A reduction from Hawk to the principal ideal problem in a quaternion algebra

Clémence Chevignard[1], Guilhem Mureau[2], Thomas Espitau[3], Alice Pellet-Mary[2], Heorhii Pliatsok[4], and Alexandre Wallet[3]

[1] Univ Rennes, Inria, CNRS, Irisa, UMR 6074, France
clemence.chevignard@inria.fr
[2] Univ Bordeaux, CNRS, Inria, Bordeaux INP, IMB, UMR 5251, Talence, France
guilhem.mureau@math.u-bordeaux.fr
alice.pellet-mary@math.u-bordeaux.fr
[3] PQ Shield Ltd., United Kingdom
thomas.espitau@pqshield.com
alexandre.wallet@pqshield.com
[4] Insitute of Mathematics, NAS of Ukraine
georgiipliatsok@icloud.com

Abstract. In this article we present a non-uniform reduction from rank-2 module-LIP over Complex Multiplication fields, to a variant of the Principal Ideal Problem, in some fitting quaternion algebra. This reduction is classical deterministic polynomial-time in the size of the inputs. The quaternion algebra in which we need to solve the variant of the principal ideal problem depends on the parameters of the module-LIP problem, but not on the problem's instance. Our reduction requires the knowledge of some special elements of this quaternion algebras, which is why it is non-uniform.

In some particular cases, these elements can be computed in polynomial time, making the reduction uniform. This is the case for the Hawk signature scheme: we show that breaking Hawk is no harder than solving a variant of the principal ideal problem in a fixed quaternion algebra (and this reduction is uniform).

1 Introduction

Two lattices L, L' are said isomorphic when there exists a linear isometry between them, and the Lattice Isomorphism Problem (LIP) asks to compute such an isometry. It has been studied in [18,25,36] as a standalone algorithmic problem, and these works achieve superexponential complexity — at best $n^{O(n)}$ time in the worst-case — for lattices of rank n. Stemming from this apparent hardness, LIP has recently been introduced as a security assumption to find cryptographic primitives in [1,3,16]. It joins other isomorphism-finding-based assumptions already in use in multivariate or code-based cryptography. Soon after, the signature

This article is the result of a merge between [9] (which had the same title) and [20].

scheme Hawk was presented [15] and submitted[5] to the second call for post-quantum digital signatures organized by the NIST. It relies on a structured variant of LIP called module-LIP. In this variant, L and L' are now module lattices (a transition identical to that of LWE to module-LWE) and an isometry compatible with the module structure must be found. The resulting scheme demonstrates slightly better efficiency and signature sizes compared to Falcon and Dilithium, the two lattice-based signatures selected by NIST during the first call [35]. Owing to its recent cryptographic introduction, the cryptanalysis of module-LIP, and thus of Hawk, is however quite young, making it an attractive target for cryptanalysts.

In the simplest version of module-LIP [15], an attacker is given a (module-compatible) rotation of $\mathcal{O}_K^2$, where $\mathcal{O}_K$ is the ring of algebraic integers of a number field K, and is asked to recover the corresponding isometry. As there may be many more symmetries linked to the algebraic structure of K, it can be hoped that finding isometries of module lattices can be an easier task than for the plain case. At Eurocrypt 2024, Mureau et al. [33] focused on the case of *totally real*[6] number fields and proposed a (heuristic) algorithm to solve module-LIP over such fields. In the special case of the module $\mathcal{O}_K^2$ and for some totally real number fields, this algorithm runs in polynomial time. On the one hand, this confirmed the intuition that module-LIP could be significantly easier than LIP (in our current state of knowledge). But, on the other hand, the current representative of module-LIP-based schemes, Hawk, is *not* designed over totally real fields. Instead, it is designed over the prevalent power-of-two cyclotomic fields, which are by nature totally imaginary. One notes that a cyclotomic field $K = \mathbb{Q}(\zeta)$ always comes with a totally real maximal subfield $F = \mathbb{Q}(\zeta + \zeta^{-1})$, but the authors of [33] could not use this to their advantage to extend their algorithm to Hawk's design. This work aims at narrowing this gap.

From LIP to quaternion ideals. In the case of Hawk, we are given a public Gram-matrix $G = \overline{C}^T C$ corresponding to a secret basis C of $\mathcal{O}_K^2$:

$$C = \begin{pmatrix} a & c \\ b & d \end{pmatrix} \quad \text{and} \quad G = \begin{pmatrix} q_{11} & q_{12} \\ \overline{q_{12}} & q_{22} \end{pmatrix}.$$

Our first observation is that the entire geometry of the problem can be fully understood by only the first secret vector $(a\ b)$. Intuitively, once we know this vector, we can examine the (scaled by q_{11}-)Gram-Schmidt orthonormalization over K of the secret matrix C, that is, the decomposition of C over the basis spanned by $(a\ b)$ and its K-orthogonal $(-\overline{b}\ \overline{a})$.

Indeed, the value of q_{22} perfectly determines the component of $(c\ d)$ over $(-\overline{b}\ \overline{a})$, while q_{12} determines its component over $(a\ b)$ itself. Thus, we may wish

[5] https://hawk-sign.info/

[6] Any number field comes with a set of embeddings into either $\mathbb{R}$ or $\mathbb{C}$. The field is said totally real (resp. totally imaginary) when all these embeddings map to $\mathbb{R}$ (resp. none of these embeddings maps to $\mathbb{R}$).

to study the problem through the lens of $(a\ b)$ (and its orthogonal $(-\overline{b}\,\overline{a})$), more specifically focusing on recovering the lattice $\mathcal{L}$ spanned by the transpose of these two vectors, i.e.,

$$\mathcal{L} := \mathcal{L}(A) \quad \text{with} \quad A := \begin{pmatrix} a & -\overline{b} \\ b & \overline{a} \end{pmatrix}^T = \begin{pmatrix} a & b \\ -\overline{b} & \overline{a} \end{pmatrix}.$$

Elementary linear algebra on the matrix A and G reveals that this lattice is also spanned by the basis $A' = \begin{pmatrix} 1 & 0 \\ q_{12} & q_{11} \end{pmatrix}$, which is entirely determined by the public knowledge of G. Note that the vector $(a\ \overline{b})$ belongs to $\mathcal{L}$, and an inspection of its K-norm (i.e., $a\overline{a} + b\overline{b} = q_{11}$) shows that it *must* be a shortest vector of $\mathcal{L}$. Consequently, we can recover it — or one of its rotations — with a single module-SVP instance. Since in the case of Hawk, the descent over $\mathbb{Z}$ of this lattice is hypercubic, advanced algorithms can even be employed to solve this instance [17]. However, the general case of rank-2 module-LIP involves less geometrically structured modules, not always of the form $\mathcal{O}_K^2$. Addressing these objects at an arithmetic level rather than a geometric one is often better suited.

To do so, let us take a step back: intuitively, the lattice $\mathcal{L}$ is spanned by $(a\ \overline{b})$ and a "conjugate" of this vector, as the image of an involution of the algebra K^2. Thus, this rank 2 lattice is effectively a *one-dimensional object*, as being determined by a single vector in K^2. Examining the matrix $M(a,b) := \begin{pmatrix} a & -b \\ \overline{b} & \overline{a} \end{pmatrix}$ reveals its similarity to the matrix representation of complex numbers, i.e., the representation $\mathbb{C} \ni x + yi \mapsto \begin{pmatrix} x & -y \\ y & x \end{pmatrix} \in \mathcal{M}_2(\mathbb{R})$. When K is totally real, we can mimic this representation and work with the element $a + ib$ in the quadratic extension $K(i)/K$. However, when K already contains the imaginary unit i, this identification is no longer possible. In such cases, we can artificially introduce an additional root of unity, say j, satisfying $j^2 = -1$, which acts as the element $\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$. Now the map $M(a,b) \mapsto a + bj$ is a morphism of algebras, with $ij = -ji$. This means that $a + bj$ belongs to a *quaternion algebra* over K, and relates the lattice $\mathcal{L}$ to the *principal* ideal $I_{\mathcal{L}}$ spanned by this element. Additionally, we are given $q_{11} = \det M(a,b) =: \mathrm{nrd}(a + bj)$, with nrd called the reduced norm of the algebra. From the geometric interpretation, the ideal $I_{\mathcal{L}}$ can be described using only public information, and we are thus solving a *principal ideal problem* (PIP) instance, given the *reduced norm* of a generator. When K is a totally real field, this problem is solved in polynomial time by the Gentry-Szydlo algorithm [23], and its generalization by Lenstra and Silverberg [27]. However, there is currently no such counterpart for noncommutative quaternion algebras. Turning to the general case in the hope for a complete reduction, we face several challenges stemming from this noncommutativity. While being somewhat common objects in isogeny-based cryptography, such structures have less exposure or involvement in lattice-based cryptography (they can be seen as a particular case of cyclic algebras, studied in a lattice context in e.g. [30]).

Contributions. Our main contribution is a general reduction from the rank-2 version of module-LIP over complex multiplication (CM) fields,[7] to the reduced-norm Principal Ideal Problem (nrdPIP). This latter problem consists of computing a particular generator of a principal ideal in a suitable quaternion algebra over K, given the reduced norm of the generator we are looking for. A notable particular case, that includes Hawk's instances, is the following:

Theorem 1.1 (Informal, special case of Corollary 4.4). *Let K be a cyclotomic field of degree d and let $G = V^*V \in \mathcal{M}_2(\mathcal{O}_K)$ with $V \in \mathbf{GL}_2(\mathcal{O}_K)$ a basis of $\mathcal{O}_K^2$. Given access to an oracle solving nrdPIP, computing a matrix $U \in \mathbf{GL}_2(\mathcal{O}_K)$ such that $U^*U = G$ can be done in time polynomial in d and the size of G, by making only one call to the nrdPIP oracle.*

The general version of this statement for all CM number fields and all rank-2 module lattices can be found in Corollary 3.9. In essence, the crux of the method is representing a module of rank 2 over a number field as an ideal — i.e. a module of rank 1 — over a ring of dimension twice larger. We stress that when the field where LIP needs to be solved is a cyclotomic field, there are *no known* polynomial time algorithm to solve the quaternionic version of the principal ideal problem with given reduced norm. In other words, this work *does not* break Hawk. Our reduction rather shows that any improvement for solving the nrdPIP problem (or SVP in ideals of quaternion algebras) would directly impact the hardness of rank-two module-LIP and the security of Hawk.

To the best of our knowledge, the best algorithm solving the nrdPIP instances generated by our reduction is due to Kirschmer and Voight [29, Alg. 6.3]. Instantiating the nrdPIP oracle with this algorithm proves in particular that a single call to an SVP solver in dimension $2d$ is sufficient to break Hawk (a fact that seemed folklore so far, but was never proven anywhere to the best of our knowledge). In the original Hawk article [15], the authors explain in Section 4.2 that the best algorithms solving (module-)LIP require *at least* one SVP call, so to be conservative they "assume that the best key recovery attack requires one to find a single shortest vector". Our result shows that this assumption is tight: one SVP call is indeed enough for a key recovery attack.

While our reduction provides an easy way to prove this fact about Hawk, it is probably overkill: the underlying module is free, has rank two, and has many orthogonal shortest vectors. There are probably more straightforward ways to show that a single SVP call in a large lattice is enough. For *general rank*-2 modules over a CM-field K, there does not seem to be a Karp reduction anymore (i.e., a reduction making only one call to an SVP oracle): instead we are only able to reduce the module-LIP problem in any *rank-2* module of K^2 to *two* instances of the nrdPIP problem (which can then be reduced to two instances of SVP in ideal lattices of a non-commutative ring of dimension $2d$).

[7] A CM field K is a totally imaginary field which is a degree 2 extension of some totally real subfield F.

Related works. The Principal Ideal Problem over a number field (say, F-PIP) has been coined as a central problem in algorithmic number theory (e.g. [10]).

In an arbitrary number field, the state-of-the-art classical algorithms are heuristic and run in subexponential time [6,4] or quantum polynomial time [5]. We note that all these algorithms reduce to the problem of computing the unit group and the class group of the underlying field. In lattice-based cryptography, F-PIP appeared in important results [12,13] on the hardness of the Ideal-SVP problem. In this article we encounter a variant of this problem over (totally definite) quaternion algebras, say $\mathcal{A}$-PIP. In this context, an algorithm to compute a generator of a principal ideal $I \subset \mathcal{A}$ is provided in [29, Alg. 6.3]. The strategy reduces to the computation of the class group of F and to a short vector computation in a rank-$2d$ $\mathbb{Z}$-lattice, where d is the degree of K.

While computing the class group may be done in quantum polynomial time, computing short vectors in lattices is believed to be hard even for quantum computers. For more general algebras, Bley et al. [8] give an algorithm solving PIP by reducing it to many subproblems, including PIP in K. We note that their work also provides algorithms to compute isomorphisms between finitely generated modules over number fields, but that these are *not* isometries between modules lattices. In other words, they are not lattice isomorphisms in the sense we are interested in.

With the additional information of the reduced norm of a generator of a principal ideal (say, F-nrdPIP), the situation can change drastically and (classical) polynomial time complexity can be achieved for CM extensions. For cyclotomic fields, this observation goes back to Gentry and Szydlo's algorithm [23] to attack NTRU encryption. Variants of this algorithm [26,22,19,21] were subsequently used to attack lattice-based signatures in several context, and a more general version was described by Lenstra and Silverberg [27], covering in particular all CM fields. On the other hand, for the quaternion variant $\mathcal{A}$-nrdPIP, there are (to our knowledge) no known polynomial time algorithms, and thus the problem is solved by using a $\mathcal{A}$-PIP solver instead.

This work can be viewed as an extension of the reduction of Mureau et al. [33] to cover the case of CM-fields, which include cyclotomic fields. Our reduction technique subsumes theirs, improving on their polynomial time algorithm to solve the problem over totally real fields, and additionally removing the need for a heuristic assumption and the dependency in ρ_K in their complexity.[8]

Luo, Jiang, Pan, and Wang published a concurrent work [31], reducing module-LIP over a CM field to the problem of finding a specific (symplectic) automorphism. With the knowledge of this automorphism, finding the correct congruence matrix reduces to a $\mathfrak{O}$-nrdPIP instance in a *commutative* ring $\mathfrak{O}$, where the polynomial time algorithm of Lenstra and Silverberg [27] applies. This "reduction to rank one" has a flavor similar to our results, and one can interpret the knowledge of the symplectic automorphism as a way to bypass the non-commutativity of quaternions. In the totally real case, that is in the context of $K(i)/K$, the

[8] ρ_K is the so-called residue at 1 of the Dedekind zeta function of the field.

authors give an alternative presentation of our reduction using the automorphism formalism (corresponding to Section 4 of [31]).

2 Preliminaries

For a ring R, we denote by $R^\times$ its set of invertible elements (that is, whose inverses are in R). The set of $n \times n$ matrices with entries in R is denoted by $\mathcal{M}_n(R)$ and the subset of invertible matrices forms the group $\mathbf{GL}_n(R)$. We use bold letters to denote vectors. For reasons of space, the proofs of the results stated further in this section without references are deferred to Appendix A.1.

2.1 Number Fields

Generalities A number field K is a finite extension of the field of rational numbers $\mathbb{Q}$. It is isomorphic to $\mathbb{Q}[X]/P(X)$, where $P(X)$ is an irreducible monic polynomial of $\mathbb{Q}[X]$. The degree $d := [K : \mathbb{Q}]$ of K over $\mathbb{Q}$ is exactly the degree of $P(X)$. A number field K of degree d has d embeddings $K \to \mathbb{C}$. Any embedding $\sigma_i : K \to \mathbb{R}$ is called a real embedding. An embedding σ_i which is not real is called complex, and it can be composed with the complex conjugation in $\mathbb{C}$ to obtain a different complex embedding $\overline{\sigma_i}$. The canonical embedding of K is defined as $\sigma(e) := (\sigma_1(e), \ldots, \sigma_d(e)) \in \mathbb{C}^d$, for all $e \in K$ and where the σ_i are all the embeddings of K. We extend it coordinate-wise to K^ℓ.

The space $K_\mathbb{R}$ is defined[9] as $K_\mathbb{R} = K \otimes_\mathbb{Q} \mathbb{R}$. Then the canonical embedding of K extends to $K_\mathbb{R}$ and its image is isomorphic to the real subspace $\mathcal{H} = \{(x_1, \ldots, x_d) \ : \ x_{d_1+i} = \overline{x_{d_1+d_2+i}} \text{ for } 1 \leq i \leq d_2\} \subset \mathbb{R}^{d_1} \times \mathbb{C}^{2d_2} \subset \mathbb{C}^d$ where d_1 is the number of real embeddings of K and $2d_2$ the number of complex embeddings. Through this identification $K_\mathbb{R}$ is equipped with a complex conjugation $\overline{\cdot}$ which amounts to taking the complex conjugate coordinate-wise. We denote by $\mathcal{U}_n(K_\mathbb{R})$ the set of $n \times n$ unitary matrices with coefficients in $K_\mathbb{R}$, *i.e.*, the set of $\Theta \in \mathcal{M}_n(K_\mathbb{R})$ such that $\Theta^*\Theta = \mathrm{Id}$, where Θ^* is the transpose-conjugate of Θ.

When all the embeddings of $K \to \mathbb{C}$ are real, we say that K is *totally real*. When none of them are, we say that K is *totally complex*. An element $a \in K$ is totally positive, resp. totally negative, if all its embeddings are positive, resp. negative real numbers (in particular, they are real numbers). The (absolute) trace and norm of $e \in K$ is $\mathrm{Tr}(e) = \sum_i \sigma_i(e)$ and $N(e) = \prod_i \sigma_i(e) \in \mathbb{Q}$.

We note $\mathcal{O}_K$ the ring of algebraic integers of a number field K. It is a free $\mathbb{Z}$-module of rank d. The (absolute) discriminant of K is $\Delta_K = |\det([\mathrm{Tr}(\beta_i\beta_j)]_{i,j})|$, for any $\mathbb{Z}$-basis $(\beta_i)_i$ of $\mathcal{O}_K$.

CM fields. A Complex Multiplication (CM) number field K is a totally complex quadratic extension of a totally real number field F (we also say K/F is a CM-extension). Equivalently, here F is a totally real number field and there exists a totally negative element $a \in F$ such that $K = F(\sqrt{a})$ [42, Page 38]. In a

[9] If $K \simeq \mathbb{Q}[X]/(P)$, then one has $K_\mathbb{R} \simeq \mathbb{R}[X]/(P)$.

CM-extension K/F, there is a unique non-trivial automorphism of K fixing F pointwise, which is called the complex conjugation. With the notation $K = F(\sqrt{a})$, it acts on K by $\tau : \sqrt{a} \mapsto -\sqrt{a}$. In particular, the *relative* norm for the extension K/F is defined by $N_{K/F}(x + y\sqrt{a}) := (x + y\sqrt{a}) \cdot \tau(x + y\sqrt{a}) = (x + y\sqrt{a})(x - y\sqrt{a}) = x^2 - ay^2$, for all $x + y\sqrt{a} \in K$. The following lemma justifies why the automorphism τ is also called complex conjugation.

Lemma 2.1 ([33, Lemma 2.7]). *Let K/F be a CM extension of number fields. For any embedding $\sigma_i : K \to \mathbb{C}$ and $x \in K$, we have $\overline{\sigma_i(x)} = \sigma_i(\tau(x))$.*

To simplify notations in the rest of this article, we write $\overline{x}$ instead of $\tau(x)$.

Roots of unity. From a celebrated result of Kronecker, one can characterize the roots of unity among algebraic integers in a number field. It is also useful to deal with equations of the form $a\overline{a} + b\overline{b} = 1$ (for a, b in the ring of integers of a CM number field K), as stated below (and proved in Appendix A). These are involved later in the paper, see Proposition 3.13 and Corollary B.23.

Proposition 2.2 ([14, Theorem K]). *Let K be a CM field and $a \in \mathcal{O}_K$. If a is non-zero and if all its conjugates have absolute value at most 1, then a is a root of unity.*

Corollary 2.3. *Let K be a CM field and $a, b \in \mathcal{O}_K$ be such that $a\overline{a} + b\overline{b} = 1$. Then either a is a root of unity and $b = 0$ or $a = 0$ and b is a root of unity.*

Computing the group $\mu(K)$ of roots of unity in a number field is handled by the following lemma. Note that this group is cyclic [34, 7.4].

Lemma 2.4 (Computing roots of unity [33, Cor. 2.11]). *Let K be a degree d number field. Then, K has at most $2d^2$ roots of unity and there is a polynomial time algorithm that given a basis of $\mathcal{O}_K$, computes the roots of unity in K.*

Ideals. An integral ideal $\mathfrak{a}$ of K is an additive subgroup of $\mathcal{O}_K$, such that for all $x \in \mathcal{O}_K$, $x\mathfrak{a} \subseteq \mathfrak{a}$. A fractional ideal $\mathfrak{a}$ of K is an additive subgroup of K such that for some $x \in K \setminus \{0\}$, $x\mathfrak{a}$ is an integral ideal. If $\mathfrak{a}$ is generated by a single element x, it is said to be *principal*, and is noted $\mathfrak{a} = x\mathcal{O}_K$. In general, fractional $\mathcal{O}_K$-ideals can all be generated using at most two elements, see [10, Proposition 4.7.7]. We will use fraktur letters to denote fractional ideals of K or F.

Let $\mathfrak{a}, \mathfrak{b}$ be two fractional ideals. The product $\mathfrak{ab}$ is the smallest ideal containing all products xy for $x \in \mathfrak{a}$ and $y \in \mathfrak{b}$. We have that $\mathfrak{a} \subseteq \mathfrak{b}$ if and only if there exists an integral ideal $\mathfrak{c}$ such that $\mathfrak{a} = \mathfrak{bc}$. When this is the case, we say that $\mathfrak{b}$ (equivalently $\mathfrak{c}$) divides $\mathfrak{a}$. An integral ideal $\mathfrak{p}$ is prime whenever $xy \in \mathfrak{p}$ implies $x \in \mathfrak{p}$ or $y \in \mathfrak{p}$. Prime ideals are the maximal ideals in $\mathcal{O}_K$. When dealing with number fields, we have a unique factorization of integral ideals into prime ideals (up to permutation). Ideals of the form $p\mathcal{O}_K$ for a prime integer p can be factored efficiently.

Lemma 2.5 ([11, Section 6.2.5]). *There exists a polynomial time algorithm that takes as input any prime integer $p \in \mathbb{Z}$ and a basis of the ring of integers $\mathcal{O}_K$ of a number field K, and computes all the prime ideals of $\mathcal{O}_K$ dividing $p \cdot \mathcal{O}_K$.*

When K/F is a CM extension and $\mathfrak{a}$ is a fractional ideal of K, the set $\overline{\mathfrak{a}} := \{\overline{x} \,|\, x \in \mathfrak{a}\}$ is again a fractional ideal of K, called the conjugate of $\mathfrak{a}$.

Modules. The main reference for this paragraph is the first chapter from [11]. Let V be a finite-dimensional vector space over a number field K. In this work, for the sake of clarity, we will call "module" any finitely generated, torsion-free $\mathcal{O}_K$-module in V. Such modules are always of the form $\mathfrak{a}_1\boldsymbol{b}_1 + \cdots + \mathfrak{a}_\ell\boldsymbol{b}_\ell$, where the $\mathfrak{a}_i$'s are fractional ideals in K and the $\mathbf{b}_i$'s are K-linearly independent vectors in V. The data of $((\boldsymbol{b}_1, \mathfrak{a}_1), \ldots, (\boldsymbol{b}_\ell, \mathfrak{a}_\ell))$ is called a pseudo-basis of M and the integer ℓ is called the rank of the module. We write $\mathbf{B} = (B, \{\mathfrak{a}_i\}_{i\leq\ell})$ where B is the (column) matrix of the $\mathbf{b}_i$'s, and we call it a pseudo-basis of the module. We use bold capital letters to denote pseudo-bases. In this work, we always consider modules with full rank, and let $\dim_K(V) = \ell$. A module $M \subset K^\ell$ (resp. $\mathcal{O}_K^\ell$) is said to be rational (resp. integer).

Two pseudo-bases $\mathbf{B} = (B, \{\mathfrak{a}_i\}_{1\leq i\leq\ell})$ and $\mathbf{C} = (C, \{\mathfrak{b}_i\}_{1\leq i\leq\ell})$ generate the same module if and only if there exists $U = (u_{i,j})_{1\leq i,j\leq\ell} \in \mathbf{GL}_\ell(K)$ such that $C = BU$ and $u_{i,j} \in \mathfrak{a}_i\mathfrak{b}_j^{-1}$ for all $1 \leq i, j \leq \ell$ and $\mathfrak{a}_1 \cdots \mathfrak{a}_\ell = (\det U)\mathfrak{b}_1 \cdots \mathfrak{b}_\ell$ ([11, Proposition 1.4.2]).

In a CM-field K, the pseudo-Gram matrix associated with a pseudo-basis $\mathbf{B} = (B, \{\mathfrak{a}_i\}_{i\leq\ell})$ is $\mathbf{G} = (B^*B, \{\mathfrak{a}_i\}_{i\leq\ell})$, where $B^* := \overline{B}^t$ is the conjugate-transpose matrix and the complex conjugation is taken coefficient-wise.[10]

Let $M, M' \subset K^\ell$ be modules and let $\Theta \in \mathcal{M}_\ell(K_\mathbb{R})$ such that $\Theta^*\Theta = Id$. When $M' = \Theta \cdot M$, we say that Θ is a *module lattice isomorphism* between M and M'. If moreover $M' = M$, we say Θ is a *module lattice automorphism*.

2.2 Quaternion algebras

We now give the background on quaternion algebras that is needed in this work. A general reference for this topic is [41], from which we borrow most of the material. For a field F, a F-algebra is an F-linear space which is also a ring (its elements can be multiplied together into another ring element). In this work, we are interested in one type of quaternion algebra, defined below. Recall that F is a totally real field, a is a totally negative element in F, so that $K = F(\sqrt{a})$ is a CM-extension.

Definition 2.6. *The quaternion algebra $\mathcal{A} := (\frac{a,-1}{F})$ is the F-algebra of dimension 4 with basis $\{1, i, j, ij\}$ and satisfying the rules*

$$i^2 = a \quad ; \quad j^2 = -1 \quad ; \quad ij = -ji.$$

[10] The pseudo-Gram matrix can be more generally defined for any number field [33, Definition 3.6], but in this work, we will only be interested in CM-field.

Because of the rule $ij = -ji$, $\mathcal{A}$ is a noncommutative algebra. Its center (the set of elements that commute with every other) is equal to F. A quaternion algebra is also equipped with an involution $\overline{\cdot}$ defined by $\overline{x + iy + jz + ijt} = x - iy - jz - ijt$. This map is F-linear and satisfies $\overline{\overline{\alpha}} = \alpha$ and $\overline{\alpha\beta} = \overline{\beta} \cdot \overline{\alpha}$ for any $\alpha, \beta \in \mathcal{A}$ (see [41, Section 3.2]). The reduced norm on $\mathcal{A}$ is the map $\mathrm{nrd} : \mathcal{A} \to F$ defined by $\alpha = x + iy + jz + ijt \mapsto \alpha\overline{\alpha} = x^2 - ay^2 + z^2 - at^2$. We have $\mathrm{nrd}(\alpha\beta) = \mathrm{nrd}(\alpha)\mathrm{nrd}(\beta)$ for all $\alpha, \beta \in \mathcal{A}$ [41, Par. 3.3.4]. The CM-extension $K = F(\sqrt{a})$ of F is included in $\mathcal{A}$, and when $x \in K$, we have $\mathrm{nrd}(x) = N_{K/F}(x)$.

Because of our choice for a, the quaternion algebras $(\frac{a,-1}{F})$ are said to be *totally definite*, as mentioned in the following proposition. In this article, we will not need to know precisely what a totally definite algebra is, but we will use results that hold only for totally definite algebras.

Proposition 2.7 ([24, Page 3], adapted). *If F is a totally real number field and $a \in F$ is totally negative, then the quaternion algebra $(\frac{a,-1}{F})$ is totally definite.*

A notable property of algebras of the form $\mathcal{A} = (\frac{a,-1}{F})$ as above is that they are division algebras — that is, all their elements are invertible, or equivalently, they are non-commutative fields. To see this, note that by definition, $-a$ is totally positive and thus so is $\mathrm{nrd}(\alpha) = x^2 - ay^2 + z^2 - at^2$. Since the embeddings of a are also non-zero, $\alpha \neq 0$ if and only if $\mathrm{nrd}(\alpha) \neq 0$. Then we can write $\alpha^{-1} = \mathrm{nrd}(\alpha)^{-1}\overline{\alpha}$, and hence conclude that for our algebras, $\mathcal{A}^\times = \mathcal{A}\backslash\{0\}$.

Quaternion orders and ideals. Let us fix a quaternion algebra $\mathcal{A} = (\frac{a,-1}{F})$ over a totally real field F. Recall that in this work modules are all finitely generated and torsion-free. We define $\mathcal{O}_F$-lattices and orders in $\mathcal{A}$.

Definition 2.8 ([41, Definition 9.3.1 and 10.2.1]). *An $\mathcal{O}_F$-lattice[11] in $\mathcal{A}$ is a $\mathcal{O}_F$-module contained in $\mathcal{A}$ and with full-rank in $\mathcal{A}$ (i.e., it is a rank-4 $\mathcal{O}_F$-module included in $\mathcal{A}$). An $\mathcal{O}_F$-order $\mathcal{O} \subseteq \mathcal{A}$ is an $\mathcal{O}_F$-lattice in $\mathcal{A}$ that is also a subring of $\mathcal{A}$ (in particular, $1 \in \mathcal{O}$). An $\mathcal{O}_F$-order of $\mathcal{A}$ is said to be maximal if it is not strictly contained in another $\mathcal{O}_F$-order.*

Similarly, one can define orders of $\mathcal{A}$ for different subrings of $\mathcal{A}$. In this article, we will only be interested in $\mathcal{O}_F$-orders, so we simply call them orders.

Lemma 2.9 ([41, Prop. 15.5.2], adapted). *In the quaternion algebra $\mathcal{A}$, there exists (at least) one maximal order, and every order $\mathcal{O}$ is contained in a maximal order.*

Contrary to the case of number fields where the ring of integers is the unique maximal order, there can be many maximal orders in a quaternion algebra.

Proposition 2.10 ([41, Lemma 10.2.7 and Definition 10.2.8]). *Let $I \subseteq \mathcal{A}$ be an $\mathcal{O}_F$-lattice. The set $\mathcal{O}_\ell(I) := \{x \in \mathcal{A} \,|\, xI \subseteq I\}$ is an order of $\mathcal{A}$, called the left order of I. Similarly, the set $\mathcal{O}_r(I) := \{x \in \mathcal{A} \,|\, Ix \subset I\}$ is an order of $\mathcal{A}$ called the right order of I.*

[11] Although related, these are not directly Euclidean lattices.

We denote by $\mathcal{O}^1$ the (multiplicative) subgroup of norm 1 elements in an order $\mathcal{O}$. In totally definite quaternion algebras, $\mathcal{O}^1$ is always finite, and $\mathcal{O}^1/\{\pm 1\}$ belongs to a known list of groups, up to isomorphism — see B.3 for details.

Given an order $\mathcal{O}$, a left (resp. right) fractional $\mathcal{O}$-ideal is an $\mathcal{O}_F$-lattice $I \subseteq \mathcal{A}$ satisfying $xI \subseteq I$ (resp. $Ix \subseteq I$) for all $x \in \mathcal{O}$. Since a left fractional $\mathcal{O}$-ideal is in particular an $\mathcal{O}_F$-lattice in $\mathcal{A}$, we can define its left order $\mathcal{O}_\ell(I)$. By definition, this order contains $\mathcal{O}$, but it can be larger. We say that I is a *sated* left fractional $\mathcal{O}$-ideal if $\mathcal{O} = \mathcal{O}_\ell(I)$ (i.e., if $\mathcal{O}$ is the largest order for which I is a left ideal) [41, Definition 16.2.11]. A similar definition holds for right $\mathcal{O}$-ideals. From these definitions, any $\mathcal{O}_F$-lattice $I \subseteq \mathcal{A}$ is a (sated) left fractional ideal for its left order $\mathcal{O}_\ell(I)$. In the rest of this section, we will review some definitions and lemmas from [41, Chapter 16], that extend similar results for ideals in number fields. These results will be stated for $\mathcal{O}_F$-lattices in $\mathcal{A}$ (but keep in mind that these are left fractional $\mathcal{O}$-ideals for some order $\mathcal{O}$, depending on the lattice).

A $\mathcal{O}_F$-lattice is principal if there exists $\alpha \in \mathcal{A}^\times$ such that $I = \alpha\mathcal{O}_r(I) = \mathcal{O}_\ell(I)\alpha$. For any $\mathcal{O}_F$-lattice I, if there exists $\alpha \in \mathcal{A}^\times$ such that $I = \alpha\mathcal{O}_r(I)$, then it also holds that $I = \mathcal{O}_\ell(I)\alpha$, see [41, 16.2.3]. Hence, to test if an $\mathcal{O}_F$-lattice is principal, it suffices to test if it is left (or right) principal. Generally, let $S \subset \mathcal{A}$ be a set generating a finitely generated module in $\mathcal{A}$ (for example, a fractional ideal in K). The left (resp. right) $\mathcal{O}$-ideal generated by S is the smallest fractional left (resp. right) $\mathcal{O}$-ideal of $\mathcal{A}$ containing the elements $\alpha \cdot s$ (resp. $s \cdot \alpha$), for $(\alpha, s) \in \mathcal{O} \times S$. It is denoted by $\mathcal{O}S$ (resp. $S\mathcal{O}$). Note that if the order $\mathcal{O}$ is maximal, $\mathcal{O}S$ is necessarily a sated left $\mathcal{O}$-ideal.

Let I and J be two $\mathcal{O}_F$-lattices in $\mathcal{A}$. The sum of I and J is defined by $I + J := \{\alpha + \beta \,|\, \alpha \in I, \beta \in J\}$ and their product IJ is the set of all finite sums $\sum_i \alpha_i\beta_i$, where $\alpha_i \in I, \beta_i \in J$. It can be checked that $I + J$ and IJ are still $\mathcal{O}_F$-lattices in $\mathcal{A}$ (the sum of two finitely generated $\mathcal{O}_F$-modules is still a finitely generated $\mathcal{O}_F$-module whose rank is at least the maximum of the ranks of the two modules; for the product see [41, p.260]). An $\mathcal{O}_F$-lattice I is *integral* if $I^2 \subset I$, and a left ideal I is integral if and only if $I \subset \mathcal{O}_\ell(I)$, if and only if $I \subset \mathcal{O}_r(I)$. For an $\mathcal{O}_F$-lattice I of $\mathcal{A}$, the reduced norm of I, denoted by $\mathrm{nrd}(I)$, is the (fractional) ideal of F generated by the set $\{\mathrm{nrd}(\alpha) \,|\, \alpha \in I\}$.The conjugate ideal is the $\mathcal{O}_F$-lattice $\overline{I} := \{\overline{\alpha} \,|\, \alpha \in I\}$.

The quasi-inverse of an $\mathcal{O}_F$-lattice $I \subset \mathcal{A}$ is the set $I^{-1} := \{\alpha \in \mathcal{A} \,|\, I\alpha I \subseteq I\}$, which is, again, an $\mathcal{O}_F$-lattice. Using the definition of the left order of an $\mathcal{O}_F$-lattice, one can check that the above definition is equivalent to $I^{-1} = \{\alpha \in \mathcal{A} \,|\, I\alpha \subseteq \mathcal{O}_\ell(I)\}$ (because for all $x \in \mathcal{A}$, we have $x \in \mathcal{O}_\ell(I)$ if and only if $xI \subseteq I$). By definition, we always have $II^{-1} \subseteq \mathcal{O}_\ell(I)$ and $I^{-1}I \subseteq \mathcal{O}_r(I)$. We say that I is invertible when the previous inclusions are in fact equalities. A left fractional $\mathcal{O}$-ideal I is invertible if it is invertible as an $\mathcal{O}_F$-lattice and if it is sated as a left $\mathcal{O}$-ideal (i.e., $\mathcal{O}_\ell(I) = \mathcal{O}$). The following lemma gives a sufficient condition for an $\mathcal{O}_F$-lattice to be invertible and an expression of its inverse.

Lemma 2.11 ([41, Prop. 16.6.15 (b)]). *Let $I \subseteq \mathcal{A}$ be an $\mathcal{O}_F$-lattice. Whenever either $\mathcal{O}_\ell(I)$ or $\mathcal{O}_r(I)$ is maximal, then both of them are maximal, and futhermore I is invertible.*

The following lemma characterizes the inverse of an invertible $\mathcal{O}_F$-lattice and the inclusion of invertible lattices (see Appendix A for proofs).

Proposition 2.12. *Let I, I' be $\mathcal{O}_F$-lattices I such that $\mathcal{O}_r(I) = \mathcal{O}_\ell(I')$ and I is invertible. Then, one has $II' = \mathcal{O}_\ell(I)$ if and only if $I' = I^{-1}$.*

Lemma 2.13. *Let $I, J \subseteq \mathcal{A}$ be $\mathcal{O}_F$-lattices with the same left order, i.e., $\mathcal{O}_\ell(I) = \mathcal{O}_\ell(J)$. If $J \subseteq I$, then the inclusion of quasi-inverses $I^{-1} \subseteq J^{-1}$ holds.*

Recall that for a given order $\mathcal{O}$, a sated fractional left $\mathcal{O}$-ideal I is an $\mathcal{O}_F$-lattice with $\mathcal{O}_\ell(I) = \mathcal{O}$. The quaternionic ideals we will consider in this work are all sated. When $\mathcal{O}$ is a maximal order, such sated left $\mathcal{O}$-ideals enjoy many nice properties akin to those of fractional ideals in number fields. A first property is that I is always invertible by Lemma 2.11. Additionally, invertible $\mathcal{O}_F$-lattices in $\mathcal{A} = (\frac{a,-1}{F})$ are locally principal ([41, Thm. 16.6.1]). A precise definition of this notion is not needed for the core of this work; rather, it is enough to know that such lattices have nice properties for the reduced norm, that we recall below — see Appendix A for the proof of the last one. We say that I is compatible with J if $\mathcal{O}_r(I) = \mathcal{O}_\ell(J)$.

Lemma 2.14 ([41, Le. 16.3.7, 16.3.5 and 16.3.8]). *Let I, J be two $\mathcal{O}_F$-lattices in $\mathcal{A}$, with I invertible.*

1. *If I is compatible with J, then $\mathrm{nrd}(IJ) = \mathrm{nrd}(I)\mathrm{nrd}(J)$.*
2. *We have $I = \mathcal{O}_\ell(I)\alpha$ if and only if $\alpha \in I$ and $\mathrm{nrd}(\alpha)\mathcal{O}_F = \mathrm{nrd}(I)$.*

Lemma 2.15 ([41, Le. 16.5.11], adapted). *Let I, J be two $\mathcal{O}_F$-lattices in $\mathcal{A}$, with I invertible. If I is compatible with J, then $\mathcal{O}_r(IJ) = \mathcal{O}_r(J)$.*

Lemma 2.16. *Let I, J be two $\mathcal{O}_F$-lattices in $\mathcal{A}$, with I invertible and $\mathcal{O}_\ell(I) = \mathcal{O}_\ell(J)$ or $\mathcal{O}_r(I) = \mathcal{O}_r(J)$. If $I \subset J$ and $\mathrm{nrd}(I) = \mathrm{nrd}(J)$, then $I = J$.*

Assume now that $\mathcal{O}$ is maximal. Then any left-$\mathcal{O}$-ideal is sated. This implies in particular that if I and J are two sated left $\mathcal{O}$-ideals, then their sum is still a sated left $\mathcal{O}$-ideal. We also have the following proposition, which gives us a description of the quasi-inverse[12] of a sum of sated left $\mathcal{O}$-ideals, and proved in Appendix A.

Proposition 2.17. *Let n be a positive integer, $\mathcal{O}$ be a maximal order in $\mathcal{A}$ and $J_1, \ldots, J_n$ be sated fractional left $\mathcal{O}$-ideals in $\mathcal{A}$. Then, the sum $I = J_1 + \cdots + J_n$ has quasi-inverse*

$$I^{-1} = J_1^{-1} \cap \cdots \cap J_n^{-1}.$$

We conclude this subsection with the *norm reduced-Principal Ideal Problem* in quaternion orders. In the commutative version of this problem, K is typically a cyclotomic number field, and the input consists in (a $\mathbb{Z}$-basis of) a principal ideal $a \cdot \mathcal{O}_K$ and the relative norm $a\overline{a}$ of one of its generator.

[12] The same result holds for sums of invertible ideals of number fields.

Definition 2.18 ($\mathcal{O}$-nrdPIP). *Let $\mathcal{O}$ be an order in $\mathcal{A}$. The $\mathcal{O}$-norm reduced Principal Ideal Problem ($\mathcal{O}$-nrdPIP) is: given as input a right $\mathcal{O}$-ideal I and an element $q \in F$ such that $\mathrm{nrd}(I) = q \cdot \mathcal{O}_F$, to compute, if it exists, an element $g \in I$ with $\mathrm{nrd}(g) = q$.*

The following lemma, proven in Appendix A, justifies that computing $g \in I$ such that $\mathrm{nrd}(g) = q$ guarantees that I is generated by g.

Lemma 2.19. *Let (I, q) be an instance of $\mathcal{O}$-nrdPIP and suppose that $g \in I$ is a solution. Then I is a principal right $\mathcal{O}$-ideal and g is a generator. Moreover, the set of solutions is precisely the set of generators of I with reduced norm q, which is equal to $g \cdot \mathcal{O}^1$.*

2.3 Module-LIP

This section introduces the problem we study, borrowing from [33].

Definition 2.20 (Congruent pseudo-Gram matrices). *Two pseudo-Gram matrices $\mathbf{G} = (G, \{\mathfrak{a}_i\}_{1 \le i \le \ell})$ and $\mathbf{G}' = (G', \{\mathfrak{b}_i\}_{1 \le i \le \ell})$ are said to be congruent if there exists a matrix $U = (u_{i,j})_{1 \le i,j \le \ell} \in \mathbf{GL}_\ell(K)$ such that:*

1. *$G' = U^* G U$.*
2. *$\forall\, i, j \in \{1, \ldots, \ell\},\ u_{i,j} \in \mathfrak{a}_i \mathfrak{b}_j^{-1}$.*
3. *$\prod_i \mathfrak{a}_i = (\det U) \prod_i \mathfrak{b}_j$.*

Such a matrix U is called a congruence matrix between $\mathbf{G}$ and $\mathbf{G}'$. The set of congruence matrices between $\mathbf{G}$ and $\mathbf{G}'$ is denoted by $\mathrm{Cong}(\mathbf{G}, \mathbf{G}')$.

Given two congruent pseudo-Gram matrices $\mathbf{G}$ and $\mathbf{G}'$, module-LIP is the task of computing the set $\mathrm{Cong}(\mathbf{G}, \mathbf{G}')$. Following the definition of [33], an instance takes as an additional parameter a pseudo-basis $\mathbf{B}$ of a module $M \subset K^\ell$, whose pseudo-Gram matrix is $\mathbf{G}$.

Definition 2.21 (wc-smodLIP$_K^{\mathbf{B}}$ [33, Definition 3.11]). *Let $\mathbf{B}$ be a pseudo-basis of a module $M \subset K^\ell$, and $\mathbf{G}$ the pseudo-Gram matrix associated to $\mathbf{B}$. Let $\mathbf{G}'$ be a pseudo-Gram matrix congruent to $\mathbf{G}$. The worst-case search module Lattice Isomorphism Problem with parameters K and $\mathbf{B}$ (wc-smodLIP$_K^{\mathbf{B}}$) and input $\mathbf{G}'$, is to compute an element of the set $\mathrm{Cong}(\mathbf{G}, \mathbf{G}')$.*

One can interpret module-LIP as the problem of computing factorizations $\mathbf{C} = (C, \{\mathfrak{b}_i\}_i)$ of $\mathbf{G}' = (G', \{\mathfrak{b}_i\}_i)$ (that is $C^* C = G'$) with the constraint that $\mathbf{C}$ is a pseudo-basis of M. In fact the equivalence between LIP and Gram factorization has already been noticed by Szydlo in [39], for rotations of $\mathbb{Z}^n$, in which case the equivalence with SVP also holds.

Lemma 2.22. *Let $\mathbf{B} = (B, \{\mathfrak{a}_i\}_{1 \le i \le \ell})$ be a pseudo-basis of a rank-ℓ module $M \subseteq K^\ell$ and with associated pseudo-Gram matrix $\mathbf{G}$. Let $\mathbf{G}' = (G', \{\mathfrak{b}_i\}_{1 \le i \le \ell})$ be a pseudo-Gram matrix congruent to $\mathbf{G}$. Then a matrix $U \in \mathbf{GL}_\ell(K)$ is in $\mathrm{Cong}(\mathbf{G}, \mathbf{G}')$ if and only if $C = BU$ satisfies $C^* C = G'$ and $\mathbf{C} = (C, \{\mathfrak{b}_i\}_{1 \le i \le \ell})$ is a pseudo-basis of M.*

See Appendix A for the proof. The relation $G' = U^*GU$ implies that $\det U$ is a solution to the norm equation $\mathrm{nrd}(x) = N_{K/F}(x) = \det G'/\det G$ in K. These equations has been studied and solved in [26], having a large number of solutions in general. The following technical lemma will be useful for the reduction and is proved in Appendix A. It tells us that all $U \in \mathrm{Cong}(\mathbf{G}, \mathbf{G}')$ have the same determinant in $K^\times/\mu(K)$. Moreover when K is a CM field, a representative of this class can be computed efficiently.

Lemma 2.23 (Computing the determinant). *Let $\mathbf{G} = (G, \{\mathfrak{a}_i\}_i)$ and $\mathbf{G}' = (G', \{\mathfrak{b}_i\}_i)$ be two congruent pseudo-Gram matrices. Congruence matrices between $\mathbf{G}$ and $\mathbf{G}'$ all have the same determinant, up to a root of unity of K. We write $\bar{\delta}(\mathbf{G}, \mathbf{G}') \in K^\times/\mu(K)$ for the equivalence class of all these determinants modulo the roots of unity of K.*

Moreover, if K is a CM field, then there is a polynomial time algorithm `ComputeDet` *that given $\mathbf{G}$ and $\mathbf{G}'$ (and a basis of $\mathcal{O}_K$), computes a representative in $K^\times$ of $\bar{\delta}(\mathbf{G}, \mathbf{G}')$.*

2.4 Representations of objects

There are standard ways to represent algebraic numbers and generally, modules over number rings through their pseudo-bases, coming together with notions of sizes. This material is common to other similar works. In particular, we borrow many tools from [33, Section 2.3]. We nonetheless present it in Appendix A.2, for the sake of completeness.

3 A reduction from modLIP to nrdPIP

Let K/F be a CM extension of number fields where $K = F(\sqrt{a})$ and $\mathcal{A}$ denotes the totally definite quaternion algebra $\mathcal{A} = (\frac{a,-1}{F})$ over F. Through this section we fix a maximal order $\mathcal{O}$ in $\mathcal{A}$ containing the order $\mathcal{O}_0 = \mathcal{O}_K \oplus \mathcal{O}_K \cdot j$. In this section, unless explicitly stated otherwise, the modules in K^2 shall be all understood with their module structure taken over $\mathcal{O}_K$.

In this section we prove the main result of this paper, namely, a polynomial time reduction from module-LIP for rank-2 modules in K^2 to nrdPIP, the problem of computing a generator of a (right) principal ideal in $\mathcal{A}$, with given reduced norm (see Definition 2.18). Thanks to Lemma 2.22, module-LIP can be reinterpreted as the task of computing the factorizations of a pseudo-Gram matrix which are also pseudo-bases of a fixed module $M \subset K^2$.

The key point is the isomorphism $\mathcal{A} = K \oplus K \cdot j \simeq K^2$ of K-vector spaces. As a consequence, to a matrix $C \in M_2(K)$ corresponds a unique pair $(\alpha, \beta) \in \mathcal{A}^2$ (applying the previous isomorphism on each column of C). We prove in Lemma 3.5 that when $C^*C = G'$ holds, then the quotient $\alpha\beta^{-1}$ can be obtained from an elementary computation involving only the entries of G' and $\det(C)$. In the setting of module-LIP, this determinant can be computed in polynomial time using Lemma 2.23, *up to a root of unity of K*.

While we are not able to recover α directly from $\alpha\beta^{-1}$, we next use it to build a principal ideal generated by α. The isomorphism $K^2 \simeq \mathcal{A}$ also associates a module $M \subset K^2$ to a $\mathcal{O}$-ideal in $\mathcal{A}$: for example, the left ideal of $\mathcal{A}$ generated by all the (images of) the vectors of M. This ideal, denoted by I_M, is efficiently computable from any pseudo-basis of M (Lemma 3.2). In Proposition 3.6, we then use the knowledge of $\alpha\beta^{-1}$ and I_M to build a principal right $\mathcal{O}'$-ideal $\alpha \cdot \mathcal{O}'$, where $\mathcal{O}'$ is some maximal order in $\mathcal{A}$, efficiently computable from I_M but different from $\mathcal{O}$ in general.

Now $\alpha \cdot \mathcal{O}'$ is known, and $\mathrm{nrd}(\alpha)$ is an entry of G': this is an instance of $\mathcal{O}'$-nrdPIP, which we handle with an oracle. Once such a generator has been computed, the other generators with the same reduced norm are its (right) multiples by the elements of $\mathcal{O}'^1$, as shown in Lemma 2.19). This leads us to the factorizations of $\mathbf{G}'$ (C with $C^*C = G'$), or equivalently to the corresponding congruence matrices, see Algorithm 1. We are however working with a determinant that may not be the correct one, since we know it up to a root of unity. Thanks to the cyclicity of $\mu(K)$, we are able to cover all the possible cases in two steps, and to recover the whole class $\mathrm{Cong}(\mathbf{G}, \mathbf{G}')$, see Algorithm 2.

Most of the objects used in the reduction depend only on the parameters of module-LIP and not on its input. Following standard practices, we assume that we are given $\mathbb{Z}$-bases of $\mathcal{O}_F$ and $\mathcal{O}_K$ and pseudo-bases of $\mathcal{O}$, I_M and $\mathcal{O}'$ (as $\mathcal{O}_F$-modules, see the previous section). We will also assume that the finite group $\mathcal{O}'^1$ has been precomputed[13]. In our situation such a group belongs to an explicit list of finite groups ([41, Chap. 32] and see also Appendix B.3).

3.1 Preliminary results

We first start by proving a few auxiliary but useful lemmata.

Embedding modules. Let us recall the setting for an instance of (rank-two) module-LIP. We are given a pseudo-basis $\mathbf{B} = (B, \mathfrak{a}_1, \mathfrak{a}_2)$ of a rank-two module $M \subset K^2$, with associated pseudo-Gram matrix $\mathbf{G}$. Then, wc-smodLIP$_K^{\mathbf{B}}$ takes as input a pseudo-Gram matrix $\mathbf{G}'$ and asks to compute the set $\mathrm{Cong}(\mathbf{G}, \mathbf{G}')$. Behind the relation $\mathcal{A} = K \oplus K \cdot j$, we have

$$\begin{aligned} \varphi : K^2 &\longrightarrow \mathcal{A} \\ (x, y) &\longmapsto x + yj \end{aligned}$$

an isomorphism of K-vectors spaces, where K acts both on K^2 and $\mathcal{A}$ by left multiplication. Recall that $\mathcal{O}$ is a maximal order of $\mathcal{A}$, fixed once and for all, and containing $\mathcal{O}_0 = \mathcal{O}_K \oplus \mathcal{O}_K \cdot j$.

Definition 3.1. *Let $M \subset K^2$ be a module. We define $I_M = \mathcal{O} \cdot \varphi(M)$, that is, the left $\mathcal{O}$-ideal generated by $\varphi(M)$.*

[13] By computation we mean an abstract finite presentation of a group G, together with an isomorphism $G \simeq \mathcal{O}^1$.

Since $dM \subset \mathcal{O}_K^2$ for some $d \in \mathbb{Z}$, the set $\varphi(M)$ is included in the $\mathcal{O}_0$-lattice $\frac{1}{d}\varphi(\mathcal{O}_K^2)$, and I_M is indeed a left $\mathcal{O}$-ideal. The content of the following lemma gives a convenient description of I_M when a pseudo-basis is known.

Lemma 3.2. *Let $\mathbf{B} = ((b_1 \mid b_2), \mathfrak{a}_1, \mathfrak{a}_2)$ be a pseudo-basis of a module $M \subset K^2$. Let $\alpha = \varphi(b_1)$ and $\beta = \varphi(b_2)$. Then we have*

$$I_M = \mathcal{O}\mathfrak{a}_1\alpha + \mathcal{O}\mathfrak{a}_2\beta.$$

Proof. Let $\{a_1, a_2\} \subset K$ be a two-elements generating set for $\mathfrak{a}_1$. Then $\mathfrak{a}_1\alpha = a_1\mathcal{O}_K \cdot \alpha + a_2\mathcal{O}_K \cdot \alpha$ is contained in a $\mathcal{O}_F$-lattice of $\mathcal{A}$.The same argument holds for $\mathfrak{a}_2\beta$. Since φ is left K-linear and $M = \mathfrak{a}_1 b_1 + \mathfrak{a}_2 b_2$ we have $\varphi(M) = \mathfrak{a}_1\alpha + \mathfrak{a}_2\beta$, so $I_M \subset \mathcal{O}\mathfrak{a}_1\alpha + \mathcal{O}\mathfrak{a}_2\beta$. Conversely, $\varphi(M)$ contains the rank one submodule $\mathfrak{a}_1\alpha$ so it holds that $\mathcal{O}\mathfrak{a}_1\alpha \subset I_M$, and in the same way $\mathcal{O}\mathfrak{a}_2\beta \subset I_M$. As I_M is stable by addition, it must contain the sum $\mathcal{O}\mathfrak{a}_1\alpha + \mathcal{O}\mathfrak{a}_2\beta$. □

By construction, the left order of I_M is $\mathcal{O}$. Its right order $\mathcal{O}' := \mathcal{O}_r(I_M)$ is a priori different from $\mathcal{O}$, except for special cases such as when $M = \mathcal{O}_K^2$. In this case, the previous lemma applied with the (pseudo)-basis $\mathbf{Id} = (\mathrm{Id}, \mathcal{O}_K, \mathcal{O}_K)$ of $\mathcal{O}_K^2$ immediately gives $I_M = \mathcal{O}$, thus $\mathcal{O}' = \mathcal{O}$ holds. This fact is stated in the following corollary, which will be useful to state a simplified version of our reduction when $M = \mathcal{O}_K^2$.

Corollary 3.3. *For $M = \mathcal{O}_K^2$, we have $I_M = \mathcal{O}$.*

Remark 3.4. By referring to the discussion at the beginning of [37, Chapter 24], the identity $I_M = \mathcal{O}$ holds whenever I_M is integral and $\mathrm{nrd}(I_M) = \mathcal{O}_F$ is verified.[14] Note that $M \subseteq \mathcal{O}_K^2$ implies $\varphi(M) \subseteq \mathcal{O}_0$ and thus $I_M \subseteq \mathcal{O}$; in other words, I_M is integral whenever $M \subseteq \mathcal{O}_K^2$.

Gram matrices and quaternions. We can identify $\mathcal{M}_2(K)$ with $K^2 \times K^2$ (taking the column vectors) and thus with $\mathcal{A}^2$, applying φ columnwise. Therefore to a matrix $C \in \mathcal{M}_2(K)$ corresponds a unique pair $(\alpha, \beta) \in \mathcal{A}^2$. In the following lemma, we prove that if C is a factorization of G', then the quotient $\alpha\beta^{-1}$ is expressible in terms of the coefficients of G' and $\det(C)$.

Lemma 3.5. *Let $C = \begin{pmatrix} x_1 & x_2 \\ y_1 & y_2 \end{pmatrix}, G' = \begin{pmatrix} q_{1,1} & q_{1,2} \\ \overline{q_{1,2}} & q_{2,2} \end{pmatrix} \in \mathbf{GL}_2(K)$ and let α, $\beta \in \mathcal{A}$ be the quaternions defined by $\alpha = \varphi(x_1, y_1)$ and $\beta = \varphi(x_2, y_2)$. Then we have the following equivalence*

$$C^*C = G' \iff \begin{cases} \mathrm{nrd}(\alpha) = q_{1,1} \\ \mathrm{nrd}(\beta) = q_{2,2} \\ \alpha\beta^{-1} = q_{2,2}^{-1}(\overline{q_{1,2}} - \det(C)j) \end{cases} \tag{1}$$

[14] Note that in [37], the norm map $N_{\mathcal{A}/F}$ is defined exclusively for normal ideals, i.e., ideals whose left and right orders are maximal. Its definition is different from the one of nrd we gave in Section 2. However Theorem 24.11 and Corollary 24.12 of [37] ensure that the identity $N_{\mathcal{A}/K}(I) = \mathrm{nrd}(I)^2$ holds for such ideals. In particular we have $N_{\mathcal{A}/K}(I) = \mathcal{O}_F$ whenever $\mathrm{nrd}(I) = \mathcal{O}_F$.

Proof. Let us write $c_1 = \binom{x_1}{y_1}$ and $c_2 = \binom{x_2}{y_2}$ for the columns of C, so that $C^*C = \left(\begin{smallmatrix} c_1^*\cdot c_1 & c_1^*\cdot c_2 \\ c_2^*\cdot c_1 & c_2^*\cdot c_2 \end{smallmatrix}\right)$. The first coefficient is $c_1^* \cdot c_1 = \overline{x_1}x_1 + \overline{y_1}y_1 = \mathrm{nrd}(\alpha)$ and in the same way, the last coefficient is $c_2^* \cdot c_2 = \mathrm{nrd}(\beta)$. For the non diagonal terms, we first compute

$$\begin{aligned}\alpha\overline{\beta} &= (x_1 + y_1 j)(\overline{x_2} - y_2 j) = \overline{x_2}x_1 + \overline{y_2}y_1 + (y_1x_2 - x_1y_2)j \\ &= c_2^* \cdot c_1 - \det(C)j,\end{aligned}$$

where we used the relations $\overline{xj} = -xj$ and $jx = \overline{x}j$ which hold for any $x \in K$. Combining with $\beta^{-1} = \overline{\beta} \cdot \mathrm{nrd}(\beta)^{-1}$, we obtain $\alpha\beta^{-1} = \mathrm{nrd}(\beta)^{-1} \cdot (c_2^* \cdot c_1 - \det(C)j)$. This gives the result. □

Next we show how to recover a principal ideal generated by α, from $\alpha\beta^{-1}$ and I_M.

Proposition 3.6. *Let $\mathbf{C} = ((c_1 \mid c_2), \mathfrak{b}_1, \mathfrak{b}_2)$ be a pseudo-basis of a module $M \subset K^2$. Let $\alpha = \varphi(c_1)$, $\beta = \varphi(c_2)$ and $\mathcal{O}' = \mathcal{O}_r(I_M)$. Then $\mathcal{O}' = I_M^{-1}I_M$ is a maximal order, and we have the following equality of right $\mathcal{O}'$-ideals*

$$\alpha\mathcal{O}' = \mathfrak{b}_1^{-1}I_M \cap \alpha\beta^{-1}\mathfrak{b}_2^{-1}I_M.$$

Proof. Recall that $\mathcal{O}_\ell(I_M) = \mathcal{O}$ is a maximal order. Then, Lemma 2.11 tells us that I_M is invertible and that $\mathcal{O}' = \mathcal{O}_r(I_M)$ is maximal too. The same argument applies to $I := \mathcal{O}\mathfrak{b}_1\alpha$ and we will show that its inverse is $I^{-1} = \alpha^{-1}\mathfrak{b}_1^{-1}\mathcal{O}$. From Proposition 2.12, it is enough to prove the equality $I(\alpha^{-1}\mathfrak{b}_1^{-1}\mathcal{O}) = \mathcal{O}_\ell(I)(= \mathcal{O})$. The inclusion $I(\alpha^{-1}\mathfrak{b}_1^{-1}\mathcal{O}) \subseteq \mathcal{O}_\ell(I)$ is clear from the definition of the product of two ideals, and because $\mathfrak{b}_1\mathfrak{b}_1^{-1} = \mathcal{O}_K$ is contained in $\mathcal{O}_\ell(I) = \mathcal{O}$. Conversely, since $\mathfrak{b}_1\mathfrak{b}_1^{-1} = \mathcal{O}_K$, there exists elements $a_1, \dots, a_k \in \mathfrak{b}_1$ and $a'_1, \dots, a'_k \in \mathfrak{b}_1^{-1}$ such that $\sum_k a_k a'_k = 1$. Since $1 \in \mathcal{O}$, we have $a_k\alpha \in I$ for all k. By definition, we also have $\alpha^{-1}a'_k x \in \alpha^{-1}\mathfrak{b}_1^{-1}\mathcal{O}$ for all $x \in \mathcal{O}$. This means that $x = \sum_k (a_k\alpha)(\alpha^{-1}a'_k x) \in I(\alpha^{-1}\mathfrak{b}_1^{-1}\mathcal{O})$, and proves the other inclusion.

Similarly we have $(\mathcal{O}\mathfrak{b}_2\beta)^{-1} = \beta^{-1}\mathfrak{b}_2^{-1}\mathcal{O}$. Using Proposition 2.17 and the definition of I_M from Lemma 3.2 yields $I_M^{-1} = (\mathcal{O}\mathfrak{b}_1\alpha)^{-1} \cap (\mathcal{O}\mathfrak{b}_2\beta)^{-1}$. Multiplying this equality by α on the left and by I_M on the right (the product of ideals is compatible), we obtain the result. □

3.2 The reduction

Now we have everything we need to prove the main result of this paper. For readability, we have cut the reduction algorithm into two smaller algorithms. A first algorithm (Algorithm 1 below) computes all matrices in $\mathrm{Cong}(\mathbf{G}, \mathbf{G}')$ with prescribed determinant, from one call to a nrdPIP oracle.[15] This algorithm contains the core techniques of the reduction. We then provide the main reduction algorithm in Algorithm 2, which computes all congruence matrices between $\mathbf{G}$

[15] This set can be empty, for example if the prescribed determinant δ for C is not a solution to the equation $\delta\overline{\delta} = \det(G')$.

and $\mathbf{G}'$, without restriction on the determinant. This algorithm calls Algorithm 1 twice. Having put the fixed determinant computation in a separated algorithm will also be useful in Section 4, where we will improve the reduction for the special module-LIP instances appearing in Hawk (still using Algorithm 1).

Algorithm 1: Computing congruence matrices of fixed determinant

Input: • $\mathbf{B} = (B, \mathfrak{a}_1, \mathfrak{a}_2)$ a pseudo-basis of a rank-2 module $M \subset K^2$, of associated pseudo-Gram matrix $\mathbf{G}$;
• $\mathbf{G}' = (G' = (q_{i,j})_{1 \leq i,j \leq 2}, \mathfrak{b}_1, \mathfrak{b}_2)$ pseudo-Gram matrix congruent to $\mathbf{G}$;
• pseudo-bases of $\mathcal{O}$, I_M and $\mathcal{O}' = \mathcal{O}_r(I_M)$ over $\mathcal{O}_F$;
• the (finite) set $\mathcal{O}'^1$;
• an oracle $\mathfrak{O}$ solving $\mathcal{O}'$-nrdPIP (outputting $\perp$ when there is no solution);
• a prescribed determinant $\delta \in K$

Output: The set of all congruence matrices in $\mathrm{Cong}(\mathbf{G}, \mathbf{G}')$ with determinant δ

$\texttt{Congruence\_mat}_\delta \leftarrow \{\}$
$\gamma \leftarrow \delta \cdot \det B$
if $\gamma \cdot \overline{\gamma} \neq \det(G')$ **then**
  Return $\{\}$ `// there are no solutions with this determinant`
$q \leftarrow q_{2,2}^{-1}(q_{2,1} - \gamma j) \in \mathcal{A}$ `// c.f., Lemma 3.5`
$I \leftarrow \mathfrak{b}_1^{-1} I_M \cap q\mathfrak{b}_2^{-1} I_M$ `// c.f., Proposition 3.6`
$\alpha' \leftarrow \mathfrak{O}(I, q_{1,1})$
if $\alpha' = \perp$ **then**
  Return $\{\}$ `// the nrd-PIP instance was invalid`
$S \leftarrow \{\alpha' \cdot x \,|\, x \in \mathcal{O}'^1\}$ `// set of all solutions to the nrdPIP instance`
for α *in* S **do**
  $\beta \leftarrow q^{-1}\alpha$
  $C \leftarrow (\varphi^{-1}(\alpha) \,|\, \varphi^{-1}(\beta))$
  if $\mathbf{C} = (C, \mathfrak{b}_1, \mathfrak{b}_2)$ *is a pseudo-basis for* M **then**
    $U \leftarrow B^{-1} \cdot C$
    $\texttt{Congruence\_mat}_\delta \leftarrow \texttt{Congruence\_mat}_\delta \cup \{U\}$
Return $\texttt{Congruence\_mat}_\delta$.

Theorem 3.7. *Let* $\mathbf{B} = (B, \mathfrak{a}_1, \mathfrak{a}_2)$ *be a pseudo-basis of a rank-2 module* $M \subset K^2$*, with associated pseudo-Gram matrix* $\mathbf{G}$*, let* $\mathbf{G}' = (G', \mathfrak{b}_1, \mathfrak{b}_2)$ *be a pseudo-Gram matrix congruent to* $\mathbf{G}$*, and let* $\delta \in K$ *be a prescribed determinant. Assume that pseudo-bases over* $\mathcal{O}_F$ *of* $\mathcal{O}$*,* I_M *and* $\mathcal{O}' = \mathcal{O}_r(I_M)$ *have been precomputed, as well as the finite group* $\mathcal{O}'^1$*. Finally, assume that we are given an oracle* $\mathfrak{O}$ *that solves* $\mathcal{O}'$*-nrdPIP. Then Algorithm 1 returns the (potentially empty) set of congruence matrices between* $\mathbf{G}$ *and* $\mathbf{G}'$ *with determinant* δ*. It makes exactly one*

call to the oracle $\mathfrak{O}$ and except for this call it runs in time

$$\mathrm{poly}(\log \Delta_K, \mathrm{size}(\mathbf{G}), \mathrm{size}(\mathbf{G}')).$$

Proof. Correctness. We prove that the algorithm outputs the set of all congruence matrices between $\mathbf{G}$ and $\mathbf{G}'$ with determinant δ. If the oracle did not fail at Step 7, then the elements α in S indeed verify $\mathrm{nrd}(\alpha) = q_{1,1}$. By definition of β, and using the fact that $\gamma \cdot \overline{\gamma} = \det(G')$ (otherwise the algorithm outputs an empty list), we can also check that $\mathrm{nrd}(\beta) = q_{2,2}$. Finally let $C = \left(\begin{smallmatrix} x_1 & x_2 \\ y_1 & y_2 \end{smallmatrix}\right)$ be a matrix computed during Step 13 and satisfying Step 14, where $\alpha = x_1 + y_1 j \in S$ and $\beta = x_2 + y_2 j$. We have to prove that $\det(C) = \gamma$, so that C will satisfy all the conditions in Lemma 3.5. First notice that $\alpha\overline{\beta} = x_1\overline{x_2} + y_1\overline{y_2} - (x_1 y_2 - y_1 x_2)j$. On one hand, $x_1\overline{x_2} + y_1\overline{y_2} = \overline{q_{1,2}}$ and $(x_1 y_2 - y_1 x_2)j = \det(C)j$, so $\alpha\overline{\beta} - (x_1\overline{x_2} + y_1\overline{y_2}) = \alpha\overline{\beta} - \overline{q_{1,2}} = -\det(C)j$. On the other hand, $\alpha\beta^{-1} = q_{2,2}^{-1}(\overline{q_{1,2}} - \gamma j)$ by Step 5, so multiplying by $\mathrm{nrd}(\beta) = q_{2,2}$, we get $\alpha\overline{\beta} - \overline{q_{1,2}} = -\gamma j$ (recall that $\beta^{-1}\mathrm{nrd}(\beta) = \overline{\beta}$). By identification in $\mathcal{A}$, we get $\det(C) = \gamma$. Hence, the coefficients of C verify all the conditions of the right-hand-side of Equation (1) in Lemma 3.5 and so the lemma implies that we must have $C^*C = G'$. Since $\mathbf{C}$ is a pseudo-basis for M, the corresponding $U = B^{-1} \cdot C$ computed at Step 15 is a pseudo-basis change between $\mathbf{B}$ and $\mathbf{C}$ and it satisfies $U^*GU = G'$, *i.e.*, $U \in \mathrm{Cong}(\mathbf{G}, \mathbf{G}')$.

Conversely, let $U_0 \in \mathrm{Cong}(\mathbf{G}, \mathbf{G}')$ with $\det(U_0) = \delta$ and let us prove that $U_0 \in \texttt{Congruence\_mat}_\delta$ by the end of the algorithm. Then, $C_0 = B \cdot U_0$ is a pseudo-basis of M with $\det(C_0) = \gamma$ and such that $C_0^* C_0 = G'$ (which implies in particular that $\gamma \cdot \overline{\gamma} = \det(G')$, so the algorithm does not terminate at Step 4). Let c_1 and c_2 be the columns of C_0, and $\alpha_0 = \varphi(c_1)$ and $\beta_0 = \varphi(c_2)$. The element q computed in Step 5 of the algorithm satisfies $q = q_{2,2}^{-1}(q_{2,1} - \gamma j) = q_{2,2}^{-1}(\overline{q_{1,2}} - \det(C)j) = \alpha_0 \cdot \beta_0^{-1}$ using Lemma 3.5. And Proposition 3.6 then implies that the ideal I computed in Step 6 satisfies $I = \alpha_0 \mathcal{O}'$. Moreover, Lemma 3.5 again tells us that $\mathrm{nrd}(\alpha_0) = q_{1,1}$, so α_0 is a solution to the nrdPIP instance $(I, q_{1,1})$. The set S computed during Step 10 of the algorithm is then non-empty and we know from Lemma 2.19 that $\alpha_0 \in S$. For this choice of $\alpha = \alpha_0$ in the loop, the matrix C computed at Step 13 is exactly C_0, so it passes Step 14, and $U = U_0$ is added to $\texttt{Congruence\_mat}_\delta$ at Step 16.

Complexity. The computation of γ at Step 2, as well as the computation of q at Step 5 and the computation of I at Step 6 can be done in polynomial time (see subsection A.2 for the computation of I). Note that the size of the right $\mathcal{O}'$-ideal I is polynomially bounded by the size of the inputs of the algorithm. At Step 7, the algorithm makes one call to the oracle $\mathfrak{O}$. Step 10 can be performed in time $\mathrm{poly}(\log \Delta_K, \mathrm{size}(\alpha'), |\mathcal{O}'^1|)$. By Lemma B.25, the size of $|\mathcal{O}'^1|$ is polynomial in the degree of F, so this step can be done in polynomial time too, and the size of S is polynomial in $[F : \mathbb{Q}]$. The for loop starting at Step 11 will then be iterated a polynomial number of times. Each computation from Step 11 to the end, including checking that the candidates $\mathbf{C}$ are indeed pseudo-bases of M, require only simple linear algebra computations, which can be performed in polynomial time. This concludes the bound on the running time of the algorithm. □

Algorithm 2: Reduction of wc-smodLIP to nrdPIP

Input: • $\mathbf{B} = (B, \mathfrak{a}_1, \mathfrak{a}_2)$ a pseudo-basis of a rank-2 module $M \subset K^2$, of associated pseudo-Gram matrix $\mathbf{G}$;
• $\mathbf{G}' = (G', \mathfrak{b}_1, \mathfrak{b}_2)$ pseudo-Gram matrix congruent to $\mathbf{G}$;
• pseudo-bases of $\mathcal{O}$, I_M and $\mathcal{O}' = \mathcal{O}_r(I_M)$ over $\mathcal{O}_F$;
• the (finite) sets $\mu(K) = \langle \mu_0 \rangle$ and $\mathcal{O}'^1$;
• an oracle $\mathfrak{O}$ solving $\mathcal{O}'$-nrdPIP (outputting $\perp$ when there is no solution)

Output: The set of all congruence matrices $\mathrm{Cong}(\mathbf{G}, \mathbf{G}')$

1 `Congruence_mat` $\leftarrow \{\}$
2 $\delta_0 \leftarrow$ `ComputeDet`$(\mathbf{G}, \mathbf{G}')$ // *c.f.,* Lemma 2.23
3 $\mu_0 \leftarrow$ A generator of $\mu(K)$
4 **for** $i \in \{0, 1\}$ **do**
5 $\delta \leftarrow \delta_0 \cdot \mu_0^i$
6 Compute $\texttt{Congruence\_mat}_\delta$ with Algorithm 1 // *c.f.,* Theorem 3.7
7 **for** $U \in \textit{Congruence\_mat}_\delta$ **do**
8 **for** $\mu \in \mu(K)$ **do**
9 $V \leftarrow \mu \cdot U$
10 `Congruence_mat` $\leftarrow$ `Congruence_mat` $\cup \{V\}$
11 Return `Congruence_mat`.

Theorem 3.8. *Let $\mathbf{B} = (B, \mathfrak{a}_1, \mathfrak{a}_2)$ be a pseudo-basis of a rank-2 module $M \subset K^2$, with associated pseudo-Gram matrix $\mathbf{G}$ and let $\mathbf{G}' = (G', \mathfrak{b}_1, \mathfrak{b}_2)$ be a pseudo-Gram matrix congruent to $\mathbf{G}$. Assume that pseudo-bases over $\mathcal{O}_F$ of $\mathcal{O}$, I_M and $\mathcal{O}' = \mathcal{O}_r(I_M)$ have been precomputed, as well as the finite groups $\mu(K)$ and $\mathcal{O}'^1$. Finally, assume that we are given an oracle $\mathfrak{O}$ that solves $\mathcal{O}'$-nrdPIP. Then Algorithm 2 returns the set $\mathrm{Cong}(\mathbf{G}, \mathbf{G}')$ of all congruence matrices between $\mathbf{G}$ and $\mathbf{G}'$. In particular it solves wc-smodLIP$^{\mathbf{B}}_K$ on input $\mathbf{G}'$. Moreover, it makes exactly two calls to the oracle $\mathfrak{O}$ and except for these calls it runs in time*

$$\mathrm{poly}(\log \Delta_K, \mathrm{size}(\mathbf{G}), \mathrm{size}(\mathbf{G}')).$$

Proof. Correctness: We want to prove that at the end of the algorithm, the variable set `Congruence_mat` contains all the congruence matrices between $\mathbf{G}$ and $\mathbf{G}'$, i.e., that `Congruence_mat` $= \mathrm{Cong}(\mathbf{G}, \mathbf{G}')$. Observe first that if $U \in \texttt{Congruence\_mat}_\delta$ is chosen at Step 7, then for all $\mu \in \mu(K)$, the matrix $V = \mu \cdot U$ satisfies the three conditions in Definition 2.20 (because U does) thus $V \in \mathrm{Cong}(\mathbf{G}, \mathbf{G}')$. This proves the inclusion `Congruence_mat` $\subseteq \mathrm{Cong}(\mathbf{G}, \mathbf{G}')$.

Let us now prove the reverse inclusion. Let $V_0 \in \mathrm{Cong}(\mathbf{G}, \mathbf{G}')$ be arbitrary, we want to prove that, by the end of the algorithm, $V_0 \in$ `Congruence_mat` holds. Let $\delta_0 \leftarrow$ `ComputeDet`$(\mathbf{G}, \mathbf{G}')$ be as in Step 2 of the algorithm. By Lemma 2.23, we know that $\det(V_0)$ is equal to δ_0 up to a root of unity of K, i.e., $\det(V_0) = \delta_0 \cdot \mu$ for some $\mu \in \mu(K)$. Since μ_0 generates $\mu(K)$, one can write in a unique way $\mu = \mu_0^i \cdot \mu_0^{2k}$, where $i \in \{0, 1\}$ and $k \in \{0, \ldots, \lfloor (|\mu(K)| + 1)/2 \rfloor\}$. Let us focus on this i-th iteration of the outer for loop. We will prove that V_0 is added to `Congruence_mat`

during this iteration. By the previous observation, $U_0 := \mu_0^{-k} \cdot V_0$ belongs to $\mathrm{Cong}(\mathbf{G}, \mathbf{G}')$ as well. Also by construction, $\det(U_0) = (\mu_0^{-k})^2 \det(V_0) = \delta_0 \mu_0^i = \delta$, where we used the fact that U_0 and V_0 are 2 by 2 matrices, so $\det(xV_0) = x^2 \det(V_0)$ for any $x \in K$. By the correctness of Algorithm 1 (Theorem 3.7), we conclude that U_0 belongs to the set $\texttt{Congruence\_mat}_\delta$ computed in Step 6. During the iteration of the inner loop corresponding to $U_0 \in \texttt{Congruence\_mat}_\delta$, V_0 is then computed at Step 9. This concludes the proof of the correctness.

Complexity: According to Lemma 2.23, a representative δ_0 of the determinant class can be computed in polynomial time. Inside the outer loop (starting at Step 4), the computation of δ at Step 5 can be done in polynomial time. Since it makes two iterations, and by Theorem 3.7, the algorithm makes exactly two calls to the nrdPIP oracle and except for these calls, Step 6 runs in polynomial time. The for loop starting at Step 7 will then be iterated a polynomial number of times (this can be made more precise, see subsection 3.3) and Lemma 2.4 tells us that there is a polynomial number of roots of unity in K, so the number of iterations of the final loop (starting at Step 8) will be polynomially bounded and each computation inside this inner loop can be performed in polynomial time. This concludes the bound on the running time of the algorithm. □

Algorithm 2 requires as input a pseudo-Gram matrix $\mathbf{G}'$ congruent to $\mathbf{G}$ (which will be the input of our module-LIP problem) but also multiple other objects: a pseudo-basis $\mathbf{B}$ of M, a maximal order $\mathcal{O}$ containing $\mathcal{O}_K + \mathcal{O}_K \cdot j$, the ideal I_M, the right order $\mathcal{O}'$ of I_M, the roots of unity $\mu(K)$ of K, and the set $\mathcal{O}'^1$ of elements of reduced norm 1 in $\mathcal{O}'$. An important observation is that all these additional objects only depend on K and $\mathbf{B}$, which are parameters of the module-LIP problem. Hence, for the purpose of reductions, one can assume that all these objects have been pre-computed somehow, and that the reduction algorithm only takes as input $\mathbf{G}'$, the input of module-LIP. This makes the reduction from module-LIP to nrdPIP *non-uniform*: for every choice of parameters K and $\mathbf{B}$ of the module-LIP problem, there exists a reduction from $\text{wc-smodLIP}_K^{\mathbf{B}}$ to nrdPIP, but there might not exist an efficient algorithm computing a description of these reductions (e.g., as a Turing Machine, or an arithmetic circuit) from the knowledge of K and $\mathbf{B}$.

Still, some of the objects from the list above can be computed efficiently. This is the case of $\mu(K)$, which can always be computed from K in polynomial time (see Lemma 2.4). If $\mathcal{O}$ has been computed, then the ideal I_M can also be computed efficiently from $\mathcal{O}$ and $\mathbf{B}$, using Lemma 3.2 and results from Appendix A.2. Once I_M has been computed, its right order $\mathcal{O}'$ can also be computed efficiently, again by Appendix A.2. The only two objects that may require effort to compute are $\mathcal{O}$ and $\mathcal{O}'^1$. If K is a cyclotomic field, then $\mathcal{O}$ becomes efficiently computable using Proposition B.18. If, in addition, we have $M = \mathcal{O}_K^2$, then $\mathcal{O}' = \mathcal{O}$ and the set $\mathcal{O}'^1 = \mathcal{O}^1$ becomes efficiently computable too, using Corollary B.27. Hence, when the field K is cyclotomic and the module M is equal to $\mathcal{O}_K^2$, all the quantities needed as input of Algorithm 2 can be computed in polynomial time from the knowledge of K and $\mathbf{B}$, and so we

obtain a *uniform* reduction from wc-smodLIP to ndrPIP (this will be detailed in Section 4, together with other improvements for this special case).

As a conclusion for this section, we prove the following statement (for general CM fields and general rank-2 modules $M \subset K^2$).

Corollary 3.9 (modLIP to $\mathcal{O}'$-nrdPIP). *There is a non-uniform polynomial time reduction from wc-smodLIP$_K^{\mathbf{B}}$ to $\mathcal{O}'$-nrdPIP, where K is any CM-field (with maximal totally real subfield F), $\mathbf{B}$ is any pseudo-basis of a rank-2 module $M \subseteq K^2$ and $\mathcal{O}'$ is a particular maximal order of a quaternion algebra over F, depending only on K and $\mathbf{B}$. The reduction makes two calls to the $\mathcal{O}'$-nrdPIP oracle.*

Proof. Let K be any CM field with maximal totally real subfield F, and let $a \in F$ totally negative such that $K = F(\sqrt{a})$. Let $\mathbf{B}$ be a pseudo-basis of a rank-2 module $M \subseteq K^2$. Let $\mathcal{A}$ be the quaternion algebra $(\frac{a,-1}{F})$ and $\mathcal{O}$ be a maximal order of $\mathcal{A}$ containing $\mathcal{O}_K + \mathcal{O}_K \cdot j$. Let I_M be the left $\mathcal{O}$-ideal of $\mathcal{A}$ associated to the module M, as in Definition 3.1, and let $\mathcal{O}' = \mathcal{O}_r(I_M)$ (note that $\mathcal{O}'$ is maximal because $\mathcal{O}$ is, using Lemma 2.11). We want to prove that there is a non-uniform polynomial time reduction from wc-smodLIP$_K^{\mathbf{B}}$ to $\mathcal{O}'$-nrdPIP. The reduction is provided by Algorithm 2. This algorithm takes as input a pseudo-Gram matrix $\mathbf{G}'$, which is the input of the wc-smodLIP$_K^{\mathbf{B}}$ problem, as well as many other inputs that only depend on K and $\mathbf{B}$, and solves wc-smodLIP$_K^{\mathbf{B}}$ on input $\mathbf{G}'$ by making two calls to a $\mathcal{O}'$-nrdPIP oracle. Since $\mathbf{B}$, $\mathcal{O}$, I_M, $\mathcal{O}'$, $\mathcal{O}'^1$ and $\mu(K)$ all depend only on K and $\mathbf{B}$, which are parameters of the module-LIP problem, we can assume that these quantities have been hardcoded into the algorithm, instead of being given as input. □

3.3 Application to the number of module lattice automorphisms of rank-2 modules

Let us fix a rank-ℓ module $M \subset K^\ell$. We clarify the link between $\mathrm{Aut}(M)$, the module lattice automorphism group of M and the full set of solutions to an instance of module-LIP. Module lattice automorphisms are represented by matrices, so that $\mathrm{Aut}(M) = \{\Theta \in \mathbf{GL}_\ell(K_\mathbb{R}) \,|\, \Theta \cdot M = M \text{ and } \Theta^*\Theta = \mathrm{Id}\}$. Note that the constraint $\Theta \cdot M = M$, together with $M \subset K^\ell$ (of rank ℓ), implies that Θ has coefficients in K. Se we can equivalently define $\mathrm{Aut}(M) = \{\Theta \in \mathbf{GL}_\ell(K) \,|\, \Theta \cdot M = M \text{ and } \Theta^*\Theta = \mathrm{Id}\}$.

Proposition 3.10. *Let $\mathbf{C} = (C, \{\mathfrak{b}_i\}_{1\le i\le \ell})$ be a pseudo-basis of a rank-ℓ module $M \subset K^\ell$ and let $G = C^*C$. We have*

$$\mathrm{Aut}(M) = \{C'C^{-1} \,|\, \mathbf{C}' = (C', \{\mathfrak{b}_i\}_{1\le i\le \ell}) \text{ is a pseudo-basis of } M \text{ and } C'^*C' = G\}.$$

Proof. If $\Theta \in \mathrm{Aut}(M)$, then $\Theta = C'C^{-1}$ for $C' = \Theta C$ which, with the coefficient ideals $\mathfrak{b}_i$, forms a pseudo-basis of M having the same Gram matrix G. Conversely, let C' be as in the right set. Then, $\Theta = C'C^{-1}$ is a K-endomorphism of K^ℓ such that $\Theta^*\Theta = (C^{-1})^*(C'^*C')C^{-1} = (C^{-1})^*(C^*C)C^{-1} = \mathrm{Id}$. Moreover, $\Theta \cdot M =$

$C'C^{-1} \cdot (C_1\mathfrak{b}_1 \oplus \cdots \oplus C_\ell\mathfrak{b}_\ell) = C'_1\mathfrak{b}_1 \oplus \cdots \oplus C'_\ell\mathfrak{b}_\ell = M$, where C_i and C'_i denote the column vectors of C and C' respectively. Hence we have proved $\Theta \in \mathrm{Aut}(M)$ and the result. □

Corollary 3.11. *Let $\mathbf{B} = (B, \{\mathfrak{a}_i\}_{1\le i\le \ell})$ be a pseudo-basis of a module $M \subset K^\ell$, with pseudo-Gram matrix $\mathbf{G}$. Consider an instance $\mathbf{G}' = (G', \{\mathfrak{b}_i\}_{1\le i\le \ell})$ of wc-smodLIP$^{\mathbf{B}}_K$. For $U_0 \in \mathrm{Cong}(\mathbf{G}, \mathbf{G}')$ arbitrary, we have*

$$\begin{aligned} \mathrm{Aut}(M) &\longrightarrow \mathrm{Cong}(\mathbf{G}, \mathbf{G}') \\ \Theta &\longmapsto (B^{-1}\Theta B) \cdot U_0 \end{aligned}$$

is a bijection. In particular, $|\mathrm{Aut}(M)| = |\mathrm{Cong}(\mathbf{G}, \mathbf{G}')|$ and from the knowledge of $\mathrm{Cong}(\mathbf{G}, \mathbf{G}')$, one can efficiently recover $\mathrm{Aut}(M)$.

Proof. We have the following sequence of equivalences

$$\begin{aligned} & U \in \mathrm{Cong}(\mathbf{G}, \mathbf{G}') \\ \iff & \mathbf{C} = (C = BU, \{\mathfrak{b}_i\}_{1\le i\le \ell}) \text{ is a pseudo-basis of } M \text{ with } C^*C = G'. \\ \iff & BU \cdot (BU_0)^{-1} \in \mathrm{Aut}(M) \iff U \in (B^{-1} \cdot \mathrm{Aut}(M) \cdot B) \cdot U_0, \end{aligned}$$

where the first equivalence comes from Lemma 2.22 and the second is a direct consequence of Proposition 3.10.

Given $\mathrm{Cong}(\mathbf{G}, \mathbf{G}')$ one can efficiently recover $\mathrm{Aut}(M)$ by fixing any $U_0 \in \mathrm{Cong}(\mathbf{G}, \mathbf{G}')$ and computing the set $\{B \cdot U \cdot U_0^{-1} \cdot B^{-1} \,|\, U \in \mathrm{Cong}(\mathbf{G}, \mathbf{G}')\}$ by linear algebra. Note that the reverse computation (computing $\mathrm{Cong}(\mathbf{G}, \mathbf{G}')$ from $\mathrm{Aut}(M)$) is a priori not easy to perform without knowing at least one element of $\mathrm{Cong}(\mathbf{G}, \mathbf{G}')$. □

Analyzing carefully Algorithms 1 and 2, we are able to give a bound on the number of solutions to a module-LIP instance, when $M \subset K^2$. In light of Corollary 3.11, this also bounds the cardinality of $\mathrm{Aut}(M)$ for such modules.

Theorem 3.12. *Let K be a CM field of degree $d > 4$ and let $M \subset K^2$ be a rank-two module. We have $|\mathrm{Aut}(M)| \le 64d^4$.*

Proof. Let $\mathbf{B}$ be any pseudo-basis of M, with associated pseudo-Gram matrix $\mathbf{G}$, and let $\mathbf{G}'$ be any instance of wc-smodLIP$^{\mathbf{B}}_K$ (for example, $\mathbf{G}' = \mathbf{G}$). By Corollary 3.11, it is enough to bound the cardinality of $\mathrm{Cong}(\mathbf{G}, \mathbf{G}')$. Looking at Algorithm 1, one observes that its output has size less than or equal to $|S| = |\mathcal{O}'^1|$. In the same way, at the end of Algorithm 2, we have $|\mathrm{Cong}(\mathbf{G}, \mathbf{G}')| = |\texttt{Congruence\_mat}| \le 2|\mathcal{O}'^1| \cdot |\mu(K)|$ (thanks to Theorem 3.8). But Proposition B.25 gives $|\mathcal{O}'^1| \le 16d^2$ and Lemma 2.4 tells us that $|\mu(K)| \le 2d^2$, hence we obtain $|\mathrm{Cong}(\mathbf{G}, \mathbf{G}')| \le 64d^4$. □

The group $\mathrm{Aut}(\mathcal{O}_K^2) := \{\Theta \in \mathbf{GL}_2(\mathcal{O}_K) \,|\, \Theta^*\Theta = \mathrm{Id}\}$ can be understood more precisely — see also Appendix B.3.

Proposition 3.13. *The group* $\mathrm{Aut}(\mathcal{O}_K^2)$ *is finite of order* $2\,|\mu(K)|^2 \leqslant 8d^4$. *Moreover,* $\Theta \in \mathrm{Aut}(\mathcal{O}_K^2)$ *is either diagonal or anti-diagonal and its non-zero coefficients are in* $\mu(K)$.

Proof. Let $\Theta = \left(\begin{smallmatrix} a & c \\ b & d \end{smallmatrix}\right) \in \mathrm{Aut}(\mathcal{O}_K^2)$. The relation $\Theta^*\Theta = \mathrm{Id}$ implies $a\overline{a} + b\overline{b} = 1$ and $c\overline{c} + d\overline{d} = 1$, looking at the conditions on the diagonal entries. Following Corollary 2.3, either a or b equals 0, and either c or d equals 0, and the other ones are roots of unity. With the conditions on the anti-diagonal entries, we see that either a and d are both 0, or b and c are. Hence Θ is either diagonal or anti-diagonal and the non zero coefficients are in $\mu(K)$. Reciprocally, one can check that any diagonal or anti-diagonal matrix with coefficients that are roots of unity is in $\mathrm{Aut}(\mathcal{O}_K^2)$.We count $|\mu(K)|^2$ diagonal matrices, and the same number of anti-diagonal matrices. Lemma 2.4 then gives $|\mathrm{Aut}(\mathcal{O}_K^2)| = 2 \cdot |\mu(K)|^2 \leq 8d^4$. □

Remark 3.14. For a random module $M \subset K^2$ one would expect $\mathrm{Aut}(M) = \{\pm\mathrm{Id}\}$. The module $\mathcal{O}_K^2$ has many more automorphisms, in the same manner that $\mathbb{Z}^n$ has $2^n n!$ automorphisms (as a Euclidean lattice, [25, Section 1.1]). If $m = 2^\ell$ and ζ_m is a m-th primitive root of unity, then $\mathcal{O}^2_{\mathbb{Q}(\zeta_m)}$ is isometric to $\mathbb{Z}^m$. This result shows that it also has much less module lattice automorphisms than plain ones.

4 The special case of Hawk

In this final section, we restrict ourselves to cyclotomic number fields K, and to the module $M = \mathcal{O}_K^2$, given by a pseudo-basis $\mathbf{B} = (B, \mathfrak{a}_1, \mathfrak{a}_2)$. These restrictions are of particular interest, as they occur in Hawk's framework [15]. Before detailing this result, we get back to our initial point of view and formalize the geometric intuition we gave in the introduction.

4.1 A geometric insight on this reduction

We hope that this informal reduction to a module-SVP instance in rank 2, which is purely geometric, sheds yet another insight into this framework. For the sake of simplicity, we restrict the presentation to the case of Hawk, when the pseudo-basis $\mathbf{B}$ is $(\mathrm{Id}, \mathcal{O}_K, \mathcal{O}_K)$ itself and its associated pseudo-Gram matrix is $(\mathrm{Id}, \mathcal{O}_K, \mathcal{O}_K)$. While *it is* possible to deal with the general case with the technique we are presenting here, it appears to be quite cumbersome and the counterpart given by quaternion arithmetic is much better suited to do so.

Simplified setting. As mentioned in the introduction, we exploit the fact, thanks to the Gram-Schmidt orthonormalization process, that we can decompose any basis in a square lattice constructed *only* from its first vector. This is an avatar of the fact that a unimodular rank-2 module is necessarily *symplectic* for the determinant form, so we can fully describe it using a single primitive vector.

Write $C = \left(\begin{smallmatrix} x_1 & x_2 \\ y_1 & y_2 \end{smallmatrix}\right) \in \mathrm{GL}_2(\mathcal{O}_K)$ the secret Gram-root of $G' = \left(\begin{smallmatrix} q_{11} & q_{12} \\ q_{12} & q_{22} \end{smallmatrix}\right)$, and assume that $\det(C) = 1$.[16] Recall that K^2 is endowed with a natural inner product defined as $\langle x, y \rangle = x^* y$.

For this inner product, the vectors $\left(\begin{smallmatrix} x_1 \\ y_1 \end{smallmatrix}\right)$ and $\left(\begin{smallmatrix} -\overline{y_1} \\ \overline{x_1} \end{smallmatrix}\right)$ form an orthogonal basis of K^2. This basis is checked to be the K-orthonormalization of C for the inner product defined above, scaled by $\sqrt{q_{11}} = \sqrt{\mathrm{nrd}(x_1 + y_1 j)}$ (the square root is taken here over $K_{\mathbb{R}}$). Let us denote by (a, b) the coordinates of the vector $\left(\begin{smallmatrix} x_2 \\ y_2 \end{smallmatrix}\right)$ in this basis, that is $\left(\begin{smallmatrix} x_2 \\ y_2 \end{smallmatrix}\right) = a \left(\begin{smallmatrix} x_1 \\ y_1 \end{smallmatrix}\right) + b \left(\begin{smallmatrix} -\overline{y_1} \\ \overline{x_1} \end{smallmatrix}\right)$.

Taking inner product with the vector $\left(\begin{smallmatrix} x_1 \\ y_1 \end{smallmatrix}\right)$ on the left-hand side and the right-hand side gives:

$$\overline{x_2} x_1 + \overline{y_2} y_1 = \left\langle a \begin{pmatrix} x_1 \\ y_1 \end{pmatrix} + b \begin{pmatrix} -\overline{y_1} \\ \overline{x_1} \end{pmatrix}, \begin{pmatrix} x_1 \\ y_1 \end{pmatrix} \right\rangle = \overline{a}(x_1 \overline{x_1} + y_1 \overline{y_1})$$

By definition, $\overline{x_2} x_1 + \overline{y_2} y_1 = \overline{q_{12}}$ and $x_1 \overline{x_1} + y_1 \overline{y_1} = q_{11}$, so we have $a = \frac{q_{12}}{q_{11}}$. Moreover, since the determinant is an alternating multilinear form over vectors, we must have $b = \frac{1}{q_{11}}$, since:

$$1 = \det \begin{pmatrix} x_1 & x_2 \\ y_1 & y_2 \end{pmatrix} = b \cdot \det \begin{pmatrix} x_1 & -\overline{y_1} \\ y_1 & \overline{x_1} \end{pmatrix} = b \cdot q_{11}$$

For this inner product, the dual basis of B is $(B^*)^{-1}$. As such, we obtained the following matrix decomposition, using the fact that the dual basis of $\left(\begin{smallmatrix} x_1 & -\overline{y_1} \\ y_1 & \overline{x_1} \end{smallmatrix}\right)$ is exactly itself scaled by the factor q_{11}:

$$\begin{pmatrix} x_1 & -\overline{y_1} \\ y_1 & \overline{x_1} \end{pmatrix}^* \cdot C = \begin{pmatrix} q_{11} & q_{12} \\ 0 & 1 \end{pmatrix} \tag{2}$$

This now entails the preliminary remark we made: it is enough to find only $(\overline{x_1}, -y_1)$ in order to fully recover the full matrix C as there is a simple linear relation given by public data derived from G'. Since the lattice spanned by C is $\mathcal{O}_K^2$, this ensures that the lattices are the same:

$$\begin{pmatrix} x_1 & -\overline{y_1} \\ y_1 & \overline{x_1} \end{pmatrix}^* \mathcal{O}_K^2 = \begin{pmatrix} q_{11} & q_{12} \\ 0 & 1 \end{pmatrix} \mathcal{O}_K^2 := \mathcal{L}$$

As such, $\mathcal{L}$ admits an orthogonal basis as $\mathcal{O}_K$-module, entailing that $(\overline{x_1}, -y_1)$ is a shortest vector of this module by the orthogonality of the basis vector. We can then recover this vector (up to an automorphism of the lattice, which is not an issue as after completing it in a basis, we will also find a valid Gram root of G') as a module-SVP instance and reconstruct the corresponding secret elements x_1 and y_1.

Remark 4.1. Lemma 3.5 is similar to Equation (2). When seeing the matrix $\left(\begin{smallmatrix} \overline{x_1} & \overline{y_1} \\ -y_1 & x_1 \end{smallmatrix}\right)$ as a left regular representation of the quaternion $\overline{x_1} - \overline{y_1} j$, this is

[16] Let us recall quickly that in Hawk, the secret key is a Gram-root C of the public key G', and both have determinant 1 by construction.

expected. Indeed, this reduction at module level is just the geometrization of the arithmetic reduction presented in this paper. As such, finding the vector $(\overline{x_1}, -y_1)$ up to an automorphism is the transcription of finding the generator of the quaternionic ideal I_M of Proposition 3.6 up to a unit. Figure 1 provides an example over the integers.

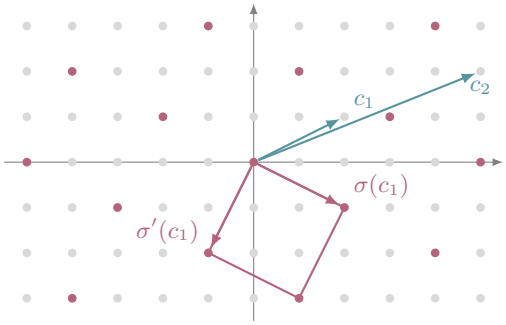

Fig. 1: Graphical depiction of the lattice involved in the attack for $C = \left(\begin{smallmatrix} 2 & 5 \\ 1 & 2 \end{smallmatrix}\right)$ in $\mathbb{Z}^2$, with c_1, c_2 its column vectors. The attack aims at recovering the *square* basis $\left(\begin{smallmatrix} \overline{x_1} & \overline{y_1} \\ -y_1 & x_1 \end{smallmatrix}\right) = \left(\begin{smallmatrix} 2 & -1 \\ -1 & -2 \end{smallmatrix}\right)$ of the lattice $\mathcal{L}$—which is described using the Gram-matrix of (c_1, c_2). This lattice corresponds to the principal ideal $I_M = (2 - j)$ in the ring $\mathbb{Z}[j]$, for $j^2 = -1$, explaining why it has such a square geometry.

Remark 4.2. We would have been able to get to this result in an even more elementary way by simply checking that:

$$\begin{aligned} -y_1 q_{12} &\equiv -y_1\overline{x_1}x_2 - y_1\overline{y_1}y_2 \pmod{q_{11}} \equiv -y_1\overline{x_1}x_2 + x_1\overline{x_1}y_2 \pmod{q_{11}} \\ &\equiv \overline{x_1}\underbrace{(x_1y_2 - y_1x_2)}_{=\det(C)=1} \pmod{q_{11}} \end{aligned}$$

entailing that $(\overline{x_1}, -y_1)$ and its conjugate belong to $\mathcal{L}$.

4.2 Instantiating our reduction for Hawk

Let us now go back to our quaternionic framework. We already mentioned that over cyclotomic fields, the reduction becomes uniform since many inputs in Algorithm 2 can be computed directly from the parameters. In addition to the uniformity, we will prove that a Karp reduction is possible in that case. This means that the same result can be achieved making only one call to the nrdPIP oracle.

Theorem 4.3. *Let $m \geq 31$ and K be the m-th cyclotomic field, of degree $d = \varphi(m)$. Let $M = \mathcal{O}_K^2$, with a pseudo-basis $\mathbf{B} = (B, \mathfrak{a}_1, \mathfrak{a}_2)$ and $\mathbf{G}$ its pseudo-Gram*

Algorithm 3: Karp reduction of wc-smodLIP to nrdPIP for $\mathcal{O}_K^2$

Input: • $K = \mathbb{Q}(\zeta_m)$ a cyclotomic field, $\mathbf{B} = (B, \mathfrak{a}_1, \mathfrak{a}_2)$ a pseudo-basis of $\mathcal{O}_K^2$;
• $\mathbf{G} = (G = B^*B, \mathfrak{a}_1, \mathfrak{a}_2)$ and $\mathbf{G}' = (G, \mathfrak{b}_1, \mathfrak{b}_2)$ congruent to $\mathbf{G}$;
• An oracle $\mathfrak{O}$ solving $\mathcal{O}$-nrdPIP.

Output: The set of all congruence matrices $\mathrm{Cong}(\mathbf{G}, \mathbf{G}')$.

1 $\mu(K) \leftarrow \langle \zeta_m \rangle \subset K^\times$
2 $\mathrm{Aut}(\mathcal{O}_K^2) \leftarrow \left\{ \left(\begin{smallmatrix} a & 0 \\ 0 & b \end{smallmatrix}\right) ; \left(\begin{smallmatrix} 0 & a \\ b & 0 \end{smallmatrix}\right) : a, b \in \mu(K) \right\}$ // *c.f.,* Corollary 3.13
3 $\mathcal{O} \leftarrow$ Run Algorithm 4 on the order $\mathcal{O}_K + \mathcal{O}_K \cdot j$ // *c.f.,* Proposition B.18
4 $I_M \leftarrow \mathcal{O}$; $\mathcal{O}' \leftarrow \mathcal{O}$ // *c.f.,* Corollary 3.3
5 $\mathcal{O}'^1 \leftarrow \langle \zeta_m, j \rangle$ // *c.f.,* Corollary B.27
6 `Congruence_mat` $\leftarrow \{\}$
7 $\delta_0 \leftarrow$ `ComputeDet`$(\mathbf{G}, \mathbf{G}')$ // *c.f.,* Lemma 2.23
8 Compute $\texttt{Congruence\_mat}_{\delta_0}$ with Algorithm 1 on input $(\mathbf{B}, \mathbf{G}', \delta_0, \mathcal{O}, I_M, \mathcal{O}', \mathcal{O}'^1)$ // *c.f.,* Theorem 3.7
9 Pick any $U_0 \in \texttt{Congruence\_mat}_{\delta_0}$
10 **for** $\Theta \in \mathrm{Aut}(\mathcal{O}_K^2)$ **do**
11 $\quad U \leftarrow B^{-1}\Theta B U_0$
12 $\quad$ `Congruence_mat` $\leftarrow$ `Congruence_mat` $\cup \{U\}$
13 Return `Congruence_mat`.

matrix. Let $\mathbf{G}' = (G', \mathfrak{b}_1, \mathfrak{b}_2)$ *be a pseudo-Gram matrix congruent to* $\mathbf{G}$*. Assume that we are given an oracle* $\mathfrak{O}$ *that solves* $\mathcal{O}$*-nrdPIP. Then Algorithm 3 returns the set* $\mathrm{Cong}(\mathbf{G}, \mathbf{G}')$ *of all congruence matrices between* $\mathbf{G}$ *and* $\mathbf{G}'$*. In particular it solves wc-smodLIP*$_K^{\mathbf{B}}$ *on input* $\mathbf{G}'$*. Moreover, it makes exactly one call to the oracle* $\mathfrak{O}$ *and except for this call it runs in time*

$$\mathrm{poly}(d, \mathrm{size}(\mathbf{G}), \mathrm{size}(\mathbf{G}')).$$

Proof. Correctness: First of all we justify that $\texttt{Congruence\_mat}_{\delta_0}$ computed at Step 8 is non empty. Since $\mathbf{G}$ and $\mathbf{G}'$ are chosen to be congruent, there exists a congruence matrix $U' \in \mathrm{Cong}(\mathbf{G}, \mathbf{G}')$, but we might have $\det(U') \neq \delta_0$ in general. However, by Lemma 2.23, it holds that $\det(U') = \mu\delta_0$ for some root of unity $\mu \in \mu(K)$. We claim that $U_0 = B^{-1} \cdot \mathrm{diag}(\mu^{-1}, 1) \cdot B \cdot U'$ belongs to $\mathrm{Cong}(\mathbf{G}, \mathbf{G}')$ too and has determinant $\det(U_0) = \mu^{-1} \cdot \det(U') = \delta_0$. The fact that $U_0 \in \mathrm{Cong}(\mathbf{G}, \mathbf{G}')$ follows from Corollary 3.11 and the fact that $\mathrm{diag}(\mu^{-1}, 1) \in \mathrm{Aut}(\mathcal{O}_K^2)$ thanks to Proposition 3.13. By the correctness of Algorithm 1, this means that $U_0 \in \texttt{Congruence\_mat}_{\delta_0}$ is non empty. For any such U_0 chosen during Step 9, Corollary 3.11 guarantees that the for loop computes iteratively exactly all the other solutions so by the end, the algorithm outputs indeed $\mathrm{Cong}(\mathbf{G}, \mathbf{G}')$.

Complexity: We need to argue that when K and M are as in the theorem, then the quantities $\mathcal{O}$, I_M, $\mathcal{O}'$, $\mathcal{O}'^1$ and $\mu(K)$ that are required as input of Algorithm 2 can be computed in polynomial time from the knowledge of K and $\mathbf{B}$. First, when K is a cyclotomic field of conductor m, then $\mathcal{O}_K = \mathbb{Z}[\zeta_m]$ can easily be computed, where ζ_m is a primitive m-th root of unity in K. Using Proposition B.18, a maximal order $\mathcal{O}$ of $\mathcal{A}$ containing $\mathcal{O}_K + \mathcal{O}_K \cdot j$ can be computed in polynomial

time. According to Corollary 3.3, we have $I_M = \mathcal{O}$ thus $\mathcal{O}' = \mathcal{O}$ and $\mathcal{O}'^1 = \mathcal{O}^1$. The latter equals $\langle \zeta_m, j \rangle$ for conductors $m \geq 31$, by Corollary B.27, so it can be computed in polynomial time. Finally, recall that $\log \Delta_K = \text{poly}(d)$ holds for cyclotomic fields. Hence, the complexity is a consequence of the above discussion, Lemma 2.23 for Step 7 and Theorem 3.7. □

Corollary 4.4 (modLIP to $\mathcal{O}$-nrdPIP, Hawk). *For any cyclotomic field K (with F its maximal totally real subfield) and pseudo-basis $\mathbf{B} = (B, \mathfrak{a}_1, \mathfrak{a}_2)$ of $\mathcal{O}_K^2$, there exists a uniform polynomial time Karp reduction from wc-smodLIP$_K^{\mathbf{B}}$ to $\mathcal{O}$-nrdPIP, where $\mathcal{O}$ is a maximal order of a quaternion algebra over F, and is efficiently computable from the parameters.*

Proof. When the conductor of K is $m \geq 31$, the reduction is provided by Algorithm 3 and the previous theorem. In that case, since Algorithm 3 makes only one call to the $\mathcal{O}$-nrdPIP oracle, the reduction is Karp. The fact that it is uniform follows from several observations, already mentioned. Indeed, $\mathcal{O}_K = \mathbb{Z}[\zeta_m]$ and $\mathcal{O} = I_M = \mathcal{O}'$ can be computed in polynomial time (see Proposition B.18), as well as the finite group $\mathcal{O}^1$ (see Corollary B.27 for conductors $m \geq 31$).

For lower conductors $m \leq 30$, we rely on a generic method that we describe below. On an input $\mathbf{G}' = (G', \mathfrak{b}_1, \mathfrak{b}_2)$ congruent to $\mathbf{G}$, one computes the "structured" Cholesky factorization of G' with coefficients in $K_{\mathbb{R}}$, that is, some $C \in \mathcal{M}_2(K_{\mathbb{R}})$ such that $C^*C = G'$ (see [33, Proposition 3.4] for more details). Observe that for any solution $U \in \text{Cong}(\mathbf{G}, \mathbf{G}')$, then $(BU)^*(BU) = G'$ is another factorization of G' (in K and *a fortiori* in $K_{\mathbb{R}}$). Thus, [33, Proposition 3.5]) ensures[17] the existence of a unitary transformation $\Theta \in \mathcal{U}_2(K_{\mathbb{R}})$ such that $C = \Theta \cdot B \cdot U$. Now from $\mathbf{B}$ and $\mathbf{C} := (C, \mathfrak{b}_1, \mathfrak{b}_2)$, we explain how to compute all such Θ, from which we will deduce the congruence matrices.

To $\mathbf{B}$ and $\mathbf{C}$ one associates the full-rank module lattices $\mathcal{L}(\mathbf{B})$, $\mathcal{L}(\mathbf{C}) \subset \mathbb{R}^{2d}$ using the canonical embedding. These two Euclidean lattices are isomorphic as module lattices and so *a fortiori* as "plain" lattices. In other words, this gives an instance of LIP as defined in [25]. Using Theorem 1.1 of [25], one computes all isomorphisms $O \in \mathcal{O}_{2d}(\mathbb{R})$ between $\mathcal{L}(\mathbf{B})$ and $\mathcal{L}(\mathbf{C})$ in time exponential in $d \leq 30$ here. Finally thanks to Corollary A.3, it is possible to check if O is a module lattice isomorphism Θ or not. When it is, we compute $U = (\Theta \cdot B)^{-1} \cdot C$ and check if $U \in \text{Cong}(\mathbf{G}, \mathbf{G}')$. Summing up, for $m \leq 30$, the algorithm we just described solves wc-smodLIP$_{\mathbf{B}}^K$ making no call to the oracle for nrdPIP, providing the claimed Karp reduction. Since all necessary structures can be computed efficiently from the parameters of the instance, it is also uniform. □

Acknowledgement. We are grateful to John Voight for helpful discussions but also for his great book on quaternion algebras, from which we learnt most of the material on this topic. We thank Alice Silverberg and Hendrik Lenstra for

[17] This is the usual way to move between two possible definitions of module-LIP (see [33, Lemma 3.10])

inspiring exchanges. Our thanks also go to Aurel Page and Renaud Coulangeon for many discussions and suggestions. Finally, we thank Pierre-Alain Fouque for insightfull comments and guidance through this work.

All the authors were supported by the France 2030 program, managed by the French National Research Agency under grant agreement No. ANR-22-PETQ-0008 PQ-TLS. Alice Pellet-Mary was supported by the CHARM ANR-NSF grant (ANR-21-CE94-0003) and the TOTORO ANR grant (ANR-23-CE48-0002).

References

1. Léo Ackermann, Adeline Roux-Langlois, and Alexandre Wallet. Public-key encryption from lip. In *International Workshop on Coding and Cryptography (WCC)*, 2024.
2. Chandrashekar Adiga, Ismail Naci Cangul, and HN Ramaswamy. On the constant term of the minimal polynomial of $\cos(2\pi n)$ over $\mathbb{Q}$. *Filomat*, 30(4):1097–1102, 2016.
3. Huck Bennett, Atul Ganju, Pura Peetathawatchai, and Noah Stephens-Davidowitz. Just how hard are rotations of $\mathbb{Z}^n$? algorithms and cryptography with the simplest lattice. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 252–281. Springer, 2023.
4. Jean-François Biasse, Thomas Espitau, Pierre-Alain Fouque, Alexandre Gélin, and Paul Kirchner. Computing generator in cyclotomic integer rings - A subfield algorithm for the principal ideal problem in L(1/2) and application to the cryptanalysis of a FHE scheme. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *Advances in Cryptology - EUROCRYPT 2017 - 36th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Paris, France, April 30 - May 4, 2017, Proceedings, Part I*, volume 10210 of *Lecture Notes in Computer Science*, pages 60–88, 2017.
5. Jean-François Biasse and Fang Song. Efficient quantum algorithms for computing class groups and solving the principal ideal problem in arbitrary degree number fields. In Robert Krauthgamer, editor, *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 893–902. SIAM, 2016.
6. Jean-François Biasse and Claus Fieker. Subexponential class group and unit group computation in large degree number fields. *LMS Journal of Computation and Mathematics*, 17(A):385–403, 2014.
7. Jean-François Biasse, Claus Fieker, and Tommy Hofmann. On the computation of the hnf of a module over the ring of integers of a number field. *Journal of Symbolic Computation*, 80:581–615, 2017.
8. Werner Bley, Tommy Hofmann, and Henri Johnston. Computation of lattice isomorphisms and the integral matrix similarity problem. *Forum of Mathematics, Sigma*, 10:e87, 2022.
9. Clémence Chevignard, Pierre-Alain Fouque, Guilhem Mureau, Alice Pellet-Mary, and Alexandre Wallet. A reduction from hawk to the principal ideal problem in a quaternion algebra. Cryptology ePrint Archive, Paper 2024/1147, 2024.
10. Henri Cohen. *A Course in Computational Algebraic Number Theory*. Springer Publishing Company, Incorporated, 2010.
11. Henri Cohen. *Advanced topics in computational number theory*, volume 193. Springer Science & Business Media, 2012.

12. Ronald Cramer, Léo Ducas, Chris Peikert, and Oded Regev. Recovering short generators of principal ideals in cyclotomic rings. In Marc Fischlin and Jean-Sébastien Coron, editors, *Advances in Cryptology - EUROCRYPT 2016*, volume 9666 of *Lecture Notes in Computer Science*, pages 559–585. Springer, 2016.
13. Ronald Cramer, Léo Ducas, and Benjamin Wesolowski. Short stickelberger class relations and application to ideal-svp. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *Advances in Cryptology - EUROCRYPT 2017*, volume 10210 of *Lecture Notes in Computer Science*, pages 324–348, 2017.
14. Artūras Dubickas and Chris Smyth. Two variations of a theorem of kronecker. *Expositiones Mathematicae*, 23(3):289–294, 2005.
15. Léo Ducas, Eamonn W. Postlethwaite, Ludo N. Pulles, and Wessel P. J. van Woerden. Hawk: Module LIP makes lattice signatures fast, compact and simple. In Shweta Agrawal and Dongdai Lin, editors, *Advances in Cryptology - ASIACRYPT 2022 - 28th International Conference on the Theory and Application of Cryptology and Information Security, Taipei, Taiwan, December 5-9, 2022, Proceedings, Part IV*, volume 13794 of *Lecture Notes in Computer Science*, pages 65–94. Springer, 2022.
16. Léo Ducas and Wessel P. J. van Woerden. On the lattice isomorphism problem, quadratic forms, remarkable lattices, and cryptography. In Orr Dunkelman and Stefan Dziembowski, editors, *Advances in Cryptology - EUROCRYPT 2022 - 41st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Trondheim, Norway, May 30 - June 3, 2022, Proceedings, Part III*, volume 13277 of *Lecture Notes in Computer Science*, pages 643–673. Springer, 2022.
17. Léo Ducas and Wessel P. J. van Woerden. On the lattice isomorphism problem, quadratic forms, remarkable lattices, and cryptography. In Orr Dunkelman and Stefan Dziembowski, editors, *Advances in Cryptology - EUROCRYPT 2022 - 41st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Trondheim, Norway, May 30 - June 3, 2022, Proceedings, Part III*, volume 13277 of *Lecture Notes in Computer Science*, pages 643–673. Springer, 2022.
18. Mathieu Dutour Sikirić, Anna Haensch, John Voight, and Wessel PJ van Woerden. A canonical form for positive definite matrices. *Open Book Series*, 4(1):179–195, 2020.
19. Thomas Espitau, Pierre-Alain Fouque, Benoît Gérard, and Mehdi Tibouchi. Side-channel attacks on BLISS lattice-based signatures: Exploiting branch tracing against strongswan and electromagnetic emanations in microcontrollers. In Bhavani Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS 2017, Dallas, TX, USA, October 30 - November 03, 2017*, pages 1857–1874. ACM, 2017.
20. Thomas Espitau and Heorhii Pliatsok. On hermitian decomposition lattices and the module-LIP problem in rank 2. Cryptology ePrint Archive, Paper 2024/1148, 2024.
21. Pierre-Alain Fouque, Paul Kirchner, Mehdi Tibouchi, Alexandre Wallet, and Yang Yu. Key recovery from gram-schmidt norm leakage in hash-and-sign signatures over NTRU lattices. In Anne Canteaut and Yuval Ishai, editors, *Advances in Cryptology - EUROCRYPT 2020 - 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, May 10-14,*

2020, Proceedings, Part III, volume 12107 of *Lecture Notes in Computer Science*, pages 34–63. Springer, 2020.
22. Sanjam Garg, Craig Gentry, and Shai Halevi. Candidate multilinear maps from ideal lattices. In Thomas Johansson and Phong Q. Nguyen, editors, *Advances in Cryptology - EUROCRYPT 2013, 32nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Athens, Greece, May 26-30, 2013. Proceedings*, volume 7881 of *Lecture Notes in Computer Science*, pages 1–17. Springer, 2013.
23. Craig Gentry and Michael Szydlo. Cryptanalysis of the revised NTRU signature scheme. In Lars R. Knudsen, editor, *Advances in Cryptology - EUROCRYPT 2002, International Conference on the Theory and Applications of Cryptographic Techniques, Amsterdam, The Netherlands, April 28 - May 2, 2002, Proceedings*, volume 2332 of *Lecture Notes in Computer Science*, pages 299–320. Springer, 2002.
24. Emmanuel Hallouin and Christian Maire. Cancellation in totally definite quaternion algebras. *Journal für die reine und angewandte Mathematik*, 2006(595):189–213, 2006.
25. Ishay Haviv and Oded Regev. On the lattice isomorphism problem. pages 391–404, 2014.
26. Nick Howgrave-Graham and Michael Szydlo. A method to solve cyclotomic norm equations. In Duncan A. Buell, editor, *Algorithmic Number Theory, 6th International Symposium, ANTS-VI, Burlington, VT, USA, June 13-18, 2004, Proceedings*, volume 3076 of *Lecture Notes in Computer Science*, pages 272–279. Springer, 2004.
27. Hendrik W. Lenstra Jr. and Alice Silverberg. Testing isomorphism of lattices over CM-orders. *SIAM J. Comput.*, 48(4):1300–1334, 2019.
28. Markus Kirschmer. Definite quadratic and hermitian forms with small class number. *Habilitation, RWTH Aachen University*, 2016.
29. Markus Kirschmer and John Voight. Algorithmic enumeration of ideal classes for quaternion orders. *SIAM Journal on Computing*, 39(5):1714–1747, 2010.
30. Cong Ling and Andrew Mendelsohn. NTRU in quaternion algebras of bounded discriminant. In Thomas Johansson and Daniel Smith-Tone, editors, *Post-Quantum Cryptography - 14th International Workshop, PQCrypto 2023, College Park, MD, USA, August 16-18, 2023, Proceedings*, volume 14154 of *Lecture Notes in Computer Science*, pages 256–290. Springer, 2023.
31. Hengyi Luo, Kaijie Jiang, Yanbin Pan, and Anyu Wang. Cryptanalysis of rank-2 module-LIP with symplectic automorphisms. Cryptology ePrint Archive, Paper 2024/1173, 2024.
32. Daniel Marcus. *Number Fields*. 01 1977.
33. Guilhem Mureau, Alice Pellet-Mary, Georgii Pliatsok, and Alexandre Wallet. Cryptanalysis of rank-2 module-lip in totally real number fields. In Marc Joye and Gregor Leander, editors, *Advances in Cryptology - EUROCRYPT 2024 - 43rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zurich, Switzerland, May 26-30, 2024, Proceedings, Part VI*, volume 14656 of *Lecture Notes in Computer Science*, pages 226–255. Springer, 2024.
34. Jürgen Neukirch. *Algebraic Number Theory*. 05 1999.
35. NIST. Round 4 standardisation results for the post-quantum cryptography standardization process, 2024.
36. Wilhelm Plesken and Bernd Souvignier. Computing isometries of lattices. *Journal of Symbolic Computation*, 24(3-4):327–334, 1997.
37. Irving Reiner. *Maximal orders*. Oxford University Press, 2003.

38. Jean-Pierre Serre. *A course in arithmetic*, volume 7. Springer Science & Business Media, 2012.
39. Michael Szydlo. Hypercubic lattice reduction and analysis of GGH and NTRU signatures. In Eli Biham, editor, *Advances in Cryptology - EUROCRYPT 2003, International Conference on the Theory and Applications of Cryptographic Techniques, Warsaw, Poland, May 4-8, 2003, Proceedings*, volume 2656 of *Lecture Notes in Computer Science*, pages 433–448. Springer, 2003.
40. John Voight. *Identifying the Matrix Ring: Algorithms for Quaternion Algebras and Quadratic Forms*, pages 255–298. Springer New York, New York, NY, 2013.
41. John Voight. *Quaternion Algebras*. Springer Nature, 01 2021.
42. Lawrence C. Washington. *Introduction to Cyclotomic Fields*. 01 1982.

A Proofs of Section 2 and algorithmic considerations

In this appendix, we present all the proofs for the results Section 2, that where delayed earlier. We also discuss in detail how to represent the mathematical structures involved, and how to carry ground-level computations.

A.1 Proofs of Section 2

Proof of Corollary 2.3. By Lemma 2.1), for all embeddings σ_i of K, $\sigma_i(a\overline{a}) = |\sigma_i(a)|^2$ and $\sigma_i(b\overline{b}) = |\sigma_i(b)|^2$ are both positive. Moreover, we have $\sigma_i(a\overline{a}) + \sigma_i(b\overline{b}) = \sigma_i(1) = 1$. Suppose that $a \neq 0$, so $0 < \sigma_i(a\overline{a}) \leq 1$ for all i's. Then Proposition 2.2 implies that $a\overline{a}$ must be a root of unity in K. But $a\overline{a}$ is totally positive so $a\overline{a} = 1$ and $b = 0$, which also implies that $|\sigma_i(a)| = 1$. Applying again Lemma 2.2 to a, we conclude that a is a root of unity. □

Proof of Proposition 2.12. Suppose that $II' = \mathcal{O}_\ell(I')$. Then multiplying on the left by I^{-1} gives $\mathcal{O}_r(I)I' = I^{-1}\mathcal{O}_\ell(I)$ but $\mathcal{O}_\ell(I') = \mathcal{O}_r(I)$ by assumption, so $\mathcal{O}_r(I)I' = I'$. By definition of the pseudo-inverse one has $\mathcal{O}_r(I^{-1}) = \mathcal{O}_\ell(I)$, so $I^{-1}\mathcal{O}_\ell(I) = I^{-1}$. Therefore one obtains $I' = I^{-1}$ as expected. The converse is contained in the definition of being an invertible $\mathcal{O}_F$-lattice. □

Proof of Lemma 2.13. Using the second definition of the quasi-inverse, we have $I^{-1} = \{\alpha \in \mathcal{A} \,|\, I\alpha \subseteq \mathcal{O}_\ell(I)\}$. Similarly $J^{-1} = \{\alpha \in \mathcal{A} \,|\, J\alpha \subseteq \mathcal{O}_\ell(J)\}$. Using the fact that $J \subseteq I$ and that $\mathcal{O}_\ell(I) = \mathcal{O}_\ell(J)$, we have that any element $\alpha \in I^{-1}$ verifies $J\alpha \subseteq I\alpha \subseteq \mathcal{O}_\ell(I) = \mathcal{O}_\ell(J)$, so $\alpha \in J^{-1}$. □

Proof of Lemma 2.16. We do the proof when $\mathcal{O}_\ell(I) = \mathcal{O}_\ell(J)$, the case where $\mathcal{O}_r(I) = \mathcal{O}_r(J)$ being analogous. Since I invertible and $I \subset J$, we have $I^{-1}I = \mathcal{O}_r(I) \subset I^{-1}J$. By hypothesis the latter is a product of compatible ideals, hence by Lemma 2.14 *1.*, $\mathrm{nrd}(I^{-1}J) = \mathrm{nrd}(I^{-1})\mathrm{nrd}(J)$. Since $\mathcal{O}_F = \mathrm{nrd}(II^{-1}) = \mathrm{nrd}(I)\mathrm{nrd}(I^{-1})$, we have $\mathrm{nrd}(I^{-1}) = \mathrm{nrd}(I)^{-1}$ and $\mathrm{nrd}(I^{-1}J) = \mathrm{nrd}(I)^{-1}\mathrm{nrd}(J) = \mathcal{O}_F$. Thus, the element $1 \in \mathcal{O}_r(I) \subset I^{-1}J$ generates $\mathrm{nrd}(1)\mathcal{O}_F = \mathcal{O}_F = \mathrm{nrd}(I^{-1}J)$, so by Lemma 2.14 *2.*, $I^{-1}J = \mathcal{O}_r(J)$ and we conclude $I = J$. □

Proof of Proposition 2.17. Since $\mathcal{O}$ is maximal, we know that I is a sated left $\mathcal{O}$-ideal, i.e., $\mathcal{O}_\ell(I) = \mathcal{O} = \mathcal{O}_\ell(J_i)$ for all i. Moreover, for any $1 \leq i \leq n$, we have $J_i \subset I$ so we can apply Lemma 2.13, which gives $I^{-1} \subset J_i^{-1}$. Therefore, $I^{-1} \subset J_1^{-1} \cap \cdots \cap J_n^{-1}$.

Conversely, let $x \in J_1^{-1} \cap \cdots \cap J_n^{-1}$. Then $Ix = (J_1 + \cdots + J_n)x = J_1x + \cdots + J_nx$. Since $x \in J_i^{-1}$ for all i, and by the second definition of the quasi-inverse, it holds that $J_ix \subseteq \mathcal{O}_\ell(J_i) = \mathcal{O}$. Thus $Ix \subset \mathcal{O} = \mathcal{O}_\ell(I)$ which means $x \in I^{-1}$. We conclude that $J_1^{-1} \cap \cdots \cap J_n^{-1} \subset I^{-1}$, as wanted. □

Proof of Lemma 2.19. Almost everything is contained in Lemma 2.14 (*2.*). The only thing we need to prove is that the solutions are all equal up to right multiplication by an element in $\mathcal{O}^1$. Let $g' \in I$ be another solution. Since $I = g\cdot\mathcal{O}$, there exists $u \in \mathcal{O}$ such that $g' = gu$. But then $\mathrm{nrd}(g) = q = \mathrm{nrd}(g')$ (and the

multiplicativity of nrd) implies $\mathrm{nrd}(u) = 1$, *i.e.,* $u \in \mathcal{O}^1$. The converse is true: if $u \in \mathcal{O}^1$ then $gu \in I$ has reduced norm q. □

Proof of Lemma 2.22. Observe that the condition *1.* in Definition 2.20 is equivalent to $G' = C^*C$, where $C = BU$ and *2., 3.* are the necessary and sufficient conditions for U to be a pseudo-base change between $\mathbf{B}$ and $\mathbf{C}$, *i.e.,* for $\mathbf{C}$ to be a pseudo-basis of the same module M. □

Proof of Lemma 2.23. Let U be a congruence matrix between $\mathbf{G}$ and $\mathbf{G'}$. Recall that $\mathbf{G}$ is a pseudo-Gram matrix associated to a pseudo-basis of a full rank module in $\mathcal{O}_K^2$. Therefore, $\det G \neq 0$. By definition U satisfies $G' = U^*GU$ so taking the determinant we see that $\det U$ is a solution to the norm equation $\overline{x}x = \det G'/\det G$. Another property of U is that $\prod_i \mathfrak{a}_i = (\det U)\prod_i \mathfrak{b}_i$. In particular, $\det U$ is a generator of the fractional ideal $I = \prod_i \mathfrak{a}_i\mathfrak{b}_i^{-1}$. Any other congruence matrix U' satisfies again these two conditions: $\det U'$ is a generator of I, so one can write $\det U' = u \cdot \det U$ with $u \in \mathcal{O}_K^\times$, and the fact that $\det U'$ is a solution to the same norm equation gives $\overline{u}u = 1$. By Kronecker's theorem, we conclude that u is a root of unity.

Knowing $\mathbf{G}$ and $\mathbf{G'}$, we can call the Lenstra-Silverberg algorithm [27, Theorem 1.3] with inputs I and relative norm $\det G'/\det G$ (and a basis of $\mathcal{O}_K$). This algorithm outputs (if it exists) a generator x of I such that $x\bar{x} = \det G'/\det G$, and runs in polynomial time. This provides us with the determinant of our congruence matrix U, up to a root of unity. □

Then, we discuss the algorithmic aspects of the reductions presented above.

A.2 Representations of objects

This section covers how we represent mathematical objects to carry actual computations. We borrow most arguments from [33, Section 2.3].

Lattices in $\mathbb{R}^\ell$. Let $1 \leqslant r \leqslant \ell$ be integers, and a fixed set of r independents vectors of $\mathbb{R}^\ell$, noted $\boldsymbol{b}_1, \ldots, \boldsymbol{b}_r$. The $\mathbb{Z}$-lattice of $\mathbb{R}^\ell$ of dimension r generated by the $\boldsymbol{b}_i$'s, is the set $\mathcal{L}(\boldsymbol{b}_1|\ldots|\boldsymbol{b}_r) := \{\sum a_i\boldsymbol{b}_i,\ a_i \in \mathbb{Z}\}$. This set is discrete and stable by addition. When $r = \ell$, we say that $\mathcal{L}$ is *full rank*.

From now on, we will only manipulate full rank lattices when dealing with lattices in $\mathbb{R}^\ell$. Consider a matrix $B \in \mathbf{GL}_\ell(\mathbb{R})$. Since B is invertible, their column vectors are independent, and span a full rank lattice $\mathcal{L}(B)$. To represent lattices in $\mathbb{R}^\ell$, we use such matrices B, in a form that is called "LLL-reduced".

Representations of ground objects. While we consider several sets of numbers, they are all built on a ground, totally real, number field F of degree d. We therefore chose this field as the base for representing all elements. Let $\alpha_1, \ldots, \alpha_d$ be a $\mathbb{Z}$-basis[18] of $\mathcal{O}_F$. An element $x \in F$ is represented by its rational coordinates

[18] Note however that computing such a basis may be an expensive task. It is a standard practice to assume that such a basis is available, at the cost of having non-uniform reductions. In most of practical usecases, a good basis is explicitely known.

in the basis $(\alpha_1, \ldots, \alpha_d)$. The size of a rational is the sum of the bit-size of its numerator and denominator, and the size of an element $x \in F$ is defined as $\mathrm{size}(x) = \sum_i \mathrm{size}(x_i)$, where x_i are the coordinates of x in the give basis of $\mathcal{O}_F$. A fractional $\mathcal{O}_F$-ideal $\mathfrak{a}$ is also a $\mathbb{Z}$-module of rank d, and admits a $\mathbb{Z}$-basis $(a_1, \ldots, a_d)$ — this includes the case of $\mathcal{O}_F$. There are many such bases for a given ideal, but we can always assume that $(\sigma(a_1), \ldots, \sigma(a_d))$ is LLL-reduced for the so-called T_2-norm $\|a\|^2 := \sum_i |\sigma_i(a)|^2$. Then the size of an ideal will be $\mathrm{size}(I) = \sum_i \mathrm{size}(a_i)$, where the a_i's are reduced in the sense above.

By LLL-reducedness and following the arguments presented in [33, Section 2.3], one can show that $\mathrm{size}(x) \leq \mathrm{poly}(\log \Delta_F, \|\sigma(x)\|)$ as well as $\|\sigma(x)\| \leq \mathrm{poly}(\log \Delta_F, \mathrm{size}(x))$ for all $x \in K$. Additionally, an integral $\mathcal{O}_F$-ideal $\mathfrak{a}$ can be represented with $\mathrm{size}(\mathfrak{a}) = \mathrm{poly}(\log \Delta_F, \log N(\mathfrak{a}))$, where $N(\mathfrak{a}) = [\mathcal{O}_F : \mathfrak{a}]$ is the algebraic norm of the ideal $\mathfrak{a}$.

Representations in extensions and of modules. Recall that we are in the setting of a totally negative $a \in F$ and a quaternion algebra $\mathcal{A} = (\frac{a,-1}{F})$. Then, the CM-extension $K = F(\sqrt{a})$ can be seen as a F-linear space of dimension 2 and basis $\{1, \sqrt{a}\}$. All $x \in K$ have coordinates $(x_1, x_2) \in F^2$ in this basis, and can thus be represented as a vector in $\mathbb{Q}^{2d}$. Likewise, since $\mathcal{A}$ is 4-dimensional over F with basis $\{1, i, j, ij\}$, every element of $\mathcal{A}$ has 4 coordinates in this basis, and corresponds to a vector in $\mathbb{Q}^{4d}$. The size of elements of K and $\mathcal{A}$ is then the sum of the sizes of their F-coordinates. For a matrix B with entries in F and ℓ columns b_i, its size is $\mathrm{size}(B) := \sum_{i \leq \ell} \mathrm{size}(b_i)$.

Fractional $\mathcal{O}_K$-ideals can be viewed as rank 2 modules over $\mathcal{O}_F$ living in $K \simeq F^2$. Similarly, quaternionic ideals in $\mathcal{A}$ are also $\mathcal{O}_F$-modules (of rank 4) in $\mathcal{A}$. Any such module has a pseudo-basis $(B, \{\mathfrak{a}_i\}_{i \leq \ell})$. According to the representation of elements above, B is a 2 by 2 or 4 by 4 matrix with entries in F and the $\mathfrak{a}_i$'s are fractional $\mathcal{O}_F$-ideals given by a LLL-reduced basis. The size of such an object M is then $\mathrm{size}(M) := \mathrm{size}(B) + \sum_i \mathrm{size}(\mathfrak{a}_i)$. Likewise, pseudo-Gram matrices $\boldsymbol{G}$ are represented by tuples $(G, \{\mathfrak{a}_i\}_{i \leq \ell})$, with G that is also a 2 by 2 or 4 by 4 matrix with entries in F, and $\mathfrak{a}_i$ fractional $\mathcal{O}_F$-ideals that supports the same assumptions as above. Therefore, $\mathrm{size}(\boldsymbol{G}) := \mathrm{size}(G) + \sum_i \mathrm{size}(\mathfrak{a}_i)$.

Computing arithmetic operations with modules. Still following [33, Section 2.3], we have $\mathrm{size}(x \cdot y) \leq \mathrm{poly}(\mathrm{size}(x), \mathrm{size}(y), \log(\Delta_F))$ for all $x, y \in F$. We now turn to ideals in F, in the CM extension K and the quaternion algebra $(\frac{a,-1}{F})$. Generally, they are all finitely generated $\mathcal{O}_F$-modules in $\mathcal{A} \simeq F^4$ of respective rank $1, 2$ or 4. This gives a convenient way to do arithmetic with them, whenever the target operation makes sense (*e.g.* compatibility for the product of quaternion ideals). Indeed, it is known that the sum, the intersection, the product of two $\mathcal{O}_F$-modules I, J can be computed from generating sets and the use of the pseudo-Hermite Normal Form algorithm [11, Section 1.5.2]. Noting that in our case the rank over $\mathcal{O}_F$ is bounded by 4, there exists version of this algorithm running in time $\mathrm{poly}(\mathrm{size}(I), \mathrm{size}(J), \log|\Delta_F|)$, see e.g. [7]. If I is invertible, computing I^{-1} can be done by using that $I^{-1} = \bar{I}\mathrm{nrd}(I)^{-1}$ [41, 16.6.14].

Selecting module lattice isomorphisms from isomorphisms of lattices. Recall that two (Euclidean) lattices $L, L' \subset \mathbb{R}^\ell$ are said to be *isomorphic* if there exists an orthogonal matrix $O \in \mathcal{O}_\ell(\mathbb{R})$ such that $L' = O \cdot L$. Such a matrix O is called an isomorphism between the lattices L and L'.

Given a module $M \subset K^\ell$ represented by a pseudo-basis $\mathbf{B}$, one can associate to it a full-rank lattice $\mathcal{L} = \mathcal{L}(\mathbf{B}) \subset \mathbb{R}^{d\ell}$, once a $\mathbb{Z}$-basis of $\mathcal{O}_K$ has been fixed. A natural question is to decide when an isomorphism between $\mathcal{L}(\mathbf{B})$ and $\mathcal{L}(\mathbf{C})$ (where $\mathbf{C}$ stands for a pseudo-basis of a module $M' \subset K^\ell$) actually corresponds to a module lattice isomorphism between M and M'. The answer is given in the following lemma.

Lemma A.1 ([28, Lemma 2.4.3, adapted]). *Let $K = \mathbb{Q}(\zeta)$ be a CM field of degree d and $M, M' \subset K^\ell$ be two modules, given by pseudo-bases $\mathbf{B}$ and $\mathbf{C}$. Suppose that $\sigma : K^\ell \to K^\ell$ is a $\mathbb{Q}$-linear map, represented by some $\Sigma \in \mathcal{M}_{d\ell}(\mathbb{Q})$ in a fixed $\mathbb{Q}$-basis of K^ℓ. The following statements are equivalent:*

1. *σ is an isomorphism of module lattices between M and M'.*
2. *$\Sigma \cdot \mathcal{L}(\mathbf{B}) = \mathcal{L}(\mathbf{C})$ and*

$$\mathrm{Tr}\Big(\alpha\sigma(\mathbf{v})^*\sigma(\mathbf{w}) + \overline{\alpha\sigma(\mathbf{v})^*\sigma(\mathbf{w})}\Big) = \mathrm{Tr}\Big(\alpha\mathbf{v}^*\mathbf{w} + \overline{\alpha\mathbf{v}^*\mathbf{w}}\Big), \tag{3}$$

for all $\mathbf{v}, \mathbf{w} \in K^\ell$ and $\alpha \in \{1, \zeta\}$.

Remark A.2. Since the form $(\mathbf{v}, \mathbf{w}) \mapsto \mathrm{Tr}(\alpha\mathbf{v}^*\mathbf{w} + \overline{\alpha\mathbf{v}^*\mathbf{w}})$ is $\mathbb{Q}$-bilinear, it is enough to check the condition *2.* on a $\mathbb{Q}$-generating set of K^ℓ.

Corollary A.3. *With the same notations as in the previous lemma, suppose that we are given an isomorphism $\Sigma \in \mathcal{O}_{d\ell}(\mathbb{R})$ between the lattices $\mathcal{L}(\mathbf{B})$ and $\mathcal{L}(\mathbf{C})$. Then there is an algorithm to determine if Σ is the ground representation of a module lattice isomorphism σ between M and M'. Moreover for fixed ℓ, this algorithm runs in polynomial time in d.*

Proof. Let us denote by $\mathcal{B}$ a fixed $\mathbb{Q}$-basis of K^ℓ, containing $d\ell$ elements. Thanks to $\mathcal{B}$, one can check in polynomial time if Σ is a $\mathbb{Q}$-endomorphism of K^ℓ. If it is not, then the algorithm returns $\perp$. Otherwise, looping over all $\mathbf{v}, \mathbf{w}$ in a $\mathcal{B} \times \mathcal{B}$, the algorithm computes $\sigma(\mathbf{v}) := \Sigma \cdot \mathbf{v}^t$ and $\sigma(\mathbf{w}) := \Sigma \cdot \mathbf{w}^t$, seen as elements of K^ℓ, and check if Equation (3) is satisfied for $\alpha \in \{1, \zeta\}$. If the condition is not satisfied, the algorithm returns $\perp$; if it finishes the loop, it returns True. Each of these computations can be done in polynomial time, and there are at most $2(d\ell)^2$ of them. The correctness is guaranteed by Lemma A.1. □

B Supplementary material

The aim of this appendix is to justify the computability of some structures used in Section 3, to prove our reduction in the case where K is a cyclotomic field and $M = \mathcal{O}_K^2$. Precisely, Algorithm 3 requires to compute a maximal order $\mathcal{O}$ containing $\mathcal{O}_K + \mathcal{O}_K \cdot j$ and the finite group $\mathcal{O}^1$.

Subsection B.2 is devoted to the computation of maximal orders in quaternion algebras $\mathcal{A}_m = (\frac{a_m,-1}{F_m})$. A preliminary step is to compute the discriminant of $\mathcal{A}_m$, which is discussed in B.1. Finally in B.3, we explicit the group $\mathcal{O}^1$ for big enough conductors m, thanks to the classification given in [41, Chapter 32].

B.1 Discriminant of a quaternion algebra

Places and ramification. The complex embeddings σ of a number field F provides absolute values $v_\sigma(x) = |\sigma(x)|$, and completing F with respect to them yields $\mathbb{R}$ or $\mathbb{C}$, depending on whether σ is real or complex — these are often called archimedian absolute values. Other absolute values can be obtained by looking at prime ideals. For a prime ideal $\mathfrak{p}$ of a number field F, the $\mathfrak{p}$-adic valuation of $x \in \mathcal{O}_F$ is the largest integer $e_\mathfrak{p}(x)$ such that $\mathfrak{p}^e | x\mathcal{O}_F$. This yields a corresponding $\mathfrak{p}$-adic absolute value $v_\mathfrak{p}(x) = N(\mathfrak{p})^{-e_\mathfrak{p}(x)}$, and accordingly a corresponding $\mathfrak{p}$-adic completion $F_\mathfrak{p}$. In a generic way, from now on we denote by v an arbitrary absolute value of F, and the completion of F at v as the field F_v. We may also call v a *place*[19] of F. Given a quaternion algebra $\mathcal{A}$ over F and a place v of F, one can extend the scalars of $\mathcal{A}$ from F to F_v, giving the quaternion algebra $\mathcal{A}_v := \mathcal{A} \otimes_F F_v$ over F_v.

Wedderburn-Artin theorem [41, Corollary 7.3.12] states that a quaternion algebra $\mathcal{A}$ over a field F is either isomorphic to $M_2(F)$, or a division algebra (i.e., a non necessarily commutative ring in which every non zero element has an inverse). In the first case, called the split case, all the completions are isomorphic to a matrix algebra : $\mathcal{A}_v \simeq M_2(F) \otimes_F F_v = M_2(F_v)$. When $\mathcal{A}$ is a division ring, $\mathcal{A}_v$ can be either a matrix algebra or again a division ring. This leads to the notion of ramification.

Definition B.1 ([41, 14.5.1 and 14.3.1]). *Let v a place of F. We say that the algebra $\mathcal{A}$ is ramified at v if $\mathcal{A}_v = \mathcal{A} \otimes_F F_v$ is a division ring, which means that every nonzero element has an inverse. Otherwise we say that $\mathcal{A}$ is split (or unramified) at v.*

We denote $\operatorname{Ram}\mathcal{A}$ the set of ramified places of $\mathcal{A}$. This set is finite [41, Lem. 14.5.3]. Analogously as the discriminant for relative extensions of number fields, the discriminant of $\mathcal{A}$ is an integral ideal of $\mathcal{O}_F$, defined as the product of the finite ramified places in $\mathcal{A}$.

$$\operatorname{disc}_F(\mathcal{A}) := \prod_{\substack{\mathfrak{p} \in \operatorname{Ram}(\mathcal{A}) \\ \mathfrak{p} \text{ finite}}} \mathfrak{p}.$$

From its definition, it is clear that the discriminant encodes the ramification at finite places. The behaviour at infinite places leads to the definition of totally

[19] Formally, the language of *places* allows to avoid explicit choices of valuations, since a place of a number field F is defined as an equivalence class of non-trivial absolute values on F.

definite and indefinite algebras. In the core of this paper we focused on the algebras $(\frac{a,-1}{F})$ where $K = F(\sqrt{a})/F$ is a CM extension. They fall into the category of totally definite quaternion algebras, an important property which implies, for example, the finiteness of the groups $\mathcal{O}^1$ (see B.3).

Definition B.2 ([41], 14.5.7). *We say that $\mathcal{A}$ is totally definite if all archimedean places of F are ramified in $\mathcal{A}$; otherwise, we say $\mathcal{A}$ is indefinite.*

Hilbert symbol. To check if a quaternion algebra $\mathcal{A}$ over F ramifies at some place v of F, one can compute a Hilbert symbol. In the following we give the definition of the Hilbert symbol and we stand some properties useful for our purpose. A standard reference for the theory of Hilbert symbol is [38, Chapter III] but all the following results can be found in [41].

Definition B.3. *Let $\mathcal{A} = (\frac{a,b}{F})$ be a quaternion algebra over a number field F and v be a place of F (either finite of infinite). The Hilbert symbol of $\mathcal{A}$ at v is*

$$\left(\frac{a,b}{v}\right) := \begin{cases} 1 \text{ if } x^2 - ay^2 - bz^2 = 0 \text{ has a non trivial solution in } (F_v)^3 \\ -1 \text{ otherwise} \end{cases}$$

Let us link the Hilbert symbol with the ramification. Recall that an element $\alpha = x + iy + jz + kt \in \mathcal{A}$ has reduced norm $\mathrm{nrd}(\alpha) = x^2 - ay^2 - bz^2 + abt^2$. In this expression, one recognizes the quadratic form involved in the definition of the Hilbert symbol, with an extra term abt^2. If there exists a non trivial solution $(x_0, y_0, z_0) \in (F_v)^3$ to $x^2 - ay^2 - bz^2 = 0$, one can consider the quaternion $\alpha_0 = x_0 + iy_0 + jz_0 \in \mathcal{A}_v$, which reduced norm is zero, by construction. Since the invertible elements in $\mathcal{A}_v$ are the ones with non zero reduced norm, we conclude that $\alpha_0 \neq 0$ is not invertible and $\mathcal{A}_v$ can't be a division ring (and so $\mathcal{A}$ does not ramify at v).

The converse is actually true, that is, any element in $\mathcal{A} \setminus \{0\}$ with reduced norm equal to zero gives a non trivial zero in $(F_v)^3$ to the quadratic form $x^2 - ay^2 - bz^2$. As a consequence, we obtain that the Hilbert symbol is non trivial exactly at the ramified places:

$$\left(\frac{a,b}{v}\right) = \begin{cases} 1 \text{ if } \mathcal{A} \text{ is split at } v \\ -1 \text{ if } \mathcal{A} \text{ is ramified at } v \end{cases}.$$

Hilbert reciprocity law states that the product $\prod_v (\frac{a,b}{v})$ over all places v of F is always equal to 1. Therefore, the set $\mathrm{Ram}(\mathcal{A}) = \{v \,|\, (\frac{a,b}{v}) = -1\}$ of ramified places has even cardinal.

Lemma B.4 (Hilbert reciprocity law, [41, 14.6.3]). *Let F be a number field and $a, b \in F^\times$. Then,*

$$\prod_v \left(\frac{a,b}{v}\right) = 1, \tag{4}$$

where the product is taken over all places v of F. In particular when F is totally real of even degree and $\mathcal{A}$ is totally definite, the same holds when the product is indexed over finite places of F.

This is a powerful result which sometimes makes us able to decide if the ramification at a place is impossible or must occur, without computing any Hilbert symbol. Finally, we state a formula for computing Hilbert symbols, in the particular case of our quaternion algebras $(\frac{a_m,-1}{F_m})$. We emphasize that the following formula does not hold for prime ideals above 2.

Lemma B.5 ([41, 12.4.10]). *Let F be a number field and $\mathcal{A} = (\frac{a,-1}{F})$. For any prime ideal $\mathfrak{p}$ of F such that $\mathfrak{p} \nmid (2)$, the Hilbert symbol of $\mathcal{A}$ at $\mathfrak{p}$ is given by*

$$\left(\frac{a,-1}{\mathfrak{p}}\right) = \left(\frac{-1}{\mathfrak{p}}\right)^{v_{\mathfrak{p}}(a)},$$

where $\left(\frac{-1}{\mathfrak{p}}\right) := \begin{cases} 1 \text{ if } -1 \text{ is a square in } (\mathcal{O}_F/\mathfrak{p})^\times \\ -1 \text{ otherwise} \end{cases}$ is the Legendre symbol of -1 at $\mathfrak{p}$ and $v_{\mathfrak{p}}(a) := \max\{e \in \mathbb{N} \mid a \in \mathfrak{p}^e\}$ is the $\mathfrak{p}$-adic valuation of a.

Algorithms. In [40], the authors gave deterministic polynomial time algorithms for computing Hilbert symbols, treating the case where $\mathfrak{p}$ is above 2 separately.

Lemma B.6 ([40, Theorem 6.1]). *Let F be a number field and let v be a place of F. There exists an algorithm to evaluate the Hilbert symbol $(\frac{a,b}{v})$ for $a, b \in F^\times$, that is deterministic polynomial time in the size of the inputs.*

Corollary B.7. *There exists an algorithm that given a quaternion algebra $\mathcal{A} = (\frac{a,-1}{F})$ and the prime factorization of $a \cdot \mathcal{O}_F$, computes $\mathrm{disc}_F(\mathcal{A})$. Moreover, this algorithm is deterministic and runs in polynomial time.*

Proof. According to Lemma B.5, it is enough to check if the prime ideals dividing $a \cdot \mathcal{O}_F$ ramify in $\mathcal{A}$, as well as the prime ideals above 2. The latters can be computed in polynomial time thanks to Lemma 2.5. For each prime ideal $\mathfrak{p}$ dividing either $a \cdot \mathcal{O}_F$ or $2 \cdot \mathcal{O}_F$, the Hilbert symbol $(\frac{a,-1}{\mathfrak{p}})$ is computed in deterministic polynomial time, using Lemma B.6. There are at most $2 \cdot [F : \mathbb{Q}]$ such ideals. □

B.2 Computing maximal orders

Relative norm for ideals in CM extensions. Here, we state some additionnal results and terminology regarding ideals in CM fields, that will be of use later. Let K/F be a CM field, and $\mathfrak{p}$ be a prime ideal of $\mathcal{O}_F$. Recall from [34, Chapter I, (8.3) and (9.1)] that $\mathfrak{p}\mathcal{O}_K$ factorizes in $\mathcal{O}_K$ either as

$$\mathfrak{p}\mathcal{O}_K = \begin{cases} \mathfrak{q}\bar{\mathfrak{q}} \text{ with } \mathfrak{q} \neq \bar{\mathfrak{q}} \text{ prime ideals (split case)} \\ \mathfrak{q}^2 \text{ with } \mathfrak{q} = \bar{\mathfrak{q}} \text{ prime ideal (ramified case)} \\ \mathfrak{q} \text{ with } \mathfrak{q} = \bar{\mathfrak{q}} \text{ prime ideal (inert case).} \end{cases} \tag{5}$$

In the split and ramified cases, we have $\mathfrak{q}\bar{\mathfrak{q}} \cap F = \mathfrak{p}\mathcal{O}_K \cap \mathcal{O}_F = \mathfrak{p}$ ([32, Chapter 3, Exercise 9 (c)]). For the inert case, $\mathfrak{q}\bar{\mathfrak{q}} \cap F = \mathfrak{p}^2\mathcal{O}_K \cap \mathcal{O}_F = \mathfrak{p}^2$. The relative norm

of a prime ideal $\mathfrak{q} \subset \mathcal{O}_K$ is then as $N_{K/F}(\mathfrak{q}) = \mathfrak{q}\overline{\mathfrak{q}} \cap F$. Thanks to the previous observation, this definition coincides with the one given in [34, Chapter III, §1]. The relative norm is then extended multiplicatively to the set of fractional ideals of K. In particular it is multiplicative, i.e., $N_{K/F}(\mathfrak{ab}) = N_{K/F}(\mathfrak{a})N_{K/F}(\mathfrak{b})$ holds. In fact $N_{K/F}(\mathfrak{a})$ is also equal to the ideal of F generated by $\{N_{K/F}(x) \,|\, x \in \mathfrak{a}\}$, see [34, Chapter III, (1.6)]. For a principal ideal $\mathfrak{a} = g \cdot \mathcal{O}_K$, we have $N_{K/F}(\mathfrak{a}) = N_{K/F}(g) \cdot \mathcal{O}_F$.

Discriminant of orders. The discriminant of an order $\mathcal{O}$ in a quaternion algebra $\mathcal{A}$ over F is the following ideal of $\mathcal{O}_F$:

$$\mathrm{disc}(\mathcal{O}) := \{\det(\mathrm{trd}(\alpha_i\alpha_j)_{1\leqslant i,j\leqslant 4}),\ \alpha_1, \ldots, \alpha_4 \in \mathcal{O}\} \cdot \mathcal{O}_F,$$

where $\mathrm{trd}(a) := a + \overline{a}$ is the reduced trace map on $\mathcal{A}$, and $\mathrm{trd}(a_ia_j)_{1\leqslant i,j\leqslant 4}$ is a 4×4 matrix with coefficients in F. Given a pseudo-basis $\mathcal{O} = \mathfrak{a}_1\alpha_1 \oplus \cdots \oplus \mathfrak{a}_4\alpha_4$, and according to [41, Corollary 15.2.7, Paragraph 15.2.8], we have

$$\mathrm{disc}(\mathcal{O}) = (\mathfrak{a}_1 \cdots \mathfrak{a}_4)^2 \cdot \det(\mathrm{trd}(\alpha_i\alpha_j)_{1\leqslant i,j\leqslant 4}) \cdot \mathcal{O}_F$$

In fact $\mathrm{disc}(\mathcal{O})$ is the square of an ideal of $\mathcal{O}_F$ (see [41, Section 15.4]) and we call reduced discriminant of $\mathcal{O}$ the ideal such that $\mathrm{discrd}(\mathcal{O})^2 = \mathrm{disc}(\mathcal{O})$. It somehow measures how far $\mathcal{O}$ is from being a maximal order, in the sense that it is a maximal order if and only if its (reduced) discriminant is equal to the one of $\mathcal{A}$.

Lemma B.8 ([41, Proposition 15.5.5]). *A quaternion order $\mathcal{O}$ in a quaternion algebra $\mathcal{A}$ is maximal if and only if* $\mathrm{discrd}(\mathcal{O}) = \mathrm{disc}(\mathcal{A})$.

Notice that relative discriminants in a CM (so quadratic) extension K/F are defined in the same fashion

$$\mathrm{disc}(\mathcal{O}) := \{\det(\mathrm{trd}(a_ia_j))_{1\leqslant i,j\leqslant 2},\ a_1, a_2 \in \mathcal{O}\} \cdot \mathcal{O}_F,$$

for any order $\mathcal{O} \subset \mathcal{O}_K$. For the maximal order $\mathcal{O} = \mathcal{O}_K$, we denote $\Delta_{K/F} := \mathrm{disc}(\mathcal{O}_K)$ the relative discriminant of K over F.

Example B.9. Consider $\mathcal{A} = (\frac{-1,-1}{\mathbb{Q}})$. According to [41, Example 15.5.7] we have $\mathrm{disc}(\mathcal{A}) = 2\mathbb{Z}$. The order $\mathcal{O} := \mathbb{Z} \oplus \mathbb{Z}i \oplus \mathbb{Z}j \oplus \mathbb{Z}k$ has basis $\{1, i, j, k\}$ and one computes $\mathrm{disc}(\mathcal{O}) = (\det \mathrm{diag}(2, -2, -2, -2)) \cdot \mathbb{Z}$ so $\mathrm{discrd}(\mathcal{O}) = 4\mathbb{Z}$ and $\mathcal{O}$ is not maximal. So this order is not maximal in $\mathcal{A}$. Now consider $\mathcal{O}' := \mathbb{Z} \oplus \mathbb{Z}i \oplus \mathbb{Z}j \oplus \mathbb{Z}\gamma$, where $\gamma = \left(\frac{1+i+j+k}{2}\right)$. Then one computes $\mathrm{disc}(\mathcal{O}') = (\det \mathrm{diag}(2, -2, -2, -1/2)) \cdot \mathbb{Z}$ so $\mathrm{discrd}(\mathcal{O}') = 2\mathbb{Z}$ and $\mathcal{O}'$ is thus maximal.

Algorithms. Before focusing on the case of the algebras $\mathcal{A}_m$, we give a generic procedure to compute a maximal order $\widetilde{O}$ containing some given order $\mathcal{O}$ in a quaternion algebra $\mathcal{A}$. As in the commutative case, the algorithm can be described iteratively. Given a prime ideal $\mathfrak{p}$ of $\mathcal{O}_F$, we say that $\mathcal{O}$ is $\mathfrak{p}$-maximal if $v_\mathfrak{p}(\mathrm{discrd}(\mathcal{O}))$ is minimal, *i.e.,* when $v_\mathfrak{p}(\mathrm{discrd}(\mathcal{O})) = v_\mathfrak{p}(\mathrm{disc}_F(\mathcal{A}))$ holds. Therefore, the maximal orders of $\mathcal{A}$ are precisely the orders which are $\mathfrak{p}$-maximal

for every prime ideal. It is enough to look at the prime ideals $\mathfrak{p}$ dividing $\mathrm{discrd}(\mathcal{O})$ (since $\mathfrak{q} \nmid \mathrm{discrd}(\mathcal{O})$ implies that $v_{\mathfrak{q}}(\mathrm{discrd}(\mathcal{O}))$ is already minimal). Once the factorization of $\mathrm{discrd}(\mathcal{O})$ is known, a $\mathfrak{p}$-maximal order containing $\mathcal{O}$ can be computed in deterministic polynomial-time. Repeating this step for each prime $\mathfrak{p} \mid \mathrm{discrd}(\mathcal{O})$ leads to a maximal order $\mathcal{O}$, as desired.

Lemma B.10 ([40, Algorithm 7.10]). *Let $\mathcal{O}$ and $\mathcal{A}$ be as above and let $\mathfrak{p}$ be a prime ideal of $\mathcal{O}_F$. There exists an algorithm that given as input a pseudo-basis of $\mathcal{O}$ and $\mathfrak{p}$, computes a pseudo-basis of a $\mathfrak{p}$-maximal order containing $\mathcal{O}$. It is deterministic and it runs in polynomial-time in $\mathrm{rank}_{\mathbb{Z}}(\mathcal{O}) = 4 \cdot [F : \mathbb{Q}]$ and in the size of $\mathcal{O}$.*

Remark B.11. The complexity of this algorithm is not mentionned in [40] however it is guaranteed to run in deterministic polynomial-time thanks to the following result.

Lemma B.12 ([40, Theorem 7.14]). *Let $\mathcal{O}$ and $\mathcal{A}$ be as above and let $\mathfrak{p}$ be a prime ideal of $\mathcal{O}_F$. There exists an algorithm that given as input a pseudo-basis of $\mathcal{O}$, computes a pseudo-basis of a maximal order $\widetilde{O} \supset \mathcal{O}$. It is deterministic polynomial-time reducible to the problem of factoring $\mathrm{discrd}(\mathcal{O})$ in $\mathcal{O}_F$.*

An explicit computation in cyclotomic fields. Let $K_m = F_m(a_m)$ be the m-th cyclotomic field with maximal totally real subfield F_m, and $\mathcal{A}_m$ be the quaternion algebra $(\frac{a_m,-1}{F_m})$ over F_m. We investigate the maximality of the order $\mathcal{O}_m = \mathcal{O}_{K_m} \oplus \mathcal{O}_{K_m} \cdot j \subset \mathcal{A}_m$, and we give a polynomial time algorithm for computing a maximal order containing it. In Corollary B.17, we prove that $\mathcal{O}_m$ is often maximal and always not far from being maximal, in the sense that $\mathrm{discrd}(\mathcal{O}_m)$ is either $\mathcal{O}_{F_m}$, a prime ideal $\mathfrak{p}$ of $\mathcal{O}_{F_m}$ or $\mathfrak{p}^2$. Since $\mathrm{discrd}(\mathcal{O}_m) \subset \mathrm{disc}_{F_m}(\mathcal{A}_m)$ holds (as for any order in $\mathcal{A}_m$) we get as a corollary the prime factorization of $\mathrm{disc}_{F_m}(\mathcal{A}_m)$. Once given the factorizations of $\mathrm{disc}_{F_m}(\mathcal{A}_m)$ and $\mathrm{discrd}(\mathcal{O}_m)$, we are then able to compute a maximal order containing $\mathcal{O}_m$ in polynomial time.

Lemma B.13 ([41, 15.2.12]). *We have the equality $\mathrm{disc}_{F_m}(\mathcal{O}_m) = \Delta^2_{K_m/F_m}$.*

Proof. Apply [41, 15.2.12] with the $\mathcal{O}_{F_m}$-order $S = \mathcal{O}_{K_m}$, whose discriminant relatively to $\mathcal{O}_{F_m}$ is by definition Δ_{K_m/F_m}. □

So, computing $\mathrm{disc}_{F_m}(\mathcal{O}_m)$ boils down to computing the factorization of Δ_{K_m/F_m}. This is done in two steps. First, we recall how this ideal can be built efficiently. Then, a property says that the prime ideals of $\mathcal{O}_{F_m}$ dividing Δ_{K_m/F_m} are the ones which ramify in $\mathcal{O}_{K_m}$ (this is in fact an equivalence, see [34, Chapter III, Corollary 2.12]). Ramification in cyclotomic CM-extensions is well-understood: Lemma B.15 recalls those ramified prime ideals. Additionally, the (relative) different ideal $\mathcal{D}_{K/F}$ is an ideal of $\mathcal{O}_K$ whose prime factors are exactly the primes of $\mathcal{O}_K$ over the ones in F that ramify. Morally, $\mathcal{D}_{K/F}$ encodes the ramification in K/F, as $\Delta_{K/F}$ does, but at the level of K. Below we recall how these ideals are linked.

Lemma B.14 ([34, Chap. 3, Prop. 2.4]). *Let $K = F(\alpha)/F$ be an extension of number fields and suppose that $\mathcal{O}_K = \mathcal{O}_F[\alpha]$. Then,*

$$\mathcal{D}_{K/F} = (T'(\alpha)) \cdot \mathcal{O}_K$$
$$\Delta_{K/F} = N_{K/F}(T'(\alpha)) \cdot \mathcal{O}_F,$$

where $T(X) \in \mathcal{O}_F[X]$ is the minimal polynomial of α over F.

In our case, $K_m = F_m(\zeta_m)$, $\mathcal{O}_{K_m} = \mathcal{O}_{F_m}[\zeta_m]$[20] and the minimal polynomial of ζ_m over F_m is $T(X) = X^2 - (\zeta_m + \zeta_m^{-1})X + 1$ so $\Delta_{K_m/F_m} = N_{K_m/F_m}(2\zeta_m - (\zeta_m + \zeta_m^{-1})) \cdot \mathcal{O}_{F_m} = N_{K_m/F_m}(\zeta_m - \zeta_m^{-1}) \cdot \mathcal{O}_{F_m} = (\zeta_m - \zeta_m^{-1})^2 \cdot \mathcal{O}_{F_m}$. Moreover, from the identity $\zeta_m^{-1} - \zeta_m = \zeta_m^{-1}(1 - \zeta_m)(1 + \zeta_m)$, we have $\mathcal{D}_{K_m/F_m} = (1 - \zeta_m)(1 + \zeta_m) \cdot \mathcal{O}_{K_m}$.

Lemma B.15 ([42, Proposition 2.15]). *If $m = p^e$ or $2p^e$ with p an odd prime, then K_m/F_m is ramified at the unique prime ideal above p and unramified everywhere else. In the other cases, K_m/F_m is unramified.*

Corollary B.16. *If $m = p^e$ or $2p^e$ with p an odd prime, then $\Delta_{K_m/F_m} = \mathfrak{p}$ where $\mathfrak{p} = (\zeta_m + \zeta_m^{-1} - 2)$ is the unique prime ideal above p. If $m = 2^e$ is a power of two (with $e > 2$), then $\Delta_{K_m/F_m} = \mathfrak{p}_2^2$ where $\mathfrak{p}_2 = (\zeta_m + \zeta_m^{-1})$ is the unique prime ideal above 2. Otherwise, $\Delta_{K_m/F_m} = \mathcal{O}_{F_m}$.*

Before proving this corollary, recall that, given a CM extension K_m/F_m, the relative norm of $a \in K_m$ over F_m is $N_{K_m/F_m}(a) = a\overline{a}$. The same notation N_{K_m/F_m} is used for the relative norm of ideals of K, as defined at the beginning of this subsection. The absolute norm of an ideal $\mathfrak{a} \subset K$ is the $\mathbb{Z}$-fractional ideal $N(\mathfrak{a})$ (equal to $|\mathcal{O}_{K_m}/\mathfrak{a}| \cdot \mathbb{Z}$ when $\mathfrak{a}$ is an integral ideal).

Proof. Thanks to [34, Chapter III, Corollary 2.3 and 2.12], the primes ideals of F_m dividing Δ_{K_m/F_m} are exactly the ramified primes in $\mathcal{O}_{K_m}$. So, by Lemma B.15, there are three cases to distinguish. If m is not a prime power, then no prime ideal ramifies so $\Delta_{K_m/F_m} = \mathcal{O}_{F_m}$. If $m = p^e$, then 2 is coprime to m and therefore $1 + \zeta_m = \frac{1-\zeta_m^2}{1-\zeta_m}$ is a cyclotomic unit, see *c.f.,* [42, §8.1]. Since K_m/F_m is ramified at the unique prime ideal above p by Lemma B.15, $1 - \zeta_m$ cannot also be a unit, and so we have $\mathcal{D}_{K_m/F_m} = (1 - \zeta_m) \cdot \mathcal{O}_{K_m}$ as the sole ideal above the prime $\mathfrak{p}$ in F_m that ramifies in K_m, and $\mathfrak{p} = (\zeta_m + \zeta_m^{-1} - 2) \cdot \mathcal{O}_{F_m}$ as claimed, by computing the relative norm of $1 - \zeta_m$. Note that the case where $m = 2p^e$ with p odd prime leads to the same result, since $K_m = K_{p^e}$.[21]

Now suppose that $p = 2$. Then both ζ_m and $-\zeta_m$ are primitive m-th roots of unity. In particular $N_{K_m/\mathbb{Q}}(1-\zeta_m) = N_{K_m/\mathbb{Q}}(1+\zeta_m)$. Using that $-\zeta_m = \zeta^{m/2+1}$, we have the identity $(1 - \zeta_m)\sum_{i=0}^{m/2} \zeta_m^i = 1 + \zeta_m$, so that $\sum_{i=0}^{m/2} \zeta_m^i \in \mathcal{O}_{K_m}$ has norm 1: it is a unit. Hence we have $(1-\zeta_m) \cdot \mathcal{O}_{K_m} = (1+\zeta_m) \cdot \mathcal{O}_{K_m}$. We compute

[20] In fact for cyclotomic rings of integers we have $\mathcal{O}_{K_m} = \mathbb{Z}[\zeta_m]$. But then $\mathcal{O}_{F_m}[\zeta_m]$ is a sub-order containing both $\mathbb{Z}$ and ζ_m, so we must have equality.

[21] Indeed, $K_{p^e} \subset K_m$ holds because $p^e \mid m$ and $\varphi(p^e) = \varphi(m)$ so the fields have same degree over $\mathbb{Q}$ and are thus equal.

Δ_{K_m/F_m} as $N_{K_m/F_m}((1-\zeta_m)^2) = (\zeta_m+\zeta_m^{-1}-2)^2$. To finish the proof, we must argue that $(\zeta_m+\zeta_m^{-1}-2)\cdot\mathcal{O}_{F_m}$ is in fact equal to $\mathfrak{p}_2$. For this, we use [2, Theorem 2.2] which implies that $N(\zeta_m+\zeta_m^{-1}) = 2$. Thus, $2 \in (\zeta_m+\zeta_m^{-1})\cdot\mathcal{O}_{F_m}$ and the inclusion $(\zeta_m+\zeta_m^{-1}-2) \subset \mathfrak{p}_2$ holds. But these two integral ideals have the same absolute norm, so they must be equal. □

Corollary B.17. *The following assertions hold:*

1. *If $m = 2^e$ (with $e > 2$) then $\mathrm{discrd}(\mathcal{O}_m) = \mathfrak{p}_2^2$, where $\mathfrak{p}_2 = (\zeta_m+\zeta_m^{-1})$ is the unique prime ideal above 2, whereas $\mathrm{disc}_{F_m}(\mathcal{A}_m) = \mathcal{O}_{F_m}$.*
2. *If $m = p^e$ or $2p^e$ with $p = 1 \pmod 4$ then $\mathrm{discrd}(\mathcal{O}_m) = \mathfrak{p}$, where $\mathfrak{p} = (\zeta_m+\zeta_m^{-1}-2)$ is the unique prime ideal above p, whereas $\mathrm{disc}_{F_m}(\mathcal{A}_m) = \mathcal{O}_{F_m}$.*
3. *If $m = p^e$ or $2p^e$ with $p = 3 \pmod 4$ then $\mathrm{discrd}(\mathcal{O}_m) = \mathfrak{p}$, where $\mathfrak{p} = (\zeta_m+\zeta_m^{-1}-2)$ is the unique prime ideal above p, whereas $\mathrm{disc}_{F_m}(\mathcal{A}_m) = \mathfrak{p}$. In particular, $\mathcal{O}_m$ is maximal.*
4. *Otherwise, $\mathrm{discrd}(\mathcal{O}_m) = \mathrm{disc}_{F_m}(\mathcal{A}_m) = \mathcal{O}_{F_m}$. In particular, $\mathcal{O}_m$ is maximal.*

Proof. In all cases we will use the inclusion of ideals $\mathrm{discrd}(\mathcal{O}_m) \subset \mathrm{disc}_{F_m}(\mathcal{A}_m) \subset \mathcal{O}_{F_m}$, so that any prime ideal dividing the second discriminant must also divide the first one.

1. If $m = 2^e$ and $e > 2$, then we have $\mathrm{disc}_{F_m}(\mathcal{O}_m) = \Delta^2_{K_m/F_m}$, by Lemma B.13 and $\Delta_{K_m/F_m} = \mathfrak{p}_2^2$ by Corollary B.16 so $\mathrm{discrd}(\mathcal{O}_m) = \mathfrak{p}_2^2$. There are $[F_m : \mathbb{Q}] = \varphi(m)/2 = 2^{e-2} \in 2\mathbb{Z}$ infinite places in F_m which all ramify in $\mathcal{A}_m$. Since $\mathrm{disc}_{F_m}(\mathcal{A}_m) \mid \mathrm{discrd}(\mathcal{O}_m)$, the unique finite place of F_m which can potentially ramify in $\mathcal{A}_m$ is $\mathfrak{p}_2$. But then by Hilbert reciprocity law (4),
$$1 = \underbrace{\prod_{v_\infty}\left(\frac{a_m,-1}{v_\infty}\right)}_{=(-1)^{deg(F_m)}=1}\cdot\prod_{\mathfrak{p}}\left(\frac{a_m,-1}{\mathfrak{p}}\right) = \left(\frac{a_m,-1}{\mathfrak{p}_2}\right)\cdot\prod_{\mathfrak{p}\nmid(2)}\underbrace{\left(\frac{a_m,-1}{\mathfrak{p}}\right)}_{=1} = \left(\frac{a_m,-1}{\mathfrak{p}_2}\right).$$
so $\mathcal{A}_m$ does not ramify at $\mathfrak{p}_2$ and $\mathrm{disc}_{F_m}(\mathcal{A}_m) = \mathcal{O}_{F_m}$.
2. If $m = p^e$ or $2p^e$ with $p = 1 \pmod 4$, then Corollary B.16 gives $\Delta_{K_m/F_m} = \mathfrak{p}$ so $\mathrm{discrd}(\mathcal{O}_m) = \mathfrak{p}$. There are $[F_m : \mathbb{Q}] = \varphi(m)/2 = (p-1)p^{e-1}/2 \in 2\mathbb{Z}$ infinite places in F_m which all ramify in $\mathcal{A}_m$. In the same way, the unique finite place of F_m which can potentially ramify in $\mathcal{A}_m$ is $\mathfrak{p}$. Again by Hilbert reciprocity law, $\mathcal{A}_m$ can't ramify at $\mathfrak{p}$ so $\mathrm{disc}_{F_m}(\mathcal{A}_m) = \mathcal{O}_{F_m}$.
3. If $m = p^e$ or $2p^e$ with $p = 3 \pmod 4$, then Corollary B.16 gives $\Delta_{K_m/F_m} = \mathfrak{p}$ so $\mathrm{discrd}(\mathcal{O}_m) = \mathfrak{p}$. There are $[F_m : \mathbb{Q}] = \varphi(m)/2 = (p-1)p^{e-1}/2 \in (2\mathbb{Z}+1)$ infinite places in F_m which all ramify in $\mathcal{A}_m$. In the same way, the unique finite place of F_m which can potentially ramify in $\mathcal{A}_m$ is $\mathfrak{p}$. Now Hilbert reciprocity law implies that $\mathcal{A}_m$ must ramify at $\mathfrak{p}$, so $\mathrm{disc}_{F_m}(\mathcal{A}_m) = \mathfrak{p}$. Finally, $\mathrm{discrd}(\mathcal{O}_m) = \mathrm{disc}_{F_m}(\mathcal{A}_m)$ so $\mathcal{O}_m$ is maximal, by Lemma B.8.
4. In all other cases, Corollary B.16 gives $\Delta_{K_m/F_m} = \mathcal{O}_{F_m}$ so $\mathrm{disc}_{F_m}(\mathcal{A}_m) = \mathrm{discrd}(\mathcal{O}_m) = \mathcal{O}_{F_m}$ and $\mathcal{O}_m$ is maximal, by Lemma B.8.

□

Algorithm 4: Computing a maximal order $\widetilde{\mathcal{O}_m} \supset \mathcal{O}_m = \mathcal{O}_{K_m} \oplus \mathcal{O}_{K_m} \cdot j$

Input: An integer $m \in \mathbb{N}_{>2}$ ($m \neq 4$), a primitive m-th root of unity ζ_m. $K_m = \mathbb{Q}(\zeta_m)$ (resp. $F_m = \mathbb{Q}(\zeta_m + \zeta_m^{-1})$). A pseudo-basis $(B, \{\mathfrak{a}, \mathfrak{b}\})$ of $\mathcal{O}_{K_m} = \mathbb{Z}[\zeta_m]$ over $\mathcal{O}_{F_m} = \mathbb{Z}[\zeta_m + \zeta_m^{-1}]$.

Output: A pseudo-basis over $\mathcal{O}_{F_m}$ of a maximal order containing $\mathcal{O}_m$.

1 Check if $m = 2^e$, p^e or $2p^e$ and if $p = 1$ or $3 \pmod 4$;

2 Compute (the prime factorization of) $\mathrm{disc}_{F_m}(\mathcal{A}_m)$ and $\mathrm{discrd}(\mathcal{O}_m)$ ▷ Thanks to Corollary B.17;

3 **if** $\mathrm{disc}_{F_m}(\mathcal{A}_m) = \mathrm{discrd}(\mathcal{O}_m)$ **then**

4 **return** $(\mathrm{diag}(B, B), \{\mathfrak{a}, \mathfrak{b}, \mathfrak{a}, \mathfrak{b}\})$

5 **else**

6 $\widetilde{\mathcal{O}_m} \leftarrow \mathfrak{p}$-maximal order containing $\mathcal{O}_m$ ▷ Using Lemma B.10;

7 **return** $(\widetilde{\mathcal{O}_m})$

Proposition B.18. *For $m \in \mathbb{N}_{>2}, m \neq 4$ and with the previous notations, Algorithm 4 computes (a pseudo-basis of) a maximal order $\widetilde{\mathcal{O}_m}$ of $\mathcal{A}_m$ containing the order $\mathcal{O}_m = \mathcal{O}_{K_m} \oplus \mathcal{O}_{K_m} \cdot j$. Moreover, it runs in polynomial time in the degree $d_m = \varphi(m) = [K_m : \mathbb{Q}]$.*

Proof. Correctness. If $\mathrm{disc}_{F_m}(\mathcal{A}_m) = \mathrm{discrd}(\mathcal{O}_m)$, then Lemma B.8 ensures that $\mathcal{O}_m$ is already maximal. Otherwise, Corollary B.17 tells us that we have $\mathrm{discrd}(\mathcal{O}_m) = \mathfrak{p}$ or $\mathfrak{p}^2$. In both cases, $v_{\mathfrak{p}}(\mathrm{disc}_{F_m}(\mathcal{A}_m)) = 0$ and $v_{\mathfrak{q}}(\mathrm{disc}_{F_m}(\mathcal{A}_m)) = v_{\mathfrak{q}}(\mathrm{discrd}(\mathcal{O}_m)) = 0$ for any prime $\mathfrak{q} \neq \mathfrak{p}$, so it is enough to build an order $\widetilde{\mathcal{O}_m} \supset \mathcal{O}_m$ which is $\mathfrak{p}$-maximal *i.e.,* such that $v_{\mathfrak{p}}(\mathrm{discrd}(\widetilde{\mathcal{O}_m}))$ is maximal. This is done in step 6, according to Lemma B.10.

Complexity. One can check if m is either of the form 2^e, p^e or $2p^e$, in polynomial time in m. Step 6 is achieved in polynomial time in $\mathrm{rank}_{\mathbb{Z}}(\mathcal{O}) = 2d$ and in $\mathrm{size}(\mathcal{O}) = \mathrm{poly}(d_m, \log \Delta_{K_m}) = \mathrm{poly}(d_m)$, as $\log \Delta_{K_m} = \mathrm{poly}(d_m)$ holds for cyclotomic fields (see [42, Proposition 2.1]). □

B.3 Units of reduced norm 1 of an order

Recall that our setting is a CM extension K/F of number fields, and a totally definite quaternion algebra $\mathcal{A} = (\frac{a,-1}{F})$, where a is such that $K = F(\sqrt{a})$. Let $\mathcal{A}^\times$ resp. $\mathcal{A}^1$ is the set of elements with non-zero reduced norm (equivalently, invertible), resp. reduced norm equal to 1. For any order $\mathcal{O}$ in $\mathcal{A}$, we let $\mathcal{O}^\times = \mathcal{A}^\times \cap \mathcal{O}$ (so, the units of $\mathcal{O}$) and $\mathcal{O}^1 = \mathcal{A}^1 \cap \mathcal{O}$. Lastly, we let $\mathcal{O}_K^\times, \mathcal{O}_K^1$ and $\mathcal{O}_F^\times, \mathcal{O}_F^1$ the intersection of $\mathcal{A}^\times$ resp. $\mathcal{A}^1$ with $\mathcal{O}_K$, resp. $\mathcal{O}_F$.

We now precise the structure of $\mathcal{O}^\times$ and $\mathcal{O}^1$.

Proposition B.19 ([41, Proposition 32.3.7]). *Let $\mathcal{O}$ be a maximal order of a definite quaternion algebra $\mathcal{A}$. Then $\mathcal{O}^1$ is a finite group.*

We are mostly concerned with the possible size of $\mathcal{O}^1$, and the goal of this section is to show that it remains small. The structure of $\mathcal{O}^1$ can sometimes

be elucidated, but following [41, Chap. 32], it is easier to understand working modulo signs. Let $P\mathcal{A}^\times := \mathcal{A}^\times/F^\times$. In the totally definite case, $\mathcal{O}^1/\{\pm 1\}$ is not only a finite subgroup of $P\mathcal{A}^1$, but its structure is also known to some extent.

A dihedral group D_m can be understood as the group of symmetry of a regular polygon with m vertices, and thus is generated by a reflexion τ and a cyclic permutation σ of order m. It is non commutative when $m > 2$, as we have $\tau\sigma\tau = \sigma^{-1}$. Recall that S_n is the group of all the permutations of n symbols, and A_n is its subgroup of even permutations.

Proposition B.20 ([41, Proposition 32.4.1]). *The finite subgroups of $P\mathcal{A}^\times$ are cyclic, dihedral, or isomorphic to a permutation group A_4, S_4, A_5. In particular, the group $\mathcal{O}^1/\{\pm 1\}$ is of this form.*

Finite groups of $P\mathcal{A}^\times$ isomorphic to a permutation group are called exceptional. Their size is constant (respectively $12, 24$ and 60), and particularly independent of the CM extension K/F. We will show that for many (and the most interesting) cases, $\mathcal{O}^1/\{\pm 1\}$ will not be an exceptional group.

There are known characterizations and even descriptions (up to isomorphism) of each of the possible situations above, proved in [41, Proposition 32.5.1, 32.5.5, 32.5.8, 32.6.6, 32.7.1]. We separate the exceptional and non-exceptional cases for clarity.

Proposition B.21 (Characterizations of non-exceptional groups).

- *$P\mathcal{A}^\times$ contains a cyclic subgroup Γ of order $m > 2$ if and only there exists a primitive m-th root of unity ζ_m in an algebraic closure of F, such that $\zeta_m + \zeta_m^{-1} \in F$ and $F(\zeta_m) = K$.*[22]
- *$P\mathcal{A}^1$ contains a cyclic subgroup of order m if and only if $P\mathcal{A}^\times$ contains one of order $2m$. In this case, it contains $\langle \zeta_{2m} \rangle$, of order m.*
- *$P\mathcal{A}^1$ contains a dihedral group of order $2m > 4$ if and only if, with the notation ζ_m as above, we have $K = F(1 + \zeta_m)$.*

When there exists a cyclic group Γ in $P\mathcal{A}^\times$, then it is conjugated to $\langle 1 + \zeta_m \rangle$, the group generated by $1 + \zeta_m$, by an element of $\mathcal{A}^\times$.

Proposition B.22 (Characterizations of exceptional groups). *The group $P\mathcal{A}^1$ contains a subgroup isomorphic to:*

- *A_4 if and only if $a^2 = -1$;*
- *S_4 if and only if $a^2 = -1$ and $\sqrt{2} \in F$;*
- *A_5 if and only if $a^2 = -1$ and $\sqrt{5} \in F$.*

Any such subgroups are isomorphic if and only if they are conjugated by an element of $\mathcal{A}^\times$.

[22] In [41, Proposition 32.5.1] the condition that K splits $\mathcal{A}$ is needed. The latter is in fact automatic for us, thanks to [41, Proposition 2.3.1].

These exceptional characterizations can be understood informally by the presence of $\frac{1}{\sqrt{2}}(1 \pm \epsilon)$, of order 4 (modulo sign) and $\frac{1}{\sqrt{2}}(\epsilon \pm \epsilon')$ of order 2 (modulo sign) when $\sqrt{2} \in F$, for ϵ, ϵ' distinct in $\{i, j, k\}$. Algebraically, one then works out the structure of S_4, or identifies these quaternions to symmetries of regular polygons. In the typical usecase where K is a power-of-two cyclotomic field, these exceptional groups appear in $P\mathcal{A}^1$. The case of $\sqrt{5}$ involves the golden ratio and can also be worked out similarly, see [41, Chap. 11].

While copies of all these well-identified groups can be explicitely written out in $P\mathcal{A}^\times$, without the knowledge of the conjugating element $\delta \in \mathcal{A}^\times$, we only know them "up to isomorphism" and cannot explicitely compute with them. We now characterize the elements of norm 1 in $\mathcal{O}_K + \mathcal{O}_K \cdot j$. Recall that $\mu(K)$ is the group of roots of unity in the number field K, which is cyclic [34, 7.4].

Corollary B.23 (Corollary of 2.3). *Let $\mathcal{O}_0 := \mathcal{O}_K + \mathcal{O}_K \cdot j$. We have $\mathcal{O}_0^1 = \langle j, \mu(K) \rangle$, that is, $\mathcal{O}_0^1$ is the group generated by j and $\mu(K)$*

Proof. Let $x = a + bj \in \mathcal{O}_0$. We have $x \in \mathcal{O}_0^1$ if and only if $\mathrm{nrd}(x) = a\overline{a} + b\overline{b} = 1$. Corollary 2.3 gives the solutions. □

This tells us that $\mathcal{O}_0^1/\{\pm 1\}$ is a dihedral group of size at least $|\mu(K)|$. When $\mu(K)$ is large enough, $\mathcal{O}^1/\{\pm 1\}$ then cannot be exceptional. We sum-up these observations in the next proposition.

Proposition B.24. *Let $\mathcal{O}$ be a maximal order containing $\mathcal{O}_0$ and $d = [F : \mathbb{Q}]$. If $|\mu(K)| \geq 61$, then $\mathcal{O}^1/\{\pm 1\}$ is dihedral and $\mathcal{O}^1$ has at most $16d^2$ elements.*

Proof. By inclusion, we have $\mathcal{O}^1 \supset \langle j, \mu(K) \rangle = \mathcal{O}_0^1$. Let G, G_0 be respectively $\mathcal{O}^1/\{\pm 1\}$ and $\mathcal{O}_0^1/\{\pm 1\}$. Because $|G_0| > 60$, neither G or G_0 can be any of the exceptional groups, thus they are cyclic or dihedral. In any of this cases, the cyclic component of G, generated by γ (say), contains the cyclic component $\mu(K)/\{\pm 1\}$ of G_0 generated by ζ. This means that γ commutes with ζ, and that $\pm\gamma^k = \pm\zeta$, for some integer $k \geq 1$. By cardinality of $\mu(K)$, we also see that $\gamma \neq -1$ and therefore $\gamma \notin F$. Now, ζ or $-\zeta$ is a primitive root of 1 in $K \setminus F$, so we have $F \subsetneq F(\zeta) \subset K$. Because K is quadratic over F, this means that we have $K = F(\zeta) \subset F(\gamma) \subset \mathcal{A}$. Since all elements in $\mathcal{A}$ have degree at most 2 over F, with minimal polynomial $T^2 - (\gamma + \overline{\gamma})T + \mathrm{nrd}(\gamma)$, $F(\gamma)$ has degree 2 over F and thus actually $F(\gamma) = K$. We deduce that γ and $-\gamma$ are roots of unity in K, and one (or both) of them any generator of $\mu(K)$. The conclusion comes from Lemma 2.4. □

A more general version of this proposition is as follows:

Proposition B.25. *Let K be a CM field, such that $K = F(\sqrt{a})$ is a quadratic extension of a totally real field F of degree $d = [F : \mathbb{Q}]$. Let $\mathcal{O}'$ be an order in $\mathcal{A} = (\frac{K,-1}{F})$. If $d > 2$, then $\mathcal{O}'^1$ has at most $16d^2$ elements.*

Proof. $\mathcal{O}'^1/\{\pm 1\}$ is a finite subgroup of $P\mathcal{A}^1$. According to Proposition B.21, and [41, Proposition 32.7.1], the finites subgroup of $P\mathcal{A}^1$ are either dihedral, cyclic, or conjugated to an exceptionnal subgroup A_4,A_5 or S_4. If $\mathcal{O}'^1/\{\pm 1\}$ falls in the latter case, considering the size of each of these groups, this means that $|\mathcal{O}'^1/\{\pm 1\}| \leqslant 60$.

Suppose now that $\mathcal{O}'^1/\{\pm 1\}$ is cyclic of order m. Then by Proposition B.21, it is conjugated to the group generated by ζ_{2m}, where ζ_{2m} is a $2m$-th root of unity in $\mathcal{A}$ (so a m-th root of -1) such that $\mathcal{A} = (\frac{F(\zeta_{2m}),-1}{F})$. Again by Proposition B.21, $\zeta_{2m} + \zeta_{2m}^{-1} \in F$, so the minimal polynomial of ζ_{2m} in $F[T]$ is $T^2 - (\zeta_{2m} + \zeta_{2m}^{-1})T + 1$. This polynomial is of degree 2, and so $[F(\zeta_{2m}) : F] = 2$. We know, according to Lemma 2.4, that $\mu(F)$ is a cyclic group of order $\leqslant 2d^2$, so $\mu(F(\zeta_{2m}))$ is a cyclic group of order $\leqslant 8d^2$, so $2m$ is at most equal to $8d^2$. To sum up, in this case, we have $|\mathcal{O}'^1/\{\pm 1\}| \leqslant 4d^2$, and $|\mathcal{O}'^1| \leqslant 8d^2$.

Finally, if $\mathcal{O}'^1/\{\pm 1\}$ is dihedral of order $2m > 4$, then by Proposition B.21, it contains a cyclic subgroup of order m. As per the same argument as above, $m \leqslant 4d^2$, $|\mathcal{O}'^1/\{\pm 1\}| \leqslant 8d^2$, and $|\mathcal{O}'^1| \leqslant 16d^2$.

Since we assumed in the Proposition that $d > 2$, we have $16d^2 > 120$, and so, to sum up, $|\mathcal{O}'^1| \leqslant 16d^2$. □

We can show that the automorphism group of $\mathcal{O}_K^2$ is the semi-direct product of $\mu(K)$ and the quaternions of norm 1 in the natural order above.

Proposition B.26. *Let F be totally real subfield with CM-extension $K = F(\sqrt{a})$, and $\mathcal{O}_0 = \mathcal{O}_K + j\mathcal{O}_K$ in $\mathcal{A} = (\frac{a,-1}{F})$. We have a split short exact sequence*

$$1 \longrightarrow \mathcal{O}_0^1 \longrightarrow Aut(\mathcal{O}_K^2) \longrightarrow \mu(K) \longrightarrow 1,$$

where the second map is the restriction of the left regular representation $\rho : \mathcal{A} \to M_2(K)$ to $\mathcal{O}_0^1$, and the third map is the restriction of the determinant to $Aut(\mathcal{O}_K^2)$.

Proof. As seen in the proof of Proposition 3.13, a direct calculation shows that matrices in $\mathrm{Aut}(\mathcal{O}_K^2)$ are either diagonal or anti-diagonal with entries in $\mu(K)$, thanks to Lemma 2.2. Looking at the group homomorphism $\det : \mathrm{Aut}(\mathcal{O}_K^2) \to \mu(K)$, we see from this description that $H := \ker \det$ is the normal subgroup where the diagonal is $x, \overline{x}$, or the antidiagonal is $-\overline{x}, x$. We now consider the left regular representation [41, Sec. 2.3.8] ρ of $\mathcal{A}$, which maps $a + jb$ to the matrix

$$\rho(x + jy) = \begin{pmatrix} x & -\overline{y} \\ y & \overline{x} \end{pmatrix}.$$

This is an homomorphism of F-algebra, and a faithful representation, so it is also an injective group homomorphism when restricted to O_0^1, it is in particular injective. Recall that $\mu(K)$ is cyclic. By examination, we see that $\rho(O_0^1) = H$, so if we let ζ be a generator of $\mu(K)$, Corollary B.23 tells us that H is generated by

$$\rho(j) = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \text{ and } \rho(\zeta) = \begin{pmatrix} \zeta & 0 \\ 0 & \overline{\zeta} \end{pmatrix}.$$

Now let C be the set of matrices in $\mathrm{Aut}(\mathcal{O}_K^2)$ which are diagonal with $a, 1$ as entries, for $a \in \mu(K)$. This is a cyclic subgroup, isomorphic to $\mu(K)$. We may always write $\mathrm{diag}(\zeta^i, \zeta^j) = \mathrm{diag}(\zeta^{i+j}, 1) \cdot \mathrm{diag}(\zeta^{-j}, \zeta^j)$, and similarly for the antidiagonal matrices, using $\rho(j)$. In other words, we have $\mathrm{Aut}(\mathcal{O}_K^2) = H \cdot C$. We also readily check that $C \cap H = \{Id\}$, which gives us that $\mathrm{Aut}(\mathcal{O}_K^2) = C \ltimes H$ by characterization of semi-direct products of groups. The result follows from the equivalence between split short exact sequences of groups and semi-direct products. □

We note that the product is certainly not direct, as C acts non trivially by conjugation over H, since $j\zeta j = -\overline{\zeta}$. A direct corollary is to recover Proposition 3.13 by inspection of the cardinality of the groups above.

Quaternion algebra over cyclotomic fields. The special case of cyclotomic CM extensions can be made explicit for large conductors, so we isolate its formulation for the sake of clarity and reusability. Recall the notation K_m for the cyclotomic field $\mathbb{Q}(\zeta_m)$, with maximal totally real subfield F_m. Denote by $\mathcal{A}_m$ the quaternion algebra $K_m + K_m \cdot j$ over F_m, with order $\mathcal{O}_m = \mathcal{O}_{K_m} + \mathcal{O}_{K_m} \cdot j$. Finally, $\widetilde{\mathcal{O}_m}$ denotes a maximal order containing $\mathcal{O}_m$. The following result explicits $\widetilde{\mathcal{O}_m}$ in all but one cases.

Corollary B.27. *Let $m \geq 2$ be an integer. If m is of the form $m = 2^e$ or $m = p^e$ or $2p^e$ with $p = 1 \pmod 4$ prime, suppose furthermore that $m \geq 31$. Then,*

$$\widetilde{\mathcal{O}_m}^1 = \mathcal{O}_m^1 = \langle \pm\zeta_m, j \rangle.$$