**CS 395T: Topics in Cryptography (Lattice-Based Cryptography)**

# Discrete Gaussian Sampling Summary

**Instructor:** David Wu

Here, we summarize some key results on sampling discrete Gaussians over lattices. Much of the material is adapted from [Pei16, GPV08, Pei10, MP12].

**Gaussians.** We define the $n$-dimensional (spherical) Gaussian function $\rho_s \colon \mathbb{R}^n \to (0,1]$ with width $s > 0$ to be the function

$$\rho_s(\mathbf{x}) := \exp(-\pi \|\mathbf{x}\|^2 / s^2).$$

For a center $\mathbf{c} \in \mathbb{R}^n$, we define the Gaussian with width $s$ centered at $\mathbf{c}$ to be the function

$$\rho_{s,\mathbf{c}} := \exp(-\pi \|\mathbf{x} - \mathbf{c}\|^2 / s^2).$$

The $n$-dimensional Gaussian function with *covariance* $\mathbf{\Sigma} \in \mathbb{R}^{n \times n}$ is the function

$$\rho_{\sqrt{\mathbf{\Sigma}}}(\mathbf{x}) := \exp(-\pi \cdot \mathbf{x}^\mathsf{T} \mathbf{\Sigma}^{-1} \mathbf{x}).$$

The covariance of a spherical Gaussian with parameter $s$ is simply $s\mathbf{I}_n$, where $\mathbf{I}_n$ is the $n \times n$ identity matrix. Note that the covariance matrix is always *positive definite* (i.e., there exists $\mathbf{B} \in \mathbb{R}^{n \times m}$ such that $\mathbf{\Sigma} = \mathbf{B}\mathbf{B}^\mathsf{T}$). If $\mathbf{x}$ is a (spherical) Gaussian with parameter $s$, then $\mathbf{R}\mathbf{x}$ is a Gaussian with covariance $\mathbf{R}\mathbf{R}^\mathsf{T}$.

**Discrete Gaussians over lattices.** Let $\mathcal{L} = \mathcal{L}(\mathbf{B})$ be a lattice. The (spherical) discrete Gaussian distribution $D_{\mathcal{L},s}$ on a lattice coset $\mathbf{c} + \mathcal{L}$ is simply the Gaussian distribution with parameter $s$ with its support restricted to $\mathbf{c} + \mathcal{L}$. Namely, for $\mathbf{x} \in \mathbf{c} + \mathcal{L}$,

$$D_{\mathbf{c}+\mathcal{L},s}(\mathbf{x}) := \frac{\rho_s(\mathbf{x})}{\rho_s(\mathbf{c} + \mathcal{L})} = \frac{\rho_s(\mathbf{x})}{\sum_{\mathbf{y} \in \mathbf{c} + \mathcal{L}} \rho_s(\mathbf{y})},$$

and for $\mathbf{x} \notin \mathbf{c} + \mathcal{L}$, $D_{\mathbf{c}+\mathcal{L},s}(\mathbf{x}) = 0$. This definition naturally extends to non-spherical Gaussians.

**Theorem 1** ([GPV08])**.** *There exists an efficient algorithm that takes as input a basis* $\mathbf{B}$ *for a lattice* $\mathcal{L} = \mathcal{L}(\mathbf{B})$*, any coset* $\mathbf{c} + \mathcal{L}$*, and any width parameter* $s \geq \|\tilde{\mathbf{B}}\| \cdot \omega(\sqrt{\log n})$ *and outputs a sample that is statistically close to* $D_{\mathbf{c}+\mathcal{L},s}$*.*

**The SIS lattice.** For a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, the SIS lattice is defined as

$$\mathcal{L}^\perp(\mathbf{A}) := \left\{ \mathbf{x} \in \mathbb{Z}_q^m : \mathbf{A}\mathbf{x} = \mathbf{0} \bmod q \right\} \supseteq q\mathbb{Z}^m.$$

For a vector $\mathbf{u} \in \mathbb{Z}_q^n$, we define

$$\mathcal{L}_\mathbf{u}^\perp(\mathbf{A}) := \left\{ \mathbf{x} \in \mathbb{Z}_q^m : \mathbf{A}\mathbf{x} = \mathbf{u} \bmod q \right\} = \mathbf{z} + \mathcal{L}^\perp(\mathbf{A}),$$

for some $\mathbf{z} \in \mathbb{Z}_q^m$ where $\mathbf{A}\mathbf{z} = \mathbf{u}$.

**Gadget trapdoors.** We say that $\mathbf{R} \in \mathbb{Z}_q^{m \times n\ell}$ is a gadget trapdoor for $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ if $\mathbf{A}\mathbf{R} = \mathbf{G}$. We can take the following approach (from Micciancio and Peikert [MP12]) to sample from $D_{\mathcal{L}_\mathbf{u}^\perp(\mathbf{A}),s}$ using a gadget trapdoor for $\mathbf{A}$:

- Set $s' = \omega(\sqrt{\log n})$. Sample a perturbation vector $\mathbf{p} \leftarrow D_{\mathbb{Z}^m, s^2\mathbf{I}_m - (s')^2\mathbf{R}\mathbf{R}^\mathsf{T}}$. We can do this as long as $s^2\mathbf{I}_m - (s')^2\mathbf{R}\mathbf{R}^\mathsf{T}$ is positive definite. Taking $s = s_1(\mathbf{R}) \cdot \omega(\sqrt{\log n})$, where $s_1(\mathbf{R}) := \max_{\|\mathbf{u}\|=1} \|\mathbf{R}\mathbf{u}\|$ denotes the largest singular value of $\mathbf{R}$ suffices here. When $s^2\mathbf{I}_m - (s')^2\mathbf{R}\mathbf{R}^\mathsf{T}$, we can decompose it as $\mathbf{L}\mathbf{L}^\mathsf{T}$ (e.g., by computing its Cholesky decomposition). Then, we can sample $\mathbf{p}$ by first sampling $\mathbf{p}' \leftarrow D_{\mathbb{Z}^m,1}$ (using Theorem 1) and setting $\mathbf{p} \leftarrow \mathbf{L}\mathbf{p}'$.

- Let $\mathbf{z} \leftarrow \mathbf{u} - \mathbf{A}\mathbf{p}$. Sample $\mathbf{y} \leftarrow D_{\mathcal{L}^{\perp}_{\mathbf{z}}(\mathbf{G}), s'}$. Recall that $\mathbf{G}$ has a basis $\mathbf{B}$ where $\|\tilde{\mathbf{B}}\| \le \sqrt{5}$ (when $q$ is a power of 2, $\|\tilde{\mathbf{B}}\| = 2$), so we can use Theorem 1 to implement this step.
- Output $\mathbf{x} \leftarrow \mathbf{R}\mathbf{y} + \mathbf{p}$.

For correctness, observe that

$$\mathbf{A}\mathbf{x} = \mathbf{A}\mathbf{R}\mathbf{y} + \mathbf{A}\mathbf{p} = \mathbf{G}\mathbf{y} + \mathbf{A}\mathbf{p} = \mathbf{z} + \mathbf{A}\mathbf{p} = \mathbf{u},$$

so $\mathbf{x} \in \mathcal{L}^{\perp}_{\mathbf{u}}(\mathbf{A})$. Consider the distribution of $\mathbf{x}$. The distribution of $\mathbf{y}$ is a discrete Gaussian with width $s'$, so $\mathbf{R}\mathbf{y}$ is a discrete Gaussian with covariance $(s')^2\mathbf{R}\mathbf{R}^{\mathsf{T}}$. The vector $\mathbf{p}$ is Gaussian with covariance $s^2\mathbf{I}_m - (s')^2\mathbf{R}\mathbf{R}^{\mathsf{T}}$, so by the Gaussian convolution lemma (see [Pei10] for a precise description), the sum $\mathbf{R}\mathbf{y} + \mathbf{p}$ is statistically close to a discrete Gaussian with covariance $(s')^2\mathbf{R}\mathbf{R}^{\mathsf{T}} + (s^2\mathbf{I}_m - (s')^2\mathbf{R}\mathbf{R}^{\mathsf{T}}) = s^2\mathbf{I}_m$. This precisely coincides with the desired distribution $D_{\mathcal{L}^{\perp}_{\mathbf{u}}(\mathbf{A}), s}$. Refer to [MP12] for more details.

**Preimage sampleable trapdoor functions.** Using the above algorithm, we can construct a preimage sampleable trapdoor function as follows:

- $\mathsf{TrapGen}(n, q)$: On input lattice parameters $n, q$, set $\bar{m} = 3n \log q$, let $t = n\lceil \log q \rceil$, and $m = \bar{m} + t$. Sample $\bar{\mathbf{A}} \xleftarrow{\text{R}} \mathbb{Z}_q^{n \times \bar{m}}$ and $\bar{\mathbf{R}} \leftarrow \{0,1\}^{m \times t}$. Construct matrices

  $$\mathbf{A} = [\bar{\mathbf{A}} \mid \mathbf{G} - \bar{\mathbf{A}}\bar{\mathbf{R}}] \in \mathbb{Z}_q^{n \times m} \qquad \mathbf{R} = \begin{bmatrix} \bar{\mathbf{R}} \\ \mathbf{I}_t \end{bmatrix} \in \mathbb{Z}_q^{m \times t}.$$

  Output the public matrix $\mathbf{A}$ and the trapdoor $\mathbf{R}$. Note that we can also sample $\mathbf{R}$ from other distributions to get smaller parameters; see [MP12].
- $\mathsf{SampleGaussian}(m, s)$: On input the dimension $m$, sample and output $\mathbf{x} \leftarrow D_{\mathbb{Z}^m, s}$ (e.g., using Theorem 1).
- $\mathsf{SamplePre}(\mathbf{A}, \mathbf{R}, \mathbf{u}, s)$: On input the public matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times 2m}$, a trapdoor $\mathbf{R} \in \mathbb{Z}_q^{2m \times m}$, and a target vector $\mathbf{u} \in \mathbb{Z}_q^{2m}$, sample and output $\mathbf{x} \leftarrow D_{\mathcal{L}^{\perp}_{\mathbf{u}}(\mathbf{A}), s}$ (using the procedure described above).

The above algorithms satisfy the following properties:

- Let $(\mathbf{A}, \mathbf{R}) \leftarrow \mathsf{TrapGen}(n, q)$. By the leftover hash lemma, the distribution of $\mathbf{A}$ is statistically close to uniform over $\mathbb{Z}_q^{n \times m}$. Since $\mathbf{R} \in \{0,1\}^{m \times t}$, we can naïvely bound $s_1(\mathbf{R})$ by $\sqrt{mt} = O(n \log q)$.
- Let $\mathbf{x} \leftarrow \mathsf{SampleGaussian}(m, s)$. If $s \ge s_1(\mathbf{R}) \cdot \omega(\sqrt{\log n})$, then the distribution of $\mathbf{A}\mathbf{x}$ is statistically close to uniform over $\mathbb{Z}_q^n$. This follows from the fact that $\eta(\mathcal{L}^{\perp}(\mathbf{A})) \le s_1(\mathbf{R}) \cdot \omega(\sqrt{\log n})$ (see [MP12, Lemma 5.3]) and the result shown from class.
- When $s \ge s_1(\mathbf{R}) \cdot \omega(\sqrt{\log n})$, the following two distributions are statistically indistinguishable:

  $$\{\mathbf{x} \leftarrow \mathsf{SampleGaussian}(m, s) : (\mathbf{x}, \mathbf{A}\mathbf{x})\} \text{ and } \left\{\mathbf{y} \xleftarrow{\text{R}} \mathbb{Z}_q^n, \mathbf{x} \leftarrow \mathsf{SamplePre}(\mathbf{A}, \mathbf{R}, \mathbf{y}, s) : (\mathbf{x}, \mathbf{y})\right\}.$$

# References


[GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pages 197–206, 2008.

[MP12] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *EUROCRYPT*, pages 700–718, 2012.

[Pei10] Chris Peikert. An efficient and parallel gaussian sampler for lattices. In *CRYPTO*, pages 80–97, 2010.

[Pei16] Chris Peikert. A decade of lattice cryptography. *Found. Trends Theor. Comput. Sci.*, 10(4):283–424, 2016.