Homomorphic Encryption and Lattices, Spring 2011	Instructor: Shai Halevi

Regev's Main Average-Case to Worst-Case Lemma

May 5, 2011 — Scribe: Ron Rothblum

We show the main part of Regev's [Reg09] proof that (under certain conditions) it is possible to relate the *average-case* hardness of the learning with errors problem (LWE) to the *worst-case* hardness of bounded distance decoding in a given lattice (BDD).

Preliminaries. We use the following parameters:
n - security parameter.
α - noise parameter ($= \frac{1}{\text{poly}(n)}$).
q - modulus ($\gg \frac{1}{\alpha}$, sometimes even $q = \exp(n)$).

Recall that for a continuous distribution D_r (with standard deviation r), $\tilde{D}_{L,r}$ denotes a discrete distribution over a lattice (or coset of a lattice) L such that every vector $\vec{z} \in L$ has probability mass proportional to $D_r(\vec{z})$.

1 The Main Lemma

In addition to an oracle that solves LWE, the reduction from BDD to LWE also needs access to an oracle that samples short vectors in Λ^* (Regev [Reg09] and Peikert [Pei09] show how to construct such an oracle in specific settings). Additionally it relies on the following properties of the LWE error distribution:

- The LWE error distribution is a projection of a spherical distribution $D_{\alpha q}$ onto its first coordinate.
- The distribution $D_{\alpha q}$ is smooth in the following sense: If Λ is some lattice (or coset of a lattice) with $\lambda_n(\Lambda) \ll \alpha q$ then if we choose $\vec{x} \leftarrow \tilde{D}_{\Lambda,r}$ and $\vec{y} \leftarrow D_s$ such that $r^2 + s^2 = (\alpha q)^2$ then the induced distribution on $\vec{x} + \vec{y}$ is close to $D_{\alpha q}$.

For example the n-dimensional discrete Gaussian has these properties (where $\Lambda_n \ll \alpha q$ means $\Lambda_n \cdot \omega(\sqrt{\log(n)}) < \alpha q$). In this case the LWE error distribution is just the one-dimensional Gaussian.

Lemma 1 ([Reg09])**.** *There is an efficient algorithm that takes as input a basis B of an n-dimensional lattice $\Lambda = \Lambda(B)$, another parameter $r \gg \frac{q}{\lambda_1(\Lambda)}$ and a point $\vec{x} \in \mathbb{R}^n$ such that* $\text{dist}(x, \Lambda) < \frac{\alpha q}{\sqrt{2}r}$ *and has access to two oracles:*

- *A "global" solver for $\mathsf{LWE}[n, \alpha, q]$ ("global" in the sense that it is unrelated to the input lattice).*
- *A "lattice specific" sampler from $D_{\Lambda^*,r}$.*

The algorithm finds (with overwhelming probability) the (unique) point $\vec{v} \in \Lambda$ closest to $\vec{x}$.

2 Proof Sketch of Lemma 1

Let $\vec{v} \in \Lambda$ be the closest point to $\vec{x}$ in Λ and let $\vec{t} \in \mathbb{Z}^n$ be the coefficients of $\vec{v}$ when expressed in basis B (i.e., $\vec{v} = B\vec{t}$) and denote $\vec{s} \stackrel{\text{def}}{=} \vec{t} \bmod q$. We show a procedure that uses the sampler for $\tilde{D}_{\Lambda^*,r}$ to generate instances of the distribution $\mathsf{LWE}_{\vec{s}}$. Then, we use the LWE solver to find $\vec{s}$. (Note that $\vec{s}$ was not chosen uniformly at random in this case, but we previously showed a random self reduction for LWE from a random $\vec{s}$ to any specific $\vec{s}$.) Later we show how from $\vec{s}$ one can find $\vec{t}$ thereby solving BDD.

LWE-Generate(B, $\vec{x}$) (With access to $\tilde{D}_{\Lambda^*,r}$)

1. Draw a sample $\vec{y} \leftarrow \tilde{D}_{\Lambda^*,r}$. Let $\vec{a}$ be the coefficients of $\vec{y}$ in basis B^* (i.e. $\vec{a} = B^T\vec{y}$).
2. Draw an error term $e \leftarrow \Phi_{\frac{\alpha}{2\sqrt{\pi}}}$.
3. Output $(\vec{a}, b = \langle \vec{x}, \vec{y} \rangle + e \bmod q)$.

Claim 1. *The output of LWE-Generate is statistically close to* $\mathsf{LWE}_{\vec{s}}$ *except that the error parameter is* $\beta \leq \alpha$.

Proof. Need to show:

(A.) $\vec{a}$ is close to uniform in $\mathbb{Z}_q^n$.

(B.) Once $\vec{a}$ is fixed, $\vec{b} = \langle \vec{s}, \vec{a} \rangle + \Phi_{\beta q}$ $(\beta \leq \alpha)$.

(A.) Consider the lattice $q \cdot \Lambda^*$ and all its q^n cosets

$$\vec{a}\text{-coset} = \{B^*\vec{a} + q\Lambda^*\} = \{B^*\vec{z} : \ \vec{z} = \vec{a} \bmod q\}$$

The vector $\vec{a}$ output by the procedure is exactly the coset of $\vec{y}$. Due to our choice of parameters, all cosets are (almost) equally likely. Indeed, since $r \gg \frac{q}{\lambda_1(\Lambda)} = \frac{q\lambda_n(\Lambda^*)}{n}$ then $\tilde{D}_{\Lambda^*,r}$ is nearly uniform among the cosets.

(B.) Conditioned on any fixed $\vec{a} \in \mathbb{Z}_q^n$, the vector $\vec{y}$ is chosen from the discrete distribution $D_{q\Lambda^*+\vec{a},r}$ on the $\vec{a}$-coset. Denoting $\vec{w} \stackrel{\text{def}}{=} \vec{x} - \vec{v}$ we have

$$\begin{aligned} \langle \vec{x}, \vec{y} \rangle &= \langle \vec{v} + \vec{w}, \vec{y} \rangle \\ &= \langle \vec{v}, \vec{y} \rangle + \langle \vec{w}, \vec{y} \rangle \\ &= \langle B\vec{t}, \vec{y} \rangle + \langle \vec{w}, \vec{y} \rangle \\ &= \langle \vec{t}, B^T\vec{y} \rangle + \langle \vec{w}, \vec{y} \rangle \\ &= \langle \vec{s}, \vec{a} \rangle + \langle \vec{w}, \vec{y} \rangle \bmod q \end{aligned}$$

hence $b = \langle \vec{s}, \vec{a} \rangle + \langle \vec{w}, \vec{y} \rangle + e \bmod q$. Notice that $\vec{s}$, $\vec{a}$ and $\vec{w}$ are fixed and the random part is just $\vec{y}$ and e.

Recall that $\Phi_{\frac{\alpha}{2\sqrt{\pi}}}$ is the projection of $D_{\frac{\alpha}{2\sqrt{\pi}}}$ onto the first coordinate, namely $\langle \vec{e_1}, D_{\frac{\alpha}{2\sqrt{\pi}}} \rangle$ and since D is spherical then this is also the same as $\langle \vec{u}, D_{\frac{\alpha}{2\sqrt{\pi}}} \rangle$ for any other unit vector $\vec{u}$. In particular, $\Phi_{\frac{\alpha}{2\sqrt{\pi}}} \equiv \langle \vec{w}, D_{\frac{\alpha}{2\sqrt{\pi}}} \rangle \frac{1}{||\vec{w}||} \equiv \langle \vec{w}, D_{\frac{\alpha}{2\sqrt{\pi}||\vec{w}||}} \rangle$.

Hence $\langle \vec{w}, \vec{y} \rangle + e \equiv \langle \vec{w}, \vec{y} \rangle + \langle \vec{w}, \vec{z} \rangle = \langle \vec{w}, \vec{y} + \vec{z} \rangle$ where $y \in_R \tilde{D}_{q\Lambda^*+\vec{a},r}$ and $z \in_R D_s$ where $s = \frac{\alpha}{2\sqrt{\pi}||\vec{w}||}$. Now $||\vec{w}||$ is "short" so s is "large". The parameters r, s are chosen large enough so that $\tilde{D}_{q\Lambda^*+\vec{a},r}$ is close to the continuous D_t where $t = \sqrt{r^2+s^2}$. Therefore $\langle \vec{w}, \vec{y} \rangle + e \approx \langle \vec{w}, D_t \rangle = \Phi_{||\vec{w}|| \cdot t}$ and the parameters are such that $||\vec{w}|| \cdot t \leq \alpha q$. □

To solve BDD for $\vec{x}$ we can apply the LWE-solver with samples from LWE-Generate to find the vector $\vec{s}$. However, to solve BDD we need to find $\vec{t}$ (recall $\vec{s} = \vec{t} \bmod q$). To do this, first observe that $\vec{v} = B\vec{t} = B\vec{s} + B(q\vec{z})$ for some $\vec{z} \in \mathbb{Z}^n$ and consider $\vec{x'} = \frac{\vec{x}-B\vec{s}}{q} = \frac{\vec{x}-\vec{v}}{q} + B\vec{z}$. Notice that by this calculation, the vector $\vec{x'}$ is at distance $\frac{||\vec{w}||}{q}$ (where $\vec{w} = \vec{x} - \vec{v}$) from the lattice (specifically the point $B\vec{z}$). If we could find the closest lattice point to $\vec{x'}$ we would have $\vec{z}$ and therefore also $\vec{v}$. To do this just repeat the above argument again and again and at each iteration the distance from the lattice is reduced by a factor of q. After n such iterations we can solve the problem by using, e.g., Babai's nearest plane algorithm.

References

[Pei09] Chris Peikert. Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In *41st Annual ACM Symposium on Theory of Computing, STOC 2009*, pages 333–342. ACM, 2009.

[Reg09] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *JACM*, 56(6), 2009.