

In Lecture 6 we proved (most of) Ostrowski's theorem for number fields, and we saw the product formula for absolute values on \mathbb{Q} . A similar product formula holds for absolute values on a number field, but in order to state and prove it we need to briefly review/introduce some standard terminology from algebraic number theory.

7.1 Field norms and traces

Let L/K be a finite field extension of degree $n = [L : K]$. Then L is an n -dimensional K -vector space, and each $\alpha \in L$ determines a linear operator $T_\alpha : L \rightarrow L$ corresponding to multiplication by α (the linearity of T_α is immediate from the field axioms).

Definition 7.1. The *trace* $\text{Tr}_{L/K}(\alpha)$ is the trace of T_α , and the *norm* $N_{L/K}(\alpha)$ is the determinant of T_α .¹

It follows immediately from this definition that the trace is additive and the norm is multiplicative, and that both take values in K .

The trace and norm can be computed as the trace and determinant of the matrix of T_α with respect to a basis, but their values are intrinsic to α and do not depend on a choice of basis. The Cayley-Hamilton theorem implies that T_α satisfies a characteristic equation

$$f_\alpha(x) = x^n + a_{n-1}x + \cdots + a_1x + a_0 = 0$$

with coefficients $a_i \in K$. We then have

$$\text{Tr}_{L/K}(\alpha) = -a_{n-1} \quad \text{and} \quad N_{L/K}(\alpha) = (-1)^n a_0,$$

equivalently, $\text{Tr}_{L/K}(\alpha)$ and $N_{L/K}(\alpha)$ are the sum and product of the roots of f_α , respectively. These roots need not lie in L , but they certainly lie in \overline{K} (in fact in the splitting field of f_α), and in any case their sum and product necessarily lie in K .

Note that α satisfies the same characteristic equation as T_α , since T_α is just multiplication by α , but f_α is not necessarily the minimal polynomial g_α of α over K (which is also the minimal polynomial of the operator T_α). We know that g_α must divide f_α , since the minimal polynomial always divides the characteristic polynomial, but f_α must be a power of g_α . This is easy (and instructive) to prove in the case that L/K is a separable extension, which includes all the cases of interest to us.²

Theorem 7.2. Let L/K be a separable field extension of degree n , let $\alpha \in L$ have minimal polynomial g_α over K and let f_α be the characteristic polynomial of T_α . Then

$$f_\alpha = g_\alpha^{n/d},$$

where $d = [K(\alpha) : K]$.

¹These are also called the *relative* trace/norm, or the trace/norm *from L down to K* to emphasize that they depend on the fields L and K , not just α .

²Recall that *separable* means that minimal polynomials never have repeated roots. In characteristic zero every finite extension is separable, and the same holds for finite fields (such fields are said to be *perfect*).

Proof. There are exactly n distinct embeddings $\sigma_1, \dots, \sigma_n$ of L into \overline{K} that fix K , and $\sigma_1(\alpha), \dots, \sigma_n(\alpha)$ are precisely the n (not necessarily distinct) roots of f_α . This list includes the d roots of g_α , since g_α divides f_α , and these d roots are distinct, since L/K is separable. But there are exactly $n/d = [L : K(\alpha)]$ distinct embeddings of L into \overline{K} that fix $K(\alpha)$, and each of these also fixes K and is hence one of the σ_i . It follows that each distinct root of f_α occurs with multiplicity at least n/d , and since f_α has at least d distinct roots, the roots of f_α are precisely the roots of g_α , each occurring with multiplicity n/d . Both f_α and g_α are monic, so $f_\alpha = g_\alpha^{n/d}$. \square

7.2 Ideal norms

Now let us fix $K = \mathbb{Q}$, so that L is a number field (a finite extension of \mathbb{Q}). Recall that the *ring of integers* of L consists of the elements in L whose minimal polynomials have integer coefficients. This subset forms a ring \mathcal{O} that is a *Dedekind domain*, an integral domain in which every nonzero proper ideal can be uniquely factored into prime ideals (equivalently, a finitely generated Noetherian ring in which every nonzero prime ideal is maximal), and L is its fraction field. The ring of integers is a free \mathbb{Z} -module of rank $n = [L : \mathbb{Q}]$, and we can pick a basis for L as an n -dimensional \mathbb{Q} -vector space that consists of elements of \mathcal{O} (such a basis is called an *integral basis*). The ring \mathcal{O} then consists of all integer linear combinations of basis elements and can be viewed as an n -dimensional \mathbb{Z} -lattice. For proofs of these facts, see any standard text on algebraic number theory, such as [1].

Definition 7.3. Let \mathfrak{a} be a nonzero \mathcal{O} -ideal. The (ideal) *norm* $N\mathfrak{a}$ of \mathfrak{a} is the cardinality of the (necessarily finite) ring \mathcal{O}/\mathfrak{a} , equivalently, the index $[\mathcal{O} : \mathfrak{a}]$ of \mathfrak{a} as a sublattice of the \mathbb{Z} -lattice \mathcal{O} .³ The norm of (0) is zero.

Remark 7.4. In a Dedekind domain every nonzero prime ideal is maximal, so for prime ideals \mathfrak{p} the ring \mathcal{O}/\mathfrak{p} is actually a field of cardinality $N\mathfrak{p} = p^f$, for some prime p and positive integer f called the *inertia degree* (also *residue degree*).

While it may not be immediately obvious from the definition, the ideal norm is multiplicative (for principal ideals this follows from Theorem 7.5 below). For an algebraic integer $\alpha \in L$ we now have two notions of norm: the field norm $N_{L/\mathbb{Q}}(\alpha)$ and the ideal norm $N(\alpha)$ of the principal \mathcal{O} -ideal generated by α . These are not unrelated.

Theorem 7.5. *Let α be an algebraic integer in a number field L . Then $N(\alpha) = |N_{L/\mathbb{Q}}(\alpha)|$.*

Proof. Fix an integral basis \mathcal{B} for L . The field norm $N_{L/\mathbb{Q}}(\alpha)$ is the determinant of the matrix of the linear operator T_α with respect to \mathcal{B} . The absolute value of this determinant is equal to the volume of a fundamental parallelepiped in the \mathbb{Z} -lattice corresponding to the principal ideal (α) as a sublattice of the \mathbb{Z} -lattice \mathcal{O} generated by \mathcal{B} , relative to the volume of a fundamental parallelepiped in \mathcal{O} . But this is precisely the index $[\mathcal{O} : (\alpha)] = N(\alpha)$. \square

³Like the field norm $N_{L/\mathbb{Q}}$, the ideal norm N depends on L , but we typically don't indicate L in the notation because N is always applied to ideals, which necessarily exist in the context of a particular ring (in our case the ring of integers of L). More generally, for any finite separable extension L/K where K is the fraction field of a Dedekind domain A , the ideal norm is defined as a map from ideals in the integral closure of A in L to A -ideals. In our setting $A = \mathbb{Z}$ is a PID, so we are effectively identifying the \mathbb{Z} -ideal $(N\mathfrak{a})$ with the integer $N\mathfrak{a}$. See [1, Ch. 4] for more details. Our definition here is also called the *absolute norm*.

7.3 Product formula for absolute values on number fields

Ostrowski's theorem for number fields classifies the absolute values on a number field up to equivalence. But in order to prove the product formula we need to properly normalize each absolute value appropriately, which we now do.

Let L be a number field with ring of integers \mathcal{O} . For each nonzero prime ideal \mathfrak{p} in \mathcal{O} we define the absolute value $|\alpha|_{\mathfrak{p}}$ on L by

$$|\alpha|_{\mathfrak{p}} = (N\mathfrak{p})^{-v_{\mathfrak{p}}(\alpha)},$$

where $v_{\mathfrak{p}}(\alpha)$ is the exponent of \mathfrak{p} in the prime factorization of the ideal (α) for nonzero $\alpha \in \mathcal{O}$, and $v_{\mathfrak{p}}(\alpha/\beta) = v_{\mathfrak{p}}(\alpha) - v_{\mathfrak{p}}(\beta)$ for any nonzero $\alpha, \beta \in \mathcal{O}$ (recall that L is the fraction field of \mathcal{O}). As usual, we let $v_{\mathfrak{p}}(0) = \infty$ and define $(N\mathfrak{p})^{-\infty} = 0$.

This addresses all the nonarchimedean absolute values of L (by Ostrowski's theorem), we now consider the archimedean ones. As a number field of degree n , there are exactly n distinct embeddings of L into $\overline{\mathbb{Q}}$, hence into \mathbb{C} . But these n embeddings do not necessarily give rise to n distinct absolute values. Let f be a defining polynomial for L over \mathbb{Q} , that is, the minimal polynomial of a primitive element θ such that $L = \mathbb{Q}(\theta)$ (such a θ exists, by the primitive element theorem). Over \mathbb{C} , the roots of f are either real (let r be the number of real roots) or come in complex-conjugate pairs (let s be the number of such pairs). We then have $n = r + 2s$ distinct embeddings of L into \mathbb{C} , each sending θ to a different root of f (the roots are distinct because every finite extension of \mathbb{Q} is separable). But there are only $r + s$ inequivalent archimedean absolute values on L , since complex-conjugate embeddings yield the same absolute value ($|z| = |\bar{z}|$).

As with \mathbb{Q} , it will be convenient to use the notation $|\cdot|_{\mathfrak{p}}$ to denote archimedean absolute values as well as nonarchimedean ones, and we may refer to the subscript \mathfrak{p} as an archimedean or "infinite" prime and write $\mathfrak{p}|\infty$ to indicate this.⁴ Using $\sigma_{\mathfrak{p}}$ to denote the embedding associated to a real archimedean prime \mathfrak{p} and $\sigma_{\mathfrak{p}}, \bar{\sigma}_{\mathfrak{p}}$ to denote the conjugate pair of complex embeddings associated to a complex archimedean prime \mathfrak{p} , we now define

$$|\alpha|_{\mathfrak{p}} = \begin{cases} |\sigma_{\mathfrak{p}}(\alpha)| & \text{if } \mathfrak{p} \text{ is a real archimedean prime,} \\ |\sigma_{\mathfrak{p}}(\alpha)| \cdot |\bar{\sigma}_{\mathfrak{p}}(\alpha)| & \text{if } \mathfrak{p} \text{ is a complex archimedean prime.} \end{cases}$$

Of course $|\sigma_{\mathfrak{p}}(\alpha)| \cdot |\bar{\sigma}_{\mathfrak{p}}(\alpha)| = |\sigma_{\mathfrak{p}}(\alpha)|^2$, but it is more illuminating to write it as above.

We now prove the product formula for absolute values on number fields.

Theorem 7.6. *Let L be a number field. For every $\alpha \in L^{\times}$ we have*

$$\prod_{\mathfrak{p}} |\alpha|_{\mathfrak{p}} = 1,$$

where \mathfrak{p} ranges over all the primes of L (both finite and infinite).

Proof. We first consider the archimedean primes. Let f_{α} be the characteristic polynomial of the linear operator on the \mathbb{Q} -vector space L corresponding to multiplication by α . If $\mathfrak{p}_1, \dots, \mathfrak{p}_r$ and $\mathfrak{p}_{r+1}, \dots, \mathfrak{p}_{r+s}$ are the real and complex archimedean primes of L , then the $n = r + 2s$ (not necessarily distinct) roots of f_{α} are precisely

$$\sigma_{\mathfrak{p}_1}(\alpha), \dots, \sigma_{\mathfrak{p}_r}(\alpha), \sigma_{\mathfrak{p}_{r+1}}(\alpha), \bar{\sigma}_{\mathfrak{p}_{r+1}}(\alpha), \dots, \sigma_{\mathfrak{p}_{r+s}}(\alpha), \bar{\sigma}_{\mathfrak{p}_{r+s}}(\alpha).$$

⁴The finite and infinite primes of L are also often referred to as *places* of L and denoted by v .

We then have

$$\prod_{\mathfrak{p}|\infty} |\alpha|_{\mathfrak{p}} = \prod_{i=1}^r |\sigma_{\mathfrak{p}_i}(\alpha)| \prod_{i=r+1}^s |\sigma_{\mathfrak{p}_i}(\alpha)| \cdot |\bar{\sigma}_{\mathfrak{p}_i}(\alpha)| = |N_{L/\mathbb{Q}}(\alpha)|,$$

since $N_{L/\mathbb{Q}}(\alpha)$ is equal to the product of the roots of f_{α} .

Now let $(\alpha) = \mathfrak{q}_1^{a_1} \cdots \mathfrak{q}_t^{a_t}$ be the prime factorization of the principal ideal (α) in the ring of integers of L . Then

$$\prod_{\mathfrak{p}<\infty} |\alpha|_{\mathfrak{p}} = \prod_{i=1}^t (N\mathfrak{q}_i)^{-a_i} = N(\alpha)^{-1} = |N_{L/\mathbb{Q}}(\alpha)|^{-1},$$

by Theorem 7.5, and therefore $\prod_{\mathfrak{p}} |\alpha|_{\mathfrak{p}} = 1$, as desired. \square

We now turn to a new topic, the *completion* of a field with respect to an absolute value.

7.4 Cauchy sequences and convergence

We begin with the usual definitions of convergence and Cauchy sequences, which apply to any field with an absolute value. Let k be a field equipped with an absolute value $\| \cdot \|$.

Definition 7.7. A sequence (x_n) of elements of k *converges* (to ℓ) if there is an element $\ell \in k$ such that for every $\epsilon > 0$ there is a positive integer N such that $\|x_n - \ell\| < \epsilon$ for all $n \geq N$. Equivalently, (x_n) converges to ℓ if $\|x_n - \ell\| \rightarrow 0$ as $n \rightarrow \infty$.⁵

The element ℓ is called the *limit* of the sequence, and if it exists, it is unique: if (x_n) converges to both ℓ and ℓ' then

$$\|\ell' - \ell\| = \|\ell' - x_n + x_n - \ell\| \leq \|\ell' - x_n\| + \|x_n - \ell\| = \|x_n - \ell'\| + \|x_n - \ell\| \rightarrow 0 + 0 = 0,$$

so $\|\ell' - \ell\| = 0$, and therefore $\ell' - \ell = 0$ and $\ell' = \ell$ (note that we used $\| -x \| = \|x\|$).

Sums and products of convergent sequences behave as expected.

Lemma 7.8. Let (x_n) and (y_n) be sequences in k that converge to x and y respectively. Then the sequences $(x_n + y_n)$ and $(x_n y_n)$ converge to $x + y$ and xy respectively.

Proof. Convergence of $(x_n y_n)$ to xy follows immediately from the multiplicativity of $\| \cdot \|$. To check $(x_n + y_n)$, for any $\epsilon > 0$ pick N so that $\|x - x_n\| < \epsilon/2$ and $\|y - y_n\| < \epsilon/2$ for all $n \geq N$. Then $\|(x_n + y_n) - (x + y)\| \leq \|x_n - x\| + \|y_n - y\| < \epsilon/2 + \epsilon/2 = \epsilon$ for all $n \geq N$. \square

We now recall a necessary condition for convergence.

Definition 7.9. A sequence (x_n) in k is a *Cauchy sequence* if for every $\epsilon > 0$ there exists a positive integer N such that $\|x_m - x_n\| < \epsilon$ for all $m, n \geq N$.

Theorem 7.10. Every convergent sequence is a Cauchy sequence.

Proof. Suppose (x_n) is a convergent sequence. For any $\epsilon > 0$ there is a positive integer N for which $\|x_n - \ell\| < \epsilon/2$ for all $n \geq N$. For all $m, n \geq N$ we then have

$$\|x_m - x_n\| = \|x_m - \ell + \ell - x_n\| \leq \|x_m - \ell\| + \|\ell - x_n\| = \|x_m - \ell\| + \|x_n - \ell\| < \epsilon/2 + \epsilon/2 = \epsilon,$$

where we have again used $\| -x \| = \|x\|$. \square

⁵The notation $\|x_n - \ell\| \rightarrow 0$ refers to convergence in \mathbb{R} in the usual sense.

The converse of Theorem 7.10 is not necessarily true, it depends on the field k .

Definition 7.11. A field k is *complete* (with respect to $\| \cdot \|$) if every Cauchy sequence in k converges (to an element of k).

Every field is complete with respect to the trivial absolute value. The field \mathbb{Q} is not complete with respect to the archimedean absolute value $| \cdot |$, but \mathbb{R} is; indeed, \mathbb{R} can be (and often is) defined as the smallest field containing \mathbb{Q} that is complete with respect to $| \cdot |$, in other words, \mathbb{R} is the completion of \mathbb{Q} . In order to formally define the completion of a field, we define an equivalence relation on sequences.

Definition 7.12. Two sequences (a_n) and (b_n) are *equivalent* if $\|a_n - b_n\| \rightarrow 0$ as $n \rightarrow \infty$.

It is easy to check that this defines an equivalence relation on the set of all sequences in k , and that any sequence equivalent to a Cauchy sequence is necessarily a Cauchy sequence. We may use the notation $[(x_n)]$ to denote the equivalence class of the sequence (x_n) .

Definition 7.13. The *completion* of k (with respect to $\| \cdot \|$) is the field \hat{k} whose elements are equivalence classes of Cauchy sequences in k , where

- (1) $0_{\hat{k}} = [(0_k, 0_k, 0_k, \dots)]$,
- (2) $1_{\hat{k}} = [(1_k, 1_k, 1_k, \dots)]$,
- (3) $[(x_n)] + [(y_n)] = [(x_n + y_n)]$ and $[(x_n)][(y_n)] = [(x_n y_n)]$.

To verify that this actually defines a field, the only nontrivial thing to check is that every nonzero element has a multiplicative inverse. So let $[(x_n)]$ be a nonzero element of \hat{k} . The Cauchy sequence (x_n) must be eventually nonzero (otherwise it would be equivalent to zero), and if we consider the element $[(y_n)] \in \hat{k}$ defined by

$$y_n = \begin{cases} x_n^{-1} & \text{if } x_n \neq 0, \\ 0 & \text{if } x_n = 0, \end{cases}$$

we see that $[(x_n)][(y_n)] = 1$, since the sequence $(x_n y_n)$ is eventually 1.

The map $x \mapsto \hat{x} = [(x, x, x, \dots)]$ is clearly a ring homomorphism from k to \hat{k} , and therefore a field embedding. We thus view \hat{k} as an extension of k by identifying k with its image in \hat{k} .

We now extend the absolute value of k to \hat{k} by defining

$$\|[(x_n)]\| = \lim_{n \rightarrow \infty} \|x_n\|.$$

This limit exists because $(\|x_n\|)$ is a Cauchy sequence of real numbers and \mathbb{R} is complete, and we must get the same limit for any Cauchy sequence (y_n) equivalent to (x_n) , so this definition does not depend on the choice of representative for the equivalence class $[(x_n)]$. Since $\|\hat{x}\| = \|x\|$ for any $x \in k$, this definition is compatible with our original $\| \cdot \|$.

We now note that any Cauchy sequence (x_n) in k can be viewed as a Cauchy sequence (\hat{x}_n) in \hat{k} , since we view k as a subfield of \hat{k} , and (\hat{x}_n) obviously converges to $[(x_n)]$ in \hat{k} . Thus every Cauchy sequence in \hat{k} that consists entirely of elements of k converges. But what about other Cauchy sequences in \hat{k} ? To show that these also converge we use the fact that k is dense in \hat{k} .

Definition 7.14. Let S be any subset of a field k with absolute value $\| \cdot \|$. The set S is *dense* in k if for every $x \in k$ and every $\epsilon > 0$ there exists $y \in S$ such that $\|x - y\| < \epsilon$.

Theorem 7.15. *Let k be a field with absolute value $\|\cdot\|$. Then k is dense in its completion \hat{k} .*

Proof. Let $x \in \hat{k}$ be the equivalence class of the Cauchy sequence (x_n) in k . For any $\epsilon > 0$ there is an x_m with the property that $\|x_m - x_n\| < \epsilon$ for all $n \geq m$. It follows that $\|x - \hat{x}_m\| < \epsilon$, where $\hat{x}_m \in k \subseteq \hat{k}$ is just the equivalence class of (x_m, x_m, x_m, \dots) . \square

Corollary 7.16. *Every Cauchy sequence in \hat{k} is equivalent to a Cauchy sequence whose elements lie in k .*

Proof. Let (z_n) be a Cauchy sequence in \hat{k} . Since k is dense in \hat{k} , for each z_n we may pick $x_n \in k \subseteq \hat{k}$ so that $\|z_n - x_n\| < 1/n$. Then for any $\epsilon > 0$ we may pick N such that $\|z_m - x_m\| < \epsilon/3$, $\|z_n - x_n\| < \epsilon/3$ and $\|z_m - z_n\| < \epsilon/3$, for all $m, n \geq N$. It then follows from the triangle inequality that $\|x_m - x_n\| < \epsilon$ for all $m, n \geq N$, hence (x_n) is Cauchy. \square

Corollary 7.17. *The completion \hat{k} of k is complete. Moreover it is the smallest complete field containing k in the following sense: any embedding of k in a complete field k' can be extended to an embedding of \hat{k} into k' .*

Proof. The first statement follows immediately from Corollary 7.16 and the discussion above. For the second, if $\pi: k \rightarrow k'$ is an embedding of k into a complete field k' , then we can extend π to an embedding of \hat{k} into k' by defining

$$\pi([(x_n)]) = \lim_{n \rightarrow \infty} \pi(x_n).$$

Such a limit always exists, since k' is complete, and the map $\pi: \hat{k} \rightarrow k'$ is a ring homomorphism (hence a field embedding) because taking limits commutes with addition and multiplication, by Lemma 7.8. \square

Remark 7.18. We could have defined \hat{k} more categorically as the field with the universal property that every embedding of k into a complete field can be extended to \hat{k} . Assuming it exists, such a \hat{k} is unique up to a canonical isomorphism (map Cauchy sequences to their limits), but we still would have to prove existence.

Finally, we note that the absolute value on the completion of k with respect to $\|\cdot\|$ is nonarchimedean if and only if the absolute value on k is nonarchimedean.

Remark 7.19. Everything we have done here applies more generally to commutative rings. For example, \mathbb{Z}_p is the completion of \mathbb{Z} with respect to the p -adic absolute value $|\cdot|_p$ on \mathbb{Z} , as we will see in the next lecture.

References

- [1] J. S. Milne, [Algebraic number theory](#), 2013.