Expressive Encryption Systems from Lattices
(abstract from the invited lecture)

Xavier Boyen

Palo Alto Research Center

Abstract. In this survey, we review a number of the many "expressive" encryption systems that have recently appeared from lattices, and explore the innovative techniques that underpin them.

1 Introduction

Lattice-based cryptosystems are becoming an increasingly popular in the research community, owing to a unique combination of factors. On the one hand, lattice systems are often conceptually simple to understand and thus easy to implement by non-specialists, at least if one makes abstraction of the finer mathematical intricacies surrounding their security analysis. On the other, their soundness is backed by strong complexity-theoretic evidence that the underlying problems are suitably "hard", of which the most often repeated are the existence of certain average-case to worst-case equivalences [7, 19] and their conjectured resistance to quantum attacks. All those factors conspire to make lattices a prime choice, if not the primary one yet, for mathematical crypto design looking out into the future.

Although empirical uses of lattices have been made in commercial cryptography, they have had a rather slow start in research circles. For more than a decade, indeed, signature schemes and basic public-key encryption have essentially remained their sole confine [7, 19]. In the past few years, however, lattices have flourished into a theoretically solid, comprehensive framework, owing to the discovery of a few key concepts and techniques. This ushered the way to the construction of ever more powerful and expressive public-key encryption systems, writ large — a whole new world of cryptographic constructions waiting to be explored and conquered.

The search for encryption systems with complex functionalities arguably originates with the field of modern cryptography itself; but it is the arrival of bilinear maps, or pairings, that truly jumpstarted it, by providing such spectacular solutions to long-standing open problems as identity-based encryption [10]. Lattices are late to this game, and currently still lag in functionality and practicality with respect to pairing-based constructions. Nevertheless, an unmistakable shift from pairings to lattices is presently occurring in the research community, driven as much as the looming threat of quantum attacks that lattices seek to alleviate, as the sheer scientific draw of tackling tough problems from wholly new directions.

In this lecture, we set out to explore some of the recent advances in that search, and distill the essential new ideas that made them possible. [1]

2 Background

A lattice is an additive subgroup of $\mathbb{R}^n$; it is therefore generated by a *basis* of n (linearly independent) vectors in $\mathbb{R}^n$. In high dimensions, many computational problems on lattices are intractable, and in some cases are even known to be NP-hard. What makes lattices useful in cryptography, is that, though all bases are equivalent from a linear algebraic point of view, bases whose vectors have *low norm* can provide easy solutions to otherwise intractable lattice problems. For instance, the "closest vector problem" (which consists of finding a lattice point within a prescribed radius from a given reference in $\mathbb{R}^n$) becomes soluble if avails a low-norm lattice basis. Without such a *good basis*, this problem and many related ones remain intractable.

Whereas this asymmetry is, of course, central to lattices' use in asymmetric cryptography, general lattices as defined above are somewhat unwieldy to work with. One often prefers to restrict oneself to a restricted class of lattices with special properties; be it for reasons of convenience or efficiency, or both.

To wit, many of the recently developed expressive cryptosystems make use of Ajtai's lattices [6]. Those are sets of vectors $\boldsymbol{x} \in \mathbb{Z}^m$ that lie in the kernel of some $\mathbf{A} \in \mathbb{Z}^{n \times m}$ modulo some prime q, *i.e.*, defined by an equation $\mathbf{A} \cdot \mathbf{x} = 0 \pmod{q}$. Aside from their definitional convenience, Ajtai's lattices are appealing for two reasons: one of security, the other of flexibility. First, they induce rich and usable cryptographic key spaces, owing to the Regev's result that random instances are just as hard as worst-case ones [19]. Second, they are closely related to error-correction codes, and in particular the matrix $\mathbf{A}$ defines a "public" computational operator that can be effectively reverted with knowledge of a "private" trapdoor, as first shown by Gentry *et al.* [15]: the map $x \mapsto \mathbf{A} \cdot \boldsymbol{x}$, restricted to for *low-norm* inputs $x \in \mathbb{Z}^m$, can be reverted, in the sense of finding a colliding pre-image $x' \in \mathbb{Z}^m$, if one knows a good basis for the implied lattice. This combination of features — easy-to-sample key spaces and a kind of invertibility — are sought for in asymmetric cryptographic constructions.

By way of comparison, we mention that Gentry's fully homomorphic encryption scheme made extensive use of a different kind of lattices, constructed from polynomial rings, whose ring structure was crucial to realize full homomorphism in his original system.

[1] Around the same time, also appeared the first realization of "fully homomorphic encryption" [14], a hugely significant breakthrough of both theoretical and (one hopes) eventual practical significance. FHE undoubtedly contributed greatly to the general surge in lattice popularity, notwithstanding the quite different flavors of problems involved. FHE has since taken a life of its own, with the most recent performance and conceptual improvements seemingly taking it away from its lattice roots, and squarely into the realm of pure number theory. We refer the interested reader to the rapidly growing literature on the subject; see [16] for pointers.

Note. Due to space contraints, we do not give formal statements of the various notions and schemes in this abstract, but refer the reader to the original papers.

2.1 Lattice Notions

We let parameters q, m, n be polynomial functions of a security parameter λ.

Lattices. Let $\mathbf{B} = \left[\begin{array}{c|c|c} \boldsymbol{b}_1 & \ldots & \boldsymbol{b}_m \end{array} \right]$ be an $m \times m$ real matrix with linearly independent column vectors. It generates an m-dimensional full-rank lattice Λ,

$$\Lambda = \mathcal{L}(\mathbf{B}) = \left\{ \boldsymbol{y} \in \mathbb{R}^m \quad \text{s.t.} \quad \exists \boldsymbol{s} = (s_1, \ldots, s_m) \in \mathbb{Z}^m \ , \quad \boldsymbol{y} = \mathbf{B}\,\boldsymbol{s} = \sum_{i=1}^{m} s_i\, \boldsymbol{b}_i \right\}$$

Of interest to us is the case of integer lattices that are invariant under translation by multiples of some integer q in each of the coordinates, or Ajtai lattices.

Ajtai lattices (and their shifts). For q prime, $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and $\boldsymbol{u} \in \mathbb{Z}_q^n$, define:

$$\Lambda_q^{\perp}(\mathbf{A}) = \left\{ \boldsymbol{e} \in \mathbb{Z}^m \quad \text{s.t.} \quad \mathbf{A}\,\boldsymbol{e} = 0 \pmod q \right\}$$
$$\Lambda_q^{\boldsymbol{u}}(\mathbf{A}) = \left\{ \boldsymbol{e} \in \mathbb{Z}^m \quad \text{s.t.} \quad \mathbf{A}\,\boldsymbol{e} = \boldsymbol{u} \pmod q \right\}$$

Ajtai [6] first showed how to sample an essentially uniform matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, along with a full-rank set $\mathbf{T}_{\mathbf{A}} \subset \Lambda^{\perp}(\mathbf{A})$ of *low-norm* vectors or points on the lattice. We state an improved version of Ajtai's basis generator, from [8].

Trapdoors for lattices. Let $n = n(\lambda), q = q(\lambda), m = m(\lambda)$ be positive integers with $q \geq 2$ and $m \geq 5n \log q$. There exists a probabilistic polynomial-time algorithm $\mathsf{TrapGen}$ that outputs a pair of $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and $\mathbf{T}_{\mathbf{A}} \in \mathbb{Z}_q^{m \times m}$ such that $\mathbf{A}$ is statistically close to uniform and $\mathbf{T}_{\mathbf{A}}$ is a basis for $\Lambda^{\perp}(\mathbf{A})$ with "Gram-Schmidt" length $L = \|\widetilde{\mathbf{T}_{\mathbf{A}}}\| \leq m \cdot \omega(\sqrt{\log m})$, with all but $n^{-\omega(1)}$ probability.

2.2 Discrete Gaussians

Central to all cryptosystems based on Ajtai lattices, is the study of the distribution of various vectors of interest (e.g., preimages to the operation $\mathbf{A}$). Multidimensional discrete Gaussian distributions are particularly useful.

Discrete Gaussians. Let m be a positive integer and Λ an m-dimensional lattice over $\mathbb{R}$. For any vector $\boldsymbol{c} \in \mathbb{R}^m$ and any positive spread parameter $\sigma \in \mathbb{R}_{>0}$, let:

$\rho_{\sigma,c}(\boldsymbol{x}) = \exp\left(-\pi \frac{\|\boldsymbol{x}-\boldsymbol{c}\|^2}{\sigma^2}\right)$: a Gaussian function of center $\boldsymbol{c}$ and parameter σ;
$\rho_{\sigma,\boldsymbol{c}}(\Lambda) = \sum_{\boldsymbol{x} \in \Lambda} \rho_{\sigma,\boldsymbol{c}}(\boldsymbol{x})$: the infinite discrete sum of $\rho_{\sigma,\boldsymbol{c}}$ over the lattice Λ;
$\mathcal{D}_{\Lambda,\sigma,\boldsymbol{c}}$: the discrete Gaussian distribution on Λ of center $\boldsymbol{c}$ and parameter σ:

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

For convenience, we abbreviate $\rho_{\sigma,0}$ and $\mathcal{D}_{\Lambda,\sigma,0}$ respectively as ρ_{σ} and $\mathcal{D}_{\Lambda,\sigma}$.

2.3 Sampling and Preimage Sampling

The public-key and secret-key functions we need for asymmetric cryptography arise from the previous notions. Specifically, while anyone can sample a discrete Gaussian preimage with no prescription on its image under $\mathbf{A}$, only with a trapdoor or short basis $\mathbf{B}$ can one sample a preimage hitting a specific target image $\boldsymbol{u}$ with the same conditional distribution. The following results are due to Gentry, Peikert, and Vaikuntanathan [15]. They first construct an algorithm for sampling from the discrete Gaussian $\mathcal{D}_{\Lambda,\sigma,\boldsymbol{c}}$, given a basis $\mathbf{B}$ for the m-dimensional lattice Λ with $\sigma \geq \|\widetilde{\mathbf{B}}\| \cdot \omega(\sqrt{\log m})$. Next they give an algorithm that given an trapdoor and a target, can sample a preimage with the same (conditional) discrete Gaussian distribution.

Sampling a discrete Gaussian. There exists a probabilistic polynomial-time algorithm, denoted $\mathsf{SampleGaussian}$, that, on input an arbitrary basis $\mathbf{B}$ of an m-dimensional full-rank lattice $\Lambda = \mathcal{L}(\mathbf{B})$, a parameter $\sigma \geq \|\widetilde{\mathbf{B}}\| \cdot \omega(\sqrt{\log m})$, and a center $\boldsymbol{c} \in \mathbb{R}^m$, outputs a sample from a distribution statistically close to $\mathcal{D}_{\Lambda,\sigma,\boldsymbol{c}}$.

Preimage sampling from trapdoor. There exists a probabilistic polynomial-time algorithm, denoted $\mathsf{SamplePre}$, that, on input a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, a short trapdoor basis $\mathbf{T}_\mathbf{A}$ for $\Lambda_q^\perp(\mathbf{A})$, a target image $\boldsymbol{u} \in \mathbb{Z}_q^n$, and a Gaussian parameter $\sigma \geq \|\widetilde{\mathbf{T}_\mathbf{A}}\| \cdot \omega(\sqrt{\log m})$, outputs a sample $\boldsymbol{e} \in \mathbb{Z}_q^m$ from a distribution within negligible statistical distance of $\mathcal{D}_{\Lambda_q^{\boldsymbol{u}}(\mathbf{A}),\sigma}$.

Micciancio and Regev [17] show that the norm of vectors sampled from discrete Gaussians is small with high probability. We omit the full statement.

2.4 Hardness Assumption

One of the classic hardness assumptions associated with Ajtai lattices, refers to the LWE — Learning With Errors — problem, first stated by [19], and since extensively studied and used. For polynomially bounded modulus q, the computational and decisional versions of the problems are polynomially reducible to each other. We give the following statement of the decisional version.

The decisional LWE problem. 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 $\boldsymbol{s} \in \mathbb{Z}_q^n$, or, a truly random sampler $\mathcal{O}_\$$, whose behaviors are respectively as follows:

$\mathcal{O}_s$: outputs noisy pseudo-random samples of the form $(\boldsymbol{w}_i, v_i) = \left(\boldsymbol{w}_i,\ \boldsymbol{w}_i^T\,\boldsymbol{s} + x_i\right) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$, where, $\boldsymbol{s} \in \mathbb{Z}_q^n$ is a uniformly distributed persistent secret key that is invariant across invocations, $x_i \in \mathbb{Z}_q$ is a freshly generated ephemeral additive noise component with distribution χ, and $\boldsymbol{w}_i \in \mathbb{Z}_q^n$ is a fresh uniformly distributed vector revealed as part of the output.

$\mathcal{O}_\$$**:** outputs truly random samples $\big(\boldsymbol{w}_i, v_i\big) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$, drawn independently uniformly at random in the entire domain $\mathbb{Z}_q^n \times \mathbb{Z}_q$.

The $(\mathbb{Z}_q, n, \chi)$-LWE problem statement, or LWE for short, allows an unspecified number of queries to be made to the challenge oracle $\mathcal{O}$, with no stated prior bound. We say that an algorithm $\mathcal{A}$ decides the $(\mathbb{Z}_q, n, \chi)$-LWE problem if $\big|\Pr[\mathcal{A}^{\mathcal{O}_s} = 1] - \Pr[\mathcal{A}^{\mathcal{O}_\$} = 1]\big|$ is non-negligible for a random $s \in \mathbb{Z}_q^n$.

Average to worst case. The confidence in the hardness of the LWE problem stems in part from a result of Regev [19] which shows that the for certain noise distributions χ, the LWE problem is as hard as (other) classic lattice problems (such as SIVP and GapSVP) in the worst case, under a quantum reduction. A non-quantum reduction with different parameters was later given by Peikert [18]. We state Regev's result for reference below.

The Regev reduction theorem. Consider a real parameter $\alpha = \alpha(n) \in (0, 1)$ and a prime $q = q(n) > 2\sqrt{n}/\alpha$. 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 $x \in \mathbb{R}$. Denote by $\bar{\Psi}_\alpha$ 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. Then, if there exists an efficient, possibly quantum, algorithm for deciding the $(\mathbb{Z}_q, n, \bar{\Psi}_\alpha)$-LWE problem, there exists a quantum $q \cdot \mathit{poly}(n)$-time algorithm for approximating the SIVP and GapSVP problems, to within $\tilde{O}(n/\alpha)$ factors in the ℓ_2 norm, in the worst case.

Since the best known algorithms for 2^k-approximations of GapSVP and SIVP run in time $2^{\widetilde{O}(n/k)}$, it follows that the LWE problem with noise ratio $\alpha = 2^{-n^\epsilon}$ ought to be hard for some constant $\epsilon < 1$.

3 Classic Constructions

We start this presentation with the systems from which all recent developments are based, starting with Regev's minimalistic public-key cryptosystem.

3.1 Regev public-key encryption

The basic principle of Regev's original public-key cryptosystem is deceptively simple, as long as one does not delve too deep in its analysis. Paradoxically, Regev's system predated the GPV trapdoor preimage sampling, and required no other machinery than a basic random Ajtai lattice, not even a short basis.

The algorithms defining the system are as follows:

Key generation. Pick a suitable modulus q, a random Ajtai matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, and a short random vector $\boldsymbol{d} \in \mathbb{Z}^m$; and let $\boldsymbol{u} = \mathbf{A} \cdot \boldsymbol{d} \bmod q \in \mathbb{Z}_q^n$. The public and secret keys are:

$$\mathsf{PK} = (q, \mathbf{A}, \boldsymbol{u}) \qquad \mathsf{SK} = \boldsymbol{e}$$

Encryption. To encrypt a bit $m \in \{0,1\}$, pick a random vector $\boldsymbol{s} \in \mathbb{Z}_q^n$, a noise scalar $y_0 \sim \psi$, and a noise vector $\boldsymbol{y}_1 \sim \psi^m$, and output:

$$\mathsf{CT} = \left(c_0 = \boldsymbol{s}^\top \boldsymbol{u} + m \cdot \lfloor \frac{q}{2} \rfloor + y_0, \quad \boldsymbol{c}_1 = \mathbf{A}^\top \boldsymbol{s} + \boldsymbol{y}_1\right)$$

Decryption. The bit m is deemed to be 0 or 1, if the following quantity is respectively closer to 0 or $\frac{q}{2}$, modulo q:

$$c_0 - \boldsymbol{c}_1^\top \boldsymbol{d} \pmod q$$

It is easy to see that all terms cancel in the decryption operation, but for the noise contributions due to y_0 and $\boldsymbol{y}_1$ and the term $m \cdot \lfloor \frac{q}{2} \rfloor$ which redundantly encodes m. The noise is chosen sufficiently small so that, even after taking the inner product of $\boldsymbol{y}_1$ with the secret key vector $\boldsymbol{d}$, the message m remains recognizable. However, for an attacker who can only find *large* preimages of $\boldsymbol{u}$, decoding will fail as the noise will completely mask the message. Technically, the noise distribution ψ is chosen according to Regev's reduction theorem, so that semantic security of the system can be reduced to a worst-case lattice hardness assumption. We refer to Regev's paper for details.

Remark. We note that in Regev's original paper [19], the roles of $\boldsymbol{d}$ and $\boldsymbol{s}$ were reversed. The above is Regev's *dual*, more conveniently extended as we now describe.

3.2 GPV identity-based encryption

Gentry *et al.* [15] first showed how to realize identity-based encryption from lattices. In IBE, the public key is arbitrary, and the corresponding secret key can be "extracted" from it by a central authority that holds a special trapdoor.

The GPV system can be viewed as an instantiation of the Regev system, where instead of having a single fixed "syndrome" vector $\boldsymbol{u}$ (see the description above), said vector is made to depend on the recipient's identity using a hash function, as in $\boldsymbol{u}_{id} = H(id)$. Since no predetermined $\boldsymbol{d}$ can serve to deduce $\boldsymbol{u}$, a central authority will need the preimage sampling trapdoor to compute a short preimage $\boldsymbol{d}_{id}$ for any desired target $\boldsymbol{u}_{id}$; the trapdoor is thus the IBE master key.

Their system is described as follows:

System setup. Pick a suitable modulus q, and sample a random Ajtai matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ with associated trapdoor $\mathbf{B} \in \mathbb{Z}^{m \times m}$. The public parameters and master secret key are:

$$\mathsf{PP} = (q, \mathbf{A}) \qquad \mathsf{MK} = \mathbf{B}$$

Private key extraction. To extract a private key corresponding to a public identity id, first compute $\boldsymbol{u}_{id} = H(id) \in Z_q^n$, and then, using the trapdoor $\mathbf{B}$, find a short preimage $\boldsymbol{d}_{id}$, i.e., a low-norm vector such that $\mathbf{A} \cdot \boldsymbol{d}_{id} = \boldsymbol{u}_{id} \pmod q$. Output the private key as:

$$\mathsf{SK}_{\mathsf{id}} = d_{id}$$

Encryption. To encrypt a bit $m \in \{0,1\}$ for an identity id, compute $\boldsymbol{u}_{id} = H(id) \in Z_q^n$ and then encrypt as in the Regev system; i.e., picking a random vector $\boldsymbol{s} \in \mathbb{Z}_q^n$, a noise scalar $y_0 \sim \psi$, and a noise vector $\boldsymbol{y}_1 \sim \psi^m$, output:

$$\mathsf{CT} = \left(c_0 = \boldsymbol{s}^\top \boldsymbol{u}_{id} + m \cdot \lfloor \frac{q}{2} \rfloor + y_0, \quad \boldsymbol{c}_1 = \mathbf{A}^\top \boldsymbol{s} + \boldsymbol{y}_1 \right)$$

Decryption. Proceed as in the Regev system, using the private key $\boldsymbol{d}_{id}$; i.e., decrypt as 0 or 1 depending on whether the following is closer to 0 or $\frac{q}{2}$, modulo q:

$$c_0 - \boldsymbol{c}_1^\top \boldsymbol{d}_{id} \pmod{q}$$

The proof of security follows readily from that of Regev's system, given the properties of trapdoor preimage sampling, in the random-oracle model.

4 Techniques and Refinements

Building upon those earlier results, a number of significant refinements were quick to appear, showing that full security reductions were possible and practical, even in the standard model.

4.1 Bit-by-bit Standard-model IBE

The first step was taken concurrently by several teams [3, 13], that quickly figured out a way to realize IBE from lattices *in the standard model*, albeit with a stiff efficiency penalty over the GPV system.

The idea was to encode the identity not in Regev's vector $\boldsymbol{u}$ as in GPV (which required a random oracle), but in the matrix $\mathbf{A}$ itself, in a binary fashion reminiscent of the pairing-based from [12]. Specifically, for an ℓ-bit identity $id = (b_1, \ldots, b_\ell) \in \{0,1\}^\ell$, the matrix $A_{id} \in \mathbb{Z}_q^{n \times (\ell+1)m}$ is defined as the following concatenation of $\ell + 1$ constant matrices of dimension $n \times m$:

$$\mathbf{A}_{id} = \left[\mathbf{A}_0 \middle| \mathbf{A}_{1,b_1} \middle| \mathbf{A}_{2,b_2} \middle| \ldots \middle| \mathbf{A}_{\ell,b_\ell} \right]$$

from which the following relationship between (a user's) public and private key will be enforced:

$$\mathbf{A}_{id} \cdot \boldsymbol{d}_{id} = \boldsymbol{u} \pmod{q}$$

It is easy to see (but harder to prove) that all that is needed to find a short solution d_{id} in the above equation, is a preimage sampling trapdoor for *any* of the matrices $A_{i,\cdot}$ intervening in $\mathbf{A}_{id}$. Accordingly, all the submatrices $\mathbf{A}_{i,b_i}$ for $i \geq 1$ can be picked at random, as merely a trapdoor $\mathbf{B}_0$ for $\mathbf{A}_0$ suffices to find short preimages under the whole of $\mathbf{A}_{id}$. Hence, such shall be the IBE master key in the real system.

The full system is described as follows:

System setup. Pick a suitable modulus q, and sample a random Ajtai matrix $\mathbf{A}_0 \in \mathbb{Z}_q^{n \times m}$ with associated trapdoor $\mathbf{B}_0 \in \mathbb{Z}^{m \times m}$. Also sample 2ℓ random matrices $\mathbf{A}_{i,b} \in \mathbb{Z}_q^{n \times m}$ for $i \in [\ell]$ and $b \in \{0,1\}$, and a random vector $\boldsymbol{u} \in \mathbb{Z}_q^n$. The public parameters and master secret key are:

$$\mathsf{PP} = (q, \mathbf{A}_0, \{\mathbf{A}_{i,b}\}, \boldsymbol{u}) \qquad \mathsf{MK} = \mathbf{B}_0$$

Private key extraction. To extract a private key corresponding to a public identity id, using the trapdoor $\mathbf{B}_0$, find a low-norm vector $\boldsymbol{d}_{id}$ such that $\mathbf{A}_{id} \cdot \boldsymbol{d}_{id} = \boldsymbol{u} \pmod q$. The private key is: $\mathsf{SK}_{\mathsf{id}} = d_{id}$.

Encryption. Proceed as in the Regev system substituting $\mathbf{A}_{id}$ for $\mathbf{A}$.

Decryption. Proceed as in the Regev system, substituting $\boldsymbol{d}_{id}$ for $\boldsymbol{d}$.

The large matrix $\mathbf{A}_{id}$ renders the system rather inefficient, but enables a security proof against "selective-identity" attacks (where the attacker reveals the target id^* in advance) *in the standard model*. One builds a simulator that can extract private keys for all identities but the pre-announced target id^*. The simulator shall set itself up with a trapdoor for every submatrix $A_{i,(1-b_i^*)}$ where b_i^* is the i-th bit of the target identity — but not $\mathbf{A}_0$ (which shall be assembled from an LWE challenge to show a reduction). This way, the resulting concatenation $\mathbf{A}_{id}$ will have one or more known trapdoors, unless $id = id^*$.

4.2 All-at-once Standard-model IBE

Just like the "bit-by-bit" construction of [3, 13], above, is a lattice analogue to the pairing-based IBE by Canetti, Halevi, and Katz [12], a similar analogy can be made from the "all-at-once" pairing-based IBE by Boneh and Boyen [9], as a more efficient way to build a provably secure IBE in the standard model. The full analysis is due to Agrawal *et al.* [1] and is quite involved, but the construction is based on a simple principle.

Here, the recipient identity is encoded into the Regev matrix $\mathbf{A}$ all at once, without decomposing it bit by bit. Specifically, the Regev encryption matrix becomes (for constant $\mathbf{A}_0, \mathbf{A}_1, \mathbf{A}_2 \in \mathbb{Z}_q^{n \times m}$),

$$\mathbf{A}_{id} = \left[\mathbf{A}_0 \middle| \mathbf{A}_1 + id \cdot \mathbf{A}_2\right]$$

when the identity $id \in \mathbb{Z}_q$, or even, in all generality,

$$\mathbf{A}_{id} = \left[\mathbf{A}_0 \middle| \mathbf{A}_1 + \mathbf{M}_{id} \cdot \mathbf{A}_2\right]$$

when the identity $id \in \mathbb{Z}_q^n$, based on a straightforward deterministic encoding into a regular square matrix $\mathbf{M}_{id} \in \mathbb{Z}_q^{n \times n}$, such that any non-trivial difference $\mathbf{M}_{id_1} - \mathbf{M}_{id_2}$ is itself non-singular.

In the real system, the central authority will have a trapdoor for A_0, and thus be able to find short solutions $\boldsymbol{d}_{id}$ for every requested id in the equation (for constant $\boldsymbol{u} \in \mathbb{Z}_q^n$):

$$\mathbf{A}_{id} \cdot \boldsymbol{d}_{id} = \boldsymbol{u} \pmod q$$

In the simulation for the security reduction, one sets things up so that the simulator can extract private keys for all identities id except the challenge id^*. The matrix $\mathbf{A}_0$ is imposed from an external LWE challenge, and thus without a trapdoor. We set $A_1 = \mathbf{A}_0 \cdot \mathbf{R} - \mathbf{M}_{id^*} \cdot \mathbf{A}_2$, for some random $\mathbf{R} \in \{-1, 1\}^{m \times m}$. It follows that for all non-challenge identities, the encryption matrix reads:

$$\mathbf{A}_{id} = \left[\mathbf{A}_0 \middle| \mathbf{A}_0 \cdot \mathbf{R} + (\mathbf{M}_{id} - \mathbf{M}_{id^*}) \cdot \mathbf{A}_2\right]$$

For the challenge identity, the factor $(\mathbf{M}_{id} - \mathbf{M}_{id^*})$ in parentheses vanishes, and what is left is:

$$\mathbf{A}_{id^*} = \left[\mathbf{A}_0 \middle| \mathbf{A}_0 \cdot \mathbf{R}\right]$$

Agrawal *et al.* [1] give an algorithm to find short preimages under matrices $\mathbf{A}_{id}$ of this form,*without* a trapdoor for $\mathbf{A}_0$, provided one knows a trapdoor for $(\mathbf{M}_{id} - \mathbf{M}_{id^*}) \cdot \mathbf{A}_2$, which will simply be that of $\mathbf{A}_2$ provided that the factor in parentheses is regular. Their algorithm exploits the appearance of multiples of $\mathbf{A}_0$ on both sides of the concatenation to engineer a cancellation. Note that the role of the matrix $\mathbf{R}$ is to blind the simulation setup, so that it looks indistinguishable from the real system to an attacker. In the case where $id = id^*$, the term in $\mathbf{A}_2$ vanishes, and so with it any beneficial use of its trapdoor.

For completeness, we describe their system as follows:

System setup. Pick a suitable modulus q, and sample a random Ajtai matrix $\mathbf{A}_0 \in \mathbb{Z}_q^{n \times m}$ with associated trapdoor $\mathbf{B}_0 \in \mathbb{Z}^{m \times m}$. Also sample two random matrices $\mathbf{A}_1, \mathbf{A}_2 \in \mathbb{Z}_q^{n \times m}$, and a random vector $\boldsymbol{u} \in \mathbb{Z}_q^n$. The public parameters and master secret key are:

$$\mathsf{PP} = (q, \mathbf{A}_0, \mathbf{A}_1, \mathbf{A}_2, \boldsymbol{u}) \qquad \mathsf{MK} = \mathbf{B}_0$$

Private key extraction. To extract a private key corresponding to a public identity id, define its matrix encoding $\mathbf{M}_{id}$ and its encryption matrix $\mathbf{A}_{id} = [\mathbf{A}_0 | \mathbf{A}_1 + \mathbf{M}_{id} \cdot \mathbf{A}_2]$. Using the trapdoor $\mathbf{B}_0$, sample a low-norm vector $\boldsymbol{d}_{id}$ solution of $\mathbf{A}_{id} \cdot \boldsymbol{d}_{id} = \boldsymbol{u} \pmod q$. The private key is: $\mathsf{SK}_{\mathsf{id}} = d_{id}$.

Encryption. Proceed as in the Regev system substituting $\mathbf{A}_{id}$ for $\mathbf{A}$.

Decryption. Proceed as in the Regev system, substituting $\boldsymbol{d}_{id}$ for $\boldsymbol{d}$.

4.3 Adaptive or "Full" Security

A drawback of the previous systems is their need to relax the security notion, from a *bona fide* adaptive-identity attack to a less realistic selective-identity one, in order to achieve a reduction in the standard model (*sans* random oracle).

In [11], we propose a scheme and accompanying proof technique that address this limitation. The general idea is to set up the simulator to fail not on one but several possible challenge queries, using an efficient key-space partitioning technique that is quite specific to lattices. The full version of [1] describes the fully secure system and its proof.

5 Delegation and Hierarchies

A classic generalization of the notion of IBE is that of hierarchical IBE, where private-key holders can serve as local authorities to issuing private keys to any identity below them in the hierarchical tree of identities.

5.1 Concatenation-based Delegation

The first inroad into HIBE from lattice is due to Cash *et al.* [13], who in the same paper leverage their bit-by-bit IBE approach into a hierarchical scheme thanks to a trapdoor delegation mechanism of their design.

The principle is as follows. Let an Ajtai matrix $\mathbf{A}_0$ and its associated "good" trapdoor $\mathbf{T}_0$. Let $\mathbf{A}_1$ be an arbitrary matrix that is dimension-compatible with $\mathbf{A}_0$. Cash *et al.* provide an algorithm that transforms $\mathbf{A}_0$'s trapdoor $\mathbf{T}_0$ into a trapdoor $\mathbf{T}$ for the concatenated matrix $\mathbf{A} = [A_0|A_1]$, and in such a way that the new trapdoor $\mathbf{T}$ has only a slightly higher norm than the originating trapdoor $\mathbf{T}_0$. (While the norm might not increase at all under a naïve delegation process, the degradation of quality is a by-product of a necessary re-randomization step to ensure that the delegated basis cannot be used to reconstruct the delegator basis).

The Cash-Hofheinz-Kiltz-Peikert HIBE. Based on this delegation algorithm, Cash *et al.* [13] extend their bit-by-bit IBE scheme into a hierarchical scheme in a straightforward manner: subordinate identities are constructed by extending an identity prefix with additional bits; the corresponding encryption matrices are likewise constructed by concatenating additional sub-matrices to the right; and the corresponding private keys are obtained by invoking the delegation algorithm for such concatenations.

The (first) Agrawal-Boneh-Boyen HIBE. Based on the same CHKP delegation algorithm, Agrawal *et al.* [1] likewise extend their all-at-once IBE scheme into a hierarchical scheme, in the same straightforward manner.

5.2 Multiplicative In-Place Delegation

A second approach to delegation and HIBE, due to Agrawal *et al.* [2], relies not on concatenation, but on multiplication by invertible low-norm matrices. They propose a delegation mechanism that operates "in place", *i.e.*, without increasing the dimensions of the lattices or the number of elements in the matrices defining them.

Given a good basis $\mathbf{T_A}$ for an Ajtai lattice $\Lambda^{\perp}(\mathbf{A})$, they show how to create a (slightly less) good basis $\mathbf{T_B}$ for another lattice $\Lambda^{\perp}(\mathbf{B})$, whose defining matrix $\mathbf{B}$ has the same dimension as $\mathbf{A}$ and can be deterministically and publicly computed from $\mathbf{A}$. The delegation mechanism furthermore ensures that given $\mathbf{A}, \mathbf{B}$ and $\mathbf{T_B}$, it is difficult to recover $\mathbf{T_A}$ or any other a short basis for $\Lambda^{\perp}(\mathbf{A})$, thus ensuring the "one-wayness" of the delegation process.

Very informally, the delegated matrix $\mathbf{B} \in \mathbb{Z}_q^{n \times m}$ is defined from the delegator matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and a low-norm invertible public delegation matrix $\mathbf{R} \in \mathbb{Z}_q^{m \times m}$, as the product:

$$\mathbf{B} = \mathbf{A} \cdot \mathbf{R}^{-1}$$

Since $\mathbf{T_A}$ is a trapdoor for $\mathbf{A}$, *i.e.*, a short basis for $\Lambda^\perp(\mathbf{A})$, it follows that $\mathbf{A} \cdot \mathbf{T_A} = \mathbf{0} \pmod q$. Hence, we also have that $(\mathbf{A} \cdot \mathbf{R}^{-1}) \cdot (\mathbf{R} \cdot \mathbf{T_A}) = \mathbf{0} \pmod q$. Hence, $\mathbf{R} \cdot \mathbf{T_A} \in \mathbb{Z}_q^{m \times m}$ defines a basis for $\Lambda^\perp(\mathbf{B})$, and a "good" one since $\mathbf{R}$ has low norm. A final re-randomization step will ensure that the delegation cannot be undone, ensuring its "one-wayness".

The (second) Agrawal-Boneh-Boyen HIBE. Equipped with their delegation tool, Agrawal *et al.* [2] construct an efficient HIBE system, with provable security from the LWE assumption [19], and where the dimension of the keys and ciphertexts does not increase with the depth of the hierarchy. In particular, for shallow hierarchies, the efficiency of their system is directly comparable to the random-oracle non-hierarchical system of [15]. For deep hierarchies, the number of private key and ciphertext elements remains the same, but the bit-size of the modulus needs to increase linearly. This results in an HIBE system whose space complexity is only linear in the depth of the hierarchy.

6 Attributes and Predicates

To conclude this tour, we note a couple of brand new results, that concurrently demonstrated that encryption systems even more expressive than (H)IBE could be constructed from lattices — thereby breaking the "IBE barrier".

Lattice-based "fuzzy IBE". One system, due to Agrawal et al. [4], is a *Fuzzy IBE* system. Fuzzy IBE, a notion originally defined and constructed from bilinear maps in [20], was the first instance of what is now referred under the umbrella of *attribute-based encryption*. In Fuzzy IBE, decryption is conditioned upon an approximate rather than exact match between recipient attributes stated in the ciphertext, and those actually present in the actual recipient's private key.

Lattice-based "fuzzy IBE". The other system, due to Agrawal et al. [5], is an instance of *Predicate-based* encryption system, where decryption is controlled by the (non-)vanishing of the inner product of two vectors of attributes: one from the ciphertext, the other from the private key.

7 Conclusion

While a great many technical and conceptual challenges remain unsolved, if there is a lesson to be drawn from the many recent exciting developments in just a few focused areas of investigation, is that lattice-based cryptography is poised to jump from the sidelines to the mainstream, and find its place into all manners of real-world applications in the coming decades. We certainly look forward to this transformation.

References

1. S. Agrawal, D. Boneh, and X. Boyen. Efficient lattice (H)IBE in the standard model. In *Advances in Cryptology—EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 553–572. Springer, 2010.
2. S. Agrawal, D. Boneh, and X. Boyen. Lattice basis delegation in fixed dimension and shorter-ciphertext hierarchical IBE. In *Advances in Cryptology—CRYPTO 2010*, volume 6223 of *LNCS*, pages 98–115. Springer, 2010.
3. S. Agrawal and X. Boyen. Identity-based encryption from lattices in the standard model. Manuscript, July 2009. http://www.cs.stanford.edu/ xb/ab09/.
4. S. Agrawal, X. Boyen, V. Vaikuntanathan, P. Voulgaris, and H. Wee. Fuzzy identity based encryption from lattices. Cryptology ePrint Archive, Report 2011/414, 2011. http://eprint.iacr.org/.
5. S. Agrawal, D. Freeman, and V. Vaikuntanathan. Functional encryption for inner product predicates from learning with errors. In *ASIACRYPT 2011*, 2011.
6. M. Ajtai. Generating hard instances of lattice problems (extended abstract). In *Proceedings of STOC 1996*, pages 99–108, New York, NY, USA, 1996. ACM.
7. M. Ajtai and C. Dwork. A public-key cryptosystem with worst-case/average-case equivalence. In *STOC*, pages 284–293, 1997.
8. J. Alwen and C. Peikert. Generating shorter bases for hard random lattices. In *STACS*, pages 75–86, 2009.
9. D. Boneh and X. Boyen. Efficient selective identity-based encryption withoutrandom oracles. *J. Cryptology*, 24(4):659–693, 2011. Abstract in EUROCRYPT 2004.
10. D. Boneh and M. Franklin. Identity-based encryption from the Weil pairing. In J. Kilian, editor, *Advances in Cryptology—CRYPTO 2001*, volume 2139 of *LNCS*, pages 213–29. Springer-Verlag, 2001.
11. X. Boyen. Lattice mixing and vanishing trapdoors – a framework for fully secure short signatures and more. In *Public Key Cryptography—PKC 2010*, volume 6056 of *LNCS*, pages 499–517, 2010.
12. R. Canetti, S. Halevi, and J. Katz. A forward-secure public-key encryption scheme. *J. Cryptology*, 20(3):265–294, 2007. Abstract in EUROCRYPT 2003.
13. D. Cash, D. Hofheinz, E. Kiltz, and C. Peikert. Bonsai trees or, how to delegate a lattice basis. In *Advances in Cryptology—EUROCRYPT 2010*, 2010.
14. C. Gentry. Fully homomorphic encryption using ideal lattices. In *STOC*, pages 169–178, 2009.
15. C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pages 197–206. ACM, 2008.
16. S. Halevi. Fully homomorphic encryption. slides from tutorial session, CRYPTO 2011, 2011. http://www.iacr.org/conferences/crypto2011/slides/Halevi.pdf.
17. D. Micciancio and O. Regev. Worst-case to average-case reductions based on gaussian measures. In *Proceedings of FOCS 2004*, pages 372–381, Washington, DC, USA, 2004. IEEE Computer Society.
18. C. Peikert. Bonsai trees (or, arboriculture in lattice-based cryptography). Cryptology ePrint Archive, Report 2009/359, 2009. http://eprint.iacr.org/.
19. O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In *Proceedings of STOC 2005*, pages 84–93, New York, NY, USA, 2005. ACM.
20. A. Sahai and B. Waters. Fuzzy identity-based encryption. In *EUROCRYPT 2005*, pages 457–473, 2005.