The ring of integers of $\mathbb{Q}(\zeta_p)$, and ideal theory

December 11, 2013

There are two things on the docket today. First we'll use the ideas from last time to show that the ring of integers of $\mathbb{Q}(\zeta_p)$ is $\mathbb{Z}[\zeta_p]$. Then we'll review, without proofs, some aspects of the theory of ideals in rings of integers.

Let $\mathcal{O}$ denote the ring of integers of $\mathbb{Q}(\zeta_p)$, and let V denote its Minkowski space. Certainly, $\mathbb{Z}[\zeta_p] \subset \mathcal{O}$. The proof that this containment is an equality will be in three steps:

1. First, we'll compute that $Vol(V/\mathbb{Z}[\zeta_p])^2 = p^{p-2}$.
2. Second, we'll show that the inclusion induces an isomorphism $\mathbb{Z}[\zeta_p]/p \xrightarrow{\sim} \mathcal{O}/p$.
3. Third, we'll put these together to deduce that $\mathbb{Z}[\zeta_p] = \mathcal{O}$.

The funny thing is that, while steps 1 and 2 give complementary information, the crucial input to both of them is the same fact: our old friend

$$p = (1-\zeta_p)(1-\zeta_p^2)\dots(1-\zeta_p^{p-1}).$$

Let's start with step 1. Recall from last time that such volumes can be computed as a determinant. Namely, if you have a number field F of degree n, and a $\mathbb{Z}$-basis $b_1, \dots, b_n$ for some abelian subgroup $L \subset F$ of rank n, then the volume $Vol(V/L)$ is equal to

$$|det(\sigma b_i)|,$$

where (σb_i) stands for the $n \times n$ matrix whose rows are indexed by the b_i and whose columns are indexed by the complex embeddings $\sigma : F \to \mathbb{C}$.

Now let's specialize to the case where $L = \mathbb{Z}[\alpha]$ for some algebraic integer α which generates L. Then our matrix is

$$(\sigma\alpha^i)$$

where i ranges from 0 to $n-1$. This is a particular kind of *Vandermonde matrix*: an $n \times n$ matrix whose n^{th} row is of the form $1, x_j, x_j^2, \dots x_j^{n-1}$, as j ranges from 0 to $n-1$ (the j corresponds to our σ), and x_j is just some variable. Let me recall that there's a nice formula for the determinant of a Vandermonde matrix: up to a sign, it is equal to

$$\prod_{j<j'} (x_{j'} - x_j).$$

I won't prove this, but let me sketch the idea. Both the Vandermonde determinant and this product are polynomials in the x_j of the same degree. On the other hand, since a determinant vanishes whenever two rows are equal, the Vandermonde determinant vanishes whenever $x_j = x_{j'}$ for some j, j', i.e. whenever $(x_{j'} - x_j)$ vanishes. From this one can argue that, up to a unit, the product of those linear forms divides the Vandermonde determinant. Then, since they have the same degree, these polynomials must be equal (up to a unit, which if we work in $\mathbb{Z}$, is just a sign.)

Now, we're actually interested in the volume squared. We can square this formula in a nice way, again up to sign: we find that the square of the Vandermonde determinant is, up to sign,

$$\prod_{j \neq j'} (x_{j'} - x_j).$$

Now let's specialize back to our case, where j corresponds to σ, and $x_j = \sigma\alpha$. We find

$$Vol(V/\mathbb{Z}[\alpha])^2 = \prod_{\sigma \neq \sigma'} (\sigma'\alpha - \sigma\alpha).$$

Now let's add the further assumption that $F/\mathbb{Q}$ is Galois. This means that all the complex embeddings $\sigma : F \to \mathbb{C}$ actually land back in F, and are just another way of describing the Galois group. Then, since multiplication by a fixed group element just permutes the group, we deduce

$$Vol(V/\mathbb{Z}[\alpha])^2 = \prod_{\sigma} \sigma \left(\prod_{\sigma' \neq 1} (\sigma'\alpha - \alpha) \right).$$

In other words, $Vol(V/\mathbb{Z}[\alpha])^2$ is just the product of all the conjugates of the number $\prod_{\sigma' \neq 1}(\sigma'\alpha - \alpha)$.

OK, let's move back to our case of interest, with $F = \mathbb{Q}(\zeta_p)$ and $\alpha = \zeta_p$. Then this product is

$$(\zeta_p^2 - \zeta_p)(\zeta_p^3 - \zeta_p) \dots (\zeta_p^{p-1} - \zeta_p).$$

Pulling out a ζ_p from each factor, again up to signs this is

$$\zeta_p^{p-2}(1 - \zeta_p)(1 - \zeta_p^2) \dots (1 - \zeta_p^{p-2}).$$

That's almost our expression for p; it's just missing a factor of $1 - \zeta_p^{p-1}$. So we can rewrite it as

$$\zeta_p^{p-2} \cdot \frac{p}{(1 - \zeta_p^{p-1})}.$$

Now, to get the volume squared, we're supposed to just take the product over all the conjugates of this number. We can do this on each of the three factors separately, then put them together.

The first factor would be easy to compute directly, but we can also do it abstractly as follows: certainly the product of the conjugates of ζ_p^{p-2} is some power of ζ_p, because Galois always sends ζ_p to a power of ζ_p. Then again this product of conjugates is invariant under Galois. But the only power of ζ_p which is invariant under Galois is $\zeta_p^0 = 1$. So the first factor gives 1.

The second factor p is fixed by Galois, so it gives p^{p-1}, since there are $p - 1$ Galois conjugates.

As for the third factor $(1-\zeta_p^{p-1})$, its conjugates are the $(1-\zeta_p^i)$ for $1 \leq i \leq p-1$, so when we take the product we just get p again, by the same old identity!

The conclusion is that

$$Vol(V/\mathbb{Z}[\zeta_p])^2 = p^{p-1}/p = p^{p-2},$$

as claimed. The whole time we worked up to a sign, but of course the volume squared is positive, so this sign, in the end, must be correct.

So that was the first step. Now let's do the second step, meaning let's show that the inclusion $\mathbb{Z}[\zeta_p] \to \mathcal{O}$ induces an isomorphism $\mathbb{Z}[\zeta_p]/p \simeq \mathcal{O}/p$.

Let's start by making some orienting remarks, which may or may not be helpful for the proof. First of all, both $\mathbb{Z}[\zeta_p]/p$ and $\mathcal{O}/p$ are finite abelian groups of the same size. In fact, both are isomorphic to $(\mathbb{Z}/p\mathbb{Z})^{p-1}$. This is simply because both $\mathbb{Z}[\zeta_p]$ and $\mathcal{O}$ are isomorphic to $\mathbb{Z}^{p-1}$. In particular, to see that $\mathbb{Z}[\zeta_p]/p \to \mathcal{O}/p$ is an isomorphism, it suffices to see either that its injective or that it's surjective.

Now, the idea behind the proof is that we somehow already know that $\mathbb{Z}[\zeta_p]$ is "big enough" from the perspective of p. This is reflected in the fact that the factorization $p = (1-\zeta_p)\dots(1-\zeta_p^{p-1})$, which is as long as it could be (length $p-1 = deg(\mathbb{Q}(\zeta_p))$), already lives in the ring $\mathbb{Z}[\zeta_p]$; no need to use other algebraic integers. To give an example in the other direction, the fact that $\mathbb{Z}[2i]$ is not the ring of integers of $\mathbb{Q}(i)$ is related to the fact that it's "missing" the factorization $2 = (1+i)(1-i)$.

OK, now let's actually prove that $\mathbb{Z}[\zeta_p]/p \to \mathcal{O}/p$ is an isomorphism. Recall from the first proof of the irreducibility of the cyclotomic polynomials that the relation $p = u\pi^{p-1}$ with u a unit in $\mathbb{Z}[\zeta_p]$ implies that the quotient $\mathbb{Z}[\zeta_p]/p$ can approached by taking $p-1$ more drastic quotients, namely the $\mathbb{Z}[\zeta_p]/\pi^k$ for $k = 0, \dots, p-2$. Furthermore, by the "third isomorphism theorem" the difference between each of these quotients is measured just by the single abelian group $\mathbb{Z}[\zeta_p]/\pi$. Thus $\mathbb{Z}[\zeta_p]/p$ is built up out of $p-1$ copies of $\mathbb{Z}[\zeta_p]/\pi$.

Now, since the factorization $p = u\pi^{p-1}$ also holds in $\mathcal{O}$ (obviously, since it holds in $\mathbb{Z}[\zeta_p]$), we get the same story with $\mathcal{O}$. And the inclusion $\mathbb{Z}[\zeta_p] \to \mathcal{O}$ respects these stories, i.e. it also induces a similar map on each quotient $\mathbb{Z}[\zeta_p]/\pi^k \to \mathcal{O}/\pi^k$.

The upshot is that to see that $\mathbb{Z}[\zeta_p]/p \to \mathcal{O}/p$ is an isomorphism, it suffices to see that

$$\mathbb{Z}[\zeta_p]/\pi \to \mathcal{O}/\pi$$

is an isomorphism. But again, since $p-1$ copies of each of these groups makes $(\mathbb{Z}/p\mathbb{Z})^{p-1}$, both of these groups have size p. Thus to see that this map is an isomorphism, it's enough to see that it's nonzero. But, this map obviously sends 1 to 1, and 1 is nonzero in both groups for the same reason: if it weren't, that would mean that π is a unit, so the groups would be trivial, whereas we know they have size p.

So we have proved that $\mathbb{Z}[\zeta_p]/p \simeq \mathcal{O}/p$, and hence we've accomplished step 2.

Now, step 3 is a formal argument. We can phrase it as follows:

Proposition 0.1. *Let F be a number field of degree d with ring of integers $\mathcal{O}$ and Minkowski space V, and let $R \subset \mathcal{O}$ be a subring which is isomorphic to $\mathbb{Z}^d$ as an abelian group. Suppose that for every*

prime p which divides $Vol^2(V/R)$, the map

$$R/p \to \mathcal{O}/p$$

induced by the inclusion is an isomorphism. Then $R = \mathcal{O}$.

The statement makes sense: recall from last time that $Vol^2(V/R)$ is necessarily an integer, since R is composed of algebraic integers.

Proof. Let's first see what $R/p \to \mathcal{O}/p$ being an isomorphism entails. Actually, let's just look at injectivity. It says that if $x \in R$ is a multiple of p in $\mathcal{O}$, then x is a multiple of p in R. In other words, letting $y = x/p$, it says that if $y \in \mathcal{O}$ with $py \in R$, then $y \in R$. Inductively, we deduce the following: if n is divisible only by primes p for which we know that $R/p \to \mathcal{O}/p$ is an isomorphism, then we have the implication

$$y \in \mathcal{O}, ny \in R \Rightarrow y \in R.$$

(We could have also directly argued that $R/n \to \mathcal{O}/n$ is an isomorphism, just as above we went from π to p.)

Now let's put that in our pocket, and see what knowing $Vol^2(V/R)$ gets us. Recall from last time that there is the relation

$$Vol^2(V/R) = [\mathcal{O} : R]^2 \cdot Vol^2(V/\mathcal{O}),$$

and that all these quantities are integers. Now, every $y \in \mathcal{O}$ satisfies $[\mathcal{O} : R]y \in R$. This is because the quotient group $\mathcal{O}/R$ has size $[\mathcal{O} : R]$, and hence each of its elements is killed by multiplication by $[\mathcal{O} : R]$. Since the above relation shows that $[\mathcal{O} : R]$ divides $Vol^2(V/R)$, it follows that every $y \in \mathcal{O}$ satisfies $Vol^2(V/R) \cdot y \in R$. Now if we check in our pocket, we find that the proof is complete. □

So we've proven that $\mathcal{O}_{\mathbb{Q}(\zeta_p)} = \mathbb{Z}[\zeta_p]$. We can summarize the proof as follows. Our volume computation shows that $\mathbb{Z}[\zeta_p]$ is pretty big (meaning, the volume of $V/\mathbb{Z}[\zeta_p]$ is small). Actually, it's as big as it could be from the perspective of every prime except possibly p. But then the long factorization $p = (1 - \zeta_p) \dots (1 - \zeta_p^{p-1})$ shows that $\mathbb{Z}[\zeta_p]$ is as big as it can be from the perspective of p as well.

Now let's move on to the second half of this lecture, about ideal theory.

The starting point is the "fundamental theorem of arithmetic": every integer can be written as a product of prime numbers, and such a product is unique up to reordering and multiplying by units. This really is a fundamental fact about arithmetic, but it's kind of annoying to state precisely. It's definitely not ideal.

Now, it turns out that the straightforward analog of this statement is false for the ring of integers $\mathcal{O}$ of a general number field F.

That's actually a subtle claim, because $\mathcal{O}$ was really trying its hardest to have this property. What I mean is that if we hadn't taken the full ring of integers, if we'd taken a subring $R \subset \mathcal{O}$ with $R \neq \mathcal{O}$, then unique prime factorization wouldn't stand a chance. One reason is the following. It's an easy consequence of unique prime factorization (in an integral domain) that, given a monic polynomial with coefficients in the domain, any root in the fraction field actually has to lie in the domain. (This gives one proof that $\mathcal{O}_{\mathbb{Q}} = \mathbb{Z}$.) So if R satisfies unique prime factorization, it must contain all the algebraic integers in $Frac(R) = F$, so it must be $\mathcal{O}$!

But nonetheless, even $\mathcal{O}$ itself need not satisfy unique prime factorization. For example, let $\mathcal{O} = \mathbb{Z}[\sqrt{-5}]$. Consider

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

I claim this is a counterexample to unique prime factorization in $\mathbb{Z}[\sqrt{-5}]$. To see this one needs to verify a bunch of things: first, that each of these factors is prime (meaning, up to units it has exactly two divisors: 1 and itself); and second, that neither of $2, 3$ is a unit multiple of $1 + \sqrt{-5}$ or $1 - \sqrt{-5}$.

This can be done by hand; after all, elements of this ring are just expressions $a+b\sqrt{-5}$ with $a, b \in \mathbb{Z}$. But it's easier to use a nice tool, the *norm*. Here's a proposition that we have all the tools to prove, thanks to your last problem set.

Proposition 0.2. *Let F be a number field with ring of integers $\mathcal{O}$ and Minkowski space V. For $\alpha \in \mathcal{O}$, the following numbers are all integers, and are all equal:*

1. *The determinant of multiplication by α acting on $\mathcal{O} \simeq \mathbb{Z}^d$;*
2. *The product $\prod_{\sigma: F \to \mathbb{C}} \sigma\alpha$;*
3. *Up to a sign, the factor by which multiplication by α scales volumes in V;*
4. *Up to a sign, the number of elements of $\mathcal{O}/\alpha$.*

The first two options define the norm of α, denoted $N(\alpha)$. Thus the last two are $|N(\alpha)|$. The important property of the norm, for present purposes, is that it's multiplicative:

$$N(\alpha\beta) = N(\alpha)N(\beta).$$

That lets us use multiplication in $\mathbb{Z}$ to check facts about multiplication in rings of integers.

For example, $\alpha \in \mathcal{O}$ is a unit if and only if $N(\alpha) = \pm 1$. The forward direction follows from the fact that if α is a unit, then so is $N(\alpha)$, by multiplicativity. But the only units in $\mathbb{Z}$ are ± 1. The backward direction follows from the fact that $N(\alpha)$ is the product of α with all its nontrivial conjugates; so if $N(\alpha) = \pm 1$, then up to a sign the product of the nontrivial conjugates of α gives an inverse to α which also lies in $\mathcal{O}$. (Alternate proof: $N(\alpha) = \pm 1$ is equivalent to $\mathcal{O}/\alpha$ being trivial, which is equivalent to α being a unit.)

Using this, we can easily check that, e.g., 2 is prime in $\mathbb{Z}[\sqrt{-5}]$. Namely, we can calculate

$$N(a + b\sqrt{-5}) = a^2 + 5b^2.$$

If $2 = \alpha\beta$, then applying N we get

$$4 = N(\alpha)N(\beta).$$

Thus either $N(\alpha)$ or $N(\beta)$ is 1, which wouldn't contradict 2 being prime, or $N(\alpha) = N(\beta) = 2$. But you can just see that 2 can't be written in the form $a^2 + 5b^2$. For unless $b = 0$, and $a = -1, 0$, or 1, this quantity is too big. And none of those possibilities give 2.

Likewise one can do all the work required to see that $6 = 2 \cdot 3 = (1 - \sqrt{-5})(1 + \sqrt{5})$ is a non-unique prime factorization.

So, unique prime factorization fails for general $\mathcal{O}$! But there's a salvage: instead of asking for unique prime factorization for elemments, ask for unique prime factorization of *ideals*.

Recall that an ideal in a ring R is just an abelian subgroup $I \subset R$ which is closed under multiplication by elements in R. That is, if $x \in I$ and $r \in R$, then $rx \in I$. Note, this is a very different concept from a subring. The first example to have in mind is the *ideal generated by an element* $\alpha \in R$. This is just the set of multiples of α:

$$(\alpha) = \alpha R = \{\alpha \cdot r \mid r \in R\}.$$

Another name for this is the *principal* ideal generated by α. So we're making a change of perspective here: instead of thinking about elements, we're thinking about ideals. Which means, instead of thinking about an element α, we think about *the collection of all elements which* α *divides*. This loses some information about α, but it turns out that it was information we wanted to forget anyway. Namely:

Proposition 0.3. *Let α and β be two elements of an integral domain R (e.g. $R = \mathcal{O}$, a ring of integers in a number field.) Then $(\alpha) = (\beta)$ if and only if α is a unit multiple of β.*

This is just the old claim that α is a unit multiple of β if and only if $\alpha|\beta$ and $\beta|\alpha$, which we already encountered in the first lecture.

Besides this fact that the ideal generated by α doesn't care if we change α by a unit, the other thing to know is that, in general, not every ideal in $\mathcal{O}$ can be generated by a single element, i.e. not every ideal is *principal*. This is what leads to the fix to the failure of unique prime factorization. For example, in $\mathcal{O} = \mathbb{Z}[\sqrt{-5}]$, we have a completely unique prime *ideal* factorization

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

Here the notation (α, β) means the ideal generated by both α and β, i.e. the set of $\mathcal{O}$-linear combinations of α and β. As for the meaning of product of ideals, it is as in the following definition:

Definition 0.4. *Let I and J be two ideals in a ring R. We define the product $I \cdot J$ to be the ideal generated by the elements xy as x ranges over I and y ranges over J.*

For example, it's easy to see that $(\alpha) \cdot (\beta) = (\alpha\beta)$, so on the principal ideals this product recovers the usual one.

Now let us state the salvage to the failure of unique prime factorization in $\mathcal{O}$. We just need one more definition, which is the analog of "prime number" for ideals. Recall that a prime number can be defined as a number which has exactly two divisors, up to units: namely, itself and 1. In ideal language, the analogous definition is the following:

Definition 0.5. *An ideal I in a ring R is called* maximal *if there are exactly two ideals J with $I \subset J$: namely, $J = I$, and $J = R$.*

So, here is the main statement, the *unique prime factorization of ideals*:

Theorem 0.6. *Let $\mathcal{O}$ be the ring of integers in a number field. Then every nonzero ideal I of $\mathcal{O}$ can be written as a product of maximal ideals of $\mathcal{O}$, and such a product description is unique up to reordering.*

We remark that nonzero ideals look completely different from the zero ideal. For example, a nonzero ideal ideal is always free of rank $d = deg(F/\mathbb{Q})$ as an abelian group. This is because if $\alpha \in I$ is nonzero, then I also contains the integer $N = N(\alpha) = \prod_\sigma \sigma(\alpha)$, which implies that $N\mathcal{O} \subset I \subset \mathcal{O}$, so I is of finite index in $\mathcal{O}$.

There is also a companion proposition to the above theorem, which is about "on the same level" (meaning as difficult to prove, and extends to the same level of generality, namely Dedekind domains):

Proposition 0.7. *Let I and J be nonzero ideals of $\mathcal{O}$, and suppose $I \subset J$. Then J divides I, meaning there exists an ideal K such that $J \cdot K = I$.*

To make this make sense in your head, think of the case of principal ideals.

Anyway, this unique prime *ideal* factorization leads to a change of perspective on the issue of uniqueness of prime *element* factorizations, which is still of arithmetic significance. Namely, by this theorem, we see that unique prime factorization of elements in $\mathcal{O}$ is equivalent to the property that every ideal in $\mathcal{O}$ is in fact principal.

So we can change our fundamental question. Instead of asking if $\mathcal{O}$ has prime element factorization, which is sort of a complicated question, we can just ask if every ideal of $\mathcal{O}$ is principal.

Of course, the answer is "no" in general, as we saw with $\mathbb{Z}[\sqrt{-5}]$. But this change in the question lets us also investigate a natural measure of "how bad" the "no" answer is, so how far every ideal is from being principal.

Namely, we can define an abelian group Cl_F, the *class group* of our number field, which measures the failure of the principal ideal property in $\mathcal{O}$. One precise claim is that Cl_F is the trivial group if and only if every ideal of $\mathcal{O}$ is principal. To define Cl_F, we start with the following remark, which defines the notion of "isomorphism of ideals" in $\mathcal{O}$:

Proposition 0.8. *Let I and J be nonzero ideals of $\mathcal{O}$. Then the following are equivalent:*

1. *There exists an isomorphism of abelian groups $\varphi : I \xrightarrow{\sim} J$ which respects the $\mathcal{O}$-multiplication, i.e. $\varphi(rx) = r\varphi(x)$ for all $r \in \mathcal{O}$ and $x \in I$;*
2. *There exists a nonzero $\alpha \in F$ such that $\alpha \cdot I = J$.*

The claim is basically just that any abstract isomorphism as in 1 actually has to be given by multiplication by some nonzero $\alpha \in F$. Note that α really has to be in F here, not in $\mathcal{O}$. For example, the inverse to multiplication by α will always be multiplication by α^{-1}.

Another example is that any two principal ideals are isomorphic. This is clear: take α to be the quotient of the generators. More generally, I and J are isomorphic if and only if there exist principlal ideals (α) and (β) (where here $\alpha, \beta \in \mathcal{O}$) such that $(\alpha)I = (\beta)J$. So the notion of isomorphism of ideals can also be thought of as a way of just neglecting the principal ideals: pretending they're not there for the purposes of multiplication.

Now we can define Cl_F.

Definition 0.9. *Let F be a number field, with ring of integers $\mathcal{O}$. Define the abelian group Cl_F to be the set of isomorphism classes of nonzero ideals of $\mathcal{O}$, under the group operation induced by the product of ideals.*

To see that this is a reasonable definition, we need to know that if I and I' are isomorphic and J is an ideal, then IJ is isomorphic to $I'J$. This is easy to verify. We also need to see that every nonzero ideal has an inverse in this group structure. That follows because if $\alpha \in I$ is nonzero, then $(\alpha) \subset I$, so by the proposition above I divides (α). But (α) is a principal ideal, so it's isomorphic to the "trivial" ideal $\mathcal{O} = (1)$, which is the unit under multiplication. So I does have an inverse in Cl_F.

Now here is the second big theorem, which, in some sense, "bounds" the failure of every ideal in $\mathcal{O}$ being principal, or of unique prime factorization in $\mathcal{O}$:

Theorem 0.10. *The class group Cl_F of a number field is finite.*

We can explain why this result is true, without however going into the details of the proof.

The immediate difficulty in coming to grips with the above theorem is that it may not be entirely clear how to picture what it's saying. I mean, what does the collection of all nonzero ideals look like? Surely it's some infinite set. Then how is it that when we factor out by isomorphisms, it collapses to a finite set?

The first important remark is that while the set of all nonzero ideals is infinite, it is in some sense "tamed" by the norm function. Namely, define the *norm* of a nonzero ideal I to be $N(I) = \#(\mathcal{O}/I)$. Note that this makes sense, since, as we saw above, I is of finite index in $\mathcal{O}$. The name is motivated by the fact that if I is principal, generated by α, then $N(I) = |N(\alpha)|$. This was one of our above definitions of $N(\alpha)$. Then:

Proposition 0.11. *Let $N \geq 0$. Then the set of nonzero ideals in $\mathcal{O}$ of norm $\leq N$ is finite.*

So while there are infinitely many ideals, if we bound the norm then we just get a finite set. The proof is not difficult, but we'll skip it. So, to prove that Cl_F is finite, we just need to see that every nonzero ideal is isomorphic to one of bounded norm. But again, recall that if α is a nonzero element of an ideal I, then $(\alpha) \subset I$, so there exists an ideal K with $IK = (\alpha)$. In other words, K is inverse to I in the ideal class group. But look: K will have small norm *provided* that α can be chosen to have a norm which is not too far from that of I. Since inversion is a bijection on any group, it follows that if we can just show that every nonzero ideal I contains a nonzero element α such that $N(I)/N(\alpha)$ is bounded, then every ideal class will be represented by an ideal of bounded norm, so Cl_F will be finite by the above proposition.

So this is really the crucial thing: to show that every nonzero ideal I, though not necessarily principal, is at most a bounded distance away from being principal.

And that kind of statement can be proven using the geometry of numbers, as in your second problem of PSet 2. Imagine I as sitting in Minkowski space. Consider the set of α in Minkowski space such that $|N(\alpha)| \leq C \cdot N(I)$, where C is (for now) some arbitrary constant. This describes some region in Minkowski space centered around the origin. What we need to see is that for some choice of C (independent of I), this region intersects I nontrivially. Minkowski's way of doing this is to find a convex centrally symmetric figure contained in this region which has large enough volume that, by choosing C appropriately, one can guarantee that some nonzero lattice point has to be contained in it. That, in the end, gives the proof.

Another way of interpreting this finiteness of the ideal class group is the following. You can think of some finite set of "bad ideals", or even of "bad maximal ideals", which, if only they were principal, we

would have unique prime factorization. In general, this is actually how you go about things in practice if you want to verify unique prime factorization for some general ring of integers. Minkowski gives a pretty good choice of C, so you only need to find generators for finitely many ideals if you want to know that all of them are principal.

Next we'll turn to some analysis and zeta functions. These are amazing objects. They are simple to write down, but if you want to study them, you're forced to do all sorts of interesting things, and some amazing facts pop out. For instance, one can, in some cases, obtain completely explicit finitary formulas for class numbers, i.e. the orders of these finite groups Cl_F, which look crazy when you first see them. (Actually, I still think they look crazy.)