ROOTS OF UNITY IN ORDERS

H. W. LENSTRA, JR. AND A. SILVERBERG

Communicated by John Cremona

ABSTRACT. We give deterministic polynomial-time algorithms that, given an order, compute the primitive idempotents and determine a set of generators for the group of roots of unity in the order. Also, we show that the discrete logarithm problem in the group of roots of unity can be solved in polynomial time. As an auxiliary result, we solve the discrete logarithm problem for certain unit groups in finite rings. Our techniques, which are taken from commutative algebra, may have further potential in the context of cryptology and computer algebra.

1. INTRODUCTION

An *order* is a commutative ring whose additive group is isomorphic to $\mathbb{Z}^n$ for some non-negative integer n. The present paper contains algorithms for computing the *idempotents* and the *roots of unity* of a given order.

In algorithms, we specify an order A by listing a system of "structure constants" $a_{ijk} \in \mathbb{Z}$ with $i, j, k \in \{1, 2, \ldots, n\}$; these determine the multiplication in A in the sense that for some $\mathbb{Z}$-basis $e_1, e_2, \ldots, e_n$ of the additive group of A, one has $e_i e_j = \sum_{k=1}^{n} a_{ijk} e_k$ for all i, j. The elements of A are then represented by their coordinates with respect to that basis.

An idempotent of a commutative ring R is an element $e \in R$ with $e^2 = e$, and we denote by $\mathrm{id}(R)$ the set of idempotents. An idempotent $e \in \mathrm{id}(R)$ is called *primitive* if $e \neq 0$ and for all $e' \in \mathrm{id}(R)$ one has $ee' \in \{0, e\}$; let $\mathrm{prid}(R)$ denote the set of primitive idempotents of R.

Orders A have only finitely many idempotents, but they may have more than can be listed by a polynomial-time algorithm; however, if one knows $\mathrm{prid}(A)$, then one implicitly knows $\mathrm{id}(A)$, since there is a bijection from the set of subsets of $\mathrm{prid}(A)$ to $\mathrm{id}(A)$ that sends $W \subset \mathrm{prid}(A)$ to $e_W = \sum_{e \in W} e \in \mathrm{id}(A)$. For $\mathrm{prid}(A)$ we have the following result.

Theorem 1.1. *There is a deterministic polynomial-time algorithm (Algorithm 6.1) that, given an order A, lists all primitive idempotents of A.*

A *root of unity* in a commutative ring R is an element of finite order of the group R^* of invertible elements of R; we write $\mu(R)$ for the set of roots of unity in R, which is a subgroup of R^*.

As with idempotents, orders A have only finitely many roots of unity, but possibly more than can be listed by a polynomial-time algorithm, and to control $\mu(A)$ we shall use generators and relations. If S is a finite system of generators for an abelian group G, then by a *set of defining relations* for S we mean a system of generators for the kernel of the surjective group homomorphism $\mathbb{Z}^S \to G$, $(m_s)_{s \in S} \mapsto \prod_{s \in S} s^{m_s}$.

Theorem 1.2. *There is a deterministic polynomial-time algorithm (Algorithm 13.2) that, given an order A, produces a set S of generators of $\mu(A)$, as well as a set of defining relations for S.*

Theorem 1.2, which provides a key ingredient in an algorithm for lattices with symmetry that was recently developed by the authors [6, 7], is our main result, and its proof occupies most of the paper. It makes use of several techniques from commutative algebra that so far have found little employment in an algorithmic context. A sketch appeared in Proposition 4.7 of [6].

We shall also obtain a solution to the discrete logarithm problem in $\mu(A)$ and all its subgroups, and more generally in all subgroups of the group $\mu(A \otimes_{\mathbb{Z}} \mathbb{Q})$, which is still finite. Note that $A \otimes_{\mathbb{Z}} \mathbb{Q}$ is a ring containing A as a subring, and that a $\mathbb{Z}$-basis for A is a $\mathbb{Q}$-basis for the additive group of $A \otimes_{\mathbb{Z}} \mathbb{Q}$. If one replaces $\mu(A)$ by $\mu(A \otimes_{\mathbb{Z}} \mathbb{Q})$ in Theorem 1.2, then it remains true, and in fact it becomes much easier to prove (Proposition 3.5). Our solution to the discrete logarithm problem in $\mu(A \otimes_{\mathbb{Z}} \mathbb{Q})$ and all of its subgroups, in particular in $\mu(A)$, reads as follows.

Theorem 1.3. *There is a deterministic polynomial-time algorithm that, given an order A, a finite system T of elements of $\mu(A \otimes_{\mathbb{Z}} \mathbb{Q})$, and an element $\zeta \in A \otimes_{\mathbb{Z}} \mathbb{Q}$, decides whether ζ belongs to the subgroup $\langle T \rangle \subset \mu(A \otimes_{\mathbb{Z}} \mathbb{Q})$ generated by T, and if so finds $(m_t)_{t \in T} \in \mathbb{Z}^T$ with $\zeta = \prod_{t \in T} t^{m_t}$.*

We shall prove Theorem 1.3 in section 7, as a consequence of the results on $\mu(A \otimes_{\mathbb{Z}} \mathbb{Q})$ in section 3 and a number of formal properties of "efficient presentations" of abelian groups that are developed in section 7.

A far-reaching generalization of Theorem 1.3, in which $\mu(A \otimes_{\mathbb{Z}} \mathbb{Q})$ is replaced by the full unit group $(A \otimes_{\mathbb{Z}} \mathbb{Q})^*$, is proven in [8].

Of the many auxiliary results that we shall use, there are two that have independent interest. The first concerns the discrete logarithm problem in certain unit groups of finite rings, and it reads as follows.

Theorem 1.4. *There is a deterministic polynomial-time algorithm that, given a finite commutative ring R and a nilpotent ideal $I \subset R$, produces a set S of generators of the subgroup $1 + I \subset R^*$, as well as a set of defining relations for S. Also, there is a deterministic polynomial-time algorithm that, given R and I as before, as well as a finite system T of elements of $1 + I$ and an element $\zeta \in R$, decides whether ζ belongs to the subgroup $\langle T \rangle \subset 1 + I$, and if so finds $(m_t)_{t \in T} \in \mathbb{Z}^T$ with $\zeta = \prod_{t \in T} t^{m_t}$.*

The proof of this theorem is given in section 11. It depends on the resemblance of $1 + I$ to the additive group I, in which the discrete logarithm problem is easy.

The second result that we single out for special mention is of a purely theoretical nature. Let R be a commutative ring. For the purposes of this paper, commutative rings have an identity element 1 (which is 0 if and only if the ring is the 0 ring). We call R *connected* if $\#\mathrm{id}(R) = 2$ or, equivalently, if $\mathrm{id}(R) = \{0, 1\}$ and $R \neq \{0\}$. A polynomial $f \in R[X]$ is called *separable* (over R) if f and its formal derivative f' generate the unit ideal in $R[X]$. For example, $f = X^2 - X$ is separable because $(f')^2 - 4f = 1$.

Theorem 1.5. *Let R be a connected commutative ring, and let $f \in R[X]$ be separable. Then $f \neq 0$ and $\#\{r \in R : f(r) = 0\} \leq \deg(f)$.*

For the elementary proof, see section 8.

While, technically, one must admit that Theorem 1.5 plays only a modest role in the paper, it does convey an important message, namely that zeroes of polynomials that are separable are easier to control than zeroes of other polynomials. Thus, $X^2 - X$ is separable over any R, while $X^m - 1$ (for $m \in \mathbb{Z}_{>0}$) is separable if and only if $m \cdot 1 \in R^*$, a condition that for a non-zero order and $m > 1$ is never satisfied; accordingly, Theorem 1.1 is much easier to prove than Theorem 1.2.

We next provide an overview of the algorithms that underlie Theorems 1.1 and 1.2. In both cases, one starts by reducing the problem, in a fairly routine manner, to the special case in which each

element of A is a zero of some separable polynomial in $\mathbb{Q}[X]$; for the rest of the introduction we assume that the latter condition is satisfied. Then the $\mathbb{Q}$-algebra $E = A \otimes_{\mathbb{Z}} \mathbb{Q}$ can be written as the product of finitely many algebraic number fields $E/\mathfrak{m}$, with $\mathfrak{m}$ ranging over the finite set $\mathrm{Spec}(E)$ of prime ideals of E; hence $\mathrm{prid}(E)$ is in bijection with $\mathrm{Spec}(E)$. The image of $A \subset E$ under the map $E \to E/\mathfrak{m}$ may be identified with the ring $A/(\mathfrak{m} \cap A)$, so that A becomes a subring of the product ring $B = \prod_{\mathfrak{m} \in \mathrm{Spec}(E)} A/(\mathfrak{m} \cap A)$; this is also an order, and it is "close" to A in the sense that the abelian group B/A is finite. The ring B has many idempotents, in the sense that $\mathrm{id}(B)$ equals all of $\mathrm{id}(E)$, and $\#\mathrm{prid}(B) = \#\mathrm{Spec}(E)$. To determine which subsets $W \subset \mathrm{prid}(B)$ give rise to idempotents that lie in A, we define a certain graph $\Gamma(A)$ with vertex set $\mathrm{Spec}(E)$ such that the connected components of $\Gamma(A)$ correspond exactly to the primitive idempotents of A. This leads to Theorem 1.1.

To prove Theorem 1.2, one likewise starts from B, generators for $\mu(B)$ being easily found by standard algorithms from algebraic number theory. However, there is no standard way of computing $\mu(A) = \mu(B) \cap A$, which is the intersection of a multiplicative group and an additive group, and we must proceed in an indirect way. For a prime number p, denote by $\mu(A)_p$ the group of roots of unity in A that are of p-power order, and likewise $\mu(B)_p$. Then $\mu(A)$ is generated by its subgroups $\mu(A)_p = \mu(B)_p \cap A$, with p ranging over the set of primes dividing $\#\mu(B)$; all these p are "small". It will now suffice to fix p and determine generators for $\mu(A)_p$. To this end, we introduce the intermediate order $A \subset C \subset B$ defined by $C = A[1/p] \cap B$. The finite abelian group B/C is of order coprime to p, and it turns out that this makes it relatively easy to determine $\mu(C)_p = \mu(B)_p \cap C$; in fact, one of the results (Proposition 8.1(b)) leading up to Theorem 1.5 stated above shows that this can be done by exploiting the graph $\Gamma(C)$ that we encountered in the context of idempotents. The passage to $\mu(A)_p = \mu(C)_p \cap A$ is of an entirely different nature, as C/A is of order a power of p. It is here that we have to invoke Theorem 1.4 for certain finite rings R that are of p-power order.

It is important to realize that the only reason that an intersection such as $\mu(A) = \mu(B) \cap A$ is hard to compute is that $\mu(B)$, though finite, may be *large*—testing each element of $\mu(B)$ for membership in A will not lead to a polynomial-time algorithm. By contrast, the *exponent* of each group $\mu(B)_p$ is *small* (Lemma 3.3(iv)), so results stating that certain subgroups of $\mu(B)_p$ are cyclic—of which there are several in the paper—are valuable in obtaining a polynomial bound for the runtime of our algorithm.

2. Definitions and examples

From now on, when we say commutative $\mathbb{Q}$-algebra we will mean a commutative $\mathbb{Q}$-algebra that is finite-dimensional as a $\mathbb{Q}$-vector space. See [1, 3] for background on commutative rings and linear algebra.

Definition 2.1. If A is an order whose additive group is isomorphic to $\mathbb{Z}^n$, we call n the **rank** of A.

If the number of idempotents in R is finite, then each idempotent is the sum of a unique subset of $\mathrm{prid}(R)$, and one has $\#\mathrm{id}(R) = 2^{\#\mathrm{prid}(R)}$.

Definition 2.2. A commutative ring R is called **connected** if $\#\{x \in R : x^2 = x\} = 2$.

Definition 2.3. If R is a commutative ring, let $\mathrm{Spec}(R)$ denote the set of prime ideals of R.

Although we do not use it, we point out that a commutative ring R is connected if and only if $R \neq 0$ and R cannot be written as a product of 2 non-zero rings. The definition is motivated by the fact that a commutative ring R is connected if and only if $\mathrm{Spec}(R)$ is connected. (A topological space is connected if and only if it has exactly 2 open and closed subsets.)

Notation 2.4. If G is a group and p is a prime number, define
$$G_p = \{g \in G : g^{p^r} = 1 \text{ for some } r \in \mathbb{Z}_{\geq 0}\}.$$

Definition 2.5. Suppose R is a commutative ring. A polynomial $f \in R[X]$ is **separable** over R if

$$R[X]f + R[X]f' = R[X],$$

where if $f = \sum_{i=0}^{t} a_i X^i$ then $f' = \sum_{i=1}^{t} i a_i X^{i-1}$.

One can show that if f is a monic polynomial over a commutative ring R, then f is separable over R if and only if its discriminant is a unit in R.

Definition 2.6. Suppose E is a commutative $\mathbb{Q}$-algebra. If $\alpha \in E$, then α is **separable** over $\mathbb{Q}$ if there exists a separable polynomial $f \in \mathbb{Q}[X]$ such that $f(\alpha) = 0$. Let E_{sep} denote the set of $y \in E$ that are separable over $\mathbb{Q}$. We say E is separable over $\mathbb{Q}$ if $E_{\mathrm{sep}} = E$.

We note that E_{sep} is a commutative $\mathbb{Q}$-algebra (see for example Theorem 1.1 of [8]).

Definition 2.7. Suppose R is a commutative ring. An element $x \in R$ is called **nilpotent** if there exists $n \in \mathbb{Z}_{>0}$ such that $x^n = 0$. An ideal I of R is called nilpotent if there exists $n \in \mathbb{Z}_{>0}$ such that $I^n = 0$, where I^n is the product of I with itself n times. The set of nilpotent elements of R is an ideal, called the **nilradical** and denoted $\sqrt{0}$ or $\sqrt{0_R}$.

Examples 2.8. The polynomial $X^2 - X$ is separable over every ring. A linear polynomial $aX + b$ is separable over R if and only if the R-ideal generated by a and b is R. If $m \in \mathbb{Z}_{\geq 0}$, then the polynomial $X^m - 1$ is separable over R if and only if $m \cdot 1$ is a unit in R.

Example 2.9. Suppose $f(X) \in \mathbb{Z}[X]$ is a monic polynomial of degree n. Then the ring $\mathbb{Z}[X]/(f)$ is an order of rank n. We remark that the map $e \mapsto \gcd(e, f)$ is a bijection from the set of idempotents of $\mathbb{Z}[X]/(f)$ to $\{g \in \mathbb{Z}[X] : g \text{ is monic}, g|f, \text{ and } R(g, f/g) = \pm 1\}$, where $R(g, f/g)$ is the resultant of g and f/g.

Example 2.10. If G is a finite group of order $2n$ with a fixed element u of order 2, then $\mathbb{Z}\langle G \rangle = \mathbb{Z}[G]/(u+1)$ is a connected order of rank n, and $\mu(\mathbb{Z}\langle G \rangle) = G$ (see Remark 16.3 of [7]).

Example 2.11. If $n \in \mathbb{Z}_{>0}$ and $A = \{(a_i)_{i=1}^{n} \in \mathbb{Z}^n : a_i \equiv a_j \mod 2 \text{ for all } i, j\}$ with componentwise addition and multiplication, then A is a connected order, $\mu(A) = \{(\pm 1, \ldots, \pm 1)\}$, and $\#\mu(A) = 2^n$. For large n, computing a set of generators for $\mu(A)$ is feasible, even when listing all elements of $\mu(A)$ is not.

Example 2.12. Suppose $A = \mathbb{Z}[\zeta_p]$, where p is a prime and ζ_p is a primitive p-th root of unity in $\mathbb{C}$. Then A has rank $p - 1$. If $p > 2$, then $\mu(A) = \langle \zeta_p \rangle \times \langle -1 \rangle$.

3. Finite $\mathbb{Q}$-algebras

The following two results are from commutative algebra. These results and basic algorithms for commutative $\mathbb{Q}$-algebras are given in [8].

Proposition 3.1. *If E is a commutative $\mathbb{Q}$-algebra, then the map*

$$E_{\mathrm{sep}} \oplus \sqrt{0} \xrightarrow{\sim} E, \quad (x, y) \mapsto x + y$$

is an isomorphism of $\mathbb{Q}$-vector spaces, and the natural map $E \to \prod_{\mathfrak{m} \in \mathrm{Spec}(E)} E/\mathfrak{m}$ induces an isomorphism of $\mathbb{Q}$-algebras

$$E_{\mathrm{sep}} \xrightarrow{\sim} \prod_{\mathfrak{m} \in \mathrm{Spec}(E)} E/\mathfrak{m}.$$

In algorithms, we specify a commutative $\mathbb{Q}$-algebra E by listing a system of structure constants $a_{ijk} \in \mathbb{Q}$ that determines the multiplication in E with respect to some $\mathbb{Q}$-basis, just as we did for orders in the introduction.

Algorithm 3.2. There is a deterministic polynomial-time algorithm that given a commutative $\mathbb{Q}$-algebra E, computes a $\mathbb{Q}$-basis for $E_{\mathrm{sep}} \subset E$, a $\mathbb{Q}$-basis for $\sqrt{0}$, the map $E \xrightarrow{\sim} E_{\mathrm{sep}} \oplus \sqrt{0}$ that is the inverse to the first isomorphism from Proposition 3.1, all $\mathfrak{m} \in \mathrm{Spec}(E)$, the fields $E/\mathfrak{m}$, and the natural maps $E \to E/\mathfrak{m}$.

Lemma 3.3. *If E is a commutative $\mathbb{Q}$-algebra, then:*
  (i) *$\mu(E) = \mu(E_{\mathrm{sep}}) \xrightarrow{\sim} \bigoplus_{\mathfrak{m}\in\mathrm{Spec}(E)} \mu(E/\mathfrak{m})$;*
  (ii) *$\mu(E)$ is finite;*
  (iii) *each $\mu(E/\mathfrak{m})$ is a finite cyclic group;*
  (iv) *if $\mu(E)$ has an element of order p^k with p a prime, then $\varphi(p^k) \leq \dim_{\mathbb{Q}}(E)$, where φ is Euler's φ-function.*

Proof. Part (i) holds by Proposition 3.1 and the fact that $X^r - 1$ is separable over $\mathbb{Q}$ for all $r \in \mathbb{Z}_{>0}$. If $\mu(E)$ has an element of prime power order p^k, then $\mathbb{Q}(\zeta_{p^k}) \subset E/\mathfrak{m}$ for some $\mathfrak{m}$, where ζ_{p^k} is a primitive p^k-th root of unity. Thus $\varphi(p^k) \leq [E/\mathfrak{m} : \mathbb{Q}] \leq \dim_{\mathbb{Q}}(E)$. Since each $E/\mathfrak{m}$ is a number field, $\mu(E/\mathfrak{m})$ is cyclic.
$\square$

Algorithm 3.4. The algorithm takes as input a commutative $\mathbb{Q}$-algebra E and produces a set of generators S of $\mu(E)$ as well as a set R of defining relations for S.

  (i) For each $\mathfrak{n} \in \mathrm{Spec}(E)$, use the algorithm in [4] to find all zeroes of $X^r - 1$ over $E/\mathfrak{n}$, for $r = 1, 2, \ldots, 2[E/\mathfrak{n} : \mathbb{Q}]^2$, let $\zeta_{\mathfrak{n}} \in (E/\mathfrak{n})^*$ be an element of maximal order among the zeroes found, and let $k(\mathfrak{n})$ be its order.
  (ii) For each $\mathfrak{n} \in \mathrm{Spec}(E)$, use linear algebra to compute the unique element $\eta_{\mathfrak{n}} \in E_{\mathrm{sep}}$ that under the second isomorphism from Proposition 3.1 maps to $(1, \ldots, 1, \zeta_{\mathfrak{n}}, 1, \ldots, 1) \in \prod_{\mathfrak{m}} \mu(E/\mathfrak{m})$ (with $\zeta_{\mathfrak{n}}$ in the $\mathfrak{n}$-th position). Output $S = \{\eta_{\mathfrak{n}} \in \mu(E) : \mathfrak{n} \in \mathrm{Spec}(E)\}$ and $R = \{(0, \ldots, 0, k(\mathfrak{n}), 0, \ldots, 0) \in \mathbb{Z}^{\mathrm{Spec}(E)} : \mathfrak{n} \in \mathrm{Spec}(E)\}$.

Proposition 3.5. *Algorithm 3.4 produces correct output and runs in polynomial time.*

Proof. If the number field $E/\mathfrak{n}$ contains a primitive r-th root of unity, then it contains the r-th cyclotomic field, which has degree $\varphi(r)$ over $\mathbb{Q}$; hence $\varphi(r) \leq [E/\mathfrak{n} : \mathbb{Q}]$ and $r \leq 2\varphi(r)^2 \leq 2[E/\mathfrak{n} : \mathbb{Q}]^2$. Together with Lemma 3.3(i), this implies that the algorithm is correct. It runs in polynomial time by [4]. $\square$

Algorithm 3.6. The algorithm takes as input a commutative $\mathbb{Q}$-algebra E, an element $\gamma \in E$, and a set $S = \{\eta_{\mathfrak{n}} \in \mu(E) : \mathfrak{n} \in \mathrm{Spec}(E)\}$ of generators for $\mu(E)$ as computed by Algorithm 3.4. It tests whether $\gamma \in \mu(E)$, and if so, finds $(a_{\mathfrak{n}})_{\mathfrak{n}\in\mathrm{Spec}(E)} \in \mathbb{Z}^{\mathrm{Spec}(E)}$ with $\gamma = \prod_{\mathfrak{n}\in\mathrm{Spec}(E)} \eta_{\mathfrak{n}}^{a_{\mathfrak{n}}}$.

  (i) Use linear algebra to test if $\gamma \in E_{\mathrm{sep}}$. If not, terminate with "no" (that is, $\gamma \notin \mu(E)$).
  (ii) Otherwise, for each $\mathfrak{n} \in \mathrm{Spec}(E)$ compute the image $\gamma_{\mathfrak{n}}$ of γ in $E/\mathfrak{n}$, and let $\zeta_{\mathfrak{n}}$ (as in Algorithm 3.4) be the image of $\eta_{\mathfrak{n}}$ in $E/\mathfrak{n}$. Try $a = 0, 1, 2, \ldots, \#\mu(E/\mathfrak{n}) - 1$ until $\gamma_{\mathfrak{n}} = \zeta_{\mathfrak{n}}^a$, and let $a_{\mathfrak{n}} = a$. If for some $\mathfrak{n}$ no $a_{\mathfrak{n}}$ exists, terminate with "no".
  (iii) Otherwise, output $(a_{\mathfrak{n}})_{\mathfrak{n}\in\mathrm{Spec}(E)}$.

That Algorithm 3.6 produces correct output and runs in polynomial time follows from Lemma 3.3, since $\mu(E/\mathfrak{n}) = \langle \zeta_{\mathfrak{n}} \rangle$.

4. ORDERS

From now on, suppose that A is an order. Let

$$E = A_{\mathbb{Q}} = A \otimes_{\mathbb{Z}} \mathbb{Q}, \qquad A_{\mathrm{sep}} = A \cap E_{\mathrm{sep}}.$$

Since $E_{\mathrm{sep}}/A_{\mathrm{sep}} \subset E/A = A_{\mathbb{Q}}/A$ is a torsion group, one has $E_{\mathrm{sep}} = (A_{\mathrm{sep}})_{\mathbb{Q}}$.

Lemma 4.1. *We have* $\mathrm{id}(E_{\mathrm{sep}}) = \mathrm{id}(E)$, $\mathrm{id}(A_{\mathrm{sep}}) = \mathrm{id}(A)$, *and* $\mu(A_{\mathrm{sep}}) = \mu(A)$.

Proof. This holds because the polynomials $X^2 - X$ and $X^r - 1$ are separable over $\mathbb{Q}$ for all $r \in \mathbb{Z}_{>0}$. $\square$

Algorithm 4.2. The algorithm takes as input an order A and it computes the $\mathbb{Q}$-algebras E and $E_{\mathrm{sep}} \subset E$, as well as the order $A_{\mathrm{sep}} = A \cap E_{\mathrm{sep}}$, giving a $\mathbb{Z}$-basis for A_{sep} expressed both in the given $\mathbb{Z}$-basis of A and in the $\mathbb{Q}$-basis for E_{sep}.

  (i) We use the given $\mathbb{Z}$-basis for A as a $\mathbb{Q}$-basis for E, with the same structure constants.
  (ii) Let $\pi_1 : A \to E_{\mathrm{sep}}$ and $\pi_2 : A \to \sqrt{0}$ be the compositions of the inclusion $A \subset E$ with the map $E \xrightarrow{\sim} E_{\mathrm{sep}} \oplus \sqrt{0}$ from Algorithm 3.2 followed by the natural projections to E_{sep} and $\sqrt{0}$, respectively. Using Algorithm 3.2, compute a $\mathbb{Q}$-basis for E_{sep} and the rational matrices describing π_1 and π_2. Applying the kernel algorithm in §14 of [5] to an integer multiple of the matrix for π_2, compute a $\mathbb{Z}$-basis for $A_{\mathrm{sep}} = \ker(\pi_2)$ expressed in the given $\mathbb{Z}$-basis for A. Applying π_1 to this $\mathbb{Z}$-basis, one obtains the same $\mathbb{Z}$-basis expressed in the $\mathbb{Q}$-basis for E_{sep}.

Algorithm 4.2 is clearly correct and polynomial time.

5. Graphs attached to rings

Lemma 5.1. *Suppose that R is a commutative ring, $\mathcal{S}$ is a finite set of ideals of R that are not R itself, and suppose that $\bigcap_{\mathfrak{a} \in \mathcal{S}} \mathfrak{a} = \{0\}$. Identify R with its image in $\prod_{\mathfrak{a} \in \mathcal{S}} R/\mathfrak{a}$. Suppose that $e = (e_{\mathfrak{a}})_{\mathfrak{a} \in \mathcal{S}} \in \{0,1\}^{\mathcal{S}} \subset \prod_{\mathfrak{a} \in \mathcal{S}} R/\mathfrak{a}$. Then $e \in R$ if and only if $e_{\mathfrak{a}} = e_{\mathfrak{b}}$ in $\{0,1\}$ for all $\mathfrak{a}, \mathfrak{b} \in \mathcal{S}$ such that $\mathfrak{a} + \mathfrak{b} \neq R$.*

Proof. First suppose $e \in R$. Suppose $\mathfrak{a}, \mathfrak{b} \in \mathcal{S}$ and $\mathfrak{a} + \mathfrak{b} \neq R$. Choose $e'_{\mathfrak{a}} \in \{0,1\} \subset R$ whose image in $R/\mathfrak{a}$ is $e_{\mathfrak{a}} = e + \mathfrak{a}$, and choose $e'_{\mathfrak{b}} \in \{0,1\} \subset R$ whose image in $R/\mathfrak{b}$ is $e_{\mathfrak{b}} = e + \mathfrak{b}$. Then $e'_{\mathfrak{a}} \equiv e$ mod $\mathfrak{a}$ and $e'_{\mathfrak{b}} \equiv e$ mod $\mathfrak{b}$, so $e'_{\mathfrak{a}} \equiv e \equiv e'_{\mathfrak{b}}$ mod $(\mathfrak{a} + \mathfrak{b})$. Since $\mathfrak{a} + \mathfrak{b} \neq R$ we have $1 \notin \mathfrak{a} + \mathfrak{b}$. Thus, $e'_{\mathfrak{a}} = e'_{\mathfrak{b}}$ in $\{0,1\}$, as desired.

Conversely, suppose that $e_{\mathfrak{a}} = e_{\mathfrak{b}}$ in $\{0,1\}$ for all $\mathfrak{a}, \mathfrak{b} \in \mathcal{S}$ with $\mathfrak{a} + \mathfrak{b} \neq R$. Let $T = \{\mathfrak{a} \in \mathcal{S} : e_{\mathfrak{a}} = 1\}$ and $U = \{\mathfrak{b} \in \mathcal{S} : e_{\mathfrak{b}} = 0\}$. Then $\mathcal{S} = T \sqcup U$. Pick $\mathfrak{a} \in T$ and $\mathfrak{b} \in U$. By our assumption, $\mathfrak{a} + \mathfrak{b} = R$. Thus, there exist $x_{\mathfrak{a},\mathfrak{b}} \in \mathfrak{a}$ and $y_{\mathfrak{a},\mathfrak{b}} \in \mathfrak{b}$ such that $1 = x_{\mathfrak{a},\mathfrak{b}} + y_{\mathfrak{a},\mathfrak{b}}$. It follows that $y_{\mathfrak{a},\mathfrak{b}} \equiv 1$ mod $\mathfrak{a}$ and $y_{\mathfrak{a},\mathfrak{b}} \equiv 0$ mod $\mathfrak{b}$. For all $\mathfrak{a} \in T$, define $z_{\mathfrak{a}} = \prod_{\mathfrak{b} \in U} y_{\mathfrak{a},\mathfrak{b}} \in R$. Then $z_{\mathfrak{a}} \equiv 1$ mod $\mathfrak{a}$ and $z_{\mathfrak{a}} \equiv 0$ modulo each $\mathfrak{b} \in U$. Define $e' = 1 - \prod_{\mathfrak{a} \in T}(1 - z_{\mathfrak{a}}) \in R$. Then $e' \equiv 1$ modulo each $\mathfrak{a} \in T$, and $e' \equiv 0$ modulo each $\mathfrak{b} \in U$. Thus, $e' \equiv e_{\mathfrak{a}}$ mod $\mathfrak{a}$ for each $\mathfrak{a} \in \mathcal{S}$, so $e' = e$. $\square$

We say that D is an order in a separable $\mathbb{Q}$-algebra if D is an order and $D_{\mathbb{Q}} = D \otimes_{\mathbb{Z}} \mathbb{Q}$ is separable.

Definition 5.2. Suppose that D is an order in a separable $\mathbb{Q}$-algebra $D_{\mathbb{Q}}$. For $\mathfrak{m}, \mathfrak{n} \in \mathrm{Spec}(D_{\mathbb{Q}})$ with $\mathfrak{m} \neq \mathfrak{n}$, let

$$n(D, \mathfrak{m}, \mathfrak{n}) = \#(D/((\mathfrak{m} \cap D) + (\mathfrak{n} \cap D))),$$

and let $\Gamma(D)$ denote the graph on $\mathrm{Spec}(D_{\mathbb{Q}})$ defined by connecting distinct vertices $\mathfrak{m}, \mathfrak{n} \in \mathrm{Spec}(D_{\mathbb{Q}})$ by an edge if and only if $n(D, \mathfrak{m}, \mathfrak{n}) > 1$.

Lemma 5.3. $n(D, \mathfrak{m}, \mathfrak{n}) \in \mathbb{Z}_{>0}$.

Proof. Let $R = D/((\mathfrak{m} \cap D) + (\mathfrak{n} \cap D))$. Then $n(D, \mathfrak{m}, \mathfrak{n}) = \#R$. Letting $-_{\mathbb{Q}} = - \otimes_{\mathbb{Z}} \mathbb{Q}$, we have

$$R_{\mathbb{Q}} = D_{\mathbb{Q}}/((\mathfrak{m}_{\mathbb{Q}} \cap D_{\mathbb{Q}}) + (\mathfrak{n}_{\mathbb{Q}} \cap D_{\mathbb{Q}})) = D_{\mathbb{Q}}/(\mathfrak{m} + \mathfrak{n}) = 0$$

so R is torsion. Since R is finitely generated as an abelian group, it is finite, so $n(D, \mathfrak{m}, \mathfrak{n}) \in \mathbb{Z}_{>0}$. $\square$

Example 5.4. Let $r \in \mathbb{Z}[X]$ be monic. Then $D = \mathbb{Z}[X]/(f)$ is an order in a separable $\mathbb{Q}$-algebra if and only if f is squarefree. Suppose f is squarefree. Then $D_{\mathbb{Q}} = \mathbb{Q}[X]/(f)$, and $\mathrm{Spec}(D_{\mathbb{Q}})$ is in bijection with the set of monic irreducible factors g of f in $\mathbb{Z}[X]$, each g corresponding to $\mathfrak{m} = (g)/(f)$. If g, h correspond to $\mathfrak{m}, \mathfrak{n}$, respectively, then $n(D, \mathfrak{m}, \mathfrak{n}) = |R(g, h)|$, with R denoting the resultant.

Suppose D is an order in a separable $\mathbb{Q}$-algebra. It is natural to ask whether the decomposition $D_{\mathbb{Q}} \xrightarrow{\sim} \prod_{\mathfrak{m} \in \mathrm{Spec}(D_{\mathbb{Q}})} D_{\mathbb{Q}}/\mathfrak{m}$ (Proposition 3.1) gives rise to a decomposition of the order D. This depends on the idempotents that are present in D. The graph $\Gamma(D)$ tells us which idempotents occur in D (see Lemma 5.1 and Proposition 5.7).

Notation 5.5. Suppose that D is an order in a separable $\mathbb{Q}$-algebra. If $W \subset \mathrm{Spec}(D_{\mathbb{Q}})$, define

$$e_W = (e_{\mathfrak{m}})_{\mathfrak{m} \in \mathrm{Spec}(D_{\mathbb{Q}})} \in \mathrm{id}(\prod_{\mathfrak{m} \in \mathrm{Spec}(D_{\mathbb{Q}})} D_{\mathbb{Q}}/\mathfrak{m}) = \{0, 1\}^{\mathrm{Spec}(D_{\mathbb{Q}})}$$

by $e_{\mathfrak{m}} = 1$ if $\mathfrak{m} \in W$ and $e_{\mathfrak{m}} = 0$ if $\mathfrak{m} \notin W$.

Algorithm 5.6. The algorithm takes an order D in a separable $\mathbb{Q}$-algebra and computes the graph $\Gamma(D)$, its connected components, and its weights $n(D, \mathfrak{m}, \mathfrak{n})$ for all $\mathfrak{m}, \mathfrak{n} \in \mathrm{Spec}(D_{\mathbb{Q}})$.

   (i) Use Algorithm 3.2 to compute $\mathrm{Spec}(D_{\mathbb{Q}})$ and the maps $D_{\mathbb{Q}} \to D_{\mathbb{Q}}/\mathfrak{m}$ for $\mathfrak{m} \in \mathrm{Spec}(D_{\mathbb{Q}})$.
   (ii) For each $\mathfrak{m} \in \mathrm{Spec}(D_{\mathbb{Q}})$ compute $\mathfrak{m} \cap D = \ker(D \to D_{\mathbb{Q}}/\mathfrak{m})$ by applying the kernel algorithm in §14 of [5].
   (iii) For all $\mathfrak{m} \neq \mathfrak{n} \in \mathrm{Spec}(D_{\mathbb{Q}})$, apply the image algorithm in §14 of [5] to compute a $\mathbb{Z}$-basis of

$$\mathrm{image}((\mathfrak{m} \cap D) \oplus (\mathfrak{n} \cap D) \to D) = (\mathfrak{m} \cap D) + (\mathfrak{n} \cap D)$$

   expressed in a $\mathbb{Z}$-basis of D, and compute $n(D, \mathfrak{m}, \mathfrak{n})$ as the absolute value of the determinant of the matrix whose columns are those basis vectors.
   (iv) Use the numbers $n(D, \mathfrak{m}, \mathfrak{n})$ to obtain the graph $\Gamma(D)$ and its connected components.

The algorithm runs in polynomial time by well-known graph algorithms (see for example [2]).

Proposition 5.7. *Suppose that D is an order in a separable $\mathbb{Q}$-algebra.*
   (i) *Suppose $e = (e_{\mathfrak{m}})_{\mathfrak{m} \in \mathrm{Spec}(D_{\mathbb{Q}})} \in \mathrm{id}(\prod_{\mathfrak{m}} D_{\mathbb{Q}}/\mathfrak{m}) = \{0, 1\}^{\mathrm{Spec}(D_{\mathbb{Q}})}$. Then the following are equivalent:*
   (a) *$e \in D$,*
   (b) *$e_{\mathfrak{m}} = e_{\mathfrak{n}}$ whenever $\mathfrak{m}$ and $\mathfrak{n}$ are connected in $\Gamma(D)$,*
   (c) *$e_{\mathfrak{m}} = e_{\mathfrak{n}}$ whenever $\mathfrak{m}$ and $\mathfrak{n}$ are in the same connected component of $\Gamma(D)$.*
   (ii) *Let Ω denote the set of connected components of the graph $\Gamma(D)$ and recall e_W from Definition 5.5. Then $W \mapsto e_W$ gives a bijection*

$$\Omega \xrightarrow{\sim} \mathrm{prid}(D) \subset D \subset \prod_{\mathfrak{m} \in \mathrm{Spec}(D_{\mathbb{Q}})} D_{\mathbb{Q}}/\mathfrak{m}.$$

Proof. Apply Lemma 5.1 with $R = D$ and $S = \{\mathfrak{m} \cap D : \mathfrak{m} \in \mathrm{Spec}(D_{\mathbb{Q}})\}$. We have $\bigcap_{\mathfrak{a} \in \mathcal{S}} \mathfrak{a} = \bigcap_{\mathfrak{m}}(\mathfrak{m} \cap D) = \{0\}$ since D injects into $\prod_{\mathfrak{m}} D_{\mathbb{Q}}/\mathfrak{m}$. Identifying $\mathrm{id}(\prod D_{\mathbb{Q}}/\mathfrak{m})$ with $\{0, 1\}^{\mathcal{S}}$, Lemma 5.1 implies that if $e = (e_{\mathfrak{m}})_{\mathfrak{m} \in \mathrm{Spec}(D_{\mathbb{Q}})} \in \mathrm{id}(\prod D_{\mathbb{Q}}/\mathfrak{m})$, then $e \in D$ if and only if $e_{\mathfrak{m}} = e_{\mathfrak{n}}$ for all $\mathfrak{m}, \mathfrak{n} \in \mathrm{Spec}(D_{\mathbb{Q}})$ that are connected in $\Gamma(D)$. It follows that for each $e = (e_{\mathfrak{m}})_{\mathfrak{m}} \in \mathrm{id}(D)$ the components $e_{\mathfrak{m}}$ are constant (0 or 1) on each connected component of $\Gamma(D)$. Part (i) now follows. It also follows that there is a bijection

$$\{\text{subsets of } \Omega\} \to \mathrm{id}(D)$$

defined by $T \mapsto \sum_{W \in T} e_W$ with inverse $e = (e_{\mathfrak{m}})_{\mathfrak{m}} \mapsto \{W \in \Omega : e_{\mathfrak{m}} = 1 \text{ for all } \mathfrak{m} \in W\}$. Under this bijection, $\mathrm{prid}(D)$ corresponds to Ω, and this gives the bijection in (ii). $\square$

Remark 5.8. In particular, by Proposition 5.7(ii) an order D is connected if and only if $\Gamma(D)$ is connected.

6. Finding idempotents

The set of idempotents of an order may be too large to compute, but the set of primitive idempotents is something that we are able to efficiently compute.

Algorithm 6.1. Given an order A, the algorithm outputs the set of primitive idempotents of A.

   (i) Use Algorithm 4.2 to compute A_{sep}.

   (ii) Use Algorithm 5.6 to compute the graph $\Gamma(A_{\mathrm{sep}})$ and its connected components.

   (iii) For each connected component W of $\Gamma(A_{\mathrm{sep}})$, with $e_W \in \{0,1\}^{\mathrm{Spec}(E)} \subset \prod_{\mathfrak{m} \in \mathrm{Spec}(E)} E/\mathfrak{m}$ as in Notation 5.5, use the inverse of the square matrix with $\mathbb{Q}$-coefficients that gives the natural map $E_{\mathrm{sep}} \xrightarrow{\sim} \prod_{\mathfrak{m} \in \mathrm{Spec}(E)} E/\mathfrak{m}$ of Proposition 3.1 to lift e_W to E_{sep}. Output these lifts.

If follows from Proposition 5.7(ii) that the lift e_W to E_{sep} is in A_{sep}, and that Algorithm 6.1 gives the desired output $\mathrm{prid}(A)$. It is clear that it runs in polynomial time.

7. Discrete logarithms

In this section, we suppose that G is a multiplicatively written abelian group with elements represented by finite bitstrings. All algorithms in the present section have G as part of their input. Thus, saying that they are polynomial-time means that their runtime is bounded by a polynomial function of the length of the parameters specifying G plus the length of the rest of the input. We suppose that polynomial-time algorithms for the group operations and for equality testing in G are available.

Definition 7.1. We say $\langle S|R \rangle$ is an **efficient presentation** for G if S is a finite set, and we have a map $f = f_S : S \to G$ satisfying:

   (a) $f(S)$ generates G, i.e., the map $g_S : \mathbb{Z}^S \to G$, $(b_s)_{s \in S} \mapsto \prod_{s \in S} f(s)^{b_s}$ is surjective,

   (b) $R \subset \mathbb{Z}^S$ is a finite set of generators for $\ker(g_S)$,

   (c) we have a polynomial-time algorithm that on input $\gamma \in G$ finds an element of $g_S^{-1}(\gamma)$ (i.e., finds $(c_s)_{s \in S} \in \mathbb{Z}^S$ such that $\gamma = \prod_{s \in S} f(s)^{c_s}$).

Notation 7.2. Suppose $\langle S|R \rangle$ is an efficient presentation for G. Define

$$\rho : \mathbb{Z}^R \to \mathbb{Z}^S, \quad \rho((m_r)_{r \in R}) = \sum_{r \in R} m_r r.$$

Suppose T is a finite set and we have a map $f_T : T \to G$. By abuse of notation we usually suppress the maps f_S and f_T and write s for $f_S(s)$ and $f_T(s)$ and write $\langle T \rangle$ for $\langle f_T(T) \rangle$. Define

$$g_T : \mathbb{Z}^T \to \langle T \rangle, \quad (b_t)_{t \in T} \mapsto \prod_{t \in T} t^{b_t}.$$

Define $h = h_T : \mathbb{Z}^T \to \mathbb{Z}^S$ by using (c) to write each $t \in T$ as $t = \prod_{s \in S} s^{c_{s,t}}$ and defining

$$h((b_t)_{t \in T}) = \left(\sum_{t \in T} b_t c_{s,t} \right)_{s \in S} \in \mathbb{Z}^S$$

so that $g_T = g_S \circ h$.

For the remainder of this section we suppose that an efficient presentation $\langle S|R \rangle$ for an abelian group G is given.

Algorithm 7.3. The algorithm takes as input G, an efficient presentation $\langle S|R \rangle$ for G, and a finite set T with a map $T \to G$, and outputs a finite set $U = U_T$ of generators for $\ker(g_T)$.

(i) Define $h - \rho : \mathbb{Z}^T \times \mathbb{Z}^R \to \mathbb{Z}^S$ by $(h - \rho)(x, y) = h(x) - \rho(y)$. Use the kernel algorithm in §14 of [5] to compute a finite set V of generators for $\ker(h - \rho)$.

(ii) Compute the image U of V under the projection map $\mathbb{Z}^T \times \mathbb{Z}^R \twoheadrightarrow \mathbb{Z}^T$, $(x, y) \mapsto x$.

Theorem 7.4. *Algorithm 7.3 produces correct output and runs in polynomial time.*

Proof. We have:

$$
\begin{aligned}
x \in \ker(g_T) \iff & h(x) \in \ker(g_S) = \mathrm{im}(\rho) \\
\iff & \exists y \in \mathbb{Z}^R \text{ such that } h(x) = \rho(y) \\
\iff & \exists y \in \mathbb{Z}^R \text{ such that } (h - \rho)(x, y) = 0 \\
\iff & \exists y \in \mathbb{Z}^R \text{ such that } (x, y) \in \langle V \rangle \\
\iff & x \in \mathrm{proj}(\langle V \rangle) = \langle \mathrm{proj}(V) \rangle = \langle U \rangle.
\end{aligned}
$$

$\square$

Algorithm 7.5. The algorithm takes as input G, an efficient presentation $\langle S | R \rangle$ for G, a finite set T with a map $T \to G$, and an element $\gamma \in G$, and decides whether $\gamma \in \langle T \rangle$, and if it is, produces an element of $g_T^{-1}(\gamma)$ (i.e., finds $(c_t)_{t \in T} \in \mathbb{Z}^T$ such that $\gamma = \prod_{t \in T} t^{c_t}$).

(i) Apply Algorithm 7.3 with $T \cup \{\gamma\}$ in place of T to find a finite set of generators $U_{T \cup \{\gamma\}} \subset \mathbb{Z}^{T \cup \{\gamma\}}$ for $\ker(g_{T \cup \{\gamma\}})$, where

$$
g_{T \cup \{\gamma\}} : \mathbb{Z}^{T \cup \{\gamma\}} = \mathbb{Z}^T \times \mathbb{Z}^{\{\gamma\}} \to G, \qquad (x, n) \mapsto g_T(x) \gamma^n.
$$

(ii) Map the elements $u \in U_{T \cup \{\gamma\}} \subset \mathbb{Z}^{T \cup \{\gamma\}} = \mathbb{Z}^T \times \mathbb{Z}^{\{\gamma\}}$ to their $\mathbb{Z}^{\{\gamma\}}$-components $u(\gamma) \in \mathbb{Z}$. If $\sum_{u \in U_{T \cup \{\gamma\}}} u(\gamma) \mathbb{Z} \neq \mathbb{Z}$ then $\gamma \notin \langle T \rangle$; if $1 = \sum_{u \in U_{T \cup \{\gamma\}}} n_u u(\gamma)$ with $(n_u)_{u \in U_{T \cup \{\gamma\}}} \in \mathbb{Z}^{U_{T \cup \{\gamma\}}}$ then $\gamma \in \langle T \rangle$ and the $\mathbb{Z}^T$-component of $-\sum_{u \in U_{T \cup \{\gamma\}}} n_u u \in \mathbb{Z}^{T \cup \{\gamma\}} = \mathbb{Z}^T \times \mathbb{Z}^{\{\gamma\}}$ is in $g_T^{-1}(\gamma)$.

Algorithm 7.6. The algorithm takes as input G, an efficient presentation $\langle S | R \rangle$ for G, and a finite set T with a map $T \to G$, and outputs an efficient presentation $\langle T | U_T \rangle$ for $\langle T \rangle$.

(i) Apply Algorithm 7.3 to obtain a set U_T of relations.

(ii) Output the presentation $\langle T | U_T \rangle$.

Theorem 7.7. *Algorithms 7.5 and 7.6 produce correct output and run in polynomial time. In particular, if one has an efficient presentation for G, and T is a finite set with a map $T \to G$, then $\langle T | U_T \rangle$ is an efficient presentation for $\langle T \rangle$.*

Proof. We have:

$$
\begin{aligned}
\gamma \in \langle T \rangle \iff & \exists x \in \mathbb{Z}^T \text{ such that } \gamma = g_T(x) \\
\iff & \exists x \in \mathbb{Z}^T \text{ such that } (-x, 1) \in \ker(g_{T \cup \{\gamma\}} : \mathbb{Z}^T \times \mathbb{Z} \to G) = \langle U_{T \cup \{\gamma\}} \rangle \\
\iff & 1 \in \mathrm{im}(\mathrm{proj} : \langle U_{T \cup \{\gamma\}} \rangle \subset \mathbb{Z}^T \times \mathbb{Z} \to \mathbb{Z}) \\
\iff & \exists (n_u)_{u \in U_{T \cup \{\gamma\}}}, \exists x \in \mathbb{Z}^T \text{ such that } \sum_u n_u u = (-x, 1)
\end{aligned}
$$

where proj is projection onto the second component. $\square$

Algorithm 7.8. The algorithm takes as input G, an efficient presentation $\langle S | R \rangle$ for G, finite sets T and T', and maps $f_T : T \to G$ and $f_{T'} : T' \to G$, and outputs a finite set of generators for the kernel of the composition $\mathbb{Z}^T \to G \to G / \langle T' \rangle$, where $\mathbb{Z}^T \to G$ is the map g_T.

(i) Apply Algorithm 7.3 to the finite set $T \sqcup T'$ and the map $T \sqcup T' \to G$ obtained from f_T and $f_{T'}$, to obtain generators for the kernel of the map

$$\mathbb{Z}^T \times \mathbb{Z}^{T'} = \mathbb{Z}^{T \sqcup T'} \to G, \qquad (x, y) \mapsto g_T(x) - g_{T'}(y).$$

(ii) Project these generators to their $\mathbb{Z}^T$-component.

Theorem 7.9. *Algorithm 7.8 produces correct output and runs in polynomial time.*

Proof. We have:

$$
\begin{aligned}
x \in \ker(\mathbb{Z}^T \to G/\langle T' \rangle) &\iff g_T(x) \in \langle T' \rangle = \operatorname{im}(g_{T'}) \\
&\iff \exists y \in \mathbb{Z}^{T'} \text{ such that } g_T(x) = g_{T'}(y) \\
&\iff \exists y \in \mathbb{Z}^{T'} \text{ such that } (x, y) \in \ker(\mathbb{Z}^T \times \mathbb{Z}^{T'} \to G) \\
&\iff x \in \operatorname{proj}(\ker(\mathbb{Z}^T \times \mathbb{Z}^{T'} \to G) \to \mathbb{Z}^T)
\end{aligned}
$$

where proj denotes projection onto the $\mathbb{Z}^T$-component. $\qquad \square$

Proof of Theorem 1.3. One starts by computing $E = A \otimes_{\mathbb{Z}} \mathbb{Q}$, using the same structure constants as for A. Algorithm 3.4 produces a presentation for $\mu(E)$, and by Algorithm 3.6 this is an efficient presentation. Given T and ζ as in Theorem 1.3, one can test whether $\zeta \in E$ by Algorithm 3.6. Now Theorem 1.3 is obtained from Algorithm 7.5, with $G = \mu(E)$ and $\gamma = \zeta$.

8. Separable polynomials over connected rings

Proposition 8.1(b) will be used to prove Proposition 10.5 below.

Proposition 8.1. *Suppose R is a connected commutative ring, $f \in R[X]$, and $R[X]f + R[X]f' = R[X]$. Then:*

   (a) *if $r, s \in R$ and $f(r) = f(s) = 0$, then $r - s \in \{0\} \cup R^*$;*
   (b) *if S is a non-zero ring and $\varphi : R \to S$ is a ring homomorphism, then the restriction of φ to $\{r \in R : f(r) = 0\}$ is injective;*
   (c) *$f \neq 0$ and $\#\{r \in R : f(r) = 0\} \leq \deg(f)$.*

Proof. Suppose $f(r) = f(s) = 0$. Write $f = (X - r)g$ and $1 = hf + kf'$ with $g, h, k \in R[X]$. Then $g(r) = f'(r) \in R^*$. Since $g(s) \equiv g(r) \bmod (r - s)R$ we can write $g(s) = g(r) + (r - s)t$ with $t \in R$. Thus, $0 = f(s) = (s - r)g(s) = (s - r)(g(r) + (r - s)t)$, so

(8.1) $$(s - r)g(r) = t(s - r)^2.$$

Thus, $t \cdot (s - r) \cdot g(r)^{-1} = (t \cdot (s - r) \cdot g(r)^{-1})^2$, an idempotent. If $t \cdot (s - r) \cdot g(r)^{-1} = 0$, then by (8.1) we have $(s - r)g(r) = 0$, and thus $r - s = 0$ since $g(r) \in R^*$. If $t \cdot (s - r) \cdot g(r)^{-1} = 1$, then $r - s \in R^*$. This gives (a).

For (b), suppose $r, s \in R$, $r \neq s$, and $f(r) = f(s) = 0$. By (a) we have $r - s \in R^*$. Since $\varphi(1) = 1 \neq 0$, we have $\varphi(r - s) \neq 0$.

For (c), let $\mathfrak{m}$ be a maximal ideal of R. Then $R \to R/\mathfrak{m}$ induces a map

$$\{r \in R : f(r) = 0\} \to \{u \in R/\mathfrak{m} : (f \bmod \mathfrak{m})(u) = 0\}$$

that is injective by (b). Since $R/\mathfrak{m}$ is a field and $f \bmod \mathfrak{m} \in (R/\mathfrak{m})[X]$ is non-zero, we have

$$\#\{r \in R : f(r) = 0\} \leq \deg(f \bmod \mathfrak{m}) \leq \deg(f).$$

$\qquad \square$

Corollary 8.2. *Suppose R is a connected commutative ring, $m \in \mathbb{Z}_{>0}$, and $m \cdot 1 \in R^*$. Then $\{\zeta \in R : \zeta^m = 1\}$ is a cyclic subgroup of R^* whose order divides m.*

Proof. Applying Proposition 8.1 with $f = X^m - 1$ gives that the subgroup has order dividing m. Applying Proposition 8.1 with $f = X^d - 1$ for each divisor d of m gives that this abelian subgroup has at most d elements of order dividing d, and thus is cyclic. $\qquad\square$

9. FROM $\mu(E)$ TO $\mu(B)$

Fix an order A. Recall that $E = A_{\mathbb{Q}} = A \otimes_{\mathbb{Z}} \mathbb{Q}$ and $A_{\mathrm{sep}} = A \cap E_{\mathrm{sep}}$. For $\mathfrak{m} \in \mathrm{Spec}(E)$, the image of A_{sep} in $E/\mathfrak{m}$ may be identified with $A_{\mathrm{sep}}/(\mathfrak{m} \cap A_{\mathrm{sep}})$; it is a ring of which the additive group is a finitely generated subgroup of the $\mathbb{Q}$-vector space $E/\mathfrak{m}$, so it is an order. We now write

$$(9.1) \qquad B = \prod_{\mathfrak{m} \in \mathrm{Spec}(E)} A_{\mathrm{sep}}/(\mathfrak{m} \cap A_{\mathrm{sep}}).$$

This is an order in $\prod_{\mathfrak{m} \in \mathrm{Spec}(E)} E/\mathfrak{m}$. We identify A_{sep} with its image in B under the map

$$E_{\mathrm{sep}} \xrightarrow{\sim} \prod_{\mathfrak{m} \in \mathrm{Spec}(E)} E/\mathfrak{m}$$

and identify B with a subring of E_{sep} using the same map. One has

$$A_{\mathrm{sep}} \subset B \subset E_{\mathrm{sep}}.$$

Since the abelian group B/A_{sep} is both torsion and finitely generated, it is finite, and one has $B_{\mathbb{Q}} = E_{\mathrm{sep}}$. The graph $\Gamma(B)$ consists of the vertices $\mathfrak{m} \in \mathrm{Spec}(E)$ and no edges.

Proposition 9.1. *There is a deterministic polynomial-time algorithm that, given an order A, computes a $\mathbb{Z}$-basis for $A_{\mathrm{sep}}/(\mathfrak{m} \cap A_{\mathrm{sep}})$ in $E/\mathfrak{m}$ for every $\mathfrak{m} \in \mathrm{Spec}(E)$, a $\mathbb{Z}$-basis for B in E_{sep}, and the index $(B : A_{\mathrm{sep}})$.*

Proof. One simply computes a $\mathbb{Z}$-basis for A_{sep} as in Algorithm 4.2, and a $\mathbb{Z}$-basis for the image of the map $A_{\mathrm{sep}} \subset E_{\mathrm{sep}} \to E/\mathfrak{m}$ using the image algorithm in §14 of [5], for each $\mathfrak{m} \in \mathrm{Spec}(E)$. Combining these bases for all $\mathfrak{m}$ and applying the inverse of the second isomorphism in Proposition 3.1 one finds a $\mathbb{Z}$-basis for B in E_{sep}. The index $(B : A_{\mathrm{sep}})$ is the absolute value of the determinant of any matrix expressing a $\mathbb{Z}$-basis for A_{sep} in a $\mathbb{Z}$-basis for B. $\qquad\square$

Proposition 9.2. *For each order A and each $\mathfrak{m} \in \mathrm{Spec}(E)$ the group $\mu(A_{\mathrm{sep}}/(\mathfrak{m} \cap A_{\mathrm{sep}}))$ is finite cyclic. Also, there is a deterministic polynomial-time algorithm that, given A and $\mathfrak{m}$, computes a generator $\theta_{\mathfrak{m}}$ of $\mu(A_{\mathrm{sep}}/(\mathfrak{m} \cap A_{\mathrm{sep}}))$, its order, the complete prime factorization of its order, and, for each prime number p a generator $\theta_{\mathfrak{m},p}$ for $\mu(A_{\mathrm{sep}}/(\mathfrak{m} \cap A_{\mathrm{sep}}))_p$.*

Proof. The first statement follows from Lemma 3.3(iii). For $\theta_{\mathfrak{m}}$ one can take the first power of the generator $\zeta_{\mathfrak{m}}$ of $\mu(E/\mathfrak{m})$ found in Algorithm 3.4 that belongs to $A_{\mathrm{sep}}/(\mathfrak{m} \cap A_{\mathrm{sep}})$, i.e., for which all coordinates on a $\mathbb{Z}$-basis of $A_{\mathrm{sep}}/(\mathfrak{m} \cap A_{\mathrm{sep}})$ (which is a $\mathbb{Q}$-basis of $E/\mathfrak{m}$) are integers. The order of $\theta_{\mathfrak{m}}$ is then easy to write down, and since the prime numbers dividing that order are, by Lemma 3.3(iv), bounded by $1 + \mathrm{rank}_{\mathbb{Z}}(A)$, it is also easy to factor into primes. If p^k is a prime power exactly dividing order($\theta_{\mathfrak{m}}$), one can take $\theta_{\mathfrak{m},p} = \theta_{\mathfrak{m}}^{\mathrm{order}(\theta_{\mathfrak{m}})/p^k}$. $\qquad\square$

Proposition 9.3. *There is a deterministic polynomial-time algorithm that, given an order A, determines all prime factors p of $\#\mu(B)$, with B as in (9.1), as well as an efficient presentation for $\mu(B)$ and, for each p, an efficient presentation for $\mu(B)_p$.*

Proof. This follows directly from Proposition 9.2 and the isomorphisms

$$\mu(B) \cong \prod_{\mathfrak{m} \in \mathrm{Spec}(E)} \mu(A_{\mathrm{sep}}/(\mathfrak{m} \cap A_{\mathrm{sep}})) \quad \text{and} \quad \mu(B)_p \cong \prod_{\mathfrak{m} \in \mathrm{Spec}(E)} \mu(A_{\mathrm{sep}}/(\mathfrak{m} \cap A_{\mathrm{sep}}))_p$$

in the same way as for $\mu(E)$ in section 3. $\qquad\square$

10. From $\mu(B)_p$ to $\mu(C)_p$

Let A, E, A_{sep}, and B be as in the previous section, and fix a prime number p. Let

$$(10.1) \qquad\qquad\qquad C = A_{\mathrm{sep}}\,[1/p] \cap B.$$

We have

$$A_{\mathrm{sep}} \subset C \subset B \subset E_{\mathrm{sep}}$$

so C is an order with $C_{\mathbb{Q}} = E_{\mathrm{sep}}$, and

$$C = \{x \in B : p^i x \in A_{\mathrm{sep}} \text{ for some } i \in \mathbb{Z}_{\geq 0}\}.$$

The group C/A_{sep} is finite of p-power order, and the group B/C is finite of order prime to p. These orders can be quickly computed from the order of B/A_{sep} computed in Proposition 9.1. We emphasize that C depends on p.

Let $t = (B : C)$. Then $C/A_{\mathrm{sep}} = t(B/A_{\mathrm{sep}})$, so $C = tB + A_{\mathrm{sep}}$, which is the image of the map $B \oplus A_{\mathrm{sep}} \to B$, $(x, y) \mapsto tx + y$. Thus one can find a $\mathbb{Z}$-basis for C from the image algorithm in §14 of [5].

Proposition 10.1. *Suppose that A is an order and p is a prime. Suppose $\mathfrak{m}, \mathfrak{n} \in \mathrm{Spec}(E)$ with $\mathfrak{m} \neq \mathfrak{n}$. Then:*

(i) *$C/((\mathfrak{m} \cap C) + (\mathfrak{n} \cap C))$ is the non-p-component of $A_{\mathrm{sep}}/((\mathfrak{m} \cap A_{\mathrm{sep}}) + (\mathfrak{n} \cap A_{\mathrm{sep}}))$;*

(ii) *$\mathfrak{m}$ and $\mathfrak{n}$ are connected in $\Gamma(C)$ if and only if $n(A_{\mathrm{sep}}, \mathfrak{m}, \mathfrak{n}) \notin p^{\mathbb{Z}_{\geq 0}}$.*

Proof. For $Z = A_{\mathrm{sep}}, B$, and C, write $\tilde{Z}$ for the finite abelian group $Z/((\mathfrak{m} \cap Z) + (\mathfrak{n} \cap Z))$ (cf. Lemma 5.3). Let $p^r = (C : A_{\mathrm{sep}})$ and $t = (B : C)$. Then $\gcd(p^r, t) = 1$. Since $\Gamma(B)$ has no edges, we have $(\mathfrak{m} \cap B) + (\mathfrak{n} \cap B) = B$, so $\tilde{B} = 0$. Consider the maps $\tilde{A}_{\mathrm{sep}} \underset{p^r}{\overset{1}{\rightleftarrows}} \tilde{C} \underset{t}{\overset{1}{\rightleftarrows}} \tilde{B} = 0$ where a map $\tilde{Z}_1 \overset{d}{\longrightarrow} \tilde{Z}_2$ is the map induced by multiplication by d on Z_1. (The maps are well-defined since $A_{\mathrm{sep}} \subset C \subset B$ and $p^r C \subset A_{\mathrm{sep}}$ and $tC \subset B$.)

Since $\tilde{B} = 0$, taking the composition $\tilde{C} \overset{1}{\longrightarrow} \tilde{B} \overset{t}{\longrightarrow} \tilde{C}$ shows that $t\tilde{C} = 0$. If $x \in \tilde{C}$ and $p^r x = 0$, then since $\gcd(p^r, t) = 1$ we have $x = 0$. Thus, the composition $\tilde{C} \overset{p^r}{\longrightarrow} \tilde{A}_{\mathrm{sep}} \overset{1}{\longrightarrow} \tilde{C}$ is an injection, and thus an automorphism α of the finite abelian group $\tilde{C}$. It follows that $\tilde{A}_{\mathrm{sep}} \overset{1}{\longrightarrow} \tilde{C}$ is surjective and $\tilde{C} \overset{p^r}{\longrightarrow} \tilde{A}_{\mathrm{sep}}$ is injective. Further, letting $\tilde{A}_{\mathrm{sep}}[p^r]$ denote the kernel of multiplication by p^r in $\tilde{A}_{\mathrm{sep}}$, we have

$$\ker(\,\tilde{A}_{\mathrm{sep}} \overset{1}{\longrightarrow} \tilde{C}\,) = \ker(\,\tilde{A}_{\mathrm{sep}} \overset{1}{\longrightarrow} \tilde{C} \overset{p^r}{\longrightarrow} \tilde{A}_{\mathrm{sep}}\,) = \tilde{A}_{\mathrm{sep}}\,[p^r]\,.$$

This gives a split short exact sequence

$$0 \longrightarrow \tilde{A}_{\mathrm{sep}}\,[p^r] \longrightarrow \tilde{A}_{\mathrm{sep}} \underset{p^r \alpha^{-1}}{\overset{1}{\rightleftarrows}} \tilde{C} \longrightarrow 0$$

with $\tilde{C}$ killed by t. Thus $\tilde{C}$ is the non-p-component of $\tilde{A}_{\mathrm{sep}}$, proving (i).

We have $n(A_{\mathrm{sep}}, \mathfrak{m}, \mathfrak{n}) \notin p^{\mathbb{Z}_{\geq 0}}$ if and only if $\tilde{A}_{\mathrm{sep}}$ is not a p-group, i.e., if and only if $\tilde{C} \neq 0$ (by (i)). But $\tilde{C} \neq 0$ if and only if $\mathfrak{m}$ and $\mathfrak{n}$ are connected in $\Gamma(C)$. This gives (ii). $\square$

One could compute $\Gamma(C)$ by applying Algorithm 5.6 with $D = C$. Thanks to Proposition 10.1 we can compute $\Gamma(C)$ without actually computing C, as follows.

Algorithm 10.2. The algorithm takes an order A and the numbers $n(A_{\mathrm{sep}}, \mathfrak{m}, \mathfrak{n})$, and computes the graph $\Gamma(C)$ and its connected components.

(i) Connect two vertices $\mathfrak{m}$ and $\mathfrak{n}$ if and only if $n(A_{\mathrm{sep}}, \mathfrak{m}, \mathfrak{n}) \notin p^{\mathbb{Z}_{\geq 0}}$.
(ii) Output the associated graph and the connected components.

Definition 10.3. If $W \subset \mathrm{Spec}(E)$, let C_W denote the image of C in the quotient

$$\prod_{\mathfrak{m} \in W} A_{\mathrm{sep}}/(\mathfrak{m} \cap A_{\mathrm{sep}})$$

of B.

Lemma 10.4. *Let Ω denote the set of connected components of the graph $\Gamma(C)$. Then the natural map $F : C \to \prod_{W \in \Omega} C_W$ is an isomorphism.*

Proof. The map F is injective, since

$$C \subset B = \prod_{W \in \Omega} \prod_{\mathfrak{m} \in W} A_{\mathrm{sep}}/(\mathfrak{m} \cap A_{\mathrm{sep}}).$$

If $f_W : C \twoheadrightarrow C_W$ is the natural map, e_W is as defined in Notation 5.5 with $D = C$, and $x = (f_W(c_W))_{W \in \Omega}$ is an arbitrary element of $\prod_{W \in \Omega} C_W$, then $F(\sum_{W \in \Omega} c_W e_W) = x$, so F is surjective. The result now follows from Proposition 5.7(ii). $\qquad\square$

Proposition 10.5. *Suppose A is an order and p is a prime number. Recall C as defined in (10.1). Fix a subset $W \subset \mathrm{Spec}(E)$ for which the induced subgraph of $\Gamma(C)$ is connected. Then:*
   (i) *the ring C_W is connected,*
   (ii) *the natural map $\mu(C_W)_p \to \mu(C_{\{\mathfrak{m}\}})_p$ is injective for all $\mathfrak{m} \in W$,*
   (iii) *the group $\mu(C_W)_p$ is cyclic,*
   (iv) *if W' is a non-empty subset of W, then the natural map $\mu(C_W)_p \to \mu(C_{W'})_p$ is injective.*

Proof. Part (i) follows from Lemma 5.1.
   Let $B_W = \prod_{\mathfrak{m} \in W} A_{\mathrm{sep}}/(\mathfrak{m} \cap A_{\mathrm{sep}})$. We have

$$\mathrm{id}(C_W\,[1/p]) \subset \mathrm{id}\left(\prod_{\mathfrak{m} \in W} E/\mathfrak{m}\right) = \mathrm{id}(B_W).$$

Recall B from (9.1). Since $(B : C)$ is coprime to p, so is $(B_W : C_W)$. Suppose $e \in \mathrm{id}(C_W\,[1/p])$. Then $e \in \mathrm{id}(B_W)$ and there exists $m \in \mathbb{Z} - p\mathbb{Z}$ such that $me \in C_W$ (e.g., $m = (B_W : C_W)$). Further, there exists $k \in \mathbb{Z}_{\geq 0}$ such that $p^k e \in C_W$. Since m and p^k are coprime, we have $e \in C_W$. Thus, $\mathrm{id}(C_W\,[1/p]) = \mathrm{id}(\bar{C}_W) = \{0, 1\}$, so $C_W\,[1/p]$ is connected. Now by Corollary 8.2 with $R = C_W\,[1/p]$ and $m = \#\mu(C_W\,[1/p])_p$, the group $\mu(C_W\,[1/p])_p$ is cyclic, so its subgroup $\mu(C_W)_p$ is cyclic as well, which is (iii). Also, by Proposition 8.1(b) with $R = C_W\,[1/p]$ and $f = X^m - 1$, the map $\mu(C_W\,[1/p])_p \to \mu(C_{W'}\,[1/p])_p$ is injective for each non-empty $W' \subset W$. This implies (iv). With $W' = \{\mathfrak{m}\}$ one obtains (ii). $\qquad\square$

Remark 10.6. If A is a connected order in a separable $\mathbb{Q}$-algebra and p is a prime number that does not divide $\#(B/A)$, then $\mu(A)_p$ is cyclic. This follows from Proposition 10.5(iii); $C = A$ since $E = E_{\mathrm{sep}}$ and $p \nmid \#(B/A)$, and one can take $C = C_W$ since A is connected.

By Proposition 10.5(ii,iii), if W is a connected component of $\Gamma(C)$, then the natural map

$$\mu(C_W)_p \to \mu(A/(\mathfrak{m} \cap A))_p$$

is injective for all $\mathfrak{m} \in W$, and $\mu(C_W)_p$ is cyclic. This gives an efficient algorithm for computing $\mu(C_W)_p$, and thus a set of generators for $\mu(C)_p$, as follows.

Algorithm 10.7. Given an order A and a prime p, the algorithm finds an efficient presentation for $\mu(C)_p$.

(i) Apply Algorithm 9.2 to compute a generator of the cyclic group $\mu(A_{\mathrm{sep}}/(\mathfrak{m} \cap A_{\mathrm{sep}}))_p$ for each $\mathfrak{m} \in \mathrm{Spec}(E)$.

(ii) Apply Algorithm 10.2 to compute $\Gamma(C)$ and its connected components W.

(iii) For each W, do the following:

    (a) Apply the image algorithm in §14 of [5] to compute a basis for the order

$$C_W = \mathrm{image}(C \to \prod_{\mathfrak{m} \in W} E/\mathfrak{m}).$$

    (b) Pick $\mathfrak{m}_1 \in W$ with $\#\mu(A_{\mathrm{sep}}/(\mathfrak{m}_1 \cap A_{\mathrm{sep}}))_p$ minimal.

    (c) Choose

$$W_1 = \{\mathfrak{m}_1\} \subset W_2 = \{\mathfrak{m}_1, \mathfrak{m}_2\} \subset \ldots \subset W$$

such that $\#W_i = i$ for all $i \geq 1$, and $W_i = W_{i-1} \cup \{\mathfrak{m}_i\}$ for all $i \geq 2$, and each $\mathfrak{m}_i$ is connected in $\Gamma(C)$ to some $\mathfrak{m}_j$ with $j < i$.

    (d) For $i = 1, 2, \ldots$ compute each $\mu(C_{W_i})_p$, and a generator for it, in succession by using that $\mu(C_{W_1})_p = \mu(A_{\mathrm{sep}}/(\mathfrak{m}_1 \cap A_{\mathrm{sep}}))_p$ is given, and for $i > 1$ listing all ordered pairs in $\mu(C_{W_{i-1}})_p \times \mu(A_{\mathrm{sep}}/(\mathfrak{m}_i \cap A_{\mathrm{sep}}))_p$ and testing whether they are in C_{W_i}, and using that

$$\mu(C_{W_i})_p = C_{W_i} \cap (\mu(C_{W_{i-1}})_p \times \mu(A_{\mathrm{sep}}/(\mathfrak{m}_i \cap A_{\mathrm{sep}}))_p).$$

This gives a generator of $\mu(C_W)_p$ for each W in the set Ω of connected components of $\Gamma(C)$. Let $\zeta_W \in \prod_{V \in \Omega} \mu(C_V)_p$ be the element with this generator as its W-th component, and all other components 1.

(iv) View the set $S = \{\zeta_W : W \in \Omega\}$ in $\mu(C)_p$ via the isomorphism $\mu(C)_p \cong \prod_{W \in \Omega} \mu(C_W)_p$ of Lemma 10.4, let $R = \{\mathrm{order}(\zeta_W)(W\text{-th basis vector})\}$, and output $\langle S|R \rangle$.

Proposition 10.8. *Algorithm 10.7 gives correct output and runs in polynomial time.*

Proof. By Lemma 10.4 we have $C \xrightarrow{\sim} \prod_W C_W$. Thus, $\mu(C)_p \xrightarrow{\sim} \bigoplus_W \mu(C_W)_p$ so the output of the algorithm is a set of generators for $\mu(C)_p$. We have

$$C_{W_i} \subset C_{W_{i-1}} \times C_{\{\mathfrak{m}_i\}}, \qquad C_{\{\mathfrak{m}_i\}} = A_{\mathrm{sep}}/(\mathfrak{m}_i \cap A_{\mathrm{sep}}).$$

Thus,

$$\mu(C_{W_i})_p \subset \mu(C_{W_{i-1}})_p \times \mu(A_{\mathrm{sep}}/(\mathfrak{m}_i \cap A_{\mathrm{sep}}))_p.$$

By Proposition 10.5, the group $\mu(C_{W_i})_p$ injects into each factor, and each factor is cyclic of prime power order. Each factor has size polynomial in the size of the algorithm's inputs (given an order of rank n and an element of order p^k, we have $\varphi(p^k) \leq n$ by Lemma 3.3, so $p^k \leq 2n$). By Proposition 10.5(ii) the natural map $\mu(C_{W_i})_p \to \mu(A_{\mathrm{sep}}/(\mathfrak{m}_1 \cap A_{\mathrm{sep}}))_p$ is injective, for all i. As i gets larger, the groups $\mu(C_{W_i})_p$ get smaller or stay the same. Thus one can list all ordered pairs, and then efficiently test whether they are in C_{W_i}. It follows from the above that the algorithm runs in polynomial time.

The presentation $\langle S|R \rangle$ is efficient by Algorithm 7.6 and Proposition 9.3, since $\mu(C)_p \subset \mu(B)_p$. $\square$

Remark 10.9. A more intelligent algorithm for step (iii)(d) is to use that each $\mu(C_{W_i})_p$ is cyclic (by Proposition 10.5(iii)), and that $\mu(C_{W_i})_p \subset \mu(C_{W_{i-1}})_p$, as follows. Starting with $i = 1$ and incrementing i, proceed as follows in place of step (d). If $\mu(C_{W_{i-1}})_p$ is trivial, stop. Otherwise, take an element $a_1 \in \mu(C_{W_{i-1}})_p$ of order p and for each of the $p - 1$ elements $b_1 \in \mu(A_{\mathrm{sep}}/(\mathfrak{m}_i \cap A_{\mathrm{sep}}))_p$ of order p test whether $(a_1, b_1) \in C_{W_i}$. If there are none, stop (the group is trivial for that W_i). If there is such a pair $(a_1, b_1) \in \mu(C_{W_i})$, if $\#\mu(C_{W_i})_p = p$ then stop with (a_1, b_1) as generator, and otherwise take each $a_2 \in \mu(C_{W_{i-1}})_p$ that is a p-th root of a_1 and for each of the p possible choices of elements $b_2 \in \mu(A_{\mathrm{sep}}/(\mathfrak{m}_i \cap A_{\mathrm{sep}}))_p$ that are a p-th root of b_1, test whether $(a_2, b_2) \in C_{W_i}$. As soon as such is found, if $\#\mu(C_{W_i})_p = p^2$ then stop with (a_2, b_2) as generator, and otherwise continue this process. Injecting into each component implies one only needs to check ordered pairs with the same

order in each component. Since $\#\mu(C_{W_i})_p$ divides $\#\mu(C_{W_{i-1}})_p$, one only needs to go up to elements of order $\#\mu(C_{W_{i-1}})_p$. The number of trials is $< p\log_p(\#\mu(C_{W_{i-1}})_p)$, since there are p choices each time, and there are $\log_p(\#\mu(C_{W_{i-1}})_p)$ steps. The final (a_j, b_j) found is a generator for $\mu(C_{W_i})_p$.

11. NILPOTENT IDEALS IN FINITE RINGS

Suppose R is a finite commutative ring and I is a nilpotent ideal of R. Algorithm 11.3 below solves the discrete logarithm problem in the multiplicative group $1 + I$, using the finite filtration:

$$1 + I \supset 1 + I^2 \supset 1 + I^4 \supset \cdots \supset 1,$$

the fact that the map $x \mapsto 1 + x$ is an isomorphism from the additive group $I^{2^i}/I^{2^{i+1}}$ to the multiplicative group $(1 + I^{2^i})/(1 + I^{2^{i+1}})$, and the fact that the discrete logarithm problem is easy in these additive groups.

We specify a finite commutative ring by giving a presentation for its additive group, i.e., a finite set of generators and a finite set of relations, and for every pair of generators their product is expressed as a $\mathbb{Z}$-linear combination of the generators.

The following result can be shown using standard methods.

Proposition 11.1. *There is a deterministic polynomial-time algorithm that, given a finite commutative ring R and 2 ideals I_1 and I_2 of R such that $I_2 \subset I_1$, computes an efficient presentation of the finite abelian group I_1/I_2.*

Lemma 11.2. *Suppose R is a finite commutative ring, I is an ideal of R such that $I \subset \sqrt{0_R}$, and for each $i \in \mathbb{Z}_{\geq 0}$ the set B_i is a subset of I^{2^i} such that $B_i \cup I^{2^{i+1}}$ generates the additive group I^{2^i}. Let $\mathcal{B} = \bigcup_{i \geq 0} B_i$. Then $1 + I = \langle 1 + b : b \in \mathcal{B} \rangle$ (as a multiplicative group).*

Proof. Since I is nilpotent, $1 + I^{2^i}$ is a multiplicative group for all $i \in \mathbb{Z}_{\geq 0}$. We have

$$I^{2^i}/I^{2^{i+1}} \xrightarrow{\sim} (1 + I^{2^i})/(1 + I^{2^{i+1}})$$

via $x \mapsto 1 + x$. Since $B_i \cup I^{2^{i+1}}$ generates the additive group I^{2^i}, we have that $B_i + I^{2^{i+1}}$ generates $I^{2^i}/I^{2^{i+1}}$. If $I^{2^{k+1}} = 0$, then B_k generates I^{2^k} and $1 + B_k$ generates the multiplicative group $1 + I^{2^k}$. It now follows that $1 + \mathcal{B}$ generates $1 + I$. $\square$

Algorithm 11.3. Given a finite commutative ring R, an ideal I of R such that $I \subset \sqrt{0}$, for each $i \in \mathbb{Z}_{\geq 0}$ a subset B_i of I^{2^i} such that $B_i \cup I^{2^{i+1}}$ generates the additive group I^{2^i}, with all but finitely many $B_i = \emptyset$, and $x \in I$, the algorithm computes $(m_b)_{b \in \mathcal{B}} \in \mathbb{Z}^{\mathcal{B}}$ with $1 + x = \prod_{b \in \mathcal{B}}(1 + b)^{m_b}$, where $\mathcal{B} = \bigcup_{i \geq 0} B_i$, as follows.

(i) Let $x_0 = x$. For $i = 0, 1, \ldots$ use Proposition 11.1 to find $(m_b)_{b \in B_i} \in \mathbb{Z}^{B_i}$ such that

$$x_i \equiv \sum_{b \in B_i} m_b b \mod I^{2^{i+1}} \ (\text{in } I^{2^i}/I^{2^{i+1}}).$$

Define $x_{i+1} \in I^{2^{i+1}}$ by

$$1 + x_{i+1} = (1 + x_i) \prod_{b \in B_i}(1 + b)^{-m_b}.$$

As soon as $x_{i+1} = 0$, terminate, setting $m_b = 0$ for all $b \in B_j$ with $j > i$ and outputting $(m_b)_{b \in \mathcal{B}} \in \mathbb{Z}^{\mathcal{B}}$.

Proposition 11.4. *Algorithm 11.3 is a deterministic algorithm that produces correct outputs in polynomial time.*

Proof. Since I is a nilpotent ideal, there exists $j \in \mathbb{Z}_{\geq 0}$ such that $I^{2^j} = 0$. Then $x_j = 0$ and the algorithm gives

$$1 + x = 1 + x_0 = \prod_{b \in \bigcup_{i < j} B_i} (1 + b)^{m_b} = \prod_{b \in \mathcal{B}} (1 + b)^{m_b}$$

as desired. $\square$

Lemma 11.5. *There is a deterministic polynomial-time algorithm that, given a finite commutative ring R, an ideal I of R such that $I \subset \sqrt{0}$, and for each $i \in \mathbb{Z}_{\geq 0}$ a subset B_i of I^{2^i} such that $B_i \cup I^{2^{i+1}}$ generates the additive group I^{2^i}, computes a $\mathbb{Z}$-basis for the kernel of the map $\mathbb{Z}^{\mathcal{B}} \to 1 + I$, $(m_b)_{b \in \mathcal{B}} \mapsto \prod_b (1 + b)^{m_b}$, where $\mathcal{B} = \bigcup_{i \geq 0} B_i$.*

Proof. Let $C_j = \bigcup_{k \geq j} B_j$. We proceed by induction on decreasing j. We have $\langle 1 + C_j \rangle = 1 + I^{2^j}$ (applying Lemma 11.2 with I^{2^j} in place of I). Assume we already have defining relations for $1 + C_j$, i.e., we have generators for the kernel of $\mathbb{Z}^{C_j} \to 1 + I^{2^j}$, $(m_b)_{b \in C_j} \mapsto \prod_{b \in C_j} (1 + b)^{m_b}$, and would like to find defining relations for $1 + C_{j-1}$. Proposition 11.1 gives an algorithm for finding a basis for the kernel of $\mathbb{Z}^{B_{j-1}} \to I^{2^{j-1}}/I^{2^j}$, $(n_b)_{b \in B_{j-1}} \mapsto \prod_{b \in B_{j-1}} n_b b + I^{2^j}$ in polynomial time. For each defining relation $(n_b)_{b \in B_{j-1}}$ for $B_{j-1} + I^{2^j}$ we have $\sum_{b \in B_{j-1}} n_b b \equiv 0 \bmod I^{2^j}$ so $\prod_{b \in B_{j-1}} (1 + b)^{n_b} \equiv 1 \bmod (1 + I^{2^j})$. Algorithm 11.3 gives a polynomial-time algorithm to find $(m_{b'})_{b' \in C_j} \in \mathbb{Z}^{C_j}$ such that $\prod_{b \in B_{j-1}} (1 + b)^{n_b} = \prod_{b' \in C_j} (1 + b')^{m_{b'}} \in 1 + I^{2^j}$. Then $((n_b)_{b \in B_{j-1}}, (-m_{b'})_{b' \in C_j})$ is in the kernel of the map $\mathbb{Z}^{C_{j-1}} \to 1 + I^{2^{j-1}}$, and these relations along with the defining relations for $1 + C_j$ form a set of defining relations for $1 + C_{j-1}$. $\square$

Theorem 11.6. *There is a deterministic polynomial-time algorithm that, given a finite commutative ring and an ideal I of R such that $I \subset \sqrt{0}$, produces an efficient presentation $\langle 1 + \mathcal{B} | \mathcal{R} \rangle$ for $1 + I$.*

Proof. Apply the algorithm in Proposition 11.1 to obtain for each $i \in \mathbb{Z}_{\geq 0}$ a set $B_i \subset I^{2^i}$ such that $B_i \cup I^{2^{i+1}}$ generates the additive group I^{2^i}. Since I is nilpotent, we can take $B_i = \emptyset$ for all but finitely many i. By Lemma 11.2 the set $\mathcal{B} = \bigcup_{i \geq 0} B_i$ has the property that $1 + \mathcal{B}$ generates $1 + I$. Defining relations $\mathcal{R}$ are given by Lemma 11.5, and part (c) of Definition 7.1 holds by Proposition 11.4. $\square$

Theorem 1.4 now follows from Theorem 11.6 and Algorithm 7.6.

Remark 11.7. Suppose R is a finite commutative ring, $I \subset R$ is a nilpotent ideal, and R' is a subring of R. Let $I' = I \cap R'$. The algorithm in Theorem 11.6 gives efficient presentations for the multiplicative groups $1 + I$ and $1 + I'$. We can apply Algorithm 7.8 with $G = 1 + I \subset R^*$, and T' a set of generators for $1 + I'$, and T a set of generators for some subgroup of $1 + I$. In the next section we will apply this to our setting.

Example 11.8. Let $R = \mathbb{Z}/p^2\mathbb{Z}$ and $I = \sqrt{0_R} = p\mathbb{Z}/p^2\mathbb{Z}$. Then $I^2 = 0$, and $1 + I$ is the order p subgroup of $(\mathbb{Z}/p^2\mathbb{Z})^* \cong \mathbb{Z}/p\mathbb{Z} \times \mathbb{Z}/(p - 1)\mathbb{Z}$. The map $1 + I \xrightarrow{\sim} \mathbb{Z}/p\mathbb{Z}$, $1 + x \mapsto x/p$ is a group isomorphism, so the discrete logarithm problem is easy in $1 + I$.

Example 11.9. Let $R = \mathbb{Z}/p^4\mathbb{Z}$ and $I = \sqrt{0_R} = p\mathbb{Z}/p^4\mathbb{Z}$. Then $I^4 = 0$. Here, the map $1 + I \xrightarrow{\sim} \mathbb{Z}/p^3\mathbb{Z}$, $1 + x \mapsto x/p$ is not a group homomorphism. The discrete logarithm problem is easy in $1 + I$ not because it is (isomorphic to) an additive group, but because there is a filtration of additive groups, namely, $(1 + I)/(1 + I^2) \cong I/I^2$ and $(1 + I^2)/(1 + I^4) \cong I^2/I^4 = I^2$.

12. FROM $\mu(C)_p$ TO $\mu(A)_p$

Let A be an order and let p be a prime. Recall C from Definition 10.1 and let

$$\mathfrak{f} = \{x \in C : xC \subset A_{\mathrm{sep}}\},$$

which is the largest ideal of C that is contained in A. We shall see that $C/\mathfrak{f}$ is a finite ring, and it has $A_{\mathrm{sep}}/\mathfrak{f}$ as a subring. Suppose we are given a set $M \subset C^*$ such that $\mu(C)_p = \langle M \rangle$. Let

$$I = \sum_{\zeta \in M} (\zeta - 1)(C/\mathfrak{f}), \qquad I' = I \cap (A_{\mathrm{sep}}/\mathfrak{f}).$$

Define

$$g_1 : \mathbb{Z}^M \twoheadrightarrow \mu(C)_p, \qquad (a_\zeta)_{\zeta \in M} \mapsto \prod_{\zeta \in M} \zeta^{a_\zeta},$$

let $g_2 : \mu(C)_p \to 1 + I$ be the natural map $\zeta \mapsto \zeta + \mathfrak{f}$, let $\hat{g} : \mu(C)_p \to (1 + I)/(1 + I')$ denote the composition of g_2 with the quotient map, define $g : \mathbb{Z}^M \to 1 + I$ by $g = g_2 \circ g_1$, and define

(12.1) $$\psi : \mathbb{Z}^M \to (1 + I)/(1 + I') \quad \text{by} \quad \psi = \hat{g} \circ g_1.$$

Proposition 12.1. *With notation as above,*

    (i) *I is a nilpotent ideal of $C/\mathfrak{f}$, i.e., $I \subset \sqrt{0_{C/\mathfrak{f}}}$;*

    (ii) *I' is a nilpotent ideal of $A_{\mathrm{sep}}/\mathfrak{f}$;*

    (iii) *$C/\mathfrak{f}$ is a finite ring of p-power order,*

    (iv) *$\mu(A)_p$ is the kernel of the map $\hat{g}$;*

    (v) *$\mu(A)_p$ is the image of $\ker(\psi)$ under the map g_1.*

Proof. Since C/A is killed by p^r for some $r \in \mathbb{Z}_{\geq 0}$, we have $p^r \in \mathfrak{f}$, so $p \in \sqrt{0_{C/\mathfrak{f}}}$, so p is in every prime ideal of $C/\mathfrak{f}$. Suppose $\zeta \in \mu(C)_p$. Then the image of ζ in every field of characteristic p is 1. Thus, $\zeta - 1$ is in every prime ideal of $C/\mathfrak{f}$, so $\zeta - 1 \in \sqrt{0_{C/\mathfrak{f}}}$. By the definition of I we have $I \subset \sqrt{0_{C/\mathfrak{f}}}$, and (i) and (ii) follow.

Since $p^r \in \mathfrak{f}$ we have $p^r C \subset \mathfrak{f}$, so $C/\mathfrak{f}$ is a quotient of $C/p^r C$, which is a finite ring of p-power order. This gives (iii).

Part (iv) follows directly from the definitions, and then (v) follows from (iv). $\qquad \square$

Algorithm 12.2. The algorithm takes as input an order A, a prime p, and a finite set of generators M for $\mu(C)_p$, and computes a finite set of generators for $\mu(A)_p$.

    (i) Compute the finite abelian group C/A_{sep} and

$$\mathrm{Hom}(C, C/A_{\mathrm{sep}}) \cong (C/A_{\mathrm{sep}}) \oplus (C/A_{\mathrm{sep}}) \oplus \cdots \oplus (C/A_{\mathrm{sep}})$$

        (with $\mathrm{rank}_{\mathbb{Z}}(C)$ summands C/A_{sep}), and compute $\mathfrak{f}$ as the kernel of the group homomorphism $A_{\mathrm{sep}} \to \mathrm{Hom}(C, C/A_{\mathrm{sep}})$ sending $x \in A_{\mathrm{sep}}$ to the map $y \mapsto xy + A_{\mathrm{sep}}$. Next compute the finite rings $A_{\mathrm{sep}}/\mathfrak{f} \subset C/\mathfrak{f}$. This entire step can be done using standard algorithms for finitely generated abelian groups.

    (ii) Apply the algorithm in Theorem 11.6 with $R = C/\mathfrak{f}$ and the I of this section to obtain an efficient presentation for $1 + I$.

    (iii) Apply the algorithm in Theorem 11.6 with $R = A_{\mathrm{sep}}/\mathfrak{f}$ and I' in place of I to obtain a finite set T' of generators for $1 + I'$.

    (iv) Apply Algorithm 7.8 with $G = 1 + I$, the efficient presentation from step (ii), $T = M$, and T' from step (iii) to obtain a finite set of generators S' for $\ker(\mathbb{Z}^T \to G/\langle T' \rangle)$.

    (v) Take the image of S' under the map $g_1 : \mathbb{Z}^M \to \mu(C)_p$.

Theorem 12.3. *Algorithm 12.2 produces correct output and runs in polynomial time.*

Proof. Since $C/\mathfrak{f}$ and $A_{\mathrm{sep}}/\mathfrak{f}$ are finite commutative rings, and I and I' are nilpotent, Theorem 11.6 is applicable in steps (ii) and (iii). The map $\mathbb{Z}^M = \mathbb{Z}^T \to G/\langle T' \rangle = (1 + I)/(1 + I')$ in step (iv) is our map ψ from (12.1). By Proposition 12.1(v), step (v) produces generators for $\mu(A)_p$. $\qquad\square$

13. Finding roots of unity

Algorithm 13.1. Given an order A, the algorithm outputs a finite set of generators for $\mu(A)$.

(i) Use Algorithm 3.2 to compute E_{sep}, all $\mathfrak{m} \in \mathrm{Spec}(E)$, the fields $E/\mathfrak{m}$, and the natural maps $E \to E/\mathfrak{m}$.

(ii) Apply Algorithm 4.2 to compute $A_{\mathrm{sep}} = A \cap E_{\mathrm{sep}}$.

(iii) Apply Algorithm 9.1 to compute for each $\mathfrak{m} \in \mathrm{Spec}(E)$ the subring $A_{\mathrm{sep}}/(\mathfrak{m} \cap A_{\mathrm{sep}})$ of $E_{\mathrm{sep}}/\mathfrak{m}$.

(iv) Apply the algorithm in Proposition 9.2 to compute, for each $\mathfrak{m} \in \mathrm{Spec}(E)$, a generator $\theta_{\mathfrak{m}}$ for $\mu(A_{\mathrm{sep}}/(\mathfrak{m} \cap A_{\mathrm{sep}}))$, its order, the prime factorization of its order, and for each prime p dividing its order a generator $\theta_{\mathfrak{m},p}$ of $\mu(A_{\mathrm{sep}}/(\mathfrak{m} \cap A_{\mathrm{sep}}))_p$.

(v) For each prime p dividing the order of at least one of the groups $\mu(A_{\mathrm{sep}}/(\mathfrak{m} \cap A_{\mathrm{sep}}))$, do the following:

(a) Use the image algorithm in §14 of [5] to compute a $\mathbb{Z}$-basis for $C = A_{\mathrm{sep}}[1/p] \cap B$ (as discussed in §10 above, just before Proposition 10.1).

(b) Apply Algorithm 10.7 to compute an efficient presentation for $\mu(C)_p$.

(c) Apply Algorithm 12.2 to compute generators for $\mu(A)_p$.

(vi) Generators for these groups $\mu(A)_p$ form a set of generators for $\mu(A)$.

That Algorithm 13.1 produces correct output and runs in polynomial time follows immediately. We can now obtain a deterministic polynomial-time algorithm that, given an order A, determines an efficient presentation for $\mu(A)$.

Algorithm 13.2. The algorithm takes an order A and produces an efficient presentation for $\mu(A)$.

(i) Apply the algorithm in Proposition 9.3 to obtain an efficient presentation $\langle S|R \rangle$ for $\mu(B)$.

(ii) Apply Algorithm 13.1 to obtain a finite set of generators for $\mu(A)$.

(iii) Apply Algorithm 7.6 with $G = \mu(B)$ to obtain an efficient presentation for $\mu(A)$.

Example 13.3. Let $A = \mathbb{Z}[X]/(X^4 - 1)$. Then with $p = 2$:
$$B = C = \mathbb{Z}[X]/(X - 1) \times \mathbb{Z}[X]/(X + 1) \times \mathbb{Z}[X]/(X^2 + 1) \cong \mathbb{Z} \times \mathbb{Z} \times \mathbb{Z}[\mathrm{i}],$$
and $(C : A) = 8$. We identify X with $(1, -1, \mathrm{i}) \in \mathbb{Z} \times \mathbb{Z} \times \mathbb{Z}[\mathrm{i}]$. Then
$$\mu(A)_2 = \mu(A) \subset \mu(B) = \mu(C)_2 = \langle (-1, 1, 1), (1, -1, 1), (1, 1, \mathrm{i}) \rangle.$$
We have
$$\mathfrak{f} = 4\mathbb{Z} \times 4\mathbb{Z} \times 2\mathbb{Z}[\mathrm{i}]$$
of index 64 in C, and
$$C/\mathfrak{f} = \mathbb{Z}/4\mathbb{Z} \times \mathbb{Z}/4\mathbb{Z} \times \mathbb{Z}[\mathrm{i}]/2\mathbb{Z}[\mathrm{i}] = \mathbb{Z}/4\mathbb{Z} \times \mathbb{Z}/4\mathbb{Z} \times \mathbb{F}_2[\varepsilon]$$
with $\varepsilon = 1 + \mathrm{i}$. The index 8 subring of $C/\mathfrak{f}$ generated by $(1, -1, 1 + \varepsilon)$ is $A/\mathfrak{f}$. Alternatively,
$$A/\mathfrak{f} = (\mathbb{Z}/4\mathbb{Z})[Y]/(2Y, Y^2)$$
where $Y = X - 1 = (0, 2, \varepsilon) \in A/\mathfrak{f}$. With $M = \{(-1, 1, 1), (1, -1, 1), (1, 1, \mathrm{i})\}$ we have
$$I = (2\mathbb{Z}/4\mathbb{Z}) \times (2\mathbb{Z}/4\mathbb{Z}) \times (\varepsilon \mathbb{F}_2[\varepsilon]) = \sqrt{0_{C/\mathfrak{f}}},$$
$I^2 = 0$, and
$$I' = I \cap (A/\mathfrak{f}) = \sqrt{0_{A/\mathfrak{f}}} = \{0, 2, Y, Y + 2\}.$$

With ψ as in (12.1), we have $\psi(a, b, c) = a + b + c + 2\mathbb{Z} \in \mathbb{Z}/2\mathbb{Z}$ and

$$\ker(\psi) = \{(a, b, c) \in \mathbb{Z}^M : a + b + c \text{ is even}\} = \mathbb{Z} \cdot (2, 0, 0) + \mathbb{Z} \cdot (1, 1, 0) + \mathbb{Z} \cdot (1, 0, 1).$$

Algorithm 13.1 outputs

$$\mu(A) = \mu(A)_2 = \langle -X^2 \rangle \times \langle -X^3 \rangle = \langle X, -1 \rangle \cong \mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/4\mathbb{Z}.$$

Example 13.4. Let $A = \mathbb{Z}[X]/(X^{12} - 1)$. Then

$$E = \mathbb{Q}[X]/(X^{12} - 1) \cong \mathbb{Q} \times \mathbb{Q} \times \mathbb{Q}(\zeta_3) \times \mathbb{Q}(i) \times \mathbb{Q}(\zeta_3) \times \mathbb{Q}(\zeta_{12})$$

and

$$\begin{aligned} B = \ & \mathbb{Z}[X]/(X - 1) \ \times \ \mathbb{Z}[X]/(X + 1) \ \times \ \mathbb{Z}[X]/(X^2 + X + 1) \\ & \times \ \mathbb{Z}[X]/(X^2 + 1) \ \times \ \mathbb{Z}[X]/(X^2 - X + 1) \ \times \ \mathbb{Z}[X]/(X^4 - X^2 + 1) \hookrightarrow E. \end{aligned}$$
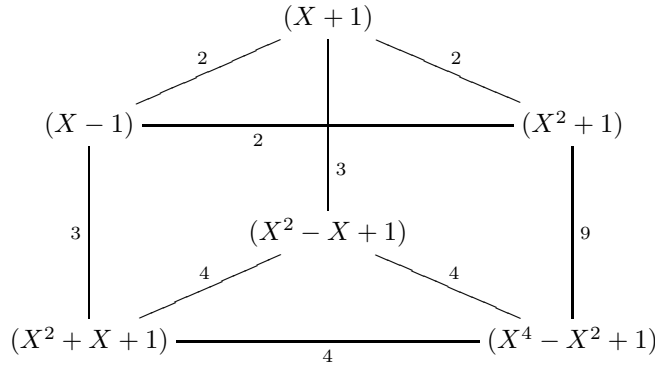
We have for the discriminants of the orders:

$$|\Delta_B| = 1 \cdot 1 \cdot 3 \cdot 4 \cdot 3 \cdot 12^2, \qquad |\Delta_A| = 12^{12},$$
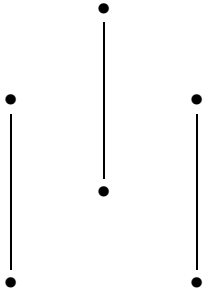
so

$$\#(B/A) = \sqrt{|\Delta_A|/|\Delta_B|} = 2^9 \cdot 3^4.$$

Thus if $p = 2$ then $(C : A) = 2^9$, while if $p = 3$ then $(C : A) = 3^4$. The graph $\Gamma(B)$ consists of 6 vertices with no edges. With the numbers $n(A, \mathfrak{m}, \mathfrak{n})$ on the edges, the graph $\Gamma(A)$ is:
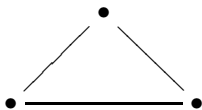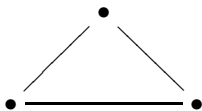

Suppose $p = 2$. Then the graph $\Gamma(C)$ is:

We have $\mu(C)_2 = \prod \mu(C_W)_2$ with the product running over the 3 connected components W. The left 2 W's give $\mu(C_W)_2 = \{\pm 1\}$, while the remaining one gives $\mu(C_W)_2 = \langle -X^3 \rangle$. This gives $-X^3, -1 \in \mu(A)_2$.

Suppose $p = 3$. Then the graph $\Gamma(C)$ is:

We have $\mu(C)_3 = \prod \mu(C_W)_3$ with the product running over the 2 connected components W. The top W has $\mu(C_W)_3 = \{1\}$, while for the bottom W one has that $\mu(C_W)_3$ is generated by the image of X^4, and this gives $X^4 \in \mu(A)_3$.

Continuing the algorithm by hand is more complicated than in the previous example. However, we note that here A is the order $\mathbb{Z}\langle G \rangle$ defined in [7] with $G = \langle -1 \rangle \times \langle X \rangle \cong \mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/12\mathbb{Z}$, and it follows from Remark 16.3 of [7] that $\mu(A) = G = \langle -1 \rangle \times \langle X \rangle$.

References

[1] M. F. Atiyah and I. G. Macdonald, Introduction to commutative algebra, Addison-Wesley Publishing Co., Reading, MA, 1969.

[2] J. Hopcroft and R. Tarjan, *Algorithm 447: efficient algorithms for graph manipulation*, Communications of the ACM, **16**, no. 6 (1973) 372–378.

[3] S. Lang, Algebra, Third edition, Graduate Texts in Mathematics **211**, Springer-Verlag, New York, 2002.

[4] A. K. Lenstra, *Factoring polynomials over algebraic number fields*, in Computer algebra (London, 1983), Lect. Notes in Comp. Sci. **162**, Springer, Berlin, 1983, 245–254.

[5] H. W. Lenstra, Jr., *Lattices*, in Algorithmic number theory: lattices, number fields, curves and cryptography, Math. Sci. Res. Inst. Publ. **44**, Cambridge Univ. Press, Cambridge, 2008, 127–181, `http://library.msri.org/books/Book44/files/06hwl.pdf`.

[6] H. W. Lenstra, Jr. and A. Silverberg, *Revisiting the Gentry-Szydlo Algorithm*, in Advances in Cryptology—CRYPTO 2014, Lect. Notes in Comp. Sci. **8616**, Springer, Berlin, 2014, 280–296.

[7] H. W. Lenstra, Jr. and A. Silverberg, *Lattices with symmetry*, to appear in Journal of Cryptology, `https://eprint.iacr.org/2014/1026`.

[8] H. W. Lenstra, Jr. and A. Silverberg, *Algorithms for commutative algebras over the rational numbers*, `http://arxiv.org/abs/1509.08843`.

Mathematisch Instituut, Universiteit Leiden, The Netherlands
E-mail address: `hwl@math.leidenuniv.nl`

Department of Mathematics, University of California, Irvine, CA 92697
E-mail address: `asilverb@math.uci.edu`