

Identity-Based, Attribute-Based and Predicate Encryption

We will start by constructing a digital signature scheme from lattices; then combine it with the dual Regev encryption scheme to construct an identity-based encryption scheme; then upgrade it to an attribute-based encryption scheme borrowing ideas from our FHE construction; and finally, a predicate encryption scheme, if time permits.

1 Digital Signatures

Here is a simple digital signature scheme. (For a definition of digital signatures and what we mean by a secure digital signature, see [Rafael Pass and abhi shelat's book](#).)

- The key generation algorithm samples a function together with a trapdoor. This would be \mathbf{A} and \mathbf{T} . The public key is \mathbf{A} and the secret key is \mathbf{T} .
- To sign a message m , first map it into the range of the function, e.g., by hashing it. That is, compute $\mathbf{v} = H(m)$. The signature is an inverse of \mathbf{v} under the function $g_{\mathbf{A}}$. That is, a short vector \mathbf{e} such that $\mathbf{A}\mathbf{e} = \mathbf{v} \pmod{q}$. This is guaranteed by the surjectivity of the function $g_{\mathbf{A}}$.
- Verification, given a message m , public key \mathbf{A} and signature \mathbf{e} , consists of checking that $\mathbf{A}\mathbf{e} = H(m) \pmod{q}$ and that $\|\mathbf{e}\|_{\infty} \leq m^2$.

Unforgeability (given no signature queries) reduces to SIS in the random oracle model, i.e., assuming that H is a random oracle.

However, given signatures on adversarially chosen messages (in fact, even random messages), this scheme is broken. The key issue is that there are many inverses of $H(m)$, and the particular inverse computed using a trapdoor \mathbf{T} leaks information about \mathbf{T} . Collecting this leakage over sufficiently many (polynomially many) signature queries enables an adversary to find \mathbf{T} , allowing her to forge signatures at will going forward.

This is most easily seen when the inversion procedure for $g_{\mathbf{A}}$ uses the inhomogenous trapdoor. Note that given \mathbf{v} , an adversary can compute $\mathbf{G}^{-}(\mathbf{v}) = \mathbf{e}'$ herself. She now gets a signature

$$\sigma = \mathbf{R} \cdot \mathbf{e}'$$

which gives her one equation on the secret \mathbf{R} . Given about m equations, she can solve linear equations and learn \mathbf{R} .

The situation remains essentially as dire even if you use the trapdoor (as opposed to the inhomogenous trapdoor). Using rounding vs the nearest plane algorithm does not help either; see the paper of [Nguyen and Regev](#) for robust attacks against this signature scheme. The fundamental difficulty seems to stem from the fact that the inversion procedure is deterministic!

To mitigate the difficulty, we need a special kind of inverter for $g_{\mathbf{A}}$. The inverter is a “pre-image sampler”; that is, it is given the trapdoor \mathbf{T} and produces a “random” pre-image. More precisely, we need the following distributions to be statistically close (computational indistinguishability is fine, but we will achieve statistical closeness):

$$\left(\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}, \mathbf{e} \leftarrow D_{\mathbb{Z}^m, s}, \mathbf{v} := \mathbf{A}\mathbf{e} \pmod{q} \right) \\ \approx_s \left(\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}, \mathbf{e} \leftarrow \text{PreSamp}(\mathbf{A}, \mathbf{T}, \mathbf{v}), \mathbf{v} \leftarrow \mathbb{Z}_q^n \right)$$

That is, the following processes produce statistically close outputs: (a) first sample \mathbf{e} from a discrete Gaussian, and deterministically set \mathbf{v} to be $\mathbf{A}\mathbf{e} \pmod{q}$; and (b) sample \mathbf{v} uniformly and use the pre-image sampler to produce an inverse of \mathbf{v} under $g_{\mathbf{A}}$ that is distributed according to the right conditional distribution. This distribution happens to be the discrete Gaussian over a coset of the lattice, that is,

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

In fact, this not quite enough; we need a multi-sample version of this. That is,

$$\begin{aligned} & \left(\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}, \{\mathbf{e}_i \leftarrow D_{\mathbb{Z}^m, s}, \mathbf{v} := \mathbf{A}\mathbf{e}_i \pmod{q}\}_{i=1}^{\text{poly}(\lambda)} \right) \\ & \approx_s \left(\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}, \{\mathbf{e} \leftarrow \text{PreSamp}(\mathbf{A}, \mathbf{T}, \mathbf{v}), \mathbf{v} \leftarrow \mathbb{Z}_q^n\}_{i=1}^{\text{poly}(\lambda)} \right) \end{aligned}$$

This is quite cumbersome to work with, so we propose an alternate stronger definition. That is, we require that for most $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$ and any trapdoor \mathbf{T} of length bounded by ℓ and $s \gg \ell$:

$$\begin{aligned} & \left(\mathbf{e} \leftarrow D_{\mathbb{Z}^m, s}, \mathbf{v} := \mathbf{A}\mathbf{e} \pmod{q} \right) \\ & \approx_s \left(\mathbf{e} \leftarrow \text{PreSamp}(\mathbf{A}, \mathbf{T}, \mathbf{v}), \mathbf{v} \leftarrow \mathbb{Z}_q^n \right) \end{aligned}$$

Proof of Security. With the one change that the inverter is replaced by a pre-image sampler, our signature scheme becomes secure in the random oracle model. We showed the proof in the class.

2 Discrete Gaussian Sampling

Throughout, we will deal with sampling from a zero-centered discrete Gaussian.

2.1 Naïve Sampling

Let us first consider sampling from a discrete Gaussian over the simplest possible lattice, namely the one-dimensional lattice of integers \mathbb{Z} . The first idea to sample from the discrete Gaussian $D_{\mathbb{Z}, s}$ is to sample from a continuous Gaussian N_s with parameter s and round to the nearest integer. Unfortunately, this is not a discrete Gaussian, not even statistically close to it. This is true even if s is much larger than the smoothing parameter. You will show this in the problem set.

For n -dimensional lattices, this statistical distance degrades with n as well making the situation much worse.

The reader may recall that the first step of Regev's worst-case to average-case reduction was sampling from a discrete Gaussian over a lattice for which Regev used the above procedure. However, he could afford to use an exponential s which makes the statistical distance small.

2.2 Sampling Discrete Gaussians over \mathbb{Z}

So, how do we sample from $D_{\mathbb{Z}, s}$ for polynomial s ? We will show that the general method of rejection sampling works. Let $Z = [-t \cdot s, t \cdot s]$ be a sufficiently large interval, where $t = \omega(\sqrt{\log \lambda})$. We do the following:

1. Sample a random integer $z \leftarrow Z$.
2. Output z with probability $\rho_s(z) := e^{-\pi z^2/s^2}$; else go to step 1 and repeat.

First of all, we will show that the probability that $D_{Z,s}$ assigns to numbers outside of the interval Z is negligible.

Lemma 1. *Let $s \geq \eta_\varepsilon(\mathbb{Z})$ for some $\varepsilon = \text{negl}(\lambda)$, and $t > 0$. We have*

$$\Pr_{x \leftarrow D_{Z,s}} [|x| > t \cdot s] \leq c \cdot e^{-\pi t^2}$$

for some absolute constant $c > 0$.

Consider the probability distribution $D'_{Z,s}$ which assigns probability $\rho_s(x)$ for all $x \in Z \cap \mathbb{Z}$ and 0 otherwise. The lemma above shows that $D'_{Z,s}$ is close to $D_{Z,s}$ if s is larger than the $\text{negl}(\lambda)$ -smoothing parameter of \mathbb{Z} , namely $\omega(\sqrt{\log n})$, and $t = \omega(\sqrt{\log n})$.

It is not hard to see that the procedure above samples from the distribution $D'_{Z,s}$ exactly. It remains to see that it terminates in polynomial time. We show two things which we leave as an exercise: (a) the probability that z sampled in step 1 lies in $[-s, s]$ is $\Omega(1/t)$ and (b) if such a z is sampled, it is output with probability $\Omega(1)$. Put together, the expected time for termination is $O(t) = \text{poly}(\lambda)$.

2.3 Sampling Discrete Gaussians over lattices: Klein-GPV algorithm

We demonstrate the sampler in two dimensions. The generalization to n dimensions follows quite naturally. We will use Figure 1 as a reference.

The lattice $\mathcal{L}(\mathbf{b}_1, \mathbf{b}_2)$ in the figure can be divided into many cosets of the lattice generated by \mathbf{b}_2 alone. The i^{th} coset, for every $i \in \mathbb{Z}$ is

$$\mathcal{L}_i = \{i \cdot \mathbf{b}_1 + y \cdot \mathbf{b}_2 : y \in \mathbb{Z}\}$$

Let's first ask what's the probability that a sample from the discrete Gaussian $D_{\mathcal{L},s,c}$ (for some center c) falls into the i^{th} coset. This is simply the Gaussian mass of \mathcal{L}_i divided by the discrete Gaussian mass of the entire lattice; that is,

$$\Pr[i] = \frac{\sum_{\mathbf{x} \in \mathcal{L}_i} \rho_s(\mathbf{x})}{\sum_{\mathbf{x} \in \mathcal{L}} \rho_s(\mathbf{x})} = \frac{\sum_{\mathbf{x} \in \mathcal{L}_i} e^{-\pi \|\mathbf{x}\|^2/s^2}}{\sum_{\mathbf{x} \in \mathcal{L}} e^{-\pi \|\mathbf{x}\|^2/s^2}} = \frac{\sum_{\mathbf{x} \in \mathcal{L}_i} e^{-\pi \|\mathbf{x}\|^2/s^2}}{\sum_{i \in \mathbb{Z}} \sum_{\mathbf{x} \in \mathcal{L}_i} e^{-\pi \|\mathbf{x}\|^2/s^2}}$$

Now, note that each vector $\mathbf{x} \in \mathcal{L}_i$ is

$$\mathbf{x} = i \cdot \mathbf{b}_1 + \mathbf{x}'$$

where $\mathbf{x}' \in \mathcal{L}(\mathbf{b}_2)$. Letting $\mathbf{b}_1 = \tilde{\mathbf{b}}_1 + \mathbf{b}_1^{\parallel}$ be the orthogonal decomposition of \mathbf{b}_1 into a component \mathbf{b}_1^{\parallel} in $\text{Span}(\mathbf{b}_2)$ plus the (Gram Schmidt) orthogonal component $\tilde{\mathbf{b}}_1$, lets us write the latter expression as

$$\frac{e^{-\pi \cdot i^2 \cdot \|\tilde{\mathbf{b}}_1\|^2/s^2} \cdot \sum_{\mathbf{x}' \in \mathcal{L}(\mathbf{b}_2)} e^{-\pi \|\mathbf{x}' + i \mathbf{b}_1^{\parallel}\|^2/s^2}}{\sum_{i \in \mathbb{Z}} e^{-\pi \cdot i^2 \cdot \|\tilde{\mathbf{b}}_1\|^2/s^2} \cdot \sum_{\mathbf{x}' \in \mathcal{L}(\mathbf{b}_2)} e^{-\pi \|\mathbf{x}' + i \mathbf{b}_1^{\parallel}\|^2/s^2}}$$

The truly key observation here is that the sum in the numerator and each of the inner sums in the denominator are nearly the same, upto negligible factors, as long as s is above the smoothing parameter of the lattice $\mathcal{L}(\mathbf{b}_1)$ (!) See Lemma 10 in Lecture 3.




Figure 1: Given a basis $(\mathbf{b}_1, \mathbf{b}_2)$ of the lattice in the picture, and a Gaussian standard deviation parameter s , the goal is to sample from a discrete Gaussian over $\mathcal{L}(\mathbf{b}_1, \mathbf{b}_2)$ with parameter s .

This vastly simplifies the expression to

$$\frac{e^{-\pi \cdot i^2 \cdot \|\tilde{\mathbf{b}}_1\|^2 / s^2}}{\sum_{i \in \mathbb{Z}} e^{-\pi \cdot i^2 \cdot \|\tilde{\mathbf{b}}_1\|^2 / s^2}}$$

which is simply a one-dimensional discrete Gaussian mass of the number $i \cdot \|\tilde{\mathbf{b}}_1\|^2$ on a lattice generated by multiples of $\tilde{\mathbf{b}}_1$ w.r.t. a Gaussian parameter s .

As we saw above, in the algorithm for sampling a discrete Gaussian over \mathbb{Z} , we can sample this efficiently as long as

$$s \gg \|\tilde{\mathbf{b}}_1\| \cdot \omega(\sqrt{\log n})$$

Once we sample the right coset of the lattice, i.e. the right \mathcal{L}_i , this now recursively reduces to the problem of sampling from a (non-zero-centered, in general) discrete Gaussian on a lattice one dimension smaller. Keep going! Sampling will work as long as

$$s \gg \max_{i \in [n]} (\|\tilde{\mathbf{b}}_i\|) \cdot \omega(\sqrt{\log n})$$

This is the Klein/GPV Gaussian sampling algorithm.

3 Identity-based Encryption

Let us think first about deploying a public-key encryption scheme on a large scale. We need a mechanism to maintain a directory of (ID, PK) pairs where ID is the identifying information of a person, say Alice's e-mail address or phone number, that other people use to send her a message. Then, when you wish to send an email to Alice, you look up her public key in the directory and encrypt to the public key.

The directory, which forms part of a public-key infrastructure (PKI), has to be authenticated and trusted. For example, an adversary should not be able to insert an entry of the form (ID_A, PK'_A) , where she presumably knows SK'_A , into the directory.

Identity-based encryption (IBE) solves the problem of having to maintain an authenticated PKI. In an IBE:

- there is a master authority who generates a master public key MPK together with a master secret key MSK , and publishes the MPK .
- To encrypt a message μ , one needs to know MPK and the identity ID (e.g., the e-mail address) of the recipient.
- Each user goes to the master authority and receives SK_{ID} after authenticating that they indeed are the owner of ID .
- Using SK_{ID} , the user can decrypt ciphertexts encrypted to the identity ID .

Let us now define the syntax of an IBE, formalizing the discussion above.

- $\text{Setup}(1^\lambda)$: is a probabilistic algorithm that generates a master public key MPK and a master secret key MSK .
- $\text{Enc}(MPK, ID, \mu)$: is a probabilistic algorithm that generates a ciphertext C of a message μ (for simplicity, we will encrypt bits but that is largely irrelevant) w.r.t. identity ID .
- $\text{KeyGen}(MSK, ID \in \{0, 1\}^*)$: is a probabilistic algorithm that generates a secret key SK_{ID} .
- $\text{Dec}(SK_{ID}, C)$: is a deterministic decryption algorithm.

You may have noticed that the master authority can decrypt *all* the ciphertexts generated in this system and is therefore *very powerful*.

Application: Access Delegation across Space. I can act as the master authority and use an IBE to delegate decryption of certain subsets of messages to other people (e.g., my administrative assistant). For example, all messages are tagged with a keyword $ID = \text{CS294}$, and I can issue the SK_{ID} to my assistant that lets him decrypt only those messages tagged with ID .

Application: Access Delegation across Time. Imagine that I go on (virtual) vacation to Cancun and want to take my laptop. However, I am worried that it will be stolen. So, I ask folks encrypting messages to me to use an IBE and tag the messages with an ID which is the current date. This allows me to generate a small set of secret keys, corresponding to the days that I am away, which allows me to decrypt only the corresponding small subset of messages. IBE lets me enjoy my vacation worry-free!

Application: Chosen-Ciphertext Security. IBE can be used in surprisingly non-trivial ways to construct other cryptographic systems, e.g., chosen ciphertext secure public-key encryption schemes and digital signature schemes (that we will describe later in this lecture).

Constructions. The first constructions used bilinear maps on elliptic curves (Boneh-Franklin'00) and quadratic residuosity (Cocks'00). We will present the third IBE scheme from the LWE assumption (Gentry-Peikert-Vaikuntanathan'08) and several variants today. Recently, Garg and Dottling have come up with a completely different scheme that relies on Diffie-Hellman groups (no need for bilinear maps!) Following up, Brakerski-Lombardi-Segev-Vaikuntanathan came up with a scheme based on learning parity with very low noise.

3.1 Definitions of Security

We imagine a PPT adversary that plays the following game with a challenger. This captures the requirement that encryptions relative to ID^* should be secure even to an adversary that can obtain secret keys for polynomially many *different* identities $ID \neq ID^*$. This is called the *adaptive security* or *full security* definition. The weaker *selective* security definition restricts the adversary to pick the identity it is attacking at the very beginning of the game (before it receives MPK).

Selectively secure IBE schemes can be generically proven to be fully secure under a sub-exponentially stronger assumption. Therefore, we will not attempt to optimize the strength of the assumption and focus on selective security for this lecture.

3.2 IBE=Signatures+Public-Key Encryption

Moni Naor observed that any IBE scheme gives us *for free* a digital signature scheme. **The intuition is that the identity secret key SK_{ID} can act as a signature for the “message” ID .** How so?

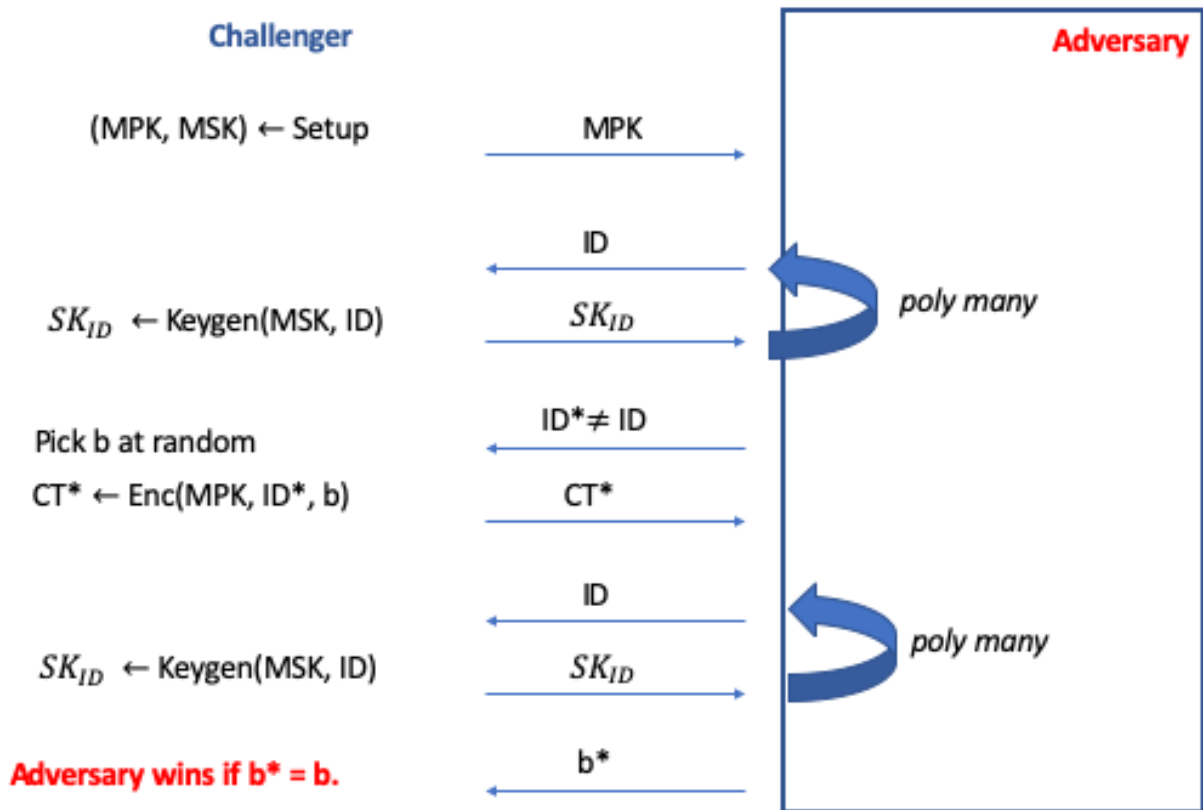
- It can be generated using the master secret key MSK (which will serve as the secret signing key.)
- It can be verified using the master public key MPK – indeed, encrypt a bunch of random messages using MPK and attempt to use the “signature” to decrypt. If decryption produces the correct message, accept the signature. Otherwise, reject.
- after receiving signatures SK_{ID} on polynomially many messages ID , being able to produce the “signature” on a different message ID^* constitutes a signature forgery; but being able to do that breaks IBE security. Conversely, in a signature scheme derived from a secure IBE scheme, it should be infeasible to do that.

Indeed, turning this around, we will use the GPV signature scheme we saw in the last class as a starting point to build an IBE scheme.

4 Recap: GPV Signatures

- $\text{KeyGen}(1^\lambda)$: Generate a random matrix $A \in \mathbb{Z}_q^{n \times m}$ and its trapdoor T by running TrapSamp .
- $\text{Sign}(\mu)$: first compute $v = H(\mu) \in \mathbb{Z}_q^n$ where H is treated as a random oracle in the analysis. Then, use Gaussian sampling (via the GPV algorithm) to compute a Gaussian solution $e \in \mathbb{Z}^m$ to the equation

$$Ae = v \pmod{q}$$



Let $\Lambda^\perp(\mathbf{A})$ denote the lattice

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

and let $\Lambda_{\mathbf{v}}^\perp(\mathbf{A})$ denote a coset of $\Lambda^\perp(\mathbf{A})$ indexed by \mathbf{v} . That is,

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

Note that the distribution of \mathbf{e} is $D_{\Lambda_{\mathbf{v}}^\perp(\mathbf{A}), \sigma}$ where $\sigma \approx \|\mathbf{T}\| \cdot \omega(\sqrt{\log n})$. (The $\omega(\sqrt{\log n})$ is so that the sampling algorithm can achieve negligible statistical distance from a true discrete Gaussian.)

- $\text{Verify}(\mathbf{A}, \mathbf{e}, \mu)$: check that (1) \mathbf{e} is short, that is $\|\mathbf{e}\| \leq \|\mathbf{T}\| \cdot \omega(\sqrt{n \log n})$; and (2) $\mathbf{A}\mathbf{e} = H(\mu) \pmod{q}$.

The key question now is how to we build an encryption algorithm whose public key is \mathbf{v} (which will be treated as $H(\text{ID})$) and the corresponding private key is \mathbf{e} as above. Indeed, we have seen precisely such a scheme in the first lecture (cf. lecture notes) called the GPV encryption scheme or more commonly, the dual-Regev encryption scheme.

But before we get there, the scheme as stated above is insecure – do you see why? Bonus points if you see how to fix it.

5 The Dual Regev Encryption Scheme

Let's recall the dual Regev scheme we saw in the first lecture (notes).

- KeyGen: the public key is an LWE matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and a random vector $\mathbf{v} \in \mathbb{Z}_q^n$. The private key is a short vector \mathbf{e} such that $\mathbf{A}\mathbf{e} = \mathbf{v} \pmod{q}$.

$$\text{pk} = (\mathbf{A}, \mathbf{v}) \quad \text{sk} = \mathbf{e}$$

- Enc(pk, μ): pick an LWE secret $\mathbf{s} \in \mathbb{Z}_q^n$ and output

$$(\mathbf{c}_1^T, c_2) := \left(\mathbf{s}^T \mathbf{A} + \mathbf{x}^T, \mathbf{s}^T \mathbf{v} + x' + m \lfloor q/2 \rfloor \right)$$

as the ciphertext. **We will call this ciphertext the dual Regev encryption of μ relative to \mathbf{A} and \mathbf{v} .**

- Dec(sk, (\mathbf{c}_1^T, c_2)): Compute

$$\tilde{\mu} := \text{Round}(c_2 - \mathbf{c}_1^T \mathbf{e})$$

where $\text{Round}(\alpha)$ outputs 1 if $|\alpha - q/2| \leq q/4$ and 0 otherwise.

We will leave the correctness and security as an exercise. (Alternatively, look at lecture 1.)

6 The GPV IBE Scheme

- Setup(1^λ): Pick the right $n = n(\lambda)$ for a security level of λ bits. Generate a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and its trapdoor $\mathbf{T} \in \mathbb{Z}^{m \times m}$ by running the trapdoor sampling algorithm.

$$(\mathbf{A}, \mathbf{T}) \leftarrow \text{TrapSamp}(1^n)$$

(The parameters m and q are picked internally by the trapdoor sampling algorithm.) The master public key is $\text{mpk} = \mathbf{A}$ and the master secret key is $\text{msk} = \mathbf{T}$.

- KeyGen(msk, ID): Compute $\mathbf{v} := H(ID) \in \mathbb{Z}_q^n$ where $H : \{0, 1\}^* \rightarrow \mathbb{Z}_q^n$ is a hash function (which, in the security analysis, will be treated as a random oracle.) Generate a short vector

$$\mathbf{e} \leftarrow \text{DGSamp}(\mathbf{A}, \mathbf{T}, \mathbf{v})$$

by running the discrete Gaussian sampling algorithm. Recall that $\mathbf{A}\mathbf{e} = \mathbf{v} \pmod{q}$. Output the secret key $\text{sk}_{ID} = \mathbf{e}$.

- Enc(mpk, ID, μ): Run the dual Regev encryption algorithm with $\text{pk} := (\mathbf{A}, \mathbf{v} = H(ID))$ and message μ and output the resulting ciphertext.
- Dec(sk_{ID}, c): Run the dual Regev decryption algorithm with $\text{sk} := \text{sk}_{ID} = \mathbf{e}$.

6.1 Proof of (Full) Security

We will come up with alternate algorithms called Setup^* , KeyGen^* and Enc^* (Dec^* will be the same as Dec) which the challenger will run. Our goal will be to show that (1) the adversary cannot distinguish between the challenger running Algorithm vs Algorithm^* and (2) Algorithms^* do not need the master secret key and moreover, a challenger using Algorithm^* can use a successful adversary to break LWE.

A crucial advantage of Algorithm^* for the GPV scheme is that it can use the programmability of the random oracle as we will see below. We will for simplicity first create algorithms for the selective security game.

- $\text{Setup}^*(ID^*, 1^\lambda)$: Sample random \mathbf{A}^* which forms the MPK^* (no need for trapdoor).
- $\text{Hash}^*(ID)$: Set $H(ID^*) = \mathbf{v}^*$, a random vector in \mathbb{Z}_q^n . For all other ID s, set $H(ID) = \mathbf{A}^* \mathbf{e}_{ID}$ where \mathbf{e}_{ID} is chosen from a Gaussian. Remember \mathbf{e}_{ID} .
- $\text{KeyGen}^*(ID)$: We know that $ID \neq ID^*$. So, we know the \mathbf{e}_{ID} by construction! **This is a consequence of working in the random oracle model!**
- $\text{Enc}^*(MPK^*, ID^*, \mu)$: return the dual Regev encryption of μ relative to \mathbf{A}^* and \mathbf{v}^* .

The Algorithm^* produce the same distribution as the original algorithms. Thus, an adversary will break the challenge ciphertext when interacting with Algorithm^* just as well as with Algorithms. By embedding the dual-Regev challenge matrix \mathbf{A} as the master public key and the dual-Regev public key \mathbf{v}^* as the hash of ID^* , we can easily turn the IBE adversary into an attack against the dual Regev public key encryption scheme.

A Note on Full Security. Since Setup^* does not know ID^* , it guesses which of the (polynomially many) hash queries will be for ID^* . (1) any adversary that succeeds has to know $H(ID^*)$ which it can only find out by making a hash query; and (2) if the guess is correct (happens with probability $1/Q$) we can translate an IBE breaker into a dual-Regev breaker just as above.

7 The CHKP IBE Scheme (Advanced Material: Not Covered in Class)

The CHKP Trick: Trapdoor Extension.

Given the trapdoor for a matrix \mathbf{A} , can you generate a trapdoor for $[\mathbf{A}||\mathbf{B}]$ where \mathbf{B} is an arbitrary matrix?

7.1 The Scheme

- $\text{Setup}(1^\lambda)$: Pick the right $n = n(\lambda)$ for a security level of λ bits. Generate matrices

$$\mathbf{A}_{1,0}, \mathbf{A}_{1,1}, \dots, \mathbf{A}_{\ell,0}, \mathbf{A}_{\ell,1} \in \mathbb{Z}_q^{n \times m}$$

where ℓ is the length of the identities. The master public key is

$$\text{mpk} = (\mathbf{A}_{i,b})_{i \in [\ell], b \in \{0,1\}}, \mathbf{v}$$

where $\mathbf{v} \in \mathbb{Z}_q^n$ is a random vector, and the master secret key is

$$\text{msk} = (\mathbf{T}_{\mathbf{A}_0}, \mathbf{T}_{\mathbf{A}_1})$$

We will never use the trapdoors for the other matrices (except in the security proof.)

- $\text{KeyGen}(\text{msk}, ID \in \{0, 1\}^\ell)$: Let

$$\mathbf{A}_{ID} := [\mathbf{A}_{1, ID_1} \| \mathbf{A}_{2, ID_2} \| \dots \| \mathbf{A}_{\ell, ID_\ell}]$$

where ID_1, \dots, ID_ℓ are the bits of ID . Generate a short vector $\mathbf{e} \leftarrow \text{DGSamp}(\mathbf{A}_{ID}, \mathbf{T}_{\mathbf{A}_{ID}}, \mathbf{v})$ by running the discrete Gaussian sampling algorithm. Recall that $\mathbf{A}_{ID} \cdot \mathbf{e} = \mathbf{v} \pmod{q}$. Output the secret key $sk_{ID} = \mathbf{e}$.

- $\text{Enc}(\text{mpk}, ID, \mu)$: Run the dual Regev encryption algorithm with $\text{pk} := (\mathbf{A}_{ID}, \mathbf{v})$ and message μ and output the resulting ciphertext.
- $\text{Dec}(sk_{ID}, c)$: Run the dual Regev decryption algorithm with $\text{sk} := sk_{ID} = \mathbf{e}$.

7.2 Proof of (Selective) Security

As before, we will come up with alternate algorithms called Setup^* , KeyGen^* and Enc^* (Dec^* will be the same as Dec) which the challenger will run. We will not be able to use random oracles here.

- $\text{Setup}^*(ID^*, 1^\lambda)$: sample random \mathbf{v}^* . sample ℓ random matrices $\mathbf{B}_1, \dots, \mathbf{B}_\ell$ and set

$$\mathbf{A}_{i, ID_i^*} = \mathbf{B}_i$$

sample ℓ matrices $\mathbf{B}'_1, \dots, \mathbf{B}'_\ell$ together with their trapdoors and set

$$\mathbf{A}_{i, 1-ID_i^*} = \mathbf{B}'_i$$

MPK^* consists of all the $\mathbf{A}_{i,b}$ and \mathbf{v}^* . MSK^* consists of the trapdoors of all $\mathbf{A}_{i, 1-ID_i^*}$.

- $\text{KeyGen}^*(ID)$: We know that $ID \neq ID^*$. Therefore, I know the trapdoor of the matrix

$$\mathbf{A}_{ID} := [\mathbf{A}_{1, ID_1} \| \dots \| \mathbf{A}_{\ell, ID_\ell}]$$

(do you see why?)

- $\text{Enc}^*(\text{MPK}^*, ID^*, \mu)$: return the dual Regev encryption of μ relative to \mathbf{A}_{ID^*} and \mathbf{v}^* . (note that MSK^* does not tell us anything about a trapdoor for \mathbf{A}_{ID^*} .)

One can also prove full security with a more sophisticated proof. In one sentence, the idea is to set up $\mathbf{A}_{i,b}$ so that Algorithm^* can generate secret keys for all the Q secret key queries and yet *not* be able to generate the secret key for ID^* .

7.3 CHKP: Pros and Cons

- PLUS: the scheme is secure without resorting to the random oracle model.
- MINUS: the public parameters are rather large, namely $O(nm \log q \cdot \ell)$ as opposed to GPV where it is $O(nm \log q)$. Consequently, also ciphertexts are large.
- PLUS: While we only showed selective security, one can augment the scheme to be adaptively (fully) secure.
- PLUS: The scheme naturally extends to a hierarchical IBE scheme, described next.

7.4 A Brief Note on Hierarchical IBE

Think of hierarchies in an organization. The CEO (the master key generator) can delegate access to the VP of Engineering who can in turn delegate to programmers and so forth (but not the other way round). In a hierarchical IBE, one can generate SK_{ID} using MSK ; in turn, the owner of SK_{ID} can generate $SK_{ID||ID'}$ etc.

The CHKP scheme has a natural hierarchical structure. Namely, if you know the trapdoor for A_{ID} , you can generate a trapdoor for $A_{ID||ID'} = [A_{ID}||A_{ID'}]$. Constructing a HIBE scheme building off of this idea is left as an exercise.

8 The ABB IBE Scheme(Advanced Material: Not Covered in Class)

The ABB Trick: Punctured Trapdoors.

Given the trapdoor for a matrix A_0 , a matrix R with small entries, and a trapdoor for G , can you generate a trapdoor for

$$[A_0||A_0R + \alpha \cdot G]$$

for an arbitrary integer $\alpha \neq 0 \pmod{q}$?

How about for $\alpha = 0 \pmod{q}$, that is, $[A_0||A_0R]$?

8.1 The Scheme

- $\text{Setup}(1^\lambda)$: Pick the right $n = n(\lambda)$ for a security level of λ bits. Generate matrices

$$A_0, A_1 \in \mathbb{Z}_q^{n \times m}$$

The master public key is

$$\text{mpk} = A_0, A_1, \mathbf{v}$$

where $\mathbf{v} \in \mathbb{Z}_q^n$ is a random vector, and the master secret key is

$$\text{msk} = T_{A_0}$$

We will never use the trapdoor for A_1 .

- $\text{KeyGen}(\text{msk}, ID \in \{0, 1\}^\ell)$: Let h be a collision-resistant hash function that maps identities to \mathbb{Z}_q^* . Define

$$A_{ID} := [A_0||A_1 + h(ID) \cdot G]$$

where G is the gadget matrix. Note that by trapdoor extension, KeyGen knows a trapdoor for A_{ID} for any ID .

Generate a short vector $\mathbf{e} \leftarrow \text{DGSamp}(A_{ID}, T_{A_{ID}}, \mathbf{v})$ by running the discrete Gaussian sampling algorithm. Recall that $A_{ID} \cdot \mathbf{e} = \mathbf{v} \pmod{q}$. Output the secret key $sk_{ID} = \mathbf{e}$.

- $\text{Enc}(\text{mpk}, ID, \mu)$: Run the dual Regev encryption algorithm with $\text{pk} := (A_{ID}, \mathbf{v})$ and message μ and output the resulting ciphertext.
- $\text{Dec}(sk_{ID}, c)$: Run the dual Regev decryption algorithm with $\text{sk} := sk_{ID} = \mathbf{e}$.

8.2 ABB: Proof of Selective Security

As before, we will come up with a bunch of alternate algorithms called Setup^* , KeyGen^* and Enc^* (Dec^* will be the same as Dec) which the challenger will run. We will not be able to use random oracles here either.

- $\text{Setup}^*(ID^*, 1^\lambda)$: sample random \mathbf{v}^* . sample a random matrix \mathbf{A}_0 and a matrix \mathbf{R} with small entries. Set

$$\mathbf{A}_1 := [\mathbf{A}_0 \parallel \mathbf{A}_0 \mathbf{R} - h(ID^*)\mathbf{G}]$$

MPK^* consists of $\mathbf{A}_0, \mathbf{A}_1$ and \mathbf{v}^* . MSK^* consists of \mathbf{R} (and the trapdoor for \mathbf{G} .)

- $\text{KeyGen}^*(ID)$: We know that $ID \neq ID^*$. Therefore, I know the trapdoor of the matrix

$$\mathbf{A}_{ID} := [\mathbf{A}_0 \parallel \mathbf{A}_1 + h(ID)\mathbf{G}] = [\mathbf{A}_0 \parallel \mathbf{A}_0 \mathbf{R} + (h(ID) - h(ID^*))\mathbf{G}]$$

(do you see why?)

- $\text{Enc}^*(MPK^*, ID^*, \mu)$: given a dual Regev encryption of μ relative to \mathbf{A}_0 and \mathbf{v}^* , compute a dual Regev encryption of μ relative to

$$\mathbf{A}_{ID^*} = [\mathbf{A}_0 \parallel (\mathbf{A}_0 \mathbf{R} - h(ID^*)\mathbf{G}) + h(ID^*)\mathbf{G}] = [\mathbf{A}_0 \parallel \mathbf{A}_0 \mathbf{R}]$$

and \mathbf{v}^* . (do you see how to do this?)

8.3 ABB: Pros and Cons

- PLUS: the scheme is secure without resorting to the random oracle model.
- PLUS: the public parameters and ciphertexts are as small as GPV, namely $O(nm \log q)$.
- PLUS: Can be extended to full security.
- PLUS: Extensible to hierarchical IBE. A different ABB paper uses additional techniques to construct a “better” HIBE (where the lattice dimension stays the same regardless of the number of levels of delegation).

9 Application: Chosen Ciphertext Secure Public-key Encryption (Not Covered in Class)

We will now show a very simple construction of a chosen ciphertext secure (CCA2-secure) public-key encryption scheme from IBE. This is due to Canetti, Halevi and Katz [?]. In fact, here we will describe a solution for the weaker notion of CCA1-security.

But first, the definition of CCA1-security. In the CCA1 game, the adversary gets the public-key PK of the encryption scheme, and can ask to get polynomially many ciphertexts decrypted. That is, a challenger will, on input c , run $\text{Dec}(SK, c)$ and return the answer to the adversary. Note that c need not be distributed like an honestly generated ciphertext, and may not even live in the range of the encryption algorithm (i.e., may not be a valid ciphertext). Eventually, the adversary gets an encryption of a random bit b under PK and is asked to guess b . CCA1 security requires that no PPT adversary can guess b with probability better than $1/2 + \text{negl}(\lambda)$.

Here is the construction.

- $\text{KeyGen}(1^\lambda)$: run $\text{IBE.Setup}(1^\lambda)$ to get an MPK_{IBE} and an MSK_{IBE} . The public key PK of the CCA scheme is MPK_{IBE} and the secret key SK is MSK_{IBE} .
- $\text{Enc}(PK, \mu)$: pick a random string ID . Run $\text{IBE.Enc}(PK = MPK_{\text{IBE}}, ID, \mu)$ and output ID together with the resulting ciphertext.
- $\text{Dec}(SK, (ID, c))$: use $SK = MSK_{\text{IBE}}$ to create SK_{ID} and run the IBE decryption algorithm $\mu = \text{IBE.Dec}(SK_{ID}, c)$.

The CCA security proof is super simple. The intuition?

- the decryption algorithm only uses SK_{ID} (and not the MSK per se) and
- the identity in the challenge ciphertext is random and hence different w.h.p. from the (adversarially chosen) identities in all the decryption queries.

Put together, IBE security should say that breaking the security of the challenge ciphertext is hard.

10 Registration-based Encryption (Advanced Material: Not Covered in Class)

We will say just a few words about RBE here. Recall from the beginning of the lecture that a major disadvantage of IBE is the power of the master key authority to decrypt all ciphertexts.

A completely orthogonal approach which does not have this problem starts from the following straw-man scheme: the master public key, curated by the authority, is the concatenation of all the users' public keys... Of course, this leads us back to exactly the PKI problem we wanted to solve. However, it is possible that the authority can publish a *short digest* of the concatenation of all public keys, which is nevertheless good enough for encryption (although it should not be clear exactly how yet!)

It turns out that this idea can be brought to fruition using the methodology of deferred encryption due to Garg et al. We refer the reader to the papers [?, ?]. The construction proceeds in a completely different way from everything we saw today, and is quite inefficient. An open problem is to come up with an RBE that is as efficient as (or more efficient than!) the IBE schemes we saw here.

11 Onto Attribute-based Encryption: The Key Lattice Equation

Let us abstract out the mathematics behind GSW into a *key lattice equation* which will guide us through constructing the rest of the primitives in this lecture, in particular an attribute-based encryption (ABE) scheme.

Recall the approximate eigenvector relation:

$$\mathbf{s}^T \mathbf{A}_i \approx \mu_i \mathbf{s}^T \mathbf{G}$$

and rewrite it as

$$\mathbf{s}^T (\mathbf{A}_i - \mu_i \mathbf{G}) \approx \mathbf{0} \tag{1}$$

Let \mathbf{A}_f be the homomorphically evaluated ciphertext for a function f . We know that

$$\mathbf{s}^T \mathbf{A}_f \approx f(\boldsymbol{\mu}) \mathbf{s}^T \mathbf{G}$$

or

$$\mathbf{s}^T(\mathbf{A}_f - f(\boldsymbol{\mu})\mathbf{G}) \approx \mathbf{0} \quad (2)$$

We will generalize this to arbitrary matrices $\mathbf{A}_1, \dots, \mathbf{A}_\ell$ – not necessarily ones that share the same eigenvector.

First, we know that \mathbf{A}_f is a function of $\mathbf{A}_1, \dots, \mathbf{A}_\ell$ and f (*but not* μ_1, \dots, μ_ℓ). Henceforth, when we say \mathbf{A}_f , we will mean a matrix obtained by the GSW homomorphic evaluation procedure. (That is, homomorphic addition of two matrices is matrix addition; homomorphic multiplication is matrix multiplication after bit-decomposing the second matrix).

Second, and very crucially, we can show that for any sequence of matrices $\mathbf{A}_1, \dots, \mathbf{A}_\ell$,

$$[\mathbf{A}_1 - \mu_1\mathbf{G} \parallel \dots \parallel \mathbf{A}_\ell - \mu_\ell\mathbf{G}] \mathbf{H}_{f,\boldsymbol{\mu}} = \mathbf{A}_f - f(\boldsymbol{\mu})\mathbf{G}$$

where $\mathbf{H}_{f,\boldsymbol{\mu}}$ is a matrix with small coefficients. We call this the key lattice equation.

To see this for addition, notice that

$$[\mathbf{A}_1 - \mu_1\mathbf{G} \parallel \mathbf{A}_2 - \mu_2\mathbf{G}] \underbrace{\begin{bmatrix} \mathbf{I} \\ \mathbf{I} \end{bmatrix}}_{\mathbf{H}_{+, \mu_1, \mu_2}} = \mathbf{A}_1 + \mathbf{A}_2 - (\mu_1 + \mu_2)\mathbf{G} = \mathbf{A}_+ - (\mu_1 + \mu_2)\mathbf{G}$$

and for multiplication,

$$[\mathbf{A}_1 - \mu_1\mathbf{G} \parallel \mathbf{A}_2 - \mu_2\mathbf{G}] \underbrace{\begin{bmatrix} \mathbf{G}^-(\mathbf{A}_2) \\ \mu_1\mathbf{I} \end{bmatrix}}_{\mathbf{H}_{\times, \mu_1, \mu_2}} = \mathbf{A}_1\mathbf{G}^-(\mathbf{A}_2) - \mu_1\mu_2\mathbf{G} = \mathbf{A}_\times - \mu_1\mu_2\mathbf{G}$$

By composition, we get that

$$[\mathbf{A}_1 - \mu_1\mathbf{G} \parallel \mathbf{A}_2 - \mu_2\mathbf{G} \parallel \dots \parallel \mathbf{A}_\ell - \mu_\ell\mathbf{G}] \mathbf{H}_{f,\boldsymbol{\mu}} = \mathbf{A}_f - f(\boldsymbol{\mu})\mathbf{G}$$

where $\mathbf{H}_{f,\boldsymbol{\mu}}$ is a matrix with small entries (roughly proportional to $m^{O(d)}$ where d is the circuit depth of f).

An Advanced Note: Given arbitrary matrices \mathbf{A}_i and \mathbf{A}_f , there exists such a small matrix \mathbf{H} ; but if \mathbf{A}_f is arbitrary, it is hard to find.

Let's re-derive FHE from the key equation:

- The ciphertexts are the matrices \mathbf{A}_i and we picked them such that

$$\mathbf{s}^T\mathbf{A} \approx \mu\mathbf{s}^T\mathbf{G}$$

- Homomorphic evaluation is computing \mathbf{A}_f starting from $\mathbf{A}_1, \dots, \mathbf{A}_\ell$.
- Correctness of homomorphic eval follows from the key equation: We know that

$$\mathbf{s}^T[\mathbf{A}_1 - \mu_1\mathbf{G} \parallel \dots \parallel \mathbf{A}_\ell - \mu_\ell\mathbf{G}] \approx \mathbf{0}$$

by the equation above that characterizes ciphertexts. Therefore, by the key equation,

$$\mathbf{s}^T[\mathbf{A}_f - f(\boldsymbol{\mu})\mathbf{G}] = \mathbf{s}^T[\mathbf{A}_1 - \mu_1\mathbf{G} \parallel \dots \parallel \mathbf{A}_\ell - \mu_\ell\mathbf{G}] \mathbf{H}_{f,\boldsymbol{\mu}} \approx \mathbf{0}$$

as well meaning that \mathbf{A}_f is an encryption of $f(\boldsymbol{\mu})$. Note that no one needs to know or compute the matrix \mathbf{H} ; it only appears in the analysis.

12 Attribute-based Encryption

Attribute-based encryption (ABE) generalizes IBE in the following way.

- Setup produces MPK, MSK .
- Enc uses MPK to encrypt a message m relative to attributes $(\mu_1, \dots, \mu_\ell) \in \{0, 1\}^\ell$.

(In an IBE scheme, $\mu = ID$.)

- KeyGen uses MSK to generate a secret key SK_f for a given Boolean function $f : \{0, 1\}^\ell \rightarrow \{0, 1\}$.

(IBE is the same as ABE where f is restricted to be the point (delta) function $f_{ID'}(ID) = 1$ iff $ID = ID'$.)

- Dec gets μ (attributes are in the clear) and uses SK_f to decrypt a ciphertext C if $f(\mu) = 1$ (true). If $f(\mu) = 0$, Dec simply outputs \perp .

Here is an ABE scheme (called the BGG+ scheme) using the key equation. It's best to view this as a generalization of the Agrawal-Boneh-Boyen IBE scheme.

- KeyGen outputs matrices A, A_1, \dots, A_ℓ and a vector \mathbf{v} and these form the MPK . The MSK is the trapdoor for A .
- Enc computes

$$\mathbf{s}^T [A \| A_1 - \mu_1 G \| \dots \| A_\ell - \mu_\ell G]$$

(plus error, of course, and we will consider that understood.) Finally, the message is encrypted as $\mathbf{s}^T \mathbf{v} + e + m \lfloor q/2 \rfloor$.

- Let's see how Dec might work. You (and in fact anyone) can compute

$$\mathbf{s}^T [A \| A_1 - \mu_1 G \| \dots \| A_\ell - \mu_\ell G] \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{H}_{f, \mu} \end{bmatrix} = \mathbf{s}^T [A \| A_f - f(\mu)G]$$

using the key equation.

If you had a short \mathbf{r} that maps $[A \| A_f - G]$ to \mathbf{v} , that is

$$[A \| A_f - G] \mathbf{r} = \mathbf{v}$$

you can decrypt and find m . (Can you fill in the blanks?)

Two notes:

- The security definition mirrors IBE exactly, and the security proof of this scheme mirrors that of the ABB IBE scheme that we did in the last lecture. I will leave it to you as an exercise. The reference is the work of Boneh et al., Eurocrypt 2014.
- One might wonder if the attributes μ need to be revealed. The answer is "NO", in fact one can construct an attribute-hiding ABE scheme (also called a predicate encryption scheme). There are two flavors of security of such a scheme, the weaker one can be realized using LWE (Gorbunov, Vaikuntanathan and Wee, Crypto 2015) and the stronger one implies indistinguishability obfuscation, a very powerful cryptographic primitive which we don't know how to construct from LWE yet. More in the next lecture.