8 Adèles

Adèles and idèles were introduced in the early 20th century as an approach to class field theory, which may be viewed as a vast generalization of quadratic reciprocity. First we introduce the notion of adèles $\mathbb{A}_\mathbb{Q}$ and idèles $\mathbb{A}_\mathbb{Q}^\times$ over $\mathbb{Q}$. Then we will discuss $\mathfrak{p}$-adic numbers for arbitrary number fields K, and use these to define the adèles $\mathbb{A}_K$ and idèles $\mathbb{A}_K^\times$ of a number field K. One defines the quotient group $\mathbb{A}_K^\times/K^\times$ to be the *idèle class group* of K. Our goal will be to show that this is essentially the ideal class group of K, and then use this to describe some of the main results of class field theory.

8.1 $\mathbb{A}_\mathbb{Q}$

The **adèles** of $\mathbb{Q}$ are the ring

$$\mathbb{A}_\mathbb{Q} = \left\{ (\alpha_2, \alpha_3, \alpha_5, \ldots ; \alpha_\infty) \in \prod_p \mathbb{Q}_p \times \mathbb{R} : \alpha_p \in \mathbb{Z}_p \text{ for a.a. } p \right\}.$$

Here "for a.a. (almost all)" p means for all but finitely many primes p.

Exercise 8.1. *With addition and multiplication defined component-wise, show $\mathbb{A}_\mathbb{Q}$ is a ring.*

Note that $\mathbb{A}_\mathbb{Q}$ puts together the information one gets from all the completions of $\mathbb{Q}$, but the whole direct product $\prod \mathbb{Q}_p \times \mathbb{R}$ is too large to work with by itself, so we only consider sequences where almost all terms are p-adic integers. This is analogous to an infinite direct sum of vector spaces V_i. Specifically, $\bigoplus_{i=1}^\infty V_i = \{(v_i) \in \prod V_i : v_i = 0 \text{ for a.a. } i\}$. For instance if each $V_i = \mathbb{R}$, then a basis for $\bigoplus V_i$ is $\{e_i\}$ where $e_i = (0, \ldots, 0, 1, 0, \ldots)$ is the vector with a 1 in the i-th coordinate and 0's elsewhere. If one removes the "for almost all i" condition, then $(1, 1, 1, \ldots) = e_1 + e_2 + e_3 + \cdots$ would be in the direct sum, but this is not a finite linear combination of basis elements.

Let us simplify now our notation slightly.

We call a nontrivial absolute value on $\mathbb{Q}$ a **place** of $\mathbb{Q}$. Hence the places of $\mathbb{Q}$ are $|\cdot|_v$ where v is either a prime p or $v = \infty$. The places $v = p$ are called **finite places**, and the place $v = \infty$ is called the **real place** or **infinite place**.[†] Let $\mathbb{Q}_v$ denote the completion of $\mathbb{Q}$ w.r.t. $|\cdot|_v$, so $\mathbb{Q}_p$ still denotes $\mathbb{Q}_p$ and now $\mathbb{Q}_\infty$ denotes $\mathbb{R}$. Let $\mathbb{Z}_v$ denote the set $\{x_v \in \mathbb{Q}_v : |x_v|_v \le 1\}$, so that $\mathbb{Z}_v = \mathbb{Z}_p$ if $v = p$ and $\mathbb{Z}_\infty = [-1, 1]$. While $\mathbb{Z}_p$ is the completion of $\mathbb{Z}$ in each $\mathbb{Q}_p$, $\mathbb{Z}_\infty$ admittedly has little to do with $\mathbb{Z}$. Nevertheless this notation is convenient. We also remark that $\mathbb{Z}_v$ is compact inside each $\mathbb{Q}_v$. In fact when $v < \infty$, i.e., $v = p$, $\mathbb{Z}_v$ is open in $\mathbb{Q}_v$.

Now we can denote the adèles as

$$\mathbb{A}_\mathbb{Q} = \left\{ (\alpha_v) \in \prod \mathbb{Q}_v : \alpha_v \in \mathbb{Z}_v \text{ for a.a. } v \right\}.$$

While the condition $\alpha_v \in \mathbb{Z}_v$ for a.a. v may at first glance look stronger than the condition $\alpha_p \in \mathbb{Z}_p$ for a.a. p because $v = \infty$ is allowed, thinking about it for a second shows they are equivalent. (Think about it for a second: $\alpha \in \mathbb{A}_\mathbb{Q}$ means the local components α_v can lie outside of $\mathbb{Z}_v$ only for v in some finite set S of places (S of course depends on α— it is like the "support" of an element

[†] It is standard to call places primes and still use the letter p, since they correspond to the usual primes and infinity. Then the ordinary primes are called finite primes, and denoted by $p < \infty$, and the infinite place is called the infinite prime $p = \infty$.

y in the infinite direct sum $R^\infty = \bigoplus_{i=1}^\infty \mathbb{R}$, which is the set of i for which the i-th component of y is nonzero. We can always add the place ∞ to S if it is not included, and S will still be finite, meaning we have the condition α_v can lie outside of $\mathbb{Z}_v$ for $v \in S \cup \{\infty\}$.)

The following examples will tell us a little bit about $\mathbb{A}_\mathbb{Q}$.

Example 8.1.1. *Let $x = \frac{a}{b} \in \mathbb{Q}$ with $a, b \in \mathbb{Z}$. Let $\alpha = (x, x, x, \ldots)$. Note that $\frac{1}{b} \in \mathbb{Z}_p$ for any p s.t. $p \nmid b$. Hence $x = \frac{a}{b} \in \mathbb{Z}_p$ for any p s.t. $p \nmid b$, i.e., $x = \alpha_v \in \mathbb{Z}_v$ for almost all v. Thus $\alpha \in \mathbb{A}_\mathbb{Q}$.*

Hence we have an (injective) ring homomorphism from $\mathbb{Q} \to \mathbb{A}_\mathbb{Q}$ given by

$$x \mapsto (x, x, x, \ldots).$$

So the additive identity of $\mathbb{A}_\mathbb{Q}$ is $(0, 0, 0, \ldots)$ and the multiplicative identity of $\mathbb{A}_\mathbb{Q}$ is $1 = (1, 1, 1, \ldots)$. We will typically identify elements of $\mathbb{Q}$ with their image in $\mathbb{A}_\mathbb{Q}$ under this map.

Example 8.1.2. *Let $\alpha = (1, 0, 0, 0, \ldots), \beta = (0, 1, 1, 1, \ldots)$. Since each component $\alpha_v, \beta_v \in \mathbb{Z}_v$ for all v, we have $\alpha, \beta \in \mathbb{A}_\mathbb{Q}$. Then $\alpha\beta = (0, 0, 0, \ldots) = 0 \in \mathbb{A}_\mathbb{Q}$. In other words, α and β are zero divisors in $\mathbb{A}_\mathbb{Q}$, so $\mathbb{A}_\mathbb{Q}$ is not an integral domain.*

Proposition 8.1.3. *For $\alpha = (\alpha_v) \in \mathbb{A}_\mathbb{Q}$, let $|\alpha| = \prod_v |\alpha_v|_v$. Then $|\cdot| : \mathbb{A}_\mathbb{Q} \to \mathbb{R}$ satisfies*

(i) $|\alpha| \geq 0$

(ii) $|\alpha\beta| = |\alpha||\beta|$

Proof. First note that $|\alpha|$ is well defined: since $\alpha = (\alpha_v) \in \mathbb{A}_\mathbb{Q}$ satisfies $\alpha_v \in \mathbb{Z}_v$ for almost all v, we have $|\alpha_v|_v \leq 1$ for almost all v, and therefore the infinite product $|\alpha| = \prod_v |\alpha|_v$ converges. It is clear that $|\alpha| \geq 0$.

Further (ii) follows immediately because it does for each $|\alpha|_v$. □

Example 8.1.4. *Taking $\alpha = (1, 0, 0, 0, \ldots)$ from the previous example, we see $|\alpha| = |1|_2 \prod_{v \neq 2} |0|_v = 0$, so $|\cdot|$ can be zero on nonzero elements. Therefore, $|\cdot|$ cannot technically be an absolute value. Of course, our earlier definition of absolute values was only for integral domains because any multiplicative homomorphism $|\cdot| : R \to \mathbb{R}$ for a non-integral domain must be 0 on some zero divisors. ($\alpha\beta = 0$ implies $|\alpha||\beta| = |\alpha\beta| = 0$ so either $|\alpha|$ or $|\beta|$ is 0.)*

However we can even find $\alpha \in \mathbb{A}_\mathbb{Q}$ which is not a zero divisor such that $|\alpha| = 0$. Namely consider $\alpha = (\alpha_v)$ where $\alpha_p = p$ and $\alpha_\infty = 1$. Each component $\alpha_v \in \mathbb{Z}_v$ so $\alpha \in \mathbb{A}_\mathbb{Q}$, but

$$\alpha = \prod_p |p|_p \cdot |1|_\infty = \prod_p \frac{1}{p} = 0.$$

In fact, another crucial property of absolute values fails also, namely the triangle inequality.

Exercise 8.2. *Find $\alpha, \beta \in \mathbb{A}_\mathbb{Q}$ such that $|\alpha + \beta| > |\alpha| + |\beta|$.*

Example 8.1.5. *Let $x \in \mathbb{Q}$ and $\alpha = (x, x, x, \ldots)$. If $x = 0$, then $|\alpha| = 0$. Otherwise $|\alpha| = 1$ by Exercise 6.11.*

The fact that $\mathbb{A}_\mathbb{Q}$ is not an integral domain makes it a little hard to work with, but the **idèles** $\mathbb{A}_\mathbb{Q}^\times = \{\alpha \in \mathbb{A}_\mathbb{Q} : \alpha \text{ invertible}\}$, namely the multiplicative subgroup of $\mathbb{A}_\mathbb{Q}$, become a nice space to work with.

Proposition 8.1.6. *The idèle group $\mathbb{A}_\mathbb{Q}^\times = \{(\alpha_v) \in \prod_v \mathbb{Q}_v^\times : \alpha_v \in \mathbb{Z}_v^\times \text{ for a.a. } v\}$.*

Note that technically we did not define $\mathbb{Z}_v^\times$ for $v = \infty$, but as above including or removing a single place $v = \infty$ from a "for all but finitely many" condition does not change anything. However, if one wishes, one can set $\mathbb{Z}_v^\times = \{\alpha_v \in \mathbb{Q}_v^\times : |\alpha_v|_v = 1\}$ so $\mathbb{Z}_v = \mathbb{Z}_p$ for $v = p$ and $\mathbb{Z}_v = \{-1, 1\}$ for $v = \infty$.

Proof. Let $\alpha = (\alpha_v) \in \mathbb{A}_\mathbb{Q}^\times$. It is clear that $\alpha_v \in \mathbb{Q}_v^\times$ for all v, otherwise some component will be zero.

Let $\beta = (\beta_v) = \alpha^{-1} \in \mathbb{A}_\mathbb{Q}$. Since $\beta_v \in \mathbb{Z}_v$ and $\alpha_v \in \mathbb{Z}_v$ for almost all v, there is a finite set S of places v such that $\alpha_v, \beta_v \in \mathbb{Z}_v$ for all $v \notin S$. Then $\alpha\beta = (1, 1, 1, \ldots)$ means $\alpha_v\beta_v = 1$ for all v, so $\alpha_v, \beta_v \in \mathbb{Z}_v^\times$ for all $v \notin S$, i.e., $\alpha_v \in \mathbb{Z}_v^\times$ for a.a. v.

This proves $\subseteq$. $\supseteq$ is straightforward—see the next exercise. □

Exercise 8.3. *Let $\alpha = (\alpha_v) \in \mathbb{A}_\mathbb{Q}$ such that $\alpha_v \neq 0$ for all v and $\alpha_v \in \mathbb{Z}_v$ for almost all v. Show there is a $\beta \in \mathbb{A}_\mathbb{Q}$ such that $\alpha\beta = 1 = (1, 1, 1, \ldots)$.*

One can use the topologies on $\mathbb{Q}_v$ and $\mathbb{Q}_v^\times$ to define topologies on the additive group of adèles and multiplicative group of idèles, to make them both into topological groups. (We already defined the topology on $\mathbb{Q}_v$ in terms of a basis of neighborhoods. One can do the same thing for $\mathbb{Q}_v^\times$, or just give $\mathbb{Q}_v^\times$ the subspace topology from $\mathbb{Q}_v^\times \subseteq \mathbb{Q}_v$. Both methods give the same topology.) To define a topology on a group, it suffices to specify a basis of open neighborhoods of the identity.

A basis of open neighborhoods of 0 in $\mathbb{A}_\mathbb{Q}$ is given by a collection of sets of the form

$$\prod_{v \in S} U_v \prod_{v \notin S} \mathbb{Z}_v \subseteq \mathbb{A}_\mathbb{Q}$$

where S is a finite set of places containing ∞ and for each $v \in S$, U_v is an open neighborhood of 0 in $\mathbb{Q}_v$. Note the requirement that $\infty \in S$ is because $\mathbb{Z}_\infty = [-1, 1]$ is a closed set in $\mathbb{Q}_\infty = \mathbb{R}$, so we do not want $v = \infty$ occurring in the product on the right.

Similarly, a basis of open neighborhoods of 1 in $\mathbb{A}_\mathbb{Q}^\times$ is given by a collections of sets of the form

$$\prod_{v \in S} U_v \prod_{v \notin S} \mathbb{Z}_v^\times \subseteq \mathbb{A}_\mathbb{Q}$$

where S is a finite set of places containing ∞ and for each $v \in S$, U_v is an open neighborhood of 1 in $\mathbb{Q}_v^\times$.

We remark that one can also form a topology of $\mathbb{A}_\mathbb{Q}$ by taking the product topology on $\prod \mathbb{Q}_v$, and put the subspace topology on $\mathbb{A}_\mathbb{Q}$. This is different than the topology we described above, and this topology induced by the product topology is too strong for our purposes. Similar remarks are true for the topology on $\mathbb{A}_\mathbb{Q}^\times$. Further, the topology on $\mathbb{A}_\mathbb{Q}^\times$ is *not* the subspace topology induced from the inclusion $\mathbb{A}_\mathbb{Q}^\times \subseteq \mathbb{A}_\mathbb{Q}$, as the open sets in the subspace topology will be too large. For example,

Exercise 8.4. *Consider the open set $U = \mathbb{R} \times \prod_p \mathbb{Z}_p \subseteq \mathbb{A}_\mathbb{Q}$. This is an open neighborhood of 1 in $\mathbb{A}_\mathbb{Q}$. Show the restriction $U \cap \mathbb{A}_\mathbb{Q}^\times$ contains but does not equal the open neighborhood $V = \mathbb{R}^\times \times \prod_p \mathbb{Z}_p^\times$ of 1 in $\mathbb{A}_\mathbb{Q}^\times$.*

A similar, but slightly more technical, argument shows that if $V = \prod_{v \in S} U_v \prod_{v \notin S} \mathbb{Z}_v^\times$ is an open neighborhood of 1 in $\mathbb{A}_\mathbb{Q}$ (where as usual S is a finite set of places), then there is no open neighborhood U of 1 in $\mathbb{A}_\mathbb{Q}$ whose restriction to $\mathbb{A}_\mathbb{Q}^\times$ will be contained in V.

Proposition 8.1.7. *$\mathbb{A}_\mathbb{Q}$ and $\mathbb{A}_\mathbb{Q}^\times$ are locally compact. Furthermore,*

(i) a subset U of $\mathbb{A}_\mathbb{Q}$ is relatively compact if and only if $U \subseteq \prod_v K_v$ where each K_v is compact in $\mathbb{Q}_v$ and $K_v = \mathbb{Z}_v$ for almost all v; and

(ii) a subset U of $\mathbb{A}_\mathbb{Q}^\times$ is relatively compact if and only if $U \subseteq \prod_v K_v$ where each K_v is compact in $\mathbb{Q}_v^\times$ and $K_v = \mathbb{Z}_v^\times$ for almost all v.

(Recall a set is called *relatively compact* if its closure is compact.)

Proposition 8.1.8. *$\mathbb{Q}$ and $\mathbb{Q}^\times$ are discrete subgroups of $\mathbb{A}_\mathbb{Q}$ and $\mathbb{A}_\mathbb{Q}^\times$.*

Proposition 8.1.9. *$\mathbb{A}_\mathbb{Q}/\mathbb{Q}$ is compact with the quotient topology, and*

$$\mathbb{A}_\mathbb{Q}/\mathbb{Q} \simeq \varprojlim_n \mathbb{R}/n\mathbb{Z} = \{(a_1, a_2, a_3, a_4, \ldots) : a_n \in \mathbb{R}/n\mathbb{Z},\ a_n \in a_m + m\mathbb{Z} \text{ if } m|n\}.$$

Thus just like $\mathbb{Z}_p = \varprojlim \mathbb{Z}/p^n\mathbb{Z}$, we can view $\mathbb{A}_\mathbb{Q}/\mathbb{Q}$ as a projective limit, i.e., as a way of putting together all the $\mathbb{R}/n\mathbb{Z}$'s in a compatible way.

See [Ramakrishnan–Valenza] for proofs.

8.2 $\mathfrak{p}$-adic fields

There are several ways to treat the theory of $\mathfrak{p}$-adic fields, just like there are several ways to treat the theory of p-adic numbers. One common way of defining them is as finite extensions of $\mathbb{Q}_p$. However, I will opt for a concrete approach via completions w.r.t. absolute values, as it is in my mind more natural.

Let K be a number field and $\mathfrak{p}$ a prime ideal of K.

In the case $K = \mathbb{Q}$ and $\mathfrak{p} = (p)$, one defines the p-adic absolute value by $|x|_p = p^{-m}$ where $x = p^m \frac{a}{b}$ and a, b are relatively prime to p. If $x = p^m a \in \mathbb{Z}$, another way to say this is that $|x|_p = p^{-m}$ where m is the highest power of p that divides x, i.e., m is the unique integer such that

$$\mathfrak{p}^m = (p)^m \supseteq (x) \not\subseteq \mathfrak{p}^{m+1} = (p)^{m+1}.$$

In fact, using fractional ideals, we can say the same thing even if $x \notin \mathbb{Z}$. Specifically, we have a filtration of $\mathbb{Q}$:

$$\cdots \supseteq \mathfrak{p}^{-2} \supseteq \mathfrak{p}^{-1} \supseteq \mathfrak{p}^0 = \mathbb{Z} \supseteq \mathfrak{p}^1 \supseteq \mathfrak{p}^2 \supseteq \cdots$$

We define $\mathrm{ord}_p(x)$ to be the largest m such that x (or (x) if you prefer) is contained in $\mathfrak{p}^m = (p)^m$, and then set $|x|_p = p^{-\mathrm{ord}_p(x)}$.

Now we return to the general case.

Definition 8.2.1. *Let K be a number field and $\mathfrak{p}$ be a prime ideal of K. For $x \in K$, define the* **$\mathfrak{p}$-adic valuation** *$\mathrm{ord}_\mathfrak{p}(x)$ to be the largest integer m such that $x \in \mathfrak{p}^m$. Then the* **$\mathfrak{p}$-adic absolute value** *on K is given by $|x|_\mathfrak{p} = N(\mathfrak{p})^{-\mathrm{ord}_\mathfrak{p}(x)}$.*

Exercise 8.5. *Let $\mathfrak{p}$ be a prime ideal of K lying above a prime p of $\mathbb{Q}$. Then for $x \in \mathbb{Q} \subseteq K$, show $|x|_\mathfrak{p} = |x|_p^f$ where $f = f(\mathfrak{p}|p)$ is the inertial degree of $\mathfrak{p}$ above p.*

As in the case $K = \mathbb{Q}$, these give, up to equivalence, all non-archimedean absolute values on K. The archimedean values are slightly more complicated then the case of $\mathbb{Q}$, and they are essentially parametrized by $\mathrm{Gal}(K/\mathbb{Q})$. This is because if we want to restrict the usual absolute value on $\mathbb{R}$ or

$\mathbb{C}$ to K, it depends upon the embedding of K into $\mathbb{R}$ or $\mathbb{C}$, and the embeddings of K into $\mathbb{C}$ was precisely our definition for the Galois group $\mathrm{Gal}(K/\mathbb{Q})$.

Let $\{\sigma_1, \ldots \sigma_r\}$ denote the set of real embeddings of K and $\{\tau_1, \overline{\tau}_1, \ldots, \tau_s, \overline{\tau}_s\}$ denote the set of complex embeddings of K. Then we define the archimedean absolute values

$$|x|_{\sigma_i} = |\sigma_i(x)|_{\mathbb{R}}$$

and

$$|x|_{\tau_j} = |\tau_j(x)|_{\mathbb{C}}$$

where $|\cdot|_{\mathbb{R}}$ is the usual absolute value on $\mathbb{R}$ and $|z|_{\mathbb{C}} = z\overline{z}$ is the *square* of the usual absolute value on $\mathbb{C}$. It is immediate from the definition of equivalence that $\Big| \, cdot|_{\mathbb{C}}$ is equivalent to the usual absolute value on $\mathbb{C}$, but it is preferable for us to take this normalization, as $|z|_{\mathbb{C}}$ is more like a norm than $\sqrt{|\cdot|_{\mathbb{C}}}$. In particular it maps $\mathbb{Z}[i]$ into $\mathbb{Z}$. For instance $|1+i|_{\mathbb{C}} = (1+i)(1-i) = 2$, but if we use the usual absolute value, then $|1+i| = \sqrt{2}$.

Example 8.2.2. *Let $K = \mathbb{Q}(\sqrt{3})$. Then $\mathrm{Gal}(K/\mathbb{Q}) = \{\sigma_1, \sigma_2\}$ where $\sigma_1(\sqrt{3}) = \sqrt{3}$ and $\sigma_2(\sqrt{3}) = -\sqrt{3}$. Then*

$$|1+\sqrt{3}|_{\sigma_1} = |1+\sqrt{3}|_{\mathbb{R}} \neq |1-\sqrt{3}|_{\mathbb{R}} = |1+\sqrt{3}|_{\sigma_2}.$$

Example 8.2.3. *Let $K = \mathbb{Q}(\sqrt{-3})$. Then $\mathrm{Gal}(K/\mathbb{Q}) = \{\tau, \overline{\tau}\}$ where $\tau(\sqrt{-3}) = i\sqrt{3}$. Then*

$$|1+\sqrt{-3}|_{\tau} = |1+i\sqrt{3}|_{\mathbb{C}} = |1-i\sqrt{3}|_{\mathbb{C}} = |1+\sqrt{-3}|_{\overline{\tau}}.$$

In general, no absolute values corresponding to two different σ_i's or τ_j's will be equivalent, but we will always have $|\cdot|_{\tau_j} = |\cdot|_{\overline{\tau}_j}$ since $|z|_{\mathbb{C}} = |\overline{z}|_{\mathbb{C}}$.

Theorem 8.2.4. *The places (equivalence classes of non-trivial absolute values) on K are precisely given by*

(i) $v = \mathfrak{p}$ where $\mathfrak{p}$ is a prime ideal of K (non-archimedean places)

(ii) $v = \sigma_i$ where σ_i is a real embedding of K (real places)

(iii) $v = \tau_j$ where τ_j runs over the set of complex embedding of K, up to complex conjugation (complex places).

Definition 8.2.5. *For a place v of K, let K_v denote the completion of K with respect to $|\cdot|_v$. Let $\mathcal{O}_{K_v} = \{x \in K_v : |x|_v \leq 1\}$ and $\mathcal{O}_{K_v}^{\times} = \{x \in K_v : |x|_v = 1\}$.*

If $v = \mathfrak{p}$ is a non-archimedean place, we call K_v the **$\mathfrak{p}$-adic numbers** *and $\mathcal{O}_{K_v}$ the* **$\mathfrak{p}$-adic integers** *over K.*

If $K = \mathbb{Q}$ and $\mathfrak{p} = (p)$ this coincides with our previous definitions. Most of the results on p-adic numbers over $\mathbb{Q}$ extend to $\mathfrak{p}$-adic numbers over K, but due to time constraints we will not explain these in detail except where we need to. A similar remark is true for the theory of adèles and idèles, which we can now define.

Definition 8.2.6. *The* **adèles** *of K are*

$$\mathbb{A}_K = \{\alpha = (\alpha_v) : \alpha_v \in K_v \text{ for all } v,\ \alpha_v \in \mathbb{Z}_v \text{ for a.a. } v\} \subset \prod_v K_v.$$

Similarly, the **idèles** *of K are*

$$\mathbb{A}_K^{\times} = \left\{\alpha = (\alpha_v) : \alpha_v \in K_v^{\times} \text{ for all } v,\ \alpha_v \in \mathbb{Z}_v^{\times} \text{ for a.a. } v\right\} \subset \prod_v K_v^{\times}.$$

In both statements, v runs over the set of places of K.

As with $\mathbb{Q}$, K and $K^\times$ embed diagonally into $\mathbb{A}_K$ and $\mathbb{A}_K^\times$, and in this way we will regard K and $K^\times$ as additive and multiplicative subgroups of $\mathbb{A}_K$ and $\mathbb{A}_K^\times$. One defines the absolute value on $\mathbb{A}_K$ or $\mathbb{A}_K^\times$ by

$$|\alpha| = |\alpha|_{\mathbb{A}_K} = \prod_v |\alpha_v|_v$$

where $\alpha = (\alpha_v) \in \mathbb{A}_K$. We will usually simply denote the absolute value on $\mathbb{A}_K$ by $|\cdot|$, but sometimes we will use $|\cdot|_{\mathbb{A}_K}$ for clarity.

Definition 8.2.7. *The* **idèle class group** *of K is $C_K = \mathbb{A}_K^\times / K^\times$.*

While $\mathbb{A}_K/K$ is compact, it is not the case that $\mathbb{A}_K^\times/K^\times$ is, owing to the fact that the open sets of $\mathbb{A}_K^\times$ are much smaller than the open sets of $\mathbb{A}_K$ (see above remarks about the difference between the topologies on $\mathbb{A}_\mathbb{Q}$ and $\mathbb{A}_\mathbb{Q}^\times$). However one can prove the following.

Theorem 8.2.8. *Let $\mathbb{A}_K^1 \subset \mathbb{A}_K^\times$ be the subgroup of idèles of K having absolute value 1. Then $K^\times \subseteq \mathbb{A}_K^1$ and the* **norm 1 idèle class group**

$$C_K^1 = \mathbb{A}_K^1 / K^\times$$

is compact.

Recall in the case $K = \mathbb{Q}$, you proved in Exercise 6.11 that the adèlic absolute value $|(x, x, x, \ldots)|_{\mathbb{A}_\mathbb{Q}} = 1$ for $x \in \mathbb{Q}^\times$. A similar argument shows that $|x|_{\mathbb{A}_K} = |(x, x, x, \ldots)|_{\mathbb{A}_K} = 1$ for any $x \in K^\times$, which means $K^\times \subseteq \mathbb{A}_K^1$.

Definition 8.2.9. *The* **∞-idèles** *are defined to be*

$$\mathbb{A}_{K,\infty}^\times = \left\{(\alpha_v) \in \mathbb{A}_K^\times : \alpha_v \in \mathcal{O}_{K_v}^\times \text{ for all } v < \infty\right\} \subseteq \mathbb{A}_K^\times$$

Exercise 8.6. *Check $\mathbb{A}_{K,\infty}^\times$ is a subgroup of $\mathbb{A}_K^\times$. Show its intersection with the subgroup $K^\times$, i.e. $\mathbb{A}_{K,\infty}^\times \cap K^\times$, is the group of units $\mathcal{O}_K^\times$ (regarded as a subgroup of $\mathbb{A}_K^\times$).*

Theorem 8.2.10. *The map*

$$\begin{aligned} \mathbb{A}_K^\times &\to \mathcal{C}l_K \\ \alpha = (\alpha_v) &\mapsto \prod \mathfrak{p}^{\mathrm{ord}_\mathfrak{p}(\alpha_\mathfrak{p})} \end{aligned}$$

is a surjective homomorphism with kernel $K^\times \cdot \mathbb{A}_{K,\infty}^\times$. In particular, this defines an isomorphism

$$C_K / \mathbb{A}_{K,\infty}^\times \simeq \mathcal{C}l_K$$

of the idèle class group mod the ∞-idèles with Dedekind's ideal class group.

Proof. Note that if $\alpha_\mathfrak{p} \in \mathcal{O}_{K_\mathfrak{p}}^\times$, then $\mathrm{ord}_\mathfrak{p}(\alpha_\mathfrak{p}) = 0$ so $\mathfrak{p}^{\mathrm{ord}_\mathfrak{p}(\alpha_\mathfrak{p})} = (1) = \mathcal{O}_K$. Since $\alpha \in \mathbb{A}_K^\times$ means $\alpha_\mathfrak{p} \in \mathcal{O}_{K_\mathfrak{p}}^\times$ for all but finitely many $\mathfrak{p}$, the product in the definition of the homomorphism is in fact a finite product so the definition makes sense. It is then obvious it is a homomorphism.

Exercise 8.7. *(i) Show for $\mathbb{A}_{K,\infty}^\times$ is in the kernel of the above map into $\mathcal{C}l_K$*

(ii) Show $K^\times$ is kernel of the above map into $\mathcal{C}l_K$. (Hint: show if $x \in \mathcal{O}_K^\times$, then the ideal $(x) = x\mathcal{O}_K = \prod \mathfrak{p}^{\mathrm{ord}_\mathfrak{p}(x)}$.)

It is not much more difficult to show that these subgroups give the whole kernel.

To complete the proof, one needs to show the above map is surjective. Let $\mathcal{I}$ be an arbitrary ideal in $\mathcal{C}l_K$, and write the prime ideal factorization as $\mathcal{I} = \prod_{\mathfrak{p} \in S} \mathfrak{p}^{e_\mathfrak{p}}$ where S is some finite set of primes. Then we can construct an idèle $\alpha = (\alpha_v)$ where $\alpha_v = 1$ if $v \notin S$ and $\alpha_\mathfrak{p} = \varpi_\mathfrak{p}^{e_\mathfrak{p}}$. Here $\varpi_\mathfrak{p}$ is a **uniformizer** of $\mathcal{O}_{K_\mathfrak{p}}$, i.e., $\mathrm{ord}_\mathfrak{p}(\varpi_\mathfrak{p}) = 1$. To see that such a $\varpi_\mathfrak{p}$ always exists, just take $\varpi_\mathfrak{p} \in \mathcal{O}_K$ such that $\varpi_\mathfrak{p} \in \mathfrak{p}$ but $\varpi_\mathfrak{p} \notin \mathfrak{p}^2$. □

This leads us to a topological proof of

Corollary 8.2.11. *The ideal class group $\mathcal{C}l_K$ is finite.*

Proof. We consider the above map restricted to $\mathbb{A}_K^1$. It is easy to see that this is still surjective—in our construction of α above, we were free to do what we want at the infinite places so we can ensure $|\alpha| = 1$. Hence we have an isomorphism

$$C_K^1/\mathbb{A}_{K,\infty}^1 \simeq \mathcal{C}l_K$$

where $\mathbb{A}_{K,\infty}^1 = \mathbb{A}_{K,\infty}^\times \cap \mathbb{A}_K^\times$. However C_K^1 is compact, and $\mathbb{A}_{K,\infty}^1$ is an open subset. Hence the quotient $\mathcal{C}l_K$ is both compact and discrete, whence finite. □

8.3 Elements of class field theory

Class field theory is regarded as the crowning acheivement of algebraic number theory, just as quadratic reciprocity was the crowning achievement of elementary number theory. Class field theory is often described as a characterization of the abelian extensions of a number field, but its explicit forms generalize quadratic and higher reciprocity laws.

What do we mean by higher reciprocity laws? Well the most basic way of thinking about quadratic reciprocity is a way to tell if something is a square mod p. Cubic reciprocity is a way to tell if something is a cube mod $\mathfrak{p}$, and similarly there are notions of biquadratic (4th power) and higher reciprocity laws. Looked at from the point of view of rings of integers, quadratic reciprocity tells us about the way primes split in quadratic extensions. So you might guess cubic reciprocity should tell us about the way primes split in (normal) cubic extensions, and so on. In general, *Artin reciprocity* (a more explicit form class field theory) tells us how primes split in abelian extensions.

Even to state the main theorems of class field theory is not so simple, and we still need to make some more definitions.

Let L/K be an extension of number fields. Let $\mathfrak{P}$ be a prime ideal of L lying above $\mathfrak{p}$, a prime ideal of K. The **decomposition group** of L/K at $\mathfrak{P}$ is

$$G(\mathfrak{P}|\mathfrak{p}) = \{\sigma \in \mathrm{Gal}(L/K) : \sigma(\mathfrak{P}) = \mathfrak{P}\}.$$

Recall that $\mathrm{Gal}(L/K)$ acts on the primes of L above $\mathfrak{p}$, so the $G(\mathfrak{P}|\mathfrak{p})$ is just the stabilizer of $\mathfrak{P}$. Each element of $G(\mathfrak{P}|\mathfrak{p})$ extends to an automorphism of the completion $L_\mathfrak{P}$ which is trivial on $K_\mathfrak{p}$. One can define Galois groups for extensions of local fields ($L_\mathfrak{P}/K_\mathfrak{p}$ is a finite extension of degree $f(\mathfrak{P}|\mathfrak{p})$) and show $\mathrm{Gal}(L_\mathfrak{P}/K_\mathfrak{p}) \simeq G(\mathfrak{P}|\mathfrak{p})$.

As in the case of number fields one can define a **norm** from $L_\mathfrak{P}$ to $K_\mathfrak{p}$ given by

$$N_{\mathfrak{P}|\mathfrak{p}}(x) = N_{L_\mathfrak{P}/K_\mathfrak{p}}(x) = \prod_{\sigma \in \mathrm{Gal}(L_\mathfrak{P}/K_\mathfrak{p})} \sigma(x).$$

One can also do something similar for the archimedean, or infinite, places. In particular, if v is an infinite place of L (i.e., an element of $\mathrm{Gal}(/\mathbb{Q})$, up to complex conjugation) and w is an infinite place of K (i.e., an element of $\mathrm{Gal}(K/\mathbb{Q})$ up to complex conjugation), we write $v|w$ if the embedding $v : L \hookrightarrow \mathbb{C}$ restricted to K gives the embedding $w : K \hookrightarrow \mathbb{C}$ (up to complex conjugation). If $v|w$, then $N_{v|w}(z) = z$ if $L_v = K_w = \mathbb{R}$ or $\mathbb{C}$, and $N_{v|w}(z) = z\overline{z}$ if $L_v = \mathbb{C}$ and $K_w = \mathbb{R}$.

Using this, one can define a **norm** from $\mathbb{A}_L^\times$ to $\mathbb{A}_K^\times$ given by

$$N_{L/K}((\alpha_v)_v) = (\prod_{v|w} N_{v|w}(\alpha_v))_w$$

Exercise 8.8. *Let $x \in L^\times$ and regard $x = (x, x, x, \dots) \in \mathbb{A}_L^\times$. Show the idèlic norm $N_{L/K}(x)$ lies in $K^\times \subseteq \mathbb{A}_K^\times$.*

For a number field K, let $\overline{K}$ denote its algebraic closure, and for a group G let G^{ab} denote is abelianization (quotient via the commutator subgroup). Note that $\mathrm{Gal}(\overline{K}/K)^{ab}$ "contains" the Galois group of any abelian extension L/K as a quotient. In fact, there is a maximal abelian extension K^{ab} of K inside $\overline{K}$ (infinite degree of course), and we will have $\mathrm{Gal}(K^{ab}/K) = \mathrm{Gal}(\overline{K}/K)^{ab}$. The extension K^{ab} contains all finite abelian extensions of K.

Now we can at least state some of the "non-explicit" assertions of class field theory:

Theorem 8.3.1. *Let K be a number field. There is a homomorphism, called the* **Artin map***,*

$$\theta_K : C_K \to \mathrm{Gal}(K^{ab}/K)$$

such that

(i) For every finite abelian extension L/K, let $\theta_{L/K}$ denote the composition of

$$\theta_{L/K} : C_K \overset{\theta_K}{\to} \mathrm{Gal}(K^{ab}/K) \to \mathrm{Gal}(L/K)$$

Then $\ker \theta_{L/K} = N_{L/K}(C_L)$, which yields an isomorphism

$$C_K/(N_{L/K}C_L) = \mathbb{A}_K^\times/(K^\times \cdot N_{L/K}(\mathbb{A}_L^\times)) \simeq \mathrm{Gal}(L/K)$$

(ii) Given any open subgroup of N of C_K of finite index, there is a finite abelian extension L of K with $N = \ker \theta_{L/K}$. Hence

$$C_K/N \simeq \mathrm{Gal}(L/K).$$

There are also some functoriality results which say how the Artin maps θ_K and θ_L are related for an extension L/K, but we will pass over these now.

Let $\zeta_n = e^{2\pi i/n}$.

Corollary 8.3.2. (Kronecker–Weber) *Every abelian extension of $\mathbb{Q}$ is contained in $\mathbb{Q}(\zeta_n)$ for some n.*

An equivalent way to state this is that the maximal abelian extension $\mathbb{Q}^{ab}$ of $\mathbb{Q}$ in $\overline{\mathbb{Q}}$ is the compositum of the extensions $\mathbb{Q}(\zeta_n)$ for all n.

The basic idea of the proof is the following. Class field theory says the abelian extensions of $\mathbb{Q}$ correspond to the open subgroups of the idèle class group $C_{\mathbb{Q}}$. To understand what these are, we want to determine the structure of $\mathbb{A}_{\mathbb{Q}}^\times$. Specifically, one can show

$$\mathbb{A}_{\mathbb{Q}}^\times \simeq \mathbb{Q}^\times \times \mathbb{R}_{>0} \times \hat{\mathbb{Z}}^\times$$

where $\hat{\mathbb{Z}}^\times = \varprojlim (\mathbb{Z}/n\mathbb{Z})^\times = \prod_p \mathbb{Z}_p^\times$. Consequently

$$C_{\mathbb{Q}} \simeq \mathbb{R}_{>0} \times \hat{\mathbb{Z}}^\times.$$

Hence if U is an open subgroup of $C_{\mathbb{Q}}$ with finite index, then U must be of the form $U \simeq \mathbb{R}_{>0} \times U'$ where U' is an open subgroup of finite index in $\mathbb{Z}^\times$. (Since there are no nontrivial open subgroups of finite index in $\mathbb{R}_{>0}$.) Then one uses a basis of neighborhoods for 1 in $\hat{\mathbb{Z}}^\times$ to show that U must contain $N_{K/\mathbb{Q}}(C_K)$ where K is some $\mathbb{Q}(\zeta_n)$. Consequently the extension corresponding to U must contain K. (The functor from open subgroups of C_K to abelian extensions of K is contravariant (i.e., inclusion-reversing), just like the functor from subgroups of the absolute Galois group of K to extensions of K.)

The above theorem of class field theory was established by Takagi, but the existence of a homomorphism θ_K was given abstractly. It was Artin who was able to give it in an explicit fashion, which we now briefly describe.

Let L/K be a Galois extension of number fields, $\mathfrak{p}$ a prime of K and $\mathfrak{P}$ a prime of L lying above $\mathfrak{p}$. Let $f = f(\mathfrak{p}|p)$ where p is the prime of $\mathbb{Q}$ lying under $\mathfrak{p}$, so the residue field $\mathcal{O}_K/\mathfrak{p}$ has order $q = p^f$. The **Frobenius map** $Fr_q : x \mapsto x^q$ generates the Galois group $\mathrm{Gal}((\mathcal{O}_L/\mathfrak{P})/(\mathcal{O}_K/\mathfrak{p}))$. The decomposition group maps to $\mathrm{Gal}((\mathcal{O}_L/\mathfrak{P})/(\mathcal{O}_K/\mathfrak{p}))$ via

$$\phi : G(\mathfrak{P}|\mathfrak{p}) \to \mathrm{Gal}((\mathcal{O}_L/\mathfrak{P})/(\mathcal{O}_K/\mathfrak{p}))$$

by

$$\sigma \mapsto (a\mathfrak{P} \mapsto \sigma(a)\mathfrak{P}).$$

This is an isomorphism if $\mathfrak{P}|\mathfrak{p}$ is unramified, and in this case and we let the **Frobenius element** $\phi_{\mathfrak{P}|\mathfrak{p}}$ of $\mathrm{Gal}(\mathfrak{P}|\mathfrak{p}) \subseteq \mathrm{Gal}(L/K)$ be the inverse image $\phi_{\mathfrak{P}|\mathfrak{p}} = \phi^{-1}(Fr_q)$ of Fr_q in $\mathrm{Gal}(\mathfrak{P}|\mathfrak{p})$.

Exercise 8.9. *Let L/K be a Galois extension of number fields. Let $\mathfrak{P}$ and $\mathfrak{P}'$ be primes of L lying above $\mathfrak{p}$. Show the decomposition groups $G(\mathfrak{P}|\mathfrak{p})$ and $G(\mathfrak{P}'|\mathfrak{p})$ are conjugate in* $\mathrm{Gal}(L/K)$.

We regard each Frobenius element $\phi_{\mathfrak{P}|\mathfrak{p}} \in \mathrm{Gal}(\mathfrak{P}|\mathfrak{p})$ as an element of $\mathrm{Gal}(L/K)$. The above exercise implies for two primes $\mathfrak{P}$ and $\mathfrak{P}'$ of L above $\mathfrak{p}$, the Frobenius elements $\phi_{\mathfrak{P}|\mathfrak{p}}$ and $\phi_{\mathfrak{P}'|p}$ are conjugate in $\mathrm{Gal}(L/K)$.

If L/K is abelian, then the conjugacy classes of $\mathrm{Gal}(L/K)$ are just single elements and the Frobenius $\phi_{\mathfrak{P}|\mathfrak{p}} \in \mathrm{Gal}(L/K)$ does not depend upon the choice of prime $\mathfrak{P}$ above $\mathfrak{p}$. Hence in this case we define the **Frobenius at $\mathfrak{p}$** to be

$$\phi_{\mathfrak{p}} = \left(\frac{L/K}{\mathfrak{p}}\right) := \phi_{\mathfrak{P}|\mathfrak{p}} \in \mathrm{Gal}(L/K)$$

where $\mathfrak{P}$ is a prime of L above $\mathfrak{p}$. The symbol $\left(\frac{L/K}{\mathfrak{p}}\right)$ is called the **Artin symbol**.

One can similarly define the Artin symbol $\left(\frac{L/K}{v}\right)$ for any place v of K, which will be some element of $\mathrm{Gal}(L/K)$. See [Neukirch] for the complete details.

Theorem 8.3.3. (Artin) *Let L/K be a finite abelian extension of number fields. Let $\varpi_{\mathfrak{p}}$ be a uniformizer for $\mathcal{O}_{K_{\mathfrak{p}}}$, i.e., an element of $\mathcal{O}_{K_{\mathfrak{p}}}$ such that $\mathrm{ord}_{\mathfrak{p}}(\varpi_{\mathfrak{p}}) = 1$. Let $x_{\mathfrak{p}} = (\alpha_v) \in \mathbb{A}_K^\times$ be the idèle such that $\alpha_v = \varpi_{\mathfrak{p}}$ for $v = \mathfrak{p}$ and $\alpha_v = 1$ otherwise.*

We may take Artin map θ_K above such that

$$\theta_{L/K}(x_{\mathfrak{p}}) = \phi_{\mathfrak{p}} = \left(\frac{L/K}{\mathfrak{p}}\right)$$

for all primes $\mathfrak{p}$ of K which are unramified in L/K.

The Artin symbol can be used to describe n-th power reciprocity laws. In order to make sense of this, we should define n-th power residue symbols $\left(\frac{a}{p}\right)_n$. This should be 1 if a is an n-th power mod p. But if this is going to be multiplicative, it can't simply be -1 if a is not an n-th power mod p (think about the case $n = 3$). What we need is $\left(\frac{a}{p}\right)_n$ should give a group homomorphism into the n-th roots of unity, such that the kernel is precisely the set of n-th powers mod p. In fact because of this, it won't make sense to define n-th power residue symbols over $\mathbb{Q}$ (or $\mathbb{Z}$), but only over number fields which contain the n-th roots of unity.

Let μ_n denote the n-th roots of unity.

Definition 8.3.4. *Let K be a number field containing μ_n and v be a place of K. The* n**-th Hilbert symbol**

$$\left(\frac{-,-}{v}\right)_n : K_v^\times \times K_v^\times \to \mu_n$$

is given by

$$\left(\frac{K_v(\sqrt[n]{b})/K_v}{v}\right)\sqrt[n]{b} = \phi_v(\sqrt[n]{b}) = \left(\frac{a,b}{v}\right)_n \sqrt[n]{b}.$$

In other words, the Frobenius $\phi_v = \left(\frac{K_v(\sqrt[n]{b})/K_v}{v}\right)$ is an element of $\mathrm{Gal}(K_v(\sqrt[n]{b})/K_v)$. But the conjugates of $\sqrt[n]{b}$ in $K_v(\sqrt[n]{b})/K_v$ are just $\sqrt[n]{b}$ times the n-th roots of unity. Hence $\phi_v(\sqrt[n]{b})$ is some n-th root of unity times $\sqrt[n]{b}$, and we let the n-th Hilbert symbol $\left(\frac{a,b}{v}\right)_n$ be that root of unity.

Theorem 8.3.5. *Let K be a number field containing μ_n and v be a place of K. For $a, b \in K^\times$, we have*

$$\prod_v \left(\frac{a,b}{v}\right)_n = 1.$$

Proof. We have

$$\prod_v \left(\frac{a,b}{v}\right)_n \sqrt[n]{b} = \prod_v \left(\frac{K_v(\sqrt[n]{b})/K_v}{v}\right)\sqrt[n]{b} = \prod_v \theta_{K(\sqrt[n]{b})/K}(a)\sqrt[n]{b}.$$

However, any element of $K^\times$ is in the kernel of the Artin map $\theta_{K(\sqrt[n]{b})/K}$, so the above must equal $\sqrt[n]{b}$ and the asserted product formula follows. □

Definition 8.3.6. *Let $a \in K^\times$ where $\mu_n \subseteq K$. For $\mathfrak{p}$ a prime of K, we define the* n**-th power residue symbol** *to be*

$$\left(\frac{a}{\mathfrak{p}}\right)_n = \left(\frac{a, \varpi_{\mathfrak{p}}}{\mathfrak{p}}\right)_n$$

where $\varpi_{\mathfrak{p}}$ is a uniformizer for $K_{\mathfrak{p}}$. If $b \in K^\times$, we set

$$\left(\frac{a}{b}\right)_n = \prod_{\mathfrak{p}_i \nmid n} \left(\frac{a}{\mathfrak{p}_i}\right)_n^{e_i},$$

where $(b) = \prod \mathfrak{p}_i^{e_i}$ ideal of K.

It is not too hard to check that

$$\left(\frac{a}{\mathfrak{p}}\right)_n = 1 \iff a \equiv x^n \bmod \mathfrak{p}$$

and more generally

$$\left(\frac{a}{\mathfrak{p}}\right)_n \equiv a^{\frac{N(\mathfrak{p})-1}{n}} \bmod \mathfrak{p}.$$

Theorem 8.3.7. **(n-th power reciprocity)** *Suppose $\mu_n \subseteq K^\times$. If $a, b \in K^\times$, then*

$$\left(\frac{a}{b}\right)_n = \left(\frac{b}{a}\right)_n \prod_{v|n\infty} \left(\frac{a,b}{v}\right)_n.$$

This follows simply from the above product formula (the previous theorem). See [Neukirch].

In particular, if a and b are prime elements of $\mathcal{O}_K$ (i.e., they generate prime ideals of $\mathcal{O}_K$), and $\prod_{v|n\infty}\left(\frac{a,b}{v}\right)_n = 1$, then a is an n-th power mod b if and only if b is an n-th power mod a.

Corollary 8.3.8. **(Quadratic Reciprocity)** *Let $K = \mathbb{Q}$ and $n = 2$. Let a, b be odd coprime integers. Then*

$$\left(\frac{a}{b}\right)_2 \left(\frac{b}{a}\right)_2 = (-1)^{\frac{a-1}{2}\frac{b-1}{2}} (-1)^{\frac{sgn(a)-1}{2}\frac{sgn(b)-1}{2}},$$

and

$$\left(\frac{-1}{b}\right)_2 = (-1)^{\frac{b-1}{2}}, \quad \left(\frac{2}{b}\right)_2 = (-1)^{\frac{b^2-1}{8}}.$$

Corollary 8.3.9. **(Cubic Reciprocity)** *Let $K = \mathbb{Q}(\zeta_3)$ and $n = 3$. Suppose p, q are primes (i.e., they generate prime ideals) in $\mathcal{O}_K$ such that $p, q \equiv \pm 1 \bmod 3$. (If (α) is prime in $\mathcal{O}_K$ which does not lie above 3, then it has 6 associates, 2 of which are $\equiv \pm 1 \bmod 3$.) Then if p and q lie above different primes of $\mathbb{Q}$, we have*

$$\left(\frac{p}{q}\right)_3 = \left(\frac{q}{p}\right)_3.$$

Hence class field theory generalizes quadratic and higher reciprocity laws.

8.4 Non-abelian class field theory

The n-th power reciprocity law says that if p and q are prime elements of $K \supset \mu_n$, then we can determine whether p is an n-th power mod q based on whether or not q is an n-th power mod p. In particular, we have

$$\left(\frac{p}{q}\right)_n = \left(\frac{q}{p}\right)_n$$

if the product $\prod_{v|n\infty}\left(\frac{p,q}{v}\right)_n = 1$. The proof of this reciprocity law is essentially to look at the Artin map

$$\theta_{L/K} : C_K \to \mathrm{Gal}(L/K)$$

for the extension L/K where $L = K(\sqrt[n]{a})$. Since this applies only to abelian extensions, we see the need for the requirement that $\mu_n \subseteq K$ from the point of view of class field theory. Specifically, assuming none of the n-roots of $x^n - a$ lie in K, the extension $K(\sqrt[n]{a})/K$ is abelian (in fact cyclic of degree n) if and only if $\mu_n \subset K$.

Hence if one wanted to extend the n-th power reciprocity law to $\mathbb{Q}$, one would want some sort of non-abelian version of class field theory. In fact, one might guess a reciprocity law roughly of the following form: *Let $f(x)$ be an irreducible polynomial over $\mathbb{Z}$. If p and q are odd primes not dividing n, then one can determine when*

$$f(x) \equiv p \textit{ is solvable mod } q$$

in terms of when

$$f(x) \equiv q \textit{ is solvable mod } p.$$

Indeed this is essentially what n-th power reciprocity says in the case $f(x) = x^n$. Though it seems likely that a "non-abelian reciprocity law" will be more complicated than this.

To put the notion of reciprocity in a little more imprecise way, recall that $x^2 \equiv q \bmod p$ has a solution, i.e., $x^2 - q$ has a root mod p, if and only if p is split in $\mathbb{Q}(\sqrt{q})$. Similarly if p and q are primes in K, then $x^n \equiv q \bmod p$ has a solution if and only if $x^n - q$ has a root mod p, which means p is split in $K(\sqrt[n]{q})$. If $K(\sqrt[n]{q})/K$ is Galois, i.e., if $\mu_n \subset K$, we can in fact say $x^n \equiv q \bmod p$ if and only if p splits completely in $K(\sqrt[n]{q})/K$. Hence we may think of n-th power reciprocity as a description of which primes split in $K(\sqrt[n]{q})/K$. Class field theory can then be thought of as a description of which primes split in an abelian extension L/K. Thus non-abelian class field theory, or a non-abelian reciprocity law, should be a description of which primes split in a non-abelian extension L/K.

Before we think about what the statement of non-abelian class field theory should look like in general, we sketch out an example.

Example 8.4.1. *(Koike, 1985) Let $f(x)$ be an irreducible polynomial of degree 3 over $\mathbb{Q}$, and let K be the splitting field of $f(x)$. Assume $\mathrm{Gal}(K/\mathbb{Q}) \simeq S_3$ and K contains an imaginary quadratic extension. One can to associate to $f(x)$ the* elliptic curve

$$E : y^2 = f(x)$$

as well as a corresponding modular form

$$F : \mathfrak{H} = \{z \in \mathbb{C} : Im(z) > 0\} \to \mathbb{C}$$

$$F(\tau) = \sum_{n=1}^{\infty} a_n e^{2\pi i n \tau},$$

where the a_n's are certain Fourier coefficients which determine the function $F(\tau)$.

Let n_p be the number of solutions $\#E(\mathbb{F}_p)$ to $y^2 \equiv f(x) \bmod p$. Then the precise correspondence between E and F is that $a_p = p + 1 - n_p$. One version of a non-abelian reciprocity law in this case say that, apart from p lying in a finite set of primes,

$$p \textit{ splits completely in } K \iff a_p = 2.$$

Hence we can describe the set of primes which split completely in K in terms of either (i) arithmetic data associated to an elliptic curve, or (ii) arithmetic data associated to a modular form.

Langlands' conjecture

In order to think about how one might set up a general non-abelian class field theory, let's go back to understanding what (abelian) class field theory says. Class field theory says there is an homomorphism from the idèle class group $C_K = \mathbb{A}_K/K^\times$ to $\mathrm{Gal}(K^{ab}/K)$, which satisfies certain properties. In particular, we have an isomorphism $C_K/N_{L/K}C_L \simeq \mathrm{Gal}(L/K)$ for any finite abelian extension L/K.

If one wants to extend this to non-abelian extensions L/K, one might look for "non-abelian class groups" $G(K)$ such that $G(K)$ is related to $\mathrm{Gal}(\overline{K}/K)$ and, for any finite Galois extension L/K, we have $G(K)/N_{L/K}(G(L)) \simeq \mathrm{Gal}(L/K)$, where $N_{L/K}(G(L))$ is a certain (normal?) subgroup of $G(K)$ associated to the "non-abelian class group" $G(L)$ of L. It is not clear how such a "non-abelian class group" could be constructed. However there are very specific conjectures for a non-abelian generalization if look at the dual picture, i.e., put things in terms of group representations and L-functions.

If G is a locally compact abelian group, we can consider the set of (unitary) characters, $\hat{G}$, consisting of continuous homomorphisms $G \to S^1$. The set $\hat{G}$ is naturally made into a locally compact abelian group, called the **dual group** of G. Pontryagin duality says that the dual group of $\hat{G}$ is isomorphic to G in a canonical way. Thus, to study C_K or $\mathrm{Gal}(K^{ab}/K)$, it is equivalent to study their dual groups. Characters $\omega : C_K \to S^1$ are called **idèle class characters** or **Hecke characters**. Characters $\chi : \mathrm{Gal}(K^{ab}/K) \to S^1$ are called 1-dimensional **Galois representations**. More generally, an n-dimensional (complex) Galois representation is a continuous homomorphism $\rho : \mathrm{Gal}(\overline{K}/K) \to \mathrm{GL}_n(\mathbb{C})$. But a 1-dimensional representation (i.e., a character) of $\chi : \mathrm{Gal}(\overline{K}/K) \to \mathrm{GL}_1(\mathbb{C}) = \mathbb{C}^\times$ will have image in S^1 factor through $\mathrm{Gal}(K^{ab}/K)$, so this agrees with our definition above.

Consider a 1-dimensional Galois representation $\chi : \mathrm{Gal}(K^{ab}/K) \to \mathbb{C}^\times$. By composition with the Artin map, we get a Hecke character

$$\chi \rightsquigarrow \omega = \chi \circ \theta_K.$$

Since the Artin map is not an isomorphism, so one does not necessarily (in fact does not) get all Hecke characters this way, but one gets all *finite order* Hecke characters this way. (A character ω is finite order if $\omega^m = 1$ for some natural number m.) Namely, continuity of χ implies χ has finite image, so it factors through (the Galois group of) a finite abelian extension $\chi : \mathrm{Gal}(L/K) \to \mathbb{C}^\times$, consequently ω will factor through C_K/N, where $N = N_{L/K}(C_L)$. For a finite abelian group G, the group of characters $\hat{G}$ is actually (non-canonically) isomorphic to G, so the above correspondence of 1-dimensional Galois representations with finite order Hecke characters gives a bijection (in fact isomorphism)

$$\left\{\omega : C_K/N \to \mathbb{C}^\times\right\} \overset{1-1}{\longleftrightarrow} \left\{\chi : \mathrm{Gal}(L/K) \to \mathbb{C}^\times\right\}.$$

This correspondence of Galois representations and finite order Hecke characters is equivalent to abelian class field theory.

Now the natural guess for a "higher dimensional," or non-abelian analogue of this would be to get a correspondence with n-dimensional representations of $\mathrm{Gal}(\overline{K}/K)$ for any n. (Again, by continuity, any given representation will factor through a finite extension $\mathrm{Gal}(L/K)$. Moreover, if $n > 1$ and the representation is irreducible, then L/K will not be abelian.) The question is, what group should we pick on the left? This was an insight of Langlands (building on the work of many before him). Note that we can view the idèle class group as

$$C_K = \mathbb{A}_K^\times / K^\times = \mathrm{GL}_1(K)\backslash\mathrm{GL}_1(\mathbb{A}_K).$$

(Read the latter as $\mathrm{GL}_1(\mathbb{A}_K)$ mod $\mathrm{GL}_1(K)$. We typically write the mod on the left as above however—we also sometimes write $K^\times\backslash\mathbb{A}_K^\times$. Of course in this case our groups our abelian, so it doesn't matter which side we mod out on, but it will for the non-abelian groups below. The reason for putting mod on the left is because we sometimes want to mod out by another subgroup on the right—of course which goes on the left and which on the right is just a matter of convention.)

Conjecture 8.4.2. (Langlands) *There is a (partial)* $1-1$ *correspondence*

$$\{\textit{automorphic representations } \pi \textit{ of } \mathrm{GL}_n(K)\backslash\mathrm{GL}_n(\mathbb{A}_K)\} \overset{1-1}{\overset{\dashrightarrow}{\longleftarrow}} \\ \left\{n\textit{-dimensional representations } \rho \textit{ of } \mathrm{Gal}(\overline{K}/K)\right\}.$$

Roughly, an **automorphic representation** of a locally compact group G is a irreducible representation of G on $L^2(G)$. The diagonal subgroup $\mathrm{GL}_n(K) \subset \mathrm{GL}_n(\mathbb{A}_K)$ is not normal (for $n > 1$), so the quotient $\mathrm{GL}_n(K)\backslash\mathrm{GL}_n(\mathbb{A}_K)$ is not actually a group. Hence this requires some explanation.

First note if G is a finite group, $L^2(G)$ is just the $\mathbb{C}$-vector space of $\mathbb{C}$-valued functions on G. We can take for a basis $\{e_g\}_{g\in G}$ where e_g is the characteristic function of g in G. Hence $L^2(G) \simeq \mathbb{C}[G]$, the group algebra, and we know $\mathbb{C}[G]$ decomposes as a direct sum of the irreducible representations of G.

When G is not finite, things are more complicated, but in any event G acts on the space $L^2(G)$ by right multiplication, i.e., $g : f(x) \to f(xg)$ for any $f \in L^2(G)$. In fact if $G = \mathrm{GL}_n(\mathbb{A}_K)$, G acts on $L^2(\mathrm{GL}_n(K)\backslash\mathrm{GL}_n(\mathbb{A}_K))$ in the same way. This representation, the *right regular representation* on $L^2(\mathrm{GL}_n(K)\backslash\mathrm{GL}_n(\mathbb{A}_K))$, decomposes into irreducible constituents. What we mean by an automorphic representation of $\mathrm{GL}_n(\mathbb{A}_K)$ (or $\mathrm{GL}_n(K)\backslash\mathrm{GL}_n(\mathbb{A}_K)$) is one of these irreducible constituents. The term *automorphic* means that the representation is realized on a space of *automorphic forms*, which are functions on $\mathrm{GL}_n(\mathbb{A}_K)$ invariant under $\mathrm{GL}_n(K)$. When $n > 1$, automorphic representations are infinite-dimensional representations, and are studied using more harmonic analysis than algebra, per say.

Langlands' conjecture states that to each n-dimensional Galois representation $\rho : \mathrm{Gal}(\overline{K}/K) \to \mathrm{GL}_n(\mathbb{C})$, there is associated (in a way we shall describe below) an automorphic representation $\pi = \pi(\rho)$ of $\mathrm{GL}_n(\mathbb{A}_K)$. However in general there will be more automorphic representations than n-dimensional Galois representations, i.e., not every automorphic representation will correspond to a Galois representation. This is indicated by the dashed arrow going from left to right in the conjecture above. This is true even when $n = 1$, the left hand side is just the set of Hecke characters of C_K, and and so one needs to restrict to finite order Hecke characters to get an honest $1-1$ correspondence between these two sets of representations in this case.

This conjecture of Langlands suggests that the conjectural group $G(K)$ should contain in some way each $\mathrm{GL}_n(K)\backslash\mathrm{GL}_n(\mathbb{A}_K)$, so that the representations of $G(K)$ correspond to all Galois representations. However this situation is even more ambiguous than the state of Langlands' conjecutre above, and in any case understanding the conjecture above would be extraordinary progress to developing a non-abelian class field theory. For these reasons, we will spend the rest of our time trying to explain what the above conjecture means.

L-functions

To describe the conjecture of Langlands above[*], one needs to specify exactly how the representations should correspond. The answer comes via the construction of L-functions associated to each representations. Let's first see what happens in the case $n = 1$.

Suppose χ is a 1-dimensional representation of $\mathrm{Gal}(\overline{K}/K)$. Then χ factors through a finite abelian extension

$$\chi : \mathrm{Gal}(L/K) \to \mathbb{C}^\times = \mathrm{GL}_1(\mathbb{C}).$$

If $\mathfrak{p}$ is a prime of K, recall we have a Frobenius element $\phi_\mathfrak{p} \in G(\mathfrak{P}|\mathfrak{p}) \subseteq \mathrm{Gal}(L/K)$ where $\mathfrak{P}$ is a prime of L above $\mathfrak{p}$. Then we define the L-function associated to χ to be

$$L(s, \chi) = \prod_{\mathfrak{p} \text{ unram}} \frac{1}{1 - \chi(\phi_\mathfrak{p}) N(\mathfrak{p})^{-s}}.$$

One can regard this as a generalization of the Dirichlet L-series, as specializing to the case $K = \mathbb{Q}$ and $L = \mathbb{Q}(\zeta_m)$ gives the Dirichlet L-functions mod m.

On the other hand, if ω is a Hecke character

$$\omega : C_K = K^\times \backslash \mathbb{A}_K^\times \to \mathbb{C}^\times,$$

we can view ω as a character of $\mathbb{A}_K^\times$ which is trivial on $K^\times$. This gives a character

$$\omega_v : K_v^\times \to \mathbb{C}^\times$$

for any place v of K simply by restricting to the v-component of $\mathbb{A}_K^\times$. Specifically $\omega_v(x_v) = \omega(1, \ldots, 1, x_v, 1 \ldots)$ where the x_v occurs in the v-th place. Then one can think of $\omega = \prod_v \omega_v$. When $v = \mathfrak{p}$, we say $\omega_\mathfrak{p}$ is **unramified** if $\omega_\mathfrak{p}$ is trivial on $\mathcal{O}_{K_\mathfrak{p}}^\times$. Then one can define the **Hecke L-function**

$$L(s, \omega) = \prod_{\omega_\mathfrak{p} \text{ unram}} \frac{1}{1 - \omega_\mathfrak{p}(\varpi_\mathfrak{p}) N(\mathfrak{p})^{-s}},$$

where $\varpi_\mathfrak{p}$ is a uniformizer for $\mathcal{O}_{K_\mathfrak{p}}$.

We say the Galois character χ and the Hecke character ω correspond if

$$L(s, \chi) = L(s, \omega),$$

i.e. if $\chi(\phi_\mathfrak{p}) = \omega_\mathfrak{p}(\varpi_\mathfrak{p})$ for each unramified $\mathfrak{p}$. This is the L-function interpretation of class field theory. This explicit correspondence of L-functions is amounts the explicit description of the Artin map.

Now we can define L-functions for higher-dimensional representations. Let

$$\rho : \mathrm{Gal}(\overline{K}/K) \to \mathrm{GL}_n(\mathbb{C})$$

be an n-dimensional Galois representation. Continuity of ρ means there is a finite extension L/K such that ρ restricted to the subgroup $\mathrm{Gal}(\overline{K}/L)$ is trivial, i.e., ρ factors through

$$\rho : \mathrm{Gal}(L/K) \to \mathrm{GL}_n(\mathbb{C}).$$

[*] This conjecture is also called the **strong Artin conjecture** or the **modularity conjecture**. Indeed Langlands made a series of far-reaching related conjectures, so if one just says "Langlands conjecture," it is not always clear which one is being referred to.

For any prime $\mathfrak{p}$ of K and $\mathfrak{P}$ of L with $\mathfrak{P}|\mathfrak{p}$, we have a surjective homomorphism

$$G(\mathfrak{P}|\mathfrak{p}) \to \mathrm{Gal}((\mathcal{O}_L/\mathfrak{P})/(\mathcal{O}_K/\mathfrak{p})).$$

Recall the group on the left, the decomposition group of $\mathfrak{P}|\mathfrak{p}$, is just the subgroup of $\mathrm{Gal}(L/K)$ which stablizes $\mathfrak{P}$. If the inertial degree $f(\mathfrak{P}|\mathfrak{p}) = 1$, in particular if $\mathfrak{p}$ is unramified in L/K, then this map is an isomorphism and the group on the right is generated by $Fr_q : x \to x^q$ where $q = N(\mathfrak{p})$. In this case, the Frobenius element $\phi_{\mathfrak{P}|\mathfrak{p}} \in G(\mathfrak{P}|\mathfrak{p}) \subseteq \mathrm{Gal}(L/K)$. Since all the primes of L lying above $\mathfrak{p}$ are conjugate in $\mathrm{Gal}(L/K)$, all the elements $\phi_{\mathfrak{P}|\mathfrak{p}}$ are conjugate as $\mathfrak{P}$ ranges over the primes above $\mathfrak{p}$. We let the **Frobenius** $\phi_{\mathfrak{p}} = \phi_{\mathfrak{P}|\mathfrak{p}}$ for some $\mathfrak{P}$, so this is well-defined up to conjugacy. Of course if L/K is abelian, each element is its own conjugacy class, and $\phi_{\mathfrak{p}}$ is well-defined as an element of $\mathrm{Gal}(L/K)$.

Since almost all primes $\mathfrak{p}$ of K are unramified, we can define the (partial) **Artin L-function** by

$$L(s, \rho) = \prod_{\mathfrak{p} \text{ unram}} \frac{1}{\det(I_n - \rho(\phi_{\mathfrak{p}})N(\mathfrak{p})^{-s})}.$$

Note even though $\phi_{\mathfrak{p}}$ is only well defined up to conjugacy in $\mathrm{Gal}(L/K)$, the quantity $\det(I_n - \rho(\phi_{\mathfrak{p}})N(\mathfrak{p})^{-s})$ is well defined because the determinant is invariant under conjugation. We say this is a partial L-function because the full or completed Artin L-function is actually defined as a product of terms over all places v (including the archimedean ones), but the partial and the full L-function only differ by a product of finitely many terms (which are well understood). For simplicity we will not define the full L-function, but just mention that at unramified primes $\mathfrak{p}$, one needs to take into account the kernel of the map $G(\mathfrak{P}|\mathfrak{p}) \to \mathrm{Gal}((\mathcal{O}_L/\mathfrak{P})/(\mathcal{O}_K/\mathfrak{p}))$, called the inertial subgroup of $G(\mathfrak{P}|\mathfrak{p})$.

Let π be an automorphic representation of $\mathrm{GL}_n(\mathbb{A}_K) \subset \prod_v \mathrm{GL}_n(K_v)$. Then $\pi = \otimes \pi_v$ where each π_v is a representation of $\mathrm{GL}_n(K_v)$. For $v = \mathfrak{p}$, we say $\pi_{\mathfrak{p}}$ is **unramified** if $\pi_{\mathfrak{p}}$ restricted to the subgroup $\mathrm{GL}_n(\mathcal{O}_{K_{\mathfrak{p}}})$ is trivial. At such a place, π_v is induced from n 1-dimensional representations $\omega_1, \ldots, \omega_n$ placed on the diagonal subgroup of $\mathrm{GL}_n(K_v)$. Set

$$A(\pi_v) = \mathrm{diag}(\omega_1(\varpi_{\mathfrak{p}}), \ldots, \omega_n(\varpi_{\mathfrak{p}})).$$

Then we define the (partial) **automorphic L-function**

$$L(s, \pi) = \prod_{\pi_{\mathfrak{p}} \text{ unram}} \frac{1}{\det(I_n - A(\pi_{\mathfrak{p}})N(\mathfrak{p})^{-s})}.$$

Now we can restate Langlands' conjecture above in more precise terms

Conjecture 8.4.3. (Langlands) *There is a (partial) $1-1$ correspondence*

$$\{\textit{automorphic representations } \pi \textit{ of } \mathrm{GL}_n(\mathbb{A}_K)\} \overset{1-1}{\overset{\dashrightarrow}{\longleftarrow}} \{\textit{n-dimensional representations } \rho \textit{ of } \mathrm{Gal}(\overline{K}/K)\}$$

such that

$$L(s, \pi) = L(s, \rho),$$

i.e., for almost all primes $\mathfrak{p}$ of K, we have

$$\det(I_n - A(\pi_{\mathfrak{p}})N(\mathfrak{p})^{-s}) = \det(I_n - \rho(\phi_{\mathfrak{p}})N(\mathfrak{p})^{-s}).$$

(If the local factors—called local L-factors or local L-functions—agree for $L(s, \pi)$ and $L(s, \rho)$ for almost all places, one can show that the local factors (in the completed L-functions) will all be the same.)

The first application of **Langlands program** (this program of attaching automorphic representations to Galois representations, which has grown into a much more general setting than what we have presented) is to the following.

Conjecture 8.4.4. (Artin) *Let $\rho : \mathrm{Gal}(\overline{K}/K)$ be an irreducible nontrivial Galois representation. Then $L(s, \rho)$ is entire.*

If ρ is trivial, then $L(s, \rho) = \zeta_K(s)$, the Dedekind zeta function of K, has a pole at $s = 1$. If ρ is not trivial, Artin conjectures $L(s, \rho)$ is entire, i.e., it has no poles. (The L-function as defined, converges for $Re(s)$ large, but is known to have meromorphic continuation to the whole complex plane.) If ρ is 1-dimensional, then this is known because $\rho = \chi$ corresponds to a Hecke character ω, and the Hecke L-functions $L(s, \omega)$ for nontrivial ω are known to be entire. It is also easy to see that if ρ is induced from a 1-dimensional representation χ, then $L(s, \rho) = L(s, \chi)$ so $L(s, \rho)$ is entire.

Not much was known about Artin's conjecture for higher dimensional representations. However, it is known that if π is a *cuspidal* automorphic representation, then $L(s, \pi)$ is entire, so if $\rho \leftrightarrow \pi$, then $L(s, \rho)$ is also entire. (Any automorphic representation corresponding to a nontrivial irreducible Galois representation will be cuspidal.) Hence Langlands conjecture implies Artin's conjecture, wherefore the above conjecture of Langlands is sometimes called the strong Artin conjecture. (In fact, the strong Artin conjecture and the Artin conjecture are known to be equivalent in the case of 2 or 3 dimensional representations. It is not clear if they should be equivalent in higher dimensions.)

The first success of the Langlands program is the following result.

Theorem 8.4.5. (Langlands, Tunnell) *Suppose $\rho : \mathrm{Gal}(L/K) \to \mathrm{GL}_2(\mathbb{C})$ is an irreducible 2-dimensional representation. If the image of ρ is solvable (a solvable subgroup of $\mathrm{GL}_2(\mathbb{C})$), then $\rho \leftrightarrow \pi$ for some cuspidal automorphic representation π of $\mathrm{GL}_2(\mathbb{A}_K)$.*

This gave new instances of Artin's conjecture. We remark that Artin's conjecture, together with the Grand Riemann Hypothesis (the analogue of the Riemann Hypothesis for more general L-functions), yields estimates for the error term in the prime number theorem.

However, there is a much more famous consequence of this theorem of Langlands and Tunnell—Fermat's Last Theorem. Very roughly, Frey, Ribet and Serre showed that Fermat's Last Theorem follows from the Taniyama–Shimura conjecture, which says that to each elliptic curve over $\mathbb{Q}$, there is an associated modular form, in the sense that their associated L-functions are equal. To prove Taniyama–Shimura, one associates to an elliptic curve E a family of *p-adic* Galois representations $\rho_p : \mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q}) \to \mathrm{GL}_2(\mathbb{Q}_p)$. This much is not difficult. Wiles essentially showed that (for "semistable" E, which is sufficient for Fermat's Last Theorem) one can (reduce to a case where one can) further associate to E a 2-dimensional *complex* Galois representation $\rho : \mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q}) \to \mathrm{GL}_2(\mathbb{C})$, where ρ has solvable image. Then Langlands–Tunnell applies, and ρ (and hence the elliptic curve E) corresponds to an automorphic representation π of $\mathrm{GL}_2(\mathbb{A}_\mathbb{Q})$. This representation π is naturally associated to some modular form f, and this gave Taniyama–Shimura (for semistable curves, which was enough for Fermat's last theorem—the general case was finished later), and hence Fermat's last theorem.