The volume of a number field

December 7, 2013

Let F be a number field, of degree $n = [F : \mathbb{Q}]$. Recall from last time the *Minkowki space*

$$F_{\mathbb{R}} = \mathbb{R}^r \oplus \mathbb{C}^s.$$

Here r is the number of real embeddings $F \to \mathbb{R}$, and s is half of the number of complex embeddings $F \to \mathbb{C}$ which are not real. Note that the total number of embeddings $F \to \mathbb{C}$ is then $r + 2s$. It follows that $r + 2s = n$. Thus we see that $dim_{\mathbb{R}} F_{\mathbb{R}} = n$, a crucial property.

Recall also that there is a canonical embedding

$$F \to R_{\mathbb{R}}$$

which assigns to an $\alpha \in F$ its Minkowski coordinates. These are defined as follows. The $\mathbb{R}$-components can be labelled by the real embeddings $\sigma : F \to \mathbb{R}$. For such a σ, the σ^{th} Minkowski coordinate of α is just $\sigma(\alpha) \in \mathbb{R}$. As for the $\mathbb{C}$-components, we can label these by complex embeddings $\sigma : F \to \mathbb{C}$, as soon as we choose one out of each pair of complex conjugate complex embeddings. Then again the σ^{th} Minkowski coordinate is $\sigma(\alpha) \in \mathbb{C}$.

One way to get rid of this strange need to make a choice of half the complex embeddings is as follows. The complex analog of Minowski space is just $\mathbb{C}^n$, where now the components are labelled by all the embeddings $\sigma : F \to \mathbb{C}$. Here the canonical map $F \to \mathbb{C}^n$ just sends an α to $\sigma(\alpha)$ in the σ^{th} component, for all $\sigma : F \to \mathbb{C}$. But we can note that the image of F lies inside the fixed points of a natural $\mathbb{Z}/2$-action on $\mathbb{C}^n$: namely, let the generator of $\mathbb{Z}/2$ act by complex conjugation simultaneously on the indexing set $\{\sigma\}$ and on the copies of $\mathbb{C}$ that make up the space. Then the canonical version of the Minkowski space and its embedding is just the induced map

$$F \to (\mathbb{C}^n)^{\mathbb{Z}/2},$$

from F the fixed points of the complex-conjugation action on complex Minkowski space.

To see what these fixed points are doing, note that for those $\sigma : F \to \mathbb{C}$ which land in $\mathbb{R}$, the action by complex conjugation on σ is trivial. So on those factors the $\mathbb{Z}/2$-action preserves each factor, and is there just given by the usual complex conjugation on $\mathbb{C}$. Thus the fixed points give a copy of $\mathbb{R}$ for every real embedding, just as we had with our old (non-canonical) Minkowski space. But for the non-real σ, now the conjugation pairs up these σ's, and we instead find a copy of the $\mathbb{Z}/2$-fixed points of $\mathbb{C} \oplus \mathbb{C}$, for the $\mathbb{Z}/2$-action which simultaneously swaps the coordinates and takes complex conjugates. These fixed points are just the set of (z, w) which are each others' conjugates. If we choose one of the factors we can identify this with $\mathbb{C}$ and thus recover the old Minkowski space, but there is no such canonical choice available.

So it seems that the canonical Minkowski space $(\mathbb{C}^n)^{\mathbb{Z}/2}$ is nicer from an abstract perspective, but the non-canonical Minkowski space $\mathbb{R}^r \oplus \mathbb{C}^s$ is easier to picture a priori.

Now, let's not lose sight here: the above discussion was just a technical side-track. But it can be helpful to try to see things in the most canonical way possible, so that intrinsic structures can come to the fore. This point will come up when we turn to the topic of this lecture, which is the geometry of Minkowski space.

So far, we've already used the Minkowski space to prove that $\mathcal{O}_F$ is free of rank n, by showing that it was a full lattice in Minkowski space. There, the only structure that was relevant was the *topology* on Minkowski space. In the proof that a full lattice is necessarily free of rank n we did use an arbitrary choice of inner product, so that we could talk about distances and balls and so on, but that was just supplementary structure, a technical artifact of the proof. It didn't enter into either the hypotheses or the conclusions.

But now we can see that Minkowski space $F_{\mathbb{R}}$ actually carries a canonical inner product — and hence, by the usual game from real geometry, a notion of distance, of angle, of volume, etc. — this all being completely determined by the original number field F.

In terms of the old Minkowski space $\mathbb{R}^r \oplus \mathbb{C}^s$, we can just imagine the usual Euclidean geometry on each factor. This is represented by the inner product

$$\langle (x_\sigma), (y_\sigma) \rangle = \sum_\sigma \left(Re(x_\sigma) Re(y_\sigma) + Im(x_\sigma) Im(y_\sigma) \right),$$

where σ labels the Minkowski coordinates, so for the complex embeddings we have to choose half. Note that on the $\mathbb{R}^r$ part, this just recovers the usual dot product of vectors in $\mathbb{R}^r$, and on the complex part, it's the same with respect to the decomposition $\mathbb{C} = \mathbb{R} \oplus \mathbb{R}$ into real and purely imaginary axes.

Even though the above description depended on our choice of half of the non-real complex embeddings, we can see that this inner product is canonical, i.e. it doesn't depend on the choice. This is because changing σ to $\overline{\sigma}$ will just add a sign to each of the Im-terms, and these signs cancel out. Or, in other words, complex conjugation preserves angles, distances, etc. on $\mathbb{C}$. So it might seem that, as far as inner products are concerned, there's nothing to gain from thinking about the canonical Minkowski space $(\mathbb{C}^n)^{\mathbb{Z}/2}$. But actually, by adopting that perspective we will get a different inner product, one which is arguably "better".

Namely, the canonical geometry on the complex Minkowski space $\mathbb{C}^n$ is the "Hermitian" one, given by

$$\langle (x_\sigma), (y_\sigma) \rangle = \sum_\sigma x_\sigma \cdot \overline{y_\sigma},$$

where now σ runs over all $F \to \mathbb{C}$. When we restrict this pairing to the canonical Minkowski space $(\mathbb{C}^n)^{\mathbb{Z}/2}$, we get a new inner product, which can be described in terms of the Minkowski coordinates as

$$\langle (x_\sigma), (y_\sigma) \rangle = \sum_{\sigma \in \text{real}} x_\sigma \cdot \overline{y_\sigma} + \sum_{\sigma \in \frac{1}{2}\text{complex}} \left(x_\sigma \cdot \overline{y_\sigma} + \overline{x_\sigma} \cdot y_\sigma \right).$$

Now again we've chosen half the complex embeddings because we're talking about Minkowski coordinates, but here we can see that the inner product doesn't depend on this choice for an easier reason: addition is commutative.

This inner product on $F_{\mathbb{R}}$ is different from the one we had before, though not substantially. Certainly, they agree on the $\mathbb{R}$-factors. On the $\mathbb{C}$-factors, you can calculate that the two inner products are related by a scalar factor of 2. (The canonical one is the "bigger" one.) Since the norm associated to a scalar product is $\sqrt{\langle z, z\rangle}$, this means that distances are scaled by $\sqrt{2}$ on each complex factor, or areas are scaled by 2 on each complex factor.

I'm going into all this because the question we're going to address is, "what is the volume of a number field?". And this question depends on which inner product we've chosen. It turns out you get a slightly nicer answer if you use the second inner product, the "canonical one".

To make the question more precise, recall that $\mathcal{O}_F$ sits as a full lattice in $F_{\mathbb{R}}$. Another way of saying this is that the pair $\mathcal{O}_F \subset F_{\mathbb{R}}$ is isomorphic to the standard pair $\mathbb{Z}^n \subset \mathbb{R}^n$. In particular, the quotient

$$F_{\mathbb{R}}/\mathcal{O}_F$$

is a torus: it's homeomorphic to $\mathbb{R}^n/\mathbb{Z}^n$. But the canonical volume on $F_{\mathbb{R}}$ coming from our inner product is translation-invariant, so it descends to a notion of volume on this torus $F_{\mathbb{R}}/\mathcal{O}_F$. Since this space is compact, it in particular makes sense to ask, what is the volume of this torus? I.e., what is

$$Vol(F_{\mathbb{R}}/\mathcal{O}_F)?$$

This is what I mean by "volume of a number field". Our number field F has canonically determined for us a certain real number, the volume of this torus. It's reasonable to ask what this number is, or how to compute it. The answer does have arithmetic significance.

Let us give a more concrete interpretation of this question, which will also help us answer it. Think back to the usual torus $\mathbb{R}^n/\mathbb{Z}^n$. The cube $[0,1]^n$ provides a nice compact set of representatives for this torus: everything in $\mathbb{R}^n$ is a translation of $[0,1]^n$ by something in $\mathbb{Z}^n$. There is only the slight redundancy that opposite edges of this cube will give the same point in the torus. But that doesn't matter for the purposes of calculating volume. Thus we see that the standard torus $\mathbb{R}^n/\mathbb{Z}^n$ has the same volume as the cube $[0,1]^n$, which is 1. Here we use the usual Euclidean inner product (dot product) on $\mathbb{R}^n$.

To move closer to our case of $\mathcal{O}_F \subset F_{\mathbb{R}}$, suppose given a full lattice Λ in $\mathbb{R}^n$ which is potentially not the standard one, but instead is given by some other basis $b_1, \ldots, b_n$. Then the analog of the above cube is the "fundamental parallelopiped"

$$\{\lambda_1 b_1 + \ldots + \lambda_n b_n \mid 0 \leq \lambda_i \leq 1\}.$$

We can understand the volume of this parallelopiped, and hence the volume of the quotient torus $\mathbb{R}^n/\Lambda$, by imagining the linear transformation T which takes the standard basis to the basis $b_1, \ldots, b_n$. In other words, T is represented by the matrix whose column vectors are the b_i. Since a linear transformation scales volumes by the factor of the absolute value of its determinant, and our parallelopiped is the image under T of the standard cube $[0,1]^n$, it follows that the volume of $\mathbb{R}^n/\Lambda$ is just

$$Vol(\mathbb{R}^n/\Lambda) = |det(T)|,$$

where again T is any matrix whose column vectors are a $\mathbb{Z}$-basis of Λ.

Now, this literally applies to our case $\mathcal{O}_F \subset F_{\mathbb{R}}$, as long as every complex embedding of F is actually real. (Fields F satisfying this condition are called "totally real"), so that the inner product on

the Minkowski space is just the standard one on $\mathbb{R}^n$. The conclusion is that if F is totally real, and if $b_1, \ldots, b_n$ denotes a $\mathbb{Z}$-basis of $\mathcal{O}_F$, then

$$Vol(F_{\mathbb{R}}/\mathcal{O}_F) = |det(\sigma b_i)|,$$

where (σb_i) stands for an $n \times n$-matrix whose columns are indexed by the basis vectors and whose rows are indexed by the embeddings $\sigma : F \to \mathbb{C}$ (which, in this case, all land in $\mathbb{R}$).

Now, I'm here to tell you that, in fact, the exact same conclusion holds for an arbitrary number field F, *as long as* you use the inner product on the canonical Minkowski space $F_{\mathbb{R}} = (\mathbb{C}^n)^{\mathbb{Z}/2}$, the one induced from the Hermitian inner product on $\mathbb{C}^n$. I won't justify this; it's just a little calculation. Because of the comparison between the naive Minkowski metric and this canonical one, we see that the analogous volume in the naive Minkowski metric will differ from this answer by a factor of 2^{-s}. But let us decide to stick with the canonical Minowski metric, so that, again, we have the formula

$$Vol(F_{\mathbb{R}}/\mathcal{O}_F) = |det(\sigma b_i)|$$

in complete generality.

Modulo the problem of finding a $\mathbb{Z}$-basis for $\mathcal{O}_F$ (which we'll come back to at the end!), this formula is actually very practical for calculating $Vol(F_{\mathbb{R}}/\mathcal{O}_F)$. It also gives us theoretic consequences, for example the following:

Proposition 0.1. *Let F be a number field, with ring of integers $\mathcal{O}_F$ embedded as a full lattice in Minkowski space $F_{\mathbb{R}} = (\mathbb{C}^n)^{\mathbb{Z}/2}$ with its canonical inner product. Then the real number $Vol(F_{\mathbb{R}}/\mathcal{O}_F)$ is the square root of an integer.*

So somehow, the integrality of $\mathcal{O}_F$ gets reflected in this volume as well. By the way, the claimed integer $Vol(F_{\mathbb{R}}/\mathcal{O}_F)^2$ is, up to a sign, a quantity called the *discriminant* of the number field F. This discriminant can also be defined by various algebraic means, independently of Minkowski space.

Proof. Consider the quantity $Vol(F_{\mathbb{R}}/\mathcal{O}_F)^2$. So we want to show $Vol(F_{\mathbb{R}}/\mathcal{O}_F)^2 \in \mathbb{Z}$. Since the only algebraic integers in $\mathbb{Q}$ are the actual integers, it suffices to see that this quantity is both an algebraic integer and lies in $\mathbb{Q}$. That it is an algebraic integer follows from the above formula

$$Vol(F_{\mathbb{R}}/\mathcal{O}_F) = |det(\sigma b_i)|.$$

Indeed, since each b_i is an algebraic integer, so is each σb_i (since σb_i must satisfy the same polynomial equations as b_i), and therefore so is the determinant, seeing as algebraic integers form a ring. Finally, the absolute value just gives a ± 1; it doesn't affect integrality.

So we've checked that $Vol(F_{\mathbb{R}}/\mathcal{O}_F)^2$ is an algebraic integer, meaning to finish we only need to see that it's rational, or, equivalently, that it's fixed by every Galois automorphism. In other words, we need to check that $Vol(F_{\mathbb{R}}/\mathcal{O}_F) = |det(\sigma b_i)|$ is fixed by Galois, up to a sign. But this follows because Galois just permutes the various σ, i.e. it permutes the rows of the matrix we're taking the determinant of. So a sign is introduced corresponding to the sign of this permutation. Then the absolute value again at most introduces another sign. So in total we have verified the claim. □

This property of the volume is actually useful in practice. To explain one reason why, let's turn to the unaddressed question of how to find a $\mathbb{Z}$-basis of $\mathcal{O}_F$.

We can always start by just randomly choosing an integral $\alpha \in F$ which generates F. Then we know that $\mathbb{Z}[\alpha]$ is a free abelian group of rank n (in fact generated by the first n powers of α). Thus the inclusion

$$\mathbb{Z}[\alpha] \subset \mathcal{O}_F$$

is an inclusion of free abelian groups of the same rank. The quotient must therefore be finite — in other words, $\mathbb{Z}[\alpha]$ is of finite index in $\mathcal{O}_F$. So in a sense, we're almost there, starting from just any random choice of α.

To get all the way and understand the whole $\mathcal{O}_F$, remark that if N denotes the index $[\mathcal{O}_F : \mathbb{Z}[\alpha]]$, then $N\mathcal{O}_F \subset \mathbb{Z}[\alpha]$ (because N kills the quotient). Thus

$$\mathbb{Z}[\alpha] \subset \mathcal{O}_F \subset \frac{1}{N}\mathbb{Z}[\alpha].$$

So we have $\mathcal{O}_F$ sandwiched in between two understood abelian groups. Furthermore, the quotient $(\frac{1}{N}\mathbb{Z}[\alpha])/\mathbb{Z}[\alpha]$ is finite, of size N^n. So we only have to go through the N^n coset representatives of this quotient, and figure out which are actually algebraic integers (e.g. by computing their minimal polynomials). You can even save some work by remembering that $\mathcal{O}_F$ is a ring. This gives a nice algorithm for constructing $\mathcal{O}_F$, *except* for the fact that we might have no idea a priori what the index $N = [\mathcal{O}_F : \mathbb{Z}[\alpha]]$ is. We don't know, given an initial guess of α, how far away $\mathbb{Z}[\alpha]$ is from the true ring of integers.

That is, we don't unless we remember the above facts about volumes. Note that the same volume discussion goes through with $\mathbb{Z}[\alpha]$ instead of $\mathcal{O}_F$: all we used about $\mathcal{O}_F$ was that it's a full lattice and consists of algebraic integers. Thus, the volume $Vol(F_{\mathbb{R}}/\mathbb{Z}[\alpha])^2$ must also be an integer. Furthermore, it's something we can compute just knowing α, by using the $\mathbb{Z}$-basis $1, \alpha, \ldots, \alpha^{n-1}$ of $\mathbb{Z}[\alpha]$.

But now, I claim that there is an equality

$$Vol(F_{\mathbb{R}}/\mathbb{Z}[\alpha]) = [\mathcal{O}_F : \mathbb{Z}[\alpha]] \cdot Vol(F_{\mathbb{R}}/\mathcal{O}_F).$$

This should make sense: it's another reflection of the idea that the difference in "size" of $\mathbb{Z}[\alpha]$ and $\mathcal{O}_F$ is measured by the index $[\mathcal{O}_F : \mathbb{Z}[\alpha]]$. This is really just a specialization of a more general fact which is also touched on in your first homework problem this week, so I'll be brief in explaining things here. One can think in one of two ways. First, if one chooses coset representatives for $\mathbb{Z}[\alpha]$ in $\mathcal{O}_F$, then one can use these to realize a fundamental paralellopiped of $\mathbb{Z}[\alpha]$ as made out of $[\mathcal{O}_F : \mathbb{Z}[\alpha]]$ copies of a fundamental parallelopiped for $\mathcal{O}_F$. This proves the claim. Second, and more canonically, one can think of the induced map of tori

$$F_{\mathbb{R}}/\mathbb{Z}[\alpha] \to F_{\mathbb{R}}/\mathcal{O}_F.$$

One sees that this is an $[\mathcal{O}_F : \mathbb{Z}[\alpha]]$-sheeted covering map, and this implies the desired claim as well.

The conclusion is that, whatever the index $[\mathcal{O}_F : \mathbb{Z}[\alpha]]$ may be, its *square* must be a divisor of the integer $Vol(F_{\mathbb{R}}/\mathbb{Z}[\alpha])^2$, which we can calculate. This generally gives a pretty good bound on $[\mathcal{O}_F : \mathbb{Z}[\alpha]]$. For example, it can even happen that you go to compute $Vol(F_{\mathbb{R}}/\mathbb{Z}[\alpha])^2$, and it turns out to be a square-free integer. Then you know that necessarily $\mathcal{O}_F = \mathbb{Z}[\alpha]$. So you made a lucky choice of α. In general, you can think like this: if the volume of $F_{\mathbb{R}}/\mathbb{Z}[\alpha]$ is small, then there's not much room for there to be a ring of algebraic integers strictly larger than $\mathbb{Z}[\alpha]$.

In the next lecture, we'll see another example of this. We'll use the above idea to prove that the ring of integers of $\mathbb{Q}(\zeta_p)$ is indeed $\mathbb{Z}[\zeta_p]$, as one might have assumed all along. Here the Vol^2 won't

turn out to be squarefree, but we'll be able to appeal to our old friend $p = (1 - \zeta_p) \dots (1 - \zeta_p^{p-1})$ to get the desired conclusion anyway.