Math 129: Number Fields

Taught by Mark Kisin

Spring 2019

These notes are scribed by Kenz Kallal (me). My email is kenz.kallal@gmail.com. Please contact me if you find any mistakes in these notes (all mistakes are mine and not the instructor's). You can also read the notes in real-time using the link [REDACTED].

Other administrative details are below:

- **Section:** SC 411, Thursday 4:30–5:45
- **CA Office hours:** Winthrop dining hall (behind the private dining room), Tuesdays 7-9 PM.
- **Office hours:** Tuesdays 1pm or by appointment: Prof. Kisin's email is kisin.mark@gmail.com
- **Textbook:** *Algebraic Theory of Numbers* by Pierre Samuel. Also Neukirch's *Algebraic Number Theory*.
- **Midterm:** Wednesday, March 13.
- **Final:** Saturday May 11th, 2pm.

I am always available outside of office hours by email to answer any questions related to the course material.

Contents

§1 January 28, 2019

§1.1 Topics

The topics we will cover are the following:

1. Unique factorization. Any $n \in \mathbf{Z}$ factors uniquely as a product of primes. For rings of integers of number fields, how is this salvaged?
2. Class groups
3. Unit groups
4. Local fields

One application of these topics is to Diophantine equations.

1. [Fermat's Christmas Theorem] If $p \equiv 1 \pmod 4$ then there are $a, b \in \mathbf{Z}$ such that $a^2 + b^2 = p$.
2. [Pell's Equation] Let d be a squarefree integer. What are the integer solutions to $a^2 - db^2 = 1$?

§1.2 Number Fields and their Rings of Integers

Important prerequisites from linear algebra:

1. Vector spaces and linear maps between them
2. Cayley–Hamilton Theorem

Definition 1.1. A **number field** is a field K of characteristic zero (i.e. $K \supseteq \mathbf{Q}$) such that K is a finite dimensional vector space over $\mathbf{Q}$. In other words, as a vector space over $\mathbf{Q}$, K is noncanonically isomorphic to $\mathbf{Q}^n$ for some $n \in \mathbf{N}$.

Example 1.2

Let d be a squarefree (possibly negative) integer, and $K = \mathbf{Q}(\sqrt{d})$. To be specific, $K = \mathbf{Q}[X]/(X^2 - d)$. The dimension of K as a $\mathbf{Q}$-vector space is

$$\dim_{\mathbf{Q}} K = 2.$$

This is because the elements 1 and $\overline{x}$ form a basis for K. As a $\mathbf{Q}$-vector space, $K = \mathbf{Q} \cdot 1 + \mathbf{Q} \cdot \overline{x}$.

We can define a number field in a similar way by adjoining a root of any irreducible polynomial over $\mathbf{Q}$.

Lemma 1.3

If K is a number field, $\alpha \in K$, then there exists a monic polynomial $f \in \mathbf{Q}[X]$ such that $f(\alpha) = 0$.

Proof 1. Let $d = [K : \mathbf{Q}]$ (this is an easier notation for $\dim_{\mathbf{Q}} K$). The powers $1, \alpha, \alpha^2, \dots, \alpha^d$ are linearly dependent over $\mathbf{Q}$ since K is d-dimensional as a $\mathbf{Q}$-vector space and there are $d+1$ elements of K in the list. This means that there exist $a_0, \dots, a_d \in \mathbf{Q}$ such that

$$a_0 + a_1\alpha + \cdots + a_d\alpha^d = 0.$$

Dividing through by a_i, where i is as large as possible such that $a_i \neq 0$, we get the desired monic polynomial. NB: this proof also guarantees the existence of f with degree at most d. □

Proof 2. Let L be a field and V a vector space over L of dimension $d < \infty$. If $\varphi : V \to V$ is an L-linear map, then the characteristic polynomial of φ, $P_\varphi(X) \in L[X]$, has degree d. The Cayley–Hamilton Theorem says that

$$P_\varphi(\varphi) = 0$$

as an element of $\mathrm{End}(V)$ [recall that polynomials in $L[X]$ can be evaluated at elements of $\mathrm{End}(V)$].

The characteristic polynomial of φ can be defined in the following way: Consider the $L[X]$-module given by $V \otimes_L L[X]$, which is identified noncanonically to the free $L[X]$-module $L[X]^d$. The determinant (i.e. the top alternating power) of the map $X - \varphi \in \mathrm{End}_{L[X]}(V \otimes_L L[X])$ is the characteristic polynomial of φ (here X stands for the map given by multiplication by X). NB: noncanonically, $X - \varphi$ is given by the matrix $X \cdot \mathrm{id} - \varphi$, since X is in the base ring.

Now we abuse Cayley–Hamilton to prove the lemma. In particular, let φ_α be the $\mathbf{Q}$-linear map $K \to K$ given by $x \mapsto x \cdot \alpha$. Cayley–Hamilton says that the characteristic polynomial of φ_α is satisfied by φ_α, i.e.

$$P_{\varphi_\alpha}(\varphi_\alpha) = 0 \in \mathrm{End}(K).$$

One can check that $P_{\varphi_\alpha}(\varphi_\alpha)$ is the $\mathbf{Q}$-linear endomorphism of K given by multiplication by $P_{\varphi_\alpha}(\alpha)$. The conclusion of Cayley–Hamilton therefore tells us that multiplication by $P_{\varphi_\alpha}(\alpha)$ kills everything in K. Since $1 \in K$, it follows that P_{φ_α} is a monic polynomial with coefficients in $\mathbf{Q}$ of degree d that kills α. □

Questions from algebra: If $K = \mathbf{Q}[X]/h(X)$ where h is irreducible, then any $\alpha \in K$ satisfies a monic polynomial. What is the polynomial for α^2, α^3? What about $\alpha + \beta$?

Definition 1.4. If K is a number field, the **ring of integers** $\mathcal{O}_K \subseteq K$ is the set of all $\alpha \in K$ such that there exists a monic polynomial $f \in \mathbf{Z}[X]$ with $f(\alpha) = 0$.

NB (from Alex Wei): it is not obvious from the definition that $\mathcal{O}_K$ is a ring. We will prove that it is later.

Lemma 1.5 (Gauss's Lemma)
If $f \in \mathbf{Z}[X]$ is monic and $f = f_1 f_2$ where $f_1, f_2 \in \mathbf{Q}[X]$ are monic, then $f_1, f_2 \in \mathbf{Z}[X]$.

Proof. Homework problem. □

Lemma 1.6

If $\alpha \in K$, the following are equivalent:

1. $\alpha \in \mathcal{O}_K$.
2. The minimal polynomial of α has integer coefficients

Proof. Recall that the *minimal polynomial* f_0 of α is the monic polynomial with the following two properties:

i. $f_0(\alpha) = 0$

ii. if $g \in \mathbf{Q}[X]$ such that $g(\alpha) = 0$, then $f | g$.

NB: f_0 is just a (uniquely determined by the condition that it is monic) generator for the ideal given by the kernel of the map $\mathbf{Q}[X] \to \mathbf{Q}[\alpha]$ given by evaluation at α.

The proof of the lemma is as follows: (ii) obviously implies (i). For the other direction, we abuse Gauss's lemma. Suppose that there exists a monic $f \in \mathbf{Z}[X]$ such that $f(\alpha) = 0$. Then the minimal polynomial f_0 divides f in $\mathbf{Q}[X]$, i.e. there is a $g \in \mathbf{Q}[X]$ such that $f = f_0 g$. Gauss's lemma guarantees that $f_0 \in \mathbf{Z}[X]$. □

§1.3 Quadratic fields

Example 1.7

Let $K = \mathbf{Q}(\sqrt{2})$. Then $\mathcal{O}_K = \mathbf{Z}[\sqrt{2}]$. The proof is as follows: if $b \neq 0$, the minimal polynomial of $a + b\sqrt{2}$ is

$$(X - (a + b\sqrt{2}))(X - (a - b\sqrt{2})) = X^2 - 2aX + a^2 - 2b^2.$$

So by Lemma 1.6, $a + b\sqrt{2} \in \mathcal{O}_K$ if and only if $2a, a^2 - 2b^2$ are integers.

NB: this isn't completely tautological. For example, the ring of integers of $\mathbf{Q}(\sqrt{5})$ is $\mathbf{Z}[(1 + \sqrt{5})/2]$. Note that the minimal polynomial of $(1 + \sqrt{5})/2$ is

$$(X - (1 + \sqrt{5})/2)(X - (1 - \sqrt{5})/2) = X^2 - X - 1 \in \mathbf{Z}[X].$$

However, in the case of $\mathbf{Q}(\sqrt{2})$ it works out: $2a \in \mathbf{Z}$ means $a \in \frac{1}{2}\mathbf{Z}$ and $a^2 \in \frac{1}{4}\mathbf{Z}$. Since $a^2 - 2b^2 \in \mathbf{Z}$, it follows that $b^2 \in \frac{1}{8}\mathbf{Z}$. This means that the denominator of b can have at most a single factor of 2, from which it follows that $b^2 \in \frac{1}{4}\mathbf{Z}$. So $2b^2 \in \frac{1}{2}\mathbf{Z}$, and from $a^2 - 2b^2 \in \mathbf{Z}$ we know $a^2 \in \frac{1}{2}\mathbf{Z}$. By the same argument as before (a cannot have any factors of 2 in the denominator or else its square would have denominator divisible by 4), $a \in \mathbf{Z}$. We conclude then that $2b^2 \in \mathbf{Z}$, hence $b \in \mathbf{Z}$ as well. The result $\mathcal{O}_K = \mathbf{Z}[\sqrt{2}]$ follows.

The example generalizes to all quadratic fields.

Proposition 1.8

If $K = \mathbf{Q}(\sqrt{d})$ for d squarefree, then

$$\mathcal{O}_K = \begin{cases} \mathbf{Z}[\sqrt{d}] \text{ if } d \not\equiv 1 \pmod 4, \\ \mathbf{Z}\left[\frac{1+\sqrt{d}}{2}\right] \text{ if } d \equiv 1 \pmod 4 \end{cases}.$$

Proof. We have $K = \mathbf{Q} + \mathbf{Q} \cdot \sqrt{d}$. If $\alpha = a + b\sqrt{d} \in K - \mathbf{Q}$ then it has minimal polynomial

$$(X - (a + b\sqrt{d}))(X - (a - b\sqrt{d})) = X^2 - 2aX + a^2 - db^2$$

so by Lemma 1.6, $\alpha \in \mathcal{O}_K$ iff $2a, a^2 - db^2 \in \mathbf{Z}$. If $\alpha \in \mathcal{O}_K$ we can write $a = A/2$ for some $A \in \mathbf{Z}$, and we have

$$A^2 - 4db^2 = A^2 - d(2b)^2 \in 4\mathbf{Z}.$$

If $a \notin \mathbf{Z}$, i.e. A is odd, then $A^2 \equiv 1 \pmod 4$ and thus $d(2b)^2 \equiv 1 \pmod 4$. So, b cannot be an integer. Since d is squarefree, it has at most one factor of 2. It follows by looking at the denominator of $d(2b)^2$ that $4b^2 \in \mathbf{Z}$ which means $b \in \frac{1}{2}\mathbf{Z}$. The fact that $b \notin \mathbf{Z}$ means that $2b$ is an odd integer and thus $(2b)^2 \equiv 1 \pmod 4$. Since $d(2b)^2 \equiv 1 \pmod 4$, it also follows that d is 1 mod 4. Part of the result falls out of this analysis: if $d \not\equiv 1 \pmod 4$, then we showed that A is even, so $a \in \mathbf{Z}$, and $db^2 \in \mathbf{Z}$ so since d is square-free, $b \in \mathbf{Z}$ as well. Hence, $\mathcal{O}_K = \mathbf{Z}[\sqrt{d}]$ (the inclusion $\mathbf{Z}[\sqrt{d}] \subseteq \mathcal{O}_K$ is obvious from computing minimal polynomials).

If $d \equiv 1 \pmod 4$, then we can check that the minimal polynomial of $(1 + \sqrt{d})/2$ has integer coefficients, so $\mathcal{O}_K$ contains $\mathbf{Z}[(1 + \sqrt{d})/2]$. Any $\alpha \in \mathcal{O}_K$ must have the form $\alpha = A/2 + b\sqrt{d}$ for some $A \in \mathbf{Z}$, so we can look at $\alpha - A(1 + \sqrt{d})/2 = (b - A/2)\sqrt{d}$. Since this is in $\mathcal{O}_K$, its minimal polynomial has integer coefficients, which means

$$(b - A/2)^2 d \in \mathbf{Z}.$$

The usual argument tells us $b - A/2 \in \mathbf{Z}$, hence b is of the form $B/2$ where B has the same parity as A. So we can write $a + b\sqrt{d}$ as $(1 + \sqrt{d})/2$ plus an element of $\mathbf{Z}[\sqrt{d}]$, from which it follows that $\mathcal{O}_K = \mathbf{Z}[(1 + \sqrt{d})/2]$. □

We used twice in this proof that $\mathcal{O}_K$ is a ring. We will show tomorrow that this is true in the general setting.

Note: it's often convenient to write that in the case $d \equiv 1 \pmod 4$, $\mathcal{O}_K$ is the set of all $(A + B\sqrt{d})/2$ such that $A, B \in \mathbf{Z}$ with the same parity. This way, it is immediate that $\mathcal{O}_K$ is a free $\mathbf{Z}$-module of rank 2. For example, one possible basis is given by

$$1, \frac{1 + \sqrt{d}}{2}.$$

Instead of using the fact that $\mathcal{O}_K$ is a ring, we could have shown directly by computing minimal polynomials that when $d \equiv 1 \pmod 4$, any $(A + B\sqrt{d})/2$ with A, B of the same parity lies in $\mathcal{O}_K$; moreover, we already showed that any element of $\mathcal{O}_K$ is of the form $(A + B\sqrt{d})/2$ where $A, B \in \mathbf{Z}$ and that if A is odd then B is odd. It suffices to show that if A is even then so is B, but coincidentally we also showed this when finishing off the $d \not\equiv 1 \pmod 4$ case (when we showed that if $a \in \mathbf{Z}$ then $b \in \mathbf{Z}$).

§2 January 30, 2019

"Were you coming from somewhere else or are you just horribly irresponsible?" – Mark Kisin on coming to class after 1:30pm.

First, we filled in the second proof of Lemma 1.3, and the proof of Proposition 1.8 from last class.

§2.1 Integral Closures and Rings of Integers

In the computation from last class (see Proposition 1.8), we showed that if K is a number field of degree 2, then $\mathcal{O}_K$ is a finitely generated free abelian group of rank 2. This is true in general, but first we will pay the debt we accrued last class by showing that $\mathcal{O}_K$ is a ring.

Note that Cayley–Hamilton also works over arbitrary commutative rings with unity. Let R be such a ring. If M is a finitely generated free R-module of rank d and φ is an R-endomorphism of M, then we can construct the characteristic polynomial of φ in the same way, and it vanishes at φ. We will show this in section on Friday.

In this class, a "ring" always refers to a commutative ring with unity.

Definition 2.1. Let $A \subseteq R$ be rings. An element $\alpha \in R$ is called **integral** over A if there exists a monic polynomial $f(X) \in A[X]$ such that $f(\alpha) = 0$. We say that R is integral over A if every $\alpha \in R$ is integral over A. Finally, we define the **integral closure** of A in R to be the set of $\alpha \in R$ which are integral over A.

Example 2.2

If K is a number field, then $\mathcal{O}_K$ is the integral closure of $\mathbf{Z}$ in K.

Proposition 2.3

If $A \subseteq R$ are rings, then the integral closure of A in R is a subring of R.

Proof. Denote by A' the integral closure of A in R. We need two important lemmas to proceed.

Lemma 2.4

Let M be a finitely generated A-module. Let $\alpha : M \to M$ be an A-linear map. Then there exists a monic polynomial over A that vanishes on α.

Proof. NB: this is stronger than the previous statement of Cayley–Hamilton over a ring, since it only assumes that M is finitely-generated instead of free.

If M is free, then the result is true by Cayley–Hamilton: we have a monic polynomial $P_\alpha \in A[X]$ which vanishes on α.

If M is finitely generated by d elements, there is a surjective A-linear map $A^d \to M$. Since a linear map out of A^d is determined by where it sends a basis, the action of α on M induces an A-linear map $\tilde{\alpha} : A^d \to A^d$ (we can send the basis element x_i to an arbitrary lift of $\alpha(\overline{x_i})$; note that the lifts are not uniquely determined). To be precise, we have a commutative diagram

$$\begin{array}{ccc} A^d & \longrightarrow & M \\ \downarrow{\scriptstyle \tilde{\alpha}} & & \downarrow{\scriptstyle \alpha} \\ A^d & \longrightarrow & M \end{array}$$

We know that $P_{\tilde{\alpha}}(\tilde{\alpha}) = 0$ by Cayley–Hamilton for free modules of finite rank. By using the same choices for lifts of images under α that we use to induce $\tilde{\alpha}$ from α, we can check that $P_{\tilde{\alpha}}(\tilde{\alpha})$ is induced by $P_{\tilde{\alpha}}(\alpha) : M \to M$. Since the diagram commutes and the sideways maps are surjective, it follows that $P_{\tilde{\alpha}}(\alpha) = 0$ as well. □

Lemma 2.5

The following are equivalent:

(1) $f(\alpha) = 0$ for some monic $f \in A[X]$.

(2) $A[\alpha] \subseteq R$ is finitely-generated as an A-module.

(3) $A[\alpha]$ is contained in a subring $M \subseteq R$ such that M is a finitely generated A-submodule of R.

Proof. If $f(\alpha) = 0$, then the map $A[X] \to A[\alpha]$ given by evaluation at α factors through $A[X]/(f)$ via the canonical projection. But $A[X]/(f)$ is generated as an A-module by $1, \dots, X^{d-1}$ where $d = \deg f$. Since the map $A[X]/(f) \to A[\alpha]$ is surjective, it follows that $A[\alpha]$ is finitely generated as an A-module. This shows that (1) implies (2). That (2) implies (3) is obvious so it suffices to show (3) implies (1).

Finally, if (3) holds, then M is closed under multiplication by α, so we have an A-linear map $\varphi_\alpha : M \to M$ given by multiplication by α. By Lemma 2.4, there exists a monic polynomial $f(X) \in A[X]$ such that $f(\varphi_\alpha) = 0$. Since $f(\varphi_\alpha)$ is the A-linear endomorphism given by multiplication by $f(\alpha)$, it follows from the fact that $1 \in M$ that $f(\alpha) = 0$, as desired. □

Lemma 2.5 lets us prove the proposition. Since $0, 1 \in A \subseteq A'$ already, it suffices to show that A' is closed under addition and multiplication. Let $\alpha, \beta \in A'$. Then by the previous lemma, $A[\alpha]$ and $A[\beta]$ are finitely-generated A-modules. Note that since $A[X] \subseteq (A[\alpha])[X]$, the fact that β is integral over A implies that it is integral over $A[\alpha]$. So $A[\alpha, \beta]$ is finitely generated as an $A[\alpha]$-module by the lemma. Since $A[\alpha]$ is finitely generated as an A-module for the same reason, the pairwise products of the generators form a generating set for $A[\alpha, \beta]$ as an A-module, which is therefore finitely generated as an A-module. Since $\alpha + \beta$ and $\alpha\beta$ are in $A[\alpha, \beta]$, part (3) of Lemma 2.5 tells us that both of them are integral over A, as desired. □

Corollary 2.6

Let K be a number field. Then $\mathcal{O}_K$ is a ring.

§3 February 4, 2019

Last time, we proved that if $A \subseteq R$ are rings, then the integral closure of A in R is a subring of R. This was Proposition 2.3, and its proof relied mostly on Lemma 2.5.

§3.1 Are Integral Closures Finitely Generated?

Let the integral closure of A in some $R \supseteq A$ be A'. It is not necessarily true that A' is finitely generated as an A-module: for example, the algebraic closure of a finite field is not finite (thus cannot be finite-dimensional). Similarly, the integral closure of $\mathbf{Z}$ in $\overline{\mathbf{Q}}$ contains $\mathbf{Z}[\{\sqrt{d}\}_{d>1}]$ which is not finitely generated as a $\mathbf{Z}$-module.

Proposition 3.1

Let $A \subseteq B \subseteq C$ be rings, where C is integral over B and B is integral over A. Then C is integral over A.

Proof. Let $\alpha \in C$. It suffices to show that α is integral over A. Since it is integral over B, there are some $b_0, \dots, b_{n-1} \in B$ such that

$$b_0 + b_1\alpha + \cdots + b_{n-1}\alpha^{n-1} + \alpha^n = 0.$$

Consider the subring $B_1 = A[b_0, \dots, b_{n-1}] \subseteq B$. Since the b_i's are integral over A (because they are in B), repeated application of part (2) of Lemma 2.5 tells us that B_1 is a finitely-generated A-module [by taking pairwise products of generators, $A[b_0, b_1]$ being a finitely-generated $A[b_0]$-module and $A[b_0]$ being a finitely-generated A module implies that $A[b_0, b_1]$ is a finitely-generated A-module]. Since α is integral over B_1 (this is evident from the polynomial α satisfies), we actually know that $B_1[\alpha]$ is finitely generated over A.

Finally, $\alpha \in B_1[\alpha]$, so part (3) of Lemma 2.5 tells us that α is integral over A as desired. □

Recall from our computation of $\mathcal{O}_K$ in the case that K is a quadratic extension of $\mathbf{Q}$ that we had a basis of two elements for K. In general this should be true [in these notes we will frequently interchange the terms "abelian group" and "$\mathbf{Z}$-module"]:

Theorem 3.2

Let K be a number field and suppose that $\mathcal{O}_K$ is a finitely generated abelian group. Then $\mathcal{O}_K$ is a finitely-generated free abelian group of rank $[K : \mathbf{Q}]$.

Proof. For now, assume that $\mathcal{O}_K$ is a finitely generated abelian group (later we will abuse the properties of the *trace* on an algebraic number field to show this fact). We will show its rank is equal to what we want it to be.

Notice that $\mathcal{O}_K$ generates K as a $\mathbf{Q}$-vector space. This is for the following reason: if $\alpha \in K$, then it satisfies a monic polynomial

$$f(X) = X^d + a_{d-1}X^{d-1} + \cdots + a_0.$$

where the a_i's are rational. Clearing denominators, we can choose $m \in \mathbf{Z}$ such that $ma_i \in \mathbf{Z}$ for all i. Then

$$m^d f(X) = (mX)^d + ma_{d-1}(mX)^{d-1} + \cdots + m^d a_0$$

so $m\alpha \in \mathcal{O}_K$. Hence, all of K is contained in the $\mathbf{Q}$-span of $\mathcal{O}_K$, as desired. In more abstract terms, the multiplication map

$$\mathcal{O}_K \otimes_{\mathbf{Z}} \mathbf{Q} \to K$$

is surjective. In fact, the map is also injective. Consider an arbitrary element $\sum \alpha_i \otimes b_i \in \mathcal{O}_K \otimes_{\mathbf{Z}} \mathbf{Q}$. If it is in the kernel, then

$$0 = \sum \alpha_i b_i.$$

Taking $m \in \mathbf{Z}$ so that $mb_i \in \mathbf{Z}$ for all i, we know

$$0 = \sum \alpha_i (b_i m).$$

In $\mathcal{O}_K \otimes_{\mathbf{Z}} \mathbf{Q}$, it follows by tensoring with 1 that

$$0 = \sum \alpha_i \otimes (b_i m) = m \sum \alpha_i \otimes b_i$$

hence $\sum \alpha_i \otimes b_i = 0$ as desired.

Since $\mathcal{O}_K$ is torsion-free, the structure theorem for modules over a PID tells us that it is isomorphic as a $\mathbf{Z}$-module to $\mathbf{Z}^d$ for some $d \in \mathbf{N}$. Since we have isomorphisms of $\mathbf{Q}$-vector spaces

$$\mathcal{O}_K \otimes_{\mathbf{Z}} \mathbf{Q} \cong K \cong \mathbf{Q}^{[K:\mathbf{Q}]}$$

and

$$\mathcal{O}_K \otimes_{\mathbf{Z}} \mathbf{Q} \cong (\mathbf{Z} \otimes_{\mathbf{Z}} \mathbf{Q})^d = \mathbf{Q}^d$$

it follows that $d = [K : \mathbf{Q}]$ by the fact that finite-dimensional vector spaces are determined uniquely up to isomorphism by their dimension (K has $\mathbf{Q}$-linearly independent sets of size d and no larger). □

§3.2 The Trace and Norm

To show that $\mathcal{O}_K$ is a finitely-generated $\mathbf{Z}$-module and generally study its properties we will frequently exploit the trace and norm on K. In general, let $A \subseteq B$ be an extension of rings where B is finitely-generated as an A-module.

Definition 3.3. Let $\alpha \in B$ and let $\varphi_\alpha : B \to B$ be the A-linear map given by $x \mapsto x \cdot \alpha$. The **trace** of α is

$$\mathrm{Tr}_{B/A}(\alpha) = \mathrm{tr}(\varphi_\alpha).$$

The **norm** of α is

$$\mathrm{N}_{B/A}(\alpha) = \det(\varphi_\alpha).$$

We also denote by P_α the characteristic polynomial of φ_α.

Lemma 3.4
If $\alpha \in \mathcal{O}_K$, then $P_\alpha(X) \in \mathbf{Z}[X]$

Proof. Let $P_{\alpha,0}$ be the minimal polynomial of α over $\mathbf{Q}$. Then we can write

$$P_{\alpha,0} = \det(X - \varphi_\alpha|_{\mathbf{Q}(\alpha)})$$

since this polynomial is monic of degree $[\mathbf{Q}(\alpha) : \mathbf{Q}]$ and the minimal polynomial has this degree and divides it (by Cayley–Hamilton). Recall that $P_\alpha = \det(X - \varphi_\alpha|_K)$. If $v_1, \dots, v_n$ is a $\mathbf{Q}(\alpha)$-basis for K where $n = [K : \mathbf{Q}(\alpha)]$, then

$$K = \mathbf{Q}(\alpha) \cdot v_1 + \cdots + \mathbf{Q}(\alpha) \cdot v_n.$$

Then φ_α acts on K by acting on each copy of $\mathbf{Q}(\alpha)$, so the matrix for φ_α is block-diagonal, with the diagonal consisting of n identical copies of the matrix for φ_α acting on $\mathbf{Q}(\alpha)$ as a vector space over $\mathbf{Q}$. Since the determinant of a block-diagonal matrix is the product of the determinants of the blocks, it follows that $P_\alpha = P_{\alpha,0}^n$ (NB: this shows that the minimal polynomial and characteristic polynomial are much more tightly related in this case than in the general theory of finite-dimensional vector spaces). Since $\alpha \in \mathcal{O}_K$, Lemma 1.6 tells us that $P_{\alpha,0}$ and hence P_α have integer coefficients, as desired. □

NB: since the trace of an operator is the X^{d-1}-coefficient of the characteristic polynomial and the determinant is the constant coefficient, it follows from our observation that $P_\alpha = P_{\alpha,0}^n$ that the trace of α is n times the X^{d-1}-coefficient of the minimal polynomial for α; and the norm of α is the n-th power of the constant coefficient of the minimal polynomial. Moreover, by definition and the multiplicativity of degrees of field extensions,

$$n = [K : \mathbf{Q}(\alpha)] = [K : \mathbf{Q}]/d.$$

We will see these facts in a different light (i.e. without using the characteristic polynomial) in section.

§4 February 6, 2019

Last class, we showed in Theorem 3.2 that if K is a number field then $\mathcal{O}_K$ spans K as a $\mathbf{Q}$-vector space and it is a free abelian group of rank $[K : \mathbf{Q}]$ (assuming it is finitely-generated). The proof that $\mathcal{O}_K$ is finitely generated as an abelian group relies on the *trace pairing*.

§4.1 The Trace Pairing and Dual Subgroup

Definition 4.1. The **trace pairing** on K is the bilinear form

$$K \times K \to \mathbf{Q}$$

given by

$$\langle x, y \rangle := \mathrm{Tr}_{K/\mathbf{Q}}(xy).$$

Checking that the trace pairing is bilinear follows from the general properties of the trace (it is additive; this is because it is the trace of the matrix corresponding to multiplication by α).

Lemma 4.2

The trace pairing is nondegenerate, i.e. it identifies $K \cong K^\vee := \mathrm{Hom}_{\mathbf{Q}}(K, \mathbf{Q})$.

Proof. We need to review the facts about the duals of vector spaces. If V is a finite-dimensional vector space over, say, $\mathbf{Q}$, then the *dual space* of V is defined to be the vector space $\mathrm{Hom}_{\mathbf{Q}}(V, \mathbf{Q})$, i.e. the set of $\mathbf{Q}$-linear maps from V to $\mathbf{Q}$ endowed with the obvious $\mathbf{Q}$-linear structure. If $e_1, \ldots, e_n$ is a basis for V, then the action of any $f \in V^\vee$ is determined by its action on the e_i's so one can check that we have a basis $e_1^\vee, \ldots, e_n^\vee$ for $V^\vee$ given by $e_i^\vee(e_j) = \delta_{ij}$ [if you haven't seen this before it is a very useful exercise to work out why this is a basis for $V^\vee$]. Hence, V and $V^\vee$ have the same dimension (they are identified, albeit noncanonically).

Recall that a pairing sends a space to its dual via $x \mapsto \langle -, x \rangle$, where $\langle -, x \rangle$ denotes the element of $K^\vee$ taking α to $\langle \alpha, x \rangle$. Since the dual space of a finite-dimensional vector space has the same dimension (we just sketched that one can construct a dual basis given by $e_i^\vee(e_j) = \delta_{ij}$), it suffices to show that this map is injective, i.e. that for all nonzero $x \in K$, the map $\langle -, x \rangle \in K^\vee$ is nonzero, i.e. that for all nonzero $x \in K$ there exists an $\alpha \in K$ such that $\mathrm{Tr}_{K/\mathbf{Q}}(\alpha x) \neq 0$. Since $x \neq 0$, we can take $\alpha = x^{-1}$, and see that

$$\mathrm{Tr}_{K/\mathbf{Q}}(\alpha x) = \mathrm{Tr}_{K/\mathbf{Q}}(1) = [K : \mathbf{Q}] \neq 0$$

since the $\mathbf{Q}$-endomorphism of K given by multiplication by 1 corresponds to the identity matrix, which has $[K : \mathbf{Q}]$ 1's along the diagonal. The fact that this doesn't vanish in $\mathbf{Q}$ follows from the fact that $\mathbf{Q}$ has characteristic zero. □

NB: The above lemma actually holds for any finite *separable* field extension though the proof cannot be done in the same way if the characteristic divides $[K : \mathbf{Q}]$.

The identification $K \cong K^\vee$ means that we may regard the elements of $K^\vee$ as elements of K. In particular, it will be useful to consider a $\mathbf{Q}$-basis $e_1, \ldots, e_n$ for K, and consider its dual basis $e_1^\vee, \ldots, e_n^\vee \in K^\vee \cong K$. Observe that

$$\delta_{ij} = e_i^\vee(e_j) = \langle e_i^\vee, e_j \rangle = \mathrm{Tr}_{K/\mathbf{Q}}(e_i^\vee \cdot e_j) \tag{4.1}$$

where the second equality is by definition of the identification $K \to K^\vee$ [as the element $e_i^\vee \in K$ must be identified with $\langle -, e_i^\vee \rangle \in K^\vee$, so the action of $e_i^\vee$ considered as an element of $K^\vee$ must be by pairing with $e_i^\vee$].

Definition 4.3. For any (additive) subgroup $L \subset K$, the **dual group** of L is defined to be

$$L^\vee = \{\alpha \in K : \mathrm{Tr}_{K/\mathbf{Q}}(\alpha x) \in \mathbf{Z} \text{ for all } x \in L\}.$$

Note that such a dual subgroup can be defined for any subgroup of a vector space with respect to any nondegenerate bilinear form. Even without a bilinear form on V to relate it to V^vee you can define the dual of a subgroup to be a subgroup of the dual vector space. In general, the dual of a lattice in a finite-dimensional vector space is a lattice in the dual space. One can see this using the dual basis:

Lemma 4.4

If $L = \mathbf{Z}e_1 + \cdots + \mathbf{Z}e_n$, then

$$L^\vee = \mathbf{Z}e_1^\vee + \cdots + \mathbf{Z}e_n^\vee.$$

Proof. Let $\beta \in K$. Then we can write $\beta = \sum a_i e_i^\vee$ using the dual basis for $V^\vee$ and the identification $V \cong V^\vee$ given by the trace pairing. Recall that $\beta \in L^\vee$ is equivalent to $\mathrm{Tr}_{K/\mathbf{Q}}(\beta \cdot x) \in \mathbf{Z}$ for all $x \in L$, which is (using the basis for L) equivalent to

$$\mathrm{Tr}_{K/\mathbf{Q}}(\beta \cdot e_i) \in \mathbf{Z}$$

for all e_i's. Of course, for any i,

$$\mathrm{Tr}_{K/\mathbf{Q}}(\beta \cdot e_i) = \mathrm{Tr}_{K/\mathbf{Q}}\left(\sum a_j e_j^\vee e_i\right) = a_i$$

by (4.1) so $\beta \in L^\vee$ if and only if all the a_i's are in $\mathbf{Z}$, which is equivalent to $\beta \in \mathbf{Z}e_1^\vee + \cdots + \mathbf{Z}e_n^\vee$, as desired. □

§4.2 Finite-Generation of the Ring of Integers

Note that from the definition, if $L_1 \subseteq L_2$ are subgroups of K, then $L_2^\vee \subseteq L_1^\vee$ (taking dual subgroups is inclusion-reversing).

Proposition 4.5

Let K be a number field. Then $\mathcal{O}_K$ is finitely-generated as an abelian group.

Proof. Recall that if $e_1, \ldots, e_n$ is a basis for K over $\mathbf{Q}$, then we can scale them by integers to obtain a basis $e_1', \ldots, e_n'$ for K over $\mathbf{Q}$ such that each $e_i' \in \mathcal{O}_K$. Take $L := \mathbf{Z}e_1 + \cdots + \mathbf{Z}e_n$. By the previous lemma, the fact that taking dual subgroups is inclusion-reversing, and the fact that $\mathcal{O}_K$ is contained in its dual (it is closed under multiplication and any element of $\mathcal{O}_K$ has integral trace), we have inclusions

$$\mathbf{Z}e_1 + \cdots + \mathbf{Z}e_n \subseteq \mathcal{O}_K \subseteq \mathcal{O}_K^\vee \subseteq (\mathbf{Z}e_1 + \cdots + \mathbf{Z}e_n)^\vee = \mathbf{Z}e_1^\vee + \cdots + \mathbf{Z}e_n^\vee.$$

Hence, $\mathcal{O}_K$ is stuck between two free abelian groups of rank n, which means it is a free abelian group of rank n (this is part of the proof of the PID structure theorem; alternatively one can take the $\mathbf{Q}$-span and look at dimensions). We have finally concluded that not only is $\mathcal{O}_K$ finitely generated as an abelian group, but we also have an alternative proof of the fact that its rank is $n = [K : \mathbf{Q}]$. This proof also tells us that $\mathcal{O}_K^\vee$ is a free abelian group of rank n. □

NB: This proof relies heavvily on the fact that the trace pairing is nondegenerate (definitely true in characteristic zero). But if we take an extension which isn't separable then this fails. For example, let $K = \mathbf{F}_p(x^{1/p})$ considered as an extension of $\mathbf{F}_p$.

K has a basis given by $1, x^{1/p}, \dots, x^{(p-1)/p}$. The trace form is determined by where it takes basis vectors. In particular,

$$\mathrm{Tr}_{K/\mathbf{F}_p}(x^{i/p}, x^{j/p}) = \mathrm{Tr}_{K/\mathbf{F}_p}(x^{(i+j)/p}).$$

If $i + j \neq 0$ modulo p, then multiplication by $x^{(i+j)/p}$ permutes the basis with no fixed points, so the diagonal of the matrix given by multiplication by $x^{(i+j)/p}$ only contains zeroes. If $i + j = 0$ modulo p, then actually $x^{(i+j)/p} = x$ and the trace of x is $px = 0$, so the trace is identically zero and thus the trace form is degenerate. That being said, one can prove by other means that the integral closure of $\mathbf{F}_p[x^{1/p}]$ in $\mathbf{F}_p(x^{1/p})$ is finitely-generated as a module over $\mathbf{F}_p[x^{1/p}]$.

§4.3 The Discriminant of a Number Field

Since $\mathcal{O}_K \subseteq \mathcal{O}_K^\vee$ both have rank $n = [K : \mathbf{Q}]$, the inclusion has finite index.

Definition 4.6. The **discriminant** of K is, up to a sign, the index

$$[\mathcal{O}_K^\vee : \mathcal{O}_K] = |\mathcal{O}_K^\vee / \mathcal{O}_K|$$

The actual definition is as follows:

Definition 4.7. Let $\alpha_1, \dots, \alpha_n \in \mathcal{O}_K$ be a basis for it over $\mathbf{Z}$ (i.e. as an abelian group). The **discriminant** of K is defined by

$$D_K = \det(\mathrm{Tr}_{K/\mathbf{Q}}(\alpha_i \alpha_j)),$$

where $\mathrm{Tr}_{K/\mathbf{Q}}(\alpha_i\alpha_j)$ is the matrix whose (i,j)-entry is $\mathrm{Tr}_{K/\mathbf{Q}}(\alpha_i\alpha_j)$.

Doesn't it depend on the choice of basis? The answer is no. It is well-defined (any basis you choose will result in the same number for the discriminant). We will prove this, along with the equivalence of the two definitions, next class.

Today, we compute an example.

Example 4.8

Let $K = \mathbf{Q}(\sqrt{d})$, d squarefree, be a quadratic number field. If $d \not\equiv 1 \pmod 4$, then $\mathcal{O}_K = \mathbf{Z}[\sqrt{d}]$ has a basis given by $1, \sqrt{d}$. So

$$D_K = \det \begin{bmatrix} \mathrm{Tr}_{K/\mathbf{Q}}(d) & \mathrm{Tr}_{K/\mathbf{Q}}(\sqrt{d}) \\ \mathrm{Tr}_{K/\mathbf{Q}}(\sqrt{d}) & \mathrm{Tr}_{K/\mathbf{Q}}(1) \end{bmatrix} = \det \begin{bmatrix} 2d & 0 \\ 0 & 2 \end{bmatrix} = 4d.$$

If $d \equiv 1 \pmod 4$, then $\mathcal{O}_K = \mathbf{Z}[(1+\sqrt{d})/2]$ has a basis given by $1, (1+\sqrt{d})/2$. So

$$D_K = \det \begin{bmatrix} \mathrm{Tr}_{K/\mathbf{Q}}((1+\sqrt{d})^2/4) & (1+\sqrt{d})/2 \\ (1+\sqrt{d})/2 & 1 \end{bmatrix} = d.$$

Actually the example is a pretty good result:

Proposition 4.9

Let $K = \mathbf{Q}(\sqrt{d})$ where d is a squarefree integer. Then

$$D_K = \begin{cases} 4d, \text{ if } d \not\equiv 1 \pmod 4, \\ d, \text{ if } d \equiv 1 \pmod 4 \end{cases}.$$

The following is a preview of things to come; none of it is rigorously justified but it is meant to give an answer to Forrest's question "why is the discriminant useful": One way in which the discriminant is useful is that it controls the *ramification* of primes. One notion of ramification is the ramification of covering maps (e.g. of Riemann surfaces) in which a covering map is "d-ramified" at a point if it has fibers (preimages) of size d locally around the point, but only a single point in the fiber at the point (one example is the polynomial map $x \mapsto x^2$ on the Riemann sphere $\mathbf{C} \cup \{\infty\}$. Most points have two preimages, except zero and infinity have 1.)

The idea is that number fields should be viewed geometrically in the same way. It might make a little more sense why this is "geometric" if you know about affine schemes but we can still formally define the *ramification* phenomena of a number field. In fact, do not ask me about affine schemes because I do not know. For every prime p in $\mathbf{Z}$, there might be many prime ideals $\mathfrak{q}$ in $\mathcal{O}_K$ that contain p. Moreover, some more of the ideals

$$\mathfrak{q} \supset \mathfrak{q}^2 \supset \mathfrak{q}^3 \supset \cdots$$

might still be big enough to contain p. When this happens, we say that p *ramifies* in K. We will see later that the ideal $p\mathcal{O}_K$ factors uniquely into prime ideals in $\mathcal{O}_K$, and containment of ideals is equivalent to divisibility, so this is equivalent to asking whether a prime factor appears more than once in the factorization of p in $\mathcal{O}_K$. The study of ramification in a number field generates a rich and useful theory which includes both D_K and $\mathcal{O}_K^\vee$. The discriminant completely controls which primes in $\mathbf{Z}$ ramify in $\mathcal{O}_K$:

Theorem 4.10

Let K be a number field and p a prime in $\mathbf{Z}$. Then p ramifies in K if and only if $p | D_K$

For this reason, one of the most convenient ways to figure out which primes ramify is to compute D_K and then factor it. Similarly, we can take the "inverse of $\mathcal{O}_K^\vee$ as a fractional ideal of K" to get the *different* ideal in $\mathcal{O}_K$. In the same way as the discriminant, the different ideal controls which primes $\mathfrak{q}$ in $\mathcal{O}_K$ are ramified over $\mathfrak{q} \cap \mathbf{Z}$.

§5 February 11, 2019

Recall from last class: let K be a number field. Then we have the *complementary module* (last time we called it the "dual subgroup")

$$\mathcal{O}_K^\vee = \{\alpha \in K : \mathrm{Tr}_{K/\mathbf{Q}}(xy) \in \mathbf{Z}\}.$$

§5.1 Revisiting Discriminants and the Complementary Module

Last class we stated but did not prove the following equivalent formulation of the discriminant.

Lemma 5.1
$D(\underline{\alpha}) = [\mathcal{O}_K^\vee : \mathcal{O}_K]$, where $\underline{\alpha} = (\alpha_1, \dots, \alpha_n)$ is a $\mathbf{Z}$-basis for $\mathcal{O}_K$.

Before getting to this, we still need to show that the discriminant of a number field is well-defined. Then we will begin to set up some more theory about the discriminant Let $\alpha_1, \dots, \alpha_n$ and $\beta_1, \dots, \beta_n$ be $\mathbf{Q}$-bases for K. Take $\underline{\alpha}, \underline{\beta}$ to be the column vectors with the bases as entries. Let

$$D(\underline{\alpha}, \underline{\beta}) = \det(\mathrm{Tr}_{K/\mathbf{Q}}(\underline{\alpha} \cdot \underline{\beta}^t)),$$

so that the discriminant is just equal to $D(\underline{\alpha}, \underline{\alpha})$ where $\underline{\alpha}$ is a $\mathbf{Z}$-basis for $\mathcal{O}_K$.

Lemma 5.2
If $\mathfrak{M}$ is an $n \times n$ matrix with entries in $\mathbf{Q}$, and take $\underline{\alpha'} = \mathfrak{M} \cdot \underline{\alpha}$. Then

$$D(\underline{\alpha'}, \underline{\beta}) = \det \mathfrak{M} \cdot D(\underline{\alpha}, \underline{\beta}).$$

Proof. We just compute

$$\begin{aligned} D(\underline{\alpha'}, \underline{\beta}) &= \det(\mathrm{Tr}_{K/\mathbf{Q}}(\mathfrak{M} \cdot \underline{\alpha} \cdot \underline{\beta}^t)) \\ &= \det(\mathfrak{M}\, \mathrm{Tr}_{K/\mathbf{Q}}(\underline{\alpha} \cdot \underline{\beta}^t)) \\ &= (\det \mathfrak{M}) \cdot \det(\mathrm{Tr}_{K/\mathbf{Q}}(\underline{\alpha} \cdot \underline{\beta}^t)). \end{aligned}$$

The key step is taking $\mathfrak{M}$ out of the trace. This is allowed because the entries of $\mathfrak{M}$ are in $\mathbf{Q}$ and the trace is $\mathbf{Q}$-linear (you can write down the details for a small matrix and see that it works if you want). □

Corollary 5.3
If $\underline{\alpha}, \underline{\beta}$ are bases for K, then then up to a sign, $D(\underline{\alpha}, \underline{\beta})$ depends only on the modules $\oplus_{i=1}^n \mathbf{Z}\alpha_i$ and $\oplus_{i=1}^n \mathbf{Z}\beta_i$.

Proof. Let $\alpha_1', \dots, \alpha_n'$ be a different basis for the same $\mathbf{Z}$-module

$$\mathbf{Z}\alpha_1 + \cdots + \mathbf{Z}\alpha_n.$$

Then $\underline{\alpha}$ and $\underline{\alpha'}$ are related by an invertible matrix $\mathfrak{M}$ with integer coefficients, where $\mathfrak{M}^{-1}$ also has integer coefficients (this comes from writing down the fact that each basis

element can be written as a $\mathbf{Z}$-linear combination of the other basis). Hence, by the multiplicativity of the determinant, $\det \mathfrak{M} = \pm 1$ (in particular it must be a unit in $\mathbf{Z}$). By Lemma 5.2, it follows that $D(\underline{\alpha}', \underline{\beta}) = \pm D(\underline{\alpha}, \underline{\beta})$. The same argument holds for the second argument β, so indeed up to a sign $D(\underline{\alpha}, \underline{\beta})$ only depends on the $\mathbf{Z}$-spans of $\underline{\alpha}, \underline{\beta}$. □

Corollary 5.4

If $\alpha_1, \ldots, \alpha_n$ is a $\mathbf{Z}$-basis for $\mathcal{O}_K$, then the discriminant $D(\underline{\alpha})$ does not depend on the choice of $\underline{\alpha}$.

Proof. If there are two bases $\underline{\alpha}, \underline{\alpha}'$ then our work in the proof of Corollary 5.3 tells they are related by some matrix $\mathfrak{M}$ with determinant ± 1, so Lemma 5.2 yields

$$D(\underline{\alpha}') = D(\underline{\alpha}', \underline{\alpha}') = (\det \mathfrak{M})^2 D(\underline{\alpha}) = D(\underline{\alpha})$$

as desired. □

Corollary 5.4 shows that the "absolute discriminant" D_K of the number field K is well-defined as defined in Definition 4.7

Lemma 5.5

Suppose $L = \mathbf{Z}\alpha_1 + \cdots + \mathbf{Z}\alpha_n$ has a submodule $L' = \mathbf{Z}\alpha_1' + \cdots \mathbf{Z}\alpha_n'$ where both spanning sets are bases for $K/\mathbf{Q}$. Then if $\mathfrak{M}$ is the matrix with integer coefficients such that $\alpha' = \mathfrak{M} \cdot \alpha$, then the index of L' in L is

$$|L/L'| = [L : L'] = \pm \det \mathfrak{M}.$$

Proof. Changing α to another basis of L doesn't change either side of the equality $[L : L'] = \pm \det \mathfrak{M}$, since it multiplies $\mathfrak{M}$ by a matrix whose determinant is ± 1. So, it's enough to prove the lemma for any choice of basis for L. From the proof of the Structure Theorem for Modules over PID (recall from the first section or from Samuel, ch. 1), we know that there are bases $e_1, \ldots, e_n \in L$ and $f_1 e_1, \ldots, f_n e_n \in L'$ where the f_i's are integers. So, the index we want to compute is

$$|L/L'| = |e_1\mathbf{Z}/e_1 f_1 \mathbf{Z}| \cdots |e_n \mathbf{Z}/e_n f_n \mathbf{Z}| = \prod f_i = \det \mathfrak{M}$$

as desired. □

Big Brain Donut Proof. We might as well take V to be the $\mathbf{Q}$-vector space given by K and do this in slightly more generality. Tensoring up by $\mathbf{R}$, we have

$$V \otimes_{\mathbf{Q}} \mathbf{R} = \mathbf{Q}^n \otimes_{\mathbf{Q}} \mathbf{R} \cong \mathbf{R}^n.$$

the submodule $L' \subseteq L$ of V is a lattice in $V \otimes_{\mathbf{Q}} \mathbf{R}$. We get a map of donuts

$$(V \otimes_{\mathbf{Q}} \mathbf{R})/L' \to (V \otimes_{\mathbf{Q}} \mathbf{R})/L.$$

The areas of the donuts (i.e. the areas of fundamental parallelotopes of the lattices) are therefore related via the matrix $\mathfrak{M}$, namely by

$$\mathrm{vol}\left((V \otimes \mathbf{R})/L'\right) = [L : L'] \mathrm{vol}\left((V \otimes \mathbf{R})/L\right).$$

The geometric properties of the determinant tell us that the same equation is true if we replace $[L : L']$ with $|\det \mathfrak{M}|$ so we are done. □

Corollary 5.6
$D_K = \pm[\mathcal{O}_K^\vee : \mathcal{O}_K]$.

Proof. Recall from Proposition 2.4 and Theorem 3.2 that $\mathcal{O}_K$ has a $\mathbf{Z}$-basis $e_1, \dots, e_n$ and therefore by Lemma 4.4 $\mathcal{O}_K^\vee$ has a dual basis $e_1^\vee, \dots, e_n^\vee$ satisfying

$$\langle e_i, e_j^\vee \rangle = \mathrm{Tr}_{K/\mathbf{Q}}(e_i e_j^\vee) = \delta_{ij}.$$

Since $\mathcal{O}_K \subseteq \mathcal{O}_K^\vee$, the dual basis is related to the basis for $\mathcal{O}_K$ via

$$\underline{\alpha} = \mathfrak{M} \cdot \underline{\alpha}^\vee.$$

And we can compute the discriminant

$$\det(\mathrm{Tr}(\underline{\alpha} \cdot \underline{\alpha}^t)) = \det(\mathrm{Tr}(\mathfrak{M} \cdot \underline{\alpha}^\vee \cdot \underline{\alpha}^t)) = \det(\mathfrak{M}\,\mathrm{Tr}(\underline{\alpha}^\vee \cdot \underline{\alpha})) = \det \mathfrak{M} = \pm[\mathcal{O}_K^\vee : \mathcal{O}_K]$$

from the previous lemma. □

§5.2 Linearly Disjoint Fields

Let K, L be number fields, and $K \cdot L$ be their composite. What is $\mathcal{O}_{K \cdot L}$?

Example 5.7
We know $\mathcal{O}_K$ for $K = \mathbf{Q}(\sqrt{d})$. What is $\mathcal{O}_K$, where $K = \mathbf{Q}(\sqrt{d_1}, \sqrt{d_2})$?

Definition 5.8. We say that number fields K, L are **linearly disjoint** if

$$[K \cdot L : \mathbf{Q}] = [K : \mathbf{Q}] \cdot [L : \mathbf{Q}],$$

i.e. if the bound $[K \cdot L : \mathbf{Q}] \leq [K : \mathbf{Q}] \cdot [L : \mathbf{Q}]$ is attained with equality.

Note: the composite $K \cdot L$ is only well-defined because we can embed K and L into a common algebraic closure $\overline{\mathbf{Q}}$ and take the field generated by them.

Example 5.9
Take $K = \mathbf{Q}(2^{1/3})$ and $L = \mathbf{Q}(2^{1/3} \cdot \zeta_3)$. Both are defined as different subfields of $\mathbf{C}$, though they are isomorphic (computing the minimal polynomial of generating elements) as

$$L \cong K \cong \mathbf{Q}[x]/(x^3 - 2).$$

On the other hand,

$$L \cdot K = \mathbf{Q}(2^{1/3}, \zeta_3)$$

is the splitting field for $x^3 - 2$ and has degree 6 (you have to check that ζ_3 is not in $\mathbf{Q}(2^{1/3})$). In this case, we see that L, K are not linearly disjoint, and despite the fact that they are isomorphic, their composite properly contains both of them.

Note that L, K are linearly disjoint if and only if $L \otimes_{\mathbf{Q}} K$ is a field, and in this case it is isomorphic to $L \cdot K$ via the map $L \otimes_{\mathbf{Q}} K \to L \cdot K$ given by multiplication. This map is always surjective by definition (check this, e.g. by writing down bases). If $K \otimes_{\mathbf{Q}} L$ is

a field, then the map is injective since it is nonzero, so the isomorphism of $\mathbf{Q}$-algebras yields an equality of dimensions

$$[L:\mathbf{Q}][K:\mathbf{Q}] = \dim_{\mathbf{Q}}(L \otimes_{\mathbf{Q}} K) = \dim_{\mathbf{Q}}(L \cdot K),$$

which means K, L are linearly disjoint. Conversely, if K, L are linearly disjoint then the equality of dimensions tells us the map is injective and we are similarly done.

The following is a useful way to compute rings of integers, and we will prove it next time:

Proposition 5.10

Let K, L be linearly disjoint number fields and let d_K, d_K be their discriminants. Then

$$\mathcal{O}_{L\cdot K} \subseteq \frac{1}{\gcd(d_K, d_L)} \mathcal{O}_K \cdot \mathcal{O}_L$$

NB: We switched from the notation D_K to d_K for the discriminant of a number field. Other notations include Δ_K, $\mathfrak{d}_K$, δ_K.

Corollary 5.11

If L, K are linearly disjoint and their discriminants are coprime, then $\mathcal{O}_{K\cdot L} = \mathcal{O}_K \cdot \mathcal{O}_L$.

Example 5.12

$\mathcal{O}_{\mathbf{Q}(\sqrt{5},\sqrt{7})} = \mathbf{Z}\left[\frac{1+\sqrt{5}}{2}, \sqrt{7}\right]$, since $d_{\mathbf{Q}(\sqrt{5})} = 5$ and $d_{\mathbf{Q}(\sqrt{7})} = 28$ are coprime by Proposition 4.9 and we (you) can check that $\mathbf{Q}(\sqrt{5}), \mathbf{Q}(\sqrt{7})$ are linearly disjoint.

Example 5.13

Computing discriminants, the proposition gives

$$\mathcal{O}_{\mathbf{Q}(\sqrt{7},\sqrt{11})} \subseteq \frac{1}{4}\mathbf{Z}[\sqrt{7}, \sqrt{11}].$$

Actually, $(\sqrt{7}-\sqrt{11})/2$ is in the ring of integers (check this by computing the minimal polynomial over $\mathbf{Z}[\sqrt{11}]$; recall from Proposition 3.1 that if α is integral over B and B is integral over A then α is integral over A; alternatively one can take the minimal polynomial and multiply it by its conjugate to get a minimal polynomial with integer coefficients by Galois theory – the elements of $\mathrm{Gal}(\mathbf{Q}(\sqrt{11})/\mathbf{Q})$ must all fix the coefficients so the coefficients are in $\mathbf{Q} \cap \mathcal{O}_{\mathbf{Q}(\sqrt{11})} = \mathbf{Z}$).

§6 February 13, 2019

Last time, we showed that the discriminant of a number field is well-defined (it does not depend on the $\mathbf{Z}$-basis of $\mathcal{O}_K$ that you use to define it). We also stated a nice result: the discriminant is equal (up to a sign) to the index $[\mathcal{O}_K^\vee : \mathcal{O}_K]$. This time, we filled in a second (donut-related) proof of a lemma related to this fact (the proof we used last time basically used the structure theorem) and wrote down the proof of the result itself. All of these are added to the notes from last class.

§6.1 Rings of Integers of Composites of Linearly Disjoint Fields

Proposition 6.1

Let L/K be number fields, and recall we have the trace map $\mathrm{Tr}_{L/K} : L \to K$. Then $\mathrm{Tr}_{L/K}(\mathcal{O}_L) \subseteq \mathcal{O}_K$.

Proof. First we show a general version of the Gauss Lemma:

Lemma 6.2 (Gauss's Lemma for $\mathcal{O}_K$)

If $f = gh$ where $f \in \mathcal{O}_K[X]$ and $g, h \in K[X]$ where all three are monic, then $g, h \in \mathcal{O}_K[X]$.

Proof. Let $\tilde{K}$ be a splitting field for f. If $\tilde{\alpha} \in \tilde{K}$ is a root of f, then $\tilde{\alpha}$ is in the integral closure of $\mathcal{O}_K$ in $\tilde{K}$. The transitivity of integral closures (see Proposition 3.1) tells us that $\tilde{\alpha}$ is in the integral closure of $\mathbf{Z}$ in $\tilde{K}$ which is $\mathcal{O}_{\tilde{K}}$. So g splits in $\tilde{K}$ and its roots are a subset of the roots of f,

$$g = \prod_{\alpha \in S} (X - \alpha) \in \mathcal{O}_{\tilde{K}}[X] \cap K[X] = \mathcal{O}_K[X]$$

and the same argument works for f. The fact that $\mathcal{O}_{\tilde{K}} \cap K = \mathcal{O}_K$ comes from the fact that $\mathcal{O}_K$ is the integral closure of $\mathcal{O}_K$ in $\tilde{K}$. □

Let $\alpha \in \mathcal{O}_K$. The trace of α is a coefficient of $P_\alpha(X) = \det_K(X - \alpha|_L)$. Recall that this is the $[L : K(\alpha)]$-th power of the minimal polynomial of α over K. Since $\alpha \in \mathcal{O}_K$, its minimal polynomial over K has coefficients in $\mathcal{O}_K$ which means that its trace is in $\mathcal{O}_K$, as desired. The fact that $P_{\alpha,0} \in \mathcal{O}_K$ comes from Gauss's lemma: we know $P_{\alpha,0} | P_{(\alpha,0)/\mathbf{Q}}$, and $P_{(\alpha,0)/\mathbf{Q}}$ has coefficients in $\mathcal{O}_K$, so the same is true of $P_{\alpha,0}$. □

Suppose K'/K are number fields, $M \subseteq K'$ an $\mathcal{O}_K$-module. M has a complementary module with respect to K:

Definition 6.3. The **complementary module** of M with respect to K is

$$M^{\vee_K} = \{\alpha \in K' : \mathrm{Tr}_{K'/K}(\alpha m) \in \mathcal{O}_K \text{ for all } m \in M\}$$

NB: If M is an $\mathcal{O}_K$-module, then so is $M^{\vee_K}$. Now we have the tools to prove the big proposition from last class.

Proposition 6.4

If K, L are linearly disjoint fields, then

$$\mathcal{O}_{K\cdot L} \subseteq \frac{1}{\gcd(d_K, d_K)}\mathcal{O}_K \cdot \mathcal{O}_L.$$

Proof. Let $e_1, \dots, e_n$ be a $\mathbf{Z}$-basis for $\mathcal{O}_K$. Then (recall $K \otimes_{\mathbf{Q}} L \cong K \cdot L$ is a field and through this isomorphism we have $\mathcal{O}_K \cdot \mathcal{O}_L \cong \mathcal{O}_K \otimes_{\mathbf{Z}} \mathcal{O}_L$; frequently we will interchange these two notations)

$$\mathcal{O}_K \otimes_{\mathbf{Z}} \mathcal{O}_L = e_1 \cdot \mathcal{O}_L + \cdots + e_n \cdot \mathcal{O}_L.$$

So

$$K \cdot L \cong K \otimes_{\mathbf{Q}} L = e_1 L + \cdots + e_n L,$$

i.e. $e_1, \dots, e_n$ is an L-basis for $K \cdot L$. Let's compute

$$(\mathcal{O}_K \otimes_{\mathbf{Z}} \mathcal{O}_L)^{\vee_K} = \{\alpha \in L \cdot K : \mathrm{Tr}_{L\cdot K/K}(\alpha\beta) \in \mathcal{O}_K \text{ for all } \beta \in \mathcal{O}_L \otimes_{\mathbf{Z}} \mathcal{O}_K\}.$$

For $\alpha \in L \cdot K$, we have

$$\begin{aligned} \alpha \in (\mathcal{O}_K \otimes_{\mathbf{Z}} \mathcal{O}_L)^{\vee_K} &\iff \mathrm{Tr}_{L\cdot K/K}(\alpha\beta) \in \mathcal{O}_K \text{ for all } \beta \in \mathcal{O}_L \otimes \mathcal{O}_K. \\ &\iff \mathrm{Tr}_{L\cdot K/K}(\alpha\beta) \in \mathcal{O}_K \text{ for all } \beta \in \mathcal{O}_L \end{aligned}$$

(this equivalence follows from our $\mathcal{O}_L$-basis for $\mathcal{O}_K \otimes \mathcal{O}_L$, NB every element of this basis is in $\mathcal{O}_K$). Note that if $\alpha = \sum \alpha_i e_i$ where $\alpha_i \in \mathcal{O}_L$, then

$$\mathrm{Tr}_{L\cdot K/K}(\alpha\beta) = \mathrm{Tr}_{L\cdot K/K}(\sum e_i \alpha_i \beta) = \sum e_i \,\mathrm{Tr}_{L\cdot K/K}(\alpha_i\beta) = \sum e_i \,\mathrm{Tr}_{L/\mathbf{Q}}(\alpha_i\beta).$$

So (using the fact that the e_i's are a $\mathbf{Z}$-basis for $\mathcal{O}_K$) all the above is equivalent to

$$\mathrm{Tr}_{L/\mathbf{Q}}(\alpha_i\beta) \in \mathbf{Z}$$

for all i (for all $\beta \in \mathcal{O}_L$), which is equivalent to $\alpha_i \in \mathcal{O}_L^\vee$ for all i. The fact that $\mathrm{Tr}_{L\cdot K/K} = \mathrm{Tr}_{L/\mathbf{Q}}$ on L comes from recalling that the trace is just the trace of the multiplication map, which has the same matrix either way.

The computation above amounts to the following:

$$(\mathcal{O}_K \otimes \mathcal{O}_L)^{\vee_K} = \mathcal{O}_K \otimes_{\mathbf{Z}} \mathcal{O}_L^\vee$$

[because we showed being in the dual is equivalent to all the L-components being in the dual of $\mathcal{O}_L$] We know that $\mathcal{O}_K \otimes \mathcal{O}_L = \mathcal{O}_K \cdot \mathcal{O}_L \subseteq \mathcal{O}_K \cdot \mathcal{O}_K$. Taking complementary modules with respect to K, we have

$$\mathcal{O}_{L\cdot K} \subseteq \mathcal{O}_{L\cdot K}^{\vee_K} \subseteq (\mathcal{O}_K \otimes \mathcal{O}_L)^{\vee_K} = \mathcal{O}_K \otimes \mathcal{O}_L^\vee.$$

Since a group is always killed by its order, we know that (by Corollary 5.6) d_L sends $\mathcal{O}_L^\vee$ to $\mathcal{O}_L$ which means $\mathcal{O}_L^\vee \subseteq \frac{1}{d_L} \cdot \mathcal{O}_L$. Applying the whole argument again except switching the roles of K and L, we get

$$\mathcal{O}_{K\cdot L} \subseteq \frac{1}{d_L}\mathcal{O}_K \cdot \mathcal{O}_L \cap \frac{1}{d_K}\mathcal{O}_K \cdot \mathcal{O}_L$$

which implies the proposition. □

§7 February 20, 2019

Today Mark Kisin is travelling, so his student Robert Cass is teaching class.

§7.1 Trace, Norm and Discriminant using Embeddings

Though the definitions of the norm and trace we have been using are somewhat more natural (see Definition 3.3), it's much more useful to define them in the following way. Let L/K be an extension of number fields of degree n, and let $\overline{K}$ be an algebraic closure of K (or indeed any algebraically closed field containing K). Then the primitive element theorem tells us $L = K[\alpha]$ for some $\alpha \in L$, so any embedding $\sigma : L \to \overline{K}$ is determined by where it sends α. Moreover, it must send α to one of the roots of $f_{\alpha,0}$ in $\overline{K}$ (this is a calculation using the fact that σ is a homomorphism and $f_{\alpha,0}(\alpha) = 0$). For each root α' we have an isomorphism

$$K[\alpha] \cong K[\alpha'] \cong K[x]/(f_{\alpha,0})$$

and these are all included in $\overline{K}$, so actually we have one embedding of $K[\alpha]$ into $\overline{K}$ for each root $\alpha' \in \overline{K}$ of $f_{\alpha,0}$. This polynomial splits completely since $\overline{K}$ is algebraically closed and has characteristic zero, which means there are exactly $[L : K]$ distinct embeddings $\sigma : L \to \overline{K}$. One way to define the trace and norm is using sums and products of these embeddings:

Definition 7.1. Let L/K be an extension of number fields and $\alpha \in L$. The **trace** of α is

$$\mathrm{Tr}_{L/K}(\alpha) := \sum_{\sigma : L \to \overline{K}} \sigma(\alpha).$$

The **norm** of α is

$$N_{L/K}(\alpha) := \prod_{\sigma : L \to \overline{K}} \sigma(\alpha)$$

Sometimes when we want to emphasize the fact that K might not be $\mathbf{Q}$ we call these the **relative trace** and **relative norm**. We still need to show that this formulation is equivalent to the one given in Definition 3.3:

Proposition 7.2

Let L/K be an extension of number fields and $\alpha \in L$. Then $\mathrm{Tr}_{L/K}(\alpha)$ is equal to the trace of the K-vector space endomorphism of L given by multiplication by α. The norm $N_{L/K}(\alpha)$ is the same as its determinant.

Proof. This will be a homework problem. □

Since we have a new definition of the trace, we can obtain a possibly more convenient definition of the discriminant.

Lemma 7.3

Let $\underline{\alpha} = (\alpha_1, \ldots, \alpha_n)$ be a basis for L over K. Then

$$D(\underline{\alpha}) = \det(\sigma_j(\alpha_i))^2$$

where σ_j runs over all n embeddings of L into $\overline{K}$.

Proof. Just compute

$$\begin{aligned} D(\underline{\alpha}) &= \det(\sum_s \sigma_s(\alpha_i\alpha_j)) \\ &= \det\left(\sum_s \sigma_s(\alpha_i)\sigma_s(\alpha_j)\right) \\ &= \det((\sigma_i(\alpha_j))^{\top}_{i,j}(\sigma_i(\alpha_j))_{i,j}) \\ &= \det(\sigma_j(\alpha_i))^2 \end{aligned}$$

by basic properties of the determinant. □

Corollary 7.4

Suppose $L = K[x]$ where $x \in L$ and x has minimal polynomial f of degree n. Then

$$D(1, \dots, x^{n-1}) = (-1)^{\frac{n(n-1)}{2}} N_{L/K}(f'(x)).$$

Proof. Relabelling $\sigma_i(x)$ as x_i, we can write

$$\det(\sigma_i(x^j))^2 = \det(x_i^j)^2.$$

This is called the Vandermonde determinant. It's an exercise in permuting rows and columns of matrices to show that this determinant is

$$\prod_{i<j}(x_i - x_j)^2.$$

[NB this is the *discriminant* of the polynomial f, since the x_i's are exactly the roots of f] We can rewrite this product as

$$(-1)^{\frac{n(n-1)}{2}} \prod_{i\neq j}(x_i - x_j) = \pm \prod_i \prod_{j\neq i}(x_i - x_j).$$

Since the x_i's are just the roots of f, we have (using the product rule)

$$f'(X) = \prod_i \prod_{j\neq i}(X - x_j).$$

So, using the fact that $x_i = \sigma_i(x)$, we know (since f' is just a polynomial with coefficients in K)

$$\sigma_i(f'(x)) = f'(x_i) = \prod_i \prod_{j\neq i}(x_i - x_j)$$

so that from before,

$$D(\underline{\alpha}) = (-1)^{\frac{n(n-1)}{2}} \prod_{i\neq j}(x_i - x_j) = (-1)^{\frac{n(n-1)}{2}} \prod_{i=1}^{n} \sigma_i(f'(x)) = (-1)^{\frac{n(n-1)}{2}} N_{L/K}(f'(x))$$

as desired. □

§7.2 Cyclotomic Fields and their Rings of Integers

Let p be an odd prime, and consider the **cyclotomic field** $K = \mathbf{Q}(\zeta_p)$, where ζ_p denotes a primitive p-th root of unity. What is $[K : \mathbf{Q}]$?

Lemma 7.5
The minimal polynomial of ζ_p is $\frac{X^p-1}{X-1} = X^{p-1} + \cdots + 1$.

Proof. We abuse Eisenstein's criterion:

Lemma 7.6 (Eisenstein's criterion)
Let $f(x) \in \mathbf{Z}[x]$ be monic and given by $f(x) = x^n + a_{n-1}x^{n-1} + \cdots a_0$. If $p|a_i$ for all i but p^2 does not divide a_0, then f is irreducible.

Proof. Suppose $f(x) = g(x)h(x)$. By Gauss's lemma, we can assume g, h are monic with integer coefficients. Reducing the coefficients mod p (applying the projection $\mathbf{Z}[x] \to \mathbf{F}_p[x]$) we have

$$\overline{f}(x) = x^n = \overline{g}(x)\overline{h}(x).$$

Since $\mathbf{F}_p$ is a field we know $\overline{g}, \overline{h}$ must be (positive) powers of x. This is impossible because the constant term of f cannot be divisible by p twice. □

To show the p-th cyclotomic polynomial is monic it suffices to show that $f(x+1)$ is irreducible. Using the binomial theorem,

$$f(x+1) = \frac{(x+1)^p - 1}{x} = \frac{x^p + px^{p-1} + \cdots + px}{x} = x^{p-1} + px^{p-2} + \cdots + p$$

which satisfies Einenstein's criterion. Hence, $f(x+1)$ and thus $f(x)$ is irreducible. □

So it is clear that $[K : \mathbf{Q}] = p - 1$.

Proposition 7.7
K is Galois with Galois group $\mathrm{Gal}(K/\mathbf{Q}) \cong (\mathbf{Z}/p\mathbf{Z})^\times$.

Proof. Note that K is generated only by roots of f. Moreover, all the roots of f are contained in K (the roots of unity are all powers of ζ_p), so K is a splitting field for f. It follows that K is a Galois extension of $\mathbf{Q}$.

Let $i \in (\mathbf{Z}/p\mathbf{Z})^\times$. We can send i to the unique automorphism of K that sends ζ_p to ζ_p^i. This yields the desired isomorphism

$$(\mathbf{Z}/p\mathbf{Z})^\times \to \mathrm{Gal}(K/\mathbf{Q}),$$

since it is clearly injective and we know that $|\mathrm{Gal}(K/\mathbf{Q})| = [K : \mathbf{Q}] = p - 1$ by the fact that K is Galois. □

What is the ring of integers of K? Actually it is what you expect.

Proposition 7.8

The ring of integers of $\mathbf{Q}(\zeta_p)$ is $\mathcal{O}_K = \mathbf{Z}[\zeta_p]$.

We are not going to prove this right away; first there will be computations. Fix an algebraic closure $\overline{\mathbf{Q}}/\mathbf{Q}$. Since $[K : \mathbf{Q}] = p - 1$ there are $p - 1$ embeddings of K into $\overline{\mathbf{Q}}$. The embeddings of ζ_p are all distinct roots of f, so we can write

$$f(x) = \prod_{\sigma : K \to \overline{\mathbf{Q}}} (x - \sigma(\zeta_p)) = \prod_{i=1}^{p-1} (x - \zeta_p^i).$$

Actually, $f(x)$ is the characteristic polynomial of the multiplication map by ζ_p (one way to see this is to recall that the characteristic polynomial is the $[K : \mathbf{Q}(\zeta_p)]$-th power of the minimal polynomial, but this power is equal to 1). So (being satisfied with this or invoking the alternate Definition 7.1 and the homework proof of its equivalence) the trace of ζ_p is just the sum of these embeddings,

$$\mathrm{Tr}_{K/\mathbf{Q}}(\zeta_p) = \sum_{i=1}^{p-1} \zeta_p^i = -1.$$

Since raising ζ_p to a (nonzero mod p) power just permutes the embeddings, we have

$$\mathrm{Tr}_{K/\mathbf{Q}}(\zeta_p^i) = -1$$

for any $i \in (\mathbf{Z}/p\mathbf{Z})^\times$. Similarly, we can compute (using the additivity of the trace)

$$\mathrm{Tr}_{K/\mathbf{Q}}(1 - \zeta_p^i) = p.$$

For the norm, we have a similar formula as the product of the Galois conjugates,

$$N_{K/\mathbf{Q}}(1 - \zeta_p^i) = \prod_{j=1}^{p-1} (1 - \zeta_p^j) = f(1) = p.$$

Lemma 7.9

Let $p \in \mathbf{Z}$ and K be as above. Then

$$(1 - \zeta_p)\mathcal{O}_K \cap \mathbf{Z} = p\mathbf{Z}.$$

Proof. We observe that

$$p = f(1) \in (1 - \zeta_p)\mathcal{O}_K \cap \mathbf{Z}$$

since $f(1) = \prod(1 - \zeta_p^i)$ and $\zeta_p^i \in \mathcal{O}_K$. Since $p\mathbf{Z}$ is a maximal ideal in $\mathbf{Z}$, it follows that $(1 - \zeta_p)\mathcal{O}_K \cap \mathbf{Z}$ is either $\mathbf{Z}$ or $p\mathbf{Z}$.

Moreover, $1 - \zeta_p$ is not a unit (its norm is $p \neq \pm 1$; you should check that something is a unit iff its norm is a unit in $\mathbf{Z}$). So, $(1 - \zeta_p)\mathcal{O}_K$ is a proper ideal, i.e. it does not contain 1, which means $(1 - \zeta_p)\mathcal{O}_K \cap \mathbf{Z}$ also does not contain 1. Hence this ideal is equal to $p\mathbf{Z}$ as desired. □

Corollary 7.10
If $y \in \mathcal{O}_K$, then $\mathrm{Tr}_{K/\mathbf{Q}}(y(1-\zeta_p)) \in p\mathbf{Z}$.

Proof. Writing it as a sum of Galois conjugates,

$$\mathrm{Tr}_{K/\mathbf{Q}}(y(1-\zeta_p)) = \sum_\sigma \sigma(y)(1-\zeta_p^i).$$

For any $i \in (\mathbf{Z}/p\mathbf{Z})^\times$, we can compute

$$\frac{1-\zeta_p^i}{1-\zeta_p} = 1 + \cdots + \zeta_p^{i-1} \in \mathcal{O}_K$$

which means that the elements $1-\zeta_p^i$ all generate the same ideal (they divide each other pairwise over $\mathcal{O}_K$). Since $\sigma(y) \in \mathcal{O}_K$ for all σ [it is integral over $\mathbf{Z}$ because it satisfies the same monic polynomial as y; it is in K because y is a $\mathbf{Q}$-linear combination of powers of ζ_p, so σ takes it to another such linear combination], we know that this trace is in the ideal $(1-\zeta_p)\mathcal{O}_K$. From Lemma 7.9, it follows that $\mathrm{Tr}_{K/\mathbf{Q}}(y(1-\zeta_p)) \in p\mathbf{Z}$ as desired. □

Now we can compute the ring of integers of K.

Proposition 7.11
The ring of integers of $K = \mathbf{Q}(\zeta_p)$ is

$$\mathcal{O}_K = \mathbf{Z}[\zeta_p].$$

Proof. Let $x = a_0 + a_1\zeta_p + \cdots + a_{p-2}\zeta_p^{p-2} \in \mathcal{O}_K$. Then

$$x(1-\zeta_p) = a_0(1-\zeta_p) + a_1\zeta_p(1-\zeta_p) + \cdots + a_{p-2}\zeta_p^{p-2}(1-\zeta_p).$$

Taking the trace,

$$\mathrm{Tr}_{K/\mathbf{Q}}(x(1-\zeta_p)) = a_0 p.$$

This is because the positive (and less than p-th) powers of ζ_p all have the same trace (they are all primitive p-th roots of unity), so the traces of all of the terms except the first one vanish. This trace has to be in $p\mathbf{Z}$ by Corollary 7.10, so $a_0 \in \mathbf{Z}$. We can repeat the same argument on $(x-a_0)\zeta_p^{-1} \in \mathcal{O}_K$ (NB ζ_p^{-1} can be written as a positive power of ζ_p so this is definitely in $\mathcal{O}_K$) to get that all of the a_i's are in $\mathbf{Z}$, which yields the desired result. □

We are interested in the following question: The Galois group of K has order $p-1$, so it has a unique subgroup of index 2 (this is a general fact about cyclic groups) and hence a quadratic subfield $\mathbf{Q}(\sqrt{d})$. What is this subfield? This question has some very nice applications for number theory but we won't get to the answer today.

§8 February 25, 2019

Let p be an odd prime (the choice $p = 2$ yields $\zeta_2 = -1$ so $\mathbf{Q}(\zeta_2) = \mathbf{Q}$ is not interesting). Last class we were looking at the number field $K = \mathbf{Q}(\zeta_p)$, and we determined its ring of integers $\mathcal{O}_K = \mathbf{Z}[\zeta_p]$. So, $\mathcal{O}_K$ has a $\mathbf{Z}$-basis $1, \dots, \zeta_p^{p-2}$. We also noticed that the discriminant of such a basis is, up to a sign, equal to $N_{K/\mathbf{Q}}(f'(\zeta_p))$ where f is the minimal polynomial for ζ_p over $\mathbf{Q}$.

Also, recall that the Galois group of $\mathbf{Q}(\zeta_p)/\mathbf{Q}$ is $(\mathbf{Z}/p\mathbf{Z})^\times$ which we know from algebra is a cyclic group of order $p-1$. When p is odd, it's a fact from algebra that this group has exactly one subgroup of index 2, so by Galois theory $\mathbf{Q}(\zeta_p)$ has a unique quadratic subfield. We would like to compute it.

§8.1 Discriminant of Cyclotomic Field

In Corollary 7.4, we can plug in the minimal polynomial for ζ_p to obtain:

Corollary 8.1

If $L = \mathbf{Q}(\zeta_p)$, then

$$d_L = \pm p^{p-2}.$$

Proof. By Proposition 7.11, $\mathcal{O}_L$ has a $\mathbf{Z}$-basis given by powers of ζ_p, namely $1, \dots, \zeta_p^{p-2}$. So the discriminant we want to compute is just the discriminant of this tuple. Recall that the minimal polynomial of ζ_p is

$$f(X) = 1 + X + \cdots + X^{p-1} = \frac{X^p - 1}{X - 1}.$$

Using the product rule on the fact that $X^p - 1 = (X-1)f(X)$, we have

$$pX^{p-1} = f'(X)(X-1) + f(X)$$

and substituting ζ_p for X yields

$$f'(\zeta_p)(\zeta_p - 1) = p\zeta_p^{p-1}.$$

Computing the norm,

$$N_{L/\mathbf{Q}}(f'(\zeta_p)) = N_{L/\mathbf{Q}}\left(\frac{p\zeta_p^{p-1}}{\zeta_p - 1}\right) = p^{p-1}\frac{N_{L/\mathbf{Q}}(\zeta_p)^{p-1}}{N_{L/\mathbf{Q}}(\zeta_p - 1)} = \pm p^{p-2}$$

from the norms we computed last class [$N_{L/\mathbf{Q}}(\zeta_p) = 1$ and $N_{L/\mathbf{Q}}(\zeta_p - 1) = p$]. From Corollary 7.4, actually we know

$$\begin{aligned} d_L &= \pm N_{L/\mathbf{Q}}(f'(\zeta_p)) \\ &= \pm p^{p-2} \end{aligned}$$

as desired. □

§8.2 More Facts about the Discriminant

The following makes sense given the analogy between ramification of number fields and branch points of Riemann surfaces: the discriminant d_L controls ramification of $\mathcal{O}_L$ over $\mathbf{Z}$, and d_K controls ramification of $\mathcal{O}_K$ over $\mathbf{Z}$. If $K \subseteq L$, then geometric intuition tells us that a "branch point" of $\mathcal{O}_K$ over $\mathbf{Z}$ is forced to also be a "branch point" of $\mathcal{O}_L$ over $\mathbf{Z}$. Since (we will see) the primes in $\mathbf{Z}$ which ramify upstairs in K are exactly those dividing d_K, and those which ramify upstairs in L are exactly those dividing d_L, we should expect from this intuition that $d_K | d_L$.

That being said, there isn't anything more than an analogy between branch points of Riemann surfaces and ramification of primes in number fields:

- "If you proved that ramificaiton of primes in number fields literally had something to do with Riemann surfaces, you would probably win a Fields medal."

As usual, for now none of this intuition is rigorously justified; it is just meant to serve as an inspiring picture:

- "I won't even bother asking if there are any questions; you are either inspired, or you are not."

First, we do a long-overdue lemma on the transitivity of the trace in towers (NB: the analogous result for norms is also true by the same proof)

Lemma 8.2

Let $K \subset L \subset M$ be fields of characteristic zero, with $[M : K] < \infty$. Then for $\alpha \in M$,

$$\mathrm{Tr}_{M/K} = \mathrm{Tr}_{L/K}(\mathrm{Tr}_{M/L}(\alpha))$$

Proof. By the homework, the trace is just the sum of the embeddings into an algebraic closure of the bottom field, so

$$\mathrm{Tr}_{M/K}(\alpha) = \sum_{\sigma : M \to \overline{K}} \sigma(\alpha).$$

Each embedding σ of M into $\overline{K}$ restricts to an embedding $\tau : L \to \overline{K}$. Hence, the embeddings $\sigma : M \to \overline{K}$ are precisely the embeddings extending each $\tau : L \to \overline{K}$. Such σ can be viewed as the embeddings of M into an algebraically closed field containing L (since τ embeds L in a fixed way into $\overline{K}$; recall that it only matters that the embeddings are to an algebraically closed field containing L. This field does not have to be "*the* algebraic closure" of L). In other words,

$$\mathrm{Tr}_{M/K}(\alpha) = \sum_{\tau : L \to \overline{K}} \sum_{\substack{\sigma : M \to \overline{K} \\ \text{restricting to } \tau}} \sigma(\alpha) = \sum_{\tau : L \to \overline{K}} \tau(\mathrm{Tr}_{M/L}(\alpha)) = \mathrm{Tr}_{L/K}\, \mathrm{Tr}_{M/L}(\alpha)$$

as desired. □

Using the transitivity of the trace in towers, we can formalize our intuition about discriminants of extensions of number fields.

Lemma 8.3

Suppose $\mathbf{Q} \subseteq K \subseteq L$ are number fields. Then $d_K | d_L$.

Proof. If $\alpha \in \mathcal{O}_K^\vee$ and $\beta \in \mathcal{O}_L$, then

$$\mathrm{Tr}_{L/\mathbf{Q}}(\alpha\beta) = \mathrm{Tr}_{K/\mathbf{Q}} \circ \mathrm{Tr}_{L/K}(\alpha\beta) = \mathrm{Tr}_{K/\mathbf{Q}}(\alpha \,\mathrm{Tr}_{L/K}(\beta))$$

since $\alpha \in K$. Since $\beta \in \mathcal{O}_L$, the trace on the inside is in $\mathcal{O}_K$ by Proposition 6.1, hence the trace we are interested in is in $\mathbf{Z}$ by the fact that $\alpha \in \mathcal{O}_K^\vee$. This means that $\alpha \in \mathcal{O}_L^\vee$, since we just showed $\mathrm{Tr}_{L/\mathbf{Q}}(\alpha\beta) \in \mathbf{Z}$ for all $\beta \in \mathcal{O}_L$; it follows that $\mathcal{O}_K^\vee \subset \mathcal{O}_L^\vee$.

Notice that

$$\mathcal{O}_L \cap \mathcal{O}_K^\vee \subseteq \mathcal{O}_L \cap K = \mathcal{O}_K$$

but $\mathcal{O}_L \cap \mathcal{O}_K^\vee$ also clearly contains $\mathcal{O}_K$ so in fact $\mathcal{O}_L \cap \mathcal{O}_K^\vee = \mathcal{O}_K$. The result is that we have an inclusion of finite abelian groups

$$\mathcal{O}_K^\vee/\mathcal{O}_K \to \mathcal{O}_L^\vee/\mathcal{O}_L$$

which tells us that the orders of these groups divide each other, i.e.

$$[\mathcal{O}_K^\vee : \mathcal{O}_K] \mid [\mathcal{O}_L^\vee : \mathcal{O}_L].$$

By Corollary 5.6, this is equivalent to the desired $d_K \mid d_L$. □

§8.3 What is the Quadratic Subfield of $\mathbf{Q}(\zeta_p)$?

We can use the machinery of the discriminant to deduce, with great ease and satisfaction, what the quadratic subfield of $\mathbf{Q}(\zeta_p)$ is. For $K = \mathbf{Q}(\sqrt{d})$, we know from Proposition 4.9 that $d_K = d$ if $d \equiv 1 \pmod 4$ and $d_K = 4d$ otherwise. If K is a subfield of $\mathbf{Q}(\zeta_p)$, then by Lemma 8.3 and Corollary 8.1 we must have $d_K \mid p^{p-2}$, so since d is squarefree, we know $d = \pm p$ and $d \equiv 1 \pmod 4$. In fact, these congruences (and the fact that $\pm p$ are distinct modulo 4) show that $d = p$ if $p \equiv 1 \pmod 4$ and $d = -p$ if $p \equiv 3 \pmod 4$.

Alternatively, we can show explicitly that $\mathbf{Q}(\sqrt{d}) \subseteq \mathbf{Q}(\zeta_p)$ by computing *Gauss sums*. The goal will be to explicitly write $\sqrt{\pm p}$ as a linear combination of powers of ζ_p. Let's define the Legendre symbol in a somewhat more algebraic way (it is clear why this definition is equivalent to the more elementary one).

Definition 8.4. The **Legendre symbol**

$$\left(\frac{\cdot}{p}\right) : \mathrm{Gal}(\mathbf{Q}(\zeta_p)/\mathbf{Q}) \to \{\pm 1\}$$

is the unique group homomorphism with kernel $H \subseteq \mathrm{Gal}(\mathbf{Q}(\zeta_p)/\mathbf{Q})$ equal to the unique subgroup of index 2.

By definition, we know $\sqrt{d}$ is fixed by H. If $\sigma \notin H$, then $\sigma(\sqrt{d}) = -\sqrt{d}$. Let

$$g = \sum_{\sigma \in G} \left(\frac{\sigma}{p}\right) \sigma(\zeta_p) \in \mathcal{O}_{\mathbf{Q}(\zeta_p)}.$$

Applying $\tau \in G$ to both sides and abusing the fact that the square of the Legendre symbol is 1,

$$\begin{aligned} \tau(g) &= \sum_{\sigma \in G} \left(\frac{\sigma}{p}\right) \tau\sigma(\zeta_p) \\ &= \left(\frac{\tau}{p}\right) \sum_{\sigma \in G} \left(\frac{\tau}{p}\right)\left(\frac{\sigma}{p}\right) \tau\sigma(\zeta_p) \\ &= \left(\frac{\tau}{p}\right) g, \end{aligned}$$

which means g^2 is fixed by $\tau(g)$ for all $\tau \in G$. By Galois theory, it follows that $g^2 \in \mathbf{Q}$. We have

$$g = \sum_{i=1}^{p-1} \left(\frac{i}{p}\right) \zeta_p^i$$

so

$$g^2 = \sum_{a,b \in \{1,\dots,p-1\}} \left(\frac{a}{p}\right)\left(\frac{b}{p}\right) \zeta_p^{a+b}$$

writing $b = ta$, we have

$$g^2 = \sum_{a,t} \left(\frac{a}{p}\right)\left(\frac{at}{p}\right) \zeta_p^{a+at} = \sum_t \left(\frac{t}{p}\right) \sum_a \zeta_p^{a(1+t)}.$$

If $t \neq -1$, then $1 + t \neq 0$ so $a(1+t)$ runs over all of $(\mathbf{Z}/p\mathbf{Z})^\times$ and the inside sum is just $\mathrm{Tr}_{\mathbf{Q}(\zeta_p)/\mathbf{Q}}(\zeta_p) = -1$. If instead $t = -1$, the inside sum is clearly $p - 1$. So

$$\begin{aligned} g^2 &= \left(\frac{-1}{p}\right)(p-1) - \sum_{t \neq -1} \left(\frac{t}{p}\right) \\ &= \left(\frac{-1}{p}\right) p - \sum_t \left(\frac{t}{p}\right). \end{aligned}$$

The sum on the right is zero since summing any nontrivial character on a group gives zero [exercise: if G is a finite group and $\chi : G \to \mathbf{C}^\times$ is a nontrivial group homomorphism, then $\sum_{g \in G} \chi(g) = 0$]. Now we have

$$g^2 = \left(\frac{-1}{p}\right) p.$$

What is sometimes called the "first supplement to the law of quadratic reciprocity" states that $\left(\frac{-1}{p}\right)$ is 1 if and only if p is 1 mod 4 [exercise: prove this fact using the fact that $(\mathbf{Z}/p\mathbf{Z})^\times$ is cyclic]. So, if $p \equiv 1 \pmod 4$, we get $\sqrt{p} \in \mathbf{Z}[\zeta_p]$, and if $p \equiv -1 \pmod 4$, we get $\sqrt{-p} \in \mathbf{Z}[\zeta_p]$. From this, we easily recover the description of the unique quadratic subfield of $\mathbf{Q}(\zeta_p)$.

§9 February 27, 2019

Today we will start talking about the generalization of unique factorization for number fields.

§9.1 Dedekind Domains

Definition 9.1. An integral domain A is called a **Dedekind domain** if it satisfies all three of the following conditions:

(1) A is Noetherian.

(2) Any nonzero prime ideal in A is maximal.

(3) A is integrally closed in its field of fractions.

Recall the definition of a *Noetherian ring*:

Definition 9.2. A ring A is called **Noetherian** if any ascending chain $I_1 \subseteq I_2 \subseteq \cdots$ is eventually stationary.

For example, $\mathbf{Z}$ is clearly Noetherian: an ascending chain of ideals is generated by a sequence of (WLOG nonnegative) integers $n_1, n_2, \ldots$ such that $n_{i+1} \mid n_i$ which must clearly be stationary since after any zeroes it is a nonincreasing sequence of positive integers. Recall, too, the definition of a *prime ideal*

Definition 9.3. An ideal $\mathfrak{p} \subseteq A$ is called a **prime ideal** if $ab \in \mathfrak{p}$ implies $a \in \mathfrak{p}$ or $b \in \mathfrak{p}$ for all $a, b \in A$.

The following are useful characterizations of prime and maximal ideals:

Lemma 9.4

Let A be a ring. An ideal $I \subseteq A$ is prime if and only if A/I is an integral domain.

Lemma 9.5

Let A be a ring. An ideal $I \subseteq A$ is maximal if and only if A/I is a field.

Not all primes are always maximal. For example, in $\mathbf{Z}[X]$, we have a prime ideal $p \cdot \mathbf{Z}[X]$ (check this by modding out and noticing that $\mathbf{F}_p[X]$ is an integral domain) which is properly contained in the proper ideal (p, X). This corresponds to the fact that $\mathbf{Z}[X]$ has Krull dimension higher than 1.

The reason why we are talking about Dedekind domains is because $\mathcal{O}_K$ is one.

Proposition 9.6

Let K be a number field. Then $\mathcal{O}_K$ is a Dedekind domain.

Proof. The hard part of this proof is the following lemma:

Lemma 9.7
If $I \subseteq \mathcal{O}_K$ is a nonzero ideal, then $|\mathcal{O}_K/I| < \infty$.

Proof. Let α be a nonzero element of I. Set $P_\alpha(X) = \det(X - \alpha|_{\mathcal{O}_K})$ where the determinant is taken using a $\mathbf{Z}$-module basis for $\mathcal{O}_K$. Then

$$0 = P_\alpha(\alpha) = \alpha^d + a_{d-1}\alpha^{d-1} + \cdots \pm N_{K/\mathbf{Q}}(\alpha).$$

Rearranging, we have that $N_{K/\mathbf{Q}}(\alpha)$ is a $\mathbf{Z}$-linear combination of powers of α, so in fact $n := N_{K/\mathbf{Q}}(\alpha) \in \mathbf{Z} \cap I$. So, we have $n \cdot \mathcal{O}_K \subseteq I$, hence a surjection

$$\mathcal{O}_K/(n \cdot \mathcal{O}_K) \to \mathcal{O}_K/I.$$

Using the fact that $\mathcal{O}_K$ is a free $\mathbf{Z}$-module of finite rank, the left hand side is finite, which forces the same to be true of $\mathcal{O}_K/I$. □

Now for the rest of the proposition. First, we show that $\mathcal{O}_K$ is Noetherian. It is isomorphic as a $\mathbf{Z}$-module to $\mathbf{Z}^m$, which we know is a Noetherian $\mathbf{Z}$-module (fact: a direct sum of Noetherian modules is Noetherian). Any increasing sequence of ideals of $\mathcal{O}_K$ is certainly an increasing sequence of $\mathbf{Z}$-submodules, so any such sequence must also be stationary. This proves the Noetherianness of $\mathcal{O}_K$ as a ring. Alternatively, we can use Lemma 9.7 to achieve surjections

$$\mathcal{O}_K/I_1 \to \mathcal{O}_K/I_2 \to \cdots$$

which must eventually be stationary because these rings are all finite.

To show that $\mathcal{O}_K$ is integrally closed in K, we just apply the transitivity of integral closures to get that the integral closure of $\mathcal{O}_K$ in K is the integral closure of $\mathbf{Z}$ in K, which is $\mathcal{O}_K$ back again.

The most interesting part of the proof is to show that all nonzero primes are maximal. Let $\mathfrak{p}$ be a nonzero prime ideal in $\mathcal{O}_K$. Lemma 9.7 tells us $\mathcal{O}_K/\mathfrak{p}$ is a finite integral domain via Lemma 9.4. Actually, any finite integral domain is a field [let C be a finite integral domain, and $\alpha \in C$ be nonzero. The map $C \to C$ given by multiplication by α is injective because α is nonzero and C is an integral domain; so actually the map is surjective and we conclude that α is invertible, hence C is a field]. So, we know $\mathcal{O}_K/\mathfrak{p}$ is a field, hence $\mathfrak{p}$ is maximal via Lemma 9.5. □

Before we see examples, we'd like to state what the general theory is. This is different from how physics is normally done, where the examples are done regardless of whether there exists a theory of which they are examples. More quotes from Professor Kisin:

- Once, Mark Kisin was interested in a paper in mathematical physics, but the whole introduction to the paper was a massive example. So, he emailed the authors asking them about exactly which vector bundle over which curve the computations were about. In the end, they had no idea.
- In physics, if whatever you are computing matches what your particle accelerator says, then you are a winner.

§9.2 Fractional Ideals and the Factorization Theorem

We are interested in factorization into primes. Let A be a Dedekind domain, K its field of fractions, and I, J additive subgroups of K. We want to be able to multiply I and J to get an additive subgroup of K. To do this, we take the subgroup generated by the products of elements of I and J, i.e.

$$I \cdot J = \{\sum_{i=1}^{n} \alpha_i \beta_i : \alpha_i \in I, \beta_i \in K\}.$$

If I, J are A-submodules of K, then $I \cdot J$ is clearly an A-submodule of K as well.

Definition 9.8. A **fractional ideal** $I \subseteq K$ is a nonzero A-submodule such that $d \cdot I \subset A$ for some nonzero $d \in A$.

The reason for this expanded definition of an ideal is that we would like to take the inverse of ideals.

Lemma 9.9

If I, J are fractional ideals, then so are $I + J$ and $I \cdot J$.

Proof. There must be $d, d' \in A$ such that $d \cdot I \subseteq A$ and $d' \cdot J \subseteq A$. It follows from the definitions that $dd' I \cdot J \subseteq A$, and $dd'(I + J) \subseteq A$. □

NB: if I is a fractional ideal of A, then $I \cdot A$ is contained in I since I is an A-module. But $1 \in A$, so it also contains I which means $I \cdot A = I$.

The result of this remark is that the fractional ideals of A form a *commutative semigroup*: The multiplication of ideals is commutative and associative (exercise), and they have an identity element given by the ideal A. It is our goal to show that actually they form a group. How can we construct the inverse of a fractional ideal? Next class we will figure this out, and we'll be able to prove the most important fact about Dedekind domains:

Theorem 9.10

Any fractional ideal I has a unique inverse I^{-1}, i.e. a unique element satisfying $II^{-1} = A$. There is also a unique factorization

$$I = \prod_{\mathfrak{p}} \mathfrak{p}^{\nu_{\mathfrak{p}}}$$

where the product runs over all nonzero primes $\mathfrak{p}$ and all but finitely many of the $\nu_{\mathfrak{p}}$ are zero. If $I' = \prod_{\mathfrak{p}} \mathfrak{p}^{\nu'_{\mathfrak{p}}}$ is some different fractional ideal, then $I \supseteq I'$ if and only if $\nu_{\mathfrak{p}} \leq \nu'_{\mathfrak{p}}$ for all $\mathfrak{p}$.

One consequence of the last part of the theorem is the following: if I is an ideal in A, then its factorization into primes has all nonnegative exponents (since A is just the product of the 0-th power of all the primes). Finally, we can do some examples of factorizations of ideals into primes.

Example 9.11

Unique factorization into irreducible elements does not always hold. Here is the standard example: let $K = \mathbf{Q}(\sqrt{-5})$ so $\mathcal{O}_K = \mathbf{Z}[\sqrt{-5}]$. Then

$$6 = 2 \cdot 3 = (1 + \sqrt{-5})(1 - \sqrt{-5}).$$

What's more, all of the elements we have factored this into are irreducible and not the same up to units. Here's how you can figure this out: if $1 + \sqrt{-5} = ab$, then using the definition of the norm as a product of Galois conjugates,

$$N_{K/\mathbf{Q}}(1 + \sqrt{-5}) = 6 = N_{K/\mathbf{Q}}(a)N_{K/\mathbf{Q}}(b).$$

It's also clear from the multiplicativity of the norm and the definition of the norm in terms of coefficients of minimal polynomials that an element of $\mathcal{O}_K$ is a unit if and only if its norm is ± 1. Also, since the norm of $x + y\sqrt{-5}$ is $x^2 + 6y^2$, it clearly cannot be 2 or 3. This means that at least one of a or b has norm ± 1, which means at least one is a unit, i.e. $1 + \sqrt{-5}$ is irreducible. The same procedure can be used to deduce that $1 - \sqrt{-5}$, 2, and 3 are irreducible. Since 2 and 3 have different norms that $1 \pm \sqrt{-5}$, we know these factorizations really are into irreducibles and really are different up to units.

Example 9.12

We can see unique factorization into primes still seems to hold in the ring $\mathcal{O}_K$ from the previous example. For example, take the prime ideals

$$\mathfrak{p} = (1 + \sqrt{-5}, 2) = (1 - \sqrt{-5}, 2)$$

and

$$\mathfrak{q}_1 = (1 + \sqrt{-5}, 3), \qquad \mathfrak{q}_2 = (1 - \sqrt{-5}, 3).$$

Then [exercise] we can check that $\mathfrak{p}\mathfrak{q}_1 = (1 + \sqrt{-5})$, $\mathfrak{p}\mathfrak{q}_2 = (1 - \sqrt{-5})$, $(3) = \mathfrak{q}_1\mathfrak{q}_2$ and $(2) = \mathfrak{p}^2$. This way, the principal ideals from last time factor further into primes and we do not arrive at a contradiction:

$$(6) = (2)(3) = (1 + \sqrt{-5})(1 - \sqrt{-5}) = \mathfrak{p}^2\mathfrak{q}_1\mathfrak{q}_2.$$

In general, we can crank the handle to figure out how to factorize in $\mathcal{O}_K$.

Example 9.13

$$\mathcal{O}_K = \mathbf{Z}[\sqrt{-5}] = \mathbf{Z}[X]/(X^2 - 5),$$

and

$$\mathcal{O}_K/2\mathcal{O}_K \cong \mathbf{F}_2[X]/(X^2 - 5) \cong \mathbf{F}_2[X]/((X + 1)^2).$$

If a prime $\mathfrak{p}$ contains 2, then we have a surjection

$$\mathcal{O}_K/2\mathcal{O}_K \to \mathcal{O}_K/\mathfrak{p}$$

so we decide to try $\mathfrak{p} = (2, X + 1) = (2, 1 - \sqrt{-5})$. The same works for (3).

Here is a preview: we can take the *ideal class group*

$$\mathrm{Cl}_A = \{\text{Fractional ideals of } A\}/\{\text{Principal fractional ideals of } A\}.$$

It will be a major goal to show that this group is finite and to develop techniques for computing it (c.f. the Minkowski bound). Moreover, this group is the obstruction to A having unique factorization into irreducible elements: if the class group is trivial, then A has unique factorization into irreducibles since it is a PID. In Example 9.12, the ideal class group has order 2 and is generated by A and $\mathfrak{p}$.

§10 March 4, 2019

§10.1 Factorization in Dedekind Domains

Last time, we had prime ideals

$$\begin{aligned}\mathfrak{p} &= (1+\sqrt{-5}, 2)\\ \mathfrak{q}_1 &= (1+\sqrt{-5}, 3)\\ \mathfrak{q}_2 &= (1-\sqrt{-5}, 3)\end{aligned}$$

in the ring of integers of $\mathbf{Q}(\sqrt{-5})$. We had factorizations into prime ideals

$$(2) = \mathfrak{p}^2, \qquad (3) = \mathfrak{q}_1\mathfrak{q}_2, \qquad (6) = \mathfrak{p}^2\mathfrak{q}_1\mathfrak{q}_2.$$

If $\mathfrak{p} = (a)$, then $\mathfrak{p}^2 = (a^2)$. So $2 = a^2 \cdot f$ for some f in the ring of integers, which is a contradiction [taking norms, $N(a)^2$ divides 4]; so the class group is not trivial (in particular $\mathfrak{p}$ is not principal).

Suppose we want to find all factorizations of 21 into irreducibles in $\mathbf{Q}(\sqrt{-17})$. To do this, we need to factorize 3 and 7.

We have

$$\mathcal{O}_K = \mathbf{Z}[\sqrt{-17}] = \mathbf{Z}[X]/(X^2+17).$$

So

$$\mathcal{O}_K/3\cdot\mathcal{O}_K = \mathbf{F}_3[X]/(X^2-1) = \mathbf{F}_3[X]/((X+1)(X-1)).$$

From this, we intuitively should expect that 3 has prime factors $(3, \sqrt{-17}\pm 1)$, and indeed we can check by hand that these are prime and that

$$(3) = (3, \sqrt{-17}+1)(3, \sqrt{-17}-1).$$

In the future we will prove a general result which will allow us to compute ideal factorizations in a somewhat broader setting (that of a number ring with a power basis).

Using the same machine,

$$\mathcal{O}_K/(7) = \mathbf{F}_7[X]/(X^2-4) = \mathbf{F}_7[X]/((X-2)(X+2))$$

so we expect

$$(7) = (7, \sqrt{-17}-2)(7, \sqrt{-17}+2).$$

We need to check that these ideals are not principal. Suppose that $\mathfrak{q}_1 = (a)$. Then $aq_2 = 3$ for some $q_2 \in \mathfrak{q}_2$. Taking the norm down to $\mathbf{Q}$, we have

$$9 = N(3) = N(a)N(q_2).$$

Since $\mathfrak{q}_2$ is a proper ideal, the norm of q_2 cannot be a unit. The same holds for (a), so the norm of both must be ± 3. In fact, both have norm 3 because this field is imaginary quadratic so (you can check that) every element has positive norm. But we can also check that no element of the ring of integers has norm 3, which means we have a contradiction. So far, we have factorized (3) and (7) as products of prime ideals which we know are not principal.

Suppose $21 = a\cdot b$. Then taking norms, we have

$$3^2\cdot 7^2 = N(a)\cdot N(b).$$

Since 3 and 7 are not the norms of any element, the only possibilities are $N(a) = 9, N(b) = 49$ or $N(a) = N(b) = 21$. In the first case, we have $9 = x^2 + 17y^2$, i.e. $x = \pm 3$ and $y = 0$. So we recover the factorization $21 = 3 \cdot 7$ (which we have to check works). In the second case, we have

$$21 = x^2 + 17y^2,$$

so $x = \pm 2$ and $y = \pm 1$ from which we recover the factorization $(2 + \sqrt{-17})(2 - \sqrt{-17})$. In terms of ideals, this factorization just gives

$$21 = (\mathfrak{q}_2\mathfrak{p}_2)(\mathfrak{q}_1\mathfrak{p}_1).$$

The fact that there are only two factorizations into irreducibles comes from the fact that the class group is $\mathbf{Z}/4\mathbf{Z}$. In particular, we can get $\gamma = [\mathfrak{p}_2] = [\mathfrak{q}_1]$ a generator for the class group, and $\gamma^3 = [\mathfrak{p}_1] = [\mathfrak{q}_2]$. This way, the only way $\mathfrak{p}_2$ can be multiplied by something to get a principal ideal is by $\mathfrak{p}_1$ or $\mathfrak{q}_2$ (hence we can only get two factorizations of (21) as a product of principal ideals since we showed all these primes are nonprincipal). This calculation also shows that the class group cannot have size 2: if it did, we could recover many more factorizations of 21 by pairing the ideals together to be principal.

§10.2 Towards the Factorization Theorem

Proposition 10.1

Let A be a Dedekind domain which is not a field. Then every maximal ideal $\mathfrak{m} \subset A$ is invertible.

Proof. The obvious candidate for an inverse is the "ideal quotient"

$$\mathfrak{m}' = \{x \in K : x \cdot \mathfrak{m} \subseteq A\}.$$

It's easy to see that $\mathfrak{m}'$ is a fractional ideal: if $0 \neq d \in \mathfrak{m}$, then $d \cdot \mathfrak{m}' \subseteq A$.

The definition immediately tells us $\mathfrak{m}\mathfrak{m}' \subseteq A$, so it remains to check the other inclusion. What's more, $\mathfrak{m}'$ contains A, so

$$\mathfrak{m} \subseteq \mathfrak{m}\mathfrak{m}' \subseteq A.$$

In particular, $\mathfrak{m}\mathfrak{m}'$ is either equal to $\mathfrak{m}$ or A. In the first case, let $x \in \mathfrak{m}'$. Then we have

$$\cdots \subset x^2 \cdot \mathfrak{m} \subset x \cdot \mathfrak{m} \subset \mathfrak{m}.$$

In particular, $x^n \in d^{-1}A$ where d clears the denominators of $\mathfrak{m}'$, so $A[x]$ is a fractional ideal in K. Since A is Noetherian, $A[x]$ is a finitely-generated A module, which means x is integral over A. Since A is integrally closed in K, actually $x \in A$. For now, this lets us conclude that $\mathfrak{m}' \subseteq A$ so actually $\mathfrak{m}' = A$.

- At Princeton, there was a speaker who was giving a talk, and he kept describing his results as: theorem 1, theorem 2, theorem 3 etc. Andrew Wiles was in the audience, and he said "I've only proved two theorems in my whole life".

Lemma 10.2

If $\mathfrak{p} \subseteq A$ is prime, and $\mathfrak{a}_1, \ldots, \mathfrak{a}_n$ are ideals in A such that $\mathfrak{a}_1 \cdots \mathfrak{a}_n \subseteq \mathfrak{p}$, then $\mathfrak{a}_i \subseteq \mathfrak{p}$ for some i.

Proof. Suppose none of them are contained in $\mathfrak{p}$. Then there exist $x_i \in \mathfrak{a}_i$ such that $x_i \notin \mathfrak{p}$. It follows that their product is not in $\mathfrak{p}$ by the fact that $\mathfrak{p}$ is prime. This contradicts the assumption that the product of the ideals is $\mathfrak{p}$. □

The crucial lemma on which ideal factorization hinges is the following:

Lemma 10.3

Every ideal of A contains a product of nonzero primes

Proof. Let Φ be the set of ideals which do not contain a product of nonzero primes. Assume for the sake of contradiction that Φ is nonempty. Then since A is Noetherian, Φ has a maximal element $\mathfrak{b}$ (one equivalent definition of a Noetherian ring is that every nonempty set of ideals has a maximal element). Suppose that $\mathfrak{b}$ is not prime. Then there are $x, y \in A$ such that $xy \in \mathfrak{b}$ and $x, y \notin \mathfrak{b}$. So $\mathfrak{b} + Ax$ and $\mathfrak{b} + Ay$ properly contain $\mathfrak{b}$. Since $\mathfrak{b}$ is maximal in Φ, it follows that $\mathfrak{b} + Ax$, $\mathfrak{b} + Ay$ contain a product of primes. In particular, their product contains the product of these two products of primes, so the ideal

$$(\mathfrak{b} + Ax)(\mathfrak{b} + Ay) = \mathfrak{b}^2 + x\mathfrak{b} + y\mathfrak{b} + xyA$$

contains a product of primes. But since $xy \in \mathfrak{b}$, this ideal is actually contained in $\mathfrak{b}$; so this contradicts the fact that $\mathfrak{b}$ does not contain a product of primes. From this we can conclude that Φ is empty, in other words every ideal contains a product of nonzero primes as desired. □

Look at a nonzero element $a \in \mathfrak{m}$. Then Lemma 10.3 tells us that aA contains a product of primes $\mathfrak{p}_1 \cdots \mathfrak{p}_n$ where n is as small as possible. By Lemma 10.2, it follows that WLOG $\mathfrak{m} \supseteq \mathfrak{p}_1$ (NB: $\mathfrak{m}$ is maximal so it is prime). All primes are maximal, so actually $\mathfrak{p}_1 = \mathfrak{m}$. Let $\mathfrak{b} = \mathfrak{p}_2 \cdots \mathfrak{p}_n$. Then aA cannot contain $\mathfrak{b}$ (this would contradict the minimality of n). In particular, there exists an element $b \in \mathfrak{b}$ such that b is not in aA. We have

$$\mathfrak{m} \cdot b \subseteq \mathfrak{m}\mathfrak{b} = \mathfrak{p}_1\mathfrak{b} \subseteq aA.$$

So $\mathfrak{m}ba^{-1} \subseteq A$, i.e. $b/a \in \mathfrak{m}' = A$. Finally this means $b \in aA$, which contradicts our choice of b. This means our original assumption that $\mathfrak{m}'\mathfrak{m} = \mathfrak{m}$ is false, hence $\mathfrak{m}'$ is the desired inverse. □

§11 March 6, 2019

Today we will finally show the main theorem on Dedekind domains. Last time we already showed one part of it, namely that the fractional ideals have inverses.

§11.1 Unique Factorization of Ideals in Dedekind Domains

Theorem 11.1

Let A be a Dedekind domain, and $\mathfrak{b}$ be a fractional ideal. Then $\mathfrak{b}$ factors uniquely as a product of (possibly negative) powers of nonzero prime ideals of A.

Proof. First, there exists a nonzero $d \in A$ such that $d\mathfrak{b} \subseteq A$. So we can write

$$\mathfrak{b} = d\mathfrak{b} \cdot (Ad)^{-1},$$

so it's enough to show existence of a factorization into primes in the case where $\mathfrak{b} \subseteq A$ (because we could then apply it to Ad and $d\mathfrak{b}$). NB: the fact that inverses are unique is a general property of groups (once we have existence, uniqueness is a consequence); this allows us to notice that the inverse of a product of primes is always the same product with the exponents negated.

Let Φ be the set of ideals which are not a product of primes. Let $\mathfrak{a}$ be a maximal element of Φ. Since $\mathfrak{a} \neq A$, it is contained in some prime $\mathfrak{p}$. By Proposition 10.1, there exists a fractional ideal $\mathfrak{p}'$ of A such that $\mathfrak{p}\mathfrak{p}' = A$. Hence,

$$\mathfrak{a} \cdot \mathfrak{p}' \subseteq \mathfrak{p}\mathfrak{p}' = A$$

Since $\mathfrak{p}'$ contains A, we actually know $\mathfrak{a}\mathfrak{p}'$ contains $\mathfrak{a}$. In fact, this containment is strict: if not, then if $x \in \mathfrak{p}'$, we know $x \cdot \mathfrak{a} \subseteq \mathfrak{a}$. So we have a descending chain

$$\mathfrak{a} \supseteq x \cdot \mathfrak{a} \supseteq x^2\mathfrak{a} \supseteq \cdots$$

Since A is Noetherian, this implies that x is integral over A, hence it is in A since A is integrally closed in K. So, $\mathfrak{p}' = A$ which is a contradiction.

In the end, we have a strict containment

$$\mathfrak{a}\mathfrak{p}' \supsetneq \mathfrak{a}.$$

Since $\mathfrak{a}$ was chosen to be maximal in Φ, we know that $\mathfrak{a}\mathfrak{p}'$ has a prime factorization, and multiplying by $\mathfrak{p}$ yields a factorization for $\mathfrak{a}$.

The uniqueness of factorizations follows from the existence of inverses. If we have two factorizations for the same ideal,

$$\prod_{\mathfrak{p}} \mathfrak{p}^{n(\mathfrak{p})} = \prod_{\mathfrak{p}} \mathfrak{p}^{n(\mathfrak{p})'}.$$

Multiplying by inverses, this means

$$\prod_{\mathfrak{p}} \mathfrak{p}^{n(\mathfrak{p})-n(\mathfrak{p})'} = A.$$

Bringing all the negative exponents to the right, it suffices to show that if

$$\mathfrak{p}_1^{\alpha_1} \cdots \mathfrak{p}_r^{\alpha_r} = \mathfrak{p}_1^{\beta_1} \cdots \mathfrak{p}_n^{\beta_n}$$

where all the exponents are positive, then all the exponents must be equal. Notice that $\mathfrak{p}_1$ contains the product of the powers of the $\mathfrak{q}_i$'s, so (by Lemma 10.3) it is equal to one of the $\mathfrak{q}_i$'s. We can proceed inductively to get that the factorizations are the same. □

Corollary 11.2 ("to contain is to divide")

If I and J are fractional ideals, then $I \supseteq J$ if and only if the exponents in the factorization of I are bounded by the exponents in the factorization of J.

Proof. $I \supseteq J$ is equivalent to

$$A \supseteq I^{-1}J = \prod_{\mathfrak{p}} \mathfrak{p}^{n(\mathfrak{p})'-n(\mathfrak{p})}$$

which we showed previously is true if and only if the exponents here are nonnegative. □

§11.2 Computing Prime Decompositions in Extensions of Number Fields

Now we give a very useful lemma for computing how a prime splits in extensions of $\mathbf{Q}$.

Lemma 11.3

Let K be a number field and p a rational prime, such that $\mathcal{O}_K$ is of the form $\mathbf{Z}[\alpha] \cong \mathbf{Z}[X]/(f_{0,\alpha})$ for some $\alpha \in \mathcal{O}_K$. Then

$$\mathcal{O}_K/(p \cdot \mathcal{O}_K) \cong \mathbf{F}_p[X]/(\bar{f}),$$

and since $\mathbf{F}_p$ is a UFD, we have a factorization into irreducible polynomials

$$\bar{f} = \bar{f}_1^{e_1} \cdots \bar{f}_n^{e_n}.$$

Then

$$p \cdot \mathcal{O}_K = \prod_i \mathfrak{p}_i^{e_i}$$

where $\mathfrak{p}_i = (p, f_i(\alpha))$ where f_i is an arbitrary lift of $\bar{f}_i$.

Proof. First, $(p, f_i(\alpha))$ is prime because

$$\mathcal{O}_K/\mathfrak{p}_i = \mathbf{F}_p[X]/(\bar{f}, \bar{f}_i) = \mathbf{F}_p[X]/(\bar{f}_i)$$

which is a field because $\bar{f}_i$ is irreducible. Setting $I = \prod_i \mathfrak{p}_i^{e_i}$, we know $I \subseteq (p)$ because under the reduction of $\mathcal{O}_K$ modulo (p), I gets sent to

$$\prod_i \bar{f}_i(\alpha)^{e_i} = 0.$$

Writing

$$p \cdot \mathcal{O}_K = \prod_i \mathfrak{p}_i^{e_i'},$$

we know

$$p \cdot \mathcal{O}_K \mapsto \prod \bar{f}_i^{e_i'}(\alpha)$$

which must be zero in the quotient by (p). So we must have a divisibility of polynomials which yields the opposite inequality $e_i' \geq e_i$.

□

In general, the direct use of this computational technique is limited: When K is not quadratic, its ring of integers does not necessarily have a power basis.

§12 March 11, 2019

§12.1 Splitting of Primes in Cyclotomic Fields

The lemma from last time applies in one other important case, namely the cyclotomic field $\mathbf{Q}(\zeta_p)$ with ring of integers $\mathbf{Z}[\zeta_p]$.

Example 12.1

We have

$$\mathcal{O}_K = \mathbf{Z}[X]/f = \mathbf{Z}[Y]/g(Y)$$

where $f(X) = (X^q - 1)/(X - 1)$ and $g(Y) = f(Y + 1)$. We can compute

$$\mathcal{O}_K/(p) = \mathbf{F}_p[Y]/Y^{p-1}$$

which means by the lemma

$$(p) = (p, \zeta_p - 1)^{p-1}$$

is the factorization of p into primes in $\mathcal{O}_K$. We may check directly that this factorization is the same as writing

$$(p) = (\zeta_p - 1)^{p-1}$$

[one way to see this is to show directly that p is in $(\zeta_p - 1)$].

All the material above will be on the midterm.

§12.2 The Ideal Norm

Now, we are starting a new topic (the finiteness of the class group and the geometry of numbers), which will not be tested on the midterm. The geometry of numbers will be our main tool, and in general when we have objects analogous to the class group (like the Tate-Shafarevich group of an elliptic curve) the fact that we do not have the geometry of numbers makes it much harder to prove the finiteness of the group.

Definition 12.2. Let $\mathfrak{a}$ be an ideal in $\mathcal{O}_K$. The **norm** of $\mathfrak{a}$ is

$$N(\mathfrak{a}) := |\mathcal{O}_K/\mathfrak{a}|.$$

Proposition 12.3

Let FI_K be the group of fractional ideals of K. There exists a homomorphism

$$N : FI_K \to \mathbf{Q}^\times$$

with the following properties:

(i) If $\mathfrak{a} \subseteq \mathcal{O}_K$ is an ideal, then $N(\mathfrak{a})$ is just the norm of $\mathfrak{a}$ as an ideal.

(ii) If x is a nonzero element of $\mathcal{O}_K$, then $N(x \cdot \mathcal{O}_K) = |N_{K/\mathbf{Q}}(x)|$.

Proof. By unique factorization, FI_K is a free abelian group on the nonzero primes of $\mathcal{O}_K$. So, there is a unique group homomorphism $N : FI_K \to \mathbf{Q}^\times$ taking $\mathfrak{p}$ to $|\mathcal{O}_K/\mathfrak{p}|$. To

prove the first desired property, it suffices to show that

$$|\mathcal{O}_K/\mathfrak{a}\mathfrak{p}| = |\mathcal{O}_K/\mathfrak{a}||\mathcal{O}_K/\mathfrak{p}|.$$

By one of the isomomorphism theorems, it suffices to show that

$$|\mathcal{O}_K/\mathfrak{p}| = |\mathfrak{a}/\mathfrak{a}\mathfrak{p}|$$

where the quotients are quotients of $\mathcal{O}_K$-modules. In fact, $\mathfrak{a}/\mathfrak{a}\mathfrak{p}$ is killed by $\mathfrak{p}$ so it is an $\mathcal{O}_K/\mathfrak{p}$-module. So it suffices to show it is one-dimensional as a $\mathcal{O}_K/\mathfrak{p}$-vector space. In particular, we want to show that the only subspaces of $\mathfrak{a}/\mathfrak{a}\mathfrak{p}$ are zero and itself. Since the $\mathcal{O}_K$-submodules of $\mathfrak{a}$ containing $\mathfrak{p}$ and those of $\mathfrak{a}/\mathfrak{p}$ are in bijection via the canonical projection, it suffices to show that the only ideals of $\mathcal{O}_K$ containing $\mathfrak{a}\mathfrak{p}$ and contained in $\mathfrak{a}$ are exactly those two ideals. This follows directly from unique factorization (recall the slogan: to contain is to divide).

Now we are left with the second desired property (the norm of a principal ideal is the absolute value of the norm of the generator). Let x be a nonzero element of $\mathcal{O}_K$. Then the norm of x is equal to the determinant of the $\mathbf{Q}$-linear operator $\varphi_x : K \to K$ given by multiplication by x.

Lemma 12.4

If $\varphi : \mathbf{Z}^n \to \mathbf{Z}^n$ is a map of abelian groups with nonzero determinant, then

$$|\det \varphi| = |\mathbf{Z}^n/\varphi(\mathbf{Z}^n)|.$$

Proof. We already proved this (see the matrix proof and the big brain donut proof of Lemma 5.5). □

From the lemma, we get

$$|N_{K/\mathbf{Q}}(x)| = |\mathcal{O}_K/x\mathcal{O}_K|$$

as desired. □

§12.3 The Geometry of Numbers and Finiteness of the Class Group

We want to prove that the class number of a number field is finite. The way we'll do this is via the Minkowski bound; the idea is the following:

Proposition 12.5

Let K be a number field. There exists a constant $c > 0$ depending only on K such that for any fractional ideal $\mathfrak{b} \subseteq \mathcal{O}_K$ there exists an $x \in \mathfrak{b}$ with $|N_{K/\mathbf{Q}}(x)| \leq cN(\mathfrak{b})$.

Corollary 12.6

The class group of K is finite.

Proof. The point is that any ideal class in Cl_K has a representative which is an ideal. In particular, any fractional ideal I is in the same ideal class as dI where d is chosen to

clear the denominators of I. If $\mathfrak{a} \subseteq \mathcal{O}_K$ is an ideal, then the proposition says that there is an $x \in \mathfrak{a}^{-1}$ such that $|N_{K/\mathbf{Q}}(x)| \leq cN(\mathfrak{a})^{-1}$. This is the same as saying

$$N(x \cdot \mathfrak{a}) \leq c.$$

In particular, the ideal class of $\mathfrak{a}$ is represented by an ideal $\subseteq \mathcal{O}_L$ (namely $x \cdot \mathfrak{a}$) with norm at most c. So, it suffices to show that there are only finitely many ideals of a given norm.

If $N(\mathfrak{b}) = n \in \mathbf{Z}$, then since a finite abelian group is killed by its order, we know $n \cdot \mathcal{O}_K \subseteq \mathfrak{b}$. If we factorize $n \cdot \mathcal{O}_K$, unique factorization into prime ideals tells us there are only finitely many choices for $\mathfrak{b}$. Alternatively, the set of such $\mathfrak{b}$ is in bijection with the set of ideals of $\mathcal{O}_K/n \cdot \mathcal{O}_K$, which is finite. □

In practice, computing the class group will follow the same lines as the proof: we'll have an explicit description of c, and we'll factor the ideal (n) into primes for all positive integers $n \leq c$. The idea of the geometry of numbers is that we have inclusions

$$\mathcal{O}_K \subseteq K \to K \otimes_{\mathbf{Q}} \mathbf{R} \cong \prod \mathbf{R} \times \prod \mathbf{C}$$

where the products are over all the real and complex embeddings of K. The reason for the isomomorphism is that by the primitive element theorem, $K = \mathbf{Q}[X]/(f)$ for some monic irreducible polynomial $f \in \mathbf{Q}[X]$, and by factoring f into irreducibles over $\mathbf{R}$ we get an isomorphism with $\prod_i \mathbf{R}[X]/f_i$, which is the same as the desired product (we'll go over this in more detail later). Moreover, $\mathcal{O}_K$ sits inside K as a free $\mathbf{Z}$-module, and when we embed it in $K \otimes_{\mathbf{Q}} \mathbf{R}$ we'll see that any ideal of $\mathcal{O}_K$ embeds in this tensor product as a lattice. It'll be the goal to bound the sizes of points in a fundamental domain of this lattice.

§13 March 25, 2019

Last time we defined the norm $N : FI_K \to \mathbf{Q}^\times$ a group homomorphism from the group of fractional ideals of $\mathcal{O}_K$ to the multiplicative group of $\mathbf{Q}$, and we noticed that for any $x \in K$ we have $|N_{K/\mathbf{Q}}(x)| = N((x))$.

§13.1 Embedding $\mathcal{O}_K$ as a Discrete Subgroup of $\mathbf{R}^n$

Let K be a number field. We have an embedding

$$\sigma : \mathcal{O}_K \to K \otimes_{\mathbf{Q}} \mathbf{R} = \mathcal{O}_K \otimes_{\mathbf{Z}} \mathbf{R}.$$

By the primitive element theorem, we can write

$$K \cong \mathbf{Q}[X]/f(X)$$

for some irreducible $f \in \mathbf{Q}[X]$. In $\mathbf{R}[X]$, f factors into irreducibles as

$$f(X) = f_1(X)\cdots f_{r_1}(X) f_{r_1+1}(X)\cdots f_{r_1+r_2}(X)$$

where $f_1, \ldots, f_{r_1}$ are of degree 1, and the rest are of degree 2. So we have an isomorphism

$$K \otimes_{\mathbf{Q}} \mathbf{R} \cong \prod \mathbf{R}[X]/f_i(X) \cong \mathbf{R}^{r_1} \times \mathbf{C}^{r_2}$$

The embedding $\sigma : \mathcal{O}_K \to \mathbf{R}^{r_1} \times \mathbf{C}^{r_2}$ is given by all the real embeddings on the real coordinates, and one choice of complex embedding from each conjugate pair on the complex coordinates (this is fine up to the automorphism of $\mathbf{C}$ given by complex conjugation). Are you happy with this or not? This is like Brexit, you have to choose one or the other.

Example 13.1

Let $K = \mathbf{Q}(\sqrt{2})$. Then $r_1 = 2$ and $r_2 = 0$. The two real embeddings are given by $\sqrt{2} \mapsto \pm\sqrt{2}$. So, the image of $\mathcal{O}_K$ in $\mathbf{R}^2$ is just the lattice spanned by $(1,1)$ and $(\sqrt{2}, -\sqrt{2})$.

Example 13.2

Let $K = \mathbf{Q}(\sqrt{-1})$. Then the image of $\mathcal{O}_K = \mathbf{Z}[i]$ in $\mathbf{C}$ is given by the usual complex embedding - it's the lattice spanned by 1 and i.

Lemma 13.3

$\sigma(\mathcal{O}_K) \subseteq \mathbf{R}^{r_1} \times \mathbf{C}^{r_2}$ is discrete.

Proof. If $x \in \mathcal{O}_K$ is nonzero, then $N_{K/\mathbf{Q}}(x) \in \mathbf{Z}$ is nonzero, and it is equal to the product of $\tau(x)$ as τ runs over the embeddings $K \to \mathbf{C}$ (this includes the "real" embeddings). It follows that

$$\left|\prod_\tau \tau(x)\right| \geq 1,$$

so there exists a τ such that $|\tau(x)| \geq 1$. This means that $B(0, 1/2) \cap \sigma(\mathcal{O}_K) = \{0\}$, where $B(0, 1/2)$ is the ball of radius $1/2$ around zero under the sup norm. Since $\sigma(\mathcal{O}_K)$ is a subgroup of $\mathbf{R}^{r_1} \times \mathbf{C}^{r_2}$, it was actually enough to show that zero was isolated (if any other point was not isolated we could subtract that point from a sequence converging to it to get a sequence converging to zero, for example). □

There's a reason why we only take one complex embedding per pair. It's because we want to take just enough embeddings so that $\sigma(\mathcal{O}_K)$ is not only discrete, but has the same rank as the dimension of the space it sits inside of. This way, $\sigma(\mathcal{O}_K)$ is a **lattice** in that vector space. Actually, the rank is the only obstruction to discrete subgroup of $\mathbf{R}^n$ being a lattice in the other sense (being generated over $\mathbf{Z}$ by a basis of $\mathbf{R}^n$).

Example 13.4

$\mathbf{Q} + \mathbf{Q} \cdot \sqrt{2}, \mathbf{Z} + \mathbf{Z} \cdot \sqrt{2} \subseteq \mathbf{R}$ are not discrete (exercise). Even though it is isomorphic to $\mathbf{Z}^2$, the second one fails to be discrete because its rank is too big.

Proposition 13.5

Let $H \subseteq \mathbf{R}^n$ be a discrete subgroup. Then H is generated as an abelian group by r linearly independent vectors for some $r \leq n$. So $H \otimes_{\mathbf{Z}} \mathbf{R} \cong \mathbf{R}^r \to \mathbf{R}^n$.

Proof. We need a lemma from topology.

Lemma 13.6

A subgroup $H \subseteq \mathbf{R}^n$ is discrete if and only if $H \cap K$ is finite for all compact $K \subseteq \mathbf{R}^n$.

Proof. If H is discrete, choose an open set U containing 0 such that $H \cap U = \{0\}$. Any Cauchy sequence $\{h_i\} \subseteq H$ must have $h_i - h_j$ eventually always in U (it's immediately in H because H is a group), which means that $h_i - h_j = 0$ for sufficiently large i, j i.e. $\{h_i\}$ is eventually constant and hence converges in H. So, H is closed. It follows that $H \cap K$ is compact and discrete. The definition of compactness tells us that $H \cap K$ is finite.

In the other direction, suppose $H \cap K$ is finite for any choice of K. Then $H \cap B(0, r)$ is finite for any r. Since finite sets of points are always isolated, for small enough $r' < r$ we know $H \cap B(0, r') = \{0\}$ so H is discrete. □

Let $\underline{e} = (e_1, \ldots, e_r)$ be a maximal set of elements of H which are $\mathbf{R}$-linearly independent. Consider the parallelepiped

$$P = \{x \in \mathbf{R}^n : x = \alpha_1 e_1 + \cdots \alpha_r e_r, \alpha_i \in [0, 1]\}.$$

P is compact, so by the lemma $P \cap H$ is finite. We should expect that the elements of $P \cap H$ generate all of H, since we expect P to be a bunch of copies of a fundamental parallelogram.

If $x \in H$, then by the maximality of $\underline{e}$ there are $\{\lambda_i\} \in \mathbf{R}$ such that $x = \sum_i \lambda_i e_i$. For $j \geq 1$, let

$$x_j := jx - \sum [j\lambda_i] e_i = \sum \{j\lambda_i\} e_i \in P,$$

i.e. a choice of representative of jx inside P. [here $[a]$ denotes the greatest integer at most a and $\{a\}$ denotes the fractional part of a]. Note that $x_j \in P \cap H$ since H is a group. Taking $j = 1$, we get

$$x = x_1 + \sum [\lambda_i] e_i$$

so every $x \in H$ is a finite $\mathbf{Z}$-linear combination of elements of $P \cap H$. Since $P \cap H$ is finite this means H is finitely generated.

The fact that $P \cap H$ also tells us that there is some $j < k$ such that $x_k = x_j$, and thus

$$(k - j)\lambda_i = [k\lambda_i] - [j\lambda_i]$$

for all i which means the λ_i's are irrational.

So, $H \subseteq \sum \mathbf{Q} \cdot e_i$, hence $H \otimes_{\mathbf{Z}} \mathbf{Q} = \sum \mathbf{Q} \cdot e_i$. Since H is finitely generated, there is a positive integer d such that

$$d \cdot H \subseteq \sum \mathbf{Z} \cdot e_i \subseteq H \subseteq \sum \mathbf{Q} \cdot e_i$$

which means the rank of H is equal to $r \leq n$. □

This is a particular example of the theory of lattices in $\mathbf{R}^n$

§13.2 Lattices

Definition 13.7. A **lattice** $H \subseteq \mathbf{R}^n$ is a discrete subgroup of rank n.

What we just proved is essentially the following.

Corollary 13.8

Let $H \subseteq \mathbf{R}^n$ be a discrete subgroup. Then the following are equivalent:

(i) H is a lattice.

(ii) H spans $\mathbf{R}^n$ over $\mathbf{R}$.

(iii) $H \otimes \mathbf{R} \cong \mathbf{R}^n$.

(iv) $\mathbf{R}^n/H$ has finite volume.

Proof. (i) is equivalent to (ii) and (iii) by the proposition. What do we mean by the volume of $\mathbf{R}^n/H$. Intuitively it's the volume of the fundamental parallelepiped of H, but we don't know yet that it is well-defined. In fact, the volume of $\mathbf{R}^n/H$ is well-defined regardless.

Around the origin $0 \in \mathbf{R}^n$ we can take a small enough open set U such that U maps bijectively onto its image via the projection to $\mathbf{R}^n/H$. We do this by taking an open set V containing 0 such that $V \cap H = \{0\}$, and then taking $U = \frac{1}{2}V$. This allows us to write down the volume of $\mathbf{R}^n/H$ by taking $\mathrm{vol}(\pi(U)) = \mathrm{vol}(U)$ for an open set U mapping bijectively onto its image.

Less canonically, we can do it via the volume of a fundamental domain. Let $e_1, \ldots, e_r$ be a $\mathbf{Z}$-basis for H. Choose $e_{r+1}, \ldots, e_n \in \mathbf{R}^n$ such that $e_1, \ldots, e_n$ is a basis for $\mathbf{R}^n$. Take P to be the fundamental domain

$$P := \{\sum \alpha_i e_i : 0 \leq \alpha_i < 1 \text{ for } i = 1, \ldots, r\}.$$

If $r < n$ then P has infinite volume and otherwise it clearly has finite volume since it is bounded. It clearly maps bijectively onto $\mathbf{R}^n/H$ so in fact $\mathbf{R}^n/H$ has volume equal to the volume of P, which is finite iff H is a lattice. □

§14 March 27, 2019

Joe Harris told a story: Ahlfors taught a class and Joe was often the only person there. Once, Joe showed up 20 minutes late and Ahlfors was there giving the lecture with zero students.

NB: when the weather starts to get warmer, students always get to class later. But it's not warm yet!!

§14.1 The Geometry of Numbers

Suppose H is a lattice, and let P_e be the fundamental domain using a partical basis e of H. Since the volume of $\mathbf{R}^n/H$ is well-defined, we've already shown that $\mathrm{vol}(P_e)$ does not depend on e. This is morally the right way to prove it, but there's a less canonical way to do it if you want to define the volume as the volume of a fundamental domain instead.

Lemma 14.1

$\mu(P_e)$ depends only on H and not on e.

Proof. If $e_1, \dots, e_n$ and $f_1, \dots, f_n$ are different $\mathbf{Z}$-bases for H, then the matrix M taking $\underline{e}$ to $\underline{f}$ has determinant ± 1 since it is invertible over $\mathbf{Z}$. By standard properties of the determinant,

$$\mu(P_f) = |\det(M)|\mu(P_e) = \mu(P_e)$$

as desired.

We'd like to prove this "standard" property of the determinant though. Actually, both the map $M \mapsto \det M$ and the map $M \mapsto \mu(P_f)/\mu(P_e)$ satisfy the same universal property: as maps $(\mathbf{R}^n)^{\otimes n} \to \mathbf{R}$ they factor through $\wedge^n \mathbf{R}^n$. So, they differ by a constant factor. Checking they have the same value on, say, the identity matrix we can be done. □

Theorem 14.2 (Minkowski)

Let $H \subseteq \mathbf{R}^n$ be a lattice, and $S \subseteq \mathbf{R}^n$ be a measurable set such that $\mu(S) > \mathrm{vol}(H)$. Then there exist $x, y \in S$ such that $x \neq y$ and $x - y \in H$.

Proof. This is the pidgeonhole principle. If $\pi : S \to \mathbf{R}^n/H$ is injective, then S maps bijectively to a subset of $\mathbf{R}^n/H$ which therefore has volume equal to $\mu(S) > \mathrm{vol}(H)$. This ic a contradiction. □

Different proof. Let $e_1, \dots, e_n$ be a $\mathbf{Z}$-basis for H. Then

$$S = \bigsqcup_{h \in H} S \cap (h + P_e).$$

By the properties of measure,

$$\begin{aligned} \mu(S) &= \sum_{h \in H} \mu(S \cap (h + P_e)) \\ &= \sum_{h \in H} \mu((-h + S) \cap P_e). \end{aligned}$$

If the sets $(-h+S)\cap P_e$ are all disjoint, then the right hand side is at most $\mu(P_e)$ which is a contradiction. So there are distinct $h,h'\in H$ such that $-h+S$ and $-h'+S$ have a common element, which means there are distinct $x,y\in S$ such that

$$x-y=h-h'\neq 0$$

and is in H since H is a group. This was the desired conclusion. □

The result of this is a very useful result, which is what people usually mean when they say "Minkowski's Theorem"

Corollary 14.3

Suppose $H\subseteq \mathbf{R}^n$ is a lattice and $S\subseteq \mathbf{R}^n$ is a measurable, symmetric and convex. If

(a) $\mu(S)>2^n\mathrm{vol}(H)$ or

(b) $\mu(S)\geq 2^n\mathrm{vol}(H)$ and S is compact,

then S contains a nonzero element of H.

Proof. First we do the case where $\mu(S)>2^n\mathrm{vol}(H)$. Let $S'=\frac{1}{2}S$. Then $\mu(S')=2^{-n}\mu(S)>\mathrm{vol}(H)$. The theorem tells us that there are $x,y\in S'$ which differ by an element of H. In particular, we have a nonzero element of H given by

$$x-y=\frac{2x-2y}{2}=\frac{2x+(-2y)}{2}\in S$$

as desired. This proves part (a).

Now consider the case where S is compact. For any $\epsilon>0$ we can take the slightly larger compact, measurable, symmetric, convex set $(1+\epsilon)S$ and use the fact that $\mu((1+\epsilon)S)>\mu(S)\geq 2^n\mathrm{vol}(H)$ to apply part (a) to $(1+\epsilon)S$. In particular,

$$(1+\epsilon)S\cap(H\smallsetminus\{0\})\neq\varnothing.$$

Each one of these intersections is finite by the compactness of $(1+\epsilon)S$. We may write $S\cap(H\smallsetminus\{0\})$ as the nested intersection

$$\bigcap_{\epsilon>0}(1+\epsilon)S\cap(H\smallsetminus\{0\}).$$

A nested intersection of finite (or indeed compact) nonempty sets must be nonempty, so we are done. □

§15 April 1, 2019

We can restrict the embedding $\mathcal{O}_K \to \mathbf{R}^{r_1} \times \mathbf{C}^{r_2}$ as a lattice to an embedding of an arbitrary ideal $\mathfrak{a} \subseteq \mathcal{O}_K$. Actually, any such ideal embeds as a sublattice of $\sigma(\mathcal{O}_K)$ since it is a submodule of finite index, and we can compute its volume.

§15.1 Volumes of Sublattices

Proposition 15.1

Let $M \subseteq K$ be a submodule of rank n. Then $\sigma(M) \subseteq \mathbf{R}^n$ has

$$\operatorname{vol}(\sigma(M)) = 2^{-r_2} |\det(\sigma_i(x_j))|$$

where the x_j's are a $\mathbf{Z}$-basis for M.

Proof. Since $\mathcal{O}_K$ spans K, there exists a positive integer d such that $d \cdot M \subseteq \mathcal{O}_K$ (take the lcm of the required d for each basis element of M). This means that $\sigma(M) \subseteq d^{-1}\sigma(\mathcal{O}_K)$. Since $\sigma(\mathcal{O}_K)$ is discrete, so is $\sigma(M)$. It has the right rank (it is a finite index subgroup of $\sigma(\mathcal{O}_K)$), so $\sigma(M)$ is a lattice in $\mathbf{R}^n \cong \mathbf{R}^{r_1} \times \mathbf{C}^{r_2} \cong \mathbf{R}^{r_1} \times \mathbf{R}^{r_2} \times (i\mathbf{R})^{r_2}$. The $2r_2$ complex embeddings come in complex conjugate pairs, and we denote them by $\sigma_{r_1+1}, \dots, \sigma_{r_1+r_2}$ and their conjugates. The basis elements map to

$$\sigma(x_i) = \Big(\sigma_1(x_i), \dots, \sigma_{r_1}(x_i), \frac{1}{2}(\sigma_{r_1+1} + \overline{\sigma}_{r_1+1})(x_i), \dots, \frac{1}{2}(\sigma_{r_1+r_2} + \overline{\sigma}_{r_1+r_2})(x_i), \\ \frac{1}{2i}(\sigma_{r_1+1} - \overline{\sigma}_{r_1+1})(x_i), \dots, \frac{1}{2i}(\sigma_{r_1+r_2} - \overline{\sigma}_{r_1+r_2})(x_i)\Big).$$

The volume of $\sigma(M)$ is the absolute value of the determinant of the matrix whose rows are the $\sigma(x_i)$'s. Taking out the factors of $1/(2i)$ in the last r_2 columns, this is equal to 2^{-r_2} times the determinant of the matrix whose rows are

$$\Big(\sigma_1(x_i), \dots, \sigma_{r_1}(x_i), \frac{1}{2}(\sigma_{r_1+1} + \overline{\sigma}_{r_1+1})(x_i), \dots, \frac{1}{2}(\sigma_{r_1+r_2} + \overline{\sigma}_{r_1+r_2})(x_i), \\ (\sigma_{r_1+1} - \overline{\sigma}_{r_1+1})(x_i), \dots, (\sigma_{r_1+r_2} - \overline{\sigma}_{r_1+r_2})(x_i)\Big).$$

subtracting $1/2$ times the last r_2 columns from the r_2 columns before them, we obtain the matrix whose rows are

$$\Big(\sigma_1(x_i), \dots, \sigma_{r_1}(x_i), \overline{\sigma}_{r_1+1}(x_i), \dots, \overline{\sigma}_{r_1+r_2}(x_i), \\ (\sigma_{r_1+1} - \overline{\sigma}_{r_1+1})(x_i), \dots, (\sigma_{r_1+r_2} - \overline{\sigma}_{r_1+r_2})(x_i)\Big).$$

Finally, adding the middle r_2 columns back to the last r_2 we get the rows $\sigma_j(x_i)$, as desired. □

Corollary 15.2

If $\mathfrak{a} \subseteq \mathcal{O}_K$ is an ideal, then $\sigma(\mathfrak{a})$ is a lattice and

$$\operatorname{vol}(\sigma(\mathfrak{a})) = 2^{-r_2} |d_k|^{1/2} N(\mathfrak{a}).$$

Proof. In the case that $\mathfrak{a} = \mathcal{O}_K$, we know from the proposition that the desired volume is $2^{-r_2}|d_k|^{1/2}$ via the definition of the discriminant (c.f. Lemma 7.3).

In general, $\sigma(\mathfrak{a})$ is a sublattice of $\sigma(\mathcal{O}_K)$ by the proposition, and the projection map

$$\mathbf{R}^n/\sigma(\mathfrak{a}) \to \mathbf{R}^n/\sigma(\mathcal{O}_K)$$

is a $N(\mathfrak{a})$-to-one covering map of manifolds, which means $\mathrm{vol}(\mathfrak{a}) = N(\mathfrak{a})\mathrm{vol}(\mathcal{O}_K)$ as desired. □

Pedestrian proof. You can use the structure theorem for modules over a PID to show that if $\Lambda \subseteq \Lambda'$ are lattices, then

$$\mathrm{vol}(\Lambda)/\mathrm{vol}(\Lambda') = [\Lambda : \Lambda'].$$

See the proofs of Lemma 5.5. □

§15.2 The Minkowski Bound and Finiteness of the Class Group

We may abuse these calculations and the machinery of the geometry of numbers to prove the finiteness of the class group.

Proposition 15.3

If $\mathfrak{a} \subseteq \mathcal{O}_K$ is an ideal, then there exists an $x \in \mathfrak{a}$ such that

$$|N_{K/\mathbf{Q}}(x)| \le \left(\frac{4}{\pi}\right)^{r_2} \frac{n!}{n^n}|d_k|^{1/2}N(\mathfrak{a}).$$

Proof. Let $t \in \mathbf{R}^+$, and consider the region

$$B_t = \{(y_1, \dots, y_{r_1}, z_1, \dots, z_{r_2}) \in \mathbf{R}^{r_1} \times \mathbf{C}^{r_2} : \sum |y_i| + 2\sum |z_i| \le t\},$$

which is easily seen to satisfy the hypotheses of Minkowski's theorem. The volume of B_t is

$$\mu(B_t) = 2^{r_1}\left(\frac{\pi}{4}\right)^{r_2}\frac{t^n}{n!}.$$

Choose t so that

$$\mu(B_t) = 2^n\mathrm{vol}(\sigma(\mathfrak{a})) = 2^{n-r_2}|d_k|^{1/2}N(\mathfrak{a}),$$

i.e.

$$t^n = 2^{n-r_1}\pi^{-r_2}n!|d_k|^{1/2}N(\mathfrak{a}).$$

So, there exists a nonzero $x \in \mathfrak{a}$ such that $x \in B_t$. Since

$$|N_{K/\mathbf{Q}}(x)| = \prod |\sigma_i(x)| = \prod_{i=1}^{r_1} |\sigma_i(x)| \cdot \prod_{i=r_1+1}^{r_1+r_2} |\sigma_i(x)|^2,$$

by AM-GM we know

$$|N_{K/\mathbf{Q}}(x)| \le \frac{1}{n}\left(\sum_{i=1}^{r_1}|\sigma_i(x)| + 2\sum_{i=r_1+1}^{r_1+r_2}|\sigma_i(x)|\right)$$

so actually the norm satisfies the inequality we wanted it to. □

Corollary 15.4

Every ideal class in Cl_K contains an ideal $\mathfrak{b} \subseteq \mathcal{O}_K$ such that $N(\mathfrak{b}) \leq \left(\frac{4}{\pi}\right)^{r_2} \frac{n!}{n^n} |d_k|^{1/2}$.

Proof. Let $\mathfrak{c} \in \mathrm{Cl}_K$. Take $\mathfrak{a} \subseteq \mathcal{O}_K$ be an ideal with $[\mathfrak{a}] = -\mathfrak{c}$. Take $x \in \mathfrak{a}$ such that $|N_{K/\mathbf{Q}}(x)|$ satisfies the bound from the proposition. The ideal

$$\mathfrak{b} = x \cdot \mathfrak{a}^{-1} \subseteq \mathcal{O}_K$$

then has norm equal to $|N_{K/\mathbf{Q}}(x)| \cdot N(\mathfrak{a})^{-1}$ which is indeed at most the Minkowski bound. □

Corollary 15.5

The class group of any number field is finite.

Proof. The Minkowski bound tells us that only finitely many positive integers can be the norm of an ideal in $\mathcal{O}_K$, and since an ideal contains its norm (if $n = N(\mathfrak{a})$ then $n \cdot \mathcal{O}_K \subset \mathfrak{a}$ since $n := |\mathcal{O}_K/\mathfrak{a}|$), there are finitely many ideals of a given norm. Hence, the class group is finite. □

§16 April 3, 2019

Today we will explain how to use the Minkowski bound to compute ideal class groups.

§16.1 Minkowski Bound Computations

Example 16.1

Let K be a quadratic field. The Minkowski bound tells us that any ideal class has a representative of norm at most $(2/\pi)|d_k|^{1/2}$ if $r_2 = 1$ and otherwise if $r_1 = 1$ then any ideal class has a representative of norm at most $(1/2)|d_K|^{1/2}$.

Example 16.2

Let $K = \mathbf{Q}(i)$. Then the Minkowski bound says that every ideal class has a representative of norm at most $4/\pi < 2$. Since there is only one ideal of norm 1, namely $\mathcal{O}_K$ it follows that the class group is trivial, hence $\mathbf{Z}[i]$ is a PID.

Example 16.3

$K = \mathbf{Q}(\sqrt{-5})$. The Minkowski bound says that every ideal class has a representative of norm at most

$$\frac{2}{\pi}\sqrt{4 \cdot 5} < 3.$$

The only ideal of norm 1 is $\mathcal{O}_K$. By the multiplicativity of the norm, it suffices to find the prime ideals of norm 2. If $N(\mathfrak{p})$ has norm 2, then $\mathfrak{p}$ divides the ideal (2). We can use the usual method to compute the prime factorization of (2):

$$(2) = (2, \sqrt{-5} - 1)^2.$$

Hence, the only prime of norm 2 is $(2, \sqrt{-5} - 1)$ which we checked is nonprincipal. So, this ideal is in a different class from the identity, hence the class group is $\mathbf{Z}/2\mathbf{Z}$.

Example 16.4

Let $K = \mathbf{Q}(\sqrt{-11})$. The Minkowski bound is

$$\frac{2}{\pi}|11|^{1/2} < 3$$

so we just need to factorize the ideal (2). It turns out to remain prime in $\mathcal{O}_K$ as a result of the fact that $X^2 - X + 1$ is irreducible over $\mathbf{F}_2$. So there are no such ideals of norm 2, hence the class group is trivial.

Example 16.5

Let $K = \mathbf{Q}(\sqrt{-17})$. The Minkowski bound tells us there is a representative of each ideal class of norm at most 5. So it suffices to find the prime ideals dividing (2), (3), and (5). The usual procedure gives

$$(2) = (2, \sqrt{-17} - 1)^2 = \mathfrak{p}$$

$$(3) = (3, \sqrt{-17} - 1)(3, \sqrt{-17} + 1) = \mathfrak{q}_1\mathfrak{q}_2$$

and that (5) remains prime. So, the representatives of ideal classes are restricted to

$$\{\mathcal{O}_K, \mathfrak{p}, \mathfrak{q}_1, \mathfrak{q}_2, \mathfrak{p}^2\}.$$

But $\mathfrak{p}^2$ is already principal (it is (2)), so we are down to

$$\{\mathcal{O}_K, \mathfrak{p}, \mathfrak{q}_1, \mathfrak{q}_2\}.$$

The class group has order between 1 and 4. NB: $\mathfrak{p}$ is not principal because if it were principal then there would be an element of $\mathcal{O}_K$ whose norm is 2 which we can check is impossible. Similarly, we may check that $\mathfrak{q}_1, \mathfrak{q}_2$ are nonprincipal. So the class group has size 2 or 4 (it can't have order 3 because $\mathfrak{p}$ has order 2). To show that the class group is $\mathbf{Z}/4\mathbf{Z}$, we just need to show that there exists an element whose order is not 1 or 2. So, we just need to show that $\mathfrak{q}_1^2$ is not principal. If $\mathfrak{q}_1^2 = (a + b\sqrt{-17})$ then we can check by computing norms that $a^2 + 17b^2 = 9$, i.e. $a = \pm 3$ and $b = 0$. But this does not satisfy the desired property since if $\mathfrak{q}_1^2 = (3)$ then $\mathfrak{q}_1 = \mathfrak{q}_2$.

- Next time you think you've solved a millenium problem, somebody should shoot you on the spot, so that you think you've solved it when you die. When Heegner's proof of the fact that there are only finitely many quadratic imaginary fields of class number 1 was found to be incorrect, he died thinking that he hadn't solved it. But Stark later found that the error was of an insignificant nature and Heegner's proof works.

Theorem 16.6 (Heegner)

There are finitely many quadratic imaginary fields K with $\mathrm{Cl}_K = \{1\}$.

In real quadratic fields, since there is a negative sign in the equation for the norm, it's somehow easier for ideals to be principal.

Example 16.7

Let $K = \mathbf{Q}(\sqrt{7})$. The Minkowski bound says that every ideal class has a representative of norm at most 2. In fact 2 factors as

$$(3 - \sqrt{7})(3 + \sqrt{7})$$

so all the relevant prime ideals are principal and the class group is trivial.

Example 16.8

Let $K = \mathbf{Q}(\sqrt{-23})$. The Minkowski bound says that every ideal class has a representative of norm at most 3. Letting $\alpha = (1+\sqrt{-23})/2$, we can factor the relevant ideals

$$(2) = (2, \alpha)(2, \alpha - 1) = \mathfrak{p}_1\mathfrak{p}_2$$

and

$$(3) = (3, \alpha)(3, \alpha - 1) = \mathfrak{q}_1\mathfrak{q}_2.$$

We may check that $\mathfrak{p}_1, \mathfrak{q}_1$ are not principal, but that $\mathfrak{p}_1\mathfrak{q}_1$ is principal. So

$$[\mathfrak{p}_1] = -[\mathfrak{q}_1] = [\mathfrak{q}_2]$$

and

$$[\mathfrak{p}_2] = -[\mathfrak{p}_1] = [\mathfrak{q}_1].$$

So the class group is equal, as a set, to

$$\{1, [\mathfrak{p}_1], [\mathfrak{p}_2]\}$$

and since $\mathfrak{p}_1^2$ is not principal we are guaranteed that the class group is $\mathbf{Z}/3\mathbf{Z}$.

§17 April 8, 2019

Today's class is taught by Lynnelle Ye, a student of Mark Kisin. The main goal will be to prove that there are only finitely many number fields of a given discriminant. Professor Kisin said we were used to a certain level of witty comments, which Lynelle says she may not be able to provide.

§17.1 Bounding the Discriminant

Proposition 17.1

Let K be a number field with $[K : \mathbf{Q}] = n$. Then

$$|d_K| \geq \frac{\pi}{3}\left(\frac{3\pi}{4}\right)^{n-1}.$$

Proof. The Minkowski bound guarantees the existence of an integral ideal $\mathfrak{b} \subseteq \mathcal{O}_K$ such that

$$1 \leq N(\mathfrak{b}) \leq \left(\frac{4}{\pi}\right)^{r_2} \frac{n!}{n^n}|d_K|^{1/2}.$$

Rearranging, we obtain

$$|d_K| \geq \left(\frac{\pi}{4}\right)^{2r_2} \frac{n^{2n}}{(n!)^2} \geq \left(\frac{\pi}{4}\right)^{n} \frac{n^{2n}}{(n!)^2}.$$

This bound grows at least exponentially, with common ratio at least

$$\frac{\pi}{4}\frac{(n+1)^{2(n+1)}}{(n+1)!^2}\frac{(n!)^2}{n^{2n}} \geq \frac{\pi}{4}\left(1+\frac{1}{n}\right)^{2n} \geq \frac{3\pi}{4}$$

(the last part is via the binomial theorem), as desired. □

Corollary 17.2

If $K \neq \mathbf{Q}$, then $|d_K| > 1$. In other words there is only one number field of discriminant 1.

This corollary can be useful because (via the upcoming connection between the discriminant and ramification) it shows that $\mathbf{Q}$ has no nontrivial unramified extension.

Theorem 17.3

Fix $d \in \mathbf{N}$. Then there exist only finitely many number fields $K \subseteq \mathbf{C}$ such that $|d_K| = d$.

Proof. By the exponential bound in Proposition 17.1, it suffices to show that there are finitely many number fields of degree n with r_1 real embeddings, r_2 complex embeddings, and discriminant d. Define $B \subseteq \mathbf{R}^{r_1} \times \mathbf{C}^{r_2}$ by the following: in case (a), there exists at least one real embedding. In that case, define

$$B = \{(y_1, \dots, y_{r_1}, z_1, \dots, z_{r_2}) \in \mathbf{R}^{r_1} \times \mathbf{C}^{r_2} : y_1 \leq 2^{n-1}\left(\frac{\pi}{2}\right)^{-r_2}|d|^{1/2}, |y_i| \leq 1/2, |z_i| \leq 1/2\}.$$

In case (b), there are no real embeddings, and we define instead

$$B = \{(z_1, \dots, z_{r_2}) \in \mathbf{C}^{r_2} : |z_1 - \overline{z_1}| \le 2^n \left(\frac{\pi}{2}\right)^{1-r_2} |d|^{1/2}, |z_1 + \overline{z_1}| \le 1/2, |z_i| \le 1/2\}.$$

We can check that Minkowski's theorem holds, since $\mu(B) = 2^n \mathrm{vol}(\sigma(\mathcal{O}_K))$. So, there is some nonzero $x \in \mathcal{O}_K$ such that $\sigma(x) \in B$. The claim is that $K = \mathbf{Q}(x)$. Once we have proven this, we can use the fact that $\sigma_i(x)$ is uniformly bounded for all K. Hence x satisfies a polynomial of degree n with integers coefficients whose coefficients are uniformly bounded. So, there are only finitely many x and hence finitely many K.

To show $K = \mathbf{Q}(x)$, the following works.

$$|N_{K/\mathbf{Q}}(x)| = \left|\prod \sigma_i(x)\right| = \prod_{i=1}^{r_1} |\sigma_i(x)| \prod_{r_1+1}^{r_1+r_2} |\sigma_i(x)|^2.$$

Since all but one of the σ_i's have absolute value less than 1, and $|N_{K/\mathbf{Q}}| \in \mathbf{N}$, it is guaranteed that $|\sigma_1(x)| > 1$ and all the rest are < 1. Actually in case (b) it's slightly different: you get one complex conjugate pair whose magnitudes are greater than 1, and you need to check from the definition of B in that case that this conjugate pair must have different values; in particular $\sigma_1(x)$ cannot be real.

Suppose K strictly contains $\mathbf{Q}(x)$. Then the map σ_1 extends in multiple different ways from $\mathbf{Q}(x)$ to K. This contradicts the fact that σ_1 has a different value on x from all the other embeddings, which means $K = \mathbf{Q}(x)$ as desired. □

§17.2 The Unit Theorem

The next topic is on the group of units of the ring of integers. It is called Dedekind's "Unit Theorem."

Theorem 17.4 (Dedekind)
As an abelian group, $\mathcal{O}_K^\times$ is isomorphic to $\mu_K \times \mathbf{Z}^{r_1+r_2-1}$, where μ_K is the group of roots of unity in K.

The machine that lets us see the structure of the unit group is that of a lattice in a subspace *logarithmic space*, times the group of roots of unity. In particular, there is a homomorphism of abelian groups

$$L : K^\times \to \mathbf{R}^{r_1+r_2}$$

given by

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

To prove the theorem, it will suffice to show that the kernel of L is equal to μ_K and that $L(\mathcal{O}_K^\times)$ is a lattice in the (r_1+r_2-1)-dimensional hyperplane W given by $\sum y_i + 2\sum z_i = 0$.

We've known deep down in our bones that the following lemma is true for a while. It's proof is not hard.

Lemma 17.5
An element $x \in \mathcal{O}_K$ is a unit if and only if $N(x) = \pm 1$.

Proof. If x is a unit, then it has an inverse $x^{-1} \in \mathcal{O}_K$. The multiplicativity of the norm says that

$$N(x)N(x^{-1}) = 1$$

so $N(x) \in \mathbf{Z}^\times = \{\pm 1\}$.

In the other direction, suppose that $N(x) = \pm 1$. Then x satisfies some integer polynomial of the form

$$x^n + a_{n-1}x^{n-1} + \cdots + x \pm 1 = 0$$

so factoring out an x and moving the ± 1 over yields an explicit inverse for x.

Alternatively, the product of the nontrivial Galois conjugates of x is the inverse of x in K. The Galois conjugates are algebraic integers in $\overline{\mathbf{Q}}$, as is their product, so the inverse of x in K is an algebraic integer, i.e. it is in $\mathcal{O}_K$.

Alternatively again, $N(x) = \pm 1$ means $\mathcal{O}_K/(x)$ has cardinality 1, so $(x) = \mathcal{O}_K$, i.e. x is a unit. □

If $x \in \mathcal{O}_K^\times$ then taking the log we get $0 = \sum_{i=1}^{r_1} \log|\sigma_i(x)| + 2\sum_{i=r_1+1}^{r_1+r_2} \log|\sigma_i(x)|$ so indeed $L(\mathcal{O}_K^\times) \subseteq W$.

§18 April 10, 2019

Today's class is also taught by Lynelle Ye.

§18.1 The Proof of the Unit Theorem

Recall that W is a distinguished hyperplane in logarithmic space $\mathbf{R}^{r_1+r_2}$, and that the units of $\mathcal{O}_K$ map into W under the logarithm map L. To prove the unit theorem, we need to first check two things:

(a) $\ker(L|_{\mathcal{O}_K^\times})$ is finite.

(b) $L(\mathcal{O}_K^\times)$ is discrete in W.

The idea is that if B is a bounded set in W and $L(x) \in B$, then all the $|\sigma_i(x)|$'s are bounded, so $\sigma(x)$ is in some bounded subset of $\mathbf{R}^{r_1} \times \mathbf{C}^{r_2}$, which we already know has finite intersection with $\sigma(\mathcal{O}_K)$, so x can take on only finitely many values [alternatively, x satisfies an integral polynomial with coefficients in a bounded range]. It follows that the image of L is discrete in W and that, since $\{0\}$ is bounded, the kernel of L is finite. In fact, it consists of the group of roots of unity in $\mathcal{O}_K$. This is because every element of the kernel has finite order so it is a root of unity.

It now remains to check that the image of L has full rank in W.

Example 18.1

Let $K = \mathbf{Q}(\sqrt{d})$ be a real quadratic field, so that $r_1 = 2$ and $r_2 = 0$. The roots of unity in K are just ± 1. The two real embeddings σ_1, σ_2 take $\sqrt{d} \to \pm\sqrt{d}$. Take

$$B_{\alpha,\lambda} = \{(x_1, x_2) \in \mathbf{R}^2 : |x_1| \le \lambda, |x_2| \le \alpha/\lambda\}$$

and note that $\mu(B_{\alpha,\lambda}) = \alpha$. Take α so that Minkowski's theorem applies to $B_{\alpha,\lambda}$ and $\sigma(\mathcal{O}_K)$. For all λ, there is a nonzero $x_\lambda \in B_{\alpha,\lambda} \cap \sigma(\mathcal{O}_K)$. There exist finitely many possible ideals (x_λ) because

$$1 \le |N(x_\lambda)| \le |\sigma_1(\lambda)\sigma_2(\lambda)| = \alpha.$$

Despite generating only finitely many ideals, this process generates infinitely many x_λ's. This is because we can keep making the rectangles thin enough to not contain any of the previous $\pm x_\lambda$'s (NB the second coordinate of x_λ must be nonzero since x_λ is nonzero).

It follows that there are infinitely many x_λ's whose ratio is not ± 1, hence $L(\mathcal{O}_K^\times)$ has rank at least 1. But $r_1 + r_2 - 1 = 1$ so we have shown it is a lattice as desired.

General Proof. It suffices to show that for any nonzero linear map $f : W \to \mathbf{R}$ we can find an $x \in \mathcal{O}_K^\times$ such that $f(L(x)) \neq 0$. Let $r = r_1 + r_2 - 1$.

For $\alpha, \lambda_1, \dots, \lambda_r \in \mathbf{R}^+$ let $\lambda_{r+1} \in \mathbf{R}^+$ be such that

$$\prod_{i=1}^{r_1} \lambda_i \prod_{i=r_1+1}^{r_1+r_2} \lambda_i^2 = \alpha$$

and $B_{\alpha,\lambda}$ be defined in the way analogous to in the example, i.e.

$$B_{\alpha,\lambda} := \{(y_1, \dots, y_{r_1}, z_1, \dots, z_{r_2}) \in \mathbf{R}^{r_1} \times \mathbf{C}^{r_2} : |y_i| \le \lambda_i, |z_j| \le \lambda_{j+r_1}\}.$$

Taking α to be big enough that Minkowski's theorem holds, we know that there exists a nonzero $x_\lambda \in B_{\alpha,\lambda} \cap \sigma(\mathcal{O}_K)$. There are only finitely many ideals (x_λ), because $|N(x_\lambda)| \le \alpha$ for all choices of λ. The point of all this is that we want to find infinitely many $f(L(x_\lambda))$'s which lie in disjoint intervals. Once we have that, we have infinitely many x_λ's with different images under $f \circ L$, which generate only finitely many ideals. So there are two distinct x_λ's which differ multiplicatively by a unit u. Hence $f(L(u)) = f(L(x_{\lambda_1})) - f(L(x_{\lambda_2})) \neq 0$ as desired.

In particular, we already saw that

$$1 \le |N_{K/\mathbf{Q}}(x_\lambda)| = \prod |\sigma_i(x_\lambda)| \le \alpha$$

by definition of x_λ. Furthermore, for any i we may check that $|\sigma_i(x_\lambda)| \ge \frac{\lambda_i}{\alpha}$. In the end, we have the bounds

$$\frac{\lambda_i}{\alpha} \le |\sigma_i(x_\lambda)| \le \lambda_i$$

and thus

$$0 \le \log \lambda_i - \log |\sigma_i(x_\lambda)| \le \log \alpha.$$

Let $f(y)$ be given by $\sum_{i=1}^r c_i y_i$ for a given $(y_1, \dots, y_{r+1}) \in W$ (summing over only r of the coordinates because the projection of W onto the first r coordinates is an isomorphism by its definition). Then

$$f(L(x_\lambda)) = \sum_{i=1}^r c_i \log |\sigma_i(x_\lambda)|$$

so from the previous inequality we obtain

$$\left| f(L(x_\lambda)) - \sum_{i=1}^r c_i \log \lambda_i \right| \le \sum_{i=1}^r |c_i| \log \alpha.$$

Constructing the appropriate λ's is done in the following way. Take β to be a positive constant strictly larger than $\sum |c_i| \log \alpha$. For a given $h \in \mathbf{N}$, select the λ_i's so that $\sum_{i=1}^r c_i \log \lambda_i = 2\beta h$ (here it's crucial that at least one of the c_i's is nonzero, i.e. that f is nonzero; this also uses the fact that this only sums to r instead of $r+1$, so that the λ_i's may all be as large as we want independently of α). Then

$$|f(L(x_\lambda)) - 2\beta h| < \beta,$$

i.e.

$$(2h-1)\beta < f(L(x_\lambda)) < (2h+1)\beta.$$

This guarantees an infinite set (one for each value of h) of x_λ's for which $f(L(x_\lambda))$ are all distinct, which by the previous discussion proves the existence of a unit $u \in \mathcal{O}_K^\times$ such that $f(L(u)) \neq 0$. This holds for any nonzero linear functional f on W, which means that $L(\mathcal{O}_K^\times)$ must be a lattice (of full rank) in W. □

§19 April 15, 2019

The Unit theorem tells you the rank of the unit group, but it's generally a tricky problem to actually compute a basis for it. Even if you find a subgroup of the unit group with the correct rank, it's completely unclear how to check that you have all of them.

§19.1 The Unit Group of a Cyclotomic Field

Example 19.1

In the case $K = \mathbf{Q}(\zeta_p)$ where p is an odd prime, there are no real embeddings, so the rank of $\mathcal{O}_K^\times$ is $r_1 + r_2 - 1 = \frac{p-3}{2}$. We have units $u_i := (\zeta_p^i - 1)/(\zeta_p - 1) = 1 + \cdots + \zeta_p^{i-1}$. Moreover,

$$u_i(-\zeta_p^{p-i}) = \frac{\zeta_p^{p-i} - 1}{\zeta_p - 1} = u_{p-i}$$

so u_i and u_{p-i} are related by a torsion element (the root of unity $-\zeta_p^{p-i}$). It turns out (though we won't prove it) that there is a basis of $\mathcal{O}_K^\times/(\mathcal{O}_K^\times)_{\text{tors}}$ given by

$$u_2, \dots, u_{(p-1)/2}$$

§19.2 Pell's Equation

Let $d \in \mathbf{N}$ be squarefree, and consider the real quadratic field $K = \mathbf{Q}(\sqrt{d})$. The units in $\mathcal{O}_K$ are the $a + b\sqrt{d} \in \mathcal{O}_K$ such that $a^2 - db^2 = \pm 1$. The roots of unity in K are just ± 1. The unit theorem (from the fact that there are 2 real embeddings) says that $\mathcal{O}_K^\times = \mathbf{Z} \times \mathbf{Z}/2\mathbf{Z}$. In particular, there exists a *fundamental unit* $u \in \mathcal{O}_K^\times$ such that u generates $\mathcal{O}_K^\times/\{\pm 1\}$, i.e. every unit is ± 1 times a (possibly negative) power of u.

The solutions to Pell's equation are not exactly the units in K. If $d \equiv 1 \pmod 4$, then some units might have half-integer coefficients.

If d is not 1 mod 4, then we just look for the smallest solution to $a^2 - db^2 = \pm 1$ to find the fundamental unit. What do we mean by "smallest" in this case? The answer is simple: observe that in this case ($d > 0$), raising $a + b\sqrt{d} \neq 1$ to a positive power strictly increases the real coordinate. Hence, a unit that generates the free part of the unit group (a *fundamental unit*) is a nontrivial one whose real coordinate is minimal. It's a finite computation to compute the smallest integer $a \in \mathbf{Z}$ such that $a^2 - db^2 = \pm 1$ has any integer solutions not given by $(a, b) = (1, 0)$.

Example 19.2

Let $d = 7$. Then the fundamental unit of $\mathbf{Q}(\sqrt{d})$ is $w = 8 + 3\sqrt{7}$.

If d is 1 mod 4 then by the description of $\mathcal{O}_{\mathbf{Q}(\sqrt{d})}$ from Proposition 1.8, we are finding integers A, B of the same parity for which $A^2 - dB^2 = \pm 4$.

Example 19.3

Let $d = 5$. The fundamental unit of $\mathbf{Q}(\sqrt{d})$ is then given by $2w = 1 + \sqrt{5}$.

In general one sees that w^3 has integer coefficients (c.f. homework 5), so in Example 19.3, $\mathbf{Z}[\sqrt{d}]^\times$ is generated by w^3 (because its index divides three and it is a proper subgroup).

This gives what is so-far an ad-hoc method of finding the solutions to Pell's Equation: you find the smallest solution, and the theory of the unit group tells you the other ones are all (up to a root of unity) a power of that solution.

§19.3 Factorization of Primes in Extensions

Let L/K be an extension of number fields. If $\mathfrak{p} \subseteq \mathcal{O}_K$ is a prime, then $\mathfrak{p} \cdot \mathcal{O}_L$ factors uniquely into primes in L. In fact, the primes $\mathfrak{q}_i$ in the factorization of $\mathfrak{p}$ in $\mathcal{O}_L$ *lie over* $\mathfrak{p}$ in the sense that $\mathfrak{q}_i \cap \mathcal{O}_K = \mathfrak{p}$ (this is because the intersection is prime and $\mathfrak{p}$ is maximal). Recall that this situation is analogous to that of a covering map of topological spaces: we have a map from the primes of $\mathcal{O}_L$ to those of $\mathcal{O}_K$ which takes a prime to its intersection with $\mathcal{O}_K$. The fibers are just the primes that something factors into, and ramification corresponds to a prime appearing more than once in the factorization.

Definition 19.4. A prime $\mathfrak{p} \subseteq \mathcal{O}_K$ is called **ramified** in L if $e_i > 1$ for some i.

There's one more important invariant corresponding to a prime lying over another one.

Definition 19.5. Let $\mathfrak{q}$ be a prime in $\mathcal{O}_L$ lying over $\mathfrak{p}$. The **inertial degree** of $\mathfrak{q}$ over $\mathfrak{p}$ is

$$f(\mathfrak{q}|\mathfrak{p}) = [\mathcal{O}_L/\mathfrak{q} : \mathcal{O}_K/\mathfrak{p}].$$

It will turn out that $\sum_{\mathfrak{q}|\mathfrak{p}} e(\mathfrak{q}|\mathfrak{p}) f(\mathfrak{q}|\mathfrak{p}) = [L : K]$. Here, $\mathfrak{q}|\mathfrak{p}$ is convenient notation for "a prime $\mathfrak{q} \subseteq \mathcal{O}_L$ lying over a prime $\mathfrak{p} \subseteq \mathcal{O}_K$". Indeed, such pairs have the property that $\mathfrak{q}|\mathfrak{p} \cdot \mathcal{O}_L$.

Here are some remarks, only the second of which we will justify in this class.

- Consider $\mathcal{O}_L \otimes \mathcal{O}_{\overline{K}}$ over $\mathcal{O}_{\overline{K}}$. In some way, this allows us to take apart all the multiple points of the covering map, making sure that every prime splits completely.
- It will turn out that only finitely many primes are ramified.
- If L/K is Galois and abelian, it's a famous theorem from number theory that all the e_i's and f_i's are 1 (the prime **splits completely**) for a proportion of primes equal to $1/[L : K]$ (to define this rigorously one uses the Dirichlet density).

§20 April 17, 2019

§20.1 The Chinese Remainder Theorem and its Consequences

Lemma 20.1 (Sunzi's "Chinese Remainder" Theorem)
Let $\mathfrak{a} \subseteq \mathcal{O}_L$ be an ideal which factors as $\prod \mathfrak{q}_i^{e_i}$. Then

$$\mathcal{O}_L/\mathfrak{a} \cong \prod \mathcal{O}_L/\mathfrak{q}_i^{e_i}.$$

Proof. We have projections $\mathcal{O}_L/\mathfrak{a} \to \mathcal{O}_L/\mathfrak{q}_i^{e_i}$ by the factorization (the ideals $\mathfrak{q}_i^{e_i}$ divide and thus contain $\mathfrak{a}$). We can put these together to get a map of finite sets

$$\mathcal{O}_L/\mathfrak{a} \to \prod \mathcal{O}_L/\mathfrak{p}_i^{e_i}$$

By the multiplicativity of the norm, these have the same size and it suffices to show that the map is injective. In a Dedekind domain, we know that if ideals $\mathfrak{a}, \mathfrak{b}$ are coprime (according to their factorization) then they are comaximal, since the ideal they generate must contain, and thus divide both of them. It's a fact from algebra that this means that $\mathfrak{a} \cap \mathfrak{b} = \mathfrak{a}\mathfrak{b}$. We can apply this fact to powers of different primes $\mathfrak{q}_i^{e_i}$. In particular, it means that

$$\mathfrak{a} = \bigcap \mathfrak{q}_i^{e_i}.$$

The kernel of our candidate isomorphism is the intersection of all the kernels of the projections modulo $\mathfrak{a}$ which is then zero. This proves injectivity, so we have the desired isomorphism (since it is a map of finite sets). □

The Chinese Remainder Theorem is the basis for some intermediate results which will prove useful in the discussion to come.

Lemma 20.2
Let k be a field and

$$0 \to V_1 \to V \to V_2 \to 0$$

be an exact sequence of k-vector spaces commuting with linear maps $\phi_i : V_i \to V_i$ and $\phi : V \to V$. Then the trace of ϕ as a map on V is the sum of the traces of ϕ_1 and ϕ_2.

Proof. By linear algebra, V is isomorphic the direct sum $V_1 \oplus V_2$, and the commutativity of the diagram means that the map $\phi : V \to V$ is equal to ϕ_1 on the first coordinate and ϕ_2 on the second coordinate. In other words, as a matrix it has two blocks along the diagonal which are ϕ_1 and ϕ_2. So its trace is indeed the sum of the traces. □

Lemma 20.3
If $\mathfrak{q} \subseteq \mathcal{O}_L$ is a prime, with

$$\mathfrak{q}^e \supseteq p \cdot \mathcal{O}_L.$$

Then for all $\alpha \in \mathcal{O}_L$,

$$\mathrm{Tr}_{\mathbf{F}_p}(\alpha|_{\mathcal{O}_L/\mathfrak{q}^e}) = e \,\mathrm{Tr}_{\mathbf{F}_p}(\alpha|_{\mathcal{O}_L/\mathfrak{q}}).$$

Proof. First of all, we have an isomorphism of $\mathcal{O}_L/\mathfrak{p}$-vector spaces $\mathfrak{p}^i/\mathfrak{p}^{i+1} \cong \mathcal{O}_L/\mathfrak{p}$. This isomorphism commutes with multiplication by any element, so in particular,

$$\mathrm{Tr}_{\mathbf{F}_p}(\alpha|_{\mathcal{O}_L/\mathfrak{p}}) = \mathrm{Tr}_{\mathbf{F}_p}(\alpha|_{\mathfrak{p}^i/\mathfrak{p}^{i+1}})$$

for any i. Note that we have an exact sequence of $\mathcal{O}_L/\mathfrak{p}$-vector spaces

$$0 \to \mathfrak{p}^{i-1}/\mathfrak{p}^i \to \mathcal{O}_L/\mathfrak{p}^i \to \mathcal{O}_L/\mathfrak{p}^{i-1} \to 0$$

with commuting maps all given by multiplication by α. Using this process and inducting, we get that

$$\mathrm{Tr}_{\mathbf{F}_p}(\alpha|_{\mathcal{O}_L/\mathfrak{p}^e}) = \sum_{i=1}^{e} \mathrm{Tr}_{\mathbf{F}_p}(\alpha|_{\mathfrak{p}^{i-1}/\mathfrak{p}^i}) = e\,\mathrm{Tr}_{\mathbf{F}_p}(\alpha|_{\mathcal{O}_L/\mathfrak{p}})$$

as desired. □

Lemma 20.4

Let $\alpha \in \mathcal{O}_L$. Then in $\mathbf{F}_p$,

$$\mathrm{Tr}_{\mathbf{F}_p}(\alpha|_{\mathcal{O}_L/(p)}) = \sum e_{\mathfrak{p}}\,\mathrm{Tr}_{\mathbf{F}_p}(\alpha|_{\mathcal{O}_L/\mathfrak{p}}).$$

Proof. The result is immediate via Sunzi's theorem and the previous two results. □

§20.2 Ramification via the Different Ideal

We want to show that finitely many primes of K ramify in L. It's enough to show this for $K = \mathbf{Q}$ because $\mathfrak{p}$ ramifies in L implies $\mathfrak{p} \cap \mathbf{Z}$ does ($\mathfrak{p} \cap \mathbf{Z}$ is the unique prime in $\mathbf{Z}$ which $\mathfrak{p}$ lies over).

Now we have reduced to the case $K = \mathbf{Q}$, the important fact that we will use is that $\mathcal{O}_L$ has a $\mathbf{Z}$-basis (NB: this is not true in general; here we are using the fact that $\mathbf{Z}$ is a PID). Our goal will be to show that the ramification data of the extension $L/\mathbf{Q}$ is contained in an ideal called the *different*.

Definition 20.5. The **different** is the ideal in $\mathcal{O}_L$ given by $\mathcal{D}_L = (\mathcal{O}_L^\vee)^{-1}$.

Theorem 20.6

Let $\mathcal{D}_L$ have prime factorization $\prod \mathfrak{p}^{m_{\mathfrak{p}}}$. Then for any $p \in \mathbf{Z}$ prime, if we have the factorization $p \cdot \mathcal{O}_L = \prod_{\mathfrak{p}} \mathfrak{p}^{e_{\mathfrak{p}}}$, then $m_{\mathfrak{p}} \geq e_{\mathfrak{p}} - 1$. If p does not divide $e_{\mathfrak{p}}$ then we have $m_{\mathfrak{p}} = e_{\mathfrak{p}} - 1$.

Proof. The result we want is equivalent to

$$\mathcal{D}_L \subseteq p \cdot \mathcal{O}_L \prod_{\mathfrak{p}|p} \mathfrak{p}^{-1} = \prod \mathfrak{p}^{e_{\mathfrak{p}}-1}$$

where, for each $\mathfrak{p}$ such that p does not divide $e_{\mathfrak{p}}$, multiplying the right hand side by $\mathfrak{p}$ causes the containment to no longer hold.

Multiplying by $\mathcal{O}_L^\vee$, we obtain this is equivalent to

$$\prod_{\mathfrak{p}|p} \mathfrak{p} \subseteq p \cdot \mathcal{O}_L^\vee,$$

where for any $\mathfrak{p}$ such that p does not divide $e_{\mathfrak{p}}$ we want to get that removing $\mathfrak{p}$ from the LHS makes the equation false. For any ideal I in $\mathcal{O}_L$, having $I \subseteq p \cdot \mathcal{O}_L^\vee$ is equivalent to $\mathrm{Tr}_{L/\mathbf{Q}}(I) \subseteq p \cdot \mathbf{Z}$ [check this from the definition of the dual module].

If $\alpha \in L$, then the trace of α is the trace of the multiplication map by α on L. Since $\mathcal{O}_L$ is a free $\mathbf{Z}$-module, we can write this as the trace of matrix given by the action of α on an arbitrary $\mathbf{Z}$-basis for $\mathcal{O}_L$. Reducing this mod p, if $\alpha \in \mathcal{O}_L$,

$$\mathrm{Tr}_{L/\mathbf{Q}}(\alpha) = \mathrm{Tr}_{\mathcal{O}_L/\mathbf{Z}}(\alpha) \equiv \mathrm{Tr}_{\mathbf{F}_p}(\alpha|_{\mathcal{O}_L/(p)}) = \sum_{\mathfrak{p}|p} e_{\mathfrak{p}} \mathrm{Tr}_{\mathbf{F}_p}(\alpha|_{\mathcal{O}_L/\mathfrak{p}}) \mod p.$$

where the last equality is by Lemma 20.4 (notice that there is a reduction modulo p happening in the middle). Setting $I = \prod \mathfrak{p}$, we have $\alpha \in I \implies \alpha \in \mathfrak{p}$ for all $\mathfrak{p}|p$. So the multiplication map by α has trace zero on each quotient, which means $\mathrm{Tr}_{L/\mathbf{Q}}(\alpha) \in p\mathbf{Z}$. Thus, we have the desired inclusion of the product of these primes into p times the dual module. This proves the desired inequality $m_{\mathfrak{p}} \geq e_{\mathfrak{p}} - 1$. Now it suffices to show that equality holds unless $p|e_{\mathfrak{p}}$.

Fix $\mathfrak{q}|p$ such that p does not divide $e_{\mathfrak{q}}$. It suffices to show that $\prod_{\mathfrak{p}\neq\mathfrak{q}} \mathfrak{p}$ is not contained in $p \cdot \mathcal{O}_K^\vee$. In other words, it suffices to find an $\alpha \in \prod_{\mathfrak{p}\neq\mathfrak{q}} \mathfrak{p}$ whose trace is nonzero mod p. But by the Chinese Remainder Theorem, we may select $\alpha \notin \mathfrak{q}$ and to be any prescribed value $\tilde{\alpha}$ modulo $\mathfrak{q}$. In particular, again via Lemma 20.4,

$$\mathrm{Tr}_{L/\mathbf{Q}}(\alpha) \equiv \sum_{\mathfrak{p}|p} e_{\mathfrak{p}} \mathrm{Tr}_{\mathbf{F}_p}(\alpha|_{\mathcal{O}_L/\mathfrak{p}}) \mod p.$$

But for all $\mathfrak{p} \neq \mathfrak{q}$, we have $\alpha \in \mathfrak{p}$, so all of those traces vanish in $\mathbf{F}_p$. By assumption, $e_{\mathfrak{q}}$ does not vanish in $\mathbf{F}_p$. So we are left with

$$\mathrm{Tr}_{L/\mathbf{Q}}(\alpha) = e_{\mathfrak{q}} \mathrm{Tr}_{\mathbf{F}_p}(\tilde{\alpha}|_{\mathcal{O}_L/\mathfrak{q}}).$$

But $(\mathcal{O}_L/\mathfrak{q})/\mathbf{F}_p$ is an extension of finite fields, so its trace down to $\mathbf{F}_p$ is not identically zero (the proof of this is postponed until next class when we will know slightly more about finite fields; see Corollary 21.4). In particular, we can select α so that $\tilde{\alpha}$ has nonzero trace, and this guarantees that α has nonzero trace mod p, as desired. □

Corollary 20.7

a prime p in $\mathbf{Z}$ ramifies in L if and only if there exists a prime on top of it dividing the different.

Corollary 20.8

Only finitely many primes in K ramify in L.

Proof. If $(p) \subseteq \mathbf{Z}$ ramifies in L, then there must exists a prime in L sitting on top of it dividing the different ideal. But the different only has finitely many divisors, and each prime in L lies over a unique prime in $\mathbf{Q}$, so there are only finitely ramified primes in $\mathbf{Q}$. But if a prime in K ramifies, then so does its intersection with $\mathbf{Z}$. And only finitely many primes in K lie over the same one in $\mathbf{Q}$. So, we can conclude that only finitely many ideals in K ramify. □

§21 April 22, 2019

§21.1 More facts about ramification

Recall that we have an *inertial degree* $f_{\mathfrak{p}} = [\mathcal{O}_L/p : \mathbf{F}_p]$.

Corollary 21.1

$$[L : \mathbf{Q}] = \sum_{\mathfrak{p}|p} e_{\mathfrak{p}} f_{\mathfrak{p}}.$$

Proof. By the Chinese Remainder Theorem,

$$\dim_{\mathbf{F}_p} \mathcal{O}_L/p = \sum_{\mathfrak{p}} \dim \mathcal{O}_L/\mathfrak{p}^{e_{\mathfrak{p}}} = \sum e_{\mathfrak{p}} f_{\mathfrak{p}}$$

where the last equality is because we saw that

$$0 \to \mathfrak{p}^{s-1}/\mathfrak{p}^s \to \mathcal{O}_L/\mathfrak{p}^s \to \mathcal{O}_L/\mathfrak{p}^{s-1} \to 0$$

is an exact sequence and by induction the dimension of $\mathcal{O}_L/\mathfrak{p}^s$ over $\mathcal{O}_L/\mathfrak{p}$ is s. Hence the dimension over $\mathbf{F}_p$ of $\mathcal{O}_L/\mathfrak{p}^e$ is $e[\mathcal{O}_L/\mathfrak{p} : \mathbf{F}_p] = ef_{\mathfrak{p}}$. □

Example 21.2

Let $L = \mathbf{Q}(\zeta_p)$, and suppose we want to find the ramified primes in L. We know that

$$\mathcal{O}_L = \mathbf{Z}[\zeta_p] = \mathbf{Z}[X]/(f(X))$$

where $f(X) = (X^p - 1)/(X - 1)$. Let ℓ be a rational prime. Then

$$\mathcal{O}_L/\ell = \mathbf{F}_\ell[X]/f(X)$$

and if f factors as a product of powers e_i of irreducible f_i's in $\mathbf{F}_\ell$ then $\ell \cdot \mathcal{O}_L = \prod(\ell, \overline{f}_i)$, and

$$\mathcal{O}_L/\ell = \prod \mathcal{O}_/\mathfrak{p}_i^{e_i}.$$

ℓ being unramified is equivalent to $\mathcal{O}_L/\ell$ being a product of fields and thus to f being separable mod ℓ. To check this, we just need to ensure that it shares no roots with its derivative. In particular,

$$f(x)(x-1) = x^p - 1$$

so

$$f'(x)(x-1) + f(x) = px^{p-1}.$$

If we plug in a root of f for x, we get that the derivative is indeed nonzero if ℓ is not equal to p. So the only ramified prime is p. It factors as

$$\prod(1 - \zeta_p^i)$$

which means that (p) is just $(1 - \zeta_p)^{p-1}$. By our characterization of the different, this means the different of the cyclotomic field is $(1 - \zeta_p)^{p-2}$.

Recall that when $L = \mathbf{Q}(\zeta_p)$ we saw that $d_L = p^{p-2}$. So we might expect the discriminant to have something to do with ramification (in fact it's common to skip the discussion on the different ideal entirely; discriminants tend to be somewhat easier to compute).

§21.2 Finite fields

First we need some finite fields review. Let $q = p^r$, and define $\mathbf{F}_q$ to be the subset of $\overline{\mathbf{F}}_p$ consisting of elements fixed by $x \mapsto x^q$.

Proposition 21.3 (Finite fields review)

Some facts about finite fields.

(i) $\mathbf{F}_q$ is a subfield of $\overline{\mathbf{F}}_p$ of degree r over $\mathbf{F}_p$.

(ii) If $q_1 = p^{r_1}$ then $\mathbf{F}_{q_1} \supseteq \mathbf{F}_q$ iff $r | r_1$.

(iii) The union of the $\mathbf{F}_q$'s is $\overline{\mathbf{F}}_p$.

(iv) The $\mathbf{F}_{q_1}/\mathbf{F}_q$ is a Galois extension with Galois group generated by the Frobenius automorphism $x \mapsto x^q$.

Proof. Since iterates of $x \mapsto x^p$ preserve addition, $\mathbf{F}_q$ is clearly a subring of $\overline{\mathbf{F}}_p$. It is a field because it is finite (a polynomial has finitely many roots; recall that finite integral domains are fields).

The backwards direction of (ii) is because $X^{p^{r_1}}$ is X raised to the p^r power r'/r times (so the inclusion can be checked directly). The opposite direction is for reasons of dimension: if $\mathbf{F}_{q_1} \supseteq \mathbf{F}_q$ we must have $[\mathbf{F}_q : \mathbf{Q}] | [\mathbf{F}_{q_1} : \mathbf{Q}]$, i.e. $r | r_1$.

For (iii) Let x be a nonzero element of $\overline{\mathbf{F}}_p$. Then $\mathbf{F}_p(x)$ is a finite extension of $\mathbf{F}_p$ (because $\overline{\mathbf{F}}_p$ is algebraic over $\overline{\mathbf{F}}_p$ by definition of the algebraic closure), so x is in one of the $\mathbf{F}_q$'s (in particular, all finite fields are embedded in $\overline{\mathbf{F}}_p$ as one of the $\mathbf{F}_q$'s since splitting fields are unique).

For (iv), since $[\mathbf{F}_{q_1} : \mathbf{F}_q] = r_1/r$, it suffices to show that the Frobenius automorphism has order r_1/r in $\mathrm{Aut}(\mathbf{F}_{q_1}/\mathbf{F}_q)$. Suppose that the order of the Frobenius automorphism was $s < r_1/r$. Then we would have

$$x^{p^{rs}} = x$$

for all $x \in \mathbf{F}_{p^{r_1}}$, in other words we would have $\mathbf{F}_{p^{r_1}} \subseteq \mathbf{F}_{p^{rs}}$ which is impossible because of their dimensions (as $rs < r_1$). This shows that the Frobenius automorphism is indeed a generator of the automorphism group, which thus has the maximal possible size $r_1/r = [\mathbf{F}_{q_1} : \mathbf{F}_q]$, so indeed the extension is Galois. □

Corollary 21.4

The trace map $\mathrm{Tr}_{L/K} : L \to K$ of an extension of finite fields is nonzero.

Proof. Since L/K is Galois, all the embeddings of L into $\overline{K}$ are automorphisms of L/K (there are exactly $[L : K] = |\mathrm{Gal}(L/K)|$ such automorphisms so all the embeddings are accounted for). □

Note that Corollary 21.4 fills in the missing step in Theorem 20.6

§21.3 Ramification and the Discriminant

Proposition 21.5
$N(\mathcal{D}_L) = |d_L|$.

Proof. We use a lemma

Lemma 21.6
Let $I \subseteq J$ be fractional ideals of L. Then

$$|J/I| = N(I)/N(J).$$

Proof. Choose $f \in \mathcal{O}_L$ so that $fJ \subseteq \mathcal{O}_L$. Then the index we want is the same as $|fJ/fI|$. Then we have

$$|\mathcal{O}_L/fI| = |\mathcal{O}_L/fJ| \cdot |fJ/fI|$$

from which the result follows. □

As a result of Lemma 21.6, we may compute

$$N(\mathcal{D}_L) = [\mathcal{D}_L^{-1} : \mathcal{O}_L] = [\mathcal{O}_L^\vee : \mathcal{O}_L] = d_L$$

where the second equality is due to Lemma 21.6 ($|\mathcal{D}_L^{-1}/\mathcal{O}_K| = 1/N(\mathcal{D}_L^{-1}) = N(\mathcal{D}_L)$), the next is by the definition of the different, and the last is by Lemma 7.3. □

Corollary 21.7
p ramifies in L if and only if $p|d_K$.

Proof. p ramifies in L if and only if $m_\mathfrak{p} \geq 1$ for some $\mathfrak{p}|p$. Since the ideal norm is multiplicative, this is equivalent to p dividing the norm of the different. □

Corollary 21.8
If $L \neq \mathbf{Q}$, then some rational prime ramifies in L.

§22 April 24, 2019

The splitting of primes in the case where L/K is Galois is much more restrictive.

§22.1 Prime Decomposition in Galois Extensions

Proposition 22.1

Suppose L/K is Galois, and G its Galois group. Then G acts transitively on the primes in $\mathcal{O}_L$ lying over a fixed prime $\mathfrak{p} \subseteq \mathcal{O}_K$.

Proof. Let $\sigma \in G$. Since σ respects polynomials, the image of any element of $\mathcal{O}_L$ is in $\mathcal{O}_L$. Factoring $\mathfrak{p}$ in $\mathcal{O}_L$, we have

$$\mathfrak{p} \cdot \mathcal{O}_L = \prod \mathfrak{q}_i^{e_i}$$

and so

$$\mathfrak{p} \cdot \mathcal{O}_L = \sigma(\mathfrak{p} \cdot \mathcal{O}_L) = \prod \sigma(\mathfrak{q}_i)^{e_i}.$$

So the action of G takes primes lying over $\mathfrak{p}$ to primes lying over $\mathfrak{p}$. It remains to show transitivity.

Let $\alpha \in \mathfrak{q}_1$ but not in the other primes lying over $\mathfrak{p}$. Such an element is guaranteed to exist by Sunzi's theorem. We have

$$\prod_\sigma \sigma(\alpha) \in \mathcal{O}_K \cap \mathfrak{q}_1 = \mathfrak{p} \subseteq \mathfrak{q}_i$$

[it is in $\mathcal{O}_K$ because it is the norm of $\alpha \in \mathcal{O}_L$. It is in $\mathfrak{q}_1$ because α is one of the terms]. So some σ takes α to $\mathfrak{q}_i$. It suffices to show that actually σ takes $\mathfrak{q}_1$ to $\mathfrak{q}_i$. Since $\sigma(\alpha) \in \mathfrak{q}_i$, it isn't in $\mathfrak{q}_j$ for any $j \neq i$ (or else $\sigma^{-1}(\sigma(\alpha)) = \alpha$ would be in two different primes). So actually $\sigma(\mathfrak{q}_1) = \mathfrak{q}_i$, and since i was arbitrary this means the action is transitive (the choice of $\mathfrak{q}_1$ was also arbitrary and taken without loss of generality). □

Corollary 22.2

If L/K is Galois, then for any prime $\mathfrak{p} \subseteq \mathcal{O}_K$, the factorization

$$\mathfrak{p} \cdot \mathcal{O}_L = \prod \mathfrak{q}_i^{e_i}$$

with inertial degrees f_i has the property that all the f_i's are equal, as are all the e_i's.

Proof. Since the Galois group acts transitively on the primes lying over $\mathfrak{p}$, we can apply σ to the prime factorization of $\mathfrak{p}$ to get that $e_i = e_j$. Moreover, the map $\sigma : \mathcal{O}_L \to \mathcal{O}_L$ gives an isomorphism $\mathcal{O}_L/\mathfrak{q}_i \to \mathcal{O}_L/\mathfrak{q}_j$ so the inertial degrees are also all the same. □

" The Minkowski bound is definitely written in God's book. Maybe he improved it, but it's definitely in his book."

"The splitting of primes in Galois extensions is kind of like grilled cheese, but the Minkowski bound is like much more refined cuisine, such as Bouillabaise"

§22.2 Decomposition and Inertia Groups

Proposition 22.3

Let L/K be a Galois extension, and $\mathfrak{q}_i$ lying over a prime $\mathfrak{p}$. Then

- $\mathcal{O}_L/\mathfrak{q}_i$ is a Galois extension of $\mathcal{O}_K/\mathfrak{p}$.
- $G_\mathfrak{q} = \{\sigma \in \mathrm{Gal}(L/K) : \sigma(\mathfrak{q}) = \mathfrak{q}\}$ maps surjectively onto $\mathrm{Gal}(\mathcal{O}_L/\mathfrak{q}/\mathcal{O}_K/\mathfrak{p})$ via the canonical map θ.
- $I_\mathfrak{q} := \ker\theta$ has order e.

Proof. Let $K' = L^{G_\mathfrak{q}}$ be the fixed field under the action of $G_\mathfrak{q}$. Let $\mathfrak{p}' = \mathfrak{q} \cap \mathcal{O}_{K'}$. Then $\mathfrak{p}' \cdot \mathcal{O}_L$ has $\mathfrak{q}$ as a prime factor. Since $G_\mathfrak{q} = \mathrm{Gal}(L/K')$ permutes the prime factors of $\mathfrak{p}' \cdot \mathcal{O}_L$ transitively but also sends $\mathfrak{q}$ to itself, we know that in fact $\mathfrak{q}$ is the only prime lying over $\mathfrak{p}'$.

$$\mathfrak{p}' \cdot \mathcal{O}_L = \mathfrak{q}^{e'}.$$

The inertial degree f' is at most f, since

$$\mathcal{O}_K/\mathfrak{p} \subseteq \mathcal{O}_{K'}/\mathfrak{p}' \subseteq \mathcal{O}_L/\mathfrak{q}.$$

We also know $e' \leq e$ by the fact that $\mathfrak{p} \subseteq \mathfrak{p}'$ as well. In fact, equality holds:

$$|\mathrm{Gal}(L/K)| = [L : K] = efg$$

so $|G_\mathfrak{q}| = ef$ (for example by the orbit-stabilizer theorem: $G_\mathfrak{q}$ is the stabilizer of $\mathfrak{q}$, which has orbit of size g). But $G_\mathfrak{q}$ is the Galois group of L/K' so its size is also equal to $e'f' = [L : K']$. Hence, $e' = e$ and $f' = f$. In particular, prime decomposition of $\mathfrak{p}$ between K and K' is trivial (all the ramification degrees and inertial degrees will be 1). So we can assume that $K = K'$ (in particular proving the statement when $K = K'$ implies the general statement because of what we have just shown).

Let $\overline{x} \in \mathcal{O}_L/\mathfrak{q}$ be a primitive element for it over $\mathcal{O}_K/\mathfrak{p}$. Let $P \in \mathcal{O}_K[X]$ be the minimal polynomial of X. Then

$$P | \prod_{\sigma \in G_\mathfrak{q}} (X - \sigma(x))$$

We can take the reduction mod $\mathfrak{p}$, $\overline{P}(X) \in \mathcal{O}_K/\mathfrak{p}$. Any root of P in $\mathcal{O}_L$ is of the form $\sigma(x)$, so any root of $\overline{P}$ in $\mathcal{O}_L/\mathfrak{q}$ is of the form $\overline{\sigma(x)} = \theta(\sigma)(\overline{x})$.

Let $Q \in \mathcal{O}_K/\mathfrak{p}[X]$ be an irreducible factor of $\overline{P}$ such that $Q(\overline{x}) = 0$. The distict roots of Q all have the form $\theta(\sigma)(\overline{x})$. Hence

$$|\mathrm{im}\theta| \geq \deg Q \geq f.$$

But f was already an upper bound on the size of the image of θ, so the extension of residue fields is Galois (because this produces the maximaul number f of automorphisms) and θ is surjective. Its image has size f, and $|G_\mathfrak{q}| = ef$ so the kernel has order e as desired. □

Corollary 22.4

$|G_\mathfrak{q}| = ef$.

If $e = 1$, then

$$G_{\mathfrak{q}} \to \mathrm{Gal}(\kappa(\mathfrak{q})/\kappa(\mathfrak{p}))$$

is an isomorphism where the group on the right is generated by the *Frobenius element* $\mathrm{Frob}_{\mathfrak{q}} : x \mapsto x^{|\kappa((\mathfrak{p}))|}$. If $\mathfrak{p} = \prod \mathfrak{q}_i$, we know (writing $\mathfrak{q}_j = \sigma(\mathfrak{q}_i)$)

$$G_{\mathfrak{q}_i} = \sigma G_{\mathfrak{q}_j} \sigma^{-1}$$

so if G is Abelian then the Frobenius element depends only on $\mathfrak{p}$. Moreover, the Frobenius element (a.k.a. "Artin symbol") corresponding to a prime $\mathfrak{p}$ downstairs is well-defined up to a conjugacy, so it defines a conjugacy class of the Galois group. This allows us to state [but not prove in Math 129] an important result from the theory of prime decomposition in Galois extensions of number fields:

Theorem 22.5 (Chebotarev Density Theorem)

Let L/K be a Galois extension of number fields and $\mathcal{C}$ a conjugacy class in $\mathrm{Gal}(L/K)$. Then the set of unramified primes $\mathfrak{p}$ in $\mathcal{O}_K$ such that $\mathrm{Frob}_{\mathfrak{p}} = \mathcal{C}$ has density $|\mathcal{C}|/|\mathrm{Gal}(L/K)|$ in the set of all primes of $\mathcal{O}_K$.

The statement is easiest to prove where the word "density" replaced with "Dirichlet density", but it is still true of the natural density. With $\mathcal{C}$ replaced with $\{1\}$, the Chebotarev density theorem implies that the set of primes which split completely has density $1/|\mathrm{Gal}(L/K)|$ (in particular there are infinitely many).

§23 April 29, 2019

Last time, we had L/K a Galois extension of number fields and $\mathfrak{p}$ an unramified prime in K (recall all but finitely many have this property). When $\mathfrak{q}|\mathfrak{p}$, we defined the *decomposition group*

$$D_{\mathfrak{q}} := \{\sigma \in G : \sigma(\mathfrak{q}) = \mathfrak{q}\} \subseteq G$$

and we showed that the projection $D_{\mathfrak{q}} \to \mathrm{Gal}(\kappa(\mathfrak{q})/\kappa(\mathfrak{p}))$ is surjective with kernal of size equal to the ramification index. When $\mathfrak{p}$ is unramified, this is therefore an isomorphism of finite groups. Since $\kappa(\mathfrak{q})/\kappa(\mathfrak{p})$ is an extension of finite fields, its Galois group is cyclic and generated by the Frobenius automorphism $x \mapsto x^{|\kappa(\mathfrak{p})|} = x^{N(\mathfrak{p})}$. It follows that the decomposition group is cyclic and generated by a distinguished element $\mathrm{Frob}_{\mathfrak{q}}$ corresponding to that particular generator.

§23.1 The Frobenius element and quadratic reciprocity

Example 23.1

Let $L = \mathbf{Q}(\zeta_\ell)$ where ℓ is a rational prime. Then the Frobenius element only depends on p since the Galois group is Abelian. Moreover, the generator of $\mathrm{Gal}(\kappa(\mathfrak{q})/\kappa(p))$ is given by raising x to the p-th power. So $\mathrm{Gal}(L/\mathbf{Q})$ has a generator Frob_p which descends to the p-th power map on $\mathbf{F}_p[\zeta_\ell]$. So, it must be equal to

$$\mathrm{Frob}_p(\zeta_\ell) = \zeta_\ell^p.$$

Let $K \subseteq K_1 \subseteq L$ be number fields where L/K is Galois, $\mathfrak{q}|\mathfrak{p}_1|\mathfrak{p}$ primes of L, K_1, K, $\mathfrak{p}$ unramified in L. Let $H = \mathrm{Gal}(L/K_1) \subseteq G = \mathrm{Gal}(L/K)$.

By the theory, we have Frobenius elements $\mathrm{Frob}_{\mathfrak{q}}^H, \mathrm{Frob}_{\mathfrak{q}}$ in the decomposition groups $G_{\mathfrak{q}}$ and $H_{\mathfrak{q}}$. Recall that

$$\mathrm{Frob}_{\mathfrak{q}}(x) \equiv x^{N(\mathfrak{p})} \mod \mathfrak{q}$$

and

$$\mathrm{Frob}_{\mathfrak{q}}^H(x) \equiv x^{N(\mathfrak{p}_1)} \mod \mathfrak{q}$$

so mod $\mathfrak{q}$, one is a power of the other (in fact the Frobenius on the smaller extension is the restriction of the one on the big extension since it is unique). In particular, if $\mathfrak{p}$ and $\mathfrak{p}_1$ are the same then the Frobenius elements have the same action on $\kappa(\mathfrak{q})$.

Let L again be the cyclotomic field $\mathbf{Q}(\zeta_\ell)$ where ℓ is an odd prime Since $\mathrm{Gal}(L/\mathbf{Q})$ is the unit group of $\mathbf{Z}/\ell$, it is cyclic of order $\ell - 1$, and thus has a unique subgroup of index two. So, there is a unique surjective group homomorphism, the *Legendre symbol*

$$\left(\frac{\cdot}{\ell}\right) : \mathrm{Gal}(L/\mathbf{Q}) \to \{\pm 1\}.$$

The famous result about the Legendre symbol is that it should depend in a simple way on the residue of the input mod ℓ. This is characterized by Gauss's quadratic reciprocity law.

Theorem 23.2

If p, ℓ are odd primes, then

$$\left(\frac{p}{\ell}\right)\left(\frac{\ell}{p}\right) = (-1)^{\frac{p-1}{2}\frac{\ell-1}{2}}.$$

Moreover, $\left(\frac{2}{\ell}\right) = (-1)^{(\ell^2-1)/8}$ and $\left(\frac{-1}{\ell}\right) = (-1)^{(\ell-1)/2}$.

Proof. This proof relies on the existence of Frobenius elements, but we'll do the easier parts first.

$\left(\frac{-1}{\ell}\right) = 1$ if and only if -1 is a square in $(\mathbf{Z}/\ell)^*$. But $(\mathbf{Z}/\ell)^*$ is cyclic of even order, so there is a generator σ of even order $\ell - 1$. As a result, we know that $\sigma^{(\ell-1)/2}$ is -1 (as it is not 1 but squares to 1), so it is a perfect square, which means it must be an even power of the generator. Hence, $(\ell-1)/2$ is even so ℓ is 1 mod 4.

Let H be the kernel of the Legendre symbol.

Lemma 23.3

For $p \neq \ell$, p a possibly even prime, $\mathrm{Frob}_p \in H$ if and only if p splits completely in $K = \mathbf{Q}(\zeta_\ell)^H = \mathbf{Q}(\sqrt{\pm\ell})$.

Proof. Let $\mathfrak{p}_1 \subseteq \mathcal{O}_K$ be a prime factor of $\mathfrak{p}$. $\mathrm{Frob}_\mathfrak{p}$ lives in $\mathrm{Gal}(L/K)$ and generates the decomposition subgroup $G_\mathfrak{q} \cong \mathrm{Gal}(\kappa(\mathfrak{q})/\kappa(\mathfrak{p}))$. In fact, $\mathrm{Frob}_\mathfrak{p} \in H$ if and only if

$$\mathrm{Frob}_\mathfrak{p} \in H \cap G_\mathfrak{q} = H_\mathfrak{q} \cong \mathrm{Gal}(\kappa(\mathfrak{q})/\kappa(\mathfrak{p}_1)).$$

Since the Frobenius generates the bigger decomposition group, it must generate the subgroup $H_\mathfrak{q}$ if it is in it, so this is equivalent to $\kappa(\mathfrak{p}_1) = \kappa(\mathfrak{p})$.

Moreover, $\kappa(\mathfrak{p}_1) = \kappa(\mathfrak{p})$ is equivalent to $f(\mathfrak{p}_1|\mathfrak{p}) = 1$. But the extension is Galois, so this is equivalent to the inertial degree of any prime in $\mathcal{O}_K$ over $\mathfrak{p}$ being 1, and since $\mathfrak{p}$ is unramified, this is equivalent to $\mathfrak{p}$ splitting completely. □

We know that $\mathrm{Frob}_p = [p] \in (\mathbf{Z}/\ell)^*$. p splitting in K is equivalent to $(p/\ell) = 1$ by the lemma. We also know that p splits in K if and only if $\mathcal{O}_K/p\cdot\mathcal{O}_K$ is an integral domain. And $\mathcal{O}_K/p\cdot\mathcal{O}_K = \mathbf{F}_p[X]/(\text{minimal poly. of } \alpha)$ where α is the generator of $\mathcal{O}_K$. This is equivalent to $\pm\ell$ being a quadratic residue mod p [important exercise: show that this is still true even when $\pm\ell$ is 1 mod 4], i.e. $\left(\frac{\pm\ell}{p}\right) = 1$. So,

$$\left(\frac{p}{\ell}\right)\left(\frac{\pm\ell}{p}\right) = 1.$$

Since $\pm\ell = (-1)^{(\ell-1)/2}\ell$ and because of the law for $(-1/p)$, this is the same as the desired reciprocity law.

Finally, we want to compute $(2/\ell)$. From the same arguments as before, $(2/p) = 1$ if and only if 2 splits in K. This is equivalent to $X^2 - X + \frac{1-(\pm\ell)}{4}$ being reducible mod 2. This is equivalent to $(1-(\pm\ell))/4$ being even, i.e. ℓ being ± 1 mod 8. □

This allows you to compute Legendre symbols with great ease and satisfaction.

Example 23.4

We can compute

$$\left(\frac{34}{101}\right) = \left(\frac{2}{101}\right)\left(\frac{17}{101}\right) = -\left(\frac{101}{17}\right) = -\left(\frac{-1}{101}\right) = -1.$$

Proposition 23.5

An odd prime p is a sum of two squares if and only if $p \equiv 1 \mod 4$.

Proof. If p is 1 mod 4 then -1 is a square mod p and p splits in $\mathcal{O}_K$. So

$$p \cdot \mathcal{O}_K = \mathfrak{p}_1\mathfrak{p}_2$$

and since the class group of $K = \mathbf{Q}(i)$ is trivial, we know these ideals are principal with norms multiplying to p^2, which means we have $x, y \in \mathbf{Z}$ such that $x^2 + y^2 = p$ as desired. □

§24 May 1, 2019

Today is the last class. We will introduce a new topic as an advertisement for the full power of algebraic number theory to extend many of our results which hold for a finite extension $K/\mathbf{Q}$ to an arbitrary extension of number fields. This is reminiscent of a quote from **Star Wars Episode VI: Return of the Jedi**:

- "Now witness the firepower of this fully armed and operational battle station" – Emperor Palpatine

(the quote is suggested by me and not Professor Kisin). The power we speak of is provided by the theory of **local fields**.

§24.1 Valuation Theory

Let K be a number field and $\mathfrak{p} \subseteq \mathcal{O}_K$ be a prime. One useful object is the *localization* of $\mathcal{O}_K$ at $\mathfrak{p}$, which yields a simpler object called a *discrete valuation ring* and allows us to get rid of the information that has to do with the primes not equal to $\mathfrak{p}$.

Definition 24.1. The **localization** of $\mathcal{O}_K$ at $\mathfrak{p}$ is

$$\mathcal{O}_{K,\mathfrak{p}} := \{ab^{-1} : a \in \mathcal{O}_K, b \in \mathcal{O}_K, b \notin \mathfrak{p}\}.$$

From a prime $\mathfrak{p}$ we also get the $\mathfrak{p}$-*adic valuation*, which is given by the following.

Definition 24.2. The **$\mathfrak{p}$-adic valuation** is the map $v_\mathfrak{p} : K \to \mathbf{Z} \cup \{\infty\}$ given by $v_\mathfrak{p}(x) = \max\{i : x \in \mathfrak{p}^i\}$.

Associated with any valuation on a field, we have a *valuation ring*, which in this case is

$$\mathcal{O}_{(v)} := \{x \in K : v(x) \geq 0\},$$

which in turns contains the "unique maximal ideal"

$$\mathfrak{m}_{(v)} := \{x \in K : v(x) > 0\}.$$

Proposition 24.3 (i) The localization of $\mathcal{O}_K$ at $\mathfrak{p}$ is the same as the valuation ring of K under $v_\mathfrak{p}$.

(ii) This is a local ring with unique maximal ideal $\mathfrak{m}_{(v)}$.

(iii) This is a discrete valuation ring. In particular, there is a *uniformizing parameter* $\pi \in \mathcal{O}_{K,\mathfrak{p}}$ such that every ideal of $\mathcal{O}_{K,\mathfrak{p}}$ is generated by some power of π.

Proof. By unique factorization, $\mathfrak{p} - \mathfrak{p}^2$ is nonempty, so there exists a $\pi \in \mathcal{O}_K$ with $v_\mathfrak{p}(\pi)$. For any $ab^{-1} \in \mathcal{O}_{K,\mathfrak{p}}$ we have

$$v_\mathfrak{p}(ab^{-1}) = v_\mathfrak{p}(a) + v_\mathfrak{p}(b^{-1})$$

but $b \notin \mathfrak{p}$ so the second term is zero which means that we have the inclusion

$$\mathcal{O}_{K,\mathfrak{p}} \subseteq \mathcal{O}_{(v)}.$$

On the other hand, if $x \in \mathcal{O}_{(v)}$, then since $x \in K$ we know $x = ab^{-1}$ for some $a, b \in \mathcal{O}_K$. Its valuation is nonnegative by the definition of the valuation ring, so $v_\mathfrak{p}(a) \geq v_\mathfrak{p}(b)$. Multiplying both numerator and denominator by $\pi^{-v_\mathfrak{p}(b)}$, we get

$$x = \left(a \cdot \pi^{-v_\mathfrak{p}(b)}\right)\left(b \cdot \pi^{-v_\mathfrak{p}(b)}\right)^{-1}.$$

The problem is now that the numerator and denominator might not be in $\mathcal{O}_K$ (their $\mathfrak{p}$-adic valuations are nonnegative but the valuations at some other prime might be negative). To deal with this, we just multiply both sides by elements of the appropriate sets $\mathfrak{p}_i^m - \mathfrak{p}_i^{m+1}$, and we have the opposite inclusion. Hence, we may use $\mathcal{O}_{K,\mathfrak{p}}$ and $\mathcal{O}_{(v)}$ to mean the same object.

Let $x \in \mathcal{O}_{(v)}$ be nonzero and let $i = v_\mathfrak{p}(x)$, so that $x\pi^{-i}$ has valuation zero and is therefore a unit in in $\mathcal{O}_{(v)}$. Therefore, every element of $\mathcal{O}_{(v)}$ can be written (uniquely since i determines its valuation) in the form $u \cdot \pi^i$ where u is a unit.

If I is an ideal in $\mathcal{O}_{(v)}$, then we can let i be the minimum of the valuations amongst elements of I. Choose $x \in I$ such that $v_\mathfrak{p}(x) = i$. Then every element of I is a constant multiplicative factor away from x, where this factor has nonnegative valuation. This is equivalent to saying that I is principally generated by x, so we are done. □

To take a *completion* of the ring of integers, one way to do it is using the algebraic definition of the completion:

Definition 24.4. The **completion** of the local ring $\mathcal{O}_{(v)}$ is

$$\widehat{\mathcal{O}}_{(v)} := \varprojlim \mathcal{O}_{(v)}/\mathfrak{m}_{(v)}^i.$$

The "inverse limit" here is the subset of $\prod \mathfrak{m}_{(v)}$ consisting of all sequences $(x_0, \ldots)$ such that x_i always maps to x_{i-1} via the canonical projection. Another canonical example of this is the inverse limit of $F[x]/x^i$, which is the ring of formal power series over F.

The valuation we wrote down is a *nonarchimedean valuation*, meaning it satisfies the following formal properties.

Definition 24.5. The map $v : K \to \mathbf{Z} \cup \{\infty\}$ is a **nonarchimedean valuation** if it satisfies

- $v(x) = \infty$ if and only if $x = 0$.
- $v(xy) = v(x) + v(y)$
- $v(x + y) \geq \min(v(x), v(y))$

From such a valuation, we can obtain an absolute value on K.

Definition 24.6. A nonarchimedian valuation v on K induces a **nonarchimedean absolute value** $|\cdot|_v$ given by $|x|_v = q^{-v(x)}$ for some fixed real constant $q > 1$.

By the properties of a nonarchimedean valuation, a nonarchimedean absolute value gives a metric on K, with respect to which we might take a completion as a metric space. In particular, we denote this completion by K_v, which is endowed with the structure of a complete topological field. It's also true that K_v is not only the field of fractions of $\widehat{O}_{(v)}$, but it is generated as a ring by $\widehat{O}_{(v)}$ and $1/\pi$.

Let L/K be an extension of number fields, and $\mathfrak{q}|\mathfrak{p}$ a prime upstairs lying over a prime downstairs. A useful operation is to take completions with respect to these two primes, and consider the extension of rings

$$\widehat{\mathcal{O}}_{L,\mathfrak{q}}/\widehat{\mathcal{O}}_{K,\mathfrak{p}}.$$

Suppose $\mathfrak{p}$ factors in L as

$$\mathfrak{p} \cdot \mathcal{O}_L = \prod \mathfrak{q}_i^{e_i}.$$

We may consider the inverse limit

$$\begin{aligned}\left(\widehat{\mathcal{O}}_L\right)_{\mathfrak{p}} &:= \varprojlim \mathcal{O}_{L,\mathfrak{p}}/\mathfrak{p}^i\mathcal{O}_{L,\mathfrak{p}} \\ &= \varprojlim \mathcal{O}_{L,\mathfrak{p}}/\prod \mathfrak{q}_j^{ie_j} \\ &= \prod \varprojlim \mathcal{O}_{L,\mathfrak{p}}/(\mathfrak{q}_j^{e_j})^i\end{aligned}$$

which is indeed just the product of the completions of $\mathcal{O}_L$ at each $\mathfrak{q}_j$. Since $(\mathcal{O}_L)_{\mathfrak{p}}$ is a finitely-generated $\mathcal{O}_{K,\mathfrak{p}}$-module and the bottom one is a DVR and thus a PID, in fact we have a finitely-generated free $\mathcal{O}_{K,\mathfrak{p}}$. The same is true of their completions. Hence, $L_{\mathfrak{q}}/K_{\mathfrak{p}}$ is a finite extension of complete fields.

Proposition 24.7

If L/K is Galois, then

$$G_{\mathfrak{q}} \cong \mathrm{Gal}(L_{\mathfrak{q}}/K_{\mathfrak{p}}).$$

Proof. If $\sigma \in G$, then it extends to the completion if and only if it is continuous with respect to the $\mathfrak{q}$-adic metric. This is equivalent to sending $\mathfrak{q}$ to itself, so indeed a σ extends to the completion if and only if it is in the decomposition group. This means we have a well-defined injection

$$G_{\mathfrak{q}} \to \mathrm{Gal}(L_{\mathfrak{q}}/K_{\mathfrak{p}})$$

given by taking the extension of the automorphism. Since these are finite sets, it now suffices to show that $[L_{\mathfrak{q}} : K_{\mathfrak{p}}] = e_{\mathfrak{q}} f_{\mathfrak{q}}$. For this you just show it for the rank of the ring of integers (in particular the ring of integers of $K_{\mathfrak{p}}$ is a PID so the upstairs ring is a finitely-generated free module over that ring). □