Theorem (Gentry-Peikert-Vaikuntanathan). There is an efficient algorithm that takes a basis B of a lattice $\mathcal{L} = \mathcal{L}(B)$, a coset $c + \mathcal{L}$ and a Gaussian width parameter $s \geq \|\tilde{B}\| \cdot \omega(\sqrt{\log n})$ and outputs a sample whose distribution is statistically close to $D_{\mathcal{L},s,c}$

distribution is independent of B!

$\|\tilde{B}\| = \max_i \|\tilde{b}_i\|_2$ where $\{\tilde{b}_i\}$ is the Gram-Schmidt basis

Approach for preimage sampling:

forward sampling: choose $x \leftarrow D_{\mathbb{Z}^m,s}$ and output (x, Ax)

need to show that for $x \leftarrow D_{\mathbb{Z}^m,s}$, $(x, Ax) \stackrel{s}{\approx} (x, u)$ where $u \xleftarrow{R} \mathbb{Z}_q^n$

backward sampling: choose $y \xleftarrow{R} \mathbb{Z}_q^n$, compute any solution $z \in \mathbb{Z}_q^m$ where $Az = y$

sample $v \leftarrow D_{\mathcal{L}^\perp(A),s,-z}$ and output $v + z$

need to show that distribution of $v + z$ is statistically close to $D_{\mathbb{Z}^m,s}$

Outputs discrete Gaussian centered at $-z$ satisfying $Av = 0$.

To reason about this precisely, we need to introduce the notion of the "smoothing parameter" of a lattice:

- Intuitively: minimum amount of Gaussian blur needed to smooth out the discrete structure of a lattice

  e.g., minimum width $s > 0$ such that every coset $c + \mathcal{L}$ has the same Gaussian mass:

  for all $c \in \mathbb{R}^n$: $\rho_{s,c}(\mathcal{L}) \in [1 - \text{negl}(n), n] \cdot \rho_s(\mathcal{L})$

  whenever $s \geq \eta(\mathcal{L})$

  smoothing parameter of $\mathcal{L}$

  $\rho_s(\mathcal{L}) = \sum_{x \in \mathcal{L}} \rho_s(x)$

  Gaussian distribution over $\mathcal{L}$ is statistically close to Gaussian distribution over any coset $c + \mathcal{L}$

  implies $\rho_s(c + \mathcal{L}) = \sum_{x \in c+\mathcal{L}} \rho_s(x) = \sum_{y \in \mathcal{L}} \rho_s(c + y)$

  $= \rho_{s,-c}(\mathcal{L}) \stackrel{s}{\approx} \rho_s(\mathcal{L})$

  [Refer to [MR07] for formal definition]

  η typically parameterized by ε, here for simplicity, we set $\eta = \text{negl}(n)$

- Smoothing parameter can be bounded as follows:

  $$\eta(\mathcal{L}) \leq \lambda_n(\mathcal{L}) \cdot \omega(\sqrt{\log n})$$

  smallest ℓ such that $\mathcal{L}$ has a basis where all vectors have norm at most ℓ

- Visually (for 2D lattices):

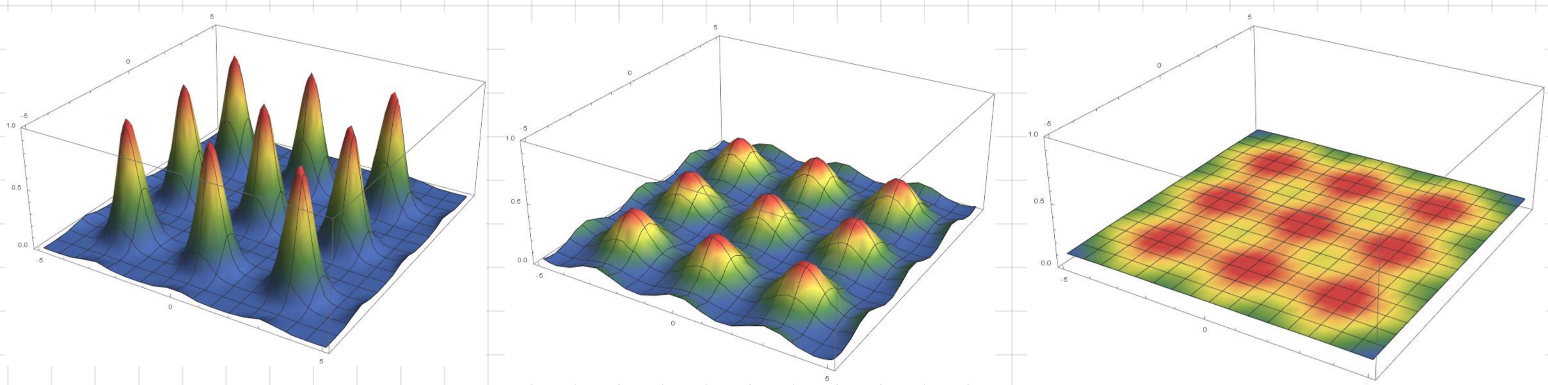

Figures show a Gaussian distribution centered at each lattice point

↳ Observe that as the width increases, distribution smoothes out and eventually looks uniform

↳ When the with is larger than the smoothing parameter, distribution over $\mathcal{L}$ and $c + \mathcal{L}$ (i.e. any translation of $\mathcal{L}$) is statistically indistinguishable

Let us now consider the distribution of Ax when $A \xleftarrow{R} \mathbb{Z}_q^{n \times m}$ and $x \leftarrow D_{\mathbb{Z}^m,s}$ where $s > \eta(\mathcal{L}^\perp(A))$ and $m > 3n \log q$

We claim this is uniform over $\mathbb{Z}_q^n$.

Proof. Recall that $\mathcal{L}^\perp(A) = \{v \in \mathbb{Z}_q^m : Av = 0\} \subseteq \mathbb{Z}_q^m$

We claim that $x \pmod{\mathcal{L}^\perp(A)}$ is statistically close to uniform over $\mathbb{Z}^m / \mathcal{L}^\perp(A)$

Take any coset $c + \mathcal{L}^\perp(A)$. Since $x \leftarrow D_{\mathbb{Z}^m,s}$,

since $s > \eta(\mathcal{L}^\perp(A))$

$$\Pr[x \in c + \mathcal{L}^\perp(A)] \propto \rho_s(c + \mathcal{L}^\perp(A)) \stackrel{s}{\approx} \rho_s(\mathcal{L}^\perp(A))$$

This holds for every coset $c + \mathcal{L}^\perp(A)$. Thus, $x \pmod{\mathcal{L}^\perp(A)} \overset{s}{\approx} \text{Uniform}(\mathbb{Z}_q^m / \mathcal{L}^\perp(A))$

$\left[\begin{array}{l} x \mapsto Ax \text{ is group homomorphism} \\ \mathcal{L}^\perp(A) \text{ is the kernel of } A \end{array}\right]$

Next, quotient group $\mathbb{Z}_q^m / \mathcal{L}^\perp(A) \cong \boxed{\mathbb{Z}_q^n}$ with isomorphism given by $x + \mathcal{L}^\perp(A) \mapsto Ax$

$\Rightarrow Ax$ is statistically close to uniform over $\mathbb{Z}_q$

$\hookrightarrow$ Technically $\mathbb{Z}_q^m / \mathcal{L}^\perp(A) \cong \text{range}(x \overset{\in \mathbb{Z}_q^m}{\mapsto} Ax)$. When $m > 3n \log q$, with prob $1-\text{negl}(n)$, the LHL says that

$$\{(A, x) : A \overset{R}{\leftarrow} \mathbb{Z}_q^{n \times m},\ x \overset{R}{\leftarrow} \{0,1\}^m\} \overset{s}{\approx} \{(A, u) : A \overset{R}{\leftarrow} \mathbb{Z}_q^{n\times m},\ u \overset{R}{\leftarrow} \mathbb{Z}_q^n\}$$

with statistical distance $\varepsilon = \frac{1}{2}\sqrt{q^n / 2^m} < q^{-n}$

so $\text{range}(x \mapsto Ax) = \mathbb{Z}_q^n$ with $1-\text{negl}(n)$ probability

Visually:

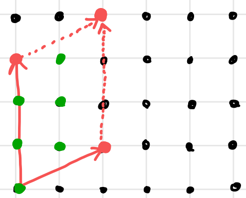

Suppose these vectors are a basis for $\mathcal{L}^\perp(A)$. For each $x \in \mathcal{L}^\perp(A)$, we have $Ax = 0 \pmod q$.

Elements in green are elements of $\mathbb{Z}^2 / \mathcal{L}^\perp(A)$. When we sample $x \leftarrow D_{\mathbb{Z}^2, s}$, with $s > \eta(\mathcal{L}^\perp(A))$, $x \pmod{\mathcal{L}^\perp(A)}$ is uniform random over $\mathbb{Z}^2 / \mathcal{L}^\perp(A)$.

Each element of $\mathbb{Z}^2 / \mathcal{L}^\perp(A)$ is associated with a value Ax.

Thus, forward sampling satisfies $(x, Ax) \overset{s}{\approx} (x, u)$ when $x \leftarrow D_{\mathbb{Z}^m, s}$, $A \overset{R}{\leftarrow} \mathbb{Z}_q^{n \times m}$, $u \overset{R}{\leftarrow} \mathbb{Z}_q^n$.

Now, we need to show that backward sampling (x, y) where $y \overset{R}{\leftarrow} \mathbb{Z}_q^n$, $x \leftarrow v + z$ where $v \leftarrow D_{\mathcal{L}^\perp(A), s, -z}$ and $Az = y$ yields the correct distribution.

Suppose we sample (x, y) using the forward sampling procedure. Then, y is uniform over $\mathbb{Z}_q^n$ so consider distribution of x conditioned on y. This is the distribution of $x \leftarrow D_{\mathbb{Z}^m, s}$ given $Ax = y$. The support of this distribution is $z + \mathcal{L}^\perp(A)$ where $z \in \mathbb{Z}_q^m$ is any solution satisfying $Az = y$. Thus, we can write

$$D(\hat{x}) = \frac{\rho_s(\hat{x})}{\rho_s(z + \mathcal{L}^\perp(A))} = \frac{\rho_{s,-z}(\hat{x} - z)}{\rho_{s,-z}(\mathcal{L}^\perp(A))} = D_{\mathcal{L}^\perp(A), s, -z}(\hat{x} - z)$$

probability of sampling any v' such that $Av = y$ (since $Az = y$)

If we write $x = z + v$, then $v = x - z$. The distribution of v is then precisely $D_{\mathcal{L}^\perp(A), s, -z}$:

$$D_v(\hat{v}) = D_x(\hat{v} + z) = D_{\mathcal{L}^\perp(A), s, -z}(\hat{v})$$

sampling value $\hat{v}$ from distribution of v

since $v = x - z$, event $v = \hat{v}$ corresponds to sampling $\hat{x} = \hat{v} + z$

It suffices now to show how to sample from $D_{\mathcal{L}, s, c}$ given a "sufficiently-good" basis B for $\mathcal{L} = \mathcal{L}(B)$. $[s > \|\tilde{B}\| \cdot \omega(\log n)]$

<u>Naïve approach</u>: sample a continuous Gaussian over $\mathbb{R}^n$ and "round" to the nearest lattice point

$\hookrightarrow$ This yields the "rounded Gaussian distribution" which is statistically far from the discrete Gaussian (even over $\mathbb{Z}$)!

To see this, suppose we want to sample from $D_{\mathbb{Z}, s}$ by sampling a continuous Gaussian with parameter s and rounding. Consider probability mass assigned to 0:

- Rounded Gaussian: $y \leftarrow \text{Gaussian}(s)$ rounds to 0 if $y \in [-\frac{1}{2}, \frac{1}{2})$.

Gauss error function
$$\text{erf}(z) = \frac{2}{\sqrt{\pi}} \int_0^z e^{-t^2} dt$$

$$\Pr[y \in [-\tfrac{1}{2}, \tfrac{1}{2})] = \frac{1}{s} \int_{-1/2}^{1/2} \rho_s(y)\, dy = \frac{2}{s} \int_0^{1/2} e^{-\pi y^2 / s^2} dy = \frac{2}{\sqrt{\pi}} \int_0^{\sqrt{\pi}/2s} e^{-t^2} dt = \text{erf}(\sqrt{\pi}/2s)$$

normalization parameter $\int_{-\infty}^{\infty} \rho_s(y)\, dy = s$

by symmetry of ρ_s

substituting $t = \sqrt{\pi}\, y / s$

Taylor expansion for $z < 1$:
$$\text{erf}(z) = \frac{2}{\sqrt{\pi}}(z - \Omega(z^3))$$

$\Rightarrow$ for large s, $\text{erf}(\sqrt{\pi}/2s) = \frac{1}{s} - \Omega\left(\frac{1}{s^3}\right)$

- Discrete Gaussian: since the Fourier transform of $\rho_s(x)$ is $\hat{\rho}_s(y) = s \cdot \rho_{1/s}(y)$, we can write by properties of the Fourier transform:

$$\rho_s(\mathbb{Z}) = \sum_{x \in \mathbb{Z}} \rho_s(x) = \sum_{y \in \mathbb{Z}} s \cdot \rho_{1/s}(y) = \sum_{y \in \mathbb{Z}} s\, e^{-\pi s^2 y^2} = s(1 + \text{negl}(\lambda))$$

[see optional addendum for proof]

whenever $s = \omega(\sqrt{\log \lambda})$ $\quad [e^{-s^2} = e^{-\omega(\log \lambda)} = \text{negl}(\lambda)]$

$\Rightarrow$ probability of sampling 0 is then negligibly close to $1/s$

Statistical distance between discrete Gaussian and a rounded Gaussian is at least $\Omega(1/s^3)$. Even larger in higher dimensions!

For applications that require pre-image sampling for security, discrete Gaussian sampling is very important. Other distributions may not be simulatable and vulnerable to attack!

An optional aside. We will show that $\sum_{x \in \mathbb{Z}} \rho_s(x) = s \cdot \sum_{y \in \mathbb{Z}} \rho_{1/s}(y)$

We say a function $f: \mathbb{R} \to \mathbb{C}$ is absolutely if $\int_{-\infty}^{\infty} |f(x)|\, dx < \infty$

For an absolutely integrable function $f: \mathbb{R} \to \mathbb{C}$, we define its Fourier transform $\hat{f}: \mathbb{R} \to \mathbb{C}$ to be

$$\hat{f}(y) = \int_{-\infty}^{\infty} f(x)\, e^{-2\pi i x y}\, dx$$

When $f, \hat{f}$ are absolutely integrable and f is continuous, then we can define the inverse Fourier transform

$$f(x) = \int_{-\infty}^{\infty} \hat{f}(y)\, e^{2\pi i x y}\, dy$$

Consider the Fourier transform of the Gaussian function ρ_s (in one dimension):

$$\hat{\rho}_s(y) = \int_{-\infty}^{\infty} \rho_s(x)\, e^{-2\pi i x y}\, dx$$

$$= s\, e^{-\pi s^2 y^2} = s \cdot \rho_{1/s}(y)$$

(See standard textbook of Fourier analysis or use Mathematica ☺)

(In particular, Fourier transform of Gaussian is Gaussian).

Suppose $f: \mathbb{R} \to \mathbb{C}$ is $\mathbb{Z}$-periodic. Namely, $f(x+y) = f(x)$ for all $x \in \mathbb{R}$ and $y \in \mathbb{Z}$. We define its Fourier series $\hat{f}: \mathbb{Z} \to \mathbb{C}$ as

$$\hat{f}(y) = \int_0^1 f(x)\, e^{-2\pi i x y}\, dx$$

The Fourier inversion formula allows us to write

$$f(x) = \sum_{y \in \mathbb{Z}} \hat{f}(y)\, e^{2\pi i x y}$$

We now show that for any well-behaved $f: \mathbb{R} \to \mathbb{C}$, it holds that

$$\sum_{x \in \mathbb{Z}} f(x) = \sum_{y \in \mathbb{Z}} \hat{f}(y)$$

Define the function $\phi(x) = \sum_{z \in \mathbb{Z}} f(x+z)$. Since ϕ is $\mathbb{Z}$-periodic:

$$\hat{\phi}(y) = \int_0^1 \phi(x)\, e^{-2\pi i x y}\, dx$$

domain of $\hat{\phi}$ is $\mathbb{Z}$

$$= \int_0^1 \sum_{z \in \mathbb{Z}} f(x+z)\, e^{-2\pi i x y}\, dx$$

$$= \sum_{z \in \mathbb{Z}} \int_0^1 f(x+z)\, e^{-2\pi i x y}\, dx$$

can interchange summation + integration if f is "well-behaved" (see Fubini's theorem for precise condition)

$$= \sum_{z \in \mathbb{Z}} \int_0^1 f(x+z)\, e^{-2\pi i (x+z) y}$$

since $yz \in \mathbb{Z}$ so $e^{-2\pi i y z} = 1$

$$= \int_{-\infty}^{\infty} f(x)\, e^{-2\pi i x y}\, dx$$

$$= \hat{f}(y)$$