Princeton University
Scribe: Boriana Gjura

Notes for Lecture 16

1 Lattices (continued)

1.1 Last time.

We defined lattices as a set of integer linear combinations of a basis.

Definition 1 *B is a basis for the lattice $\mathfrak{L}$ if the columns of B are linearly independent and*

$$\mathfrak{L} = \{\, B \cdot x \mid x \in \mathbb{Z}^n \,\}.$$

We say that $\mathfrak{L} := \mathfrak{L}(B)$ is the lattice spanned by B.

We discussed the following two computational problems.

(SVP) Shortest Vector Problem. Given a basis $B \in \mathbb{Z}^{n \times n}$, find the shortest (nontrivial) vector in $\mathfrak{L}(B) \backslash \{0\}$.

(CVP) Closest Vector Problem. Given a basis $B \in \mathbb{Z}^{n \times n}$, and a target vector $t \in \mathbb{Z}^n$, where t is not necessarily in $\mathfrak{L}(B)$, find the closest point to t in $\mathfrak{L}(B)$.

We also defined gap versions of the above problems. We will continue by analyzing some special classes of lattices, discussing the Learning with Errors assumption and looking at some applications.

1.2 Some special classes of lattices

From now on, we will only consider lattices in $\mathbb{Z}^n$. This is useful, because finite precision will not be an issue. Moreover, any basis $B \in \mathbb{Z}^{n \times m}$ defines a lattice, even it its columns are not linearly independent, which is not the case in $\mathbb{R}^n$.

Let $q \geq 2$ be an integer, and let $m, n \in \mathbb{Z}$, with $m > n$. Let $A \in \mathbb{Z}_q^{n \times m}$ be a wide matrix. We will consider two special classes of lattices.

1.2.1 $\Lambda_q^{\perp}(A) = \{\, x \in \mathbb{Z}^m \mid Ax = 0 \pmod{q} \,\}$

This is indeed a lattice, since adding any two vectors in the set yields another element in the set.

The null space of A is an $(m-n)$ dimensional object. Let $C \in \mathbb{Z}^{m\times(m-n)}$ be such that $AC = 0 \pmod{q}$. Since $\Lambda_q^{\perp}(A)$ is m-dimensional, C alone will not suffice as a spanning basis. We fix this by adding vectors, to get

$$\Lambda_q^{\perp}(A) = \mathfrak{L}(C \mid qI_m),$$

where I is the identity matrix.

1.2.2 $\Gamma_q(A) = \{\, x \in \mathbb{Z}^n \mid \exists\, r : x = A^T r \pmod{q} \,\}$

We can easily check that this a lattice. If $x_1, x_2 \in \Gamma_q(A)$, then there exist r_1, r_2 such that $x_1 = A^T r_1 \pmod{q}$ and $x_2 = A^T r_2 \pmod{q}$. Then $x_1 + x_2 = A^t(r_1 + r_2) \pmod{q}$, hence $x_1 + x_2 \in \Gamma_q(A)$. Then as before, we get

$$\Gamma_q^{\perp}(A) = \mathfrak{L}(A \mid qI_m)$$

We will analyze hard problems on the first lattice.

1.3 Some hard lattice problems

(SIS) Short Integer Solution Problem

Let q, m, β be functions of n, where n will play the role of our security parameter.

The problem is as follows.

(SIS) Given $A \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$, find $x \in \mathbb{Z}^m$ such that:

$$(i)\ \|x\|_2 \leq \beta \qquad (ii)\ Ax = 0 \pmod{q}$$[1]

Fun Fact. With high probability over the choice of A, if m is large enough over n ($m \geq \Omega(n \log q)$ suffices), there exists an $x \in \Gamma_q^{\perp}(A) \cap \{0,1\}^m$.

We will not prove this fact, which states that the short vectors are 0,1 vectors with high probability. This implies that, for $\gamma = \frac{\beta}{\sqrt{m}}$, we have

$$SIS_{q,m,b} \approx SVP_{\gamma}.$$

[1] That is, $x \in \Lambda_q^{\perp}(A)$

Assumption. $SIS_{q,m,\beta}$ is hard, i.e. every probabilistic polyonimial time algorithm only has a negligible probability of giving a SIS solution.

This assumption allows us to construct hash functions.

Let $A \in \mathbb{Z}_q^{n\times m}$, and define $f_A : \{0,1\}^m \to \mathbb{Z}_q^n$ as $f_A(x) = Ax \pmod q$, $\forall x \in \{0,1\}^m$.

If SIS is hard, f_A is collision resistant.

The idea is to turn a collision into a SIS solution. Assume we can find $x, y \in \{0,1\}^m$ such that $Ax = Ay \pmod q$, then $A(x-y) = 0$ with $(x-y) \in \{-1,0,1\}^m$ short.

Based on a previous homework, if m is sufficiently large the function is compressing, so in addition to being collision resistant, it is also a one-way function. Moreover, f_A is fast to compute.

1.3.1 (LWE) Learning with errors

Discrete Gaussian

We want to get a probability distribution over $\mathbb{Z}$, which is proportional to the probability distribution of the continuous Gaussian. We start with

$$D'_{\sigma,c}(x) = Pr[x : x \leftarrow D_{\sigma,c}] = e^{-\pi|x-c|^2/\sigma^2},$$

which is slightly different from the usual definition of a Gaussian [2], but this will simplify some of the calculations. To actually get a probability distribution, we need to normalize:

$$D_{\sigma,c}(x) = \frac{e^{-\pi|x-c|^2/\sigma^2}}{\sum_x e^{-\pi|x-c|^2/\sigma^2}}$$

Note. We will take for granted that this distribution can be sampled efficiently.

Learning with Errors Problem

Let $A \in \mathbb{Z}_q^{n\times m}$ be a wide matrix, i.e. $m > n$. Given $s^T A$, it is easy to find s using linear algebra. However, adding a noise makes finding s had.

LWE$_{q,m,s}$: Given random $A \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$, let $U = s^T A + e^T \pmod q$, where $s \xleftarrow{\$} \mathbb{Z}_q^n$, $e \xleftarrow{\$} D_{\sigma,0}^m$.

We can define two versions of the problem.

Search. Find s.

[2] The usual definition is $e^{-|x-c|^2/2\sigma^2}$

Decisional. Distinguish (A, U) from (A, u), where $u \xleftarrow{\$} \mathbb{Z}_q^m$.

We make the computational assumption that these problems are hard.

Note. Solving LWE is similar to solving SVP for $\Lambda_q^{\perp}(A)$. The idea is that if we can recover s from $s^T A + e$, this corresponds to finding a closest vector to U in the lattice.

1.4 Public Key Encryption using LWE [Regev '05]

We describe a public key encryption protocol using the Learning With Errors assumption, first introduced by Regev.

Gen():

Randomly choose A, s, e: $\quad A \xleftarrow{\$} \mathbb{Z}_q^{n\times m},\quad s \xleftarrow{\$} \mathbb{Z}_q^n,\quad e \xleftarrow{\$} \mathbb{D}_{\sigma,0}^m$ [3]
Set the secret and public key: $\quad sk = (A, s),\quad pk = (A, s^T A + e)$

Enc(pk = (A, u), m):
Let $m \in \{0,1\}$ (encrypt bit by bit). Choose $x \xleftarrow{\$} \{0,1\}^n$, encrypt as
$(Ax \pmod q), ux + \lceil \frac{q}{2} \rfloor m \pmod q)$.

Dec(sk = (A, s), (y, z)):
Compute $z - sy \pmod q$ to get

$$\begin{aligned} z - sy \pmod q &= (s^T A + e)x + \left\lceil \frac{q}{2} \right\rfloor m - s^T Ax \pmod q \\ &= ex + \left\lceil \frac{q}{2} \right\rfloor m \pmod q \end{aligned}$$

e is chosen by Discrete Gaussian of width σ - entries are roughly of order σ. Then $\|ex\|_2 \le \sqrt{m}\sigma n^{o(1)} \pmod q$ with high probability. If $m = 0$, $\lceil \frac{q}{2} \rfloor m$ is short, and if $m = 1$, then $ex + \lceil \frac{q}{2} \rfloor m$ is close to $\lceil \frac{q}{2} \rfloor m$. So decrypt bit as 0 if closer to 0 then $\lceil \frac{q}{2} \rfloor$, 1 otherwise.

Security Theorem. If LWE holds, then (Gen, Enc, Dec) is CPA-secure.

Proof Idea.

Assume toward contradiction that we have an adversary E (eavesdropper), and define the CPA-experiment as per usual. The challenger uses $Gen()$ to generate a (sk, pk) pair using the above procedure, then outputs the public key. The eavesdropper then sends m_0, m_1 to the challenger, who later outputs $c = Enc(pk, m_b)$. E then ouptputs

[3] σ is polynomial on m, the problem can be easy if it is a constant.

b', and wins if $b = b'$ with probability greater than half. Define the following hybrids:

Hybrid 0. CPA-Exp for random bit b. (E outputs b with probability $1/2 + \varepsilon$).

Hybrid 1. Same experiment, except that now the public key $pk = (A, u)$ is generated at random, that is, $A \xleftarrow{\$} \mathbb{Z}^{n\times m}$, $u \xleftarrow{\$} \mathbb{Z}^m$.

Assuming decisional LWE holds, in the second hybrid, E must output b with probability $1/2 + \varepsilon - negl$. We show that this yields a contradiction.

E sees $\left(Ax, ux + \lceil \frac{q}{2} \rfloor b\right)$. If m is big enough relative to n, say $m \geq \Omega(n \log q)$, an entropy argument gives

$$\left(\frac{A}{u}, \frac{A}{u}x\right) \simeq_s \big(\text{random matrix, random vector}\big)$$

The view of the eavesdropper then is statistically close to (random, random + $\lceil \frac{q}{2} \rfloor$), which is statistically close to (random, random). Hence the view of E is independent of b, and it outputs b with probability at most $1/2 + negl$. Contradiction.

1.5 Lattice Trapdoors

Choose random $x \xleftarrow{\$} \{0,1\}^m$. Choose $A \leftarrow \mathbb{Z}_q^{n\times m}$ such that $Ax = 0 \pmod q$.

We claim that A is statistically close to a random matrix, if x is hidden.

Knowing x allows us to solve decisional LWE for A. Indeed, to distinguish between $(A, s^T A + E)$ from (A, u), where u is chosen at random, we can compute $ux \pmod q$, which should be random, and $s^T Ax + ex = ex \pmod q$, which is small. We call x a trapdoor, since without knowing x decisional LWE remains hard. We have the following result.

Theorem 2 (Ajtai '99) *We can sample $T \leftarrow \mathbb{Z}^{m\times m}, A \leftarrow \mathbb{Z}^{n\times m}$ such that:*
(i) T is short
(ii) T is full rank over $\mathbb{Z}$
(iii) $AT \pmod q = 0$
(iv) A is statistically close to random

Knowing T yields a solution for the search LWE. Indeed, given $(A, T, s^T A + e^T)$, we have $(s^T A + e^T)T = e^T T \pmod q$. Since $e^T T$ is short, this holds over $\mathbb{Z}$. We can use the fact that T is full rank over $\mathbb{Z}$ to recover $e = (e^T T)(T^{-1})$. So we have $e^T = ((e^T A + e^T)T \pmod q)T^{-1}$, and we have reduced the problem to the no-error case, which we can solve using linear algebra.