IAS/Park City Mathematics Series
Volume 00, Pages 000–000
S 1079-5634(XX)0000-0

Polynomial-time algorithms in algebraic number theory

Daniël M. H. van Gent

Abstract.

Contents

1 Introduction 2
  1.1 Algorithms 2
  1.2 Basic computations 3
  1.3 Commutative algebra and number theory 3
  1.4 Exercises 4
2 Coprime base factorization 5
  2.1 Coprime base factorization in number fields 11
3 Finitely generated abelian groups 12
  3.1 Lattices and short bases 14
  3.2 The LLL-algorithm 17
  3.3 The kernel-image algorithm 18
  3.4 Applications of the kernel-image algorithm 19
  3.5 Homomorphism groups of finitely generated abelian groups 21
  3.6 Structure theorem for finitely generated abelian groups 23
4 Computing symbols 24
  4.1 Jacobi symbols in number rings 26
  4.2 Signs of permutations 26
  4.3 Computing signs of group automorphisms 28
  4.4 Computing Jacobi symbols in number rings 29
5 Finite commutative rings 30
  5.1 Polynomial-time reductions 32
  5.2 Trace radicals 33
  5.3 Discriminants 34
  5.4 Proof of the main theorems 36
  5.5 Exercises 37
6 Inverting ideals 38
  6.1 Blowing up 41

2010 *Mathematics Subject Classification.* Primary 14Dxx; Secondary 14Dxx.
Key words and phrases. Park City Mathematics Institute.

6.2 Inverting subgroups with two generators 45
6.3 Multiplicative inversion 47
6.4 Unit products 49
7 Tame orders 49

1. Introduction

In these notes we study methods for solving algebraic computational problems, mainly those involving number rings, which are fast in a mathematically precise sense. Our motivating problem is to decide for a number field K, elements $a_1, \ldots, a_t \in K^*$ and $n_1, \ldots, n_t \in \mathbb{Z}$ whether $\prod_i a_i^{n_i} = 1$, which turns out to be non-trivial. From basic operations such as addition and multiplication of integers we work up to computation in finitely generated abelian groups and algebras of finite type. Here we encounter lattices and Jacobi symbols. Some difficult problems in computational number theory are known, for example computing a prime factorization or a maximal order in a number field. We will however focus mainly on the positive results, namely the problems for which fast computational methods are known.

1.1. Algorithms To be able to talk about polynomial-time algorithms, we should first define what an algorithm is. We equip the set of natural numbers $\mathbb{N} = \mathbb{Z}_{\geqslant 0}$ with a length function $l : \mathbb{N} \to \mathbb{N}$ that sends n to the number of digits of n in base 2, with $l(0) = 1$.

Definition 1.1. A *problem* is a function $f : I \to \mathbb{N}$ for some set of inputs $I \subseteq \mathbb{N}$ and we call f a *decision problem* if $f(I) \subseteq \{0, 1\}$. An *algorithm* for a problem $f : I \to \mathbb{N}$ is a 'method' to compute $f(x)$ for all $x \in I$. An algorithm for f is said to run in *polynomial time* if there exist $c_0, c_1, c_1 \in \mathbb{R}_{>0}$ such that for all $x \in I$ the time required to compute $f(x)$, called the *run-time*, is at most $(c_0 + c_1 l(x))^{c_2}$. We say a problem f is *computable* if there exists an algorithm for f.

This definition is rather empty: we have not specified what a 'method' is, nor have we explained how to measure run-time. We will briefly treat this more formally. The reader for which the above definition is sufficient can freely skip the following paragraph. The main conclusion is that we will not heavily rely on the formal definition of run-time in these notes.

In these notes, the word algorithm will be synonymous with the word *Turing machine*. For an extensive treatment of Turing machines, see [4]. A Turing machine is a model of computation described by Alan Turing in 1936 that defines an abstract machine which we these days think of as a computer. The main differences between a Turing machine and a modern day computer is that the memory of a Turing machine is a tape as opposed to random-access memory, and that a Turing machine has infinite memory. The run-time of a Turing machine is then measured as the number of elementary tape operations: reading a symbol on the tape, writing a symbol on the tape and moving the tape one unit forward or backward. It is then immediately clear that it is expensive for a Turing machine to move the tape around much to look up data, as opposed to the random-access memory model where the cost of a memory lookup is constant, regardless of where the data are stored in memory. This also poses a problem for our formal treatment of run-time, as it may depend on our model of computation. However, both models of computation are able to emulate each other in such a way that it preserves the property of computability in polynomial time, even though the constants c_0, c_1 and c_2 as in Definition 1.1 may increase drastically. We use this as an excuse to be informal in these notes about determining the run-time of an algorithm.

1.2. Basic computations In these notes we build up our algorithms from basic building blocks. First and foremost, we remark that the basic operations in $\mathbb{Z}$ and $\mathbb{Q}$ are fast. Addition, subtraction, multiplication and division (with remainder in the case of $\mathbb{Z}$) can be done in polynomial time, as well as checking the sign of a number and whether numbers are equal. We assume here that we represent a rational number by a pair of integers, a numerator and a denominator. We may even assume the numerator and denominator are coprime: Given $a, b \in \mathbb{Z}$ we can compute their greatest common divisor $\gcd(a, b)$ and solve the Bézout equation $ax + by = \gcd(a, b)$ for some $x, y \in \mathbb{Z}$ using the (extended) Euclidean algorithm in polynomial time (we define $\gcd(0, 0) = 0$). Applying these techniques in bulk we can also do addition, subtraction and multiplication of integer and rational matrices in polynomial time. Least trivially of our building blocks, using the theory of lattices we can compute bases for the kernel and the image of an integer matrix in polynomial time, which is the topic of Section 3.

1.3. Commutative algebra and number theory

Definition 1.2. Let R be a commutative ring. We write R^* for the set of units and $\operatorname{spec} R$ for the set of prime ideals of R. We say R is *local* if R has a unique maximal ideal. For an ideal $I \subseteq R$ we define the *radical* $\sqrt{I}$ of I to be the ideal $\{x \in R \mid (\exists n \in \mathbb{Z}_{>0})\, x^n \in I\}$. We call $\operatorname{nil}(R) = \sqrt{0R}$ the *nilradical* of R and we call the elements of $\operatorname{nil}(R)$ the *nilpotents* of R. We say R is *reduced* when $\operatorname{nil}(R) = 0$. We say ideals $I, J \subseteq R$ are *coprime* when $I + J = R$.

Definition 1.3. A *number field* is a field K containing the field of rational numbers $\mathbb{Q}$ such that the dimension of K over $\mathbb{Q}$ as vector space is finite. A *number ring* is a subring of a number field.

For a number ring R, write $R_\mathfrak{p}$ for the localization at the prime ideal $\mathfrak{p}$ of R.

Definition 1.4. Let R be a number ring and let $K = R_{0R}$ be its field of fractions. A *fractional ideal* of R is a finitely generated non-zero R-submodule of K. A fractional ideal of R is *integral* if it is contained in R. A fractional ideal of R is *principal* if it is of the form xR for some $x \in K$. We write $I + J$ and $I \cdot J$ for the R-modules generated by $\{i + j \mid i \in I, j \in J\}$ respectively $\{i \cdot j \mid i \in I, j \in J\}$. A fractional ideal I of R is *invertible* if there exists some fractional ideal J of R such that $I \cdot J$ is principal. A non-invertible prime ideal is called *singular*. For fractional ideals I and J of R write $I : J = \{x \in K \mid xJ \subseteq I\}$.

Definition 1.5. An *order* is a commutative ring whose additive group is isomorphic to $\mathbb{Z}^n$ for some $n \in \mathbb{Z}_{\geqslant 0}$. We say R is an order of a number field K if R is an order such that $R \subseteq K$ and $R_{0R} = K$.

Any reduced order is contained in some finite product of number fields. In our algorithms we encode an order by first specifying its rank n, and then writing down its $n \times n$ multiplication table for the standard basis vectors. By distributivity this completely and uniquely defines a multiplication on the order. We encode a number field K simply by an order R such that $R_{0R} = K$.

1.4. Exercises

Exercise 1.6. Let $\mathbb{Q}^{\text{alg}}$ be some algebraic closure of $\mathbb{Q}$.

- **a.** Show that $\#\mathbb{Q}^{\text{alg}} = \#\mathbb{N}$.
- **b.** Show that there are precisely $\#\mathbb{N}$ number fields contained in $\mathbb{Q}^{\text{alg}}$.
- **c.** Show that there are precisely $\#\mathbb{R}$ subfields of $\mathbb{Q}^{\text{alg}}$.
- **d.** The same question as (a) and (b), but counting the fields up to isomorphism.
- **e.** Do there exist $\#\mathbb{R}$ subfields of $\mathbb{Q}^{\text{alg}}$ that are pairwise isomorphic?
- **f.** Let $K \neq \mathbb{Q}$ be a number field. Show that there are precisely $\#\mathbb{N}$ orders and $\#\mathbb{R}$ subrings in K with field of fractions K. What differs for $K = \mathbb{Q}$?
- **g.** Argue why it is natural to restrict the input of our algorithms to orders and number fields as opposed to general number rings.

Exercise 1.7. Show that there exists a polynomial-time algorithm that, given a cube integer matrix, determines whether it encodes an order.

Exercise 1.8. Let R be a commutative ring. Show that R is a local ring if and only if $R \setminus R^*$ is an additive subgroup of R.

Exercise 1.9. Let R be an order of a number field K. Prove that $R = \mathcal{O}_K$ if and only if the sum of every two invertible ideals is again invertible.

Exercise 1.10. Let R be a commutative ring. Show that $\text{nil}(R)$ is equal to the intersection of all prime ideals of R. Moreover, show that if $\text{nil}(R)$ is finitely generated, then $\text{nil}(R)$ is *nilpotent*, i.e. $\text{nil}(R)^n = 0$ for some $n > 0$.

Exercise 1.11. Show that any finite commutative domain is a field. Conclude that in a general commutative ring prime ideals of finite index are maximal.

Exercise 1.12. Let R be a commutative ring and let $I_1, \ldots, I_m, J_1, \ldots, J_n \subseteq R$ be ideals such that for all i, j we have $I_i + J_j = R$. Show that $I_1 \cdots I_m + J_1 \cdots J_n = R$.

Conclude that for any two distinct maximal ideals $\mathfrak{m}, \mathfrak{n} \subseteq R$ and any $m, n \in \mathbb{Z}_{\geqslant 0}$ we have $\mathfrak{m}^m + \mathfrak{n}^n = R$.

Exercise 1.13 (Chinese remainder theorem for ideals)**.** Let R be a commutative ring and let $I_1, \dots, I_n \subseteq R$ be pairwise coprime ideals. Show that $\bigcap_{i=1}^n I_i = \prod_{i=1}^n I_i$ and prove that the natural homomorphism

$$R/\Big(\bigcap_{i=1}^n I_i\Big) \to \prod_{i=1}^n (R/I_i)$$

is an isomorphism.

Exercise 1.14. Let I, J be fractional ideals of a number ring R.

a. Show that if $IJ = R$, then $J = R : I$.

b. Show that if I is invertible, then $I : I = R$.

c. Show that IJ is invertible if and only if I and J are invertible.

Exercise 1.15. Let I be a non-zero ideal in a number ring R. Show that if all prime ideals $I \subseteq \mathfrak{p} \subseteq R$ are invertible, then I factors as a product of invertible prime ideals and in particular is invertible.

Exercise 1.16. Let I be a non-zero ideal in a number ring R.

a. Show that there exist prime ideals $\mathfrak{p}_1, \dots, \mathfrak{p}_n \subseteq R$ that contain I and satisfy $\mathfrak{p}_1 \cdots \mathfrak{p}_n \subseteq I$.

b. Suppose I is a product of prime ideals. Show that $I : I = R$ if and only if I is invertible.

Hint: First suppose I is prime. If I is non-invertible and $a \in \mathfrak{p}$ is non-zero, then $\mathfrak{p}_1 \cdots \mathfrak{p}_t \subseteq aR \subseteq I$ and without loss of generality $\mathfrak{p}_1 = I$. Any $b \in \mathfrak{p}_2 \cdots \mathfrak{p}_n \setminus aR$ satisfies $bI \subseteq aR$.

Exercise 1.17. Let α be an algebraic integer of degree at least 3 and let p be a prime number. Show that $R = \mathbb{Z} + p\mathbb{Z}[\alpha]$ is a domain, and that $I = \alpha\mathbb{Z} + R$ is a fractional R-ideal. Moreover, prove that $I : I = R$ and that I is not invertible.
Note: This provides a counter-example to the converse of Exercise 1.14.b.

2. Coprime base factorization

In this section we treat the following problem, which will be the motivation for the *coprime base algorithm*.

Theorem 2.1. *There is a polynomial-time algorithm that on input* $t \in \mathbb{N}$, $q_1, \dots, q_t \in \mathbb{Q}^*$ *and* $n_1, \dots, n_t \in \mathbb{Z}$ *decides whether*

$$\prod_{i=1}^t q_i^{n_i} = 1. \tag{2.1}$$

It is clear we can determine whether such a product has the correct sign: Simply take the sum of all n_i for which $q_i < 0$ and check whether the result is even. It is then sufficient to prove the following theorem instead.

Theorem 2.2. *There is a polynomial-time algorithm that on input* $t \in \mathbb{N}$, $a_1, \dots, a_t, b_1, \dots, b_t \in \mathbb{Z}_{>0}$ *and* $n_1, \dots, n_t, m_1, \dots, m_t \in \mathbb{Z}$ *decides whether*

$$\prod_{i=1}^{t} a_i^{n_i} = \prod_{i=1}^{t} b_i^{m_i}. \tag{2.2}$$

In this form, the problem looks deceptively easy. Consider for example the most straightforward method to decide (2.2).

Method 2.3. Compute $\prod_{i=1}^{t} a_i^{n_i}$ and $\prod_{i=1}^{t} b_i^{m_i}$ explicitly and compare the results.

This method is certainly correct in that it is able to decide (2.2). However, it fails to run in polynomial time even when $t = 1$. For $n \in \mathbb{Z}_{>0}$ we have that $l(2^n) = n + 1 \approx 2^{l(n)}$. Hence the length of 2^n is not bounded by any polynomial in $l(n)$. We wouldn't even have enough time to write down the number regardless of our proficiency in multiplication because the number is too long.

Another method uses the fundamental theorem of arithmetic, also known as unique prime factorization in $\mathbb{Z}$.

Method 2.4. Factor $a_1, \dots, a_t, b_1, \dots, b_t$ into primes and for each prime that occurs compute the number of times it occurs in the products $\prod_{i=1}^{t} a_i^{n_i}$ and $\prod_{i=1}^{t} b_i^{m_i}$ and compare the results.

It is true that once we have factored all integers into primes only a polynomial number of steps remains. If we write x_{ip} for the exponent of the prime p in a_i, then we may compute $\sum_{i=1}^{t} n_i x_{ip}$, the exponent of p in $\prod_{i=1}^{t} a_i^{n_i}$, in polynomial time. Moreover, the number of prime factors of $n \in \mathbb{Z}_{>0}$ is at most $l(n)$, so the number of primes occurring is at most $\sum_{i=0}^{t}(l(a_i) + l(b_i))$, which is less than the length of the combined input. The problem lies in the fact that we have not specified how to factor integers into primes. As of July 21, 2022, nobody has been able to show that we can factor integers in polynomial time. Until this great open problem is solved, Method 2.4 is out the window.

An interesting observation is that the main obstruction in Method 2.3 lies in the exponents being large, while for Method 2.4 the obstruction is in the bases. Our proof for Theorem 2.2 will be to slightly tweak Method 2.4. Namely, observe that we do not need to factor into prime elements but that it suffices to factor into pairwise coprime elements. The following lemma follows readily from unique prime factorization.

Lemma 2.5 (Unique coprime factorization)**.** *Let* $s \in \mathbb{N}$ *and let* $c_1, \dots, c_s \in \mathbb{Z}_{>1}$ *be pairwise coprime. If for* $n_1, \dots, n_s, m_1, \dots, m_s \in \mathbb{Z}_{\geqslant 0}$ *we have*

$$\prod_{i=1}^{s} c_i^{n_i} = \prod_{i=1}^{s} c_i^{m_i}, \tag{2.3}$$

then $n_i = m_i$ *for all* i. □

We now propose the following algorithm for deciding (2.2).

Method 2.6. Factor $a_1, \dots, a_t, b_1, \dots, b_t$ into pairwise coprime $c_1, \dots, c_s \in \mathbb{Z}_{>0}$. For each c_i compute the number of times it occurs in $\prod_{i=1}^t a_i^{n_i}$ and $\prod_{i=1}^t b_i^{m_i}$ and compare the results.

Now to prove Theorem 2.2 and in turn Theorem 2.1 it suffices to prove the following.

Theorem 2.7 (Coprime base factorization)**.** *There is a polynomial-time algorithm that on input* $t \in \mathbb{N}$ *and* $a_1, \dots, a_t \in \mathbb{Z}_{>0}$ *computes* $s \in \mathbb{N}$, $c_1, \dots, c_s \in \mathbb{Z}_{>1}$ *and* $(n_{ij}) \in \mathbb{Z}_{\geqslant 0}^{t \times s}$ *such that* $c_1, \dots, c_s$ *are pairwise coprime and* $a_i = \prod_{j=1}^s c_j^{n_{ij}}$ *for all* i.

We state the algorithm first and prove the theorem later.

Method 2.8. Construct a complete simple graph G and label the vertices with $a_1, \dots, a_s$. We call it a labeling because the map sending a vertex to its label need not be injective. While there are edges in G, repeat the following 5 steps:

1. (1) Choose an edge $\{U, V\}$ of G, delete it from the graph, and let u and v be the labels of U respectively V.
2. (2) Compute $w = \gcd(u, v)$ using the Euclidean algorithm.
3. (3) Add a vertex W labeled w to G and connect it to U, V and those vertices which are neighbors of both U and V.
4. (4) Update the labels of U and V to u/w and v/w respectively.
5. (5) For each $S \in \{U, V, W\}$, if the label of S is 1, then delete S and its incident edges from G.

Now $V = \{c_1, \dots, c_s\}$ consists of the required pairwise coprime elements. The remaining output can now be computed in polynomial time.

In the graph that we construct and update, the edges represent the pairs of numbers of which we do not yet know whether they are coprime, while a missing edge denotes that we know the pair to be coprime.

Example 2.9. We apply Method 2.8 to $(a_1, a_2) = (4500, 5400)$. Since there are only two vertices our graphs will fit on a single line. We denote the edge we choose in each iteration with a bullet and edges we have to erase are dotted. On the right we show how to keep track of the factorization of 4500 with minimal bookkeeping by writing it as a product of vertices in the graph.

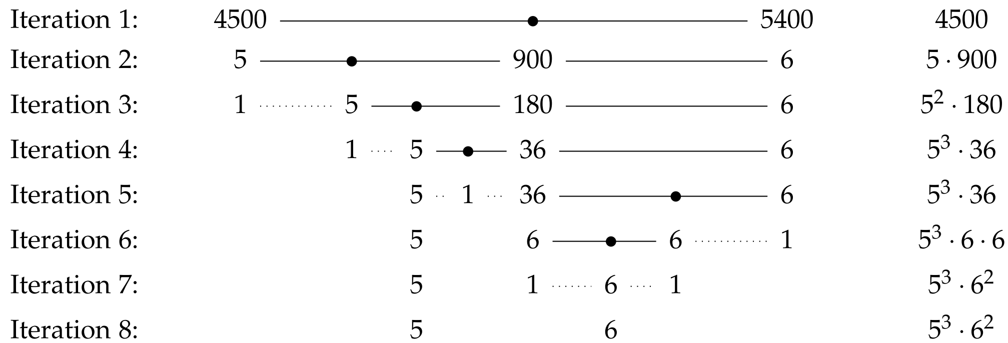

We obtain $(c_1, c_2) = (5, 6)$ and $4500 = 5^3 \cdot 6^2$. By trial division we obtain $5400 = 5^2 \cdot 6^3$.

Example 2.10. We apply Method 2.8 to $(a_1, a_2, a_3) = (15, 21, 35)$.

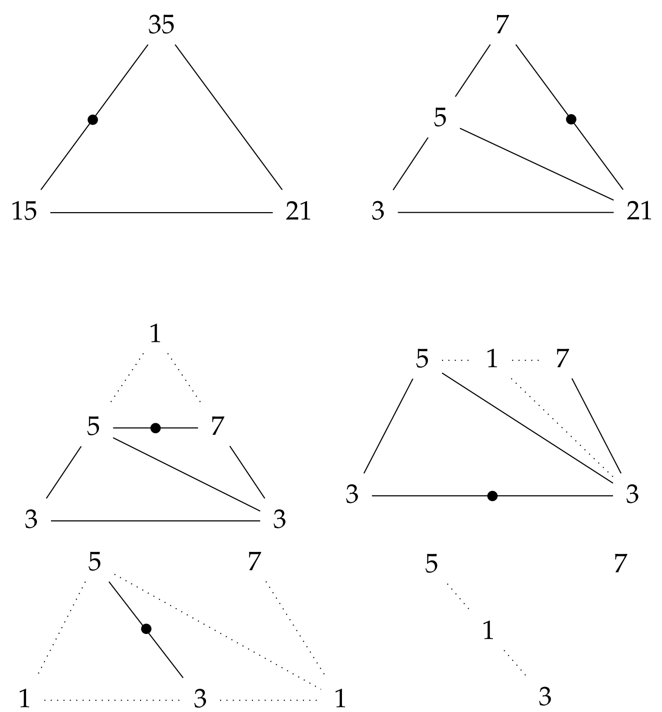

The resulting coprime base is $(c_1, c_2, c_3) = (3, 5, 7)$. In the fifth graph something interesting happens. Vertex 7 suddenly becomes disconnected from the graph because we know it is coprime to one of the 3's.

Proof of Theorem 2.7. We claim Method 2.8 is correct and runs in polynomial time. One can show inductively that throughout the algorithm two vertices in the graph G are coprime when there is no edge between them. When the algorithm terminates because there are no edges in the graph, we may conclude that $c_1, \ldots, c_s$ are coprime. Additionally, one shows inductively that the numbers $a_1, \ldots, a_t$ can be written as some product of the vertices of G. Hence $c_1, \dots, c_s$ forms a coprime base for $a_1, \dots, a_t$, so Method 2.8 is correct. It remains to show that it is fast.

Write P_n for the product of all labels of the vertices in the graph at step n. Note that $P_0 = a_1 \cdots a_s$ and that $P_{n+1} \mid P_n$ for all $n \in \mathbb{N}$. Since P_0 has at most $B = l(P_0)$ prime factors counting multiplicities, there are at most B steps n for which $P_{n+1} < P_n$ and at most B vertices in G. The steps for which $P_n = P_{n+1}$ are those where the edge we chose is between coprime integers, meaning no vertices or edges are added to the graph and one edge is deleted. As the number of edges is at most B^2, then so is the number of consecutive steps for which $P_n = P_{n+1}$. Hence the total number of steps is at most B^3, which is polynomial in the length of the input. Lastly, note that each step takes only polynomial time because the values of the vertices are bounded from above by B and the Euclidean algorithm runs in polynomial time. Hence Method 2.8 runs in polynomial time. □

The speed of this algorithm heavily depends on how fast the product of all vertices decreases. In each step we want to choose our edge $\{u, v\}$ such that $\gcd(u, v) \gg 1$. A heuristic for this could be to choose edges between large numbers. For those interested in the efficiency of coprime base factorization we refer to [2] for a provably faster algorithm.

Exercise 2.11. Show that there exists a polynomial-time algorithm that, given $a, b, c, d \in \mathbb{Z}_{>0}$ such that $ab = cd$, computes $w, x, y, z \in \mathbb{Z}_{>0}$ such that $(a, b, c, d) = (wx, yz, wz, xy)$.

Exercise 2.12 (Modified Euclidean algorithm)**.** Recall that $\gcd(0, 0) = 0$.

- **a.** Show that for all $a, b \in \mathbb{Z}$ with $b \neq 0$ there exist $r, q \in \mathbb{Z}$ with $a = qb + r$ and $|r| \leqslant |b|/2$.
- **b.** Show that for all $q, b, r \in \mathbb{Z}$ with $a = qb + r$ we have $\gcd(a, b) = \gcd(b, r)$ and $\gcd(a, 0) = |a|$.
- **c.** Prove that there exists a polynomial-time algorithm that, given $a, b \in \mathbb{Z}$, computes $\gcd(a, b)$ as well as $x, y \in \mathbb{Z}$ such that $ax + by = \gcd(a, b)$.
- **d.** Conclude that there exists a polynomial-time algorithm that, given $a, n \in \mathbb{Z}$ with $n > 1$, decides whether $a \in (\mathbb{Z}/n\mathbb{Z})^*$ and if so computes some $a' \in \mathbb{Z}$ such that $aa' \equiv 1 \bmod n$.
- **e.** Prove that there exists a polynomial-time algorithm that, given $a, b, m, n \in \mathbb{Z}$ with $n, m > 1$, decides whether there exists some $c \in \mathbb{Z}$ such that $c \equiv a \bmod m$ and $c \equiv b \bmod n$ and if so computes such a c.

Exercise 2.13. Show that there exists a polynomial-time algorithm that, given $k, a_1, \dots, a_k, n \in \mathbb{Z}$ satisfying $n \geqslant 1$ and $a_i^2 \equiv 1 \bmod n$ for all i, decides whether there exists some non-empty subset $I \subseteq \{1, \dots, k\}$ such that $\prod_{i \in I} a_i \equiv 1 \bmod n$ and if so computes one such I. *Hint:* Factor n.

Exercise 2.14. We equip $\mathbb{Q}^2 \setminus \{(0, 0)\}$ with an equivalence relation $\sim$ where $(x_1, y_1) \sim (x_2, y_2)$ if and only if there exists some $\lambda \in \mathbb{Q}^*$ such that $(\lambda x_1, \lambda y_1) = (x_2, y_2)$.

Write $\mathbb{P}^1(\mathbb{Q}) = (\mathbb{Q}^2 \setminus \{(0,0)\})/\sim$ for the *projective line* and write $(x : y)$ for the image of (x, y) in $\mathbb{P}^1(\mathbb{Q})$. Let $a, b \in \mathbb{Z}_{>0}$ and let $c_1, \ldots, c_n$ be the coprime base for a and b produced by Method 2.8.

a. For $p \mid ab$ prime write $f(p) = (\mathrm{ord}_p(a) : \mathrm{ord}_p(b)) \in \mathbb{P}^1(\mathbb{Q})$. Show that every c_i naturally corresponds to a fiber of f and give the prime factorization of c_i.

b. Suppose $n = 7$. Show that $ab \geqslant 1485890406000$ and equality holds for exactly 8 pairs (a, b).

c. (difficult) Give an asymptotic formula for the minimum of ab in terms of n.

Exercise 2.15. Let $n \in \mathbb{Z}_{>0}$. We encode matrices $\overline{M} = (\overline{m}_{ij})_{i,j}$ over $\mathbb{Z}/n\mathbb{Z}$ as a matrix $M = (m_{ij})_{i,j}$ over $\mathbb{Z}$ such that $0 \leqslant m_{ij} < n$ and $m_{ij} \equiv \overline{m}_{ij} \bmod n$ for all i, j. Show that there exist polynomial-time algorithms for the following problems:

a. given $n \in \mathbb{Z}_{\geqslant 0}$ and matrices M and N over $\mathbb{Z}/n\mathbb{Z}$, compute $M + N$ and $M \cdot N$ if well-defined;

b. given $n, k \in \mathbb{Z}_{>0}$ and a square matrix M over $\mathbb{Z}/n\mathbb{Z}$, compute M^k;
Note: An algorithm that takes k steps is not polynomial-time!

c. given $n \in \mathbb{Z}_{>0}$ and a matrix M over $\mathbb{Z}/n\mathbb{Z}$, compute a row-echelon form of M;

d. given $n \in \mathbb{Z}_{>0}$ and a square matrix M over $\mathbb{Z}/n\mathbb{Z}$, compute $\det(M)$ and $\mathrm{Tr}(M)$;

e. given $n \in \mathbb{Z}_{>0}$ and a matrix M over $\mathbb{Z}/n\mathbb{Z}$, decide whether M^{-1} exists and if so compute it.

You may use the following fact: For every $k, B \in \mathbb{Z}_{>0}$ and matrix $M = (m_{ij})_{i,j} \in \mathbb{Z}^{k \times k}$ with $|m_{ij}| \leqslant B$ for all i, j it holds that $|\det(M)| \leqslant B^k \cdot k^{k/2}$ (see Hadamard's inequality, Exercise 3.10).

f. Show that there exists a polynomial-time algorithm that, given a square integer matrix M, computes $\det(M)$ and $\mathrm{Tr}(M)$.

Exercise 2.16. Show that there exists a polynomial-time algorithm that, given $a, b, k, n \in \mathbb{Z}$ with $k, n > 0$ and $a^k \equiv 1 \bmod n$ and $b^k \equiv -1 \bmod n$, computes some $c \in \mathbb{Z}$ such that $a \equiv c^2 \bmod n$. *Hint:* First consider n odd and k a power of 2.

Exercise 2.17. Show that there exist polynomial-time algorithms for the following problems:

a. given $a, p, q \in \mathbb{Z}$ with p and q prime and $\gcd(a, p) = 1$, compute $u, e \in \mathbb{Z}$ such that $\gcd(u, q) = 1$ and such that the order of a in $(\mathbb{Z}/p\mathbb{Z})^*$ equals uq^e;

b. given $a, p \in \mathbb{Z}$ with p prime, decide whether a is a square modulo p;

c. given $a, b, p \in \mathbb{Z}$ with p prime, a a square modulo p and b not a square modulo p, compute $c \in \mathbb{Z}$ such that $c^2 \equiv a \bmod p$;

d. given $a, b, p \in \mathbb{Z}$ with p prime, compute $c \in \mathbb{Z}$ such that c^2 equals a, b or ab modulo p.

2.1. Coprime base factorization in number fields We would like to generalize Theorem 2.1 to arbitrary number fields, by which we mean that there is an additional input K, a number field, and that we take $q_1, \ldots, q_t \in K^*$. The theorem we will prove in the final sections is the following.

Theorem 2.18. *There exists a polynomial-time algorithm that, given a number field* K*, an* $n \in \mathbb{Z}_{\geqslant 0}$ *and* $a_1, \ldots, a_n \in K^*$*, computes the kernel of the map* $\mathbb{Z}^n \to \langle a_1, \ldots, a_n \rangle$ *given by* $(k_1, \ldots, k_n) \mapsto \prod_i a_i^{k_i}$.

In Section 3 we will actually define what it means to compute the kernel of a linear map. When we try to prove this theorem by generalizing the theorems from the previous section, we run into some classic problems in (computational) number theory.

Theorem 2.2, to which we reduce, is a statement about integers. Hence we replace $\mathbb{Z}$ by an order R in K. One problem is that R^* will generally contain more than just $\{\pm 1\}$, and it is not obvious how to pick a set $R_{>0} \subseteq R \setminus \{0\}$ of representatives of $(R \setminus \{0\})/R^*$ like the positive integers for $\mathbb{Z}$. Another problem is that we would like to at least compute the set R^*, which is a finitely generated abelian group of known rank by Dirichlet's unit theorem, but it is not known how to do this in polynomial time. Even if we disregard run-time issues, we want a Lemma 2.5 for orders. However, generally R will not be a UFD like $\mathbb{Z}$, making Theorem 2.7 hard to generalize.

The 'correct' way to generalize the theory is to translate it into a theorem about ideals, since the maximal order $\mathcal{O}_K$ of K has unique ideal factorization. Moreover, as opposed to the elements of $\mathcal{O}_K \setminus \{0\}$ themselves, the ideals are invariant under multiplication by units. However, computing $\mathcal{O}_K$ is also difficult. Luckily this is something we can work around. First we generalize Lemma 2.5.

Lemma 2.19 (Unique coprime factorization for ideals)**.** *Let* $s \in \mathbb{N}$*, let* R *be an order and let* $\mathfrak{c}_1, \ldots, \mathfrak{c}_s \subsetneq R$ *be pairwise coprime invertible integral ideals. If for* $n_1, \ldots, n_s$*,* $m_1, \ldots, m_s \in \mathbb{Z}_{\geqslant 0}$ *we have*

$$\prod_{i=1}^{s} \mathfrak{c}_i^{n_i} = \prod_{i=1}^{s} \mathfrak{c}_i^{m_i}, \tag{2.4}$$

then $n_i = m_i$ *for all* i.

Proof. Since the ideals are invertible we may divide out $\mathfrak{c}_i^{\min\{n_i, m_i\}}$ and thus assume without loss of generality that $n_i = 0$ or $m_i = 0$ for all i. But then the product on the left hand side of (2.4) and the product on the right hand side of (2.4) are coprime, so the products equal R. By the Chinese remainder theorem for ideals we get $0 = R/(\prod_{i=1}^s \mathfrak{c}_i^{n_i}) = \prod_{i=1}^s (R/\mathfrak{c}_i^{n_i})$, so $\mathfrak{c}_i^{n_i} = R$ for all i. If $n_i > 0$ we have $\mathfrak{c}_i \supseteq \mathfrak{c}_i^{n_i} = R$, so $\mathfrak{c}_i = R$, a contradiction. Thus $n_i = m_i = 0$ for all i. □

We would now like to prove the following theorem.

Theorem 2.20. *There exists a polynomial-time algorithm that, given an order* R*, an* $n \in \mathbb{Z}_{\geqslant 0}$ *and non-zero ideals* $\mathfrak{a}_1, \dots, \mathfrak{a}_n \subseteq R$*, computes either an order* $S \supsetneq R$ *or a coprime base* $\mathfrak{c}_1, \dots, \mathfrak{c}_m \subsetneq R$ *of invertible ideals for* $\mathfrak{a}_1, \dots, \mathfrak{a}_n$*.*

To prove this theorem we we are required to do some more work. First of all we require some definitions on how to encode orders and ideals. Then, we should construct algorithms to do arithmetic on ideals in polynomial time. More generally, we will study algorithms for finitely generated abelian groups in the next section.

Exercise 2.21. Suppose we can compute a basis of the image and kernel of a morphism $\mathbb{Z}^n \to \mathbb{Z}^m$ encoded by an integer matrix in polynomial time (we can, but this is non-trivial). For an order R of rank n we encode a fractional ideal $\mathfrak{a}$ of R as an injective linear map $\mathbb{Z}^n \to K$ with image $\mathfrak{a}$. Show that for an order R and fractional ideals $\mathfrak{a}$ and $\mathfrak{b}$ of R we may compute $\mathfrak{a} + \mathfrak{b}$, $\mathfrak{a} \cdot \mathfrak{b}$ and $\mathfrak{a} : \mathfrak{b}$ and decide whether $\mathfrak{a} = \mathfrak{b}$ in polynomial time.

Exercise 2.22. Suppose that we may compute for ideals $\mathfrak{a}$ and $\mathfrak{b}$ the ideals $\mathfrak{a} + \mathfrak{b}$, $\mathfrak{a} \cdot \mathfrak{b}$ and $\mathfrak{a} : \mathfrak{b}$ and decide whether $\mathfrak{a} = \mathfrak{b}$. Show that there exists an algorithm that, given an order R, $n \in \mathbb{Z}_{\geqslant 0}$ and non-zero ideals $\mathfrak{a}_1, \dots, \mathfrak{a}_n \subseteq R$, computes either an order $S \supsetneq R$ or a coprime base $\mathfrak{c}_1, \dots, \mathfrak{c}_m \subsetneq R$ of invertible ideals for $\mathfrak{a}_1, \dots, \mathfrak{a}_n$. Show that your algorithm runs in polynomial time when the input is restricted to ideals of the form aR with $a \in \mathbb{Z}_{>0}$. *Hint:* You can assume every ideal you encounter is invertible.

3. Finitely generated abelian groups

In this section we treat algorithms on finitely generated abelian groups. Many important objects in algebraic number theory are finitely generated abelian groups. For example, the additive group of orders R in number fields, as well as finitely generated modules over R, notably its ideals I and quotients R/I. In this section we will use additive notation for our abelian groups and we will use [3] as our reference. Other finitely generated abelian groups of interest are unit groups of finite commutative rings like $\mathbb{Z}/n\mathbb{Z}$ or $\mathbb{F}_q$, or an elliptic curve. However, as we will soon see, there is an obstruction in working with these groups.

We begin by specifying a representation for our finitely generated groups. Recall that every finitely generated abelian group A fits in some exact sequence

$$\mathbb{Z}^m \xrightarrow{\alpha} \mathbb{Z}^n \xrightarrow{f} A \longrightarrow 0.$$

Namely, we obtain n and f by writing down some generators $a_1, \dots, a_n \in A$ for A and let f map the i-th standard basis vector to a_i. For m and α we repeat the procedure with A replaced by $\ker(f)$. Note that α, being a morphism between free $\mathbb{Z}$-modules, has a natural representation as a matrix with integer coefficients. By the isomorphism theorem $A \cong \mathbb{Z}^n/\ker(f) = \mathbb{Z}^n/\operatorname{im}(\alpha) = \operatorname{coker}(\alpha)$, so A is completely defined by α. Thus we choose to encode A as the matrix corresponding to α. A morphism $f : A \to B$ of finitely generated abelian groups in terms of this representation gives a commutative diagram of exact sequences

$$\begin{array}{ccccccc} \mathbb{Z}^k & \xrightarrow{\alpha} & \mathbb{Z}^l & \longrightarrow & A & \longrightarrow & 0 \\ & & \downarrow{\scriptstyle\varphi} & & \downarrow{\scriptstyle f} & & \\ \mathbb{Z}^m & \xrightarrow{\beta} & \mathbb{Z}^n & \longrightarrow & B & \longrightarrow & 0. \end{array} \tag{3.1}$$

Here φ is any morphism that makes the diagram commute. We encode f by the matrix representing φ. Important to note is that not every φ defines a morphism $f : A \to B$. It defines a morphism precisely when $\mathrm{im}(\varphi \circ \alpha) \subseteq \mathrm{im}(\beta)$, however it is not immediately obvious how to test this. Computing the composition of morphisms and evaluating morphisms in this form is straightforward, as it is just matrix multiplication.

To work with abelian groups in our algorithms it takes more than just to specify an encoding. The following is a list in no particular order of operations we would like to be able to perform in polynomial time.

(1) decide whether a matrix encodes a morphism of given groups;
(2) compute kernels, images and cokernels of group homomorphisms;
(3) test if a group homomorphism is injective/surjective and if bijective compute an inverse;
(4) decide if two group homomorphisms are equal;
(5) compute an element in the preimage of a given group element under a group homomorphism;
(6) compute direct sums, tensor products and homomorphism groups of pairs of groups;
(7) compute the order of a given group element;
(8) compute the order/exponent of a finite group;
(9) split exact sequences;
(10) compute the torsion subgroup of a group;
(11) write a group as a direct sum of cyclic groups.

We will spend this section working up to the last element of this list: An algorithmic version of the fundamental theorem of finitely generated abelian groups.

All algorithms for the above problems will be very straight-forward. The only serious complication arises at the fundamentals, namely that when doing linear algebra over the integers you need your coefficients to remain small after every manipulation. For this, lattice basis reduction helps.

Finally, we address an important subtlety that arises from our choice of encoding.

Lemma 3.1. *Assuming the above problems have polynomial-time algorithms, we may solve the* discrete logarithm problem *in polynomial time. That is, given an abelian* *group* A *and elements* $a, b \in A$*, decide whether there exists some positive integer* n *such that* $na = b$ *and if so compute such* n*.*

It is well known that the discrete logarithm problem for $\mathbb{F}_q^*$ or elliptic curves over finite fields is difficult, i.e. not known to be solvable in polynomial time, even though both are finitely generated abelian groups. The difficulty is representing $\mathbb{F}_q^*$ and its elements in our encoding. For starters, we need to write down generators for $\mathbb{F}_q^*$ and subsequently write the input to our algorithms in terms of these generators. Doing so is almost equivalent to the discrete logarithm problem.

Exercise 3.2. Prove Lemma 3.1.

Exercise 3.3. Show that there exists a polynomial-time algorithm that, given finitely generated abelian groups A and B, computes the group $A \times B$ and the corresponding inclusions and projections.

3.1. Lattices and short bases To understand general finitely generated abelian groups, we first need to understand the simplest instances, the free abelian groups. It will turn out to be fruitful to consider free abelian groups together with an inner product. This will allow us later to compute images and kernels of linear maps.

Definition 3.4. A *Euclidean (vector) space* is a finite-dimensional real inner product space. For an element x in an inner product space we will write $q(x) = \langle x, x \rangle$. A *lattice* is a discrete subgroup of a Euclidean space.

A Euclidean space we naturally encounter for any number field K is $K \otimes_{\mathbb{Z}} \mathbb{R}$, which we equip with the inner product

$$\langle x, y \rangle = \frac{1}{[K : \mathbb{Q}]} \sum_{\sigma : K \otimes_{\mathbb{Z}} \mathbb{R} \to \mathbb{C}} \sigma(x) \cdot \overline{\sigma(y)}$$

where the sum ranges over all $\mathbb{R}$-algebra homomorphisms. In this Euclidean space every order of K is a lattice.

Proposition 3.5. *A lattice* Λ *in a Euclidean space* V *is a free* $\mathbb{Z}$*-module with* $\operatorname{rk} \Lambda \leqslant \dim V$ *and the restriction of the inner product to* Λ *is* $\mathbb{Z}$*-bilinear, real-valued, symmetric and satisfies* $\inf\{\langle x, x \rangle \mid x \in \Lambda \setminus \{0\}\} > 0$*. Conversely, every free* $\mathbb{Z}$*-module* Λ *of finite rank equipped with a* $\mathbb{Z}$*-bilinear, real-valued, symmetric form* φ *for which* $\inf\{\varphi(x, x) \mid x \in \Lambda \setminus \{0\}\} > 0$ *can be embedded in a Euclidean vector space such that the inner product restricted to* Λ *equals* φ*.* □

This proposition shows that we have an equivalent definition of a lattice that does not require an ambient vector space. The bilinear form φ in the proposition is again naturally given by a matrix $F = (\varphi(b_i, b_j))_{1 \leqslant i,j \leqslant n}$, the *Gram-matrix*, where $(b_1, \ldots, b_n)$ is the basis encoding Λ. For our computational purposes it is practical to restrict to matrices with rational entries. This will be our encoding for lattices.

An algorithmic problem we will encounter is computing a 'short basis' for a lattice $\Lambda \subseteq \mathbb{Z}^n$.

Definition 3.6. Let Λ be a lattice and let $(b_1, \dots, b_n)$ be a basis of Λ. Consider the matrix $B = (\langle b_i, b_j \rangle)_{1 \leqslant i,j \leqslant n}$. We define the *determinant* of Λ to be $\det(\Lambda) = |\det(B)|^{1/2}$.

Exercise 3.7. Show that the determinant of a lattice does not depend on the choice of basis.

The determinant $\det(\Lambda)$ also equals the volume of the parallelepiped spanned by a basis of Λ. Since the determinant is an invariant, finding a 'shorter' basis is equivalent to finding a 'more orthogonal' basis. In the lattice $\mathbb{Z}b_1 + \mathbb{Z}b_2$ below we can find a better basis.

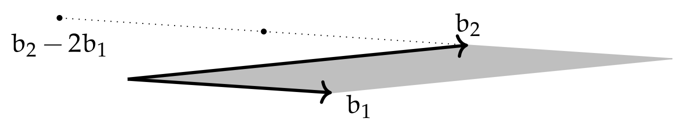

Taking $c_1 = b_1$ and $c_2 = b_2 - 2b_1$ produces the following basis:

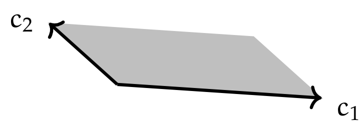

For a Euclidean space the Gram–Schmidt algorithm transforms a basis into an orthogonal one as follows.

Definition 3.8. Let V be a Euclidean space with basis $B = (b_1, \dots, b_n)$. We iteratively define

$$\mu_{ij} = \frac{\langle b_i, b_j^* \rangle}{\langle b_j^*, b_j^* \rangle} \quad \text{for } 1 \leqslant j < i \text{ and} \quad b_i^* = b_i - \sum_{j<i} \mu_{ij} b_j^* \quad \text{for } 1 \leqslant i \leqslant n.$$

We call $B^* = (b_1^*, \dots, b_n^*)$ and $(\mu_{ij})_{j<i}$ the *Gram–Schmidt basis* respectively *Gram–Schmidt coefficients* corresponding to B.

When interpreting $M = (\mu_{ij})_{j<i}$ as an upper-triangular matrix, we note that that $(\mathrm{id} + M)B^* = B$. In particular $\det(B) = \det(B^*)$. Since $b_1^*, \dots, b_n^*$ are indeed pairwise orthogonal, they form an orthogonal basis of V. Sadly the Gram–Schmidt coefficients will generally not be integers, meaning that if B is a basis for a lattice Λ, then generally B^* will not be. This is quite unsurprising, as not every lattice even has an orthogonal basis. A possible solution is to round the Gram-Schmidt coefficients to integers in every step, so that we are guaranteed to obtain a basis for Λ. However, this does not yield the necessary bounds on our basis.

Exercise 3.9. (difficult) We say a basis $B = (b_1, \dots, b_n)$ is *Gram–Schmidt reduced* if $|\mu_{ij}| \leqslant \frac{1}{2}$ holds for all Gram–Schmidt coefficients μ_{ij} of B. Show that the following algorithm is guaranteed to terminate, and hence computes a Gram–Schmidt reduced basis: Let $(b_1, \dots, b_n)$ be a basis.

(1) Compute the Gram–Schmidt coefficients $(\mu_{ij})_{j<i}$ of $(b_1, \dots, b_n)$.
(2) If $|\mu_{ij}| \leqslant \frac{1}{2}$ for all i, j, then return $(b_1, \dots, b_n)$ and terminate.
(3) Choose any i, j such that $|\mu_{ij}| > \frac{1}{2}$ and replace b_i by $b_i - \lceil \mu_{ij} \rfloor b_j$.
(4) Go to step 1.

The algorithm of Exercise 3.9 will not run in polynomial time. In the next section we will state the existence of a better algorithm.

Exercise 3.10. Let $(b_1, \dots, b_n)$ be a basis of a lattice Λ in a Euclidean space V. Write $\Lambda_k = \sum_{j \leqslant k} \mathbb{Z} b_i$ for all $0 \leqslant k \leqslant n$.

a. Show that

$$q(b_i) \geqslant q(b_i^*) = \Big(\frac{\det(\Lambda_i)}{\det(\Lambda_{i-1})}\Big)^2.$$

b. Conclude *Hadamard's inequality*: For B the matrix with columns $b_1, \dots, b_n$ we have

$$|\det(B)| \leqslant \prod_{i=1}^{n} \|b_i\|$$

with equality if and only if the b_i are pairwise orthogonal.

Exercise 3.11. Let V be a Euclidean space. For a subspace $W \subseteq V$ we write $W^\perp = \{v \in V \mid \langle v, W \rangle = 0\}$.

a. Show that for all $W \subseteq V$ the natural map $W^\perp \to V/W$ is an isomorphism of vector spaces.

We equip V/W with the natural Euclidean vector space structure induced by $W^\perp$. Suppose $\Lambda \subseteq V$ is a lattice with a sublattice Λ' such that Λ/Λ' is a torsion free group.

b. Show that the natural map $\Lambda/\Lambda' \to V/\mathbb{R}\Lambda'$ is injective and that its image is a lattice.

c. Show that $\det(\Lambda) = \det(\Lambda/\Lambda') \cdot \det(\Lambda')$.

Exercise 3.12. Let Λ be a lattice and define the $\Lambda^\dagger$ to be the group $\mathrm{Hom}(\Lambda, \mathbb{Z})$ together with the map $\langle \cdot, \cdot \rangle : \Lambda^\dagger \times \Lambda^\dagger \to \mathbb{R}$ given by

$$\langle f, g \rangle = \sup_{x \in \Lambda \setminus \{0\}} \frac{f(x) g(x)}{\langle x, x \rangle}.$$

a. Show that $\Lambda^\dagger$ is a lattice. We will call $\Lambda^\dagger$ the *dual lattice* of Λ.

b. Suppose $\Lambda \subseteq \mathbb{R}^n$ has rank n and let $\Lambda' = \{x \in \mathbb{R}^n \mid \langle x, \Lambda \rangle \subseteq \mathbb{Z}\}$. Show that Λ' is a lattice.

c. Show that $\Lambda^\dagger \cong \Lambda'$ and $(\Lambda^\dagger)^\dagger \cong \Lambda$.

d. Show that $\det(\Lambda^\dagger) = \det(\Lambda)^{-1}$.

e. For a homomorphism $\varphi : \Lambda_1 \to \Lambda_2$ of groups write $\varphi^\dagger : \Lambda_2^\dagger \to \Lambda_1^\dagger$ for the map $f \mapsto f \circ \varphi$. Show that

$$\det(\ker(\varphi)) \cdot \det(\mathrm{im}(\varphi)) \cdot \det(\Lambda_1^\dagger) = \det(\ker(\varphi^\dagger)) \cdot \det(\mathrm{im}(\varphi^\dagger)) \cdot \det(\Lambda_2).$$

3.2. The LLL-algorithm In this section we state the existence the LLL-algorithm, which produces a 'small' basis for a given lattice. We will not prove the correctness of the algorithm, nor will we actually describe the algorithm. What we will do is define what we mean by 'small' bases in the context of the LLL-algorithm and derive some of their properties. In this section we will use [6] as our reference. A reference on the LLL-algorithm we do not draw upon which may be of interest to a reader is [8].

Definition 3.13. Let Λ be a lattice with basis $B = (b_1, \ldots, b_n)$ and let $(b_1^*, \ldots, b_n^*)$ and $(\mu_{ij})_{j<i}$ be its corresponding Gram–Schmidt basis and coefficients as defined in Definition 3.8. Let $\frac{4}{3} \geqslant c$. We say B is c-reduced if

(1) For all $1 \leqslant j < i \leqslant n$ we have $|\mu_{ij}| \leqslant \frac{1}{2}$.
(2) For all $1 \leqslant k < n$ we have $c\,q(b_{k+1}^*) \geqslant q(b_k^*)$.

Note that the first condition states that B is Gram–Schmidt reduced in the sense of Exercise 3.9. We may compute a c-reduced basis, and in particular it always exists, by Exercise 3.20. That we may in fact compute it in polynomial time is non-trivial.

Theorem 3.14 (LLL-algorithm)**.** *Let $c > \frac{4}{3}$. There exists a polynomial-time algorithm that, given a lattice Λ, produces a c-reduced basis of Λ.* □

Although the algorithm does not fundamentally differ when we modify c, it cannot be part of the input because then the algorithm would no longer run in polynomial time. We should warn the reader that the literature contains various definitions of a 'reduced basis', and even in the context of the LLL-algorithm there are at least two.

Definition 3.15. Let Λ be a lattice of rank n. For $0 < i \leqslant n$ we define the i*-th successive minimum* to be the value

$$\lambda_i(\Lambda) = \min\{r \in \mathbb{R}_{\geqslant 0} \mid \operatorname{rk}\langle x \in \Lambda \mid q(x) \leqslant r\rangle \geqslant i\}.$$

A basis of vectors attaining the successive minima is the gold standard of 'small' bases, although it does not always exist. The following proposition gives bounds on how far away a reduced basis can be from these minima.

Proposition 3.16. *Let $c \geqslant \frac{4}{3}$ and suppose $(b_1, \ldots, b_n)$ is a c-reduced basis for a lattice Λ. Then for all $0 < i \leqslant n$ we have $c^{1-n} \cdot q(b_i) \leqslant \lambda_i(\Lambda) \leqslant c^{i-1} \cdot q(b_i)$.*

Proof. See Exercise 3.18. □

Exercise 3.17. Show that there exists a lattice Λ for which no basis $b_1, \ldots, b_n$ attains the successive minima, i.e. satisfies $q(b_i) = \lambda_i(\Lambda)$ for all i. *Hint:* Consider $2\mathbb{Z}^n \subseteq \Lambda \subseteq \mathbb{Z}^n$.

Exercise 3.18. Let $c \geqslant \frac{4}{3}$ and $0 < i \leqslant n$ and suppose $(b_1, \ldots, b_n)$ is a c-reduced basis of Λ.

a. Show that $q(b_j^*) \leqslant c^{i-j} q(b_i^*)$ for all $j \leqslant i$.

b. Recall that $b_i = b_i^* + \sum_{j<i} \mu_{ij} b_j^*$. Show that $q(b_i) \leqslant c^{i-1} q(b_i^*)$.

c. Show that $q(b_j) \leqslant c^{i-1} q(b_i^*) \leqslant c^{i-1} q(b_i)$ for all $j \leqslant k$.

d. Conclude that $\lambda_i(\Lambda) \leqslant \max\{q(b_j) \mid j \leqslant i\} \leqslant c^{i-1} q(b_i)$.

Write $\Lambda_k = \sum_{j \leqslant k} \mathbb{Z} b_j$ for all $0 \leqslant k \leqslant n$.

e. Prove that for all $0 < k \leqslant n$ and $x \in \Lambda_k \setminus \Lambda_{k-1}$ we have $q(x) \geqslant q(b_k^*)$.

Write $S = \{x \in \Lambda \mid q(x) \leqslant \lambda_i(\Lambda)\}$ and let k be minimal such that $S \subseteq \Lambda_k$.

f. Show that $k \geqslant \mathrm{rk}\langle S \rangle \geqslant i$.

g. Conclude that $\lambda_i(\Lambda) \geqslant q(b_k^*) \geqslant c^{1-n} q(b_i)$.

Exercise 3.19. Let $c \geqslant \frac{4}{3}$ and suppose $(b_1, \ldots, b_n)$ is a c-reduced basis of a lattice Λ. Show that

$$\det(\Lambda)^2 \leqslant \prod_{i=1}^{n} q(b_i) \leqslant c^{\binom{n}{2}} \det(\Lambda)^2.$$

Exercise 3.20. Let $c > \frac{4}{3}$. Show that the following algorithm is guaranteed to terminate, and hence computes a c-reduced basis: Let $(b_1, \ldots, b_n)$ be a basis.

(1) Compute the Gram–Schmidt basis $(b_1^*, \ldots, b_n^*)$ and coefficients $(\mu_{ij})_{j<i}$ of $(b_1, \ldots, b_n)$.

(2) If there exists some $1 \leqslant k < n$ such that $c q(b_{k+1}^*) < q(b_k^*)$ and $|\mu_{k+1,k}| \leqslant \frac{1}{2}$, choose any such k, swap b_{k+1} and b_k and go to step 1.

(3) If there exist some $1 \leqslant j < i \leqslant n$ such that $|\mu_{ij}| > \frac{1}{2}$, choose any such i, j, replace b_i by $b_i - \lceil \mu_{ij} \rfloor b_j$ and go to step 1.

(4) Return $(b_1, \ldots, b_n)$ and terminate.

Hint: Let $\Lambda_k = \sum_{i=1}^{k} \mathbb{Z} b_i$. What can you say about $\prod_{k=1}^{n} \det(\Lambda_k)$ in step 2?

3.3. The kernel-image algorithm The LLL-algorithm allows us to prove the kernel-image algorithm, from which most of the algorithms for finitely generated abelian groups from the beginning of this section follow without much effort. We first need a theorem from linear algebra.

Theorem 3.21 (Cramer's rule)**.** *Suppose $A \in \mathbb{R}^{n \times n}$ is an invertible matrix and let $b \in \mathbb{R}^n$. Write A_i for the matrix obtained from A by replacing the i-th column with b. Then there exists a unique $x \in \mathbb{R}^n$ such that $Ax = b$, and it is given by* $x = \det(A)^{-1} \cdot (\det(A_1), \ldots, \det(A_n))$. □

Exercise 3.22. Let $n \geqslant 0$. Suppose $N \subseteq M \subseteq \mathbb{Z}^n$ are subgroups such that $N \oplus P = \mathbb{Z}^n$ for some $P \subseteq \mathbb{Z}^n$. Show that $\mathrm{rk}(M) = \mathrm{rk}(N)$ if and only if $M = N$.

Theorem 3.23 (Kernel-image algorithm)**.** *There exists a polynomial-time algorithm that, given a linear map $\varphi : \mathbb{Z}^n \to \mathbb{Z}^m$, computes the rank r of φ and injective linear maps $\iota : \mathbb{Z}^r \to \mathbb{Z}^n$ and $\kappa : \mathbb{Z}^{n-r} \to \mathbb{Z}^n$ such that* $\mathrm{im}(\varphi \circ \iota) = \mathrm{im}(\varphi)$ *and* $\mathrm{im}(\kappa) = \ker(\varphi)$.

Proof. Write B for the largest absolute value of a coefficient of the matrix defining φ. Compute

$$\omega = 2^{n-1} \cdot n^{n+1} \cdot B^{2n} + 1$$

and note that its length is polynomially bounded by the size of the input. Hence we can consider the lattice $L = \mathbb{Z}^n$ together with the bilinear form given by $q(x) = \|x\|^2 + \omega\|\varphi(x)\|^2$. Using the LLL-algorithm, compute a 2-reduced basis $(b_1, \ldots, b_n)$ of L. We will show that this basis has the following properties:

(a) $q(b_i) < \omega$ for $0 < i \leqslant n - r$;
(b) $(b_1, \ldots, b_{n-r})$ is a basis for $\ker(\varphi)$;
(c) $q(b_i) \geqslant \omega$ for $n - r < i \leqslant n$;
(d) $(\varphi(b_{n-r+1}), \ldots, \varphi(b_n))$ is a basis for $\mathrm{im}(\varphi)$.

Once we have shown this, it is clear how to compute r and the maps ι and κ in polynomial time.

<u>*Claim:*</u> For all $0 < i \leqslant n - r$ we have $\lambda_i(\Lambda) \leqslant n^{n+1}B^{2n}$.

Proof. By Cramer's rule, we may find linearly independent vectors $a_1, \ldots, a_{n-r} \in \ker(\varphi)$ for which the coefficients are determinants of $r \times r$ submatrices of F. Then by Hadamard's inequality (Exercise 3.10), each coefficient is bounded in absolute value by $r^{r/2}B^r \leqslant n^{n/2}B^n$. Hence $q(a_i) = \|a_i\|^2 \leqslant n^{n+1}B^{2n}$ for all $0 < i \leqslant n - r$. The claim now follows from the independence of the a_i. ■

From Proposition 3.16 and the claim it follows that

$$q(b_i) \leqslant 2^{n-1} \cdot \lambda_i(\Lambda) \leqslant 2^{n-1} \cdot n^{n+1} \cdot B^{2n} < \omega$$

for all $0 < i \leqslant n - r$, proving (a). Clearly for all $x \in \Lambda$ such that $\omega > q(x) = \|x\|^2 + \omega\|\varphi(x)\|^2$ we have $\|\varphi(x)\|^2 = 0$ and thus $x \in \ker(\varphi)$. In particular, we have linearly independent $b_1, \ldots, b_{n-r} \in \ker(\varphi)$. From Exercise 3.22 we may conclude it is in fact a basis for $\ker(\varphi)$, proving (b). It follows from (b) that $b_i \notin \ker(\varphi)$ and thus $q(b_i) \geqslant \omega$ for all $n - r < i \leqslant n$, proving (c). Lastly, (d) follows from (b) and the homomorphism theorem. □

Note that in the proof of Theorem 3.23 we could have constructed a c-reduced basis for values of c other than 2. Moreover, the exact value of ω is not important, as long as it is sufficiently large (while still being computable in polynomial time). Exercise 3.33 will prove a version of the kernel image algorithm for general finitely generated abelian groups.

Exercise 3.24. Show that for a matrix $\varphi : \mathbb{Z}^n \to \mathbb{Z}^n$ we have $\#\operatorname{coker}(\varphi) = |\det(\varphi)|$ if $\det(\varphi) \neq 0$. Conclude that there exist a polynomial-time algorithm that, given an abelian group A, decides whether A is finite and if so computes #A. *Note:* The matrix representing A need not be injective.

3.4. Applications of the kernel-image algorithm In this subsection we provide a polynomial-time algorithm for most problems in the beginning of this section. An immediate consequence of the kernel-image algorithm is the following.

Corollary 3.25. *There exists a polynomial-time algorithm that, given a linear map* $\varphi : A \to B$ *of finitely generated free abelian groups, decides whether* φ *is injective/surjective.*

Proof. Injectivity is obvious. For surjectivity add Exercise 3.24. □

Proposition 3.26. *There exists a polynomial-time algorithm that, given linear maps $\varphi : A \to C$ and $\psi : B \to C$ of finitely generated free abelian groups, decides whether $\mathrm{im}(\psi) \subseteq \mathrm{im}(\varphi)$.*

Proof. Using Theorem 3.23 compute the kernel $\kappa : K \to A \times B$ of $A \times B \to C$ as in the following diagram.

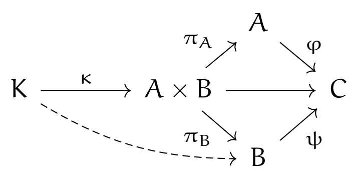

The image of $\pi_B \circ \kappa$ is precisely the set of elements $b \in B$ for which there exists an $a \in A$ such that $\varphi(a) = \psi(b)$. Hence it suffices to decide using Corollary 3.25 whether $\pi_B \circ \kappa$ is surjective. □

Corollary 3.27. *There exists a polynomial-time algorithm that, given finitely generated abelian groups A and B, represented by linear maps $\alpha : A_0 \to A_1$ respectively $\beta : B_0 \to B_1$, and a linear map $\varphi : A_1 \to B_1$, decides whether φ represents a morphism $f : A \to B$.*

Proof. Recall φ represents a morphism precisely when $\mathrm{im}(\varphi \circ \alpha) \subseteq \mathrm{im}(\beta)$. □

Corollary 3.28. *There exists a polynomial-time algorithm that, given finitely generated abelian groups A and B and morphisms $f, g : A \to B$, decides whether $f = g$.*

Proof. Considering $h = f - g$ it suffices to be able to decide whether a morphism is zero. With η the matrix representing h and β the matrix representing B, we have $h = 0$ precisely when $\mathrm{im}(\eta) \subseteq \mathrm{im}(\beta)$. □

Proposition 3.29. *There exists a polynomial-time algorithm that, given a morphism $f : A \to B$ of finitely generated abelian groups, decides whether f is injective/surjective.*

Proof. Write $\alpha : A_0 \to A_1$ and $\beta : B_0 \to B_1$ for the representatives of A respectively B and φ for the representative of f. Note that f is surjective if and only if $B = \mathrm{im}(f) = \mathrm{im}(\varphi)/\mathrm{im}(\beta)$ if and only if $B_1 = \mathrm{im}(\varphi) + \mathrm{im}(\beta)$. It suffices to decide whether the map $A_1 \times B_0 \to B_1$ induced by φ and β is surjective, for which we have Corollary 3.25. Note that f is injective if and only if $\ker(\varphi) \subseteq \mathrm{im}(\alpha)$. Using Theorem 3.23 we compute a linear map $\kappa : K \to A_1$ with $\mathrm{im}(\kappa) = \ker(\varphi)$, and apply Proposition 3.26 to decide $\mathrm{im}(\kappa) \subseteq \mathrm{im}(\alpha)$. □

Proposition 3.30. *There exists a polynomial-time algorithm that, given a linear map $\varphi : A \to B$ of free abelian groups A and B and $b \in B$, decides whether an element $a \in A$ exists such that $\varphi(a) = b$, and if so computes one.*

Proof. Consider the linear map $\psi : A \times \mathbb{Z} \to B$ that sends (a, x) to $\varphi(a) + xb$.

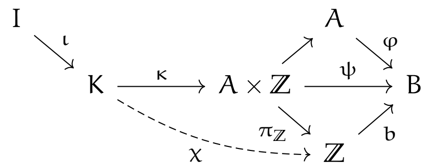

Compute using Theorem 3.23 the kernel $\kappa : K \to A \times \mathbb{Z}$ of ψ and in turn $\iota : I \to K$ a preimage of $\chi = \pi_{\mathbb{Z}} \circ \kappa$. Note that $\varphi(a) = b$ has a solution precisely when -1 is in the image of χ. Moreover, if we find $k \in K$ such that $\chi(k) = -1$, then its image under $K \to A \times \mathbb{Z} \to A$ gives an element $a \in A$ such that $\varphi(a) = b$. Note that $I \subseteq \mathbb{Z}$. If $I = 0$ no solution to $\chi(k) = -1$ exists, and otherwise $k \in \iota(\{\pm 1\})$ gives a solution if it exists. □

Corollary 3.31. *There exists a polynomial-time algorithm that, given a homomorphism* $f : A \to B$ *of finitely generated abelian groups* A *and* B *and* $b \in B$*, decides whether an element* $a \in A$ *exists such that* $f(a) = b$*, and if so computes one.*

Proof. Let $\varphi : A_1 \to B_1$ be the representative of f. Simply apply Proposition 3.30 to φ and some representative of b in B_1. □

Exercise 3.32. Give a direct proof of Proposition 3.26 or Proposition 3.30 by giving an algorithm that applies the LLL-algorithm only once, similar to Theorem 3.23.

Exercise 3.33. Show that there exists a polynomial-time algorithm that, given a morphism $f : A \to B$ of finitely generated abelian groups, computes injective morphisms $k : K \to A$ and $i : I \to B$ such that $\mathrm{im}(k) = \ker(f)$ and $\mathrm{im}(i) = \mathrm{im}(f)$.

3.5. Homomorphism groups of finitely generated abelian groups Although we can algorithmically work with individual morphisms $f : A \to B$ of finitely generated abelian groups, we have yet to treat the group of homomorphisms $\mathrm{Hom}(A, B)$ as a whole. Certainly, we would like to compute $\mathrm{Hom}(A, B)$, but we have to decide what that means. Firstly, we have to give an abelian group H represented by $H_0 \to H_1$ such that $H \cong \mathrm{Hom}(A, B)$. Secondly, we want to evaluate elements of $\mathrm{Hom}(A, B)$ at elements of A. For this we give two possibilities:

- Note that morphisms are already encoded as matrices in $\mathbb{Z}^{n \times m}$, so we simply give a map $H_1 \to \mathbb{Z}^{n \times m}$ that maps H_1 to matrices representing morphisms $A \to B$ and H_0 to the zero morphisms.
- We give a bilinear map $A \times H \to B$, i.e. a linear map $A \otimes H \to B$, which corresponds to the evaluation map $A \times \mathrm{Hom}(A, B) \to B$ under the isomorphism $H \cong \mathrm{Hom}(A, B)$.

Regardless of which representation we choose, when we talk about computing $\mathrm{Hom}(A, B)$ we mean computing both a group isomorphic to $\mathrm{Hom}(A, B)$ as well as a way to evaluate its elements in A.

Exercise 3.34. Show that the above representations for $\mathrm{Hom}(A, B)$ are 'polynomially equivalent', i.e. there exist polynomial-time algorithm that transforms one representation of $\mathrm{Hom}(A, B)$ into the other.

The moral of Exercise 3.34 is that as long as it is easy to describe how elements from H correspond to homomorphisms, it does not matter how we encode this.

Theorem 3.35. *There exists a polynomial-time algorithm that, given finitely generated abelian groups* A *and* B*, computes* $\mathrm{Hom}(A, B)$.

Proof. Consider the case where A is a free abelian group and B is a finitely generated abelian group represented by $\beta : B_0 \to B_1$. We may compute the matrix $\beta_* : \mathrm{Hom}(A, B_0) \to \mathrm{Hom}(A, B_1)$ given by $f \mapsto \beta \circ f$. Since A is free we get an exact functor $\mathrm{Hom}(A, \_)$ such that

$$\left[\ B_0 \xrightarrow{\beta} B_1 \to B \to 0\ \right] \xRightarrow{\mathrm{Hom}(A,\_)} \left[\ \mathrm{Hom}(A, B_0) \xrightarrow{\beta_*} \mathrm{Hom}(A, B_1) \to \mathrm{Hom}(A, B) \to 0\ \right],$$

and thus $\mathrm{Hom}(A, B) \cong \mathrm{coker}(\beta_*)$. Evaluating an element of $\mathrm{Hom}(A, B)$ in A reduces to evaluating an element of $\mathrm{Hom}(A, B_1)$ in A, which is just matrix multiplication.

Now consider the general case where A and B are general finitely generated abelian groups represented by $\alpha : A_0 \to A_1$ respectively β. We note that the functor $\mathrm{Hom}(\_\_, B)$ is left-exact and contravariant. Applied to the exact sequence of A we get

$$\left[\ A_0 \xrightarrow{\alpha} A_1 \to A \to 0\ \right] \xRightarrow{\mathrm{Hom}(\_\_,B)} \left[\ 0 \to \mathrm{Hom}(A, B) \to \mathrm{Hom}(A_1, B) \xrightarrow{\alpha^*} \mathrm{Hom}(A_0, B)\ \right].$$

Hence $\mathrm{Hom}(A, B) \cong \ker(\alpha^*)$. By the previous case we may compute $\mathrm{Hom}(A_1, B)$ and $\mathrm{Hom}(A_0, B)$ since A_1 and A_0 are free. It is not difficult to show we may then compute α^* and in turn its kernel using Exercise 3.33. Evaluation in A of elements in $\mathrm{Hom}(A, B)$ reduces to evaluation in A_1 of elements in $\mathrm{Hom}(A_1, B)$, which we may also do by the previous case. □

Exercise 3.36 (Group exponent)**.** Show that there exists a polynomial-time algorithm that, given a finitely generated abelian group A

(a) and an element $a \in A$, decides whether a is torsion and if so computes $\mathrm{ord}(a)$;

(b) decides whether A is finite and if so computes its exponent and an $a \in A$ with that order.

Exercise 3.37 (Splitting exact sequences)**.** Show that there exists a polynomial-time algorithm that, given morphisms $f : A \to B$ and $g : B \to C$ of finitely generated abelian groups,

(a) decides whether the sequence $0 \to A \xrightarrow{f} B \xrightarrow{g} C \to 0$ is exact;

(b) if so, decides whether the sequence is split;

(c) if so, produces a left-inverse of f and a right-inverse of g.

Hint: Consider the map $g_* : \mathrm{Hom}(C, B) \to \mathrm{Hom}(C, C)$.

Note that by taking $C = 0$ in Exercise 3.37 we may conclude that there exists a polynomial-time algorithm that, given a morphism $f : A \to B$ of finitely generated abelian groups, decides whether f is an isomorphism and if so computes its inverse.

3.6. Structure theorem for finitely generated abelian groups The structure theorem for finitely generated abelian groups is the following.

Theorem 3.38. *Suppose* A *is a finitely generated abelian group. Then there exists a unique sequence* $(r, m, n_1, n_2, \ldots, n_m)$ *of integers with* $r, m \geqslant 0$ *and* $n_1, \ldots, n_m > 1$ *such that* $n_m \mid \cdots \mid n_2 \mid n_1$*, for which*

$$A \cong \mathbb{Z}^r \times \prod_{k=1}^{m} (\mathbb{Z}/n_m\mathbb{Z}).$$

In this section we will prove its algorithmic counterpart.

Exercise 3.39. Let $n > 0$ and $M \subseteq \mathbb{Q}^n$. For subgroups $H \subseteq M$ write $H^\perp = \{x \in M \mid \langle x, H\rangle = 0\}$. Show that for all subgroups $N \subseteq M$ we have $(N^\perp)^\perp = (\mathbb{Q}N) \cap M$. *Hint:* First consider the case where M and N are $\mathbb{Q}$-vector spaces.

Lemma 3.40. *There exists a polynomial-time algorithm that, given a finitely generated abelian group* A*, computes its torsion subgroup.*

Proof. For a subgroup H of $\mathbb{Z}^n$ write $H^\perp = \{x \in \mathbb{Z}^n \mid \langle x, H\rangle = 0\}$, or equivalently for a map $h : H \to \mathbb{Z}^n$ write $h^\perp : H^\perp \to \mathbb{Z}^n$ for the kernel of the map $\mathbb{Z}^n \to \mathrm{Hom}(H, \mathbb{Z})$ given by $x \mapsto (y \mapsto \langle x, h(y)\rangle)$. Note that using Theorem 3.23 we may compute $h^\perp$ in polynomial time. In particular, we may compute $(\alpha^\perp)^\perp : T \to A_1$ for the representative $\alpha : A_0 \to A_1$ of A. It follows from Exercise 3.39 that $(\alpha^\perp)^\perp$ is the torsion subgroup of A: Its image is precisely the set of those elements of A_1 for which a positive integer multiple is in $\alpha(A_0)$. □

Theorem 3.41. *There exists a polynomial-time algorithm that, given a finitely generated abelian group* A*, computes integers* $(r, m, n_1, \ldots, n_m)$ *with* $r, m \geqslant 0$ *and* $n_1, \ldots, n_m > 1$ *such that* $n_m \mid \cdots \mid n_1$ *and computes for* A *an isomorphism*

$$A \cong \mathbb{Z}^r \times \prod_{k=1}^{m} (\mathbb{Z}/n_m\mathbb{Z}),$$

i.e. projections to and inclusions from the individual factors on the right hand side.

Proof. We may compute using Lemma 3.40 the torsion subgroup T of A. Using the image algorithm of Exercise 3.33 we may compute an isomorphism $\mathbb{Z}^r \to A/T$. We have an exact sequence $0 \to T \to A \to A/T \to 0$ which splits, hence by Exercise 3.37 we may compute maps $A \to T$ and $A/T \to A$ such that $A \cong T \times (A/T)$. Replacing A by T we may now assume A is torsion. If $A = 0$ we are done. Using Exercise 3.36 we may compute an element $a \in A$ with order equal to the exponent e of A. Again we have an exact sequence $0 \to \mathbb{Z}a \to A \to A/(\mathbb{Z}a) \to 0$ which is split to which we apply Exercise 3.37. We proceed recursively with A replaced by $A/(\mathbb{Z}a)$. Note that the exponent of $A/(\mathbb{Z}a)$ is a divisor of the exponent of A, so indeed we will get $n_m \mid \cdots \mid n_1$. □

Corollary 3.42. *There exists a polynomial-time algorithm that, given a finitely generated abelian group* A *and a set* S *of integers, computes* $r, m \in \mathbb{Z}_{\geqslant 0}$ *and* $c_1, \ldots, c_m \in \mathbb{Z}_{>1}$

and $n_1, \dots, n_m \in \mathbb{Z}_{>0}$ *such that any two* c_i *are either coprime or a power of the same integer, and every* c_i *either divides some power of an element of* S *or is coprime to all elements of* S*, and computes for* A *an isomorphism*

$$A \cong \mathbb{Z}^r \times \prod_{k=1}^{m} (\mathbb{Z}/c_k\mathbb{Z})^{n_k}.$$

Proof. Apply Theorem 3.41 and compute a coprime basis from $S \cup \{n_1, \dots, n_m\}$ using Theorem 2.7. Write every n_i in terms of this basis and proceed as in Theorem 3.41. □

Exercise 3.43. For a sequence of integers $n_1, \dots, n_k > 0$ we define

$$\operatorname{rex}(n_1, \dots, n_k) = \prod_p p^{\gcd(\operatorname{ord}_p n_1, \dots, \operatorname{ord}_p n_k)}.$$

a. Show that $\operatorname{rad}(n_1 \cdots n_k) \mid \operatorname{rex}(n_1, \dots, n_k)$.

b. Show that there exists a polynomial-time algorithm that computes rex.

For a finite abelian group A isomorphic to $\prod_{i=1}^{k} (\mathbb{Z}/n_i\mathbb{Z})$ write $\operatorname{rex}(A) = \operatorname{rex}(n_1, \dots, n_k)$.

c. Show that $\operatorname{rex}(A)$ is well-defined (i.e. does not depend on the choice of n_i) and can be computed in polynomial time.

4. Computing symbols

In algebraic number theory we find lots of 'symbols'. It is not well-defined what a symbol is, but generally they are maps which encode algebraic properties of its parameters which also satisfies some reciprocity law. In this section we will use [5] as reference. In this section we will define and give algorithms for computing some of these symbols. The father of all symbols is the Legendre symbol.

Definition 4.1. Let p be an odd prime and let a be an integer coprime to p. We define the *Legendre symbol*

$$\left(\frac{a}{p}\right) = \begin{cases} +1 & a \text{ is a square in } \mathbb{Z}/p\mathbb{Z} \\ -1 & \text{otherwise} \end{cases}$$

or equivalently $\left(\frac{a}{p}\right) \equiv a^{\frac{p-1}{2}} \bmod p$.

It is easy to see that we can compute the Legendre symbol directly from the (equivalent) definition in polynomial time using a square-and-multiply algorithm modulo p.

Theorem 4.2 (Quadratic reciprocity)**.** *Suppose* p *and* q *are distinct odd primes. Then*

$$\left(\frac{p}{q}\right)\left(\frac{q}{p}\right) = (-1)^{\frac{p-1}{2} \cdot \frac{q-1}{2}}. \quad \square$$

Definition 4.3. Let b be a positive odd integer and a an integer coprime to b. In terms of the prime factorization $b = p_1^{k_1} \cdots p_n^{k_n}$ of b we define the *Jacobi symbol*

$$\left(\frac{a}{b}\right) = \left(\frac{a}{p_1}\right)^{k_1} \cdots \left(\frac{a}{p_n}\right)^{k_n},$$

where $\left(\frac{a}{p_i}\right)$ is the Legendre symbol defined previously.

Note that the Jacobi symbol extends the Legendre symbol, which justifies using the same notation for both. To compute the Jacobi symbol directly from the definition we need to be able to factor b, which is unfeasible. Quadratic reciprocity for the Jacobi symbol gives us a better method.

Proposition 4.4 (Quadratic reciprocity)**.** *Suppose a and b are coprime positive odd integers. Then*

$$\left(\frac{a}{b}\right)\left(\frac{b}{a}\right) = (-1)^{\frac{a-1}{2}\cdot\frac{b-1}{2}}.$$

Proof. Note that for fixed b the maps $a \mapsto \left(\frac{a}{b}\right)\left(\frac{b}{a}\right)$ and $a \mapsto (-1)^{\frac{a-1}{2}\cdot\frac{b-1}{2}}$ are multiplicative. Hence it suffices to prove the proposition for a prime. Applying the same reasoning to b we have reduced to Theorem 4.2. □

Exercise 4.5. Let p be an odd prime. Show that 2 is a square in $\mathbb{F}_p$ if and only if $p \equiv \pm 1 \bmod 8$. Conclude that $\left(\frac{2}{b}\right) = (-1)^{(b^2-1)/8}$ for all odd $b > 0$. *Hint:* Write $\sqrt{2}$ in terms of 8-th roots of unity.

Exercise 4.6. Let $x_0, x_1, x_2, x_3 \in \mathbb{Z}$ such that x_1 is odd and $x_{n+2} \equiv x_n \bmod x_{n+1}$ for all n. Show that when the expressions are well-defined we have the following equalities.

$$\textbf{a.}\ \left(\frac{x_0}{x_1}\right) = (-1)^{(x_1-1)(x_2-1)/4}\cdot\left(\frac{x_1}{x_2}\right) \quad \text{if } x_2 \equiv 1 \bmod 2,$$

$$\textbf{b.}\ \left(\frac{x_0}{x_1}\right) = \left(\frac{x_2}{x_3}\right) \quad \text{if } x_2 \equiv 0 \bmod 4$$

$$\text{and}\quad \textbf{c.}\ \left(\frac{x_0}{x_1}\right) = (-1)^{(x_1x_3-1)(x_1x_3+x_2-1)/8}\cdot\left(\frac{x_2}{x_3}\right) \quad \text{if } x_2 \equiv 2 \bmod 4$$

Theorem 4.7. *There exists a polynomial-time algorithm that, given coprime positive odd integers a and b, computes the Jacobi symbol $\left(\frac{a}{b}\right)$.*

Proof. We define the *gcd sequence* of positive integers x_0 and x_1 to be the sequence $(x_0, \ldots, x_N)$ where x_{n+2} is the unique integer such that $0 \leqslant x_{n+2} < x_{n+1}$ and $x_{n+2} \equiv x_n \bmod x_{n+1}$ and with $x_N = 0$. The proof of the Euclidean algorithm shows that this sequence contains only linearly many elements and can be computed in polynomial time. Moreover, $x_{N-1} = \gcd(x_0, x_1)$.

Compute the gcd sequence of a and b. As $\gcd(x_n, x_{n+1}) = \gcd(a, b) = 1$ for all $n < N$, the symbols $\left(\frac{x_n}{x_{n+1}}\right)$ are defined. Using Exercise 4.6 we may express $\left(\frac{a}{b}\right) = s_n \cdot \left(\frac{x_n}{x_{n+1}}\right)$ for some $s_n \in \{\pm 1\}$ iteratively for n with x_{n+1} odd. Clearly we may compute these s_n in polynomial time. As $\left(\frac{x_{N-2}}{x_{N-1}}\right) = \left(\frac{x_{N-2}}{1}\right) = 1$, we simply return s_{N-1}. □

Exercise 4.8. The Euclidean algorithm implied by Exercise 2.12 produces a different type of 'gcd sequence' than those used in Theorem 4.7, namely those where $|x_{n+2}| \leqslant |x_{n+1}|/2$ and $x_{n+2} \equiv x_n \bmod x_{n+1}$ for all n. Give a proof of Theorem 4.7 using such gcd sequences.

The Jacobi symbol is defined on a subset of the integers. As is the theme in this document, we will 'extend' the Jacobi symbol to number rings.

Remark 4.9. One is sometimes interested in a more general Jacobi symbol $\left(\frac{a}{b}\right)$ where b is not necessarily odd or positive, by defining

$$\left(\frac{a}{2}\right) = (-1)^{\frac{a^2-1}{8}} \quad \text{and} \quad \left(\frac{a}{-1}\right) = \frac{a}{|a|} = \begin{cases} +1 & \text{if } a > 0 \\ -1 & \text{if } a < 0 \end{cases}.$$

This is called the *Kronecker symbol*. The Kronecker symbol can be computed in polynomial time by writing $b = uc2^k$ where $u = \pm 1$ and c is odd and positive, and applying Theorem 4.7 to the factor $\left(\frac{a}{c}\right)$ in $\left(\frac{a}{b}\right) = \left(\frac{a}{u}\right)\left(\frac{a}{c}\right)\left(\frac{a}{2}\right)^k$.

4.1. Jacobi symbols in number rings First we define the Legendre symbol for a general number ring.

Definition 4.10. Let $\mathfrak{p}$ be a prime ideal in a number ring R of odd index $n_\mathfrak{p} = (R : \mathfrak{p})$ and let $a \in R$ such that $aR + \mathfrak{p} = R$. We define the *Legendre symbol*

$$\left(\frac{a}{\mathfrak{p}}\right) = \begin{cases} +1 & \text{if } a \text{ is a square in } R/\mathfrak{p} \\ -1 & \text{otherwise} \end{cases},$$

or equivalently $\left(\frac{a}{\mathfrak{p}}\right) = a^{\frac{n_\mathfrak{p}-1}{2}} \bmod \mathfrak{p}$.

Extending the definition to general ideals as for the Jacobi symbol cannot be done similarly, unless R is a Dedekind domain, because it would require prime factorization of ideals. Instead, we consider the following.

Definition 4.11. Suppose $\mathfrak{b}$ is an ideal of a number ring R. For a prime $\mathfrak{p}$ we define $l_\mathfrak{p}(\mathfrak{b})$ such that $(R_\mathfrak{p} : \mathfrak{b}_\mathfrak{p}) = (R : \mathfrak{p})^{l_\mathfrak{p}(\mathfrak{b})}$. For $(R : \mathfrak{b})$ odd and $a \in R$ with $aR + \mathfrak{b} = R$ we define the *Jacobi symbol* by

$$\left(\frac{a}{\mathfrak{b}}\right) = \prod_{\mathfrak{p} \in \operatorname{spec} R} \left(\frac{a}{\mathfrak{p}}\right)^{l_\mathfrak{p}(\mathfrak{b})}.$$

Theorem 4.12. *There exists a polynomial-time algorithm that, given an order R, an ideal $\mathfrak{b}$ of R such that $(R : \mathfrak{b})$ is odd and $a \in R$ such that $aR + \mathfrak{b} = R$, computes the Jacobi symbol $\left(\frac{a}{\mathfrak{b}}\right)$.*

We will prove this theorem by expressing the Jacobi symbol in terms of yet another symbol.

Exercise 4.13. Let $\mathfrak{p}$ and $\mathfrak{b}$ be ideals in a number ring R with $\mathfrak{p}$ prime and take any composition series $0 = M_0 \subsetneq M_1 \subsetneq \cdots \subsetneq M_n = R/\mathfrak{b}$ of $R/\mathfrak{b}$ as R-module. Show that $l_\mathfrak{p}(\mathfrak{b})$ equals the number of quotients M_{i+1}/M_i that is isomorphic to $R/\mathfrak{p}$ as an R-module.

4.2. Signs of permutations In this section we will consider the following symbol.

Definition 4.14. Let B be a finite abelian group and let $\sigma \in \operatorname{Aut}(B)$. We define (σ, B) to be the sign of σ as element of the permutation group on B.

Lemma 4.15. *Suppose A and C are sets and $\alpha \in \operatorname{Aut}(A)$ and $\gamma \in \operatorname{Aut}(C)$ are permutations. Write $\alpha \sqcup \gamma$ respectively $\alpha \times \gamma$ for the induced permutation on the disjoint union $A \sqcup C$ and product $A \times C$. Then $\operatorname{sgn}(\alpha \sqcup \gamma) = \operatorname{sgn}(\alpha) \cdot \operatorname{sgn}(\gamma)$ and $\operatorname{sgn}(\alpha \times \gamma) = \operatorname{sgn}(\alpha)^{\#C} \cdot \operatorname{sgn}(\gamma)^{\#A}$.*

Proof. That $\operatorname{sgn}(\alpha \sqcup \gamma) = \operatorname{sgn}(\alpha) \cdot \operatorname{sgn}(\gamma)$ follows from the fact that $\alpha \sqcup \gamma = \alpha\gamma$ when $\operatorname{Aut}(A)$ and $\operatorname{Aut}(C)$ are naturally mapped to $\operatorname{Aut}(A \sqcup C)$.

For the second part write $\alpha' = \alpha \times \mathrm{id}_C$ and $\gamma' = \mathrm{id}_A \times \gamma$ and note that $\alpha \times \gamma = \alpha' \cdot \gamma'$. Now α' acts as α on $\#C$ disjoint copies of A, hence by the previous $\operatorname{sgn}(\alpha') = \operatorname{sgn}(\alpha)^{\#C}$. Mutatis mutandis we obtain the same for γ', and the lemma follows from multiplicativity of the sign. □

Proposition 4.16. *Suppose B is a finite abelian group and $\beta \in \operatorname{Aut}(B)$. Suppose we have an exact sequence $0 \to A \to B \to C \to 0$ such that β restricts to an automorphism α of A. Then β induces an automorphism γ of C such that the following diagram commutes*

$$\begin{array}{ccccccccc} 0 & \longrightarrow & A & \xrightarrow{f} & B & \xrightarrow{g} & C & \longrightarrow & 0 \\ & & \downarrow{\scriptstyle\alpha} & & \downarrow{\scriptstyle\beta} & & \downarrow{\scriptstyle\gamma} & & \\ 0 & \longrightarrow & A & \xrightarrow{f} & B & \xrightarrow{g} & C & \longrightarrow & 0 \end{array}$$

and if $\#C$ is odd we have $(\beta, B) = (\alpha, A) \cdot (\gamma, C)^{\#A}$.

Proof. The map γ exists by a diagram chasing argument. Since $\#C$ is odd we may write $C = \{0\} \sqcup D \sqcup (-D)$ for some subset $D \subseteq C$. Choosing any section $D \to B$ of g (which is not a group homomorphism!), we may extend it to a section $h : C \to B$ in such a way that $h(-c) = -h(c)$. Now the maps f and h together give a bijection of sets $A \times C \to B$ and let β' be the induced action of β on $A \times C$. By Lemma 4.15 we have that $(\alpha \times \gamma, A \times C) = (\alpha, A)^{\#C} \cdot (\gamma, C)^{\#A}$, hence to prove the proposition it suffices to show that $\sigma = (\alpha \times \gamma)^{-1} \cdot \beta'$ is an even permutation. For all $d \in D$ the action of σ restricts to $A \times \{d\}$. Note that σ commutes with -1, hence the action of σ on $A \times \{-d\}$ is isomorphic to the action on $A \times \{d\}$. Hence the restriction of σ to $A \times (C \setminus \{0\})$ is even. Finally, σ is the identity on $A \times \{0\}$, so we conclude σ is even, as was to be shown. □

The exact sequence $0 \to 2\mathbb{Z}/4\mathbb{Z} \to \mathbb{Z}/4\mathbb{Z} \to \mathbb{Z}/2\mathbb{Z} \to 0$ resists application of Proposition 4.16. If we choose the non-trivial automorphism σ given by $x \mapsto -x$ on $\mathbb{Z}/4\mathbb{Z}$ we see that its sign is -1, while the induced maps on the other terms are trivial. Hence $(\sigma, \mathbb{Z}/4\mathbb{Z})$ is not an $\mathbb{F}_2$-linear combination of $(\sigma, 2\mathbb{Z}/4\mathbb{Z})$ and $(\sigma, \mathbb{Z}/2\mathbb{Z})$.

Exercise 4.17. Show that for $k \in \mathbb{Z}_{\geqslant 2}$ and $a \in (\mathbb{Z}/2^k\mathbb{Z})^*$ we have

$$(x \mapsto ax, \mathbb{Z}/2^k\mathbb{Z}) = (-1)^{\frac{a-1}{2}}.$$

Hint: Write $\mathbb{Z}/2^k\mathbb{Z} = (\mathbb{Z}/2^k\mathbb{Z})^* \sqcup (2\mathbb{Z}/2^k\mathbb{Z})$ and show $(x \mapsto ax, (\mathbb{Z}/2^k\mathbb{Z})^*) = -1$ if and only if a generates $(\mathbb{Z}/2^k\mathbb{Z})^*$.

Exercise 4.18. Suppose B is an abelian group. For $b \in B$ write $\lambda_b : B \to B$ for the map $x \mapsto x + b$.

a. Show that $\operatorname{sgn}(\lambda_b) = -1$ if and only if $(B : \langle b \rangle)$ is odd and $2 \mid \#B$.

Suppose B has order 2^k for some $k \geqslant 1$ and let $\beta \in \operatorname{Aut}(B)$.

b. Show that there exists a subgroup $A \subseteq B$ such that $\beta(A) = A$ and $(B : A) = 2$.

c. Let $b \in B$ such that $B = A \cup (b + A)$. Show that $(\beta, B) = -1$ if and only if $A = \langle \beta(b) - b \rangle$.

d. Suppose $(\beta, B) = -1$. Show that $B = \mathbb{Z}/2^k\mathbb{Z}$ or $B = (\mathbb{Z}/2\mathbb{Z})^2$.

e. Show that there exists a polynomial-time algorithm that, given a finite abelian group B and $\beta \in \operatorname{Aut}(B)$ such that $2 \mid \#B$, computes (β, B).

4.3. Computing signs of group automorphisms In this section we will prove we can compute the sign of group automorphisms in polynomial time. We will need an elementary lemma about determinants that mirrors Proposition 4.16.

Lemma 4.19. *Let $\mathbb{F}$ be a field and let $0 \to A \to B \to C \to 0$ be an exact sequence of finite dimensional $\mathbb{F}$-vector spaces together with automorphism α, β and γ such that the diagram*

$$\begin{array}{ccccccccc} 0 & \longrightarrow & A & \longrightarrow & B & \longrightarrow & C & \longrightarrow & 0 \\ & & \downarrow{\scriptstyle\alpha} & & \downarrow{\scriptstyle\beta} & & \downarrow{\scriptstyle\gamma} & & \\ 0 & \longrightarrow & A & \longrightarrow & B & \longrightarrow & C & \longrightarrow & 0 \end{array}$$

commutes. Then $\det(\beta) = \det(\alpha) \cdot \det(\gamma)$. □

In terms of matrices, the above lemma simply states that for square matrices A and C and a matrix P that fits, the block matrix $B = \left(\begin{smallmatrix} A & P \\ 0 & C \end{smallmatrix}\right)$ satisfies $\det(B) = \det(A) \cdot \det(C)$.

Theorem 4.20. *Suppose $b \in \mathbb{Z}_{>0}$ is odd and B is a free $(\mathbb{Z}/b\mathbb{Z})$-module of finite rank. Then for all $\sigma \in \operatorname{Aut}(B)$ we have $(\sigma, B) = \left(\frac{\det(\sigma)}{b}\right)$.*

Proof. For $B = 0$ the theorem clearly holds, so assume $b \neq 1$ and that B has rank at least 1.

First suppose b is prime and B has rank 1. Then σ is given by multiplication with $a \in (\mathbb{Z}/b\mathbb{Z})^*$. If a generates $(\mathbb{Z}/b\mathbb{Z})^*$, then the corresponding permutation fixes 0 and acts transitively on the $b-1$ remaining elements of B, so that $(\sigma, B) = -1 = \left(\frac{a}{b}\right) = \left(\frac{\det(\sigma)}{b}\right)$. By multiplicativity of both symbols in σ, this also proves the case where a is not a generator.

Now we prove using induction the case for general rank of B. Suppose σ is given by an upper or lower triangular matrix. Then there exists a subspace $0 \subsetneq A \subseteq B$ such that σ restricts to A. Hence we have a split exact sequence $0 \to A \to B \to C \to 0$ with $C = B/A$ and let α and γ be the induced maps on A respectively C. Then by Proposition 4.16, the induction hypothesis and

Lemma 4.19 we get

$$(\sigma, B) = (\alpha, A) \cdot (\gamma, C) = \left(\frac{\det(\alpha)}{b}\right) \cdot \left(\frac{\det(\gamma)}{b}\right) = \left(\frac{\det(\alpha)\det(\gamma)}{b}\right) = \left(\frac{\det(\sigma)}{b}\right).$$

Since every matrix can be written as a product of upper and lower triangular matrices, the case for general α follows.

Now we prove the theorem for general b with induction to the number of divisors of b. We have just proven the induction base with b prime. For b not prime we may take a divisor $1 < d < b$ of b. Let $A = dB$ and $C = B/A$ and note that they are free modules over $\mathbb{Z}/\frac{b}{d}\mathbb{Z}$ respectively $\mathbb{Z}/d\mathbb{Z}$. Moreover, σ induces maps α and γ on A respectively C that make the usual diagram commute. It follows from the definition of the determinant that $\det(\alpha) \equiv \det(\sigma) \bmod \frac{b}{d}$ and $\det(\gamma) \equiv \det(\sigma) \bmod d$. Then

$$(\sigma, B) = (\alpha, A)(\gamma, C) = \left(\frac{\det(\alpha)}{b/d}\right)\left(\frac{\det(\gamma)}{d}\right) = \left(\frac{\det(\sigma)}{b/d}\right)\left(\frac{\det(\sigma)}{d}\right) = \left(\frac{\det(\sigma)}{b}\right).$$

The theorem now follows by induction. □

Theorem 4.21. *There exists a polynomial-time algorithm that, given an finite abelian group B and an automorphism σ of B, computes the symbol (σ, B).*

Proof. If $2 \mid \#B$ we have Exercise 4.18, so suppose B has odd order. Using Theorem 3.41, write B as a product $\prod_{k=1}^{m}(\mathbb{Z}/n_k\mathbb{Z})$ of non-trivial cyclic groups such that $n_j \mid n_k$ for all $j > k$. Note that B fits in an exact sequence $0 \to A \to B \to C \to 0$ with $A = n_m B$ and $C = B/A$, such that σ restricts to A and C. Then

$$A = \prod_{k=1}^{m}(n_m\mathbb{Z}/n_k\mathbb{Z}) \cong \prod_{k=1}^{m-1}\left(\mathbb{Z}\Big/\frac{n_k}{n_m}\mathbb{Z}\right) \quad \text{and} \quad C \cong (\mathbb{Z}/n_m\mathbb{Z})^m.$$

Since C is a free $(\mathbb{Z}/n_m\mathbb{Z})$-module, we may compute (σ, C) using Theorem 4.20 in polynomial time. Note that A is a product of strictly less cyclic groups than B, as well as having smaller order. While $A \neq 0$ we compute (σ, A) recursively and apply Proposition 4.16 to compute (σ, B). Since m is polynomially bounded in the length of the input, there is only polynomially many recursive steps and the algorithm runs in polynomial time. □

4.4. Computing Jacobi symbols in number rings To compute Jacobi symbols in polynomial time it now suffices to reduce to Theorem 4.21.

Lemma 4.22. *Suppose $\mathfrak{b}$ is an ideal in a number ring R odd index $(R : \mathfrak{b})$, and suppose $a \in R$ satisfies $aR + \mathfrak{b} = R$. Then $\left(\frac{a}{\mathfrak{b}}\right) = (\alpha, R/\mathfrak{b})$ where the map α is multiplication by a.*

Proof. First suppose $\mathfrak{b}$ is a prime ideal and suppose a generates $(R/\mathfrak{b})^*$. Then $\left(\frac{a}{\mathfrak{b}}\right) = -1$. As a acts transitively on an even number of elements $(R/\mathfrak{b}) \setminus \{0\}$, we conclude that $(x \mapsto ax, R/\mathfrak{b}) = -1 = \left(\frac{a}{\mathfrak{b}}\right)$. The case for general a follows from multiplicativity of both symbols.

Now consider the case of general $\mathfrak{b}$. Choose some composition series $0 = M_0 \subsetneq M_1 \subsetneq \cdots \subsetneq M_n = R/\mathfrak{b}$ of $R/\mathfrak{b}$ as R-module. Consider the exact sequence

$0 \to M_i \to M_{i+1} \to M_{i+1}/M_i \to 0$ and note that $M_{i+1}/M_i \cong R/\mathfrak{p}$ as R-modules for some prime ideal $\mathfrak{p}_i$ of R. By applying Proposition 4.16 inductively we obtain $(\alpha, B) = \prod_{i=1}^{n}(\alpha, M_{i+1}/M_i) = \prod_{i=1}^{n}\left(\frac{\alpha}{\mathfrak{p}_i}\right)$. It follows from Exercise 4.13 that the latter equals $\left(\frac{\alpha}{\mathfrak{b}}\right)$. □

Proof of Theorem 4.12. Compute $R/\mathfrak{b}$ and the map $\alpha : R/\mathfrak{b} \to R/\mathfrak{b}$ given by $x \mapsto ax$. Using Theorem 4.21 compute $(\alpha, R/\mathfrak{b})$, which equals $\left(\frac{a}{\mathfrak{b}}\right)$ by Lemma 4.22. □

5. Finite commutative rings

Often problems in algebraic number theory can be reduced to a problem concerning finite commutative rings. In this section we will state theorems that show a relation between computational questions about rings of integers and finite commutative rings, namely that they are 'equally hard' in some precise sense. Some important examples of finite commutative rings are the following:

- the ring $\mathbb{Z}/m\mathbb{Z}$, for some $m \in \mathbb{Z}_{>0}$;
- the field $\mathbb{F}_q$ of cardinality q, for some prime-power q;
- the ring R/I, where R is a number ring and I is a non-zero R-ideal;
- the ring $\mathbb{F}_q[X]/(g)$, where $g \in \mathbb{F}_q[X]$ is a non-zero polynomial;
- the group ring $\mathbb{F}_q[G]$, where G is a finite abelian group.

To apply algorithms to finite commutative rings, we will first need to be able to encode them. We encode them as finitely generated abelian groups A, together with a multiplication morphism $x \mapsto ax$ for every generator a of A.

Theorem 5.1 (Structure Theorem for Finite Commutative Rings)**.** *Let A be a finite commutative ring. Then for all $n \in \mathbb{Z}_{>0}$ sufficiently large, the natural map $A \to \prod_{\mathfrak{p} \in \operatorname{spec} A} A/\mathfrak{p}^n$ given by $a \mapsto (a + \mathfrak{p}^n)_{\mathfrak{p}}$ is an isomorphism of rings. Hence, A is isomorphic to a finite product of local rings with nilpotent maximal ideals.*

Proof. By Exercise 1.10, Exercise 1.11 and Exercise 1.12 we have

$$A \supseteq \operatorname{nil}(A) = \bigcap_{\mathfrak{p} \in \operatorname{spec} A} \mathfrak{p} = \prod_{\mathfrak{p} \in \operatorname{spec} A} \mathfrak{p}.$$

It is clear that $\operatorname{nil}(A)$ is finitely generated, and thus $\operatorname{nil}(A)$ is nilpotent by Exercise 1.10. Hence, we have $\prod_{\mathfrak{m}} \mathfrak{m}^n = 0$ for all n sufficiently large. By the Chinese remainder theorem (Exercise 1.13) and Exercise 1.12 we get

$$A \cong A/\prod_{\mathfrak{p}} \mathfrak{p}^n \cong \prod_{\mathfrak{p}} A/\mathfrak{p}^n.$$

Notice that for each $\mathfrak{p} \in \operatorname{spec} A$ the unique maximal ideal $\mathfrak{p}/\mathfrak{p}^n$ of $A/\mathfrak{p}^n$ is nilpotent. □

Proposition 5.2. *There exists a polynomial-time algorithm that, given a finite commutative ring A and a sequence of coprime integers $m_1, \ldots, m_n$ such that* $\exp(A^+) \mid$

$\mathfrak{m}_1 \cdots \mathfrak{m}_n$, *computes for* A *a ring isomorphism*

$$A \cong \prod_{i=1}^{n} (A/\mathfrak{m}_i A).$$

Proof. Note that $\bigcap_{i=1}^{n}(\mathfrak{m}_i A) = \prod_{i=1}^{n}(\mathfrak{m}_i A) \subseteq \exp(A^+) \cdot A = 0$ since the ideals $\mathfrak{n}_i A$ are coprime. Then by the Chinese remainder theorem (Exercise 1.13) the natural isomorphism of abelian groups $A \cong \prod_{i=1}^{n}(A/\mathfrak{m}_i A)$ is a ring homomorphism. We may compute this isomorphism using Corollary 3.42. □

We may apply Theorem 5.1 to the ring $\mathbb{Z}/m\mathbb{Z}$ for some $m \in \mathbb{Z}_{>0}$. We have $\operatorname{spec}(\mathbb{Z}/m\mathbb{Z}) = \{p\mathbb{Z}/m\mathbb{Z} \mid p \text{ prime}, p|m\}$ and get a decomposition $\mathbb{Z}/m\mathbb{Z} \cong \prod_p (\mathbb{Z}/p^{k_p}\mathbb{Z})$ as rings, where p ranges over the primes and m_p is the order of p in m. Since we have no algorithm to factor integers, this also implies that we have no polynomial-time algorithm to compute decompositions as in Theorem 5.1. However, it is possible to decide whether an integer is a prime number in polynomial-time using an AKS primality test [1] [7]. Similarly, we can decide whether a finite commutative ring is a local ring (see Exercise 5.6). We are able to trivially factor prime powers into a product of primes by taking roots, and similarly we get the following for local rings.

Theorem 5.3. *There is a polynomial-time algorithm that, given a finite commutative ring* A *with* $\#A$ *a prime power, computes* $\operatorname{nil}(A)$.

Lemma 5.4. *Let* A *be a commutative ring and let* $n, k \in \mathbb{Z}_{\geqslant 1}$. *Then* $\operatorname{nil}(A/n^k A)$ *is the inverse image of* $\operatorname{nil}(A/nA)$ *under the natural map* $A/n^k A \to A/nA$.

Proof. We have $(nA)^k \equiv 0 \bmod n^k A$, so $nA \subseteq \operatorname{nil}(A/n^k A)$. □

Algorithm 5.5. Let A be a finite commutative ring of prime power order.

(1) Compute $\#A$ and find p prime and $n \in \mathbb{Z}_{\geqslant 0}$ such that $\#A = p^n$.
(2) Compute $B = A/pA$, which is an $\mathbb{F}_p$-algebra, and the least $t \in \mathbb{Z}_{\geqslant 0}$ such that $\dim_{\mathbb{F}_p} B < p^t$.
(3) Let $F : B \to B$ be the Frobenius map $b \mapsto b^p$ and compute $K = \ker(F^t)$.
(4) Return the inverse image of K under the natural map $\pi : A \to B$.

Proof of Theorem 5.3. We will show Algorithm 5.5 satisfies our requirements. Note that all steps in the algorithm involve only basic computations as described in Section 1.2 and Section 3. Most importantly, we use that F and thus F^t is a linear map, to compute K. Hence the algorithm runs in polynomial time. We will thus show it is correct. For each $x \in \operatorname{nil}(B)$ and $m \geqslant \#A$ we have $x^m = 0$, so $x \in \operatorname{nil}(B)$ if and only if $F^t(x) = 0$. Then $K = \operatorname{nil}(A)$ by Lemma 5.4, so the algorithm is correct. □

Exercise 5.6. Suppose there exists polynomial-time primality test.

a. Show that a finite commutative ring R is local if and only $\operatorname{nil}(R)$ is a maximal ideal.

b. Show that a finite reduced commutative ring R is a field if and only if $\#R = p^k$ for some prime p and integer $k > 0$ and for no $0 < m < k$ the homomorphism $x \mapsto x^{p^m}$ is the identity map.

c. Conclude that there exist a polynomial-time algorithm that, given a finite commutative ring R, decides whether R is local and if so computes its maximal ideal.

5.1. Polynomial-time reductions Intuitively, a problem f is easier than a problem g if you can solve f once you know how to solve g, or more specific to algorithms, when you can solve them in polynomial time. We can make this more formal.

Definition 5.7. Let f and g be two problems. We say f *can be reduced to* g, or symbolically $f \preceq g$, if there exists an algorithm X for f that takes as additional input and algorithm Y for g which X may use as a sub-algorithm, and such that X spends only polynomial time not running Y. We say f and g are *equally hard*, or symbolically $f \approx g$, if $f \preceq g \preceq f$.

Note that if there exists a polynomial-time algorithm for f, then we always have $f \preceq g$ for any problem g. This shows the concept of reductions is not very useful for problems for which we know a polynomial-time algorithm exists.

Example 5.8. An example of a reduction is given after Theorem 5.1: We can reduce factoring integers to decomposing finite commutative rings into a product of local rings with nilpotent maximal ideals. We can construct our algorithm X as in Definition 5.7 as follows. For an integer input m we call our algorithm Y for decomposing finite commutative rings on $\mathbb{Z}/m\mathbb{Z}$. The definition of a reduction then grants us polynomial-time to transform the output $(A_i)_{i\in I}$ of Y satisfying $\mathbb{Z}/m\mathbb{Z} \cong \prod_{i\in I} A_i$, into a prime factorization of m. As noted before, we only have to find primes p_i and integers m_i in polynomial time such that $\#A_i = p_i^{m_i}$, which we can do by naively taking roots.

The *radical* or *square-free part* of a positive integer m, written $\mathrm{rad}(m)$, is the product of the prime divisors of m. We say m is *square-free* if $\mathrm{rad}(m) = m$. We call $\mathrm{rad}(m)$ the radical because $\mathrm{rad}(m)\mathbb{Z} = \sqrt{m\mathbb{Z}}$. We say a commutative ring R is *reduced* if $\mathrm{nil}(R) = 0$. We may now state the two main theorems of this section, which we prove in Section 5.4.

Theorem 5.9. *The following three computational problems are equally hard:*

(1) *Given a number field* K*, compute* $\mathcal{O}_K$*.*
(2) *Given a finite ring* A*, compute* $\mathrm{nil}(A)$*.*
(3) *Given an integer* $m > 0$*, compute* $\mathrm{rad}(m)$*.*

Each of the problems in Theorem 5.9 has an associated decision problem, which we will show are all equally hard as well.

Theorem 5.10. *The following three decision problems are equally hard:*

(1) *Given an order* R *in a number field* K*, decide whether* $R = \mathcal{O}_K$.
(2) *Given a finite commutative ring* A*, decide whether* A *is reduced.*
(3) *Given an integer* $m > 0$*, decide whether* m *is square-free.*

As of July 21, 2022, for none of the six above problems there are known polynomial-time algorithms. Note that each decision problem trivially reduces to its corresponding computational problem.

5.2. Trace radicals Let $A \subseteq B$ be commutative rings such that B is *free of finite rank* over A, meaning that $B \cong A^n$ as A-modules for some $n \in \mathbb{Z}_{\geqslant 0}$. Then the ring $\mathrm{End}_A(B)$ of A-module endomorphisms of B is isomorphic to $\mathrm{Mat}_n(A)$, the ring of $n \times n$-matrices with coefficients in A. This isomorphism induces a determinant $\det : \mathrm{End}_A(B) \to A$ and trace $\mathrm{Tr} : \mathrm{End}_A(B) \to A$ on $\mathrm{End}_A(B)$ which are multiplicative respectively A-linear. We have a natural map $B \to \mathrm{End}_A(B)$ given by $b \mapsto (x \mapsto bx)$ which induces the A-module homomorphism $\mathrm{Tr}_{B/A} : B \to A$ when composed with Tr, which we call the *trace of* B *over* A.

Exercise 5.11. Show that $\mathrm{Tr} : \mathrm{End}_A(B) \to A$ and $\det : \mathrm{End}_A(B) \to A$ are independent of the chosen A-basis of B.

Exercise 5.12. Suppose B is a free A-module and let n and m be coprime integers such that $nm = \#A$. Then $B \cong (B/nB) \times (B/mB)$ and B/nB is a free A/nA-module. Show that Tr and det commute with this isomorphism when B has finite rank.

Definition 5.13. Let $A \subseteq B$ be commutative rings such that B is free over A as A-module. Then we define the *trace radical* of B over A as

$$\mathrm{Trad}(B/A) = \{x \in B \mid \mathrm{Tr}_{B/A}(xB) = 0\} = \ker(x \mapsto (y \mapsto \mathrm{Tr}_{B/A}(xy))).$$

It follows trivially that $\mathrm{Trad}(B/A)$ is an ideal of B. In fact, it is the largest ideal contained in $\ker(\mathrm{Tr}_{B/A})$. We may compute $\mathrm{Trad}(B/A)$ in polynomial time using linear algebra over the integers. As the following exercise in linear algebra shows, the trace function can in some sense test nilpotency. Similarly, we hope to compute $\mathrm{nil}(B)$ from $\mathrm{Trad}(B/A)$.

Exercise 5.14. Let A be a field and let B be a finite dimensional A-vector space. Suppose A has characteristic 0 or p for some prime $p > \dim_A(B)$. Show that $M \in \mathrm{End}_A(B)$ is nilpotent if and only if $\mathrm{Tr}(M^n) = 0$ for all $n \in \mathbb{Z}_{>0}$. Conclude that $\mathrm{Trad}(B/A) = \mathrm{nil}(B)$ when B is an A-algebra. What happens when we replace A by a local ring and B by a free A-module?
Hint: Consider the characteristic polynomial of M and use Newton's identities.

Proposition 5.15. *Let* m *be a square-free integer and a free module* B *over* $A = \mathbb{Z}/m\mathbb{Z}$*. If* $p > \mathrm{rk}_A(B)$ *for all primes* p *dividing* m*, then* $\mathrm{Trad}(B/A) = \mathrm{nil}(B)$.

Proof. We have $A \cong \prod_{p|m}(\mathbb{Z}/p\mathbb{Z})$ and $B \cong \prod_{p|m}(B/pB)$, and without loss of generality we assume this is an equality. Clearly $\mathrm{nil}(B) = \prod_{p|m} \mathrm{nil}(B/pB)$, and from Exercise 5.12 it follows that also $\mathrm{Trad}(B/A) = \prod_{p|m} \mathrm{Trad}((B/pB)/(A/pA))$. Note that $p > \mathrm{rk}_A(B) \geqslant \mathrm{rk}_{A/pA}(B/pB)$ for all $p \mid m$. Hence it suffices to prove the proposition for m prime, which is Exercise 5.14. □

With this proposition, we may prove the following.

Theorem 5.16. *There exists a polynomial-time algorithm that, given a finite commutative ring* A *and the integer* $\mathrm{rad}(\#A)$, *computes* $\mathrm{nil}(A)$.

Algorithm 5.17. Let A be a finite commutative ring and let $r = \mathrm{rad}(\#A)$ and $l = \log_2(\#A)$.

(1) Apply Algorithm 2.8 to $\{r, \#A\} \cup \{p \mid p \leqslant l, p \text{ prime}\}$ to compute a coprime basis.
(2) Factor $\#A = \prod_{i=1}^{n} c_i^{k_i}$ uniquely with the c_i in the coprime basis and $k_i \geqslant 1$.
(3) Apply Proposition 5.2 to compute the natural surjection $A \to \prod_{i=1}^{n}(A/c_iA)$.
(4) For each c_i compute $\mathrm{nil}(A/c_iA)$ as follows:
    (a) If $c_i \leqslant l$ is a prime, then apply Algorithm 5.5.
    (b) Otherwise, compute $\mathrm{nil}(A/c_iA)$ as $\mathrm{Trad}((A/c_iA)/(\mathbb{Z}/c_i\mathbb{Z}))$ using linear algebra.
(5) Return the inverse image of $\prod_{i=1}^{n} \mathrm{nil}(A/c_iA)$ under the map $A \to \prod_{i=1}^{n}(A/c_iA)$.

Proof of Theorem 5.16. Clearly Algorithm 5.17 runs in polynomial time if it is correct. Because r is square-free, the coprime basis consists of square-free integers. In particular all c_i are square-free. Moreover, every c_i is either a prime at most l, or all primes dividing it are greater than l. In the first case, Algorithm 5.5 correctly computes $\mathrm{nil}(A/c_iA)$ by Theorem 5.3. In the second case, A/c_iA over $\mathbb{Z}/c_i\mathbb{Z}$ indeed satisfies the conditions to Proposition 5.15. By Lemma 5.4 we indeed compute $\mathrm{nil}(A)$ in the final step. □

5.3. Discriminants We now move from the trace radical to the discriminant. Let $A \subseteq B$ be commutative rings such that B is free over A of rank n. Then the *discriminant* of B over A is $\Delta_{B/A} = \det(M) \in A$, where $M = (\mathrm{Tr}_{B/A}(e_ie_j))_{1 \leqslant i,j \leqslant n}$ is a matrix and $(e_1, \ldots, e_n)$ is a basis for B over A. This definition depends on a choice of basis for B (see Exercise 5.18). However, for $A = \mathbb{Z}$ the discriminant is uniquely determined, and we write $\Delta(B) = \Delta_{B/\mathbb{Z}}$ in this case. If $C \subseteq B$ is of finite index and both are free over $\mathbb{Z}$, then $|\Delta(B)| = (B : C)^2 \cdot |\Delta(C)|$. In this section we will use the following two elementary facts from number theory regarding discriminants.

Exercise 5.18. Prove that $\Delta_{B/A}$ is defined up to multiples of $(A^*)^2$. Conclude that $\Delta_{B/A}$ is uniquely defined for $A = \mathbb{Z}$.

Exercise 5.19. Suppose $d \in \mathbb{Z}$ is non-zero and let $K = \mathbb{Q}(\sqrt{d})$.

(1) Show that if $d \equiv 1 \bmod 4$ and d is square-free, then $\mathcal{O}_K = \mathbb{Z}[(\sqrt{d}+1)/2]$.
(2) Show that if $d \equiv 2, 3 \bmod 4$ and d is square-free, then $\mathcal{O}_K = \mathbb{Z}[\sqrt{d}]$.

(3) Show that we can write $d = \square(d)^2 \cdot \boxtimes(d)$ uniquely with $\square(d) \in \mathbb{Z}_{>0}$ and $\boxtimes(d)$ square-free.
(4) Conclude that $\Delta(\mathcal{O}_K)/\boxtimes(d) \in \{1, 4\}$.

Exercise 5.20. Let R be an order of a number field K and let p be a prime number. Show that $p \mid (\mathcal{O}_K : R)$ if and only if R has a singular (i.e. non-invertible) prime ideal containing p.

Proposition 5.21. *Let R be an order of a number field K and let $a \in \mathbb{Z}_{>1}$. Let $\mathfrak{a}$ be the product of all prime ideals of R containing a. Then the following are equivalent:*
(1) The integer a is coprime to $(\mathcal{O}_K : R)$;
(2) The ideal $\mathfrak{a}$ is invertible;
(3) We have a strict inclusion $R \subsetneq \mathfrak{a} : \mathfrak{a}$.

Proof. That (1) and (2) are equivalent follows from Exercise 5.20 and Exercise 1.14.c. That (2) and (3) are equivalent follows from Exercise 1.16. □

The result of Exercise 5.20 as stated is algorithmically impractical, because it requires us to compute $\mathcal{O}_K$. However, we have that $(\mathcal{O}_K : R)^2 \mid \Delta(R)$ and the latter we can compute. In this form we recover a single implication: If R is singular above p, i.e. R has a singular prime ideal containing p, then $p \mid \Delta(R)$. To get an implication in the reverse direction we use the reduced discriminant.

Definition 5.22. Let R be an order in a number field K. We define the *trace dual* $R^\dagger$ of R as the fractional R-ideal

$$R^\dagger = \{x \in K \mid \mathrm{Tr}_{K/\mathbb{Q}}(xR) \subseteq \mathbb{Z}\}.$$

We define the *reduced discriminant* $\delta(R)$ to be the exponent of the finite abelian group $R^\dagger/R$.

Note that we compute the trace dual and reduced discriminant in polynomial time.

Theorem 5.23. *Let R be an order in a number field K and $p > [K : \mathbb{Q}]$ be a prime number. Then there exists a singular prime ideal $\mathfrak{p} \subseteq R$ over p if and only if $p^2 \mid \delta(R)$.*

Proof. See Exercise 5.24 and Exercise 5.25. □

Exercise 5.24. Let R be an order of a number field K and let $p > [K : \mathbb{Q}]$ be a prime number. Suppose that $p \nmid (\mathcal{O}_K : R)$.
a. Show that $R_{(p)}$ is Dedekind and that $\mathrm{Trad}((R/pR)/\mathbb{F}_p) = (pR^\dagger \cap R)/pR$.
Since $R_{(p)}$ is Dedekind there exists a unique $m : \operatorname{spec} R_{(p)} \to \mathbb{Z}$ such that $R^\dagger_{(p)} = \prod_{\mathfrak{p}} \mathfrak{p}^{m(\mathfrak{p})}$.
b. Show that for all $\mathfrak{p} \in \operatorname{spec} R_{(p)}$ we have $m(\mathfrak{p}) = 1 - e(\mathfrak{p})$.
c. Conclude that $pR^\dagger \subseteq R$ and thus $p^2 \nmid \delta(R)$.

Exercise 5.25. Let R be an order of a number field K and let $p > [K : \mathbb{Q}]$ be a prime number. Write $\mathfrak{a}$ for the product of all prime ideals $\mathfrak{p} \subseteq R$ containing p.

a. Show that $\mathrm{Trad}((R/pR)/\mathbb{F}_p) = \mathfrak{a}/pR = \mathrm{rad}(R/pR)$.

b. Show that $\mathrm{Tr}_{R/\mathbb{Z}}$ induces a perfect pairing $\psi : R/\mathfrak{a} \otimes R/\mathfrak{a} \to \mathbb{F}_p$ of $\mathbb{F}_p$-vector spaces given by

$$(x+\mathfrak{a}, y+\mathfrak{a}) \mapsto (\mathrm{Tr}_{R/\mathbb{Z}}(xy) + p\mathbb{Z}).$$

In other words, show that ψ is a well-defined $\mathbb{F}_p$-linear map and that for any $y \in R/\mathfrak{a}$ the map $R/\mathfrak{a} \to \mathrm{Hom}(R/\mathfrak{a}, \mathbb{F}_p)$ given by $x \mapsto [y \mapsto \psi(x,y)]$ is an isomorphism of $\mathbb{F}_p$-vector spaces.

c. Show that there is a short exact sequence

$$0 \longrightarrow ((pR^\dagger) \cap (\mathfrak{a} : \mathfrak{a}))/\mathfrak{a} \longrightarrow (\mathfrak{a} : \mathfrak{a})/\mathfrak{a} \longrightarrow \mathrm{Hom}(R/\mathfrak{a}, \mathbb{F}_p) \longrightarrow 0.$$

Hint: The map $(\mathfrak{a} : \mathfrak{a})/\mathfrak{a} \to \mathrm{Hom}(R/\mathfrak{a}, \mathbb{F}_p)$ can be obtained by naturally extending the isomorphism $R/\mathfrak{a} \to \mathrm{Hom}(R/\mathfrak{a}, \mathbb{F}_p)$ from part (b).

d. Suppose that $p \mid (\mathcal{O}_K : R)$. Show that $R^\dagger/R$ contains an element of order p^2 and conclude that $p^2 \mid \delta(R)$.

5.4. Proof of the main theorems First we give a tool to compute the ring of integers $\mathcal{O}_K$ of a number field K given an order R of K and enough information about the singular primes of R.

Theorem 5.26. *There is a polynomial-time algorithm that, given an order R in a number field K and square-free integer d contained in all singular prime ideals, decides whether $R = \mathcal{O}_K$. Moreover, if $R \neq \mathcal{O}_K$, then it additionally computes an order $S \supsetneq R$ of K.*

Note that a possible d in the above is any (square free) multiple of $\mathrm{rad}(\Delta(R))$ by Exercise 5.20. To compute the ring of integers we repeatedly apply the algorithm to S if R is not already the ring of integers.

Algorithm 5.27. Let R be an order in a number field K and let d be square-free.

(1) Compute $\mathrm{rad}(R/dR)$ using Algorithm 5.17.
(2) Compute the inverse image $\mathfrak{a}$ of $\mathrm{rad}(R/dR)$ under the projection $R \to R/dR$.
(3) If $\mathfrak{a} : \mathfrak{a} = R$, return true. Otherwise, return false and $\mathfrak{a} : \mathfrak{a}$.

Proof of Theorem 5.26. We will show Algorithm 5.27 does as we want. Clearly it runs in polynomial time because of Theorem 5.16. Note that if $\mathfrak{a}$, because it is an ideal of R, indeed satisfies $\mathfrak{a} : \mathfrak{a} \neq R$, then clearly $R \subsetneq (\mathfrak{a} : \mathfrak{a}) \subseteq \mathcal{O}_K$. Then for correctness, it suffices to show that $R = \mathcal{O}_K$ when $\mathfrak{a} : \mathfrak{a} = R$.

We have

$$\mathrm{rad}(R/dR) = \prod_{\mathfrak{p} \in \mathrm{spec}(R/dR)} \mathfrak{p} = \prod_{\substack{\mathfrak{p} \in \mathrm{spec}(R) \\ dR \subseteq \mathfrak{p}}} (\mathfrak{p} + dR),$$

hence $\mathfrak{a}$ is the product of all prime ideals containing d. If $\mathfrak{a} : \mathfrak{a} = R$, then d is coprime to $(\mathcal{O}_K : R)$ by Proposition 5.21, hence $(\mathcal{O}_K : R) = 1$ by definition of d. Thus the algorithm is correct. □

We can now prove the main theorems.

Proof of Theorem 5.9. If we have an algorithm to compute $\mathrm{rad}(\#A)$, then Theorem 5.16 gives us an algorithm to compute $\mathrm{nil}(A)$. Conversely, we have that $\mathrm{rad}(m) = \#\mathrm{nil}(\mathbb{Z}/m\mathbb{Z})$. Hence we may compute $\mathrm{rad}(m)$ if we have an algorithm to compute $\mathrm{nil}(\mathbb{Z}/m\mathbb{Z})$. Therefore problem (2) and (3) are equally hard.

Suppose we have an algorithm for (1). We compute $\mathcal{O}_K$ for $K = \mathbb{Q}(\sqrt{m})$ and in turn we compute $\Delta(\mathcal{O}_K)$. Then from Exercise 5.19.d we obtain $\boxtimes(m)$ and in turn $\square(m)$. Finally, we compute $\mathrm{rad}(m) = \mathrm{lcm}(\boxtimes(m), \square(m))$, so (3) reduces to (1).

Suppose we have an algorithm for (3). Let R be the order encoding K. Then we may compute a positive integer d that is a multiple of $\mathrm{rad}\,|\Delta(R)|$, for example $d = |\Delta(R)|$. By Theorem 5.26 and Exercise 5.20 we have an algorithm that either tells us that $R = \mathcal{O}_K$ or gives an order $S \supsetneq R$. If $R = \mathcal{O}_K$ we are done, so suppose we have an order $S \supsetneq R$. Then $|\Delta(R)| = (R : S)^2 \cdot |\Delta(S)|$, so we may recursively apply our algorithm to S with the same d. In each recursive step $|\Delta(R)|$ decreases by a factor of at least 4, so the number of steps is bounded by $\log_4 |\Delta(R)|$, which in turn is bounded by a polynomial in the input size. Hence we have a polynomial-time algorithm for (1). □

Proof of Theorem 5.10. As before, (3) reduces to (2) by considering $A = \mathbb{Z}/m\mathbb{Z}$. For the converse, we compute $d = \exp(A^+)$ and note that if A is to be reduced we need d to be square-free. Assuming we have a polynomial-time algorithm for (3) we can check this, so we may assume d is square-free. However, in this case $d = \mathrm{rad}(\#A)$, so we may apply Theorem 5.16 to compute $\mathrm{nil}(A)$. Hence (2) reduces to (3). The reduction from (3) to (1) follows from Exercise 5.19 as in the proof of Theorem 5.9.

Now suppose we have a polynomial-time algorithm for (3). Then we may simply check whether $\delta(R)$ is square-free to conclude that either $R \neq \mathcal{O}_K$ or that the singular primes of R must lie over primes $p \leqslant [K : \mathbb{Q}]$ by Theorem 5.23. We then take d to be the product of all such primes that divide $\Delta_{R/\mathbb{Z}}$ and apply Theorem 5.26 with this choice of d. Hence (1) reduces to (3). □

5.5. Exercises

Exercise 5.28. Let q be a prime power, let C_6 be a cyclic group of order 6, and consider the group ring $\mathbb{F}_q[C_6]$.
a. Prove $\#\mathrm{spec}(\mathbb{F}_q[C_6]) = 6, 2, 2, 3, 4$ for $q \equiv 1, 2, 3, 4, 5 \bmod 6$ respectively.
b. Write $\mathbb{F}[C_6]$ explicitly as a product of local rings for $q = 2, 3, 4, 5, 7$.

Exercise 5.29. The reductions of algorithmic problems in we consider are *Cook reductions*, where $f \preceq g$ if we can construct an algorithm for f using an 'oracle' for g. We can also consider a different type of reduction, more similar to a *Karp reduction*, where the oracle can only be consulted once. Prove a version of Theorem 5.9 for Karp reductions.

Exercise 5.30. Let $n, m \in \mathbb{Z}_{>0}$. Prove that there exists a number field K such that $[K : \mathbb{Q}] = n$ and $\gcd(\Delta_K, m) = 1$.

Exercise 5.31. Let $A \subseteq B$ be commutative rings such that B is free over A. Show that the map $B \to \mathrm{Hom}_A(B, A)$ given by $x \mapsto (y \mapsto \mathrm{Tr}_{B/A}(xy))$ is an isomorphism of B-modules if and only $\Delta(B/A)$ is a unit in A.

Exercise 5.32. Let A be a finite commutative ring, and let $m \in \mathbb{Z}_{>0}$ such that $mA = 0$.

a. Prove that for some $t, b \in \mathbb{Z}_{>0}$ the following is true: for all sequences $(d_i)_{i=1}^t$ consisting of t integers $d_i \geqslant b$ there exists a sequence $(e_i)_{i=1}^t$ of integers satisfying $0 \leqslant e_i < d_i$ such that for all sequences $(f_i)_{i=1}^t$ of monic polynomials $f_i \in \mathbb{Z}[X]$ satisfying $f_i \equiv X^{d_i} - X^{e_i} \bmod m\mathbb{Z}[X]$ there exists a surjective ring homomorphism

$$\mathbb{Z}[X_1, \ldots, X_t]/(f_1(X_1), \ldots, f_t(X_t)) \to A.$$

b. Prove that there exists an order R in some number field and an ideal $I \subseteq R$ such that $R/I \cong A$ as rings.

Note: This 'proves' a converse to the statement in the very beginning of this section. Questions about finite commutative rings can be reduced to questions about (ideals in) number rings. We do however say nothing about the quality of such a reduction in the context of algorithms.

Exercise 5.33. Here we show that the conclusions to Exercise 5.24 and Exercise 5.25 fail to hold when we drop the restriction that the prime is at most the degree of the number field.

a. Show that there exists a quadratic order R for which $2 \mid (\mathcal{O}_{R_{0R}} : R)$ and $2^2 \nmid \delta(R)$.

b. Show that there exists a quadratic order R for which $2 \nmid (\mathcal{O}_{R_{0R}} : R)$ and $2^2 \mid \delta(R)$.

6. Inverting ideals

Ideals do not exist without a ring which they are ideals of. However, we should be prepared to encounter many orders of a fixed number field in our algorithms, since many of these algorithms have a 'back door' that allow them to output a strictly larger order than it has gotten as input. As a consequence, the ideals put in the algorithm should also change together with these orders. It will be fruitful to consider a concept more general than ideals.

Definition 6.1. Let S be a commutative ring. For subsets $I, J \subseteq S$ we write

$$I + J = \{x + y : x \in I,\ y \in J\}$$

$$I \cdot J = \Big\{ \sum_{i=1}^{n} x_i y_i : n \geqslant 0,\ x_1, \ldots, x_n \in I,\ y_1, \ldots, y_n \in J \Big\}$$

$$I : J = (I : J)_S = \big\{ x \in S : xJ \subseteq I \big\}.$$

We also inductively define $I^0 = (I : I)_S$ and $I^{n+1} = I^n \cdot I$ for $n \geqslant 0$. Suppose $R \subseteq S$ is a subring and I is also an R-module. We say I is *invertible over* R if there exists some R-submodule $I^{-1} \subseteq S$ such that $I \cdot I^{-1} = R$. We write $\mathcal{I}(R) = \mathcal{I}_S(R)$ for the group of subgroups of S that are invertible over R.

Although the operations of Definition 6.1 are defined for arbitrary subsets of S, we will mostly restrict to (additive) subgroups. Note that if $I, J \subseteq S$ are subgroups, so are $I + J$, $I \cdot J$ and $I : J$.

Lemma 6.2. *Let* S *be a commutative ring and* $H, I, J \subseteq S$ *subgroups. Then* (i) $I : I$ *is a subring of* S*;* (ii) $I^1 = I$*;* (iii) $I : I = I$ *if and only if* I *is a subring of* S*;* (iv) *if* I *is finitely generated and* $J \cap S^* \neq \emptyset$*, then* $I : J$ *is finitely generated;* (v) $I : J \subseteq (HI) : (HJ)$*, with equality if* H*,* I *and* J *are* R*-modules for some subring* $R \subseteq S$ *and* H *is invertible over* R*;* (vi) *if* $S \subseteq A$ *is a ring extension and* $J \cap S^* \neq \emptyset$*, then* $(I : J)_A = (I : J)_S$.

Proof. (i) This follows directly from the definition. (ii) Since $1 \in I : I$ by (i), we have $I \subseteq (I : I) \cdot I$, while $(I : I) \cdot I \subseteq I$ by definition of the quotient. Thus $I = (I : I) \cdot I = I^1$. (iii) The implication ($\Rightarrow$) follows from (i). For ($\Leftarrow$), suppose $I \subseteq S$ is a subring. Then $I \cdot I \subseteq I$ and consequently $I \subseteq I : I$. For $x \in I : I$ we have $xI \subseteq I$, so in particular $x = x \cdot 1 \in I$ and thus $I : I \subseteq I$. (iv) Let $j \in J$ be invertible. Then $j^{-1}I$ is finitely generated and hence a Noetherian $\mathbb{Z}$-module. Thus $I : J \subseteq j^{-1}I$ is finitely generated. (v) The inclusion follows directly from the definition. If H is invertible we get $I : J \subseteq (HI) : (HJ) \subseteq (H^{-1}HI) : (H^{-1}HJ) = (RI) : (RJ) = I : J$, from which equality follows. (vi) Clearly $(I : J)_S \subseteq (I : J)_A$ because $S \subseteq A$. Let $j \in J$ be invertible. If $x \in (I : J)_A$, then $x \in j^{-1}I \subseteq S$, so $x \in (I : J)_S$. □

Exercise 6.3. Let S be a commutative ring and $R \subseteq S$ a subring and $I \subseteq S$ an R-module. Show that if I is invertible over R, then $R : I$ is its inverse and that I is a finitely generated R-module.

Exercise 6.4. Let K be a number field and $I \subseteq K$ be a subgroup. Then $I^2 = I$ if and only if I is a ring. *Hint:* If $I^2 = I$, then I is a finitely generated ideal of the ring $\mathbb{Z} + I$.

Exercise 6.5. Let $K \subseteq L$ be a finite degree field extension. Show for all subgroups $I, J \subseteq L$ that $(KI) : (KJ) \supseteq K(I : J)$, with equality when J is finitely generated. Give an example where equality does not hold.

Exercise 6.6. Let R be an order in a number field and recall the definition of $R^\dagger$ from Definition 5.22. Show that the following are equivalent: (1) $R^\dagger$ is invertible; (2) For all fractional ideals I of R it holds that $I : I = R$ if and only I is invertible.

Proposition 6.7. *Let* S *be a commutative ring and* $I \subseteq S$ *a subgroup. Then the* blowup *at* I,

$$\mathrm{Bl}(I) = \mathrm{Bl}_S(I) := \bigcup_{k \geqslant 0} (I^k : I^k)_S,$$

is a ring. If RI *is an invertible* R*-ideal for some subring* $R \subseteq S$*, then* $\mathrm{Bl}(I) \subseteq R$*. If* S *is a number field and* I *is finitely generated and non-zero, then* $\mathrm{Bl}(I) = (I^n : I^n)$ *for some* $n \geqslant 0$ *and* $\mathrm{Bl}(I)$ *is an order.*

Proof. It follows from Lemma 6.2 for all $k \geqslant 0$ that $I^k : I^k$ is a ring and that $I^k : I^k \subseteq I^{k+1} : I^{k+1}$. Hence $\mathrm{Bl}(I)$, the direct limit of rings, is a ring.

Suppose RI is an invertible R-ideal. Then $RI^k = (RI)^k$ is invertible for all $k \geqslant 0$. In particular, by Lemma 6.2 we have $I^k : I^k \subseteq (RI^k) : (RI^k) = R : R = R$ for all k, so $\mathrm{Bl}(I) \subseteq R$.

Suppose S is a number field and I is finitely generated and non-zero. Then $I^k : I^k$ is both a subring and a finitely generated subgroup of S by Lemma 6.2, hence it is an order. Because $I^k : I^k \subseteq \mathcal{O}_S$ and $\mathcal{O}_S$ is a Noetherian $\mathbb{Z}$-module, there must exist some $n \geqslant 0$ such that $I^n : I^n = I^k : I^k$ for all $k \geqslant n$. Consequently $I^n : I^n = \mathrm{Bl}(I)$. □

We will show later that, in the case of orders in number fields, $\mathrm{Bl}(I)I$ is invertible over $\mathrm{Bl}(I)$.

Exercise 6.8. Let S be a commutative ring and $I, J \subseteq S$ subgroups. Show that $\mathrm{Bl}(I) \cdot \mathrm{Bl}(J) \subseteq \mathrm{Bl}(IJ)$. Show that $\mathrm{Bl}(I) = I$ if and only if I is a subring of S.

Exercise 6.9. Let $R \subseteq S$ be a commutative subrings and and $I, J \subseteq S$ fractional R-ideals. If I is invertible, then $\mathrm{Bl}(IJ) = \mathrm{Bl}(J)$.

Exercise 6.10. Give an example of a field K and a finitely generated subgroup $I \subseteq K$ such that $I^n : I^n = I^{n+1} : I^{n+1} \neq \mathrm{Bl}(I)$ for some $n \geqslant 1$.

Proposition 6.11. *Let $R \subseteq S$ be subrings such that R is semi-local and let $I \subseteq S$ be an R-module. Then I is invertible if and only if there exists some $x \in S^*$ such that $I = xR$.*

Proof. The implication ($\Leftarrow$) is obvious, since xR has inverse $x^{-1}R$.

Now suppose I is invertible with inverse J. Let $\mathfrak{m} \subseteq R$ be maximal. As $IJ = R \not\subseteq \mathfrak{m}$ there exist $x_\mathfrak{m} \in I$ and $y_\mathfrak{m} \in J$ such that $x_\mathfrak{m} y_\mathfrak{m} \in R \setminus \mathfrak{m}$. Since R is semi-local, there exist $\lambda_\mathfrak{m} \in R$ for each maximal $\mathfrak{m} \subseteq R$ such that $\lambda_\mathfrak{m} \in \mathfrak{m}'$ if and only if $\mathfrak{m} \neq \mathfrak{m}'$: Namely, let $r_{\mathfrak{m}'} \in \mathfrak{m}' \setminus \mathfrak{m}$, which exist since the ideals are maximal, and take $\lambda_\mathfrak{m} = \prod_{\mathfrak{m}' \neq \mathfrak{m}} r_{\mathfrak{m}'}$. Consider $x = \sum_\mathfrak{m} \lambda_\mathfrak{m} x_\mathfrak{m}$ and $y = \sum_\mathfrak{m} \lambda_\mathfrak{m} y_\mathfrak{m}$. Then

$$xy = \sum_{\mathfrak{m},\mathfrak{m}'} \lambda_\mathfrak{m} \lambda_{\mathfrak{m}'} x_\mathfrak{m} y_{\mathfrak{m}'}.$$

For all $\mathfrak{m}$ there is precisely one term not contained in $\mathfrak{m}$, which is $\lambda_\mathfrak{m} \lambda_\mathfrak{m} x_\mathfrak{m} y_\mathfrak{m}$, hence $xy \notin \mathfrak{m}$. Since $x_\mathfrak{m} y_{\mathfrak{m}'} \in IJ = R$ we have $xy \in R$. It follows that xy is a unit of R, and x is a unit of S. Finally, $xR \subseteq I = xyI \subseteq xJI = xR$ and $I = xR$. □

Corollary 6.12. *Let $R \subseteq S$ be subrings such that R is semi-local. Then $\mathcal{I}_S(R) = S^*/R^*$.* □

Exercise 6.13. Let R be a subring of a number field K. Show that **a.** R/aR is finite for all non-zero $a \in R$; **b.** R is Noetherian; **c.** every non-zero prime ideal of R is maximal; and **d.** if R has a local subring, then R is semi-local. *Hint:* First suppose $a \in \mathbb{Z}$ is a prime (power).

Lemma 6.14. *Let* R *be a domain and* K *its field of fractions. Then for every localization* S *of* R *the natural map* $S \to K$ *is injective, and* $I = \bigcap_{\mathfrak{m} \in \operatorname{maxspec} R} I_{\mathfrak{m}}$ *for every fractional ideal* I *of* R.

Proof. The maps $S \to K$ are injective since R and hence S is a domain. Clearly $I \subseteq I_{\mathfrak{m}}$. Suppose $x \in \bigcap_{\mathfrak{m}} I_{\mathfrak{m}}$ and $x \notin I$. Consider $J = (I : xR)_R \subsetneq R$ is an ideal not containing 1. Let $J \subseteq \mathfrak{m} \subseteq R$ be a maximal ideal. Since $x \in I_{\mathfrak{m}}$ there exists some $r \in R \setminus \mathfrak{m}$ such that $rx \in I$. Then $r \in J$ by definition of J, which contradicts $J \subseteq \mathfrak{m} \not\ni r$. Hence $I = \bigcap_{\mathfrak{m}} I_{\mathfrak{m}}$. □

Exercise 6.15. Let $R \subseteq K$ be subrings such that K is a number field and the field of fractions of R. Let $\mathfrak{p} \subseteq R$ be a maximal ideal. Show that if $I_{\mathfrak{p}} \subseteq R_{\mathfrak{p}}$ is an invertible ideal over $R_{\mathfrak{p}}$, then $J = I_{\mathfrak{p}} \cap R$ is an invertible ideal over R such that $J_{\mathfrak{p}} = I_{\mathfrak{p}}$, and $J_{\mathfrak{q}} = R_{\mathfrak{q}}$ for all maximal $\mathfrak{q} \subseteq R$ not equal to $\mathfrak{p}$.

Theorem 6.16. *Let* K *be a number field and* $R \subseteq S \subseteq K$ *be subrings such that* K *is the field of fractions of* R. *Then we have a group isomorphism* $\mathcal{I}_K(S) \to \bigoplus_{\mathfrak{p} \in \operatorname{maxspec} R} \mathcal{I}_K(R_{\mathfrak{p}} S)$ *given by* $I \mapsto (R_{\mathfrak{p}} I)_{\mathfrak{p}}$, *with inverse* $(I_{\mathfrak{p}})_{\mathfrak{p}} \mapsto \bigcap_{\mathfrak{p}} I_{\mathfrak{p}}$. □

Proof. The diagram constructed from maps as in the theorem is commutative.

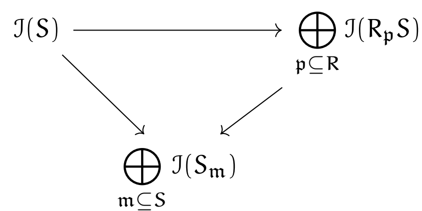

Hence to show that the horizontal map is an isomorphism, it suffices to prove this for the downward maps. These maps are both injective with the appropriate left inverse by Lemma 6.14, while they are surjective by Exercise 6.15. □

Corollary 6.17. *Let* R *be an order in a number field. Then a fractional* R*-ideal is invertible if and only if it is locally principal.* □

Lemma 6.18. *Let* K *be a number field and* $R \subseteq S \subseteq K$ *subrings such that* K *is the field of fractions of* R. *Then there is a surjective group homomorphism* $\sigma : \mathcal{I}(R) \to \mathcal{I}(S)$ *given by* $I \mapsto SI$.

Proof. The only non-trivial part is surjectivity. Using Theorem 6.16 one verifies that the induced map $\sigma : \bigoplus_{\mathfrak{p} \subseteq R} \mathcal{I}(R_{\mathfrak{p}}) \to \bigoplus_{\mathfrak{p} \subseteq R} \mathcal{I}(R_{\mathfrak{p}} S)$ is the direct sum of maps $\sigma_{\mathfrak{p}} : \mathcal{I}(R_{\mathfrak{p}}) \to \mathcal{I}(R_{\mathfrak{p}} S)$ given by $I \mapsto SI$. Thus it suffices to prove the theorem when R is local, and hence S is semi-local by Exercise 6.13.d. This case follows immediately from Corollary 6.12 □

6.1. Blowing up We will prove that for a number field K and a finitely generated subgroup $I \subseteq K$ the group $\mathrm{Bl}(I) \cdot I$ is an invertible $\mathrm{Bl}(I)$-ideal.

Definition 6.19. Let R be a commutative ring and M an R-module. We define

$$\gamma_R(M) = \inf\Big\{\#X : X \subseteq M, \sum_{x\in X} Rx = M\Big\}.$$

Note that M is finitely generated if and only if $\gamma_R(M) < \infty$.

Lemma 6.20. *Let R be a commutative ring and M an R-module. Then*

(1) *we have $\gamma_R(M/N) \leqslant \gamma_R(M)$ for all submodules $N \subseteq M$;*
(2) *if R is a principal ideal domain, then $\gamma_R(N) \leqslant \gamma_R(M)$ for all submodules $N \subseteq M$;*
(3) *if M is finitely generated and $I \subseteq \mathrm{Jac}(R)$, then $\gamma_R(M) = \gamma_{R/I}(M/IM)$;*
(4) *if A is a commutative R-algebra and N an A-module, then*

$$\gamma_R(N) \leqslant \gamma_R(A)\cdot\gamma_A(N) \quad \textit{and} \quad \gamma_A(M\otimes_R N) \leqslant \gamma_R(M)\cdot\gamma_A(N),$$

with equality if A and R are fields.

Proof. (i) We may simply use a generating set of M for M/N. (ii) It suffices to prove this when M is free, since we may replace M by $R^{(X)}$ for any generating set $X \subseteq M$. Since R is a principal ideal domain, N is free of rank at most #X. Hence $\gamma_R(N) \leqslant \gamma_R(M)$. (iii) Write $J = \mathrm{Jac}(R)$. Then $\gamma_R(M) \geqslant \gamma_{R/I}(M/IM) \geqslant \gamma_{R/J}(M/JM)$ by (i). The inequality $\gamma_{R/J}(M/JM) \geqslant \gamma_R(M)$ follows from Nakayama's lemma. (iv) One can verify for both inequalities that the product of two generating sets on the modules on the right always gives a generating set for the module on the left. If A and R are fields, then the remaining statement follows from linear algebra. □

Exercise 6.21. Let R be a commutative ring and M an R-module. Show that

$$\sup\{\gamma_{R/\mathfrak{m}}(M/\mathfrak{m}M) : \mathfrak{m} \in \operatorname{maxspec} R\} \leqslant \gamma_R(M),$$

with equality if R is semi-local and M is finitely generated.

Exercise 6.22. Let $f : R \to S$ be a morphism of commutative rings such that S is finitely generated when interpreted as R-module. Show that the map $\operatorname{maxspec} S \to \operatorname{maxspec} R$ given by $\mathfrak{m} \mapsto f^{-1}\mathfrak{m}$ is well-defined and has fibers of size at most $\gamma_R(S)$.

Lemma 6.23. *Let R be a local ring with maximal ideal $\mathfrak{p}$, and S a commutative R-algebra. Let $M \subseteq S$ be a R-submodule such that $M \not\subseteq \mathfrak{m}$ for all maximal ideals $\mathfrak{m} \subseteq S$. If $\#\operatorname{maxspec} S < \#(R/\mathfrak{p}) < \infty$, then $M \cap S^* \neq \emptyset$.*

Proof. Since $S \setminus S^* = \bigcup_{\mathfrak{m}} \mathfrak{m}$, where the union ranges over the maximal ideals of S, it suffices to show that $\bigcup_{\mathfrak{m}}(M \cap \mathfrak{m}) \neq M$. We will prove the stronger statement that an R-module M cannot be written as a union of k proper submodules for any $1 \leqslant k < \#(R/\mathfrak{p})$. For $k = 1$ this is trivial, so suppose $k \geqslant 2$. For the sake of contradiction suppose $M = \bigcup_{i\in k} M_i$ for some submodules $M_i \subsetneq M$. By removing submodules we may assume there exist no inclusion relations between these modules. In particular, for each i there exists some $m_i \in M_i$ that is not contained in any other subspace. Consider the map $\ell : R \to M$ given by $t \mapsto$

$tm_1 + (1-t)m_2$. Since $k < \#(R/\mathfrak{p})$ there exist $s, t \in R$ such that $s \not\equiv t \bmod \mathfrak{p}$ and $\ell(s), \ell(t) \in M_i$ for some i. Note that $s - t \in R^*$ because $s - t \not\equiv 0 \bmod \mathfrak{p}$. We have $(s-t)(m_1 - m_2) = \ell(s) - \ell(t) \in M_i$, so $m_1 - m_2 \in M_i$. Hence $\mathrm{im}(\ell) \subseteq M_i$ and $m_1, m_2 \in M_i$, which is a contradiction. □

Proposition 6.24. *Let Z be a principal ideal domain and $R \subseteq S$ commutative Z-algebras such that R is local and S is a Noetherian Z-module. Let $M \subseteq S$ be an R-submodule such that $M \not\subseteq \mathfrak{m}$ for all maximal ideals $\mathfrak{m} \subseteq S$. If $M^n : M^n = R$ for some integer $n \geqslant \gamma_R(S) - 1$, then $M = Ru$ for some $u \in S^*$.*

Proof. For a local ring A with unique maximal ideal $\mathfrak{m}$ write $\kappa(A) = A/\mathfrak{m}$ for the residue field. By localizing Z at the maximal ideal of R we may assume without loss of generality that Z is local.

We will first show that there exists an R-algebra A which is local and free of finite rank d such that $\#\kappa(A) > d \cdot \#\operatorname{maxspec} S$. Because S is semi-local (see Exercise 6.22), we may choose $A = R$ if $\kappa(R)$ is infinite. Thus assume $\kappa(R)$ is finite, hence $\kappa(R) = \mathbb{F}_q$ for some power q of a prime p. Choose d sufficiently large such that $\#\kappa(R)^d > d \cdot \#\operatorname{maxspec} S$. Then consider the field extension $\kappa(R) \subseteq \mathbb{F}_{q^d} = \mathbb{F}_q(\alpha)$ for some $\alpha \in \mathbb{F}_{q^d}$. Let $f \in R[X]$ be some monic lift of the minimal polynomial of α over $\kappa(R)$. Clearly $A = R[X]/(f)$ is free of rank d as R-module. Since R has a unique maximal ideal $\mathfrak{p}$, the maximal ideals of A correspond to those of $A/\mathfrak{p}A = \mathbb{F}_{q^d}$ by Exercise 6.22, which is a field. Hence A is local with $\kappa(A) = \mathbb{F}_{q^d}$. Finally, $\#\kappa(A) = \#\kappa(R)^d > d \cdot \#\operatorname{maxspec} S$, as was to be shown.

Write $(-)_A$ for the functor $- \otimes_R A$ and $\mathfrak{q}$ for the maximal ideal of A. We want to apply Lemma 6.23 to M_A, for which we now verify the conditions. Note that $(-)_A$ is exact since A is free as R-module. In particular, $M_A^n : M_A^n = (M^n : M^n)_A = R_A = A$ and $S \to S_A$ is injective. If $M_A \subseteq \mathfrak{n}$ for some maximal $\mathfrak{n} \subseteq S_A$ then also $M = M_A \cap S$ is contained in the maximal ideal of $\mathfrak{n} \cap S$ of S by Exercise 6.22, which is a contradiction. Lastly, $\#\operatorname{maxspec} S_A \leqslant d \cdot \#\operatorname{maxspec} S < \#\kappa(A)$. Thus there exists some $u \in M_A \cap S_A^*$ by Lemma 6.23.

We will first show $M_A = Au$, or equivalently show that $I := M_A u^{-1}$ equals A. Consider $I^\infty = \bigcup_{k \geqslant 0} I^k \subseteq S_A$, which is the smallest subring containing I. Using Lemma 6.20 we obtain

$$\begin{aligned}\gamma_{Z/(\mathfrak{q}\cap Z)}(\kappa(A)) \cdot \dim_{A/\mathfrak{q}}(I^\infty/\mathfrak{q}I^\infty) &= \gamma_{Z/(\mathfrak{q}\cap Z)}(I^\infty/\mathfrak{q}I^\infty) \leqslant \gamma_Z(I^\infty) \leqslant \gamma_Z(S_A)\\ &\leqslant \gamma_Z(A) \cdot \gamma_A(S_A) = \gamma_{Z/(\mathfrak{q}\cap Z)}(\kappa(A)) \cdot \gamma_A(S_A) \leqslant \gamma_{Z/(\mathfrak{q}\cap Z)}(\kappa(A)) \cdot \gamma_R(S),\end{aligned}$$

and thus $\dim_{A/\mathfrak{q}}(I^\infty/\mathfrak{q}I^\infty) \leqslant \gamma_R(S) \leqslant n+1$. From $1 \in I$ it follows that $I^0 \subseteq I^1 \subseteq I^2 \subseteq \cdots$, so $(I^0 + \mathfrak{q}I^\infty)/\mathfrak{q}I^\infty \subseteq (I^1 + \mathfrak{q}I^\infty)/\mathfrak{q}I^\infty \subseteq \cdots$ with $\dim_{A/\mathfrak{q}}((I^0 + \mathfrak{q}I^\infty)/\mathfrak{q}I^\infty) = 1$. We conclude that $(I^r + \mathfrak{q}I^\infty)/\mathfrak{q}I^\infty = (I^{r+1} + \mathfrak{q}I^\infty)/\mathfrak{q}I^\infty$ for some $r \leqslant n$. It follows that $(I^n + \mathfrak{q}I^\infty)/\mathfrak{q}I^\infty = I^\infty/\mathfrak{q}I^\infty$. Equivalently $I^n + \mathfrak{q}I^\infty = I^\infty$, so by Nakayama's lemma we have $I^n = I^\infty$. As I^n is a ring we have $I \subseteq I^n = I^n : I^n = A \subseteq I$ by Lemma 6.2, so $I = A$.

It follows that

$$1 = \gamma_A(M_A) = \gamma_{A/\mathfrak{q}}(M_A/\mathfrak{q}M_A) = \gamma_{A/\mathfrak{q}}((M/\mathfrak{p}M) \otimes_{R/\mathfrak{p}} (A/\mathfrak{q}))$$
$$= \gamma_{R/\mathfrak{p}}(M/\mathfrak{p}M) = \gamma_R(M).$$

Hence $M = Rv$ for some $v \in M$. If $v \notin S^*$, then $v \in \mathfrak{m}$ for some maximal $\mathfrak{m} \subseteq S$ and $M = vR \subseteq \mathfrak{m}$, a contradiction. Thus v is a unit. □

Definition 6.25. Let $R \subseteq S$ be orders in a number field. Then the *conductor* $\mathfrak{f}_{S/R}$ of S over R is the largest ideal of S contained in R.

Clearly the conductor is well-defined: It is the sum of all ideals of S contained in R, or equivalently $(R : S)_S$. As S/R is finite we have $\#(S/R) \cdot S \subseteq R$. Thus $\#(S/R) \cdot S \subseteq \mathfrak{f}_{S/R}$ and the latter is non-zero.

Theorem 6.26. *Let K be a number field and $I \subseteq K$ a finitely generated non-zero subgroup. Then $\mathbb{Q}(\mathrm{Bl}(I)) = \mathbb{Q}(Ix^{-1})$ for any non-zero $x \in I$, and $\mathrm{Bl}(I) = I^n : I^n$ for any $n \geqslant [K : \mathbb{Q}] - 1$. For all subrings $R \subseteq K$, the ideal RI is invertible if and only if $\mathrm{Bl}(I) \subseteq R$.*

Proof. Write $A = I^n : I^n$ and $F = \mathbb{Q}(\mathrm{Bl}(I))$. We will first show that $\mathrm{Bl}(I) = A$.

Clearly $A \subseteq \mathrm{Bl}(I)$ by definition of the blowup, so by Proposition 6.7 it suffices to show that AI is an invertible A-ideal. Since $\mathcal{O}_F I$ is an invertible $\mathcal{O}_F$-ideal there exists some invertible ideal H of A such that $\mathcal{O}_F I = \mathcal{O}_F H$ by Lemma 6.18. Consider $J = I(A : H)_F$, which is an A-ideal. To prove AI is an invertible A-ideal it suffices by Lemma 6.2 to show J is an invertible A-ideal. Note that $(J^n : J^n)_F = (I^n : I^n)_F = A$ since $(A : H)_F$ is invertible. Thus we may assume without loss of generality that $\mathcal{O}_F I = \mathcal{O}_F$ by replacing I with J.

Consider the conductor $\mathfrak{f} = \mathfrak{f}_{\mathcal{O}_F/A}$. We have $\mathfrak{f} = \mathfrak{f}\mathcal{O}_F = \mathfrak{f}\mathcal{O}_F I = \mathfrak{f}I \subseteq AI = I$ and $I \subseteq \mathcal{O}_F$. From $\mathcal{O}_F I = \mathcal{O}_F$ it also follows that $I \not\subseteq \mathfrak{m}$ for all maximal ideals $\mathfrak{m} \subseteq \mathcal{O}_F$. Now apply Proposition 6.24 with $A/\mathfrak{f}$, $\mathcal{O}_F/\mathfrak{f}$ and $I/\mathfrak{f}$ in the place of respectively R, S and M, using $\gamma_R(S) \leqslant \gamma_{\mathbb{Z}}(\mathcal{O}_F) \leqslant [F : \mathbb{Q}] \leqslant n + 1$. Then $I = Ax + \mathfrak{f}$ for some $x \in \mathcal{O}_F$ coprime to $\mathfrak{f}$. Thus $I_\mathfrak{p} = A_\mathfrak{p} x$ if $\mathfrak{p} + \mathfrak{f} = A$ and $I_\mathfrak{p} = A_\mathfrak{p} 1 + \mathfrak{f}_\mathfrak{p} = A_\mathfrak{p}$ if $\mathfrak{f} \subseteq \mathfrak{p}$. Hence I is locally principal and thus invertible by Corollary 6.17, as was to be shown.

Proposition 6.7 states that for all subrings $R \subseteq K$, the ideal RI is invertible if and only if $\mathrm{Bl}(I) \subseteq R$.

It remains to be shown that $\mathbb{Q}(\mathrm{Bl}(I)) = \mathbb{Q}(Ix^{-1})$ for any non-zero $x \in I$. Since $\mathrm{Bl}(I) = \mathrm{Bl}(aI)$ for any $a \in K^*$, we may replace I by Ix^{-1}, so that $1 \in I$, and show that $\mathbb{Q}(\mathrm{Bl}(I)) = \mathbb{Q}(I)$. Since $I\mathbb{Q}(I) = \mathbb{Q}(I)$ is invertible over $\mathbb{Q}(I)$ we conclude that $\mathrm{Bl}(I) \subseteq \mathbb{Q}(I)$. As $1 \in I$ we have a chain $\mathbb{Q}I^0 \subseteq \mathbb{Q}I^1 \subseteq \mathbb{Q}I^2 \subseteq \cdots$, and since K is finite dimensional it stabilizes at say $\mathbb{Q}I^k$. Then $I \subseteq (\mathbb{Q}I^k : \mathbb{Q}I^k) = \mathbb{Q}(I^k : I^k) \subseteq \mathbb{Q}(\mathrm{Bl}(I))$ by Exercise 6.5, from which $\mathbb{Q}(I) = \mathbb{Q}(\mathrm{Bl}(I))$ follows. □

Corollary 6.27. *There exists a polynomial-time algorithm that, given a number field K and a finitely generated subgroup $I \subseteq K$, computes $\mathrm{Bl}(I)$.* □

Exercise 6.28. Let K be a number field and $R \subseteq K$ a subring. Show that for all finitely generated subgroups $I, J \subseteq K$ we have $\mathrm{Bl}(IJ) = \mathrm{Bl}(I) \cdot \mathrm{Bl}(J)$. Conclude $\mathrm{Bl}(R\alpha + R\beta) = R \cdot \mathrm{Bl}(\mathbb{Z}\alpha + \mathbb{Z}\beta)$ for all $\alpha, \beta \in K^*$.

Exercise 6.29. Let K be a number field and $I \subseteq K$ a finitely generated non-zero subgroup. If RI is invertible over $R = I^n : I^n$, then $R = \mathrm{Bl}(I)$.

Exercise 6.30. Let R be a number ring and I a fractional ideal of R. Show that $\mathrm{Bl}(I)$ is a subring of $(R : I) : (R : I)$.

Exercise 6.31. Let $R \subseteq S$ be orders of maximal rank in a number field. Show that $\ker(\mathcal{I}(R) \to \mathcal{I}(S)) = (S/\mathfrak{f}_{S/R})^*/(R/\mathfrak{f}_{S/R})^*$. Conclude that this kernel equals $\mathcal{I}(R)^{\mathrm{tors}}$ when S is the maximal order.

Exercise 6.32. Let R be an order and let K be its field of fractions. Show that that the following are equivalent:

1) There exist a finite *Boolean ring* B, i.e. a ring isomorphic to $\mathbb{F}_2^I$ for some set I, a subring $A \subseteq B$ and a surjection $f : \mathcal{O}_K \to B$ such that $R = f^{-1}A$.

2) There exists an equivalence relation $\sim$ on the set $\mathrm{Hom}_{\mathrm{Ring}}(\mathcal{O}_K, \mathbb{F}_2)$ such that $R = \{x \in \mathcal{O}_K : f \sim g \Rightarrow f(x) = g(x)\}$.

3) The group $\mathcal{I}(R)$ is torsion free.

Give an example of a non-maximal order such that $\mathcal{I}(R)$ is torsion free.

6.2. Inverting subgroups with two generators We will now give an explicit way to compute $\mathrm{Bl}(I)$ for subgroups I of number fields K of the form $I = \alpha\mathbb{Z} + \beta\mathbb{Z}$ for $\alpha, \beta \in K^*$.

Exercise 6.33. Let $f \in \mathbb{Z}[X]$ be a *primitive* polynomial, i.e. a polynomial for which the only integer divisors are ± 1. Show that $(f\mathbb{Q}[X]) \cap \mathbb{Z}[X] = f\mathbb{Z}[X]$. *Hint:* Gauss' lemma states that for $f, g \in \mathbb{Z}[X]$ the product fg is primitive if and only if f and g are primitive.

Theorem 6.34. *Let $K = \mathbb{Q}(\gamma)$ be a number field of degree $n \geqslant 2$ and let $f = a_n X^n + \cdots + a_0 X^0 \in \mathbb{Z}[X]$ be irreducible with $f(\gamma) = 0$. Consider for $0 \leqslant j \leqslant n$ the elements*

$$p_j = a_n\gamma^{j-0} + a_{n-1}\gamma^{j-1} + \cdots + a_{n-j} \quad \textit{and} \quad q_j = a_0\gamma^{0-j} + a_1\gamma^{1-j} + \cdots + a_j$$

and the groups

$$D = \mathbb{Z}p_0 + \mathbb{Z}p_1 + \cdots + \mathbb{Z}p_{n-1} \quad \textit{and} \quad N = \mathbb{Z}q_0 + \mathbb{Z}q_1 + \cdots + \mathbb{Z}q_{n-1}.$$

Then (i) $A = \mathbb{Z} + D = \mathbb{Z} + N$ *is a ring;* (ii) N *and* D *are coprime invertible ideals of* A *such that* $\gamma A = N : D$ *and such that we have ring isomorphisms* $A/N \cong \mathbb{Z}/a_0\mathbb{Z}$ *and* $A/D \cong \mathbb{Z}/a_n\mathbb{Z}$; (iii) $A = M : M = \mathbb{Z}[\gamma] \cap \mathbb{Z}[\gamma^{-1}] = \mathrm{Bl}(\mathbb{Z} + \gamma\mathbb{Z})$ *for* $M = \mathbb{Z} + \mathbb{Z}\gamma + \cdots + \mathbb{Z}\gamma^{n-1}$.

Proof. Observe that there is an involution $-^*$ on set of inputs to this theorem that sends $\gamma \mapsto \gamma^{-1}$ and $a_n X^n + \cdots + a_0 X^0 \mapsto a_0 X^n + \cdots + a_n X^0$. From this we obtain

$$p_j^* = a_n^*(\gamma^*)^{j-0} + \cdots + a_{n-j}^* = a_0\gamma^{0-j} + \cdots + a_j = q_j$$

and $N^* = D$.

(i) Note that $p_n = f(\gamma) = 0$. We additionally define $p_j = 0$ for $j > n$ and $a_j = 0$ for $j < 0$. Then note that $p_{j+1} = \gamma p_j + a_{n-j-1}$ for all $j \geqslant 0$. We will show by induction on $i \geqslant 0$ that for all $j \geqslant 0$ we have $p_i p_j \in D$. For $i = 0$ we have $p_i p_j = a_n p_j \in \mathbb{Z} p_j \subseteq D$ for all j. For arbitrary i we have

$$p_{i+1}p_j - p_i p_{j+1} = a_{n-i-1}p_j - a_{n-j-1}p_i \in D.$$

As $p_i p_{j+1} \in D$ by the induction hypothesis we obtain $p_{i+1}p_j \in D$, as was to be shown. It now follows that $D \cdot D \subseteq D$, so $\mathbb{Z} + D$ is a ring.

Now note that for $0 \leqslant j \leqslant n$ we have

$$\begin{aligned} p_j + q_{n-j} &= \gamma^{j-n}(a_n\gamma^n + \cdots + a_{n-j}\gamma^{n-j}) + \gamma^{j-n}(a_0 + a_1\gamma + \cdots + a_{n-j}\gamma^{n-j}) \\ &= \gamma^{j-n}f(\gamma) + a_{n-j} = a_{n-j}. \end{aligned}$$

In particular $q_{n-j} \in \mathbb{Z} + D$ for all $j \geqslant 0$ and thus $\mathbb{Z} + N \subseteq \mathbb{Z} + D$. By involution $\mathbb{Z} + N = \mathbb{Z} + D$.

(ii) From $D \cdot D \subseteq D$ it follows that D is an ideal of $\mathbb{Z} + D = A$. By involution N is an ideal of $\mathbb{Z} + N = A$. From the identity $p_j + q_{n-j} = a_{n-j}$ for all $j \geqslant 0$ it follows that $a_0, \ldots, a_n \in N + D$. As f is irreducible we have $\mathbb{Z}a_0 + \cdots + \mathbb{Z}a_n = \mathbb{Z}$, so $1 \in N + D$ and N and D are coprime. It follows from $\gamma p_j = p_{j+1} - a_{n-j-1} = q_{n-j-1}$ for all $0 \leqslant j < n$ that $\gamma D = N$. Hence $D \cdot (A + A\gamma) = D + N = A$, so D is invertible and so is N by symmetry. It then follows that $\gamma A = N \cdot D^{-1} = N : D$. The map $\mathbb{Z} \to A/D$ is surjective and has kernel $\mathbb{Z} \cap D$. It follows from the definition of the p_j that $\mathbb{Z} \cap D = \mathbb{Z}p_0 = a_n\mathbb{Z}$, proving $\mathbb{Z}/a_n\mathbb{Z} \cong A/D$. By symmetry $\mathbb{Z}/a_0\mathbb{Z} \cong A/N$.

(iii) It follows from $\gamma p_j = p_{j+1} + a_{n-j-1} \in D$ that $\gamma D \subseteq D$ and thus $MD \subseteq D$. As $p_j \in M$ for all j we have $D \subseteq M$. Hence $AM = M + DM \subseteq M$ and $A \subseteq M : M$. We have $M : M \subseteq M \subseteq \mathbb{Z}[\gamma]$ since $1 \in M$. As $M : M = (\gamma^{1-n}M) : (\gamma^{1-n}M)$ and $\gamma^{1-n}M = \mathbb{Z} + \mathbb{Z}\gamma^{-1} + \cdots + \mathbb{Z}(\gamma^{-1})^{n-1}$ we get $M : M \subseteq \mathbb{Z}[\gamma^{-1}]$ by symmetry. We conclude $A \subseteq M : M \subseteq \mathbb{Z}[\gamma] \cap \mathbb{Z}[\gamma^{-1}]$.

Now suppose there exists some $x \in \mathbb{Z}[\gamma] \cap \mathbb{Z}[\gamma^{-1}]$ such that $x \notin A$, and let $x = \sum_{i=0}^{B} b_i\gamma^i = \sum_{i=0}^{C} c_i\gamma^{-i}$ be such an element for which B is minimal. Then $g = X^C(\sum_{i=0}^{B} b_iX^i - \sum_{i=0}^{C} c_iX^{-i}) \in \mathbb{Z}[X]$. As $g(\gamma) = \gamma^C(x - x) = 0$ we conclude $g \in f\mathbb{Q}[X]$ since f is irreducible. Hence $g = fh$ for some $h \in \mathbb{Z}[X]$ by Exercise 6.33. It follows that $a_n \mid b_B$, so we may subtract $(b_B/a_n)p_B \in A \subseteq \mathbb{Z}[\gamma] \cap \mathbb{Z}[\gamma^{-1}]$ from x, violating minimality of B. From the contradiction, we obtain $A = M : M = \mathbb{Z}[\gamma] \cap \mathbb{Z}[\gamma^{-1}]$. Note that $(\mathbb{Z} + \gamma\mathbb{Z})^{n-1} = M$, so $\mathrm{Bl}(\mathbb{Z} + \gamma\mathbb{Z}) = (\mathbb{Z} + \gamma\mathbb{Z})^{n-1} : (\mathbb{Z} + \gamma\mathbb{Z})^{n-1} = M : M$ by Theorem 6.26 □

Example 6.35. Let K be a number field of degree n, R a subring of K, and $\alpha, \beta \in K^*$. We will give an explicit expression for $\mathrm{Bl}(R\alpha + R\beta)$. This is particularly interesting given the fact that if $\mathbb{Q}R = K$, every invertible R-ideal can be generated by two elements (Exercise 6.37). With $\gamma = \alpha/\beta$ we may write

$$\mathrm{Bl}(\alpha R + \beta R) = \mathrm{Bl}(R + \gamma R) = R \cdot \mathrm{Bl}(\mathbb{Z} + \gamma\mathbb{Z}) = R + p_1R + \cdots + p_{n-1}R,$$

where the p_i are as in Theorem 6.34.

Exercise 6.36. Let K be a number field and $\alpha_1, \ldots, \alpha_n \in K^*$. Consider $I = \alpha_1\mathbb{Z} + \cdots + \alpha_n\mathbb{Z} \subseteq K$ and $R_i = \mathbb{Z}[\alpha_1/\alpha_i, \alpha_2/\alpha_i, \ldots, \alpha_n/\alpha_i]$. Show that $\mathrm{Bl}(I) = \bigcap_i R_i$.

Exercise 6.37. Let R be a number ring.

a. Show that every invertible ideal is generated by two elements.

b. It follows from (a) that, if R is a Dedekind domain, then every ideal is generated by two elements. Show that the converse fails for $R = \mathbb{Z}[\sqrt{3}]$.

c. Show that if every ideal of R generated by two elements is invertible, then every ideal is invertible.

d. (*) Can you give a global proof of (c)?

6.3. Multiplicative inversion

Definition 6.38. Let R be a number ring in a number field K and let X be a set of fractional ideals of R. We write $C(X)$, the *closure* of X, for the set of ideals of the form $I : J$, where I and J are sums of products of ideals of X. Write $B(X) = \max\{\mathfrak{a} \in C(X) : \mathfrak{a} \subsetneq R\}$ for the maximal elements with respect to inclusion.

Lemma 6.39. *Let R be a number ring and X a set of fractional ideals such that $C(X) \subseteq \mathcal{I}(R)$. Then $C(X)$ is closed under addition, multiplication and division. If $X \subseteq Y$ is some set of ideals closed under addition, multiplication and division, then $C(X) \subseteq Y$.*

Proof. For $(I : J), (I' : J') \in C(X)$ all of $(I : J) + (I' : J') = (IJ' + I'J) : (JJ')$ and $(I : J) \cdot (I' : J') = (II') : (JJ')$ and $(I : J) : (I' : J') = (IJ' : I'J)$ are in $C(X)$. The second statement follows trivially from the definition of $C(X)$. □

Exercise 6.40. Let R be a number ring and $X \subseteq \mathcal{I}(R)$ be a subset that is closed under addition, multiplication and division. Show that X is closed under intersection.

Exercise 6.41. Let R be a number ring and let X be a set of fractional ideals of R. Show that the elements of $B(X)$ are pairwise coprime and that the multiplication map $\mathbb{Z}^{(B(X))} \to C(X)$ is injective. Give an example where the map is not surjective.

Proposition 6.42. *Let R be a number ring and let X be a set of integral ideals of R. If $C(X) \subseteq \mathcal{I}(R)$, then $C(X)$ is a free group under multiplication with basis $B(X)$.*

Proof. The map $f : \mathbb{Z}^{(B(X))} \to C(X)$ is injective by Exercise 6.41. By Lemma 6.39 the set $C(X)$ is closed under addition and inversion. Hence it suffices to shown using Noetherian induction for each integral $\mathfrak{a} \in C(X)$ that $\mathfrak{a} \in \mathrm{im}(f)$, from which the proposition follows. For $\mathfrak{a} = R$ this is clear. For $\mathfrak{a} \subsetneq R$, consider $\mathfrak{b} \in B(X)$ such that $\mathfrak{a} \subseteq \mathfrak{b}$, which exists by Zorn's lemma. Then $\mathfrak{a} : \mathfrak{b} \in C(X)$ is integral and strictly contains $\mathfrak{a}$. Hence $\mathfrak{a} : \mathfrak{b} \in \mathrm{im}(f)$ by assumption and thus $\mathfrak{a} \in \mathrm{im}(f)$. It follows that f is surjective. □

Definition 6.43. Let R be a subring of a number field K and let X be a set of fractional ideals of R. A *coprime base* for X is a set of integral ideals B of R such that for all $\mathfrak{b}, \mathfrak{c} \in B$ the ideal $\mathfrak{b}$ is invertible and $\mathfrak{b} + \mathfrak{c} = R$, and for all $\mathfrak{a} \in X$ there exist $(k_\mathfrak{b})_\mathfrak{b} \in \mathbb{Z}^{(B)}$ such that $\mathfrak{a} = \prod_{\mathfrak{b} \in B} \mathfrak{b}^{k_\mathfrak{b}}$.

Exercise 6.44. Let R be a number ring and X, Y be sets of fractional R ideals. Show that if $X \subseteq C(Y) \subseteq \mathcal{I}(R)$, then $C(X) \subseteq C(Y)$. Conclude that if X consists of only integral ideals and has a coprime base, then $B(X)$ is a coprime base.

Theorem 6.45. *There exists a polynomial-time algorithm that, given an order* R *and ideals* $\mathfrak{a}_1, \ldots, \mathfrak{a}_m \subseteq R$*, computes the smallest order* $R \subseteq S$ *for which there exists a coprime base of* $S\mathfrak{a}_1, \ldots, S\mathfrak{a}_m$ *and computes* $B(S\mathfrak{a}_1, \ldots, S\mathfrak{a}_m)$*. Equivalently, this* S *is the smallest ring for which* $C(S\mathfrak{a}_1, \ldots, S\mathfrak{a}_m) \subseteq \mathcal{I}(S)$*.*

Algorithm 6.46. Construct a complete simple graph G with m vertices and label the vertices $\mathfrak{a}_1, \ldots, \mathfrak{a}_m$. Set S equal to R. Then, for each i replace S by $\mathrm{Bl}_S(\mathfrak{a}_i)$. While there are edges in G, repeat the following 5 steps:

(1) Choose an edge $\{a, b\}$ of G and let $\mathfrak{a}$ and $\mathfrak{b}$ be the labels of a and b.
(2) Compute $\mathfrak{c} = \mathfrak{a} + \mathfrak{b}$, replace S by $\mathrm{Bl}_S(\mathfrak{c})$ and compute $\mathfrak{c}^{-1} = S : \mathfrak{c}$.
(3) Add a vertex c labeled $\mathfrak{c}$ to G and connect it to a, b and those vertices which are neighbors of both a and b.
(4) Update the labels of a and b to $\mathfrak{a}\mathfrak{c}^{-1}$ and $\mathfrak{b}\mathfrak{c}^{-1}$ respectively.
(5) For each $s \in \{a, b, c\}$, if the label of s is S, then delete s and its incident edges from G.

Writing $V = \{\mathfrak{c}_1, \ldots, \mathfrak{c}_n\}$, the required coprime base is $S\mathfrak{c}_1, \ldots, S\mathfrak{c}_n$. The remaining output can now be computed in polynomial time.

Proof of Theorem 6.45. Most of the proof is analogous to the proof of Theorem 2.7. Write S_{al} for S produced by the algorithm, S_{cl} for the minimal $R \subseteq S$ such that $C(S\mathfrak{a}_1, \ldots, S\mathfrak{a}_m) \subseteq \mathcal{I}(S)$ and S_{cb} for the minimal $R \subseteq S$ such that $S\mathfrak{a}_1, \ldots, S\mathfrak{a}_m$ has a coprime base. It remains to show that $S_{\text{al}} = S_{\text{cl}} = S_{\text{cb}}$.

From the correctness of the algorithm it follows that $S_{\text{cb}} \subseteq S_{\text{al}}$. Inductively, one can show that every ideal in the graph of the algorithm at any step is in $C(T\mathfrak{a}_1, \ldots, T\mathfrak{a}_m)$ for some $R \subseteq T \subseteq S_{\text{cl}}$. Hence $S_{\text{al}} \subseteq S_{\text{cl}}$. Every ideal of $C(S_{\text{cb}}\mathfrak{a}_1, \ldots, S_{\text{cb}}\mathfrak{a}_m)$ is invertible: Since each $S_{\text{cb}}\mathfrak{a}_i$ can be written as a power-product of the coprime base, so can every sum, product and quotient of such ideals, and thus all such ideals are invertible. Hence $S_{\text{cl}} \subseteq S_{\text{cb}}$. We conclude that $S_{\text{al}} = S_{\text{cl}} = S_{\text{cb}}$. □

Theorem 6.47. *There exists a polynomial-time algorithm that, given an order* R *in a number field* K *and a finitely generated subgroup* $X = \langle x_1, \ldots, x_n \rangle \subseteq K^*$*, computes an order* $R \subseteq S$*, minimal with respect to inclusion, such that for all finite subsets* $Y \subseteq X$ *the ideal* SY *is invertible.*

Proof. First compute using Theorem 6.26 the smallest ring S_0 for which all ideals of the form $S_0 + x_i S_0$ and $S_0 + x_i^{-1} S_0$ are invertible. One can then show locally that $D_i = (S_0 + x_i S_0)^{-1}$ and $N_i = (S_0 + x_i^{-1} S_0)^{-1}$ are integral ideals such that $x_i S_0 = N_i : D_i$. Now apply Theorem 6.45 to find the minimal ring $S_0 \subseteq S$ for which $C(\{SN_i, SD_i : i\}) \subseteq \mathcal{I}(S)$. Note that Sx_i for all i, hence Sx for all $x \in X$, and hence SY for all finite subsets $Y \subseteq X$, is in $C(\{SN_i, SD_i : i\})$ and therefore is invertible. It remains to show this S is minimal.

Suppose $R \subseteq T$ is minimal such that TY is invertible for all finite $Y \subseteq X$. Equivalently T is minimal such that $C(\{x_i T : i\}) \subseteq \mathcal{I}(T)$. Clearly $S_0 \subseteq T$ since $T + x_i T$ and $T + x_i^{-1} T$ have to be invertible. We have $TN_i, TD_i \in C(\{Tx_i : i\})$ for all i by definition of N_i and D_i. Thus $C(\{TN_i, TD_i : i\}) \subseteq C(\{Tx_i : i\}) \subseteq \mathcal{I}(T)$ by Exercise 6.44. Hence $S \subseteq T$. □

6.4. Unit products The coprime base algorithm for ideals gives us algorithms for the following decision problems.

Theorem 6.48. *There exists a polynomial time algorithm that, given an order R in a number field K, fractional ideals $\mathfrak{a}_1, \ldots, \mathfrak{a}_m$ and integers $k_1, \ldots, k_m \in \mathbb{Z}$, decides whether there exists an order $R \subseteq S \subseteq K$ such that $S \prod_i \mathfrak{a}_i^{k_i} = S$.*

Proof. Write $\mathfrak{a}_i$ as a quotient of integral ideals $N_i : D_i$ for all i. Compute an order $R \subseteq S$ and a coprime base $\mathfrak{c}_1, \ldots, \mathfrak{c}_n$ for $SN_1, \ldots, SN_m, SD_1, \ldots, SD_m$, as well as constants such that $N_i = \prod_j \mathfrak{c}_j^{n_{ij}}$ and $D_i = \prod_j \mathfrak{c}_j^{d_{ij}}$. Then return whether $\sum_i k_i \sum_j (n_{ij} - d_{ij}) = 0$. □

Corollary 6.49. *There exists a polynomial time algorithm that, given elements $\alpha_1, \ldots, \alpha_m$ in a number field K and integers $k_1, \ldots, k_m \in \mathbb{Z}$, decides whether $\prod_i \alpha_i^{k_i} \in \mathcal{O}_K^*$.*

Proof. Pick any order R of K and apply Theorem 6.48 to $\alpha_1 R, \ldots, \alpha_m R$. □

We can consider Corollary 6.49 as a weak generalization of Theorem 2.1 to number fields. The main deficiency of this corollary is that does not decide whether the product equals 1 but whether it equals a unit. Even when applied to $\mathbb{Q}$ the corollary is weaker.

7. Tame orders

In this section we will define a class of orders which has nice algorithmic properties.

Definition 7.1. For an order R of degree n write $\delta'(R)$ for the maximal divisor of $\delta(R)$ coprime to all primes $p \leqslant n$.

Observe that we can compute $\delta(R)$ and hence $\delta'(R)$ in polynomial time.

Definition 7.2. An order R in a number field K of degree n is *tame* if the following holds:

(1) For all primes $p \leqslant n$ it holds that $p \nmid \#(\mathcal{O}_K/R)$;
(2) The R-ideal $R^\dagger$ is invertible;
(3) The $\mathbb{Z}/\delta'(R)\mathbb{Z}$-module $(R^\dagger/R)/\delta'(R)(R^\dagger/R)$ is projective.

Lemma 7.3. *For any number field K the order $\mathcal{O}_K$ is tame.*

Proof. Conditions 1 and 2 are trivially satisfied. For condition 3 it is sufficient to show that $\delta(\mathcal{O}_K)$, and hence q, is square-free. This follows from Exercise 5.24. □

Lemma 7.4. *Let R be a tame order in a number field K and p a prime. Then $p \mid \#(\mathcal{O}_K/R)$ if and only if $p^2 \mid \delta'(R)$. Moreover, $R = \mathcal{O}_K$ if and only if $\delta'(R)$ is square-free.*

Proof. This follows from condition 1 combined with Theorem 5.23. □

Lemma 7.5. *There exists a polynomial-time algorithm that, given an order R and a square-free integer $d > 0$, computes an order $R \subseteq R'$ such that all prime ideals $\mathfrak{p} \subseteq R'$ with $d \in \mathfrak{p}$ are invertible.*

Proof. We have $\operatorname{rad}(R/dR) = \prod_{\mathfrak{p} \in \operatorname{spec}(R/dR)} \mathfrak{p}$, and its inverse image in R is $\mathfrak{a} = \prod_{\mathfrak{p} \in \operatorname{spec} R,\, d \in \mathfrak{p}} \mathfrak{p}$. Then $\mathfrak{a}$ is invertible in $\operatorname{Bl}(\mathfrak{a})$, and so is each prime factor. We may compute $\operatorname{rad}(R/dR)$ using Theorem 5.16, from which we easily compute $\mathfrak{a}$. Using Corollary 6.27 we compute $R' = \operatorname{Bl}(\mathfrak{a})$. □

Lemma 7.6. *There exists a polynomial-time algorithm that, given a finite commutative ring R and an R-module M, computes an isomorphism $R^n \to M$ for some $n \geqslant 0$ if it exists, or computes an ideal $0 \subsetneq I \subsetneq R$.*

Proof. If $M = 0$, then $R^0 \to M$ gives an isomorphism. Suppose $M \neq 0$. Compute some non-zero $m \in M$ and the homomorphism $\lambda : R \to M$ given by $r \mapsto rm$. If $\ker(\lambda) \neq 0$ we may return $I = \ker(\lambda)$ and we are done. Thus suppose λ is injective. Then we apply this algorithm to $M' = M/\operatorname{im}(\lambda)$ recursively, which has strictly smaller cardinality. Hence we either obtain some non-trivial ideal I of R, in which case we are done, or some isomorphism $R^n \to M'$. It lifts to some map $R^n \to M$, and the sum with λ induces a map $R^{n+1} \to M$. This gives two exact sequences and compatible maps down of which the outer ones are isomorphisms:

$$\begin{array}{ccccccccc} 0 & \longrightarrow & R & \longrightarrow & R^{n+1} & \longrightarrow & R^n & \longrightarrow & 0 \\ & & \downarrow & & \downarrow & & \downarrow & & \\ 0 & \longrightarrow & \operatorname{im}(\lambda) & \longrightarrow & M & \longrightarrow & M' & \longrightarrow & 0 \end{array}$$

Hence $R^{n+1} \to M$ is an isomorphism. Note that the algorithm requires at most $\#M/\#R$ recursive steps, and therefore runs in polynomial time. □

Note that the algorithm of Lemma 7.6 is not functorial. If $R = M = \mathbb{Z}/6\mathbb{Z}$ and $m \in M$ is chosen to be equal to 1, then the isomorphism $R \to M$ is found, but if $m = 2$ we return the ideal $3\mathbb{Z}/6\mathbb{Z}$ instead.

Lemma 7.7. *There exists a* functorial *polynomial-time algorithm that, given an integer $q > 0$ and a $\mathbb{Z}/q\mathbb{Z}$-module M, decides whether M is projective and if not computes an integer $0 < d < q$ such that $\operatorname{rad}(q) \mid d \mid q$.*

Proof. Compute $d = \text{rex}(M)$ as defined in Exercise 3.43. Write $M \cong \prod_i (\mathbb{Z}/d_i\mathbb{Z})$. It holds that $d = q \Leftrightarrow$ for all primes p we have $\gcd(\text{ord}_p d_1, \ldots, \text{ord}_p d_k) = \text{ord}_p q$ $\Leftrightarrow$ for all i and primes p we have $\text{ord}_p d_i \in \{0, \text{ord}_p q\} \Leftrightarrow M$ is locally free. Hence if $d = q$, then M is projective, and otherwise it is not projective and $\text{rad}(q) \mid d \mid q$ with $0 < d < q$. □

Algorithm 7.8. Let R be an order in a number field K.

(1) Let $q = \delta'(R)$ and $S = R$.
(2) Compute $\mathfrak{a} = (\frac{1}{q}S) \cap S^\dagger$. If $\mathfrak{a}$ is not invertible, update S to $\text{Bl}(\mathfrak{a})$ and q to $\text{rex}(q, \delta'(S))$, and repeat step 2.
(3) Using Lemma 7.7 test if $\mathfrak{a}/S$ is projective over $\mathbb{Z}/q\mathbb{Z}$. If not, obtain some integer $0 < d < q$ such that $\text{rad}(q) \mid d \mid q$, update q to d and go to step 2.
(4) Do as in step 3 with $S/\mathfrak{a}^{-1}$ in the place of $\mathfrak{a}/S$.

Proposition 7.9. *There exists a functorial polynomial-time algorithm that, given an order* R *in a number field* K*, computes an order* $R \subseteq R'$ *such that* R' *is tame.*

References

[1] M. Agrawal, N. Kayal, and N. Saxena, *PRIMES is in P*, Ann. of Math. **160** (2004), no. 2, 781–793. ←31

[2] D. J. Bernstein, *Factoring into coprimes in essentially linear time*, Journal of Algorithms **54** (2005), no. 1, 1–30. ←9

[3] I. Ciocanea Teodorescu, *Algorithms for finite rings*, Ph.D. Thesis, 2016. ←12

[4] J. E. Hopcroft, R. Motwani, and J. D. Ullman, *Introduction to automata theory, languages and computations*, 3rd ed., Pearson, 2006. ←2

[5] H. W. Lenstra Jr., *Computing jacobi symbols in algebraic number fields*, Nieuw Archief voor Wiskunde **14** (1995), no. 3, 421–426. ←24

[6] H. W. Lenstra Jr., *Lattices*, Algorithmic Number Theory **44** (2008), 127–181. ←17

[7] H. W. Lenstra Jr. and C. Pomerance, *Primality testing with gaussian periods*, J. Eur. Math. Soc. **21** (2019), no. 4, 1229–1269. ←31

[8] P. Q. Nguyen and B. Vallée (eds.), *The lll algorithm, survey and applications*, Springer, 2010. ←17

Leiden University, Niels Bohrweg 1, 2333 CA Leiden
Email address: `d.m.h.van.gent@math.leidenuniv.nl`