the reduction algorithm programs $\text{H}(\text{ID})$ in the above manner, and for the i -th query, it programs the output of $\text{H}(\text{ID}^*)$ to be the vector \mathbf{u} contained in the LWE instance that is given as the challenge. Specifically, the reduction algorithm sets the challenge user’s identity vector \mathbf{u}_{ID^*} as the random vector \mathbf{u} contained in the LWE instance. If the guess is correct, then it can embed the LWE instance into the challenge ciphertexts c_0^* and \mathbf{c}_1^* ; in case it is a valid LWE instance, then (c_0^*, \mathbf{c}_1^*) is properly set to $(\mathbf{u}_{\text{ID}^*}^\top \mathbf{s} + x + M \lfloor q/2 \rfloor, \mathbf{A}^\top \mathbf{s} + \mathbf{x})$ as in the real scheme. Therefore, the challenge ciphertext can be switched to random due to the LWE assumption. After this switch, M is perfectly hidden and thus the security of GPV-IBE is reduced to the LWE assumption. Since the reduction algorithm programs the random oracle in the same way except for the challenge identity, this type of proof methodology is often times referred to as the “all-but-one programming”.

Security Proof in QROM in [Zha12b]. Unfortunately, the above proof cannot be simply extended to a proof in the QROM. The reason is that in the QROM, even a single hash query can be a superposition of *all* the identities. In such a case, to proceed with the above all-but-one programming approach, the reduction algorithm would have to guess a single identity out of all the possible identities which he hopes that would be used as the challenge identity ID^* by the adversary. Obviously, the probability of the reduction algorithm being right is negligible, since the number of possible identities is exponentially large. This is in sharp contrast with the ROM setting, where the reduction algorithm was allowed to guess the single identity out of the polynomially many (classical) random oracle queries made by the adversary. Therefore, the all-but-one programming as in the classical case cannot be used in the quantum case. To overcome this barrier, Zhandry [Zha12b] introduced a useful lemma regarding what he calls the semi-constant distribution. The semi-constant distribution with parameter $0 < p < 1$ is a distribution over functions from \mathcal{X} to \mathcal{Y} such that a function chosen according to the distribution gives the same fixed value for random p -fraction of all inputs, and behaves as a random function for the rest of the inputs. He proved that a function according to the semi-constant distribution with parameter p and a random function cannot be distinguished by an adversary that makes Q oracle queries with advantage greater than $\frac{8}{3}Q^4p^2$. In the security proof, the reduction algorithm partitions the set of identities into controlled and uncontrolled sets. The uncontrolled set consists of randomly

chosen p -fraction of all identities, and the controlled set is the complement of it. The reduction algorithm embeds an LWE instance into the uncontrolled set, and programs the hash values for the controlled set so that the decryption keys for identities in the controlled set can be extracted efficiently. Then the reduction algorithm works as long as the challenge identity falls inside the uncontrolled set and all identities for secret key queries fall inside the controlled set (otherwise it aborts). By appropriately setting p , we can argue that the probability that the reduction algorithm does not abort is non-negligible, and thus the security proof is completed. Though this technique is very general and useful, a huge reduction loss is inherent as long as we take the above strategy because the reduction algorithm has to abort with high probability. It may be useful to point out for readers who are familiar with IBE schemes in the standard model that the above technique is conceptually very similar to the partitioning technique which is often used in the context of adaptively secure IBE scheme in the standard model [Wat05, ABB10, CHKP10]. The reason why we cannot make the proof tight is exactly the same as that for the counterparts in the standard model.

Our Tight Security Proof in QROM. As discussed above, we cannot obtain a tight reduction as long as we use a partitioning-like technique. Therefore we take a completely different approach, which is rather similar to that used in the public key encryption scheme of Cramer and Shoup [CS98], which has also been applied to the pairing-based IBE construction of Gentry [Gen06]. The idea is that we simulate in a way so that we can create exactly one valid secret key for every identity. Note that this is opposed to the partitioning technique (and the all-but-one programming technique) where the simulator cannot create a secret key for an identity in the uncontrolled set. To create the challenge ciphertext, we use the one secret key we know for that challenge identity. If the adversary can not tell which secret key the ciphertext was created from and if there are potentially many candidates for the secret key, we can take advantage of the entropy of the secret key to statistically hide the message.

In more detail, the main observation is that the secret key \mathbf{e} , i.e. a short vector \mathbf{e} such that $\mathbf{A}\mathbf{e} = \mathbf{u}$, retains plenty of entropy even after fixing the public values \mathbf{A} and \mathbf{u} . Therefore, by programming the hash value \mathbf{u} of an identity, we can easily create a situation where the simulator knows exactly one secret key out of the many possible candidates. Furthermore, the simulator knowing a secret key \mathbf{e}_{ID^*} such that $\mathbf{A}\mathbf{e}_{\text{ID}^*} = \mathbf{u}_{\text{ID}^*}$, can simulate the challenge ciphertext by creating $c_0^* = \mathbf{e}_{\text{ID}^*}^\top \mathbf{c}_1^* + M_{\lfloor q/2 \rfloor}$ and $\mathbf{c}_1^* = \mathbf{A}^\top \mathbf{s} + \mathbf{x}$. Here, the key observation is that we no longer require the LWE instance $(\mathbf{u}_{\text{ID}^*}, \mathbf{u}_{\text{ID}^*}^\top \mathbf{s} + x)$ to simulate the challenge ciphertext. Though the distribution of c_0^* simulated as above is slightly different from that of the real ciphertext due to the difference in the noise distributions, we ignore it in this overview. In the real proof, we overcome this problem by using the noise rerandomization technique by Katsumata and Yamada [KY16]. Then we use the LWE assumption to switch \mathbf{c}_1^* to random. Finally, we argue that $\mathbf{e}_{\text{ID}^*}^\top \mathbf{c}_1^*$ is almost uniform if the min-entropy of \mathbf{e}_{ID^*} is high and \mathbf{c}_1^* is uniformly random due to the leftover hash lemma. Therefore, all information of the message M is hidden and thus the proof is completed.

Finally, we observe that the above proof naturally fits in the QROM setting. The crucial difference from the partitioning technique is that in our security proof we program the random oracle in the same way for all identities. Therefore even if an adversary queries a superposition of all identities, the simulator can simply quantumly perform the programming procedure for the superposition. Thus the proof in the classical ROM can be almost automatically converted into the one in the QROM.

Tight Security in Multi-Challenge Setting. Unfortunately, the above idea does not extend naturally to the tightly-secure multi-challenge setting. One can always prove security in the multi-

challenge setting starting from a scheme that is single-challenge secure via a hybrid argument, however, as mentioned by Hofheinz and Jager [HJ12], this type of reduction does not preserve tightness. A careful reader may think that the above programming technique can be extended to the multi-challenge setting, hence bypassing the hybrid argument. We briefly explain why this is not the case. Informally, in the above proof, the reduction algorithm embeds its given LWE instance $(\mathbf{A}, \mathbf{A}^\top \mathbf{s} + \mathbf{x})$ into the challenge ciphertext by creating $(c_0^* = \mathbf{e}_{\text{ID}^*}^\top \mathbf{c}_1^* + M \lfloor q/2 \rfloor, \mathbf{c}_1^* = \mathbf{A}^\top \mathbf{s} + \mathbf{x})$, where \mathbf{e}_{ID^*} is the secret key of the challenge user \mathbf{u}_{ID^*} . Therefore, since the \mathbf{c}_1^* component of every ciphertext is an LWE instance for the same public matrix \mathbf{A} , to simulate multiple challenge ciphertexts in the above manner, the reduction algorithm must be able to prepare a special type of LWE instance $(\mathbf{A}, \{\mathbf{A}^\top \mathbf{s}^{(k)} + \mathbf{x}^{(k)}\}_{k \in [N]})$, where $N = \text{poly}(\lambda)$ is the number of challenge ciphertext queried by the adversary. It can be easily seen that this construction is tightly-secure in the multi-challenge setting with the same efficiency as the single-challenge setting, *if* we assume that this special type of LWE problem is provided to the reduction algorithm as the challenge. However, unfortunately, we still end up losing a factor of N in the reduction when reducing the standard LWE problem to this special LWE problem. In particular, we only shifted the burden of having to go through the N hybrid arguments to the assumption rather than to the scheme. As one may have noticed, there is a way to bypass the problem of going through the N hybrid arguments by using conventional techniques (See [Reg05, Reg10]) of constructing an unlimited number of fresh LWE instances given a fixed number of LWE instances. However, this techniques requires the noise of the newly created LWE instances to grow proportionally to the number of created instances. In particular, to create the above special LWE instance from a standard LWE instance, we require the size of the noise $\mathbf{x}^{(k)}$ to grow polynomially with N , where recall that N can be an arbitrary polynomial. Hence, although we can show a tightly secure reduction in the multi-challenge setting, for the concrete parameters of the scheme to be independent of N , we need to assume the LWE assumption with super-polynomial modulus size to cope with the super-polynomial noise blow up. This is far more inefficient than in the single-challenge setting where we only had to require an LWE assumption with polynomial modulus size.

To overcome this problem, we use the “lossy mode” of the LWE problem. It is well known that the secret vector \mathbf{s} is uniquely defined given an LWE instance $(\mathbf{A}, \mathbf{A}^\top \mathbf{s} + \mathbf{x})$ for large enough samples. A series of works, e.g., [GKPV10, BKPW12, AKPW13, LSSS17] have observed that if we instead sample \mathbf{A} from a special distribution that is computationally indistinguishable from the uniform distribution, then $(\mathbf{A}, \mathbf{A}^\top \mathbf{s} + \mathbf{x})$ leaks almost no information of the secret \mathbf{s} , hence the term “lossy mode”. This idea can be leveraged to prove (almost) tight security of the above single-challenge construction, where the reduction loss is independent of the number of challenge ciphertext. A first attempt of using this idea is as follows: During the security proof of the GPV-IBE, we first change the public matrix \mathbf{A} to a lossy matrix $\tilde{\mathbf{A}}$ and generate the secret keys and program the random oracle in the same way as before. To create the challenge ciphertexts, the reduction algorithm honestly samples $\mathbf{s}^{(k)}, x^{(k)}, \mathbf{x}^{(k)}$ and sets $(c_0^* = \mathbf{u}_{\text{ID}^*}^\top \mathbf{s}^{(k)} + x^{(k)} + M \lfloor q/2 \rfloor, \mathbf{c}_1^* = \tilde{\mathbf{A}}^\top \mathbf{s}^{(k)} + \mathbf{x}^{(k)})$. Now, it may seem that owing to the lossy mode of LWE, we can rely on the entropy of the secret vector $\mathbf{s}^{(k)}$ to argue that c_0^* is distributed uniformly random via the leftover hash lemma. The main difference between the previous single-challenge setting is that we can rely on the entropy of the secret vector $\mathbf{s}^{(k)}$ rather than on the entropy of the secret key \mathbf{e}_{ID^*} . Since each challenge ciphertext is injected with fresh entropy and we can argue statistically that a single challenge ciphertext is not leaking any information on the message, the reduction loss will be independent of the number of challenge ciphertext query N .

Although the above argument may seem correct at first glance, it incurs a subtle but a fatal flaw, thus bringing us to our proposed construction. The problem of the above argument is how

we use the leftover hash lemma. To use the lemma correctly, the vector \mathbf{u}_{ID^*} viewed as a hash function is required to be universal. This is true in case \mathbf{u}_{ID^*} is set as $\mathbf{A}\mathbf{e}_{\text{ID}^*}$, where $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$ and \mathbf{e}_{ID^*} is sampled from some appropriate distribution. However, this is *not* true anymore once we change \mathbf{A} to a lossy matrix $\tilde{\mathbf{A}}$, since $\tilde{\mathbf{A}}$ now lives in an exponentially small subset of $\mathbb{Z}_q^{n \times m}$, and \mathbf{u}_{ID^*} is no longer universal. Specifically, we can no longer rely on the entropy of $\mathbf{s}^{(k)}$ to statistically hide the message. To overcome this problem, our final idea is to use the Katz-Wang [KW03] technique. Specifically, we slightly alter the encryption algorithm of GPV-IBE to output the following instead:

$$c_0 = \mathbf{u}_{\text{ID}||0}^\top \mathbf{s} + x_0 + \text{M}\lfloor q/2 \rfloor, \quad c_1 = \mathbf{u}_{\text{ID}||1}^\top \mathbf{s} + x_1 + \text{M}\lfloor q/2 \rfloor, \quad \text{and} \quad \mathbf{c}_2 = \mathbf{A}^\top \mathbf{s} + \mathbf{x},$$

where $\mathbf{u}_{\text{ID}||b} = H(\text{ID}||b)$ for $b \in \{0, 1\}$. During the security proof, the reduction algorithm sets $\mathbf{u}_{\text{ID}||0}$ and $\mathbf{u}_{\text{ID}||1}$ so that one of them is uniformly random over \mathbb{Z}_q^n and the other is constructed as $\mathbf{A}\mathbf{e}_{\text{ID}}$. Then, for the ciphertext c_b corresponding to the uniformly random vector $\mathbf{u}_{\text{ID}||b}$, we can correctly use the leftover hash lemma to argue that c_b statistically hides the message M . By going through one more hybrid argument, we can change both c_0, c_1 into random values that are independent of the message M . Note that instead of naively using the Katz-Wang technique, by reusing the \mathbf{c}_2 component, the above GPV-IBE variant only requires one additional element in \mathbb{Z}_q compared to the original GPV-IBE. Furthermore, in the actual construction, we do not require the noise terms x_0, x_1 in c_0, c_1 since we no longer rely on the LWE assumption to change c_0, c_1 into random values. Our construction and security reduction does not depend on the number of challenge ciphertext query N and in particular, can be proven under the LWE assumption with polynomial modulus size. In addition, due to the same reason as the single-challenge setting, our classical ROM proof can be naturally converted to a QROM proof.

1.4 Discussion.

Similar Techniques in Other Works. The idea to simulate GPV-IBE in a way so that we can create exactly one valid secret key for every secret key query is not new. We are aware of few works that are based on this idea. Gentry, Peikert and Vaikuntanathan [GPV08] mentioned that by using this technique, they can prove the security of the GPV-IBE in the standard model based on a non-standard interactive variant of the LWE (I-LWE) assumption which requires a hash function to define. Here since the hash function is given to the adversary, a quantum adversary may query quantum states to the hash functions on its own. Therefore, in addition with the fact that the I-LWE assumption is made in the standard model, the statement made by [GPV08] would hold in the QROM as well. However, they only gave a sketch of the proof, and did not give a formal proof. Alwen et al. [ADN⁺10] use the idea to construct an identity-based hash proof system (IB-HPS) based on the mechanism of GPV-IBE. We note that they assume the modulus q to be super-polynomial. Outside the context of identity-based primitives, Applebaum et al. [ACPS09] and Bourse et al. [BDPMW16] provide an analysis of rerandomizing LWE samples which can be seen as a refinement of the idea mentioned in [GPV08]. [ACPS09] constructs a KDM-secure cryptosystem based on the LWE problem and [BDPMW16] shows a simple method for constructing circuit private fully homomorphic encryption schemes (FHE) based on the lattice-based FHE scheme of Gentry et al. [GSW13]. Both of their analysis only requires the modulus q to be polynomial. In summary, though similar ideas have been used, all of the previous works are irrelevant to tight security or the security in the QROM.

On Parameter-Tightness of Our Schemes. In the above overview, we focused on the tightness of the security proof. Here, we provide some discussions on how the parameters compare to the

original GPV-IBE scheme [GPV08]. For the single challenge setting, our parameters are only a small factor worse than the GPV-IBE scheme. This is because the only difference is using the noise rerandomization technique of [KY16], which only slightly degrades the noise-level.¹ For the multi-challenge setting, the situation is more different. In this case, the parameters are much worse than the original (single-challenge secure) GPV-IBE scheme. This is because we have to go through the lossy-mode of LWE which requires for larger parameters. The concrete parameters are provided in Sec. 4.3.

Relation to CCA-Secure PKE. By applying the Canetti-Halevi-Katz transformation [CHK04] to our single-challenge-secure IBE scheme, we obtain a public key encryption (PKE) scheme secure against chosen ciphertext attacks (CCA) that is tightly secure in the single-challenge setting under the LWE assumption in the QROM. We note that Saito et al. [SXY18] already proposed such a PKE scheme in the single-challenge setting that is more efficient than the scheme obtained by the above transformation.

On Running Time of Reductions. In the above overview, we ignore the running time of reductions. Though it seems that the above described reductions run in nearly the same time as the adversaries, due to a subtle problem of simulating random oracles against quantum adversaries, there is a significant blowup by a square factor of the number of queries the adversaries make. In the classical ROM, when we simulate a random oracle in security proofs, we usually sample a random function in a lazy manner. That is, whenever an adversary queries a point that has not been queried before, a reduction algorithm samples a fresh randomness and assigns it as a hash value for that point. However, this cannot be done in the QROM because an adversary may query a superposition of all the inputs in a single query. Therefore a reduction algorithm has to somehow commit to the hash values of all inputs at the beginning of the simulation.

Zhandry [Zha12b] proved that an adversary that makes Q queries cannot distinguish a random function and a $2Q$ -wise independent hash function via quantum oracle accesses. Therefore we can use a $2Q$ -wise independent hash to simulate a random oracle. However, if we take this method, the simulator has to evaluate a $2Q$ -wise independent hash function for each hash query, and this is the reason why the running time blowups by $\Omega(Q^2)$.

One possible way to avoid this huge blowup is to simulate a random oracle by a PRF secure against quantum accessible adversaries. Since the time needed to evaluate a PRF is some fixed polynomial in the security parameter, the blowup for the running time can be made $Q \cdot \text{poly}(\lambda)$ which is significantly better than $\Omega(Q^2)$. However, in order to use this method, we have to additionally assume the existence of quantumly secure PRFs. Such PRFs can be constructed based on any quantumly-secure one-way function [Zha12a], and thus they exist if the LWE assumption holds against quantum adversaries. However, the reduction for such PRFs are non-tight and thus we cannot rely on them in the context of tight security. Our suggestion is to use a real hash function to implement PRFs and to assume that it is a quantumly secure PRF. We believe this to be a natural assumption if we are willing to idealize a hash function as a random oracle. (See also the discussion in Sec. 2.2.)

1.5 Related Work

Schemes in QROM. Boneh et al. [BDF⁺11] introduced the QROM, and gave security proofs for the GPV-signature [GPV08] and a hybrid variant of the Bellare-Rogaway encryption [BR93]

¹ Our parameter selection in the main body may seem much worse compared to GPV-IBE, but this is only because we choose the parameters conservatively. Specifically, we can set the parameters to be only slightly worse than GPV-IBE by setting them less conservatively as in [GPV08]. Please, see end of Sec. 3.2 for more details.

in the QROM. We note that their security proof for the GPV-signature is tight. Zhandry [Zha12b] proved that GPV-IBE and full-domain hash signatures are secure in the QROM. Targhi and Unruh [TU16] proposed variants of Fujisaki-Okamoto transformation and OAEP that are CCA secure in the QROM. Jiang et al. [JZC⁺18] gave tighter security proofs for some variants of Fujisaki-Okamoto transformation. Some researchers studied the security of the Fiat-Shamir transform in the QROM [ARU14, Unr15, Unr17]. Unruh [Unr14b] proposed a revocable quantum timed-release encryption scheme in the QROM. Unruh [Unr14a] proposed a position verification scheme in the QROM. Recently, some researchers studied tight securities in the QROM. Alkim et al. [ABB⁺17] proved that the signature scheme known as TESLA [BG14] is tightly secure under the LWE assumption. Saito et al. [SXY18] proposed a tightly CCA secure variant of the Bellare-Rogaway encryption. Kiltz et al. [KLS18] gave a tight reduction for the Fiat-Shamir transform in the QROM.

Tightly Secure IBEs. The first tightly secure IBE scheme from lattices in the single challenge setting and in the standard model was proposed by Boyen and Li [BL16]. Recently, they extended their results to the multi-challenge setting [BL18]. While both constructions are theoretically interesting and elegant, they are very inefficient and requires LWE assumption with super-polynomial approximation factors. As for the construction from bilinear maps, the first tightly secure IBE from standard assumptions in the single challenge setting and in the random oracle model was proposed by Katz and Wang [KW03]. Coron [Cor09] gave a tight reduction for a variant of the original Boneh-Franklin IBE [BF01]. Later, the first realization in the standard model was proposed by Chen and Wee [CW13]. In the subsequent works, it is further extended to the multi-challenge setting [HKS15, AHY15, GDCC16]. They are efficient but are not secure against quantum computers.

2 Preliminaries

Notations. For $n \in \mathbb{N}$, denote $[n]$ as the set $\{1, \dots, n\}$. For a finite set S , we let $U(S)$ denote the uniform distribution over S . For a distribution D and integer $k > 0$, define $(D)^k$ as the distribution $\prod_{i \in [k]} D$. For a distribution or random variable X we write $x \leftarrow X$ to denote the operation of sampling a random x according to X . For a set S , we write $s \leftarrow S$ as a shorthand for $s \leftarrow U(S)$. Let X and Y be two random variables over some finite set S_X, S_Y , respectively. The *statistical distance* $\Delta(X, Y)$ between X and Y is defined as $\Delta(X, Y) = \frac{1}{2} \sum_{s \in S_X \cup S_Y} |\Pr[X = s] - \Pr[Y = s]|$. The *min-entropy* of a random variable X is defined as $\mathbf{H}_\infty(X) = -\log(\max_x \Pr[X = x])$, where the base of the logarithm is taken to be 2 throughout the paper. For a bit $b \in \{0, 1\}$, \bar{b} denotes $1 - b$. For sets \mathcal{X} and \mathcal{Y} , $\text{Func}(\mathcal{X}, \mathcal{Y})$ denotes the set of all functions from \mathcal{X} to \mathcal{Y} . For a vector $\mathbf{v} \in \mathbb{R}^n$, denote $\|\mathbf{v}\|$ as the standard Euclidean norm. For a matrix $\mathbf{R} \in \mathbb{R}^{n \times n}$, denote $\|\mathbf{R}\|$ as the length of the longest column and $\|\mathbf{R}\|_{\text{GS}}$ as the longest column of the Gram-Schmidt orthogonalization of \mathbf{R} .

2.1 Quantum Computation

We briefly give some backgrounds on quantum computation. We refer to [NC00] for more details. A state $|\psi\rangle$ of n qubits is expressed as $\sum_{x \in \{0,1\}^n} \alpha_x |x\rangle \in \mathbb{C}^{2^n}$ where $\{\alpha_x\}_{x \in \{0,1\}^n}$ is a set of complex numbers such that $\sum_{x \in \{0,1\}^n} |\alpha_x|^2 = 1$ and $\{|x\rangle\}_{x \in \{0,1\}^n}$ is an orthonormal basis on \mathbb{C}^{2^n} (which is called a computational basis). If we measure $|\psi\rangle$ in the computational basis, then the outcome is a classical bit string $x \in \{0,1\}^n$ with probability $|\alpha_x|^2$, and the state becomes $|x\rangle$. An evolution of quantum state can be described by a unitary matrix U , which transforms $|x\rangle$ to

$U|x\rangle$. A quantum algorithm is composed of quantum evolutions described by unitary matrices and measurements. We also consider a quantum oracle algorithm, which can quantumly access to certain oracles. The running time $\text{Time}(\mathcal{A})$ of a quantum algorithm \mathcal{A} is defined to be the number of universal gates (e.g., Hadamard, phase, CNOT, and $\pi/8$ gates) and measurements required for running \mathcal{A} . (An oracle query is counted as a unit time if \mathcal{A} is an oracle algorithm.) Any efficient classical computation can be realized by a quantum computation efficiently. That is, for any function f that is classically computable, there exists a unitary matrix U_f such that $U_f|x, y\rangle = |x, f(x) \oplus y\rangle$, and the number of universal gates to express U_f is linear in the size of a classical circuit that computes f .

Quantum random oracle model. Boneh et al. [BDF⁺11] introduced the quantum random oracle model (QROM), which is an extension of the usual random oracle model to the quantum setting. Roughly speaking, the QROM is an idealized model where a hash function is idealized to be a quantumly accessible oracle that simulates a random function. More precisely, in security proofs in the QROM, a random function $H : \mathcal{X} \rightarrow \mathcal{Y}$ is uniformly chosen at the beginning of the experiment, and every entity involved in the system is allowed to access to an oracle that is given $\sum_{x,y} \alpha_{x,y} |x, y\rangle$ and returns $\sum_{x,y} \alpha_{x,y} |x, H(x) \oplus y\rangle$. We denote a quantum algorithm \mathcal{A} that accesses to the oracle defined as above by $\mathcal{A}^{(H)}$. In the QROM, one query to the random oracle is counted as one unit time. As in the classical case, we can implement two random oracles H_0 and H_1 from one random oracle H by defining $H_0(x) := H(0||x)$ and $H_1(x) := H(1||x)$. More generally, we can implement n random oracles from one random oracle by using the $\lfloor \log n \rfloor$ -bit prefix of an input as indices of random oracles.

As shown by Zhandry [Zha12b], a quantum random oracle can be simulated by a family of $2Q$ -wise independent hash functions against an adversary that quantumly accesses the oracle at most Q times. As a result, he obtained the following lemma.

Lemma 1. ([Zha12b, Theorem 6.1].) *Any quantum algorithm \mathcal{A} making quantum queries to random oracles can be efficiently simulated by a quantum algorithm \mathcal{B} , which has the same output distribution, but makes no queries. Especially, if \mathcal{A} makes at most Q queries to a random oracle $H : \{0, 1\}^a \rightarrow \{0, 1\}^b$, then $\text{Time}(\mathcal{B}) \approx \text{Time}(\mathcal{A}) + Q \cdot T_{a,b}^{2Q\text{-wise}}$ where $T_{a,b}^{2Q\text{-wise}}$ denotes the time to evaluate a $2Q$ -wise independent hash function from $\{0, 1\}^a$ to $\{0, 1\}^b$.*

The following lemma was shown by Boneh et al. [BDF⁺11]. Roughly speaking, this lemma states that if an oracle outputs independent and almost uniform value for all inputs, then it is indistinguishable from a random oracle even with quantum oracle accesses.

Lemma 2. ([BDF⁺11, Lem. 3].) *Let \mathcal{A} be a quantum algorithm that makes at most Q oracle queries, and \mathcal{X} and \mathcal{Y} be arbitrary sets. Let \mathcal{H} be a distribution over $\text{Func}(\mathcal{X}, \mathcal{Y})$ such that when we take $H \leftarrow \mathcal{H}$, for each $x \in \mathcal{X}$, $H(x)$ is identically and independently distributed according to a distribution D whose statistical distance is within ϵ from uniform. Then for any input z , we have*

$$\Delta(\mathcal{A}^{(\text{RF})}(z), \mathcal{A}^{(H)}(z)) \leq 4Q^2\sqrt{\epsilon}$$

where $\text{RF} \leftarrow \text{Func}(\mathcal{X}, \mathcal{Y})$ and $H \leftarrow \mathcal{H}$.

2.2 Pseudorandom Function.

We review the definition of quantum-accessible pseudorandom functions (PRFs) [BDF⁺11].

Definition 1 (Quantum-accessible PRF). *We say that a function $F : \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}$ is a quantum-accessible pseudorandom function if for all quantum polynomial time adversaries \mathcal{A} , its advantage defined below is negligible:*

$$\text{Adv}_{\mathcal{A},F}^{\text{PRF}}(\lambda) := \left| \Pr[\mathcal{A}^{|\text{RF}}(1^\lambda) = 1] - \Pr[\mathcal{A}^{F(K,\cdot)}(1^\lambda) = 1] \right|$$

where $\text{RF} \leftarrow \text{Func}(\mathcal{X}, \mathcal{Y})$ and $K \leftarrow \mathcal{K}$.

Zhandry [Zha12a] proved that some known constructions of classical PRFs including the tree-based construction [GGM86] and lattice-based construction [BPR12] are also quantum-accessible PRFs. However, these reductions are non-tight, and thus we cannot rely on these results when aiming for tight security. Fortunately, we can use the following lemma which states that we can use a quantum random oracle as a PRF similarly to the classical case.

Lemma 3. ([SXY18, Lem. 2.2]) *Let ℓ be an integer. Let $\text{H} : \{0, 1\}^\ell \times \mathcal{X} \rightarrow \mathcal{Y}$ and $\text{H}' : \mathcal{X} \rightarrow \mathcal{Y}$ be two independent random functions. If an unbounded time quantum adversary \mathcal{A} makes a query to H at most Q_{H} times, then we have*

$$\left| \Pr[\mathcal{A}^{|\text{H}, \text{H}'(K,\cdot)}(1^\lambda) = 1 \mid K \leftarrow \{0, 1\}^\ell] - \Pr[\mathcal{A}^{|\text{H}, \text{H}'}(1^\lambda) = 1] \right| \leq Q_{\text{H}} \cdot 2^{-\frac{\ell+1}{2}}.$$

2.3 Identity-Based Encryption

Syntax. We use the standard syntax of IBE [BF01]. Let \mathcal{ID} be the ID space of the scheme. If a collision resistant hash function $\text{CRH} : \{0, 1\}^* \rightarrow \mathcal{ID}$ is available, one can use an arbitrary string as an identity. An IBE scheme is defined by the following four algorithms.

Setup(1^λ) \rightarrow (**mpk**, **msk**): The setup algorithm takes as input a security parameter 1^λ 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 **ID** $\in \mathcal{ID}$. It outputs a private key **sk_{ID}**. We assume that **ID** is implicitly included in **sk_{ID}**.

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

Decrypt(**mpk**, **sk_{ID}**, **C**) \rightarrow **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 λ , all **ID** $\in \mathcal{ID}$, and all **M** in the specified message space,

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

holds, where the probability is taken over the randomness used in $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda)$, $\text{sk}_{\text{ID}} \leftarrow \text{KeyGen}(\text{mpk}, \text{msk}, \text{ID})$, 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} . Let $\text{CTSAm}(\cdot)$ be a sampling algorithm that takes as input a master public key of the scheme and outputs an element in the ciphertext space.

- **Setup.** At the outset of the game, the challenger runs $\text{Setup}(1^\lambda) \rightarrow (\text{mpk}, \text{msk})$ and gives **mpk** to \mathcal{A} . The challenger also picks a random coin $\text{coin} \leftarrow \{0, 1\}$ and keeps it secretly. After given

mpk, \mathcal{A} can adaptively make the following two types of queries to the challenger. These queries can be made in any order and arbitrarily many times.

Secret Key Queries. If \mathcal{A} submits $ID \in \mathcal{ID}$ to the challenger, the challenger returns $sk_{ID} \leftarrow \text{KeyGen}(\text{mpk}, \text{msk}, ID)$.

Challenge Queries. If \mathcal{A} submits a message M^* and an identity $ID^* \in \mathcal{ID}$ to the challenger, the challenger proceeds as follows. If $\text{coin} = 0$, it runs $\text{Encrypt}(\text{mpk}, ID^*, M^*) \rightarrow C^*$ and gives the challenge ciphertext C^* to \mathcal{A} . If $\text{coin} = 1$, it chooses the challenge ciphertext C^* from the distribution $\text{CTSAm}(\text{mpk})$ as $C^* \leftarrow \text{CTSAm}(\text{mpk})$ at random and gives it to \mathcal{A} .

We prohibit \mathcal{A} from making a challenge query for an identity ID^* such that it has already made a secret key query for the same $ID = ID^*$ and vice versa.

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

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

We say that Π is adaptively-anonymous secure, if there exists efficiently sampleable distribution $\text{CTSAm}(\text{mpk})$ and the advantage of any PPT \mathcal{A} is negligible in the above game. The term anonymous captures the fact that the ciphertext does not reveal the identity for which it was sent to. (Observe that $\text{CTSAm}(\text{mpk})$ depends on neither of ID^* nor M^* .)

Single Challenge Security. We can also consider a variant of the above security definition where we restrict the adversary to make the challenge query only once during the game. We call this security notion “single challenge adaptive anonymity”, and call the notion without the restriction “multi challenge security”. By a simple hybrid argument, we can show that these definitions are in fact equivalent in the sense that one implies another. However, the proof that the former implies the latter incurs a huge security reduction loss that is linear in the number of challenge queries. Since the focus of this paper is on tight security reductions, we typically differentiate these two notions.

Remark 1. *We say that an IBE scheme is stateful if the key generation algorithm has to record all previously issued secret keys, and always outputs the same secret key for the same identity. By the technique by Goldreich [Gol86], a stateful scheme can be converted to a stateless one (in which the key generation algorithm need not remember previous executions) by using PRFs. Since PRFs exist in the QROM without assuming any computational assumption as shown in Lem. 3, if we make the key size of PRFs sufficiently large, this conversion hardly affects the tightness. Therefore in this paper, we concentrate on constructing tightly secure stateful IBE schemes for simplicity.*

2.4 Background on Lattices

A (full-rank-integer) m -dimensional lattice Λ in \mathbb{Z}^m is a set of the form $\{\sum_{i \in [m]} x_i \mathbf{b}_i \mid x_i \in \mathbb{Z}\}$, where $\mathbf{B} = \{\mathbf{b}_1, \dots, \mathbf{b}_m\}$ are m linearly independent vectors in \mathbb{Z}^m . We call \mathbf{B} the basis of the lattice Λ . For any positive integers n, m and $q \geq 2$, a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and a vector $\mathbf{u} \in \mathbb{Z}_q^n$, we define $\Lambda^\perp(\mathbf{A}) = \{\mathbf{z} \in \mathbb{Z}^m \mid \mathbf{A}\mathbf{z} = \mathbf{0} \pmod{q}\}$, and $\Lambda_{\mathbf{u}}^\perp(\mathbf{A}) = \{\mathbf{z} \in \mathbb{Z}^m \mid \mathbf{A}\mathbf{z} = \mathbf{u} \pmod{q}\}$.

Gaussian Measures. For an m -dimensional lattice Λ , the discrete Gaussian distribution over Λ with center \mathbf{c} and parameter σ is defined as $D_{\Lambda, \sigma, \mathbf{c}}(\mathbf{x}) = \rho_{\sigma, \mathbf{c}}(\mathbf{x}) / \rho_{\sigma, \mathbf{c}}(\Lambda)$ for all $\mathbf{x} \in \Lambda$, where $\rho_{\sigma, \mathbf{c}}(\mathbf{x})$ is a Gaussian function defined as $\exp(-\pi \|\mathbf{x} - \mathbf{c}\|^2 / \sigma^2)$ and $\rho_{\sigma, \mathbf{c}}(\Lambda) = \sum_{\mathbf{x} \in \Lambda} \rho_{\sigma, \mathbf{c}}(\mathbf{x})$. Further for an m -dimensional shifted lattice $\Lambda + \mathbf{t}$, we define the Gaussian distribution $D_{\Lambda + \mathbf{t}, \sigma}$ with parameter

σ as the process of adding the vector \mathbf{t} to a sample from $D_{\Lambda, \sigma, -\mathbf{t}}$. Finally, we call D a B -bounded distribution, if all the elements in the support of D have absolute value smaller than B .

Discrete Gaussian Lemmas. The following lemmas are used to manipulate and obtain meaningful bounds on discrete Gaussian vectors.

Lemma 4 (Adopted from [GPV08], Lem. 5.2). *Let n, m, q be positive integers such that $m \geq 2n \log q$ and q a prime. Let σ be any positive real such that $\sigma \geq \sqrt{n + \log m}$. Then for all but $2^{-\Omega(n)}$ fraction of $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, we have that the distribution of $\mathbf{u} = \mathbf{A}\mathbf{e} \bmod q$ for $\mathbf{e} \leftarrow D_{\mathbb{Z}^m, \sigma}$ is $2^{-\Omega(n)}$ -close to uniform distribution over \mathbb{Z}_q^n . Furthermore, for a fixed $\mathbf{u} \in \mathbb{Z}_q^n$, the conditional distribution of $\mathbf{e} \leftarrow D_{\mathbb{Z}^m, \sigma}$, given $\mathbf{A}\mathbf{e} = \mathbf{u} \bmod q$ is $D_{\Lambda_{\mathbf{u}}^\perp(\mathbf{A}), \sigma}$.*

The following lemma is obtained by combining Lem. 4.4 in [MR07] and Lem. 5.3 in [GPV08].

Lemma 5 ([MR07], [GPV08]). *Let $\sigma > 16\sqrt{\log 2m/\pi}$ and \mathbf{u} be any vector in \mathbb{Z}_q^n . Then, for all but q^{-n} fraction of $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, we have that*

$$\Pr_{\mathbf{x} \leftarrow D_{\Lambda_{\mathbf{u}}^\perp(\mathbf{A}), \sigma}} [\|\mathbf{x}\| > \sigma\sqrt{m}] < 2^{-(m-1)}.$$

The following lemma can be obtained by a straightforward combination of Lem. 2.6, Lem. 2.10, and Lem. 5.3 in [GPV08] (See also [PR06, Pei07]).

Lemma 6 ([PR06, Pei07, GPV08]). *Let $\sigma > 16\sqrt{\log 2m/\pi}$ and \mathbf{u} be any vector in \mathbb{Z}_q^n . Then, for all but q^{-n} fraction of $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, we have*

$$\mathbf{H}_\infty(D_{\Lambda_{\mathbf{u}}^\perp(\mathbf{A}), \sigma}) \geq m - 1.$$

The following is a useful lemma used during the security proof. It allows the simulator to create new LWE samples from a given set of LWE samples (i.e., the LWE challenge provided to the simulator) for which it does not know the associating secret vector.² We would like to note that the following lemma is built on top of many previous results [Reg05, Pei10, BLP⁺13] and is formatted in a specific way to be useful in the security proof for LWE-based cryptosystems.

Lemma 7 (Noise Rerandomization, [KY16], Lem. 1). *Let q, ℓ, m be positive integers and r a positive real satisfying $r > \Omega(\sqrt{n})$. Let $\mathbf{b} \in \mathbb{Z}_q^m$ be arbitrary and \mathbf{z} chosen from $D_{\mathbb{Z}^m, r}$. Then there exists a PPT algorithm ReRand such that for any $\mathbf{V} \in \mathbb{Z}^{m \times \ell}$ and positive real $\sigma > s_1(\mathbf{V})$, the output of $\text{ReRand}(\mathbf{V}, \mathbf{b} + \mathbf{z}, r, \sigma)$ is distributed as $\mathbf{b}' = \mathbf{V}^\top \mathbf{b} + \mathbf{z}' \in \mathbb{Z}_q^\ell$ where the distribution of \mathbf{z}' is within $2^{-\Omega(n)}$ statistical distance of $D_{\mathbb{Z}^\ell, 2r\sigma}$.*

Sampling Algorithms. The following lemma states useful algorithms for sampling short vectors from lattices. In particular, the second preimage sampler is the exact gaussian sampler of [BLP⁺13], Lem. 2.3.

Lemma 8. ([GPV08, MP12, BLP⁺13]) *Let $n, m, q > 0$ be integers with $m > 3n \lceil \log q \rceil$.*

- $\text{TrapGen}(1^n, 1^m, q) \rightarrow (\mathbf{A}, \mathbf{T}_\mathbf{A})$: *a PPT algorithm that outputs a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and a full-rank matrix $\mathbf{T}_\mathbf{A} \in \mathbb{Z}^{m \times m}$, where $\mathbf{T}_\mathbf{A}$ is a basis for $\Lambda^\perp(\mathbf{A})$, the distribution of \mathbf{A} is $2^{-\Omega(n)}$ -close to uniform and $\|\mathbf{T}_\mathbf{A}\|_{\text{GS}} = O(\sqrt{n \log q})$.*

² Compared to [KY16] our choice of parameter is more conservative since we consider $2^{-\Omega(n)}$ statistical distance rather than $2^{-\omega(\log n)}$.

- $\text{SamplePre}(\mathbf{A}, \mathbf{T}_\mathbf{A}, \mathbf{u}, \sigma)$: a PPT algorithm that, given a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, a basis $\mathbf{T}_\mathbf{A} \in \mathbb{Z}^{m \times m}$ for $\Lambda^\perp(\mathbf{A})$, a vector $\mathbf{u} \in \mathbb{Z}_q^n$ and a Gaussian parameter $\sigma > \|\mathbf{T}_\mathbf{A}\|_{\text{CS}} \cdot \sqrt{\log(2m+4)/\pi}$, outputs a vector $\mathbf{e} \in \mathbb{Z}^m$ sampled from a distribution $2^{-\Omega(n)}$ -close to $D_{\Lambda_{\frac{1}{4}}^\perp(\mathbf{A}), \sigma}$.
- $\text{SampleZ}(\sigma)$: a PPT algorithm that, given a Gaussian parameter $\sigma > 16(\sqrt{\log 2m/\pi})$, outputs a vector $\mathbf{e} \in \mathbb{Z}^m$ sampled from a distribution $2^{-\Omega(n)}$ -close to $D_{\mathbb{Z}^m, \sigma}$.

Hardness Assumptions. We define the Learning with Errors (LWE) problem introduced by Regev [Reg05].

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

$$\text{Adv}_{\mathcal{A}}^{\text{LWE}_{n,m,q,\chi}} = \left| \Pr [\mathcal{A}(\mathbf{A}, \mathbf{A}^\top \mathbf{s} + \mathbf{z}) = 1] - \Pr [\mathcal{A}(\mathbf{A}, \mathbf{w} + \mathbf{z}) = 1] \right|$$

where $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$, $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, $\mathbf{w} \leftarrow \mathbb{Z}_q^m$, $\mathbf{z} \leftarrow \chi^m$. We say that the LWE assumption holds if $\text{Adv}_{\mathcal{A}}^{\text{LWE}_{n,m,q,\chi}}$ is negligible for all quantum polynomial time algorithm \mathcal{A} .

The (decisional) $\text{LWE}_{n,m,q,D_{\mathbb{Z},\alpha q}}$ for $\alpha q > 2\sqrt{n}$ has been shown by Regev [Reg05] via a quantum reduction to be as hard as approximating the worst-case SIVP and GapSVP problems to within $\tilde{O}(n/\alpha)$ factors in the ℓ_2 -norm in the worst case. In the subsequent works, (partial) dequantization of the reduction were achieved [Pei09, BLP⁺13].

We also define the LWE assumption against adversaries that can access to a quantum random oracle as is done by Boneh et al. [BDF⁺11].

Definition 3 (Learning with Errors relative to Quantum Random Oracle). Let n, m, q and χ be the same as in Def. 2, and a, b be some positive integers. For a quantum polynomial time algorithm \mathcal{A} , the advantage for the learning with errors problem $\text{LWE}_{n,m,q,\chi}$ of \mathcal{A} relative to a quantum random oracle is defined as follows:

$$\text{Adv}_{\mathcal{A}, \text{QRO}_{a,b}}^{\text{LWE}_{n,m,q,\chi}}(\lambda) = \left| \Pr [\mathcal{A}^{|\text{H}}(\mathbf{A}, \mathbf{A}^\top \mathbf{s} + \mathbf{z}) = 1] - \Pr [\mathcal{A}^{|\text{H}}(\mathbf{A}, \mathbf{w} + \mathbf{z}) = 1] \right|$$

where $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$, $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, $\mathbf{w} \leftarrow \mathbb{Z}_q^m$, $\mathbf{z} \leftarrow \chi^m$, $\text{H} \leftarrow \text{Func}(\{0, 1\}^a, \{0, 1\}^b)$. We say that the LWE assumption relative to an (a, b) -quantum random oracle holds if $\text{Adv}_{\mathcal{A}, \text{QRO}_{a,b}}^{\text{LWE}_{n,m,q,\chi}}(\lambda)$ is negligible for all quantum polynomial time algorithm \mathcal{A} .

It is easy to see that the LWE assumption relative to a quantum random oracle can be reduced to the LWE assumption with a certain loss of the time for the reduction by Lem. 1. Alternatively, if we assume the existence of a quantumly-accessible PRF, then the reduction loss can be made smaller. Namely, we have the following lemmas.

Lemma 9. For any n, m, q, χ, a, b , and a quantum polynomial time algorithm \mathcal{A} making at most Q oracle queries, there exists a quantum polynomial time algorithm \mathcal{B} such that

$$\text{Adv}_{\mathcal{A}, \text{QRO}_{a,b}}^{\text{LWE}_{n,m,q,\chi}}(\lambda) = \text{Adv}_{\mathcal{B}}^{\text{LWE}_{n,m,q,\chi}}(\lambda)$$

and $\text{Time}(\mathcal{B}) \approx \text{Time}(\mathcal{A}) + Q \cdot T_{a,b}^{2Q\text{-wise}}$ where $T_{a,b}^{2Q\text{-wise}}$ denotes the time to evaluate a $2Q$ -wise independent hash function from $\{0, 1\}^a$ to $\{0, 1\}^b$.

Lemma 10. *Let $F : \mathcal{K} \times \{0, 1\}^a \rightarrow \{0, 1\}^b$ be a quantumly-accessible PRF. For any n, m, q, χ, a, b and a quantum polynomial time algorithm \mathcal{A} making at most Q oracle queries, there exists quantum polynomial time algorithms \mathcal{B} and \mathcal{C} such that*

$$\text{Adv}_{\mathcal{A}, \text{QRO}_{a,b}}^{\text{LWE}_{n,m,q,\chi}}(\lambda) \leq \text{Adv}_{\mathcal{B}}^{\text{LWE}_{n,m,q,\chi}}(\lambda) + \text{Adv}_{\mathcal{C}, F}^{\text{PRF}}(\lambda)$$

and $\text{Time}(\mathcal{B}) \approx \text{Time}(\mathcal{A}) + Q \cdot T_F$ and $\text{Time}(\mathcal{C}) \approx \text{Time}(\mathcal{A})$ where T_F denotes the time to evaluate F .

In this paper, we give reductions from the security of IBE schemes to the LWE assumption relative to a quantum random oracle. Given such reductions, we can also reduce them to the LWE assumption or to the LWE assumption plus the security of quantumly-accessible PRFs by Lem. 9 or 10, respectively. The latter is tighter than the former at the cost of assuming the existence of quantumly-accessible PRFs.

Remark 2. *A keen reader may wonder why we have to require the extra assumption on the existence of PRFs when we are working in the QROM, since as we mentioned earlier in Sec. 2.2, it seems that we can use a QRO as a PRF. The point here is that during the security reduction, the simulator (which is given the classical LWE instance) must simulate the QRO query to the adversary against the LWE problem relative to a quantum random oracle query, hence, the simulator is not in possession of the QRO. Note that the reason why we are able to use the QRO as a PRF as mentioned in Rem. 1 is because the simulator is aiming to reduce the LWE problem relative to a quantum random oracle query to the IBE scheme. Specifically, in this case the simulator can use the QRO provided by its challenge to simulate a PRF.*

3 Tightly Secure Single Challenge GPV-IBE

In this section, we show that we can give a tight security proof for the original GPV-IBE [GPV08] in the single-challenge setting if we set the parameters appropriately. Such proofs can be given in both the classical ROM and QROM settings.

3.1 Construction

Let the identity space \mathcal{ID} of the scheme be $\mathcal{ID} = \{0, 1\}^{\ell_{\text{ID}}}$, where $\ell_{\text{ID}}(\lambda)$ denotes the identity-length. Let also $H : \{0, 1\}^{\ell_{\text{ID}}} \rightarrow \mathbb{Z}_q^n$ be a hash function treated as a random oracle during security analysis. The IBE scheme GPV is given as follows. For simplicity, we describe the scheme as a stateful one. As remarked in Rem. 1, we can make the scheme stateless without any additional assumption in the QROM.

Setup(1^λ): On input 1^λ , it first chooses a prime q , positive integers n, m , and Gaussian parameters α', σ , where all these values are implicitly a function of the security parameter λ . The precise parameter selection is specified in the following section. It then runs $(\mathbf{A}, \mathbf{T}_{\mathbf{A}}) \leftarrow \text{TrapGen}(1^n, 1^m, q)$ to generate a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ with a trapdoor $\mathbf{T}_{\mathbf{A}} \in \mathbb{Z}^{m \times m}$ such that $\|\mathbf{T}_{\mathbf{A}}\|_{\text{GS}} \leq O(n \log q)$. Then it outputs

$$\text{mpk} = \mathbf{A} \quad \text{and} \quad \text{msk} = \mathbf{T}_{\mathbf{A}}$$

KeyGen(mpk, msk, ID): If sk_{ID} is already generated, then this algorithm returns the same sk_{ID} . Otherwise it computes $\mathbf{u}_{\text{ID}} = H(\text{ID})$ and samples $\mathbf{e}_{\text{ID}} \in \mathbb{Z}^m$ such that

$$\mathbf{A} \mathbf{e}_{\text{ID}} = \mathbf{u}_{\text{ID}} \pmod{q}$$

using $\mathbf{e}_{\text{ID}} \leftarrow \text{SamplePre}(\mathbf{A}, \mathbf{T}_{\mathbf{A}}, \mathbf{u}_{\text{ID}}, \sigma)$. It returns $\text{sk}_{\text{ID}} = \mathbf{e}_{\text{ID}}$ as the secret key.

$\text{Enc}(\text{mpk}, \text{ID}, M)$: To encrypt a message $M \in \{0, 1\}$, it first samples $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, $\mathbf{x} \leftarrow D_{\mathbb{Z}^m, \alpha'q}$ and $x \leftarrow D_{\mathbb{Z}, \alpha'q}$. Then it sets $\mathbf{u}_{\text{ID}} = \mathbf{H}(\text{ID})$ and computes

$$c_0 = \mathbf{u}_{\text{ID}}^\top \mathbf{s} + x + M \lfloor q/2 \rfloor, \quad \mathbf{c}_1 = \mathbf{A}^\top \mathbf{s} + \mathbf{x}.$$

Finally, it outputs the ciphertext $C = (c_0, \mathbf{c}_1) \in \mathbb{Z}_q \times \mathbb{Z}_q^m$.

$\text{Dec}(\text{mpk}, \text{sk}_{\text{ID}}, C)$: To decrypt a ciphertext $C = (c_0, \mathbf{c}_1)$ with a secret key sk_{ID} , it computes $w = c_0 - \mathbf{c}_1^\top \mathbf{e}_{\text{ID}} \in \mathbb{Z}_q$ and outputs 0 if w is closer to 0 than to $\lfloor q/2 \rfloor$ modulo q . Otherwise it outputs 1.

3.2 Correctness and Parameter Selection

The following shows correctness of the above IBE scheme.

Lemma 11 (Correctness). *Suppose the parameters q , σ , and α' are such that*

$$\sigma > \|\mathbf{T}_{\mathbf{A}}\|_{\text{GS}} \cdot \sqrt{\log(2m+4)/\pi}, \quad \alpha' < 1/8\sigma m.$$

Let $\mathbf{e}_{\text{ID}} \leftarrow \text{KeyGen}(\mathbf{A}, \mathbf{T}_{\mathbf{A}}, \text{ID})$, $C \leftarrow \text{Enc}(\mathbf{A}, \text{ID}', M \in \{0, 1\})$ and $M' \leftarrow \text{Dec}(\mathbf{A}, \mathbf{e}_{\text{ID}}, C)$. If $\text{ID} = \text{ID}'$, then with overwhelming probability we have $M' = M$.

Proof. When the Dec algorithm operates as specified, we have

$$w = c_0 - \mathbf{e}_{\text{ID}}^\top \mathbf{c}_1 = M \lfloor q/2 \rfloor + \underbrace{x + \mathbf{e}_{\text{ID}}^\top \mathbf{x}}_{\text{error term}}.$$

By Lem. 8 and the condition posed on the choice of σ , we have that the distribution of \mathbf{e}_{ID} is $2^{-\Omega(n)}$ close to $D_{\Lambda_{\text{H}}^\perp(\mathbf{A}), \sigma}$. Therefore, by Lem. 5, we have $|x| \leq \alpha'q\sqrt{m}$, $\|\mathbf{x}\| \leq \alpha'q\sqrt{m}$, and $\|\mathbf{e}_{\text{ID}}\| \leq \sigma \cdot \sqrt{m}$ except for $2^{-\Omega(n)}$ probability. Then, the error term is bounded by

$$|x + \mathbf{e}_{\text{ID}}^\top \mathbf{x}| \leq |x| + \|\mathbf{e}_{\text{ID}}\| \cdot \|\mathbf{x}\| \leq 2\alpha'q\sigma m.$$

Hence, for the error term to have absolute value less than $q/4$, it suffices to choose q and α' as in the statement of the lemma. \square

Parameter Selection. For the system to satisfy correctness and make the security proof work, we need the following restrictions. Note that we will prove the security of the scheme under the LWE assumption whose noise rate is α , which is lower than α' that is used in the encryption algorithm.

- The error term is less than $q/4$ (i.e., $\alpha' < 1/8m\sigma$ by Lem. 11)
- TrapGen operates properly (i.e., $m > 3n \log q$ by Lem. 8)
- Samplable from $D_{\Lambda_{\text{H}}^\perp(\mathbf{A}), \sigma}$ (i.e., $\sigma > \|\mathbf{T}_{\mathbf{A}}\|_{\text{GS}} \cdot \sqrt{\log(2m+4)/\pi} = O(\sqrt{n \log m \log q})$ by Lem. 8),
- σ is sufficiently large so that we can apply Lem. 4 and 6 (i.e., $\sigma > \sqrt{n + \log m}$, $16\sqrt{\log 2m/\pi}$),
- We can apply Lem. 7 (i.e., $\alpha'/2\alpha > \sqrt{n(\sigma^2 m + 1)}$),
- $\text{LWE}_{n, m, q, D_{\mathbb{Z}, \alpha q}}$ is hard (i.e., $\alpha q > 2\sqrt{n}$).

To satisfy these requirements, for example, we can set the parameters $m, q, \sigma, \alpha, \alpha'$ as follows:

$$\begin{aligned} m &= n^{1+\kappa}, & q &= 10n^{3.5+4\kappa}, & \sigma &= n^{0.5+\kappa}, \\ \alpha'q &= n^{2+2\kappa}, & \alpha q &= 2\sqrt{n}, \end{aligned}$$

where $\kappa > 0$ is a constant that can be set arbitrarily small. To withstand attacks running in time 2^λ , we may set $n = \tilde{\Omega}(\lambda)$. In the above, we round up m to the nearest integer and q to the nearest largest prime. We remark that though the above parameter is worse compared to the original GPV-IBE scheme, this is due to our conservative choice of making the statistical error terms appearing in the reduction cost $2^{-\Omega(n)}$ rather than the standard negligible notion $2^{-\omega(\log \lambda)}$. The latter choice of parameters will lead to better parameters, which may be as efficient as the original GPV-IBE.

3.3 Security Proof in ROM

The following theorem addresses the security of GPV in the classical ROM setting. Our analysis departs from the original one [GPV08] and as a consequence results in a much tighter reduction.

Theorem 1. *The IBE scheme GPV is adaptively-anonymous single-challenge secure in the random oracle model assuming the hardness of $\text{LWE}_{n,m,q,D_{\mathbb{Z},\alpha q}}$. Namely, for any classical PPT adversary \mathcal{A} making at most Q_H random oracle queries to H and Q_{ID} secret key queries, there exists a classical PPT algorithm \mathcal{B} such that*

$$\text{Adv}_{\mathcal{A},\text{GPV}}^{\text{IBE}}(\lambda) \leq \text{Adv}_{\mathcal{B}}^{\text{LWE}_{n,m,q,D_{\mathbb{Z},\alpha q}}}(\lambda) + (Q_H + Q_{\text{ID}}) \cdot 2^{-\Omega(n)}$$

and

$$\text{Time}(\mathcal{B}) = \text{Time}(\mathcal{A}) + (Q_H + Q_{\text{ID}}) \cdot \text{poly}(\lambda).$$

Proof of Theorem 1. Let $\text{CTSAm}(\text{mpk})$ be an algorithm that outputs a random element from $\mathbb{Z}_q \times \mathbb{Z}_q^m$ and \mathcal{A} be a classical PPT adversary that attacks the adaptively-anonymous security of the IBE scheme. Without loss of generality, we make some simplifying assumptions on \mathcal{A} . First, we assume that whenever \mathcal{A} queries a secret key or asks for a challenge ciphertext, the corresponding ID has already been queried to the random oracle H . Second, we assume that \mathcal{A} makes the same query for the same random oracle at most once. Third, we assume that \mathcal{A} does not repeat secret key queries for the same identity more than once. We show the security of the scheme via the following games. In each game, we define X_i as the event that the adversary \mathcal{A} wins in Game_i .

Game₀ : This is the real security game. At the beginning of the game, $(\mathbf{A}, \mathbf{T}_{\mathbf{A}}) \leftarrow \text{TrapGen}(1^n, 1^m, q)$ is run and the adversary \mathcal{A} is given \mathbf{A} . The challenger then samples $\text{coin} \leftarrow \{0, 1\}$ and keeps it secret. During the game, \mathcal{A} may make random oracle queries, secret key queries, and the challenge query. These queries are handled as follows:

- When \mathcal{A} makes a random oracle query to H on ID , the challenger chooses a random vector $\mathbf{u}_{\text{ID}} \leftarrow \mathbb{Z}_q^n$ and locally stores the tuple $(\text{ID}, \mathbf{u}_{\text{ID}}, \perp)$, and returns \mathbf{u}_{ID} to \mathcal{A} .
- When the adversary \mathcal{A} queries a secret key for ID , the challenger computes $\mathbf{e}_{\text{ID}} = \text{SamplePre}(\mathbf{A}, \mathbf{T}_{\mathbf{A}}, \mathbf{u}_{\text{ID}}, \sigma)$ and returns \mathbf{e}_{ID} to \mathcal{A} .
- When the adversary makes the challenge query for ID^* and a message M^* , the challenger returns $(c_0, \mathbf{c}_1) \leftarrow \text{Encrypt}(\text{mpk}, \text{ID}, M)$ if $\text{coin} = 0$ and $(c_0, \mathbf{c}_1) \leftarrow \text{CTSAm}(\text{mpk})$ if $\text{coin} = 1$.

At the end of the game, \mathcal{A} outputs a guess $\widehat{\text{coin}}$ for coin . Finally, the challenger outputs $\widehat{\text{coin}}$. By definition, we have $|\Pr[X_0] - \frac{1}{2}| = |\Pr[\widehat{\text{coin}} - \text{coin}] - \frac{1}{2}| = \text{Adv}_{\mathcal{A}, \text{GPV}}^{\text{IBE}}(\lambda)$.

Game₁ : In this game, we change the way the random oracle queries to H are answered. When \mathcal{A} queries the random oracle H on ID , the challenger generates a pair $(\mathbf{u}_{\text{ID}}, \mathbf{e}_{\text{ID}})$ by first sampling $\mathbf{e}_{\text{ID}} \leftarrow \mathcal{D}_{\mathbb{Z}^m, \sigma}$ and setting $\mathbf{u}_{\text{ID}} = \mathbf{A}\mathbf{e}_{\text{ID}}$. Then it locally stores the tuple $(\text{ID}, \mathbf{u}_{\text{ID}}, \perp)$, and returns \mathbf{u}_{ID} to \mathcal{A} . Here, we remark that when \mathcal{A} makes a secret key query for ID , the challenger returns $\mathbf{e}'_{\text{ID}} \leftarrow \text{SamplePre}(\mathbf{A}, \mathbf{T}_{\mathbf{A}}, \mathbf{u}_{\text{ID}}, \sigma)$, which is independent from \mathbf{e}_{ID} that was generated in the simulation of the random oracle H on input ID . Note that in this game, we only change the distribution of \mathbf{u}_{ID} for each identity. Due to Lem. 4, the distribution of \mathbf{u}_{ID} in **Game₂** is $2^{-\Omega(n)}$ -close to that of **Game₁** except for $2^{-\Omega(n)}$ fraction of \mathbf{A} since we choose $\sigma > \sqrt{n + \log m}$. Therefore, the statistical distance between the view of \mathcal{A} in **Game₁** and **Game₂** is $2^{-\Omega(n)} + Q_H \cdot 2^{-\Omega(n)} < Q_H \cdot 2^{-\Omega(n)}$. Therefore, we have $|\Pr[X_1] - \Pr[X_2]| = Q_H \cdot 2^{-\Omega(n)}$.

Game₂ : In this game, we change the way secret key queries are answered. By the end of this game, the challenger will no longer require the trapdoor $\mathbf{T}_{\mathbf{A}}$ to generate the secret keys. When \mathcal{A} queries the random oracle on ID , the challenger generates a pair $(\mathbf{u}_{\text{ID}}, \mathbf{e}_{\text{ID}})$ as in the previous game. Then it locally stores the tuple $(\text{ID}, \mathbf{u}_{\text{ID}}, \mathbf{e}_{\text{ID}})$ and returns \mathbf{u}_{ID} to \mathcal{A} . When \mathcal{A} queries a secret key for ID , the challenger retrieves the unique tuple $(\text{ID}, \mathbf{u}_{\text{ID}}, \mathbf{e}_{\text{ID}})$ from local storage and returns \mathbf{e}_{ID} . For any fixed $\mathbf{u}_{\text{ID}} \in \mathbb{Z}_q^n$, let $\mathbf{e}_{\text{ID}, \mathbf{u}_{\text{ID}}}^{(1)}$ and $\mathbf{e}_{\text{ID}, \mathbf{u}_{\text{ID}}}^{(2)}$ be random variables that are distributed according to the distributions of sk_{ID} conditioning on $H(\text{ID}) = \mathbf{u}_{\text{ID}}$ in **Game₁** and **Game₂**, respectively. Due to Lem. 8, we have $\Delta(\mathbf{e}_{\text{ID}, \mathbf{u}_{\text{ID}}}^{(1)}, D_{\Lambda_{\mathbf{u}}^{\perp}(\mathbf{A}), \sigma}) \leq 2^{-\Omega(n)}$. On the other hand, due to Lem. 4, we have $\Delta(\mathbf{e}_{\text{ID}, \mathbf{u}_{\text{ID}}}^{(2)}, D_{\Lambda_{\mathbf{u}}^{\perp}(\mathbf{A}), \sigma}) \leq 2^{-\Omega(n)}$. Since \mathcal{A} obtains at most Q_{ID} user secret keys sk_{ID} , we have $|\Pr[X_1] - \Pr[X_2]| = Q_{\text{ID}} \cdot 2^{-\Omega(n)}$.

Game₃ : In this game, we change the way the matrix \mathbf{A} is generated. Concretely, the challenger chooses $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$ without generating the associated trapdoor $\mathbf{T}_{\mathbf{A}}$. By Lem. 8, this makes only $2^{-\Omega(n)}$ -statistical difference. Since the challenger can answer all the secret key queries without the trapdoor due to the change we made in the previous game, the view of \mathcal{A} is altered only negligibly. Therefore, we have $|\Pr[X_2] - \Pr[X_3]| = 2^{-\Omega(n)}$.

Game₄ : In this game, we change the way the challenge ciphertext is created when $\text{coin} = 0$. Recall in the previous games when $\text{coin} = 0$, the challenger created a valid challenge ciphertext as in the real scheme. In this game, to create the challenge ciphertext for identity ID^* and message bit M^* , the challenger first retrieves the unique tuple $(\text{ID}^*, \mathbf{u}_{\text{ID}^*}, \mathbf{e}_{\text{ID}^*})$ from local storage. Then the challenger picks $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, $\bar{\mathbf{x}} \leftarrow D_{\mathbb{Z}^m, \alpha q}$ and computes $\mathbf{v} = \mathbf{A}^{\top} \mathbf{s} + \bar{\mathbf{x}} \in \mathbb{Z}_q^m$. It then runs

$$\text{ReRand}([\mathbf{e}_{\text{ID}^*} | \mathbf{I}_m], \mathbf{v}, \alpha q, \frac{\alpha'}{2\alpha}) \rightarrow \mathbf{c}' \in \mathbb{Z}_q^{m+1}$$

from Lem. 7, where \mathbf{I}_m is the identity matrix with size m . Let $c'_0 \in \mathbb{Z}_q$ denote the first entry of \mathbf{c}' and $\mathbf{c}_1 \in \mathbb{Z}_q^m$ denote the remaining entries of \mathbf{c}' . Finally, the challenger outputs the challenge ciphertext as

$$C^* = (c_0 = c'_0 + M^* \lfloor q/2 \rfloor, \mathbf{c}_1). \quad (1)$$

We now proceed to bound $|\Pr[X_3] - \Pr[X_4]|$. We apply the noise rerandomization lemma (Lem. 7) with $\mathbf{V} = [\mathbf{e}_{\text{ID}^*} | \mathbf{I}_m]$, $\mathbf{b} = \mathbf{A}^{\top} \mathbf{s}$ and $\mathbf{z} = \bar{\mathbf{x}}$ to see that the distribution of \mathbf{c}' is negligibly close to the following:

$$\mathbf{c}' = \mathbf{V}^{\top} \mathbf{b} + \mathbf{x}' = \left(\mathbf{A} \cdot [\mathbf{e}_{\text{ID}^*} | \mathbf{I}_m] \right)^{\top} \mathbf{s} + \mathbf{x}' = [\mathbf{u}_{\text{ID}^*} | \mathbf{A}]^{\top} \mathbf{s} + \mathbf{x}'$$

where the distribution of \mathbf{x}' is $2^{-\Omega(n)}$ -close to $D_{\mathbb{Z}^{m+1}, \alpha'q}$. Here, the last equality follows from $\mathbf{A}\mathbf{e}_{\text{ID}^*} = \mathbf{u}_{\text{ID}^*}$ and we can appropriately apply the noise rerandomization lemma since we have the following for our parameter selection:

$$\alpha'/2\alpha > \sqrt{n(\sigma^2 m + 1)} \geq \sqrt{n(\|\mathbf{e}_{\text{ID}^*}\|^2 + 1)} \geq \sqrt{n} \cdot s_1([\mathbf{e}_{\text{ID}^*} | \mathbf{I}_m]),$$

where the second inequality holds with $1 - 2^{-\Omega(n)}$ probability. It can be seen that the challenge ciphertext is distributed statistically close to that in Game_3 . Therefore, we may conclude that $|\Pr[X_3] - \Pr[X_4]| = 2^{-\Omega(n)}$.

Game₅ : In this game, we further change the way the challenge ciphertext is created when $\text{coin} = 0$. If $\text{coin} = 0$, to create the challenge ciphertext the challenger first picks $\mathbf{b} \leftarrow \mathbb{Z}_q^m$, $\bar{\mathbf{x}} \leftarrow D_{\mathbb{Z}^m, \alpha q}$ and computes $\mathbf{v} = \mathbf{b} + \bar{\mathbf{x}} \in \mathbb{Z}_q^m$. It then sets $\mathbf{V} = [\mathbf{e}_{\text{ID}^*} | \mathbf{I}_m]$ and runs the **ReRand** algorithm as in Game_3 . Finally, it sets the challenge ciphertext as in Eq. (1). We claim that $|\Pr[X_4] - \Pr[X_5]|$ is negligible assuming the hardness of the $\text{LWE}_{n,m,q,D_{\mathbb{Z},\alpha q}}$ problem. To show this, we use \mathcal{A} to construct an LWE adversary \mathcal{B} as follows:

\mathcal{B} is given a problem instance of LWE as $(\mathbf{A}, \mathbf{v} = \mathbf{b} + \bar{\mathbf{x}}) \in \mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^m$ where $\bar{\mathbf{x}} \leftarrow D_{\mathbb{Z}^m, \alpha q}$. The task of \mathcal{B} is to distinguish whether $\mathbf{b} = \mathbf{A}^\top \mathbf{s}$ for some $\mathbf{s} \leftarrow \mathbb{Z}_q^n$ or $\mathbf{b} \leftarrow \mathbb{Z}_q^m$. \mathcal{B} sets the master public key mpk to be the LWE matrix \mathbf{A} . Note that unlike the real IBE scheme, \mathcal{B} does not require the master secret key $\mathbf{T}_\mathbf{A}$ due to the modification we made in Game_3 . To generate the challenge ciphertext, \mathcal{B} first picks $\text{coin} \leftarrow \{0, 1\}$. If $\text{coin} = 0$, it generates the challenge ciphertext as in Eq. (1) using \mathbf{v} , and returns it to \mathcal{A} . We emphasize that all \mathcal{B} needs to do to generate the ciphertext is to run the **ReRand** algorithm, which it can do without the knowledge of the secret randomness \mathbf{s} and $\bar{\mathbf{x}}$. If $\text{coin} = 1$, \mathcal{B} returns a random ciphertext using $\text{CTSAm}(\text{mpk})$. At the end of the game, \mathcal{A} outputs $\widehat{\text{coin}}$. Finally, \mathcal{B} outputs 1 if $\widehat{\text{coin}} = \text{coin}$ and 0 otherwise. It can be seen that if \mathbf{A}, \mathbf{v} is a valid LWE sample (i.e., $\mathbf{v} = \mathbf{A}^\top \mathbf{s}$), the view of the adversary corresponds to Game_4 . Otherwise (i.e., $\mathbf{v} \leftarrow \mathbb{Z}_q^m$), it corresponds to Game_5 . We therefore conclude that assuming the hardness of $\text{LWE}_{n,m,q,D_{\mathbb{Z},\alpha q}}$ problem we have $|\Pr[X_4] - \Pr[X_5]| = \text{negl}$.

Game₆ : In this game, we change the way the challenge ciphertext is created once more. If $\text{coin} = 0$, to create the challenge ciphertext the challenger first picks $\mathbf{b} \leftarrow \mathbb{Z}_q^m$, $\bar{\mathbf{x}}' \leftarrow D_{\mathbb{Z}^m, \alpha'q}$ and computes

$$\mathbf{c}' = [\mathbf{e}_{\text{ID}^*} | \mathbf{I}_m]^\top \mathbf{b} + \bar{\mathbf{x}}'.$$

It then parses \mathbf{c}' into c'_0 and \mathbf{c}_1 (as in Game_4) and sets the challenge ciphertext as Eq. (1). Similarly to the change from Game_3 to Game_4 , we have $|\Pr[X_5] - \Pr[X_6]| = 2^{-\Omega(n)}$ by Lem. 7.

It remains to show that no adversary has negligible chance in winning Game_6 . Notice that when $\text{coin} = 0$, the challenge ciphertext can be written as

$$c_0 = \mathbf{e}_{\text{ID}^*}^\top \mathbf{b} + x'_0 + \text{M}[q/2], \quad \mathbf{c}_1 = \mathbf{b} + \mathbf{x}'_1,$$

where x'_0 is the first entry of $\bar{\mathbf{x}}'$ and \mathbf{x}'_1 is the remaining entries. It suffices to show that the joint distribution of $(\mathbf{b}, \mathbf{e}_{\text{ID}^*}^\top \mathbf{b})$ is negligibly close to the uniform distribution over $\mathbb{Z}_q^m \times \mathbb{Z}_q$, conditioned on \mathbf{u}_{ID^*} . From the view of \mathcal{A} , \mathbf{e}_{ID^*} is distributed as $D_{\Lambda_{\mathbf{u}_{\text{ID}^*}}^\perp(\mathbf{A}), \sigma}$. By Lem. 6, we have

$$\mathbf{H}_\infty(\mathbf{e}_{\text{ID}^*}) \geq m - 1$$

for all but $2^{-\Omega(n)}$ fraction of \mathbf{A} . Now we can apply the leftover hash lemma since \mathbf{b} is distributed uniformly at random over \mathbb{Z}_q^m and conclude that $(\mathbf{b}, \mathbf{e}_{\text{ID}^*}^\top \mathbf{b})$ is $\sqrt{q}/2^{m-1}$ -close to the uniform distribution. Hence, we have $\Pr[X_6] \leq 2^{-\Omega(n)} + \sqrt{q}/2^{m-1} < 2^{-\Omega(n)}$.

Therefore, combining everything together, the theorem is proven. \square

3.4 Security Proof in QROM

As we explained in the introduction, our analysis in the ROM can be easily be extended to the QROM setting. We can prove the following theorem that addresses the security of the GPV-IBE scheme in the QROM setting. The analysis here is different from that by Zhandry [Zha12b], who gave the first security proof for the GPV-IBE scheme in the QROM setting and our analysis here is much tighter.

Theorem 2. *The IBE scheme GPV is adaptively-anonymous single-challenge secure assuming the hardness of $\text{LWE}_{n,m,q,D_{\mathbb{Z},\alpha q}}$ in the quantum random oracle model. Namely, for any quantum adversary \mathcal{A} making at most Q_H queries to $|H\rangle$ and Q_{ID} secret key queries, there exists a quantum algorithm \mathcal{B} making $Q_H + Q_{ID}$ quantum random oracle queries such that*

$$\text{Adv}_{\mathcal{A},\text{GPV}}^{\text{IBE}}(\lambda) \leq \text{Adv}_{\mathcal{B},\text{QRO}_{\ell_{ID},\ell_r}}^{\text{LWE}_{n,m,q,D_{\mathbb{Z},\alpha q}}}(\lambda) + (Q_H^2 + Q_{ID}) \cdot 2^{-\Omega(n)}$$

and

$$\text{Time}(\mathcal{B}) = \text{Time}(\mathcal{A}) + (Q_H + Q_{ID}) \cdot \text{poly}(\lambda)$$

where ℓ_r denotes the length of the randomness for $\text{Sample}\mathbb{Z}$.

Proof of Theorem 2. Let $\text{CTSAm}(\text{mpk})$ be an algorithm that outputs a random element from $\mathbb{Z}_q \times \mathbb{Z}_q^m$ and \mathcal{A} be a quantum adversary that attacks the adaptively-anonymous security of the IBE scheme. Without loss of generality, we can assume that \mathcal{A} makes secret key queries on the same identity at most once. We show the security of the scheme via the following games. In each game, we define X_i as the event that the adversary \mathcal{A} wins in Game_i .

Game₀ : This is the real security game for the adaptively-anonymous security. At the beginning of the game, the challenger chooses a random function $H : \{0, 1\}^{\ell_{ID}} \rightarrow \mathbb{Z}_q^n$. Then it generates $(\mathbf{A}, \mathbf{T}_A) \leftarrow \text{TrapGen}(1^n, 1^m, q)$ and gives \mathbf{A} to \mathcal{A} . Then it samples $\text{coin} \leftarrow \{0, 1\}$ and keeps it secret. During the game, \mathcal{A} may make (quantum) random oracle queries, secret key queries, and a challenge query. These queries are handled as follows:

- When \mathcal{A} makes a random oracle query on a quantum state $\sum_{ID,y} \alpha_{ID,y} |ID\rangle |y\rangle$, the challenger returns $\sum_{ID,y} \alpha_{ID,y} |ID\rangle |H(ID) \oplus y\rangle$.
- When \mathcal{A} makes a secret key query on ID , the challenger samples $\mathbf{e}_{ID} = \text{SamplePre}(\mathbf{A}, \mathbf{T}_A, \mathbf{u}_{ID}, \sigma)$ and returns \mathbf{e}_{ID} to \mathcal{A} .
- When \mathcal{A} makes a challenge query for ID^* and a message M^* , the challenger returns $(c_0, c_1) \leftarrow \text{Encrypt}(\text{mpk}, ID, M)$ if $\text{coin} = 0$ and $(c_0, c_1) \leftarrow \text{CTSAm}(\text{mpk})$ if $\text{coin} = 1$.

At the end of the game, \mathcal{A} outputs a guess $\widehat{\text{coin}}$ for coin . Finally, the challenger outputs $\widehat{\text{coin}}$. By definition, we have $|\Pr[X_0] - \frac{1}{2}| = |\Pr[\widehat{\text{coin}} - \text{coin}] - \frac{1}{2}| = \text{Adv}_{\mathcal{A},\text{GPV}}^{\text{IBE}}(\lambda)$.

Game₁ : In this game, we change the way the random oracle H is simulated. Namely, the challenger first chooses another random function $\widehat{H} \leftarrow \text{Func}(\{0, 1\}^{\ell_{ID}}, \{0, 1\}^{\ell_r})$. Then we define $H(ID) := \mathbf{A}\mathbf{e}_{ID}$ where $\mathbf{e}_{ID} := \text{Sample}\mathbb{Z}(\sigma; \widehat{H}(ID))$, and use this H throughout the game. For any fixed ID , the distribution of $H(ID)$ is identical and its statistical distance from the uniform distribution is $2^{-\Omega(n)}$ for all but $2^{-\Omega(n)}$ fraction of \mathbf{A} due to Lem. 4 since we choose $\sigma > \sqrt{n + \log m}$. Note that in this game, we only change the distribution of \mathbf{u}_{ID} for each identity, and the way we create secret keys are

unchanged. Then due to Lem. 2, we have $|\Pr[X_0] - \Pr[X_1]| = 2^{-\Omega(n)} + 4Q_H^2 \sqrt{2^{-\Omega(n)}} = Q_H^2 \cdot 2^{-\Omega(n)}$.

Game₂ : In this game, we change the way secret key queries are answered. By the end of this game, the challenger will no longer require the trapdoor \mathbf{T}_A to generate the secret keys. When \mathcal{A} queries a secret key for ID , the challenger returns $\mathbf{e}_{\text{ID}} := \text{SampleZ}(\sigma; \widehat{\mathbf{H}}(\text{ID}))$. For any fixed $\mathbf{u}_{\text{ID}} \in \mathbb{Z}_q^n$, let $\mathbf{e}_{\text{ID}, \mathbf{u}_{\text{ID}}}^{(1)}$ and $\mathbf{e}_{\text{ID}, \mathbf{u}_{\text{ID}}}^{(2)}$ be random variables that are distributed according to the distributions of \mathbf{e}_{ID} conditioning on $\mathbf{H}(\text{ID}) = \mathbf{u}_{\text{ID}}$ in **Game₁** and **Game₂**, respectively. Due to Lem. 8, we have $\Delta(\mathbf{e}_{\text{ID}, \mathbf{u}_{\text{ID}}}^{(1)}, D_{\Lambda_{\mathbf{u}_{\text{ID}}}^\perp(\mathbf{A}), \sigma}) \leq 2^{-\Omega(n)}$. On the other hand, due to Lem. 4, we have $\Delta(\mathbf{e}_{\text{ID}, \mathbf{u}_{\text{ID}}}^{(2)}, D_{\Lambda_{\mathbf{u}_{\text{ID}}}^\perp(\mathbf{A}), \sigma}) \leq 2^{-\Omega(n)}$. Since \mathcal{A} obtains at most Q_{ID} user secret keys \mathbf{e}_{ID} , we have $|\Pr[X_1] - \Pr[X_2]| = Q_{\text{ID}} \cdot 2^{-\Omega(n)}$.

Game₃ : In this game, we change the way the matrix \mathbf{A} is generated. Concretely, the challenger chooses $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$ without generating the associated trapdoor \mathbf{T}_A . By Lem. 8, the distribution of \mathbf{A} differs at most by $2^{-\Omega(n)}$. Since the challenger can answer all the secret key queries without the trapdoor due to the change we made in the previous game, the view of \mathcal{A} is altered only by $2^{-\Omega(n)}$. Therefore, we have $|\Pr[X_2] - \Pr[X_3]| = 2^{-\Omega(n)}$.

Game₄ : In this game, we change the way the challenge ciphertext is created when $\text{coin} = 0$. Recall in the previous games when $\text{coin} = 0$, the challenger created a valid challenge ciphertext as in the real scheme. In this game, to create the challenge ciphertext for identity ID^* and message bit M^* , the challenger first computes $\mathbf{e}_{\text{ID}^*} := \text{SampleZ}(\sigma; \widehat{\mathbf{H}}(\text{ID}^*))$ and $\mathbf{u}_{\text{ID}^*} := \mathbf{A} \mathbf{e}_{\text{ID}^*}$. Then the challenger picks $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, $\bar{\mathbf{x}} \leftarrow D_{\mathbb{Z}^m, \alpha q}$ and computes $\mathbf{v} = \mathbf{A}^\top \mathbf{s} + \bar{\mathbf{x}} \in \mathbb{Z}_q^m$. It then runs

$$\text{ReRand}([\mathbf{e}_{\text{ID}^*} | \mathbf{I}_m], \mathbf{v}, \alpha q, \frac{\alpha'}{2\alpha}) \rightarrow \mathbf{c}' \in \mathbb{Z}_q^{m+1}$$

from Lem. 7, where \mathbf{I}_m is the identity matrix with size m . Let $c'_0 \in \mathbb{Z}_q$ denote the first entry of \mathbf{c}' and $\mathbf{c}_1 \in \mathbb{Z}_q^m$ denote the remaining entries of \mathbf{c}' . Finally, the challenger outputs the challenge ciphertext as

$$C^* = (c_0 = c'_0 + M^* \lfloor q/2 \rfloor, \mathbf{c}_1). \quad (2)$$

We now proceed to bound $|\Pr[X_3] - \Pr[X_4]|$. We apply the noise rerandomization lemma (Lem. 7) with $\mathbf{V} = [\mathbf{e}_{\text{ID}^*} | \mathbf{I}_m]$, $\mathbf{b} = \mathbf{A}^\top \mathbf{s}$ and $\mathbf{z} = \bar{\mathbf{x}}$ to see that the following equation holds:

$$\mathbf{c}' = \mathbf{V}^\top \mathbf{b} + \mathbf{x}' = \left(\mathbf{A} \cdot [\mathbf{e}_{\text{ID}^*} | \mathbf{I}_m] \right)^\top \mathbf{s} + \mathbf{x}' = [\mathbf{u}_{\text{ID}^*} | \mathbf{A}]^\top \mathbf{s} + \mathbf{x}'$$

where \mathbf{x}' is distributed according to a distribution whose statistical distance is at most $2^{-\Omega(n)}$ from $D_{\mathbb{Z}^{m+1}, \alpha' q}$. Here, the last equality follows from $\mathbf{A} \mathbf{e}_{\text{ID}^*} = \mathbf{u}_{\text{ID}^*}$ and we can appropriately apply the noise rerandomization lemma since we have the following for our parameter selection:

$$\alpha'/2\alpha > \sqrt{n(\sigma^2 m + 1)} \geq \sqrt{n(\|\mathbf{e}_{\text{ID}^*}\|^2 + 1)} \geq \sqrt{n} \cdot s_1([\mathbf{e}_{\text{ID}^*} | \mathbf{I}_m]),$$

where the second inequality holds with $1 - 2^{-\Omega(n)}$ probability. It therefore follows that the statistical distance between the distributions of the challenge ciphertext in **Game₃** and **Game₄** is at most $2^{-\Omega(n)}$. Therefore, we may conclude that $|\Pr[X_3] - \Pr[X_4]| = 2^{-\Omega(n)}$.

Game₅ : In this game, we further change the way the challenge ciphertext is created when $\text{coin} = 0$. If $\text{coin} = 0$, to create the challenge ciphertext the challenger first picks $\mathbf{b} \leftarrow \mathbb{Z}_q^m$, $\bar{\mathbf{x}} \leftarrow D_{\mathbb{Z}^m, \alpha q}$

and computes $\mathbf{v} = \mathbf{b} + \bar{\mathbf{x}} \in \mathbb{Z}_q^m$. It then runs the ReRand algorithm as in Game_4 . Finally, it sets the challenge ciphertext as in Eq. (2). We claim that $|\Pr[X_4] - \Pr[X_5]|$ is negligible assuming the hardness of the $\text{LWE}_{n,m,q,D_{\mathbb{Z},\alpha q}}$ problem relative to a quantum random oracle $|\widehat{\mathbf{H}}\rangle : \{0,1\}^{\ell_{\text{ID}}} \rightarrow \{0,1\}^{\ell_r}$. To show this, we use \mathcal{A} to construct an adversary \mathcal{B} that breaks the LWE assumption relative to $|\widehat{\mathbf{H}}\rangle$.

\mathcal{B} is given a problem instance of LWE as $(\mathbf{A}, \mathbf{v} = \mathbf{b} + \bar{\mathbf{x}}) \in \mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^m$ where $\bar{\mathbf{x}} \leftarrow D_{\mathbb{Z}^m, \alpha q}$. The task of \mathcal{B} is to distinguish whether $\mathbf{b} = \mathbf{A}^\top \mathbf{s}$ for some $\mathbf{s} \leftarrow \mathbb{Z}_q^n$ or $\mathbf{b} \leftarrow \mathbb{Z}_q^m$. First, we remark that \mathcal{B} can simulate the quantum random oracle $|\mathbf{H}\rangle$ for \mathcal{A} by using its own random oracle $|\widehat{\mathbf{H}}\rangle$ because \mathbf{H} is programmed as $\mathbf{H}(\text{ID}) := \mathbf{A}\mathbf{e}_{\text{ID}}$ where $\mathbf{e}_{\text{ID}} := \text{SampleZ}(\sigma; \mathbf{H}(\text{ID}))$ by the modification we made in Game_1 . \mathcal{B} sets the master public key mpk to be the LWE matrix \mathbf{A} . Note that unlike the real IBE scheme, \mathcal{B} does not require the master secret key $\mathbf{T}_{\mathbf{A}}$ due to the modification we made in Game_3 . Namely, when \mathcal{A} queries ID for the key oracle, \mathcal{B} just returns $\mathbf{e}_{\text{ID}} := \text{SampleZ}(\sigma; \widehat{\mathbf{H}}(\text{ID}))$. To generate the challenge ciphertext, \mathcal{B} first picks $\text{coin} \leftarrow \{0,1\}$. If $\text{coin} = 0$, it generates the challenge ciphertext as in Eq. (2) using \mathbf{v} , and returns it to \mathcal{A} . We emphasize that all \mathcal{B} needs to do to generate the ciphertext is to run the ReRand algorithm, which it can do without the knowledge of the secret randomness \mathbf{s} and $\bar{\mathbf{x}}$. If $\text{coin} = 1$, \mathcal{B} returns a random ciphertext using $\text{CTSAm}(\text{mpk})$. At the end of the game, \mathcal{A} outputs $\widehat{\text{coin}}$. Finally, \mathcal{B} outputs 1 if $\widehat{\text{coin}} = \text{coin}$ and 0 otherwise.

It can be seen that if \mathbf{A}, \mathbf{v} is a valid LWE sample (i.e., $\mathbf{v} = \mathbf{A}^\top \mathbf{s}$), the view of the adversary corresponds to Game_4 . Otherwise (i.e., $\mathbf{v} \leftarrow \mathbb{Z}_q^m$), it corresponds to Game_5 . Therefore we have $|\Pr[X_4] - \Pr[X_5]| = \text{Adv}_{\mathcal{B}, \text{QRO}_{\ell_{\text{ID}}, \ell_r}}^{\text{LWE}_{n,m,q,D_{\mathbb{Z},\alpha q}}}(\lambda)$. As for the running time, we have $\text{Time}(\mathcal{B}) = \text{Time}(\mathcal{A}) + (Q_{\text{H}} + Q_{\text{ID}}) \cdot \text{poly}(\lambda)$ since all \mathcal{B} has to do is to run \mathcal{A} once plus to compute some additional computations that can be done in a fixed polynomial time whenever \mathcal{A} makes a quantum random oracle or secret key query.

Game_6 : In this game, we further change the way the challenge ciphertext is created. If $\text{coin} = 0$, to create the challenge ciphertext the challenger first picks $\mathbf{b} \leftarrow \mathbb{Z}_q^m$, $\mathbf{x}' \leftarrow D_{\mathbb{Z}^m, \alpha' q}$ and computes

$$\mathbf{c}' = [\mathbf{e}_{\text{ID}^*} | \mathbf{I}_m]^\top \mathbf{b} + \mathbf{x}'.$$

It then parses \mathbf{c}' into c'_0 and \mathbf{c}_1 (as in Game_4) and sets the challenge ciphertext as Eq. (2). Similarly to the change from Game_3 to Game_4 , we have $|\Pr[X_5] - \Pr[X_6]| = 2^{-\Omega(n)}$ by Lem. 7.

It remains to show that no adversary has non-negligible chance in winning Game_6 . Notice that when $\text{coin} = 0$, the challenge ciphertext can be written as

$$c_0 = \mathbf{e}_{\text{ID}^*}^\top \mathbf{b} + x'_0 + \text{M}\lfloor q/2 \rfloor, \quad \mathbf{c}_1 = \mathbf{b} + \mathbf{x}'_1,$$

where x'_0 is the first entry of \mathbf{x}' and \mathbf{x}'_1 is the remaining entries. It suffices to show that the joint distribution of $(\mathbf{b}, \mathbf{e}_{\text{ID}^*}^\top \mathbf{b})$ is statistically close to the uniform distribution over $\mathbb{Z}_q^m \times \mathbb{Z}_q$, conditioned on \mathbf{u}_{ID^*} . From the view of \mathcal{A} , \mathbf{e}_{ID^*} is distributed as $D_{\Lambda_{\mathbf{u}(\text{ID}^*)}^\perp(\mathbf{A}), \sigma}$ because all information of \mathbf{e}_{ID^*} revealed to \mathcal{A} is $\mathbf{H}(\text{ID}^*) = \mathbf{A}\mathbf{e}_{\text{ID}^*}$ where $\widehat{\mathbf{e}}_{\text{ID}^*} = \text{SampleZ}(\sigma; \widehat{\mathbf{H}}(\text{ID}^*))$ and $\widehat{\mathbf{H}}(\text{ID}^*)$ is completely random from the view of \mathcal{A} . (Remark that $\widehat{\mathbf{H}}(\text{ID}^*)$ is used in the game only when \mathcal{A} queries ID^* to the key generation oracle, which is prohibited in the adaptively-anonymous security game.) By Lem. 6, we have

$$\mathbf{H}_\infty(\mathbf{e}_{\text{ID}^*}) \geq m - 1$$

for all but $2^{-\Omega(n)}$ fraction of \mathbf{A} . Now we can apply the leftover hash lemma since \mathbf{b} is distributed uniformly at random over \mathbb{Z}_q^m and conclude that $(\mathbf{b}, \mathbf{e}_{\text{ID}^*}^\top \mathbf{b})$ is $\sqrt{q/2^{m-1}}$ -close to the uniform distribution by the leftover hash lemma. Hence, we have $\Pr[X_6] \leq 2^{-\Omega(n)} + \sqrt{q/2^{m-1}} < 2^{-\Omega(n)}$.

Therefore, combining everything together, the theorem is proven. \square

4 (Almost) Tightly Secure Multi-Challenge IBE

In this section, we propose an IBE scheme that is (almost) tightly secure in the multi-challenge setting. The security of the scheme is proven both in the classical ROM and QROM settings. Our construction is obtained by applying the Katz-Wang [KW03] technique to the original GPV-IBE scheme. Since the proofs require some previous results on random extractions and lossy mode LWE, we first review them below.

4.1 Randomness Extraction

We recap some definitions and results on randomness extraction. As we have already introduced in the preliminaries, the min-entropy of a random variable X was defined as $\mathbf{H}_\infty(X) = -\log(\max_x \Pr[X = x])$. A similar notion called the *average min-entropy*, as introduced by Dodis et al. [DORS04], is defined as follows:

$$\tilde{\mathbf{H}}_\infty(X|I) = -\log(\mathbb{E}_{i \leftarrow I}[2^{-\mathbf{H}_\infty(X|I=i)}]).$$

The average min-entropy corresponds to the optimal probability of guessing X , given knowledge of I . Min-entropy is a rather fragile notion, since a single high-probability element can ruin the min-entropy of an otherwise good distribution. Therefore, it is often more beneficial to work with the ϵ -smooth min-entropy introduced by Renner and Wolf [RW04], which considers all distributions that are ϵ -close to X , but which has higher entropy:

$$\mathbf{H}_\infty^\epsilon(X) = \max_{Y: \Delta(X,Y) \leq \epsilon} \mathbf{H}_\infty(Y).$$

Similarly, a smooth version of average min-entropy can be defined as follows:

$$\tilde{\mathbf{H}}_\infty^\epsilon(X|I) = \max_{(Y,J): \Delta((X,I),(Y,J)) \leq \epsilon} \tilde{\mathbf{H}}_\infty(Y|J).$$

We recall the definition of universal hash functions and provide an elementary construction of them that will be used in our construction of multi-insatnce secure IBE schemes.

Definition 4 (Universal Hash Functions). *A family of functions $\mathcal{H} = \{h : \mathcal{X} \rightarrow \mathcal{D}\}_h$ is called a family of universal hash functions, if for all $x, x' \in \mathcal{X}$ with $x \neq x'$, we have $\Pr_{h \leftarrow \mathcal{H}}[h(x) = h(x')] \leq \frac{1}{|\mathcal{D}|}$.*

Fact 1. *Let $q > 2$. Let $\mathcal{H} = \{\mathbf{u} : \mathbb{Z}_q^n \rightarrow \mathbb{Z}_q\}_{\mathbf{u} \in \mathbb{Z}_q^n}$ be a family of hash functions, where $\mathbf{u}(\mathbf{s})$ is defined as $\mathbf{u}(\mathbf{s}) = \mathbf{u}^\top \mathbf{s} \pmod q$. Then, \mathcal{H} is a family of universal hash functions.*

The following lemma gives a lower bound for the smooth average min-entropy of some random variable when partial related information is leaked.

Lemma 12 ([AKPW13], Lem. 2.4). *Let X, Y , and Z be correlated random variables and \mathcal{Z} be some set such that $\Pr[Z \in \mathcal{Z}] \leq \epsilon$ and $|\mathcal{Z}| \leq 2^z$. Then, for any $\epsilon' > 0$, $\tilde{\mathbf{H}}_\infty^{\epsilon+\epsilon'}(X|(Y,Z)) \geq \tilde{\mathbf{H}}_\infty^{\epsilon'}(X|Y) - z$.*

The following is a generalization of the leftover hash lemma due to [DORS04]. Here, we provide the smoothed-variant of the lemma. Roughly, this relates to the number of extractable bits that look nearly uniform to the adversary who knows some value that is ϵ -close to the random variable I .

Lemma 13 (Generalized Leftover Hash Lemma). *Let $\mathcal{H} = \{h : \mathcal{X} \rightarrow \mathcal{D}\}$ be a family of universal hash functions. Let X be an independent random variable with values in \mathcal{X} , let I be any random variable. Then, for any $\epsilon \geq 0$, we have*

$$\Delta((h, h(X), I), (h, U(\mathcal{D}), I)) \leq 2\epsilon + \frac{1}{2} \cdot \sqrt{2^{-\tilde{\mathbf{H}}_\infty^\epsilon(X|I)} \cdot |\mathcal{D}|}.$$

Proof. Let (Y, J) be random variables such that

$$(Y, J) = \arg \max_{(Y, J): \Delta((X, I), (Y, J)) \leq \epsilon} \tilde{\mathbf{H}}_\infty(Y|J).$$

Here, note that we have $\Delta(I, J) \leq \epsilon$. Then,

$$\begin{aligned} & \Delta((h, h(X), I), (h, U(\mathcal{D}), I)) \\ & \leq \Delta((h, h(X), I), (h, h(Y), J)) + \Delta((h, h(Y), J), (h, U(\mathcal{D}), J)) \\ & \quad + \Delta((h, U(\mathcal{D}), J), (h, U(\mathcal{D}), I)) \\ & \leq 2\epsilon + \frac{1}{2} \cdot \sqrt{2^{-\tilde{\mathbf{H}}_\infty(Y|J)} \cdot |\mathcal{D}|} \end{aligned} \tag{3}$$

In the above derivation, Eq. (3) follows from the standard generalized leftover hash lemma of [DORS04] and the definition of smooth average min-entropy. This completes the proof. \square

Lossy Mode for LWE. It is well known that the secret vector $\mathbf{s} \in \mathbb{Z}_q^n$ is uniquely defined (with all but negligible probability) given an LWE instance $(\mathbf{A}, \mathbf{A}^\top \mathbf{s} + \mathbf{x}) \in \mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^m$ when \mathbf{A} is uniformly chosen from $\mathbb{Z}_q^{n \times m}$ for sufficiently large m . On the other hand, if we sample \mathbf{A} from a special distribution which is computationally indistinguishable from the uniform distribution over $\mathbb{Z}_q^{n \times m}$, then the pair $(\mathbf{A}, \mathbf{A}^\top \mathbf{s} + \mathbf{x})$ leaks almost no information of the secret \mathbf{s} . Since the LWE problem of this version does not reveal much information about the secret \mathbf{s} , this instance is often referred to as the “lossy mode”. A series of works, e.g., [GKPV10, BKPW12, AKPW13, LSSS17] have investigated the lossy nature of the problem. We first formally describe the procedure `SampleLossy` to sample \mathbf{A} in the lossy mode. Let n, m, ℓ be positive integers, and χ be a distribution over \mathbb{Z} .

`SampleLossy`(n, m, ℓ, χ): It samples $\mathbf{C} \leftarrow U(\mathbb{Z}_q^{n \times \ell})$, $\mathbf{B} \leftarrow U(\mathbb{Z}_q^{\ell \times m})$, and $\mathbf{F} \leftarrow \chi^{n \times m}$, and outputs $\mathbf{A} = \mathbf{CB} + \mathbf{F}$.

It is easy to see that \mathbf{A} in the lossy mode is indistinguishable from random under the LWE assumption by a standard hybrid argument.

Lemma 14. *For any quantum polynomial time algorithm \mathcal{A} , there exists a quantum polynomial algorithm \mathcal{B} such that*

$$\left| \Pr[\mathcal{A}(\mathbf{A}_0) = 1] - \Pr[\mathcal{A}(\mathbf{A}_1) = 1] \right| \leq n \cdot \text{Adv}_{\mathcal{B}}^{\text{LWE}_{\ell, m, q, \chi^m}}(\lambda)$$

and $\text{Time}(\mathcal{B}) \approx \text{Time}(\mathcal{A})$ where $\mathbf{A}_0 \leftarrow \mathbb{Z}_q^{n \times m}$ and $\mathbf{A}_1 \leftarrow \text{SampleLossy}(n, m, \ell, \chi)$. If \mathcal{A} has access to a quantum random oracle from $\{0, 1\}^a$ to $\{0, 1\}^b$, then the right hand side is replaced by $n \cdot \text{Adv}_{\mathcal{B}, \text{QRO}_{a,b}}^{\text{LWE}_{\ell, n, q, \chi}}(\lambda)$, and the number of quantum random oracle queries by \mathcal{A} and \mathcal{B} are the same.

We will be using the following lemma slightly adapted from the works of [AKPW13]. In particular, we consider the smooth *average* min-entropy rather than the smooth min-entropy.

Lemma 15 (Adapted from [AKPW13], Lem. B.4). *Let n, ℓ, m, q, γ be positive integer parameters, α, β be some Gaussian parameters, and χ be a distribution (all parameterized by the security parameter λ), such that $\Pr_{x \leftarrow \chi}[|x| \geq \beta q] \leq \text{negl}(\lambda)$ and $\alpha \geq \beta \gamma n m$. Let \mathbf{s} and \mathbf{e} be random variables distributed according to $U([- \gamma, \gamma]^n)$ and $D_{\mathbb{Z}^m, \alpha q}$, respectively. Furthermore, let \mathbf{A} be a matrix sampled by $\text{SampleLossy}(n, m, \ell, \chi)$. Then, for any $\epsilon \geq 2^{-\lambda}$, we have the following:*

$$\tilde{\mathbf{H}}_{\infty}^{\epsilon}(\mathbf{s} | \mathbf{A}, \mathbf{A}^{\top} \mathbf{s} + \mathbf{e}) \geq \mathbf{H}_{\infty}(\mathbf{s}) - (\ell + 2\lambda) \log q - \text{negl}(\lambda).$$

Remark 3. *In [AKPW13], they do not consider the average over $(\mathbf{A}, \mathbf{A}^{\top} \mathbf{s} + \mathbf{e})$, i.e., they only prove the above lemma for $\mathbf{H}_{\infty}^{\epsilon}(\mathbf{s} | \mathbf{A}, \mathbf{A}^{\top} \mathbf{s} + \mathbf{e})$. However, their proof actually works for the above statement as well. This change will be useful during the security proof when we apply the (smoothed) generalized leftover hash lemma, which works for smooth-average min-entropies.*

4.2 Construction

Let the identity space \mathcal{ID} of the scheme be $\mathcal{ID} = \{0, 1\}^{\ell_{\text{ID}}}$, where $\ell_{\text{ID}}(\lambda)$ denotes the identity-length. Let also $\mathbf{H} : \{0, 1\}^{\ell_{\text{ID}}+1} \rightarrow \mathbb{Z}_q^n$ be a hash function treated as a random oracle during the security analysis where ℓ_{ID} denotes the identity-length. The IBE scheme GPV_{mult} is given as follows. For simplicity, we describe the scheme as a stateful one. As remarked in Rem. 1, we can make the scheme stateless without any additional assumption in the QROM.

Setup(1^λ): On input 1^λ , it first chooses a prime q , positive integers n, m, γ , and Gaussian parameters α, σ , where all these values are implicitly a function of the security parameter λ . The precise parameter selection is specified in the following section. It then runs $(\mathbf{A}, \mathbf{T}_{\mathbf{A}}) \leftarrow \text{TrapGen}(1^n, 1^m, q)$ to generate a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ with a trapdoor $\mathbf{T}_{\mathbf{A}} \in \mathbb{Z}^{m \times m}$ such that $\|\mathbf{T}_{\mathbf{A}}\|_{\text{GS}} \leq O(n \log q)$. Then it outputs

$$\text{mpk} = \mathbf{A} \quad \text{and} \quad \text{msk} = \mathbf{T}_{\mathbf{A}}$$

KeyGen($\text{mpk}, \text{msk}, \text{ID}$): If sk_{ID} is already generated, then this algorithm returns it. Otherwise it picks $b_{\text{ID}} \leftarrow \{0, 1\}$, computes $\mathbf{u}_{\text{ID} \| b_{\text{ID}}} = \mathbf{H}(\text{ID} \| b_{\text{ID}})$, and samples $\mathbf{e}_{\text{ID} \| b_{\text{ID}}} \in \mathbb{Z}^m$ such that

$$\mathbf{A} \mathbf{e}_{\text{ID} \| b_{\text{ID}}} = \mathbf{u}_{\text{ID} \| b_{\text{ID}}} \pmod{q}$$

as $\mathbf{e}_{\text{ID} \| b_{\text{ID}}} \leftarrow \text{SamplePre}(\mathbf{A}, \mathbf{T}_{\mathbf{A}}, \mathbf{u}_{\text{ID} \| b_{\text{ID}}}, \sigma)$. It returns $\text{sk}_{\text{ID}} = (b_{\text{ID}}, \mathbf{e}_{\text{ID} \| b_{\text{ID}}})$ as the secret key.

Enc($\text{mpk}, \text{ID}, \text{M}$): To encrypt a message $\text{M} \in \{0, 1\}$, it first samples $\mathbf{s} \leftarrow U([- \gamma, \gamma])$, $\mathbf{x} \leftarrow D_{\mathbb{Z}^m, \alpha q}$. Then it computes $\mathbf{u}_{\text{ID} \| 0} = \mathbf{H}(\text{ID} \| 0)$ and $\mathbf{u}_{\text{ID} \| 1} = \mathbf{H}(\text{ID} \| 1)$ and sets the ciphertext as

$$c_0 = \mathbf{u}_{\text{ID} \| 0}^{\top} \mathbf{s} + \text{M} \lfloor q/2 \rfloor, \quad c_1 = \mathbf{u}_{\text{ID} \| 1}^{\top} \mathbf{s} + \text{M} \lfloor q/2 \rfloor, \quad \mathbf{c}_2 = \mathbf{A}^{\top} \mathbf{s} + \mathbf{x}.$$

Finally, it outputs the ciphertext $C = (c_0, c_1, \mathbf{c}_2) \in \mathbb{Z}_q \times \mathbb{Z}_q \times \mathbb{Z}_q^m$.

Dec($\text{mpk}, \text{sk}_{\text{ID}}, C$): To decrypt a ciphertext $C = (c_0, c_1, \mathbf{c}_2)$ with a secret key sk_{ID} , it computes $w = c_{b_{\text{ID}}} - \mathbf{c}_2^{\top} \mathbf{e}_{\text{ID} \| b_{\text{ID}}} \in \mathbb{Z}_q$ and outputs 0 if w is closer to 0 than to $\lfloor q/2 \rfloor$ modulo q . Otherwise it outputs 1.

4.3 Correctness and Parameter Selection

The following shows correctness of the above IBE scheme.

Lemma 16 (Correctness). *Suppose the parameters q , σ , and α are such that*

$$\sigma > \|\mathbf{T}_A\|_{\text{GS}} \cdot \sqrt{\log(2m+4)/\pi}, \quad \alpha < 1/4\sigma m.$$

Let $\mathbf{e}_{\text{ID}\|b_{\text{ID}}} \leftarrow \text{KeyGen}(\mathbf{A}, \mathbf{T}_A, \text{ID})$, $C \leftarrow \text{Enc}(\mathbf{A}, \text{ID}', M \in \{0, 1\})$ and $M' \leftarrow \text{Dec}(\mathbf{A}, \mathbf{e}_{\text{ID}\|b_{\text{ID}}}, C)$. If $\text{ID} = \text{ID}'$, then with overwhelming probability we have $M' = M$.

Proof. When the Dec algorithm operates as specified, we have

$$w = c_{b_{\text{ID}}} - \mathbf{e}_{\text{ID}\|b_{\text{ID}}}^\top \mathbf{c}_2 = M \lfloor q/2 \rfloor + \underbrace{\mathbf{e}_{\text{ID}\|b_{\text{ID}}}^\top \mathbf{x}}_{\text{error term}}.$$

By Lem. 8 and the condition posed on the choice of σ , we have that the distribution of $\mathbf{e}_{\text{ID}\|b_{\text{ID}}}$ is $2^{-\Omega(n)}$ close to $D_{\Lambda_{\mathbf{u}_{\text{ID}\|b_{\text{ID}}}}^\perp(\mathbf{A}), \sigma}$. Therefore, by Lem. 5, we have $\|\mathbf{x}\| \leq \alpha q \sqrt{m}$, and $\|\mathbf{e}_{\text{ID}\|b_{\text{ID}}}\| \leq \sigma \cdot \sqrt{m}$ except for $2^{-\Omega(n)}$ probability. Then, the error term is bounded by

$$|\mathbf{e}_{\text{ID}\|b_{\text{ID}}}^\top \mathbf{x}| \leq \|\mathbf{e}_{\text{ID}\|b_{\text{ID}}}\| \cdot \|\mathbf{x}\| \leq \alpha q \sigma m.$$

Hence, for the error term to have absolute value less than $q/4$, it suffices to choose q and α as in the statement of the lemma. \square

Parameter Selection. For the system to satisfy correctness and make the security proof work, we need the following restrictions. Note that we will prove the security of the scheme under $\text{LWE}_{\ell, m, q, \chi^m}$, where ℓ is specified in the following.

- The error term is less than $q/4$ (i.e., $\alpha < 1/4m\sigma$ by Lem. 16)
- TrapGen operates properly (i.e., $m > 3n \log q$ by Lem. 8)
- Samplable from $D_{\Lambda_{\mathbf{u}_{\text{ID}\|b_{\text{ID}}}}^\perp(\mathbf{A}), \sigma}$ (i.e., $\sigma > \|\mathbf{T}_A\|_{\text{GS}} \cdot \sqrt{\log(2m+4)/\pi} = O(\sqrt{n \log m \log q})$ by Lem. 8),
- $\text{LWE}_{\ell, n, q, \chi}$ is hard so that we can use the lossy mode in the proof by Lem. 14 (i.e., $\chi = D_{\mathbb{Z}, 2\sqrt{\ell}}$),
- σ is sufficiently large so that we can apply Lem. 4 (i.e., $\sigma > \sqrt{n + \log m}$),
- we can apply to Lem. 15 in the proof (i.e., $\alpha > \beta \gamma n m$ for β such that $\Pr_{x \leftarrow \chi}[|x| \geq \beta q] \leq \text{negl}(\lambda)$),
- we can apply the generalized leftover hash lemma (Lem. 13) in the proof (i.e., $n \log(2\gamma) - (\ell + 3\lambda) \log q \geq \log q + \Omega(n)$).

To satisfy these requirements, for example, we can set the parameters $\ell, n, m, q, \sigma, \alpha, \beta, \gamma$ as follows:

$$\begin{aligned} n &= 25\ell, & m &= n^{1+\kappa}, & \sigma &= n^{0.5+\kappa}, & q &= 5n^{5.5+3\kappa}, \\ \alpha q &= n^{4+\kappa}, & \beta q &= n, & \gamma &= n, \end{aligned}$$

where $\kappa > 0$ is a constant that can be set arbitrarily small. To withstand attacks running in time 2^λ , we may set $\ell = \tilde{\Omega}(\lambda)$. In the above, we round up m to the nearest integer and q to the nearest largest prime. As the case with the single-challenge setting, if we make the more aggressive choice of using the negligible notion $2^{-\omega(\log \lambda)}$, we will be able to obtain better parameter selections.

4.4 Security Proof in ROM

We can (almost) tightly prove the security of our IBE scheme GPV_{mult} both in the classical ROM and QROM settings. The following theorem addresses the security of GPV_{mult} in the classical ROM setting.

Theorem 3. *The IBE scheme GPV_{mult} is adaptively-anonymous multi-challenge secure assuming the hardness of $\text{LWE}_{\ell,m,q,\chi}$ in the random oracle model, where $\chi = D_{\mathbb{Z},\alpha q}$. Namely, for any classical PPT adversary \mathcal{A} making at most Q_{H} queries to H , Q_{ch} challenge queries, and Q_{ID} secret key queries, there exists a classical PPT algorithm \mathcal{B} such that*

$$\text{Adv}_{\mathcal{A},\text{GPV}_{\text{mult}}}^{\text{IBE}}(\lambda) \leq 3n \cdot \text{Adv}_{\mathcal{B}}^{\text{LWE}_{\ell,m,q,D_{\mathbb{Z},\alpha q}}}(\lambda) + (Q_{\text{H}} + Q_{\text{ID}} + Q_{\text{ch}}) \cdot 2^{-\Omega(n)}$$

and

$$\text{Time}(\mathcal{B}) = \text{Time}(\mathcal{A}) + (Q_{\text{H}} + Q_{\text{ID}} + Q_{\text{ch}}) \cdot \text{poly}(\lambda).$$

Proof of Theorem 3. Let $\text{CTS}_{\text{sam}}(\text{mpk})$ be an algorithm that outputs (c_0, c_1, \mathbf{c}_2) such that $c_0 \leftarrow \mathbb{Z}_q$, $c_1 \leftarrow \mathbb{Z}_q$, and $\mathbf{c}_2 = \mathbf{A}^\top \mathbf{s} + \mathbf{x}$ for $\mathbf{s} \leftarrow U([- \gamma, \gamma])$ and $\mathbf{x} \leftarrow D_{\mathbb{Z}^m, \alpha q}$. Let also \mathcal{A} be a classical PPT adversary that attacks the (multi-challenge) adaptively-anonymous security of the IBE scheme. Without loss of generality, we make some simplifying assumptions on \mathcal{A} . First, we assume that whenever \mathcal{A} queries a secret key or asks for a challenge ciphertext, $\text{ID}||0$ and $\text{ID}||1$ for the corresponding ID has already been queried to the random oracle H . We also assume that whenever \mathcal{A} queries H on input $\text{ID}||b$, it also queries H on input $\text{ID}||\bar{b}$ as well. Furthermore, we assume that \mathcal{A} makes the same query for the random oracle or secret key oracle at most once. We show the security of the scheme via the following games. In each game, we define X_i as the event that the adversary \mathcal{A} wins in Game_i .

Game₀ : This is the real security game. At the beginning of the game, $(\mathbf{A}, \mathbf{T}_{\mathbf{A}}) \leftarrow \text{TrapGen}(1^n, 1^m, q)$ is run and the adversary \mathcal{A} is given \mathbf{A} . The challenger then samples $\text{coin} \leftarrow \{0, 1\}$ and keeps it secret. During the game, \mathcal{A} may make random oracle queries, secret key queries, and challenge queries. These queries are handled as follows:

- When \mathcal{A} queries the random oracle H on $\text{ID}||0$ and $\text{ID}||1$, the challenger samples $\mathbf{u}_{\text{ID}||0} \leftarrow \mathbb{Z}_q^n$ and $\mathbf{u}_{\text{ID}||1} \leftarrow \mathbb{Z}_q^n$ and returns $\mathbf{u}_{\text{ID}||0}$ and $\mathbf{u}_{\text{ID}||1}$.
- When \mathcal{A} queries a secret key for ID, the challenger first computes $b_{\text{ID}} \leftarrow \{0, 1\}$ and returns $\mathbf{e}_{\text{ID}||b_{\text{ID}}} = \text{SamplePre}(\mathbf{A}, \mathbf{T}_{\mathbf{A}}, \mathbf{u}_{\text{ID}||b_{\text{ID}}}, \sigma)$ to \mathcal{A} where $\mathbf{u}_{\text{ID}||b_{\text{ID}}}$ is the vector chosen during the simulation of the random oracle.
- When the adversary makes the challenge query for ID^* and a message M^* , the challenger returns $(c_0, c_1, \mathbf{c}_2) \leftarrow \text{Encrypt}(\text{mpk}, \text{ID}^*, M^*)$ if $\text{coin} = 0$ and $(c_0, c_1, \mathbf{c}_2) \leftarrow \text{CTS}_{\text{sam}}(\text{mpk})$ if $\text{coin} = 1$. It then returns $C^* = (c_0, c_1, \mathbf{c}_2)$ to \mathcal{A} .

At the end of the game, \mathcal{A} outputs a guess $\widehat{\text{coin}}$ for coin . Finally, the challenger outputs $\widehat{\text{coin}}$. By definition, we have $|\Pr[X_0] - \frac{1}{2}| = |\Pr[\widehat{\text{coin}} - \text{coin}] - \frac{1}{2}| = \text{Adv}_{\mathcal{A},\text{GPV}_{\text{mult}}}^{\text{IBE}}(\lambda)$.

Game₁ : In this game, we change the game so that b_{ID} is chosen when \mathcal{A} queries the random oracle H on input $\text{ID}||0$ and $\text{ID}||1$ rather than when \mathcal{A} queries a secret key for ID. It is clear that this is only a conceptual change and we have $\Pr[X_0] = \Pr[X_1]$.

Game₂ : In this game, we change the way the random oracle queries to H are answered. When \mathcal{A} queries H on inputs $\text{ID}||0$ and $\text{ID}||1$, the challenger first samples $b_{\text{ID}} \leftarrow \{0, 1\}$ as specified

in the previous game. Then, it generates a pair $(\mathbf{u}_{\text{ID}\|b_{\text{ID}}}, \mathbf{e}_{\text{ID}\|b_{\text{ID}}})$ by first sampling $\mathbf{e}_{\text{ID}\|b_{\text{ID}}} \leftarrow \mathcal{D}_{\mathbb{Z}^m, \sigma}$ and setting $\mathbf{u}_{\text{ID}\|b_{\text{ID}}} = \mathbf{A}\mathbf{e}_{\text{ID}\|b_{\text{ID}}}$. It also samples $\mathbf{u}_{\text{ID}\|\bar{b}_{\text{ID}}} \leftarrow \mathbb{Z}_q^n$. Then it locally stores the tuples $(\text{ID}\|b_{\text{ID}}, \mathbf{u}_{\text{ID}\|b_{\text{ID}}}, \perp)$ and $(\text{ID}\|\bar{b}_{\text{ID}}, \mathbf{u}_{\text{ID}\|\bar{b}_{\text{ID}}}, \perp)$, and returns $\mathbf{u}_{\text{ID}\|0}$ and $\mathbf{u}_{\text{ID}\|1}$ to \mathcal{A} . Here, we remark that when \mathcal{A} makes a secret key query for ID, the challenger returns $\mathbf{e}'_{\text{ID}\|b_{\text{ID}}} \leftarrow \text{SamplePre}(\mathbf{A}, \mathbf{T}_{\mathbf{A}}, \mathbf{u}_{\text{ID}\|b_{\text{ID}}}, \sigma)$, which is independent from $\mathbf{e}_{\text{ID}\|b_{\text{ID}}}$ that was generated in the simulation of the random oracle H on input ID||0 and ID||1. Note that in this game, we only change the distribution of $\mathbf{u}_{\text{ID}\|b_{\text{ID}}}$ for each identity. Due to Lem. 4, the distribution of $\mathbf{u}_{\text{ID}\|b_{\text{ID}}}$ in Game_2 is $2^{-\Omega(n)}$ -close to that of Game_1 except for $2^{-\Omega(n)}$ fraction of \mathbf{A} since we choose $\sigma > \sqrt{n + \log m}$. Therefore, the statistical distance between the view of \mathcal{A} in Game_1 and Game_2 is $2^{-\Omega(n)} + Q_{\text{H}} \cdot 2^{-\Omega(n)} < Q_{\text{H}} \cdot 2^{-\Omega(n)}$. Therefore, we have $|\Pr[X_1] - \Pr[X_2]| = Q_{\text{H}} \cdot 2^{-\Omega(n)}$.

Game₃ : In this game, we change the way secret key queries are answered. By the end of this game, the challenger will no longer require the trapdoor $\mathbf{T}_{\mathbf{A}}$ to generate the secret keys. When \mathcal{A} queries H on ID||0 and ID||1, the challenger generates a pair $(\mathbf{u}_{\text{ID}\|b_{\text{ID}}}, \mathbf{e}_{\text{ID}\|b_{\text{ID}}})$ and $\mathbf{u}_{\text{ID}\|\bar{b}_{\text{ID}}}$ as in the previous game. Then it locally stores the tuple $(\text{ID}\|b_{\text{ID}}, \mathbf{u}_{\text{ID}\|b_{\text{ID}}}, \mathbf{e}_{\text{ID}\|b_{\text{ID}}})$ and $(\text{ID}\|\bar{b}_{\text{ID}}, \mathbf{u}_{\text{ID}\|\bar{b}_{\text{ID}}}, \perp)$, and returns $\mathbf{u}_{\text{ID}\|b_{\text{ID}}}$ and $\mathbf{u}_{\text{ID}\|\bar{b}_{\text{ID}}}$ to \mathcal{A} . When \mathcal{A} queries a secret key for ID, the challenger retrieves the unique tuple $(\text{ID}\|b_{\text{ID}}, \mathbf{u}_{\text{ID}\|b_{\text{ID}}}, \mathbf{e}_{\text{ID}\|b_{\text{ID}}})$ from local storage and returns $\mathbf{e}_{\text{ID}\|b_{\text{ID}}}$. For any fixed $\mathbf{u}_{\text{ID}\|b_{\text{ID}}} \in \mathbb{Z}_q^n$, let $\mathbf{e}_{\text{ID}\|b_{\text{ID}}}^{(2)}$ and $\mathbf{e}_{\text{ID}\|b_{\text{ID}}}^{(3)}$ be random variables that are distributed according to the distributions of $\text{sk}_{\text{ID}\|b_{\text{ID}}}$ conditioning on $\text{H}(\text{ID}\|b_{\text{ID}}) = \mathbf{u}_{\text{ID}\|b_{\text{ID}}}$ in Game_2 and Game_3 , respectively. Due to Lem. 8, we have $\Delta(\mathbf{e}_{\text{ID}\|b_{\text{ID}}}^{(2)}, D_{\Lambda_{\mathbf{u}_{\text{ID}\|b_{\text{ID}}}}(\mathbf{A}, \sigma)}) \leq 2^{-\Omega(n)}$. On the other hand, due to Lem. 4, we have $\Delta(\mathbf{e}_{\text{ID}\|b_{\text{ID}}}^{(3)}, D_{\Lambda_{\mathbf{u}_{\text{ID}\|b_{\text{ID}}}}(\mathbf{A}, \sigma)}) \leq 2^{-\Omega(n)}$. Since \mathcal{A} obtains at most Q_{ID} user secret keys, we have $|\Pr[X_2] - \Pr[X_3]| = Q_{\text{ID}} \cdot 2^{-\Omega(n)}$.

Game₄ : In this game, we change the way the matrix \mathbf{A} is generated. Concretely, the challenger chooses $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$ without generating the associated trapdoor $\mathbf{T}_{\mathbf{A}}$. By Lem. 8, this makes only $2^{-\Omega(n)}$ -statistical difference. Since the challenger can answer all the secret key queries without the trapdoor due to the change we made in the previous game, the view of \mathcal{A} is altered only negligibly. Therefore, we have $|\Pr[X_3] - \Pr[X_4]| = 2^{-\Omega(n)}$.

Game₅ : In this game, we change \mathbf{A} to the lossy mode (See Lem. 15). We claim that $|\Pr[X_4] - \Pr[X_5]|$ is negligible assuming the hardness of the $\text{LWE}_{\ell, m, q, D_{\mathbb{Z}, \alpha q}}$ problem. To show this, we use \mathcal{A} to construct an adversary \mathcal{B}' that distinguishes random \mathbf{A} from that in lossy mode. This can be done by a straightforward reduction since \mathcal{B}' does not require the master secret key $\mathbf{T}_{\mathbf{A}}$ to simulate the game due to the modification we made in Game_3 . We therefore by Lem. 14 conclude that there exists \mathcal{B} such that $|\Pr[X_4] - \Pr[X_5]| = n \cdot \text{Adv}_{\mathcal{B}}^{\text{LWE}_{\ell, m, q, D_{\mathbb{Z}, \alpha q}}}(\lambda)$.

Game₆ : In this game, we change the way the challenge ciphertexts are generated. Recall that by our assumption, \mathcal{A} makes queries for H on inputs ID*||0 and ID*||1 before making a challenge query for identity ID*. When \mathcal{A} makes a challenge query for $(\text{ID}^*, \text{M}^*)$ it samples a ciphertext as $C^* \leftarrow \text{CTSam}(\text{mpk})$ and returns C^* to \mathcal{A} if $\text{coin} = 1$. If $\text{coin} = 0$, it generates the ciphertext as

$$c_{b_{\text{ID}^*}} = \mathbf{u}_{\text{ID}^*\|b_{\text{ID}^*}}^{\top} \mathbf{s} + \text{M}^* \lfloor q/2 \rfloor, \quad c_{\bar{b}_{\text{ID}^*}} \leftarrow \mathbb{Z}_q, \quad \mathbf{c}_2 = \mathbf{A}^{\top} \mathbf{s} + \mathbf{x}.$$

for $\mathbf{s} \leftarrow U([- \gamma, \gamma])$, where b_{ID^*} is generated when the hash queries for ID*||0 and ID*||1 were made. It then rearranges the terms if necessary and returns (c_0, c_1, \mathbf{c}_2) to \mathcal{A} .

We argue that the view of \mathcal{A} is statistically close to that in the previous game. To prove this, we make a hybrid argument over all the challenge ciphertexts and change $c_{\bar{b}_{\text{ID}^*}}$ to be random one-by-one (when $\text{coin} = 0$). To conclude, it suffices to show that each ciphertext $c_{\bar{b}_{\text{ID}^*}}$ is distributed

$2^{-\Omega(n)}$ close to the uniform distribution over \mathbb{Z}_q . Specifically, for each ciphertext, for any $\epsilon' = 2^{-\lambda}$, we have

$$\begin{aligned}
\tilde{\mathbf{H}}_\infty^{\epsilon'}(\mathbf{s}|\mathbf{A}, c_{b_{\text{ID}^*}}, \mathbf{c}_2) &\geq \tilde{\mathbf{H}}_\infty^{\epsilon'}(\mathbf{s}|\mathbf{A}, \mathbf{c}_2) - \log q \\
&\geq \mathbf{H}_\infty(\mathbf{s}) - (\ell + 2\lambda + 1) \log q - \text{negl}(\lambda) \\
&= n \log(2\gamma) - (\ell + 3\lambda) \log q \\
&\geq \log q + \Omega(n)
\end{aligned}$$

where the first inequality follows by applying Lem. 12 with $\mathcal{Z} = \mathbb{Z}_q$ and $\epsilon = 0$, the second inequality follows from Lem. 15, and the last inequality follows from our parameter choice. This implies that $\Delta(\mathbf{u}_{\text{ID}^* \| \bar{b}_{\text{ID}^*}}, \mathbf{u}_{\text{ID}^* \| \bar{b}_{\text{ID}^*}}^\top \mathbf{s}, |\mathbf{A}, c_{b_{\text{ID}^*}}, \mathbf{c}_2) \leq 2^{-\Omega(n)}$ for $\mathbf{u}_{\text{ID}^* \| \bar{b}_{\text{ID}^*}} \leftarrow \mathbb{Z}_q^n$ by Lem. 13 together with Fact 1. Therefore, we have $|\Pr[X_5] - \Pr[X_6]| = Q_{\text{ch}} \cdot 2^{-\Omega(n)}$.

From Game₇ to Game₁₀ in the following, we undo the changes we added from Game₂ to Game₅.

Game₇ : In this game, \mathbf{A} is sampled as $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$. Similarly to the change from Game 4 to

Game 5, there exists \mathcal{B} such that we have $|\Pr[X_6] - \Pr[X_7]| = n \cdot \text{Adv}_{\mathcal{B}}^{\text{LWE}_{\ell, m, q, D_{\mathbb{Z}}, \alpha q}}(\lambda)$.

Game₈ : In this game, \mathbf{A} is sampled with a trapdoor as $(\mathbf{A}, \mathbf{T}_{\mathbf{A}}) \leftarrow \text{TrapGen}(1^n, 1^m, q)$. Similarly to the change from Game 3 to Game 4, we have $|\Pr[X_7] - \Pr[X_8]| = 2^{-\Omega(n)}$.

Game₉ : In this game, we change the way secret key queries are answered. When \mathcal{A} makes a secret key query for ID, the challenger returns $\mathbf{e}'_{\text{ID} \| b_{\text{ID}}} \leftarrow \text{SamplePre}(\mathbf{A}, \mathbf{T}_{\mathbf{A}}, \mathbf{u}_{\text{ID} \| b_{\text{ID}}}, \sigma)$, where b_{ID} and \mathbf{u}_{ID} are chosen when random oracle queries on ID||0 and ID||1 are made. Similarly to the change from Game 2 to Game 3, we have $|\Pr[X_8] - \Pr[X_9]| = Q_{\text{H}} \cdot 2^{-\Omega(n)}$.

Game₁₀ : In this game, we change the way the random oracle queries to H are answered. When \mathcal{A} queries the random oracle on ID||0 and ID||1, the challenger samples $\mathbf{u}_{\text{ID} \| 0}, \mathbf{u}_{\text{ID} \| 1} \leftarrow \mathbb{Z}_q^n$ and locally stores the tuples $(\text{ID} \| 0, \mathbf{u}_{\text{ID} \| 0}, \perp)$ and $(\text{ID} \| 1, \mathbf{u}_{\text{ID} \| 1}, \perp)$, and returns $\mathbf{u}_{\text{ID} \| 0}$ and $\mathbf{u}_{\text{ID} \| 1}$ to \mathcal{A} . These $\mathbf{u}_{\text{ID} \| 0}$ and $\mathbf{u}_{\text{ID} \| 1}$ are also used when answering the secret key queries. Similarly to the change from Game 1 to Game 2, we have $|\Pr[X_{10}] - \Pr[X_9]| = Q_{\text{H}} \cdot 2^{-\Omega(n)}$.

Because of the changes we introduced in Game₇ to Game₁₀, we can add the following change.

Game₁₁ : In this game, we change the way the challenge ciphertexts are generated. When \mathcal{A} makes a challenge query for $(\text{ID}^*, \mathbf{M}^*)$, it returns a random ciphertext if $\text{coin} = 1$. If $\text{coin} = 0$, it generates the ciphertext as

$$c_{b_{\text{ID}^*}} \leftarrow \mathbb{Z}_q, \quad c_{\bar{b}_{\text{ID}^*}} = \mathbf{u}_{\text{ID}^* \| \bar{b}_{\text{ID}^*}}^\top \mathbf{s} + \mathbf{M}^* \lfloor q/2 \rfloor, \quad \mathbf{c}_2 = \mathbf{A}^\top \mathbf{s} + \mathbf{x}.$$

for $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, where b_{ID^*} is generated when the hash queries for ID*||0 and ID*||1 were made. It then rearranges the terms if necessary and returns (c_0, c_1, \mathbf{c}_2) to \mathcal{A} . We claim that this change is only conceptual. Note that the distribution of the ciphertexts in this game corresponds to that in the previous game, if we were to flip the value of b_{ID^*} for every challenge identity ID*. However, since \mathcal{A} never makes a secret key query for ID* and $\mathbf{u}_{\text{ID}^* \| 0}$ and $\mathbf{u}_{\text{ID}^* \| 1}$ are sampled from exactly the same distribution, the value of b_{ID^*} is information theoretically hidden from \mathcal{A} . This implies that the distributions are the same in this and the previous game. Therefore, we have $\Pr[X_{11}] = \Pr[X_{10}]$.

Game₁₂ : In this game, we further change the way the challenge ciphertexts are generated. When \mathcal{A} makes a challenge query for $(\text{ID}^*, \mathbf{M}^*)$, it returns a ciphertext sampled from $\text{CTSam}(\text{mpk})$ if $\text{coin} = 1$. If $\text{coin} = 0$, it generates the ciphertext as

$$c_{b_{\text{ID}^*}} \leftarrow \mathbb{Z}_q, \quad c_{\bar{b}_{\text{ID}^*}} \leftarrow \mathbb{Z}_q, \quad \mathbf{c}_2 = \mathbf{A}^\top \mathbf{s} + \mathbf{x}.$$

for $\mathbf{s} \leftarrow \mathbb{Z}_q^n$ and $\mathbf{x} \leftarrow D_{\mathbb{Z}^m, \alpha q}$. It then returns (c_0, c_1, \mathbf{c}_2) to \mathcal{A} . We claim that the change is unnoticed by \mathcal{A} assuming the LWE assumption. This can be shown by adding changes to Game_{11} that are almost the same as those we introduced from Game_2 to Game_6 . The only difference is that $c_{b_{\text{ID}^*}}$ is always sampled as $c_{b_{\text{ID}^*}} \leftarrow \mathbb{Z}_q$ here. By the similar analysis, there exists \mathcal{B} such that we have $|\Pr[X_{11}] - \Pr[X_{12}]| \leq n \cdot \text{Adv}_{\mathcal{B}'}^{\text{LWE}_{\ell, m, q, D_{\mathbb{Z}, \alpha q}}}(\lambda) + (Q_{\text{H}} + Q_{\text{ID}} + Q_{\text{ch}}) \cdot 2^{-n}$.

Finally, we observe that the challenge ciphertexts are sampled from $\text{CTSAm}(\text{mpk})$ regardless of whether $\text{coin} = 0$ or 1 . Therefore, we have $\Pr[X_{12}] = 1/2$. Putting things together, the theorem readily follows. \square

4.5 Security Proof in QROM

As we explained in the introduction, our analysis in the ROM can be easily extended to the QROM setting. Specifically, we prove the following theorem that addresses the security of GPV_{mult} in the QROM.

Theorem 4. *The IBE scheme GPV_{mult} is adaptively-anonymous multi-challenge secure assuming the hardness of $\text{LWE}_{\ell, m, q, \chi}$ in the quantum random oracle model, where $\chi = D_{\mathbb{Z}, \alpha q}$. Namely, for any quantum polynomial time adversary \mathcal{A} making at most Q_{H} quantum random oracle queries, Q_{ch} challenge queries, and Q_{ID} secret key queries, there exists a quantum polynomial time algorithm \mathcal{B} making at most $3Q_{\text{H}} + 2Q_{\text{ID}} + 6Q_{\text{ch}}$ quantum random oracle queries such that*

$$\text{Adv}_{\mathcal{A}, \text{GPV}_{\text{mult}}}^{\text{IBE}}(\lambda) \leq 3n \cdot \text{Adv}_{\mathcal{B}, \text{QRO}_{\ell_{\text{ID}}+2, \max\{\ell_r, (\lfloor \log q \rfloor + n) \times n\}}}^{\text{LWE}_{\ell, m, q, D_{\mathbb{Z}, \alpha q}}}(\lambda) + (Q_{\text{H}} + Q_{\text{ID}} + Q_{\text{ch}}) \cdot 2^{-\Omega(n)}$$

and

$$\text{Time}(\mathcal{B}) = \text{Time}(\mathcal{A}) + (Q_{\text{H}} + Q_{\text{ID}} + Q_{\text{ch}}) \cdot \text{poly}(\lambda)$$

where ℓ_r denotes the length of the randomness for $\text{Sample}\mathbb{Z}$.

Proof of Theorem 4. Let $\text{CTSAm}(\text{mpk})$ be an algorithm that outputs (c_0, c_1, \mathbf{c}_2) such that $c_0 \leftarrow \mathbb{Z}_q$, $c_1 \leftarrow \mathbb{Z}_q$, and $\mathbf{c}_2 = \mathbf{A}^\top \mathbf{s} + \mathbf{x}$ for $\mathbf{s} \leftarrow U([-\gamma, \gamma])$ and $\mathbf{x} \leftarrow D_{\mathbb{Z}^m, \alpha q}$. Let also \mathcal{A} be a quantum adversary that attacks the (multi-challenge) adaptively-anonymous security of the IBE scheme. Without loss of generality, we can assume that \mathcal{A} makes secret key queries on the same identity at most once.

Game₀ : This is the real security game. At the beginning of the game, the challenger chooses a random function $\mathbf{H} \leftarrow \text{Func}(\{0, 1\}^{\ell_{\text{ID}}+1}, \mathbb{Z}_q^n)$, which is used to simulate the random oracle. The challenger generates $(\mathbf{A}, \mathbf{T}_{\mathbf{A}}) \leftarrow \text{TrapGen}(1^n, 1^m, q)$ and the adversary \mathcal{A} is given \mathbf{A} . The challenger then samples $\text{coin} \leftarrow \{0, 1\}$ and keeps it secret. During the game, \mathcal{A} may make (quantum) random oracle queries, secret key queries, and challenge queries. These queries are handled as follows:

- When \mathcal{A} makes a random oracle query on a quantum state $\sum_{\text{ID}, b, y} \alpha_{\text{ID}, b, y} |\text{ID}\|b\rangle |y\rangle$, then the challenger returns $\sum_{\text{ID}, b, y} \alpha_{\text{ID}, b, y} |\text{ID}\|b\rangle |\mathbf{H}(\text{ID}\|b) \oplus y\rangle$.
- When \mathcal{A} makes a secret key query on ID , the challenger first chooses $b_{\text{ID}} \leftarrow \{0, 1\}$, computes $\mathbf{u}_{\text{ID}\|b_{\text{ID}}} = \mathbf{H}(\text{ID}\|b_{\text{ID}})$, and returns $\mathbf{e}_{\text{ID}\|b_{\text{ID}}} = \text{SamplePre}(\mathbf{A}, \mathbf{T}_{\mathbf{A}}, \mathbf{u}_{\text{ID}\|b_{\text{ID}}}, \sigma)$ to \mathcal{A} .
- When \mathcal{A} makes the challenge query for ID^* and a message M^* , the challenger returns $(c_0, c_1, \mathbf{c}_2) \leftarrow \text{Encrypt}(\text{mpk}, \text{ID}^*, M^*)$ if $\text{coin} = 0$ and $(c_0, c_1, \mathbf{c}_2) \leftarrow \text{CTSAm}(\text{mpk})$ if $\text{coin} = 1$. It then returns $C^* = (c_0, c_1, \mathbf{c}_2)$ to \mathcal{A} .

At the end of the game, \mathcal{A} outputs a guess $\widehat{\text{coin}}$ for coin . Finally, the challenger outputs $\widehat{\text{coin}}$. By definition, we have $|\Pr[X_0] - \frac{1}{2}| = |\Pr[\widehat{\text{coin}} - \text{coin}] - \frac{1}{2}| = \text{Adv}_{\mathcal{A}, \text{GPV}_{\text{mult}}}^{\text{IBE}}(\lambda)$.

Game₁ : In this game, we change the way b_{ID} is chosen in simulations of secret key queries. Namely, in this game, the challenger first picks a random function $H' \leftarrow \text{Func}(\{0, 1\}^{\ell_{\text{ID}}}, \{0, 1\})$. When \mathcal{A} makes a secret key query ID , the challenger uses $b_{\text{ID}} = H'(\text{ID})$ instead of randomly choosing it. We remark that an oracle access to H' is *not* given to \mathcal{A} . Since $H'(\text{ID})$ is an independently and uniformly random bit for all ID , we have $\Pr[X_0] = \Pr[X_1]$.

Game₂ : In this game, we change the way the random oracle H is simulated. Namely, the challenger first chooses additional random functions $\widehat{H}_0 \leftarrow \text{Func}(\{0, 1\}^{\ell_{\text{ID}}}, \{0, 1\}^{\ell_r})$ and $\widehat{H}_1 \leftarrow \text{Func}(\{0, 1\}^{\ell_{\text{ID}}}, \{0, 1\}^{(\lceil \log q \rceil + n) \times n})$. Let $\iota : \{0, 1\}^{(\lceil \log q \rceil + n) \times n} \rightarrow \mathbb{Z}_q^n$ denote a natural embedding function. More precisely, it is given $(a_1, \dots, a_n) \in \{0, 1\}^{(\lceil \log q \rceil + n) \times n}$ as an input, and outputs $(\tilde{a}_1 \bmod q, \dots, \tilde{a}_n \bmod q)^T$ where \tilde{a}_i denotes a positive integer whose binary representation is a_i . It is easy to see that if we sample $(a_1, \dots, a_n) \leftarrow \{0, 1\}^{(\lceil \log q \rceil + n) \times n}$, then the statistical distance between the distribution of $\iota(a_1, \dots, a_n)$ and the uniform distribution over \mathbb{Z}_q^n is $2^{-\Omega(n)}$. Then the challenger defines H as follows:

$$H(\text{ID} \| b) := \begin{cases} \mathbf{A} \mathbf{e}_{\text{ID} \| b} & \text{If } b = H'(\text{ID}) \\ \iota(\widehat{H}_1(\text{ID})) & \text{Otherwise} \end{cases} \quad (4)$$

where $\mathbf{e}_{\text{ID} \| b} := \text{SampleZ}(\sigma; \widehat{H}_0(\text{ID}))$. We remark that oracle accesses to \widehat{H}_0 and \widehat{H}_1 are *not* given to \mathcal{A} . For any fixed ID , for the case of $b \neq H'(\text{ID})$, the distribution of $H(\text{ID} \| b)$ is identical for all ID and its statistical distance from the uniform distribution is $2^{-\Omega(n)}$ as remarked above. For the case of $b = H'(\text{ID})$, the distribution of $H(\text{ID} \| b)$ is identical for all ID and its statistical distance from the uniform distribution is $2^{-\Omega(n)}$ for all but $2^{-\Omega(n)}$ fraction of \mathbf{A} due to Lem. 4 since we choose $\sigma > \sqrt{n + \log m}$. Then due to Lem. 2, we have $|\Pr[X_1] - \Pr[X_2]| = 2^{-\Omega(n)} + 4Q_{\text{H}}^2 \sqrt{2^{-\Omega(n)}} + 4Q_{\text{H}}^2 \sqrt{2^{-\Omega(n)}} = Q_{\text{H}}^2 \cdot 2^{-\Omega(n)}$.

Game₃ : In this game, we change the way secret key queries are answered. By the end of this game, the challenger will no longer require the trapdoor $\mathbf{T}_{\mathbf{A}}$ to generate the secret keys. When \mathcal{A} queries a secret key for ID , the challenger returns $(b_{\text{ID}}, \mathbf{e}_{\text{ID}})$ where $b_{\text{ID}} = H'(\text{ID})$ and $\mathbf{e}_{\text{ID} \| b_{\text{ID}}} := \text{SampleZ}(\sigma; \widehat{H}_0(\text{ID}))$. Since b_{ID} is unchanged from the previous game, we only have to prove that the distribution of $\mathbf{e}_{\text{ID} \| b_{\text{ID}}}$ differs negligibly from that in the previous game. For any fixed $\mathbf{u}_{\text{ID} \| b_{\text{ID}}} \in \mathbb{Z}_q^n$, let $\mathbf{e}_{\text{ID} \| b_{\text{ID}}}^{(2)}$ and $\mathbf{e}_{\text{ID} \| b_{\text{ID}}}^{(3)}$ be random variables that are distributed according to the distributions of $\mathbf{e}_{\text{ID} \| b_{\text{ID}}}$ conditioning on $H(\text{ID} \| b_{\text{ID}}) = \mathbf{u}_{\text{ID} \| b_{\text{ID}}}$ in **Game₂** and **Game₃**, respectively. Due to Lem. 8, we have $\Delta(\mathbf{e}_{\text{ID} \| b_{\text{ID}}}^{(2)}, D_{\Lambda_{\mathbf{u}_{\text{ID} \| b_{\text{ID}}}}^{\perp}(\mathbf{A}), \sigma}) \leq 2^{-\Omega(n)}$. On the other hand, due to Lem. 4, we have $\Delta(\mathbf{e}_{\text{ID} \| b_{\text{ID}}}^{(3)}, D_{\Lambda_{\mathbf{u}_{\text{ID} \| b_{\text{ID}}}}^{\perp}(\mathbf{A}), \sigma}) \leq 2^{-\Omega(n)}$. Since \mathcal{A} obtains at most Q_{ID} user secret keys \mathbf{e}_{ID} , we have $|\Pr[X_2] - \Pr[X_3]| = Q_{\text{ID}} \cdot 2^{-\Omega(n)}$.

Game₄ : In this game, we change the way the matrix \mathbf{A} is generated. Concretely, the challenger chooses $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$ without generating the associated trapdoor $\mathbf{T}_{\mathbf{A}}$. By Lem. 8, the distribution of \mathbf{A} differs at most by $2^{-\Omega(n)}$. Since the challenger can answer all the secret key queries without the trapdoor due to the change we made in the previous game, the view of \mathcal{A} is altered only by $2^{-\Omega(n)}$. Therefore, we have $|\Pr[X_3] - \Pr[X_4]| = 2^{-\Omega(n)}$.

Game₅ : In this game, we change \mathbf{A} to the lossy mode (See Lem. 15). We claim that $|\Pr[X_4] - \Pr[X_5]|$ is negligible assuming the hardness of the $\text{LWE}_{\ell, n, q, D_{\mathbb{Z}}, \alpha, q}$ problem relative to a quantum

random oracle $|\tilde{\mathbf{H}}\rangle : \{0, 1\}^{\ell_{\text{ID}}+2} \rightarrow \{0, 1\}^{\max\{\ell_r, (\lfloor \log q \rfloor + n) \times n\}}$. For this purpose, we construct an algorithm \mathcal{B}' that distinguishes uniform \mathbf{A} and that in the lossy mode relative to $|\tilde{\mathbf{H}}\rangle$. As remarked in Sec. 2.1, we can implement three independent quantum random oracles $|\mathbf{H}'\rangle : \{0, 1\}^{\ell_{\text{ID}}} \rightarrow \{0, 1\}$, $|\hat{\mathbf{H}}_0\rangle : \{0, 1\}^{\ell_{\text{ID}}} \rightarrow \{0, 1\}^{\ell_r}$, and $|\hat{\mathbf{H}}_1\rangle : \{0, 1\}^{\ell_{\text{ID}}} \rightarrow \{0, 1\}^{(\lfloor \log q \rfloor + n) \times n}$ by using $|\tilde{\mathbf{H}}\rangle$. Therefore we assume that \mathcal{B}' can access to three random oracles $|\mathbf{H}'\rangle$, $|\hat{\mathbf{H}}_0\rangle$ and $|\hat{\mathbf{H}}_1\rangle$.

\mathcal{B}' is given a matrix \mathbf{A} , and its task is to distinguish whether $\mathbf{A} \leftarrow \mathbb{Z}_q^n$ or $\mathbf{A} \leftarrow \text{SampleLossy}(n, m, \ell, \chi)$. First, we remark that \mathcal{B}' can simulate the quantum random oracle $|\mathbf{H}\rangle$ by using its own random oracles $|\mathbf{H}'\rangle$, $|\hat{\mathbf{H}}_0\rangle$ and $|\hat{\mathbf{H}}_1\rangle$ because $|\mathbf{H}\rangle$ is programmed based on these three oracles by the modification made in Game_2 . \mathcal{B}' sets the master public key to be the LWE matrix \mathbf{A} . Note that unlike the real IBE scheme, \mathcal{B}' does not require the master secret key $\mathbf{T}_{\mathbf{A}}$ due to the modification we made in Game_4 . Namely, when \mathcal{A} queries ID for the secret key oracle, \mathcal{B}' just returns $(b_{\text{ID}}, \mathbf{e}_{\text{ID}})_{b_{\text{ID}}}$ where $b_{\text{ID}} = \mathbf{H}'(\text{ID})$ and $\mathbf{e}_{\text{ID}} := \text{SampleZ}(\sigma; \hat{\mathbf{H}}_0(\text{ID}))$. When \mathcal{A} makes a challenge query for $(\text{ID}^*, \mathbf{M}^*)$, it samples a ciphertext as $C^* \leftarrow \text{CTSAm}(\text{mpk})$ if $\text{coin} = 1$, and honestly generates a ciphertext as $C^* \leftarrow \text{Encrypt}(\text{mpk}, \text{ID}^*, \mathbf{M}^*)$ if $\text{coin} = 0$, and returns C^* to \mathcal{A} . At the end of the game, \mathcal{A} outputs $\widehat{\text{coin}}$. Finally, \mathcal{B} outputs 1 if $\widehat{\text{coin}} = \text{coin}$ and 0 otherwise.

It can be seen that if $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$ the view of the adversary corresponds to Game_5 , and if $\mathbf{A} \leftarrow \text{SampleLossy}(n, m, \ell, \chi)$ the view of the adversary corresponds to Game_6 . Therefore we can bound $|\Pr[X_5] - \Pr[X_6]|$ by the distinguishing advantage of \mathcal{B}' , which makes at most $3Q_{\text{H}} + 2Q_{\text{ID}} + 6Q_{\text{ch}}$ quantum random oracle queries. As for the running time, we have $\text{Time}(\mathcal{B}') = \text{Time}(\mathcal{A}) + (Q_{\text{H}} + Q_{\text{ID}} + Q_{\text{ch}}) \cdot \text{poly}(\lambda)$ since all \mathcal{B}' has to do is to run \mathcal{A} once plus to compute some additional computations that can be done in a fixed polynomial time whenever \mathcal{A} makes any query. Then by Lem. 14 we conclude that there exists \mathcal{B} such that we have $|\Pr[X_4] - \Pr[X_5]| = n \cdot \text{Adv}_{\mathcal{B}, \text{QRO}}^{\text{LWE}_{\ell, m, q, D_{\mathbb{Z}}, \alpha q}}(\lambda)$ and the running time of \mathcal{B} is almost the same as that of \mathcal{B}' .

Game_6 : In this game, we change the way the challenge ciphertexts are generated. When \mathcal{A} makes a challenge query for $(\text{ID}^*, \mathbf{M}^*)$, it samples a ciphertext as $C^* \leftarrow \text{CTSAm}(\text{mpk})$ and returns C^* to \mathcal{A} if $\text{coin} = 1$. If $\text{coin} = 0$, it generates the ciphertext as

$$c_{b_{\text{ID}^*}} = \mathbf{u}_{\text{ID}^*}^{\top} \mathbf{s} + \mathbf{M}^* \lfloor q/2 \rfloor, \quad c_{\bar{b}_{\text{ID}^*}} \leftarrow \mathbb{Z}_q, \quad \mathbf{c}_2 = \mathbf{A}^{\top} \mathbf{s} + \mathbf{x}.$$

for $\mathbf{s} \leftarrow U([-\gamma, \gamma])$, where $b_{\text{ID}^*} = \mathbf{H}'(\text{ID}^*)$. It then rearranges the terms if necessary and returns (c_0, c_1, \mathbf{c}_2) to \mathcal{A} .

We argue that the view of \mathcal{A} is statistically close to that in the previous game. To prove this, we make a hybrid argument over all the challenge ciphertexts and change $c_{\bar{b}_{\text{ID}^*}}$ to be random one-by-one (when $\text{coin} = 0$). To conclude, it suffices to show that each ciphertext $c_{\bar{b}_{\text{ID}^*}}$ is distributed $2^{-\Omega(n)}$ close to the uniform distribution over \mathbb{Z}_q . Specifically, for each ciphertext, for any $\epsilon' = 2^{-\lambda}$, we have

$$\begin{aligned} \tilde{\mathbf{H}}_{\infty}^{\epsilon'}(\mathbf{s} | \mathbf{A}, c_{b_{\text{ID}^*}}, \mathbf{c}_2) &\geq \tilde{\mathbf{H}}_{\infty}^{\epsilon'}(\mathbf{s} | \mathbf{A}, \mathbf{c}_2) - \log q \\ &\geq \mathbf{H}_{\infty}(\mathbf{s}) - (\ell + 2\lambda + 1) \log q - \text{negl}(\lambda) \\ &= n \log(2\gamma) - (\ell + 3\lambda) \log q \\ &\geq \log q + \Omega(n) \end{aligned}$$

where the first inequality follows by applying Lem. 12 with $\mathcal{Z} = \mathbb{Z}_q$ and $\epsilon = 0$, the second inequality follows from Lem. 15, and the last inequality follows from our parameter choice. This implies that $\Delta(\mathbf{u}_{\text{ID}^*}^{\top} \mathbf{s}, \mathbf{u}_{\text{ID}^*}^{\top} \mathbf{s}, \mathbf{A}, c_{b_{\text{ID}^*}}, \mathbf{c}_2) \leq 2^{-\Omega(n)}$ for $\mathbf{u}_{\text{ID}^*} \leftarrow \mathbb{Z}_q^n$ by Lem. 13 together with Fact 1. Therefore, we have $|\Pr[X_5] - \Pr[X_6]| = Q_{\text{ch}} \cdot 2^{-\Omega(n)}$.

From Game_7 to Game_{10} in the following, we undo the changes we added from Game_2 to Game_5 .

Game_7 : In this game, \mathbf{A} is sampled as $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$. Similarly to the change from Game_4 to Game_5 , there exists \mathcal{B}' such that we have $|\Pr[X_6] - \Pr[X_7]| = n \cdot \text{Adv}_{\mathcal{B}', \text{QRO}_{\ell_{\text{ID}}+2, \max\{\ell_r, (\lfloor \log q \rfloor + n) \times n\}}}^{\text{LWE}_{\ell, m, q, D_{\mathbb{Z}, \alpha q}}}(\lambda)$.

Game_8 : In this game, \mathbf{A} is sampled with a trapdoor as $(\mathbf{A}, \mathbf{T}_{\mathbf{A}}) \leftarrow \text{TrapGen}(1^n, 1^m, q)$. Similarly to the change from Game_3 to Game_4 , we have $|\Pr[X_7] - \Pr[X_8]| = 2^{-\Omega(n)}$.

Game_9 : In this game, we change the way secret key queries are answered. When \mathcal{A} makes a secret key query for ID , the challenger returns $\mathbf{e}'_{\text{ID} \parallel b_{\text{ID}}} \leftarrow \text{SamplePre}(\mathbf{A}, \mathbf{T}_{\mathbf{A}}, \mathbf{u}_{\text{ID} \parallel b_{\text{ID}}}, \sigma)$, where $b_{\text{ID}} = \text{H}'(\text{ID})$, $\mathbf{e}_{\text{ID} \parallel b_{\text{ID}}} = \text{SampleZ}(\sigma; \widehat{\text{H}}_0(\text{ID}))$ and $\mathbf{u}_{\text{ID}} = \mathbf{A} \mathbf{e}_{\text{ID} \parallel b_{\text{ID}}}$. Similarly to the change from Game_2 to Game_3 , we have $|\Pr[X_8] - \Pr[X_9]| = Q_{\text{H}} \cdot 2^{-\Omega(n)}$.

Game_{10} : In this game, we change the way the random oracle queries to H are answered. Namely, the challenger simply uses a random function $\text{H} : \{0, 1\}^{\ell_{\text{ID}}+1} \rightarrow \mathbb{Z}_q^n$ to simulate $|\text{H}\rangle$ instead of programming it as in Eq. (4). Similarly to the change from Game_1 to Game_2 , we have $|\Pr[X_{10}] - \Pr[X_9]| = Q_{\text{H}} \cdot 2^{-\Omega(n)}$.

Because of the changes we introduced in Game_7 to Game_{10} , we can add the following change.

Game_{11} : In this game, we change the way the challenge ciphertexts are generated. When \mathcal{A} makes a challenge query for $(\text{ID}^*, \text{M}^*)$, it returns a ciphertext sampled from $\text{CTSAm}(\text{mpk})$ if $\text{coin} = 1$. If $\text{coin} = 0$, it generates the ciphertext as

$$c_{b_{\text{ID}^*}} \leftarrow \mathbb{Z}_q, \quad c_{\bar{b}_{\text{ID}^*}} = \mathbf{u}_{\text{ID}^* \parallel b_{\text{ID}^*}}^\top \mathbf{s} + \text{M} \lfloor q/2 \rfloor, \quad \mathbf{c}_2 = \mathbf{A}^\top \mathbf{s} + \mathbf{x}.$$

for $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, where $b_{\text{ID}^*} = \text{H}'(\text{ID}^*)$. It then rearranges the terms if necessary and returns (c_0, c_1, \mathbf{c}_2) to \mathcal{A} . We claim that this change is only conceptual. Note that the distribution of the ciphertexts in this game corresponds to that in the previous game if we flip the value of $\text{H}'(\text{ID}^*)$ for all challenge identities ID^* . Since \mathcal{A} never makes a secret key query for ID^* , the value of $\text{H}'(\text{ID}^*)$ is information theoretically hidden from \mathcal{A} . This implies that even if we flip values of $\text{H}'(\text{ID}^*)$ for all challenge identities ID^* , \mathcal{A} cannot notice it at all. Therefore, we have $\Pr[X_{11}] = \Pr[X_{10}]$.

Game_{12} : In this game, we further change the way the challenge ciphertexts are generated. When \mathcal{A} makes a challenge query for $(\text{ID}^*, \text{M}^*)$, it returns a random ciphertext sampled from $\text{CTSAm}(\text{mpk})$ if $\text{coin} = 1$. If $\text{coin} = 0$, it generates the ciphertext as

$$c_{b_{\text{ID}^*}} \leftarrow \mathbb{Z}_q, \quad c_{\bar{b}_{\text{ID}^*}} \leftarrow \mathbb{Z}_q, \quad \mathbf{c}_2 = \mathbf{A}^\top \mathbf{s} + \mathbf{x}.$$

for $\mathbf{s} \leftarrow \mathbb{Z}_q^n$. It then returns (c_0, c_1, \mathbf{c}_2) to \mathcal{A} . We claim that the change is unnoticed by \mathcal{A} assuming the LWE assumption. This can be shown by adding changes to Game_{11} that are almost the same as those we introduced from Game_2 to Game_6 . The only difference is that $c_{b_{\text{ID}^*}}$ is always sampled as $c_{b_{\text{ID}^*}} \leftarrow \mathbb{Z}_q$ here. By similar analysis, there exists \mathcal{B} such that we have $|\Pr[X_{11}] - \Pr[X_{12}]| \leq n \cdot \text{Adv}_{\mathcal{B}, \text{QRO}_{\ell_{\text{ID}}+2, \max\{\ell_r, (\lfloor \log q \rfloor + n) \times n\}}}^{\text{LWE}_{\ell, m, q, D_{\mathbb{Z}, \alpha q}}}(\lambda) + (Q_{\text{H}} + Q_{\text{ID}} + Q_{\text{ch}}) \cdot 2^{-n}$.

Finally, we observe that the challenge ciphertexts are sampled from $\text{CTSAm}(\text{mpk})$ regardless of whether $\text{coin} = 0$ or 1. Therefore, we have $\Pr[X_{12}] = 1/2$. Putting things together, the theorem readily follows. \square

Acknowledgement. The first author was partially supported by JST CREST Grant Number JPMJCR1302 and JSPS KAKENHI Grant Number 17J05603. The second author was supported by JST CREST Grant No. JPMJCR1688 and JSPS KAKENHI Grant Number 16K16068.

References

- [ABB10] Shweta Agrawal, Dan Boneh, and Xavier Boyen. Efficient lattice (h) ible in the standard model. In *EUROCRYPT*, pages 553–572. Springer, 2010. [6](#)
- [ABB⁺17] Erdem Alkim, Nina Bindel, Johannes A. Buchmann, Özgür Dagdelen, Edward Eaton, Gus Gutoski, Juliane Krämer, and Filip Pawlega. Revisiting TESLA in the quantum random oracle model. In *PQCrypto*, pages 143–162. Springer, 2017. [10](#)
- [ACPS09] Benny Applebaum, David Cash, Chris Peikert, and Amit Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In *CRYPTO*, pages 595–618. Springer, 2009. [8](#)
- [ADN⁺10] Joël Alwen, Yevgeniy Dodis, Moni Naor, Gil Segev, Shabsi Walfish, and Daniel Wichs. Public-key encryption in the bounded-retrieval model. In *EUROCRYPT*, pages 113–134. Springer, 2010. [8](#)
- [AHY15] Nuttapon Attrapadung, Goichiro Hanaoka, and Shota Yamada. A framework for identity-based encryption with almost tight security. In *ASIACRYPT*, pages 521–549. Springer, 2015. [10](#)
- [AKPW13] Joël Alwen, Stephan Krenn, Krzysztof Pietrzak, and Daniel Wichs. Learning with rounding, revisited. In *CRYPTO*, pages 57–74. Springer, 2013. [7](#), [24](#), [25](#), [26](#)
- [ARU14] Andris Ambainis, Ansis Rosmanis, and Dominique Unruh. Quantum attacks on classical proof systems: The hardness of quantum rewinding. In *FOCS*, pages 474–483. IEEE, 2014. [10](#)
- [BDF⁺11] Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. Random oracles in a quantum world. In *Eurocrypt*, pages 41–69. Springer, 2011. [3](#), [9](#), [11](#), [15](#)
- [BDPMW16] Florian Bourse, Rafaël Del Pino, Michele Minelli, and Hoeteck Wee. The circuit privacy almost for free. In *CRYPTO*, pages 62–89. Springer, 2016. [8](#)
- [BF01] Dan Boneh and Matt Franklin. Identity-based encryption from the weil pairing. In *CRYPTO*, pages 213–229. Springer, 2001. [10](#), [12](#)
- [BG14] Shi Bai and Steven D. Galbraith. An improved compression technique for signatures based on learning with errors. In *CT-RSA*, pages 28–47, 2014. [10](#)
- [BKPW12] Mihir Bellare, Eike Kiltz, Chris Peikert, and Brent Waters. Identity-based (lossy) trapdoor functions and applications. In *EUROCRYPT*, pages 228–245. Springer, 2012. [7](#), [25](#)
- [BL16] Xavier Boyen and QinYi Li. Towards tightly secure lattice short signature and id-based encryption. In *ASIACRYPT(2)*, pages 404–434. Springer, 2016. [10](#)
- [BL18] Xavier Boyen and Qinyi Li. Almost tight multi-instance multi-ciphertext identity-based encryption on lattices. In *ACNS*, pages 535–553. Springer, 2018. [10](#)

- [BLP⁺13] Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of learning with errors. In *STOC*, pages 575–584, 2013. [14](#), [15](#)
- [BPR12] Abhishek Banerjee, Chris Peikert, and Alon Rosen. Pseudorandom functions and lattices. In *EUROCRYPT*, pages 719–737. Springer, 2012. [12](#)
- [BR93] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *CCS*, pages 62–73. ACM, 1993. [3](#), [9](#)
- [CHK04] Ran Canetti, Shai Halevi, and Jonathan Katz. Chosen-ciphertext security from identity-based encryption. In *EUROCRYPT*, pages 207–222. Springer, 2004. [9](#)
- [CHKP10] David Cash, Dennis Hofheinz, Eike Kiltz, and Chris Peikert. Bonsai trees, or how to delegate a lattice basis. In *EUROCRYPT*, pages 523–552. Springer, 2010. [6](#)
- [Cor09] Jean-Sébastien Coron. A variant of Boneh-Franklin IBE with a tight reduction in the random oracle model. *Designs, Codes and Cryptography*, 50(1):115–133, 2009. [10](#)
- [CS98] Ronald Cramer and Victor Shoup. A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In *CRYPTO*, pages 13–25. Springer, 1998. [6](#)
- [CW13] Jie Chen and Hoeteck Wee. Fully, (almost) tightly secure IBE and dual system groups. In *CRYPTO*, pages 435–460. Springer, 2013. [10](#)
- [DORS04] Yevgeniy Dodis, Rafail Ostrovsky, Leonid Reyzin, and Adam Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. In *EUROCRYPT*, pages 523–540. Springer, 2004. [24](#), [25](#)
- [GDCC16] Junqing Gong, Xiaolei Dong, Jie Chen, and Zhenfu Cao. Efficient IBE with tight reduction to standard assumption in the multi-challenge setting. In *ASIACRYPT*, pages 624–654. Springer, 2016. [10](#)
- [Gen06] Craig Gentry. Practical identity-based encryption without random oracles. In *EUROCRYPT*, pages 445–464. Springer, 2006. [6](#)
- [GGM86] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *J. ACM*, 33(4):792–807, 1986. [12](#)
- [GKPV10] Shafi Goldwasser, Yael Kalai, Chris Peikert, and Vinod Vaikuntanathan. Robustness of the learning with errors assumption. *ICS*, pages 230–240, 2010. [7](#), [25](#)
- [Gol86] Oded Goldreich. Two remarks concerning the Goldwasser-Micali-Rivest signature scheme. In *CRYPTO*, pages 104–110. Springer, 1986. [13](#)
- [GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pages 197–206. ACM, 2008. [3](#), [4](#), [5](#), [8](#), [9](#), [14](#), [16](#), [18](#)
- [GSW13] Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In *CRYPTO*, pages 75–92. Springer, 2013. [8](#)

- [HJ12] Dennis Hofheinz and Tibor Jager. Tightly secure signatures and public-key encryption. In *CRYPTO*, pages 590–607. Springer, 2012. [4](#), [7](#)
- [HKS15] Dennis Hofheinz, Jessica Koch, and Christoph Striecks. Identity-based encryption with (almost) tight security in the multi-instance, multi-ciphertext setting. In *PKC*, pages 799–822. Springer, 2015. [10](#)
- [JZC⁺18] Haodong Jiang, Zhenfeng Zhang, Long Chen, Hong Wang, and Zhi Ma. Ind-cca-secure key encapsulation mechanism in the quantum random oracle model, revisited. In *CRYPTO*, pages 96–125. Springer, 2018. [10](#)
- [KLS18] Eike Kiltz, Vadim Lyubashevsky, and Christian Schaffner. A concrete treatment of fiat-shamir signatures in the quantum random-oracle model. In *EUROCRYPT III*, pages 552–586. Springer, 2018. [10](#)
- [KW03] Jonathan Katz and Nan Wang. Efficiency improvements for signature schemes with tight security reductions. In *Computer and Communications Security*, pages 155–164. ACM, 2003. [8](#), [10](#), [24](#)
- [KY16] Shuichi Katsumata and Shota Yamada. Partitioning via non-linear polynomial functions: more compact ibes from ideal lattices and bilinear maps. In *ASIACRYPT*, pages 682–712. Springer, 2016. [6](#), [9](#), [14](#)
- [LSSS17] Benoît Libert, Amin Sakzad, Damien Stehlé, and Ron Steinfeld. All-but-many lossy trapdoor functions and selective opening chosen-ciphertext security from lwe. In *CRYPTO*, pages 332–364. Springer, 2017. [7](#), [25](#)
- [MP12] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *EUROCRYPT*, pages 700–718. Springer, 2012. [14](#)
- [MR07] Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on gaussian measures. *SIAM J. Comput.*, 37(1):267–302, 2007. [14](#)
- [NC00] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000. [10](#)
- [Pei07] Chris Peikert. Limits on the hardness of lattice problems in ℓ_p norms. In *Conference on Computational Complexity*, pages 333–346. IEEE, 2007. [14](#)
- [Pei09] Chris Peikert. Public-key cryptosystems from the worst-case shortest vector problem. In *STOC*, pages 333–342. ACM, 2009. [15](#)
- [Pei10] Chris Peikert. An efficient and parallel gaussian sampler for lattices. In *CRYPTO*, pages 80–97. Springer, 2010. [14](#)
- [PR06] Chris Peikert and Alon Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In *TCC*, pages 145–166. Springer, 2006. [14](#)
- [Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In *STOC*, pages 84–93. ACM Press, 2005. [3](#), [7](#), [14](#), [15](#)
- [Reg10] Oded Regev. The learning with errors problem. *Invited survey in CCC*, 2010. [7](#)

- [RW04] Renato Renner and Stefan Wolf. Smooth rényi entropy and applications. In *International Symposium on Information Theory – ISIT*, pages 233–233. IEEE, 2004. [24](#)
- [Sho94] Peter W. Shor. Algorithms for quantum computation: Discrete logarithms and factoring. In *FOCS*, pages 124–134. IEEE, 1994. [3](#)
- [SXY18] Tsunekazu Saito, Keita Xagawa, and Takashi Yamakawa. Tightly-secure key-encapsulation mechanism in the quantum random oracle model. In *EUROCRYPT III*, pages 520–551. Springer, 2018. [9](#), [10](#), [12](#)
- [TU16] Ehsan Ebrahimi Targhi and Dominique Unruh. Post-quantum security of the fujisaki-okamoto and OAEP transforms. In *TCC II*, pages 192–216. Springer, 2016. [10](#)
- [Unr14a] Dominique Unruh. Quantum position verification in the random oracle model. In *CRYPTO II*, pages 1–18. Springer, 2014. [10](#)
- [Unr14b] Dominique Unruh. Revocable quantum timed-release encryption. In *EUROCRYPT*, pages 129–146. Springer, 2014. [10](#)
- [Unr15] Dominique Unruh. Non-interactive zero-knowledge proofs in the quantum random oracle model. In *EUROCRYPT II*, pages 755–784. Springer, 2015. [10](#)
- [Unr17] Dominique Unruh. Post-quantum security of fiat-shamir. In *ASIACRYPT I*, pages 65–95. Springer, 2017. [10](#)
- [Wat05] Brent Waters. Efficient identity-based encryption without random oracles. In *EUROCRYPT*, pages 114–127. Springer, 2005. [6](#)
- [Zha12a] Mark Zhandry. How to construct quantum random functions. In *FOCS*, pages 679–687. IEEE, 2012. [4](#), [9](#), [12](#)
- [Zha12b] Mark Zhandry. Secure identity-based encryption in the quantum random oracle model. In *CRYPTO*, pages 758–775. Springer, 2012. [3](#), [4](#), [5](#), [9](#), [10](#), [11](#), [21](#)