

Analysis and implementation of algorithms in number theory

Preliminary Arizona Winter School, 2025

Juanita Duque-Rosero

Department of Mathematics and Statistics, Boston University

juanita@bu.edu

Introduction

These lecture notes accompany the lectures for the Preliminary Arizona Winter School 2025: Algorithms in number theory. The main references I used are [Coh93, Har21]. Other useful references are [Coh00, Ste, Voi21, vzGG13].

The course consists of an exploration of the algorithms and computational ideas that power modern algebra and number theory. We will start with the basics: analyzing what makes an algorithm *efficient*, and working through classic methods in integer arithmetic and linear algebra. These techniques will come up again and again in the rest of the lectures.

From there, we will study algebraic numbers and number fields. We will see how to represent them and do arithmetic with them. We then move on to working with rings of integers, discriminants, and integral bases. Then, we go back to algorithmic linear algebra to look at the LLL Algorithm and its applications to the study of number fields. Finally, we study ideals, class groups, and units and we pull everything together with examples from imaginary quadratic fields.

Explicit computations

The ideal way to follow these notes is to try examples in your favorite computer algebra system. I personally use [Magma](#), but you are welcome to use whatever you prefer ([SageMath](#), [PARI/GP](#), [Oscar](#), etc.). You can find [Magma](#) examples [here](#). Also, [here](#) is a list of random tricks that I have compiled and [here](#) is a scavenger hunt to get you started.

Prerequisites

This course will assume fluency with algebra at a beginning graduate level and familiarity with the basic objects of algebraic number theory (such as number fields and their rings of integers). Some good references are [Mil, Lan94].

A note from the author

Despite my best efforts, these notes will contain typos. If you spot any, please feel free to email me, I appreciate your help!

An (irrelevant) note from the author

When I think about foundational algorithms that have really made a difference in my own number theory research, *Gröbner bases* are on top of my list. They are very useful for arithmetic geometry, and they have “saved” my work more than once. Unfortunately, I could not find space in these lecture notes to cover them. If you are curious, the book [CLO15] is beautifully written, and has a lot of the relevant theory.

Acknowledgments

I thank John Voight for insightful discussions while I was preparing these notes. Also, thank you to David Harvey for his helpful feedback on an earlier draft. Thanks to Jennifer Balakrishnan, Jerson Caro, Aashraya Jha, and Padmavathi Srinivasan for their valuable input. Finally, thanks to the problem session leaders Thomas Bouchet, Kate Finnerty, Asimina S. Hamakiotes, and Yongyuan Huang for the careful reading and feedback.

Lecture 1: Arithmetic and linear algebra

Even when you are doing advanced computation, you will end up using basic algorithms in arithmetic and linear algebra. In this first lecture, we explore some of those algorithms to compute basic arithmetic, greatest common divisors, and matrices normal forms. They will be useful in the rest of the course. This lecture follows [Har21, §2] and [vzGG13, Chapter 2]. Other useful references are [BZ11] and [Coh93].

1.1 Analysis of algorithms

When analyzing the effectiveness of an algorithm, we can consider many factors, such as the amount of memory used, or the number of operations required for completion. We start by introducing some useful notation.

1.1.1 Big-O notation

Definition 1.1. Let $f(n)$ and $g(n)$ be functions defined on the natural numbers. We say that $f(n)$ is *big-O* of $g(n)$, and write

$$f(n) = O(g(n)),$$

if there exist constants $C > 0$ and $n_0 \in \mathbb{N}$ such that

$$|f(n)| \leq C |g(n)|$$

for all $n \geq n_0$.

You can think of $O(g(n))$ as describing a multiplicative bound on the growth of $f(n)$, for large n . That is, $f(n)$ does not grow faster than a constant multiple of $g(n)$ for sufficiently large inputs.

Lemma 1.2. Let $f(n)$, $g(n)$, $a(n)$, and $b(n)$ be functions satisfying $f(n) = O(g(n))$ and $a(n) = O(b(n))$. Then

$$f(n) + a(n) = O(\max(|g(n)|, |b(n)|))$$

and

$$f(n)a(n) = O(g(n)b(n)).$$

Exercise 1.3. Prove Lemma 1.2.

In this lecture, we will use the term *input size* for an algorithm, which depends on how the data is represented. Precise analysis sometimes requires care in defining what counts as a single operation (for example, adding two numbers versus multiplying two numbers), but for most algorithms in arithmetic and linear algebra we will focus on basic operations such as integer addition, multiplication, and comparisons.

Lemma 1.4. *Let $f(n)$ be a function defined over \mathbb{N} . Then $f(n) = O(\log(n))$ if and only if $f(n) = O(\log_k(n))$ for any $k > 1$. In this case, we write $O(\log(n)) = O(\log_k(n))$.*

Proof. We recall the relation between logarithms: $\log(n) = \log_k(n)/\log_k(10)$. This shows that $f(n) = O(\log(n))$ if and only if there is a constant C such that for all large enough n ,

$$|f(n)| \leq C|\log(n)| = C \frac{|\log_k(n)|}{|\log_k(10)|} = \frac{C}{\log_k(10)} |\log_k(n)| = C' |\log_k(n)|.$$

This shows that $f(n) = O(\log(n))$ if and only if $f(n) = O(\log_k(n))$. □

1.1.2 Computational models

When deciding on the complexity of an algorithm, it is of vital importance to decide which computational model we are considering. That is, setting a formal framework for what it means to “compute”. We have many choices available, for example, Turing machine, Random Access Machine, or even quantum computer. Each choice makes some analysis cleaner or more awkward. One good reference to learn about this is [Pap94]. In this course, we follow the choice of [Har21] and use the deterministic multitape Turing model. Even though this will not be the focus of this course, we give an intuition of what this model does.

(Deterministic) Turing machines

We can think of a (deterministic) Turing machine as an infinite tape together with a head. The head moves along the tape and can read and modify the contents of each position. For a given Turing machine, you need to pick the alphabet (possible symbols in the tape), a set of states (what is the machine doing at a given time), and a function describing the behavior of the machine given the state and symbol.

Example 1.5. Just for fun, we can consider a very simple Turing machine that adds $4 + 3$. We represent the numbers $4 = ||||$ and $3 = |||$. The head always starts at the start symbol \triangleright . It moves right until it encounters the symbol \circ . Then, it deletes \circ , changes its state to “adding”, and moves to the right.




Figure 1.1: A Turing machine to add $4 + 3$.

While in this state, whenever we encounter a symbol $|$, we delete it, move to the left, add a symbol $|$, and move twice to the right. When the head reads a blank at this state, the new state is “halt” since we are done with the addition. You can see a picture of this in Figure 1.1. You can check that the total number of steps this machine takes to get to “halt” is 15.

Exercise 1.6. Can you describe a Turing machine that adds $4 + 3$ using less steps than in Example 1.5 but using the same alphabet $\{\triangleright, \square, \circ, |\}$?

Exercise 1.7. Can you change the alphabet and describe a Turing machine that adds $4 + 3$ using less steps than in Example 1.5?

Deterministic multitape Turing machines

Instead of working with one tape in a Turing machine, we can consider the case when we have finitely many tapes and one head that reads one position on each tape. It turns out that multitape Turing machines are as capable as Turing machines, but faster. For instance, one can describe a multitape Turing machine that performs the operation of Example 1.5 in 9 less steps!

Definition 1.8. The complexity of a multitape Turing machine model refers to the number of steps executed by the machine over the course of a computation. Each step that a Turing machine takes is called a **bit operation**.

Remark 1.9. The complexity of an algorithm depends on what the *alphabet* of the multitape Turing machine is. It is not the same to have a machine that reads any integer, to a machine that only reads $|$ and \circ , so we need to be specific about the alphabet.

In practice, we describe the complexity of Turing Machines by writing the number of steps in using Big- O notation for functions on the size (number of bits) of the input. We also note that another useful thing to consider is the space complexity, i.e., the amount of memory used by a computation. We will only focus on the time complexity, i.e. the number of steps that it takes to terminate. For certain algorithms the space complexity becomes the main bottleneck in practice, so it is good to remember that this might be a problem.

Definition 1.10. Given an algorithm that takes an input of n -bits and requires at most $O(f(n))$ bit operations to complete, we say the algorithm runs in **time** $O(f(n))$ or has **complexity** $O(f(n))$.

Remark 1.11. Because we want to access the number theory and not the computational complexity, through these lectures, we will not be very specific about the particular construction of Turing machines for each algorithm.

1.1.3 An example: addition of integers.

The first question we need to answer is how to represent the objects that we want to input, or what is the alphabet of our Turing machine. In the next section, we discuss representing integers more generally; here, we assume integers are given in binary.

Example 1.12. The integer 431 is represented by the 9-digit binary number **110101111** since

$$431 = 1 \cdot 2^8 + 1 \cdot 2^7 + 0 \cdot 2^6 + 1 \cdot 2^5 + 0 \cdot 2^4 + 1 \cdot 2^3 + 1 \cdot 2^2 + 1 \cdot 2^1 + 1 \cdot 2^0.$$

Now we are ready to add two integers together. Without loss of generality, we can assume that the length of the expansions is n (we can always pick the maximum length, and then add zeros to the shortest one). Let's look at Algorithm 1.13.

Algorithm 1.13 (Integer addition [HvdH21, Algorithm 2.1]).

The input is two binary expansions $a = (a_{n-1}, \dots, a_1, a_0)$ and $b = (b_{n-1}, \dots, b_1, b_0)$ of length n . This algorithm outputs the binary expansion for the integer c that is the sum $a + b$.

1. Set $\gamma_0 := 0$.
2. For $i = 0, \dots, n-1$ do
 - (a) Set $c_i := a_i + b_i + \gamma_i$.
 - (b) If $c_i \geq 2$, set $c_i := c_i - 2$ and $\gamma_{i+1} := 1$; otherwise, set $\gamma_{i+1} := 0$.
3. Set $c_n := \gamma_n$.

Return $c = (c_n, \dots, c_1, c_0)$.

Remark 1.14. The description of the algorithm does not explicitly give a multitape Turing machine (§1.1.2), but you should convince yourself that it gives you all the information you need to rigorously define the machine.

Theorem 1.15. *The complexity of Algorithm 1.13 is $O(n)$.*

Proof. We analyze the complexity step by step. It is useful to recall arithmetic of big- O notation from Lemma 1.2. Step 1 requires one bit operation, so its time is $O(1)$. Step 2 iterates n times. Each iteration performs two bit addition operations, potentially subtracts 2, and sets the carry bit. Each of these operations is a constant-time word operation, so each iteration takes $O(1)$ time. In total, the complexity of Step 2 is $O(n)O(1) = O(n)$. Finally, Step 3 sets one bit, which is time $O(1)$. Altogether, the total complexity of the algorithm is

$$O(1) + O(n) + O(1) = O(n).$$

□

Remark 1.16. The constant n in Theorem 1.15 denotes the length of the inputs, so if a and b are the integers we want to add, then

$$n = \log_2(\max\{a, b\}) = O(\log(\max\{a, b\})),$$

where the last equality follows from Lemma 1.4.

1.2 Integer arithmetic

With basic notation established, we now explore algorithms for integer arithmetic. First, how do we represent integers? Then we move on to addition, multiplication, division, and greatest common divisors.

1.2.1 Representing integers

We will represent numbers in binary, with another bit to represent the sign.

That could be the only line of this subsection. The following is a small detour relates to expressing integers in actual computers.

Modern laptops use a 64-bit processor. That means that each integer can be up to $2^{64} - 1$ in value (unsigned), or from -2^{63} to $2^{63} - 1$ (signed). The CPU can process (add, multiply, etc.) two of these 64-bit numbers in one operation.

We can represent larger integers by an array of 64-bit words as follows

$$a = (-1)^s \sum_{i=0}^n a_i 2^{64i}, \quad (1.17)$$

where $s \in \{0, 1\}$, $0 \leq n + 1 < 2^{63}$, and $a_i \in \{0, \dots, 2^{64} - 1\}$. The numbers a_i are the *digits in base 2^{64}* of a .

Definition 1.18. For an integer a , its **standard representation** is given by the array

$$(s \cdot 2^{63} + n + 1, a_0, \dots, a_n),$$

where s and a_i are as in (1.17) and a_n is nonzero if $a \neq 0$. If the standard representation of a has length n we call a an n -bit integer.

1.2.2 Addition

To add integers using their binary representation, we can just use Algorithm 1.13 that has complexity $O(n)$, where n is the number of bits of the integers (in binary). We usually just pick the largest number of bits and set that as n . By Remark 1.16, the complexity of addition is the same as $O(\log(\max(a, b)))$, where a and b are the integers we want to add.

If you went on the detour about the standard representation, you can modify Step 2 (b) of Algorithm 1.13 by

$$\text{If } c_i \geq 2^{64}, \text{ then set } c_i = c_i - 2^{64} \text{ and } \gamma_{i+1} = 1.$$

1.2.3 Multiplication

With addition, we noted that the naive algorithm to add integers (Algorithm 1.13) has complexity $O(n)$, where n is the number of bits. The naive algorithm that we use to multiply integers runs in time $O(n^2)$. This is usually fine for smaller integers, but more efficiency is needed for larger integers.

Exercise 1.19. You can check the time that it takes Magma to run one line by writing `time` at the beginning of the line:

```
> time 2^115032204*3^473444585;
Time: 312.990
```

Can you find two large integers (but maybe not as large as above) for which multiplication takes longer than 0 seconds? How big are your integers? How long does it take to add them?

Exercise 1.20. Write and analyze an algorithm that implements naive multiplication for integers. Your algorithm should use $O(mn)$ word operations, where m and n are the number of bits of the integers you are multiplying.

The theoretical state of the art is the Algorithm presented in [HvdH21]. This algorithm has complexity $O(n \log n)$ and is believed to be optimal but maybe not practical.

In practice, many computer algebra systems use the **GMP library**, which is a “free library for arbitrary precision arithmetic”. GMP implements an algorithm whose theoretical complexity comes very close to $O(n \log n)$ [Har21, Remark 2.12].

1.2.4 Division

Division of large integers is typically accomplished using algorithms based on repeated subtraction, long division, or more advanced methods such as Newton-Raphson iteration for reciprocal approximation. For most practical purposes, the classical long division algorithm suffices.

In general, since \mathbb{Z} is an Euclidean domain, given $a, b \in \mathbb{Z}$, the division algorithm to divide a by b should return the unique integers q and r with $0 \leq r < b$ such that $a = qb + r$. We call r the **remainder** of dividing a by b , or a modulo b , and denote it as $\text{rem}(a, b)$.

Exercise 1.21. Describe the classical long division algorithm for binary integers. Prove that the algorithm takes time $O((n - m)m)$, where the integers have m and n bits, respectively.

Just like with multiplication, we can improve this bound. A division algorithm that combines fast multiplication with Newton’s method gives the same complexity as multiplication: $O(n \log n)$.

1.2.5 Greatest common divisors

The greatest common divisor (gcd) of two integers a and b , denoted $\text{gcd}(a, b)$, is the largest integer that divides both. The standard method to compute gcd ’s is the Euclidean Algorithm (Algorithm 1.22).

Algorithm 1.22 (Euclidean Algorithm for gcd).

Given integers $a \geq b > 0$, compute $g := \gcd(a, b)$.

1. Set $r_0 := a$ and $r_1 := b$.
2. Set $i := 1$ and while $r_i \neq 0$, do the following
 - (a) Set $r_{i+1} := \text{rem}(r_{i-1}, r_i)$ and $i := i + 1$.

Return r_{i-1} .

Exercise 1.23. Explain why Algorithm 1.22 terminates and correctly computes the greatest common divisor.

Theorem 1.24 ([vzGG13, Theorem 3.13]). *Algorithm 1.22 for positive n -bit and m -bit integers has complexity $O(mn)$.*

Exercise 1.25. Use Exercises 1.20 and 1.21 to prove Theorem 1.24.

Remark 1.26. As you might know, an essential property of the greatest common divisor is that it corresponds to the smallest positive linear combination of a and b . That is, there exist $x, y \in \mathbb{Z}$ such that $\gcd(a, b) = ax + by$. Running the Euclidean Algorithm for gcd (Algorithm 1.22) is almost enough for computing x and y , we just need to “undo” the operations. An optimized version of this algorithm runs in time $O(n \log^2 n)$, where n is again the number of bits of a and b .

Remark 1.27. In particular, Remark 1.26 implies that we can find the inverse of an n -bit integer in $\mathbb{Z}/M\mathbb{Z}$ for $M \geq 2$ in time $O(n \log^2 n)$.

1.2.6 Large powers

We will look at a basic (but useful) algorithm for computing powers g^k , in the general case when g is an element of any group G and k is an integer.

Algorithm 1.28 (Exponentiation algorithm [Coh93, Algorithm 1.2.1]).

The input is an element g of a multiplicative group G and an integer k . This algorithm computes g^k in G .

1. Set $y := 1_G$. If $k = 0$, output y and terminate. If $k < 0$, let $K := -k$ and $z := g^{-1}$. Otherwise, set $K := k$ and $z := g$.
 2. If K is odd set $y := z \cdot y$.
 3. Set $K := \lfloor K/2 \rfloor$. If $K = 0$, output y as the answer and terminate. Otherwise, set $z = z \cdot z$ and go to Step 2.
-

Exercise 1.29. Prove that Algorithm 1.28 computes g^k using $O(\log |k|)$ group multiplications. In particular, when g is an integer, prove that the algorithm runs in $O(n \log n)$ time, where n is the number of bits of the input g .

1.2.7 Summary

We end this section with a summary of results on integer arithmetic in Table 1.1.

Operation (of n -bit integers)	Naive Algorithm	Optimized Algorithm
Addition	$O(n)$	$O(n)$
Multiplication	$O(n^2)$	$O(n \log n)$
Division	$O(n^2)$	$O(n \log n)$
GCD	$O(n^2)$	$O(n \log^2 n)$
Exponentiation (to k -bit integer)	$O(n^2)$	$O(n \log n)$

Table 1.1: Summary of complexity for basic arithmetic algorithms.

1.3 More arithmetic

We will use what we learned about integer arithmetic to study the complexity of modular and polynomial arithmetic.

1.3.1 Modular arithmetic

Let $M \geq 2$. Elements of $\mathbb{Z}/M\mathbb{Z}$ can be represented as integers $x \in \{0, \dots, M-1\}$. In particular, elements of $\mathbb{Z}/M\mathbb{Z}$ occupy at most $\log(M)$ bits of space. We describe the basic arithmetic operations in this ring and record a summary in Table 1.2.

- To add two elements of $\mathbb{Z}/M\mathbb{Z}$, we can add their representatives and subtract M if the result is $\geq M$. This algorithm has complexity $O(\log(M))$.
- To multiply, we multiply the representatives and then take the remainder modulo M , so the complexity of multiplication is $O(\log M \log \log M)$.
- Division is achieved by running the Euclidean Algorithm as in Remark 1.27 to invert the denominator (complexity $O(\log M (\log \log M)^2)$), and then multiplying by the numerator (complexity $O(\log M \log \log M)$). In total, the complexity of division is $O(\log M (\log \log M)^2)$.
- For exponentiation, one can use Algorithm 1.28 in time $O(\log M (\log \log M) \log k)$.

Remark 1.30. One can estimate the time of performing any arithmetic operation over $\mathbb{Z}/M\mathbb{Z}$ by $O(\log^{1+\epsilon} M)$. This is sometimes enough information for complexity computations.

Operation over $\mathbb{Z}/M\mathbb{Z}$	Running Time
Addition	$O(\log M)$
Multiplication	$O(\log M \log \log M)$
Division	$O(\log M \log^2(\log M))$
Exponentiation (x^k)	$O(\log M \log \log M \log k)$

Table 1.2: Summary of running times for modular arithmetic.

1.3.2 Polynomial arithmetic

We can use what we have studied about integer arithmetic to determine the complexity of arithmetic in $\mathbb{Z}[x]$ or $\mathbb{Q}[x]$. In general, polynomials in $\mathbb{Z}[x]$ can be represented as polynomials over finite rings $\mathbb{Z}/M\mathbb{Z}$ for large enough M . Polynomials in $\mathbb{Z}/M\mathbb{Z}[x]$ of degree $< n$ can be represented as a sequence of n coefficients. Also, note that each coefficient can be represented using $O(\log M)$ bits, so to encode a polynomial in $\mathbb{Z}/M\mathbb{Z}[x]$ of degree $< n$, we need space $O(n \log M)$. The running times for algorithms for polynomial multiplication are related to the ones for integer multiplication. Let $M_{\text{int}}(n)$ denote the cost of multiplying n -bit integers. Then, the cost of polynomial arithmetic is summarized in Table 1.3. The interested reader can look at [Har21, §2].

Operation over $\mathbb{Z}/M\mathbb{Z}[x]$ of deg n	Running Time
Addition	$O(n \log M)$
Multiplication	$O(M_{\text{int}}(n \log(nM)))$
Division	$O(M_{\text{int}}(n \log(nM)))$

Table 1.3: Summary of running times for polynomial arithmetic.

1.4 Linear algebra

Now that we have looked at basic arithmetic, we move on to linear algebra, an area that also allows us to make (fast) explicit computations. The complexity of arithmetic operations on matrices depends on the complexity of arithmetic over the base field or ring. As we saw in §1.2 and §1.3, this complexity varies and that is why we focus on the number of ring operations (multiplications/divisions) needed for each algorithm. For basic arithmetic of matrices, one can easily find (upper bounds) for the complexity of adding and multiplying matrices with \mathbb{Z} coefficients, so we leave this as an exercise. We will skip the basic arithmetic and move on to more interesting matrix manipulations.

Exercise 1.31. Compute a function $f(n)$ such that $O(f(n))$ represents the number of field multiplications needed to compute the product of two $n \times n$ matrices. Why can you ignore the number of addition operations?

Remark 1.32. The best known bound for the number of operations (over a field) needed to multiply two $n \times n$ matrices is $O(n^{2.3728596})$ [Har21, Remark 2.5.1].

1.4.1 Gaussian elimination

This is perhaps one of the most useful and widely used algorithms in linear algebra. It gives us a way of solving linear systems, compute determinants, and find inverses and pseudo-inverses, etc. Because of the applications, we focus on square matrices.

The number of multiplications/divisions needed in Algorithm 1.33 is $O(n^3/3)$. Now we can look at a couple applications of the ideas from this algorithm.

Algorithm 1.33 (Gaussian Elimination for square matrices [Coh93, Algorithm 2.2.1]).

The input is an $n \times n$ matrix M with entries in a field and a vector B of length n . This algorithm returns a vector X such that $MX = B$ if M is invertible and **false** if not.

1. Set $j := 0$.
2. Let $j := j + 1$. If $j > n$, then go to Step 6.
3. If $m_{i,j} = 0$ for all $i \geq j$, then return **false** and terminate. Otherwise, let $i \geq j$ be some index such that $m_{i,j} \neq 0$.
4. If $i > j$, for $l = j, \dots, n$ exchange $m_{i,l}$ and $m_{j,l}$ and then exchange b_i and b_j .
5. Note that $m_{j,j} \neq 0$. Set $d := m_{j,j}^{-1}$ and for all $k > j$ set $c_k := dm_{k,j}$. For all $k > j$ and $l > j$ set $m_{k,l} := m_{k,l} - c_k m_{j,l}$. Finally, for $k > j$ set $b_k := b_k - c_k b_j$ and go to Step 2.
6. Note that M is now upper-triangular! For $i = n, n-1, \dots, 1$ set

$$x_i := \left(b_i - \sum_{i < j \leq n} m_{i,j} x_j \right) / m_{i,i},$$

output $X = (x_i)_{1 \leq i \leq n}$, and terminate.

Exercise 1.34. Modify Algorithm 1.33 to compute the inverse of a square matrix. Check how many multiplications/divisions does it take to run your algorithm. Can you get the number of operations to be asymptotic to $4n^3/3$?

Exercise 1.35. Modify Algorithm 1.33 to compute the determinant of a square matrix. Check how many multiplications/divisions does it take to run your algorithm. Can you get the number of operations to be asymptotic to $n^3/3$?

Remark 1.36. Algorithm 1.33 hinges on being able to invert $m_{j,j}$ in Step 5. This is an obstacle for computing determinants of matrices with coefficients in integral domains but not fields (which will be essential in the following lectures). The reader can check [Coh93, Algorithm 2.2.6] for an example of an algorithm to solve this. The algorithm takes $O(n^3)$ operations.

1.4.2 Normal forms and picking a basis

Gaussian elimination allows us to represent a matrix by a similar matrix that is simpler. This is definitely not the only normal form for a matrix. In §5.1.1, we will describe and use the Hermite normal form, which works over the integers \mathbb{Z} , and allows us to represent ideals in orders.

Other normal forms that we will not focus on here, but which are very useful, include the Smith normal form (which helps with module and abelian group structure computations) and the Jordan normal form.

Finding normal forms is just finding a new basis for your space, in which the linear operator represented by the matrix can be written in a simpler way. The last special basis that we will explore in Lecture 4 is the LLL-reduced basis, for which there is a highly efficient algorithm to compute.

Lecture 2: Algebraic numbers and number fields

Number fields are at the heart of number theory. They are a natural setting where arithmetic and algebraic structures meet. You can find them in the study of Diophantine equations, class field theory, Galois representations, etc. In this lecture, we will focus on studying the basics of number fields and algebraic numbers from a computational perspective. Topics are mostly from [Coh93, Chapters 3 and 4]. For further reading, the reader can consult, for example, [Mil], [Lan94], or [Ste].

2.1 Factoring in $\mathbb{Z}[x]$

As a warm-up, we first go back to the arithmetic of polynomials. Throughout these lectures, we will mainly be working with irreducible polynomials over \mathbb{Z} . In this section, we study the difficult problem of factoring polynomials over $\mathbb{Z}[x]$ or certifying that said polynomial is irreducible. Many approaches rely on being able to factor over finite fields and the use of Hensel's lemma. Another way of solving the problem uses the LLL algorithm, which we will study in Lecture 4. In this section, we will investigate a method that uses field operations (which, in practice, gives a faster approach).

Let $f(x)$ be a non-zero polynomial in $\mathbb{Z}[x]$. We may assume that the greatest common divisor of the coefficients of $f(x)$ is 1 by dividing by any common factor; such a polynomial is called **primitive**. By Gauss's Lemma, deciding if $f(x)$ is irreducible over $\mathbb{Q}[x]$ is equivalent to deciding if it is irreducible over $\mathbb{Z}[x]$, so there is nothing to gain by considering operations over \mathbb{Q} .

We sketch the steps to decide if a primitive polynomial $f(x) \in \mathbb{Z}[x]$ is irreducible, or find a nontrivial factorization (for details, see [Coh93, Algorithm 3.5.7]). The method relies on picking a suitable prime number p .

Step 1: reduce to squarefree polynomials

Deciding if $f(x)$ is squarefree reduces to computing $\gcd(f(x), f'(x))$, as shown in Exercise 2.1. If we want to factor the polynomial, then we can then factor $f(x)/\gcd(f(x), f'(x))$, which is squarefree by construction.

Exercise 2.1. Let $f(x) \in \mathbb{Z}[x]$ be a monic polynomial. Show that $f(x)$ is squarefree in $\mathbb{Z}[x]$ if and only if $\gcd(f(x), f'(x)) = 1$.

Step 2: bound the coefficients of the factors

Assume $f(x) \in \mathbb{Z}[x]$ is primitive and squarefree. For integer polynomials $f(x)$ and $g(x)$ with $g(x)|f(x)$, there is an explicit bound for the absolute value of the coefficients of $g(x)$ in terms of the coefficients of $f(x)$ (see [Coh93, Theorem 3.5.1]). This allows us to find $B > 0$

bounding the coefficients of all irreducible factors of $f(x)$ of degree $\leq \deg(f(x))/2$. Let $\ell(f(x))$ be the leading coefficient of $f(x)$. Choose e to be the smallest exponent for which $p^e > 2\ell(\bar{f})B$.

Step 3: find a factorization over a finite field

Let \mathbb{F}_p be the field $\mathbb{Z}/p\mathbb{Z}$. In this step, we find a factorization of $f(x)$ over $\mathbb{F}_p[x]$. There are many algorithms to do this, but we pick one that works well for small p to show the main ideas (see Algorithm 2.3). We will not show that the algorithm is correct, but the interested reader can look at Exercise 2.4. The main idea of this algorithm relies on the following proposition.

Proposition 2.2 ([Coh93, Proposition 3.4.9]). *Let $\bar{f}(x) \in \mathbb{F}_p(x)$ be squarefree and assume that its decomposition into irreducibles is $\bar{f}(x) = \prod_{1 \leq i \leq r} T_i(x)$. The polynomials $T_i(x) \in \mathbb{F}_p[x]$ with $\deg(T_i(x)) < \deg(\bar{f}(x))$ for which for each i with $q \leq i \leq r$ there exists $s_i \in \mathbb{F}_p$ with $T_i(x) \equiv s_i \pmod{\bar{f}(x)}$, are exactly the p^r polynomials $T_i(x)$ such that $\deg(T_i(x)) < \deg(\bar{f}(x))$ and $T_i(x)^p \equiv T_i(x) \pmod{\bar{f}(x)}$.*

Algorithm 2.3 (Berlekamp for small p [Coh93, Algorithm 3.4.10]).

The input is a squarefree polynomial $\bar{f} \in \mathbb{F}_p[x]$ of degree n , this algorithm computes the factorization of $\bar{f}(x)$ into irreducible factors.

1. Compute inductively for $0 \leq k < n$ values $q_{i,k} \in \mathbb{F}_p$ such that

$$x^{pk} \equiv \sum_{0 \leq i < n} q_{i,k} x^i \pmod{\bar{f}(x)}.$$

2. Let $Q := [q_{i,k}]_{i,k}$. Find a basis V_1, \dots, V_r of the kernel $Q - I$ such that V_1 is the column vector $(1, 0, \dots, 0)^t$. Set $E := \{\bar{f}(x)\}$, $k := 1$, and $j := 1$.
 3. If $k = r$, output E as the set of irreducible factors of $\bar{f}(x)$ and terminate. Otherwise, set $j := j + 1$, and let $T(x) := \sum_{0 \leq i < n} (V_j)_i x^i$.
 4. For each polynomial $g(x) \in E$ such that $\deg(g(x)) > 1$ do the following. For each $s \in \mathbb{F}_p$ compute $\gcd(g(x), T(x) - s)$. Let F be the set of such gcd's whose degree is greater than or equal to 1. Set $E := (E - g(x)) \cup F$ and $k := k - 1 + |F|$. If in the course of this computation we reach $k = r$, output E and terminate the algorithm. Otherwise, go to Step 3.
-

Exercise 2.4. Show that Algorithm 2.3 terminates and correctly computes the factorization of \bar{f} into irreducibles. You can follow the following steps.

1. As a warm-up, let $\bar{f}(x) \in \mathbb{F}_p(x)$ be a polynomial of degree n . Show that $\bar{f}(x)$ is irreducible if and only if

- (i) $x^{p^n} \equiv x \pmod{\bar{f}(x)}$; and
 - (ii) for each prime $q|n$, $\gcd(x^{p^{n/q}} - x, \bar{f}(x)) = 1$.
2. Prove Proposition 2.2.
 3. Using the notation of Step 2 of the algorithm, show that any polynomial $T(x)$ in the kernel of $Q - I$ holds that $T(x)^p \equiv T(x) \pmod{\bar{f}(x)}$.
 4. Explain why the dimension of $\ker(Q - I)$ is exactly r and why the column vector $(1, 0, \dots, 0)^t$ belongs to the kernel.
 5. Let $T(x)$ be a polynomial corresponding to a V_j . Explain why the polynomials F from Step 4 of the algorithm correspond to irreducible factors once we have $k = r$.

Step 4: lift the factorization

This step shows a very useful technique when trying to approximate a solution by its residues modulo p . This is done, for example, when working with p -adic integers. The idea is to use Hensel's Lemma, which we recall since it is a fundamental result.

Lemma 2.5 (Hensel's Lemma for integers). *Let $f(x) \in \mathbb{Z}[x]$ and let p be a prime. Suppose there exists $a_0 \in \mathbb{Z}$ such that*

$$f(a_0) \equiv 0 \pmod{p}, \quad \text{and} \quad f'(a_0) \not\equiv 0 \pmod{p}.$$

Then for every $k \geq 1$, there exists an integer a_k such that

$$f(a_k) \equiv 0 \pmod{p^k} \quad \text{and} \quad a_k \equiv a_0 \pmod{p}.$$

There is a similar version for polynomials

Lemma 2.6 (Hensel's Lemma for polynomials). *Let p be a prime and $f(x) \in \mathbb{Z}[x]$. Suppose $f(x) \equiv g_0(x)h_0(x) \pmod{p}$, where $g_0(x), h_0(x) \in \mathbb{Z}[x]$ are monic and coprime modulo p . Then, for each $k \geq 1$, there exist monic polynomials $g_k(x), h_k(x) \in \mathbb{Z}[x]$ such that*

$$f(x) \equiv g_k(x)h_k(x) \pmod{p^k}, \quad g_k(x) \equiv g_0(x) \pmod{p}, \quad h_k(x) \equiv h_0(x) \pmod{p},$$

and $g_k(x), h_k(x)$ remain coprime modulo p .

This last Lemma allows us to lift the factorizations from Step 3 to

$$f(x) \equiv \ell(f)\tilde{T}_1(x) \cdots \tilde{T}_r(x) \pmod{p^e},$$

where e is as in Step 2 and the polynomials $\tilde{T}_i(x)$ are monic.

Step 5: combine multiple factors

We now repeat for every $d \in \{1, \dots, r/2\}$. For every combination of factors $\bar{U} := \tilde{T}_{i_1} \cdots \tilde{T}_{i_d}$, where we take $i_d := 1$ if $d = 1/2r$, compute the unique polynomial $U \in \mathbb{Z}[x]$ such that all the coefficients of U are in $[-1/2p^e, 1/2p^e)$, and satisfying

$$\begin{aligned} U &\equiv \ell(f)\bar{U} \pmod{p^e}, & \text{if } \deg(U) \leq 1/2 \deg(f), \\ U &\equiv f/U \pmod{p^e}, & \text{if } \deg(U) > 1/2 \deg(f). \end{aligned}$$

If U divides $\ell(f)f$ in $\mathbb{Z}[x]$, output the factor $F = U/\gcd(u_i)$, set $f(x) = f(x)/F$, and remove the corresponding \tilde{T}_i from the list of factors modulo p^e . If $d \geq (1/2)r$, terminate the algorithm by outputting $f(x)$.

If we are done looking at all the possible combinations, we have shown that $f(x)$ is irreducible.

Exercise 2.7. Look through the sketch of the steps of the algorithm to factor primitive polynomials $f(x) \in \mathbb{Z}[x]$ and decide what primes are good candidates to run the algorithm on.

Exercise 2.8. You can generate random integer polynomials of degree d with coefficients in $[-b, b]$ in Magma by running the script

```
R<t,y> := PolynomialRing(Integers(),2);
S<x> := PolynomialRing(Integers());
f := Evaluate(Random(d,R,b),[x,1]);
```

Create a polynomial that has 10 random divisors of degree 3. Run the steps of the algorithm with at least two primes and compare results.

2.2 Number fields

We are now ready to embark on our study of algebraic number theory! To establish notation, we first review the basic definitions for number fields.

Definition 2.9. A number field K is a field such that K is a finite-dimensional \mathbb{Q} -vector space.

Definition 2.10. Let K be a number field, the **degree** of K , denoted $[K : \mathbb{Q}]$, is the dimension of K as a \mathbb{Q} -vector space.

A really good place to find examples of number fields is the [number field database of the LMFDB](#). They even have pictures(!), like the one in Figure 2.1.

Also, if you want to work with number fields in **Magma**, you can download the database [Anf.tar.gz](#), “comprising over 2.6 million number fields of degrees between 2 and 9 (inclusive)”.

Now, we are ready to look at some examples.




Figure 2.1: Number field 6.0.9747.1.

Example 2.11. The field of rational numbers \mathbb{Q} is a number field (of degree 1).

Example 2.12. Let $d \in \mathbb{Z}$ be a nonzero integer that is not a square. The vector space

$$\mathbb{Q}(\sqrt{d}) := \{a + b\sqrt{d} : a, b \in \mathbb{Q}\}$$

is a number field. A \mathbb{Q} -basis is $\{1, \sqrt{d}\}$, so the number field has degree 2. These number fields are called **quadratic fields**. To create them in **Magma**, you can type

```
K<s> := QuadraticField(d);
```

The variable **s** represents \sqrt{d} (or $-\sqrt{d}$, they are indistinguishable). Indeed, you can check

```
assert s^2 eq d;
```

Exercise 2.13. Prove that $\mathbb{Q}(\sqrt{d})$ is the smallest field containing \sqrt{d} .

Example 2.14. Let $n \geq 2$ be an integer and let $\zeta_n := \exp(2\pi i/n)$ be a primitive n -th root of unity. Then

$$\mathbb{Q}(\zeta_n) := \left\{ \sum_{i=0}^{n-1} a_i \zeta_n^i : a_i \in \mathbb{Q} \right\}$$

is a number field and it is called the n -th **cyclotomic field**. The degree $[\mathbb{Q}(\zeta_n) : \mathbb{Q}] = \varphi(n)$, where $\varphi(n)$ is Euler's totient function. You can define this number field in **Magma** as

```
K<z> := CyclotomicField(n);
```

and just as with quadratic fields, **z** represents ζ_n (or any power ζ_n^k with $\gcd(k, n) = 1$).

Example 2.15. For a more general example, consider an irreducible polynomial $f(x) \in \mathbb{Q}[x]$ of degree d . The quotient $\mathbb{Q}[x]/(f)$ is called an **algebraic extension** of \mathbb{Q} and is a number field of degree d .

Let $f(x) \in \mathbb{Q}[x]$ be an irreducible polynomial. Let α be a root of $f(x)$. The field $\mathbb{Q}(\alpha)$ denotes the smallest field that contains both \mathbb{Q} and α . Then, we have $\mathbb{Q}[x]/(f(x)) \simeq \mathbb{Q}(\alpha)$ via $x \mapsto \alpha$. In this sense, note that it does not matter which root of $f(x)$ we pick: they all produce isomorphic fields. When we pick a root α , we pick an explicit embedding

$$\mathbb{Q}[x]/(f(x)) \simeq \mathbb{Q}(\alpha) \subseteq \mathbb{C}. \quad (2.16)$$

Exercise 2.17. With the notation above, convince yourself that $\mathbb{Q}[\alpha] = \mathbb{Q}(\alpha)$. Note that it is enough to write the inverse of α as a polynomial (over \mathbb{Q}) in α . What is the degree of $\mathbb{Q}(\alpha)$ in terms of the degree of $f(x)$?

Lemma 2.18. *If K and L are number fields with $K \subseteq L$, then $[K : \mathbb{Q}]$ divides $[L : \mathbb{Q}]$.*

Proof. This follows from basic facts about vector spaces. By definition, $[K : \mathbb{Q}]$ is the dimension of K as a \mathbb{Q} -vector space, and $[L : \mathbb{Q}]$ is the dimension of L as a \mathbb{Q} -vector space. Also, because $K \subseteq L$, then L is a K -vector space of finite dimension n . Choose a \mathbb{Q} -basis k_1, \dots, k_m for K , where $m = [K : \mathbb{Q}]$, and a K -basis ℓ_1, \dots, ℓ_n for L . Then the set

$$\{\ell_i k_j \mid 1 \leq i \leq n, 1 \leq j \leq m\}$$

forms a \mathbb{Q} -basis for L . Therefore, $[L : \mathbb{Q}] = nm$. □

Examples 2.12 and 2.14 are particular instances of Example 2.15. Indeed,

$$\mathbb{Q}(\sqrt{d}) \simeq \mathbb{Q}[x]/(x^2 - d) \quad \text{and} \quad \mathbb{Q}(\zeta_n) \simeq \mathbb{Q}[x]/(\Phi_n(x)),$$

where $\Phi_n(x)$ is the cyclotomic polynomial of degree $\varphi(n)$. It turns out that we have described all number fields just by looking at Example 2.15!

Theorem 2.19. *If K is a number field, then K is an algebraic extension of \mathbb{Q} .*

Exercise 2.20. Prove Theorem 2.19.

Remark 2.21. Given a number field $K = \mathbb{Q}(\alpha)$ of degree n over \mathbb{Q} , the choices of roots of $m_\alpha(x)$ provide n distinct embeddings of K in \mathbb{C} , as constructed in (2.16). In fact, we will show that every number field can be written as $\mathbb{Q}(\theta)$, so every number field of degree n comes equipped with n distinct embeddings

$$\sigma_i: K \hookrightarrow \mathbb{C}, \quad \sigma_i(\alpha) = \alpha_i.$$

In fact, every element α of a number field K is a root of a polynomial in $\mathbb{Q}[x]$, which implies that $\mathbb{Q}(\alpha)$ is a number field too. This invites the following definition.

Definition 2.22. Let $\alpha \in \mathbb{C}$, we say that α is an **algebraic number** if α is a root of a nonzero polynomial in $\mathbb{Q}[x]$. The **minimal polynomial** of α is a monic polynomial $m_\alpha(x) \in \mathbb{Q}[x]$ of minimal degree. We call all the roots of $m_\alpha(x)$ the **conjugates** of α . We denote the set of algebraic numbers as $\bar{\mathbb{Q}}$.

Remark 2.23. If $\alpha \in \mathbb{C}$ is an algebraic number, then we can copy the definition from Example 2.15 to see that $\mathbb{Q}(\alpha) \simeq \mathbb{Q}[x]/(m_\alpha(x))$, where $m_\alpha(x)$ is the minimal polynomial of α .

Example 2.24. Let $\alpha_1, \dots, \alpha_k$ be algebraic numbers. The number field $\mathbb{Q}(\alpha_1, \dots, \alpha_k)$ is the smallest field that contains \mathbb{Q} and $\alpha_1, \dots, \alpha_k$. Equivalently, it is the *compositum* of the number fields $\mathbb{Q}(\alpha_1), \dots, \mathbb{Q}(\alpha_k)$.

Example 2.25. Working with number fields in **Magma** can be tricky. Let's consider the number fields $K := \mathbb{Q}(\sqrt{6})$ and $L := \mathbb{Q}(\sqrt{2})\mathbb{Q}(\sqrt{3}) = \mathbb{Q}(\sqrt{2}, \sqrt{3})$, the compositum of $\mathbb{Q}(\sqrt{2})$ and $\mathbb{Q}(\sqrt{3})$:

```
K<z> := QuadraticField(6);
L := Compositum(QuadraticField(2), QuadraticField(3));
```

The variable z is representing $\sqrt{6}$. Since $\sqrt{6} = \sqrt{2}\sqrt{3}$, we could argue that z is an element of L . However, asking something like: $z \text{ in } L$; gives an error. The problem? You do not know yet an embedding from K to L , so **Magma** does not know in advance that they are related. You can first ask `IsSubfield(K,L)`; which stores the field embedding $K \subset L$, and allows you to write z as an element of K . Now, $z \text{ in } L$; returns `true`. Moreover, trying `L!z`; writes z in the basis of L .

Lemma 2.26. Let α is an algebraic number with minimal polynomial $m_\alpha(x)$. Assume that $f(x) \in \mathbb{Q}[x]$ satisfies $f(\alpha) = 0$. Then $m_\alpha(x)$ divides $f(x)$ in $\mathbb{Q}[x]$.

Exercise 2.27. Prove Lemma 2.26 by using the Euclidean Algorithm.

The set of algebraic numbers, $\bar{\mathbb{Q}}$, is strictly contained in the complex numbers. The numbers in the difference of these sets are called **transcendental numbers**. You can find an interesting Quanta Magazine article on the history of transcendental numbers [here](#).

Another structural fact about algebraic numbers is that $\bar{\mathbb{Q}}$ forms a field. You might have seen proofs of this, that follow directly from using Theorem 2.19. But what if we care about representing $\alpha + \beta$, $\alpha\beta$, and α/β as algebraic numbers for $\alpha, \beta \in \bar{\mathbb{Q}}$? Let's see what we mean.

2.3 Representing algebraic numbers

From an algorithmic perspective, it is useful to be able to represent algebraic numbers efficiently. We will discuss various explicit representations of algebraic numbers and analyze their computational properties, particularly with respect to arithmetic operations such as addition, multiplication, and inversion.

2.3.1 Using minimal polynomials

Let α be an algebraic number with minimal polynomial $m_\alpha(x)$. Remark 2.23 gives an isomorphism $\mathbb{Q}(\alpha) \simeq \mathbb{Q}[x]/(m_\alpha(x))$, so α can be represented as the class of x in the quotient ring $\mathbb{Q}[x]/(m_\alpha(x))$. We now need to figure out how to add, multiply, and divide algebraic numbers using this representation.

Given two algebraic numbers α and β , the minimal polynomials of $\alpha + \beta$, $\alpha\beta$, or α/β can be computed using *resultants*.

Definition 2.28. Let R be an integral domain with fraction field K , and let \bar{K} be the algebraic closure of K . Let $A(x), B(x) \in R[x]$ be polynomials of degree m and n , respectively. Decompose $A(x) = a \prod_{i=1}^m (x - \alpha_i)$ and $B(x) = b \prod_{i=1}^n (x - \beta_i)$ in \bar{K} , so $\alpha_1, \dots, \alpha_m$ are the

roots of A and β_1, \dots, β_n are the roots of B . The **resultant** of $A(x)$ and $B(x)$, denoted $\text{Res}(A(x), B(x))$, is given by one of the equivalent formulas

$$\begin{aligned} \text{Res}(A(x), B(x)) &= a^n B(\alpha_1) \cdots B(\alpha_m) \\ &= (-1)^{mn} b^m A(\beta_1) \cdots A(\beta_n) \\ &= a^n b^m \prod_{\substack{1 \leq i \leq m \\ 1 \leq j \leq n}} (\alpha_i - \beta_j). \end{aligned}$$

Equivalently, the resultant of two polynomials $A(x) = \sum_{i=0}^m a_i x^i$ and $B(x) = \sum_{i=0}^n b_i x^i$ is the determinant of the *Sylvester matrix* associated to $A(x)$ and $B(x)$:

$$\text{Res}(A(x), B(x)) = \det \begin{pmatrix} a_m & a_{m-1} & \cdots & a_0 & 0 & 0 & 0 \\ 0 & a_m & a_{m-1} & \cdots & a_0 & 0 & 0 \\ \vdots & \vdots & \ddots & \ddots & \ddots & \ddots & \vdots \\ 0 & 0 & 0 & a_m & a_{m-1} & \cdots & a_0 \\ b_n & b_{n-1} & \cdots & b_0 & 0 & 0 & 0 \\ 0 & b_n & b_{n-1} & \cdots & b_0 & 0 & 0 \\ \vdots & \vdots & \ddots & \ddots & \ddots & \ddots & \vdots \\ 0 & 0 & 0 & b_n & b_{n-1} & \cdots & b_0 \end{pmatrix}. \quad (2.29)$$

As we saw in Remark 1.36, computing the determinant of the matrix has cost $O((m+n)^3)$ integer multiplications.

Remark 2.30. If you want to explore a more efficient algorithm for computing resultants, you can check [Coh93, Algorithm 3.3.7].

Now we are ready to compute minimal polynomials! Recall our goal: for two algebraic numbers α and β , we want to find the minimal polynomials of $\alpha + \beta$, $\alpha\beta$, and α/β from the information of the minimal polynomials $m_\alpha(x)$ and $m_\beta(x)$. Let $\alpha_1, \dots, \alpha_m$ and β_1, \dots, β_n be the roots of $m_\alpha(x)$ and $m_\beta(x)$, respectively.

We add an auxiliary variable y , then compute

$$m_\alpha(x - y) = \prod_{i=1}^m (x - y - \alpha_i) = (-1)^m \prod_{i=1}^m (y - (x - \alpha_i)),$$

so $x - \alpha_i$ are the roots of $m_\alpha(x - y)$, seen as a polynomial in y . Consequently, by definition, the resultant of $m_\alpha(x - y)$ and $m_\beta(y)$ seen as polynomials in y is

$$\text{Res}_y(m_\alpha(x - y), m_\beta(y)) = \prod_{\substack{1 \leq i \leq m \\ 1 \leq j \leq n}} ((x - \alpha_i) - \beta_j) = \prod_{\substack{1 \leq i \leq m \\ 1 \leq j \leq n}} (x - (\alpha_i + \beta_j)).$$

By construction, this polynomial in x is monic and has $\alpha + \beta$ as a root. It also has coefficients in \mathbb{Q} by (2.29). If the polynomial is irreducible, then it must be the minimal polynomial of $\alpha + \beta$. If not, then the minimal polynomial $m_{\alpha+\beta}(x)$ must divide it by Lemma 2.26. We can then factor it using [Coh93, Algorithm 3.5.7].

Exercise 2.31. Show that $\text{Res}_y(y^m m_\alpha(x/y), m_\beta(y))$ has $\alpha\beta$ as a root, so factoring this polynomial will result on finding the minimal polynomial of $\alpha\beta$. Similarly, show that you can recover the minimal polynomial of α/β from $\text{Res}_y(m_\alpha(xy), m_\beta(y))$.

2.3.2 Primitive element theorem and the standard representation

A much easier situation occurs when we work inside a number field $\mathbb{Q}(\alpha)$, where we can find representations by using information from $m_\alpha(x)$. This is easier to see for quadratic number fields, as shown by the following exercise.

Exercise 2.32. Consider the quadratic number field $K := \mathbb{Q}(\sqrt{-7})$. Note that $\sqrt{-7} + 1$ is an element of K . Can you find its minimal polynomial? How is it related to the minimal polynomial of $\sqrt{-7}$? Can you now find an algorithm to compute the minimal polynomial of any element $a + b\sqrt{-7} \in K$? Can you generalize this to any quadratic number field?

In general, we have the following theorem.

Theorem 2.33 (Primitive Element Theorem). *Let K be a number field, then there exists an element $\theta \in K$ such that $K = \mathbb{Q}(\theta)$. We say that θ is a **primitive element**.*

Sketch of the proof. Let $K = \mathbb{Q}(\alpha_1, \alpha_2, \dots, \alpha_m)$ be a number field generated over \mathbb{Q} by algebraic numbers α_i . We will show that there is $\theta \in K$ such that $K = \mathbb{Q}(\theta)$. It suffices to show this for $K = \mathbb{Q}(\alpha, \beta)$ (by induction on the number of generators). If α and β are both in \mathbb{Q} , then $K = \mathbb{Q}$. Otherwise, consider the elements $\theta = \alpha + c\beta$ for $c \in \mathbb{Q}$. It turns out that $K = \mathbb{Q}(\theta)$ for all but finitely many c (proving this fact requires using automorphisms of K , so we skip it for brevity). Pick one of the infinitely many c so $K = \mathbb{Q}(\theta)$ ¹. \square

Exercise 2.34. Consider the biquadratic number field $\mathbb{Q}(\sqrt{a}, \sqrt{b})$. Follow the proof of Theorem 2.33 to find a primitive element θ such that $\mathbb{Q}(\sqrt{a}, \sqrt{b}) = \mathbb{Q}(\theta)$. Can you find a way to compute the minimal polynomial of θ ? Can you write \sqrt{a} and \sqrt{b} as polynomials in θ ? If you want, you can pick specific values for a and b .

Remark 2.35. The proof of Theorem 2.33 gives an explicit algorithm for computing a primitive element for any number field K . Given $K = \mathbb{Q}(\alpha_1, \dots, \alpha_m)$, you need to try combinations $\theta_{(c)} := \sum_{i=1}^m c_i \alpha_i$ for $c_i \in \mathbb{Q}$ until you get $K = \mathbb{Q}(\theta_{(c)})$. This process will terminate since there are only finitely many vectors (c) for which $\mathbb{Q}(\theta_{(c)}) \subsetneq K$. However, this algorithm depends on being able to compute if two number fields are equal. We will study this problem in Lecture 4.

Going back to our problem of representing algebraic numbers, Theorem 2.33 gives us a way to perform easier arithmetic when the algebraic numbers belong to the same number field (Given $\alpha, \beta \in \bar{\mathbb{Q}}$, we have that $\mathbb{Q}(\alpha, \beta)$ is indeed a number field containing all the relevant quantities we care about).

If K is a number field of degree n , then we can find a primitive element θ with minimal polynomial $m_\theta(x)$ of degree n . Then, all elements $\alpha \in K = \mathbb{Q}(\theta)$ can be represented uniquely as $\sum_{i=0}^{n-1} b_i \theta^i$ for $b_i \in \mathbb{Q}$. Note that this just means that $\{1, \theta, \dots, \theta^{n-1}\}$ is a \mathbb{Q} -basis of K .

¹In practice, trying $\alpha + \beta$ is always a good choice.

Taking d as the (positive) greatest common divisor of the rational numbers b_i , we arrive to the representation

$$\alpha = \frac{\sum_{i=0}^{n-1} a_i \theta^i}{d}, \quad d > 0, a_i \in \mathbb{Z}, \text{ and } \gcd(a_0, \dots, a_{n-1}, d) = 1,$$

called the standard representation of α with respect to θ .

Remark 2.36. Magma represents elements of number fields using the standard representation with respect to a primitive element which is stored when you create the number field.

Example 2.37. We can consider the number field $K = \mathbb{Q}(\sqrt{2}, \zeta_3)$. we can create this field in Magma by writing it as the compositum of $\mathbb{Q}(\sqrt{2})$ and $\mathbb{Q}(\zeta_3)$. Every time you create a number field in Magma, it comes with a primitive element θ that we can recover

```
N1<s> := QuadraticField(2);
N2<z> := CyclotomicField(3);
K<theta> := Compositum(N1, N2);
```

Then we can check that `z+s eq theta;`, so $s + z$ is chosen as the primitive element of K . We can also check that the elements of K are written in the standard representation with respect to θ . For example, `-(1/50)*s+1+3*z^2+(3/2)*z;` returns

$$1/550*(-148*\theta^3 - 222*\theta^2 - 159*\theta - 730)$$

Let's understand the complexity of working with this standard representation. Adding elements of K reduces to (basically) vector addition in \mathbb{Q}^{n+1} , so it takes $O(n)$ integer operations. To study multiplication, let $m_\theta(x) = \sum_{i=0}^n t_i x^i \in \mathbb{Q}[x]$ be the minimal polynomial of θ , so $t_n = 1$. We note that we can reduce

$$\theta^n = -t_{n-1}\theta^{n-1} - \dots - t_0,$$

and the use recursion to reduce any power $k \geq n$. To make multiplication more efficient, we can precompute and store those reductions. Let $k \geq n$ and write

$$\theta^{n+k} = \sum_{i=0}^{n-1} r_{i,k} \theta^i \tag{2.38}$$

so $r_{i,n} = -t_i$ and

$$r_{k+1,i} = \begin{cases} r_{k,i-1} - t_i r_{k,n-1} & \text{if } i \geq 1, \\ -t_0 r_{k,n-1} & \text{if } i = 0. \end{cases}$$

Exercise 2.39. Show that precomputing the constants $r_{i,k}$ as in (2.38) takes $O(kn)$ field operations.

Once we know the coefficients $r_{i,k}$, we can compute the product of any two elements $\beta, \gamma \in K$ using Algorithm 2.40.

Exercise 2.41. Study the complexity of Algorithm 2.40 in terms of the number of integer operations. You might find the estimates of § 1.3.2 useful.

Algorithm 2.40 (Multiplication in standard representation).

The input is two algebraic numbers in $K = \mathbb{Q}(\theta)$, written in the standard representation $\beta = \frac{1}{d_\beta} \sum_{i=0}^{n-1} b_i \theta^i$ and $\gamma = \frac{1}{d_\gamma} \sum_{j=0}^{n-1} c_j \theta^j$, where $b_i, c_j \in \mathbb{Z}$. We also input the precomputed values $r_{i,k}$ up to $k = (n-1)^2$.

1. Set $d := d_\beta d_\gamma$.
2. Compute the product polynomial: $h(x) := \left(\sum_{i=0}^{n-1} b_i x^i \right) \left(\sum_{j=0}^{n-1} c_j x^j \right) = \sum_{i=0}^{(n-1)^2} h_i x^i$.
3. Set $a_i := h_i$ for $i \in \{1, \dots, n-1\}$.
4. For k in $n, \dots, (n-1)^2$
 - (a) Set $a_i := a_i + r_{i,k}$ for $i = 0, \dots, n-1$.
5. Compute $g := \gcd(a_1, \dots, a_{n-1}, d)$
6. Set $d := d/g$ and the coefficients $a_i := a_i/g$.

Return the standard representation of the product: $\frac{1}{d} \sum_{i=0}^{n-1} a_i \theta^i$.

Remark 2.42. As you can see in Algorithm 2.40, multiplying two algebraic numbers in the standard representation is equivalent to computing a product of two polynomials and then reducing that product modulo $m_\theta(x)$. You can learn more about this in [Coh93, Chapter 3]. However, precomputing the coefficients $r_{i,k}$ makes the algorithm more efficient.

For division, we can use our idea of representing elements of K using polynomials in $\mathbb{Z}[x]$ (plus another integer for the denominator). Then, finding the quotient of β by $\gamma \neq 0$ is equivalent to computing the quotient of the corresponding polynomials modulo $m_\theta(x)$. This can be done as follows. Let $B(x)$ and $C(x)$ be the polynomials associated to β and γ , respectively. The polynomial $C(x)$ is coprime to $m_\theta(x)$ since $\gamma \neq 0$ and the degree of $C(x)$ is at most $n-1$. Then we can explicitly compute $U(x)$ such that

$$U(x)C(x) + V(x)m_\theta(x) = \gcd(C(x), m_\theta(x)) = 1$$

as part of the computation of the gcd (see §1.3.2). Combined with the multiplication algorithm, we can then obtain the standard representation of β/γ .

2.3.3 Other representations

Using minimal polynomials or the standard representation are not the only ways to represent algebraic numbers. We will not discuss more ways because of time. If you are curious, you can look at, for example [Coh93, § 4.2].

Bibliography

- [BZ11] Richard P. Brent and Paul Zimmermann. *Modern computer arithmetic*, volume 18 of *Cambridge Monographs on Applied and Computational Mathematics*. Cambridge University Press, Cambridge, 2011. ↑1.
- [CLO15] David A. Cox, John Little, and Donal O’Shea. *Ideals, varieties, and algorithms*. Undergraduate Texts in Mathematics. Springer, Cham, fourth edition, 2015. An introduction to computational algebraic geometry and commutative algebra. ↑ii.
- [Coh93] Henri Cohen. *A course in computational algebraic number theory*, volume 138 of *Graduate Texts in Mathematics*. Springer-Verlag, Berlin, 1993. ↑ii, 1, 7, 10, 12, 13, 19, 22, 25, 28, 30, 33, 35, 36, 39, 40.
- [Coh00] Henri Cohen. *Advanced topics in computational number theory*, volume 193 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 2000. ↑ii.
- [Har21] David Harvey. Counting points on hyperelliptic curves over finite fields, 2021. *IAS/Park City Mathematics Series*. ↑ii, 1, 2, 6, 9.
- [HvdH21] David Harvey and Joris van der Hoeven. Integer multiplication in time $O(n \log n)$. *Ann. of Math. (2)*, 193(2):563–617, 2021. ↑4, 6.
- [Lan94] Serge Lang. *Algebraic number theory*, volume 110 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 1994. ↑ii, 12.
- [LLL82] A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász. Factoring polynomials with rational coefficients. *Math. Ann.*, 261(4):515–534, 1982. ↑32.
- [Mil] J. S. Milne. Algebraic number theory. <https://www.jmilne.org/math/CourseNotes/ant.html>. ↑ii, 12.
- [Pap94] Christos H. Papadimitriou. *Computational complexity*. Addison-Wesley Publishing Company, Reading, MA, 1994. ↑2.
- [Sta67] H. M. Stark. A complete determination of the complex quadratic fields of class-number one. *Michigan Math. J.*, 14:1–27, 1967. ↑39.
- [Ste] William Stein. Algebraic number theory, a computational approach. <https://wstein.org/books/ant/>. ↑ii, 12, 30.

- [Voi21] John Voight. *Quaternion algebras*, volume 288 of *Graduate Texts in Mathematics*. Springer, Cham, [2021] ©2021. [↑ii](#).
- [vzGG13] Joachim von zur Gathen and Jürgen Gerhard. *Modern computer algebra*. Cambridge University Press, Cambridge, third edition, 2013. [↑ii](#), [1](#), [7](#).