21 Class field theory: ray class groups and ray class fields

In the previous lecture we proved that every abelian extension L of $\mathbb{Q}$ is contained in a cyclotomic field $\mathbb{Q}(\zeta_m)$. The isomorphism $\mathrm{Gal}(\mathbb{Q}(\zeta_m)/\mathbb{Q}) \simeq (\mathbb{Z}/m\mathbb{Z})^\times$ then allows us to view $\mathrm{Gal}(L/\mathbb{Q})$ as a quotient of $(\mathbb{Z}/m\mathbb{Z})^\times$. We would like to replace the base field $\mathbb{Q}$ with an arbitrary number field K, but we need analogs of the cyclotomic fields $\mathbb{Q}(\zeta_m)$ and the abelian Galois groups $(\mathbb{Z}/m\mathbb{Z})^\times$. These analogs are *ray class fields*, and their Galois groups are isomorphic to *ray class groups*. Ray class fields are not, in general, cyclotomic extensions of K, their construction is rather more complicated. Before defining them, let us first recall some properties of the Artin map we defined in Lecture 7.

21.1 The Artin map

Let L/K be a finite Galois extension of global fields, and let $\mathfrak{p}$ be a prime of K. Recall that the Galois group $\mathrm{Gal}(L/K)$ acts on the set $\{\mathfrak{q}|\mathfrak{p}\}$ (primes $\mathfrak{q}$ of L lying above $\mathfrak{p}$) and the stabilizer of $\mathfrak{q}|\mathfrak{p}$ is the decomposition group $D_\mathfrak{q} \subseteq \mathrm{Gal}(L/K)$. By Proposition 7.9, we have surjective homomorphism

$$\begin{aligned}\pi_\mathfrak{q}\colon D_\mathfrak{q} &\to \mathrm{Gal}(\mathbb{F}_\mathfrak{q}/\mathbb{F}_\mathfrak{p})\\ \sigma &\mapsto \overline{\sigma} := (\overline{a} \mapsto \overline{\sigma(a)}),\end{aligned}$$

where $a \in \mathcal{O}_L$ is any lift of $\overline{a} \in \mathbb{F}_\mathfrak{q} := \mathcal{O}_L/\mathfrak{q}$ to $\mathcal{O}_L$ and $\overline{\sigma(a)}$ is the reduction of $\sigma(a) \in \mathcal{O}_L$ to $\mathbb{F}_\mathfrak{q}$; kernel of $\pi_\mathfrak{q}$ is the inertia group $I_\mathfrak{q}$. If $\mathfrak{q}$ is unramified then $I_\mathfrak{q}$ is trivial and $\pi_\mathfrak{q}$ is an isomorphism. The *Artin symbol* (Definition 7.18) is defined by

$$\left(\frac{L/K}{\mathfrak{q}}\right) := \sigma_\mathfrak{q} := \pi_\mathfrak{q}^{-1}(x \mapsto x^{\#\mathbb{F}_\mathfrak{p}}),$$

where $(x \mapsto x^{\#\mathbb{F}_\mathfrak{p}}) \in \mathrm{Gal}(\mathbb{F}_\mathfrak{q}/\mathbb{F}_\mathfrak{p})$ is the Frobenius automorphism, a canonical generator for the cyclic group $\mathrm{Gal}(\mathbb{F}_\mathfrak{q}/\mathbb{F}_\mathfrak{p})$. Equivalently, $\sigma_\mathfrak{q}$ is the unique element of $\mathrm{Gal}(L/K)$ for which

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

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

$$\left(\frac{L/K}{\mathfrak{p}}\right) := \sigma_\mathfrak{p}.$$

Now assume L/K is abelian, let $\mathfrak{m}$ be an $\mathcal{O}_K$-ideal divisible by every ramified prime of K, and let $\mathcal{I}_K^\mathfrak{m}$ denote the subgroup of fractional ideals $I \in \mathcal{I}_K$ for which $v_\mathfrak{p}(I) = 0$ for all $\mathfrak{p}|\mathfrak{m}$. The Artin map (Definition 7.21) is the homomorphism

$$\begin{aligned}\psi_{L/K}^\mathfrak{m}\colon \mathcal{I}_K^\mathfrak{m} &\to \mathrm{Gal}(L/K)\\ \prod_{\mathfrak{p}\nmid\mathfrak{m}} \mathfrak{p}^{n_\mathfrak{p}} &\mapsto \prod_{\mathfrak{p}\nmid\mathfrak{m}} \left(\frac{L/K}{\mathfrak{p}}\right)^{n_\mathfrak{p}}.\end{aligned}$$

A key ingredient of class field theory (which we will prove in this lecture) is that the Artin map $\psi_{L/K}^\mathfrak{m}$ is surjective. We can then identify $\mathrm{Gal}(L/K)$ with the quotient of $\mathcal{I}_K^\mathfrak{m}/\ker\psi_{L/K}^\mathfrak{m}$.

The primes $\mathfrak{p} \in \ker^{\mathfrak{m}}_{L/K}$ are unramified and have the property that the Frobenius elements $\sigma_{\mathfrak{q}}$ for $\mathfrak{q}|\mathfrak{p}$ are all trivial, meaning that the residue field extension $\mathbb{F}_{\mathfrak{q}}/\mathbb{F}_{\mathfrak{p}}$ is trivial; this means that $\mathfrak{p}$ *splits completely* in L (it is unramified and primes above it have residue degree one). Conversely, every prime $\mathfrak{p} \in \mathcal{I}^{\mathfrak{m}}_K$ that splits completely in L lies in $\ker \psi^{\mathfrak{m}}_{L/K}$.

Proposition 21.1. *Let $K \subseteq L \subseteq M$ be a tower of finite abelian extension of global fields and let $\mathfrak{m}$ be an $\mathcal{O}_K$-ideal divisible by all primes $\mathfrak{p}$ of K that ramify in M. We have a commutative diagram*

$$\begin{array}{ccc} \mathcal{I}^{\mathfrak{m}}_K & \xrightarrow{\psi^{\mathfrak{m}}_{M/K}} & \mathrm{Gal}(M/K) \\ & \searrow^{\psi^{\mathfrak{m}}_{L/K}} & \downarrow \mathrm{res} \\ & & \mathrm{Gal}(L/K) \end{array}$$

where the vertical map is the homomorphism $\sigma \to \sigma_{|_L}$ induced by restriction.

Proof. It suffices to check commutativity at primes $\mathfrak{p} \nmid \mathfrak{m}$, which are necessarily unramified. The proposition then follows from Proposition 7.20. ☐

21.2 Class field theory for $\mathbb{Q}$

We now specialize to $K = \mathbb{Q}$. The Kronecker-Weber theorem tells us that every abelian extension L/K lies in a cyclotomic field $\mathbb{Q}(\zeta_m)$. Each $\sigma \in \mathrm{Gal}(\mathbb{Q}(\zeta_m)/\mathbb{Q})$ is determined by its action on ζ_m, and we have an isomorphism

$$\omega\colon \mathrm{Gal}(\mathbb{Q}(\zeta_m)/\mathbb{Q}) \xrightarrow{\sim} (\mathbb{Z}/m\mathbb{Z})^\times$$

defined by $\sigma(\zeta_m) = \zeta_m^{\omega(\sigma)}$. The primes p that ramify in $\mathbb{Q}(\zeta_m)$ are precisely those that divide m (by Corollary 10.20). For each prime $p \not| m$ the Frobenius element σ_p is the unique $\sigma \in \mathrm{Gal}(\mathbb{Q}(\zeta_m)/\mathbb{Q})$ for which $\sigma(x) \equiv x^p \bmod \mathfrak{q}$ for any (equivalently, all) $\mathfrak{q}|(p)$. Thus $\omega(\sigma_p) = p \bmod m$, and it follows that the Artin map induces an inverse isomorphism $(\mathbb{Z}/m\mathbb{Z})^\times \to \mathrm{Gal}(\mathbb{Q}(\zeta_m)/\mathbb{Q})$: for every integer a coprime to m we have $(a) \in \mathcal{I}^m_{\mathbb{Q}}$ and

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

where $\bar{a} = a \bmod m$. As you showed on Problem Set 4, the surjectivity of the Artin map follows immediately, since a ranges over all integers coprime to m.

Now let L be a subfield of $\mathbb{Q}(\zeta_m)$. We cannot apply ω to $\mathrm{Gal}(L/\mathbb{Q})$, since $\mathrm{Gal}(L/\mathbb{Q})$ is a quotient of $\mathrm{Gal}(\mathbb{Q}(\zeta_m)/\mathbb{Q})$, not a subgroup, but the Artin map $\mathcal{I}^m_{\mathbb{Q}} \to \mathrm{Gal}(L/\mathbb{Q})$ is available; notice that the modulus m works for L as well as $\mathbb{Q}(\zeta_m)$, since any primes that ramify in L also ramify in $\mathbb{Q}(\zeta_m)$ and therefore divide m. By Proposition 21.1, the Artin map factors through the surjective homomorphism $\mathrm{Gal}(\mathbb{Q}(\zeta_m)/\mathbb{Q}) \to \mathrm{Gal}(L/\mathbb{Q})$ induced by restriction and thus induces a surjective homomorphism $(\mathbb{Z}/m\mathbb{Z})^\times \to \mathrm{Gal}(L/\mathbb{Q})$.

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

- **Existence**: for each integer m we have a *ray class field* $\mathbb{Q}(\zeta_m)$: an abelian extension ramified only at $p|m$ with Galois group isomorphic to the *ray class group* $(\mathbb{Z}/m\mathbb{Z})^\times$.
- **Completeness**: every abelian extension of $\mathbb{Q}$ lies in a ray class field $\mathbb{Q}(\zeta_m)$.

- **Reciprocity**: if L is an abelian extension of $\mathbb{Q}$ contained in the ray class field $\mathbb{Q}(\zeta_m)$, the Artin map $\mathcal{I}_\mathbb{Q}^m \to \mathrm{Gal}(L/\mathbb{Q})$ induces a surjective homomorphism from the ray class group $(\mathbb{Z}/m\mathbb{Z})^\times$ to $\mathrm{Gal}(L/\mathbb{Q})$, letting us view $\mathrm{Gal}(L/\mathbb{Q})$ as a quotient of $(\mathbb{Z}/m\mathbb{Z})^\times$.

All of these statements will be made more precise; in particular, we will refine the first two statements so that ray class fields are uniquely determined by the modulus m, and we will give an explicit description of the kernel of the Artin map that allows us to identify $\mathrm{Gal}(L/\mathbb{Q})$ with a quotient of $(\mathbb{Z}/m\mathbb{Z})^\times$. But let us first consider how to generalize these statements to number fields other than $\mathbb{Q}$ and define the terms *ray class field*, and *ray class group*. In order to do so, we first need to make the role of the integer m more precise by introducing the notion of a *modulus*.

21.3 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 finite places, those arising from a discrete valuation associated to a prime $\mathfrak{p}$ of K (a nonzero prime ideal of $\mathcal{O}_K$), we may write $\mathfrak{p}$ in place of v. We write $v|\infty$ to indicate that v is an infinite place (one not arising from a prime of K); recall that when K is a number field all infinite places are archimedean, and they may be real ($K_v \simeq \mathbb{R}$) or complex ($K_v \simeq \mathbb{C}$).

Definition 21.2. Let K be a number field. A *modulus* (or *cycle*) $\mathfrak{m}$ for K is a function $M_K \to \mathbb{Z}_{\geq 0}$ with finite support such that for $v|\infty$ we have $\mathfrak{m}(v) \leq 1$ with $\mathfrak{m}(v) = 0$ unless v is a real place. We view $\mathfrak{m}$ as a formal product $\prod v^{\mathfrak{m}(v)}$ over M_K, which we may factor as

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

where $\mathfrak{m}_0$ is an $\mathcal{O}_K$-ideal and $\mathfrak{m}_\infty$ 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 $\mathrm{lcm}(\mathfrak{m}, \mathfrak{n})$ in the obvious way. The zero function is the *trivial modulus*, with $\mathfrak{m}_0 = (1)$ and $\#\mathfrak{m}_\infty = 0$. We use $\mathcal{I}_K$ to denote the ideal class group of $\mathcal{O}_K$ and define the following notation:[1]

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

The groups $\mathcal{R}_K^\mathfrak{m}$ are called *rays* or *ray groups*.

Definition 21.3. Let $\mathfrak{m}$ be a modulus for a number field K. The *ray class group* for the modulus $\mathfrak{m}$ is the quotient

$$\mathrm{Cl}_K^\mathfrak{m} := \mathcal{I}_K^\mathfrak{m}/\mathcal{R}_K^\mathfrak{m}.$$

[1] This 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.

A finite abelian extension L/K that is unramified at all places[2] not in the support of $\mathfrak{m}$ for which the kernel of the Artin map $\psi^{\mathfrak{m}}_{L/K}\colon \mathcal{I}^{\mathfrak{m}}_K \to \operatorname{Gal}(L/K)$ is equal to the ray group $\mathcal{R}^{\mathfrak{m}}_K$ is a *ray class field* for the modulus $\mathfrak{m}$.

When $\mathfrak{m}$ is the trivial modulus, the ray class group is the same as the usual class group $\mathrm{Cl}_K := \mathrm{cl}(\mathcal{O}_K)$, but in general the class group Cl_K is a quotient of the ray class group $\mathrm{Cl}^{\mathfrak{m}}_K$ (as we will prove shortly). While not immediately apparent from the definition, we will see that ray class fields are uniquely determined by $\mathfrak{m}$, so it makes sense to speak of *the* ray class field for the modulus $\mathfrak{m}$ (assuming existence).

Remark 21.4. The definitions above make sense for any global field, but in our ideal-theoretic treatment of class field theory we will mostly restrict our attention to number fields. Our adelic/idelic formulation of class field theory will address all global fields.

Remark 21.5. If $\mathfrak{m}(v) = 1$ for every real place v of K then $\mathrm{Cl}^{\mathfrak{m}}_K$ is a *narrow ray class group*. The narrow ray class group with $\mathfrak{m}_0 = (1)$ is the *narrow class group*; the usual class group $\mathrm{Cl}_K = \mathrm{cl}\,\mathcal{O}_K$ is sometimes called the *wide class group* to distinguish the two. Note that the wide class group is a quotient of the narrow class group, thus smaller in general; this terminology can be confusing, but the thing to remember is that narrow equivalence is *stronger* than ordinary equivalence, so there are *more* narrow equivalence classes, in general. Of course for number fields with no real places (imaginary quadratic fields, in particular) there is no distinction.

Example 21.6. For $K = \mathbb{Q}$ with the modulus $\mathfrak{m} = (5)$ we have $K^{\mathfrak{m}} = \{a/b : a, b \not\equiv 0 \bmod 5\}$ and $K^{\mathfrak{m},1} = \{a/b : a \equiv b \not\equiv 0 \bmod 5\}$. Thus

$$\begin{aligned}\mathcal{I}^{\mathfrak{m}}_K &= \{(1), (1/2), (2), (1/3), (2/3), (3/2), (3), (1/4), (3/4), (4/3), (4), (1/6), (6), \ldots\},\\ \mathcal{R}^{\mathfrak{m}}_K &= \{(1), (2/3), (3/2), (1/4), (4), (6), (1/6), (2/7), (7/2), \ldots\}.\end{aligned}$$

You might not have expected $(2/3) \in \mathcal{R}^{\mathfrak{m}}_K$, since $2/3 \notin K^{\mathfrak{m},1}$, but note that $-2/3 \in K^{\mathfrak{m},1}$ and $(-2/3) = (2/3)$. The ray class group is

$$\mathrm{Cl}^{\mathfrak{m}}_K = \mathcal{I}^{\mathfrak{m}}_K/\mathcal{R}^{\mathfrak{m}}_K = \{[(1)], [(2)]\} \simeq (\mathbb{Z}/5\mathbb{Z})^\times/\{\pm 1\},$$

which is isomorphic to the Galois group of the totally real subfield $\mathbb{Q}(\zeta_5)^+$ of $\mathbb{Q}(\zeta_5)$, which is the ray class field for this modulus. If we change the modulus to $\mathfrak{m} = (5)\infty$ we instead get $\mathcal{R}^{\mathfrak{m}}_K = \{(1), (6), (1/6), (2/7), (7/2), \ldots\}$, $\mathrm{Cl}^{\mathfrak{m}}_K \simeq (\mathbb{Z}/5\mathbb{Z})^\times$, and the ray class field is $\mathbb{Q}(\zeta_5)$.

Lemma 21.7. *Let A be a Dedekind domain and let $\mathfrak{a}$ be an A-ideal. Every ideal class in $\mathrm{cl}(A)$ contains an A-ideal coprime to $\mathfrak{a}$.*

Proof. Let I be a nonzero fractional ideal of A. For each prime $\mathfrak{p}|\mathfrak{a}$ we can pick $\pi_\mathfrak{p} \in \mathfrak{p}$ such that $v_\mathfrak{q}(\pi_\mathfrak{p}) = v_\mathfrak{q}(\mathfrak{p})$ for all $\mathfrak{q}|\mathfrak{a}$, by Corollary 3.21. If we then put $\alpha := \prod_{\mathfrak{p}|\mathfrak{a}} \pi_\mathfrak{p}^{-v_\mathfrak{p}(I)}$, then $v_\mathfrak{p}(\alpha I) = 0$ for all $\mathfrak{p}|\mathfrak{a}$; thus αI is coprime to $\mathfrak{a}$ and $[\alpha I] = [I]$.

Now let S be the finite set of primes $\mathfrak{p}$ for which $v_\mathfrak{p}(\alpha I) < 0$ and pick $\pi_\mathfrak{p} \in \mathfrak{p}$ such that $v_\mathfrak{q}(\pi_\mathfrak{p}) = v_\mathfrak{q}(\mathfrak{p})$ for all $\mathfrak{q} \in S$ and $\mathfrak{q}|\mathfrak{a}$ (again using Corollary 3.21). If we now put $a := \prod_{\mathfrak{p}\in S} \pi_\mathfrak{p}^{-v_\mathfrak{p}(\alpha I)} \in A$, then $v_\mathfrak{p}(a\alpha I) \geq 0$ for all $\mathfrak{p}$ and $v_\mathfrak{p}(a\alpha I) = 0$ for all $\mathfrak{p}|\mathfrak{a}$. Thus $a\alpha I$ is an A-ideal coprime to $\mathfrak{a}$ and $[a\alpha I] = [I]$. ☐

[2]A real place v of K is unramified in L if every place of L above v is also a real place. But if L is unramified at all $\mathfrak{p} \nmid \mathfrak{m}_0$ (necessary for $\psi^{\mathfrak{m}}_{L/K}$ to be defined), and if $\ker \psi^{\mathfrak{m}}_{L/K} = \mathcal{R}^{\mathfrak{m}}_K$, then L will necessarily be unramified at all $v|\mathfrak{m}_\infty$; so in the definition it is enough for L to be unramified away from $\mathfrak{m}_0$.

Theorem 21.8. *Let $\mathfrak{m}$ be a modulus for a number field K. We have an exact sequence*

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

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 $\alpha \mapsto (\alpha)$:

$$K^{\mathfrak{m},1} \xrightarrow{f} K^{\mathfrak{m}} \xrightarrow{g} \mathcal{I}_K^{\mathfrak{m}}.$$

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 $\mathrm{Cl}_K^{\mathfrak{m}} = \mathcal{I}_K^{\mathfrak{m}}/\mathcal{R}_K^{\mathfrak{m}}$ (by definition), and the cokernel of g is Cl_K (by Lemma 21.7). Applying the snake lemma (see [2, Lemma 5.13], for example) to the following commutative diagram with exact rows

$$\begin{array}{ccccccccc}
1 & \longrightarrow & K^{\mathfrak{m},1} & \overset{f}{\hookrightarrow} & K^{\mathfrak{m}} & \longrightarrow & K^{\mathfrak{m}}/K^{\mathfrak{m},1} & \longrightarrow & 1 \\
 & & \downarrow{\scriptstyle g\circ f} & & \downarrow{\scriptstyle g} & & \downarrow{\scriptstyle \pi} & & \\
1 & \longrightarrow & \mathcal{I}_K^{\mathfrak{m}} & \xrightarrow{\sim} & \mathcal{I}_K^{\mathfrak{m}} & \longrightarrow & 1 & &
\end{array}$$

yields the exact sequence $\ker g \circ f \to \ker g \to \ker \pi \to \operatorname{coker} g \circ f \to \operatorname{coker} g \to \operatorname{coker} \pi$:

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

where the initial 1 follows from the fact that f is injective (and $\ker \pi = \operatorname{coker} f$).

We can write each $\alpha \in K^{\mathfrak{m}}$ as $\alpha = a/b$ with $a, b \in \mathcal{O}_K$ such that (a) and (b) are coprime to $\mathfrak{m}_0$ and to each other. The ideals (a) and (b) are uniquely determined by α, since $a/b = a'/b' \Rightarrow ab' = a'b \Rightarrow (a)(b') = (a')(b)$, and since (a) and (b) are coprime we must have $(a) = (a')$ and $(b) = (b')$ (by unique factorization of ideals).

We now define the homomorphism

$$\varphi\colon K^{\mathfrak{m}} \to \left(\prod_{v|\mathfrak{m}_\infty} \{\pm 1\}\right) \times (\mathcal{O}_K/\mathfrak{m}_0)^\times$$
$$\alpha \mapsto \left(\prod_{v|\mathfrak{m}_\infty} \operatorname{sgn}(\alpha_v)\right) \times (\bar{\alpha}),$$

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 coprime to $\mathfrak{m}_0$). The ring $(\mathcal{O}_K/\mathfrak{m}_0)^\times$ is isomorphic to $\prod_{\mathfrak{p}|\mathfrak{m}_0}(\mathcal{O}_K/\mathfrak{p}^{\mathfrak{m}(\mathfrak{p})})^\times$, by the Chinese remainder theorem, and weak approximation (Theorem 8.5) implies that φ is surjective. The kernel of φ is clearly $K^{\mathfrak{m},1}$, thus φ induces an isomorphism $K^{\mathfrak{m}}/K^{\mathfrak{m},1} \simeq \{\pm\}^{\#\mathfrak{m}_\infty} \times (\mathcal{O}_K/\mathfrak{m}_0)^\times$. This isomorphism is canonical, because $\bar{\alpha}$ depends only on the uniquely determined ideals (a) and (b) (if we replace a with $a' = au$ for some $u \in \mathcal{O}_K^\times$ we must replace b with $b' = bu$). ∎

Corollary 21.9. *Let K be a number field and let $\mathfrak{m}$ be a modulus for K. The ray class group $\mathrm{Cl}_K^{\mathfrak{m}}$ is a finite abelian group whose cardinality $h_K^{\mathfrak{m}} := \#\mathrm{Cl}_K^{\mathfrak{m}}$ is given by*

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

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

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

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

Proof. The exact sequence implies $\phi(\mathfrak{m})/[\mathcal{O}_K^\times : \mathcal{O}_K^\times \cap K^{\mathfrak{m},1}] = h_K^{\mathfrak{m}}/h_K$, and that both sides of this equality are integers. □

Computing the ray class number $h_K^{\mathfrak{m}}$ is not a trivial problem, but there are algorithms for doing so; see [1], which considers this problem in detail.

21.4 Polar density

We now want to prove the surjectivity of the Artin map for finite abelian extensions L/K of number fields (as noted in §21.2, we already know this for $K = \mathbb{Q}$). In order to do so we first 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. Throughout this section and the next, all number fields are assumed to lie in some fixed algebraic closure of $\mathbb{Q}$.

Definition 21.10. 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 complex function

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

which converges to a holomorphic function on $\mathrm{Re}(s) > 1$ (by the same argument we used for $\zeta_K(s)$ in Lecture 18).

If S is finite then $\zeta_{K,S}(s)$ is certainly holomorphic (and nonzero) on a neighborhood of 1. If S contains all but finitely many primes of K then it differs from $\zeta_K(s)$ by a holomorphic factor and therefore extends to a meromorphic function with a simple pole at $s = 1$, by Theorem 19.12.

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) to measure the density of S.

Definition 21.11. If for some integer $n \ge 1$ the function $\zeta_{K,S}^n$ extends to a meromorphic function on a neighborhood of 1, the *polar density* of S is defined by

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

(so m is the order of the pole at $s = 1$, if one is present). Note that if $\zeta_{K,S}^{n_1}$ and $\zeta_{K,s}^{n_2}$ both extend to a meromorphic function on a neighborhood of 1 then we necessarily have

$$n_2\mathrm{ord}_{s=1}\zeta_{K,S}^{n_1}(s) = \mathrm{ord}_{s=1}\zeta_{K,S}^{n_1n_2} = n_1\mathrm{ord}_{s=1}\zeta_{K,S}^{n_2},$$

which implies that $\rho(S)$ does not depend on the choice of n. We will show below that (whenever it is defined) $\rho(S)$ is a rational number in the interval $[0, 1]$.

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

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

(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*

$$\delta(S) := \lim_{x\to\infty} \frac{\#\{\mathfrak{p}\in S : \mathrm{N}(\mathfrak{p}) \le x\}}{\#\{\mathfrak{p} : \mathrm{N}(\mathfrak{p}) \le x\}}.$$

On Problem Set 9 you proved 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 21.12. *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. In particular, $\rho(S) \in [0, 1]$ whenever it is defined.*

Proof. Suppose S has polar density $\rho(S) = m/n$. By taking the Laurent series expansion of $\zeta_{K,S}^n(s)$ at $s = 1$ and factoring out the leading nonzero term we can write

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

for some $a \in \mathbb{C}^\times$. 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 $\lim_{s\to 1^+}(s-1)^m \zeta_{K,S}(s)^n$ is a positive real number. Taking logs of both sides yields

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

which implies that S has Dirichlet density $d(S) = m/n$ (note that $\log(a) = O(1)$ plays no role, since $-m\log(s-1) \to \infty$ as $s \to 1^+$). ☐

Corollary 21.13. *Let S be a set of primes of a number field K. If S has both a polar density and a natural density then the two coincide.*

We should note that not every set of primes with a natural density has a polar density, since the later is always a rational number while the former need not be.

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 $\mathrm{N}(\mathfrak{p}) = [\mathcal{O}_K : \mathfrak{p}] = \#\mathbb{F}_{\mathfrak{p}}$ is prime.

Proposition 21.14. *Let S and T denote sets of primes in a number field K, let $\mathcal{P}$ be the set of all primes of K, and let $\mathcal{P}_1$ be the set of degree-1 primes of K. The following hold:*

(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 $\rho(S) \le \rho(T)$.*

(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) *We have $\rho(\mathcal{P}_1) = 1$, and $\rho(S \cap \mathcal{P}_1) = \rho(S)$ whenever S has a polar density.*

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 functions $f(s)$ and $g(s)$ holomorphic on $\mathbb{C}$. 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 $\mathrm{ord}_{s=1}\zeta_K(s)=-1$, by Theorem 19.12).

Part (b) follows from the analogous statement for Dirichlet density proved on Problem Set 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 $\mathcal{P}_2:=\mathcal{P}-\mathcal{P}_1$ so that $\mathcal{P}=\mathcal{P}_1\sqcup\mathcal{P}_2$. For each rational prime p there are at most $n:=[K:\mathbb{Q}]$ (in fact $n/2$) primes $\mathfrak{p}|p$ in $\mathcal{P}_2$, each of which has absolute norm $\mathrm{N}(\mathfrak{p})\geq p^2$. It follows by comparison with $\zeta(2s)^n$ that the product defining $\zeta_{K,S_2}(s)$ converges absolutely to a holomorphic function on $\mathrm{Re}(s)>1/2$ and is therefore holomorphic (and nonvanishing, since it is an Euler product) on a neighborhood of 1; thus $\rho(\mathcal{P}_2)=0$ and $\rho(\mathcal{P}_1)=1$. We therefore have $\rho(S\cap\mathcal{P}_2)=0$, so $\rho(S)=\rho(S\cap\mathcal{P}_1)$ whenever $\rho(S)$ exists, by (c). □

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

Theorem 21.15. *Let L/K be a Galois extension of number fields of degree n. Then*

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

Proof. Let S be the set of degree-1 primes of K that split completely in L; it suffices to show $\rho(S)=1/n$, by Proposition 21.14. 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 $\mathfrak{q}$ of L that lie above some $\mathfrak{p}\in S$. For each $\mathfrak{q}\in T$ lying above $\mathfrak{p}\in S$ we have $\mathrm{N}_{L/K}(\mathfrak{q})=\mathfrak{p}^{f_\mathfrak{p}}=\mathfrak{p}$, so $\mathrm{N}(\mathfrak{q})=\mathrm{N}(\mathrm{N}_{L/K}(\mathfrak{q}))=\mathrm{N}(\mathfrak{p})$, thus $\mathfrak{q}$ is a degree-1 prime, since $\mathfrak{p}$ is.

On the other hand, if $\mathfrak{q}$ is any unramified degree-1 prime of L and $\mathfrak{p}=\mathfrak{q}\cap\mathcal{O}_K$, then $N(\mathfrak{q})=\mathrm{N}(\mathrm{N}_{L/K}(\mathfrak{q}))=\mathrm{N}(\mathfrak{p}^{f_\mathfrak{p}})$ is prime, so we must have $f_\mathfrak{p}=1$, and $e_\mathfrak{p}=1$ since $\mathfrak{q}$ is unramified, which implies that $\mathfrak{p}$ is a degree-1 prime that splits completely in L and is thus an element of S. Only finitely many primes ramify, so all but finitely many of the degree-1 primes in L lie in T, thus $\rho(T)=1$, by Proposition 21.14. Each $\mathfrak{p}\in S$ has exactly n primes $\mathfrak{q}\in T$ lying above it (since $\mathfrak{p}$ splits completely), and we have

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

It follows that $\rho(S)=\frac{1}{n}\rho(T)=\frac{1}{n}$ as desired. □

Corollary 21.16. *If L/K is a finite extension of number fields with Galois closure M/K of degree n, then $\rho(\mathrm{Spl}(L))=\rho(\mathrm{Spl}(M))=1/n$*

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. □

Corollary 21.17. *Let L/K be a finite Galois extension of number fields with Galois group $G:=\mathrm{Gal}(L/K)$ and let H be a normal subgroup of G. The set S of primes for which $\mathrm{Frob}_\mathfrak{p}\subseteq H$ has polar density $\rho(S)=\#H/\#G$.*

Proof. Let $F = L^H$; then F/K is Galois (since H is normal) and $\mathrm{Gal}(F/K) \simeq G/H$. For each unramified prime $\mathfrak{p}$ of K, the Frobenius class $\mathrm{Frob}_\mathfrak{p}$ lies in H if and only if every $\sigma_\mathfrak{q} \in \mathrm{Frob}_\mathfrak{p}$ acts trivially on $L^H = F$, which occurs if and only if $\mathfrak{p}$ splits completely in F. By Theorem 21.15, the density of this set of primes is $1/[F:K] = \#H/\#G$. □

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 $S \precsim T \Leftrightarrow S \sim S \cap T$ (in other words, $S - T$ is finite). If S and T have polar densities, then $S \precsim T$ implies $\rho(S) \leq \rho(T)$, by Proposition 21.14.

Theorem 21.18. *If L/K and M/K are two finite Galois extensions of number fields then*

$$L \subseteq M \iff \mathrm{Spl}(M) \precsim \mathrm{Spl}(L) \iff \mathrm{Spl}(M) \subseteq \mathrm{Spl}(L),$$
$$L = M \iff \mathrm{Spl}(M) \sim \mathrm{Spl}(L) \iff \mathrm{Spl}(M) = \mathrm{Spl}(L),$$

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

Proof. The implications $L \subseteq M \Rightarrow \mathrm{Spl}(M) \subseteq \mathrm{Spl}(L) \Rightarrow \mathrm{Spl}(M) \precsim \mathrm{Spl}(L)$ are clear, so it suffices to show that $\mathrm{Spl}(M) \precsim \mathrm{Spl}(L) \Rightarrow L \subseteq M$.

A prime $\mathfrak{p}$ of K splits completely in the compositum LM if and only if it splits completely in both L and M: the forward implication is clear and for the reverse, note that if $\mathfrak{p}$ splits completely in both L and M then it certainly splits completely in $L \cap M$, so we may assume $K = L \cap M$; we then have $\mathrm{Gal}(LM/K) \simeq \mathrm{Gal}(L/K) \times \mathrm{Gal}(M/K)$, and if the decomposition subgroups of all primes above $\mathfrak{p}$ are trivial in both $\mathrm{Gal}(L/K)$ and $\mathrm{Gal}(M/K)$ then the same applies in $\mathrm{Gal}(LM/K)$. Thus $\mathrm{Spl}(LM) = \mathrm{Spl}(L) \cap \mathrm{Spl}(M)$.

It follows that $\mathrm{Spl}(M) \precsim \mathrm{Spl}(L) \Rightarrow \mathrm{Spl}(LM) \sim \mathrm{Spl}(M)$. By Theorem 21.15, we have $\rho(\mathrm{Spl}(M)) = 1/[M:K]$ and $\rho(\mathrm{Spl}(LM) = 1/[LM:K]$, thus $\mathrm{Spl}(LM) \sim \mathrm{Spl}(M)$ implies

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

in which case $LM = M$ and $L \subseteq M$. This proves $\mathrm{Spl}(M) \precsim \mathrm{Spl}(L) \Rightarrow L \subseteq M$, so the three conditions in the first line of biconditionals are all equivalent, and this immediately implies the second line of biconditionals. The last statement of the theorem is clear, since $\mathrm{Spl}(L)$ has positive polar density, by Theorem 21.15. □

21.5 Ray class fields and Artin reciprocity

As a special case of Corollary 21.16, 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 $F = K$. We will use this fact to prove that the Artin map is surjective.

Theorem 21.19. *Let L/K be a finite abelian extension of number fields and let $\mathfrak{m}$ be a modulus for K that is divisible by all primes of K that ramify in L. Then the Artin map $\psi^\mathfrak{m}_{L/K} \colon \mathcal{I}^\mathfrak{m}_K \to \mathrm{Gal}(L/K)$ is surjective.*

Proof. Let $H \subseteq \mathrm{Gal}(L/K)$ be the image of $\psi^\mathfrak{m}_{L/K}$ and let $F = L^H$ be its fixed field, which we note is a Galois extension of K, since H is normal (because $\mathrm{Gal}(L/K)$ is abelian). For each prime $\mathfrak{p} \in \mathcal{I}^\mathfrak{m}_K$ the automorphism $\psi^\mathfrak{m}_{L/K}(\mathfrak{p}) \in H$ acts trivially on $F = L^H$, 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. Thus $[F:K]=1$ and $H=\operatorname{Gal}(L/K)$, by Corollary 21.16. □

We now show that the kernel of the Artin map $\psi_{L/K}^{\mathfrak{m}}$ uniquely determines the field L.

Theorem 21.20. *Let $\mathfrak{m}$ be a modulus for a number field K and let L and M be finite abelian extensions of K unramified at all primes not in the support of $\mathfrak{m}$. If $\ker\psi_{L/K}^{\mathfrak{m}}=\ker\psi_{M/K}^{\mathfrak{m}}$ then $L=M$. In particular, ray class fields are unique whenever they exist.*

Proof. Let S be the set of primes of K that do not divide $\mathfrak{m}$. Each prime $\mathfrak{p}$ in S is unramified in both L and M, and $\mathfrak{p}$ splits completely in L (resp. M) if and only if it lies in the kernel of $\psi_{L/K}^{\mathfrak{m}}$ (resp. $\psi_{M/K}^{\mathfrak{m}}$). If $\ker\psi_{L/K}^{\mathfrak{m}}=\ker\psi_{M/K}^{\mathfrak{m}}$ then

$$\operatorname{Spl}(L)\sim(S\cap\ker\psi_{L/K}^{\mathfrak{m}})=(S\cap\ker\psi_{M/K}^{\mathfrak{m}})\sim\operatorname{Spl}(M),$$

and therefore $L=M$, by Theorem 21.18. □

Theorem 21.19 implies that we have an exact sequence

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

One of the key results of class field theory is that for a suitable choice of the modulus $\mathfrak{m}$, we have $\mathcal{R}_K^{\mathfrak{m}}\subseteq\ker\psi_{L/K}^{\mathfrak{m}}$. This implies that the Artin map induces an isomorphism between $\operatorname{Gal}(L/K)$ and a quotient of the ray class group $\operatorname{Cl}_K^{\mathfrak{m}}=\mathcal{I}_K^{\mathfrak{m}}/\mathcal{R}_K^{\mathfrak{m}}$. When L is the ray class field for the modulus $\mathfrak{m}$, the Artin map allows us to relate subfields of L to quotients of the ray class group $\operatorname{Cl}_K^{\mathfrak{m}}\simeq\operatorname{Gal}(L/K)$ in a way that we will make more precise in the next lecture; this is known as *Artin reciprocity*.

References

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

[2] Allen Altman and Steven Kleiman, *A term of commutative algebra*, Worldwide Center of Mathematics, 2013.