Short Bases of Lattices over Number Fields

Claus Fieker[1] and Damien Stehlé[1,2]

[1] Magma Computer Algebra Group, School of Mathematics and Statistics, University of Sydney, NSW 2006, Australia.
[2] CNRS and Macquarie University.
claus.fieker@sydney.edu.au, damien.stehle@gmail.com

Abstract. Lattices over number fields arise from a variety of sources in algorithmic algebra and more recently cryptography. Similar to the classical case of $\mathbb{Z}$-lattices, the choice of a nice, "short" (pseudo)-basis is important in many applications. In this article, we provide the first algorithm that computes such a "short" (pseudo)-basis. We utilize the LLL algorithm for $\mathbb{Z}$-lattices together with the Bosma-Pohst-Cohen Hermite Normal Form and some size reduction technique to find a pseudo-basis where each basis vector belongs to the lattice and the product of the norms of the basis vectors is bounded by the lattice determinant, up to a multiplicative factor that is a field invariant. As it runs in polynomial time, this provides an effective variant of Minkowski's second theorem for lattices over number fields.

1 Introduction

Let K be a number field and $\mathcal{O}_K$ be its maximal order. An $\mathcal{O}_K$-module M is a finitely generated set of elements which is closed under addition and multiplication by elements in $\mathcal{O}_K$. Frequently, we have $M \subseteq K^m$ for some m. In the case of K being $\mathbb{Q}$, we have $\mathcal{O}_K = \mathbb{Z}$, thus $\mathcal{O}_K$-modules are just the classical $\mathbb{Z}$-lattices. Since $\mathbb{Z}$ is a principal ideal domain, every (torsion free) module is free, thus there exists a basis $b_1, \ldots, b_n \in M$ for some $n \leq m$ such that $M = \oplus_{i \leq n} \mathbb{Z} b_i$. Any two bases $(b_i)_i$ and $(c_i)_i$ have the same cardinality and are linked by some unimodular matrix $T \in \mathrm{GL}(n, \mathbb{Z})$. The choice of a *good* basis is crucial for almost all computational problems attached to M. Generally one tries to find a basis whose vectors have short Euclidean norms, using, for example, the LLL algorithm [15].

Replacing $\mathbb{Z}$ by the maximal order $\mathcal{O}_K$ makes the classification more complicated since $\mathcal{O}_K$ may no longer be a principal ideal domain. However, since $\mathcal{O}_K$ is still a Dedekind domain, the modules $M \subseteq K^m$ have a well known structure ([7, Cor. 1.2.25], [23, Th. 81:3]): there exist linearly independent elements $\mathbf{b}_1, \ldots, \mathbf{b}_n \in K^m$ and (non-zero fractional) ideals $\mathfrak{b}_1, \ldots, \mathfrak{b}_n$ such that $M = \oplus_{i \leq n} \mathfrak{b}_i \mathbf{b}_i$, i.e., every $\mathbf{b} \in M$ has a unique representation as $\mathbf{b} = \sum_{i \leq n} x_i \mathbf{b}_i$ with $x_i \in \mathfrak{b}_i$ for all $i \leq n$. Such a representation is commonly called a *pseudo-basis*. It should be noted that $\mathbf{b}_i$ may not belong to M, and in fact $\mathbf{b}_i \in M$ if and only if $1 \in \mathfrak{b}_i$. Similarly to the case of $\mathbb{Z}$-lattices, different pseudo-bases share the same cardinality, and it is known how to move from a pseudo-basis to another.

As for $\mathbb{Z}$-lattices, the choice of the pseudo-basis is of utmost importance. However, a key difference is that no analogue of LLL is known, as repeatedly noted in [7]. There have been attempts [10, 22, 11] but the algorithms are either limited to certain fields or give no guaranteed bounds on the output size. While every $\mathcal{O}_K$-module is also a $\mathbb{Z}$-lattice and can thus be analyzed with all the tools available over $\mathbb{Z}$, for many applications the additional structure as an $\mathcal{O}_K$-module is important. This structure is typically lost when applying techniques over $\mathbb{Z}$.

Originally, $\mathcal{O}_K$-modules mainly came from the study of finite extensions of K but now they occur in a wider range of problems from group theory (matrix groups and representations [9]) to applications in geometry (automorphism algebras of Abelian varieties). $\mathcal{O}_K$-modules also occur in lattice-based cryptography [17, 19, 24–26], and in that context the module rank n is usually poly-logarithmic in the degree of the number field. Cryptography based on $\mathcal{O}_K$-modules is increasingly popular, as on one side they lead to compact representations and to fast operations, and on the other side they enjoy a worst-case to average-case reduction for variants of the shortest vector problem, which allows the cryptographic security to be based on worst-case hardness assumptions.

As diverse as the applications are the requirements: only one (or more) short module element(s) may be needed, or a short (pseudo)-basis may be required, some applications rely on canonical representations, while any representation may suffice for others. We note that canonical representations tend to have components that are much larger than short representations as obtained by lattice reduction or our techniques. To find one short element it suffices to consider the underlying $\mathbb{Z}$-module (of dimension nd with $d = [K : \mathbb{Q}]$). For $\mathbb{Z}$-lattices contained in $\mathbb{Q}^m$, a canonical representation is the Hermite Normal Form (HNF). It has been generalized (BPC-HNF) to $\mathcal{O}_K$-modules contained in K^m by Bosma and Pohst [4] and Cohen [7, Chap. 1.4] (see also [12]).

Our results. In the present work, we describe an algorithm that computes a pseudo-basis made of short vectors. Given an arbitrary pseudo-basis $[(\mathbf{a}_i)_i, (\mathfrak{a}_i)_i]$ of a module $M \subseteq K^m$, it returns a pseudo-basis $[(\mathbf{b}_i)_i, (\mathfrak{b}_i)_i]$ such that:

$$\forall i \leq n: \quad \mathbf{b}_i \in M,\ \mathcal{N}(\mathfrak{b}_i) \in [2^{-O(d^2)}, 1] \ \text{ and } \|\mathbf{b}_i\| \leq 2^{O(dn)} \lambda_i(M),$$

where the $O(\cdot)$'s depend only on the field K and the choice of a given LLL-reduced integral basis, the euclidean norm $\|\cdot\|$ is a module extension of the T_2-norm over K, and the $\lambda_i(M)$'s correspond to the module minima. We refer to Corollary 1 for a precise statement. Overall, this provides a module equivalent to LLL-reduced bases of $\mathbb{Z}$-lattices in the sense that the vectors cannot be arbitrarily longer than the minima. Since it runs in polynomial time, it can also be interpreted as an effective approximate variant of the adaptation to $\mathcal{O}_K$-modules of Minkowski's second theorem (given in Theorem 2). We also study the representation of one-dimensional $\mathcal{O}_K$-modules, i.e., modules that are isomorphic to ideals of $\mathcal{O}_K$. We show how to modify Belabas' 2-element representation algorithm [2, Alg. 6.15] so that the output is provably small. Combining the latter and our module pseudo-reduction algorithm leads to compact representations of $\mathcal{O}_K$-modules.

The most natural approach to obtain reduced pseudo-bases consists in trying to generalize LLL, but as mentioned earlier all previous attempts have only partially succeeded. In contrast, we start by viewing the $\mathcal{O}_K$-module as a high-dimensional $\mathbb{Z}$-lattice. We find short module elements by applying LLL to a basis of the latter lattice and interpreting the output as module elements. At this point, we have a pseudo-basis (the input) and a full-rank set of short module vectors (produced by LLL). If we had a $\mathbb{Z}$-lattice instead of an $\mathcal{O}_K$-module, we would then use a technique common in the lattice-based cryptography community (see, e.g., [20, Le. 7.1]), consisting in using the HNF to convert a full rank set of short lattice vectors to a short basis. We adapt this technique to number fields, using the BPC-HNF and introducing a size-reduction algorithm for pseudo-bases.

Let us compare (pseudo-)LLL-reduced and BPC-HNF pseudo-bases. A theoretical advantage of the LLL approach is that it is not restricted to K^m but also works in a continuous extension (similarly to LLL-reduction being well-defined for real lattices). It should also be significantly more efficient to work with pseudo-bases made of short vectors because smaller integers and polynomials of smaller degrees are involved. On the other side, (pseudo-)LLL-reduced pseudo-bases are far from being unique, and seem more expensive to obtain.

Road-map. In Section 2, we give some reminders and elementary results on lattices, number fields and modules. In Section 3, we modify Belabas' 2-element representation algorithm for ideals of $\mathcal{O}_K$, as described above. We then give our module reduction algorithm in Section 4. Finally, in Section 5 we describe our implementation and give some examples.

Implementation. The algorithms have been implemented in the Magma computer algebra system [3, 18] and are available on request. They will be part of upcoming releases.

2 Preliminaries

We assume the reader is familiar with the geometry of numbers and algebraic number theory. We refer to [16, 20], [5, 21] and [7, Chap. 1] for introductions to the computational aspects of lattices, elementary algebraic number theory and to modules over Dedekind domains, respectively.

2.1 Lattices

In this work, we will call any finitely generated free $\mathbb{Z}$-module L a lattice. A usual lattice corresponds to the case where L is a discrete additive subgroup of $\mathbb{R}^n$ for some n. Any lattice can be written $L = \oplus_{i \leq d} \mathbb{Z} b_i$. If the b_i's are $\mathbb{Z}$-free, they are called a basis of L. A given lattice may have infinitely many bases but their cardinality d is constant and called rank. Any two bases are related by a unimodular transformation, i.e., one is obtained from the other by multiplying by a matrix in $\mathbb{Z}^{d \times d}$ of determinant ± 1.

If $L \subseteq \mathbb{Q}^n$ is of rank d, then there exists a basis $B = (\mathbf{b}_i)_i \in \mathbb{Q}^{n \times d}$ of L such that $\mu_j = \min\{i : B_{i,j} \neq 0\}$ (strictly) increases with j, and for all $j > k$ we

have $B_{\mu_j,j} > B_{\mu_j,k} \geq 0$. If $d = n$, this means that B is a row-wise diagonally strictly dominant lower triangular matrix and that its entries are non-negative. This basis is unique and called the Hermite Normal Form (HNF) of L. It can be computed in polynomial time from any basis [13].

In order to quantify the smallness of an element of a lattice L, we associate to L a positive definite bilinear form $q : L_{\mathbb{R}} \times L_{\mathbb{R}} \mapsto \mathbb{R}$. We use it to map a basis $(b_i)_i$ to its Gram matrix $G_q(b_1, \ldots, b_d) := (q(b_i, b_j))_{i,j}$. We denote $\sqrt{q(b,b)}$ by $\|b\|_q$, and may omit the subscript if it is clear from the context. The determinant of L, defined as $\det_q(L) = \det(G_q(b_1, \ldots, b_d))^{1/2}$, does not depend on the particular choice of the basis of L. Note that if $L \subseteq \mathbb{R}^n$ and q is the euclidean inner product, then $\det(L)$ is the d-dimensional volume of the parallelepiped $\{\sum_i y_i b_i : y_i \in [0,1]\}$. We define the lattice minima as follows:

$$\forall i \leq d,\ \lambda_{i,q}(L) = \min\{r : \exists c_1, \ldots, c_i \in L \text{ free},\ \max_{k \leq i} \|c_k\|_q \leq r\}.$$

Minkowski's second theorem states that $\prod_{i \leq d} \lambda_{i,q}(L) \leq \sqrt{d}^d \det_q(L)$. Frequently one tries to represent a lattice L by a basis that approximates the minima. In this article, we assume that we have an algorithm `LatRed` that takes as input an arbitrary basis of L and returns a reduced basis satisfying $\|b_i\| \leq \gamma \lambda_i(L)$, for all $i \leq d$. For example, if we use the LLL algorithm [15], then we can take $\gamma = 2^{d/2}$. We proceed as follows: compute the Gram matrix G of the input basis; use the Gram matrix LLL algorithm (see, e.g., [5, p. 88]), to find U unimodular such that $U^t G U$ is reduced; apply U to the input lattice basis. If the arithmetic over L is efficient, and if q can be efficiently computed or approximated with high accuracy, then this provides an efficient algorithm. Apart from being well-defined for more general lattices (not only for lattices on a rational vector space), a significant advantage of the LLL-reduction over the HNF is that it provides small lattice elements. However, it seems more expensive to obtain and the uniqueness of the representation is lost. Taking the HKZ-reduction instead of the LLL-reduction allows one to take $\gamma = 1/2\sqrt{d+3}$ (see [14]), but the complexity of the best algorithm for computing it [1] is exponential in d.

Let $(b_i)_{i \leq d}$ be a lattice basis. For any $i > j$, we define $\mu_{i,j} = q(b_i, b_j^*)/q(b_j^*, b_j^*)$, where $b_i^* = \text{argmin} \|b_i + \sum_{j<i} \mathbb{R} b_j\|$ thus $\|b_i^*\| = \min\{\|b_i + x\| : x \in \sum_{j<i} \mathbb{R} b_j\}$. We call the $\mu_{i,j}$'s and the b_i^*'s the Gram-Schmidt orthogonalisation (GSO) of the b_i's. If the b_i's are LLL-reduced, then $\|b_i^*\| \geq 2^{-d/2} \|b_i\|$ for all i. In the following, we will assume that `LatRed`-reduced bases also satisfy this property. Size-reduction of a vector $b \in \sum_{i \leq d} \mathbb{R} b_i$ with respect to $(b_i)_{i \leq j}$ consists in subtracting from b integer multiples of these b_i's so that the magnitudes of the first j coordinates of the output vector c when written as a linear combination of all the b_i^*'s belong to $[-1/2, 1/2)$. The latter uniquely defines c, and if $j = d$ we have $\|c\|^2 \leq \sum_{i \leq d} \|b_i^*\|^2 \leq d \max_{i \leq d} \|b_i\|^2$. We call size-reduction of the basis $(b_i)_i$ the process of size-reducing each b_i with respect to the previous b_j's for increasing i. The output remains a basis of the lattice spanned by the b_i's.

A standard technique in the lattice-based cryptography community (see, e.g., [20, Le. 7.1]) allows one to derive a short lattice basis from an arbitrary basis $(a_i)_i$ and a full-rank free set of short lattice vectors $(s_i)_i$. As we will adapt this tech-

nique to modules, we describe it briefly. Since the s_i's belong to the lattice, there exists $T \in \mathbb{Z}^{d\times d}$ such that $(s_i)_i = (a_i)_i \cdot T$. We compute the HNF T'^t of T^t: $T'^t = T^t(U^{-1})^t$ with U unimodular. We thus have $(s_i)_i = (b_i)_i \cdot T'$ where $(b_i)_i := (a_i)_i \cdot U$ is a lattice basis and T' is upper triangular with diagonal entries ≥ 1. The shape of T' implies that for any i we have $\|b_i^*\| \leq \|s_i^*\|$. Size-reducing the basis $(b_i)_i$ leads to a basis $(b'_i)_i$ such that $\max\|b'_i\| \leq \sqrt{d}\max\|s_i^*\| \leq \sqrt{d}\max\|s_i\|$. It can be checked that if $L \subseteq \mathbb{Q}^n$, then all the computations may be performed in polynomial time.

2.2 Number fields

Let K be a number field of degree d, with real and complex embeddings $(\theta_i)_{i\leq s_1}$, $(\theta_i)_{s_1<i\leq s_1+2s_2}$. Its maximal order $\mathcal{O}_K$ is a lattice: there exists a free set $(r_i)_i \in \mathcal{O}_K^d$ such that $\mathcal{O}_K = \oplus_i \mathbb{Z}r_i$. The r_i's form an integral basis of K, and we have $K = \mathcal{O}_K \otimes \mathbb{Q}$. We define $K_\mathbb{R} = K \otimes \mathbb{R}$, which is isomorphic (as rings) to $\mathbb{R}^{s_1} \times \mathbb{C}^{s_2}$, and extend the θ_i's to $K_\mathbb{R}$. Many quadratic forms may be associated to $K_\mathbb{R}$, but the most natural one derives from $q(x, x') = T_2(x, x') := \sum \theta_i(x)\bar{\theta}_i(x')$. The discriminant of K is defined as $\Delta_K = \det^2_{T_2}(\mathcal{O}_K)$. Note that for any $x, x' \in K_\mathbb{R}$, we have $\|xx'\| \leq \|x\| \cdot \|x'\|$ where $\|x\| := T_2(x)^{1/2}$ is the induced norm. The (field) norm of an element $x \in K_\mathbb{R}$ is defined as $\mathcal{N}(x) = \prod_i |\theta_i(x)|$. Note that with our definition, the norm cannot be negative.

A (fractional) ideal I is any finitely generated $\mathcal{O}_K$-module contained in K. An integral ideal I is a fractional ideal contained in $\mathcal{O}_K$. For any fractional ideal I there exists $r \in \mathbb{Z}$ such that rI is an integral ideal. If $r \in K$, we let (r) denote the (principal) ideal $r\mathcal{O}_K$. The product $IJ = \langle ij : i \in I, j \in J\rangle$ and the sum $I + J = \{i + j : i \in I, j \in J\}$ of two ideals are also ideals. A non-zero integral ideal is said to be prime if it is divisible only by $\mathcal{O}_K$ and itself. As $\mathcal{O}_K$ is a Dedekind domain, any non-zero fractional ideal can be uniquely decomposed as a product of (possibly negative) powers of prime ideals. If $\mathfrak{p}$ is a prime ideal, we define $\nu_\mathfrak{p}(I) = \max(k \in \mathbb{Z} : \mathfrak{p}^k|I)$. The norm of I is defined as $\mathcal{N}(I) = \det(I)/\det(\mathcal{O}_K)$. If $I \neq 0$ is integral, then this is exactly the index of I in $\mathcal{O}_K$, defined as $[\mathcal{O}_K : I] = |\mathcal{O}_K/I|$. We define $\mathcal{N}(0) = 0$, which allows us to assert that $\mathcal{N}(IJ) = \mathcal{N}(I)\mathcal{N}(J)$ for any ideals I and J. Note that if $I = (r)$ is principal, then $\mathcal{N}(I) = \mathcal{N}(r)$. The inverse $I^{-1} = \{r \in K : rI \subseteq \mathcal{O}_K\}$ of a non-zero fractional ideal I is also a fractional ideal, and we have $II^{-1} = \mathcal{O}_K$. Note that the arithmetic over the ideals can be performed in polynomial time (e.g., see [2]).

Any non-zero ideal, including the maximal order, is naturally a free $\mathbb{Z}$-module of rank d thus a lattice under the T_2-norm. By fixing an integral basis for K, we also fix a $\mathbb{Z}$-lattice structure for $\mathcal{O}_K$ that we can then reduce. We say that a basis of a non-zero fractional ideal I is in HNF if the (rational) matrix of the coefficients with respect to a fixed integral basis of K is in HNF. This provides a unique representation for any ideal. In the following, we assume that we know an integral basis $(r_i)_i$ of K that is `LatRed`-reduced with respect to T_2. It can be known for particular K's (e.g., cyclotomic number fields, with $\max\|r_i\|^2 = d$), or can be computed by reducing an arbitrary integral basis. As it is computed

once and for all, it may prove interesting to strongly reduce it. We have the following result.

Lemma 1. *If $(r_i)_i$ is a* `LatRed`*-reduced integral basis of K, then* $\max \|r_i\| \leq \sqrt{d}\gamma^d \sqrt{\Delta_K}$.

Proof. Using the reducedness and Minkowski's second theorem, we get $\prod \|r_i\|^2 \leq \gamma^{2d} d^d \Delta_K$. The arithmetic-geometric inequality gives $1 \leq \mathcal{N}(r_i)^{2/d} \leq \|r_i\|^2/d$ for all i, which provides the result. □

The bounds of our main results involve the quantity $\max \|r_i\|$. Lemma 1 allows one to express them with field invariants only. We choose to keep $\max \|r_i\|$ in our bounds since it can be much smaller, as in the case of cyclotomic number fields.

With our a choice of integral basis, any element of $\mathcal{O}_K$ with small T_2-norm can be represented with a small number of bits.

Lemma 2. *Assume that $(r_i)_i$ is a* `LatRed`*-reduced integral basis of K. If $x = \sum x_i r_i \in K$, then* $\max |x_i| \leq 2^{3d/2}\|x\|$.

Proof. We show by induction of i that

$$\forall i : |x_i| \leq 2^{d-i} \frac{\|x\|}{\min_j \|r_j^*\|}.$$

First, we have $\|x\| \geq |x_d| \|r_d^*\|$. Suppose now that $i < d$ and that the result holds for any $j > i$. The GSO of the r_i's shows that $\|x\| \geq |x_i + \sum_{j>i} \mu_{j,i} x_j| \|r_i^*\|$. Therefore, we have $|x_i| \leq \|x\|/\|r_i^*\| + \sum_{j>i} |x_j|$, which gives the bound. To complete the proof, note that the reducedness of the r_i's gives $\min_j \|r_j^*\| \geq 2^{-d/2} \min_j \|r_j\|$, and that $\|r_j\| \geq \sqrt{d}$ for all j. □

2.3 $\mathcal{O}_K$-modules

Let $\mathbf{b}_1, \ldots, \mathbf{b}_n \in K_{\mathbb{R}}^m$ with $n = \mathrm{rank}_K(\mathbf{b}_i)_i$, and $\mathfrak{b}_1, \ldots, \mathfrak{b}_n$ be fractional ideals of $\mathcal{O}_K$. The $\mathcal{O}_K$-module $M[(\mathbf{b}_i)_i, (\mathfrak{b}_i)_i]$ spanned by the pseudo-basis $[(\mathbf{b}_i)_i, (\mathfrak{b}_i)_i]$ is $\sum \mathfrak{b}_i \mathbf{b}_i$. The $\mathfrak{b}_i$'s are called the coefficient ideals. As each $\mathfrak{b}_i$ is a $\mathbb{Z}$-lattice, so is M. More precisely, if $\mathfrak{b}_i = \sum_{j \leq d} \mathbb{Z}\beta_i^{(j)}$, then $M = \sum_{i,j} \mathbb{Z}\beta_i^{(j)} \mathbf{b}_i$. Two pseudo-bases $[(\mathbf{b}_i)_i, (\mathfrak{b}_i)_i]$ and $[(\mathbf{c}_i)_i, (\mathfrak{c}_i)_i]$ represent the same $\mathcal{O}_K$-module M if and only if there exists a non-singular $U \in K^{n \times n}$ with ([23, §81 C]):

1. $(\mathbf{c}_1, \ldots, \mathbf{c}_n) = (\mathbf{b}_1, \ldots, \mathbf{b}_n)U$;
2. For all i, j, we have $U_{i,j} \in \mathfrak{b}_i \mathfrak{c}_j^{-1}$;
3. For all i, j, we have $U'_{i,j} \in \mathfrak{c}_i \mathfrak{b}_j^{-1}$, where $U' = U^{-1}$.

Cohen [6] generalized the HNF to modules in K^m. The algorithm of [4] may also be interpreted as such a generalization. We refer to [12, Chap. 4] for a detailed exposure and comparison.

Theorem 1. *Let $M \subseteq K^m$ be an $\mathcal{O}_K$-module of rank n. There exists a pseudo-basis $[(\mathbf{b}_i)_i, (\mathfrak{b}_i)_i]$ of M such that $\mu_j = \min\{i : B_{i,j} \neq 0\}$ (strictly) increases with j, for all j we have $B_{\mu_j,j} = 1$ and for all $j > k$ the entry $B_{\mu_j,k} \in K$ is size-reduced modulo the HNF of $\mathfrak{b}_j\mathfrak{b}_k^{-1}$. This unique pseudo-basis is called the HNF of M. It can be computed in polynomial time from any pseudo-basis of M.*

Similarly to the HNF for lattices, the above HNF can only handle $\mathcal{O}_K$-modules $M \subseteq K^m$ (as opposed to $K_{\mathbb{R}}^m$) and does not necessarily contain small elements of M. We now define the concept of small-ness for elements of $K_{\mathbb{R}}^m$. For any two vectors $\mathbf{b} = (b_1, \ldots, b_m)^t, \mathbf{b}' = (b'_1, \ldots, b'_m)^t \in K_{\mathbb{R}}^m$, we define $T_2^{\otimes m}(\mathbf{b}, \mathbf{b}') = \sum_{i \leq m} T_2(b_i, b'_i)$, and we denote $\sqrt{T_2^{\otimes m}(\mathbf{b}, \mathbf{b})}$ by $\|\mathbf{b}\|$. Notice that for any $(r, \mathbf{b}) \in K_{\mathbb{R}} \times K_{\mathbb{R}}^m$, we have $\|r\mathbf{b}\| \leq \|r\| \cdot \|\mathbf{b}\|$. With this definition at hand, we can define the minima of M:

$$\forall i \leq n,\ \lambda_i(M) = \min\{r : \exists \mathbf{c}_1, \ldots, \mathbf{c}_i \in M, \mathrm{rank}_K(\mathbf{c}_k)_k = i \ \text{ and } \ \max \|\mathbf{c}_k\| \leq r\}.$$

Let $[(\mathbf{b}_i)_i, (\mathfrak{b}_i)_i]$ be a pseudo-basis of an $\mathcal{O}_K$-module $M \subseteq K_{\mathbb{R}}^m$. Assume that $\mathfrak{b}_i = \sum_{j \leq d} \mathbb{Z}\beta_i^{(j)}$. We define $\det(M)$ as the square root of the determinant of the $nd \times nd$ symmetric positive definite matrix $T_2^{\otimes m}(\beta_i^{(j)}\mathbf{b}_i, \beta_{i'}^{(j')}\mathbf{b}_{i'})_{i,j;i',j'}$. This is a module invariant. When M is a non-zero fractional ideal of $\mathcal{O}_K$, this matches $\det_{T_2}(M)$. It should be noted that $\det(M)$ is not immediately related to the (Steinitz) class of M nor to the maximal exterior power of M. The following is a direct consequence of Minkowski's second theorem over $\mathbb{Z}$-lattices.

Theorem 2. *Let $M \subseteq K_{\mathbb{R}}^m$ be an $\mathcal{O}_K$-module of rank n. Then $\prod_{i \leq n} \lambda_i(M) \leq \sqrt{dn}^n \det(M)^{1/d}$.*

Proof. The module M can be seen as a lattice L of dimension nd, with $\det(M) = \det(L)$. Minkowski's second theorem asserts that $\prod_{i \leq nd} \lambda_i(L) \leq \sqrt{dn}^{dn} \det(L)$. Let $c_1, \ldots, c_{nd} \in M$ be free over the integers such that $\|c_i\| = \lambda_i(L)$ holds for all i. For all $i \leq n$, let $\phi(i) = \min(j : \mathrm{rank}_K(c_1, \ldots, c_j) = i)$. As $\mathcal{O}_K$ has rank d as a $\mathbb{Z}$-module, we have $\phi(i) \leq (i-1)d + 1$. We conclude with the following sequence of inequalities:

$$\prod_{i \leq n} \lambda_i(M) \leq \prod_{i \leq n} \|c_{\phi(i)}\| \leq \prod_{i \leq n} \lambda_{(i-1)d+1}(L) \leq \prod_{i \leq dn} \lambda_i(L)^{\frac{1}{d}} \leq \sqrt{dn}^n \det(M)^{\frac{1}{d}}. \quad \square$$

We now extend the concept of GSO. Let $[(\mathbf{b}_i)_i, (\mathfrak{b}_i)_i]$ be a pseudo-basis of an $\mathcal{O}_K$-module M. We define $\mathbf{b}_i^* = \mathrm{argmin}\|\mathbf{b}_i + \sum_{j<i} K_{\mathbb{R}}\mathbf{b}_j\|$ for all $i \leq n$, and let $\mu_{i,1}, \ldots, \mu_{i,i-1} \in K_{\mathbb{R}}$ be such that $\mathbf{b}_i = \mathbf{b}_i^* + \sum_{j<i} \mu_{i,j}\mathbf{b}_j^*$.

3 Small 2-element representation of an ideal

We start our study of $\mathcal{O}_K$-modules by the one-dimensional case, i.e., fractional ideals of K. There are several ways of representing an ideal $I \neq 0$. A natural approach is to provide a basis $(b_i)_{i \leq d} \in K^d$, or the coordinates matrix

of a basis with respect to an integral basis $(r_i)_i$ of K. This coordinates matrix belongs to $\mathbb{Q}^{d\times d}$, and it may prove interesting to find the basis of kI such that the coordinates matrix is in HNF, for the smallest non-zero integer k such that kI is integral. This representation requires a space of $O(d\log\mathcal{N}(kI)+\log k+d^2) = O(d\log\mathcal{N}(I)+d^2+d^2\log k)$ bits. Alternatively, one may use the so-called two-element representation: any ideal I may be written $I=(x_1)+(x_2)$ for some $x_1,x_2\in I$. A classical way to obtain such a representation consists in taking an arbitrary $x_1\in I$ and then choosing x_2 uniformly in I modulo (x_1) (the latter being a full-rank sublattice of the former). This succeeds with probability $\geq\prod(1-1/\mathcal{N}(\mathfrak{p}))$, where the product is taken over the prime ideals $\mathfrak{p}$ that divide $(x_1)/I$ (see [2, Le. 6.14]). If $\mathcal{N}(x_1)/\mathcal{N}(I)$ is small and if there do not exist too many prime ideals of small norm, then the success probability is large. Belabas [2, Alg. 6.15] proposed a probabilistic polynomial time variant, which always succeeds with high probability. However, the obtained representation of I may still be of bit-size $\Omega(d\log\mathcal{N}(I)+d+d^2\log k)$.

We modify Belabas' algorithm to provide a 2-element representation made of small elements: $I=(x_1)+(x_2)$ with both $\|x_1\|$ and $\|x_2\|$ small. For instance, the first element x_1 is chosen to be the first element of a `LatRed`-reduced basis of I. This may be seen as a rigorous variant of [7, Alg. 1.3.15], in which smallness was provided but the success probability could be small. Although our analysis is close to Belabas', we give a full proof, as there are quite a few small differences.

Theorem 3. *Let $(r_i)_i$ be an integral basis of a number field K. There exists a probabilistic polynomial time algorithm that takes as inputs a $\mathbb{Z}$-basis of a non-zero fractional ideal I of $\mathcal{O}_K$ and a success parameter t (in unary), and returns $x_1,x_2\in I$ such that $I=(x_1)+(x_2)$ holds with probability $1-2^{-t}$, and:*

$$\|x_1\|,\|x_2\| \leq 4\gamma^8\Delta_K^{\frac{4}{d}}\max\|r_i\|^4\cdot\mathcal{N}(I)^{\frac{4}{d}}, \tag{1}$$

where $\|\cdot\|$ corresponds to the T_2 norm and γ is the `LatRed` *approximation constant. As a consequence, the ideal I may be represented on $5\log_2\mathcal{N}(I)+O(\log\Delta_K+d(d+\log k+\log\max\|r_i\|))$ bits, where k is the smallest non-zero integer such that kI is integral and the r_i's are assumed* `LatRed`*-reduced.*

Let us comment on (1). The quantity $4\gamma^8\Delta_K^{\frac{4}{d}}$ is an invariant of the field, and $\max\|r_i\|^4$ is independent from I (and can be bounded using Lemma 1). The only term that is not an invariant is $\mathcal{N}(I)^{\frac{4}{d}}$. If x_1 and x_2 were basis vectors of a reduced basis of I, we would expect $\mathcal{N}(I)^{\frac{1}{d}}$ instead of $\mathcal{N}(I)^{\frac{4}{d}}$ (see (2) below). We do not know how to reach this bound for x_2.

Let us now prove Theorem 3. Since the smallest integer k such that kI is integral can be computed efficiently, we assume that I is integral. As the ideal I is given by a $\mathbb{Z}$-basis, we can find a basis of it that is `LatRed`-reduced (for T_2). The algorithm of Figure 1 is an adaptation of [2, Alg. 6.15]. We follow the algorithm step by step. The reducedness of the input directly gives that $\|x_1\|\leq$

Inputs: A `LatRed`-reduced basis of a non-zero integral ideal I of $\mathcal{O}_K$;
a success parameter t.
Output: $x_1, x_2 \in I$ such that $I = (x_1) + (x_2)$, or `Fail`.

1. Let x_1 be the first basis element; $\mathfrak{a} := (x_1)$. If $I = \mathfrak{a}$, return x_1 and $x_2 := 0$.
2. Find y such that $y \log y = \log \mathcal{N}(\mathfrak{a})$; $S := \{\mathfrak{p} \text{ prime} : \mathcal{N}(\mathfrak{p}) \leq y\}$.
3. $\mathfrak{a}_0 := \prod_{\mathfrak{p} \in S} \mathfrak{p}^{\nu_\mathfrak{p}(\mathfrak{a})}$; $I_0 := \prod_{\mathfrak{p} \in S} \mathfrak{p}^{\nu_\mathfrak{p}(I)}$; $\mathfrak{a}_1 := \mathfrak{a}\mathfrak{a}_0^{-1}$; $I_1 := I I_0^{-1}$.
4. For $i := 1$ to $2t$ do
5. Sample π_1 uniformly in $I_1/\mathfrak{a}_1$. If $I_1 = \mathfrak{a}_1 + (\pi_1)$, then go to Step 7.
6. Return `Fail`.
7. Let b be the first element of a `LatRed`-reduced basis of $\mathfrak{a}_1$.
8. Size-reduce π_1 with respect to the $b \cdot r_i$'s.
9. Using [2, Alg. 6.8], find $\pi_0 \in \mathcal{O}_K$ such that $\nu_\mathfrak{p}(\pi_0) = \nu_\mathfrak{p}(I_0)$ for all $\mathfrak{p} \in S$.
10. Let b be the first element of a `LatRed`-reduced basis of $\prod_{\mathfrak{p} \in S} \mathfrak{p}^{\nu_\mathfrak{p}(I_0)+1}$.
11. Size-reduce π_0 with respect to the $b \cdot r_i$'s.
12. Using [2, Alg. 5.4], find $\alpha_0 \in \mathfrak{a}_0$ and $\alpha_1 \in \mathfrak{a}_1$ such that $\alpha_0 + \alpha_1 = 1$.
13. Let b be the first element of a `LatRed`-reduced basis of $\mathfrak{a}$.
14. Size-reduce α_0 and α_1 with respect to the $b \cdot r_i$'s.
15. Return x_1 and $x_2 := (\pi_0\alpha_1 + \alpha_0)(\pi_1\alpha_0 + \alpha_1)$.

Fig. 1. Computing a small 2-element representation of an integral ideal.

$\gamma \Delta_K^{1/2d} \mathcal{N}(I)^{1/d}$. By using the arithmetic-geometric inequality, we obtain:

$$\mathcal{N}(\mathfrak{a})^{\frac{1}{d}} = \mathcal{N}(x_1)^{\frac{1}{d}} \leq \frac{1}{\sqrt{d}} \|x_1\| \leq \frac{\gamma \Delta_K^{\frac{1}{2d}}}{\sqrt{d}} \mathcal{N}(I)^{\frac{1}{d}}. \tag{2}$$

As a consequence, the variable y of Step 2, can be bounded by a polynomial in d, $\log \mathcal{N}(I)$ and $\log \Delta_K$. This ensures that the computation of S can be done in polynomial time. At Step 3, the computations of $\mathfrak{a}_0$, I_0, $\mathfrak{a}_1$ and I_1 can be performed in polynomial time: this follows from the above study of S. We have $\mathfrak{a} = \mathfrak{a}_0\mathfrak{a}_1$ and $I = I_0 I_1$. We also have $I_i | \mathfrak{a}_i$ and $I_i + \mathfrak{a}_{1-i} = \mathcal{O}_K$ for $i \in \{0, 1\}$.

As $\mathfrak{a}_1$ is a full-rank sublattice of I_1, sampling π_1 uniformly in $I_1/\mathfrak{a}_1$ can be done in polynomial time. The equality $I_1 = \mathfrak{a}_1 + (\pi_1)$ can also be tested in polynomial time (see, e.g., [20, Prop. 8.2]). By adapting the analysis of [2, Le. 6.1], we obtain:

$$\Pr[I_1 = \mathfrak{a}_1 + (\pi_1)] \geq \prod_{\mathfrak{p} \text{ prime},\ \mathfrak{p} | \mathfrak{a}_1} \left(1 - \frac{1}{\mathcal{N}(\mathfrak{p})}\right) \geq \left(1 - \frac{1}{y}\right)^{\log_y \mathcal{N}(\mathfrak{a})} \geq \frac{1}{\mathrm{e}}.$$

As a consequence, the algorithm returns `Fail` at Step 6 with probability $\leq 2^{-t}$.

At Step 8, the $b \cdot r_i$'s are a basis of a sublattice of $\mathfrak{a}_1$. Therefore, after Step 8, we still have $I_1 = \mathfrak{a}_1 + (\pi_1)$. After the size-reduction of π_1 with respect to the $b \cdot r_i$'s, we have:

$$\|\pi_1\| \leq \sqrt{d} \max_i \|b r_i\| \leq \sqrt{d} \|b\| \max_i \|r_i\| \leq \sqrt{d} \gamma \Delta_K^{\frac{1}{2d}} \mathcal{N}(\mathfrak{a}_1)^{\frac{1}{d}} \max_i \|r_i\|.$$

It is shown in [2] that Step 9 can be performed in polynomial time. The bounds on S imply that Step 10 can be done in polynomial time. Step 11 ensures that

$$\|\pi_0\| \leq \sqrt{d}\gamma\Delta_K^{\frac{1}{2d}}\mathcal{N}\left(\prod_{\mathfrak{p}\in S}\mathfrak{p}^{\nu_\mathfrak{p}(I_0)+1}\right)^{\frac{1}{d}}\max_i\|r_i\| \leq \sqrt{d}\gamma\Delta_K^{\frac{1}{2d}}\mathcal{N}(I_0)^{\frac{2}{d}}\max_i\|r_i\|.$$

After Step 11, we still have that $\nu_\mathfrak{p}(\pi_0) = \nu_\mathfrak{p}(I_0)$, for all $\mathfrak{p} \in S$, and thus $I_0 = \mathfrak{a}_0 + (\pi_0)$. It is shown in [2] that Step 12 can be performed in polynomial time. Step 14 ensures that $\|\alpha_0\|, \|\alpha_1\| \leq \sqrt{d}\gamma\Delta_K^{\frac{1}{2d}}\mathcal{N}(\mathfrak{a})^{\frac{1}{d}}\max_i\|r_i\|$. Since $\mathfrak{a} = \mathfrak{a}_0\mathfrak{a}_1$, we still have $\alpha_i \in \mathfrak{a}_i$ after Step 14, for $i \in \{0,1\}$. At Step 15, we have:

$$\begin{aligned}\|x_2\| &\leq (\|\pi_0\|\|\alpha_1\| + \|\alpha_0\|)(\|\pi_1\|\|\alpha_0\| + \|\alpha_1\|)\\ &\leq d^2\gamma^4\Delta_K^{\frac{2}{d}}\max_i\|r_i\|^4\mathcal{N}(\mathfrak{a})^{\frac{2}{d}}\left(\mathcal{N}(I_0)^{\frac{2}{d}}+1\right)\left(\mathcal{N}(\mathfrak{a}_1)^{\frac{1}{d}}+1\right)\\ &\leq 4d^2\gamma^4\Delta_K^{\frac{2}{d}}\max_i\|r_i\|^4\mathcal{N}(\mathfrak{a})^{\frac{4}{d}},\end{aligned}$$

where we used the fact that $\mathcal{N}(\mathfrak{a}_1) = \mathcal{N}(\mathfrak{a})/\mathcal{N}(\mathfrak{a}_0) \leq \mathcal{N}(\mathfrak{a})/\mathcal{N}(I_0)$. Combining the latter with (2) provides the upper bound on $\|x_2\|$ from Theorem 3.

Also, we have that $\pi_i' := \pi_i\alpha_{1-i} + \alpha_i$ is congruent to π_i modulo $\mathfrak{a}_i$ and to 1 modulo $\mathfrak{a}_{1-i}$, for $i \in \{0,1\}$. Therefore, we have $I_i = \mathfrak{a}_i + (\pi_i')$ and $I_i + (\pi_{i-1}') = \mathcal{O}_K$. Finally, we obtain $I = I_0I_1 = \mathfrak{a}_0\mathfrak{a}_1 + (\pi_0'\pi_1') = (x_1) + (x_2)$, thus proving the correctness of the algorithm.

We now consider the amount of space needed to represent the coordinates of x_1 and x_2 with respect to the integral basis $(r_i)_i$. We write $x_j = \sum y_i^{(j)}r_i$ with $y_i^{(j)} \in \mathbb{Z}$ and $j \in \{1,2\}$. Using Lemma 2, we have that each $y_i^{(j)}$ may be stored on $\log_2\|x_j\| + O(d)$ bits. Combining the latter with (2) and (1) provides the result. □

4 Computing short pseudo-bases

In this section, we (constructively) show that any $\mathcal{O}_K$-module $M \subseteq K_\mathbb{R}^m$ always has a pseudo-basis $[(\mathbf{b}_i)_i, (\mathfrak{b}_i)_i]$ such that the $\mathbf{b}_i$'s belong to M and are not much longer than the module minima.

4.1 From a short basis of a submodule to a short pseudo-basis

We are going to generalize to $\mathcal{O}_K$-modules the technique we mentioned at the end of Section 2.1, that takes as inputs a basis of a lattice L and a short basis of a full-rank sub-lattice of L, and returns a short basis of L. We split the algorithm into several smaller ones that may be of independent interest.

The algorithm of Figure 2 takes as inputs a pseudo-basis $[(\mathbf{a}_i)_i, (\mathfrak{a}_i)_i]$ of an $\mathcal{O}_K$-module $M \subseteq K_\mathbb{R}^m$ and a full-rank set of short module vectors $(\mathbf{s}_i)_i$, and returns a pseudo-basis $[(\mathbf{b}_i)_i, (\mathfrak{b}_i)_i]$ of M such that $\mathbf{b}_i \in \mathrm{span}_{j\leq i}\,\mathbf{s}_j$. This can be

Inputs: A pseudo-basis $[(\mathbf{a}_i)_i, (\mathfrak{a}_i)_i]$ of an $\mathcal{O}_K$-module $M \subseteq K_{\mathbb{R}}^m$, a full-rank set $(\mathbf{s}_i)_i$ of vectors in M.
Output: A pseudo-basis of M.
1. Compute $T \in K^{n \times n}$ such that $(\mathbf{s}_1, \ldots, \mathbf{s}_n) = (\mathbf{a}_1, \ldots, \mathbf{a}_n)T$.
2. Let $\mathbf{t}_1, \ldots, \mathbf{t}_n$ be the columns of T^t.
3. Compute the BPC-HNF $[(\mathbf{t}'_i)_i, (\mathfrak{b}_i^{-1})_i]$ of the pseudo-basis $[(\mathbf{t}_i)_i, (\mathfrak{a}_i^{-1})_i]$.
4. Let T' be the matrix whose rows are the $(\mathbf{t}'_i)^t$'s, and $U = T(T')^{-1} \in K^{n \times n}$.
5. Let $(\mathbf{b}_1, \ldots, \mathbf{b}_n) = (\mathbf{a}_1, \ldots, \mathbf{a}_n)U$.
6. Return $[(\mathbf{b}_i)_i, (\mathfrak{b}_i)_i]$.

Fig. 2. Constructing a pseudo-basis with small GSO.

interpreted as a constructive variant of [23, Th. 81.3]. The HNF over lattices is replaced by the BPC-HNF (Theorem 1), with special care being taken for the coefficient ideals.

Theorem 4. *If given as inputs a pseudo-basis $[(\mathbf{a}_i)_i, (\mathfrak{a}_i)_i]$ of a module $M \subseteq K_{\mathbb{R}}^m$ and a full-rank set $(\mathbf{s}_i)_i$ of vectors in M, then the algorithm of Figure 2 returns a pseudo-basis $[(\mathbf{b}_i)_i, (\mathfrak{b}_i)_i]$ of M, which satisfies, for all $i \leq n$: $\mathbf{b}_i \in M$; $\mathbf{b}_i \in \mathrm{span}_{j \leq i} \mathbf{s}_j$; $\mathbf{b}_i^* = \mathbf{s}_i^*$. If $M \subseteq K^m$, then it terminates in polynomial time.*

Proof. We first prove that $[(\mathbf{b}_i)_i, (\mathfrak{b}_i)_i]$ is a pseudo-basis of M. We have $(\mathbf{b}_i)_i = (\mathbf{a}_i)_i \cdot U$, with $U \in K^{n \times n}$ non-singular. It therefore suffices to prove that for any i, j, we have $U_{i,j} \in \mathfrak{a}_i \mathfrak{b}_j^{-1}$ and $U'_{i,j} \in \mathfrak{b}_i \mathfrak{a}_j^{-1}$, where $U' = U^{-1}$. This is ensured by Theorem 1: as the pseudo-bases $[(\mathbf{t}'_i)_i, (\mathfrak{b}_i^{-1})_i]$ and $[(\mathbf{t}_i)_i, (\mathfrak{a}_i^{-1})_i]$ span the same module, we have $U'_{j,i} \in \mathfrak{a}_i^{-1}\mathfrak{b}_j$ and $U_{j,i} \in \mathfrak{b}_i^{-1}\mathfrak{a}_j$, for any i, j.

Because of the definitions of T, T', U and $(\mathbf{b}_i)_i$, we have $(\mathbf{s}_i)_i = (\mathbf{b}_i)_i \cdot T'$. Furthermore, by Theorem 1, the matrix T' is upper triangular with diagonal coefficients equal to 1. We thus have $\mathbf{b}_i \in \mathrm{span}_{j \leq i} \mathbf{s}_j$, for all i. In fact, we even have $\mathbf{b}_i + \sum_{j<i} K_{\mathbb{R}} \mathbf{b}_j = \mathbf{s}_i + \sum_{j<i} K_{\mathbb{R}} \mathbf{s}_j$, which gives $\mathbf{b}_i^* = \mathbf{s}_i^*$. Finally, the shape of T' gives that $\mathbf{s}_i = \mathbf{b}_i + \sum_{j<i} T'_{j,i} \mathbf{b}_j$. As the $\mathbf{s}_i$'s belong to M, so must the $\mathbf{b}_i$'s (the decomposition of $\mathbf{s}_i$ as an element of $\sum_j K\mathbf{b}_j$ is unique). □

The algorithm of Figure 3 generalizes size-reduction to $\mathcal{O}_K$-modules.

Theorem 5. *If given as input a pseudo-basis $[(\mathbf{a}_i)_i, (\mathfrak{a}_i)_i]$ of a module $M \subseteq K_{\mathbb{R}}^m$, then the algorithm of Figure 3 returns a pseudo-basis $[(\mathbf{b}_i)_i, (\mathfrak{b}_i)_i]$ of M, such that for all i we have $\mathbf{b}_i^* = \mathbf{a}_i^*$, $\mathfrak{b}_i = \mathfrak{a}_i$ and*

$$\|\mathbf{b}_i\| \leq \sqrt{dn}\gamma \Delta_K^{\frac{1}{2d}} \max_k \|r_k\| \left(\frac{\max_{j \leq i} \mathcal{N}(\mathfrak{b}_j)}{\mathcal{N}(\mathfrak{b}_i)} \right)^{\frac{1}{d}} \max_{j \leq i} \|\mathbf{a}_j^*\|.$$

If $M \subseteq K^m$ and `LatRed` *is LLL, then it terminates in polynomial time.*

Proof. The operations performed on the pseudo-basis can be checked to preserve the generated module and the $\mathbf{b}_i^*$'s. Steps 2, 6 and 7 ensure that the $\mu_{i,j}$'s of the output pseudo-basis satisfy $\|\mu_{i,j}\| \leq \sqrt{d}\gamma \Delta_K^{\frac{1}{2d}} \mathcal{N}(\mathfrak{b}_i^{-1}\mathfrak{b}_j)^{\frac{1}{d}} \max_k \|r_k\|$. Pythagoras' theorem then provides the result. □

Input: A pseudo-basis $[(\mathbf{a}_i)_i, (\mathfrak{a}_i)_i]$ of an $\mathcal{O}_K$-module $M \subseteq K_{\mathbb{R}}^m$.
Output: A pseudo-basis of M.
1. $[(\mathbf{b}_i)_i, (\mathfrak{b}_i)_i] := [(\mathbf{a}_i)_i, (\mathfrak{a}_i)_i]$.
2. For $j \leq i$, let $x_{i,j}$ be the first element of a `LatRed` basis of $\mathfrak{b}_i^{-1}\mathfrak{b}_j$.
3. For i from 2 to n, do
4. For j from $i-1$ to 1, do
5. Compute the GSO decomposition $\mathbf{b}_i = \mathbf{b}_i^* + \sum_{j<i} \mu_{i,j}\mathbf{b}_j^*$,
6. Let y be the size-reduction of $\mu_{i,j}$ with respect to the $x_{i,j}r_k$'s,
7. $\mathbf{b}_i := \mathbf{b}_i - (\mu_{i,j} - y)\mathbf{b}_j$.
8. Return $[(\mathbf{b}_i)_i, (\mathfrak{b}_i)_i]$.

Fig. 3. Size-reducing a pseudo-basis of an $\mathcal{O}_K$-module.

The adaptation to $\mathcal{O}_K$-modules of [20, Le. 7.1] is given in Figure 4. The aim of Steps 2–4 is to allow us to bound the term $\frac{\max_{j\leq i}\mathcal{N}(\mathfrak{b}_j)}{\mathcal{N}(\mathfrak{b}_i)}$ from Theorem 5.

Inputs: A pseudo-basis $[(\mathbf{a}_i)_i, (\mathfrak{a}_i)_i]$ of an $\mathcal{O}_K$-module $M \subseteq K_{\mathbb{R}}^m$,
a free full-rank set $(\mathbf{s}_i)_i$ of vectors in M.
Output: A pseudo-basis of M.
1. Use the algorithm of Figure 2 to obtain a pseudo-basis $[(\mathbf{b}_i)_i, (\mathfrak{b}_i)_i]$ of M.
2. For any $i \leq n$,
3. Let $x \in \mathfrak{b}_i$ be the first vector of a `LatRed` basis of $\mathfrak{b}_i$,
4. $\mathfrak{b}_i := (x)^{-1}\mathfrak{b}_i$; $\mathbf{b}_i := x\mathbf{b}_i$.
5. Return the output of the algorithm of Figure 3, given $[(\mathbf{b}_i)_i, (\mathfrak{b}_i)_i]$ as input.

Fig. 4. From small vectors to a small pseudo-basis

Theorem 6. *If given as inputs a pseudo-basis $[(\mathbf{a}_i)_i, (\mathfrak{a}_i)_i]$ of an $\mathcal{O}_K$-module $M \subseteq K_{\mathbb{R}}^m$ and a full-rank set $(\mathbf{s}_i)_i$ of vectors in M, then the algorithm of Figure 4 returns a pseudo-basis $[(\mathbf{b}_i)_i, (\mathfrak{b}_i)_i]$ of M, such that for all i: $\mathbf{b}_i \in M$, $\mathrm{span}_{j\leq i}\mathbf{b}_j = \mathrm{span}_{j\leq i}\mathbf{s}_j$, $\|\mathbf{b}_i^*\| \leq \gamma\Delta_K^{\frac{1}{2d}}\|\mathbf{s}_i^*\|$, $\mathcal{N}(\mathfrak{b}_i) \in \left[\left(\frac{\sqrt{d}}{\gamma}\right)^d \frac{1}{\sqrt{\Delta_K}}, 1\right]$ and*

$$\|\mathbf{b}_i\| \leq \sqrt{n}\gamma^3\Delta_K^{\frac{3}{2d}} \max_k \|r_k\| \cdot \max_{j\leq i}\|\mathbf{s}_j\|.$$

If $M \subseteq K^m$ and `LatRed` is LLL, then it terminates in polynomial time.

Proof. The fact that the algorithm returns a pseudo-basis of M is easy to check. Also, at the end of Step 1, we have that $\mathbf{b}_i \in M$, for all i. Since the x of Step 3 belongs to $\mathfrak{b}_i$, the latter fact is preserved throughout the rest of the execution. The equality $\mathrm{span}_{j\leq i}\mathbf{b}_j = \mathrm{span}_{j\leq i}\mathbf{s}_j$ directly derives from Theorems 4 and 5.

At any time after Step 1, we have $\mathcal{O}_K \subseteq \mathfrak{b}_i$ and thus $\mathcal{N}(\mathfrak{b}_i) \leq 1$. At Step 3, we have $\|x\| \leq \gamma\Delta_K^{\frac{1}{2d}}\mathcal{N}(\mathfrak{b}_i)^{\frac{1}{d}}$. This gives that after Step 4 we have $\|\mathbf{b}_i^*\| \leq$

$\gamma \Delta_K^{\frac{1}{2d}} \|\mathbf{s}_i^*\|$, which is preserved throughout Step 5. Also, the arithmetic-geometric inequality implies that $\mathcal{N}(x) \leq (\gamma/\sqrt{d})^d \sqrt{\Delta_K} \mathcal{N}(\mathfrak{b}_i)$. Therefore, after Step 4 the quantity $\mathcal{N}(\mathfrak{b}_i)$ has been divided by $\mathcal{N}(x)$ and we have $\mathcal{N}(\mathfrak{b}_i) \geq \left(\frac{\sqrt{d}}{\gamma}\right)^d \frac{1}{\sqrt{\Delta_K}}$. Using Theorem 5, this allows us to derive that at the end of the execution we have:

$$\|\mathbf{b}_i\| \leq \sqrt{dn} \gamma \Delta_K^{\frac{1}{2d}} \max_k \|r_k\| \left(\frac{\sqrt{\Delta_K}}{(\sqrt{d}/\gamma)^d}\right)^{\frac{1}{d}} \cdot \left(\gamma \Delta_K^{\frac{1}{2d}} \max_{j \leq i} \|\mathbf{s}_j^*\|\right).$$

The inequalities $\|\mathbf{s}_j^*\| \leq \|\mathbf{s}_j\|$ lead to the result. □

4.2 Computing a short pseudo-basis

Suppose we have a pseudo-basis of an $\mathcal{O}_K$-module M of rank n. We can expand it to obtain a basis of M as a $\mathbb{Z}$-module. By LLL-reducing the latter with respect to T_2, we obtain dn module vectors whose integer linear combinations span M. By using linear algebra over K, it is possible to select n module vectors $\mathbf{s}_1, \ldots, \mathbf{s}_n$ among these dn vectors, such that $\mathrm{rank}_K(\mathbf{s}_i) = n$. Furthermore, thanks to the initial reduction, these vectors are also small, and we can apply Theorem 6.

Corollary 1. *There exists an algorithm that takes as input a pseudo-basis of an $\mathcal{O}_K$-module $M \subseteq K_{\mathbb{R}}^m$ and returns a pseudo-basis $[(\mathbf{b}_i)_i, (\mathfrak{b}_i)_i]$ of M, such that for all i: $\mathbf{b}_i \in M$, $\mathcal{N}(\mathfrak{b}_i) \in \left[\left(\frac{\sqrt{d}}{\gamma}\right)^d \frac{1}{\sqrt{\Delta_K}}, 1\right]$ and*

$$\|\mathbf{b}_i\| \leq 2^{\frac{dn}{2}} \sqrt{n} \gamma^3 \Delta_K^{\frac{3}{2d}} \max_k \|r_k\|^2 \cdot \lambda_i(M).$$

Therefore:

$$\prod_i \|\mathbf{b}_i\| \leq 2^{\frac{dn^2}{2}} (\sqrt{dn})^n \gamma^{3n} \Delta_K^{\frac{3n}{2d}} \max_k \|r_k\|^{2n} \cdot (\det(M))^{\frac{1}{d}}.$$

If $M \subseteq K^m$ and `LatRed` *is LLL, then it terminates in polynomial time, and the output may be stored on a number of bits bounded by*

$$m \log_2 \det(M) + O\left(md^2n^2 + nm \log \Delta_K + mdn \log \max_k \|r_k\|\right).$$

Proof. Let L denote M when considered as a lattice. Let $(\mathbf{s}_i)_{i \leq dn}$ be a LLL-reduced basis of L. We have $\|\mathbf{s}_i\| \leq 2^{dn/2} \lambda_i(L)$, for all i. Let $\psi(i) = \min(j : \mathrm{rank}_K(\mathbf{s}_k)_{k \leq j} = i)$. Since K has degree d, we have $\psi(i) \leq d(i-1)+1$, for all i. We use the $\mathbf{s}_{\psi(i)}$'s as input to the algorithm of Figure 4. The first statement on the $\|\mathbf{b}_i\|$'s derives from Theorem 6 and the fact that $\lambda_{\psi(i)}(L) \leq \max_k \|r_k\| \cdot \lambda_{\lceil \psi(i)/d \rceil}(M) \leq \max_k \|r_k\| \cdot \lambda_i(M)$. By combining Theorem 2 and the latter, we obtain the second statement on the $\|\mathbf{b}_i\|$'s.

We now consider the bit-size of the representation when $M \subseteq K^m$. Using Lemma 2 for the first component of the pseudo-basis, we obtain that the bit-size of the latter is $\leq md(\log_2 \prod \|\mathbf{b}_i\| + O(nd))$. To represent the ideal coefficients, we use Theorem 3 with the inverses of the ideals. The latter are integral, and have norms $\leq 2^{d^2} \Delta_K$. Therefore, each of these can be represented on $O(d^2 + \log \Delta_K + d \log \max_k \|r_k\|)$ bits. This completes the proof of the theorem. □

Note that the norm bound on the ideals depends only on the field and the choice for `LatRed` and is, in particular, independent of M.

By applying Corollary 1 with $m = n = 1$, we obtain yet another compact representation of ideals of K. If I is an ideal and k is the smallest non-zero integer such that kI is integral, then we see that I can be represented on $\log_2 \mathcal{N}(I) + O(\log \Delta_K + d(d + \log k + \log \max_i \|r_i\|))$ bits. If $\mathcal{N}(I)$ is large, this representation is smaller than the one from Theorem 3, but for a small $\mathcal{N}(I)$, this is the opposite as the $O(\cdot)$ constant is larger. Considering $((x_1)+(x_2))$ instead of its inverse leads to a representation whose bit-size grows faster with respect to d.

4.3 Short almost free pseudo-bases

A common strengthening of the properties of a pseudo-basis is to pass to an almost free (or Steinitz) representation: For any M, there exists a pseudo-basis $[(\mathbf{b}_i)_i, (\mathfrak{b}_i)_i]$ of M with $\mathfrak{b}_i = \mathcal{O}_K$ for $i < n$. We explain here how to obtain an almost free pseudo-basis consisting of short vectors. We first use Corollary 1 to find a "short" pseudo-basis. We then use the following lemma, from [7, Prop. 1.3.12, Alg. 1.3.16], which allows us to pass from a module with coefficient ideals $(\mathfrak{a}, \mathfrak{b})$ to a representation of this module with ideals $(1, \mathfrak{ab})$.

Lemma 3. *Let $\mathfrak{a}$ and $\mathfrak{b}$ be non-zero fractional ideals. There exists a polynomial-time algorithm that finds $a \in \mathfrak{a}$, $b \in \mathfrak{b}$, $x \in \mathfrak{a}^{-1}$, $y \in \mathfrak{b}^{-1}$ such that $ax - by = 1$.*

One can use Lemma 3 to progressively change the short pseudo-basis obtained in Corollary 1 into a short almost free pseudo-basis, collecting all the coefficient ideals into the last one. The corresponding algorithm is given in Figure 5. It can be checked that the output is an almost free pseudo-basis of the input module.

Input: A pseudo-basis $[(\mathbf{a}_i)_i, (\mathfrak{a}_i)_i]$ of an $\mathcal{O}_K$-module $M \subseteq K_{\mathbb{R}}^m$.
Output: An almost free pseudo-basis of M.
1. For $i = 1$ to $n - 1$ do
2. Use Lemma 3 with $\mathfrak{a} := \mathfrak{a}_i$, $\mathfrak{b} := \mathfrak{a}_{i+1}$ to find a, b, x, y as indicated,
3. Replace $\mathbf{a}_i$ by $a\mathbf{a}_i + b\mathbf{a}_{i+1}$ and $\mathbf{a}_{i+1}$ by $y\mathbf{a}_i + x\mathbf{a}_{i+1}$,
4. Replace $\mathfrak{a}_{i+1}$ by $\mathfrak{a}_i\mathfrak{a}_{i+1}$ and $\mathfrak{a}_i$ by $\mathcal{O}_K$.

Fig. 5. From a pseudo-basis to an almost free pseudo-basis.

Furthermore, if the input of the algorithm is a module pseudo-basis such as in Corollary 1, then during the execution, Lemma 3 is applied to ideals whose

norms can be bounded independently of the module M. As a consequence, the obtained transformation coefficients a, b, x, y have T_2-norms that can be bounded independently of M. At the end of the execution, we still have $\mathbf{a}_i \in M$ for all i, and the quantity $\prod_i \|\mathbf{a}_i\|$ (resp. each $\|\mathbf{a}_i\|$) remains bounded by $\det(M)^{\frac{1}{d}}$ (resp. by the corresponding $\lambda_i(M)$) up to a multiplicative factor that is independent of M. Similarly, the norm of the non-trivial coefficient ideal can also be bounded independently of M.

Finally, it should be noted that the basis generated by the algorithm of Figure 5 satisfies $\mathbf{b}_i \in \mathrm{span}_{j \leq i+1} \mathbf{a}_j$ for $i < n$, and thus can be compared to the results from [8].

5 Examples

We start by some example coming from group theory, focusing only on the use of lattice reduction. Representations of finite groups give easy access to non-trivial and interesting lattices. In general starting with a finite subgroup $G < \mathrm{GL}(m, K)$ and any $\mathcal{O}_K$-module N we obtain a G-invariant $\mathcal{O}_K$-module M via $M := \sum_{g \in G} Ng$. Next we change G to act on M, $G \in \mathrm{GL}(M)$ and, fixing a complex conjugation on K, obtain a G-invariant Hermitean form on K^m from $H := \sum_{g \in G} g^* g$. The main application is to find a reduced (short) pseudo-basis $S = MT$ for M and then replace G by $G^T = \{T^{-1}gT : g \in G\}$ to find an isomorphic version of G where the elements are (hopefully) "smaller".

Let G be the quaternion group Q_8 with 8 elements. As a subgroup of $\mathrm{GL}(2, K)$ for $K := \mathbb{Q}(i)$, it can be generated by

$$\frac{1}{5}\begin{pmatrix} i+2 & 2i-6 \\ 2i+4 & -i-2 \end{pmatrix} \quad \text{and} \quad \frac{1}{2}\begin{pmatrix} -i-1 & 3i+1 \\ i-1 & i+1 \end{pmatrix}.$$

Computing the $\mathcal{O}_K$-module generated by $g\begin{pmatrix} 1 \\ 0 \end{pmatrix}$ for all $g \in G$, we use $M := \mathcal{O}_K \begin{pmatrix} 1 \\ 0 \end{pmatrix} + (\frac{1+3i}{10}\mathcal{O}_K) \begin{pmatrix} 3 \\ 1 \end{pmatrix}$. As a Hermitean form, we compute $\sum_{g \in G} gg^*$ where g^* denotes the transposed complex conjugate. We then normalize the matrix to have 1 as the top left entry and obtain

$$H := \frac{1}{5}\begin{pmatrix} 5 & i+2 \\ -i+2 & 3 \end{pmatrix}.$$

We reduce the corresponding $\mathbb{Z}$-lattice and use the following short $\mathbb{Q}(i)$-independent basis elements:

$$\frac{1}{10}\begin{pmatrix} -3i-1 \\ -i+3 \end{pmatrix} \quad \text{and} \quad -\frac{1}{5}\begin{pmatrix} 2i-1 \\ -i+3 \end{pmatrix}.$$

The two elements can be seen to freely generate the module. Using the transformation to change G, we now get

$$\begin{pmatrix} i & 0 \\ 0 & -i \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}.$$

which is a "nicer" version of G.

Let $G := \mathrm{S}Z_8$ the 8th Suzuki group with 29 120 elements. This group has 11 characters, and we consider the second among them. The latter defines a representation of degree 14 over some field containing i. For theoretical reasons, the representation can be defined over $\mathbb{Q}(i)$, but it is initially computed over $\mathbb{Q}(\zeta_{52})$, of degree 24. A complicated procedure will now find a representation over $\mathbb{Q}(i)$, i.e., we have three matrices (one for each generator) over $\mathbb{Q}(i)$ generating G. The coefficients of the original matrix entries over $\mathbb{Q}(\zeta_{52})$ have about 100 digits each, and over $\mathbb{Q}(i)$ this increases to about 200 digits. In this representation the group G fixes a Hermitean form M which has again entries with about 200 digits each. Since the representation is absolutely irreducible, the quadratic form is unique up to multiplication by scalars. We normalized the form to have 1 as the entry in position $(1, 1)$. After application of our reduction technique, the form as well as the representation now have only 1 digit entries. The module used here is generated by $Ge_1 \subseteq \mathbb{Q}(i)^2$.

We used the following Magma code to generate the second example:

```
> G := Sz(8);
> T := CharacterTable(G);
> M := GModule(T[2]:SparseCyclo := false);
> N := AbsoluteModuleOverMinimalField(M);
> IsAlmostIntegral(N); //computes the module
true
> _ := InvariantForm(N);  // compute the form
> SetVerbose("RLLL", 1);
> O := Nice(N);
> #Sprint(ActionGenerators(M));
1359862
> #Sprint(ActionGenerators(N));
327378
> #Sprint(ActionGenerators(O));
4577
```

The function `Nice` implements the procedure outlined above. Note that the actual result can vary substantially as several parts use randomized algorithms. The `Sprint` statements are only used as a very crude indication of the output size, they simply give the number of characters neccessary to write the generating matrices for G.

References

1. M. Ajtai, R. Kumar, and D. Sivakumar. A sieve algorithm for the shortest lattice vector problem. In *Proc. STOC 2001*, pages 601–610. ACM, 2001.
2. K. Belabas. Topics in computational algebraic number theory. *J. théorie des nombres de Bordeaux*, 16:19–63, 2004.

3. W. Bosma, J. Cannon, and C. Playoust. The Magma algebra system. I. The user language. *J. Symbolic Comput.*, 24(3–4):235–265, 1997.
4. W. Bosma and M. Pohst. Computations with finitely generated modules over Dedekind domains. In *Proc. ISSAC'91*, pages 151–156. ACM, 1991.
5. H. Cohen. *A Course in Computational Algebraic Number Theory*. Springer, 1995.
6. H. Cohen. Hermite and Smith normal form algorithms over Dedekind domains. *Math. Comp.*, 65:1681–1699, 1996.
7. H. Cohen. *Advanced topics in computational number theory*. Springer, 2000.
8. J.-H. Evertse. Reduced bases of lattices over number fields. *Indag. Mathem. N.S.*, 2(3):153–168, 1992.
9. C. Fieker. Minimizing representations over number fields II: Computations in the Brauer group. *J. Algebra*, 3(322):752–765, 2009.
10. C. Fieker and M. E. Pohst. Lattices over number fields. In *Proc. ANTS II*, volume 1122 of *LNCS*, pages 147–157. Springer, 1996.
11. Y. H. Gan, C. Ling, and W. H. Mow. Complex lattice reduction algorithm for low-complexity full-diversity MIMO detection. *IEEE Trans. Signal Processing*, 57:2701–2710, 2009.
12. A. Hoppe. *Normal forms over Dedekind domains, efficient implementation in the computer algebra system KANT*. PhD thesis, Technical University of Berlin, 1998.
13. R. Kannan and A. Bachem. Polynomial algorithms for computing the Smith and Hermite normal forms of an integer matrix. *SIAM J. Comput.*, 8(4):499–507, 1979.
14. J. C. Lagarias, H. W. Lenstra, Jr., and C. P. Schnorr. Korkine-Zolotarev bases and successive minima of a lattice and its reciprocal lattice. *Combinatorica*, 10:333–348, 1990.
15. A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász. Factoring polynomials with rational coefficients. *Math. Ann.*, 261:515–534, 1982.
16. L. Lovász. *An Algorithmic Theory of Numbers, Graphs and Convexity*. SIAM, 1986. CBMS-NSF Regional Conference Series in Applied Mathematics.
17. V. Lyubashevsky and D. Micciancio. Generalized compact knapsacks are collision resistant. In *Proc. ICALP (2) 2006*, volume 4052 of *LNCS*, pages 144–155. Springer, 2006.
18. Magma. The Magma computational algebra system for algebra, number theory and geometry. Available at `http://magma.maths.usyd.edu.au/magma/`.
19. D. Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions. *Comput. Complexity*, 16(4):365–411, 2007.
20. D. Micciancio and S. Goldwasser. *Complexity of lattice problems: a cryptographic perspective*. Kluwer Academic Press, 2002.
21. R. A. Mollin. *Algebraic Number Theory*. Chapman and Hall/CRC Press, 1999.
22. H. Napias. A generalization of the LLL-algorithm over Euclidean rings or orders. *J. théorie des nombres de Bordeaux*, 2:387–396, 1996.
23. O. T. O'Meara. *Introduction to Quadratic Forms*, volume 117 of *Grundlehren der Mathematischen Wissenschaften*. Springer, 1963.
24. C. Peikert and A. Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In *Proc. TCC 2006*, volume 3876 of *LNCS*, pages 145–166. Springer, 2006.
25. C. Peikert and A. Rosen. Lattices that admit logarithmic worst-case to average-case connection factors. In *Proc. STOC 2007*, pages 478–487. ACM, 2007.
26. D. Stehlé, R. Steinfeld, K. Tanaka, and K. Xagawa. Efficient public key encryption based on ideal lattices. In *Proc. Asiacrypt 2009*, volume 5912 of *LNCS*, pages 617–635. Springer, 2009.