Math 223a: Algebraic Number Theory

Taught by Fabian Gundlach

Fall 2019

These notes are scribed by Kenz Kallal (me). My email is kenz.kallal@gmail.com. Please contact me if you find any mistakes in these notes (all mistakes are mine and not the instructor's). You can also read the notes in real-time using the link [REDACTED].

Other administrative details are below:

- **Section:** Thursdays 4:30-5:30pm, SC411.
- **CA Office hours:** Wednesdays 7-10pm, Winthrop basement room S009 (later moved to Sundays)
- **Office hours:** TTh 2-3pm, SC233.
- **Email:** gundlach@math.harvard.edu (Fabian Gundlach)
- **Textbook:** For infinite Galois theory, see *Algebra: From the Viewpoint of Galois Theory* by Bosch [1]. Some good references for algebraic number theory and class field theory are Neukrich's *Algebraic Number Theory* [5], Lang's *Algebraic Number Theory* [2], and Milne's notes entitled *Class Field Theory* [4].
- **Midterm:** None.
- **Final:** There will be a final paper (5–10 pages).

I am always available outside of office hours by email to answer any questions related to the course material.

The main topics of this course are as follows:

- Infinite Galois theory
- Local fields
- Local class field theory via Lubin–Tate formal groups
- Galois cohomology, especially of local fields
- The main statements of global class field theory.

As such, it has the following prerequisites:

- Galois theory (e.g. math 123)
- Basic algebraic number theory (e.g. math 129)
- Basic topology (e.g. math 131).

Grades will be determined by:

- 70 percent homework, due Thursdays at noon. The 2 lowest grades will be dropped.
- 30 percent final paper.

Contents

§1 September 3, 2019

The main goal of this course is *class field theory*. Though we won't prove the global version, it is useful to start out with a concrete example to motivate the theory and give an idea of what it is about.

§1.1 Classical reciprocity laws

One of the main statements of global class field theory, namely Artin's reciprocity law, is a broad generalization of the classical reciprocity laws, of which quadratic reciprocity is the simplest example. It has many important applications to algebraic number theory, but the simplest way to view it is as a generalization of the classical reciprocity laws.

Let p be an odd prime. The law of quadratic reciprocity concerns the *Legendre symbol mod p*, which is the unique surjective group homomorphism

$$\left(\frac{\cdot}{p}\right) : \mathbf{F}_p^\times \to \{\pm 1\}.$$

For an arbitrary integer a, the Legendre symbol $\left(\frac{a}{p}\right)$ is 0 if $p|a$ and otherwise is ± 1 depending on whether a is a perfect square ("quadratic residue") mod p.

Lemma 1.1 (Euler's criterion)

Let $a \in \mathbf{Z}$. The Legendre symbol $\left(\frac{a}{p}\right)$ is congruent to $a^{(p-1)/2} \mod p$. In particular, a is a quadratic residue mod p if and only if $a^{(p-1)/2} \equiv 1 \mod p$.

Proof. One way to do this is to exploit the fact that $\mathbf{F}_p^\times$ is cyclic. However, it can be done directly. If a is a quadratic residue mod p, then $a \equiv x^2 \mod p$ for some $x \in \mathbf{F}_p^\times$, so (working in $\mathbf{F}_p$)

$$a^{(p-1)/2} = x^{p-1} = 1.$$

But $a^{(p-1)/2}$ is always ± 1 if $a \neq 0$ (since its square is 1 and the polynomial $x^2 - 1$ has the two distinct roots ± 1). Moreover, $x^{(p-1)/2} - 1$ has at most $(p-1)/2$ roots in $\mathbf{F}_p^\times$, and $\mathbf{F}_p^\times$ contains $(p-1)/2$ quadratic residues (for example because each nonzero perfect square in $\mathbf{F}_p$ has exactly two square roots), so this proves that in fact $x^{(p-1)/2} - 1$ has the maximal number $(p-1)/2$ of roots in $\mathbf{F}_p^\times$, and that the quadratic residues are exactly these roots. As a result, $a^{(p-1)/2} = -1$ if a is a quadratic non-residue mod p and is 1 if a is a quadratic residue mod p. □

For a fixed p, the Legendre symbol $\left(\frac{a}{p}\right)$ depends only on $a \mod p$. On the other hand, it turns out that the Legendre symbol is also periodic in p for a fixed a (its value only depends on p modulo a fixed integer n depending only on a). This is essentially the *law of quadratic reciprocity*. Without proving the full result, we can see a few easy examples.

Example 1.2

Let $a = 0$. Then $\left(\frac{a}{p}\right) = 0$ for all p.

Example 1.3

Let $a = 1$. Then $\left(\frac{a}{p}\right) = 1$ for all p.

Example 1.4

Let $a = -1$. Then $\left(\frac{a}{p}\right) = (-1)^{(p-1)/2}$ so -1 is a QR mod p if and only if $p \equiv 1 \mod 4$.

So indeed we see that at least in these special cases the Legendre symbol $\left(\frac{a}{p}\right)$ is periodic in p. One obvious generalization is to look at cubic residues.

Example 1.5

Let $a = 5$. It turns out that whether a is a cubic residue mod p cannot be made to depend on congruence conditions on p.

Example 1.6

The number of roots of $x^3 - 3x + 1$ mod p depends only on p mod 9.

In this course we will explain the differing behavior between $x^3 - 5$ and $x^3 - 3x + 9$[1]. We might also explain how this generalizes to base field equal to an arbitrary number field. In general, this question can take the following form:

Question 1.7. Let K be a number field and $f(X) \in \mathcal{O}_K[X]$ a polynomial. Is there a convenient description of the splitting behavior of $f \mod \mathfrak{p}$ depending on the nonzero prime ideal $\mathfrak{p} \subseteq \mathcal{O}_K$?

Of course this is highly related[2] to a slightly more abstract question

Question 1.8. Let L/K be an extension of number fields. Is there a convenient description of the splitting behavior of the primes of K in L? In particular, which primes have some prescribed splitting type?

For example, the law of quadratic reciprocity amounts to (in the language of Question 1.8) a concrete description of which primes $p \in \mathbf{Z}$ split completely in the quadratic extension $\mathbf{Q}(\sqrt{q})$, or alternatively (in the language of Question 1.7) a concrete description of the splitting behavior of the polynomial $X^2 - q$ mod p. Class field theory will be concerned with certain special cases of these types of questions (though we will see that it still contains a massive generalization of Gauss' law of quadratic reciprocity), namely (in the langauge of Question 1.8) when L/K is Galois with abelian Galois group.

§1.2 Galois theory

Let K be a number field (or more generally an arbitrary global or local field[3]). The main results of class field theory also contain a convenient description of the finite abelian extensions of K. To this end, it is useful to consider all the finite Galois extensions of K

[1] Hint: one of them has abelian Galois group over $\mathbf{Q}$ and one of them doesn't.

[2] See [3, Theorem 27] or [5, I.8.3]

[3] See [5, II.5]

together by looking at the Galois[4] group $\operatorname{Gal}(\overline{K}/K)$. Since $\overline{K}/K$ is in general not finite (for example when K is a number field it has finite extensions of arbitrarily large degree which we could construct by adjoining square roots for instance), we need to spend some time developing the theory of infinite Galois extensions.

Recall the fundamental theorem of finite Galois theory:

Theorem 1.9

Let M/K be a finite Galois extension, and $G = \operatorname{Gal}(M/K)$. The subfields of M containing K are in inclusion-reversing bijective correspondence with the subgroups $H \subseteq G$ via $H \mapsto M^H$, the fixed field of H. The inverse map is just $L \mapsto \operatorname{Gal}(M/L)$ with the canonical inclusion into $\operatorname{Gal}(M/K)$. Moreover, $L = M^H$ is a Galois extension of K if and only if H is normal in G, in which case every element of $\operatorname{Gal}(M/K)$ restricts to an element of $\operatorname{Gal}(L/K)$. So the (surjective) restriction homomorphism $\operatorname{Gal}(M/K) \to \operatorname{Gal}(L/K)$ induces an isomorphism $\operatorname{Gal}(L/K) \cong \operatorname{Gal}(M/K)/\operatorname{Gal}(M/L) = \operatorname{Gal}(M/K)/H$.

Proof. See [1, §4.1]. □

Although we can do finite Galois theory just using the definition $|\operatorname{Aut}(L/K)| = [L:K]$, it is necessary in our case to use a more general definition:

Definition 1.10. A **Galois** extension L/K is an algebraic extension which is normal and separable.

Remark 1.11. Theorem 1.9 is false as stated for infinite Galois extensions (see future Example 2.1).

If M/K is a (possibly infinite) Galois extension, then every $x \in M$ lies in some finite Galois extension of K (take the splitting field of the minimal polynomial of x over K). So specifying an element of $\operatorname{Gal}(M/K)$ amounts to specifying its action on each finite Galois subextension[5]. This can be stated in a useful way using an inverse limit:

Theorem 1.12

Let M/K be Galois and $\mathcal{L}$ the set of finite Galois subextensions L/K. For $L \in \mathcal{L}$, the groups $\operatorname{Gal}(L/K)$ form an inverse system, equipped with the projection maps $\operatorname{Gal}(L/K) \to \operatorname{Gal}(L'/K)$ whenever $L' \subseteq L$. We have

$$\operatorname{Gal}(M/K) \cong \varprojlim_{L\in\mathcal{L}} \operatorname{Gal}(L/K) := \left\{ (\sigma_L)_{L\in\mathcal{L}} \in \prod_{L\in\mathcal{L}} \operatorname{Gal}(L/K) : L' \subseteq L \implies \sigma_L|_{L'} = \sigma_{L'} \right\},$$

where the isomorphism is given by restriction to each factor $\operatorname{Gal}(L/K)$.

Proof. Let $\sigma \in \operatorname{Gal}(M/K)$. It's obvious that restricting σ to each subextension $L \in \mathcal{L}$ yields a valid group homomorphism

$$\operatorname{Gal}(M/K) \to \varprojlim_{L\in\mathcal{L}} \operatorname{Gal}(L/K).$$

[4] When K (local or global) has positive characteristic, it is not true that $\overline{K}/K$ is Galois, since it isn't separable (this is because K will not be perfect). In this situation, we can replace $\overline{K} = K^{\mathrm{alg}}$ with the *separable closure* $K^{\mathrm{sep}} \subset K^{\mathrm{alg}}$ of K to get the Galois extension K^{sep}/K. This sublety will never be important for us in this course so we might frequently write $\overline{K}$ when this will need to be understood in some cases as K^{sep}.

[5] This morally why infinite Galois theory is often said to be a "formal consequence of the finite case".

If σ is trivial on each $L \in \mathcal{L}$, then $\sigma(x) = x$ for all $x \in M$ (each such x is contained in some $L \in \mathcal{L}$ as already mentioned), which means $\sigma = \mathrm{id}$. So the homomorphism is injective. To prove surjectivity, let $(\sigma_L)_{L\in\mathcal{L}} \in \varprojlim \operatorname{Gal}(L/K)$. Now define $\sigma \in \operatorname{Gal}(M/K)$ by $\sigma(x) = \sigma_L(x)$ for any $L \in \mathcal{L}$ containing x. We need to check two things:

1. The definition of $\sigma(x)$ does not depend on the choice of $L \in \mathcal{L}$ containing x.

2. σ is a bona fide element of $\operatorname{Gal}(M/K)$.

The first point is guaranteed by the definition of the inverse limit (since that definition guarantees that all the σ_L's agree on the Galois closure K' of $K(x)/K$ because they must all restrict to $\sigma_{K'}$). For the second one, let $x, y \in M$, and let L be a finite Galois extension of K contained in M containing both x and y (just take the Galois closure of $K(x,y)$). Then $\sigma(xy) = \sigma_L(xy) = \sigma_L(x)\sigma_L(y) = \sigma(x)\sigma(y)$ as desired. □

Now we can discuss two basic examples: finite fields and cyclotomic extensions of $\mathbf{Q}$.

Example 1.13

Let q be a prime power, and consider the extension $\overline{\mathbf{F}_q}/\mathbf{F}_q$. This extension is clearly Galois (since $\mathbf{F}_p$ is perfect and by definition of algebraic closure the extension is normal). Recall that the finite extensions of $\mathbf{F}_q$ are all Galois, all equal to $\mathbf{F}_{q^n}$ for some n, and with cyclic Galois group $\mathbf{Z}/n\mathbf{Z}$ generated by the Frobenius automorphism $x \mapsto x^q$. As a result, by the general description of the Galois group of an infinite extension in Theorem 1.12,

$$\operatorname{Gal}(\overline{\mathbf{F}_q}/\mathbf{F}_q) \cong \varprojlim_{n\in\mathbf{N}} \mathbf{Z}/n\mathbf{Z},$$

where the inverse system is given by the natural modulo-m projection $\mathbf{Z}/n\mathbf{Z} \to \mathbf{Z}/m\mathbf{Z}$ whenever $m|n$, since we have $\mathbf{F}_{q^m} \subseteq \mathbf{F}_{q^n}$ if and only if $m|n$ and the Frobenius in $\operatorname{Gal}(\mathbf{F}_{q^n}/\mathbf{F}_q)$ restricts to the Frobenius in $\operatorname{Gal}(\mathbf{F}_{q^m}/\mathbf{F}_q)$ (they are both defined by exponentiation by q), and therefore the restriction maps are the ones $\mathbf{Z}/n\mathbf{Z} \to \mathbf{Z}/m\mathbf{Z}$ for $m|n$ sending $1 + n\mathbf{Z}$ to $1 + m\mathbf{Z}$, as claimed.

By the Chinese remainder theorem, we can see that[a]

$$\varprojlim_{n\in\mathbf{N}} \mathbf{Z}/n\mathbf{Z} = \prod_p \varprojlim_{n\in\mathbf{N}} \mathbf{Z}/p^n\mathbf{Z} = \prod_p \mathbf{Z}_p$$

where the product is over all rational primes p. Next lecture we will see why this provides a counterexample to the infinite case of Theorem 1.9.

[a] It's a simple-enough exercise to exhibit the isomorphism and prove that it works.

$\operatorname{Gal}(\overline{\mathbf{F}}_q/\mathbf{F}_q) \cong \varprojlim \mathbf{Z}/n\mathbf{Z} \cong \prod_p \mathbf{Z}_p$ is called the *profinite completion* of $\mathbf{Z}$, also denoted $\hat{\mathbf{Z}}$. In general, the profinite completion of a group G is the inverse limit of the G/H (with the obvious inverse system structure), where H runs over all normal subgroups of finite index.

Example 1.14

In the previous example, we took $\overline{\mathbf{F}}_q$, which is a union of finite extensions of degree n with Galois group $\mathbf{Z}/n\mathbf{Z}$. In this example, we can consider the union of the fields $\mathbf{Q}(\zeta_n)$ where ζ_n is a primitive n-th root of unity. The union over all n of $\mathbf{Q}(\zeta_n)$ is a field because if $x \in \mathbf{Q}(\zeta_n)$ and $y \in \mathbf{Q}(\zeta_m)$ then $xy, x+y \in \mathbf{Q}(\zeta_n, \zeta_m) = \mathbf{Q}(\zeta_{nm})$. Letting this union of fields be $\mathbf{Q}(\zeta_\infty)$, we know

$$\mathrm{Gal}(\mathbf{Q}(\zeta_\infty)/\mathbf{Q}) \cong \varprojlim_n (\mathbf{Z}/n\mathbf{Z})^\times,$$

since each finite subfield of $\mathbf{Q}(\zeta_\infty)/\mathbf{Q}$ is contained in some $\mathbf{Q}(\zeta_n)/\mathbf{Q}$ (for instance, by the primitive element theorem, any such finite subfield is of the form $\mathbf{Q}(\alpha)$ for some $\alpha \in \mathbf{Q}(\zeta_\infty)$, which means by definition $\alpha \in \mathbf{Q}(\zeta_n)$ for some n and therefore this finite subfield is contained in $\mathbf{Q}(\zeta_n)$). By the Chinese remainder theorem, this is the same as $\prod_p \varprojlim_n (\mathbf{Z}/p^n\mathbf{Z})^\times = \prod_p \mathbf{Z}_p^\times = \widehat{\mathbf{Z}}^\times$. Here we view $\widehat{\mathbf{Z}}$ as a ring and take its unit group.

§2 September 5, 2019

Today we will fix the fundamental theorem of Galois theory to work for infinite extensions.

§2.1 More infinite Galois theory

First, we finish off the counterexample from last time to the usual statement.

Example 2.1

Let $G = \mathrm{Gal}(\overline{\mathbf{F}_q}/\mathbf{F}_q) \cong \widehat{\mathbf{Z}}$ as discussed in Example 1.13, and consider the subgroup of G given by $H = \langle \varphi_q \rangle \cong \mathbf{Z}$, where φ_p denotes the Frobenius automorphism, given by each coordinate equal to 1 (i.e. φ_q acts on each finite extension $\mathbf{F}_{q^n}/\mathbf{F}_q$ via $x \mapsto x^q$). The inclusion $H \subset G$ produces the canonical inclusion $\mathbf{Z} \subseteq \widehat{\mathbf{Z}}$.

The fixed field of H is the set of all $x \in \overline{\mathbf{F}}_q$ such that $x^q = x$. The polynomial $x^q - x$ has degree q, so its roots are precisely the elements of $\mathbf{F}_q$. So the fixed field is $\mathbf{F}_q$. But H is not all of G (for instance because the inclusion $\mathbf{Z} \to \widehat{\mathbf{Z}} = \prod_p \mathbf{Z}_p$ is given by inclusion into each coordinate, and $\mathbf{Z}_p$ properly contains $\mathbf{Z}$)[a]. This means that Theorem 1.9 does not extend as-stated to the infinite case.

[a] The fact that $\mathbf{Z} \subsetneq \mathbf{Z}_p$ is easy to see in all sorts of ways. So far we have defined $\mathbf{Z}_p = \varprojlim \mathbf{Z}/p^r\mathbf{Z}$, and we can just observe that there is no integer congruent to $\sum_{i=0}^{r-1}(-p)^i \mod p^r$ since such an integer a would have the property that $(p+1)a \equiv 1 \mod p^r$ for all r, and therefore $(p+1)a = 1$ which is impossible if $a \in \mathbf{Z}$. This is a special case of the fact that $\mathbf{Z}_p$ contains the localization of $\mathbf{Z}$ away from (p). In fact we have strict inclusions $\mathbf{Z} \subsetneq \mathbf{Z}_{(p)} \subsetneq \mathbf{Z}_p$.

The rest of the lecture will be devoted to stating and proving the fundamental theorem of Galois theory for infinite extensions, via a convenient topology, called the *Krull topology*, on the Galois group that comes from the inverse limit structure given in Theorem 1.12.

Let M/K be a (possibly infinite) Galois extension. For intermediate extensions $K \subseteq L \subseteq M$, we have

$$\mathrm{Gal}(M/L) = \bigcap_{L' \subseteq L} \mathrm{Gal}(M/L'),$$

where L' runs over the finite extensions of K contained in L. This motivates us to look for a topology on $\mathrm{Gal}(M/K)$ such that $\mathrm{Gal}(M/L) \subseteq \mathrm{Gal}(M/K)$ is closed for all such L (since the arbitrary intersection of closed sets is closed). In particular, for any subgroup $H \subseteq G$, we want $\mathrm{Gal}(M/M^H)$ to be the closure of H. We also want $\mathrm{Gal}(M/K)$ to be a topological group, so we can't just take these to be exactly the closed subgroups. Instead we take the unique topological group structure on $\mathrm{Gal}(M/K)$ with closed subgroups around the identity given by the $\mathrm{Gal}(M/L)$. This will be a compact topological group, so (as we will see) the open subgroups are the same as the closed ones of finite index. Thus we should take as a basis of open subgroups around id the subgroups of the form $\mathrm{Gal}(M/L)$ where L/K is finite. In fact it's convenient to restrict the elements of the base around id to where L/K is finite and Galois (so that infinite Galois theory becomes more visibly connected to the finite case).

Definition 2.2. Define the **Krull topology** on $\mathrm{Gal}(M/K)$ using the base of open sets[6]

$$U_{\sigma,L} = \sigma\,\mathrm{Gal}(M/L) = \{\tau \in \mathrm{Gal}(M/K) : \tau|_L = \sigma|_L\}$$

[6] It's easy to check that this is a valid basis of open sets. In particular, $U_{\sigma_1,L_1} \cap U_{\sigma_2,L_2}$ is just the set of all $\tau \in \mathrm{Gal}(M/K)$ such that τ agrees with σ_1 on L_1 and with σ_2 on L_2. If this set is nonempty, it is equal to $U_{\sigma',L_1\cdot L_2}$ for any $\sigma' \in \mathrm{Gal}(M/K)$ restricting to σ_1 on L_1 and to σ_2 on L_2.

running over all $\sigma \in \operatorname{Gal}(M/K)$ and finite Galois[7] extensions L/K contained in M.

Roughly speaking, we should consider σ and τ to be close if they agree on a large finite Galois subextension. Here again it's useful that L/K is Galois so that σ and τ really do restrict to bona fide elements of $\operatorname{Gal}(L/K)$.

Example 2.3

If M/K is a finite Galois extension, then for any σ, we can set $L = M$ to see that $\operatorname{Gal}(M/K)$ has the discrete topology (every subset is closed and open).

In fact, another way to define the Krull topology on $\operatorname{Gal}(M/K)$ is as the inverse limit

$$\varprojlim_{L\in\mathcal{L}} \operatorname{Gal}(L/K)$$

in the category of topological groups where $\operatorname{Gal}(L/K)$ is given the discrete topology (here $\mathcal{L}$ has the same meaning as it does in Theorem 1.12). That is to say, the Krull topology on $\operatorname{Gal}(M/K)$ is equal to the induced topology from the inclusion

$$\operatorname{Gal}(M/K) = \varprojlim \operatorname{Gal}(L/K) \subseteq \prod_{L\in\mathcal{L}} \operatorname{Gal}(L/K).$$

This is because the isomorphism from Theorem 1.12 $\operatorname{Gal}(M/K) \cong \varprojlim \operatorname{Gal}(L/K)$ sends the open set $U_{\sigma,F}$ to the set of elements of $\varprojlim \operatorname{Gal}(L/K)$ whose F-coordinate equals $\sigma|_F$. This open set is the intersection with $\varprojlim \operatorname{Gal}(L/K)$ of the open set of $\prod_L \operatorname{Gal}(L/K)$ given by

$$\left(\prod_{L\neq F} \operatorname{Gal}(L/K)\right) \times \{\sigma|_F\}.$$

So the isomorphism sends a basis for the Krull topology to a basis for the induced topology on the projective limit (since each finite Galois group $\operatorname{Gal}(L/K)$ has the discrete topology). So the resulting topology on $\varprojlim \operatorname{Gal}(L/K)$ is the same as the induced one from the product. Now armed with some basic knowledge about the Krull topology, we can proceed to prove the infinite Galois correspondence.

Proposition 2.4

If L is a subfield of M/K, then $M^{\operatorname{Gal}(M/L)} = L$.

Proof. Suppose that $x \in M \smallsetminus L$ is fixed by every $\sigma \in \operatorname{Gal}(M/L)$. Let L_x be a finite Galois extension of L containing x (e.g. the Galois closure of $L(x)/L$). By the finite Galois theory of L_x/L, there is some element $\overline{\sigma} \in \operatorname{Gal}(L_x/L)$ which does not fix x. But $\overline{\sigma}$ extends to an element of $\operatorname{Gal}(M/L)$ by Zorn's lemma, which contradicts the fact that x is fixed by $\operatorname{Gal}(M/L)$. So no $x \in M \smallsetminus L$ is fixed by $\operatorname{Gal}(M/L)$, which proves the only nontrivial part of the result, namely $M^{\operatorname{Gal}(M/L)} \subseteq L$. □

One part of the "Fundamental theorem of infinite Galois theory" is as follows:

[7] As noted before, the requirement that L/K is Galois doesn't change the topology. For any finite L/K, we have $\operatorname{Gal}(M/L) = \bigcup_{\sigma\in\operatorname{Gal}(M/L)} U_{\sigma,L'}$ where L' is the Galois closure of L/K so $\operatorname{Gal}(M/L)$ is still open under the Krull topology.

Theorem 2.5

Let H be a subgroup of $\operatorname{Gal}(M/K)$. Then

$$\operatorname{Gal}(M/M^H) = \overline{H}.$$

Proof. Let $\sigma \in \operatorname{Gal}(M/M^H)$. To show that $\operatorname{Gal}(M/M^H) \subseteq \overline{H}$, we need to show that any open neighborhood of σ has nontrivial intersection with H. Such an open neighborhood contains one of the form[8] $\sigma \operatorname{Gal}(M/T) = U_{\sigma,T}$ for some finite Galois T/K. So fix a finite Galois extension T/K. Then $M^H \cap T = T^H$, so

$$\sigma|_T \in \operatorname{Gal}(T/T^H).$$

Since $\operatorname{Gal}(T/K)$ is finite, we know $\operatorname{Gal}(T/T^H)$ is the set of restrictions of H to T (every element of $\operatorname{Gal}(T/K)$ is a restriction from $\operatorname{Gal}(M/K)$; by finite Galois theory, the elements of $\operatorname{Gal}(T/T^H)$ are exactly the elements of $\operatorname{Gal}(T/K)$ which are restrictions from H). So there is an element $\tau \in H$ which agrees with σ on T. Hence, $\tau \in H \cap U_{\sigma,T}$ which means that (since T was chosen arbitrarily amongst finite Galois intermediate extensions of M/K) any open neighborhood of $\sigma \in \operatorname{Gal}(M/M^H)$ has nonempty intersection with H. This proves the inclusion $\operatorname{Gal}(M/M^H) \subseteq \overline{H}$.

For the other inclusion, let $\sigma \in \operatorname{Gal}(M/K) \setminus \operatorname{Gal}(M/M^H)$. So there is some $x \in M^H$ such that $\sigma(x) \neq x$. We can take $T \subseteq M$ to be a finite Galois extension of K containing x. Then $U_{\sigma,T}$ has no intersection with H, because any element of $U_{\sigma,T}$ has to agree with σ when restricted to T, and no element of H does this (x is fixed by H but not by σ). Therefore, $\operatorname{Gal}(M/M^H)$ is closed. Combined with the inclusions

$$H \subseteq \operatorname{Gal}(M/M^H) \subseteq \overline{H}$$

we already have, it follows that

$$\operatorname{Gal}(M/M^H) = \overline{H}$$

as desired. □

So in the more general Galois correspondence (with the maps defined in the usual way), the closed subgroups of $\operatorname{Gal}(M/K)$ correspond to the intermediate extensions. The remaining part of the Galois correspondence is also still true in this scenario, and for the exact same reasons (in fact we're done dealing with the Krull topology; everything from now is a restatement of the usual arguments). Suppose that a closed subgroup $H \subseteq \operatorname{Gal}(M/K)$ has $L = M^H$ and thus $\operatorname{Gal}(M/L) = H$. The embeddings $\sigma : L \to M$ fixing K correspond to cosets of H in $\operatorname{Gal}(M/K)$. Since M/K is Galois, we know L/K is separable, so if L is finite then $L = K(\alpha)$ for some $\alpha \in L \subset M$. In this case the minimal polynomial for α splits completely into $[L:K]$ distinct factors over M (since M/K is normal), so in fact there are $[L:K]$ such embeddings σ and we have $[\operatorname{Gal}(M/K) : H] = [L:K]$. If L/K is infinite then by looking at finite subextensions (and extending the embeddings thereof) we see that $[\operatorname{Gal}(M/K) : H] = \infty$ as well. So

$$[\operatorname{Gal}(M/K) : H] = [L:K]$$

in the extended sense. So in fact, under the Galois correspondence the closed subgroups of finite index correspond to the finite intermediate extensions. Finally, the normal subgroups still correspond to the Galois subextensions, for the reason that the conjugate subgroup $\sigma H \sigma^{-1}$ corresponds to $\sigma(L)$. So we have proved:

[8] It must contain some basis element $\sigma \in U_{\tau,T}$ and thus $U_{\tau,T} = U_{\sigma,T}$ by the definition of these sets.

Theorem 2.6 (Infinite Galois correspondence)

Let M/K be a Galois extension. There is a natural inclusion-reversing bijection of sets

$$\{\text{closed subgroups of } \mathrm{Gal}(M/K)\} \to \{\text{intermediate fields } K \subseteq L \subseteq M\}$$

given by

$$H \mapsto M^H,$$

with inverse given by

$$L \mapsto \mathrm{Gal}(M/L).$$

Under this bijection, the closed subgroups of finite index correspond to the finite extensions L/K (more precisely the degree of the extension is equal to the index of the subgroup), and H is normal in $\mathrm{Gal}(M/K)$ if and only if the corresponding L/K is Galois, in which case we have $\mathrm{Gal}(L/K) \cong \mathrm{Gal}(M/K)/H$.

§2.2 Galois groups as topological groups

In fact, the Krull topology also gives $\mathrm{Gal}(M/K)$ the structure of a *topological group*.

Definition 2.7. A group G with a topology is a **topological group** if the multiplication map $G \times G \to G$ and the inversion map $G \to G$ are continuous.

Someone asked whether multiplication $G \times G \to G$ being continuous is the same as the multiplication by g map $G \to G$ being continuous for each $g \in G$. It turns out there is a counterexample (on Terence Tao's blog[9])

Example 2.8

Let $G = (\mathbf{R}, +)$ equipped with the cocompact topology (where the open sets are the sets whose complements are compact, plus the empty set). Then the addition by a fixed element is clearly continuous but you can check directly that the map $G \times G \to G$ is not.

Lemma 2.9

Let M/K be an algebraic field extension. Then $\mathrm{Gal}(M/K)$ is a topological group under the Krull topology.

Proof. It suffices to show that the preimage of any basis element is open, under the multiplication and inverse maps. First the inverse map. We claim that the preimage of $U_{\sigma,L}$ under the inverse map is $U_{\sigma^{-1},L}$. This is clear because an element $\tau \in \mathrm{Gal}(M/K)$ has the property that $\tau|_L = \sigma|_L$ if and only if $\tau^{-1}|_L = \sigma^{-1}|_L$. For the multiplication map, note that the $\tau_1, \tau_2 \in \mathrm{Gal}(M/K)$ such that $\tau_1\tau_2|_L = \sigma|_L$ are exactly the (τ_1, τ_2) restricting to $(\tau_1|_L, \tau_2|_L)$ such that $(\tau_1|_L)(\tau_2|_L) = \sigma|_L$. In particular, we want the preimage of $\sigma|_L$ under the multiplication map $\mathrm{Gal}(L/K) \times \mathrm{Gal}(L/K) \to \mathrm{Gal}(L/K)$ composed on the right with restriction from $\mathrm{Gal}(M/K) \times \mathrm{Gal}(M/K)$. Since these Galois groups have the discrete topology, in fact $\mathrm{Gal}(L/K) \times \mathrm{Gal}(L/K)$ is also discrete so the multiplication map

$$\mathrm{Gal}(L/K) \times \mathrm{Gal}(L/K) \to \mathrm{Gal}(L/K)$$

[9] See https://math.stackexchange.com/questions/812513

is forced to be continuous. It remains now to show the continuity of the restriction map

$$\mathrm{Gal}(M/K) \times \mathrm{Gal}(M/K) \to \mathrm{Gal}(L/K) \times \mathrm{Gal}(L/K).$$

But this is obvious from the definition[10] of the Krull topology, which essentially says that the restriction homomorphism $\mathrm{Gal}(M/K) \to \mathrm{Gal}(L/K)$ is continuous. □

Lemma 2.10

Let G be a topological group, and $U \subseteq G$ an arbitrary open subset. Then gU is open for any $g \in G$, as is U^{-1} (the set of inverses of elements of U).

Proof. This follows directly from the fact that the inverse and multiplication maps are continuous, and the fact that G has inverses (so we can consider the multiplication map by g^{-1} for any $g \in G$). □

Proposition 2.11

If G is a compact topological group, then a subgroup $H \subseteq G$ is open if and only if H is closed and has finite index.

Proof. If H is open, then it is closed because $G = \bigcup_{g \in G} gH$ where this union can be made disjoint by only taking distinct cosets of H. In particular, H is one of these cosets, so complement of H is a union of open sets (by Lemma 2.10) which means H is closed (this is a general fact about topological groups and does not depend on compactness). Since G is compact, it must also have finite index (since we have written G as a disjoint union of open sets which are cosets of H, so any finite subcover must be the entire cover and therefore there are finitely many cosets). Conversely, suppose H is closed of finite index. Then the same coset decomposition says that the complement of H is a finite union of closed subsets (by Lemma 2.10), so H is open. □

Proposition 2.12

$G = \mathrm{Gal}(M/K)$ is Hausdorff, compact, and totally disconnected

Proof. G is Hausdorff because any two distinct σ and σ' in G are separated by the open sets specified by a finite extension on which they disagree. In particular, there is some $x \in M$ such that $\sigma(x) \neq \sigma'(x)$, so letting L be the Galois closure of $K(x)$ we have $U_{\sigma,L} \cap U_{\sigma',L} = \varnothing$.

To show it is totally disconnected, it suffices to show that the basis elements $U_{\sigma,L}$ are clopen[11]. This is because $U_{\sigma,L} = \sigma U_{1,L}$ and $U_{1,L}$ is an open subgroup and therefore[12] closed by Proposition 2.11. So by Lemma 2.10, $U_{\sigma,L}$ is clopen as desired.

[10] It's definitely obvious if you accept the definition of the Krull topology as being an inverse limit in the category of topological groups. Even using the original Definition 2.2 it's obvious, because the preimage of a single $\sigma \in \mathrm{Gal}(L/K)$ is $U_{\sigma',L}$ for a $\sigma' \in \mathrm{Gal}(M/K)$ extending σ.

[11] Since G is Hausdorff, this means that any two points are in different connected components

[12] The fact that open implies closed doesn't depend on compactness, but it doesn't matter since we are about to prove G is compact in the next paragraph.

Finally, by Tychonoff's theorem, $\prod_L \operatorname{Gal}(L/K)$ is compact, so it suffices to show that $\varprojlim \operatorname{Gal}(L/K)$ is closed in this product[13]. But it is defined just by imposing on

$$(\sigma_L)_{L\in\mathcal{L}} \in \prod_{L\in\mathcal{L}} \operatorname{Gal}(L/K)$$

conditions of the form $\sigma_{L_2}|_{L_1} = \sigma_{L_1}$ whenever $L_1 \subseteq L_2$. Since arbitrary intersections of closed sets are closed, it suffices to show that imposing one such restriction results in a closed set. So consider such a subset

$$S = \left\{ (\sigma_L)_{L\in\mathcal{L}} \in \prod_{L\in\mathcal{L}} : \sigma_{L_2}|_{L_1} = \sigma_{L_1} \right\}$$

for some fixed choice of $L_1 \subseteq L_2$ in $\mathcal{L}$. Let $(\sigma_L)_{L\in\mathcal{L}} \in \prod_L \operatorname{Gal}(L/K) \smallsetminus S$. This means that

$$\sigma_{L_2}|_{L_1} \neq \sigma_{L_1}.$$

The open set

$$\left(\prod_{L\in\mathcal{L}\smallsetminus\{L_1,L_2\}} \operatorname{Gal}(L/K) \right) \times \{\sigma_{L_1}\} \times \{\sigma_{L_2}\},$$

where the $\{\sigma_{L_1}\}$ is in the L_1-coordinate and the $\{\sigma_{L_2}\}$ is in the L_2-coordinate, contains $(\sigma_L)_{L\in\mathcal{L}}$ and has no intersection with S since $\sigma_{L_2}|_{L_1} \neq \sigma_{L_1}$. This proves that S is closed, and therefore $\operatorname{Gal}(M/K)$ is compact as desired. □

So in the infinite Galois correspondence we see that the closed subgroups of $\operatorname{Gal}(M/K)$ correspond to the subextensions, and the open subgroups (i.e. closed of finite index) correspond to the finite subextensions.

§2.3 Example: algebraic extensions of a finite field

Now we will observe the Galois correspondence in action by seeing what it says about finite fields, namely intermediate extensions of $\overline{\mathbf{F}}_q/\mathbf{F}_q$. As already remarked, by Theorem 2.6 and Proposition 2.11, the closed subgroups of $\operatorname{Gal}(\overline{\mathbf{F}}_q/\mathbf{F}_q)$ correspond to the intermediate extensions, and the open subgroups correspond to the intermediate extensions which are finite over $\mathbf{F}_q$. Recall from Example 1.13 and the fact from earlier today that the isomorphism from Theorem 1.12 is a homeomorphism as well that we have an isomorphism of topological groups[14]

$$\operatorname{Gal}(\overline{\mathbf{F}}_q/\mathbf{F}_q) \cong \widehat{\mathbf{Z}} \cong \prod_p \mathbf{Z}_p.$$

So to understand the topology on $\widehat{\mathbf{Z}}$ better, we should really go back to $\mathbf{Z}_p$. It's clear that $\mathbf{Z}$ is dense in $\mathbf{Z}_p$ (for example if $\alpha \in \mathbf{Z}_p$ for any $n \geq 0$ there is an $a_n \in \mathbf{Z}$ which reduces to the same thing mod p^n as the projection of α onto $\mathbf{Z}/p^n\mathbf{Z}$ and the a_n's converge to α in $\mathbf{Z}_p$ by definition). We want to determine the closed subgroups of $\mathbf{Z}_p$. If a subgroup contains $x \in \mathbf{Z}_p$, then (since $\mathbf{Z}_p$ is an abelian group) it contains $x\mathbf{Z} \subseteq \mathbf{Z}_p$. The closure of this set is $x\mathbf{Z}_p$ (if $y \in \mathbf{Z}_p$ and $a_i \to y$ is a sequence of elements of $\mathbf{Z}$ converging to y, then xa_i converges to xy[15]). Moreover, $x\mathbf{Z}_p = p^{v_p(x)}\mathbf{Z}_p$ since any element of $\mathbf{Z}_p$ not divisible

[13] This is because we have already noted that the Krull topology on $\operatorname{Gal}(M/K)$ is the same as the topology induced from the topology on $\varprojlim \operatorname{Gal}(L/K)$ via the isomorphism of Theorem 1.12.

[14] We technically haven't justified the fact that the group isomorphism $\widehat{\mathbf{Z}} \cong \prod_p \mathbf{Z}_p$ is also a homeomorphism, but this is also a simple exercise.

[15] This uses the topological ring structure of $\mathbf{Z}_p$, which comes from the definition as an inverse limit of rings

by p is a unit in $\mathbf{Z}_p$ (since it is a unit in each $\mathbf{Z}/p^n\mathbf{Z}$). Note that it's possible that $x = 0$, in which case $v_p(x) = \infty$ and $x\mathbf{Z}_p = 0$ which is how we should think of $p^{v_p(x)}\mathbf{Z}_p$ when $v_p(x) = \infty$. So if H is a closed subgroup then it equal to

$$\bigcup_{x \in H} p^{v_p(x)}\mathbf{Z}_p,$$

which is already of the form $p^{v_p(x)}\mathbf{Z}_p$ for some x, namely any $x \in H$ that minimizes $v_p(x)$. So the closed subgroups of $(\mathbf{Z}_p, +)$ are just the subgroups of the form $p^v\mathbf{Z}_p$ where $v \in \mathbf{Z}_{\geq 0} \cup \{\infty\}$. We really want to know what the closed subgroups of $\widehat{\mathbf{Z}} = \prod_p \mathbf{Z}_p$ are. We can repeat the same argument[16] to conclude what we want about $\widehat{\mathbf{Z}}$:

Lemma 2.13

The closed subgroups of $\widehat{\mathbf{Z}} \cong \prod_p \mathbf{Z}_p$ are precisely those of the form

$$\prod_p p^{e_p}\mathbf{Z}_p$$

where each $e_p \in \mathbf{Z}_{\geq 0} \cup \{\infty\}$.

Proof. Let H be a closed subgroup of $\widehat{\mathbf{Z}} \cong \prod_p \mathbf{Z}_p$. For any $x \in H$, H contains the closure of $x\mathbf{Z} \subseteq \widehat{\mathbf{Z}}$. Since $\widehat{\mathbf{Z}}$ is a topological ring and $\mathbf{Z}$ is dense in $\widehat{\mathbf{Z}}$ (by essentially the same argument for why $\mathbf{Z}$ is dense in $\mathbf{Z}_p$), the closure of $x\mathbf{Z}$ is

$$x\widehat{\mathbf{Z}} \cong x \prod_p \mathbf{Z}_p.$$

Representing $x \in \prod_p \mathbf{Z}_p$ as the tuple $(x_p)_p$, this is just

$$\prod_p p^{v_p(x_p)}\mathbf{Z}_p.$$

So since it is closed, H is the subgroup of $\prod_p \mathbf{Z}_p$ generated by the subgroups $\prod_p p^{v_p(x_p)}\mathbf{Z}_p$ where x runs over all $x \in H$. But in fact

$$\prod_p p^{a_p}\mathbf{Z}_p + \prod_p p^{b_p}\mathbf{Z}_p = \prod_p p^{\min(a_p,b_p)}\mathbf{Z}_p$$

which proves that H is of the desired form (in particular it is $\prod_p p^{e_p}\mathbf{Z}_p$ where e_p is the minimum of all the x_p's for $x \in H$). □

Now we can apply the Galois correspondence Theorem 2.6. The finite-index closed subgroups, in the language of Lemma 2.13, are those for which all all of the e_p are finite and all but finitely many of them are zero (since $|\mathbf{Z}_p/p^e\mathbf{Z}_p| = p^e$). In other words, they are of the form $n\widehat{\mathbf{Z}}$ for integers n. This is consistent with what the Galois correspondence says: the finite extensions of $\mathbf{F}_q$ are the fields $\mathbf{F}_{q^n}$. And the fixed field of $n\widehat{\mathbf{Z}}$ is $\mathbf{F}_{q^n}$ since $n\widehat{\mathbf{Z}}$ is the closure of $n\mathbf{Z}$, whose fixed field we know is $\mathbf{F}_{q^n}$ since $\mathbf{Z} \subset \widehat{\mathbf{Z}}$ is taken to be generated by the Frobenius and $\mathbf{F}_{q^n}$ is indeed defined to be the set of elements $x \in \overline{\mathbf{F}}_q$ such that $x^{q^n} = x$ (or alternatively because $n\widehat{\mathbf{Z}}$ has index n in $\widehat{\mathbf{Z}}$ so by Theorem 2.6 the fixed field is the unique degree-n extension of $\mathbf{F}_q$).

[16] As far as I can tell there's no way to deduce this directly from the description of the closed subgroups of $\mathbf{Z}_p$. In general the direct product of groups can of course have subgroups which don't come from taking the direct product of subgroups.

Moreover, via the isomorphism $\widehat{\mathbf{Z}} \cong \prod_p \mathbf{Z}_p$, a tuple $(x_p)_p \in \prod_p \mathbf{Z}_p$ acts on $\mathbf{F}_{q^{p^n}} \subseteq \overline{\mathbf{F}}_q$ via $\alpha \mapsto \alpha^{q^{x_p \bmod p^n}}$. So the fixed field of

$$\prod_p p^{e_p} \mathbf{Z}_p$$

is the compositum of all the $\mathbf{F}_{q^{p^{e_p}}}$, which is understood to mean $\bigcup_{n \geq 0} \mathbf{F}_{q^{p^n}}$ when $e_p = \infty$. So this is what the algebraic extensions of $\mathbf{F}_q$ are (of course this is perfectly clear without the use of the infinite Galois correspondence, but serves as an instructive example of how the theorem works).

§3 September 10, 2019

Today we will start to review some basic notions from introductory algebraic number theory (e.g. math 129), though perhaps in a slightly more general setting than usual.

§3.1 Algebraic number theory basics: Dedekind domains

Definition 3.1. A **Dedekind domain** is an integral domain in which every nonzero ideal factors (uniquely) as a product of nonzero prime ideals.

In particular, in a Dedekind domain R, every fractional ideal is invertible and factors as a product of integer powers of nonzero primes.

Example 3.2

All PID's are Dedekind domains. For example $\mathbf{Z}$ and $k[T]$ are Dedekind.

Example 3.3

Let K be a number field. The ring of integers $\mathcal{O}_K$ is a Dedekind domain (see [3, Theorem 14] or [5, I.3.1])

Recall that a Dedekind domain is integrally closed in its field of fractions. Indeed, recall that another equivalent[17] definition of Dedekind domain is as follows:

Definition 3.4. An integral domain R is a **Dedekind domain** if and only if all of the following conditions are met:

1. R is integrally closed in its field of fractions.
2. R is Noetherian.
3. R has Krull dimension 1 (all nonzero primes are maximal).

We will typically work in the following setup: let K be a field, $\mathcal{O}_K$ a Dedekind domain with field of fractions K, L/K an algebraic extension of K, and $\mathcal{O}_L$ the integral closure of $\mathcal{O}_K$ in L.

Lemma 3.5 (Krull–Akizuki Theorem)

If L/K is finite, then $\mathcal{O}_L$ is a Dedekind domain.

Proof. It's surprising that this is true without any further conditions on L/K, but it is: see [5, I.8.1]. This is easier to see when L/K is separable, So we'll recall the proof conditional on the separability of L/K. By Definition 3.4, it suffices to show that $\mathcal{O}_L$ is Noetherian and has Krull dimension 1 (it is the integral closure of $\mathcal{O}_K$ in L, so it already satisfies the condition of being integrally closed in its field of fractions[18]). Let $\mathfrak{P}$ be a nonzero prime ideal in $\mathcal{O}_L$. Since $1 \notin \mathfrak{P}$, we know that $\mathfrak{p} = \mathcal{O}_K \cap \mathfrak{P}$ is a proper ideal in $\mathcal{O}_K$. Also,

[17] The fact that Definition 3.4 implies Definition 3.1 is a basic fact from math 129; see for instance [5, I.3.3] or [3, Theorem 16]. The other direction is more obscure. A proof can be found in `http://www.math.uchicago.edu/~may/MISC/Dedekind.pdf`.

[18] This is because of the standard fact from commutative algebra that if $A \subset B \subset C$ are rings and $c \in C$ is integral over B and B is integral over A, then C is integral over A.

it is a nonzero ideal of $\mathcal{O}_K$ because there exists a nonzero $\alpha \in \mathfrak{P}$, which is integral over $\mathcal{O}_K$, so there is a monic irreducible polynomial $f(X) = X^n + a_{n-1}X^{n-1} + \cdots + a_0 \in \mathcal{O}_K[X]$ such that $f(\alpha) = 0$. Since α is nonzero and f is irreducible, f has nonzero constant term, and therefore

$$\alpha^n + a_{n-1}\alpha^{n-1} + \cdots + a_1\alpha = -a_0 \neq 0.$$

This means $-a_0 \in \mathfrak{P}$ since α is, but since a_0 is a nonzero element of $\mathcal{O}_K$ it follows that $\mathfrak{p}$ is a nonzero prime of $\mathcal{O}_K$. Since $\mathcal{O}_K$ is Dedekind, we know $\mathfrak{p}$ is a maximal ideal of $\mathcal{O}_K$. So in the inclusion

$$\mathcal{O}_K/\mathfrak{p} \subset \mathcal{O}_L/\mathfrak{P}$$

we know that $\mathcal{O}_K/\mathfrak{p}$ is a field, and that all the elements of the integral domain $\mathcal{O}_L/\mathfrak{P}$ are algebraic over the field $\mathcal{O}_K/\mathfrak{p}$. This is enough to conclude that $\mathcal{O}_L/\mathfrak{P}$ is a field and therefore $\mathfrak{P}$ is maximal. One argument is that adjoining an algebraic element to a field always yields a field, so $\mathcal{O}_L/\mathfrak{P}$ is a direct limit of fields and is therefore a field. Alternatively, if $\mathcal{O}_L/\mathfrak{P}$ is not a field, then it has a nonzero maximal ideal $\mathfrak{m}$, and the fact that the extension is algebraic shows via the same argument as before that $\mathfrak{m} \cap \mathcal{O}_K/\mathfrak{p}$ is a nonzero prime ideal of $\mathcal{O}_K/\mathfrak{p}$, but since this is a field we have a contradiction. Notice that the Krull dimension of $\mathcal{O}_L$ being 1 did not depend on the finiteness or separability of L/K.

Now it remains to show that $\mathcal{O}_L$ is Noetherian, which is where the properties we assumed of L/K come in. Recall that when L/K is separable, we have a nondegenerate trace pairing $\langle\cdot,\cdot\rangle : L \times L \to K$ given by $\langle x, y\rangle$ (this is easiest to see in characteristic zero but is actually a characterizing property of separable extensions[19]). By clearing denominators, we have a basis $\alpha_1, \ldots, \alpha_n \in \mathcal{O}_L$ for L/K. Since the trace pairing is nondegenerate, this basis has a dual basis $\alpha_1^\vee, \ldots, \alpha_n^\vee \in L^\vee$, so we have inclusions

$$\alpha_1\mathcal{O}_K \oplus \cdots \oplus \alpha_n\mathcal{O}_K \subset \mathcal{O}_L \subset \mathcal{O}_L^\vee \subset \alpha_1^\vee\mathcal{O}_K \oplus \cdots \oplus \alpha_n^\vee\mathcal{O}_K$$

from which it follows that $\mathcal{O}_L$ is stuck between two free $\mathcal{O}_K$-modules of rank n. This does not mean that $\mathcal{O}_L$ is free unless $\mathcal{O}_K$ is a PID[20], but just using the fact that $\mathcal{O}_L$ is a submodule of a free $\mathcal{O}_K$-module, we immediately see that $\mathcal{O}_L$ is Noetherian (since direct sums and submodules of Noetherian modules are Noetherian, so $\mathcal{O}_L$ is Noetherian as an $\mathcal{O}_K$-module and therefore also as a $\mathcal{O}_L$-module). □

Example 3.6

When L/K is an extension of number fields and $\mathcal{O}_K$ is the ring of integers of K, we know that the integral closure of $\mathcal{O}_K$ is the ring of integers $\mathcal{O}_L$ in L, and is therefore Dedekind as stated in Example 3.3.

§3.2 Algebraic number theory basics: Decomposition of primes

Let $\mathfrak{p}$ be a prime of $\mathcal{O}_K$ (from now on prime always means nonzero prime ideal). Then there are possibly infinitely many primes of $\mathcal{O}_L$ containing $\mathfrak{p}$, if L/K is infinite. If L/K is Galois with Galois group G, then G acts on the primes of L lying over $\mathfrak{p}$. It is a basic result from math 129 when L/K is finite, G acts transitively on the primes lying over $\mathfrak{p}$ (see [5, I.9.1]). We prove the general case now, assuming the finite case.

[19] see for instance https://kconrad.math.uconn.edu/blurbs/galoistheory/separable2.pdf

[20] see https://kconrad.math.uconn.edu/blurbs/gradnumthy/notfree.pdf for some good counterexamples

Theorem 3.7

G acts transitively on the primes lying over $\mathfrak{p}$.

Proof. Let $\mathfrak{P}_1, \mathfrak{P}_2$ be primes lying over $\mathfrak{p}$. Suppose for the sake of contradiction that $\sigma(\mathfrak{P}_1) \neq \mathfrak{P}_2$ for all $\sigma \in G$. So there is some $x \in \mathcal{O}_L$ which is in $\mathfrak{P}_2 \setminus \sigma(\mathfrak{P}_1)$, and taking $F_\sigma \subset L$ to be a finite Galois extension of K containing x, we have for all $\sigma \in G$,

$$\sigma(\mathfrak{P}_1 \cap F_\sigma) \neq \mathfrak{P}_2 \cap F_\sigma$$

since the right hand side contains x but the left hand side does not. Using the language of Definition 2.2 to refer to the basis of open sets of G, for any $\sigma' \in U_{\sigma, F_\sigma}$, since σ' agrees with σ on F_σ we know the same equation above holds for σ' as well. By the compactness of G (Proposition 2.12), we can write

$$G = \bigcup_{i=1}^{n} U_{\sigma_i, F_{\sigma_i}}$$

for some finite collection of $\sigma_i \in G$ (since $U_{\sigma_i, F_{\sigma_i}}$ is an open set containing σ_i and therefore if we take σ_i to run over all elements of G we have an open cover which must have a finite subcover).

Let F be the compositum of the F_{σ_i}. Since F_{σ_i} is Galois and finite over K, we know that F/K is also a finite Galois extension. The point is that the compactness basically allows us use F uniformly for all σ instead of using an F_σ for each σ, which then allows us to invoke the finite version of the result. If $\tau \in \mathrm{Gal}(F/K)$, then it is the restriction to F of some $\tilde{\tau} \in G$, which means (by our expression for G as a finite union) that $\tilde{\tau}$ is in one of the $U_{\sigma_i, F_{\sigma_i}}$ and by our previous remarks this implies that $\tau(\mathfrak{P}_1 \cap F_{\sigma_i}) \neq \mathfrak{P}_2 \cap F_{\sigma_i}$, and thus

$$\tau(\mathfrak{P}_1 \cap F) \neq \mathfrak{P}_2 \cap F.$$

This applies for all $\tau \in \mathrm{Gal}(F/K)$, and $\mathfrak{P}_1 \cap F$ and $\mathfrak{P}_2 \cap F$ are primes in $\mathcal{O}_F$ (we will frequently just call these "primes in F"), so this contradicts the finite version of the theorem. □

So for any prime $\mathfrak{p}$ in K, we have a transitive action of $G = \mathrm{Gal}(L/K)$ on the primes of L lying over $\mathfrak{p}$. The next logical step is to look at the stabilizers of this action.

Definition 3.8. Let $\mathfrak{P}|\mathfrak{p}$. Then we define the **Decomposition group**

$$D(\mathfrak{P}|\mathfrak{p}) := \{\sigma \in G : \sigma(\mathfrak{P}) = \mathfrak{P}\}.$$

Lemma 3.9

For any $\mathfrak{P}|\mathfrak{p}$, the decomposition group $D(\mathfrak{P}|\mathfrak{p})$ is closed in G.

Proof. As usual, we just look at the finite Galois subextensions. The important claim is

$$D_{L/K}(\mathfrak{P}|\mathfrak{p}) = \bigcap_{F} \{\sigma \in \mathrm{Gal}(L/K) : \sigma(\mathfrak{P} \cap F) = \mathfrak{P} \cap F)\}$$

where the intersection is taken over all finite Galois intermediate extensions $K \subset F \subset L$ (where F/K is finite Galois). The inclusion of the left hand side into the right hand side is simply because any $\sigma \in \mathrm{Gal}(L/K)$ restricts to an element of $\mathrm{Gal}(F/K)$, and in particular it takes elements of F to elements of F and elements of $L - F$ to $L - F$. So if σ takes $\mathfrak{P}$ to $\mathfrak{P}$, it must take $\mathfrak{P} \cap F$ to itself as well.

For the inclusion of the right hand side into the left hand side, as usual the theme is that every $\alpha \in L$ is contained in one of the F's that the intersection is being taken over. Using this, the fact that $\sigma(\mathfrak{P} \cap F) = \mathfrak{P} \cap F$ for all such F immediately implies both inclusions of $\sigma(\mathfrak{P}) = \mathfrak{P}$.

Anyway, now we have written $D_{L/K}(\mathfrak{P}|\mathfrak{p})$ as an intersection of subsets which are clearly closed as a result of the fact that the restriction map $\mathrm{Gal}(L/K) \to \mathrm{Gal}(F/K)$ is continuous (a fact which we used in the proof of Lemma 2.9, where we observed that it is basically what the definition of the Krull topology says): the set $\{\sigma \in \mathrm{Gal}(L/K) : \sigma(\mathfrak{P} \cap F) = \mathfrak{P} \cap F)\}$ is the preimage under the restriction map of the closed set $D_{F/K}(\mathfrak{P}|\mathfrak{p}) \subset \mathrm{Gal}(F/K)$. □

By the theory of group actions, the primes lying over $\mathfrak{p}$ are in natural bijection with the coset space $G/D(\mathfrak{P}|\mathfrak{p})$ via $gD(\mathfrak{P}|\mathfrak{p}) \mapsto g(\mathfrak{P})$. Moreover, the decomposition group is equipped with a natural map given by projection

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

In math 129, we showed that if the residue fields are finite and L/K is finite, then this projection is surjective. In fact, it is true even without either of those conditions. We prove this right now, first in the case where L/K is finite and then in general.

Proposition 3.10

Let L/K be a finite Galois extension with group G, $\mathcal{O}_K$ a Dedekind domain with field of fractions K, and $\mathcal{O}_L$ its integral closure in L. For any prime $\mathfrak{P}$ in $\mathcal{O}_L$ lying over the prime $\mathfrak{p}$ in $\mathcal{O}_K$, the extension $\kappa(\mathfrak{P})/\kappa(\mathfrak{p}) = (\mathcal{O}_L/\mathfrak{P})/(\mathcal{O}_K/\mathfrak{p})$ is normal and finite, and the natural map $D(\mathfrak{P}|\mathfrak{p}) \to \mathrm{Aut}(\kappa(\mathfrak{P})/\kappa(\mathfrak{p}))$ is surjective.

Proof. First, since L/K is finite, we know by (the easy part of) Lemma 3.5 that $\mathcal{O}_L$ is a Dedekind domain, so $\mathfrak{P}$ is maximal and thus $\kappa(\mathfrak{P})$ is a field. As remarked in the proof of Lemma 3.5, we know $\mathcal{O}_L$ is a finitely-generated module over $\mathcal{O}_K$, so $\kappa(\mathfrak{P})/\kappa(\mathfrak{p})$ is finite and thus algebraic. To show it is normal, consider an arbitrary element $\overline{\alpha} \in \kappa(\mathfrak{P})$ with $\alpha \in \mathcal{O}_L$. Since L/K is finite, we know α satisfies some monic irreducible polynomial $f(X) \in \mathcal{O}_K[X]$. Since L/K is Galois and f has the root $\alpha \in L$, actually f splits completely into linear factors over L. All of its roots are in $\mathcal{O}_L$, so in fact we can just reduce each linear factor mod $\mathfrak{P}$ to see that $\overline{f}$ splits completely into linear factors over $\kappa(\mathfrak{P})$. So the minimal polynomial of $\overline{\alpha}$, which divides $\overline{f}$, also splits completely. As a result, the extension $\kappa(\mathfrak{P})/\kappa(\mathfrak{p})$ is normal. So far, this is the same as the usual proof. The problem is that now that we haven't assumed that $\kappa(\mathfrak{p})$ is perfect (or finite as the usual assumption goes), we don't know that $\kappa(\mathfrak{P})/\kappa(\mathfrak{p})$ is separable. Luckily, this is fixed with a little bit of extra maneuvering.

Let F be the separable closure of $\kappa(\mathfrak{p})$ inside $\kappa(\mathfrak{P})$. Then every element of $\mathrm{Gal}(F/\kappa(\mathfrak{p}))$ has exactly one extension[21] to $\mathrm{Aut}(\kappa(\mathfrak{P})/\kappa(\mathfrak{p}))$, which means that it suffices to show that every element of $\mathrm{Aut}(\kappa(\mathfrak{P})/\kappa(\mathfrak{p}))$ with a particular restriction to $\mathrm{Gal}(F/\kappa(\mathfrak{p}))$ lifts

[21] Since fields of characteristic zero are perfect, there is nothing to check in that case, and we can assume we are in characteristic $p > 0$. Let $a \in \kappa(\mathfrak{P}) - F$ so that the minimal polynomial $f(X) \in \kappa(\mathfrak{p})[X]$ for a splits completely over $\kappa(\mathfrak{P})$ but is not separable (because the roots are not distinct). So f shares a root with its derivative, which means that $f'(X) = 0$. This means that $f = g(X^p)$ for some irreducible polynomial $g \in \kappa(\mathfrak{p})[X]$ of smaller degree, and thus we can repeat this finitely many times to get $f = g(X^{p^n})$ for some n, where g is separable. Therefore, a^{p^n} is separable over $\kappa(\mathfrak{p})$ so $a^{p^n} \in F$. This shows that every $a \in \kappa(\mathfrak{P}) - F$ is the p^n-th root of some element of F, so it satisfies the

to an element of $D(\mathfrak{P}|\mathfrak{p})$. Then the usual proof works[22]. In particular, $F/\kappa(\mathfrak{p})$ is finite and separable, so $F = (\kappa(\mathfrak{p}))(a)$ for some $a \in F$, and so it suffices to show that the roots of the minimal polynomial for a over $\kappa(\mathfrak{p})$ are all of the form $\overline{\sigma}(a)$ where $\overline{\sigma}$ denotes the projection to the residue field of some $\sigma \in D(\mathfrak{P}|\mathfrak{p})$. If we can show that the polynomial

$$f(X) = \prod_{\sigma \in D(\mathfrak{P}|\mathfrak{p})} (X - \overline{\sigma}(a)) \in F[X]$$

has coefficients in $\kappa(\mathfrak{p})$, then we are done because that would imply that f divides a power of the minimal polynomial for a over $\kappa(\mathfrak{p})$ (since all the roots of f are also roots of this minimal polynomial), so since the minimal polynomial is irreducible, f is itself a power of the minimal polynomial which means that indeed the Galois conjugates of a over $\kappa(\mathfrak{p})$ are all of the form $\overline{\sigma}(a)$ as desired. To prove that $f(X)$ has coefficients in $\kappa(\mathfrak{p})$, let $\alpha \in \mathcal{O}_L$ be such that $\overline{\alpha} = a$, and $\alpha \in \mathfrak{P}'$ for all primes $\mathfrak{P}'$ of $\mathcal{O}_L$ other than $\mathfrak{P}$ (such an α exists by the Chinese remainder theorem, since from Lemma 3.5 we know $\mathcal{O}_L$ is a Dedekind domain and there are therefore finitely many such primes). Then $\sigma(\alpha) \in \mathfrak{P}$ for all $\sigma \in G - D(\mathfrak{P}|\mathfrak{p})$, so the polynomial

$$\prod_{\sigma \in G} (X - \sigma(\alpha)) \in \mathcal{O}_K[X]$$

reduces mod $\mathfrak{p}$ to

$$X^{|G - D(\mathfrak{P}|\mathfrak{p})|} f(X)$$

which shows that indeed f has coefficients in $\kappa(\mathfrak{p})$ as desired. □

We may deduce the infinite case directly from the finite case, with a little input from the topological group structure of $\text{Gal}(L/K)$ and one small lemma.

Lemma 3.11

Let L/K be a Galois extension, with an arbitrary intermediate extension $K \subset F \subset L$ such that F/K is Galois. Then the restriction map

$$D_{L/K}(\mathfrak{P}|\mathfrak{p}) \to D_{F/K}(\mathfrak{P} \cap F|\mathfrak{p})$$

is surjective.

polynomial $X^{p^n} - \alpha$ for some $\alpha \in F$ with $\alpha = a^{p^n}$. This polynomial is equal to $(X-a)^{p^n}$ since we are in characteristic p, and as a result the minimal polynomial for a over F has a single root, namely a (but that root is not simple). It follows that any element of $\text{Aut}(\kappa(\mathfrak{P})/\kappa(\mathfrak{p}))$ that fixes F must also fix a, which means the kernel of the restriction map $\text{Aut}(\kappa(\mathfrak{P})/\kappa(\mathfrak{p})) \to \text{Gal}(F/\kappa(\mathfrak{p}))$ is trivial and thus (by the surjectivity, which we already know in general) we have the desired isomorphism

[22] In math 129, we did this by first moving K up to the fixed field of $D(\mathfrak{P}|\mathfrak{p})$ so that we could assume that $G = D(\mathfrak{P}|\mathfrak{p})$ (see [6, Ch. 6, §6.2, Proposition 2]); In particular, $\text{Aut}(\kappa(\mathfrak{P} \cap L^{D(\mathfrak{P}|\mathfrak{p})})/\kappa(\mathfrak{p})) = 1$, which means it suffices to lift elements of $\text{Aut}(\kappa(\mathfrak{P})/\kappa(\mathfrak{P} \cap L^{D(\mathfrak{P}|\mathfrak{p})}))$ to elements of $\text{Gal}(L/L^{D^{\mathfrak{P}|\mathfrak{p}}}) = D(\mathfrak{P}|\mathfrak{p})$. Even after this, the argument is essentially the same but we avoid the step where we need to use the Chinese remainder theorem. Also, it takes more work to do this when L/K is infinite: we need to show that $\kappa(\mathfrak{P} \cap L^D) = \kappa(\mathfrak{p})$, which is evident from degree considerations in the finite case given $|D| = ef$ and the multiplicativity in towers of ramification and inertial degrees. To do this, we show that any finite subextension is trivial, which can be done just by taking a subextension $K \subset F \subset L^{D(\mathfrak{P}|\mathfrak{p})}$ such that F/K is finite. Let F' be the Galois closure of F/K. The fact that F is fixed by $D(\mathfrak{P}|\mathfrak{p})$ is equivalent to F being fixed by $D(\mathfrak{P}'|\mathfrak{p})$ since the first group surjects onto the second by restriction (see Lemma 3.11). Now we are back in the finite case: F is a finite extension of K which is fixed by the decomposition group of $\mathfrak{P} \cap F'$ where F' is a finite Galois extension of K containing F.

Proof. We've already noticed in for example Lemma 3.9 that the restriction homomorphism really does map to the decomposition group downstairs. To check it is surjective, let $\sigma_F \in D_{F/K}(\mathfrak{P} \cap F|\mathfrak{p})$. Since F/K is Galois, we know that there is some $\sigma_L \in \mathrm{Gal}(L/K)$ that restricts to σ_F on F. We just need to modify it so that it still restricts to σ_F but lands in $D_{L/K}(\mathfrak{P}|\mathfrak{p})$.

We know that $\sigma_L(\mathfrak{P})$ is a prime in $\mathcal{O}_L$. By Theorem 3.7, there is a $\tau \in \mathrm{Gal}(L/F)$ such that $\tau(\sigma_L(\mathfrak{P})) = \mathfrak{P}$, so actually $\tau \circ \sigma_L \in D(\mathfrak{P}|\mathfrak{p})$ and restricts to σ_F on F since τ fixes F, as desired. □

NB: we could have done this just using the finite case of Theorem 3.7 and Zorn's lemma. Or better yet, we could avoid use of transitivity of the action completely, just by (in the finite case) using the fact that the restriction map $D(\mathfrak{P}|\mathfrak{p}) \to D(\mathfrak{P} \cap F|\mathfrak{p})$ has kernel

$$D(\mathfrak{P}|\mathfrak{p}) \cap \mathrm{Gal}(L/F) = D(\mathfrak{P}|\mathfrak{P} \cap F)$$

so this induces an injection

$$\frac{D(\mathfrak{P}|\mathfrak{p})}{D(\mathfrak{P}|\mathfrak{P} \cap F)} \to D(\mathfrak{P} \cap F|\mathfrak{p}).$$

By the multiplicativity of ramification and inertial degrees, both sides have cardinality $e(\mathfrak{P} \cap F|\mathfrak{p})f(\mathfrak{P} \cap F|\mathfrak{p})$. Once we have the finite case, the Zorn's lemma argument is the same as usual. Consider the poset P consisting of pairs (F, σ) where $F \in \mathcal{F}$ and $\sigma \in D(\mathfrak{P} \cap F|\mathfrak{p})$, where the ordering is just $(F_1, \sigma_1) \le (F_2, \sigma_2)$ if $F_1 \subset F_2$ and σ_2 extends σ_1. Any chain has an upper bound (just take the union of all the fields and use the σ's to define the element of the decomposition group where they are defined, which is valid by the proof of Lemma 3.9). So by Zorn's lemma, it has a maximal element, which must have $F = L$ by the finite case of the surjectivity (if $F \neq L$ then we can take an intermediate extension $F \subset F' \subset L$ which is finite over F and Galois over K; the finite case tells us we can extend the $\sigma \in D(\mathfrak{P} \cap F|\mathfrak{p})$ to $D(\mathfrak{P} \cap F'|\mathfrak{p})$).

Theorem 3.12

Let L/K be a (possibly infinite) Galois extension, $\mathcal{O}_K$ a Dedekind domain with field of fractions K, and $\mathcal{O}_L$ its integral closure in L. For any prime $\mathfrak{P}$ in $\mathcal{O}_L$ and $\mathfrak{p}$ in $\mathcal{O}_K$, the extension $\kappa(\mathfrak{P})/\kappa(\mathfrak{p}) = (\mathcal{O}_L/\mathfrak{P})/(\mathcal{O}_K/\mathfrak{p})$ is normal, and the natural map

$$D(\mathfrak{P}|\mathfrak{p}) \to \mathrm{Aut}(\kappa(\mathfrak{P})/\kappa(\mathfrak{p}))$$

is surjective.

Proof. First, note that the proof of the fact that the primes of $\mathcal{O}_L$ are maximal from Lemma 3.5 did not depend on either the finiteness or separability of L/K, and therefore in this context we really can say that $\kappa(\mathfrak{P})$ is a field. Also, the extension $\kappa(\mathfrak{P})/\kappa\mathfrak{p}$ is normal by same argument as in the finite case Proposition 3.10, basically from the fact that L/K is normal. First of all, $\kappa(\mathfrak{P})/\kappa(\mathfrak{p})$ is algebraic because any $\alpha \in \kappa(\mathfrak{P})$ is the reduction mod $\mathfrak{P}$ of some $\tilde{\alpha} \in \mathcal{O}_L$, which has a minimal polynomial $f \in \mathcal{O}_K[X]$. So $\overline{f}(\alpha)$ is 0 in $\kappa(\mathfrak{p})$ as well where $\overline{f}$ is the reduction mod $\mathfrak{p}$ of f, which means that α is algebraic over $\kappa(\mathfrak{p})$. Moreover, the minimal polynomial of α over $\kappa(\mathfrak{p})$ divides $\overline{f}$, and $\overline{f}$ splits into linear factors over $\kappa(\mathfrak{P})$ since f does over L (all of the roots of f are in $\mathcal{O}_L$ since f is monic with coefficients in $\mathcal{O}_K$). This means the minimal polynomial of x splits into linear factors over $\kappa(\mathfrak{p})$, and therefore the residue field extension is normal[23] as claimed.

For the surjectivity, we use the finite case Proposition 3.10. Recall from the proof of Lemma 3.9 that

$$D(\mathfrak{P}|\mathfrak{p}) = \bigcap_{F\in\mathcal{F}} \{\sigma \in \mathrm{Gal}(L/K) : \sigma|_F \in D_{F/K}(\mathfrak{P}\cap F|\mathfrak{p})\}$$

where $\mathcal{F}$ is the collection of all intermediate fields $K \subset F \subset L$ such that F/K is finite and Galois. So under the isomorphism from Theorem 1.12, the subgroup $D_{L/K}(\mathfrak{P}|\mathfrak{p}) \subset \mathrm{Gal}(L/K)$ is mapped isomorphically to the subgroup of $\varprojlim_{F\in\mathcal{F}} \mathrm{Gal}(F/K)$ given by all $(\sigma_F)_{F\in\mathcal{F}}$ such that $\sigma_F \in D_{F/K}(\mathfrak{P}\cap F|\mathfrak{p})$ for each F. So we have an isomorphism

$$D_{L/K}(\mathfrak{P}|\mathfrak{p}) \to \varprojlim_{F\in\mathcal{F}} D_{F/K}(\mathfrak{P}\cap F|\mathfrak{p}),$$

where the inverse system is given by the restriction maps $D_{F_1/K}(\mathfrak{P}\cap F_1|\mathfrak{p}) \to D_{F_2/K}(\mathfrak{P}\cap F_2|\mathfrak{p})$ whenever $F_2 \subset F_1$ (we saw in the proof of Lemma 3.9 that the restriction to $\mathrm{Gal}(F_2/K)$ lands in the decomposition group), given by sending

$$\sigma \mapsto (\sigma|_F)_{F\in\mathcal{F}}.$$

On the other side of things, let $\tau \in \mathrm{Aut}(\kappa(\mathfrak{P})/\kappa(\mathfrak{p}))$. By the normality of $\kappa(\mathfrak{P}\cap F)/\kappa(\mathfrak{p})$ (proved in the same way as before but replacing $\mathcal{O}_L$ with $\mathcal{O}_F$), we know that τ restricts to a valid element of $\mathrm{Aut}(\kappa(\mathfrak{P}\cap F)/\kappa(\mathfrak{p}))$ for each finite Galois intermediate extension F. Moreover, an element $\alpha \in \kappa(\mathfrak{P})$ is the reduction mod $\mathfrak{P}$ of some $\tilde{\alpha} \in \mathcal{O}_L$. As usual, we know $\tilde{\alpha}$ is contained in some $F \in \mathcal{F}$ (for the usual reasons: take the Galois closure of $K(\alpha)/K$), so in fact $\tilde{\alpha} \in \mathcal{O}_L \cap F = \mathcal{O}_F$. This means $\alpha \in \kappa(\mathfrak{P}\cap F)$. So by Theorem 1.12[24],

$$\mathrm{Aut}(\kappa(\mathfrak{P})/\kappa(\mathfrak{p})) \cong \varprojlim_{F\in\mathcal{F}} \mathrm{Aut}(\kappa(\mathfrak{P}\cap F)/\kappa(\mathfrak{p}))$$

given by

$$\sigma \mapsto (\sigma|_{\kappa(\mathfrak{P}\cap F)})_{F\in\mathcal{F}}$$

where the inverse system is the usual one given by the restriction homomorphisms $\mathrm{Aut}(\kappa(\mathfrak{P}\cap F_1)/\kappa(\mathfrak{P})) \to \mathrm{Aut}(\kappa(\mathfrak{P}\cap F_2)/\kappa(\mathfrak{P}))$ (well-defined by the remarks earlier in this paragraph) whenever $F_2 \subset F_1$. Under these isomorphisms, the projection

$$D_{L/K}(\mathfrak{P}|\mathfrak{p}) \to \mathrm{Aut}(\kappa(\mathfrak{P})/\kappa(\mathfrak{p}))$$

becomes the map

$$\varphi : \varprojlim_{F\in\mathcal{F}} D_{F/K}(\mathfrak{P}\cap F|\mathfrak{p}) \to \varprojlim_{F\in\mathcal{F}} \mathrm{Aut}(\kappa(\mathfrak{P}\cap F)/\kappa(\mathfrak{p}))$$

given in the F-coordinate by the projection $D_{F/K}(\mathfrak{P}\cap F|\mathfrak{p}) \to \mathrm{Aut}(\kappa(\mathfrak{P}\cap F)/\kappa(\mathfrak{p}))$, which is surjective by the finite case Proposition 3.10. Moreover, φ is continuous[25]. We want

[23] NB: the residue field extension is NOT always separable. One of the advantages of the additional assumption that $\kappa(\mathfrak{p})$ is finite (which is true when K is a function field over a finite field or a number field) is that it implies that the residue field extension is separable and thus Galois.

[24] At least by the modified version we've used already, which has the exact same proof except we replace the word "Galois" with "normal," and $\mathcal{L}$ is only required to contain enough finite normal intermediate extensions to hit every point in the big field $\kappa(\mathfrak{P})$.

[25] Indeed, the topology on $\mathrm{Aut}(\kappa(\mathfrak{P})/\kappa(\mathfrak{p}))$ has a basis running over all $F \in \mathcal{F}$ consisting of the preimages under restriction to $\mathrm{Aut}(\kappa(\mathfrak{P}\cap F)/\kappa(\mathfrak{p}))$ of any element τ. The preimage under φ of this restriction is the set of all $\sigma \in D(\mathfrak{P}|\mathfrak{p})$ that restrict to a preimage of τ in $D_{F/K}(\mathfrak{P}\cap F|\mathfrak{p})$, and is therefore an open set in the topology on $D(\mathfrak{P}|\mathfrak{p})$ since it is a (finite) union of basis elements. NB: it's easy to check that as long as our set of finite Galois subextensions $\mathcal{F}$ hits every element of L, the topology from $\varprojlim_{F\in\mathcal{F}} \mathrm{Gal}(F/K)$ is always the same (the proof is the same as the remarks after Example 2.3, where we originally made this observation), so all the isomorphisms we have of decomposition groups or automorphism groups of residue field extensions with inverse limits really are topological isomorphisms.

to check that for any $(\tau_F)_{F\in\mathcal{F}} \in \varprojlim_{F\in\mathcal{F}} \operatorname{Aut}(\kappa(\mathfrak{P}\cap F)/\kappa(\mathfrak{p}))$, $\varphi^{-1}((\tau_F)_{F\in\mathcal{F}})$ is nonempty. Of course, $(\tau_F)_{F\in\mathcal{F}}$ corresponds to the element $\tau \in \operatorname{Aut}(\kappa(\mathfrak{P})/\kappa(\mathfrak{p}))$ that restricts to τ_F on each $F \in \mathfrak{F}$. In particular,

$$\{\tau\} = \bigcap_{F\in\mathcal{F}} \{g \in \operatorname{Aut}(\kappa(\mathfrak{P})/\kappa(\mathfrak{p})) : g|_F = \tau_F\},$$

and therefore

$$\varphi^{-1}(\tau) = \bigcap_{F\in\mathcal{F}} \left\{\sigma \in D_{L/K}(\mathfrak{P}|\mathfrak{p}) : \overline{\sigma|_F} = \tau_F\right\}.$$

Each of the sets here is nonempty by Proposition 3.10 and Lemma 3.11, and closed in $D_{L/K}(\mathfrak{P}|\mathfrak{p})$ by definition of the Krull topology. So $\varphi^{-1}(\tau)$ is the intersection of a collection of nonempty closed subsets of $D(\mathfrak{P}|\mathfrak{p})$. Since $D(\mathfrak{P}|\mathfrak{p})$ is a closed subset of a compact Hausdorff space $\operatorname{Gal}(L/K)$ by Proposition 2.12 and Lemma 3.9, it follows that $\varphi^{-1}(\tau)$ is an intersection of nonempty closed compact sets, and is therefore nonempty, as desired. □

Remark 3.13. In most cases, we will assume that the residue field $\mathcal{O}_K/\mathfrak{p}$ is finite or at least perfect (for instance this is the case when $\mathcal{O}_K$ is the ring of integers in a number field K, or when $\mathcal{O}_K = \mathbf{F}[X]$ for a finite field $\mathbf{F}$). For now, we do not assume this unless we say so explicitly.

Armed with this reduction map, it's natural to consider

Definition 3.14. The kernel of the map $D(\mathfrak{P}|\mathfrak{p}) \to \operatorname{Aut}(\kappa(\mathfrak{P})/\kappa(\mathfrak{p}))$ is called the **inertia subgroup** $I(\mathfrak{P}|\mathfrak{p})$.

By Theorem 3.12, if L/K is a Galois extension with group G, and $\mathfrak{P}$ a prime of $\mathcal{O}_L$ lying over a prime $\mathfrak{p}$ of $\mathcal{O}_K$, then

$$D(\mathfrak{P}|\mathfrak{p})/I(\mathfrak{P}|\mathfrak{p}) \cong \operatorname{Aut}(\kappa(\mathfrak{P})/\kappa(\mathfrak{p})).$$

In fact, this is also an isomorphism of topological spaces, where the right hand side is given the Krull topology (NB the definition of the Krull topology does not depend on separability). This is because the projection $D(\mathfrak{P}|\mathfrak{p}) \to \operatorname{Aut}(\kappa(\mathfrak{P})/\kappa(\mathfrak{p}))$ is continuous (which we remarked in the proof of Theorem 3.12) and surjective (part of the statement of Theorem 3.12). It is also an open map, since it takes the basis element

$$\{\sigma \in D(\mathfrak{P}|\mathfrak{p}) : \sigma|_F = \sigma_F\}$$

to the basis element

$$\{\tau \in \operatorname{Aut}(\kappa(\mathfrak{P})|\kappa(\mathfrak{p})) : \tau|_{\kappa(\mathfrak{P}\cap F)} = \varphi(\sigma_F)\}.$$

As a result, the induced map on the quotient is bijective, continuous, and open, and is therefore an isomorphism of topological groups.

Lemma 3.15
$I(\mathfrak{P}|\mathfrak{p})$ is a closed subgroup of G.

Proof. The argument is almost identical to that of Lemma 3.9. In particular, any $a \in \kappa(\mathfrak{P})$ lifts to an $\alpha \in \mathcal{O}_L$, which is contained in $\mathcal{O}_L \cap F = \mathcal{O}_F$ for some intermediate field $K \subset F \subset L$ where F/K is finite and Galois. As a result, $\alpha \in \kappa(\mathfrak{P}\cap F)$ which is a finite normal extension of $\kappa(\mathfrak{p})$ (by the first part of Proposition 3.10). So the intermediate extensions $\kappa(\mathfrak{P}\cap F)/\kappa(\mathfrak{p})$ for F/K finite, Galois and contained in L are enough so that

$$I(\mathfrak{P}|\mathfrak{p}) = \bigcap_F \{\sigma \in G : \sigma|_F \in I_{F/K}(\mathfrak{P}\cap F|\mathfrak{p})\}$$

since any element of the right hand side is in $D(\mathfrak{P}|\mathfrak{p})$ (by the intersection formula from Lemma 3.9), and therefore induces an element of $\mathrm{Aut}(\kappa(\mathfrak{P})/\kappa(\mathfrak{p}))$ that restricts to the identity on $\kappa(\mathfrak{P}\cap F)$ for each F. This implies it is the identity on all of $\kappa(\mathfrak{P})$ since all $\alpha \in \kappa(\mathfrak{P})$ are in $\kappa(\mathfrak{P}\cap F)$ for some F, which proves the inclusion of the right hand side in the left. The inclusion of the left hand side into the right is obvious. So $I(\mathfrak{P}|\mathfrak{p})$ is an intersection of closed sets (which are closed by the same argument we used in Lemma 3.9) and is therefore closed as desired. □

Remark 3.16. Alternatively, Lemma 3.15 is true just because the inertia group is the kernel of the continuous homomorphism from the decomposition group to the automorphism group of the residue field extension, which is Hausdorff under the Krull topology.

Since G acts transitively on the primes $\mathfrak{P}|\mathfrak{p}$ (Theorem 3.7), we can get all the decomposition and inertia groups by taking conjugates of just one, as

$$D(\sigma(\mathfrak{P})|\mathfrak{p}) = \sigma D(\mathfrak{P}|\mathfrak{p})\sigma^{-1}$$

and

$$I(\sigma(\mathfrak{P})|\mathfrak{p}) = \sigma I(\mathfrak{P}|\mathfrak{p})\sigma^{-1}$$

(the first one is trivial, and the second one follows from the fact that $\kappa(\mathfrak{P}) \cong \kappa(\sigma\mathfrak{P})$ via $\overline{\alpha} \mapsto \overline{\sigma(\alpha)}$). So if G is abelian (or if these subgroups are just normal), then the decomposition and inertia groups only depend on the prime downstairs. Also, $I(\mathfrak{P}|\mathfrak{p})$ is normal in $D(\mathfrak{P}|\mathfrak{p})$ since for any $\sigma \in D(\mathfrak{P}|\mathfrak{p})$ we have

$$\sigma I(\mathfrak{P}|\mathfrak{p})\sigma^{-1} = I(\sigma\mathfrak{P}|\mathfrak{p}) = I(\mathfrak{P}|\mathfrak{p}).$$

Definition 3.17. If $I(\mathfrak{P}|\mathfrak{p})$ is trivial, then we say that $\mathfrak{p}$ is **unramified** in L.

So if L/K is a Galois extension with prime $\mathfrak{P}$ in $\mathcal{O}_L$ lying over $\mathfrak{p}$ in $\mathcal{O}_K$ with $\mathfrak{p}$ unramified, then by Theorem 3.12,

$$D(\mathfrak{P}|\mathfrak{p}) \cong \mathrm{Aut}(\kappa(\mathfrak{P})/\kappa(\mathfrak{p})).$$

If we further stipulate that $\kappa(\mathfrak{p})$ is finite and L/K is finite, then since by the first part of Proposition 3.10, $\kappa(\mathfrak{P})/\kappa(\mathfrak{p})$ is finite, the group $\mathrm{Aut}(\kappa(\mathfrak{P})/\kappa(\mathfrak{p}))$ is the Galois group of an extension of finite fields and is therefore cyclic, with canonical choice of generator given by $x \mapsto x^{|\kappa(\mathfrak{p})|}$. So we may define we can pull the Frobenius map back to $D(\mathfrak{P}|\mathfrak{p})$ through the isomorphism to associate a generator of a cyclic group

$$\mathrm{Frob}_{\mathfrak{P}} \in D(\mathfrak{P}|\mathfrak{p})$$

to each prime $\mathfrak{P}|\mathfrak{p}$. In particular, our isomorphism $D(\mathfrak{P}|\mathfrak{p}) \cong \mathrm{Gal}(\kappa(\mathfrak{P})/\kappa(\mathfrak{p}))$ tells us that the following definition is well-defined:

Definition 3.18. Let L/K be a finite Galois extension, $\mathcal{O}_K$ a Dedekind domain with field of fractions K with integral closure $\mathcal{O}_L$ in L. Let $\mathfrak{P}$ be a prime in $\mathcal{O}_L$ lying over $\mathfrak{p}$ in $\mathcal{O}_K$, and suppose that $\kappa(\mathfrak{p})$ is finite and $\mathfrak{p}$ is unramified in L. The **Frobenius element** corresponding to $\mathfrak{P}|\mathfrak{p}$ is the unique element $\mathrm{Frob}_{\mathfrak{P}} \in \mathrm{Gal}(L/K)$ satisfying

$$\mathrm{Frob}_{\mathfrak{P}}(x) \equiv x^{|\kappa(\mathfrak{p})|} \mod \mathfrak{P}$$

for all $x \in L$.

Again, because any $\sigma \in \mathrm{Gal}(L/K)$ provides an isomorphism $\kappa(\mathfrak{P}) \to \kappa(\sigma(\mathfrak{P}))$, we have

$$\mathrm{Frob}_{\sigma\mathfrak{P}} = \sigma \mathrm{Frob}_{\mathfrak{P}} \sigma^{-1},$$

so for a prime $\mathfrak{p}$ of $\mathcal{O}_K$, the Frobenius defines a conjugacy class of $\mathrm{Gal}(L/K)$. When $\mathrm{Gal}(L/K)$ is abelian this means the Frobenius element depends only on $\mathfrak{p}$, and we can extend linearly to get a group homomorphism $I(K) \to \mathrm{Gal}(L/K)$, where $I(K)$ is the group of fractional ideals of $\mathcal{O}_K$. This is what Artin's global reciprocity law involves.

Remark 3.19. Suppose we drop the assumption that L/K is finite. Now, $\kappa(\mathfrak{P})/\kappa(\mathfrak{p})$ is a possibly infinite (separable) algebraic extension of a finite field, so its Galois group is a quotient of $\widehat{\mathbf{Z}}$ by a closed subgroup (see Example 1.13, Example 2.1, and Lemma 2.13 and the discussion following it). Now there is still a canonical "Frobenius" element of $\mathrm{Gal}(\kappa(\mathfrak{P})/\kappa(\mathfrak{p}))$, namely the one given by $x \mapsto x^{|\kappa(\mathfrak{p})|}$. We recall that the subgroup generated by this Frobenius element (which corresponds to $\mathbf{Z} \subset \widehat{\mathbf{Z}}$) is dense in $\widehat{\mathbf{Z}} = \mathrm{Gal}(\overline{\kappa(\mathfrak{p})}/\kappa(\mathfrak{p}))$, which one can check directly (by the definition of $\widehat{\mathbf{Z}}$ or by the fact that the finite extensions of $\kappa(\mathfrak{p})$ are cyclic with Galois group generated by the Frobenius). So when we take quotients and pull back under the isomorphism with $D(\mathfrak{P}|\mathfrak{p})$, we still obtain a canonical choice of Frobenius element in the decomposition group, but it does not generate the whole thing; instead it generates a dense subgroup.

Since by now we've already assumed L/K is finite and $\kappa(\mathfrak{p})$ is finite and $\mathfrak{p}$ unramified in order to define the Frobenius element, we should make the usual remarks about what this has to do with the splitting of primes. Now, we only assume L/K is finite and Galois. Since G acts transitively on the primes lying over $\mathfrak{p}$, the residue field extensions $\kappa(\mathfrak{P})/\kappa(\mathfrak{p})$ are all isomorphic for the primes $\mathfrak{P}$ lying over $\mathfrak{p}$. They have the same degree, namely $f_{\mathfrak{p}}$ (the *inertial degree* of $\mathfrak{P}|\mathfrak{p}$), and again by the transitivity we have

$$\mathfrak{p}\mathcal{O}_L = \Big(\prod_i \mathfrak{P}_i\Big)^{e_{\mathfrak{p}}},$$

where the $\mathfrak{P}_i$ are the primes lying over $\mathfrak{p}$ and the positive integer $e_{\mathfrak{p}}$ is the *ramification index* of $\mathfrak{P}|\mathfrak{p}$. We've already remarked that since G acts transitively on the $\mathfrak{P}_i$, we know that the cosets of $D(\mathfrak{P}|\mathfrak{p})$ in G are in bijection with the $\mathfrak{P}_i$, so the number of prime factors is $[G : D(\mathfrak{P}_i|\mathfrak{p})]$ (any $\mathfrak{P}_i$ works). We showed in math 129[26] the "fundamental identity" $\sum e_i f_i = [L : K]$, where e_i and f_i are the ramification indices of the primes lying

[26] In math 129 we probably didn't prove this "fundamental identity" in the full generality, but it isn't hard to do. Recall that the three ingredients are (1) the Chinese remainder theorem, (2) the fact that $\mathcal{O}_L/\mathfrak{P}^e$ is e-dimensional as a $\kappa(\mathfrak{P})$-vector space, and (3) the fact that $\mathcal{O}_L/\mathfrak{p}\mathcal{O}_L$ is an $[L : K]$-dimensional vector space over $\kappa(\mathfrak{p})$. The Chinese remainder theorem still applies: since L/K is finite, Lemma 3.5 implies that $\mathcal{O}_L$ is Dedekind, so the ideals $\mathfrak{P}_i^{e_i}$ are pairwise comaximal. And we still have the exact sequence of $\kappa(\mathfrak{P})$-vector spaces $0 \to \mathcal{O}_L/\mathfrak{P}^{e-1} \to \mathcal{O}_L/\mathfrak{P}^e \to \mathcal{O}_L/\mathfrak{P} \to 0$, so that part of the argument is unchanged. The only part that requires modification is the proof that $\mathcal{O}_L/\mathfrak{p}\mathcal{O}_L$ is an $[L : K]$-dimensional vector space over $\kappa(\mathfrak{p})$. To do that, we can check directly that using the multiplicative set $S = \mathcal{O}_K - \mathfrak{p}$, the inclusion $\mathcal{O}_K \to S^{-1}\mathcal{O}_K$ induces an isomorphism of fields $\kappa(\mathfrak{p}) = \mathcal{O}_K/\mathfrak{p}\mathcal{O}_K \to S^{-1}\mathcal{O}_K/\mathfrak{p}S^{-1}\mathcal{O}_K$ which results in an isomorphism of $\kappa(\mathfrak{p})$-vector spaces $\mathcal{O}_L/\mathfrak{p}\mathcal{O}_L \to S^{-1}\mathcal{O}_L/\mathfrak{p}S^{-1}\mathcal{O}_L$, so since localization

over $\mathfrak{p}$ (this holds for arbitrary finite separable L/K). So in our case (L/K is Galois so definitely separable), we have

$$[L:K] = |G| = e_{\mathfrak{p}} f_{\mathfrak{p}} [G : D(\mathfrak{P}|\mathfrak{p})]$$

and as a result

$$D(\mathfrak{P}|\mathfrak{p}) = e_{\mathfrak{p}} f_{\mathfrak{p}}.$$

Since $D(\mathfrak{P}|\mathfrak{p})/I(\mathfrak{P}|\mathfrak{p}) \cong \operatorname{Aut}(\kappa(\mathfrak{P})/\kappa(\mathfrak{p}))$ by Proposition 3.10, if we assume further that $\kappa(\mathfrak{P})/\kappa(\mathfrak{p})$ is separable, then we know it is a finite Galois extension by Proposition 3.10, and thus

$$f_{\mathfrak{p}} = [\kappa(\mathfrak{P}) : \kappa(\mathfrak{p})] = |\operatorname{Gal}(\kappa(\mathfrak{P})/\kappa(\mathfrak{p}))| = [D(\mathfrak{P}|\mathfrak{p}) : I(\mathfrak{P}|\mathfrak{p})] = \frac{e_{\mathfrak{p}} f_{\mathfrak{p}}}{I(\mathfrak{P}|\mathfrak{p})}$$

so the ramification index is $e_{\mathfrak{p}} = |I(\mathfrak{P}_i|\mathfrak{p})|$. This summarizes what the decomposition and inertia groups have to do with the splitting of primes in finite Galois extensions, when the residue field extension is separable. It also explains why we say that $\mathfrak{p}$ is unramified if its inertia groups are trivial, and the following definition:

Definition 3.20. Let L/K be a Galois extension, $\mathfrak{P}$ a prime in $\mathcal{O}_L$ lying over $\mathfrak{p}$ in $\mathcal{O}_K$. Then $\mathfrak{p}$ **splits completely** in L if $D(\mathfrak{P}|\mathfrak{p})$ is trivial.

Now we return to the case of a general Galois extension, first proving the analogue of Lemma 3.11 for the inertia group.

Lemma 3.21

Let L/K be a Galois extension, $\mathcal{O}_K$ a Dedekind domain with field of fractions K, and $\mathcal{O}_L$ its integral closure in L. For any prime $\mathfrak{P}$ in $\mathcal{O}_L$ lying over $\mathfrak{p}$ in $\mathcal{O}_K$ and intermediate extension $K \subset F \subset L$ such that F/K is Galois, the restriction map

$$I(\mathfrak{P}|\mathfrak{p}) \to I(\mathfrak{P} \cap F|\mathfrak{p})$$

is surjective.

Proof. We have a commutative diagram

preserves integral closures and Dedekind domains, and the localization of a Dedekind domain at a prime is a DVR (and in particular a PID), and $\mathfrak{p}S^{-1}\mathcal{O}_K$ is the prime in $S^{-1}\mathcal{O}_K$, it suffices to prove the result in the case where $\mathcal{O}_K$ is a PID. In that case, the argument is the same as in 129: since L/K is separable, as remarked in the proof of Lemma 3.5, we know that $\mathcal{O}_L$ is a finitely-generated $\mathcal{O}_K$-module, and since it is an integral domain and $\mathcal{O}_K$ is a PID, it is actually a free $\mathcal{O}_K$-module of rank $[L:K]$. Writing $\mathcal{O}_L = \bigoplus_{i=1}^{[L:K]} \omega_i \mathcal{O}_K$, we have $\mathfrak{p}\mathcal{O}_L = \bigoplus_{i=1}^{[L:K]} \omega_i \mathfrak{p}\mathcal{O}_K$, So $\mathcal{O}_L/\mathfrak{p}\mathcal{O}_L$ is a direct sum of $[L:K]$ copies of $\kappa(\mathfrak{p})$. Alternatively, one can avoid the local argument and apply Nakayama's lemma (see [5, I.8.2]).

$$\begin{tikzcd}
 & 1 \arrow[d] & 1 \arrow[d] & 1 \arrow[d] & \\
1 \arrow[r] & I_{L/F}(\mathfrak{P}|\mathfrak{P}\cap F) \arrow[r] \arrow[d] & I_{L/K}(\mathfrak{P}|\mathfrak{p}) \arrow[r] \arrow[d] & I_{F/K}(\mathfrak{P}\cap F|\mathfrak{p}) \arrow[d] & \\
1 \arrow[r] & D_{L/F}(\mathfrak{P}|\mathfrak{P}\cap F) \arrow[r] \arrow[d] & D_{L/K}(\mathfrak{P}|\mathfrak{p}) \arrow[r] \arrow[d] & D_{F/K}(\mathfrak{P}\cap F|\mathfrak{p}) \arrow[r] \arrow[d] & 1 \\
1 \arrow[r] & \mathrm{Aut}(\kappa(\mathfrak{P})/\kappa(\mathfrak{P}\cap F)) \arrow[r] \arrow[d] & \mathrm{Aut}(\kappa(\mathfrak{P})/\kappa(\mathfrak{p})) \arrow[r] \arrow[d] & \mathrm{Aut}(\kappa(\mathfrak{P}\cap F)/\kappa(\mathfrak{p})) \arrow[r] \arrow[d] & 1 \\
 & 1 & 1 & 1 &
\end{tikzcd}$$

in which the first row is exact because the kernel of the restriction map $I(\mathfrak{P}|\mathfrak{p}) \to I(\mathfrak{P}\cap F|\mathfrak{p})$ is

$$I(\mathfrak{P}|\mathfrak{p}) \cap \mathrm{Gal}(L/F) = I(\mathfrak{P}|\mathfrak{P}\cap F),$$

the second row is exact because the same is true with I replaced with D and because of Lemma 3.11, the third row is exact because $\kappa(\mathfrak{P})/\kappa(\mathfrak{p})$ is normal and so is $\kappa(\mathfrak{P}\cap F)/\kappa(\mathfrak{p})$ by the first part of Theorem 3.12, and the columns are all exact by the second part of Theorem 3.12. Moreover, the inertia subgroups are all normal subgroups of the decomposition groups, so we can apply the snake lemma where we think of the inertia groups as kernels of the vertical maps and the trivial groups on the bottom as cokernels (beware, the snake lemma does not hold in general for groups; you need to check that the appropriate subgroups are normal). Thus we obtain a snake map which provides an exact sequence

$$1 \to I_{L/F}(\mathfrak{P}|\mathfrak{P}\cap F) \to I_{L/K}(\mathfrak{P}|\mathfrak{p}) \to I_{F/K}(\mathfrak{P}\cap F|\mathfrak{p}) \to 1$$

thus proving the desired surjectivity. □

Now we collect everything we have learned about how inertia and decomposition groups behave under taking subextensions.

Theorem 3.22

Let L/K be a Galois extension, $\mathcal{O}_K$ a Dedekind domain with field of fractions K, and $\mathcal{O}_L$ its integral closure in L with a prime $\mathfrak{P}$ lying over a prime $\mathfrak{p}$ in $\mathcal{O}_K$. For any closed subgroup $H \subset \mathrm{Gal}(L/K)$, we have

$$D_{L/K}(\mathfrak{P}|\mathfrak{p}) \cap H = D_{L/L^H}(\mathfrak{P}|\mathfrak{P}\cap L^H)$$

and

$$I_{L/K}(\mathfrak{P}|\mathfrak{p}) \cap H = I_{L/L^H}(\mathfrak{P}|\mathfrak{P}\cap L^H).$$

If H is also a normal subgroup, then we have surjective restriction maps $D_{L/K}(\mathfrak{P}|\mathfrak{p}) \to D_{L^H/K}(\mathfrak{P}\cap L^H|\mathfrak{p})$ and $I_{L/K}(\mathfrak{P}|\mathfrak{p}) \to I_{L^H/K}(\mathfrak{P}\cap L^H|\mathfrak{p})$ which means that the restriction map of Galois groups $\mathrm{Gal}(L/K) \to \mathrm{Gal}(L^H/K)$ induces isomorphisms

$$D_{L/K}(\mathfrak{P}|\mathfrak{p})/D_{L/L^H}(\mathfrak{P}|\mathfrak{P}\cap L^H) \cong D_{L^H/K}(\mathfrak{P}\cap L^H|\mathfrak{p})$$

and

$$I_{L/K}(\mathfrak{P}|\mathfrak{p})/I_{L/L^H}(\mathfrak{P}|\mathfrak{P}\cap L^H) \cong I_{L^H/K}(\mathfrak{P}\cap L^H|\mathfrak{p}).$$

Proof. The Galois correspondence Theorem 2.6 implies that running over all closed $H \subset \operatorname{Gal}(L/K)$ is the same as running over all intermediate extensions $K \subset F \subset L$ and setting $H = \operatorname{Gal}(L/F)$, and running over all normal closed $H \subset \operatorname{Gal}(L/K)$ is the same as running all over intermediate extensions F with F/K Galois. The first part of the statement regarding the compatibility under taking intersections (i.e. moving the base field up) is something which we have already taken to be true a few times (for instance in the proof of Lemma 3.21 because it is pretty simple to check:

$$D_{L/K}(\mathfrak{P}|\mathfrak{p}) \cap \operatorname{Gal}(L/F) = D_{L/F}(\mathfrak{P}|\mathfrak{P} \cap F),$$

since both are defined to be the subset of $\operatorname{Gal}(L/F)$ consisting of all σ such that $\sigma(\mathfrak{P}) = \mathfrak{P}$, and

$$I_{L/K}(\mathfrak{P}|\mathfrak{p}) \cap \operatorname{Gal}(L/F) = I_{L/F}(\mathfrak{P}|\mathfrak{P} \cap F),$$

since both are defined to be the subset of $D_{L/F}(\mathfrak{P}|\mathfrak{P} \cap F)$ consisting of all σ which induce the trivial automorphism on $\kappa(\mathfrak{P})$. The remainder of the statement (the surjectivity of the restriction maps of inertia and decomposition groups) is just Lemma 3.11 and Lemma 3.21. □

So the theme is that the inertia and decomposition groups are compatible with taking subextensions.

Example 3.23

Consider the cyclotomic extension $\mathbf{Q}(\zeta_n)/\mathbf{Q}$, and suppose $n = p^k m$ where p is prime and $(m, p) = 1$. Since the Galois group is abelian, the decomposition groups and inertia groups only depend on the downstairs prime, so we call them $I(p)$ and $D(p)$. Much of this computation is review from math 129, but I'll repeat it anyway. We begin with a special case, namely when $m = 1$ so that $n = p^k$. In this case, we claim that the ring of integers of $\mathbf{Q}(\zeta_{p^k})$ has exactly one prime lying over p, with full ramification and no inertia. The construction is done explicitly: we claim that the prime lying over p is given by $(1 - \zeta_{p^k})$. First, we compute

$$\begin{aligned} \mathrm{N}(1-\zeta_{p^k}) &= N_{\mathbf{Q}(\zeta_{p^k})/\mathbf{Q}}(1-\zeta_{p^k}) \\ &= \prod_{a\in(\mathbf{Z}/p^k\mathbf{Z})^\times} (1-\zeta_{p^k}^a) \\ &= \Phi_{p^k}(1) \end{aligned}$$

where Φ_{p^k} denotes the p^k-th cyclotomic polynomial. Luckily,

$$\Phi_{p^k}(X) = \frac{X^{p^k}-1}{X^{p^{k-1}}-1} = 1 + X^{p^{k-1}} + X^{2p^{k-1}} + \cdots + X^{(p-1)p^{k-1}}$$

so in fact $\mathrm{N}(1-\zeta_{p^k}) = p$, which proves that $(1-\zeta_{p^k})$ is a prime in the ring of integers of $\mathbf{Q}(\zeta_{p^k})$ with trivial inertia. Moreover, if $a \in (\mathbf{Z}/p^k\mathbf{Z})^\times$,

$$\frac{1-\zeta_{p^k}^a}{1-\zeta_{p^k}} = 1 + \zeta_{p^k} + \cdots + \zeta_{p^k}^{a-1} \in \mathcal{O}_{\mathbf{Q}(\zeta_{p^k})}$$

which means that the Galois conjugates of $(1-\zeta_{p^k})$ are all equal as ideals, so p does not split at all in $\mathbf{Q}(\zeta_{p^k})$, has trivial inertia, and full ramification. This means that $D_{\mathbf{Q}(\zeta_{p^k})/\mathbf{Q}}(p) = I_{\mathbf{Q}(\zeta_{p^k})/\mathbf{Q}}(p) = (\mathbf{Z}/p^k\mathbf{Z})^\times$. Now let's consider the opposite extreme, namely $(n,p) = 1$, so $n = m$. We know p is unramified in $\mathbf{Q}(\zeta_m)$ by a discriminant computation for example (see the discussion after [3, Theorem 8]), so it has a well-defined Frobenius which is given by $\zeta_m \mapsto \zeta_m^p$ (see [6, Théorème 6.4.1]), i.e. the residue class of p in $(\mathbf{Z}/m\mathbf{Z})^\times$. So $D_{\mathbf{Q}(\zeta_m)/\mathbf{Q}}(p) = \langle p\rangle$ and $I_{\mathbf{Q}(\zeta_m)/\mathbf{Q}}(p) = 1$. Luckily, the field we are interested in is

$$\mathbf{Q}(\zeta_{p^k m}) = \mathbf{Q}(\zeta_{p^k})\mathbf{Q}(\zeta_m)$$

and the elements of the Galois group $(\mathbf{Z}/n\mathbf{Z})^\times$ are determined by their restrictions to the Galois groups of the two smaller cyclotomic extensions. By Theorem 3.22, $I(p)$ restricts to $\mathbf{Q}(\zeta_{p^k})$ as the full Galois group $(\mathbf{Z}/p^k\mathbf{Z})^\times$ and to $\mathbf{Q}(\zeta_m)$ as the trivial subgroup. This means that every element of $I(p)$ has $(\mathbf{Z}/m\mathbf{Z})^\times$-coordinate 1, and yet every possible $(\mathbf{Z}/p^k\mathbf{Z})^\times$-coordinate is attained, so in fact

$$I(p) = (\mathbf{Z}/p^k\mathbf{Z})^\times \times \{1\} \subset \mathrm{Gal}(\mathbf{Q}(\zeta_n)/\mathbf{Q}) \cong (\mathbf{Z}/p^k\mathbf{Z})^\times \times (\mathbf{Z}/m\mathbf{Z})^\times.$$

Similarly, $D(p)$ restricts to all of $(\mathbf{Z}/p^k\mathbf{Z})^\times$ in the first coordinate, and to $\langle p\rangle \subset (\mathbf{Z}/m\mathbf{Z})^\times$ in the second coordinate, but it also contains $I(p) = (\mathbf{Z}/p^k\mathbf{Z})^\times \times \{1\}$, which means in fact

$$I(p) = (\mathbf{Z}/p^k\mathbf{Z})^\times \times \langle p\rangle \subset \mathrm{Gal}(\mathbf{Q}(\zeta_n)/\mathbf{Q}) \cong (\mathbf{Z}/p^k\mathbf{Z})^\times \times (\mathbf{Z}/m\mathbf{Z})^\times.$$

In other words, for cyclotomic extensions of $\mathbf{Q}$ we can determine the entire splitting behavior (or at least the information of the decomposition and inertia groups) of a prime based on its residue mod n. The same is therefore true for subfields of $\mathbf{Q}(\zeta_n)$.

The inertia and decomposition groups are also useful for splitting up L/K into extensions where only one type of splitting happens at a time on the way from $\mathfrak{p}$ to $\mathfrak{P}$.

Let $Z = L^{D(\mathfrak{P}|\mathfrak{p})}$ and $T = L^{I(\mathfrak{P}|\mathfrak{p})}$, so that

$$K \subset Z \subset T \subset L$$

and T/Z is Galois since $I(\mathfrak{P}|\mathfrak{p})$ is normal in $D(\mathfrak{P}|\mathfrak{p})$ and both of these are closed (by Lemma 3.9 and Lemma 3.15), so

$$\mathrm{Gal}(L/Z) = D(\mathfrak{P}|\mathfrak{p}), \qquad \mathrm{Gal}(L/T) = I(\mathfrak{P}|\mathfrak{P}).$$

So by Theorem 3.22,

$$D_{L/Z}(\mathfrak{P}|\mathfrak{P}\cap Z) = D_{L/K}(\mathfrak{P}|\mathfrak{p}).$$

Let's go back to the case from math 129 where L/K is finite. Then this implies that

$$e(\mathfrak{P}|\mathfrak{P}\cap Z)f(\mathfrak{P}|\mathfrak{P}\cap Z) = e(\mathfrak{P}|\mathfrak{P})f(\mathfrak{p}|\mathfrak{p})$$

and therefore by the multiplicativity of inertial degree and ramification index, we see that

$$e(\mathfrak{P}\cap Z|\mathfrak{p}) = f(\mathfrak{P}\cap Z|\mathfrak{p}) = 1.$$

So there is no growth in residue field corresponding to $\mathfrak{P}$ when we go from K to Z. But if we increase the size of Z at all, then the decomposition group upstairs must shrink (since it is already all of $\mathrm{Gal}(L/Z)$), so Z is the maximal subextension such that $\mathfrak{P}\cap Z|\mathfrak{p}$ has trivial ramification and inertia. Similarly, by Theorem 3.22, $I(\mathfrak{P}|\mathfrak{P}\cap T) = I(\mathfrak{P}|\mathfrak{p})$, so if we also assume that the residue field extension is separable so that $|I| = e$, then T is the maximal subextension in which $\mathfrak{P}\cap T|\mathfrak{p}$ has trivial ramification. Since I is normal in D, it is also the maximal subextension of L/Z such that $\mathfrak{P}\cap Z$ is unramified in L. By looking at the inertial degrees, we see that the only growth in the residue field happens from $\kappa(\mathfrak{P}\cap Z)$ to $\kappa(\mathfrak{P}\cap T)$. So the slogan is as follows: as we go up from $\mathfrak{p}$ to $\mathfrak{P}$, all of the splitting occurs between K and Z (though there might be other primes in Z lying over $\mathfrak{p}$ which do not have trivial ramification and inertia), all of the growth in the residue field occurs between Z and T, and all of the ramification happens between T and L.

We can also get this to work when L/K is infinite. Let $F \subset L^{D(\mathfrak{P}|\mathfrak{p})}$ be finite over K, and let F' be its Galois closure over K. Then by Theorem 3.22, every element of $D_{F'/K}(\mathfrak{P}\cap F'|\mathfrak{p})$ is the restriction of some element of $D_{L/K}(\mathfrak{P}|\mathfrak{p})$, and therefore F is fixed by $D_{F'/K}(\mathfrak{P}\cap F'|\mathfrak{p})$. Since F'/K is finite, by the finite case we know that $\kappa(\mathfrak{P}\cap F) = \kappa(\mathfrak{p})$. Since this holds for all finite subextensions F of Z, and the union of all the $\kappa(\mathfrak{P}\cap F)$ is $\kappa(\mathfrak{P}\cap Z)$ by the same argument as in Theorem 3.12, it follows that $\kappa(\mathfrak{P}\cap Z) = \kappa(\mathfrak{p})$. So even in the infinite case, the residue field does not grow from $\mathfrak{p}$ to $\mathfrak{P}\cap Z$. It's also true that no ramification happens from $\mathfrak{p}$ to $\mathfrak{P}\cap Z$, since it doesn't happen in any finite subextension (Z/K is not Galois, so this is the only definition that makes sense).

Again, by Theorem 3.22, we know that T is the maximal Galois subextension of L/Z such that $I_{T/Z}(\mathfrak{P}\cap T|\mathfrak{P}\cap Z) = 1$, and therefore by Definition 3.17, it is the maximal Galois subextension in which $\mathfrak{P}\cap Z$ is unramified. Since $\mathrm{Gal}(L/T) = I_{L/T}(\mathfrak{P}|\mathfrak{P}\cap T)$, we know by Theorem 3.12 that $\mathrm{Aut}(\kappa(\mathfrak{P})/\kappa(\mathfrak{P}\cap T)) = 1$. If the residue field extension is separable this implies the residue field does not grow from $\mathfrak{P}\cap T$ to $\mathfrak{P}$.

Remark 3.24. Notice that in a lot of this we really only cared about one prime at a time, or in the relative setting, one prime upstairs and its intersection downstairs. So it will be useful (or at least convenient) to pass to the local setting (where much of the content of this class will take place), which will be explained soon.

§4 September 12, 2019

Let M/K be a finite Galois extension with group G, $\mathcal{O}_K$ a Dedekind domain with field of fractions K, $\mathcal{O}_L$ its integral closure in L, and $\mathfrak{p}$ a nonzero prime ideal in K which is unramified in M. Let $\mathfrak{P}$ be any prime of M lying over $\mathfrak{p}$. The decomposition group $D(\mathfrak{P}|\mathfrak{p})$ tells us all the primes lying over $\mathfrak{p}$, since the elements G/D correspond to prime ideals by the orbit-stabilizer theorem and the fact that G acts transitively on the primes $\mathfrak{P}|\mathfrak{p}$. We also know that $|D| = f(\mathfrak{P}|\mathfrak{p})$ and therefore $\mathfrak{p}$ splits completely in M if and only if D is trivial. We can consider an arbitrary subgroup $H \subseteq G$ and the fixed field $L = M^H$. We obtain a prime in $\mathcal{O}_L$ by taking $\mathfrak{P} \cap \mathcal{O}_L$. By the transitivity of the action of G on the primes of M lying over $\mathfrak{p}$, this shows that the primes of L lying over K are in bijection with the double cosets (of H and D) in G. Moreover, the intertial degrees are

$$[\kappa(\sigma\mathfrak{P} \cap L) : \kappa(\mathfrak{p})] = \frac{|D|}{|\sigma D \sigma^{-1} \cap H|}.$$

$\mathfrak{p}$ splits completely in L if and only if $G/D/H = [L : K] = |G/H|$. This is equivalent to saying that $D \subseteq gHg^{-1}$ for all $g \in G$. So this characterizes the splitting of primes in an arbitrary extension in terms of their splitting in a bigger Galois extensions.

Lemma 4.1

Let M be the Galois closure of L/K. Then $\mathfrak{p}$ splits completely in L if and only if it splits completely in M.

Proof. We know that $\mathfrak{p}$ splits completely in M iff $D = 1$ and that it splits completely in L iff $D \subseteq \bigcap_{g \in G} gHg^{-1} = H'$. So it suffices to show that the second one implies the first. Since H' is a normal subgroup of G, $M^{H'}$ we have

$$L \subseteq M^{H'} \subseteq M.$$

and $M^{H'}/K$ is Galois. If H' is nontrivial this would contradict the fact that M is the smallest extension of L which is Galois over K (that is the definition of the Galois closure). □

Theorem 4.2

Let K be a number field and $n \geq 1$. Then the following are equivalent:

1. For any two prime numbers $p \equiv p' \mod n$, then (p) and (p') split in the same way in K.
2. $K \subseteq \mathbf{Q}(\zeta_n)$.
3. For any primes $p, p' \equiv 1 \mod n$, either both of them split completely in K or neither of them do.

Proof. First, we show that $2 \implies 1$ (as usual we will exploit the nice structure of the frobenius for cyclotomic fields). First, suppose that p and p' divide n. Then they must be equal (they are congruent mod n). So we can consider the case where p and p' are both coprime to n. Therefore, p, p' are both unramified in $\mathbf{Q}(\zeta_n)$. They have Frobenius element $\mathrm{Frob}_p = p \mod n \in (\mathbf{Z}/n\mathbf{Z})^\times$ (in particular they have the same Frobenius element). Since the Frobenius generates the decomposition group, this implies that p and p' have the same splitting type in $\mathbf{Q}(\zeta_n)$ and therefore in K by the previous machinery. To prove that $3 \implies 2$, since a prime splits completely if and only if it splits completely in the Galois closure, we can replaces K with its Galois closure and assume $K/\mathbf{Q}$ is Galois. An unramified prime p splits completely if and only if its Frobenius is trivial in $\mathrm{Gal}(K/\mathbf{Q})$. Let $L = K(\zeta_n) = K \cdot \mathbf{Q}(\zeta_n)$.

$$\begin{tikzcd}
 & L \arrow[dl, dash] \arrow[dr, dash] & \\
K \arrow[dr, dash] & & \mathbf{Q}(\zeta_n) \arrow[dl, dash] \\
 & \mathbf{Q} &
\end{tikzcd}$$

Then $\mathrm{Gal}(L/\mathbf{Q})$ surjects by restriction onto $\mathrm{Gal}(K/\mathbf{Q})$ and $\mathrm{Gal}(\mathbf{Q}(\zeta_n)/\mathbf{Q}) = (\mathbf{Z}/n\mathbf{Z})^\times$. If K is not inside $\mathbf{Q}(\zeta_n)$, then $\mathrm{Gal}(L/K)$ does not contain $\mathrm{Gal}(L/\mathbf{Q}(\zeta_n))$, so there exists some $\sigma \in \mathrm{Gal}(L/\mathbf{Q})$ such that σ restricts to id on $\mathbf{Q}(\zeta_n)$ but not on K. By the Chebotarev density theorem, there is a prime p of $\mathbf{Q}$ with $\mathrm{Frob}_p(\mathbf{Q}(\zeta_n)) = \mathrm{id}$ (i.e. $p \equiv 1 \mod n$) and not splitting completely in K, as well as a prime p that splits completely in both. This proves $3 \implies 2$.

□

For those unfamiliar with the statement, we restate the Chebotarev Density Theorem below:

Theorem 4.3

Let L/K be a finite Galois extension of number fields. Let $\mathcal{C}$ be a conjugacy class in $\mathrm{Gal}(L/K)$. Then, the set of (unramified) primes $\mathfrak{p}$ of K with Frobenius element in $\mathcal{C}$ has Dirichlet density $|\mathcal{C}|/|G|$.

Recall that Dirichlet density is defined with respect to the ordering of the prime ideals according to their norm. In fact, one can show that the result is also true for natural density with respect to this ordering.

Example 4.4

Let $G = S_3$. Then the conjugacy classes of G are just the identity, the two-cycles, and the three-cycles. Let p be an unramified prime. If Frob_p is the identity then p splits completely. So the density of the primes that split completely is $1/6$. If Frob_p is a two-cycle, then D is generated by a two-cycle so it is just $\mathbf{Z}/2\mathbf{Z}$, which means it splits into three different ideals with probability $1/2$. If Frob_p is a three-cycle, then $|D| = 3$, so p splits into 2 primes with probability $1/3$.

If L/K is infinite, then there are many parts that need to be salvaged. For example, it is possible that infinitely many primes ramify, and the fraction $|\mathcal{C}|/|G|$ is not well-defined. Even if we take a normalized Haar measure on G, most of the time the measure of $\mathcal{C}$ is zero.

Now we return to the problem of quadratic residues. Now that we see that the periodicity is related to being contained in a cyclotomic extension, we can view quadratic reciprocity as a special case of that.

Example 4.5

We will see that $\mathbf{Q}(\sqrt{n}) \subseteq \mathbf{Q}(\zeta_{4n})$. So the splitting behavior of p in $\mathbf{Q}(\sqrt{n})$ depends only on $p \mod n$. It suffices to show the inclusion when $n = q$ is prime. The case $p = 2$ is done separately. We can just compute directly that $\sqrt{2} = \zeta_8 + \zeta_8^{-1}$. Otherwise, we know that the quadratic subextensions of $\mathbf{Q}(\zeta_q)$ correspond to subgroups of index 2 of $(\mathbf{Z}/q\mathbf{Z})^\times$ which is cyclic of even order $q - 1$ (this is why we need to deal with q even as a separate case). Since this is even there is exactly one subgroup of even order (the subgroup H containing all the quadratic residues mod q), and by computing the discriminants we can get that this subfield is $\mathbf{Q}(\sqrt{q}) \subseteq \mathbf{Q}(\zeta_q)$. Another way to see this is is by computing the Gauss sum

$$\alpha = \sum_{x \in (\mathbf{Z}/q\mathbf{Z})^\times} \left(\frac{x}{q}\right) \zeta_q^x.$$

A Galois conjugate of α looks like

$$\sum_{x \in (\mathbf{Z}/q\mathbf{Z})^\times} \left(\frac{x/y}{q}\right) \zeta_q^x$$

which is the same as

$$\sum_{x \in (\mathbf{Z}/q\mathbf{Z})^\times} \left(\frac{x}{q}\right)\left(\frac{y}{q}\right) \zeta_q^x$$

which is $\pm\alpha$, so the Gauss sum should have degree 2 over $\mathbf{Q}$ (it only has two Galois conjugates). We can even compute α^2 directly, to get

$$\alpha^2 = \sum_{x_1, x_2} \left(\frac{x_1 x_2}{q}\right) \zeta_q^{x_1 + x_2} = \left(\frac{-1}{q}\right) q = \pm q$$

depending on the value of $q \mod 4$. So in particular $\mathbf{Q}(\sqrt{\pm q}) \subseteq \mathbf{Q}(\zeta_q)$. Since $\sqrt{-1} \in \mathbf{Q}(\zeta_4)$ this shows that $\mathbf{Q}(\sqrt{q}) \subseteq \mathbf{Q}(\zeta_q)$.

§5 September 17, 2019

Remember that we found a criterion for the splitting type of a polynomial mod p to depend only on p mod n, namely that the splitting field is contained in $\mathbf{Q}(\zeta_n)$ for some n. It turns out that this is the same thing as the splitting field being abelian.

Theorem 5.1 (Kronecker–Weber)

A finite Galois field extension $K/\mathbf{Q}$ is abelian if and only if $K \subset \mathbf{Q}(\zeta_n)$ for some n. As a result, the maximal abelian extension of $\mathbf{Q}$ is $\mathbf{Q}^{\mathrm{ab}} = \mathbf{Q}(\zeta_\infty)$ which we saw has Galois group $\widehat{\mathbf{Z}}^\times$.

Definition 5.2. The least n such that $K \subseteq \mathbf{Q}(\zeta_n)$ is called the **conductor** for K. Note that it is the gcd of all n that work.

Example 5.3

Let $K = \mathbf{Q}(\sqrt{a})$. Then K is abelian with Galois group $\mathbf{Z}/2\mathbf{Z}$. It can be seen that the conductor is $|d_K|$.

Given the success of Kronecker–Weber in the case of $\mathbf{Q}$, we can ask

Question 5.4. Let K be a number field. What are the abelian extensions of K? What is K^{ab}? What is $\mathrm{Gal}(K^{\mathrm{ab}}/K)$?

The determination of the abelian extensions of K is unknown except for when K is $\mathbf{Q}$ (Kronecker–Weber) or an imaginary quadratic field (the theory of complex multiplication on elliptic curves).

One useful strategy for approaching this problem is a "local–to–global" approach, which starts with the corresponding local problem involving the field of p-adic numbers $\mathbf{Q}_p$.

Theorem 5.5 (Local Kronecker–Weber)

Let $\mathbf{Q}_p$ be the field of fractions of $\mathbf{Z}_p$. Then

$$\mathbf{Q}_p^{\mathrm{ab}} = \bigcup_{n\geq 0} \mathbf{Q}_p(\zeta_n).$$

More generally, if K is a number field we can obtain an explicit description of $K_{\mathfrak{p}}^{\mathrm{ab}}$ and its Galois group.

Why study $\mathbf{Q}_p$ first?

Example 5.6

Let $f(X_1, \dots, X_r) \in \mathbf{Z}[X_1, \dots, X_r]$. Does f have a root in $\mathbf{Z}^r$? One way to show that this doesn't have a solution is to show that it doesn't have a solution mod n for some $n \in \mathbf{Z}$. Another possible obstruction to f having a root in $\mathbf{Z}$ would be because it does not even have a root in $\mathbf{R}$. Checking whether f has a root in $\mathbf{R}$ belongs to the realm of numerical analysis. By the Chinese remainder theorem, f has a solution mod all n if and only if it has a solution in $\mathbf{Z}_p^r$ for all p. This implies that f has a root mod p. In fact, most of the time roots in $\mathbf{F}_p$ lift to roots in $\mathbf{Z}_p$ (this is "Hensel's Lemma"). So if we want to show that f has no solutions, then we should try to show that it has no solutions mod p for any p. If these are equivalent for a class of polynomials, that class is said to have the *Hasse Local-Global Principle*. For example, Hasse and Minkowski showed that the binary quadratic forms satisfy the local-global principle.

In general, we can expect $\mathbf{F}_p$ to be easier than $\mathbf{Z}_p$, and $\mathbf{Z}_p$ to be easier than $\mathbf{Z}$. So it is very useful (if some kind of local-global principle holds) to study the simpler rings/fields instead.

Definition 5.7. Let K be a field. A **discrete valuation** on K is a map $v : K \to \mathbf{R} \cap \{\infty\}$ such that

- $v(x) = \infty$ if and only if $x = 0$.
- $v(xy) = v(x) + v(y)$.
- $v(x + y) \geq \min(v(x), v(y))$.
- $v(K)$ has a smallest positive element $s \in \mathbf{R}$, so that $v(K) = s\mathbf{Z}$.

We can always multiply a discrete valuation by $\frac{1}{s}$ to get a **normalized discrete valuation**, a discrete valuation whose **value group** (the image of K under v) is equal to $\mathbf{Z} \cup \{\infty\}$.

Example 5.8

Let $\mathcal{O}_K$ be a Dedekind domain with field of fractions K and a nonzero prime ideal $\mathfrak{p} \subseteq \mathcal{O}_K$. Then the $\mathfrak{p}$-adic valuation on K is defined by

$$v_{\mathfrak{p}}(x) = \sup\{n \in \mathbf{Z} : x \in \mathfrak{p}^n\}.$$

In other words, it is the exponent of $\mathfrak{p}$ in the unique expansion of the fractional ideal $x \cdot \mathcal{O}_K$ as a product of prime powers.

Note that if v is a discrete valuation on K then

$$v(1) = v(-1) = 0,$$

$$v(x/y) = v(x) - v(y),$$

and (most importantly) if $v(x) \neq v(y)$ then in fact we have the equality

$$v(x + y) = \min(v(x), v(y)).$$

If K is a number field, then there is a famous result, due to Ostrowski, which states that all the discrete valuations are $\mathfrak{p}$-adic.

Definition 5.9. Let v be a normalized discrete valuation on K. Then the **valuation ring** of K is

$$\mathcal{O}_{K,v} = \{x \in K : v(x) \geq 0\}$$

(it is a principal ideal domain), and it has group of units

$$\mathcal{O}_{K,v}^{\times} = \{x \in K : v(x) = 0\}.$$

The non-units in $\mathcal{O}_{K,v}$ form an ideal. As a result, $\mathcal{O}_{K,v}$ is a local ring with maximal ideal $\mathfrak{p}_v$ equal to the set of non-units. Since $\mathcal{O}_{K,v}$ is Dedekind, every ideal is therefore of the form $\mathfrak{p}_v^n$ for some $n \geq 0$. So for any **uniformizer** for v, namely a $\pi_v \in K$ with $v(\pi_v) = 1$ [in other words $\pi_v \in \mathfrak{p}_v - \mathfrak{p}_v^2$], we can write that any ideal in $\mathcal{O}_{K,v}$ is of the form (π_v^n) for some $n \geq 0$. If $a \leq b$, then we have an isomorphism of $\mathcal{O}_{K,v}$-modules

$$\mathfrak{p}_v^a / \mathfrak{p}^b \cong \mathcal{O}_{K,v} / \mathfrak{p}_v^{b-a}.$$

From the fact (proved inside the definition oops) that the valuation ring is a DVR, we have the following as well:

Lemma 5.10

If I is the ideal generated by $\alpha_1, \ldots, \alpha_r$ in $\mathcal{O}_{K,v}$, then I is generated by any α_i with smallest valuation.

Example 5.11

Let $K = \mathbf{Q}$ and $v = v_p$ for a rational prime p. Then

$$\mathcal{O}_{K,v} = \mathbf{Z}_{(p)} = \{\frac{a}{b} : b \not\equiv 0 \mod p\}$$

is just the localization of $\mathbf{Z}$ at (p).

If $\mathcal{O}_K$ is a DVR, then $\mathcal{O}_{K,v} = \mathcal{O}_K$. In general, if you start with a Dedekind domain which is not a DVR, then by taking the valuation ring of its field of fractions at a particular $\mathfrak{p}$-adic valuation you end up with something larger than $\mathcal{O}_K$. Moreover, taking valuation rings leaves important invariants intact. For example, we have a ring isomorphism

$$\mathcal{O}_K / \mathfrak{p}^n \cong \mathcal{O}_{K,v} / \mathfrak{p}_v^n.$$

Now back to the question of classifying the valuations on a number field, and ultimately the proof of Ostrowki's theorem.

Theorem 5.12

Any normalized discrete valuation on $\mathbf{Q}$ is v_p for some rational prime p.

Proof. We know $\mathcal{O}_{\mathbf{Q},v} \subseteq \mathbf{Q}$ is a subring, so $\mathbf{Z} \subseteq \mathcal{O}_{\mathbf{Q},v}$. Since $\mathfrak{p}_v \cap \mathbf{Z}$ is a maximal ideal containing only elements of positive valuation (in particular not 1), we know

$$\mathfrak{p}_v \cap \mathbf{Z} = p\mathbf{Z}$$

for some rational prime p. As a result, $v(p) \geq 1$, and $v(q) = 0$ for all $q \neq p$. Because of how valuations work on products and by unique factorization of $\mathbf{Z}$, we can deduce that $v = v_p$ (there must be some element of valuation 1, so $v(p) = 1$, then proceed by prime factorization). □

The other distinguished example of this theory (other than number fields) is that of function fields over a finite field.

Lemma 5.13

A finite field $\mathbf{F}_q$ has no discrete valuations.

Proof. Any nonzero element of $\mathbf{F}_q$ is a root of unity, so $v(x) = 0$ for all $x \in \mathbf{F}_q$, which means v cannot be a legitimate valuation. □

Lemma 5.14

An algebraically closed field K has no discrete valuations.

Proof. without loss of generality v is normalized, so there is an $x \in K$ such that $v(x) = 1$. Then $v(\sqrt{x}) = 1/2$ which contradicts the fact that v is normalized. □

Theorem 5.15

Let k be a finite field or an algebraically closed field. Then any normalized discrete valuation on $K = k(T)$ is $v = v_f$ for some irreducible polynomial $f \in k[T]$, or $v = v_{\deg}$, the valuation given by taking $v(a/b) = \deg(b) - \deg(a)$.

Proof. Restrict v to k. Then $v|_{k^\times} = 0$ by the previous two lemmas. In particular, $k \subseteq \mathcal{O}_{K,v}$. If $v(T) \geq 0$, then $k[T] \subseteq \mathcal{O}_{K,v}$. So $\mathfrak{p}_v \cap k[T]$ is a nonzero prime ideal $(f(T))$ and it follows that $v = v_{f(T)}$. On the other hand, if $v(T) < 0$ then $k[1/T] \subseteq \mathcal{O}_{K,v}$. As a result, $\mathfrak{p}_v \cap k[1/T]$ is a nonzero prime ideal containing $1/T$, and thus v is the $(1/T)$-adic valuatoin which is the same as $v_{\deg}$, as desired. □

Back to number fields. Notice that for example $\mathbf{Z}_{(p)} \subseteq \mathbf{Z}_p$, but $\mathbf{Z}_p$ is a lot nicer. The construction of $\mathbf{Z}_p$ can be generalized to take the completion of any field with respect to a discrete valuation.

Definition 5.16. Let K be a field with a discrete valuation v. Then the **completion** of $\mathcal{O}_{K,v}$ with respect to v is

$$\widehat{\mathcal{O}}_{K,v} = \varprojlim_{n \geq 0} \mathcal{O}_{K,v}/\mathfrak{p}_v^n.$$

Moreover, $\mathcal{O}_{K,v}$ is **complete** if it is equal to its completion via the embedding that takes an element to its residue in each coordinate. The **completion** of K with respect to v is just the field of fractions of the completion of $\mathcal{O}_{K,v}$.

Example 5.17

The completion of $\widehat{\mathbf{Z}}_{(p)}$ is $\mathbf{Z}_p$, whose field of fractions is $\mathbf{Q}_p$.

Definition 5.18. Let v be a discrete valuation on K with finite residue field $\kappa(v) = \mathcal{O}_v/\mathfrak{p}_v \cong \mathbf{F}_q$. Then we define the **norm** (also called an **absolute value**) on K coming from v to be given by

$$|x|_v = q^{-v(x)}.$$

This is the norm that comes naturally out of the Haar measure on K. Intuitively, this definition means that $x, y \in K$ are close if they are divisible by a high power of $\mathfrak{p}_v$. In fact, the norm coming from v is a norm on K as a $\mathbf{Q}$-vector space, and thus gives it the structure of a metric space. It satisfies the axioms of a **nonarchimedean norm**. The completion of K and $\mathcal{O}_{K,v}$ with respect to the norm coming from v actually coincide with the definitions above.

Lemma 5.19

$\widehat{\mathcal{O}}_{K,v}$ is the completion of $\mathcal{O}_{K,v} \subseteq K$ with respect to $|\cdot|_v$, and analogously for K and K_v.

Lemma 5.20

If $S \subseteq \mathcal{O}_v$ is a set of representatives for the residue field, then each $a \in \widehat{\mathcal{O}}_{K,v}$ can be written uniquely as a convergent series

$$a = \sum_{n=0}^{\infty} \alpha_i \pi_v^i$$

where $\alpha_i \in S$ and π_v is a fixed uniformizer for v. Each $a \in K_v$ may be written uniquely as

$$a = \pi_v^m \sum_{i=0}^{\infty} a_i \pi_v^i$$

where $a_0 \notin \mathfrak{p}_v$.

Note that v extends to a discrete valuation on K_v via

$$v(\lim a_n) = \lim v(a_n)$$

(by this definition the resulting valuation is discrete and normalized if the original one was). Because of this, the valuation ring of the completion is the same as the completion of the valuation ring, and same for the maximal ideal. A canonical isomorphism between residue fields can also be given.

§6 September 19, 2019

Definition 6.1. A **local field** is a field K which is complete with respect to a discrete valuation and whose residue field is finite.

Example 6.2

$\mathbf{Q}_p$ is a local field. Its valuation ring is $\mathbf{Z}_p$ and its residue field is $\mathbf{F}_p$.

For any local field K, the valuation ring $\mathcal{O}_{K,v}$ is compact because $\mathcal{O}_{K,v} \subseteq \prod_{n\geq 1} \mathcal{O}_{K,v}/\mathfrak{p}_v$.

Remember that every element of $\mathbf{Z}_p$ can be written as a convergent power series $\sum a_i p^i$. The units of $\mathbf{Z}_p$ are the elements whose first digit (a_0) is nonzero.

Theorem 6.3

Let v be a normalized discrete valuation on K. For any $n \geq 1$, consider the **higher unit group**

$$U^{(n)} = 1 + \mathfrak{p}_v^n = \{x \in K : v(x-1) \geq n\} \subseteq \mathcal{O}_{K,v}^\times.$$

Then $\mathcal{O}_v^\times / U^{(n)} \cong (\mathcal{O}_{K,v}/\mathfrak{p}_v^n)^\times$ and $U^{(n)}/U^{(n+1)} \cong \mathcal{O}_{K,v}/\mathfrak{p}_v$.

Proof. The first isomorphism can be made natural. In particular we have a natural map $\mathcal{O}_{K,v}^\times \to (\mathcal{O}_{K,v}/\mathfrak{p}_v^n)^\times$, and its kernel is $U^{(n)}$. The second isomorphism is not natural. We use the surjective map $U^{(n)} \to \mathcal{O}_{K,v}/\mathfrak{p}_v$ given by

$$1 + \pi_v^n t \mapsto t \mod \mathfrak{p}_v$$

and see that its kernel is equal to $U^{(n+1)}$. This isomorphism is not canonical because it depends on a choice of uniformizer. □

Lemma 6.4 (Hensel's Lemma, most basic version)

Let K be complete with respect to the discrete valuation v. Let $f(X) \in \mathcal{O}_{K,v}[X]$ and $\overline{a} \in \mathcal{O}_{K,v}/\mathfrak{p}_v$ be a simple root of $\overline{f(X)}$. Then there is exactly one root f in $\mathcal{O}_{K,v}$ which reduces to $\overline{a}$ mod $\mathfrak{p}_v$.

Proof. To prove the existence, we need to construct a sequence of elements $\alpha_1, \ldots \in \mathcal{O}_{K,v}$ such that $f(\alpha_n) \to 0$ as $n \to \infty$ and α_n converges in K. To do this is suffices to ensure that

$$|\alpha_{n+1} - \alpha_n|_v \leq |f(\alpha_n)|_v$$

because this ensures that the sequence is Cauchy. For α_0, take any $\alpha_0 \equiv \overline{\alpha} \mod \mathfrak{p}_v$. Then $|f(\alpha_0)| < 1$ and $|f'(\alpha_0)| = 1$ since α_0 is a simple root mod $\mathfrak{p}$. The procedure looks a lot like Newton's method, but in our case it will always succeed. In particular, we take

$$\alpha_{n+1} = \alpha_n - \frac{f(\alpha_n)}{f'(\alpha_n)},$$

so that

$$|\alpha_{n+1} - \alpha_n| = |f(\alpha_n)|.$$

By Taylor expansion of polynomials, we know

$$f(\alpha_n + t) = f(\alpha_n) + f'(\alpha_n)t \mod t^2,$$

so in fact (applying this to $t = -f(\alpha_n)/f'(\alpha_n)$) we have

$$|f(\alpha_{n+1})| = |f(\alpha_n + t)| \le |t|^2 = |f(\alpha_n)|^2.$$

Using the construction, it is clear that $f(\alpha_n) \to 0$ since $|f(\alpha_0)| < 1$. We also have

$$|\alpha_{n+1} - \alpha_n| < |f(\alpha_n)| \to 0,$$

so by telescoping a sum and using the nonarchimedean triangle inequality it follows that α_n is Cauchy. □

Example 6.5

Let $a \in \mathbf{Z}$ which is a quadratic residue mod p, where a is nonzero mod p and $p \neq 2$. Then $a = x^2$ for some $x \in \mathbf{Z}_p$, by applying Hensel's lemma to the polynomial $f(X) = X^2 - a$. Then $f'(X) = 2X$ so if $p \neq 2$ and $a \neq 0 \mod p$ then any root of this mod p is simple and Hensel's lemma lets us lift the root to $\mathbf{Z}_p$. But lifting the root can fail in both cases we left out. For example, p is clearly not a square in $\mathbf{Z}_p$, and 3 is not a square in $\mathbf{Z}/4$.

Example 6.6

$\mathbf{Z}_p$ contains all $(p-1)$-th roots of unity. This is because you can apply Hensel's lemma to $f(X) = X^{p-1} - 1$, which has $p - 1$ distinct simple roots in $\mathbf{F}_p^\times$.

There is also a stronger version of Hensel's lemma

Theorem 6.7 (Hensel's lemma for factors)

Let K be complete with respect to the normalized discrete valuation v. If $f(X) \in \mathcal{O}_{K,v}[X]$ factors in the residue field as

$$\overline{f(X)} = \overline{h(X)g(X)} \mod \mathfrak{p},$$

where $\overline{h}, \overline{g}$ are coprime. Then there are $h(X), g(X) \in \mathcal{O}_{K,v}$ such that

$$f(X) = g(X)h(X)$$

and $\overline{g} \cong g \mod \mathfrak{p}$ and analogously for h, plus $\deg(h) = \deg(\overline{h})$ and same for g.

Proof. See Neukirch, Theorem II.4.6 □

There are also useful tricks for dealing with the valuations of roots of polynomials given their coefficients (the general idea is that of **Newton polygons**)

Lemma 6.8

Let $f(X) = a_nX^n + \cdots + a_0 \in K[X]$ be irreducible. Then $v(a_i) \ge \min(v(a_n), v(a_0))$ for all i.

Proof. Without loss of generality we can clear denominators to force $v(a_i) \geq 0$, and $v(a_i) = 0$ for some i. Let i be the minimal index such that $v(a_i) = 0$, so that

$$f(X) = X^i(a_nX^{n-i} + \cdots + a_i) \mod \mathfrak{p}_v.$$

Unless one of these polynomials is constant mod $\mathfrak{p}_v$ (i.e. $i = 0$ or n), Hensel's lemma says that $f(X)$ is not irreducible. In the exceptional cases, we always have $v(a_j) \geq 0 = \min(v(a_n), v(a_0))$. □

§7 September 24, 2019

Remember the lemma from last class:

Lemma 7.1

Let $f(x) = a_0 + a_1x + \cdots a_nx^n$ be an irreducible polynomial over a complete discretely valued field. Then $v(a_i) \geq \min(v(a_n), v(a_0))$.

One reason it is relevant is because of

Theorem 7.2

Let K be complete with respect to the discrete valuation v, and L/K a finite extension of degree n. Then there is exactly one extension v' of L that extends K. It is given by $v'(x) = \frac{1}{n}v(N_{L/K}x)$

Proof. First of all, v' at least restricts to v on K (obvious from its definition). $N_{L/K}x = 0$ if and only if $x = 0$, so the fact that $v'(x) = 0$ if and only if $x = 0$ follows from the same property of v. It is also clear that $v'(xy) = v'(x) + v'(y)$ by the multiplicativity of the norm.

Lemma 7.3

$x \in L$ is integral over $\mathcal{O}_{K,v}$ if and only if $v'(x) \geq 0$

Proof. Let $f(X) = a_0 + a_1X + \cdots + a_{t-1}X^{t-1} + X^t$ be the minimal polynomial of x. Then $t = [K(x) : K]$, and

$$N_{L/K} = \pm a_0^{[L:K(x)]}.$$

If x is integral over $\mathcal{O}_{K,v}$ then all the coefficients of the minimal polynomial are in $\mathcal{O}_{K,v}$, so $N_{L/K}(x)$ has nonnegative valuation, and hence $v'(x) = (1/n)v(Nx) \geq 0$. Conversely, if $v'(x) \geq 0$, then $v(a_0) \geq 0$, and by the lemma $v(a_i) \geq 0$ for all i, and thus f has coefficients in the valuation ring. □

The nonarchimedean triangle inequality can be proved in the following way: WLOG assume that $v'(x) \geq v'(y)$. Then $v'(x/y) \geq 0$, so x/y is integral over $\mathcal{O}_{K,v}$. The same is true of $x/y + 1$ and therefore $v'(x + y) \geq v'(y) = \min(v'(x), v'(y))$ as desired. We should also show that L is complete under v'. The fact that L is complete with respect to v' is a special case of the general fact that any finite dimensional normed vector space over a complete field is complete (in particular it has nothing to do with how we defined v').

It remains to show that v' is unique. Let v'' be some other extension of v to L. Then $\mathcal{O}_{L,v''}$ is an integrally closed PID containing $\mathcal{O}_v$, so by the lemma

$$\mathcal{O}_{v'} \subseteq \mathcal{O}_{v''}.$$

So $\mathfrak{p}_{v''} \cap \mathcal{O}_{v'}$ is a nonzero prime ideal of $\mathcal{O}'_v$, so actually

$$\mathfrak{p}_{v'} \subseteq \mathfrak{p}_{v''}.$$

If $v''(x) \geq 0$, then $v''(1/x) \leq 0$, so $1/x \notin \mathfrak{p}_{v''}$ and thus $1/x \notin \mathfrak{p}_{v'}$ which means $v'(1/x) \leq 0$, i.e. $v'(x) \geq 0$. This means that in fact $\mathcal{O}_{v'} = \mathcal{O}_{v''}$ and thus $v' = v''$, as desired. If $v''(x) \geq 0$, then $v''(1/x) \leq 0$, so $1/x \notin \mathfrak{p}_{v''}$ and thus $1/x \notin \mathfrak{p}_{v'}$ which means $v'(1/x) \leq 0$, i.e. $v'(x) \geq 0$. This means that in fact $\mathcal{O}_{v'} = \mathcal{O}_{v''}$ and thus $v' = v''$, as desired. □

The fact that v' is unique tells us that $v' \circ \sigma = v'$ (in the Galois case), so taking the sum over all σ tells us $v'(N_{L/K}) = nv'(\alpha)$. This is what motivates the construction of v' in the theorem.

Corollary 7.4

Finite extensions of local fields are local fields.

Proof. Let $\mathfrak{p}_v \mathcal{O}_{L,v'} = \mathfrak{p}_{v'}^e$. So $v'(\pi'_v) = \frac{1}{e} v'(\pi_v)$, and thus $v'(L^\times) = \frac{1}{e} v(K^\times)$ so v' is discrete, which is all we had left to show. □

Corollary 7.5

There is exactly one valuation v' on $\overline{K}$ extending v, but it is not discrete (in fact its value group is dense in $\mathbf{R}$).

It might also not be true that $\overline{K}$ is complete with respect to v'. But its completion at v' is still algebraically closed. Still, the valuation ring at v' is the integral closure of $\mathcal{O}_{K,v}$ and it has a single nonzero prime which is the set of elements of positive valuation.

If K is complete wrt a discrete valuation v, then we call the corresponding normalized discrete valuation v_K. We also denote the extension to $\overline{K}$ by v_K. Also, we might as well stop writing $\mathcal{O}_{K,v}$ and write $\mathcal{O}_K$ when K is a local field, etc.

If L/K is finite, then $v_L(x) = ev_K(x)$ where e is the ramification index of L/K.

Corollary 7.6

If $f(X) \in K[X]$ is irreducible, then all of its roots have the same valuation.

Remark 7.7. Let $f(X) \in \mathbf{R}[X]$ be irreducible. Then in fact f has degree 1 or 2, and all the roots have the same valuation (if there are more than one then they are conjugates). Note that $|x|_{\mathbf{C}} = \sqrt{|x\overline{x}|_{\mathbf{R}}}$ is the only extension of the absolute value to $\mathbf{C}$.

Where did all the primes go? Let $\mathcal{O}_K$ be a Dedekind domain with field of fractions K, $\mathfrak{p} \subseteq \mathcal{O}_K$ a nonzero prime, L/K a finite extension, and $\mathcal{O}_L$ the integral closure of $\mathcal{O}_K$ in L. Then

$$\mathfrak{p}\mathcal{O}_L = \mathfrak{p}_1^{e_1} \cdots \mathfrak{p}_r^{e_r}.$$

We can show that $L \otimes K_{\mathfrak{p}} \cong L_{\mathfrak{p}_1} \times \cdots \times L_{\mathfrak{p}_r}$, and the analogous isomorphism holds for the rings of integers.

Proof. Recall that $\widehat{\mathcal{O}}_{\mathfrak{p}} = \varprojlim \mathcal{O}_K/\mathfrak{p}_K^n$, so when we tensor by $\mathcal{O}_L$ we get

$$\prod_{i=1}^{r} \varprojlim \mathcal{O}_L/\mathfrak{p}_i^{e_i n},$$

which is the same as the product of the $\widehat{\mathcal{O}}_{p_i}$, because the tensor product commutes with inverse limits (this works if L/K is separable, in which case $\mathcal{O}_L$ is a finitely-generated $\mathcal{O}_K$-module). The fact for the fields follows from this. □

Recall the general form of the Dedekind-Kummer theorem: If $\mathfrak{p}|[\mathcal{O}_L : \mathcal{O}_K[\alpha]]$ for some $\alpha \in L$, then its minimal polynomial factors in $\widehat{\mathcal{O}}_{\mathfrak{p}}[X]$ as $f_1 \cdots f_r$ where f_i is irreducible of degree $[L_{\mathfrak{p}_i} : K_{\mathfrak{p}}]$. Moreover, f_i factors mod $\mathfrak{p}$ as $f_i(X) = g_i(X)^{e_i}$ with g_i irreducible of degree equal to $f(\mathfrak{p}_i|\mathfrak{p})$.

Theorem 7.8

Let K be a field with valuation v and $r_1, \dots, r_n \in K$. WLOG assume $v(r_1) \le v(r_n)$. Then the coefficients of

$$\prod_i (X - r_i)$$

satisfy $v(a_{n-i}) \ge v(r_1) + \cdots v(r_i)$. If $i = n$ or $v(r_i) < v(r_{i+1})$.

Proof. By "Vieta", we know a_{n-i} is (up to a sign) a sum of products of i of the roots. The smallest possible valuation is the term coming from the first i roots (since they come in order). So the result follows from the nonarchimedean triangle inequality. □

This basically begins the theory of Newton polygons. The theorem we just proved says that the **Newton polygon** of $f(X)$, namely the lower convex hull of the points $(i, v(a_i))$, encodes the valuations of the roots via the slopes of the line segments comprising it.

Corollary 7.9

If f is irreducible, its Newton polygon is just a line segment since all roots have the same valuation.

The Newton polygon of the product of polynomials can be written down in terms of the two Newton polygons (since you just take the union of the roots).

Corollary 7.10 (Corollary of Corollary)

If the Newton polygon of f is a line segment which contains no integer points except the endpoints, then f must be irreducible.

Corollary 7.11 (Eisenstein's criterion; Corollary of Corollary of Corollary)

$f(X)$ is irreducible if its Newton polygon is the line segment $[(0,1),(n,0)]$.

Newton polygons are a good tool for checking how a polynomial factors, but it isn't always sufficient for determining even whether is is reducible.

§8 September 26, 2019

Let $\mathcal{O}_K$ be a Dedekind domain with field of fractions K, L/K a Galois extension, and $\mathfrak{P} \subseteq \mathcal{O}_L$ a prime lying over $\mathfrak{p} \subseteq \mathcal{O}_K$. In that case, there is a canonical identification

$$\operatorname{Gal}(L_{\mathfrak{P}}/K_{\mathfrak{p}}) \cong D(\mathfrak{P}|\mathfrak{p}).$$

If L/K is already complete then the decomposition group is the entire Galois group (there is only one prime upstairs). The inertia group in the completion is the same as the inertia group, too (since taking completions leaves the residue fields invariant). If L/K is a finite Galois extension, and

$$\mathfrak{p}\mathcal{O}_L = (\mathfrak{p}_1 \cdots \mathfrak{p}_r)^e,$$

then $\operatorname{Gal}(L/K)$ permutes the fields in $L \otimes_K K_{\mathfrak{p}} = \prod_{i=1}^r L_{\mathfrak{q}}$ according to the way it permutes the primes. This is another technical thing you can use to translate between local and global.

Theorem 8.1

Let K be complete with respect to a discrete valuation, and L/K a finite extension of ramification index e, inertial degree f, and degree n. Then $n = ef$.

Proof. Let $\omega_1, \ldots, \omega_f \in \mathcal{O}_L$ be so that their reductions mod $\mathfrak{P}_L$ are a basis for κ_L/κ_K. Then we claim that $\{\omega_i \pi_L^j\}_{1 \le i \le f, 0 \le j \le e}$ is a basis for $\mathcal{O}_L/\mathcal{O}_K$. For linear independence (which we can just do over K), suppose

$$0 = \sum_{i,j} a_{ij} \omega_i \pi_L^j = 0.$$

For any j, we have

$$v_L\left(\sum_i a_{ij}\omega_i\right) = \min_i\{v_L(a_{ij})\} = \min_i\{e v_K(a_{ij})\}$$

which is divisible by e unless it is ∞. The reason why it is equal to the minimum of the valuations is because of the fact that the $\omega_i \mod \mathfrak{p}_L$ are a basis for κ_L. If a_{1j} has the smallest valuation, then

$$v_L\left(\sum_i a_{ij}\omega_i\right) = v_L(a_1 j) + v_L\left(\sum_i \frac{a_{ij}}{a_{1j}}\omega_i\right),$$

where we know $a_{ij}/a_{1j} \in \mathcal{O}_K$. If the sum on the RHS has positive valuation, then the linear independence of the ω_i mod $\mathfrak{p}_L$ means that all the a_{ij}/a_{1j} have positive valuation which is a contradiction (take $i = 1$).

So in the big sum, each term $\sum_i a_{ij}\omega_i\pi_L^j$ has valuation congruent to j mod e if it is finite. In particular, no two summands have the same valuation, so the sum cannot be zero unless all the a_{ij} were zero in the first place.

It remains to show that the basis elements span $\mathcal{O}_L$. To do this we will construct sequences $a_{ij}^{(t)} \in \mathcal{O}_K$ such that

$$\sum_{i,j} a_{ij}^{(t)} \omega_i \pi_L^j \to x$$

as $t \to \infty$. Take $a_{ij}^{(0)} = 0$. Let

$$v_L(x - \sum a_{ij}\omega_i\pi_L^j) = j_0 + es.$$

Then take $a_{ij}^{(t+1)} = a_{ij}^{(t)}$ for $j \neq j_0$ and take $a_{ij_0}^{(t)}$ so that

$$\frac{x - \sum a_{ij}^{(t)}\omega_i\pi_L^j}{\pi_L^{j_0+es}} \cong \sum_i \frac{a_{ij_0}^{(t+1)}}{\pi_K^s}\omega_i\frac{\pi_L^{j_0}}{\pi_L^{j_0}} \mod \mathfrak{p}_L,$$

so

$$v_L(x - \sum a_{ij}^{(t+1)}\omega_i\pi_L^j) \geq v_L(x - \sum a_{ij}^{(t)}\omega_i\pi_L^j) + 1.$$

□

Theorem 8.2

The local fields are the following:

- The finite extensions of $\mathbf{Q}_p$
- $\mathbf{F}_q((T))$ the field of Laurent series over $\mathbf{F}_q$ in a single variable.

Proof. Let K be a local field of residue field $\kappa_K = \mathbf{F}_q$. We will separate between the case of zero and nonzero characteristic. Suppose K has characteristic zero, so that $\mathbf{Q} \subseteq K$. Then $p \in K$ but p is zero in the residue field, so $v(p) > 0$ and hence v is a multiple of the p-adic valuation on $\mathbf{Q}$. Since K is complete, this means $\mathbf{Q}_p \subseteq K$. Since the residue field of K is finite and v is discrete on K, the previous lemma tells us that $[K : \mathbf{Q}_p]$ is finite.

If the characteristic of K is nonzero, then it must be p (since we can't have p and p' both zero in the residue field or else the residue field would have only one element). So we have $\mathbf{F}_p \subseteq K$. The polynomial $X^q - X$ has q distinct roots in $\mathbf{F}_q$. By Hensel's lemma, this polynomial also has q distinct roots in K, so in fact $\mathbf{F}_q \subseteq K$. Since every element of K can be written as a convergent Laurent series in a uniformizer whose coefficients are in a system of representatives for the residue field (which in this case can just be taken to be the copy of $\mathbf{F}_q$ inside of K), this proves that K is the set of the formal Laurent series. □

Now we will begin the study of class field theory. First, we state "class field theory for finite fields". As usual, this begins with the statement of an "Artin reciprocity law". Let $k = \mathbf{F}_q$. The injective homomorphism $\theta_k : \mathbf{Z} \to \mathrm{Gal}(\overline{\mathbf{F}}_q/\mathbf{F}_q)$ which takes 1 to the Frobenius map has dense image. That is because $\mathbf{Z}$ is dense in $\hat{\mathbf{Z}}$. The finite extensions ℓ/k correspond to finite index subgroups $U = n\mathbf{Z}$ of $\mathbf{Z}$, by taking $U = \theta_k^{-1}(\mathrm{Gal}(\overline{\mathbf{F}}_q/\ell))$. So the Frobenius map gives us an isomorphism

$$\theta_{\ell/k} : \mathbf{Z}/U \to \mathrm{Gal}(\ell/k).$$

There is a similar isomorphism called the "local reciprocity law",

Theorem 8.3

Let K be a local field. There is an injective homomorphism $\theta_K : K^\times \to \mathrm{Gal}(K^{ab}/K)$ with dense image. The finite abelian extensions of K correspond to open subgroups U of $K^\times$ of finite index via $U = \theta^{-1}(\mathrm{Gal}(K^{ab}/L))$, and

$$\theta_{L/K} : K^\times/U \to \mathrm{Gal}(L/K).$$

The key fact is that the kernel of $\theta_{L/K}$ is actually $N_{L/K}(L^\times)$.

The proof of the local reciprocity law will be one of the main goals of this class. It's also true that the local and finite reciprocity maps are compatible, in the following way: Also,

$$\begin{tikzcd}
K^\times \arrow[r, hook, "\theta_K"] \arrow[d, two heads, "v_K"] & \mathrm{Gal}(K^{ab}/K) \arrow[d, two heads] \\
\mathbf{Z} \arrow[r, hook, "\theta"] & \mathrm{Gal}(\overline{k}/k)
\end{tikzcd}$$

the "norm limitation theorem" says that this kind of approach will only be helpful for abelian extensions. In particular, for a nonabelian extension, the norm subgroup is the same as the norm subgroup of the maximal abelian subextension.

Let K be a local field with residue field κ, and K^{ur} the maximal unramified extension of K (it turns out that this is separable). Then this is the fixed field of the inertia subgroup $I(\mathfrak{p}_{(}K^{sep}/\mathfrak{p}_K)$, so

$$\mathrm{Gal}(K^{ur}/K) \cong \mathrm{Gal}(\kappa_{K^{sep}}/\kappa_K).$$

The fact that X^{q^f} is separable means that it splits completely in K^{sep}, and thus $\mathbf{F}_{q^f} \subseteq \kappa_{K^{sep}}$, which means $\kappa_{K^{sep}} = \overline{\mathbf{F}}_q$. So $\mathrm{Gal}(K^{sep}/K) \cong \mathrm{Gal}(\overline{\mathbf{F}}_q/\mathbf{F}) = \widehat{\mathbf{Z}}$ (NB this implies K^{ur}/K is abelian). By Hensel's lemma, $X^{q^f} - X$ has q^f roots in K^{ur}. So actually the maximal unramified extension of K is easy to construct: it is equal to

$$K^{ur} = \bigcup_{f \geq 1} K(\zeta_{q^f} - 1) = \bigcup_{\gcd(m,q)=1} K(\zeta_m).$$

Under the isomorphism $\mathrm{Gal}(K^{ab}/K) \cong \mathcal{O}_K \times \widehat{\mathbf{Z}}$, the inertia subgroup corresponds to $\mathcal{O}_K^\times$. In particular, the field fixed by H is unramified if and only if H contains $\mathcal{O}_K^\times$, and we have an equivalence of categories between the unramified extensions of K and the finite field extensions of the residue field of K.

§9 October 1, 2019

First we corrected some mistakes from previous lectures. The first was about the tensor product $\mathcal{O}_L \otimes \widehat{\mathcal{O}}_{K,\mathfrak{p}}$ only decomposing as the desired product if L/K is separable. The other was that if L/K is an extension of local fields, then we can't always write $\pi_L^e = \pi_K$ since π_L^e won't necessarily be in K.

Last, time we showed that $\mathrm{Gal}(K^{ur}/K) \cong \widehat{\mathbf{Z}}$. This time, we will prove $I(K^{ab}/K) \cong \mathcal{O}_K^\times$. Consider a totally ramified finite extension L/K of local fields. Then from our previous description specialized to the case $f = 1$, we know that

$$1, \pi_L, \ldots, \pi^{e-1}$$

is an $\mathcal{O}_K$-basis for $\mathcal{O}_L$.

Lemma 9.1

The minimal polynomial of π_L is an Eisenstein polynomial

Proof. The degree of the minimal polynomial is e, and $v_K(\pi_L) = 1/e$, so by the theory of Newton polygons the minimal polynomial is Eisenstein. □

Conversely, if α is a root of an Eisenstein polynomial then $K(\alpha)/K$ is totally ramified by the same type of argument.

According to the local Artin reciprocity law, the open subgroups $H \subseteq \mathrm{Gal}(K^{ab}/K) \cong \mathcal{O}_K^\times \times \widehat{Z}$ correspond to the abelian extensions of K, and the ones corresponding to the totally ramified ones are those for which $H \to \widehat{\mathbf{Z}}$ is surjective.

WARNING: the compositum of the totally ramified abelian extensions of K need not be totally ramified. For example, $X^2 - \pi_K$ is Eisenstein of degree 2 and so produces a totally ramified extension of degree 2. So is $X^2 - t\pi_K$. But the compositum contains $\sqrt{t}$, which means (if t is not a perfect square mod $|\kappa(K)|$) the residue field has to increase in size and hence the compositum is not totally ramified. So, despite the fact that we like to think of THE maximal unramified (abelian) extension, we will never have a canonical choice of maximal totally ramified extension; though it will be useful in local class field theory to make a specific choice of maximal totally ramified extension.

The analysis of where $I(K^{ab}/K)$ is mapped to to under the Artin map leads to the theory of **higher ramification groups**. We want to show that it is mapped to by $\mathcal{O}_K^\times$. Remember we had a filtration of $\mathcal{O}_K^\times$

$$\mathcal{O}_K^\times \supseteq U_K^{(1)} \supseteq \cdots$$

of higher unit groups. The corresponding subgroups of the Galois group are called the higher ramification groups:

Definition 9.2. Let $\mathcal{O}_K$ be a Dedekind domain with field of fractions K, L/K a finite Galois extension, and $\mathcal{O}_L$ the integral closure of $\mathcal{O}_K$ in L. The s-th **ramification group** of a prime $\mathfrak{P} \subseteq \mathcal{O}_L$ is defined to be

$$I_s(\mathfrak{P}|\mathfrak{p}) := \{\sigma \in D(\mathfrak{P}|\mathfrak{p}) : \sigma(a) \equiv a \mod \mathfrak{P}^{s+1} \text{ for all } a \in \mathcal{O}_L\},$$

in other words it is the set of all $\sigma \in D(\mathfrak{P}|\mathfrak{p})$ such that $i_{L/K}(\sigma) \geq s+1$, where $i_{L/K}(\sigma) := \min_{a \in \mathcal{O}_L}(v_{\mathfrak{P}}(\sigma(a) - a))$.

Example 9.3

I_{-1} is the decomposition group, and I_0 is the inertia subgroup.

The ramification groups give a filtration of the Decomposition group, and they are trivial for sufficiently large s since L/K is finite and $i_{L/K}$ is finite for nontrivial elements. Moreover, all of the ramification groups are normal subgroups of the decomposition group.

It is also enough just to look at a set of generators of $\mathcal{O}_L$ as an $\mathcal{O}_K$-algebra. In the case where L/K are local fields, there will always be a single such generator.

Let F/K be a subextension of L. Then $\operatorname{Gal}(L/F) \subseteq \operatorname{Gal}(L/K)$ and

$$I_s(\mathfrak{P}|\mathfrak{P} \cap F) = I_s(\mathfrak{P}|\mathfrak{p}) \cap \operatorname{Gal}(L/F).$$

We'd like this to also be compatible with taking quotients. The problem is that the valuation gets rescaled when we go downstairs, so we can't just do it directly. Instead we need to renormalize the numbering, to obtain the **upper numbering** of the ramification groups.

Lemma 9.4

Let L/K be a finite Galois extension of local fields. Then

$$I_s(L/K) := \{\sigma \in I(L/K) : \sigma(\pi_L) \equiv \pi_L \mod \mathfrak{P}^{s+1}\}$$

i.e.

$$\{\sigma \in I(L/K) : \sigma(\pi_L)/\pi_L \in U_L^{(s)}\}$$

Proof. Let F be the maximal unramified subextension. Then (since it is the fixed field of the inertia group)

$$I_s(L/K) = I_s(L/F)$$

and since L/F is totally ramified, we have $\mathcal{O}_L = \mathcal{O}_F[\pi_L]$, so this is what we want it to be by the above remark about generators. □

Suppose again that $\mathcal{O}_K$ and $\mathcal{O}_L$ are Dedekind domains as before. Then $D(\mathfrak{P}|\mathfrak{p}) = \operatorname{Gal}(L_{\mathfrak{P}}/K_{\mathfrak{p}})$ and the ramification groups of $\mathfrak{P}|\mathfrak{p}$ are the same as the $I_s(L_{\mathfrak{P}}/K_{\mathfrak{p}})$. This is a consequence of the definitions (taking completions doesn't change the value groups or residue fields).

Theorem 9.5

Let L/K be a finite Galois extension of local fields. We obtain injective group homomorphisms $I/I_1 \to \mathcal{O}_L^\times/U_L^{(1)}$ and $I_s/I_{s+1} \to U_L^{(s)}/U_L^{(s+1)}$.

Proof. The homomorphisms are constructed by taking

$$\sigma I_{s+1} \mapsto \frac{\sigma(\pi_L)}{\pi_L} U_L^{(s+1)}.$$

First, this homomorphism does not depend on the choice of uniformizer, because if u is a unit in $\mathcal{O}_L$ and $\sigma \in I_s$ then $v_L(\sigma(u))/u \in U_L^{(s+1)}$. Checking these are injective group homomorphisms is straightforward from the same ideas. □

$D/I = \mathrm{Gal}(\kappa_L/\kappa_K) = \mathbf{Z}/f\mathbf{Z}$, and I/I_1 embeds into $\mathcal{O}_L^\times/U_L^{(1)} = \kappa_L^\times$, while I_1/I_2 embeds into κ_L. If p is the residue characteristic, then this means $|I_1/I_2|$ is a power of p, and $|I/I_1|$ is coprime to p. The fact that these quotients are abelian (they inject into abelian things) means that G is solvable. And I_1 is the (unique) p-Sylow subgroup of I_0.

It's another useful fact about local fields (useful in our case for talking about ramification groups) that the ring extension is always monogenic

Lemma 9.6

If L/K is a finite extension of local fields, then $\mathcal{O}_L = \mathcal{O}_K[\alpha]$ for some α.

Proof. Let $\mathbf{F}_{q^n}/\mathbf{F}_q$ be the residue field extension. By Hensel's lemma, $\zeta_{q^n-1} \in \mathcal{O}_L$. If this is not already a generator, just consider

$$\alpha = \zeta_{q^n-1} + \pi_L.$$

Then mod $\mathfrak{p}_L^2$,

$$\begin{aligned}\alpha^{q^n} - \alpha &= (\zeta + \pi)^{q^n} - (\zeta + \pi)\\ &= \zeta^{q^n} - \zeta - \pi\\ &= -\pi\end{aligned}$$

so $v_L(\alpha^{q^n} - \alpha) = 1$ and therefore we can get a uniformizer from α. This uniformizer plus ζ_{q^n-1} generate $\mathcal{O}_L$ as an algebra, so we are done. □

§10 October 3, 2019

The higher ramification groups of an extension of local fields actually are enough to determine the different ideal of the extension. We will state this in terms of the discriminant:

Theorem 10.1

Let L/K be a finite Galois extension of local fields. Then

$$v_K(d_{L/K}) = f(L/K) \cdot \sum_{\sigma \neq 1 \in \mathrm{Gal}(L/K)} i_{L/K}(\sigma) = f(L/K) \sum_{s=0}^{\infty} (|I_s| - 1).$$

Proof. We've shown that actually $\mathcal{O}_L$ has a power basis over $\mathcal{O}_K$, namely $1, \alpha, \alpha^2, \ldots, \alpha^{[L:K]-1}$, and by the usual Vandermonde matrix computation, the discriminant of the extension is the ideal generated by $N_{L/K}(f'(\alpha))$ where f is the minimal polynomial of α. In particular,

$$\begin{aligned} v_K(d_{L/K}) &= v_K(\prod_{\sigma,\tau}(\sigma\alpha - \tau\alpha)) \\ &= v_K(\prod_{\sigma} \sigma \prod_{\tau}(\alpha - \tau\alpha)) \\ &= [L:K] \sum_{\tau \neq 1} v_K(\alpha - \tau\alpha) \end{aligned}$$

so by renormalizing we have

$$f(L/K) \sum_{\tau \neq 1} v_L(\alpha - \tau\alpha)$$

as claimed. □

Example 10.2

Consider the extension $\mathbf{Q}_p(\sqrt{p})/\mathbf{Q}_p$, whose ring of integers is $\mathbf{Z}_p[\sqrt{p}]$ since 2 is invertible in $\mathbf{Z}_p$ when p is odd. Then $\mathrm{Gal}(\mathbf{Q}_p(\sqrt{p})/\mathbf{Q}_p) = \mathbf{Z}/2\mathbf{Z}$, and if σ is the nontrivial element of the Galois group then we have

$$i_{L/K}(\sigma) = v_{\mathbf{Q}_p(\sqrt{p})}(-\sqrt{p} - \sqrt{p}) = 1.$$

So I_0 is the entire Galois group and the rest of the ramification groups are trivial.

Example 10.3

Consider $\mathbf{Q}_2(\sqrt{p})/\mathbf{Q}_2$, where p is 3 mod 4. Then $i(\sigma)$ is the valuation of $-2\sqrt{p}$ which is 2. So the first two ramification groups are $\mathbf{Z}/2$ and the rest are trivial.

Example 10.4

Take $\mathbf{Q}_2(\sqrt{2})/\mathbf{Q}_2$. Then $i(\sigma) = v(-2\sqrt{2}) = 3$.

All of these objects are compatible with the corresponding ones for global fields (compare the valuations we just computed with the discriminants of the quadratic number fields). What we actually care about is the cyclotomic extensions, since we want to study abelian extensions of local fields (in this case by local Kronecker-Weber).

Example 10.5

The Galois group of $K_n = \mathbf{Q}_p(\zeta_{p^n})/\mathbf{Q}_p$. We know that $(\zeta_{p^n} - 1)$ is the only prime in $\mathbf{Q}(\zeta_{p^n})$. So

$$\mathrm{Gal}(K_n/\mathbf{Q}_p) = \mathrm{Gal}(\mathbf{Q}(\zeta_{p^n})/\mathbf{Q}) = (\mathbf{Z}/p^n\mathbf{Z})^\times$$

and therefore

$$\mathrm{Gal}(K_n/K_m) = \{r \in (\mathbf{Z}/p^n\mathbf{Z})^\times : r \equiv 1 \mod p^m\}.$$

Let $1 \neq r \in (\mathbf{Z}/p^n\mathbf{Z})^\times$. Then

$$i(K_n/\mathbf{Q}_p)(\phi_r) = v_{K_n}(\zeta_{p^n}^r - \zeta_{p^n}) = v_{K_n}(\zeta_{p^n}^{r-1} - 1) = v_{K_n}(\phi_r(\zeta_{p^n}^{p^t} - 1))$$

where $r - 1 = p^t k$ (take $t = v_p(r-1)$), which is just

$$v_{K_n}(\zeta_{p^{n-t}} - 1) = e(K_n/K_{n-t}) v_{K_{n-t}}(\zeta_{p^{n-t}} - 1) = p^t.$$

This shows that $I_s(K_n/\mathbf{Q}_p) = \mathrm{Gal}(K_n/K_t)$, where $t \geq 0$ is the minimal integer $\leq n$ such that $s \leq p^t - 1$ or $t = n$.

To convert from the ramification groups and those of a subextension, there is a whole technical setup involving the following function, used to convert between the upper and lower indexing schemes.

Definition 10.6. Let $\eta_{L/K}(s) = \int_0^s \frac{dx}{[I_0 : I_x]} = \frac{1}{e(L/K)} \sum_\sigma \min(i_{L/K}(\sigma), s+1)$ (you can check these are the same by looking at the derivative).

The y-coordinate of the places where $\eta_{L/K}$ bends are all integers, as long as L/K is abelian (this is called the Hase-Liu-Arf theorem).

Definition 10.7. Let $I^t(L/K) = I_{\eta_{L/K}^{-1}(t)}(L/K)$.

Theorem 10.8 (Herbrand's Theorem)

Let $K \subseteq L' \subseteq L$ be a tower of Galois extensions, where $G = \mathrm{Gal}(L/K)$, $H = \mathrm{Gal}(L/L')$, and $G/H = \mathrm{Gal}(L'/L)$. Then $I^t(L'/K)$ is the image of $I^t(L/K)$ under the projection $G \to G/H$.

Proof. We'll need some technical lemmas first.

Lemma 10.9

For any $\sigma \in G$, $i_{L'/K}(\sigma|_{L'}) = \frac{1}{e(L/L')} \sum_{\tau \in H} i_{L/K}(\sigma\tau)$.

Proof. If $\sigma|_{L'} = \mathrm{id}$, then both sides are ∞. Now suppose it is not id. Then for a choice of generator $\alpha_{L/K}$,

$$i_{L/K}(\sigma\tau) = v_L(\sigma\tau\alpha_{L/K} - \alpha_{L/K})$$

and

$$i_{L'/K}(\sigma|_{L'}) = v_L(\sigma\alpha_{L'/K} - \alpha_{L'/K})$$

so we need to show that

$$v_L(\sigma\alpha_{L'/K} - \alpha_{L'/K}) = \sum_{\tau\in H} v_L(\sigma\tau\alpha_{L/K} - \alpha_{L/K}) = v_L(\prod_{\tau\in H}(\sigma\tau\alpha_{L/K} - \alpha_{L/K})).$$

Let the thing inside the valuation on the LHS be a, and the thing inside the valuation on the RHS be b. We will first show that $v_L(a) \le v_L(b)$. Let $f(X) = \prod_{\tau\in H}(X - \tau\alpha_{L/K}) \in \mathcal{O}_{L'}[X]$ be the minimal polynomial of $\alpha_{L/K}$ over L'. Then

$$(\sigma f)(X) = \prod_{\tau\in H}(X - \sigma\tau\alpha_{L/K})$$

and thus

$$(\sigma f - f)(\alpha_{L/K}) = (\sigma f)(\alpha_{L/K}) - f(\alpha_{L/K})$$

which is just $\pm b$ since $f(\alpha_{L/K}) = 0$. The coefficients of $\sigma f - f$ are of the form $\sigma c - c$ for $c \in \mathcal{O}_{L'}$, so they all have valuation at least $v_L(a)$. As a result, $v_L(b) \ge v_L(a)$. Now we show the opposite inequality. The fact that $\alpha_{L/K}$ is a generator for $\mathcal{O}_L$ as an $\mathcal{O}_K$-algebra means that $\alpha_{L'/K} = g(\alpha_{L/K})$ for some $g \in \mathcal{O}_K[X]$. So $\alpha_{L/K}$ is a root of the polynomial $g(X) - \alpha_{L'/K} \in \mathcal{O}_{L'}[X]$, which means

$$f(X) | (g(X) - \alpha_{L'/K})$$

in $\mathcal{O}_{L'}[X]$. After applying σ and plugging in $\alpha_{L/K}$, we get

$$(\sigma f)(\alpha_{L/K}) | (\sigma g)(\alpha_{L/K}) - \sigma\alpha_{L'/K}.$$

The LHS is just $\pm b$ and the RHS is $\alpha_{L'/K} - \sigma\alpha_{L'/K} = -a$, which proves that $v_L(b) \le v_L(a)$. □

□

Corollary 10.10

Even for infinite extensions, we can define

$$I^t(L/K) = \{\sigma \in \mathrm{Gal}(L/K) : \forall L \supseteq L'/K \text{ finite Galois subextension }, \sigma|_{L'} \in I^t(L'/K)\}.$$

Example 10.11

$\mathrm{Gal}(K_\infty/\mathbf{Q}_p) = \mathbf{Z}_p^\times$, and in upper numbering the t-th higher ramification group of this extension is just the subset of elements of $\mathbf{Z}_p^\times$ congruent to 1 mod (p^t), i.e. the higher unit group $U^{(t)}$. This suggests that there is some general framework in which the higher unit groups correspond via a natural isomorphism (the Artin map!!!) with the higher ramification groups.

§11 October 8, 2019

Lemma 11.1

$i_{L'/K}(\sigma|_{L'}) = \frac{1}{e(L/L')}\sum_{\tau\in\mathrm{Gal}(L/L')} i_{L/K}(\sigma\tau)$.

Lemma 11.2

$I_{\eta_{L/L'}(s)}(L'/K)$ is the image of $I_s(L/K)$ under the quotient $\mathrm{Gal}(L/K)\to\mathrm{Gal}(L'/K)$.

Proof. We need to check that for any $\sigma\in\mathrm{Gal}(L/K)$,

$$i_{L/K}(\sigma|_{L'}) - 1 = \eta_{L/L'}(\max_{\tau\in H}\{i_{L/K}(\sigma\tau)-1\}).$$

The left hand side is

$$\max\{s|\sigma|_{L'}\in I_s(L'/K)\}$$

and the maximum in the right hand side is

$$\max\{s:\sigma\tau\in I_s(L/K)\}.$$

WLOG we can assume the maximum occurs at $\tau=\mathrm{id}$, and thus (since I_s is a group),

$$i_{L/K}(\sigma\tau)=\min(i_{L/K}(\sigma), i_{L/K}(\tau)).$$

By the previous lemma, we have

$$\begin{aligned} i_{L'/K}(\sigma|_{L'/K}) &= \frac{1}{e(L/L')}\sum_{\tau\in H} i_{L/K}(\sigma\tau)\\ &= \frac{1}{e(L/L')}\sum_{\tau\in H}\min(i_{L/K}(\sigma), i_{L/K}(\tau))\\ &= \eta_{L/L'}(i_{L/K}(\sigma)-1)\end{aligned}$$

By the characterization of η from last class. □

This almost proves Herbrand's theorem but we need one more lemma.

Lemma 11.3

$\eta_{L/K}=\eta_{L'/K}\circ\eta_{L/L'}$

Proof. The kernel of $I_s(L/K)\to I_{\eta_{L/L'}(s)}(L'/K)$ is $I_s(L/K)\cap\mathrm{Gal}(L/L')=I_s(L/L')$. So

$$|I_s(L/K)| = |I_{\eta_{L/L'}(s)}(L'/K)|\cdot|I_s(L/L')|.$$

Since $I_0(L/K)$ is just the intertia group, we also know that

$$|I_0(L/K)| = e(L/K) = e(L'/K)e(L/L') = |I_0(L'/K)|\cdot|I_0(L/L')|.$$

So using the definition of η, we get

$$\eta'_{L/K}(s)=\frac{1}{[I_0:I_s(L/K)]}=\frac{1}{[I_0:I_{\eta_{L/L'}(s)}(L'/K)]}\frac{1}{[I_0:I_s(L/L')]}=\eta'_{L'/K}(\eta_{L/L'}(s))\eta'_{L/L'}(s)$$

which is what we want by the chain rule (checking at zero and integrating gives the desired result) □

Proof of Herbrand's Theorem. From the previous two lemmas, we can just see that

$$I^t(L'/K) = I_{\eta(L'/K)^{-1}(t)}(L/L)$$

is the image of

$$I_{\eta_{L/L'}^{-1}(\eta_{L'/K}^{-1}(t))} = I_{\eta_{L/K}^{-1}(t)}(L/K) = I^t(L/K)$$

as desired. □

Now a few words about the Hasse–Arf theorem.

Theorem 11.4 (Hasse–Arf)

If L/K is an abelian extension of local fields, the corners of the graph of $\eta_{L/K}$ have integer coordinates.

More generally, if L/K is nonabelian, consider the class function

$$a_{L/K} : G \to \mathbf{Z}$$

which takes $\sigma \neq \text{id}$ to $-f(L/K)i_{L/K}(\sigma)$ and the identity to $f(L/K)\sum_{\tau\neq\text{id}} i_{L/K}(\tau)$. This is called the *Artin representation* of L/K, and Artin showed that it is the character of a representation of G over $\mathbf{C}$. The number of times the trivial representation shows up in the Artin representation is

$$\langle a_{L/K}, \text{triv}\rangle = \frac{1}{|G|}\sum_{\sigma\in G} a_{L/K}(\sigma) = 0.$$

We can also compute that

$$\langle a_{L/K}, \text{reg}_G\rangle = v_K(d_{L/K})$$

from the formula involving higher ramification groups and the valuation of the discriminant.

Lemma 11.5

$$a_{L/K} = \sum_{s=0}^{\infty} \frac{1}{[I_0 : I_s]}(\text{reg}_G - \text{Ind}_{I_s}^G \text{triv}_{I_s}),$$

where the induced character of the trivial character is (since I_s is normal) $[G : I_s]$ when $\sigma \in I_s$ and 0 otherwise. In upper numbering, we have

$$a_{L/K} = \int_{-1}^{\infty} (\text{reg}_G - \text{Ind}_{I^t}^G \text{triv}_{I^t})\, dt.$$

Using this one can deduce the Hasse-Arf theorem. There is also a relationship in the other direction using Brauer's theorem from representation theory.

§11.1 Back to class field theory

In local CFT, so far we have computed $\text{Gal}(K^{ur}/K) \cong \widehat{\mathbf{Z}}$. The kernel of the projection $\text{Gal}(K^{ab}/K) \to \text{Gal}(K^{ur}/K)$ is $I(K^{ab}/K)$. We want to construct an Artin reciprocity isomorphism from $\mathcal{O}_K^\times$ to $I(K^{ab}/K)$ which takes the higher unit groups to the higher ramification groups.

Definition 11.6. A Galois extension L/K of local fields is **tamely ramified** if $I_1(L/K) = 1$.

If L/K is finite, then L/K is tamely ramified iff $e(L/K)$ is not divisible by the characteristic of the residue field, p. That is because I_1 is the p-Sylow subgroup of I. In particular, if p does not divide $[L:K]$ then L/K is tamely ramified. The maximal tamely ramified subextension of L/K is the fixed field of I_1 (the lower index is not a typo). The lower index is okay because for any $L \supseteq L' \supseteq K$ with L'/K Galois, we have

$$I_1(L'/K) = I_\epsilon(L'/K) = I^{\epsilon'}(L'/K)$$

which is the image under the quotient of $I^{\epsilon'}(L/K)$, which is the same as the image of $I_{\epsilon''}$. If ϵ is made small enough, all of these will be less than 1 so we are good.

Theorem 11.7

The maximal tamely ramified extension of a local field K with residue field $\mathbf{F}_q$ is

$$K^{\text{tame}} = \bigcup_{\substack{m \geq 1 \\ (m,q)=1}} K^{\text{ur}}(\pi_K^{1/m}).$$

For any $\alpha \in K^{tame}$, $X^m - \alpha$ has m distinct roots in K^{tame}. In particular, write $\alpha = \beta \cdot \pi_K^C$ with $\beta \in \mathcal{O}_K^\times$. Then by taking derivatives we see that $X^m - \beta$ has m distinct roots in the residue field of K^{tame}, and therefore in K^{tame} as well by Hensel's lemma.

§12 October 10

First, a warning. If K is a local field, then K^{ur} is an infinite extension of K and is not complete with respect to the unique extension of the valuation of K. So a priori we cannot apply Hensel's lemma to it. On the other hand, we CAN apply Hensel's lemma because Hensel's lemma only involves finitely many elements so you can restrict to a finite Galois subextension and apply it there.

Theorem 12.1

Let K be a local field of residue field $\mathbf{F}_q$. Then the maximal tamely ramified extension of K is

$$K^{tame} = \bigcup_{m\geq 1,(m,q)=1} K(\zeta_m, \pi_K^{1/m}) = \bigcup_{t\geq 1} K(\zeta_{q^t-1}, \pi_K^{1/(q^t-1)}).$$

Proof. First we need to show that this is actually tamely ramified. To do this it suffices to show it for each of the fields whose union is being taken (by the properties of the upper numbering). Notice that the tower $K \subseteq K(\zeta_m) \subseteq K(\zeta_m, \pi_K^{1/m})$ has first step which is unramified, and second step totally ramified. The ramification index is $e = m$ so in fact the second step is tamely ramified and we are done.

Next we need to show that if L/K is a finite Galois tamely ramified extension, then $L \subseteq K(\zeta_m, \pi_K^{1/m})$ for some m. Splitting L into its unramified part $L' = L \cap K^{ur} \subseteq K(\zeta_m)$ and its totally ramified part L/L'. We know from before that $L = L'(\pi_L)$. Let $f(X) \in \mathcal{O}_{L'}[X]$ be the (Eisenstein) minimal polynomial of π_L over L'. It suffices to show that $f(X)$ splits completely in $K^{ur}(\pi_K^{1/e})$. Recall that (i.e. f is Eisenstein) f has very bad Newton polygon (it reduces mod p to something with a massive root). However, the roots r_i of f have valuation $1/e$ (since π_L^e has the same valuation as π_K). So instead we take $r_i/\pi_K^{1/e}$, which are the roots of $g(X) = f(\pi_K^{1/e}X)$. The Newton polygon of g is now just flat, and all the coefficients have valuation at least one. So divide by π_K to get that $\frac{1}{\pi_K} f(\pi_K^{1/e}X)$ has Newton polygon which sits on the x-axis. By construction,

$$g(X) \in \mathcal{O}_{K^{ur}(\pi_K^{1/e})}[X]$$

and

$$g(X) \equiv X^e + c \mod \mathfrak{p}$$

which means g actually has e simple roots mod $\mathfrak{p}$, and thus by Hensel's lemma (using the fact that $(e,q) = 1$) these roots lift to e distinct roots of g. It follows that our original polynomial f splits completely in $K^{ur}(\pi_K^{1/e})$, as desired [because this implies that $\pi_L \in K^{ur}(\pi_K^{1/e})$]. □

Now for the Galois group. Recall that the maximal unramified extension of K has Galois group $\widehat{Z}$ which is the closure of the copy of $\mathbf{Z}$ generated by the Frobenius. Moreover,

$$\operatorname{Gal}(K^{ur}(\pi_K^{1/m})/K^{ur}) \cong \mathbf{Z}/m\mathbf{Z}$$

where this is just generated by the automorphism sending $\pi_K^{1/m} \mapsto \zeta_m \pi_K^{1/m}$ (this is fine because K^{ur} contains ζ_m). As a result,

$$\operatorname{Gal}(K^{tame}/K^{ur}) = \varprojlim_{(m,q)=1} \operatorname{Gal}(K^{ur}(\pi_K^{1/m})/K^{ur}) = \varprojlim \mathbf{Z}/m\mathbf{Z}$$

and you can check what the restrictions have to look like, but this will be the same as (via the Chinese remainder theorem)

$$\prod_{\ell \neq p} \mathbf{Z}_\ell.$$

It is topologically generated by τ, the automorphism that corresponds to 1 in each copy of $\mathbf{Z}/m\mathbf{Z}$. You can also lift the Frobenius to K^{tame} by taking

$$\phi_q(\pi_K^{1/m}) = \pi_K^{1/m}$$

and

$$\phi_q(\zeta_m) = \zeta_m^q,$$

making sure to choose the m-th roots in a way in which this is actually well-defined. In fact, $\mathrm{Gal}(K^{tame}/K)$ is topologically generated by τ and ϕ_q in the sense that

$$\mathrm{Gal}(K^{tame}/K) = \overline{\langle\tau\rangle} \rtimes \overline{\langle\phi_q\rangle}$$

where of course $\phi_q \tau \phi_q^{-1} = \tau^q$.

Example 12.2

Now let's look at a finite extension. Take $L = K(\zeta_{q^t-1}, \pi_K^{1/(q^t-1)})$. Our arguments show that

$$\mathrm{Gal}(L/K) = \mathbf{Z}/(q^t - 1) \rtimes \mathbf{Z}/t.$$

The open subgroups of $\mathrm{Gal}(K^{tame}/K)$ correspond to the finite subextensions and are the same as the finite index subgroups. In particular, $\mathrm{Gal}(K^{tame}/K)^n$ has finite index in $\mathrm{Gal}(K^{tame}/K)$ and so K has only finitely many tamely ramified Galois extensions of degree n. On the homework we showed that in characteristic zero there are finitely many (arbitrarily ramified) Galois extensions of given degree. But this is false in characteristic p:

Example 12.3

$K = \mathbf{F}_{p^k}((T))$ has infinitely many Galois extensions with Galois group $\mathbf{Z}/p\mathbf{Z}$.

Theorem 12.4

The maximal tamely ramified abelian extension of K is

$$K^{tame,ab} = K^{ur}(\pi_K^{1/(q-1)})$$

and its Galois group is $\mathbf{Z}/(q-1)\mathbf{Z} \times \widehat{\mathbf{Z}}$.

Proof. To do this we can just mod out by the commutator to compute the abelianization of the Galois group. Or just look explicitly at the presentation for the Galois group. In particular, take

$$\overline{\tau} = \tau|_{K^{ur}(\pi_K^{1/(q-1)})}$$

and

$$\overline{\phi}_q = \phi_q|_{K^{ur}(\pi_K^{1/(q-1)})}.$$

These commute because conjugating $\overline{\tau}$ by $\overline{\phi}_q$ gives you $\overline{\tau}^q$ which is the same as $\overline{\tau}$. So at least $K^{tame,ab}$ is abelian. The other direction is almost the same. Let L/K be a finite tamely ramified abelian extension. Then defining $\overline{\tau}$ and $\overline{\phi}_q$ just by restriction to L, they must commute, which implies that $\overline{\tau} = \overline{\tau}^q$, i.e. $\overline{\tau}$ has order dividing $q-1$. So $K^{tame,ab}$ has degree at most $q-1$ over K^{ur}, which means it is actually equal to the thing we wrote before. □

Remark 12.5. This is expected because by Artin reciprocity $\mathcal{O}_K^\times \times \widehat{Z} = \mathrm{Gal}(K^{ab}/K)$ so we expect to get $\mathrm{Gal}(K^{ab,tame}/K)$ by modding out by $I^1(K^{ab}/K)$ which maps under the reciprocity mapping to $U_K^{(1)}$. Then recall that $\mathcal{O}_K^\times/U_K^{(1)} \cong \mathbf{F}_q^\times$.

Remark 12.6. We proved the following on the homework: if $f(X) \in \mathcal{O}_K[X]$ is Eisenstein of degree $q-1$ and $\alpha \in \overline{K}$ is a root of f, then $K^{ur}(\alpha) = K^{tame,ab}$.

§13 October 15, 2019

If L/K is a finite Galois extension of local fields with residue field extension $\mathbf{F}_{q^n}/\mathbf{F}_q$, then we have canonical embeddings $I/I_1 \to \mathcal{O}_L^\times/U_L^{(1)} \cong \mathbf{F}_{q^n}^\times$ and $I_n/I_{n+1} \to U_L^{(n)}/U_L^{(n+1)} \cong \mathbf{F}_{q^n}$, which are given by $\sigma \mapsto \sigma(\pi_L)/\pi_L$ (and are independent of the choice of uniformizer).

Lemma 13.1

If L/K is abelian, then we have embeddings $I/I_1 \to \mathbf{F}_q^\times$ and $I_n/I_{n+1} \to \mathbf{F}_q$.

Fake proof. Let $\tilde{\varphi}_q \in \operatorname{Gal}(L/K)$ be a lift of the Frobenius automorphism. For any $\sigma \in I_n$, we have (mod $U_L^{(n+1)}$)

$$\tilde{\varphi}_q(\sigma(\pi_L)/\pi_L) = \frac{\tilde{\varphi}_q\sigma(\pi_L)}{\tilde{\varphi}_q(\pi_L)} = \frac{\sigma\tilde{\varphi}_q(\pi_L)}{\tilde{\varphi}_q(\pi_L)}$$

which has to be the same mod $U_L^{(n+1)}$ as the ratio of the q-th powers.

This proof is fake because the isomorphism with $\mathbf{F}_{q^n}$ is noncanonical. So from this proof we only conclude that $I/I_1 \to \mathbf{F}_q^\times$. □

Legit proof. For $n \geq 1$, if $\sigma \in I_n$ we can write $\sigma(\pi_L)/\pi_L = 1 + \pi_L^n x$, and

$$\begin{aligned} \tilde{\varphi}_q(1+\pi_L x) &\equiv 1 + \pi_L^n x \mod U_L^{(n+1)} \\ \tilde{\varphi}_q(\pi_L)^n \tilde{\varphi}_q(x) &\equiv \pi_L^n x \mod \mathfrak{p}_L^{n+1} \\ \frac{\tilde{\varphi}_q(\pi_L)}{\pi_L^n}\tilde{\varphi}_q(x) &\equiv x \mod \mathfrak{p}_L \\ \frac{\tilde{\varphi}_q(\pi_L)^n}{\pi_L^n} x^q &\equiv x \mod \mathfrak{p}_L. \end{aligned}$$

This means x satisfies this particular polynomial over $\mathbf{F}_{q^n}$, which has at most q roots. So, the image of $I_n/I_{n+1} \to \mathbf{F}_{q^n}$ has size at most q, which gives an (EXTREMELY NONCANONICAL) injection $I_n/I_{n+1} \to \mathbf{F}_q$. □

What's important is that we have a bound on the sizes of quotients of successive ramification groups.

Theorem 13.2 (Local Kronecker–Weber)

$\mathbf{Q}_p^{ab} = \mathbf{Q}_p(\zeta_\infty) = \bigcup_{m\geq 1,(p,m)=1} \mathbf{Q}_p(\zeta_m) \cdot \bigcup_{k\geq 0} \mathbf{Q}_p(\zeta_{p^k})$.

Proof assuming Hasse–Arf. Assume that $\mathbf{Q}_p^{ab}$ properly contains $\mathbf{Q}_p(\zeta_\infty)$. Then

$$H = \operatorname{Gal}(\mathbf{Q}_p^{ab}/\mathbf{Q}_p(\zeta_m))$$

is nontrivial. It's a general fact about the upper numbering for infinite extensions that

$$\bigcap_{n=0}^{\infty} I^n(\mathbf{Q}_p^{ab}/\mathbf{Q}_p) = 1$$

[because in any finite extension $K/\mathbf{Q}_p$, the upper ramification groups are eventually zero].

Let $n \geq 0$ be the least nonnegative integer such that H is not contained in $I^n(\mathbf{Q}_p^{ab}/\mathbf{Q}_p)$. If $n = 0$, then

$$(\mathbf{Q}_p^{ab})^H \not\supseteq (\mathbf{Q}_p^{ab})^{I_0}$$

where the LHS is $\mathbf{Q}_p(\zeta_\infty)$ and the RHS is an unramified extension of $\mathbf{Q}_p$. This is a contradiction. Now suppose $n \geq 1$. We have a map

$$I^{n-1}(\mathbf{Q}_p^{ab}/\mathbf{Q}_p)/I^n(\mathbf{Q}_p^{ab}/\mathbf{Q}_p) \to I^{n-1}(\mathbf{Q}_p(\zeta_\infty)/\mathbf{Q}_p)/I^n(\mathbf{Q}_p^{ab}/\mathbf{Q}_p).$$

By the minimality of n,

$$H \subseteq I^{n-1}(\mathbf{Q}_p^{ab}/\mathbf{Q}_p),$$

and the map above is given by reduction mod H, so the kernel contains $H/I^n(\mathbf{Q}_p^{ab}/\mathbf{Q}_p)$. Since H is not inside of this (by definition of n), the kernel is nontrivial, so

$$|I^{n-1}/I^n(\mathbf{Q}_p^{ab}/\mathbf{Q}_p)| > |I^{n-1}/I^n(\mathbf{Q}_p(\zeta_\infty)/\mathbf{Q}_p)|.$$

But in fact the exact opposite of this is true by the previous lemma. □

Corollary 13.3

$\mathrm{Gal}(\mathbf{Q}_p^{ab}/\mathbf{Q}_p) \cong \widehat{\mathbf{Z}} \times \mathbf{Z}_p^\times$. Also, $I^n(\mathbf{Q}_p^{ab}/\mathbf{Q}_p) = \mathrm{Gal}(\mathbf{Q}_p^{ab}/\bigcup_{\substack{m\geq 1\\ p^{m+1}\nmid m}} \mathbf{Q}_p(\zeta_m))$ [this is because of the homework problem concerning the upper numbering of the ramification groups for cyclotomic extensions of $\mathbf{Q}_p$].

More generally, let

$$K_0 \subseteq K_1 \subseteq \cdots \subseteq \cdots$$

be abelian extensions of K such that $K_0 = K^{ur}$, $I^n(K_n/K) = 1$, and $[K_{n+1} : K_n] = q - 1$ if $n = 0$ and q if $n \geq 1$. Then Hasse–Arf implies that $\bigcup_{n\geq 0} K_n$.

In Lubin–Tate theory, we will see exactly how to construct these fields to get "explicit local class field theory," by noticing that the roots of unity are the torsion points of a particular group structure upstairs, and generalizing this to the appropriate group. The method of constructing these groups is by *formal group laws*. Basically, the group law should be given by power series.

Definition 13.4. A (commutative) **formal group** over a ring R is a power series $F(X,Y) \in R[[X,Y]]$ such that

1. $F(X,Y) = F(Y,X)$ [commutativity]
2. $F(F(X,Y),Z) = F(X,F(Y,Z))$ [associativity]
3. $F(X,Y) \equiv X + Y \mod \text{degree } 2$.

It's often the case in formal group laws that other important properties (which in the usual group theory you would have to specify) follow from the very basic ones. For example, it will follow that $F(0,X) = X$.

Example 13.5

Take the additive group $\mathbf{G}_a = X + Y$.

Example 13.6

There's also the multiplicative group

$$\mathbf{G}_m = (1+X)(1+Y) = 1 + X + Y + XY.$$

This is basically saying multiplication is linear near 1.

Definition 13.7. A **homomorphism of formal groups** $f : F \to G$ is a power series $f(X) \in R[[X]]$ with $f(0) = 0$ such that

$$f(F(X,Y)) = G(f(X), f(Y)).$$

Lemma 13.8

If $f(X) = \sum_{n=1}^{\infty} a_n X^n \in R[[X]]$, then the following are equivalent:

1. There exists some $g \in R[[X]]$ such that $f(g(X)) = g(f(X)) = X$
2. $a_1 \in R^{\times}$.

Proof. A lot of these things are just bash with power series. Let $g(X) = b_0 + b_1 X + \cdots$. Then

$$X = g(f(X) = b_0 + (\text{deg } \geq 1)$$

so $b_0 = 0$ at least. Moreover,

$$X = f(g(x)) = b_0 + a_1 b_1 X + (\text{deg } \geq 1)$$

so $a_1 \in R^{\times}$. In fact, to go the other way we can directly recursively construct a right inverse one coefficient at a time.

Suppose $f(g(X)) = X$. Then $b_0 = 0$ and $b_1 = 1/a_1$. Assume we have coefficients $b_0, \ldots, b_{n-1}$ so that $f(b_1 X + \cdots + b_{n-1} X^{n-1})$ agrees with X until degree n. This will force the choice of b_n. In particular, take b_n so that

$$f(b_1 X + \cdots + b_{n-1} X^{n-1}) = X - a_1 b_n X^n + (\text{deg} \geq n+1)$$

and it works out. Similarly, there is an h such that $g(h(X)) = X$, and thus

$$g(f(X)) = g(f(g(h(X)))) = g(h(X)) = X.$$

so g is also a right inverse to f. □

Definition 13.9. A **formal module** F over R is a formal group $F \in R[[X,Y]]$ plus a ring homomorphism $R \to \operatorname{End}_R(F)$ satisfying $a \mapsto [a]_F(X)$, where

$$[a]_F(X) = aX + (\text{deg. } \geq 2)$$

for all $a \in R$, where the endomorphism ring $\operatorname{End}_R(F)$ is defined to be the send of formal group endomorphisms of F, where we need to remember that adding two endomorphisms is done using the group law F.

§14 October 17, 2019

Let K be a local field.

Definition 14.1. A **Lubin–Tate series** for a uniformizer π_K is a power series $e(X) \in \mathcal{O}_K[[X]]$ such that

1. $e(X) \equiv X^q \mod \pi_K$
2. $e(X) \equiv \pi_K X \mod \text{deg. } 2$.

Example 14.2

In reality we'll only need one such series, and the most convenient examples are polynomials. For example, $e(X) = X^q + \pi_K X$ or more generally $e(X) = Xf(X)$ for any Eisenstein polynomial f with constant coefficient π_K.

Example 14.3

When $K = \mathbf{Q}_p$, $e(X) = (X+1)^p - 1$ works. As far as I know this is the important motivating example.

Theorem 14.4

Let $e(X)$ be a Lubin-Tate series for π_K. Then there is a unique formal $\mathcal{O}_K$-module F_e, called the **Lubin-Tate module for $e(x)$**, such that $[\pi]_{F_e}(X) = e(X)$.

Example 14.5

Let $K = \mathbf{Q}_p$, $\pi = p$, $e(X) = (X+1)^p - 1$. The unique formal $\mathcal{O}_K$-module with this property is $F_e = \mathbf{G}_m$, since $[a]_{F_e}(X) = (X+1)^a - 1$, which specializes to $e(X)$ when $a = \pi$.

The key technical lemma that lets you prove things like this (especially uniqueness statements about Lubin-Tate modules or power series) is as follows:

Lemma 14.6

Let $e(X), \tilde{e}(X)$ be two Lubin–Tate series for π_K, and let $a_1, \ldots, a_r \in \mathcal{O}_K$. Then there is a unique power series $\phi(x_1, \ldots, x_r) \in \mathcal{O}_K[[x_1, \ldots, x_r]]$ such that $\phi(x_1, \ldots, x_r) = a_1 x_1 + \cdots + a_r x_r + (\text{deg.} \geq 2)$ and $e(\phi(x_1, \ldots, x_r)) = \phi(\tilde{e}(x_1), \ldots, \tilde{e}(x_r))$.

Proof. We know the linear terms of ϕ, so we can inductively construct the terms of higher degree (and this choice we will see is unique and actually works). Write $X = (x_1, \ldots, x_r)$, $e(X) = (e(x_1), \ldots, e(x_r))$, $\phi(X) = \phi(x_1, \ldots, x_r)$ to save space. We will inductively construct ϕ_n, the homogeneous degree-n part of ϕ. Suppose for the inductive hypothesis that we have constructed $\phi_1, \ldots, \phi_{n-1}$ such that

$$e(\phi_1(X) + \cdots + \phi_{n-1}(X)) = \phi_1(\tilde{e}(X)) + \cdots + \phi_{n-1}(\tilde{e}(X)) + \deg \geq n.$$

For example (in the base case), setting $n = 2$ we just need to look at linear terms of everything, and we see that this is satisfied by the choice of ϕ_1 because

$$a_1\tilde{e}(X_1) + \cdots + a_r\tilde{e}(X_r) = a_1\pi X_1 + \cdots a_r\pi X_r \mod \deg \geq 2.$$

In general, we have (just by looking at degrees of terms and what can possibly contribute to the sum mod higher degree)

$$e(\phi_1(X) + \cdots + \phi_n(X)) = e(\phi_1(X) + \cdots + \phi_{n-1}(X)) + \pi\phi_n(X) + (\deg \geq n+1)$$

while

$$\phi_1(\tilde{e}(X)) + \cdots + \phi_n(\tilde{e}(x)) = (\phi_1 + \cdots + \phi_{n-1})(\tilde{e}(X)) + \pi^n\phi_n(X) + (\deg \geq n+1).$$

So we are forced to take ϕ_n to be the homogeneous degree n part of

$$\frac{e(\phi_1(X), \dots, \phi_{n-1}(X)) - (\phi_1 + \cdots + \phi_{n-1})(\tilde{e}(X))}{\pi^n - \pi}.$$

It remains to check that this has coefficients in $\mathcal{O}_K$, which is actually forced by the condition that $e(X) \equiv X^q \mod \pi_K$ and the same for $\tilde{e}$. In particular, it shows that

$$e((\phi_1 + \phi_{n-1})(X)) \equiv (\phi_1 + \cdots + \phi_{n-1})(\tilde{e}(X)) \mod \pi_K.$$

□

Proof of the Theorem using this lemma. By the lemma, there is a unique $F_e(X, Y)$ such that $F_e(X, Y) \equiv X + Y \mod \deg 2$ and

$$e(F_e(X, Y)) = F_e(e(X), e(Y))$$

(so a unique group law for which e is an endomorphism). From the same lemma, for any a there is a unique choice of $[a]_{F_e}(X) = aX + \deg \geq 2$ such that $e([a]_{F_e}(X)) = [a]_{F_e}(e(X))$. So there is at most one choice of Lubin–Tate module structure (given by these choices). We need to check that this is actually a valid Lubin–Tate module. So we need to check F_e is a bona fide group law, and that our choice of $[a]_{F_e}$ (the action of $\mathcal{O}_K$ on the group) makes it into a bona-fide $\mathcal{O}_K$ module. All of these are routine applications of the uniqueness statement of the lemma. For example, take the associative law. The power series $F_e(X, F_e(Y, Z))$ and $F_e(F_e(X, Y), Z)$ both have linear terms $X + Y + Z$, and they both commute with e (since F_e does). So the lemma tells us they are the unique solution described. Consider also the distributive law $[a]_{F_e}(F_e(X, Y)) = F_e([a]_{F_e}(X), [a]_{F_e}(X))$ for $a \in \mathcal{O}_K$. Both are the unique solution to $\phi(X, Y) = aX + aY \mod \deg \geq 2$ which commute with e. All the others follow the exact same scheme. □

We said in the example that we "only need one Lubin–Tate series". That is because of another consequence of this stuff, which says that the group laws F_e are all isomorphic if the same choice of π_K is always made.

Theorem 14.7

If $e(X), \tilde{e}(X)$ are Lubin–Tate series for the same uniformizer π_K of K, then F_e and $F_{\tilde{e}}$ are isomorphic as formal $\mathcal{O}_K$-modules.

Proof. By the lemma, there is a unique $f(X) \equiv X \mod \deg \geq 2$ such that $f(e(X)) = \tilde{e}(f(X))$. The uniqueness statement of the lemma shows that f is a homomorphism $F_e \to F_{\tilde{e}}$ (in exactly the same way as the previous theorem), and this f is invertible because it has linear coefficient 1. So we have produced the desired isomorphism. □

After having developed the theory of formal modules, what we actually want is to use these group laws to construct points which we will adjoin in order to get totally ramified extensions of K (this is what we want in local CFT). To do this we will use one of the formal modules F_e to construct a bona fide $\mathcal{O}_K$-module structure on the maximal ideal of the separable closure of K. First of all, by the nonarchimedean triangle inequality it's easy to see that for any formal group law F over $\mathcal{O}_K$, we can define an actual group law on $\mathfrak{p}_K$ by $x + y := F(x, y)$ [this power series will always converge to an element of the maximal ideal]. Similarly, a formal $\mathcal{O}_K$-module also produces an $\mathcal{O}_K$-module structure on $\mathfrak{p}_K$.

Example 14.8

Take $F = \mathbf{G}_a$. Then F induces the standard additive $\mathcal{O}_K$-module structure on $\mathfrak{p}_K$.

Example 14.9

Take $F = \mathbf{G}_m$. Then F induces the $\mathcal{O}_K$-module structure on $\mathfrak{p}_K$ given by the multiplicative group structure on $1 + \mathfrak{p}_K$.

But we want to find elements in $\overline{K}$ to adjoin. So extend the valuation on K to every finite extension of K. The power series involved all converge on $\mathfrak{p}_{\overline{K}}$, because we can always reduce to the finite extension $K(x, y)$. This we can use to construct the abelian extensions K_n from before such that

$$K^{ab} = \bigcup_{n \geq 0} K_n.$$

Choose a uniformizer π_K and consider one of the Lubin-Tate modules F for π_K, with $e(X) = [\pi]_F(X)$. This is all in analogy to the fact that over $\mathbf{Q}_p$ (from local Kronecker–Weber) that we should get the K_n's from adjoining p-th roots of unity. For a general group law, these should be the torsion points of the action of powers of $[\pi]_F$. So we consider the sets of torsion points

$$F(n) = \{\lambda \in \mathfrak{p}_{\overline{K}} : [\pi^n]_F(\lambda) = 0\}$$

so we can let $K_{\pi,n}$ be given by adjoining all the elements of $F(n)$ to K. We have a tower of extensions

$$K = K_{\pi,0} \subset K_{\pi,1} \subset \cdots$$

We need to prove (as we saw was necessary) that $K_{\pi,n}$ is a totally ramified extension of K of appropriate degree.

§15 October 22, 2019

Let $e(X) \in \mathcal{O}_K[[X]]$ be a Lubin–Tate series (recall that this means $e(X) \equiv X^q \mod \mathfrak{p}_K$ and $X \equiv \pi_K X \mod \deg \geq 2$). We obtained a unique Lubin–Tate module F such that

$$[\pi_K]_F(X) = e(X)$$

(the isomorphism class of F depends on the choice of π_K but not on F). Moreover, we turned F into a bona fide group by plugging in points of positive valuation in $\overline{K}$. The π_K^n-torsion points of F are then

$$F(n) = \{\lambda \in \mathfrak{p}_{\overline{K}} | [\pi_K^n]_F(\lambda) = 0\},$$

and we will consider the fields $K_{\pi,n}$ obtained by adjoining $F(n)$ to K.

Lemma 15.1

Let $n \geq 1$. Then $|F(n)| \leq q^n$. Also, $F(n) \supsetneq F(n-1)$. For any $\lambda_n \in F(n) \smallsetminus F(n-1)$, actually $K(\lambda_n)/K$ is a totally ramified extension of degree $q^{n-1}(q-1)$ with uniformizer λ_n.

Proof. Since up to an isomorphism of actual groups, this doesn't depend on the choice of e, we will always just take $e(X) = X^q + \pi X$. Then $\lambda \in F(n)$ if and only if $e^n(\lambda) = 0$. But e^n is a polynomial of degree q^n, so it has at most q^n roots. This establishes the bound $|F(n)| \leq q^n$. For the remaining statements, we use induction over n.

First, take $n = 1$. Then $\lambda \in F(1)$ if and only if $\lambda^q + \pi\lambda = 0$, and $\lambda \in F(0)$ if and only if $\lambda \neq 0$. So in fact

$$\lambda^{q-1} + \pi = 0$$

so λ satisfies an Eisenstein polynomial, forcing the extension $K(\lambda_n)/K$ to be totally ramified of degree $q-1$, and with uniformizer λ_n. For the inductive step, suppose the statement is established for $n-1$. Let $\lambda' = e(\lambda)$. Then $\lambda \in F(n)$ if and only if $\lambda' \in F(n-1)$ and $\lambda \notin F(n-1)$ if and only if $\lambda' \notin F(n-2)$. By the inductive hypothesis there exists a $\lambda' \in F(n-1) \smallsetminus F(n-2)$, and its valuation (by the fact that this is totally ramified) is

$$v_K(\lambda') = \frac{1}{q^{n-2}(q-1)}.$$

The polynomial $e(X) - \lambda' = X^q + \pi X - \lambda'$ is Eisenstein since λ' is a uniformizer for $K(\lambda')$, so by taking λ to be a root of it, we are done. In particular, $K(\lambda', \lambda)/K(\lambda')$ is totally ramified of degree q, but $K(\lambda', \lambda) = K(\lambda)$ [λ' is a polynomial in λ], so in the end we have a tower of totally ramified extensions where the degree gets multiplied by q, and thus $K(\lambda)/K$ is a totally ramified extension of the specified degree. □

Theorem 15.2

The $\mathcal{O}_K$-module $F(n)$ is isomorphic to $\mathcal{O}_K/\mathfrak{p}_K^n$. So in particular its size is exactly q^n.

Proof. By the previous lemma, choose $\lambda \in F(n) \smallsetminus F(n-1)$. Analogously to the case where $K = \mathbf{Q}$, where you get all the roots of unity just by exponenting a primitive root of unity, we'll get all the λ's by hitting λ with $[a]_F$ for each $a \in \mathcal{O}_K$. In particular we have an $\mathcal{O}_K$-module homomorphism

$$\mathcal{O}_K \to F(n)$$

given by

$$a \mapsto [a]_F(\lambda).$$

The kernel contains $\mathfrak{p}_K^m$ if and only if $[\pi_K^m]_F(\lambda) = 0$ which is equivalent to $m \geq n$. So in fact the kernel of this map equals $\mathfrak{p}^n$. It remains to show surjectivity. But this is obvious just from computing sizes. We have so far an injective homomorphism

$$\mathcal{O}_K/\mathfrak{p}_K^n \to F(n),$$

where the LHS has size q^n and the RHS we already know has size at most q^n. So this forces the RHS to have size exactly q^n so we are done. □

Corollary 15.3

$K_{\pi,n}/K$ is a Galois extension with Galois group $(\mathcal{O}_K/\mathfrak{p}_K^n)^\times$.

Proof. WLOG $e(X) = X^q + \pi_K X$. By definition, $K_{\pi,n}$ is the splitting field of $e^n(X)$, and we saw earlier (just by the description in terms of the polynomials involved in the induction) that $K_{\pi,n}/K$ is separable. So in fact it is Galois. Also, any element $\sigma \in \operatorname{Gal}(K_{\pi,n}/K)$ induces an $\mathcal{O}_K$-module automorphism of $F(n)$. We have

$$|\operatorname{Gal}(K_{\pi,n}/K)| = q^{n-1}(q-1)$$

and (from the identification of elements of the Galois group with module automorphisms) an injective map

$$\operatorname{Gal}(K_{pi,n}/K) \to \operatorname{Aut}_{\mathcal{O}_K}(F(n)) \cong \operatorname{Aut}_{\mathcal{O}_K}(\mathcal{O}_K/\mathfrak{p}_K^n) \cong (\mathcal{O}_K/\mathfrak{p}_K^n)^\times.$$

The LHS and RHS have the same size, so in fact we have an isomorphism, as desired. □

Corollary 15.4

$K_\pi := \bigcup_{n\geq 0} K_{\pi,n}$ is a Galois extension of K with Galois group $\mathcal{O}_K^\times$.

Proof. The theorem tells us that

$$\operatorname{Gal}(K_\pi/K) \cong \varprojlim (\mathcal{O}_K/\mathfrak{p}_K^n)^\times$$

which is exactly what we claim it is. □

Theorem 15.5

$I^t(K_{\pi,n}/K) = \operatorname{Gal}(K_{\pi,n}/K_{\pi,t}) \cong U_K^{(t)}/U_K^{(n)}$.

Proof. Consider a nontrivial automorphism $\sigma \in \mathrm{Gal}(K_{\pi,n}/K)$ corresponding to $a \in \mathcal{O}_K^\times/U_K^{(n)}$. If a does not fix $K_{\pi,1}$, i.e. $a \notin U_K^{(1)}$, then

$$\begin{aligned} i_{K_{\pi,n}/K}(\sigma) &= v_{K_{\pi,n}}(\sigma(\lambda_n) - \lambda_n) \\ &= v_{K_{\pi,n}}([a]_F(\lambda_n) - \lambda_n) \\ &= 1 \end{aligned}$$

because $[a]_F(\lambda_n)$ is divisible by λ_n exactly once (by definition of a Lubin–Tate power series) and λ_n is a uniformizer for $K_{\pi,n}$. More generally, if $\sigma \in \mathrm{Gal}(K_{\pi,n}/K_{\pi,t}) \smallsetminus \mathrm{Gal}(K_{\pi,n}/K_{\pi,t+1})$, i.e. $a \in U_K^{(t)}/U_K^{(n)} \smallsetminus U_K^{(t+1)}/U_K^{(n)}$, then (writing $a = 1 + b\pi_K^t$ for b a unit)

$$\begin{aligned} i_{K_{\pi,n}}(\sigma) &= v_{K_{\pi,n}}(\sigma(\lambda_n) - \lambda_n) \\ &= v_{K_{\pi,n}}([a]_F(\lambda_n) - \lambda_n) \\ &= v_{K_{\pi,n}}(\lambda_{n-t}) \\ &= q^t \end{aligned}$$

Using the fact that $K_{\pi,n}/K_{\pi,n-t}$ is totally ramified. This shows that in lower numbering,

$$I_0 \cong \mathcal{O}_K^\times/U_K^{(n)},$$

$$I_1, \dots, I_{q-1} \cong U_K^{(1)}/U_K^{(n)},$$

$$I_q, \dots, I_{q^2-1} \cong U_K^{(2)}/U_K^{(n)},$$

et cetera until

$$I_{q^n-1} = 1$$

so by computing the indices of these groups inside each other and applying the definition of the upper numbering, we get the desired statement. □

Corollary 15.6

$I^t(K_\pi/K) = \mathrm{Gal}(K_\pi/K_{\pi,t}) \cong U_K^{(t)} \subseteq \mathcal{O}_K^\times \cong \mathrm{Gal}(K_\pi/K)$.

Corollary 15.7 (Local CFT, assuming Hasse–Arf)

The maximal abelian extension of K is $K^{ur} \cdot K_\pi$, and thus $\mathrm{Gal}(K^{ab}/K) \cong \hat{\mathbf{Z}} \times \mathcal{O}_K^\times$.

§16 October 24, 2019

Today the goal is to actually construct the local Artin reciprocity map. Last time we constructed the maximal abelian extension of a local field K, but the construction depended on the choice of π_K.

Theorem 16.1 (Construction of the local Artin map)
The map

$$\theta : K^\times \to \mathrm{Gal}(K^{ab}/K) = \mathrm{Gal}(K^{ur}K_\pi/K) \cong \widehat{\mathbf{Z}} \times \mathcal{O}_K^\times$$

given by $a\pi_K^n \mapsto (n, a)$ is injective with dense image and independent of the choice of uniformizer π_K.

Proof. The fact that θ is injective with dense image is obvious from the definition, as $\mathbf{Z}$ is dense in $\widehat{\mathbf{Z}}$. To show that it is independent of π is more annoying. Take two uniformizers π_K and π'_K for K. They have Lubin–Tate modules F and F', respectively. They are nonisomorphic as formal modules, but the $\mathcal{O}_K$-modules they induce on the completion of K^{ur}. This can be shown using almost the same technique as for the proof that any Lubin–Tate series for the same uniformizer induce isomorphic Lubin–Tate modules. The proof is in Neukirch's *Algebraic Number Theory*, Chapter V, Thm. 2.2, Corollary 2.3. See Theorem 5.5 also. □

Corollary 16.2
The subgroup $\mathrm{Gal}(K^{ab}/K_\pi)$ corresponds under the Artin reciprocity map to $\langle \pi \rangle \subseteq K^\times$. So, $K_\pi \neq K_{\pi'}$ when π and π' are distinct uniformizers for K.

Now we turn to the question of how to go from local to global. One example is how to go from local Kronecker–Weber to the global version.

Theorem 16.3
$\mathbf{Q}^{ab} = \mathbf{Q}(\zeta_\infty)$.

Proof. Recall (for example as a consequence of Minkowski's theorem) that $\mathbf{Q}$ has no nontrivial unramified extensions. Actually we will do it now. Recall that $K/\mathbf{Q}$ being unramified means that its discriminent $d_K \in \mathbf{Z}$ is ± 1. By Minkowski's theorem, there exists a nonzero $a \in \mathcal{O}_K$ such that

$$|N_{K/\mathbf{Q}}(a)| \le \frac{n!}{n^n} \left(\frac{4}{\pi}\right)^{n/2} \sqrt{|D|} < 1$$

which is impossible since the norm of any nonzero element of $\mathcal{O}_K$ is a nonzero rational integer.

Let $K/\mathbf{Q}$ be a finite abelian extension. Changing a prime upstairs just conjugates the corresponding ramification group, so since it is abelian actually all the ramification groups $I(\mathfrak{p}|p)$ are the same if p remains fixed. For ease of notation, write $I^t(p)$ to be this well-defined subgroup of $\mathrm{Gal}(K/\mathbf{Q})$. Let $a_p \ge 0$ be minimal such that $I^{a_p}(p)$ is trivial. Let $n = \prod_p p^{a_p}$. We want to show that $K \subseteq \mathbf{Q}(\zeta_n)$. Replace K with $K \cdot \mathbf{Q}(\zeta_n)$. Note that (from the homework) $I^{a_p}_{\mathbf{Q}(\zeta_n)/\mathbf{Q}}(p) = 1$ implies that $I^{a_p}(K \cdot \mathbf{Q}(\zeta_n)/\mathbf{Q})$ is still 1. To show the desired equality, it now suffices to show the degree is what it should be. We know that $K_{\mathfrak{p}}/\mathbf{Q}_p$ has trivial I^{a_p}, so

$$K_{\mathfrak{p}} \subseteq (\mathbf{Q}_p^{ab})^{I^{a_p}(\mathbf{Q}_p^{ab}/\mathbf{Q}_p)} = \mathbf{Q}_p^{ur}(\zeta_{p^{a_p}}).$$

Moreover, $I^0(K_{\mathfrak{p}}/\mathbf{Q}_p)$ has a surjection (given by the quotient) from $I^0(\mathbf{Q}_p(\zeta_{p^{a_p}})/\mathbf{Q}_p) = (\mathbf{Z}/p^{a_p}\mathbf{Z})^\times$, so taking the product over all p we get

$$\left|\prod_p I^0(p)\right| \leq |(\mathbf{Z}/n\mathbf{Z})^\times|.$$

The subfield of K fixed by $\sum_p I^0(p) \subseteq \operatorname{Gal}(K/\mathbf{Q})$ is unramified, so actually it is equal to $\mathbf{Q}$ [here by the sum of the groups we mean the thing generated by it; this is smaller in size than the product]. Putting this into the previous inequality, we see that

$$[K : \mathbf{Q}] \leq [\mathbf{Q}(\zeta_n) : \mathbf{Q}]$$

so we are done. □

Now we will do some Kummer Theory.

Theorem 16.4 (Hilbert 90)

Let L/K be a cyclic extension with Galois group generated by σ. Then $\alpha \in L^\times$ has norm 1 iff it is of the form $\beta/\sigma(\beta)$ for some $\beta \in L^\times$.

Proof. Only one direction is not obvious. For the other direction, you just take

$$\beta = t + \alpha\sigma(t) + \cdots + \alpha\sigma(\alpha)\cdots\sigma^{n-2}(\alpha)\sigma^{n-1}(t),$$

so that

$$\alpha\sigma(\beta) = \beta.$$

You can use linear independence of characters to show that t can be chosen to make β nonzero. □

Corollary 16.5 (Kummer Theory, cyclic case)

Let K be a field containing n distinct n-th roots of unity, i.e. $\zeta_n \in K$ and the characteristic of K does not divide n. Then each Galois extension L/K with cyclic Galois group of order n is obtained from K by adjoining an n-th root.

Proof. This is a direct application of Hilbert's theorem 90 to ζ_n. We have $N_{L/K}(\zeta_n) = \zeta_n^n = 1$, so by Hilbert 90 there is a $b \in L^\times$ such that $\zeta_n = b/\sigma(b)$. Raising this to the n-th power, we see that $b^n = \sigma(b^n)$. So actually $b^n \in K^\times$. So we predict that in fact $L = K(b)$. To do that it suffices to check that b is not fixed by any power of σ which is not the trivial automorphism. But we know that $b/\sigma(b) = \zeta_n$, so we can induct to see that this is indeed impossible. □

Theorem 16.6 (Integral Hilbert 90)

Let L/K be a Galois extension of local fields with $\mathrm{Gal}(L/K) \cong \mathbf{Z}/n\mathbf{Z}$ generated by σ. Then there exists $r \geq 1$ such that for all $a \in U_L^{(r)}$, $N(a) = 1$ if and only if there exists $b \in U_L^{(1)}$ such that $a = b/\sigma(b)$.

Proof. As before take $b = t + a\sigma(t) + \cdots + a\sigma(a)\cdots\sigma^{n-2}(a)\sigma^{n-1}(t)$. Now take $t \in L^\times$ with trace 1 (this is possible by scaling any element of nonzero trace), and let $r = -v_K(t) + 1$. Then you can check that

$$v_K(b-1) \geq 1$$

as desired. □

§17 October 29, 2019

Today we continue the discussion of Kummer theory. Suppose K contains n distinct n-th roots of unity. Then we saw that every cyclic degree n extension of K is of the form $K(b^{1/n})$ for some $b \in K$. Its Galois group is $\mathbf{Z}/n\mathbf{Z}$ with generator corresponding to the automorphism $b^{1/n} \mapsto \zeta_n b^{1/n}$.

Definition 17.1. Let K be a local field containing n distinct n-th roots of unity. The group of n-th roots of unity in K is denoted $\mu_n = \{\zeta_n^i\} \cong \mathbf{Z}/n\mathbf{Z}$.

Recall that the local reciprocity map

$$\rho_K : K^\times \to \mathrm{Gal}(K^{ab}/K)$$

has dense image. For any $a, b \in K^\times$, we can define the **Hilbert symbol** (a, b) to be the n-th root of unity ω such that $\rho_K(a)(b^{1/n}) = \omega b^{1/n}$. Note that this is independent of the choice of n-th root of b because the n-th roots of unity all lie in K and are therefore fixed by $\rho_K(a)$. We know from local CFT (the kernel of the local Artin map for a finite extension is just the norms from upstairs) that

$$(a, b) = 1 \iff a \in N_{K(b^{1/n})/K}(K(b^{1/n})^\times).$$

Since the degree divides n, it's definitely true that the n-th powers of elements of $K^\times$ are norms. So we can check that

1. $(c^n, b) = 1$ for all $b, c \in K^\times$.
2. $(a, c^n) = 1$ for all $a, c \in K^\times$.

It's just as clear that the Hilbert symbol is multiplicatively bilinear (this is just definition-pushing). Basically all this is is a system for keeping track of what the Artin symbol is on a Kummer extension of local fields. Such an extension is a natural thing to consider for example if you want to generate the law of quadratic reciprocity. To summarize, the Hilbert symbol is a bilinear pairing

$$(\cdot, \cdot) : K^\times/(K^\times)^n \times K^\times/(K^\times)^n \to \mu_n$$

Notice that $K^\times/(K^\times)^n$ is finite. To show this it suffices to show that $\mathcal{O}_K^\times/(\mathcal{O}_K^\times)^n$ is finite because of the noncanonical isomorphism $\mathcal{O}_K^\times \times \mathbf{Z} \cong K^\times$. By Hensel's lemma applied to the polynomial $X^n - t$, we can conclude that for $t \in \mathcal{O}_K^\times$ sufficiently close to 1 this polynomial has a root, and therefore the set of n-th powers is open. Technically this requires a slightly stronger form of Hensel's lemma than the one we actually proved in class (I did it in section). Open subgroups of a compact topological group like $\mathcal{O}_K^\times$ are finite index so we are done.

Lemma 17.2

For all $x \in K$ and $b \in K^\times$ such that $x^n \neq b$,

$$(x^n - b, b) = 1.$$

Proof. You just need to show that $x^n - b$ is a norm from $K(b^{1/n})/K$. For simplicity assume the Galois group is all of $\mathbf{Z}/n\mathbf{Z}$. The norm of $x - b^{1/n}$ is

$$\prod_{i=0}^{n-1} (x - \zeta_n^i b^{1/n}) = x^n - b$$

so we are done in this case. In general if $[K(b^{1/n}) : K] = n/d$, in which case $b = c^d$ for some $c \in K^\times$ and we still get what we want by taking the norm of

$$\prod_{i=0}^{d-1}(x - \zeta_n^i b^{1/n}).$$

□

Corollary 17.3

$(a, 1-a) = 1$ for all $a \neq 0, 1$, and $(a, -a) = 1$ for $a \neq 0$.

Corollary 17.4

The Hilbert symbol is skew-symmetric.

Proof. We can directly compute

$$(a,b)(b,a) = (a,-a)(a,b)(b,a)(b,-b) = (a,-ab)(b,-ab) = (ab,-ab) = 1$$

□

Corollary 17.5

a is a norm from $K(b^{1/n})$ if and only if b is a norm from $K(a^{1/n})$.

It's also true that the Hilbert pairing is nondegenerate: if $(a,b) = 1$ for all b, then a is an n-th power (it suffices to prove this by skew-symmetry). This is obvious because if a is not an n-th power then $[K(a^{1/n}) : K] > 1$ and the Chebotarev density theorem (or if you want the full strength of local Artin reciprocity) implies that there exists a b for which the Artin symbol of b is not the identity.

Corollary 17.6

Let $b_1, \dots, b_r \in K^\times$ be a system of representative for $K^\times/(K^\times)^n$. Then $L = K(b_1^{1/n}, \dots, b_r^{1/n})$ is the maximal abelian extension of K of exponent dividing n. Its norms in $K^\times$ are exactly $(K^\times)^n$.

Proof. We know that

$$\mathrm{Gal}(K^{ab}/K) = \bigcap_i \mathrm{Gal}(K^{ab}/K(b_i^{1/n}))$$

so under Artin reciprocity we get

$$\begin{aligned} N_{L/K}L^\times &= \bigcap_i N_{K(b_i^{1/n})/K}(K(b_i^{1/n})^\times) \\ &= \bigcap_i \{a \in K^\times : (a, b_i) = 1\} \\ &= (K^\times)^n \end{aligned}$$

by the nondegeneracy of the Hilbert symbol. □

Theorem 17.7

Assume that $\mathfrak{p}$ does not divide n, where $\mathfrak{p}$ is the maximal ideal of K. Let the residue field of K be $\mathbf{F}_q$. Then

$$(a,b) \equiv \left((-1)^{v(a)v(b)} \frac{b^{v(a)}}{a^{v(b)}} \right)^{\frac{q-1}{n}} \mod \mathfrak{p}$$

(the fact that $\mu_n \subseteq K$ and $\mathfrak{p}$ does not divide n means that $\mu_n \subseteq \mathbf{F}_q$ and therefore $n|(q-1)$; so the exponent is actually an integer; the same fact also guarantees that this formula uniquely determines the Hilbert symbol, since it is an n-th root of unity, which we just argued remain distinct in $\mathbf{F}_q = \mathcal{O}_K/\mathfrak{p}$).

Proof. Using the fact that both sides are bilinear in a and b, it suffices to do it in the following cases:

1. $a = \pi, b = -\pi$ for π a uniformizer of K.
2. $b \in \mathcal{O}_K^\times$ and a is a unit or a uniformizer.

but actually we proved (1) earlier (skew-symmetry). For the second part, remember that the reciprocity map $K^\times \to \mathrm{Gal}(K^{ab}/K)$ is compatible via the valuation map $K^\times \to \mathbf{Z}$ and the projection $\mathrm{Gal}(K^{ab}/K) \to \mathrm{Gal}(\overline{\mathbf{F}}_q/\mathbf{F}_q)$ with the frobenius $\mathbf{Z} \to \mathrm{Gal}(\overline{\mathbf{F}}_q/\mathbf{F}_q) = \widehat{\mathbf{Z}}$. So actually

$$(\rho_K(a) \mod \mathfrak{p}_{K^{ab}}) = \rho_{\mathbf{F}_q}(v_k(a)) = \varphi_q^{v_k(a)}.$$

As a result,

$$\rho_K(a)(b^{1/n}) = (b^{1/n})^{q^{v(a)}} = (a,b)b^{1/n}$$

and rearranging this yields

$$(a,b) = (b^{1/n})^{q^{v(a)}-1} = b^{\frac{q^{v(a)}-1}{n}}$$

which is actually the same mod $\mathfrak{p}$ to the desired $(b^{v(a)})^{(q-1)/n}$ because $v(a) = 0$ or 1. □

Corollary 17.8

The **Legendre symbol** is defined by the unique n-th root of unity in K satisfying

$$\left(\frac{u}{\mathfrak{q}}\right)_n \equiv u^{(q-1)/n} \mod \mathfrak{p}$$

Theorem 17.9

$\left(\frac{u}{\mathfrak{q}}\right)_n = 1 \iff (u \mod \mathfrak{p}) \in (\mathbf{F}_q^\times)^n$.

Proof. Same as usual, use the fact that $\mathbf{F}_q^\times$ is cyclic of order $q-1$. □

Note that this Legendre symbol is actually defined to be (π, u) for any choice of uniformizer $\pi \in K$.

Theorem 17.10

Let $n = 2$. Then $(a, b) = 1$ if and only if $ax^2 + by^2 = z^2$ has a nontrivial solution over K.

Proof. The Hilbert symbol being 1 is equivalent to a being a norm from $K(\sqrt{b})^\times$, i.e. (if this extension actually has degree 2)

$$a = z^2 - by^2.$$

and we have a solution with $x = 1$. The rest is just checking. □

§18 October 31, 2019

Last time, recall that we had to assume that a is a unit or a uniformizer in case (2) of the computation of the Hilbert symbol. Actually this is not necessary. Suppose $u \in \mathbf{F}_q^\times$, $v \in \mathbf{Z}$ and n divides $q-1$. We want to show that

$$u^{(q^v-1)/n} \equiv u^{v\frac{q-1}{n}} \mod q$$

for which by Fermat's little theorem it suffices to show that

$$\frac{q^v-1}{n} \equiv v\frac{q-1}{n} \mod q-1$$

which is obvious by writing the LHS as

$$\frac{q-1}{n}(1+\cdots+q^{v-1}).$$

Now we begin the topic of group cohomology.

§18.1 Group cohomology

Let G be a finite group.

Definition 18.1. A (left) **G-module** is an abelian group A with a left action of G, where this action is compatible with the group operation of A.

Group cohomology is something which is applied to a pair (G, A) consisting of a group G and G-module A.

Example 18.2

Any abelian group A and any group G with the trivial action on A (take all of G to the identity in $\operatorname{End}(A)$) form a valid G-module. Typically $\mathbf{Z}$ and $\mathbf{Z}/n\mathbf{Z}$ will have the trivial action of G. We'll see that group cohomology applied to $\mathbf{Z}$ will let us read off the size of G.

Example 18.3

The objects of obvious interest are the Galois extensions L/K with the action of $\operatorname{Gal}(L/K)$ on the abelian group $(L,+)$ or $(L^\times, \times)$. The subgroup $\mu_n(L)$ consisting of the n-th roots of unity also works. Similarly, the actions above of $\operatorname{Gal}(L/K)$ restrict to actions on $\mathcal{O}_L$ and $\mathcal{O}_L^\times$. If L and K are number fields, then the group of fractional ideals of L, denoted $\mathfrak{I}(L)$, is equipped with an action of $\operatorname{Gal}(L/K)$ which obviously descends to an action on the class group of L. If E is an elliptic curve defined over K, $E(L)$ has an obvious action of $\operatorname{Gal}(L/K)$ too.

The G-modules form a category.

Definition 18.4. Let A, B be G-modules. A **G-module homomorphism** from A to B is an abelian group homomorphism $f : A \to B$ with the property that $f(g \cdot a) = g \cdot f(a)$ for all $a \in A$ and $g \in G$.

Choosing the G-modules as objects and G-module homomorphisms as morphisms makes the G-modules into a bona fide category. The abelian groups $A \oplus B$ and B/A (when $A \subseteq B$) are G-modules in the obvious way.

A G-module is also the same thing as a module over the group ring $\mathbf{Z}[G]$ (for example because a ring map from $\mathbf{Z}[G]$ to $\operatorname{End}(A)$ is determined by a choice of image for each element of G).

Definition 18.5. The **group of invariants** of A is

$$A^G := \{a \in A : g \cdot a = a \forall a \in A\}.$$

Note that $-^G$ is a functor from the category of G-modules to the category of abelian groups.

Definition 18.6. The **group of covariants** of A is

$$A_G := A/B$$

where B is the subgroup of A generated by the set of elements of the form $ga - a$ for $g \in G$ and $a \in A$. Taking covariants is also a functor from the category of G-modules to the category of abelian groups.

Notice that the group of invariants of A is the largest submodule with trivial induced G-action, and the group of covariants is the largest such quotient.

There is a canonical element $N \in \mathbf{Z}[G]$ given by the sum of all $g \in G$. Applying it to an element $a \in A$ is just taking the "norm" $\sum_g ga$.

The group of invariants is of obvious interest given the examples above. Let's compute some.

Example 18.7

Take $\mathbf{Z}$ with the trivial action of G. Then $\mathbf{Z} = \mathbf{Z}^G = \mathbf{Z}_G$. For any $x \in \mathbf{Z}$ we have $Nx = |G|x$.

Example 18.8

Take L/K a Galois extension and $G = \operatorname{Gal}(L/K)$, and take the usual action of G on L. Then for $x \in L$ we have $Nx = \sum_{g \in G} gx = \operatorname{Tr}_{L/K} x$. Similarly, if we take the action of G on $L^\times$ instead, $Nx = N_{L/K}x$ for $x \in L^\times$.

Example 18.9

$L^G = K$, $(L^\times)^G = K^\times$, $(\mathcal{O}_L^\times)^G = \mathcal{O}_K$, etc.

Example 18.10

$\mathfrak{I}(L)^G \supseteq \mathfrak{I}(K)$, with equality iff L/K is unramified.

Example 18.11

Let E be an elliptic curve over K. Then $E(L)^G = E(K)$.

Lemma 18.12 (The beginning of the exact hexagon I did in section)

If $0 \to A \to B \to C \to 0$ is an exact sequence of G-modules, these maps restrict to an exact sequence

$$0 \to A^G \to B^G \to C^G.$$

Proof. Straightforward. □

Example 18.13

Let L/K be a Galois extension of local fields with group G. We have an exact sequence of G-modules

$$1 \to \mathcal{O}_L^\times \to L^\times \to \frac{1}{e(L/K)} \cdot \mathbf{Z} \to 0$$

but when we take invariants we get only the exact sequence

$$1 \to \mathcal{O}_K^\times \to K^\times \to \frac{1}{e(L/K)} \cdot \mathbf{Z}$$

where the last map is not surjective since its image is just $\mathbf{Z}$ (unless L/K was unramified to begin with).

Example 18.14

Suppose $G = \{e, \sigma\}$ is cyclic of order 2. This has a nontrivial action on $\mathbf{Z}$ where σ is given by negation. We call this G-module $\tilde{\mathbf{Z}}$. We have an exact sequence

$$0 \to \tilde{\mathbf{Z}} \xrightarrow{\cdot 2} \tilde{\mathbf{Z}} \to \mathbf{Z}/2\mathbf{Z} \to 0$$

and taking G-invariants gives

$$0 \to 0 \to 0 \to \mathbf{Z}/2\mathbf{Z}$$

where the last map is clearly not surjective.

Instead of writing down the definition of cohomology, we'll first explicitly write down the definition of H^1.

Definition 18.15. The abelian group of **inhomogeneous 1-cochains** is

$$C^1(G, A) := \{(a_g)_{g \in G} : a_g \in A \forall g \in G\}.$$

The **inhomogeneous 1-cocycles** are

$$Z^1(G, A) := \{(a_g)_{g \in G} : a_{gh} = a_g + g a_h \forall g, h \in G\}$$

and the **inhomogeneous 1-coboundaries** are

$$B^1(G, A) := \{(ga - a)_{g \in G} : a \in A\} \subseteq Z^1(G, A).$$

Then as usual you define the **first cohomology group**

$$H^1(G, A) := Z^1(G, A)/B^1(G, A).$$

As usual H^1 is a functor from the category of G-modules to the category of abelian groups.

Theorem 18.16 (the next part of the exact hexagon)
If $0 \to A \to B \to C \to 0$ is a short exact sequence of G-modules, then we can extend the previous exact sequence to

$$0 \to A^G \to B^G \to C^G \overset{\delta}{\to} H^1(G,A) \to H^1(G,B) \to H^1(G,C).$$

Proof. There's no point explaining this too much since in reality it follows from generalities of homological algebra. Let's go over the definition of δ. For any $c \in C^G$, choose $b \in B$ mapping to c. Then $gb - b$ reduces mod A to

$$g(b \mod A) - (b \mod A) = gc - c = 0$$

so $gb-b \in A$, and hence $(gb-b)_{g\in G} \in Z^1(G,A)$. It remains to check that this is well-defined, i.e. it gives the same result (modulo $B^1(G,A)$) regardless of the choice of b. But this is clear from the fact that b is unique mod A.

To check exactness at C^G, we need to compute the kernel of δ. Suppose $c \in C^G$ and $\delta(c) = 0$. This means there exists $b \in B$ such that $b \in A$ and $gb - b - 0$, i.e. $b \in B^G$ and $c = \delta(b)$ as desired. □

Definition 18.17. A G-module A is **free** if it satisfies the following (equivalent) conditions:

1. A is a free $\mathbf{Z}[G]$-module.
2. $A \cong \bigoplus_{i\in I} \mathbf{Z}[G]$ for some set I.
3. A is a $\mathbf{Z}[G]$-module with a basis.

Definition 18.18. A G-modules A is **coinduced** if it satisfies one of the following (equivalent) statements:

1. $A \cong \mathrm{Hom}_{\mathbf{Z}}(\mathbf{Z}[G], X)$ for some group X with trivial G-action. Here the G-action on A is just given by
$$(g \cdot f)(h) = f(g^{-1}h)$$
2. A is just the set of maps from G to X with the obvious action of G.
3. $A \cong \{(x_g)_{g\in G}\}$.
4. A is the group of formal sums $\sum x_g g$ where $x_g \in X$.

Definition 18.19. A G-module A is **induced** if it satisfies one of the following (equivalent) statements:

1. $A \cong \mathbf{Z}[G] \otimes_{\mathbf{Z}} X$ for some abelian group X where G acts by acting on the first tensor factor.
2. $A \cong \{\sum x_g g : x_g \in X, x_g = 0 \text{ for all but finitely many } g\}$

For finite groups G this is the same as coinduced anyway.

Example 18.20

$(\mathbf{Z}/2\mathbf{Z})[G]$ is coinduced and induced but not free.

The cohomology groups are actually the unique groups satisfying a set of axioms (this should be familiar from algebraic topology).

Theorem 18.21

There is a unique family of **cohomology functors** $H^i(G, -)$ from the category of G modules to the category of abelian groups satisfying the following axioms:

1. If A is coinduced then $H^i(G, A) = 0$ for all A and $i \geq 1$.
2. If $0 \to A \to B \to C \to 0$ is a short exact sequence of G-modules, then there is a long exact sequence
$$0 \to A \to B \to C \to H^1(G, A) \to H^1(G, B) \to H^1(G, C) \to H^2(G, A) \to \cdots$$
3. The construction of the long exact sequence is natural in the usual sense (a map of short exact sequences induces a map of long exact sequences).

By convention, you set $H^0(G, A) = A^G$.

§19 November 5, 2019

Recall the theorem from last time:

Theorem 19.1

There is a unique family of **cohomology functors** $H^i(G,-)$ from the category of G modules to the category of abelian groups satisfying the following axioms:

1. If A is coinduced then $H^i(G,A)=0$ for all A and $i \geq 1$.
2. If $0 \to A \to B \to C \to 0$ is a short exact sequence of G-modules, then there is a long exact sequence
$$0 \to A \to B \to C \to H^1(G,A) \to H^1(G,B) \to H^1(G,C) \to H^2(G,A) \to \cdots$$
3. The construction of the long exact sequence is natural in the usual sense (a map of short exact sequences induces a map of long exact sequences).

Proof. First we prove uniqueness. This is accomplished via a dimension-shifting argument (we will prove uniqueness for H^i given uniqueness for H^{i-1}). Consider the injective homomorphism of G-modules

$$A \to \{\text{maps } G \to A\} = A^*$$

given by

$$a \mapsto (g \mapsto g^{-1}a).$$

This gives a short exact sequence

$$0 \to A \to A^* \to A^*/A \to 0.$$

Also the definition of A^* tells us immediately that A^* is coinduced. The G-module structure on A^* is given by

$$(gf)(h) = f(g^{-1}h).$$

So we have a long exact sequence

$$0 \to A^G \to (A^*)^G \to (A^*/A)^G \to H^1(A) \to H^1(A^*) \to H^1(A^*/A) \to H^2(A) \to \cdots$$

and since the cohomology of coinduced modules needs to be trivial we know that $H^i(A^*)=1$ for all $i \geq 1$, so

$$H^1(G,A) = \operatorname{coker}((A^*)^G \to (A^*/A)^G)$$

and for $i \geq 1$ it's even better:

$$H^{i+1}(G,A) = H^i(G,A^*/A),$$

which inductively proves that the cohomology groups are uniquely defined by these conditions. A similar argument shows that the action on morphisms is uniquely defined. I guess this also gives an inductive construction of the cohomology groups.

It's much more useful to use the derived functor construction instead. To do it we're going to pick a free resolution of $\mathbf{Z}$ as a G-module with the trivial action (it will be convenient when doing computations to explicitly choose a free resolution). This is an exact sequence of G-modules

$$\cdots \to P_2 \to P_1 \to P_0 \to \mathbf{Z} \to 0$$

where all P_i are free. Then you apply the $\mathrm{Hom}_G(-, A)$ functor to this sequence (after omitting the $\mathbf{Z}$ as usual) to get a cochain complex

$$0 \overset{d_0}{\to} \mathrm{Hom}_G(P_0, A) \overset{d^1}{\to} \mathrm{Hom}_G(P_1, A) \overset{d^2}{\to} \cdots$$

whose cohomology we can take, as usual defining

$$H^i(G, A) := \ker(d^{i+1})/\mathrm{im}(d^i).$$

Notice that

$$H^0(G, A) = \ker(d^1) = \mathrm{Hom}_G(P_0/d^1(P_1), A) = \mathrm{Hom}_G(\mathbf{Z}, A) = A^G.$$

Now we explain the construction of the long exact sequence. If $0 \to A \to B \to C \to 0$ is a short exact sequence of G-modules, then the fact that the P_i's are free means that

$$0 \to \mathrm{Hom}_G(P_i, A) \to \mathrm{Hom}_G(P_i, B) \to \mathrm{Hom}_G(P_i, C) \to 0$$

is exact, and maps via the induced maps to the exact sequence

$$0 \to \mathrm{Hom}_G(P_{i+1}, A) \to \mathrm{Hom}_G(P_{i+1}, B) \to \mathrm{Hom}_G(P_{i+1}, C) \to 0$$

and then an application of the snake lemma gives the map. □

Example 19.2

Let $G = \mathbf{Z}/n\mathbf{Z}$ generated by σ. Then there's an obvious free resolution

$$\cdots \to \mathbf{Z}[G] \to \mathbf{Z}[G] \to \mathbf{Z}[G] \to \mathbf{Z} \to 0$$

where the map $\mathbf{Z}[G] \to \mathbf{Z}$ is the map taking $\sum_g a_g g \to \sum_g a_g$ and the other ones alternate between $\sigma - 1$ and the norm $N : \mathbf{Z}[G] \to \mathbf{Z}[G]$. The induced cochain complex is

$$0 \to \mathrm{Hom}_G(\mathbf{Z}[G], A) \to \mathrm{Hom}_G(\mathbf{Z}[G], A) \to \cdots$$

But all of these are equal to A because a G-module homomorphism from $\mathbf{Z}[G]$ to A is determined by the image of σ, and the maps just alternate between the action of $\sigma - 1$ and $N_G : A \to A$. Then we can see that

$$H^0(G, A) = A^G,$$

and the rest are periodic:

$$H^i(G, A) = \frac{\ker(N)}{(\sigma - 1)A}$$

when i is odd, and

$$H^i(G, A) = \frac{A^G}{N_G A}$$

when i is even and positive.

The cyclic case is the only one we'll actually need when we do class field theory (pretty much anything about abelian extensions can be deduced from the cyclic case). The computation we did above shows that the cohomology groups are actual things we are interested in from class field theory (so that hopefully cohomology will help us prove them).

Example 19.3 (example of example)

Let $G = \mathrm{Gal}(L/K)$ where L/K is cyclic. If $A = L^\times$, then $A^G = K^\times$, $\ker N_G = \{\alpha \in L^\times : N_{L/K}\alpha = 1\}$, and
$$\mathrm{im}(\sigma - 1) = \{\sigma(y)/y : y \in L^\times\}.$$
By Hilbert 90, $\ker N_G = \mathrm{im}(\sigma - 1)$, so $H^1(G, L^\times)$ is trivial (this is actually true regardless of whether G is cyclic and is also called Hilbert's theorem 90 even though it was proved by Emmy Noether). So then
$$H^2(G, L^\times) = K^\times / N_{L/K} L^\times$$
which is of obvious interest for class field theory.

Example 19.4

Now set $A = L$ (so now everything is additive, and the "norm" coming from the action of G is actually the trace on L/K). We see that $A^G = K$, $\ker(N) = \mathrm{im}(\sigma - 1)$ by additive Hilbert 90, so $H^1(G, L) = 0$, and $NA = \mathrm{Tr}\, L = K$, so actually all the cohomology groups are trivial (as long as G is cyclic) except $H^0 = K$.

For general groups G, the standard free resolution is given by taking P_i to be the free $\mathbf{Z}$-module on G^{i+1} equipped with the diagonal action of G. This is a free $\mathbf{Z}[G]$-module of rank $|G|^i$, which we can give basis
$$\{(1, g_1, \dots, g_i) : g_1, \dots, g_i \in G\}.$$
these somehow correspond to the standard simplices in singular (co-)homology. The maps $d^i : P_i \to P_{i-1}$ are given on the $\mathbf{Z}[G]$-basis elements by
$$(g_0, \dots, g_i) \mapsto \sum_{j=0}^{i} (-1)^j (g_0, \dots, \hat{g}_j, \dots, g_i).$$
The fact that $d^i \circ d^{i-1} = 0$ is standard. We still need to show exactness, which can be done by showing the maps $h^i : P_i \to P_{i+1}$ given by $(g_0, \dots, g_i) \mapsto (1, \dots, g_i)$ are chain homotopies.

We can view the G-module $\mathrm{Hom}_G(P_i, A)$ as the set of tuples $(a_{g_0,\dots,g_i} : (g_0, \dots, g_i) \in G^{i+1})$ with the property that
$$a_{gg_0,\dots,gg_i} = g a_{g_0,\dots,g_i}.$$
This is called the group of **homogeneous** i-cochains. The differentials are just given by
$$d^i(a_{g_0,\dots,g_i}) = \sum_{j=0}^{i} (-1)^j a_{g_0,\dots,\hat{g}_0,\dots,g_i}.$$

§20 November 7, 2019

Recall we had defined the space of i-homogeneous cochains

$$\tilde{C}^i(G,A) := \operatorname{Hom}_G(P_i, A) = \{\tilde{f} : G^{i+1} \to A : \tilde{f}(gg_0, \dots, gg_i) = g\tilde{f}(g_0, \dots, g_i)\}.$$

The differential on the homogeneous cochains was defined the same way it usually is, e.g. for singular homology:

$$d^i : \tilde{C}^{i-1}(G,A) \to \tilde{C}^i(G,A)$$

by

$$(d^i \tilde{f})(g_0, \dots, g_i) = \sum_{j=0}^{j} (-1)^j \tilde{f}(g_0, \dots, \hat{g}_j, \dots, g_i)$$

and then $H^i(G,A) = \ker(d^{i+1})/\operatorname{im}(d^i)$. In practice it's better to reduce dimensions by 1.

Definition 20.1. The group of **inhomogeneous i-cochains** is

$$C^i(G,A) = \{(a_{g_1,\dots,g_i})_{g_1,\dots,g_i} : a_{g_1,\dots,g_i} \in A\}.$$

There is an isomorphism (this is maybe not the most obvious one)

$$\tilde{C}^i(G,A) \cong C^i(G,A)$$

via

$$\tilde{f} \mapsto (\tilde{f}(1, g_1, g_1g_2, \dots, g_1 \cdots g_i))_{g_1,\dots,g_i \in G}.$$

You can write down an explicit formula for the differentials, but we'll just write it down for the first few.

$$d^1 : C^0(G,A) \cong A \to C^1(G,A)$$

is given by

$$a \mapsto (ga - a)_{g \in G}.$$

Meanwhile,

$$d^2 : C^1(G,A) \to C^2(G,A)$$

is given by

$$(a_g)_{g \in G} \mapsto (g_1 a_{g_2} + a_{g_1} - a_{g_1g_2})_{g_1,g_2 \in G}$$

and (though we hopefully will nver have to use it)

$$d^3 : C^2(G,A) \to C^3(G,A)$$

is given by

$$(a_{g_1,g_2})_{g_1,g_2 \in G} \mapsto (g_1 a_{g_2,g_3} - a_{g_1g_2,g_3} + a_{g_1g_2,g_3} + a_{g_1,g_2g_3} - a_{g_1,g_2}).$$

And the homology groups can be defined as usual.

Example 20.2

From our computation of d^1, we have

$$H^0(G,A) = \ker(d^1) = \{a \in A : ga - a = 0 \forall g \in G\} = A^G.$$

Similarly, $\ker(d^2)$ is the group of "crossed homomorphisms", namely the set of maps $f : G \to A$ such that $f(g_1g_2) = f(g_1) + g_1 f(g_2)$, and $\operatorname{im}(d^1)$ is the group of "principal crossed homomorphisms, namely the set of maps $g : G \to A$ of the form $g \mapsto ga - a$ for some $a \in A$. So $H^1(G,A)$ is the crossed homomorphisms modulo the principal crossed homomorphisms.

Now we move on to the most important (for us) application of group cohomology, namely the case where G is the Galois group of a field extension.

Theorem 20.3

Let L/K be a finite Galois extension with Galois group G. Then L is a coinduced G-module.

Proof. The normal basis theorem says that there exists an $x \in L$ such that $\{gx : g \in G\}$ form a K-basis for L. This implies that $L \cong K[G]$ as a G-module (the K-basis for L is in bijection with G, and the action of G on it is the same as the action of G on itself). □

This results in something which is maybe somewhat stronger than additive Hilbert 90:

Corollary 20.4

$H^i(G, L) = 0$ for all $i \geq 1$.

The multiplicative Hilbert 90 theorem is generalized by the cohomological version of it. Without any cyclic hypothesis on G:

Theorem 20.5 ("Hilbert 90", due to Noether)

$H^1(G, L^\times) = 1$.

Proof. Let $(a_g)_{g\in G} \in Z^1(G, L^\times)$, so it is a crossed homomorphism. Then we know (using multiplicative notation since we are in $L^\times$)

$$a_{gh} = (a_g)(g \cdot a_h)$$

for all $g, h \in G$. The proof goes pretty much like that of Hilbert 90. In particular, let $x \in L$ (this is analogous to the variable t from our proof in the cyclic case)

$$b = \sum_{g\in G} a_g g(x)$$

(which can be made nonzero by appropriate choice of x via the linear independence of characters). Then it's immediate from the definition that

$$a_g g(b) = b \forall g \in G$$

as desired (this proves all the 1-cocycles are 1-coboundaries and therefore H^1 is trivial). □

§20.1 Functoriality in G

Now back to group cohomology in general.

Definition 20.6. Let A be a G-module and A' a G'-module. We call group homomorphisms $\mu : G' \to G$ and $f : A \to A'$ **compatible (for cohomology)** if

$$f(\mu(g')a) = g' f(a) \forall g' \in G', a \in A.$$

Given two such compatible homomorphisms, we get homomorphisms

$$\tilde{C}^n(G,A) \to \tilde{C}^n(G',A')$$

given by

$$(a_{g_0,\dots,g_n})_{g_0,\dots,g_n\in G} \mapsto (f(a_{\mu(g_0'),\dots,\mu(g_n')})_{g_0',\dots,g_n'\in G'}.$$

Example 20.7

If $G = G'$ with $\mu = \mathrm{id}$, then this is the map $H^n(G,A) \to H^n(G,A')$ coming from the functor $H^n(G,-)$.

Definition 20.8. For $H \subseteq G$ a subgroup, the inclusion $\mu : H \to G$ plus $f = \mathrm{id} : A \to A$ induce (via the above construction) the **restriction homomorphism**

$$\mathrm{Res} : H^n(G,A) \to H^n(H,A).$$

Example 20.9

Let $n = 0$. The restriction homomorphism $H^0(G,A) \to H^0(H,A)$ is just the canonical inclusion of A^G into A^G.

Notice that a resolution of $\mathbf{Z}$ by free G-modules is also a resolution by free H-modules, so these things can all be computed using the same resolution.

Definition 20.10. For $H \subseteq G$ a normal subgroup and any G-module A, the projection

$$G \to G/H$$

and inclusion $A^H \to A$ (where A^H has the obvious G/H-module structure) induce the **inflation homomorphism**

$$\mathrm{Inf} : H^n(G/H, A^H) \to H^n(G,A).$$

Definition 20.11. For $H \subseteq G$ and any H-module A, the **induced G-module** of A is

$$\mathrm{Ind}_H^G A = \mathbf{Z}[G] \otimes_{\mathbf{Z}[H]} A.$$

Here $\mathbf{Z}[H]$ acts on $\mathbf{Z}[G]$ on the right, but in the G-module structure for the tensor product, G acts on the $\mathbf{Z}[G]$-coordinate on the left. Alternatively, it is the set of maps $\phi : G \to A$ such that $\phi(hg) = h\phi(g)$.

Here a map $\phi : G \to A$ corresponds to

$$\sum_{g \in G^{\setminus H}} g^{-1} \otimes \phi(g).$$

Example 20.12

If $H = 1$, then

$$\mathrm{Ind}_H^G A = \mathbf{Z}[G] \otimes_{\mathbf{Z}} A = \{\text{maps } \phi : G \to A\}.$$

Similarly to in representation theory we have Frobenius reciprocity.

Theorem 20.13

Let $H \subseteq G$ be a subgroup, A a G-module, and B an H-module. Then

$$\operatorname{Hom}_G(A, \operatorname{Ind}_H^G B) \cong \operatorname{Hom}_H(A, B).$$

Proof. The isomorphism is given by sending the G-module homomorphism $a \mapsto (g \mapsto \phi(a)g)$ to the H-module homomorphism $a \mapsto \phi(a)1$. The inverse is given by sending $a \mapsto f(a)$ to $a \mapsto (g \mapsto f(ga))$. □

The functor Ind_H^G from the category of H-modules to the category of G-modules is actually exact.

Theorem 20.14 (Shapiro's Lemma)

Let $H \subseteq G$ be a subgroup, A an H-module. Then there is a canonical isomorphism

$$H^n(G, \operatorname{Ind}_H^G A) \cong H^n(H, A).$$

Proof. Let $\cdots \to P_1 \to P_0 \to \mathbf{Z} \to 0$ be a free resolution of G-modules. By a previous remark it is also a free resolution of H-modules. By Frobenius reciprocity,

$$\operatorname{Hom}_G(P_i, \operatorname{Ind}_H^G A) \cong \operatorname{Hom}_H(P_i, A)$$

which commutes with the differentials, which is exactly what is necessary. □

Example 20.15

Let $n = 0$. Then Shapiro's lemma shows that

$$(\operatorname{Ind}_H^G A)^G \cong A^H.$$

§21 November 12, 2019

Last time we discussed the functoriality of group cohomology. Let $H \subseteq G$ be a subgroup. This inclusion induces the restriction and inflation homomorphisms

$$\mathrm{Res} : H^n(G, A) \to H^n(H, A)$$

and

$$\mathrm{Inf} : H^n(G/H, A) \to G^n(G, A)$$

under the additional hypothesis that H is normal in G. We also stated Shapiro's lemma:

Lemma 21.1

If $H \subseteq G$ is a subgroup, and A is an H-module, then

$$H^n(G, \mathrm{Ind}_H^G A) \cong H^n(H, A)$$

which was a straightforward consequence of Frobenius reciprocity.

Definition 21.2. For $H \subseteq G$ of finite index, and any G-module A, the homomorphism

$$\mathrm{Ind}_H^G A \to A$$

given by (using the definition $\mathrm{Ind}_H^G A = A \otimes_{\mathbf{Z}[H]} \mathbf{Z}[G]$)

$$g \otimes a \mapsto g \cdot a.$$

This homomorphism induces a homomorphism of cohomology groups

$$\mathrm{Cor} : H^n(H, A) \cong H^n(G, \mathrm{Ind}_H^G A) \to H^n(G, A).$$

This is called the **corestriction** homomorphism.

Lemma 21.3

The composition $\mathrm{Cor} \circ \mathrm{Res} : H^n(G, A) \to H^n(G, A)$ is multiplication by $[G : H]$.

Proof. Let P_i be a free resolution of G-modules for $\mathbf{Z}$. Cor∘Res is induced by $\mathrm{Hom}_G(P_i, A) \to \mathrm{Hom}_H(P_i, A) \cong \mathrm{Hom}_G(P_i, \mathrm{Ind}_H^G A)$. Using the definition of $\mathrm{Ind}_H^G A$ as the set of maps from G to A, this takes $f \in \mathrm{Hom}_G(P_i, A)$ to the element of $\mathrm{Hom}_G(P_i, \mathrm{Ind}_H^G A)$ which maps

$$p \mapsto (g \mapsto f(gp))$$

(this is just tracing through the Frobenius reciprocity isomorphism). The map $[g \mapsto f(gp)] \in \mathrm{Ind}_H^G A$ is (using the tensor product definition) the same as $\sum_{g \in G/H} g^{-1} \otimes f(gp)$, and under corestriction this goes to the element of $\mathrm{Hom}_G(P_i, A)$

$$p \mapsto \sum_{g \in G/H} g^{-1} f(gp) = \sum_{g \in G/H} f(p) = [G : H] f(p)$$

as desired. □

Remark 21.4. $H^n(1, A)$ is A if $n = 0$ and 0 in higher dimensions.

Corollary 21.5

$|G| \cdot H^n(G, A) = 0$ for all $n \geq 1$.

Proof. Apply the lemma with $H = 1$, so that $H^n(H, A) = 0$ and the lemma says that multiplying by $|G|$ is the zero map. □

Remark 21.6. If the multiplication-by-$|G|$ map $A \to A$ is an isomorphism, then this means $H^n(G, A) = 0$ for all $n \geq 1$.

Theorem 21.7

Let $H \subseteq G$ be a normal subgroup and A a G-module. Let $n \geq 1$. If $H^i(H, A) = 0$ for all $i = 1, \ldots, n-1$, then

$$0 \to H^n(G/H, A^H) \overset{\text{Inf}}{\to} H^n(G, A) \overset{\text{Res}}{\to} H^n(H, A)$$

is exact (it is called the **inflation-restriction exact sequence**).

Proof. The proof is by induction on n. For $n = 1$, you just need to compute the maps on the level of inhomogeneous cocycles.

To go from n to $n+1$, let $A^* = \mathrm{Hom}_{\mathbf{Z}}(\mathbf{Z}[G], A) = \{\text{maps } G \to A\}$. The short exact sequence

$$0 \to A \to A^* \to A^*/A \to 0$$

of G-modules induces a long exact sequence, a section of which looks like

$$0 = H^k(G, A^*) \to H^k(G, A^*/A) \to H^{k+1}(G, A) \to H^{k+1}(G, A^*) = 0$$

for $k \geq 1$ (since A^* is coinduced). So this gives an isomorphism $H^k(G, A^*/A) \cong H^{k+1}(G, A)$. Now taking $k = 0$, we get fewer zeroes, and the exact sequence is

$$0 \to A^H \to (A^*)^H \to (A^*/A)^H \to H^1(H, A) = 0$$

and roughly as before

$$H^{n+1}(G/H, A^H) \cong H^n(G/H, (A^*/A)^H).$$

By induction the exact sequence

$$0 \to H^n(G/H, (A^*/A)^H) \overset{\text{Inf}}{\to} H^n(G, A^*/A) \overset{\text{Res}}{\to} H^n(H, A^*/A)$$

is isomorphic (via the isomorphisms we just described) to

$$0 \to H^{n+1}(G/H, A^H) \overset{\text{Inf}}{\to} H^{n+1}(G, A^*/A) \overset{\text{Res}}{\to} H^{n+1}(H, A^*/A).$$

We didn't check that the diagram actually commutes. □

§21.1 The cup product

Let A, B be G-modules and $r, s \geq 0$. Then $A \otimes_{\mathbf{Z}} B$ is also a G-module.

Definition 21.8. The **cup product**

$$\smile : H^r(G, A) \times H^s(G, B) \to H^{r+s}(G, A \otimes B)$$

is defined on inhomogenous cocycles by

$$\tilde{f}_1 \smile \tilde{f}_2(g_0, \dots, g_{r+s}) = \tilde{f}_1(g_0, \dots, g_r) \otimes \tilde{f}_2(g_r, \dots, g_{r+s}).$$

Example 21.9

Let $r = s = 0$. Then the cup product is the map

$$\smile : A^G \times B^G \to (A \otimes B)^G$$

given by

$$(a, b) \mapsto a \otimes b.$$

Remark 21.10. You can check a bunch of basic properties for the cup product:

1. $(x \smile y) \smile z = x \smile (y \smile z)$
2. $x \smile y = (-1)^{rs} y \smile x$ if we identify $A \otimes B \cong B \otimes A$.
3. For $H \subseteq G$,

   $$\operatorname{Res}(x \smile y) = \operatorname{Res}(x) \smile \operatorname{Res}(y)$$

   and

   $$\operatorname{Cor}(x \smile \operatorname{Res}(y)) = \operatorname{Cor}(x) \smile y$$

§21.2 Group homology

Remember that $A_G := A/\langle ga - a : g \in G, a \in A\rangle$ is the largest quotient of A on which G acts trivially.

Definition 21.11. The kernel of the **augmentation map** $\mathbf{Z}[G] \to G$ given by $\sum a_g g \mapsto \sum a_g$ is called the **augmentation ideal** I (sometimes called I_G).

Remark 21.12. $A_G = A/IA$. Also I has $\mathbf{Z}$-basis given by $\{g - 1\}_{g \neq 1}$

Theorem 21.13

There is a unique family of **group homology functors**

$$H_i(G,-):\mathsf{Mod}_G\to\mathsf{Grp}$$

satisfying the following axioms:

1. If $0\to A\to B\to C\to 0$ is a short exact sequence of G-modules, then we get a long exact sequence of groups
$$\dots\to H_1(G,A)\to H_1(G,B)\to H_1(G,C)\to A_G\to B_G\to C_G\to 0.$$
2. If A is induced, then $H_i(G,A)=0$ for all $i\ge 1$.
3. A map of short exact sequences induces a map of long exact sequences.

Proof. The proof is the same as for cohomology, but for the explicit construction it's obtained from $\operatorname{Hom}_G(P_i,A)$ instead of the tensor product. □

Theorem 21.14

$H_1(G,\mathbf{Z})\cong G^{\mathrm{ab}}$.

Proof. The short exact sequence

$$0\to I\to\mathbf{Z}[G]\to\mathbf{Z}\to 0$$

induces a long exact sequence, a part of which looks like

$$H_1(G,\mathbf{Z}[G])\to H_1(G,\mathbf{Z})\to I_G\to\mathbf{Z}[G]_G\to\mathbf{Z}_G\to 0$$

where I_G means the maximal quotient of I on which G acts trivially (not I). Since $\mathbf{Z}[G]$ is induced and $I_G=I/I^2$ and $\mathbf{Z}[G]_G=\mathbf{Z}[G]/I$ and $\mathbf{Z}_G=\mathbf{Z}$, this is really an exact sequence

$$0\to H_1(G,\mathbf{Z})\to I/I^2\to\mathbf{Z}[G]/I\to\mathbf{Z}\to 0,$$

except the map $I/I^2\to\mathbf{Z}[G]/I$ is the zero map so $H_1(G,\mathbf{Z})\cong I/I^2$. The isomorphism

$$G^{\mathrm{ab}}\cong I/I^2$$

is given by $g\mapsto g-1 \mod I^2$. This map is clearly surjective. Moreover, since I has basis $\{g-1\}$, we know that I^2 has basis

$$\{(g_1-1)(g_2-1)\}=\{(g_1g_2-1)-(g_1-1)-(g_2-1)\}.$$

□

§22 November 14, 2019

Today we will talk about Tate cohomology. This will allow us to combine the definitions of homology and cohomology, and to fix the issues that happen in dimension zero (recall that in the cyclic case the long exact sequence was almost periodic; this will force it to be periodic with period 2)

Definition 22.1. The **Tate cohomology groups** (sometimes denoted $\hat{H}$ too) are $H_T^n(G,A) := H^n(G,A)$ for $n \geq 1$,

$$H_T^0(G,A) = A^G/NA,$$

$$H_T^{-1}(G,A) = \ker(N)/IA,$$

and $H^{-n}(G,A) = H_{n-1}(G,A)$ for $n \geq 1$.

Theorem 22.2

A short exact sequence of G-modules $0 \to A \to B \to C \to 0$ induces a long exact sequence of groups

$$\begin{tikzcd}
H_T^{-1}(G,A) \arrow[r] & H_T^{-1}(G,B) \arrow[r] & H_T^{-1}(G,C) \arrow[dll] \\
H_T^{0}(G,A) \arrow[r] & H_T^{0}(G,B) \arrow[r] & H_T^{0}(G,C) \arrow[dll] \\
H_T^{0}(G,A) \arrow[r] & H_T^{0}(G,B) \arrow[r] & H_T^{0}(G,C)
\end{tikzcd}$$

Proof. Both sides of the sequence come from the long exact sequences for homology and cohomology, but the H_T^0's have been engineered to make them fit together. In particular, the last part of the long exact sequence for homology

$$H_0(G,A) \to H_0(G,B) \to H_0(G,C) \to 0$$

maps (by taking norms) to the first part of the long exact sequence for cohomology

$$0 \to H^0(G,A) \to H^0(G,B) \to H^0(G,C)$$

so if we restrict to the kernel of the norm on the H_0's and mod out by the norm on the H^0's we are good for exactness. □

Another way to construct the Tate cohomology groups is to take a free resolution of G-modules

$$0 \leftarrow \mathbf{Z} \leftarrow P_0 \leftarrow P_1 \leftarrow P_2 \leftarrow \cdots$$

and the dual resolution

$$0 \to \mathbf{Z} \to P_{-1} \to P_{-2} \to \cdots$$

obtained by taking $P_{-i} = \operatorname{Hom}(P_i, \mathbf{Z})$ to get a long exact sequence

$$\cdots \leftarrow P_{-2} \leftarrow P_{-1} \leftarrow P_0 \leftarrow P_1 \leftarrow P_2 \leftarrow \cdots$$

and then take the cohomology of this cochain complex.

Theorem 22.3

If A is a coinduced G-module, then all the Tate cohomology groups vanish.

Proof. We only need to check this in dimensions $n = 0, -1$, and this is on the homework. □

The restriction and corestriction maps, as well as cup products, can be extended to the Tate cohomology groups by dimension shifting (maybe there is some problem with extending inflation to negative indices).

Example 22.4

Remember that when G is cyclic the cohomology groups are particularly nice (the long exact sequence will be periodic). The Tate cohomology is even nicer, and its long exact sequence becomes the "exact hexagon." Let $G \cong \mathbf{Z}/n\mathbf{Z}$ with generator σ. Take the free resolution of G-modules

$$\mathbf{Z} \overset{\epsilon}{\leftarrow} \mathbf{Z}[G] \overset{\sigma-1}{\leftarrow} \mathbf{Z}[G] \overset{N}{\leftarrow} \mathbf{Z}[G] \leftarrow \cdots.$$

When we take Homs we get $\mathrm{Hom}(\mathbf{Z}[G], \mathbf{Z}) \cong \mathbf{Z}[G]$ in the usual way, so computing the cohomology of the big cochain complex we see that $H_T^n(G, A) = A^G/NA$ for all even n and $H_T^n(G, A) = \ker(N)/IA$ for all odd n.

Example 22.5

$H_T^{-2}(G, \mathbf{Z}) = H_1(G, \mathbf{Z}) = G^{\mathrm{ab}}$ (we proved this earlier). The corestriction map

$$H_T^{-2}(H, \mathbf{Z}) \to H_T^{-2}(G, \mathbf{Z})$$

for a subgroup $H \subseteq G$ is given by the inclusion of $H^{\mathrm{ab}} \to G^{\mathrm{ab}}$. Also the restriction map

$$H_T^{-2}(G, \mathbf{Z}) \to H_T^{-2}(H, \mathbf{Z})$$

is the map $G^{\mathrm{ab}} \to H^{\mathrm{ab}}$ induced by the map

$$G \to H^{\mathrm{ab}}$$

given by

$$g \mapsto \prod_i h_i$$

where $x_i \in G$ are a choice of right coset representatives in G/H and h_i are such that $gx_i = x_{j_i} h_i$ for $h_i \in H$.

For local class field theory, we want to construct the Artin reciprocity isomorphism

$$\rho_{L/K} : K^\times / N_{L/K}(L^\times) \to \mathrm{Gal}(L/K)$$

which after taking inverse limits is an injective map

$$K^\times \to \mathrm{Gal}(K^{\mathrm{sep}}/K)^{\mathrm{sep}}$$

with dense image. To do this we might use the cup product that comes from

$$H_T^{-2} \times H_T^2 \to H_T^0$$

since $H_T^0(\mathrm{Gal}(L/K), L^\times) = K^\times/NL^\times$, $H_T^{-2}(G, \mathbf{Z}) = G^{\mathrm{ab}}$.

Theorem 22.6

If L/K is a finite Galois extension of local fields with group, then there is a canonical isomorphism

$$H^1(G, \mathcal{O}_L^\times) \cong \frac{1}{e}\mathbf{Z}/\mathbf{Z}$$

Proof. Properly normalizing the valuation on L, take the short exact sequence

$$1 \to \mathcal{O}_L^\times \to L^\times \overset{v_L}{\to} \frac{1}{e}\mathbf{Z} \to$$

which induces a long exact sequence on cohomology

$$1 \to \mathcal{O}_K^\times \to K^\times \to \frac{1}{e}\mathbf{Z} \to H^1(G, \mathcal{O}_L^\times) \to H^1(G, L^\times).$$

The last term is trivial by Hilbert 90, and the map from $K^\times \to \frac{1}{e}\mathbf{Z}$ is still the valuation so it is the inclusion $\mathbf{Z} \to \frac{1}{e}\mathbf{Z}$, and thus by exactness we have the desired isomorphism. □

Theorem 22.7

If L/K is an unramified Galois extension of local fields with group G, then $H_T^k(G, \mathcal{O}_L) = 1$ for all $k \in \mathbf{Z}$.

Proof. Since L/K is unramified and Galois, it is cyclic, and thus H_T^1 is trivial. We also know in the cyclic case that the Tate cohomology groups are all equal to H_T^0 or H_T^1, so it remains to show H_T^0 is trivial. But we know

$$H_T^0(G, \mathcal{O}_L^\times) = \mathcal{O}_K^\times/N_{L/K}\mathcal{O}_L^\times = 1$$

by an old homework problem. □

Theorem 22.8

If L/K is unramified of degree $n = [L:K]$, there is a canonical isomorphism

$$H_T^k(G, L^\times) \cong \begin{cases} 1, & k \equiv 1 \mod 2 \\ \frac{1}{n}\mathbf{Z}/\mathbf{Z}, & k \equiv 0 \mod 2 \end{cases}.$$

Proof. Hilbert 90 gives the result in odd dimension (again using the 2-periodicity of the Tate cohomology groups for cyclic groups). For H_T^2 (though we technically did H_T^0 on the homework), again take the short exact sequence

$$1 \to \mathcal{O}_L^\times \to L^\times \to \mathbf{Z} \to 0$$

and part of its induced long exact sequence

$$H^2(G, \mathcal{O}_L^\times) \to H^2(G, L^\times) \to H^2(G, \mathbf{Z}) \to 1.$$

From the previous theorem, $H^2(G, \mathcal{O}_L^\times) = 1$, and on the homework we will see that

$$H^2(G, \mathbf{Z}) \cong H^2(G, \mathbf{Q}/\mathbf{Z}) \cong \frac{1}{n}\mathbf{Z}/\mathbf{Z},$$

so putting this into the exact sequence gives the desired automorphism. □

The isomorphism $H^2(G, L^\times) \cong \frac{1}{n}\mathbf{Z}/\mathbf{Z}$ is the **invariant map**, denoted $\mathrm{inv}_{L/K}$. This H^2 is called the **Brauer group**, and it classifies central simple algebras over K.

Remark 22.9. Taking $n = 0$, we see that $H_T^0(G, L^\times) = K^\times / NL^\times \cong \mathbf{Z}/n\mathbf{Z}$.

Lemma 22.10

If $K \subseteq L \subseteq M$ is unramified, then the following diagram commutes:

$$\begin{tikzcd}
H^2(\mathrm{Gal}(L/K), L^\times) \arrow[r, "\sim"] \arrow[d, "\mathrm{Inf}"] & \frac{1}{[L:K]}\mathbf{Z}/\mathbf{Z} \arrow[d] \\
H^2(\mathrm{Gal}(M/L), M^\times \arrow[r, "\sim"] & \frac{1}{[M:K]}\mathbf{Z}/\mathbf{Z}
\end{tikzcd}$$

For infinite extensions L/K, we have an isomorphism

$$\mathrm{Gal}(L/K) = \varprojlim_{F/K \text{ finite subext.}} \mathrm{Gal}(F/K)$$

and if A is a continuous $\mathrm{Gal}(L/K)$-module

$$H^n(\mathrm{Gal}(L/K), A) = \varinjlim_{F/K \text{ finite subext.}} H^n(\mathrm{Gal}(F/K), A^{\mathrm{Gal}(L/F)})$$

where the maps in the direct system are just the inflation maps.

Example 22.11

$H^1(\mathrm{Gal}(L/K), L^\times) = 1$ for all infinite extensions L/K (take the limit of Hilbert 90). Also $H^n(\mathrm{Gal}(L/K), L) = 0$ for all $n \geq 1$, and

$$H^2(\mathrm{Gal}(K^{\mathrm{ur}}/K), (K^{\mathrm{ur}})^\times) = \varinjlim_{n \geq 1} \frac{1}{n}\mathbf{Z}/\mathbf{Z} = \mathbf{Q}/\mathbf{Z}$$

§23 November 19, 2019

Recall that last time we sketched how to do cohomology for infinite extensions:

$$H^n(L|K,A) = \varinjlim_{F|K \text{fin.Gal.}} H^n(F|K, A^{\mathrm{Gal}(F|K)}) = \coprod_{F|K} H^n(F|K, A^{\mathrm{Gal}(F|K)})/\sim,$$

where $\sim$ is generated by the relations $c_1 \sim c_2$ iff $F_1 \supseteq F_2$ and $c_1 = \mathrm{Inf}(c_2)$, where $c_1 \in H^n(F_1/K, A^{\mathrm{Gal}(F_1/K)})$ and $c_2 \in H^n(F_2/K, A^{\mathrm{Gal}(F_2/K)})$.

Remember the Inf-Res exact sequence: if $H^i(F_1/F_2, A^{\mathrm{Gal}(L/F_1)}) = 0$ for $i = 1, \ldots, n-1$, then $\mathrm{Inf} : H^n(F_1/K, \cdots) \to H^n(F_2/K, \cdots)$ is injective. By Hilbert 90, $H^1(L/K, L^\times) = 1$, and the additive version says that $H^n(L/K, L) = 0$ for all $n \geq 1$. So the inflation map

$$H^2(L/K, L^\times) \to H^2(K^{\mathrm{ur}}/K, (K^{\mathrm{ur}})^\times)$$

is injective and corresponds under the $^{-1}$ isomorphisms to the inclusion of $\frac{1}{n}\mathbf{Z}/\mathbf{Z}$ into $\mathbf{Q}/\mathbf{Z}$.

Now to compute H^2 of K^{sep}. The Inf-Res exact sequence is

$$0 \to H^2(K^{\mathrm{ur}}/K, (K^{\mathrm{ur}})^\times) \overset{\mathrm{Inf}}{\to} H^2(K^{\mathrm{sep}}/K, (K^{\mathrm{sep}})^\times) \overset{\mathrm{Res}}{\to} H^2(K^{\mathrm{sep}}/K^{\mathrm{ur}}, (K^{\mathrm{sep}})^\times).$$

We will compute

$$H^2(K^{\mathrm{sep}}/K^{\mathrm{ur}}, (K^{\mathrm{sep}})^\times) = 1,$$

thus yielding an isomorphism

$$H^2(K^{\mathrm{sep}}/K, (K^{\mathrm{sep}})^\times) \cong H^2(K^{\mathrm{ur}}/K, (K^{\mathrm{ur}})^\times) \cong \mathbf{Q}/\mathbf{Z}.$$

To compute this H^2, the most convenient way uses the identification with the Brauer group.

§23.1 Crash course on the Brauer group and Division algebras

Definition 23.1. Let K be a field. A **division K-algebra** is a finite-dimensional K-algebra D in which every nonzero element has an inverse. It is **central** if its center is equal to K. The **Brauer group** $\mathrm{Br}(K)$ is the set of isomorphism classes of central division K-algebras.

Lemma 23.2

if K is algebraically closed, $\mathrm{Br}(K) = \{K\}$.

Proof. Let D be a division algebra over K, and let $x \in D$. The powers of x in D are linearly dependent, since D is finite-dimensional. Thus x is algebraic over K, and since K is algebraically closed this means $x \in K$. This is because $K(x)$ is a **field** extension of K, since multiplication in there is clearly commutative. □

Lemma 23.3

$\mathrm{Br}(\mathbf{R}) = \{\mathbf{R}, \mathbf{H}\}$ where $\mathbf{H}$ denotes the real quaternion algebra. There is one more non-central algebra over $\mathbf{R}$, namely $\mathbf{C}$.

Proof. later. □

If L/K is a field extension and D is a central division K-algebra, then

$$D \otimes_K L \cong M_{n\times n}(E)$$

for some $n \geq 1$, where E is some central division L-algebra E. Taking D to E defines a map

$$\mathrm{Br}(K) \to \mathrm{Br}(L),$$

which once we define the group structure on the Brauer group will be a homomorphism. So we define the **relative Brauer group**

$$\mathrm{Br}(L/K) = \{D \in \mathrm{Br}(K) : \exists n \geq 1, D \otimes L \cong M_n(L)\}.$$

Example 23.4

$\mathbf{H} \otimes_{\mathbf{R}} \mathbf{C} = M_{2\times 2}(\mathbf{C})$, so

$$\mathrm{Br}(\mathbf{C}/\mathbf{R}) = \mathrm{Br}(\mathbf{R}).$$

More generally, $\mathrm{Br}(K^{\mathrm{sep}}/K) = \mathrm{Br}(K)$.

Corollary 23.5

For any central division K-algebra D, actually $\dim_K D = n^2$ for some positive integer n.

Proof. $D \otimes_K K^{\mathrm{sep}} \cong M_n(K^{\mathrm{sep}})$ for some n, and the dimension of this over K^{sep}, which is n^2, is the same as $\dim_K D$. □

Remark 23.6. The maximal subfields L of D all have degree $n = \sqrt{\dim_K D}$.

D **splits** (this just means $D \otimes_K L = M_{n\times n}(L)$) over any such maximal subfield L.

Definition 23.7. The group structure on $\mathrm{Br}(K)$ is defined as follows: The product of D and D' is the central division K-algebra E such that

$$D \otimes_K D' = M_{n\times n}(E).$$

Example 23.8

K is the identity element in $\mathrm{Br}(K)$, since tensoring by it doesn't do anything.

Example 23.9

$\mathbf{H} \otimes_{\mathbf{R}} \mathbf{H} \cong M_{4\times 4}(\mathbf{R})$ (since we know E must be $\mathbf{R}$ from the fact that the Brauer group has size 2, and the dimension must be 16).

The only perhaps non-obvious part of proving the Brauer group is a group is showing the existence of inverses. To do that, we define the **opposite** K-algebra $D^{\mathrm{opp}} \in \mathrm{Br}(K)$, which is the same as an additive group but the multiplication is defined so that the new xy is what yx used to be.

Theorem 23.10 (Wedderburn)
$\operatorname{Br}(\mathbf{F}_q) = \{\mathbf{F}_q\}$.

Proof sketch. All maximal subfields of $D \in \operatorname{Br}(K)$ have degree n, so they're all isomorphic as fields to $\mathbf{F}_{q^n}$. By Noether–Skolem, they are all conjugate, and hence $D^\times$ is a union of conjugates of a single subgroup, which we saw on the homework is impossible. □

Theorem 23.11
Let K be a local field. The maximal unramified extension has

$$\operatorname{Br}(K^{\mathrm{ur}}) = \{K^{\mathrm{ur}}\}.$$

Proof. Let $D \in \operatorname{Br}(K^{\mathrm{ur}})$ be nontrivial. Then $\dim_K D = n^2$. The valuation on K extends uniquely to K^{ur}. This valuation then has a unique extension to D, though this is maybe not completely obvious. One way to write it down is as

$$v_K(x) = \frac{1}{n^2} v_K(N_{D/K}(x))$$

where as usual $N_{D/K}(x)$ denotes the determinant of the matrix for left-multiplication by x in D. As usual we can take the valuation ring $\mathcal{O}_D$, and $\mathfrak{p}_D$, generated by a uniformizer π_D of valuation 1, etc.

Moreover, if $x \in D^\times$, then we know that $v_K(x) \in \frac{1}{n}\mathbf{Z}$, since the degree of $K(x)$ divides n and the ramification index divides that. So

$$v_K(D^\times) \subset \frac{1}{n}\mathbf{Z}.$$

Since $\kappa(D) = \mathcal{O}_D/\mathfrak{p}_D$ is a central division $\kappa(K^{\mathrm{ur}}) = \overline{\kappa(K)}$-algebra, it must be equal to $\overline{\kappa(K)}$. It follows that $(1, \pi_D, \dots, \pi_D^{n-1})$ generates D as a K-vector space, which is a contradiction. □

Remark 23.12. If K is a local field and D in its Brauer group of degree n^2, then the ramification index $[v_K(D^\times) : v_K(K^\times)]$ and the residue field degree $[\kappa(D) : \kappa(K)]$ are both equal to n.

§24 November 21, 2019

Today we will explain why the Brauer group is the same as H^2.

Let $D \in \operatorname{Br}(L/K)$ split over L, so that $D \otimes_K L = M_{n\times n}(L)$. Let $\sigma \in \operatorname{Gal}(L/K)$ acts on $D \otimes_K L$ where the action on pure tensors is specified by leaving D fixed and acting on L/K. By Noether–Skolem, the automorphisms of $M_{n\times n}(L)$ are given by conjugation by an element of $GL_n(L)$, and hence this action provides a map

$$\operatorname{Gal}(L/K) \to GL_n(L)$$

(σ gets mapped to a choice of $a_\sigma \in GL_n(L)$ such that the action of σ on $D \otimes L$ corresponds in $M_{n\times n}(L)$ to conjugation by $a\sigma$). So for any $\sigma, \tau \in \operatorname{Gal}(L/K)$ and $x \in D \otimes L$,

$$\sigma\tau(x) = a_{\sigma\tau} x a_{\sigma\tau}^{-1} = a_\sigma a_\tau x a_\tau^{-1} a_\sigma^{-1}$$

so the element

$$b_{\sigma,\tau} := a_{\sigma\tau} a_\tau^{-1} a_\sigma^{-1}$$

is in the center of $GL_n(L)$ which is just $L^\times$. The tuple

$$(b_{\sigma,\tau})_{\sigma,\tau\in\operatorname{Gal}(L/K)}$$

is a valid inhomogeneous cocycle in $C^2(\operatorname{Gal}(L/K), L^\times)$ which is well-defined up to 2-coboundaries. This is how the map $\operatorname{Br}(L/K) \to H^2(\operatorname{Gal}(L/K), L^\times)$ is defined.

Remark 24.1. For local fields, the isomorphism $\operatorname{Br}(K) \cong \mathbf{Q}/\mathbf{Z}$ can be described as follows: Let $D \in \operatorname{Br}(K)$ be of dimension n^2. Let $L \subseteq D$ be a maximal subfield, which must have degree n. We can also choose L to be unramified. Because D is central, by Noether–Skolem, the K-automorphisms of $L \subseteq D$ are given by conjugation by some element of $D^\times$, unique up to multiplication by $y \in Z(D^\times) = K^\times$. Since L/K is unramified, we should be concerned with the Frobenius automorphism $\varphi_{L/K}$. There is some $y \in D^\times$ such that $\varphi_{L/K}(x) = yxy^{-1}$, and $D \in \operatorname{Br}(K)$ corresponds to $v_K(y) \in \mathbf{Q}/\mathbf{Z}$ (remember we had a canonical way to extend the valuation to D).

To go the other way, it suffices to see how to go from $1/n$ to a division ring. Let L/K be the degree-n unramified extension of K. Define the K-algebra

$$D = \bigoplus_{i=0}^{n-1} Le_i$$

(here L is not necessarily in the center) with multiplication satisfying

$$e_i x e_i^{-1} = \varphi_q^i(x)$$

and

$$e_i e_j = \begin{cases} e_{i+j}, & i+j \le n-1 \\ \pi_K \cdot e_{i+j-n}, & i+j \ge n \end{cases}$$

Lemma 24.2

For a degree n extension L/K of local fields, the diagram

$$\begin{tikzcd} \operatorname{Br}(K) \arrow[r, "\sim"] \arrow[d] & \mathbf{Q}/\mathbf{Z} \arrow[d, "\times n"] \\ \operatorname{Br}(L) \arrow[r, "\sim"] & \mathbf{Q}/\mathbf{Z} \end{tikzcd}$$

commutes. Thus, $\operatorname{Br}(L/K) \cong \frac{1}{n}\mathbf{Z}/\mathbf{Z}$.

Lemma 24.3

Let D be a central division K-algebra of degree n^2 over a local field K. Then D contains every degree-n field extension of K.

Proof. By looking at the definition of $\mathrm{Br}(K) \cong \mathbf{Q}/\mathbf{Z}$, we see that D is an n-torsion element, and hence for any degree n extension L/K, $D \in \mathrm{Br}(L/K)$, and thus D splits over L. By something we didn't prove I think, this means $L \subseteq D$. □

§24.1 Back to Galois cohomology

Let L/K be a degree n extension of local fields. The Tate cohomology groups were

$$H_T^{-2}(L/K, \mathbf{Z}) = H_T^{-1}(L/K, \mathbf{Z}) \cong \mathrm{Gal}(L/K)^{\mathrm{ab}},$$

$$H_T^0(L/K, L^\times) \cong \frac{K^\times}{N_{L/K}L^\times}$$

and

$$H_T^2(L/K, L^\times) \cong \mathrm{Br}(L/K) = \frac{1}{n}\mathbf{Z}/\mathbf{Z}.$$

So we get a bilinear product map

$$H_T^{-2}(L/K, \mathbf{Z}) \times H_T^2(L/K, L^\times) \to H_T^0(L/K, L^\times)$$

which is just a cup product

$$\mathrm{Gal}(L/K)^{\mathrm{ab}} \times \frac{1}{n}\mathbf{Z}/\mathbf{Z} \to K^\times/N_{L/K}L^\times.$$

It turns out that the homomorphism

$$\mathrm{Gal}(L/K)^{\mathrm{ab}} \to K^\times/N_{L/K}L^\times$$

given by $x \mapsto x \smile (1/n \mod \mathbf{Z})$ is actually the inverse of the CFT isomorphism, but we haven't proven it yet. This was discovered by Tate:

Theorem 24.4 (Tate's Theorem)

Let G be a finite group and A a G-module. Assume that for all $H \subseteq G$, we have $H^1(H, A) = 0$ and $H^2(H, A)$ is cyclic of size $|H|$ (these two things are guaranteed in the number theory situation). Let γ be a generator of $H^2(G, A) = \mathbf{Z}/|G|\mathbf{Z}$. Then the cup product map

$$H_T^r(G, A) \to H_T^{r+2}(G, A)$$

given by $x \mapsto x \smile \gamma$ is an isomorphism.

Proof. The proof is long and starts with a lemma.

Lemma 24.5

Let G be a finite group, and A a G-module. Assume that $H^1(G, A) = 0$ and that $H^2(G, A) = 0$ for all $H \subseteq G$. Then all the Tate cohomology groups vanish.

Proof. In the case where G is cyclic, we already saw that the Tate cohomology groups were all equal to H^1 or H^2 so without needing this for all subgroups we are immediately done. So we might hope to prove it when G is solvable. Let $H \subsetneq G$ be a proper normal subgroup, where we can assume that G/H is cyclic of prime order. By induction, we can assume that $H_T^r(H,A) = 0$ for all r, since H is solvable. So we can use the inflation-restriction exact sequence

$$0 \to H^r(G/H, A^H) \to H^r(G,A) \to H^r(H,A) = 0$$

which tells us that in fact

$$H^r(G/H, A^H) \cong H^r(G,A)$$

so since G/H is cyclic, and we already know the lemma for cyclic groups, it follows that $H^r(G,A) = 0$ (after verifying the hypotheses for G/H) for all $r \geq 1$ (since the inflation-restriction sequence only works in this case). Let's check it directly for $r = 0$. Let $x \in H^0(G,A) = A^G$. We just need to show that x is a norm with respect to G. We know that

$$A^G/N_{G/H}(A^H) = H_T^0(G/H, A^H) = 0,$$

so there exists $y \in A^H$ such that

$$N_{G/H}y = \sum_{g \in G/H} gy = x.$$

And since $A^H/N_H A = H_T^0(H,A) = 0$, we know that this y is an H-norm from A, i.e. there exists a $z \in A$ such that $\sum_{h \in H} hz = y$. Putting these together, we have

$$N_G z = N_{G/H} N_H z = x$$

as desired. This proves that $H_T^0(G,A) = 0$. For $r < 0$ we exploit dimension-shifting. Take the short exact sequence

$$0 \to A' \to \mathbf{Z}[G] \otimes_{\mathbf{Z}} A \to A \to 0$$

where the one in the middle is H-induced for all $H \subset G$, so the corresponding long exact sequence is actually just a bunch of isomorphisms (the Tate cohomology of the middle one is always zero) and we have isomorphisms

$$H_T^{r+1}(H,A') \cong H_T^r(H,A)$$

for all $H \subseteq G$ and $r \in \mathbf{Z}$. In particular, $H_T^1(H,A') \cong H_T^0(H,A) = 0$. Also, $H_T^2(H,A') \cong H_T^1(H,A) = 0$ by the hypothesis of the lemma. Then proceeding by induction, the lemma is proved.

Somehow the case of general groups is a formal consequence of the case of solvable groups. In general, any p-Sylow subgroup G_p of G is solvable, so $H_T^r(G_p,A) = 0$ for all $r \in \mathbf{Z}$. Recall that we had restriction and corestriction maps between $H_T^r(G,A)$ and $H_T^r(G_p,A)$ such that

$$0 = \mathrm{Cor} \circ \mathrm{Res} = [G : G_p] \cdot \mathrm{id}.$$

But if x is a nonzero element of $H_T^r(G,A)$ has order divisible by p, then since p does not divide $[G : G_p]$ this is a contradicton (multiplication by this is supposed to kill x). □

Note that $\mathrm{Res}(\gamma)$ generates $H^2(H,A)$ for any subgroup $H \subseteq G$, because of the fact that restriction composed with corestriction is multiplication by $[G : H]$ (so you get a contradiction if the image of restriction has fewer than $|H|$ elements). Let $\varphi = (\varphi_{\sigma,\tau})_{\sigma,\tau\in G}$ be a cocycle representing $\gamma \in H^2(G, A)$. Consider the G-module

$$A(\varphi) = A \oplus \bigoplus_{\sigma\neq 1} \mathbf{Z}e_\sigma$$

with G-action given by

$$\sigma e_\tau = e_{\sigma\tau} - e_\sigma - \varphi_{\sigma,\tau}$$

where the $e_{\sigma\tau}$ is replaces by $\varphi_{1,1}$ if $\sigma\tau = 1$. This is defined to force (from the definition of the cocycles, recall the "crossed homomorphisms" definition)

$$d^1((e_\sigma)_{\sigma\in G}) = (\varphi_{\sigma,\tau})_{\sigma,\tau\in G}$$

so the map $H^2(G, A) \to H^2(G, A(\varphi))$ sends γ to 0. For any $H \subseteq G$, the exact sequence

$$0 \to I \to \mathbf{Z}[G] \to \mathbf{Z} \to 0$$

induces the long exact sequence

$$0 \to I^H \to \mathbf{Z}[G]^H \to \mathbf{Z} \to H^1(H, I) \to H^1(H, \mathbf{Z}[G]) \to H^1(H, \mathbf{Z}) \to H^2(H, I) \to H^2(H, \mathbf{Z}[G])$$

where we know by hypothesis that $H^1(H, \mathbf{Z}[G]) = H^2(H, \mathbf{Z}[G]) = 0$, so actually

$$H^2(H, I) \cong H^1(H, I) \cong \mathrm{Hom}(H, I) = 0$$

and

$$H^1(H, I) = \mathrm{coker}(\mathbf{Z}[G]^H \to \mathbf{Z}) = \mathbf{Z}/|H|\mathbf{Z}.$$

We can also exploit the short exact sequence

$$0 \to A \to A(\varphi) \to I \to 0$$

where the map $A(\varphi) \to I$ is given by 0 on A and $e_\sigma \mapsto \sigma - 1$. This induces the long exact sequence

$$0 = H^1(H, A) \to H^1(H, A(\varphi)) \to H^1(H, I) \to H^2(H, A) \to H^2(H, A(\varphi)) \to H^2(H, I) = 0$$

where we also know $H^1(H, I) = \mathbf{Z}/|H|\mathbf{Z}$ and $H^2(H, A) \cong \mathbf{Z}/|H|\mathbf{Z}$ is generated by $\mathrm{Res}(\gamma)$. The image of $\mathrm{Res}(\gamma)$ in $H^2(H, A(\varphi))$, by the compatibility with all these maps of Res, is just Res of the image of γ in $H^2(G, A(\varphi))$, which must be 0 (we proved this earlier in this proof). So the map $H^2(H, A) \to H^2(H, A(\varphi))$ is zero, and hence

$$H^2(H, A(\varphi)) = 0.$$

Moreover, the surjectivity of the map $H^1(H, I) \to H^2(H, A)$ proves it is an isomorphism since both have size $|H|$, and thus

$$H^1(H, A(\varphi)) = 0$$

since it is the kernel. So the lemma proves that $H_T^r(G, A(\varphi)) = 0$ for all $r \in \mathbf{Z}$.

Now we take some more exact sequences, namely

$$0 \to A \to A(\varphi) \to I \to 0$$

which by the long exact sequence gives $H_T^{r+2}(G, A) \cong H_T^{r+1}(G, I)$, and

$$0 \to I \to \mathbf{Z}[G] \to \mathbf{Z} \to 0$$

which by the long exact sequence gives $H_T^{r+1}(G, I) \cong H_T^r(G, \mathbf{Z})$. So in the end

$$H_T^{r+2}(G, A) \cong H_T^r(G, A)$$

which you can check is the same map that is asked for in the statement of the theorem. □

How to determine D from $R = M_n(D)$? Let $I \neq 0$ be a minimal right R-ideal. Then $I \cong D^n$ (the space of row vectors), and $D \cong \mathrm{End}(I)$ where this isomorphism is given by multiplication on the left by $x \in D$.

§25 November 26, 2019

Last time we talked about Tate's theorem on Galois cohomology, which we had proved modulo the proof of a certain lemma. I've filled it in in yesterday's notes.

§25.1 The local existence theorem

Let K be a local field. Recall that the class field correspondence is supposed to be a bijection: for any finite index open subgroup U of $K^\times$, there is supposed to be a unique abelian L/K with $NL^\times = U$. This is called the "local existence theorem."

Lemma 25.1

If $U \subseteq K^\times$ is a norm subgroup, then so is any subgroup containing U.

Proof. Let $U = N_{L/K}(L^\times)$, and L/K abelian. The Artin reciprocity map provides an isomorphism

$$\theta_{L/K} : K^\times / N_{L/K} L^\times \to \mathrm{Gal}(L/K)$$

so if we let $H = \theta_{L/K}(V \mod U)$ it's clear that $M := L^H$ has norm group V. □

Lemma 25.2

If $U_1, U_2 \subseteq K^\times$ are norm subgroups, then so is $U_1 \cap U_2$.

Proof. From the basic properties of the Artin map it's clearly the norm subgroup corresponding to the compositum of the class fields for U_1 and U_2. □

Now we are ready to prove the local existence theorem

Theorem 25.3

Every open subgroup of finite index in $K^\times$ is a norm subgroup.

Proof. Since $K^\times \cong \mathcal{O}_K^\times \times \mathbf{Z}$, any finite-index open subgroup of it must contain a power of a uniformizer, so

$$U \supseteq U_K^{(r)} \cdot \pi_K^s$$

for some positive integers s, r. But by Fermat's little theorem,

$$U_K^{(r)} \supseteq (\mathcal{O}_K^\times)^{q^r - 1}$$

which means that we actually have a further containment of the $n = (q^r - 1)$-th powers

$$U \supseteq U_K^{(r)} \cdot \pi_K^s \supseteq (K^\times)^n.$$

This means (by one of the lemmas above) it suffices to show that the group of n-th powers contains a norm subgroup. When the characteristic of K does not divide n and $\zeta_n \in K$, we are immediately done via Kummer theory (you can obtain L by adjoining the n-th roots of everything in $K^\times$).

Now in the general case (still assuming that the characteristic of K doesn't divide n), we need to adjoin the appropriate root of unity ζ_n. Let $K' = K(\zeta_n)$ and (using the above thing) take L'/K' to be the class field for $(K'^\times)^n$. Let M/K be the Galois closure of L' over K. This is a class field for a subgroup contained inside the n-th powers (by the transitivity of the norm) so we are done by the lemma. □

§25.2 Class field theory

So far we have a bunch of statements for class field theory.

Theorem 25.4 (Finite (i.e. unramified) class field theory)

The Frobenius map $\theta_k : \mathbf{Z} \to \operatorname{Gal}(\overline{\mathbf{F}}_q/\mathbf{F}_q)$ has dense image. Let $k = \mathbf{F}_q$. For any finite extension ℓ/k, we have a commutative diagram

$$\begin{tikzcd}
\mathbf{Z} \arrow[r] \arrow[d] & \operatorname{Gal}(\overline{\ell}/\ell)) \arrow[d] \\
\mathbf{Z} \arrow[r] & \operatorname{Gal}(\overline{k}/k)
\end{tikzcd}$$

where the first vertical map is just multiplication by n, and induces an isomorphism

$$\mathbf{Z}/n\mathbf{Z} \cong \operatorname{Gal}(\ell/k).$$

Theorem 25.5 (local Artin reciprocity)

Let K be a local field. Then we have a reciprocity map

$$K^\times \to \operatorname{Gal}(K^{ab}/K)$$

with dense image. Also, the reciprocity map takes the higher unit groups to the upper-indexed higher ramification groups. Uniformizers in $K^\times$ get mapped to Frobenii in $\operatorname{Gal}(K^{ab}/K)$. If k denotes the residue field of K, then the following diagram commutes:

$$\begin{tikzcd}
K^\times \arrow[r] \arrow[d] & \operatorname{Gal}(K^{\mathrm{ab}}/K) \arrow[d] \\
\mathbf{Z} \arrow[r] & \operatorname{Gal}(\overline{k}/k) \cong \hat{\mathbf{Z}}
\end{tikzcd}$$

(corresponding to the fact that the unramified part of the Artin map is the one given by the v_k-power of the Frobenius; the $\operatorname{Gal}(\overline{k}/k)$ is really code for $\operatorname{Gal}(K^{\mathrm{ur}}/K)$). In general, for finite extensions abelian extensions L/K, the diagram

$$\begin{tikzcd}
L^\times \arrow[r] \arrow[d] & \operatorname{Gal}(L^{\mathrm{ab}}/L) \arrow[d] \\
K^\times \arrow[r] & \operatorname{Gal}(K^{ab}/K)
\end{tikzcd}$$

commutes and induces an isomorphism $K^\times/N_{L/K}(L^\times) \to \operatorname{Gal}(L/K)$.

I discussed global CFT in section, and we'll talk about it briefly in the last lecture.

§26 December 3, 2019

First, we confess to lying about the definition of a local field. In fact, we have only been working with *nonarchimedean local fields*. There are two more (archimedean) ones, namely $\mathbf{R}$ and $\mathbf{C}$. The maximal abelian extensions of $\mathbf{R}$ and $\mathbf{C}$ are both $\mathbf{C}$, and in each case we want a surjective map

$$K^\times \to \operatorname{Gal}(\overline{K}/K)$$

et cetera, and the kernel has index 2 when $K = \mathbf{R}$, so it is actually $\mathbf{R}_{>0}$ and the Artin map for $\mathbf{R}$ is just the reduction to $\mathbf{R}^\times/\mathbf{R}_{\geq 0} = \mathbf{Z}/2\mathbf{Z}$. For $\mathbf{C}$ it is just the constant map to the trivial group.

The local fields are important because they arise as completions of global fields.

Definition 26.1. A **global field** is either

(a) A number field (finite extension of $\mathbf{Q}$)

(b) A global function field (finite extension of $\mathbf{F}_q(T)$).

A global field K comes with a ring of integers $\mathcal{O}_K$, which is the integral closure of $\mathbf{Z}$ if K is a number field, and the integral closure of $\mathbf{F}_q[T]$ if K is a global function field.

For each prime $\mathfrak{p}$ in $\mathcal{O}_K$, we get a local field $K_\mathfrak{p}$ (by taking completions with respect to $v_\mathfrak{p}$) with an embedding $K \to K_\mathfrak{p}$.

Definition 26.2. A **place** of K is either an embedding of K into $\mathbf{R}$ or $\mathbf{C}$, or $K \to K_\mathfrak{p}$ for a prime $\mathfrak{p}$.

As usual, each place v comes with a *norm*

$$\|\cdot\|_v : K \to \mathbf{R}_{\geq 0}$$

which for $v = v_\mathfrak{p}$ is given by

$$\|x\|_\mathfrak{p} = (|\kappa(\mathfrak{p})|)^{-v_\mathfrak{p}(x)},$$

and for embeddings $v : K \to \mathbf{R}$ is given by $|v(x)|$, and for embeddings $v : K \to \mathbf{C}$ is given by $|v(x)|^2$. The reason for the way we have normalized and squared (even though this makes the triangle inequality fail), is to force a *product formula* to hold, namely

Theorem 26.3

For any $x \in K^\times$, $\prod_v \|x\|_v = 1$.

Example 26.4

Let $K = \mathbf{Q}$. Then there is one place for every rational prime, and one infinite place v corresponding to the standard embedding of $\mathbf{Q}$ into $\mathbf{R}$. If $x = \pm \prod_p p^{a_p}$, then

$$\|x\|_p = p^{-a_p}$$

so $\prod_p \|x\|_p = \prod_p p^{-a_p} = \|x\|_\infty^{-1}$ which means the product formula holds for $\mathbf{Q}$.

Actually the proof of the product formula in general relies mostly on the above example.

Definition 26.5. Let K be a global field. The **ring of adeles** on K is

$$\mathbb{A}_K = \prod_v{}' K_v = \left\{ (x_v)_v \in \prod K_v : x_v \in \mathcal{O}_v \text{ for almost all } v \right\}.$$

N.B. $\mathcal{O}_v$ is not well-defined for v infinite.

Definition 26.6. The **idèle group** is

$$\mathbb{A}_K^\times = \prod_v{}' K_v^\times = \left\{ (x_v)_v \in \prod K_v^\times : x_v \in \mathcal{O}_v^\times \text{ for almost all } v \right\}.$$

The adeles and idéles also come with a topology, with a WARNING: the topology on the idéles should NOT be the same as the subspace topology for them as a subset of the adeles. The topology on $\mathbb{A}_K$ is given by the basis of open sets

$$\prod_v U_v,$$

where all $U_v \subset K_v$ are open and all but finitely many of them are equal to $\mathcal{O}_v$.

Similarly, we can define a topology on $\mathbb{A}_K^\times$ where you replace K_v with $K_v^\times$ and $\mathcal{O}_v$ with $\mathcal{O}_v^\times$.

The embeddings $K \to K_v$ give rise to diagonal embeddings $K \to \mathbb{A}_K$ and $K^\times \to \mathbb{A}_K^\times$. This map is useful for studying local-global principles (we used it in section to prove a local-global principle for norms in cyclic extensions).

The product formula implies that $K^\times$ is discrete in $\mathbb{A}_K^\times$. Of importance to class field theory is the quotient group (equipped with the quotient topology)

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

which is called the **idèle class group**. An automorphism σ of K induces automorphisms of $\mathbb{A}_K$ and $\mathbb{A}_K^\times$ (it permutes the places though), and this action agrees with the action of σ on $K^\times \subseteq \mathbb{A}_K^\times$. As a result, we actually have an action on C_K. Most importantly, if L/K is a finite Galois extension, there is a norm map

$$N_{L/K} : C_L \to C_K$$

given by taking the product of the actions of all $\sigma \in \operatorname{Gal}(L/K)$.

In fact, it's possible to do this without the Galois restriction. For any finite extension L/K, we have

$$\mathbb{A}_L \cong \mathbb{A}_K \otimes_K L$$

so multiplication by x can be written down as a matrix with entries in $\mathbb{A}_K$ and we can take its determinant.

$$\begin{tikzcd}
C_L \arrow[r] \arrow[d, "N_{L/K}"] & \mathrm{Gal}(L^{\mathrm{ab}}/L) \arrow[d] \\
C_K \arrow[r] & \mathrm{Gal}(K^{\mathrm{ab}}/K)
\end{tikzcd}$$

Theorem 26.7 (Global Class Field Theory)
There is a homomorphism

$$C_K = \mathbb{A}_K^\times / K^\times \overset{\theta_K}{\to} \mathrm{Gal}(K^{ab}/K)$$

with dense image. If K is a number field,

- θ_K is surjective
- θ_K has kernel $\prod_{v:K\to\mathbf{R}} \mathbf{R}_{>0} \times \prod_{v:K\to\mathbf{C}} \mathbf{C}_{>0}$.

If K is a function field, there are no infinite places, and

- θ_K is injective
- θ_K has dense image, namely $\{\sigma \in \mathrm{Gal}(K^{\mathrm{ab}}/K) : \sigma|_{\overline{\mathbf{F}}_q} \in \mathbf{Z} \subseteq \widehat{\mathbf{Z}} = \mathrm{Gal}(\overline{\mathbf{F}}_q/\mathbf{F}_q)\}$.

Moreover, the open finite-index subgroups of C_K correspond bijectively to the Galois groups $\mathrm{Gal}(K^{\mathrm{ab}}/L)$ for finite abelian L/K.

In fact, the global map is compatible with the local ones. The local Artin reciprocity map was originally (I think) defined just by taking the embedding

$$K_v^\times \to \mathbb{A}_K^\times / K^\times$$

and then taking the global reciprocity map (and proving that the image is in the decomposition group of v). This compatibility also means that you could define the global map from the local ones, but it isn't actually clear that the resulting map only depends on the residue mod $K^\times$.

For finite extensions L/K, we have a commutative diagram from which it follows that the reciprocity map is actually an isomorphism

$$C_K / N_{L/K}(C_L) \to \mathrm{Gal}(L/K).$$

§26.1 The Brauer group of a global field

Recall that for a local field K_v, we had a bijection

$$\mathrm{Br}(K_v) \overset{inv}{\to} \mathbf{Q}/\mathbf{Z}.$$

How about the Brauer group of a global field?

Theorem 26.8 (Fundamental exact sequence)
For any global field K, there is an exact sequence

$$0 \to \mathrm{Br}(K) \to \bigoplus_v \mathrm{Br}(K_v) \to \mathbf{Q}/\mathbf{Z} \to 0$$

where the first map is given in each coordinate v by $D \mapsto E_v$ where

$$D \otimes_K K_v \cong M_{K_v}(E_v)$$

and the second map is just the sum of the isomorphisms $\mathrm{Br}(K_v) \to \mathbf{Q}/\mathbf{Z}$.

§26.2 Applications

For any $a, b \in \mathbf{Q}_v^\times$, recall that the degree-2 Hilbert symbol is

$$(a,b)_v = \begin{cases} 1, & ax^2 + by^2 = z^2 \text{ has a nontrivial solution} \\ -1 & \text{else} \end{cases}$$

and that it is multiplicatively bilinear and skew-symmetric. If $p \neq 2$ and $a, b \in \mathbf{Z}_p^\times$ then $(a,b)_p = 1$. If $p \neq 2$ and $a \in \mathbf{Z}_p^\times$, then $(a,b)_p$ is the Legendre symbol $\left(\frac{a}{p}\right)$.

Also, remember that for any $a, b \in \mathbf{Q}_p^\times$, the invariant of the quaternion algebra $(a,b)_{\mathbf{Q}_p}$ is

$$\mathrm{inv}_p(a,b)_{\mathbf{Q}_p} = \begin{cases} 0 \mod \mathbf{Z}, & \text{if } ax^2 + by^2 = z^2 \text{ has a nontrivial solution} \\ 1/2 \mod \mathbf{Z}, & \text{else} \end{cases}$$

so actually this invariant contains the same information as the Hilbert symbol.

Corollary 26.9
Let $a, b \in \mathbf{Q}^\times$. Then by the fundamental exact sequence,

$$\sum_v \mathrm{inv}_v(a,b)_{\mathbf{Q}_v} = 0$$

and thus $\prod_v (a,b)_v = 1$.

Corollary 26.10
Quadratic reciprocity (this uses the above corollary and the computation on problem set 8 problem 2).

Theorem 26.11
The equation

$$y^2 + z^2 = (3 - x^2)(x^2 - 2)$$

has no rational solution.

proof sketch. The Hilbert symbol this with −1 is

$$(3 - x^2, -1)_v (x^2 - 2, -1)_v = (y^2 + z^2, -1)_v = 1.$$

So $(3-x^2, -1)_v = (x^2-2, -1)_v = \pm 1$, and in particular the numbers $a_v = \mathrm{inv}_v(3-x^2, -1)_{\mathbf{Q}_v} = \mathrm{inv}_v(x^2 - 2, -1)_{\mathbf{Q}_v}$ need to sum to zero by the fundamental exact sequence. You can compute that $a_\infty = a_p = 0$ for all odd p, and that $a_2 = 1/2$, which is a contradiction. □

References

[1] S. Bosch. *Algebra: From the Viewpoint of Galois Theory*. Springer, 2018.

[2] S. Lang. *Algebraic number theory*, volume 110 of *Graduate Texts in Mathematics*. Springer Science & Business Media, 2013.

[3] D.A. Marcus. *Number fields*, volume 8 of *Universitext*. Springer, 1977.

[4] J.S. Milne. Class field theory, 1997. URL: http://www.math.lsa.umich.edu/jmilne.

[5] J. Neukirch. *Algebraic number theory*, volume 322 of *Graduate Texts in Mathematics*. Springer Science & Business Media, 2013.

[6] Pierre Samuel. *Théorie algébrique des nombres*. Hermann, 1967.