

Algorithms on Ideal over Complex Multiplication order

Paul Kirchner

ABSTRACT

We show in this paper that the Gentry-Szydlo algorithm for cyclotomic orders, previously revisited by Lenstra-Silverberg, can be extended to complex-multiplication (CM) orders, and even to a more general structure. This algorithm allows to test equality over the polarized ideal class group, and finds a generator of the polarized ideal in polynomial time. Also, the algorithm allows to solve the norm equation over CM orders and the recent reduction of principal ideals to the real suborder can also be performed in polynomial time. Furthermore, we can also compute in polynomial time a unit of an order of any number field given a (not very precise) approximation of it.

Our description of the Gentry-Szydlo algorithm is different from the original and Lenstra-Silverberg's variant and we hope the simplifications made will allow a deeper understanding.

Finally, we show that the well-known speed-up for enumeration and sieve algorithms for ideal lattices over power of two cyclotomics can be generalized to any number field with many roots of unity.

1. Introduction

Recently, an algorithmic study of lattices was made necessary by new cryptographic proposals. Indeed, lattice-based cryptography has several advantages : it seems post-quantum secure, allows to build a lot of cryptosystems and enjoys a worst-case/average-case reduction over any lattice problem [1]. Yet, the schemes are slow and have large keys so that most designers turned towards *ideal*-lattice based cryptography, based on the Ring-LWE [2]. However, ideal lattices are less studied and the added algebraic structure might allow significant gains with respect to the same problem for a random lattice. Some suggested to remove a part of the algebraic structure by choosing a polynomial with a large Galois group [†].

We show that many algorithms previously discovered for cyclotomic fields can be generalized to CM orders with little loss. In particular, a large part of the present paper is dedicated to the Gentry-Szydlo algorithm [3]. This algorithm, given an ideal generated by some $v \in \mathbb{Z}[X]/(X^n - 1)$ and the autocorrelation of v , finds v up to the (few) root of unity in polynomial time. It was first used to break the NTRU signature scheme [3]. Lenstra and Silverberg [4] then extended it to (essentially) product of cyclotomic rings, and made it rigorous. According to [4], Gentry referred to Gentry-Szydlo's algorithm as "a rather crazy, unusual combination of LLL with more 'algebraic' techniques", while Smart viewed it as magic. We hope the simplifications made to the algorithms[‡], as well as our effort to see the exact conditions under which Gentry-Szydlo's algorithm can be run will help to removed this "dark magic" aspect.

This new algebraic structure extends the CM order, and can be applied to any order. We give here a particular case of our main theorem Theorem 4 :

2000 *Mathematics Subject Classification* 00000.

[†]<http://blog.cr.yp.to/20140213-ideal.html>

[‡]In a recent conference, no less than 9 hours and 33 minutes were dedicated to Gentry-Szydlo's algorithm and Lenstra-Silverberg's modifications.

Theorem 1. *Under GRH or using randomness, we can test in polynomial time if (I, r) where I is an ideal of a CM order is equal to some $((v), v\bar{v})$ for some invertible v and find such a v .*

It has several applications, that we extend in section 4. The most useful is the reduction of searching for short vectors in ideals over CM orders to ideals over their real suborder, which is enough to attack the Smart-Vercauteren FHE scheme [5, 6], the GGH multilinear map scheme [7], Soliloquy [8] and the Boneh-Freeman homomorphic signatures [9]. Remark that a quantum algorithm was recently discovered which breaks all these schemes in polynomial time [10]. This problem can be mitigated by switching to other number fields, since the attack crucially relies on having a very orthogonal basis of the unit lattice [11]. It (usually) divides by two the dimension of the lattice, which corresponds to reduce the running time to its square root. For most cryptosystems based on ideal lattices, the ideals are *not* principal, so this attack does not apply [†]. However, if it can be modified to attack Ring-LWE, which is decoding in the lattice $\begin{pmatrix} q & a \\ 0 & 1 \end{pmatrix} \mathcal{O}^2$ for some uniform $a \in \mathcal{O}/q$, this would break almost all published design in practical time [‡]. Another application is to solve the norm equation, though not in polynomial time. A last application is the heuristic ability to solve bounded distance decoding in polynomial time over the unit lattice with approximation factor around $n^{O(\log(\log n))}$, even though the unit lattice is not known. Note however that units may not be of size polynomial in the discriminant. A recent attack against GGH uses Gentry-Szydło's algorithm in a similar way, our extension allows it to work over any number field [12].

While CM orders and their polarized class group were introduced for the study of abelian varieties, such as some elliptic curves, this paper does not use any algebraic geometry. Also, this result may indicate that polarized class group are more tractable than the class group from a computational point of view. In particular, given some generators of a subgroup of the polarized class group, either we have an incremental multiset hash function [13] (or a hash function), or there is an efficient bijection between this abelian group and its normal form. Finally, testing equality in the polarized class group is related to isotopy of knots, see [14, Knot theory].

In appendix, we show that for ideal lattices, we can accelerate standard algorithms for searching the shortest vector by a polynomial factor with respect to their (exponential) complexity on generic lattices.

2. Preliminaries

2.1. Basics

The norm of a vector or a matrix $\|A\|$ is the Frobenius form, i.e. $\|A\|^2 = \sum_{i,j} A_{i,j}^2$. The binary logarithm is denoted by \log and \ln is the neperian logarithm. All indices start from zero.

DEFINITION 1. An order \mathcal{O} is a commutative unitary ring, whose additive group is isomorphic to \mathbb{Z}^n . We denote \mathcal{O}^* the group of units in the order. The trace of a , $\text{tr}(a)$, is the trace of the endomorphism $x \mapsto ax$.

[†]Remark that the schemes are proven to be at least as secure as ideal lattices, current attacks need to find short vectors over a module of rank two.

[‡]We consider therefore that a wise design based on ideal lattices should have at least 256 bits of security if the users want to keep data secure for several decades with high probability.

Notice that usually, an order is defined with respect to an algebra, this is not the case here.

For the rest of the paper, we will work with an order, always denoted by \mathcal{O} , whose corresponding algebra $\mathbb{Q} \otimes \mathcal{O}$ is denoted by E . We refer to elements of $E^* \cap \mathcal{O}$ as *invertible*. An order is given to an algorithm by n , followed by the $n(n+1)/2$ products of two basis elements, written as integer linear combinations of basis elements. We define the height function of $x \in \mathbb{Z}$ to be $h(x) = 2 \log(2 + |x|)$, and extend it to the rationals $h(p/q) = h(p) + h(q)$ where p and q are coprime. The height of a matrix or a list is the sum of the heights of the components. It represents the number of bits used to describe an object.

DEFINITION 2. An integral ideal I over an order \mathcal{O} is an \mathcal{O} -submodule. Its norm $N(I)$ is defined as the cardinal of \mathcal{O}/I . A (fractional) ideal is an integral ideal, up to a rational number. An ideal is said to be *invertible* if there exists a fractional ideal J such that $IJ = \mathcal{O}$.

An ideal is described by a \mathbb{Z} -basis. It is well known that ideals can be multiplied in polynomial time [15, Section 4.7]. It is clear that if $v \in \mathcal{O}$ is invertible, then (v) is invertible.

DEFINITION 3. A positive-definite quadratic form $f : E \rightarrow \mathbb{Q}$ over a \mathbb{Q} vector space E is a function such that $(a, b) \mapsto (f(a+b) - f(a-b))/4$ is bilinear, and $f(x) > 0$ for all non-zero x . Its determinant is the determinant of the corresponding bilinear form.

While most algorithms in the litterature are presented with lattices, they can usually be transformed into algorithms over quadratic forms and we will be forced here to use the quadratic form.

NOTATION 1. The group of the roots of units in an order \mathcal{O} is denoted $\mu(\mathcal{O})$.

NOTATION 2. The p -Sylow of a finite abelian group G is denoted by G_p .

NOTATION 3. Given a number field F , we denote its ring of integers by \mathbb{Z}_F and its discriminant by d_F .

Theorem 2. Let E be a commutative algebra of dimension n over the field K . Then there is a unique $E' = \prod_i F_i$ where F_i is a finite extension of K such that for all $x \in E$, there is a unique (a, b) with a nilpotent, $a + b = x$ and $b \in E'$.

Proof. This is a direct consequence of the Artin-Wedderburn theorem. □

On first read, the reader may assume that $\mathcal{O} = \mathbb{Z}[\alpha]$ for some algebraic α .

An algorithm has a negligible probability of failure if the probability of failure is bounded by $2^{-\Omega(n)}$.

2.2. Advanced definitions

DEFINITION 4. A split CM order is an order \mathcal{O} , a norm function $N : E \rightarrow \mathcal{N}$ where \mathcal{N} is a commutative semigroup, and a trace function $\text{tr} : \mathcal{N} \rightarrow \mathbb{Q}$ such that :

- N is a morphism as a semigroup, i.e. $N(xy) = N(x)N(y)$ for all $x, y \in E$
- $x \mapsto \text{tr}(N(x))$ is a positive-definite quadratic form.

Without loss of generality, we will impose furthermore that $(\text{tr}(N(x+y)) - \text{tr}(N(x-y)))/4$ is an integer for all $x, y \in \mathcal{O}$.

Norms will be used in our algorithms in a black-box manner, so we define the height of a norm to be the number of bits used to represent it. We describe a split CM order with functions which run in polynomial time of the height of the input and $h(\mathcal{O})$ for N , multiplication and inversion in \mathcal{N} , as well as the trace.

We now give several examples of split CM orders. The simplest one, not interesting for our purpose, is where the norm is the identity function and the trace is any positive-definite quadratic form.

DEFINITION 5. A CM order is an order \mathcal{O} equipped with an automorphism $x \mapsto \bar{x}$ which is an involution, and such that $\text{tr}(x\bar{x})$ is a positive-definite quadratic form. The real suborder \mathcal{O}^+ is defined as $\{x = \bar{x}; x \in \mathcal{O}\}$, the imaginary lattice \mathcal{O}^- is $\{x = -\bar{x}; x \in \mathcal{O}\}$.

We can easily build a split CM order from a CM order by taking $N(x) = x\bar{x}$ and the trace function as the standard trace. For $\alpha \in \mathbb{C}$ some algebraic number, if $\bar{\alpha} \in \mathbb{Z}[\alpha]$, then $\mathbb{Z}[\alpha]$ equipped with the conjugation is a CM order. In particular, $\mathbb{Z}[\zeta_n]$ where $\zeta_n = \exp(2i\pi/n)$ and n a positive integer is a CM order. In these cases, the norm function is the *algebraic norm* over the real subfield, which must not be confused with the corresponding *geometric norm*. We define a cyclotomic field to be any $\mathbb{Q}[\zeta_n]$. Also, $\mathbb{Z}[X]/(X^n - 1)$ equipped with $X \mapsto X^{n-1}$ is a CM order and the norm is the autocorrelation.

Finally, $\mathbb{Z}[\alpha]$ with the norm $\alpha \mapsto \alpha\bar{\alpha}$ is a *split* CM order. We can generalize this construction using Theorem 2 to give a non-trivial split CM order from any order.

DEFINITION 6. A *polarized ideal* is a pair (I, r) where I is a fractional invertible ideal and $r \in \mathcal{N}^*$. The determinant of a polarized ideal (I, r) is the determinant of $x \mapsto \text{tr}(N(x)/r)$ defined over I . The *polarized ideal group* is the group of all polarized ideals. The *principal polarized ideal group* is the group of all $((v), N(v))$ for all $v \in E^*$, and is a subgroup of the polarized ideal group.

We can now state the *informal* problem we (partly) solve :

PROBLEM 1. Given a split CM order and a polarized ideal (I, r) , determine all invertible solutions $v \in \mathcal{O}$ to $(I, r) = ((v), N(v))$.

Theorem 3. *If $N(v) = N(w)$ for $v, w \in \mathcal{O}$ invertibles, then $v/w \in \mu(\mathcal{O})$.*

Proof. $N((v/w)^k) = 1$ for all integer k , but since $x \mapsto \text{tr}(N(x))$ is definite positive, $(v/w)^k$ can take only a bounded number of values. Hence, v/w is a root of unity. \square

This implies that v is defined up to a group of roots of unity.

Though we will not need the following definition, it gives a nice interpretation of the Gentry-Szydlo algorithm.

DEFINITION 7. Let Δ be the determinant of $(\mathcal{O}, N(1))$. The polarized ideal class group is the maximum group of polarized ideals (I, r) of determinant Δ , modulo the principal polarized ideal group.

Lemma 1. *The only nilpotent of a CM order \mathcal{O} is 0.*

Proof. Let x be a nilpotent. Then $x\bar{x}$ is also a nilpotent, so that $\text{tr}(x\bar{x}) = 0$. Since $x \mapsto \text{tr}(x\bar{x})$ is definite, $x = 0$. \square

The following definition is the main conceptual novelty with respect to previous descriptions of Gentry-Szydło’s algorithm, and allows to greatly simplify its exposition.

DEFINITION 8. The formal products over E , $\mathbb{Z}^{(E)}$, is the additive group of function from E to \mathbb{Z} which are non-zero on a finite set, called support. The evaluation of $f \in \mathbb{Z}^{(E)}$, denoted by $[f]$, is defined by

$$\prod_{x \in E} x^{f(x)}.$$

We call the elements of the support the base, and $f(x)$ is called the exponent of x . The group law is denoted multiplicatively.

For clarity reasons, we will consider an element $x \in E$ to be also the function of $\mathbb{Z}^{(E)}$ equal to zero everywhere except in x where it is equal to one.

DEFINITION 9. We call a group of polarized ideals to be *reducible* if for all x , $\text{tr}^{-1}(\{x\})$ is finite.

We call a group of polarized ideals to be *poly-reducible* if there is a constant c such that for all (I, r) in the group with I integral, $h(r) \leq (h(\mathcal{O}) + \log(\text{tr}(r)))^c$.

It can be checked that in all our examples of split CM orders, the full group is poly-reducible. From a mathematical perspective, reducibility is more pertinent, as it is enough to prove that a reducible group modulo the principal polarized ideals is compact. Yet, we need the stronger condition for efficient computation over this finite group, see subsection 3.2. However, if we extend the norms to $\mathcal{N} \times \mathbb{Q}$, the polarized ideal class group is clearly non compact.

The polarized class group of a number field with the norm $x\bar{x}$ is in bijection with the product of the class group and $\mathbb{R}^{r+s-1}/\Lambda$ where Λ is the log unit lattice with $E \otimes \mathbb{R} \simeq \mathbb{R}^r \mathbb{C}^s$. It has also been named Arakelov class group by others [16].

Lemma 2. *The principal polarized ideal group is poly-reducible.*

Proof. Let G be the Gram matrix of $x \mapsto \text{tr}(N(x))$ over \mathcal{O} . Let $G = L^t L$ be its corresponding Cholesky decomposition, that is L is an upper triangular matrix. Then, we can write $L_{i,i}^2$ as a quotient of two non-zero integer determinants of submatrices of G , so that the Hadamard bound gives $L_{i,i} \geq \|G\|^{-n/2}$. Therefore, for any $v \in \mathcal{O}$ which is invertible, there exists i such that

$$|v_i|^2 \geq v^t G v \|G\|^{-1/n} = \text{tr}(N(v)) \|G\|^{-n}.$$

This implies that $h(v) \geq h(\text{tr}(N(v))) - h(\|G\|^n)$. Now, $x \mapsto \text{tr}(N(x))$ is computed in time polynomial in $h(\mathcal{O}) + h(x)$, so there exists c_0 such that $h(\|G\|^n) \leq h(\mathcal{O})^{c_0}$. Also, there exists c_1 such that $h(1/N(v)) \leq h(v)^{c_1}$. We conclude that there exists c_2 such that $h(N(1/v)) \leq (h(\mathcal{O}) + h(\text{tr}(N(v))))^{c_2}$. \square

Lemma 3. *A group is poly-reducible if and only if, there exists a constant c such that for all polarized ideals (I, r) in the group with $\mathcal{O} \subset I$, $h(r) \leq (h(\mathcal{O}) + \log(\text{tr}(1/r)))^c$.*

Proof. Remark that $I \subset \mathcal{O}$, $1 \in I$, and I^{-1} is integral are equivalent. Thus, because inversion runs in polynomial time, $h(1/r) = h(r)^{\Theta(1)}$ and the result follows. \square

3. Gentry-Szydlo algorithm

We recall that an order \mathcal{O} has a corresponding algebra $\mathbb{Q} \otimes \mathcal{O}$ denoted by E , which contains a maximal product of number fields E' . This section is devoted to prove the following theorem : **Theorem 4.** *Given a polarized ideal (I, r) of a split CM order, if E' is a product of cyclotomic field, or the Generalized Riemann Hypothesis (GRH) is true, or we have access to randomness, we can find $v\omega \in \mu(E)\mathcal{O}$ such that $I = (v)$ or prove there is no such $v\omega$ with $(I, r) = ((v), N(v))$. We can also do this unconditionnally in time $2^{n^{(1+o(1))/\log \log n}}$. Furthermore, if \mathcal{O} is a CM order, we can find $v \in \mathcal{O}$ such that $(I, r) = (v, N(v))$ under the same conditions.*

In the next subsections, all algorithms are authorized to fail if there is no solution to $(I, r) = ((v), N(v))$. Since the output of the algorithm can easily be checked, we may assume that there is some solution v , except for the analysis of the complexity. Also, in the case of a CM order, we can assume v is invertible by working in $\mathcal{O}/(\mathbb{Q} \otimes I)$.

We fix Δ to be the determinant of $(\mathcal{O}, N(1))$. Remark that $h(\Delta) = h(\mathcal{O})^{O(1)}$.

The hero of our story will be the following group :

DEFINITION 10. Given a poly-reducible group G where all $(I, r) \in G$ are of determinant Δ , its *compactification* is the group of all (I, r, s) where $s \in \mathbb{Z}^{(E)}$ such that $[s]$ is invertible ; modulo the subgroup of all $((1/x), N(1/x), x)$ for all $x \in E^*$.

Indeed, it replaces Gentry-Szydlo's cumbersome "polynomial chains", and Lenstra-Silverberg's chain of tensor multiplication maps. Compactification is to be understood in its computer science meaning, i.e. a short representation, and not in a topological sense.

In subsection 3.1, we show that using LLL, we can reduce the description of I and r to a polynomial value, independent of I and r . In subsection 3.2, we show how to compute a power of $(I, r, 1)$ in polynomial time. We then use this powering algorithm in subsection 3.3 to compute the image over a field of E' of some high power of v , as a formal product. By combining various high powers, we show in subsection 3.4 how to compute the image over a field of v , up to a root of unity. Finally, we explain in subsection 3.5 how to compute the nilpotent part of v and one root of unity.

The four different cases in the theorem are introduced in subsection 3.4.

3.1. Reduction

Theorem 5. *Given a positive-definite integer matrix G of dimension n , we can compute in polynomial time a unimodular integer matrix U such that $U^t G U$ has entries bounded by $n2^n \det(G)$.*

Proof. See [17] where the Gram-Schmidt orthogonalization is replaced by Cholesky decomposition. Its output verifies $U^t G U = L^t L$ where L is an upper-triangular matrix, such that $L_{i+1, i+1} \geq L_{i, i}/\sqrt{2}$ and $|L_{i, j}| \leq L_{j, j}$ for all i, j . Let $m = \arg \max_i L_{i, i}$. Then, $\prod_{i < m} L_{i, i}^2$ is a positive integer, since it is the determinant of the corresponding upper-left submatrix of $U^t G U$. Also, for any $j \geq m$, $L_{j, j} \geq L_{m, m} 2^{(m-j)/2}$. It implies that

$$\Delta = \prod_i L_{i, i}^2 \geq L_{m, m}^{2(n-m)} 2^{-(n-m)(n-m-1)/2}.$$

We deduce that $L_{m, m} \leq 2^{(n-m-1)/2} \Delta^{1/(2(n-m))} \leq 2^{n/2} \sqrt{\Delta}$. Using $|L_{i, j}| \leq L_{j, j}$ gives the result. \square

Lemma 4. *Given m matrices A_i in $M_n(\mathbb{Q})$ such that a linear combination is invertible, we can find $x_i \in \mathbb{Z}$ and $|x_i| \leq n$ such that $\sum_i A_i x_i$ is invertible in polynomial time.*

Proof. If $m = 1$, we output $x_0 = 1$. Else, we compute r the rank of A_0 and a set I of r rows and J a set of r columns such that the restriction of A_0 to these lines and columns is invertible. We then recursively find the $x_i, i \geq 1$ corresponding to B_i , the restriction of A_i to the complements of I and J . Finally, we search through all x_0 from 0 to n and output the first solution.

Without loss of generality, we may analyze this algorithm by assuming A_0 is diagonal, with r ones followed by zeroes on the diagonal. By assumption $\det(\sum_i A_i X_i) = \sum_j X_0^j P_j(X_1, \dots, X_{m-1})$ is non-zero. Then, $P_r = \det(\sum_{i \geq 1} B_i X_i)$, and is also non-zero. By our choice of x_i , that is $\det(\sum_{i \geq 1} B_i x_i) \neq 0$, we have that $\det(A_0 X + \sum_{i \geq 1} A_i x_i) = \sum_j X_0^j P_j(x_1, \dots, x_{m-1})$ is a non-zero univariate polynomial of degree at most n , so that it has at most n roots, which guarantees the algorithm will find a solution. \square

Theorem 6. *Given (I, r) in a poly-reducible group of determinant D , we can find in polynomial time $x \in E^*$ and a basis C of I/x , such that $h(C) + h(r/N(x)) = (h(\mathcal{O}) + h(D))^{O(1)}$. Also, $\mathcal{O} \subset I$.*

Proof. Without loss of generality, we can assume that I is an integral ideal, of basis $A = (e_i)_{i < n}$. We then compute the Gram matrix corresponding to $x \mapsto \text{tr}(N(x)/r)$, that is $G_{i,j} = (\text{tr}(N(e_i + e_j)/r) - \text{tr}(N(e_i - e_j)/r))/4$, which we can do in polynomial time. We use Theorem 5 to compute U such that $U^t G U$ is bounded by $n2^n D$. Now, we compute $B = AU = (b_i)_{i < n}$, and use Lemma 4 with the multiplication matrices of b_i , to find $y = \sum_i x_i b_i$ invertible. Finally, we return y and B/y .

The running time is clear. Because $II^{-1} = \mathcal{O}$, there exists $z \in \mathbb{Z}^n$ such that Az is invertible, so that the condition of Lemma 4 is fulfilled.

We now have $\text{tr}(N(y)/r) = \sum x_i^2 \text{tr}(N(x_i)/r) \leq n^4 2^n D$. Since $h(\mathcal{O}) \geq n$ and $y \in I$ implies $I/y \subset \mathcal{O}$, $h(r/N(y)) = (h(\mathcal{O}) + h(D))^{O(1)}$ follows from Lemma 3. Also, for all i , we have $\text{tr}(N(b_i/y)N(y)/r) \leq n2^n D$ so that $h(N(b_i/y)N(y)/r) = (h(\mathcal{O}) + h(D))^{O(1)}$. Now we can compute $N(b_i/y)$ from $N(b_i/y)N(y)/r$ and $r/N(y)$, and hence $\text{tr}(N(b_i/y))$ in time $(h(\mathcal{O}) + h(D))^{O(1)}$. It implies $\log \text{tr}(N(b_i/y)) = (h(\mathcal{O}) + h(D))^{O(1)}$. Then, with $H = L^t L$ the Gram matrix of $x \mapsto \text{tr}(N(x))$ over \mathcal{O} and its Cholesky decomposition, we have $\log(L(b_i/y)_j) = (h(\mathcal{O}) + h(D))^{O(1)}$ for all j . Since $L_{j,j} \geq \|H\|^{-n/2}$ we have $L_{j,j}^{-1} \leq \|H\|^{n/2}$, we deduce $h(b_i/y) = (h(\mathcal{O}) + h(D))^{O(1)}$. \square

3.2. Powering

Theorem 7. *Given $(I, r) = ((v), N(v))$ a principal polarized ideal and an integer e , we can compute $(I, r, 1)^e$ over the compactification of the principal polarized ideal group in polynomial time. Furthermore, the norm which is outputted have a height $h(\mathcal{O})^{O(1)}$ and the ideal contains \mathcal{O} .*

Proof. If $e = 0$, we return $(\mathcal{O}, N(1), 1)$. If e is even, we recursively compute $(I, r, 1)^{e/2} = (K, u, s)$, use Theorem 6 with (K^2, u^2) which returns an ideal C and $x \in E^*$, and outputs $(C, u^2/N(x), s^2 x)$. Else, we recursively compute $(I, r, 1)^{e-1} = (K, u, s)$, use Theorem 6 with (KI, ur) which returns an ideal C and $x \in E^*$, and outputs $(C, ur/N(x), sx)$. If at any point, the height is too large with respect to the bounds given by Theorem 6, we fail.

Correctness is clear. By induction, the output is reduced so its height without the formal product is bounded by $h(\mathcal{O})^{O(1)}$. It implies that the height of the bases in the formal product is bounded by $h(\mathcal{O})^{O(1)}$, while the exponents are bounded by e and the cardinal of the support is bounded by $O(\log(e))$. Therefore, the algorithm runs in polynomial time. \square

Note that we can, in fact, compute any circuit over the compactification of any poly-reducible group in polynomial time. In particular, we may use shorter addition chains.

3.3. Recovery of a high power of v

We fix in this subsection a maximal ideal \mathfrak{m} of E , and suppose a \mathbb{Q} -basis is given. Remember that E/\mathfrak{m} is a number field, which we denote F . We define \mathcal{O}' as a suborder of \mathcal{O}/\mathfrak{m} .

Lemma 5. *Given an invertible integral ideal \mathfrak{a} of \mathcal{O}' , we have for all $x \in \mathfrak{a}^k$, $x \neq 0$, $h(x)/n \geq k \log(N(\mathfrak{a})) - O(h(\mathcal{O}'))$.*

Proof. It is well known that the norm is multiplicative for invertible ideals, see for example [15, Proposition 4.6.8]. Also, for $x \in \mathfrak{a}^k$, $x \neq 0$, we have $(x) \subset \mathfrak{a}^k$ so that $N((x)) \geq N(\mathfrak{a})^k$. Then, $N((x))$ is also the absolute value of the determinant of the multiplication by x . Using the Hadamard bound, we have $\log(N((x))) \leq n(\log(\|x\|) + h(\mathcal{O}'))$. \square

Lemma 6. *Given an invertible matrix A with $\lambda_1 = \min_{x \in \mathbb{Z}^n - \{0\}} \|Ax\|$, and c such that there exists $y \in \mathbb{Z}^n$ with $\|Ay - c\| \leq 2^{-n} \lambda_1$, we can recover y in polynomial time.*

Proof. This was proven by Babai [18, Theorem 3.1], as an application of LLL. \square

Lemma 7. *Given a prime ideal \mathfrak{p} of \mathcal{O}' with a prime number $p \in \mathfrak{p}$, we have $x^{(N(\mathfrak{p})-1)p^k} \in 1 + \mathfrak{p}^{k+1}$ for any invertible $x \in \mathcal{O}' - \mathfrak{p}$ and k a positive integer.*

Proof. We use induction on k . $\mathcal{O}'/\mathfrak{p}$ is a field, so $x^{N(\mathfrak{p})-1} \in 1 + \mathfrak{p}$ for any invertible $x \in \mathcal{O}' - \mathfrak{p}$. Now, let $y \in \mathfrak{p}^k$. If we develop $(1+y)^p - 1 - y^p$, $p \in \mathfrak{p}$ divides all binomial coefficients and the power of y is at least one, so $(1+y)^p - 1 - y^p \in (p)\mathfrak{p}^k \subset \mathfrak{p}^{k+1}$. Since $p \geq 2$, we also have $y^p \in \mathfrak{p}^{k+1}$ and hence $(1+y)^p \in 1 + \mathfrak{p}^{k+1}$. \square

Theorem 8. *Given $(I, r) = ((v), N(v))$ a principal polarized ideal and \mathfrak{p} be an invertible prime ideal of \mathcal{O}' with $p \in \mathfrak{p}$ a prime integer and $v \notin \mathfrak{p} + \mathfrak{m}$. Then, we can output k and s in polynomial time such that $v^{(N(\mathfrak{p})-1)p^k} = [s]$ modulo \mathfrak{m} .*

Proof. Let $e = (N(\mathfrak{p}) - 1)p^k$ for some integer k . We first compute using Theorem 7 $(I, r, 1)^e = (J, a, s)$. We know that $1/a = N([s]/v^e)$, $h(1/a) = h(\mathcal{O})^{O(1)}$ and J^{-1} is integral, and therefore $[s]/v^e$ is an invertible integer, whose height is in $h(\mathcal{O})^{O(1)}$. Using Lemma 5, there exists a constant c such that with $k = h(\mathcal{O})^c$, any non zero element in \mathfrak{p}^{k+1} has a coordinate larger than 2^n times any coordinate of $[s]/v^e$.

Remark that we can compute in polynomial time \mathfrak{p}^{k+1} . We then run Lemma 6 with the basis of \mathfrak{p}^{k+1} and $[s]$ modulo \mathfrak{m} and \mathfrak{p}^{k+1} ; we call the result lifted to \mathcal{O} , c . Because of the previous lemma, since $v \notin \mathfrak{p} + \mathfrak{m}$, $v^e = 1$ modulo \mathfrak{m} and \mathfrak{p}^{k+1} . Hence, c and $[s]/v^e$ differs by an element of \mathfrak{p}^{k+1} modulo \mathfrak{m} but by definition of k , Lemma 5 and Lemma 6, it must be zero. Therefore, we return s/c . \square

Though v may not be unique, the given power is.

3.4. Recovery of v over a field

Lemma 8. *Given an order \mathcal{O}' over a number field, we can compute $\mathcal{O}' \subset \mathcal{O}'_p$ and integral invertible ideals \mathfrak{p}_i of \mathcal{O}'_p such that*

$$p\mathcal{O}'_p = \prod_i \mathfrak{p}_i^{e_i}$$

in time which is polynomial in the size of the input, and p . Using randomness, we can do the same in polynomial time.

Proof. See [15, Sections 6.1 and 6.2], where the only randomness used is for factoring polynomials modulo p . \square

Lemma 9. *Let e_i be integers, s_i be formal products such that $[s_i] = v^{e_i}$ modulo \mathfrak{m} and $I = (v)$. Let g be the greatest common divisor of the e_i . Then, we can find in time polynomial in the size of the input and g , an element w such that w/v reduced modulo \mathfrak{m} is a root of unity.*

Proof. We can compute in polynomial time by applying a Hermite normal form algorithm over e , a vector of integers u_i such that $\sum_i e_i u_i = g$. We search for some prime p such that the bases in the support of all s_i are invertible modulo p . Since $a \in \mathcal{O}$ is invertible modulo p if and only if its norm is divisible by p , there exists a p which is bounded by a polynomial of the size of the input which works and we can find it in polynomial time. We then compute $s = \prod_i s_i^{u_i}$, and evaluate $[s]$ modulo \mathfrak{m} and some sufficiently high power of p (but polynomial in $h(I)^g$), so that we recover y , congruent to v^g modulo \mathfrak{m} . Finally, we factor $X^g - y$ in polynomial time over F (see [15, Section 3.6.2]), and if there exists a linear factor $X - w$, we output w . Else, we fail.

Remark that v reduced modulo \mathfrak{m} is a root of $X^g - y$, and the quotient of two roots must be a g -th root of unity. \square

Lemma 10. *If F is a cyclotomic field, we can choose in polynomial time two primes such that the gcd of the corresponding exponents in Theorem 8 is polynomial.*

Proof. If $F = \mathbb{Q}[\zeta_m]$, then we choose the first two primes which split in linear factors, which is equivalent to being congruent to one modulo m . We use the latest version of Linnik's theorem, which says that the smallest prime congruent to a modulo k is $O(k^5)$ [19, Theorem 2.1]. Then, the smallest prime congruent to one mod m , p verifies $O(m^5)$. Let r be the smallest prime which does not divide $p - 1$ or m , we know that $r = O(\log(m))$. We then define a to be the element of $\mathbb{Z}/(rm)$ congruent to one modulo m and to $1 + p$ modulo r . Thus, we can define q to be the smallest prime congruent to a modulo rm , and $q = O((m \log m)^5)$, $q > p$. Finally, for any α, β , $\gcd((p - 1)p^\alpha, (q - 1)q^\beta) = \gcd((p - 1)p^\alpha, q - 1) < q$. \square

Remark that we can make the gcd equal to m using the technique of [4, Proposition 4.5], but the exponent then becomes 50.

Lemma 11. *Let $H = F[\zeta_m]$ be a Galois extension of F and $m \notin \mathfrak{p}$ be some invertible prime ideal of \mathcal{O}'_p . Then, $m|N(\mathfrak{p}) - 1$ if and only if $\mathfrak{p}\mathbb{Z}_K$ splits completely over H .*

Proof. Using the properties of the conductor ideal, it is a standard fact that without loss of generality, we can assume $\mathcal{O}'_p = \mathbb{Z}_K$. Also, \mathfrak{p} does not ramify. Remark that for any prime \mathfrak{q} above \mathfrak{p} , we have $x^{N(\mathfrak{q})} = x$ over $\mathbb{Z}_H/\mathfrak{q}$. Hence, \mathfrak{p} splits completely is equivalent to $x^{N(\mathfrak{p})}$ fixes $\mathbb{Z}_H/\mathfrak{p}$. But $\zeta_m \notin \mathfrak{p}$ so that it is equivalent to fixing ζ_m , which is $m|N(\mathfrak{p}) - 1$. \square

Theorem 9. *Given a polarized principal ideal $(I, r) = ((v), N(v))$ and \mathfrak{m} , if F is a cyclotomic field, or GRH is true, or we have access to randomness, we can find w such that v/w modulo \mathfrak{m} is a root of unity in polynomial time. We can also do this unconditionnally in time polynomial in $2^{n^{(1+o(1))/\log \log n}}$ and the size of the input.*

Proof. The algorithm consists in applying Lemma 8 to generate the input of Theorem 8, and we combine the outputs using Lemma 9. The crux of the matter is to bound the greatest common divisor of the exponents used. The case of cyclotomic field is easily treated with Lemma 10. Indeed, for each prime p , then either the image of v is divisible by p so the factor can then be removed and this happens at most a polynomial number of times, or we can use some \mathfrak{p} above p .

Under GRH, we show that using all prime ideals \mathfrak{p} of inertia degree one above all primes p smaller than a polynomial will work. Using the previous lemma and [20, Théorème 4], we have that there are a polynomial number of primes smaller than some polynomial who have a prime ideal above it of inertia degree one. Hence, we can find two amongst them which does not divide v , and m will be $\gcd((p - 1)p^f, (q - 1)q^e) < q$ if $q > p$. Therefore, the algorithm runs in polynomial time.

For the other algorithms, we first start by trying for all the $2n^2 + n + 1$ smallest primes any ideal above it such that the image of v is not in the ideal. Each time a prime is detected as dividing the image of v , it can be factored out. Let m be the current greatest common divisor of the exponent used. Remark that there are n inertia degree possible, so that by the pigeonhole principle, there exist one degree d with $2n + 2$ corresponding primes p_i . Let $p^k \mid m$ with k positive and p prime. By removing p from the list of p_i , we have that for $2n + 1$ distinct primes p_i , $p^k \mid p_i^d - 1$. This implies that either $p^k = O(n^2 \log n)$, or there are $2n + 1$ elements of order dividing d in $(\mathbb{Z}/(p^k))^*$. If p is odd, the group is cyclic so that $2n + 1 \leq d$ which is absurd. Else $p = 2$, the group has at most $2d$ elements of order dividing d , and $2n + 1 \leq 2d$ which is also absurd. Hence at this point, $p^k = O(n^2 \log n)$ for all $p^k \mid m$.

For the unconditional algorithm, we continue to do so for the first primes. Let P be the largest prime used, which we will fix later to some function in $2^{n^{\Theta(1/\log \log n)}}$. Using the same argument, we have that for the new m and some d , that either $m \leq O(Pn^2 \log n)$ or there exist a $m' \mid m$ with $m' = O(Pn^2 \log n)$ and odd such that all primes considered except possibly $O(n^2 \log n)$ of them are of order dividing d in $(\mathbb{Z}/(m'))^*$, and these primes are distinct elements modulo m' . Thus, there are $\Omega(P/n/\log P)$ elements of order dividing d in $(\mathbb{Z}/(m'))^*$, which is a proportion of $\Omega(1/n^3/\log^3(n)n^{-\Theta(1/\log \log n)}) = \Omega(1/n^4)$. Decomposing $(\mathbb{Z}/(m'))^*$ as a product of cyclic groups, we first consider the groups where all elements are of order dividing d , that is $(p-1)p^k \mid d$, for $(p-1)p^k \mid m'$ with p prime. Let $(q-1)q^r$ be another cyclic group of equal order. Then, without loss of generality $q \geq p$ and if k is positive, $q \mid (p-1)p^k$ which is not possible. Hence q is unique with respect to (p, k) and using Wiegert's theorem, we deduce that there are at most $d^{(1+o(1))/\log \log d}$ cyclic groups where all elements are of order at most d . But because of Lagrange's theorem, there are at most $O(\log(n))$ cyclic groups where not all elements are of order dividing d . Hence, for

$$P \geq (n^2 \log n)^{O(\log(n)) + d^{(1+o(1))/\log \log d}}$$

this is absurd. The total running time is therefore in $2^{n^{(1+o(1))/\log \log n}}$.

Using randomness, we sample a polynomial number of integers smaller than $B^{1/f}$ with $B = \exp(O(\log(d) \log(\log d) \log(\log \log d)))$ where $d = d_{F[\zeta_q]}$ for some prime power q dividing m such that F has no q -th root of unity. If the number divides m , which happens with negligible probability, we restart. Else, if the number is prime, we use all prime ideals above this prime. The probability that some ideal divides v is negligible. We use all f from one to n .

[20, Théorème 3] shows that the likelihood of a complete split is then $\epsilon \leq c/[F[\zeta_q] : F]$ for a uniform prime of norm below B and some universal constant c . Assume $\epsilon < 1/2$. Then, there are $k \geq 1$ inertia degrees f such that the likelihood for a prime below $B^{1/f}$ to have prime above it with an inertia degree f is at least $(1 - \epsilon)/n$. Further, the likelihood that a uniform prime of norm below B with one of these inertia degrees not to completely split is at least $(1 - \epsilon)k/n$. Therefore, there exist an inertia degree f such that a uniform prime of norm below B of inertia degree f will not completely split with probability at least $(1 - \epsilon)/n$, and a uniform prime below $B^{1/f}$ has a probability at least $(1 - \epsilon)/n$ to have a prime of inertia degree f . We deduce that the above procedure takes polynomial time to find a prime ideal which does not split completely with high probability. In case we find a prime not have a complete split, because $p \nmid m$, we have $\gcd(m, (p^f - 1)p^e) \leq m/2$ so that with high probability, after a polynomial number of tests, we have that $p^k \mid m$ only if F has a $p^{k/c}$ root of unity. We deduce then that $m = O(n \log \log n)^c$, so that the running time is polynomial. \square

3.5. Recovery of v

Theorem 10. *Given E , we can compute \mathfrak{m}_i such that E is the sum of a nilpotent vector space and $\prod_i E/\mathfrak{m}_i$ where \mathfrak{m}_i are maximal ideals of E in polynomial time. We can also compute*

the generators of group of roots of unity of E and \mathcal{O} in polynomial time and e_i such that $\mathcal{O} = \prod_i e_i \mathcal{O}$ and $e_i \mathcal{O}$ has only trivial idempotents.

Proof. See [21, Theorem 1.2] for the first algorithm, which starts by expressing the product of number fields as $\mathbb{Q}[x]/(f)$ and then compute the factorization of f by LLL. The second part is quite involved and is proven in [22]. \square

The following method for recovering the root of unity is heavily inspired by [4].

Lemma 12. *Let $w \in \mu(\mathcal{O})$ where \mathcal{O} is a CM order. Then $w\bar{w} = 1$.*

Proof. With Lemma 1 and Theorem 2, E is a product of fields. Hence, $E \otimes \mathbb{C}$ is isomorphic to \mathbb{C}^n , and the conjugation can be projected to a conjugation over \mathbb{C} which has the same properties. Therefore, it is the standard conjugation over \mathbb{C} , and all roots of unity ζ over \mathbb{C} verifies $\zeta\bar{\zeta} = 1$. Since the projection of a root of unity $w \in \mathcal{O}$ is a root of unity, we have $w\bar{w} = 1$. \square

Lemma 13. *For any $a \in \mathcal{O}$ and \mathcal{O} a CM order, $\text{tr}(a\bar{a}) \geq r$ where r is the dimension of $a\bar{a}\mathcal{O}$.*

Proof. Without loss of generality, $a \neq 0$. Consider the application $x \mapsto xa\bar{a}$ over $a\bar{a}\mathcal{O}$. Its determinant is a non-zero integer since there is $a\bar{a}$ is not nilpotent, so that the inequality of arithmetic and geometric means over the eigenvalues gives the result. \square

Lemma 14. *If A and B are CM orders, then $A \otimes B$ is a CM order and if $\sum_i a_i \otimes b_i \in \mu(A \otimes B)$, we have $A = \prod_i a_i \bar{a}_i A \otimes b_i \bar{b}_i B$.*

Proof. The only difficult point in the first statement is to show that $x \mapsto \text{tr}(x\bar{x})$ is positive-definite. This comes from the fact that the corresponding Gram matrix is the Kronecker product of the two Gram matrices corresponding to A and B , which can be diagonalized thanks to the spectral theorem.

Then, if $\sum_i a_i \otimes b_i \in \mu(A \otimes B)$ where the sum is finite and $a_i, b_i \neq 0$, we have $(\sum_i a_i \otimes b_i)(\sum_i \bar{a}_i \otimes \bar{b}_i) = 1$ with Lemma 12. Therefore, $\sum_i a_i \bar{a}_i \otimes b_i \bar{b}_i = 1$. We deduce $\sum_i \text{tr}(a_i \bar{a}_i) \text{tr}(b_i \bar{b}_i) = \text{tr}(1)$ and using the previous lemma, $A \otimes B = \prod_i a_i \bar{a}_i A \otimes b_i \bar{b}_i B$ as product of suborders. \square

Lemma 15. *$B = \mathbb{Z}[X]/(X^n - 1)$ equipped with $X \mapsto X^{n-1}$ is a CM order and $\mu(B)$ is generated by X and -1 . Its idempotents are zero and one.*

Proof. Remark that for any $\omega \in \mu(B)$, we have with Lemma 12 $\omega\bar{\omega} = 1$. Then, with $\omega = \sum_{i=0}^{n-1} a_i X^i$, $\text{tr}(\omega\bar{\omega}) = n \sum_{i=0}^{n-1} a_i^2$. Therefore, $\omega = \pm X^i$ and the converse is clear.

If e is an idempotent, then $e\bar{e}$ is also an idempotent. But if $e \neq 0$, $\text{tr}(e\bar{e}) \geq n$ so that $e\bar{e} = 1$. Hence e is invertible, so that $e = 1$. \square

Theorem 11. *Given a CM order \mathcal{O} , an ideal $I = \omega\mathcal{O}$ with $\omega \in \mu(\mathbb{Q} \otimes \mathcal{O})$, we can find $\zeta \in \mu(\mathbb{Q} \otimes \mathcal{O})$ such that $\omega\zeta \in \mathcal{O}$ in polynomial time.*

Proof. Without loss of generality, we can assume $\omega \in \mu(\mathbb{Q} \otimes \mathcal{O})_p$ and we know some $e \leq 2n$ such that $\omega^e = 1$. We now compute all the primitive idempotents e_i of \mathcal{O} and by combining the results for all $I/e_i\mathcal{O}$ over $e_i\mathcal{O}$, we can assume \mathcal{O} has only trivial idempotents.

We then build the CM order $\mathcal{O} \otimes \mathbb{Z}[X]/(X^e - 1)$, by concatenating the basis of $I^i = (\omega^i)$. We now use [22, Theorem 1.2] to find the generators of the roots of unity of this order. Because of the previous lemmata, they are of the form $w \otimes X^i$ with $w \in \mu(\mathcal{O})$. By combining the generators, we can deduce a root of unity of the form $w \otimes X$. Hence, $w \in \omega\mu(\mathcal{O})$ so we can output $1/w$. \square

Theorem 12. *Given (I, r) , if E' is a product of cyclotomic field, or GRH is true, or we have access to randomness, we can find $v\omega \in \mu(E)\mathcal{O}$ such that $I = (v)$ or prove there is no such $v\omega$ with $N(v) = r$. We can also do this unconditionally in time $2^{n^{(1+o(1))/\log \log n}}$. Furthermore, if \mathcal{O} is a CM order, we can find $v \in \mathcal{O}$ such that $(I, r) = ((v), N(v))$.*

Proof. We first compute all \mathfrak{m}_i , apply Theorem 9 for each \mathfrak{m}_i and recover some x using the Chinese remainder theorem. Then, we compute $J = I/x = (\omega + a)$ where $a^n = 0$ and $\omega \in E$ is a root of unity. From the knowledge of the group of roots of unity of E , we can deduce in polynomial time e such that $\omega^e = 1$. We may then compute $J^{e2^k} = (1 + b)^{2^k}$ with $b = (\omega + a)^e - 1$, which is easily seen to be a nilpotent. Since $1 - x \mapsto \sum_{i=1}^n -x^i/i$ for x nilpotent is a morphism, whose inverse is $x \mapsto \sum_{i=0}^n x^k/k!$ (see [21, Proposition 8.1]), this takes time which is polynomial in $k \log(e)$.

We can compute \mathcal{O}' , the largest order which contains \mathcal{O} and all roots of unity of E in polynomial time, and we have $b \in \mathcal{O}'$. Therefore, we have $(1 + b)^2 = 1 + b^2$ in $\mathcal{O}'/(2)$ so that $(1 + b)^{2^{\lceil \log n \rceil}} \in 1 + 2\mathcal{O}'$. We deduce $(1 + b)^{2^k} = 1 + 2^{k - \lfloor \log n \rfloor} \mathcal{O}'$. We apply Theorem 6 to J^{e2^k} to produce an invertible y such that $y/(1 + b)^{2^k} \in \mathcal{O}'$, with $h(y/(1 + b)^{2^k}) = h(\mathcal{O})^{\mathcal{O}(1)}$. Hence, we choose k sufficiently large so that $y/(1 + b)^{2^k}$ is equal to the lift of y modulo $2^{k - \lfloor \log n \rfloor} \mathcal{O}'$. We can therefore compute $(1 + b)^{2^k}$ and using the two morphisms, deduce $z = x(1 + a/\omega)$.

We finally apply the previous theorem with $I/(z)$ to recover ω . □

3.6. Comments on the algorithm

One problem with the given algorithm is that it is *not explicit*. In particular, we need an upper-bound on the constants of the algorithms dealing with the norms (tr , N , multiplication and inversion in \mathcal{N}). However, this seems to be an unavoidable consequence of our black-box model, as slower algorithms mean a possibly larger set of solutions. Also, for any application exposed at the beginning, these constants are explicit. Furthermore, if we impose that there is a solution, one can simply increase the constant until we reach the solution.

Another difficulty is the sheer complexity of the algorithm, both in term of code length and running time. However, a large part of this complexity can be removed. Indeed, in practice, as soon as we combine information given by a couple of exponents (typically two, if we manage to find small primes having a prime ideal of inertia degree one above them), the greatest common divisor becomes tiny. It can be explained by a heuristic application of Chebotarev's theorem : if p^k divides the current greatest common divisor but does not divide the number of roots, the probability that k will not decrease is $O(1/p)$. Hence, we can simply ignore all primes p where $p\mathcal{O}$ is not invertible, or p divides the discriminant of the polynomial ; and beside the exponentiation, we only need to factor the polynomial defining the number field to produce the prime ideals. Also, applications can generally cope with finding the solution up to root of unity of E , since they usually work with an order in a number field with a known polynomial, which contains few roots of unity, and no nilpotents. Root extraction can be efficiently computed if we know an inert prime ; a Newton-Hensel iteration may also work. Ideal multiplication can be accelerated by compressing the lattice, see [23, Section 4]. While the exponent needed might seem to be huge, it is usually fairly small. For example, when $\mathcal{O} = \mathbb{Z}[\zeta_m]$, the precision needed is exactly the size of a typical LLL reduction of a lattice of determinant one, which is in practice 1.022^n in dimension n [24]. Finally, we explain in subsection 4.3 that under plausible heuristics, the exponent is bounded by $O(\log(h(\mathcal{O})) \log(\log(h(\mathcal{O}))))$, so that the resulting complexity is in general a couple of lattice reductions. We add that we can save an ideal powering using a Hensel iteration :

Lemma 16. *Let s be a formal product such that $[s] = v^e$ for some known e , and $v \in \mathcal{O}'$. Given a bound on $h(v)$ and a prime invertible ideal \mathfrak{p} of inertia degree f above the prime p ,*

with $d = \gcd(e, p^f - 1)$, $p \nmid e$ and $v \notin \mathfrak{p}$, we can compute v in time polynomial in the size of the input, p and d .

Proof. We select some k and compute $[s] \bmod \mathfrak{p}^k$ and c the inverse of e/d modulo $p^f - 1$. We can then factor $X^d - [s]$ in the finite field $\mathcal{O}'/\mathfrak{p}$ in time polynomial in p , d and the size of the input. For each root r , we have r^c a root of $X^e - [s] \bmod \mathfrak{p}$, which we can extend (since $p \nmid e$) using Hensel lifting to a root of $X^e - [s] \bmod \mathfrak{p}^k$. Using Lemma 5 and Lemma 6, for some polynomially large k , we can recover v from $v \bmod \mathfrak{p}^k$. \square

4. Applications

4.1. Dimension halving

In this subsection we define \mathcal{O} to be a CM order. The algorithm was first evoked in Gentry's dissertation [25, Section 6.2] before being developed in GGH [7, Section 8.8.1]. We correct here two benign mistakes in the algorithm. The first is that we should prove the existence of a short non-zero vector. The second is that the Gentry-Szydlo does not give a unique solution.

Lemma 17. *Given an integral invertible ideal $I = (v)$ of \mathcal{O} and some $x \in I - \{0\}$ which minimizes $\text{tr}(x\bar{x})$, then there exists $y \neq 0$ in $v\mathcal{O}^+$ or $v\mathcal{O}^-$ such that $0 < \text{tr}(y\bar{y}) \leq 2 \text{tr}(x\bar{x})$. Furthermore, these two lattices are included in I .*

Proof. Let $z = x/v$. Remark that $\text{tr}(v\bar{z}v\bar{z}) = \text{tr}(vz\bar{v}\bar{z}) = \text{tr}(x\bar{x})$. Since $x \mapsto \text{tr}(x\bar{x})$ is a quadratic form, there exists $s \in \{-1, 0, 1\}$ such that $0 < \text{tr}(v(z + s\bar{z})\overline{v(z + s\bar{z})}) \leq 2 \text{tr}(x\bar{x})$ and $s = 0$ only if $z = \bar{z}$. If $s = 0$, then $vz \in v\mathcal{O}^+$. If $s = 1$, then $v(z + \bar{z}) \in v\mathcal{O}^+$. Else $s = -1$ and $v(z - \bar{z}) \in v\mathcal{O}^-$. \square

Theorem 13. *Given a CM order \mathcal{O} included in a product of k fields a principal ideal I , using one call to Theorem 4, having access to an oracle finding a non-zero vector in a lattice at most γ times larger than the shortest non-zero vector, and time polynomial in the size of the input and 2^k , we can find a non-zero vector at most $\gamma\sqrt{2}$ larger than the shortest non-zero vector of the ideal. Furthermore, all calls to the oracle are of dimension at most $\max(\dim \mathcal{O}^+, \dim \mathcal{O}^-)$.*

Proof. Without loss of generality, I is invertible and integral. Let $v \in \mathcal{O}$ such that $I = (v)$. We first compute \bar{I}/I and run Theorem 4, which returns some $\omega\bar{v}/v$ and $\omega \in \mu(\mathcal{O})$. Then, for all $\zeta \in \mu(\mathcal{O})/\mu(\mathcal{O})^2$, we deduce $J = I(1 + \zeta\omega\bar{v}/v) = (v + \zeta\omega\bar{v})$. For some ζ , we will have $\zeta\omega = \bar{w}^2$ and $w \in \mu(\mathcal{O})$. Thus with Lemma 12, $J = (vw + \bar{w}\bar{w})$ and $I = (vw)$. We can then compute $J \cap \mathcal{O}^+ = (vw + \bar{w}\bar{w})\mathcal{O}^+$ since $vw + \bar{w}\bar{w} \in \mathcal{O}^+$. Dividing by $1 + \bar{w}\bar{w}^2/v$, we get a basis of $v\mathcal{O}^+$. Now, the direct sum of $v\mathcal{O}^+$ and $v\mathcal{O}^-$ is $2I$ so we can compute a basis of $v\mathcal{O}^-$. Using Lemma 17, we just need to call the oracle on these two lattices.

The complexity is given by the fact that there are at most 2^k different ζ . \square

Usually $k = 1$ so that the algorithm is efficient. Then, either $\mathcal{O}^+ = \mathcal{O}$ and nothing happens or $\dim \mathcal{O}^+ = \frac{1}{2} \dim \mathcal{O}$ and the algorithm halve the dimension for a moderate cost.

It is easy to show that considering only $\zeta = 1$ does not work. For example, with $I = (1 + i)\mathbb{Z}[i]$, we have $\bar{I}/I = -i\mathbb{Z}[i] = \mathbb{Z}[i]$ so that we may recover 1 with Gentry-Szydlo's algorithm. Then $I(1 + 1) = (2 + 2i)\mathbb{Z}[i]$ is not generated by an element of $\mathbb{Z}[i]^+ = \mathbb{Z}$.

4.2. Solving the norm equation

PROBLEM 2. We are given a CM order \mathcal{O} and $r \in \mathcal{O}^+$. We want to know all $x \in \mathcal{O}$ such that $x\bar{x} = r$.

This is the norm equation problem, in the case of a CM order. The following algorithm was introduced by Howgrave-Graham and Szydło [26]. See [27] for a more general technique. Remark that there may be many solutions since $(x\bar{y})(\bar{x}y) = (xy)(\overline{xy})$, and possibly more than a polynomial. Also, this case seems to show that in a way, we are factoring a number, and hence, discovering factors of its algebraic norm. Hence, it is plausible that the following algorithm is close to optimal.

Theorem 14. *Let \mathcal{O} be a CM order over a number field. Given the factorisation of the algebraic norm over \mathbb{Q} of $r \in \mathcal{O}^+$, with d the number of divisors and $r \in \mathcal{O}^+$, we can compute all $x \in \mathcal{O}$ such that $x\bar{x} = r$, in time polynomial in the size of the input, d and calls to Theorem 4.*

Proof. Without loss of generality, using [15, Section 6.1] we can assume that all primes involved are invertible in \mathcal{O} and \mathcal{O}^+ . We may then find the factorisation of $r\mathcal{O}^+$ using [15, Sections 4.8.3 and 6.2]. Then, $\mathbb{Q} \otimes \mathcal{O}$ is a Galois extension of $\mathbb{Q} \otimes \mathcal{O}^+$ so that a prime ideal \mathfrak{p} of \mathcal{O}^+ is inert or factored into $\mathfrak{p}\bar{\mathfrak{p}}$. If it is inert, then the \mathfrak{p} valuation of $x\mathcal{O}$ must be half the \mathfrak{p} valuation of $r\mathcal{O}^+$. Else, the \mathfrak{p} valuation of $x\mathcal{O}$ must be less than the \mathfrak{p} valuation of $r\mathcal{O}^+$, and the $\bar{\mathfrak{p}}$ valuation is uniquely determined by it. Therefore, there are at most d $x\mathcal{O}$ distinct, and we can find all of them. Using Theorem 4, we may then obtain x . \square

4.3. Lowering the exponent and applications

In many cases, the norm, traces and exponentiation are in fact smooth functions. We can leverage this property by trying to run Gentry-Szydło's algorithm with an approximate norm. Indeed, what we need is that the last reduction in our powering algorithm (Theorem 8) gives a meaningful result. Of course, the quality of the approximation depends on the exponent used. We show here that heuristically, we can use tiny exponents. The idea comes from GH [7, Section 8.6]. Since all algorithms of this subsection use the following strong heuristic, we will also allow them to use randomness and GRH.

HEURISTIC 1. Let F be a number field of \mathbb{Q} -dimension n with exactly m roots of unity. Then, the expected value of the number of prime ideals above $p = am + 1$ of inertia degree one is $\Omega\left(\frac{\phi(m)}{\log(p)}\right)$ for a random a .

Proof. Note that a prime p having a prime ideal above it of inertia degree one must be of the form $am + 1$. The density of prime numbers among integers of this form is $\frac{m}{\phi(m)\log(p)}$. p always factors over $\mathbb{Q}[\zeta_m]$ into $\phi(m)$ ideals of inertia degree one. The sum of k times the density of prime ideals of $\mathbb{Q}[\zeta_m]$ having k prime ideals above it of inertia degree one is one, by Chebotarev theorem and Burnside's lemma. Assuming the two results occur somewhat randomly, and independently implies the heuristic. \square

Theorem 15. *Let \mathcal{O} be a split CM order with no nilpotents beside zero. If the heuristic assumption is true for each number fields, we can in polynomial time, given $(I, r) = (v, N(v))$, find some solution $v \in \mathcal{O}$ if it exists. Furthermore, the exponent used in calls to Theorem 8 is $e = h(\mathcal{O})^{O(\log(h(\mathcal{O})))}$. Therefore, if $\text{tr}(N(x)(\tilde{r}/r)^e) / \text{tr}(N(x)) \in [1/2; 2]$ for all $x \in \mathcal{O} - \{0\}$, we can find some solution $v \in \mathcal{O}$ to $(I, r) = (v, N(v))$ given I and \tilde{r} .*

Proof. Using Theorem 6, we can assume that $h(v) = h(\mathcal{O})^{O(1)}$. Let p_i be the sequence of prime numbers. For some k , we let $e = m \prod_{i=0}^{k-1} p_i$. We then proceed just like in Theorem 8 with $(I, r, 1)^e$. Remark that there are at least 2^k divisors of e of the form km , and $\log(e) = O(k \log k)$. We therefore expect $\Omega\left(\frac{2^k \phi(m)}{k \log k}\right)$ of the $1 + d$ with d divisor of e to be prime with a prime ideal $\mathfrak{p}_{a,j}$ of inertia degree one above it. Only $O(\log(h(\mathcal{O})))$ of these can divide v . The inverse of the returned ideal is generated by a small integer in $1 + \prod_j \mathfrak{p}_j$. Hence, if the determinant of $\prod_j \mathfrak{p}_{a,j}$

is exponential in $h(\mathcal{O})^{O(1)}$, we can proceed. Then, using another prime of inertia degree one, we can finish in polynomial time with Lemma 16. Our condition is then $\Omega(2^k \phi(m)) = h(\mathcal{O})^{O(1)}$, so that some $k = O(\log(h(\mathcal{O})))$ works. \square

In case our number field is $\mathbb{Z}[\zeta_n]$, we need only the product of the primes to be above the LLL approximation factor, which in practice is $\approx 1.022^{\phi(n)}$ [24]. For $n = 2^{16}$, we can use $k = 8$ so that $e = 635678883840$ and the product is around 2^{1048} which is larger than the required 2^{1029} . This implies that only ≈ 50 lattice reductions are needed. The sum of $2/\ln(d)$ for all $n \mid d \mid e$ is ≈ 27 , and there are 38 $d + 1$ which are primes. Taking $k = 21$ leads to 168076 primes instead of the predicted 98361, and the product has more than 10 million bits while $e < 2^{112}$.

This theorem can be used to recover a unit u from its approximation \tilde{u} by calling it with $(\mathcal{O}, \tilde{u}\bar{\tilde{u}})$ within the corresponding split CM order. If \mathcal{O} is in a number field, $\mathbb{R} \otimes \mathcal{O} \simeq \mathbb{R}^r \mathbb{C}^s$ and by applying some complex logarithm, the image of \mathcal{O}^* is a lattice of dimension $r + s - 1$. Now, the precision required in this basis is simpler to express : the error should be at most $n^{-O(\log(\log n))}$ on each coordinate of the image of $\tilde{u}\bar{\tilde{u}}$.

The following theorem can be seen as a way to compute a greatest common divisor.

Theorem 16. *Let \mathcal{D} be a samplable distribution over the number field E such that for all embedding $\psi : \mathcal{O} \rightarrow \mathbb{C}$, $\log(|\psi(a)|)$ has standard deviation at most σ . Given (v) and k samples $s_i = va_i$ where the a_i are independent and sampled from \mathcal{D} , if the heuristic holds, we can recover $v\mu(\mathcal{O})$ in polynomial time if $k \geq \sigma^2 n^{O(\log \log n)}$ except with negligible probability.*

Proof. We fix an embedding of E into \mathbb{C} and then we can define the split CM order using the norm $x \mapsto x\bar{x}$. We compute \tilde{r} the average of $s_i \bar{s}_i$ divided by the average of $a_i \bar{a}_i$ computed by sampling from \mathcal{D} . Using Chebyshev inequality, we can prove that if we use $\sigma^2 n^{O(\log \log n)}$ samples, then for all embedding $\psi : E \mapsto \mathbb{C}$,

$$\psi(\tilde{r})/|\psi(v)| \in [1 - n^{-O(\log \log n)}; 1 + n^{-O(\log \log n)}]$$

with probability at least $1/2$. Since $\text{tr}(N(x)) = \sum_{\psi} |\psi(x)|^2$, we can use the previous theorem. \square

The original Gentry-Szydlo attack on NTRU signatures [3] is essentially an application of this theorem. It improves on it by remarking that if \mathcal{O} is a CM order we can compute $(v\bar{v})$, reduce this basis, and use it to decode \tilde{r} and recover $v\bar{v}$. Another possibility which works for any order is to decode \tilde{r} over a basis of \mathcal{N} by truncating the coefficients, which has the advantage of being polynomial time. It gives a proven algorithm which is polynomial, and needs a number of samples which is about the maximum coefficient of $v\bar{v}$.

Note that taking the ideal generated by all s_i should get (v) for most applications so that the hardest condition to achieve is the possibility of sampling from \mathcal{D} .

Theorem 17. *Let \mathcal{O} be in a number field, and for some embedding in \mathbb{C} , we define $N(x) = x\bar{x}$.*

Given (I, r) a polarized ideal of determinant Δ , if the heuristic holds, we can determine if there exists a $v \in I$ such that for all embedding ψ into \mathbb{C} , we have $|\psi(v)|^2/\psi(r) \leq 1 + 1/e$ and find it, for some $e = \log(h(\mathcal{O}))^{O(\log(h(\mathcal{O})))}$.

Proof. We use $e = m \prod_{i=0}^{k-1} p_i$ and compute $(I, r, 1)^e = (J, p, s)$. Now, $\det(I^e) = \det(I)^e$ using [15, Proposition 4.6.8], so that the determinant of $(I, r)^e$ is also Δ . Without loss of generality, we can assume $h(I) = h(\mathcal{O})^{O(1)}$. We select a random subset of half the prime ideals of inertia degree one above p with $p - 1 \mid e$, so that with high probability $v^e = 1 \pmod{K}$ where $K \cap I$ has no non-zero vector shorter than $2^n 3n$. We deduce that $v^e/[s] \in J \cap (1/[s] + K)$ and $\text{tr}(N(v^e/[s])/p) = \text{tr}(N(v^e)/r^e) \leq 3n$. Therefore, we can apply Babai's algorithm Lemma 6 on $J \cap (1/[s] + K)$ equipped with the norm $x \mapsto \text{tr}(N(x)/p)$ and recover v^e as a formal product. We now find another small prime ideal and using Lemma 16, we recover v . \square

Note that this implies that finding the shortest vector (for some norm) of invertible ideals is easy if it is almost as small as it can be (1). If $r = N(v)$, then $\det((v))/\det(I) < 2$ so that $I = (v)$ which is the standard case.

Acknowledgement

We thank Pierre-Alain Fouque for his comments allowing to improve a draft of this paper.

References

1. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. pages 84–93, 2005.
2. Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. pages 1–23, 2010.
3. Craig Gentry and Michael Szydlo. Cryptanalysis of the revised NTRU signature scheme. pages 299–320, 2002.
4. H. W. Lenstra and A. Silverberg. Revisiting the Gentry-Szydlo algorithm. pages 280–296, 2014.
5. Nigel P. Smart and Frederik Vercauteren. Fully homomorphic encryption with relatively small key and ciphertext sizes. pages 420–443, 2010.
6. Craig Gentry and Shai Halevi. Implementing Gentry’s fully-homomorphic encryption scheme. pages 129–148, 2011.
7. Sanjam Garg, Craig Gentry, and Shai Halevi. Candidate multilinear maps from ideal lattices. pages 1–17, 2013.
8. Peter Campbell, Michael Groves, and Dan Shepherd. Soliloquy: A cautionary tale. In *ETSI 2nd Quantum-Safe Crypto Workshop*, 2014.
9. Dan Boneh and David Mandell Freeman. Homomorphic signatures for polynomial functions. pages 149–168, 2011.
10. Jean-François Biasse and Fang Song. Efficient quantum algorithms for computing class groups and solving the principal ideal problem in arbitrary degree number fields. pages 893–902, 2016.
11. Ronald Cramer, Léo Ducas, Chris Peikert, and Oded Regev. Recovering short generators of principal ideals in cyclotomic rings. Cryptology ePrint Archive, Report 2015/313, 2015. <http://eprint.iacr.org/2015/313>.
12. Martin Albrecht, Shi Bai, and Léo Ducas. A subfield lattice attack on overstretched ntru assumptions: Cryptanalysis of some fhe and graded encoding schemes. 2016.
13. Dwaine E. Clarke, Srinivas Devadas, Marten van Dijk, Blaise Gassend, and G. Edward Suh. Incremental multiset hash functions and their application to memory integrity checking. pages 188–207, 2003.
14. Eva Bayer-Fluckiger. Ideal lattices. *A panorama of number theory or the view from Baker’s garden (Zurich, 1999)*, pages 168–184, 2002.
15. Henri Cohen. *A course in computational algebraic number theory*, volume 138. Springer Science & Business Media, 2013.
16. René Schoof. Computing arakelov class groups. *arXiv preprint arXiv:0801.3835*, 2008.
17. Arjen Klaas Lenstra, Hendrik Willem Lenstra, and László Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261(4):515–534, 1982.
18. László Babai. On lovász’lattice reduction and the nearest lattice point problem. *Combinatorica*, 6(1):1–13, 1986.
19. Triantafyllos Xylouris. *Über die Nullstellen der Dirichletschen L-Funktionen und die kleinste Primzahl in einer arithmetischen Progression*. PhD thesis, Universitäts- und Landesbibliothek Bonn, 2011.
20. Jean-Pierre Serre. Quelques applications du théorème de densité de chebotarev. *Publications Mathématiques de l’IHES*, 54:123–201, 1981.
21. HW Lenstra Jr and A Silverberg. Algorithms for commutative algebras over the rational numbers. *arXiv preprint arXiv:1509.08843*, 2015.
22. HW Lenstra Jr and A Silverberg. Roots of unity in orders. *arXiv preprint arXiv:1509.02612*, 2015.
23. Zhuliang Chen and Arne Storjohann. A blas based c library for exact linear algebra on integer matrices. In *Proceedings of the 2005 international symposium on Symbolic and algebraic computation*, pages 92–99. ACM, 2005.
24. Nicolas Gama and Phong Q. Nguyen. Predicting lattice reduction. pages 31–51, 2008.
25. Craig Gentry. *A fully homomorphic encryption scheme*. PhD thesis, Stanford University, 2009.
26. Nick Howgrave-Graham and Mike Szydlo. A method to solve cyclotomic norm equations $f^* \setminus \bar{\{f\}}$. In *Algorithmic Number Theory*, pages 272–279. Springer, 2004.
27. Denis Simon. Solving norm equations in relative number fields using s-units. *Mathematics of computation*, 71(239):1287–1305, 2002.
28. Nicolas Gama, Phong Q. Nguyen, and Oded Regev. Lattice enumeration using extreme pruning. pages 257–278, 2010.

29. Guillaume Hanrot, Xavier Pujol, and Damien Stehlé. Algorithms for the shortest and closest lattice vector problems. In *Coding and Cryptology*, pages 159–190. Springer, 2011.
30. Thijs Laarhoven. *Search problems in cryptography*. PhD thesis, Eindhoven University of Technology, 2015.
31. Anja Becker and Thijs Laarhoven. Efficient (ideal) lattice sieving using cross-polytope LSH. Cryptology ePrint Archive, Report 2015/823, 2015. <http://eprint.iacr.org/2015/823>.
32. Joppe W. Bos, Michael Naehrig, and Joop van de Pol. Sieving for shortest vectors in ideal lattices: a practical perspective. Cryptology ePrint Archive, Report 2014/880, 2014. <http://eprint.iacr.org/2014/880>.

Appendix A. Exploiting roots of unity

We assume here that E is a number field with m roots of unity in \mathcal{O} . We show how their presence allows to accelerate standard lattice algorithms when the geometric norm is $\|x\|^2 = \text{tr}(N(x)) = \text{tr}(x\bar{x})$. We define P as $X^{m/2} + 1$ if $4 \mid m$, $X^m - 1$ else; and let m' be the degree of P . Remark that there is a natural bijection between the roots of unity of \mathcal{O} and $\pm X^i$ modulo P .

Theorem 18. *For all $x \in I$, with I an ideal of \mathcal{O} , we have $x\omega \in I$ and $\|x\| = \|x\omega\|$ for any $\omega \in \mu(\mathcal{O})$.*

Proof. The first property stems from I being an ideal, the second from $N(\omega) = \omega\bar{\omega} = 1$. \square

This implies that I has at least $|\mu(\mathcal{O})|$ non-zero shortest vectors, making the extreme pruning algorithm [28] about $|\mu(\mathcal{O})|/2$ times faster than on a "random" lattice, since $x\mu(\mathcal{O})$ is somewhat uniform over the sphere.

Also, sieving algorithms (see [29, 30][†] for surveys) can take advantage of this by reducing the size of the list of vectors by a factor of $|\mu(\mathcal{O})|/2$ for the same reason. A recent algorithm [31] works by introducing a hash function h which for a vector returns the index of the largest coordinate, as well as its sign. It is then randomized to $h_a(x) = h(ax)$ for a Gaussian a to produce a locality-sensitive hash function H by concatenating outputs of several h_a .

We can improve on this by embedding I and E in $\mathbb{Q}[X]/(P(X))[Y]/(Q(Y))$ for some irreducible polynomial $Q \in \mathbb{Q}[\zeta_m][Y]$ of degree $n/\phi(m)$, so that they have the same geometry. We can now choose $h_a(x) = h(ax)$ and observe that $h_a(x\omega)$ for ω a root of unity is simply a rotation of h_a . Hence, we can build H as the concatenation of $h_{a_0}, h_{a_1}, \dots, h_{a_k}$ where the output of h_{a_0} is forced to be on a positive monomial of the form Y^i by considering the unique root of unity which allows this. The algorithm then has to compute the shortest element among $x + \omega y$ for all $\omega \in \mu(E)$. We now show that this can be computed efficiently.

Theorem 19. *Given $x, y \in \mathbb{Q}[X]/(P(X))[Y]/(Q(Y))$ with P and Q defined as above, we can compute $\arg \min_{\omega \in \mu(\mathcal{O})} \|x + \omega y\|$ in $O((n/\phi(m))^2 m \log m)$ arithmetic operations.*

Proof. We denote $x = \sum_{i=0}^{n/\phi(m)-1} x_i Y^i$ for any x . Now $\langle x, y \rangle = \sum_{i,j} G_{i,j} \langle x_i, y_j \rangle$ for some Gram matrix G with the scalar product $\langle x_i, y_j \rangle$ corresponding to the norm over the CM order. Hence, we only need to show how to compute $\langle a, bX^i \rangle$ for $a, b \in \mathbb{Q}[X]/(P(X))$ and all i in $O(m \log m)$ operations.

Since the norm over $\mathbb{Q}[X]/(P(X))$ is $x \mapsto \text{tr}(x\bar{x})$, we have $\langle a, b\omega \rangle = \text{tr}(a\bar{b}\bar{\omega})$. We can therefore compute $a\bar{b}$ with a Fourier transform in time $O(m \log m)$. Finally, $\text{tr}(aX^{-i})$ is exactly the i -th coefficient of a . \square

This implies an overall speed-up of $m^{1.43+o(1)}$, while [31] gives a speed-up of only $O(m)$, and the ideals were required to be over a ring of the form $X^m \pm 1$. The use of Fourier transform for accelerating geometric computations was first introduced by [32].

[†]Beware that the literature often uses different way for expressing multiplication, multiplication by a root of unity or conjugation, such as (nega)cyclic matrices, rotation and reflex polynomial.