Algebraic Number Theory

Brian Osserman

Contents

CHAPTER 1

Introduction

In these notes, we will cover the basics of what is called algebraic number theory. Just as number theory is often described as the study of the integers, algebraic number theory may be loosely described as the study of certain subrings of fields K with $[K : \mathbb{Q}] < \infty$; these rings, known as "rings of integers", tend to act as natural generalizations of the integers. However, although algebraic number theory has evolved into a subject in its own right, we begin by emphasizing that the subject evolved naturally as a systematic method of treating certain classical questions about the integers themselves.

1.1. An elementary question

The following question neatly illustrates the natural evolution of algebraic number theory from elementary number theory:

Question 1.1.1. For a fixed $n \in \mathbb{N}$, which natural numbers may be written as $x^2 + ny^2$ for $x, y \in \mathbb{Z}$?

This question could be approached, and solved in special cases, in an elementary manner, but the full solution requires the full machinery of modern algebraic number theory. The basic observation is that

$$x^2 + ny^2 = (x + y\sqrt{-n})(x - y\sqrt{-n}), \tag{1.1.1.1}$$

so the question is essentially one of how different integers factor in the ring

$$\mathbb{Z}[\sqrt{-n}] := \{x + y\sqrt{-n} : x, y \in \mathbb{Z}\}. \tag{1.1.1.2}$$

We already see the benefits of taking the algebraic number theory point of view on the problem, as the following becomes a simple observation:

Lemma 1.1.2. *The set of solutions to Question 1.1.1 for a given n is closed under multiplication.*

Of course, this could have been proved in an elementary but more complicated way, but the statement itself would have been substantially more difficult to discover.

Because of this multiplicativity, it is a natural starting point to restrict our question to the case of which primes are of the form $x^2 + ny^2$. To start with, we will focus on the case $n = 1$. As you may well already know, there is the following theorem:

Theorem 1.1.3. *A prime number p may be written in the form $x^2 + y^2$ for $x, y \in \mathbb{Z}$ if and only if $p = 2$ or $p \equiv 1 \pmod 4$.*

One direction of this is completely elementary: since $0, 1$ are the only squares modulo 4, any odd sum of squares must be 1 modulo 4. However, the converse is far deeper. We observe that if $p = x^2 + y^2$, we have $(\frac{x}{y})^2 \equiv -1 \pmod{p}$, so -1 has a square root modulo p. We will use the following easy result from classical number theory:

Proposition 1.1.4. *-1 has a square root modulo p if and only if $p = 2$ or $p \equiv 1 \pmod{4}$.*

However, it is not immediately obvious that the existence of a square root of -1 mod p should yield a solution to $p = x^2 + y^2$. This is where algebraic number theory comes in. The key point is that if $z^2 \equiv -1 \pmod{p}$, we have that $p|(z^2 + 1)$, and thus $p|(z + i)(z - i)$ in the ring $\mathbb{Z}[i]$. It is clear that p does not divide $z + i$ or $z - i$, so what we would like to know is the analogue for $\mathbb{Z}[i]$ of the following statement over the integers:

Theorem 1.1.5. *Suppose p is a prime number, and $p|ab$. Then either $p|a$ or $p|b$.*

Indeed, this statement does generalize to $\mathbb{Z}[i]$ by Exercise 1.3, so it follows that p cannot be prime in $\mathbb{Z}[i]$, and then by Exercise 1.2 we have $p = a^2 + b^2$, as desired.

It is possible to give a completely elementary proof of which primes are sums of squares, as in fact Euler did (see [**3**, pp. 10-11]). The proof is not even excessively long, although it is rather intricate. However, the proof lacks the clear conceptual structure of the algebraic number theory proof we have discussed, and therefore, although generalizable to some degree on a case-by-case basis, could not hope to answer the question for general n, as we will ultimately be able to.

1.2. A new vocabulary

The problem we have just discussed makes it clear that in order to approach problems in elementary number theory, it will often be helpful to know which rings satisfy Theorem 1.1.5. In order to phrase the question properly, we should first state carefully what we should mean by "prime" number:

Definition 1.2.1. Let R be a ring (always commutative, with identity). Then $x \in R$ is a **unit** if there exists $y \in R$ with $xy = 1$, and x is **irreducible** if x is not a unit, and for any $y, z \in R$ with $x = yz$, either y or z is a unit. Finally, x is **prime** if it satisfies the condition on p of Theorem 1.1.5.

We are thus interested in which rings have the property that every irreducible element is prime. We immediately see that the distinction in terminology is justified:

Example 1.2.2. In the ring $\mathbb{Z}[\sqrt{-5}]$, 2 is irreducible, but not prime, since $2|(1 + \sqrt{-5})(1 - \sqrt{-5})$, but $2 \not| (1 \pm \sqrt{-5})$.

If we attempt to generalize the proof of the $n = 1$ case of Question 1.1.1, we can use quadratic reciprocity to completely describe primes p such that $-n$ is a perfect square mod p, and we already saw that any p which is not irreducible in $\mathbb{Z}[\sqrt{-n}]$ may be written as $x^2 + ny^2$. Our previous argument gives us the following:

Proposition 1.2.3. *Fix $n \in \mathbb{N}$. Then if a prime p may be written in the form $x^2 + ny^2$, it follows that $-n$ is a square mod p, and conversely, if $-n$ is a square mod p, it follows that p is not prime in $\mathbb{Z}[\sqrt{-n}]$.*

Thus, the only obstruction to obtaining an if and only if statement is that in general in $\mathbb{Z}[\sqrt{-n}]$, if p is not prime, it may still be irreducible. However, the machinery of class field theory will allow us to prove the following theorem:

THEOREM 1.2.4. *Fix $n \in \mathbb{N}$. Then there exists a monic irreducible polynomial $f_n(x) \in \mathbb{Z}[x]$ such that for primes p not dividing $2n \operatorname{disc} f_n(x)$, we have that $p = x^2 + ny^2$ for some $x, y \in \mathbb{Z}$ if and only if $-n$ is a perfect square mod p, and $f_n(x)$ has a root mod p.*

Here $\operatorname{disc} f_n(x)$ denotes the discriminant of $f_n(x)$. In the case that $\mathbb{Z}[\sqrt{-n}]$ has every irreducible element prime, we have that the $f_n(x)$ of the theorem could simply be x. Unfortunately, there are only finitely many such n, and indeed the degree of $f_n(x)$ goes to infinity as n does. Finally, we remark that there are explicit methods of constructing the polynomial $f_n(x)$, but these rely on the theory of elliptic curves with complex multiplication, and are thus beyond the scope of our treatment.

1.3. Other motivating problems

We go on to mention a couple other classical problems which algebraic number theory addresses in one way or another.

First, there is Pell's equation (note that apparently, Pell had nothing to do with this problem; see [**4**, p. 33]):

QUESTION 1.3.1. For a fixed $n \in \mathbb{N}$, which is not a perfect square, what are the integer solutions to the equation $x^2 - ny^2 = 1$?

Solutions to this problem are constructive, classical and well-known, but algebraic number theory places the problem in a more conceptual and general framework, with the following observation:

PROPOSITION 1.3.2. *For $x, y \in \mathbb{Z}$, we have $x^2 - ny^2 = \pm 1$ if and only if $x + y\sqrt{n}$ is a unit in $\mathbb{Z}[\sqrt{n}]$.*

In particular, because the units in any ring necessarily form a group under multiplication, we find immediately that there is a natural group structure on the solutions to Pell's equation, and the following theorem, at least in the case of square-free n, will be a special case of a far more general result.

THEOREM 1.3.3. *For any $n \in \mathbb{N}$ not a perfect square, the group of solutions to Pell's equation is isomorphic to the additive group $\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$.*

Finally, we mention the following rather well-known problem:

QUESTION 1.3.4. For a fixed $n > 2$, are there non-zero integers solutions to the equation $x^n + y^n = z^n$?

Fermat famously asserted that there are no such solutions for any n, and because it was the last of his assertions to remain unproved, it came to be known as "Fermat's last theorem". Although this has now been answered for all n through extremely deep work of Frey, Serre, Mazur, Ribet, Wiles, and Taylor, prior progress had been made by Kummer and others via more accessible techniques of algebraic number theory.

In 1847, Lamé announced a proof of Fermat's last theorem, but it was flawed in that it incorrectly assumed that the ring $\mathbb{Z}[\zeta_p]$, where $\zeta_p = e^{2\pi i/p}$ is a pth root of unity, has the property that every irreducible element is prime. Although this was false, Kummer was able to refine the ideas in it to give an elegant partial solution.

To properly state Kummer's theorem, we need the following definition from classical number theory:

DEFINITION 1.3.5. We define the **Bernoulli numbers**, $B_m \in \mathbb{Q}$, inductively as follows: $B_0 = 1$, and for all m, $B_m = \frac{-1}{m+1}\sum_{k=0}^{m-1}\binom{m+1}{k}B_k$.

Kummer developed the theory of ideals of rings in part to prove the following remarkable theorem:

THEOREM 1.3.6. *Suppose that p is an odd prime number such that p does not divide any of the numerators of $B_2, B_4, \dots, B_{p-3}$ (in particular, $p = 3$ is acceptable). Then the equation $x^p + y^p = z^p$ has no solutions in non-zero integers.*

Of course, not all primes satisfy the hypotheses of the theorem, but the statement is sufficiently constructive that it is trivial to check whether or not a given prime does. Unfortunately, although numerically it seems that roughly 60% of primes satisfy the hypotheses of the theorem, and it has been shown that infinitely many do not satisfy them, it is still not known whether infinitely many primes satisfy the hypotheses. However, later refinements of Kummer's techniques resulted in algorithms to verify that $x^p + y^p = z^p$ has no non-zero solutions for any given prime p, and these were used to check Fermat's Last Theorem for primes up to 4,000,000 before Wiles announced his proof of the general statement.

Kummer's theorem is highly non-trivial, but we will be able to give a proof using the theory of ideal class groups and the analytic class number formula.

These examples together present a strong case that even if one only wishes to study problems in elementary number theory, it is often natural and important to consider more general number systems than the integers or $\mathbb{Z}/n\mathbb{Z}$. By considering number systems that include roots of carefully-chosen equations, but still have many properties in common with the integers, we will develop powerful tools that apply even to very elementary questions.

1.4. Rings of integers

Since we wish to consider what happens when we allow ourselves to work with roots of particular polynomials, the first definition of algebraic number theory is the following:

DEFINITION 1.4.1. A **number field** K is an extension of $\mathbb{Q}$ of finite degree.

By the primitive element theorem, this is the same as a field obtained by adjoining a root of a single polynomial to $\mathbb{Q}$.

What is the best way to generalize rings such as $\mathbb{Z}[i]$? Although the definition may appear unmotivated at first, we will mainly consider the following rings:

DEFINITION 1.4.2. Given a number field K, the **ring of integers** $\mathscr{O}_K$ of K is defined to be the set of elements $\alpha \in K$ such that α satisfies a monic polynomial equation with coefficients in $\mathbb{Z}$; that is, there exists $f(x) \in \mathbb{Z}[x]$ monic and such that $f(\alpha) = 0$.

We will see that rings of integers have many important properties in common with $\mathbb{Z}$. The ring of integers of a number field is intended to generalize the ring $\mathbb{Z}$

inside of $\mathbb{Q}$. Thus, we ought to check that $\mathbb{Z}$ is in fact the ring of integers of $\mathbb{Q}$. We recall some basic terminology which will be directly relevant.

DEFINITION 1.4.3. We say that an integral domain R is a **unique factorization domain**, or UFD, if every non-zero element may be factored uniquely as a product of irreducible elements, up to multiplication by units. In particular, every irreducible element is prime.

DEFINITION 1.4.4. Let R be an integral domain with field of fractions K, and L a field containing K. We say that $\alpha \in L$ is **integral** over R if it is the root of a monic polynomial with coefficients in R. The **integral closure** of R inside L is the set of all elements of L which are integral over R. Finally, we say that R is **integrally closed** (in its field of fractions) if it is its own integral closure inside K.

Exercise 1.4 shows that in the definition of integral we can require that the monic minimal polynomial of α has coefficients in R.

The following lemma will allow us to check that $\mathbb{Z}$ is the ring of integers of $\mathbb{Q}$.

LEMMA 1.4.5. *Let R be a unique factorization domain. Then R is integrally closed in its field of fractions.*

PROOF. Given $\alpha = r/s$ in the field of fractions of R, we may assume that r, s have no non-unit common factors. Suppose that α satisfies a monic polynomial $x^m + a_{m-1}x^{m-1} + \cdots + a_1 x + a_0$, with $a_i \in R$ for all i. Because R is a UFD, if $r/s \notin R$, there is some prime p such that $p|s$, whence $p \not| r$. But substituting r/s for x in the polynomial, and multiplying through by s^m, we find $r^m + a_{m-1}r^{m-1}s + \cdots + a_1 r s^{m-1} + s^m = 0$, so $p|r^m$, and hence $p|r$, which is a contradiction. □

It then follows immediately from the definitions and the fact that $\mathbb{Z}$ is a UFD that:

COROLLARY 1.4.6. *$\mathbb{Z}$ is integrally closed, hence is the ring of integers of $\mathbb{Q}$.*

Thus, rings of integers indeed generalize the classical integers. Moreover, by Exercise 1.5 they generalize the ring $\mathbb{Z}[i]$ as well.

The next statement we ought to check is that our terminology is justified: a ring of integers is actually a ring.

PROPOSITION 1.4.7. *A ring of integers $\mathscr{O}_K$ is an integral domain (in particular, a ring), and integrally closed in K.*

More generally, if R is an integral domain, and L a field containing R, then the set S of elements of L which are integral over R is an integral domain, and integrally closed in L.

The proof of the first part of the proposition is very similar to the proof that if a field is extended by a finite number of algebraic elements, the resulting field extension is algebraic. In fact, the statement for fields follows as a special case. We need the following lemma.

LEMMA 1.4.8. *Let R be an integral domain, L any field containing R, and $W \subseteq L$ a non-zero finitely-generated R-module. Given any $\alpha \in L$, suppose that $\alpha W \subseteq W$. Then α is integral over R; that is, it satisfies a monic polynomial with coefficients in R. Conversely, if α is integral over R, then $R[\alpha]$ is a W as above.*

In particular, $R[\alpha]$ is a finitely-generated R-module if and only if α is integral over R.

Proof. Let $m_\alpha : W \to W$ denote the map given by multiplication by α, which we observe is R-linear. Given a set of generators $(w_1, \dots, w_n)$ for W over R, we find that m_α is represented by an $n \times n$ matrix $(m_{i,j})_{i,j}$ with coefficients in R. Now, inside of L, we have $\alpha w_1 = \sum_{i=1}^n m_{i1} w_i$, by definition. But it follows that the map $\alpha I_n - (m_{i,j})_{i,j} : L^n \to L^n$ has w_1 in its kernel, hence the determinant is 0. Expanding the determinant, we obtain a monic polynomial for α with coefficients in R.

Conversely, if $f(\alpha) = 0$ for $f(x) \in R[x]$ monic of degree d, it is clear that $1, \alpha, \dots, \alpha^{d-1}$ generates $R[\alpha]$ as a R-module. □

Proof of Proposition. Since S is contained in the field L, it is clear that if it is a ring, it is an integral domain. It also is clear that if α is integral, then so is $-\alpha$, since we can just alternately change the signs in any monic integral polynomial satisfied by α. It therefore suffices to check that if α_1, α_2 are integral over R, then so are $\alpha_1 + \alpha_2$ and $\alpha_1\alpha_2$. Now, if α_1, α_2 satisfy monic polynomials in $R[x]$ of degrees d_1, d_2 respectively, consider the R-module $W = \left\langle (\alpha_1^i \alpha_2^j)_{0 \leqslant i < d_1, 0 \leqslant j < d_2} \right\rangle$. One checks that because of the monic polynomials satisfied by the α_i, we have $\alpha_i W \subseteq W$ for $i = 1, 2$, and hence $\alpha_1\alpha_2 W \subseteq W$ and $(\alpha_1 + \alpha_2)W \subseteq W$, so the fact that S is a ring follows from the lemma.

Next, suppose $\alpha \in L$ satisfies a polynomial $x^m + a_{m-1}x^{m-1} + \cdots + a_1 x + a_0$, with $a_i \in S$ for all i. Then consider the ring $W := R[\{a_i\}_i, \alpha] \subseteq L$ as an R-module. By Lemma 1.4.8 we have that $R[\{a_i\}_i]$ is a finitely-generated R-module because the a_i are integral over R, and W is a finitely-generated $R[\{a_i\}_i]$-module for the same reason. It easily follows that W is a finitely-generated R-module by taking products of generators. Since $\alpha W \subseteq W$, another application of Lemma 1.4.8 give us that α is integral over R, so $\alpha \in S$ and S is integrally closed in L. □

Thus, we have seen that our definition gives a family of rings which generalizes $\mathbb{Z}$ and $\mathbb{Z}[i]$. The rest of this chapter will be devoted to studying other ways in which rings of integers generalize properties of $\mathbb{Z}$.

1.5. Norm, trace, and discriminant

The norm, trace and discriminant are fundamental concepts for dealing with extensions of rings, and are important tools in understanding rings of integers. We treat these concepts in full generality, as the more general forms will be useful the chapters which follow. We suppose that R and S are rings (not necessarily integral domains), with S containing R, and such that S is a free R-module of rank n.

Notation 1.5.1. Given any $\alpha \in S$, we denote by $m_\alpha : S \to S$ the map given by multiplication by α.

Observe that m_α is an R-linear endomorphism of the R-module S.

Definition 1.5.2. Given S over R, we define the **norm** $N_{S/R}(\alpha)$ to be $\det m_\alpha$, the **trace** $\mathrm{Tr}_{S/R}(\alpha)$ to be $\mathrm{Tr}\, m_\alpha$, and the **discriminant** $D_{S/R}((\alpha_i)_i)$ to be $\det((\mathrm{Tr}_{S/R}(\alpha_i\alpha_j))_{i,j})$, given any $(\alpha_1, \dots, \alpha_n)$ in S.

Applying basic linear algebra, we note the following.

Observation 1.5.3. *The trace is additive and R-linear, and the norm is multiplicative.*

We have the following change-of-basis formula for the discriminant.

PROPOSITION 1.5.4. *Let* $(x_1, \ldots, x_n)$ *be in* S, *and* $M : S \to S$ *an* R*-linear map. Write* $y_i := M(x_i)$. *Then*

$$D_{S/R}(y_1, \ldots, y_n) = (\det M)^2 D_{S/R}(x_1, \ldots, x_n).$$

PROOF. Recalling that the trace is additive, we have

$$\mathrm{Tr}_{S/R}(y_p y_q) = \mathrm{Tr}_{S/R}(\sum_{i,j} m_{p,i} m_{q,j} x_i x_j) = \sum_{i,j} m_{p,i} m_{q,j} \mathrm{Tr}_{S/R}(x_i x_j).$$

This gives the matrix identity

$$(\mathrm{Tr}_{S/R}(y_p y_q))_{p,q} = (m_{pi})_{p,i} (\mathrm{Tr}_{S/R}(x_i x_j))(m_{q,j})_{j,q},$$

and taking determinants and using that determinants are invariant under transpose, we get the desired identity. □

Suppose L/K is separable, as will always be the case in characteristic 0, so in particular for number fields. Fix an algebraic closure $\bar{K}$, and let $\sigma_1, \ldots, \sigma_n$ be the distinct imbeddings of L into $\bar{K}$ fixing K (if L/K is Galois, these are all related by automorphism of L). We have the following:

PROPOSITION 1.5.5. *With* L/K *separable, we have*

$$\mathrm{Tr}_{L/K}(\alpha) = \sum_{i=1}^{n} \sigma_i(\alpha),$$

$$N_{L/K}(\alpha) = \prod_{i=1}^{n} \sigma_i(\alpha),$$

and

$$D_{L/K}((\alpha_i)_i) = (\det(\sigma_i(\alpha_j))_{i,j})^2.$$

If L/K *inseparable, then* $\mathrm{Tr}_{L/K}(\alpha) = 0$ *for any* α, *and hence* $D_{L/K}((\alpha_i)_i) = 0$ *for any* $(\alpha_i)_i$.

PROOF. All but the discriminant assertion for separable extensions is proved in Exercise 1.6 below. Applying the trace formula, we have

$$\mathrm{Tr}_{L/K}(\alpha_i \alpha_j) = \sum_{\ell=1}^{n} \sigma_\ell(\alpha_i) \sigma_\ell(\alpha_j).$$

It follows that we obtain the matrix identity

$$(\mathrm{Tr}_{L/K}(\alpha_i \alpha_j))_{i,j} = (\sigma_\ell(\alpha_i))_{i,\ell} (\sigma_\ell(\alpha_j))_{\ell,j},$$

and taking the determinant of both sides, and using that the determinant is invariant under matrix transposition, we get the desired identity. □

A useful corollary of the proposition is the integrality of traces, norms, and discriminants:

COROLLARY 1.5.6. *In the situation of the proposition, suppose that* α *and the* α_i *are all integral over some ring* $R \subseteq K$. *Then* $\mathrm{Tr}_{L/K}(\alpha), N_{L/K}(\alpha)$, *and* $D_{L/K}((\alpha_i)_i)$ *are integral over* R.

PROOF. First note that if $\omega \in L$ is integral over R, then $\sigma_i(\omega)$ is still integral over R for any i, and in fact satisfies the same monic integral polynomial, since by hypothesis applying σ fixes the coefficients, as they are in $R \subseteq K$. Since the trace, norm and discriminant are obtained as sums and products of such elements, by Proposition 1.4.7 we obtain integrality over R. □

The terminology of discriminants is not a coincidence. Indeed, it agrees with discriminants of polynomials in a natural way:

LEMMA 1.5.7. *If L/K is a field extension of degree n, with $L = K(\alpha)$ for some α, and $f(x)$ the monic minimal polynomial for α over K, then $D_{L/K}(1, \alpha, \ldots, \alpha^{n-1}) = \operatorname{disc} f(x)$. Recall that we define*

$$\operatorname{disc} f(x) := \prod_{i<j} (\alpha_i - \alpha_j)^2,$$

where α_i are the roots of $f(x)$ in a splitting field, counted with multiplicity.

PROOF. If $f(x)$ is inseparable, then $\operatorname{disc} f(x) = 0$ by definition, and we saw in Proposition 1.5.5 that $D_{L/K}((\beta)_i = 0$ for any β_i. On the other hand, if L/K is separable, Proposition 1.5.5 expresses $D_{L/K}(1, \alpha, \ldots, \alpha^{n-1})$ as the square of the Vandermonde determinant for $\sigma_i(\alpha)$, which is $\prod_{i>j}(\sigma_i(\alpha) - \sigma_j(\alpha))$. Therefore, the $D_{L/K}(1, \alpha, \ldots, \alpha^{n-1}) = \operatorname{disc} f(x)$. □

Recall that there are general formulas for the discriminant of a polynomial in terms of its coefficients: the discriminant formula is clearly a symmetric polynomial in the roots of $f(x)$, and the coefficients of $f(x)$ are the elementary symmetric polynomials in the roots, so the discriminant for any given degree may be written as a polynomial in the coefficients.

EXAMPLE 1.5.8. In the case that $f(x) = x^2 + a_1 x + a_2$ with roots α_1, α_2, then $a_1 = -\alpha_1 - \alpha_2$, and $a_2 = \alpha_1 \alpha_2$. We have the discriminant is $(\alpha_2 - \alpha_1)^2 = \alpha_2^2 - 2\alpha_2\alpha_1 + \alpha_1^2 = a_1^2 - 4a_2$.

Finally, we also see that the discriminant tests for bases of separable extensions L over K.

LEMMA 1.5.9. *If L/K is a field extension of degree n, and we are given $\alpha_1, \ldots, \alpha_n \in L$, we have $D_{L/K}((\alpha_i)_i) \neq 0$ if and only if $(\alpha_i)_i$ forms a basis for L over K, and L is separable over K.*

PROOF. If L is separable over K, by the primitive element theorem we can set $L = K(\alpha)$ for some $\alpha \in L$ with separable minimal polynomial. Lemma 1.5.7 shows that $D_{L/K}(1, \alpha, \ldots, \alpha^{n-1}) \neq 0$. The fact that the discriminant is non-zero for any basis of L then follows from the change-of-basis formula of Proposition 1.5.4.

On the other hand, any K-linear dependence of the α_i gives the same K-linear dependence of the $\sigma_j(\alpha_i)$ for any fixed j, so we obtain a linear dependence on the columns of the matrix $(\sigma_j(\alpha_i))_{j,i}$, which then has determinant 0, implying by Proposition 1.5.5 that the discriminant is 0.

Finally, by the last part of Proposition 1.5.5, if L/K is inseparable, we always have $D_{L/K}((\alpha_i)_i) = 0$. □

1.6. First properties of rings of integers

We now begin to examine rings of integers in earnest, showing that $\mathscr{O}_K$ has field of fractions K, and that for any non-zero ideal $\mathscr{O}_K/I$ is finite. We also introduce the discriminant of a number field, using its ring of integers.

The discriminant gives us a somewhat constructive proof of the following fact:

Proposition 1.6.1. *Let I be a non-zero ideal in a ring of integers $\mathscr{O}_K$, where $[K : \mathbb{Q}] = n$. Then I is a free $\mathbb{Z}$-module of rank n, and generated by a $\mathbb{Q}$-basis for K.*

More precisely, I contains a $\mathbb{Q}$-basis for K, and among such bases, any basis whose discriminant has minimal absolute value freely generates I as a $\mathbb{Z}$-module.

Proof. Let $(\alpha_1, \dots, \alpha_n)$ be any $\mathbb{Q}$-basis for K; we first claim there for each i, there exists a non-zero $d_i \in \mathbb{Z}$ such that $d_i\alpha_i \in \mathscr{O}_K$. Indeed, it is easy to check that it is enough to let d_i be the leading term of any integer polynomial satisfied by α_i. Thus, for any non-zero $\beta \in I$, we find that $(\beta d_1\alpha_1, \dots, \beta d_n\alpha_n)$ is a $\mathbb{Q}$-basis for K contained in I.

It remains to show that such a basis with minimal discriminant over $\mathbb{Q}$ is in fact a $\mathbb{Z}$-basis for I. Note that because the discriminant of elements in $\mathscr{O}_K$ is an integer, there is a minimal one. Take any $\alpha \in I$, and write $\alpha = \sum_{i=1}^n a_i\alpha_i$, for $a_i \in \mathbb{Q}$. We want to show that $a_i \in \mathbb{Z}$ for all i. Suppose not; without loss of generality, we may assume $a_1 \notin \mathbb{Z}$. Write $a_1 = m+\epsilon$, with $0 < \epsilon < 1$ and $m = \lfloor a_1 \rfloor$, and replace α_1 by $\alpha_1' = \alpha - m\alpha_1$, which is still in I. We then obtain a new basis inside I, and the determinant of the transition matrix is ϵ, so by the change-of-basis formula of the previous section, this new basis has smaller discriminant, contradicting minimality. □

We can immediately conclude:

Corollary 1.6.2. *We can define the* **discriminant** *D_K of a number field K in terms of any generators of $\mathscr{O}_K$ over $\mathbb{Z}$ as in the proposition. It is a well-defined, non-zero integer.*

Proof. The existence of such a basis follows from the proposition, and the fact that it is a well-defined, non-zero integer follows from the last three results of the previous section, noting that the transition matrix between any two $\mathbb{Z}$-bases must have determinant ±1. □

The discriminant will be a basic and important tool for studying number fields. We also find:

Corollary 1.6.3. *The field of fractions of $\mathscr{O}_K$ is K.*

The proposition also allows us to conclude the following important fact:

Corollary 1.6.4. *Let I be a non-zero ideal in a ring of integers $\mathscr{O}_K$. Then $\mathscr{O}_K/I$ is finite.*

Proof. Suppose we know that I contains some non-zero integer $m \in \mathbb{Z}$. Then $\mathscr{O}_K/I$ is a quotient of $\mathscr{O}_K/(m)$, which from the proposition is isomorphic as a $\mathbb{Z}$-module to $(\mathbb{Z}/m\mathbb{Z})^n$, hence both are finite. The corollary then follows from the easy lemma below. □

Lemma 1.6.5. *Let R be an integral domain, and S an integral domain containing R. Let I be an ideal of S, and suppose I contains a non-zero element α satisfying a non-zero polynomial $f(x) \in R[x]$. Then $I \cap R \neq 0$.*

Proof. Because S is an integral domain, we may factor out any powers of x dividing $f(x)$, and can therefore assume that $f(0) \neq 0$. But $\alpha | (f(\alpha) - f(0))$, so $-f(0) = f(\alpha) - f(0) \in I \cap R$, as desired. □

This finiteness result will play a key role in what follows, and indeed in all of algebraic number theory. It allows us to make the following definition:

Definition 1.6.6. If I is an ideal of a ring of integers $\mathscr{O}_K$, we define the **norm** of I to be $N(I) := \#\{\mathscr{O}_K/I\}$.

The norm can be useful in recognizing prime ideals:

Lemma 1.6.7. *Suppose $\mathfrak{p}$ is an ideal of a ring of integers $\mathscr{O}_K$, and $N(\mathfrak{p}) = p$ for some integer prime $p \in \mathbb{Z}$. Then $\mathfrak{p}$ is prime in $\mathscr{O}_K$.*

Proof. Indeed, if $\mathscr{O}_K/\mathfrak{p}$ has p elements, since the natural map $\mathbb{Z} \to \mathscr{O}_K/\mathfrak{p}$ cannot be uniformly zero, it must have kernel (p), and we find $\mathbb{Z}/p \overset{\sim}{\to} \mathscr{O}_K/\mathfrak{p}$, so in particular the latter is an integral domain and $\mathfrak{p}$ is prime. □

Of course, many prime ideals do not have prime norms. The following lemma generalizes the statement that $|\det M| = 1$ if and only if M is invertible over $\mathbb{Z}$, and we can use it to relate norms of elements to norms of principal ideals.

Lemma 1.6.8. *Let W, V be free $\mathbb{Z}$-modules of rank n, and $M : W \to V$ an injective $\mathbb{Z}$-linear map. Then $|\det M| = \#V/M(W)$.*

Proof. Identify V with $\mathbb{Z}^n \subseteq \mathbb{R}^n$, so that $M(W)$ forms a sublattice of V. Taking any fundamental domain S for V, which necessarily has volume 1, it is clear that we can construct a fundamental domain for $M(W)$ as unions of translates of S, having volume $\#V/M(W)$. But the volume of a fundamental parallelepiped for $M(W)$ is $|\det M|$, and since any two fundamental domains have the same volume, we obtain the identity. □

An immediate corollary is the following:

Corollary 1.6.9. *If $I = (x)$ is a principal ideal, then $|N_{K/\mathbb{Q}}(x)| = N((x))$.*

Proof. We simply apply the lemma to the case that $W = \mathscr{O}_K$, and M is multiplication by x. □

1.7. The discriminant and rings of integers

Having defined and studied some of the most basic properties of rings of integers, we start to develop techniques for computing what the ring of integers of a given field is. The following is helpful when D_K happens to be square-free.

Lemma 1.7.1. *Let $\mathscr{O}_K$ be a ring of integers, and $(x_1, \dots, x_n) \in \mathscr{O}_K$. Suppose that $D_{K/\mathbb{Q}}((x_i)_i)$ is square-free. Then as a $\mathbb{Z}$-module, $\mathscr{O}_K = \langle x_1, \dots, x_n \rangle$.*

Proof. By hypothesis, $\langle x_1, \dots, x_n \rangle$ is contained as a $\mathbb{Z}$-sub-module of $\mathscr{O}_K$. If $(y_1, \dots, y_n)$ are a basis of $\mathscr{O}_K$, we have an integer matrix M expressing the x_i in terms of the y_i, and by Proposition 1.5.4, we have that $D_{K/\mathbb{Q}}((x_i)_i) = (\det M)^2 D_{K/\mathbb{Q}}((y_i)_i)$. Since $D_{K/\mathbb{Q}}((x_i)_i)$ is square free, we have $|\det M| = 1$, so M is invertible over $\mathbb{Z}$, and we have the desired statement. □

We can conclude:

Corollary 1.7.2. *Let $f(x) \in \mathbb{Z}[x]$ be an irreducible monic polynomial, and α a root of $f(x)$. If* $\operatorname{disc} f(x)$ *is square-free, then $\mathscr{O}_{\mathbb{Q}(\alpha)} = \mathbb{Z}[\alpha]$.*

Proof. By Lemma 1.5.7, we have $\operatorname{disc} f(x) = D_{\mathbb{Q}(\alpha)/\mathbb{Q}}(1, \alpha, \dots, \alpha^{n-1})$, where $n = \deg f(x)$. But by the lemma, this means that if $\operatorname{disc} f(x)$ is square-free, $\mathbb{Z}[\alpha] = \left\langle 1, \alpha, \dots, \alpha^{n-1} \right\rangle$ is equal to $\mathscr{O}_{\mathbb{Q}(\alpha)}$. □

We conclude with an example of using the discriminant to compute a ring of integers of a quintic field extension.

Example 1.7.3. We consider the number field $\mathbb{Q}(\alpha)$, where α is a root of the polynomial $f(x) = x^5 - x + 1$. One can check by direct computation that this polynomial is irreducible mod 5, so it is irreducible over $\mathbb{Q}$, and $K = \mathbb{Q}(\alpha)$ has degree 5 over $\mathbb{Q}$, and up to isomorphism is independent of the choice of α.

The formula for the discriminant of a polynomial of the form $x^5 + ax + b$ is $5^5b^4 + 2^8a^5$. Plugging in $a = -1, b = 1$, we find that our $f(x)$ has discriminant 2869. This factors as $19 \cdot 151$, so it is square-free, and by Corollary 1.7.2, we find that $\mathscr{O}_K = \mathbb{Z}[\alpha]$.

1.8. Exercises

Exercise 1.1. Show that an element α of a ring of integers $\mathscr{O}_K$ is a unit if and only if its norm over $\mathbb{Q}$ is ± 1. Show that in the case that K is Galois over $\mathbb{Q}$, this still holds for any ring $R \subseteq \mathscr{O}_K$ which is Galois invariant.

Conclude that in particular, $x + y\sqrt{n}$ is a unit in $\mathbb{Z}[\sqrt{n}]$ if and only if x, y is a solution to either the Pell equation $x^2 - ny^2 = 1$ or the equation $x^2 - ny^2 = -1$ (observe, however, that if $x+\sqrt{-n}y$ is a solution to $x^2-ny^2 = -1$, then $(x+\sqrt{-n}y)^2$ is a solution to the Pell equation).

Exercise 1.2. Let $p \in \mathbb{Z}$ be a prime number, and suppose that $p = \alpha\beta$ for some non-units $\alpha, \beta \in \mathbb{Z}[\sqrt{-n}]$. Using the first part of Exercise 1.1, show that $p = x^2 + ny^2$ for some $x, y \in \mathbb{Z}$.

Recall that an integral domain R is said to be a **Euclidean domain** if there exists a norm map $N : R \to \mathbb{N} \cup \{0\}$ satisfying:

(i) $N(r) = 0$ if and only if $r = 0$;
(ii) For all $a, b \in R$, with b non-zero, there exist $q, r \in R$ such that $a = bq + r$, and $N(r) < N(b)$.

Recall that a Euclidean domain is a principal ideal domain, hence a unique factorization domain.

Exercise 1.3. Show that $\mathbb{Z}[\sqrt{n}]$ is a Euclidean domain, and hence satisfies the statement of Theorem 1.1.5, if $n = -2, -1, 2$ or 3.

Exercise 1.4. Let R be integrally closed with fraction field K, and S an integral domain containing R, and suppose that $x \in S$ is integral over R. Then the monic minimal polynomial of x in $K[t]$ in fact lies in $R[t]$.

Exercise 1.5. Using integrality of the norm and trace, check that for $n \in N$ square-free, the ring of integers of $\mathbb{Q}(\sqrt{n})$ is given as $\{a + b\omega : a, b \in \mathbb{Z}\}$, and where ω is given as follows:

(I) if $n \not\equiv 1 \pmod 4$, then $\omega = \sqrt{n}$;
(II) if $n \equiv 1 \pmod 4$, then $\omega = \frac{1+\sqrt{n}}{2}$.

Exercise 1.6. Prove the first two parts of Proposition 1.5.5 in the following steps:

(1) Show that if $\alpha \in L$, and $f(x) \in K(x)$ is the monic minimal polynomial for α, with degree d, then $\det(xI - m_\alpha) = f(x)^{n/d}$.
(2) Using the fact that if F/E is a separable field extension of finite degree, then any imbedding $E \to \bar{F}$ has exactly $[F : E]$ extensions to imbeddings $F \to \bar{F}$, show that in the above notation, $f(x)^{n/d} = \prod_{i=1}^{n}(x - \sigma_i(\alpha))$.
(3) Conclude the statements on the norm and trace by comparing with the appropriate coefficients of $\det(xI - m_\alpha)$.

Next, prove the assertion on inseparable extensions using (1) and (3), together with the following: if L/K is a finite inseparable extension, then there is a tower $L/E/K$, with E containing all elements separable over K, and every element of L inseparable over E. Moreover, the characteristic must be p for some p, and the degree $[L : E]$ is a multiple of p.

Exercise 1.7. Fix an $n \in \mathbb{N}$. Observe that if $m = p_1^{e_1} \cdots p_\ell^{e_\ell}$, and for each i, either e_i is even, or p_i can be written in the form $x^2 + ny^2$, then m can be written in this form. Show that the converse is also true if $n = 1$ or 2.

On the other hand, show by counterexample that the converse is false if $n = 5$.

CHAPTER 2

Dedekind Domains

We next introduce the important concept of Dedekind domains, which we study in this chapter and the next, before turning to a more focused study of rings of integers. Rings of integers are an important class of Dedekind domains, but other examples include rings of polynomial functions on smooth algebraic curves. In this chapter, we prove two fundamental facts about Dedekind domains: every nonzero ideal can be factored uniquely as a product of prime ideals; and the set of "fractional ideals" form a group under multiplication. We prove these statements by examining the local structure of Dedekind domains. The group structure allows us to introduce the "ideal class group", which measures how far a Dedekind domain is from being a UFD, and plays a fundamental role throughout algebraic number theory.

2.1. The failure of unique factorization

We motivate our results on Dedekind domains by recalling our study of primes of the form x^2+ny^2. We saw in Exercise 1.2 that $p=x^2+ny^2$ if and only if p can be factored in $\mathbb{Z}[\sqrt{-n}]$. However, this is most useful when $\mathbb{Z}[\sqrt{-n}]$ has unique prime factorization, as in the case of $\mathbb{Z}[i]$, where we were able to analyze precisely when p factors. Similarly, we recall that Lamé gave a proof of Fermat's Last Theorem which mistakenly relied on unique factorization in $\mathbb{Z}[\zeta_p]$.

It turns out that unique factorization doesn't hold very often in either the imaginary quadratic or cyclotomic case. However, if instead of considering factorization of elements, we consider factorization of ideals, we will find that we do have unique factorization into prime ideal in rings of integers. Indeed, this was what led Dedekind to introduce the notion and terminology of ideals, in an admittedly primitive form. In order to study irreducibility of elements, it then makes sense to study the factorizations into prime ideals, and then to analyze when the prime ideals that appear are principal, leading to factorizations into elements. A prototypical restatement is the following elementary fact:

Proposition 2.1.1. *A prime p is of the form x^2+ny^2 if and only if*

(i) *$p\mathbb{Z}[\sqrt{-n}]$ factors as $\mathfrak{p}_1\mathfrak{p}_2$ for two (not necessarily distinct) prime ideals of $\mathbb{Z}[\sqrt{-n}]$;*
(ii) *Both $\mathfrak{p}_1, \mathfrak{p}_2$ are principal ideals.*

Proof. We already saw in Exercise 1.2 that $p=x^2+ny^2$ if and only if p is reducible in $\mathbb{Z}[\sqrt{-n}]$, and that in this case, both factors would have norm p. Thus, by Corollary 1.6.9 and Lemma 1.6.7 they generate prime ideals, so the proposition is simply a rephasing of the result of the exercise. □

We note that this situation doesn't quite fall into the situation of rings of integers or more generally Dedekind domains; see Exercise 2.3 below. However, it will turn out to remain well-behaved for p not dividing $2n$.

Factorization into prime ideals turns out to be relatively straightfoward; indeed, we will see already in the next chapter how to analyze when (i) occurs. However, (ii) is far subtler, and will require the full machinery of class field theory to study effectively. That said, by shifting our attention from elements to ideals, we will ultimately be able to obtain a good answer to our original question. Similarly, by studying the structure of ideals in $\mathbb{Z}[\zeta_p]$, Kummer was able to give his partial results on Fermat's Last Theorem, and his arguments can be extended to yield an algorithm which allows one to check in finite time, for any fixed prime p, whether or not $x^p + y^p = z^p$ has solutions.

2.2. Dedekind domains

We begin by defining:

Definition 2.2.1. An integral domain A is a **Dedekind domain** if it satisfies:

(i) A is Noetherian;
(ii) Every non-zero prime ideal of A is maximal;
(iii) A is integrally closed in its field of fractions.

Recall that for A to be Noetherian means that every ascending chain of ideals stabilizes.

We will use the following to see that every ring of integers is a Dedekind domain.

Proposition 2.2.2. *Suppose that A is an integral domain, integrally closed in its field of fractions, and that for any non-zero ideal I, we have that A/I is finite. Then A is a Dedekind domain.*

Proof. Since A is integrally closed by hypothesis, we need only check that it is Noetherian and that every non-zero prime ideal is maximal. But if I is a non-zero ideal, since A/I is finite, and the ideals of A containing I are in bijection with the ideals of A/I, there can only be finitely many such ideals, and any ascending chain containing I stabilizes. Thus, A is Noetherian.

Similarly, if $\mathfrak{p}$ is a non-zero prime ideal of A, then $A/\mathfrak{p}$ is a finite integral domain, and it is a general fact that any finite integral domain is a field, so that $\mathfrak{p}$ is maximal. Indeed, let A be a finite integral domain, and $a \in A$ a non-zero element. Then we must have $a^{k_1} = a^{k_2}$ for some $k_1 > k_2$, by finiteness. Because A is an integral domain, we find $a^{k_1-k_2} = 1$, so a is invertible, and A is a field. □

Using Proposition 1.4.7 and Corollary 1.6.4 to apply the previous proposition to the case of rings of integers, we conclude:

Corollary 2.2.3. *A ring of integers $\mathscr{O}_K$ is a Dedekind domain.*

Algebraic Geometry Remark 2.2.4. From an algebraic geometry perspective, we see that the definition of a Dedekind domains means that it has dimension one and is normal; i.e., we can think of it as a nonsingular curve. Subrings of rings of integers such as $\mathbb{Z}[\sqrt{-3}]$ are still curves, but have singularities, and their containment in the ring of integers corresponds to the normalization map. As in the geometric situation, in order to study the singular curve it is frequently helpful to start by studying the normalization, so we focus primarily on the rings of integers themselves.

2.3. Properties of Dedekind domains

From a number-theoretic point of view, the two most basic properties of Dedekind domains involve the behavior of their ideals under multiplication. We state the main results, and use them to define the ideal class group. The first is:

THEOREM 2.3.1. *Every non-zero ideal of a Dedekind domain may be uniquely factored as a product of prime ideals, up to reordering.*

The second requires a definition:

DEFINITION 2.3.2. Let R be an integral domain with fraction field K. We say that $I \subseteq K$ is a **fractional ideal** of R if it is closed under addition and under scalar multiplication by elements of R, and if there exists a non-zero $d \in R$ such that $dI \subseteq R$. A fractional ideal is **principal** if it is of the form αR, for some $\alpha \in K$.

Given fractional ideals I, J of R, the product IJ is defined to be

$$\{\alpha \in L : \alpha = \sum_{\ell} i_\ell j_\ell, i_\ell \in I, j_\ell \in J\}$$

The product of two fractional ideals is easily seen to be a fractional ideal.

THEOREM 2.3.3. *The set of fractional ideals of a Dedekind domain R form a group under multiplication, with R as the identity.*

We will prove these two theorems in §2.5 and §2.6 below. However, we observe that the second theorem may be equivalently stated as saying that for any ideal I of R, there exists another ideal J such that IJ is principal. Equivalently, ideals modulo principal ideals form a group, called the "ideal class group". This is therefore the first step in understanding the relationship between all ideals and principal ideals. However, having gone to the trouble to define fractional ideals, we make the definition as follows:

DEFINITION 2.3.4. Given a Dedekind domain R, the **ideal class group** of R is defined to be the group of fractional ideals modulo the group of principal fractional ideals.

Thus, the ideal class group measures how far a Dedekind domain is from being a principal ideal domain. We claim in our context, this is equivalent to measuring how far away every irreducible element is from being prime. Specifically:

PROPOSITION 2.3.5. *Let R be an integral domain.*

(i) *If R is Noetherian, then R is a unique factorization domain if and only if every irreducible element is prime.*

(ii) *R is a principal ideal domain if and only if R is a Dedekind domain and a unique factorization domain.*

Note that although this is intended as motivation for Theorems 2.3.1 and 2.3.3, the only part which requires either one is the $\Leftarrow$ direction of (ii); we will use the other statements in the proofs of the theorems.

PROOF. For (i), we have already observed that a UFD has every irreducible element prime, by definition. Conversely, given an element $x \in R$, we claim we can write x as a product of irreducibles: if not, it is clear that we can write $x = x_1 y$, with neither of x_1, y a unit, and where x_1 cannot be written as a product of irreducibles.

Repeating this inductively, we obtain a sequence x_i with $x_0 = x$, and $x_{i+1}|x_i$ for each i, with $\frac{x_{i+1}}{x_i}$ not a unit in R. But then the ideals generated by the x_i form an infinite ascending chain, contradicting the hypothesis that R is Noetherian. Thus, x may be factored into irreducibles, and it is easy to check inductively that this factorization is unique, using that every irreducible is prime.

For (ii), it follows from Exercise 2.2 that every PID is Noetherian. Next, recall that in a PID, every irreducible element is prime: if x is irreducible, and $\mathfrak{m}$ is any prime ideal containing x, because $\mathfrak{m}$ is principal it must be equal to (x). This shows by (i) that any PID is a UFD. But the same argument, if x is a generator for any non-zero prime ideal, shows that (x) is maximal. We already showed that any UFD is integrally closed, so every PID is also a Dedekind domain, as desired.

For the converse, note that by Theorem 2.3.1, it suffices to show that every prime ideal is principal. But given a non-zero prime ideal $\mathfrak{p}$, we must have that $\mathfrak{p}$ contains an irreducible element (x), by starting with any non-zero element, factoring it into irreducibles, and applying the definition of prime ideal inductively. But then this element is prime, and since every non-zero prime ideal is maximal, and $(x) \subseteq \mathfrak{p}$, we conclude that $(x) = \mathfrak{p}$, as desired. □

Finally, we mention that Exercises 2.3 and 2.4 give counterexamples to the two theorems for rings which are not Dedekind domains, but which are relatively close.

2.4. Dedekind domains and DVRs

There are several proofs of Theorems 2.3.1 and 2.3.3. The most classical approach only works for rings of integers, and first proves that the ideal class group is finite, and concludes these theorems. See [**5**, §12.2]. A slightly less direct, but more general proof is given in [**7**, §1.6]. We will take a more technology-heavy approach of proving these theorems via study of local rings. This is a bit longer, but has the advantage of introducing local rings and the concept of constructing global data from local data.

We now explore the properties of local rings of Dedekind domains. Recall the following definitions:

Definition 2.4.1. Let R be an integral domain with field of fractions K, and S a multiplicatively closed subset not containing 0. We define $S^{-1}R \subseteq K$ to be the subring of the form $\{\frac{r}{s} : r \in R, s \in S\}$. If $\mathfrak{p}$ is a prime ideal of R, we define $R_{\mathfrak{p}}$, the **local ring** of R at $\mathfrak{p}$, to be $S^{-1}R$ with $S = R \smallsetminus \mathfrak{p}$.

Definition 2.4.2. An integral domain R' is said to be a **discrete valuation ring** or DVR if it is a principal ideal domain with a unique maximal ideal.

We have already seen in Proposition 2.3.5 that every PID is Dedekind, so in particular every DVR is Dedekind. It follows that the only prime ideals of a DVR are (0) and the maximal ideal.

We next prove two converses to the statement that a DVR is Dedekind. The first is the following.

Lemma 2.4.3. *Any Dedekind domain R' with a unique maximal ideal is a discrete valuation ring. Every non-zero ideal of a discrete valuation ring is a power of the maximal ideal.*

Proof. From the definitions, it suffices to check that R' is a PID. We first claim that the maximal ideal $\mathfrak{m}$ is principal: choose $a \in \mathfrak{m}$ non-zero; by Exercise

2.5, $\exists n \in \mathbb{N}$ such that $\mathfrak{m}^n \subseteq (a)$, but $\mathfrak{m}^{n-1} \not\subseteq (a)$. Choose $b \in \mathfrak{m}^{n-1} \smallsetminus (a)$, and consider $x = \frac{a}{b} \in K$, the field of fractions of R'. From the construction, we see that $x^{-1} \notin R'$, but $x^{-1}\mathfrak{m} \subseteq R'$. Now, since R' is integrally closed, we find that x^{-1} is not integral over R', so since $\mathfrak{m}$ is finitely generated, $x^{-1}\mathfrak{m} \not\subseteq \mathfrak{m}$, by Lemma 1.4.8 But $x^{-1}\mathfrak{m} \subseteq R'$ means it is an ideal of R', so if it is not contained in $\mathfrak{m}$, it must be equal to R', and we conclude that $\mathfrak{m} = (x)$.

We next show that $\mathfrak{m} = (x)$ implies that every ideal is principal. We first claim that every irreducible element is prime, and more precisely, of the form xu for some unit u. But given $y \in R'$ irreducible, because y is not a unit, $y \subseteq \mathfrak{m}$, so $x|y$, and by the definition of irreducibility, $y = xu$ for some unit u, as desired. By Proposition 2.3.5 (i), it follows that R' is a UFD, and we see that every element of R' may be written as $x^n u$ for some $n \in \mathbb{N} \cup \{0\}$, u a unit. If I is a non-zero ideal of R', let $n \in \mathbb{N} \cup \{0\}$ be $\min_{a \in I}\{n' : a = x^{n'}u\}$; it is then clear that $I = (x)^n$. Thus R' is a DVR, as desired.

We have further shown that in R', every non-zero ideal is a power of the maximal ideal, and since we already saw that every DVR is Dedekind, we conclude that this holds in every DVR. □

We can now prove the following stronger statement, characterizing Dedekind domains in terms of their local rings.

PROPOSITION 2.4.4. *Let R be a Noetherian integral domain which is not a field. Then R is a Dedekind domain if and only if for all non-zero prime ideals $\mathfrak{p}$, the local ring $R_{\mathfrak{p}}$ is a discrete valuation ring.*

We first give a lemma in more generality than is needed here, for later use. The general form requires the following definition:

DEFINITION 2.4.5. Given a fractional ideal $I \subseteq K$, and $\mathfrak{p}$ a non-zero prime ideal of R, denote by $I_{\mathfrak{p}} \subseteq K$ the fractional ideal of $R_{\mathfrak{p}}$ described by $I_{\mathfrak{p}} = R_{\mathfrak{p}} I$ (observe that any denominator for I is a denominator for $I_{\mathfrak{p}}$).

We find that in general, a fractional ideal is determined by its local images.

LEMMA 2.4.6. *For an integral domain R, if I is a fractional ideal of R, then $I = \cap_{\mathfrak{p}} I_{\mathfrak{p}}$, where the intersection is taken over all prime ideals of R.*

PROOF. Take $x = a/b \in \cap_{\mathfrak{p}} I_{\mathfrak{p}}$. Let $J = \{y \in R : ya \in bI\}$. This is certainly an ideal of R, and we claim it cannot be contained in any prime ideal $\mathfrak{p}$ of R: indeed, since $x \in I_{\mathfrak{p}}$, we can write $a/b = c/d$ for some $d \notin \mathfrak{p}$, and $c \in I$, so by definition, we have $d \in J \smallsetminus \mathfrak{p}$. Since J is not contained in any prime ideal, it must be all of R, and in particular, $1 \cdot a \in bI$, so $x = a/b \in I$. □

We now have all the pieces to prove the proposition.

PROOF OF PROPOSITION. Suppose that R is Dedekind, and take $\mathfrak{p}$ a non-zero prime ideal of R. Then by Exercise 2.1, we have that $R_{\mathfrak{p}}$ is Noetherian, and moreover the only non-zero prime ideal of $R_{\mathfrak{p}}$ is necessarily $\mathfrak{p}$, and in particular is maximal. Finally, $R_{\mathfrak{p}}$ must be integrally closed: denote by K the field of fractions of R, and suppose $x \in K$ is integral over $R_{\mathfrak{p}}$. Now x satisfies some monic polynomial $x^n + a_{n-1}x^{n-1} + \dots a_0 = 0$ with each $a_i \in R_{\mathfrak{p}}$, and we can therefore clear denominators to get an equation $sx^n + a'_{n-1}x^{n-1} + \dots a'_0 = 0$ where $s \in R \smallsetminus \mathfrak{p}$, and

$a'_i \in R$. Multiplying through by s^{n-1} gives that sx is integral over R and hence in R, so x is in $R_\mathfrak{p}$ as desired.

Conversely, let $\mathfrak{p}$ be any non-zero prime ideal of R, and $\mathfrak{m}$ a maximal ideal containing $\mathfrak{p}$; by hypothesis, $R_\mathfrak{m}$ is a DVR, so $\mathfrak{p} = \mathfrak{m}$ in $R_\mathfrak{m}$ and hence in R by Exercise 2.1. Next, an element $x \in K$ integral over R is integral over every $R_\mathfrak{p}$, so must actually be an element of $R_\mathfrak{p}$ by hypothesis. The previous lemma, in the case $I = R$, then completes the proof. □

2.5. The group of fractional ideals

We are now ready to prove Theorem 2.3.3, that the set of fractional ideals of a Dedekind domain form a group under under multiplication.

Definition 2.5.1. For fractional ideals I, J of R, denote by $(I : J)$ the set $\{x \in K : xJ \subseteq I\}$. Also denote by I' the set $(R : I)$.

Note that by Exercise 2.6, when R is Noetherian we have that $(I : J)$ is again a fractional ideal.

Lemma 2.5.2. *Let I, J be fractional ideals of an integral domain R. Then if $I_\mathfrak{p} J_\mathfrak{p} = R_\mathfrak{p}$ for all non-zero $\mathfrak{p}$, we have $IJ = R$.*

Proof. From Exercise 2.8 and Lemma 2.4.6, we have

$$IJ = \cap_\mathfrak{p} (IJ)_\mathfrak{p} = \cap_\mathfrak{p} I_\mathfrak{p} J_\mathfrak{p} = \cap_\mathfrak{p} R_\mathfrak{p} = R.$$

□

We finally prove that the fractional ideals of a Dedekind domain R form a group under multiplication.

Proof of Theorem 2.3.3. We claim that if I is a fractional ideal of R, then I' is the inverse of I. Indeed, by Exercises 2.7 and 2.8, $I'_\mathfrak{p} I_\mathfrak{p} = (I_\mathfrak{p})' I_\mathfrak{p} = R_\mathfrak{p}$, since we know that $R_\mathfrak{p}$ is a PID. Then by the lemma, $I'I = R$, as desired. □

2.6. Unique factorization of ideals

We now prove Theorem 2.3.1, that every ideal in a Dedekind domain can be factored uniquely into a product of prime ideals. We assume throughout this section that R is a Dedekind domain.

Lemma 2.6.1. *If $\mathfrak{p} \supset \mathfrak{p}_1 \cdots \mathfrak{p}_k$, then $\mathfrak{p} = \mathfrak{p}_i$ for some i.*

Proof. $R/\mathfrak{p}$ is an integral domain, and in it, $\mathfrak{p}_1 \cdots \mathfrak{p}_k = 0$. Now, the product of non-zero ideals can never be 0 in an integral domain, so some $\mathfrak{p}_i = 0$, and is contained in $\mathfrak{p}$. But $\mathfrak{p}_i$ is maximal, so we actually have $\mathfrak{p}_i = \mathfrak{p}$, as desired. □

Definition 2.6.2. If I is an ideal of R, the **valuation** of I at $\mathfrak{p}$, $\nu_\mathfrak{p}(I)$, is defined to be k, where $IR_\mathfrak{p}$ is the kth power of the maximal ideal of $R_\mathfrak{p}$. We have $\nu_\mathfrak{p}(IJ) = \nu_\mathfrak{p}(I) + \nu_\mathfrak{p}(J)$, so for I a fractional ideal of R with $dI \subseteq R$, we can define $\nu_\mathfrak{p}(I) = \nu_\mathfrak{p}(dI) - \nu_\mathfrak{p}((d))$.

$\nu_\mathfrak{p}$ thus gives a homomorphism from the group of fractional ideals to $\mathbb{Z}$. The existence of this map justifies the terminology discrete valuation ring.

Lemma 2.6.3. *Given a non-zero $x \in R$, x is contained in only finitely many prime ideals.*

Proof. Note that any descending chain of ideals containing x must stabilize: if $R \supset I_1 \supset I_2 \supset \ldots (x)$, since the fractional ideals form a group we get an inverse chain (preserving strict inclusion) $R \subseteq I_1^{-1} \subseteq I_2^{-1} \subseteq \ldots (x^{-1})$. But $(x^{-1}) \cong R$ as R-modules, and is in particular Noetherian, so this chain must stabilize.

Now, if $\mathfrak{p}_1, \mathfrak{p}_2, \mathfrak{p}_3, \cdots \supset (x)$, then we get $\mathfrak{p}_1 \supset \mathfrak{p}_1 \cap \mathfrak{p}_2 \supset \mathfrak{p}_1 \cap \mathfrak{p}_2 \cap \mathfrak{p}_3 \supset \ldots (x)$, which must stabilize, meaning for some k, and all $i > k$, $\mathfrak{p}_i \supset \mathfrak{p}_1 \cap \mathfrak{p}_2 \cap \ldots \mathfrak{p}_k \supset \mathfrak{p}_1 \cdot \mathfrak{p}_2 \cdots \mathfrak{p}_k$, which, by the above lemma, implies that $\mathfrak{p}_i = \mathfrak{p}_j$ for some $j \leqslant k$. □

Corollary 2.6.4. *For any fractional ideal I of R, $\nu_{\mathfrak{p}}(I) = 0$ for all but finitely many $\mathfrak{p}$.*

Proof. First suppose that I is an actual ideal of R. It is clear that $\nu_{\mathfrak{p}} > 0 \Leftrightarrow \mathfrak{p} \supset I$, and choosing any non-zero element $x \in I$, applying the above lemma and using $I \supset (x)$ immediately gives the desired result. But now if I is a fractional ideal, choose $d \in R$ so that $dI \subseteq R$, and we find that both $\nu_{\mathfrak{p}}(d)$ and $\nu_{\mathfrak{p}}(dI)$ are non-zero at all but finitely many places, and then it follows that $\nu_{\mathfrak{p}}(I)$ also has finite support. □

This gives us all the necessary tools to prove that every ideal in a Dedekind domain can be factored uniquely into a product of prime ideals.

Proof of Theorem 2.3.1. Given a fractional ideal I, and setting $e_i = \nu_{\mathfrak{p}_i}(I)$, we know by the previous corollary that only finitely many of the e_i are non-zero, so we can define $J = \prod \mathfrak{p}_i^{e_i}$. By Lemma 2.4.6, a fractional ideal of R is determined by its local images, so we see that $I = J$, and I is a product of prime ideals, as desired. □

We conclude with a simple corollary of the theorem.

Corollary 2.6.5. *$I \subseteq J$ if and only if $J|I$.*

Proof. Containment is preserved by localization, so since the multiplicity of each prime factor of I and J is determined by localizing, each prime factor of J has exponent less than or equal to that of I, so $J|I$. The converse is trivial. □

2.7. Exercises

Exercise 2.1. If one considers the map from ideals of $R_{\mathfrak{p}}$ to ideals of R given by $I \mapsto I \cap R$, this map is injective, and if restricted to prime ideals, gives a bijection between prime ideals of $R_{\mathfrak{p}}$ and prime ideals of R contained in $\mathfrak{p}$.

In particular, if R is Noetherian, then $R_{\mathfrak{p}}$ is Noetherian.

Exercise 2.2. Show that a ring R is Noetherian if and only if every ideal is finitely generated.

Exercise 2.3. Show that in $\mathbb{Z}[\sqrt{-3}]$, the ideal generated by 2 cannot be written as a product of prime ideals. Show that the ideal $(2, 1 + \sqrt{-3})$ does not have any fractional ideal inverse.

Exercise 2.4. Show that in $k[x, y]$, where k is any field, the ideal (x, y^2) cannot be written as a product of prime ideals. Show that (x, y) does not have any fractional ideal inverse.

Exercise 2.5. Show that if R is a Noetherian ring, and I an ideal, such that there exists a unique prime ideal $\mathfrak{p}$ containing I, then I contains some power of $\mathfrak{p}$.

Hint: first show that if an element x is contained in every prime ideal of a ring, then $x^n = 0$ for some n. Do this by considering the family of all ideals such that no power of x lies in them. Next, consider the ring R/I.

Exercise 2.6. Show that if K is the field of fractions of a Noetherian integral domain R, and $I \subseteq K$ is an R-module, then I is a fractional ideal of R if and only if I is finitely generated.

Show also that in this situation, if I and J are fractional ideals of R, then $(I : J)$ is a fractional ideal of R.

Exercise 2.7. Show that in any integral domain R, if a fractional ideal I has an inverse, it must be given by I'.

Check that if R is a PID, then I' is the inverse of I.

Exercise 2.8. Check that for fractional ideals I, J, and a prime ideal $\mathfrak{p}$, we have $I_{\mathfrak{p}} J_{\mathfrak{p}} = (IJ)_{\mathfrak{p}}$, and if R is Noetherian, $(I_{\mathfrak{p}} : J_{\mathfrak{p}}) = (I : J)_{\mathfrak{p}}$.

CHAPTER 3

Extensions of Dedekind domains

We now shift the focus of our study from individual Dedekind domains to extensions of Dedekind domains, and in particular the behavior of prime ideals under such extensions. This is one of the central ideas in number theory: for instance, we have already seen that the factorization of an integer prime p in the ring $\mathbb{Z}[i]$ determines whether or not p is a sum of squares. We introduce the important related notion of ramification of prime ideals in an extension, and relate ramification to discriminants.

3.1. Degrees and prime factorizations

In this chapter, we will frequently find ourselves in the following:

Situation 3.1.1. Let R be a Dedekind domain with fraction field K, and L an extension of K of degree n. Let S be a Dedekind domain containing R, with fraction field L, and finitely generated as an R-module.

The most basic result governing factorization of prime ideals in extensions of Dedekind domains is the following:

Theorem 3.1.2. *In Situation 3.1.1, let $\mathfrak{p}$ be a prime ideal of R, and let*

$$\mathfrak{p}S = \prod_{i=1}^{m} \mathfrak{q}_i^{e_i}$$

be the factorization of $\mathfrak{p}S$ into prime ideals in S. Then we have a ring isomorphism

$$S/\mathfrak{p}S \cong \prod_{i=1}^{m} S/\mathfrak{q}_i^{e_i},$$

and if we denote by f_i the dimension of $S/\mathfrak{q}_i$ over $R/\mathfrak{p}_i$, then each f_i is finite, and we have

$$n = \sum_i e_i f_i.$$

In order to prove the theorem, we formally recall the following theorem classifying modules over a PID:

Theorem 3.1.3. *Let R be a PID, and M a finitely generated R-module. Then $M \cong \bigoplus_{i=1}^{m} R/(x_i) \oplus R^{\ell}$ for some m, ℓ, and $x_i \in R$.*

We also recall the general version of the Chinese remainder theorem:

Theorem 3.1.4. *Let R be a ring, and $I_1, \dots, I_n$ ideals which are pairwise coprime (i.e., $I_i + I_j = R$ for all $i \neq j$). Then the natural map*

$$R/(I_1 \cdots I_n) \to \prod_{i=1}^{n} R/I_i$$

is an isomorphism.

Proposition 3.1.5. *In Situation 3.1.1, we have that $R_{\mathfrak{p}}S \subseteq L$ is a free $R_{\mathfrak{p}}$-module of rank n.*

Note that $R_{\mathfrak{p}}S$ is not a local ring unless $\mathfrak{p}$ has only one prime factor in S.

Proof. We first need to observe that since S is assumed to be finitely generated over R, it follows that $R_{\mathfrak{p}}S$ is finitely generated over $R_{\mathfrak{p}}$. Therefore, we can apply Theorem 3.1.3, and we only need to check that the rank is correct, and to see that $R_{\mathfrak{p}}S$ cannot have any summands of the form $R_{\mathfrak{p}}/(x)$ for $x \in R_{\mathfrak{p}}$. But for any non-zero $x \in R_{\mathfrak{p}}$, since all rings in the picture are integral domains, multiplication by x is injective, so $R_{\mathfrak{p}}S$ cannot have a summand of the form $R_{\mathfrak{p}}/(x)$. Only the rank assertion remains: we have $R_{\mathfrak{p}}S \cong R_{\mathfrak{p}}^{n'}$ for some n', and we want to see that $n' = n$. But suppose that $x_1, \dots, x_{n'}$ are a basis of $R_{\mathfrak{p}}S$ over $R_{\mathfrak{p}}$; we claim that they are also a basis for L over K. Indeed, the x_i must span L over K, since $K[x_1, \dots, x_{n'}]$ is a field contained in L and containing S. However, we also see that they are linearly independent, since any non-trivial linear relation over K yields a non-trivial relation over $R_{\mathfrak{p}}$ by clearing denominators. Thus, the x_i form a basis of L over K, and we conclude $n' = n$, as desired. □

We can now prove the theorem.

Proof of Theorem 3.1.2. The isomorphism follows immediately from the Chinese remainder theorem, as soon as we verify that $\mathfrak{q}_i^{e_i} + \mathfrak{q}_j^{e_j} = S$ for all $i \neq j$. But this is clear: $\mathfrak{q}_i$ is the only maximal ideal containing $\mathfrak{q}_i^{e_i}$, and $\mathfrak{q}_j$ is the only maximal ideal containing $\mathfrak{q}_j^{e_j}$, so no maximal ideal contains $\mathfrak{q}_i^{e_i} + \mathfrak{q}_j^{e_j}$.

It therefore remains only to prove that the f_i are finite, and $n = \sum_i e_i f_i$. Noting that $R/\mathfrak{p} \cong R_{\mathfrak{p}}/\mathfrak{p}$ and $S/\mathfrak{p}S \cong R_{\mathfrak{p}}S/\mathfrak{p}R_{\mathfrak{p}}S$, we see from the proposition that $S/\mathfrak{p}S$ is an n-dimensional vector space over $R/\mathfrak{p}$. In particular, since each $S/\mathfrak{q}_i$ is a quotient of $S/\mathfrak{p}S$, we have that the f_i are finite as asserted. Furthermore, using our ring isomorphism we see it is enough to check that $S/\mathfrak{q}_i^{e_i}$ has dimension $e_i f_i$ over $R/\mathfrak{p}$ for each i.

Slightly more generally, we prove that if we have $\mathfrak{q} \subseteq S$ prime and $e > 0$ such that $\mathfrak{p}S \subseteq \mathfrak{q}^e$, then $S/\mathfrak{q}^e$ has dimension ef over $R/\mathfrak{p}$, where $f := \dim_{R/\mathfrak{p}} S/\mathfrak{q}$. The proof proceeds by induction on e: the case $e = 1$ is clear, so we assume the statement is known for $e - 1$. It then suffices to see that $\mathfrak{q}^{e-1}/\mathfrak{q}^e$ has dimension f over $R/\mathfrak{p}$, but here we use that $S/\mathfrak{q} \cong S_{\mathfrak{q}}/\mathfrak{q}S_{\mathfrak{q}}$ and $\mathfrak{q}^{e-1}/\mathfrak{q}^e \cong \mathfrak{q}^{e-1}S_{\mathfrak{q}}/\mathfrak{q}^e S_{\mathfrak{q}}$, and that $S_{\mathfrak{q}}$ is a DVR, so $\mathfrak{q}S_{\mathfrak{q}} = (\pi)$ for some $\pi \in S_{\mathfrak{q}}$. We then see that multiplication by π^{e-1} gives an isomorphism $S_{\mathfrak{q}}/\mathfrak{q}S_{\mathfrak{q}} \overset{\sim}{\to} \mathfrak{q}^{e-1}S_{\mathfrak{q}}/\mathfrak{q}^e S_{\mathfrak{q}}$, so we conclude that the dimensions over $R/\mathfrak{p}$ are the same, as desired. □

Specializing to the case of rings of integers, we conclude the following:

Theorem 3.1.6. *Let K be a number field. Given a prime ideal $\mathfrak{p}$ in $\mathscr{O}_K$, and L/K of degree n, write*

$$\mathfrak{p}\mathscr{O}_L = \prod_{i=1}^{m} \mathfrak{q}_i^{e_i}.$$

Then

$$\mathscr{O}_L/\mathfrak{p}\mathscr{O}_L \cong \prod_{i=1}^{m} \mathscr{O}_L/\mathfrak{q}_i^{e_i}$$

and

$$\sum_{i=1}^{m} e_i f_i = n,$$

where f_i is determined by $N(\mathfrak{q}_i) = N(\mathfrak{p})^{f_i}$.

Proof. This follows from Theorem 3.1.2: we need only check that $\mathscr{O}_L$ is finitely generated as an $\mathscr{O}_K$-module and has fraction field L, but we know from Proposition 1.6.1 that the fraction field is L and that in fact $\mathscr{O}_L$ is finitely generated as a $\mathbb{Z}$-module, so this is clear. □

We will frequently consider this sort of situation, so we define:

Definition 3.1.7. In Situation 3.1.1, we say that a non-zero prime ideal $\mathfrak{q}$ of S **lies above** $\mathfrak{p} \subseteq R$ if $\mathfrak{p} = \mathfrak{q} \cap R$. We say that a prime ideal $\mathfrak{p}$ of R is **ramified** in S if some $\mathfrak{q}$ in S occurs with multiplicity greater than 1 in the factorization of $\mathfrak{p}S$, or if $S/\mathfrak{q}$ is an inseparable extension of $R/\mathfrak{p}$. Finally, if $\mathfrak{p}$ is unramified, we say that $\mathfrak{p}$ is **inert** in S if there is a unique prime $\mathfrak{q}$ lying over $\mathfrak{p}$, and $\mathfrak{p}$ **splits completely** if for every $\mathfrak{q}$ lying over $\mathfrak{p}$, we have $S/\mathfrak{q} = R/\mathfrak{p}$.

Note that if $S/\mathfrak{q}$ and $R/\mathfrak{p}$ are finite, as in the case of rings of integers, then $S/\mathfrak{q}$ cannot be inseparable over $R/\mathfrak{p}$, so $\mathfrak{p}$ is ramified if and only if some $\mathfrak{q}$ occurs with multiplicity greater than 1 in the factorization of $\mathfrak{p}S$.

We have:

Proposition 3.1.8. *In Situation 3.1.1, every non-zero prime ideal $\mathfrak{q}$ of S lies above a unique prime ideal $\mathfrak{p}$ of R. Moreover, $\mathfrak{q}$ lies above a non-zero $\mathfrak{p}$ if and only if $\mathfrak{q}$ occurs in the factorization of $\mathfrak{p}S$, if and only if $\mathfrak{q}$ contains $\mathfrak{p}$.*

Proof. Clearly, if $\mathfrak{p} = \mathfrak{q} \cap R$, then $\mathfrak{q}$ contains $\mathfrak{p}$. Furthermore, $\mathfrak{q}$ contains $\mathfrak{p}$ if and only if $\mathfrak{q}$ occurs in the factorization of $\mathfrak{p}S$, by Corollary 2.6.5. Finally, if $\mathfrak{q}$ contains $\mathfrak{p}$, then because $\mathfrak{p}$ is maximal, the following lemma shows both that $\mathfrak{q}$ lies over $\mathfrak{p}$, and that every $\mathfrak{q}$ does lie over a (necessarily unique) prime $\mathfrak{p}$. □

Lemma 3.1.9. *Let $R \subseteq S$ be an extension of rings, and $\mathfrak{q}$ a prime ideal of S. Then $\mathfrak{q} \cap R$ is a prime ideal of R.*

Proof. Given $x, y \in R$ with $xy \in \mathfrak{q} \cap R$, since $\mathfrak{q}$ is prime we have one of x, y in $\mathfrak{q}$, hence in $\mathfrak{q} \cap R$, so $\mathfrak{q} \cap R$ is prime. □

Example 3.1.10. Let's consider the case of $\mathbb{Q}[i]/\mathbb{Q}$. We know that a prime p remains prime in $\mathbb{Z}[i]$ if and only if $p \equiv 3 \pmod 4$. In this case, we have $e_1 = 1$, $f_1 = 2$, as $N(p\mathbb{Z}[i]) = p^2 = N((p))^2$. Otherwise we can write p as $(a+bi)(a-bi)$ for some $a, b \in \mathbb{Z}$. $N((a+bi)) = N((a-bi)) = p$, so we have $f_i = 1$ in these cases. If $p = 2$, observe that $(1+i) = (1-i)$, so the number of prime ideals lying above 2 is 1, and $e_1 = 2$. Otherwise, we claim that $(a+bi) \neq (a-bi)$. This is equivalent to the assertion that $a - bi$ doesn't divide $a + bi$ in $\mathbb{Z}[i]$. We can write $(a+bi)/(a-bi) = (a^2 - b^2 + 2abi)/p$, so if $a - bi$ divides $a + bi$, we must have $p|2ab$. But p can't divide a or b, so if $p > 2$, this can't happen.

Notation 3.1.11. In Situation 3.1.1, fix $\mathfrak{p} \subseteq \mathfrak{q}$ in $R \subseteq S$. Then we denote by $e_{\mathfrak{q}/\mathfrak{p}}$ the ramification index of $\mathfrak{q}$ over $\mathfrak{p}$, and set $f_{\mathfrak{q}/\mathfrak{p}}$ to be the dimension of $S/\mathfrak{q}$ over $R/\mathfrak{p}$.

We conclude with a statement on multiplicativity in towers:

Proposition 3.1.12. *Given $S/T/R$ Dedekind domains with fraction fields $L/E/K$, with S a finitely generated T-module and T a finitely generated R-module, and $\mathfrak{p} \subseteq \mathfrak{q} \subseteq \mathfrak{r}$ prime ideals of $R \subseteq T \subseteq S$, we have*

$$e_{\mathfrak{r}/\mathfrak{p}} = e_{\mathfrak{r}/\mathfrak{q}} e_{\mathfrak{q}/\mathfrak{p}}, \text{ and } f_{\mathfrak{r}/\mathfrak{p}} = f_{\mathfrak{r}/\mathfrak{q}} f_{\mathfrak{q}/\mathfrak{p}}.$$

Proof. Both identities follow easily from the definitions. For the first, we compare $\mathfrak{p}S$ to $(\mathfrak{p}T)S$. For the second, we have that $S/\mathfrak{r}$ is $f_{\mathfrak{r}/\mathfrak{q}}$-dimensional over $T/\mathfrak{q}$, and $T/\mathfrak{q}$ is $f_{\mathfrak{q}/\mathfrak{p}}$-dimensional over $R/\mathfrak{p}$, so $S/\mathfrak{r}$ is $(f_{\mathfrak{r}/\mathfrak{q}} f_{\mathfrak{q}/fp})$-dimensional over $R/\mathfrak{p}$. □

3.2. Relative discriminants

We want to show that discriminants can be used to understand ramification in rings of integers. To do this, it will be helpful to have more general notions of discriminants, such as a relative discriminant $D_{L/K}$ for an extension of number fields L/K. However, in our arguments we will need more general definitions. The definition which we will apply directly to extensions of rings of integers (and which works more generally for extensions of Dedekind domains) is the following:

Definition 3.2.1. Suppose that S contains R, and both are integral domains, with fraction fields L and K respectively. The **discriminant** $D_{S/R} \subseteq R$ is the ideal generated by elements of the form $D_{L/K}((x_i)_i)$ as the x_i are allowed to range over all elements of S.

We will also be interested in modding out by certain ideals, which will not in general yield integral domains. We therefore make the following additional definition.

Definition 3.2.2. Suppose S contains R, and is free over R. The **discriminant** of S/R is the ideal $D_{S/R} \subseteq R$ generated by elements of the form $D_{S/R}((x_i)_i)$ as the x_i are allowed to range over all elements of S.

We should check that these definitions do not conflict in the case that both are defined, and that they are compatible in an appropriate sense with our earlier definitions.

Proposition 3.2.3. *We have the following basic properties of the discriminant:*

(i) *In the case that S is free over R and they are both integral domains, the two definitions of $D_{S/R}$ agree.*
(ii) *In the case that S is free over R, we have that $D_{S/R}$ is the principal ideal generated by $D_{S/R}(x_1, \ldots, x_n)$, with the x_i any basis of S over R.*
(iii) *Suppose that $K = \mathbb{Q}$, and L is a number field. Then $D_{\mathscr{O}_L/\mathbb{Z}}$ is the ideal generated by the absolute discriminant D_L.*

Proof. (i) is clear from the definitions, noting only that if $x \in S$, we have $\mathrm{Tr}_{L/K}(x) = \mathrm{Tr}_{S/R}(x)$, since any basis of S over R is in particular a basis of L of K, so multiplication by x can be represented by the same matrix.

(ii) Certainly, $D_{S/R}$ contains $D_{S/R}(x_1, \ldots, x_n)$, but it follows from Proposition 1.5.4 that $D_{S/R}(y_1, \ldots, y_n)$ for any y_i is a multiple of $D_{S/R}(x_1, \ldots, x_n)$, so we conclude that $D_{S/R}$ is indeed generated by $D_{S/R}(x_1, \ldots, x_n)$.

(iii) follows immediately from (ii). □

From (iii), we see that in applying our relative definition to number fields, all we lost was the information of the sign of D_K. From the point of view of ramification, this won't matter. Note however that we are not asserting that $D_{\mathscr{O}_L/\mathscr{O}_K}$ is a principal ideal in general, when $\mathscr{O}_L$ is not free over $\mathscr{O}_K$.

3.3. The discriminant and ramification

A key tool in algebraic number theory is the relationship between the discriminant and ramification. This relationship allows us to work much more effectively in using the discriminant to carry out a range of computations.

We first review some of the basic properties of the discriminant, which we will use to relate it to ramification. We continue with the notation of the preceding section.

Recall from Lemma 1.5.9 that if L is a separable extension of K, we have $D_{L/K} \neq 0$.

We will also want to see the following two properties:

LEMMA 3.3.1. *If S is a direct sum of rings S_1 and S_2, each free over R, then*

$$D_{S/R} = D_{S_1/R} D_{S_2/R}.$$

PROOF. We apply (i) of the previous proposition. Let $x_1, \dots, x_{n_1}$ and $y_1, \dots, y_{n_2}$ be bases for S_1 and S_2 respectively; then since $x_i y_j = 0$ in S, we have $\mathrm{Tr}_{S/R}(x_i y_j) = 0$, and also $\mathrm{Tr}_{S/R}(x) = \mathrm{Tr}_{S_i/R}(x)$ for $x \in S_i$. We thus see that the matrix whose determinant defines the discriminant is block diagonal, with blocks consisting of $(\mathrm{Tr}_{S_1/R}(x_i x_j))_{i,j}$ and $(\mathrm{Tr}_{S_2/R}(y_i y_j))_{i,j}$, so its determinant is the product of $D_{S_1/R}$ and $D_{S_2/R}$, as desired. □

LEMMA 3.3.2. *If R is a field, and S has any nilpotent elements (i.e., non-zero x with $x^m = 0$ for some m), then $D_{S/R} = 0$.*

PROOF. Given x such that $x^m = 0$, since R is a field, we can make a basis $x = x_1, x_2, \dots, x_n$ of S over R. We claim that $\mathrm{Tr}_{S/R}(x_1 x_i) = 0$ for all i: indeed, since repeated multiplication by $x_1 x_i$ gives 0, all eigenvalues of $m_{x_1 x_i}$ must be 0, and it follows that the trace is 0. Thus, the first row of the matrix $\mathrm{Tr}_{S/R}(x_i x_j)$ is 0, so the determinant, and hence $D_{S/R}$, is 0. □

We also need to know that discriminants commute with modding out by ideals.

LEMMA 3.3.3. *Let R be a ring, and S a ring containing R, and free over R of rank n. Let I be an ideal of R, and denote reduction mod I (or mod IS) by $x \mapsto \bar{x}$. Then $D_{(S/IS)/(R/I)} = \bar{D}_{S/R}$.*

PROOF. It suffices to observe that if $x_1, \dots, x_n$ are a basis for S over R, then $\bar{x}_i, \dots, \bar{x}_n$ are a basis for S/IS over R/I. The statement follows by definition-chasing. □

Finally, discriminants behave well with respect to localization:

LEMMA 3.3.4. *Let $R \subseteq S$ be an extension of integral domains, and $\mathfrak{p}$ a prime ideal of R. Then $D_{R_\mathfrak{p} S/R_\mathfrak{p}} = R_\mathfrak{p} D_{S/R}$.*

PROOF. Certainly, we have $R_\mathfrak{p} D_{S/R} \subseteq D_{R_\mathfrak{p} S/R_\mathfrak{p}}$, since for any $(x_1, \dots, x_n)$ in S, we have $D_{S/R}(x_1, \dots, x_n) = D_{R_\mathfrak{p} S/R_\mathfrak{p}}(x_1, \dots, x_n)$. On the other hand, given

$(y_1, \dots, y_n)$ in $R_{\mathfrak{p}}S$, for each i we have $y_i = \frac{x_i}{z_i}$ with $x_i \in S$, and $z_i \in R \smallsetminus \mathfrak{p}$. By Proposition 1.5.4 we have

$$D_{R_{\mathfrak{p}}S/R_{\mathfrak{p}}}(y_1, \dots, y_n) = (\prod_i \frac{1}{z_i})^2 D_{R_{\mathfrak{p}}S/R_{\mathfrak{p}}}(x_1, \dots, x_n) = (\prod_i \frac{1}{z_i})^2 D_{S/R}(x_1, \dots, x_n) \in R_{\mathfrak{p}} D_{S/R},$$

so we obtain the opposite containment as well. □

With these basic properties out of the way, we now return to the case of extensions of Dedekind domains, where we prove:

Theorem 3.3.5. *In Situation 3.1.1, let $\mathfrak{p}$ be a prime of R. Then $\mathfrak{p}$ is ramified in S if and only if $\mathfrak{p}$ divides $D_{S/R}$.*

Proof. Write $\mathfrak{p}S = \prod_i \mathfrak{q}_i^{e_i}$ for distinct primes $\mathfrak{q}_i$. Also, recall that $R_{\mathfrak{p}}/\mathfrak{p}R_{\mathfrak{p}} = R/\mathfrak{p}$, and $R_{\mathfrak{p}}S/\mathfrak{p}R_{\mathfrak{p}}S = S/\mathfrak{p}S$, and that $R_{\mathfrak{p}}S$ is a free $R_{\mathfrak{p}}$-module of rank n. We first claim that $D_{S/R}$ is contained in $\mathfrak{p}$ if and only if $D_{(S/\mathfrak{q}_i^{e_i})/(R/\mathfrak{p})} = 0$ for some i. Since $D_{(S/\mathfrak{p}S)/(R/\mathfrak{p})} = \prod_i D_{(S/\mathfrak{q}_i^{e_i})/(R/\mathfrak{p})}$, the latter is equivalent to $D_{(S/\mathfrak{p}S)/(R/\mathfrak{p})} = 0$, and by the previous lemmas, this is equivalent to $D_{R_{\mathfrak{p}}S/R_{\mathfrak{p}}}$ being contained in $\mathfrak{p}R_{\mathfrak{p}}$, which is equivalent to $D_{S/R}$ being contained in $\mathfrak{p}$.

Thus we conclude that $D_{S/R}$ is contained in $\mathfrak{p}$ if and only if $D_{(S/\mathfrak{q}_i^{e_i})/(R/\mathfrak{p})} = 0$ for some i.

Now, suppose that $\mathfrak{p}$ is unramified: then all the e_i are 1, and each $S/\mathfrak{q}_i$ is separable over $R/\mathfrak{p}$. Hence $S/\mathfrak{p}S$ is a product of fields, each separable over $R/\mathfrak{p}$, and we thus have $D_{(S/\mathfrak{q}_i^{e_i})/(R/\mathfrak{p})} = D_{(S/\mathfrak{q}_i)/(R/\mathfrak{p})} \neq 0$ for all i by Lemma 1.5.9. Thus, we conclude by the above that $D_{S/R}$ is not contained in $\mathfrak{p}$.

Conversely, suppose that some $e_i > 1$; then $S/\mathfrak{q}_i^{e_i}$ has nilpotent elements (any $x \in \mathfrak{q}_i S \smallsetminus \mathfrak{q}_i^{e_i}$), so by Lemma 3.3.2, $D_{(S/\mathfrak{q}_i^{e_i})/(R/\mathfrak{p})} = 0$. On the other hand, if every $e_i = 1$ but some $S/\mathfrak{q}_i$ is inseparable over $R/\mathfrak{p}$, we again have $D_{(S/\mathfrak{q}_i^{e_i})/(R/\mathfrak{p})} = D_{(S/\mathfrak{q}_i)/(R/\mathfrak{p})} = 0$ by Lemma 1.5.9. Either way, by the above we have that $D_{S/R}$ is contained in $\mathfrak{p}$. □

We conclude the following basic fact:

Corollary 3.3.6. *In situation 3.1.1, if further L is separable over K, there are only finitely many prime ideals of R ramified in S.*

Proof. This follows from the theorem, and the fact that $D_{S/R}$ is not the zero ideal, since L over K is separable and S contains a basis of L over K. □

3.4. Explicit factorization

The goal of this section is to provide explicit formulas for factoring prime ideals in extensions, in terms of minimal polynomials of elements in the extension ring. These formulas work in many, but not all cases.

We will need the following simple fact relating discriminants of minimal polynomials to discriminants of ring extensions.

Lemma 3.4.1. *In situation 3.1.1, suppose we have $\alpha \in S$ such that $L = K(\alpha)$, and let $f(x)$ be the monic minimal polynomial for α over K. Then* $\operatorname{disc} f(x) \in D_{S/R}$.

PROOF. We already proved in Lemma 1.5.7 that $\operatorname{disc} f(x) = D_{L/K}(1, \alpha, \ldots, \alpha^{n-1})$, and since the α^i are in S, by the definition of $D_{S/R}$ we find that it contains $\operatorname{disc} f(x)$. □

In order to make some pretense of avoiding notational clutter, we we introduce the following:

NOTATION 3.4.2. We write $S_{\mathfrak{p}} = R_{\mathfrak{p}} S \subseteq L$.

The main theorem is then the following:

THEOREM 3.4.3. *In situation 3.1.1, given a prime ideal $\mathfrak{p}$ of R, suppose there exists $\alpha \in S$ such that $S_{\mathfrak{p}} = R_{\mathfrak{p}}[\alpha]$, and let $f(x)$ be the monic minimal polynomial for α. Factor $\bar{f}(x) = \prod_{i=1}^m \bar{f}_i(x)^{e_i}$ in $(R/\mathfrak{p})[x]$, with the $\bar{f}_i$ distinct and irreducible. Then we have:*

(i) *$\mathfrak{p}S$ factors as $\mathfrak{p}S = \prod_{i=1}^m \mathfrak{q}_i^{e_i}$, with $S/\mathfrak{q}_i$ of dimension $\deg \bar{f}_i$ over $R/\mathfrak{p}$.*
(ii) *The $\mathfrak{q}_i$ are explicitly given by $\mathfrak{q}_i = (\mathfrak{p}, \tilde{f}_i(\alpha))$, where $\tilde{f}_i(x)$ is any polynomial in $R[x]$ whose reduction mod $\mathfrak{p}$ is $\bar{f}_i(x)$.*

Finally, given $\mathfrak{p}$, suppose instead that we have an $\alpha \in S$ such that its monic minimal polynomial $f(x)$ over K satisfies $((\operatorname{disc} f) : D_{S/R})$ is not contained in $\mathfrak{p}$. Then $S_{\mathfrak{p}} = R_{\mathfrak{p}}[\alpha]$, so the above conclusions hold.

PROOF. We first observe:

$$S/(\mathfrak{p}) \cong S_{\mathfrak{p}}/(\mathfrak{p}) \cong R_{\mathfrak{p}}[\alpha]/(f(\alpha), \mathfrak{p}) \cong (R/\mathfrak{p})[\bar{\alpha}]/(\bar{f}(\bar{\alpha})).$$

It follows that the given $\mathfrak{q}_i$ in (ii) are distinct, since the only relations on $1, \bar{\alpha}, \ldots, \bar{\alpha}^i, \ldots$ over $R/\mathfrak{p}$ are generated by $\bar{f}(\bar{\alpha})$, so for any $i \neq j$, since the degrees of $\bar{f}_i(x)$ and $\bar{f}_j(x)$ are less than the degree of $\bar{f}(x)$, we find that $\bar{f}_i(\bar{\alpha}) \neq \bar{f}_j(\bar{\alpha})$, so $(\mathfrak{p}, \tilde{f}_i(\alpha)) \neq (\mathfrak{p}, \tilde{f}_j(\alpha))$.

We also see that the $\mathfrak{q}_i$ are in fact prime ideals, $S/\mathfrak{q}_i$ of dimension $\deg \bar{f}_i(x)$ over $R/\mathfrak{p}$. Indeed, we have

$$S/(\mathfrak{p}, \tilde{f}_i(\alpha)) \cong (R/\mathfrak{p})[\alpha]/(\bar{f}(\alpha), \bar{f}_i(\alpha)) \cong (R/\mathfrak{p})[\alpha]/(\bar{f}_i(\alpha)),$$

and since $\bar{f}_i(x)$ is irreducible in $R/\mathfrak{p}$ by hypothesis, this is a field extension of $R/\mathfrak{p}$, of degree equal to $\deg \bar{f}_i(x)$, as desired.

We next note that if $\mathfrak{q}$ is prime and contains $\mathfrak{p}$, then it must be one of the $\mathfrak{q}_i$. Indeed, we have $S/\mathfrak{q} = (R/\mathfrak{p})[\bar{\alpha}]/(\bar{\mathfrak{q}}, \bar{f}(\bar{\alpha}))$, and since $(R/\mathfrak{p})[\bar{\alpha}]$ is a PID, we have $(\bar{\mathfrak{q}}, \bar{f}(\bar{\alpha})) = \bar{g}(\bar{\alpha})$ for some $\bar{g}(x) \in (R/\mathfrak{p})[x]$. Since $\mathfrak{q}$ is prime, $\bar{g}(x)$ must be irreducible so that the quotient is an integral domain. But then we have $\bar{g}(x) | \bar{f}(x)$ and is irreducible, so it must be $\bar{f}_i(x)$ for some i, and it follows that $\mathfrak{q} = \ker(S \to S/\mathfrak{q}) = (\mathfrak{p}, \tilde{f}_i(\alpha))$, as desired.

Thus, we know that we have $\mathfrak{p}S = \prod_{i=1}^m \mathfrak{q}_i^{e_i'}$ for some e_i', and it remains to see that $e_i' = e_i$. Note that

$$\prod_{i=1}^m \mathfrak{q}_i^{e_i} \subseteq (\mathfrak{p}, \prod_{i=1}^m \tilde{f}_i(\alpha)^{e_i}) \subseteq \mathfrak{p}S,$$

with the last inclusion because $\prod_{i=1}^m \tilde{f}_i(\alpha)^{e_i} \equiv f(\alpha) = 0 \pmod{\mathfrak{p}}$. Hence $e_i \geqslant e_i'$ for all i, using unique factorization into prime ideals, and the containment relation.

Finally, the statement that $e_i = e'_i$ for all i follows because we have

$$\sum_{i=1}^m e_i f_i = \deg f(x) = [L : K] = \sum_{i=1}^m e'_i f_i.$$

It remains to check the last assertion, that if we have α satisfying $f(x)$ with $((\operatorname{disc} f) : D_{S/R})$ not contained in $\mathfrak{p}$, that $S_\mathfrak{p} = R_\mathfrak{p}[\alpha] \subseteq L$. Recall from Lemma 3.4.1 that $\operatorname{disc} f \in D_{S/R}$, so $((\operatorname{disc} f) : D_{S/R})$ is an ideal of R. Also, by Lemma 3.3.4 we have $R_\mathfrak{p} D_{S/R} = D_{S_\mathfrak{p}/R_\mathfrak{p}}$. Thus, the condition that $((\operatorname{disc} f) : D_{S/R})$ is not contained in $\mathfrak{p}$ is equivalent to the condition that $((\operatorname{disc} f(x)) : D_{S_\mathfrak{p}/R_\mathfrak{p}}) = R_\mathfrak{p}$. We then have $(D_{S_\mathfrak{p}/R_\mathfrak{p}}(1, \alpha, \dots, \alpha^{n-1})) = (\operatorname{disc} f(x)) = D_{S_\mathfrak{p}/R_\mathfrak{p}}$, and applying the usual argument involving the change of basis formula, we find that $1, \alpha, \dots, \alpha^{n-1}$ is a basis of $S_\mathfrak{p}$ over $R_\mathfrak{p}$, as desired. □

Note that it follows from this argument that $(\mathfrak{p}, \tilde{f}_i(\alpha))^{e_i} \supset \mathfrak{p}$, a fact by no means obvious.

The second part of theorem is often applied when $\operatorname{disc} f$ itself is not in $\mathfrak{p}$, in which case one need not know anything about S itself. We restate this case as a corollary:

Corollary 3.4.4. *In situation 3.1.1, given a prime ideal $\mathfrak{p}$ of R, suppose there exists $\alpha \in S$ such that $L = K(\alpha)$, and let $f(x)$ be the monic minimal polynomial for α. Factor $\bar{f}(x) = \prod_{i=1}^m \bar{f}_i(x)^{e_i}$ in $(R/\mathfrak{p})[x]$, and suppose further that $\bar{f}(x)$ is separable. Then we have:*

(i) *$\mathfrak{p}S$ factors as $\mathfrak{p}S = \prod_{i=1}^m \mathfrak{q}_i$, with $S/\mathfrak{q}_i$ of dimension $\deg \bar{f}_i$ over $R/\mathfrak{p}$.*
(ii) *The $\mathfrak{q}_i$ are explicitly given by $\mathfrak{q}_i = (\mathfrak{p}, \tilde{f}_i(\alpha))$, where $\tilde{f}_i(x)$ is any polynomial in $R[x]$ whose reduction mod $\mathfrak{p}$ is $\bar{f}_i(x)$.*

In particular, $\mathfrak{p}$ is unramified in S.

Proof. The first part of the corollary is a direct application of the theorem, since $\bar{f}(x)$ is separable if and only if $\operatorname{disc} f \not\in \mathfrak{p}$. We also note that if $\bar{f}(x)$ is separable, we have $e_i = 1$ for all i, and each $\bar{f}_i(x)$ is separable. To see that $\mathfrak{p}$ is unramified, we need only see that each $S/\mathfrak{q}_i$ is separable over $R/\mathfrak{p}$, which follows from the assertion of the theorem that $S_\mathfrak{p} = R_\mathfrak{p}[\alpha]$, and the hypothesis that $\bar{f}_i$ is separable. □

We conclude with an example in the case of cyclotomic extensions. We recall some preliminary definitions.

Definition 3.4.5. Let K be any field, with algebraic closure $\bar{K}$, and fix $n \geqslant 1$ relatively prime to $\operatorname{char} K$, if the latter is non-zero. Let μ_n be the multiplicative group of nth roots of unity in $\bar{K}$, and $\zeta \in \mu_n$ a primitive nth root of unity (i.e., a generator of μ_n). We define $\Phi_n(x)$ by

$$\Phi_n(x) = \prod_{0<i<n:(i,n)=1} (x - \zeta^i).$$

We also recall the following basic properties of cyclotomic extensions:

Proposition 3.4.6. *We have*

$$\prod_{d|n} \Phi_d(x) = x^n - 1.$$

The polynomial $\Phi_n(x)$ is separable, independent of the choice of ζ, and lies in $K[x]$. Indeed, when $K = \mathbb{Q}$ we have $\Phi_n(x) \in \mathbb{Z}[x]$, and for any other field $\Phi_n(x)$ is given by the natural map $\mathbb{Z} \to K$.

$K(\zeta)$ is a Galois extension of K, and there is a natural injection $\mathrm{Gal}(K(\zeta)/K) \hookrightarrow \mathrm{Aut}(\mu_n) \cong (\mathbb{Z}/n\mathbb{Z})^\times$.

Proof. $\Phi_n(x)$ is independent of the choice of ζ because it may be rewritten as the product over all primitive nth roots of unity ζ' of $(x-\zeta')$. We note that $\Phi_n(x)$ divides $x^n - 1 = \prod_{i=1}^n (x-\zeta)$, and that if we separate the product over $d = n/(i,n)$, we obtain the asserted factorization of $x^n - 1$. Furthermore, $x^n - 1$ is separable, since its derivative is nx^{n-1} and we have assumed that n is non-zero in K. Thus, $\Phi_n(x)$ is separable. Furthermore, the minimal polynomial of ζ must divide $x^n - 1$, so all of its roots lie in $K(\zeta)$, and we conclude that $K(\zeta)$ is both separable and normal, hence Galois.

The natural injection $\mathrm{Gal}(K(\zeta)/K) \hookrightarrow \mathrm{Aut}(\mu_n)$ is clear: $K(\zeta)$ contains μ_n, so there is a homomorphism $\mathrm{Gal}(K(\zeta)/K) \to \mathrm{Aut}(\mu_n)$, but this map must be injective because an automorphism of $K(\zeta)$ over K is determined by where it sends ζ. We also note that the isomorphism $\mathrm{Aut}(\mu_n) \cong (\mathbb{Z}/n\mathbb{Z})^\times$ is determined by the choice of ζ: since ζ generates μ_n, an element of $\mathrm{Aut}(\mu_n)$ is uniquely determined by where it sends ζ, and ζ must be sent to ζ^i for $(i,n) = 1$ in order to determine an automorphism.

Finally, we see that $\Phi_n(x)$ lies in $K[x]$: any element of the Galois group of $K(\zeta)/K)$ permutes the primitive nth roots of unity, so leaves $\Phi_n(x)$ invariant. Thus the coefficients of $\Phi_n(x)$ are fixed by $\mathrm{Gal}(K(\zeta)/K)$, and lie in K. In the case $K = \mathbb{Q}$, since each ζ^i is integral, we find that $\Phi_n(x) \in \mathbb{Z}[x]$. For the final assertion, that in general $\Phi_n(x)$ is obtained from the coefficients in the case $K = \mathbb{Q}$, we use induction on n, with the base case $n = 1$ trivial. But using the factorization $\prod_{d|n} \Phi_d(x) = x^n - 1$, which holds both over $\mathbb{Z}$ and in K, and applying the induction hypothesis to all $d < n$, we immediately conclude the desired statement for n as well. □

The proposition implies that the minimal polynomial of ζ over K divides $\Phi_n(x)$. However, they need not be equal: for instance if K is algebraically closed, $\zeta \in K$ already, while if $K = \mathbb{F}_{p^r}$, we have that $\mathrm{Gal}(K(\zeta)/K)$ is generated by Frobenius, so we see that the image of $\mathrm{Gal}(K(\zeta)/K)$ in $(\mathbb{Z}/n\mathbb{Z})^\times$ is the subgroup generated by p. However, we will see in Corollary 3.8.6 below that if $K = \mathbb{Q}$, then we do have that $\Phi_n(x)$ is irreducible, so that $[K(\zeta) : K] = \varphi(n)$. In our proof of this fact, a key step is the following example of our results thus far.

Example 3.4.7. Fix $n > 1$. Then we've seen that for any p not dividing n, we have that $\Phi_n(x)$ remains separable modulo p, so it follows from Corollary 3.4.4 that p is unramified in $\mathscr{O}_{\mathbb{Q}(\zeta_p)}$.

3.5. Factorization of primes in Galois extensions

We begin by refining our analysis of factorization of prime ideals in the case that an extension is Galois. This will relate the factorization of prime ideals more closely to Galois theory, which will play an important role in our analysis of Dirichlet characters. Throughout this section, we suppose we are in Situation 3.1.1, with the additional hypothesis that L/K is a Galois extension.

Proposition 3.5.1. *For any prime ideal $\mathfrak{p} \subseteq R$, if we write*

$$\mathfrak{p}S = \prod_{i=1}^{g} \mathfrak{q}_i^{e_i}$$

and set f_i to be the dimension of $S/\mathfrak{q}_i$ over $R/\mathfrak{p}$, we have that $\operatorname{Gal}(L/K)$ acts transitively on the $\mathfrak{q}_i$, and $e_i = e_j$ and $f_i = f_j$ for all i, j. If we write $e = e_i$, $f = f_i$ for all i, we have $n = efg$.

Lemma 3.5.2. *Any $\sigma \in \operatorname{Gal}(L/K)$ induced an automorphism of S.*

Proof. We claim that S is the integral closure of R in L (without any Galois assumptions). Indeed, since S is integrally closed in L, it certainly contains the integral closure of R. On the other hand, since S is a finitely generated R-module, by Lemma 1.4.8 it consists entirely of elements integral over R, and must be equal to the integral closure.

The assertion that σ induces an automorphism of S then follows, since σ must send elements integral over R to elements integral over R, and the same is true of σ^{-1}. □

Proof of proposition. First, we claim that for $\sigma \in \operatorname{Gal}(L/K)$, we have $\sigma(\mathfrak{q}_1) = \mathfrak{q}_i$ for some i. Indeed, from the lemma we know that σ maps S to itself as a ring automorphism, so we can write $\sigma(\mathfrak{q}_1) = (\sigma^{-1})^{-1}(\mathfrak{q}_1)$ (where the second $^{-1}$ denotes preimage of a set), and by the usual argument this is a prime ideal. Moreover, $\mathfrak{q}_1 \cap R = \mathfrak{p}$ is fixed by σ, so $\sigma(\mathfrak{q}_1)$ contains $\mathfrak{p}$ and must be one the $\mathfrak{q}_i$.

To see transitivity, suppose that for some j, there does not exist σ such that $\mathfrak{q}_j = \sigma(\mathfrak{q}_1)$; then, recalling that we had

$$S/\mathfrak{p}S \cong \prod_{i=1}^{g} S/(\mathfrak{q}_i^{e_i}),$$

we can choose $x \in \mathfrak{q}_j$ such that $x \notin \sigma^{-1}(\mathfrak{q}_1)$ for all $\sigma \in \operatorname{Gal}(L/K)$. Then $N_{L/K}(x)$ is a product of elements not in $\mathfrak{q}_1$, hence is not in $\mathfrak{p} = \mathfrak{q}_1 \cap R$. On the other hand, since $x \in \mathfrak{q}_j$, we have $N_{L/K}(x) \in \mathfrak{q}_j \cap R$, which is a contradiction.

Finally, the transitivity easily implies that all $e_i = e_j$ and $f_i = f_j$: the first follows by applying $\sigma \in \operatorname{Gal}(L/K)$ to the factorization of $\mathfrak{p}S$ and using uniqueness of the factorization, while the second follows because we must have $S/\mathfrak{q}_i \cong S/\mathfrak{q}_j$ for all i, j. □

This justifies a simplification of notation:

Notation 3.5.3. Given $\mathfrak{p} \subseteq \mathfrak{q}$ in $R \subseteq S$, since $e_{\mathfrak{q}/\mathfrak{p}}$ and $f_{\mathfrak{q}/\mathfrak{p}}$ are independent of $\mathfrak{q}$, we denote them by $e_{\mathfrak{p}}$ and $f_{\mathfrak{p}}$ (as long as there is no possibility of confusion about the extension ring S). We denote by $g_{S,\mathfrak{p}}$ or $g_{\mathfrak{p}}$ be the number of primes lying above $\mathfrak{p}$ in S.

3.6. Decomposition and inertia groups

To better relate the behavior of primes to Galois theory, we define the notions of decomposition and inertia groups. Throughout this section, we continue with the Galois hypotheses of the previous section, except where explicitly specified.

Definition 3.6.1. Let $\mathfrak{p}$ be a prime ideal of R, and $\mathfrak{q}$ a prime ideal of S lying above $\mathfrak{p}$. Then we define the **decomposition group** $D_{\mathfrak{q}/\mathfrak{p}} \subseteq \operatorname{Gal}(L/K)$ of $\mathfrak{q}$ to be the group of σ such that $\sigma(\mathfrak{q}) = \mathfrak{q}$, or equivalently, such that σ induces a map $S/\mathfrak{q} \to S/\mathfrak{q}$. We define the **inertia group** $I_{\mathfrak{q}/\mathfrak{p}} \subseteq D_{\mathfrak{q}/\mathfrak{p}}$ of $\mathfrak{q}$ to be the group of σ such that $\sigma : S/\mathfrak{q} \to S/\mathfrak{q}$ is the identity map.

The fundamental results on the decomposition and inertia groups are the following:

Theorem 3.6.2. *Given $\mathfrak{p}$, for any $\mathfrak{q}$ lying over $\mathfrak{p}$, we have the following:*

(i) *Let $L_{D,\mathfrak{q}}$ be the fixed field of $D_{\mathfrak{q}/\mathfrak{p}}$, and set $S_{D,\mathfrak{q}} = S \cap L_{D,\mathfrak{q}}$ Then $L_{D,\mathfrak{q}}$ is the minimal field E between K and L such that $\mathfrak{q}$ is the only prime ideal lying over $\mathfrak{q} \cap E$ in $S \cap E$. Furthermore, $[L : L_{D,\mathfrak{q}}] = e_{\mathfrak{p}} f_{\mathfrak{p}}$, and if we denote by $\mathfrak{q}_0$ the ideal $\mathfrak{q} \cap S_{D,\mathfrak{q}}$, we have $e_{\mathfrak{q}_0} = e_{\mathfrak{p}}$ and $f_{\mathfrak{q}_0} = f_{\mathfrak{p}}$.*

(ii) *Suppose $S/\mathfrak{q}$ is separable over $R/\mathfrak{p}$. Then there is a natural map $D_{\mathfrak{q}/\mathfrak{p}} \to$ $\operatorname{Gal}((S/\mathfrak{q})/(R/\mathfrak{p}))$ which is surjective, with kernel $I_{\mathfrak{q}/\mathfrak{p}}$. In particular, $I_{\mathfrak{q}/\mathfrak{p}}$ is a normal subgroup of $D_{\mathfrak{q},\mathfrak{p}}$, with $D_{\mathfrak{q}/\mathfrak{p}}/I_{\mathfrak{q}/\mathfrak{p}}$ having order $f_{\mathfrak{p}}$, and $I_{\mathfrak{q}/\mathfrak{p}}$ has order $e_{\mathfrak{p}}$.*

(iii) *Let $L_{I,\mathfrak{q}}$ be the fixed field of $I_{\mathfrak{q}/\mathfrak{p}}$, and set $S_{I,\mathfrak{q}} = S \cap L_{D,\mathfrak{q}}$. Then $L_{I,\mathfrak{q}}$ is the minimal field E between K and L such that $\mathfrak{q} \cap E$ is totally ramified in L. Furthermore, $[L : L_{I,\mathfrak{q}}] = e_{\mathfrak{p}}$, and if we denote by $\mathfrak{q}_1$ the ideal $\mathfrak{q} \cap S_{I,\mathfrak{q}}$, we have $e_{\mathfrak{q}_1} = e_{\mathfrak{p}}$ and $f_{\mathfrak{q}_1} = 1$.*

Because we treat Dedekind domains in full generality, we need the following lemma, which is unnecessary in the case of rings of integers.

Lemma 3.6.3. *In Situation 3.1.1, given a field E between K and L, the ring $S \cap E$ is a Dedekind domain with fraction field E, and is a finitely generated R-module, with S a finitely generated $(S \cap E)$-module.*

Proof. It is immediate that S is a finitely generated $(S \cap E)$-module. The key point is that since R is Noetherian, and S is a finitely generated R-module, the submodule $S \cap E$ must also be finitely generated over R, and is in particular a Noetherian R-module. Ideals of $S \cap E$ are in particular R-submodules, so then satisfy the ascending chain condition, so we find that $S \cap E$ is Noetherian.

Since S is integrally closed in L, it is clear that $S \cap E$ is integrally closed in E. Since any $x \in E$ may be written as $\frac{y}{z}$, with y integral over R and $z \in R$, we have that E is the fraction field of $S \cap E$. Any non-zero prime $\mathfrak{q}$ of S must lie over a non-zero prime of $S \cap E$, since it lies over a non-zero prime of R. Hence, $S \cap E$ is not a field. Finally, by the going up theorem applied to $S \cap E$ inside S, we see that every non-zero prime ideal $S \cap E$ is maximal. □

Proof of theorem. (i) We have $\operatorname{Gal}(L/L_{D,\mathfrak{q}}) = D_{\mathfrak{q}/\mathfrak{p}}$, which fixes $\mathfrak{q}$ by definition. By the transitivity statement applied to $L/L_{D,\mathfrak{q}}$, we conclude that $\mathfrak{q}$ is the unique prime lying over $\mathfrak{q} \cap S_{D,\mathfrak{q}}$. Now suppose E has $\mathfrak{q}$ as the only prime lying over $\mathfrak{q} \cap E$. Then $\operatorname{Gal}(L/E) \subseteq \operatorname{Gal}(L/K)$ fixes $\mathfrak{q}$, so $\operatorname{Gal}(L/E) \subseteq D_{\mathfrak{q}/\mathfrak{p}}$ by definition, and $L_{D,\mathfrak{q}} \subseteq E$, as desired. For the degree statement, we note that because $\operatorname{Gal}(L/K)$ acts transitively on the primes over $\mathfrak{p}$, $D_{\mathfrak{q}/\mathfrak{p}}$ has $g_{\mathfrak{p}}$ cosets in $\operatorname{Gal}(L/K)$. Since $|\operatorname{Gal}(L/K)| = e_{\mathfrak{p}} f_{\mathfrak{p}} g_{\mathfrak{p}}$, we get $|D_{\mathfrak{q}/\mathfrak{p}}| = [L : L_{D,\mathfrak{q}}] = e_{\mathfrak{p}} f_{\mathfrak{p}}$. In particular, since $\mathfrak{q}$ is the unique prime lying over $\mathfrak{q}_0$, we have $e_{\mathfrak{q}_0} f_{\mathfrak{q}_0} = e_{\mathfrak{p}} f_{\mathfrak{p}}$, and by Proposition 3.1.12, we have $e_{\mathfrak{q}_0} \leqslant e_{\mathfrak{p}}$ and $f_{\mathfrak{q}_0} \leqslant f_{\mathfrak{p}}$, so we must have equality.

(ii) We first note that $S/\mathfrak{q}$ is normal, hence Galois over $R/\mathfrak{p}$: indeed, it is generated by elements of the form $\bar{\alpha}$, with $\alpha \in S \subseteq L$, and if $f(x) \in R[x]$ is the monic minimal polynomial of α over K, since L/K every root of $f(x)$ must lie in L, hence in S since S is integrally closed. Now, if $\bar{f}(x)$ is the monic minimal polynomial of $\bar{\alpha}$, we have that $\bar{f}(x)$ divides the reduction modulo $\mathfrak{q}$ of $f(x)$, so all its roots are images of roots of $f(x)$, and hence are contained in $S/\mathfrak{q}$, giving the desired normality.

By definition, the elements of $D_{\mathfrak{q}/\mathfrak{p}}$ act on $S/\mathfrak{q}$, and they must also fix $R/\mathfrak{p}$. Moreover, $I_{\mathfrak{q}/\mathfrak{p}}$ is the kernel of this map by definition. Thus, we need only prove surjectivity. To see this, in the notation of (i), we first observe that it suffices to prove the same statement for $L/L_{D,\mathfrak{q}}$ and $\mathfrak{q}$ lying over $\mathfrak{q}_0$, since by (i) $D_{\mathfrak{q}/\mathfrak{p}} = \operatorname{Gal}(L/L_{D,\mathfrak{q}}) = D_{\mathfrak{q}/\mathfrak{q}_0}$, and $S_{D,\mathfrak{q}}/\mathfrak{q}_0 = R/\mathfrak{p}$. Replacing K by $L_{D,\mathfrak{q}}$ and $\mathfrak{p}$ by $\mathfrak{q}_0$, we may therefore assume that $\mathfrak{q}$ is the unique prime ideal lying over $\mathfrak{p}$.

Now, $S/\mathfrak{q}$ is separable over $R/\mathfrak{p}$, so there is some $\bar{\alpha} \in S/\mathfrak{q}$ generating it as a field over $R/\mathfrak{p}$. Denote by $\bar{f}(x)$ the minimal polynomial of $\bar{\alpha}$, and choose $\alpha \in S$ mapping to $\bar{\alpha}$. If $f(x)$ is the minimal polynomial of α, we have $\bar{f}(x)|\overline{f(x)}$, so every root of $\bar{f}(x)$ is the image of a root of $f(x)$. An automorphism of $S/\mathfrak{q}$ is determined by sending α to another root of $\bar{f}(x)$, and $\operatorname{Gal}(L/K)$ acts transitively on the roots of $f(x)$, so we have that $D_{\mathfrak{q}/\mathfrak{p}} = \operatorname{Gal}(L/K)$ surjects onto $\operatorname{Gal}((S/\mathfrak{q})/(R/\mathfrak{p}))$, as desired. Note that in fact this argument implies that $\overline{f(x)}$ is a power of $\bar{f}(x)$, since we find that $\operatorname{Gal}((S/\mathfrak{q})/(R/\mathfrak{p}))$ acts transitively on the roots of $\overline{f(x)}$.

In particular, returning to the general situation, $I_{\mathfrak{q}/\mathfrak{p}}$ must be normal with $D_{\mathfrak{q}/\mathfrak{p}}/I_{\mathfrak{q}/\mathfrak{p}} = \operatorname{Gal}((S/\mathfrak{q})/(R/\mathfrak{p}))$, so since the latter has order $f_{\mathfrak{p}}$, and since we already knew that $|D_{\mathfrak{q}/\mathfrak{p}}| = e_{\mathfrak{p}} f_{\mathfrak{p}}$, we have the desired assertions.

(iii) First observe that part of the definition of a prime being totally ramified is that a unique prime lies above it, so by (i), any E with $\mathfrak{q} \cap E$ totally ramified in L must contain $L_{D,\mathfrak{q}}$, and is the fixed field of a subgroup of $D_{\mathfrak{q}/\mathfrak{p}}$. Let $G \subseteq D_{\mathfrak{q}/\mathfrak{p}}$ be any subgroup, and denote by L_G the fixed field of G, and $\mathfrak{q}_G$ the prime ideal $\mathfrak{q} \cap L_G$. By (i), we have that $\mathfrak{q}$ is the unique prime lying over $\mathfrak{q}_G$, and we have $G \twoheadrightarrow \operatorname{Gal}((S/\mathfrak{q})/((S \cap L_G)/\mathfrak{q}_G))$, with kernel clearly equal to $G \cap I_{\mathfrak{q}/\mathfrak{p}}$. Thus, $\mathfrak{q}_G$ is totally ramified if and only if $S/\mathfrak{q} = (S \cap L_G)/\mathfrak{q}_G$ if and only if $G/(G \cap I_{\mathfrak{q}/\mathfrak{p}}) = (1)$, if and only if $I_{\mathfrak{q}/\mathfrak{p}} \subseteq G$. In particular, $L_{I,\mathfrak{q}}$ has $\mathfrak{q} \cap S_{I,\mathfrak{q}}$ totally ramified in L, and is the minimal field with this property. □

3.7. Fixed fields in the abelian case

The case of abelian extensions is frequently important in number theory. It contains quadratic and cyclotomic extensions, is the subject of class field theory, and a number of important theorems (including the Tchebotarev density theorem) can be reduced to it. Accordingly, we will want to give precise descriptions of $L_{D,\mathfrak{q}}$ and $L_{I,\mathfrak{q}}$ in the case that $\operatorname{Gal}(L/K)$ is abelian. We continue with the Galois hypotheses of the previous two sections.

Easy Facts 3.7.1. Given $\mathfrak{q}_1, \mathfrak{q}_2$ lying over $\mathfrak{p}$, choose $\sigma \in \operatorname{Gal}(L/K)$ such that $\sigma(\mathfrak{q}_1) = \mathfrak{q}_2$. Then:

(i) $D_{\mathfrak{q}_1/\mathfrak{p}} = \sigma^{-1} D_{\mathfrak{q}_2/\mathfrak{p}} \sigma$.
(ii) $I_{\mathfrak{q}_1/\mathfrak{p}} = \sigma^{-1} I_{\mathfrak{q}_2/\mathfrak{p}} \sigma$.

We see that in this case, $D_{\mathfrak{q}/\mathfrak{p}}$ and $I_{\mathfrak{q}/\mathfrak{p}}$ depend only on $\mathfrak{p}$, so we denote them by $D_{\mathfrak{p}}$, $I_{\mathfrak{p}}$, and their fixed fields by $L_{D,\mathfrak{p}}$, $L_{I,\mathfrak{p}}$.

Theorem 3.7.2. *Suppose* $\mathrm{Gal}(L/K)$ *is abelian. Then we have:*

(i) $L_{D,\mathfrak{p}}$ *is the maximal extension of* K *contained in* L *in which* $\mathfrak{p}$ *splits completely.*

(ii) $L_{I,\mathfrak{p}}$ *is the maximal extension of* K *contained in* L *in which* $\mathfrak{p}$ *is unramified.*

Proof. (i) By (i) of the earlier theorem, we find that the primes over $\mathfrak{p}$ in $S_{D,\mathfrak{p}}$ correspond to those over $\mathfrak{p}$ in L, with all e and f occurring in $L/L_{D,\mathfrak{p}}$. By the multiplicativity of the e and f, we see that $e_{\mathfrak{q}/\mathfrak{p}} = f_{\mathfrak{q}/\mathfrak{p}} = 1$ for all $\mathfrak{q} \subseteq S_{D,\mathfrak{p}}$ lying over $\mathfrak{p}$, so we have that $\mathfrak{p}$ splits completely in $L_{D,\mathfrak{p}}$. Thus, the maximal extension in which $\mathfrak{p}$ splits completely must contain $L_{D,\mathfrak{p}}$.

Conversely, given E such that $\mathfrak{p}$ splits completely, let $\mathfrak{q}$ be a prime of S lying over $\mathfrak{p}$. We claim that $D_{\mathfrak{q},\mathfrak{p}} \subseteq \mathrm{Gal}(L/E)$. Indeed, $D_{\mathfrak{q}/\mathfrak{p}}$ must fix $\mathfrak{q} \cap E$, but $D_{(\mathfrak{q}\cap E)/\mathfrak{p}}$ is trivial since $\mathfrak{p}$ splits completely in E. Thus $D_{\mathfrak{q}/\mathfrak{p}}$ maps to (1) in $\mathrm{Gal}(E/K) = \mathrm{Gal}(L/K)/\mathrm{Gal}(L/E)$, and we get the desired assertion.

(ii) It follows immediately from (iii) of the previous theorem that $\mathfrak{p}$ is unramified over $L_{I,\mathfrak{p}}$, since for each $\mathfrak{q}$ over $\mathfrak{p}$, all the ramification of $\mathfrak{p}$ occurs in the extension L over $L_{I,\mathfrak{q}} = L_{I,\mathfrak{p}}$. It is likewise clear that any field containing $L_{D,\mathfrak{p}}$ in which $\mathfrak{p}$ is unramified must be contained in $L_{I,\mathfrak{p}}$. We can therefore finish the proof by arguing that if $\mathfrak{p}$ is unramified in E, it will remain unramified in $EL_{I,\mathfrak{p}}$, from which we conclude that for E to be maximal, it must contain and hence be equal to $L_{I,\mathfrak{p}}$. Therefore, the following proposition completes the proof of the theorem. □

Proposition 3.7.3. *Given* E, E' *intermediate fields between* L *and* K*, Galois over* K*, with* $EE' = L$*. Given also* $\mathfrak{p}$ *a prime of* K*, suppose* $\mathfrak{p}$ *is unramified in* $S \cap E$ *and* $S \cap E'$*. Then* $\mathfrak{p}$ *is unramified in* S*.*

Proof. First note that since ramification is multiplicative in towers, we can replace K by $E \cap E'$ and R by $S \cap E \cap E'$ without affecting the statement. Since E and E' are Galois over K, they are Galois over $E \cap E'$, and hence linearly disjoint over it. By hypothesis, there exist $(x_1, \dots, x_n)$ in $S \cap E$ generating E over K and such that $D_{E/K}(x_1, \dots, x_n) \notin \mathfrak{p}$. By linear disjointness, $(x_1, \dots, x_n)$ is also a basis of L over E', so $D_{L/E'}(x_1, \dots, x_n) = D_{E/K}(x_1, \dots, x_n)$, and we see that for any $\mathfrak{q}$ lying over $\mathfrak{p}$, $D_{L/E'}$ is not contained in $\mathfrak{q}$. Thus, none of the primes lying over $\mathfrak{p}$ in E' are ramified in L, and since $\mathfrak{p}$ is unramified in E', it follows that $\mathfrak{p}$ is unramified in L, as desired. □

Remark 3.7.4. In fact, the proposition holds without the Galois hypothesis, and in a sharper form. But we defer further discussion of this until we develop the tools of local fields.

3.8. Frobenius elements and cyclotomic fields

We now specialize our situation somewhat, to the case that residue fields are finite. This allows us to define special elements of Galois groups known as Frobenius elements, which generate decomposition groups at unramified primes. As an application of Frobenius elements, we will prove that the cyclotomic polynomials $\Phi_n(x)$ defined earlier are irreducible over $\mathbb{Q}$. We assume the following situation:

Situation 3.8.1. Let R be a Dedekind domain with fraction field K, and L an extension of K of degree n. Let S be a Dedekind domain containing R, with fraction field L, and finitely generated as an R-module. Further assume that for every non-zero prime ideal $\mathfrak{p} \subseteq R$, we have $R/\mathfrak{p}$ finite.

In this situation, if $\mathfrak{q}$ lies over $\mathfrak{p}$, the extension $S/\mathfrak{q}$ over $R/\mathfrak{p}$ is always separable, because finite field are perfect. We can therefore apply Theorem 3.6.2 to see that given $\mathfrak{q}$ lying over $\mathfrak{p}$, the decomposition group $D_{\mathfrak{q}/\mathfrak{p}}$ surjects onto $\operatorname{Gal}((\mathscr{O}_L/\mathfrak{q})/(\mathscr{O}_K/\mathfrak{p}))$, with kernel $I_{\mathfrak{q}/\mathfrak{p}}$. Because Galois groups of extensions of finite fields are cyclic, generated by Frobenius, we have:

Corollary 3.8.2. *The quotient group $D_{\mathfrak{q}/\mathfrak{p}}/I_{\mathfrak{q}/\mathfrak{p}}$ is cyclic, canonically generated by the element mapping to the Frobenius automorphism in* $\operatorname{Gal}((\mathscr{O}_L/\mathfrak{q})/(\mathscr{O}_K/\mathfrak{p}))$.

This motivates the following definition:

Definition 3.8.3. Let $\mathfrak{q}$ be a prime of S lying over $\mathfrak{p}$, and suppose that $e_{\mathfrak{p}} = 1$. Then the **Frobenius element** $\operatorname{Fr}(\mathfrak{q}/\mathfrak{p})$ is defined to be the (unique) element of $D_{\mathfrak{q}/\mathfrak{p}}$ whose image in $\operatorname{Gal}((S/\mathfrak{q})/(R/\mathfrak{p}))$ is the Frobenius automorphism. If $\operatorname{Gal}(L/K)$ is abelian, this depends only on $\mathfrak{p}$, and we denote it by $\operatorname{Fr}(\mathfrak{p})$.

The last part is justified by Easy Facts 3.7.1 and the following:

Easy Fact 3.8.4. Given $\mathfrak{q}_1, \mathfrak{q}_2$ lying over $\mathfrak{p}$, choose $\sigma \in \operatorname{Gal}(L/K)$ such that $\sigma(\mathfrak{q}_1) = \mathfrak{q}_2$, and $e_{\mathfrak{p}} = 1$. $\operatorname{Fr}(\mathfrak{q}_1/\mathfrak{p}) = \sigma^{-1} \operatorname{Fr}(\mathfrak{q}_2/\mathfrak{p})\sigma$.

In the case of an abelian extension, using unique factorization of ideals as products of prime ideals, we therefore obtain map from ideals of R to $\operatorname{Gal}(L/K)$; in the number field case, one can show that this map is surjective, and study its kernel in terms of generalized ideal class groups. This is the subject of class field theory. While it was, appropriately enough, Frobenius who had the idea assigning elements of Galois groups to prime ideals, it was Artin who saw the relationship to class field and ultimately used it to give a very elegant formulation of the subject. We will return to this topic in a later chapter.

We now turn briefly to the study of cyclotomic extensions.

Example 3.8.5. Fix $n > 1$, and ζ a primitive nth root of unity in $\bar{\mathbb{Q}}$. Consider the case $L = \mathbb{Q}(\zeta)$, $K = \mathbb{Q}$, and p any prime not dividing n. By Example 3.4.7, we know p is unramified in $\mathscr{O}_L$. Let $\mathfrak{q}$ be any prime lying above p, and let $\bar{\zeta}$ be the image of ζ in $\mathscr{O}_L/\mathfrak{q}$; we then have an injection $\operatorname{Gal}(\mathbb{Q}(\zeta)/\mathbb{Q}) \hookrightarrow (\mathbb{Z}/n\mathbb{Z})^\times$ determined by which power of itself ζ is sent to, and we claim that $\operatorname{Fr}(p) = p$ under this injection.

Indeed, we have from Proposition 3.4.6 that $\bar{\zeta}$ is a root of $\Phi_n(x)$ for $\mathscr{O}_L/\mathfrak{q}$, hence still a primitive nth root of unity. So we have an injection $\operatorname{Gal}((\mathscr{O}_L/\mathfrak{q}))/(\mathbb{Z}/p\mathbb{Z})) \hookrightarrow (\mathbb{Z}/n\mathbb{Z})^\times$ compatible with the above injection under the isomorphism $D_p \xrightarrow{\sim} \operatorname{Gal}((\mathscr{O}_L/\mathfrak{q}))/(\mathbb{Z}/p\mathbb{Z}))$. But the the Frobenius automorphism in $\operatorname{Gal}((\mathscr{O}_L/\mathfrak{q}))/(\mathbb{Z}/p\mathbb{Z}))$ sends $\bar{\zeta}$ to its pth power by definition, so its image in $(\mathbb{Z}/n\mathbb{Z})^\times$ is p, and we have proved the claim, and described the Frobenius elements for cyclotomic extensions.

We can use this to conclude the following basic result:

Corollary 3.8.6. *For $K = \mathbb{Q}$, the cyclotomic polynomial $\Phi_n(x)$ is irreducible. Equivalently, the map*

$$\operatorname{Gal}(\mathbb{Q}(\zeta)/\mathbb{Q}) \hookrightarrow (\mathbb{Z}/n\mathbb{Z})^\times$$

induced by ζ is surjective.

Proof. We have seen that for p prime to n, we have $\mathrm{Fr}(p) \in \mathrm{Gal}(\mathbb{Q}(\zeta)/\mathbb{Q})$ represented by $p \in (\mathbb{Z}/n\mathbb{Z})^{\times}$. But any integer prime to n is a product of primes prime to n, so such p generate $(\mathbb{Z}/n\mathbb{Z})^{\times}$, proving the desired surjectivity. This implies that the degree $[\mathbb{Q}(\zeta) : \mathbb{Q}]$ is $\varphi(n)$, which is also the degree of $\Phi_n(x)$ by definition, so $\Phi_n(x)$ is the minimal polynomial for ζ and in particular irreducible. □

Of course, there are more elementary proofs of this fact, but these tend to require substantially more work. A related consequence is the following:

Corollary 3.8.7. *Let K be any number field. Then K contains only finitely many roots of unity.*

Proof. Let n be the degree of K over $\mathbb{Q}$. In order for K to contain a primitive dth root of unity, we see from the previous corollary that we must have $n \leqslant \varphi(d)$. But for any fixed n, only finitely many d satisfy this inequality. □

3.9. Orders and the imaginary quadratic case

Before we end our discussion of factorization of primes in extensions, we recall that our original motivation, due to Proposition 2.1.1, was to study when p factors as a product of prime ideals in $\mathbb{Z}[\sqrt{-n}]$. But in n is not square-free or is congruent to 3 modulo 4, we have that $\mathbb{Z}[\sqrt{-n}]$ is strictly contained in $\mathscr{O}_{\mathbb{Q}(\sqrt{-n})}$, so as of yet the theory we have developed does not apply to the case we are interested in. However, if we are willing to restrict our attention to primes p not dividing $2n$, we will not have any problems. We make the following definition:

Definition 3.9.1. We say that a subring $R \subseteq \mathscr{O}_K$ is an **order** if R contains a basis for K over $\mathbb{Q}$.

The basic result we need is the following:

Proposition 3.9.2. *Let R be an order of index m in some $\mathscr{O}_L$ and containing some $\mathscr{O}_K$, and $\mathfrak{p}$ a prime of $\mathscr{O}_K$ not containing m. Then:*

(i) *$R_{\mathfrak{p}} := \mathscr{O}_{K,\mathfrak{p}}R$ is equal to $\mathscr{O}_{L,\mathfrak{p}}$ inside L.*

(ii) *$\mathfrak{p}R$ factors as a product of primes ideals, with the factorization having the same e_i, f_i as in $\mathscr{O}_L$, and in fact being obtained by factoring $\mathfrak{p}\mathscr{O}_L$ and intersecting with R.*

To prove this, it will be helpful to know the following commutative algebra facts:

Commutative Algebra Facts 3.9.3. Fix an integral domain R and multiplicately closed subset S. Then:

(i) For any ideal $I \subseteq S^{-1}R$, we have $S^{-1}(I \cap R) = I$.

(ii) For ideals $I_1, I_2 \subseteq R$, we have $S^{-1}(I_1 I_2) = (S^{-1}I_1)(S^{-1}I_2)$.

(iii) Suppose that $I \subseteq R$ satisfies the property that for any $\mathfrak{p} \subseteq R$ prime such that $\mathfrak{p} \cap S \neq \emptyset$, we have $I_{\mathfrak{p}} = R_{\mathfrak{p}}$. Then $(S^{-1}I) \cap R = I$.

(iv) Given $I_1, \dots, I_g \subseteq S^{-1}R$, suppose that $(I_i \cap R)_{\mathfrak{p}} = R_{\mathfrak{p}}$ for any $\mathfrak{p} \subseteq R$ such that $\mathfrak{p} \cap S \neq \emptyset$. Then $(I_1 \cdots I_g) \cap R = (I_1 \cap R) \cdots (I_g \cap R)$.

Proof. (i) and (ii) are exactly as in Problem Set 2; the arguments there did not depend on having $S = R \smallsetminus \mathfrak{p}$ for a prime ideal $\mathfrak{p}$.

For (iii), it is clear that $I \subseteq (S^{-1}I) \cap R$ without any hypotheses. Thus, for any $\mathfrak{p} \subseteq R$, we have $I_\mathfrak{p} \subseteq ((S^{-1}I) \cap R)_\mathfrak{p}$. We will argue the opposite inclusion for all $\mathfrak{p}$, and conclude that $I = (S^{-1}I) \cap R$ because all ideals are determined locally by Lemma 2.4.6. First, if $\mathfrak{p} \cap S = \emptyset$, then $((S^{-1}I) \cap R)_\mathfrak{p} = (S^{-1}(S^{-1}I \cap R))_\mathfrak{p} = (S^{-1}I)_\mathfrak{p} = I_\mathfrak{p}$, using (i). On the other hand, if $\mathfrak{p} \cap S \neq \emptyset$, $((S^{-1}I) \cap R)_\mathfrak{p} \subseteq R_\mathfrak{p} = I_\mathfrak{p}$ by hypothesis. Thus, we are done.

For (iv), since $(I_i \cap R)_\mathfrak{p} = R_\mathfrak{p}$ for all $\mathfrak{p}$ with $\mathfrak{p} \cap S \neq \emptyset$, we have

$$((I_1 \cap R) \cdots (I_g \cap R))_\mathfrak{p} = (I_1 \cap R)_\mathfrak{p} \cdots (I_g \cap R)_\mathfrak{p} = R_\mathfrak{p}$$

for all such $\mathfrak{p}$. Then by (iii),

$$\begin{aligned}(I_1 \cdots I_g) \cap R &= ((S^{-1}(I_1 \cap R)) \cdots (S^{-1}(I_g \cap R))) \cap R \\ &= (S^{-1}((I_1 \cap R) \cdots (I_g \cap R))) \cap R = (I_1 \cap R) \cdots (I_g \cap R),\end{aligned} \tag{3.9.3.1}$$

as desired. □

Proof of Proposition. For (i), observe that because $\mathscr{O}_L/R$ (as an abelian group under addition) has order m, we have $m\mathscr{O}_L \subseteq R$, or equivalently, $\mathscr{O}_L \subseteq \frac{1}{m}R$. But because $\mathfrak{p}$ doesn't contain m, we have $\frac{1}{m} \in \mathscr{O}_{K,\mathfrak{p}}$, so $\mathscr{O}_{L,\mathfrak{p}} \subseteq \frac{1}{m}R_\mathfrak{p} = R_\mathfrak{p} \subseteq \mathscr{O}_{L,\mathfrak{p}}$, as desired.

For (ii), write $\mathfrak{p}\mathscr{O}_L = \prod_{i=1}^g \mathfrak{q}_i^{e_i}$. Then we also have $\mathfrak{p}\mathscr{O}_{L,\mathfrak{p}} = \prod_{i=1}^g (\mathfrak{q}_i\mathscr{O}_{L,\mathfrak{p}})^{e_i}$, so $\mathfrak{p}R_\mathfrak{p} = \prod_{i=1}^g (\mathfrak{q}_i R_\mathfrak{p})^{e_i}$, and we have $\mathfrak{q}'_i := (\mathfrak{q}_i R_\mathfrak{p}) \cap R$ is prime. We claim that the $\mathfrak{q}_i R_\mathfrak{p}$ satisfy the conditions of (iii) above, so that we get a factorization $(\mathfrak{p}R_\mathfrak{p}) \cap R = \prod_{i=1}^g \mathfrak{q}_i'^{e_i}$. Indeed, $\mathfrak{q}_i \supset \mathfrak{p}$, and given $\mathfrak{p}'$ prime in R and such that $\mathfrak{p}' \cap S \neq \emptyset$, with $S = \mathscr{O}_K \smallsetminus \mathfrak{p}$, this means that $\mathfrak{p}' \cap \mathscr{O}_K \neq \mathfrak{p}$. Since $\mathfrak{p}$ and $\mathfrak{p}' \cap \mathscr{O}_K$ are both non-zero primes, we have $\mathfrak{p}' \not\supset \mathfrak{p}$, so choose $s \in \mathfrak{p}$ but not in $\mathfrak{p}'$. Then $s \in \mathfrak{q}_i$, so $s \in (\mathfrak{q}_i R_\mathfrak{p}) \cap R$, and since s in invertible in $R_{\mathfrak{p}'}$, we have $((\mathfrak{q}_i R_\mathfrak{p}) \cap R)_{\mathfrak{p}'} = R_{\mathfrak{p}'}$, which is the desired condition of (iii) above.

The same argument as above implies that because $\mathfrak{p}R$ contains $\mathfrak{p}$, it satisfies the condition for (ii) above, so $(\mathfrak{p}R_\mathfrak{p}) \cap R = \mathfrak{p}R$, and we have a factorization of $\mathfrak{p}R$ into prime ideals $\mathfrak{q}'_i$. We need to check now that in fact $\mathfrak{q}'_i = \mathfrak{q}_i \cap R$. Certainly, $\mathfrak{q}_i \cap R \subseteq (\mathfrak{q}_i R_\mathfrak{p}) \cap R = \mathfrak{q}'_i$, and for the converse, since $R \subseteq \mathscr{O}_L$, $\mathfrak{q}_i R_\mathfrak{p} = \{\frac{x}{s} : x \in \mathfrak{q}_i, s \in \mathscr{O}_K \smallsetminus \mathfrak{p}\}$. Given such an $\frac{x}{s}$ which is also in R, we have $\frac{x}{s} = y \in R \subseteq \mathscr{O}_L$, so $sy \in \mathfrak{q}_i$, and since $\mathfrak{q}_i$ is prime and $\mathfrak{q}_i \cap \mathscr{O}_K = \mathfrak{p}$, we find $y \in \mathfrak{q}_i$, giving the desired equality.

Finally, since $\mathfrak{q}_i R_\mathfrak{p} = \mathfrak{q}'_i R_\mathfrak{p}$,

$$\#(\mathscr{O}_L/\mathfrak{q}_i) = \#(\mathscr{O}_{L,\mathfrak{p}}/(\mathfrak{q}_i)) = \#(R_\mathfrak{p}/(\mathfrak{q}_i)) = \#(R/\mathfrak{q}'_i),$$

so the f_i are the same. □

3.10. Exercises

Exercise 3.1. a) For all square-free integers n, let $K = \mathbb{Q}(\sqrt{n})$, and compute D_K, and then carefully and completely describe how every prime $p \in \mathbb{Z}$ factors in $\mathscr{O}_K$.

b) Given $n \in \mathbb{Z}$, $n \neq 0, 1$, show that for p prime to $2n$, we have $p\mathbb{Z}[\sqrt{-n}] = \mathfrak{p}_1\mathfrak{p}_2$ in $\mathbb{Z}[\sqrt{-n}]$ if and only if $-n$ has a square root mod p.

Exercise 3.2. Give an elementary proof of the irreducibility of $\Phi_n(x)$ over $\mathbb{Q}$ when n is a prime power by applying the Eisenstein criterion to $\Phi_n(x+1)$ to show that $\Phi_n(x)$ is irreducible over $\mathbb{Z}$, and concluding it is irreducible over $\mathbb{Q}$ by Gauss' lemma.

Exercise 3.3. Fix $n \in \mathbb{N}$: describe in elementary terms which primes $p \in \mathbb{Z}$ not dividing n split completely in $\mathscr{O}_{\mathbb{Q}(\zeta_n)}$.

It is a remarkable consequence of the density theorems that a Galois extension of $\mathbb{Q}$ is completely determined by which primes split completely in it. However, the proof of the following special case of the Dirichlet density theorem is considerably easier:

Exercise 3.4. Show that for any n, there are infinitely many primes congruent to 1 mod n.

Hint: argue by contradiction, and evaluate $\Phi_n(x)$ at large, cleverly chosen values.

Notes

The argument of §3.3 is adapted from Samuel, §5.3 of [**8**].

CHAPTER 4

Lattice techniques: ideal class groups and the structure of the group of units

We now begin in earnest our study of algebraic number theory, proving results on ideal class groups and the group of units that do not hold for general Dedekind domains.

4.1. The immediate goal

We recall that Theorem 2.3.3 allows us to define the ideal class group of a Dedekind domain, and in particular of a ring of integers, as the group of fractional ideals modulo the subgroup of principal ideals. We will prove that in the case of a ring of integers, the ideal class group is finite. In fact, we will shortly give a stronger statement due to Minkowski. Using similar techniques, we will also study the structure of the group of units in a ring of integers.

Before we state our theorems, we introduce some notation:

Definition 4.1.1. Given a number field K, denote by r_1 and $2r_2$ the number of imbeddings of K into $\mathbb{R}$ and into $\mathbb{C}$ (but not $\mathbb{R}$) respectively.

Note that r_2 is an integer, because complex conjugation acts as an involution without fixed points on the set of complex (and not real) imbeddings. Also, we have $r_1 + 2r_2 = n := [K : \mathbb{Q}]$. It is frequently easy to compute r_1 and r_2, due to the following observation:

Lemma 4.1.2. *Let $f(x) \in \mathbb{Q}[x]$ be an irreducible polynomial, with root α, and let $K = \mathbb{Q}(\alpha)$. Then r_1 is the number of real roots of $f(x)$, and r_2 is half the number of non-real roots.*

Proof. r_1 is by definition the number of real imbeddings of K, while r_2 is half the number of complex imbeddings. But $K \hookrightarrow \bar{\mathbb{Q}}$ is contained in $\mathbb{R}$ if and only if α maps to a real number, and the imbeddings of K in $\bar{\mathbb{Q}}$ are determined completely by sending α to each root of $f(x)$. □

Recall also that D_K denotes the absolute discriminant of the number field K, which is a non-zero integer.

Our theorems are then the following:

Theorem 4.1.3. *(Minkowski) Given any fractional ideal I of a ring of integers $\mathscr{O}_K$, there exists some ideal $J \subseteq \mathscr{O}_K$, and $x \in K$, with $I = xJ$, and*

$$N(J) \leqslant \frac{n!}{n^n}\left(\frac{4}{\pi}\right)^{r_2} |D_K|^{1/2}.$$

We will see how this implies:

Theorem 4.1.4. *The ideal class group of a ring of integers $\mathscr{O}_K$ is finite.*

However, the advantage of the first theorem is that it gives a sufficiently effective bound for the norm of ideals in each ideal class that in at least some examples, one can explicitly compute the ideal class group.

Because of Theorem 4.1.4, we can make the following fundamental definition:

Definition 4.1.5. Let K be a number field. Then the **class number** h_K of K is the order of the ideal class group.

We will use a similar set of techniques to prove Dirichlet's unit theorem. We will give a slightly more general form, but the main content is:

Theorem 4.1.6. *The group of units of $\mathscr{O}_K$ is of the form*

$$\mathbb{Z}/m\mathbb{Z} \times \mathbb{Z}^{r_1+r_2-1},$$

with the elements of finite order corresponding to the roots of unity lying in $\mathscr{O}_K$.

4.2. The structure of the finiteness proof

We give an overview of the proof of Theorem 4.1.3. The starting point is to consider the imbedding $K \to \mathbb{R}^n$ described as follows: let $\sigma_1, \dots, \sigma_{r_1}$ be the real imbeddings of K, and $\sigma_{r_1+1}, \dots, \sigma_{r_1+r_2}$ be representatives for each pair of complex imbeddings. Via the usual identification $\mathbb{C} \cong \mathbb{R}^2$, we get a map $\varphi : K \to \mathbb{R}^n$ by $x \mapsto (\sigma_1(x), \dots, \sigma_{r_1+r_2}(x))$.

The fundamental idea will be to show that given any non-zero ideal I of $\mathscr{O}_K$, we can find a non-zero element $x \in I$ with bounded norm:

Theorem 4.2.1. *Given a non-zero ideal I of a ring of integers $\mathscr{O}_K$, there exists some non-zero $x \in I$ such that:*

$$|N_{K/\mathbb{Q}}(x)| \leqslant \frac{n!}{n^n}\left(\frac{4}{\pi}\right)^{r_2} N(I)|D_K|^{1/2}.$$

This theorem can in turn be broken up into three steps. First, using discriminants, we show:

Theorem 4.2.2. *Given a non-zero ideal I of a ring of integers $\mathscr{O}_K$, $\varphi(I)$ is a lattice of full rank n in $\mathbb{R}^n$, and the volume of a fundamental parallelepiped for $\varphi(I)$ is given by*

$$2^{-r_2}N(I)|D_K|^{1/2}.$$

Next, we describe the region of $\mathbb{R}^n$ in which elements of $\mathscr{O}_K$ have bounded norm:

Theorem 4.2.3. *Given $t > 0$, there exists a compact region R_t of $\mathbb{R}^n$ which is convex and symmetric about the origin, such that for $x \in K$ with $\varphi(x) \in R_t$, we have $|N_{K/\mathbb{Q}}(x)| < \frac{t^n}{n^n}$, and with volume*

$$\mathrm{vol}(R_t) = 2^{r_1-r_2}\pi^{r_2}t^n/n!.$$

Finally, we have the following result, known as Minkowski's Theorem:

Theorem 4.2.4. *Let $R \subseteq \mathbb{R}^n$ be compact, convex and symmetric about the origin, let $L \subseteq \mathbb{R}^n$ be a lattice of full rank, with fundamental parallelepiped of volume V, and suppose that*

$$\mathrm{vol}\, R \geqslant 2^n V.$$

Then R contains a non-zero point of L.

Putting the last three theorems together for a non-zero ideal I, we find that $\varphi(I)$ is a lattice in $\mathbb{R}^n$, and if we set

$$t = (n!(\frac{4}{\pi})^{r_2} N(I)|D_K|^{1/2})^{1/n},$$

we find that R_t has the correct volume to apply the last theorem to produce a non-zero element of the norm required by Theorem 4.2.1. Thus, the last three theorems imply the first one.

We conclude by showing that Theorem 4.2.1 implies Theorem 4.1.3, and that Theorem 4.1.3 implies Theorem 4.1.4.

Proof of Theorem 4.1.3. Given a fractional ideal I, choose any non-zero $x \in I$; then xI^{-1} is contained in, hence an ideal of, $\mathscr{O}_K$. By Theorem 4.2.1, we have some non-zero $y \in xI^{-1}$ with

$$|N_{K/\mathbb{Q}}(y)|N(xI^{-1})^{-1} \leqslant \frac{n!}{n^n}\left(\frac{4}{\pi}\right)^{r_2}|D_K|^{1/2}.$$

If we let $J = \frac{y}{x}I$, we find by Corollary 1.6.9 and Exercise 4.1 that

$$N(J) = |N_{K/\mathbb{Q}}(y)|N(xI^{-1})^{-1} \leqslant \frac{n!}{n^n}\left(\frac{4}{\pi}\right)^{r_2}|D_K|^{1/2},$$

as desired. □

Proof of Theorem 4.1.4. It suffices to show that given any positive integer M, there are finitely many ideals of $\mathscr{O}_K$ having norm M. We observe that if I has $N(I) = M$, then $M \in I$: indeed, if $\mathscr{O}_K/I$ has cardinality M, then it must have characteristic dividing M, by basic group theory. But in fact there are only finitely many ideals of $\mathscr{O}_K$ containing M: we already know that there are finitely many prime ideals $\mathfrak{p}_1, \ldots, \mathfrak{p}_m$ containing M and we also know that ideals containing M satisfy the descending chain condition, so for each $\mathfrak{p}_i$ there is some e_i such that for $e > e_i$, $\mathfrak{p}^{e_i}$ does not contain M. Now, since every ideal can be factored as a product of prime ideals, and recalling that for ideals I, J, we have $IJ \subseteq I \cap J$, we find that the only ideals which could possibly contain M are the finitely many of the form $\prod_i \mathfrak{p}_i^{e'_i}$, with $e'_i \leqslant e_i$ for all i. This completes the proof that there are only finitely many ideals of a given norm, and we thus conclude Theorem 4.1.4 from Theorem 4.1.3. □

4.3. Ideals as lattices

Recall that we had the imbedding $K \to \mathbb{R}^n$ described as follows: let $\sigma_1, \ldots, \sigma_{r_1}$ be the real imbeddings of K, and $\sigma_{r_1+1}, \ldots, \sigma_{r_1+r_2}$ be representatives for each pair of complex imbeddings. Via the usual identification $\mathbb{C} \cong \mathbb{R}^2$, we get a map $\varphi : K \to \mathbb{R}^n$ by $x \mapsto (\sigma_1(x), \ldots, \sigma_{r_1+r_2}(x))$.

We define:

Definition 4.3.1. A **lattice** L of rank m in $\mathbb{R}^n$ is an additive subgroup of the form $L = \{a_1v_1 + \cdots + a_mv_m : a_i \in \mathbb{Z}\}$ for some linearly independent set of $v_i \in \mathbb{R}^n$. A lattice is said to be of **full rank** if $m = n$.

We begin by using discriminants to prove Theorem 4.2.2. In fact, we will see that it is not difficult to relate the volume of the fundamental parallelepiped for

$\varphi(I)$ to the discriminant of I, and the proof will follow. First, we observe we can make the following definition:

Definition 4.3.2. Given a non-zero ideal I of a ring of integers $\mathscr{O}_K$, we define the **absolute discriminant** $D(I)$ of I to be $D_{K/\mathbb{Q}}((x_i)_i)$, where $(x_i)_i$ is any $\mathbb{Z}$-basis of I.

As in the definition of D_K, we find that $D(I)$ is a well-defined non-zero integer. We have:

Lemma 4.3.3. *Let $\bar{x} = (x_1, \dots, x_n) \in K^n$, and let $\varphi_{\bar{x}}$ be the real matrix with rows given by $\varphi(x_i)$. Then*

$$D_{K/\mathbb{Q}}((x_i)_i) = (-4)^{r_2} (\det \varphi_{\bar{x}})^2.$$

In particular, if the x_i are linearly independent, the discriminant is positive if and only if r_2 is even.

Proof. Let $\sigma_{\bar{x}}$ be the complex matrix with rows

$$(\sigma_1(x_i), \dots, \sigma_{r_1}(x_i), \sigma_{r_1+1}(x_i), \sigma_{r_1+r_2+1}(x_i), \dots, \sigma_{r_1+r_2}(x_i), \sigma_{r_1+2r_2}(x_i)).$$

We first show that

$$\det(\varphi_{\bar{x}}) = (-2i)^{-r_2} \det \sigma_{\bar{x}}.$$

Recall that we ordered the σ_i so that $\sigma_{r_1+i} = \bar{\sigma}_{r_1+r_2+i}$, so the ordering of $\sigma_{\bar{x}}$ was chosen so that the (r_1+2i-1)st and (r_1+2i)th columns are complex conjugates, for $i > 0$. Starting from $\sigma_{\bar{x}}$, we can then obtain $\varphi_{\bar{x}}$ as follows: the first r_1 columns already agree; for each $i > 0$, if we add the (r_1+2i)th column to the (r_1+2i-1)st, and divide the (r_1+2i-1)st column by 2, then the (r_1+2i-1)st column consists of $\Re\sigma_{r_1+i}(x_i)$, which is the $(r1+2i-1)$st column of $\varphi_{\bar{x}}$. This divides the determinant by 2. Similarly, substracting the new (r_1+2i-1)st column from the (r_1+2i)th removes the real part, and dividing by $-i$ we obtain the (r_1+2i)th row of $\varphi_{\bar{x}}$, while dividing the determinant by $-i$. Doing this for each $0 < i \leqslant r_2$, we obtain $\varphi_{\bar{x}}$ after dividing the determinant by $(-2i)^{r_s}$, as desired.

But now we are done: by our earlier discriminant formula, we have that $D_{K/\mathbb{Q}}((x_i)_i) = (\det \sigma_{\bar{x}})^2$, giving the desired formula. □

The theorem now follows easily:

Proof of Theorem 4.2.2. We simply apply the lemma to the case that the x_i are a $\mathbb{Z}$-basis of the ideal I. We know that $D(I)$ is non-zero, because by Proposition 1.6.1 we know I contains a $\mathbb{Q}$-basis of K, so by Lemma 1.5.9 the discriminant is non-zero. It follows that $\det \varphi_{\bar{x}}$ is non-zero, so $\varphi(I)$ spans $\mathbb{R}^n$, and $\varphi(I)$ must be a full lattice as desired. Furthermore, by the standard result that determinants compute volume , $|\varphi_{\bar{x}}|$ is then the volume of the fundamental parallelepiped of $\varphi(I)$, and we obtain the desired statement, with the exception that we have $|D(I)|^{1/2}$ instead of $N(I)|D_K|^{1/2}$. But recall that $N(I)$ is by definition the index of I in $\mathscr{O}_K$, so the volume of a fundamental parallelepiped of $\varphi(I)$ is $N(I)$ times the volume for $\varphi(\mathscr{O}_K)$, giving the desired statement. □

Corollary 4.3.4. *We have $D(I) = N(I)^2 D_K$.*

Proof. We saw in the preceding proof that this equation holds up to sign, but by the lemma the signs are determined by K, independently of I. □

4.4. A region with bounded norm

We now move on to the region R_t:

PROOF OF THEOREM 4.2.3. Explicitly, we set

$$R_t = \{(x_1, \ldots, x_n) \in \mathbb{R}^n : |x_1| + \cdots + |x_{r_1}| + 2\sqrt{x_{r_1+1}^2 + x_{r_1+2}^2} + \cdots + 2\sqrt{x_{r_1+2r_2-1}^2 + x_{r_1+2r_2}^2} \leqslant t\}.$$

This is visibly compact and symmetric about the origin; it is convex because each term in the sum is linear under scaling, and satisfies the triangle inequality.

We now check the assertion that if $\varphi(x) \in R_t$, then $|N_{K/\mathbb{Q}}(x)| \leqslant \frac{t^n}{n^n}$. By our earlier formulas, $|N_{K/\mathbb{Q}}(x)| = \prod_{i=1}^n |\sigma_i(x)|$, where $|z|$ is the absolute value on $\mathbb{C}$ given by $|a + bi| = \sqrt{a^2 + b^2}$. But by the arithmetic-geometric mean inequality for non-negative real numbers, we have

$$(\prod_{i=1}^n |\sigma_i(x)|)^{1/n} \leqslant \frac{\sum_{i=1}^n |\sigma_i(x)|}{n},$$

and taking nth powers we obtain the desired inequality.

Finally, it remains to compute the volume of R_t. This may be carried out by multivariable integration in polar coordinates; see [**7**, V, §3, Lem. 3, p. 117]. □

4.5. Minkowski's theorem

The third result we needed to prove was Theorem 4.2.4, known as Minkowski's theorem.

We will need to know the following two lemmas, in which we ignore questions of existence of volume:

LEMMA 4.5.1. *Let R be a bounded region, and L a lattice of full rank, with fundamental parallelepiped of volume V. If all translates of R under L are disjoint, we have* $\operatorname{vol} R \leqslant V$.

PROOF. Let S be a (half-open) fundamental parallelepiped for L; by definition, the translates of S under L are disjoint, and cover $\mathbb{R}^n$; it follows that $R = \cup_{v \in L}(R \cap (S + v)) = \cup_{v \in L}((R - v) \cap S) + v$. Since the translates of R are disjoint, in particular we have that $\cup_{v \in L}((R - v) \cap S)$ are disjoint, and contained in S, so

$$\operatorname{vol} R = \operatorname{vol} \cup_{v \in L}((R - v) \cap S) \leqslant \operatorname{vol} S =: V.$$

□

LEMMA 4.5.2. *For any lattice $L \subseteq \mathbb{R}^n$, there are only finitely many points of L in any bounded region of $\mathbb{R}^n$.*

PROOF. Let $v_1, \ldots, v_m$ be a basis for L, and extend to a basis $v_1, \ldots, v_n$ of $\mathbb{R}^n$. Denote by C_t the real cube of vectors $v = (c_1, \ldots, c_n) \in \mathbb{R}^n$ such that $|c_i| < t$ for all i. Also denote by $C_{v,t}$ the region of vectors of the form $v = \sum_i c_i' v_i$ with $c_i' < t'$ for all i. We claim that given $t > 0$, there exists a $t' > 0$ such that $C_{v,t'} \subseteq C_t$. Indeed, let M be the inverse of the matrix having v_i as its ith column. Then $(c_i')_i = Mv$, so if m denotes the maximal coordinate of M, it is easily checked that $t' = nmt$ suffices. Now, any bounded region is contained in C_t for some t, hence in $C_{v,t'}$ for some t'. Since elements of L must be of the form $\sum_i c_i' v_i$ for $c_i' \in \mathbb{Z}$, we have $L \cap C_{v,t'}$ is finite, as desired. □

Proof of Theorem 4.2.4. We first prove the statement in the case that $\operatorname{vol} R > 2^n V$, where the compactness hypothesis will not be necessary. By the first lemma, we have that some two translates of $\frac{1}{2}R$ have a point in common: i.e., $\exists x, v_1, v_2 \in \mathbb{R}^n$ with $v_1, v_2 \in L$ distinct, and $x \in (\frac{1}{2}R + v_1) \cap (\frac{1}{2}R + v_2)$, or equivalently, $x - v_1, x - v_2 \in \frac{1}{2}R$. By symmetry about the origin, we have $v_2 - x \in \frac{1}{2}R$, and then by convexity, we have $\frac{x-v_1+v_2-x}{2} = \frac{v_2-v_1}{2} \in \frac{1}{2}R$, so $v_2 - v_1 \in R$ is the desired non-zero lattice point.

We can now conclude the statement in the compact case: for all $\epsilon > 0$, we have that $(1+\epsilon)R$ contains a non-zero point of L, and only finitely many such points by the second lemma. As ϵ goes to 0, one non-zero point P of L must be in $(1+\epsilon)R$ for all ϵ, so that $\frac{P}{1+\epsilon} \in R$ for all ϵ, and it follows by compactness of R that the point lies in R. □

This completes the proof of Theorem 4.1.3, and hence of Theorem 4.1.4.

4.6. Ideals of bounded norm

Since every ideal in a ring of integers factors into prime ideals, and since the norm is multiplicative by Exercise 4.1, the main step of finding all ideals of norm below a given bound is to find the prime ideals of norm below the same bound.

We have the following lemma:

Lemma 4.6.1. *If $\mathfrak{p}$ is a non-zero prime ideal in a ring of integers $\mathscr{O}_K$, then $N(\mathfrak{p}) = p^n$ for some prime $p \in \mathbb{Z}$, and $n \in \mathbb{N}$. If $(p) = \prod_i \mathfrak{p}_i^{e_i}$ in $\mathscr{O}_K$, then the $\mathfrak{p}_i$ are precisely the prime ideals with norm equal to a power of p.*

Proof. We know that any prime ideal $\mathfrak{p}$ of $\mathscr{O}_K$ lies above a unique (p) of $\mathbb{Z}$, so the second statement will imply the first. We now show the second statement: suppose that $N(\mathfrak{p}) = p^n$ for some n; then we saw that $p^n \in \mathfrak{p}$, so $\mathfrak{p}|(p^n)$, and since $\mathfrak{p}$ is prime $\mathfrak{p}|(p)$. Conversely, taking norms of both sides and using multiplicativity, it is clear that if $(p) = \prod_i \mathfrak{p}_i^{e_i}$, we have $N(\mathfrak{p}_i) = p^{n_i}$ for some $n_i \in \mathbb{N}$. □

We can summarize as follows:

Proposition 4.6.2. *For a natural number N, let $p_1, \dots, p_m$ be the primes less than or equal to N. Given a ring of integers $\mathscr{O}_K$, for each i write $(p_i) = \prod_{j=1}^{m_j} \mathfrak{p}_{i,j}$ in $\mathscr{O}_K$. Then the ideals of $\mathscr{O}_K$ with norm at most N are precisely those of the form $\prod_{i,j} \mathfrak{p}_{i,j}^{e_{i,j}}$, with $\prod_{i,j} N(\mathfrak{p}_{i,j})^{e_{i,j}} \leqslant N$.*

We make one further proposition:

Proposition 4.6.3. *Let $K = \mathbb{Q}(\alpha)$, where α is a root of an irreducible monic polynomial $f(x) \in \mathbb{Z}[x]$. Let $\mathfrak{p}$ be an ideal of $\mathscr{O}_K$, and suppose that $N(\mathfrak{p}) = p$ is prime. Then $f(x)$ has a root mod p.*

Proof. If $N(\mathfrak{p}) = p$, then the natural ring homomorphism $\mathbb{Z}/p \to \mathscr{O}_K/\mathfrak{p}$ is an isomorphism. Then the preimage of α must be a root of $f(x)$ in $\mathbb{Z}/p$. □

We now give an example of an ideal class group computation for a ring of integers in a quintic number field.

Example 4.6.4. We again consider the number field $\mathbb{Q}(\alpha)$, where α is a root of the polynomial $f(x) = x^5 - x + 1$. We saw in Example 1.7.3 that $K = \mathbb{Q}(\alpha)$ has degree 5 over $\mathbb{Q}$, and the ring of integers of $\mathbb{Q}(\alpha)$ is $\mathbb{Z}[\alpha]$.

We apply Minkowski's bound to compute the class group. We claim that $f(x)$ has a single real root: indeed, the derivative is $5x^4 - 1$, so $f(x)$ has two real critical points, one between -1 and 0, and one between 0 and 1. It is easy to see that $f(x)$ is strictly positive between -1 and 1 (and hence for $x > 1$ as well), so we conclude that its only real root occurs for $x < -1$. Thus, by the lemma, we have $r_2 = 2$, and we compute $\frac{n!}{n^n}\left(\frac{4}{\pi}\right)^{r_2}|D_K|^{1/2} < 4$. Thus, the only ideals we need to consider are ideals having norm 2 or 3. By Proposition 4.6.3, such ideals can only occur if $f(x)$ has a root mod 2 or 3, and it is easy to check that it doesn't. Hence, the only ideal class is the class of the ideal having norm 1, which is the trivial ideal.

4.7. The canonical mapping on units

We will prove the unit theorem for arbitrary orders inside rings of integers. Dirichlet's unit theorem states the following:

THEOREM 4.7.1. *Let R be an order in $\mathscr{O}_K$. The group of units of R is of the form*

$$\mathbb{Z}/m\mathbb{Z} \times \mathbb{Z}^{r_1+r_2-1},$$

with the elements of finite order corresponding to the roots of unity lying in R.

The main idea is a canonical mapping much like the imbedding used in the proof of the finiteness of the ideal class group. As in that case, we order the imbeddings σ_i of K into $\bar{\mathbb{Q}}$ so that $\sigma_1, \dots, \sigma_{r_1}$ are real, and $\sigma_{r_1+1}, \dots, \sigma_{r_1+r_2}$ contain one representive of each conjugate pair of complex (i.e., non-real) imbeddings. The map we consider is then the following:

$$\psi : K^* \to \mathbb{R}^{r_1+r_2}$$

given by

$$x \to (\log|\sigma_1(x)|, \dots, \log|\sigma_{r_1}(x)|, 2\log|\sigma_{r_1+1}(x)|, \dots, 2\log|\sigma_{r_1+r_2}(x)|).$$

Here we again define $|a + bi| := \sqrt{a^2 + b^2}$.

Since $|\cdot|$ is multiplicative, it follows that this map is a homomorphism from K^* under multiplication to $\mathbb{R}^{r_1+r_2}$ under addition.

We reduce the unit theorem to the following:

THEOREM 4.7.2. *The image $\psi(R^*) \subseteq \mathbb{R}^{r_1+r_2}$ is a lattice of rank $r_1 + r_2 - 1$, spanning the hyperplane H defined by $x_1 + \cdots + x_{r_1+r_2} = 0$.*

LEMMA 4.7.3. *Theorem 4.7.2 implies Theorem 4.7.1.*

PROOF. Theorem 4.7.2 implies that $R^*/\ker\psi \equiv \mathbb{Z}^{r_1+r_2-1}$; we check that the kernel of ψ is exactly the roots of unity of R^*. Certainly, the roots of unity are in the kernel, since $\mathbb{R}^{r_1+r_2}$ as an additive group has no non-zero elements of finite order. Conversely, if $x \in K^*$ has $\psi(x) = 0$, this means that $|\sigma_i(x)| = 1$ for all i. Note that this also implies the same for x^j for all j. In particular $\varphi(x^j)$ lies in a bounded region of $\mathbb{R}^n$ for all j. Since $x \in \mathscr{O}_K$, and we had that $\varphi(\mathscr{O}_K)$ was a lattice in $\mathbb{R}^n$, it follows that $\varphi(x^j)$ is a finite set, so x^j is a finite set, and x is a root of unity.

The statement then follows from Corollary 3.8.7 and the following proposition, which implies that the roots of unity in K form a cyclic group. □

PROPOSITION 4.7.4. *Let G be a finite multiplicative subgroup of a field. Then G is cyclic.*

Proof. Let m be the maximal order of the elements of G. We will show that $m = |G|$. Of course, it is clear that $m \leqslant |G|$. We next claim that the order of every element of G divides m: it suffices to show that given elements x_1, x_2 of orders m_1, m_2, there exists an element of order $\operatorname{lcm}(m_1, m_2)$. Indeed, finding m_1', m_2' which are relatively prime, divide m_1 and m_2 respectively, and have $m_1', m_2' = \operatorname{lcm}(m_1, m_2)$, it is easy to check that $x_1^{m_1/m_1'} x_2^{m_2/m_2'}$ has the desired order. But elements of order dividing m are roots of $x^m - 1$, so there can be at most m of them, and we have $|G| \leqslant m$, so $|G| = m$, as desired. $\square$

Remark 4.7.5. Note that the full $\ker \psi$ on K^* is larger: for instance, in $\mathbb{Q}[i]$ we have $\psi(\frac{3+4i}{5}) = 0$.

We first prove the easier portion of Theorem 4.7.2:

Proposition 4.7.6. *The image $\psi(R^*) \subseteq \mathbb{R}^{r_1+r_2}$ is a lattice, and contained in the hyperplane defined by $x_1 + \cdots + x_{r_1+r_2} = 0$.*

We use the following lemma:

Lemma 4.7.7. *Let $L \subseteq \mathbb{R}^n$ be an additive subgroup such that for any bounded region $R \subseteq \mathbb{R}^n$, we have that $R \cap L$ is a finite set. Then L is a lattice.*

Proof. Let $v_1, \ldots, v_m$ be a maximal linearly independent subset of L, and let L^0 be the lattice spanned by the v_i. We first show that L/L^0 is finite. In fact, if we fix a fundamental parallelepiped S for L^0, we may write any element of L as $v + w$, with $v \in L$, and $w \in S$. But then $w = (v + w) - v \in L$, and since S is bounded, we find there are finitely many possibilities for S.

Hence, L/L^0 is a finite group of some order d. It follows that $\frac{1}{d}L^0$ contains L. But it is a standard theorem that an abelian group which both contains and is contained in groups isomorphic to $\mathbb{Z}^m$ is itself isomorphic to $\mathbb{Z}^m$, so $L = \mathbb{Z} < v_1', \ldots, v_m' >$. for some v_i'. It remains only to note that $\operatorname{span} L^0 = \operatorname{span} \frac{1}{d}L^0 = \operatorname{span}(v_1, \ldots, v_m)$, and $\operatorname{span} L^0 \subseteq \operatorname{span} L \subseteq \operatorname{span} \frac{1}{d}L^0$, so the v_i' must span an m-dimensional space, and are therefore linearly independent. $\square$

Proof of proposition. We first claim that $\psi(R^*)$ is contained in the hyperplane defined by $x_1 + \cdots + x_n = 0$. But since $R \subseteq \mathscr{O}_K$, we have $R^* \subseteq \mathscr{O}_K^*$, so if $x \in R^*$, we have $|N_{K/\mathbb{Q}}(x)| = 1$. But $|N_{K/\mathbb{Q}}(x)| = \prod_{i=1}^{r_1+2r_2} |\sigma_i(x)| = \prod_{i=1}^{r_1} |\sigma_i(x)| \prod_{i=r_1+1}^{r_2} |\sigma_i(x)|^2$, since $|y| = |\bar{y}|$, and taking log of both sides gives the desired statement.

Next, we need to show that $\psi(R^*)$ is a lattice, or equivalently, that it contains only finitely many points in any bounded region. But note that if we restrict each coordinate $\log |\sigma_i(x)|$ or $2 \log |\sigma_i(x)|$, depending on whether or not $i \leqslant r_1$, to be bounded by some t, this is equivalent to requiring that $|\sigma_i(x)| < e^t$ or $|\sigma_i(x)|^2 < e^t$, so it follows that the images $\varphi(x)$ are bounded in $\mathbb{R}^n$. Since $R^* \subseteq \mathscr{O}_K$, there are only finitely many $x \in R^*$ with $\varphi(x)$ in the given region in $\mathbb{R}^n$, and hance only finitely many with $\psi(x)$ in the given region in $\mathbb{R}^{r_1+r_2}$, proving that $\psi(R^*)$ is a lattice, as desired. $\square$

The theorem also allows us to give a definition which plays a role for units analogous to the role of discriminant for the full ring of integers.

Definition 4.7.8. The **regulator** of R is the volume of a fundamental parallelepiped for $\psi(R^*)$ in H.

The regulator is not necessary for the proof of the unit theorem, but will play an important role in the analytic class number formula and its applications.

4.8. Properties of orders

To prove the harder portion of Theorem 4.7.2, we have to briefly examine the basic properties of orders.

Lemma 4.8.1. *We have:*

(i) *$\varphi(R)$ is a lattice of full rank in $\mathbb{R}^n$; in particular, the additive group of R is a free abelian group of rank n.*
(ii) *R/I is finite for any non-zero ideal I.*
(iii) *R is Noetherian, and has the property that every non-zero prime ideal is maximal.*
(iv) *For any non-zero $x \in R$, there exist only finitely many ideals of R containing x.*

Proof. For (i), we have that $\varphi(R) \subseteq \varphi(\mathscr{O}_K)$ and is still an additive subgroup, so it is certainly a lattice. But since R contains a $\mathbb{Q}$-basis of K, it follows by the same discriminant argument used in the case of $\mathscr{O}_K$ that it must span all of $\mathbb{R}^n$, and must therefore be of full rank.

For (ii), by (i) and as in the case of $\mathscr{O}_K$, it suffices to see that I contains a non-zero integer. But this still follows from Lemma 1.6.5, since it was stated for arbitrary integral domains.

For (iii), the proof follows from (ii) just as in the case of the full ring of integers.

For (iv), it suffices to show that $R/(x)$ has only finitely many ideals. and this follows from the fact that $R/(x)$ is finite. □

Remark 4.8.2. R is not in general a Dedekind domain, as it will not satisfy the condition of being integrally closed. In fact, $\mathscr{O}_K$ is the integral closure of R in K.

4.9. Fullness

Our aim now is to prove the harder half of Theorem 4.7.2:

Proposition 4.9.1. *The image $\psi(R^*)$ spans the hyperplane defined by $x_1 + \cdots + x_{r_1+r_2} = 0$.*

To do so, we use the following criterion:

Lemma 4.9.2. *Let $L \subseteq \mathbb{R}^m$ be a lattice. Then L is of full rank if and only if there exists a bounded region $S \subseteq \mathbb{R}^n$ such that the translates $S + L$ cover $\mathbb{R}^m$.*

Proof. Certainly, if L is full, then the translates of a fundamental parallelepiped of L cover $\mathbb{R}^n$. Conversely, let S be a bounded region of $\mathbb{R}^n$; say that the maximal distance of a point in S from $(0, \dots, 0)$ is t. If L is not full, then its span is a proper subspace of $\mathbb{R}^m$, and there exist vectors of distance $> t$ from any point in this subspace. It follows that such vectors cannot be in any translate of S by L. □

Our argument will use the following structures:

Definition 4.9.3. We have the multiplication map $\cdot : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}^n$ given by considering $\mathbb{R}^n$ as $\mathbb{R}^{r_1} \times \mathbb{C}^{r_2}$, and multiplying coordinates.

Definition 4.9.4. We define the norm map $N : \mathbb{R}^n \to \mathbb{R}_{\geqslant 0}$ by the formula

$$N(x_1, \dots, x_n) = \prod_{i=1}^{r_1} |x_i| \prod_{i=1}^{r_2} (x_{r_1+2i-1}^2 + x_{r_1+2i}^2).$$

Let $\chi : \mathbb{R}^n \to \mathbb{R}^{r_1+r_2}$ be the map such that $\psi = \chi \circ \varphi$. Explicitly,

$$\chi(x_1, \dots, x_n) = (\log |x_1|, \dots, \log |x_{r_1}|, \log(x_{r_1+1}^2 + x_{r_1+2}^2), \dots, \log(x_{n-1}^2 + x_n^2)).$$

We then have the following compatibilities with existing ideas:

(i) $\varphi(xy) = \varphi(x) \cdot \varphi(y)$ for $x, y \in K$.
(ii) $\chi(\bar{x} \cdot \bar{y}) = \chi(\bar{x}) + \chi(\bar{y})$ for $\bar{x}, \bar{y} \in \mathbb{R}^n$.
(iii) $N(\bar{x} \cdot \bar{y}) = N(\bar{x}) N(\bar{y})$ for $\bar{x}, \bar{y} \in \mathbb{R}^n$.
(iv) $\log N(\bar{x})$ is given by the sum of the coordinates of $\chi(\bar{x})$, for $\bar{x} \in \mathbb{R}^n$.

We also want the following description of a region of $\mathbb{R}^n$:

Lemma 4.9.5. *Given positive $\bar{c} = (c_1, \dots, c_{r_1+r_2})$, the region $S_{\bar{c}}$ of $\mathbb{R}^n$ consisting of $(x_1, \dots, x_n)$ with $|x_i| \leqslant c_i$ for $i \leqslant r_1$ and $x_{r_1+2i-1}^2 + x_{r_1+2i}^2 \leqslant c_{r_i+i}$ for $i \leqslant r_2$ is compact, convex, symmetric about the origin, and has volume $2^{r_1} \pi^{r_2} \prod_i c_i$.*

Proof. The region is clearly compact and symmetric about the origin, and convexity is easily checked from the definition. For the volume, note that the region is simply a product of r_1 intervals of length $2c_i$ for $i \leqslant r_1$, and r_2 discs in the plane of radius $\sqrt{c_i}$ for $i > r_1$. The volume is therefore $\prod_{i=1}^{r_1}(2c_i) \prod_{i=r_1+1}^{r_1+r_2} (\pi(\sqrt{c_i})^2) = 2^{r_1} \pi^{r_2} \prod_i c_i$. □

We are now ready to complete our proof.

Proof of Proposition 4.9.1. We wish to show that there exists a bounded region S of the hyperplane H defined by $\sum_i x_i = 0$ in $\mathbb{R}^{r_1+r_2}$ such that the translates of S under $\psi(R^*)$ cover H. We begin by translating this into an equivalent statement in terms of $\varphi(R^*)$ and $\mathbb{R}^n$.

Recalling the norm we had defined last time on $\mathbb{R}^n$, we have that $x \in R$ is a unit if and only if $N(\varphi(x)) = 1$. Now, let $U \subseteq \mathbb{R}^n$ be the subset of elements with $N(\bar{x}) = 1$; then $\varphi(R^*) \subseteq U$, and $\chi(U)$ is precisely H inside $\mathbb{R}^{r_1+r_2}$. We claim that if $\tilde{S}$ is a bounded region of U, then $\chi(\tilde{S})$ is a bounded region of H: if we choose C so that for all $(x_1, \dots, x_n) \in \tilde{S}$, we have $|x_i| < C$ for $i \leqslant r_1$, and $x_{r_1+2i-1}^2 + x_{r_1+2i}^2 < C$ for $i \leqslant r_2$, then we see that the coordinates of $\chi(U)$ are all bounded above by $\log C$, and since $\chi(U)$ lies on H, the coordinates must also be bounded below, explicitly by $-(r_1 + r_2 - 1) \log C$.

Since χ sends the multiplication map on $\mathbb{R}^n$ to translation on $\mathbb{R}^{r_1+r_2}$, we find that in order to find a bounded $S \subseteq H$ such that the translates of S under $\psi(R^*)$ cover H, it is sufficient to find a bounded $\tilde{S}$ in U such that for any $v \in U$, there exists $x \in R^*$ with $\varphi(x) \cdot v \in \tilde{S}$. Setting $S = \chi(\tilde{S})$ then has the desired properties.

We now describe $\tilde{S}$. Let V denote the volume of a fundamental parallelepiped for $\varphi(R)$, and choose positive $\bar{c} = (c_1, \dots, c_{r_1+r_2})$ such that $C = \prod_i c_i > \left(\frac{4}{\pi}\right)^{r_2} V$. Let $S_{\bar{c}}$ be the region of the lemma, so that we have $\operatorname{vol} S_{\bar{c}} = 2^{r_1} \pi^{r_2} C > 2^n V$, and if a point is in $S_{\bar{c}}$, its norm is at most C. Now, there are only finitely many ideals of R of norm at most C, and in particular finitely many principal ideals, so denote these by $(\alpha_1), \dots, (\alpha_m)$ for some m. We claim that the region

$$\tilde{S} = U \cap \left(\cup_{i=1}^m \varphi(\alpha_i^{-1}) \cdot S_{\bar{c}} \right)$$

has the desired property: that is, it is bounded, and for any $y \in U$, there exists an $\alpha \in R^*$ such that $\varphi(\alpha) \cdot y \in \tilde{S}$.

It is easy to check that since $S_{\bar{c}}$ is bounded, each $\varphi(\alpha_i) \cdot S_{\bar{c}}$ is also bounded, so we have that $\tilde{S}$ is bounded. Next, let y be any element of U; recall that this meant that $N(y) = 1$. Now, since the multiplication by $\varphi(y)$ is linear on $\mathbb{R}^n$, given by a block-diagonal matrix, one easily checks that its determinant is $N(y) = 1$. Thus, the lattice $y \cdot (\varphi(R))$ has the same volume V as $\varphi(R)$, and it follows by Minkowski's theorem that $y \cdot \varphi(R)$ has a non-zero point in $S_{\bar{c}}$, say $\varphi(\alpha')$. Recalling that the norm we have defined on $\mathbb{R}^n$ is multiplicative with respect to our multiplication map, since $N(y) = 1$ we have $|N_{K/\mathbb{Q}}(\alpha')| = N(\varphi(\alpha')) \leqslant C$, so it follows that $(\alpha') = (\alpha_i)$ for some i, and $\alpha' = \alpha\alpha_i$ for some unit $\alpha \in R^*$. Unwinding the definitions, we have $y \cdot \varphi(\alpha') \in S_{\bar{c}}$, so $y \in \varphi(\alpha'^{-1}) \cdot S_{\bar{c}} = \varphi(\alpha^{-1}) \cdot \varphi(\alpha_i^{-1}) \cdot S_{\bar{c}} \subseteq \varphi(\alpha^{-1}) \cdot \tilde{S}$, as desired. □

This completes the proof of Theorem 4.7.2, and hence of Theorem 4.7.1, Dirichlet's unit theorem.

4.10. Quadratic fields

As an immediate corollary of the unit theorem, we can prove the theorem on Pell's equation which we initially mentioned in our motivation for studying number fields.

THEOREM 4.10.1. *Given $n \in \mathbb{N}$ not a perfect square, the solutions to the equation $x^2 - ny^2 = 1$ naturally form an abelian group, isomorphic to $\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}$.*

PROOF. We know that $a + b\sqrt{n}$ gives a solution to $a^2 - nb^2 = \pm 1$ if and only if it is a unit in $\mathbb{Z}[\sqrt{n}]$. We have $r_1 = 2, r_2 = 0$, and the only roots of unity in the ring (since it is contained in $\mathbb{R}$) are ± 1, so by the unit theorem, this unit group is isomorphic to $\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}$. Let α be a generator for the units modulo ± 1. Then $N_{\mathbb{Q}(\sqrt{n})/\mathbb{Q}}(\alpha) = \pm 1$; if it is 1, then all units must have norm 1, and we are done. If it is -1, then α^2 has norm 1, and we see that the group of solutions to Pell's equation is generated by $-1, \alpha^2$. □

We conclude with some brief remarks on what is known about class numbers of quadratic fields. It turns out that the imaginary case is quite different from the real case; this is closely related to the fact that in the real case, the regulator is non-trivial (by convention, the regular in the case that $r_1 + r_2 - 1 = 0$ is defined to be 1). We will understand this more clearly after studying the analytic class number formula.

Gauss conjectured (in terms of binary quadratic forms) that the class numbers of imaginary quadratic fields go to infinity as the discriminant goes to infinity, and this was proved in the 1930's by Hecke, Deuring, Mordell, and Heilbronn (see [**5**, §20.6] for this remarkable story).

The situation for real quadratic fields is quite different. Here, numerically it appears that about 80% of real quadratic fields have class number 1, but it has not been shown that infinitely many do. In fact, it is not even known that infinitely many number fields have class number 1. In the case of real quadratic fields, the main difficulty lies in predicting the size of the regulator, which is to say, the size of the smallest solutions to Pell's equation.

4.11. Exercises

EXERCISE 4.1. Given non-zero ideals I, J of a ring of integers $\mathscr{O}_K$, show that $N(IJ) = N(I)N(J)$.

Hint: reduce to the case that $I = \mathfrak{p}$ is a non-zero prime, then write $\#\mathscr{O}_K/(\mathfrak{p}J) = \#\mathscr{O}_K/J \cdot \#J/(\mathfrak{p}J)$, and consider $J/\mathfrak{p}J$ as a $\mathscr{O}_K/\mathfrak{p}$-vector space.

EXERCISE 4.2. Using Minkowski's bounds on the norm of representatives of each ideal class, compute the ideal class group of $\mathbb{Z}[\sqrt{-5}]$. Follow the method from Example 4.6.4.

The following isn't relevant to algebraic number theory, but gives an addition application of Minkowski's theorem.

EXERCISE 4.3. Use Minkowski's theorem on convex regions and lattices to give a different proof that if -1 is a square mod a prime p, then $p = x^2 + y^2$ for some $x, y \in \mathbb{Z}$.

Hint: consider the lattice generated by $(p, 0)$ and $(x, 1)$ in $\mathbb{R}^2$, where $x^2 \equiv -1 \pmod{p}$.

CHAPTER 5

Cyclotomic fields and Fermat's Last Theorem

We have now developed nearly enough theory to begin a more detailed study of cyclotomic fields, and prove the first portion of Kummer's results on Fermat's Last Theorem. Our first order of business will be to develop a few more general results relevant to computing rings of integers. We then apply these results to compute the rings of integers of cyclotomic fields, and we can finally use the structure of the ring of integers to approach Fermat's Last Theorem.

5.1. Ramification and the local case

We extend the results of Chapter 3 by examining in more detail the case that a single prime $\mathfrak{q}$ lies over a given prime $\mathfrak{p}$.

Theorem 5.1.1. *In Situation 3.1.1, let $\mathfrak{p} \subseteq R$ be a prime ideal, and suppose there is a unique $\mathfrak{q} \subseteq S$ lying above $\mathfrak{p}$, and $S/\mathfrak{q}$ is separable over $R/\mathfrak{p}$. Then:*

(i) *$S_\mathfrak{p} = S_\mathfrak{q}$ is a PID.*

(ii) *There exist $u \in S$ which generates $S/\mathfrak{q}$ as a field over $R/\mathfrak{p}$, and $q \in S$ which generates $\mathfrak{q}S_\mathfrak{p}$, and for any such u, q we have $S_\mathfrak{p} = R_\mathfrak{p}[u, q]$.*

Proof. (i) Since $\mathfrak{p} = \mathfrak{q} \cap R$, it is clear that $S_\mathfrak{p} \subseteq S_\mathfrak{q}$. Conversely, given $\frac{x}{s} \in S_\mathfrak{q}$, with $s \in S \smallsetminus \mathfrak{q}$, we claim that $(s) \cap R$ is not contained in $\mathfrak{p}$: indeed, write $(s) = \prod_i \mathfrak{q}_i^{e_i}$, with all the q_i different from $\mathfrak{q}$. Thus $\mathfrak{p}_i := \mathfrak{q}_i \cap R \neq \mathfrak{p}$ for any i, so we have $\mathfrak{p} \not\supset \prod_i \mathfrak{p}_i e^i \subseteq (s) \cap R$, as desired. Thus, $S_\mathfrak{p} = S_\mathfrak{q}$, and we already knew that the latter is a PID.

(ii) The asserted u exists by the primitive element theorem, since $S/\mathfrak{q}$ is assume to be separable over $R/\mathfrak{p}$. Next, q exists because $S_\mathfrak{p}$ is a PID by (i). We claim that $\mathfrak{p}\S_\mathfrak{p} + R_\mathfrak{p}[q, u] = S_\mathfrak{p}$, or equivalently, the elements q, u generate $S_\mathfrak{p}/\mathfrak{p}S_\mathfrak{p}$ over $R_\mathfrak{p}$. Since $\mathfrak{q}$ is the unique prime ideal lying above $\mathfrak{p}$, we have $\mathfrak{p} = \mathfrak{q}^e = (q)^e$ for some e in $S_\mathfrak{p}$. Given $x \in S_\mathfrak{p}$, suppose that $x \in (q)^m$ but not $(q)^{m+1}$; we prove the desired statement by downward induction on m, noting that if $m \geqslant e$, we have $x \in \mathfrak{p}S_\mathfrak{p}$. Now since $S_\mathfrak{p}$ is a PID with unique maximal ideal (q), $x = q^m u'$ for some unit u'. Since u generates $S/\mathfrak{q}$ over $R/\mathfrak{p}$, we have $u' = p(u) + q'$ for some $p(x) \in R[x]$, and $q' \in \mathfrak{q}$. We have thus written $x = q^m p(u) + q^m q'$, so $q^m q' \in \mathfrak{q}^{m+1}$, and by the induction hypothesis, is in $\mathfrak{p}S_\mathfrak{p} + S_\mathfrak{p}[q, u]$. We have thus proved the claim.

Now, the claim may be rephrased as $\mathfrak{p}(S_{L,\mathfrak{p}}/R_\mathfrak{p}[q, u]) = S_\mathfrak{p}/R_\mathfrak{p}[q, u]$. Since we know that $S_\mathfrak{p}$ is free of rank $[L : K]$ over $R_\mathfrak{p}$, it is in particular finitely generated, so the following version of Nakayama's lemma completes the proof of the theorem. □

Commutative Algebra Fact 5.1.2. (Nakayama's Lemma) Let I be an ideal contained in all maximal ideals of a ring R, and M a finite-generated R-module. If $IM = M$, then $M = 0$.

See [**7**, p. 8] for the proof.

5.2. The Minkowski bound and discriminants

Returning to the bound obtained by Minkowski in studying the ideal class group, we find we also obtain a lower bound on the discriminant of a number field, which will be a powerful tool for us:

Theorem 5.2.1. *For any number field $K \neq \mathbb{Q}$, we have $|D_K| > 1$.*

Proof. By the bound of Theorem 4.1.3, we have in particular that $1 \leqslant \frac{n!}{n^n}\left(\frac{4}{\pi}\right)^{r_2}|D_K|^{1/2}$, so

$$|D_K| \geqslant \frac{n^n}{n!}\left(\frac{\pi}{4}\right)^{r_2} \geqslant \frac{n^n}{n!}\left(\frac{\pi}{4}\right)^{n/2}.$$

Denoting the right hand side by a_n, we have

$$\frac{a_{n+1}}{a_n} = \left(\frac{\pi}{4}\right)^{1/2}\left(1+\frac{1}{n}\right)^n,$$

which is > 1 for all $n > 1$. Noting that $a_2 > 1$, we get the desired result. □

As an immediate consequence of this theorem and Theorem 3.3.5, we find:

Corollary 5.2.2. *For any non-trivial extension K of $\mathbb{Q}$, some prime p is ramified in $\mathscr{O}_K$.*

Thus, we can prove two number fields to be disjoint over $\mathbb{Q}$ by studying their discriminants:

Corollary 5.2.3. *Let K, L be number fields, such that $(D_K, D_L) = 1$. Then $K \cap L = \mathbb{Q}$, so in particular if L is Galois over $\mathbb{Q}$, we have $[KL : \mathbb{Q}] = [K : \mathbb{Q}][L : \mathbb{Q}]$.*

Proof. By Theorem 3.3.5, there is no prime p which ramifies in both $\mathscr{O}_K$ and $\mathscr{O}_L$. By Proposition 3.1.12, this means that no p can ramify in $\mathscr{O}_{K\cap L}$, so by the previous corollary, $K \cap L = \mathbb{Q}$.

The second assertion then follows by standard field theory: if $x \in L$, with minimal polynomial $f(t)$ over $\mathbb{Q}$, then because L is Galois over $\mathbb{Q}$, all roots of $f(t)$ lie in L, and in particular if $f(t)$ factors non-trivially in K, there is some coefficient α appearing in the factorization which is in $K \smallsetminus \mathbb{Q}$, and this coefficient must also lie in L, so we have $K \cap L \neq \mathbb{Q}$. If we choose x to be a primitive element for L over $\mathbb{Q}$, we find its degree over K is the same as its degree over $\mathbb{Q}$, giving the desired assertion. □

We can use the corollary to obtain a statement entirely in terms of minimal polynomials over $\mathbb{Q}$.

Theorem 5.2.4. *Choose $\alpha, \beta \in \bar{\mathbb{Q}}$, and suppose that α and β satisfy monic integral polynomials $f(x), g(x) \in \mathbb{Z}[x]$, and that there is no prime p for which both $f(x)$ and $g(x)$ have repeated roots mod p. Then $\mathbb{Q}(\alpha) \cap \mathbb{Q}(\beta) = \mathbb{Q}$, and if $\mathbb{Q}(\alpha)$ is Galois over $\mathbb{Q}$, we have $[\mathbb{Q}(\alpha, \beta) : \mathbb{Q}] = [\mathbb{Q}(\alpha) : \mathbb{Q}][\mathbb{Q}(\beta) : \mathbb{Q}]$.*

Proof. Let $f_0(x)$ and $g_0(x)$ be the monic minimal polynomials for α and β; then $f_0(x)|f(x)$ and $g_0(x)|g(x)$, so it follows that there is no prime p for which $f_0(x)$ and $g_0(x)$ have repeated roots mod p. By the definition of the polynomial discriminant, this is equivalent to $(\operatorname{disc} f_0(x), \operatorname{disc} g_0(x)) = 1$, which by Lemma 3.4.1 implies that $(D_{\mathbb{Q}(\alpha)}, D_{\mathbb{Q}(\beta)}) = 1$. The previous corollary then gives the desired result. □

5.3. Discriminants and rings of integers

Theorem 5.3.1. *Let K be a number field, contained in number fields L and E. Suppose $E \cap L = K$, and E and L are Galois over K. Then*

$$D_{\mathscr{O}_E/\mathscr{O}_K}\mathscr{O}_{EL} \subseteq \mathscr{O}_E\mathscr{O}_L \subseteq EL.$$

If in addition we have $D_{\mathscr{O}_L/\mathscr{O}_K} + D_{\mathscr{O}_E/\mathscr{O}_K} = \mathscr{O}_K$, then we have $\mathscr{O}_{EL} = \mathscr{O}_E\mathscr{O}_L$.

The proof of Lemma 2.4.6 gives the following more general fact.

Commutative Algebra Fact 5.3.2. Given R an integral domain, and R-modules $M_1, M_2 \subseteq L$ for L some field containing R, then $M_1 \subseteq M_2$ if and only if $R_{\mathfrak{p}}M_1 \subseteq R_{\mathfrak{p}}M_2$ for all prime ideals of R.

More specifically, we have $M_1 = \cap_{\mathfrak{p}} R_{\mathfrak{p}} M_1$ and similarly for M_2.

We will need the following linear algebra lemma.

Lemma 5.3.3. *Let $K \subseteq L$ be a separable field extension of degree n, and let $(x_1, \dots, x_n)$ be a basis for L over K, and $(y_1, \dots, y_n) \subseteq L$ be a dual basis to $(x_1, \dots, x_n)$ with respect to the map $\mathrm{Tr}_{L/K}$: this is, $\mathrm{Tr}_{L/K}(x_iy_j) = \delta_{ij}$, where $\delta_{ij} = 1$ if $i = j$ and 0 if $i \neq j$. Then:*

(i) *if $x = \sum_{i=1}^n \alpha_i y_i$, we have $\alpha_i = \mathrm{Tr}_{L/K}(xx_i)$.*

(ii) *if $M_{\bar{x}}$ is the matrix with i,jth coefficient $\mathrm{Tr}_{L/K}(x_ix_j)$ and $M_{\bar{y}}$ is the same for the y_i, we have*

$$M_{\bar{x}}M_{\bar{y}} = I_n.$$

Proof. For (i), for any j we find

$$\mathrm{Tr}_{L/K}(xx_j) = \sum_{i=1}^n \mathrm{Tr}_{L/K}(\alpha_i y_i x_j) = \sum_{i=1}^n \alpha_i \mathrm{Tr}_{L/K}(x_j y_i) = \alpha_j,$$

as desired.

For (ii), we may then write $x_i = \sum_{j=1}^n \mathrm{Tr}_{L/K}(x_ix_j)y_j$. By definition, the (i,j)th coordinate of $M_{\bar{x}}M_{\bar{y}}$ is

$$\sum_{k=1}^n \mathrm{Tr}_{L/K}(x_ix_k)\,\mathrm{Tr}_{L/K}(y_ky_j) = \sum_{k=1}^n \mathrm{Tr}_{L/K}(\mathrm{Tr}_{L/K}(x_ix_k)y_ky_j) = \mathrm{Tr}_{L/K}(y_j\sum_{k=1}^n \mathrm{Tr}_{L/K}(x_ix_k)y_k)$$

by K-linearity of $\mathrm{Tr}_{L/K}$, and this can be rewritten

$$\mathrm{Tr}_{L/K}(y_jx_i) = \delta_{i,j}$$

by substitution from above, so we get the desired identity. □

Proof of theorem. We first conclude the second part of the theorem from the first. It is clear that $\mathscr{O}_E\mathscr{O}_L \subseteq \mathscr{O}_{EL}$, so we apply the first part of the theorem twice, reversing the roles of L and E, and see:

$$\mathscr{O}_{EL} = D_{\mathscr{O}_E/\mathscr{O}_K}\mathscr{O}_{EL} + D_{\mathscr{O}_L/\mathscr{O}_K}\mathscr{O}_{EL} \subseteq \mathscr{O}_E\mathscr{O}_L,$$

as desired.

Thus, it remains to prove the first part of the theorem. By the commutative algebra fact, and the fact from Lemma 3.3.4 that $D_{\mathscr{O}_{L,\mathfrak{p}}/\mathscr{O}_{K/\mathfrak{p}}} = (D_{\mathscr{O}_L,\mathscr{O}_K})_{\mathfrak{p}}$, it suffice to prove the statement after localizing at each prime of $\mathscr{O}_K$. Let $(x_1, \dots, x_n)$ be a basis of $\mathscr{O}_{E,\mathfrak{p}}$ over $\mathscr{O}_{K,\mathfrak{p}}$, and let $(y_1, \dots, y_n) \subseteq E$ be the dual basis with respect to $\mathrm{Tr}_{E/K}$. Our first task is to prove that $y_i \in \frac{1}{D_{\mathscr{O}_{E,\mathfrak{p}}/\mathscr{O}_{K,\mathfrak{p}}}}\mathscr{O}_{E,\mathfrak{p}}$ for all i.

By the second part of the lemma, we find

$$D_{\mathscr{O}_{E,\mathfrak{p}}/\mathscr{O}_{K,\mathfrak{p}}} D_{E/K}(y_1,\dots,y_n) = D_{E/K}(x_1,\dots,x_n)D_{E/K}(y_1,\dots,y_n) = \det M_{\bar{x}} \det M_{\bar{y}} = 1,$$

and in particular, $M_{\bar{y}}$ has coefficients in $\frac{1}{D_{\mathscr{O}_{E,\mathfrak{p}}/\mathscr{O}_{K,\mathfrak{p}}}}\mathscr{O}_E$. By the first part of the lemma, if we consider $M_{\bar{y}}$ to be the matrix of a linear map $E \to E$ with respect to the basis $\bar{x}$, we find

$$y_i = M_{\bar{y}}(x_i) \in \frac{1}{D_{\mathscr{O}_{E,\mathfrak{p}}/\mathscr{O}_{K,\mathfrak{p}}}}\mathscr{O}_{E,\mathfrak{p}},$$

as desired.

Now, note that the hypothesis that $E \cap L = K$ implies that $(y_1,\dots,y_n)$ are also a basis for EL over L. Choose $x \in \mathscr{O}_{EL,\mathfrak{p}}$, and write $x = \sum_i \alpha_i y_i$ for $\alpha_i \in L$. Using the first part of the lemma again, we have $\alpha_i = \operatorname{Tr}_{EL/L}(xx_i) \in \mathscr{O}_{L,\mathfrak{p}}$, so we conclude that $x \in \frac{1}{D_{\mathscr{O}_{E,\mathfrak{p}}/\mathscr{O}_{K,\mathfrak{p}}}}\mathscr{O}_{E,\mathfrak{p}}\mathscr{O}_{L,\mathfrak{p}}$, and since x was arbitrary, we have $D_{\mathscr{O}_{E,\mathfrak{p}}/\mathscr{O}_{K,\mathfrak{p}}}\mathscr{O}_{EL,\mathfrak{p}} \subseteq \mathscr{O}_{E,\mathfrak{p}}\mathscr{O}_{L,\mathfrak{p}}$, completing the proof of the theorem. □

We can also conclude the following corollary, as we already saw in Corollary 5.2.3 that over $\mathbb{Q}$, the condition $D_{\mathscr{O}_K/\mathbb{Z}} + D_{\mathscr{O}_L/\mathbb{Z}} = \mathbb{Z}$ implies that $K \cap L$ is unramified at every prime, and hence equal to $\mathbb{Q}$.

Corollary 5.3.4. *Given number fields K, L, if $D_{\mathscr{O}_K/\mathbb{Z}} + D_{\mathscr{O}_L/\mathbb{Z}} = \mathbb{Z}$, we have $\mathscr{O}_{KL} = \mathscr{O}_K\mathscr{O}_L$.*

We conclude with a simple application:

Theorem 5.3.5. *Given $m, n \in \mathbb{Z}$ square-free and not equal 1, with $m \equiv 1 \pmod 4$ and $(m,n) = 1$, then*

$$\mathscr{O}_{\mathbb{Q}(\sqrt{m},\sqrt{n})} = \mathscr{O}_{\mathbb{Q}(\sqrt{m})}\mathscr{O}_{\mathbb{Q}(\sqrt{n})}.$$

Proof. We already calculated that $D_{\mathscr{O}_{\mathbb{Q}(\sqrt{m})}/\mathbb{Z})} = m$. If $n \equiv 1 \pmod 4$ we have $D_{\mathscr{O}_{\mathbb{Q}(\sqrt{n})}/\mathbb{Z})} = n$; otherwise, we calculate

$$D_{\mathscr{O}_{\mathbb{Q}(\sqrt{n})}/\mathbb{Z})} = \mathscr{O}_{\mathbb{Q}(\sqrt{n})/\mathbb{Q}}(1,\sqrt{n}) = 4n.$$

Either way, the discriminants are relatively prime, so the previous theorem implies the desired result. □

However, the condition that $(m,n) = 1$ in the previous theorem is not sufficient; see Exercise 5.1.

5.4. The cyclotomic ring of integers

Since ζ_m satisfies $x^m - 1 = 0$, we have $\mathbb{Z}[\zeta_m] \subseteq \mathscr{O}_{\mathbb{Q}(\zeta_m)}$. We can now prove the following theorem fairly easily:

Theorem 5.4.1. *If ζ_m is a primitive mth root of unity, then $\mathscr{O}_{\mathbb{Q}(\zeta_m)} = \mathbb{Z}[\zeta_m]$.*

To do so, we need the following lemma:

Lemma 5.4.2. *Suppose $m = p^n$. Then in $\mathscr{O}_{\mathbb{Q}(\zeta_{p^n})}$, we have:*

(i) *The elements $(1 - \zeta_{p^n}^i)$ all generate the same ideal for $(i,p) = 1$;*
(ii) *If $p > 2$, the element $1 + \zeta_{p^n}$ is a unit;*
(iii) *We have the ideal factorization $(p) = (1 - \zeta_{p^n})^{(p-1)p^{n-1}}$*

Proof. (i) We have $(1-\zeta_{p^n}^i)/(1-\zeta_{p^n}) = 1+\zeta_{p^n}+\cdots+\zeta_{p^n}^{i-1} \in \mathbb{Z}[\zeta_{p^n}]$. On the other hand, if $ii' \equiv 1 \pmod{p^n}$, we have $(1-\zeta_{p^n})/(1-\zeta_{p^n}^i) = (1-\zeta_{p^n}^{ii'})/(1-\zeta_{p^n}^i) = 1+\zeta_{p^n}^i+\cdots+\zeta_{p^n}^{i(i'-1)} \in \mathbb{Z}[\zeta_{p^n}]$, so we find that $(1-\zeta_{p^n}^i)$ and $(1-\zeta_{p^n})$ divide one another in $\mathbb{Z}[\zeta_{p^n}]$ and hence in $\mathscr{O}_{\mathbb{Q}(\zeta_{p^n})}$. (ii) By (i), $1+\zeta_{p^n} = (1-\zeta_{p^n}^2)/(1-\zeta_{p^n})$ is a unit in $\mathbb{Z}[\zeta_{p^n}]$, hence in $\mathscr{O}_{\mathbb{Q}(\zeta_{p^n})}$. (iii) We have $\Phi_{p^n}(x) = (x^{p^n}-1)/(x^{p^{n-1}}-1) = 1+x^{p^{n-1}}+\cdots+x^{(p-1)p^{n-1}} = \prod_{(i,p)=1}(x-\zeta_{p^n}^i)$, so $\Phi_{p^n}(1) = p = \prod_{(i,p)=1}(1-\zeta_{p^n}^i)$. By (i), the right side generates the same ideal as $(1-\zeta_{p^n})^{(p-1)p^{n-1}}$. □

We now prove the theorem.

Proof of the theorem. We first prove the case that $m = p^n$ for some n. It is enough to show that for every prime q, we have $\mathscr{O}_{\mathbb{Q}(\zeta_{p^n}),q} = \mathbb{Z}[\zeta_{p^n}]_q$. Now, for $q \neq p$, we know that $\Phi_{p^n}(x)$ is separable mod q, so $\operatorname{disc}\Phi_{p^n}(x) = D_{\mathbb{Q}(\zeta_{p^n})/\mathbb{Q}}(1, \zeta_{p^n}, \ldots, \zeta_{p^n}^{(p-1)p^{n-1}})$ is a unit in $\mathbb{Z}_(q)$, and it follows that $\mathscr{O}_{\mathbb{Q}(\zeta_{p^n}),q} = \mathbb{Z}[\zeta_{p^n}]_q$. It remains to consider the case that $q = p$. By (iii) above, we have that $(1-\zeta_{p^n})$ is the unique prime ideal lying above p, and we also observe by the degree formula that we must have $\mathscr{O}_{\mathbb{Q}(\zeta_{p^n})}/(1-\zeta_{p^n}) \cong \mathbb{Z}/(p)$. Thus, by Theorem 5.1.1, we have

$$\mathscr{O}_{\mathbb{Q}(\zeta_{p^n}),p} = \mathbb{Z}[1, 1-\zeta_{p^n}]_p = \mathbb{Z}[\zeta_{p^n}]_p.$$

This completes the proof for the $m = p^n$ case.

But for the general case, we induct on the number of prime factors; as in the proof of the irreducibility of $\Phi_m(x)$, we have that $(D_{\mathbb{Q}(\zeta_{p^m})}, D_{\mathbb{Q}(\zeta_r)}) = 1$ if p is prime to r, and then by Theorem 5.3.1, we have

$$\mathscr{O}_{\mathbb{Q}(\zeta_{p^m r})} = \mathscr{O}_{\mathbb{Q}(\zeta_{p^m}, \zeta^r)} = \mathscr{O}_{\mathbb{Q}(\zeta_{p^m})}\mathscr{O}_{\mathbb{Q}(\zeta_r)},$$

and since $\mathbb{Z}[\zeta_{p^m}, \zeta_r] = \mathbb{Z}[\zeta_{p^m}]\mathbb{Z}[\zeta_r]$, we are able to carry out the induction. □

Example 5.4.3. This example is due to Dedekind. Let α be a root of $x^3 - x^2 - 2x - 8$, and consider $\mathbb{Q}(\alpha)$. One can show that $\mathscr{O}_{\mathbb{Q}(\alpha)}$ is generated by α and $\frac{4}{\alpha}$. By analyzing $\mathscr{O}_{\mathbb{Q}(\alpha)}$ modulo 2, one can show that it is not generated over $\mathbb{Z}$ by any single element.

5.5. Fermat's Last Theorem

We will now prove certain cases of Fermat's Last Theorem, by studying the arithmetic of $\mathbb{Z}[\zeta_p]$. We introduce the following (non-standard) terminology:

Definition 5.5.1. We say that an odd prime p is **h-regular** if p does not divide the class number (i.e., the order of the ideal class group) of $\mathbb{Q}(\zeta_p)$. We say that p is **strongly** h-regular if it is h-regular, and if it satisfies the further property that for all units $u \in \mathbb{Z}[\zeta_p]^\times$ such that $u \equiv m \pmod{p\mathbb{Z}[\zeta_p]}$ for some $m \in \mathbb{Z}$, then $u = u'^p$ for some $u' \in \mathbb{Z}[\zeta_p]$.

Remark 5.5.2. We will see later, by class field theory, the fact (known as Kummer's lemma) that a prime p being h-regular implies that it is strongly h-regular. We will also give an explicitly computable criterion for h-regularity in terms of Bernoulli numbers.

Fermat's Last Theorem for prime exponents is classically broken up into two cases: case I is the situation that all terms are prime to p, whereas in case II, one term may be divisible by p. Case II is the harder one. While both cases begin with a factorization of the left side in $\mathbb{Z}[\zeta_p]$, case I then shows that the factors are prime, and produces a linear relation among different powers of ζ_p modulo p. Case II involves subtler analysis, and an induction argument on the power of p dividing the appropriate term of the equation.

We leave case I as exercise; see Exercise 5.3. Our aim is thus to prove case II of Fermat's Last Theorem for strongly h-regular primes.

Theorem 5.5.3. *Let p be strongly h-regular. Then there are no non-zero integer solutions to*

$$x^p + y^p = z^p$$

with p dividing (at least) one of x, y, z.

Proof. Suppose to the contrary that a solution exists. We first note that we may assume that $p|z$, but p is prime to x, y: indeed, we may certainly assume that x, y, z have no common factors; because p is odd, we may write $x^p + y^p + (-z)^p = 0$, so the situation is symmetric in x, y, z, and we assume that $p|z$. But then, if p divided x or y, it would have to divide all three, contradicting relative primality.

Let $p^{n'}$ be the largest power of p dividing z. We will prove by induction on n' that no such solution is possible. However, we will induct on the following slightly stronger statement: there do not exist $\alpha, \beta, \gamma \in \mathbb{Z}[\zeta_p], u \in \mathbb{Z}[\zeta_p]^\times,$ and $n \in \mathbb{N}$ such that

$$\alpha^p + \beta^p + u(1-\zeta_p)^{pn}\gamma^p = 0, \tag{5.5.3.1}$$

and $(1-\zeta_p)$ does not divide $\alpha\beta\gamma$. (Note that n here is $n'(p-1)$.) Before beginning the induction, we make some general observations. If we had such a solution, we would obtain an identity of ideals:

$$\prod_{j=0}^{p-1}(\alpha + \zeta_p^j\beta) = (1-\zeta_p)^{pn}(\gamma)^p. \tag{5.5.3.2}$$

Note that $\alpha + \zeta_p^j\beta \equiv \alpha + \beta \pmod{1-\zeta_p}$, so since $1-\zeta_p$ must divide at least one factor on the left, it must in fact divide all of them.

We now observe that all the $\alpha + \zeta_p^j\beta$ must be distinct modulo $(1-\zeta_p^2)$. If we had $\alpha + \zeta_p^j\beta \equiv \alpha + \zeta_p^{j'}\beta \pmod{(1-\zeta_p)^2}$ for $j' > j$, we would have $(1-\zeta_p)^2|(\zeta_p^j(1-\zeta_p^{j'-j})\beta)$ in $\mathbb{Z}[\zeta_p]$. Since $(1-\zeta_p^{j'-j})$ is a unit multiple of $1-\zeta_p$, this would imply $(1-\zeta_p)|\beta$, contradicting our initial hypothesis.

The base case for our induction is $n = 1$, and we claim that we have already proved that this cannot occur. Indeed, we need only recall that $\mathbb{Z}[\zeta_p]/(1-\zeta_p) \cong \mathbb{Z}/(p)$, and we then conclude that there are p multiples of $(1-\zeta_p)$ modulo $(1-\zeta_p)^2$, hence $p-1$ non-zero multiples of $(1-\zeta_p)$, so in order for all p terms of the form $\alpha + \zeta_p^j\beta$ to be distinct modulo $(1-\zeta_p)^2$, one of them must be congruent to 0. But since we observed earlier that $(1-\zeta_p)$ must divide all p terms, we find that we cannot have $n = 1$.

We now prove the induction step, by showing that a solution for a given $n > 1$ would yield a solution for $n-1$. The basic idea is to use the following relation between our terms:

$$\zeta_p(\alpha + \zeta_p^{p-1}\beta) + (\alpha + \zeta_p\beta) - (1+\zeta_p)(\alpha+\beta) = 0 \tag{5.5.3.3}$$

along with the hypotheses of strong h-regularity. If $\alpha + \zeta_p^{j_0}\beta$ is the term which is 0 modulo $(1-\zeta_p)^2$, we may replace β by $\zeta_p^{j_0}\beta$ in the statement we are trying to prove, so we may assume that $(1-\zeta_p)^2$ divides $\alpha+\beta$, but does not divide $\alpha+\zeta_p^j\beta$ for $0 < j < p$.

Denote by $\mathfrak{d}$ the ideal generated by (α, β). Certainly, $\mathfrak{d}$ is a common divisor of any two of the terms $\alpha + \zeta_p^j\beta$. Since $(1-\zeta_p)$ is prime and does not divide $\alpha\beta$, but does divide each term, it is easy to check, as is done for case I, that the ideal gcd of any two terms is precisely $\mathfrak{d}(1-\zeta_p)$. If we write $\alpha + \zeta_p\beta = \mathfrak{d}(1-\zeta_p)I_j$ for $j > 0$, and $\alpha + \zeta_p\beta = \mathfrak{d}(1-\zeta_p)^{(n-1)p+1}I_0$, we have that no two of the I_j have any common ideal factor, so by (5.5.3.2), we find that each I_j is a pth power, say $I_j = J_j^p$.

Dividing any two such terms, we see that J_j^p/J_0^p must be a principal fractional ideal. Because we assumed that p doesn't divide the class number, we have that J_j/J_0 is a principal fractional ideal, so we can choose $x_j \in \mathbb{Q}(\zeta_p)$ such that $(x_j) = J_j/J_0$; we have in particular that the prime factorization of x_j has no $(1-\zeta_p)$ term in it. If we choose appropriate units $u_j \in \mathbb{Z}[\zeta_p]^\times$, we can pass from ideals to elements to get:

$$\frac{\alpha + \zeta_p^j\beta}{\alpha+\beta} = \frac{u_j x_j^p}{(1-\zeta_p)^{p(n-1)}}.$$

If we divide (5.5.3.3) through by $(\alpha+\beta)$, and substitute in the above formula, and clear the $(1-\zeta_p)^{p(n-1)}$ from the denominator, we obtain:

$$\zeta_p u_{p-1} x_{p-1}^p + u_1 x_1^p - (1+\zeta_p)(1-\zeta_p)^{p(n-1)}.$$

Now, the x_j are in $\mathbb{Q}(\zeta_p)$, but we can write them as quotients of elements of $\mathbb{Z}[\zeta_p]$, and since $(1-\zeta_p)$ is prime and doesn't appear in the factorization of the x_j, we can require the numerators and denominators to be prime to $(1-\zeta_p)$. Clearing denominators and then dividing through by $\zeta_p u_{p-1}$, we find

$$\alpha'^p + \frac{u_1}{\zeta_p u_{p-1}}\beta'^p - \frac{(1+\zeta_p)}{\zeta_p u_{p-1}}(1-\zeta_p)^{p(n-1)}\gamma'^p,$$

with $\alpha', \beta', \gamma' \in \mathbb{Z}[\zeta_p]$ and not divisible by $(1-\zeta_p)$.

The last step is to consider the above equation modulo p: one first notes that for any $z \in \mathbb{Z}[\zeta_p]$, we have $z^p \equiv m \pmod{p}$ for some $m \in \mathbb{Z}$, so in particular, we have $\alpha'^p \equiv m_1 \pmod{p}$ and $\beta'^p \equiv m_2 \pmod{p}$, both non-zero. Since $n > 1$, the last term is 0 modulo p, so we have

$$m_1 + \frac{u_1}{\zeta_p u_{p-1}} m_2 \pmod{p},$$

and if $m_2' m_2 = 1 \in \mathbb{Z}/(p)$, we find $\frac{u_1}{\zeta_p u_{p-1}} \equiv -m_1 m_2' \pmod{p}$. Thus, we can finally apply the hypothesis of strong h-regularity to conclude that $\frac{u_1}{\zeta_p u_{p-1}} = u^p$ for some $u \in \mathbb{Z}[\zeta_p]^\times$, and if we replace β' by $u\beta'$, we have that

$$\alpha'^p + \beta'^p - \frac{(1+\zeta_p)}{\zeta_p u_{p-1}}(1-\zeta_p)^{p(n-1)}\gamma'^p$$

is of the same form as (5.5.3.1), but with n replaced by $n-1$, yielding the desired contradiction by induction. □

5.6. Exercises

EXERCISE 5.1. Find an element in $\mathscr{O}_{\mathbb{Q}(\sqrt{3},\sqrt{-1})}$ which is not in $\mathscr{O}_{\mathbb{Q}(\sqrt{3})}\mathscr{O}_{\mathbb{Q}(\sqrt{-1})}$.

In the next two exercises, $\bar{z}$ denotes the complex conjugate of z.

EXERCISE 5.2. Prove that in $\mathbb{Z}[\zeta_p]$, if $u \in \mathbb{Z}[\zeta_p]^\times$, then $u/\bar{u} = \pm\zeta_p^i$ for some i. (Note: if you like, you may make your life easier on the next problem by proving that in fact, $u/\bar{u} = \zeta_p^i$, by considering the situation modulo $1 - \zeta_p$)

EXERCISE 5.3. Recall that we say an odd prime p is h-regular if p does not divide the order of the ideal class group of $\mathscr{O}_{\mathbb{Q}(\zeta_p)} = \mathbb{Z}(\zeta_p)$. Prove the following theorem: given an h-regular prime p, there do not exist integer solutions to the equation

$$x^p + y^p = z^p,$$

with x, y, z all prime to p.

Proceed in the following steps:

a) Supposing to the contrary a solution to the equation, factor the left side in $\mathbb{Z}[\zeta_p]$, and show that the factors are all relatively prime, in the sense that they generate the unit ideal. Conclude that each factor on the left is a pth power, when considered as an ideal.
b) Using the h-regularity hypothesis, conclude that each factor on the left is of the form $u\alpha^p$, for $u \in \mathbb{Z}[\zeta_p]^\times$ and $\alpha \in \mathbb{Z}[\zeta_p]$.
c) Pick a particular factor on the left, and show that modulo p, the above α satisfies $\alpha^p \equiv \bar{\alpha}^p$.
d) Determine the relationship between $u\alpha^p$ and $\bar{u}\bar{\alpha}^p$ modulo p, and use this to obtain a linear relation modulo p between different powers of ζ_p.
e) Show that $1, \zeta_p, \ldots, \zeta_p^{p-2}$ are linearly independent in $\mathbb{Z}[\zeta_p]/(p)$.
f) Working case by case, conclude that for $p \geqslant 5$, the only possibility is that $x \equiv y \pmod{p}$.
g) Noting that p is odd, use a substitution in the initial equation to conclude that in fact, we must also have $x \equiv -z \pmod{p}$, and derive a contradiction in the $p \geqslant 5$ case.
h) Show that the case $p = 3$ is not an issue by considering the equation modulo 9.

You will need to use structure results on $\mathscr{O}_{\mathbb{Q}(\zeta_p)}$ from Lemma 5.4.2.

Notes

The treatment of Fermat's Last Theorem for regular primes was taken from Keith Conrad's notes [**2**].

CHAPTER 6

The Analytic Class Number Formula

We now prove the analytic class number formula, which relates invariants of number fields, including the class number and regulator, to a generalization of the Riemann zeta function called the Dedekind zeta function. In this chapter, we prove the general formula, while in the next chapter, we will specialize to subfields of cyclotomic field, and use Dirichlet L-series to prove more explicit formulas in that case (including for quadratic extensions).

6.1. Dedekind zeta functions

We will only need the most basic properties of Dedekind zeta functions, but we take the opportunity to discuss briefly some deeper conjectures and consequences that could be studied in a course in analytic number theory. Recall the definition of the Riemann zeta function:

Definition 6.1.1. The **Riemann zeta function** is defined by the sum

$$\zeta(s) = \sum_{n=1}^{\infty} \frac{1}{n^s}.$$

Since these are algebraic number theory notes, we could observe that this could be described as a sum over all non-zero ideals of $\mathbb{Z}$, with n being the norm of each ideal. This naturally leads to the generalization:

Definition 6.1.2. Let K be a number field. The **Dedekind zeta function** of K is defined by the sum

$$\zeta_K(s) = \sum_{I} \frac{1}{N(I)^s},$$

where I runs over all non-zero ideals of $\mathscr{O}_K$.

We will prove the following statements about Dedekind zeta functions:

Theorem 6.1.3. *Let K be a number field. Then:*

(i) *The sum for $\zeta_K(s)$ converges for $s > 1$.*

(ii) *(Analytic class number formula) $\zeta_K(s)$ goes to infinity at $s = 1$, but we have*

$$\lim_{s \to 1^+} (s-1)\zeta_K(s) = \frac{2^{r_1+r_2}\pi^{r_2} h_K R_K}{m_K |D_K|^{1/2}}.$$

(iii) *(Euler product) For $s > 1$, we can also write*

$$\zeta_K(s) = \prod_{\mathfrak{p}} (\frac{1}{1 - \frac{1}{N(\mathfrak{p})^s}}),$$

where $\mathfrak{p}$ ranges over non-zero prime ideals of $\mathscr{O}_K$.

Recall our notation: r_1 is the number of real imbeddings of K into $\mathbb{C}$, and r_2 is the number of conjugate pairs of non-real imbeddings. h_K is the class number, R_K is the regulator, m_K is the number of roots of unity contained in K, and D_K is the absolute discriminant.

In fact, the proofs of (iii) and (ii) are unrelated, except that both make use of (i). However, (iii) plays a key role in our applications of (ii), since together they establish a subtle relationship between the behavior of primes in $\mathscr{O}_K$, and the constants associated to K. An explicit example is worked out in the exercises, and we work more generally with sub-cyclotomic fields in the next chapter. The analytic class number formula is frequently named after Dirichlet, who worked out a version in the case of quadratic extensions. However, he had difficulty extending this formula more generally, and it was Kummer who developed the ideas that led to the more general version.

For the analytic class number formula and its applications, we will not need any deeper properties of $\zeta_K(s)$ than those described by the theorem. However, we note that a lot more is known, and conjectured, and we make a brief survey of these results. For instance:

Theorem 6.1.4. *The function $\zeta_K(s)$ has a functional equation*

$$\zeta_K(1-s)\Gamma(\frac{1-s}{2})^{r_1}\Gamma(1-s)^{r_2} = \zeta_K(s)\Gamma(\frac{s}{2})^{r_1}\Gamma(s)^{r_2}(4^{-r_2}\pi^{-n}|D_K|)^{s-1}.$$

This may be used to extend $\zeta_K(s)$ to an analytic function on $\mathbb{C}$, except for a simple pole at $s=1$.

Here, $\Gamma(s)$ is defined for $\Re s > 0$ by $\Gamma(s) = \int_0^\infty x^s e^{-x} dx$, and for other $s \in \mathbb{C}$ by recursively applying the identity $\Gamma(s+1) = s\Gamma(s)$.

Using this theorem, the zeroes of $\zeta_K(s)$ outside the region $0 \leqslant \Re s \leqslant 1$ (known as the **critical strip**) are easily analyzed. However, it is the zeroes of $\zeta_K(s)$ inside the critical strip which turn out to be particularly important. We have:

Conjecture 6.1.5. *(Riemann, extended[1]) The zeroes of $\zeta_K(s)$ inside the critical strip all lie on the line $\Re s = \frac{1}{2}$.*

Riemann conjectured the $K = \mathbb{Q}$ case of this, and discovered the close relationship between the location of zeroes in the critical strip and the distribution of prime numbers. In fact, we have the following remarkable theorem:

Theorem 6.1.6. *The Riemann hypothesis for $K = \mathbb{Q}$ is equivalent to the estimate*

$$\pi(x) = \operatorname{li} x + O(x^{\frac{1}{2}} \log x),$$

where $\pi(x) = \#\{p \text{ prime}, p \leqslant x\}$, and $\operatorname{li} x = \int_2^x \frac{dy}{\log y} = \frac{x}{\log x} + O(\frac{x}{(\log x)^2})$. More generally, given $t_0 \in [\frac{1}{2}, 1)$, all the zeroes of $\zeta_{\mathbb{Q}}(s)$ in the critical strip have $\Re s \leqslant t_0$ if and only if

$$\pi(x) = \operatorname{li} x + O(x^{t_0} \log x).$$

The extended Riemann hypothesis implies similar results for distribution of prime ideals in rings of integers, but it also has some applications of a different flavor.

[1] There is some disagreement of terminology about the names of the various specific generalizations of the Riemann hypothesis, but referring to this as the extended Riemann hypothesis seems reasonable.

One application of the extended Riemann hypothesis was given by Stark. We showed earlier that Minkowski's formulas yielded a lower bound for the discriminant of a number field; by Stirling's approximation for $n!$, this bound may be rewritten as

$$\log|D_K| \geqslant \alpha n + \beta r_1 - o(n).$$

However, the α and β that one derives from Minkowski's bounds are not optimal: they are $\alpha = 2 - \log(4/\pi) \sim 1.76$ and $\beta = \log(4/\pi) \sim .24$. Assuming the extended Riemann hypothesis, Stark showed the following:

THEOREM 6.1.7. *(Stark) Suppose that the extended Riemann hypothesis holds. Then*

$$\log|D_K| \geqslant \alpha n + \beta r_1 - o(n)$$

with $\alpha = \log(8\pi) - \Gamma'(1) \sim 3.80$, *and* $\beta = \log(2/\pi) \sim 1.57$.

The extended Riemann hypothesis also gives, for cyclotomic fields, the following stronger statement combining the prime number theorem with Dirichlet's theorem on primes in arithmetic sequences:

THEOREM 6.1.8. *For a given n, the extended Riemann hypothesis, applied to the case of $\mathbb{Q}(\zeta_n)$, implies that for all k with $(k, n) = 1$, we have*

$$\pi_{k,n}(x) = \frac{1}{\phi(n)} \operatorname{li} x + O(x^{\frac{1}{2}} \log x),$$

where $\pi_{k,n}(x) = \#\{p \text{ prime}, p \leqslant x, p \equiv k \pmod{n}\}$.

This is frequently stated as a consequence of the Riemann hypothesis for Dirichlet L-series, but that follows from the extended Riemann hypothesis for $\mathbb{Q}(\zeta_n)$ because the zeta functions factor as finite products including the Dirichlet L-series.

6.2. Overview of argument

The proofs of (i) and (ii) are established by defining:

DEFINITION 6.2.1. For an ideal class C of $\mathscr{O}_K$, define

$$\zeta_{K,C}(s) = \sum_{I,C} \frac{1}{N(I)^s},$$

where I runs over all ideals of $\mathscr{O}_K$ in the class C.

We then want to prove:

THEOREM 6.2.2. *Let K be a number field, and C an ideal class of $\mathscr{O}_K$. Then:*

(i) *The sum for $\zeta_{K,C}(s)$ converges for $\Re s > 1$.*

(ii) *$\zeta_{K,C}(s)$ has a simple pole at $s = 1$, with residue given by*

$$\lim_{s \to 1^+} (s-1)\zeta_{K,C}(s) = \frac{2^{r_1+r_2}\pi^{r_2} R_K}{m_K |D_K|^{1/2}}.$$

Given this theorem, it is clear that

$$\zeta_K(s) = \sum_C \zeta_{K,C}(s),$$

and since this is a finite sum with h_K terms, we obtain (i) and (ii) of the prior theorem.

Thus, from now on we fix an ideal class C. The strategy is to rewrite the sum for $\zeta_{K,C}(s)$ as a sum over certain elements of $\mathscr{O}_K$, and use what we know about the lattices $\varphi(\mathscr{O}_K)$ and $\psi(\mathscr{O}_K^*)$ to give a more explicit form, and evaluate it.

We start by fixing an ideal $I_0 \in C^{-1}$, so that a fractional ideal I is an ideal in C if and only if $I_0 I = (\alpha)$ for some $\alpha \in K$ with $\alpha \in I_0$. Thus, we can write

$$\zeta_{K,C}(s) = \sum_{(\alpha):\alpha\in I_0} \frac{N(I_0)^s}{N((\alpha))^s} = N(I_0)^s \sum_{(\alpha):\alpha\in I_0} \frac{1}{N((\alpha))^s}.$$

Note that the above sum is over principal ideals, and not over the elements α themselves. In order to evaluate the sum, we wish to give a more explicit description in terms of the elements α. To do this, we fix a set of fundamental units $u_1, \dots, u_{r_1+r_2+1}$ for $\mathscr{O}_K$; i.e., units such that the $\psi(u_i)$ form a basis for the lattice $\psi(\mathscr{O}_K^*)$. We write $\bar{x}_0 = (\underbrace{1,\dots,1}_{r_1},\underbrace{2,\dots,2}_{r_2})$, and set $\bar{x}_i = \psi(u_i)$ for $i > 0$. Since the $\psi(u_i)$ span the hyperplane H with coordinates summing to 0, and $\bar{x}_0$ lies outside H, we have that the $\bar{x}_i$ together form a basis for $\mathbb{R}^{r_1+r_2}$.

We can now define:

Definition 6.2.3. Given the choice of fundamental units u_i, we define a **fundamental domain** $X_{K,\bar{u}}$ in $\mathbb{R}^n$ to be the set of $\bar{z} \subseteq \mathbb{R}^n$ satisfying:

(i) $N(\bar{z}) \neq 0$;
(ii) $\chi(\bar{z}) = \sum_{i=0}^{r_1+r_2-1} c_i \bar{x}_i$ with $0 \leqslant c_i < 1$ for $i > 0$.
(iii) $0 \leqslant \arg z_1 < 2\pi/m_K$ for $r_1 = 0$, and $z_1 > 0$ for $r_1 > 0$, where z_1 is the 1st coordinate of $\bar{z} \in \mathbb{R}^n = \mathbb{R}^{r_1} \times \mathbb{C}^{r_2}$.

We will see:

Theorem 6.2.4. *The fundamental domain $X_{K,\bar{u}}$ is a cone in $\mathbb{R}^n$, such that*

$$\mathrm{vol}(X_{K,\bar{u}} \cap \{\bar{z} \in \mathbb{R}^n : N(\bar{z}) \leqslant 1\}) = \frac{2^{r_1}\pi^{r_2} R_K}{m_K},$$

and for every $\bar{z} \in \mathbb{R}^n$, there is a unique $u \in \mathscr{O}_K^*$ such that $\varphi(u) \cdot \bar{z} \in X_{K,\bar{u}}$.

Here the notation $\cdot$ is as in Definition 4.9.3. Using the compatibility of the norms on K and $\mathbb{R}^n$, the theorem allows us to write

$$\zeta_{K,C}(s) = N(I_0)^s \sum_{\bar{z}\in\varphi(I_0)\cap X_{K,\bar{u}}} \frac{1}{N(\bar{z})^s}.$$

We will then prove the following general theorem:

Theorem 6.2.5. *Let $X \subseteq \mathbb{R}^n$ be any cone containing only points $\bar{z}$ with $N(\bar{z}) \neq 0$, and such that $X \cap \{\bar{z} \in \mathbb{R}^n : N(\bar{z}) \leqslant 1\}$ is bounded and has some volume v_X. Let $L \subseteq \mathbb{R}^n$ be a lattice of full rank with volume v_L. Define*

$$\zeta_{X,L}(s) = \sum_{\bar{z}\in X\cap L} \frac{1}{N(\bar{z})^s}.$$

Then this sum converges for $s > 1$, and we have

$$\lim_{s\to 1^+} (s-1)\zeta_{X,L}(s) = \frac{v_X}{v_L}.$$

Recalling that in Theorem 4.2.2 we had already calculated

$$\operatorname{vol}\varphi(I_0)=\frac{N(I_0)|D_K|^{1/2}}{2^{r_2}},$$

we can put together the previous two theorems to find that the sum for $\zeta_{K,C}(s)$ converges for all $s>1$, and we have

$$\lim_{s\to 1^+}(s-1)\zeta_{K,C}(s)=N(I_0)\frac{2^{r_1}\pi^{r_2}R_K}{m_K}\frac{2^{r_2}}{N(I_0)|D_K|^{1/2}}=\frac{2^{r_1+r_2}\pi^{r_2}R_K}{m_K|D_K|^{1/2}},$$

as desired. We therefore see that to prove the analytic class number formula, it suffices to prove the two theorems above.

6.3. The Euler product

We will now prove (iii) of Theorem 6.1.3, assuming (i). More specifically, we have:

PROPOSITION 6.3.1. *Given (i) of Theorem 6.1.3, for $s>1$ and $N\in\mathbb{N}$ we have*

$$|\prod_{\mathfrak{p}:N(\mathfrak{p})\leqslant N}(1-\frac{1}{N(\mathfrak{p})^s})^{-1}-\zeta_K(s)|<\sum_{I:N(I)>N}\frac{1}{N(I)^s},$$

and the right side goes to 0 *as* $N\to\infty$*. In particular, (iii) of Theorem 6.1.3 follows.*

PROOF. We note that the right hand side goes to 0 simply by (i) of Theorem 6.1.3. (iii) of Theorem 6.1.3 will follow simply by taking the limit as $N\to\infty$. It thus remains to prove the desired bound.

Let $\mathfrak{p}_1,\dots,\mathfrak{p}_r$ be the prime ideals of $\mathscr{O}_K$ with norm at most N. Then

$$(1-\frac{1}{N(\mathfrak{p}_i)^s})^{-1}=1+\frac{1}{N(\mathfrak{p}_i)^s}+\frac{1}{N(\mathfrak{p}_i)^{2s}}+\dots,$$

so

$$\prod_{\mathfrak{p}:N(\mathfrak{p})\leqslant N}(1-\frac{1}{N(\mathfrak{p})^s})^{-1}=\sum_{k_1,\dots,k_r\geqslant 0}\frac{1}{N(\mathfrak{p}_1^{k_1}\cdots\mathfrak{p}_r^{k_r})^s}=\sum_{I\in S}\frac{1}{N^s},$$

where S is the set of ideals divisible only by $\mathfrak{p}_1,\dots,\mathfrak{p}_r$. Now, as we have seen, S includes all ideals I with $N(I)\leqslant N$, so in fact we have

$$0<\zeta_K(s)-\prod_{\mathfrak{p}:N(\mathfrak{p})\leqslant N}(1-\frac{1}{N(\mathfrak{p})^s})^{-1}<\sum_{I:N(I)>N}\frac{1}{N(I)^s},$$

giving the desired result. □

6.4. The fundamental domain

Recall we had defined:

DEFINITION 6.4.1. For an ideal class C of $\mathscr{O}_K$, define

$$\zeta_{K,C}(s)=\sum_{I,C}\frac{1}{N(I)^s},$$

where I runs over all ideals of $\mathscr{O}_K$ in the class C.

We had also fixed a set of fundamental units $u_1, \ldots, u_{r_1+r_2+1}$ for $\mathscr{O}_K$; i.e., units such that the $\psi(u_i)$ form a basis for the lattice $\psi(\mathscr{O}_K^*)$. We write $\bar{x}_0 = (\underbrace{1, \ldots, 1}_{r_1}, \underbrace{2, \ldots, 2}_{r_2})$, and set $\bar{x}_i = \psi(u_i)$ for $i > 0$. Since the $\psi(u_i)$ span the hyperplane H with coordinates summing to 0, and $\bar{x}_0$ lies outside H, we have that the $\bar{x}_i$ together form a basis for $\mathbb{R}^{r_1+r_2}$.

We can define:

Definition 6.4.2. Given the choice of fundamental units u_i, we define a **fundamental domain** $X_{K,\bar{u}}$ in $\mathbb{R}^n$ to be the set of $\bar{z} \subseteq \mathbb{R}^n$ satisfying:

(i) $N(\bar{z}) \neq 0$;
(ii) $\chi(\bar{z}) = \sum_{i=0}^{r_1+r_2-1} c_i \bar{x}_i$ with $0 \leqslant c_i < 1$ for $i > 0$.
(iii) $0 \leqslant \arg z_1 < 2\pi/m_K$ for $r_1 = 0$, and $z_1 > 0$ for $r_1 > 0$, where z_1 is the 1st coordinate of $\bar{z} \in \mathbb{R}^n = \mathbb{R}^{r_1} \times \mathbb{C}^{r_2}$.

We now wish to prove the following:

Theorem 6.4.3. *The fundamental domain $X_{K,\bar{u}}$ is a cone in $\mathbb{R}^n$, such that the region $X_{K,\bar{u}} \cap \{\bar{z} \in \mathbb{R}^n : N(\bar{z}) \leqslant 1\}$ is bounded, with volume $\frac{2^{r_1}\pi^{r_2}R_K}{m_K}$, and for every $\bar{z} \in \mathbb{R}^n$, with $N(\bar{z}) \neq 0$, there is a unique $u \in \mathscr{O}_K^*$ such that $\varphi(u) \cdot \bar{z} \in X_{K,\bar{u}}$.*

Proof. First, we verify that $X_{K,\bar{u}}$ is a cone; i.e., that if $\bar{z} \in X_{K,\bar{u}}$, that for any $t \in \mathbb{R}_{>0}$, we have $t\bar{z} \in X_{K,\bar{u}}$. We certainly have that conditions (i) and (iii) for $X_{K,\bar{u}}$ are preserved under multiplication by a positive real number. For (ii), note that from the definition that $\chi(t\bar{z}) = (\log t + \log |z_1|, \ldots, \log t + \log |z_{r_1}|, 2\log t + \log |z_{r_1+1}|, \ldots 2\log t + \log |z_{r_1+r_2}|) = (\log t)\bar{x}_0 + \chi(\bar{z})$. Thus, only the $\bar{x}_0$ coefficient is affected by multiplication by t, and condition (ii) is preserved as well.

We next verify the last assertion. Let U denote the subgroup of $\mathscr{O}_K$ generated by $u_1, \ldots, u_{r_1+r_2-1}$. Then immediately from the definitions, we see that for any $\bar{z} \in \mathbb{R}^n$ with $N(\bar{z}) \neq 0$ (which is the condition for $\chi(\bar{z})$ to be defined), , there is a unique element $u \in U$ such that $u \cdot \bar{z}$ satisfies condition (ii) of $X_{K,\bar{u}}$. Now, write μ_K for the group of roots of unity in K, so that $\mathscr{O}_K^* \cong \mu_K \times U$; we note that multiplication by elements of μ_K will not afffect condition (ii). In the case that $r_1 > 0$, we have $\mu_K = \pm 1$, and we want $z_1 > 0$; otherwise, $z_1 \in \mathbb{C}$, and we want $0 \leqslant \arg z_1 < 2\pi/m_K$; either way, it is clear that for any $\bar{z}' \in \mathbb{R}^n$ with $N(\bar{z}') \neq 0$, we have a unique $u_0 \in \mu_K$ with $u_0 \cdot \bar{z}'$ satisfying condition (iii). Thus, taking $u_0 \cdot u \cdot \bar{z}$, we obtain a point of $X_{K,\bar{u}}$. And we see that u is uniquely determined for $u_0 \cdot u \cdot \bar{z}$ to satisfy condition (ii), and then u_0 is uniquely determined to satisfy condition (iii), so since $\mathscr{O}_K^* \cong \mu_K \times U$, we obtain the desired uniqueness statement, as well.

What remains is the statement on boundedness and volume of $S := X_{K,\bar{u}} \cap \{\bar{z} : N(\bar{z}) \leqslant 1\}$. Note that the boundedness is not a triviality, as $\{\bar{z} : N(\bar{z}) \leqslant 1\}$ is not in fact bounded. The trick is to consider the subset $S_0 := X_{K,\bar{u}} \cap \{\bar{z} : N(\bar{z}) = 1\}$; it is clear that $S = \{tS_0 : 0 < t \leqslant 1\}$, so if we show that S_0 is bounded, it follows that S is bounded. But note that $\chi(S_0) \subseteq H$, and in fact is equal to the fundamental parallelepiped for $u_1, \ldots, u_{r_1+r_2-1}$, which is bounded in $\mathbb{R}^{r_1+r_2-1}$, and as we observed in the proof of the unit theorem, a subset of $\mathbb{R}^n$ whose image under χ is bounded must itself be bounded. According to our established modus operandi, we do not carry out the volume computation, which is another exercise in multivariable integration with polar coordinates; see [**7**, pp. 134-135], or [**1**, pp. 317-320] for a readable presentation. □

The theorem allows us to write

$$\zeta_{K,C}(s) = N(I_0)^s \sum_{\bar{z} \in \varphi(I_0) \cap X_{K,\bar{u}}} \frac{1}{N(\bar{z})^s}.$$

6.5. The zeta function

To complete the proof of the analytic class number formula, we now prove the following general theorem:

THEOREM 6.5.1. *Let $X \subseteq \mathbb{R}^n$ be any cone containing only points $\bar{z}$ with $N(\bar{z}) \neq 0$, and such that $X \cap \{\bar{z} \in \mathbb{R}^n : N(\bar{z}) \leqslant 1\}$ is bounded and has some volume v_X. Let $L \subseteq \mathbb{R}^n$ be a lattice of full rank with volume v_L. Define*

$$\zeta_{X,L}(s) = \sum_{\bar{z} \in X \cap L} \frac{1}{N(\bar{z})^s}.$$

Then this sum converges for $s > 1$, and we have

$$\lim_{s \to 1^+} (s-1)\zeta_{X,L}(s) = \frac{v_X}{v_L}.$$

PROOF. Denote by S the region $X \cap \{\bar{z} : N(\bar{z}) \leqslant 1\}$. For any $t \in \mathbb{R}_{t>0}$, denote by L_t the lattice $\frac{1}{t}L$, which has volume $\frac{1}{t^n}v_L$. Denote also by $P(t)$ the number of points of $L_t \cap S$. Approximating $\operatorname{vol} S$ by looking at the parallelepipeds of L_t inside S and taking the limit as $t \to \infty$, we see

$$v_X = \operatorname{vol} S = \lim_{t \to \infty} P(t) \frac{v_L}{t^n} = v_L \lim_{t \to \infty} \frac{P(t)}{t^n}.$$

Observe that $P(t)$ is also the number of points of $L \cap tS$, and $tS = X \cap \{\bar{z} : N(\bar{z}) \leqslant t^n\}$, in particular, the number of points of $X \cap L$ with bounded norm is finite, so the norms of points of $X \cap L$ are discrete. We can thus arrange the points of $L \cap X$ by non-decreasing norm, so that we have $L \cap X = \{x_1, x_2, \dots\}$ with $N(x_i) \leqslant N(x_{i+1})$ for all i. We claim that

$$\lim_{k \to \infty} \frac{k}{N(x_k)} = \frac{v_X}{v_L}.$$

By the above (replacing t^n with t), it suffices to prove that

$$\lim_{k \to \infty} \frac{k}{N(x_k)} = \lim_{t \to \infty} \frac{\#\{\bar{z} \in L \cap X : N(\bar{z}) \leqslant t\}}{t}.$$

For each k, and any $\epsilon > 0$, we see that

$$\#\{\bar{z} \in L : N(\bar{z}) \leqslant N(x_k) - \epsilon\} < k \leqslant \#\{\bar{z} \in L : N(\bar{z}) \leqslant N(x_k)\},$$

so

$$\frac{\#\{\bar{z} \in L \cap X : N(\bar{z}) \leqslant N(x_k) - \epsilon\}}{N(x_k) - \epsilon} \frac{N(x_k) - \epsilon}{N(x_k)} < \frac{k}{N(x_k)} \leqslant \frac{\#\{\bar{z} \in L \cap X : N(\bar{z}) \leqslant N(x_k)\}}{N(x_k)}.$$

Since the $N(x_k)$ go to infinity, we have $\lim_{k \to \infty} \frac{N(x_k) - \epsilon}{N(x_k)} = 1$, so taking limits and considering $t = N(x_k)$ and $t = N(x_k) - \epsilon$ gives the desired statement.

We now take as well-known that $\zeta_{\mathbb{Q}}(s)=\sum_{n=1}^{\infty}\frac{1}{n^s}$ converges for $s>1$ (this follows trivially from the integral test for convergence). We then have $\zeta_{X,L}(s)=\sum_{k=1}^{\infty}\frac{1}{N(x_k)^s}$, so it follows that $\zeta_{X,L}(s)$ converges for $s>1$, since

$$\lim_{k\to\infty}\frac{k^s}{N(x_k)^s}=\left(\frac{v_X}{v_L}\right)^s,$$

which is non-zero.

For any $\epsilon>0$, from the same equation we also see that for k_0 sufficiently large, for all $k\geqslant k_0$ we have

$$\left(\frac{v_X}{v_L}-\epsilon\right)\frac{1}{k}<\frac{1}{N(x_k)}<\left(\frac{v_X}{v_L}+\epsilon\right)\frac{1}{k}.$$

Thus, for $s>1$ we have

$$\left(\frac{v_X}{v_L}-\epsilon\right)^s\sum_{k=k_0}^{\infty}\frac{1}{k^s}<\sum_{k=k_0}^{\infty}\frac{1}{N(x_k)^s}<\left(\frac{v_X}{v_L}+\epsilon\right)^s\sum_{k=k_0}^{\infty}\frac{1}{k^s}.$$

Multiplying by $(s-1)$ and letting s go to 1 from above, we note that the middle sum goes to $\lim_{s\to 1^+}(s-1)\zeta_{X,L}(s)$, since $\lim_{s\to 1^+}(s-1)\sum_{k=1}^{k_0-1}\frac{1}{N(k)^s}=0$, and similarly, the sums on the left and right go to $\lim_{s\to 1^+}(s-1)\zeta_{\mathbb{Q}}(s)$. We thus obtain

$$\left(\frac{v_X}{v_L}-\epsilon\right)\lim_{s\to 1^+}(s-1)\zeta_{\mathbb{Q}}(s)\leqslant\liminf_{s\to 1^+}(s-1)\zeta_{X,L}(s)$$
$$\leqslant\limsup_{s\to 1^+}(s-1)\zeta_{X,L}(s)\leqslant\left(\frac{v_X}{v_L}+\epsilon\right)\lim_{s\to 1^+}(s-1)\zeta_{\mathbb{Q}}(s),$$

so if we show that $\lim_{s\to 1^+}(s-1)\zeta_{\mathbb{Q}}(s)=1$, we conclude the desired statement for $\zeta_{X,L}(s)$.

But this also follows from a slightly refined version of the integral test: we note that

$$\frac{1}{s-1}=\int_1^{\infty}\frac{dx}{x^s}<\zeta_{\mathbb{Q}}(s)<1+\int_1^{\infty}\frac{dx}{x^s}=1+\frac{1}{s-1},$$

and multiplying by $(s-1)$ and taking the limit as s goes to 1 from above gives the desired statement. □

REMARK 6.5.2. In fact, this proof does not use in an essential way the definition of the function $N(\bar{z})$, and goes through if $N(\bar{z})$ is replaced by any function $F:X\to\mathbb{R}_{>0}$ such that $F(t\bar{z})=t^nF(\bar{z})$ for any $\bar{z}\in X$ and $t\in\mathbb{R}_{>0}$, and such that $X\cap\{\bar{z}:F(\bar{z})\leqslant 1\}$ is bounded with volume $v_{X,F}$; in the formula of the theorem, we replace v_X by $v_{X,F}$.

As noted earlier, the previous two theorems together complete the proof of the analytic class number formula.

REMARK 6.5.3. We will discuss several applications of the analytic class number formula, but we also remark that it is the template for the famous conjecture of Birch and Swinnerton-Dyer, when one replaces the zeta function of a number field by the L-function of an elliptic curve.

6.6. Exercises

Exercise 6.1. Given a sequence of $a_n \in \mathbb{C}$ such that there exists $C \in \mathbb{R}$ with $|\sum_{n=1}^{m} a_n| \leqslant C$ for all m, then the sum $\sum_{n=1}^{\infty} \frac{a_n}{n^s}$ converges for all $s > 0$. In fact, for any $\epsilon > 0$, we have uniform convergence for $s \geqslant \epsilon$, so the sum is continuous for $s > 0$.

Hint: given $\epsilon > 0, \delta > 0$, choose n_0 such that for all $n \geqslant n_0$, we have $\frac{1}{n^\epsilon} < \delta$. Use telescoping sums to show that for any $n_2 > n_1 \geqslant n_0$, and $s \geqslant \epsilon$, we have $|\sum_{k=n_1}^{n_2} \frac{a_k}{k^s}| < 2C\delta$.

Exercise 6.2. Consider the case of $K = \mathbb{Q}(i)/\mathbb{Q}$. Using the Euler products, give a series expression for $s > 1$ for $\zeta_K(s)/\zeta(s)$, and then use the previous exercise to give an expression for $\lim_{s \to 1^+}(1-s)\zeta_K(s)$. Compare to the series expression for arctan and use the analytic class number formula to give an alternate proof that the class number of K is 1.

Notes

The contents of this chapter were taken from Chapter 5 of Borevich and Shafarevich [**1**].

CHAPTER 7

Dirichlet L-series

We will now begin a study of the case of subfields of cyclotomic fields, with the goal of concrete applications of the analytic class number formula to both the cyclotomic and quadratic case (and some elementary consequences of each). The first step will be to use the Euler product to give a different description of the zeta function in this situation. We will need to develop some background on Dirichlet characters and L-functions.

7.1. Dirichlet characters

Definition 7.1.1. A **Dirichlet character** is a multiplicative homomorphism $\chi : \mathbb{Z}/f\mathbb{Z}^* \to \mathbb{C}^*$, such that for any $d|f$, χ does not induce a map $\mathbb{Z}/d\mathbb{Z}^* \to \mathbb{C}^*$; then f is called the **conductor** of χ. We will also consider χ as a multiplicative map $\mathbb{Z} \to \mathbb{C}$ by defining

$$\chi(x) = \begin{cases} \chi(\bar{x}) : (x,f) = 1 \\ 0 : \text{otherwise} \end{cases}.$$

Remark 7.1.2. In fact, these characters are usually called “primitive”, and our convention of always working with primitive characters differs some other sources, which prefer to work with groups of characters of a fixed modulus. However, for our purposes, the formulas will be simpler with this convention. It does, however, mean that we will have to make a slightly more complicated convention of multiplication of characters.

Although $\mathbb{Z}/1\mathbb{Z}^*$ is not a very good concept, we set the following convention:

Notation 7.1.3. We denote by χ_1 the trivial character, which is the unique character considered to be of conductor 1, and corresponds to the constant map $\mathbb{Z} \to \mathbb{C}$ with value 1.

Given a Dirichlet character χ, we also denote by $\bar{\chi}$ the complex conjugate character defined by $\bar{\chi}(a) = \overline{\chi(a)}$.

Definition 7.1.4. We say that χ is a **Dirichlet character modulo** n if it is a Dirichlet character of conductor f with $f|n$; we then have that χ induces a multiplicative map $\mathbb{Z}/n\mathbb{Z}^* \to \mathbb{C}^*$. Conversely, given a multiplicative map $\mathbb{Z}/n\mathbb{Z}^* \to \mathbb{C}^*$, we obtain a unique Dirichlet character (which will be modulo n) as $\chi : \mathbb{Z}/f\mathbb{Z}^* \to \mathbb{C}^*$, where $f|n$ is the smallest integer such that the original map remains defined modulo f.

Definition 7.1.5. Given characters χ and ψ of conductors f_χ and f_ψ, we obtain maps $\chi, \psi : \mathbb{Z}/(\mathrm{lcm}(f_\chi, f_\psi))\mathbb{Z}^* \to \mathbb{C}^*$, so by taking products we obtain a map $\mathbb{Z}/(\mathrm{lcm}(f_\chi, f_\psi))\mathbb{Z}^* \to \mathbb{C}^*$. We thus let the **product character** $\chi\psi$ be the Dirichlet character associated to this map.

Easy Facts 7.1.6. We make the following observations about Dirichlet characters:

(i) If χ, ψ are Dirichlet characters modulo n, then $\chi\psi$ is a Dirichlet character modulo n.
(ii) For any Dirichlet character χ, we have $\chi\bar{\chi} = \chi_1$.
(iii) The set of Dirichlet characters modulo n form a group under multiplication.

From our part of view, the significance of Dirichlet characters is that they may also be viewed as characters on the Galois group of cyclotomic fields:

Lemma 7.1.7. $\mathrm{Gal}(\mathbb{Q}(\zeta_n)/\mathbb{Q}) = \mathbb{Z}/n\mathbb{Z}^*$ *(i.e., there is a canonical isomorphism). Thus, every Dirichlet character modulo n may be viewed as a homomorphism* $\mathrm{Gal}(\mathbb{Q}(\zeta_n)/\mathbb{Q}) \to \mathbb{C}^*$.

Proof. Although the group of nth roots of unity is non-canonically isomorphic to $\mathbb{Z}/n\mathbb{Z}$, and in particular, the choice of ζ_n is non-canonical, for any choice of ζ_n, we see that an automorphism of $\mathbb{Q}(\zeta_n)/\mathbb{Q}$ is uniquely determined by $\zeta_n \mapsto \zeta_n^a$ for $(a, n) = 1$. Moreover, since we then have $\zeta_n^i \mapsto (\zeta_n^i)^a$, we see that a is independent of the choice of ζ_n. □

Therefore, we can make the following definition:

Definition 7.1.8. Let G be a subgroup of the group of Dirichlet characters modulo n. We define the associated field $K_G \subseteq \mathbb{Q}(\zeta_n)$ to be the fixed field of $G' \subseteq \mathrm{Gal}(\mathbb{Q}(\zeta_n)/\mathbb{Q})$, where G' is the subgroup of elements in the kernel of every $\chi \in G$.

We can also go in the other direction:

Definition 7.1.9. Given $K \subseteq \mathbb{Q}(\zeta_n)$, we can define a group G_K of Dirichlet characters modulo n by taking all characters $\chi : \mathrm{Gal}(\mathbb{Q}(\zeta_n)/\mathbb{Q}) \to \mathbb{C}^*$ which are trivial when restricted to $\mathrm{Gal}(\mathbb{Q}(\zeta_n)/K) \subseteq \mathrm{Gal}(\mathbb{Q}(\zeta_n)/\mathbb{Q})$.

Basic theory of characters of finite abelian groups, together with basic Galois theory, give:

Proposition 7.1.10. *The maps $G \mapsto K_G$ and $K \mapsto G_K$ are inverse to one another, giving a natural (inclusion-preserving) bijection between groups of Dirichlet characters modulo n and subfields of $\mathbb{Q}(\zeta_n)$.*

For the purpose of examples, it will be useful to consider the following:

Definition 7.1.11. We say that a Dirichlet character χ is **odd** if $\chi(-1) = -1$; otherwise, $\chi(-1) = 1$, and we say χ is **even**.

We conclude by discussing three important special cases, which we will return to later:

Example 7.1.12. If G is the full group of Dirichlet characters modulo n, one easily checks that G' is trivial, so that $K_G = \mathbb{Q}(\zeta_n)/\mathbb{Q}$.

If G is the subgroup of even Dirichlet characters modulo n, then G' consists of the identity and complex conjugation, and $K_G = \mathbb{Q}(\zeta_n + \zeta_n^{-1})$.

Finally, if n is odd and square-free, and G consists only of χ_1 and χ_n, which we set to be the Jacobi symbol modulo n, we have that G' has index 2 in the group of Dirichlet characters modulo n, and $K_G = \mathbb{Q}(\sqrt{\chi_n(-1)n})$. In particular, K_G is real if and only if G consists of even characters.

7.2. Dirichlet L-series

Let χ be a Dirichlet character of conductor f: that is, a multiplicative map $(\mathbb{Z}/f\mathbb{Z})^* \to \mathbb{C}^*$ which does not remained well-defined mod d for any $d|f$, $d \neq f$. Recall that we also consider χ as a map $\mathbb{Z} \to \mathbb{C}$ by defining

$$\chi(x) = \begin{cases} \chi(\bar{x}) : (x, f) = 1 \\ 0 : \text{otherwise} \end{cases}.$$

We can then define:

Definition 7.2.1. The **Dirichlet L-series** associated to χ is

$$L(s, \chi) := \sum_{n=1}^{\infty} \frac{\chi(n)}{n^s}.$$

The most basic properties of $L(s, \chi)$ are the following.

Proposition 7.2.2. *Given a Dirichlet character χ, we have:*

(i) *The sum for $L(s, \chi)$ converges absolutely for $s > 1$.*

(ii) *There is a product expansion for $s > 1$*

$$L(s, \chi) = \prod_p \left(1 - \frac{\chi(p)}{p^s}\right)^{-1}.$$

Proof. (i) Since $|\frac{\chi(n)}{n^s}| = \begin{cases} \frac{1}{n^s} : (n, f) = 1 \\ 0 : \text{otherwise} \end{cases}$, the absolute convergence for $s > 1$ follows from the convergence of $\zeta_{\mathbb{Q}}(s)$ for $s > 1$.

(ii) Given the convergence for $s > 1$, the proof of the product expansion is identical to the case $\zeta_{\mathbb{Q}}(s)$. □

We also observe that $L(s, \chi_1) = \zeta_{\mathbb{Q}}(s)$. As we will see, for non-trivial characters, the L-series behave rather differently.

We will analyze these functions further, but first we want to demonstrate their relevance to us:

Theorem 7.2.3. *Given a number field $K \subseteq \mathbb{Q}(\zeta_n)$, let G_K be the corresponding group of Dirichlet characters modulo n. Then for $s > 1$ we have*

$$\zeta_K(s) = \prod_{\chi \in G_K} L(s, \chi).$$

The proof of the theorem will require a rather long detour, so we put it off for the moment and examine some consequences.

Remark 7.2.4. The theorem shows that the Riemann hypothesis for the Dedekind zeta function $\zeta_{\mathbb{Q}(\zeta_n)}(s)$ implies the Riemann hypothesis for Dirichlet L-series modulo n, since the zeroes of $\zeta_{\mathbb{Q}(\zeta_n)}$ will be the union of the zeroes of the $L(s, \chi)$.

Lemma 7.2.5. *For χ non-trivial, we have that $L(s, \chi)$ converges to a continuous function for $s > 0$. In particular, $L(1, \chi)$ is finite.*

Proof. We first observe that if f is the conductor of χ, then $|\sum_{i=1}^{n}\chi(i)| \leqslant f$ for all n. Indeed, if we choose $x \in (\mathbb{Z}/f\mathbb{Z})^*$ with $\chi(x) \neq 1$, we have

$$(1-\chi(x))(\sum_{i=k+1}^{k+f}\chi(i)) = (1-\chi(x))(\sum_{i\in(\mathbb{Z}/f\mathbb{Z})^*}\chi(i)) = (\sum_{i\in(\mathbb{Z}/f\mathbb{Z})^*}\chi(i)) - (\sum_{i\in(\mathbb{Z}/f\mathbb{Z})^*}\chi(xi)) = 0$$

for any k, so $\sum_{i=k+1}^{k+f}\chi(i) = 0$, and we find $\sum_{i=1}^{mf}\chi(i) = 0$ for all m. Since $|\chi(i)| = 1$, we obtain the desired bound. The statement of the lemma then follows from Exercise 6.1. □

We can therefore draw the following corollary from the theorem:

Corollary 7.2.6. *Given the above theorem, for any non-trivial Dirichlet character χ, we have*

$$L(1,\chi) \neq 0.$$

Proof. By the analytic class number formula, $\lim_{s\to 1^+}(s-1)\zeta_{\mathbb{Q}(\zeta_n)}(s) \neq 0$. If we apply the product formula of the theorem, and recall that $\lim_{s\to 1^+}(s-1)\zeta_{\mathbb{Q}}(s) = 1$, we have

$$0 \neq \lim_{s\to 1^+}(s-1)\zeta_{\mathbb{Q}(\zeta_n)}(s) = \prod_{\chi\neq\chi_1} L(1,\chi),$$

where the product is over all non-trivial Dirichlet characters modulo n, and each term on the right is finite by the lemma. We thus conclude that each term on the right is non-zero, as desired. □

Remark 7.2.7. In fact, this corollary is the basis of the standard proof of the Dirichlet theorem on primes in arithmetic progressions, and the theorem follows without too much additional work. However, rather than pursue this now, we will see how it follows, in the context of more general density theorems, from class field theory.

We will see that it is always possible to give explicit formulas for $L(1,\chi)$, so that for subfields of cyclotomic fields, the analytic class number formula always gives a computable expression.

7.3. Prime factorization and characters

We now return to the case of sub-cyclotomic fields, and use the theory of decomposition and inertia groups to relate groups of Dirichlet characters to factorization of prime ideals. The main theorem is the following:

Theorem 7.3.1. *Let G be a group of Dirichlet characters modulo n, and K_G the associated field. Fix a prime number p, and let $H_1 \subseteq H_2 \subseteq G$ be the groups of χ such that $\chi(p) = 1$ or $\chi(p) \neq 0$ respectively. Then $G/H_2 \cong I_{K_G,p} \subseteq \mathrm{Gal}(K_G/\mathbb{Q})$, and $G/H_1 \cong D_{K_G,p}$ (note that these are non-canonical isomorphisms).*

We will need the following lemma:

Lemma 7.3.2. *In the situation of the theorem, p is ramified in K if and only if $G/H_2 \neq (1)$.*

PROOF. $G/H_2 \neq (1)$ if and only if there exists $\chi \in G$ with $\chi(p) = 0$ if and only if there exists $\chi \in G$ of conductor f with $p|f$. Let us write $n = \prod_{i=1}^{k} p_i^{e_i}$; by Exercise 7.2, for any $\chi \in G$ of conductor f (necessarily, $f|n$), we have $\chi = \prod_{p_i} \chi_{p_i}$ with each χ_{p_i} of conductor $p_i^{e'_i}$, where $f = \prod_{i=1}^{k} p_i^{e'_i}$. In particular, $e'_i \leqslant e_i$ for each i, so we may consider χ_{p_i} to be a Dirichlet character modulo n, and we can define G_p to be the group of Dirichlet characters modulo p^{e_i} consisting of the χ_p as χ ranges over the characters G. We thus see that $G/H_2 \neq (1)$ if and only if $G_p \neq (1)$.

We now consider what it means for G_p to be trivial or not. Write $n = p^e m$ with $p \not| m$. We have that $K_{G_p} \subseteq \mathbb{Q}(\zeta_{p^e})$, so p is totally ramified in K_{G_p}, by Lemma 5.4.2 (iii). Thus, p is ramified in K_{G_p} if and only if $G_p \neq (1)$. The key idea is to consider the field $L = K_G(\zeta_m)$, the compositum of K_G with $\mathbb{Q}(\zeta_m)$. Then $L = K_{GG_m}$, where G_m denotes the full group of characters modulo m. It is easy to check that $GG_m = G_p G_m$, so in fact L is the compositum of K_{G_p} wih $K_{G_m} = \mathbb{Q}(\zeta_m)$. Now, p is unramified in $\mathbb{Q}(\zeta_m)$, so by the Proposition on unramified extensions from last time, p is ramified in L if and only if it is ramified in K_{G_p} if and only if $G/H_2 \neq (1)$. But L is also the compositum of K_G with $\mathbb{Q}(\zeta_m)$ by definition, so we see that p is ramified in K_G if and only if $G/H_2 \neq (1)$, as desired. □

Before completing the proof of the theorem, we extend our correspondence between groups of Dirichlet characters modulo n and subfields of $\mathbb{Q}(\zeta_n)$ to a slightly more general case. Suppose we are given a group G of Dirichlet characters modulo n, and the corresponding field $K_G \subseteq \mathbb{Q}(\zeta_n)$. We claim we have a correspondence directly between subgroups of G and subfields of K_G in the same way as before.

The main point is that any $\chi \in G$ will still give a map $\mathrm{Gal}(K_G/\mathbb{Q}) \to \mathbb{C}$: indeed, $\mathrm{Gal}(K_G/\mathbb{Q})$ is a quotient of $(\mathbb{Z}/n\mathbb{Z})^*$ by the group G' of elements in the kernel of every $\chi \in G$, so by definition any $\chi \in G$ is still defined on $\mathrm{Gal}(K_G/\mathbb{Q}) = (\mathbb{Z}/n\mathbb{Z})^*/G'$. Thus, given $H \subseteq G$, we can define H' to be the subgroup of $\mathrm{Gal}(K_G/\mathbb{Q})$ on which every $\chi \in H$ is trivial, and define K_H as a subfield of K_G in the same way as before, and it is easy to check that this agrees with our earlier definition of K_H as a subfield of $\mathbb{Q}(\zeta_n)$.

We now give the proof of the theorem.

PROOF OF THEOREM 7.3.1. It follows from the lemma that K_{H_2} is the maximal extension of $\mathbb{Q}$ inside K_G in which p is unramified. Thus, by Theorem 3.7.2, we have $\mathrm{Gal}(K_G/K_{H_2}) = I_{K_G,p}$. Under the correspondence between groups of Dirichlet characters modulo n and subfields of $\mathbb{Q}(\zeta_n)$, we obtain compatible isomorphisms $\mathrm{Gal}(K_G/\mathbb{Q}) \cong G$ and $\mathrm{Gal}(K_{H_2}/\mathbb{Q}) \cong H_2$, so $I_{K_G,p} = \mathrm{Gal}(K_G/K_{H_2}) \cong G/H_2$.

We now shift our focus to K_{H_2}, and show that K_{H_1} is the fixed field of the decomposition group $D_{K_{H_2},p}$. Let m be the least common multiple of the conductors of $\chi \in H_2$; we have by definition that $p \not| m$. We then have $K_{H_2} \subseteq \mathbb{Q}(\zeta_m)$, and p is unramified in either extension, so $D_{K_{H_2},p}$ is cyclic, generated by $\mathrm{Fr}(K_{H_2}, p)$. Thus, if we show that H_1 is precisely the subgroup of H_2 on which $\mathrm{Fr}(K_{H_2}, p)$ is trivial, we obtain the desired statement under the correspondence between groups of characters and subfields of $\mathbb{Q}(\zeta_m)$. We know that $\mathrm{Fr}(\mathbb{Q}(\zeta_m), p)$ is just the class of p in $\mathrm{Gal}(\mathbb{Q}(\zeta_m)/\mathbb{Q}) = (\mathbb{Z}/m\mathbb{Z})^*$, and $\mathrm{Gal}(K_{H_2}/\mathbb{Q})$ is an appropriate quotient of $(\mathbb{Z}/m\mathbb{Z})^*$. In fact, $\mathrm{Fr}(K_{H_2}, p)$ is also the class of p in this quotient, because it is clear from the definition that the restriction of $\mathrm{Fr}(\mathbb{Q}(\zeta_m), p)$ to $\mathrm{Gal}(K_{H_2}/\mathbb{Q})$ must be $\mathrm{Fr}(K_{H_2}, p)$. Thus $\chi(\mathrm{Fr}(K_{H_2}), p) = \chi(p)$, and by definition H_1 is the subgroup of characters trivial on $\mathrm{Fr}(K_{H_2}, p)$, and K_{H_1} is the fixed field of $D_{K_{H_2},p}$ inside K_{H_2}, as desired.

It finally remains to conclude from this that K_{H_1} is the fixed field of $D_{K_G,p}$ inside K_G. But this is a general fact for abelian extensions: the decomposition group of the inertia field is simply the quotient of the decomposition group of the original field, so they have the same fixed field. Thus we have $D_{K_G,p} = \mathrm{Gal}(K_G/K_{H_1}) \cong G/H_1$, as desired. □

7.4. Factorization of zeta functions

We are now ready to prove Theorem 7.2.3, expressing Dedekind zeta functions of sub-cyclotomic fields as products of Dirichlet L-series. We recall the statement:

THEOREM 7.4.1. *Given a number field $K \subseteq \mathbb{Q}(\zeta_n)$, let G_K be the corresponding group of Dirichlet characters modulo n. Then for $s > 1$ we have*

$$\zeta_K(s) = \prod_{\chi \in G_K} L(s, \chi).$$

PROOF. Since we have absolute convergence for $s > 1$, we can freely rearrange terms in the product expansions. Recall the Euler product expansion

$$\zeta_K(s) = \prod_{\mathfrak{p} \in \mathscr{O}_K} \left(1 - \frac{1}{N(\mathfrak{p})^s}\right)^{-1};$$

we will compare the product expansions for each integer prime p. If we write $p\mathscr{O}_K = \prod_{i=1}^{g} \mathfrak{p}_i^e$, with $N(\mathfrak{p}_i) = p^f$, then the Euler factors for $\zeta_K(s)$ corresponding to p are simply $(1 - \frac{1}{p^{fs}})^{-g}$. Recall that we had the Euler product

$$L(s, \chi) = \prod_{p} \left(1 - \frac{\chi(p)}{p^s}\right)^{-1}$$

for $L(s, \chi)$; the terms corresponding to p are then $\prod_{\chi \in G}(1 - \frac{\chi(p)}{p^s})^{-1}$.

In the notation of the theorem from last time, only the $\chi \in H_2$ will contribute to the product. We know moreover from the theorem that H_2/H_1 is cyclic of order f, and H_1 has order g, so choose $\chi_0 \in H_2$ with image generating H_2/H_1; it is clear that $\chi_0(p) = \zeta_f$ is some primitive fth root of unity. We then have

$$\prod_{i=0}^{f-1}(1 - \frac{\chi_0^i(p)}{p^s})^{-1} = \prod_{i=0}^{f-1}(1 - \frac{\zeta_f^i}{p^s})^{-1} = (1 - \frac{1}{p^{fs}})^{-1}.$$

But taking the product over all $\chi \in G$ is the same as taking it over all $\chi \in H_2$ is the same as taking it g times over the powers of χ_0, since $\chi(p) = 1$ for any $\chi \in H_1$, and χ_0 generates H_2/H_1. So the Euler factor at p for the product of L-series is $(1 - \frac{1}{p^{fs}})^{-g}$, just as for $\zeta_K(s)$, completing the proof of the theorem. □

In particular, the factorization theorem together with the analytic class number formula mean that we will be interested in evaluating $L(1, \chi)$ for non-trivial Dirichlet characters χ. Before we do this, we briefly discuss another application.

7.5. Functional equations and the conductor-discriminant formula

For this application, we will assume the functional equations for $\zeta_K(s)$ and $L(s,\chi)$. These are:

Theorem 7.5.1. *Given a number field K, the function $\zeta_K(s)$ can be analytically extended to the entire complex plane except for a simple pole at $s=1$, and satisfies the functional equation:*

$$\zeta_K(s)\Gamma(\frac{s}{2})^{r_1}\Gamma(s)^{r_2} = \zeta_K(1-s)\Gamma(\frac{1-s}{2})^{r_1}\Gamma(1-s)^{r_2}(4^{-r_2}\pi^{-n}|D_K|)^{\frac{1}{2}-s}.$$

Given a non-trivial Dirichlet character χ, the function $L(s,\chi)$ can be analytically extended to the entire complex plane, and satisfies the functional equation:

$$L(s,\chi)\Gamma(\frac{s}{2}) = L(1-s,\bar{\chi})\Gamma(\frac{1-s}{2})\frac{\tau(\chi)}{\sqrt{f}}(\frac{f}{\pi})^{\frac{1}{2}-s}$$

for χ even, and

$$L(s,\chi)\Gamma(\frac{s+1}{2}) = L(1-s,\bar{\chi})\Gamma(1-\frac{s}{2})\frac{\tau(\chi)}{i\sqrt{f}}(\frac{f}{\pi})^{\frac{1}{2}-s}$$

for χ odd, where $\tau(\chi) = \sum_{k=1}^{f}\chi(k)\zeta_f^k$ with $\zeta_f = e^{2\pi i/f}$ is the Gauss sum for χ (note that $\tau(\chi_1)=1$).

From the functional equations and the factorization theorem, we can conclude two formulas; the first provides further insight into the relationship between groups of characters and sub-cyclotomic fields, while the second will be useful in studying the quadratic case.

Theorem 7.5.2. *Let G be a group of Dirichlet characters modulo n, and $K_G \subseteq \mathbb{Q}(\zeta_n)$ the associated field. For each $\chi \in G$, denote by f_χ the conductor. Then we have:*

$$\prod_{\chi\in G} f_\chi = (-1)^{r_2}D_{K_G}, \tag{7.5.2.1}$$

and

$$\prod_{\chi\in G} \tau(\chi) = i^{r_2}\sqrt{|D_K|}. \tag{7.5.2.2}$$

Note that because K_G is Galois, either $r_1 = |G| = [K_G:\mathbb{Q}]$ and $r_2=0$ or $r_1=0$ and $r_2=|G|/2$.

Proof. We first consider the case that $r_2=0$, so K_G is real. Then all the characters in G are even, and multiplying the functional equations for $L(s,\chi)$ over all $\chi\in G$, rewriting in terms of $\zeta_{K_G}(s)$, and applying the functional equation of $\zeta_{K_G}(s)$, we obtain:

$$\prod_{\chi\in G}\frac{\tau(\chi)}{f_\chi^s} = \frac{\sqrt{|D_{K_G}|}}{|D_{K_G}|^s},$$

for all s. Setting $s=0$ we obtain $\prod_{\chi\in G}\tau(\chi) = \sqrt{|D_{K_G}|}$, and then cancelling these terms we find $\prod_{\chi\in G} f_\chi = |D_{K_G}|$. Recall from Lemma 4.3.3 that $D_{K_G}>0$ if and only if r_2 is even, so in this case we have $|D_{K_G}| = D_{K_G}$, and we shown the desired formulas.

In the case $r_1 = 0$, observe that exactly half the characters of G must be odd, and half even. The argument is then the same, except that we have to apply the identity

$$\Gamma\left(\frac{s}{2}\right)\Gamma\left(\frac{s+1}{2}\right) = 2^{1-s}\sqrt{\pi}\Gamma(s),$$

and obtain $\prod_{\chi\in G}\tau(\chi) = i^{r_2}\sqrt{|D_{K_G}|}$ and $\prod_{\chi\in G} f_\chi = |D_{K_G}|$. Because we know that $|D_{K_G}| = (-1)^{r_2}D_{K_G}$, this gives the desired formulas. □

Note that the first formula implies, in particular, that the quadratic fields of discriminant D in fact correspond to characters of conductor $|D|$, as suggested by Exercise 7.4. From the analytic class number formula, we find:

Corollary 7.5.3. *If K is a quadratic field corresponding to a quadratic character χ, we have:*

$$h_K = \begin{cases} \frac{\sqrt{D_K}}{2\log\epsilon}L(1,\chi) : D_K > 0 \\ \frac{m_K\sqrt{|D_K|}}{2\pi}L(1,\chi) : D_K < 0 \end{cases},$$

where ϵ generates the units of $\mathscr{O}_K$ modulo ± 1, and is uniquely determined by the condition $\epsilon > 1$.

Recall that $m_K := \#\{$roots of unity $\in K\}$ is 2 unless $K = \mathbb{Q}(\sqrt{-1})$ or $\mathbb{Q}(\sqrt{-3})$, where it is 4 or 6 respectively.

Proof. For the case $D_K > 0$, we know that $m_K = 2$, so using the analytic class number formula and the factorization formula in terms of L-series, the only point is to note that $\log\epsilon = R_K$: by definition, we have $R_K = |\log|u||$ where u generates $\mathscr{O}_K^*$ modulo ± 1; u is determined up to $\pm u^{\pm 1}$, so we can always choose $\epsilon > 1$, and for this choice no absolute values are required.

For the case $D_K < 0$, the only point is to recall that R_K is defined to be 1, as the only units are roots of unity. □

7.6. The value of $L(1,\chi)$

We now begin the final portion of our application of the analytic class number formula by evaluating $L(1,\chi)$ for $\chi \neq \chi_1$. The theorem we wish to prove is the following:

Theorem 7.6.1. *Let χ be a Dirichlet character of conductor $f > 1$. Then for χ even, we have*

$$\begin{aligned} L(1,\chi) &= -\tfrac{\tau(\chi)}{f}\textstyle\sum_{k=1}^{f-1}\bar{\chi}(k)\log|1-\zeta_f^k| \\ &= -\tfrac{\tau(\chi)}{f}\textstyle\sum_{k=1}^{f-1}\bar{\chi}(k)\log\sin(\tfrac{\pi k}{f}). \end{aligned}$$

For χ odd, we have

$$L(1,\chi) = \frac{\pi i\tau(\chi)}{f^2}\sum_{k=1}^{f-1}\bar{\chi}(k)k.$$

Plugging this result into the previous one, we obtain:

Corollary 7.6.2. *If K is a quadratic field corresponding to a quadratic character χ, we have:*

$$h_K = \begin{cases} -\frac{1}{\log\epsilon}\sum_{1\leqslant k<D_K/2}\chi(k)\log\sin(\frac{\pi k}{D_K}) : D_K > 0 \\ -\frac{m_K}{2|D_K|}\sum_{k=1}^{|D_K|-1}\chi(k)k : D_K < 0 \end{cases},$$

with ϵ as in the previous corollary.

PROOF. Substituting the formulas of the theorem in the previous corollary, we have $\chi = \bar{\chi}$ since χ is quadratic. The main point is to recall that $f_\chi = |D_K|$, and to note that by (7.5.2.2), since $\tau(\chi_1) = 1$, we find

$$\tau(\chi) = \begin{cases} \sqrt{D_K} : D_K > 0 \\ i\sqrt{|D_K|} : D_K < 0 \end{cases}.$$

In addition, since $D_K > 0$ if and only if $\chi(-1) = 1$, we note that in this case $\chi(k) = \chi(D_K - k)$ for all k, so we may sum over $k < D_K$ and divide by 2. □

We next show that $L(1,\chi)$ may be written as the following finite sum:

PROPOSITION 7.6.3. *For $\chi \neq \chi_1$, we have*

$$L(1,\chi) = -\frac{1}{f}\sum_{k=1}^{f-1} \tau_k(\chi)\log(1-\zeta_f^{-k}),$$

where $\tau_k(\chi) := \sum_{j=1}^{f} \chi(j)\zeta_f^{jk}$.

PROOF. We first observe that we may write

$$\begin{aligned} L(s,\chi) &= \textstyle\sum_{(j,f)=1}\left(\chi(j)\sum_{n\equiv j \pmod f}\frac{1}{n^s}\right) \\ &= \textstyle\sum_{(j,f)=1}\left(\chi(j)\sum_{n=1}^{\infty}\frac{\delta_{j,n}}{n^s}\right), \end{aligned}$$

where $\delta_{j,n} = \begin{cases} 1 : n \equiv j \pmod f \\ 0 : \text{ otherwise} \end{cases}$. Now, observe that we can write

$$\delta_{j,n} = \frac{1}{f}\sum_{k=0}^{f-1}\zeta_f^{(j-n)k}.$$

Thus, we obtain:

$$\begin{aligned} L(s,\chi) &= \textstyle\sum_{(j,f)=1}\left(\chi(j)\sum_{n=1}^{\infty}\frac{1}{f}\sum_{k=0}^{f-1}\zeta_f^{(j-n)k}\frac{1}{n^s}\right) \\ &= \textstyle\frac{1}{f}\sum_{k=0}^{f-1}\left(\sum_{(j,f)=1}\chi(j)\zeta_f^{kj}\right)\sum_{n=1}^{\infty}\frac{\zeta_f^{-nk}}{n^s} \\ &= \textstyle\frac{1}{f}\sum_{k=0}^{f-1}\tau_k(\chi)\sum_{n=1}^{\infty}\frac{\zeta_f^{-nk}}{n^s} \\ &= \textstyle\frac{1}{f}\sum_{k=1}^{f-1}\tau_k(\chi)\sum_{n=1}^{\infty}\frac{\zeta_f^{-nk}}{n^s}, \end{aligned}$$

with the last equality following from the observation that $\tau_0(\chi) = \sum_{j=1}^{f}\chi(j) = 0$.

Now, observe that $|\sum_{n=1}^{N}\zeta_f^{-nk}|$ remains bounded by f for any N, so by Exercise 6.1 we find that $\sum_{n=1}^{\infty}\frac{\zeta_f^{-nk}}{n^s}$ converges to a continuous function for any $s > 0$, and in particular we may evaluate the limit as s approaches 1 from above simply by setting $s = 1$. We thus conclude

$$L(1,\chi) = \frac{1}{f}\sum_{k=1}^{f-1}\tau_k(\chi)\sum_{n=1}^{\infty}\frac{\zeta_f^{-nk}}{n}.$$

So it remains to see that $\sum_{n=1}^{\infty} \frac{\zeta_f^{-nk}}{n} = -\log(1-\zeta_f^{-k})$. This follows almost immediately from the power series expansion $\log(1-x) = -\sum_{n=1}^{\infty} \frac{x^n}{n}$, but since the radius of convergence for this series is 1, and $|\zeta_f^{-k}| = 1$, we also need to use that we already know that the sum $\sum_{n=1}^{\infty} \frac{\zeta_f^{-nk}}{n}$ converges, so we may apply Abel's theorem. This says (after some rephrasing) that if we obtain a convergent series by evaluating a power series at a point on the boundary of the disc of convergence, then the sum of this series is the limit of the sum of the power series. Thus we obtain the desired identity. □

We next show that this may be rewritten again as follows:

Proposition 7.6.4. *For $\chi \neq \chi_1$, we have*

$$L(1,\chi) = -\frac{\tau(\chi)}{f}\sum_{k=1}^{f-1} \bar{\chi}(k)\log(1-\zeta_f^{-k}).$$

Proof. This follows easily from the previous proposition by two observations. First, that $\tau_k(\chi) = 0$ if $(k,f) > 1$, and second, that $\tau_k(\chi) = \chi(k)^{-1}\tau(\chi) = \bar{\chi}(k)\tau(\chi)$ if $(k,f) = 1$.

The first observation depends on the fact that if $d|f$ and $d \neq f, 1$, there exists m with $(m,f) = 1$ and $m \equiv 1 \pmod{d}$ and such that $\chi(m) \neq 1$. Indeed, if not we can check directly that χ is in fact defined modulo d: given a, b with $(a,f) = (a,f) = 1$ and $a \equiv b \pmod{d}$, then $a \equiv mb \pmod{f}$ for some m, and we see that $m \equiv 1 \pmod{d}$, so $\chi(a) = \chi(m)\chi(b) = \chi(b)$, as desired.

Given this fact, if $\frac{f}{d} = (k,f)$, we choose $m \equiv 1 \pmod{d}$ with $(m,f) = 1$ and such that $\chi(m) \neq 1$; then multiplication by m permutes the elements of $(\mathbb{Z}/f\mathbb{Z})^*$, and we have

$$\chi(m)\tau_k(\chi) = \chi(m)\sum_{(j,f)=1} \chi(j)\zeta_f^{jk} = \sum_{(j,f)=1} \chi(mj)\zeta_f^{jk} = \sum_{(j,f)=1} \chi(mj)\zeta_f^{mjk} = \tau_k(\chi),$$

with $\zeta_f^{jk} = \zeta_f^{mjk}$ because $m \equiv 1 \pmod{d}$ and $(\zeta_f^k)^d = 1$, since $\frac{f}{d}$ divides k. Thus we conclude $\tau_k(\chi) = 0$.

The second observation is simpler: when $(k,f) = 1$, multiplication by k permutes the elements of $(\mathbb{Z}/f\mathbb{Z})^*$, so

$$\chi(k)\tau_k(\chi) = \chi(k)\sum_{(j,f)=1} \chi(j)\zeta_f^{jk} = \sum_{(j,f)=1} \chi(jk)\zeta_f^{jk} = \tau(\chi).$$

Hence we conclude the formula of the proposition from the previous one. □

Finally, rewriting once more and considering the even and odd cases separately, we can finish the proof of the theorem:

Proof of the theorem. We wish to analyze the sum

$$S_\chi := \sum_{k=1}^{f-1} \bar{\chi}(k)\log(1-\zeta_f^{-k}).$$

We first note that we have

$$1-\zeta_f^{-k} = 2\sin\left(\frac{\pi k}{f}\right)\left(e^{i\left(\frac{\pi}{2}-\frac{\pi k}{f}\right)}\right),$$

so if $0 < k < f$, we have

$$\log|1 - \zeta_f^{-k}| = \log 2 \sin\left(\frac{\pi k}{f}\right),$$

and

$$\log(1 - \zeta_f^{-k}) = \log|1 - \zeta_f^{-k}| + i\pi\left(\frac{1}{2} - \frac{k}{f}\right).$$

We then also find

$$\log(1 - \zeta_f^{k}) = \log|1 - \zeta_f^{k}| - i\pi\left(\frac{1}{2} - \frac{k}{f}\right).$$

Let us consider the case that χ (and hence $\bar{\chi}$) is even. Then by replacing k with $-k$, we find $S_\chi = \sum_{k=1}^{f-1} \bar{\chi}(k) \log(1 - \zeta_f^k)$, so we have

$$2S_\chi = \sum_{k=1}^{f-1} \bar{\chi}(k)(\log(1 - \zeta_f^{-k}) + \log(1 - \zeta_f^k)),$$

so by the above formulas,

$$\begin{aligned} S_\chi \qquad & = \textstyle\sum_{k=1}^{f-1} \bar{\chi}(k) \log|1 - \zeta_f^k| \qquad = \displaystyle\sum_{k=1}^{f-1} \bar{\chi}(k) \log 2 \sin(\frac{\pi k}{f}) \\ & = \textstyle\sum_{k=1}^{f-1} \bar{\chi}(k) \log 2 + \sum_{k=1}^{f-1} \bar{\chi}(k) \log \sin(\frac{\pi k}{f}) \\ & = \log 2 \tau_0(\bar{\chi}) + \textstyle\sum_{k=1}^{f-1} \bar{\chi}(k) \log \sin(\frac{\pi k}{f}) \\ & = \textstyle\sum_{k=1}^{f-1} \bar{\chi}(k) \log \sin(\frac{\pi k}{f}). \end{aligned}$$

Thus, we obtain the desired formula for even characters.

The odd case proceeds similarly: $S_\chi = -\sum_{k=1}^{f-1} \bar{\chi}(k) \log(1 - \zeta_f^k)$, so we have

$$2S_\chi = \sum_{k=1}^{f-1} \bar{\chi}(k)(\log(1 - \zeta_f^{-k}) - \log(1 - \zeta_f^k)),$$

and we find

$$\begin{aligned} S_\chi \quad & = \textstyle\sum_{k=1}^{f-1} \bar{\chi}(k) \pi i \left(\frac{1}{2} - \frac{k}{f}\right) \\ & = \tfrac{\pi i}{2} \tau_0(\bar{\chi}) - \tfrac{\pi i}{f} \textstyle\sum_{k=1}^{f-1} \bar{\chi}(k) k \\ & = -\tfrac{\pi i}{f} \textstyle\sum_{k=1}^{f-1} \bar{\chi}(k) k. \end{aligned}$$

This completes the proof of the theorem. □

Ultimately, in addition to the application to quadratic fields, we will apply this result to recover Kummer's criterion for regular primes in terms of Bernoulli numbers. However, we put this off until we discuss class field theory, which will expedite the middle steps of the proof.

7.7. Exercises

Exercise 7.1. Show that if χ, ψ are Dirichlet characters of conductors f_χ, f_ψ, and $(f_\chi, f_\psi) = 1$, then the conductor of $\chi \cdot \psi$ is $f_\chi f_\psi$.

Exercise 7.2. Show that if $f = \prod_{i=1}^n p_i^{e_i}$, and χ is a character of conductor f, then χ may be written uniquely as $\chi_{p_1} \cdots \chi_{p_n}$, where χ_{p_i} is a character of conductor $p_i^{e_i}$.

Exercise 7.3. If G_1, G_2 are two groups of Dirichlet characters, then $K_{G_1 G_2} = K_{G_1} K_{G_2}$.

Exercise 7.4. A Dirichlet character χ is called **quadratic** if its square is χ_1, or equivalently, if χ takes values in ± 1. Show explicitly that there is a one-to-one correspondence between quadratic Dirichlet characters and quadratic fields as follows (note that you are not required relate this to the more general equivalence described in §7.1 between groups of Dirichlet characters modulo n and subfields of $\mathbb{Q}(\zeta_n)$):

a) Show that there is a unique quadratic character of conductor n when n is an odd prime or 4, and that there are exactly 2 quadratic characters of conductor 8, one odd and one even.
b) Show that if $D \in \mathbb{N}$ is the conductor of a quadratic character, then D must be of the form dm, with m odd and square-free, and $d = 1, 4, 8$.
c) Show the same for $D \in \mathbb{N}$ with $D = |D_K|$ for K a quadratic field.
d) With $D = dm$ as above, show that if $d < 8$, there exists a unique quadratic character of conductor D and a unique quadratic field K with $|D_K| = D$, and note that the character is even if and only if K is real.
e) With $D = dm$ as above, if $d = 8$ show that there exist exactly 2 quadratic characters of conductor D, one odd and one even, and exactly 2 quadratic fields K with $|D_K| = D$, one imaginary and one real.

In the following exercises, for p an odd prime, and $j > 0$, $\left(\frac{j}{p}\right)$ denotes the Legendre symbol, defined to be 1 if j is a quadratic residue modulo p and -1 otherwise.

Exercise 7.5. Show that if $p \equiv 1 \pmod 4$, that

$$\prod_{0<j<\frac{p}{2}:\left(\frac{j}{p}\right)=-1} \left(\sin \frac{\pi j}{p}\right) > \prod_{0<j<\frac{p}{2}:\left(\frac{j}{p}\right)=1} \left(\sin \frac{\pi j}{p}\right).$$

Note that the number of quadratic residues and non-residues in $(0, \frac{p}{2})$ is the same. Since sin is monotone increasing in $[0, \frac{\pi}{2}]$, this exercise shows that quadratic residues cluster in the first half of the interval $(0, \frac{p}{2})$.

Exercise 7.6. Show that for $p \equiv 3 \pmod 4$, and $p > 3$, the class number formula for $\mathbb{Q}(\sqrt{-p})$ may be rewritten as

$$h_K = \frac{1}{2 - \left(\frac{2}{p}\right)} \sum_{0<j<\frac{p}{2}} \left(\frac{j}{p}\right),$$

and conclude that there are more quadratic residues in the interval $(0, \frac{p}{2})$ than non-quadratic residues (of course, this conclusion holds also for $p = 3$).

Hint: use the behavior of $\left(\frac{j}{p}\right)$ under $j \mapsto p-j$, and compare the sums obtained by separating j first by size, and then by parity.

CHAPTER 8

Local fields

When attempting to determine whether a Diophantine equation has solutions over the integers, it is only natural to consider first whether it has solutions modulo p for different p. Similarly, when we try to analyze rings of integers, we frequently benefit from analyzing one prime at a time by considering the local rings. The theory of local fields may be viewed as lying between these two approaches: we will obtain more information than we would just working modulo p, but not as much as working in the local ring of the ring of integers itself. Conversely, while the structure of local fields is more complex than that of finite fields, it is simpler, or at least easier to analyze, than the structure of the number fields themselves.

The tool for achieving this is completion. Just as $\mathbb{R}$ is the completion of $\mathbb{Q}$ with respect to the standard absolute value, and it is easier to analyze roots of polynomials and algebraic extensions in $\mathbb{R}$ than in $\mathbb{Q}$, local fields will be completions of number fields with respect to p-adic absolute values, and it is easier to analyze roots of polynomials and algebraic extensions in the local field case.

8.1. The p-adic integers

We describe three constructions of $\mathbb{Z}_p$, the p-adic integers, and see that they are equivalent. We fix a prime number p for the remainder of the discussion.

Definition 8.1.1. The ring $\mathbb{Z}_p$ is defined to be the set of integers written in base p, and allowed to have infinitely many digits, under the addition and multiplication obtained by the usual formulas.

Since the usual laws for addition and multiplication give formulas for the nth digit in terms of previous digits, we can use them just as well to add and multiply numbers with infinitely many digits.

Definition 8.1.2. The ring $\mathbb{Z}_p$ is defined to be $\{(a_0, a_1, a_2, \dots) : a_i \in \mathbb{Z}/p^{i+1}\mathbb{Z}, a_{i+1} \equiv a_i \pmod{p^{i+1}}\}$, with coordinatewise addition and multiplication.

Remark 8.1.3. This is the same thing as saying that $\mathbb{Z}_p$ is the **inverse limit** over the rings $\mathbb{Z}/p^n\mathbb{Z}$ under the usual quotient maps.

The final definition requires some preliminary terminology.

Definition 8.1.4. We define the **p-adic valuation** $\nu_p : \mathbb{Q}^* \to \mathbb{Z}$ by $\nu_p(\frac{x}{y}) = \mathrm{ord}_p(x) - \mathrm{ord}_p(y)$. The **$p$-adic absolute value** $||\cdot||_p : \mathbb{Q} \to \mathbb{R}_{\geqslant 0}$ is defined by $||z||_p = p^{-\nu_p(z)}$ for $z \neq 0$, and $||0||_p = 0$.

Thus, under the p-adic absolute value, the more powers of p that are in a number, the "smaller" it is.

Easy Facts 8.1.5. ν_p satisfies:

(i) $\nu_p(-1) = 0$;
(ii) $\nu_p(z_1 z_2) = \nu_p(z_1) + \nu_p(z_2)$;
(iii) $\nu_p(z_1 + z_2) \geqslant \min\{\nu_p(z_1), \nu_p(z_2)\}$.

$|| \cdot ||_p$ satisfies:

(i) $||z||_p = 0$ if and only if $z = 0$;
(ii) $|| - 1||_p = 1$;
(iii) $||z_1 z_2||_p = ||z_1||_p ||z_2||_p$;
(iv) $||z_1 + z_2||_p \leqslant \max\{||z_1||_p, ||z_2||_p\}$.

In particular, $|| \cdot ||_p$ satisfies the triangle inequality, and induces a metric on $\mathbb{Q}$, and hence on $\mathbb{Z}$.

Definition 8.1.6. The ring $\mathbb{Z}_p$ is the completion of $\mathbb{Z}$ with respect to the metric $d(x, y) := ||x - y||_p$.

Proposition 8.1.7. *The three definitions of $\mathbb{Z}_p$ above are equivalent.*

Proof. It is easy to see that the first two definitions are the same: indeed, an element of $\mathbb{Z}/p^n\mathbb{Z}$ may be written uniquely as an integer base p with at most n digits, and the number obtained from this by taking the image in $\mathbb{Z}/p^{n-1}\mathbb{Z}$ is simply obtained by dropping the highest digit. Thus we obtain a natural bijection.

To see that the second and third definitions agree, we note that sending $a_i \in \mathbb{Z}/p^{i+1}\mathbb{Z}$ to any representative of it in $\mathbb{Z}$ gives a Cauchy sequence under the metric $d(\cdot, \cdot)$, and one checks that for different choices of representatives, we obtain Cauchy sequences whose differences tend to 0, and are therefore equivalent in the completion, so we obtain a well-defined map from the second set to the third. We obtain a map back via the observation that if $(b_0, b_1, b_2, \dots)$ is a Cauchy sequence of integers, then for any i, there exists N such that for all $n_1, n_2 \geqslant N$, $b_{n_1} \equiv b_{n_2} \pmod{p^{i+1}}$. We can then set $a_i = b_N$, and doing this for each i gives a map from the third set to the second. It is easy to check that these maps are mutually inverse, and hence define a bijection.

Since the addition and multiplication in all definitions are obtained from that of the integers, it is easy to see that they agree. □

Under the second definition of $\mathbb{Z}_p$, it is easy to check that it is an integral domain, since if two elements are non-zero in the ith and jth places respectively, their product must be non-zero in the $(i + j)$th place. We can then define $\mathbb{Q}_p$ to be the field of fractions of $\mathbb{Z}_p$.

At this point, it seems reasonable to wonder what makes $\mathbb{Z}_p$ simpler in any sense that the local ring $\mathbb{Z}_{(p)} \subseteq \mathbb{Q}$. One answer is provided by the following, which is elementary to prove, and which will be subsumed by the stronger statement of Theorem 8.4.1.

Lemma 8.1.8. *(Hensel's) Suppose that $f(x) \in \mathbb{Z}[x]$ is such that there exists an $x_0 \in \mathbb{Z}$ with $f(x_0) \equiv 0 \pmod{p}$, and $f'(x_0) \not\equiv 0 \pmod{p}$. Then there exists a root of $f(x)$ in $\mathbb{Z}_p$ agreeing with x_0 modulo p.*

8.2. $\mathfrak{p}$-adic completions of number fields

Let K be a number field, and $\mathfrak{p}$ a prime ideal of $\mathscr{O}_K$. Recall that for $x \in K^*$, we had defined $\operatorname{ord}_{\mathfrak{p}}(x)$ to be the number of powers of $\mathfrak{p}$ (possibly negative) occurring in the prime factorization of the fractional ideal $x\mathscr{O}_K$. In keeping with the preceding, we shall write $\nu_{\mathfrak{p}}(x) := \operatorname{ord}_{\mathfrak{p}}(x)$.

Definition 8.2.1. We define the **$\mathfrak{p}$-adic absolute value** $||\cdot||_p : K \to \mathbb{R}_{\geqslant 0}$ by $||x||_\mathfrak{p} = N(\mathfrak{p})^{-\nu_\mathfrak{p}(x)}$ for $x \neq 0$, and $||0||_\mathfrak{p} = 0$.

Easy Facts 8.2.2. $||\cdot||_\mathfrak{p}$ satisfies:

(i) $||x||_\mathfrak{p} = 0$ if and only if $x = 0$;
(ii) $||-1||_\mathfrak{p} = 1$;
(iii) $||x_1 x_2||_\mathfrak{p} = ||x_1||_\mathfrak{p} ||x_2||_\mathfrak{p}$;
(iv) $||x_1 + x_2||_\mathfrak{p} \leqslant \max\{||x_1||_\mathfrak{p}, ||x_2||_\mathfrak{p}\}$.

Thus, as before, we get a metric $d(x_1, x_2) := ||x_1 - x_2||_\mathfrak{p}$, and we can define:

Definition 8.2.3. $K_\mathfrak{p}$ is the completion of K with respect to the metric $d(\cdot,\cdot)$. $\hat{\mathscr{O}}_{K,\mathfrak{p}}$ is the completion of $\mathscr{O}_K$ with respect to the same metric.

The notation $\hat{\mathscr{O}}_{K,\mathfrak{p}}$ is somewhat cumbersome, but we have already defined $\mathscr{O}_{K,\mathfrak{p}}$ to be the standard local ring, and this notation is consistent with the standard notation of algebraic geometry.

Proposition 8.2.4. *We have:*

(i) *$K_\mathfrak{p}$ is the field of fractions of $\hat{\mathscr{O}}_{K,\mathfrak{p}}$.*
(ii) *$\hat{\mathscr{O}}_{K,\mathfrak{p}}$ is the subset of elements of $K_\mathfrak{p}$ with absolute value at most* 1.
(iii) *$\hat{\mathscr{O}}_{K,\mathfrak{p}}$ is also the completion of $\mathscr{O}_{K,\mathfrak{p}}$ with respect to the metric $d(\cdot,\cdot)$.*
(iv) *$\hat{\mathscr{O}}_{K,\mathfrak{p}}$ is the inverse limit of $\mathscr{O}_K/\mathfrak{p}^n$ over $n \in \mathbb{N}$.*

Proof. The main content is actually in (iii), which is equivalent to the statement that $\mathscr{O}_K$ is dense in $\mathscr{O}_{K,\mathfrak{p}}$ under the metric $d_\mathfrak{p}$. For this, we need to use the fact that $\mathscr{O}_K/\mathfrak{p}^n \xrightarrow{\sim} \mathscr{O}_{K,\mathfrak{p}}/\mathfrak{p}^n$ for all n, which we have used before, but never justified. We certainly get a map induced by $\mathscr{O}_K \hookrightarrow \mathscr{O}_{K,\mathfrak{p}}$, and the injectivity is easy to check, so we need to check surjectivity. Given $\frac{x}{s} \in \mathscr{O}_{K,\mathfrak{p}}$, with $x, s \in \mathscr{O}_K$, and $s \notin \mathfrak{p}$, we want to show that there exists $y \in \mathscr{O}_K$ with $ys \equiv x \pmod{\mathfrak{p}^n}$; it clearly suffices to see that s is a unit in $\mathscr{O}_K/\mathfrak{p}^n$. Since the latter is a finite ring, there exist $i < j$ with $s^i = s^j$ modulo $\mathfrak{p}^n$, so $s^i(1 - s^{j-i}) \in \mathfrak{p}^n$. Since $s \notin \mathfrak{p}$, by unique factorization into prime ideals, we find $(1 - s^{j-i}) \in \mathfrak{p}^n$, and s is a unit modulo $\mathfrak{p}^n$, as desired. We thus obtain the isomorphism $\mathscr{O}_K/\mathfrak{p}^n \xrightarrow{\sim} \mathscr{O}_{K,\mathfrak{p}}/\mathfrak{p}^n$.

It then follows easily that $\mathscr{O}_K$ is dense in $\mathscr{O}_{K,\mathfrak{p}}$: given $\frac{x}{s} \in \mathscr{O}_{K,\mathfrak{p}}$, for each n we can find $y_n \in \mathscr{O}_K$ such that $y_n \equiv \frac{x}{s} \pmod{\mathfrak{p}^n}$, which then gives a sequence in $\mathscr{O}_K$ converging to $\frac{x}{s}$. This completes the proof of (iii).

We then see (ii) because $\mathscr{O}_{K,\mathfrak{p}}$ is nearly defined to be the subset of K with absolute value at most 1; certainly, we have $||z||_\mathfrak{p} \leqslant 1$ for any $z \in \mathscr{O}_{K,\mathfrak{p}}$. Conversely, given $\frac{x}{y} \in K^*$ with $||\frac{x}{y}||_\mathfrak{p} \leqslant 1$, we claim we can write $\frac{x}{y} = \frac{x'}{s}$ with $s \notin \mathfrak{p}$. Indeed, let t be a generator of $\mathfrak{p}$ in $\mathscr{O}_{K,\mathfrak{p}}$, which we know is a DVR. Then we can write $x = t^e \frac{s_x}{s'_x}$, $y = t^{e'} \frac{s_y}{s'_y}$, with $s_x, s'_x, s_y, s'_y \notin \mathfrak{p}$, and $e \geqslant e'$ since $||\frac{x}{y}||_\mathfrak{p} \leqslant 1$. Thus $\frac{x}{y} = \frac{t^{e-e'} s_x s'_y}{s'_x s_y} \in \mathscr{O}_{K,\mathfrak{p}}$, as desired. It remains to see that this description is maintained under completion. But here we observe that $||\cdot||_\mathfrak{p}$ takes on discrete values away from 0, so any sequence in K converging to a point with $||\cdot||_\mathfrak{p} \leqslant 1$ must, after a finite number of terms, have every element with $||\cdot||_\mathfrak{p} \leqslant 1$, and hence contained in $\mathscr{O}_{K,\mathfrak{p}}$.

(i) then follows easily from (ii), as we see that for any $z \in K$, either $z \in \hat{\mathscr{O}}_{K,\mathfrak{p}}$ or $1/z$ is.

Finally, (iv) follows easily from the definition of $\hat{\mathscr{O}}_{K,\mathfrak{p}}$, just as in the case of $\mathbb{Z}_p$. □

Algebraic Geometry Remark 8.2.5. The rings $\hat{\mathscr{O}}_{K,\mathfrak{p}}$ are thus completions of local rings of rings of integers, in the sense used in algebraic geometry. As a result, they are closely analogous to rings of power series, and exhibit much of the same behavior.

We give two basic structural corollaries.

Corollary 8.2.6. *For any number field K and prime $\mathfrak{p}$, the ring $\hat{\mathscr{O}}_{K,\mathfrak{p}}$ is compact for the $\mathfrak{p}$-adic metric.*

Proof. Let $x_1, x_2, x_3, \dots$ be any sequence in $\hat{\mathscr{O}}_{K,\mathfrak{p}}$. We show for any $\epsilon > 0$ that there exists a subsequence $x_{i_{\epsilon,1}}, x_{i_{\epsilon,2}}, x_{i_{\epsilon,3}}, \dots$ with the property that any two elements in it have distance at most ϵ. We use the description of $\hat{\mathscr{O}}_{K,\mathfrak{p}}$ from (iv) of the proposition, noting that if we choose n large enough, two elements of $\hat{\mathscr{O}}_{K,\mathfrak{p}}$ which agree in $\mathscr{O}_K/\mathfrak{p}^n$ will be within ϵ of one another. Since there are only finitely elements of $\mathscr{O}_K/\mathfrak{p}^n$, at least one element must be hit infinitely often by the x_i, and we use this to construct our subsequence. If we then inductively take subsequences with ϵ tending to 0, we obtain a sequence of subsequences of $x_1, x_2, x_3, \dots$, and choosing a diagonal subsequence will give us a Cauchy sequence, which then converges in $\hat{\mathscr{O}}_{K,\mathfrak{p}}$ by completeness. □

Corollary 8.2.7. *For any number field K and prime $\mathfrak{p}$, the ring $\hat{\mathscr{O}}_{K,\mathfrak{p}}$ is a DVR.*

Proof. Indeed, since $\hat{\mathscr{O}}_{K,\mathfrak{p}} = \{z \in K_\mathfrak{p} : ||z||_\mathfrak{p} \leqslant 1\}$, we see that $\mathfrak{m}_\mathfrak{p} := \{z \in K_\mathfrak{p} : ||z||_\mathfrak{p} < 1\}$ must be the unique maximal ideal. Furthermore, given any ideal I, since $||\cdot||_\mathfrak{p}$ is discrete away from 0, there is some $t \in I$ with $||t||_\mathfrak{p}$ maximal, and we then have that t generates I. □

8.3. Local fields

For our purposes, a **local field** is $\mathbb{R}$, $\mathbb{C}$, or $K_\mathfrak{p}$ for some number field K and prime ideal $\mathfrak{p} \subseteq \mathscr{O}_K$. We say that the first category is **Archimedean**, while the second category is **non-Archimedean** (the terminology arises from the fact that for the second class, we have $||z_1 + z_2|| \leqslant \max\{||z_1||, ||z_2||\}$.

This definition of local field may seem rather ad-hoc, but from the proper point of view, it is quite natural. Indeed, we have the following basic theorem:

Theorem 8.3.1. *Let K be a number field. Then every absolute value $||\cdot||$ on K is equivalent to either $|\cdot| \circ \sigma$ for some $\sigma : K \to \mathbb{C}$, or $||\cdot||_\mathfrak{p}$ for some $\mathfrak{p} \subseteq \mathscr{O}_K$.*

Here, an absolute value $||\cdot|| : K \to \mathbb{R}_{\geqslant 0}$ is defined to be any map satisfying:

(i) $||z|| = 0$ *if and only if* $z = 0$*;*
(ii) $||z_1 z_2|| = ||z_1|| \cdot ||z_2||$*;*
(iii) $||z_1 + z_2|| \leqslant ||z_1|| + ||z_2||$.

The equivalence condition for absolute values is that the associated metric $d(z_1, z_2) := ||z_1 - z_2||$ induce the same topology on K.

Since this result is stated only for context, we do not give the proof; see [**1**, Thm. 4.4.1].

We denote the set of absolute values described in the theorem by M_K. Thus, every completion of K with respect to an absolute value arises as the completion with respect to an element of M_K, which is one of the local fields we have discussed. This is reflected by the following basic fact:

PROPOSITION 8.3.2. *Given $z \in K^*$, we have the identity:*

$$\prod_{||\cdot|| \in M_K} ||z|| = 1.$$

PROOF. We first treat the case $z \in \mathscr{O}_K$ (and still non-zero). We then have that $||z||_{\mathfrak{p}} = 1$ for all but finitely many $\mathfrak{p}$, so the product makes sense. Indeed, if we write $(z) = \prod_i \mathfrak{p}_i^{e_i}$, we see that $||z||_{\mathfrak{p}_i} = N(\mathfrak{p}_i)^{-e_i}$, so

$$\prod_{\mathfrak{p}} ||z||_{\mathfrak{p}} \frac{1}{N((z))} = \frac{1}{|N_{K/\mathbb{Q}}(z)|}.$$

On the other hand, $\prod_{\sigma_i} |\sigma_i(z)| = |N_{K/\mathbb{Q}}(z)|$, and since every absolute value in M_K is one of these two possibilities, we get the desired identity. The statement for all $z \in K^*$ then follows by the multiplicativity of all $|| \cdot || \in M_K$. □

Finally, we remark that the class of local fields itself is represented by a rather natural condition:

THEOREM 8.3.3. *Let K be a topological field complete with respect to some absolute value, and locally compact. Then if K has characteristic* 0*, it is a local field in our sense.*

In fact, one typically defines a local field by this condition; the case of characteristic p is closely analogous, and every field K of characteristic p satisfying the conditions of the theorem is of the form $\mathbb{F}_{p^e}((t))$, with the absolute value $||z|| = 2^{-\operatorname{ord}_t z}$ for any $2 \in \mathbb{R}_{>1}$.

In any case, in order to study roots of polynomials over K, we see that it is rather natural to study roots over each $K_{\mathfrak{p}}$.

8.4. Hensel's lemma

There are many variations on Hensel's lemma, considering factorization rather than roots of polynomials, or multiple polynomials in multiple variables. We give the version given in [**7**], a variation on Newton's method for finding roots of real polynomials and then conclude a corollary for multivariate polynomials.

THEOREM 8.4.1. *Let $f(x) \in \mathscr{O}_K[x]$, and $x_0 \in \mathscr{O}_K$ such that $||f(x_0)||_{\mathfrak{p}} < ||f'(x_0)^2||_{\mathfrak{p}}$. Then the sequence determined by $x_i = x_{i-1} - \frac{f(x_{i-1})}{f'(x_{i-1})}$ converges to a root of $f(x)$ in $\hat{\mathscr{O}}_{K,\mathfrak{p}}$, which agrees with x_0 modulo $\mathfrak{p}$.*

Before proving the theorem, we note the following more general corollary.

COROLLARY 8.4.2. *Let $f(x_1, \dots, x_n)$ be a polynomial with coefficients in $\mathscr{O}_K$, and suppose that $(y_1, \dots, y_n) \in \mathscr{O}_K^n$ is such that*

$$||f(y_1, \dots, y_n)||_{\mathfrak{p}} < ||\frac{\partial f}{\partial x_i}(y_1, \dots, y_n)^2||_{\mathfrak{p}}$$

for some i. Then there exists $y_i' \in \hat{\mathscr{O}}_{K,\mathfrak{p}}$ such that $(y_1, \dots, y_i', \dots, y_n)$ is a root of f.

Proof. Indeed, we may substitute y_j for x_j for all $j \neq i$, and treat f as a polynomial in the single variable x_i. The statement then follows directly from the previous theorem. $\square$

To prove Theorem 8.4.1, we start by observing the following slightly stronger version of the "triangle inequality" for $K_{\mathfrak{p}}$:

Easy Fact 8.4.3. Given $z_1, z_2 \in K_{\mathfrak{p}}$, we have

$$||z_1 + z_2||_{\mathfrak{p}} \leqslant \max\{||z_1||_{\mathfrak{p}}, ||z_2||_{\mathfrak{p}}\},$$

with equality whenever $||z_1||_{\mathfrak{p}} \neq ||z_2||_{\mathfrak{p}}$.

Proof of theorem. Write $C := \frac{||f(x_0)||_{\mathfrak{p}}}{||f'(x_0)^2||_{\mathfrak{p}}}$. We show the following by induction:

(i) $||x_i||_{\mathfrak{p}} \leqslant 1$;
(ii) $||x_i - x_0|| \leqslant C$;
(iii) $||f'(x_i)||_{\mathfrak{p}} = ||f'(x_{i-1})||_{\mathfrak{p}}$ for $i > 0$;
(iv) $\frac{||f(x_i)||_{\mathfrak{p}}}{||f'(x_i)^2||_{\mathfrak{p}}} \leqslant C^{2^i}$;

We then see that the sequence converges, because

$$||x_i - x_{i-1}||_{\mathfrak{p}} = \frac{||f(x_{i-1})||_{\mathfrak{p}}}{||f'(x_{i-1})||_{\mathfrak{p}}} \leqslant \frac{||f(x_{i-1})||_{\mathfrak{p}}}{||f'(x_{i-1})^2||_{\mathfrak{p}}},$$

using that $||f'(x_{i-1})||_{\mathfrak{p}} \leqslant 1$ by (i), and this goes to 0 by (iii). Then, (iii) and (iv) together imply that the sequence converges to a root of f in $K_{\mathfrak{p}}$, (i) implies that the root lies in $\hat{\mathscr{O}}_{K,\mathfrak{p}}$, and (ii) implies that it agrees with x_0 modulo $\mathfrak{p}$.

The $i = 0$ case is simply the stated hypotheses. For $i > 0$, we observe that from the statement for $i-1$, $||x_i - x_{i-1}||_{\mathfrak{p}} = \frac{||f(x_{i-1})||_{\mathfrak{p}}}{||f'(x_{i-1})||_{\mathfrak{p}}} \leqslant \frac{||f(x_{i-1})||_{\mathfrak{p}}}{||f'(x_{i-1})^2||_{\mathfrak{p}}} \leqslant C^{2^{i-1}} \leqslant C < 1$, so since $||x_{i-1}||_{\mathfrak{p}} \leqslant 1$, we have $||x_i||_{\mathfrak{p}} \leqslant 1$ as well, giving (i) (and also showing by induction that $x_i \in \mathscr{O}_{K,\mathfrak{p}}$). We also have $||x_i - x_0||_{\mathfrak{p}} \leqslant \max\{||x_i - x_{i-1}||_{\mathfrak{p}}, ||x_{i-1} - x_0||_{\mathfrak{p}}\} \leqslant C$, giving (ii).

Noting that $\frac{1}{n!}f^{(n)}(x) \in \mathscr{O}_K[x]$ for any n, we Taylor-expand $f'(x)$ about x_{i-1} and find

$$f'(x_i) = f'(x_{i-1}) + y' \frac{f(x_{i-1})}{f'(x_{i-1})}$$

for some $y' \in \mathscr{O}_{K,\mathfrak{p}}$. By the (iv) in the induction hypothesis, $||f'(x_{i-1})||_{\mathfrak{p}} > \frac{||f(x_{i-1})||_{\mathfrak{p}}}{||f'(x_{i-1})||_{\mathfrak{p}}}$, and since $||y'||_{\mathfrak{p}} \leqslant 1$, we find $||f'(x_i)||_{\mathfrak{p}} = ||f'(x_{i-1})||_{\mathfrak{p}}$, giving (iii).

Finally, if we Taylor-expand $f(x)$ about x_{i-1}, we find

$$f(x_i) = f(x_{i-1}) - f'(x_{i-1})\frac{f(x_{i-1})}{f'(x_{i-1})} + y\left(\frac{f(x_{i-1})}{f'(x_{i-1})}\right)^2$$

for some $y \in \mathscr{O}_{K,\mathfrak{p}}$. Cancelling the first two terms, and using $||y||_{\mathfrak{p}} \leqslant 1$, we find $||f(x_i)||_{\mathfrak{p}} \leqslant \left|\left|\frac{f(x_{i-1})}{f'(x_{i-1})}\right|\right|_{\mathfrak{p}}^2$. Taking absolute values and substituting (iii) and the induction hypothesis of (iv) for $i-1$, we obtain (iv) for i. $\square$

Although it is not directly related, we also observe the following:

Proposition 8.4.4. *Let $f(x_1, \ldots, x_n)$ be a polynomial with coefficients in $\mathscr{O}_K$, and suppose that f has roots modulo $\mathfrak{p}^n$ for all n. Then f has roots in $\hat{\mathscr{O}}_{K,\mathfrak{p}}$.*

Proof. We claim that there exists a sequence $(x_{1,1},\dots,x_{1,n}),(x_{2,1},\dots,x_{2,n}),\dots$ of tuples in $\mathscr{O}_K^n$ such that $f(x_{i,1},\dots,x_{i,n}) \equiv 0 \pmod{\mathfrak{p}^i}$ for $i \geqslant 1$ and $(x_{i,1},\dots,x_{i,n}) \equiv (x_{i-1,1},\dots,x_{i-1,n}) \pmod{\mathfrak{p}^{i-1}}$ for $i > 1$; it follows immediately from the definitions that this gives an element of $\hat{\mathscr{O}}_{K,\mathfrak{p}}^n$ which is a root of f.

The proof is by induction on i, constructing the $(x_{i,1},\dots,x_{i-1,n})$ under the further hypothesis that for an infinite number of $j > i$, there exists $(x_{j,1},\dots,x_{j,n})$ such that $f(x_{j,1},\dots,x_{j,n}) \equiv 0 \pmod{\mathfrak{p}^j}$ and $(x_{j,1},\dots,x_{j,n}) \equiv (x_{i,1},\dots,x_{i,n}) \pmod{\mathfrak{p}^i}$. By hypothesis, for every j we have some $(x_{j,1},\dots,x_{j,n})$ satisfying the first condition, and each of these gives a root of f modulo i for any $i < j$.

Since there are only finitely many possibilities for $(x_{1,1},\dots,x_{1,n})$ modulo $\mathfrak{p}$, this infinite sequence of roots modulo $\mathfrak{p}^j$ for each j must give some root modulo $\mathfrak{p}$ infinitely often, so we choose this for $(x_{1,1},\dots,x_{1,n})$. Similarly, given the sequence up to $(x_{i,1},\dots,x_{i,n})$, by hypothesis there are infinitely many j such that there exists $(x_{j,1},\dots,x_{j,n})$ satisfying $f(x_{j,1},\dots,x_{j,n}) \equiv 0 \pmod{\mathfrak{p}^j}$ and $(x_{j,1},\dots,x_{j,n}) \equiv (x_{i,1},\dots,x_{i,n}) \pmod{\mathfrak{p}^i}$. Each of these gives a root modulo $\mathfrak{p}^{i+1}$ agreeing with our chosen one modulo $\mathfrak{p}^i$, and since there are only finitely many possible roots modulo $\mathfrak{p}^{i+1}$, one of them must be given infinitely often as j varies, and we choose this for $(x_{i+1,1},\dots,x_{i+1,n})$. □

8.5. The Hasse-Minkowski theorem and Hasse principle

In many arguments on number fields, it often turns out to be useful to first analyze the situation for local fields, and then to deduce the desired statement for number fields. One might hope to do this in finding roots of polynomials. The Hasse-Minkowski theorem gives one example where this is possible:

Theorem 8.5.1. *(Hasse-Minkowski) A multivariate quadratic polynomial f with coefficients in a number field K has a root in K if and only if it has a root in every local field arising as a completion of K.*

We will not give a complete proof of this theorem, although we will revisit it in the context of more general results of class field theory. For a complete proof in the case $K = \mathbb{Q}$, see [**9**, Thm. 8, p. 41].

We observe that the hypothesis that the polynomial have roots even over $\mathbb{R}$ is necessary:

Example 8.5.2. Consider the polynomial $f(u,v,x,y) = u^2+v^2+x^2+y^2+1$; this has no roots in $\mathbb{R}$, hence no roots in $\mathbb{Q}$. On the other hand, it is a theorem that every integer is a sum of four squares, so f has roots modulo p^n for every p, n and hence in every $\mathbb{Q}_p$, by the proposition of the previous section.

The **Hasse principle** is the philosophy that if something is true in every completion of a number field, it should be true in the number field. The simplest case to consider is whether a polynomial has a root, and the Hasse-Minkowski theorem gives a non-trivial case in which the principle holds.

However, the principle fails almost immediately beyond the situation treated by the theorem: it is known (but rather non-trivial to prove) that the equation $3x^3 + 4y^3 + 5 = 0$ has roots in every $\mathbb{Q}_p$ and in $\mathbb{R}$, but not in $\mathbb{Q}$. In this case of cubics in two variables, this phenomenon is measured by the **Tate-Shafarevich group**, the finiteness of which is still a major open conjecture in elliptic curve theory, closely tied to the conjecture of Birch and Swinnerton-Dyer.

Nonetheless, there are a number of important cases beyond polynomials having roots for which the Hasse principle holds more generally, and we will mention some of these during our discussion of class field theory.

8.6. Preliminaries for field extensions

Our next topic is the study of extensions of local fields; for instance, the relation between the Galois group of an extension of number fields and an associated extensions of local fields. To prepare for this, we discuss some preliminary points.

The first is a renormalized p-adic absolute value. The map $||\cdot||_{\mathfrak{p}}$ works very well when analyzing all the completions of a given number field, as demonstrated by the product formula, but it is not compatible with extensions of local fields. We thus introduce:

Definition 8.6.1. Given K, $\mathfrak{p} \subseteq \mathscr{O}_K$, if $p = \mathfrak{p} \cap \mathbb{Z}$, we let $|\cdot|_{\mathfrak{p}} : K \to \mathbb{R}_{\geqslant 0}$ be the (unique) absolute value on K extending the p-adic absolute value on $\mathbb{Q}$: namely, we set

$$|z|_{\mathfrak{p}} := p^{-\nu_{\mathfrak{p}}(z)/e_{\mathfrak{p}/p}}.$$

It is easy to check that $|\cdot|_{\mathfrak{p}}$ induces the same topology on K as $||\cdot||_{\mathfrak{p}}$. We see that $|z|_{\mathfrak{p}} \leqslant 1$ if and only if $||z||_{\mathfrak{p}} \leqslant 1$, so we can still describe $\mathscr{O}_{K,\mathfrak{p}}$ as the subset of K with $|\cdot|_{\mathfrak{p}} \leqslant 1$.

We will invoke the two following general theorems on extensions:

Theorem 8.6.2. *Let $K_{\mathfrak{p}}$ be a non-Archimedean local field, and $\hat{L}$ a finite extension. Then the absolute value $|\cdot|_{\mathfrak{p}}$ extends uniquely to $\hat{L}$.*

For the proof, see [**7**, pp. 32-33].

Theorem 8.6.3. *Let $K_{\mathfrak{p}}$ be a non-Archimedean local field, and $\hat{L}$ a finite extension. Then $\hat{L} = L_{\mathfrak{q}}$ for some finite extension L of K and $\mathfrak{q}$ lying above $\mathfrak{p}$.*

The main idea of the proof is to show that if one starts with a generator α for $\hat{L}$ over $K_{\mathfrak{p}}$, and if one approximates it sufficiently closely by some β, then β generates the same field extension. One can then approximate α by an element algebraic over K to get the desired result. See [**7**, pp. 43-44].

8.7. Unramified extensions

We now begin to consider extensions in earnest. We suppose we are given an extension L/K of number fields, and $\mathfrak{q} \in \mathscr{O}_L$ lying over $\mathfrak{p} \in \mathscr{O}_K$.

Proposition 8.7.1. *$\hat{\mathscr{O}}_{K,\mathfrak{p}}$ and $\hat{\mathscr{O}}_{L,\mathfrak{q}}$ are Dedekind domains, with $\hat{\mathscr{O}}_{L,\mathfrak{q}}$ a finitely generated module over $\hat{\mathscr{O}}_{K,\mathfrak{p}}$.*

Proof. Indeed, we have already seen in Corollary 8.2.7 that $\hat{\mathscr{O}}_{K,\mathfrak{p}}$ and $\hat{\mathscr{O}}_{L,\mathfrak{q}}$ are DVRs, hence they are Dedekind domains. We claim that if $x_1, \dots, x_n$ generate $\mathscr{O}_L$ as a module over $\mathscr{O}_K$, then they also generate $\hat{\mathscr{O}}_{L,\mathfrak{q}}$ as a module over $\hat{\mathscr{O}}_{K,\mathfrak{p}}$. An arbitrary element of $\hat{\mathscr{O}}_{L,\mathfrak{q}}$ can be written as a Cauchy sequence $z_1, z_2, z_3, \dots$ of elements $z_i \in \mathscr{O}_L$, and we can then write each $z_i = c_{i,1}x_1 + \dots c_{i,n}x_n$ for some $c_i \in \mathscr{O}_K$. By compactness of $\hat{\mathscr{O}}_{K,\mathfrak{p}}$, we can choose a subsequence of the z_i in which the $c_{i,1}$ form a Cauchy sequence, and then a subsequence of that in which the $c_{i,2}$ form a Cauchy sequence, and so forth, until we have a subsequence of the z_i in which $c_{i,j}$ is a Cauchy sequence for all $j = 1, \ldots, n$. We may then think of the $c_{i,j}$ as giving elements of $\hat{\mathscr{O}}_{K,\mathfrak{p}}$ for each j, so that the element of $\hat{\mathscr{O}}_{L,\mathfrak{q}}$ described by the z_i can be written as a $\hat{\mathscr{O}}_{K,\mathfrak{p}}$-linear combination of the x_i, as desired. □

We then easily check the following, using the above and Theorem 3.1.2, together with the observation that since $\hat{\mathscr{O}}_{L,\mathfrak{q}}$ is a DVR, it has a unique prime ideal lying over $\mathfrak{p}\hat{\mathscr{O}}_{K,\mathfrak{p}}$, to conclude (iii):

Easy Facts 8.7.2. We have:

(i) $f_{\mathfrak{q}/\mathfrak{p}} = \dim_{\hat{\mathscr{O}}_{K,\mathfrak{p}}/\mathfrak{p}} \hat{\mathscr{O}}_{L,\mathfrak{q}}/\mathfrak{q}$;

(ii) $\mathfrak{p}\hat{\mathscr{O}}_{L,\mathfrak{q}} = \mathfrak{q}^{e_{\mathfrak{q}/\mathfrak{p}}}$;

(iii) $[L_\mathfrak{q} : K_\mathfrak{p}] = e_{\mathfrak{q}/\mathfrak{p}} f_{\mathfrak{q}/\mathfrak{p}}$.

We can now show:

Theorem 8.7.3. *Write $q := \#(\hat{\mathscr{O}}_{K,\mathfrak{p}}/(\mathfrak{p}))$. Every finite extension $\hat{L}$ of $K_\mathfrak{p}$ in which $\mathfrak{p}$ is unramified can be obtained by adjoining a $(q^f - 1)$st root of unity, with $f = [\hat{L} : K_\mathfrak{p}]$, and every such extension has $\mathfrak{p}$ unramified.*

In particular, there is exactly one unramified extension of $K_\mathfrak{p}$ of each degree, and it is Galois.

The key point is the following:

Lemma 8.7.4. *Let $L_\mathfrak{q}$ be a non-Archimedean local field, and $n \in N$ such that if we write $(p) := \mathfrak{q} \cap \mathbb{Z}$, we have $p \not| n$. Then $L_\mathfrak{q}$ contains a primitive nth root of unity if and only if $n | (\#(\hat{\mathscr{O}}_{L,\mathfrak{q}})/(\mathfrak{q}) - 1)$.*

Proof. Write $\mathbb{F}_\mathfrak{q} = (\hat{\mathscr{O}}_{L,\mathfrak{q}})/(\mathfrak{q})$, and $m = \#\mathbb{F}_\mathfrak{q} - 1 = \#\mathbb{F}_\mathfrak{q}^*$. Then $\mathbb{F}_\mathfrak{q}$ contains a primitive nth root of unity if and only if $n|m$.

If $L_\mathfrak{q}$ contains a primitive nth root of unity ζ_n, then $\Phi_n(x)$ splits in $L_\mathfrak{q}$, and hence in $\mathbb{F}_\mathfrak{q}$. We have seen that a root of $\Phi_n(x)$ is a primitive nth root of unity over any field of characteristic prime to n, so it follows that $\mathbb{F}_\mathfrak{q}$ has a primitive nth root of unity, and $n|m$.

Conversely, if $n|m$, we have that $\Phi_n(x)$ has a root over $\mathbb{F}_\mathfrak{q}$, and since n is prime to p, $\Phi_n(x)$ has no multiple roots in $\bar{\mathbb{F}}_\mathfrak{q}$, and we can apply Hensel's lemma to conclude that $\Phi_n(x)$ has a root in $\hat{\mathscr{O}}_{L,\mathfrak{q}}$. □

Proof of theorem. By the theorem, $\hat{L}$ is $L_\mathfrak{q}$ for some L and $\mathfrak{q}$, and note that because the extension is unramified, and by the easy facts above, we have $f = f_{\mathfrak{q}/\mathfrak{p}}$. By the lemma, $K_\mathfrak{p}$ contains a primitive $(q-1)$st root of unity, and no higher roots of unity, while $L_\mathfrak{q}$ contains a primitive $(q^f - 1)$st root of unity ζ. It follows that the residue field extension of $K_\mathfrak{p}(\zeta) \subseteq L_\mathfrak{q}$ has degree f, so $[K_\mathfrak{p}(\zeta) : K_\mathfrak{p}] \geqslant f$, and we must have $K_\mathfrak{p}(\zeta) = L_\mathfrak{q}$, as desired.

Of course, any such extension has p unramified because $q^f - 1$ is prime to $p := N(\mathfrak{p} \cap \mathbb{Z})$. □

Corollary 8.7.5. *Let E, L be extensions of a number field K, and $\mathfrak{p} \subseteq \mathscr{O}_K$ a prime ideal unramified in both E and L. Then $\mathfrak{p}$ is unramified in EL.*

Proof. Let $\mathfrak{q}$ be a prime lying over $\mathfrak{p}$ in EL; by the easy facts above, it is enough to check that $\mathfrak{p}$ is unramified in $EL_\mathfrak{q}$ over $K_\mathfrak{p}$. Furthermore, if $\mathfrak{q}_L := \mathfrak{q} \cap \mathscr{O}_L$ and $\mathfrak{q}_E := \mathfrak{q} \cap \mathscr{O}_E$, then the natural inclusions $E_{\mathfrak{q}_E} \to EL_\mathfrak{q}$ and $L_{\mathfrak{q}_L} \to EL_\mathfrak{q}$ induce an isomorphism $E_{\mathfrak{q}_E} L_{\mathfrak{q}_L} \overset{\sim}{\to} EL_{\mathfrak{q}}$ by Exercise 8.2. By the theorem, each of $E_{\mathfrak{q}E}$ and $L_{\mathfrak{q}_L}$ is generated by a root of unity of order prime to $p := N(\mathfrak{p} \cap \mathbb{Z})$, so $E_{\mathfrak{q}_E} L_{\mathfrak{q}_L}$ is also generated by a root of unity of order prime to p, and is hence unramified over $K_{\mathfrak{p}}$, as desired. □

8.8. Decomposition groups as Galois groups

We conclude with a brief discussion of the role local fields play in the Galois case. Thus, we assume that L/K is Galois. Recall that we had defined $D_{\mathfrak{q}/\mathfrak{p}}$ and $I_{\mathfrak{q}/\mathfrak{p}}$, the decomposition and inertia groups of $\mathfrak{q}$ over $\mathfrak{p}$, with the decomposition group being the subgroup of elements of $\mathrm{Gal}(L/K)$ which induce automorphisms of $\mathscr{O}_L/\mathfrak{q}$ over $\mathscr{O}_K/\mathfrak{p}$, and the inertia group being the subgroup acting trivially on $\mathscr{O}_L/\mathfrak{q}$. We had shown that $\mathrm{Gal}((\mathscr{O}_L/\mathfrak{q})/(\mathscr{O}_K/\mathfrak{p})) = D_{\mathfrak{q}/\mathfrak{p}}/I_{\mathfrak{q},\mathfrak{p}}$.

One can view the content of the previous theorem as saying that for unramified extensions, all the information in the extension of local fields is captured in the corresponding extension of finite fields. However, for ramified extensions this is not the case, with the local fields containing the information on the ramification as well. This is made precise in the Galois case by the following.

Theorem 8.8.1. *We have that $L_{\mathfrak{q}}$ is Galois over $K_{\mathfrak{p}}$, and* $\mathrm{Gal}(L_{\mathfrak{q}}/K_{\mathfrak{p}}) = D_{\mathfrak{q}/\mathfrak{p}}$ *under a natural map.*

Proof. By definition, any $\sigma \in D_{\mathfrak{q},\mathfrak{p}}$ fixes $\mathfrak{q}$, and hence preserves $|\cdot|_{\mathfrak{q}}$. Thus σ preserves Cauchy sequences, and induces an endomorphism of $L_{\mathfrak{q}}$ which one easily checks is in fact an automorphism which fixes $K_{\mathfrak{p}}$, hence in $\mathrm{Aut}(L_{\mathfrak{q}}/K_{\mathfrak{p}})$. Since $L_{\mathfrak{q}}$ contains L, this gives an injection $D_{\mathfrak{q}/\mathfrak{p}} \to \mathrm{Aut}(L_{\mathfrak{q}}/K_{\mathfrak{p}})$. But by the easy fact (iii) above,

$$|\,\mathrm{Aut}(L_{\mathfrak{q}}/K_{\mathfrak{p}})| \leqslant [L_{\mathfrak{q}} : K_{\mathfrak{p}}| = e_{\mathfrak{q}/\mathfrak{p}} f_{\mathfrak{q}/\mathfrak{p}} = |D_{\mathfrak{q}/\mathfrak{p}}|,$$

so we must have equality, and it follows that the map is an isomorphism and $L_{\mathfrak{q}}$ is Galois over $K_{\mathfrak{p}}$. □

Exercises

Exercise 8.1. Show that the polynomial $(x^2-2)(x^2-3)(x^2-6)$ has roots modulo p for all p, and in $\mathbb{Q}_p$ for all $p > 3$, but not in $\mathbb{Q}_2$ or $\mathbb{Q}_3$.

Exercise 8.2. Let E, L be extensions of a number field K, and $\mathfrak{p} \subseteq \mathscr{O}_K$ a prime ideal, with $\mathfrak{q} \in \mathscr{O}_{EL}$ lying over $\mathfrak{p}$. If $\mathfrak{q}_L := \mathfrak{q} \cap \mathscr{O}_L$ and $\mathfrak{q}_E := \mathfrak{q} \cap \mathscr{O}_E$, show that the natural inclusions $E_{\mathfrak{q}_E} \to EL_{\mathfrak{q}}$ and $L_{\mathfrak{q}_L} \to EL_{\mathfrak{q}}$ induce an isomorphism $E_{\mathfrak{q}_E} L_{\mathfrak{q}_L} \overset{\sim}{\to} EL_{\mathfrak{q}}$.

CHAPTER 9

Class field theory: an overview

Class field theory is one of the most fundamental breakthroughs in 20th century number theory. It relates the extrinsic data of abelian extensions of number fields to the intrinsic data of ideal classes groups, and in the process yields new information about both. The main statements of class field theory are relatively simple, but the proofs are notoriously difficult, machinery-heavy and unenlightening, so we settle here for describing the statements and a number of applications.

9.1. Frobenius elements and the Artin map

Class field theory relates abelian extensions of a given number field K to certain generalized ideal class groups of K. The fundamental tool for doing this is Frobenius elements, and the Artin map obtained from them.

Recall Definition 3.8.3 if L/K is an abelian extension (i.e., a Galois extension with abelian Galois group), and $\mathfrak{p}$ is a prime of $\mathscr{O}_K$ unramified in $\mathscr{O}_L$, then we defined $\operatorname{Fr}(\mathfrak{p}) \in \operatorname{Gal}(L/K)$ to be the unique element acting as Frobenius on $\mathscr{O}_L/\mathfrak{q}$ over $\mathscr{O}_K/\mathfrak{p}$ for any $\mathfrak{q}$ lying over $\mathfrak{p}$. In particular, it generates $D_{\mathfrak{q}/\mathfrak{p}}$, which has order $e_{\mathfrak{q}/\mathfrak{p}} f_{\mathfrak{q}/\mathfrak{p}}$, so we observe:

Easy Fact 9.1.1. $\operatorname{Fr}(\mathfrak{p}) = 1$ if and only if $\mathfrak{p}$ splits completely in L.

If we denote by $I_{L/K}$ the subgroup of the group of fractional ideals of $\mathscr{O}_K$ generated by prime ideals which are unramified in L (which we know includes all but a finite number of prime ideals of $\mathscr{O}_K$), we can define the Artin map:

Definition 9.1.2. The Artin map $\operatorname{Art} : I_{L/K} \to \operatorname{Gal}(L/K)$ is the unique homomorphism extending the Frobenius element map. That is,

$$\operatorname{Art}(\prod_i \mathfrak{p}_i^{e_i}) := \prod_i \operatorname{Fr}(\mathfrak{p}_i)^{e_i},$$

which is well-defined because $\operatorname{Gal}(L/K)$ is abelian.

The main theorems of class field theory come in two parts: the "reciprocity law" says that given any abelian extension, the Artin map is surjective, and describes the kernel. The "existence theorem" states a converse, that given a group of ideals of the appropriate form, and a subgroup which could potentially be the kernel of an Artin map, then there exists an abelian extension L of K whose Artin map has the desired kernel. To state these results precisely, we will introduce the notion of a modulus and of generalized ideal class groups.

9.2. Generalized ideal class groups

Recall that we denote by M_K the set of all absolute values on K, up to equivalence. These can be broken into M_K^0, the set of non-Archimedean absolute values, which correspond to the primes ideals $\mathfrak{p} \subseteq \mathscr{O}_K$, and M_K^∞, the Archimedean absolute values, which correspond to imbeddings $\sigma : K \to \mathbb{C}$, up to complex conjugation. We define:

Definition 9.2.1. A **modulus** in K is a map $\mathfrak{m} : M_K \to \mathbb{Z}_{\geqslant 0}$ such that:

(i) $\mathfrak{m}(\varpi) = 0$ for all but finitely many $\varpi \in M_K$;
(ii) $\mathfrak{m}(M_K^\infty) \subseteq \{0, 1\}$.
(iii) $\mathfrak{m}(\varpi) = 0$ if ϖ is a complex Archimedean absolute value.

Given a modulus $\mathfrak{m}$, we can restrict $\mathfrak{m}$ to M_K^0 to obtain $\mathfrak{m}_0 : M_K^0 \to \mathbb{Z}_{\geqslant 0}$, which we identify with the ideal of $\mathscr{O}_K$ given by $\prod_{\mathfrak{p} \in M_K^0} \mathfrak{p}^{\mathfrak{m}_0(\mathfrak{p})}$. We also denote by $\mathfrak{m}_1$ the set of prime ideals with $\mathfrak{m}_0(\mathfrak{p}) > 0$. Finally, the restriction of $\mathfrak{m}$ to M_K^∞ gives a map $\mathfrak{m}_\infty : M_K^\infty \to \{0, 1\}$. Given a modulus $\mathfrak{m}$ in K, we write $I_K(\mathfrak{m}) := I_K(\mathfrak{m}_1)$. Let $P_{K,1}(\mathfrak{m})$ be the subgroup of $I_K(\mathfrak{m})$ of principal ideals generated by (α) with:

(i) $\alpha \in \mathscr{O}_K$;
(ii) $\alpha \equiv 1 \pmod{\mathfrak{m}_0}$;
(iii) $\sigma(\alpha) > 0$ for all $\sigma \in \mathfrak{m}_\infty$.

We say that $H \subseteq I_K(S)$ is a **congruence subgroup** for a modulus $\mathfrak{m}$ if $\mathfrak{m}_1 = S$ and $P_{K,1}(\mathfrak{m}) \subseteq H$. We then call $I_K(S)/H$ a **generalized ideal class group** of K, with modulus $\mathfrak{m}$.

Finally, we define $I_K(\mathfrak{m})/P_{K,1}(\mathfrak{m})$ to be the **ray class group** for the modulus $\mathfrak{m}$.

Thus, every generalized ideal class group is a quotient of a ray class group.

Example 9.2.2. When $\mathfrak{m}$ is the zero map, we have $I_K(\mathfrak{m})$ the entire group of fractional ideals, and $P_{K,1}(\mathfrak{m})$ the entire group of principal fractional ideals, so $I_K(\mathfrak{m})/P_{K,1}(\mathfrak{m})$ is the standard ideal class group, and generalized ideal class groups with modulus $\mathfrak{m}$ correspond to quotients of the usual ideal class group.

Example 9.2.3. We consider the case of $\mathbb{Q}$. There is only one real imbedding of $\mathbb{Q}$, so we write $\mathfrak{m} = m$ or $m\infty$ for some $m \in \mathbb{N}$, depending on whether $\mathfrak{m}(\sigma) = 0$ or 1 for the unique $\sigma : \mathbb{Q} \to \mathbb{R}$ in $M_\mathbb{Q}^\infty$. In either case, we have that $I_\mathbb{Q}(\mathfrak{m})$ is the set of principal ideals of the form $\left(\frac{x}{y}\right)$ with $x, y \in \mathbb{Z}$, and $(x, m) = (y, m) = 1$. We observe that we have a surjective map from $I_\mathbb{Q}(\mathfrak{m})$ to $(\mathbb{Z}/m\mathbb{Z})^*$. Indeed, since the generator of a fractional ideal of $\mathbb{Q}$ is unique up to ± 1, if we always require $x, y > 0$, every element of $I_\mathbb{Q}(\mathfrak{m})$ has a unique generator, and since x, y are units in $\mathbb{Z}/m\mathbb{Z}$, we can send $\left(\frac{x}{y}\right)$ to $\frac{x}{y}$ in $(\mathbb{Z}/m\mathbb{Z})^*$.

The kernel of this map is precisely the principal ideals generated by $\frac{x}{y}$ with $x \equiv y \pmod{m}$, and $x, y \in \mathbb{N}$. This is the same as $P_{K,1}(\mathfrak{m})$ when $\mathfrak{m} = m\infty$, so in this case the ray class group is simply $(\mathbb{Z}/m\mathbb{Z})^*$. However, when $\mathfrak{m} = m$, the situation is slightly different: $P_{K,1}(\mathfrak{m})$ consists of all principal ideals generated by $\frac{x}{y}$ with $x, y \equiv 1 \pmod{m}$, or equivalently, the principal ideals which can be generated by $\frac{x}{y}$ with $x \equiv y \pmod{m}$ (given such x, y, if $z \in \mathbb{N}$ is such that $xz \equiv 1 \pmod{m}$, we can replace x, y with xz, yz). We then see that a fractional ideal in $I_\mathbb{Q}(\mathfrak{m})$ is in $P_{K,1}(\mathfrak{m})$ if and only if it maps to ± 1 in $(\mathbb{Z}/m\mathbb{Z})^*$, since the two possible generators are $\frac{x}{y}$ and $-\frac{x}{y}$, so it is in $P_{K,1}(\mathfrak{m})$ if and only if either $x \equiv y \pmod{m}$ or $-x \equiv y \pmod{m}$.

9.3. The main theorems

We are now state preliminary versions of the main theorems of class field theory. We emphasize that although the statements are simple, the proofs are quite deep, and would require the better part of a semester-long course to cover in full.

THEOREM 9.3.1. *(Artin reciprocity) Let L/K be an abelian extension of number fields, and S the set of primes of K ramified in L. Then there exists a modulus $\mathfrak{m}$ such that $\mathfrak{m}_1 = S$, and the Artin map*

$$I_K(\mathfrak{m}) \to \mathrm{Gal}(L/K)$$

is surjective, with kernel a congruence subgroup for $\mathfrak{m}$.

We will see applications of this theorem that will justify, at least to some degree, the terminology “reciprocity”.

THEOREM 9.3.2. *(Existence theorem) Let K be a number field, and $\mathfrak{m}$ a modulus in K. For any congruence subgroup H for $\mathfrak{m}$, there exists a unique number field L which is an abelian extension of K, which is ramified only at primes contained in $\mathfrak{m}_1$, and for which the Artin map*

$$I_K(\mathfrak{m}) \to \mathrm{Gal}(L/K)$$

has kernel exactly H.

Thus we obtain a bijection between the intrinsic data of generalized ideal class groups of K, and the extrinsic data of abelian extensions of K.

9.4. Class groups revisited

If $\mathfrak{m}, \mathfrak{m}'$ are two moduli in K, we say $\mathfrak{m}|\mathfrak{m}'$ if $\mathfrak{m}(\varpi) \leqslant \mathfrak{m}'(\varpi)$ for all $\varpi \in M_K$; this is equivalent to the condition that $\mathfrak{m}_0|\mathfrak{m}'_0$ and $\mathfrak{m}_\infty \subseteq \mathfrak{m}'_\infty$. We then define $\gcd(\mathfrak{m}, \mathfrak{m}')$ in the analogous way. We also write $(\mathscr{O}_K/\mathfrak{m})^*$ to denote the group $(\mathscr{O}_K/fm_0)^* \oplus \{\pm 1\}^{|\mathfrak{m}_\infty|}$.

We write $K_\mathfrak{m}$ for the set of elements $z \in K^*$ such that $|z|_\mathfrak{p} = 1$ for all $\mathfrak{p} \in \mathfrak{m}_1$, and we have a natural map $K_\mathfrak{m} \to (\mathscr{O}_K/\mathfrak{m})^*$ sending z to its image modulo $\mathfrak{m}_0$, together with the sign of z under each real imbedding of $\mathfrak{m}_\infty$. We denote the kernel by $K_{\mathfrak{m},1}$, and note that $P_{K,1}(\mathfrak{m})$ is precisely the set of principal ideals generated by $K_{\mathfrak{m},1}$. We also write $U_\mathfrak{m}$ for the image of $\mathscr{O}_K^*$ in $(\mathscr{O}_K/\mathfrak{m})^*$.

If we have $\mathfrak{m}|\mathfrak{m}'$, we have a natural map $(\mathscr{O}_K/\mathfrak{m}')^* \to (\mathscr{O}_K/\mathfrak{m})$ obtained by modding out the first summand, and forgetting coordinates corresponding to elements of $\mathfrak{m}'_\infty$ not in $\mathfrak{m}_\infty$.

We discuss the relationship between ideal class groups of different moduli. Let $\mathfrak{m}, \mathfrak{m}'$ be moduli in K with $\mathfrak{m}|\mathfrak{m}'$. Our first claim is that the natural map $I_K(\mathfrak{m}') \to I_K(\mathfrak{m})$ gives a well-defined map

$$I_K(\mathfrak{m}')/P_{K,1}(\mathfrak{m}') \to I_K(\mathfrak{m})/P_{K,1}(\mathfrak{m}).$$

Indeed, this follows immediately from the definitions. More substantively, we have:

PROPOSITION 9.4.1. *For $\mathfrak{m}|\mathfrak{m}'$, the natural map*

$$\varphi : I_K(\mathfrak{m}')/P_{K,1}(\mathfrak{m}') \to I_K(\mathfrak{m})/P_{K,1}(\mathfrak{m})$$

is surjective, with finite kernel of order

$$|(\mathscr{O}_K/\mathfrak{m}')^*||U_\mathfrak{m}|/|(\mathscr{O}_K/\mathfrak{m})^*||U_{\mathfrak{m}'}|.$$

To prove this, we will need a more fundamental result:

Theorem 9.4.2. *(Approximation theorem) Given $|\cdot|_1, \dots, |\cdot|_n \in M_K$ distinct, and $x_1, \dots, x_n \in K$, then for any $\epsilon > 0$ there exists a $y \in K$ such that $|y - x_i|_i < \epsilon$ for all i.*

The proof is not difficult, but we refer to [**7**, Thm. 1, p. 35].

Proof of proposition. We first show surjectivity. We suppose we have a fractional ideal $I \in I_K(\mathfrak{m})$, and wish to show that for some $z \in P_{K,1}(\mathfrak{m})$, we have $zI \in I_K(\mathfrak{m}')$. Let $\mathfrak{p}_1, \dots, \mathfrak{p}_m$ be the primes of $\mathfrak{m}_0'$ not in $\mathfrak{m}$; we thus want zI be to prime to each $\mathfrak{p}_i$, which is equivalent to saying that (z) should have the same number of powers of $\mathfrak{p}_i$ as I^{-1}, for all i. Choose $x_1, \dots, x_m$ in K with (x_i) having the same number of powers of $\mathfrak{p}_i$ as I^{-1}; we thus want $|z|_{\mathfrak{p}_i} = |x_i|_{\mathfrak{p}_i}$ for all i. Now, let $|\cdot|_i = |\cdot|_{\mathfrak{p}_i}$ for $i = 1, \dots, m$, and then set $|\cdot|_{m+1}, \dots, |\cdot|_n$ to be the remaining absolute values of $\mathfrak{m}$. Set $x_{m+1}, \dots, x_n$ equal to 1. Now, let z be the element given to us by the approximation theorem for ϵ small. For $i = 1, \dots, m$ we have $|x_i - z|_i < \epsilon$ implies $|x_i|_i = |z|_i$ as long as $\epsilon < |x_i|$, so $zI \in I_K(\mathfrak{m}')$. Furthermore, because $|x_i - z|_i < \epsilon$ for $i > m$ corresponding to primes of $\mathfrak{m}_0$, we have $z \equiv 1 \pmod{\mathfrak{m}_0}$ as long as ϵ is small relative to $1/N(\mathfrak{m}_0)$. Finally, the remaining i correspond to real absolute values, and we have $|x_i - z|_i < \epsilon$ which implies $z > 0$ as long as $\epsilon < 1$. We thus conclude that $z \in P_{K,1}(\mathfrak{m})$, as desired, completing the proof of surjectivity.

It is clear that $\ker \varphi = (P_{K,1}(\mathfrak{m}) \cap I_K(\mathfrak{m}'))/P_{K,1}(\mathfrak{m}')$. We consider the natural map

$$r : (\mathscr{O}_K/\mathfrak{m}')^* \to (\mathscr{O}_K/\mathfrak{m})^*.$$

Using the approximation theorem, we see that $K_{\mathfrak{m}'}$ surjects onto $(\mathscr{O}_K/\mathfrak{m}')^*$, so we can view $\ker r$ as $K_{\mathfrak{m},1} \cap K_{\mathfrak{m}'}/K_{\mathfrak{m}',1}$. One also checks that we can write $I_K(\mathfrak{m}') \cap P_{K,1}(\mathfrak{m})$ as $K_{\mathfrak{m}'} \cap K_{\mathfrak{m},1}/\mathscr{O}_K^* \cap K_{\mathfrak{m},1}$. We thus obtain a diagram

$$\begin{tikzcd}
\mathscr{O}_K^* \cap K_{\mathfrak{m}',1} \arrow[r, hook] \arrow[d, hook] & K_{\mathfrak{m}',1} \arrow[r, two heads] \arrow[d, hook] & P_{K,1}(\mathfrak{m}') \arrow[d, hook] \\
\mathscr{O}_K^* \cap K_{\mathfrak{m},1} \arrow[r, hook] \arrow[d, "="] & K_{\mathfrak{m},1} \cap K_{\mathfrak{m}'} \arrow[r, two heads] \arrow[d] & I_K(\mathfrak{m}') \cap P_{K,1}(\mathfrak{m}) \arrow[d] \\
\mathscr{O}_K^* \cap K_{\mathfrak{m},1} \arrow[r] & \ker r \arrow[r, two heads] & \ker \varphi
\end{tikzcd},$$

where everything is exact except for the left column, and *a priori* the bottom row. However, a diagram chase confirms that the bottom row is also exact. Now, the kernel of the lower left map is $\mathscr{O}_K^* \cap K_{\mathfrak{m}',1}$, and we observe that

$$\mathscr{O}_K^* \cap K_{\mathfrak{m},1}/\mathscr{O}_K^* \cap K_{\mathfrak{m}',1} \cong (\mathscr{O}_K^*/\mathscr{O}_K^* \cap K_{\mathfrak{m}',1})/(\mathscr{O}_K^*/\mathscr{O}_K^* \cap K_{\mathfrak{m},1}) = U_{\mathfrak{m}'}/U_{\mathfrak{m}},$$

so we conclude that $\ker \varphi \cong (\ker r)/(U_{\mathfrak{m}'}/U_{\mathfrak{m}})$, giving the desired formula. □

Using the proposition with $\mathfrak{m}$ equal to the zero map, we conclude:

Corollary 9.4.3. *Every generalized ideal class group is finite.*

We conclude with a discussion of the equivalences obtained by considering class groups for different moduli.

Corollary 9.4.4. *Given an ideal class group of modulus $\mathfrak{m}$, corresponding to a congruence subgroup H, and given $\mathfrak{m}'$ with $\mathfrak{m}|\mathfrak{m}'$, then the preimage H' of H in $I_K(\mathfrak{m}')$ is a congruence subgroup for $\mathfrak{m}'$, and the corresponding ideal class groups are naturally isomorphic.*

Proof. The statement that H' is a congruence subgroup is immediate from the definitions, while the isomorphism of class groups follows from the surjectivity of the proposition. □

The proof of the following is just a definition-chase:

Easy Fact 9.4.5. Given $\mathfrak{m}'|\mathfrak{m}$, if H is a congruence subgroup for $\mathfrak{m}$, and H contains $I_K(\mathfrak{m}) \cap P_{K,1}(\mathfrak{m}')$, and if we write $H' = HP_{K,1}(\mathfrak{m}')$, then we have $H = I_K(\mathfrak{m}) \cap H'$.

This motivates the following definition:

Definition 9.4.6. Let H be a congruence subgroup for a modulus $\mathfrak{m}$, and $\mathfrak{m}'|\mathfrak{m}$. We say that H is **defined** for $\mathfrak{m}'$ if H contains $I_K(\mathfrak{m}') \cap P_{K,1}(\mathfrak{m})$.

Corollary 9.4.7. *Every congruence subgroup has a smallest modulus for which it is defined.*

Proof. Suppose that H is a congruence subgroup for $\mathfrak{m}$, and $\mathfrak{m}', \mathfrak{m}''|\mathfrak{m}$ are such that H is defined for $\mathfrak{m}'$ and $\mathfrak{m}''$. It suffices to see that H is defined for $\gcd(\mathfrak{m}', \mathfrak{m}'')$. This is another application of the approximation theorem, which we leave to the reader. □

Definition 9.4.8. The smallest modulus for which a congruence subgroup H is defined is called the **conductor** of H, or of the associated generalized ideal class group.

9.5. The main theorems revisited

We now state in more detail the main theorems of class field theory. However, we need two more preliminary definitions:

Definition 9.5.1. Given an extension L/K of number fields, and a fractional ideal I of L, we define the fractional ideal $N_{L/K}(I)$ of K as follows: write $I = \prod_i \mathfrak{q}_i^{e_i}$, let $\mathfrak{p}_i = \mathfrak{q}_i \cap \mathscr{O}_K$, and let $N_{L/K}(I) = \prod_i \mathfrak{p}_i^{e_i f_i}$.

Definition 9.5.2. Let $\sigma : K \to \mathbb{R}$ be a real imbedding. We say that σ is **ramified** in an extension L if σ extends to a complex imbedding of L.

We now state the main theorems of class field theory with more precision than before.

Theorem 9.5.3. *(Artin reciprocity) Let L/K be an abelian extension of number fields, and S the set of primes of K ramified in L (including real imbeddings). Then there exists a modulus $\mathfrak{m}$ such that $\mathfrak{m}_1 = S \cap M_K^0$, $\mathfrak{m}_\infty = M_K^\infty$ and the Artin map*

$$\mathrm{Art} : I_K(\mathfrak{m}) \to \mathrm{Gal}(L/K)$$

is surjective, with kernel a congruence subgroup for $\mathfrak{m}$. More precisely, the kernel is equal to the subgroup of $I_K(\mathfrak{m})$ generated by $P_{K,1}(\mathfrak{m})$ and by $N_{L/K}(I)$ as I ranges over fractional ideals of L which are prime to $\mathfrak{m}_0$. Finally, we may set $\mathfrak{m}$ to be the conductor of ker Art*, which is the smallest modulus for which* ker Art *is a congruence subgroup, and which we call also the* **conductor** *of L/K.*

This is [**7**, Thm. 3, p. 205], together with the discussion on pp. 146-7 on admissible cycles.

Theorem 9.5.4. *(Existence theorem) Let K be a number field, and $\mathfrak{m}$ a modulus in K. For any congruence subgroup H for $\mathfrak{m}$, there exists a unique number field L which is an abelian extension of K, which has conductor dividing $\mathfrak{m}$ (in particular, it is ramified only at primes contained in $\mathfrak{m}_1 \cup \mathfrak{m}_\infty$), and for which the Artin map*

$$I_K(\mathfrak{m}) \to \mathrm{Gal}(L/K)$$

has kernel exactly H. In fact, the conductor of L is equal to the conductor of H.

Definition 9.5.5. Given any modulus $\mathfrak{m}$ for K, **ray class field** $K_{\mathfrak{m}}$ is the abelian extension obtained from $I_K(\mathfrak{m})/P_{K,1}(\mathfrak{m})$ by the existence theorem.

9.6. Partial orders

Now we have seen that we can consider a congruence subgroup for a modulus $\mathfrak{m}$ in terms of a modulus $\mathfrak{m}'$ with $\mathfrak{m}|\mathfrak{m}'$, without affecting the associated generalized ideal class group. We thus have a natural partial ordering on generalized ideal class groups, determined by whether there exists a common modulus in which one congruence subgroup is contained in the other. Of course, we also have the inclusion partial ordering on the collection of abelian extensions of a given number field, and we now see that the two partial orderings correspond to one another under the bijection obtained by class field theory.

We have:

Easy Fact 9.6.1. Let $E/L/K$ be number fields, with E (and hence L) abelian over K. Then for any $\mathfrak{p} \in K$ which is unramified in E, we have that $\mathrm{Fr}_{L/K}(\mathfrak{p})$ is simply the restriction of $\mathrm{Fr}_{E/K}(\mathfrak{p})$ to $\mathrm{Gal}(L/K)$.

Corollary 9.6.2. *Let E, L be finite abelian extensions of K. Then $L \subseteq E$ if and only if there exists a modulus $\mathfrak{m}$ in K such that* $\ker \mathrm{Art}$ *for E and L, which we denote by H_E and H_L respectively, are congruence subgroups for $\mathfrak{m}$, and $H_E \subseteq H_L$. Moreover, it suffices to take $\mathfrak{m}$ to be the least modulus divisible by the conductor of both fields, which will be the conductor of E if $L \subseteq E$.*

Proof. If $L \subseteq E$, we apply Artin reciprocity to obtain a modulus $\mathfrak{m}$ such that H_E is a congruence subgroup for $\mathfrak{m}$; Then, the easy fact and multiplicativity of the Artin map gives us that the Artin map for L is obtained by restricting the Artin map for E. So if $I \in \ker \mathrm{Art}_{E/K}$, then we must have $I \in \ker \mathrm{Art}_{L/K}$ as well, and $H_E \subseteq H_L$; in particular, H_L is also a congruence subgroup for $\mathfrak{m}$.

Conversely, given $\mathfrak{m}$ as in the statement, let G be the image of H_L under the Artin map for E over K, and let $\tilde{L} \subseteq E$ be the fixed field of G. We claim that $\ker \mathrm{Art}_{L/K} = \ker \mathrm{Art}_{\tilde{L}/K}$. Indeed, if $I \in \ker \mathrm{Art}_{L/K} = H_L$, then $\mathrm{Art}_{\tilde{L}/K}(I) = \mathrm{Art}_{E/K}(I)|_{\tilde{L}} \in G|_{\tilde{L}} = 1$ by definition. Conversely, if $I \in \ker \mathrm{Art}_{\tilde{L}/K}$ then by Galois theory $\mathrm{Art}_{E/K}(I) \in G = \mathrm{Art}_{E/K}(H_K)$, so $I \in H_L + H_E = H_L$ because $H_E \subseteq H_L$ by hypothesis. By the uniqueness in the existence theorem, it follows that $L = \tilde{L}$, and $L \subseteq E$. □

Corollary 9.6.3. *Let S be a collection of moduli in K such that for any modulus $\mathfrak{m}$, there is a $\mathfrak{m}' \in S$ with $\mathfrak{m}|\mathfrak{m}'$. Then every abelian extension is contained in a ray class field $K_{\mathfrak{m}}$ with $\mathfrak{m} \in S$.*

PROOF. Let L be an abelian extension of K. By Artin reciprocity, there exists a modulus $\mathfrak{m}$ for which $\ker \mathrm{Art}_{L/K}$ is a congruence subgroup. Let $\mathfrak{m}' \in S$ be such that $\mathfrak{m}|\mathfrak{m}'$; then we have that $\ker \mathrm{Art}_{L/K}$ for the modulus $\mathfrak{m}'$ is also a congruence subgroup, and necessarily contains $P_{K,1}(\mathfrak{m}')$. It follows from the previous corollary that $L \subseteq K_{\mathfrak{m}'}$. □

9.7. Examples and applications

EXAMPLE 9.7.1. We consider the case of $\mathbb{Q}(\zeta_n)$ over $\mathbb{Q}$. For notational simplicity, we assume that $n \not\equiv 2 \pmod 4$, since otherwise we could use $\mathbb{Q}(\zeta_{n/2})$ instead. Since $\mathbb{Q}$ is real and $\mathbb{Q}(\zeta_n)$ is complex, the infinite prime is ramified, so we will work with the modulus $\mathfrak{m} = n\infty$. We know that for p prime to n, $\mathrm{Fr}(p) = [p] \in (\mathbb{Z}/n\mathbb{Z})^* \cong \mathrm{Gal}(\mathbb{Q}(\zeta_n)/\mathbb{Q})$. Note that here we have chosen the positive generator of the ideal (p). Thus, $\mathrm{Art}_{\mathbb{Q}(\zeta_n)/\mathbb{Q}}((x)) = [x] \in (\mathbb{Z}/n\mathbb{Z})^*$ if x is the positive generator of (x), and the kernel is precisely the set of principal ideals generated by x such that $x \equiv 1 \pmod n$ and $x > 0$, which is the definition of $P_{K,1}(\mathfrak{m})$. Thus, we see that $\mathbb{Q}(\zeta_n)$ is the ray class field $\mathbb{Q}_{\mathfrak{m}}$.

Putting the example together with the previous corollary, we immediately find:

THEOREM 9.7.2. *(Kronecker-Weber) Every abelian extension of $\mathbb{Q}$ is contained in a cyclotomic extension.*

Of course, there are proofs of this theorem not requiring class field theory, but the fact that it falls out immediately as a special case of far more general results illustrates some of the power of the theory.

We conclude by showing that quadratic reciprocity follows easily from Artin reciprocity, perhaps somewhat justifying the terminology.

THEOREM 9.7.3. *We have $\left(\frac{*p}{q}\right) = \left(\frac{q}{p}\right)$, where $* = (-1)^{\frac{p-1}{2}}$.*

PROOF. Consider the fields $\mathbb{Q}(\zeta_p)$ and $\mathbb{Q}(\sqrt{*p})$; by our earlier work on fixed fields of Dirichlet characters, we know that $\mathbb{Q}(\sqrt{*p})$ has discriminant $*p$ and is therefore contained in $\mathbb{Q}(\zeta_p)$. We also know that $\mathbb{Q}(\zeta_p)$ is the ray class field for $\mathfrak{m} = p\infty$.

We consider the Artin map for $\mathbb{Q}(\sqrt{*p})$ and the modulus $\mathfrak{m}$. We see directly from the definitions that evaluated on primes q, it is simply $\left(\frac{*p}{q}\right)$, since the quadratic character determines whether q splits in $\mathbb{Q}(\sqrt{*p})$.

On the other hand, from class field theory the Artin map has kernel a congruence subgroup, and in particular containing $P_{K,1}(\mathfrak{m})$. Another way to say this is that for $\mathfrak{m}$, $\mathrm{Art}_{\mathbb{Q}(\sqrt{*p})}$ factors through $\mathrm{Art}_{\mathbb{Q}(\zeta_p)}$, which factors through $\mathrm{Gal}(\mathbb{Q}(\zeta_p)/\mathbb{Q}) \cong (\mathbb{Z}/p\mathbb{Z})^*$. Thus $\mathrm{Art}_{\mathbb{Q}(\sqrt{*p})}$ induces a (non-trivial) map $(\mathbb{Z}/p\mathbb{Z})^*$ to (± 1), and since quadratic characters modulo p, we conclude that this map is $\left(\frac{q}{p}\right)$. □

9.8. The Hilbert class field

We now discuss the important example of the Hilbert class field, which we obtain by considering the case of the classical ideal class group of K. We have:

DEFINITION 9.8.1. Given a number field K, the **Hilbert class field** is the abelian extension H_K of K such that the Artin map of H_K/K induces an isomorphism between the ideal class group of K and $\mathrm{Gal}(H_K/K)$.

It follows immediately from the definitions that H_K has the property that a prime ideal $\mathfrak{p}$ of $\mathscr{O}_K$ splits completely in $\mathscr{O}_{H_K}$ if and only it is principal.

Corollary 9.8.2. *Let K be a number field. Then there exists a monic irreducible polynomial $f(x) \in \mathscr{O}_K[x]$ such that for all primes $\mathfrak{p}$ of $\mathscr{O}_K$ with $\mathfrak{p}$ not dividing* $\operatorname{disc} f(x)$*, we have that $\mathfrak{p}$ is principal if and only if $f(x)$ has a root modulo $\mathfrak{p}$.*

Proof. Let $f(x)$ be the minimal polynomial for an element $\alpha \in \mathscr{O}_{H_K}$ generating H_K over K. We know that if $\mathfrak{p}$ is prime to $\operatorname{disc} f(x)$, we have that α generates $\mathscr{O}_{H_K,\mathfrak{p}}$ over $\mathscr{O}_{K,\mathfrak{p}}$ and (using also that H_K over K is Galois) that $\mathfrak{p}$ splits completely if and only if $f(x)$ has a root modulo $\mathfrak{p}$. □

We also have:

Theorem 9.8.3. *H_K is the maximal everywhere unramified abelian extension of K.*

Proof. By the statement of the existence theorem, since we are working with the trivial modulus, H_K is everywhere unramified over K. Moreover, any abelian extension which is everywhere unramified has trivial conductor, so corresponds to an ideal class group which is a quotient of the standard one, and by our result on orderings from last time, is a subfield of H_K. □

In particular, we see that for every number field K, there are only finitely many abelian everywhere unramified extensions of K. Recall that when $K = \mathbb{Q}$, we had by the Minkowski estimates that there are no unramified extensions of any kind, even without considering ramification at the infinite prime. However, this is a highly atypical situation, and Golod and Shafarevich showed even for many quadratic fields, if one drops the abelian requirement, there are infinitely many everywhere unramified extensions, which can be obtained by taking successive Hilbert class fields. Such towers are called class field towers, and now discuss them in more detail.

9.9. Class field towers

For a number field K, we will write $H_K^{(i)}$ to denote the ith Hilbert class field of K, so that $H_K^{(i)}$ is K for $i = 0$, and $H_{H_K^{(i-1)}}$ for $i > 0$. This is called the **class field tower** of K. Clearly, if $H_K^{(i)}$ never stabilizes, then K is contained in infinitely many (no longer abelian) everywhere unramified extensions. It seems that the case that $H_K^{(i)}$ stabilizes is quite special, as it only happens when some $H_K^{(i)}$ has trivial class group. However, we see that the situation is not quite as special as it seems. We first note:

Lemma 9.9.1. *Let L be any extension of K with trivial class group. Then $H_K \subseteq L$.*

Proof. We first claim that $H_K L$ is abelian and unramified over L. The abelian assertion follows from Galois theory, as $H_K L$ will still be Galois over L, with Galois group a subgroup of $\operatorname{Gal}(H_K/K)$. The unramified assertion follows by considering local fields just as in the argument for Corollary 2.4 of lecture 31 that the compositum of two unramified fields is unramified, except that we also need to check that the infinite primes remain unramified, which is more or less the statement that if

K is real, then H_K is also real. Hence $H_K L \subseteq H_L$, and since L has trivial class group, $H_L = L$ and it follows that $H_K \subseteq L$. □

We conclude:

Corollary 9.9.2. *The class field tower of K stabilizes if and only if K is contained in some L with trivial class number.*

Proof. If the tower stabilizes, then some $H_K^{(i)}$ itself has trivial class number. Conversely, by inductively applying the lemma, we see that every $H_K^{(i)}$ is contained in L, so since L is a finite extension of K, the tower must stabilize at some finite i. □

Golod and Shafarevich studied class fields towers, and showed:

Theorem 9.9.3. *When sufficiently many primes of $\mathbb{Q}$ are ramified in K, then $H_K^{(i)}$ does not stabilize.*

This is not an abstract existence result. For instance, they showed that $\mathbb{Q}(\sqrt{-2 \cdot 3 \cdot 5 \cdot 7 \cdot 11 \cdot 13})$ has an infinite class field tower, as well as $\mathbb{Q}(\sqrt{d})$ where d is square-free with at least 8 distinct prime factors.

9.10. Further applications

Although it is certainly not the case that H_K has trivial class number in general, there is a theorem conjectured by Hilbert, reduced to group theory by Artin, and proved by Furtwangler, which states:

Theorem 9.10.1. *Every ideal of $\mathscr{O}_K$ becomes principal in $\mathscr{O}_{H_K}$.*

The idea is to reduce to a group-theoretic argument by considering $H_K^{(2)}/H_K/K$, since one can measure whether an ideal in H_K is principal by whether it is in the kernel of the Artin map of $H_K^{(2)}$ over H_K. One then proves the theorem using a group-theoretic tool called the transfer.

The reason that H_K doesn't necessarily have trivial class number is that if $\mathfrak{p} \in \mathscr{O}_K$ factors in $\mathscr{O}_{H_K}$, its factors may not be principal, but may have principal product.

Thus, the relationship between the class number of a field K and some extension L is rather subtle; some ideals can become principal in L, while new non-principal classes can arise in L. However, in one case we have a very nice relationship:

Theorem 9.10.2. *Let L be an extension of K such that $L \cap H_K = K$ (for instance, if some prime $\mathfrak{p}$ of K is totally ramified in L). Then $h_K | h_L$.*

Proof. First note that any $\mathfrak{p}$ being totally ramified implies that $L \cap H_K = K$, since $\mathfrak{p}$ is totally ramified in any subfield of L, and unramified in any subfield of H_K. Next, as in the earlier lemma, we have $H_K L$ abelian and unramified over L, so we have

$$h_K = [H_K : K] = [H_K L : L] | [H_L : L] = h_L.$$

□

We consider this theorem in a special case:

Corollary 9.10.3. *Let $L = \mathbb{Q}(\zeta_n)$, and $K = \mathbb{Q}(\zeta_n + \bar{\zeta}_n)$ be the maximal real subfield. Then $h_K | h_L$.*

PROOF. Since K is totally real and L is totally complex, and $[L : K] = 2$, every infinite prime of K is totally ramified in L, and one checks that the argument of the theorem works equally well in this context. □

CHAPTER 10

Ring class fields and $p = x^2 + ny^2$

Recall that we had shown that given $n \in \mathbb{N}$, and p a prime number, then $p = x^2 + ny^2$ if and only if $(p) = \mathfrak{p}\bar{\mathfrak{p}}$ in $\mathbb{Z}[\sqrt{-n}]$, and $\mathfrak{p}$ is principal. The first condition is easy to analyze: we saw that at least for p not dividing $2n$, $(p) = \mathfrak{p}\bar{\mathfrak{p}}$ if and only if $\left(\frac{-n}{p}\right) = 1$, which in turn can be described very concretely for any given n by quadratic reciprocity. The condition that $\mathfrak{p}$ be principal is subtler, and requires class field theory.

In the case that $\mathbb{Z}[\sqrt{-n}] = \mathscr{O}_{\mathbb{Q}(\sqrt{-n})}$, we know that $\mathfrak{p}$ is principal if and only if it splits completely in the Hilbert class field of $\mathbb{Q}(\sqrt{-n})$, and this lets us understand the situation concretely, at least from an abstract point of view. However, for the general case, we will need to see how class field theory interacts with orders in imaginary quadratic fields. This is treated by the theory of ring class fields.

10.1. Ring class groups and ring class fields

The main result we will prove is the following:

Theorem 10.1.1. *Given $n \in \mathbb{N}$, there exists an irreducible monic polynomial $f_n(t) \in \mathbb{Z}[t]$ such that for any prime p not dividing $2n \operatorname{disc} f_n(t)$, we can write $p = x^2 + ny^2$ for some $x, y \in \mathbb{Z}$ if and only if $(-n/p) = 1$ and $f_n(t)$ has a root modulo p.*

Recall the notion of orders:

Definition 10.1.2. An **order** $\mathscr{O}$ of **conductor** f in an imaginary quadratic field K is the unique subring of $\mathscr{O}_K$ of index f, given explicitly as $\mathscr{O} = \{x + yf\omega : x, y \in \mathbb{Z}\}$, where ω is such that $\mathscr{O}_K = \mathbb{Z}[\omega]$.

It turns out that the key to dealing with ideals in orders is to restrict to ideals prime to the conductor; i.e., $I \subseteq \mathscr{O}$ such that $I + (f) = \mathscr{O}$.

Definition 10.1.3. The **ideal class group** of $\mathscr{O}$ is the group obtained by considering ideals prime to f modulo principal ideals prime to f.

It will follow that this is a group from the following:

Proposition 10.1.4. *Given $\mathscr{O}$ of conductor f in $\mathscr{O}_K$, the maps $I \mapsto I\mathscr{O}_K$ and $I \mapsto I \cap \mathscr{O}$ are mutually inverse on ideals prime to f, and induce a multiplicative bijection between ideals of $\mathscr{O}$ prime to f and ideals of $\mathscr{O}_K$ to prime to f.*

The image in $\mathscr{O}_K$ of the principal ideals of $\mathscr{O}$ prime to f is exactly $P_{K,\mathbb{Z}}(f)$, the group of principal ideals of $\mathscr{O}_K$ generated by x with $x \equiv n \pmod{f\mathscr{O}_K}$, for some $n \in \mathbb{Z}$ (and relatively prime to f).

See [**3**, Prop. 7.20, Prop. 7.22] for the proof.

We can thus define:

Definition 10.1.5. If $\mathcal{O}$ is the order in $\mathcal{O}_K$ of conductor f, we define the **ring class group** to be $I_K(f)/P_{K,\mathbb{Z}}(f)$, which is naturally isomorphic to the group of ideals of $\mathcal{O}$ prime to f modulo principal ideals of $\mathcal{O}$ prime to f. We define the **class number** $h_{\mathcal{O}}$ of $\mathcal{O}$ to be the order of the ring class group.

The proposition gives a natural isomorphism between the ideal class group of $\mathcal{O}$ and the ring class group of $\mathcal{O}$.

Remark 10.1.6. Note that having trivial class group does not imply that $\mathcal{O}$ is a PID. In fact, for $f > 1$, $\mathcal{O}$ is never Dedekind, and in particular never a PID.

Observe that $P_{K,\mathbb{Z}}(f)$ is a congruence subgroup of conductor f (note that there is no real infinite place in K, so a conductor $\mathfrak{m}$ is determined by $\mathfrak{m}_0$). We can therefore make the following definition:

Definition 10.1.7. Given $\mathcal{O}$, the **ring class field** $K_{\mathcal{O}}$ is the abelian extension of K associated by the existence theorem of class field theory to the ring class group of $\mathcal{O}$.

Relating splitting of primes in $\mathcal{O}$ to splitting of primes in $\mathcal{O}_K$, it is not hard to see that $K_{\mathcal{O}}$ has the following property generalizing the Hilbert class field:

Theorem 10.1.8. *A prime $\mathfrak{p}$ of $\mathcal{O}$ which is prime to f is principal if and only if it splits completely in $K_{\mathcal{O}}$.*

One then shows that $\mathbb{Q}(\sqrt{-n})_{\mathbb{Z}[\sqrt{-n}]}$ is Galois over $\mathbb{Q}$, and using Proposition 2.1.1, it is then straightforward to deduce:

Theorem 10.1.9. *Fix $n > 0$. Then for p not dividing $2n$, we have*

$$p = x^2 + ny^2 \Leftrightarrow p \text{ splits completely in } \mathbb{Q}(\sqrt{-n})_{\mathbb{Z}[\sqrt{-n}]}.$$

From here, one finishes the proof of Theorem 10.1.1 by showing that the theorem is satisfied by setting $f_n(x)$ to be the minimal polynomial of a primitive element for $\mathbb{Q}(\sqrt{-n})_{\mathbb{Z}[\sqrt{-n}]} \cap \mathbb{R}$ over $\mathbb{Q}$. See [**3**, §9 A].

Remark 10.1.10. In fact, the condition that p is prime to disc $f_n(x)$ in Theorem 10.1.1 can be dropped if $f_n(x)$ is chosen appropriately; this is a consequence of the explicit methods discussed below.

The theory we have discussed thus far is sufficient to compute a number of examples. The basic idea is to use that we can compute the degree of $f_n(x)$ as a class number, and to use our knowledge of where the ring class field is ramified, to reduce down to a finite set of possibilities for the ring class field, and then check one by one whether they agree with the theorem. For instance:

Example 10.1.11. In the theorem, for $n = 27$, the polynomial $f_n(x)$ may be taken to be $x^3 - 2$, while for $n = 64$, we may take $f_n(x) = x^4 - 2$. For the method of finding and proving the correctness of these polynomials, see [**3**, §9 B].

10.2. The theory of complex multiplication

Recall that we used abstract class field theory to show the Kronecker-Weber theorem, that every abelian extension of $\mathbb{Q}$ is a subfield of some cyclotomic extension. This is the first case of an *explicit* class field theory, where the abelian extensions, described abstractly by class field theory, are somehow made explicit.

We could rephrase the Kronecker-Weber theorem as saying rational values of the function $e^{2\pi ix}$ generate abelian extensions of $\mathbb{Q}$, and every abelian extension of $\mathbb{Q}$ is contained in the field generated by some rational value. If we want to have a more constructive form of Theorem 10.1.1, we need an explicit class field theory for imaginary quadratic fields, and that is what we now discuss.

In the discussion that follows, we always take our lattices to be of full rank. Recall that any order may be viewed as a lattice in $\mathbb{R}^n$; in our case, if $\mathscr{O}$ is an order of an imaginary quadratic field, it is naturally a lattice inside $\mathbb{C}$. The same is true of any non-zero ideal of $\mathscr{O}$ prime to the conductor, and it is easy to check that two such ideals $I, J \in \mathscr{O}$ are equivalent in the class group if and only if $\exists \alpha \in \mathbb{C}$ such that $\alpha I = J$ (the main point to check is that $\alpha I = J$ implies that $\alpha \in K$).

Definition 10.2.1. We say that two lattices $L, L' \subseteq \mathbb{C}$ are **homothetic** if $\exists \alpha \in \mathbb{C}$ with $\alpha L = L'$.

We define certain functions on complex lattices as follows:

Definition 10.2.2. Let $L \subseteq \mathbb{C}$ be a lattice. Then we define:

$$g_2(L) = 60 \sum_{\omega \in L \smallsetminus \{0\}} \frac{1}{\omega^4},$$

$$g_3(L) = 140 \sum_{\omega \in L \smallsetminus \{0\}} \frac{1}{\omega^6},$$

$$\Delta(L) = g_2(L)^3 - 27 g_3(L)^2,$$

and

$$j(L) = 1728 \frac{g_2(L)^3}{\Delta(L)}.$$

It is not too difficult to prove:

Theorem 10.2.3. *For two lattices $L, L' \in \mathbb{C}$, we have $j(L) = j(L')$ if and only if L, L' are homothetic.*

It is then easy to establish a relationship between ideal classes of $\mathscr{O}$ and values $j(I)$ for I an ideal of $\mathscr{O}$ prime to the conductor. It is not too hard to prove:

Theorem 10.2.4. *Given $\sigma \in \mathrm{Aut}(\mathbb{C})$, and I an ideal of $\mathscr{O}$ prime to the conductor, then $\sigma(j(I)) = j(I')$ for I' some other ideal of $\mathscr{O}$.*

Corollary 10.2.5. *$j(\mathscr{O})$ generates an extension of $\mathbb{Q}$ of degree at most equal to $h_{\mathscr{O}}$.*

We thus start to suspect a relationship between the j-function and the ring class field of $\mathscr{O}$. Indeed, we have the following far deeper theorem:

Theorem 10.2.6. *If K is an imaginary quadratic field, and $\mathscr{O}$ an order of $\mathscr{O}_K$, then $j(\mathscr{O})$ generates the ring class field of $\mathscr{O}$ over K.*

More precisely, $j(\mathscr{O})$ is an algebraic integer, and we may take the $f_n(x)$ of Theorem 10.1.1 to be the minimal polynomial of $j(\mathscr{O})$, which is given explicitly as $\prod(x - j(I))$ as I ranges over the ideal classes of $\mathscr{O}$. Moreover, for this choice of $f_n(x)$, the hypothesis of Theorem 10.1.1 that p not divide $\operatorname{disc} f_n(x)$ *may be dropped.*

See [**3**, Thm. 11.1, Exer. 11.1, Thm. 9.2, Prop. 13.2, Thm. 13.23].

One rather involved algorithm for finding the minimal polynomial of $j(\mathscr{O})$, called the **class equation**, is described in [**3**, §13 A,B]. However, it is also possible to compute more efficiently by finding representatives I_i for each ideal class of $\mathscr{O}$, computing each $j(I_i)$ to within a certain precision, and taking advantage of the fact that the minimal polynomial of $j(\mathscr{O})$ has integer coefficients.

This discussion has produced an explicit class field theory for ring class fields. We conclude by mentioning that a slight generalization gives the full collection of ray class fields of imaginary quadratic fields.

Definition 10.2.7. Given a lattice $L \subseteq \mathbb{C}$, let

$$\wp(z, L) = \frac{1}{z^2} + \sum_{\omega \in L \smallsetminus \{0\}} \left(\frac{1}{(z-\omega)^2} - \frac{1}{\omega^2} \right)$$

be the Weierstass $\wp$-function, and

$$\tau(z, L) = \begin{cases} \frac{g_2(L)^2}{\Delta(L)} \wp(z, L)^2 : & g_3(L) = 0; \\ \frac{g_3(L)}{\Delta(L)} \wp(z, L)^3 : & g_2(L) = 0; \\ \frac{g_2(L) g_3(L)}{\Delta(L)} \wp(z, L) : & \text{otherwise} \end{cases}$$

be the Weber function.

We then have the following explicit class field theory for imaginary quadratic fields:

Theorem 10.2.8. *Given K an imaginary quadratic field and $N \in \mathbb{N}$, the ray class field of K of conductor N is generated by $j(\mathscr{O}_K), \tau(1/N, \mathscr{O}_K)$.*

In particular, these are all abelian extensions of K, and every abelian extension of K is contained in one of these.

Remark 10.2.9. The entire theory of complex multiplication can be understood in the context of elliptic curves. Indeed, elliptic curves over $\mathbb{C}$ are closely related to lattices in $\mathbb{C}$, and our j-function simply becomes the j-invariant of the elliptic curve in question. The order $\mathscr{O}$ corresponds to the endomorphism ring of the elliptic curve. Ultimately, we find that we are generating ring class fields by adjoining the j-invariants of elliptic curves having endomorphism ring $\mathscr{O}$, and that to generate all ray class fields for K, we further adjoin the x-coordinates of N-torsion points of the curve.

CHAPTER 11

Density theorems

The classical Dirichlet density theorem asserts that for any $n \in \mathbb{N}$, there exist prime numbers p congruent modulo n to any given value m, as long as $(m, n) = 1$, and that in fact the primes p not dividing n are distributed with equal density among the $\phi(n)$ possible values of m. A much more general statement is the Tchebotarev density theorem, which looks at the distribution of prime ideals with different Frobenius elements in a given extension. The Tchebotarev density theorem may be proved independently of class field theory, but it is also possible to prove (a slightly weaker version of) it using the results of class field theory, and we present that argument here.

11.1. The density theorems

We consider the case that L/K is a Galois, but not necessarily abelian, extension of number fields. We will state the strongest form of the Tchebotarev density theorem, and show some consequences of the theorem.

Recall that in this case, for a fixed prime $\mathfrak{p}$ of $\mathscr{O}_K$, because any two primes lying over $\mathfrak{p}$ are related by an automorphism of $\mathrm{Gal}(L/K)$, we have that the Frobenius elements $\mathrm{Fr}(\mathfrak{q}/\mathfrak{p})$ range through a conjugacy class of $\mathrm{Gal}(L/K)$ as $\mathfrak{q}$ ranges over primes lying over $\mathfrak{p}$. For this reason, we will still write $\mathrm{Fr}(\mathfrak{p})$, but it will denote a conjugacy class rather than an element of $\mathrm{Gal}(L/K)$.

We also need a notion of density:

Definition 11.1.1. Given a collection S of prime ideals of $\mathscr{O}_K$, we say that S has **density** C if

$$C = \lim_{n\to\infty} \frac{\#\{\mathfrak{p} \in S : N(\mathfrak{p}) \leqslant n\}}{\#\{\mathfrak{p} \subseteq \mathscr{O}_K : N(\mathfrak{p}) \leqslant n\}}.$$

Note that a set of primes need not have a density, but if it does, it is a uniquely defined real number between 0 and 1.

In this context, the most general density theorem is the following:

Theorem 11.1.2. *(Tchebotarev density) Let L/K be Galois of degree n, and fix a conjugacy class $\mathfrak{c}$ of* $\mathrm{Gal}(L/K)$*, with m elements. Then the set of primes $\mathfrak{p}$ of K such that* $\mathrm{Fr}(\mathfrak{p}) = \mathfrak{c}$ *has density $\frac{m}{n}$.*

This theorem has some powerful consequences. For instance, one very special case is:

Corollary 11.1.3. *(Dirichlet density) Given $k, n \in \mathbb{N}$ with $(k, n) = 1$, there are infinitely many primes congruent to k modulo n. More precisely, the set of primes congruent to k modulo n has density $\frac{1}{\phi(n)}$.*

Proof. Consider $K = \mathbb{Q}$, and $L = \mathbb{Q}(\zeta_n)$. We saw that $\mathrm{Fr}(p)$ is just the image of p in $\mathrm{Gal}(L/K) \cong (\mathbb{Z}/n\mathbb{Z})^*$, so the condition that $p \equiv k \pmod{n}$ is precisely specifying $\mathrm{Fr}(p)$, and we simply apply Tchebotarev density. □

Another important consequence is that Galois extensions of K are uniquely characterized by which primes split completely in K. We introduce some notation to give the strongest form of the theorem.

Notation 11.1.4. Given an extension L/K, we denote by $S_{L/K}$ the set of primes of K splitting completely in L.

Given sets S, S' of primes of K, we write $S \prec S'$ if S is contained in S' except for a set of primes of density 0.

Applying the Tchebotarev density theorem to the class of the identity, we conclude:

Corollary 11.1.5. *Let L/K be Galois of degree n. Then $S_{L/K}$ has density $\frac{1}{n}$.*

Theorem 11.1.6. *Let L/K be Galois, and E any extension of K. Then $E \subseteq L$ if and only if $S_{L/K} \prec S_{E/K}$.*

Before sketching the proof, we observe that the Galois hypothesis is necessary, because of the following:

Lemma 11.1.7. *Given L/K, let $\tilde{L}$ be the Galois closure of L over K. Then $S_{L/K} = S_{\tilde{L}/K}$.*

Proof. Certainly any prime which splits completely in $\tilde{L}$ must split completely in L. For the converse, we note that for a prime $\mathfrak{p}$ to split completely in $\tilde{L}$ is equivalent to having $\tilde{L}_{\mathfrak{q}} = K_{\mathfrak{p}}$ for one (equivalently, for all) $\mathfrak{q}$ lying over $\mathfrak{p}$. If we fix an imbedding $\bar{K} \to \bar{K}_{\mathfrak{p}}$, this is the same as saying that $\tilde{L}$ is contained in $K_{\mathfrak{p}}$, which is the same as saying that every conjugate of L is contained in $K_{\mathfrak{p}}$. But if $\mathfrak{p}$ splits in L, it splits in every conjugate of L, and so every conjugate is contained in $K_{\mathfrak{p}}$, and $\mathfrak{p}$ splits in $\tilde{L}$, as desired. □

Proof of the theorem. It is clear that if $E \subseteq L$, then $S_{L/K} \prec S_{E/K}$, so we assume that $S_{L/K} \prec S_{E/K}$. By the lemma, we may assume E is Galois over K. We next reduce to the case that $L \subseteq E$: indeed, if we consider EL, which is Galois over K, we have

$$S_{EL/K} = S_{E/K} \cap S_{L/K} \succ S_{L/K},$$

so if we can conclude that $EL \subseteq L$, we have $EL = L$ and $E \subseteq L$, as desired.

Finally, in the case $L \subseteq E$, we use the density theorem: if $S_{L/K} \prec S_{E/K}$, then the density of $S_{L/K}$ is at most the density of $S_{E/K}$, and we conclude that $[L : K] \geqslant [E : K]$, and hence $L = E$. □

From this, we can prove another Hasse-type theorem:

Corollary 11.1.8. *Let $f(x) \in K[x]$ be an irreducible polynomial, with roots in $K_{\mathfrak{p}}$ for all primes $\mathfrak{p}$ of K except possibly a set of density 0; then $f(x)$ has a root in K, and is therefore linear.*

PROOF. (Sketch) Let $L = K(\alpha)$, and $\tilde{L}$ the Galois closure of L over K. Let S be the given set of primes $\mathfrak{p}$ of K such that $f(x)$ has a root in $K_{\mathfrak{p}}$, which has density 1 by hypothesis. We wish to show that every prime of S which is unramified in $\tilde{L}$ splits, which by the previous theorem would then imply that $\tilde{L} = L = K$, and give the desired statement.

Fix $\mathfrak{p}$ unramified in $\tilde{L}$; by hypothesis, $\mathfrak{p}$ has a root over $K_{\mathfrak{p}}$, so there is some prime $\mathfrak{q}$ of L lying over $\mathfrak{p}$, and such that $f_{\mathfrak{q}/\mathfrak{p}} = 1$. Choose $\tilde{\mathfrak{q}}$ in $\tilde{L}$ lying over $\mathfrak{q}$. We claim that $\mathrm{Fr}(\tilde{\mathfrak{q}}/\mathfrak{p}) \in \mathrm{Gal}(\tilde{L}/L) \subseteq \mathrm{Gal}(\tilde{L}/K)$: indeed, since $f_{\mathfrak{q}/\mathfrak{p}} = 1$, we know that $1 = \mathrm{Fr}(\mathfrak{q}/\mathfrak{p}) = \mathrm{Fr}(\tilde{\mathfrak{q}}/\mathfrak{p})|_L$, so $\mathrm{Fr}(\tilde{\mathfrak{q}}/\mathfrak{p})$ holds L fixed. Thus, we see that for all $\mathfrak{p} \in S$ and unramified in $\tilde{L}$, which is a set of density 1, we have $\mathrm{Fr}(\mathfrak{p}) \cap \mathrm{Gal}(\tilde{L}/L)$ non-empty.

By Tchebotarev's theorem, every conjugacy class of $\mathrm{Gal}(\tilde{L}/K)$ arises on one of these $\mathrm{Fr}(\mathfrak{p})$, so we find that $\mathrm{Gal}(\tilde{L}/L)$ meets every conjugacy class of $\mathrm{Gal}(\tilde{L}/K)$, or equivalently, $\mathrm{Gal}(\tilde{L}/K)$ is the union of conjugates of $\mathrm{Gal}(\tilde{L}/L)$. An elementary counting argument shows that this is impossible unless $\mathrm{Gal}(\tilde{L}/L) = \mathrm{Gal}(\tilde{L}/K)$, so $L = \tilde{L}$ was Galois to start with, and every prime $\mathfrak{p}$ of S unramified in $\tilde{L}$, having one $\mathfrak{q}$ lying over it with $f_{\mathfrak{q}/\mathfrak{p}} = 1$, must be split in $\tilde{L} = L$, as desired. □

We observe that Exercise 8.1 shows that the irreducibility hypothesis on $f(x)$ in the corollary is necessary.

11.2. Tchebotarev: sketch of a proof

A slightly weaker formulation of the Tchebotarev density theorem follows easily from class field theory, and we sketch this here. The weakening is due to a slightly different notion of density:

DEFINITION 11.2.1. Let S be a set of primes of K. Then we say S has **Dirichlet density** C if

$$C = \lim_{s \to 1^+} \frac{\sum_{\mathfrak{p} \in S} N(\mathfrak{p})^{-s}}{\log \frac{1}{s-1}}.$$

As with the previous notion of density, it need not exist, but (one checks) that if it does, it is a real number between 0 and 1. One also checks that if the density exists, then the Dirichlet density exists and is the same. However, a set may have a Dirichlet density without having a density, which is why statements in terms of Dirichlet density are slightly weaker.

We can prove the Tchebotarev density theorem with Dirichlet density in place of density:

THEOREM 11.2.2. *(Tchebotarev density) Let L/K be Galois of degree n, and fix a conjugacy class $\mathfrak{c}$ of* $\mathrm{Gal}(L/K)$*, with m elements. Then the set of primes $\mathfrak{p}$ of K such that* $\mathrm{Fr}(\mathfrak{p}) = \mathfrak{c}$ *has Dirichlet density $\frac{m}{n}$.*

All the applications we have described go through unchanged with this weaker form, as long as we replace density with Dirichlet density throughout.

The main idea is to use class field theory and a slight generalization of the analytic class number formula in order to prove the density theorem for abelian extensions, and then (as is typically done) to deduce the non-abelian case from the abelian case.

One basic tool which is used is a simultaneous generalization of the Dirichlet L-series and Dedekind ζ function:

DEFINITION 11.2.3. Given a modulus $\mathfrak{m}$ in a number field K, and a character χ from the ray class group $I_K(\mathfrak{m})/P_{K,1}(\mathfrak{m})$ to $\mathbb{C}^*$, we define:

$$L_{\mathfrak{m}}(s,\chi) := \sum_{I \subseteq \mathscr{O}_K : I \in I_K(\mathfrak{m})} \frac{\chi(I)}{N(I)^s}.$$

Note that when $\mathfrak{m}$ and χ are trivial, we recover $\zeta_K(s)$. The relationship to Dirichlet series is obtained by recalling that the ray class groups of $\mathbb{Q}$ are precisely $(\mathbb{Z}/n\mathbb{Z})^*$. One checks as in the classical case that we have the Euler product formula:

$$L_{\mathfrak{m}}(s,\chi) = \prod_{\mathfrak{p} \nmid \mathfrak{m}} (1 - \frac{\chi(\mathfrak{p})}{N(\mathfrak{p})^s})^{-1}.$$

In order to analyze Dirichlet density, we introduce:

NOTATION 11.2.4. Given functions $f(s), g(s)$, which are defined for $s > 1$ but go to infinity for $s = 1$, we write

$$f(s) \sim g(s)$$

if $f(s) - g(s)$ has a finite limit as s goes to 1 from above.

This is related to Dirichlet density by the observation that a set S of primes has Dirichlet density C if

$$C \log \frac{1}{s-1} \sim \sum_{\mathfrak{p} \in S} N(\mathfrak{p})^{-s}.$$

One checks that $\log \zeta_K(s) \sim \log \frac{1}{s-1}$ using the analytic class number formula. Without class field theory, one can also show using the identity

$$\log L_{\mathfrak{m}}(s,\chi) \sim \sum_{C \in I_K(\mathfrak{m})/P_{K,1}(\mathfrak{m})} \chi(C) \sum_{\mathfrak{p} \in C} \frac{1}{N(\mathfrak{p})^2}$$

that $L_{\mathfrak{m}}(1,\chi)$ is non-zero for non-trivial χ and any ray class group arising from an abelian extension (this is [**7**, pp. 164-165]). The main point is then that the existence theorem of class field theory gives us the ray class field for any ray class group, so that one can conclude that $L_{\mathfrak{m}}(1,\chi) \neq 0$ for any $\mathfrak{m}$, and non-trivial χ, and from there, starting from the above identity, multiplying through by $\chi(C_0^{-1})$ for any fixed ray class C_0, and adding over all χ, we find that

$$\frac{1}{h_{\mathfrak{m}}} \log \frac{1}{s-1} \sim \sum_{\mathfrak{p} \in C_0} \frac{1}{N(\mathfrak{p})^s},$$

where $h_{\mathfrak{m}}$ is the order of the ray class group; see [**7**, p. 166].

Using Dirichlet density throughout, we then conclude density statements for ray classes, and hence for classes in arbitrary generalized ideals classes. Using Artin reciprocity, this gives the statement of Tchebotarev density for abelian extensions, and a standard (and rather short) algebraic argument deduces the general case from the abelian case; see [**7**, pp. 169-170].

CHAPTER 12

Fermat's Last Theorem revisited

We now return to Kummer's work on Fermat's Last Theorem. Although class field is not required, it greatly expedites a number of key steps of the argument.

12.1. The class number factors and regularity

Let us fix an odd prime p. Having shown in Theorem 5.5.3 that Fermat's Last Theorem holds for p when it is strongly h-regular, we wish to give a more computable criterion for when this holds. The final result will be:

Theorem 12.1.1. *An odd prime p is strongly h-regular if and only if p does not divide any of the numerators of the Bernoulli numbers $B_2, B_4, \ldots, B_{p-3}$.*

We recall:

Definition 12.1.2. p is h-regular if p does not divide the class number h of $\mathbb{Q}(\zeta_p)$. p is strongly h-regular if it is h-regular, and if in addition we have that for every unit $u \in \mathbb{Z}[\zeta_p]^*$, if $u \equiv n \pmod{p}$ for some $n \in \mathbb{Z}$, then u is a pth power.

We also recall that in Corollary 9.10.3, we saw that the class number of the maximal real subfield $\mathbb{Q}(\zeta_p + \bar{\zeta}_p)$ divides the class number of $\mathbb{Q}(\zeta_p)$. We thus write

$$h = h^- \cdot h^+$$

where h is the class number of $\mathbb{Q}(\zeta_p)$, h^+ is the class number of $\mathbb{Q}(\zeta_p + \bar{\zeta}_p)$, and $h^- \in \mathbb{N}$ (these are often called the first and second factors).

The proof of the theorem may be broken into two broad steps: first, with the aid of class field theory, we show that p is strongly h-regular if and only if p does not divide h^-; second, we use our work on the analytic class number formula and evaluation of L-series to describe when p divides h^-.

The first step is itself in two parts: we will prove Kummer's lemma, that if p is regular, then in fact p is necessarily strongly h-regular; that is, the condition on the units follows from the condition on the class number. We then show using Kummer theory that if p divides h, it necessarily divides h^-.

We remark that in fact it is a conjecture of Vandiver that p never divides the factor h^+, and this has been checked for primes into the millions. Washington [**10**, pp. 158-159] produces heuristics however showing that without any assumptions on p tending not to divide h^+, one would only expect the number of p less than 4 million such that p divides h^+ to be roughly 1.36. Thus, he views the (superficially very substantial) numerical evidence for the conjecture to be unconvincing.

12.2. Kummer's lemma

We first address the statement that if p is h-regular, it must in fact be strongly h-regular:

THEOREM 12.2.1. *(Kummer's lemma) If p is h-regular, then it is strongly h-regular.*

PROOF. The basic idea is to suppose that we have a unit $u \in \mathbb{Z}[p]^*$ such that $u \equiv n \mod p$, and to explicitly construct an abelian everywhere unramified extension of $\mathbb{Q}(\zeta_p)$ of degree p; this would have to be contained in the Hilbert class field of $\mathbb{Q}(\zeta_p)$, and would therefore imply that the class number h of $\mathbb{Q}(\zeta_p)$ is a multiple of p, as desired.

For notational convenience, we write $\pi = \zeta_p - 1$; recall that $(p) = (\pi)^{p-1}$ as ideals in $\mathbb{Z}[\zeta_p]$.

We first note that without loss of generality, we can assume that u is congruent to 1 modulo p: indeed, if we raise u to the $(p-1)$st power, this will hold, and u^{p-1} is a pth power if and only if u is (suppose $(u')^p = u^{p-1}$; then $(u/u')^p = u$). We thus replace u by u^{p-1}, and have reduced to the situation that u is congruent to 1 modulo p.

We suppose that u is not a pth power. Then the extension $\mathbb{Q}(\zeta_p, u^{1/p})$ of $\mathbb{Q}(\zeta_p)$ is abelian of degree p, and we wish to show that it is everywhere unramified, so that we can conclude that p is not regular by class field theory. The infinite places cannot be ramified, since $\mathbb{Q}(\zeta_p)$ has no real imbeddings. Furthermore, the polynomial $X^p - u$ is inseparable only for primes lying above p, and π is the only prime lying above p, so we conclude that no prime other than π can be ramified in $\mathbb{Q}(\zeta_p, u^{1/p})$. It remains to show that π is unramified.

To do this, we claim that our conditions on u imply that it is actually congruent to 1 modulo π^p, and not merely modulo π^{p-1} as was supposed. Because π^{p-1}/p is a unit, and because every element of $\mathbb{Z}[\zeta_p]$ is congruent to an integer modulo π, we may write

$$u = 1 + pn + p\pi x, \text{ for some } n \in \mathbb{Z}, x \in \mathbb{Z}[\zeta_p].$$

We then compute

$$\pm 1 = N_{\mathbb{Q}(\zeta_p)/\mathbb{Q}}(u) \equiv (1+pn)^{p-1} \equiv 1 + (p-1)pn \equiv 1 - pn \pmod{p\pi},$$

so we see that the only possibility is that $N_{\mathbb{Q}(\zeta_p)/\mathbb{Q}}(u) = 1$ and $\pi | n$, so we had $u \equiv 1 \pmod{p\pi}$ to start with, and u is congruent 1 modulo $(\pi)^p$, as desired.

We see that this implies that the polynomial

$$f(x) = \frac{(\pi x - 1)^p + u}{\pi^p}$$

is actually monic, with coefficients in $\mathbb{Z}[\zeta_p]$. We observe that the roots of $f(x)$ are precisely $\frac{1-\zeta_p^i u^{1/p}}{\pi}$, so they generate the same field extension as $u^{1/p}$. Any two of these roots differ by $\frac{(\zeta^i - \zeta^j)u^{1/p}}{\pi}$ which is a unit of $\mathbb{Z}[\zeta_p]$ times $u^{1/p}$, and we conclude that the discriminant of $f(x)$ is a unit in $\mathbb{Z}[\zeta_p]$, so our extension is everywhere unramified, and we conclude that p divides the class number of $\mathbb{Q}(\zeta_p)$, as desired. □

REMARK 12.2.2. In fact, a more careful analysis of the preceding argument allows one to show that u must be real, from which it then follows that $u \equiv 1 \pmod{\pi^{p+1}}$, and by analyzing the same polynomial $f(x)$ and using Hensel's lemma one sees that π is not only unramified, but splits completely in $\mathbb{Q}(\zeta_p, u^{1/p})$. However, we do not need this sharper result for our application.

12.3. Kummer theory

Kummer theory will play an important role in our analysis of the class number factors. We therefore discuss Kummer extensions and the Kummer pairing. Throughout this section, we fix a field K and an integer n, and suppose that K contains the group of nth roots of unity, which we denote by μ_n.

We say an abelian group has **exponent** n if every element has order dividing n. We say that an extension L/K is **cyclic** (respectively, abelian of exponent n) if it is Galois, with Galois group cyclic (respectively, abelian of exponent n).

Let H be a finite subgroup of $K^*/(K^*)^n$; then H is abelian of exponent n, and it is not hard to see that $K(H^{1/n})$ is a Galois extension of K with $Gal(L/K) \cong H$. Such an extension is called a **Kummer n-extension**. We have:

THEOREM 12.3.1. *An extension L/K is a Kummer n-extension if and only if it is abelian of exponent n.*

In other words, there is a one-to-one correspondence between finite subgroups H of $K^*/(K^*)^n$ and extensions L/K which are abelian of exponent n, and $H \cong \operatorname{Gal}(L/K)$ under this correspondence. Explicitly, $H \mapsto L = K(H^{1/n})$, and $L \mapsto H = K^* \cap (L^*)^n$.

PROOF. Given L/K abelian of exponent n and Galois group G, denote by $H' \subseteq L^*$ the set of elements $x \in L^*$ with $x^n \in K$. If we denote by $\hat{G}$ the group of characters $G \to \mu_n$, we define a map $\psi : H'/K^* \to \hat{G}$ by $\psi(x)(g) = g(x)/x$; we claim it suffices to show that this is an isomorphism. Indeed, denote by H the image of H'/K^* in $K^*/(K^*)^n$ under the nth power map, and let $L' \subseteq L$ be $K(H^{1/n}) = K(H')$. As in the above discussion, L' is Galois with group H, so if we can show that $H \cong G$, it follows that $L' = L$, and L is a Kummer extension. However, it is easy to see from the definitions that the map $H'/K^* \to H$ is an isomorphism, and $G \cong \hat{G}$ (non-canonically), so if ψ is an isomorphism, we get $H \cong G$ and $L = L'$, as desired.

We first note that ψ is injective: an element of the kernel is some $x \in L^*$ such that for any $g \in G = \operatorname{Gal}(L/K)$, we have $g(x) = x$; by Galois theory, we conclude $x \in K^*$. Surjectivity is a deeper matter. Let us suppose that the image of ψ is a proper subgroup of $\hat{G}$. One then shows that there exists some $\sigma \in G$ which is not the identity such that $\psi(x)(\sigma) = 1$ for all $x \in H'$.

If we decompose G appropriately, we may write $G = \langle\tau\rangle \times G_0$, where $\sigma \notin G_0$, so that $\sigma = \tau^i g$ with $\tau^i \neq 1$, and $g \in G_0$. Let $E \subseteq L$ be the fixed field of G_0, so we have $\operatorname{Gal}(E/K) = \langle\tau\rangle$, and denote by $d = \operatorname{ord}\tau$ the degree of E over K. We have $d|n$ because G was originally supposed to have exponent n. Then $\zeta_d \in K$, and we claim it is enough to see that there exists $x \in E$ such that $\zeta_d = \tau(x)/x$. Indeed, in this case we have x^d fixed by τ, so that $x^d \in K$, and $x \in H'$. We thus have by hypothesis that $1 = \psi(x)(\sigma)$. On the other hand, we have

$$\psi(x)(\sigma) = \psi(x)(\tau^i g) = \psi(x)(\tau)^i \psi(x)(g) = \psi(x)(\tau)^i = \zeta_d^i \neq 1,$$

because $\psi(x)(g) = 1$ for any $g \in G_0$, since x is in the fixed field of G_0. This is a contradiction, so it is enough to show that the desired $x \in L$ exists.

But finally, we observe that $N_{E/K}(\zeta_d) = \zeta_d^d = 1$, so the following theorem will complete the proof. □

Theorem 12.3.2. *(Hilbert's Theorem 90) Let L be a cyclic extension of K, with σ a generator of* $\mathrm{Gal}(L/K)$. *Then for any $x \in L$ with $N_{L/K}(x) = 1$, there exists some $y \in L$ with $x = y/\sigma y$.*

For a proof, see [**6**, Cor. A.4, p. 263]. However, we remark that Hilbert's theorem 90, and by extension Kummer theory, is understood most elegantly in the context of Galois cohomology.

An important related tool is the Kummer pairing: given L/K a Kummer extension associated to $H \subseteq K^*/(K^*)^n$, and with Galois group $G \cong H$, the isomorphism is described as follows:

Definition 12.3.3. The **Kummer pairing** is a non-degenerate pairing

$$H \times G \to \mu_n,$$

defined by

$$\langle x, \sigma \rangle = \frac{\sigma(x^{1/n})}{x^{1/n}}.$$

The non-degeneracy statement is the main content of the above proof.

One checks easily that the pairing doesn't depend on the choice of $x^{1/n}$, since we have assumed that $\mu_n \subseteq K$. We thus obtain a natural isomorphism between H and the group of characters of G; this is non-canonically isomorphic to G itself, so we obtain the isomorphism $H \cong G$.

Finally, we need to analyze certain natural group actions which arise in this setting. For this, we assume that we are given some $K_0 \subseteq K$ such that K/K_0 is Galois with group G_0. Let us suppose further that $\mathrm{Gal}(K/K_0)$ sends any representative $x \in K^*$ with $[x] \in H \subseteq K^*/(K^*)^n$ to another representative of an element of H. We then have the following G_0 actions:

(i) G_0 acts on $H \subseteq K^*/(K^*)^n$ because it sends nth powers to nth powers, and sends representatives of H to representatives of H by hypothesis;
(ii) G_0 acts on G by conjugation: given $g_0 \in G_0$ and $g \in G$, extend g_0 to some $\tilde{g}_0 \in \mathrm{Gal}(L/K_0)$, and send $g \to \tilde{g}_0 g \tilde{g}_0^{-1}$; because $\mathrm{Gal}(L/K)$, one checks that this is independent of the choice of $\tilde{g}_0$.
(iii) G_0 acts on μ_n by definition.

The point is that one then easily checks:

Easy Fact 12.3.4. The actions of G_0 on G, H, μ_n commute with the Kummer pairing: i.e., for any $g_0 \in G_0$, and $g \in G$, $h \in H$, we have $\langle g_0(h), g_0(g) \rangle = g_0(\langle h, g \rangle)$.

12.4. Comparison of units and ideal classes

In addition to Kummer theory, we will also need the following background result comparing units in $\mathbb{Q}(\zeta_n)$ and $\mathbb{Q}(\zeta_n + \bar{\zeta}_n)$:

Proposition 12.4.1. *Let $U := \mathbb{Z}[\zeta_n]^*$, $U^+ := \mathbb{Z}[\zeta_n + \bar{\zeta}_n]^*$, and define*

$$Q := [U : \mu_n U^+].$$

Then $Q = 1$ or 2.

Proof. Define a map $U \to \mu_n$ by $u \mapsto u/\bar{u}$; this gives a well-defined map by exercise 2 of homework 5. Compose with the quotient $\mu_n \to \mu_n/\mu_n^2$, and denote the resulting map by $\phi : U \to \mu_n/\mu_n^2$. It is enough to show that $\ker \phi = \mu_n U^+$, since μ_n/μ_n^2 has order 2.

Certainly $\mu_n U^+ \subseteq \ker \phi$, since $\phi(\zeta u^+) = [\zeta^2] = 0$ in μ_n/μ_n^2 for $u^+ \in U^+$ and $\zeta \in \mu_n$. Conversely, suppose $\phi(u) = 0$ for some $u \in U$: then $u/\bar{u} = \zeta^2$ for some $\zeta \in \mu_n$, and $\zeta^{-1}u$ satisfies $\zeta^{-1}u/\bar{\zeta}^{-1}\bar{u} = 1$, so must be real, and $u \in \mu_n U^+$, as desired. □

REMARK 12.4.2. In fact, one can show (see [**10**, Cor. 4.13]) that $Q = 1$ if and only if n is a prime power; however, we will not need this finer result.

PROPOSITION 12.4.3. *An ideal I of $\mathbb{Z}[\zeta_{p^n} + \bar{\zeta}_{p^n}]$ is principal if and only if $I\mathbb{Z}[\zeta_p]$ is principal.*

PROOF. The only if direction is obvious, so let's assume that $I\mathbb{Z}[\zeta_p] = (x)$ for some $x \in \mathbb{Z}[\zeta_p]$, and show that x is of the form $x_1 u$ for some $u \in \mathbb{Z}[\zeta_p]^*$ and $x_1 \in \mathbb{Z}[\zeta_p + \bar{\zeta}_p]$. Now, $(\bar{x}/x) = (\bar{I}/I) = \mathbb{Z}[\zeta_p]$, so $\bar{x}/x$ is a unit of absolute value 1, and by the same argument as exercise 2 of problem set 5, we have that it is a root of unity. Now, write $\pi = \zeta_{p^n} - 1$. Then $\pi/\bar{\pi} = -\zeta_{p^m}$ generates the roots of unity of K, so $\bar{x}/x = (\pi/\bar{\pi})^d$ for some d. Note that this implies that $x\pi^d$ is real.

Let us denote by π^+ the ideal $(\pi) \cap \mathbb{Z}[zeta_p + \bar{\zeta}_p]$; since p is totally ramified in $\mathbb{Q}(\zeta_p)$, with (π) the only ideal lying over it, we see that we must have $\pi^+\mathbb{Z}[\zeta_p] = (\pi)^2$, and any ideal I of $\mathbb{Z}[\zeta_p]$ coming from $\mathbb{Z}[\zeta_p + \bar{\zeta}_p]$ must have an even power of π in it, which is to say that $\nu_\pi(I)$ must be even.

We now observe that $d = \nu_\pi(x\pi^d) - \nu_\pi(x) = \nu_\pi(x\pi^d) - \nu_\pi(I)$ must be even, since $x\pi^d$ and I are real. So we have that $\bar{x}/x = \zeta^2 = \zeta/\bar{\zeta}$ for some root of unity ζ, and we have that $x\zeta$ is real, so $I = (x\zeta)$ in $\mathbb{Z}[\zeta_p + \bar{\zeta}_p]$. □

12.5. The eigenspaces and p-rank

We are now ready to use use Kummer theory to prove:

THEOREM 12.5.1. *If p divides h^+, it must also divide h^-. Hence, p is h-regular if and only if p does not divide h^-.*

The basic idea is to use class field theory to relate Kummer theory to the ideal class group, and to study eigenspaces for the action of the complex conjugation map τ on the p-part of the ideal class group of $\mathbb{Q}(\zeta_p)$. We note that τ acts on fractional ideals of K, and sends principal ideals to principal ideals, so acts on the ideal class group.

Let C_p denote the p-Sylow subgroup of the ideal class group of K; then τ acts on C_p, and we may write $C_p^{\pm} = \{I \in C_p : \tau(I) = I^{\pm 1}\}$. One easily checks that $C_p^{\pm} = \frac{1 \pm \tau}{2} C_p$, and hence (noting that $C_p^+ \cap C_p^-$ consists of classes with trivial square, which must then be trivial since p is odd) that $C_p = C_p^+ \oplus C_p^-$. We claim:

PROPOSITION 12.5.2. *C_p^+ is isomorphic to the p-Sylow subgroup of the ideal class group of K^+ under the natural map $I \mapsto I\mathscr{O}_K$.*

PROOF. Certainly, the map sends an ideal I of p-power order to an element of C_p^+. This map is injective by the previous proposition, so we need only check surjectivity. Since K is quadratic over K^+, every prime ideal $\mathfrak{p}$ of K^+ either remains prime in K, or factors in the form $\mathfrak{q}\tau(\mathfrak{q})$. We see that an ideal I of K, as long as it is divisible only by primes unramified over K^+, comes from K^+ if and only if it has the same number of factors of $\mathfrak{q}$ and $\tau(\mathfrak{q})$ for every prime $\mathfrak{q}$ of K. But if $I \in C_p^+$, we can multiply it by a principal ideal so that it is prime to any primes ramified over K^+, and we can write $I = \frac{1+\tau}{2} I'$ for some ideal I', at least if we allow ourselves to multiply I by an appropriate principal ideal. It is then clear that $\frac{1+\tau}{2}I'$ satisfies the condition on its prime factor to come from an ideal of K^+. □

We now bring class field theory into the picture in order to study C_p in more detail. We have that C_p/C_p^p is the largest quotient (in the sense of modding out by the smallest subgroup) of C_p of exponent p; this is also the largest quotient of the ideal class group of K itself having exponent p. By class field theory, there is a field L with $\mathrm{Gal}(L/K)$ naturally isomorphic to C_p/C_p^p under the Artin map, and by Kummer theory, we have a subgroup $H \subseteq K^*/(K^*)^p$ such that $L = K(H^{1/n})$, and $H \cong C_p/C_p^p$. Class field theory tells us that L is the maximal abelian everywhere unramified extension of K of exponent p, so we see that H is described as the set of elements $x \in K^*$ such that $K(x^{1/p})$ is everywhere unramified over K. Thus, H is preserved by automorphisms of K, and in particular it has an action of τ, so we denote the eigenspaces by H^+ and H^- respectively, and similarly for $\mathrm{Gal}(L/K)$. We have the Kummer pairing as

$$\langle , \rangle : H \times \mathrm{Gal}(L/K) \to \mu_p,$$

and it is easily checked that $\langle \tau(x), \tau(g) \rangle = \tau(\langle x, g \rangle)$.

We also note that under the Artin map, the ± eigenspaces of τ acting on C_p are sent to the corresponding eigenspaces of $\mathrm{Gal}(L/K)$.

Finally, we have a map

$$\phi : H \to C_{p,1} := \{I \in C_p : I^p = 1\}$$

defined by sending $x \in H$ to the pth root of (x) as an ideal: this exists because $K(x^{1/p})$ is everywhere unramified, so the fact that (x) is a pth power of an ideal in $K(x^{1/p})$ implies that it is the pth power of an ideal in K, simply by considering prime factorizations. It is clear that ϕ commutes with the τ action.

We have:

Proposition 12.5.3. *$\ker \phi$ is naturally contained in $\mathscr{O}_K^*/(\mathscr{O}_K^*)^p$ under the map $\mathscr{O}_K^*/(\mathscr{O}_K^*)^p \to K^*/(K^*)^p$.*

Proof. We first observe that we have an isomorphism $\mathscr{O}_K^*/(\mathscr{O}_K^*)^p \xrightarrow{\sim} \mathscr{O}_K^*(K^*)^p/(K^*)^p$. Next, suppose that $\phi(x) = 1$. Then $(x) = (y)^p$ for some $y \in K^*$, and we conclude that $x = uy^p$ for some $u \in \mathscr{O}_K^*$, so that $x \in \mathscr{O}_K^*(K^*)^p/(K^*)^p$, as desired. □

Recall that the p-rank of a finite abelian group is defined to be the number of factors in a cyclic decomposition having order a power of p, so the p-rank of C_p is the minimal number of generators.

The main theorem is:

Theorem 12.5.4. *The p-rank of C_p^+ is less than or equal to the p-rank of C_p^-.*

From this theorem, we immediately conclude that if $p|h^+$, we must also have $p|h^-$, so Theorem 12.5.1 follows immediately.

Proof. We first note that $\langle H^+, \mathrm{Gal}(L/K)^+ \rangle = \langle H^-, \mathrm{Gal}(L/K)^- \rangle = 1$, since $\langle x, g \rangle = \langle \pm x, \pm g \rangle = \langle \tau(x), \tau(g) \rangle = \tau \langle x, g \rangle$, and $\langle , \rangle$ takes values in the pth roots of unity, which (because p is odd) are all complex except for 1. Thus, since the Kummer pairing is non-degenerate, we conclude that

$$\langle , \rangle : H^- \times \mathrm{Gal}(L/K)^+ \to \mu_p$$

must be non-degenerate, so the $p-\operatorname{rank} H^- = p-\operatorname{rank} \operatorname{Gal}(L/K)^+$, and by the class field theory isomorphism, we have

$$p-\operatorname{rank} C_p^+ = p-\operatorname{rank}\operatorname{Gal}(L/K)^+ = p-\operatorname{rank} H^-.$$

It remains to bound the p-rank of H^- in terms of the p-rank of C_p^-.

But ϕ induces a map

$$\phi^- : H^- \to C_{p,1}^-,$$

which we claim is injective. We know that the kernel is a subgroup of $(\mathscr{O}_K^*/(\mathscr{O}_K^*)^p)^-$, and we claim that we have $(\mathscr{O}_K^*/(\mathscr{O}_K^*)^p)^- \xleftarrow{\sim} \mu_p$. Indeed, the natural map is certainly injective, and we remarked (but didn't prove) earlier that $\mathscr{O}_K^* = \mu_p(\mathscr{O}_K^*)^+$, i.e. that $Q = 1$. One checks that this means that any $u \in \mathscr{O}_K^*$ with image in $(\mathscr{O}_K^*/(\mathscr{O}_K^*)^p)^-$ is of the form ζu_1 with $\zeta \in \mu_p$ and $u_1^2 = u_0^p$, so that in $(\mathscr{O}_K^*/(\mathscr{O}_K^*)^p)^-$, u is equivalent to $\pm\zeta$, and hence to ζ. Therefore, the kernel of ϕ^- is contained in $\mu_p \cap H$, which is trivial because if $\zeta \in \mu_p$ is non-trivial, we have $K(\zeta^{1/p}) = \mathbb{Q}(\zeta_{p^2})$ is totally ramified over p, whereas $K(H^{1/p})$ is everywhere unramified over K by hypothesis. Thus, ϕ^- is injective, and the p-rank of H^- is bounded by the p-rank of $C_{p,1}^-$, which is easily checked to be the same as the p-rank of C_p^-. This gives the desired bound. □

We finally conclude the desired statement, that if $p|h^+$, we must also have $p|h^-$, so we have proved that p is regular if and only if p does not divide h^-. We are thus reduced to the study of when p divides h^-, which we will carry out with the analytic class number formula.

12.6. A regulator calculation

Since the analytic class number formula uses regulators, we have to compare the regulators of $K = \mathbb{Q}(\zeta_p)$ and $K^+ = \mathbb{Q}(\zeta_p + \bar{\zeta}_p)$.

Recall the notation $Q = [U : \mu_p U^+]$, where $U = \mathbb{Z}[\zeta_p]^*$ and $U^+ = \mathbb{Z}[\zeta_p + \bar{\zeta}_p]^*$. Recall also that we showed that $Q = 1$ or 2. We will show:

Proposition 12.6.1. *The regulators are related by*

$$R_K/R_{K^+} = \frac{1}{Q}2^{(p-3)/2}.$$

Proof. Let $u_1, \ldots, u_m$ be a basis for U^+ modulo ± 1 (which are the roots of unity in K^+); by definition, R_{K^+} is the volume of the lattice spanned by $\psi_{K^+}(u_i) \subseteq \mathbb{R}^{m+1}$, where $m + 1 = (p-1)/2$ and the lattice has rank $m = (p-3)/2$. But now consider the lattice $\psi_K(u_i)$; this has index Q in $\psi_K(U)$, so its volume is QR_K. But the only difference between $\psi_K(u_i)$ and $\psi_{K^+}(u_i)$ is that every coordinate has been multiplied by 2, since every imbedding was real for K^+ but is complex for K. Thus

$$QR_K = \operatorname{vol}\psi_K(u_i) = 2^m \operatorname{vol}\psi_{K^+}(u_i) = 2^m R_{K^+},$$

and we get the desired formula. □

12.7. h^- and the analytic class number formula

We now give a formula for h^- using the analytic class number formula and our evaluations of Dirichlet L-series. Considering $K = \mathbb{Q}(\zeta_p)$, and varying characters over the entire group of Dirichlet characters modulo p, we see:

$$\frac{2^{r_1+r_2}\pi^{r_2}h_K R_K}{m_K|D_K|^{1/2}} = \frac{2^{(p-1)/2}\pi^{(p-1)/2}h_K R_K}{2p|D_K|^{1/2}} \lim_{s\to 1^+}(s-1)\zeta_K(s) = \prod_{\chi\neq\chi_1} L(1,\chi),$$

and

$$\frac{2^{r_1+r_2}\pi^{r_2}h_{K^+} R_{K^+}}{m_{K^+}|D_{K^+}|^{1/2}} = \frac{2^{(p-1)/2}h_{K^+} R_{K^+}}{2|D_{K^+}|^{1/2}} \lim_{s\to 1^+}(s-1)\zeta_{K^+}(s) = \prod_{\chi\neq\chi_1,\chi \text{ even}} L(1,\chi).$$

Dividing through and using our regulator calculation, we find

$$\frac{h^- 2^{(p-3)/2}\pi^{(p-1)/2}}{pQ|D_K/D_{K^+}|^{1/2}} = \prod_{\chi\neq\chi_1,\chi \text{ odd}} L(1,\chi).$$

We now apply our evaluation of

$$L(1,\chi) = \frac{\pi i \tau(\chi)}{f_\chi^2}\sum_{k=1}^{f-1}\bar{\chi}(k)k$$

and we find

$$\frac{h^- 2^{(p-3)/2}\pi^{(p-1)/2}}{pQ|D_K/D_{K^+}|^{1/2}} = \pi^{(p-1)/2} i^{(p-1)/2} \prod_{\chi \text{ odd}} \frac{\tau(\chi)}{f_\chi^2}\sum_{k=1}^{f-1}\bar{\chi}(k)k.$$

Recall our formulas for the product of conductors and for the product of the Gauss sums:

$$\prod_{\chi\in G} f_\chi = |D_{K_G}|$$

$$\prod_{\chi\in G} \tau(\chi) = i^{r_2}\sqrt{|D_{K_G}|}$$

so applying the formulas to K and K^+ and taking ratios, we find that

$$\prod_{\chi \text{ odd}} \frac{\tau(\chi)}{f_\chi} = \frac{i^{(p-1)/2}|D_K|^{1/2}}{|D_{K^+}|^{1/2}}\frac{|D_{K^+}|}{|D_K|} = \frac{i^{(p-1)/2}|D_{K^+}|^{1/2}}{|D_K|^{1/2}}.$$

Substituting and cancelling terms, we find:

$$h^- = \frac{(-1)^{(p-1)/2}pQ}{2^{(p-3)/2}} \prod_{\chi \text{ odd}} \frac{1}{f_\chi}\sum_{k=1}^{f-1}\chi(k)k.$$

If denote denote $\frac{1}{f_\chi}\sum_{k=1}^{f-1}\chi(k)k$ by $B_{1,\chi}$ (this is called a **generalized Bernoulli number**, we can ignore factors of any numbers prime to p for our purposes, and we have proved:

THEOREM 12.7.1. *h^- is divisible by p if and only if*

$$p\prod_{\chi \text{ odd}} B_{1,\chi}$$

is divisible by p.

The $B_{1,\chi}$ are closely related to the classical Bernoulli numbers B_n. There are various ways to pursue this relationship: one can view the characters as taking values in a p-adic, rather than complex, field, and prove a certain congruence between each $B_{1,\chi}$ and some B_n, as in [**10**, Cor. 5.15], or equivalently, one can analyze the product of the $B_{1,\chi}$ in $\mathbb{Q}(\zeta_{p-1})$ and relate it to a product of sums of powers of integers, which are closely related to Bernoulli numbers, as in [**1**, §5.4]. In either case, one exploits this relationship to prove Kummer's theorem:

THEOREM 12.7.2. *$p|h^-$ if and only if p divides one of the numerators of the Bernoulli numbers $B_2, \dots, B_{p-3}$. In particular, if p does not divide a numerator of any $B_2, \dots, B_{p-3}$, then*

$$x^p + y^p = z^p$$

has no non-zero integer solutions.

Bibliography

1. Z. I. Borevich and I. R. Shafarevich, *Number theory*, Academic Press, 1966.
2. Keith Conrad, *Fermat's last theorem for regular primes*.
3. David A. Cox, *Primes of the form $x^2 + ny^2$*, Wiley-Interscience, 1989.
4. Harold M. Edwards, *Fermat's last theorem*, Springer-Verlag, 1977.
5. Kenneth Ireland and Michael Rosen, *A classical introduction to modern number theory*, second ed., Springer-Verlag, 1990.
6. Gerald Janusz, *Algebraic number fields*, Academic Press, 1990.
7. Serge Lang, *Algebraic number theory*, second ed., Springer-Verlag, 1994.
8. Pierre Samuel, *Algebraic theory of numbers*, Houghton Mifflin, 1970, translated from the French by Allan J. Silberger.
9. J. P. Serre, *A course in arithmetic*, Graduate Texts in Mathematics, no. 7, Springer-Verlag, 1973.
10. Lawrence Washington, *An introduction to cyclotomic fields*, second ed., Graduate Texts in Mathematics, no. 83, Springer-Verlag, 1997.