

Gaussian Sums

Instructor: Daniele Micciancio

UCSD CSE

Gaussian distributions and harmonic analysis play a fundamental role both in the design of lattice-based cryptographic functions, and the theoretical study of lattice problems. A typical problem that arises in lattice-based cryptography and in the study of the average-case complexity of lattice problems is to determine how much “random noise” $\mathcal{E}$ should be added to a lattice Λ , so that the resulting distribution $\Lambda + \mathcal{E}$ is (statistically close to) uniform. One can study this problem with respect to any noise distribution $\mathcal{E}$, e.g., the uniform distribution over a sufficiently large ball. Fourier analysis is a powerful technique to study this type of problems, which naturally leads to setting $\mathcal{E}$ to a gaussian distribution. Using fourier analysis, we will define a fundamental parameter associated to any lattice Λ , called the “smoothing parameter”, which precisely determines the amount of gaussian noise to be added to Λ in order to smooth out its discrete structure, and obtain the uniform distribution. We begin by developing the basics of fourier analysis in the case of finite groups, which are the most common setting in lattice cryptography. Then, we will extend the theory to functions that are periodic modulo a lattice, and prove the Poisson summation formula, a central tool in harmonic analysis. Finally, with these tools at our disposal, we will go back to our original motivation, and look at some applications in cryptography and complexity.

1 Fourier transform on finite groups

Harmonic analysis in $\mathbb{R}^n$ requires sensible assumptions on the function f to guarantee convergence and the validity of the fourier inversion formula. Fortunately, in lattice cryptography, we deal primarily with discrete probability distributions and the fourier transform over finite groups. This allows for a much simpler treatment, using linear algebra.

For any lattice Λ and full rank sublattice $\Gamma \subseteq \Lambda$, the quotient

$$\mathbf{G} = \Lambda / \Gamma = \{\mathbf{x} + \Gamma \mid \mathbf{x} \in \Lambda\}$$

is a finite (additive) group of order $|\mathbf{G}| = \det(\Gamma) / \det(\Lambda)$. (In fact, any finite additive group can be represented this way.)

Exercise 1. Show that any finitely generated additive group $\mathbf{G}$ can be represented as (i.e., is isomorphic to) the quotient Λ / Γ of two lattices, with $\Gamma \subseteq \Lambda$. Moreover, $\mathbf{G}$ is finite if and only if Λ and Γ have the same rank. [Hint: Let $\{\mathbf{g}_1, \dots, \mathbf{g}_n\}$ a set of group elements that generate $\mathbf{G} = \sum_i \mathbf{g}_i \cdot \mathbb{Z}$, and consider the lattices $\Lambda = \mathbb{Z}^n$ and $\Gamma = \{\mathbf{x} \in \mathbb{Z}^n \mid \sum_i x_i \cdot \mathbf{g}_i = 0\}$.]

We define the dual of the group $\mathbf{G}=\Lambda/\Gamma$ as the quotient

$$\hat{\mathbf{G}} = \Gamma^*/\Lambda^*$$

where Λ^* and Γ^* are the dual lattices of Λ and Γ . Since $\Gamma \subseteq \Lambda$, we also have $\Lambda^* \subseteq \Gamma^*$ and the quotient Γ^*/Λ^* is well defined. Moreover, it is easy to see that the dual group has exactly the same size as the original group:

$$|\hat{\mathbf{G}}| = \frac{\det(\Lambda^*)}{\det(\Gamma^*)} = \frac{1/\det(\Lambda)}{1/\det(\Gamma)} = \frac{\det(\Gamma)}{\det(\Lambda)} = |\mathbf{G}|.$$

Notice that a finitely generated abelian group can be represented as a quotient of two lattices $\mathbf{G}=\Lambda/\Gamma$ in many different ways, possibly using lattices Λ of different rank. (The rank of a group $\mathbf{G}$ is the smallest n such that $\mathbf{G}$ can be represented as the quotient of a rank n lattice Λ , and it equals the size of the smallest generating set $g_1, \dots, g_n$.) The following exercise shows that the additive group structure of the dual group does not depend on the choice of representation.

Exercise 2. Let $\mathbf{G}_0=\Lambda_0/\Gamma_0$ and $\mathbf{G}_1=\Lambda_1/\Gamma_1$ be isomorphic as additive groups, i.e., assume there is a bijection $\varphi: \mathbf{G}_0 \rightarrow \mathbf{G}_1$ such that $\varphi(g+g') = \varphi(g) + \varphi(g')$. Show that the dual groups $\hat{\mathbf{G}}_0 = \Gamma_0^*/\Lambda_0^*$ and $\hat{\mathbf{G}}_1 = \Gamma_1^*/\Lambda_1^*$ are also isomorphic.

In fact, for finite groups, $\mathbf{G}$ and $\hat{\mathbf{G}}$ are also isomorphic, though this is harder to see, and the same is not true in general when $\mathbf{G}=\Lambda/\Gamma$ is not finite, i.e., when Λ and Γ have different rank. For example, as we will see in the next section, the dual group of $\mathbb{Z}$ is the unit interval $\mathbb{T} = [0, 1)$ with addition modulo 1. So, for the rest of this section, we will restrict our attention to finite groups $\mathbf{G}=\Lambda/\Gamma$ defined by a full rank sublattice $\Gamma \subseteq \Lambda$.

Exercise 3. Show that the dual of the dual of a finite additive group $\mathbf{G}=\Lambda/\Gamma$ equals the original group $\hat{\hat{\mathbf{G}}} = \mathbf{G}$, giving an explicit isomorphism $\varphi: \Lambda \rightarrow \Lambda^{**}$ such that $\varphi(\Gamma) = \Gamma^{**}$.

The set of functions $V = (\mathbf{G} \rightarrow \mathbb{C})$ is a $|\mathbf{G}|$ -dimensional vector space over the field $\mathbb{C}$ of complex numbers, where each function $f: \mathbf{G} \rightarrow \mathbb{C}$ may be identified with the vector $(f(\mathbf{x}))_{\mathbf{x} \in \mathbf{G}} \in \mathbb{C}^{|\mathbf{G}|}$. This vector space has an inner product defined as

$$\langle f, g \rangle = \mathbb{E}_{\mathbf{x} \in \mathbf{G}} [f(\mathbf{x}) \cdot \overline{g(\mathbf{x})}] = \frac{1}{|\mathbf{G}|} \sum_{\mathbf{x} \in \mathbf{G}} f(\mathbf{x}) \cdot \overline{g(\mathbf{x})}$$

where $\overline{a+ib} = a - ib \in \mathbb{C}$ is the complex conjugation operation, for $a, b \in \mathbb{R}$.

We note that the division by $|\mathbf{G}|$ in the definition of $\langle f, g \rangle$ is just a scaling factor, and we could have defined the scalar product simply as the standard (unscaled) dot product between vectors $\sum_{\mathbf{x} \in \mathbf{G}} f(\mathbf{x}) \cdot \overline{g(\mathbf{x})}$. But it is customary to scale the product by $|\mathbf{G}|$, as it avoids the need to use other normalization factors elsewhere, when defining an orthonormal basis for V .

Exercise 4. Show that $\langle \cdot, \cdot \rangle$ is indeed an inner product, i.e., it satisfies the following properties: $\langle a\mathbf{x}, \mathbf{y} \rangle = a \cdot \langle \mathbf{x}, \mathbf{y} \rangle$, $\langle \mathbf{x} + \mathbf{z}, \mathbf{y} \rangle = \langle \mathbf{x}, \mathbf{y} \rangle + \langle \mathbf{z}, \mathbf{y} \rangle$, $\langle \mathbf{x}, \mathbf{y} \rangle = \overline{\langle \mathbf{y}, \mathbf{x} \rangle}$, and $\langle \mathbf{x}, \mathbf{x} \rangle \geq 0$, with equality $\langle \mathbf{x}, \mathbf{x} \rangle = 0$ if and only if $\mathbf{x} = \mathbf{0}$.

Consider the family of functions $\chi_{\mathbf{v}}: \mathbf{G} \rightarrow \mathbb{C}$ indexed by $\mathbf{v} \in \hat{\mathbf{G}}$ and mapping $\mathbf{x} \in \mathbf{G}$ to the complex number

$$\chi_{\mathbf{v}}(\mathbf{x}) = e^{2\pi i \langle \mathbf{x}, \mathbf{v} \rangle}$$

where the term $\langle \mathbf{x}, \mathbf{v} \rangle$ in the exponent is just the standard scalar product in $\mathbb{R}^n$ between any two representatives of $\mathbf{x} \in \Lambda / \Gamma$ and $\mathbf{v} \in \Gamma^* / \Lambda^*$. Notice that these functions are well defined, because for any representatives $\mathbf{x} \in \Lambda$ and $\mathbf{v} \in \Gamma^*$, the inner product

$$\langle \mathbf{x} + \Gamma, \mathbf{v} + \Lambda^* \rangle = \langle \mathbf{x}, \mathbf{v} \rangle + \langle \mathbf{x}, \Lambda^* \rangle + \langle \Gamma, \mathbf{v} \rangle + \langle \Gamma, \Lambda^* \rangle \subseteq \langle \mathbf{x}, \mathbf{v} \rangle + \mathbb{Z}$$

is well defined modulo $\mathbb{Z}$, and the function $r \mapsto e^{2\pi i r}$ is periodic modulo 1. So, for any integer $z \in \mathbb{Z}$

$$e^{2\pi i (\langle \mathbf{x}, \mathbf{v} \rangle + z)} = e^{2\pi i \langle \mathbf{x}, \mathbf{v} \rangle} \cdot e^{2\pi i z} = e^{2\pi i \langle \mathbf{x}, \mathbf{v} \rangle}.$$

The following lemma shows that this set of functions is an orthonormal basis for V .

Lemma 1. For any $\mathbf{v}, \mathbf{w} \in \hat{\mathbf{G}} = \Gamma^* / \Lambda^*$,

$$\langle \chi_{\mathbf{v}}, \chi_{\mathbf{w}} \rangle = \begin{cases} 1 & \text{if } \mathbf{v} = \mathbf{w} \\ 0 & \text{otherwise} \end{cases}$$

Proof. If $\mathbf{v} = \mathbf{w}$, then $\langle \mathbf{v} - \mathbf{w}, \mathbf{x} \rangle = \langle \mathbf{0}, \mathbf{x} \rangle = 0$, and

$$\langle \chi_{\mathbf{v}}, \chi_{\mathbf{w}} \rangle = \mathbb{E}_{\mathbf{x} \in \mathbf{G}} [e^{2\pi i \langle \mathbf{v} - \mathbf{w}, \mathbf{x} \rangle}] = \mathbb{E}_{\mathbf{x} \in \mathbf{G}} [e^0] = 1.$$

On the other hand, if $\mathbf{v} \neq \mathbf{w}$ (in $\hat{\mathbf{G}} = \Gamma^* / \Lambda^*$), then $\mathbf{u} = \mathbf{v} - \mathbf{w} \notin \Lambda^*$, and, by definition of dual lattice, there must exist a vector $\mathbf{y} \in \Lambda$ such that $\langle \mathbf{u}, \mathbf{y} \rangle \notin \mathbb{Z}$. It follows that $e^{2\pi i \langle \mathbf{u}, \mathbf{y} \rangle} \neq 1$. Let $f: \mathbf{G} \rightarrow \mathbf{G}$ be the function $f(\mathbf{x}) = \mathbf{x} + \mathbf{y}$, and notice that $f(\mathbf{G}) = \mathbf{G}$ because f is bijective, with inverse function $f^{-1}(\mathbf{z}) = \mathbf{z} - \mathbf{y}$. It follows that

$$\begin{aligned} \langle \chi_{\mathbf{v}}, \chi_{\mathbf{w}} \rangle &= \mathbb{E}_{\mathbf{x} \in \mathbf{G}} [\chi_{\mathbf{v}}(\mathbf{x}) \cdot \overline{\chi_{\mathbf{w}}(\mathbf{x})}] \\ &= \mathbb{E}_{\mathbf{x} \in \mathbf{G}} [e^{2\pi i \langle \mathbf{u}, \mathbf{x} \rangle}] \\ &= \mathbb{E}_{\mathbf{x} \in \mathbf{G}} [e^{2\pi i \langle \mathbf{u}, f(\mathbf{x}) \rangle}] \\ &= \mathbb{E}_{\mathbf{x} \in \mathbf{G}} [e^{2\pi i \langle \mathbf{u}, \mathbf{x} \rangle} \cdot e^{2\pi i \langle \mathbf{u}, \mathbf{y} \rangle}] \\ &= \langle \chi_{\mathbf{v}}, \chi_{\mathbf{w}} \rangle \cdot e^{2\pi i \langle \mathbf{u}, \mathbf{y} \rangle}. \end{aligned}$$

Since $e^{2\pi i \langle \mathbf{u}, \mathbf{y} \rangle} \neq 1$, it must be $\langle \chi_{\mathbf{v}}, \chi_{\mathbf{w}} \rangle = 0$. □

The functions $\chi_{\mathbf{v}}$ with $\mathbf{v} \in \hat{\mathbf{G}}$, are called the *characters* of the group $\mathbf{G}$, and, since there are precisely $|\hat{\mathbf{G}}| = |\mathbf{G}|$ of them, they form an orthonormal basis for $V = \mathbf{G} \rightarrow \mathbb{C}$, called the *fourier basis* of V . It follows that any function $f \in V$ can be expressed as a linear combination of characters

$$f(\mathbf{x}) = \sum_{\mathbf{v}} \langle f, \chi_{\mathbf{v}} \rangle \chi_{\mathbf{v}}(\mathbf{x}). \quad (1)$$

The coefficients $\langle f, \chi_{\mathbf{v}} \rangle$ of this linear combination are called the *fourier coefficients* of f , and the function

$$\hat{f}(\mathbf{v}) = \langle f, \chi_{\mathbf{v}} \rangle$$

mapping each $\mathbf{v} \in \hat{\mathbf{G}}$ to the corresponding coefficient is called the (discrete) *fourier transform* of f . Using the definition of $\langle f, \chi_{\mathbf{v}} \rangle$, the fourier transform can be equivalently be rewritten in the more familiar form

$$\hat{f}(\mathbf{v}) = \frac{1}{|\mathbf{G}|} \sum_{\mathbf{x} \in \mathbf{G}} f(\mathbf{x}) \cdot e^{-2\pi i \langle \mathbf{x}, \mathbf{v} \rangle} \quad (2)$$

and expressing the function f in terms of the fourier basis (1) gives the fourier inversion formula

$$f(\mathbf{x}) = \sum_{\mathbf{v} \in \hat{\mathbf{G}}} \hat{f}(\mathbf{v}) \cdot e^{2\pi i \langle \mathbf{x}, \mathbf{v} \rangle} \quad (3)$$

Exercise 5. Give special cases of the fourier transform and inversion formula for the groups $\mathbf{G} = \mathbb{Z}_q = \mathbb{Z} / (q \mathbb{Z})$, and $\mathbf{G}^n = \mathbb{Z}_q^n = \mathbb{Z}^n / (q \mathbb{Z}^n)$.

2 Periodic functions and Poisson summation formula

As we are interested in lattices, it is natural to consider functions $f: \Lambda \rightarrow \mathbb{C}$ defined over a lattice Λ , as well as functions that are periodic modulo a lattice Λ , i.e., $\varphi(\mathbf{t} + \mathbf{v}) = \varphi(\mathbf{t})$ for every $\mathbf{v} \in \Lambda$. The canonical example of a periodic function is the distance

$$\varphi(\mathbf{t}) = \inf_{\mathbf{x} \in \Lambda} \|\mathbf{t} - \mathbf{x}\|$$

of a target point $\mathbf{t}$ to the lattice Λ . Clearly, shifting the target $\mathbf{t}$ by a lattice point $\mathbf{v}$, does not affect the distance of $\mathbf{t} + \mathbf{v}$ to the lattice. So, $\varphi(\mathbf{t})$ is periodic modulo Λ . Another example of periodic function is obtained by taking an arbitrary function $f: \mathbb{R} \rightarrow \mathbb{C}$, and defining its *periodization*

$$\varphi(\mathbf{x}) = f(\mathbf{x} + \Lambda) = \sum_{\mathbf{y} \in \Lambda} f(\mathbf{x} + \mathbf{y}),$$

assuming $\lim_{\mathbf{x} \rightarrow \infty} f(\mathbf{x}) = 0$ and the summation converges.

Let $\Lambda \subset \mathbb{R}^n$ be a full rank lattice, and $\varphi: \mathbb{R}^n / \Lambda \rightarrow \mathbb{C}$ a periodic function modulo Λ . (All this can be easily extended to lattices that are not full rank by using $\text{span}(\Lambda)$ instead of $\mathbb{R}^n$.) Assume φ is a sufficiently nice function, in the sense that it can be reasonably approximated by the values of $\varphi(\mathbf{x})$ on a fine grid $\mathbf{x} \in \frac{1}{q}\Lambda$, for a large integer q . Restricting φ to $\frac{1}{q}\Lambda$ gives a function $\varphi_q: \mathbf{G} \rightarrow \mathbb{C}$ defined over a finite group $\mathbf{G}_q = (\frac{1}{q}\Lambda) / \Lambda \equiv \mathbb{Z}_q^n$ of size $|\mathbf{G}_q| = q^n$. So, we can take the (finite) fourier transform of this function

$$\hat{\varphi}_q(\mathbf{v}) = \mathbb{E}_{\mathbf{x} \in \mathbf{G}_q} [\varphi_q(\mathbf{x}) \cdot \overline{\chi_{\mathbf{v}}(\mathbf{x})}] = \frac{1}{q^n} \sum_{\mathbf{x} \in \mathbf{G}_q} \varphi_q(\mathbf{x}) e^{-2\pi i \langle \mathbf{v}, \mathbf{x} \rangle} \quad (4)$$

mapping each $\mathbf{v} \in \hat{\mathbf{G}} = \Lambda^* / (q\Lambda^*)$ to a corresponding fourier coefficient $\hat{\varphi}_q(\mathbf{v}) \in \mathbb{C}$. By the fourier inversion formula (over finite groups), for all $\mathbf{x} \in \frac{1}{q}\Lambda$ we have

$$\varphi(\mathbf{x}) = \varphi_q(\mathbf{x}) = \sum_{\mathbf{v} \in \hat{\mathbf{G}}_q} \hat{\varphi}_q(\mathbf{v}) \cdot \chi_{\mathbf{v}}(\mathbf{x}). \quad (5)$$

For any $\mathbf{v} \in \Lambda^*$, we may think of (4) as an approximation of the expectation

$$\hat{\varphi}(\mathbf{v}) = \mathbb{E}_{\mathbf{x} \in \mathbb{R}^n / \Lambda} [\varphi(\mathbf{x}) \cdot \overline{\chi_{\mathbf{v}}(\mathbf{x})}] = \frac{1}{\det(\Lambda)} \cdot \int_{\mathbf{x} \in \mathbb{R}^n / \Lambda} \varphi(\mathbf{x}) \cdot e^{-2\pi i \langle \mathbf{v}, \mathbf{x} \rangle} d\mathbf{x} \quad (6)$$

computed over the entire domain of the function. In fact, assuming the function φ is sufficiently nice (e.g., continuous, Lipschitz, etc.), the value of (6) equals the limit $\hat{\varphi}(\mathbf{v}) = \lim_{q \rightarrow \infty} \hat{\varphi}_q(\mathbf{v})$ of the finite approximations, for $q \rightarrow \infty$.

This function $\hat{\varphi}: \Lambda^* \rightarrow \mathbb{C}$ is the discrete fourier transform of the periodic function $\varphi: \mathbb{R}^n / \Lambda \rightarrow \mathbb{C}$. Notice that while $\hat{\varphi}_q(\mathbf{v})$ is defined over the finite group $\Lambda^* / (q\Lambda^*)$, and repeats periodically modulo $q\Lambda^*$, when we take the limit for $q \rightarrow \infty$ we get a (non-periodic) function $\hat{\varphi}: \Lambda^* \rightarrow \mathbb{C}$ defined over the entire lattice Λ^* . As one may expect, taking $q \rightarrow \infty$ in (5) gives the inversion formula

$$\varphi(\mathbf{x}) = \sum_{\mathbf{v} \in \Lambda^*} \hat{\varphi}(\mathbf{v}) \cdot \chi_{\mathbf{v}}(\mathbf{x}) = \sum_{\mathbf{v} \in \Lambda^*} \hat{\varphi}(\mathbf{v}) \cdot e^{2\pi i \langle \mathbf{x}, \mathbf{v} \rangle} \quad (7)$$

which is valid under appropriate assumptions on φ . We will not be concerned with the precise conditions required to prove the validity of the fourier inversion formula for periodic functions, as we will use it only for one, particularly nice function: the periodization of the gaussian function $\rho(\mathbf{x}) = e^{-\pi \|\mathbf{x}\|^2}$. However, it is useful to make some general observations about the general theory of the discrete fourier transform.

Notice that in the finite case, the domain of a function $\varphi_q: \mathbf{G} \rightarrow \mathbb{C}$ and its fourier transform $\hat{\varphi}_q: \hat{\mathbf{G}} \rightarrow \mathbb{C}$ are finite sets of the same size $|\mathbf{G}| = |\hat{\mathbf{G}}|$, and, in fact, these two sets are isomorphic as additive groups. In the case of periodic functions and the discrete fourier transform, $\varphi: (\mathbb{R}^n/\Lambda) \rightarrow \mathbb{C}$ and $\hat{\varphi}: \Lambda^* \rightarrow \mathbb{C}$ are defined over very different sets: the domain of φ is an uncountably infinite, compact set $\mathbb{R}^n/\Lambda$, while the domain of $\hat{\varphi}$ is a countable, discrete set of points Λ^* . In particular, the (infinite) group $\mathbb{R}^n/\Lambda$ and its dual Λ^* are quite different. This is a special case of the general theory of duality for topological (locally compact, abelian) groups (called “Pontryagin duality”), where the dual of a compact group is discrete, and vice versa. As an example, the dual of $\mathbb{Z}$ is the unit interval $\mathbb{T} = \mathbb{R}/\mathbb{Z} = [0, 1)$. None of this will be used in what follows. But note that this is very different from the definition of dual lattice, where the lattice dual of $\mathbb{Z}$ is $\mathbb{Z}$ itself.

We will also use the fourier transform of functions $f: \mathbb{R}^n \rightarrow \mathbb{C}$ that are not periodic. For any function $f: \mathbb{R}^n \rightarrow \mathbb{C}$ such that $\int_{\mathbf{x} \in \mathbb{R}^n} |f(\mathbf{x})| d\mathbf{x} < \infty$, the fourier transform of f is defined as

$$\hat{f}(\mathbf{v}) = \int_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}) \cdot \overline{\chi_{\mathbf{v}}(\mathbf{x})} d\mathbf{x} = \int_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}) \cdot e^{-2\pi i \langle \mathbf{v}, \mathbf{x} \rangle} d\mathbf{x}.$$

So, the fourier transform is also a function $\hat{f}: \mathbb{R}^n \rightarrow \mathbb{C}$ from the euclidean space $\mathbb{R}^n$ to the complex numbers. This Fourier transform also admits an inversion formula $f(\mathbf{x}) = \int_{\mathbf{v} \in \mathbb{R}^n} \hat{f}(\mathbf{v}) \cdot \chi_{\mathbf{v}}(\mathbf{x}) d\mathbf{v}$, but we will not need it. In fact, the only thing we will use about the real fourier transform is that the fourier transform of the gaussian function ρ is the gaussian function itself $\hat{\rho}(\mathbf{x}) = \rho(\mathbf{x})$.

Lemma 2. *The gaussian function $\rho(\mathbf{x}) = e^{-\pi \|\mathbf{x}\|^2}$ equals its own fourier transform*

$$\hat{\rho}(\mathbf{x}) = \rho(\mathbf{x})$$

and it satisfies

$$\int_{\mathbf{x}} \rho(\mathbf{x}) d\mathbf{x} = 1.$$

Proof. It is enough to prove the statement in dimension $n = 1$, as the general statement follows by

$$\begin{aligned} \hat{\rho}(\mathbf{y}) &= \int_{\mathbf{x} \in \mathbb{R}^n} \rho(\mathbf{x}) e^{-2\pi i \langle \mathbf{x}, \mathbf{y} \rangle} d\mathbf{x} \\ &= \int_{\mathbf{x} \in \mathbb{R}^n} \prod_k \rho(x_k) e^{-2\pi i x_k y_k} d\mathbf{x} \\ &= \prod_k \int_{x \in \mathbb{R}} \rho(x) e^{-2\pi i x y_k} dx \\ &= \prod_k \hat{\rho}(y_k) = \rho(\mathbf{y}). \end{aligned}$$

So, let $\rho(x) = e^{-\pi x^2}$ the one-dimensional gaussian. We compute

$$\begin{aligned}
\hat{\rho}(y) &= \int_{x \in \mathbb{R}} \rho(x) e^{-2\pi i x y} dx \\
&= \int_{x \in \mathbb{R}} e^{-\pi(x^2 + 2ixy)} dx \\
&= e^{-\pi y^2} \int_y e^{-\pi(x+iy)^2} dx \\
&= \rho(y) \int_{x \in \mathbb{R}+iy} \rho(x) dx.
\end{aligned}$$

Finally, we observe that $\int_{x \in \mathbb{R}+iy} \rho(x) dx = \int_{x \in \mathbb{R}} \rho(x) dx$ by Cauchy's theorem, and

$$\begin{aligned}
\int_{x \in \mathbb{R}} \rho(x) dx &= \sqrt{\int_{x_1 \in \mathbb{R}} \rho(x_1) dx_1 \cdot \int_{x_2 \in \mathbb{R}} \rho(x_2) dx_2} \\
&= \sqrt{\int_{\mathbf{x} \in \mathbb{R}^2} \rho(\mathbf{x}) d\mathbf{x}} = \sqrt{\int_{r=0}^{\infty} 2\pi r \rho(r) dr} = 1
\end{aligned}$$

where the last equality follows from the fact that $\rho'(r) = -2\pi r \cdot \rho(r)$. $\square$

We note that in other settings the gaussian distribution is often defined as $g(x) = e^{-\frac{1}{2}x^2}$, which is the same as ρ , but with a different scaling factor. Using $g(x)$ corresponds to normalizing the *standard deviation* $\sqrt{\int_x g(x)^2 dx} = 1$, but introduces a scaling factor when taking the fourier transform of g . As we will make extensive use of the fourier transform, in lattice-based cryptography (or, more broadly, in measure theory) it is typically preferable to use $\rho(x) = e^{-\pi x^2}$ as the “standard” gaussian function, so that $\hat{\rho} = \rho$. The standard deviation of ρ is $\sqrt{\int_{\mathbf{x} \in \mathbb{R}^n} \rho(\mathbf{x})^2 d\mathbf{x}} = \sqrt{\frac{n}{2\pi}}$. Since the probability distribution with density function proportional to $g(x) = e^{-\frac{1}{2}x^2}$ is called the *normal distribution*, we may call our function $\rho(x) = e^{-\pi x^2}$ the *abnormal distribution* and refer to the scaled standard deviation $\sqrt{2\pi} \sigma$ as the *nonstandard deviation*, and say that the abnormal distribution has nonstandard deviation 1.

We are now ready to prove a result known as Poisson summation formula.

Theorem 3. (Poisson summation formula) *Let $f: \mathbb{R}^n \rightarrow \mathbb{C}$ be a function and $\hat{f}: \mathbb{R}^n \rightarrow \mathbb{C}$ its fourier transform. For any full rank lattice $\Lambda \subset \mathbb{R}^n$ and point $\mathbf{x} \in \mathbb{R}^n$, we have*

$$f(\mathbf{x} + \Lambda) = \det(\Lambda^*) \sum_{\mathbf{z} \in \Lambda^*} \hat{f}(\mathbf{z}) e^{2\pi i \langle \mathbf{z}, \mathbf{x} \rangle}.$$

In particular, for $\mathbf{x} = \mathbf{0}$, we get

$$f(\Lambda) = \det(\Lambda^*) \hat{f}(\Lambda^*).$$

Proof. Let $f: \mathbb{R}^n \rightarrow \mathbb{C}$ and define the periodic function $\varphi(\mathbf{x}) = f(\mathbf{x} + \Lambda) = \sum_{\mathbf{y} \in \Lambda} f(\mathbf{x} + \mathbf{y})$. Notice that $\varphi(\mathbf{x}) = \varphi(\mathbf{x} + \mathbf{y})$ for any $\mathbf{y} \in \Lambda$, i.e., φ is periodic modulo Λ . The fourier series of φ is

$$\begin{aligned}
\hat{\varphi}(\mathbf{z}) &= \det(\Lambda^*) \int_{\mathbf{x} \in \mathbb{R}^n / \Lambda} \varphi(\mathbf{x}) \cdot e^{-2\pi i \langle \mathbf{x}, \mathbf{z} \rangle} d\mathbf{x} \\
&= \det(\Lambda^*) \int_{\mathbf{x} \in \mathbb{R}^n / \Lambda} \sum_{\mathbf{w} \in \Lambda} f(\mathbf{x} + \mathbf{w}) \cdot e^{-2\pi i \langle \mathbf{x}, \mathbf{z} \rangle} d\mathbf{x} \\
&= \det(\Lambda^*) \int_{\mathbf{x} \in \mathbb{R}^n / \Lambda} \sum_{\mathbf{w} \in \Lambda} f(\mathbf{x} + \mathbf{w}) \cdot e^{-2\pi i \langle \mathbf{x} + \mathbf{w}, \mathbf{z} \rangle} d\mathbf{x} \\
&= \det(\Lambda^*) \int_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}) \cdot e^{-2\pi i \langle \mathbf{x}, \mathbf{z} \rangle} d\mathbf{x} \\
&= \det(\Lambda^*) \hat{f}(\mathbf{z})
\end{aligned}$$

where we have used the fact that, for any $\mathbf{z} \in \Lambda^*$, and $\mathbf{w} \in \Lambda$, the scalar product $\langle \mathbf{w}, \mathbf{z} \rangle \in \mathbb{Z}$ is an integer, $e^{-2\pi i \langle \mathbf{x}, \mathbf{z} \rangle} = e^{-2\pi i \langle \mathbf{x} + \mathbf{w}, \mathbf{z} \rangle}$. Finally, using the fourier inversion formula, we get

$$f(\mathbf{x} + \Lambda) = \varphi(\mathbf{x}) = \sum_{\mathbf{z} \in \Lambda^*} \hat{\varphi}(\mathbf{z}) \cdot e^{2\pi i \langle \mathbf{z}, \mathbf{x} \rangle}. \quad \square$$

3 Gaussian sums

Let $\rho(\mathbf{x}) = \exp(-\pi \|\mathbf{x}\|^2)$ the gaussian function. We are interested in gaussian sums over (subsets of) lattice cosets $\rho(A) = \sum_{\mathbf{x} \in A} \rho(\mathbf{x})$, where $A \subseteq \Lambda + \mathbf{c}$. We begin by bounding the one-dimensional gaussian sum $\rho(s\mathbb{Z})$.

Lemma 4. *For any $s > 0$, we have*

$$1 + \frac{2}{\exp(\pi s^2)} \leq \rho(s\mathbb{Z}) \leq 1 + \frac{2}{\exp(\pi s^2) - 1}.$$

Proof. For the lower bound, we restrict the summation to the integers $\{-1, 0, 1\} \subset \mathbb{Z}$:

$$\rho(s\mathbb{Z}) \geq \rho(s\{-1, 0, 1\}) = 1 + 2\rho(s) = 1 + \frac{2}{\exp(\pi s^2)}.$$

For the upper bound, we extend the summation over $\sqrt{\mathbb{Z}} = \{\pm\sqrt{n}: n \in \mathbb{N}\} \supset \mathbb{Z}$:

$$\rho(s\mathbb{Z}) \leq \rho(s\sqrt{\mathbb{Z}}) = 1 + 2 \sum_{k \geq 1} \exp(-\pi s^2 k) = 1 + \frac{2}{\exp(\pi s^2) - 1}$$

where the last step uses the equality $\sum_{k \geq 1} x = (x^{-1} - 1)^{-1}$. $\square$

The above lemma provides a tight estimate of $\rho(s\mathbb{Z})$ when s is large. For small s , the following exercise provides better bounds.

Exercise 6. Show that for any s , we have $\frac{1}{s} - 1 \leq \rho(s\mathbb{Z}) \leq \frac{1}{s} + 1$. [Hint: use the fact that $\int_{x \in \mathbb{R}} \rho(x) dx = 1$ and the fact that $\rho(x)$ is monotonically decreasing in $|x|$.]

Next we use Poisson summation formula to bound gaussian sums over scaled or shifted lattices.

Lemma 5. For any lattice Λ and real $s \geq 1$,

$$\rho(\Lambda/s) \leq s^n \rho(\Lambda)$$

Proof. By the Poisson summation formula (Theorem 3), and using $\rho(s\mathbf{x}) \leq \rho(\mathbf{x})$, we get

$$\rho(\Lambda/s) = \det(s\Lambda^*) \rho(s\Lambda^*) \leq s^n \det(\Lambda^*) \rho(\Lambda^*) = s^n \rho(\Lambda). \quad \square$$

Lemma 6. For any lattice coset $\Lambda + \mathbf{u}$,

$$\rho(\Lambda + \mathbf{u}) \leq \rho(\Lambda).$$

Proof. Using Poisson summation formula (Theorem 3) and triangle inequality we get

$$\begin{aligned} \rho(\Lambda + \mathbf{u}) &= |\det(\Lambda^*) \sum_{\mathbf{y} \in \Lambda^*} \rho(\mathbf{y}) \cdot \exp(2\pi i \langle \mathbf{y}, \mathbf{u} \rangle)| \\ &\leq \det(\Lambda^*) \sum_{\mathbf{y} \in \Lambda^*} \rho(\mathbf{y}) = \rho(\Lambda). \end{aligned} \quad \square$$

The next lemma bounds gaussian sums over a lattice coset when restricted to a halfspace.

Lemma 7. For any vector $\mathbf{h}$, let H be the set of all points closer to the origin than to $\mathbf{h}$. For any lattice coset $\Lambda + \mathbf{c}$,

$$\rho((\Lambda + \mathbf{c}) \setminus H) \leq \rho(\mathbf{h}/2) \cdot \rho(\Lambda + \mathbf{c} - \mathbf{h}/2).$$

Proof. If $I_H(\mathbf{x})$ the indicator function of H , then we have

$$\begin{aligned}
\rho((\Lambda + \mathbf{c}) \setminus H) &= \sum_{\mathbf{x} \in \Lambda + \mathbf{c}} \rho(\mathbf{x}) \cdot (1 - I_H(\mathbf{x})) \\
&\leq \sum_{\mathbf{x} \in \Lambda + \mathbf{c}} \rho(\mathbf{x}) \cdot \frac{\exp(2\pi \langle \mathbf{x}, \mathbf{h}/2 \rangle)}{\exp(2\pi \|\mathbf{h}/2\|^2)} \\
&= \rho(\mathbf{h}/2) \sum_{\mathbf{x} \in \Lambda + \mathbf{c}} \rho(\mathbf{x} - \mathbf{h}/2) \\
&= \rho(\mathbf{h}/2) \rho(\Lambda + \mathbf{c} - \mathbf{h}/2).
\end{aligned}$$

□

The last bound is proved similarly to Lemma 7, but using the Poisson summation formula to evaluate a gaussian sum over a lattice coset restricted to the points outside a ball of radius $O(\sqrt{n})$.

Theorem 8. For any lattice coset $\Lambda + \mathbf{c}$ and $\alpha \geq 1$,

$$\rho((\Lambda + \mathbf{c}) \setminus \mathcal{B}(\alpha \sqrt{n/(2\pi)})) \leq \left(\frac{\alpha^2}{\exp(\alpha^2 - 1)} \right)^{n/2} \rho(\Lambda)$$

where $\mathcal{B}(r) = \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\| \leq r\}$.

Proof. If $I_r(\mathbf{x})$ is the characteristic function of $B(r)$, then for any $0 < t < \pi$ and $r = \alpha \sqrt{n/(2\pi)}$ we have

$$\begin{aligned}
\rho((\Lambda + \mathbf{c}) \setminus B(r)) &= \sum_{\mathbf{x} \in \Lambda + \mathbf{c}} \rho(\mathbf{x}) (1 - I_r(\mathbf{x})) \\
&\leq \sum_{\mathbf{x} \in \Lambda + \mathbf{c}} \rho(\mathbf{x}) \exp(t(\|\mathbf{x}\|^2 - r^2)) \\
&= \exp(-tr^2) \rho(\sqrt{1 - t/\pi} \cdot (\Lambda + \mathbf{c})) \\
&\leq \exp(-tr^2) \rho(\sqrt{1 - t/\pi} \cdot (\Lambda)) \\
&\leq \exp(-tr^2) \cdot (1 - t/\pi)^{-n/2} \cdot \rho(\Lambda)
\end{aligned}$$

This function is minimized at $t = \pi - n/(2r^2) \geq 0$, which gives the bound in the theorem. □

Notice that the base $\alpha^2/\exp(\alpha^2 - 1)$ of the exponential factor in the above theorem is monotonically decreasing for $\alpha \geq 1$, and equals $\alpha^2/\exp(\alpha^2 - 1) = 1$ when $\alpha = 1$.

4 The smoothing parameter

Informally, the smoothing parameter is the amount of (gaussian) noise that needs to be added to a lattice to obtain a uniformly distributed point in space. Technically, the smoothing parameter is defined using gaussian sums over the dual lattice.

Definition 9. For any lattice Λ , the ϵ -smoothing parameter of the lattice $\eta_\epsilon(\Lambda)$ is the smallest real $s > 0$ such that $\rho(s\Lambda^*) \leq 1 + \epsilon$, i.e., the gaussian sum over $s\Lambda^*$ is concentrated on the origin.

The next theorem shows that this technical definition captures the informal intuition.

Theorem 10. For any lattice coset $\Lambda + \mathbf{u}$, if $\eta_\epsilon(\Lambda) \leq 1$ then

$$\rho(\Lambda + \mathbf{u}) \in [1 \pm \epsilon] \cdot \det(\Lambda^*).$$

Proof. Assume $\eta_\epsilon(\Lambda) \leq 1$, or, equivalently, $\rho(\Lambda^*) \leq 1 + \epsilon$. Then

$$\begin{aligned} |\rho(\Lambda + \mathbf{u}) - \det(\Lambda^*)| &= \det(\Lambda^*) \left| \sum_{\mathbf{y} \in \Lambda^*} \rho(\mathbf{y}) \cdot \exp(2\pi i \langle \mathbf{y}, \mathbf{u} \rangle) - 1 \right| \\ &= \det(\Lambda^*) \left| \sum_{\mathbf{y} \in \Lambda^* \setminus \{\mathbf{0}\}} \rho(\mathbf{y}) \cdot \exp(2\pi i \langle \mathbf{y}, \mathbf{u} \rangle) \right| \\ &\leq \det(\Lambda^*) \sum_{\mathbf{y} \in \Lambda^* \setminus \{\mathbf{0}\}} \rho(\mathbf{y}) \\ &= \det(\Lambda^*) \cdot (\rho(\Lambda^*) - 1) \leq \epsilon \cdot \det(\Lambda^*). \end{aligned}$$

□

It is clear from the definition that for any lattice Λ and scalar $c > 0$, $\eta_\epsilon(c\Lambda) = c\eta_\epsilon(\Lambda)$. Next, we turn to evaluating the smoothing parameter of a lattice. We begin with the integer lattice.

Lemma 11. For any $\epsilon > 0$, we have

$$\sqrt{\frac{\ln(2/\epsilon)}{\pi}} \leq \eta_\epsilon(\mathbb{Z}) \leq \sqrt{\frac{\ln(1+2/\epsilon)}{\pi}}.$$

Proof. We use Lemma 4 to estimate the Gaussian sum $\rho(s\mathbb{Z})$. Setting the bound to $1 + \epsilon$ and solving for s gives upper and lower bounds on $\eta_\epsilon(\mathbb{Z})$. Similarly, one can bound $\eta_\epsilon(\mathbb{Z})$ using Exercise 6 instead of Lemma 4. □

In order to bound the smoothing parameter of arbitrary lattices, we look at how the smoothing parameter interacts with orthogonalization. Notice that for any mutually orthogonal lattice $\langle \Lambda_1, \Lambda_2 \rangle =$, the dual of the sum $(\Lambda_1 + \Lambda_2)^*$ equals the sum of the duals $\Lambda_1^* + \Lambda_2^*$, and therefore

$$\rho((\Lambda_1 + \Lambda_2)^*) = \rho(\Lambda_1^* + \Lambda_2^*) = \rho(\Lambda_1^*) \rho(\Lambda_2^*).$$

So, if $s = \eta_{\epsilon_1}(\Lambda_1) = \eta_{\epsilon_2}(\Lambda_2)$, then $s = \eta_\epsilon(\Lambda_1 + \Lambda_2)$ for $\epsilon = \epsilon_1 + \epsilon_2 + \epsilon_1 \epsilon_2$. This already gives a way to bound the smoothing parameter of lattices that have an orthogonal basis. But most lattices do not have such a basis, so we need one more tool. For the general case, we use the fact that orthogonalizing a lattice can only increase its smoothing parameter.

Theorem 12. *For any lattice basis $\mathbf{B}$ with Gram-Schmidt orthogonalization $\mathbf{B}^*$ and $\epsilon > 0$,*

$$\eta_\epsilon(\mathcal{L}(\mathbf{B})) \leq \eta_\epsilon(\mathcal{L}(\mathbf{B}^*)).$$

The theorem is proved using Lemma 6, by induction on the dimension, using the relation between the orthogonalization of $\mathbf{B}$ and the orthogonalization (in reverse order) of its dual basis.

This theorem allows to bound the smoothing parameter in terms of $\tilde{\lambda}_n$.

Corollary 13. *For any n -dimensional lattice Λ and $\epsilon > 0$, $\eta_\epsilon(\Lambda) \leq \eta_\epsilon(\mathbb{Z}^n) \cdot \tilde{\lambda}_n(\Lambda)$.*

Proof. Let $\mathbf{B}$ be a basis of Λ achieving $\tilde{\lambda}_n$. Then, we have

$$\eta_\epsilon(\Lambda) \leq \eta_\epsilon(\mathbf{B}^* \mathbb{Z}^n) \leq \eta_\epsilon(\tilde{\lambda}_n \cdot \mathbb{Z}^n) = \eta_\epsilon(\mathbb{Z}^n) \cdot \tilde{\lambda}_n. \quad \square$$

It is also possible to show that λ_n is never much bigger than the smoothing parameter. In fact, one can use the smoothing parameter to bound the covering radius of a lattice μ . An upper bound on λ_n follows using the inequality $\lambda_n \leq 2\mu$.

Theorem 14. *For any $\alpha > 1$, $\epsilon > 0$, dimension n such that*

$$\left(\frac{\alpha^2}{\exp(\alpha^2 - 1)} \right)^{n/2} \leq \frac{1 - \epsilon}{1 + \epsilon},$$

the covering radius of any n -dimensional lattice Λ satisfies

$$\mu(\Lambda) \leq \alpha \sqrt{\frac{n}{2\pi}} \eta_\epsilon(\Lambda).$$

In particular, for any $\alpha > 1$ there is an $\epsilon > 0$ such that this is true for any n .

Proof. Assume without loss of generality that $\eta_\epsilon(\Lambda) = 1$. We want to prove that $\mu \leq \alpha \sqrt{\frac{n}{2\pi}}$. Assume for contradiction that this is not true, and let $\mathbf{h}$ be a deep hole of the lattice, i.e., a point in space at distance larger than $r = \alpha \sqrt{\frac{n}{2\pi}}$ from any lattice point. It follows that the lattice coset $\Lambda - \mathbf{h}$ has no points inside a ball of radius r . By Theorem 8,

$$\rho(\Lambda - \mathbf{h}) = \rho((\Lambda - \mathbf{h}) \setminus \mathcal{B}(r)) < \beta^{n/2} \rho(\Lambda)$$

where $\beta = \alpha^2 / \exp(\alpha^2 - 1) < 1$. But from Theorem 10 we also have $\rho(\Lambda - \mathbf{h}) \geq (1 - \epsilon) \det(\Lambda^*)$ and $\rho(\Lambda) \leq (1 + \epsilon) \det(\Lambda^*)$. It follows that

$$\beta^{n/2} > \frac{\rho(\Lambda - \mathbf{h})}{\rho(\Lambda)} \geq \frac{1 - \epsilon}{1 + \epsilon},$$

contradicting the assumption in the first part of the theorem.

The second part follows from the fact that $\beta < 1$. So, $\beta^{n/2}$ is maximized at $n = 1$, and one can choose, for example, $\epsilon = (1 - \sqrt{\beta})/2$. $\square$

Similarly, one can relate the smoothing parameter and the minimum distance of the dual lattice.

Theorem 15. *For any $\alpha > 1$, $\epsilon > 0$, and dimension n such that*

$$\left(\frac{\alpha^2}{e^{\alpha^2-1}} \right)^{n/2} \leq \frac{\epsilon}{1+\epsilon},$$

the smoothing parameter of any n -dimensional lattice Λ satisfies

$$\lambda_1(\Lambda^*) \cdot \eta_\epsilon(\Lambda) \leq \alpha \sqrt{\frac{n}{2\pi}}.$$

In particular, this is true for any $\alpha > 1$, $\epsilon > 0$ and sufficiently high dimension n .

Proof. Assume without loss of generality that $\eta_\epsilon = 1$. We need to prove that $\lambda_1(\Lambda^*)$ is at most $r = \alpha \sqrt{n/2\pi}$. Assume for contradiction that this is not the case. Then, the only dual lattice vector within a ball of radius r is the origin, and by Theorem 8 we have

$$\rho(\Lambda^* \setminus \{\mathbf{0}\}) = \rho(\Lambda^* \setminus \mathcal{B}(r)) \leq \beta^{n/2} \rho(\Lambda^*)$$

where $\beta = \alpha^2 / \exp(\alpha^2 - 1) < 1$. Substituting $\rho(\Lambda^*) = 1 + \epsilon$ and $\rho(\Lambda^* \setminus \{\mathbf{0}\}) = \epsilon$, we get $\epsilon \leq \beta^{n/2} (1 + \epsilon)$, a contradiction. $\square$

Combining the last two bounds we obtain Banaszczyk's transference theorem.

Theorem 16. *For any $c > 1/(2\pi)$ and sufficiently high dimension n , any n -dimensional lattice Λ satisfies*

$$\mu(\Lambda) \cdot \lambda_1(\Lambda^*) \leq c \cdot n.$$

Moreover, the smallest c such that this is true for any dimension n is precisely $c = 1/2$.

Proof. Let $\alpha = \sqrt{2\pi c} > 1$, $\beta = \alpha^2 / \exp(\alpha^2 - 1)$, and $\epsilon = 1/2$. Notice that $\beta < 1$. So, $\lim_{n \rightarrow \infty} \beta^{n/2} = 0$ and for all n large enough, we have

$$\beta^{n/2} \leq \frac{1}{3} = \frac{1-\epsilon}{1+\epsilon} = \frac{\epsilon}{1+\epsilon}$$

and the hypotheses of Theorem 14 and Theorem 15 are satisfied. Using the bounds in those two theorems, we see that the product of $\mu(\Lambda)$ and $\lambda_1(\Lambda^*) \eta_\epsilon(\Lambda)$ is at most $\alpha^2 \frac{n}{2\pi} \eta_\epsilon$. Dividing by η_ϵ , we get $\lambda_1^* \cdot \mu \leq \alpha^2 n / (2\pi) = c \cdot n$.

For the second part, notice that when $n = 1$, we may assume $\Lambda = \mathbb{Z}$, which satisfies

$$\mu(\Lambda) \lambda_1(\Lambda^*) = \mu(\mathbb{Z}) \lambda_1(\mathbb{Z}) = \frac{1}{2} \cdot 1 = \frac{n}{2}.$$

For other values of n , notice that when $c = 1/2$, we have $\alpha = \sqrt{\pi}$, $\beta = \pi / \exp(\pi - 1)$. So, for $n \geq 3$, we have

$$\beta^{n/2} \leq \beta^{3/2} \approx 0.22 \leq \frac{1}{3}.$$

The case $n = 2$ is proved in the following exercise. □

Exercise 7. Show that for any 2-dimensional lattice Λ , $\mu(\Lambda) \cdot \lambda_1(\Lambda^*) \leq \sqrt{(1/4) + (4/\pi^2)} < 1$. [Hint: Assume $\det(\Lambda) = 1$, and show that Λ has a basis $[\mathbf{b}_1, \mathbf{b}_2]$ such that $\|\mathbf{b}_1\| = \lambda_1^*$ and $\|\mathbf{b}_2\| = 1/\lambda_1^*$. Bound μ using this basis, and λ_1^* using Minkowski's convex body theorem.]

Using the inequalities $\lambda_n \leq 2\mu$ and $\lambda_n \lambda_1^* \geq 1$, we see that for any n -dimensional lattice we have

$$1 \leq \lambda_n(\Lambda) \lambda_1(\Lambda^*) \leq n.$$

This is called a “transference” bound because it allows to use minima λ_1, λ_n of a lattice to estimate the minima λ_1^*, λ_n^* of its dual, within a factor n .

Similarly, one can use Theorem 15 and the following lemma to connect the minimum dual distance and the smoothing parameter

$$\eta_\epsilon(\mathbb{Z}) \leq \eta_\epsilon(\Lambda) \lambda_1(\Lambda^*) \leq \sqrt{\frac{n}{2\pi}}$$

within a smaller factor $O(\sqrt{n})$.

Lemma 17. For all $\epsilon > 0$, and n -dimensional lattice Λ ,

$$\eta_\epsilon(\Lambda) \geq \frac{\eta_\epsilon(\mathbb{Z})}{\lambda_1(\Lambda^*)}.$$

Proof. Let $\eta = \eta_\epsilon(\Lambda)$ and $\mathbf{v} \in \Lambda^*$ any nonzero vector in the dual lattice. Since $\mathbf{v} \cdot \mathbb{Z} \subseteq \Lambda$, we have

$$\begin{aligned} 1 + \epsilon &= \rho(\eta \cdot \Lambda^*) \\ &\geq \rho(\eta \cdot \mathbf{v} \cdot \mathbb{Z}) \\ &= \rho(\eta \cdot \|\mathbf{v}\| \cdot \mathbb{Z}). \end{aligned}$$

So, by definition of smoothing parameter, it must be $\eta \cdot \|\mathbf{v}\| \geq \eta_\epsilon(\mathbb{Z})$. Choosing a vector of length $\|\mathbf{v}\| = \lambda_1(\Lambda^*)$, gives $\eta_\epsilon(\Lambda) \geq \eta_\epsilon(\mathbb{Z}) / \lambda_1(\Lambda^*)$. □

All the bounds we proved on the parameters of a lattice can be summarized with the following chain of inequalities

$$\sqrt{\frac{2\pi}{n}} \eta_\epsilon \leq \frac{1}{\lambda_1^*} \leq \lambda_n \leq 2\mu \leq \sqrt{\frac{2n}{\pi}} \eta_\epsilon.$$

where the first and last one hold asymptotically.

5 The Discrete Gaussian distribution

We know from Lemma 2 that $\int_{\mathbf{x}} \rho(\mathbf{x}) d\mathbf{x} = 1$. So, the gaussian function can be interpreted as a probability distribution over $\mathbb{R}^n$. Lattice cryptography uses a discrete version of this probability distribution which selects points $\mathbf{x} \in \Lambda$ from a lattice with probability proportional to $\rho(\mathbf{x})$.

Definition 18. *For any lattice Λ , the discrete gaussian probability distribution D_Λ is the probability distribution over Λ that selects each lattice point $\mathbf{x} \in \Lambda$ with probability $\Pr\{\mathbf{x}\} = \rho(\mathbf{x}) / \rho(\Lambda)$, where $\rho(\Lambda) = \sum_{\mathbf{x} \in \Lambda} \rho(\mathbf{x})$ is a normalization factor.*

Informally (or even formally, with some effort) one can think of the discrete gaussian distribution D_Λ as the conditional distribution of the continuous gaussian distribution $\Pr\{\mathbf{x}\} = \rho(\mathbf{x})$ on $\mathbb{R}^n$, conditioned on the event that $\mathbf{x} \in \Lambda$ is a lattice point. The reason we state this informally is that the event of picking a lattice point $\mathbf{x} \in \Lambda$ when choosing $\mathbf{x}$ according to a continuous probability distribution has zero probability. So, giving a formal definition of conditional distribution is a bit tricky. Luckily this is not needed, as we will work only with discrete distributions, and use Definition 18 directly.

An immediate consequence of Lemma 6 is that for any lattices $\Gamma \subset \Lambda$, the gaussian probability distribution $(D_\Lambda \bmod \Gamma)$ over the quotient group Λ/Γ is maximized at $0(\bmod \Gamma)$.

Using these bounds we can easily prove several tail inequalities on the norm of vectors chosen according to a discrete Gaussian distribution.

Using a union bound over all standard unit vectors $\pm \mathbf{e}_i$, gives a tail inequality on the infinity norm of a vector chosen according to the discrete gaussian distribution from an n -dimensional lattice.

Corollary 19. *For any n -dimensional lattice Λ , if $x \leftarrow D_\Lambda$ then*

$$\Pr \leq 2n \exp(-\pi t^2).$$

As a special case, using the scaled integer lattice $\Lambda = \mathbb{Z}/s$, we get a tail bound for gaussian samples from $\mathbb{Z}$.

Corollary 20. *If $x \leftarrow D_{\mathbb{Z},s}$, then $\Pr\{|x| \geq st\} \leq 2 \exp(-\pi t^2)$.*

Corollary 21. *For any $\alpha \geq 1$, if $\mathbf{x} \leftarrow D_\Lambda$ then*

$$\Pr \leq \left(\frac{\alpha^2}{\exp(\alpha^2 - 1)} \right)^{n/2}.$$

Also this corollary can be used to bound the probability that $x \leftarrow D_{\mathbb{Z},s}$ is bigger than $|x| > st$.