A search to distinguish reduction for the isomorphism problem on direct sum lattices

Daniël van Gent[1] and Wessel van Woerden[2]

[1] CWI, Amsterdam, the Netherlands
[2] PQShield, the Netherlands

Abstract. At Eurocrypt 2003, Szydlo presented a search to distinguish reduction for the Lattice Isomorphism Problem (LIP) on the integer lattice $\mathbb{Z}^n$. Here the search problem asks to find an isometry between $\mathbb{Z}^n$ and an isomorphic lattice, while the distinguish variant asks to distinguish between a list of auxiliary lattices related to $\mathbb{Z}^n$.
In this work we generalize Szydlo's search to distinguish reduction in two ways. Firstly, we generalize the reduction to any lattice isomorphic to Γ^n, where Γ is a fixed *base lattice*. Secondly, we allow Γ to be a *module lattice* over any number field. Assuming the base lattice Γ and the number field K are fixed, our reduction is polynomial in n.
As a special case we consider the module lattice $\mathcal{O}^2$ used in the module-LIP based signature scheme HAWK, and we show that one can solve the search problem, leading to a full key recovery, with less than $2d^2$ distinguishing calls on two lattices each, where d is the degree of the power-of-two cyclotomic number field and $\mathcal{O}$ its ring of integers.

1 Introduction

The search variant of the Lattice Isomorphism Problem (sLIP), which asks to find an isometry between two isomorphic lattices, has recently found applications on the constructive side of cryptology [9,4,14,7,1,2,12,19], highlighted by the efficient signature scheme HAWK [8]. However, the first occurrence of sLIP in cryptology was on the cryptanalytic side. In 2003, Szydlo showed that the leakage coming from NTRUsign or GGH signatures could be used to fully recover their secret keys, *if* one could solve sLIP for the integer lattice $\mathbb{Z}^n$ [23].

Szydlo proceeded by giving a reduction from sLIP on $\mathbb{Z}^n$, to a distinguishing variant of LIP on some auxiliary lattices, which was deemed easier to solve. This distinguishing variant ΔLIP asks, given lattices $\Lambda_1, \ldots, \Lambda_k$ and Λ to return an index b such that $\Lambda \cong \Lambda_b$, promised that such an index exists. As ΔLIP can be reduced to either sLIP, or to a decisional variant that simply asks if a pair of lattices is isomorphic, it can indeed be seen as the weakest variant of LIP. At the same time, ΔLIP is exceptionally useful in security proofs, as it allows one, in the security game, to replace the lattice used in a scheme by another *indistinguishable* lattice, as done in [9].

Szydlo had some heuristic ideas to solve the ΔLIP problem on these auxiliary lattices, but actually instantiating these results in attacks worse than directly

solving sLIP for $\mathbb{Z}^n$. This gives way for a new interpretation of Szydlo's work: the search to distinguish reduction gives evidence that ΔLIP on these auxiliary lattices is just as hard as sLIP on $\mathbb{Z}^n$, strengthening our understanding in the security of ΔLIP. In this work we extend this further, by generalizing Szydlo's search to distinguish reduction beyond the integer lattice $\mathbb{Z}^n$.

Szydlo's search to distinguish reduction for $\mathbb{Z}^n$. The main idea behind Szydlo's reduction is that the set of index 2 superlattices of $\mathbb{Z}^n$ is small *up to the many automorphisms* of $\mathbb{Z}^n$. Most importantly, the automorphism group of $\mathbb{Z}^n$ contains all (signed) coordinate permutations. Therefore, while there are $2^n - 1$ superlattices $\frac{1}{2}\boldsymbol{v}\mathbb{Z} + \mathbb{Z}^n \supset \mathbb{Z}^n$ of index 2, corresponding to nonzero binary vectors $\boldsymbol{v} \in \mathbb{Z}^n/2\mathbb{Z}^n$, they fall into only n isomorphism classes determined solely by the number of zero coordinates in $\boldsymbol{v}$ mod 2.

Any isometry $f : \Lambda \to \mathbb{Z}^n$ also induces a bijection $\frac{1}{2}\boldsymbol{v}\mathbb{Z} + \Lambda \mapsto \frac{1}{2}f(\boldsymbol{v})\mathbb{Z} + \mathbb{Z}^n$ between their sets of index-2 superlattices, which maps such a superlattice of Λ to an isomorphic one of $\mathbb{Z}^n$. By distinguishing which of the n possible superlattices of $\mathbb{Z}^n$ the lattice $\frac{1}{2}\boldsymbol{v}\mathbb{Z} + \Lambda$ is isomorphic to, we can therefore gain information about the number of zero coordinates in $f(\boldsymbol{v})$ mod 2.

Szydlo's reduction exploits this information further to recover an isometry 'modulo 2'. From there on the reduction proceeds inductively: using the information from an isometry modulo 2^k, one considers a specific small set of superlattices of index 2^{k+1}, and uses that information to build an isometry modulo 2^{k+1}. For large enough k, LLL-reduction can be used to recover the full isometry.

1.1 Contributions

In this work we generalize Szydlo's search to distinguishing reduction in two ways.

Firstly, instead of only considering the integer lattice $\mathbb{Z}^n$, we generalize the reduction to lattices Λ isomorphic to Γ^n, where Γ is a fixed *base lattice* of rank $r = \mathrm{rk}(\Gamma)$. A base lattice is any lattice whose minimal vectors generate a full-rank indecomposable sublattice, which is a technical condition that we need to have a good understanding of the superlattices and automorphisms of Γ^n. Examples of such base lattices are $\mathbb{Z}$, the highly symmetric root lattices A_r ($r \geqslant 2$), D_r ($r \geqslant 3$), E_6, E_7, E_8, Λ_{24} [6], and many more interesting low-dimensional lattices such as the large family of perfect lattices [20]. In the most costly step of the algorithm, the lattices we have to distinguish between are the superlattices of Λ corresponding to subspaces of $\Lambda/p\Lambda \cong \Gamma^n/p\Gamma^n$ of rank at most r. Because the base lattice Γ, and hence its rank r, are considered fixed, the number of such superlattices up to isomorphism is polynomial in n, resulting in a polynomial-time reduction.

Secondly, we generalize the reduction from classical lattices to $\mathcal{O}$-module lattices over any number field K with ring of integers $\mathcal{O}$ (simply referred to as *lattice* in this work). The prime p naturally gets generalized to any prime ideal $\mathfrak{p}$ of $\mathcal{O}$, and the auxiliary lattices correspond to rank r subspaces of the $(\mathcal{O}/\mathfrak{p})$-vector space $\Lambda/\mathfrak{p}\Lambda$. This leads to the following main result.

Theorem 7.3 (Main result). *Fix a number field with maximal order $\mathcal{O}$, a base $\mathcal{O}$-lattice Γ and a prime $\mathfrak{p}$ of $\mathcal{O}$, and suppose we have an oracle for $\mathcal{O}$-ΔLIP. Then there exists a polynomial-time algorithm that, given an $\mathcal{O}$-lattice Λ isomorphic to Γ^n for some n, computes such an isomorphism in $O(n^3)$ oracle calls on $O(1)$ lattices each and $O(n^2)$ oracle calls on $O(n^{e-1})$ lattices each, where $q = \mathrm{N}(\mathfrak{p})$, $r = \operatorname{rk}\Gamma$ and $e = q^{r^2}$.*

While the reduction we present is polynomial in n assuming $\mathcal{O}$, $\mathfrak{p}$ and Γ are fixed, the exponent can be rather large in terms of the rank of Γ and the norm of $\mathfrak{p}$. Therefore, in Section 8 we present some practical improvements to the exponent that make use of the automorphisms of Γ itself. For example for $\Gamma = A_2$, this allows us to reduce the total number of lattices in a batch to distinguish between from $O(n^{e-1}) = O(n^{15})$ down to $O(n)$, which is asymptotically the same as for $\Gamma = \mathbb{Z}$.

As a special case, we consider the lattice of HAWK [8], where K the power-of-two cyclotomic field of degree $d = 2^{\ell-1}$ and the lattice is $\mathcal{O}^2$, which as a $\mathbb{Z}$-lattice is isomorphic to $\mathbb{Z}^{2d}$. Along with the Gentry–Szydlo algorithm, we obtain a search to distinguish reduction that requires only $O(d^2)$ distinguishing oracle calls on 2 lattices each.

Theorem 9.3. *There is a polynomial-time algorithm that, given the order $\mathcal{O} = \mathbb{Z}[\zeta_{2^\ell}]$ of degree $d = 2^{\ell-1}$ for some $\ell > 0$, an $\mathcal{O}$-lattice Λ isomorphic to $\mathcal{O}^2$, and an oracle for $\mathcal{O}$-ΔLIP, computes an isomorphism $\Lambda \to \mathcal{O}^2$ using at most $4^{\ell-1} - \ell < 2d^2$ oracle calls on 2 lattices each.*

Our algorithms are rigorous and specialize to those of Szydlo, including the time complexity in the reduction, when taking $\Gamma = \mathbb{Z}$. In particular, we provide in Section 5 a non-trivial algorithm to choose which oracle calls to make in the first step of the algorithm, which was left to the reader in [23].

1.2 Open Questions

In this work we present a generalization of Szydlo's search to distinguish reduction to (module) lattices of the form Γ^n under some technical conditions on Γ. This generalization is possible because Γ^n has a large automorphism group whose structure we understand well. Given that a large automorphism group is a requirement to limit the number of orbits of small index superlattices, one could wonder if there are other families of highly symmetric lattices for which one could make such a reduction work. For example, it should take little effort to generalize our main result to lattices Λ isomorphic to $\bigoplus_{\Gamma\in\mathcal{B}} \Gamma^{n_\Gamma}$ for some fixed finite set $\mathcal{B}$ of base lattices, although this would hardly be more insightful. It could be of interest to consider and optimize the reductions both in the asymptotic sense for large dimensional highly symmetric lattice families, or simply for specific highly symmetric lattices such as E_8 or the Leech lattice.

We leave further optimizations of our reduction as another open question, in particular those that could reduce the polynomial exponent further in terms of

the rank r of the base lattice. The large exponent is mainly because in the worst-case we have to distinguish between a large set of (overlattices corresponding to the) subspaces of rank r. One possibility to reduce the exponent is using either an average-case analysis to show that such oracle calls are needed rarely, or a rerandomization step that could avoid the case altogether.

1.3 Overview

We give here a short overview of this work. The necessary preliminaries on (module-)lattices, number theory, and the lattice isomorphism problems are discussed in Section 2. In Section 3 we define the family of lattices and their superlattices that occur in our reduction. In Section 4 we consider the inductive step, that lifts information about an isometry modulo $\mathfrak{a}^k\Lambda$ to information modulo $\mathfrak{a}^{k+1}\Lambda$. The more complex base case will follow in Section 6 after some technical results in Section 5. In Section 7 the inductive step and the base case are combined to give our main result. Following this we consider some practical improvements to our main result in Section 8, and we consider the special case of HAWK in Section 9.

1.4 Acknowledgments

The first author is supported by the Quantum Software Consortium project (file number 024.003.037) of the Zwaartekracht research programme which is (partly) financed by the Dutch Research Council (NWO). The second author was supported by the CHARM ANR-NSF grant (ANR-21-CE94-0003). A significant part of the work on this paper was executed while the second author was under employment at IMB in Bordeaux. The authors would also like to thank Alice Pellet–Mary for insightful discussions.

2 Preliminaries

2.1 Notation

Vectors $\boldsymbol{x}, \boldsymbol{y}$ are denoted in lower-case bold and should be interpreted as column vectors. Matrices $\boldsymbol{A}, \boldsymbol{B}$ are denoted in upper-case bold. We write $\mathbb{1}$ for the *indicator function*, which maps a proposition to 1 if it is true and to 0 if it is false.

2.2 Lattices

A $\mathbb{Z}$-*lattice* Λ is a discrete subgroup of $\mathbb{R}^\ell$. A $\mathbb{Z}$-lattice is always of the form $\Lambda = \boldsymbol{b}_1\mathbb{Z} + \ldots + \boldsymbol{b}_k\mathbb{Z}$ for some $\mathbb{R}$-linearly independent *basis* $\boldsymbol{b}_1, \ldots, \boldsymbol{b}_k \in \mathbb{R}^\ell$. We call the number of basis vectors $k \leqslant \ell$ the *rank* or *dimension* of the lattice. If $k = \ell$ we say that the $\mathbb{Z}$-lattice has *full rank*. The *volume* of a $\mathbb{Z}$-lattice Λ is given

by the (co)volume $\mathrm{Vol}(\Lambda) = \mathrm{Vol}_{\mathbb{R}}(\mathrm{span}\,\Lambda/\Lambda)$. The i-th minimum $\lambda_i(\Lambda)$ of a $\mathbb{Z}$-lattice Λ is the minimum radius $\lambda > 0$ such that its vectors of Euclidean length at most λ span a subspace of rank at least i. In particular, the *first minimum* is given by $\lambda_1(\Lambda) = \min_{\boldsymbol{v}\in\Lambda\setminus\{0\}} \|\boldsymbol{v}\|_2$.

Remark 2.1. Throughout this work we will explicitly refer to these classical lattices as $\mathbb{Z}$-lattices, and use the general term *lattice* for out later to be defined module lattices (which contain the classical lattices as a special case).

2.3 Number fields

Throughout this document we fix some number field $K = \mathbb{Q}[X]/(P)$ for a monic irreducible polynomial $P \in \mathbb{Q}[X]$ of degree d. We denote $\mathcal{O}$ for the ring of integers of K, i.e., the unique maximal order in K. A number field K of degree d comes with d distinct embeddings $\sigma : K \to \mathbb{C}$, where X is mapped to a (complex) root of P. We denote the set of these d embeddings by Σ. For every embedding $\sigma \in \Sigma$ of K, its conjugation $\overline{\sigma}(x) := \overline{\sigma(x)}$ is also an embedding of K. We call σ a *real embedding* if $\sigma(K) \subset \mathbb{R}$, or equivalently if $\overline{\sigma} = \sigma$, and otherwise we call σ a *complex embedding*. We denote $\boldsymbol{\sigma} : K \to \mathbb{C}^{\Sigma} \cong \mathbb{C}^d$ for the concatenation $\boldsymbol{\sigma}(x) := (\sigma(x))_{\sigma\in\Sigma}$ of all embeddings, also known as the *canonical embedding* of K.

The space $K_{\mathbb{R}}$. Note that $\boldsymbol{\sigma}(K)$ is a $\mathbb{Q}$-vector subspace of $\mathbb{C}^{\Sigma}$ of rank d. We can extend this to an $\mathbb{R}$-vector space by considering $K_{\mathbb{R}} := K \otimes \mathbb{R}$ and by naturally extending the domain of $\boldsymbol{\sigma}$. Hereby $\boldsymbol{\sigma} : K_{\mathbb{R}} \to \mathbb{C}^{\Sigma}$ becomes a ring-homomorphism whose image $\boldsymbol{\sigma}(K_{\mathbb{R}})$ is the d-dimensional real subspace

$$\mathcal{H} = \{(y_\sigma)_\sigma \in \mathbb{C}^{\Sigma} : y_{\overline{\sigma}} = \overline{y_\sigma} \text{ for } \sigma \in \Sigma\}$$

of $\mathbb{C}^{\Sigma}$ interpreted as a real vector space of dimension $2d$. We extend the (absolute) *trace* and *norm* map of K to $K_{\mathbb{R}}$ by writing $\mathrm{Tr}(x) := \sum_{\sigma\in\Sigma} \sigma(x) \in \mathbb{R}$ and $\mathrm{N}(x) := \prod_{\sigma\in\Sigma} \sigma(x) \in \mathbb{R}$ respectively. We have $\mathrm{Tr}(x), \mathrm{N}(x) \in \mathbb{Q}$ (resp. $\mathbb{Z}$) for $x \in K$ (resp. $\mathcal{O}_K$).

Note that $\mathcal{H}$ is closed under complex conjugation and therefore naturally defines a conjugation map on $K_{\mathbb{R}}$. For $x \in K_{\mathbb{R}}$, its conjugate $\overline{x} \in K_{\mathbb{R}}$ is the unique element such that $\sigma(\overline{x}) = \overline{\sigma(x)}$. We extend the canonical embedding $\boldsymbol{\sigma}$ coefficient-wise on $K_{\mathbb{R}}^{\ell}$ and define two $\mathbb{R}$-bilinear forms on this space. Firstly, we consider the form on $K_{\mathbb{R}}^{\ell}$ given by

$$\langle \boldsymbol{x}, \boldsymbol{y} \rangle_K := \boldsymbol{x}^* \boldsymbol{y} = \sum_{i=1}^{\ell} \overline{\boldsymbol{x}_i} \boldsymbol{y}_i \in K_{\mathbb{R}} \text{ for all } \boldsymbol{x}, \boldsymbol{y} \in K_{\mathbb{R}}^{\ell}.$$

With respect to the conjugation just defined on $K_{\mathbb{R}}$, this form is, by slight abuse of terminology, a *$K_{\mathbb{R}}$-Hermitian inner product*. Secondly, we consider the form on $K_{\mathbb{R}}^{\ell}$ induced by the canonical embedding $\boldsymbol{\sigma}$,

$$\langle \boldsymbol{x}, \boldsymbol{y} \rangle_{\mathbb{Q}} := \langle \boldsymbol{\sigma}(\boldsymbol{x}), \boldsymbol{\sigma}(\boldsymbol{y}) \rangle = \mathrm{Tr}(\langle \boldsymbol{x}, \boldsymbol{y} \rangle_K) \in \mathbb{R} \text{ for all } \boldsymbol{x}, \boldsymbol{y} \in K_{\mathbb{R}}^{\ell}.$$

Without any abuse of terminology, this is a (real) inner product, also known as the *trace inner product*. Note that for $K = \mathbb{Q}$ we have $\mathcal{O} = \mathbb{Z}$, $K_{\mathbb{R}} = \mathbb{R}$, and both forms fall back to the standard Euclidean inner product.

Modules. Modules can be seen as a generalization of vector spaces, where the field is replaced by a ring. Let R be a ring and let M be an additively written abelian group. Suppose we have a multiplication operation $\cdot : R \times M \to M$ that is both distributive (over both additions) and compatible with the ring multiplication, i.e., $(rs)\cdot x = r\cdot(s\cdot x)$. Then we call M with such a multiplication operation an R-module. If $M' \subset M$ is also an R-module (under the restricted multiplication operation of M) we call M' an *R-submodule* of M. A *generating set* of M is a subset $S \subseteq M$ such that $M = \{\sum_{m\in S} r_m \cdot m : r \in R^S\}$, and we say that M is *finitely generated* if it has a finite generating set. A subset $S \subseteq M$ is *linearly independent* if the only $x \in R^S$ with $\sum_{m\in S} r_m \cdot m = 0$ is the zero vector. A finite generating set that is also linearly independent is a *basis*, and we call a module *free* if it has a basis. As an example one can consider the free R-module $M = R^\ell$ with the natural coefficient-wise multiplication operation $r\cdot(x_1,\ldots,x_\ell) := (rx_1,\ldots,rx_\ell)$. Similarly, any $\mathbb{Z}$-lattice $\Lambda \subset \mathbb{R}^\ell$ is a $\mathbb{Z}$-submodule of $\mathbb{R}^\ell$ and $\mathcal{O}^\ell$ is an $\mathcal{O}$-submodule of K^ℓ.

Ideals. A *fractional ideal* of $\mathcal{O}$ is a non-zero finitely generated $\mathcal{O}$-submodule of K. An *integral ideal* of $\mathcal{O}$ is a fractional ideal of $\mathcal{O}$ contained in $\mathcal{O}$. A fractional ideal generated by a single element of K is called *principal*. The product $I \cdot J$ of two fractional ideals I, J of $\mathcal{O}$ is the smallest $\mathcal{O}$-submodule of K containing all products xy for $x \in I$ and $y \in J$, which is again a fractional ideal. Because $\mathcal{O}$ is not just any order but a Dedekind domain, there exists for each fractional ideal I a unique fractional ideal, which we call its *inverse* and denote I^{-1}, such that $I \cdot I^{-1} = \mathcal{O}$. The *norm* of an integral ideal $\mathfrak{a}$ is $\mathrm{N}(\mathfrak{a}) = |\mathcal{O}/\mathfrak{a}|$, and we extend this definition multiplicatively to fractional ideals. For a principal ideal we have $\mathrm{N}(a\mathcal{O}) = |\mathrm{N}(a)|$ for $a \in K$. A *prime* of $\mathcal{O}$ is an integral ideal $\mathfrak{p}$ such that $ab \in \mathfrak{p}$ implies $a \in \mathfrak{p}$ or $b \in \mathfrak{p}$. Note that this definition excludes $\mathfrak{p} = 0$. For primes $\mathfrak{p}$, the quotient ring $\mathcal{O}/\mathfrak{p}$ is a finite field.

2.4 Module lattices

Informally, a module lattice is a module over some order along with a positive definite bilinear form that determines its geometry. For a general reference on module lattices we refer to [5, Chapter 1]. In this work we consider $\mathcal{O}$-submodules of $K_{\mathbb{R}}^\ell$ of the form $\Lambda = \boldsymbol{b}_1\mathfrak{a}_1 + \ldots + \boldsymbol{b}_k\mathfrak{a}_k \subset K_{\mathbb{R}}^\ell$ where $\boldsymbol{b}_1,\ldots,\boldsymbol{b}_k$ are $K_{\mathbb{R}}$-linearly-independent vectors and $\mathfrak{a}_1,\ldots,\mathfrak{a}_k$ are integral ideals. Here $K_{\mathbb{R}}$ linear independence means that there are no $x_1,\ldots,x_k \in K_{\mathbb{R}}$ not all zero such that $\sum_{i=1}^k x_i\boldsymbol{b}_i = 0$. We call $k \leqslant \ell$ the *(module) rank* of Λ. The rank is independent of representation: the $\mathcal{O}$-module Λ is locally a free module of constant rank k. For the geometric properties of our module-lattice we will in this work always consider the Hermitian inner product $\langle \boldsymbol{x}, \boldsymbol{y}\rangle_K$ for any $\boldsymbol{x}, \boldsymbol{y} \in \Lambda$. Beyond these preliminaries we will refer to these $\mathcal{O}$-module lattices simply by the term *lattice*

as we assume the number field K is fixed. We will explicitly refer to classical ($\mathbb{Z}$-module) lattices by the term $\mathbb{Z}$-lattice.

Example 2.2. Since $K_{\mathbb{R}}$ is generally not a field, $K_{\mathbb{R}}$-linear independence of vectors $\boldsymbol{b}_1, \ldots, \boldsymbol{b}_k \in K_{\mathbb{R}}^{\ell}$ is meaningful even when $k = 1$. Take $K = \mathbb{Q}(\sqrt{2})$, so that $K_{\mathbb{R}} \cong \mathbb{R}^2$ as rings. Then the single vector $(1, 0) \in K_{\mathbb{R}}^1$ is not $K_{\mathbb{R}}$-linearly-independent, since there is a non-zero scalar $(0, 1) \in K_{\mathbb{R}}$ such that $(0, 1) \cdot (1, 0) = 0$.

Via the trace inner product one can also interpret a ($\mathcal{O}$-)lattice $\Lambda \subset K_{\mathbb{R}}^{\ell}$ of rank k as a $\mathbb{Z}$-lattice $\boldsymbol{\sigma}(\Lambda) \subset \mathcal{H}^{\ell} \subset \mathbb{C}^{dl} \cong \mathbb{R}^{2dl}$ of rank dk. We therefore write $\mathrm{Vol}(\Lambda)$ or $\lambda_i(\Lambda)$ for $\mathrm{Vol}(\boldsymbol{\sigma}(\Lambda))$ and $\lambda_i(\boldsymbol{\sigma}(\Lambda))$ respectively. Generally however, we will refrain from going down to the underlying $\mathbb{Z}$-lattice and its geometry, and instead generalize the properties we need with respect to the Hermitian inner product $\langle ., . \rangle_K$, which is required for some of our proofs.

Definition 2.3 (First minimum $\lambda_1^{\mathcal{O}}$). *For a rank-1 lattice Φ consider*

$$\sum_{\mathbf{x} \in \Phi} \mathbb{Z} \cdot \mathrm{N}(\langle \mathbf{x}, \mathbf{x} \rangle_K).$$

It is a $\mathbb{Z}$-submodule of $\mathbb{R}$ of rank 1, and thus of the form $\mathbb{Z}r$ for some $r \in \mathbb{R}_{>0}$. We write $\mathrm{N}(\Phi) = \sqrt{r}$. For a lattice Λ write $\min(\Lambda)$ for the set of rank-1 sublattices $\Phi \subseteq \Lambda$ for which $\mathrm{N}(\Phi)$ is minimal. We write $\lambda_1^{\mathcal{O}}(\Lambda) = \min_{\Phi \subset \Lambda} \mathrm{N}(\Phi)^{1/d}$ where d is the degree of K.

Note that if $\Phi = \mathcal{O}\boldsymbol{x}$ for some $\boldsymbol{x} \in \Lambda \backslash \{0\}$, then $\mathrm{N}(\Phi) = \mathrm{N}(\boldsymbol{x})$. If we interpret any fractional ideal of $\mathcal{O}$ as a rank-1 lattice, then the map N in Definition 2.3 and the map N defined in the previous section agree. Recall that generally the rank-1 submodules need not be of the form $\mathcal{O}\boldsymbol{x}$, since $\mathcal{O}$ need not be a principal ideal domain like $\mathbb{Z}$. If $K = \mathbb{Q}$, then the rank-1 submodules of Λ are of the form $\mathbb{Z}\boldsymbol{x}$ for $\boldsymbol{x} \in \Lambda \backslash \{0\}$. The ones of those contained in $\min(\Lambda)$ are precisely those for which $\boldsymbol{x}$ is a shortest vector in the lattice, and $\lambda_1^{\mathbb{Z}}$ is the classical first minimum.

Lemma 2.4. *Let Φ be a rank 1-lattice and I be a fractional ideal of $\mathcal{O}$. Then $\mathrm{N}(I\Phi) = \mathrm{N}(I) \cdot \mathrm{N}(\Phi)$. In particular, $\lambda_1^{\mathcal{O}}(I\Lambda) = \mathrm{N}(I)^{1/d} \cdot \lambda_1^{\mathcal{O}}(\Lambda)$ for any $\mathcal{O}$-module lattice Λ, where d is the degree of K.*

Proof. It suffices to show that

$$\sum_{\boldsymbol{x} \in I\Phi} \mathbb{Z} \cdot \mathrm{N}(\langle \boldsymbol{x}, \boldsymbol{x} \rangle_K) = \mathrm{N}(I)^2 \cdot \sum_{\boldsymbol{x} \in \Phi} \mathbb{Z} \cdot \mathrm{N}(\langle \boldsymbol{x}, \boldsymbol{x} \rangle_K).$$

We may prove this equality of $\mathbb{Z}$-modules locally at the primes of $\mathbb{Z}$. Let p be such a prime and write $-_p$ for the localization. Then $\mathcal{O}_p$ is semi-local, hence $I = \mathcal{O}_p \cdot a$ for some $a \in I$. Therefore $I\Phi_p = a\Phi_p$ and

$$\mathbb{Z}_p \cdot \mathrm{N}(\langle a\boldsymbol{x}, a\boldsymbol{x} \rangle_K) = \mathbb{Z}_p \cdot \mathrm{N}(a\overline{a}\langle \boldsymbol{x}, \boldsymbol{x} \rangle_K) = \mathrm{N}(a)^2 \cdot \mathbb{Z}_p \cdot \mathrm{N}(\langle \boldsymbol{x}, \boldsymbol{x} \rangle_K),$$

from which the lemma follows. □

Definition 2.5. *For a finite collection of lattices $\{\Lambda_i\}_{i\in I}$ write $\bigoplus_{i\in I}\Lambda_i$ for their* orthogonal sum. *We simply write $\Lambda_1\oplus\Lambda_2$ for $\bigoplus_{i\in\{1,2\}}\Lambda_i$ and $\Lambda^n=\bigoplus_{i\in\{1,\dots,n\}}\Lambda$ for the n-fold orthogonal sum.*

The norm-like map N we defined for rank-1 lattices interacts as expected with orthogonal projections, e.g., non-trivial projections decrease the norm.

Lemma 2.6. *Let $\Lambda=\Lambda_1\oplus\Lambda_2$ be a lattice and consider the orthogonal projection $\pi:\Lambda\to\Lambda$ with image Λ_1. If $\Phi\subseteq\Lambda$ is a rank-1 sublattice, then $\mathrm{N}(\pi(\Phi))\leqslant\mathrm{N}(\Phi)$, with equality if and only if $\Phi\subseteq\Lambda_1$.*

Proof. First assume that $\Phi=\mathcal{O}\boldsymbol{x}$ for some $\boldsymbol{x}\in\Lambda$. Then

$$\mathrm{N}(\pi(\Phi))^2=\mathrm{N}(\langle\pi(\boldsymbol{x}),\pi(\boldsymbol{x})\rangle_K)\leqslant\mathrm{N}(\langle\boldsymbol{x},\boldsymbol{x}\rangle_K)=\mathrm{N}(\Phi)^2,$$

where the inequality holds because

$$\sigma(\langle\pi(\boldsymbol{x}),\pi(\boldsymbol{x})\rangle_K)\leqslant\sigma(\langle\pi(\boldsymbol{x}),\pi(\boldsymbol{x})\rangle_K)+\sigma(\langle\boldsymbol{x}-\pi(\boldsymbol{x}),\boldsymbol{x}-\pi(\boldsymbol{x})\rangle_K)=\sigma(\langle\boldsymbol{x},\boldsymbol{x}\rangle_K)$$

holds for each $\sigma:K\to\mathbb{C}$. For general Φ, there exists some fractional ideal I such that $I\Phi=\mathcal{O}\boldsymbol{x}$ for some $\boldsymbol{x}$. The lemma then follows from the previous combined with Lemma 2.4. □

2.5 Module-LIP.

Let $\Lambda\subset K_{\mathbb{R}}^{\ell}$ and $\Lambda'\subset K_{\mathbb{R}}^{\ell'}$ be two lattices for some $\ell,\ell'\geqslant 0$. An *isometry* between Λ and Λ' is an $\mathcal{O}$-linear isomorphism $f:\Lambda\to\Lambda'$ such that

$$\langle f(\boldsymbol{x}),f(\boldsymbol{y})\rangle_K=\langle\boldsymbol{x},\boldsymbol{y}\rangle_K\text{ for all }\boldsymbol{x},\boldsymbol{y}\in\Lambda.$$

If there exists an isometry between two lattices, then we call them *isomorphic*. Note that we do not require that the lattices live in an ambient space of the same dimension. We denote the set of all isometries $f:\Lambda\to\Lambda'$ by $\mathrm{Isom}(\Lambda,\Lambda')$ and the set of all isometries $f:\Lambda\to\Lambda$, also called *automorphisms* of Λ, by $\mathrm{Aut}(\Lambda)$. Note that $\mathrm{Aut}(\Lambda)$ is a group under composition. If $f:\Lambda\to\Lambda'$ is an isometry, then $g\mapsto f\circ g\circ f^{-1}$ is a group isomorphism $\mathrm{Aut}(\Lambda)\to\mathrm{Aut}(\Lambda')$, and $\mathrm{Isom}(\Lambda,\Lambda')=f\circ\mathrm{Aut}(\Lambda)=\{f\circ g:g\in\mathrm{Aut}(\Lambda)\}$ and $\mathrm{Isom}(\Lambda,\Lambda')=\mathrm{Aut}(\Lambda')\circ f$.

We now define the two main lattice isomorphism problems that play a central role in this work. We defer details about input and output representation of these problems to the next subsection.

Definition 2.7 (Search-LIP (sLIP)). *The* sLIP *problem asks, given two isomorphic lattices, to compute an isometry between them.*

Definition 2.8 (Distinguish-LIP (ΔLIP)). *The* ΔLIP *problem asks, given lattices $\Lambda_1,\dots,\Lambda_k$ for some $k\geqslant 1$ and a lattice Λ that is isomorphic to at least one of the Λ_b, to compute some b such that $\Lambda\cong\Lambda_b$.*

We now consider the specific sLIP problem that is of interest to us.

Definition 2.9 (Γ-LIP). *Fix a non-zero lattice Γ. The Γ-LIP problem asks, given a lattice Λ isomorphic to Γ^n for some n, to compute an isometry $\Lambda \to \Gamma^n$.*

For the $\mathbb{Z}$-lattice $\Gamma = \mathbb{Z}$ we obtain the familiar $\mathbb{Z}$-LIP problem. As Szydlo showed, there is a many-to-one reduction from $\mathbb{Z}$-LIP to ΔLIP, which we will generalize in this work.

Remark 2.10. Note that Definition 2.8 of ΔLIP allows for some (or all) of the input lattices to be isomorphic. In the case that Λ' is isomorphic to multiple Λ_b the definition allows to return any of them. Additionally, the problem is worst-case in the sense that there is no prior distribution on which lattice Λ' is isomorphic to, and there is no notion of a distinguishing advantage. In particular, we will assume that an oracle solving ΔLIP will always return a correct index. Note that contrary to other distinguishing problems in cryptography (for example for LWE) this is a reasonable definition as there is no natural overlap between the cases, i.e., lattices are isomorphic or they are not (and they can't be isomorphic with some non-trivial probability).

Remark 2.11. Note that the difficulty of the ΔLIP problem can depend on the number of lattices one has to distinguish between. Therefore, we will quantify this number in our reduction. In particular, for our reduction applied to lattices isomorphic to Γ^n for some fixed Γ, it will always be polynomial in n

2.6 Representation of objects

Here we shortly discuss the representation of our objects, which is needed for the actual computations. For the sake of a clear presentation we will mostly abstract away from this in the rest of this work, but the representations and tools here are required to make the typical computations efficient.

We represent the field K by some irreducible polynomial $P \in \mathbb{Q}[X]$ such that $K \cong \mathbb{Q}[X]/(P)$. We represent the maximal order $\mathcal{O}$ of K by a $\mathbb{Z}$-basis $R = [\boldsymbol{r}_1, \ldots, \boldsymbol{r}_d] \in K^d$ that is LLL-reduced with respect to the trace inner product. Elements of $\mathcal{O}$ and K are represented by integer and rational vectors with respect to R respectively. Integral ideals of $\mathcal{O}$ are represented by an integer basis in Hermite Normal Form with respect to the basis of $\mathcal{O}$. Fractional ideals I are represented by an integral ideal $\mathfrak{a}$ along with a scalar $b \in \mathcal{O}$ such that $I = \mathfrak{a}/b$.

Ignoring the representation of elements in $K_\mathbb{R}$ for now, lattices $\Lambda \subset K_\mathbb{R}^\ell$ can be represented by a *pseudo-basis* given by $K_\mathbb{R}$ linearly-independent vectors $\boldsymbol{b}_1, \ldots, \boldsymbol{b}_k \in K_\mathbb{R}^\ell$ and (integral) ideals $\mathfrak{a}_1, \ldots, \mathfrak{a}_k$ such that $\Lambda = \sum_{i=1}^k \boldsymbol{b}_i \mathfrak{a}_i$. More natural for LIP however is to represent a pseudo-basis up to an isometry by considering a *pseudo-Gram matrix* $G = ((\langle \boldsymbol{b}_i, \boldsymbol{b}_j \rangle_K)_{i,j}, (\mathfrak{a}_i)_i)$. Lattice vectors $\boldsymbol{x} \in \Lambda$ are then simply represented by a vector $(x_i)_i \in K^k$ with $x_i \in \mathfrak{a}_i$. Only for metric information we have to resolve to computations in $K_\mathbb{R}$ by $\langle \boldsymbol{x}, \boldsymbol{y} \rangle_K = \sum_i \sum_j \overline{x_i} y_j \langle \boldsymbol{b}_i, \boldsymbol{b}_j \rangle_K$. For the latter we assume that $(\langle \boldsymbol{b}_i, \boldsymbol{b}_j \rangle_K)_{i,j}$, which is given as the input representation, is part of some subring of $K_\mathbb{R}$ that

contains K, is closed under conjugation, and that allows for efficient representation and standard computations[3]. For example when $\Lambda \subset K^\ell$ one could consider the symbolic subring $\overline{K}K \subset K_{\mathbb{R}}$ generated by all pairs $\overline{\boldsymbol{r}_i} \cdot \boldsymbol{r}_j$.

In this setting, (module) LIP and isometries can then be fully phrased in terms of and represented by a basis transformation $U \in \mathrm{GL}_k(K)$ with some additional constraints only depending on the ideals in the pseudo-Gram representation (besides the inner product preserving condition). We refer to [22] for a much more extensive overview about these representations, and the equivalence for LIP between the pseudo-basis and pseudo-Gram point of view.

3 Indecomposable and p-ary lattices

Recall that we have fixed a number field K with maximal order $\mathcal{O}$, and that by lattice we mean $\mathcal{O}$-module lattice. In this section we define the family of lattices of the form Γ^n that our reduction works for and for each the family of superlattices of size $\mathrm{poly}(n)$ that our algorithm will make distinguishing calls between.

For our algorithm to work without additional issues we do not want the lattice Γ^n to have any unexpected additional automorphisms. The simplest way to achieve this is by assuming that Γ is indecomposable. Recall that we have defined orthogonality with respect of the Hermitian inner product $\langle ., . \rangle_K$.

Definition 3.1. *Let Λ be a lattice. An* orthogonal summand *of Λ is a sublattice $\Lambda_1 \subseteq \Lambda$ to which there exists an orthogonal sublattice $\Lambda_2 \subseteq \Lambda$ such that $\Lambda_1 \oplus \Lambda_2 = \Lambda$. We say Λ is* indecomposable *if it has exactly two orthogonal summands, or equivalently, $\Lambda \neq 0$, and 0 and Λ are the only orthogonal summands. Write $I(\Lambda)$ for the set of indecomposable orthogonal summands of Λ.*

Lemma 3.2. *Let I be a fractional ideal of $\mathcal{O}$ and Λ be a lattice. Then Λ is indecomposable if and only if $I\Lambda$ is indecomposable.*

Proof. If $\Lambda = \Lambda_1 \oplus \Lambda_2$, then $I\Lambda = I\Lambda_1 \oplus I\Lambda_2$. In particular, if $I\Lambda$ is indecomposable, then so is Λ. The converse follows from the previous argument applied to I^{-1} and $I\Lambda$ in place of I and Λ. □

We now extend a classical result by Eichler [10, Satz 2]. The case that K is a CM-field was already treated by [13, Theorem 3.4].

Theorem 3.3. *Each lattice is the orthogonal sum of its indecomposable orthogonal summands.*

Proof. Let V be the set of indecomposable orthogonal summands of Λ when interpreted as a $\mathbb{Z}$-lattice, and for each $\Gamma \in V$ let π_Γ be the corresponding projection. Because the theorem holds in the classical case by Eichler [10, Satz 2], every element of $I(\Lambda)$ is of the form $\Lambda_S = \sum_{\Gamma \in S} \Gamma$ for some $S \subseteq V$. Let

[3] Note that one can even have these technical representation problems for classical $\mathbb{Z}$-lattices $\Lambda \subset \mathbb{R}^n$ when their Gram matrix is not rational.

$E = \{\{\Gamma_1, \Gamma_2\} \in \binom{V}{2} \mid \pi_{\Gamma_1}(\mathcal{O}\Gamma_2) + \pi_{\Gamma_2}(\mathcal{O}\Gamma_1) \neq 0\}$. Then (V, E) is a graph, and let C be its set of connected components. Note that Λ_S is an $\mathcal{O}$-sublattice of Λ if and only if S is a union of connected components. Since $\Lambda_{\cup X} = \bigoplus_{c \in X} \Lambda_c$ for any $X \subseteq C$, it follows that $I(\Lambda) = \{\Lambda_c \mid c \in C\}$, from which the theorem follows. □

For an indecomposable lattice Γ, we can directly relate the automorphism group of Γ^n to that of Γ.

Corollary 3.4. *Let Γ be an indecomposable lattice and write* Sym_n *for the symmetric group on $n \in \mathbb{Z}_{\geqslant 0}$ symbols. Then we have a bijection* $\mathrm{Aut}(\Gamma)^n \times \mathrm{Sym}_n \to \mathrm{Aut}(\Gamma^n)$ *given by*

$$(f_1, \ldots, f_n, \tau) \mapsto \Big[(\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n) \mapsto (f_1(\boldsymbol{x}_{\tau^{-1}(1)}), \ldots, f_n(\boldsymbol{x}_{\tau^{-1}(n)}))\Big].$$

Proof. It is easy to see that the map is injective. Conversely, each isometry g of Γ^n must respect the set of indecomposable orthogonal summands of Γ^n, which by Theorem 3.3 are precisely the n copies of Γ. Hence g is a permutation of the coordinates followed by a coordinate-wise isometry. □

We will generalize the results of Szydlo to lattices isomorphic to Γ^n, as opposed to $\mathbb{Z}^n$, for certain lattices Γ. In the remainder of this section we will specify the conditions on Γ and derive the properties of Γ we will use.

Definition 3.5. *A* base lattice *is a lattice Γ for which the sublattice*

$$\sum_{\Phi \in \min(\Gamma)} \Phi$$

*generated by its minimal rank-*1 *sublattices is of full rank and indecomposable. For a base lattice Γ and integral ideal $\mathfrak{a}$ of $\mathcal{O}$, an $\mathfrak{a}$-*ary Γ-lattice *is a sublattice $\mathfrak{a}\Gamma^n \subseteq B \subseteq \Gamma^n$ for some $n \geqslant 0$.*

Note that $\mathcal{O}$ is a base lattice, as $\min(\mathcal{O}) = \{\mathcal{O}\}$. For $\mathcal{O} = \mathbb{Z}$ and $n \geqslant 2$ and $n \geqslant 3$ respectively, the root lattices

$$A_n = \Big\{\boldsymbol{x} \in \mathbb{Z}^{n+1} \,\Big|\, \sum_i x_i = 0\Big\} \quad \text{and} \quad D_n = \Big\{\boldsymbol{x} \in \mathbb{Z}^n \,\Big|\, \sum_i x_i \equiv 0 \pmod 2\Big\}$$

are also base lattices: both are generated by their shortest vectors and are indecomposable.[4] The lattice D_n is also an example of a $2\mathbb{Z}$-ary $\mathbb{Z}$-lattice. More generally, any *perfect lattice* [20] is a base lattice, which includes all lattices that reach an optimal packing density.

We think of base lattices as being 'strongly indecomposable'. In particular, they are indecomposable.

[4] The root lattice D_2 is a rare exception here as it attains an orthogonal basis given by the vectors $(1, 1)^\top$ and $(1, -1)^\top$. However, it is isomorphic to $\sqrt{2}\mathbb{Z}^2$ so one can use the reduction for $\Gamma = \sqrt{2}\mathbb{Z}$.

Lemma 3.6. *A base lattice is indecomposable.*

Proof. Suppose $\Gamma_0 = \Gamma_1 \oplus \Gamma_2$ with $\Gamma_1, \Gamma_2 \neq 0$, and write $\Gamma_i' = \sum_{\Phi \in \min \Gamma_i} \Phi$. It follows from Lemma 2.6 that $\min(\Gamma_0) \subseteq \min(\Gamma_1) \sqcup \min(\Gamma_2)$. If $\lambda_1^{\mathcal{O}}(\Gamma_1) \neq \lambda_1^{\mathcal{O}}(\Gamma_2)$, say $\lambda_1^{\mathcal{O}}(\Gamma_1) < \lambda_1^{\mathcal{O}}(\Gamma_2)$, then $\Gamma_0' \subseteq \Gamma_1'$ is not of full rank in Γ_0, and Γ_0 is not a base lattice. Otherwise, we have a decomposition $\Gamma_0' = \Gamma_1' \oplus \Gamma_2'$, so again Γ_0 is not a base lattice. □

A desirable property of Γ would be that for each $\mathfrak{a}$-ary Γ-lattice $\mathfrak{a}\Gamma^n \subseteq B \subseteq \Gamma^n$, every automorphism of B is a restriction of an automorphism of Γ^n, which by Corollary 3.4 means that the 'coordinate system' of B is preserved. The following example shows that this need not be the case.

Example 3.7. Let $(\mathcal{O}, \mathfrak{a}, \Gamma) = (\mathbb{Z}, 2\mathbb{Z}, D_8)$. Consider the lattice E_8 given by

$$E_8 = \Big\{ (x_0, \dots, x_7) \in \mathbb{Z}^8 \;\Big|\; \sum_{i=0}^{7} x_i \equiv 2x_0 \equiv 2x_1 \equiv \cdots \equiv 2x_7 \pmod 4 \Big\}.$$

Note that $\mathfrak{a}D_8 \subseteq \mathbb{Z} \cdot (1, \dots, 1) + \mathfrak{a}D_8 = E_8 \subseteq D_8$, and thus E_8 is an $\mathfrak{a}$-ary D_8-lattice. However, the reflection $\boldsymbol{x} \mapsto \boldsymbol{x} - \frac{1}{4}\langle \boldsymbol{x}, \boldsymbol{y} \rangle \boldsymbol{y}$ for $\boldsymbol{y} = (1,1,1,1,-1,-1,-1,-1)^\top$ preserves E_8 but not D_8.

However, for us it will suffice that zeros modulo $\mathfrak{a}$ are preserved by isomorphisms.

Proposition 3.8. *Let Γ be a base lattice and $\mathfrak{a}$ an integral ideal of $\mathcal{O}$. If B is an $\mathfrak{a}$-ary Γ-lattice, then*

$$z(B) = \#\{ i \mid (\forall v \in B)\ v_i \in \mathfrak{a}\Gamma \}$$

is equal to the number of indecomposable factors of B isomorphic to $\mathfrak{a}\Gamma$.

Proof. Write $f_i : \mathfrak{a}\Gamma \to B$ for the embedding in the i-th coordinate. Observe that for $1 \leqslant i \leqslant n$ we have $i \in \{ j \mid (\forall v \in B)\ v_j \in \mathfrak{a}\Gamma \}$ if and only if $f_i(\mathfrak{a}\Gamma)$ is an indecomposable factor of B. Thus it suffices to show that if $B = B_1 \oplus B_2$ with $B_1 \cong \mathfrak{a}\Gamma$, then $B_1 = f_i(\mathfrak{a}\Gamma)$ for some i.

We claim that for each i and $\Phi \in \min(\mathfrak{a}\Gamma)$ we have $f_i(\Phi) \subseteq B_1$ or $f_i(\Phi) \subseteq B_2$. Write $\pi_1 : B \to B$ for the orthogonal projection with image B_1. If $f_i(\Phi) \nsubseteq B_1$, then by Lemma 2.6 we have that $\mathrm{N}(\pi_1(f_i(\Phi)))^{1/d} < \mathrm{N}(f_i(\Phi))^{1/d} = \lambda_1^{\mathcal{O}}(\mathfrak{a}\Gamma)$, yet $\pi_1(f_i(\Phi)) \subseteq B_1 \cong \mathfrak{a}\Gamma$. Hence $\pi_1(f_i(\Phi)) = 0$, and $f_i(\Phi) \subseteq B_2$.

It follows that Γ', the sublattice of Γ generated by $\min(\Gamma)$, satisfies

$$\mathfrak{a}\Gamma' = (\mathfrak{a}\Gamma' \cap f_i^{-1}(B_1)) \oplus (\mathfrak{a}\Gamma' \cap f_i^{-1}(B_2))$$

for each i. Namely, $\mathfrak{a}\Gamma'$ is generated by $\min(\mathfrak{a}\Gamma)$, and each $\Phi \in \min(\mathfrak{a}\Gamma)$ is contained in one of the summands by the previous paragraph. However, $\mathfrak{a}\Gamma'$ is indecomposable by the assumption on Γ. Hence $f_i(\mathfrak{a}\Gamma') \subseteq B_1$ or $f_i(\mathfrak{a}\Gamma') \subseteq B_2$. Since $B_1 \neq 0$ and $\sum_i f_i(\mathfrak{a}\Gamma') \subseteq B$ has full rank, we must have $f_i(\mathfrak{a}\Gamma') \subseteq B_1$ for at least one i. It follows that $B_1 = f_i(\mathfrak{a}\Gamma)$, as was to be shown. □

It follows from Proposition 3.8 that $z(B)$ is an isomorphism invariant. Therefore it makes sense to extend the definition of z to lattices that are only isomorphic to an $\mathfrak{a}$-ary Γ-lattice. Although a lattice can be isomorphic to an $\mathfrak{a}$-ary Γ-lattice for multiple values of $\mathfrak{a}$ and Γ, we omit them from the notation when they are understood from the context. Similarly, if $\mathfrak{a}$ is prime, then the dimension $d(B)$ of the $\mathcal{O}/\mathfrak{a}$-vector space $B/\mathfrak{a}\Gamma^n$ is preserved by isomorphisms, because $\#(B/\mathfrak{a}\Gamma^n) = \det(\mathfrak{a}\Gamma^n)/\det(B)$ is preserved, where the determinants are taken as $\mathbb{Z}$-lattices. We may compute z in polynomial time using a ΔLIP oracle as follows.

Proposition 3.9. *Fix a base lattice Γ, a prime $\mathfrak{p}$ of $\mathcal{O}$ and an integer $k \geqslant 0$, and suppose we have an oracle for ΔLIP. Then there exists a polynomial-time algorithm that, given a lattice A isomorphic to a $\mathfrak{p}$-ary Γ-lattice in Γ^n such that $d(A) \leqslant k$, computes $z(A)$ using a single oracle call on at most n^{e-1} lattices, where $e = \#(\Gamma/\mathfrak{p}\Gamma)^k$ is constant.*

We will apply Proposition 3.9 with k at most the rank of Γ.

Proof. We will show that there are, up to isomorphism, at most n^e possible $\mathfrak{p}$-ary Γ-lattices B with $d(B) \leqslant k$. Then we can decide which of those A is isomorphic to using the oracle and simply read off $z(B)$.

For each matrix $\mathbf{B} \in (\Gamma/\mathfrak{p}\Gamma)^{n\times k}$, the columns of $\mathbf{B}$ together with $\mathfrak{p}\Gamma^n$ generate a $\mathfrak{p}$-ary Γ-lattice B with $d(B) \leqslant k$, and each such $\mathfrak{p}$-ary Γ-lattice can be obtained in this way. By Corollary 3.4, permuting the rows of $\mathbf{B}$ results in an isomorphic lattice. Thus, to specify $\mathbf{B}$ up to isomorphism, it suffices to specify how often each row occurs. Since there are e possible rows, which can each occur at most n times, there are at most n^{e-1} non-isomorphic M. □

It is important to note that, although the reduction runs in polynomial time in n, the exponent is 'galactic' unless $\mathrm{N}(\mathfrak{p})$, $\operatorname{rk}\Gamma$ and k are all very small. In Section 8 we will make some practical improvements to the algorithm to lower this exponent, and in Section 9 we will show that the reduction can be highly efficient for the structured lattices common in cryptography.

4 Inductive step

Our algorithm to reduce Γ-LIP to ΔLIP proceeds by inductively computing the indecomposable orthogonal summands of the lattice modulo powers of an integral ideal. In this section we will treat the inductive step of this algorithm. The more complex base case will follow in Section 6 after some technical results in Section 5.

Definition 4.1. *For an integral ideal $\mathfrak{a}$ of $\mathcal{O}$ and a lattice Λ we define $I_{\mathfrak{a}}(\Lambda) = \{L+\mathfrak{a}\Lambda \mid L \in I(\Lambda)\}$, where $I(\Lambda)$ is the set of indecomposable orthogonal summands of Λ as in Definition 3.1.*

We consider the elements of $I_{\mathfrak{a}}(\Lambda)$ to be the approximations of the indecomposable orthogonal summands of Λ. Note that the natural map $I(\Lambda) \to I_{\mathfrak{a}}(\Lambda)$ is a bijection if $\mathfrak{a} \neq \mathcal{O}$.

The following lemma shows how we may gain information on the coordinates of a vector using a test for lattice isomorphism. In particular, we assume that we have (partial) information modulo $\mathfrak{b}\Gamma$ and lift this to (partial) information modulo $\mathfrak{a}\mathfrak{b}\Gamma$ for some integral ideals $\mathfrak{b} \subseteq \mathfrak{a}$. Later we will set $\mathfrak{b} = \mathfrak{a}^k$ to go from $\mathfrak{a}^k$ to $\mathfrak{a}\mathfrak{b} = \mathfrak{a}^{k+1}$ for $k > 0$.

Lemma 4.2. *Let Γ be a base lattice, $\mathfrak{b} \subseteq \mathfrak{a} \subsetneq \mathcal{O}$ be integral ideals and $n > 0$. Let $K_i, L_i \subseteq \Gamma^n$ be the kernel and image respectively of the projection onto the i-th component. Suppose $\boldsymbol{x} = (\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n) \in \Gamma^n$ satisfies $\boldsymbol{x}_2, \ldots, \boldsymbol{x}_n \in \mathfrak{b}\Gamma$. Then*

($i = 1$) for $\boldsymbol{y} \in \Gamma$ we have $\mathcal{O}\boldsymbol{y} + \mathfrak{a}\Gamma \cong \mathcal{O}\boldsymbol{x}_1 + \mathfrak{a}\Gamma$ if and only if

$$(\mathcal{O}\boldsymbol{y} + \mathfrak{a}\Gamma) \oplus \Gamma^{n-1} \cong \mathcal{O}\boldsymbol{x} + K_1 + \mathfrak{a}\Gamma^n.$$

($i > 1$) for $\boldsymbol{z} \in \mathfrak{b}L_i + \mathfrak{a}\mathfrak{b}\Gamma^n$ we have $\boldsymbol{z}_i + \mathfrak{a}\mathfrak{b}\Gamma = \boldsymbol{x}_i + \mathfrak{a}\mathfrak{b}\Gamma$ if and only if

$$\mathcal{O}(\boldsymbol{x} - \boldsymbol{z}) + \mathfrak{a}K_i + \mathfrak{a}\mathfrak{b}\Gamma^n \cong (\mathcal{O}\boldsymbol{x}_1 + \mathfrak{a}\Gamma) \oplus \mathfrak{a}\Gamma^{n-2} \oplus \mathfrak{a}\mathfrak{b}\Gamma.$$

Note that although we can only obtain $\boldsymbol{x}_1 \pmod{\mathfrak{a}\Gamma}$ up to equivalence, we can obtain $\boldsymbol{x}_i \pmod{\mathfrak{a}\mathfrak{b}\Gamma}$ for $i > 1$ exactly, which will be the key to the inductive algorithm.

Proof. (i) We observe that $\mathcal{O}\boldsymbol{x} + K_1 + \mathfrak{a}\Gamma^n \cong (\mathcal{O}\boldsymbol{x}_1 + \mathfrak{a}\Gamma) \oplus \Gamma^{n-1}$. Then note that $(\mathcal{O}\boldsymbol{x}_1 + \mathfrak{a}\Gamma) \oplus \Gamma^{n-1} \cong (\mathcal{O}\boldsymbol{y} + \mathfrak{a}\Gamma) \oplus \Gamma^{n-1}$ if and only if $\mathcal{O}\boldsymbol{x}_1 + \mathfrak{a}\Gamma \cong \mathcal{O}\boldsymbol{y} + \mathfrak{a}$.

(ii) Without loss of generality we assume $i = n \geqslant 2$. Observe that $\mathfrak{a}K_n + \mathfrak{a}\mathfrak{b}\Gamma^n = \mathfrak{a}\Gamma^{n-1} \oplus \mathfrak{a}\mathfrak{b}\Gamma$ and $\boldsymbol{x} - \boldsymbol{z} \equiv (\boldsymbol{x}_1, 0, \ldots, 0, \boldsymbol{x}_n - \boldsymbol{z}_n) \pmod{\mathfrak{a}\Gamma^{n-1} \oplus \mathfrak{a}\mathfrak{b}\Gamma}$. The forward implication follows trivially. Conversely, if $\mathcal{O}(\boldsymbol{x}-\boldsymbol{z}) + \mathfrak{a}K_n + \mathfrak{a}\mathfrak{b}\Gamma^n \cong (\mathcal{O}\boldsymbol{x}_1 + \mathfrak{a}\Gamma) \oplus \mathfrak{a}\Gamma^{n-2} \oplus \mathfrak{a}\mathfrak{b}\Gamma$, then by Proposition 3.8 we may conclude that the left hand side has exactly 1 coordinate in $\mathfrak{a}\mathfrak{b}\Gamma$. This can only be the n-th coordinate, from which the reverse implication follows. □

Using Lemma 4.2 we can now lift an approximation of the indecomposable orthogonal summands modulo some power $\mathfrak{a}^k$ to a higher modulus $\mathfrak{a}^{k+1}$.

Proposition 4.3. *Fix a base lattice Γ and an integral ideal $\mathfrak{a}$, and suppose we have an oracle for ΔLIP. Then there exists a polynomial-time algorithm that, given a lattice Λ isomorphic to Γ^n for some $n \geqslant 0$, and the set $I_{\mathfrak{a}^k}(\Lambda)$ for some $k \geqslant 1$, computes $I_{\mathfrak{a}^{k+1}}(\Lambda)$ in at most rn^2 oracle calls on at most q^r lattices each, where $q = \mathrm{N}(\mathfrak{a})$ and $r = \operatorname{rk}\Gamma$.*

Proof. We may assume $\mathfrak{a} \neq \mathcal{O}$ and $n > 0$. Write $\mathfrak{b} = \mathfrak{a}^k$, enumerate $I(\Lambda) = \{L_1, \ldots, L_n\}$ and let $\pi_i : \Lambda \to L_i$ be the natural projection. Although the L_i and π_i are not algorithmically available, the sublattices $\overline{L}_i = L_i + \mathfrak{b}\Lambda$ and $\overline{K}_i = \sum_{j \neq i} L_j + \mathfrak{b}\Lambda$ of A are. To prove the proposition, it suffices to give a polynomial-time algorithm that, given i and $\boldsymbol{x} \in \overline{L}_i$, computes some lifting $\boldsymbol{x}' \in \overline{L}_i$ such that

$\boldsymbol{x}' \equiv \boldsymbol{x} \pmod{\mathfrak{b}}$ and $\pi_j(\boldsymbol{x}') \equiv 0 \pmod{\mathfrak{a}\mathfrak{b}}$ for all $j \neq i$, in at most n oracle calls. Namely, we may simply apply this algorithm for each i to a generating set of $\overline{L}_i/\mathfrak{b}\Lambda$ and obtain a generating set for $L_i + \mathfrak{a}\mathfrak{b}\Lambda$.

Suppose $\boldsymbol{x} \in \overline{L}_1$ is given and write $(\boldsymbol{x}_1, \dots, \boldsymbol{x}_n) \in \Gamma^n$ for the image of $\boldsymbol{x}$ under some isomorphism $\Lambda \cong \Gamma^n$ that maps L_i to the i-th component. Again, the $\boldsymbol{x}_i$ are not algorithmically available. However, by enumerating $\Gamma/\mathfrak{a}\Gamma$ we compute using Lemma 4.2 for $i = 1$ and an oracle call some $\boldsymbol{y} \in \Gamma$ such that $\mathcal{O}\boldsymbol{x}_1 + \mathfrak{a}\Gamma \cong \mathcal{O}\boldsymbol{y} + \mathfrak{a}\Gamma$. Subsequently, by enumerating $\mathfrak{b}(\overline{L}_i + \mathfrak{a}\Lambda)/\mathfrak{a}\mathfrak{b}\Lambda$ we compute using $\boldsymbol{y}$, Lemma 4.2 and an oracle for each $i > 1$ some $\boldsymbol{z}_i \in \mathfrak{b}\overline{L}_i$ such that $\pi_i(\boldsymbol{z}_i) \equiv \boldsymbol{x}_i \pmod{\mathfrak{a}\mathfrak{b}}$. It is easy to see that we may take $\boldsymbol{x}' = \boldsymbol{x} - \sum_{i>1} \boldsymbol{z}_i$. □

Note that the computation of $\boldsymbol{y}$ in the proposition can often be eliminated in subsequent iterative steps: $\boldsymbol{y}$ only depends on the value of $\boldsymbol{x} \pmod{\mathfrak{a}}$, which is unchanged if we keep lifting the same vector throughout the algorithm. However, for practical reasons one might want to apply a basis reduction algorithm at some point and thus changing $\boldsymbol{x} \pmod{\mathfrak{a}}$. Moreover, keeping track of $\boldsymbol{y}$ has no impact on the asymptotics.

5 Combinatorics over finite fields

In this section we prove a combinatorial theorem on vector spaces over finite fields, to be used in the base case of our inductive main algorithm.

Definition 5.1. *Let V be a finite-dimensional vector space and let r be an integer. We write $\mathbb{G}_r(V)$ and $\mathbb{G}^r(V)$ for the set of all subspaces of V of dimension at most r and of codimension at most r respectively.*

We will work to prove the following theorem. The subset $T \subseteq \mathbb{G}^r(V)$ in the statement will later determine the precise oracle queries we have to make in the base case of our main algorithm.

Theorem 5.2. *Fix a prime-power $q > 1$ and $r \in \mathbb{Z}_{\geqslant 0}$. There exists a polynomial-time algorithm that, given a finite-dimensional $\mathbb{F}_q$-vector space V and a subset $S \subseteq \mathbb{G}_r(V)$, computes a subset $T \subseteq \mathbb{G}^r(V)$ such that the matrix given by $(\mathbb{1}(A \subseteq B))_{(A,B)\in S\times T}$ is invertible over $\mathbb{Q}$.*

The main difficulty here is that the sets $\mathbb{G}_r(V)$ and $\mathbb{G}^r(V)$ can be exponentially large in the dimension of V, while we want an algorithm that is polynomial in $\#S$. We will first show in Proposition 5.8 that the matrix $(\mathbb{1}(A \subseteq B))_{(A,B)\in S\times T}$ is invertible for $S = \mathbb{G}_r(V)$ and $T = \mathbb{G}^r(V)$.

Definition 5.3. *For $\ell \in \mathbb{Z}_{\geqslant 0}$ and $k \in \mathbb{Z}$ we define the* Gaussian binomial coefficient

$$\binom{\ell}{k}_q = \frac{(q^\ell - 1)(q^\ell - q)\cdots(q^\ell - q^{k-1})}{(q^k - 1)(q^k - q)\cdots(q^k - q^{k-1})} \in \mathbb{Z}[q]$$

if $0 \leqslant k \leqslant \ell$ and $\binom{\ell}{k}_q = 0$ otherwise, where q is a formal variable.

Note that if we evaluate at $q = 1$, we obtain the usual binomial coefficient. The following is a generalization of an identity of binomial coefficients that is easily verified.

Lemma 5.4. *For all $\ell, k \in \mathbb{Z}$ with $\ell > 0$ we have $q^{\ell-k}\binom{\ell-1}{k-1}_q + \binom{\ell-1}{k}_q = \binom{\ell}{k}_q$.*

Similar to how binomial coefficients count the size-k subsets of a size n-set, Gaussian binomial coefficients count subspaces of vector spaces.

Lemma 5.5. *Let $q > 1$ be a prime power and let $U \subseteq V$ be $\mathbb{F}_q$-vector spaces of finite dimensions m and and ℓ respectively. Then*

1. *for $k \in \mathbb{Z}$, the number of k-dimensional subspaces of V equals $\binom{\ell}{k}_q$;*
2. *for $k_1, k \in \mathbb{Z}$, the number of k-dimensional subspaces W of V such that $\dim(U \cap W) = k_1$ equals $\binom{m}{k_1}_q \binom{\ell-m}{k-k_1}_q q^{(m-k_1)(k-k_1)}$.*

Proof. (i) Suppose $0 \leqslant k \leqslant \ell$, otherwise the claim is trivial. There are $(q^\ell - 1)\cdots(q^\ell - q^{k-1})$ ways to (iteratively) choose a sequence of k linearly independent vectors in an ℓ-dimensional $\mathbb{F}_q$-vector space. Each subspace of dimension k is generated by $(q^k - 1)\cdots(q^k - q^{k-1})$ distinct such sequences.

(ii) Fix a projection $\pi : V \to U$. By (i), there are $\binom{m}{k_1}_q \binom{\ell-m}{k-k_1}_q$ ways to choose $W_1 \subseteq U$ of dimension k_1 and $W_2 \subseteq V/U$ of dimension $k - k_1$. Each $W \subseteq V$ with $W \cap U = W_1$ and $W \equiv W_2 \pmod U$ corresponds to a linear map $f : W_2 \to U/W_1$, namely the map that makes the following diagram commute:

$$\begin{array}{ccc} W & \xrightarrow{\pi} & U \\ {\scriptstyle \bmod U}\downarrow & & \downarrow{\scriptstyle \bmod W_1} \\ W_2 & \xrightarrow{f} & U/W_1 \end{array}$$

Note that there are $q^{\dim(U/W_1)\cdot\dim(W_2)} = q^{(m-k_1)(k-k_1)}$ such linear maps. □

Lemma 5.6. *For all integers $c, d \geqslant 0$ the determinant of the matrix*

$$\boldsymbol{A}_{c,d} = \left(\binom{c+a}{b}_q\right)_{0\leqslant a,b<d}$$

is a power of q.

This lemma is well-known for $q = 1$ and $c = 0$, where we would call $\boldsymbol{A}_{c,d}$ a *Pascal matrix*.

Proof. For $d = 0$ we have $\det(\boldsymbol{A}_{c,d}) = 1$. Let $\boldsymbol{L} = 1 - (\mathbb{1}(a = a' + 1))_{0\leqslant a,a'<d}$, which is a matrix of determinant 1, and thus we have $\det(\boldsymbol{L}\boldsymbol{A}_{c,d}) = \det(\boldsymbol{A}_{c,d})$. By Lemma 5.4 we have for $a > 0$ that

$$(\boldsymbol{L}\boldsymbol{A}_{c,d})_{a,b} = \binom{c+a}{b}_q - \binom{c+a-1}{b}_q = q^{c+a-b}\binom{c+a-1}{b-1}_q.$$

As $(\boldsymbol{L}\boldsymbol{A}_{c,d})_{a,0} = \mathbb{1}(a=0)$, the determinant of $\boldsymbol{A}_{c,d}$ is equal to the determinant of the matrix $\boldsymbol{L}\boldsymbol{A}_{c,d}$ with the first row and column removed. However, this matrix is equal to

$$\big(\mathbb{1}(a=a')\cdot q^{c+a}\big)_{0\leqslant a,a'<d}\cdot \boldsymbol{A}_{c,d-1}\cdot\big(\mathbb{1}(b=b')\cdot q^{-b}\big)_{0\leqslant b,b'<d},$$

whose determinant is a power of q times the determinant of $\boldsymbol{A}_{c,d-1}$. We can now conclude by induction on d. □

Proposition 5.7. *Let $r,\ell,t\geqslant 0$ be integers. Consider the sets*

$$\mathcal{S}=\mathcal{S}_{r,\ell,t}=\{(x_1,x_2)\in\mathbb{Z}^2_{\geqslant 0}\,|\,x_1\leqslant t,\ x_2\leqslant \ell-t,\ x_1+x_2\leqslant r\}\quad\text{and}$$
$$\mathcal{U}=\mathcal{U}_{r,\ell,t}=\{(x_1,x_2)\in\mathbb{Z}^2_{\geqslant 0}\,|\,x_1\leqslant t,\ x_2\leqslant \ell-t,\ x_1+x_2\geqslant r\}.$$

Then the determinant of the matrix

$$\boldsymbol{B}=\boldsymbol{B}_{r,\ell,t}=\left(q^{(u_1-s_1)s_2}\cdot\binom{u_1}{s_1}_q\cdot\binom{u_2}{s_2}_q\right)_{((u_1,u_2),(s_1,s_2))\in\mathcal{U}\times\mathcal{S}}$$

is (defined up to sign) a power of q.

Proof. Consider the maps $\mathcal{S}_{r-1,\ell-1,t}\to\mathcal{S}$ and $\mathcal{U}_{r-1,\ell-1,t}\to\mathcal{U}$ both given by $(x_1,x_2)\mapsto(x_1,x_2+1)$, and let $\mathcal{S}_1$ and $\mathcal{U}_1$ be their respective images and $\mathcal{S}_2$ and $\mathcal{U}_2$ their respective complements in $\mathcal{S}$ and $\mathcal{U}$. Consider the matrix

$$\boldsymbol{L}=1-\Big(\mathbb{1}(u_1=u_1',\,u_2=u_2'+1)\Big)_{\boldsymbol{u},\boldsymbol{u}'\in\mathcal{U}},$$

which has determinant 1. Let $\boldsymbol{u}\in\mathcal{U}$ and $\boldsymbol{s}\in\mathcal{S}$. Suppose that $\boldsymbol{u}\in\mathcal{U}_1$. Then $(u_1,u_2-1)\in\mathcal{U}$ and

$$\begin{aligned}(\boldsymbol{L}\boldsymbol{B})_{\boldsymbol{u},\boldsymbol{s}}&=\boldsymbol{B}_{(u_1,u_2),\boldsymbol{s}}-\boldsymbol{B}_{(u_1,u_2-1),\boldsymbol{s}}\\&=q^{(u_1-s_1)s_2}\cdot\binom{u_1}{s_1}_q\cdot\left(\binom{u_2}{s_2}_q-\binom{u_2-1}{s_2}_q\right)\\&=q^{(u_1-s_1)s_2}\cdot\binom{u_1}{s_1}_q\cdot q^{u_2-s_2}\cdot\binom{u_2-1}{s_2-1}_q\\&=q^{(u_1+u_2)-(s_1+s_2)}\cdot\boldsymbol{B}_{(u_1,u_2-1),(s_1,s_2-1)}.\end{aligned}$$

If $\boldsymbol{s}\in\mathcal{S}_2$, then $s_2=0$ and the above becomes 0. Hence the determinant of $\boldsymbol{B}$ is (up to sign) a power of q if the same holds for the blocks $\mathcal{U}_1\times\mathcal{S}_1$ and $\mathcal{U}_2\times\mathcal{S}_2$ of $\boldsymbol{L}\boldsymbol{B}$. The block corresponding to $\mathcal{U}_1\times\mathcal{S}_1$ is equal to

$$\big(\mathbb{1}(\boldsymbol{u}=\boldsymbol{u}')\cdot q^{u_1+u_2}\big)_{\boldsymbol{u},\boldsymbol{u}'\in\mathcal{U}_1}\cdot B_{r-1,\ell-1,t}\cdot\big(\mathbb{1}(\boldsymbol{s}=\boldsymbol{s}')\cdot q^{-(s_1+s_2)}\big)_{\boldsymbol{s},\boldsymbol{s}'\in\mathcal{S}_1},$$

so we may resolve this case by induction. For $\boldsymbol{u}\in\mathcal{U}_2$ and $\boldsymbol{s}\in\mathcal{S}_2$ we have $s_2=u_2=0$, hence $(\boldsymbol{L}\boldsymbol{B})_{\boldsymbol{u},\boldsymbol{s}}=\boldsymbol{B}_{\boldsymbol{u},\boldsymbol{s}}=\binom{u_1}{s_1}_q$, so $\mathcal{U}_2\times\mathcal{S}_2$ follows from Lemma 5.6. □

Proposition 5.8. *Let $q > 1$ be a prime-power and $r, \ell \in \mathbb{Z}_{\geqslant 0}$. Then the matrix*

$$\boldsymbol{M} = \boldsymbol{M}_{q,\ell,r} = (\mathbb{1}(A \subseteq B))_{(A,B)\in \mathbb{G}_r(\mathbb{F}_q^\ell)\times \mathbb{G}^r(\mathbb{F}_q^\ell)}$$

is invertible over $\mathbb{Q}$.

Proof. The matrix $\boldsymbol{B}_{r,\ell,t}$ from Proposition 5.7 is invertible over $\mathbb{Q}$ for any non-zero value of q, so let $v_t \in \mathbb{Q}^{\mathcal{S}_{r,\ell,t}}$ be such that $\boldsymbol{B}_{r,\ell,t} \cdot v_t = \big(\mathbb{1}(\boldsymbol{u} = (t,0))\big)_{\boldsymbol{u}\in\mathcal{U}_{r,\ell,t}}$. Now define

$$\boldsymbol{N} = \big((v_{\dim(T)})_{(\dim(S\cap T),\dim(S)-\dim(S\cap T))}\big)_{(T,S)\in\mathbb{G}^r\times\mathbb{G}_r}.$$

Let $T, U \in \mathbb{G}^r$ with $(t, u_1, u_2) = (\dim T, \dim T \cap U, \dim U - \dim T \cap U)$. Using Lemma 5.5 we compute

$$\begin{aligned}
(\boldsymbol{NM})_{T,U} &= \sum_{S\in\mathbb{G}_r} (v_t)_{(\dim(S\cap T),\dim S-\dim(S\cap T))} \cdot \mathbb{1}(S \subseteq U) \\
&= \sum_{(s_1,s_2)\in\mathcal{S}_{r,\ell,t}} (v_t)_{(s_1,s_2)} \cdot \#\{S \subseteq U \mid \dim(S) = s_1 + s_2,\ \dim(S \cap T) = s_1\} \\
&= \sum_{(s_1,s_2)\in\mathcal{S}_{r,\ell,t}} (v_t)_{(s_1,s_2)} \cdot q^{(u_1-s_1)s_2} \cdot \begin{bmatrix} u_1 \\ s_1 \end{bmatrix}_q \cdot \begin{bmatrix} u_2 \\ s_2 \end{bmatrix}_q \\
&= (B_{r,\ell,t} \cdot v_t)_{(u_1,u_2)} = \mathbb{1}((u_1,u_2) = (t,0)) = \mathbb{1}(T = U).
\end{aligned}$$

Hence $\boldsymbol{N}$ is the inverse of $\boldsymbol{M}$. □

To move to a submatrix of $\boldsymbol{M}_{q,\ell,r}$ we have the following elementary result from linear algebra.

Lemma 5.9. *There is a polynomial-time algorithm that, given a field $\mathbb{F}$ which is $\mathbb{Q}$ or finite, an invertible matrix $\boldsymbol{M} = (M_{i,j})_{i,j} \in \mathbb{F}^{I\times J}$ for sets I and J, and a subset $S \subseteq I$, computes a subset $T \subseteq J$ such that the submatrix $(M_{s,t})_{(s,t)\in S\times T}$ is invertible.*

Proof. Suppose $\mathbb{F}$ is finite. Consider the projection $\pi : \mathbb{F}^I \to \mathbb{F}^S$. By row-operations on $\pi \circ \boldsymbol{M}$, i.e. composing by invertible maps on the left, we obtain a matrix $\alpha \circ \pi \circ \boldsymbol{M}$ in row-echelon form. Since $\mathbb{F}$ is finite, arithmetic operations do not suffer from coefficient blowup, and the computation can trivially be done in polynomial time. We then choose T to be the set of pivot columns. The resulting matrix $\alpha \circ \pi \circ \boldsymbol{M} \circ \iota$ for $\iota : \mathbb{F}^T \to \mathbb{F}^J$ is square and of full rank, so it and $\pi \circ M \circ \iota$ are invertible. If $\mathbb{F} = \mathbb{Q}$, we may pick a prime p such that $\boldsymbol{M}$ is well-defined and invertible modulo p, and reduce to the previous case. □

Note that Theorem 5.2 does not immediately follow from Proposition 5.8 and Lemma 5.9, because we generally cannot compute the matrix $\boldsymbol{M}_{q,n,r}$ in polynomial time as the number of rows $\#\mathbb{G}_r(V)$ and the number of columns $\#\mathbb{G}^r(V)$ can be too large. Using the special structure of the matrix we can however reduce to the case where $\boldsymbol{M}_{q,n,r}$ is small enough.

Proof (Theorem 5.2). Consider the set

$$E = \{\boldsymbol{x} \in V \mid (\forall A, A' \in S)\ A = A' \Leftrightarrow A + \mathbb{F}_q\boldsymbol{x} = A' + \mathbb{F}_q\boldsymbol{x}\}.$$

We will first show that its complement is polynomially bounded. Suppose $\boldsymbol{x} \notin E$. Then there exist distinct $A, A' \in S$ such that $A + \mathbb{F}_q\boldsymbol{x} = A' + \mathbb{F}_q\boldsymbol{x}$. It follows that $\boldsymbol{x} \in A + A'$ and $\dim(A + A') \leqslant r + 1$, so there are at most q^{r+1} possible values for $\boldsymbol{x}$. Hence $\#(V \backslash E) \leqslant \binom{\#S}{2} \cdot q^{r+1}$, as was to be shown.

Hence we may compute, in polynomial time, a nonzero element $\boldsymbol{x} \in E$ or decide that no such element exists. Suppose first that no such $\boldsymbol{x}$ exists. Then $V = V \backslash E \cup \{0\}$ is polynomially bounded in size, and so are $\mathbb{G}_r(V)$ and $\mathbb{G}^r(V)$. Hence we may compute the entire matrix from Proposition 5.8 and apply Lemma 5.9.

Suppose now we have some nonzero $\boldsymbol{x} \in E$. Then we compute the projection $\pi : V \to V' = V/(\mathbb{F}_q \cdot \boldsymbol{x})$ and $S' = \{\pi A \mid A \in S\}$. Note that by definition of E the natural map $S \to S'$ is injective. We proceed recursively and obtain a set T' of subspaces of V' of codimension at most r such that the matrix $\boldsymbol{M'} = (\mathbb{1}(A' \subseteq B'))_{(A',B') \in S' \times T'}$ is invertible over $\mathbb{Q}$. We then return $T = \{\pi^{-1}B' \mid B' \in T'\}$. Note that again $T' \to T$ is a bijection. In fact, with these bijections $\boldsymbol{M'} = (\mathbb{1}(A \subseteq B))_{(A,B) \in S \times T}$, so in particular the latter is invertible. □

If we apply duality to to this theorem, we obtain the following.

Corollary 5.10. *Fix a prime-power $q > 1$ and $r \in \mathbb{Z}_{\geqslant 0}$. There exists a polynomial-time algorithm that, given a finite-dimensional $\mathbb{F}_q$-vector space V and a subset $T \subseteq \mathbb{G}^r(V)$, computes a subset $S \subseteq \mathbb{G}_r(V)$ such that the matrix given by $(\mathbb{1}(A \subseteq B))_{(A,B) \in S \times T}$ is invertible over $\mathbb{Q}$.* □

6 Base case

To finish the proof of the main theorem, it remains to prove the base case of the inductive algorithm. Namely, we will compute $I_{\mathfrak{p}}(\Lambda)$ (see Definition 4.1) for some prime $\mathfrak{p}$ of $\mathcal{O}$. We may phrase this problem in terms of vector spaces over finite fields.

Definition 6.1. *Let V be a finite vector space over a finite field $\mathbb{F}$. A* decomposition *of V is a set I of subspaces of V such that the natural map $\prod_{L \in I} L \to V$ is an isomorphism. For $L \in I$ write $L^{\perp} = \sum_{M \in I,\, M \neq L} M$ for the kernel of the projection $V \to L$. An* oracle *to I is an algorithm that, given a subspace $A \subseteq V$ of dimension at most $\max\{\dim L \mid L \in I\}$, outputs the number of $L \in I$ such that $A \subseteq L^{\perp}$.*

Proposition 6.2. *Fix a finite field $\mathbb{F}_q$ and an integer $r \in \mathbb{Z}_{\geqslant 0}$. There exists a polynomial-time algorithm that, given finite $\mathbb{F}_q$-vector spaces $U' \subseteq U \subseteq V$ with $\dim(U/U') = 1$, an oracle to a decomposition I of V with $\dim L \leqslant r$ for all $L \in I$, and the set $\{L^{\perp} \cap U' \mid L \in I\}$, computes $\{L^{\perp} \cap U \mid L \in I\}$ in at most $(1 + (q-1)q^r) \cdot n$ oracle calls, where $n = \#I$.*

Proof. For $A, B \subseteq V$ we define the variables

$$X_A = \#\{L \in I \mid A \subseteq L^\perp\} \quad \text{and}$$
$$Y_B = \#\{L \in I \mid L^\perp \cap U = B\}.$$

Although L is algorithmically unavailable, and hence X_A and Y_B not (yet) known, the oracle allows us to compute X_A if $A \in \mathbb{G}_r(U)$. On the other hand, knowing $D = \{B \in \mathbb{G}^r(U) \mid Y_B \neq 0\}$ finishes the algorithm, since each B with $Y_B \neq 0$ satisfies $\dim(U/B) \leqslant r$. We will use the theory from the previous section to establish a relation between X and Y.

Write $D' = \{L^\perp \cap U' \mid L \in I\}$, which is given as input. Now define

$$T = \{B' \in D' \mid \dim(U/B') \leqslant r\} \cup \bigcup_{B' \in D'} \{B \subseteq U \mid B' \subseteq B \nsubseteq U', \dim(B/B') = 1\}.$$

By construction we have $D \subseteq T$, since each $B \in D$ contains $B' = B \cap U' \in D'$ with $\dim(B/B') \leqslant \dim(U/U') = 1$. Moreover, we have $T \subseteq \mathbb{G}^r(U)$ and

$$\#T \leqslant \#D' + \#D' \cdot \#\{\boldsymbol{x} \in \mathbb{F}_q^{r+1} \mid \boldsymbol{x} \notin \mathbb{F}_q^r\} \leqslant n + n(q-1)q^r$$

is polynomially bounded as function of n. Hence we compute T in polynomial time. It suffices to compute the vector $(Y_B)_{B \in T}$.

We compute some set $S \subseteq \mathbb{G}_r(U)$ such that the matrix $M = (\mathbb{1}(A \subseteq B))_{(A,B) \in S \times T}$ is invertible using Corollary 5.10. In particular $\#S = \#T \leqslant (1 + (q-1)q^r) \cdot n$. For $A \in S$ we have

$$X_A = \sum_{B \in T} \#\{L \in I \mid A \subseteq B,\ B = L^\perp \cap U\} = \sum_{B \in T} \mathbb{1}(A \subseteq B) \cdot Y_B$$

We compute the $(X_A)_{A \in S}$ using the oracle. The invertibility of M then allows us to compute $(Y_B)_{B \in T}$ using linear algebra, as was to be shown. □

Corollary 6.3. *Fix a base lattice Γ and prime $\mathfrak{p}$ of $\mathcal{O}$, and suppose we have an oracle for ΔLIP. Then there exists a polynomial-time algorithm that, given a lattice Λ isomorphic to Γ^n for some $n \geqslant 0$, computes $I_\mathfrak{p}(\Lambda)$ using $(1+(q-1)q^r) \cdot n^2$ oracle calls on at most n^{e-1} lattices, where $q = \mathrm{N}(\mathfrak{p})$, $r = \operatorname{rk} \Gamma$ and $e = q^{r^2}$.*

Note that for $(\mathcal{O}, \Gamma, \mathfrak{p}) = (\mathbb{Z}, \mathbb{Z}, 2\mathbb{Z})$ we recover the result from Szydlo which makes $O(n^2)$ oracle calls on n lattices each to compute an isomorphism to $\mathbb{Z}^n$.

Proof. Using Proposition 3.9 we may construct an oracle to the decomposition $I = \{\overline{L}/\mathfrak{p}\Lambda \mid \overline{L} \in I_\mathfrak{p}(\Lambda)\}$ of $V = \Lambda/\mathfrak{p}\Lambda$, where each oracle call requires n^{e-1} calls to the oracle for ΔLIP, where $e = q^{r \cdot r}$. Choose some filtration $U_0 \subset U_1 \subset \cdots \subset U_n = V$ such that $\dim U_i = i$ for all i, and iteratively apply Proposition 6.2 to compute I and consequently $I_\mathfrak{p}(\Lambda)$. □

We can reduce the number of queries to the oracle in Corollary 6.3 by reusing computations of X_A's and Y_B's from earlier iterations. As for Proposition 4.3 these optimizations do not impact the asymptotic run time, as in, they do not decrease the constant in the polynomial exponent of n.

7 Main theorem

We will now combine all our results to prove the main theorem. We rely on basis reduction to reduce the computation of $I(\Lambda)$ to that of $I_{\mathfrak{p}^k}(\Lambda)$ for some large enough power of $\mathfrak{p}$.

Theorem 7.1 (Lenstra–Lenstra–Lovász, Section 10 in [16]). *Fix $c > 4/3$. Then there exists a polynomial-time algorithm that, given a $\mathbb{Z}$-lattice Λ, computes a basis $(\boldsymbol{b}_1, \dots, \boldsymbol{b}_\ell)$ of Λ such that $\|\boldsymbol{b}_i\|^2 \leqslant c^{\ell-1} \cdot \lambda_i(\Lambda)^2$ for all i.* □

Lemma 7.2. *Write $d = [K : \mathbb{Q}]$. If Λ is a lattice, then $\lambda_1(\Lambda) \geqslant \sqrt{d} \cdot \lambda_1^{\mathcal{O}}(\Lambda)$.*

Proof. Let $\boldsymbol{x} \in \Lambda$ be such that $\langle \boldsymbol{x}, \boldsymbol{x} \rangle_{\mathbb{Q}} = \lambda_1(\Lambda)^2$. By the inequality of the arithmetic and geometric mean we have

$$\lambda_1^{\mathcal{O}}(\Lambda)^2 \leqslant \mathrm{N}(\langle \boldsymbol{x}, \boldsymbol{x} \rangle_K)^{1/d} \leqslant (1/d) \operatorname{Tr}(\langle \boldsymbol{x}, \boldsymbol{x} \rangle_K) = (1/d) \cdot \langle \boldsymbol{x}, \boldsymbol{x} \rangle_{\mathbb{Q}},$$

from which the lemma follows immediately. □

Theorem 7.3 (Main result). *Fix a number field with maximal order $\mathcal{O}$, a base $\mathcal{O}$-lattice Γ and a prime $\mathfrak{p}$ of $\mathcal{O}$, and suppose we have an oracle for $\mathcal{O}$-ΔLIP. Then there exists a polynomial-time algorithm that, given an $\mathcal{O}$-lattice Λ isomorphic to Γ^n for some n, computes such an isomorphism in $O(n^3)$ oracle calls on $O(1)$ lattices each and $O(n^2)$ oracle calls on $O(n^{e-1})$ lattices each, where $q = \mathrm{N}(\mathfrak{p})$, $r = \operatorname{rk} \Gamma$ and $e = q^r$.*

Proof. Write $d = [K : \mathbb{Q}]$ and fix some $c > 4/3$. Combining Corollary 6.3 and Proposition 4.3, we may compute $I_{\mathfrak{p}^k}(\Lambda)$ for k the smallest integer satisfying

$$k > \frac{1}{2}(ndr - 1)d \cdot \log_q(c) + d \log_q \left(\frac{\lambda_{dr}(\Gamma)}{\sqrt{d} \cdot \lambda_1^{\mathcal{O}}(\Gamma)} \right), \tag{1}$$

which is linear in n. The number of oracle calls required is as claimed.

Let $\overline{L} \in I_{\mathfrak{p}^k}(\Lambda)$ and compute a $\mathbb{Z}$-basis $(\boldsymbol{b}_1, \dots, \boldsymbol{b}_{ndr})$ for $\overline{L}$ using Theorem 7.1. We claim that $L' := \sum_{i=1}^{dr} \mathbb{Z}\boldsymbol{b}_i \in I(\Lambda)$ and $L' + \mathfrak{p}^k\Lambda = \overline{L}$. If true, we have computed the decomposition $\Lambda \to \prod_{L \in I(\Lambda)} L$. In turn, we may compute an isomorphism $L \cong \Gamma$ in polynomial time by brute force, since Γ is of constant rank and K is fixed[5]. It remains to prove the claim.

Let $L \in I(\Lambda)$ be such that $L + \mathfrak{p}^k\Lambda = \overline{L}$. Fix $1 \leqslant i \leqslant dr$ and write $\boldsymbol{b}_i = \sum_{M \in I(\Lambda)} \boldsymbol{x}_M$ for $\boldsymbol{x}_M \in M$. We have $\lambda_i(\overline{L}) \leqslant \lambda_i(L) = \lambda_i(\Gamma)$, because $\Gamma \cong L \subseteq \overline{L}$. Using Lemma 2.4 and Lemma 7.2 we obtain

$$\begin{aligned} \|\boldsymbol{x}_M\|^2 \leqslant \|\boldsymbol{b}_i\|^2 &\leqslant c^{ndr-1} \lambda_i(\Lambda)^2 \leqslant c^{ndr-1} \lambda_{dr}(\Gamma)^2 \\ &< d \cdot N(\mathfrak{p})^{2k/d} \cdot \lambda_1^{\mathcal{O}}(\Gamma)^2 = d \cdot \lambda_1^{\mathcal{O}}(\mathfrak{p}^k\Gamma)^2 \leqslant \lambda_1(\mathfrak{p}^k M)^2, \end{aligned}$$

where the strict inequality is Equation 1. For $M \neq L$ we have $\boldsymbol{x}_M \in \mathfrak{p}^k M$, so $\boldsymbol{x}_M = 0$. Hence $\boldsymbol{b}_i \in L$ and $L' = \sum_{i=1}^{dr} \mathbb{Z}\boldsymbol{b}_i \subseteq L$. Since $\operatorname{rk} L' = \operatorname{rk} L$ and L' is a summand of Λ, we conclude that $L' = L \in I(\Lambda)$, as was to be shown. □

[5] While K is fixed these rank r instances over K might still be difficult to solve in practice, for example for HAWK where $\dim(K) \geqslant 512$. To treat that case we will later use that rank $r = 1$ instances can always be solved efficiently.

8 Practical considerations

We will fix a prime $\mathfrak{p}$ of $\mathcal{O}$ and write $\mathbb{F} = \mathcal{O}/\mathfrak{p}$. In this section we will give an improved version of Proposition 3.9 that, with some precomputation, achieves a better exponent in n as function of Γ and $\mathfrak{p}$.

We recall the general structure of Proposition 3.9. Given a lattice $\mathfrak{p}\Lambda \subseteq A \subseteq \Lambda$, we compute a list of candidates of lattices $\mathfrak{p}\Gamma^n \subseteq B \subseteq \Gamma^n$ such that A must be isomorphic to one of the lattices in the list. Using the ΔLIP oracle, one finds such a B in the list, and because of how B was constructed, we may read off several isomorphism invariants of B and hence of A.

First, we observe that some of the lattices in the list may be isomorphic and hence superfluous. In particular, for any basis $\mathbf{B} \in (\Gamma/\mathfrak{p}\Gamma)^{n\times k}$ of $B/\mathfrak{p}\Gamma^n$ we can apply an automorphism of Γ coordinate-wise on each row $C_i \in (\Gamma/\mathfrak{p}\Gamma)^k$ independently and permute the rows.

Definition 8.1. *For $k \geqslant 0$ we define the set of* row classes

$$\mathrm{rc}_k = \mathrm{rc}_k(\Gamma, \mathfrak{p}) = \mathrm{Aut}(\Gamma)\backslash(\Gamma/\mathfrak{p}\Gamma)^k,$$

which is the set of orbits of $\mathrm{Aut}(\Gamma)$ acting coordinate-wise on $(\Gamma/\mathfrak{p}\Gamma)^k$.

Example 8.2. Consider the ring $\mathcal{O} = \mathbb{Z}$ with prime $\mathfrak{p} = 2\mathbb{Z}$ and the hexagonal base lattice $\Gamma = A_2$, which has a basis consisting of

$$\boldsymbol{a} = (1, -1, 0)^\top \quad \text{and} \quad \boldsymbol{b} = (0, 1, -1)^\top,$$

and has a rather large automorphism group of order 12. If we denote by $[\boldsymbol{g}_1, \ldots, \boldsymbol{g}_k]$ the class of $(\boldsymbol{g}_1, \ldots, \boldsymbol{g}_k)$ in rc_k, then

$$\begin{aligned}
\mathrm{rc}_0 &= \{[\,]\},\\
\mathrm{rc}_1 &= \{[\boldsymbol{0}], [\boldsymbol{a}]\},\\
\mathrm{rc}_2 &= \{[\boldsymbol{0}, \boldsymbol{0}], [\boldsymbol{0}, \boldsymbol{a}], [\boldsymbol{a}, \boldsymbol{0}], [\boldsymbol{a}, \boldsymbol{a}], [\boldsymbol{a}, \boldsymbol{b}]\}.
\end{aligned}$$

The sublattices $\mathfrak{p}\Gamma^4 \subseteq B_i \subseteq \Gamma^4$ generated by $\mathfrak{p}\Gamma^4$ and the images of

$$q_1 = \begin{pmatrix} \boldsymbol{a}+\boldsymbol{b} & \boldsymbol{0} \\ \boldsymbol{a}+\boldsymbol{b} & \boldsymbol{a} \\ \boldsymbol{b} & \boldsymbol{b} \\ \boldsymbol{a} & \boldsymbol{0} \end{pmatrix}, \quad q_2 = \begin{pmatrix} \boldsymbol{a} & \boldsymbol{0} \\ \boldsymbol{a} & \boldsymbol{b} \\ \boldsymbol{a} & \boldsymbol{a} \\ \boldsymbol{a} & \boldsymbol{0} \end{pmatrix} \quad \text{and} \quad q_3 = \begin{pmatrix} \boldsymbol{a} & \boldsymbol{0} \\ \boldsymbol{a} & \boldsymbol{0} \\ \boldsymbol{a} & \boldsymbol{a} \\ \boldsymbol{a} & \boldsymbol{b} \end{pmatrix}$$

respectively are isomorphic. Namely, we may obtain q_2 from q_1 by applying a (possibly different) automorphism of Γ to each row, and q_3 from q_2 by reordering the rows.

Two bases $\mathbf{B}, \mathbf{B}' \in (\Gamma/\mathfrak{p}\Gamma)^{n\times k}$ have the same row classes (with multiplicity), if and only if the corresponding $\mathfrak{p}$-ary lattices are isomorphic.

Lemma 8.3. *Let Γ be an indecomposable lattice and $k \geqslant 0$. We have a natural map*

$$\mathrm{rep}_k : (\Gamma/\mathfrak{p}\Gamma)^{n\times k} \to \mathbb{Z}^{\mathrm{rc}_k},$$

counting row classes, that induces a bijection

$$\mathrm{Aut}(\Gamma^n)\backslash(\Gamma/\mathfrak{p}\Gamma)^{n\times k} \leftrightarrow \left\{(x_c)_c \in \mathbb{Z}^{\mathrm{rc}_k} \;\middle|\; x_c \geqslant 0, \sum_c x_c = n\right\}.$$

This allows us to give a sharper version of Proposition 3.9.

Proposition 8.4. *Fix a base lattice Γ and an integer $k \geqslant 0$, and suppose we have an oracle for ΔLIP. Then there exists a polynomial-time algorithm that, given a lattice A isomorphic to a $\mathfrak{p}$-ary Γ-lattice in Γ^n such that $d(A) \leqslant k$, computes $z(A)$ in one oracle call on $\binom{n+e-1}{e-1} = \Theta(n^{e-1})$ lattices, where $e = \#\mathrm{rc}_k$ is constant.*

Proof. The map from Lemma 8.3 and its inverse are easily computable in polynomial time. Thus we may construct a full set of representatives in $(\Gamma/\mathfrak{p}\Gamma)^{n\times k}$ up to the action of $\mathrm{Aut}(\Gamma^n)$, and in turn corresponding lattices $\mathfrak{p}\Gamma^n \subseteq B \subseteq \Gamma^n$. By construction, A is isomorphic to at least one of those lattices. Note that there are exactly $\binom{n+e-1}{e-1}$ such representatives, and the algorithm proceeds as in Proposition 3.9. □

We will apply Proposition 8.4 with k at most the rank of Γ. Unlike Proposition 3.9, the runtime of Proposition 8.4 depends on more properties of Γ than just its rank. We want the number of row classes to be small, which correlates with Γ having a large automorphism group. Further optimizations can be made: not only does $\mathrm{Aut}(\Gamma^n)$ act on the left of $(\Gamma/\mathfrak{p}\Gamma)^{n\times k}$, but $\mathrm{GL}_k(\mathbb{F})$ also acts on the right. Since the latter group is constant, taking into account this action does not yield a better exponent.

In the remainder of this section, we will decrease the complexity of higher dimensional queries, i.e., computing $z(A)$ when $d(A)$ is large, at the expense of more lower dimensional queries using recursive techniques. However, to make the recursion work, we do not compute z, but rep_k instead. By definition, rep_k is invariant under automorphisms of Γ^n, and thus extends to lattices Λ isomorphic to Γ^n.

Remark 8.5. Although the definition of rep_k extends to the case where Γ^n is replaced by a lattice Λ isomorphic to Γ^n, take care to note that rep_k is not a function on lattices isomorphic to $\mathfrak{p}$-ary Γ-lattices. Although we can effortlessly go from a basis $\mathbf{A} \in (\Lambda/\mathfrak{p}\Lambda)^k$ to a lattice A isomorphic to a $\mathfrak{p}$-ary Γ-lattice, to go from a lattice A to a basis $\mathbf{A}$, we need to know how A embeds in Λ. Example 3.7 shows that this can be done in non-trivial ways. In particular, to test whether $\mathrm{rep}_k(\mathbf{A}) = \mathrm{rep}_k(\mathbf{B})$, it is not sufficient to test whether the corresponding lattices A and B are isomorphic.

Although Remark 8.5 presents an obstruction, the following algorithm allows us to compare $\mathrm{rep}_k(\mathbf{A})$ and $\mathrm{rep}_k(\mathbf{B})$. However, this comes at the cost of increasing the rank of the lattice.

Proposition 8.6. *Fix an integer $k \geqslant 0$ and a rational $R > 0$. There exists a polynomial-time algorithm that, given a lattice Λ with $\lambda_1(\Lambda) > R$ and a $\mathbf{A} \in (\Lambda/\mathfrak{p}\Lambda)^k$, computes a lattice $\mathrm{aug}_k(\mathbf{A})$ of rank $\operatorname{rk}\Lambda + k$ satisfying the following: For any two inputs $(\Lambda_1, \mathbf{A}_1)$ and $(\Lambda_2, \mathbf{A}_2)$, there exists an isometry $f : \Lambda_1 \to \Lambda_2$ such that $f(\mathbf{A}_1) = \mathbf{A}_2$ if and only if $\mathrm{aug}_k(\mathbf{A}_1) \cong \mathrm{aug}_k(\mathbf{A}_2)$.*

Proof. We may precompute some rank-k lattice Ξ such that $\mathrm{Aut}(\Xi) = \mathrm{Aut}(\mathcal{O})$, a condition which is satisfied by almost all lattices, and an isomorphism $\mathbf{B} : \mathbb{F}^k \to \Xi/\mathfrak{p}\Xi$. Then by scaling we may assume that $\mathfrak{p}\Xi$ is generated by its vectors of length at most R. We return the lattice

$$\mathrm{aug}_k(\mathbf{A}) = \{(\mathbf{A}\boldsymbol{x}, \mathbf{B}\boldsymbol{x}) \mid \boldsymbol{x} \in \mathbb{F}^k\} + \mathfrak{p}(\Lambda \oplus \Xi) \subseteq \Lambda \oplus \Xi.$$

Now let $(\Lambda_1, \mathbf{A}_1)$ and $(\Lambda_2, \mathbf{A}_2)$ be two inputs, and write $M_i = \mathrm{aug}_k(\mathbf{A}_i)$.

($\Rightarrow$) Suppose $f : \Lambda_1 \to \Lambda_2$ is an isomorphism such that $f(\mathbf{A}_1) = \mathbf{A}_2$. Then the map $g : M_1 \to M_2$ given by $(\boldsymbol{y}, \boldsymbol{z}) \mapsto (f(\boldsymbol{y}), \boldsymbol{z})$ is an isomorphism.

($\Leftarrow$) Suppose $g : M_1 \to M_2$ is an isomorphism. For each $(\boldsymbol{y}, \boldsymbol{z}) \in M_i$ such that $\boldsymbol{y} \neq 0$ we have $\|(\boldsymbol{y}, \boldsymbol{z})\| \geqslant \lambda_1(\Lambda_i) > R$, so by definition of Ξ we get

$$\mathrm{span}\{\boldsymbol{x} \in M_i \mid \|\boldsymbol{x}\| \leqslant R\} = 0 \oplus \mathfrak{p}\Xi.$$

In particular, $g(0 \oplus \mathfrak{p}\Xi) = 0 \oplus \mathfrak{p}\Xi$, i.e., g induces an automorphism of Ξ. Since $\mathrm{Aut}(\Xi) = \mathrm{Aut}(\mathcal{O})$, we may compose g with some $\alpha \in \mathrm{Aut}(\mathcal{O}) \subseteq \mathcal{O}$ to assume g induces the identity on Ξ. Because $\mathbf{B}$ is injective, we also obtain $\mathfrak{p}\Lambda_i$ within M_i as the orthogonal complement of Ξ, so g similarly induces an isomorphism $f : \Lambda_1 \to \Lambda_2$.

Let $\boldsymbol{x}_1 \in \mathbb{F}^k$. Then $g(\mathbf{A}_1\boldsymbol{x}_1, \mathbf{B}\boldsymbol{x}_1) \equiv (\mathbf{A}_2\boldsymbol{x}_2, \mathbf{B}\boldsymbol{x}_2) \pmod{\mathfrak{p}}$ for some $\boldsymbol{x}_2 \in \mathbb{F}^k$. Since g acts trivially on Ξ by assumption, we have $\mathbf{B}\boldsymbol{x}_1 \equiv \mathbf{B}\boldsymbol{x}_2 \pmod{\mathfrak{p}}$ and thus $\boldsymbol{x}_1 \equiv \boldsymbol{x}_2 \pmod{\mathfrak{p}}$. Hence $f(\mathbf{A}_1\boldsymbol{x}_1) \equiv \mathbf{A}_2\boldsymbol{x}_1 \pmod{\mathfrak{p}}$, as was to be shown. □

The following generalization of Proposition 8.4 shows how we can compute rep_k using Proposition 8.6.

Proposition 8.7. *Fix an indecomposable lattice Γ and an integer $k \geqslant 0$, and suppose we have an oracle for ΔLIP. Then there exists a polynomial-time algorithm that, given a lattice Λ isomorphic to Γ^n and $\mathbf{A} \in (\Lambda/\mathfrak{p}\Lambda)^k$, computes $\mathrm{rep}_k(\mathbf{A})$ in one oracle call on $\binom{n+e-1}{e-1} = \Theta(n^{e-1})$ lattices of rank $\operatorname{rk}\Lambda + k$, where $e = \#\mathrm{rc}_k$ is constant.*

Proof. We compute a full set of representatives $Q \subseteq (\Gamma/\mathfrak{p}\Gamma)^{n \times k}$ up to the left action of $\mathrm{Aut}(\Gamma^n)$, of which there are exactly $\binom{n+e-1}{e-1}$. Using Proposition 8.6 and the oracle, we may find a $\mathbf{B} \in Q$ for which there exists some isomorphism $f : \Lambda \to \Gamma^n$ such that $\mathbf{B} = f(\mathbf{A})$, and return $\mathrm{rep}_k(\mathbf{B})$. □

One important observation is that Γ no longer needs to be a base lattice, but only be indecomposable. In fact, we may use Proposition 8.6 to similarly weaken this condition in the main theorem, at the additional cost of larger rank ΔLIP queries. However, the practical condition that $\mathrm{Aut}(\Gamma)$ is large often forces Γ to be a base lattice, in which case such a modified theorem is unnecessary.

To compute $z(A)$, we may choose some basis $\mathbf{A} \in (\Lambda/\mathfrak{p}\Lambda)^k$ for A, compute $\mathrm{rep}_k(\mathbf{A})$ and return the coefficient at $[0, \dots, 0]$. It remains to give an algorithm to compute rep_k recursively and more efficiently than Proposition 8.7.

Definition 8.8. *Fix an indecomposable lattice Γ and let $k > 0$. We define the integer matrix*

$$M_k = \left(\mathbb{1}(\boldsymbol{c}_k \mathbf{P} = \boldsymbol{c}_{k-1})\right)_{((\mathbf{P}, \boldsymbol{c}_{k-1}), \boldsymbol{c}_k) \in (\mathbb{F}^{k \times (k-1)} \times \mathrm{rc}_{k-1}) \times \mathrm{rc}_k}.$$

This matrix encodes for each representative $\boldsymbol{r} \in \mathbb{Z}^{\mathrm{rc}_k}$ the representatives of all of its restrictions induced by the elements of $\mathbb{F}^{k \times (k-1)}$. Conversely, it captures how much information on $\boldsymbol{r}$ can be extracted from the restrictions of $\boldsymbol{r}$.

Example 8.9. We continue Example 8.2 of $\Gamma = A_2$ for $k = 2$. Then M_2 is as follows:

		$[\mathbf{0},\mathbf{0}]$	$[\mathbf{0},\boldsymbol{a}]$	$[\boldsymbol{a},\mathbf{0}]$	$[\boldsymbol{a},\boldsymbol{a}]$	$[\boldsymbol{a},\boldsymbol{b}]$
$\binom{0}{0}$	$[\mathbf{0}]$	1	1	1	1	1
	$[\boldsymbol{a}]$	0	0	0	0	0
$\binom{1}{0}$	$[\mathbf{0}]$	1	1	0	0	0
	$[\boldsymbol{a}]$	0	0	1	1	1
$\binom{0}{1}$	$[\mathbf{0}]$	1	0	1	0	0
	$[\boldsymbol{a}]$	0	1	0	1	1
$\binom{1}{1}$	$[\mathbf{0}]$	1	0	0	1	0
	$[\boldsymbol{a}]$	0	1	1	0	1

Note that M_2 has rank 4, with a (right) kernel span by

$$\boldsymbol{t} = -[\mathbf{0},\mathbf{0}] + [\mathbf{0},\boldsymbol{a}] + [\boldsymbol{a},\mathbf{0}] + [\boldsymbol{a},\boldsymbol{a}] - 2[\boldsymbol{a},\boldsymbol{b}].$$

Suppose now that $\mathbf{A} \in (\Lambda/2\Lambda)^2$ is a query to rep_2. Then we obtain 4 distinct lower dimension queries $\mathbf{AP}$ with $\mathbf{P} \in \mathbb{F}^{2\times 1}$. For each $\mathbf{P}$, we compute $\mathrm{rep}_1(\mathbf{AP})$ in 1 call to a ΔLIP oracle on n lattices using Proposition 8.7. We compute some $\boldsymbol{s} \in \mathbb{Z}^{\mathrm{rc}_2}$ such that $M_2\boldsymbol{s}$ is consistent with the $\mathrm{rep}_1(\mathbf{AP})$ using linear algebra. Then $\mathrm{rep}_2(\mathbf{A}) = \boldsymbol{s} + \lambda\boldsymbol{t}$ for some $\lambda \in \mathbb{Z}$. Together with the restriction that the coefficients of $\mathrm{rep}_2(\mathbf{A})$ are non-negative, this leaves at most n possible values for λ, and we may compute $\mathrm{rep}_2(\mathbf{A})$ in at most 1 additional oracle call on n lattices. The total number of oracle calls to compute $\mathrm{rep}_2(\mathbf{A})$ is at most 5 on n lattices each, compared to 1 on n^4 lattices in Proposition 8.7. In particular, up to constant factors, the Γ-LIP to ΔLIP reduction algorithm for $\Gamma = A_2$ has the same complexity as for $\Gamma = \mathbb{Z}$.

We could even do with less queries: $\mathrm{rep}_1(\mathbf{A}\binom{0}{0})$ actually gives us no additional information given $\mathrm{rep}_1(\mathbf{AP})$ for all $\mathbf{P} \neq 0$. Hence we may improve 5 to 4 queries.

Applying the techniques shown in the example recursively, we obtain the following proposition.

Proposition 8.10. *Fix an indecomposable lattice Γ, a prime $\mathfrak{p}$ of $\mathcal{O}$ and an integer $k \geqslant 0$, and suppose we have an oracle for ΔLIP. Then there exists a polynomial-time algorithm that, given a lattice Λ isomorphic to a $\mathfrak{p}$-ary Γ-lattice in Γ^n and an $\mathbf{A} \in (\Lambda/\mathfrak{p}\Lambda)^k$, computes $\operatorname{rep}_k(\mathbf{A})$ in, for each $0 < i \leqslant k$ combined, at most*

$$\prod_{j=i+1}^{k} \operatorname{rk} M_j$$

oracle calls on at most $n^{\#\mathrm{rc}_i - \operatorname{rk} M_i}$ lattices of rank $\operatorname{rk}\Lambda + i$ each.

Proof. We proceed with induction on k. If $k = 0$, then we may return $\operatorname{rep}_0(\mathbf{A}) = n \cdot [\,]$ without any oracle calls. Suppose now that $k > 0$ and that the proposition holds for all integers less than k. We precompute a set of $\operatorname{rk} M_k$ rows $S \subseteq \mathbb{F}^{k \times k-1} \times \mathrm{rc}_{k-1}$ of M_k such that the induced submatrix of M_k has the same rank as M_k. In turn, we obtain a set $P_k = \{\mathbf{P} \mid (\mathbf{P}, c) \in S\}$ of cardinality at most $\operatorname{rk} M_k$.

Suppose now we have some input $(\Lambda, \mathbf{A})$ to the algorithm. Compute using the induction hypothesis $\operatorname{rep}_{k-1}(\mathbf{AP})$ for each $\mathbf{P} \in P_k$. The number of remaining possible solutions to $\operatorname{rep}_k(\mathbf{A})$ then is at most $n^{\#\mathrm{rc}_k - \operatorname{rk} M_k}$, which we may compute and test using the oracle by Proposition 8.6. One now simply verifies that the number of oracle calls is as claimed. □

Some final non-asymptotic improvements can be extracted from the inner products as shown by the following example.

Example 8.11. Consider the case where $\mathcal{O} = \Gamma = \mathbb{Z}$ and $\mathfrak{p} = 2\mathbb{Z}$. Then for each $\mathbf{x} = (x_1, \ldots, x_n) \in \Gamma^n$ we have

$$\langle \mathbf{x}, \mathbf{x} \rangle = \sum_i x_i^2 \equiv \#\{i \mid x_i \equiv 1 \pmod{\mathfrak{p}}\} \equiv n - z(\mathbf{x}) \pmod 4.$$

Hence we may disregard roughly $\frac{3}{4}$ of all candidate lattices B when testing for isomorphism without using the ΔLIP oracle. It turns out, by the following proposition, that the value of $\langle \Lambda', \Lambda' \rangle \pmod 4$ in fact determines the genus of Λ' modulo $\mathfrak{p}$ with 2 exceptions.

Proposition 8.12. *Let $\mathbb{Z}_2$ be the 2-adic integers and G the orthogonal group of $\mathbb{Z}_2^n$. Then the orbits of $\mathbb{F}_2^n$ under G are $O_t = \{\boldsymbol{x} \in \mathbb{F}_2^n \setminus \{\mathbf{0}, \mathbf{1}\} \mid \langle \boldsymbol{x}, \boldsymbol{x} \rangle \equiv t \ (4)\}$ for $t \in \mathbb{Z}/4\mathbb{Z}$, $\{\mathbf{0}\}$ and $\{\mathbf{1}\}$.*

Proof. It is clear that $\{\mathbf{0}\}$ and the O_t are respected. For $\{\mathbf{1}\}$, note that $\mathbf{1}$ is the unique vector $\boldsymbol{y} \in \mathbb{F}_2^n$ satisfying $\langle \boldsymbol{y}, \boldsymbol{x} \rangle = \langle \boldsymbol{x}, \boldsymbol{x} \rangle$ for all $\boldsymbol{x} \in \mathbb{F}_2^n$, and thus is fixed by G. It remains to show that G acts transitively on O_t. For $\boldsymbol{z} \in \mathbb{Z}_2^n$, the reflection $\rho(\boldsymbol{z})$ given by $\boldsymbol{y} \mapsto \boldsymbol{y} - 2(\langle \boldsymbol{y}, \boldsymbol{z} \rangle / \langle \boldsymbol{z}, \boldsymbol{z} \rangle)\boldsymbol{z}$ respects $\mathbb{Z}_2^n$ if $\langle \boldsymbol{z}, \boldsymbol{z} \rangle \not\equiv 0 \pmod 4$. Write $\boldsymbol{y}_i = (1, \ldots, 1, 0, \ldots, 0)$, where i is the number of ones. Then for $0 < i < j < n$ with $i \equiv j \pmod 4$ note that $\rho(\boldsymbol{y}_{j+1} - \boldsymbol{y}_{i-1})(\boldsymbol{y}_i) \equiv \boldsymbol{y}_{j+1} - \mathbf{e}_i \pmod 2$. Hence G maps a vector with support size i to a vector with support size j. We are done because G contains the full symmetric group permuting the coordinates. □

The above example can be generalized to any ΔLIP oracle call by first computing the genus invariant of the input lattice in polynomial-time, to reduce the number of possible isomorphic candidates.

9 Example: HAWK

So far we have considered the number ring $\mathcal{O}$ and the base $\mathcal{O}$-lattice Γ fixed and only considered the complexity of our reduction in terms of the number of orthogonal copies n, i.e., for $\Lambda \cong \Gamma^n$. It is common in cryptographic schemes, however, to consider lattices of low rank over large degree number fields for efficiency purposes. In this section we provide a similar polynomial-time search to distinguish reduction for the module lattice family used in the signature scheme HAWK. This reduction hinges on two important facts: one can pick $\mathfrak{p} \subset \mathcal{O}$ such that the residue field $\mathcal{O}/\mathfrak{p}$ does not change as $\mathcal{O}$ varies, and Γ has rank 1 over $\mathcal{O}$. In particular, the constant e in Theorem 7.3 is unchanged. A more subtle issue is that Theorem 7.3 currently assumes that one can solve sLIP for the base lattice (the case $n = 1$). This was fine for our reduction focusing on the complexity in terms of n, but this step can quickly become infeasible as $\mathcal{O}$ grows. When Γ has rank 1, there is a solution that makes no use of the oracle: the Gentry-Szydlo algorithm [11] and its extensions by Lenstra and Silverberg [15,17,18] can solve sLIP for rank-1 lattices in polynomial time when $\mathcal{O}$ is a cyclotomic ring or more generally, a CM-order, respectively. Furthermore, under a light number-theoretic heuristic the Gentry-Szydlo algorithm is extended to any number field in the full version of [3].

Let $(\mathcal{O}, \mathfrak{p}, \Gamma) = (\mathbb{Z}[\zeta_{2^\ell}], \mathcal{O}\pi, \mathcal{O})$, where ζ_{2^ℓ} is a primitive root of unity of order 2^ℓ for some (variable) $\ell > 0$, and $\pi = 1 - \zeta_{2^\ell}$ generates the unique prime $\mathfrak{p}$ of $\mathcal{O}$ above 2, which has norm 2. Note that K has degree $d = 2^{\ell-1}$. First we show that for HAWK the base case and the first few induction steps require no oracle calls.

Lemma 9.1. *There exists a polynomial-time algorithm that, given $\ell > 0$ and a lattice Λ isomorphic to $\mathcal{O}^2$ for $\mathcal{O} = \mathbb{Z}[\zeta_{2^\ell}]$, computes $I_{\mathfrak{p}^\ell}(\Lambda)$.*

Proof. As in Example 8.11 we may compute for $\boldsymbol{x} \in \Lambda$ the number of $L \in I_{\mathfrak{p}}(\Lambda)$ such that $\boldsymbol{x} \in L$. Hence we may compute a basis $(\boldsymbol{x}_1, \boldsymbol{x}_2)$ such that $\mathcal{O}\boldsymbol{x}_i + \mathfrak{p}\Lambda \in I_{\mathfrak{p}}(\Lambda)$. Since $\mathrm{Aut}(\mathcal{O}) = \{\zeta_{2^\ell}^i \mid i \in \mathbb{Z}\}$ maps surjectively to the unit group of $\mathcal{O}/\mathfrak{p}^\ell$, we may assume $\boldsymbol{x}_1 \equiv (1, y_1) \pmod{\mathfrak{p}^\ell}$ and $\boldsymbol{x}_2 \equiv (y_2, 1) \pmod{\mathfrak{p}^\ell}$ for some unknown $y_1, y_2 \in \mathfrak{p}$. Then, from the inner products one may iteratively compute the $y_i \pmod{\mathfrak{p}^k}$ and proceed as in Proposition 4.3. □

Proposition 9.2. *Suppose we have an oracle for ΔLIP. There exists a polynomial-time algorithm that, given $\ell, k > 0$, a lattice Λ isomorphic to $\mathcal{O}^2$ for $\mathcal{O} = \mathbb{Z}[\zeta_{2^\ell}]$, and $I_{\mathfrak{p}^k}(\Lambda)$, computes $I_{\mathfrak{p}^{k+1}}(\Lambda)$ in one oracle call on two lattices.*

Proof. Write $I_{\mathfrak{p}^k}(\Lambda) = \{\mathcal{O}\boldsymbol{x}_1 + \mathfrak{p}^k\Lambda, \mathcal{O}\boldsymbol{x}_2 + \mathfrak{p}^k\Lambda\}$. Then exactly one $\boldsymbol{x}_1' \in \{\boldsymbol{x}_1, \boldsymbol{x}_1 + \pi^k\boldsymbol{x}_2\}$ satisfies $\mathcal{O}\boldsymbol{x}_1' + \mathfrak{p}^{k+1}\Lambda \cong \mathcal{O}(1,0) + \mathfrak{p}^{k+1}\mathcal{O}^2$, which we test using the oracle.

Hence we may assume $\mathcal{O}\boldsymbol{x}_1 + \mathfrak{p}^{k+1}\Lambda \in I_{\mathfrak{p}^{k+1}}(\Lambda)$. To do the same for $\boldsymbol{x}_2$ we now do not need the oracle: we may take $\boldsymbol{x}_2' \in \{\boldsymbol{x}_2, \boldsymbol{x}_2 + \pi^k \boldsymbol{x}_1\}$ for which $\langle \boldsymbol{x}_1, \boldsymbol{x}_2' \rangle_K \equiv 0 \pmod{\mathfrak{p}^{k+1}}$. □

Lemma 9.1 and Proposition 9.2 allow us to prove the following analog of Theorem 7.3.

Theorem 9.3. *There is a polynomial-time algorithm that, given the order $\mathcal{O} = \mathbb{Z}[\zeta_{2^\ell}]$ of degree $d = 2^{\ell-1}$ for some $\ell > 0$, an $\mathcal{O}$-lattice Λ isomorphic to $\mathcal{O}^2$, and an oracle for $\mathcal{O}$-ΔLIP, computes an isomorphism $\Lambda \to \mathcal{O}^2$ using at most $4^{\ell-1} - \ell < 2d^2$ oracle calls on 2 lattices each.*

Proof. We proceed as in Theorem 7.3. For $4/3 < c < 2$, we may take $k = 4^{\ell-1}$ in Equation 1. We compute $I_{\mathfrak{p}^\ell}(\Lambda)$ using Lemma 9.1, and subsequently $I_{\mathfrak{p}^k}(\Lambda)$ using Proposition 9.2 with $k - \ell$ oracle calls. Then we may compute $I(\Lambda) = \{L_1, L_2\}$ using Theorem 7.1. Finally, since Γ has rank 1 and $\mathcal{O}$ is a cyclotomic ring, we may use the Gentry–Szydlo algorithm [11,15,17,18,3] to efficiently compute isomorphisms $L_i \to \Gamma$ and in turn the isomorphism $\Lambda \to \Gamma^2$. □

Remark 9.4. The two lattices passed to the oracle in Proposition 9.2 are non-trivial to distinguish for $k > 1$. In particular, they lie in the same genus and even in the same special genus. The distinguisher from [21] in combination with our search to distinguish reduction does therefore not give an efficient attack on HAWK. For $k = 1$ one can distinguish the two lattices similar to Example 8.11, as already used in the proof of Lemma 9.1.

References

1. L. Ackermann, A. Roux-Langlois, and A. Wallet. Public-key encryption from the lattice isomorphism problem. In *WCC 2024-The Thirteenth International Workshop on Coding and Cryptography*, pages 1–11, 2024.
2. M. R. Albrecht, B. Benčina, and R. W. Lai. Hollow LWE: A new spin, unbounded updatable encryption from LWE and PCE. In *EUROCRYPT 2025: Annual International Conference on the Theory and Applications of Cryptographic Techniques*, 2025.
3. B. Allombert, A. Pellet-Mary, and W. van Woerden. Cryptanalysis of rank-2 module-LIP: a single real embedding is all it takes. In *EUROCRYPT 2025: Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 184–212. Springer, 2025. Full version at https://eprint.iacr.org/2025/280.
4. H. Bennett, A. Ganju, P. Peetathawatchai, and N. Stephens-Davidowitz. Just how hard are rotations of $\mathbb{Z}^n$? algorithms and cryptography with the simplest lattice. In *EUROCRYPT 2023: Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 252–281. Springer, 2023.
5. H. Cohen. *Advanced topics in computational number theory*, volume 193. Springer Science & Business Media, 2012.
6. J. H. Conway and N. J. A. Sloane. *Sphere packings, lattices and groups*, volume 290. Springer Science & Business Media, 2013.

7. G. de Castro Biage, G. Zambonin, T. B. Idalino, D. Panario, and R. Custodio. A concrete LIP-based KEM with simple lattices. *IEEE Access*, 12:16408–16420, 2024.
8. L. Ducas, E. W. Postlethwaite, L. N. Pulles, and W. van Woerden. Hawk: Module LIP makes lattice signatures fast, compact and simple. In *ASIACRYPT 2022: International Conference on the Theory and Application of Cryptology and Information Security*, pages 65–94. Springer, 2022.
9. L. Ducas and W. van Woerden. On the lattice isomorphism problem, quadratic forms, remarkable lattices, and cryptography. In *EUROCRYPT 2022: Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 643–673. Springer, 2022.
10. M. Eichler. Note zur Theorie der Kristallgitter. *Mathematische Annalen*, 125(1):51–55, 1952.
11. C. Gentry and M. Szydlo. Cryptanalysis of the revised NTRU signature scheme. In *EUROCRYPT 2002: International conference on the theory and applications of cryptographic techniques*, pages 299–320. Springer, 2002.
12. K. Jiang, A. Wang, H. Luo, G. Liu, T. Gang, Y. Pan, and X. Wang. Re-randomize and extract: A novel commitment construction framework based on group actions. In *EUROCRYPT 2025: Annual International Conference on the Theory and Applications of Cryptographic Techniques*, 2025.
13. V. Karemaker, A. Tamagawa, and C.-F. Yu. Uniqueness of indecomposable idempotents in algebras with involution. *arXiv preprint arXiv:2402.09323*, 2024.
14. X. T. Khuc, A. T. Ta, W. Susilo, D. H. Duong, F. Guo, K. Fukushima, and S. Kiyomoto. Logarithmic-size (linkable) ring signatures from lattice isomorphism problems. In *International Conference on Security, Privacy, and Applied Cryptography Engineering*, pages 214–241. Springer, 2023.
15. H. W. Lenstra and A. Silverberg. Revisiting the Gentry–Szydlo algorithm. In *Advances in Cryptology–CRYPTO 2014: 34th Annual Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2014, Proceedings, Part I 34*, pages 280–296. Springer, 2014.
16. H. W. Lenstra Jr. Lattices. In *J. P. Buhler and P. Stevenhagen (eds.), Algorithmic Number Theory*, volume 44 of *MSRI Publications*, pages 127–181. Cambridge University Press, 2008. http://library.msri.org/books/Book44/.
17. H. W. Lenstra Jr and A. Silverberg. Lattices with symmetry. *Journal of Cryptology*, 30:760–804, 2017.
18. H. W. Lenstra Jr and A. Silverberg. Testing isomorphism of lattices over CM-orders. *SIAM Journal on Computing*, 48(4):1300–1334, 2019.
19. H. Luo, K. Jiang, Y. Pan, and A. Wang. Commitment schemes based on module-LIP. *Cryptology ePrint Archive*, 2025.
20. J. Martinet. *Perfect lattices in Euclidean spaces*, volume 327. Springer Science & Business Media, 2013.
21. G. Mureau. Special genera of hermitian lattices and applications to HAWK. *Cryptology ePrint Archive*, 2025.
22. G. Mureau, A. Pellet-Mary, G. Pliatsok, and A. Wallet. Cryptanalysis of rank-2 module-LIP in totally real number fields. In *EUROCRYPT 2024: Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 226–255. Springer, 2024.
23. M. Szydlo. Hypercubic lattice reduction and analysis of GGH and NTRU signatures. In *EUROCRYPT 2003: International Conference on the Theory and Applications of Cryptographic Techniques*, pages 433–448. Springer, 2003.