Solving SIS in any norm via Gaussian sampling

Maiara F. Bollauf[1], Amaury Pouly[2] and Yixin Shen[2]

[1] University of Tartu, Estonia
[2] University of Rennes, Inria, CNRS, IRISA, Rennes, France

Abstract. The Short Integer Solution (SIS) problem plays an important role in lattice-based cryptography. In this paper, we construct a natural and simple provable algorithm that solves the SIS problem for any norm when the norm bound ℓ is less than half the modulus q. The algorithm consists of using a discrete Gaussian sampler on the SIS q-ary lattice to obtain many lattice vectors, and requires estimating the probability that one of them is nonzero and falls into a ball of radius ℓ in the given norm. For the latter, we improve upon previous analysis of a random q-ary lattice by obtaining tight bounds on the expected value and variance of the Gaussian mass of the entire lattice and of an ℓ_p-norm ball, for any $p \in (0, \infty]$. These bounds require new technical results on the discrete Gaussian distribution and on the ratio of two Gaussian mass functions of $\mathbb{Z}$. This allows us to show that the proposed algorithm is provably correct. When instantiated with a Markov chain Monte Carlo (MCMC)-based discrete Gaussian sampler, the complexity of the algorithm can be estimated precisely. In particular, we show that its performance is at least 50 bits faster than the version of Wagner's algorithm of Ducas, Engelberts, and Loyer (Crypto 2025) for solving SIS with the infinity norm across all security levels of Dilithium, but still far from heuristic estimates underlying the security of this scheme.

1 Introduction

The *Short Integer Solution* problem (SIS) is a central computational problem in lattice-based cryptography. It asks to find a short nonzero lattice vector in a randomly generated q-ary lattice, with respect to a given norm. The standard case of the ℓ_2-norm was first introduced by Ajtai [Ajt96] and enjoys a worst-case-to-average-case reduction in certain parameter regimes [Ajt96, MR04, GPV08] . Several follow-up works have further investigated its hardness in specific parameter regimes. For instance, the cryptanalysis of SIS with small parameters was first studied in [MP13], and more recently, its inhomogeneous version (ISIS) for small modulus q was discussed in [DEP23], together with potential applications to Falcon and Mitaka [EFG+22].

SIS has numerous applications in cryptography, including the construction of collision-resistant hash functions, digital signatures, identification schemes, and, more recently, zero-knowledge proofs. In practice, designers of cryptographic schemes are tempted to use SIS instances that are outside the asymptotic coverage of the worst-case to average-case reduction regime in order to gain efficiency. In many cases, they also consider variants of the problem. This is, for example, the case of the recently NIST standardized signature scheme Dilithium [DKL+18], whose hardness relies on the hardness of the SIS problem under the ℓ_∞-norm.

There are only a few studies on the harness of SIS under the ℓ_∞-norm before this work. An efficient quantum polynomial-time algorithm was presented in [CLZ22] when the norm

E-mail: maiara.bollauf@ut.ee (Maiara F. Bollauf), amaury.pouly@cnrs.fr (Amaury Pouly), yixin.shen@inria.fr (Yixin Shen)

This work is licensed under a "CC BY 4.0" license.
Date of this document: 2026-05-22.

bound is very close to $q/2$. This algorithm was later dequantized in [KOW26]. In [CLZ22], the authors also mentioned a classical polynomial-time algorithm by Regev for the case where the modulus q is composite and m is very large.

Recently, a randomized variant of Wagner's algorithm [Wag02], interpreted as an approximate discrete Gaussian sampler over q-ary lattices, was proposed in [DEL25] to (provably) solve the ℓ_∞-norm variant of SIS. Its complexity was shown to be subexponential in an extreme regime outside cryptographic interest. This line of work was motivated by an analogous technique for solving LWE, and the authors observe that their resulting algorithm has no practical impact on the concrete security of Dilithium.

A common feature underlying the reductions in [MR04], the randomized algorithm of [DEL25] for solving SIS_∞, and other fundamental constructions in lattice-based cryptography [GPV08, FHK+19] is the *discrete Gaussian sampling*. This technique enables the random selection of lattice points from a discrete Gaussian distribution and serves as the core of the general SIS solver proposed in this work.

1.1 Contributions

We provide a provable algorithm to solve the SIS problem under arbitrary norms by applying discrete Gaussian sampling over the implicit random lattice, in the case where the norm bound ℓ is smaller than half the modulus q. To achieve this goal, we derive several bounds on the expected value and variance of the Gaussian mass over random lattices and over ℓ_p-balls, which may be of independent interest in lattice theory. We compare the runtime of our algorithm with that in [DEL25], showing large improvements, but not enough to pose a threat to the security of Dilithium. More precisely, our contributions are as follows:

1. We use a discrete Gaussian sampler on the SIS q-ary lattice to sample lattice vectors and we estimate the probability that one of them is nonzero and falls into a ball of radius ℓ in the given norm. If this is the case, then the sampled vector will be a solution of the SIS problem.

2. As a consequence of the probabilistic analysis of our algorithm, we derive several theoretical bounds of independent interest. In particular, we obtain tight bounds on the expected value and variance of the Gaussian mass of an ℓ_p-norm ball, for any $p \in (0, \infty]$ (Lemma 10), as well as for the entire lattice (Lemma 8, Corollary 2). Additionally, we establish a probabilistic bound on the smoothing parameter of a random q-ary lattice.

3. Finally, we estimate the complexity of this SIS solver using the Markov chain Monte Carlo (MCMC)-based discrete Gaussian sampler proposed by Wang and Ling [WL19], and conclude about its superior performance when compared to [DEL25] in certain parameter regimes, but which still remains far from the best heuristic attacks. More details on the comparison with other provable and heuristic attacks on SIS can be found in Section 5.2.

1.2 Technical Overview

Sampling points inside a ball of radius ℓ. In a nutshell, our norm-invariant algorithm for solving SIS can be explained as follows: we use a discrete Gaussian sampler over the SIS q-ary lattice to generate lattice vectors and estimate the probability that a nonzero sample lies within a ball of radius ℓ in the given norm. In this case, the sampled vector is a solution to the SIS problem. Our algorithm is then demonstrated to be correct and its complexity can be estimated precisely for any norm ℓ_p according to the following (informal) statement.

Theorem 1 (Informal, see Theorem 5 and Remark 9)**.** *There is a randomized algorithm that for every $p \in [1,\infty]$, $d \in \mathbb{N}$, $1 \leqslant k \leqslant d$, odd prime number q, $s > 0$ and $0 \leqslant \ell < q/2$, solves $\mathrm{SIS}^{p}_{d-k,d,q,\ell}$ on a fraction at least $1 - \mathsf{negl}(d)$ of lattices. This algorithm requires oracle access to a discrete Gaussian sampler at width s, makes $N = {\rho_s(B_d^{p,\mathbb{Z}}(\ell)\setminus\{0\})}/{\rho_s(\mathbb{Z}^d)}$ queries to the sampler and runs in time polynomial in N and d.*

Asymptotic versions of the complexity for $p = 2$ (Theorem 7) and $p = \infty$ (Theorem 6) turn out to be considerably simpler.

Bounds on key measures for random lattices. A central point in the analysis of the algorithm is to estimate the probability that a discrete Gaussian sample falls into a ball with a given norm (Corollary 4). This, in turn, requires a more thorough analysis of key measures for random lattices, such as the Gaussian mass and the smoothing parameter.

We improve upon the previous analysis of a random q-ary lattice by obtaining tight bounds on the expected value and variance of the Gaussian mass of an ℓ_p-norm of a ball, for any $p \in (0,\infty]$ (Lemma 10). We also prove similar bounds for the mass of the entire lattice (Lemma 8 and Corollary 2), in addition to a fundamental result about the quotient of two Gaussian mass functions of the $\mathbb{Z}$ lattice, i.e., we show that $\rho_{\gamma s}(\mathbb{Z})/\rho_s(\mathbb{Z})$ is strictly increasing for $\gamma > 1$.

This allow us to derive the following probabilistic bound on the smoothing parameter of a random q-ary lattice, which may be of independent interest.

Corollary 1 (Informal, see Corollary 3)**.** *Fix k, q, ε and assume that both $q^{1-k/d}\varepsilon^{2/d}$ and $q^{k/d}\varepsilon^{-2/d}$ grow faster than $\sqrt{\log d}$. Then, for a random lattice $L \sim \mathcal{L}_{d,k,q}$, the smoothing parameter $\eta_\varepsilon(L)$ is at most s_ε with probability $1-2^{-d/2}$, where $s_\varepsilon = q^{1-k/d}\left({(2\varepsilon+q+\sqrt{4\varepsilon q+q^2})}/{2\varepsilon^2}\right)^{1/d}$.*

Although our bounds on the variance of the Gaussian mass are essentially optimal (Remark 5 and Remark 8), this does not translate into an optimal bound for the smoothing parameter for very small values of ε (Remark 7). This suggests that more work is needed to better understand the smoothing parameter of random q-ary lattices, for example, by computing the higher moments of the Gaussian mass.

Complexity estimates and comparison. To assess the performance of the SIS solver proposed in Theorem 1, access to a generic discrete Gaussian sampler is required. For concrete applications, we use the Markov chain Monte Carlo (MCMC)-based discrete Gaussian sampler proposed by Wang and Ling [WL19] whose complexity was analyzed in [PS25] when the basis of the lattice is obtained via the BKZ algorithm, and under the Geometric Series Assumption (GSA). Notice that optimized samplers may be available for structured lattices and eventual restrictions on the choices of s [Kle00, Pei10, EWY23, BLL25], which would consequently imply improvements for the SIS solver.

We evaluate our algorithm on the parameters of Dilithium (see Section 5.2) and verify that it does not compromise its security. We then provide some comparisons with provable and heuristic algorithms.

2 Preliminaries

2.1 Notation

We denote vectors and matrices in boldface. We denote by $\mathbf{x}^T$ the transpose of the (column) vector $\mathbf{x}$, which is therefore a row vector. For any vector $\mathbf{x} \in \mathbb{R}^n$, we denote by $\|\mathbf{x}\|$ its Euclidean norm and by $\|\mathbf{x}\|_p$ its p-norm for any $p \in [1,\infty]$. For any* $p \in [1,\infty]$, $n \geqslant 1$

*Using a slight abuse of notation, we mean either a real $p \geqslant 1$ or $p = \infty$.

and $r > 0$, we denote by B_n^p the open p-norm ball of radius 1 in $\mathbb{R}^n$, and by $B_n^p(r)$ the open p-norm ball of radius r. Finally, we let $B_n^{p,\mathbb{Z}}(\ell) = B_n^p(\ell) \cap \mathbb{Z}^n$. For $p = 2$, we omit the subscript and just write B_n, $B_n(r)$ and $B_n^{\mathbb{Z}}(r)$.

For any finite set X, we denote by $\mathcal{U}(X)$ the uniform distribution over X. As usual, if P and Q are two probability distributions over X and Y respectively, we denote by PQ the product distribution over $X \times Y$. For any two distributions P and Q, we denote by $\mathrm{d_{TV}}(P, Q)$ the statistical distance (or total variation distance) between P and Q. We say that two distributions P and Q are ε-close if $\mathrm{d_{TV}}(P, Q) \leqslant \varepsilon$.

We denote by erf the error function, defined for any $x \geqslant 0$ by $\mathrm{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} \,\mathrm{d}\, t$.

We will use the standard fact that for any $d \in \mathbb{N}$ and $r \geqslant 0$,

$$\mathrm{vol}(B_d(r - \tfrac{1}{2}\sqrt{d})) \leqslant \left|\left\{\mathbf{x} \in \mathbb{Z}^d : \|\mathbf{x}\| \leqslant r\right\}\right| \leqslant \mathrm{vol}(B_d(r + \tfrac{1}{2}\sqrt{d})) \tag{1}$$

where the first inequality holds when $r \geqslant \frac{1}{2}\sqrt{d}$. Recall that this can be obtained by placing a unit open hypercube $(-\frac{1}{2}, \frac{1}{2})^d$ centered on each point of the set on the left-hand side and arguing that those disjoint hypercubes all belong to the ball $B_d(r + \frac{1}{2}\sqrt{d})$ and contain the ball $B_d(r - \frac{1}{2}\sqrt{d})$. These estimates are asymptotically tight when $r = \omega(\sqrt{d})$.

2.2 Lattices

A *lattice* $\mathcal{L}$ is a discrete additive subgroup of $\mathbb{R}^d$. Equivalently, it is the set

$$\mathcal{L}(\mathbf{b}_1, \ldots, \mathbf{b}_n) = \left\{ \sum_{i=1}^{n} x_i \mathbf{b}_i \ : \ x_i \in \mathbb{Z} \right\}$$

of all integer combinations of n linearly independent vectors $\mathbf{b}_1, \ldots, \mathbf{b}_n \in \mathbb{R}^d$, for $n \leq d$. Such $\mathbf{b}_i$'s form a *basis* of $\mathcal{L}$ and are usually collected in matrix form $\mathbf{B} = \begin{bmatrix} \mathbf{b}_1 & \cdots & \mathbf{b}_n \end{bmatrix}$. The lattice $\mathcal{L}$ is said to be *full-rank* if $n = d$. The *first minimum* of $\mathcal{L}$, denoted by $\lambda_1(\mathcal{L})$ is the length of a shortest nonzero vector of $\mathcal{L}$.

The *volume* of $\mathcal{L}$ is given by $\mathrm{vol}(\mathcal{L}) = \sqrt{\det(\mathbf{B}^{\mathrm{T}}\,\mathbf{B})}$ for any basis $\mathbf{B}$ of $\mathcal{L}$. In particular, for a full-rank lattice $\mathcal{L}$, $\mathrm{vol}(\mathcal{L}) = \det(\mathbf{B})$. Recall that for two lattices $\mathcal{L} \subseteq \mathcal{L}'$, the *index* of $\mathcal{L}$ in $\mathcal{L}'$ is the size of the quotient $\mathcal{L}' \,/\, \mathcal{L}$.

For a rank n lattice $\mathcal{L} \subset \mathbb{R}^d$, the *dual lattice*, denoted $\widehat{\mathcal{L}}$, is defined as the set of all points in $\mathrm{span}(\mathcal{L})$ that have integer inner products with all lattice points,

$$\widehat{\mathcal{L}} = \{\mathbf{w} \in \mathsf{span}(\mathcal{L}) : \forall \mathbf{y} \in \mathcal{L}, \langle \mathbf{w}, \mathbf{y} \rangle \in \mathbb{Z}\}.$$

Similarly, for a lattice basis $\mathbf{B} = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$, we define the dual basis $\mathbf{B}^* = (\mathbf{b}_1^*, \ldots, \mathbf{b}_n^*)$ to be the unique set of vectors in $\mathsf{span}(\mathcal{L})$ satisfying $\langle \mathbf{b}_i^*, \mathbf{b}_j \rangle = 1$ if $i = j$, and 0, otherwise. It is easy to show that $\widehat{\mathcal{L}}$ is itself a rank n lattice, that $\mathrm{vol}(\widehat{\mathcal{L}}) = \frac{1}{\mathrm{vol}(\mathcal{L})}$, and $\mathbf{B}^*$ is a basis of $\widehat{\mathcal{L}}$. Given a lattice $\mathbf{B} = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$, we denote $\| \mathbf{B} \|_2 = \max_i \|\mathbf{b}_i\|$.

Let $d \in \mathbb{N}$, $1 \leqslant k \leqslant d$ and q be a prime number. We say that a lattice L is a d-dimensional *q-ary lattice* if $q\mathbb{Z}^d \subseteq L \subseteq \mathbb{Z}^d$. Given a matrix $\mathbf{A} \in \mathbb{Z}^{d \times k}$, we consider the following d-dimensional q-ary lattices:

$$L_q(\mathbf{A}) = \left\{\mathbf{x} \in \mathbb{Z}^d : \exists \mathbf{s} \in \mathbb{Z}^k,\ \mathbf{A}\mathbf{s} = \mathbf{x} \bmod q\right\},$$
$$L_q^{\perp}(\mathbf{A}) = \left\{\mathbf{x} \in \mathbb{Z}^d : \mathbf{A}^T\mathbf{x} = \mathbf{0} \bmod q\right\}.$$

It is well-know that for any q-ary lattice L, there exists $\mathbf{A}$ and $\mathbf{B}$ such that $L = L_q(\mathbf{A}) = L_q^{\perp}(\mathbf{B})$, and that $\widehat{L_q^{\perp}(\mathbf{A})} = \frac{1}{q} L_q(\mathbf{A})$. Furthermore $\mathrm{vol}(L_q(\mathbf{A})) = q^{d - \mathrm{rk}\,\mathbf{A}} \geqslant q^{d-k}$ and therefore $\mathrm{vol}(L_q^{\perp}(\mathbf{A})) = q^{\mathrm{rk}\,\mathbf{A}} \leqslant q^k$. Finally, since $\mathbb{Z}_q$ is a field, a random matrix $\mathbf{A}$ has full rank (equal to k) with probability at least $1 - kq^{k-1-d}$.

We refer the reader to [ELZ05], [ZNKB14, § 2.5.1], [MR09] or [PS24] for more details on those constructions and why these lattices play a crucial role in lattice-based cryptography, in particular due to its connection with the LWE and SIS problems.

2.3 Random q-ary lattices

We will consider the distributions $\mathcal{L}_{d,k,q}$ and $\mathcal{L}^{\perp}_{d,k,q}$ of q-ary lattices defined over the set of integer lattices by

$$\mathcal{L}_{d,k,q}(L) = \Pr_{\mathbf{A}\sim\mathcal{U}(\mathbb{Z}_q^{d\times k})}[L = L_q(\mathbf{A})],$$
$$\mathcal{L}^{\perp}_{d,k,q}(L) = \Pr_{\mathbf{A}\sim\mathcal{U}(\mathbb{Z}_q^{d\times (d-k)})}\big[L = L_q^{\perp}(\mathbf{A})\big].$$

In other words, the distribution is obtained by taking a matrix $\mathbf{A} \in \mathbb{Z}_q^{d\times k}$ with uniform and independently distributed entries, and looking at the q-ary lattice generated by $\mathbf{A}$; and similarly for the orthogonal version. When neither k nor $d-k$ are too small, those two distributions are very close [PS24, Lemma 5].

Those distributions satisfy good uniformity properties when q goes to infinity. In particular, the following theorem shows that we can compute statistical properties of lattices sampled according to $\mathcal{L}^{\perp}_{d,k,q}$. See [PS24, § 2.5] for more context.

In this paper, we will only make use of the following special case to compute the variance of a sum over a lattice.

Theorem 2 ([PS24, Theorem 3])**.** *Let $d \in \mathbb{N}$, $1 \leqslant k \leqslant d$ and q be a prime number. For any $f : \mathbb{Z}^d \to \mathbb{R}$,*

$$\mathbb{E}_{L\sim\mathcal{L}^{\perp}_{d,k,q}}\left[\sum_{\mathbf{x}\in L} f(\mathbf{x})\right] = \sum_{\mathbf{x}\in q\mathbb{Z}^d} f(\mathbf{x}) + q^{k-d}\sum_{\mathbf{x}\in\mathbb{Z}^d\setminus q\mathbb{Z}^d} f(\mathbf{x}),$$
$$\mathbb{V}_{L\sim\mathcal{L}^{\perp}_{d,k,q}}\left[\sum_{\mathbf{x}\in L} f(\mathbf{x})\right] = (q^{k-d} - q^{2(k-d)})\sum_{\mathbf{x}\in\mathbb{Z}^d\setminus q\mathbb{Z}^d}\sum_{\mathbf{u}\in\mathbb{Z}^d}\sum_{\alpha\in\mathbb{Z}_q\setminus\{0\}} f(\mathbf{x})f(\alpha\mathbf{x}+q\mathbf{u}).$$

2.4 Short integer solution (SIS) problem

We will study the following lattice problem in this paper.

Definition 1 (SIS)**.** Let $m \geqslant n \geqslant 1$, q be integer, $\ell > 0$ and $p \in (0,\infty]$. The search problem $\mathrm{SIS}^p_{n,m,q,\ell}$ is defined as follows. The input is a random matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$. The goal is to output a vector $\mathbf{x} \in \mathbb{Z}^m$ such that $\mathbf{A}\mathbf{x} = \mathbf{0} \bmod q$ and $0 < \|\mathbf{x}\|_p \leqslant \ell$.

Using the notation from Section 2.2, solving $\mathrm{SIS}^p_{n,m,q,\ell}$ is equivalent to finding a nonzero vector in $L_q^{\perp}(\mathbf{A}^T) \cap B_n^p(\ell)$ with $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ sampled uniformly at random. Using the notation from Section 2.3, this is equivalent to finding a nonzero vector in $L \cap B_n^p(\ell)$ with $L \sim \mathcal{L}^{\perp}_{m,m-n,q}$. Note that "$k = m - n$" here because our definition of $\mathcal{L}^{\perp}_{d,k,q}$ samples a matrix in $\mathbb{Z}_q^{d\times(d-k)}$.

When $p = 2$, the SIS problem is known to be as hard as solving approximate SVP for all lattices [Ajt96]. In practice, a standard lattice-reduction algorithm (BKZ) is used on the lattice $L_q^{\perp}(\mathbf{A}^T)$ to estimate the complexity of the problem.

Due to its simplicity, the case of $p = \infty$ is more often used in cryptography. However, when the parameters are set so that $\ell\sqrt{m} > q$, there is no known worst-case problem that is reducible to the average case $\mathrm{SIS}^{\infty}_{n,m,q,\ell}$. This parameter setting is used in the recent NIST-standardized signature scheme, Dilithium [DKL+18]. In its documentation, the authors proposed a heuristic model to estimate the cost of $\mathrm{SIS}^{\infty}_{n,m,q,\ell}$, by assuming that the coordinates of the short vectors generated by a sieving algorithm follow a Gaussian

distribution. This model was not backed by experiments as far as we know. More generally, there are very few algorithms known for solving $\mathrm{SIS}^{\infty}_{n,m,q,\ell}$.

In [CLZ22], the authors mentioned a polynomial-time classical algorithm from Regev that solves the problem when q is a composite number and m is very large. They also proposed a polynomial-time quantum algorithm (that was later dequantized in [KOW26]) that solves $\mathrm{SIS}^{\infty}_{n,m,q,\frac{q-c}{2}}$ when $c \in \mathbb{Z}$ is a constant and m is polynomial in n. More recently, [DEL25] showed a subexponential algorithm for $\mathrm{SIS}^{\infty}_{n,m,q,\ell}$ when $q = \mathsf{poly}(n)$, $m = n + \omega(n/\log\log(n))$ and $\ell = q/\operatorname{polylog}(n)$. However, when instantiating their algorithm on the concrete parameters of Dilithium, the results are far from threatening Dilithium's security.

For convenience, when we discuss the running time of the algorithms solving the problems above, we ignore polynomial factors in the bit-length of the individual input basis vectors, i.e., we assume the input basis has a bit-size polynomial in the ambient dimension d.

2.5 BKZ

BKZ is a famous algorithm to reduce a lattice basis [Sch87]. We refer the reader to [HPS11] recent work on this topic, and [LN20] for a formal analysis of its running time.

Let $\mathbf{B}$ be a BKZ-β reduced basis lattice (of rank d) in $\mathbb{R}^d$. We will denote by $\widetilde{\mathbf{b}}_1, \ldots, \widetilde{\mathbf{b}}_d$ its Gram-Schmidt vectors. Recall that the root Hermite factor $\delta_{\mathbf{B}}$ is defined by

$$\|\mathbf{b}_1\| = \delta_{\mathbf{B}}^{d-1} \operatorname{vol}(L)^{1/d}.$$

Experimentally, it has been verified [Che13] that

$$\delta_{\mathbf{B}} \approx H_\beta \coloneqq \left(\frac{\beta}{2\pi e}(\pi\beta)^{1/\beta}\right)^{1/2(\beta-1)} \tag{2}$$

See also [EJK20] for more information. We will estimate the asymptotic runtime of BKZ with block size β by $2^{0.292\beta+o(\beta)}$ [BDGL16].

We will make use of the Geometric Series Assumption (GSA) [Sch03] to obtain more precise information on the shape of the basis.

Heuristic 1 (Geometric Series Assumption (GSA))**.** *Let $\mathbf{b}_1, \ldots, \mathbf{b}_d$ be a BKZ-β reduced basis for a lattice L and $\widetilde{\mathbf{b}}_1, \ldots, \widetilde{\mathbf{b}}_d$ be the corresponding Gram-Schmidt vectors. Then for all $i = 1, \ldots, d$,*

$$\|\widetilde{\mathbf{b}}_i\| = \|\mathbf{b}_1\| H_\beta^{-2(i-1)}, \qquad \|\mathbf{b}_1\| = H_\beta^{d-1} \operatorname{vol}(L)^{1/d}.$$

The GSA is known to be reasonably accurate when β is larger than 50 and much smaller than d, although it does not correctly model what happens in the last $d - \beta$ coordinates, see e.g. [AD21].

2.6 Discrete Gaussian distribution

Let $d \in \mathbb{N}$ and $s > 0$. For any $\mathbf{x} \in \mathbb{R}^d$, we let $\rho_s(\mathbf{x}) = e^{-\pi\|\mathbf{x}\|^2/s^2}$. We extend ρ_s to sets by

$$\rho_s(X) = \sum_{\mathbf{x}\in X} \rho_s(\mathbf{x})$$

for any set X. For any lattice $\mathcal{L} \subset \mathbb{R}^d$, we denote the *discrete Gaussian distribution* over $\mathcal{L}$ by $D_{\mathcal{L},s}(\mathbf{x}) = \frac{\rho_s(\mathbf{x})}{\rho_s(\mathcal{L})}$ for any $\mathbf{x} \in \mathcal{L}$. We denote $D_{\mathcal{L},1}$ by $D_{\mathcal{L}}$ for simplicity. It is well-known by the Poisson summation formula that for any lattice $\mathcal{L}$ and any $s > 0$,

$$\rho_{1/s}(\widehat{\mathcal{L}}) = \operatorname{vol}(\mathcal{L}) s^{-d} \rho_s(\mathcal{L}).$$

It easily follows from the Poisson summation formula that $\mathbf{t} \in \mathbb{R}^d$, $\rho_s(\mathbf{t}+\mathcal{L}) \leqslant \rho_s(\mathcal{L})$ [Ban93]. See e.g. [Ste17] for a good introduction on this topic. The discrete Gaussian distribution plays an essential role in lattice-based cryptography and an important problem is to be able to sample efficiently from it: this is known as the *discrete Gaussian sampling (DGS)* problem.

Definition 2 (DGS)**.** For $\delta = \delta(d) \geq 0$, s a function that maps lattices to non-negative real numbers, and $m = m(d) \in \mathbb{N}$, δ-DGS_s^m is defined as follows. The input is a basis $\mathbf{B}$ for a lattice $\mathcal{L} \subset \mathbb{R}^d$ and a parameter $s > s(\mathcal{L})$. The goal is to output a sequence of m vectors whose joint distribution is δ-close to m independent samples from $D_{\mathcal{L},s}$.

We omit the parameter δ if $\delta = 0$, and the parameter m if $m = 1$. We stress that δ bounds the statistical distance between the *joint* distribution of the output vectors and m independent samples from $D_{\mathcal{L},s}$.

In general, the smaller s is, the harder it is to construct a sampler for $D_{\mathcal{L},s}$. The notion of smoothing parameter [MR04] captures the idea that sampling for a value of s above this threshold is significantly easier than sampling below because the distribution looks more like a continuous Gaussian. Formally, for any $\varepsilon > 0$, the smoothing parameter of a lattice $\mathcal{L}$ is defined by

$$\eta_\varepsilon(\mathcal{L}) = \inf\Big\{ s > 0 : \rho_{1/s}(\widehat{\mathcal{L}}) \leqslant 1+\varepsilon \Big\}.$$

Observe that $\eta_\varepsilon(\mathcal{L})$ is a decreasing function of ε.

There are many algorithms to sample above the smoothing parameter [Kle00, GPV08] [BLP+13], including a time-space trade-off [ACKS21]. Sampling below the smoothing parameter is much more challenging and usually inefficient [ADRS15], unless the lattice possesses a well-defined structure [BLL25]. At the extreme, sampling for sufficiently small values of s allows one to solve the *shortest vector problem* (SVP) [ADRS15] which is known to be NP-hard under randomized reduction [Ajt98], and also the *closest vector problem* (CVP) [ADSD15]. The Monte Carlo Markov Chain based algorithm of [WL19] (see Theorem 3 below) works for all values of s but the complexity significantly depends on s and the shape of the basis [PS24]. Finally, it should be noted that the problem of estimating the smoothing parameter is provably harder than BDD and LWE in certain parameter ranges [CDLP13].

Theorem 3 ([WL19])**.** *There is an algorithm that given a basis of a lattice $L \subset \mathbb{R}^d$, any vector $\mathbf{t} \in \mathbb{R}^n$, any $\varepsilon > 0$ and any $s > 0$, returns a sample according to some distribution $\mathcal{D}_{L,s,\mathbf{t},\varepsilon}$ such that $\mathrm{d_{TV}}(\mathcal{D}_{L,s,\mathbf{t},\varepsilon}, D_{L,s,\mathbf{t}}) \leqslant \varepsilon$. This algorithm runs in time $\ln\left(\frac{1}{\varepsilon}\right) \cdot \frac{1}{\Delta} \cdot \mathsf{poly}(d)$ where $\frac{1}{\Delta} = \frac{1}{\rho_s(\mathbf{t}+L)} \prod_{i=1}^k \rho_{s/\|\widetilde{\mathbf{b}}_i\|}(\mathbb{Z}) \leqslant \prod_{i=1}^k \rho_{\|\widetilde{\mathbf{b}}_i\|/s}(\mathbb{Z})$ and $\widetilde{\mathbf{b}}_1, \dots, \widetilde{\mathbf{b}}_k$ are the Gram-Schmidt vectors of the basis.*

The main difficulty in using this theorem is to estimate $\prod_{i=1}^k \rho_{\|\widetilde{\mathbf{b}}_i\|/s}(\mathbb{Z})$ which strongly depends on the basis. In [PS25], the authors gave an approximation formula for this quantity when the basis is BKZ-reduced, and assuming that the Geometric Series Assumption (GSA) holds.

Theorem 4 ([PS25, Theorem 3]†)**.** *Let $0 < d \leqslant n$ and $0 < \beta \leqslant d$. Let $\mathbf{b}_1, \dots, \mathbf{b}_d \in \mathbb{R}^n$ be a BKZ-β reduced basis of a lattice L and $\widetilde{\mathbf{b}}_1, \dots, \widetilde{\mathbf{b}}_d$ be the Gram-Schmidt vectors of the basis. Let $s > 0$ and $\alpha = \|\mathbf{b}_1\|/s$. If the GSA holds, then for any odd number $p \geqslant 1$,*

$$\ln\left(\prod_{i=1}^d \rho_{\|\widetilde{\mathbf{b}}_i\|/s}(\mathbb{Z})\right) \leqslant A + \sum_{\ell=1}^p \frac{2^\ell(-1)^{\ell+1}}{\ell}(B_\ell + C_\ell)$$

†There seems to be a small error in the statement in [PS25]: the two $E_1(\cdot)$ terms in C_ℓ should be swapped, otherwise it would be negative.

where

$$A = (d_0+1)\ln\frac{\alpha}{H_\beta^{d_0}} + d\varepsilon, \qquad B_\ell = \frac{E_1\left(\pi\ell\alpha^2 H_\beta^{-4(d_0+\frac{1}{2})}\right) - E_1(\pi\ell\alpha^2 H_\beta^{-2})}{2\ln(H_\beta)},$$
$$C_\ell = \frac{E_1\left(\frac{\pi}{\alpha^2}\ell H_\beta^{4(d_0+\frac{1}{2})}\right) - E_1\left(\frac{\pi}{\alpha^2}\ell H_\beta^{4(d-\frac{1}{2})}\right)}{2\ln(H_\beta)}, \qquad H_\beta = \left(\frac{\beta}{2\pi e}(\pi\beta)^{1/\beta}\right)^{1/2(\beta-1)},$$

where $d_0 = \max\left(-1, \min\left(d-1, \left\lfloor \frac{\ln(\alpha)}{2\ln(H_\beta)} \right\rfloor\right)\right)$ *and* $\varepsilon \leqslant 6.974685811 \times 10^{-6}$*, and* E_1 *refers to the exponential integral.*

3 Technical results on the discrete Gaussian distribution

In this section, we give some technical results on the discrete Gaussian. They are relevant in our work, but can also be of independent interest.

3.1 On the ratio of the Gaussian mass of $\mathbb{Z}$

We start by demonstrating a preliminary result, which serves as a crucial step towards our main contribution. It describes the behavior of the ratio of Gaussian masses of $\mathbb{Z}$, $\rho_{\gamma s}(\mathbb{Z})/\rho_s(\mathbb{Z})$ for $\gamma > 1$, illustrated in Figure 1. The idea of the proof comes from the fact that this ratio can also be rewritten in terms of two Jacobi theta functions, i.e. $\rho_{\gamma s}(\mathbb{Z})/\rho_s(\mathbb{Z}) = \vartheta_3(i/(\gamma s)^2)/\vartheta_3(i/s^2)$, and one can use a similar argument to [BLY23, Lemma 38] to analyse its monotonicity.

Lemma 1. *For any* $\gamma > 1$*, the ratio* $\frac{\rho_{\gamma s}(\mathbb{Z})}{\rho_s(\mathbb{Z})}$ *is increasing for* $s \geqslant 0$*.*

Proof. Observe that $\vartheta_3(z) = \sum_{\mathbf{x}\in\mathbb{Z}} e^{\pi i z\|\mathbf{x}\|^2}$, for $\mathrm{Im}(z) > 0$. Then,

$$\frac{\rho_{\gamma s}(\mathbb{Z})}{\rho_s(\mathbb{Z})} = \frac{\vartheta_3(i/(\gamma s)^2)}{\vartheta_3(i/s^2)} = \frac{\vartheta_3((1/\gamma^2)(i/s^2))}{\vartheta_3(i/s^2)}, \tag{3}$$

where following the definition of $\vartheta_3(z)$, we consider $z = i/(\gamma s)^2$ in the numerator and $z = i/s^2$ in the denominator.

By the expansion of $\vartheta_3(z)$ from [CS99, p. 105], we have that

$$\vartheta_3(z) = \prod_{m=1}^{\infty}(1-q^{2m})(1+q^{2m-1})^2,$$

where $q = e^{\pi i z}$, for $\mathrm{Im}(z) > 0$. To simplify notation, we will denote $a_\gamma = \pi/(\gamma s)^2$ and $a_s = \pi/s^2$. Now, returning to (3) and making the corresponding choices of z, we get

$$\begin{aligned}\frac{\rho_{\gamma s}(\mathbb{Z})}{\rho_s(\mathbb{Z})} &= \frac{\prod_{m=1}^{\infty}(1-e^{-2ma_\gamma})(1+e^{-(2m-1)a_\gamma})^2}{\prod_{m=1}^{\infty}(1-e^{-2ma_s})(1+e^{-(2m-1)a_s})^2} \\ &\overset{(a)}{=} \prod_{m=1}^{\infty}\frac{8e^{-(3m-1)a_\gamma}\sinh(ma_\gamma)\cosh^2((m-1/2)a_\gamma)}{8e^{-(3m-1)a_s}\sinh(ma_s)\cosh^2((m-1/2)a_s)} \\ &= \prod_{m=1}^{\infty}\underbrace{e^{-(3m-1)(a_\gamma-a_s)}}_{f_1(s)}\underbrace{\frac{\sinh(ma_\gamma)}{\sinh(ma_s)}}_{f_2(s)}\underbrace{\frac{\cosh^2((m-1/2)a_\gamma)}{\cosh^2((m-1/2)a_s)}}_{f_3^2(s)},\end{aligned} \tag{4}$$

where in (a), we have used both in the numerator and denominator that $1 - e^{-2ma} = e^{-ma}(e^{ma} - e^{-ma}) = 2e^{-ma}\sinh(ma)$, and analogously, $1 + e^{-(2m-1)a} = e^{-(m-1/2)a}(e^{(m-1/2)a} - e^{-(m-1/2)a}) = 2e^{-(m-1/2)a}\cosh((m-\frac{1}{2})a)$, for $a = a_\gamma$ and $a = a_s$, respectively.

To show that $\rho_{\gamma s}(\mathbb{Z})/\rho_s(\mathbb{Z})$ is an increasing function in s, we will show that the component functions $f_1(s), f_2(s)$ and $f_3^2(s)$ in (4) are positive and increasing for $s \in (0,\infty)$ and all $m \in \mathbb{Z}$, and consequently, their product is also increasing.

Notice that $f_1(s)$ is increasing and positive for $s \in (0,\infty)$, since $a_\gamma - a_s < 0$ and $m \geq 1$.

Moving forward, we will study $f_2(s) = \frac{\sinh(ma_\gamma)}{\sinh(ma_s)} = \frac{\sinh(\alpha/s^2)}{\sinh(\beta/s^2)}$, where $\alpha = \frac{1}{\gamma^2}\beta$ and $\beta = m\pi$. Since $\gamma > 1$, $\alpha < \beta$. Thus, we have that,

$$f_2'(s) = \frac{2}{s^3} f_2(s)\big(\beta\coth(\beta/s^2) - \alpha\coth(\alpha/s^2)\big). \tag{5}$$

Define $g(x) = x\coth(x)$, then

$$\beta\coth(\beta/s^2) - \alpha\coth(\alpha/s^2) = s^2(g(\beta/s^2) - g(\alpha/s^2))$$

and since $g'(x) = \coth(x) - x\,\mathrm{csch}^2(x) > 0$ for every $x > 0$, $g(x)$ also is increasing on $(0,\infty)$. Since $0 < \alpha < \beta$, it holds that $s^2(g(\beta/s^2) - g(\alpha/s^2)) > 0$ and since $f_2(s) > 0$ and $2/s^3 > 0$, it follows that $f_2'(s) > 0$ for all $s > 0$. Consequently, $f_2(s)$ is strictly increasing and positive for $s \in (0,\infty)$.

We now proceed with a similar analysis for $f_3(s) = \frac{\cosh((m-1/2)a_\gamma)}{\cosh((m-1/2)a_s)} = \frac{\cosh(\alpha/s^2)}{\cosh(\beta/s^2)}$, where now, $\alpha = \frac{1}{\gamma^2}\beta$ and $\beta = (m-1/2)\pi$. Analogous to the previous case, $\alpha < \beta$ since $\gamma > 1$. Hence,

$$f_3'(s) = \frac{2}{s^3} f_3(s)\big(\beta\tanh(\beta/s^2) - \alpha\tanh(\alpha/s^2)\big). \tag{6}$$

If we set $h(x) = x\tanh(x)$, then

$$\beta\tanh(\beta/s^2) - \alpha\tanh(\alpha/s^2) = s^2(h(\beta/s^2) - h(\alpha/s^2)),$$

and $h'(x) = \tanh(x) + x\,\mathrm{sech}^2(x) > 0$ for $x > 0$. Consequently, $h(x)$ is strictly increasing and for $0 < \alpha < \beta$, it holds that $s^2(h(\beta/s^2) - h(\alpha/s^2)) > 0$. Moreover, $f_3(s) > 0$ and $2/s^3 > 0$, which implies that $f_3(s)$ and, consequently $f_3^2(s)$ are both strictly increasing, for $s \in (0,\infty)$.

Therefore, since the factors all functions $f_1(s), f_2(s)$ and $f_3(s)$ are strictly increasing and positive functions from $s \in (0,\infty)$, so is their product. Since this holds for all $m \in \mathbb{N}$, we can conclude from (4) that $\frac{\rho_{\gamma s}(\mathbb{Z})}{\rho_s(\mathbb{Z})}$ is increasing for all $s \in (0,\infty)$ and $\gamma > 1$, and the proof is complete. □

3.2 Bounds on certain sums over shifted copies of $\mathbb{Z}^d$

When estimating the variance of the Gaussian mass of a lattice, we will need to bound certain sums involving the Gaussian mass of the shifted lattice $\mathbb{Z}^d$. More precisely, consider, for a given $\alpha \in \mathbb{Z}_q \setminus \{0\}$, the sum

$$S(\alpha) = \sum_{\mathbf{x}\in\mathbb{Z}^d} \rho_s(\mathbf{x})\rho_s(\alpha\mathbf{x} + q\mathbb{Z}^d). \tag{7}$$

The claim is that this sum should be maximum for $\alpha = 1$. The intuition is that each term $\rho_s(\mathbf{x})$ is "weighted down" by $\rho_s(\alpha\mathbf{x} + q\mathbb{Z}^d)$ which, for most values of α and $\mathbf{x}$, will be essentially close to $\rho_s(\alpha\mathbf{x})$, and therefore we want to take α as small as possible. Informally,

$$S(\alpha) \approx \sum_{\mathbf{x}\in\mathbb{Z}^d} \rho_s(\mathbf{x})\rho_s(\alpha\mathbf{x}) = \sum_{\mathbf{x}\in\mathbb{Z}^d} \rho_{s/\sqrt{1+\alpha^2}}(\mathbf{x}) = \rho_{s/\sqrt{1+\alpha^2}}(\mathbb{Z}^d)$$

which is indeed maximum for $\alpha = 1$. The next lemma formalizes this intuition.

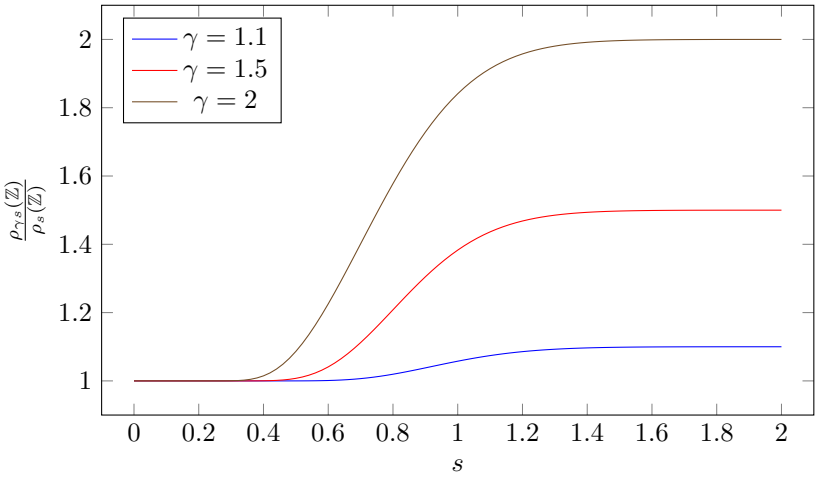


Figure 1: Plot of $\frac{\rho_{\gamma s}(\mathbb{Z})}{\rho_s(\mathbb{Z})}$ for several values of γ.

Lemma 2. *For any $\alpha \in \mathbb{Z}_q \setminus \{0\}$, $d \geqslant 1$ and $q \geqslant 3$,*

$$\sum_{\mathbf{x} \in \mathbb{Z}^d} \rho_s(\mathbf{x}) \rho_s(\alpha \mathbf{x} + q\mathbb{Z}^d) \leqslant \rho_{s/\sqrt{2}}(\mathbb{Z}^d) \rho_{s\sqrt{2}/q}(\mathbb{Z}^d).$$

Remark 1. Note that for $\alpha = 1$, and by [DRSD14, Lemma 2.14], the sum is equal to

$$\begin{aligned}\sum_{\mathbf{u} \in \mathbb{Z}^d} \rho_{s\sqrt{2}}(q\mathbf{u}) \rho_{s\sqrt{2}}(\mathbf{q}u + 2\mathbb{Z}^d) &\geqslant \sum_{\mathbf{u} \in \mathbb{Z}^d} \rho_{s\sqrt{2}}(q\mathbf{u}) \rho_{s\sqrt{2}}(\mathbf{q}u \bmod 2) \rho_{s\sqrt{2}}(2\mathbb{Z}^d) \\ &\geqslant \sum_{\mathbf{u} \in \mathbb{Z}^d} \rho_{s\sqrt{2}}(q\mathbf{u}) \rho_{s\sqrt{2}}(\sqrt{d}) \rho_{s/\sqrt{2}}(\mathbb{Z}^d) \\ &= \rho_{s\sqrt{2}}(\sqrt{d}) \rho_{s\sqrt{2}/q}(\mathbb{Z}^d) \rho_{s/\sqrt{2}}(\mathbb{Z}^d).\end{aligned}$$

Therefore Lemma 2 is tight[‡] when $\sqrt{d}/s$ is not large.

Proof. First, observe that for any $\alpha \in \mathbb{Z}_q$,

$$F_{d,q}(\alpha) = \sum_{\mathbf{x}, \mathbf{u} \in \mathbb{Z}^d} \rho_s(\mathbf{x}) \rho_s(\alpha \mathbf{x} + q\mathbf{u}) = \rho_s(\mathcal{L})$$

where $\mathcal{L} = \{(\mathbf{x}, \alpha \mathbf{x} + q\mathbf{u}) : \mathbf{x}, \mathbf{u} \in \mathbb{Z}^d\}$. Since ρ_s only depends on the 2-norm of its argument, we have that $F_{d,q}(\alpha) = \rho_s(R\,\mathcal{L})$ for any rotation R. Consider the rotation R defined by its matrix:

$$R = \begin{bmatrix} \frac{\alpha}{\sqrt{1+\alpha^2}} I_d & \frac{-1}{\sqrt{1+\alpha^2}} I_d \\ \frac{1}{\sqrt{1+\alpha^2}} I_d & \frac{\alpha}{\sqrt{1+\alpha^2}} I_d \end{bmatrix}$$

where I_n is the identity matrix in dimension d. Then check that

$$R\,\mathcal{L} = \tfrac{1}{\sqrt{1+\alpha^2}} \{(-q\mathbf{u}, (1+\alpha^2)\mathbf{x} + \alpha q \mathbf{u}) : \mathbf{x}, \mathbf{u} \in \mathbb{Z}^d\}.$$

It follows that

$$F_{d,q}(\alpha) = \sum_{\mathbf{x}, \mathbf{u} \in \mathbb{Z}^d} \rho_{s\sqrt{1+\alpha^2}}(-q\mathbf{u}) \rho_{s\sqrt{1+\alpha^2}}((1+\alpha^2)\mathbf{x} + \alpha q \mathbf{u})$$

[‡]An upper bound independent of α is sufficient, it will become clear why in Section 4.1.

$$
\begin{aligned}
&= \sum_{\mathbf{u}\in\mathbb{Z}^d} \rho_{s\sqrt{1+\alpha^2}}(-q\mathbf{u})\rho_{s\sqrt{1+\alpha^2}}(\alpha q\mathbf{u} + (1+\alpha^2)\mathbb{Z}^d) \\
&\leqslant \sum_{\mathbf{u}\in\mathbb{Z}^d} \rho_{s\sqrt{1+\alpha^2}}(-q\mathbf{u})\rho_{s\sqrt{1+\alpha^2}}((1+\alpha^2)\mathbb{Z}^d) \\
&= \rho_{s\sqrt{1+\alpha^2}}(q\mathbb{Z}^d)\rho_{s\sqrt{1+\alpha^2}}((1+\alpha^2)\mathbb{Z}^d) = \rho_{s\sqrt{1+\alpha^2}/q}(\mathbb{Z}^d)\rho_{s/\sqrt{1+\alpha^2}}(\mathbb{Z}^d).
\end{aligned}
$$

Now observe that $F_{d,q}$ is q-periodic so we can assume that $\alpha \in \left\{-\frac{q-1}{2}, \dots, \frac{q-1}{2}\right\} \setminus \{0\}$ and therefore $1+\alpha^2 \leqslant q^2/4$ since $q \geqslant 3$. It follows that $\sqrt{1+\alpha^2}/q \leqslant 1/2 \leqslant 1/\sqrt{2}$. We can use Lemma 1 for $\gamma = \frac{\sqrt{1+\alpha^2}}{\sqrt{2}} \geqslant 1$ (since $\alpha \neq 0$) to get that

$$
\begin{aligned}
F_{d,q}(\alpha) &= \frac{\rho_{s\sqrt{1+\alpha^2}/q}(\mathbb{Z}^d)}{\rho_{s\sqrt{2}/q}(\mathbb{Z}^d)}\rho_{s\sqrt{2}/q}(\mathbb{Z}^d)\rho_{s/\sqrt{1+\alpha^2}}(\mathbb{Z}^d) \\
&\leqslant \frac{\rho_{s/\sqrt{2}}(\mathbb{Z}^d)}{\rho_{s/\sqrt{1+\alpha^2}}(\mathbb{Z}^d)}\rho_{s\sqrt{2}/q}(\mathbb{Z}^d)\rho_{s/\sqrt{1+\alpha^2}}(\mathbb{Z}^d) = \rho_{s/\sqrt{2}}(\mathbb{Z}^d)\rho_{s\sqrt{2}/q}(\mathbb{Z}^d).
\end{aligned}
$$

□

3.3 Bounds on certain sums in a ball

In the next section, we will need to bound a similar sum to (7) but restricted to an ℓ_p-ball. This is quite difficult to do in the general case when the ball contains some q-vectors. However we show that if the ball is not too large, we can obtain a tight bound. We start with a straightforward combinatorial lemma.

Lemma 3. *Let $k \geqslant 1$, τ be a permutation of $\{1,\dots,k\}$ and $x_1,\dots,x_k \in \mathbb{R}$. Then $\sum_{i=1}^k x_i x_{\tau(i)} \leqslant \sum_{i=1}^k x_i^2$.*

Proof. By a straightforward calculation: $\sum_{i=1}^k x_i x_{\tau(i)} \leqslant \frac{1}{2}\sum_{i=1}^k \left(x_i^2 + x_{\tau(i)}^2\right) = \sum_{i=1}^k x_i^2$ since τ is a permutation. □

Similar to Lemma 2, we want to show that the sum is maximum for $\alpha = 1$.

Lemma 4. *For any $p \in [1,\infty]$, prime q, $\alpha \in \mathbb{Z}_q \setminus \{0\}$, $d \geqslant 1$ and $0 \leqslant \ell \leqslant q/2$,*

$$
\sum_{\mathbf{x}\in B_d^{p,\mathbb{Z}}(\ell)} \sum_{\substack{\mathbf{u}\in\mathbb{Z}^d: \\ \alpha\mathbf{x}+q\mathbf{u}\in B_d^p(\ell)}} \rho_s(\mathbf{x})\rho_s(\alpha\mathbf{x}+q\mathbf{u}) \leqslant \rho_{s/\sqrt{2}}(B_d^{p,\mathbb{Z}}(\ell)).
$$

Remark 2. This upper bound cannot be improved because for $\alpha = 1$, the sum is

$$
\sum_{\mathbf{x}\in B_d^{p,\mathbb{Z}}(\ell)} \sum_{\substack{\mathbf{u}\in\mathbb{Z}^d: \\ \mathbf{x}+q\mathbf{u}\in B_d^p(\ell)}} \rho_s(\mathbf{x})\rho_s(\mathbf{x}+q\mathbf{u}) = \sum_{\mathbf{x}\in B_d^{p,\mathbb{Z}}(\ell)} \rho_s(\mathbf{x})\rho_s(\mathbf{x}) = \rho_{s/\sqrt{2}}(B_d^{p,\mathbb{Z}}(\ell))
$$

since if $\mathbf{x} \in B_d^{p,\mathbb{Z}}(\ell)$ then the only value of $\mathbf{u}$ for which $\mathbf{x}+q\mathbf{u} \in B_d^p(\ell)$ is $\mathbf{u} = 0$.

Proof. Note that the sum is a q-periodic function of α, so we can assume without loss of generality that $\alpha \in \{1,\dots,q-1\}$ in the rest of the proof.

For any $\mathbf{x} \in \mathbb{Z}^d$, let $U_\alpha(\mathbf{x}) := \left\{\mathbf{u} \in \mathbb{Z}^d : \alpha\mathbf{x}+q\mathbf{u} \in B_d^p(\ell)\right\}$. We claim that $|U_\alpha(\mathbf{x})| \leqslant 1$. Indeed, let $\mathbf{u},\mathbf{v} \in U_\alpha(\mathbf{x})$ and observe that

$$
\|\alpha\mathbf{x}+q\mathbf{u}\|_p + \|\alpha\mathbf{x}+q\mathbf{v}\|_p < 2\ell \leqslant q
$$

by definition of $U_\alpha(\mathbf{x})$. On the other hand,

$$\|\alpha\mathbf{x} + q\mathbf{u}\|_p + \|\alpha\mathbf{x} + q\mathbf{v}\|_p \geqslant \|(\alpha\mathbf{x} + q\mathbf{u}) - (\alpha\mathbf{x} + q\mathbf{v})\|_p = q\|\mathbf{u} - \mathbf{v}\|_p.$$

Since $\mathbf{u}, \mathbf{v} \in \mathbb{Z}^d$, either $\mathbf{u} = \mathbf{v}$ or $\|\mathbf{u} - \mathbf{v}\|_p \geqslant 1$, but the latter contradicts the above equations.

For any $\mathbf{x} \in B_d^{p,\mathbb{Z}}(\ell)$, let $f(\mathbf{x}) = \alpha\mathbf{x} + q\mathbf{u}$ if $U_\alpha(\mathbf{x}) = \{\mathbf{u}\}$, and $f(\mathbf{x})$ be undefined if $U_\alpha(\mathbf{x})$ is empty. Then by definition, f is a map from $B_d^{p,\mathbb{Z}}(\ell)$ to itself. Furthermore, we claim that f is injective. Indeed, if we have $f(\mathbf{x}) = f(\mathbf{y})$ for some $\mathbf{x}, \mathbf{y} \in B_d^{p,\mathbb{Z}}(\ell)$ then $\alpha(\mathbf{x} - \mathbf{y}) = q(\mathbf{u} - \mathbf{v})$ where $U_\alpha(\mathbf{x}) = \{\mathbf{u}\}$ and $U_\alpha(\mathbf{y}) = \{\mathbf{v}\}$. Recall that $\mathbf{x}, \mathbf{y}, \mathbf{u}, \mathbf{v}$ have integer coordinates and q is prime. Therefore, q divides $\alpha(\mathbf{x} - \mathbf{y})$. But q is prime and $\alpha \in \{1, \ldots, q-1\}$ so q divides $\mathbf{x} - \mathbf{y}$. But since $\mathbf{x}, \mathbf{y} \in B_d^{p,\mathbb{Z}}(\ell)$, we have that $\|\mathbf{x} - \mathbf{y}\|_p < 2\ell \leqslant q$. Therefore the only possibility is that $\mathbf{x} = \mathbf{y}$.

Finally, observe that since f is injective from $B_d^{p,\mathbb{Z}}(\ell)$ into intself, we can extend f to an arbitrary permutation τ of $B_d^{p,\mathbb{Z}}(\ell)$. We then have that

$$\begin{aligned}\sum_{\mathbf{x}\in B_d^{p,\mathbb{Z}}(\ell)} \sum_{\mathbf{u}\in U_\alpha(\mathbf{x})} \rho_s(\mathbf{x})\rho_s(\alpha\mathbf{x} + q\mathbf{u}) &= \sum_{\mathbf{x}\in \mathrm{Supp}(f)} \rho_s(\mathbf{x})\rho_s(f(\mathbf{x})) \\ &\overset{(a)}{\leqslant} \sum_{\mathbf{x}\in B_d^{p,\mathbb{Z}}(\ell)} \rho_s(\mathbf{x})\rho_s(\tau(\mathbf{x})) \\ &\overset{(b)}{\leqslant} \sum_{\mathbf{x}\in B_d^{p,\mathbb{Z}}(\ell)} \rho_s(\mathbf{x})\rho_s(\mathbf{x}) = \rho_{s/\sqrt{2}}(B_d^{p,\mathbb{Z}}(\ell)),\end{aligned}$$

where (a) holds since ρ_s is non-negative, and (b) is derived from Lemma 3. □

3.4 Other useful bounds

The following results are useful for the discrete Gaussian mass of lattices. The first result gives a tight bound on the Gaussian mass of the interval $[-\ell, \ell]$ (in dimension one). This will be very useful when studying SIS in the infinity norm where we need bounds on the Gaussian mass of the ℓ_∞-ball of radius ℓ which is $[-\ell, \ell]^d$, see Theorem 6.

Lemma 5. *For any $s > 0$ and $\ell \geqslant 1$, we have that*

$$1 + s\left(\mathrm{erf}\left(\frac{\ell\sqrt{\pi}}{s}\right) - \mathrm{erf}\left(\frac{\sqrt{\pi}}{s}\right)\right) \leqslant \rho_s([-\ell, \ell]) \cap \mathbb{Z}) \leqslant 1 + s\,\mathrm{erf}\left(\frac{\ell\sqrt{\pi}}{s}\right).$$

Proof. Let $N = \lfloor \ell \rfloor$ so that $\rho_s([-\ell, \ell] \cap \mathbb{Z}) = \sum_{k=-N}^{N} \rho_s(k) = 1 + 2\sum_{k=1}^{N} \rho_s(k)$. Then check that since ρ_s is decreasing over $[0, \infty]$,

$$\sum_{k=1}^{N} \rho_s(k) \leqslant \sum_{k=1}^{N} \int_{k-1}^{k} \rho_s(t)\,\mathrm{d}\,t = \int_0^k \rho_s(t)\,\mathrm{d}\,t = \frac{s}{2}\,\mathrm{erf}\left(\frac{N\sqrt{\pi}}{s}\right)$$

and we conclude since $N \leqslant \ell$ and erf is increasing. Similarly,

$$\sum_{k=1}^{N} \rho_s(k) \geqslant \int_1^{k+1} \rho_s(t)\,\mathrm{d}\,t = \frac{s}{2}\left(\mathrm{erf}\left(\frac{(N+1)\sqrt{\pi}}{s}\right) - \mathrm{erf}\left(\frac{\sqrt{\pi}}{s}\right)\right)$$

and we conclude since $N \geqslant \ell - 1$ and erf is increasing. □

Our second result gives some intuitive asymptotics on the Gaussian mass of $\mathbb{Z}^d$ for extreme values of the width, specifically that $\rho_s(\mathbb{Z})^d$ should be very close to s^d for large s and very close to 1 for small s.

Lemma 6. *Let $s = s(d) > 0$. If $s = \omega(\sqrt{\log d})$ then $\rho_s(\mathbb{Z})^d = s^d(1 + O(de^{-\pi s^2})) = s^d(1 + o(1))$ as $d \to \infty$. If $s = \omega(\sqrt{\log d})$ then $\rho_{1/s}(\mathbb{Z})^d = 1 + O(de^{-\pi s^2}) = 1 + o(1)$ as $d \to \infty$.*

Proof. Recall that by the Poisson formula, for $s \to \infty$,

$$\rho_s(\mathbb{Z}) = s\left(1 + 2\sum_{k=1}^{\infty} e^{-\pi k^2 s^2}\right) \leqslant s + 2s\sum_{k=1}^{\infty} e^{-\pi k s^2} = s(1 + 2e^{-\pi s^2} + o(e^{-\pi s^2})).$$

Therefore,

$$\begin{aligned}\rho_s(\mathbb{Z})^d &\leqslant s^d e^{d\log\left(1+2e^{-\pi s^2}+o(e^{-\pi s^2})\right)} \leqslant s^d e^{2de^{-\pi s^2}+o(de^{-\pi s^2})} \\ &= s^d(1 + O(de^{-\pi s^2})) = s^d(1 + o(1))\end{aligned}$$

if $s = \omega(\sqrt{\log d})$. Similarly, for $s \to \infty$,

$$\rho_{1/s}(\mathbb{Z}) \leqslant 1 + 2\sum_{k=1}^{\infty} e^{-\pi k s^2} = 1 + 2e^{-\pi s^2} + o(e^{-\pi s^2})$$

and therefore

$$\rho_{1/s}(\mathbb{Z})^d \leqslant e^{2de^{-\pi s^2}+o(de^{-\pi s^2})} = 1 + O(de^{-\pi s^2})$$

if $s = \omega(\sqrt{\log d})$. □

Finally, the following lemma gives relatively tight asymptotic bounds on the Gaussian mass $\rho_s(B_d^{\mathbb{Z}}(\ell))$ of the integers points in an ℓ_2-ball when the radius is neither too small nor too large. The former constraint is required so that we can accurately approximate the number of integer points in an ℓ_2-ball (see (1)). The latter is a technical condition used to simplify the calculation and essentially guarantee that we are in the "heavy boundary" regime. This means that most of the Gaussian weight is concentrated on the boundary of $B_d^{\mathbb{Z}}(\ell)$. For sufficiently large radii ℓ, this becomes false as the weight of the boundary decreases with ℓ.

Lemma 7. *Let $\ell = \ell(d)$ and $s = s(\ell)$. If $\ell^2 \leqslant \frac{d(4s^2+\pi-\sqrt{\pi^2+8\pi s^2})}{8\pi}$ and $\ell = \omega(\sqrt{d})$ then, as $d \to \infty$,*

$$\ell^d e^{-\pi\ell^2/s^2}\left(\frac{2\pi e}{d}\right)^{d/2} e^{o(d)} \leqslant \rho_s(B_d^{\mathbb{Z}}(\ell)) \leqslant \ell^{d+2} e^{-\pi\ell^2/s^2}\left(\frac{2\pi e}{d}\right)^{d/2} e^{o(d)}.$$

Proof. To simplify the proof, we will assume that ℓ is an integer. This can be formalised by replacing ℓ by $\lfloor\ell\rfloor$, which is asymptotically equivalent. Then,

$$\begin{aligned}\rho_s(B_d^{\mathbb{Z}}(\ell)) &\leqslant \sum_{k=0}^{\ell^2} e^{-\pi k/s^2} |\{\mathbf{x} \in \mathbb{Z}^d : \|\mathbf{x}\|^2 \leqslant k\}| \\ &\overset{(1)}{\leqslant} \sum_{k=0}^{\ell^2} e^{-\pi k/s^2} \operatorname{vol}(B_d(\sqrt{k} + \tfrac{1}{2}\sqrt{d})) = \operatorname{vol}(B_d)\sum_{k=0}^{\ell^2} e^{-\pi k/s^2}(\sqrt{k} + \tfrac{1}{2}\sqrt{d})^d.\end{aligned}$$

Let $f(t) = e^{-\pi t/s^2}(\sqrt{t} + \frac{1}{2}\sqrt{d})^d$ be defined for any $t > 0$. A rather tedious analysis shows that f is increasing over $[0, t_0]$ where

$$t_0 = \frac{d(4s^2 + \pi - \sqrt{\pi^2 + 8\pi s^2})}{8\pi}.$$

It follows that if $\ell^2 \leqslant t_0$ and $\ell = \omega(\sqrt{d})$ then as $d \to \infty$

$$\rho_s(B_d^{\mathbb{Z}}(\ell)) \leqslant (1+\ell^2) f(\ell^2) \operatorname{vol}(B_d) = \ell^2 e^{-\pi\ell^2/s^2} \ell^d e^{o(d)} \frac{1}{\sqrt{d\pi}} \left(\frac{2\pi e}{d}\right)^{d/2}.$$

The lower bound is simpler since

$$\begin{aligned}
\rho_s(B_d^{\mathbb{Z}}(\ell)) &= \sum_{k=0}^{\ell^2} e^{-\pi k/s^2} |\{\mathbf{x} \in \mathbb{Z}^d : \|\mathbf{x}\|^2 = k\}| \\
&\geqslant e^{-\pi\ell^2/s^2} \sum_{k=0}^{\ell^2} |\{\mathbf{x} \in \mathbb{Z}^d : \|\mathbf{x}\|^2 = k\}| = e^{-\pi\ell^2/s^2} |\operatorname{vol}(B_d^{\mathbb{Z}}(\ell))| \\
&\geqslant e^{-\pi\ell^2/s^2} \operatorname{vol}(B_d(\ell - \tfrac{1}{2}\sqrt{d})) = e^{-\pi\ell^2/s^2} (\ell - \tfrac{1}{2}\sqrt{d})^d \operatorname{vol}(B_d)
\end{aligned}$$

and the rest of the computation is analogous. □

4 On the Gaussian mass of random q-ary lattices

In this section, we give probabilistic estimates on the value of $\rho_s(L)$ when L is a random q-ary lattice (see Section 2.3 for more details). These bounds also allow us to bound the smoothing parameter of random q-ary lattices.

A similar result was shown for "standard" random q-ary lattices (i.i.d. from uniform entries) in [LLBS14, Lemma 3] but only gives the expected value, whereas we also bound the variance. A much weaker bound on the variance is given in [PS25]. A closely related result is available in [KNSW20] which studies matrices with each entry independently and identically distributed from an integer Gaussian distribution. Similarly, [CPS+20, Appendix A], [LPR13, § 7] and [SS11, Theorem 2] analyzes the Gaussian mass of a random q-ary lattice over cyclotomic fields. Finally, a very similar result appeared in [PS26] for real random lattices.

4.1 Mass of the entire lattice

The following gives a bound on the expected value and variance of the Gaussian of a random q-ary lattices. We follow the same approach as [PS25, Lemma 4] by using Theorem 2 but we obtain a much stronger bound using the results of the previous section, see Remark 4 for more details.

Lemma 8. *For any $d \in \mathbb{N}$, $1 \leqslant k \leqslant d$, an odd prime number q and $s > 0$, it holds that*

$$\begin{aligned}
\mathbb{E}_{L \sim \mathcal{L}^{\perp}_{d,k,q}}[\rho_s(L)] &= \rho_{s/q}(\mathbb{Z}^d) + q^{k-d}\rho_s(\mathbb{Z}^d \setminus q\mathbb{Z}^d), \\
\mathbb{V}_{L \sim \mathcal{L}^{\perp}_{d,k,q}}[\rho_s(L)] &\leqslant q^{1+k-d}\rho_{s\sqrt{2}/q}(\mathbb{Z}^d)\rho_{s/\sqrt{2}}(\mathbb{Z}^d), \\
\mathbb{V}_{L \sim \mathcal{L}^{\perp}_{d,k,q}}[\rho_s(L)] &\geqslant (q^{k-d} - q^{2(k-d)})X
\end{aligned}$$

where

$$X = \rho_{s\sqrt{2}}(\sqrt{d})\rho_{s\sqrt{2}/q}(\mathbb{Z}^d)\rho_{s/\sqrt{2}}(\mathbb{Z}^d) - (q-1)\rho_{s/q}(\mathbb{Z}^d)^2.$$

Proof. We use Theorem 2 to get that

$$\mathbb{E}_{L \sim \mathcal{L}^{\perp}_{d,k,q}}[\rho_s(L)] = \rho_s(q\mathbb{Z}^d) + q^{k-d}\rho_s(\mathbb{Z}^d \setminus q\mathbb{Z}^d)$$

To estimate the variance, we use Theorem 2 again to get that

$$\mathbb{V}_{L \sim \mathcal{L}^{\perp}_{d,k,q}}[\rho_s(L)] = (q^{k-d} - q^{2(k-d)}) \sum_{\mathbf{x} \in \mathbb{Z}^d \setminus q\mathbb{Z}^d} \sum_{\mathbf{u} \in \mathbb{Z}^d} \sum_{\alpha \in \mathbb{Z}_q \setminus \{0\}} \rho_s(\mathbf{x})\rho_s(\alpha\mathbf{x} + q\mathbf{u})$$

$$= (q^{k-d} - q^{2(k-d)}) \sum_{\mathbf{x}\in\mathbb{Z}^d\setminus q\mathbb{Z}^d} \sum_{\alpha\in\mathbb{Z}_q\setminus\{0\}} \rho_s(\mathbf{x})\rho_s(\alpha\mathbf{x}+q\mathbb{Z}^d)$$
$$= (q^{k-d} - q^{2(k-d)})(A-B)$$

where

$$A = \sum_{\mathbf{x}\in\mathbb{Z}^d} \sum_{\alpha\in\mathbb{Z}_q\setminus\{0\}} \rho_s(\mathbf{x})\rho_s(\alpha\mathbf{x}+q\mathbb{Z}^d), \qquad B = \sum_{\mathbf{x}\in q\mathbb{Z}^d} \sum_{\alpha\in\mathbb{Z}_q\setminus\{0\}} \rho_s(\mathbf{x})\rho_s(\alpha\mathbf{x}+q\mathbb{Z}^d).$$

By Lemma 2, we have that

$$A \leqslant (q-1)\rho_{s\sqrt{2}}(q\mathbb{Z}^d)\rho_{s/\sqrt{2}}(\mathbb{Z}^d).$$

The upper bound follows since $B \geqslant 0$ and $q^{k-d} - q^{2(k-d)} \leqslant q^{k-d}$. For the lower bound, first note that by Remark 1, we have (by keeping only the term $\alpha = 1$)

$$A \geqslant \rho_{s\sqrt{2}}(\sqrt{d})\rho_{s\sqrt{2}}(q\mathbb{Z}^d)\rho_{s/\sqrt{2}}(\mathbb{Z}^d).$$

Furthermore,

$$B \leqslant (q-1)\sum_{\mathbf{x}\in q\mathbb{Z}^d} \rho_s(\mathbf{x})\rho_s(q\mathbb{Z}^d) = (q-1)\rho_{s/q}(\mathbb{Z}^d)^2.$$

This concludes for the lower bound. □

Remark 3. If $1 \leqslant s \leqslant q$ then by Lemma 6, $\mathbb{E}_{L\sim\mathcal{L}^\perp_{d,k,q}}[\rho_s(L)] \approx 1 + q^{k-d}s^d$. Furthermore, $\det(L) = q^{d-k}$ for almost all $L \sim \mathcal{L}^\perp_{d,k,q}$ (Sections 2.2 and 2.3). Hence, we recover a similar result to that of real lattices [PS26], namely that

$$\rho_s(L) \approx 1 + s^d/\mathrm{vol}(L).$$

Remark 4. [PS24, Lemma 4] bounds the same variance by $q^{1+k-d}\rho_s(q\mathbb{Z}^d)\rho_s(\mathbb{Z}^d)$. To see why Lemma 8 is a significant improvement, observe that for the regime of interest where q is much larger than s but s is not small (here $s \geqslant \sqrt{2}$ suffices), we have $\rho_s(q\mathbb{Z}^d)\rho_s(\mathbb{Z}^d) \approx s^d$, whereas

$$\rho_{s\sqrt{2}}(q\mathbb{Z}^d)\rho_{s/\sqrt{2}}(\mathbb{Z}^d) \approx 2^{-d/2}\cdot s^d.$$

This allows us to significantly improve the probabilistic bound on ρ_s of a random q-ary lattice, as shown in the corollary below.

Remark 5. The upper bound on the variance is essentially optimal when $1 \ll s \ll q$ and q^{d-k} is large. Indeed, in this case $q^{k-d} - q^{2(k-d)} \approx q^{k-d}$ and $X \approx \rho_{s\sqrt{2}/q}(\mathbb{Z}^d)\rho_{s/\sqrt{2}}(\mathbb{Z}^d)$ which matches the upper bound up to a factor q.

Corollary 2. *For any $d \in \mathbb{N}$, $1 \leqslant k \leqslant d$, an odd prime number q, $\alpha > 0$ and $s > 0$, it holds that*

$$\Pr_{L\sim\mathcal{L}^\perp_{d,k,q}}[|\rho_s(L) - \mu| > \alpha] \leqslant \frac{q^{1+k-d}\rho_{s\sqrt{2}/q}(\mathbb{Z}^d)\rho_{s/\sqrt{2}}(\mathbb{Z}^d)}{\alpha^2}$$

where $\mu = \rho_{s/q}(\mathbb{Z}^d) - q^{k-d}\rho_s(\mathbb{Z}^d \setminus q\mathbb{Z}^d)$.

Proof. This is a direct application of Lemma 8 and Chebyshev's inequality. □

A direct consequence of these results is that we can obtain a bound on the smoothing parameter of a random lattice. Note here that since the definition of the smoothing parameter concerns the dual of the lattice, the result is on the distribution $\mathcal{L}_{d,k,q}$. This distribution is known to be close to $\mathcal{L}^\perp_{d,k,q}$ when $1 \ll k \ll d$ [PS24, Lemma 5].

Lemma 9. *Let $d \in \mathbb{N}$, $1 \leqslant k \leqslant d$, an odd prime number q, $s > 0$ and $\varepsilon > 0$. Define $\alpha = 1 + \varepsilon - \rho_{1/s}(\mathbb{Z}^d) - q^{-k}\rho_{q/s}(\mathbb{Z}^d \setminus q\mathbb{Z}^d)$. For $\alpha > 0$,*

$$\Pr_{L\sim\mathcal{L}_{d,k,q}}[\eta_\varepsilon(L) > s] \leqslant \frac{q^{1-k}\rho_{\sqrt{2}/s}(\mathbb{Z}^d)\rho_{q/s\sqrt{2}}(\mathbb{Z}^d)}{\alpha^2}.$$

Proof. Recall that $\eta_\varepsilon(L) > s$ if and only $\rho_{1/s}(\widehat{L}) > 1 + \varepsilon$. Furthermore, if $L \sim \mathcal{L}_{d,k,q}$ then $\widehat{L} \sim \frac{1}{q}\mathcal{L}^\perp_{d,d-k,q}$. Therefore,

$$\Pr_{L\sim\mathcal{L}_{d,k,q}}[\eta_\varepsilon(L) > s] = \Pr_L\Big[\rho_{1/s}(\tfrac{1}{q}L) > 1 + \varepsilon\Big] = \Pr_L\big[\rho_{q/s}(L) > 1 + \varepsilon\big].$$

By Corollary 2, for any $\alpha > 0$ and $\sigma > 0$,

$$\begin{aligned}\Pr_{L\sim\mathcal{L}^\perp_{d,d-k,q}}\big[\rho_\sigma(L) > \alpha + \rho_{\sigma/q}(\mathbb{Z}^d) + q^{-k}\rho_\sigma(\mathbb{Z}^d \setminus q\mathbb{Z}^d)\big]&\\ \leqslant \frac{q^{1-k}\rho_{\sigma\sqrt{2}/q}(\mathbb{Z}^d)\rho_{\sigma/\sqrt{2}}(\mathbb{Z}^d)}{\alpha^2}&.\end{aligned}$$

Therefore, by our assumption, we can let $\sigma = q/s$ and $\alpha = 1+\varepsilon-\rho_{\sigma/q}(\mathbb{Z}^d)-q^{k-d}\rho_\sigma(\mathbb{Z}^d\setminus q\mathbb{Z}^d)$ to get the result. □

The previous lemma is difficult to use due to the nontrivial quantities involved in the bound. The following corollary gives an asymptotic bound under reasonable conditions (see Remark 6).

Corollary 3. *Let $k = k(d)$, $q = q(d)$, $\varepsilon = \varepsilon(n)$. Assume that $q^{1-k/d}\varepsilon^{2/d} = \omega(\sqrt{\log d})$ and $q^{k/d}\varepsilon^{-2/d} = \omega(\sqrt{\log d})$. Let $s_\varepsilon = q^{1-k/d}\left(\frac{2\varepsilon+q+\sqrt{4\varepsilon q+q^2}}{2\varepsilon^2}\right)^{1/d}$. Then, it holds that*

$$\Pr_{L\sim\mathcal{L}_{d,k,q}}[\eta_\varepsilon(L) > s_\varepsilon] \leqslant 2^{-d/2}(1 + o(1)).$$

Remark 6. In a typical setting where $q = \mathsf{poly}(d)$ and $1 \ll k \ll d$, we trivially have $q^{1-k/d} = \omega(\sqrt{\log d})$ and $q^{k/d} = \omega(\sqrt{\log d})$. Therefore the technical conditions hold true as soon as $\varepsilon^{2/d} = \Omega(1)$ and $\varepsilon^{-2/d} = \Omega(1)$ which in particular is true when both $\varepsilon(d)$ and $1/\varepsilon(d)$ are at most subexponential in d, a rather weak condition. A much better analysis is possible with precise parameter values.

Proof. Let $s = s(n)$ be such that $s = \omega(\sqrt{\log d})$ and $q/s = \omega(\sqrt{\log d})$. By our assumption that $s = \omega(\sqrt{\log d})$ and $q/s = \omega(\sqrt{\log d})$, together with Lemma 6, it follows that

$$\begin{aligned}\rho_{1/s}(\mathbb{Z})^d &= 1 + O(de^{-\pi s^2}), & \rho_{\sqrt{2}/s}(\mathbb{Z})^d &= 1 + O(de^{-\pi s^2/2}),\\ \rho_{q/s}(\mathbb{Z})^d &= q^d s^{-d}(1 + o(1)), & \rho_{q/s\sqrt{2}}(\mathbb{Z})^d &= q^d 2^{-d/2} s^{-d}(1 + o(1)).\end{aligned}$$

Therefore, using the notation of Lemma 9

$$\begin{aligned}\alpha &= 1 + \varepsilon - \rho_{1/s}(\mathbb{Z}^d) - q^{-k}\rho_{q/s}(\mathbb{Z}^d \setminus q\mathbb{Z}^d)\\ &= \varepsilon - O(de^{-\pi s^2}) - q^{d-k}s^{-d}(1 + o(1)) = \tilde{\varepsilon} - 1/x\end{aligned}$$

where $\tilde{\varepsilon} = \varepsilon + O(de^{-\pi s^2})$ and $x = q^{k-d}s^d(1 + o(1))$. Furthermore, if $\alpha > 0$ then $\Pr_{L\sim\mathcal{L}_{d,k,q}}[\eta_\varepsilon(L) > s] \leqslant p$ where

$$p = \frac{q^{1-k}\rho_{\sqrt{2}/s}(\mathbb{Z}^d)\rho_{q/s\sqrt{2}}(\mathbb{Z}^d)}{\alpha^2} = \frac{q^{1-k}(1 + o(1))q^d 2^{-d/2}(1 + o(1))}{s^d\alpha^2}$$

$$= \frac{2^{-d/2} q(1+o(1))}{x(\tilde{\varepsilon} - 1/x)^2}.$$

Solving for $x(\tilde{\varepsilon} - 1/x)^2 = q$, we get that the only solution meaningful solution is $x = (2\tilde{\varepsilon} + q + \sqrt{4\tilde{\varepsilon}q + q^2})/2\tilde{\varepsilon}^2$, assuming that $\tilde{\varepsilon} > 0$. From this, we get that $s = q^{1-k/d}x^{1/d}$ which verifies our initial assumptions that $s = \omega(\sqrt{\log d})$ and $q/s = \omega(\sqrt{\log d})$ by the assumptions of the lemma. Therefore $p \leqslant 2^{-d/2}(1 + o(1))$. □

Remark 7. This result is very similar to that of [PS26, Corollary 3]. Indeed, the extra $q^{1-k/d}$ factor in the definition of s_ε accounts for the volume q^{d-k} of almost all q-ary lattices in $\mathcal{L}_{d,k,q}$, but the second part is otherwise almost identical. In particular, for similar reasons to [PS26, Section 2], we expect this bound to be quite good for constant ε but not asymptotically optimal as $\varepsilon \to 0$. More precisely, our asymptotic will be $q^{1-k/d}\varepsilon^{-2/d}$ when we expect $q^{1-k/d}\varepsilon^{-1/d}$.

4.2 Gaussian mass of lattice points inside a ball

For the application to SIS, we not only require a bound on $\rho_s(L)$ but also a bound on $\rho_s(L \cap B_n^p(\ell))$, for any norm $\ell \in [1, \infty]$. The following lemma follows a similar strategy for this quantity. Due to some technical limitations in our proof, we can only handle radii ℓ up to $q/2$, so this lemma does not allow us to recover the result on $\rho_s(L)$ by letting ℓ tend to infinity. The main result, Corollary 4, informally shows that for a random lattice $L \sim \mathcal{L}^\perp_{d,k,q}$,

$$\rho_s(L \cap B_d^p(\ell)) \approx 1 + q^{k-d}\rho_s(B_d^{p,\mathbb{Z}}(\ell) \setminus \{0\})$$

and quantifies the deviation from the average.

Lemma 10. *For any $p \in [1, \infty]$, $d \in \mathbb{N}$, $1 \leqslant k \leqslant d$, an odd prime number q, $0 \leqslant \ell \leqslant q/2$ and $s > 0$,*

$$\begin{aligned}
&\mathbb{E}_{L\sim\mathcal{L}^\perp_{d,k,q}}[\rho_s(L \cap B_d^p(\ell))] = 1 + q^{k-d}\rho_s(B_d^{p,\mathbb{Z}}(\ell) \setminus \{0\}),\\
&\mathbb{V}_{L\sim\mathcal{L}^\perp_{d,k,q}}[\rho_s(L \cap B_d^p(\ell))] \leqslant (q-1)q^{k-d}\rho_s(B_d^{p,\mathbb{Z}}(\ell) \setminus \{0\}),\\
&\mathbb{V}_{L\sim\mathcal{L}^\perp_{d,k,q}}[\rho_s(L \cap B_d^p(\ell))] \geqslant (q^{k-d} - q^{2(k-d)})\rho_s(B_d^{p,\mathbb{Z}}(\ell) \setminus \{0\}).
\end{aligned}$$

Proof. We use Theorem 2 to get that

$$\begin{aligned}
\mathbb{E}_{L\sim\mathcal{L}^\perp_{d,k,q}}[\rho_s(L \cap B_d^p(\ell))] &= \rho_s(q\mathbb{Z}^d \cap B_d^p(\ell)) + q^{k-d}\rho_s((\mathbb{Z}^d \setminus q\mathbb{Z}^d) \cap B_d^p(\ell))\\
&= 1 + q^{k-d}\rho_s(B_d^{p,\mathbb{Z}}(\ell) \setminus \{0\})
\end{aligned}$$

where we have used that $\ell \leqslant q$ to conclude that $q\mathbb{Z}^d \cap B_d^p(\ell) = \{0\}$ and $(\mathbb{Z}^d \setminus q\mathbb{Z}^d) \cap B_d^p(\ell) = (\mathbb{Z}^d \setminus \{0\}) \cap B_d^p(\ell) = B_d^{p,\mathbb{Z}}(\ell) \setminus \{0\}$. To estimate the variance, let $f(\mathbf{x}) = \rho_s(\mathbf{x})\mathbb{1}[\mathbf{x} \in B_d^p(\ell)]$. We use Theorem 2 again to get that

$$\begin{aligned}
&\mathbb{V}_{L\sim\mathcal{L}^\perp_{d,k,q}}[\rho_s(L \cap B_d^p(\ell))]\\
&= (q^{k-d} - q^{2(k-d)}) \sum_{\mathbf{x}\in\mathbb{Z}^d\setminus q\mathbb{Z}^d} \sum_{\mathbf{u}\in\mathbb{Z}^d} \sum_{\alpha\in\mathbb{Z}_q\setminus\{0\}} f(\mathbf{x})f(\alpha\mathbf{x} + q\mathbf{u})\\
&\leqslant q^{k-d} \sum_{\alpha\in\mathbb{Z}_q\setminus\{0\}} \sum_{\substack{\mathbf{x}\in\mathbb{Z}^d\setminus q\mathbb{Z}^d:\\ \mathbf{x}\in B_d^p(\ell)}} \sum_{\substack{\mathbf{u}\in\mathbb{Z}^d:\\ \alpha\mathbf{x}+q\mathbf{u}\in B_d^p(\ell)}} \rho_s(\mathbf{x})\rho_s(\alpha\mathbf{x} + q\mathbf{u})\\
&= q^{k-d} \sum_{\alpha\in\mathbb{Z}_q\setminus\{0\}} \big(F(\alpha, \mathbb{Z}^d) - F(\alpha, q\mathbb{Z}^d)\big)
\end{aligned}$$

where

$$F(\alpha, X) = \sum_{\substack{\mathbf{x}\in X:\\ \mathbf{x}\in B_d^p(\ell)}} \sum_{\substack{\mathbf{u}\in\mathbb{Z}^d:\\ \alpha\mathbf{x}+q\mathbf{u}\in B_d^p(\ell)}} \rho_s(\mathbf{x})\rho_s(\alpha\mathbf{x}+q\mathbf{u}).$$

By Lemma 4, $F(\alpha, \mathbb{Z}^d) \leqslant \rho_{s/\sqrt{2}}(B_d^{p,\mathbb{Z}}(\ell))$ for any $\alpha \in \mathbb{Z}_q \setminus \{0\}$ since $\ell < q/2$. Furthermore, $F(\alpha, q\mathbb{Z}^d) = 1$ since the only term in the sum is $\mathbf{x} = \mathbf{u} = 0$ due to the constraint that $\ell \leqslant q$. Finally, we get

$$\mathbb{V}_{L\sim\mathcal{L}_{d,k,q}^{\perp}}[\rho_s(L \cap B_d^p(\ell))] \leqslant q^{k-d}(q-1)\left(\rho_{s/\sqrt{2}}(B_d^{p,\mathbb{Z}}(\ell)) - 1\right).$$

For the lower bound, by Remark 2, we have that $F(1, \mathbb{Z}^d) \geqslant \rho_{s/\sqrt{2}}(B_d^{p,\mathbb{Z}}(\ell))$ and

$$\begin{aligned}
\mathbb{V}_{L\sim\mathcal{L}_{d,k,q}^{\perp}}[\rho_s(L \cap B_d^p(\ell))] &= (q^{k-d} - q^{2(k-d)}) \sum_{\alpha\in\mathbb{Z}_q\setminus\{0\}} \left(F(\alpha, \mathbb{Z}^d) - F(\alpha, q\mathbb{Z}^d)\right) \\
&= (q^{k-d} - q^{2(k-d)}) \sum_{\alpha\in\mathbb{Z}_q\setminus\{0\}} \left(F(\alpha, \mathbb{Z}^d) - 1\right) \\
&\geqslant (q^{k-d} - q^{2(k-d)})(F(1, \mathbb{Z}^d) - 1)) \\
&\geqslant (q^{k-d} - q^{2(k-d)})\left(\rho_{s/\sqrt{2}}(B_d^{p,\mathbb{Z}}(\ell)) - 1\right).
\end{aligned}$$

□

Remark 8. The upper bound on the variance is essentially optimal when q^{d-k} is large since the lower bound matches the upper bound, up to a factor of q.

Corollary 4. *For any $d \in \mathbb{N}$, $1 \leqslant k \leqslant d$, an odd prime number q, $0 \leqslant \ell < q/2$ and $s > 0$,*

$$\Pr_{L\sim\mathcal{L}_{d,k,q}^{\perp}}[|\rho_s(L \cap B_d^p(\ell)) - 1 - \mu_{s,\ell}| > \alpha] \leqslant q \cdot \frac{\mu_{s,\ell}}{\alpha^2}.$$

where $\mu_{s,\ell} = q^{k-d}\rho_s(B_d^{p,\mathbb{Z}}(\ell) \setminus \{0\})$.

Proof. This is a direct application of Lemma 10 and Chebyshev's inequality. □

5 Application to the SIS problem

In this section, we apply the results of the previous section to solve the SIS problem. The algorithm is simple: given access to a Gaussian sampler, we sample many vectors and estimate the probability of obtaining a short enough vector. The analysis, however, is not trivial and requires estimating, in particular, the Gaussian mass of a random q-ary lattice and the Gaussian mass of a random q-ary lattice intersected with a ball. An interesting side-effect of this approach is that we can handle an arbitrary ℓ_p-norm, including $p = \infty$.

Before going to the main result, it is worth making some remarks on the parameters. Consider $\mathrm{SIS}_{d-k,d,q,\ell}^p$ for some $d, q \in \mathbb{N}$, $1 \leqslant k \leqslant d$ and $\ell > 0$. This is equivalent (Section 2.4) to finding a nonzero vector in $L \cap B_n^p(\ell)$ with $L \sim \mathcal{L}_{d,k,q}^{\perp}$. Most such lattices satisfy [PS24, Corollary 2] that

$$\lambda_1(L) \approx \mathrm{GH}(L) = \mathrm{vol}(L)^{1/d}\sqrt{\tfrac{d}{2\pi e}} = q^{1-k/d}\sqrt{\tfrac{d}{2\pi e}}$$

since $\mathrm{vol}(L) = q^{d-k}$ with probability close to 1. In particular, this means that the $\mathrm{SIS}_{d-k,d,q,\ell}^p$ problem only makes sense for

$$\ell = \alpha q^{1-k/d}\sqrt{\tfrac{d}{2\pi e}} \quad \text{and} \quad \alpha \geqslant 1. \tag{8}$$

In cryptographic applications, we do not expect α to be too close to 1 since this would essentially be equivalent to the SVP problem. Nevertheless, α can still be a relatively small constant such as in the Kyber and Dilithium cryptosystems (Section 5.2) where α is around 20 depending on the security level. It will also be useful to note that in all cryptographic applications, q usually grows at least polynomially with d and $1 \ll k \ll d$. For our analysis, we will only require the much weaker conditions, i.e.,

$$q^{k/d} = \omega(\sqrt{\log d}), \qquad q^{1-k/d} = \omega(\sqrt{\log d}). \tag{9}$$

Our algorithm works by sampling vectors from the discrete Gaussian distribution until we get one that is short enough. Since sampling is harder for smaller values of s, we do not want the value of s to be too small. On the other hand, a higher value of s means that samples will be longer on average, hence increasing the expected number of samples before obtaining a sufficiently short one. Therefore, there is a nontrivial trade-off between the complexity of the sampler and the number of samples. Recall that the expected length of such a vector sampled from the discrete Gaussian at width s is $s\sqrt{\frac{d}{2\pi}}$ [MR04, Lemma 4.3]. We want this quantity to be at least equal to ℓ, so by (8), we always want to take

$$s \geqslant \frac{\alpha}{\sqrt{e}} q^{1-k/d}. \tag{10}$$

5.1 General result

We provide next a general statement on the complexity of the algorithm and its success probability. Informally, our result says that if we have access to a Gaussian sampler for a random lattice $L \sim \mathcal{L}^{\perp}_{d,k,q}$ at width s, then there is an algorithm which solves $\mathrm{SIS}^{p}_{d-k,d,q,\ell}$ and whose complexity essentially depends on the Gaussian weight of the ℓ_p-ball of radius ℓ, i.e., $\rho_s(B_d^{p,\mathbb{Z}}(\ell))/\rho_s(\mathbb{Z}^d)$.

The formal statement below is significantly more complicated and an informal explanation of the quantities appearing in the theorem is available below in Remark 9 and to derive the simplified Theorem 1. We also give asymptotic statements for $p = 2$ and $p = \infty$ after the remark.

Theorem 5 (General result for SIS in any norm)**.** *There is a randomized algorithm that for every $p \in [1, \infty]$, $d \in \mathbb{N}$, $1 \leqslant k \leqslant d$, odd prime number q, $s > 0$ and $0 \leqslant \ell < q/2$, and on a fraction at least $1 - \varepsilon_{s,\ell} - \frac{\theta_s}{\nu_s^2}$ of random lattices $L \sim \mathcal{L}^{\perp}_{d,k,q}$, outputs in time $N\mathsf{poly}(d)$ and polynomial space a nonzero vector of L of ℓ_p-norm at most ℓ with probability at least $1/2$, where*

$$N = \Omega(1/x), \qquad x = \frac{\mu_{s,\ell}}{4\nu_s},$$
$$\nu_s = \rho_{s/q}(\mathbb{Z}^d) + q^{k-d}\rho_s(\mathbb{Z}^d \setminus q\mathbb{Z}^d), \qquad \varepsilon_{s,\ell} = \tfrac{4q}{\mu_{s,\ell}},$$
$$\theta_s = q^{1+k-d}\rho_{s\sqrt{2}/q}(\mathbb{Z}^d)\rho_{s/\sqrt{2}}(\mathbb{Z}^d), \qquad \mu_{s,\ell} = q^{k-d}\rho_s(B_d^{p,\mathbb{Z}}(\ell) \setminus \{0\}).$$

This algorithm requires access to a discrete Gaussian sampler for L at width s and makes N queries to the sampler. In particular, this algorithm solves $\mathrm{SIS}^{p}_{d-k,d,q,\ell}$ on a fraction at least $1 - \varepsilon_{s,\ell} - \theta_s/\nu_s^2$ of random lattices.

Remark 9 (Informal description). Let us describe some of the quantities appearing above and give some intuitions about their behavior.

- $\mu_{s,\ell}$ is the Gaussian mass of the ball of radius ℓ multiplied by q^{k-d}. For large enough ℓ, this quantity grows as $\ell^d q^{k-d}$ and $\ell \geqslant \sqrt{\frac{d}{2\pi e}} q^{1-k/d}$ by (8), making $\mu_{s,\ell}$ super-exponentially large in d.

- $\varepsilon_{s,\ell}$ measures the probability that a particular lattice L is "bad" and that the algorithm will fail on it. By the above reasoning for $\mu_{s,\ell}$, we expect $\varepsilon_{s,\ell}$ to be negligible.
- ν_s is the expected value of the Gaussian mass of a random q-ary lattice. By (10), s will never be small so by Lemma 6, ν_s is asymptotically equal to $q^{k-d}\rho_s(\mathbb{Z}^d)$ which is at least $q^{k-d}s^d \geqslant (\frac{\alpha}{\sqrt{e}})^d$. Therefore ν_s is expected to be exponentially large as soon as $\alpha > \sqrt{e}$.
- $\frac{\theta_s}{\nu_s^2}$ also measures the probability that a particular lattice L is "bad" and that the algorithm will fail on it. By the above reasoning, ν_s is asymptotically close to $q^{k-d}s^d$ and $\rho_{s/\sqrt{2}}(\mathbb{Z}^d)$ is close to $s^d 2^{-d/2}$ by a similar reasoning. Therefore $\frac{\theta_s}{q\nu_s^2}$ is asymptotically at most $\frac{2^{-d/2}}{q^{k-d}s^d}\rho_{s\sqrt{2}/q}(\mathbb{Z}^d)$. There are two regimes for this quantity:
  - If $s \geqslant q$ then, by Lemma 6, this quantity is close to q^{-k} which is super-exponentially small by (9).
  - Otherwise, by Lemma 6, this quantity is close to $\frac{2^{-d/2}}{q^{k-d}s^d} \leqslant \alpha^{-d}$ by (10) which we therefore expect to be exponentially small.

  Hence is all cases, we expect $\frac{\theta_s}{\nu_s^2}$ to be negligible.

By putting everything together, we obtain Theorem 1 which is a simplified version of Theorem 5.

Proof. Let N to be fixed later. Consider the following algorithm:

- Sample N vectors independently according to $D_{L,s}$.
- Return the shortest nonzero vector.

We will analyze this algorithm. First the running time is clearly $N\mathsf{poly}(d)$ with N queries to the sampler, and it only needs to store the shortest vector found so far. It remains to analyze the correctness.

Let $\mu_{s,\ell} = q^{k-d}\rho_s(B_d^{p,\mathbb{Z}}(\ell) \setminus \{0\})$. By Corollary 4 for $\alpha = \frac{\mu_{s,\ell}}{2}$, and with probability at least $1 - \varepsilon_{s,\ell}$ where $\varepsilon_{s,\ell} = q\frac{\mu_{s,\ell}}{\alpha^2} = \frac{4q}{\mu_{s,\ell}}$ over the choice of L we have

$$|\rho_s(L \cap B_d^p(\ell)) - 1 - \mu_{s,\ell}| < \alpha$$

so in particular

$$\rho_s(L \cap B_d^p(\ell)) \geqslant 1 + \mu_{s,\ell} - \alpha = 1 + \frac{\mu_{s,\ell}}{2}. \tag{11}$$

Let $\nu_s = \rho_{s/q}(\mathbb{Z}^d) + q^{k-d}\rho_s(\mathbb{Z}^d \setminus q\mathbb{Z}^d)$ and $\theta_s = q^{1+k-d}\rho_{s\sqrt{2}/q}(\mathbb{Z}^d)\rho_{s/\sqrt{2}}(\mathbb{Z}^d)$. By Corollary 2 for $\beta = \nu_s$, and with probability at least $1 - \frac{\theta_s}{\beta^2}$ over the choice of L we have $|\rho_s(L) - \nu_s| < \nu_s$ so in particular

$$\rho_s(L) \leqslant 2\nu_s. \tag{12}$$

By a union-bound, we get that (11) and (12) hold simulatenously with probability at least $1 - \varepsilon_{s,\ell} - \frac{\theta_s}{\nu_s^2}$. In particular, we have that

$$\frac{\rho_s((L \setminus \{0\}) \cap B_d^p(\ell))}{\rho_s(L)} \geqslant x := \frac{\mu_{s,\ell}}{4\nu_s}.$$

We now analyze the algorithm. Since the samples are independent, the probability that each sample is nonzero and of length at most ℓ is x. Hence, if $N = \Omega(1/x)$, then with constant probability there is a nonzero vector in $L \cap B_d^p(\ell)$. □

The previous result is not easy to use because it depends on many nontrivial quantities to estimate, such as $\rho_{s/\sqrt{2}}(\mathbb{Z}^d)$. For precise values of all parameters, these quantities can be evaluated with good precision, but it makes it difficult to understand the behavior in general. The next theorems give asymptotic values for the parameters in a regime of interest of cryptography, for $p = 2$ and $p = \infty$, which are the two cases of most practical interest.

Theorem 6 (Asymptotic complexity for $p = \infty$)**.** *Let $q := q(d)$, $k := k(d)$, $s := s(d)$ and $\ell := \ell(d)$ be such that $s = \omega(\sqrt{\log d})$, $q/s = \omega(\sqrt{\log d})$, $s > q^{1-k/d}$, $q = 2^{o(d)}$, and $\ell = \omega(1)$ as $d \to \infty$. Then the algorithm from Theorem 5 solves $\mathrm{SIS}^{\infty}_{d-k,d,q,\ell}$ on a fraction at least $1 - \frac{q^{1+d-k}}{s^d \operatorname{erf}\left(\frac{\ell\sqrt{\pi}}{s}\right)^d e^{o(d)}} - 2^{-d/2+o(d)}$ of random lattices $L \sim \mathcal{L}^{\perp}_{d,k,q}$ using $N = \frac{e^{o(d)}}{\operatorname{erf}\left(\frac{\ell\sqrt{\pi}}{s}\right)^d}$ calls to a Gaussian sampler for L at width s.*

Remark 10. For similar reasons to Remark 9, all technical conditions except $q/s = \omega(\sqrt{\log d})$ are trivially satisfied for all cryptographic applications.

Proof. In what follows, we use notations from Theorem 5. By our assumption that $s = \omega(\sqrt{d})$ and that $q/s = \omega(\sqrt{\log d})$ and Lemma 6 it follows that

$$\rho_s(\mathbb{Z})^d = s^d(1+o(1)), \qquad \rho_{s/\sqrt{2}}(\mathbb{Z})^d = s^d 2^{-d/2}(1+o(1)),$$
$$\rho_{s/q}(\mathbb{Z})^d = 1+o(1), \qquad \rho_{s\sqrt{2}/q}(\mathbb{Z})^d = 1+o(1).$$

Therefore

$$\begin{aligned}\theta_s &= q^{1+k-d}\rho_{s\sqrt{2}/q}(\mathbb{Z}^d)\rho_{s/\sqrt{2}}(\mathbb{Z}^d) = q^{1+k-d}2^{-d/2}s^d(1+o(1)), \quad \text{and}\\ \nu_s &= \rho_{s/q}(\mathbb{Z}^d) + q^{k-d}\rho_s(\mathbb{Z}^d \setminus q\mathbb{Z}^d) = 1+o(1)+q^{k-d}s^d(1+o(1))\\ &= q^{k-d}s^d(1+o(1))\end{aligned}$$

where the last equality comes from the fact that $s > q^{1-k/d}$. It follows that

$$\frac{\theta_s}{\nu_s^2} = 2^{-d/2}q(1+o(1)) \leqslant 2^{-d/2+o(d)}$$

since $q = 2^{o(d)}$. Furthermore, by Lemma 5,

$$\begin{aligned}\mu_{s,\ell} &= q^{k-d}\rho_s(B_n^{\infty,\mathbb{Z}}(\ell)\setminus\{0\}) = q^{k-d}\left(\rho_s([-\ell,\ell]\cap\mathbb{Z})^d - 1\right)\\ &\leqslant q^{k-d}\left(\left(1+s\operatorname{erf}\left(\frac{\ell\sqrt{\pi}}{s}\right)\right)^d - 1\right).\end{aligned}$$

We now claim that $s\operatorname{erf}\left(\frac{\ell\sqrt{\pi}}{s}\right) = \Omega(\ell)$. Indeed[§], if $\ell = o(s)$ then $s\operatorname{erf}\left(\frac{\ell\sqrt{\pi}}{s}\right) \sim s\frac{2}{\sqrt{x}}\frac{\ell\sqrt{\pi}}{s} = 2\ell$, and if $\ell = \Omega(s)$ then $s\operatorname{erf}\left(\frac{\ell\sqrt{\pi}}{s}\right) = \Omega(s) = \Omega(\ell)$ by the properties of the erf function. since $\ell = \omega(1)$, it follows that

$$\mu_{s,\ell} = q^{k-d}s^d\operatorname{erf}\left(\frac{\ell\sqrt{\pi}}{s}\right)^d e^{o(d)} \quad \text{and} \quad p = \frac{\mu_{s,\ell}}{4\nu_s} = \operatorname{erf}\left(\frac{\ell\sqrt{\pi}}{s}\right)^d e^{o(d)}.$$

Finally,

$$\varepsilon_{s,\ell} = \frac{4q}{\mu_{s,\ell}} = \frac{q^{1+d-k}}{s^d\operatorname{erf}\left(\frac{\ell\sqrt{\pi}}{s}\right)^d e^{o(d)}}.$$

□

[§]Here to simplify the proof, we are assuming that either $\ell = o(s)$ or $\ell = \Omega(s)$. This is not technically always true but a more formal reasoning can be done if necessary to obtain the same result.

Theorem 7 (Asymptotic complexity for $p = 2$)**.** *Let $q := q(d)$, $k := k(d)$, $s := s(d)$ and $\ell := \ell(d)$ be such that $s = \omega(\sqrt{\log d})$, $q/s = \omega(\sqrt{\log d})$, $s > q^{1-k/d}$, $q = 2^{o(d)}$, $\ell^2 \leqslant \frac{d(s^2+\pi-\sqrt{\pi^2+2\pi s^2})}{2\pi}$ and $\ell = \omega(\sqrt{d})$ as $d \to \infty$. Then the algorithm from Theorem 5 solves $\mathrm{SIS}^2_{d-k,d,q,\ell}$ on a fraction at least $1 - q^{1+d-k}e^{\pi\ell^2/s^2+o(d)}\left(\frac{d}{2\pi e\ell^2}\right)^{d/2}$ of random lattices $L \sim \mathcal{L}^{\perp}_{d,k,q}$ using $N = e^{\pi\ell^2/s^2+o(d)}\left(\frac{ds^2}{2\pi e\ell^2}\right)^{d/2}$ calls to a Gaussian sampler for L at width s.*

Proof. In what follows, we use notations from Theorem 5. By the same reasoning as in Theorem 6, we have that

$$\rho_s(\mathbb{Z})^d = s^d(1+o(1)), \qquad \rho_{s/\sqrt{2}}(\mathbb{Z})^d = s^d 2^{-d/2}(1+o(1)),$$
$$\rho_{s/q}(\mathbb{Z})^d = 1+o(1), \qquad \rho_{s\sqrt{2}/q}(\mathbb{Z})^d = 1+o(1)$$

and

$$\nu_s = q^{k-d}s^d(1+o(1)), \qquad \frac{\theta_s}{\nu_s^2} \leqslant 2^{-d/2+o(d)}.$$

Furthermore, by Lemma 7,

$$\mu_{s,\ell} = q^{k-d}\rho_s(B_n^{2,\mathbb{Z}}(\ell) \setminus \{0\}) \geqslant q^{k-d}\ell^d e^{-\pi\ell^2/s^2}\left(\frac{2\pi e}{d}\right)^{d/2} e^{o(d)}.$$

Therefore

$$p = \frac{\mu_{s,\ell}}{4\nu_s} \geqslant \left(\frac{\ell}{s}\right)^d e^{-\pi\ell^2/s^2}\left(\frac{2\pi e}{d}\right)^{d/2} e^{o(d)}.$$

Finally,

$$\varepsilon_{s,\ell} = \frac{4q}{\mu_{s,\ell}} \leqslant \frac{q}{q^{k-d}\ell^d e^{-\pi\ell^2/s^2}\left(\frac{2\pi e}{d}\right)^{d/2} e^{o(d)}}.$$

□

5.2 Application to Dilithium

The parameter of SIS^∞ for the Dilithium signature scheme [DKL+18] are in Figure 2.

We have applied Theorem 5 to those instances. In order to do so, we use the MCMC Gaussian sampler (Theorem 3) on a β-BKZ reduced basis of the lattice. The complexity of this sampler can be estimated (Theorem 4) and the complexity of BKZ using standard results in the literature. A key factor affecting the choice of the sampler is that we need to be able to sample the discrete Gaussian distribution at a relatively low width s. Samplers such as the Klein sampler run in polynomial time but place a limit to how small s can be depending on the reduction quality of the basis (see Section 2.6). This is not ideal for our attack because it does not give a lot of freedom in the choices of parameters and leads to poor attack complexities. The MCMC sampler, in contrast, allows us to sample at arbitrary values s at the cost of a higher sampling time. This opens up more interesting choices where we can lower s, resulting in a smaller number of required samples (N) but an increased cost of producing each sample.

Since the choice of β and s is not clear, we wrote an optimizer to find the optimal choice in each instance. The code can be found at https://github.com/pamaury/random_sis The results are given in Figure 3. We recall that these estimates rely on the GSA and also ignore polynomial factors, those figures should be taken as rough estimates of the complexity of the problem. Nevertheless, we point out that the GSA is only used to estimate the complexity (and not the correctness) of the sampler, and its use is standard in security estimates.

NIST Level	m	n	q	ℓ
2	$256 \cdot 9$	$256 \cdot 4$	8380417	350209
3	$256 \cdot 12$	$256 \cdot 6$	8380417	724481
5	$256 \cdot 16$	$256 \cdot 8$	8380417	769537

Figure 2: SIS^∞ parameters for the Dilithium cryptosystem.

Level	β	s	$\log_2 N$	$\log_2 T_{\text{sample}}$	$\log_2 T_{\text{BKZ}}$	$\log_2 T_{\text{attack}}$	[DEL25, § A.1]	[APS15]
2	609	346702	39	159	201	**202**	269.9	152.2
3	919	723749	55	230	288	**289**	343.0	210.5
5	1314	769533	74	323	400	**400**	450.2	288.2

Figure 3: Estimated time cost of our attack on SIS^∞ for Dilithium. The T_{sample} column corresponds to the cost of the MCMC sampler (to produce one sample) and T_{BKZ} is the cost of reducing the basis. N is the number of Gaussian samples used in the attack. The total cost of the attack is $T_{\text{attack}} = T_{\text{BKZ}} + N \cdot T_{\text{sample}}$, ignoring polynomial factors. Note that $T_{\text{BKZ}} \approx N \cdot T_{\text{sample}}$, meaning that the cost of BKZ is roughly balanced with the total cost of sampling the vectors. Columns β and s give the parameters of the attacks (BKZ block size and width of the Gaussian sampler). The column [DEL25, § A.1] corresponds to the cost associated with Wagner's algorithm for solving SIS^∞ (see comparison with provable attacks). The column [APS15] corresponds to the estimated cost of the heuristic, sieving-based attack described in [DKL+21, Appendix C.3] and implemented in the lattice estimator (see comparison with heuristic attacks).

5.2.1 Comparison with provable attacks

Except for the use of the GSA in the complexity estimates of the sampler, our algorithm is provably correct. The closest provable attack to our work is [DEL25], which analyses a provable variant of Wagner's algorithm and also gives complexity estimates for Dilithium. The estimates in Figure 3 show an improvement of roughly 50 bits on all security levels compared to [DEL25]. We note that this comparison might be slightly unfair due to our use of the GSA and the fact that Wagner's algorithm is only efficient in more extreme regimes, far away from that of Dilithium [DEL25, § A.2]. Although there are other provable attacks on SIS, we are not aware of any that apply to the regime of the parameter regime of Dilithium.

5.2.2 Comparison with heuristic attacks

As explained above, our algorithm is provably correct. Therefore, we do not expect that such an algorithm can compete with the heuristic of sieving-based algorithms. The security estimates used in the Dilithium specification [DKL+21, Appendix C.3] were obtained using the following attack. First run BKZ on the lattice basis and use lattice sieving for the SVP solver to obtain many short vectors (in the ℓ_2 norm). The difficulty lies in the fact that one needs to estimate the probability that one of those vectors has a short ℓ_∞ norm. To estimate the ℓ_∞ norm, the idea was to model heuristically the coordinates of the short vectors as independent Gaussian of standard deviation σ, which depends on the shape of the basis[¶]. Using the approach, one can therefore estimate the probability that a given

[¶]There is an additional issue, depending on whether the Euclidean norm of the short vectors is bigger than q or not. If this is the case, then we need to account for the q-vectors in the BKZ basis, which complicate the estimates. As noted in [DKL+21, Appendix C.3], this is the case for the parameters of

vector solves the SIS^∞ problem and therefore the number of vectors necessary to mount an attack.

We point out, however, that this heuristic may be overly optimistic. Indeed, the same heuristic was used by [MAT22] to show that dual attacks on LWE are competitive with primal attacks. However, this heuristic was later shown to be wrong [DP23], at least in the context of dual attacks, and several works addressed this issue by using a different, or significantly more complex analysis [PS24, CMHST25, QX25, DP26]. It would therefore be interesting to see if similar techniques can be used to either confirm or disprove the use of this heuristic in sieving for SIS^∞.

Finally, it is worth noting that using DGS to solve SIS may prove advantageous on quantum computers. Indeed, the best quantum algorithms for sieving run in time $2^{0.2563n+o(n)}$ [BCSS23, CL21], compared to $2^{0.292n+o(n)}$ on a classical computer [BDGL16], far from a quadratic speedup. Further, those algorithms require exponential quantum memory. By contrast, some recent work [CSS26] has shown a quadratic[‖] speedup on DGS sampling using only a polynomial number of qubits, and obtained some significant speedup on Dilithium specifically.

References

[ACKS21] Divesh Aggarwal, Yanlin Chen, Rajendra Kumar, and Yixin Shen. Improved (provable) algorithms for the shortest vector problem via bounded distance decoding. In Markus Bläser and Benjamin Monmege, editors, *38th International Symposium on Theoretical Aspects of Computer Science, STACS 2021, March 16-19, 2021, Saarbrücken, Germany (Virtual Conference)*, volume 187 of *LIPIcs*, pages 4:1–4:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPIcs.STACS.2021.4.

[AD21] Martin Albrecht and Léo Ducas. Lattice attacks on NTRU and LWE: A history of refinements. Cryptology ePrint Archive, Paper 2021/799, 2021. URL: https://eprint.iacr.org/2021/799.

[ADRS15] Divesh Aggarwal, Daniel Dadush, Oded Regev, and Noah Stephens-Davidowitz. Solving the shortest vector problem in 2^n time using discrete Gaussian sampling: Extended abstract. In Rocco A. Servedio and Ronitt Rubinfeld, editors, *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015*, pages 733–742. ACM, 2015. doi:10.1145/2746539.2746606.

[ADSD15] Divesh Aggarwal, Daniel Dadush, and Noah Stephens-Davidowitz. Solving the closest vector problem in 2^n time – the discrete Gaussian strikes again! In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pages 563–582, 2015. doi:10.1109/FOCS.2015.41.

[Ajt96] M. Ajtai. Generating hard instances of lattice problems (extended abstract). In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '96, page 99–108, New York, NY, USA, 1996. Association for Computing Machinery. doi:10.1145/237814.237838.

[Ajt98] Miklós Ajtai. The shortest vector problem in L2 is NP-hard for randomized reductions (extended abstract). In *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing*, STOC '98, page 10–19, New York, NY,

Dilithium and the authors suggest, as an alternative, to re-randomize the basis prior to running BKZ to avoid the issue and obtain better estimates.

[‖]Excluding any preprocessing, such as lattice reduction, which can be done classically.

USA, 1998. Association for Computing Machinery. doi:10.1145/276698.276705.

[APS15] Martin R. Albrecht, Rachel Player, and Sam Scott. On the concrete hardness of learning with errors. *Journal of Mathematical Cryptology*, 9:169 – 203, 2015. URL: https://github.com/malb/lattice-estimator.

[Ban93] Wojciech Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 296:625–635, 1993.

[BCSS23] Xavier Bonnetain, André Chailloux, André Schrottenloher, and Yixin Shen. Finding many collisions via reusable quantum walks: Application to lattice sieving. In Carmit Hazay and Martijn Stam, editors, *Advances in Cryptology – EUROCRYPT 2023, Part V*, volume 14008 of *Lecture Notes in Computer Science*, pages 221–251, Lyon, France, April 23–27, 2023. Springer, Cham, Switzerland. doi:10.1007/978-3-031-30589-4_8.

[BDGL16] Anja Becker, Léo Ducas, Nicolas Gama, and Thijs Laarhoven. New directions in nearest neighbor searching with applications to lattice sieving. In Robert Krauthgamer, editor, *27th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 10–24, Arlington, VA, USA, January 10–12, 2016. ACM-SIAM. doi:10.1137/1.9781611974331.ch2.

[BLL25] Maiara F. Bollauf, Maja Lie, and Cong Ling. On Gaussian sampling for q-ary lattices and linear codes with Lee weight. In Yael Tauman Kalai and Seny F. Kamara, editors, *Advances in Cryptology – CRYPTO 2025*, pages 321–352, Cham, 2025. Springer Nature Switzerland.

[BLP+13] Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of learning with errors. In *Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing*, STOC '13, page 575–584, New York, NY, USA, 2013. Association for Computing Machinery. doi:10.1145/2488608.2488680.

[BLY23] Maiara F. Bollauf, Hsuan-Yin Lin, and Øyvind Ytrehus. Formally unimodular packings for the Gaussian wiretap channel. *IEEE Transactions on Information Theory*, 69(12):7755–7776, 2023. doi:10.1109/TIT.2023.3322771.

[CDLP13] Kai-Min Chung, Daniel Dadush, Feng-Hao Liu, and Chris Peikert. On the lattice smoothing parameter problem. In *2013 IEEE Conference on Computational Complexity*, pages 230–241, 2013. doi:10.1109/CCC.2013.31.

[Che13] Yuanmi Chen. *Réduction de réseau et sécurité concrète du chiffrement complètement homomorphe*. PhD thesis, 2013. Thèse de doctorat dirigée par Nguyen, Phong Q. Informatique Paris 7 2013. URL: http://www.theses.fr/2013PA077242.

[CL21] André Chailloux and Johanna Loyer. Lattice sieving via quantum random walks. In Mehdi Tibouchi and Huaxiong Wang, editors, *Advances in Cryptology – ASIACRYPT 2021, Part IV*, volume 13093 of *Lecture Notes in Computer Science*, pages 63–91, Singapore, December 6–10, 2021. Springer, Cham, Switzerland. doi:10.1007/978-3-030-92068-5_3.

[CLZ22] Yilei Chen, Qipeng Liu, and Mark Zhandry. Quantum algorithms for variants of average-case lattice problems via filtering. In Orr Dunkelman and Stefan Dziembowski, editors, *Advances in Cryptology – EUROCRYPT 2022*, pages 372–401, Cham, 2022. Springer International Publishing.

[CMHST25] Kevin Carrier, Charles Meyer-Hilfiger, Yixin Shen, and Jean-Pierre Tillich. Assessing the impact of a variant of MATZOV's dual attack on Kyber. In Yael Tauman Kalai and Seny F. Kamara, editors, *Advances in Cryptology – CRYPTO 2025*, pages 444–476, Cham, 2025. Springer Nature Switzerland.

[CPS+20] Chitchanok Chuengsatiansup, Thomas Prest, Damien Stehlé, Alexandre Wallet, and Keita Xagawa. Modfalcon: Compact signatures based on module-NTRU lattices. In *Proceedings of the 15th ACM Asia Conference on Computer and Communications Security*, ASIA CCS '20, page 853–866, New York, NY, USA, 2020. Association for Computing Machinery. doi:10.1145/3320269.3384758.

[CS99] John H. Conway and Neil J. A. Sloane. *Sphere Packings, Lattices and Groups*. New York, NY, USA, 3rd edition, 1999. doi:10.1007/978-1-4757-6568-7.

[CSS26] Clémence Chevignard, André Schrottenloher, and Yixin Shen. Quantum algorithm for discrete Gaussian sampling. Cryptology ePrint Archive, Paper 2026/984, 2026. URL: https://eprint.iacr.org/2026/984.

[DEL25] Léo Ducas, Lynn Engelberts, and Johanna Loyer. Wagner's algorithm provably runs in subexponential time for SIS$^\infty$. In *Advances in Cryptology – CRYPTO 2025: 45th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 17–21, 2025, Proceedings, Part I*, page 353–384, Berlin, Heidelberg, 2025. Springer-Verlag. doi:10.1007/978-3-032-01855-7_12.

[DEP23] Léo Ducas, Thomas Espitau, and Eamonn W. Postlethwaite. Finding short integer solutions when the modulus is small. In Helena Handschuh and Anna Lysyanskaya, editors, *Advances in Cryptology – CRYPTO 2023, Part III*, volume 14083 of *Lecture Notes in Computer Science*, pages 150–176, Santa Barbara, CA, USA, August 20–24, 2023. Springer, Cham, Switzerland. doi:10.1007/978-3-031-38548-3_6.

[DKL+18] Léo Ducas, Eike Kiltz, Tancrède Lepoint, Vadim Lyubashevsky, Peter Schwabe, Gregor Seiler, and Damien Stehlé. CRYSTALS-Dilithium: A lattice-based digital signature scheme. *Transactions on Cryptographic Hardware and Embedded Systems*, 2018, Issue 1:238–268, 2018. doi:10.13154/tches.v2018.i1.238-268.

[DKL+21] Léo Ducas, Eike Kiltz, Tancrède Lepoint, Vadim Lyubashevsky, Peter Schwabe, Gregor Seiler, and Damien Stehlé. CRYSTALS-Dilithium specification (round 3). https://pq-crystals.org/dilithium/data/dilithium-specification-round3-20210208.pdf, 2021.

[DP23] Léo Ducas and Ludo N. Pulles. Does the dual-sieve attack on learning with errors even work? *IACR Cryptol. ePrint Arch.*, page 302, 2023. URL: https://eprint.iacr.org/2023/302.

[DP26] Léo Ducas and Ludo N. Pulles. Accurate score prediction for dual-sieve attacks. *J. Cryptol.*, 39(1):8, 2026. URL: https://doi.org/10.1007/s00145-025-09560-7, doi:10.1007/S00145-025-09560-7.

[DRSD14] Daniel Dadush, Oded Regev, and Noah Stephens-Davidowitz. On the closest vector problem with a distance guarantee. In *2014 IEEE 29th Conference on Computational Complexity (CCC)*, pages 98–109, 2014. doi:10.1109/CCC.2014.18.

[EFG+22] Thomas Espitau, Pierre-Alain Fouque, François Gérard, Mélissa Rossi, Akira Takahashi, Mehdi Tibouchi, Alexandre Wallet, and Yang Yu. Mitaka: A simpler, parallelizable, maskable variant of Falcon. In Orr Dunkelman and Stefan Dziembowski, editors, *Advances in Cryptology – EUROCRYPT 2022, Part III*, volume 13277 of *Lecture Notes in Computer Science*, pages 222–253, Trondheim, Norway, May 30 – June 3, 2022. Springer, Cham, Switzerland. doi:10.1007/978-3-031-07082-2_9.

[EJK20] Thomas Espitau, Antoine Joux, and Natalia Kharchenko. On a dual/hybrid approach to small secret LWE - A dual/enumeration technique for learning with errors and application to security estimates of FHE schemes. In Karthikeyan Bhargavan, Elisabeth Oswald, and Manoj Prabhakaran, editors, *Progress in Cryptology - INDOCRYPT 2020: 21st International Conference in Cryptology in India*, volume 12578 of *Lecture Notes in Computer Science*, pages 440–462, Bangalore, India, December 13–16, 2020. Springer, Cham, Switzerland. doi:10.1007/978-3-030-65277-7_20.

[ELZ05] U. Erez, S. Litsyn, and R. Zamir. Lattices which are good for (almost) everything. *IEEE Transactions on Information Theory*, 51(10):3401–3416, 2005. doi:10.1109/TIT.2005.855591.

[EWY23] Thomas Espitau, Alexandre Wallet, and Yang Yu. On Gaussian sampling, smoothing parameter and application to signatures. In Jian Guo and Ron Steinfeld, editors, *Advances in Cryptology – ASIACRYPT 2023, Part VII*, volume 14444 of *Lecture Notes in Computer Science*, pages 65–97, Guangzhou, China, December 4–8, 2023. Springer, Singapore, Singapore. doi:10.1007/978-981-99-8739-9_3.

[FHK+19] Pierre-Alain Fouque, Jeffrey Hoffstein, Paul Kirchner, Vadim Lyubashevsky, Thomas Pornin, Thomas Prest, Thomas Ricosset, Gregor Seiler, William Whyte, and Zhenfei Zhang. Falcon: Fast-fourier lattice-based compact signatures over NTRU. 2019. URL: https://api.semanticscholar.org/CorpusID:231637439.

[GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*, STOC '08, page 197–206, New York, NY, USA, 2008. Association for Computing Machinery. doi:10.1145/1374376.1374407.

[HPS11] Guillaume Hanrot, Xavier Pujol, and Damien Stehlé. Analyzing blockwise lattice algorithms using dynamical systems. In *Advances in Cryptology - CRYPTO 2011 - 31st Annual Cryptology Conference*, volume 6841 of *Lecture Notes in Computer Science*, page 441. Springer, 2011. URL: https://www.iacr.org/archive/crypto2011/68410441/68410441.pdf, doi:10.1007/978-3-642-22792-9_25.

[Kle00] Philip Klein. Finding the closest lattice vector when it's unusually close. In *Proceedings of the Eleventh Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '00, page 937–941, USA, 2000. Society for Industrial and Applied Mathematics.

[KNSW20] Elena Kirshanova, Huyen Nguyen, Damien Stehlé, and Alexandre Wallet. On the smoothing parameter and last minimum of random orthogonal lattices. *Des. Codes Cryptogr.*, 88(5):931–950, 2020. doi:10.1007/S10623-020-00719-W.

[KOW26] Robin Kothari, Ryan O'Donnell, and Kewen Wu. No exponential quantum speedup for SIS$^\infty$ anymore, 2026. URL: https://arxiv.org/abs/2510.07515, arXiv:2510.07515.

[LLBS14] Cong Ling, Laura Luzzi, Jean-Claude Belfiore, and Damien Stehlé. Semantically secure lattice codes for the Gaussian wiretap channel. *IEEE Transactions on Information Theory*, 60(10):6399–6416, 2014. doi:10.1109/TIT.2014.2343226.

[LN20] Jianwei Li and Phong Q. Nguyen. A complete analysis of the BKZ lattice reduction algorithm. *IACR Cryptol. ePrint Arch.*, page 1237, 2020. URL: https://eprint.iacr.org/2020/1237.

[LPR13] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. A toolkit for Ring-LWE cryptography. In Thomas Johansson and Phong Q. Nguyen, editors, *Advances in Cryptology – EUROCRYPT 2013*, pages 35–54, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.

[MAT22] MATZOV. Report on the Security of LWE: Improved Dual Lattice Attack, April 2022. doi:10.5281/zenodo.6412487.

[MP13] Daniele Micciancio and Chris Peikert. Hardness of SIS and LWE with small parameters. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology – CRYPTO 2013, Part I*, volume 8042 of *Lecture Notes in Computer Science*, pages 21–39, Santa Barbara, CA, USA, August 18–22, 2013. Springer Berlin Heidelberg, Germany. doi:10.1007/978-3-642-40041-4_2.

[MR04] D. Micciancio and O. Regev. Worst-case to average-case reductions based on Gaussian measures. In *45th Annual IEEE Symposium on Foundations of Computer Science*, pages 372–381, 2004. doi:10.1109/FOCS.2004.72.

[MR09] Daniele Micciancio and Oded Regev. *Lattice-based Cryptography*, pages 147–191. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009. doi:10.1007/978-3-540-88702-7_5.

[Pei10] Chris Peikert. An efficient and parallel Gaussian sampler for lattices. In Tal Rabin, editor, *Advances in Cryptology – CRYPTO 2010*, volume 6223 of *Lecture Notes in Computer Science*, pages 80–97, Santa Barbara, CA, USA, August 15–19, 2010. Springer Berlin Heidelberg, Germany. doi:10.1007/978-3-642-14623-7_5.

[PS24] Amaury Pouly and Yixin Shen. Provable dual attacks on learning with errors. In *Advances in Cryptology – EUROCRYPT 2024: 43rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zurich, Switzerland, May 26–30, 2024, Proceedings, Part VII*, page 256–285, Berlin, Heidelberg, 2024. Springer-Verlag. doi:10.1007/978-3-031-58754-2_10.

[PS25] Amaury Pouly and Yixin Shen. Discrete Gaussian sampling for BKZ-reduced basis. In *Post-Quantum Cryptography: 16th International Workshop, PQCrypto 2025, Taipei, Taiwan, April 8–10, 2025, Proceedings, Part II*, page 63–88, Berlin, Heidelberg, 2025. Springer-Verlag. doi:10.1007/978-3-031-86602-9_3.

[PS26] Amaury Pouly and Yixin Shen. Solving the shortest vector problem in $2^{0.63269n+o(n)}$ time on random lattices. In *Advances in Cryptology – EUROCRYPT 2026*, pages 92–123, Cham, 2026. Springer Nature Switzerland. doi:10.1007/978-3-032-25327-9_4.

[QX25] Hongyuan Qu and Guangwu Xu. On the provable dual attack for LWE by modulus switching. In Goichiro Hanaoka and Bo-Yin Yang, editors, *Advances in Cryptology – ASIACRYPT 2025, Part III*, volume 16247 of *Lecture Notes in Computer Science*, pages 34–64, Melbourne, VIC, Australia, December 8–12, 2025. Springer, Singapore, Singapore. doi:10.1007/978-981-95-5099-9_2.

[Sch87] C.P. Schnorr. A hierarchy of polynomial time lattice basis reduction algorithms. *Theor. Comput. Sci.*, 53(2):201–224, jun 1987.

[Sch03] Claus Peter Schnorr. Lattice reduction by random sampling and birthday methods. In Helmut Alt and Michel Habib, editors, *STACS 2003*, pages 145–156, Berlin, Heidelberg, 2003. Springer Berlin Heidelberg.

[SS11] Damien Stehlé and Ron Steinfeld. Making NTRU as secure as worst-case problems over ideal lattices. In Kenneth G. Paterson, editor, *Advances in Cryptology – EUROCRYPT 2011*, pages 27–47, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg.

[Ste17] Noah Stephens-Davidowitz. *On the Gaussian measure over lattices*. Phd thesis, New York University, 2017.

[Wag02] David Wagner. A generalized birthday problem. In Moti Yung, editor, *Advances in Cryptology – CRYPTO 2002*, volume 2442 of *Lecture Notes in Computer Science*, pages 288–303, Santa Barbara, CA, USA, August 18–22, 2002. Springer Berlin Heidelberg, Germany. doi:10.1007/3-540-45708-9_19.

[WL19] Zheng Wang and Cong Ling. Lattice Gaussian sampling by Markov chain Monte Carlo: Bounded distance decoding and trapdoor sampling. *IEEE Transactions on Information Theory*, 65(6):3630–3645, 2019. doi:10.1109/TIT.2019.2901497.

[ZNKB14] Ram Zamir, Bobak Nazer, Yuval Kochman, and Ilai Bistritz. *Lattice Coding for Signals and Networks: A Structured Coding Approach to Quantization, Modulation and Multiuser Information Theory*. Cambridge University Press, 2014. doi:10.1017/CBO9781139045520.