

CS 294. Worst-case to Average-case Reduction for SIS and LWE

In this lecture, we will see a few basic theorems on lattices and a property called smoothing. We will then use smoothing as a tool to come up with worst-case to average-case reductions for SIS and LWE. In a nutshell, the worst-case to average-case reductions show how to transform any algorithm that solves SIS/LWE *on the average* into an algorithm that solves “approximate short vector problems” on lattices *in the worst case*.

1 Lattices and Minkowski’s Theorem

We define the lattice $\mathcal{L}(\mathbf{B})$ and its fundamental parallelepiped as follows:

$$\mathcal{L}(\mathbf{B}) = \mathbf{B}\mathbb{Z}^n \text{ and } \mathcal{P}(\mathbf{B}) = \mathbf{B}[0, 1]^n$$

As additive groups, $\mathcal{P}(\mathbf{B}) = \mathbb{R}^n / \mathcal{L}(\mathbf{B})$. (Think of the one-dimensional analogy: the torus $[0, 1) = \mathbb{R} / \mathbb{Z}$.)

We define the determinant of the lattice as the volume of the fundamental parallelepiped. While the parallelepiped is itself defined by a basis and is basis-dependent, its volume is a lattice-invariant. Analytically, the determinant $\det(\mathcal{L})$ of full-rank lattices can be computed as the determinant of the basis matrix \mathbf{B} .

We define $\lambda_1(\mathcal{L})$ to be the length of the shortest non-zero vector of the lattice. Alternatively, and apparently somewhat more convolutedly,

$$\lambda_1(\mathcal{L}) = \inf_{r \in \mathbb{R}} \dim(\text{Span}(\mathcal{L} \cap \mathcal{B}(\mathbf{0}, r))) \geq 1$$

where $\mathcal{B}(\mathbf{0}, r)$ denotes the ball of radius r centered at $\mathbf{0}$. In words, the smallest r such that the space generated by lattice vectors of length at most r has dimension at least 1. This naturally leads us to the definition of the i -th minimum $\lambda_i(\mathcal{L})$:

$$\lambda_i(\mathcal{L}) = \inf_{r \in \mathbb{R}} \dim(\text{Span}(\mathcal{L} \cap \mathcal{B}(\mathbf{0}, r))) \geq i$$

Our intuition tells us that as the determinant of the lattice gets smaller, the lattice gets denser, and therefore has shorter vectors. Minkowski’s theorems formalize this intuition.

Lemma 1 (Minkowski). *For any rank- n lattice \mathcal{L} , we have*

- $\lambda_1(\mathcal{L}) \leq \sqrt{n} \cdot \det(\mathcal{L})^{1/n}$; and even stronger,
- $\det(\mathcal{L}) \leq \prod_{i=1}^n \lambda_i(\mathcal{L}) \leq n^{n/2} \cdot \det(\mathcal{L})$.

Computational Problems. The γ -approximate shortest vector problem (SVP) is to find a non-zero vector $\mathbf{v} \in \mathcal{L}(\mathbf{B})$ given a basis \mathbf{B} such that $\|\mathbf{v}\| \leq \gamma \cdot \lambda_1(\mathcal{L})$. The γ -approximate shortest independent vectors problem (SIVP) is to find n linearly independent vectors $\mathbf{v}_1, \dots, \mathbf{v}_n$ such that $\|\mathbf{v}_i\| \leq \gamma \cdot \lambda_n(\mathcal{L})$.

2 Lattice Smoothing

2.1 Lattice Duality

For a rank- n lattice \mathcal{L} , its dual denoted \mathcal{L}^* is defined as

$$\mathcal{L}^* = \{\mathbf{x} \in \mathbb{R}^n : \forall \mathbf{y} \in \mathcal{L}, \langle \mathbf{x}, \mathbf{y} \rangle \in \mathbb{Z}\}$$

Indeed, each dual lattice vector \mathbf{x} corresponds to a linear function $\phi_{\mathbf{x}} : \mathcal{L} \rightarrow \mathbb{Z}$ and the dual lattice corresponds to a basis of the space of such linear functions.

Let us start with examples and some properties.

- In one dimension, the only possible lattices are $k\mathbb{Z}$. Its dual is $(1/k) \cdot \mathbb{Z}$.
- The dual of \mathbb{Z}^n is \mathbb{Z}^n itself.
- If $\mathcal{L} = \mathcal{L}(\mathbf{B})$ for a basis matrix $\mathbf{B} \in \mathbb{R}^{n \times n}$, then \mathcal{L}^* is generated by the columns of \mathbf{B}^{-T} , its transposed inverse. Indeed, the pairwise inner products of the basis vectors and the dual basis vectors is captured in the matrix $\mathbf{B} \cdot (\mathbf{B}^{-T})^T = \mathbf{I}$.

The determinant of the dual lattice is immediately seen to be the inverse of the determinant of the lattice. In an intuitive sense, as the lattice gets sparser (the determinant gets larger), the dual lattice gets denser (its determinant gets smaller). This leads us to the following lemma.

Lemma 2. *For any rank- n lattice \mathcal{L} , $\lambda_1(\mathcal{L}^*) \cdot \lambda_1(\mathcal{L}) \leq n$.*

Proof. We know from Minkowski that

$$\lambda_1(\mathcal{L}) \leq \sqrt{n}(\det(\mathcal{L}))^{1/n} \text{ and } \lambda_1(\mathcal{L}^*) \leq \sqrt{n}(\det(\mathcal{L}^*))^{1/n}$$

Multiplying the two, we get

$$\lambda_1(\mathcal{L}) \cdot \lambda_1(\mathcal{L}^*) \leq n$$

as desired. □

In fact, using far more advanced tools, we can show something stronger, namely that $\lambda_1(\mathcal{L}^*) \cdot \lambda_n(\mathcal{L}) \leq n$. The following lemma goes in the other direction, has an elementary proof, and we will find it useful later on.

Lemma 3. *For any rank- n lattice \mathcal{L} , $\lambda_n(\mathcal{L}^*) \cdot \lambda_1(\mathcal{L}) \geq 1$.*

Proof. Let $\mathbf{x} \in \mathcal{L}$ be the shortest non-zero vector. Let $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathcal{L}^*$ be linearly independent. At least one of the \mathbf{v}_i has a non-zero inner product with \mathbf{x} , say $\langle \mathbf{v}_i, \mathbf{x} \rangle > 0$. Since the inner products of lattice vectors and dual vectors are integers, $\langle \mathbf{v}_i, \mathbf{x} \rangle \geq 1$. Therefore,

$$\lambda_1(\mathcal{L}) = \|\mathbf{x}\| \geq 1/\|\mathbf{v}_i\| \geq 1/\lambda_n(\mathcal{L}^*) .$$

□

2.2 Gaussians

The Gaussian function over \mathbb{R} with (zero mean and) parameter s is defined as

$$\rho_s(x) = e^{-\pi x^2/s^2}$$

We note that

$$\int_{-\infty}^{\infty} \rho_s(x) dx = \int_{-\infty}^{\infty} e^{-\pi x^2/s^2} dx = \frac{s}{\sqrt{\pi}} \int_{-\infty}^{\infty} e^{-z^2} dz = s$$

where the second equality is by a change of variables and the third by using value of the Gaussian integral ($= \sqrt{\pi}$). This fact can be used to turn the Gaussian function into a probability distribution over the reals by scaling ρ_s by $1/s$.

Something very similar can be done in n dimensions. That is, the n -dimensional Gaussian function over \mathbb{R}^n is defined as

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

This can again be turned into a probability distribution after scaling by $1/s^n$.

2.3 Basic Fourier Analysis

We call a function $f : \mathbb{R}^n \rightarrow \mathbb{C}$ “nice” if it is absolutely integrable, that is, $\int_{\mathbb{R}^n} |f(\mathbf{x})| d\mathbf{x} < \infty$.

Definition 4 (Fourier Transform). *For a nice function $f : \mathbb{R}^n \rightarrow \mathbb{C}$, we define its Fourier transform $\hat{f} : \mathbb{R}^n \rightarrow \mathbb{C}$ as*

$$\hat{f}(\mathbf{y}) = \int_{\mathbb{R}^n} f(\mathbf{x}) e^{-2\pi i \langle \mathbf{x}, \mathbf{y} \rangle} d\mathbf{x}$$

If f, \hat{f} are nice and f is continuous, we can recover a function from its Fourier transform using the inverse formula:

$$f(\mathbf{x}) = \int_{\mathbb{R}^n} \hat{f}(\mathbf{y}) e^{2\pi i \langle \mathbf{x}, \mathbf{y} \rangle} d\mathbf{y}$$

Lemma 5 (Fourier Transform of the Gaussian function). *Let $\hat{\rho}_s$ denote the Fourier transform of the Gaussian function ρ_s . Then,*

$$\hat{\rho}_s(\mathbf{x}) = s^n \cdot \rho_{1/s}(\mathbf{x})$$

Proof. We provide a proof in one dimension.

$$\begin{aligned} \hat{\rho}_s(\mathbf{y}) &= \int_{\mathbb{R}^n} \rho_s(\mathbf{x}) e^{-2\pi i \langle \mathbf{x}, \mathbf{y} \rangle} d\mathbf{x} \\ &= \int_{\mathbb{R}^n} e^{-\pi \|\mathbf{x}\|^2/s^2} e^{-2\pi i \langle \mathbf{x}, \mathbf{y} \rangle} d\mathbf{x} \\ &= e^{-\pi s^2 \|\mathbf{y}\|^2} \int_{\mathbb{R}^n} e^{-\pi \|\mathbf{x}/s + i s \mathbf{y}\|^2} d\mathbf{x} \end{aligned}$$

The latter integral, on a complex change of variables becomes $s^n \cdot \int_{\mathbb{R}^n} e^{-\pi \|\mathbf{z}\|^2} d\mathbf{z}$ which is simply s^n . So,

$$\hat{\rho}_s(\mathbf{y}) = s^n e^{-\pi s^2 \|\mathbf{y}\|^2} = s^n \cdot \rho_{1/s}(\mathbf{y})$$

□

For periodic functions, we have the closely related notion of Fourier series.

Definition 6 (Fourier Series). *We will define Fourier series for periodic functions. For a “nice enough” function $f : \mathbb{R}^n \rightarrow \mathbb{C}$ that is \mathcal{L} -periodic, that is, $f(\mathbf{x} + \mathbf{y}) = f(\mathbf{x})$ for all $\mathbf{x} \in \mathbb{R}^n$ and $\mathbf{y} \in \mathcal{L}$, we have its Fourier series $\hat{f} : \mathcal{L}^* \rightarrow \mathbb{C}$ defined as*

$$\hat{f}(\mathbf{y}) = \frac{1}{\det(\mathcal{L})} \cdot \int_{\mathcal{P}(\mathcal{L})} f(\mathbf{x}) e^{-2\pi i \langle \mathbf{x}, \mathbf{y} \rangle} d\mathbf{x}$$

We will state the Fourier inversion formula below without proof.

Lemma 7 (Fourier Inversion). $f(\mathbf{x}) = \sum_{\mathbf{y} \in \mathcal{L}^*} \hat{f}(\mathbf{y}) e^{2\pi i \langle \mathbf{x}, \mathbf{y} \rangle}$.

An important fact that connects a function f and its Fourier transform is the Poisson Summation formula. The proof of this formula goes via the Fourier series.

Lemma 8 (Poisson Summation). *Given $f : \mathbb{R}^n \rightarrow \mathbb{C}$, and any full-rank lattice \mathcal{L} , we have*

$$\sum_{\mathbf{x} \in \mathcal{L}} f(\mathbf{x}) = \frac{1}{\det(\mathcal{L})} \cdot \sum_{\mathbf{y} \in \mathcal{L}^*} \hat{f}(\mathbf{y}) = \det(\mathcal{L}^*) \cdot \sum_{\mathbf{y} \in \mathcal{L}^*} \hat{f}(\mathbf{y})$$

Proof. Although f is not periodic, the proof of Poisson summation goes through the Fourier series of a “periodized” f . In particular, consider the function

$$\phi(\mathbf{x}) = \sum_{\mathbf{z} \in \mathcal{L}} f(\mathbf{x} + \mathbf{z})$$

Clearly ϕ is periodic over \mathcal{L} , therefore $\hat{\phi}$ is defined over \mathcal{L}^* . For any $\mathbf{y} \in \mathcal{L}^*$, we have

$$\begin{aligned} \hat{\phi}(\mathbf{y}) &= \det(\mathcal{L}^*) \int_{\mathbf{x} \in \mathcal{P}(\mathcal{L})} \phi(\mathbf{x}) e^{-2\pi i \langle \mathbf{x}, \mathbf{y} \rangle} d\mathbf{x} \\ &= \det(\mathcal{L}^*) \int_{\mathbf{x} \in \mathcal{P}(\mathcal{L})} \left(\sum_{\mathbf{z} \in \mathcal{L}} f(\mathbf{x} + \mathbf{z}) \right) e^{-2\pi i \langle \mathbf{x}, \mathbf{y} \rangle} d\mathbf{x} \\ &= \det(\mathcal{L}^*) \sum_{\mathbf{z} \in \mathcal{L}} \int_{\mathbf{x} \in \mathcal{P}(\mathcal{L})} f(\mathbf{x} + \mathbf{z}) e^{-2\pi i \langle \mathbf{x}, \mathbf{y} \rangle} d\mathbf{x} \\ &= \det(\mathcal{L}^*) \sum_{\mathbf{z} \in \mathcal{L}} \int_{\mathbf{x} \in \mathcal{P}(\mathcal{L})} f(\mathbf{x} + \mathbf{z}) e^{-2\pi i \langle \mathbf{x} + \mathbf{z}, \mathbf{y} \rangle} d\mathbf{x} \\ &= \det(\mathcal{L}^*) \int_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}) e^{-2\pi i \langle \mathbf{x}, \mathbf{y} \rangle} d\mathbf{x} \\ &= \det(\mathcal{L}^*) \hat{f}(\mathbf{y}) \end{aligned}$$

where the first equality used the definition of the Fourier series for ϕ , the second used the definition of ϕ , the third used the “niceness” of f to switch the integral and summation, the fourth used the fact that $\langle \mathbf{y}, \mathbf{z} \rangle \in \mathbb{Z}$, and the final one used the definition of the Fourier transform of f .

Now use Fourier inversion for ϕ to show that

$$\sum_{\mathbf{x} \in \mathcal{L}} f(\mathbf{x}) = \phi(\mathbf{0}) = \sum_{\mathbf{y} \in \mathcal{L}^*} \hat{\phi}(\mathbf{y}) = \det(\mathcal{L}^*) \sum_{\mathbf{y} \in \mathcal{L}^*} \hat{f}(\mathbf{y})$$

□

2.4 Smoothing Lemma and Proof

Let ϕ_s denote the distribution obtained by picking a vector from the (continuous) Gaussian distribution defined by ρ_s and reducing it modulo the parallelepiped $\mathcal{P}(\mathbf{B})$. Thus,

$$\phi_s(\mathbf{x}) = 1/s^n \cdot \sum_{\mathbf{y} \in \mathcal{L}(\mathbf{B})} \rho_s(\mathbf{x} + \mathbf{y}) := 1/s^n \cdot \rho_s(\mathbf{x} + \mathcal{L}(\mathbf{B}))$$

Now, since ϕ_s is clearly a periodic function over the lattice $\mathcal{L}(\mathbf{B})$, we can compute it alternatively using the Poisson summation formula. For any $\mathbf{x} \in \mathcal{P}(\mathbf{B})$, we have

$$\begin{aligned} \phi_s(\mathbf{x}) &= \sum_{\mathbf{y} \in \mathcal{L}^*} \hat{\phi}_s(\mathbf{y}) e^{2\pi i \langle \mathbf{x}, \mathbf{y} \rangle} \\ &= \det(\mathcal{L}^*) \cdot (1/s^n) \sum_{\mathbf{y} \in \mathcal{L}^*} \hat{\rho}_s(\mathbf{y}) e^{2\pi i \langle \mathbf{x}, \mathbf{y} \rangle} \\ &= s^n \cdot \det(\mathcal{L}^*) \cdot (1/s^n) \sum_{\mathbf{y} \in \mathcal{L}^*} \rho_{1/s}(\mathbf{y}) e^{2\pi i \langle \mathbf{x}, \mathbf{y} \rangle} \\ &= \det(\mathcal{L}^*) \cdot \left(1 + \sum_{\mathbf{y} \in \mathcal{L}^* \setminus \{\mathbf{0}\}} \rho_{1/s}(\mathbf{y}) \cdot e^{2\pi i \langle \mathbf{y}, \mathbf{x} \rangle} \right) \end{aligned}$$

where the first equality is by the definition of Fourier inversion, the second by the definition of ϕ_s and by the linearity of the Fourier transform, the third by the Fourier transform of the Gaussian function (Lemma 5), and the final one just by grouping terms together.

We will use this formulation to compute the statistical distance of ϕ_s from the uniform distribution over the parallelepiped whose density function is $U_{\mathcal{P}(\mathbf{B})}(\mathbf{x}) = 1/\det(\mathcal{L}) = \det(\mathcal{L}^*)$.

$$\begin{aligned} \Delta(\phi_s, U_{\mathcal{P}(\mathbf{B})}) &= \int_{\mathcal{P}(\mathbf{B})} |\phi_s(\mathbf{x}) - U_{\mathcal{P}(\mathbf{B})}(\mathbf{x})| d\mathbf{x} \\ &= \det(\mathcal{L}^*) \int_{\mathcal{P}(\mathbf{B})} \left| \sum_{\mathbf{y} \in \mathcal{L}^* \setminus \{\mathbf{0}\}} \rho_{1/s}(\mathbf{y}) \cdot e^{2\pi i \langle \mathbf{y}, \mathbf{x} \rangle} \right| d\mathbf{x} \\ &= \det(\mathcal{L}^*) \cdot \det(\mathcal{L}) \cdot \max_{\mathbf{x} \in \mathcal{P}(\mathbf{B})} \left| \sum_{\mathbf{y} \in \mathcal{L}^* \setminus \{\mathbf{0}\}} \rho_{1/s}(\mathbf{y}) \cdot e^{2\pi i \langle \mathbf{y}, \mathbf{x} \rangle} \right| \\ &\leq \sum_{\mathbf{y} \in \mathcal{L}^* \setminus \{\mathbf{0}\}} \rho_{1/s}(\mathbf{y}) := \rho_{1/s}(\mathcal{L}^* \setminus \{\mathbf{0}\}) \end{aligned} \tag{1}$$

In other words, we established $\rho_{1/s}(\mathcal{L}^* \setminus \{\mathbf{0}\})$ as the quantity that governs the variation (or statistical) distance between the continuous Gaussian reduced modulo $\mathcal{P}(\mathbf{B})$ and the uniform distribution over $\mathcal{P}(\mathbf{B})$. We will now bound this quantity.

Bounding the Gaussian Weight of Non-Zero (Dual) Lattice Vectors. Let us first try to build some intuition for why we should expect to bound the Gaussian weight $\rho_{1/s}(\mathcal{L}^*)$ by something close to 1. First of all, the heaviest vector is the zero vector that gets a weight of 1. Secondly, if $\lambda_1(\mathcal{L}^*) \gtrsim (1/s\sqrt{2\pi}) \cdot \omega(\sqrt{\log n})$, then the next heaviest vector has weight $e^{-\omega(\log n)}$ which is negligible in n . However, there could be exponentially many vectors of that length which could

make the collective contribution much larger. We have to balance these two effects: the fact that a large λ_1 results in the Gaussian weight of each individual non-zero lattice vector to be tiny, versus the fact that there may be exponentially many lattice vectors of a given length.

First, let us come up with a simple upper bound on the number of lattice vectors of a given length using a packing argument.

Lemma 9. *Let \mathcal{L} be a rank- n lattice. The number of lattice vectors of length at most r is at most $\left(1 + \frac{2r}{\lambda_1(\mathcal{L})}\right)^n$.*

Proof. Draw balls of radius $\lambda_1/2$ around each lattice point. These balls do not intersect. As long as the length of each such lattice point is at most r , these balls are all contained in the ball of radius $r + \lambda_1/2$ around the origin. By a volume argument, we have

$$\text{vol}_n\left(r + \frac{\lambda_1}{2}\right) \geq N_r \cdot \text{vol}_n\left(\frac{\lambda_1}{2}\right)$$

where N_r is the number of lattice vectors of length at most r . Put together, we get

$$N_r \leq \frac{\text{vol}_n\left(r + \frac{\lambda_1}{2}\right)}{\text{vol}_n\left(\frac{\lambda_1}{2}\right)} = \left(\frac{r + \frac{\lambda_1}{2}}{\frac{\lambda_1}{2}}\right)^n = \left(1 + \frac{2r}{\lambda_1(\mathcal{L})}\right)^n$$

□

We now use this to bound the sum $\sum_{\mathbf{y} \in \mathcal{L}} \rho_s(\mathbf{y})$. The proof is due to Noah Stephens-Davidowitz.

Lemma 10. *Let \mathcal{L} be a rank- n lattice. $\sum_{\mathbf{y} \in \mathcal{L}} \rho_s(\mathbf{y}) = 1 + 2^{-O(n)}$ as long as $\lambda_1 > Cs \cdot \sqrt{n/2\pi e}$ for some absolute constant $C \approx 3$.*

Proof. Using a ‘‘Lebesgue integral trick’’ (mentioned in class), we have

$$\begin{aligned} \sum_{\mathbf{y} \in \mathcal{L}} \rho_s(\mathbf{y}) &= \int_0^1 \left| \{\mathbf{y} \in \mathcal{L} : \rho_s(\mathbf{y}) \geq t\} \right| dt \\ &= \int_0^1 \left| \{\mathbf{y} \in \mathcal{L} : e^{-\pi\|\mathbf{y}\|^2/s^2} \geq t\} \right| dt \end{aligned}$$

Now, we do a change of variables $t = e^{-\pi r^2/s^2}$, we get:

$$\begin{aligned} &= \frac{2\pi}{s^2} \int_0^\infty \left| \{\mathbf{y} \in \mathcal{L} : \|\mathbf{y}\| \leq r\} \right| r e^{-\pi r^2/s^2} dr \\ &\leq \frac{2\pi}{s^2} \cdot \left(\int_0^{\lambda_1} + \int_{\lambda_1}^\infty \right) \left| \{\mathbf{y} \in \mathcal{L} : \|\mathbf{y}\| \leq r\} \right| r e^{-\pi r^2/s^2} dr \\ &\leq (1 - e^{-\pi\lambda_1^2/s^2}) + \frac{2\pi}{s^2} \int_{\lambda_1}^\infty \left| \{\mathbf{y} \in \mathcal{L} : \|\mathbf{y}\| \leq r\} \right| r e^{-\pi r^2/s^2} dr \\ &\leq 1 + \frac{2\pi}{s^2} \int_{\lambda_1}^\infty \left(\frac{3r}{\lambda_1}\right)^n r e^{-\pi r^2/s^2} dr \\ &\leq 1 + \frac{2\pi C^n}{s^2 \lambda_1^n} \int_{\lambda_1}^\infty r^{n+1} e^{-\pi r^2/s^2} dr \end{aligned}$$

where $C = 3$. After another change of variables ($w = \pi r^2/s^2$), we can bound this by

$$1 + \left(\frac{sC}{\lambda_1 \sqrt{\pi}} \right)^n \Gamma(n/2)$$

where $\Gamma(\cdot)$ is the gamma function. Applying the bound on gamma functions, we get

$$1 + \left(\frac{sC}{\lambda_1} \sqrt{\frac{n}{2\pi e}} \right)^n$$

As long as $\lambda_1 > Cs \cdot \sqrt{n/2\pi e}$, we get a sum that is exponentially close to 1. \square

Finally, applying this to our scenario, where the lattice is \mathcal{L}^* and the function is $\rho_{1/s}$, we get that the sum $\sum_{\mathbf{y} \in \mathcal{L}^*} \rho_{1/s}(\mathbf{y})$ is exponentially close to 1 as long as

$$s \geq \lambda_n(\mathcal{L}) \cdot C \cdot \sqrt{\frac{n}{2\pi e}}$$

Indeed, if s is so large, we have $\lambda_1(\mathcal{L}^*) \geq 1/\lambda_n(\mathcal{L}) \geq \frac{C}{s} \cdot \sqrt{\frac{n}{2\pi e}}$ where the first inequality is by Lemma 3.

3 Worst-case to Average-case Reduction for SIS

The reduction is due to Ajtai originally, but our presentation follows the work of Micciancio and Regev, and borrows from Regev’s lecture notes.

We first illustrate the intuition behind the worst-case to average-case reduction by showing how to reduce the approximate-SIVP problem to a variant of SIS over the torus $\mathbb{T} = \mathbb{R}/\mathbb{Z}$, $\text{SIS}_{\mathbb{T}}$. $\text{SIS}_{\mathbb{T}}$ is exactly as in SIS, except that you are given a matrix $\mathbf{A} \in \mathbb{T}^{n \times m}$ and you are asked to find a small integer linear combination that sums to zero. That is, find $\mathbf{x} \in \mathbb{Z}^m$ such that $\mathbf{A}\mathbf{x} = 0$ and $\|\mathbf{x}\|$ is “small”.

How would such a reduction look like? On the one hand, the reduction has to generate a *uniformly random* $\text{SIS}_{\mathbb{T}}$ instance from a given lattice \mathcal{L} ; therefore, the SIS instance “forgets” the lattice \mathcal{L} that was used to generate it. On the other hand, a solution to the SIS instance has to somehow be mapped back to a non-trivially short vector in \mathcal{L} . This (apparent) conundrum is common to all worst-case to average-case reductions, and the answer is that the reduction knows some information connecting the lattice to the SIS instance which, together with the SIS solution, helps it generate short vectors in \mathcal{L} .

The reduction first generates a random vector $\mathbf{v} \in \mathcal{P}(\mathbf{B})$ in the parallelepiped associated to the given basis. It does so by sampling a vector $\mathbf{x} \leftarrow \rho_s$ from the (zero-centered) Gaussian with standard deviation parameter $s \geq \eta_{\epsilon}(\mathcal{L})$, the smoothing parameter for some negligible function $\epsilon = \epsilon(n)$, and setting

$$\mathbf{v} = \mathbf{x} \pmod{\mathcal{P}(\mathbf{B})}$$

By the smoothing lemma, \mathbf{v} is (close to) random over the parallelepiped. The first column of the SIS matrix \mathbf{A} is then set to

$$\mathbf{a} = \mathbf{B}^{-1}\mathbf{v} \in \mathbb{T}^n$$

which is (close to) random over $[0, 1]^n$. Repeat this process independently m times to generate the statistically close to uniform $\text{SIS}_{\mathbb{T}}$ matrix $\mathbf{A} \in \mathbb{T}^{n \times m}$ where

$$\mathbf{A} = \mathbf{B}^{-1}\mathbf{V}$$

Call the Gaussian matrix corresponding to \mathbf{V} as \mathbf{X} . The reduction will keep \mathbf{X} to itself.

Assume now that there is a $\text{SIS}_{\mathbb{T}}$ algorithm that gives us a non-zero integer vector $\mathbf{x} \in \mathbb{Z}^m$ such that $\mathbf{A}\mathbf{x} \in \mathbb{Z}^n$. (this is what it means for $\mathbf{A}\mathbf{x}$ to be $\mathbf{0} \pmod{1}$.) Then we know that $\mathbf{B}^{-1}\mathbf{V}\mathbf{x} \in \mathbb{Z}^n$ and therefore, $\mathbf{V}\mathbf{x} \in \mathcal{L}(\mathbf{B})$ is a lattice vector. Now, since $\mathbf{X} \equiv \mathbf{V} \pmod{\mathcal{P}(\mathbf{B})}$, we know that $\mathbf{X}\mathbf{x} \in \mathcal{L}(\mathbf{B})$ is also a lattice vector.

We now argue that it is short. We know that $\|\mathbf{X}\mathbf{x}\| \approx s\|\mathbf{x}\|\sqrt{n} \approx \lambda_n\sqrt{mn}$. Here, the first equality is because each column of \mathbf{X} is a continuous Gaussian with parameter s and therefore $\mathbf{X}\mathbf{x}$ has parameter $s\|\mathbf{x}\|$ and therefore length $s\|\mathbf{x}\|\sqrt{n}$ w.h.p. The second equality is using the smoothing lemma, substituting λ_n for s upto logarithmic factors and \sqrt{m} as the norm of \mathbf{x} , assuming it is a 0-1 vector.

This seems to work, except that we are uncomfortable working with real numbers. Furthermore, it is unclear that a “random” matrix $\mathbf{A} \in \mathbb{T}^{n \times m}$ will have an SIS solution at all. We therefore discretize.

Discretization. Consider splitting each entry into a multiple of $1/q$ (for some sufficiently large value of q that we will set shortly) and an error term. That is,

$$\mathbf{A} = \mathbf{Q} + \mathbf{E} \pmod{1}$$

where $q\mathbf{Q} \in \mathbb{Z}^{n \times m}$ and $\|\mathbf{E}\|_{\infty} \leq 1/2q$.

Our first try is to feed the SIS algorithm with the matrix $q\mathbf{Q}$ which is uniformly random mod q . The adversary returns an \mathbf{x} such that $q\mathbf{Q}\mathbf{x} = \mathbf{0} \pmod{q}$. This gives us

$$\mathbf{0} = \mathbf{Q}\mathbf{x} = (\mathbf{A} - \mathbf{E})\mathbf{x} = \mathbf{B}^{-1}(\mathbf{V} - \mathbf{B}\mathbf{E})\mathbf{x} = \mathbf{B}^{-1}(\mathbf{X} - \mathbf{B}\mathbf{E})\mathbf{x} \pmod{1}$$

and therefore, $(\mathbf{X} - \mathbf{B}\mathbf{E})\mathbf{x}$ is a lattice vector. We would, in analogy to before, show that these are short lattice vectors.

$$\|\mathbf{X}\mathbf{x} - \mathbf{B}\mathbf{E}\mathbf{x}\| \leq \|\mathbf{X}\mathbf{x}\| + \|\mathbf{B}\mathbf{E}\mathbf{x}\| \leq s\|\mathbf{x}\|\sqrt{n} + \frac{\|\mathbf{x}\|_1}{q} \cdot \max_i \|\mathbf{b}_i\|_2$$

So, this does not give us short vectors, rather it reduces the length of the longest vector in the basis by a factor of $q/\|\mathbf{x}\|_1 \geq q/m$ (roughly, assuming SIS produces 0-1 vectors). So, as long as $q \gg m \approx n \log q$, we get an improvement. Repeat this iteratively many times to get to roughly $s\sqrt{mn} \approx \lambda_n\sqrt{mn} \approx \lambda_n \cdot \tilde{O}(n)$.

We are stuck at solving n -approximate SIVP given a solver for SIS. Can we improve this?

Open Problem 3.1. Show a reduction from \sqrt{n} -SIVP (or better) to average-case SIS.

In the regime of exponential reductions, we show such reductions in a recent joint work with Brakerski and Stephens-Davidowitz.

Another question is to improve the values of q for which one can show SIS average-case hard.

Open Problem 3.2. Show a reduction from approximate SIVP to SIS with modulus $q = O(1)$.

Why do we get a non-zero vector, again? There is one important issue that we overlooked. We showed that the reduction produces a short lattice vector, but why is the vector non-zero? Relatedly, when the reduction produces many shorter vectors that form a new basis to iterate on, why do we have the guarantee that we get n linearly independent vectors from the reduction?

We will now show non-zero-ness formally, but here is the intuition: we need to think of the SIS algorithm as the adversary who is trying to send us a vector \mathbf{x} which is somehow cleverly designed so that $(\mathbf{X} - \mathbf{BE})\mathbf{x}$ is the zero vector. What does the SIS algorithm see? It possibly sees $\mathbf{V} = \mathbf{X} \pmod{\mathcal{P}(\mathbf{B})}$ but (a) it never sees \mathbf{X} itself; and (b) given \mathbf{V} , there are multiple possible values of \mathbf{X} , which is a consequence of smoothing-type arguments. In other words, the adversary is trying to force $(\mathbf{X} - \mathbf{BE})\mathbf{x}$ to be $\mathbf{0}$, but it does not know what \mathbf{X} is. We then argue that information-theoretically, it cannot succeed.

We omit the formal argument, but refer the reader to [Regev's lecture notes](#) for the full proof.

3.1 Other Open Problems

Vinod finds it rather bothersome that Ajtai's reduction (and essentially every other known reduction) that demonstrates average-case hardness of SIS starts from the *SIVP* problem, rather than the more natural *SVP*. This motivates the following open problem.

Open Problem 3.3. Show a reduction from worst-case SVP to (average-case) SIS.

In fact, he would ideally like a reduction from worst-case SIS to average-case SIS, bypassing lattices altogether. Indeed, observe that solving SIS is the same as finding a short vector in a lattice (namely, the lattice $\Lambda^\perp(\mathbf{A}) := \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{A}\mathbf{x} = 0 \pmod{q}\}$). However, even viewing through these lens, what we have demonstrated is an algorithm that finds vectors of length related to λ_n , and not λ_1 . That is, if the worst-case SIS lattice has a short vector but no n linearly independent short vectors, then the reduction will miss finding the short vector (!!)

We view this as a deficiency in our understanding of SIS and worst-case to average-case reductions. Therefore, a related problem is:

Open Problem 3.4. Show a reduction from worst-case SIS to average-case SIS without going through lattices.