10/28/24

Lecture 15: Cryptographic hardness

$(q, m, n, \mathcal{D}, \mathcal{E})$-LWE: Given secret vector s sampled from a distribution $\mathcal{D}$ over $\mathbb{R}^n$ (usually over $(\mathbb{Z}/q\mathbb{Z})^n$), samples $(x_1, y_1), \dots, (x_m, y_m)$ generated iid via

$$x \sim \mathrm{unif}((\mathbb{Z}/q\mathbb{Z})^n)$$
$$y = (\langle x, s \rangle + e) \bmod q, \quad e \sim \mathcal{E}$$

Standard choice of noise distribution: $\mathcal{E} = D_{\mathbb{Z},\sigma}$ where

Discrete Gaussian $D_{\mathbb{Z},\sigma}$: distribution over $\mathbb{Z}$ w/ density

this normalization just makes passing to Fourier more convenient

$$D_{\mathbb{Z},\sigma}(x) = \frac{1}{C}\exp(-\pi x^2/\sigma^2) \quad \text{for} \quad C = \sum_{y \in \mathbb{Z}} \exp(-\pi y^2/\sigma^2)$$

> **Decisional LWE**: distinguish whether samples came from $(q, m, n, \mathcal{D}, D_{\mathbb{Z},\sigma})$-LWE or from $\mathcal{D} \times \mathrm{Unif}(\mathbb{Z}/q\mathbb{Z})$.

$(\beta, \gamma, m, n, \mathcal{D})$-CLWE: Given secret vector w sampled from a distribution $\mathcal{D}$ over $\mathbb{S}^{n-1}$, samples $(x_1, y_1), \dots, (x_m, y_m)$ given by

$$x \sim N(0, \mathrm{Id}_n)$$
$$y = (\gamma \langle x, w \rangle + e) \bmod 1, \quad e \sim N(0, \beta^2)$$

> **Decisional CLWE**: distinguish whether samples came from $(\beta, \gamma, m, n, \mathcal{D})$-CLWE or from $\mathcal{D} \times N(0,1)$

Homogeneous CLWE $\equiv$ (infinite parallel pancakes)

Define the event $\mathcal{E}: \gamma z + N(0,\beta^2) \in \mathbb{Z}$ for $z = \langle w, x\rangle \sim N(0,1)$

$$\Pr[z \mid \mathcal{E}] \propto \Pr[x] \cdot \sum_{k\in\mathbb{Z}} \exp\left(-\frac{(k-\gamma x)^2}{2\beta^2}\right)$$

$$\propto \sum_{k\in\mathbb{Z}} \exp\left(-\frac{x^2}{2}\left(1+\frac{\gamma^2}{\beta^2}\right) + \frac{\gamma x k}{\beta^2}\right) \exp\left(-\frac{k^2}{2\beta^2}\right)$$

$$= \sum_{k\in\mathbb{Z}} \exp\left(-\frac{1+\gamma^2/\beta^2}{2}\left(x - \frac{\gamma k}{\beta^2+\gamma^2}\right)^2\right) \exp\left(-\frac{k^2}{2\beta^2}\right)$$

$$\propto \sum_{k\in\mathbb{Z}} \exp\left(-\frac{k^2}{2\beta^2}\right) \cdot N\left(\frac{\gamma k}{\beta^2+\gamma^2}, \frac{\beta^2}{\beta^2+\gamma^2}\right).$$

So in direction of w, x distributed as mixture of Gaussians of width $\frac{\beta}{\sqrt{\beta^2+\gamma^2}}$ centered at multiples of $\frac{\gamma}{\beta^2+\gamma^2}$ with exponentially decaying mixing weights

Thm [Gupte-Vafa-Vaikuntanathan '22]:
Let $\mathcal{D}$ be any distribution over vectors in $\mathbb{Z}^n$ w/ norm r.
Let $\mathcal{E}$ be the discrete Gaussian distribution $D_{\mathbb{Z},\sigma}$.
Let $\gamma = \tilde{O}(r)$ and $\beta = O(\frac{\sigma}{q})$. Provided $\sigma >> r$, there is a poly-time reduction from decisional $(q, m, n, \mathcal{D}, \mathcal{E})$-LWE to decisional $(\beta, \gamma, m, n, \mathcal{D}')$-CLWE for $\mathcal{D}'$ the distribution given by rescaling $\mathcal{D}$ to unit norm vectors.

Steps of proof:

0. preliminaries about Gaussians on lattices
1. discrete noise $\to$ continuous noise
2. discrete x's $\to$ x's $\sim \mathrm{Unif}(\mathbb{R}/q\mathbb{Z})$
3. $\mathrm{Unif}(\mathbb{R}/q\mathbb{Z}) \to N(0, \mathrm{Id})$

0 Preliminaries:

Lattice of rank n: set $\boxed{\Lambda} \subseteq \mathbb{Z}^n$ of all integer linear combinations of n linearly independent vectors $\mathcal{B} = \{b_1, \dots, b_n\}$ ("basis")

Dual lattice $\boxed{\Lambda^*}$: $y \in \mathbb{R}^n$ s.t. $\langle x, y\rangle \in \mathbb{Z}$ $\forall$ $x \in \Lambda$

If Λ has basis $\mathcal{B}$, Λ^* has basis $(\mathcal{B}^T)^{-1}$

Denote by $\rho_{s,c}(x) = \exp(-\pi \|x-c\|^2/s^2)$ (Gaussian density)

so that $\mathcal{N}(c, \frac{s^2}{2\pi}, x) = \rho_{s,c}(x)/s^n \triangleq \boxed{D_{s,c}(x)}$

when $c = 0$, denote by $\rho_s(x)$.

Given $T \subseteq \mathbb{R}^n$, denote

$$\rho_s(T) \triangleq \sum_{y \in T} \rho_s(y)$$

and define general discrete Gaussian over $\Lambda + c$ w/ spread s:

$$\boxed{D_{\Lambda+c,s}(x)} = \frac{\rho_s(x)}{\rho_s(\Lambda+c)}$$

where $\Lambda + c \triangleq \{c+y : y \in \Lambda\}$.

Fundamental parallelepiped: Given basis $B = \{b_1, \dots, b_n\} \subset \mathbb{Z}^n$

$$\boxed{P(B)} \triangleq \{\sum_i x_i b_i : 0 \le x_i < 1 \text{ for } 1 \le i \le n\}.$$

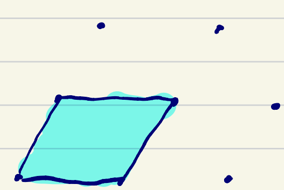

Given $x \in \mathbb{R}^n$, $x \bmod P(B)$ is the unique point $z \in P(B)$ s.t. $x - z \in \Lambda$.

$\boxed{\det(\Lambda)} \triangleq \det(B) = \text{volume}(P(B))$ (note: $\det(\Lambda^*) = \frac{1}{\det(\Lambda)}$)

Smoothing parameter: Given lattice Λ,

$$\boxed{\eta_\varepsilon(\Lambda)} \triangleq \inf\{s : \rho_{1/s}(\Lambda^* \setminus \{0\}) \le \varepsilon\}$$

Intuition (made formal below): amount of Gaussian noise g needed for $g \bmod P(B)$ to be approximately distributed as uniform over $P(B)$.

(Most important lemmas highlighted in green)

Lemma 1: For any $\varepsilon, s > 0$, $c \in \mathbb{R}^n$, rank-n lattice Λ with basis $\mathcal{B}$,

$$TV\left(D_{s,c} \bmod \mathcal{P}(\mathcal{B}),\ \mathrm{Unif}(\mathcal{P}(\mathcal{B}))\right) \leq \varepsilon/2$$

provided $s \geq \eta_\varepsilon(\mathcal{B})$.

Proof: Denote by $Y(\cdot)$ the density of $D_{s,c} \bmod \mathcal{P}(\mathcal{B})$, so

$$Y(x) = \frac{1}{s^n} \sum_{u \in \Lambda} \rho_{s,c}(u+x)$$

$$= \frac{1}{s^n} \rho_{s,\, c-x}(\Lambda)$$

(Lemma 2 (poisson summation) below)

$$= \det(\Lambda^*) \sum_{w \in \Lambda^*} e^{-2\pi i \langle w, c-x \rangle} \rho_{1/s}(w)$$

$$= \det(\Lambda^*) \left(1 + \sum_{w \in \Lambda^* \setminus \{0\}} e^{-2\pi i \langle w, c-x \rangle} \rho_{1/s}(w)\right)$$

If $U(\cdot)$ denotes uniform density on $\mathcal{P}(\mathcal{B})$,

$$TV(Y, U) \triangleq \frac{1}{2} \int \left| \frac{Y(x)}{U(x)} - 1 \right| dU(x)$$

$$\leq \frac{1}{2} \sum_{w \in \Lambda^* \setminus \{0\}} \rho_{1/s}(w)$$

$$= \frac{1}{2} \rho_{1/s}(\Lambda^* \setminus \{0\}) \leq \varepsilon/2. \qquad \square$$

Lemma 2 (Poisson summation): for any "nice"

(i.e. infinitely differentiable w/ at least polynomially decaying tails)

function $f : \mathbb{R}^n \to \mathbb{C}$,

$$\sum_{y \in \Lambda} f(y) = \det(\Lambda^*) \cdot \sum_{w \in \Lambda^*} \hat{f}[w],$$

where $\hat{f}[w] \triangleq \int_{\mathbb{R}^n} f(y)\, e^{-2\pi i \langle w, y \rangle}\, dy$.

Proof sketch: Let's just show this for $\Lambda = \beta \cdot \mathbb{Z}$, for which $\Lambda^* = \beta^{-1}\mathbb{Z}$. Define $F(x) = \sum_{y=-\infty}^{\infty} f(x + \beta y)$, which as a function $\mathbb{R}/\alpha\cdot\mathbb{Z} \to \mathbb{C}$ has Fourier series

$$F(x) = \sum_{n=-\infty}^{\infty} a_n\, e^{2\pi i n x/\beta}$$

for

$$\begin{aligned}
a_n &= \frac{1}{\beta} \int_0^{\beta} F(x)\, e^{-2\pi i n x/\beta}\, dx \\
&= \frac{1}{\beta} \sum_{y=-\infty}^{\infty} \int_0^{\beta} f(x + \beta y)\, e^{-2\pi i n x/\beta}\, dx \\
&= \frac{1}{\beta} \sum_{y=-\infty}^{\infty} \int_0^{\beta} f(x + \beta y)\, e^{-2\pi i n (x + \beta y)/\beta}\, dx \\
&= \frac{1}{\beta} \int_{-\infty}^{-\infty} f(x)\, e^{-2\pi i n x/\beta}\, dx = \frac{1}{\beta} \hat{f}[n/\beta],
\end{aligned}$$

so

$$F(0) = \sum_{n=-\infty}^{\infty} a_n = \frac{1}{\beta} \sum_{n=-\infty}^{\infty} \hat{f}[n/\beta] = \det(\Lambda^*) \sum_{w \in \Lambda^*} \hat{f}[w]. \qquad \square$$

Can apply Lemma 2 in proof of Lemma 1 to $f = \rho_{s, c-x}$, noting that

$$\hat{f} = \hat{\rho}_{s,c-x} = e^{-2\pi i \langle c-x, \cdot \rangle} \hat{\rho}_s$$

(Fourier transform of Gaussian is Gaussian)

$$= e^{-2\pi i \langle c-x, \cdot \rangle} s^n \cdot \rho_{1/s}$$

Next lemma says that adding a wide enough discrete Gaussian to a wide enough continuous Gaussian results in distribution very close to a continuous Gaussian w/ the expected variance:

Lemma 3: Suppose $r, s > 0$ satisfy

$$\frac{rs}{\sqrt{r^2+s^2}} \geq \eta_\varepsilon(\Lambda).$$

Then $TV\left(D_{\Lambda+c, r} * D_s, D_{\sqrt{r^2+s^2}}\right) \lesssim \varepsilon$.

Pf: Let $Y(\cdot)$ denote density of $D_{\Lambda+c,r} * D_s$. Then

$$Y(x) = \frac{1}{s^n \rho_r(\Lambda+c)} \sum_{y \in \Lambda+c} \rho_r(y) \rho_s(x-y)$$

(some tedious algebra)

$$= \frac{1}{s^n} \rho_{\sqrt{r^2+s^2}}(x) \cdot \frac{\rho_{\frac{rs}{\sqrt{r^2+s^2}}, \frac{r^2}{r^2+s^2}x - c}(\Lambda)}{\rho_{r,-c}(\Lambda)}$$

$$\overset{\text{(Poisson summation)}}{=} \frac{1}{s^n}\rho_{\sqrt{r^2+s^2}}(x) \cdot \frac{\widehat{\rho_{\frac{rs}{\sqrt{r^2+s^2}},\, \frac{r^2}{r^2+s^2}x-c}}(\Lambda^*)}{\widehat{\rho_{r,-c}}(\Lambda^*)}$$

$$= \underbrace{\frac{1}{\left(\sqrt{r^2+s^2}\right)^n}\rho_{\sqrt{r^2+s^2}}(x)}_{D_{\sqrt{r^2+s^2}}(x)} \cdot \frac{\left(\frac{\sqrt{r^2+s^2}}{rs}\right)^n \widehat{\rho_{\frac{rs}{\sqrt{r^2+s^2}},\, \frac{r^2}{r^2+s^2}x-c}}(\Lambda^*)}{(1/r)^n\ \widehat{\rho_{r,-c}}(\Lambda^*)}$$

For any $w \in \Lambda^*$,

$$\text{(A)}\quad \left(\frac{\sqrt{r^2+s^2}}{rs}\right)^n \cdot \widehat{\rho_{\frac{rs}{\sqrt{r^2+s^2}},\, \frac{r^2}{r^2+s^2}x-c}}(w) = e^{-2\pi i\langle w,\, \frac{r^2}{r^2+s^2}x-c\rangle}\rho_{\frac{\sqrt{r^2+s^2}}{rs}}(w)$$

$$\text{(B)}\quad \frac{1}{r^n}\widehat{\rho_{r,-c}}(w) = e^{-2\pi i\langle w, -c\rangle}\rho_{1/r}(w)$$

$$\left|1-\sum_{w\in\Lambda^*}\text{(A)}\right| = \left|\sum_{w\in\Lambda^*\setminus\{0\}} e^{-2\pi i\langle w,\, \frac{r^2}{r^2+s^2}x-c\rangle}\rho_{\frac{\sqrt{r^2+s^2}}{rs}}(w)\right|$$

$$\leq \rho_{\frac{\sqrt{r^2+s^2}}{rs}}(\Lambda^*\setminus\{0\}) \leq \varepsilon$$

Similarly,

$$\left|1-\sum_{w\in\Lambda^*}\text{(B)}\right| \leq \rho_{1/r}(\Lambda^*\setminus\{0\}) \leq \varepsilon \quad \left(\text{b/c } 1/r \leq \frac{\sqrt{r^2+s^2}}{rs}\right).$$

So $Y(x) = \left(1 \pm O(\varepsilon)\right) \cdot D_{\sqrt{r^2+s^2}}(x)\ \forall x$, from which the lemma follows. $\square$

Corollary: For $z, c \in \mathbb{R}^n$, $r, \alpha > 0$, if

$$\frac{1}{\sqrt{1/r^2 + (\|z\|/\alpha)^2}} \geq \eta_\varepsilon(L),$$

then for $v \sim D_{\Lambda+c, r}$ and $e \sim D_\alpha$,

$$TV\left(\mathrm{law}(\langle z, v\rangle + e),\ D_{\sqrt{(r\|z\|)^2 + \alpha^2}}\right) \lesssim \varepsilon.$$

How big is $\eta_\varepsilon(\Lambda)$?

Lemma 4: $\eta_\varepsilon(\Lambda) \leq \sqrt{\frac{\ln(2n(1+1/\varepsilon))}{\pi}}\ \lambda_n(\Lambda)$

where $\lambda_n(\Lambda)$ is smallest radius r s.t. ball of radius r around 0 contains n linearly independent vectors in Λ.

Proof idea: Let $s = \sqrt{\frac{\ln(2n(1+1/\varepsilon))}{\pi}} \cdot \lambda_n(\Lambda)$. Idea is to show that for any v of norm $\leq \lambda_n(\Lambda)$, almost all of the mass for $D_{\Lambda^*, 1/s}$ comes from points in Λ^* orthogonal to v. So if there are n such v's that are linearly independent, then almost all mass comes from the origin, i.e. $\rho_{1/s}(\Lambda^* \setminus \{0\})$ will be small (in fact, exponentially small in s^2).

We are now ready to prove the main theorem.

① Discrete noise $\to$ continuous noise:

$\boxed{\text{Given a sample } (x,y), \text{ sample } e' \sim D_{\sigma'} \text{ for } \sigma' \gtrsim \eta_\delta(\mathbb{Z}) = \Theta(\lg 1/\delta) \text{ and form sample } (x, y+e').}$

Claim: If (x,y)'s $\sim$ LWE over $(\mathbb{Z}/q\mathbb{Z})^n$ with discrete Gaussian noise (resp if (x,y)'s $\sim$ $\mathrm{unif}((\mathbb{Z}/q\mathbb{Z})^n) \times \mathrm{Unif}(\mathbb{Z}/q\mathbb{Z})$), then distribution over $(x, y+e')$'s is $O(\delta m)$-close in TV to LWE with continuous Gaussian noise with slightly higher variance (resp $O(\delta m)$-close in TV to $\mathrm{unif}((\mathbb{Z}/q\mathbb{Z})^n) \times \mathrm{unif}(\mathbb{R}/q\mathbb{Z})$)

Pf: Ⓐ: if (x,y)'s $\sim$ LWE w/ discrete noise, then $y = (\langle x, w\rangle + e + e') \bmod q$. $TV(\mathrm{law}(e+e'), D_{\sqrt{\sigma^2+\sigma'^2}}) \lesssim \delta$ by Corollary above.

Ⓑ: if (x,y)'s uniform, then $TV(\mathrm{law}(e' \bmod 1), \mathrm{Unif}([0,1])) \lesssim \delta$ by Lemma 2 above, so $TV(\mathrm{law}(y+e' \bmod q), \mathrm{Unif}(\mathbb{R}/q\mathbb{Z})) \lesssim \delta$.

Claim follows by union bound over m samples. $\square$

② Discrete x's $\to$ uniform continuous x's:

$\boxed{\text{Given a sample } (x,y), \text{ sample } g \sim D_{\sigma'} \text{ for } \sigma' \gtrsim \eta_\delta(\mathbb{Z}^n) = \Theta(\lg 1/\delta) \text{ and form sample } ((x+g) \bmod q, y)}$

Claim: If (x,y)'s $\sim$ LWE over $(\mathbb{Z}/q\mathbb{Z})^n$ with Gaussian noise (resp if (x,y)'s $\sim \mathrm{Unif}((\mathbb{Z}/q\mathbb{Z})^n) \times \mathrm{Unif}(\mathbb{R}/q\mathbb{Z})$), then distribution over $((x+g) \bmod q, y)$'s is $O(\delta m)$-close in TV to LWE over $D = \mathrm{Unif}((\mathbb{R}/q\mathbb{Z})^n)$ with continuous Gaussian noise (resp. $O(\delta m)$-close in TV to $\mathrm{Unif}((\mathbb{R}/q\mathbb{Z})^n) \times \mathrm{Unif}(\mathbb{R}/q\mathbb{Z})$.

Pf: (A): if (x,y)'s $\sim$ LWE, then

$$y = (\langle x, w\rangle + e) \bmod q$$

$$= (\langle \underbrace{x+g}_{\substack{\text{close to}\\ \mathrm{Unif}((\mathbb{R}/q\mathbb{Z})^n)\\ \text{by Lemma 1}}}, w\rangle + e - \langle g, w\rangle) \bmod q$$

continuous Gaussian

conditioned on $(x+g) \bmod q$, distributed as discrete Gaussian

close to continuous Gaussian noise by Corollary above

(B) if (x,y)'s uniform, then

$(x+g) \bmod q$ distributed approximately as

$\mathrm{Unif}((\mathbb{R}/q\mathbb{Z})^n)$ by Lemma 1.

(3) Uniform continuous x's $\to$ Gaussian x's:

Given $x \sim \mathrm{Unif}([0,1])$,

for $\tau = \omega(\sqrt{\lg n})$, form $x' \sim D_{\mathbb{Z}+x,\tau}$

Note: $\langle x, w\rangle \equiv \langle x', w\rangle \bmod 1$ for $w \in \mathbb{Z}^n$.

because $D_{\mathbb{Z}+x,\tau}$ supported on points $x' \equiv x \bmod 1$

Let $\gamma(\cdot)$ denote density for x'.

$$\gamma(x') = \int_0^1 D_{\mathbb{Z}+x,\tau}(x)\,dx$$

$$= \int_0^1 \mathbb{1}[x'-x\in\mathbb{Z}]\cdot\frac{\rho_\tau(x')}{\rho_\tau(\mathbb{Z}+x)}\,dx$$

$$= \frac{\rho_\tau(x')}{\rho_\tau(\mathbb{Z}+x')} \quad \rho_\tau(\mathbb{Z}+x') \approx \rho_\tau(\mathbb{Z})$$

$$\approx \rho_\tau(x'),$$

so $\mathrm{law}(x')$ close to D_τ as desired.