Efficient Lattice (H)IBE in the Standard Model*

Shweta Agrawal[1], Dan Boneh[2,**], and Xavier Boyen[3]

[1] University of Texas, Austin
[2] Stanford University
[3] Université de Liège, Belgium

Abstract. We construct an efficient identity based encryption system based on the standard learning with errors (LWE) problem. Our security proof holds in the standard model. The key step in the construction is a family of lattices for which there are two distinct trapdoors for finding short vectors. One trapdoor enables the real system to generate short vectors in all lattices in the family. The other trapdoor enables the simulator to generate short vectors for all lattices in the family except for one. We extend this basic technique to an adaptively-secure IBE and a Hierarchical IBE.

1 Introduction

Identity-Based Encryption (IBE) provides a public-key encryption mechanism where a public key is an arbitrary string such as an email address or a telephone number. The corresponding private key can only be generated by a Private-Key Generator (PKG) who has knowledge of a master secret. Identity-based encryption was first proposed by Shamir [28], however, it is only recently that practical implementations were proposed. Boneh and Franklin [8] define a security model for identity-based encryption and give a construction based on the Bilinear Diffie-Hellman (BDH) problem. Cocks [13] describes a construction using quadratic residues modulo a composite (see also [9]) and Gentry et al. [16] give a construction using lattices. The security of all these systems requires cryptographic hash functions that are modeled as random oracles.

For pairing-based systems, the structure of pairing groups enabled several secure IBE systems in the standard model [11,6,7,31,17,32]. For systems based on quadratic residuosity it is still not known how to build a secure IBE in the standard model.

In this paper we focus on lattice-based IBE. Cash et al. [12], Peikert [24] and Agrawal et al. [3] recently showed how to construct secure IBE in the standard model from the learning with errors (LWE) problem [27]. Their constructions view an identity as a sequence of bits and then assign a matrix to each bit. The resulting systems, while quite elegant, are considerably less efficient than the underlying random-oracle system of [16] on which they are built.

* A full version of this paper is available at [1].

** Supported by NSF and the Packard Foundation.

H. Gilbert (Ed.): EUROCRYPT 2010, LNCS 6110, pp. 553–572, 2010.

1.1 Our Results

We construct a lattice-based IBE in the standard model whose performance is comparable to the performance of the random-oracle system from [16]. In particular, we process identities as one chunk rather than bit-by-bit resulting in lattices whose dimension is similar to those in the random oracle system.

Lattices in our system are built from two parts called "right" and "left" lattices. A trapdoor for the left lattice is used as the master secret in the real system and enables one to generate private keys for all identities. A trapdoor for the right lattice is only used in the proof of selective security and enables the simulator to generate private keys for all identities except for one. We use a "low norm" randomization matrix R to ensure that an attacker cannot distinguish between the real world and a simulation.

In pairing-based IBE systems one uses large groups G and therefore identities can be encoded as integers in the range $1 \dots |G|$. In contrast, lattice systems are typically defined over a relatively small field $\mathbb{Z}_q$ and consequently encoding identities as integers in $1 \dots q$ would result in too few identities for the system. Instead, we represent identities as matrices in $\mathbb{Z}_q^{n \times n}$ for some n. More precisely, we represent identities as elements in $\mathbb{Z}_q^n$ (for a total of q^n identities) and then use an encoding function $H : \mathbb{Z}_q^n \to \mathbb{Z}_q^{n \times n}$ to map identities to matrices. Our security proof requires that for all $\mathsf{id}_1 \neq \mathsf{id}_2$ the matrix $H(\mathsf{id}_1) - H(\mathsf{id}_2) \in \mathbb{Z}_q^{n \times n}$ is invertible. We present an encoding function H that has this property and expect this encoding to be useful in other lattice-based constructions. A similar function H was developed by Cramer and Damgard [14] in an entirely different context.

Full IBE. In the full version of the paper [1] we show that our base construction extends to an adaptively-secure IBE using a lattice analog of the Waters IBE [31]. Our base construction requires that the underlying field $\mathbb{Z}_q$ satisfy $q > Q$ where Q is the number of private key queries issued by the adversary. This requirement can be relaxed using the framework of Boyen [10].

Hierarchical IBE (HIBE). In the full version of the paper [1] we show how to extend our base IBE to an HIBE using the basis delegation technique from [12,24]. The construction assigns a matrix to each level of the hierarchy and the resulting lattice dimension is linear in the recipient identity's depth. Since we do not process identities bit-by-bit we obtain an efficient HIBE where the lattice dimension is much smaller than in [12,24]. We note that a recent result of [2] uses a different basis delegation mechanism to construct an improved HIBE where the lattice dimension is fixed for the entire hierarchy.

2 Preliminaries

Notation. Throughout the paper we say that a function $\epsilon : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is negligible if $\epsilon(n)$ is smaller than all polynomial fractions for sufficiently large n. We say that an event happens with overwhelming probability if it happens with

probability at least $1 - \epsilon(n)$ for some negligible function ϵ. We say that integer vectors $v_1, \ldots, v_n \in \mathbb{Z}^m$ are $\mathbb{Z}_q$-linearly independent if they are linearly independent when reduced modulo q.

2.1 IBE and Hierarchical IBE

Recall that an Identity-Based Encryption system (IBE) consists of four algorithms [28,8]: Setup, Extract, Encrypt, Decrypt. The Setup algorithm generates system parameters, denoted by PP, and a master key MK. The Extract algorithm uses the master key to extract a private key corresponding to a given identity. The encryption algorithm encrypts messages for a given identity (using the system parameters) and the decryption algorithm decrypts ciphertexts using the private key.

In a Hierarchical IBE [20,18], identities are vectors, and there is a fifth algorithm called Derive. A vector of dimension ℓ represents an identity at depth ℓ. Algorithm Derive takes as input an identity $\mathsf{id} = (\mathsf{I}_1, \ldots, \mathsf{I}_\ell)$ at depth ℓ and the private key $\mathsf{SK}_{\mathsf{id}|\ell-1}$ of the parent identity $\mathsf{id}_{|\ell-1} = (\mathsf{I}_1, \ldots, \mathsf{I}_{\ell-1})$ at depth $\ell - 1 \geq 0$. It outputs the private key $\mathsf{SK}_{\mathsf{id}}$ for identity id. We sometimes refer to the master key as the private key at depth 0, given which the algorithm Derive performs the same function as Extract. The Setup algorithm in an HIBE scheme takes the maximum depth of the hierarchy as input.

Selective and Adaptive ID Security. The standard IBE security model of [8] defines the indistinguishability of ciphertexts under an adaptive chosen-ciphertext and chosen-identity attack (IND-ID-CCA2). A weaker notion of IBE security given by Canetti, Halevi, and Katz [11] forces the adversary to announce ahead of time the public key it will target, which is known as a selective-identity attack (IND-sID-CCA2).

As with regular public-key encryption, we can deny the adversary the ability to ask decryption queries (for the target identity), which leads to the weaker notions of indistinguishability of ciphertexts under an adaptive chosen-identity chosen-plaintext attack (IND-ID-CPA) and under a selective-identity chosen-plaintext attack (IND-sID-CPA) respectively.

Security Game. We define IBE and HIBE selective security using a game that captures a strong privacy property called *indistinguishable from random* which means that the challenge ciphertext is indistinguishable from a random element in the ciphertext space. This property implies both semantic security and recipient anonymity, and also implies that the ciphertext hides the public parameters (PP) used to create it. This can make the IBE more resistant to subpoenas since an observer cannot tell from the ciphertext which authority holds the corresponding master secret. For a security parameter λ, we let $\mathcal{M}_\lambda$ denote the message space and let $\mathcal{C}_\lambda$ denote the ciphertext space. The game, for a hierarchy of maximum depth d, proceeds as follows.

Init: The adversary is given the maximum depth of the hierarchy d and outputs a target identity $\mathsf{id}^* = (\mathsf{I}_1^*, \ldots, \mathsf{I}_k^*), k \leq d$.

Setup: The challenger runs $\mathsf{Setup}(1^\lambda, 1^d)$ (where $d = 1$ for IBE) and gives the adversary the resulting system parameters PP. It keeps the master key MK to itself.

Phase 1: The adversary issues queries $q_1, \ldots, q_m$ where the i-th query q_i is a query on id_i, where $\mathsf{id}_i = (\mathsf{I}_1, \ldots, \mathsf{I}_u)$ for some $u \leq d$. We require that id_i is not a prefix of id^*, (i.e., it is not the case that $u \leq k$ and $\mathsf{I}_i = \mathsf{I}_i^*$ for all $i = 1, \ldots, u$). The challenger responds by running algorithm $\mathsf{Extract}$ to obtain a private key d_i for the public key id_i. It sends d_i to the adversary. All queries may be made adaptively, that is, the adversary may ask q_i with knowledge of the challenger's responses to $q_1, \ldots, q_{i-1}$.

Challenge: Once the adversary decides that Phase 1 is over it outputs a plaintext $M \in \mathcal{M}_\lambda$ on which it wishes to be challenged. The challenger picks a random bit $r \in \{0, 1\}$ and a random ciphertext $C \in \mathcal{C}_\lambda$. If $r = 0$ it sets the challenge ciphertext to $C^* := \mathsf{Encrypt}(\mathsf{PP}, \mathsf{id}^*, M)$. If $r = 1$ it sets the challenge ciphertext to $C^* := C$. It sends C^* as the challenge to the adversary.

Phase 2: The adversary issues additional adaptive queries $q_{m+1}, \ldots, q_n$ where q_i is a private-key extraction query on id_i, where id_i is not a prefix of id^*. The challenger responds as in Phase 1.

Guess: Finally, the adversary outputs a guess $r' \in \{0, 1\}$ and wins if $r = r'$.

We refer to such an adversary $\mathcal{A}$ as an $\mathsf{IND\text{–}sID\text{-}CPA}$ adversary. We define the advantage of the adversary $\mathcal{A}$ in attacking an IBE or HIBE scheme $\mathcal{E}$ as

$$\mathrm{Adv}_{d,\mathcal{E},\mathcal{A}}(\lambda) = \left| \Pr[r = r'] - 1/2 \right|$$

The probability is over the random bits used by the challenger and the adversary.

Definition 1. *We say that an IBE or a depth d HIBE system $\mathcal{E}$ is selective-identity, indistinguishable from random if for all $\mathsf{IND\text{–}sID\text{-}CPA}$ PPT adversaries $\mathcal{A}$ we have that $Adv_{d,\mathcal{E},\mathcal{A}}(\lambda)$ is a negligible function. We abbreviate this by saying that $\mathcal{E}$ is $\mathsf{IND\text{–}sID\text{-}CPA}$ secure for depth d.*

2.2 Statistical Distance

Let X and Y be two random variables taking values in some finite set Ω. Define the *statistical distance*, denoted $\Delta(X; Y)$, as

$$\Delta(X; Y) := \frac{1}{2} \sum_{s \in \Omega} \left| \Pr[X = s] - \Pr[Y = s] \right|$$

We say that X is δ-uniform over Ω if $\Delta(X; U_\Omega) \leq \delta$ where U_Ω is a uniform random variable over Ω.

Let $X(\lambda)$ and $Y(\lambda)$ be ensembles of random variables. We say that X and Y are statistically close if $d(\lambda) := \Delta(X(\lambda); Y(\lambda))$ is a negligible function of λ.

2.3 Integer Lattices

Let $B = \left[b_1 \middle| \ldots \middle| b_m \right] \in \mathbb{R}^{m \times m}$ be an $m \times m$ matrix whose columns are linearly independent vectors $b_1, \ldots, b_m \in \mathbb{R}^m$. The m-dimensional full-rank lattice Λ generated by B is the set,

$$\Lambda = \mathcal{L}(B) = \left\{ y \in \mathbb{R}^m \quad \text{s.t.} \quad \exists s \in \mathbb{Z}^m \ , \ \ y = B\,s = \sum_{i=1}^{m} s_i\, b_i \right\}$$

Here, we are interested in integer lattices, i.e, when L is contained in $\mathbb{Z}^m$.

Definition 2. *For q prime, $A \in \mathbb{Z}_q^{n \times m}$ and $u \in \mathbb{Z}_q^n$, define:*

$$\begin{aligned}
\Lambda_q(A) &:= \left\{ e \in \mathbb{Z}^m \quad s.t. \quad \exists s \in \mathbb{Z}_q^n \text{ where } A^\top s = e \pmod q \right\} \\
\Lambda_q^\perp(A) &:= \left\{ e \in \mathbb{Z}^m \quad s.t. \quad A\,e = 0 \pmod q \right\} \\
\Lambda_q^u(A) &:= \left\{ e \in \mathbb{Z}^m \quad s.t. \quad A\,e = u \pmod q \right\}
\end{aligned}$$

Observe that if $t \in \Lambda_q^u(A)$ then $\Lambda_q^u(A) = \Lambda_q^\perp(A) + t$ and hence $\Lambda_q^u(A)$ is a shift of $\Lambda_q^\perp(A)$.

2.4 The Gram-Schmidt Norm of a Basis

Let S be a set of vectors $S = \{s_1, \ldots, s_k\}$ in $\mathbb{R}^m$. We use the following notation:

- $\|S\|$ denotes the L_2 length of the longest vector in S, i.e. $\|S\| := \max_i \|s_i\|$ for $1 \le i \le k$.
- $\widetilde{S} := \{\tilde{s}_1, \ldots, \tilde{s}_k\} \subset \mathbb{R}^m$ denotes the Gram-Schmidt orthogonalization of the vectors $s_1, \ldots, s_k$ taken in that order.

We refer to $\|\widetilde{S}\|$ as the Gram-Schmidt norm of S.

Micciancio and Goldwassser [22] showed that a full-rank set S in a lattice Λ can be converted into a basis T for Λ with an equally low Gram-Schmidt norm.

Lemma 1 ([22, Lemma 7.1]). *Let Λ be an m-dimensional lattice. There is a deterministic polynomial-time algorithm that, given an arbitrary basis of Λ and a full-rank set $S = \{s_1, \ldots, s_m\}$ in Λ, returns a basis T of Λ satisfying*

$$\|\widetilde{T}\| \le \|\widetilde{S}\| \quad \textit{and} \quad \|T\| \le \|S\|\sqrt{m}/2$$

Ajtai [4] showed how to sample an essentially uniform matrix $A \in \mathbb{Z}_q^{n \times m}$ with an associated basis S_A of $\Lambda_q^\perp(A)$ with low Gram-Schmidt norm. We use an improved sampling algorithm from Alwen and Peikert [5]. The following follows from Theorem 3.2 of [5] taking $\delta := 1/3$.

Theorem 1. *Let $q \geq 3$ be odd and $m := \lceil 6n \log q \rceil$.*
There is a probabilistic polynomial-time algorithm $\mathsf{TrapGen}(q, n)$ that outputs a pair $(A \in \mathbb{Z}_q^{n \times m},\ S \in \mathbb{Z}^{m \times m})$ such that A is statistically close to a uniform matrix in $\mathbb{Z}_q^{n \times m}$ and S is a basis for $\Lambda_q^\perp(A)$ satisfying

$$\|\widetilde{S}\| \leq O(\sqrt{n \log q}) \quad \textit{and} \quad \|S\| \leq O(n \log q)$$

with all but negligible probability in n.

We will also need the following simple lemma about the effect of matrix multiplication on the Gram-Schmidt norm.

Lemma 2. *Let R be a matrix in $\mathbb{R}^{\ell \times m}$ and $S = \{s_1, \ldots, s_k\} \subset \mathbb{R}^m$ a linearly independent set. Let $S_R := \{Rs_1, \ldots, Rs_k\}$. Then*

$$\|\widetilde{S_R}\| \leq \max_{1 \leq i \leq k} \|R\tilde{s}_i\|$$

Proof. We show that for all $i = 1, \ldots, k$ the i-th Gram-Schmidt vector of S_R has L_2 norm less than $\|R\tilde{s}_i\|$. This will prove the lemma.

For $i \in \{1, \ldots, k\}$ let $V := \mathrm{span}_{\mathbb{R}}(Rs_1, \ldots, Rs_{i-1})$. Set $v := s_i - \tilde{s}_i$. Then $v \in \mathrm{span}_{\mathbb{R}}(s_1, \ldots, s_{i-1})$ and therefore $Rv \in V$. Let u be the projection of $R\tilde{s}_i$ on V and let $z := R\tilde{s}_i - u$. Then z is orthogonal to V and

$$Rs_i = Rv + R\tilde{s}_i = Rv + u + z = (Rv + u) + z\ .$$

By construction, $Rv + u \in V$ and hence, since z is orthogonal to V, this z must be the i-th Gram-Schmidt vector of S_R. Since z is the projection of $R\tilde{s}_i$ on $V^\perp$ we obtain that $\|z\| \leq \|R\tilde{s}_i\|$. Hence, for all $i = 1, \ldots, k$ the i-th Gram-Schmidt vector of S_R has L_2 norm less than $\|R\tilde{s}_i\|$ which proves the lemma. □

2.5 Discrete Gaussians

Let L be a subset of $\mathbb{Z}^m$. For any vector $c \in \mathbb{R}^m$ and any positive parameter $\sigma \in \mathbb{R}_{>0}$, define:

$\rho_{\sigma,c}(x) = \exp\left(-\pi \frac{\|x-c\|^2}{\sigma^2}\right)$: a Gaussian-shaped function on $\mathbb{R}^m$ with center c and parameter σ,
$\rho_{\sigma,c}(L) = \sum_{x \in L} \rho_{\sigma,c}(x)$: the (always converging) sum of $\rho_{\sigma,c}$ over L,
$\mathcal{D}_{L,\sigma,c}$: the discrete Gaussian distribution over L with parameters σ and c,

$$\forall y \in L \quad , \quad \mathcal{D}_{L,\sigma,c}(y) = \frac{\rho_{\sigma,c}(y)}{\rho_{\sigma,c}(L)}$$

We abbreviate $\rho_{\sigma,0}$ and $\mathcal{D}_{L,\sigma,0}$ as ρ_σ and $\mathcal{D}_{L,\sigma}$. We write ρ to denote ρ_1. The distribution $\mathcal{D}_{L,\sigma,c}$ will most often be defined over the lattice $L = \Lambda_q^\perp(A)$ for a matrix $A \in \mathbb{Z}_q^{n \times m}$ or over a coset $L = t + \Lambda_q^\perp(A)$ where $t \in \mathbb{Z}^m$.

Properties. The following lemma from [24] captures standard properties of these distributions. The first two properties follow from Lemma 4.4 of [23] and Corollary 3.16 of [27] respectively (using Lemma 3.1 from [16] to bound the smoothing parameter). We state in property (2) a stronger version of Regev's Corollary 3.16 found in [2]. The last two properties are algorithms from [16].

Lemma 3. *Let $q \geq 2$ and let A be a matrix in $\mathbb{Z}_q^{n \times m}$ with $m > n$. Let T_A be a basis for $\Lambda_q^\perp(A)$ and $\sigma \geq \|\widetilde{T_A}\| \, \omega(\sqrt{\log m})$. Then for $c \in \mathbb{R}^m$ and $u \in \mathbb{Z}_q^n$:*

1. $\Pr\left[x \sim \mathcal{D}_{\Lambda_q^\perp(A),\sigma} \ : \ \|x\| > \sqrt{m}\,\sigma \right] \leq \mathrm{negl}(n)$.
2. *A set of $O(m \log m)$ samples from $\mathcal{D}_{\Lambda_q^\perp(A),\sigma}$ contains a full rank set in $\mathbb{Z}^m$, except with negligible probability.*
3. *There is a PPT algorithm $\mathsf{SampleGaussian}(A, T_A, \sigma, c)$ that returns $x \in \Lambda_q^\perp(A)$ drawn from a distribution statistically close to $\mathcal{D}_{\Lambda,\sigma,c}$.*
4. *There is a PPT algorithm $\mathsf{SamplePre}(A, T_A, u, \sigma)$ that returns $x \in \Lambda_q^u(A)$ sampled from a distribution statistically close to $\mathcal{D}_{\Lambda_q^u(A),\sigma}$.*

Recall that if $\Lambda_q^u(A)$ is not empty then $\Lambda_q^u(A) = t + \Lambda_q^\perp(A)$ for some $t \in \Lambda_q^u(A)$. Algorithm $\mathsf{SamplePre}(A, T_A, u, \sigma)$ works by calling $\mathsf{SampleGaussian}(A, T_A, \sigma, t)$ and subtracts t from the result.

2.6 The LWE Hardness Assumption

Security of all our constructions reduces to the LWE (learning with errors) problem, a classic hard problem on lattices defined by Regev [27].

Definition 3. *Consider a prime q, a positive integer n, and a distribution χ over $\mathbb{Z}_q$, all public. An $(\mathbb{Z}_q, n, \chi)$-LWE problem instance consists of access to an unspecified challenge oracle $\mathcal{O}$, being, either, a noisy pseudo-random sampler $\mathcal{O}_s$ carrying some constant random secret key $s \in \mathbb{Z}_q^n$, or, a truly random sampler $\mathcal{O}_\$$, whose behaviors are respectively as follows:*

$\mathcal{O}_s$: *outputs samples of the form $(u_i, v_i) = \left(u_i, \ u_i^T\, s + x_i\right) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$, where, $s \in \mathbb{Z}_q^n$ is a uniformly distributed persistent value invariant across invocations, $x_i \in \mathbb{Z}_q$ is a fresh sample from χ, and u_i is uniform in $\mathbb{Z}_q^n$.*

$\mathcal{O}_\$$: *outputs truly uniform random samples from $\mathbb{Z}_q^n \times \mathbb{Z}_q$.*

The $(\mathbb{Z}_q, n, \chi)$-LWE problem allows repeated queries to the challenge oracle $\mathcal{O}$. We say that an algorithm $\mathcal{A}$ decides the $(\mathbb{Z}_q, n, \chi)$-LWE problem if $\left| \Pr[\mathcal{A}^{\mathcal{O}_s} = 1] - \Pr[\mathcal{A}^{\mathcal{O}_\$} = 1] \right|$ is non-negligible for a random $s \in \mathbb{Z}_q^n$.

Regev [27] shows that for certain noise distributions χ, denoted $\overline{\Psi}_\alpha$, the LWE problem is as hard as the worst-case SIVP and GapSVP under a quantum reduction (see also [25]).

Definition 4. *Consider a real parameter $\alpha = \alpha(n) \in (0, 1)$ and a prime q. Denote by $\mathbb{T} = \mathbb{R}/\mathbb{Z}$ the group of reals $[0, 1)$ with addition modulo 1. Denote*

by Ψ_α the distribution over $\mathbb{T}$ of a normal variable with mean 0 and standard deviation $\alpha/\sqrt{2\,\pi}$ then reduced modulo 1. Denote by $\lfloor x \rceil = \lfloor x + \frac{1}{2} \rfloor$ the nearest integer to the real $x \in \mathbb{R}$. We denote by $\overline{\Psi}_\alpha$ the discrete distribution over $\mathbb{Z}_q$ of the random variable $\lfloor q\,X \rceil \mod q$ where the random variable $X \in \mathbb{T}$ has distribution Ψ_α.

Theorem 2 ([27]). *If there exists an efficient, possibly quantum, algorithm for deciding the $(\mathbb{Z}_q, n, \overline{\Psi}_\alpha)$-LWE problem for $q > 2\sqrt{n}/\alpha$ then there exists an efficient quantum algorithm for approximating the SIVP and GapSVP problems, to within $\tilde{O}(n/\alpha)$ factors in the ℓ_2 norm, in the worst case.*

If we assume the hardness of approximating the SIVP or GapSVP problems in lattices of dimension n to within approximation factors that are polynomial in n, then it follows from Lemma 2 that deciding the LWE problem is hard when n/α is polynomial in n.

3 Randomness Extraction

We will need the following lemma which follows directly from a generalization of the left over hash lemma due to Dodis et al. [15].

Lemma 4. *Suppose that $m > (n+1)\log_2 q + \omega(\log n)$ and that q is prime. Let A, B be matrices chosen uniformly in $\mathbb{Z}_q^{n\times m}$ and let R be an $m \times m$ matrix chosen uniformly in $\{1, -1\}^{m\times m} \bmod q$. Then, for all vectors w in $\mathbb{Z}_q^m$, the distribution $(A,\ AR,\ R^\top w)$ is statistically close to the distribution $(A,\ B,\ R^\top w)$.*

To prove the lemma recall that for a prime q the family of hash functions $h_A : \mathbb{Z}_q^m \to \mathbb{Z}_q^n$ for $A \in \mathbb{Z}_q^{n\times m}$ defined by $h_A(x) = A\,x$ is universal. Therefore, when the columns of R are sampled independently and have sufficient entropy, the left over hash lemma (e.g. as stated in [29, Theorem 8.38]) shows that the distributions $(A,\ AR)$ and $(A,\ B)$ are statistically close. A generalization by Dodis et al. [15] (Lemma 2.2b and 2.4) shows that the same holds even if some small amount of information about R is leaked. In our case $R^\top w$ is leaked which is precisely the settings of Dodis et al. We provide the complete proof of Lemma 4 in the full version of the paper [1].

3.1 Random Subset Sums

We will also need the following simple lemma.

Lemma 5. *Let R be an $m \times m$ matrix chosen at random from $\{-1, 1\}^{m\times m}$. Then for all vectors $u \in \mathbb{R}^m$ we have*

$$\Pr\left[\ \|Ru\| > \|u\|\ \sqrt{m}\cdot\omega(\sqrt{\log m})\ \right] < \mathrm{negl}(m)\ .$$

Proof. Let $r \in \{-1,1\}^m$ be a row vector of the matrix R. Then $r \cdot u$ can be written as $r^\top u = \sum_{i=1}^m x_i$ where $x_i = r_i u_i$. We know that $E[x_i] = 0$ and that $x_i \in [-u_i, u_i]$ for all $i = 1, \ldots, m$. Then, by the Hoeffding bound [19, Theorem 2] we obtain that

$$\Pr\left[\ |r \cdot u| > \|u\|\ \omega(\sqrt{\log m})\ \right] < \mathrm{negl}(m)$$

The lemma now follows since an m-vector whose entries are less than some bound B has L_2 norm less than $\sqrt{m}B$. □

4 Sampling Algorithms

Let A and B be matrices in $\mathbb{Z}_q^{n\times m}$ and let R be a matrix in $\{-1,1\}^{m\times m}$. Our construction makes use of matrices of the form $F = (A \mid AR + B) \in \mathbb{Z}_q^{n\times 2m}$ and we will need to sample short vectors in $\Lambda_q^u(F)$ for some u in $\mathbb{Z}_q^n$. We show that this can be done using either a trapdoor for $\Lambda_q^\perp(A)$ or a trapdoor $\Lambda_q^\perp(B)$. More precisely, we define two algorithms:

1. SampleLeft takes a basis for $\Lambda_q^\perp(A)$ (the left side of F) and outputs a short vector $e \in \Lambda_q^u(F)$.
2. SampleRight takes a basis for $\Lambda_q^\perp(B)$ (the right side of F) and outputs a short vector $e \in \Lambda_q^u(F)$.

We will show that, with appropriate parameters, the distributions on e produced by these two algorithms are statistically indistinguishable.

4.1 Algorithm SampleLeft

Algorithm $\mathsf{SampleLeft}(A, M_1, T_A, u, \sigma)$:
Inputs:

$$\begin{array}{l}\text{a rank } n \text{ matrix } A \text{ in } \mathbb{Z}_q^{n\times m} \text{ and a matrix } M_1 \text{ in } \mathbb{Z}_q^{n\times m_1},\\ \text{a "short" basis } T_A \text{ of } \Lambda_q^\perp(A) \text{ and a vector } u \in \mathbb{Z}_q^n,\\ \text{a gaussian parameter } \sigma > \|\widetilde{T_A}\| \cdot \omega(\sqrt{\log(m+m_1)}).\end{array} \quad (1)$$

Output: Let $F_1 := (A \mid M_1)$. The algorithm outputs a vector $e \in \mathbb{Z}^{m+m_1}$ sampled from a distribution statistically close to $\mathcal{D}_{\Lambda_q^u(F_1),\sigma}$. In particular, $e \in \Lambda_q^u(F_1)$.

The algorithm appears in Theorem 3.4 in [12] and also in the signing algorithm in [24]. For completeness, we briefly review the algorithm.

1. sample a random vector $e_2 \in \mathbb{Z}^{m_1}$ distributed statistically close to $\mathcal{D}_{\mathbb{Z}^{m_1},\sigma}$,
2. run $e_1 \xleftarrow{R} \mathsf{SamplePre}(A, T_A, y, \sigma)$ where $y = u - (M_1 \cdot e_2) \in \mathbb{Z}_q^n$,
   note that $\Lambda_q^y(A)$ is not empty since A is rank n,
3. output $e \leftarrow (e_1, e_2) \in \mathbb{Z}^{m+m_1}$

Clearly $(A \,|\, M_1) \cdot e = u \bmod q$ and hence $e \in \Lambda_q^u(F_1)$. Theorem 3.4 in [12] shows that the vector e is sampled from a distribution statistically close to $\mathcal{D}_{\Lambda_q^u(F_1),\sigma}$.

Peikert's basis extension method [24] gives an alternate way to view this. Given the basis T_A of $\Lambda_q^\perp(A)$ Peikert shows how to build a basis T_{F_1} of $\Lambda_q^\perp(F_1)$ with the same Gram-Schmidt norm as T_A. Then calling $\mathsf{SamplePre}(F_1, T_{F_1}, u, \sigma)$ generates a vector e sampled from a distribution close to $\mathcal{D}_{\Lambda_q^u(F_1),\sigma}$. We summarize this in the following theorem.

Theorem 3. *Let $q > 2$, $m > 2n \log q$ and $\sigma > \|\widetilde{T_A}\| \cdot \omega(\sqrt{\log(m + m_1)})$. Then Algorithm $\mathsf{SampleLeft}(A, M_1, T_A, u, \sigma)$ taking inputs as in (1), outputs a vector $e \in \mathbb{Z}^{m+m_1}$ distributed statistically close to $\mathcal{D}_{\Lambda_q^u(F_1),\sigma}$ where $F_1 := (A \mid M_1)$.*

4.2 Algorithm SampleRight

Algorithm $\mathsf{SampleRight}(A, B, R, T_B, u, \sigma)$.

Inputs:

$$\begin{array}{l} \text{matrices } A, B \text{ in } \mathbb{Z}_q^{n\times m} \text{ where } B \text{ is rank } n, \\ \text{a uniform random matrix } R \in \{-1, 1\}^{m\times m}, \\ \text{a basis } T_B \text{ of } \Lambda_q^\perp(B) \text{ and a vector } u \in \mathbb{Z}_q^n, \\ \text{a parameter } \sigma > \|\widetilde{T_B}\| \cdot \sqrt{m} \cdot \omega(\log m). \end{array} \tag{2}$$

Output: Let $F_2 := (A \mid AR+B)$. The algorithm outputs a vector $e \in \mathbb{Z}^{2m}$ sampled from a distribution statistically close to $\mathcal{D}_{\Lambda_q^u(F_2),\sigma}$. In particular, $e \in \Lambda_q^u(F_2)$.

The algorithm uses the basis growth method of Peikert [24, Sec. 3.3] and works in three steps:

1. First, it constructs a set T_{F_2} of $2m$ linearly independent vectors in $\Lambda_q^\perp(F_2)$ such that
$$\|\widetilde{T_{F_2}}\| < \|\widetilde{T_B}\| \cdot \sqrt{m} \cdot \omega(\sqrt{\log m}) < \sigma/\omega(\sqrt{\log m})$$
with overwhelming probability over the choice of R.
2. Next, if needed it uses Lemma 1 to convert T_{F_2} into a basis T'_{F_2} of $\Lambda_q^\perp(F_2)$ with the same Gram-Schmidt norm as T_{F_2}.
3. Finally, it invokes $\mathsf{SamplePre}(F_2, T'_{F_2}, u, \sigma)$ to generate a vector $e \in \Lambda_q^u(F_2)$. Since $\sigma > \|\widetilde{T_{F_2}}\| \, \omega(\sqrt{\log m})$ w.h.p, this e is distributed as $\mathcal{D}_{\Lambda_q^u(F_2),\sigma}$, as required.

Step 1 is the only step that needs explaining. Let $T_B = \{b_1, \ldots, b_m\} \in \mathbb{Z}^{m\times m}$ be the given basis of $\Lambda_q^\perp(B)$. We construct the $2m$ vectors in $\Lambda_q^\perp(F_2)$ as follows:

1. for $i = 1, \ldots, m$ set $t_i := (-Rb_i \mid b_i) \in \mathbb{Z}^{2m}$ and view it as a column vector; then clearly $F_2 \cdot t_i = B\, b_i = 0 \bmod q$ and therefore t_i is in $\Lambda_q^\perp(F_2)$.
2. for $i = 1, \ldots, m$ let w_i be the i-th column of the identity matrix I_m. Let u_i be an arbitrary vector in $\mathbb{Z}^m$ satisfying $Aw_i + Bu_i = 0 \bmod q$. This u_i exists since B is rank n. Set t_{i+m} to be
$$t_{i+m} := \begin{bmatrix} w_i - Ru_i \\ u_i \end{bmatrix} \in \mathbb{Z}^{2m}$$
Then $F_2 \cdot t_{i+m} = Aw_i + Bu_i = 0 \bmod q$ and hence, $t_{i+m} \in \Lambda_q^\perp(F_2)$.

We show that $T_{F_2} := \{t_1, \ldots, t_{2m}\}$ are linearly independent in $\mathbb{Z}^{2m}$. First, observe that the first m vectors are linearly independent and span the linear space V of vectors of the form $(-Rx \mid x)$ where $x \in \mathbb{Z}_q^m$. For all $i > m$, the vector t_i is the sum of the unit vector $(w_i \mid 0^m)$ plus a vector in V. It follows that T_{F_2} is a linearly independent set. This also means that for $i > m$ the i-th Gram-Schmidt vector of T_{F_2} cannot be longer than $(w_i \mid 0^m)$ and therefore has norm at most 1. Hence, to bound $\|\widetilde{T_{F_2}}\|$ it suffices to bound the Gram-Schmidt norm of the first m vectors $\{t_1, \ldots, t_m\}$.

Let $W \in \mathbb{Z}^{2m \times m}$ be the matrix $(-R^\top \mid I_m)^\top$. Then $t_i = Wb_i$ for $i = 1, \ldots, m$. Since R is uniform in $\{-1, 1\}^{m \times m}$ we know by Lemma 5 that for all vectors $x \in \mathbb{R}^m$ we have w.h.p

$$\|W\,x\| \leq \|R\,x\| + \|x\| \leq \|x\|\sqrt{m} \cdot \omega(\sqrt{\log m}) + \|x\| \leq \|x\|\sqrt{m} \cdot \omega(\sqrt{\log m})$$

Now, since $t_i = Wb_i$ for $i = 1, \ldots, m$, applying Lemma 2 to the matrix W gives a bound on the Gram-Schmidt norm of $\{t_1, \ldots, t_m\}$ (and hence also on $\|\widetilde{T_{F_2}}\|$):

$$\begin{aligned}\|\widetilde{T_{F_2}}\| &\leq \max_{1 \leq i \leq m} \|W\,\tilde{b}_i\| \leq \max_{1 \leq i \leq m} \|\tilde{b}_i\| \cdot \sqrt{m} \cdot \omega(\sqrt{\log m}) \\ &\leq \|\widetilde{T_B}\| \cdot \sqrt{m} \cdot \omega(\sqrt{\log m})\end{aligned}$$

Thus, we built $2m$ linearly independent vectors in $\Lambda_q^\perp(F_2)$ that w.h.p. have a short Gram-Schmidt norm as required for Step 1. This completes the description of algorithm SampleRight. We summarize this in the following theorem.

Theorem 4. *Let $q > 2, m > n$ and $\sigma > \|\widetilde{T_B}\| \cdot \sqrt{m} \cdot \omega(\log m)$. Then Algorithm, SampleRight$(A, B, R, T_B, u, \sigma)$ taking inputs as in (2), with R uniform in $\{1, -1\}^{m \times m}$, outputs a vector $e \in \mathbb{Z}^{2m}$ distributed statistically close to $\mathcal{D}_{\Lambda_q^u(F_2),\sigma}$ where $F_2 := (A \mid AR + B)$.*

5 Encoding Identities as Matrices

Our construction uses an encoding function $H : \mathbb{Z}_q^n \to \mathbb{Z}_q^{n \times n}$ to map identities in $\mathbb{Z}_q^n$ to matrices in $\mathbb{Z}_q^{n \times n}$. Our proof of security requires that the map H satisfy a strong notion of injectivity, namely that, for any two distinct inputs id_1 and id_2, the difference between the outputs $H(\mathsf{id}_1)$ and $H(\mathsf{id}_2)$ is never singular, i.e., $\det(H(\mathsf{id}_1) - H(\mathsf{id}_2)) \neq 0$.

Definition 5. *Let q be a prime and n a positive integer. We say that a function $H : \mathbb{Z}_q^n \to \mathbb{Z}_q^{n \times n}$ is an* **encoding with full-rank differences** *(FRD) if:*

1. *for all distinct $u, v \in \mathbb{Z}_q^n$, the matrix $H(u) - H(v) \in \mathbb{Z}_q^{n \times n}$ is full rank; and*
2. *H is computable in polynomial time (in $n \log q$).*

Clearly the function H must be injective since otherwise, if $u \neq v$ satisfies $H(u) = H(v)$, then $H(u) - H(v)$ is not full-rank and hence H cannot be FRD.

The function H in Definition 5 has domain of size q^n which is the largest possible for a function satisfying condition 1 of Definition 5. Indeed, if H had domain larger than q^n then its image is also larger than q^n. But then, by pigeonhole, there are two distinct inputs u, v such that the matrices $H(u)$ and $H(v)$ have the same first row and therefore $H(u) - H(v)$ is not full rank. It follows that our definition of FRD, which has domain of size of q^n, is the largest possible.

An Explicit FRD Construction. We construct an injective FRD encoding for the exponential-size domain $\mathsf{id} \in \mathbb{Z}_q^n$. A similar construction is described in [14]. Our strategy is to construct an additive subgroup $\mathbb{G}$ of $\mathbb{Z}_q^{n\times n}$ of size q^n such that all non-zero matrices in $\mathbb{G}$ are full-rank. Since for all distinct $A, B \in \mathbb{G}$ the difference $A - B$ is also in $\mathbb{G}$, it follows that $A - B$ is full-rank.

While our primary interest is the finite field $\mathbb{Z}_q$ we describe the construction for an arbitrary field $\mathbb{F}$. For a polynomial $g \in \mathbb{F}[X]$ of degree less than n define $\text{coeffs}(g) \in \mathbb{F}^n$ to be the n-vector of coefficients of g (written as a row-vector). If g is of degree less than $n-1$ we pad the coefficients vector with zeroes on the right to make it an n-vector. For example, for $n = 6$ we have $\text{coeffs}(x^3 + 2x + 3) = (3, 2, 0, 1, 0, 0) \in \mathbb{F}^6$. Let f be some polynomial of degree n in $\mathbb{F}[X]$ that is irreducible. Recall that for a polynomial $g \in \mathbb{F}[X]$ the polynomial $g \bmod f$ has degree less than n and therefore $\text{coeffs}(g \bmod f)$ is a vector in $\mathbb{F}^n$.

Now, for an input $u = (u_0, u_1, \ldots, u_{n-1}) \in \mathbb{F}^n$ define the polynomial $g_u(X) = \sum_{i=0}^{n-1} u_i x^i \in \mathbb{F}[X]$. Define $H(u)$ as

$$H(u) := \begin{pmatrix} \text{coeffs}(\ g_u\ ) \\ \text{coeffs}(\ X \cdot g_u \bmod f\ ) \\ \text{coeffs}(\ X^2 \cdot g_u \bmod f\ ) \\ \vdots \\ \text{coeffs}(\ X^{n-1} \cdot g_u \bmod f\ ) \end{pmatrix} \in \mathbb{F}^{n\times n} \tag{3}$$

This completes the construction. Since for all primes q and integers $n > 1$ there are (many) irreducible polynomials in $\mathbb{Z}_q[X]$ of degree n, the construction can accommodate any pair of q and n.

The following theorem proves that the function H in (3) is an FRD. The proof, given in [14], is based on the observation that the matrix $H(u)^\top$ corresponds to multiplication by a constant in the number field $K = \mathbb{F}[X]/(f)$ and is therefore invertible when the matrix is non-zero. We note that similar matrix encodings of ring multiplication were previously used in [26,21].

Theorem 5. *Let $\mathbb{F}$ be a field and f a polynomial in $\mathbb{F}[X]$. If f is irreducible in $\mathbb{F}[X]$ then the function H defined in (3) is an encoding with full-rank differences (or FRD encoding).*

An example. Let $n = 4$ and $f(X) = x^4 + x - 1$. The function H works as follows:

$$H\big(u = (u_0, u_1, u_2, u_3)\big) \;:=\; \begin{pmatrix} u_0 & u_1 & u_2 & u_3 \\ u_3 & u_0 - u_3 & u_1 & u_2 \\ u_2 & u_3 - u_2 & u_0 - u_3 & u_1 \\ u_1 & u_2 - u_1 & u_3 - u_2 & u_0 - u_3 \end{pmatrix}$$

Theorem 5 shows that the map H is FRD for all primes q where $x^4 + x - 1$ is irreducible in $\mathbb{Z}_q[X]$ (e.g. $q = 19, 31, 43, 47$).

6 The Main Construction: An Efficient IBE

The system uses parameters q, n, m, σ, α specified in Section 6.3. Throughout the section, the function H refers to the FRD map $H : \mathbb{Z}_q^n \to \mathbb{Z}_q^{n \times n}$ defined in Section 5. We assume identities are elements in $\mathbb{Z}_q^n$. The set of identities can be expanded to $\{0,1\}^*$ by hashing identities into $\mathbb{Z}_q^n$ using a collision resistant hash.

6.1 Intuition

The public parameters in our system consist of three random $n \times m$ matrices over $\mathbb{Z}_q$ denoted by A_0, A_1 and B as well as a vector $u \in \mathbb{Z}_q^n$. The master secret is a trapdoor T_{A_0} (i.e. a basis with a low Gram-Schmidt norm) for the lattice $\Lambda_q^\perp(A_0)$.

The secret key for an identity id is a short vector $e \in \mathbb{Z}^{2m}$ satisfying $F_{\mathsf{id}} \cdot e = u$ in $\mathbb{Z}_q$ where

$$F_{\mathsf{id}} := (A_0 \mid A_1 + H(\mathsf{id})\, B) \quad \in \mathbb{Z}_q^{n \times 2m}$$

The vector e is generated using algorithm $\mathsf{SampleLeft}$ (Theorem 3) and the trapdoor T_{A_0}.

In a selective IBE security game the attacker announces an identity id^* that it plans to attack. We need a simulator that can respond to private key queries for $\mathsf{id} \neq \mathsf{id}^*$, but knows nothing about the private key for id^*. We do so by choosing the public parameters A_0 and B at random as before, but choosing A_1 as

$$A_1 := A_0\, R - H(\mathsf{id}^*)\, B$$

where R is a random matrix in $\{1, -1\}^{m \times m}$. We show that $A_0\, R$ is uniform and independent in $\mathbb{Z}_q^{n \times m}$ so that A_1 is distributed as required. We provide the simulator with a trapdoor T_B for $\Lambda_q^\perp(B)$, but no trapdoor for $\Lambda_q^\perp(A_0)$.

Now, to respond to a private key query for an identity id, the simulator must produce a short vector e satisfying $F_{\mathsf{id}} \cdot e = u$ in $\mathbb{Z}_q$ where

$$F_{\mathsf{id}} := \big(A_0 \mid A_0 \cdot R + B'\big) \in \mathbb{Z}_q^{n \times 2m} \quad \text{and} \quad B' := \big(H(\mathsf{id}) - H(\mathsf{id}^*)\big) \cdot B\ .$$

When $\mathsf{id} \neq \mathsf{id}^*$ we know that $H(\mathsf{id}) - H(\mathsf{id}^*)$ is full rank by construction and therefore T_B is also a trapdoor for the lattice $\Lambda_q^\perp(B')$. The simulator can now generate e using algorithm $\mathsf{SampleRight}$ and the basis T_B.

When $\mathsf{id} = \mathsf{id}^*$ the matrix F_{id} no longer depends on B and the simulator's trapdoor disappears. Consequently, the simulator can generate private keys for

all identities other than id^*. As we will see, for id^* the simulator can produce a challenge ciphertext that helps it solve the given LWE challenge.

6.2 The Basic IBE Construction

Setup(λ)**:** On input a security parameter λ, set the parameters q, n, m, σ, α as specified in Section 6.3 below. Next do:

1. Use algorithm $\mathsf{TrapGen}(q, n)$ to select a uniformly random $n \times m$-matrix $A_0 \in \mathbb{Z}_q^{n \times m}$ with a basis T_{A_0} for $\Lambda_q^{\perp}(A_0)$ such that $\|\widetilde{T_{A_0}}\| \leq O(\sqrt{n \log q})$
2. Select two uniformly random $n \times m$ matrices A_1 and B in $\mathbb{Z}_q^{n \times m}$.
3. Select a uniformly random n-vector $u \xleftarrow{R} \mathbb{Z}_q^n$.
4. Output the public parameters and master key,

$$\mathsf{PP} = \begin{pmatrix} A_0, A_1, B, u \end{pmatrix} \quad ; \quad \mathsf{MK} = \begin{pmatrix} T_{A_0} \end{pmatrix} \in \mathbb{Z}^{m \times m}$$

Extract$(\mathsf{PP}, \mathsf{MK}, \mathsf{id})$**:** On input public parameters PP, a master key MK, and an identity $\mathsf{id} \in \mathbb{Z}_q^n$, do:

1. Sample $e \in \mathbb{Z}^{2m}$ as $e \leftarrow \mathsf{SampleLeft}(A_0, \quad A_1 + H(\mathsf{id})\, B, \quad T_{A_0}, \quad u, \quad \sigma)$ where H is an FRD map as defined in Section 5.
   Note that A_0 is rank n w.h.p as explained in Section 6.3.
2. Output $\mathsf{SK}_{\mathsf{id}} := e \in \mathbb{Z}^{2m}$

Let $F_{\mathsf{id}} := \big(A_0 \mid A_1 + H(\mathsf{id})\, B\big)$, then $F_{\mathsf{id}} \cdot e = u$ in $\mathbb{Z}_q$ and e is distributed as $D_{\Lambda_q^u(F_{\mathsf{id}}),\sigma}$ by Theorem 3.

Encrypt$(\mathsf{PP}, \mathsf{id}, b)$**:** On input public parameters PP, an identity id, and a message $b \in \{0, 1\}$, do:

1. Set $F_{\mathsf{id}} \leftarrow \big(A_0 \mid A_1 + H(\mathsf{id}) \cdot B\big) \in \mathbb{Z}_q^{n \times 2m}$
2. Choose a uniformly random $s \xleftarrow{R} \mathbb{Z}_q^n$
3. Choose a uniformly random $m \times m$ matrix $R \xleftarrow{R} \{-1, 1\}^{m \times m}$
4. Choose noise vectors $x \xleftarrow{\bar{\Psi}_\alpha} \mathbb{Z}_q$ and $y \xleftarrow{\bar{\Psi}_\alpha^m} \mathbb{Z}_q^m$, and set $z \leftarrow R^\top y \in \mathbb{Z}_q^m$ (the distribution $\bar{\Psi}_\alpha$ is as in Definition 4),
5. Set $c_0 \leftarrow u^\top s + x + b \lfloor \frac{q}{2} \rfloor \in \mathbb{Z}_q$ and $c_1 \leftarrow F_{\mathsf{id}}^\top s + \begin{bmatrix} y \\ z \end{bmatrix} \in \mathbb{Z}_q^{2m}$
6. Output the ciphertext $\mathsf{CT} := (c_0, c_1) \in \mathbb{Z}_q \times \mathbb{Z}_q^{2m}$.

Decrypt$(\mathsf{PP}, \mathsf{SK}_{\mathsf{id}}, \mathsf{CT})$**:** On input public parameters PP, a private key $\mathsf{SK}_{\mathsf{id}} := e_{\mathsf{id}}$, and a ciphertext $\mathsf{CT} = (c_0, c_1)$, do:

1. Compute $w \leftarrow c_0 - e_{\mathsf{id}}^\top c_1 \in \mathbb{Z}_q$.
2. Compare w and $\lfloor \frac{q}{2} \rfloor$ treating them as integers in $\mathbb{Z}$. If they are close, i.e., if $\left| w - \lfloor \frac{q}{2} \rfloor \right| < \lfloor \frac{q}{4} \rfloor$ in $\mathbb{Z}$, output 1, otherwise output 0.

The matrix R. The matrix R used in encryption plays an important role in the security proof. Note that the matrix is only used as a tool to sample the noise vector (y, z) from a specific distribution needed in the simulation.

6.3 Parameters and Correctness

When the cryptosystem is operated as specified, we have,

$$w = c_0 - e_{\mathsf{id}}^{\top} c_1 = b \lfloor \frac{q}{2} \rfloor + \underbrace{x - e_{\mathsf{id}}^{\top} \begin{bmatrix} y \\ z \end{bmatrix}}_{\text{error term}}$$

In the full paper we show that the error term is bounded by $\tilde{O}(q\alpha\sigma m)$ w.h.p. This follows from the same analysis as in [16, Lemma 8.2] plus Lemma 5 to bound $\|z\|$.

To ensure that the error term is less than $q/5$, that σ is sufficiently large for SampleLeft and SampleRight, that TrapGen can operate (i.e. $m > 6n \log q$), and that Regev's reduction applies (i.e. $q > 2\sqrt{n}/\alpha$), we set the parameters (q, m, σ, α) as follows, taking n to be the security parameter:

$$\begin{aligned} m = 6\, n^{1+\delta} \quad &, \quad q = m^2 \sqrt{n} \cdot \omega(\log n) \\ \sigma = m \cdot \omega(\log n) \quad &, \quad \alpha = [m^2 \cdot \omega(\log n)]^{-1} \end{aligned} \tag{4}$$

and round up m to the nearest larger integer and q to the nearest larger prime. Here we assume that δ is such that $n^{\delta} > \lceil \log q \rceil = O(\log n)$.

Since the matrices A_0, B are random in $\mathbb{Z}_q^{n \times m}$ and $m > n \log q$, with overwhelming probability both matrices will have rank n. Hence, calling SampleLeft in algorithm Extract succeeds w.h.p.

6.4 Security Reduction

We show that the basic IBE construction is indistinguishable from random under a selective identity attack as in Definition 1. Recall that indistinguishable from random means that the challenge ciphertext is indistinguishable from a random element in the ciphertext space. This property implies both semantic security and recipient anonymity.

Theorem 6. *The basic IBE system with parameters $(q, n, m, \sigma, \alpha)$ as in (4) is INDr–sID-CPA secure provided that the $(\mathbb{Z}_q, n, \bar{\Psi}_\alpha)$-LWE assumption holds.*

Proof. The proof proceeds in a sequence of games where the first game is identical to the INDr–sID-CPA game from Definition 1. In the last game in the sequence the adversary has advantage zero. We show that a PPT adversary cannot distinguish between the games which will prove that the adversary has negligible advantage in winning the original INDr–sID-CPA game. The LWE problem is used in proving that Games 2 and 3 are indistinguishable.

Game 0. This is the original INDr–sID-CPA game from Definition 1 between an attacker $\mathcal{A}$ against our scheme and an INDr–sID-CPA challenger.

Game 1. Recall that in Game 0 the challenger generates the public parameters PP by choosing three random matrices A_0, A_1, B in $\mathbb{Z}_q^{n \times m}$ such that a trapdoor

T_{A_0} is known for $\Lambda_q^{\perp}(A_0)$. At the challenge phase the challenger generates a challenge ciphertext CT^*. We let $R^* \in \{-1,1\}^{m\times m}$ denote the random matrix generated for the creation of CT^* (in step 3 of Encrypt).

In Game 1 we slightly change the way that the challenger generates A_1 in the public parameters. Let id^* be the identity that $\mathcal{A}$ intends to attack. The Game 1 challenger chooses R^* at the setup phase and constructs A_1 as

$$A_1 \leftarrow A_0\, R^* - H(\mathsf{id}^*)\, B \tag{5}$$

The remainder of the game is unchanged.

We show that Game 0 is statistically indistinguishable from Game 1 by Lemma 4. Observe that in Game 1 the matrix R^* is used only in the construction of A_1 and in the construction of the challenge ciphertext where $z \leftarrow (R^*)^{\top} y$. By Lemma 4 the distribution $(A_0,\ A_0\, R^*,\ z)$ is statistically close to the distribution $(A_0,\ A_1',\ z)$ where A_1' is a uniform $\mathbb{Z}_q^{n\times m}$ matrix. It follows that in the adversary's view, the matrix $A_0\, R^*$ is statistically close to uniform and therefore A_1 as defined in (5) is close to uniform. Hence, A_1 in Games 0 and 1 are indistinguishable.

Game 2. We now change how A_0 and B in PP are chosen. In Game 2 we generate A_0 as a random matrix in $\mathbb{Z}_q^{n\times m}$, but generate B using algorithm TrapGen so that B is a random matrix in $\mathbb{Z}_q^{n\times m}$, but the challenger has a trapdoor T_B for $\Lambda_q^{\perp}(B)$. The choice of A_1 remains as in Game 1, i.e. $A_1 = A_0 \cdot R^* - H(\mathsf{id}^*) \cdot B$.

The challenger responds to private key queries using the trapdoor T_B. To respond to a private key query for $\mathsf{id} \neq \mathsf{id}^*$ the challenger needs a short $e \in \Lambda_q^u(F_{\mathsf{id}})$ where

$$F_{\mathsf{id}} := (A_0 \mid A_1 + H(\mathsf{id}) \cdot B) = \big(A_0 \mid A_0 R^* + \big(H(\mathsf{id}) - H(\mathsf{id}^*)\big)B\big)\ .$$

By construction, $[H(\mathsf{id}) - H(\mathsf{id}^*)]$ is non-singular and therefore T_B is also a trapdoor for $\Lambda_q^{\perp}(B')$ where $B' := \big(H(\mathsf{id}) - H(\mathsf{id}^*)\big)B$. Moreover, since B is rank n w.h.p, so is B'. The challenger can now respond to the private key query by running

$$e \leftarrow \mathsf{SampleRight}\big(A_0,\ \ \big(H(\mathsf{id}) - H(\mathsf{id}^*)\big)B,\ \ R^*,\ T_B,\ \ u,\ \ \sigma\big) \in \mathbb{Z}_q^{2m}$$

and sending $\mathsf{SK}_{\mathsf{id}} := e$ to $\mathcal{A}$. Since the σ used in the system is sufficiently large, this e is distributed close to $D_{\Lambda_q^u(F_{\mathsf{id}}),\sigma}$, as in Game 1 by Theorem 4.

Game 2 is otherwise the same as Game 1. Since A_0, B and responses to private key queries are statistically close to those in Game 1, the adversary's advantage in Game 2 is at most negligibly different from its advantage in Game 1.

Game 3. Game 3 is identical to Game 2 except that the challenge ciphertext (c_0^*, c_1^*) is *always* chosen as a random independent element in $\mathbb{Z}_q \times \mathbb{Z}_q^{2m}$. Since the challenge ciphertext is always a fresh random element in the ciphertext space, $\mathcal{A}$'s advantage in this game is zero.

It remains to show that Game 2 and Game 3 are computationally indistinguishable for a PPT adversary, which we do by giving a reduction from the LWE problem.

Reduction from LWE. Suppose $\mathcal{A}$ has non-negligible advantage in distinguishing Games 2 and 3. We use $\mathcal{A}$ to construct an LWE algorithm $\mathcal{B}$.

Recall from Definition 3 that an LWE problem instance is provided as a sampling oracle $\mathcal{O}$ which can be either truly random $\mathcal{O}_\$$ or a noisy pseudorandom $\mathcal{O}_s$ for some secret $s \in \mathbb{Z}_q^n$. The simulator $\mathcal{B}$ uses the adversary $\mathcal{A}$ to distinguish between the two, and proceeds as follows:

Instance. $\mathcal{B}$ requests from $\mathcal{O}$ and receives, for each $i = 0, \ldots, m$, a fresh pair $(u_i, v_i) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$.

Targeting. $\mathcal{A}$ announces to $\mathcal{B}$ the identity id^* that it intends to attack.

Setup. $\mathcal{B}$ constructs the system's public parameters PP as follows:

1. Assemble the random matrix $A_0 \in \mathbb{Z}_q^{n \times m}$ from m of the previously given LWE samples by letting the i-th column of A_0 be the n-vector u_i for all $i = 1, \ldots, m$.
2. Assign the zeroth LWE sample (so far unused) to become the public random n-vector $u_0 \in \mathbb{Z}_q^n$.
3. The remainder of the public parameters, namely A_1 and B, are constructed as in Game 2 using id^* and R^*.

Queries. $\mathcal{B}$ answers each private-key extraction query as in Game 2.

Challenge. $\mathcal{B}$ prepares, when prompted by $\mathcal{A}$ with a message bit $b^* \in \{0, 1\}$, a challenge ciphertext for the target identity id^*, as follows:

1. Let $v_0, \ldots, v_m$ be entries from the LWE instance. Set $v^* = \begin{bmatrix} v_1 \\ \vdots \\ v_m \end{bmatrix} \in \mathbb{Z}_q^m$.
2. Blind the message bit by letting $c_0^* = v_0 + b^* \lfloor \frac{q}{2} \rceil \in \mathbb{Z}_q$.
3. Set $c_1^* = \begin{bmatrix} v^* \\ (R^*)^\top v^* \end{bmatrix} \in \mathbb{Z}_q^{2m}$.
4. Choose a random bit $r \xleftarrow{R} \{0, 1\}$. If $r = 0$ send $\mathsf{CT}^* = (c_0^*, c_1^*)$ to the adversary. If $r = 1$ choose a random $(c_0, c_1) \in \mathbb{Z}_q \times \mathbb{Z}_q^{2m}$ and send (c_0, c_1) to the adversary.

We argue that when the LWE oracle is pseudorandom (i.e. $\mathcal{O} = \mathcal{O}_s$) then CT^* is distributed exactly as in Game 2. First, observe that $F_{\mathsf{id}^*} = (A_0 \mid A_0 R^*)$. Second, by definition of $\mathcal{O}_s$ we know that $v^* = A_0^\top s + y$ for some random noise vector $y \in \mathbb{Z}_q^m$ distributed as $\bar{\Psi}_\alpha^m$. Therefore, c_1^* defined in step (3) above satisfies

$$c_1^* = \begin{bmatrix} A_0^\top s + y \\ (R^*)^\top A_0^\top s + (R^*)^\top y \end{bmatrix} = \begin{bmatrix} A_0^\top s + y \\ (A_0 R^*)^\top s + (R^*)^\top y \end{bmatrix} = (F_{\mathsf{id}^*})^\top s + \begin{bmatrix} y \\ (R^*)^\top y \end{bmatrix}$$

and the quantity on the right is precisely the c_1 part of a valid challenge ciphertext in Game 2. Also note that $v_0 = u_0^\top s + x$, just as the c_0 part of the challenge ciphertext in Game 2.

When $\mathcal{O} = \mathcal{O}_\$$ we have that v_0 is uniform in $\mathbb{Z}_q$ and v^* is uniform in $\mathbb{Z}_q^m$. Therefore c_1^* as defined in step (3) above is uniform and independent in $\mathbb{Z}_q^{2m}$ by the standard left over hash lemma (e.g. Theorem 8.38 of [29]) where the hash function is defined by the matrix $(A_0^\top | v^*)$. Consequently, the challenge ciphertext is always uniform in $\mathbb{Z}_q \times \mathbb{Z}_q^{2m}$, as in Game 3.

Guess. After being allowed to make additional queries, $\mathcal{A}$ guesses if it is interacting with a Game 2 or Game 3 challenger. Our simulator outputs $\mathcal{A}$'s guess as the answer to the LWE challenge it is trying to solve.

We already argued that when $\mathcal{O} = \mathcal{O}_s$ the adversary's view is as in Game 2. When $\mathcal{O} = \mathcal{O}_\$$ the adversary's view is as in Game 3. Hence, $\mathcal{B}$'s advantage in solving LWE is the same as $\mathcal{A}$'s advantage in distinguishing Games 2 and 3, as required. This completes the description of algorithm $\mathcal{B}$ and completes the proof.

6.5 Multi-bit Encryption

We briefly note that, as in [16], it is possible to reuse the same ephemeral encryption randomness s to encrypt multiple message bits. An N-bit message can thus be encrypted as N components c_0 plus a single component c_1, where the same ephemeral $s \in \mathbb{Z}_q^n$ is used throughout. The total ciphertext size with this technique is 1 element of $\mathbb{Z}_q$ for each bit of the message, plus a constant $2\,m$ elements of $\mathbb{Z}_q$ regardless of the message length. The ciphertext size is thus $(N + 2\,m)$ elements of $\mathbb{Z}_q$.

7 Extensions: HIBE and Adaptively-Secure IBE

In the full version of the paper [1] we show two extensions of the basic IBE construction from Section 6.2.

Adaptively secure IBE. Recall that Waters [31] showed how to convert the selectively-secure IBE in [6] to an adaptively secure IBE. We show that a similar technique, also used in Boyen [10], can convert our basic IBE construction to an adaptively secure IBE. We treat an identity id as a sequence of ℓ bits $\mathsf{id} = (b_1, \ldots, b_\ell)$ in $\{1, -1\}^\ell$. Then during encryption we use the matrix

$$F_{\mathsf{id}} := \left(A_0 \mid C + \sum_{i=1}^{\ell} b_i\, A_i \right) \quad \in \mathbb{Z}_q^{n \times 2m}$$

where $A_0, A_1, \ldots, A_\ell, C$ are matrices in the public parameters. The result is an adaptively secure lattice IBE, simpler and with shorter ciphertexts than the recent construction of Cash et al. [12].

Hierarchical IBE. We show how the basis delegation techniques from [12,24] can convert the basic IBE construction to an HIBE. For an identity $\mathsf{id} = (\mathsf{id}_1, \ldots, \mathsf{id}_\ell)$ at depth ℓ the matrix F_{id} used in encryption is defined as follows:

$$F_{\mathsf{id}} := \left(A_0 \mid A_1 + H(\mathsf{id}_1)B \mid \ \ldots \ \mid A_\ell + H(\mathsf{id}_\ell)B \right) \quad \in \mathbb{Z}_q^{n \times (\ell+1)m}$$

where $A_0, A_1, \ldots, A_\ell, B$ are matrices in the public parameters. We note that a recent HIBE construction in [2] gives a lattice-based HIBE where the lattice dimension does not grow with the identity's depth in the hierarchy.

8 Conclusion and Open Problems

We constructed an efficient identity-based encryption scheme and proven its security in the standard model from the LWE assumption (which is itself implied by worst-case lattice assumptions). In the full paper [1] we extend the basic selective-ID secure scheme to provide full adaptive-ID security, and to support a delegation mechanism to make it hierarchical.

It would be interesting to improve these constructions by adapting them to ideal lattices [30]. Another open problem is to construct an adaptively secure lattice-based IBE in the standard model where all the data is short (including the public parameters).

Acknowledgments

We are grateful to Chris Peikert for suggesting that we use the basis extension method from [24] to simplify the analysis of algorithm SampleLeft. This suggestion also let us to remove the matrix R from the master secret. We also thank Ron Rivest for pointing out that indistinguishability from random can help IBE systems resist subpoenas.

References

1. Shweta Agrawal, Dan Boneh, and Xavier Boyen. Efficient lattice (H)IBE in the standard model (2010); Full version of this paper. Available on the authors' web page
2. Agrawal, S., Boneh, D., Boyen, X.: Lattice basis delegation in fixed dimension and shorter-ciphertext hierarchical IBE (2010) (manuscript)
3. Agrawal, S., Boyen, X.: Identity-based encryption from lattices in the standard model (2009) (manuscript), http://www.cs.stanford.edu/~xb/ab09/
4. Ajtai, M.: Generating hard instances of the short basis problem. In: Wiedermann, J., Van Emde Boas, P., Nielsen, M. (eds.) ICALP 1999. LNCS, vol. 1644, pp. 1–9. Springer, Heidelberg (1999)
5. Alwen, J., Peikert, C.: Generating shorter bases for hard random lattices. In: Proc. of STACS 2009, pp. 75–86 (2009)
6. Boneh, D., Boyen, X.: Efficient selective-id secure identity-based encryption without random oracles. In: Cachin, C., Camenisch, J.L. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 223–238. Springer, Heidelberg (2004)
7. Boneh, D., Boyen, X.: Secure identity based encryption without random oracles. In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 443–459. Springer, Heidelberg (2004)
8. Boneh, D., Franklin, M.: Identity-based encryption from the Weil pairing. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 213–229. Springer, Heidelberg (2001)
9. Boneh, D., Gentry, C., Hamburg, M.: Space-efficient identity based encryption without pairings. In: Proc. of FOCS 2007, pp. 647–657 (2007)
10. Boyen, X.: Lattices niçoises and vanishing trapdoors: A framework for fully secure short signatures and more. In: PKC 2010. LNCS. Springer, Heidelberg (to appear, 2010)
11. Canetti, R., Halevi, S., Katz, J.: A forward-secure public-key encryption scheme. J. Cryptol. 20(3), 265–294 (2007)
12. Cash, D., Hofheinz, D., Kiltz, E.: How to delegate a lattice basis. Cryptology ePrint Archive, Report 2009/351 (2009), http://eprint.iacr.org/

13. Cocks, C.: An identity based encryption scheme based on quadratic residues. In: Proceedings of the 8th IMA Conference, pp. 26–28 (2001)
14. Cramer, R., Damgard, I.: On the amortized complexity of zero-knowledge protocols. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 177–191. Springer, Heidelberg (2009)
15. Dodis, Y., Ostrovsky, R., Reyzin, L., Smith, A.: Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. SIAM Journal on Computing 38(1), 97–139 (2008)
16. Gentry, C., Peikert, C., Vaikuntanathan, V.: Trapdoors for hard lattices and new cryptographic constructions. In: Proc. of STOC 2008, pp. 197–206 (2008)
17. Gentry, C.: Practical identity-based encryption without random oracles. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 445–464. Springer, Heidelberg (2006)
18. Gentry, C., Silverberg, A.: Hierarchical id-based cryptography. In: Zheng, Y. (ed.) ASIACRYPT 2002. LNCS, vol. 2501, pp. 548–566. Springer, Heidelberg (2002)
19. Hoeffding, W.: Probability inequalities for sums of bounded random variables. Journal of the American Statistical Association 58(301), 13–30 (1963)
20. Horwitz, J., Lynn, B.: Toward hierarchical identity-based encryption. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 466–481. Springer, Heidelberg (2002)
21. Lyubashevsky, V., Micciancio, D.: Generalized compact knapsacks are collision resistant. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) ICALP 2006. LNCS, vol. 4052, pp. 144–155. Springer, Heidelberg (2006)
22. Micciancio, D., Goldwasser, S.: Complexity of Lattice Problems: a cryptographic perspective, vol. 671. Kluwer Academic Publishers, Dordrecht (2002)
23. Micciancio, D., Regev, O.: Worst-case to average-case reductions based on gaussian measures. In: Proc. of FOCS 2004, pp. 372–381 (2004)
24. Peikert, C.: Bonsai trees (or, arboriculture in lattice-based cryptography). Cryptology ePrint Archive, Report 2009/359 (2009), http://eprint.iacr.org/
25. Peikert, C.: Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In: Proc. of STOC 2009, pp. 333–342. ACM, New York (2009)
26. Peikert, C., Rosen, A.: Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In: Halevi, S., Rabin, T. (eds.) TCC 2006. LNCS, vol. 3876, pp. 145–166. Springer, Heidelberg (2006)
27. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. In: Proc. of STOC 2005, pp. 84–93 (2005)
28. Shamir, A.: Identity-based cryptosystems and signature schemes. In: Blakely, G.R., Chaum, D. (eds.) CRYPTO 1984. LNCS, vol. 196, pp. 47–53. Springer, Heidelberg (1985)
29. Shoup, V.: A Computational Introduction to Number Theory and Algebra, 2nd edn. Cambridge University Press, Cambridge (2008)
30. Stehle, D., Steinfeld, R., Tanaka, K., Xagawa, K.: Efficient public-key encryption based on ideal lattices. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 617–635. Springer, Heidelberg (2009)
31. Waters, B.: Efficient identity-based encryption without random oracles. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 114–127. Springer, Heidelberg (2005)
32. Waters, B.: Dual key encryption: Realizing fully secure IBE and HIBE under simple assumption. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 619–636. Springer, Heidelberg (2009)