

Average-Case Hardness of Lattice Problems in Special Genera of Fixed Rank

Koen de Boer¹, Maxime Bombar², Aurel Page³, and Wessel van Woerden⁴

¹ Independent Researcher, the Netherlands

² Univ. Bordeaux, CNRS, Bordeaux INP, IMB, UMR 5251, F-33400, Talence, France

³ Inria, Univ. Bordeaux, CNRS, Bordeaux INP, IMB, UMR 5251, F-33400, Talence, France

⁴ PQShield, the Netherlands

Abstract. The distinguishing variant Δ LIP of the Lattice Isomorphism Problem underlies the security proofs of a growing family of lattice-based schemes. As the special genus is the strongest efficiently computable invariant of a lattice, such a proof may replace the lattice used by the scheme with a random lattice of the same genus. This raises the question of how hard lattice problems actually are on such random lattices.

We answer it, under GRH, with a worst-case to average-case reduction inside any special genus $\mathcal{G}$ of Hermitian lattices of fixed rank $r \geq 2$ over a cyclotomic field K of degree d : for each of (H)SVP, HCVP, (Δ) BDD, DSP and DGS, solving the problem on a random class of $\mathcal{G}$, drawn from its natural distribution, is as hard as solving it on the *worst-case* lattice of $\mathcal{G}$, at the cost of a factor $p = \text{poly}(d) \cdot \det_{\mathbb{Q}}(\Lambda)^{O_r(1)/d}$ in the approximation factor. The reduction is a random walk on the Kneser $\mathfrak{p}$ -neighbour graph of $\mathcal{G}$: neighbours stay close enough to transfer problem instances, while the walk equidistributes quickly, a fact we deduce from spectral bounds for the associated Hecke operators.

As an application we give a direct and tight security proof for a minimal LIP-based KEM, resting on two independent assumptions: that the lattice it uses is indistinguishable from a random lattice of its genus, and that Δ BDD is hard on that genus. Unlike existing LIP-based proofs, this requires neither an ad-hoc lattice of the form $q\Lambda \oplus (q+1)\Lambda$, nor the tightness loss caused by the dense sublattice such proofs rely on.

1 Introduction

The Lattice Isomorphism Problem (LIP) [PS97, HR14, LJS19] asks, given two isometric lattices, to recover an isometry mapping one onto the other. It has emerged as an alternative foundation for lattice-based cryptography [DvW22, BG-PSD23, DPPvW22], and has since driven work on search-to-distinguish reductions [Szy03, vGvW25, BW25], on the genus and other invariants [BDG23, LLM24, vW24, vG25, Mur25], and an active cryptanalytic effort on module-LIP and its instantiation in Hawk [DG23, MPPW24, LJPW24, APMvW25, vGP25, CME⁺25, CM25, SW26, MPM26]. Alongside it one considers its distinguishing variant Δ LIP: given a sample $\Lambda \leftarrow [\Lambda_b]$ from one of two distinct isomorphism classes $[\Lambda_1], [\Lambda_2]$,

recover $b \in \{1, 2\}$. Just as for LWE, the search problem often appears as the underlying key-recovery or decryption problem⁵, while the distinguishing variant of the problem allows for GPV [GPV08] and lossy-trapdoor-like [PW08] security proofs. Indeed, in this way the distinguishing variant allows for basic encryption and signature schemes [DvW22, BGPSD23, BZI⁺24, ARLW24], but also for advanced cryptographic constructions like identity-based encryption [BLSD26] and (fully) homomorphic encryption [BMM25, LRvW25, BLSD26].

Invariants are the natural way to attack Δ LIP. The determinant is the simplest example: whenever $\det([A_1]) \neq \det([A_2])$, computing $\det(A)$ already solves the problem. The strongest efficiently computable invariant is the (spinor) *genus* $\text{gen}(A)$ of a lattice: two integral lattices lie in the same genus when they are equivalent over the p -adic integers $\mathbb{Z}_p$ for every prime p . The notion extends to module lattices over a number field K , replacing $\mathbb{Z}_p$ by the $\mathfrak{p}$ -adic integers $\mathcal{O}_{K,\mathfrak{p}}$ for every prime ideal $\mathfrak{p}$ of $\mathcal{O}_K$. For a CM field it sharpens further to the *special genus* $\text{sgen}(A)$ of a Hermitian lattice A , which was investigated in the context of Hawk in [Mur25]. Since no stronger efficient invariant is known, Δ LIP over a CM field is typically assumed to be hard when the two classes it asks to distinguish lie in the same special genus.

Such a Δ LIP assumption buys freedom in a security proof: the lattice A_1 used by the scheme may be replaced by any lattice $A_2 \in \text{sgen}(A_1)$. That freedom is finite, as a (special) genus contains only a *finite* number of isomorphism classes $[A_1], \dots, [A_N]$, but N can grow rather large, and the genus comes with a natural distribution $\mathcal{H}$ over them. One idea is thus to replace A_1 by such a *random lattice* $A_2 \leftarrow \mathcal{H}(\text{sgen}(A_1))$.

For $K = \mathbb{Q}$ it was shown in [vW24] that lattices drawn from this distribution behave like truly random ones: their minima, covering radius and smoothing parameter all match those of a uniformly random lattice. In particular, they provably follow the Gaussian Heuristic. Geometrically, one therefore expects such a random lattice to be among the hardest in its genus: for example that finding a short vector in it is at least as hard as in any other, worst-case, lattice of the same genus. It is precisely this that we would like to assume once the replacement in a security proof has been made. Since the works of Ajtai [Ajt96] and Regev [Reg09], the standard way to justify an assumption of this shape is a worst-case to average-case reduction, which raises the question: can the intuition above be formalized into such a reduction within a genus?

Contribution 1: a worst-case to average-case reduction in special genera of fixed rank We answer this question affirmatively, for any special genus of *fixed rank* $r \geq 2$ and a large variety of lattice problems. This result rests on sharp spectral bounds for the Hecke operators attached to Kneser $\mathfrak{p}$ -neighbours, which we obtain from deep results in the representation theory of p -adic groups when $r \geq 3$, and in the theory of automorphic forms when $r = 2$.

⁵ The search LIP variant also led to interesting constructions based on cryptographic group actions, but these are outside the scope of this work.

Theorem 1.1 (Informal). *Let the rank $r \geq 2$ be fixed, let K be a cyclotomic number field of degree d , and assume the generalized Riemann hypothesis (GRH). For the lattice problems $\mathcal{P} \in \{(\text{H})\text{SVP}, \text{HCVP}, (\Delta)\text{BDD}, \text{DSP}, \text{DGS}\}$ (see Section 4), solving $\mathcal{P}$ on a rank r Hermitian lattice Λ efficiently reduces to solving $\mathcal{P}$ on a random $\Lambda' \leftarrow \mathcal{H}(\mathcal{G})$, where $\mathcal{G} = \mathcal{G}(\Lambda)$ is the special genus of Λ and $\mathcal{H}(\mathcal{G})$ the Haar-uniform distribution on it. The reduction loses a factor $p = \text{poly}(d) \cdot \det_{\mathbb{Q}}(\Lambda)^{O_r(1)/d}$ in the approximation factor of $\mathcal{P}$.*

The reduction is polynomial-time in the degree d of K and the normalized rational determinant of the lattice, which also accounts for the discriminant of the field. The loss is squared for those problems whose approximation factor scales with the lattice itself, for example for SVP and BDD; Section 4 treats both cases. We state the result for cyclotomic K out of technical convenience. This covers most module lattices used in lattice cryptography, and in particular it covers the rank 2 and power-of-two cyclotomic case of Hawk [DPPvW22]. The machinery we develop applies to any CM field, but there is a small caveat regarding the existence and efficient computability of a suitable rational prime p (see Remark 3.7), which is easily resolved for most CM fields. GRH plays the same limited role: we assume it only to find such a p efficiently (Lemma 3.6), the reduction itself being unconditional.

Contribution 2: a direct and tight security proof for LIP-based schemes

Our second contribution shows what such a reduction buys on the cryptographic side. The security proofs of current LIP-based schemes [DvW22, BGPSD23, BZI⁺24, ARLW24, BMM25, BLS26] all reduce to a *single* ΔLIP instance. The trick is that the lattice $\Lambda_1 = q\Lambda \oplus (q+1)\Lambda$ used by the scheme is replaced, in the proof, by the auxiliary lattice $\Lambda_2 = \Lambda \oplus q(q+1)\Lambda$ of the same genus, which contains the dense sublattice Λ ; the existence of that dense sublattice then gives information-theoretic security. While this leads to a reduction to a single LIP problem, it comes with two significant drawbacks:

1. **Ad-hoc.** The scheme has to use the lattice $\Lambda_1 = q\Lambda \oplus (q+1)\Lambda$, rather than a lattice chosen for efficiency or functionality;
2. **Tightness loss.** The auxiliary lattice Λ_2 contains a dense sublattice and is therefore *easier* to distinguish than the lattice Λ_1 actually used, so the assumption one has to make is stronger than the one the scheme appears to rely on.

The tightness loss can be improved upon using [vW24], but the ad-hoc instantiation remains. Following a suggestion made informally in [LRvW25, Appendix A.2], we propose instead to split the security of a LIP-based scheme into two *independent* assumptions:

- (i) the lattice used by the scheme is hard to distinguish from a *random* lattice of its (special) genus, and
- (ii) the underlying lattice problem is hard on a random lattice of that genus.

Assumption (i) matches the state of the art on invariants, namely that the (spinor/special) genus is the strongest efficiently computable one, and it constrains the scheme only through the genus it lives in. Assumption (ii) is exactly what Contribution 1 turns into worst-case hardness. In Section 5 we make this formal for a minimal key encapsulation mechanism Π_0 .

Theorem 1.2 (Informal, see Theorem 5.3 and Corollary 5.5). *The KEM Π_0 built on a Hermitian lattice $\Lambda_0 \in \mathcal{G}$ with an efficient decoder of radius ρ is IND-CPA secure, assuming (i) that a random class of $\mathcal{G}$ is indistinguishable from $[\Lambda_0]$, and (ii) that ΔBDD is hard in the worst case on every lattice of $\mathcal{G}$, at a decoding distance within a $\text{poly}(d)$ factor of ρ .*

There is no dense sublattice left, and no unexplained gap between the instance the scheme uses and the instance assumed to be hard: Π_0 may be built on any lattice admitting an efficient decoder, and the proof loses nothing beyond the two assumptions themselves. The reduction is moreover tight with respect to the attacks we know: the natural way to break such a scheme is either to solve LIP through ΔLIP , which is (i), or to solve the underlying decoding problem in a lattice that cannot be distinguished from a random one, which is (ii).

Technical details The main technical tool of this work is the use of $\mathfrak{p}$ -neighbours (Section 3), first introduced by Kneser [Kne57]. This notion turns the finite set of isomorphism classes $[\Lambda_1], \dots, [\Lambda_N]$ inside a genus into the vertex set of a *connected graph*, where two classes share an edge if they are $\mathfrak{p}$ -neighbours. This is useful because of three properties:

- (1) one can efficiently construct a random $\mathfrak{p}$ -neighbour of a given lattice;
- (2) neighbour lattices are “sufficiently similar” to allow the transfer of solutions to algorithmic problems from one to the other;
- (3) the random walk given by $\mathfrak{p}$ -neighbours has good equidistribution properties.

Property (1) is classical and has had algorithmic applications for a long time: see [Sch98, KV10, GV14, AC25] for some examples; see also Algorithm 1. Property (2), although natural, requires a careful choice of the neighbouring relation. Indeed, the neighbours are not sublattices of the starting lattice; instead they contain a common sublattice, possibly with large index. We iterate neighbours over all prime ideals above a single rational prime p to arrive at a lattice Λ' : this implies the crucial containment $p\Lambda' \subset \Lambda$ while allowing us to keep p small. It also has the additional benefit of preserving the special genus. Finally, Property (3) rests on the interpretation of $\mathfrak{p}$ -neighbours in terms of the Hecke operators appearing in the theory of automorphic forms. We therefore benefit from a rich theory, developed in part to answer Ramanujan’s conjecture [Ram16] and its generalizations [BB13]. In our setting, we derive the eigenvalue bounds from Oh’s local bound [Oh02] when $r \geq 3$, and when $r = 2$ from an exceptional isomorphism, the Jacquet–Langlands correspondence [JL70] and Livné’s generalization [Liv01] of Deligne’s bound [Del73] to Hilbert modular forms. The proof of the equidistribution result is given in Section B.

Related works The idea of performing a random walk on a graph or continuous space, combined with fast mixing and the transfer of a problem instance along the walk, has been fruitful and forms the foundation of many worst-case to average-case results in cryptography. The closest related work to our setting might not be found in lattice cryptography, but instead in the world of isogeny-based cryptography. Indeed, Jao et al. [JMV09] and Jetchev and Wesolowski [JW15] showed self-reducibility results where the objects were elliptic curves or abelian surfaces, and where two objects share an edge if they are related by an ℓ -isogeny. In fact, for ordinary objects the class group of an order in a CM field acts on the vertices of this graph through complex multiplication, closely paralleling the way the isomorphism classes of a special genus of Hermitian lattices over a CM field form the vertices of the Kneser $\mathfrak{p}$ -neighbour graph. In the supersingular case the analogous dictionary is the Deuring correspondence, under which isogenies translate into quaternion ideals, i.e. into lattices again.

On the lattice side, however, the closest in spirit to our work are the worst-case to average-case reductions for SVP on ideal lattices [dBDPMW20] and, recently, for SIVP on module lattices of *fixed rank* [dBPTW26]. Because the space of ideal lattices or module lattices is *continuous*, these reductions were, however, rather difficult to formalize; one must consider not only discrete steps (passing to renormalized index- p sublattices), but also continuous ones. Moreover, the space of fixed rank $r \geq 2$ module lattices is far from compact, which means that some lattices lie far away from the average-case lattices in the space.

Both of these difficulties completely disappear for the special genera of fixed rank we consider: we have a finite and discrete space with discrete steps, and moreover it comes with natural bounds on how far from average a lattice can be; e.g., integrality already bounds all successive minima from below by 1. While we thus rely on deep mathematical results for the mixing properties, the remainder of the reduction in this work is surprisingly elementary, and allows for rather clean proofs and a clear presentation.

2 Preliminaries

2.1 Notation

We use standard algebraic number theory throughout; see Section A for a self-contained recollection of the facts below, and [NS13] for a standard textbook on this material. We now set the notation used throughout this work.

Let K be a number field of degree $d = [K : \mathbb{Q}]$, with ring of integers $\mathcal{O}_K$ and d complex embeddings $\sigma : K \rightarrow \mathbb{C}$. For $a \in K$ we write $\text{Tr}(a) = \sum_{\sigma} \sigma(a)$ and $N(a) = \prod_{\sigma} \sigma(a)$, both of which lie in $\mathbb{Q}$. We say that K is a *CM field* with maximal totally real subfield F , and write K/F for the corresponding *CM extension*, when K is totally complex and $K = F(\sqrt{a})$ for some totally negative $a \in F$. Complex conjugation then induces a non-trivial automorphism of K fixing F , which we denote by $x \mapsto \bar{x}$; it satisfies $\sigma(\bar{x}) = \overline{\sigma(x)}$ for every embedding σ . Our running example is the cyclotomic field $K = \mathbb{Q}(\zeta_m)$, which is CM with $F = \mathbb{Q}(\zeta_m + \zeta_m^{-1})$.

A *fractional ideal* $\mathfrak{a}$ of K is a nonzero $\mathcal{O}_K$ -submodule of K with $z\mathfrak{a} \subseteq \mathcal{O}_K$ for some $z \in K^\times$; it is *integral* when $\mathfrak{a} \subseteq \mathcal{O}_K$, in which case its norm is $N(\mathfrak{a}) \stackrel{\text{def}}{=} |\mathcal{O}_K/\mathfrak{a}|$. We write $\bar{\mathfrak{a}} \stackrel{\text{def}}{=} \{\bar{x} \mid x \in \mathfrak{a}\}$ for its complex conjugate. As $\mathcal{O}_K$ is a Dedekind domain, every integral ideal factors uniquely into prime ideals, and the nonzero primes are exactly the maximal ideals, so that $\mathcal{O}_K/\mathfrak{p}$ is a finite field for every prime $\mathfrak{p}$. A prime ideal $\mathfrak{q}$ of F is *split* in K when $\mathfrak{q}\mathcal{O}_K = \mathfrak{p}\bar{\mathfrak{p}}$ for a prime $\mathfrak{p} \neq \bar{\mathfrak{p}}$ of $\mathcal{O}_K$, in which case $\mathcal{O}_K/\mathfrak{p} \simeq \mathcal{O}_K/\bar{\mathfrak{p}} \simeq \mathcal{O}_F/\mathfrak{q} \simeq \mathbb{F}_q$; it is *inert* when $\mathfrak{q}\mathcal{O}_K$ remains prime, and *ramified* when $\mathfrak{q}\mathcal{O}_K = \mathfrak{p}^2$. Only finitely many primes of F ramify.

Finally, we write $F_{\mathfrak{q}}$ for the completion of F at a prime $\mathfrak{q}$, with valuation ring $\mathcal{O}_{F,\mathfrak{q}}$, and $K_{\mathfrak{q}} \stackrel{\text{def}}{=} K \otimes_F F_{\mathfrak{q}}$, to which complex conjugation extends by $k \otimes f \mapsto \bar{k} \otimes f$. When $\mathfrak{q} = \mathfrak{p}\bar{\mathfrak{p}}$ splits in K there is a ring isomorphism

$$K_{\mathfrak{q}} \simeq F_{\mathfrak{q}} \times F_{\mathfrak{q}}, \quad k \otimes f \mapsto (kf, \bar{k}f),$$

under which complex conjugation acts by exchanging the two factors.

2.2 Hermitian Spaces

Let K/F be a CM extension, and let V be a K -vector space of dimension r . By choosing a K -basis of V , we will silently identify an element of V with the (column) vector in K^r of its K -coordinates. We say that V is a Hermitian space over K when it is equipped with a nondegenerate map $\langle \cdot, \cdot \rangle$ such that for any $\mathbf{x}, \mathbf{x}', \mathbf{y}, \mathbf{y}' \in V$, and any $a, b \in K$, we have

$$\begin{aligned} \langle a\mathbf{x} + b\mathbf{x}', \mathbf{y} \rangle &= \bar{a}\langle \mathbf{x}, \mathbf{y} \rangle + \bar{b}\langle \mathbf{x}', \mathbf{y} \rangle, \\ \langle \mathbf{x}, a\mathbf{y} + b\mathbf{y}' \rangle &= a\langle \mathbf{x}, \mathbf{y} \rangle + b\langle \mathbf{x}, \mathbf{y}' \rangle, \\ \langle \mathbf{y}, \mathbf{x} \rangle &= \overline{\langle \mathbf{x}, \mathbf{y} \rangle}. \end{aligned}$$

In particular, for every $\mathbf{v} \in V$, we have $\langle \mathbf{v}, \mathbf{v} \rangle = \overline{\langle \mathbf{v}, \mathbf{v} \rangle}$, i.e. $\langle \mathbf{v}, \mathbf{v} \rangle \in F$. The form is said to be totally positive definite if for every embedding $\sigma: F \rightarrow \mathbb{R}$, and every $v \neq 0$ we have $\sigma(\langle \mathbf{v}, \mathbf{v} \rangle) > 0$. A Hermitian space is called *isotropic* if it contains a nonzero element $x \in V$ such that $\langle x, x \rangle = 0$.

If $b_1, \dots, b_r$ is a K -basis of V , its *Gram matrix* is $G \stackrel{\text{def}}{=} (\langle b_i, b_j \rangle)_{1 \leq i, j \leq r}$. This matrix is Hermitian, i.e. it is equal to its transpose-conjugate $G^* = \bar{G}^T$.

Example 2.1. Let $K \stackrel{\text{def}}{=} \mathbb{Q}(\zeta)$ be a cyclotomic field, which is CM when regarded as a quadratic extension of its totally real subfield $F \stackrel{\text{def}}{=} \mathbb{Q}(\zeta + \zeta^{-1})$. The standard Hermitian form on K^r is defined as

$$\langle \mathbf{v}, \mathbf{w} \rangle \stackrel{\text{def}}{=} \sum_{i=1}^r \bar{v}_i w_i.$$

Hermitian forms allow one to define isometries of V .

Definition 2.2 (Unitary Groups). Let $(V, \langle \cdot, \cdot \rangle)$ be a Hermitian space. We define the unitary group

$$U(V) \stackrel{\text{def}}{=} \{\sigma \in \mathbf{GL}_K(V) \mid \langle \sigma(\mathbf{x}), \sigma(\mathbf{y}) \rangle = \langle \mathbf{x}, \mathbf{y} \rangle, \text{ for all } \mathbf{x}, \mathbf{y} \in V\}.$$

Elements of $U(V)$ are called global isometries of $(V, \langle \cdot, \cdot \rangle)$. The subgroup of determinant 1 isometries is called the special unitary group of V and is denoted by

$$SU(V) \stackrel{\text{def}}{=} \{\sigma \in U(V) \mid \det \sigma = 1\}.$$

Local Isometries In this work, we will frequently need to look at what happens locally above a specific finite prime of F . More precisely, let $\mathfrak{q}$ be a prime ideal of F , and let $F_{\mathfrak{q}}$ be the completion of F at $\mathfrak{q}$. We define the $K_{\mathfrak{q}}$ -module $V_{\mathfrak{q}} \stackrel{\text{def}}{=} V \otimes_F F_{\mathfrak{q}}$. The Hermitian form $\langle \cdot, \cdot \rangle$ extends naturally to $V_{\mathfrak{q}}$, and we let

$$U(V_{\mathfrak{q}}) \stackrel{\text{def}}{=} \{\sigma \in \mathbf{GL}_{K_{\mathfrak{q}}}(V_{\mathfrak{q}}) \mid \langle \sigma(\mathbf{x}), \sigma(\mathbf{y}) \rangle = \langle \mathbf{x}, \mathbf{y} \rangle, \text{ for all } \mathbf{x}, \mathbf{y} \in V_{\mathfrak{q}}\}$$

denote the group of local isometries at $\mathfrak{q}$. Similarly, $SU(V_{\mathfrak{q}})$ will denote the subgroup of determinant one local isometries at $\mathfrak{q}$. The structure of these groups depends on the splitting behaviour of $\mathfrak{q}$.

Assume that $\mathfrak{q}$ splits in K as $\mathfrak{q}\mathcal{O}_K = \mathfrak{p}\bar{\mathfrak{p}}$. Under the ring isomorphism $K_{\mathfrak{q}} \simeq F_{\mathfrak{q}} \times F_{\bar{\mathfrak{q}}}$, let $e_{\mathfrak{p}} \stackrel{\text{def}}{=} (1, 0)$ and $e_{\bar{\mathfrak{p}}} \stackrel{\text{def}}{=} (0, 1)$. They are idempotents of $K_{\mathfrak{q}}$ such that $e_{\mathfrak{p}} + e_{\bar{\mathfrak{p}}} = 1$ and $e_{\mathfrak{p}}e_{\bar{\mathfrak{p}}} = 0$. Moreover, complex conjugation exchanges them. Define $W_{\mathfrak{p}} \stackrel{\text{def}}{=} e_{\mathfrak{p}}V_{\mathfrak{q}}$ and $W_{\bar{\mathfrak{p}}} \stackrel{\text{def}}{=} e_{\bar{\mathfrak{p}}}V_{\mathfrak{q}}$. They are both r -dimensional $F_{\mathfrak{q}}$ -vector spaces and we have the direct sum decomposition $V_{\mathfrak{q}} = W_{\mathfrak{p}} \oplus W_{\bar{\mathfrak{p}}}$. The choice of a particular basis of $V_{\mathfrak{q}}$ allows us to obtain the matrix representation specified in the following lemma:

Lemma 2.3 (Local isometries at a split prime). Let K/F be a CM extension and let $(V, \langle \cdot, \cdot \rangle)$ be a Hermitian space of dimension r over K . Let $\mathfrak{q}$ be a nonzero prime ideal of F which splits in K . Then, the local isometries can be identified with invertible matrices over the local field $F_{\mathfrak{q}}$

$$\begin{aligned} U(V_{\mathfrak{q}}) &\simeq \mathbf{GL}_r(F_{\mathfrak{q}}) \\ SU(V_{\mathfrak{q}}) &\simeq \mathbf{SL}_r(F_{\mathfrak{q}}) \end{aligned}$$

Proof. Decompose $V_{\mathfrak{q}} \simeq W_{\mathfrak{p}} \oplus W_{\bar{\mathfrak{p}}}$. A simple computation shows that $\langle W_{\mathfrak{p}}, W_{\mathfrak{p}} \rangle = \langle W_{\bar{\mathfrak{p}}}, W_{\bar{\mathfrak{p}}} \rangle = 0$, and that $\langle W_{\mathfrak{p}}, W_{\bar{\mathfrak{p}}} \rangle \subset \{0\} \times F_{\mathfrak{q}} \simeq F_{\mathfrak{q}}$. Let $(e_1, \dots, e_r)$ be an $F_{\mathfrak{q}}$ -basis of $W_{\mathfrak{p}}$. Since $\langle \cdot, \cdot \rangle$ is nondegenerate, there exists a unique corresponding basis $(f_1, \dots, f_r)$ of $W_{\bar{\mathfrak{p}}}$ such that $\langle e_i, f_j \rangle = (0, \delta_{i,j})$. Now, for $\mathbf{x}, \mathbf{y} \in V_{\mathfrak{q}}$ written in this adapted basis as $\mathbf{x} \stackrel{\text{def}}{=} (x_1, x_2)$ and $\mathbf{y} = (y_1, y_2)$, then $\langle \mathbf{x}, \mathbf{y} \rangle \in K_{\mathfrak{q}}$ identifies with $(x_2^{\top} y_1, x_1^{\top} y_2) \in F_{\mathfrak{q}} \times F_{\mathfrak{q}}$. Now, a $K_{\mathfrak{q}}$ linear isometry $\sigma \in U(V_{\mathfrak{q}})$ has the form $\sigma = (A, B)$ where A (resp. B) is an $F_{\mathfrak{q}}$ -linear transformation acting on the first factor $W_{\mathfrak{p}}$ (resp. the second factor $W_{\bar{\mathfrak{p}}}$). In particular, for $\mathbf{x} = (x_1, x_2) \in V_{\mathfrak{q}}$ and $\mathbf{y} = (y_1, y_2) \in V_{\mathfrak{q}}$, we must have

$$\langle \sigma(\mathbf{x}), \sigma(\mathbf{y}) \rangle = (x_2^{\top} B^{\top} A y_1, x_1^{\top} A^{\top} B y_2).$$

Since σ is an isometry of $(V_{\mathfrak{q}}, \langle \cdot, \cdot \rangle)$, this yields the relation $B = (A^{-1})^\top$. In particular

$$U(V_{\mathfrak{q}}) = \{(A, (A^{-1})^\top) \mid A \in \mathbf{GL}_r(F_{\mathfrak{q}})\} \simeq \mathbf{GL}_r(F_{\mathfrak{q}}).$$

Furthermore, $\sigma = (A, (A^{-1})^\top) \in SU(V_{\mathfrak{q}})$ if and only if $\det A = 1$, *i.e.*

$$SU(V_{\mathfrak{q}}) = \{(A, (A^{-1})^\top) \mid A \in \mathbf{SL}_r(F_{\mathfrak{q}})\} \simeq \mathbf{SL}_r(F_{\mathfrak{q}}).$$

□

Remark 2.4. The above lemma shows that when $\mathfrak{q}$ is a split prime of F , the Hermitian space $V_{\mathfrak{q}}$ is isotropic.

2.3 Hermitian Lattices

Let K/F be a CM extension, and let $(V, \langle \cdot, \cdot \rangle)$ be a Hermitian space. We recall some well-known facts about lattices in V , and we refer the reader to [Coh99] for more details. A *Hermitian lattice* $\Lambda \subset V$ is a finitely generated, torsion-free $\mathcal{O}_K$ -module which spans V over K . Any such module is always of the form $\Lambda = \mathfrak{a}_1 b_1 + \cdots + \mathfrak{a}_r b_r$ where $(b_1, \dots, b_r)$ is a K -basis of V and the $\mathfrak{a}_i$'s are fractional ideals in K . Let B be the column matrix of the b_i 's. The collection $\mathbf{B} = (B, (\mathfrak{a}_i))$ of B and the fractional ideals $\mathfrak{a}_i$'s is called a *pseudo-basis* of Λ . Two pseudo-bases $\mathbf{B} = (B, (\mathfrak{a}_i))$ and $\mathbf{B}' = (B', (\mathfrak{b}_i))$ define the same lattice Λ if and only if there exists $U \in GL_r(K)$ whose entries satisfy $u_{i,j} \in \mathfrak{a}_i \mathfrak{b}_j^{-1}$, such that $B' = BU$ and $\mathfrak{a}_1 \cdots \mathfrak{a}_r = (\det U) \mathfrak{b}_1 \cdots \mathfrak{b}_r$. The lattice Λ is *integral* when $\langle \Lambda, \Lambda \rangle \subset \mathcal{O}_K$. For a prime $\mathfrak{q}$ of F , we write $\Lambda_{\mathfrak{q}} \stackrel{\text{def}}{=} \Lambda \otimes_{\mathcal{O}_F} \mathcal{O}_{F, \mathfrak{q}} \subset V_{\mathfrak{q}}$ for the completion of Λ at $\mathfrak{q}$.

Definition 2.5. Two lattices $\Lambda, \Lambda' \subset V$ are said to be *isometric* if there exists a global isometry $\sigma \in U(V)$ such that $\Lambda' = \sigma(\Lambda)$. The *isometry class* of Λ is

$$[\Lambda] \stackrel{\text{def}}{=} \{\text{Hermitian lattices } \Lambda' \text{ in } V \mid \Lambda \simeq \Lambda'\}.$$

The *automorphism group* $\text{Aut}(\Lambda)$ of Λ is the subgroup of global isometries given by $\text{Aut}(\Lambda) \stackrel{\text{def}}{=} \{\sigma \in U(V) \mid \sigma(\Lambda) = \Lambda\}$.

Remark 2.6. Since $\langle \cdot, \cdot \rangle$ is positive-definite, the unitary group $U(V)$ is compact while the matrices which preserve Λ form a discrete subgroup. In particular, $\text{Aut}(\Lambda)$ is finite.

Given a pseudo-basis $\mathbf{B} = (B, (\mathfrak{a}_i)_i)$ of Λ , the pair $\mathcal{Q} = (G, (\mathfrak{a}_i)_i)$ with $G = (\langle b_i, b_j \rangle)_{i,j}$ is its *pseudo-Gram matrix*, which for $V = K^r$ with the standard form is the notion of [MPPW24, Def. 3.6]. This is the representation we use *implicitly* whenever we refer to a lattice Λ , both in problem statements and as the input that an algorithm actually receives: every Λ has one, and it is integral exactly when Λ is. Pseudo-Gram matrices of two lattices $\Lambda, \Lambda' \subset V$ are

related by some $U \in GL_r(K)$ satisfying the ideal conditions above if and only if $\Lambda' = \sigma(\Lambda)$ for some $\sigma \in U(V)$ [MPPW24, Lem. 3.10]; hence a pseudo-Gram matrix determines exactly the isometry class $[\Lambda]$ and the group $\text{Aut}(\Lambda)$. We will later define a distribution on the equivalence class $[\Lambda]$ which will (implicitly) sample such a pseudo-Gram matrix representation of bounded size.

The *determinant* of Λ is the fractional ideal

$$\det(\Lambda) \stackrel{\text{def}}{=} (\det G) \mathfrak{a}_1 \overline{\mathfrak{a}_1} \cdots \mathfrak{a}_r \overline{\mathfrak{a}_r}$$

of $\mathcal{O}_K$, where $\mathcal{Q} = (G, (\mathfrak{a}_i)_i)$ is any pseudo-Gram matrix of Λ . It does not depend on the choice of pseudo-basis, and it is integral whenever Λ is. As $\det G \in F$ and $\mathfrak{a}_i \overline{\mathfrak{a}_i} = N_{K/F}(\mathfrak{a}_i) \mathcal{O}_K$, it is the extension to $\mathcal{O}_K$ of an ideal of $\mathcal{O}_F$; in particular, for a prime $\mathfrak{q}$ of F the condition $\mathfrak{q} \nmid \det(\Lambda)$ is unambiguous, and it holds if and only if $\mathfrak{p} \nmid \det(\Lambda)$ for the primes $\mathfrak{p}$ of K above $\mathfrak{q}$.

When two lattices Λ, Λ' are (globally) isometric, then they are isometric after completion at any prime $\mathfrak{q}$ of F . However, the converse does not hold. This discrepancy is captured by the notion of the *genus* of a lattice.

Definition 2.7 (Genus of a lattice). *Let $\Lambda \subset V$ be a Hermitian lattice. Its genus is defined as*

$$\text{gen}(\Lambda) \stackrel{\text{def}}{=} \{\Lambda' \subset V \mid \Lambda_{\mathfrak{q}} \simeq \Lambda'_{\mathfrak{q}} \quad \text{for every nonzero prime ideal } \mathfrak{q} \subset F\}.$$

Remark 2.8. The genus $\text{gen}(\Lambda)$ can be partitioned into a finite number of disjoint global isometry classes.

In particular, the genus is an invariant under isometries. A finer invariant known as the *special genus* has recently seen a resurgence of interest in lattice-based cryptography, following the work of Mureau [Mur25]. We recall its definition.

Definition 2.9 (Special Genus of a Hermitian Lattice [Mur25, Definition 2.1]). *Let K/F be a CM extension, and let $(V, \langle \cdot, \cdot \rangle)$ be a Hermitian space. Fix a Hermitian lattice $\Lambda \subset V$. A lattice $M \subset V$ belongs to the special genus $\text{sgen}(\Lambda)$ of Λ if there exists one global isometry $\sigma \in U(V)$ such that, for every prime $\mathfrak{q}$ of F , there exists a determinant-one local isometry $\varphi_{\mathfrak{q}} \in SU(V_{\mathfrak{q}})$ for which $M_{\mathfrak{q}} = \sigma(\varphi_{\mathfrak{q}}(\Lambda_{\mathfrak{q}}))$.*

In the above definition, the local isometries $\varphi_{\mathfrak{q}}$ may vary with $\mathfrak{q}$ as long as they remain in the local special unitary group. However, the global isometry σ must work at every prime. This immediately provides

$$[\Lambda] \subset \text{sgen}(\Lambda) \subset \text{gen}(\Lambda).$$

Furthermore, belonging to the same special genus is an equivalence relation on a genus.

2.4 Probabilities

Throughout this article, we measure the closeness between two random variables X, Y taking values in the same finite space $\mathcal{E}$ with the so-called *statistical distance*, defined as

$$\Delta(X, Y) \stackrel{\text{def}}{=} \frac{1}{2} \sum_{a \in \mathcal{E}} |\mathbb{P}(X = a) - \mathbb{P}(Y = a)| = \max_{A \subseteq \mathcal{E}} |\mathbb{P}(X \in A) - \mathbb{P}(Y \in A)|.$$

It only depends on the *probability distributions* of X and Y . More precisely, if f and g denote the distribution of X and Y respectively, we have

$$\Delta(X, Y) = \frac{1}{2} \|f - g\|_1.$$

A crucial property of the statistical distance is the *data processing inequality*: for any deterministic or randomized transformation φ ,

$$\Delta(\varphi(X), \varphi(Y)) \leq \Delta(X, Y). \quad (1)$$

For this inequality to hold, the only requirement is that the internal randomness of φ must be independent of X and Y . In particular, it implies that the success probability of any algorithm $\mathcal{A}$ for inputs distributed according to X or Y can only differ by at most $\Delta(X, Y)$.

Mass distribution on a special genus Since the genus, and the special genus, of lattices consist of finitely many isometry classes, one may be tempted to put the uniform distribution over them. However, it turns out that this is not the most mathematically natural way of defining a random isometry class. More precisely, we should take into account the symmetries induced by automorphisms within a given class. Intuitively, on average a lattice should have few automorphisms, and therefore the smaller the automorphism group of a class is, the more probable this isometry class should be under this distribution. This choice can be completely formalized by considering Haar measures, but we will not go into that detail here. This is the standard probability distribution over the genus, and more generally over a finite set of isometry classes (see for instance [vW24]). The naturality will also come from the fact that this distribution will appear as the limit distribution of the random walk considered in this article.

Definition 2.10. *Let X be the finite collection of isometry classes within a special genus. We define*

$$\text{Mass}(X) \stackrel{\text{def}}{=} \sum_{x \in X} \frac{1}{|\text{Aut}(\Lambda_x)|}, \quad \text{and } \pi_X(x) \stackrel{\text{def}}{=} \frac{|\text{Aut}(\Lambda_x)|^{-1}}{\text{Mass}(X)}$$

where Λ_x denotes any representative of the isometry class x .

In the main body of this work we only consider special genera, and thus for a lattice Λ we simply write $\mathcal{G}(\Lambda) \stackrel{\text{def}}{=} \text{sgen}(\Lambda)$. For a special genus $\mathcal{G}$ we write $\mathcal{H}(\mathcal{G}) \stackrel{\text{def}}{=} \pi_{\mathcal{G}}$ and call it the *Haar-uniform* distribution on $\mathcal{G}$. Note that it is not uniform on classes, but weights each class x by $|\text{Aut}(\Lambda_x)|^{-1}$.

2.5 Hermitian Lattices as Euclidean Module Lattices

The Hermitian form of Section 2.2 takes values in K , so it does not by itself make Λ a Euclidean lattice. We now make this extension, as the lattice problems we will consider in our reductions are mostly stated relative to the Euclidean metric: norms, successive minima, distances, determinants and Gaussians. These will only be used in Sections 4 and 5; the main random-walk machinery behind the reductions and the special genus only rely on the Hermitian form.

Let K/F be a CM extension with $d = [K : \mathbb{Q}]$, let $(V, \langle \cdot, \cdot \rangle)$ be a totally positive definite Hermitian space of dimension r over K , and let $\Lambda \subset V$ be a Hermitian lattice. We write $V_{\mathbb{R}} \stackrel{\text{def}}{=} V \otimes_{\mathbb{Q}} \mathbb{R}$ for the real ambient space and set $\text{span}(\Lambda) \stackrel{\text{def}}{=} V_{\mathbb{R}}$, the $\mathbb{R}$ -span of Λ . For $x \in V_{\mathbb{R}}$ we define

$$\|x\|^2 \stackrel{\text{def}}{=} \text{Tr}_{K/\mathbb{Q}} \langle x, x \rangle = \sum_{\sigma: K \rightarrow \mathbb{C}} \sigma(\langle x, x \rangle). \quad (2)$$

Since $\langle x, x \rangle \in F$ and the form is totally positive definite, every term is a positive real number, so Equation (2) is a positive definite quadratic form on $V_{\mathbb{R}}$. Hence Λ is a Euclidean lattice of rank

$$n \stackrel{\text{def}}{=} rd$$

in the n -dimensional space $\text{span}(\Lambda)$, and we write $\det_{\mathbb{Q}}(\Lambda)$ for its covolume with respect to Equation (2). Computing it on a pseudo-basis gives $\det_{\mathbb{Q}}(\Lambda)^2 = |\Delta_K|^r N_{K/\mathbb{Q}}(\det \Lambda)$, with $\det(\Lambda)$ the determinant ideal of Section 2.3. If Λ is integral, this is an integer, and a rational prime $p \nmid \det_{\mathbb{Q}}(\Lambda)^2$ is then unramified in K and satisfies $\mathfrak{p} \nmid \det(\Lambda)$ for every prime $\mathfrak{p} \mid p$ of $\mathcal{O}_K$. We also write $\lambda_i(\Lambda)$ for the i -th successive minimum of Λ , so in particular $\lambda_1(\Lambda) = \min_{v \in \Lambda \setminus \{0\}} \|v\|$, and $\text{dist}(t, \Lambda) \stackrel{\text{def}}{=} \min_{v \in \Lambda} \|t - v\|$ for $t \in \text{span}(\Lambda)$. Throughout, n always denotes this $\mathbb{Z}$ -rank rd , never the module rank r or the degree d alone.

Remark 2.11. For $V = K^r$ with the standard Hermitian form $\langle v, w \rangle = \sum_i \overline{v_i} w_i$, the form Equation (2) is exactly $\|\iota(x)\|^2$ for the canonical embedding $\iota: K \rightarrow \mathbb{C}^d$ applied coefficient-wise, since $\text{Tr}_{K/\mathbb{Q}}(\overline{v_i} v_i) = \sum_{\sigma} |\sigma(v_i)|^2$. This is the T_2 -norm of [MPPW24, §2]. Note that $\text{Tr}_{K/\mathbb{Q}}(a) = 2\text{Tr}_{F/\mathbb{Q}}(a)$ for $a \in F$, so Equation (2) is twice the form obtained from the trace of the totally real subfield; we use the $\text{Tr}_{K/\mathbb{Q}}$ normalization throughout. We also do *not* rescale by $1/d$, unlike [DPPvW22], whose normalization is tailored to the coefficient embedding of power-of-two cyclotomic fields.

2.6 Gaussians

For a lattice Λ of rank n and $\sigma > 0$ we write $\mathcal{N}_{\sigma}$ for the continuous Gaussian on $\text{span}(\Lambda)$ and, for a coset $c + \Lambda$ with $c \in \text{span}(\Lambda)$, $\mathcal{N}_{\sigma}(c + \Lambda)$ for the discrete Gaussian on that coset, all with density proportional to $\exp(-\pi \|x\|^2 / \sigma^2)$. We abbreviate $\mathcal{N}_{\sigma}(\Lambda) \stackrel{\text{def}}{=} \mathcal{N}_{\sigma}(0 + \Lambda)$, and write $\mathcal{N}_{\sigma}(S) \stackrel{\text{def}}{=} \sum_{x \in S} \exp(-\pi \|x\|^2 / \sigma^2)$ for the Gaussian mass of a countable set S . Whether $\mathcal{N}_{\sigma}(\cdot)$ denotes a distribution

or a mass is always determined by its argument: a coset gives the distribution supported on it, an arbitrary countable set gives the number. Here $\|\cdot\|$ and $\text{span}(\Lambda)$ are those of Section 2.5 whenever Λ is a module lattice, so that $n = rd$. All Gaussian parameters in this work (such as s , σ and σ') follow this normalization, in which the standard deviation per coordinate is $\sigma/\sqrt{2\pi}$, the coordinates being taken with respect to any orthonormal basis of $\text{span}(\Lambda)$ for the form (2). For $\delta \in (0, 1)$ we write $\eta_\delta(\Lambda)$ for the smoothing parameter of Λ [MR07]. Recall that $\eta_\delta(c\Lambda) = c\eta_\delta(\Lambda)$ for $c > 0$, and that $\eta_\delta(\Lambda) \geq \eta_\delta(\Lambda')$ whenever $\Lambda \subseteq \Lambda'$.

Remark 2.12. Both facts above concern scaling by a rational $c > 0$, and every reduction in Section 4 uses them with $c = p$ a rational prime. Multiplication by a prime ideal $\mathfrak{p}$ does not preserve the Euclidean geometry, so neither identity has an ideal analogue. This is why in Algorithm 2 we will make sure to apply the paired walk of Section B.1 to all prime ideals above a totally split p consecutively, which combines to multiplication by the rational p in our proofs.

Lemma 2.13 ([Ban93, Lem. 1.5(i)]). *For every lattice Λ of rank n and every $\sigma > 0$ we have $\mathbb{P}_{x \leftarrow \mathcal{N}_\sigma(\Lambda)}[\|x\| > \sigma\sqrt{n}] \leq 2^{-n}$.*

Lemma 2.14 ([MR07, Lem. 4.4]). *Let Λ' be a lattice, $\delta \in (0, 1)$ and $\sigma \geq \eta_\delta(\Lambda')$. Then all coset masses $\mathcal{N}_\sigma(v + \Lambda')$ with $v \in \text{span}(\Lambda')$ agree up to a factor $\frac{1+\delta}{1-\delta}$. In particular, if Λ' has finite index in a lattice $\Lambda \supseteq \Lambda'$, then the coset of $x \leftarrow \mathcal{N}_\sigma(c + \Lambda)$ in $(c + \Lambda)/\Lambda'$ is within statistical distance 2δ of uniform, for any $c \in \text{span}(\Lambda)$.*

For a $\mathbb{Z}$ -basis $B = (b_1, \dots, b_n)$ of Λ (or equivalently a Gram matrix) we write $\|B^*\| \stackrel{\text{def}}{=} \max_i \|b_i^*\|$ for its largest Gram–Schmidt norm. Given B one can efficiently sample from $\mathcal{N}_\sigma(\Lambda)$ for any $\sigma \geq \|B^*\| \sqrt{\ln(2n+4)}/\pi$ [BLP⁺13, Lem. 2.3], and we have the bounds

$$\begin{aligned} \eta_\delta(\Lambda) &\leq \|B^*\| \sqrt{\ln(2n(1+1/\delta))}/\pi && [\text{GPV08, Lem. 3.1}], \\ \eta_\delta(\Lambda) &\leq \lambda_n(\Lambda) \sqrt{\ln(2n(1+1/\delta))}/\pi && [\text{MR07, Lem. 3.3}]. \end{aligned}$$

The module counterpart of this statement is that when Λ is given by a pseudo-Gram matrix rather than a basis, one can sample $\mathcal{N}_\sigma(\Lambda)$ directly from that pseudo-Gram matrix by [MPPW24, Lem. 3.8], at the cost of a threshold on σ depending on Δ_K and the norms of the ideals $\mathfrak{a}_i$; so the representation of Section 2.3 suffices for everything below.

2.7 Average-case Distribution on the Isometry Class

To state average-case problems we need to also consider a distribution $\mathcal{D}_s([A])$ on every isometry class $[A]$ of Section 2.3, and in particular on the representation of its elements. This directly allows one to define average-case versions of the lattice isomorphism problem, but also average-case versions (over representatives of $[A]$) of the more standard lattice problems. For more details on instantiating

such a distribution we refer to [DvW22] (for $\mathbb{Z}$ -lattices) and [MPPW24] (for $\mathcal{O}_K$ -lattices). We will simply characterize these distributions by the properties we will use in our work.

Definition 2.15 (Average-case class distribution). *Let $s > 0$ and let Λ be a Hermitian lattice of $\mathbb{Z}$ -rank n . A distribution $\mathcal{D}_s([\Lambda])$ over lattices $\Lambda' \in [\Lambda]$, with output given as pseudo-Gram matrices (Section 2.3), is called an average-case class distribution for s if*

- (i) *it depends only on the isometry class $[\Lambda]$;*
- (ii) *(efficient sampling) there is a probabilistic polynomial-time algorithm that, given any $\Lambda' \in [\Lambda]$ in a representation from which $\mathcal{N}_s(\Lambda')$ can be sampled efficiently, outputs $\Lambda'' \leftarrow \mathcal{D}_s([\Lambda])$ together with an isometry $\sigma \in U(V)$ with $\sigma(\Lambda') = \Lambda''$;*
- (iii) *(Gaussian sampling) except with probability $300n \cdot 2^{-n}$ over $\Lambda'' \leftarrow \mathcal{D}_s([\Lambda])$, one can sample $\mathcal{N}_\sigma(\Lambda'')$ exactly from the representation of Λ'' , for every $\sigma \geq s\sqrt{n \ln(2n+4)}/\pi$.*

Here $[\Lambda]$ is the $U(V)$ -isometry class, not the $\mathbb{Z}$ -isometry class of [DvW22], which is what makes the weights $|\text{Aut}(\Lambda_i)|^{-1}$ of the distribution $\mathcal{H}(\mathcal{G})$ (Definition 2.10) well defined.

The parameter s measures the quality of the average-case representative (it determines the length of the (pseudo-)basis vectors): (ii) needs s at the quality of the representative one starts from, and by (iii) the sampled one has quality $s\sqrt{n \ln(2n+4)}/\pi$, so it reveals nothing beyond $[\Lambda]$ and s . Since every representative allows one to sample $\mathcal{N}_s$ for $s \geq 2^{O(n)}\lambda_n(\Lambda)$ after LLL reduction (see above), (ii) gives the worst-case to average-case reduction in that regime: $\mathcal{D}_s([\Lambda])$ can be sampled from *every* $\Lambda' \in [\Lambda]$, so a problem on $\mathcal{D}_s([\Lambda])$ is as hard as on its hardest representative (cf. [DvW22, Lem. 3.9]). We simply write $\mathcal{D}([\Lambda])$ if s is fixed to such a sufficiently large value $s \geq 2^{\Theta(n)}\lambda_n(\Lambda)$ ⁶. In the rest of this work we will use this distribution in our problem statements for simplicity, but one could as well state them more precisely and track the $O(p \cdot \text{poly}(n))$ loss on s throughout the reduction.

Remark 2.16. For $\mathbb{Z}$ -lattices and $s \geq \max\{\lambda_n(\Lambda), \eta_{1/2}(\Lambda)\}$ the Gaussian form distribution of [DvW22, Def. 3.3, Alg. 1] is such a distribution, (ii) being [DvW22, Lem. 3.5] and (iii) [DvW22, Lem. 3.6] with $\delta = 1/2$. For module lattices, (i) and (ii) are [MPPW24, Alg. A.1, Lem. A.5, Prop. A.7] for any rank, any (not necessarily free) Λ and any number field; their algorithm only needs samples from $\mathcal{N}_s(\Lambda')$, so no factor Δ_K enters. For rank 2 over power-of-two cyclotomic fields this is AC_σ of [DPPvW22, Alg. 6, Lem. 4–5]. Property (iii) is not stated in [MPPW24], but the argument of [DvW22, Lem. 3.6] adapts: the coordinates of the Gaussian samples in the canonical pseudo-basis are class-invariant by the

⁶ Technically, the sufficiently large parameter s depends on the value of λ_n , but this is either assumed to be known (for the lattice we use), or assumed to be sufficiently bounded (for the random genus lattices we encounter) following (a structured version of) [vW24].

proof of [MPPW24, Lem. A.5], so they can be part of the representation and give $O(r)$ vectors of norm at most $s\sqrt{n}$ from which $\mathcal{N}_\sigma(A'')$ is sampled as in [DvW22].

3 Kneser Neighbours

The key ingredient in our work is the notion of neighbour lattices, which was first introduced by Kneser [Kne57] for integral $\mathbb{Z}$ -lattices in a quadratic space, and later generalized to lattices in arbitrary CM extensions. We follow the description of [Sch98, GV14, Kir16].

3.1 Neighbour Lattices

Let K/F be a CM extension, and Λ a Hermitian lattice.

Definition 3.1 ($\mathfrak{p}$ -neighbours). *Let $\mathfrak{p}$ be an unramified prime ideal in $\mathcal{O}_K$ which does not divide $\det(\Lambda)$. An $\mathcal{O}_K$ -lattice M in V is called a $\mathfrak{p}$ -neighbour of Λ iff the following holds: M is integral and we have the isomorphisms of $\mathcal{O}_K$ -modules*

$$M/(\Lambda \cap M) \simeq \mathcal{O}_K/\mathfrak{p}, \quad \Lambda/(\Lambda \cap M) \simeq \mathcal{O}_K/\bar{\mathfrak{p}}.$$

Reversing the relation exchanges $\mathfrak{p}$ and its conjugate.

The above description immediately implies that if M is a $\mathfrak{p}$ -neighbour of Λ , then

$$\mathfrak{p}M \subset \Lambda, \quad \bar{\mathfrak{p}}\Lambda \subset M.$$

The number of $\mathfrak{p}$ -neighbours of a given lattice Λ is always finite, and depends only on the prime $\mathfrak{q} = \mathfrak{p} \cap \mathcal{O}_F$ below $\mathfrak{p}$. When it is split, [Sch98, Lemma 3.3] shows that there are $D_{\mathfrak{q}} = \frac{q^r - 1}{q - 1}$ many such neighbours, where $q = |\mathcal{O}_F/\mathfrak{q}|$.

The main property of neighbours that will be crucial in this article is that taking a $\mathfrak{p}$ -neighbour of a lattice Λ only modifies Λ locally. In fact, every $\mathfrak{p}$ -neighbour of Λ is of the form

$$\Lambda' = \mathfrak{p}^{-1}x + \{y \in \Lambda : \langle y, x \rangle \in \mathfrak{p}\},$$

where $x \in \Lambda \setminus \mathfrak{p}\Lambda$ satisfies $\langle x, x \rangle \in \mathfrak{p}\bar{\mathfrak{p}}$, and they can easily be constructed (see Algorithm 1).

Lemma 3.2. *Let $\mathfrak{q} = \mathfrak{p}\bar{\mathfrak{p}}$ be a prime ideal of $\mathcal{O}_F$ which splits in K . If M is a $\mathfrak{p}$ -neighbour of Λ , then $M_{\mathfrak{t}} = \Lambda_{\mathfrak{t}}$ for every prime $\mathfrak{t} \neq \mathfrak{q}$.*

Conversely, a global lattice is determined by all its completions at finite primes:

$$\Lambda = \{x \in V : x \in \Lambda_{\mathfrak{t}} \text{ for every prime ideal } \mathfrak{t} \text{ of } \mathcal{O}_F\}.$$

Therefore, a neighbour M of Λ can equivalently be described by how it differs after completion. Choosing one local lattice $M_{\mathfrak{q}}$ defines

$$M = \{x \in V : x \in M_{\mathfrak{q}} \text{ and } x \in \Lambda_{\mathfrak{t}} \text{ for every } \mathfrak{t} \neq \mathfrak{q}\}.$$

It is well-known that two neighbour lattices belong to the same genus (see [Sch98, Lemma 2.8]). On the other hand, the special genus need not be preserved by a single step of the walk. More precisely, [Sch98, Theorem 2.10] computes exactly how it affects the special genus. The precise statement is beyond the scope of this article, but as a corollary, Schiemann shows that all $\mathfrak{p}$ -neighbours of all lattices in a fixed special genus X_0 lie in one special genus, denoted by X_1 . Conversely, all $\bar{\mathfrak{p}}$ -neighbours of a lattice in X_1 lie in X_0 . This motivates the following walk.

Definition 3.3 (Paired Random Walk). *Let $\mathfrak{q} = \mathfrak{p}\bar{\mathfrak{p}}$ be a prime ideal of $\mathcal{O}_F$ which splits in K . Given a lattice Λ_0 , one step of our paired random walk consists of the concatenation of the following two steps:*

1. *Sample a uniformly random $\mathfrak{p}$ -neighbour Λ_1 of Λ_0 .*
2. *Sample a uniformly random $\bar{\mathfrak{p}}$ -neighbour Λ_2 of Λ_1 .*

The resulting lattice is Λ_2 .

The previous discussion can be summarized as the following proposition.

Proposition 3.4. *Let Λ be an $\mathcal{O}_K$ -lattice and let M be obtained from the paired random walk. Then Λ and M belong to the same special genus. Moreover, M and Λ are locally equal at all primes $\mathfrak{t} \neq \mathfrak{q}$ and*

$$\mathfrak{q}\Lambda_2 \subset \Lambda_0, \quad \mathfrak{q}\Lambda_0 \subset \Lambda_2.$$

3.2 Finding Random Neighbours

Lemma 3.5. *Algorithm 1 runs in polynomial time and outputs a uniformly random $\mathfrak{p}$ -neighbour N of Λ .*

Proof. Throughout we use that $\mathfrak{p}, \bar{\mathfrak{p}} \nmid \det(\Lambda)$, and therefore the Hermitian form induces a *non-degenerate* pairing

$$\Lambda/\bar{\mathfrak{p}}\Lambda \times \Lambda/\mathfrak{p}\Lambda \longrightarrow \mathcal{O}_K/\mathfrak{p}, \quad (x, y) \mapsto \langle x, y \rangle \bmod \mathfrak{p}$$

between two r -dimensional $\mathbb{F}_q$ -vector spaces, where $q = |\mathcal{O}_K/\mathfrak{p}| = |\mathcal{O}_K/\bar{\mathfrak{p}}|$. Note that this pairing is well defined, as $\langle \alpha x, y \rangle = \bar{\alpha} \langle x, y \rangle$ and $\alpha \in \bar{\mathfrak{p}}$ implies $\bar{\alpha} \in \mathfrak{p}$.

The algorithm is well defined and efficient. Since $u \neq 0$, the map $x \mapsto \langle x, u \rangle \bmod \mathfrak{p}$ is a non-zero linear form on $\Lambda/\bar{\mathfrak{p}}\Lambda$ by the non-degeneracy above, and thus its kernel has size $q^{r-1} > 0$, independent of u ; in particular a v as in line 2 exists. For the ℓ obtained in line 3 we have (since the left argument is the ‘conjugate’

Algorithm 1 Finding a uniformly random neighbour of an integral lattice over $\mathcal{O}_K$ (split case)

Require:

- A lattice Λ of rank r with integral inner products over $\mathcal{O}_K$, the ring of integers of a CM field K .
- A prime ideal $\mathfrak{p}$ of $\mathcal{O}_K$, lying above a prime $\mathfrak{q}$ of $\mathcal{O}_F$, satisfying $\mathfrak{p} \nmid 2$, $\mathfrak{p} \nmid \det(\Lambda)$ and $\bar{\mathfrak{p}} \nmid \det(\Lambda)$, such that $\mathfrak{q}$ splits in K as $\mathfrak{q} = \mathfrak{p}\bar{\mathfrak{p}}$.

Ensure: A uniformly random $\mathfrak{p}$ -neighbour lattice N of Λ .

- 1: Sample a non-zero $u \in \Lambda/\mathfrak{p}\Lambda$ uniformly.
 - 2: Compute any $v \in \Lambda/\bar{\mathfrak{p}}\Lambda$ satisfying $\langle v, u \rangle \equiv 0 \pmod{\mathfrak{p}}$.
 - 3: Compute any $\ell \in \Lambda$ such that $\ell \equiv u$ modulo $\mathfrak{p}\Lambda$ and $\ell \equiv v$ modulo $\bar{\mathfrak{p}}\Lambda$. (This implies that $\langle \ell, \ell \rangle \equiv 0 \pmod{\mathfrak{p}\bar{\mathfrak{p}}}$).
 - 4: Construct $M = \{x \in \Lambda \mid \langle x, \ell \rangle \equiv 0 \pmod{\mathfrak{p}}\}$.
 - 5: Return $N = M + \ell \cdot \mathfrak{p}^{-1}$.
-

argument) that $\langle \ell, \ell \rangle \equiv \langle v, u \rangle \equiv 0 \pmod{\mathfrak{p}}$, and $\langle \ell, \ell \rangle \equiv \langle u, v \rangle \equiv 0 \pmod{\bar{\mathfrak{p}}}$, since $\langle v, u \rangle \equiv 0 \pmod{\mathfrak{p}}$. Therefore, as claimed in line 3,

$$\langle \ell, \ell \rangle \equiv 0 \pmod{\mathfrak{p}\bar{\mathfrak{p}}}. \quad (3)$$

Lines 1–3 are linear algebra over $\mathcal{O}_K/\mathfrak{p} \simeq \mathcal{O}_K/\bar{\mathfrak{p}} \simeq \mathbb{F}_q$ plus a CRT lift, hence polynomial time. Line 4 is done by first computing $M/\mathfrak{p}\Lambda$ by means of linear algebra over the field $\mathcal{O}_K/\mathfrak{p}$, then extending it by $\mathfrak{p}\Lambda$, and applying a basis reduction algorithm to remove linear dependencies to arrive at a pseudo-basis of M . Line 5 applies again such a linear-dependency-removing algorithm, e.g. [BF12].

N is a $\mathfrak{p}$ -neighbour of Λ . We first check that $M = N \cap \Lambda$. The inclusion $M \subseteq N \cap \Lambda$ is clear. Conversely, let $m + w\ell \in \Lambda$ with $m \in M$ and $w \in \mathfrak{p}^{-1}$. If $w \notin \mathcal{O}_K$ then $w\mathcal{O}_K + \mathcal{O}_K = \mathfrak{p}^{-1}$, so $\mathfrak{p}^{-1}\ell \subseteq \Lambda$ and thus $\ell \in \mathfrak{p}\Lambda$, contradicting $u \neq 0$. Hence $w \in \mathcal{O}_K$ and $N \cap \Lambda = M$.

Next, M is the kernel of the $\mathcal{O}_K$ -linear map $\Lambda \rightarrow \mathcal{O}_K/\bar{\mathfrak{p}}$ with $x \mapsto \langle \ell, x \rangle \pmod{\bar{\mathfrak{p}}}$, which is the conjugate of the defining condition $\langle x, \ell \rangle \equiv 0 \pmod{\mathfrak{p}}$. This map is non-zero by non-degeneracy and $\ell \not\equiv 0$ in $\Lambda/\mathfrak{p}\Lambda$, hence surjective, and so $\Lambda/M \simeq \mathcal{O}_K/\bar{\mathfrak{p}}$.

By Equation (3) we have $\mathcal{O}_K\ell \subseteq M$, and $\mathfrak{p}^{-1}\ell \cap M = \mathcal{O}_K\ell$ by the first paragraph. Hence $N/M = (M + \ell \cdot \mathfrak{p}^{-1})/M = (\ell \cdot \mathfrak{p}^{-1})/(\ell \cdot \mathfrak{p}^{-1} \cap M) = (\ell \cdot \mathfrak{p}^{-1})/(\ell \cdot \mathcal{O}_K) \simeq \mathcal{O}_K/\mathfrak{p}$.

Finally, N is integral: for $w, w' \in \mathfrak{p}^{-1}$ we have $\langle w\ell, w'\ell \rangle = \bar{w}w'\langle \ell, \ell \rangle \in \bar{\mathfrak{p}}^{-1}\mathfrak{p}^{-1}\mathfrak{p}\bar{\mathfrak{p}} \subseteq \mathcal{O}_K$ by Equation (3); the inner products between $x \in M$ and $w\ell$ are integral since $\langle x, w\ell \rangle = w\langle x, \ell \rangle \in \mathfrak{p}^{-1}\mathfrak{p} \subseteq \mathcal{O}_K$ by the definition of M ; and the inner products among elements of M alone are integral since M is a sublattice of Λ .

N is uniformly random. Both M and N depend on ℓ only through the line (‘projective class’) $L_u = \{\alpha u \in \Lambda/\mathfrak{p}\Lambda \mid \alpha \in \mathcal{O}_K/\mathfrak{p}\}$ it spans in $\Lambda/\mathfrak{p}\Lambda$. By [Sch98, Lemma 3.1] and [Sch98, (1.) after Lemma 3.1], the map $L_u \mapsto N$ is a bijection

between these lines and the $\mathfrak{p}$ -neighbours of Λ , of which there are $D_{\mathfrak{q}} = \frac{q^r - 1}{q - 1}$ many. Since every line contains exactly $q - 1$ non-zero elements, the uniformly random non-zero u of line 1 yields a uniformly random line, and hence a uniformly random $\mathfrak{p}$ -neighbour N . $\square$

3.3 The random walk procedure

Algorithm 2 The random walk procedure $\mathcal{W}$

Require: An integral module lattice Λ of rank r over a CM field K of degree d , and hence of $\mathbb{Z}$ -rank $n = rd$, an $\varepsilon \in (0, 1)$, and optionally an auxiliary modulus $Q \in \mathbb{Z}_{\geq 1}$.
Ensure: A prime $p \nmid 2Q \det_{\mathbb{Q}}(\Lambda)^2$ and a p -neighbour Λ_k of Λ with $\Delta([\Lambda_k], \mathcal{H}(\mathcal{G}(\Lambda))) \leq \varepsilon$.

- 1: Compute a prime number p , totally split in K , with $p \geq \max(\text{poly}(d) \cdot \det_{\mathbb{Q}}(\Lambda)^{O_r(1)/n} \cdot (1/\varepsilon)^{2/d}, d^3, d \log(2Q))$ as in Corollary B.18 and Lemma 3.6, and coprime to $2Q \det_{\mathbb{Q}}(\Lambda)^2$.
 - 2: $\Lambda_0 = \Lambda$
 - 3: Factor $p\mathcal{O}_K = \prod_{i=1}^k \mathfrak{p}_i$
 - 4: **for** $i = 1$ to k **do**
 - 5: Sample a random $\mathfrak{p}_i$ -neighbour Λ_i of Λ_{i-1} via Algorithm 1
 - 6: **end for**
 - 7: Output Λ_k, p .
-

Lemma 3.6. *Let $K = \mathbb{Q}(\zeta_m)$ be a cyclotomic field, let $R \geq 1$ be an integer, and let $N \geq C \max(\phi(m)^3, \phi(m) \log R)$ for a suitable absolute constant C . Assume GRH for the Dirichlet L -functions modulo m . Then there exists a probabilistic algorithm that computes a prime number $p \in [N, 2N]$ that splits completely over $\mathcal{O}_K$ and satisfies $p \nmid R$, in expected time $\text{poly}(\log N, \log R)$.*

Proof. For the m -th cyclotomic field $\mathbb{Q}(\zeta_m)$, a prime number p splits completely if and only if $p \equiv 1 \pmod{m}$. Assuming GRH for these L -functions, one can derive that [MV06, Corollary 13.8]

$$|\pi_m(X) - \text{Li}(X) \cdot \phi(m)^{-1}| \leq O(\sqrt{X} \log(mX)),$$

where $\pi_m(X) \stackrel{\text{def}}{=} |\{p \text{ prime} \mid p \leq X, p \equiv 1 \pmod{m}\}|$, ϕ is the Euler totient and where Li is the logarithmic integral. Looking at $\pi_m(2N) - \pi_m(N)$ the main term satisfies $(\text{Li}(2N) - \text{Li}(N)) \cdot \phi(m)^{-1} \geq \frac{N}{\phi(m) \log(2N)}$, while the error term has size $O(\sqrt{N} \log(mN))$. For $N \geq \Omega(\phi(m)^3)$ these are of size $\Omega(N^{2/3} / \log(N))$ and $O(\sqrt{N} \log(N))$ respectively. For a suitably large absolute constant C and $N \geq C\phi(m)^3$, the error term is thus at most half the main term on $[N, 2N]$, so this interval contains at least $\frac{N}{2\phi(m) \log(2N)}$ primes $p \equiv 1 \pmod{m}$. At most $\log_N R$ primes $p \geq N$ divide R , which for $N \geq C\phi(m) \log R$ is at most half of them. Hence, sampling $k \in \mathbb{N}$ uniformly with $1 + km \in [N, 2N]$, a fraction

$\Omega(1/\log N)$ of the numbers $1 + km$ are primes not dividing R . Finding such a prime therefore takes expected time $\text{poly}(\log N, \log R)$, by using an efficient prime testing algorithm, e.g., [AKS04], and a gcd computation with R . $\square$

In Algorithm 2 we apply Lemma 3.6 with $R = 2Q \det_{\mathbb{Q}}(A)^2$, where $\phi(m) = d$. Since $\log \det_{\mathbb{Q}}(A)^2 \leq 2n \det_{\mathbb{Q}}(A)^{1/n}$, the condition $N \geq Cd \log R$ only adds the term $d \log(2Q)$ to the lower bound on p , up to the constants already hidden in $\text{poly}(d) \cdot \det_{\mathbb{Q}}(A)^{O_r(1)/n}$. This term is $\text{poly}(d)$ whenever $\log Q = \text{poly}(n)$. This covers the application in Section 5, where the auxiliary modulus Qt is exponentially large but $\log(Qt) = \text{poly}(n)$.

Remark 3.7. Recall that K is a CM field with maximal real subfield F . For the random walk algorithm to proceed, it is required that each prime ideal $\mathfrak{q}_i$ occurring in the decomposition in $\mathcal{O}_F$: $p\mathcal{O}_F = \prod_{i=1}^t \mathfrak{q}_i$ splits in $\mathcal{O}_K$, i.e. $\mathfrak{q}_i\mathcal{O}_K = \mathfrak{p}_i\bar{\mathfrak{p}}_i$. Indeed, in this paper the random walk is *paired*, i.e. every sub-step consists of randomization with respect to a prime ideal together with its conjugate prime ideal.

Using Chebotarev-like arguments [NS13, Chapter 7, §13], the probability that a prime $\mathfrak{q}_i$ in $\mathcal{O}_F$ splits in $\mathcal{O}_K$ is $1/2$ for a ‘random prime’. For totally real fields F without Galois structure the probability for an unramified prime p to ‘split completely over K/F ’ (as described above) can thus be as small as 2^{-t} , where t is the number of prime ideals in the decomposition of $p\mathcal{O}_F$.

For the case of an ‘arbitrary’ real field F , the Galois group of the normal closure is (close to) the symmetric group, and similar Chebotarev-like arguments can be used to show that (totally) inert primes (i.e. $t = 1$) occur often (a proportion of $1/[F : \mathbb{Q}]$). Hence, in such cases, primes of the form $p = \mathfrak{p}\bar{\mathfrak{p}}$ in K occur often as well. Note, however, that the Riemann hypothesis bound on the density might depend on the discriminant of this normal closure of F (which has degree $O([F : \mathbb{Q}]!)$), which might cause the first such inert p to be too large for our use case. A case where t is too large on average is that of real multiquadratic fields $F = \mathbb{Q}(\sqrt{m_1}, \dots, \sqrt{m_k})$, for which $t \approx [F : \mathbb{Q}]/2$, causing 2^{-t} to be exponentially small in the degree.

Summarizing, making the technique in Algorithm 2 work for any CM field requires addressing a number of technical details and edge cases. Using cyclotomic fields, we circumvent these difficulties entirely since the splitting behaviour of p in $\mathbb{Q}(\zeta_m)$ is totally determined by the order of p in the group $(\mathbb{Z}/m\mathbb{Z})^\times$. The condition $p \equiv 1 \pmod{m}$ is equivalent to p being completely split in $\mathbb{Z}[\zeta_m]$, and $-1 \notin \langle \bar{p} \rangle$ is equivalent to p being ‘split’ in $\mathbb{Q}(\zeta_m)/\mathbb{Q}(\zeta_m)^+$, where $\langle \bar{p} \rangle \subseteq (\mathbb{Z}/m\mathbb{Z})^\times$ is the subgroup generated by $p \pmod{m}$. For simplicity, we just use primes p that are totally split in $\mathbb{Z}[\zeta_m]$.

Lemma 3.8. *For $K = \mathbb{Q}(\zeta_m)$, and assuming GRH, Algorithm 2 is efficient, and outputs a p -neighbouring lattice Λ_k of Λ .*

Proof. As computing the determinant of Λ , sampling the prime p (see Lemma 3.6), factoring a prime into prime ideals [Coh93, §6.2] and computing neighbours (Lemma 3.5) can be done efficiently, Algorithm 2 is efficient. $\square$

3.4 Equidistribution of the Kneser p -neighbouring walk

We briefly instantiate and summarize here the equidistribution results from Section B.1.

Theorem 3.9 (Equidistribution theorem). *Let $r \geq 2$ be fixed and let $\varepsilon \in (0, 1)$. Let Λ be an integral Hermitian lattice of rank r , over a cyclotomic field K of degree d , and let $\mathcal{G}$ be its special genus.*

Then, assuming GRH, there exists an efficient procedure $\mathcal{W}$ (see Algorithm 2) which on input (Λ, ε) returns (Λ_k, p) such that $\Lambda_k \in \mathcal{G}$ is a p -neighbour of Λ and $p = \text{poly}(d) \cdot \det_{\mathbb{Q}}(\Lambda)^{O_r(1)/n} \cdot (1/\varepsilon)^{2/d}$. Moreover, the distribution of $[\Lambda_k]$ is ε -close to $\mathcal{H}(\mathcal{G}(\Lambda))$ with respect to the statistical distance.

Proof. The procedure $\mathcal{W}$ is Algorithm 2. By Lemma 3.5 each step samples a uniformly random neighbour. As the corresponding Hecke operators commute, the order of the steps does not matter, so grouping each $\mathfrak{p}$ -step with its $\bar{\mathfrak{p}}$ -step the class $[\Lambda_k]$ is distributed as $\prod_i P_{\mathfrak{q}_i} \delta_{\Lambda}$ (Section B.1); since p satisfies the requirements of Corollary B.18, $\Delta([\Lambda_k], \mathcal{H}(\mathcal{G}(\Lambda))) \leq \varepsilon$. Finally, each such pair of steps preserves the special genus [Sch98, Theorem 2.10], so $\Lambda_k \in \mathcal{G}$. $\square$

4 Reduction

We now describe the worst-case to average-case reduction for a variety of lattice problems. The general strategy, which we sketch here, is the same for each reduction. Given a worst-case instance on a Hermitian lattice Λ , we run the paired Kneser p -neighbouring walk $\mathcal{W}$ of Algorithm 2 to obtain a lattice Λ_k in the same special genus $\mathcal{G} = \mathcal{G}(\Lambda)$, whose class is ε -close to the Haar-uniform distribution $\mathcal{H}(\mathcal{G})$ by Theorem 3.9. We then call the average-case oracle on Λ_k , appropriately transforming the instance to Λ_k and its answer back to Λ . Both steps are possible because the walk is *reversible up to a factor p* : the two lattices contain rescaled copies of one another,

$$p\Lambda \subseteq \Lambda_k \quad \text{and} \quad p\Lambda_k \subseteq \Lambda,$$

so multiplying by p maps a vector or a target from either lattice into the other. This accounts for every approximation loss below: one factor p is lost in transporting the instance or its solution, and a second one exactly when the quality of a solution is measured against a quantity that itself changes under the walk, such as λ_1 or $\text{dist}(t, \cdot)$, rather than against the walk-invariant $\det_{\mathbb{Q}}(\Lambda)^{1/n}$. Absolute or Hermite-type problems therefore lose a factor p , and their relative counterparts a factor p^2 .

Lemma 4.1. *Let $(\Lambda_k, p) \leftarrow \mathcal{W}(\Lambda, \varepsilon)$ be the output of Algorithm 2. Then, for every $\delta \in (0, 1)$ and every $t \in \text{span}(\Lambda)$,*

- (i) $\text{span}(\Lambda_k) = \text{span}(\Lambda)$ and $\det_{\mathbb{Q}}(\Lambda_k) = \det_{\mathbb{Q}}(\Lambda)$;
- (ii) $p\Lambda \subseteq \Lambda_k$ and $p\Lambda_k \subseteq \Lambda$;

- (iii) $\lambda_1(\Lambda) \leq p \lambda_1(\Lambda_k) \leq p^2 \lambda_1(\Lambda)$;
- (iv) $\text{dist}(pt, \Lambda_k) \leq p \text{dist}(t, \Lambda)$ and $\text{dist}(pt, \Lambda) \leq p \text{dist}(t, \Lambda_k)$;
- (v) $\eta_\delta(\Lambda_k) \leq p \eta_\delta(\Lambda)$.

Proof. Write $p\mathcal{O}_K = \prod_{i=1}^k \mathfrak{p}_i$ and let $\Lambda = \Lambda_0, \dots, \Lambda_k$ be the intermediate lattices of the walk. Each step satisfies $\mathfrak{p}_i \Lambda_i \subseteq \Lambda_{i-1}$ and $\bar{\mathfrak{p}}_i \Lambda_{i-1} \subseteq \Lambda_i$ (Section 3.1), and has $[\Lambda_{i-1} : \Lambda_{i-1} \cap \Lambda_i] = [\Lambda_i : \Lambda_{i-1} \cap \Lambda_i] = N(\mathfrak{p}_i)$, whence (i). Composing the inclusions over $i = 1, \dots, k$, and using that conjugation permutes the primes above the totally split p , both $\prod_i \mathfrak{p}_i$ and $\prod_i \bar{\mathfrak{p}}_i$ equal $p\mathcal{O}_K$; this gives (ii), and is where we use that the walk takes a step at every prime above p (Remark 2.12). Items (iii)–(v) are immediate from (ii), as λ_1 , $\text{dist}(t, \cdot)$ and η_δ do not increase when passing to a superlattice and scale by p on $p\Lambda$ and $p\Lambda_k$. $\square$

Conventions. Throughout this section $K = \mathbb{Q}(\zeta_m)$ is a cyclotomic field of degree d , which is in particular a CM field with totally real subfield F ; see Remark 4.4 for more general CM fields. Lattices are Hermitian $\mathcal{O}_K$ -lattices of fixed rank $r \geq 2$ with integral inner products, of $\mathbb{Z}$ -rank $n = rd$, and $\mathcal{G} = \mathcal{G}(\Lambda)$ denotes the special genus of Λ . Let $\varepsilon \in (0, 1)$. We write $(\Lambda_k, p) \leftarrow \mathcal{W}(\Lambda, \varepsilon)$ for the output of Algorithm 2 (with $Q = 1$ by default), so that by Theorem 3.9 the class $[\Lambda_k]$ is ε -close to $\mathcal{H}(\mathcal{G})$ and $p = \text{poly}(d) \cdot \det_{\mathbb{Q}}(\Lambda)^{O_r(1)/n} \cdot (1/\varepsilon)^{2/d}$. Average-case problems always receive their input from $\mathcal{D}(\mathcal{H}(\mathcal{G}))$, with the rerandomization of Remark 4.3 applied implicitly, and $\beta \in [0, 1]$ denotes the success probability (or advantage) of the oracle. For every problem $\mathcal{P}$ with parameter θ defined below we write $\theta\text{-ac}\mathcal{P}^{\mathcal{G}}$ for its average-case variant: the very same problem, but on input $\Lambda \leftarrow \mathcal{D}(\mathcal{H}(\mathcal{G}))$ instead of a worst-case Λ . As the walk is independent of the instance, the data processing inequality (Equation (1)) turns the ε -closeness of $[\Lambda_k]$ into success probability $\beta - \varepsilon$ on Λ_k , which is the source of the $\beta - \varepsilon$ below. The proofs below therefore only show that a correct oracle answer on Λ_k is turned into a correct answer on Λ . Finally, *efficient* means polynomial time in the input size and $\log(1/\varepsilon)$ for fixed rank r and growing degree d ; every reduction below is a call to $\mathcal{W}$, an oracle call and a multiplication by p , hence efficient by Lemma 3.8.

Theorem 4.2. *Let the rank $r \geq 2$ be fixed, let K be a cyclotomic number field of degree d , and assume GRH. For each problem $\mathcal{P}$ listed below, solving $\mathcal{P}$ with the worst-case parameter on a rank r Hermitian lattice Λ efficiently reduces, for every $\varepsilon \in (0, 1)$ and at the cost of ε in the success probability, to solving $\mathcal{P}$ with the average-case parameter on $\Lambda' \leftarrow \mathcal{D}(\mathcal{H}(\mathcal{G}))$, where $\mathcal{G} = \mathcal{G}(\Lambda)$ is the special genus of Λ and $\mathcal{D}(\mathcal{H}(\mathcal{G}))$ is an efficient representation of the Haar-uniform distribution on $\mathcal{G}$.*

$\mathcal{P}$	<i>worst-case average-case parameter measured against</i>		
<i>HSVP</i>	$p\gamma$	γ	$\det_{\mathbb{Q}}(\Lambda)^{1/n}$
<i>HCVP</i>	$p\gamma$	γ	$\det_{\mathbb{Q}}(\Lambda)^{1/n}$
<i>DSP</i>	$p\gamma$	γ	$\det_{\mathbb{Q}}(\Lambda)^{1/n}$
<i>SVP</i>	$p^2\gamma$	γ	$\lambda_1(\Lambda)$
<i>BDD</i>	α/p^2	α	$\lambda_1(\Lambda)$
Δ <i>BDD</i>	σ	$p\sigma$	<i>Gaussian width</i>
<i>DGS</i>	$p\sigma$	σ	<i>Gaussian width</i>

Proof. Each row is one of the reductions of this section: Lemma 4.7 for HSVP and SVP, Lemma 4.9 for HCVP, Lemma 4.18 for DSP, Lemma 4.12 for BDD, Lemma 4.14 for Δ BDD and Lemma 4.21 for DGS. The value of p and the loss ε are those of Theorem 3.9, as fixed in the conventions above. $\square$

Remark 4.3. Our average-case problems obtain instances from $\mathcal{D}(\mathcal{H}(\mathcal{G}))$ which samples a random class $[A']$ from the genus $\mathcal{G}$ and a random representative $A'' \leftarrow \mathcal{D}([A'])$. As the latter distribution can always be sampled from (see Definition 2.15), along with an isometry $f : A' \rightarrow A''$, we can always assume without loss of generality that any average-case oracle applies this rerandomization internally. For example, given an average-case oracle for BDD, it samples (A'', f) and applies the oracle to $(A'', f(t))$ for a target t , and returns $f^{-1}(c)$ for the solution c . Therefore, to avoid needless repetitions, we refrain from making this rerandomization explicit in the reductions below. If one wants to replace $\mathcal{D}$ by $\mathcal{D}_s$ for some parameter $s > 0$, one can use that the neighbouring procedure always returns a representative which is a factor at most p worse than the starting lattice, which allows one to sample at $s = O(p) \cdot s_0$, where s_0 was the parameter of the starting instance.

Remark 4.4. As already pointed out in Remark 3.7, Theorem 4.2 can also be made valid for CM fields for which there is an efficient way of sampling prime numbers p of polynomial size whose prime ideal divisors in $\mathcal{O}_F$ split in $\mathcal{O}_K$. Also note that p depends on $\det_{\mathbb{Q}}(\Lambda)$ which implicitly introduces a discriminant factor $|\Delta_K|^{O(1)/d}$ into the approximation loss.

4.1 (Hermite) SVP

Definition 4.5 (γ -SVP). *The γ -approximate Shortest Vector Problem is the problem that, given an input lattice Λ , asks to find a non-zero vector $v \in \Lambda$ that satisfies $\|v\| \leq \gamma \lambda_1(\Lambda)$.*

Definition 4.6 (γ -HSVP). *The γ -approximate Hermite Shortest Vector Problem is the problem that, given an n -dimensional input lattice Λ , asks to find a non-zero vector $v \in \Lambda$ that satisfies $\|v\| \leq \gamma \det_{\mathbb{Q}}(\Lambda)^{1/n}$.*

Reduction algorithm from γ' -(H)SVP to γ -ac(H)SVP $^{\mathcal{G}}$

- $\Lambda_0 = \Lambda$

- $\Lambda_k, p \leftarrow \mathcal{W}(\Lambda_0, \varepsilon)$
- $v \leftarrow \mathbb{O}_{\gamma\text{-ac(H)SVP}^g}(\Lambda_k)$.
- Return $p \cdot v$.

Lemma 4.7. *The above reduction algorithm solves γ' -SVP (resp. γ' -HSVP) with probability $\beta - \varepsilon$ with $\gamma' = p^2 \cdot \gamma$ (resp. $\gamma' = p\gamma$), if the oracle $\mathbb{O}_{\gamma\text{-ac(H)SVP}^g}$ succeeds with probability $\beta \in [0, 1]$.*

Proof. We distinguish the cases of Hermite SVP and ordinary SVP, where we start with the latter. The output v of the oracle satisfies $\|v\| \leq \gamma \lambda_1(\Lambda_k)$. Hence, we have

$$\|pv\| \leq \gamma \lambda_1(p\Lambda_k) = p\gamma \lambda_1(\Lambda_k) \leq p^2 \gamma \lambda_1(\Lambda),$$

using $\lambda_1(\Lambda_k) \leq p\lambda_1(\Lambda)$ from Lemma 4.1(iii). We use $p\Lambda_k \subseteq \Lambda$ (Lemma 4.1(ii)) to show that $pv \in \Lambda$, since $v \in \Lambda_k$. So we conclude that pv is indeed a solution to γ' -SVP on Λ with $\gamma' = p^2 \cdot \gamma$.

For the Hermite case, the output v of the oracle satisfies $\|v\| \leq \gamma \det_{\mathbb{Q}}(\Lambda_k)^{1/n} = \gamma \det_{\mathbb{Q}}(\Lambda)^{1/n}$ by Lemma 4.1(i), and $pv \in p\Lambda_k \subseteq \Lambda$. Hence, $\|pv\| \leq p\gamma \det_{\mathbb{Q}}(\Lambda)^{1/n}$ is a solution to γ' -HSVP for $\gamma' = p\gamma$. $\square$

4.2 Hermite CVP

Definition 4.8 (γ -HCVP). *The γ -approximate Hermite Closest Vector Problem is the problem that, given an n -dimensional input lattice Λ and a target $t \in \text{span}(\Lambda)$, asks to find a vector $v \in \Lambda$ that satisfies $\|v - t\| \leq \gamma \det_{\mathbb{Q}}(\Lambda)^{1/n}$.*

Reduction algorithm from γ' -HCVP to $\gamma\text{-acHCVP}^g$

- $\Lambda_0 = \Lambda$, target t .
- $\Lambda_k, p \leftarrow \mathcal{W}(\Lambda_0, \varepsilon)$
- $v \leftarrow \mathbb{O}_{\gamma\text{-ac-HCVP}^g}(\Lambda_k, t/p)$
- Return $p \cdot v$.

Lemma 4.9. *The above reduction algorithm solves γ' -HCVP with probability $\beta - \varepsilon$ with $\gamma' = p\gamma$, if the oracle $\mathbb{O}_{\gamma\text{-acHCVP}^g}$ succeeds with probability $\beta \in [0, 1]$.*

Proof. The output v of the oracle satisfies $\|v - t/p\| \leq \gamma \det_{\mathbb{Q}}(\Lambda_k)^{1/n}$. Hence $\|pv - t\| \leq p\gamma \det_{\mathbb{Q}}(\Lambda_k)^{1/n} = p\gamma \det_{\mathbb{Q}}(\Lambda)^{1/n}$ by Lemma 4.1(i), while $pv \in p\Lambda_k \subseteq \Lambda$ by Lemma 4.1(ii), so that $p \cdot v \in \Lambda$ is a solution to γ' -HCVP for $\gamma' = p\gamma$. $\square$

Remark 4.10. We do not obtain a reduction for approximate CVP relative to $\text{dist}(t, \Lambda)$, as we either cannot always map our solution back or we lose the bound on the distance. For example, mapping $t \mapsto \frac{t-z}{p}$ for some $z \in \Lambda$, and returning $py + z$, always maps a solution back, but the quality is then measured relative to $\text{dist}(t, z + p\Lambda_k) \geq \text{dist}(t, \Lambda)$, for which we have no better bound than the covering radius of $p\Lambda_k$. Mapping $t \mapsto pt$ instead does preserve the promise, as $\text{dist}(pt, \Lambda_k) \leq p \text{dist}(t, \Lambda)$ by Lemma 4.1(iv), but a solution $w \in \Lambda_k$ only satisfies $w/p \in \frac{1}{p}\Lambda_k \not\subseteq \Lambda$. This problem disappears as soon as the solution is unique, which is the regime of Lemma 4.12.

4.3 Bounded Distance Decoding

Definition 4.11 (α -BDD). Let $\alpha \in (0, 1/2)$. The α -Bounded Distance Decoding problem is the problem that, given an input lattice Λ and a target $t \in \text{span}(\Lambda)$ with the promise that $\text{dist}(t, \Lambda) \leq \alpha \lambda_1(\Lambda)$, asks to find the (unique) $v \in \Lambda$ with $\|v - t\| \leq \alpha \lambda_1(\Lambda)$.

Reduction algorithm from α' -BDD to α -acBDD $^{\mathcal{G}}$

- $\Lambda_0 = \Lambda$, target t .
- $\Lambda_k, p \leftarrow \mathcal{W}(\Lambda_0, \varepsilon)$
- $w \leftarrow \mathbb{O}_{\alpha\text{-acBDD}^{\mathcal{G}}}(\Lambda_k, p \cdot t)$
- Return w/p .

Lemma 4.12. Let Λ be a lattice with special genus $\mathcal{G}$. The above reduction algorithm efficiently solves α' -BDD on Λ with probability $\beta - \varepsilon$ with $\alpha' = \alpha/p^2$, if the oracle $\mathbb{O}_{\alpha\text{-acBDD}^{\mathcal{G}}}$ succeeds with probability $\beta \in [0, 1]$.

Proof. Let $v \in \Lambda$ be the solution of the α' -BDD instance, so $\|v - t\| \leq \alpha' \lambda_1(\Lambda)$. Since $p\Lambda \subseteq \Lambda_k$ we have $pv \in \Lambda_k$, and $\|pv - pt\| \leq p\alpha' \lambda_1(\Lambda) \leq p^2 \alpha' \lambda_1(\Lambda_k) = \alpha \lambda_1(\Lambda_k)$, using $\lambda_1(\Lambda) \leq p \lambda_1(\Lambda_k)$ from Lemma 4.1(iii). So (Λ_k, pt) is a valid α -BDD instance, and as $\alpha < 1/2$ its solution is unique, so a correct oracle output is $w = pv$. We conclude that $w/p = v \in \Lambda$ is a solution to the original instance. $\square$

4.4 Distinguishing BDD

We now consider a distinguishing version Δ BDD of the Bounded Distance Decoding problem which can be seen as a natural analogue of decision-LWE. In Section 5 we will demonstrate how the IND-CPA security of a simple KEM can be reduced to a Δ LIP and Δ BDD problem.

Definition 4.13 (σ - Δ BDD). Let $\sigma > 0$. The σ -Distinguishing BDD problem is the problem that, given an input lattice Λ and a sample $t \in \text{span}(\Lambda)/\Lambda$, asks to distinguish between the cases $t \leftarrow \mathcal{N}_\sigma \bmod \Lambda$ and $t \leftarrow \mathcal{U}(\text{span}(\Lambda)/\Lambda)$, where $\mathcal{N}_\sigma$ denotes the continuous Gaussian of width σ on $\text{span}(\Lambda)$.

Reduction algorithm from σ - Δ BDD to $p\sigma$ -ac Δ BDD $^{\mathcal{G}}$

- $\Lambda_0 = \Lambda$, sample t .
- $\Lambda_k, p \leftarrow \mathcal{W}(\Lambda_0, \varepsilon)$
- Return $\mathbb{O}_{p\sigma\text{-ac}\Delta\text{BDD}^{\mathcal{G}}}(\Lambda_k, p \cdot t \bmod \Lambda_k)$.

Lemma 4.14. The above reduction algorithm solves σ - Δ BDD with advantage $\beta - \varepsilon$, if the oracle $\mathbb{O}_{p\sigma\text{-ac}\Delta\text{BDD}^{\mathcal{G}}}$ has advantage β .

Proof. It suffices to show that both cases are mapped exactly to the corresponding case for Λ_k . If $t = e + \ell$ with $e \leftarrow \mathcal{N}_\sigma$ and $\ell \in \Lambda$, then $pt = pe + p\ell$ with $p\ell \in p\Lambda \subseteq \Lambda_k$, so $pt \equiv pe \bmod \Lambda_k$ with $pe \leftarrow \mathcal{N}_{p\sigma}$. If $t \leftarrow \mathcal{U}(\text{span}(\Lambda)/\Lambda)$, then $pt \leftarrow \mathcal{U}(\text{span}(\Lambda)/p\Lambda)$. Moreover, reducing modulo $\Lambda_k \supseteq p\Lambda$ is a surjective homomorphism, so $pt \bmod \Lambda_k$ is uniform on $\text{span}(\Lambda)/\Lambda_k$. Since the walk is independent of t , the advantage transfers up to the statistical distance ε of $[\Lambda_k]$ to $\mathcal{H}(\mathcal{G})$. $\square$

Remark 4.15. The gap variant of BDD does not follow from such a reduction. The ‘close’ case is preserved as in the BDD reduction above, but we cannot give any guarantees on the ‘far’ case, which after multiplication by p could be arbitrarily close to the neighbouring lattice Λ_k .

Remark 4.16. We stated the Δ BDD problem with a target living in the continuous ambient space $\text{span}(\Lambda)$. We can also adapt the reduction to a discrete ambient space like $\frac{1}{Q}\Lambda$ for an integer $Q > 1$, with $t \in \frac{1}{Q}\Lambda/\Lambda$, discrete Gaussian error $e \leftarrow \mathcal{N}_\sigma(\frac{1}{Q}\Lambda)$, uniform case $t \leftarrow \mathcal{U}(\frac{1}{Q}\Lambda/\Lambda)$, and the average-case problem stated over $\frac{1}{Q}\Lambda_k$, provided that

1. p is sampled coprime to $Q \det_{\mathbb{Q}}(\Lambda)^2$ in Algorithm 2, and
2. the reduction calls the oracle on $pt + f \bmod \Lambda_k$ for a fresh $f \leftarrow \mathcal{N}_s(\frac{1}{Q}\Lambda_k)$ with $s \geq \frac{p}{Q}\eta_\delta(\Lambda)$.

It then reduces σ - Δ BDD to σ' -ac Δ BDD^G with $\sigma' = \sqrt{p^2\sigma^2 + s^2}$, at the cost of an additional $O(\delta)$ loss in the advantage. In particular, for $Q \geq \eta_\delta(\Lambda)/\sigma$ one can take $s = p\sigma$ and $\sigma' = \sqrt{2}p\sigma$.

4.5 Dense Submodule Problem (DSP)

Definition 4.17 (γ -DSP). *The γ -approximate Dense Submodule Problem is the problem that, given an n -dimensional⁷ input module lattice Λ and submodule rank $\ell \in \{1, \dots, r\}$, asks to find a rank- ℓ submodule M satisfying $\det_{\mathbb{Q}}(M)^{1/(d\ell)} \leq \gamma \det_{\mathbb{Q}}(\Lambda)^{1/n}$.*

Reduction algorithm from γ' -DSP to γ -acDSP^G

- $\Lambda_0 = \Lambda$, submodule rank ℓ .
- $\Lambda_k, p \leftarrow \mathcal{W}(\Lambda_0, \varepsilon)$
- $M \leftarrow \mathbb{O}_{\gamma\text{-acDSP}^G}(\Lambda_k, \ell)$
- Return $p \cdot M$.

Lemma 4.18. *The above reduction algorithm solves γ' -DSP with probability $\beta - \varepsilon$ with $\gamma' = p\gamma$, if the oracle $\mathbb{O}_{\gamma\text{-acDSP}^G}$ succeeds with probability $\beta \in [0, 1]$.*

Proof. The output M of the oracle is a rank- ℓ submodule of Λ_k satisfying $\det_{\mathbb{Q}}(M)^{1/(d\ell)} \leq \gamma \det_{\mathbb{Q}}(\Lambda_k)^{1/n} = \gamma \det_{\mathbb{Q}}(\Lambda)^{1/n}$, using Lemma 4.1(i). Since $p\Lambda_k \subseteq \Lambda$ (Lemma 4.1(ii)) the scaled module pM is a rank- ℓ submodule of Λ , and $\det_{\mathbb{Q}}(pM)^{1/(d\ell)} = p \det_{\mathbb{Q}}(M)^{1/(d\ell)} \leq p\gamma \det_{\mathbb{Q}}(\Lambda)^{1/n}$, so it is a solution to γ' -DSP for $\gamma' = p\gamma$. $\square$

Remark 4.19. Alternatively we can measure the density of a rank- ℓ submodule M algebraically by $N(M) \stackrel{\text{def}}{=} N_{K/\mathbb{Q}}(\det M)^{1/2}$. However this gives nothing new: since $\det_{\mathbb{Q}}(M)^2 = |\Delta_K|^\ell N(M)^2$, the algebraic variant of DSP coincides with DSP. As a special case $\ell = 1$, we obtain a reduction relative to the algebraic first minimum: $\lambda_1^K(\Lambda) \stackrel{\text{def}}{=} \min_I N(I)^{1/d}$ over rank-1 submodules $I \subseteq \Lambda$, which loses a factor p^2 , exactly as for SVP versus HSVP.

⁷ The module lattice Λ has rank r over a degree d number field, so that $n = rd$.

4.6 Discrete Gaussian Sampling

We conclude with Discrete Gaussian Sampling, the standard sampling primitive underlying many worst-case to average-case reductions. We state it, as is common, for an arbitrary coset $c + \Lambda$; the case $c = 0$ is the usual problem of sampling $\mathcal{N}_\sigma(\Lambda)$ itself.

Definition 4.20 ((ϵ, σ) -DGS). *Let $\epsilon \in (0, 1]$ and $\sigma > 0$. The (ϵ, σ) -Discrete Gaussian Sampling problem is the problem that, given an input lattice Λ and a coset representative $c \in \text{span}(\Lambda)$, asks to output a sample from a distribution that is ϵ -close to $\mathcal{N}_\sigma(c + \Lambda)$ in statistical distance.*

Reduction algorithm from (ϵ_1, σ_1) -DGS to (ϵ_0, σ_0) -acDGS $^\mathcal{G}$

- $\Lambda_0 = \Lambda$, coset representative c .
- $\Lambda_k, p \leftarrow \mathcal{W}(\Lambda_0, \varepsilon)$
- Sample $z \leftarrow \Lambda/p\Lambda_k$ uniformly
- $y \leftarrow \mathbb{O}_{(\epsilon_0, \sigma_0)\text{-acDGS}^\mathcal{G}}(\Lambda_k, \frac{c+z}{p})$
- Return $p \cdot y$.

The reduction is an instance of the general fact that a sampler on a full-rank sublattice $\Lambda' \subseteq \Lambda$ induces one on Λ above $\eta_\delta(\Lambda')$, applied to $\Lambda' = p\Lambda_k$.

Lemma 4.21. *Let $\delta \in (0, 1)$ and $\sigma_1 = p\sigma_0$ with $\sigma_0 \geq \eta_\delta(\Lambda_k)$, for which $\sigma_1 \geq p^2\eta_\delta(\Lambda)$ is sufficient. The above reduction algorithm solves (ϵ_1, σ_1) -DGS with $\epsilon_1 = \epsilon_0 + \varepsilon + 2\delta$, if the oracle $\mathbb{O}_{(\epsilon_0, \sigma_0)\text{-acDGS}^\mathcal{G}}$ solves (ϵ_0, σ_0) -acDGS $^\mathcal{G}$.*

Proof. By Lemma 4.1(i),(ii) the lattice $p\Lambda_k \subseteq \Lambda$ is a sublattice of index $[\Lambda : p\Lambda_k] = p^n$, so a uniform $z \in \Lambda/p\Lambda_k$ can be sampled from bases of Λ and Λ_k . The sufficiency of the condition on σ_1 is Lemma 4.1(v). Since the walk is independent of c , it is enough to show that for a fixed Λ_k the algorithm outputs a sample within statistical distance $\epsilon_0 + 2\delta$ of $\mathcal{N}_{\sigma_1}(c + \Lambda)$. The additional ε is the statistical distance of $[\Lambda_k]$ to $\mathcal{H}(\mathcal{G})$. Now consider the partition into cosets of $p\Lambda_k$:

$$c + \Lambda = \bigsqcup_{z \in \Lambda/p\Lambda_k} (c + z + p\Lambda_k).$$

Sampling $x \leftarrow \mathcal{N}_{\sigma_1}(c + \Lambda)$ is the same as first picking the coset z with probability proportional to $\mathcal{N}_{\sigma_1}(c + z + p\Lambda_k)$ and then sampling $x \leftarrow \mathcal{N}_{\sigma_1}(c + z + p\Lambda_k)$. As $\sigma_1 \geq p\eta_\delta(\Lambda_k) = \eta_\delta(p\Lambda_k)$, Lemma 2.14 shows that the first step is within statistical distance 2δ of the uniform choice made by the algorithm. For the second step, note that $x \mapsto x/p$ is a bijection from $c + z + p\Lambda_k$ onto $\frac{c+z}{p} + \Lambda_k$ with $\exp(-\pi \|x\|^2 / \sigma_1^2) = \exp(-\pi \|x/p\|^2 / \sigma_0^2)$, so

$$\mathcal{N}_{\sigma_1}(c + z + p\Lambda_k) = p \cdot \mathcal{N}_{\sigma_0}\left(\frac{c+z}{p} + \Lambda_k\right)$$

as distributions, the latter being exactly what the oracle is called on, up to ϵ_0 . Finally $p \cdot y \in c + z + p\Lambda_k \subseteq c + \Lambda$, so the output always lies in the correct coset. $\square$

Remark 4.22. The loss is only a factor p on the width σ . Relative to the smoothing parameter it is p^2 , as the hypothesis $\sigma_0 \geq \eta_\delta(\Lambda_k)$ is only guaranteed by $\sigma_1 \geq p^2 \eta_\delta(\Lambda)$.

Remark 4.23. The same random shift gives a reduction for CVP with a *uniformly random* target, the problem underlying the security of GPV-style hash-and-sign signatures, where a forgery is a lattice vector close to the hash of a message: given Λ and $t \leftarrow \mathcal{U}(\text{span}(\Lambda)/\Lambda)$, find $v \in \Lambda$ with $\|v - t\| \leq \gamma \det_{\mathbb{Q}}(\Lambda)^{1/n}$. Calling the oracle on $\frac{t-z}{p}$ and returning $py + z \in z + p\Lambda_k \subseteq \Lambda$ loses a factor $\gamma' = p\gamma$, exactly as for γ -HCVP. Just as for Δ BDD, one can also directly consider a reduction for the discretized version of this problem, but without any smoothing constraint on the modulus $Q \geq 2$.

5 Applications

We now make formal the two-assumption security reduction sketched in the introduction, which replaces the ad-hoc dense-sublattice instantiation of current LIP-based schemes [DvW22, BGPSD23, BZI⁺24, ARLW24, BMM25, BLSD26] by the two independent steps proposed informally in [LRvW25]: (1) the lattice used by the scheme is hard to distinguish from a random lattice of its genus, and (2) the underlying lattice problem is hard for a random lattice of that genus. The first is the state-of-the-art assumption that the (spinor/special) genus is the strongest efficiently computable invariant, while the second is exactly what the reductions of Section 4 turn into worst-case hardness.

We make this explicit for a minimal key encapsulation mechanism Π_0 , a variant of the KEM of [DvW22] in which the lossy dense-sublattice instance is replaced by a random lattice of the genus, and the randomness extractor by a residue modulo $t\Lambda$. Throughout, Λ_0 is a $\mathcal{O}_K$ -lattice of rank r over a number field K of degree d , and we write $n = rd$ and $\mathcal{G} = \text{sgen}(\Lambda_0)$.

Definition 5.1 (s -ac Δ LIP $^{\Lambda_0, \mathcal{G}}$). *Let $s > 0$. The problem s -ac Δ LIP $^{\Lambda_0, \mathcal{G}}$ asks, given a lattice Λ , to distinguish between $\Lambda \leftarrow \mathcal{D}_s([\Lambda_0])$ and $\Lambda \leftarrow \mathcal{D}_s(\mathcal{H}(\mathcal{G}))$.*⁸

This is the analogue of the problem ac- Δ LIP $_s$ of [DvW22, Thm. 5.2], with the lossy class there replaced by a random class of the genus. We work in the regime $s \geq 2^n \max_{\Lambda \in \mathcal{G}} \lambda_n(\Lambda)$ ⁹, in which $\mathcal{D}_s$ is an average-case class distribution for every class of $\mathcal{G}$ (Definition 2.15 and Remark 2.16) and the problem is equivalent to its worst-case variant (cf. [DvW22, Lem. 3.9, 3.12]): by Definition 2.15(ii) one can sample $\mathcal{D}_s$ of a class from any of its representatives, so only the class matters. The second assumption is precisely the problem σ' -ac Δ BDD $^{\mathcal{G}}$ of Section 4 in

⁸ Note that $\Lambda_0 \in \mathcal{G}$, so the two cases overlap; the problem can only be hard if the class $[\Lambda_0]$ has negligible mass in $\mathcal{G}$.

⁹ Note that in our wc-ac reduction any sampled lattice $\Lambda \in \mathcal{G}$ satisfies $\lambda_n(\Lambda) \leq p \cdot \lambda_n([\Lambda_0])$ by construction. Alternatively this can be replaced by concentration bounds such as from [vW24].

its discrete form (Remark 4.16) with ambient space $\frac{1}{Q_t}\Lambda$: given $\Lambda \leftarrow \mathcal{D}(\mathcal{H}(\mathcal{G}))$, distinguish $\mathcal{N}_{\sigma'}(\frac{1}{Q_t}\Lambda) \bmod \Lambda$ from $\mathcal{U}(\frac{1}{Q_t}\Lambda/\Lambda)$.

The scheme Π_0 . In contrast to [DvW22], we state the scheme here in its lattice version instead of in terms of quadratic forms to follow the notation in the rest of this work: the public key is the lattice Λ itself in the representation output by $\mathcal{D}_s$, and Encaps is efficient by Definition 2.15(iii); nothing in the statements depends on the choice of representation. For an efficient instantiation one preferably uses integral quadratic forms as in [DvW22]. Fix an integral lattice Λ_0 with an efficient decoder $\text{Decode}(\Lambda_0, \cdot)$ of radius $\rho < \lambda_1(\Lambda_0)/2$, a key modulus $t \geq 2$, and a parameter $s \geq 2^n \max_{\Lambda \in \mathcal{G}} \lambda_n(\Lambda)$, and set

$$\sigma \stackrel{\text{def}}{=} s\sqrt{n \ln(2n+4)/\pi}, \quad Q \stackrel{\text{def}}{=} \lceil \sigma\sqrt{n}/\rho \rceil + 1, \quad \sigma' \stackrel{\text{def}}{=} \sigma/(Qt),$$

where Q is adjusted so that $\gcd(Q, t) = 1$.

- **KeyGen.** Sample $\Lambda \leftarrow \mathcal{D}_s([\Lambda_0])$ together with the isometry $O \in U(V)$ with $O\Lambda = \Lambda_0$; output $pk \stackrel{\text{def}}{=} \Lambda$ and $sk \stackrel{\text{def}}{=} O$.
- **Encaps(Λ).** Sample $x^* \leftarrow \mathcal{N}_\sigma(\Lambda)$ and output the ciphertext $c \stackrel{\text{def}}{=} x^*/Q \bmod \Lambda \in \frac{1}{Q}\Lambda/\Lambda$ together with the key $k \stackrel{\text{def}}{=} x^* \bmod t\Lambda \in \Lambda/t\Lambda$.
- **Decaps(O, c).** Lift c to $\tilde{x} \in \frac{1}{Q}\Lambda$, compute $w \stackrel{\text{def}}{=} \text{Decode}(\Lambda_0, O\tilde{x})$ and $x^* \stackrel{\text{def}}{=} Q(\tilde{x} - O^{-1}w) \in \Lambda$, and output $x^* \bmod t\Lambda$.

Lemma 5.2 (Correctness). $\mathbb{P}[\text{Decaps}(O, c) \neq k] \leq (300n + 1)2^{-n}$ over the randomness of *KeyGen* and *Encaps*.

Proof. By Definition 2.15(iii) and Lemma 2.13, except with probability $(300n + 1)2^{-n}$ Encaps samples exactly from $\mathcal{N}_\sigma(\Lambda)$ and we have $\|x^*\| \leq \sigma\sqrt{n} < Q\rho$. Any lift $\tilde{x}$ of c satisfies $\tilde{x} = x^*/Q + z$ with $z \in \Lambda$ and $\|x^*/Q\| < \rho < \lambda_1(\Lambda)/2$, so z is the unique lattice point within distance ρ of $\tilde{x}$. As O is an isometry from Λ onto Λ_0 the decoder returns $w = Oz$, hence $x^* = Q(\tilde{x} - z)$ is recovered exactly, and so is k . $\square$

Since $\gcd(Q, t) = 1$, the Chinese remainder theorem gives $\Lambda/Qt\Lambda \cong \Lambda/Q\Lambda \times \Lambda/t\Lambda$, so the pair (c, k) is nothing but the single point $x^*/(Qt) \bmod \Lambda$ of $\frac{1}{Qt}\Lambda/\Lambda$, and both directions of this correspondence are computable from Λ alone. In other words, the statement that the key is pseudorandom given the ciphertext is the problem σ' -ac Δ BDD $^{\mathcal{G}}$.

Theorem 5.3. Write $\text{Adv}^{\text{cpa}}(\mathcal{A}) \stackrel{\text{def}}{=} |\mathbb{P}[\mathcal{A}(pk, c, k) = 1] - \mathbb{P}[\mathcal{A}(pk, c, u) = 1]|$ for $u \leftarrow \mathcal{U}(\Lambda/t\Lambda)$. For every PPT adversary $\mathcal{A}$ against the IND-CPA security of Π_0 there are PPT distinguishers $\mathcal{B}_1, \mathcal{B}_2$ against s -ac Δ LIP $^{\Lambda_0, \mathcal{G}}$ and σ' -ac Δ BDD $^{\mathcal{G}}$, of essentially the same running time, with

$$\text{Adv}^{\text{cpa}}(\mathcal{A}) \leq 2 \text{Adv}^{\Delta\text{LIP}}(\mathcal{B}_1) + 2 \text{Adv}^{\Delta\text{BDD}}(\mathcal{B}_2) + 600n 2^{-n}.$$

Proof. Replacing $\Lambda \leftarrow \mathcal{D}_s([A_0])$ by $\Lambda \leftarrow \mathcal{D}_s(\mathcal{H}(\mathcal{G}))$ costs $\text{Adv}^{\Delta\text{LIP}}(\mathcal{B}_1)$ in each of the two probability terms, so we may assume $\Lambda \leftarrow \mathcal{D}_s(\mathcal{H}(\mathcal{G}))$; we must then bound the distance between (Λ, c, k) and (Λ, c, u) . By Definition 2.15(iii), except with probability $300n2^{-n}$ Encaps samples exactly from $\mathcal{N}_\sigma(\Lambda)$; replacing Encaps by this ideal distribution in both tuples costs $600n2^{-n}$ in total. On input a lattice Λ_1 and a sample $y \in \frac{1}{Q_t}\Lambda_1/\Lambda_1$, the distinguisher $\mathcal{B}_2$ samples $\Lambda \leftarrow \mathcal{D}_s([A_1])$ itself together with an isometry $\Lambda_1 \rightarrow \Lambda$ (Definition 2.15(ii)), maps y to $\frac{1}{Q_t}\Lambda/\Lambda$, splits it by the above correspondence into a pair (c, k) , and runs $\mathcal{A}$ on (Λ, c, k) . For $y \leftarrow \mathcal{N}_{\sigma'}(\frac{1}{Q_t}\Lambda_1) \bmod \Lambda_1$ this is exactly the ideal distribution above; for y uniform both components are uniform and independent. Hence (Λ, c, k) is within $\text{Adv}^{\Delta\text{BDD}}(\mathcal{B}_2)$ of (Λ, u_1, u) with u_1 uniform. Projecting $\frac{1}{Q_t}\Lambda/\Lambda$ onto $\frac{1}{Q}\Lambda/\Lambda$ maps both cases of σ' -ac $\Delta\text{BDD}^\mathcal{G}$ to the corresponding cases of $\frac{\sigma}{Q}$ -ac $\Delta\text{BDD}^\mathcal{G}$, so the latter is at least as hard, and appending an independent uniform u shows that (Λ, u_1, u) is within a further $\text{Adv}^{\Delta\text{BDD}}(\mathcal{B}_2)$ of (Λ, c, u) . $\square$

Remark 5.4. The theorem only asserts that k is pseudorandom *given* c , which is all that IND-CPA requires. In our regime the key is in addition statistically close to uniform on $\Lambda/t\Lambda \cong (\mathbb{Z}/t\mathbb{Z})^n$ for any $t = \text{poly}(n)$: since $\eta_\epsilon(\Lambda) \leq \lambda_n(\Lambda)\sqrt{\ln(2n(1+1/\epsilon))}/\pi$ [MR07, Lem. 3.3], we have $\sigma \geq t\eta_\epsilon(\Lambda) = \eta_\epsilon(t\Lambda)$ already for $\epsilon = 2^{-n}$, so all coset masses $\mathcal{N}_\sigma(v + t\Lambda)$ agree up to a factor $\frac{1+\epsilon}{1-\epsilon}$ by Lemma 2.14 and k is within 2ϵ of uniform.

Corollary 5.5. *Suppose that in Algorithm 2 the prime p is additionally sampled coprime to Qt . Then the IND-CPA security of Π_0 follows from s -ac $\Delta\text{LIP}^{\Lambda_0, \mathcal{G}}$ together with the worst-case hardness of $\frac{\sigma}{\sqrt{2p}Qt}$ - ΔBDD on every lattice of $\mathcal{G}$, up to an additional loss of $O(\delta) + 2\epsilon + 600n2^{-n}$ with $\delta = 2^{-n}$ and ϵ the statistical distance of Theorem 3.9.*

Proof. Apply the reduction of Remark 4.16 with ambient space $\frac{1}{Q_t}\Lambda$ and worst-case parameter $\frac{\sigma}{\sqrt{2p}Qt}$, so that the average-case parameter is $\sqrt{2}p \cdot \frac{\sigma}{\sqrt{2p}Qt} = \sigma'$. Its conditions hold: p is coprime to $Qt \det_{\mathbb{Q}}(\Lambda)^2$, and the flooding condition $\sigma \geq \sqrt{2}p\eta_\delta(\Lambda)$ is met for every $\Lambda \in \mathcal{G}$ as $\sigma \geq 2^n\lambda_n(\Lambda)\sqrt{n\ln(2n+4)}/\pi$ while $\eta_\delta(\Lambda) \leq \lambda_n(\Lambda)\sqrt{\ln(2n(1+1/\delta))}/\pi$ [MR07, Lem. 3.3] and $p = \text{poly}(n)$. Combining with Theorem 5.3 gives the claim. $\square$

Remark 5.6. The exponentially large s is chosen so that both assumptions match the problems of Section 4 exactly. It makes Q , and with it the ciphertext coefficients, exponentially large as well, so Π_0 is a theoretical (polynomial-time) example rather than a practical scheme. The worst-case problem of Corollary 5.5 is unaffected by this: since $\sigma\sqrt{n}/\rho \leq Q \leq 2\sigma\sqrt{n}/\rho$, its parameter satisfies $\frac{\sigma}{\sqrt{2p}Qt} = \Theta(\frac{\rho}{pt\sqrt{n}})$ independently of s and Q . In particular, by Lemma 2.13 the worst-case error has norm within a $\text{poly}(n)$ factor $O(pt)$ of the decoding radius ρ . The exponential parameters only enlarge the ambient space $\frac{1}{Q_t}\Lambda/\Lambda$. A polynomial s , and hence $Q = \text{poly}(n)$, is possible by tracking the distribution $\mathcal{D}_s$ through the reductions and Kneser neighbouring steps as in [DvW22]. The assumptions

and problem statements are then stated relative to $\mathcal{D}_s$ and one should expect a factor $p = \text{poly}(n)$ loss on s in the worst-case to average-case reductions.

References

- AC25. Bill Allombert and Gaëtan Chenevier. Unimodular hunting. II. *Forum Math. Sigma*, 13:18, 2025. Id/No e136. doi:[10.1017/fms.2025.10058](https://doi.org/10.1017/fms.2025.10058).
- Ajt96. Miklós Ajtai. Generating hard instances of lattice problems. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 99–108, 1996.
- AKS04. Manindra Agrawal, Neeraj Kayal, and Nitin Saxena. PRIMES is in P. *Annals of Mathematics*, 160(2):781–793, 2004. URL: <http://www.jstor.org/stable/3597229>.
- APMvW25. Bill Allombert, Alice Pellet-Mary, and Wessel van Woerden. Cryptanalysis of rank-2 module-LIP: a single real embedding is all it takes. In *EUROCRYPT 2025, 44th Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 184–212. Springer, 2025.
- ARLW24. Léo Ackermann, Adeline Roux-Langlois, and Alexandre Wallet. Public-key encryption from the lattice isomorphism problem. In *WCC 2024: The Thirteenth International Workshop on Coding and Cryptography*, 2024. Extended abstract.
- Ban93. Wojciech Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 296(1):625–635, 1993.
- BB13. Valentin Blomer and Farrell Brumley. The role of the Ramanujan conjecture in analytic number theory. *Bull. Am. Math. Soc., New Ser.*, 50(2):267–320, 2013. doi:[10.1090/S0273-0979-2013-01404-6](https://doi.org/10.1090/S0273-0979-2013-01404-6).
- BDG23. Peter Bruin, Léo Ducas, and Shane Gibbons. Genus distribution of random q -ary lattices. *Banach Center Publications*, 126:137–159, 2023.
- dBdPMW20. Koen de Boer, Léo Ducas, Alice Pellet-Mary, and Benjamin Wesolowski. Random self-reducibility of ideal-SVP via Arakelov random walks. In *CRYPTO*, 2020.
- BF12. Jean-François Biasse and Claus Fieker. A polynomial time algorithm for computing the HNF of a module over the integers of a number field. In *Proceedings of the 37th International Symposium on Symbolic and Algebraic Computation*, ISSAC ’12, page 75–82, New York, NY, USA, 2012. Association for Computing Machinery. doi:[10.1145/2442829.2442844](https://doi.org/10.1145/2442829.2442844).
- 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 *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 252–281. Springer, 2023.
- BLP⁺13. Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of learning with errors. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 575–584, 2013.
- BLSD26. Huck Bennett, Zhengnan Lai, and Noah Stephens-Davidowitz. Advanced cryptography from lattice isomorphism—new constructions of IBE and FHE. In *Annual International Cryptology Conference*, pages 446–476. Springer, 2026.
- BMM25. Pedro Branco, Giulio Malavolta, and Zayd Maradni. Fully-homomorphic encryption from lattice isomorphism. In *Theory of Cryptography Conference*, pages 220–252. Springer, 2025.

- dBPTW26. Koen de Boer, Aurel Page, Radu Toma, and Benjamin Wesolowski. Average hardness of SIVP for module lattices of fixed rank. In Aditya Bhaskara and Artur Czumaj, editors, *Proceedings of the 58th Annual ACM Symposium on Theory of Computing, STOC 2026, Salt Lake City, UT, USA, June 22–26, 2026*, pages 814–822. ACM, 2026. doi:10.1145/3798129.3800796.
- BW25. Huck Bennett and Kaung Myat Htay Win. Relating code equivalence to other isomorphism problems. *Designs, Codes and Cryptography*, 93(3):701–723, 2025.
- BZI⁺24. Gustavo C. Biage, Gustavo Zambonin, Thaís B. Idalino, Daniel Panario, and Ricardo Custódio. A concrete LIP-based KEM with simple lattices. *IEEE Access*, 2024.
- Car86. Henri Carayol. On the bad reduction of Shimura curves. *Compos. Math.*, 59:151–230, 1986. URL: <https://eudml.org/doc/89787>.
- Cay46. A. Cayley. Sur quelques propriétés des déterminants gauches. *J. Reine Angew. Math.*, 32:119–123, 1846. URL: <https://eudml.org/doc/147336>, doi:10.1515/crll.1846.32.119.
- Che51. Claude Chevalley. Deux théorèmes d’arithmétique. *J. Math. Soc. Japan*, 3:36–44, 1951. doi:10.2969/jmsj/00310036.
- Cho16. Sungmun Cho. Group schemes and local densities of ramified Hermitian lattices in residue characteristic 2, Part I. *Algebra & Number Theory*, 10(3):451–532, 2016. doi:10.2140/ant.2016.10.451.
- CM25. Clémence Chevnard and Guilhem Mureau. Ideally HAWKward: how not to break module-LIP. *Cryptology ePrint Archive*, 2025. URL: <https://eprint.iacr.org/2025/1095>.
- CME⁺25. Clémence Chevnard, Guilhem Mureau, Thomas Espitau, Alice Pellet-Mary, Heorhii Pliatsok, and Alexandre Wallet. A reduction from Hawk to the principal ideal problem in a quaternion algebra. In Serge Fehr and Pierre-Alain Fouque, editors, *Advances in Cryptology - EUROCRYPT 2025 - 44th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Madrid, Spain, May 4–8, 2025, Proceedings, Part II*, volume 15602 of *Lecture Notes in Computer Science*, pages 154–183. Springer, 2025. doi:10.1007/978-3-031-91124-8_6.
- Coh93. Henri Cohen. *A Course in Computational Algebraic Number Theory*. Springer, 1993.
- Coh99. Henri Cohen. *Advanced Topics in Computational Number Theory*. Graduate Texts in Mathematics. Springer New York, 1999.
- COU01. Laurent Clozel, Hee Oh, and Emmanuel Ullmo. Hecke operators and equidistribution of Hecke points. *Invent. Math.*, 144(2):327–351, 2001. doi:10.1007/s002220100126.
- CST08. Tullio Ceccherini-Silberstein, Fabio Scarabotti, and Filippo Tolli. *Harmonic Analysis on Finite Groups: Representation Theory, Gelfand Pairs and Markov Chains*, volume 108 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, 2008. doi:10.1017/CB09780511619823.
- Del73. Pierre Deligne. La conjecture de Weil. I. *Publ. Math., Inst. Hautes Étud. Sci.*, 43:273–307, 1973. URL: <https://eudml.org/doc/103930>, doi:10.1007/BF02684373.
- DG23. Léo Ducas and Shane Gibbons. Hull attacks on the lattice isomorphism problem. In *PKC 2023, 26th 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 *Advances in Cryptology - ASIACRYPT 2022, Part IV*, volume 13794 of *Lecture Notes in Computer Science*, pages 65–94. Springer, 2022. URL: <https://eprint.iacr.org/2022/1155>, doi: [10.1007/978-3-031-22972-5_3](https://doi.org/10.1007/978-3-031-22972-5_3).
- DvW22. Léo Ducas and Wessel van Woerden. On the lattice isomorphism problem, quadratic forms, remarkable lattices, and cryptography. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 643–673. Springer, 2022.
- vG25. Daniël MH van Gent. A note on the genus of the HAWK lattice. *Cryptology ePrint Archive*, 2025. URL: <https://eprint.iacr.org/2025/215>.
- vGP25. Daniël MH van Gent and Ludo N. Pulles. HAWK: having automorphisms weakens key. *IACR Communications in Cryptology*, 2(2):20, 2025.
- GPV08. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 197–206, 2008.
- GV14. Matthew Greenberg and John Voight. Lattice methods for algebraic modular forms on classical groups. In *Computations with modular forms*, volume 6 of *Contrib. Math. Comput. Sci.*, pages 147–179. Springer, Cham, 2014.
- vGvW25. Daniël van Gent and Wessel van Woerden. A search to distinguish reduction for the isomorphism problem on direct sum lattices. In *ASIACRYPT 2025, 31st International Conference on the Theory and Application of Cryptology and Information Security*, pages 3–33. Springer, 2025.
- 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.
- JL70. H. Jacquet and R. P. Langlands. *Automorphic forms on GL(2)*, volume 114 of *Lect. Notes Math.* Springer, Cham, 1970. doi:[10.1007/BFb0058988](https://doi.org/10.1007/BFb0058988).
- JMV09. David Jao, Stephen D Miller, and Ramarathnam Venkatesan. Expander graphs based on GRH with an application to elliptic curve cryptography. *Journal of Number Theory*, 129(6):1491–1504, 2009.
- JW15. Dimitar Jetchev and Benjamin Wesolowski. On graphs of isogenies of principally polarizable abelian surfaces and the discrete logarithm problem. *CoRR*, abs/1506.00522, 2015. URL: <http://arxiv.org/abs/1506.00522>, arXiv:1506.00522.
- Kir16. Markus Kirschmer. Definite quadratic and Hermitian forms with small class number. Habilitationsschrift, RWTH Aachen University, 2016.
- KMRT98. Max-Albert Knus, Alexander Merkurjev, Markus Rost, and Jean-Pierre Tignol. *The book of involutions. With a preface by J. Tits*, volume 44 of *Colloq. Publ., Am. Math. Soc.* Providence, RI: American Mathematical Society, 1998.
- Kne57. Martin Kneser. Klassenzahlen definiter quadratischer Formen. *Arch. Math.*, 8:241–250, 1957.
- KV10. Markus Kirschmer and John Voight. Algorithmic enumeration of ideal classes for quaternion orders. *SIAM J. Comput.*, 39(5):1714–1747, 2010. doi:[10.1137/080734467](https://doi.org/10.1137/080734467).

- Liv01. R. Livné. Communication networks and Hilbert modular forms. In *Applications of algebraic geometry to coding theory, physics and computation. Proceedings of the NATO advanced research workshop, Eilat, Israel, February 25–March 1, 2001*, pages 255–270. Dordrecht: Kluwer Academic Publishers, 2001.
- LJPW24. Hengyi Luo, Kaijie Jiang, Yanbin Pan, and Anyu Wang. Cryptanalysis of rank-2 module-LIP with symplectic automorphisms. In *ASIACRYPT 2024, 30th International Conference on the Theory and Application of Cryptology and Information Security*, pages 359–385. Springer, 2024.
- LJS19. Hendrik W Lenstra Jr and Alice Silverberg. Testing isomorphism of lattices over CM-orders. *SIAM Journal on Computing*, 48(4):1300–1334, 2019.
- LLM24. Cong Ling, Jingbo Liu, and Andrew Mendelsohn. On the spinor genus and the distinguishing lattice isomorphism problem. In *ASIACRYPT 2024, 30th International Conference on the Theory and Application of Cryptology and Information Security*, pages 329–358. Springer, 2024.
- Lou00. Stéphane Louboutin. Explicit bounds for residues of Dedekind zeta functions, values of L -functions at $s = 1$, and relative class numbers. *Journal of Number Theory*, 85(2):263–282, 2000.
- LRvW25. Alberto Leporati, Lorenzo Rovida, and Wessel van Woerden. Beyond LWE: a lattice framework for homomorphic encryption. *Cryptology ePrint Archive*, 2025.
- MPM26. Guilhem Mureau and Alice Pellet-Mary. Dimension reduction for SVP in Hawk: A trace-zero approach. *Cryptology ePrint Archive*, Paper 2026/1560, 2026. URL: <https://eprint.iacr.org/2026/1560>.
- MPPW24. Guilhem Mureau, Alice Pellet-Mary, Heorhii Pliatsok, and Alexandre Wallet. Cryptanalysis of rank-2 module-LIP in totally real number fields. In *Advances in Cryptology - EUROCRYPT 2024, Lecture Notes in Computer Science*. Springer, 2024. URL: <https://eprint.iacr.org/2024/441>.
- MR07. Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on Gaussian measures. *SIAM journal on computing*, 37(1):267–302, 2007.
- Mur25. Guilhem Mureau. Special genera of Hermitian lattices and applications to HAWK. In *Theory of Cryptography Conference, TCC 2025*, Lecture Notes in Computer Science. Springer Verlag, December 2025. URL: <https://eprint.iacr.org/2025/940>.
- MV06. Hugh L. Montgomery and Robert C. Vaughan. *Multiplicative Number Theory I: Classical Theory*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 2006.
- NS13. J. Neukirch and N. Schappacher. *Algebraic Number Theory*. Grundlehren der mathematischen Wissenschaften. Springer Berlin Heidelberg, 2013. URL: <https://books.google.nl/books?id=hS3qCAAQBAJ>.
- Oh02. Hee Oh. Uniform pointwise bounds for matrix coefficients of unitary representations and applications to Kazhdan constants. *Duke Math. J.*, 113(1):133–192, 2002. doi:10.1215/S0012-7094-02-11314-3.
- PR94. Vladimir Platonov and Andrei Rapinchuk. *Algebraic groups and number theory. Transl. from the Russian by Rachel Rowen*, volume 139 of *Pure Appl. Math., Academic Press*. Boston, MA: Academic Press, 1994.
- PS97. Wilhelm Plesken and Bernd Souvignier. Computing isometries of lattices. *Journal of Symbolic Computation*, 24(3-4):327–334, 1997.

- PW08. Chris Peikert and Brent Waters. Lossy trapdoor functions and their applications. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 187–196, 2008.
- PW24. Aurel Page and Benjamin Wesolowski. The supersingular endomorphism ring and one endomorphism problems are equivalent. In Marc Joye and Gregor Leander, editors, *Advances in Cryptology - EUROCRYPT 2024 - 43rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zurich, Switzerland, May 26-30, 2024, Proceedings, Part VI*, volume 14656 of *Lecture Notes in Computer Science*, pages 388–417. Springer, 2024. doi:10.1007/978-3-031-58751-1_14.
- Ram16. Srinivasa Ramanujan. On certain arithmetical functions. *Trans. Camb. Philos. Soc.*, 22:159–184, 1916. Id/No 9.
- Reg09. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *Journal of the ACM (JACM)*, 56(6):1–40, 2009.
- Sch98. Alexander Schiemann. Classification of Hermitian forms with the neighbour method. *J. Symbolic Comput.*, 26(4):487–508, 1998.
- Ser07. Jean-Pierre Serre. Bounds for the orders of the finite subgroups of $G(k)$. *Group representation theory*, pages 405–450, 2007.
- Shi64. Goro Shimura. Arithmetic of unitary groups. *Annals of Mathematics*, 79(2):369–409, 1964.
- SW26. Zygimantas Straznickas and Stephen A. Weis. HAWK- n key recovery reduces to SVP in dimension $n/2 + 1$. Cryptology ePrint Archive, Paper 2026/1593, 2026. URL: <https://eprint.iacr.org/2026/1593>.
- Szy03. Michael Szydło. Hypercubic lattice reduction and analysis of GGH and NTRU signatures. In *International Conference on the Theory and Applications of Cryptographic Techniques*, pages 433–448. Springer, 2003.
- vW24. Wessel P. J. van Woerden. Dense and smooth lattices in any genus. In Kai-Min Chung and Yu Sasaki, editors, *Advances in Cryptology - ASIACRYPT 2024 - 30th International Conference on the Theory and Application of Cryptology and Information Security, Kolkata, India, December 9-13, 2024, Proceedings, Part IV*, volume 15487 of *Lecture Notes in Computer Science*, pages 386–417. Springer, 2024. doi:10.1007/978-981-96-0894-2_13.

A Number-theoretic Preliminaries

We recall here, for completeness, the standard facts on number fields, CM extensions, ideals and completions that are summarized in Section 2.1. We refer to [NS13] for proofs.

A number field K is a finite extension of the field $\mathbb{Q}$ of rational numbers. In other words, it is of the form $K \stackrel{\text{def}}{=} \mathbb{Q}[X]/(P)$ where $P \in \mathbb{Q}[X]$ is an irreducible polynomial. The degree $d = [K : \mathbb{Q}]$ of the extension is by definition the dimension of K as a $\mathbb{Q}$ -vector space, and is equal to the degree of the polynomial P defining the extension. A number field K of degree d has exactly d complex embeddings $\sigma_i : K \rightarrow \mathbb{C}$. Any embedding whose image is included in $\mathbb{R}$ is called a real embedding; otherwise it is called a complex embedding. When all the complex embeddings of a number field K have their image in $\mathbb{R}$, the field K is called totally real. Conversely, if K has no real embeddings, it is called totally complex.

Similarly, an element $a \in K$ is called totally real if $\sigma(a) \in \mathbb{R}$ for every complex embedding σ of K . Moreover, it is called totally positive if all its embeddings are real and positive. For an element $a \in K$, we define its trace and norm as

$$\mathrm{Tr}(a) = \sum_{\sigma: K \rightarrow \mathbb{C}} \sigma(a), \quad \text{and} \quad N(a) \stackrel{\text{def}}{=} \prod_{\sigma: K \rightarrow \mathbb{C}} \sigma(a).$$

They are always rational numbers. For a number field K , we denote by $\mathcal{O}_K$ its ring of integers, *i.e.*, the integral closure of $\mathbb{Z}$ in K . It is a free $\mathbb{Z}$ -module of rank d .

CM fields and Cyclotomic Extensions A number field K is a CM field¹⁰ if K is totally complex, and there is a subfield F which is totally real and such that $K = F(\sqrt{a})$ for some totally negative $a \in F$, where $\sqrt{a}$ denotes any of the complex roots of $X^2 - a$. When F needs to be made explicit, we also say that K/F is a CM extension. An important property of CM extensions is that complex conjugation induces a non-trivial automorphism τ of $K = F(\sqrt{a})$ which acts by fixing F and sending $\sqrt{a}$ to the other root $-\sqrt{a}$. More precisely, for every complex embedding σ of K , and for every $x \in K$, we have $\overline{\sigma(x)} = \sigma(\tau(x))$. For that reason, τ is also called *complex conjugation* of K , and we will use the notation $\bar{x}$ to denote $\tau(x)$.

Important examples of CM fields are the cyclotomic fields $K = \mathbb{Q}(\zeta_m)$ obtained by adjoining to $\mathbb{Q}$ a primitive m -th root of unity. The maximal totally real subfield of such a field is $F = \mathbb{Q}(\zeta_m + \bar{\zeta}_m) = \mathbb{Q}(\zeta_m + \zeta_m^{-1})$.

Ideals and Splitting Behaviour An integral ideal $\mathfrak{p}$ of K is by definition an ideal of the ring $\mathcal{O}_K$, *i.e.*, an additive subgroup of $\mathcal{O}_K$ such that for every $x \in \mathcal{O}_K$ we have $x\mathfrak{p} \subset \mathfrak{p}$. An additive subgroup $\mathfrak{a} \neq \{0\}$ of K is a fractional ideal of K if there exists a non-zero element $z \in K^\times$ such that $z\mathfrak{a}$ is an integral ideal. When $\mathfrak{a}$ is a fractional ideal of a CM field K , the set $\bar{\mathfrak{a}} \stackrel{\text{def}}{=} \{\bar{x} \mid x \in \mathfrak{a}\}$ is also a fractional ideal of K called its complex conjugate.

When $\mathfrak{a}$ and $\mathfrak{b}$ are two fractional ideals of K , the product $\mathfrak{a}\mathfrak{b}$ is the fractional ideal of K generated by the products xy for $x \in \mathfrak{a}$ and $y \in \mathfrak{b}$. We say that $\mathfrak{b}$ divides $\mathfrak{a}$ when $\mathfrak{a} \subset \mathfrak{b}$, which is equivalent to the existence of an integral ideal $\mathfrak{c}$ such that $\mathfrak{a} = \mathfrak{b}\mathfrak{c}$. An integral ideal $\mathfrak{p}$ is said to be prime whenever $xy \in \mathfrak{p}$ implies $x \in \mathfrak{p}$ or $y \in \mathfrak{p}$. Equivalently, an integral ideal is prime when $\mathcal{O}_K/\mathfrak{p}$ is an integral domain. Furthermore, $\mathfrak{p}$ is said to be maximal when $\mathcal{O}_K/\mathfrak{p}$ is even a field. When $\mathfrak{p}$ is a prime ideal of $\mathcal{O}_K$, there is a unique prime number $p \in \mathbb{Z}$ such that $\mathfrak{p} \cap \mathbb{Z} = p\mathbb{Z}$. In this case, we say that p is the prime below $\mathfrak{p}$, and the residue field $\mathcal{O}_K/\mathfrak{p}$ is a finite extension $\mathbb{F}_{p^f}$ of $\mathbb{F}_p$. For any nonzero integral ideal $\mathfrak{a} \subset \mathcal{O}_K$, its norm is defined as $N(\mathfrak{a}) \stackrel{\text{def}}{=} |\mathcal{O}_K/\mathfrak{a}|$. When $\mathfrak{a}$ is a prime ideal, its norm is a prime power.

When K is a number field, its ring of integers $\mathcal{O}_K$ is a Dedekind domain. In particular, the nonzero prime ideals of $\mathcal{O}_K$ are exactly the maximal ideals, and

¹⁰ CM stands for complex multiplication.

the quotients $\mathcal{O}_K/\mathfrak{p}$ are finite fields. Furthermore, every integral ideal admits a factorization as a product of prime ideals.

When K/F is CM, let $\mathfrak{q} \subset \mathcal{O}_F$ be a prime ideal of F . We say that $\mathfrak{q}$ splits in K when the ideal $\mathfrak{q}\mathcal{O}_K$ factorizes as $\mathfrak{q}\mathcal{O}_K = \mathfrak{p}\bar{\mathfrak{p}}$ where $\mathfrak{p} \neq \bar{\mathfrak{p}}$ is a prime ideal of $\mathcal{O}_K$. In this case, the fields $\mathcal{O}_K/\mathfrak{p}$, $\mathcal{O}_K/\bar{\mathfrak{p}}$ and $\mathcal{O}_F/\mathfrak{q}$ are all isomorphic to the same finite field $\mathbb{F}_q$. The characteristic of $\mathbb{F}_q$ is the unique prime below $\mathfrak{q}$. When $\mathfrak{q}$ remains prime in $\mathcal{O}_K$, then $\mathfrak{q}$ is said to be *inert* and $\mathcal{O}_K/\mathfrak{q}$ is isomorphic to $\mathbb{F}_{q^2}$ where $\mathbb{F}_q = \mathcal{O}_F/\mathfrak{q}$. Finally, we say that $\mathfrak{q}$ *ramifies* in K when $\mathfrak{q}\mathcal{O}_K = \mathfrak{p}^2$ for some prime ideal $\mathfrak{p}$ of $\mathcal{O}_K$. Only finitely many prime ideals of F ramify.

Completion of CM fields Let K/F be a CM extension, and let $\mathfrak{q}$ be a prime ideal of F . For any nonzero $x \in F$, the subgroup $x\mathcal{O}_F$ is a fractional ideal of F and there exists $a \in \mathbb{Z}$ such that $\mathfrak{q}^a$ divides (x) but $\mathfrak{q}^{a+1}$ does not, which is denoted by $v_{\mathfrak{q}}(x)$. More specifically the $\mathfrak{q}$ -valuation of F is the map $v_{\mathfrak{q}}: F \rightarrow \mathbb{Z} \cup \{\infty\}$ defined by $v_{\mathfrak{q}}(0) = \infty$ and $v_{\mathfrak{q}}(x) = \max\{a \in \mathbb{Z} \mid x\mathcal{O}_F \subset \mathfrak{q}^a\}$ when $x \neq 0$. It is well-known that $v_{\mathfrak{q}}$ defines an absolute value on F and we denote by $F_{\mathfrak{q}}$ the completion of F with respect to that absolute value. The valuation map $v_{\mathfrak{q}}$ extends canonically to $F_{\mathfrak{q}}$ and its *valuation ring* is $\mathcal{O}_{F,\mathfrak{q}} \stackrel{\text{def}}{=} \{x \in F_{\mathfrak{q}} \mid v_{\mathfrak{q}}(x) \geq 0\}$. It is a local Dedekind domain which contains $\mathcal{O}_F$, with unique prime ideal $\mathfrak{q}\mathcal{O}_{F,\mathfrak{q}} \stackrel{\text{def}}{=} \{x \in F_{\mathfrak{q}} \mid v_{\mathfrak{q}}(x) > 0\}$, which we will also denote by $\mathfrak{q}$. This ideal is principal, and a generator ϖ is called a *uniformizer* of $F_{\mathfrak{q}}$. It holds that every fractional ideal of $F_{\mathfrak{q}}$ is a power of $\mathfrak{q}$ and in particular is also principal.

For a prime ideal $\mathfrak{q}$ of F , we denote by $K_{\mathfrak{q}} \stackrel{\text{def}}{=} K \otimes_F F_{\mathfrak{q}}$. Complex conjugation on K extends naturally to $K_{\mathfrak{q}}$ by $k \otimes f \mapsto \bar{k} \otimes f$. Assume that $\mathfrak{q} = \mathfrak{p}\bar{\mathfrak{p}}$ splits in K . In this case, the completions $K_{\mathfrak{p}}$ and $K_{\bar{\mathfrak{p}}}$ are both field extensions of degree 1 of $F_{\mathfrak{q}}$, i.e. both are isomorphic to $F_{\mathfrak{q}}$, which induces a ring isomorphism $K_{\mathfrak{q}} \simeq F_{\mathfrak{q}} \times F_{\mathfrak{q}}, k \otimes f \mapsto (kf, \bar{k}f)$. In particular, complex conjugation simply acts on $K_{\mathfrak{q}}$ by exchanging the factors.

B Equidistribution

B.1 Random Walk over a Special Genus

Our worst-case to average-case reductions will follow a similar path as prior works such as [dBDPMW20, PW24, dBPTW26]. More precisely, the average-case distribution for the problems we consider is the mass distribution π_X over the special genus X of an input lattice Λ . Starting from a fixed lattice $\Lambda \in X$, we look for a random related lattice Λ' which should be easily computable from Λ and such that an answer to some lattice problem over Λ' can be transported back to an answer to the problem for Λ with small loss in quality. Our goal is to find such a random-walk procedure which converges towards the mass distribution π_X .

From now on, we fix a lattice Λ and a prime ideal $\mathfrak{q} = \mathfrak{p}\bar{\mathfrak{p}}$ of $\mathcal{O}_F$ which splits in K . Let X_0 be the special genus of Λ , and let X_1 be the common special genus of all $\mathfrak{p}$ -neighbours of lattices in X_0 . Recall from Section 3.1 that each lattice

in X_0 has $D_{\mathfrak{q}} = \frac{q^r - 1}{q - 1}$ such neighbours. In order to analyze how this random process modifies an input distribution, we formalize it as an operator acting on probability distributions over X_i , or more generally on the space $L^2(X_i)$ of complex-valued functions on X_i equipped with the inner-product

$$\langle f, g \rangle_{\pi_i} \stackrel{\text{def}}{=} \sum_{x \in X_i} \pi_i(x) \overline{f(x)} g(x),$$

where $\pi_i = \pi_{X_i}$ is the mass distribution on the special genus X_i from Section 2.4.

Let $n_{\mathfrak{p}}(x, y)$ be the number of local $\mathfrak{p}$ -neighbour choices from a representative of $x \in X_0$ whose endpoint lies in the class $y \in X_1$. Define similarly $n_{\bar{\mathfrak{p}}}(y, x)$. These numbers are independent of the chosen representatives and

$$\sum_{y \in X_1} n_{\mathfrak{p}}(x, y) = D_{\mathfrak{q}} = \sum_{x \in X_0} n_{\bar{\mathfrak{p}}}(y, x). \quad (4)$$

Consider the operator $U_{\mathfrak{p}} : L^2(X_1) \rightarrow L^2(X_0)$ defined for each $f \in L^2(X_1)$ by

$$(U_{\mathfrak{p}}f)(x) = \frac{1}{D_{\mathfrak{q}}} \sum_{y \in X_1} n_{\mathfrak{p}}(x, y) f(y), \quad \text{for } x \in X_0.$$

This operator averages over all $\mathfrak{p}$ -neighbours. Consider as well the operator $U_{\bar{\mathfrak{p}}} : L^2(X_0) \rightarrow L^2(X_1)$ corresponding to the $\bar{\mathfrak{p}}$ -step, and defined analogously:

$$(U_{\bar{\mathfrak{p}}}f)(y) = \frac{1}{D_{\mathfrak{q}}} \sum_{x \in X_0} n_{\bar{\mathfrak{p}}}(y, x) f(x), \quad \text{for } y \in X_1.$$

Finally, let $P_{\mathfrak{q}} \stackrel{\text{def}}{=} U_{\mathfrak{p}} U_{\bar{\mathfrak{p}}} : L^2(X_0) \rightarrow L^2(X_0)$ be the operator representing a complete round of the paired random walk. It can be interpreted as follows. If $f \in L^2(X_0)$ denotes a probability density function, then it is readily seen that $P_{\mathfrak{q}}f$ is again a probability density function, which corresponds to sampling a lattice with density f and applying one round of the $(\mathfrak{p}, \bar{\mathfrak{p}})$ -random walk. More generally, if $\nu \in \mathbb{N}$ is a non-zero integer, $P_{\mathfrak{q}}^{\nu}f$ denotes the probability density function of the endpoint of the walk after ν steps.

The following lemma shows that $U_{\mathfrak{p}}$ and $U_{\bar{\mathfrak{p}}}$ are adjoint to each other.

Lemma B.1. *Let $\mathfrak{q} = \mathfrak{p}\bar{\mathfrak{p}}$ be a prime ideal of $\mathcal{O}_F$ which splits in K . For $x \in X_0$ and $y \in X_1$,*

$$\frac{n_{\mathfrak{p}}(x, y)}{|\text{Aut}(\Lambda_x)|} = \frac{n_{\bar{\mathfrak{p}}}(y, x)}{|\text{Aut}(\Lambda_y)|}$$

In particular,

$$\text{Mass}(X_0) = \text{Mass}(X_1), \quad \text{and} \quad U_{\bar{\mathfrak{p}}} = U_{\mathfrak{p}}^*.$$

Proof. Fix lattice representatives Λ_x and Λ_y of the isometry classes $x \in X_0$ and $y \in X_1$. Let $\mathcal{A}_{x,y}$ be the set of pairs (M, σ) such that M is a $\mathfrak{p}$ -neighbour of Λ_x

in the class y ; and $\sigma \in U(V)$ is an isometry which satisfies $\sigma M = \Lambda_y$. For each such M , there are $|\text{Aut}(\Lambda_y)|$ choices for σ . Therefore,

$$|\mathcal{A}_{x,y}| = n_{\mathfrak{p}}(x,y) |\text{Aut}(\Lambda_y)|.$$

Similarly, let $\mathcal{B}_{y,x}$ be the set of pairs (N, τ) such that N is a $\bar{\mathfrak{p}}$ -neighbour of Λ_y in the class x and $\tau N = \Lambda_x$. Then,

$$|\mathcal{B}_{y,x}| = n_{\bar{\mathfrak{p}}}(y,x) |\text{Aut}(\Lambda_x)|.$$

By definition of $\mathfrak{p}$ -neighbours, for $(M, \sigma) \in \mathcal{A}_{x,y}$ we have

$$M/(M \cap \Lambda_x) \simeq \mathcal{O}_K/\mathfrak{p}, \quad \text{and} \quad \Lambda_x/(M \cap \Lambda_x) \simeq \mathcal{O}_K/\bar{\mathfrak{p}}$$

Moreover, σ induces isomorphisms

$$M/(M \cap \Lambda_x) \simeq \sigma M/(\sigma M \cap \sigma \Lambda_x)$$

and

$$\Lambda_x/(M \cap \Lambda_x) \simeq \sigma \Lambda_x/(\sigma M \cap \sigma \Lambda_x).$$

Hence,

$$\Lambda_y/(\Lambda_y \cap \sigma \Lambda_x) \simeq \mathcal{O}_K/\mathfrak{p}, \quad \text{and} \quad \sigma \Lambda_x/(\Lambda_y \cap \sigma \Lambda_x) \simeq \mathcal{O}_K/\bar{\mathfrak{p}}$$

i.e. $\sigma \Lambda_x$ is a $\bar{\mathfrak{p}}$ -neighbour of Λ_y , in the isometry class x . Consequently, the map

$$\begin{cases} \mathcal{A}_{x,y} & \longrightarrow & \mathcal{B}_{y,x} \\ (M, \sigma) & \longmapsto & (\sigma \Lambda_x, \sigma^{-1}) \end{cases}$$

is well defined, and is bijective. Its inverse is $(N, \tau) \mapsto (\tau \Lambda_y, \tau^{-1})$. Therefore, $n_{\mathfrak{p}}(x,y) |\text{Aut}(\Lambda_y)| = n_{\bar{\mathfrak{p}}}(y,x) |\text{Aut}(\Lambda_x)|$ or equivalently,

$$\frac{n_{\mathfrak{p}}(x,y)}{|\text{Aut}(\Lambda_x)|} = \frac{n_{\bar{\mathfrak{p}}}(y,x)}{|\text{Aut}(\Lambda_y)|}$$

The end of the proof follows by summing over $x \in X_0$ and $y \in X_1$ and using (4). $\square$

As a consequence of Lemma B.1, the walk-operator $P_{\mathfrak{q}}$ is self-adjoint and positive semidefinite on $L^2(X_0)$. Standard results on random walks on finite sets imply that π_0 is a stationary measure for the random walk defined by $P_{\mathfrak{q}}$ (see for example [CST08, Chapter 1]; strong approximation implies that this stationary measure is unique, see below). Moreover, the decomposition $L^2(X_0) = \mathbb{C}\mathbf{1} \oplus L_0^2(X_0)$ where

$$L_0^2(X_0) \stackrel{\text{def}}{=} (\mathbb{C}\mathbf{1})^\perp = \left\{ h \in L^2(X_0) \mid \sum_{x \in X_0} \pi_0(x) h(x) = 0 \right\},$$

is a $P_{\mathfrak{q}}$ -stable orthogonal decomposition, such that $P_{\mathfrak{q}}$ acts as the identity on $\mathbb{C}\mathbf{1}$. In order to get a quantitative convergence rate of the walk defined by $P_{\mathfrak{q}}$, it remains to bound the norm of $P_{\mathfrak{q}}$ on $L_0^2(X_0)$. More precisely, we can reformulate [CST08, Corollary 1.9.4] in our setting.

Proposition B.2. *Assume that $\|P_{\mathfrak{q}}|_{L_0^2(X_0)}\| \leq \rho$ for some $0 < \rho < 1$. Then for every probability distribution μ on X_0 , and every $\nu \in \mathbb{N}$,*

$$\Delta(P_{\mathfrak{q}}^\nu \mu, \pi_0) \leq \frac{1}{2} \rho^\nu \left\| \frac{\mu}{\pi_0} - \mathbf{1} \right\|_{2, \pi_0}.$$

In particular, if $\mu = \delta_x$ denotes the Dirac distribution centered at any isometry class $x \in X_0$, we get

$$\Delta(P_{\mathfrak{q}}^\nu \delta_x, \pi_0) \leq \frac{1}{2} \rho^\nu \sqrt{\pi_0(x)^{-1} - 1} = \frac{1}{2} \rho^\nu \sqrt{|\text{Aut}(\Lambda_x)| \text{Mass}(X_0) - 1}.$$

Proof. Let $h \stackrel{\text{def}}{=} \frac{\mu}{\pi_0}$. Observe that $h - \mathbf{1} \in L_0^2(X_0)$. Therefore, the Cauchy-Schwarz inequality yields

$$\begin{aligned} 2\Delta(P_{\mathfrak{q}}^\nu \mu, \pi_0) &= \sum_{x \in X_0} \pi_0(x) |P_{\mathfrak{q}}^\nu(h - \mathbf{1})(x)| \\ &\leq \|P_{\mathfrak{q}}^\nu(h - \mathbf{1})\| \leq \rho^\nu \|h - \mathbf{1}\|. \end{aligned}$$

If $\mu = \delta_x$, a simple computation shows that $\|h - \mathbf{1}\|^2 = \pi_0(x)^{-1} - 1$. $\square$

The goal of the following section will be to get an explicit bound ρ .

The $(\mathfrak{p}, \bar{\mathfrak{p}})$ -random walk produces global lattices in the same Hermitian space. Recall from Lemma 3.2 that a lattice Λ and the lattice M obtained after one half-step of the walk are locally equal at every prime of F except at $\mathfrak{q}$. More precisely, each half-step $U_{\mathfrak{p}}$ only modifies the completion at $\mathfrak{q}$. However, the $\mathfrak{p}$ -neighbour step sends classes in the special genus X_0 of Λ to classes in a single special genus X_1 , while the reverse $\bar{\mathfrak{p}}$ -step sends classes in X_1 back to X_0 .

The goal of this section is to realize the corresponding operator $U_{\mathfrak{p}}$ on a larger Hilbert space on which number-theoretic estimates are available. We construct a finite class set $\mathcal{X}_{\mathfrak{q}}$ containing both special genera X_0 and X_1 , and an operator $T_{\mathfrak{p}}$ on $L^2(\mathcal{X}_{\mathfrak{q}})$ which realizes $U_{\mathfrak{p}}$. The operator $T_{\mathfrak{p}}$ is known as the *normalized Hecke operator* associated with the local $\mathfrak{p}$ -neighbour. These operators are well studied in the literature and many estimates are known. Hecke operators have recently gained significant interest in several other cryptographic reductions [dBDPMW20, PW24, dBPTW26].

B.2 Local Description of $\mathfrak{p}$ -Neighbours

In this section, we follow the description of [GV14]. Let V be a Hermitian space of dimension $r \geq 2$ and let $\Lambda \subset V$ be a lattice in V . Let $\mathfrak{q} = \mathfrak{p}\bar{\mathfrak{p}}$ be a prime of F which splits in K with $\mathfrak{p} \nmid 2$ and $\mathfrak{p} \nmid \det(\Lambda)$. Let $F_{\mathfrak{q}}$ be the completion of F at $\mathfrak{q}$ and let ϖ be a uniformizer. Recall from Lemma 2.3 that the unitary groups identify with groups of invertible matrices over $F_{\mathfrak{q}}$:

$$\begin{aligned} U(V_{\mathfrak{q}}) &\simeq \mathbf{GL}_r(F_{\mathfrak{q}}) \\ SU(V_{\mathfrak{q}}) &\simeq \mathbf{SL}_r(F_{\mathfrak{q}}). \end{aligned}$$

Let M be a $\mathfrak{p}$ -neighbour of Λ . After completion at $\mathfrak{q} = \mathfrak{p}\bar{\mathfrak{p}}$, write

$$\Lambda_{\mathfrak{q}} = \Lambda_{\mathfrak{p}} \oplus \Lambda_{\bar{\mathfrak{p}}}, \quad M_{\mathfrak{q}} = M_{\mathfrak{p}} \oplus M_{\bar{\mathfrak{p}}}.$$

By definition,

$$\Lambda/(\Lambda \cap M) \simeq \mathcal{O}_K/\bar{\mathfrak{p}}$$

and therefore looking only at the component at $\mathfrak{p}$ we have

$$(\Lambda/(\Lambda \cap M))_{\mathfrak{p}} = 0.$$

In particular, $\Lambda_{\mathfrak{p}} = (\Lambda \cap M)_{\mathfrak{p}} \subset M_{\mathfrak{p}}$. On the other hand,

$$(M/(\Lambda \cap M))_{\mathfrak{p}} \simeq \mathcal{O}_{F,\mathfrak{q}}/\varpi.$$

In other words, the neighbour conditions imply

$$\Lambda_{\mathfrak{p}} \subset M_{\mathfrak{p}} \subset \varpi^{-1}\Lambda_{\mathfrak{p}}, \quad M_{\mathfrak{p}}/\Lambda_{\mathfrak{p}} \simeq \mathcal{O}_{F,\mathfrak{q}}/(\varpi).$$

Following [GV14, Section 4], the theory of elementary divisors therefore shows that every matrix taking $\Lambda_{\mathfrak{p}}$ to $M_{\mathfrak{p}}$ belongs to

$$\underbrace{\mathbf{GL}_r(\mathcal{O}_{F,\mathfrak{q}}) \begin{pmatrix} \varpi^{-1} & 0 & \cdots & 0 \\ 0 & 1 & & \\ \vdots & & \ddots & 0 \\ 0 & \cdots & 0 & 1 \end{pmatrix}}_{\stackrel{\text{def}}{=} a_{\mathfrak{p}}} \mathbf{GL}_r(\mathcal{O}_{F_{\mathfrak{q}}}) = \bigsqcup_{j=1}^{D_{\mathfrak{q}}} h_j \mathbf{GL}_r(\mathcal{O}_{F,\mathfrak{q}}) \quad (5)$$

where (h_j) form a set of coset-representatives. Conversely, [GV14] shows that every $\mathfrak{p}$ -neighbour can be locally described by cosets of $a_{\mathfrak{p}}$. More precisely, we may attach a global lattice to any local isometry in the following way.

Definition B.3. *Let Λ be a fixed reference lattice, and let $g \in U(V_{\mathfrak{q}})$ be a local isometry. We associate to g the unique (global) lattice $\Lambda(g)$ defined by*

$$\Lambda(g)_{\mathfrak{t}} = \Lambda_{\mathfrak{t}} \text{ for } \mathfrak{t} \neq \mathfrak{q}, \quad \text{and} \quad \Lambda(g)_{\mathfrak{q}} = g\Lambda_{\mathfrak{q}}.$$

In particular, $\Lambda(g)$ may differ from Λ only at $\mathfrak{q}$, and is locally isometric to Λ .

With this definition, we can state the main proposition relating global $\mathfrak{p}$ -neighbours of Λ to the local matrices $a_{\mathfrak{p}}$.

Proposition B.4 (Adapted from [GV14, Proposition 5.6 and Eq. (14)]).

Let $\mathcal{K}_{\mathfrak{q}} \stackrel{\text{def}}{=} \mathbf{GL}_r(\mathcal{O}_{F,\mathfrak{q}})$ and let $a_{\mathfrak{p}} = \text{diag}(\varpi^{-1}, 1, \dots, 1)$ be as above. Write $\mathcal{K}_{\mathfrak{q}} a_{\mathfrak{p}} \mathcal{K}_{\mathfrak{q}} = \bigsqcup h_j \mathcal{K}_{\mathfrak{q}}$.

Then, the map

$$h_j \mathcal{K}_{\mathfrak{q}} \mapsto \Lambda(h_j)$$

is a bijection between the set of (global) $\mathfrak{p}$ -neighbours of Λ , and the coset-representatives in $\mathcal{K}_{\mathfrak{q}} a_{\mathfrak{p}} \mathcal{K}_{\mathfrak{q}}$.

B.3 Hecke Operators

Fix a lattice $\Lambda \subset V$. Let $\mathfrak{q} = \mathfrak{p}\bar{\mathfrak{p}}$ be a prime of F which splits in K . In the previous section we showed how we may associate a specific coset of a local isometry to any $\mathfrak{p}$ -neighbour. In this section we generalize this approach and will define a larger double-quotient parametrizing all isometry classes of global lattices which are reachable by iterating $\mathfrak{p}$ -neighbours and $\bar{\mathfrak{p}}$ -neighbours. The corresponding Hecke operator will then be a generalization of the half-step walk operators $U_{\mathfrak{p}}$ and $U_{\bar{\mathfrak{p}}}$.

Let $\mathcal{K}_{\mathfrak{q}} \stackrel{\text{def}}{=} \mathbf{GL}_r(\mathcal{O}_{F,\mathfrak{q}})$ be the stabilizer of $\Lambda_{\mathfrak{q}}$ under the action of $U(V_{\mathfrak{q}})$ so that two local isometries which differ by the right multiplication of an element in $\mathcal{K}_{\mathfrak{q}}$ define the same global lattice. Finally, let

$$\Gamma_{\mathfrak{q}} \stackrel{\text{def}}{=} \{\gamma \in U(V) : \gamma\Lambda_{\mathfrak{t}} = \Lambda_{\mathfrak{t}} \text{ for every prime } \mathfrak{t} \neq \mathfrak{q}\}$$

be the group of global isometries which preserve Λ locally away from $\mathfrak{q}$. By completion at $\mathfrak{q}$, it can be identified with a subgroup of $U(V_{\mathfrak{q}})$. In particular, $\Gamma_{\mathfrak{q}}$ acts on the latter by left multiplication $\gamma \cdot g \stackrel{\text{def}}{=} \gamma_{\mathfrak{q}}g$.

We consider the double-coset set corresponding to both actions on the left and on the right:

$$\mathcal{X}_{\mathfrak{q}} \stackrel{\text{def}}{=} \Gamma_{\mathfrak{q}} \backslash U(V_{\mathfrak{q}}) / \mathcal{K}_{\mathfrak{q}}.$$

More precisely, two local isometries $g, g' \in U(V_{\mathfrak{q}})$ represent the same class in $\mathcal{X}_{\mathfrak{q}}$ when $g' = \gamma_{\mathfrak{q}}gk$ for some $\gamma \in \Gamma_{\mathfrak{q}}$ and $k \in \mathcal{K}_{\mathfrak{q}}$.

By the following proposition, $\mathcal{X}_{\mathfrak{q}}$ parametrizes the (global) isometry classes of lattices of the form $\Lambda(g)$ equal to Λ away from $\mathfrak{q}$ and locally isometric to Λ .

Proposition B.5. *The map*

$$\Gamma_{\mathfrak{q}}g\mathcal{K}_{\mathfrak{q}} \mapsto [\Lambda(g)]$$

is a bijection from $\mathcal{X}_{\mathfrak{q}}$ onto the set of isometry classes admitting a representative which equals Λ locally away from $\mathfrak{q}$ and is locally isometric to Λ at $\mathfrak{q}$. In particular, $\mathcal{X}_{\mathfrak{q}}$ is finite.

Proof. Let $g \in U(V_{\mathfrak{q}})$. Right multiplication by $\mathcal{K}_{\mathfrak{q}}$ does not change $g\Lambda_{\mathfrak{q}}$. On the other hand, if $\gamma \in \Gamma_{\mathfrak{q}}$, then $\Lambda(\gamma g) = \gamma\Lambda(g)$, so left multiplication induces (globally) isometric lattices.

Conversely, let $g, g' \in U(V_{\mathfrak{q}})$ and suppose a global isometry γ maps $\Lambda(g)$ to $\Lambda(g')$. Away from $\mathfrak{q}$, both lattices are equal to Λ and therefore γ must belong to $\Gamma_{\mathfrak{q}}$. At $\mathfrak{q}$, we have the following equality

$$\gamma_{\mathfrak{q}}g\Lambda_{\mathfrak{q}} = g'\Lambda_{\mathfrak{q}}$$

which implies $g'^{-1}\gamma_{\mathfrak{q}}g \in \mathcal{K}_{\mathfrak{q}}$. Hence, g and g' belong to the same double-coset.

Finiteness comes from the fact that (the image of) $\mathcal{X}_{\mathfrak{q}}$ is contained in the genus of Λ . $\square$

We next relate this double-quotient to the finite set carrying our random walk. We will follow the description of [Kir16, Section 5].

Definition B.6 ([Kir16, Definition 5.2.6]). *We denote by $\mathcal{N}(\Lambda, \mathfrak{p})$ the set of lattices M for which there exist a global isometry $\sigma \in U(V)$ and a sequence of lattices*

$$\Lambda = \Lambda_0, \Lambda_1, \dots, \Lambda_s = \sigma(M),$$

such that Λ_i is a $\mathfrak{p}$ -neighbour of Λ_{i-1} for every $1 \leq i \leq s$.

Proposition B.7 ([Kir16, Proposition 5.2.7]). *The set $\mathcal{N}(\Lambda, \mathfrak{p})$ coincides with*

$$\{M \in \text{gen}(\Lambda) \mid \text{there exists } M' \in [M] \text{ such that } \Lambda_{\mathfrak{t}} = M'_{\mathfrak{t}} \text{ for } \mathfrak{t} \neq \mathfrak{q}\}.$$

When taking isometry classes in Proposition B.7, this precisely coincides with the lattice interpretation described in Proposition B.5. More specifically we have the following corollary.

Corollary B.8. $\mathcal{X}_{\mathfrak{q}}$ *is the set of isometry classes of lattices reachable from any lattice in $[\Lambda]$ by applying $\mathfrak{p}$ -neighbours iteratively.*

The space of complex-valued functions on $\mathcal{X}_{\mathfrak{q}}$ comes equipped with the action of specific averaging operators related to $\mathfrak{p}$ -neighbours and known as *Hecke operators*. Indeed, recall the local correspondence between $\mathfrak{p}$ -neighbours and local isometries from Proposition B.4; with $a_{\mathfrak{p}} \stackrel{\text{def}}{=} \text{diag}(\varpi^{-1}, 1, \dots, 1)$ as earlier, write the decomposition

$$\mathcal{K}_{\mathfrak{q}} a_{\mathfrak{p}} \mathcal{K}_{\mathfrak{q}} = \bigsqcup_{j=1}^{D_{\mathfrak{q}}} h_j \mathcal{K}_{\mathfrak{q}},$$

where $\mathcal{K}_{\mathfrak{q}} = \mathbf{GL}_r(\mathcal{O}_{F, \mathfrak{q}})$. Then, for every complex-valued function f on $\mathcal{X}_{\mathfrak{q}}$, define the operator $T_{\mathfrak{p}}$ as

$$(T_{\mathfrak{p}} f)(\Gamma_{\mathfrak{q}} g \mathcal{K}_{\mathfrak{q}}) \stackrel{\text{def}}{=} \frac{1}{D_{\mathfrak{q}}} \sum_{j=1}^{D_{\mathfrak{q}}} f(\Gamma_{\mathfrak{q}} g h_j \mathcal{K}_{\mathfrak{q}}). \quad (6)$$

It is readily seen that this definition does not depend on the choice of representative for the different cosets involved. Hence, $T_{\mathfrak{p}}$ is well-defined. Furthermore, under the lattice interpretation of $\mathcal{X}_{\mathfrak{q}}$ and Proposition B.4, it can be seen as an averaging operator over $\mathfrak{p}$ -neighbours.

Proposition B.9. *For every complex-valued function f on $\mathcal{X}_{\mathfrak{q}}$ and every $[M] = [\Lambda(g)] \in \mathcal{X}_{\mathfrak{q}}$,*

$$(T_{\mathfrak{p}} f)([M]) = \frac{1}{D_{\mathfrak{q}}} \sum_{N \sim_{\mathfrak{p}} M} f([N]),$$

where the sum ranges over all lattices N which are $\mathfrak{p}$ -neighbours of M , and isometric neighbours are counted with multiplicities.

It remains to show that this provides a suitable ambient space to analyze our random walk. More precisely, we will identify the special genus X_0 of our starting lattice Λ , and the common special genus X_1 of each of its $\mathfrak{p}$ -neighbours, as contained in $\mathcal{X}_{\mathfrak{q}}$. This stems from an important result in the theory of algebraic groups known as *strong approximation*, which is due to Shimura in the case of $SU(V)$ [Shi64]. We refer the reader to [Sch98, Section 2], [GV14, Section 5] and [Kir16, Section 5.1] for precise statements. We only need a consequence of this strong approximation theorem.

Proposition B.10 ([Kir16, Corollary 5.2.8]). *Let Λ be a lattice in V , and let $\mathfrak{q}$ be a prime of F and $\mathfrak{p}$ a prime of K above $\mathfrak{q}$. If $\dim_F(V) \geq 3$ and the local space $V_{\mathfrak{q}}$ is isotropic, then the entire special genus of Λ is reachable by a sequence of $\mathfrak{p}$ -neighbours.*

Corollary B.11. *Let X_0 be the special genus of our initial lattice Λ , and let X_1 be the common special genus of any of its $\mathfrak{p}$ -neighbours. Then $X_0 \cup X_1 \subset \mathcal{X}_{\mathfrak{q}}$.*

Proof. Our Hermitian space V has K -dimension $r \geq 2$, and therefore $\dim_F(V) = 2r \geq 4 > 3$. Moreover, since $\mathfrak{q} = \mathfrak{p}\bar{\mathfrak{p}}$ splits in K , the completion $V_{\mathfrak{q}}$ is isotropic (see Remark 2.4). The above Proposition B.10 precisely implies that $X_0 \subset \mathcal{X}_{\mathfrak{q}}$. It suffices to consider first a $\mathfrak{p}$ -neighbour Λ_1 of Λ to get the inclusion $X_1 \subset \mathcal{X}_{\mathfrak{q}}$. $\square$

Finally, we equip $\mathcal{X}_{\mathfrak{q}}$ with the natural mass distribution. For each $x \in \mathcal{X}_{\mathfrak{q}}$, let Λ_x be any lattice in the isometry class corresponding to x , and set

$$\text{Mass}(\mathcal{X}_{\mathfrak{q}}) \stackrel{\text{def}}{=} \sum_{x \in \mathcal{X}_{\mathfrak{q}}} \frac{1}{|\text{Aut}(\Lambda_x)|}, \quad \pi_{\mathcal{X}_{\mathfrak{q}}}(x) \stackrel{\text{def}}{=} \frac{|\text{Aut}(\Lambda_x)|^{-1}}{\text{Mass}(\mathcal{X}_{\mathfrak{q}})}.$$

This measure is compatible with the mass measures π_0 and π_1 we put on the special genera X_0 and X_1 . Let $L^2(\mathcal{X}_{\mathfrak{q}})$ be the corresponding space of complex-valued functions on $\mathcal{X}_{\mathfrak{q}}$. With all the discussion in this section, it is now readily seen that $L^2(X_0)$ and $L^2(X_1)$ both embed isometrically in $L^2(\mathcal{X}_{\mathfrak{q}})$ by extending functions on X_i by 0 outside of their domain. In particular, for every $v \in L^2(X_1)$ and $w \in L^2(X_0)$, we have

$$\langle U_{\mathfrak{p}}v, w \rangle_{\pi_{X_0}} = \langle T_{\mathfrak{p}}v, w \rangle_{\pi_{\mathcal{X}_{\mathfrak{q}}}},$$

where we still denote by v, w their extensions by 0 to $\mathcal{X}_{\mathfrak{q}}$. In particular,

$$\|U_{\mathfrak{p}}: L^2_0(X_1) \rightarrow L^2_0(X_0)\| = \sup_{\substack{v \in L^2_0(X_1), w \in L^2_0(X_0) \\ \|v\|=1, \|w\|=1}} |\langle T_{\mathfrak{p}}v, w \rangle|. \quad (7)$$

B.4 Spectral Estimates

In order to bound the operator norm in (7), we will follow the route of [COU01]. The final ingredient that we require will be an expression of $T_{\mathfrak{p}}$ in terms of *unitary representations* of $U(V_{\mathfrak{q}}) \simeq \mathbf{GL}_r(F_{\mathfrak{q}})$.

Recall the definition

$$\Gamma_{\mathfrak{q}} \stackrel{\text{def}}{=} \{\gamma \in U(V) : \gamma \Lambda_{\mathfrak{t}} = \Lambda_{\mathfrak{t}} \text{ for every prime } \mathfrak{t} \neq \mathfrak{q}\},$$

and consider the set

$$\mathcal{Y}_{\mathfrak{q}} \stackrel{\text{def}}{=} \Gamma_{\mathfrak{q}} \backslash U(V_{\mathfrak{q}})$$

of orbits of $U(V_{\mathfrak{q}})$ under left multiplication by $\Gamma_{\mathfrak{q}}$. Observe that $\mathcal{K}_{\mathfrak{q}} \simeq \mathbf{GL}_r(\mathcal{O}_{F,\mathfrak{q}})$ is a compact-open subgroup of $U(V_{\mathfrak{q}})$. Since $\mathcal{X}_{\mathfrak{q}}$ is finite, $\mathcal{Y}_{\mathfrak{q}}$ is a finite union of compact sets, and is therefore compact as well. Furthermore, since $\mathcal{K}_{\mathfrak{q}}$ is the stabilizer of $\Lambda_{\mathfrak{q}}$, it follows that $\Gamma_{\mathfrak{q}} \cap \mathcal{K}_{\mathfrak{q}} = \text{Aut}(\Lambda)$ is finite. Therefore, there exists an open neighbourhood $W \subset \mathcal{K}_{\mathfrak{q}}$ of 1 which avoids the finitely many nonidentity elements. Then, $W \cap \Gamma_{\mathfrak{q}} = W \cap (\Gamma_{\mathfrak{q}} \cap \mathcal{K}_{\mathfrak{q}}) = \{1\}$ which exactly means that $\Gamma_{\mathfrak{q}}$ is a discrete subgroup of $U(V_{\mathfrak{q}})$.¹¹

Moreover, the group $U(V_{\mathfrak{q}})$ is locally compact, and hence carries a Haar measure. We denote by $\mu_{\mathfrak{q}}$ the corresponding normalized Haar measure on $\mathcal{Y}_{\mathfrak{q}}$, and let $\mathcal{H}_{\mathfrak{q}} \stackrel{\text{def}}{=} L^2(\mathcal{Y}_{\mathfrak{q}}, \mu_{\mathfrak{q}})$. This Hilbert space naturally carries an action of $U(V_{\mathfrak{q}})$, by right multiplication:

$$(a \cdot f)(\gamma_{\mathfrak{q}} g) \stackrel{\text{def}}{=} f(\gamma_{\mathfrak{q}} g a), \quad a \in U(V_{\mathfrak{q}}), f \in L^2(\mathcal{Y}_{\mathfrak{q}}). \quad (8)$$

Denote by $\rho(a)$ the corresponding operator on $\mathcal{H}_{\mathfrak{q}}$. Since $\mu_{\mathfrak{q}}$ is invariant by translation, this action preserves the L^2 norm, in other words ρ is a *unitary representation* of $U(V_{\mathfrak{q}})$.

Furthermore, observe that the functions on $\mathcal{Y}_{\mathfrak{q}}$ which are fixed by $\mathcal{K}_{\mathfrak{q}}$ are exactly the functions which descend to $\mathcal{X}_{\mathfrak{q}}$. In particular, under this identification, the Hecke operator defined in (6) acts on $\mathcal{K}_{\mathfrak{q}}$ -invariant functions on $\mathcal{Y}_{\mathfrak{q}}$ by

$$(T_{\mathfrak{p}} f) \stackrel{\text{def}}{=} \frac{1}{D_{\mathfrak{q}}} \sum_{j=1}^{D_{\mathfrak{q}}} \rho(h_j)(f). \quad (9)$$

Therefore, for any unit vectors $v, w \in \mathcal{H}_{\mathfrak{q}}^{\mathcal{K}_{\mathfrak{q}}}$,

$$\langle T_{\mathfrak{p}} v, w \rangle = \frac{1}{D_{\mathfrak{q}}} \sum_{j=1}^{D_{\mathfrak{q}}} \langle \rho(h_j) v, w \rangle.$$

Furthermore, each h_j is of the form $h_j = k_j a_{\mathfrak{p}}$ with $k_j \in \mathcal{K}_{\mathfrak{q}}$. Hence,

$$\langle \rho(h_j) v, w \rangle = \langle \rho(k_j) \rho(a_{\mathfrak{p}}) v, w \rangle = \langle \rho(a_{\mathfrak{p}}) v, \rho(k_j^{-1}) w \rangle = \langle \rho(a_{\mathfrak{p}}) v, w \rangle,$$

where in the last equality we used the fact that $\rho(k_j^{-1})$ corresponds to right multiplication by $k_j^{-1} \in \mathcal{K}_{\mathfrak{q}}$, and w is $\mathcal{K}_{\mathfrak{q}}$ invariant. Taking the average yields

$$|\langle T_{\mathfrak{p}} v, w \rangle| = |\langle \rho(a_{\mathfrak{p}}) v, w \rangle|. \quad (10)$$

¹¹ With the terminology of [COU01], $\Gamma_{\mathfrak{q}}$ is a *lattice* in $U(V_{\mathfrak{q}})$, in the topological group meaning. It is an example of an *S-arithmetic group*.

In the literature on Hecke operators, these quantities are referred to as *matrix coefficients* of the unitary representation ρ and are thoroughly studied in [Oh02]. However, there are two obstructions to applying the results in our setting. Indeed, the bound requires ρ to not have any $\mathbf{SL}_r(F_q)$ -invariant vectors. But more importantly, it requires $r \geq 3$. We will therefore consider the two situations separately.

B.4.1 Rank at least 3: Applying Oh's local bound Assume $r \geq 3$. The bound of [Oh02, Theorem 1.1] would precisely answer our problem. Before stating it we will need to restrict to unitary representations which do not contain $\mathbf{SL}_r(F_q)$ -invariant vectors. However, the full representation on $\mathcal{H}_q$ does not satisfy this hypothesis.

Define

$$\mathcal{H}_q^{\mathbf{SL}_r(F_q)} \stackrel{\text{def}}{=} \{z \in \mathcal{H}_q : \rho(s)z = z \text{ for every } s \in \mathbf{SL}_r(F_q)\}.$$

The space $\mathcal{H}_q^{\mathcal{K}_q} \cap \mathcal{H}_q^{\mathbf{SL}_r(F_q)}$ may be interpreted as the set of functions which are precisely constant on each special genus in $\mathcal{X}_q$. Consider instead

$$\mathcal{Z}_q \stackrel{\text{def}}{=} \left(\mathcal{H}_q^{\mathbf{SL}_r(F_q)}\right)^\perp \cap \mathcal{H}_q^{\mathcal{K}_q}.$$

A simple computation shows that $\mathcal{Z}_q$ is stable under the action of the Hecke operator T_p , and the remark above shows that the spaces $L_0^2(X_i)$ both embed in $\mathcal{Z}_q$. Consequently, we obtain

$$\begin{aligned} \|U_p : L_0^2(X_1) \rightarrow L_0^2(X_0)\| &\leq \|T_p|_{\mathcal{Z}_q}\| \\ &= \sup_{\substack{v, w \in \mathcal{Z}_q \\ \|v\|=1, \|w\|=1}} |\langle \rho(a_p)v, w \rangle| \end{aligned} \quad (11)$$

We are now ready to use the results from [Oh02].

Theorem B.12 (Adapted from [Oh02, Theorem 1.1]). *Let $r \geq 3$, and let ρ be a unitary representation of $\mathbf{GL}_r(F_q)$ with no nonzero $\mathbf{SL}_r(F_q)$ -invariant vector. Let $q = N_{F/\mathbb{Q}}(\mathfrak{q})$ be the absolute norm of $\mathfrak{q}$. If v, w are fixed by $\mathbf{GL}_r(\mathcal{O}_{F, \mathfrak{q}})$, then for $\varepsilon \in \{-1, 1\}$,*

$$|\langle \rho(\text{diag}(\varpi^\varepsilon, 1, \dots, 1))v, w \rangle| \leq \Xi_q(1) \|v\| \|w\| = \frac{2\sqrt{q}}{q+1} \|v\| \|w\|.$$

The function Ξ_q is the rank-one Harish-Chandra function (see [Oh02] for more details), defined for $j \geq 0$ by

$$\Xi_q(j) \stackrel{\text{def}}{=} \frac{1}{\sqrt{q}} \frac{j(q-1) + (q+1)}{q+1}.$$

We finally get the following bound on the operator norm of the $\mathfrak{p}$ -neighbour walk.

Proposition B.13. *Let K/F be a CM extension, $r \geq 3$, and let $\mathfrak{q}$ be a prime of F which splits in K as $\mathfrak{q} = \mathfrak{p}\bar{\mathfrak{p}}$. Denote by $q = N_{F/\mathbb{Q}}(\mathfrak{q}) = |\mathcal{O}_F/\mathfrak{q}|$ the absolute norm of $\mathfrak{q}$. Assume that $\mathfrak{q}$ does not divide $\det(\Lambda)$. Then,*

$$\|U_{\mathfrak{p}} : L_0^2(X_1) \rightarrow L_0^2(X_0)\| \leq \frac{2\sqrt{q}}{q+1}.$$

B.4.2 The Rank 2 Case When $r = 2$, no local bound like Theorem B.12 can exist, so we must use global arguments. We derive the bound we need from known results via an exceptional isomorphism (see also [Car86, §2.1 and §2.2] for a similar argument) and the Jacquet–Langlands correspondence (see also [KV10, Section 7] for a similar argument), reducing to Hecke operators on Hilbert modular cusp forms.

To formulate these global arguments we need to introduce adelic language. Let $\hat{F} = \prod_{\mathfrak{l}} F_{\mathfrak{l}} = F \otimes_{\mathbb{Z}} \hat{\mathbb{Z}}$ denote the topological ring of finite adeles of F . For various algebraic groups $\mathbb{G}$ and compact open subgroups $\mathcal{K} \subset \mathbb{G}(\hat{F})$, we will consider the sets

$$\mathcal{X}(\mathbb{G}, \mathcal{K}) = \mathbb{G}(F) \backslash \mathbb{G}(\hat{F}) / \mathcal{K}.$$

The L^2 functions on $\mathcal{X}(\mathbb{G}, \mathcal{K})$ admit an action of Hecke operators induced by the double cosets of $\mathcal{K}$.

Proposition B.14. *Let K/F be a CM extension, $r = 2$, and let $\mathfrak{q}$ be a prime of F which splits in K as $\mathfrak{q} = \mathfrak{p}\bar{\mathfrak{p}}$. Denote by $q = N_{F/\mathbb{Q}}(\mathfrak{q}) = |\mathcal{O}_F/\mathfrak{q}|$ the absolute norm of $\mathfrak{q}$. Let p be the rational prime below $\mathfrak{q}$ and assume that $p \nmid \det_{\mathbb{Q}}(\Lambda)^2$. Then,*

$$\|U_{\mathfrak{p}} : L_0^2(X_1) \rightarrow L_0^2(X_0)\| \leq \frac{2\sqrt{q}}{q+1}.$$

Proof. Let $\hat{\Lambda} = \hat{\mathbb{Z}} \otimes \Lambda$.

By Chevalley’s theorem [Che51, Théorème 1] on the congruence subgroup problem for units, there exists an ideal $\mathfrak{N}$ of $\mathcal{O}_F$ coprime to p such that the kernel $\mathcal{O}_F^{\times}(\mathfrak{N})$ of the map $\mathcal{O}_F^{\times} \rightarrow (\mathcal{O}_F/\mathfrak{N})^{\times}$ satisfies $\mathcal{O}_F^{\times}(\mathfrak{N}) \subset (\mathcal{O}_F^{\times})^2$.

Let $\mathbb{G}_1 = \mathrm{U}(V)$, let $\mathcal{K}_1$ be the stabilizer of $\hat{\Lambda}$ in $\mathbb{G}_1(\hat{F})$, and consider the genus $\mathcal{X}(\mathbb{G}_1, \mathcal{K}_1)$ of Λ . Let $\mathbb{G}_2 = \mathrm{GU}(V)$, *i.e.*

$$\mathbb{G}_2 = \{\sigma \in \mathrm{GL}(V) \mid \exists \lambda, \langle \sigma(x), \sigma(y) \rangle = \lambda \langle x, y \rangle \text{ for all } x, y \in V\}.$$

Let $\lambda : \mathbb{G}_2 \rightarrow \mathbb{G}_m$ denote the similitude factor morphism, so that $\mathbb{G}_1 = \ker(\lambda)$. Let $\mathcal{K}'_2$ be the stabilizer of $\hat{\Lambda}$ in $\mathbb{G}_2(\hat{F})$, and let $\mathcal{K}_2 = \ker(\lambda : \mathcal{K}'_2 \rightarrow (\mathcal{O}_F/\mathfrak{N})^{\times})$. We have $\mathbb{G}_1 \subset \mathbb{G}_2$ and $\mathcal{K}_1 \subset \mathcal{K}_2$. We claim that the natural map

$$\mathcal{X}(\mathbb{G}_1, \mathcal{K}_1) \rightarrow \mathcal{X}(\mathbb{G}_2, \mathcal{K}_2)$$

is injective. Indeed, if $x, y \in \mathbb{G}_1(\hat{F})$ have the same image in $\mathcal{X}(\mathbb{G}_2, \mathcal{K}_2)$, then there exists $\gamma \in \mathbb{G}_2(F)$ such that $y = \gamma x \bmod \mathcal{K}_2$. This implies that $\lambda(\gamma) = 1 \bmod \mathfrak{N}$, and therefore $\lambda(\gamma) = \alpha^2$ with $\alpha \in \mathcal{O}_F^{\times}$, so that $\lambda(\gamma/\alpha) = 1$, *i.e.* $\gamma' \in \mathbb{G}_1(F)$

where $\gamma' = \gamma/\alpha$. Since $\mathbb{G}_1(\hat{F}) \cap \alpha\mathcal{K}_2 \subset \mathcal{K}_1$, this proves that x, y have the same image in $\mathcal{X}(\mathbb{G}_1, \mathcal{K}_1)$, proving the claim.

We therefore have an injection $L^2(\mathcal{X}(\mathbb{G}_1, \mathcal{K}_1)) \subset L^2(\mathcal{X}(\mathbb{G}_2, \mathcal{K}_2))$ by extension by 0, compatible with the action of the Hecke operators that stabilize the first space.

We let $\mathbb{G}_3$ be the unit group of K as an algebraic group over F (i.e. $\mathbb{G}_3 = \text{Res}_{K/F}(\mathbb{G}_m)$), and we introduce the algebraic group

$$\mathbb{G}_4 = \{\sigma \in \text{GL}(V) \mid \langle \sigma(x), \sigma(y) \rangle = \det(\sigma) \langle x, y \rangle \text{ for all } x, y \in V\}.$$

We have natural inclusions $\mathbb{G}_3 \subset \mathbb{G}_2$ and $\mathbb{G}_4 \subset \mathbb{G}_2$. For $i = 3, 4$, let $\mathcal{K}_i = \mathbb{G}_i(\hat{F}) \cap \mathcal{K}_2$. Multiplication induces a surjection

$$\mathbb{G}_3 \times \mathbb{G}_4 \twoheadrightarrow \mathbb{G}_2;$$

this surjection also holds on points over any extension of F (by Hilbert's theorem 90), so we get a surjection

$$\mathcal{X}(\mathbb{G}_3, \mathcal{K}_3) \times \mathcal{X}(\mathbb{G}_4, \mathcal{K}_4) \twoheadrightarrow \mathcal{X}(\mathbb{G}_2, \mathcal{K}_2).$$

We therefore obtain an injection

$$L^2(\mathcal{X}(\mathbb{G}_2, \mathcal{K}_2)) \hookrightarrow L^2(\mathcal{X}(\mathbb{G}_3, \mathcal{K}_3)) \hat{\otimes} L^2(\mathcal{X}(\mathbb{G}_4, \mathcal{K}_4)),$$

compatible with the Hecke action. Since $\mathbb{G}_3$ is a torus, the corresponding L^2 space decomposes as a direct sum of Hecke eigenspaces corresponding to one-dimensional characters, and the corresponding eigenvalues are roots of unity. So we only need to examine the Hecke eigenvalues in the L^2 space of $\mathbb{G}_4$.

There exists a quaternion algebra A over F containing K , whose reduced norm nrd is a Hermitian form such that (A, nrd) is isometric to $(V, \langle \cdot, \cdot \rangle)$. Since V is isotropic at $\mathfrak{q}$, the algebra A is not ramified at $\mathfrak{q}$. It is well-known that the action by left multiplication of $A^\times$ on A then induces an isomorphism $\iota: \mathbb{G}_4 \cong \mathbb{G}_5$, where $\mathbb{G}_5$ is the unit group of A , so that

$$\mathcal{X}(\mathbb{G}_4, \mathcal{K}_4) \cong \mathcal{X}(\mathbb{G}_5, \mathcal{K}_5),$$

where $\mathcal{K}_5 = \iota(\mathcal{K}_4)$, and this bijection is compatible with the Hecke action.

By the Jacquet–Langlands correspondence [JL70, Theorem 14.4, see also Corollary 14.3], every Hecke eigenvalue system in $L^2(\mathcal{X}(\mathbb{G}_5, \mathcal{K}_5))$ that does not factor through the reduced norm is the same as one appearing in a corresponding space of Hilbert modular cusp forms of parallel weight 2.

Therefore, by [Liv01, Theorem 0.1], every eigenvalue λ of $U_{\mathfrak{p}}$ that is not a root of unity satisfies the Ramanujan bound $|\lambda| \leq \frac{2\sqrt{q}}{q+1}$; here we use that $p \nmid \Delta_F$, that the compact group $\mathcal{K}_5$ is maximal at every prime above p , and that the algebra is unramified at every prime above p .

Finally, let $X_0, X_1, \dots, X_m$ be the special genera that can be reached by iterating $\mathfrak{p}$ -neighbours. By strong approximation, every $U_{\mathfrak{p}}$ -eigenfunction in $\bigoplus_i L^2(X_i)$ with eigenvalue a root of unity is constant on each X_i . Therefore, no eigenvalue of $U_{\mathfrak{p}}$ on $\bigoplus_i L^2_0(X_i)$ is a root of unity. As $U_{\mathfrak{p}}$ acts as a normal operator on $\bigoplus_i L^2_0(X_i)$ (Lemma B.1), its norm equals its spectral radius and the claimed bound follows. $\square$

B.5 The Equidistribution Theorem

We are now ready to wrap up everything and give an explicit bound for the convergence rate of our paired $(\mathfrak{p}, \bar{\mathfrak{p}})$ -random walk.

Theorem B.15. *We fix a CM extension K/F . Let Λ be a Hermitian lattice of rank $r \geq 2$ and let $X_0 \stackrel{\text{def}}{=} \text{sgen}(\Lambda)$ be its special genus. Let $\mathfrak{q}$ be a prime ideal of $\mathcal{O}_F$ which splits as $\mathfrak{q}\mathcal{O}_K = \mathfrak{p}\bar{\mathfrak{p}}$ and put $q = |\mathcal{O}_F/\mathfrak{q}|$. Let p be the rational prime below $\mathfrak{q}$ and suppose that $p \nmid 2 \det_{\mathbb{Q}}(\Lambda)^2$. Let $P_{\mathfrak{q}} = U_{\mathfrak{p}}U_{\bar{\mathfrak{p}}}$ be the paired walk operator. Then*

$$\|P_{\mathfrak{q}}|_{L_0^2(X_0)}\| \leq \frac{4q}{(q+1)^2}.$$

Proof. Recall that $U_{\bar{\mathfrak{p}}} = U_{\mathfrak{p}}^*$ and therefore $\|P_{\mathfrak{q}}\| = \|U_{\mathfrak{p}}\|^2$. Propositions B.13 and B.14 therefore give

$$\|P_{\mathfrak{q}}|_{L_0^2(X_0)}\| \leq \left(\frac{2\sqrt{q}}{q+1} \right)^2 = \frac{4q}{(q+1)^2}.$$

□

By combining Theorem B.15 with Proposition B.2 we immediately deduce the bound on the statistical distance:

Corollary B.16. *For every isometry class $x \in \text{sgen}(\Lambda)$, denote by δ_x the Dirac distribution centered at x . Let $p \nmid 2 \det_{\mathbb{Q}}(\Lambda)^2$ be a rational prime with $p\mathcal{O}_F = \prod_{i=1}^t \mathfrak{q}_i$ and each $\mathfrak{q}_i$ split in K . The random-walk distribution resulting from sequentially applying the (paired) Hecke operators for the $\mathfrak{q}_i$, starting at a lattice in x , is $\prod_i P_{\mathfrak{q}_i} \delta_x$. The bound of Theorem B.15 gives the following bound on the statistical distance:*

$$\Delta \left(\prod_{i=1}^t P_{\mathfrak{q}_i} \delta_x, \pi_0 \right) \leq \frac{1}{2} \left(\prod_{i=1}^t \frac{4N(\mathfrak{q}_i)}{(N(\mathfrak{q}_i) + 1)^2} \right) \sqrt{|\text{Aut}(\Lambda_x)| \text{Mass}(X_0) - 1},$$

where Λ_x is any lattice representative of x , and π_0 is the mass measure on the special genus X_0 .

Lemma B.17. *Let $\mathfrak{q}$ be a prime of F with $q = N(\mathfrak{q})$ and $e = v_{\mathfrak{q}}(2)$, and let $\Lambda_{\mathfrak{q}}$ be an integral Hermitian lattice of rank r over $\mathcal{O}_K \otimes_{\mathcal{O}_F} \mathcal{O}_{F,\mathfrak{q}}$. Then its local density satisfies $\beta(\Lambda_{\mathfrak{q}}) \geq q^{-(e+1)r^2}$.*

Proof. For all but one type of prime the local density is known explicitly, and the claimed bound can be read off from the relevant formula.

At the *good* primes in the sense of [Kir16, Definition 3.5.1], that is at every $\mathfrak{q} \nmid 2$ and at every $\mathfrak{q} \mid 2$ which is not ramified in K/F , the density is given by [Kir16, Theorem 4.2.4]. At the remaining (ramified) primes $\mathfrak{q} \mid 2$ it is known when $\Lambda_{\mathfrak{q}}$ is square-free by [Kir16, Theorems 4.5.2 and 4.5.5], and it is known for an arbitrary $\Lambda_{\mathfrak{q}}$ whenever $F_{\mathfrak{q}}/\mathbb{Q}_2$ is unramified, by [Cho16, Theorem 5.2] and its

sequel. The one case left open in the literature is that of a prime $\mathfrak{q} \mid 2$ ramified in K/F with $F_{\mathfrak{q}}/\mathbb{Q}_2$ ramified and $\Lambda_{\mathfrak{q}}$ not square-free, which occurs for instance when K is cyclotomic of power-of-two conductor.

For the sake of completeness we provide a proof that covers all cases. A first version of this proof was produced with the assistance of Fable; we then rewrote and verified it, and identified the relevant references. Let $E = \text{End}(\Lambda_{\mathfrak{q}})$. Following Kirschmer [Kir16, Def. 4.2.2], we let H be the local integral group scheme such that

$$H(R) = \{\sigma \in \text{End}(\Lambda_{\mathfrak{q}} \otimes R) \mid \langle \sigma(x), \sigma(y) \rangle = \langle x, y \rangle \text{ for all } x, y \in \Lambda_{\mathfrak{q}} \otimes R\}.$$

We then have

$$\beta(\Lambda_{\mathfrak{q}}) = \lim_{N \rightarrow \infty} \frac{\#H(\mathcal{O}_{F,\mathfrak{q}}/\mathfrak{q}^N)}{q^{Nr^2}}.$$

We use the classical Cayley transform ([Cay46], see also [PR94, proof of Proposition 7.4]) to construct local points of H . Let $W \subset E$ be the subset of antisymmetric elements; this is a free $\mathcal{O}_{F,\mathfrak{q}}$ -module of rank r^2 [KMRT98, Proposition (2.17)] which is saturated in E . Let $C: W \rightarrow \text{U}(V_{\mathfrak{q}})$ be defined by

$$C: X \mapsto (1 - X)^{-1}(1 + X).$$

Then we have $C(\mathfrak{q}W) \subset E^{\times}$. Let $N \geq e + 1$. The identity

$$X - X' = \frac{1}{2}(1 - X)(C(X) - C(X'))(1 - X')$$

holds for all $X, X' \in \mathfrak{q}W$. Therefore, if $C(X) \equiv C(X') \pmod{\mathfrak{q}^N}$ then $X - X' \in \mathfrak{q}^{N-e}E \cap \mathfrak{q}W = \mathfrak{q}^{N-e}W$ (since W is saturated). The Cayley transform therefore induces an injection

$$C: \mathfrak{q}W/\mathfrak{q}^{N-e}W \hookrightarrow H(\mathcal{O}_{F,\mathfrak{q}}/\mathfrak{q}^N).$$

This shows that

$$\#H(\mathcal{O}_{F,\mathfrak{q}}/\mathfrak{q}^N) \geq \#(\mathfrak{q}W/\mathfrak{q}^{N-e}W) = q^{(N-e-1)r^2}.$$

Passing to the limit as $N \rightarrow \infty$ gives the lower bound on $\beta(\Lambda_{\mathfrak{q}})$. $\square$

Corollary B.18. *Let $r \geq 2$ be fixed, let Λ be an integral Hermitian lattice of rank r over a CM field K of degree d , put $n = rd$, and let $\varepsilon \in (0, 1)$. Let $p \nmid 2 \det_{\mathbb{Q}}(\Lambda)^2$ be a rational prime for which $p\mathcal{O}_F = \prod_{i=1}^t \mathfrak{q}_i$ with each $\mathfrak{q}_i$ split in K . If*

$$p \geq \text{poly}(d) \cdot \det_{\mathbb{Q}}(\Lambda)^{O_r(1)/n} \cdot (1/\varepsilon)^{2/d},$$

then $\Delta(\prod_{i=1}^t P_{\mathfrak{q}_i} \delta_{\Lambda}, \pi_0) \leq \varepsilon$, where π_0 is the mass measure on the special genus X_0 of Λ , that is, $\pi_0 = \mathcal{H}(X_0)$ (Definition 2.10).

Proof. Since $\prod_{i=1}^t N(\mathfrak{q}_i) = p^{d/2}$, $1 \leq t \leq d/2$ and $4N(\mathfrak{q}_i)/(N(\mathfrak{q}_i)+1)^2 \leq 4/N(\mathfrak{q}_i)$ for all i , Corollary B.16 gives

$$\Delta \leq \frac{1}{2}(4/p)^{d/2} \sqrt{|\text{Aut}(\Lambda)| \text{Mass}(X_0)},$$

so it suffices that $p^d \geq 4^{d-1} |\text{Aut}(\Lambda)| \text{Mass}(X_0) \varepsilon^{-2}$. Looking at the embedding of Λ as $\mathbb{Z}$ -lattice we obtain that $\text{Aut}(\Lambda)$ is a finite subgroup of $\text{GL}_n(\mathbb{Z})$, hence $|\text{Aut}(\Lambda)| \leq 2^{O(n \log n)}$ [Ser07] and $|\text{Aut}(\Lambda)|^{1/d} = d^{O(r)} = \text{poly}(d)$ which is sufficient for our application. For the mass, $\text{Mass}(X_0) \leq \text{Mass}(\text{gen } \Lambda)$, and the Hermitian mass formula [Kir16, Thm. 4.2.3] expresses the latter as

$$2^{O_r(d)} |\Delta_K|^{O(r^2)} N_{K/\mathbb{Q}}(\det \Lambda)^{O(r)} \cdot \prod_{\mathfrak{q}} \beta(\Lambda_{\mathfrak{q}})^{-1},$$

where $\mathfrak{q}$ runs over the primes of F and $\beta(\Lambda_{\mathfrak{q}})$ is the local density. Let $S \stackrel{\text{def}}{=} \{\mathfrak{q} : \mathfrak{q} \mid 2 \det(\Lambda)\}$, a set of at most $d/2 + \log_2 N_{K/\mathbb{Q}}(\det \Lambda)$ primes. For $\mathfrak{q} \notin S$ the lattice $\Lambda_{\mathfrak{q}}$ is unimodular and $\beta(\Lambda_{\mathfrak{q}})^{-1}$ equals the Euler factor at $\mathfrak{q}$ of $\prod_{i=1}^r L_F(\chi^i, i)$, with χ the quadratic character of K/F [Kir16, Prop. 4.2.7]. Completing this partial product to the full L -values over all primes multiplies it by $\prod_{\mathfrak{q} \in S} \prod_{i=1}^r (1 - \chi^i(\mathfrak{q}) N(\mathfrak{q})^{-i}) \leq 2^{r|S|}$. For $\mathfrak{q} \in S$ instead Lemma B.17 gives $\beta(\Lambda_{\mathfrak{q}})^{-1} \leq N(\mathfrak{q})^{(v_{\mathfrak{q}}(2)+1)r^2}$, and since $\prod_{\mathfrak{q} \mid 2} N(\mathfrak{q})^{v_{\mathfrak{q}}(2)+1} \leq 2^d$ and $\prod_{\mathfrak{q} \mid \det(\Lambda)} N(\mathfrak{q}) \leq N_{K/\mathbb{Q}}(\det \Lambda)$, the primes of S contribute at most $2^{O(r^2 d)} N_{K/\mathbb{Q}}(\det \Lambda)^{O(r^2)}$ in total. Here $L_F(\chi^i, i) \leq \zeta(i)^{d/2} = 2^{O(d)}$ for $i \geq 2$ and $L_F(\chi, 1) \leq (O(\log |\Delta_K|))^{d/2}$ [Lou00]. So

$$\text{Mass}(X_0)^{1/d} \leq \text{poly}(d) |\Delta_K|^{O(r^2)/d} N_{K/\mathbb{Q}}(\det \Lambda)^{O(r^2)/d}.$$

Finally $\det_{\mathbb{Q}}(\Lambda)^2 = |\Delta_K|^r N_{K/\mathbb{Q}}(\det \Lambda)$ with $N_{K/\mathbb{Q}}(\det \Lambda) \geq 1$, so both factors are $\det_{\mathbb{Q}}(\Lambda)^{O_r(1)/n}$, which gives the claim. $\square$