Dense and smooth lattices in any genus

Wessel van Woerden[0000−0002−5565−4015]

Univ. Bordeaux, CNRS, Inria, Bordeaux INP, IMB, Talence, France
wessel.van-woerden@math.u-bordeaux.fr

Abstract. The Lattice Isomorphism Problem (LIP) was recently introduced as a new hardness assumption for post-quantum cryptography. The strongest known efficiently computable invariant for LIP is the genus of a lattice. To instantiate LIP-based schemes one often requires the existence of a lattice that (1) lies in some fixed genus, and (2) has some good geometric properties such as a high packing density or small smoothness parameter.

In this work we show that such lattices exist. In particular, building upon classical results by Siegel (1935), we show that essentially any genus contains a lattice with a close to optimal packing density, smoothing parameter and covering radius. We present both how to efficiently compute concrete existence bounds for any genus, and asymptotically tight bounds under weak conditions on the genus.

The introduction of the lattice isomorphism problem (LIP) as a hardness assumption for cryptography raises a new family of interesting questions [DvW22, BGPSD23, DPPvW22]. LIP asks to determine if two lattices are isomorphic, i.e., if one is an orthonormal transformation of the other. One way to answer this question in the negative is using invariants, and the *genus* of a lattice, gives the strongest known efficiently computable invariant for LIP. If two lattices fall into distinct genera LIP thus becomes easy. Therefore, in the context of LIP, one often works inside a certain genus or a family of genera. In particular, notions like randomness, reductions and hardness questions are suddenly restricted to within a genus.

In this work we study the geometric properties of random lattices in a fixed genus $\mathcal{G}$ and use that to show the existence of a lattice $\mathcal{L} \in \mathcal{G}$ with a good packing density, smoothing parameter or covering radius. We show that the strong condition of being in a fixed genus does in fact not change much to the behavior we are used to from random lattices. This is both interesting from a theoretic perspective, but in addition it allows to tightly instantiate cryptographic schemes that are based on LIP. Previously, the existence of such lattices was assumed heuristically [DvW22, BGPSD23, ARLW24].

Random lattices and the existence of good packings. Within cryptography families of random lattices play an important role, for example we often consider random q-ary lattices $q\mathbb{Z}^n \subset \mathcal{L} \subset \mathbb{Z}^n$ which are related to LWE, NTRU or SIS. Within cryptanalysis random lattices often play the role of worst-case instances and therefore their properties are used in heuristic analysis of algorithms. For example, heuristically we assume that lattices we encounter follow the Gaussian Heuristic which in full generality says that the number of nonzero lattice points in a 'nice' volume S is about $\mathrm{vol}(S)/\mathrm{vol}(\mathcal{L})$. What is precisely meant here by 'random' is often not so important as these heuristics give good estimates in practice.

On the mathematical side however, the notion of a random lattice is more explicitly defined. Here we look at the whole space of lattices $\mathcal{L}_{[n,D]}$ of some fixed dimension $n \geq 2$ and (co)volume $D > 0$. There exists a natural and finite Haar measure on $\mathcal{L}_{[n,D]}$, which thereby induces a natural probability distribution $\mathcal{D}(\mathcal{L}_{[n,D]})$. While the space $\mathcal{L}_{[n,D]}$ and the distribution $\mathcal{D}(\mathcal{L}_{[n,D]})$ can be quite complicated to understand, it allows for remarkably clean and provable statements about the expected behavior of lattices following this distribution. For example, for any star-shaped volume S the expected number $\mathbb{E}[|S \cap \mathcal{L} \setminus \{0\}|]$ of nonzero lattice points in S over $\mathcal{D}(\mathcal{L}_{[n,D]})$ is precisely equal to $\mathrm{vol}(S)/\mathrm{vol}(\mathcal{L})$, i.e., what the Gaussian Heuristic prescribes. However, in this case it is a provable result.

Now by choosing $S \subset \mathbb{R}^n$ to be a ball of radius λ we can determine the expected number of nonzero lattice points of length at most λ. If this expectation is strictly less than 2 we know that there must exist a lattice $\mathcal{L} \in \mathcal{L}_{[n,D]}$ that has strictly less than 2 vectors of length λ. Because any lattice point occurs in a pair $\pm x$ of the same length we thus know that this lattice contains no lattice points of length at most λ and thus that its minimum distance $\lambda_1(\mathcal{L}) := \min\{\|v\| : v \in \mathcal{L}\}$ satisfies $\lambda_1(\mathcal{L}) > \lambda$. By picking the appropriate λ and by slightly refining this argument one gets the Minkowski-Hlawka Theorem, which says that there exists a lattice $\mathcal{L} \subset \mathbb{R}^n$ with $\lambda_1(\mathcal{L}) \geq (2\zeta(n)\,\mathrm{vol}(\mathcal{L})/\omega_n)^{1/n} \approx \sqrt{n/2\pi e} \cdot \mathrm{vol}(\mathcal{L})^{1/n}$, where ω_n is the volume of the n-dimensional unit ball. This is close to optimal as Minkowski's Theorem says that $\lambda_1(\mathcal{L}) \leq \mathrm{mk}(\mathcal{L}) := 2 \cdot (\mathrm{vol}(\mathcal{L})/\omega_n)^{1/n}$.

So from the average-case behavior of random lattices one can show the existence of a lattice with a large minimum distance, i.e., that of a good lattice packing. More generally, random lattices are known to have other good geometric properties. Besides a good packing density, they also have a large covering radius and a small smoothing parameter in expectation. And again, this immediately results in a proof of existence for lattices with such good properties.

Random lattices in a fixed genus. Our question is if the same can be said when we add the seemingly strong restriction of falling in some fixed genus. Two (full-rank) integral lattices $\mathcal{L}_1, \mathcal{L}_2 \subset \mathbb{R}^n$ fall into the same genus if they are equivalent over the p-adic integers $\mathbb{Z}_p$ for all primes p, i.e., if $\mathcal{L}_1 \otimes_{\mathbb{Z}} \mathbb{Z}_p \cong \mathcal{L}_2 \otimes_{\mathbb{Z}} \mathbb{Z}_p$[1] Computing if two lattices are equivalent over $\mathbb{Z}_p$ is efficient, essentially because one can (block) diagonalize (gram) matrices over this local ring, and the equivalence

[1] One could view this as formally replacing the lattice $B \cdot \mathbb{Z}^n$ by $B \cdot \mathbb{Z}_p^n$, and the standard Euclidean inner product with image $\mathbb{Z}$ becomes the bi-linear form $(x, y) \mapsto \sum_i x_i y_i$ with image $\mathbb{Z}_p$. An isomorphism has to preserve both the $\mathbb{Z}_p$ structure as the bi-linear form.

class can then simply be read from the (block) diagonalized form. Furthermore, assuming that $\mathrm{vol}(\mathcal{L}_1) = \mathrm{vol}(\mathcal{L}_2)$, one only has to check this equivalence over the finite number of primes p dividing $2\,\mathrm{vol}(\mathcal{L}_i)^2$. Given the prime factorization of $\mathrm{vol}(\mathcal{L}_i)^2$ the genus equivalence is thus efficiently computable.

Minkowski showed that any genus only contains a finite number of isomorphism classes $[\mathcal{L}_1], \ldots, [\mathcal{L}_m]$ [Min85]. Furthermore, if one restrict the usual Haar measure to a fixed genus $\mathcal{G}$ we obtain a natural distribution $\mathcal{D}(\mathcal{G})$ on these classes, where each $[\mathcal{L}_i]$ is sampled relative to its *mass* $m([\mathcal{L}_i]) = 1/|\mathrm{Aut}(\mathcal{L}_i)|$, where $\mathrm{Aut}(\mathcal{L}_i)$ is the automorphism group of $\mathcal{L}_i$.

Now that we have a natural notion of randomness on a genus $\mathcal{G}$, the question is if we can say something about the expected behavior of certain lattice properties. It turns out that the answer is yes, and in fact most of the theory for this was already developed almost 90 years ago by Siegel [Sie35] in the form of *mass formulas.*

For an integer $k \geq 1$ and an integral lattice $\mathcal{L}$ we denote the number of lattice point with squared norm k by $N_{\mathcal{L}}(k) := |\{x \in \mathcal{L} : \|x\|^2 = k\}|$. Generally, computing $N_{\mathcal{L}}(k)$ is a very hard problem, however Siegel showed that its expected value over a genus is essentially equal to a converging product of local densities at each prime p. Each local density is efficiently computable and similarly as for the genus one only really has to compute them for the primes p dividing $2k\,\mathrm{vol}(\mathcal{G})^2$. Siegel's mass formula thus implies, that for a genus $\mathcal{G}$ we can efficiently compute the expectation $N_{\mathcal{G}}(k) := \mathbb{E}_{[\mathcal{L}]\leftarrow\mathcal{D}(\mathcal{G})}[N_{\mathcal{L}}(k)]$ (given the prime factorization of $2k\,\mathrm{vol}(\mathcal{G})^2$).

We can now make a similar argument as for the Minkowski-Hlawka Theorem. For an integer $\lambda \geq 1$ consider the sum $S_\lambda := \sum_{k=1}^{\lambda} N_{\mathcal{G}}(k)$. This sum represents the expected number of nonzero lattice vectors of squared norm at most λ for a lattice $\mathcal{L}$ sampled from $\mathcal{D}(\mathcal{G})$. Now if $S_\lambda < 2$ we know that there must exist a lattice $\mathcal{L} \in \mathcal{G}$ with strictly less than 2 and thus precisely 0 nonzero vectors of squared norm less than λ. So we have that $\lambda_1(\mathcal{L})^2 > \lambda$, and by appropriately picking λ this can show the existence of a good lattice packing in $\mathcal{G}$. This reasoning was already used in an unpublished work by Conway and Thompson, and written down by Milnor [MH$^+$73], to show the existence of odd unimodular lattices, integral lattices $\mathcal{L}$ with volume 1 which contain vectors of odd squared norm, with $\lambda_1(\mathcal{L})^2 \geq \lfloor(\frac{3}{5}\omega_n)^{-2/n}\rceil$. For the case of even unimodular lattices, that only exist when $8|n$, Milnor [MH$^+$73], using computations of Serre [Ser73], claims a similar bound of $\lambda_1(\mathcal{L})^2 \geq 2\lfloor\frac{1}{2}(\frac{3}{5}\omega_n)^{-2/n}\rceil$. Note that for both odd and even unimodular lattices the existence bound is only slightly weaker than the Minkowski-Hlawka Theorem, and they quickly converge to each other for large n. The additional restriction of falling in a fixed genus therefore does not seem strong enough to influence the good geometric properties of random lattices too much.

Contributions. The first aim of this work is to survey these classical and maybe surprising results related to the genus. The existing literature however only seems to consider the unimodular case and the packing density[2]. In this work, we therefore extend these results in two ways that are of interest to cryptography.

Firstly, we extend the Minkowski-Hlawka-like Theorem to almost any genus. In particular, under light conditions[3] on the genus $\mathcal{G}$, we show the existence of a good lattice packing $\mathcal{L} \in \mathcal{G}$, with minimum distance equivalent to the Minkowski-Hlawka Theorem up to a small factor $O(1)^{1/n}$. We achieve this by bounding the local densities and thus the expected number of lattice vectors $N_{\mathcal{G}}(k)$ of squared norm k.

Theorem 1 (General packing). *For any integral genus $\mathcal{G}$ in dimension $n \geq 6$ such that $\mathrm{rk}_p(\mathcal{G}) \geq 6$ for all primes p, and any constant $0 < c \leq 1$, we have*

$$\Pr_{[\mathcal{L}] \leftarrow \mathcal{D}(\mathcal{G})}\left[\lambda_1(\mathcal{L})^2 \geq \left\lceil c^2 \cdot \left(\frac{7\zeta(3)}{9\zeta(2)} \cdot \frac{\mathrm{vol}(\mathcal{L})}{\omega_n}\right)^{2/n}\right\rceil\right] > 1 - c^n.$$

In particular, there exists a lattice $\mathcal{L} \in \mathcal{G}$ with

$$\lambda_1(\mathcal{L})^2 \geq \left\lceil \left(\frac{7\zeta(3)}{9\zeta(2)} \cdot \frac{\mathrm{vol}(\mathcal{L})}{\omega_n}\right)^{2/n}\right\rceil \approx n/2\pi e \cdot \mathrm{vol}(\mathcal{L})^{2/n}.$$

Note that in fact we show something stronger, i.e., by roughly lowering the bound on the first minimum by a constant factor c we show that it is attained with a probability of at least $1 - c^n$ over $\mathcal{D}(\mathcal{G})$. This follows directly from an application of Markov's inequality in the proof and the result closely matches the behavior of the first minimum for random lattices [AEN19]. Furthermore, this allows us to show the existence of a lattice $\mathcal{L} \in \mathcal{G}$ for which both $\mathcal{L}$ and its dual $\mathcal{L}^*$ have a good packing density.

Secondly, we show that the reasoning can be extended to prove the existence of lattices with a good covering radius $v(\mathcal{L}) := \min\{\lambda > 0 : \mathrm{dist}(\mathcal{L}, x) \leq \lambda \ \forall x \in \mathbb{R}^n\}$ and a good smoothing parameter $\eta_\varepsilon(\mathcal{L})$, even for the relatively large values $\varepsilon \gg e^{-n}$ that are of interest in cryptography.

Theorem 2 (General smoothing). *For any integral genus $\mathcal{G}$ in dimension $n \geq 6$ such that $\mathrm{rk}_p(\mathcal{G}) \geq 6$ for all primes p, constants $C = 26.1$ and $0 < c \leq 1$, and $\epsilon \geq C \cdot (ce)^{-n} \cdot \mathrm{vol}(\mathcal{G})^{-1}$, we have*

$$\Pr_{[\mathcal{L}] \leftarrow \mathcal{D}(\mathcal{G})}\left[\eta_\epsilon(\mathcal{L}^*) \leq \frac{1}{c} \cdot \left(\frac{C \cdot \mathrm{vol}(\mathcal{L}^*)}{\epsilon}\right)^{1/n}\right] > 1 - c^n.$$

In particular, there exists a lattice $\mathcal{L} \in \mathcal{G}$ such that $\eta_\epsilon(\mathcal{L}^*) \leq (C \cdot \mathrm{vol}(\mathcal{L}^*)/\epsilon)^{1/n}$.

[2] As far as we know.

[3] We require that the rank $\mathrm{rk}_p(\mathcal{G})$ of a Gram matrix $G \bmod p$ over $\mathbb{F}_p$ of any lattice $\mathcal{L} \in \mathcal{G}$ is at least 6. Note that this property only has to be checked for primes $p | \det(G)$, and is (after normalization) true for most integral lattices of sufficiently large dimension. In particular it is true for $\mathbb{Z}^n$, SIS, LWE and NTRU lattices with a sufficiently high dimension and number of samples.

Theorem 3 (General covering radius). *For any integral genus $\mathcal{G}$ in dimension $n \geq 6$ such that $\mathrm{rk}_p(\mathcal{G}) \geq 6$ for all primes p, and constants $C = 26.1$ and $e^{-1}(2C/(3\,\mathrm{vol}(\mathcal{G})))^{1/n} \leq c \leq 1$, we have*

$$\Pr_{[\mathcal{L}]\leftarrow\mathcal{D}(\mathcal{G})}\left[v(\mathcal{L}^*) \leq \frac{1}{c}\cdot\left(\sqrt{n/2\pi}+1\right)\cdot\left(\frac{2}{3}C\cdot\mathrm{vol}(\mathcal{L}^*)\right)^{1/n}\right] > 1 - c^n.$$

In particular, when additionally $n \geq 7$, there exists a lattice $\mathcal{L} \in \mathcal{G}$ such that

$$v(\mathcal{L}^*) \leq \left(\sqrt{n/2\pi}+1\right)\cdot\left(\frac{2}{3}C\cdot\mathrm{vol}(\mathcal{L}^*)\right)^{1/n} \approx \sqrt{e}\cdot\sqrt{n/2\pi e}\cdot\mathrm{vol}(\mathcal{L}^*)^{1/n}.$$

Besides these essentially tight asymptotic bounds we also explain how to efficiently compute concrete existence bounds for any fixed genus.

Applications. Finally, we give some applications of these results for the instantiation of LIP-based schemes. Suppose we have an efficiently decodable lattice $\mathcal{L}$ with unique decoding radius $\rho = \Theta(\lambda_1(\mathcal{L}))$, and suppose that

$$\mathrm{gap}(\mathcal{L}) := \max\{\mathrm{mk}(\mathcal{L})/\lambda_1(\mathcal{L}), \mathrm{mk}(\mathcal{L}^*)/\lambda_1(\mathcal{L}^*)\} \leq f,$$

i.e., the primal and dual minimum distances and the decoding radius are within a factor $O(f)$ from optimal. Heuristically, the larger f is the easier it is to decode or find short vectors in this lattice (or its dual). Given such a lattice [DvW22] shows how to instantiate an encryption scheme where the security is solely based on distinguishing between some isomorphism classes $[\mathcal{L}_1], [\mathcal{L}_2]$ constructed from $\mathcal{L}$ that lie in the same genus. However in this construction the geometric gaps blow up to $\mathrm{gap}(\mathcal{L}_i) = O(f^3)$ which reduces the concrete security significantly. Our results show the existence of a lattice $\mathcal{L}'$ in the same genus as $\mathcal{L}$ but such that $\mathrm{gap}(\mathcal{L}') = O(1)$. This can in turn be used to create a suitable pair $\mathcal{L}_1, \mathcal{L}_2$ for which $\mathrm{gap}(\mathcal{L}_i) = O(f)$, and thus we reduce the cubic loss to only a small constant loss. The encryption scheme from [BZI$^+$24] based on the same framework benefits from the same improvement. Similarly, we show how to instantiate the signature scheme from [DvW22] with a constant loss $O(f)$ instead of a quadratic loss $O(f^2)$.

Another interesting work [BGPSD23] introduces constructions based on LIP for the unimodular lattice $\mathbb{Z}^n$. To instantiate their scheme the authors assume that there exists a lattice $\mathcal{L}$ in the odd unimodular genus $\mathcal{G}_{\mathrm{odd}}$ of $\mathbb{Z}^n$ with $\lambda_1(\mathcal{L}) \geq \Omega(n/\log(n))$ and $\eta_\varepsilon(\mathbb{Z}^n) \leq \eta_\varepsilon(\mathbb{Z}^n)/\sqrt{\log(n)}$ for $\varepsilon < n^{-\omega(1)n}$. Similarly, the encryption scheme [ARLW24] based on LIP requires the existence of an even unimodular lattice $\mathcal{L}$ such that $\lambda_1(\mathcal{L}) \geq \sqrt[4]{72n}$, and this is conjectured to be true for $n \geq 85$. We raise that the first claim for the first minimum is already answered by [Ser73, MH$^+$73], and the second claim for the smoothing parameter follows from Lemma 4. In fact, this shows a much stronger result than required.

Acknowledgements. W. van Woerden was supported by the CHARM ANR-NSF grant (ANR-21-CE94-0003).

1 Preliminaries

Notation. Vectors are *column vectors.* For a ring R we denote $\mathcal{GL}_n(R)$ as the general linear group of $n \times n$ invertible matrices over R. For $R \subset \mathbb{R}$ we denote $\mathcal{S}_n^{>0}(R)$ as the space of positive-definite symmetric matrices over R. We denote $\mathcal{O}_n(\mathbb{R})$ for the group of orthonormal transformations over the reals $\mathbb{R}$. We denote $\zeta(\cdot)$ for the Riemann zeta function.

1.1 Lattices and quadratic forms

Lattices. A lattice $\mathcal{L}$ is a discrete additive subgroup of the Euclidean space $\mathbb{R}^n$. We call the dimension of the real span $\operatorname{Span}_{\mathbb{R}}(\mathcal{L}) \subset \mathbb{R}^n$ the *rank* $\operatorname{rk}(\mathcal{L})$ of a lattice, and say that $\mathcal{L} \subset \mathbb{R}^n$ has *full-rank* if $\operatorname{rk}(\mathcal{L}) = n$. In this work we restrict ourselves to full-rank lattices. Full-rank lattices $\mathcal{L} \subset \mathbb{R}^n$ can be represented by a full-rank basis $B \in \mathcal{GL}_n(\mathbb{R})$ such that

$$\mathcal{L} = \mathcal{L}(B) := B \cdot \mathbb{Z}^n = \{Bx : x \in \mathbb{Z}^n\}.$$

Such a basis representation is not unique, i.e., for any basis $B \in \mathcal{GL}_n(\mathbb{R})$ and any unimodular matrix $U \in \mathcal{GL}_n(\mathbb{Z})$ we have $\mathcal{L}(B) = \mathcal{L}(B \cdot U)$.

For a basis B we call $G_B := B^\top B \in \mathcal{S}_n^{>0}(\mathbb{R})$ the *gram* matrix of B and a gram matrix of the lattice $\mathcal{L}(B)$. Note that a gram matrix does not uniquely define a lattice, in particular for any basis $B \in \mathcal{GL}_n(\mathbb{R})$ and any orthonormal transformation $O \in \mathcal{O}_n(\mathbb{R})$ we have $G_B = G_{OB}$ while $\mathcal{L}(B)$ and $\mathcal{L}(OB)$ are usually distinct. From a gram matrix $G \in \mathcal{S}_n^{>0}(\mathbb{R})$ one can always construct a corresponding lattice basis by computing the unique Cholesky decomposition $G = C^\top C$ where C is an upper-triangular matrix with positive diagonal. Generally however, the Cholesky decomposition of G_B does not return the basis B of $\mathcal{L}(B)$, but some basis $C = O \cdot B$ of $O \cdot \mathcal{L}(B)$ for some $O \in \mathcal{O}_n(\mathbb{R})$.

For a lattice $\mathcal{L}$ we write $\mathcal{L}^*$ for its *dual* lattice given by

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

As expected we have that $(\mathcal{L}^*)^* = \mathcal{L}$. If B is a basis and G a gram matrix of $\mathcal{L}$, then $(B^{-1})^\top$ is a basis and G^{-1} a gram matrix of $\mathcal{L}^*$.

Lattice properties. Due to the discrete and additive nature of a lattice there exists a positive minimum (Euclidean) distance $\lambda_1(\mathcal{L})$ called the *first minimum* between any two distinct lattice points. Equivalently, this can be defined as $\lambda_1(\mathcal{L}) := \min_{x \in \mathcal{L} \setminus \{0\}} \|x\|$. For a (full-rank) lattice $\mathcal{L} = \mathcal{L}(B)$ we define its *(co)volume* $\operatorname{vol}(\mathcal{L})$ as $|\det(B)|$ which is independent of the chosen basis. Equivalently, we have $\operatorname{vol}(\mathcal{L}) = \det(G)^{\frac{1}{2}}$ for any gram matrix $G \in \mathcal{S}_n^{>0}(\mathbb{R})$ of $\mathcal{L}$. Furthermore, note that $\operatorname{vol}(\mathcal{L}^*) = \operatorname{vol}(\mathcal{L})^{-1}$. The first minimum and the volume of a lattice $\mathcal{L} \subset \mathbb{R}^n$ are related to each-other by Minkowski's Theorem which says that

$$\lambda_1(\mathcal{L}) \leq \operatorname{mk}(\mathcal{L}) := 2 \cdot \frac{\operatorname{vol}(\mathcal{L})^{1/n}}{\omega_n^{1/n}} \approx \sqrt{2n/\pi e} \cdot \operatorname{vol}(\mathcal{L})^{1/n},$$

where ω_n is the volume of the n-dimensional unit ball. The *covering radius* $v(\mathcal{L})$ of a (full-rank) lattice $\mathcal{L} \subset \mathbb{R}^n$ is the minimum radius $r > 0$ such that any target $t \in \mathcal{L}$ is at distance at most r from the lattice, i.e., such that $\mathcal{L} + r\mathcal{B}^n = \mathbb{R}^n$.

We call a lattice *integral* or *rational*, if all for all pair-wise $x, y \in \mathcal{L}$ the inner product $\langle x, y \rangle \in \mathbb{Z}$ is integer or rational, respectively. Equivalently this means that for any basis B of $\mathcal{L}$ the gram matrix $G_B = B^\top B$ has integer or rational coefficients, i.e., $G_B \in \mathcal{S}_n^{>0}(\mathbb{Z})$ or $G_B \in \mathcal{S}_n^{>0}(\mathbb{Q})$, respectively. Note that a lattice being *integral* is a weaker condition than being *integer* $\mathcal{L} \subset \mathbb{Z}^n$, and many well-known lattices are integral but not integer. We define the *scale* of a rational lattice $\mathcal{L}$ by $\operatorname{scale}(\mathcal{L}) := \max\{0 < s < \infty : \frac{1}{s}\mathcal{L} \text{ is integral}\}$, which can efficiently be computed from any gram matrix G of $\mathcal{L}$. We call an integral lattice $\mathcal{L}$ *normalized* if $\operatorname{scale}(\mathcal{L}) = 1$. Note that every rational lattice can be normalized to an integral lattice as $\mathcal{L}/\operatorname{scale}(\mathcal{L})$. If a lattice $\mathcal{L}$ is rational, then its dual is also rational and thus integral up to scaling. We define the *parity* of an integral lattice $\mathcal{L}$ by $\operatorname{par}(\mathcal{L}) := \gcd(\{\|x\|^2 : x \in \mathcal{L}\}) / \gcd(\{\langle x, y \rangle : x, y \in \mathcal{L}\}) \in \{1, 2\}$, which can efficiently be computed from any gram matrix G of $\mathcal{L}$. For an integral lattice $\mathcal{L} \subset \mathbb{R}^n$ with gram matrix $G \in \mathcal{S}_n^{>0}(\mathbb{Z})$ and any prime p we define its p-rank by $\operatorname{rk}_p(\mathcal{L}) := \operatorname{rk}_{\mathbb{F}_p}(G)$, which is independent of the choice of gram matrix.

Besides the first minimum $\lambda_1(\mathcal{L})$ of a lattice we can also ask how many lattice vectors exists of a certain length. This information is usually denoted by the theta series of a lattice.

Definition 1 (Theta series). *Let $\mathcal{L}$ be a lattice, the* theta series $\theta_{\mathcal{L}}(q)$ *of $\mathcal{L}$ is the formal q-series*

$$\theta_{\mathcal{L}}(q) = \sum_{v \in \mathcal{L}} q^{\|v\|^2}.$$

For integral lattices $\mathcal{L}$ we obtain the formal power series $\theta_{\mathcal{L}}(q) = 1 + \sum_{k=1}^{\infty} N_{\mathcal{L}}(k) \cdot q^k$, where $N_{\mathcal{L}}(k) := |\{v \in \mathcal{L} : \|v\|^2 = k\}|$ is the number of vectors with squared norm k.

Note that for an integral lattice $\mathcal{L}$ we have $N_{\mathcal{L}}(k) = 0$ for all $0 < k < \lambda_1(\mathcal{L})^2$, i.e., the first $\lambda_1(\mathcal{L})^2 - 1$ non-trivial coefficients of the theta series are 0. Theta series and their relation to the first minimum of a lattice will play an important role in this work. Another property that is important in lattice-based cryptography is the smoothing parameter.

Definition 2 (smoothing parameter). *For $\varepsilon > 0$ the smoothing parameter $\eta_\varepsilon(\mathcal{L})$ of a lattice $\mathcal{L}$ is given by the minimum $s > 0$ such that $\theta_{\mathcal{L}^*}(\exp(-\pi s^2)) = 1 + \varepsilon$.*

While this might not immediately be clear from the definition, the smoothing parameter indicates how large the standard deviation of a centered (continuous)

Gaussian must be such that it becomes ε-close to uniform over the quotient $\mathbb{R}^n/\mathcal{L}$. The latter property can for example be used in security proofs of signature schemes to show that the signatures sampled from a discrete Gaussian with standard deviation $\sigma \geq \eta_\varepsilon(\mathcal{L})$ do not leak any information. Preferably we thus want the smoothing parameter $\eta_\varepsilon(\mathcal{L})$ to be small, something which we informally call *good smoothing.* A lattice with a good smoothing automatically also has a small covering radius.

Lemma 1 ([RSD24, Lemma 6.1]). *For any lattice $\mathcal{L} \subset \mathbb{R}^n$ we have*

$$v(\mathcal{L}) \leq \left(\sqrt{n/2\pi} + 1\right) \cdot \eta_{\frac{3}{2}}(\mathcal{L}).$$

Note that computing the first minimum, (part of) the theta series, the covering radius, or the smoothing parameter is generally a hard problem for which the best algorithms take at least $2^{\Omega(n)}$ time.

Quadratic forms. In the literature on post-quantum cryptography and cryptanalysis it is common to work with bases and lattices. On the contrary, in the mathematical study of lattices it is quite common to work with gram matrices and (positive definite) quadratic forms. We discuss how those are related and how they essentially give a different view on the same object.

Consider a basis $B \in \mathcal{GL}_n(\mathbb{R})$ and its gram matrix $G = B^\top B$. A gram matrix is positive definite and naturally defines a positive definite real quadratic form

$$f_G : \mathbb{R}^n \to \mathbb{R}, \quad x \mapsto x^\top G x = \sum_{i=1}^{n} \sum_{j=1}^{n} G_{ij} x_i x_j.$$

From now on we will simply identify f_G with G and call a gram matrix G a *quadratic form* or simply a *form.* Due to the positive-definiteness such a quadratic form defines a norm by $\|x\|_G := \sqrt{x^\top G x}$. Note that for $v = Bx$ we have the following identity:

$$\|v\|_2^2 = (Bx)^\top Bx = x^\top B^\top Bx = x^\top Gx =: \|x\|_G^2.$$

More generally, G defines an inner product $\langle x, y\rangle_G := x^\top G y$, and for Bx, By we have that

$$\langle Bx, By\rangle = (Bx)^\top By = x^\top B^\top By = x^\top Gy = \langle x, y\rangle_G.$$

In terms of geometry it is thus equivalent to consider the vector Bx under the Euclidean geometry or the vector x under the geometry induced by G. Every lattice point of $\mathcal{L}(B)$ can be written as Bx for an integer vector $x \in \mathbb{Z}^n$, thus on the quadratic form side we always consider the lattice $\mathbb{Z}^n$ (but we change its geometry). For a quadratic form G one could thus similarly define its first minimum by $\lambda_1(G) = \min_{x \in \mathbb{Z}^n \setminus \{0\}} \|x\|_G$, or its (co)volume $\mathrm{vol}(G) := \sqrt{\det(G)}$ matching those of any corresponding lattice.

Throughout this work when we talk about a lattice we always implicitly assume it is represented by some basis or some gram matrix. Generally, we will stick to the lattice terminology to present our main results, but for the proofs we will often switch to quadratic forms as they are more natural to work with in this setting.

1.2 Random lattices

The space of full-rank lattices in $\mathbb{R}^n$ of volume 1 can be identified by the quotient $\mathcal{L}_{[n]} := \mathcal{GL}_n(\mathbb{R})/\mathcal{GL}_n(\mathbb{Z})$. Here $\mathcal{GL}_n(\mathbb{R})$ represents all the bases of volume 1, and $\mathcal{GL}_n(\mathbb{Z})$ represents the basis transformations that turn one basis into another basis of the same lattice.

The space $\mathcal{GL}_n(\mathbb{R})$ has a natural invariant Haar measure, and Siegel proved in 1945 [Sie45] that the mass of $\mathcal{L}_{[n]}$ is finite under the projection of this Haar measure. After normalization this yields a probability distribution μ_n over $\mathcal{L}_{[n]}$. By construction this distribution is invariant under both orthonormal and basis transformations, i.e., for any measurable set $\mathcal{A} \subset \mathcal{L}_{[n]}$ and all $O \in \mathcal{GL}_n(\mathbb{R})$, and $U \in \mathcal{GL}_n(\mathbb{Z})$ we have $\mu_n(O\mathcal{A}U) = \mu_n(\mathcal{A})$. A *random lattice* is thus a unit lattice $\mathcal{L} \in \mathcal{L}_{[n]}$ sampled under the probability distribution μ_n. More generally, simply by scaling, we also speak of random lattices of some fixed volume $D > 0$.

In 1943, Hlawka proved the following, maybe surprising result about the expectation of a function $f : \mathbb{R}^n \to \mathbb{R}$ over a random lattice.

Theorem 4 ([Hla43, Sie45]). *Let $n \geq 2$ and let $f : \mathbb{R}^n \to \mathbb{R}$ be an Riemann-integrable function such that $\|x\|^{n+c} f(x)$ is bounded on $\mathbb{R}^n$ for some fixed $c > 0$. Then*

$$\int\limits_{\mathcal{L} \in \mathcal{L}_{[n]}} \sum_{x \in \mathcal{L} \setminus \{0\}} f(x) d\mu_n = \int\limits_{\mathbb{R}^n} f(x) dx.$$

In particular, for a star-shaped volume S the expected number of nonzero lattice vectors in S for a random lattice of volume D is $\mathrm{vol}(S)/D$. Furthermore, the expected number of primitive lattice vectors is $\frac{\mathrm{vol}(S)}{\zeta(n)D}$.

Knowing the expected number of (primitive) lattice points in a certain volume is enough to show the existence of a lattice with a good packing.

Corollary 1 (Minkowski-Hlawka theorem [Min10, Hla43]). *For any dimension n and volume D there exists a lattice $\mathcal{L} \in \mathcal{L}_{[n,D]}$ with*

$$\lambda_1(\mathcal{L}) \geq (2\zeta(n)\,\mathrm{vol}(\mathcal{L})/\omega_n)^{1/n} \approx \sqrt{n/2\pi e} \cdot \mathrm{vol}(\mathcal{L})^{1/n}.$$

Proof. For any $(2\zeta(n)D/\omega_n)^{1/n} > \varepsilon > 0$ let $\lambda = (2\zeta(n)D/\omega_n)^{1/n} - \varepsilon > 0$ and let $S_\lambda \subset \mathbb{R}^n$ be the n-dimensional ball with radius λ. Then by construction

$$\frac{\mathrm{vol}(S_\lambda)}{\zeta(n)D} = \frac{\lambda^n \cdot \omega_n}{\zeta(n)D} < 2,$$

and thus the expected number of primitive lattice vectors in S_λ is strictly less than 2. There thus exists a lattice $\mathcal{L} \in \mathcal{L}_{[n,D]}$ such that $|S_\lambda \cap \mathcal{L}| < 2$ which implies that $|S_\lambda \cap \mathcal{L}| = 0$ as any lattice vectors occurs as a pair $\pm x$. So $\lambda_1(\mathcal{L}) > (2\zeta(n)\operatorname{vol}(\mathcal{L})/\omega_n)^{1/n} - \epsilon$. In particular, letting $\varepsilon \to 0$ shows that

$$\sup_{\mathcal{L} \in \mathcal{L}_{[n,D]}} \lambda_1(\mathcal{L}) \geq (2\zeta(n)\operatorname{vol}(\mathcal{L})/\omega_n)^{1/n}.$$

It is a classical result (see e.g. [Wat60, p. 29-31]) that this supremum is attained by some lattice $\mathcal{L} \in \mathcal{L}_{[n,D]}$, from which the Theorem follows. □

Ignoring the small factor[4] $(2\zeta(n))^{1/n}$ the quantity $\operatorname{gh}(\mathcal{L}) := (\operatorname{vol}(\mathcal{L})/\omega_n)^{1/n} \approx \sqrt{n/2\pi e} \cdot \operatorname{vol}(\mathcal{L})^{1/n}$ is often called the *Gaussian heuristic* of a lattice $\mathcal{L}$. Besides the existence of a lattice $\mathcal{L} \subset \mathbb{R}^n$ with $\lambda_1(\mathcal{L}) \geq \operatorname{gh}(\mathcal{L})$ one can also show concentration results that show that the first minimum of a random lattice becomes heavily concentrated around $\operatorname{gh}(\mathcal{L})$ for growing n (see [AEN19] for a survey). In cryptanalysis, this is often also heuristically assumed to be the case for 'random' lattices in a more broader sense, hence the name.

Beyond the existence of a good packing one can also show the existence of a lattice with a good (small) smoothing parameter for any $\varepsilon > 0$.

Corollary 2 (Random smoothing). *For any dimension n and volume D and any $\epsilon > 0$ there exists a lattice $\mathcal{L} \in \mathcal{L}_{[n,D]}$ with*

$$\eta_\epsilon(\mathcal{L}) \leq \left(\frac{\operatorname{vol}(\mathcal{L})}{\epsilon}\right)^{\frac{1}{n}}.$$

Proof. Without loss of generality we normalize to have volume $D = 1$, and consider the function $f(x) = e^{-\pi s^2 \|x\|^2}$ for $s > 0$ and for which $\|x\|^{n+1} f(x)$ is clearly bounded on $\mathbb{R}^n$. Applying Theorem 4 we obtain that

$$\int\limits_{\mathcal{L}^* \in \mathcal{L}_{[n]}} \sum_{x \in \mathcal{L}^* \setminus \{0\}} e^{-\pi s^2 \|x\|^2} d\mu_n = \int\limits_{\mathbb{R}^n} e^{-\pi s^2 \|x\|^2} dx = s^{-n}.$$

Let $s := \varepsilon^{-\frac{1}{n}}$, by the above there exists a lattice $\mathcal{L}^* \in \mathcal{L}_{[n]}$ such that

$$\theta_{\mathcal{L}^*}(\exp(-\pi s^2)) = 1 + \sum_{x \in \mathcal{L}^* \setminus \{0\}} e^{-\pi s^2 \|x\|^2} \leq 1 + s^{-n} = 1 + \varepsilon.$$

Then by definition for the dual $\mathcal{L} \in \mathcal{L}_{[n]}$ of $\mathcal{L}^*$ we have $\eta_\varepsilon(\mathcal{L}) \leq s$. □

Note that Corollary 2 goes beyond just combining Corollary 1 with bounds based in the (dual) minimal distance like $\eta_\varepsilon(\mathcal{L}) \leq \sqrt{\ln(1/\varepsilon)}/\lambda_1(\mathcal{L}^*)$ for $\epsilon \in (0, e^{-n}]$. In particular, it gives a better and tighter bound for large $\varepsilon > e^{-n}$, which is precisely the regime interesting for cryptography. This bound is represented in a different setting in [DADRT23, Proposition 4.].

[4] One could argue that for the actual Gaussian Heuristic this factor should not be neglected, but it is often ignored as it quickly converges to 1 as n grows.

1.3 Lattice Isomorphism Problem

The lattice isomorphism problem asks if two lattices $\mathcal{L}_1$ and $\mathcal{L}_2$ are related to each-other by an orthonormal transformation. In terms of bases this means there exists both an orthonormal transformation on the left and a unimodular (basis) transformation on the right that transforms one basis into the other. In the setting of gram matrices or quadratic forms the orthonormal transformation is irrelevant and only the unimodular transformation remains but is applied (transposed) to both sides.

Definition 3 (Lattice Isomorphism). *We call two full-rank lattices $\mathcal{L}_1, \mathcal{L}_2 \subset \mathbb{R}^n$* isomorphic *and write $\mathcal{L}_1 \cong \mathcal{L}_2$ if there exists an orthonormal transformation $O \in \mathcal{O}_n(\mathbb{R})$ such that $O \cdot \mathcal{L}_1 = \mathcal{L}_2$. If $\mathcal{L}_i = \mathcal{L}(B_i)$ for bases $B_1, B_2 \in \mathcal{GL}_n(\mathbb{R})$ then $\mathcal{L}_1 \cong \mathcal{L}_2$ if and only if:*

1. *there exist $O \in \mathcal{O}_n(\mathbb{R}), U \in \mathcal{GL}_n(\mathbb{Z})$ such that $O \cdot B_1 \cdot U = B_2$, or equivalently,*
2. *there exist $U \in \mathcal{GL}_n(\mathbb{Z})$ such that $U^\top G_1 U = G_2$ where $G_i = B_i^\top B_i$.*

In the quadratic form setting the gram matrices G_1, G_2 are called $\mathbb{Z}$-equivalent or simply *equivalent* if they represent isomorphic lattices. Given two isomorphic lattices it is computationally a hard problem to find the isomorphism between them.

Definition 4 (Search LIP). *Given a pair of isomorphic lattices $\mathcal{L}_1, \mathcal{L}_2 \subset \mathbb{R}^n$, compute an orthonormal transformation $O \in \mathcal{O}_n(\mathbb{R})$ such that $O \cdot \mathcal{L}_1 = \mathcal{L}_2$.*

We allow for the lattice to be either represented by a basis or by a gram matrix. In the case of a gram matrices G_1, G_2 we rephrase search LIP as finding an unimodular transformation $U \in \mathcal{GL}_n(\mathbb{Z})$ such that $U^\top G_1 U = G_2$. This makes the orthonormal transformation irrelevant. Furthermore note that if we restrict to integral lattices, then this formulation of LIP only involves integer arithmetic.

The best provable algorithm to solve search LIP runs in time $n^{O(n)}$ [HR14]. Furthermore, the general algorithms for solving LIP [PP85, PS97, DSHVvW20] require as a first step the computation of short lattice vectors, which takes time $2^{O(n)}$. The high complexity of these algorithm is what makes LIP interesting as a hardness assumption.

Often however, the security proof of LIP based cryptographic schemes is not based on the search variant, but on a distinguishing variant. For any lattice $\mathcal{L}$ we will denote its isomorphism class by $[\mathcal{L}] = \{O \cdot \mathcal{L} : O \in \mathcal{O}_n(\mathbb{R})\}$. We can then ask to distinguish between different isomorphism classes.

Definition 5 (Distinguish LIP). *Let $\mathcal{L}_1, \mathcal{L}_2$ be non-isomorphic lattices. Given any lattice $\mathcal{L} \in [\mathcal{L}_b]$ for a uniformly random $b \leftarrow \mathcal{U}(\{1, 2\})$. Recover b.*

2 The genus of a lattice.

For any two isomorphic lattices $\mathcal{L}_1 \cong \mathcal{L}_2$ we have that $\mathrm{vol}(\mathcal{L}_1) = \mathrm{vol}(\mathcal{L}_2)$. The (co)volume of the lattice is thus an efficiently computable *invariant* for lattice isomorphisms. Other examples of this are the $\mathrm{scale}(\mathcal{L})$ or parity $\mathrm{par}(\mathcal{L})$ of an integral lattice. If two lattices have distinct invariants we can use that to solve distinguish LIP, simply by computing the same invariant for $\mathcal{L} \in [\mathcal{L}_b]$ and see if it matches that of $\mathcal{L}_1$ or $\mathcal{L}_2$. When instantiating a cryptographic scheme based on distinguishing LIP [DvW22, BGPSD23, BZI$^+$24, ARLW24] we thus should make sure that the non-isomorphic lattices $\mathcal{L}_1, \mathcal{L}_2$ match on all efficiently computable invariants.

In terms of quadratic forms the lattice isomorphism problem boils down to $\mathbb{Z}$-equivalence, which is seemingly hard. It is natural however to look at weaker forms of equivalence over larger rings $R \supset \mathbb{Z}$ which might be more efficient to compute.

Definition 6 (R-equivalence). *Let $R \supset \mathbb{Z}$ be any ring containing $\mathbb{Z}$. We say that two integral lattices $\mathcal{L}_1, \mathcal{L}_2 \subset \mathbb{R}^n$ are R-equivalent if $\mathcal{L}_1 \otimes_{\mathbb{Z}} R \cong \mathcal{L}_2 \otimes_{\mathbb{Z}} R$. Alternatively, two integral quadratic forms $G_1, G_2 \in \mathcal{S}_n^{>0}(\mathbb{Z})$ are R-equivalent if there exists a $U \in \mathcal{GL}_n(R)$ such that $U^\top G_1 U = G_2$ over R.*

One such weaker form of equivalence is that over the p-adic integers $\mathbb{Z}_p$. In contrast to $\mathbb{Z}$-equivalence it is efficient to compute if two integral lattices are $\mathbb{Z}_p$-equivalent for any prime p. In short, it follows from the fact that forms are (block-)diagonalizable over $\mathbb{Z}_p$, after which the equivalence is relatively easy to determine. See [CS99, Chapter 15.7] for more information on this computation and how to determine a complete set of invariants for $\mathbb{Z}_p$-equivalence. Furthermore, assuming that $\mathrm{vol}(\mathcal{L}_1) = \mathrm{vol}(\mathcal{L}_2)$, we only have to focus on those primes p that divide $2\,\mathrm{vol}(\mathcal{L}_i)^2$, as for all other primes the $\mathbb{Z}_p$-equivalence follows directly. Assuming that we know the factorization of $\mathrm{vol}(\mathcal{L}_i)^2$ we can thus determine the $\mathbb{Z}_p$-equivalence of $\mathcal{L}_1$ and $\mathcal{L}_2$ for all primes p.

Definition 7 (Genus [CS99, Chapter 15]). *The genus $\mathrm{gen}(\mathcal{L})$ of an integral lattice $\mathcal{L} \subset \mathbb{R}^n$ consists of all (integral) lattices of dimension n that are $\mathbb{Z}_p$-equivalent to $\mathcal{L}$ for all primes p. Given an integral lattice $\mathcal{L}$, and the prime factorization of $\mathrm{vol}(\mathcal{L})^2$, we can efficiently compute a canonical label of the genus it corresponds to.*

In case we are not only considering full-rank lattices there is an extra condition that the lattices must be equivalent over the reals. However, two full-rank lattices of the same dimension are always equivalent over $\mathbb{R}$ so we can safely ignore this condition.

Two lattices $\mathcal{L}_1 \cong \mathcal{L}_2$ that are $\mathbb{Z}$-equivalent are also $\mathbb{Z}_p$ equivalent for any prime p given that $\mathbb{Z} \subset \mathbb{Z}_p$, and thus $\mathrm{gen}(\mathcal{L}_1) = \mathrm{gen}(\mathcal{L}_2)$. In particular, if we have lattices $\mathcal{L}_1, \mathcal{L}_2$ such that $\mathrm{gen}(\mathcal{L}_1) \neq \mathrm{gen}(\mathcal{L}_2)$ it follows directly that they cannot be equivalent, and this can be efficiently computed. In this way, the genus of a lattice gives us a strong invariant for lattice isomorphisms. Note that as a result it is also well defined to speak about the genus of a $\mathbb{Z}$-equivalence class $[\mathcal{L}]$, and denote $\mathrm{gen}([\mathcal{L}]) := \mathrm{gen}(\mathcal{L})$. As far as we know the genus covers all the known efficiently computable invariants, which makes it interesting for us to study. As a result we can also simply define $\mathrm{vol}(\mathcal{G}) := \mathrm{vol}(\mathcal{L}), \mathrm{scale}(\mathcal{G}) := \mathrm{scale}(\mathcal{L})$ and $\mathrm{rk}_p(\mathcal{G}) := \mathrm{rk}_p(\mathcal{L})$ for $\mathcal{L} \in \mathcal{G}$ which is independent of the chosen representative.

Remark 1. For simplicity we only consider here the genus of *integral* lattices. Because the structure of the genus is invariant under integer scaling one could easily extend these notions to rational lattices. In particular, $\mathrm{scale}(\mathcal{L})$ is an invariant of the (rational) genus. More generally, similar notions exist for quadratic forms over the ring of integers of number fields [Wei65, Kir16].

2.1 Randomness over the Genus

Every genus consists of a finite number of $\mathbb{Z}$-equivalence classes [Min85] and thus one could consider a uniform distribution $\mathcal{U}(\mathcal{G})$ over it. However, mathematically, this is not the most natural distribution and we have to give each equivalence class a slightly different weight depending on the size of their automorphism group.

Definition 8 (Randomness over a genus). *For a genus $\mathcal{G}$ we define the probability distribution $\mathcal{D}(\mathcal{G})$ which samples $[\mathcal{L}] \in \mathcal{G}$ with relative* mass $m(\mathcal{L}) := 1/|\operatorname{Aut}(\mathcal{L})|$. *In particular, for any $[\mathcal{L}] \in \mathcal{G}$ we have*

$$\Pr_{[\mathcal{L}'] \leftarrow \mathcal{D}(\mathcal{G})}[\mathcal{L}' \cong \mathcal{L}] = \frac{m(\mathcal{L})}{\sum\limits_{[\mathcal{L}'] \in \mathcal{G}} m(\mathcal{L}')}.$$

Just as in the case of fully random lattice the measure $m(\mathcal{L}) = 1/\operatorname{Aut}(\mathcal{L})$ again follows naturally, this time from the Haar measure on $\mathcal{O}_n(\mathbb{R})$. More precisely, equip $\mathcal{O}_n(\mathbb{R})$ with its volume 1 Haar measure. The isometry class $[\mathcal{L}]$ is then endowed with an $\mathcal{O}_n(\mathbb{R})$-invariant measure m with total measure $m([\mathcal{L}]) = m(\mathcal{O}_n(\mathbb{R}) \cdot \mathcal{L}) = \frac{1}{|\operatorname{Aut}(\mathcal{L})|}$, because $\mathcal{L}$ is left invariant by precisely $|\operatorname{Aut}(\mathcal{L})|$ orthonormal transformations. Furthermore, this is precisely the distribution one gets when restricting the general probability distribution $\mu_{n,D}$ to the genus $\mathcal{G}$. In particular, if we sample a random $\mathcal{L}' \in \mu_{n,D}$ under the restriction that $\mathcal{L}' \in \mathcal{G}$, then the probability that $\mathcal{L}' \cong \mathcal{L}$ for some $\mathcal{L} \in \mathcal{G}$ is precisely $m(\mathcal{L})/\sum_{[\mathcal{L}'']} m(\mathcal{L}'')$.

Given one representative $\mathcal{L} \in \mathcal{G}$ in a genus there is a natural notion of p-neighbours within the same genus for any prime $p \nmid 2\operatorname{vol}(\mathcal{G})^2$, namely all those lattices $\mathcal{L}'$ in the same genus for which $\mathcal{L} \cap \mathcal{L}'$ has index p in both $\mathcal{L}$ and $\mathcal{L}'$. These connections turn the genus into a graph where the nodes are isomorphism classes $[\mathcal{L}]$ and the edges are p-neighbours. For large enough primes p this graph is furthermore connected[5]. By picking any of those (finite) p-neighbours uniformly at random one obtains a random walk over this graph. Another reason for the distribution $\mathcal{D}(\mathcal{G})$ to be natural is that for large enough p it is the natural limit or stationary distribution for this random walk, and this precisely allows us to sample efficiently from $\mathcal{D}(\mathcal{G})$ as stepping through this graph is efficient [Hei16]. In fact for large enough p, a single step is enough to be negligibly close to the distribution of $\mathcal{D}(\mathcal{G})$ [Che21], i.e., any isomorphism class is reached with relative weight $w(\mathcal{L})$.

Theorem 5 ([Hei16, Che21]). *There exists an efficient algorithm to sample from $\mathcal{D}(\mathcal{G})$.*

[5] We ignore here the rare case that the genus splits into multiple spinor-genera.

2.2 Siegel's Mass formulas

Given that the genus is an invariant under $\mathbb{Z}$-equivalence we can view any particular genus as a set of $\mathbb{Z}$-equivalence classes. This set is always finite, but more surprisingly we can even compute a notion of its size, i.e., its *mass*.

Theorem 6 (Smith-Siegel-Minkowski Mass formula [Sie35]). *Let*

$$M(\mathcal{G}) := \sum_{[\mathcal{L}] \in \mathrm{gen}(G)} \frac{1}{|\operatorname{Aut}(\mathcal{L})|}$$

be the mass of $\mathcal{G}$, where the sum is over all equivalence classes in the genus. Given as input a gram matrix G of any lattice $\mathcal{L} \in \mathcal{G}$, and the prime factorization of $\det(\mathcal{G})^2$*, the mass $M(\mathcal{G})$ can be computed in polynomial time in the input, n and* $\log(\mathrm{vol}(\mathcal{G}))$.

This directly gives an estimate for the number of equivalence classes.

Corollary 3. *Let $\mathcal{G}$ be a non-empty genus of dimension n and let $M(\mathcal{G})$ be its mass. Then the genus contains at least $2M(\mathcal{G})$ and for $n > 10$ at most $2^n \cdot n! \cdot M(\mathcal{G})$ distinct equivalence classes.*

Proof. Clearly $\{\pm I_n\} \subset \operatorname{Aut}(\mathcal{L})$ and thus $|\operatorname{Aut}(\mathcal{L})| \geq 2$ for any lattice. Furthermore, Feit [Fei96] showed in 1996 that with some exception for $n \leq 10$, $\mathbb{Z}^n$ has the largest automorphism group for any n-dimensional lattice $\mathcal{L}$ and thus $|\operatorname{Aut}(\mathcal{L})| \leq |\operatorname{Aut}(\mathbb{Z}^n)| = 2^n \cdot n!$ for any $n > 10$. The latter is one of the first significant results based on the classification of finite simple groups in an unpublished manuscript from 1984 of Weisfeiler [Wei84]. The result follows immediately from these bounds and the definition of the mass formula. □

We remark that asymptotically most lattices have a trivial automorphism group and thus we expect the number of equivalence classes to be quite close to $2M(\mathcal{G})$. While the existence of such a mass formula might already be surprising, one can go even further then this.

Note that computing the first minimum $\lambda_1(\mathcal{L})^2 = \arg\min_{k \geq 1}\{N_{\mathcal{L}}(k) > 0\}$ is already a hard problem. So to say anything about the theta series of a given lattice is very hard. However, once we start to look at the expected theta series over a genus, we suddenly are able to compute it efficiently.

Theorem 7 (Siegel's mass formula [Sie35]). *For a non-empty genus $\mathcal{G}$ we define the expected theta series as*

$$\Theta_{\mathcal{G}}(q) = \mathbb{E}_{[\mathcal{L}] \leftarrow \mathcal{D}(\mathcal{G})}\left[\theta_{\mathcal{L}}(q)\right] = \frac{\sum_{[\mathcal{L}] \in \mathcal{G}} \frac{1}{|\operatorname{Aut}(\mathcal{L})|} \cdot \theta_{\mathcal{L}}(q)}{\sum_{[\mathcal{L}] \in \mathcal{G}} \frac{1}{|\operatorname{Aut}(\mathcal{L})|}} =: 1 + \sum_{k=1}^{\infty} N_{\mathcal{G}}(k) \cdot q^k.$$

Given the prime factorizations of $\mathrm{vol}(\mathcal{G})^2$ *and $k > 0$ the coefficient $N_{\mathcal{G}}(k)$ of $\Theta_{\mathcal{G}}(q)$ can be efficiently computed.*

Remark 2. While these mass formulas are indeed computable in polynomial time, it is not necessarily and easy task to do it. In particular, the computations are very prone to errors. For an extensive explanation on how to compute the Smith-Siegel-Minkowski mass formula see [CS88]. In Section 4.1 we explain (partly) how to compute Siegel's mass formula. Both mass formulas have been implemented in Sagemath [The24]. For example, the total mass of a genus can be computed by calling `Q.conway_mass()` on a `QuadraticForm Q`.

3 On the existence of lattices with good properties

3.1 Lattices with good properties

The theta series give a lot of information about the geometric properties of a lattice. Due to this connection we can also hope that from the expected theta series over a genus, we can derive the existence of a lattice with good geometric properties within this genus.

One such example is to derive a good lattice packing. Given that the genus already fixed the determinant this means we want to find a lattice with a large first minimum $\lambda_1(\mathcal{L})$. Note that for such a lattice the theta series coefficients $N_1(\mathcal{L}), \ldots, N_{\lambda_1(\mathcal{L})^2-1}(\mathcal{L})$ are zero. In case the first few coefficients of the expected theta series are small we can conclude by a counting argument that at least one of the lattices must have only zeros there. Beyond existence, recall that for a non-negative random variable X and $a > 0$, Markov's inequality states that $\Pr[X < a] \geq 1 - \frac{\mathbb{E}[X]}{a}$, which can directly gives us a lower bound on the probability density of such lattices. This leads to a density analogue of the Minkowski-Hlawka Theorem restricted to a fixed genus.

Lemma 2 (Good packing density). *Let $\mathcal{G}$ be a genus with expected theta series $\Theta_{\mathcal{G}}(q) = 1 + \sum_{k=1}^{\infty} N_{\mathcal{G}}(k) q^k$. If $\sum_{k=1}^{\lambda-1} N_{\mathcal{G}}(k) < 2r$ for some $0 < r \leq 1$ and some integer $\lambda \geq 1$, then $\Pr_{[\mathcal{L}] \leftarrow \mathcal{D}(\mathcal{G})}[\lambda_1(\mathcal{L})^2 \geq \lambda] > 1 - r$. In particular, then there exists a lattice $\mathcal{L} \in \mathcal{G}$ such that $\lambda_1(\mathcal{L})^2 \geq \lambda$.*

Proof. We consider the non-negative random variable $\sum_{k=1}^{\lambda-1} N_{\mathcal{L}}(k)$ where $[\mathcal{L}] \leftarrow \mathcal{D}(\mathcal{G})$. By definition its expectation is given by $\sum_{k=1}^{\lambda-1} N_{\mathcal{G}}(k)$. By Markov's inequality we then obtain

$$\Pr_{[\mathcal{L}] \leftarrow \mathcal{D}(\mathcal{G})}\left[\sum_{k=1}^{\lambda-1} N_{\mathcal{L}}(k) < 2\right] \geq 1 - \frac{\sum_{k-1}^{\lambda-1} N_{\mathcal{G}}(k)}{2} > 1 - \frac{2r}{2} = 1 - r,$$

from which the result follows as $\sum_{k=1}^{\lambda-1} N_{\mathcal{L}}(k) < 2$ if and only if $\lambda_1(\mathcal{L})^2 \geq \lambda$. For the existence result note that the probability is strictly positive for $r \leq 1$. □

Remark 3. One can observe that Lemma 2 proves a stronger statement than merely the existence of a good packing. It also gives a lower bound on the probability that any lattice sampled from $\mathcal{D}(\mathcal{G})$ achieves a certain minimum distance. Moreover, this can be turned into a quantitative statement on the number of such lattices by considering the mass $M(\mathcal{G})$ of the genus. In particular, if $\sum_{k=1}^{\lambda-1} N_{\mathcal{G}}(k) < 2r$ for $0 < r \leq 1$ then there exist at least $2(1-r)M(\mathcal{G})$ non-isomorphic lattices $\mathcal{L} \in \mathcal{G}$ such that $\lambda_1(\mathcal{L})^2 \geq \lambda$.

Remark 4. Just as for general random lattices there exists a variant of Siegel's mass formula that computes the number of *primitive* vectors of squared norm k (see e.g. [Han04]). Clearly, Lemma 2 works just a well with these quantities. For large n however, it does not seem to make a large difference (just as for the Hlawka-Minkowski Theorem), so we do not consider this small improvement.

To obtain a good dual lattice packing one can apply the same Lemma to the (scaled) dual theta series. Note that one could even apply the same proof to the primal and dual theta series simultaneously to obtain a single lattice with both a good primal and dual packing.

Lemma 3 (Good primal and dual packing). *Let $\mathcal{G}$ be a genus with expected theta series $\Theta_{\mathcal{G}}(q) = 1 + \sum_{k=1}^{\infty} N_{\mathcal{G}}(k)q^k$, and let $c\mathcal{G}^{-1}$ for $c = \mathrm{scale}(\mathcal{G}^{-1})^{-1} \in \mathbb{Q}$ be the integral scaled dual genus with expected theta series $\Theta_{c\mathcal{G}^{-1}}(q) = 1 + \sum_{k=1}^{\infty} N_{c\mathcal{G}^{-1}}(k)q^k$. Let $0 < r \leq 1$ and let $\lambda, \lambda' \geq 1$ be integers. If $\sum_{k=1}^{\lambda-1} N_{\mathcal{G}}(k) + \sum_{k=1}^{\lambda'-1} N_{c\mathcal{G}^{-1}}(k) < 2r$, then $\Pr_{[\mathcal{L}]\leftarrow\mathcal{D}(\mathcal{G})}\left[\lambda_1(\mathcal{L})^2 \geq \lambda \text{ and } \lambda_1(\mathcal{L}^*)^2 \geq \frac{\lambda'}{c}\right] > 1-r$. In particular, then there exists a lattice $\mathcal{L} \in \mathcal{G}$ with $\lambda_1(\mathcal{L})^2 \geq \lambda$ and $\lambda_1(\mathcal{L}^*)^2 \geq \frac{\lambda'}{c}$.*

The existence of a good dual packing immediately also implies the existence of a lattice with good smoothing as for $\varepsilon \in (0, e^{-n}]$ we have $\eta_\varepsilon(\mathcal{L}) \leq \sqrt{\ln(1/\varepsilon)}/\lambda_1(\mathcal{L}^*)$. Usually however, we are interested in the smoothing for larger values of ε. In that case we can consider the following result.

Lemma 4 (Good smoothing parameter). *Let $\varepsilon > 0$, $0 < r \leq 1$ and let $s > 0$ be such that $\Theta_{\mathcal{G}}(\exp(-\pi s^2)) < 1 + r\varepsilon$, then $\Pr_{[\mathcal{L}]\leftarrow\mathcal{D}(\mathcal{G})}[\eta_\varepsilon(\mathcal{L}^*) < s] > 1-r$. In particular, then there exists a lattice $\mathcal{L} \in \mathcal{G}$ such that $\eta_\varepsilon(\mathcal{L}^*) < s$.*

Proof. Note that $\Theta_{\mathcal{G}}(\exp(-\pi s^2)) - 1$ is the expectation of the non-negative random variable $\Theta_{\mathcal{L}}(\exp(-\pi s^2)) - 1$ where $[\mathcal{L}] \leftarrow \mathcal{D}(\mathcal{G})$. By Markov's inequality we then obtain

$$\Pr_{[\mathcal{L}]\leftarrow\mathcal{D}(\mathcal{G})}\left[\Theta_{\mathcal{L}}(\exp(-\pi s^2)) - 1 < \varepsilon\right] \geq 1 - \frac{\Theta_{\mathcal{G}}(\exp(-\pi s^2)) - 1}{\varepsilon} > 1 - \frac{r\varepsilon}{\varepsilon} = 1 - r,$$

from which the result follows as by definition $\Theta_{\mathcal{L}}(\exp(-\pi s^2)) < 1 + \varepsilon$ if and only if $\eta_\varepsilon((\mathcal{L}')^*) < s$. For the existence result note that the probability is strictly positive for $r \leq 1$. □

3.2 Example: unimodular lattices

Let us consider the easiest, but in some sense also most interesting genera, those of unimodular lattices. These lattices are self-dual, which protects them from Hull attacks like [DG23]. In addition, due to their small determinant one could in principle describe them using small matrix entries, potentially leading to smaller keys. Because unimodular lattices have determinant 1 we only have to focus on the p-adic equivalence for $p = 2$. Furthermore, because 2 does not divide the determinant the full 2-adic equivalence is determined by the parity $\mathrm{par}(\mathcal{L}) \in \{1, 2\}$ of these unimodular lattices. We thus obtain two genera, the *odd* and *even* one. Here we consider as an example the even case for which all the vectors $v \in \mathcal{L}$ have an even squared norm $\|v\|^2$.

The even case includes the famous root lattice E_8 and the Leech lattice Λ_{24}. Even unimodular lattices only exist in dimensions that are a multiple of 8 (see e.g. [Ser73, p. 53]). For even unimodular lattices the expected theta series is given by a rational scaling of the q-expansion of the Eisenstein series [MH$^+$73]. To be more precise, let $n \geq 8$ with $8|n$, and let $\mathcal{G}_{n,e}$ be the genus of even unimodular lattices of dimension $n = 8m$, then we have

$$\Theta_{\mathcal{G}_{8m,e}}(q) = E_{4m}(q^2) = 1 + \frac{-8m}{B_{4m}} \sum_{k=1}^{\infty} \sigma_{4m-1}(k) q^{2k},$$

where B_i is the i-th Bernoulli number, and $\sigma_z(m) = \sum_{d|m} d^z$ is the sum of positive divisors function.

Using the above expected theta series and Lemma 2 we can prove the existence of an even unimodular lattice with first minimum essentially as indicated by the Gaussian Heuristic.

Lemma 5 (Even packing). *Let $n = 8m \geq 8$ with $m \in \mathbb{N}$, then there exists an n-dimensional even unimodular lattice $\mathcal{L}$ with $\lambda_1(\mathcal{L})^2 \geq 2 \left\lceil \frac{1}{2} \cdot \left(\frac{3\zeta(n/2)}{2\omega_n}\right)^{2/n} \right\rceil \approx n/2\pi e$.*

Proof. Let $k' = \left\lceil \frac{1}{2} \cdot \left(\frac{3\zeta(4m)}{2\omega_{8m}}\right)^{1/4m} \right\rceil$. To apply Lemma 2 we need to show that the sum of the first $k' - 1$ non-trivial coefficients of $\Theta_{\mathcal{G}_{8m,e}}(q)$ is bounded by 2. Recall that these have values $\frac{-8m}{B_{4m}} \sigma_{4m-1}(k)$ for $k = 1, \ldots, k' - 1$. First, note that

$$\begin{aligned} \sum_{k=1}^{k'-1} \sigma_{4m-1}(k) = \sum_{k=1}^{k'-1} \sum_{d|k} d^{4m-1} &= \sum_{d=1}^{k'-1} \left\lfloor \frac{k'-1}{d} \right\rfloor d^{4m-1} \leq \sum_{d=1}^{k'-1} (k'-1) \cdot d^{4m-2} \\ &\leq (k'-1) \cdot \frac{1}{4m-1} \left(k' - \frac{1}{2}\right)^{4m-1} < \frac{(k' - \frac{1}{2})^{4m}}{4m-1} \\ &\leq \frac{3\zeta(4m) \cdot \omega_{8m}^{-1}}{2^{4m+1} \cdot (4m-1)}, \end{aligned}$$

where we use that $d^{4m-2} \leq \int_{d-\frac{1}{2}}^{d+\frac{1}{2}} x^{4m-2} dx$. Secondly, by using the common identity for even Bernoulli numbers and the volume ω_{8m} of an $8m$-dimensional unit ball, we get

$$\frac{-8m}{B_{4m}} = 8m\omega_{8m} \cdot \frac{2^{4m-1}}{\zeta(4m)}.$$

Combining the two we get that

$$\frac{-8m}{B_{4m}} \cdot \sum_{k=1}^{k'-1} \sigma_{4m-1}(k) < 8m \cdot \omega_{8m} \cdot \frac{2^{4m-1}}{\zeta(4m)} \cdot \frac{3\zeta(4m) \cdot \omega_{8m}^{-1}}{2^{4m+1} \cdot (4m-1)} = \frac{3}{2} \cdot \frac{4m}{4m-1} \leq 2.$$

We can conclude by Lemma 2 and the even parity that there exists a lattice $\mathcal{L} \in \mathcal{G}_{n,e}$ with minimum $\lambda_1(\mathcal{L})^2 \geq 2k'$. □

Remark 5. We want to emphasize that this result is not novel. The bound in Lemma 5 is essentially the same as claimed by Milnor [MH+73, p. 47], where a lower bound of $2 \cdot \lceil \frac{1}{2}(\frac{3}{5}\omega_n)^{-2/n} \rfloor$ is given based on computations in [Ser73]. The proof here uses a different representation of the Eisenstein series and is more elementary. For a concrete comparion see Fig. 1. For odd unimodular lattices Milnor [MH+73, p. 46] gives a full proof for a lower bound of $\lambda_1(\mathcal{L})^2 \geq \lceil (\frac{3}{5}\omega_n)^{-2/n} \rfloor$.

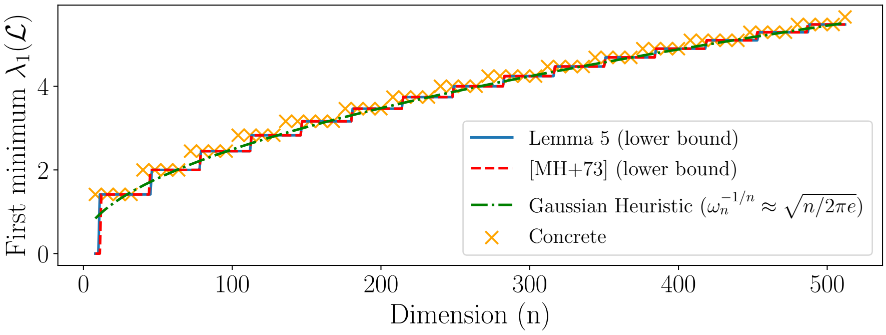

Fig. 1. First minimum guarantee for even unimodular lattices as given by *Lemma* 5 and [MH+73], compared to concrete values directly computed using $\Theta_{\mathcal{G}_{n,e}}$ and *Lemma* 2.

For the smoothing parameter we can also prove a result that is essentially tight, and similar to Corollary 2 for random lattices. We restrict ourselves to $\varepsilon \geq \Omega(e^{-n})$ as the general bounds based on the dual minimum distance often fail to give tight results in this important regime for cryptography.

First we require two small technical lemmas.

Lemma 6 (Technical lemma I). *For $x \geq 2$ and any integer $y \geq 1$ we have $\sigma_x(y) \leq \zeta(x) \cdot y^x$, where $\sigma_x(y) = \sum_{d|y} d^x$ is the sum of positive divisors function.*

Proof. Let $y = p_1^{a_1} \cdots p_n^{a_n}$, be the prime factorization of y with $p_1, \ldots, p_n$ distinct primes and $n \geq 0, a_i > 0$. Now for any prime power p^a we have $\sigma_x(p^a) = \sum_{d|p^a} d^x = 1 + p^x + \ldots + p^{ax} = p^{ax} \cdot \frac{1-p^{-(a+1)x}}{1-p^{-x}} \leq p^{ax} \cdot \frac{1}{1-p^{-x}}$. The sum of divisors function σ_x is multiplicative for coprime inputs and thus we get

$$\sigma_x(y) = \prod_{i=1}^{n} \sigma_x(p_i^{a_i}) \leq \prod_{i=1}^{n} p_i^{a_i x} \cdot \frac{1}{1-p^{-x}} \leq y^x \cdot \prod_{p \text{ prime}} \frac{1}{1-p^{-x}} = \zeta(x) \cdot y^x.$$

□

Lemma 7 (Technical lemma II). *Let $0 < c \leq C$ for some constant $C > 0$, then we have for $x \geq 2$ that*

$$\mathrm{Li}_{-x}(\exp(-c)) \leq \left(1 + 2 \cdot \sum_{k=1}^{\infty} (1 + 4k^2\pi^2/C^2)^{-3/2})\right) \cdot \Gamma(1+x) \cdot c^{-x-1},$$

where $\mathrm{Li}_y(z) = \sum_{k=1}^{\infty} \frac{z^k}{k^y}$ is the polylogarithm function.

Proof. For negative $y < 0$ we have the following identity by Wood [Woo92, (13.1)]:

$$\begin{aligned}\mathrm{Li}_{-x}(\exp(-c)) &= \Gamma(x+1) \cdot \sum_{k=-\infty}^{\infty} (2k\pi i + c)^{-x-1} \\ &= \Gamma(x+1) \cdot c^{-x-1} \cdot \sum_{k=-\infty}^{\infty} (2k\pi i/c + 1)^{-x-1}.\end{aligned}$$

The summation is real-valued as the terms $\pm k$ are conjugates, and we have

$$\begin{aligned}\sum_{k=-\infty}^{\infty} (2k\pi i/c + 1)^{-x-1} &\leq 1 + 2 \cdot \sum_{k=1}^{\infty} |2k\pi i/c + 1|^{-x-1} \\ &= 1 + 2 \cdot \sum_{k=1}^{\infty} (1 + 4k^2\pi^2/c^2)^{-(x+1)/2} \\ &\leq 1 + 2 \cdot \sum_{k=1}^{\infty} (1 + 4k^2\pi^2/C^2)^{-1.5}.\end{aligned}$$

□

Lemma 8 (Even smoothing). *Let $n = 8m \geq 8$ with $m \in \mathbb{N}$, $C = 17.8$, and let $\varepsilon > C \cdot e^{-n}$, then there exists an n-dimensional even unimodular lattice $\mathcal{L} \in \mathcal{G}_{n,e}$ such that $\eta_\epsilon(\mathcal{L}) \leq (C/\epsilon)^{1/n}$.*

Proof. Let $s = (C/\epsilon)^{1/8m} \leq e$. To apply Lemma 4 we have to show that the following sum is bounded by ε:

$$\mu = \frac{-8m}{B_{4m}} \cdot \sum_{k=1}^{\infty} \sigma_{4m-1}(k) \cdot \exp(-2\pi s^2 k).$$

First, using Lemma 6 we have the bound $\sigma_{4m-1}(k) \leq \zeta(4m-1) \cdot k^{4m-1} \leq \zeta(3)k^{4m-1}$. This gives us that

$$\mu \leq \frac{-8m\zeta(3)}{B_{4m}} \cdot \sum_{k=1}^{\infty} k^{4m-1} \cdot \exp(-2\pi s^2 k) = \frac{-8m\zeta(3)}{B_{4m}} \mathrm{Li}_{1-4m}(e^{-2\pi s^2}),$$

where $\mathrm{Li}_p(z) = \sum_{k=1}^{\infty} \frac{z^k}{k^p}$ is the polylogarithm function. From Lemma 7 we get that $\mathrm{Li}_{1-4m}(e^{-2\pi s^2}) \leq 14.78 \cdot \Gamma(4m) \cdot (2\pi s^2)^{-4m}$, which combined with the identity for the even Bernoulli numbers gives us

$$\mu \leq \frac{-8m\zeta(3)}{B_{4m}} \mathrm{Li}_{1-4m}(e^{-2\pi s^2}) \leq \zeta(3) \cdot 8m \cdot \frac{(2\pi)^{4m}}{2 \cdot (4m)!} \cdot 14.78 \cdot \Gamma(4m) \cdot (2\pi s^2)^{-4m}$$
$$= \zeta(3) \cdot 14.78 \cdot s^{-8m} \leq C \cdot \frac{\varepsilon}{C} = \varepsilon.$$

We conclude by Lemma 4 and the fact that unimodular lattices are self-dual. □

We note that in principle the above Lemma is not restrained to $\varepsilon \geq Ce^{-n}$ and could easily be adapted to handle $\varepsilon < C \cdot e^{-n}$ at the cost of a larger constant C. In particular, by slightly adapting Lemma 7 we can obtain a bound of $(3\zeta(3)/\varepsilon)^{1/(n-2)}$ that is valid for any $3\zeta(3) \geq \varepsilon > 0$. Numerical evidence indicates that the bound from Lemma 7 could be improved further leading to a lower constant C both here and for Theorem 2. However in this regime the bound obtained from a good dual packing is better than the one given here. In Fig. 2 we can see that the bound in Lemma 8 is rather tight compared to a direct application of Lemma 4.

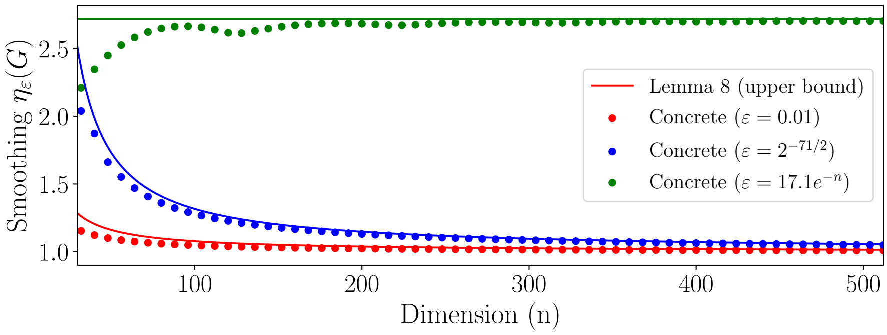

Fig. 2. Smoothing bound for even unimodular lattices as given by *Lemma* 8, compared to concrete values directly computed using $\Theta_{\mathcal{G}_{n,e}}$ and *Lemma* 4. The value $\varepsilon = 2^{-71/2} = 1/\sqrt{q_s \cdot \lambda}$ is common in hash-and-sign schemes with $\lambda = 128$ bits of security that can sign 2^{64} signatures.

4 A general result

We now consider the general case for almost all genera. By our knowledge existing literature only show packing results for the (simpler) unimodular case, but the results do generalize.

Theorem 1 (General packing). *For any integral genus $\mathcal{G}$ in dimension $n \geq 6$ such that $\mathrm{rk}_p(\mathcal{G}) \geq 6$ for all primes p, and any constant $0 < c \leq 1$, we have*

$$\Pr_{[\mathcal{L}] \leftarrow \mathcal{D}(\mathcal{G})}\left[\lambda_1(\mathcal{L})^2 \geq \left[c^2 \cdot \left(\frac{7\zeta(3)}{9\zeta(2)} \cdot \frac{\mathrm{vol}(\mathcal{L})}{\omega_n}\right)^{2/n}\right]\right] > 1 - c^n.$$

In particular, there exists a lattice $\mathcal{L} \in \mathcal{G}$ with

$$\lambda_1(\mathcal{L})^2 \geq \left\lceil \left(\frac{7\zeta(3)}{9\zeta(2)} \cdot \frac{\mathrm{vol}(\mathcal{L})}{\omega_n}\right)^{2/n} \right\rceil \approx n/2\pi e \cdot \mathrm{vol}(\mathcal{L})^{2/n}.$$

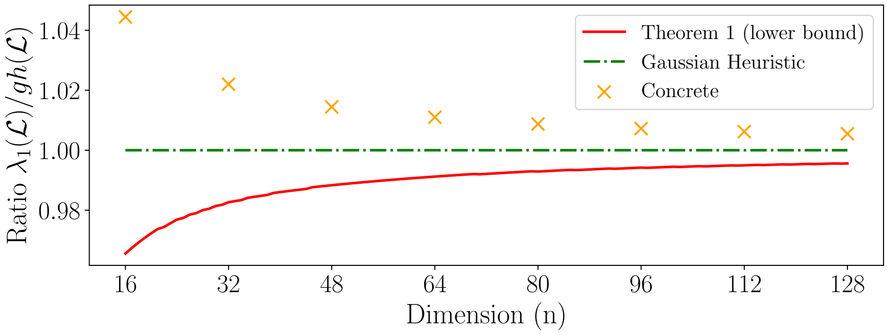

Fig. 3. The asymptotic existence bound of Theorem 1 versus a concrete computation for the genus of the lattice $\mathbb{Z}^k \oplus 521\mathbb{Z}^k$. This genus contains a large class of random q-ary lattices [BDG23] or of NTRU lattices with $q = 521$. All values are normalized by the Gaussian Heuristic.

In Fig. 3 we demonstrate that our asymptotic existence bound in Theorem 1 is close to the Gaussian Heuristic and to concrete bounds obtained via computing Siegel's mass formula exactly. In particular, their ratio quickly goes to 1 when the dimension n increases.

Similarly, we can show the existence of a lattice with a good smoothing parameter and covering radius in any genus.

Theorem 2 (General smoothing). *For any integral genus $\mathcal{G}$ in dimension $n \geq 6$ such that $\mathrm{rk}_p(\mathcal{G}) \geq 6$ for all primes p, constants $C = 26.1$ and $0 < c \leq 1$,* *and* $\epsilon \geq C \cdot (ce)^{-n} \cdot \mathrm{vol}(\mathcal{G})^{-1}$*, we have*

$$\Pr_{[\mathcal{L}] \leftarrow \mathcal{D}(\mathcal{G})} \left[\eta_\epsilon(\mathcal{L}^*) \leq \frac{1}{c} \cdot \left(\frac{C \cdot \mathrm{vol}(\mathcal{L}^*)}{\epsilon} \right)^{1/n} \right] > 1 - c^n.$$

In particular, there exists a lattice $\mathcal{L} \in \mathcal{G}$ *such that* $\eta_\epsilon(\mathcal{L}^*) \leq (C \cdot \mathrm{vol}(\mathcal{L}^*)/\epsilon)^{1/n}$ *.*

The proofs of Theorems 1 and 2 are stated in Section 4.3 after some preliminary definitions and results in Sections 4.1 and 4.2.

Theorem 3 (General covering radius). *For any integral genus* $\mathcal{G}$ *in dimension* $n \geq 6$ *such that* $\mathrm{rk}_p(\mathcal{G}) \geq 6$ *for all primes* p*, and constants* $C = 26.1$ *and* $e^{-1}(2C/(3\,\mathrm{vol}(\mathcal{G})))^{1/n} \leq c \leq 1$*, we have*

$$\Pr_{[\mathcal{L}] \leftarrow \mathcal{D}(\mathcal{G})} \left[v(\mathcal{L}^*) \leq \frac{1}{c} \cdot \left(\sqrt{n/2\pi} + 1 \right) \cdot \left(\frac{2}{3} C \cdot \mathrm{vol}(\mathcal{L}^*) \right)^{1/n} \right] > 1 - c^n.$$

In particular, when additionally $n \geq 7$*, there exists a lattice* $\mathcal{L} \in \mathcal{G}$ *such that*

$$v(\mathcal{L}^*) \leq \left(\sqrt{n/2\pi} + 1 \right) \cdot \left(\frac{2}{3} C \cdot \mathrm{vol}(\mathcal{L}^*) \right)^{1/n} \approx \sqrt{e} \cdot \sqrt{n/2\pi e} \cdot \mathrm{vol}(\mathcal{L}^*)^{1/n}.$$

Proof. We combine Theorem 2 for $\varepsilon = \frac{3}{2}$ and Lemma 1. The constraint on c is equivalent to the constraint on $\varepsilon = \frac{3}{2}$ in *Theorem* 2. For the existence claim note that if $n \geq 7$ then $e^{-1}(2C/(3\,\mathrm{vol}(\mathcal{G})))^{1/n} < 1$ and the result follows from the strictly positive probability for $c = 1$. □

Remark 6. We expect that the condition on $\mathrm{rk}_p(\mathcal{G})$ could be removed at the cost of a minor loss in the bound and a more tedious proof. Note that the condition is not satisfied by a genus with $\mathrm{scale}(\mathcal{G}) > 1$, however one can always circumvent this by first normalizing the genus before applying the result.

Remark 7. By choosing $c = 3^{-n}$ in Theorems 1 to 3 we get a probability of strictly more than $\frac{2}{3}$ for each property. In particular, this implies the existence of a lattice $\mathcal{L} \in \mathcal{G}$ having a good packing density and good dual smoothing and covering

4.1 Computing Siegel's mass formula

What makes testing equivalence over $\mathbb{Z}_p$ easy is the fact that we can efficiently (block) diagonalize forms over $\mathbb{Z}_p$.

Lemma 9 ([CS99, p.370]). *For* $p \neq 2$ *every integral form* $G \in \mathcal{S}_n^{>0}(\mathbb{Z})$ *is* $\mathbb{Z}_p$*-equivalent to a diagonal matrix. For* $p = 2$ *every integral form* $G \in \mathcal{S}_n^{>0}(\mathbb{Z})$ *is* $\mathbb{Z}_2$*-equivalent to a block diagonal matrix with blocks*

$$\begin{pmatrix} qx \end{pmatrix}, \quad \begin{pmatrix} qa & qb \\ qb & qc \end{pmatrix},$$

where q is a power of 2, a and c are divisible by 2, but x, b and $d = ac - b^2$ are not. As a corollary, for any prime p the form $G \in \mathcal{S}_n^{>0}(\mathbb{Z})$ is equivalent over $\mathbb{Z}_p$ to a decomposition

$$G_1 \oplus pG_p \oplus p^2 G_{p^2} \oplus \ldots \oplus qG_q \oplus \ldots$$

where $p \nmid \det(G_q)$ for all $q = p^i$, all but a finite number of the G_{p^i} have dimension 0, and each G_q is (block) diagonalized.

The (block) diagonalization is not necessarily unique but can be made canonical with some additional rules [CS99]. Testing for $\mathbb{Z}_p$-equivalence then simply becomes testing for equality. Note that because scale(.) is a genus invariant we can normalize (all forms in) a genus just as we can normalize individual forms. Then for for any form in a normalized genus $\mathcal{G}$ the first block G_1 has by construction a nonzero dimension which coincides precisely with having a p-rank $\mathrm{rk}_p(\mathcal{G}) \geq 1$. For our main results we will require that $\mathrm{rk}_p(\mathcal{G}) \geq 6$, which simply states that the first block isn't too small.

Recall that the k-th coefficient of Siegel's mass formula computes the expected number of integer solutions to $x^\top Gx$ for a random form G in a fixed genus. By a local-global principle the number of such solutions is related to the density of such solutions over the localization $\mathbb{Z}_p$.

Definition 9 (Local density [Sie35]). *For an integral form $G \in \mathcal{S}_n^{>0}(\mathbb{Z})$, prime p and integers $k, j \geq 0$ we denote*

$$N_G(k \bmod p^j) := |\{x \in (\mathbb{Z}/p^j\mathbb{Z})^n : x^\top Gx \equiv k \bmod p^j\}|.$$

for the number of distinct solutions of $x^\top Gx = k \bmod p^j$. The average number of solutions over all values $k = 0, \ldots, p^j - 1$ is given by $p^{(n-1)j}$ and we denote

$$\delta_G(k \bmod p^j) := \frac{N_G(k \bmod p^j)}{p^{(n-1)j}}$$

for the relative density of solutions. For $k \geq 1$ the following limit exists and is finite

$$\delta_{G,p}(k) := \lim_{j \to \infty} \delta_G(k \bmod p^j).$$

For $k = 0$ the limit does not always exist so we define

$$\delta_{G,p}(0) := \limsup_{j \to \infty} \delta_G(0 \bmod p^j),$$

which might be ∞. We call $\delta_{G,p}(k)$ the local density over $\mathbb{Z}_p$ at k.

The number of local solutions and therefore the density over $\mathbb{Z}_p$ is invariant under $\mathbb{Z}_p$-equivalence. We therefore also denote $\delta_{\mathcal{G},p}(k) := \delta_{G,p}(k)$ for any form G in the genus $\mathcal{G}$. We also consider one last local density over the reals, which is often also called 'the prime at infinity' or at -1.

Definition 10 (Local density at $p = \infty$). *For an integral genus $\mathcal{G}$ of dimension n, and any $k \geq 1$ we denote*

$$\delta_{\mathcal{G},\infty}(k) = \mathrm{vol}(\mathcal{G})^{-1} \cdot \frac{1}{2} n \omega_n k^{n/2-1},$$

for the local density over the reals $\mathbb{R}$ (or $p = \infty$).

We can now state the local-global result of Siegel that allows us to express the coefficients of Siegel's mass formula in terms of the local densities.

Theorem 8 (Siegel [Sie35]). *Let $\mathcal{G}$ be a genus of dimension $n \geq 3$, and let $\Theta_{\mathcal{G}}(q) = 1 + \sum_{k=1}^{\infty} N_{\mathcal{G}}(k) q^k$ be the expected theta series over $\mathcal{G}$. Then for all $k \geq 1$ we have*

$$N_{\mathcal{G}}(k) = \prod_{p=2,3,5,\dots,\infty} \delta_{\mathcal{G},p}(k).$$

This product converges and is 0 if and only if at least one of the factors is 0. Furthermore, this product is efficiently computable given the prime factorization of k and $\mathrm{vol}(\mathcal{G})^2$.

Recall that to determine if two integral forms lie in the same genus we only have to compute something for each prime divisor of $2\,\mathrm{vol}(\mathcal{G})^2$, as they are automatically equivalent over the other primes (if their volume matches). We have a similar property here that we only have to compute the local densities for $p = \infty$ and the primes dividing $2k\,\mathrm{vol}(\mathcal{G})^2$. For these primes we can get a (block) diagonalized representative over $\mathbb{Z}_p$, and from this there are efficient recursive formulas to compute the local density $\delta_{\mathcal{G},p}(k)$ (see [Han04]). For the other primes the local density is easy to express and their total infinite product can be computed efficiently using a series identity.

Siegel's Theorem is also valid when restricting to primitive (global and local) solutions which could in theory give slightly better packing bounds. For simplicity we do not consider this case.

Remark 8. These formulas are implemented in Sagemath [The24] by Hanke [Han04] and others. For a `QuadraticForm` object `Q`, one can call `Q.local_density(p, k)` to compute $\delta_{Q,p}(k)$.[6] Furthermore, the whole product at $k \geq 1$ is computed as `Q.siegel_product(k)`. Note, for lattices with even parity this actually computes the local density for $Q/2$ because Sagemath uses a different normalization. These functions allow to compute Siegel's mass formula explicitly and get better concrete bounds directly based on Lemmas 2 and 4.

[6] The current Sagemath implementation for computing the local density $\delta_{Q,p}(k)$ at $p = 2$ follows a naive brute-force approach and therefore becomes infeasible to compute for dimensions as low as $n \geq 8$. In our artifact available at `https://github.com/WvanWoerden/siegel_asiacrypt_artifact` we supply a patch that resolves this issue and we aim at integrating this fix into Sagemath.

4.2 Bounding the local densities

We will use Theorem 8 to bound the coefficients $N_{\mathcal{G}}(k)$ of the expected theta series over the genus $\mathcal{G}$. We will show that under light conditions the local density at each finite prime is bounded sufficiently such that the magnitude of the product of local densities is mostly driven by $\delta_{\mathcal{G},\infty}$.

The general idea is to use the orthogonal decomposition we get from Lemma 9 and show that if we have $G = G_1 \oplus G_2$ we only need to bound the local density of G_1 to bound the local density of G. The reason for this is that all solutions to $x^\top Gx = k \bmod p^j$ come from solutions $x_i^\top G_i x_i = k_i \bmod p^j$ for $k_1 + k_2 = k \bmod p^j$. The local densities of G are thus an averaged out version of the local densities of G_1 and G_2.

Lemma 10 (Decompose and conquer). *Let p be a prime, G a form of dimension n over the p-adic integers, and suppose that $G = G_1 \oplus G_2$ can be written as an orthogonal sum of non-trivial G_1 and G_2. Then for a constant $C > 0$ we have*

$$\forall \text{ integers } k' \geq 0, \delta_{G_1,p}(k') \leq C \implies \forall \text{ integers } k \geq 0, \delta_{G,p}(k) \leq C.$$

Proof. Note that we can express the number of solutions of G in terms of G_1 and G_2 as follows

$$N_G(k \bmod p^j) = \sum_{\substack{k_1,k_2 \in \mathbb{Z}/p^j\mathbb{Z}, s.t. \\ k_1+k_2 \equiv k \bmod p^j}} N_{G_1}(k_1 \bmod p^j) \cdot N_{G_2}(k_2 \bmod p^j).$$

Dividing both sides by $p^{(n-1)j}$ gives us

$$\begin{aligned} \delta_G(k \bmod p^j) &= p^{-j} \sum_{\substack{k_1,k_2 \in \mathbb{Z}/p^j\mathbb{Z}, s.t. \\ k_1+k_2 \equiv k \bmod p^j}} \delta_{G_1}(k_1 \bmod p^j) \cdot \delta_{G_2}(k_2 \bmod p^j) \\ &\leq \max_{k_1} \delta_{G_1}(k_1 \bmod p^j) \cdot \left(p^{-j} \sum_{k_2 \in \mathbb{Z}/p^j\mathbb{Z}} \delta_{G_2}(k_2 \bmod p^j) \right) \\ &= \max_{k_1 \in \mathbb{Z}/p^j\mathbb{Z}} \delta_{G_1}(k_1 \bmod p^j) \end{aligned}$$

Taking the limit $j \to \infty$ we obtain a supremum over $\delta_{G_1,p}(k_1)$ for all $k_1 \geq 0$ each of which is bounded by C. So $\delta_G(k) \leq C$. □

Lemma 11 (Classification of local normal forms [CS99, Chapter 15.7]). *Let p be a prime. Let $G \in \mathcal{S}_n^{>0}(\mathbb{Z})$ be an integral form of dimension $n \geq 4$ and such that $p \nmid \det(G)$. Then G is equivalent over $\mathbb{Z}_p$ to*

$$\left(I_n\right), or \quad \left[\begin{array}{c|c} u & \\ \hline & I_{n-1} \end{array}\right], \quad with \ \left(\frac{u}{p}\right) = -1, when\ p\ is\ an\ odd\ prime,\ or$$

$$\left[\begin{array}{c|c} -I_3 & \\ \hline & I_{n-3} \end{array}\right], \textit{or} \left[\begin{array}{cc|c} a & 0 & \\ 0 & b & \\ \hline & & I_{n-2} \end{array}\right], \textit{with } a, b \in \{\pm 1, \pm 3\}, \textit{when } p = 2, \operatorname{par}(G) = 1, \textit{or}$$

$$\begin{bmatrix} \boxed{B} & & & & \\ & \boxed{B} & & & \\ & & \boxed{A} & & \\ & & & \ddots & \\ & & & & \boxed{A} \end{bmatrix}, \textit{or} \begin{bmatrix} \boxed{B} & & & \\ & \boxed{A} & & \\ & & \ddots & \\ & & & \boxed{A} \end{bmatrix}, \quad \begin{array}{c} \textit{when } p = 2, \textit{ and } \operatorname{par}(G) = 2, \textit{ where} \\ A := \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \quad B := \begin{bmatrix} 2 & 1 \\ 1 & 2 \end{bmatrix}. \end{array}$$

Proof. This follows from the classification of a complete set of invariants for $\mathbb{Z}_p$-equivalence extensively discussed in [CS99, Chapter 15.7]. We shortly repeat it here for the case that $p \nmid \det(G)$. If $p \neq 2$ then the $\mathbb{Z}_p$ equivalence is fully determined by the dimension and the *sign* given by the Legendre symbol $\left(\frac{\det(G)}{p}\right) \in \{\pm 1\}$. For any dimension n we thus have two cases which can be represented by the forms in the statement.

The case $p = 2$ is a bit more complicated, first the sign is 1 if $\det(G) = \pm 1 \bmod 8$ and -1 if $\det(G) = \pm 3 \bmod 8$. Then we have the parity of G which is 1 if there is at least one odd entry on the diagonal, and otherwise 2. If the parity is 1 then we can fully diagonalize the form and we can consider the *oddity* t that is the sum of the diagonal modulo 8. Note that $t = n \bmod 2$ as the diagonal consists of odd entries, so there are 4 possibilities left. Combined with the sign there are thus 8 pairs of values for each dimension n and one can check that each of those are attained by the representatives in the statement.

Finally, for parity 2 we only have to consider the sign, giving the two cases in the statement. □

Now if we assume that $\operatorname{rk}_p(\mathcal{G}) \geq 4$ then we can assume without loss of generality that our form G takes the shape $G = G_1 \oplus G_2$ where G_1 is one of the forms given in Lemma 11. What remains is to bound the local densities of these forms. For this we have to make a distinction between the prime $p = 2$ and primes $p \geq 3$.

Lemma 12 ($p = 2$). *For all $k \geq 0$ we have $\delta_{B,2}(k) \leq 3$ and $\delta_{I_2,2}(k) \leq 2$.*

Proof. For I_2 and $k \geq 1$ Milnor [MH$^+$73, Lemma 9.1, p. 43] states that $\delta_{I_2,2}(k) \in \{0, 2\}$ and thus is bounded by 2. What remains is the case $k = 0$. For every solution to $x^2 + y^2 \equiv 0 \bmod 2^j$ for $j \geq 3$ one can verify that necessarily $x, y \equiv 0 \bmod 2^{\lfloor j/2 \rfloor}$. This leaves at most $(2^{j - \lfloor j/2 \rfloor})^2 \leq 2^{j+1}$ solutions and thus a density of at most $2^{j+1}/2^j = 2$. We conclude that $\delta_{I_2,2}(k) \leq 2$ for all $k \geq 0$.

We now consider B, i.e., the number of solutions to the equation $2x^2 + 2xy + 2y^2 \equiv k \bmod 2^j$ for $j \geq 3$. Clearly $\delta_{B,2}(k) = 0$ if $k \equiv 1 \bmod 2$. So we assume that $k = 2k'$ and normalize by 2 to obtain the equation $f(x, y) = x^2 + xy + y^2 = k' \bmod 2^{j-1}$. Note that if $2 \mid x, y$, then $4 \mid x^2 + xy + y^2$, and thus we need $k' \equiv 0 \bmod 4$. Else either $2 \nmid x$ or $2 \nmid y$, and we see modulo 2 that $k' \equiv 1 \bmod 2$. If $k' \equiv 2 \bmod 4$ there are thus no solutions. We now consider the case that

$k' \equiv 1 \bmod 2$. In this case the Jacobian $(2x+y, 2y+x) \equiv (y,x) \bmod 2$ of f at any solution (x,y) is nonzero modulo 2. So by a quantitive Hensel's Lemma every solution modulo 2 lifts to precisely 2^{j-1} solutions modulo 2^j, i.e., the density remains unchanged. One can simply count 3 solutions $(0,1),(1,0),(1,1)$ modulo 2 and thus we have a density of 1.5 for all $k' \equiv 1 \bmod 2$. Now we consider the case that $k' \equiv 0 \bmod 4$. Note that x, y are both divisible by 2 and thus we can divide both sides of the equation by 4. The density of the number of solutions is thus equal to that of $f(x,y) \equiv k'/4 \bmod p^{j-2}$. More generally we can divide by a power of 4 until we obtain a number equal to ± 1 or 2 modulo 4. By the previous result we thus obtain that the density is 1.5 if $v_2(k') \equiv 0 \bmod 2$, and 0 if $v_2(k') \equiv 1 \bmod 2$. Lastly, for $k' = 0$ and $j = 2j'$ we have the 2^j solutions $(a \cdot 2^{j'}, b \cdot 2^{j'}) \bmod 2^j$, and thus a density of 1. Note that we have only shown lower bounds for the densities, but one can quickly verify that we have accounted for all solutions as for $j = 2j'$ there are 2^{j-2i-1} elements in $1, \ldots, 2^j - 1$ with valuation $2i$, and thus we obtain a total density modulo 2^j of

$$1 + \sum_{i=0}^{j'-1} 2^{2j'-2i-1} \cdot \frac{3}{2} = 2^j.$$

To conclude we note that the density scales by a factor 2 after scaling back. □

Lemma 13 ($p \geq 3$). *Let p be an odd prime and consider the 6-dimensional form $D_u = uI_1 \oplus I_5$ with $\left(\frac{u}{p}\right) \in \{\pm 1\}$. Then for all $k \geq 0$ we have*

$$\delta_{D_u,p}(k) \leq \frac{1-p^{-3}}{1-p^{-2}}.$$

Additionally,

$$\prod_{p'=3,5,7,\ldots} \frac{1-p^{-3}}{1-p^{-2}} = \frac{6\zeta(2)}{7\zeta(3)} \leq 1.173.$$

Proof. Note that $\det(D_u) = u$, and let $\epsilon = \left(\frac{-u}{p}\right) \in \{\pm 1\}$. Let $k = p^l v \geq 1$ where p^l is the highest power of p dividing k. Then by [Sie35, Hilfssatz 16,p.544] we have for $j > l$ that

$$\delta_{D_u,p}(p^l v \bmod p^j) = (1 - \epsilon p^{-3})(1 + \epsilon p^{-2} + \epsilon^2 p^{-4} + \ldots + \epsilon^l p^{-2l}).$$

Note that the above local density is maximized if $\epsilon = 1$ and $l \to \infty$, i.e., we have

$$\delta_{D_u,p}(p^l v \bmod p^j) \leq (1 - p^{-3}) \sum_{i=0}^{\infty} p^{-2i} = \frac{1-p^{-3}}{1-p^{-2}},$$

and thus in particular $\delta_{D_u,p}(k) \leq \frac{1-p^{-3}}{1-p^{-2}}$. Furthermore, because the density only depends on the largest power p^l dividing $k > 0$ we also have $\delta_{D_u,p}(0) = \lim_{l\to\infty} \delta_{D_i,p}(p^l) = \frac{1-\epsilon p^{-3}}{1-\epsilon p^{-2}} \leq \frac{1-p^{-3}}{1-p^{-2}}$. Finally, we use the identity $\prod_{p=2,3,5,\ldots}(1 - p^{-i}) = 1/\zeta(i)$ for $i \geq 2$, and that $\frac{1-2^{-3}}{1-2^{-2}} = 7/6$. □

Corollary 4 (Finite places are bounded). *For any integral genus $\mathcal{G}$ of dimension $n \geq 6$ such that $\mathrm{rk}_p(\mathcal{G}) \geq 6$ for all primes p, we have for all $k \geq 0$*

$$\prod_{p=2,3,5,\ldots} \delta_{\mathcal{G},p}(k) \leq \frac{18\zeta(2)}{7\zeta(3)} < 3.52.$$

Proof. Let G be a form in $\mathcal{G}$ and consider a finite prime p. Because $\mathrm{rk}_p(\mathcal{G}) \geq 6$ we assume by Lemma 9 without loss of generality that G decomposes as $G = G_1 \oplus pG_2$ with $p \nmid \det(G_1)$ and $\dim(G_1) \geq 6$. For $p \neq 2$ an odd prime Lemma 11 shows that we can assume without loss of generality that $G_1 = D_u \oplus I_l$ where $D_u = uI_1 \oplus I_5$ for a unit $u \in \mathbb{Z}_p$. Lemma 13 gives that $\delta_{D_u,p}(k') \leq \frac{1-p^{-3}}{1-p^{-2}}$ for all $k' \geq 0$ and thus we can conclude that $\delta_{G,p}(k) \leq \frac{1-p^{-3}}{1-p^{-2}}$ by Lemma 10.

For $p = 2$ we either have $G_1 = I_2 \oplus G_1'$ or $G_1 = B \oplus G_1'$ and we can again conclude that $\delta_{G,2}(k) \leq 3$ by Lemmas 10 and 12. □

4.3 Proving the main result

Taking Theorem 8 and *Corollary* 4 together gives a bound on the coefficients $N_{\mathcal{G}}(k)$ of the expected theta series over a genus $\mathcal{G}$. This bound is sufficient to prove our main results.

Proof (Theorem 1). Let $\mathcal{G}$ and $0 < c \leq 1$ be as in the theorem statement and let $\Theta_{\mathcal{G}}(q) = 1 + \sum_{k=1}^{\infty} N_{\mathcal{G}}(k) q^k$ be the expected theta series of $\mathcal{G}$. By Theorem 8, Corollary 4 and Definition 10 we have for all $k \geq 1$ that

$$N_{\mathcal{G}}(k) = \prod_{p=2,3,\ldots,\infty} \delta_{\mathcal{G},p}(k) \leq \frac{9\zeta(2)}{7\zeta(3)} \cdot \frac{n\omega_n}{\mathrm{vol}(\mathcal{G})} \cdot k^{n/2-1}.$$

Let $\lambda = \left\lceil c^2 \cdot \left(\frac{7\zeta(3)}{9\zeta(2)} \omega_n^{-1} \mathrm{vol}(\mathcal{G})\right)^{2/n} \right\rceil$, then using the inequality $j^k \leq \int\limits_{j-\frac{1}{2}}^{j+\frac{1}{2}} t^k dt$ for $k \geq 1$ we get

$$\sum_{k=1}^{\lambda-1} N_{\mathcal{G}}(k) < \frac{9\zeta(2)}{7\zeta(3)} \cdot \frac{n\omega_n}{\mathrm{vol}(\mathcal{G})} \cdot \int\limits_0^{\lambda-\frac{1}{2}} t^{n/2-1} dt = \frac{18\zeta(2)}{7\zeta(3)} \cdot \frac{\omega_n}{\mathrm{vol}(\mathcal{G})} \cdot (\lambda - \tfrac{1}{2})^{n/2},$$

which by the choice of λ is bounded by $2c^n$. The result then follows from Lemma 2. □

Proof (Theorem 2). Let $\mathcal{G}$ be as in the theorem statement and let $\Theta_{\mathcal{G}}(q) = 1 + \sum_{k=1}^{\infty} N_{\mathcal{G}}(k) q^k$ be the expected theta series of $\mathcal{G}$. By Theorem 8, Corollary 4 and Definition 10 we have for all $k \geq 1$ that

$$N_{\mathcal{G}}(k) = \prod_{p=2,3,\ldots,\infty} \delta_{\mathcal{G},p}(k) \leq \frac{9\zeta(2)}{7\zeta(3)} \cdot \frac{n\omega_n}{\mathrm{vol}(\mathcal{G})} \cdot k^{n/2-1}.$$

Let $C = 26.1$, $\varepsilon \geq C \cdot (ce)^{-n} \cdot \mathrm{vol}(\mathcal{G})^{-1}$ and $s = \frac{1}{c} \cdot \left(\frac{C}{\varepsilon \cdot \mathrm{vol}(\mathcal{G})}\right)^{1/n} \leq e$. Then we have

$$\begin{aligned}
\Theta_{\mathcal{G}}(\exp(-\pi s^2)) - 1 &= \sum_{k=1}^{\infty} N_{\mathcal{G}}(k) \cdot \exp(-\pi s^2 k) \\
&\leq \sum_{k=1}^{\infty} \frac{9\zeta(2)}{7\zeta(3)} \cdot \frac{n\omega_n}{\mathrm{vol}(\mathcal{G})} \cdot k^{n/2-1} \cdot \exp(-\pi s^2 k) \\
&= \frac{9\zeta(2)}{7\zeta(3)} \cdot \frac{n\omega_n}{\mathrm{vol}(\mathcal{G})} \cdot \mathrm{Li}_{1-n/2}(\exp(-\pi s^2)) = (*).
\end{aligned}$$

Now by Lemma 7 and using that $\omega_n = \frac{\pi^{n/2}}{\Gamma(n/2+1)}$ we get

$$\begin{aligned}
(*) &\leq 7.39 \cdot \frac{9\zeta(2)}{7\zeta(3)} \cdot \frac{n}{\mathrm{vol}(\mathcal{G})} \cdot \frac{\pi^{n/2}}{\Gamma(n/2+1)} \cdot \Gamma(n/2) \cdot (\pi s^2)^{-n/2} \\
&= 7.39 \cdot \frac{9\zeta(2)}{7\zeta(3)} \cdot \mathrm{vol}(\mathcal{G})^{-1} \cdot 2 \cdot s^{-n} < C \cdot \mathrm{vol}(\mathcal{G})^{-1} \cdot s^{-n} = c^n \varepsilon.
\end{aligned}$$

So $\Theta_{\mathcal{G}}(\exp(-\pi s^2)) < 1 + c^n \varepsilon$ and we can conclude the proof by applying Lemma 4, where we recall that $\mathrm{vol}(\mathcal{L}^*) = \mathrm{vol}(\mathcal{G})^{-1}$. □

5 Applications

Instantiating [DvW22] without increasing the geometric gap. The LIP framework introduced in [DvW22] turns a decodable lattice $\mathcal{L}$ into a Key Encapsulation Scheme based on the hardness of distinguishing LIP between two auxiliary lattices $\mathcal{L}_1, \mathcal{L}_2$ in the same genus. The concrete hardness of this distinguishing problem is directly related to the geometry of these lattices, in particular, assuming the lattices are normalized to have determinant 1, the best known attacks seems to be driven by the *gap* $\max\{\mathrm{mk}(\mathcal{L})/\lambda_1(\mathcal{L}), \mathrm{mk}(\mathcal{L}^*)/\lambda_1(\mathcal{L}^*)\}$ between their first minimum (or their dual's) and the Minkowski bound. In the example instantiation given in [DvW22], an efficiently decodable lattice with primal and dual distance, and decoding gap bounded by f, leads to auxiliary lattices with geometric gaps bounded by $O(f^3)$, and thus this leads to a significant reduction in concrete security. The authors already hint that knowledge of good lattices in the same genus can decrease this blowup in the construction. Here we show how Theorem 1 could be used to only have a constant blowup, from f to $O(f)$.

Lemma 14. *Let $\mathcal{L}$ be any $n \geq 6$-dimensional lattice with primal, dual and efficient decoding gap bounded by $O(f)$ and such that $\mathrm{rk}_p(\mathcal{L}) \geq 6$ for all primes p. Then there exists a KEM which security reduces to a $2n$-dimensional instance of distinguish LIP with geometric gaps bounded by $O(f)$.*

Proof. Let $g \in \mathbb{Z}_{>0}$ be some scaling factor to be determined later. We need to construct an appropriate pair of lattices $\mathcal{L}_S, \mathcal{L}_Q$ to instantiate [DvW22, Theorem 5.2]. Let $\mathcal{L}' \in \mathrm{gen}(\mathcal{L})$ be a lattice with $\mathrm{gap}(\mathcal{L}') = O(1)$ as exists according to Theorem 1. We set $\mathcal{L}_S := g\mathcal{L} \oplus (g+1)\mathcal{L}$ as our well decodable lattice with decoding radius $\rho' = O(g/f \cdot \mathrm{gh}(\mathcal{L})$, and $\mathcal{L}_Q = \mathcal{L}' \oplus g(g+1)\mathcal{L}'$ as our auxilary lattice that has a dense sublattice $\mathcal{L}' \subset \mathcal{L}_Q$. Note that $\mathrm{gap}(\mathcal{L}_S) = O(f)$ and $\mathrm{gap}(\mathcal{L}_Q) = O(g)$. To satisfy the conditions of [DvW22, Theorem 5.2] we require that $\eta_{\frac{1}{2}}(\mathcal{L}') \leq \rho'/(2\sqrt{2n})$. Now note that because $\mathrm{gap}(\mathcal{L}') = O(1)$ we have

$$\eta_{\frac{1}{2}}(\mathcal{L}') \leq \eta_{2^{-n}}(\mathcal{L}') \leq \frac{\sqrt{n}}{\lambda_1((\mathcal{L}')^*)} = \theta(\det(\mathcal{L}')^{1/n}).$$

Furthermore, we have $\rho'/2\sqrt{2n} = \theta(g/f \cdot \mathrm{gh}(\mathcal{L})/\sqrt{n})$ and thus it is sufficient to pick g that satisfies

$$\theta(\det(\mathcal{L}')^{1/n}) \leq \theta(g/f \cdot \mathrm{gh}(\mathcal{L})/\sqrt{n}) = \theta(g/f \cdot \det(\mathcal{L}')^{1/n}),$$

from which it is clear that $g = \Theta(f)$ sufficies. We conclude by noting that then $\mathrm{gap}(\mathcal{L}_S) = O(f)$ and $\mathrm{gap}(\mathcal{L}_Q) = \Theta(g) = O(f)$. □

The encryption scheme from [BZI+24] based on the same framework benefits from the same improvement. For the signature scheme from [DvW22] we also improve the blowup from $O(f^2)$ to $O(f)$.

Lemma 15. *Let $\mathcal{L}$ be any $n \geq 6$-dimensional lattice with primal, dual and efficient Gaussian sampling gap bounded by $O(f)$ and such that $\mathrm{rk}_p(\mathcal{L}) \geq 6$ for all primes p. Then there exists a signature scheme which security reduces to a $2n$-dimensional instance of distinguish LIP with geometric gaps bounded by $O(f)$.*

Proof. The proof follows similarly as the construction in [DvW22] and the proof of Lemma 14, but with $\mathcal{L}_S := g\mathcal{L} \oplus (g+1)\mathcal{L}$ and $\mathcal{L}_{Q^{-1}} = \mathcal{L}' \oplus g(g+1)\mathcal{L}'$. □

Instantiating for $\mathbb{Z}^n$. Another interesting concurrent work [BGPSD23], that arrived shortly after [DvW22], introduces some cryptographic constructions based on LIP for $\mathbb{Z}^n$. In this work the authors raised some open questions about the instantiability of their schemes, as for this they require the existence of a lattice with certain geometric properties in the same genus as $\mathbb{Z}^n$. In particular, it is asked if there exist a lattice $\mathcal{L}$ in the genus of $\mathbb{Z}^n$ that have $\lambda_1^2(\mathcal{L}) \geq \Omega(n/\log n)$, or $\eta_\epsilon(\mathcal{L}) \leq \eta_\epsilon(\mathbb{Z}^n)/\sqrt{\log n} \approx O(\sqrt{\log(1/\epsilon)/\log n})$ for $\epsilon < n^{-\omega(1)}$. Note that the genus of $\mathbb{Z}^n$ is that of all odd unimodular lattices of dimension n. Therefore, the first question is in fact already answered by [MH+73], which shows that there exists an odd unimodular lattice with $\lambda_1(\mathcal{L})^2 \geq \Omega(n) > \Omega(n/\log n)$ for growing n. For the second, we can instantiate Theorem 2 with for example $\epsilon = 2^{-n}$, which implies the existence of $\mathcal{L}$ in the genus of $\mathbb{Z}^n$ with $\eta_\epsilon(\mathcal{L}) = O(1) < \eta_\epsilon(\mathbb{Z}^n)/\sqrt{\log(n)} \approx \Theta(\sqrt{n/\log(n)})$.

Instantiation of the Encryption scheme of [ARLW24]. In [ARLW24] a public-key encryption scheme based on LIP is presented. To instantiate the scheme however an auxiliary lattice is needed that satisfies some geometric properties. In the instantiation the authors would like to use unimodular lattices and they conjecture that for $n \geq 85$ there exists at least one unimodular lattice of dimension n such that $\lambda_1(\mathcal{L}) \geq \sqrt[4]{72n}$. We see that Lemma 5 is enough to answer this conjecture positively, and even show the existence of much better packings. More generally, given a decodable lattice $\mathcal{L}$ of dimension n they require another lattice $\mathcal{L}' \in \mathcal{G}$ with large minimum distance $\lambda_1(\mathcal{L}')$. For most genera Theorem 1 shows the existence of such a lattice.

6 Open questions

We discuss some open questions that could be interesting for further work.

In this work we have restricted ourselves for simplicity to unstructured lattices. However, in cryptography we often use structured module lattices to decrease storage and improve efficiency. The genus theory extends to these cases and mass formulas can also be extended in certain settings [Wei65, Kir16]. One has to be careful when considering the existing literature as it often considers the quadratic form $B^\top B$ for some module-lattice basis B. Except for the case of totally real number fields, the matrix $B^\top B$ does not correspond to the geometry of the module lattice via the canonical embedding. For example, in the common case of CM-fields, which is relevant for Hawk [DPPvW22], this information is instead captured by the Hermitian form $B^* B$. Luckily, most of the genus theory and the Smith-Siegel-Minkowski mass formula has been generalized to hermitian forms over CM-fields by [Kir16]. We therefore expect that Siegel's mass formula can also be generalized, leading to similar claims as in this work for module lattices over CM-fields.

Furthermore, the results of Siegel go further than just the expectation of the number of vectors of some squared norm. In fact, they can also count the expected number of higher rank constellations of multiple vectors with certain norm and inner product within each lattice. More precisely, for forms $G \in \mathcal{S}_n^{>0}(\mathbb{Z})$ and $K \in \mathcal{S}_m^{>0}(\mathbb{Z})$ for $1 \leq m \leq n$, it counts the expected number of solutions $X \in \mathbb{Z}^{n \times m}$ such that $X^\top G X = K$ when varying G over a genus. This might open up the possibility to study more advanced geometric properties of random lattices over a genus.

References

AEN19. Yoshinori Aono, Thomas Espitau, and Phong Q. Nguyen. Random lattices: Theory and practice, 2019. Available at `https://espitau.github.io/bin/random_lattice.pdf`.

ARLW24. Léo Ackermann, Adeline Roux-Langlois, and Alexandre Wallet. Public-key encryption from the lattice isomorphism problem. *WCC 2024: The Thirteenth International Workshop on Coding and Cryptography*, 2024.

Extended abstract available at `https://wcc2024.sites.dmi.unipg.it/WCC_proceedings.pdf`. Full version to appear.

BDG23. Peter Bruin, Léo Ducas, and Shane Gibbons. Genus distribution of random q-ary lattices. *Banach Center Publications*, 126:137–159, December 2023.

BGPSD23. Huck Bennett, Atul Ganju, Pura Peetathawatchai, and Noah Stephens-Davidowitz. Just how hard are rotations of $\mathbb{Z}^n$? algorithms and cryptography with the simplest lattice. In *EUROCRYPT 2023: Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 252–281. Springer, 2023.

BZI$^+$24. Gustavo C Biage, Gustavo Zambonin, Thaís B Idalino, Daniel Panario, and Ricardo Custodio. A concrete LIP-based KEM with simple lattices. *IEEE Access*, 2024.

Che21. Gaëtan Chenevier. Statistics for kneser p-neighbors. *arXiv preprint arXiv:2104.06846*, 2021.

CS88. John Horton Conway and Neil JA Sloane. Low-dimensional lattices. IV. The mass formula. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, 419(1857):259–286, 1988.

CS99. J. H. Conway and N. J. A. Sloane. *Sphere Packings, Lattices and Groups.* Springer New York, 1999.

DADRT23. Thomas Debris-Alazard, Léo Ducas, Nicolas Resch, and Jean-Pierre Tillich. Smoothing codes and lattices: Systematic study and new bounds. *IEEE Transactions on Information Theory*, 69(9):6006–6027, 2023.

DG23. Léo Ducas and Shane Gibbons. Hull attacks on the lattice isomorphism problem. In *PKC 2023: IACR International Conference on Public-Key Cryptography*, pages 177–204. Springer, 2023.

DPPvW22. Léo Ducas, Eamonn W Postlethwaite, Ludo N Pulles, and Wessel van Woerden. Hawk: Module LIP makes lattice signatures fast, compact and simple. In *ASIACRYPT 2022: International Conference on the Theory and Application of Cryptology and Information Security*, pages 65–94. Springer, 2022.

DSHVvW20. Mathieu Dutour Sikirić, Anna Haensch, John Voight, and Wessel van Woerden. A canonical form for positive definite matrices. *ANTS XIV*, 4(1):179–195, 2020.

DvW22. Léo Ducas and Wessel van Woerden. On the lattice isomorphism problem, quadratic forms, remarkable lattices, and cryptography. In *EUROCRYPT 2022: Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 643–673. Springer, 2022.

Fei96. W. Feit. Orders of finite linear groups. In *Proceedings of the First Jamaican Conference on Group Theory and its Applications*, pages 9–11. Univ. West Indies, Kingston, 1996.

Han04. Jonathan Hanke. Local densities and explicit bounds for representability by a quadratic form. *Duke Mathematical Journal*, 124(2):351 – 388, 2004.

Hei16. Jeffery P. Hein. *Orthogonal modular forms: an application to a conjecture of Birch, algorithms and computations.* PhD thesis, Dartmouth College Library Press, 2016.

Hla43. Edmund Hlawka. Zur geometrie der zahlen. *Mathematische Zeitschrift*, 49(1):285–312, 1943.

HR14. Ishay Haviv and Oded Regev. On the lattice isomorphism problem. In *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*, pages 391–404. SIAM, 2014.

Kir16. Markus Kirschmer. Definite quadratic and hermitian forms with small class number. *Habilitation, RWTH Aachen University*, 2016. available at `https://www.math.rwth-aachen.de/~Markus.Kirschmer/papers/herm.pdf`.

MH$^+$73. John Willard Milnor, Dale Husemoller, et al. *Symmetric bilinear forms*, volume 5. Springer, 1973.

Min85. Hermann Minkowski. *Untersuchungen über quadratische Formen: Bestimmung der Anzahl verschiedener Formen, welche ein gegebenes Genus enthält.* E. Erlatis, 1885.

Min10. Hermann Minkowski. *Geometrie der Zahlen.* B.G. Teubner, 1910.

PP85. Wilhelm Plesken and Michael Pohst. Constructing integral lattices with prescribed minimum. I. *Mathematics of computation*, 45(171):209–221, 1985.

PS97. Wilhelm Plesken and Bernd Souvignier. Computing isometries of lattices. *Journal of Symbolic Computation*, 24(3-4):327–334, 1997.

RSD24. Oded Regev and Noah Stephens-Davidowitz. A reverse Minkowski theorem. *Annals of Mathematics*, 199(1):1–49, 2024.

Ser73. Jean-Pierre Serre. *A Course in Arithmetic.* Springer New York, 1973.

Sie35. Carl Ludwig Siegel. Uber die analytische theorie der quadratischen formen. *Annals of Mathematics*, pages 527–606, 1935.

Sie45. Carl Ludwig Siegel. A mean value theorem in geometry of numbers. *Annals of Mathematics*, pages 340–347, 1945.

The24. The Sage Developers. *SageMath, the Sage Mathematics Software System (Version 10.4)*, 2024. `https://www.sagemath.org`.

Wat60. G.L. Watson. *Integral Quadratic Forms, By G.L. Watson.* Cambridge Tracts in Mathematics and Mathematical Physics, No. 51. 1960.

Wei65. André Weil. Sur la formule de siegel dans la théorie des groupes classiques. *Acta math*, 113(1-87):2, 1965.

Wei84. Boris Weisfeiler. On the size and structure of finite linear groups. available at `https://arxiv.org/pdf/1203.1960`, 1984.

Woo92. David C. Wood. The computation of polylogarithms. Technical report, University of Kent, Computing Laboratory, University of Kent, Canterbury, UK, 1992.