

20 Class field theory, ray class groups and ray class fields

Having proved the Kronecker-Weber theorem, we are in a position to present the *class field theory* of \mathbb{Q} , which gives an essentially complete description of the abelian extensions K/\mathbb{Q} . We know that every such K corresponds to a subfield of a cyclotomic extension $\mathbb{Q}(\zeta_m)/\mathbb{Q}$, and its Galois group $\text{Gal}(K/\mathbb{Q})$ is thus isomorphic to a quotient of $(\mathbb{Z}/m\mathbb{Z})^\times \simeq \text{Gal}(\mathbb{Q}(\zeta_m)/\mathbb{Q})$; conversely, every quotient H of $(\mathbb{Z}/m\mathbb{Z})^\times$ corresponds to an abelian extension K/\mathbb{Q} that is a subfield of $\mathbb{Q}(\zeta_m)$ for which $H \simeq \text{Gal}(K/\mathbb{Q})$. We would like to make the isomorphism $H \simeq \text{Gal}(K/\mathbb{Q})$ explicit; we can do so via the Artin map that we defined in Lecture 8 which we recall below.

20.1 The Artin map

Let L/K be a finite Galois extension of global fields. For each prime \mathfrak{p} of K , the Galois group $\text{Gal}(L/K)$ acts on the set $\{\mathfrak{q}|\mathfrak{p}\}$ of primes lying above \mathfrak{p} , and for each $\mathfrak{q}|\mathfrak{p}$ the stabilizer of \mathfrak{q} under this action is the decomposition group $D_{\mathfrak{q}}$, which has cardinality $e_{\mathfrak{q}}f_{\mathfrak{q}}$, where $e_{\mathfrak{q}}$ is the ramification index and $f_{\mathfrak{q}}$ is the residue field degree. The decomposition groups $D_{\mathfrak{q}}$ for $\mathfrak{q}|\mathfrak{p}$ are all conjugate, and if L/K is abelian they are all equal and we may unambiguously write $D_{\mathfrak{p}}$ instead of $D_{\mathfrak{q}}$.

Let $\mathbb{F}_{\mathfrak{q}} := \mathcal{O}_L/\mathfrak{q}$ and $\mathbb{F}_{\mathfrak{p}} := \mathcal{O}_K/\mathfrak{p}$ denote the residue fields of \mathfrak{q} and \mathfrak{p} respectively. Then $\mathbb{F}_{\mathfrak{q}}/\mathbb{F}_{\mathfrak{p}}$ is a cyclic Galois extension of degree $f_{\mathfrak{q}}$, and have a surjective homomorphism $\pi_{\mathfrak{q}}: D_{\mathfrak{q}} \rightarrow \text{Gal}(\mathbb{F}_{\mathfrak{q}}/\mathbb{F}_{\mathfrak{p}})$ that sends $\sigma \in D_{\mathfrak{q}}$ to the automorphism $\bar{x} \mapsto \overline{\sigma(x)}$ (here \bar{x} denotes the image of $x \in \mathcal{O}_L$ in $\mathcal{O}_L/\mathfrak{q}$). The kernel of $\pi_{\mathfrak{q}}$ is (by definition) the inertia group $I_{\mathfrak{q}}$, which has cardinality $e_{\mathfrak{q}}$, and we have a short exact sequence

$$1 \longrightarrow I_{\mathfrak{q}} \longrightarrow D_{\mathfrak{q}} \xrightarrow{\pi_{\mathfrak{q}}} \text{Gal}(\mathbb{F}_{\mathfrak{q}}/\mathbb{F}_{\mathfrak{p}}) \longrightarrow 1.$$

When \mathfrak{q} is unramified, $e_{\mathfrak{q}} = 1$ and the inertia group is trivial, so we have an isomorphism $D_{\mathfrak{q}} \simeq \text{Gal}(\mathbb{F}_{\mathfrak{q}}/\mathbb{F}_{\mathfrak{p}})$. The Galois group $\text{Gal}(\mathbb{F}_{\mathfrak{q}}/\mathbb{F}_{\mathfrak{p}})$ is generated by the Frobenius automorphism $x \mapsto x^{\#\mathbb{F}_{\mathfrak{p}}}$, and its inverse image under the isomorphism $\pi_{\mathfrak{q}}$ is the Frobenius element $\sigma_{\mathfrak{q}} \in D_{\mathfrak{q}}$. Equivalently, $\sigma_{\mathfrak{q}}$ is the unique element of $\text{Gal}(L/K)$ for which

$$\sigma_{\mathfrak{q}}(x) \equiv x^{\#\mathbb{F}_{\mathfrak{p}}} \pmod{\mathfrak{q}}$$

for all $x \in \mathcal{O}_L$. The Frobenius elements $\sigma_{\mathfrak{q}}$ for $\mathfrak{q}|\mathfrak{p}$ are all conjugate (they form the Frobenius class $\text{Frob}_{\mathfrak{p}}$), and when L/K is abelian they necessarily coincide, in which case we may write $\sigma_{\mathfrak{p}}$ instead of $\sigma_{\mathfrak{q}}$ (or use $\text{Frob}_{\mathfrak{p}} = \{\sigma_{\mathfrak{p}}\}$ to denote $\sigma_{\mathfrak{p}}$).

For each unramified prime \mathfrak{q} of L the Artin symbol is defined by $\left(\frac{L/K}{\mathfrak{q}}\right) := \sigma_{\mathfrak{q}}$, and the *Artin map* sends each unramified prime \mathfrak{p} of K to the corresponding Frobenius class $\text{Frob}_{\mathfrak{p}}$. When L/K is abelian we may define $\left(\frac{L/K}{\mathfrak{p}}\right) := \sigma_{\mathfrak{p}}$, and if \mathfrak{m} is an \mathcal{O}_K -ideal divisible by all the ramified primes of K (the discriminant, for example), we define the *Artin map*

$$\psi_{L/K}^{\mathfrak{m}}: \mathcal{I}_K^{\mathfrak{m}} \rightarrow \text{Gal}(L/K)$$

$$\prod_{\mathfrak{p} \nmid \mathfrak{m}} \mathfrak{p}^{n_{\mathfrak{p}}} \mapsto \prod_{\mathfrak{p} \nmid \mathfrak{m}} \left(\frac{L/K}{\mathfrak{p}}\right)^{n_{\mathfrak{p}}}.$$

where $\mathcal{I}_K^{\mathfrak{m}}$ is the subgroup of the ideal group \mathcal{I}_K (fractional ideals of \mathcal{O}_K) generated by the primes \mathfrak{p} of K that do not divide \mathfrak{m} ; the ideal \mathfrak{m} is a *modulus* (not a maximal ideal), a term we will define more generally in what follows, and all but finitely many of the exponents $n_{\mathfrak{p}}$ appearing in the products above are zero.

20.2 Class field theory for \mathbb{Q}

We now specialize to the case $K = \mathbb{Q}$. If $L = \mathbb{Q}(\zeta_m)$ is a cyclotomic extension then the primes that ramify in L are precisely the primes that divide m . Each $\sigma \in \text{Gal}(\mathbb{Q}(\zeta_m)/\mathbb{Q})$ is determined by its action on ζ_m , and we have an isomorphism $\omega: \text{Gal}(\mathbb{Q}(\zeta_m)/\mathbb{Q}) \rightarrow (\mathbb{Z}/m\mathbb{Z})^\times$ defined by $\sigma(\zeta_m) = \zeta_m^{\omega(\sigma)}$. For each prime $p \nmid m$ the Frobenius element σ_p is the unique $\sigma \in \text{Gal}(\mathbb{Q}(\zeta_m)/\mathbb{Q})$ for which $\sigma(x) \equiv x^p \pmod{\mathfrak{q}}$ for any (equivalently, all) $\mathfrak{q}|p$. It follows that we must have $\omega(\sigma_p) = p \pmod{m}$.

The Artin map $p \mapsto \sigma_p$ induces an inverse of the isomorphism ω : for each $a \in (\mathbb{Z}/m\mathbb{Z})^\times$ we have $\omega^{-1}(a) = \sigma_p$ for every prime $p \equiv a \pmod{m}$ (of which there are infinitely many such, by Dirichlet's theorem on primes in arithmetic progressions). But in fact the situation is even simpler. We have $a \in (\mathbb{Z}/m\mathbb{Z})^\times$, so $(a) \in \mathcal{I}_{\mathbb{Q}}^m$ and

$$\omega^{-1}(a) = \left(\frac{\mathbb{Q}(\zeta_m)/\mathbb{Q}}{(a)} \right).$$

Now if L is subfield of $\mathbb{Q}(\zeta_m)$, we cannot directly apply ω to $\text{Gal}(L/\mathbb{Q})$, since $\text{Gal}(L/\mathbb{Q})$ is a quotient of $\text{Gal}(\mathbb{Q}(\zeta_m)/\mathbb{Q})$, but for any $a \in (\mathbb{Z}/m\mathbb{Z})^\times$ we still have the Artin map $\mathcal{I}_{\mathbb{Q}}^m \rightarrow \text{Gal}(L/\mathbb{Q})$ available (notice that the modulus m is not specific to the field $\mathbb{Q}(\zeta_m)$, it works perfectly well for any subfield, since any primes that ramify in a subfield will ramify in $\mathbb{Q}(\zeta_m)$ and must divide m). Moreover, the Artin map factors through the quotient map $\text{Gal}(\mathbb{Q}(\zeta_m)/\mathbb{Q}) \rightarrow \text{Gal}(L/\mathbb{Q}) \simeq \text{Gal}(\mathbb{Q}(\zeta_m)/\mathbb{Q})/\text{Gal}(\mathbb{Q}(\zeta_m)/L)$ which is induced by restriction (send each $\sigma \in \text{Gal}(\mathbb{Q}(\zeta_m)/\mathbb{Q})$ to $\sigma|_L \in \text{Gal}(L/\mathbb{Q})$). That is, for any $a \in (\mathbb{Z}/m\mathbb{Z})^\times$ we have

$$\left(\frac{L/\mathbb{Q}}{(a)} \right) = \left(\frac{\mathbb{Q}(\zeta_m)/\mathbb{Q}}{(a)} \right) \Big|_L.$$

To see this, write $L = \mathbb{Q}(\alpha)$ with $\alpha \in \mathcal{O}_L$; then $\alpha \in \mathcal{O}_{\mathbb{Q}(\zeta_m)} = \mathbb{Z}[\zeta_m]$, so $\alpha = f(\zeta_m)$ for some $f \in \mathbb{Z}[x]$. If we pick a prime $p \equiv a \pmod{m}$ and put $\sigma = \left(\frac{L/\mathbb{Q}}{(p)} \right)$ and $\tau = \left(\frac{\mathbb{Q}(\zeta_m)/\mathbb{Q}}{(p)} \right)$, then σ is determined by its action on α and we have

$$\sigma(\alpha) = \sigma(f(\zeta_m)) \equiv_{\mathfrak{q}} f(\zeta_m)^p = f(\zeta_m^p) \equiv_{\mathfrak{q}} f(\tau(\zeta_m)) = \tau(f(\zeta_m)),$$

where the equivalences are modulo any prime $\mathfrak{q}|p$ of L (note that $f(\tau(\zeta_m)) = \tau(f(\zeta_m))$ is conjugate to $\alpha = f(\zeta_m)$ and therefore lies in \mathcal{O}_L so this makes sense). This implies that the Artin map surjects onto $\text{Gal}(L/\mathbb{Q})$, since it surjects onto $\text{Gal}(\mathbb{Q}(\zeta_m)/\mathbb{Q})$ and the quotient map $\text{Gal}(\mathbb{Q}(\zeta_m)/\mathbb{Q}) \rightarrow \text{Gal}(L/\mathbb{Q})$ is surjective.

To sum up, we can now say the following about abelian extensions of $K = \mathbb{Q}$:

- **Existence:** for every modulus m there is an abelian extension $K(m)/K$ ramified only at primes $p|m$ for which $\text{Gal}(K(m)/K) \simeq (\mathbb{Z}/m\mathbb{Z})^\times$ (namely, $K(m) = \mathbb{Q}(\zeta_m)$).
- **Completeness:** every abelian extension of K lies in one of the fields $K(m)$ (this is the Kronecker-Weber theorem).
- **Reciprocity:** for each abelian extension L of K contained in $K(m)$ the Artin map induces a surjective homomorphism from $(\mathbb{Z}/m\mathbb{Z})^\times$ to $\text{Gal}(L/K)$.

All of these statements can be made more precise; in particular, we can refine the first two statements so that the fields are uniquely determined up to isomorphism, and we will give an explicit description of the kernel of the Artin map that allows us to identify $\text{Gal}(L/K)$ with a quotient of $(\mathbb{Z}/m\mathbb{Z})^\times$. But before we do any of this, let us first consider how we to generalize our setup to base fields K other than \mathbb{Q} .

20.3 Moduli and ray class groups

Recall that for a global field K we use M_K to denote its set of places (equivalence classes of absolute values). We generically denote places by the symbol v , but for nonarchimedean places we may use \mathfrak{p} to denote both a prime of K (a nonzero prime ideal of \mathcal{O}_K) and the place corresponding to the absolute value $|\cdot|_{\mathfrak{p}}$. We write $v|\infty$ to indicate that v is an archimedean place, which we recall may be real (arising from embeddings $K \rightarrow \mathbb{R}$) or complex (arising from a conjugate pair of embeddings $K \rightarrow \mathbb{C}$).

Definition 20.1. Let K be a number field. A *modulus* (or *cycle*) \mathfrak{m} for K is a function $M_K \rightarrow \mathbb{Z}_{\geq 0}$ with finite support such that $\mathfrak{m}(v) \leq 1$ for all real v , and $\mathfrak{m}(v) = 0$ for all complex v . We view \mathfrak{m} as a formal product $\prod v^{\mathfrak{m}(v)}$ over M_K , which we may factor as

$$\mathfrak{m} = \mathfrak{m}_0 \mathfrak{m}_{\infty}, \quad \mathfrak{m}_0 := \prod_{\mathfrak{p}/\infty} \mathfrak{p}^{\mathfrak{m}(\mathfrak{p})}, \quad \mathfrak{m}_{\infty} := \prod_{v|\infty} v^{\mathfrak{m}(v)},$$

where \mathfrak{m}_0 corresponds to an \mathcal{O}_K -ideal and \mathfrak{m}_{∞} represents a subset of the real places of K ; we use $\#\mathfrak{m}_{\infty}$ to denote the number of real places in the support of \mathfrak{m} . If \mathfrak{m} and \mathfrak{n} are two moduli for K we say that \mathfrak{m} divides \mathfrak{n} if $\mathfrak{m}(v) \leq \mathfrak{n}(v)$ for all $v \in M_K$ and define $\gcd(\mathfrak{m}, \mathfrak{n})$ and $\text{lcm}(\mathfrak{m}, \mathfrak{n})$ in the obvious way. We use 1 to denote the trivial modulus (the zero function).

We use \mathcal{I}_K to denote the ideal class group of \mathcal{O}_K and introduce the following notation:¹

- a fractional ideal $\mathfrak{a} \in \mathcal{I}_K$ is *prime to \mathfrak{m}* (or *coprime to \mathfrak{m}*) if $v_{\mathfrak{p}}(\mathfrak{a}) = 0$ for all $\mathfrak{p}|\mathfrak{m}_0$.
- $\mathcal{I}_K^{\mathfrak{m}} \subseteq \mathcal{I}_K$ is the subgroup of fractional ideals prime to \mathfrak{m} .
- $K^{\mathfrak{m}} \subseteq K^{\times}$ is the subgroup of elements $\alpha \in K^{\times}$ for which $(\alpha) \in \mathcal{I}_K^{\mathfrak{m}}$.
- $K^{\mathfrak{m},1} \subseteq K^{\mathfrak{m}}$ is the subgroup of elements $\alpha \in K^{\mathfrak{m}}$ for which $v_{\mathfrak{p}}(\alpha - 1) \geq v_{\mathfrak{p}}(\mathfrak{m}_0)$ for $\mathfrak{p}|\mathfrak{m}_0$ and $\alpha_v > 0$ for $v|\mathfrak{m}_{\infty}$ (here $\alpha_v \in \mathbb{R}$ is the image of α under the real-embedding v).
- $\mathcal{P}_K^{\mathfrak{m}} \subseteq \mathcal{I}_K^{\mathfrak{m}}$ is the subgroup of principal fractional ideals $(\alpha) \in \mathcal{I}_K^{\mathfrak{m}}$ with $\alpha \in K^{\mathfrak{m},1}$.

The groups $\mathcal{P}_K^{\mathfrak{m}}$ are sometimes called *rays*, inspired by the example $\mathcal{P}_{\mathbb{Q}}^{\infty}$, which is the set of positive rational numbers.

Definition 20.2. The *ray class group* of K for the modulus \mathfrak{m} is the quotient

$$\text{Cl}_K^{\mathfrak{m}} := \mathcal{I}_K^{\mathfrak{m}} / \mathcal{P}_K^{\mathfrak{m}}.$$

When $\mathfrak{m} = 1$ is the trivial modulus, this is just the usual class group $\text{Cl}_K := \text{cl}(\mathcal{O}_K)$.

Example 20.3. For $K = \mathbb{Q}$ and $\mathfrak{m} = (5)$ we have we have $K^{\mathfrak{m}} = \{a/b : a, b \not\equiv 0 \pmod{5}\}$, $K^{\mathfrak{m},1} = \{a/b : a \equiv b \pmod{5}\}$, and

- $\mathcal{I}_K^{\mathfrak{m}} = \{(1), (1/2), (2), (1/3), (2/3), (3/2), (3), (1/4), (3/4), (4/3), (4), (1/6), (6), \dots\}$.
- $\mathcal{P}_K^{\mathfrak{m}} = \{(1), (2/3), (3/2), (1/4), (4), (6), (1/6), (2/7), (7/2), \dots\}$.

You might not have expected $(2/3) \in \mathcal{P}_K^{\mathfrak{m}}$, but note that $-2/3 \in K^{\mathfrak{m},1}$ and $(-2/3) = (2/3)$. The ray class group $\text{Cl}_K^{\mathfrak{m}} = \mathcal{I}_K^{\mathfrak{m}} / \mathcal{P}_K^{\mathfrak{m}} = \{[(1)], [(2)]\} \simeq (\mathbb{Z}/5\mathbb{Z})^{\times} / \{\pm 1\}$. But for the modulus $\mathfrak{m} = (5)_{\infty}$ we have $\mathcal{P}_K^{\mathfrak{m}} = \{(1), (6), (1/6), (2/7), (7/2), \dots\}$ and $\text{Cl}_K^{\mathfrak{m}} \simeq (\mathbb{Z}/5\mathbb{Z})^{\times}$.

¹This notation varies from author to author; there is unfortunately no universally accepted notation for these objects (in particular, many authors put some but not all of the \mathfrak{m} 's in subscripts). Things will improve when we come to the adelic/idelic formulation of class field theory where there is more consistency.

Lemma 20.4. *Let A be a Dedekind domain and let \mathfrak{a} be an A -ideal. Every ideal class in $\text{cl}(A)$ can be represented by an ideal coprime to \mathfrak{a} .*

Proof. Let $I = \prod_{\mathfrak{p}} \mathfrak{p}^e$ be the unique factorization of some nonzero fractional ideal I of A . We may write $I = I_1 I_2$ with I_1 coprime to $I_2 = \prod_{\mathfrak{p}_i | \mathfrak{a}} \mathfrak{p}_i^{e_i}$. If we choose a uniformizer π_i for each \mathfrak{p}_i and put $\alpha = \prod_{\mathfrak{p}_i | \mathfrak{a}} \pi_i^{-e_i}$ then $[\alpha I] = [I]$ and αI is coprime to \mathfrak{a} . \square

Theorem 20.5. *Let \mathfrak{m} be a modulus for a number field K . We have an exact sequence*

$$1 \longrightarrow \mathcal{O}_K^\times / (\mathcal{O}_K^\times \cap K^{\mathfrak{m},1}) \longrightarrow K^{\mathfrak{m}} / K^{\mathfrak{m},1} \longrightarrow \text{Cl}_K^{\mathfrak{m}} \longrightarrow \text{Cl}_K \longrightarrow 1$$

and a canonical isomorphism

$$K^{\mathfrak{m}} / K^{\mathfrak{m},1} \simeq \{\pm 1\}^{\#\mathfrak{m}_\infty} \times (\mathcal{O}_K / \mathfrak{m}_0)^\times.$$

Proof. Let us consider the composition of the maps $K^{\mathfrak{m},1} \subseteq K^{\mathfrak{m}}$ and $\alpha \mapsto (\alpha)$:

$$K^{\mathfrak{m},1} \xrightarrow{f} K^{\mathfrak{m}} \xrightarrow{g} I_K^{\mathfrak{m}}.$$

The kernel of f is trivial, the kernel of $g \circ f$ is $\mathcal{O}_K^\times \cap K^{\mathfrak{m},1}$ (since $(\alpha) = (1) \iff \alpha \in \mathcal{O}_K^\times$), the kernel of g is \mathcal{O}_K^\times , the cokernel of f is $K^{\mathfrak{m}} / K^{\mathfrak{m},1}$, the cokernel of $g \circ f$ is $\text{Cl}_K^{\mathfrak{m}} = I_K^{\mathfrak{m}} / \mathcal{P}_K^{\mathfrak{m}}$ (by definition), and the cokernel of g is Cl_K (by Lemma 20.4). Applying the snake lemma to the commutative diagram

$$\begin{array}{ccccccc} K^{\mathfrak{m},1} & \xhookrightarrow{f} & K^{\mathfrak{m}} & \longrightarrow & K^{\mathfrak{m}} / K^{\mathfrak{m},1} & \longrightarrow & 1 \\ & & \downarrow g \circ f & & \downarrow g & & \downarrow \\ 1 & \longrightarrow & I_K^{\mathfrak{m}} & \xrightarrow{\sim} & I_K^{\mathfrak{m}} & \longrightarrow & 1 \end{array}$$

yields the kernel-cokernel exact sequence

$$1 \longrightarrow \mathcal{O}_K^\times \cap K^{\mathfrak{m},1} \longrightarrow \mathcal{O}_K^\times \longrightarrow K^{\mathfrak{m}} / K^{\mathfrak{m},1} \longrightarrow \text{Cl}_K^{\mathfrak{m}} \longrightarrow \text{Cl}_K \longrightarrow 1$$

(the last 1 follows from the surjectivity of the last map in bottom row of the diagram above). The desired exact sequence appearing in the theorem follows immediately.

From the unique prime factorization of (α) in \mathcal{I}_K , it is clear that we can write each $\alpha \in K^{\mathfrak{m}}$ as $\alpha = a/b$ with $a, b \in \mathcal{O}_K$ and (a) and (b) coprime to \mathfrak{m}_0 and to each other, with the ideals (a) and (b) uniquely determined by α (even though a and b are not). We now define the homomorphism

$$\begin{aligned} \phi: K^{\mathfrak{m}} &\rightarrow \left(\prod_{v | \mathfrak{m}_\infty} \{\pm 1\} \right) \times (\mathcal{O}_K / \mathfrak{m}_0)^\times \\ \alpha &\mapsto \left(\prod_{v | \mathfrak{m}_\infty} \text{sgn}(\alpha_v) \right) \times (\bar{\alpha}), \end{aligned}$$

where $\bar{\alpha} = \bar{a}\bar{b}^{-1} \in (\mathcal{O}_K / \mathfrak{m}_0)^\times$ (here \bar{a}, \bar{b} are the images of $a, b \in \mathcal{O}_K$ in $\mathcal{O}_K / \mathfrak{m}_0$, and they both lie in $(\mathcal{O}_K / \mathfrak{m}_0)^\times$ because (a) and (b) are prime to \mathfrak{m}_0). The ring $(\mathcal{O}_K / \mathfrak{m}_0)^\times$ is isomorphic to $\prod_{\mathfrak{p} | \mathfrak{m}_0} (\mathcal{O}_K / \mathfrak{p}^{m(\mathfrak{p})})^\times$, by the Chinese remainder theorem, and weak approximation (Theorem 3.26) implies that ϕ is surjective. The kernel of ϕ is clearly $K^{\mathfrak{m},1}$, thus ϕ induces an isomorphism $K^{\mathfrak{m}} / K^{\mathfrak{m},1} \simeq \{\pm 1\}^{\#\mathfrak{m}_\infty} \times (\mathcal{O}_K / \mathfrak{m}_0)^\times$. This isomorphism is canonical, because \bar{a} and \bar{b} (and therefore $\bar{\alpha}$) depend only on the uniquely determined ideals (a) and (b) . \square

Corollary 20.6. Let K be a number field and let \mathfrak{m} be a modulus for K . The ray class group $\text{Cl}_K^{\mathfrak{m}}$ is a finite abelian group whose cardinality $h_K^{\mathfrak{m}} := \#\text{Cl}_K^{\mathfrak{m}}$ is given by

$$h_K^{\mathfrak{m}} = h_K \frac{\phi(\mathfrak{m})}{[\mathcal{O}_K^{\times} : \mathcal{O}_K^{\times} \cap K^{\mathfrak{m},1}]},$$

where $h_K := \#\text{Cl}_K$ and $\phi(\mathfrak{m}) := \#(K^{\mathfrak{m}}/K^{\mathfrak{m},1}) = \phi(\mathfrak{m}_{\infty})\phi(\mathfrak{m}_0)$, with

$$\phi(\mathfrak{m}_{\infty}) = 2^{\#\mathfrak{m}_{\infty}}, \quad \phi(\mathfrak{m}_0) = \#(\mathcal{O}_K/\mathfrak{m}_0)^{\times} = N(\mathfrak{m}_0) \prod_{\mathfrak{p}|\mathfrak{m}_0} (1 - N(\mathfrak{p})^{-1}).$$

In particular, $h_K | h_K^{\mathfrak{m}}$ and $h_K^{\mathfrak{m}} | (h_K \phi(\mathfrak{m}))$.

Explicitly computing the integer $h_{\mathfrak{m}}$ is not a trivial problem, but there are algorithms for doing so; see [1], which considers this problem in detail.

20.4 Polar density

We now want to prove the surjectivity of the Artin map for finite abelian extensions L/K of number fields; as explained in §20.2, we already know this for the case $K = \mathbb{Q}$. In order to do this we want to introduce a new way to measure the density of a set of primes that is defined in terms of a generalization of the Dedekind zeta function.

Definition 20.7. Let K be a number field and let S be a set of primes of K . The *partial Dedekind zeta function* associated to S is the function

$$\zeta_{K,S}(s) := \prod_{\mathfrak{p} \in S} (1 - N(\mathfrak{p})^{-s})^{-1}.$$

(when S contains all primes of K then $\zeta_{K,S}(s)$ is the Dedekind zeta function $\zeta_K(s)$).

The product defining $\zeta_{K,S}(s)$ converges for $\text{Re}(s) > 1$ and thus $\zeta_{K,S}(s)$ is holomorphic and nonzero in this region (since $\zeta_K(s)$ is). It is clear that if S is finite then $\zeta_{K,S}(s)$ is holomorphic (and nonzero) on a neighborhood of 1, and if S contains all but finitely many primes then it differs from $\zeta_K(s)$ by a holomorphic function and thus extends to a meromorphic function with a simple pole at $s = 1$ (by Theorem 18.13).

Between these two extremes the function $\zeta_{K,S}(s)$ may or may not extend to a function that is meromorphic on a neighborhood of 1, but if it does, or more generally, if some power of it does, then we can use the order of the pole at 1 (or the absence of a pole) as a convenient and simple way to measure the density of S .

Definition 20.8. If for some integer $n \geq 1$ the function $\zeta_{K,S}^n$ extends to a meromorphic function on a neighborhood of 1, the *polar density* of S is defined by

$$\rho(S) := \frac{m}{n}, \quad m = -\text{ord}_{s=1} \zeta_{K,S}^n$$

(so m is the order of the pole at $s = 1$, if one is present).

In Lecture 17 we encountered two other notions of density, the *Dirichlet density*

$$d(S) := \lim_{s \rightarrow 1^+} \frac{\sum_{\mathfrak{p} \in S} N(\mathfrak{p})^{-s}}{\sum_{\mathfrak{p}} N(\mathfrak{p})^{-s}} = \lim_{s \rightarrow 1^+} \frac{\sum_{\mathfrak{p} \in S} N(\mathfrak{p})^{-s}}{\log \frac{1}{s-1}},$$

(the equality of the two expressions for $d(S)$ follows from the fact that $\zeta_K(s)$ has a simple pole at $s = 1$, see Problem Set 9), and the *natural density*

$$\delta(S) := \lim_{x \rightarrow \infty} \frac{\#\{\mathfrak{p} \in S : N(\mathfrak{p}) \leq x\}}{\#\{\mathfrak{p} : N(\mathfrak{p}) \leq x\}},$$

and you proved on Problem Set 9 that if S has a natural density then it has a Dirichlet density and the two coincide. We now show that the same is true of the polar density.

Proposition 20.9. *Let S be a set of primes of a number field K . If S has a polar density then it has a Dirichlet density and the two are equal.*

Proof. Suppose S has polar density $\rho(S) = m/n$. Near 1 we can write $\zeta_{K,S}(s)^n$ in the form

$$\zeta_{K,S}(s)^n = \frac{a}{(s-1)^m} \left(1 + \sum_{n>1} a_n (s-1)^n \right),$$

with $a \in \mathbb{C}$ nonzero. We must have $a \in \mathbb{R}_{>0}$, since $\zeta_{K,S}(s) \in \mathbb{R}_{>0}$ for $s \in \mathbb{R}_{>1}$ and therefore $\lim_{s \rightarrow 1} (s-1)^m \zeta_{K,S}(s)^n$ is a positive real number. Taking logs of both sides yields

$$n \sum_{\mathfrak{p} \in S} N(\mathfrak{p})^{-s} \sim m \log \frac{1}{s-1} \quad (\text{as } s \rightarrow 1^+),$$

which implies that S has Dirichlet density $d(S) = m/n$. □

This proposition lets us apply all the properties we already know for Dirichlet density to polar density. Recall that a degree-1 prime in a number field K is a prime with residue field degree 1 over \mathbb{Q} , equivalently, a prime \mathfrak{p} whose absolute norm $[\mathcal{O}_K : \mathfrak{p}]$ is prime.

Proposition 20.10. *Let S and T denote sets of primes in a number field K and let \mathcal{P} be the set of all primes of K .*

- (a) *If S is finite then $\rho(S) = 0$; if $\mathcal{P} - S$ is finite then $\rho(S) = 1$.*
- (b) *If $S \subseteq T$ both have polar densities, then $\rho(S) \leq \rho(T)$.*
- (c) *If two sets S and T have finite intersection and any two of the sets S , T , and $S \cup T$ have polar densities then so does the third and $\rho(S \cup T) = \rho(S) + \rho(T)$.*
- (d) *The polar density of all degree-1 primes is 1 and the polar density of any set of primes is determined by its subset of degree-1 primes.*

Proof. We first note that for any finite set S , the function $\zeta_{K,S}(s)$ is a finite product of nonvanishing entire functions and therefore holomorphic and nonzero everywhere (including at $s = 1$). If the symmetric difference of S and T is finite, then $\zeta_{K,S}(s)f(s) = \zeta_{K,T}(s)g(s)$ for some nonvanishing entire functions $f(s)$ and $g(s)$. Thus if S and T differ by a finite set, then $\rho(S) = \rho(T)$ whenever either set has a polar density

Part (a) follows, since $\rho(\emptyset) = 0$ and $\rho(\mathcal{P}) = 1$ (note that $\zeta_{K,\mathcal{P}}(s) = \zeta_K(s)$, and $\text{ord}_{s=1} \zeta_K(s) = -1$, by Theorem 18.13).

Part (b) follows from the analogous statement for Dirichlet density proved on Problem Set 9, by Proposition 20.9.

For (c) we may assume S and T are disjoint (by the argument above), in which case $\zeta_{K,S \cup T}(s)^n = \zeta_{K,S}(s)^n \zeta_{K,T}(s)^n$ for all $n \geq 1$, and the claim follows.

For (d), let S_1 be the set of all degree-1 primes and let $S_2 = \mathcal{P} - S$, so that $\mathcal{P} = S_1 \sqcup S_2$. Let $n = [K : \mathbb{Q}]$. As $s \rightarrow 1^+$ we may bound $\zeta_{K,S_2}(s)$ from above by $\zeta(2)^n$, so $\zeta_{K,S_2}(s)$ cannot have a pole at $s = 1$; this implies $\rho(S_2) = 0$, provided S_2 has a polar density (we cannot have $\rho(S_2) < 0$ by (a) and (b)). As $s \rightarrow 1^+$ we have $\log \zeta_{K,S_2}(s) \sim \sum_{\mathfrak{p}} N(\mathfrak{p})^{-s}$, and the sum on the RHS converges absolutely by comparison with the series for $\log \zeta(2)^n$: for each prime p there are at most n primes $\mathfrak{p}|p$ in T (in fact at most $n/2$), and for each we have $N(\mathfrak{p}) \geq p^2$. It follows that $\zeta_{K,S_2}(s)$ is holomorphic at $s = 1$, so $\rho(S_2) = 0$, and then $\rho(S) = \rho(\mathcal{P}) - \rho(S_2) = 1$, by (c). The same argument applies to any subset of S_2 , and (c) then implies that the polar density of any set S is determined by its subset of degree-1 primes S' , since $S' - S \subseteq S_2$ has polar density 0; that is, if either S or S' has a polar density then $\rho(S) = \rho(S')$. \square

The proposition implies that the polar densities are rational numbers in $[0, 1]$ whenever they exist. If S and T are sets of primes whose symmetric difference is finite, then either $\rho(S) = \rho(T)$ or neither set has a polar density. Let us write $S \sim T$ to indicate that two sets of primes have finite symmetric difference (this is clearly an equivalence relation), and partially order sets of primes by defining $S \lesssim T \Leftrightarrow S \sim S \cap T$ (all but finitely many of the primes in S lie in T). If S and T have polar densities, then $S \lesssim T$ implies $\rho(S) \leq \rho(T)$.

For a finite Galois extension of number fields L/K , let $\text{Spl}(L/K)$ denote the set of primes of K that split completely in L . When K is clear from context we may just write $\text{Spl}(L)$.

Theorem 20.11. *Let L/K be a Galois extension of number fields of degree n . Then*

$$\rho(\text{Spl}(L)) = 1/n.$$

Proof. Let S be the set of degree-1 primes that split completely in L ; it suffices to show $\rho(S) = 1/n$, by Proposition 20.10. Recall that \mathfrak{p} splits completely in L if and only if both the ramification index $e_{\mathfrak{p}}$ and residue field degree $f_{\mathfrak{p}}$ are equal to 1. Let T be the set of primes $\mathfrak{q}|\mathfrak{p}$ for some $\mathfrak{p} \in S$. For each $\mathfrak{q} \in T$ we have $N_{L/K}(\mathfrak{q}) = \mathfrak{p}^{f_{\mathfrak{p}}} = \mathfrak{p}$, so $N(\mathfrak{q}) = N(\mathfrak{p})$ and \mathfrak{q} is a degree-1 prime. On the other hand, if \mathfrak{q} is an unramified degree-1 prime of L and $\mathfrak{p} = \mathfrak{q} \cap \mathcal{O}_K$, then $N(\mathfrak{q}) = N(N_{L/K}(\mathfrak{q})) = N(\mathfrak{p}^{f_{\mathfrak{p}}})$ is prime, so we must have $f_{\mathfrak{p}} = 1$, and $e_{\mathfrak{p}} = 1$ since \mathfrak{q} is unramified, so \mathfrak{p} splits completely in L . Only finitely many primes ramify, so all but finitely many of the degree-1 primes in L lie in T and therefore $\rho(T) = 1$. Each $\mathfrak{p} \in S$ has exactly n primes $\mathfrak{q} \in T$ lying above it, thus

$$\zeta_{L,T}(s) = \prod_{\mathfrak{q} \in T} (1 - N(\mathfrak{q})^{-s})^{-1} = \prod_{\mathfrak{q} \in T} (1 - N(\mathfrak{p})^{-s})^{-1} = \prod_{\mathfrak{p} \in S} (1 - N(\mathfrak{p})^{-s})^{-n} = \zeta_{K,S}(s)^n.$$

Therefore $\rho(S) = \frac{1}{n} \rho(T) = \frac{1}{n}$ as desired. \square

Corollary 20.12. *If L/K is a finite extension of number fields with Galois closure M/K of degree n , then $\rho(\text{Spl}(L)) = \rho(\text{Spl}(M)) = 1/n$*

Proof. A prime \mathfrak{p} of K splits completely in L if and only if it splits completely in all the conjugates of L in M ; the Galois closure M is the compositum of the conjugates of L , so \mathfrak{p} splits completely in L if and only if it splits completely in M . \square

Corollary 20.13. *Let L/K be a finite Galois extension of number fields with Galois group $G := \text{Gal}(L/K)$ and let H be a normal subgroup of G . The set S of primes for which $\text{Frob}_{\mathfrak{p}} \subseteq H$ has polar density $\rho(S) = \#H/\#G$.*

Proof. Let $F = L^H$; then F/K is Galois (since H is normal) and $\text{Gal}(F/K) \simeq G/H$. For each unramified prime \mathfrak{p} of K , the Frobenius class $\text{Frob}_{\mathfrak{p}}$ lies in H if and only if every $\sigma_{\mathfrak{q}} \in \text{Frob}_{\mathfrak{p}}$ acts trivially on $L^H = F$, which occurs if and only if \mathfrak{p} splits completely in F . By Theorem 20.11, the density of this set of primes is $1/[F : K] = \#H/\#G$. \square

Theorem 20.14. *If L/K and M/K are two Galois extensions of number fields then*

$$\begin{aligned} L \subseteq M &\iff \text{Spl}(M) \lesssim \text{Spl}(L) \iff \text{Spl}(M) \subseteq \text{Spl}(L), \\ L = M &\iff \text{Spl}(M) \sim \text{Spl}(L) \iff \text{Spl}(M) = \text{Spl}(L), \end{aligned}$$

and the map $L \mapsto \text{Spl}(L)$ is an injection from the set of finite Galois extensions of K (inside some fixed \overline{K}) to sets of primes of K that have a positive polar density.

Proof. If $L \subseteq M$ then clearly $\text{Spl}(M) \subseteq \text{Spl}(L)$ and $\text{Spl}(M) \lesssim \text{Spl}(L)$ both hold. It is easy to see that $\text{Spl}(LM) = \text{Spl}(L) \cap \text{Spl}(M)$, so if $\text{Spl}(M) \lesssim \text{Spl}(L)$ then $\text{Spl}(LM) \sim \text{Spl}(M)$. We have $\rho(\text{Spl}(M)) = 1/[M : K]$ and $\rho(\text{Spl}(LM)) = 1/[LM : K]$, by Theorem 20.11, and if $\text{Spl}(LM) \sim \text{Spl}(M)$ then

$$[LM : K] = \rho(\text{Spl}(LM)) = \rho(\text{Spl}(M)) = [M : K],$$

which implies $LM = M$ and therefore $L \subseteq M$. Thus the three conditions in the first line of implications are all equivalent, and this immediately implies the second line of implications. The last statement is clear, since $\text{Spl}(L)$ has positive polar density, by Theorem 20.11. \square

20.5 Ray class fields and Artin reciprocity

As a special case of Corollary 20.12, if F/K is a finite extension of number fields in which all but finitely many primes split completely, then $[F : K] = 1$ and therefore $F = K$. This implies that the Artin map is surjective.

Theorem 20.15. *If L/K is a finite abelian extension of number fields and \mathfrak{m} is a modulus for K divisible by all ramified primes, the Artin map $\psi_{L/K}^{\mathfrak{m}}: \mathcal{I}_K^{\mathfrak{m}} \rightarrow \text{Gal}(L/K)$ is surjective.*

Proof. Let $H \subseteq \text{Gal}(L/K)$ be the image of $\psi_{L/K}^{\mathfrak{m}}$ and let $F = L^H$ be its fixed field. For each prime $\mathfrak{p} \in I_K^{\mathfrak{m}}$ the automorphism $\psi_{L/K}^{\mathfrak{m}}(\mathfrak{p})$ acts trivially on F , which implies that $\psi_{F/K}^{\mathfrak{m}}(\mathfrak{p}) = 1$ and therefore \mathfrak{p} splits completely in F . The group $\mathcal{I}_K^{\mathfrak{m}}$ contains all but finitely many primes \mathfrak{p} of K , so the polar density of the set of primes of K that split completely in F is 1, therefore $[F : K] = 1$ and $H = \text{Gal}(L/K)$. \square

The theorem implies that we have an exact sequence

$$1 \rightarrow \ker \psi_{L/K}^{\mathfrak{m}} \rightarrow \mathcal{I}_K^{\mathfrak{m}} \rightarrow \text{Gal}(L/K) \rightarrow 1.$$

One of the key results of class field theory is that for a suitable choice of the modulus \mathfrak{m} , we have $\mathcal{P}_K^{\mathfrak{m}} \subseteq \ker \psi_{L/K}^{\mathfrak{m}}$. This implies that the Artin map induces an isomorphism between $\text{Gal}(L/K)$ and a quotient of the ray class group $\text{Cl}_K^{\mathfrak{m}} = \mathcal{I}_K^{\mathfrak{m}}/\mathcal{P}_K^{\mathfrak{m}}$. If $\mathcal{P}_K^{\mathfrak{m}} = \ker \psi_{L/K}^{\mathfrak{m}}$, then we have an isomorphism $\text{Gal}(L/K) \simeq \text{Cl}_K^{\mathfrak{m}}$. Such a field L is called the *ray class group* of K for the modulus \mathfrak{m} , denoted $K(\mathfrak{m})$.

When $K = \mathbb{Q}$ the ray class group for $\mathfrak{m} = (m)\infty$ is the cyclotomic field $\mathbb{Q}(\zeta_m)$. For $K \neq \mathbb{Q}$ it is far from obvious that ray class fields exist, but this is indeed the case; this is another key result of class field theory. Once we have a ray class field $K(\mathfrak{m})$, the Artin map

allows us to relate subfields of $K(\mathfrak{m})$ to quotients of the ray class group $\text{Cl}_K^{\mathfrak{m}} \simeq \text{Gal}(K(\mathfrak{m})/K)$ in a way that we will make more precise in the next lecture; this is known as *Artin reciprocity*.

The ray class field for the trivial modulus $\mathfrak{m} = 1$ has a special name; it is called the *Hilbert class field* of K . As we will prove in the next lecture, it is the maximal unramified abelian extension of K (which is the usual way to define the Hilbert class field), and it is unique up to isomorphism. The Hilbert class field L of K has the remarkable property that $\text{Gal}(L/K) \simeq \text{Cl}_K$; it is a Galois extension of K whose Galois group is canonically isomorphic to the class group of K .

References

- [1] Henri Cohen, *Advanced topics in computational number theory*, Springer, 2000.