Towards Tightly Secure Short Signature and IBE

Xavier Boyen, Qinyi Li

Queensland University of Technology, Brisbane, Australia
xavier.boyen@qut.edu.au
qinyi.li@student.qut.edu.au

Abstract. Constructing short signatures with tight security from standard assumptions is a long-standing open problem. We present an adaptively secure, short (and stateless) signature scheme, featuring a constant security loss relative to a conservative hardness assumption, Short Integer Solution (SIS), and the security of a concretely instantiated pseudorandom function (PRF). This gives a class of tightly secure short lattice signature schemes whose security is based on SIS and the underlying assumption of the instantiated PRF.

Our signature construction further extends to give a class of tightly and adaptively secure "compact" Identity-Based Encryption (IBE) schemes, reducible with constant security loss from Regev's vanilla Learning With Errors (LWE) hardness assumption and the security of a concretely instantiated PRF. Our approach is a novel combination of a number of techniques, including Katz and Wang signature, Agrawal et al. lattice-based secure IBE, and Boneh et al. key-homomorphic encryption.

Our results, at the first time, eliminate the dependency between the number of adversary's queries and the security of short signature/IBE schemes in the context of lattice-based cryptography. They also indicate that tightly secure PRFs (with constant security loss) would imply tightly, adaptively secure short signature and IBE schemes (with constant security loss).

1 Introduction

Short signatures are useful and desirable for providing data authenticity in low-bandwidth and/or high-throughput applications where many signatures have to be processed very quickly. Most digital signature schemes are based on computationally hard problems on specific algebraic groups, e.g., finite fields, curves, and lattices. A signature is "short" if the signature consists in a (small) constant number of group elements (e.g., field elements or lattice points).

Although bare-bones signatures can be obtained from very weak assumptions (e.g., collision-resistant hash functions), constructing efficient short signatures satisfying standard security requirements (e.g., existential unforgeability under adaptively chosen-message attacks), from reasonable assumptions, appears to be a challenging task. Some of the existing short signature schemes use random oracles, e.g., [20,10,46,35,48], or rely on non-standard computational assumptions (strong, interactive assumptions, and/or q-type parametric assumptions), e.g., [33,29,32,16,25], or require signers to maintain state across signatures, e.g., [43].

The first short signature scheme from a reasonable and non-parametric assumption without random oracles was proposed by Waters [54]. Hohenberger and Waters later proposed a short signature scheme from standard RSA [44]. Lattice-based short signatures from the very mild SIS assumption in the standard model were proposed in [21,49]. Recently, the "confined guessing" technique developed by Böhl et al. [13] has produced short signatures from standard RSA and bilinear-group CDH assumptions, and also from the ring-SIS/SIS assumption in combination with lattice techniques [31,4] with very loose reductions.

Despite these elegant constructions, signature schemes that are *short* and enjoy *tight security* reductions to *standard assumptions* in the *standard model* (without random oracle), remain unknown. Existing tightly secure signature schemes either have large signature size, e.g., [41,1,11], or merely have heuristic security arguments based on random oracles, e.g., [46,38]. We have not been able to ascertain the earliest occurrence of this long-standing folklore problem in cryptography, but here [11] is one recent formulation:

Open Problem #1—**Tightly Secure Short Signatures**

> "Construct a tightly secure and short (in the sense that the signature contains constant number of group elements or vectors and the security loss is a constant) signature scheme from standard assumptions." — Blazy, Kakvi, Kiltz, Pan (2015)

1.1 Tight Security

The reductionist approach to cryptographic security algorithms seeks to prove theorems along the lines of: "If a t-time adversary attacks the scheme with successful probability ϵ, then a t'-time algorithm can be constructed to break some computational problem with success probability $\epsilon' = \epsilon/\theta$ and $t' = k \cdot t + o(t)$.". The parameters $\theta \geq 1$ and $k \geq 1$, or more simply the product $k \cdot \theta$, measures how tightly the security of the cryptographic scheme is related to the hardness of the underlying computational problem. Alternatively, when $k \approx 1$ as is the case in many reductions, θ measures the security loss of the security reduction of our cryptographic scheme from the underlying assumption. A cryptographic scheme is *tightly secure* if θ is a small constant that in particular does not depend on parameters under the adversary's control, such as the adversary's own success probability ϵ, the number of queries it chooses to make, and even the scheme's security parameter. The reduction phrases "almost tight security" from the literature refers to the case where θ only depends on a small polynomial of the security parameter.

Tight reduction is an elegant notion from a theoretical point of view. A tight reductionist proof (with respect to a well-defined security model) indicates that the security of a cryptographic scheme is (extremely) closely related to the hardness of the underlying hard problem, which is the optimal case we expect from provable security theory. On the other hand, it is also a determinant factor to the practicality of real-world security. Its opposite, loose security, means that in order to realise a desired "real" target security level, one has to increase the "apparent" security level inside the construction to compensate for the loose reduction. This inflates the size of data atoms by some polynomial, with in turn increases the running time of cryptographic operations by another polynomial, combining multiplicatively.

1.2 Identity-Based Encryption with Tight Security

Digital signatures and identity-based encryption (IBE) are closely connected, which suggests that techniques that improve upon the security of signatures might also improve upon the security of IBE. In this work, we also investigate the problem of constructing tightly secure IBE from standard assumptions (without random oracles).

In an IBE system, any random string that uniquely represents a user's identity, such as email address or driver license number, can act as a public key (within a certain domain or realm). Encryption uses this identity, together with some common domain-specific public parameters, to encrypt messages. Users are issued private decryption keys corresponding to their public identities, by a trusted authority (or distributed authorities) called Private Key Generator (PKG) which hold(s) (shares of) the master secret key for a domain. Decryption succeeds if the identity associated with the ciphertext matches the identity associated with the private key, in the same domain.

The strongest, most natural and most widely accepted notion of security for IBE is the *adaptive* security model or *full* security model, formally defined in [18]. In this model, the adversary is able to announce its target (the challenge identity it wants to attack) at any time during the course of its adaptive interaction with the system. Without the luxury of random oracles, an easier security model to achieve was the *selective* security model, where the adversary must announce its target identity at the onset of its interaction with the system.

In the last fifteen years, a great many IBE schemes have been proposed, with varying efficiency, security models, hardness assumptions, and other features. In the standard model (i.e., without random oracles or other idealised oracles), we mention several notable IBE schemes which have been constructed from bilinear maps in the selective model [26,14] and the adaptive model [15,54,34,55,28,12], and from lattices in the adaptive model [2,27,5]. It is fair to say that, by now, the art of selectively secure IBE has been well honed. However, adaptively secure IBE schemes from standard assumptions with tight security (in the sense that the security loss is a small constant) remain unknown. The best known adaptively secure IBE schemes in terms of tight reduction are based on linear assumptions over pairings and achieve almost tight security (e.g., [28,12,6,42]). Waters [54] states this open problem as follows:

Open Problem #2—**Tight Adaptively Secure IBE**

> "Construct a tightly, adaptively secure IBE scheme from standard computational hardness assumptions without random oracles." —Waters (2005)

Furthermore, for all known directly constructed adaptively secure IBE scheme from standard post-quantum assumption (specifically the LWE assumption), i.e. [2,27,5], their security loss during reduction depends on the number adversary's of queries. That is there is current no even "almost tightly" secure adaptive IBE scheme based on standard computational problems which are conjectured to be hard under quantum attacks. The following problem is still open.

Open Problem #3—**"Almost" Tight Adaptively Secure, Post-Quantum IBE**
"Construct an "almost" tightly, adaptively secure IBE scheme from standard post-quantum assumptions without random oracles."

1.3 Our Results

Our work uses pseudorandom functions (PRFs) that are efficiently computable by Boolean circuits with up to polynomial depth in their input length. Recall a PRF is a function: $\mathsf{PRF} : \mathcal{K} \times \mathcal{D} \to \mathcal{R}$ with the following security property. For random $K \overset{\$}{\leftarrow} \mathcal{K}$, $\mathsf{PRF}(K, \cdot)$ is computationally indistinguishable from a random function $\Omega : \mathcal{D} \to \mathcal{R}$, given oracle access to either $\mathsf{PRF}(K, \cdot)$ or Ω. PRFs can be constructed from general assumptions (e.g., the existence of pseudo-random number generators [39]), number-theoretic assumptions (e.g., the DDH/k-LIN assumption [51,30,45]), and lattice assumptions LWE [9,8].

Our contribution is a construction of a class of adaptively secure short signature schemes/IBE schemes in the standard model. The schemes' security is tightly related to SIS/LWE and the security of an instantiated PRF PRF in the sense that the security loss is a nearly optimal constant factor. More precisely, let ϵ and ϵ' be the advantage of an adversary in attacking our signature and IBE schemes respectively, ϵ_{SIS} and ϵ_{LWE} be the security level of the SIS and LWE assumptions on which our schemes are based, and ϵ_{PRF} is the security level of the PRF instantiation PRF. Our constructions provide the following: $\epsilon \approx 2(\epsilon_{\mathsf{SIS}} + \epsilon_{\mathsf{PRF}})$, $\epsilon' \approx 2(\epsilon_{\mathsf{LWE}} + \epsilon_{\mathsf{PRF}})$, and the (polynomial) runtime of reduction is approximately the same as attacker's runtime.

Note that, depending on the underlying hardness assumption and the reduction of PRF, underlying assumptions and tightness of our signature/IBE scheme vary. By instantiating existing lattice-based/number theoretic-based PRFs, we obtain the following improvements upon known results:

- By instantiating the low-depth LWE-based PRFs by Banerjee et al. [9] or Banerjee and Peikert [8], we obtain the first "almost" tightly secure short signature/IBE schemes from LWE (which is stronger than SIS) whose security does not depend on the number of adversarial queries. Previously, the known lattice signature schemes either enjoy short signatures but loose reduction (such as [21,49,31]) or have tight reduction but rather large signatures ([11]), and the known adaptively secure lattice-based IBE schemes ([2,5]) have loose reductions. This, at the first time, eliminates the dependency between the number of adversary's queries and the security of lattice-based *short* signature scheme/IBE scheme. It also allows us to answer the open problem Open Problem #3.

- If we relex the requirment of quantum recistance, by instantiating the (black-box) tightly secure PRF based on DDH or k-LIN, whose security loss is only $O(\log^2 \lambda)$ for security parameter λ, due to Jager [45], we obtain the IBE scheme with tightest security reduction so far: a factor of $O(\log^2 \lambda)$. Previous IBE schemes with almost tight security [28,12] have a factor of $O(\lambda)$ of security loss. This improvement brings us closer again to answering the Open Problem #1 and #2.

Meanwhile, an interesting and independent contribution of our work is that it indicates that tightly secure PRFs, which are efficiently computable by Boolean circuits (with depth from constant (i.e. TC^0) up to polynomial), from standard computational assumptions are sufficient for us to build tightly, adaptively secure lattice signature/IBE from SIS/LWE (whose average-case hardness is equivalent to classic worst-case lattice problems with approximation factors from polynomial value up to sub-exponential value).

Finally, we note that if we instantiate PRF with efficient PRFs within shallow circuits class TC^0 (e.g., the efficient LWE-based PRFs from [9] and DDH-based PRF from [51,30,45]), the parameter size of instantiated PRFs in our constructions will barely affect the asymptotic efficiency of the signature/IBE scheme at all. This means compensating the PRFs security by increasing the parameter size of the instantiated PRFs only incurs small amount of additional overhead to the original signature/IBE scheme. Although our constructions do not reach the final goal of getting constant security loss (because of lack of tightly secure PRFs currently), the tight reductions of our signature/IBE schemes do provide asymptotically optimal suggestions for parameter selection (under specific SIS/LWE assumptions). We discuss this in section 3.2.

Table 1 provides a comparison between our signature scheme with a LWE-based PRF instantiation (from [9]) and a representative sample of the prominent lattice-based (quantum-safe) signature schemes from the literature. Note, Katz and Wang did not propose a SIS-based signature scheme in [46]. The scheme we refer to is a straightforward application of Katz-Wang's proof technique to GPV'08 signature scheme. Table 2 provides a comparison between our signature scheme with DDH PRF instantiation from [45], which only looses a factor $O(\log^2 \lambda)$ in security proof, and the representative signature schemes from traditional number-theoretic assumptions, including (strong) RSA, Dlog and linear assumptions over pairings. All of those assumptions are not conjectured to be quantum-safe. In each case, the two tables refer to conjectured quantum safe and quantum-unsafe constructions respectively.

Table 3 gives a comparison between our IBE scheme (with both PRF instantiation from LWE [9] and DDH instantiation from [45]) and a representative selection of existing IBE schemes from the literature. We note the LWE-based PRF instantiation from [9] requires a somewhat large modulus which accounts for security assumption of PRF. (See the discussion in section 4.2.) But is does not impact the asymptotic efficiency of our schemes much.

Table 1: Comparison between diverse signature schemes from SIS and Ring-SIS assumptions. Here, λ is the security parameter, $n = n(\lambda)$ is the lattice hardness parameter, q_s is the number of signing queries, and β is the SIS parameter. For DM'14, the ring $\mathcal{R} = \mathbb{Z}_q[X]/(f(X))$ for some cyclotomic polynomial f of degree n and $q \geq \beta\sqrt{n}\omega(\sqrt{\log n})$. For Alperin'15, δ is a value satisfying $2q_s^2/\epsilon < 2^{\lfloor c^d \rfloor}$ for the security level of the signature scheme in [4] and arbitrary constant $c > 1$. q_{hash} the number of random-oracle queries (if applicable)

Scheme	Signature size	Security loss	Assumption(s)	Standard model?
KW'03 [46]	$O(1) \times \mathbb{Z}^m$	$O(1)$	SIS, $\beta = \tilde{\Omega}(n^{3/2})$	ROM
GPV'08 [35]	$O(1) \times \mathbb{Z}^m$	$O(q_{\text{hash}})$	SIS, $\beta = \tilde{\Omega}(n^{3/2})$	ROM
Boyen'10 [21]	$O(1) \times \mathbb{Z}^m$	$O(\lambda q_s)$	SIS, $\beta = \tilde{\Omega}(n^{7/2})$	✔
Lyu'12 [48]	$O(1) \times \mathbb{Z}^m$	$O(\lambda q_s)$	SIS, $\tilde{\Omega}(n)$	ROM
MP'12 [49]	$O(1) \times \mathbb{Z}^m$	$O(\lambda q_s)$	SIS, $\beta = \tilde{\Omega}(n^{5/2})$	✔
BHJKSS'13 [13]	$O(\log \lambda) \times \mathbb{Z}^m$	$O(\lambda q_s)$	SIS, $\beta = \tilde{\Omega}(n^{5/2})$	✔
DM'14 [31]	$O(1) \times \mathcal{R}_q^{O(\log q)}$	$O(\lambda q_s)$	Ring-SIS, $\beta = \tilde{\Omega}(n^{7/2})$	✔
BKKP'15 [11]	$O(\lambda) \times \mathbb{Z}^m$	$O(1)$	SIS, $\beta = \tilde{\Omega}(n^{3/2})$	✔
Alperin'15 [4]	$O(1) \times \mathbb{Z}^m$	$O(\lambda q_s)$	SIS, $\beta = \tilde{\Omega}(\delta^{2\delta} n^{11/2})$	✔
Ours	$O(1) \times \mathbb{Z}^m$	$O(\lambda)$	SIS+LWE, $\beta = \tilde{\Omega}(n^{O(1)})^\star$	✔

* The exponent $O(1)$ of parameter n accounts for the constant circuit depth of the instantiated PRF from [9].

Efficiency Consideration. Though we focus on tightness of reduction in the context of short signature and IBE, we do not hide the inefficiency of our schemes, particularly with comparison to the adptively secure lattice-based signature/IBE scheme obained from the "complexity leveraging" [14] of efficient selectively secure lattice-based signature/IBE scheme such as [2]. Although complexity leveraing is not very satisfactory from a theoretical perspective, it indeed often leads to the most practical secure cryptographic schemes. In the context of IBE, we have seen that the adaptively secure IBE scheme levearaged from selective DBDH-based IBE scheme in [14] has higher real-world efficiency than the adaptively secure Waters IBE scheme [54] (as well as the subsequent adaptive IBE schemes from similar standard pairing assumptions without random oracles) for the same security level. This may seem counter-intuitive, but to design adaptively secure IBE schemes one needs to carefully embed some specially crafted complex structures into the scheme, to provide enough freedom for the security reduction. This makes directly constructed adaptive IBE schemes rather bulky. Therefore, our current results are of more theoretic value. One the other hand, directly constructing adaptively secure schemes from standard assumptions usually requires new proof ideas and techniques which advance the state-of art and

Table 2: Comparison between diverse signature schemes from various quantum-unsafe assumptions. Here, λ is the security level, $n = n(\lambda)$ is the lattice hardness parameter, q_s the number of signing queries, N the RSA modulus, m the lattice dimension, k a non-adversary-query-dependent parameter of the LIN assumption, for the KW'03, $|D|$ the domain size of the instantiated claw-free permutation, which is abbreviated as CFP, and D-I hash stands for division-intractable hash.

Scheme	Sig. size	Sec. loss	Assumption(s)	Standard model?
GHR'99 [33]	$O(1) \times \mathbb{Z}_N$	$O(1)$	Strong-RSA + D-I Hash	✔
BLS'01 [20]	$O(1) \times \mathbb{G}$	$O(\lambda q_s)$	CDH	ROM
KW'03 [46]	$O(1) \times \lvert D\rvert$	$O(1)$	CFP	ROM
BB'04 [16]	$O(1) \times \mathbb{G}$	$O(1)$	q_s-SDH	✔
Waters'05 [54]	$O(1) \times \mathbb{G}$	$O(\lambda q_s)$	CDH	✔
HW'09 [44]	$O(1) \times \mathbb{Z}_N$	$O(\lambda q_s)$	RSA	✔
BHJKSS'13 [13]	$O(1) \times \mathbb{G}$	$O(\lambda q_s)$	DLog	✔
BHJKSS'13 [13]	$O(1) \times \mathbb{Z}_N$	$O(\lambda q_s)$	RSA	✔
ADKMO'13 [1]	$O(\lambda) \times \mathbb{G}$	$O(1)$	DLIN	✔
CW'13 [28]	$O(k) \times \mathbb{G}$	$O(\lambda)$	k-LIN	✔
BKP'14 [12]	$O(k) \times \mathbb{G}$	$O(\lambda)$	k-LIN	✔
BKKP'15 [11]	$O(\lambda) \times \mathbb{G}$	$O(1)$	DLog	✔
BKKP'15 [11]	$O(\lambda) \times \mathbb{Z}_N$	$O(1)$	RSA,FAC	✔
Ours	$O(1) \times \mathbb{Z}^m$	$O(\log^2 \lambda)$	SIS+DDH, $\beta = \tilde{\Omega}(n^{O(1)})$ ⋆	✔

⋆ The exponent $O(1)$ of parameter n accounts for the constant circuit depth of the instantiated PRF from [9].

lead to further applications. Trying to get tighter reduction for the directly constructed adaptively secure schemes should be always welcome as it remains a very promising way of bridging the efficiency gap.

1.4 Overview of Our Approach

Construction Outline. Our constructions use pseudorandom functions (PRFs). Recall a pseudorandom function $\mathsf{PRF} : \{0,1\}^k \times \{0,1\}^t \to \{0,1\}$ takes as input a truly random secret key from $\{0,1\}^k$ and a string from $\{0,1\}^t$, and deterministically outputs a bit which is computationally indistinguishable from a random bit. In our signature scheme, apart from the "left" matrix $\mathbf{A}$ typical of all SIS/LWE based constructions, we set another $4 + k$ random matrices from $\mathbb{Z}_q^{n\times m}$, comprising: two "signature subspace selection" matrices $\mathbf{A}_0, \mathbf{A}_1$, some "PRF secret key" matrices $\{\mathbf{B}_i\}_{i\in[k]}$, and two "message representation" matrices $\mathbf{C}_0, \mathbf{C}_1$. The key generation algorithm further chooses a secure pseudo-random function PRF, which is expressed as a Boolean circuit, as a part of the public parameters or perhaps a common reference string. The signing key is a "short" basis $\mathbf{T}_{\mathbf{A}}$of $\mathbf{A}$ and a PRF key $K \xleftarrow{\$} \{0,1\}^k$ for PRF.

The signer takes three steps to generate the signature of message $\mathsf{M} = x_1x_2\dots x_t \in \{0,1\}^t$. Firstly, it uses the key-homomorphic algorithm from [19] to compute the unique matrix $\mathbf{A}_{\mathsf{PRF},\mathsf{M}}$ from the circuit of PRF and the $k + t$

Table 3: Comparison between adaptively secure IBE schemes from various assumptions in both standard and random-oracle models. Here, λ is the security level, q_{id} the number of private key queries and q_{hash} the number of random-oracle queries (if applicable).

Scheme	Security loss	Assumption	Standard model?	Quantum-safe
BF'01 [18]	$O(q_{\mathsf{id}})$	BDH	ROM	✗
KW'03 [46]	$O(1)$	BDH	ROM	✗
BB'04a [14]	$O(2^{\lambda})$	DBDH, q_{id}-BDHI	✔	✗
BB'04b [15]	$O(\lambda q_{\mathsf{id}})$	DBDH	✔	✗
Waters'05 [54]	$O(\lambda q_{\mathsf{id}})$	DBDH	✔	✗
Gentry'06 [34]	$O(1)$	q_{id}-ABDHE	✔	✗
GPV'08 [35]	$O(q_{\text{hash}})$	LWE	ROM	✔
Waters'09 [55]	$O(q_{\mathsf{id}})$	DBDH	✔	✗
CHKP'10 [27]	$O(\lambda q_{\mathsf{id}})$	LWE	✔	✔
ABB'10 [2]	$O(\lambda q_{\mathsf{id}})$	LWE	✔	✔
LW'12 [47]	$O(q)$	DLIN	✔	✗
CW'13 [28]	$O(\lambda)$	k-LIN	✔	✗
BKP'14 [12]	$O(\lambda)$	k-LIN	✔	✗
Ours	$O(1)$	LWE	✔	✔

matrices $\{\mathbf{B}_i\}_{i\in[k]}, \mathbf{C}_{x_1}, \mathbf{C}_{x_2}, \ldots, \mathbf{C}_{x_t}$. Then it computes $b = \mathsf{PRF}(K, \mathsf{M})$ and sets the matrix $\mathbf{F}_{\mathsf{M},1-b} = [\mathbf{A} \mid \mathbf{A}_{1-b} - \mathbf{A}_{\mathsf{PRF},\mathsf{M}}] \in \mathbb{Z}_q^{n\times 2m}$. Finally, it applies the trapdoor $\mathbf{T}_{\mathbf{A}}$ to generate the signature: a low-norm non-zero vector $\mathbf{d}_{\mathsf{M}} \in \mathbb{Z}^{2m}$ such that $\mathbf{F}_{\mathsf{M},b} \cdot \mathbf{d}_{\mathsf{M}} = \mathbf{0} \pmod q$. The verification algorithm checks whether the signature is non-zero and has low-norm, and whether $\mathbf{F}_{\mathsf{M},b} \cdot \mathbf{d}_{\mathsf{M}} = \mathbf{0} \pmod q$ or $\mathbf{F}_{\mathsf{M},1-b} \cdot \mathbf{d}_{\mathsf{M}} = \mathbf{0} \pmod q$.

Our IBE scheme works as follows. The public parameters contain matrices $\mathbf{A}$, $\mathbf{A}_0$, $\mathbf{A}_1$, $\{\mathbf{B}_i\}_{i\in[k]}$, $\mathbf{C}_0, \mathbf{C}_1$, a Boolean circuit representation of a secure PRF PRF, and a random syndrome vector $\mathbf{u} \in \mathbb{Z}_q^n$ which is used to hide messages. The trapdoor basis $\mathbf{T}_{\mathbf{A}}$ serves as master secret key. The PKG generates the private key of identity $\mathsf{id} = x_1x_2\ldots x_t \in \{0,1\}^t$ through a similar procedure as the signing algorithm of our signature. It uses the key-homomorphic algorithm to compute the unique matrix $\mathbf{A}_{\mathsf{PRF},\mathsf{id}}$ from the circuit of PRF and the $k+t$ matrices $\{\mathbf{B}_i\}_{i\in[k]}, \mathbf{C}_{x_1}, \mathbf{C}_{x_2}, \ldots, \mathbf{C}_{x_t}$. It then sets the "function" matrix to $\mathbf{F}_{\mathsf{id},1-b} = [\mathbf{A} \mid \mathbf{A}_{1-b} - \mathbf{A}_{\mathsf{PRF},\mathsf{id}}] \in \mathbb{Z}_q^{n\times 2m}$ for a random fair coin $b \xleftarrow{\$} \{0,1\}$, and uses the master secret key to sample a Gaussian vector $\mathbf{d}_{\mathsf{id}} \in \mathbb{Z}^{2m}$ as private identity key such that $\mathbf{F}_{\mathsf{id},b} \cdot \mathbf{d}_{\mathsf{id}} = \mathbf{u} \pmod q$.

To encrypt a message $\mathsf{Msg} \in \{0,1\}$ with an identity id, the encryptor computes $\mathbf{A}_{\mathsf{PRF},\mathsf{id}}$ and sets two "function" matrices $\mathbf{F}_{\mathsf{id},b} = [\mathbf{A} \mid \mathbf{A}_b - \mathbf{A}_{\mathsf{PRF},\mathsf{id}}]$ and $\mathbf{F}_{\mathsf{id},1-b} = [\mathbf{A} \mid \mathbf{A}_{1-b} - \mathbf{A}_{\mathsf{PRF},\mathsf{id}}]$. It generates two independent GPV-style ciphertexts [35]. The first one uses $\mathbf{F}_{\mathsf{id},b}$:

$$\begin{cases} c_{b,0} = \mathbf{s}_b^\top \mathbf{u} + \nu_{b,0} + \mathsf{Msg} \cdot \lfloor q/2 \rfloor \\ \mathbf{c}_{b,1}^\top = \mathbf{s}_b^\top \mathbf{F}_{\mathsf{id},b} + \boldsymbol{\nu}_{b,1}^\top \end{cases}$$

and the second is based on $\mathbf{F}_{\mathsf{id},1-b}$:

$$\begin{cases} c_{1-b,0} &= \mathbf{s}_{1-b}^{\top}\mathbf{u} + \nu_{1-b,0} + \mathsf{Msg} \cdot \lfloor q/2 \rfloor \\ \mathbf{c}_{1-b,1}^{\top} &= \mathbf{s}_{1-b}^{\top}\mathbf{F}_{\mathsf{id},1-b} + \boldsymbol{\nu}_{1-b,1}^{\top} \end{cases}$$

for random vectors $\mathbf{s}_b, \mathbf{s}_{1-b} \xleftarrow{\$} \mathbb{Z}_q^n$, two small noise scalars $\nu_{b,0}, \nu_{1-b,0}$, and two low-norm noise vectors $\boldsymbol{\nu}_{b,1}, \boldsymbol{\nu}_{1-b,1}$.

The decryption algorithm uses $\mathbf{d}_{\mathsf{id}}$ to try both ciphertexts; one of them should work. Here as a technical caveat, we need some redundant information in the messages in order to check whether a recovered message is well-formed. To this end, one option is to apply the standard way of encrypting multiple bits in GPV-style ciphertexts without affecting the security analysis. That is, instead of using just a vector $\mathbf{u} \in \mathbb{Z}_q^n$ in the public key, we use a matrix $\mathbf{U} \in \mathbb{Z}_q^{n\times z}$ allowing us to encrypt z bits. A second option, which costs nothing if hybrid encryption is being used, is to use multi-bit GPV-like encryption to encrypt a symmetric session key without redundancy, again using a matrix $\mathbb{Z}_q^{n\times z}$ and rely on downstream symmetric integrity checks or MACs to weed out the incorrect ciphertexts.

Proof Outline. The security reduction of our signature scheme uses an efficient adversary to solve a of SIS problem instance $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$: a short vector $\mathbf{e} \in \mathbb{Z}$ such that $\mathbf{Ae} = \mathbf{0} \pmod q$. The reduction embeds a randomly picked secret key K for PRF in verification key. Along with K, PRF is applied to select the bit value b on M from queries. More specifically, the reduction selects low-norm matrices $\mathbf{R}_{\mathbf{A}_0}$, $\mathbf{R}_{\mathbf{A}_1}$, $\{\mathbf{R}_{\mathbf{B}_i}\}_{i\in[k]}$, $\mathbf{R}_{\mathbf{C}_0}$, $\mathbf{R}_{\mathbf{C}_1} \xleftarrow{\$} \{1,-1\}^{m\times m}$, a PRF secret key $K = s_1 s_2 \ldots s_k \xleftarrow{\$} \{0,1\}^k$ and sets $\mathbf{A}_0 = \mathbf{A}\mathbf{R}_{\mathbf{A}_0}$, $\mathbf{A}_1 = \mathbf{A}\mathbf{R}_{\mathbf{A}_1} + \mathbf{G}$, $\{\mathbf{B}_i = \mathbf{A}\mathbf{R}_{\mathbf{B}_i} + s_i\mathbf{G}\}_{i\in[k]}$, $\mathbf{C}_0 = \mathbf{A}\mathbf{R}_{\mathbf{C}_0}$ and $\mathbf{C}_1 = \mathbf{A}\mathbf{R}_{\mathbf{C}_1} + \mathbf{G}$. Here, K is completely hidden from adversary's view. For answering a signing query on message M, the reduction computes $\mathbf{A}_{\mathsf{PRF},\mathsf{M}} = \mathbf{AR} + \mathsf{PRF}(K,\mathsf{M})\mathbf{G}$ for some known low-norm $m \times m$ matrix $\mathbf{R}$ that depends on $\mathbf{R}_{\mathbf{A}_0}, \mathbf{R}_{\mathbf{A}_1}, \{\mathbf{R}_{\mathbf{B}_i}\}_{i\in[k]}, \mathbf{R}_{\mathbf{C}_0}, \mathbf{R}_{\mathbf{C}_1}$, K and M. Let $\mathsf{PRF}(K,\mathsf{M}) = b$, the reduction sets $\mathbf{F}_{\mathsf{M},1-b} = [\mathbf{A} \mid \mathbf{A}_{1-b} - \mathbf{A}_{\mathsf{PRF},\mathsf{M}}] = [\mathbf{A} \mid \mathbf{AR} + (1-2b)\mathbf{G}]$ and uses the trapdoor from $\mathbf{G}$ to compute the decryption key. Note, we use PRF to select the matrix $\mathbf{A}_b$ which is the same as the real scheme, for the message M^* of a forgery, since $b = \mathsf{PRF}(K,\mathsf{M}^*)$ is unpredictable for the adversary. With essentially probability $1/2$ the forged non-zero signature $\mathbf{d}_{\mathsf{M}^*}$ satisfies $\mathbf{F}_{\mathsf{M}^*,b}\mathbf{d}_{\mathsf{M}^*} = \mathbf{0} \pmod q$ leading to a valid SIS solution.

The security reduction for our IBE scheme is similar to the reduction of the signature scheme. Basically, the reduction answers key generation queries the same way as answering signing queries in the signature scheme reduction. To construct the challenge ciphertext for a challenge identity id^*, the LWE challenge is embedded in the function matrix $\mathbf{F}_{\mathsf{id}^*,b} = [\mathbf{A} \mid \mathbf{AR}]$ for which the simulator cannot produce private key. Another ciphertext based on $\mathbf{F}_{\mathsf{id}^*,1-b} = [\mathbf{A} \mid \mathbf{AR} + (1-2b)\mathbf{G}]$ is generated as in the real scheme. With half probability, the adversary will choose the ciphertext under $\mathbf{F}_{\mathsf{id}^*,b}$ to attack giving out useful information for solving the LWE challenge. We refer to the full details in the rest of the paper.

Related Works. In the related and concurrent work by Brakerski and Vaikuntanathan [24], a similar idea of embedding PRFs into encryption schemes has been used to construct the first semi-adaptively secure attribute-based encryption scheme from lattices supporting an a priori unbounded number of attributes. The recent work by Bai et al. [7] addresses the problem of improving efficiency of lattice-based cryptographic schemes via a different but novel way. Their proposal is about using Rényi divergence instead of statistical distance in the context of lattice-based cryptography which leads to (sometimes simpler) security proofs for more efficient lattice-based schemes.

2 Preliminaries

Notation. 'PPT' abbreviates "probabilistic polynomial-time". If S is a set, we denote by $a \overset{\$}{\leftarrow} S$ the uniform sampling of a random element of S. For a positive integer n, we denote by $[n]$ the set of positive integers no greater than n. We use bold lowercase letters (e.g. $\mathbf{a}$) to denote vectors and bold capital letters (e.g. $\mathbf{A}$) to denote the matrices. For a positive integer $q \geq 2$, let $\mathbb{Z}_q$ be the ring of integers modulo q. We denote the group of $n \times m$ matrices in $\mathbb{Z}_q$ by $\mathbb{Z}_q^{n\times m}$. Let $\mathbf{I}_m$ be the $m \times m$ identity matrix. Vectors are treated as column vectors and the transpose of a matrix $\mathbf{A}$ is denoted by $\mathbf{A}^\top$. For $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ and $\mathbf{B} \in \mathbb{Z}_q^{n\times m'}$, let $[\mathbf{A}|\mathbf{B}] \in \mathbb{Z}_q^{n\times(m+m')}$ be the concatenation of $\mathbf{A}$ and $\mathbf{B}$. We denote by $\|\mathbf{A}\|$ or $\|\mathbf{a}\|$ the Euclidean norm of a matrix $\|\mathbf{A}\|$ or vector $\|\mathbf{a}\|$. We denote by $\tilde{\mathbf{A}}$ the Gram-Schmidt ordered orthogonalization of $\|\mathbf{A}\|$, and its Euclidean norm by $\|\tilde{\mathbf{A}}\|$. The inner product of two vectors $\mathbf{x}$ and $\mathbf{y}$ is written $\langle \mathbf{x}, \mathbf{y}\rangle$. For a security parameter λ, a function $\mathsf{negl}(\lambda)$ is negligible in λ if it is smaller than all polynomial fractions for a sufficiently large λ.

Randomness Extractor. We recall the following generalisation of the left-over hash lemma.

Lemma 1 ([2], Lemma 4). *Suppose that $m > (n+1)\log q + \omega(\log n)$ and that $q > 2$ is prime. Let $\mathbf{R}$ be an $m\times k$ matrix chosen uniformly in $\{1,-1\}^{m\times k} \mod q$ where $k = k(n)$ is polynomial in n. Let $\mathbf{A}$ and $\mathbf{B}$ be matrices chosen uniformly in $\mathbb{Z}_q^{n\times m}$ and $\mathbb{Z}_q^{n\times k}$ respectively. Then, for all vectors $\mathbf{w} \in \mathbb{Z}_q^m$, the distribution $(\mathbf{A}, \mathbf{AR}, \mathbf{R}^\top\mathbf{w})$ is statistically close to the distribution $(\mathbf{A}, \mathbf{B}, \mathbf{R}^\top\mathbf{w})$.*

Norm of a Random Matrix. Let $\mathbf{S}^m$ be the m-sphere $\{\mathbf{x} \in \mathbb{R}^{m+1} : \|\mathbf{x}\| = 1\}$. We define $s_\mathbf{R} \overset{\mathsf{def}}{=} \|\mathbf{R}\| = \sup_{\mathbf{x}\in\mathbb{R}^{m+1}} \|\mathbf{R}\cdot\mathbf{x}\|$.

Lemma 2 ([2], Lemma 5). *Let $\mathbf{R}$ be a random chosen matrix from $\{1,-1\}^{m\times m}$, then* $\Pr[\|\mathbf{R} > 12\sqrt{m}\|] < e^{-m}$.

2.1 Lattice Background

Lattice Definitions

Definition 1. *Let a basis* $\mathbf{B} = [\mathbf{b}_1 \mid \ldots \mid \mathbf{b}_m] \in (\mathbb{R}^m)^m$ *of linearly independent vectors. The lattice generated by* $\mathbf{B}$ *is defined as* $\Lambda = \{\mathbf{y} \in \mathbb{R}^m \ : \exists s_i \in \mathbb{Z},\ \mathbf{y} = \sum_{i=1}^m s_i \mathbf{b}_i\}$. *The dual lattice* Λ^* *of* Λ *is defined as* $\Lambda^* = \{\mathbf{z} \in \mathbb{R}^m \ : \forall \mathbf{y} \in \Lambda,\ \langle \mathbf{z}, \mathbf{y} \rangle \in \mathbb{Z}\}$.

Definition 2. *For* q *prime,* $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ *and* $\mathbf{u} \in \mathbb{Z}_q^n$, *we define the* m*-dimensional (full-rank) random integer lattice* $\Lambda_q^{\perp}(\mathbf{A}) = \{\mathbf{e} \in \mathbb{Z}^m \ : \mathbf{A}\mathbf{e} = \mathbf{0} \pmod q\}$, *and the "shifted lattice" as the coset* $\Lambda_q^{\mathbf{u}}(\mathbf{A}) = \{\mathbf{e} \in \mathbb{Z}^m \ : \mathbf{A}\mathbf{e} = \mathbf{u} \pmod q\}$.

Trapdoors of Lattices and Discrete Gaussians It is shown in [3,49] how to sample a "nearly" uniform random matrix $\mathbf{A} \in \mathbb{Z}^{n\times m}$ along with a trapdoor matrix $\mathbf{T}_{\mathbf{A}} \in \mathbb{Z}^{m\times m}$ which is a short or low-norm basis of the induced lattice $\Lambda_q^{\perp}(\mathbf{A})$. We refer to this procedure as $\mathsf{TrapGen}$.

Lemma 3. *There is a PPT algorithm* $\mathsf{TrapGen}$ *that takes as input integers* $n \geq 1$, $q \geq 2$ *and a sufficiently large* $m = O(n \log q)$, *outputs a matrix* $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ *and a trapdoor matrix* $\mathbf{T}_{\mathbf{A}} \in \mathbb{Z}^{m\times m}$, *such that* $\mathbf{A} \cdot \mathbf{T}_{\mathbf{A}} = 0$, *the distribution of* $\mathbf{A}$ *is statistically close to the uniform distribution over* $\mathbb{Z}_q^{n\times m}$ *and* $\|\tilde{\mathbf{T}}_{\mathbf{A}}\| = O(\sqrt{n \log q})$.

Discrete Gaussians. Let $m \in \mathbb{Z}_{>0}$ be a positive integer and $\Lambda \subset \mathbb{Z}^m$. For any real vector $\mathbf{c} \in \mathbb{R}^m$ and positive parameter $\sigma \in \mathbb{R}_{>0}$, let the Gaussian function $\rho_{\sigma,\mathbf{c}}(\mathbf{x}) = \exp\left(-\pi\|\mathbf{x} - \mathbf{c}\|^2/\sigma^2\right)$ on $\mathbb{R}^m$ with center $\mathbf{c}$ and parameter σ. Define the discrete Gaussian distribution over Λ with center $\mathbf{c}$ and parameter σ as $D_{\Lambda,\sigma} = \rho_{\sigma,\mathbf{c}}(\mathbf{y})/\rho_\sigma(\Lambda)$ for $\forall \mathbf{y} \in \Lambda$, where $\rho_\sigma(\Lambda) = \sum_{\mathbf{x}\in\Lambda} \rho_{\sigma,\mathbf{c}}(\mathbf{x})$. For notational convenience, $\rho_{\sigma,\mathbf{0}}$ and $D_{\Lambda,\sigma,\mathbf{0}}$ are abbreviated as ρ_σ and $D_{\Lambda,\sigma}$.

The following lemma bounds the length of a discrete Gaussian vector with sufficiently large Gaussian parameter.

Lemma 4 ([50]). *For any lattice* Λ *of integer dimension* m *with basis* $\mathbf{T}$, $\mathbf{c} \in \mathbb{R}^m$ *and Gaussian parameter* $\sigma \geq \|\tilde{\mathbf{T}}\|\omega(\sqrt{\log m})$, *we have* $\Pr[\|\mathbf{x} - \mathbf{c}\| > \sigma\sqrt{m} : \mathbf{x} \leftarrow D_{\Lambda,\sigma,\mathbf{c}}] \leq \mathsf{negl}(n)$.

Smoothing Parameter. We recall the very important notion of smoothing parameter of a lattice Λ. It is the smallest value of s such that the discrete Gaussian $D_{\Lambda,s}$ "behaves" like a continuous Gaussian.

Definition 3 ([50]). *For any lattice* Λ *and positive real tolerance* $\epsilon > 0$, *the smoothing parameter* $\eta_\epsilon(\Lambda)$ *is the smallest real* $s > 0$ *such that* $\rho_{1/s}(\Lambda^* \backslash \{\mathbf{0}\}) < \epsilon$.

We will make use of the following lemma, which is a special case of Corollary 3.10 from [53].

Lemma 5 (special case of Corollary 3.10 of [53]). *Let* $\mathbf{r} \in \mathbb{Z}^m$ *be a vector and* $r, \alpha > 0$ *be reals. Assume that* $1/\sqrt{1/r^2 + (\|\mathbf{r}\|/\alpha)^2} \geq \eta_\epsilon(\mathbb{Z}^m)$ *for some* $\epsilon < 1/2$. *Let* $\mathbf{y}$ *be a vector with distribution* $D_{\mathbb{Z}^m,r}$ *and* e *be a scalar with distribution* $D_{\mathbb{Z},\alpha}$. *The distribution of* $\langle \mathbf{r}, \mathbf{y} \rangle + e$ *is statistically close to* $D_{\mathbb{Z},\sqrt{(r\|\mathbf{r}\|)^2+\alpha^2}}$.

Lattice Sampling Algorithms Our constructions make use of the "two-sided trapdoor" framework from [2,21] which consists of two sampling algorithms SampleLeft and SampleRight.

$$\textit{Algorithm } \mathsf{SampleLeft}(\mathbf{A}, \mathbf{B}, \mathbf{T}_{\mathbf{A}}, \mathbf{u}, s) \tag{1}$$

Inputs: a full-rank matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and a short basis $\mathbf{T}_{\mathbf{A}}$ of $\Lambda_q^{\perp}(\mathbf{A})$, a matrix $\mathbf{B} \in \mathbb{Z}_q^{n \times m_1}$, a vector $\mathbf{u} \in \mathbb{Z}_q^n$, and a Gaussian parameter s.

Output: Let $\mathbf{F} = \begin{bmatrix} \mathbf{A} \mid \mathbf{B} \end{bmatrix}$. The algorithm outputs a vector $\mathbf{d} \in \mathbb{Z}^{m+m_1}$ in the set $\Lambda_q^{\mathbf{u}}(\mathbf{F})$.

Theorem 1 ([2,27]). *Let $q > 2$, $m > n$ and $s > \|\tilde{\mathbf{T}}_{\mathbf{A}}\| \cdot \omega(\sqrt{\log(m + m_1)})$. Then SampleLeft$(\mathbf{A}, \mathbf{B}, \mathbf{T}_{\mathbf{A}}, \mathbf{u}, s)$ taking inputs as in (1), outputs a vector $\mathbf{d} \in \mathbb{Z}^{m+m_1}$ distributed statistically close to $D_{\Lambda_q^{\mathbf{u}}(\mathbf{F}), s}$.*

$$\textit{Algorithm } \mathsf{SampleRight}(\mathbf{A}, \mathbf{B}, \mathbf{R}, \mathbf{T}_{\mathbf{B}}, \mathbf{u}, s) \tag{2}$$

Inputs: matrices $\mathbf{A} \in \mathbb{Z}_q^{n \times k}$ and $\mathbf{R} \in \mathbb{Z}^{k \times m}$, a full-rank matrix $\mathbf{B} \in \mathbb{Z}_q^{n \times m}$, a short basis $\mathbf{T}_{\mathbf{B}}$ of $\Lambda_q^{\perp}(\mathbf{B})$, a vector $\mathbf{u} \in \mathbb{Z}_q^n$, and a Gaussian parameter s.

Output: Let $\mathbf{F} = \begin{bmatrix} \mathbf{A} \mid \mathbf{A}\mathbf{R} + \mathbf{B} \end{bmatrix}$; the algorithm outputs a vector $\mathbf{d} \in \mathbb{Z}^{m+m_1}$ in the set $\Lambda_q^{\mathbf{u}}(\mathbf{F})$

Theorem 2 ([2], Theorem 19). *Let $q > 2$, $m > n$. Let $s_{\mathbf{R}} := \|\mathbf{R}\|$ and $s > \|\tilde{\mathbf{T}}_{\mathbf{B}}\| \cdot s_{\mathbf{R}} \cdot \omega(\sqrt{\log m})$. Then SampleRight$(\mathbf{A}, \mathbf{B}, \mathbf{R}, \mathbf{T}_{\mathbf{B}}, \mathbf{u}, s)$ taking inputs as in (2), outputs a vector $\mathbf{d} \in \mathbb{Z}^{m+k}$ distributed statistically close to $D_{\Lambda_q^{\mathbf{u}}(\mathbf{F}), s}$.*

Gadget Matrix In our construction, instead of using a random matrix $\mathbf{B}$ in the algorithm SampleRight, we will use the "gadget matrix" $\mathbf{G}$ defined in [49]. We recall the following two facts.

Lemma 6 ([49], Theorem 1). *Let q be a prime, and n, m be integers with $m = n \log q$. There is a fixed full-rank matrix $\mathbf{G} \in \mathbb{Z}_q^{n \times m}$ such that the lattice $\Lambda_q^{\perp}(\mathbf{G})$ has a publicly known trapdoor matrix $\mathbf{T}_{\mathbf{G}} \in \mathbb{Z}^{n \times m}$ with $\|\tilde{\mathbf{T}}_{\mathbf{G}}\| \leq \sqrt{5}$.*

Lemma 7 ([19], Lemma 2.1). *There is a deterministic algorithm, denoted $\mathbf{G}^{-1}(\cdot) : \mathbb{Z}_q^{n \times m} \to \mathbb{Z}^{m \times m}$, that takes any matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ as input, and outputs the preimage $\mathbf{G}^{-1}(\mathbf{A})$ of $\mathbf{A}$ such that $\mathbf{G} \cdot \mathbf{G}^{\mathbf{-1}}(\mathbf{A}) = \mathbf{A} \pmod q$ and $\|\mathbf{G}^{-1}(\mathbf{A})\| \leq 2\sqrt{m} \leq m$.*

Computational Assumptions We recall the two most mainstream and conservative average-case computational assumptions for lattice problems.

The Learning with Errors (LWE) Assumption. The learning with errors problem was first proposed by Regev [53]. For a vector $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n$ and a noise distribution χ over $\mathbb{Z}_q$, let $A_{\mathbf{s},\chi}$ be the distribution over $\mathbb{Z}_q^n \times \mathbb{Z}_q$ by taking $\mathbf{a} \xleftarrow{\$} \mathbb{Z}_q^n$ and $x \leftarrow \chi$, and outputting $(\mathbf{a}, \mathbf{s}^\top \mathbf{a} + x) \pmod q$. Usually, χ is a discrete Gaussian $D_{\mathbb{Z},\alpha q}$ for some $\alpha < 1$, reduced modulo q. We refer to [53] for further details.

Definition 4. *For a security parameter Λ, let a positive integer $n = n(\lambda)$, a prime $q = q(\lambda)$, and a distribution χ over $\mathbb{Z}_q$. The learning with errors problem $LWE_{n,q,\chi}$ is to distinguish the oracle $\mathcal{O}_{\mathbf{s}}$, which outputs samples from the distribution $A_{\mathbf{s},\chi}$, from the oracle $\mathcal{O}_{\$}$, which outputs samples from the uniform distribution over $\mathbb{Z}_q^n \times \mathbb{Z}_q$, for an unspecified polynomial number of queries. We define the advantage (in the security parameter λ) of an algorithm $\mathcal{A}$ in solving the $LWE_{n,q,\chi}$ problem as*

$$\mathsf{Adv}_{\mathcal{A}}^{LWE_{n,q,\chi}}(\lambda) = \left|\Pr[\mathcal{A}^{\mathcal{O}_{\mathbf{s}}}(1^\lambda) = 1] - \Pr[\mathcal{A}^{\mathcal{O}_{\$}}(1^\lambda) = 1]\right|$$

We say that the $(t, \epsilon_{\mathsf{LWE}})$-$\mathsf{LWE}_{n,q,\chi}$ assumption holds if no t-time algorithm $\mathcal{A}$ that has advantage at least ϵ_{LWE} in solving the $\mathsf{LWE}_{n,q,\chi}$ problem.

For polynomial size q in λ, there are known quantum [53] and classical [22] reductions from the average-case $\mathrm{LWE}_{n,q,\chi}$ assumption to many standard worst-case lattice problems (e.g., GapSVP). [1] This further strengthens the appeal of the LWE assumption. Peikert [52] also gave a classic reduction that applies (only) for exponential moduli q in λ.

The Short Integer Solution (SIS) Assumption. The security of our adaptively secure signature scheme is based on the SIS problem, which can be seen as an average-case approximate shortest vector problem on random integer lattices, or also as the decoding problem for random linear codes. In a sense, SIS is the computational counterpart to the decisional LWE.

Definition 5. *For a security parameter λ, let $n = n(\lambda)$, $m = m(\lambda)$, and $\beta = \beta(\lambda)$. Let q be a prime integer. The short integer solution problem $\mathsf{SIS}_{n,q,\beta,m}$ is as follows. Given a uniform random matrix $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$, find a non-zero vector $\mathbf{e} \in \mathbb{Z}^m$ such that $\mathbf{Ae} = \mathbf{0} \pmod q$ and $\|\mathbf{e}\| \le \beta$. We define the advantage (function of the security parameter λ) of an algorithm $\mathcal{A}$ in solving the $\mathsf{SIS}_{n,q,\beta,m}$ problem as*

$$\mathsf{Adv}_{\mathcal{A}}^{\mathsf{SIS}_{n,q,\beta,m}}(\lambda) = \left[\begin{matrix}\mathbf{Ae} = \mathbf{0} \pmod q \\ \text{and } \|\mathbf{e}\| \le \beta, \\ \text{and } \mathbf{e} \ne \mathbf{0}.\end{matrix} \;:\; \begin{matrix}\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m} \\ \mathbf{e} \leftarrow \mathcal{A}(1^\lambda, \mathbf{A})\end{matrix}\right]$$

We say the $(t, \epsilon_{\mathsf{SIS}})$-$\mathsf{SIS}_{n,q,\beta,m}$ assumption holds if no t-time algorithm $\mathcal{A}$ that has advantage at least ϵ_{SIS} in solving the $\mathsf{SIS}_{n,q,\beta,m}$ problem.

[1] Equivalently, this is to say that many classic worst-case lattice *problems* reduce *to* the average-case LWE *problem*, for suitable parameters.

It has been shown in [50] that solving the average-case instances of the $\mathsf{SIS}_{n,q,\beta,m}$ problem for certain parameters is as hard as solving worst-case instances of the approximate Shortest Independent Vector Problem (SIVP).

2.2 Key-Homomorphic Evaluation Algorithm

We recall the key-homomorphic algorithm Eval developed by Gentry et al. [37] and Boneh et al. [19]. The deterministic algorithm Eval takes as input a Boolean circuit $C : \{0,1\}^k \to \{0,1\}$ with k input wires which have matrix encoding $\mathbf{A}_1, \mathbf{A}_2, \ldots, \mathbf{A}_k \in \mathbb{Z}_q^{n\times m}$, and outputs a matrix $\mathbf{A}_C \in \mathbb{Z}_q^{n\times m}$ encoding the output wire of C. Without loss of generality, assume the Boolean circuit C is a composition of NAND gates. For every NAND gate $g(u,v;w)$ with input wires u,v and output wire w, assume we have already computed the matrices $\mathbf{A}_u$ and $\mathbf{A}_v$ encoding the wires u and v respectively. The algorithm Eval defines $\mathbf{A}_w = \mathbf{G} - \mathbf{A}_u \cdot \mathbf{G}^{-1}(\mathbf{A}_v) \in \mathbb{Z}_q^{n\times m}$ to be the matrix encoding of output wire w. Following this computation Eval finally obtains the matrix $\mathbf{A}_C$. [2]

In the simulation, we will assign actual inputs to C by constructing $\mathbf{A}_i = \mathbf{A}\mathbf{R}_i + x_i\mathbf{G}$ for i-th binary input $x_i \in \{0,1\}$ and random low-norm matrices $\mathbf{R}_i \xleftarrow{\$} \{1,-1\}^{m\times m}$. For a NAND gate $g(u,v;w)$ where the input wires take input bits x_u, x_v, we let $\mathbf{A}_u = \mathbf{A}\mathbf{R}_u + x_u\mathbf{G}$ and $\mathbf{A}_v = \mathbf{A}\mathbf{R}_v + x_v\mathbf{G}$ be two matrices according to u and v. We have

$$\begin{aligned}
\mathbf{A}_w &= \mathbf{G} - \mathbf{A}_u \cdot \mathbf{G}^{-1}(\mathbf{A}_v) \\
&= \mathbf{G} - (\mathbf{A}\mathbf{R}_u + x_u\mathbf{G}) \cdot \mathbf{G}^{-1}(\mathbf{A}\mathbf{R}_v + x_v\mathbf{G}) \\
&= \mathbf{G} - \mathbf{A}\left(\mathbf{R}_u \cdot \mathbf{G}^{-1}(\mathbf{A}_v) + x_u\mathbf{R}_v\right) - x_u x_v\mathbf{G} \\
&= \mathbf{A}\left(-\mathbf{R}_u \cdot \mathbf{G}^{-1}(\mathbf{A}_v) - x_u\mathbf{R}_v\right) + (1 - x_u x_v)\mathbf{G} \\
&= \mathbf{A}\mathbf{R}_g + (1 - x_u x_v)\mathbf{G}
\end{aligned}$$

where $1 - x_u x_v \overset{\mathsf{def}}{=} \mathsf{NAND}(x_u, x_v)$, and $\mathbf{R}_g = -\mathbf{R}_u \cdot \mathbf{G}^{-1}(\mathbf{A}_v) - x_u\mathbf{R}_v$ has low-norm if $\mathbf{R}_u, \mathbf{R}_v$ have low-norm. Inductively applying above procedure, given k input matrices $\mathbf{A}_i = \mathbf{A}\mathbf{R}_i + x_i\mathbf{G}$ where $\{x_i\}_{i\in[k]}$ are actual inputs of the Boolean circuit C, we can deterministically compute $\mathbf{A}_C = \mathbf{A}\mathbf{R}_C + C(x_1,\ldots,x_k)\mathbf{G}$ where $C(x_1,\ldots,x_k)$ is the output bit of C on the arguments $x_1,\ldots,x_k$.

Noise Growth. Let $d_{\max}$ be the depth of circuit C. Consider the computation on the NAND gate $g(u,v;w)$ above. We have

$$\begin{aligned}
\|\mathbf{R}_g\| &\le \|\mathbf{R}_u\| \cdot \|\mathbf{G}^{-1}(\mathbf{A}_v)\| + |x_u| \cdot \|\mathbf{R}_v\| \\
&\le \max(\|\mathbf{R}_u\|, \|\mathbf{R}_v\|) \cdot (1 + m)
\end{aligned}$$

[2] Note, for random $\mathbf{A}_u$, $\mathbf{A}_v$, and $\mathbf{A}'_v$ where $\mathbf{A}_v \neq \mathbf{A}'_v$, $\mathbf{G} - \mathbf{A}_u \cdot \mathbf{G}^{-1}(\mathbf{A}_v) = \mathbf{A}_w = \mathbf{A}'_w = \mathbf{G} - \mathbf{A}_u \cdot \mathbf{G}^{-1}(\mathbf{A}'_v)$ with negligible probability. Finding such $\mathbf{A}_v$ and $\mathbf{A}'_v$ is equivalent to solving SIS on random $\mathbf{A}_u$ since $\mathbf{G}^{-1}(\cdot)$ outputs matrices with small entries.

The second inequality holds due to the Lemma 7. Thus, starting from the noise level of the inputs of C, by induction the output noise will be $\|\mathbf{R}_C\| \leq (m+1)^{d_{\max}} \cdot \max(\|\mathbf{R}_1\|, \ldots, \|\mathbf{R}_k\|)$; therefore by Lemma 2 we get the bound $\|\mathbf{R}_C\| \leq m^{O(d_{\max})}$. (This translates into $O(d_{\max} \log m)$ bits of noise.)

2.3 Pseudorandom Functions

Definition 6 (Pseudorandom Functions). *Let $\lambda > 0$ be the security parameter, and let $k = k(\lambda)$, $t = t(\lambda)$ and $\ell = \ell(\lambda)$. A pseudorandom function $\mathsf{PRF}\colon \{0,1\}^k \times \{0,1\}^t \to \{0,1\}^\ell$ is an efficiently computable, deterministic two-input function where the first input, denoted by K, is the key. Let Ω be the set of all functions that map t bits strings to ℓ bits strings. We define the advantage (in the security parameter λ) of an adversary $\mathcal{A}$ in attacking the PRF as*

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

where the probability is taken over a uniform choice of key $K \xleftarrow{\$} \{0,1\}^k$ and $F \xleftarrow{\$} \Omega$, and the randomness of $\mathcal{A}$. We say that PRF is $(t_{\mathsf{PRF}}, \epsilon_{\mathsf{PRF}})$-secure if for all t_{PRF}-time adversaries $\mathcal{A}$, $\mathsf{Adv}_{\mathsf{PRF},\mathcal{A}}(\lambda) \leq \epsilon_{\mathsf{PRF}}$.

2.4 Digital Signatures

A digital signature scheme consists of three PPT algorithms: KeyGen, Sign, and Ver. The algorithm KeyGen takes as input a security parameter and generates a public verification key Vk and a private signing key Sk. The signing algorithm Sign takes as input the signing key Vk and a massage M, and outputs the signature Sig of M. The verification algorithm Ver takes as input a signature-message pair $(\mathsf{Sig}, \mathsf{M})$ as well as the verification key Vk. It outputs 1 if Sig is valid, or 0 if Sig is invalid.

We review the standard security notion of digital signature schemes. The existential unforgeability under chosen-message attack ($\mathsf{EUF\text{-}CMA}$) of a digital signature scheme Π is defined through the following security game between an adversary $\mathcal{A}$ and a challenger $\mathcal{B}$.

Setup. $\mathcal{B}$ runs $\mathsf{Setup}(1^\lambda) \to (\mathsf{Sk}, \mathsf{Vk})$, and passes Vk to $\mathcal{A}$.

Query. $\mathcal{A}$ adaptively selects messages $\mathsf{M}_1, \ldots, \mathsf{M}_{q_s}$ to ask for the corresponding signatures under Vk from $\mathcal{B}$. For the query M_i, $\mathcal{B}$ responds with a signature $\mathsf{Sig}_i \leftarrow \mathsf{Sign}(\mathsf{Sk}, \mathsf{M}_i)$.

Forge. $\mathcal{A}$ outputs a pair $(\mathsf{Sig}^*, \mathsf{M}^*)$ and wins if

1. $\mathsf{M}^* \notin \{\mathsf{M}_1, \ldots, \mathsf{M}_{q_s}\}$, and
2. $\mathsf{Ver}(\mathsf{Vk}, \mathsf{Sig}^*, \mathsf{M}^*) \to 1$.

We refer to such an adversary $\mathcal{A}$ as $\mathsf{EUF\text{-}CMA}$ adversary. We define the advantage (in the security parameter λ) $\mathsf{Adv}_{\Pi,\mathcal{A}}(\lambda)$ of $\mathcal{A}$ in attacking a digital signature scheme Π to be the probability that $\mathcal{A}$ wins above game.

Definition 7. *For a security parameter λ, let $t = t(\lambda)$, $q_s = q_s(\lambda)$ and $\epsilon = \epsilon(\lambda)$. We say that a digital signature scheme Π is (t, q_s, ϵ)-EUF-CMA secure if for any t time EUF-CMA adversary $\mathcal{A}$ that makes at most q_s signing queries and has $\mathsf{Adv}_{\Pi,\mathcal{A}}(\lambda) \leq \epsilon$.*

2.5 Identity-Based Encryption

An Identity-Based Encryption system (IBE) consists of four PPT algorithms: Setup, KeyGen, Encrypt, and Decrypt. The algorithm Setup takes as input a security parameter and generates public parameters Pub and a master secret key Msk. The algorithm KeyGen uses the master secret key Msk to produce an identity private key $\mathsf{Sk}_{\mathsf{id}}$ corresponding to an identity id. The algorithm Encrypt takes the public parameters Pub to encrypt messages for any given identity id. The algorithm Decrypt decrypts ciphertexts using the identity private key if the identity of the ciphertext matches the identity of the private key.

We review the security model of IBE proposed in [18], which defines the indistinguishability of ciphertexts under an adaptive chosen-ciphertext and adaptive chosen-identity attack (IND-ID-CCA2). The IND-ID-CCA2 security of IBE is defined through the following game between an adversary $\mathcal{A}$ and a challenger $\mathcal{B}$. For a security parameter λ, let $\mathcal{M}_\lambda$ be the message space, and $\mathcal{C}_\lambda$ be the ciphertext space.

Setup. $\mathcal{B}$ runs $\mathsf{Setup}(1^\lambda) \rightarrow (\mathsf{Pub}, \mathsf{Msk})$, passes the public parameters Pub to $\mathcal{A}$, and keeps the master secret Msk.

Phase 1. $\mathcal{A}$ adaptively issues queries $q_1, \ldots, q_m$ where each query q_i is one of:

- Private key query for identity id_i. $\mathcal{B}$ runs KeyGen to generate $\mathsf{Sk}_{\mathsf{id}_i}$ and sends it to $\mathcal{A}$.
- Decryption query for a ciphertext $\mathsf{Ctx}_{\mathsf{id}_i}$ under identity id_i. $\mathcal{B}$ runs KeyGen to generate $\mathsf{Sk}_{\mathsf{id}_i}$. Then, $\mathcal{B}$ runs the decryption algorithm to decrypt $\mathsf{Ctx}_{\mathsf{id}_i}$ and returns the message to $\mathcal{A}$.

Challenge. When $\mathcal{A}$ decides the Phase 1 is over, it outputs a challenge identity id^*, which is not been queried during Phase 1, and two equal length messages $\mathsf{Msg}_0, \mathsf{Msg}_1 \in \mathcal{M}_\lambda$. $\mathcal{B}$ flips a fair coin $\gamma \xleftarrow{\$} \{0, 1\}$ and sets $\mathsf{Ctx}_{\mathsf{id}^*} \leftarrow \mathsf{Encrypt}(\mathsf{Pub}, \mathsf{Msg}_\gamma, \mathsf{id}^*)$. Finally $\mathcal{A}$ passes $\mathsf{Ctx}_{\mathsf{id}^*}$ to $\mathcal{A}$.

Phase 2. $\mathcal{A}$ adaptively issues more queries $q_{m+1}, \ldots, q_n$ where q_i is one of

- Private key query for identity $\mathsf{id}_i \neq \mathsf{id}^*$.
- Decryption query for a ciphertext $\mathsf{Ctx}_{\mathsf{id}_i} \neq \mathsf{Ctx}_{\mathsf{id}^*}$.

In both cases, $\mathcal{B}$ responds as in Phase 1.

Guess. $\mathcal{A}$ outputs $\gamma' \in \{0, 1\}$ and it wins if $\gamma' = \gamma$.

We refer to such an adversary $\mathcal{A}$ as an IND-ID-CCA2 adversary. We define the advantage (in the security parameter λ) of $\mathcal{A}$ in attacking an IBE scheme $\mathcal{E}$ as $\mathsf{Adv}_{\mathcal{E},\mathcal{A}}(\lambda) = |\Pr[\gamma' = \gamma] - 1/2|$.

Definition 8. *For a security parameter λ, let $t = t(\lambda)$, $q_{id} = q_{id}(\lambda)$, $q_{Ctx} = q_{Ctx}(\lambda)$, and $\epsilon = \epsilon(\lambda)$. We say that an IBE system $\mathcal{E}$ is $(t, q_{id}, q_{Ctx}, \epsilon)$-IND-ID-CCA2 secure if for any t-time IND-ID-CCA2 adversary $\mathcal{A}$ that makes at most q_{id} private key queries and at most q_{Ctx} decryption queries, we have $\mathsf{Adv}_{\mathcal{E},\mathcal{A}}(\lambda) \leq \epsilon$.*

Chosen-Plaintext Security. We define the chosen-plaintext security (IND-ID-CPA) for IBE systems as in above security game, except the adversary is not allowed to issue decryption queries. The adversary is still able to adaptively make private key queries.

Definition 9. *We say that an IBE system $\mathcal{E}$ is $(t, q_{\mathsf{id}}, \epsilon)$-IND-ID-CPA secure if $\mathcal{E}$ is $(t, q_{\mathsf{id}}, 0, \epsilon)$-IND-ID-CCA2 secure.*

Selective Security. A weaker and less realistic security model of IBE system, introduced in [26], is the selective security model in which adversary is required to commit to the challenge identity even before seeing the public parameters. We note that under computational assumptions with sub-exponential hardness, a selectively secure IBE is also adaptively secure through a standard "complexity leveraging" argument from [14]; however, complexity leveraging incurs a rather severe loss of tightness in the security reduction, causing the resulting scheme to suffer from a possibly large loss of efficiency per a similar argument as discussed in the introduction.

3 Signature Scheme with Tight Security

3.1 Constructions

$\mathsf{KeyGen}(1^\lambda)$ The key generation algorithm does the following.

1. Sample a matrix $\mathbf{A}$ along with a trapdoor basis of lattice $\Lambda_q^\perp(\mathbf{A})$ by $\mathsf{TrapGen}$.
2. Select matrices $\mathbf{A}_0$, $\mathbf{A}_1$, "PRF key" matrices $\mathbf{B}_1, \ldots, \mathbf{B}_k$, and "PRF input" matrices $\mathbf{C}_0$, $\mathbf{C}_1$ from $\mathbb{Z}_q^{n\times m}$ uniformly at random.
3. Select a secure pseudorandom function $\mathsf{PRF} : \{0,1\}^k \times \{0,1\}^t \to \{0,1\}$, which is expressed as a Boolean circuit with depth $d = d(\lambda)$, and a PRF key $K = s_1 s_2 \ldots s_k \xleftarrow{\$} \{0,1\}^k$.
4. Output the verification key and signing key as:

$$\mathsf{Vk} = \big(\mathbf{A}, \{\mathbf{A}_0, \mathbf{A}_1\}, \{\mathbf{B}_i\}_{i\in[k]}, \{\mathbf{C}_0, \mathbf{C}_1\}, \mathsf{PRF}, \beta\big), \quad \mathsf{Sk} = (\mathbf{T}_{\mathbf{A}}, K)$$

$\mathsf{Sign}(\mathsf{Vk}, \mathsf{Sk}, \mathsf{M})$ The signing algorithm takes as input the public verification key Vk, the signing key Sk and a message $\mathsf{M} = m_1 m_2 \ldots m_t \in \{0,1\}^t$. It does:

1. Compute $\mathbf{A}_{\mathsf{PRF},\mathsf{M}} = \mathsf{Eval}(\mathsf{PRF}, \{\mathbf{B}_i\}_{i\in[k]}, \mathbf{C}_{m_1}, \mathbf{C}_{m_2}, \ldots, \mathbf{C}_{m_t}) \in \mathbb{Z}_q^{n\times m}$.
2. Compute bit value $b = \mathsf{PRF}(K, \mathsf{M})$ and set $\mathbf{F}_{\mathsf{M},1-b} = \big[\mathbf{A} \mid \mathbf{A}_{1-b} - \mathbf{A}_{\mathsf{PRF},\mathsf{M}}\big]$.
3. Run $\mathsf{SampleLeft}$ to sample $\mathbf{d}_{\mathsf{M}} \in \mathbb{Z}^{2m}$ with distribution $D_{\Lambda_q^\perp(\mathbf{F}_{\mathsf{M},1-b}),s}$.
4. Output the signature $\mathsf{Sig} = \mathbf{d}_{\mathsf{M}}$.

$\mathsf{Ver}(\mathsf{Vk}, \mathsf{M}, \mathsf{Sig})$ The verification algorithm takes as input the verification key Vk, message M and the signature of M, verifies as follows:

1. Assume $\mathsf{Sig} = \mathbf{d}$. It checks if $\mathbf{d} \neq \mathbf{0}$ and $\|\mathbf{d}\| \leq s\sqrt{2m}$.
2. Check if $\mathbf{F}_{\mathsf{M},b}\mathbf{d} = \big[\mathbf{A} \mid \mathbf{A}_b - \mathbf{A}_{\mathsf{PRF},\mathsf{M}}\big]\,\mathbf{d} = \mathbf{0} \pmod q$ for $b = 0$ or 1.

If all above verifications pass, accept the signature; otherwise, reject.

3.2 Parameters Selection and Discussion

Let λ be the security parameter, we set $n = n(\lambda)$, let the message length be $t = t(\lambda)$, and let the circuit depth of PRF be d. To ensure we can run $\mathsf{TrapGen}$ in the Lemma 3, we set $m = O(n \log q)$. To run $\mathsf{SampleLeft}$ and $\mathsf{SampleRight}$ in the real scheme and simulation per Theorem 2, we set s sufficiently large such that $s > \|\tilde{\mathbf{T}}_{\mathbf{G}}\| \cdot m^{O(d)} \cdot \omega(\sqrt{\log m})$. So we set $s = O(m^{O(d)})$. To ensure the applicability of the average-case to worst-case reduction for SIS, we need $q \geq \beta \cdot \tilde{O}(\sqrt{n})$. So we set $q = O(n^{3/2} \log n)^{O(d)}$.

Note that the internal parameter d is the depth of the circuit which is used to compute the PRF in our constructions. Let l be the length of input of the PRF. If we instantiate our scheme by using efficient LWE-based PRFs from [9,8] and DDH-based PRFs from [51,30,45], then we are dealing with circuits in class TC^0 (which consists of polynomial-size circuits of constant depth $O(1)$). With such instantiations, we end up with a polynomial modulus q.

We also note that for the efficient PRFs instantiations within class TC^0 (such as the the ones from [9,30,45]), increasing the parameter size (length of input) of PRFs barely increases the asymptotic complexity of our scheme. Consider a instantiated PRF $\mathsf{PRF} : \{0,1\}^k \times \{0,1\}^t \to \{0,1\}$ where k is the bit length of the secret key. We use $k+2$ matrices to encode $k + t = l$ bits input of PRF, leading to $O(k)$ matrices over $\mathbb{Z}_q^{n\times m}$. These $O(k)$ matrices can be expressed by $O(kn^2 \log^2 q) = O(kn^2 \log^2 (n^{3/2} \log n))$ bits where the exponent that accounts for constant depth is absorbed by the big O symbol (recall $m = O(n \log q)$ and $q = O(n^{3/2} \log n)^{O(1)}$). Increasing k linearly will exponentially increase the security of PRF, compensating for any security loss in the PRF reduction very fast. If the circuit depth of PRF is constant, increasing k linearly only increases the bit length of public matrices linearly and does not affect q. Therefore, the tight reduction we obtain, from the schemes to SIS assumption and the security of PRFs does provide optimal parameter selection. This discussion applies also to our IBE scheme.

3.3 Security of the Signature Scheme

The security of our signature scheme is stated by the following theorem.

Theorem 3. *If the $(t_{\mathsf{SIS}}, \epsilon_{\mathsf{SIS}})$-$\mathsf{SIS}_{n,q,\beta,m}$ assumption holds and the PRF used in the signature scheme is $(t_{\mathsf{PRF}}, \epsilon_{\mathsf{PRF}})$-secure, the signature scheme is (t, q_s, ϵ)-EUF-CMA secure where $\epsilon_{\mathsf{SIS}} \geq \epsilon/2 - \epsilon_{\mathsf{PRF}} - \mathsf{negl}(\lambda)$, for some negligible statistical error $\mathsf{negl}(\lambda)$, and $\max(t_{\mathsf{PRF}}, t_{\mathsf{LWE}}) \leq t + O(q_s \cdot (T_S + N \cdot T_M))$ where q_s is the number of signing query, T_S is the maximum running time of sampling a Gaussian vector in $\mathbb{Z}^m$, N is the number of circuit gate of PRF, and T_M is the time of computing the multiplication between a n-by-m matrix and a m-by-m matrix in $\mathbb{Z}_q$.*

Proof. Consider the following security game between an adversary $\mathcal{A}$ and a simulator $\mathcal{B}$. Upon receiving a $\mathsf{SIS}_{n,q,\beta,m}$ challenge $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$, the challenger $\mathcal{B}$ prepares Vk as follows:

1. Select $k+4$ matrices $\mathbf{R}_{\mathbf{A}_0}$, $\mathbf{R}_{\mathbf{A}_1}$, $\{\mathbf{R}_{\mathbf{B}_i}\}_{i\in[k]}$, $\mathbf{R}_{\mathbf{C}_0}$, $\mathbf{R}_{\mathbf{C}_1} \xleftarrow{\$} \{1,-1\}^{m\times m}$.
2. Select a secure pseudorandom function $\mathsf{PRF} : \{0,1\}^k \times \{0,1\}^t \to \{0,1\}$ which is expressed as a Boolean circuit with depth d.
3. Select a PRF key $K = s_1 s_2 \ldots s_k \xleftarrow{\$} \{0,1\}^k$.
4. Set $\mathbf{A}_b = \mathbf{A}\mathbf{R}_{\mathbf{A}_b} + b\mathbf{G}$ and $\mathbf{C}_b = \mathbf{A}\mathbf{R}_{\mathbf{C}_b} + b\mathbf{G}$ for $b = 0, 1$.
5. Set $\mathbf{B}_i = \mathbf{A}\mathbf{R}_{\mathbf{B}_i} + s_i\mathbf{G}$ for $i \in [k]$.
6. Publish $\mathsf{Vk} = \big(\mathbf{A}, \{\mathbf{A}_0, \mathbf{A}_1\}, \{\mathbf{B}_i\}_{i\in[k]}, \{\mathbf{C}_0, \mathbf{C}_1\}, \mathsf{PRF}\big)$.

In the query phase, the adversary $\mathcal{A}$ adaptively issues messages for inquiring the corresponding signatures. Consider a message $\mathsf{M} = m_1 m_2 \ldots m_t \in \{0,1\}^t$. $\mathcal{B}$ does the following to prepare the signature:

1. Compute $\mathbf{A}_{\mathsf{PRF}} = \mathbf{A}\mathbf{R}_{\mathsf{PRF},\mathsf{M}} + \mathsf{PRF}(K, \mathsf{M})\mathbf{G} \in \mathbb{Z}_q^{n\times m}$ by $\mathsf{Eval}(\mathsf{PRF}, \{\mathbf{B}_i\}_{i\in[k]}, \mathbf{C}_{m_1}, \mathbf{C}_{m_2}, \ldots, \mathbf{C}_{m_t})$.
2. Let $b = \mathsf{PRF}(K, \mathsf{M})$, it sets

$$\begin{aligned}\mathbf{F}_{\mathsf{M},1-b} &= \big[\mathbf{A} \mid \mathbf{A}_{1-b} - \mathbf{A}_{\mathsf{PRF},\mathsf{M}}\big]\\ &= \big[\mathbf{A} \mid \mathbf{A}(\mathbf{R}_{\mathbf{A}_1} - \mathbf{R}_{\mathsf{PRF},\mathsf{M}}) + (1-2b)\mathbf{G}\big]\end{aligned}$$

   and runs $\mathsf{SampleRight}$ to generate the signature $\mathsf{Sig} = \mathbf{d}_{\mathsf{M}} \sim D_{\Lambda_q^{\perp}(\mathbf{F}_{\mathsf{M},1-b}),s}$.

Finally, $\mathcal{A}$ output a forgery $(\mathbf{d}^*, \mathsf{M}^*)$. Let $\mathsf{PRF}(K, \mathsf{M}^*) = b$. If $\|\mathbf{d}\| > s\sqrt{2m}$ or $\big[\mathbf{A} \mid \mathbf{A}_{1-b} - \mathbf{A}_{\mathsf{PRF},\mathsf{M}^*}\big]\,\mathbf{d}^* = 0 \pmod q$, $\mathcal{B}$ aborts. Otherwise, we have $\big[\mathbf{A} \mid \mathbf{A}_b - \mathbf{A}_{\mathsf{PRF},\mathsf{M}^*}\big]\,\mathbf{d}^* = 0 \pmod q$. Let $\mathbf{d}^* = [\mathbf{d}_1^\top \mid \mathbf{d}_2^\top]^\top \in \mathbb{Z}^{2m}$. $\mathcal{B}$ outputs $\mathbf{e} = \mathbf{d}_1 + (\mathbf{R}_{\mathbf{A}_b} - \mathbf{R}_{\mathsf{PRF},\mathsf{M}^*})\mathbf{d}_2$ as a solution for the $\mathsf{SIS}_{n,q,\beta,m}$ problem instance.

We show that Vk output by $\mathcal{B}$ has the correct distribution. In the real scheme, the matrix $\mathbf{A}$ is generated by $\mathsf{TrapGen}$. In the simulation, $\mathbf{A}$ is has uniform distribution in $\mathbb{Z}_q^{n\times m}$ as it comes from the SIS challenge. By the Lemma 3, $\mathbf{A}$ generated in the simulation has right distribution except a negligibly small statistical error $\mathsf{negl}(\lambda)$. Secondly, the matrices $\mathbf{A}$, $\{\mathbf{A}_0, \mathbf{A}_1\}$, $\{\mathbf{B}_i\}_{i\in[k]}$, and $\{\mathbf{C}_0, \mathbf{C}_1\}$ computed in the simulation have the distribution that is statistically close to uniform distribution in $\mathbb{Z}_q^{n\times m}$ by the special case (no leakage of the low-norm matrices) of the Lemma 1. In particular, the PRF secret key $\{s_i\}_{i\in[k]}$ is information-theoretically concealed by $\{\mathbf{B}_i\}_{i\in[k]}$.

In the query phase, for any message, instead of randomly picking the bit value b to select matrix $\mathbf{A}_{1-b}$, $\mathcal{B}$ computes b by PRF depending on the message and the secret PRF key. By the definition of PRFs, $\mathcal{A}$ has advantage ϵ_{PRF} in detecting this change. Meanwhile, the signatures replied to $\mathcal{A}$ have the correct distribution under the predefined conditions. Indeed, by the Theorem 2, for sufficient large Gaussian parameter s, the the distribution of signatures generated in the simulation by $\mathsf{SampleRight}$ is statistically close to $D_{\Lambda_q^{\perp}(\mathbf{F}_{\mathsf{M}}),s}$ where the distribution of signatures generated in the real scheme by $\mathsf{SampleLeft}$ is also statistically close to $D_{\Lambda_q^{\perp}(\mathbf{F}_{\mathsf{M}}),s}$.

In the forge phase, $\mathcal{A}$ will have at most ϵ_{PRF} advantage in predicting the bit value b based on its forgery. Therefore, if $\mathcal{A}$ can not distinguish PRF from random functions, it will randomly pick either of the matrices $\mathbf{A}_0$ or $\mathbf{A}_1$ to make a forgery. With $\frac{1}{2}$ chance it will pick the one that $\mathcal{B}$ will be able to use to solve the SIS problem. So we have $\epsilon_{\mathsf{SIS}} \geq \epsilon/2 - \epsilon_{\mathsf{PRF}} - \mathsf{negl}(\lambda)$.

To argue that $\mathbf{e} = \mathbf{d}_1 + (\mathbf{R}_{\mathbf{A}_1} - \mathbf{R}_{\mathsf{PRF},\mathsf{M}^*})\mathbf{d}_2$ is a valid solution of the $\mathsf{SIS}_{n,q,\beta,m}$ problem instance, we need to show $\mathbf{e}$ is sufficiently short, and non-zero except with negligible probability. First of all, we have

$$\begin{aligned}\left[\mathbf{A} \mid \mathbf{A}_b - \mathbf{A}_{\mathsf{PRF},\mathsf{M}^*}\right] \mathbf{d}^* &= \left[\mathbf{A} \mid \mathbf{A}(\mathbf{R}_{\mathbf{A}_b} - \mathbf{R}_{\mathsf{PRF},\mathsf{M}^*})\right] \mathbf{d}^* \\ &= \mathbf{A}\mathbf{d}_1 + \mathbf{A}(\mathbf{R}_{\mathbf{A}_b} - \mathbf{R}_{\mathsf{PRF},\mathsf{M}^*})\mathbf{d}_2 \\ &= \mathbf{A}\left(\mathbf{d}_1 + (\mathbf{R}_{\mathbf{A}_b} - \mathbf{R}_{\mathsf{PRF},\mathsf{M}^*})\mathbf{d}_2\right) \\ &= \mathbf{0} \pmod q\end{aligned}$$

By the Lemma 2, we have $\|\mathbf{R}_{\mathbf{A}_b}\| \leq 12\sqrt{m}$. By the analysis of noise growth of the key homomorphic computation, we have $\|\mathbf{R}_{\mathsf{PRF},\mathsf{M}^*}\| \leq m^{O(d)}$. So $\|\mathbf{R}_{\mathsf{PRF},\mathsf{M}^*} - \mathbf{R}_{\mathbf{A}_b}\|$ is still less than $m^{O(d)}$ by absorbing the constant by the big O notation. Since $\mathbf{d}_1, \mathbf{d}_2$ have distribution $D_{\mathbb{Z}^m,s}$ with condition $\mathbf{d} \in \Lambda_q^{\perp}(\mathbf{F}_{\mathsf{M},b})$, by the Lemma 4, $\mathbf{d}_1, \mathbf{d}_2 \leq s\sqrt{m}$. So $\|\mathbf{e}\| \leq \|\mathbf{d}_1\| + (\|\mathbf{R}_{\mathsf{PRF},\mathsf{M}^*}\| + \|\mathbf{R}_{\mathbf{A}_b}\|) \cdot \|\mathbf{d}_2\| \leq O(m)^{O(d)}$. Let $\beta = O(m)^{O(d)}$ is sufficient.

Let $\mathbf{R} = \mathbf{R}_{\mathsf{PRF},\mathsf{M}^*} - \mathbf{R}_{\mathbf{A}_b}$ (note $\mathbf{R} = \mathbf{0}$ happens with only negligible probability), we show $\mathbf{d}_1 \neq \mathbf{R} \cdot \mathbf{d}_2$ with all but negligible probability. Suppose $\mathbf{d}_2 \neq \mathbf{0}$, we have $\mathbf{e} \neq \mathbf{0}$ since $\mathbf{d} \neq \mathbf{0}$. On the other hand, we have $\mathbf{d}_2 = (d_1, \ldots, d_m)^\top \neq \mathbf{0}$ and, thus, at least one coordinate of $\mathbf{d}_2$, say d_j, is not 0. We write $\mathbf{R} = (\mathbf{r}_1, \ldots, \mathbf{r}_m)$ and so

$$\mathbf{R} \cdot \mathbf{d}_2 = \mathbf{r}_j \cdot d_j + \sum_{i=1, i \neq j}^{m} \mathbf{r}_i \cdot d_i$$

Observe that for the fixed message M^* on which $\mathcal{A}$ made the forgery, $\mathbf{R}$ (therefore $\mathbf{r}_j$) depends on the low-norm matrices $\mathbf{R}_{\mathbf{A}_0}, \mathbf{R}_{\mathbf{A}_1}, \{\mathbf{R}_{\mathbf{B}_i}\}_{i \in [k]}, \mathbf{R}_{\mathbf{C}_0}, \mathbf{R}_{\mathbf{C}_1}$ and the secret key of PRF. The only information about $\mathbf{r}_j$ for $\mathcal{A}$ is from the public matrices in Vk, i.e. $\{\mathbf{A}_0, \mathbf{A}_1\}, \{\mathbf{B}_i\}_{i \in [k]}, \{\mathbf{C}_0, \mathbf{C}_1\}$. So by the pigeonhole principle there is a (exponentially) large freedom to pick a value to $\mathbf{r}_j$ which is compatible with $\mathcal{A}$'s view, i.e. $\mathbf{A}\mathbf{r}'_j = \mathbf{A}\mathbf{r}''_j \pmod q$ for admissible (low-norm) $\mathbf{r}'_j, \mathbf{r}''_j$ where $\mathbf{r}'_j \neq \mathbf{r}''_j$. (In fact, here we have more freedom than the case in [21] where $\mathbf{R}$ is picked from $\{1, -1\}^{m \times m}$).

Finally, to answer one signing query, $\mathcal{B}$'s running time is bounded by $O(T_S + N \cdot T_M)$. So the total running time of $\mathcal{B}$ in the simulation is bounded by $O(Q(T_S + N \cdot T_M))$ This concludes the proof.

4 IBE with Tight Security

4.1 Construction with CPA Security

Setup(1^λ) The setup algorithm takes as input a security parameter λ. It does the following:

1. Sample a random matrix $\mathbf{A}$ along with a trapdoor basis $\mathbf{T}_{\mathbf{A}}$ of lattice $\Lambda_q^{\perp}(\mathbf{A})$ by running $\mathsf{TrapGen}$.
2. Select random matrices $\mathbf{A}_0$, $\mathbf{A}_1$, random "PRF key" matrices $\mathbf{B}_1, \ldots, \mathbf{B}_k$, and random "PRF input" matrices $\mathbf{C}_0$, $\mathbf{C}_1$ from $\mathbb{Z}_q^{n \times m}$ uniformly at random.

3. Select a random vector $\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^n$.
4. Select a secure pseudorandom function $\mathsf{PRF} : \{0,1\}^k \times \{0,1\}^t \to \{0,1\}$, which is expressed as a Boolean circuit with depth $d = d(\lambda)$, and a PRF key $K = s_1 s_2 \ldots s_k \xleftarrow{\$} \{0,1\}^k$.
5. Output the public parameters
$$\mathsf{Pub} = \big(\mathbf{A}, \{\mathbf{A}_0, \mathbf{A}_1\}, \{\mathbf{B}_i\}_{i \in [k]}, \{\mathbf{C}_0, \mathbf{C}_1\}, \mathbf{u}, \mathsf{PRF}\big)$$
and the master secret key $\mathsf{Msk} = (\mathbf{T_A})$.

KeyGen$(\mathsf{Pub}, \mathsf{Msk}, \mathsf{id})$ Upon an input identity $\mathsf{id}{=}x_1 x_2 \ldots x_t \in \{0,1\}^t$, the key generation algorithm does the following:

1. Compute $b = \mathsf{PRF}(K, \mathsf{id})$.
2. Deterministically compute $\mathbf{A}_{\mathsf{PRF},\mathsf{id}} \in \mathbb{Z}_q^{n \times m}$ as
$$\mathbf{A}_{\mathsf{PRF},\mathsf{id}} = \mathsf{Eval}(\mathsf{PRF}, \{\mathbf{B}\}_{i \in [k]}, \mathbf{C}_{x_1}, \mathbf{C}_{x_2}, \ldots, \mathbf{C}_{x_t})$$
3. Set $\mathbf{F}_{\mathsf{id},1-b} = \big[\mathbf{A} \mid \mathbf{A}_{1-b} - \mathbf{A}_{\mathsf{PRF},\mathsf{id}}\big] \in \mathbb{Z}_q^{n \times 2m}$.
4. Run $\mathsf{SampleLeft}$ to sample $\mathbf{d}_{\mathsf{id}}$ from the discrete Gaussian distribution $D_{\Lambda_q^{\mathbf{u}}(\mathbf{F}_{\mathsf{id},1-b}),s}$ hence $\mathbf{F}_{\mathsf{id},1-b}\mathbf{d}_{\mathsf{id}} = \mathbf{u} \pmod q$. Output $\mathsf{Sk}_{\mathsf{id}} = \mathbf{d}_{\mathsf{id}}$.

Encrypt$(\mathsf{Pub}, \mathsf{id}, \mathsf{Msg})$ To encrypt a message $\mathsf{Msg} \in \{0,1\}$ with respect to an identity $\mathsf{id} = x_1 x_2 \ldots x_t \in \{0,1\}^t$:

1. Compute $\mathbf{A}_{\mathsf{PRF},\mathsf{id}} \in \mathbb{Z}_q^{n \times m}$ as
$$\mathbf{A}_{\mathsf{PRF},\mathsf{id}} = \mathsf{Eval}(\mathsf{PRF}, \{\mathbf{B}_i\}_{i \in [k]}, \mathbf{C}_{x_1}, \mathbf{C}_{x_2}, \ldots, \mathbf{C}_{x_t})$$
2. Set $\mathbf{F}_{\mathsf{id},b} = \big[\mathbf{A} \mid \mathbf{A}_b - \mathbf{A}_{\mathsf{PRF},\mathsf{id}}\big] \in \mathbb{Z}_q^{n \times 2m}$ for $b = 0, 1$.
3. Select two random vectors $\mathbf{s}_0, \mathbf{s}_1 \xleftarrow{\$} \mathbb{Z}_q^n$.
4. Select two noise scalars $\nu_{0,0}, \nu_{1,0} \leftarrow D_{\mathbb{Z},\sigma_{\mathsf{LWE}}}$ and four noise vectors $\hat{\boldsymbol{\nu}}_{0,1}, \hat{\boldsymbol{\nu}}_{1,1} \leftarrow D_{\mathbb{Z}^m,\sqrt{2}\sigma_{\mathsf{LWE}}}$ and $\check{\boldsymbol{\nu}}_{0,1}, \check{\boldsymbol{\nu}}_{1,1} \leftarrow D_{\mathbb{Z}^m,\sigma}$ for sufficiently large σ_{LWE} and σ such that $\sigma > O(m^{O(d)}) \cdot \sigma_{\mathsf{LWE}}$.
5. Compute the ciphertext $\mathsf{Ctx}_{\mathsf{id}} = (c_{0,0}, \mathbf{c}_{0,1}, c_{1,0}, \mathbf{c}_{1,1})$ as:
$$\begin{cases} c_{0,0} &= \big(\mathbf{s}_0^\top \mathbf{u} + \nu_{0,0} + \mathsf{Msg}\lfloor q/2 \rfloor\big) \bmod q \\ \mathbf{c}_{0,1}^\top &= \big(\mathbf{s}_0^\top \mathbf{F}_{\mathsf{id},0} + [\hat{\boldsymbol{\nu}}_{0,1}^\top \mid \check{\boldsymbol{\nu}}_{0,1}^\top]\big) \bmod q \end{cases}$$
$$\begin{cases} c_{1,0} &= \big(\mathbf{s}_1^\top \mathbf{u} + \nu_{1,0} + \mathsf{Msg}\lfloor q/2 \rfloor\big) \bmod q \\ \mathbf{c}_{1,1}^\top &= \big(\mathbf{s}_1^\top \mathbf{F}_{\mathsf{id},1} + [\hat{\boldsymbol{\nu}}_{1,1}^\top \mid \check{\boldsymbol{\nu}}_{1,1}^\top]\big) \bmod q \end{cases}$$

Decrypt$(\mathsf{Pub}, \mathsf{Sk}_{\mathsf{id}}, \mathsf{Ctx}_{\mathsf{id}})$ The decryption algorithm uses the key $\mathbf{d}_{\mathsf{id}}$ to try to decrypt both $(c_{0,0}, \mathbf{c}_{0,1})$ and $(c_{1,0}, \mathbf{c}_{1,1})$ [3]. W.l.o.g., assume that $(c_{b,0}, \mathbf{c}_{b,1})$ is the correct ciphertext. The decryption algorithm computes
$$\tau = \big(c_{b,0} - \mathbf{c}_{b,1}^\top \mathbf{d}_{\mathsf{id}}\big) \bmod q$$
View τ as an integer in $(-q/2, q/2]$. If τ is closer to 0 than $\pm q/2$, the output is $\mathsf{Msg} = 0$. Otherwise, it is $\mathsf{Msg} = 1$.

[3] To ensure correct decryption, the message should contain some redundancy to weed out the incorrect ciphertext. It is a standard technique to encrypt multiple bits in

4.2 Correctness and Parameters Selection

Following the decryption algorithm, let $\mathbf{d}_{\mathsf{id}} = [\mathbf{d}_1^\top \mid \mathbf{d}_2^\top]^\top$. We have

$$\begin{aligned} \tau &= \left(c_{b,0} - \mathbf{c}_{b,1}^\top \mathbf{d}_{\mathsf{id}}\right) \bmod q \\ &= \left(\mathsf{Msg}\lfloor q/2 \rfloor + \nu_{b,0} - \hat{\boldsymbol{\nu}}_{0,1}^\top \mathbf{d}_1 - \check{\boldsymbol{\nu}}_{0,1}^\top \mathbf{d}_2\right) \bmod q \end{aligned}$$

Recall, the norm of $\mathbf{d}_1$ and $\mathbf{d}_2$ is bounded by $s\sqrt{m}$, and the norm of $\hat{\boldsymbol{\nu}}_{b,1}$ and $\check{\boldsymbol{\nu}}_{b,1}$ is bounded by $\sigma_{\mathsf{LWE}}\sqrt{2m}$ and $\sigma\sqrt{m}$ respectively, by Lemma 4. To ensure correctness of decryption, we need

$$\begin{aligned} |\tau| &= |c_{b,0} - \hat{\boldsymbol{\nu}}_{b,1}^\top \mathbf{d}_1 - \check{\boldsymbol{\nu}}_{0,1}^\top \mathbf{d}_2| \\ &\leq |c_{b,0}| + \|\hat{\boldsymbol{\nu}}_{0,1}\| \cdot \|\mathbf{d}_1\| + \|\hat{\boldsymbol{\nu}}_{0,1}\| \cdot \|\mathbf{d}_2\| \\ &\leq O(s \cdot \sigma_{\mathsf{LWE}} \cdot m^{O(d)}) \\ &\leq q/4 \end{aligned}$$

Accordingly, it is enough to set q such that $O(s \cdot \sigma_{\mathsf{LWE}} \cdot m^{O(d)}) \leq q/4$.

Parameter Selection. We now discuss a consistent parameter instantiation that achieves both correctness and security. We set the LWE dimension $n = n(d)$. To ensure we can run $\mathsf{TrapGen}$ in the Lemma 3, we set $m = O(n \log q)$. To make sure $\mathsf{SampleLeft}$ and $\mathsf{SampleRight}$ have the same distribution per Theorem 2, we need a sufficiently large Gaussian parameter $s > \|\tilde{\mathbf{T}}_{\mathbf{G}}\| \cdot m^{O(d)} \cdot \omega(\sqrt{\log m})$. We can set $s = O(m^{O(d)}) = O(n \log q)^{O(d)}$. To ensure the applicability of Regev's [53] LWE reductions from worst-case lattice problems, we set the Gaussian parameter of LWE noise distribution to be $\sigma_{\mathsf{LWE}} = \sqrt{n}$, therefore the LWE noise distribution is $(D_{\mathbb{Z},\sqrt{n}}) \bmod q$. (Here we may also wish to consider Peikert's [52] LWE reduction, although this may require to pick an unnecessarily large modulus q, to the detriment of size and efficiency.) Finally, to ensure correctness condition of decryption, we set $q = O(n^{3/2} \log n)^{O(d)}$.

Tight Reduction and Hardness of LWE. One feature of our IBE scheme (and the signature scheme it induces) is that depending on different PRFs circuits instantiations, the LWE assumption we use for our tight reduction may vary. More precisely, let B be the maximum magnitude of the LWE noise added to the ciphertext. We have seen the LWE modulus q depends on the circuit depth d of instantiated PRF. The hardness of the LWE problem depends on the ratio q/B. The LWE problem becomes easier when this ratio grows. In this regard, the appeal of our tight reduction varies: tight reduction to harder LWE problem is more preferable than tight reduction to easier LWE problem. This is true

Regev-like encryption, by replacing $\mathbf{u}$ with a matrix $\mathbf{U} \in \mathbb{Z}_q^{n \times z}$ in Pub with which we can now independently encrypt $z > 1$ bits without change to the security analysis. If hybrid encryption is used, the multiple bits can be used to encrypt a symmetric key *without* redundancy, deferring the integrity check to the symmetric realm where it can be performed at minimal cost.

particularly when one considers the average-case hardness of LWE to worst-case hardness of classic lattice problems, e.g. GapSVP and SIVP, reduction [53,52,22] where ratio q/B is smaller, the solutions for classic lattice problems are better. This fact indicates using lower depth PRFs is desirable, not only for better efficiency, but also for higher security.

In our case, thanks to the constant circuit-depth PRFs from [9,51,30,45], the LWE modulus q in our scheme is a polynomial, in the asymptotic sense (It could be actually large since the hidden constant of circuit depth could be large.). As a consequence, we are dealing with *asymptotically* small ratios q/B. On the other hand, we believe, however, that our tight reduction is still very valuable even for large ratio q/B. Firstly, it shows that, at the first time, we actually can eliminate the dependency between the number of adversary's queries and the security of lattice-based IBE scheme (as well as *short* lattice signature scheme). This is very important since the number of adversary's queries can be quite large, which will negatively impact the schemes' security seriously. Secondly, the average-case to worst-case reduction does provide some security confidence for the LWE assumption, but this is not the whole story. For certain parameters, many classic lattice problems are NP-hard. However, those parameters have no direct connection to lattice-based cryptography. (There is even evidence that the classic lattice problems with parameters relevant cryptography are not NP-hard.) On the other hand, the LWE problem (with various parameters) may be assured to be a hard problem in its own right. It has shown robustness against various attacks in a relatively long-term period. This has made LWE widely accepted as standard assumption and for use in cryptography. For instance, even for sub-exponentially large ratios $q/B = 2^{O(n^c)}$ where n is the LWE dimension and $0 < c < 1/2$, the LWE problem is still believed to be hard and leads to powerful cryptographic schemes which we were not able to obtain by other means, including fully homomorphic encryption, e.g. [23], attribute-based encryption for circuits, e.g. [36,19,24], and predicate encryption for circuits [40].

4.3 Proof of Security

The security of our IBE scheme with respect to the Definition 9 can be stated by the following theorem.

Theorem 4. *Suppose there exists a t-time $\mathsf{IND\text{-}ID\text{-}CPA}$ adversary $\mathcal{A}$ who makes q_{id} private key queries against our IBE scheme with advantage ϵ, there exists adversaries $\mathcal{B}_1$ with running time t_{PRF}, $\mathcal{B}_2$ with running time t_{LWE} respectively against the PRF PRF, and $\mathsf{LWE}_{n,q,\chi}$ (where χ is the distribution $(D_{\mathbb{Z},\sigma_{\mathsf{LWE}}}) \bmod q$) with respective advantages ϵ_{PRF} and ϵ_{LWE}, such that $\epsilon \leq 2(\epsilon_{\mathsf{PRF}} + \epsilon_{\mathsf{LWE}}) + \mathsf{negl}(\lambda)$ for some negligible function $\mathsf{negl}(\lambda)$ and $\max(t_{\mathsf{PRF}}, t_{\mathsf{LWE}}) \leq t + O\left(q_{\mathsf{id}} \cdot (T_S + N \cdot T_M)\right)$ where T_S is the maximum time of sampling a Gaussian vector from $\mathbb{Z}^m$, N is the number of circuit gate of PRF, and T_M is the time of computing the multiplication between a n-by-m matrix and a m-by-m matrix in $\mathbb{Z}_q$.*

We prove above theorem through a sequence of indistinguishable security games. The first game is identical to the $\mathsf{IND\text{-}ID\text{-}CPA}$ game. In the last game, the adversary has no advantage. We will show that a PPT adversary will not be able to distinguish the neighboring games which will prove that the adversary has only negligibly small advantage in wining the first (real) game.

Firstly, we define the following simulation algorithms $\mathsf{Sim.Setup}$, $\mathsf{Sim.KeyGen}$ and $\mathsf{Sim.Encrypt}$.

$\mathsf{Sim.Setup}(1^\lambda)$ On input of the security parameter λ, the algorithm does the following:

1. Select matrix $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$.
2. Select $k{+}4$ low-norm matrices $\mathbf{R}_{\mathbf{A}_0}, \mathbf{R}_{\mathbf{A}_1}, \{\mathbf{R}_{\mathbf{B}_i}\}_{i\in[k]}, \mathbf{R}_{\mathbf{C}_0}, \mathbf{R}_{\mathbf{C}_1} \xleftarrow{\$} \{1,-1\}^{m\times m}$.
3. Select a secure pseudorandom function $\mathsf{PRF} : \{0,1\}^k \times \{0,1\}^t \to \{0,1\}$ which is expressed as a Boolean circuit with depth $d = d(\lambda)$.
4. Select a uniformly random string $K = s_1 s_2 \dots s_k \xleftarrow{\$} \{0,1\}^k$.
5. Set $\mathbf{A}_b = \mathbf{A}\mathbf{R}_{\mathbf{A}_b} + b\mathbf{G}$ and $\mathbf{C}_b = \mathbf{A}\mathbf{R}_{\mathbf{C}_b} + b\mathbf{G}$ for $b = 0, 1$.
6. Set $\mathbf{B}_i = \mathbf{A}\mathbf{R}_{\mathbf{B}_i} + s_i\mathbf{G}$ for $i \in [k]$.
7. Select vector $\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^n$.
8. Publish $\mathsf{Pub} = (\ \mathbf{A},\ \{\mathbf{A}_0, \mathbf{A}_1\},\ \{\mathbf{B}_i\}_{i\in[k]},\ \{\mathbf{C}_0, \mathbf{C}_1\},\ \mathbf{u}, \mathsf{PRF})$

$\mathsf{Sim.KeyGen}(\mathsf{Pub}, \mathsf{Msk}, \mathsf{id})$ Upon an input identity $\mathsf{id} = x_1 x_2 \dots x_t \in \{0,1\}^t$, the algorithm uses the parameters generated from $\mathsf{Sim.Setup}$ to do the following:

1. Compute $\mathbf{A}_{\mathsf{PRF,id}} = \mathbf{A}\mathbf{R}_{\mathsf{PRF,id}} + \mathsf{PRF}(K, \mathsf{id})\mathbf{G} \leftarrow \mathsf{Eval}(\mathsf{PRF}, \{\mathbf{B}_i\}_{i\in[k]}, \mathbf{C}_{x_1}, \dots, \mathbf{C}_{x_t})$.
2. Let $\mathsf{PRF}(K, \mathsf{id}) = b \in \{0,1\}$. Set

$$\begin{aligned}\mathbf{F}_{\mathsf{id},1-b} &= \left[\mathbf{A} \mid \mathbf{A}_{1-b} - \mathbf{A}_{\mathsf{PRF,id}^*}\right] \\ &= \left[\mathbf{A} \mid \mathbf{A}(\mathbf{R}_{\mathbf{A}_{1-b}} - \mathbf{R}_{\mathsf{PRF,id}^*}) + (1-2b)\mathbf{G}\right].\end{aligned}$$

3. Run $\mathsf{SampleRight}$ to sample $\mathbf{d}_{\mathsf{id}} \in D_{\Lambda_q^{\mathbf{u}}(\mathbf{F}_{\mathsf{id},1-b}),s}$ as the private key $\mathsf{Sk}_{\mathsf{id}}$.

$\mathsf{Sim.Encrypt}(\mathsf{Pub}, \mathsf{id}^*, \mathsf{Msg})$ To encrypt a message $\mathsf{Msg}^* \in \{0,1\}$ with respect to an identity id^*:

1. Compute $b = \mathsf{PRF}(K, \mathsf{id}^*)$.
2. Set

$$\begin{aligned}\mathbf{F}_{\mathsf{id}^*,b} &= \left[\mathbf{A} \mid \mathbf{A}_b - \mathbf{A}_{\mathsf{PRF,id}^*}\right] \\ &= \left[\mathbf{A} \mid \mathbf{A}(\mathbf{R}_{\mathbf{A}_b} - \mathbf{R}_{\mathsf{PRF,id}^*})\right]\end{aligned}$$

   and

$$\begin{aligned}\mathbf{F}_{\mathsf{id}^*,1-b} &= \left[\mathbf{A} \mid \mathbf{A}_{1-b} - \mathbf{A}_{\mathsf{PRF,id}^*}\right] \\ &= \left[\mathbf{A} \mid \mathbf{A}(\mathbf{R}_{\mathbf{A}_{1-b}} - \mathbf{R}_{\mathsf{PRF,id}^*}) + (1-2b)\mathbf{G}\right].\end{aligned}$$

3. Select random vectors $\mathbf{s}_b, \mathbf{s}_{1-b} \xleftarrow{\$} \mathbb{Z}_q^n$.
4. Select noise scalars $\nu_{b,0}, \nu_{1-b,0} \leftarrow D_{\mathbb{Z},\sigma_{\mathsf{LWE}}}$.

5. Sample noise vectors $\mathbf{x}, \mathbf{y} \leftarrow D_{\mathbb{Z}^m,\sigma_{\mathsf{LWE}}}$ for sufficiently large Gaussian parameter σ_{LWE} ($\sigma_{\mathsf{LWE}} \geq \eta_\varepsilon(\mathbb{Z}^m)$ for some small $\varepsilon > 0$). Set $\hat{\boldsymbol{\nu}}_{b,1} = \mathbf{x} + \mathbf{y}$. By the Lemma 5, $\hat{\boldsymbol{\nu}}_{b,1}$ has distribution which is statistically close to $D_{\mathbb{Z}^m,\sqrt{2}\sigma_{\mathsf{LWE}}}$.
6. Let $\mathbf{R} = \mathbf{R}_b - \mathbf{R}_{\mathsf{PRF},\mathsf{id}^*}$ and $\mathbf{r}_i$ be the i-th column of $\mathbf{R}$. We sample the noise vector $\mathbf{z} = (z_1, z_2, \ldots, z_m) \in \mathbb{Z}^m$ with $z_i \leftarrow D_{\mathbb{Z},\sigma_{1,i}}$ for the Gaussian parameter $\sigma_{1,i} = \sqrt{\sigma^2 - 2(\|\mathbf{r}_i\| \cdot \sigma_{\mathsf{LWE}})^2}$. Set $\check{\boldsymbol{\nu}}_{b,1} = \mathbf{R}^\top \cdot (\mathbf{x} - \mathbf{y}) + \mathbf{z}$. By the Lemma 5, for sufficiently large $\sigma_{1,i}$ and σ_{LWE}, $\check{\boldsymbol{\nu}}_{b,1}$ has distribution which is statistically close to $D_{\mathbb{Z}^m,\sigma}$.
7. Select noise vectors $\hat{\boldsymbol{\nu}}_{1-b,1} \leftarrow D_{\mathbb{Z}^m,\sqrt{2}\sigma_{\mathsf{LWE}}}$, $\check{\boldsymbol{\nu}}_{1-b,1} \leftarrow D_{\mathbb{Z}^m,\sigma}$.
8. Set the challenge ciphertext $\mathsf{Ctx}_{\mathsf{id}^*} = (c_{b,0},\, \mathbf{c}_{b,1},\, c_{1-b,0},\, \mathbf{c}_{1-b,1})$ as:

$$\begin{cases} c_{b,0} = \left(\mathbf{s}_b^\top \mathbf{u} + \nu_{b,0} + \mathsf{Msg}\lfloor q/2 \rfloor\right) \bmod q \\ \mathbf{c}_{b,1}^\top = \left(\mathbf{s}_b^\top \mathbf{F}_{\mathsf{id}^*,b} + [\hat{\boldsymbol{\nu}}_{b,1}^\top \mid \check{\boldsymbol{\nu}}_{b,1}^\top]\right) \bmod q \end{cases}$$

$$\begin{cases} c_{1-b,0} = \left(\mathbf{s}_1^\top \mathbf{u} + \nu_{1-b,0} + \mathsf{Msg}\lfloor q/2 \rfloor\right) \bmod q \\ \mathbf{c}_{1-b,1}^\top = \left(\mathbf{s}_{1-b}^\top \mathbf{F}_{\mathsf{id}^*,1-b} + [\hat{\boldsymbol{\nu}}_{1-b,1}^\top \mid \check{\boldsymbol{\nu}}_{1-b,1}^\top]\right) \bmod q \end{cases}$$

Now we define a series of games and prove that the neighboring games are either statistically indistinguishable, or computationally indistinguishable.

Game 0 This is the real $\mathsf{IND\text{-}ID\text{-}CPA}$ game from the definition. All the algorithms are the same as the real scheme.

Game 1 This game is the same as **Game 0** except it runs $\mathsf{Sim.Setup}$ and $\mathsf{Sim.KeyGen}$ instead of Setup and KeyGen.

Game 2 This game is the same as **Game 1** except that the challenge ciphertext is generated by $\mathsf{Sim.Encrypt}$ instead of $\mathsf{Encrypt}$.

Game 3 This game is the same as **Game 2** except that during preparation of the challenge ciphertext for identity id^*, it samples $(c_{b,0}, \mathbf{c}_{b,1})$ uniformly random from $\mathbb{Z}_q \times \mathbb{Z}_q^{2m}$ for $b = \mathsf{PRF}(K, \mathsf{id}^*)$. Another part of the challenge ciphertext $(c_{1-b,0}, \mathbf{c}_{1-b,1})$ is computed by $\mathsf{Sim.Encrypt}$ as in **Game 2**.

Game 4 This game is the same as **Game 3** except for $b = \mathsf{PRF}(K, \mathsf{id}^*)$ it runs real encryption algorithm $\mathsf{Encrypt}$ to generate $(c_{1-b,0}, \mathbf{c}_{1-b,1})$ of the challenge ciphertext instead of using $\mathsf{Sim.Encrypt}$.

Game 5 This game is the same as **Game 4** except it runs Setup and KeyGen to generate Pub and private identity keys.

Game 6 This game is the same as **Game 5** except that for $b = \mathsf{PRF}(K, \mathsf{id}^*)$, the challenge ciphertext part $(c_{b,0}, \mathbf{c}_{b,1})$ is generated by $\mathsf{Encrypt}$ instead of choosing it randomly, and $(c_{1-b,0}, \mathbf{c}_{1-b,1})$ is chosen randomly.

Game 7 This game is the same as **Game 6** except that it runs $\mathsf{Sim.Setup}$ and $\mathsf{Sim.KeyGen}$ to generate Pub and private identity keys.

Game 8 This game is the same as **Game 7** except that for the bit value $b = \mathsf{PRF}(K, \mathsf{id}^*)$, it computes the challenge ciphertext $(c_{b,0}, \mathbf{c}_{b,1})$ by $\mathsf{Sim.Encrypt}$.

Game 9 This game is the same as **Game 8** except that the whole challenge ciphertext is sampled uniformly at random from the ciphertext space. Therefore, in **Game 5** the adversary has no advantage in wining the game.

In **Game** i, we let S_i be the event that $\gamma' = \gamma$ at the end of the game. The adversary's advantage in **Game** i is $|\Pr[S_i] - \frac{1}{2}|$. We prove the following lemmas to prove the Theorem 4.

Lemma 8. **Game 1** *and* **Game 0** *are statistically indistinguishable, so* $|\Pr[S_0] - \Pr[S_1]| \leq \mathsf{negl}(\lambda)$ *for some negligible function* $\mathsf{negl}(\lambda)$*.*

Proof. See Appendix A.1

Lemma 9. **Game 2** *and* **Game 1** *are statistically indistinguishable, so* $|\Pr[S_1] - \Pr[S_2]| \leq \mathsf{negl}(\lambda)$ *for some negligible function* $\mathsf{negl}(\lambda)$*.*

Proof. See Appendix A.2

Lemma 10. *If* $(t, \epsilon_{\mathsf{LWE}})$*-*$\mathsf{LWE}_{n,q,\chi}$ *assumption holds where* χ *stands for the distribution* $D_{\mathbb{Z},\sigma_{\mathsf{LWE}}}$ *reduced modulo* q*, then* $|\Pr[S_2] - \Pr[S_3]| \leq \epsilon_{\mathsf{LWE}}$*.*

Proof. See Appendix A.3

Lemma 11. $|\Pr[S_3] - \Pr[S_4]| = 0$*.*

Proof. See Appendix A.4

Lemma 12. **Game 5** *and* **Game 4** *are statistically indistinguishable, so* $|\Pr[S_4] - \Pr[S_5]| \leq \mathsf{negl}(\lambda)$ *for some negligible function* $\mathsf{negl}(\lambda)$*.*

Proof. See Appendix A.5

Lemma 13. *If the PRF* PRF *is* $(t, \epsilon_{\mathsf{PRF}})$*-secure, then* $|\Pr[S_5] - \Pr[S_6]| \leq 2\epsilon_{\mathsf{PRF}}$*.*

Proof. See Appendix A.6

Lemma 14. **Game 7** *and* **Game 6** *are statistically indistinguishable, so* $|\Pr[S_6] - \Pr[S_7]| \leq \mathsf{negl}(\lambda)$ *for some negligible function* $\mathsf{negl}(\lambda)$*.*

Proof. See Appendix A.7

Lemma 15. **Game 8** *and* **Game 7** *are statistically indistinguishable, so* $|\Pr[S_7] - \Pr[S_8]| \leq \mathsf{negl}(\lambda)$ *for some negligible function* $\mathsf{negl}(\lambda)$*.*

Proof. See Appendix A.8

Lemma 16. *If $(t, \epsilon_{\mathsf{LWE}})$-$\mathsf{LWE}_{n,q,\chi}$ assumption holds where χ stands for the distribution $D_{\mathbb{Z},\sigma_{\mathsf{LWE}}}$ reduced modulo q, then $|\Pr[S_8] - \Pr[S_9]| \leq \epsilon_{\mathsf{LWE}}$.*

Proof. See Appendix A.9

Now we prove the Theorem 4 by the established lemmas.

Proof. Based on the lemmas that show the difference between the sequence of games, we have $\epsilon = |\Pr[S_0] - 1/2| \leq 2(\epsilon_{\mathsf{PRF}} + \epsilon_{\mathsf{LWE}}) + \mathsf{negl}(\lambda)$ for some negligibly small statistical error $\mathsf{negl}(\lambda)$. The running time of $\mathcal{B}$ is dominated by answering q_{id}, $\mathcal{A}$'s private key generation queries. For answering one such query, $\mathcal{B}$ needs to apply the key-homomorphic algorithm on the circuit of PRF. Each gate requires a matrix multiplication and a addition over $\mathbb{Z}_q$. Besides that, $\mathcal{B}$ needs to sample Gaussian vectors for constructing the private keys (for private key queries) and constructing the challenge ciphertext (in the challenge phase). Therefore, for one query, $\mathcal{B}$ roughly runs $O(T_S + N \cdot T_M)$ time. For all q_{id} queries and constructing the challenge ciphertext, the total time is bounded by $O\left(q_{\mathsf{id}} \cdot (T_S + N \cdot T_M)\right)$. So if an adversary $\mathcal{A}$ has running time t, $\max(t_{\mathsf{LWE}}, t_{\mathsf{PRF}}) \leq t + O(q_{\mathsf{id}} \cdot (T_S + N \cdot T_M))$.

4.4 Adaptively CCA-Secure IBE and CCA-Secure PKE

The extension to tightly CCA-secure IBE and CCA-secure PKE scheme from our IBE construction is discussed in Appedix B.

5 Conclusions

In this paper, we propose a short adaptively secure lattice signature scheme and a "compact" adaptively secure IBE scheme in the standard model. Our constructions make use of PRFs in a novel way by combining several recent techniques in the area of lattice-based cryptography. The security of our signature and IBE scheme is tightly related to the conservative lattice assumptions SIS and LWE, respectively, and the security of an instantiated PRF, with a constant loss factor. By instantiating the existing efficient PRFs from lattice and number-theoretic assumptions which can be implemented by very shallow circuits, we obtain the first "almost" tightly secure lattice-based short signature/IBE scheme (relies on lattice assumptions only), and an adaptively secure IBE scheme with the tightest security reduction so far, i.e. with only $O(\log^2 \lambda)$ factor of security loss for the security parameter λ, based on a novel combination of lattice and number-theoretic assumptions.

The problem of constructing a tightly and adaptively secure IBE scheme from standard assumptions (in the sense that the security loss of reduction is a constant) remains open. Our work suggests that constructing tightly secure PRFs, which is another important open problem left by [30,45], would solve it. We leave as a fascinating open problem the question of employing similar (or different) techniques to construct compact and (almost) tightly secure signature and encryption schemes with increased expressiveness, such as hierarchical and attribute-based encryption scheme, or homomorphic signatures. Another interesting open question is to construct an almost tightly secure IBE scheme from LWE in the multi-user setting.

References

1. Abe, M., David, B., Kohlweiss, M., Nishimaki, R., Ohkubo, M.: Tagged one-time signatures: tight security and optimal tag size. In: Kurosawa, K., Hanaoka, G. (eds.) PKC 2013, LNCS, vol. 7778, pp. 312–331, Springer (2013)
2. Agrawal, S., Boneh, D., Boyen, X.: Efficient lattice (H)IBE in the standard model. In: Gilbert, H. (ed.) EUROCRYPT 2010, LNCS, vol. 6110, pp. 553–572. Springer (2010)
3. Ajtai, M.: Generating hard instances of lattice problems (extended abstract). In: STOC 1996, pp. 99–108. ACM (1996)
4. Alperin-Sheriff, J.: Short signatures with short public keys from homomorphic trapdoor functions. In: Katz, J. (ed.) PKC 2015, LNCS, vol. 9020, pp. 236–255. Springer (2015)
5. Apon, D., Fan, X., Liu, F.H.: Fully-secure lattice-based IBE as compact as PKE. Cryptology ePrint Archive, Report 2016/125 (2016)
6. Attrapadung, N., Hanaoka, G., Yamada, S.: A framework for identity-based encryption with almost tight security. In: Iwata, T., Cheon, J. (eds.) ASIACRYPT 2015, LNCS, vol. 9452, pp. 521–549. Springer (2015)
7. Bai, S., Langlois, A., Lepoint, T., Stehlé, D., Steinfeld, R.: Improved security proofs in lattice-based cryptography: using the Rényi divergence rather than the statistical distance. In: ASIACRYPT 2015, pp. 3–24. Springer (2015)
8. Banerjee, A., Peikert, C.: New and improved key-homomorphic pseudorandom functions. In: Garay, J., Gennaro, R. (eds.) CRYPTO 2014, LNCS, vol. 8616, pp. 353–370. Springer (2014)
9. Banerjee, A., Peikert, C., Rosen, A.: Pseudorandom functions and lattices. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012, LNCS, vol. 7237, pp. 719–737. Springer (2012)
10. Bellare, M., Rogaway, P.: Random oracles are practical: A paradigm for designing efficient protocols. In: CCS 1993, pp. 62–73. ACM (1993)
11. Blazy, O., Kakvi, S., Kiltz, E., Pan, J.: Tightly-secure signatures from chameleon hash functions. In: Katz, J. (ed.) PKC 2015, LNCS, vol. 9020, pp. 256–279. Springer (2015)
12. Blazy, O., Kiltz, E., Pan, J.: (Hierarchical) identity-based encryption from affine message authentication. In: Garay, J., Gennaro, R. (eds.) CRYPTO 2014, LNCS, vol. 8616, pp. 408–425. Springer (2014)
13. Böhl, F., Hofheinz, D., Jager, T., Koch, J., Seo, J.H., Striecks, C.: Practical signatures from standard assumptions. In Johansson, T., Nguyen, P. (eds.) EUROCRYPT 2013, LNCS, vol. 7881, pp. 461–485. Springer (2013)
14. Boneh, D., Boyen, X.: Efficient selective-id secure identity-based encryption without random oracles. In: Cachin, C., Camenisch, J. (eds.) EUROCRYPT 2004, LNCS, vol. 3027, pp. 223–238. Springer (2004)
15. Boneh, D., Boyen, X.: Secure identity based encryption without random oracles. In: Franklin, M. (ed.) CRYPTO 2004, LNCS, vol. 3152, pp. 443–459. Springer (2004)

16. Boneh, D., Boyen, X.: Short signatures without random oracles. In: Cachin, C., Camenisch, J. (eds.) EUROCRYPT 2004, LNCS, vol. 3027, pp. 56–73. Springer (2004)
17. Boneh, D., Canetti, R., Halevi, S., Katz, J.: Chosen-ciphertext security from identity-based encryption. SIAM J. Comput. 36(5), 1301–1328, (2006)
18. Boneh, D., Franklin, M.: Identity-based encryption from the weil pairing. In: Kilian, J. (ed.) CRYPTO 2001, LNCS, vol. 2139, pp. 213–229. Springer (2001)
19. Boneh, D., Gentry, C., Gorbunov, S., Halevi, S., Nikolaenko, V., Segev, G., Vaikuntanathan, V., Vinayagamurthy, D.: Fully key-homomorphic encryption, arithmetic circuit abe and compact garbled circuits. In: Nguyen, P., Oswald, E. (eds.) EUROCRYPT 2014, LNCS, vol. 8441, pp. 533–556. Springer (2014)
20. Boneh, D., Lynn, B., Shacham, H.: Short signatures from the weil pairing. Journal of Cryptology 17(4), 297–319 (2004)
21. Boyen, X.: Lattice mixing and vanishing trapdoors: A framework for fully secure short signatures and more. In: Nguyen, P., Pointcheval, D. (eds.) PKC 2010, LNCS, vol. 6056, pp. 499–517. Springer (2010)
22. Brakerski, Z., Langlois, A., Peikert, C., Regev, O., Stehlé, D.: Classical hardness of learning with errors. In: STOC 13. pp. 575–584. ACM (2013)
23. Brakerski, Z., Vaikuntanathan, V.: Efficient fully homomorphic encryption from (standard) LWE. In: FOCS 2011. pp. 97–106. IEEE (2011)
24. Brakerski, Z., Vaikuntanathan, V.: Circuit-ABE from LWE: Unbounded attributes and semi-adaptive security. Cryptology ePrint Archive, Report 2016/118 (2016)
25. Camenisch, J., Lysyanskaya, A.: Signature schemes and anonymous credentials from bilinear maps. In: Franklin, M. (eds.) CRYPTO 2004, LNCS, vol. 3152, pp. 56–72. Springer (2004)
26. Canetti, R., Halevi, S., Katz, J.: A forward-secure public-key encryption scheme. In: Biham, E. (ed.) EUROCRYPT 2003, LNCS, vol. 2656, pp. 255–271. Springer (2003)
27. Cash, D., Hofheinz, D., Kiltz, E., Peikert, C.: Bonsai trees, or how to delegate a lattice basis. Journal of Cryptology 25(4), 601–639 (2012)
28. Chen, J., Wee, H.: Fully, (almost) tightly secure IBE and dual system groups. In: Canetti, R., Garay, J. (eds.) CRYPTO 2013, LNCS, vol. 8043, pp. 435–460. Springer (2013)
29. Cramer, R., Shoup, V.: Signature schemes based on the strong RSA assumption. ACM Trans. Inf. Syst. Secur. 3(3), 161–185 (2000)
30. Döttling, N., Schröder, D.: Efficient pseudorandom functions via on-the-fly adaptation. In: CRYPTO 2015, LNCS, vol. 9215, pp. 329–350. Springer (2015)
31. Ducas, L., Micciancio, D.: Improved short lattice signatures in the standard model. In: Garay, J., Gennaro, R (eds.) CRYPTO 2014. LNCS, vol. 8616, pp. 335–352. Springer (2014)
32. Fischlin, M.: The Cramer-Shoup strong-RSA signature scheme revisited. In: Oswald, E., Fischer, M (eds.) PKC 2003, LNCS, vol 2567, pp. 116–129. Springer (2003)
33. Gennaro, R., Halevi, S., Rabin, T.: Secure hash-and-sign signatures without the random oracle. In: Stern, J (eds.) EUROCRYPT 1999, LNCS, vol. 1592, pp. 123–139. Springer (1999)
34. Gentry, C.: Practical identity-based encryption without random oracles. In: Vaudenay, S. (ed.) EUROCRYPT 2006, LNCS, vol. 4004, pp. 445–464. Springer (2006)
35. Gentry, C., Peikert, C., Vaikuntanathan, V.: Trapdoors for hard lattices and new cryptographic constructions. In: STOC 2008. pp. 197–206. ACM (2008)

36. Gorbunov, S., Vaikuntanathan, V., Wee, H.: Attribute-based encryption for circuits. In: STOC 2013. pp. 545–554. ACM (2013)
37. Gentry, C., Sahai, A., Waters, B.: Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In: Canetti, R., Garay, J. (eds.) CRYPTO 2013, LNCS, vol. 8042, pp. 75–92. Springer (2013)
38. Goh, E.J., Jarecki, S.: In: Biham, E (eds.) EUROCRYPT 2003, LNCS, vol. 2656, pp. 401–415. Springer (2003)
39. Goldreich, O., Goldwasser, S., Micali, S.: How to construct random functions. J. ACM 33(4), 792–807 (Aug 1986)
40. Gorbunov, S., Vaikuntanathan, V., Wee, H.: Predicate encryption for circuits from LWE. In: Gennaro, R., Robshaw, M. (eds.) CRYPTO 2015, LNCS, vol. 9216, pp. 503–523. Springer (2015)
41. Hofheinz, D., Jager, T.: Tightly Secure Signatures and Public-Key Encryption. In: Safavi-Naini, R., Canetti, R (eds.) CRYPTO 2012, LNCS, vol. 7417, pp. 590–607. Springer (2012)
42. Hofheinz, D., Koch, J., Striecks, C.: Identity-based encryption with (almost) tight security in the multi-instance, multi-ciphertext setting. In: Katz, J. (ed.) PKC 2015, LNCS, vol. 9020, pp. 799–822. Springer (2015)
43. Hohenberger, S., Waters, B.: Realizing Hash-and-Sign Signatures under Standard Assumptions. In: Joux, A (eds.) EUROCRYPT 2009, LNCS, vol. 5479, pp. 333–350. Springer (2009)
44. Hohenberger, S., Waters, B.: Short and stateless signatures from the RSA assumption. In: Halevi, S. (ed.) CRYPTO 2009, LNCS, vol. 5677, pp. 654–670. Springer (2009)
45. Jager, T.: Tightly-secure pseudorandom functions via work factor partitioning. Cryptology ePrint Archive, Report 2016/121 (2016)
46. Katz, J., Wang, N.: Efficiency improvements for signature schemes with tight security reductions. In: CCS 2003. pp. 155–164. CCS '03, ACM (2003)
47. Lewko, A., Waters, B.: New proof methods for attribute-based encryption: Achieving full security through selective techniques. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012, LNCS, vol. 7417, pp. 180–198. Springer (2012)
48. Lyubashevsky, V.: Lattice signatures without trapdoors. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012, LNCS, vol. 7237, pp. 738–755. Springer (2012)
49. Micciancio, D., Peikert, C.: Trapdoors for lattices: Simpler, tighter, faster, smaller. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012, LNCS, vol. 7237, pp. 700–718. Springer (2012)
50. Micciancio, D., Regev, O.: Worst-case to average-case reductions based on gaussian measures. SIAM J. Comput. 37(1), 267–302 (2007)
51. Naor, M., Reingold, O.: Number-theoretic constructions of efficient pseudo-random functions. J. ACM 51(2), 231–262 (2004)
52. Peikert, C.: Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In: STOC 2009. pp. 333–342. ACM (2009)
53. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. In: STOC 2005. pp. 84–93. ACM (2005)
54. Waters, B.: Efficient identity-based encryption without random oracles. In: Cramer, R. (ed.) EUROCRYPT 2005, LNCS vol. 3494, pp. 114–127. Springer (2005)
55. Waters, B.: Dual system encryption: Realizing fully secure ibe and hibe under simple assumptions. In: Halevi, S. (ed.) CRYPTO 2009, LNCS, vol. 5677, pp. 619–636. Springer (2009)

A Supplemental Proofs

A.1 Proof of Lemma 8

Proof (Proof of Lemma 8). We analyse the differences between **Game 0** and **Game 1**:

1. In **Game 0**, the matrix $\mathbf{A}$ is generated by $\mathsf{TrapGen}$, and in **Game 2**, the matrix $\mathbf{A}$ is chosen uniformly random. By the Lemma 3, the distributions of these two ways of constructing the matrix $\mathbf{A}$ are statistically close.
2. In **Game 0**, the matrices $\{\mathbf{A}_0, \mathbf{A}_1\}$, $\{\mathbf{B}_i\}_{i\in[k]}$, $\{\mathbf{C}_0, \mathbf{C}_1\}$ are chosen uniformly at random from $\mathbb{Z}_q^{n\times m}$. In **Game 2**, They are computed as $\mathbf{A}_b = \mathbf{A}\mathbf{R}_{\mathbf{A}_b} + b\mathbf{G}$, $\mathbf{C}_b = \mathbf{A}\mathbf{R}_{\mathbf{C}_b} + b\mathbf{G}$ for $b = 0, 1$, and $\mathbf{B}_i = \mathbf{A}\mathbf{R}_{\mathbf{B}_i} + s_i\mathbf{G}$ for $i \in [k]$ for random and secret low-norm matrices $\mathbf{R}_{\mathbf{A}_0}$, $\mathbf{R}_{\mathbf{A}_1}$, $\{\mathbf{R}_{\mathbf{B}_i}\}_{i\in[k]}$, $\mathbf{R}_{\mathbf{C}_0}, \mathbf{R}_{\mathbf{C}_1}$ from $\{1, -1\}^{m\times m}$. By the special case (no leakage of the low-norm matrices) of the Lemma 1, the distributions of these two ways of generating these public matrices are statistically close. In particular, the PRF secret key $\{s_i\}_{i\in[k]}$ is information-theoretically concealed by $\{\mathbf{B}_i\}_{i\in[k]}$.
3. We note that in both **Game 0** and **Game 1**, the use of $\mathbf{A}_0$ or $\mathbf{A}_1$ of the key generation algorithms is decided by $b = \mathsf{PRF}(K, \mathsf{id})$. For a private key query on id in **Game 2**, let

$$\begin{aligned}\mathbf{F}_{\mathsf{id},1-b} &= \left[\mathbf{A} \mid \mathbf{A}_{1-b} - \mathbf{A}_{\mathsf{PRF},\mathsf{id}}\right] \\ &= \left[\mathbf{A} \mid \mathbf{A}(\mathbf{R}_{\mathbf{A}_{1-b}} - \mathbf{R}_{\mathsf{PRF},\mathsf{id}}) + (1-2b)\mathbf{G}\right].\end{aligned}$$

   Note that the publicly known trapdoor of $\Lambda_q^{\perp}(\mathbf{G})$ is also a trapdoor of $\Lambda_q^{\perp}((1-2b)\mathbf{G})$. In **Game 1**, the identity key $\mathbf{d}_{\mathsf{id}} \in \Lambda_q^{\mathbf{u}}(\mathbf{F}_{\mathsf{id},1-b})$ is generated by $\mathsf{SampleLeft}$ with the trapdoor basis $\mathbf{T}_{\mathbf{A}}$ of $\Lambda_q^{\perp}(\mathbf{A})$. In **Game 2**, $\mathbf{d}_{\mathsf{id}}$ is generated by $\mathsf{SampleRight}$ with the trapdoor of $\Lambda_q^{\perp}((1-2b)\mathbf{G})$. By the Theorems 1 and 2, for sufficient large Gaussian parameter s, the identity key $\mathbf{d}_{\mathsf{id}}$ will have the same distribution $D_{\Lambda_q^{\mathbf{u}}(\mathbf{F}_{\mathsf{id},1-b}),s}$ up to a negligibly small statistical difference.

Summing up, the distributions of **Game 0** and **Game 1** are statistically close, and thus $|\Pr[S_0] - \Pr[S_1]| \leq \mathsf{negl}(\lambda)$ for some negligible function $\mathsf{negl}(\lambda)$.

A.2 Proof of Lemma 9

Proof (Proof of Lemma 9). Let $\mathbf{R} = \mathbf{R}_{\mathbf{A}_b} - \mathbf{R}_{\mathsf{PRF},\mathsf{id}^*}$ in the $\mathsf{Sim.Encrypt}$ algorithm. The difference between **Game 1** and **Game 2** is the way of generating the challenge ciphertext. In **Game 1**, the challenge ciphertext is generated by $\mathsf{Encrypt}$, and the noise vectors are sampled from some discrete Gaussian distributions that are independent of Pub. In **Game 2** the challenge ciphertext is generated by $\mathsf{Sim.Encrypt}$, and $\mathbf{R}$, where $\mathbf{R}$ is computed from $\mathbf{R}_{\mathbf{A}_0}, \mathbf{R}_{\mathbf{A}_1}, \{\mathbf{R}_{\mathbf{B}_i}\}_{i\in[k]}, \mathbf{R}_{\mathbf{C}_0}, \mathbf{R}_{\mathbf{C}_1}$, PRF's key K, and id^*.

Firstly, we note by the construction of the challenge ciphertext in **Game 2**,

$$\begin{aligned}\mathbf{c}_{b,1}^\top &= \left(\mathbf{s}_b^\top \mathbf{F}_{\mathsf{id}^*,b} + [\hat{\boldsymbol{\nu}}_{b,1}^\top \mid \check{\boldsymbol{\nu}}_{b,1}^\top]\right) \bmod q \\ &= \mathbf{s}_0^\top \left[\mathbf{A}|\mathbf{A}(\mathbf{R}_{\mathbf{A}_b} - \mathbf{R}_{\mathsf{PRF},\mathsf{id}^*})\right] +[(\mathbf{x}+\mathbf{y})^\top \mid \mathbf{R}(\mathbf{x}-\mathbf{y})^\top + \mathbf{z}^\top]\big) \bmod q \\ &= \left(\mathbf{s}_0^\top \left[\mathbf{A}|\mathbf{A}\mathbf{R}\right] + [(\mathbf{x}+\mathbf{y})^\top|\mathbf{R}(\mathbf{x}-\mathbf{y})^\top + \mathbf{z}^\top]\right) \bmod q\end{aligned}$$

By the Lemma 1 (the generalised left-over hash lemma), with $\mathbf{R}$ appearing in the challenge ciphertext, the public matrices $\mathbf{A}_0, \mathbf{A}_1, \{\mathbf{B}_i\}_{i\in[k]}, \mathbf{C}_0, \mathbf{C}_1$ still have distribution which is statistically close to the uniform distribution on $\mathbb{Z}_q^{n\times m}$.

Secondly, we argue that the noise vectors have correct distributions. Since the Gaussian vectors $\mathbf{x}, \mathbf{y}$ have distribution $D_{\mathbb{Z}^m,\sigma_{\mathsf{LWE}}}$, for sufficiently large σ_{LWE}, $\hat{\boldsymbol{\nu}}_{b,1} = \mathbf{x}+\mathbf{y}$ and $\mathbf{x}-\mathbf{y}$ essentially have distribution $D_{\mathbb{Z}^m,\sqrt{2}\sigma_{\mathsf{LWE}}}$ by the Lemma 5. Again, by the Lemma 5, with adding the "smoothing" error $\mathbf{z}$, the distribution of $\check{\boldsymbol{\nu}}_{b,1} = \mathbf{R}^\top\cdot(\mathbf{x}-\mathbf{y})+\mathbf{z}$ is statistically close to the Gaussian distribution $D_{\mathbb{Z}^m,\sigma}$.

Finally, we note that by the constructions, the noise vectors $\hat{\boldsymbol{\nu}}_{b,1}$ and $\check{\boldsymbol{\nu}}_{b,1}$ are statistically uncorrelated. Since the covariance of two random variables according to the Gaussian distributions of $\hat{\boldsymbol{\nu}}_{b,1}$, $\check{\boldsymbol{\nu}}_{b,1}$ is set to be 0 by the simulation.

Summing up, **Game 1** and **Game 2** are statistically indistinguishable and the lemma follows.

A.3 Proof of Lemma 10

Proof (Proof of Lemma 10). We show a simulation algorithm $\mathcal{B}$ that uses its LWE challenge to simulate either **Game 2** or **Game 3** for an adversary $\mathcal{A}$. At the beginning, $\mathcal{B}$ receives its LWE challenge $(\mathbf{W}, \mathbf{v}) \in \mathbb{Z}_q^{n\times m} \times \mathbb{Z}_q^m$ and $(\mathbf{w}, v) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$ which is either from $\mathcal{O}_\$$ or $\mathcal{O}_\mathbf{s}$.

Setup. $\mathcal{B}$ prepares the public parameters for $\mathcal{A}$ as follows:

1. Set $\mathbf{A} \leftarrow \mathbf{W}$ and $\mathbf{u} \leftarrow \mathbf{v}$. We note $\mathbf{A}, \mathbf{u}$ have uniform distribution.
2. Set other public parameters as **Game 2**.

Phase 1. $\mathcal{B}$ answers private key queries like **Game 2**.
Challenge. $\mathcal{B}$ prepares the challenge ciphertext of identity id^* as follows.

1. Let $b = \mathsf{PRF}(K, \mathsf{id}^*)$. $\mathcal{B}$ sets

$$\begin{aligned}\mathbf{F}_{\mathsf{id}^*,1-b} &= \left[\mathbf{A} \mid \mathbf{A}_{1-b} - \mathbf{A}_{\mathsf{PRF},\mathsf{id}^*}\right] \\ &= \left[\mathbf{A} \mid \mathbf{A}(\mathbf{R}_{\mathbf{A}_{1-b}} - \mathbf{R}_{\mathsf{PRF},\mathsf{id}^*}) + (1-2b)\mathbf{G}\right]\end{aligned}$$

2. Let $\mathbf{R} = \mathbf{R}_{\mathbf{A}_0} - \mathbf{R}_{\mathsf{PRF},\mathsf{id}^*}$. To construct $(c_{b,0}, \mathbf{c}_{b,1})$, $\mathcal{B}$ samples $\mathbf{y} \leftarrow D_{\mathbb{Z}^m,\sigma_{\mathsf{LWE}}}$ and $\mathbf{z}$ like $\mathsf{Sim.Encrypt}$. Then it sets

$$\begin{cases} c_{b,0} = (v + \mathsf{Msg}^*\lfloor q/2\rfloor) \bmod q \\ \mathbf{c}_{b,1}^\top = \left([\mathbf{v}^\top|\mathbf{v}^\top\mathbf{R}] + [\mathbf{y}^\top| -\mathbf{y}^\top\mathbf{R} + \mathbf{z}^\top]\right) \bmod q \end{cases}$$

3. $\mathcal{B}$ sets $(c_{1-b,0}, \mathbf{c}_{1-b,1})$ the same as **Game 2**.

Phase 2. $\mathcal{B}$ replies the private key queries as in **Game 2**.

Guess. Finally, $\mathcal{A}$ outputs whether it is interacting with **Game 2** or **Game 3**. If $\mathcal{A}$ says **Game 2**, $\mathcal{B}$ decides its LWE challenge is from $\mathcal{O}_{\mathbf{s}}$. Otherwise, $\mathcal{B}$ decides the LWE challenge is from $\mathcal{O}_{\$}$.

If $\mathcal{B}$ gets the LWE challenge from the oracle $\mathcal{O}_{\mathbf{s}}$, there exists a secret vector $\mathbf{s} \in \mathbb{Z}_q^n$, a noise scalar x with distribution $D_{\mathbb{Z},\sigma_{\mathsf{LWE}}}$, a noise vector $\mathbf{x} \in \mathbb{Z}^m$ with distribution $D_{\mathbb{Z}^m,\sigma_{\mathsf{LWE}}}$ such that $\mathbf{v}^\top = \mathbf{s}^\top \mathbf{A} + \mathbf{x}^\top$ and $v = \mathbf{s}^\top \mathbf{w} + x$. Rewrite the ciphertext we have

$$\begin{aligned} c_{b,0} &= (v + \mathsf{Msg}^* \lfloor q/2 \rfloor) \bmod q \\ &= \big(\mathbf{s}^\top \mathbf{w} + x + \mathsf{Msg}^* \lfloor q/2 \rfloor\big) \bmod q \\ &= \big(\mathbf{s}_b^\top \mathbf{u} + \nu_{b,0} + \mathsf{Msg}^* \lfloor q/2 \rfloor\big) \bmod q \end{aligned}$$

and

$$\begin{aligned} \mathbf{c}_{b,1}^\top &= \big([\mathbf{v}^\top | \mathbf{v}^\top \mathbf{R}] + [\mathbf{y}^\top | -\mathbf{y}^\top \mathbf{R} + \mathbf{z}^\top]\big) \bmod q \\ &= \big([\mathbf{s}^\top \mathbf{A} + \mathbf{x}^\top | (\mathbf{s}^\top \mathbf{A} + \mathbf{x}^\top)\mathbf{R}] + [\mathbf{y}^\top | -\mathbf{y}^\top \mathbf{R} + \mathbf{z}^\top]\big) \bmod q \\ &= \big(\mathbf{s}^\top [\mathbf{A} | \mathbf{A}\mathbf{R}] + [\mathbf{x}^\top + \mathbf{y}^\top | (\mathbf{x}^\top - \mathbf{y}^\top)\mathbf{R} + \mathbf{z}^\top]\big) \bmod q \\ &= \big(\mathbf{s}_b^\top \mathbf{F}_{\mathsf{id}^*,b} + [\hat{\boldsymbol{\nu}}_{b,1}^\top | \check{\boldsymbol{\nu}}_{b,1}^\top]\big) \bmod q \end{aligned}$$

They are valid challenge ciphertext parts in **Game 2**. Therefore, in this case $\mathcal{B}$ simulates **Game 2** for $\mathcal{A}$. On the other hand, if $\mathcal{B}$ gets samples from $\mathcal{O}_{\$}$, $(c_{b,0}, \mathbf{c}_{b,1})$ constructed above will be random, which is the case of **Game 3**, and $\mathcal{B}$ simulates **Game 3**. $|\Pr[S_2] - \Pr[S_3]| \leq \epsilon_{\mathsf{LWE}}$ follows.

A.4 Proof of Lemma 11

Proof (Proof of Lemma 11). Note for generating $(c_{1-b,0}, \mathbf{c}_{1-b,1})$ of the challenge ciphertext, Encrypt and Sim.Encrypt behave the same. So **Game 4** and **Game 3** are the same.

A.5 Proof of Lemma 12

Proof (Proof of Lemma 12). The proof is essentially the same as the proof for the Lemma 8. We omit the details.

A.6 Proof of Lemma 13

Proof (Proof of Lemma 13). We recall the difference between **Game 6** and **Game 5**. let $b = \mathsf{PRF}(K, \mathsf{id}^*)$ for the challenge identity id^*. In **Game 5**, the ciphertext component $(c_{b,0}, \mathbf{c}_{b,1})$ is uniformly random and $(c_{1-b,0}, \mathbf{c}_{1-b,1})$ is computed by Encrypt. In **Game 6**, the ciphertext component $(c_{b,0}, \mathbf{c}_{b,1})$ is computed by Encrypt and $(c_{1-b,0}, \mathbf{c}_{1-b,1})$ is uniformly random. To prove the indistinguishably between **Game 6** and **Game 5**, three additional security games are added.

Firstly we define **Game 5.1** which is same as **Game 5** except that it samples $b \xleftarrow{\$} \{0,1\}$ to select matrix $\mathbf{A}_b$ for generating private keys and challenge ciphertext instead of using PRF to compute it. Also, if same identity is queried multiple times, the same bit b will be used (For simulation, we simply let the simulator keep a state remembering the bit for each identity.). Obviously, a distinguisher between **Game 5** and **Game 5.1** leads to a attacker for PRF. So $|\Pr[S_5] - \Pr[S_{5.1}]| \leq \epsilon_{\mathsf{PRF}}$.

Secondly, we define **Game 5.2** which is the same as **Game 5.1** except for randomly sampled bit b for id^*, it runs $\mathsf{Encrypt}$ to produce $(c_{b,0}, \mathbf{c}_{b,1})$ and samples $(c_{1-b,0}, \mathbf{c}_{1-b,1})$ uniformly random from $\mathbb{Z}_q \times \mathbb{Z}_q^{2m}$. While here b is uniformly random, we must have $|\Pr[S_{5.1}] - \Pr[S_{5.2}]| = 0$.

Finally, as **Game 6** is the same as **Game 5.2** except the bit value b is computed via PRF in key generation query phase and challenge phase, so we have $|\Pr[S_{5.2}] - \Pr[S_6]| \leq \epsilon_{\mathsf{PRF}}$.

$|\Pr[S_5] - \Pr[S_6]| \leq 2\epsilon_{\mathsf{PRF}}$ follows.

A.7 Proof of Lemma 14

Proof (Proof of Lemma 14). The proof is essentially the same as the proof for Lemma 8. We omit the details.

A.8 Proof of Lemma 15

Proof (Proof of Lemma 15). The proof is essentially the same as the proof for Lemma 9. We omit the details.

A.9 Proof of Lemma 16

Proof (Proof of Lemma 16). The proof is essentially the same as the proof for Lemma 10. We omit the details.

B Adaptively CCA-Secure IBE and CCA-Secure PKE

Boneh at al. [17] showed a $\ell + 1$-depth CPA-secure Hierarchical IBE (HIBE) scheme ($\ell \geq 0$) can be tightly transferred into an ℓ-depth CCA-secure HIBE scheme with small additional overhead (known as the BCHK transformation). In particularly, a 1-depth HIBE scheme is an IBE scheme and a 0-depth HIBE scheme is a public-key encryption scheme PKE. Generally, in HIBE, identities are arranged in a directed tree. A user with identity of a father node can issue private keys for the users with identities of children nodes. This process is called delegation. Ideally, we would like to have HIBE schemes supporting identity trees with polynomial depth. Unfortunately, directly applying our technique will result in an HIBE scheme with only log-depth identity tree. On the other hand, our technique particularly works for 2-depth HIBE scheme. So by applying the

BCHK transformation, we obtain a IND-ID-CCA2 secure IBE scheme from the 2-depth IND-ID-CPA HIBE scheme and a IND-CCA2 secure PKE scheme from our IND-ID-CPA secure IBE scheme[4].

[4] This transformation does not require us to add new computational assumptions. The SIS assumption, which is weaker than the LWE assumption, is enough.