(Some) perspectives on class field theory

Avi Zeff

This is an expository note on several different formulations of class field theory, starting from elementary motivations in §1. We will generally omit all proofs, and will make no attempt to build up the necessary machinery: instead different sections will be targeted at different audiences, with §1.1 - 1.2 requiring no mathematical background at all (besides basic arithmetic); §1.3 - 3.3 requiring some abstract algebra (groups, rings, ideals, modules, and Galois theory), as well as some light discussion involving concepts from algebraic geometry towards the end of §3.3; and §4 requiring a fair amount of algebraic geometry and topology (e.g. (étale) sheaf theory, ∞-categories, K-theory, topological cyclic homology). None of the material is original; the first three sections are fairly standard (one excellent reference is [4], which many of the precise statements in §3.1 are drawn from), while §4 is drawn almost entirely from [1].

For simplicity, we will mostly stick with the number field case, though the function field case is also handled in the same way in many parts and will be touched on in §3.3.

Contents

1.1 Prime numbers and irreducible polynomials

Let's start with prime numbers: a prime number is a positive integer that is divisible only by two numbers, itself and 1.[1] For example, 5 is prime, while 6 is divisible by 2 and 3 as well as 1 and 6. There is a lot to say about prime numbers, but let's first try to generalize.

In any situation where we can multiply things together, we can check divisibility in a similar way. A particularly important example is polynomials, combinations of constants with a given variable by addition, subtraction, and multiplication (but not division): for example, $4x + 3 \cdot (5 - x)^2$.[2] By expanding everything out, one can always write polynomials as $a_0 + a_1x + a_2x^2 + \cdots + a_nx^n$ for some constants $a_0, a_1, \ldots, a_n$; for example, $4x + 3 \cdot (5 - x)^2$ above is equal to $75 - 26x + 3x^2$. The number n here is called the degree of the polynomial Note that polynomials include the set of integers as the constant polynomials, e.g. $f(x) = 3$; these are the polynomials of degree 0.

One can again talk about divisibility: for example, $x^2 + 3x + 2$ can be written as $(x + 1)(x + 2)$, so it is divisible by both $x + 1$ and $x + 2$. It's again especially interesting to look at the polynomials which aren't divisible by any smaller polynomial other than 1 (or, since we want to allow negative coefficients as well, -1); to avoid confusion with prime numbers, we call these polynomials irreducible. Any linear polynomial $a_0 + a_1x$ must be irreducible unless a_0 and a_1 have a common factor;[3] a nonlinear example is $x^2 + 1$. Another interesting example is $x^2 - 2$; if we allowed all real numbers as coefficients, this would factor as $(x + \sqrt{2})(x - \sqrt{2})$, but since we're only allowing integer coefficients it is irreducible. This is worth remembering: how a polynomial can factor depends in part on what coefficients we allow.

Let's take a brief aside to talk about modular arithmetic. Modular arithmetic is where, for some fixed positive number n, we count from 0 up to $n - 1$ and then start over; in other words n is considered to be the same as 0, so $n + 1$ is the same as 1, and so on. This is sometimes called "clock arithmetic": if $n = 12$ this is very much like how a clock works, counting from 1 to 12 and then starting over: one hour after 12 is 1, not 13, etc.[4] One can also think of this in terms of the remainder after dividing by n: for example, 7 mod 5 is the remainder of 7 after dividing by 5, which is 2; so we say $7 \equiv 2 \pmod 5$.

The case of interest for us is where we work modulo a prime number p; the set of integers modulo p is written as $\mathbb{F}_p$. We can still talk about polynomials in this setting: these are things of the form $a_0 + a_1x + \cdots + a_nx^n$ where each a_i is only defined modulo p, i.e. if we change a_i to $a_i + p$ or $a_i - p$ (or $a_i + 2p$, etc.) the polynomial is considered to be the same. One way to get such a polynomial is to start with a polynomial $f(x)$ with coefficients in the integers, and then take $f(x) \bmod p$. We can do this for different primes p to see different behavior: for example, if $f(x) = 5x^2 + 1$, then $f(x) \equiv 2x^2 + 1 \pmod 3$, while $f(x) \equiv 1 \pmod 5$.

[1] Note that 1 is *not* prime, since it is only divisible by one number—1!

[2] For now, we take all the constants to be integers, so something like $\frac{1}{2} - x$ or πx is not allowed.

[3] We generally don't worry too much about constant factors, since they're easy to pull out; we're more concerned with divisors of degree at least 1.

[4] Whether we count from 0 to $n - 1$ or 1 to n doesn't matter very much; mathematicians like to use 0, while clock faces prefer to say 12:00 than 0:00.

Given that we care about irreducible polynomials, we might ask: if $f(x)$ is an irreducible polynomial over the integers, is $f(x)$ mod p irreducible over $\mathbb{F}_p$? What about the other way around?

1.2 Splitting types

One direction is relatively easy: if $f(x)$ is irreducible mod p, it must also be irreducible over the integers. Indeed, if $f(x)$ were *not* irreducible over the integers, we could write it as $f(x) = g(x)h(x)$ for g and h of degree at least 1; and then that factorization would still be true modulo p, so it couldn't be irreducible modulo p.[5]

The other direction however turns out to be very complicated. Consider $f(x) = x^2 + 1$, which we mentioned is irreducible over the integers. Modulo p, we have $x^2 + 1 \equiv x^2 - 4 = (x + 2)(x - 2) \pmod 5$, so f is *not* irreducible modulo 5 even though it is over the integers. This doesn't always occur, though: e.g. $f(x)$ is still irreducible modulo 7. How can we possibly say which behavior occurs when?

In fact, we could ask for even more information. Given a polynomial $f(x)$ of degree d over the integers, modulo p it might still be irreducible, or it might factor into a product of irreducible polynomials $f(x) \equiv f_1(x) \cdot f_2(x) \cdots f_r(x) \pmod p$. If $d_i = \deg f_i$, the degrees of each factor sum up to the total degree, so $d_1 + \cdots + d_r = d$; the order of the factors doesn't matter, so we might as well assume the d_i are in descending order. Thus given the factorization of f modulo p, we can get out a *partition* of d.

Here a partition of d means a way of writing d as a sum (disregarding order): e.g. the partitions of $d = 4$ are 4, $3 + 1$, $2 + 2$, $2 + 1 + 1$, and $1 + 1 + 1 + 1$. The partition of d associated to the factorization of f modulo p is called the *splitting type* of $f(x)$ modulo p. At the one extreme, if $f(x)$ is still irreducible modulo p then the splitting type is just d; on the other hand if it's completely reducible, splitting into d linear factors, then the splitting type is $1 + 1 + \cdots + 1$ (in this case we say that it splits completely at p).

There are many questions we could ask about splitting types: given $f(x)$ and p, how can one determine what the splitting type will be (other than laboriously factoring $f(x)$ modulo p every time)? For any given $f(x)$, will there be infinitely many primes p with a given splitting type? For a given f of interest, how do the splitting types relate to other questions about prime numbers?

It's perhaps worth dwelling a little more on this last question, if only for motivation. After all, while intriguingly mysterious this question of splitting type seems fairly obscure, and not obviously closely related to more charismatic questions often posed about prime numbers. It turns out however that at least certain such questions are closely related. For example, consider the following problem: which prime numbers can be written in the form $x^2 + y^2$ for integers x, y?

By playing around with some examples, you can work out that the answer is some but

[5] The careful reader will note that strictly speaking one could carefully choose p such that this is not true: for example, $5x^2 + x$ is reducible over the integers (as $x(5x+1)$), but modulo p it's just x, which is irreducible. This is a sort of pathological scenario though, and other than this sort of thing where p divides the leading coefficient it will not occur; if you like you can think of this statement as being true for all sufficiently large p. In general it turns out to be very natural to restrict to polynomials where the leading coefficient a_n is equal to 1; such polynomials are called monic.

not all primes: for example $2 = 1^2 + 1^2$, $5 = 1^2 + 2^2$, $13 = 2^2 + 3^2$, but there is no such expression for 3, 7, or 11. To go further, we can try some modular arithmetic: if $x^2 + y^2 = p$, then certainly $x^2 + y^2$ is divisible by p, but neither x nor y is (since if they were, both would have to be at least p and so $x^2 + y^2$ would be at least $2p^2$, which is clearly greater than p). In other words, $x^2 + y^2 \equiv 0 \pmod{p}$, but x and y are nonzero modulo p. Subtracting, this means that $x^2 \equiv -y^2 \pmod{p}$, or equivalently $(x/y)^2 \equiv -1 \pmod{p}$.[6] Thus in order for p to be of this form, there must exist some z such that $z^2 \equiv -1 \pmod{p}$, i.e. $z^2 + 1 \equiv 0 \pmod{p}$. If $f(x) = x^2 + 1$, this is precisely equivalent to $f(x)$ splitting modulo p: if $f(x) = x^2 + 1 \equiv (x - z_1)(x - z_2) \pmod{p}$, then either z_1 or z_2 satisfy $z^2 + 1 \equiv 0 \pmod{p}$.[7] Thus p can be written in the form $x^2 + y^2$ if and only if $f(x) = x^2 + 1$ splits modulo p; so if we can solve the latter problem we can solve the former.

In fact, this reformulation in terms of whether p is of the right form is useful not just for motivation but also for solving the original splitting question. The right way to formulate this in general is Dedekind's theorem, which turns this sort of question into one that the methods of class field theory can address. For example in the situation above, it turns out that $x^2 + 1$ splits modulo p (or equivalently, p can be written in the form $x^2 + y^2$) if and only if $p \equiv 1 \pmod{4}$.[8][9]

In general, the consequences of class field theory for polynomials over the integers can be summarized as follows: if $f(x)$ is an irreducible polynomial satisfying a certain important condition,[10] then there exists a positive integer c, called the conductor, such that the splitting behavior of $f(x)$ at p is determined by the value of p modulo c. For example if $p \equiv 1 \pmod{c}$ then $f(x)$ will split completely at p, though the converse need not be true.

As we'll see, this is a consequence of the Kronecker–Weber theorem; but first we need to reinterpret our problem in language more accessible to these kinds of tools.

1.3 Dedekind's theorem

The "right" way to reformulate the splitting problem (using some abstract algebra) turns out to be as follows.

Theorem (Dedekind's theorem)**.** *Let $f(x)$ be an irreducible monic*[11] *polynomial with integer coefficients, and $\mathcal{O} = \mathbb{Z}[x]/(f(x))$ be the associated number ring.*[12][13] *Then for any prime p,*

[6] While x/y may not be an integer, it can still make sense modulo p: we define $\frac{1}{y}$ to be the unique element modulo p such that $\frac{1}{y} \cdot y \equiv 1 \pmod{p}$. It takes some work to check that this exists and makes sense for every nonzero y, but it is in fact true.

[7] Checking the converse claim, as well as the fact that $x^2 + 1$ splitting modulo p is not just necessary but also sufficient, is left as an exercise to the reader.

[8] With the exception of $p = 2$, which is a bit of a special case for this example.

[9] This particular claim doesn't actually need all the machinery of class field theory, and was known to Fermat; class field theory gives a way of vastly generalizing this sort of statement.

[10] Namely that the number field generated by the roots of this polynomial is abelian over $\mathbb{Q}$.

[11] Recall this means that the leading coefficient is 1.

[12] Here $\mathbb{Z}$ is the ring of integers $\{\ldots, -2, -1, 0, 1, 2, \ldots\}$.

[13] This means that we adjoin an element x to $\mathbb{Z}$, subject to the condition that $f(x) = 0$. For example, if $f(x) = x^2 + 1$, the condition $x^2 + 1 = 0$ is equivalent to $x^2 = -1$, the defining property of the imaginary unit i; so $\mathcal{O} = \mathbb{Z}[i]$, the set of numbers of the form $a + bi$ where a and b are integers. These are called the Gaussian integers, and are a useful example to keep in mind.

the splitting type of f at p is the same as the splitting type of the ideal (p) in $\mathcal{O}$ when factored as a product of prime ideals: if

$$(p) = \mathfrak{p}_1^{e_1} \cdot \mathfrak{p}_2^{e_2} \cdots \mathfrak{p}_r^{e_r},$$

then

$$f(x) \equiv f_1(x)^{e_1} \cdot f_2(x)^{e_2} \cdots f_r(x)^{e_r} \pmod{p}$$

where the $f_i(x)$ are distinct irreducible polynomials modulo p.

One can also read off the degrees of the polynomial factors from the corresponding prime ideals, but this is a little more intricate.

One detail here which we have not touched on before is the exponents e_i. Generically, we expect that the factorization ought to be into different polynomials/prime ideals; but it can happen that multiple factors are in fact the same, so we combine them into $f_i(x)^{e_i}$; this phenomenon is called ramification, and only occurs at finitely many primes.[14] Dedekind's theorem then tells us that the correspondence between factorization of polynomials modulo p and the ideal (p) into prime ideals also preserves ramification.

It's worth pausing to remark on what we mean by all this: if p is prime, isn't the ideal (p) automatically prime as well? The key is that we're looking at the ideal (p) not in $\mathbb{Z}$, where it is indeed prime, but in $\mathcal{O}$, i.e. $p\mathcal{O}$ rather than $p\mathbb{Z}$, where it may not be. For example, taking $f(x) = x^2 + 1$ again so that $\mathcal{O} = \mathbb{Z} \oplus i\mathbb{Z}$ is the ring of Gaussian integers, if $p = 5$ then $(5) = (2+i) \cdot (2-i)$, so (5) is no longer a prime ideal; instead it factors into two prime ideals of $\mathcal{O}$ that are not themselves defined over the usual integers, and only make sense in $\mathcal{O}$. On the other hand if $p = 7$ then (7) remains prime in $\mathcal{O}$. So one way of understanding Dedekind's theorem is as changing the splitting problem into the study of whether prime ideals remain prime under an extension of number rings.

We can now also relate this interpretation back to the problem from §1.2: if p splits completely as a product of prime ideals $\mathfrak{p}_i$, then the norm of each ideal $\mathfrak{p}_i$ must be p.[15] For example in the case above, $N(2+i) = N(2-i) = 2^2 + 1^2 = 5$. In general for the Gaussian integers the norm of an ideal generated by $a + bi$ is $a^2 + b^2$; so a prime number splits completely in the Gaussian integers if and only if it can be written in the form $a^2 + b^2$. This gives a more abstract way of formulating the equivalence from the previous section; and indeed we can now see how to generalize it already, e.g. a prime number can be written in the form $a^2 + 2b^2$ if and only if it splits completely in $\mathbb{Z}[\sqrt{-2}]$. But of course, we still don't know in general how to determine the factorization of a prime ideal in a number ring; a partial solution to this problem is the essence of class field theory.

2. Class field theory I: ideals

(Note that we'll put off most of the real content of the theorems of class field theory until §3, as the formulation is nicer in the adelic language (for a suitable notion of niceness). In

[14] This is, for example, why the prime 2 behaves a little differently from the rest of the primes for $f(x) = x^2 + 1$, in addition to being even: while for all other primes the factors of $f(x)$ mod p will always be distinct, $x^2 + 1 \equiv (x+1)^2 \pmod{2}$, i.e. 2 is ramified for this polynomial.

[15] The norm of an ideal I of a ring R can be defined as the order of the quotient R/I (when it exists); in this case, we could also think of it as the product of the Galois conjugates. (In general these are special cases of broader definitions which do not in general agree.)

the meantime we'll develop some of the concepts underlying the theorems.)

The main theorems of class field theory can generally be phrased in terms of a "reciprocity map"[16] between two groups, one of which (the class group) classifies ideals in some sense and the other of which (the Galois group) in some sense determines the splitting behavior. Thus in principle once we understand this map we should be able to understand the splitting behavior of a prime ideal by interpreting its image under the reciprocity map.

Before we can discuss the map, we need to understand the source and target. I am going to assume some basic knowledge of Galois theory here, and so won't spend much time talking about the Galois group except as it arises; but let's discuss the class group.

2.1 The class group

Given two ideals I, J of a number ring $\mathcal{O}$, we can multiply them together to get a third ideal IJ; this gives a commutative monoid structure on the set of ideals, but this is not a group since we don't have inverses: for example in $\mathbb{Z}$, we have a multiplicative unit $(1) = \mathbb{Z}$, but for any other ideal (n) (with $n \neq \pm 1$) there is no ideal I such that $I \cdot (n) = (1)$. To fix this, we introduce fractional ideals, which include things such as $\frac{1}{n} \cdot \mathbb{Z}$. More precisely, if $\mathcal{O}$ is our number ring and $K = \operatorname{Frac} \mathcal{O}$ is its field of fractions (a number field, i.e. a finite extension of the rational numbers $\mathbb{Q} = \operatorname{Frac} \mathbb{Z}$), just as one can define ideals of $\mathcal{O}$ to be the $\mathcal{O}$-submodules of $\mathcal{O}$, we define the fractional ideals to be the $\mathcal{O}$-submodules of K.[17]

Now we can hope to invert ideals, with the obvious exception of (0): for any ideal I, $(0) \cdot I = (0)$. But it turns out that if we remove the zero ideal, everything works: the set of nonzero fractional ideals of $\mathcal{O}$ forms a group under multiplication, which we call $\mathcal{I}$.[18]

Now, $\mathcal{I}$ is really too large a group to work with: for convenience, we like to have our groups finite, or at least finitely generated, while for $\mathcal{O} = \mathbb{Z}$ just the principal fractional ideals $q \cdot \mathbb{Z}$ for $q \in \mathbb{Q}^\times$ give an embedding of the infinitely generated group $\mathbb{Q}^\times \hookrightarrow \mathcal{I}$. In fact this is true in general: the group of principal fractional ideals $q^\times \cdot \mathcal{O}$ for $q \in K^\times$ gives an embedding $K^\times \hookrightarrow \mathcal{I}$. This means that $\mathcal{I}$ must be inconveniently large, but it also gives a hint as to how we could find something more manageable that preserves most of the interesting structure: if the principal ideals are relatively easy to understand, then the rest of the content of $\mathcal{I}$ should be essentially described by the cokernel $\mathcal{I}/K^\times$. It turns out that this group is always finite: we call it the class group of K (or of $\mathcal{O}$), and write it as $\operatorname{Cl}(K)$.

How should we understand $\operatorname{Cl}(K)$? Well, by construction it is the fractional ideals of $\mathcal{O}$ modulo the principal ideals, so its nontrivial elements in some sense describe the non-principal fractional ideals. The fact that we're quotienting by all of $K^\times$ instead of just nonzero elements of $\mathcal{O}$ suggests that actually this description should hold for genuine ideals, rather than just fractional ones, i.e. $\operatorname{Cl}(K)$ describes the non-principal ideals of $\mathcal{O}$. This is true: in particular, $\operatorname{Cl}(K)$ is trivial (i.e. $|\operatorname{Cl}(K)| = 1$) if and only if every ideal of $\mathcal{O}$ is principal.

To see why this is an interesting tool, let's turn for a moment to the notion of unique factorization. For the usual integers $\mathbb{Z}$, the fundamental theorem of arithmetic says that

[16] Often called the Artin reciprocity map, for Emil Artin.

[17] Strictly speaking this differs slightly from the broader definition; we are using that $\mathcal{O}$ is a number ring, so in particular a Dedekind domain, meaning that its ideals uniquely factor into products of prime ideals.

[18] Again we are relying on the fact that $\mathcal{O}$ is a Dedekind domain; this isn't always true otherwise.

every integer has a unique factorization (up to order and the units $\pm 1 \in \mathbb{Z}^\times$) as a product of prime numbers. One might hope that this will hold for our number rings $\mathcal{O}$ as well; for our standard example of the Gaussian integers, this turns out to be true.

However, it is not true in general. Consider $\mathcal{O} = \mathbb{Z}[\sqrt{-5}]$. In this ring, similarly to the Gaussian integers, a prime p splits if and only if it can be written in the form $a^2 + 5b^2$ for integers a and b; so we can check easily that e.g. 2 and 3 do not split (they are instead said to be inert). Therefore we expect 2 and 3 to still be prime; so $2 \cdot 3 = 6$ should be the only factorization of 6 into primes, if unique factorization is to hold. But note that we can also write $5 = (1 + \sqrt{-5})(1 - \sqrt{-5})$, and certainly neither factor is divisible by 2 or 3; in fact both of these factors are also "prime" in the sense of having no nontrivial factors. So 6 has two distinct prime factorizations in $\mathbb{Z}[\sqrt{-5}]$: it is not a unique factorization domain!

However, all of our number rings $\mathcal{O}$ are Dedekind domains, the key feature of which is that the *ideals* satisfy unique prime factorization: every *ideal* factors uniquely as a product of prime ideals. If every ideal is principal, then this is equivalent to regular unique factorization; but if there are non-principal ideals, this can be different. In fact it is not too hard to check that a Dedekind domain $\mathcal{O}$ will be a unique factorization domain if and only if it is a principal ideal domain. So since the class group measures the failure of $\mathcal{O}$ to be a principal ideal domain, it equivalently measures the failure of $\mathcal{O}$ to have unique factorization—which is pretty clearly of arithmetic interest![19]

2.2 The Frobenius elements

We can now try and say what the reciprocity map should be—although as we'll see there are reasons to be dissatisfied with this formulation. It'll be useful to fix an extension L/K of number fields, which induces a corresponding extension of number rings $\mathcal{O}_L/\mathcal{O}_K$. Given a prime ideal $\mathfrak{p}$ of $\mathcal{O}_K$, we want to know how it splits in $\mathcal{O}_L$. To start with, as $\mathfrak{p}$ is an ideal we can view it as an element of $\mathcal{I}_K$, and so it has an image $[\mathfrak{p}]$ in $\mathrm{Cl}(K)$. Indeed, since the prime ideals (freely) generate $\mathcal{I}_K$ it suffices to define our reciprocity map on prime ideals.

On the other side, we said we expect to have a Galois group; the natural one is $\mathrm{Gal}(L/K)$. Now we need an important condition: since $\mathrm{Cl}(K)$ is always abelian, we'd like the Galois group to be abelian as well (for symmetry, if nothing else), and indeed we won't be able to say much about non-abelian Galois groups; so we require that L/K have abelian Galois group.

The automorphisms $\sigma \in \mathrm{Gal}(L/K)$ of L preserve the ring of integers $\mathcal{O}_L$, so we may view them as automorphisms of $\mathcal{O}_L$ fixing $\mathcal{O}_K$. For each prime ideal $\mathfrak{p}$ of $\mathcal{O}_L$ and prime ideal $\mathfrak{q}$ of $\mathcal{O}_K$ lying over $\mathfrak{p}$,[20] quotienting gives an automorphism of $\mathcal{O}_L/\mathfrak{q}$ fixing the subring $\mathcal{O}_K/\mathfrak{p}$ provided that the initial automorphism preserves $\mathfrak{q}$. Since $\mathfrak{p}$ and $\mathfrak{q}$ are maximal ideals, both of the quotient rings are fields; in fact they are finite fields, called the residue fields of L and K respectively at $\mathfrak{q}$ and $\mathfrak{p}$. The construction above gives, for every choice of $\mathfrak{p}$ and $\mathfrak{q}$, a map from the stabilizer of $\mathfrak{q}$ in $\mathrm{Gal}(L/K)$ to $\mathrm{Gal}((\mathcal{O}_L/\mathfrak{q})/(\mathcal{O}_K/\mathfrak{p}))$. Provided $\mathfrak{q}$ over $\mathfrak{p}$ is unramified, this is even an isomorphism.

[19] After all, unique factorization for $\mathbb{Z}$—equivalent to the statement that $\mathrm{cl}(\mathbb{Q}) = \{1\}$—is called the fundamental theorem of arithmetic for a reason.

[20] Meaning that $\mathfrak{q} \cap \mathcal{O}_K = \mathfrak{p}$.

The Galois theory of finite fields is well understood: if $k \subset \ell$ is an extension of finite fields, then $\mathrm{Gal}(\ell/k)$ is cyclic (of order depending on the degree of the extension). It has a canonical generator: if k has order q, then $x \mapsto x^q$ is an automorphism of ℓ which fixes k (by Fermat's little theorem).[21] The preimage of this generator under the isomorphism above with the stabilizer of $\mathfrak{q}$ is called the Frobenius element of $\mathfrak{q}$, written as $\sigma_{\mathfrak{q}}$. It is not too hard to check that the dependence on $\mathfrak{q}$, as opposed to just $\mathfrak{p}$, is fairly light: for different choices of $\mathfrak{q}$ over $\mathfrak{p}$, the resulting elements $\sigma_{\mathfrak{q}}$ of $\mathrm{Gal}(L/K)$ will be conjugate.

Now, our situation is particularly good because we have assumed that $\mathrm{Gal}(L/K)$ is abelian: in abelian groups, conjugate elements are equal.[22] Therefore the Frobenius elements do not depend on $\mathfrak{q}$ at all but only on $\mathfrak{p}$, and so we can safely call them $\sigma_{\mathfrak{p}}$.

So we have constructed, for each prime ideal $\mathfrak{p}$ of $\mathcal{O}_K$, an element $\sigma_{\mathfrak{p}}$ in $\mathrm{Gal}(L/K)$. Since the prime ideals freely generate $\mathcal{I}_K$, this gives a map $\mathcal{I}_K \to \mathrm{Gal}(L/K)$, which we might hope to have good properties; e.g. perhaps it descends to the class group (which recall we defined as a quotient of $\mathcal{I}_K$), and in an ideal world might even be an isomorphism.

Unfortunately such hopes are doomed to disappointment, at least for the moment. One obvious issue is that the left-hand side depends only on K, while the right-hand side depends on both L and K; so to say anything about the Galois groups in general, we would need some further data on the left. Another is most easily demonstrated in the simplest case where $K = \mathbb{Q}$: here, the fundamental theorem of arithmetic says that $\mathrm{Cl}(K)$ is the trivial group, while $\mathrm{Gal}(L/\mathbb{Q})$ can certainly be very nontrivial (and always will be for $L \neq \mathbb{Q}$), so if the reciprocity map factors through the class group then it can't possibly tell us anything interesting!

There are several ways of fixing these issues, in particular by introducing the notion of a modulus, which makes the class group much more flexible; for a suitable modulus $\mathfrak{m}$ one finds that the reciprocity map $\mathrm{Cl}_{\mathfrak{m}}(K) \to \mathrm{Gal}(L/K)$ is surjective, and its kernel can be described explicitly in terms of the extension L/K. Even with trivial modulus, for general K there are still highly nontrivial statements: one is the existence of the Hilbert class field H of K, which is the maximal abelian unramified extension of K; it satisfies the property that $\mathrm{Cl}(K) \simeq \mathrm{Gal}(H/K)$. Since $\mathrm{Cl}(\mathbb{Q})$ is trivial, the Hilbert class field of $\mathbb{Q}$ is just $\mathbb{Q}$; but in general H has many interesting properties.

Rather than delve into the details of the correct ideal-theoretic formulation, though, now that we have the key tools I want to reformulate everything in terms of adeles, which allow for much cleaner statements at the cost of introducing some new machinery—but this machinery turns out to be very useful in number theory. First, though, there are two natural questions I want to address: if the class group is too insensitive to be the right source for the reciprocity map without modification, why is it a useful object to define at all? What would a reciprocity map (from a suitably modified class group) would tell us about prime splitting?

[21] Although this would not be a ring homomorphism in characteristic 0, it is in characteristic p!

[22] One might guess that this is why we restrict to abelian extensions; but actually here this is only a matter of convenience, and it is entirely possible to talk about conjugacy classes of Frobenius elements for non-abelian extensions. The issue is rather that the reciprocity map can only describe the abelianization of the Galois group, as we'll see in §3.1.

2.3 Class number formula

One motivation for discussing the class group in its original form is its property of measuring the failure unique factorization; another is the way it generalizes to the adelic setting, as we'll see in §3.1. However, it is additionally a deep invariant of number fields in its own right, with various mysterious incarnations. One such is the class number formula.

The first ingredient in the class number formula is the Dedekind zeta function. The classical example is the Riemann zeta function, which for $s \in \mathbb{C}$ with real part greater than 1 can be defined as follows:

$$\zeta(s) = \sum_{n \geq 1} n^{-s}.$$

Euler's product formula (which is essentially an analytic formulation of the fundamental theorem of arithmetic) states that we can also write

$$\zeta(s) = \prod_{p} (1 - p^{-s})^{-1}.$$

Many of the early theorems of analytic number theory (e.g. the prime number theorem) come from analyzing this relationship: the first definition gives $\zeta(s)$ in terms of positive integers, while the second in terms of prime numbers, so by carefully exploiting this relationship (using a lot of complex analysis) we can estimate e.g. the frequency of prime numbers, given the frequency of positive integers (which is of course much simpler).

We can make similar definitions over any number field K. We no longer have unique factorization in general, but we do have unique factorization of ideals; so we should replace the positive integers above by the ideals of $\mathcal{O}_K$, and the primes by the (nonzero) prime ideals of $\mathcal{O}_K$. Thus we define the Dedekind zeta function of K to be

$$\zeta_K(s) = \sum_{\mathfrak{n}} N(\mathfrak{n})^{-s} = \prod_{\mathfrak{p}} (1 - N(\mathfrak{p})^{-s})^{-1},$$

where $\mathfrak{n}$ ranges over ideals of $\mathcal{O}_K$ and $\mathfrak{p}$ over nonzero prime ideals, and $N(-)$ is the norm. Thus $\zeta_{\mathbb{Q}}(s)$ is the classical Riemann zeta function $\zeta(s)$.

One could hope to study the distribution of prime ideals using the Dedekind zeta function, similar to classical analytic number theory; indeed this is possible, but now we need as input an estimate of the number of ideals of norm at most x. When our ideals are just positive integers, i.e. for $K = \mathbb{Q}$, this is trivial: there are x (or $\lfloor x \rfloor$) positive integers less than or equal to x. But for ideals in larger number rings it is much more complicated.

Nevertheless, it is possible to get good bounds on the distribution of ideals, using a theory often called the "geometry of numbers"; this has consequences for the distribution of prime ideals, which we will not get into too much, but also lets us reproduce analogues of other important features of the Riemann zeta function. For instance, for s near 1 one can estimate

$$\zeta(s) = \sum_{n \geq 1} n^{-s} \approx \int_1^\infty t^{-s}\, dt = \frac{1}{s-1},$$

and with a little more work one can verify that this is a good approximation: $\zeta(s) = \frac{1}{s-1} + O(1)$ near $s = 1$, i.e. $\zeta(s)$ extends to a meromorphic function on a neighborhood

surrounding $s = 1$[23] with a simple pole at $s = 1$ with residue 1; equivalently, its Laurent series near $s = 1$ is of the form

$$\zeta(s) = \frac{1}{s-1} + c_0 + c_1(s-1) + c_2(s-1)^2 + \cdots .$$

One can also show, using rough bounds on the count of ideals, that $\zeta_K(s)$ has a simple pole at $s = 1$; however its residue $c_{-1} = \lim_{s\to 1^+}(s-1)\zeta_K(s)$, appearing in the Laurent expansion

$$\zeta_K(s) = \frac{c_{-1}}{s-1} + c_0 + c_1(s-1) + c_2(s-1)^2 + \cdots ,$$

is no longer necessarily 1. Instead, its value is predicted by the class number formula:

Theorem (Class number formula)**.** *For any number field K, $\zeta_K(s)$ extends to a meromorphic function on a neighborhood of $s = 1$ with a simple pole at $s = 1$ with residue*

$$\lim_{s\to 1}(s-1)\zeta_K(s) = \frac{2^{r_1}(2\pi)^{r_2} R_K |\operatorname{Cl}(K)|}{w_K \sqrt{|\operatorname{disc}(K)|}}.$$

Here r_1 is the number of real places of K, r_2 is the number of complex places, R_K is the regulator, w_K is the number of roots of unity in K, and $\operatorname{disc}(K)$ is the discriminant.

Now, there are many elements of this formula which we have not previously defined. They will not generally be very important to us so I don't want to spend too much time on them, but briefly: the real places of a number field K are the embeddings $K \hookrightarrow \mathbb{R}$, complex places are embeddings $K \hookrightarrow \mathbb{C}$ up to complex conjugation which do not factor through $\mathbb{R}$,[24] R_K is the (generally irrational) determinant of a certain $r \times r$ matrix related to the units of $\mathcal{O}_K$ where $r = r_1 + r_2 - 1$, and the discriminant (a rational number) is the square of the determinant of a matrix related to bases of $\mathcal{O}_K$ as a module over $\mathbb{Z}$. Although the definitions in general can be somewhat complicated, typically the most difficult thing to compute on the right-hand side is the class number $|\operatorname{Cl}(K)|$, so we often view this formula as giving a relationship between the complicated algebraic quantity $|\operatorname{Cl}(K)|$ and the analytic quantity $\lim_{s\to 1}(s-1)\zeta_K(s)$, up to some "simpler" factors given by the other terms on the right.

For an example, consider $K = \mathbb{Q}$. We earlier mentioned that in this case the residue on the left is just 1. On the right, we only have one embedding of $\mathbb{Q}$ into $\mathbb{R}$, and none into $\mathbb{C}$ that don't factor through $\mathbb{R}$, so $r_1 = 1$ and $r_2 = 0$; thus $r = r_1 + r_2 - 1 = 0$ and so $R_{\mathbb{Q}}$ is the determinant of a 0×0 matrix, which is the empty matrix and so by convention has determinant 1; the only roots of unity in $\mathbb{Q}$ are 1 and -1, so $w_{\mathbb{Q}} = 2$; and $\operatorname{disc}(\mathbb{Q}) = 1$, so the class number formula says that

$$1 = \frac{2^1 \cdot (2\pi)^0 \cdot 1 \cdot |\operatorname{Cl}(\mathbb{Q})|}{2 \cdot 1} = |\operatorname{Cl}(\mathbb{Q})|.$$

[23] In fact it extends to a meromorphic function on the entire complex plane, but we don't want to bother pinning down its poles, or those of $\zeta_K(s)$ in general.

[24] So for example $\mathbb{Q}[x]/(x^2-2)$ has two real places, depending on whether x is sent to $\sqrt{2}$ or $-\sqrt{2}$, and no complex places; $\mathbb{Q}[x]/(x^2+1)$ has one complex place, given by the two conjugate embeddings sending $x \mapsto \pm i$, and no real places; and $\mathbb{Q}[x]/(x^3+2)$ has one real place given by $x \mapsto -\sqrt[3]{2}$ and one complex place given by the conjugate pair $x \mapsto 2^{-2/3}(1 \pm i\sqrt{3})$. In general if $K/\mathbb{Q}$ has degree d then $d = r_1 + 2r_2$.

Thus we recover the statement that $\mathrm{Cl}(\mathbb{Q})$ is trivial—which recall is equivalent to the fundamental theorem of arithmetic!

More generally, we can view the class number formula as a first example of a class of formulas relating "special values" of L-functions[25] to some arithmetic quantities, often involving a regulator term as well as the order of some complicated group. Often—as here—the right thing to study is not the value of the function, which may be trivial or may, as here, not make sense, but the leading coefficient of its Taylor or Laurent series. In this case, we could reformulate the class number formula using the functional equation for the Dedekind zeta function[26] to give a formula for the leading coefficient of the Taylor series of $\zeta_K(s)$ around $s = 0$, which is a more typical form for these sorts of formulas; a vast generalization is the Beilinson–Bloch conjecture, which subsumes related conjectures such as the Birch–Swinnerton-Dyer conjecture (which gives a similar formula for L-functions of elliptic curves).

The class number also has consequences for classical analytic and algebraic number theory: the fact that $\zeta_K(s)$ has a simple pole at $s = 1$ with finite residue can be used to show that Dirichlet L-functions are nonzero at $s = 1$, a crucial step in proving Dirichlet's theorem on the infinitude of primes in arithmetic progressions; and the proof of the class number formula yields along the way the finiteness of the class group.[27]

Hopefully this justifies our interest in the class group beyond the realm of pure class field theory; finally, before reinterpreting in terms of adeles let's discuss what the Frobenius element means in terms of prime splitting.

2.4 Prime splitting: reprise

Let's start with our standard example of the Gaussian integers, $\mathcal{O} = \mathbb{Z}[i]$, $K = \mathbb{Q}(i)$. Since this is a quadratic extension of $\mathbb{Q}$, the Galois group is just $\mathrm{Gal}(K/\mathbb{Q}) \simeq \mathbb{Z}/2$, with the nontrivial element acting by complex conjugation.

For each prime p, there are three possibilities: either p is ramified, i.e. it factors as $\mathfrak{p}^2$ for some prime ideal $\mathfrak{p}$ of $\mathcal{O}$ (as we've seen, this happens only for $p = 2$); it is split, i.e. it factors as $\mathfrak{p}_1 \cdot \mathfrak{p}_2$ for distinct primes $\mathfrak{p}_1$ and $\mathfrak{p}_2$, which will be Galois conjugate (we've claimed that this occurs for $p \equiv 1 \pmod 4$); or it is inert, i.e. (p) remains prime in $\mathcal{O}$ (we've claimed that this occurs for $p \equiv 3 \pmod 4$). Ramification makes things a little bit complicated, but it only ever occurs for finitely many primes so we won't worry about it too much; let's think about the Frobenius element in the split and inert cases.

If $(p) = \mathfrak{p}_1 \cdot \mathfrak{p}_2$ is split, complex conjugation interchanges the two primes $\mathfrak{p}_1$ and $\mathfrak{p}_2$, so the only element of $\mathrm{Gal}(K/\mathbb{Q})$ that fixes $\mathfrak{p}_i$ is the identity 1; so this must be the Frobenius element.

On the other hand, if p is inert, then since it is defined over $\mathbb{Z}$ it is fixed by every element of the Galois group; we have $\mathcal{O}/p \simeq \mathbb{F}_p(i) \simeq \mathbb{F}_{p^2}$, while $\mathbb{Z}/p = \mathbb{F}_p$, and the Frobenius element is a generator and so must be the nontrivial element of the group $\mathrm{Gal}(\mathbb{F}_{p^2}/\mathbb{F}_p) \simeq \{1, \sigma\}$.

[25] A large class of analytic functions of arithmetic interest, including zeta functions as the prototypical examples.

[26] Which relates $\zeta_K(s)$ and $\zeta_K(1-s)$.

[27] This of course depends on how you do it—but the tools involved are very similar.

Therefore in this case the Frobenius is nontrivial. Thus, assuming p is unramified, it is split if and only if $\sigma_p = 1$: so we can read the splitting off of the reciprocity map.

The same principle holds in general: for a prime $\mathfrak{p}$ in $\mathcal{O}_K$, its splitting in $\mathcal{O}_L$ is governed by $\sigma_{\mathfrak{p}} \in \mathrm{Gal}(L/K)$. If $\sigma_{\mathfrak{p}} = 1$, then $\mathfrak{p}$ is completely split; if $\sigma_{\mathfrak{p}}$ has maximal order, then $\mathfrak{p}$ is inert.[28] One can also deduce the precise splitting behavior for intermediate splitting: $\mathrm{Gal}(L/K)$ can be thought of as permuting the roots of the polynomial $f(x)$ over K defining L, so if L/K is degree d then $\mathrm{Gal}(L/K)$ embeds naturally into the permutation group S_d on d items. Permutations of d items have associated cycle types, which can be thought of as partitions of d; the cycle type of $\sigma_{\mathfrak{p}}$ in $\mathrm{Gal}(L/K)$ is the splitting type of $\mathfrak{p}$ in $\mathcal{O}_L$. Thus the splitting problem finally boils down to understanding the reciprocity map $\mathfrak{p} \mapsto \sigma_{\mathfrak{p}}$.

3. Class field theory II: adeles

3.1 Adeles

To get a "better" formulation of class field theory, we first need to introduce a new object: the ring of adeles. We do need some topology in this section, although we'll mostly handwave it. We also assume some familiarity with p-adic numbers.

Consider the simplest number field $\mathbb{Q}$. One can make $\mathbb{Q}$ into a topological field in many ways; if we require that the topology be induced by an absolute value, then it's possible to classify these metrics. This is Ostrowski's theorem: the absolute values on $\mathbb{Q}$, up to a certain equivalence relation, are all either the standard absolute value $|\cdot|$ induced from the embedding into the real numbers or the p-adic absolute value $|\cdot|_p$. The idea is to treat these all on the same level: we call each of these *places*,[29] with the standard absolute value $|\cdot|$ being the "infinite place." Thus we think of augmenting the prime numbers with an extra "infinite prime."

To each place v we can naturally associate an embedding $\mathbb{Q} \hookrightarrow \mathbb{Q}_v$ where $\mathbb{Q}_v$ is the completion of $\mathbb{Q}$ with respect to $|\cdot|_v$: for $v = p$ these are the p-adic numbers $\mathbb{Q}_p$, while for $v = \infty$ the completion is $\mathbb{Q}_\infty = \mathbb{R}$. For the "finite places," i.e. $v = p$, these completions come equipped with rings of integers $\mathbb{Z}_p$, the p-completion of $\mathbb{Z}$; we can think of $\mathbb{Z}_p$ as the subset of $\mathbb{Q}_p$ with $|x|_p \leq 1$, which has a ring structure due to the fact that $|\cdot|_p$ is *nonarchimedean*, i.e. it satisfies the strong triangle inequality: $|x+y|_p \leq \max(|x|_p, |y|_p)$. The infinite place has no such structure, since the subset of $\mathbb{R}$ with $|x|_\infty \leq 1$ is not a ring: this is because $|\cdot|_\infty$ is *archimedean*, i.e. it satisfies the weak triangle inequality $|x+y|_\infty \leq |x|_\infty + |y|_\infty$ but not the strong one.

To get a version of the integers together with their completion at every finite place, one can just form the product $\widehat{\mathbb{Z}} := \prod_p \mathbb{Z}_p$; this can also be defined as the inverse limit $\varprojlim_n \mathbb{Z}/n\mathbb{Z}$, with n ranging over the integers partially ordered by divisibility. This is well-behaved: for example every $\mathbb{Z}_p$ is compact, and so so is $\widehat{\mathbb{Z}}$.

On the other hand, the analogous thing $\prod_p \mathbb{Q}_p$ (or better, to incorporate the infinite place, $\prod_v \mathbb{Q}_v$) is not very well-behaved: each $\mathbb{Q}_v$ is locally compact, which makes it possible e.g. to do Fourier analysis on the real or p-adic numbers; but the infinite product no longer

[28] Again, we neglect ramification.

[29] The real and complex places from §2.3 refer to the same notion.

is.

As a hint towards how to fix this, we can try looking at how $\mathbb{Q}$ embeds into this infinite product; surely even if we restrict the product somehow it should still naturally contain $\mathbb{Q}$. For any rational number $\frac{a}{b}$ (in reduced form), there are only finitely many primes that divide b, so these are the only ones for which the image of $\frac{a}{b}$ in $\mathbb{Q}_p$ will have $|\frac{a}{b}| > 1$; for all other primes, the image of $\frac{a}{b}$ will be in $\mathbb{Z}_p$. Thus if we restrict the product to require that all but finitely many components live in $\mathbb{Z}_p$, the result will still contain the image of $\mathbb{Q}$.

This turns out to be very advantageous, because this restricted product is now locally compact! Thus we define the ring of adeles $\mathbb{A}_\mathbb{Q}$ to be the "restricted product"

$$\mathbb{A}_\mathbb{Q} = \prod_v{}'(\mathbb{Q}_v, \mathbb{Z}_v),$$

i.e. the subring of the product $\prod_v \mathbb{Q}_v$ consisting of elements (x_v) such that all but finitely many of the x_v are contained in $\mathbb{Z}_v$. Of course for $v = \infty$ the subring Z_∞ doesn't exist, so at the infinite place the condition is trivial; but there is only one infinite place so this is okay.

If we were to split off the infinite place, letting $\mathbb{A}_\mathbb{Q}^\infty$ be the restricted product only over the finite places then $\mathbb{A}_\mathbb{Q} = \mathbb{R} \times \mathbb{A}_\mathbb{Q}^\infty$; and in fact we can express $\mathbb{A}_\mathbb{Q}^\infty$ much more simply as $\widehat{\mathbb{Z}} \otimes_\mathbb{Z} \mathbb{Q}$, since the factor of $\mathbb{Q}$ allows finitely many primes in the denominator and everything else in $\widehat{\mathbb{Z}} = \prod_p \mathbb{Z}_p$. Thus we can more compactly write $\mathbb{A}_\mathbb{Q} = \mathbb{R} \times \widehat{\mathbb{Z}} \otimes_\mathbb{Z} \mathbb{Q}$; but we still want to think of this in terms of the restricted product above.

By construction, we have a natural embedding $\mathbb{Q} \hookrightarrow \mathbb{A}_\mathbb{Q}$. If we take a nonzero element $x \in \mathbb{Q}$, its image in the pth factor of $\mathbb{A}_\mathbb{Q}$ is given by the order of p in x, together with some unit in $\mathbb{Z}_p^\times$; thus the embedding records the prime factorization of x together with some data about p-adic units. If we only want to recover the prime factorization, this suggests quotienting by $\mathbb{Z}_p^\times$; doing this for every finite place p gives a map $\mathbb{Q}^\times \hookrightarrow (\mathbb{A}_\mathbb{Q}^\infty)^\times / \widehat{\mathbb{Z}}^\times$, with the right-hand side describing possible orders at every finite place.

This is strongly reminiscent of the definition of the class group: the right-hand side looks like the definition of the group of fractional ideals, while the left-hand side corresponds to principal ideals. This suggests taking the cokernel $\mathbb{Q}^\times \backslash (\mathbb{A}_\mathbb{Q}^\infty)^\times / \widehat{\mathbb{Z}}^\times$;[30] and in fact this is isomorphic to the class group $\mathrm{Cl}(\mathbb{Q})$, and thus in this case trivial.

If it's trivial, why describe it in this way at all? Because we can generalize all of the above to any number field K, rather than just $\mathbb{Q}$: the finite places correspond to nonzero prime ideals $\mathfrak{p}$, and the infinite places are real and complex places (the fact that there is only one infinite place is special to $\mathbb{Q}$ (and imaginary quadratic fields), but there are always only finitely many). We always have a map $K^\times \to (\mathbb{A}_K^\infty)^\times$, and always have an isomorphism $K^\times \backslash (\mathbb{A}_K^\infty)^\times / \widehat{\mathcal{O}}_K^\times \simeq \mathrm{Cl}(K)$, where $\mathcal{O}_K = \prod_\mathfrak{p} \mathcal{O}_{K,\mathfrak{p}}$ is the product of the $\mathfrak{p}$-completions of $\mathcal{O}_K$; in general this group is no longer trivial.

As we've seen, the class group is in general not flexible enough to use as the source for the reciprocity map. However, this construction of the class group as $K^\times \backslash (\mathbb{A}_K^\infty)^\times / \widehat{\mathcal{O}}_K$ suggests a natural source to use instead: simply don't quotient by $\widehat{\mathcal{O}}_K$! In fact, our philosophy of

[30] The double quotient here means quotienting by two group actions, one on the right and one on the left; since everything is abelian it doesn't matter too much, but avoids some confusion in putting them on the same side, and there are generalizations to nonabelian cases (as we'll discuss).

treating the finite and infinite places on the same level suggests that the infinite places should also be relevant;[31] so we'll take as our adelic class group[32] the quotient $K^\times \backslash \mathbb{A}_K^\times$.

Given an adele (x_v) for v ranging over all places of K, at each finite place $\mathfrak{p}$ we can take the $\mathfrak{p}$-adic valuation $v_\mathfrak{p}(x_\mathfrak{p})$[33] and thence the ideal $\mathfrak{p}^{v_\mathfrak{p}(x_\mathfrak{p})}$. Taking the product over all finite places gives a map $\mathbb{A}_K^\times \to \mathcal{I}_K$ sending

$$(x_v) \mapsto \prod_\mathfrak{p} \mathfrak{p}^{v_\mathfrak{p}(x_\mathfrak{p})}.$$

Since only finitely many of the $v_\mathfrak{p}(x_\mathfrak{p})$ are nonzero (as any element of $\mathcal{O}_{K,\mathfrak{p}}^\times$ has valuation 0), the product has only finitely many nontrivial terms and so is well-defined. Further since we can choose the $v_\mathfrak{p}(x_\mathfrak{p})$ arbitrarily for each $\mathfrak{p}$, this map is clearly surjective; its kernel consists of adeles (x_v) for which $x_\mathfrak{p} \in \mathcal{O}_{K,\mathfrak{p}}^\times$ for all finite places $\mathfrak{p}$, i.e. $\widehat{\mathcal{O}}_K^\times \times \prod_{v|\infty} K_v$ where the product is over all archimedean places v of K.

We have a natural embedding $K^\times \hookrightarrow \mathbb{A}_K^\times$, so we can consider the image of $x \in K^\times$ under this map. It will be the fractional ideal with prime factorization given by the valuation of x at each prime ideal: in other words, it is the principal ideal generated by x! So the natural map $\mathbb{A}_K^\times \to \mathcal{I}_K$ above induces, upon quotienting by $K^\times$ on the left (via the natural embedding) and the right (as the space of principal fractional ideals) induces a map

$$K^\times \backslash \mathbb{A}_K^\times \to K^\times \backslash \mathcal{I}_K.$$

The quotient on the right, by definition, is the class group $\mathrm{Cl}(K)$; the quotient on the left is an adelic version of it, sometimes called the adelic or idele class group C_K. Thus we have a comparison map

$$C_K \to \mathrm{Cl}(K).$$

We know from above that it is surjective; and we can describe its kernel explicitly as $K^\times \backslash \left(\widehat{\mathcal{O}}_K^\times \times \prod_{v|\infty} K_v \right)$. If we instead took the finite adeles and quotiented by $\mathcal{O}_K^\times$, we would get an isomorphism

$$K^\times \backslash (\mathbb{A}_K^\infty)^\times / \widehat{\mathcal{O}}_K^\times \simeq \mathrm{Cl}(K),$$

generalizing the statement for $K = \mathbb{Q}$ above.

Our claim is that the right way to "augment" $\mathrm{Cl}(K)$ as a source for our reciprocity map is to replace it by the idele class group C_K. We've spent a lot of effort working out the right source for our reciprocity map, but our target also has to change: we haven't specified any extension L here, so the target should also be some sort of "absolute" Galois group! Indeed, we'll take as a target the group $\mathrm{Gal}_K := \mathrm{Gal}(\overline{K}/K)$ where $\overline{K}$ is an algebraic closure of K;

[31] In practice the dependence on the infinite places can be subtle; if you find them confusing I encourage you to ignore them, at least at first pass.

[32] Sometimes the group of units of the ring of adeles is called the ideles, and so this is called the idele class group (parallel to the ideal class group).

[33] The $\mathfrak{p}$-adic valuation is related to the $\mathfrak{p}$-adic absolute value by $|x|_\mathfrak{p} = N(\mathfrak{p})^{-v_\mathfrak{p}(x)}$, so by convention $v_\mathfrak{p}(0) = \infty$. (One can replace $N(\mathfrak{p})$ with any other constant greater than 1 and get a topologically equivalent absolute value, but this is a standard normalization.)

this is called the absolute Galois group of K.[34] Again we really can only target the abelian part of this group $\mathrm{Gal}_{\mathbb{Q}}^{\mathrm{ab}}$,[35] which can be understood as the Galois group $\mathrm{Gal}(K^{\mathrm{ab}}/K)$ where K^{ab} is the maximal abelian extension of K; this is again an infinite extension.

So we hope to have a reciprocity map

$$C_K \to \mathrm{Gal}_K^{\mathrm{ab}},$$

with some good properties as gestured towards in §2. We won't get into the construction of this map, but we do want to talk more about what properties it should have. In particular, say we do fix a finite abelian extension L/K. Given our algebraic closure $\overline{K}/K$, L gives an intermediate extension and thus a quotient $\mathrm{Gal}_K \to \mathrm{Gal}(L/K)$; since the target is abelian, this quotient factors through $\mathrm{Gal}_K^{\mathrm{ab}}$.[36] Can we find a similar quotient of C_K to act as a source of a finite-level reciprocity map as in §2.2?

We can get a hint from the Galois side: the kernel of $\mathrm{Gal}_K = \mathrm{Gal}(\overline{K}/K) \to \mathrm{Gal}(L/K)$ is $\mathrm{Gal}(\overline{K}/L) = \mathrm{Gal}_L$ and similarly on abelianizations.[37] Of course, the whole story should be true for L just as well as for K, so we expect to have a reciprocity map $C_L \to \mathrm{Gal}_L^{\mathrm{ab}}$; as $\mathrm{Gal}_L^{\mathrm{ab}} \hookrightarrow \mathrm{Gal}_K^{\mathrm{ab}}$ with cokernel $\mathrm{Gal}(L/K)$, we might hope that something similar happens on the adelic side. First, we need a map $C_L \to C_K$, going the opposite direction from the inclusion $K \hookrightarrow L$. Such a map is given by the norm: for an adele $(x_w) \in \mathbb{A}_L^\times$, at each place w of L we can take the product of the Galois conjugates $\sigma(x_w)$ for $\sigma \in \mathrm{Gal}(L_w/K_v)$, where v is the place of K lying under w, to get the relative $N(x_w) \in K_v^\times$. Doing this for every place and taking the product gives the norm map

$$N : \mathbb{A}_L^\times \to \mathbb{A}_K^\times,$$

and if (x_w) is the image of an element of $L^\times$ in $\mathbb{A}_L^\times$ then its norm is just the field norm $N : L^\times \to K^\times$, defined in the same way, and so the norm map descends to

$$N : C_L \to C_K.$$

However we cannot expect it to be an injection: for example, any two Galois-conjugate elements of $\mathbb{A}_L^\times$ have the same norm in $\mathbb{A}_K^\times$. Nevertheless, we can consider its cokernel $C_K/N(C_L)$: this should be the "finite level" source for the reciprocity map

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

When we discussed this sort of map in §2, we viewed prime ideals of $\mathcal{O}_K$ as having an equivalence class on the left and being sent under this map to their Frobenius, which

[34] The Galois theory of infinite extensions is a little more complicated than classical Galois theory and needs to be augmented by viewing Gal_K as a *topological* group, but we won't worry too much about this sort of thing for the moment; given a suitable infinite Galois theory, one can understand the study of number fields over K as the study of Gal_K, or in an absolute sense the study of all number fields as the study of $\mathrm{Gal}_{\mathbb{Q}}$.

[35] The abelianization, which can be defined as the maximal abelian quotient of $\mathrm{Gal}_{\mathbb{Q}}$.

[36] Alternatively, we could view L as a subfield of L^{ab}.

[37] Since L/K is algebraic and $\overline{K}$ is algebraically closed and contains L, it is also the algebraic closure of L, so $\mathrm{Gal}(\overline{K}/L) = \mathrm{Gal}(\overline{L}/L) = \mathrm{Gal}_L$.

determines their splitting type in $\mathcal{O}_L$. The description should match up here: a representative of a prime $\mathfrak{p}$ of K is given by an adele (x_v) where $x_\mathfrak{p}$ is an element of $K_\mathfrak{p}$ with $v_\mathfrak{p}(x_\mathfrak{p}) = 1$, and for all $v \neq \mathfrak{p}$ we have $x_v = 1$. Such an element $x_\mathfrak{p}$ of $K_\mathfrak{p}$ is called a uniformizer, often denoted ϖ; for example in $\mathbb{Q}_p$ a standard choice of uniformizer is p, though $u \cdot p$ for any $u \in \mathbb{Z}_p^\times$ would also work. We could construct the adele (x_v) as the image of a uniformizer under the inclusion $K_\mathfrak{p}^\times \hookrightarrow \mathbb{A}_K^\times \to C_K$. On the other hand, at least when $\mathfrak{p}$ is unramified in L the claim is that the image of a uniformizer under this composition should be the Frobenius element of $\mathfrak{p}$. To describe the rest of the composition $K_\mathfrak{p}^\times \to C_K/N(C_L) \to \mathrm{Gal}(L/K)$, we need to know what happens to $\mathcal{O}_{K,\mathfrak{p}}^\times$, since together with the uniformizer this generates the whole multiplicative group; at least in the unramified situation, we're used to quotienting out by the group of units, so we expect that $\mathcal{O}_{K,\mathfrak{p}}^\times$ should be in the kernel of this map.

One might also ask about the composition $K_v^\times \to C_K/N(C_L) \to \mathrm{Gal}(L/K)$ when v is archimedean. In this case, since $\mathrm{Gal}(L/K)$ is a discrete group and $K_v^\times$ is a topological group with either one or two connected components, it must always send the connected component of the identity to 1; if v is unramified, it will kill all of $K_v^\times$. If v is ramified (so $K_v^\times = \mathbb{R}^\times$ has two connected components), the component of the identity goes to 1, while the component of -1 maps to complex conjugation on L/K.

In particular, a lot of the requirements we'd like to have on our reciprocity map seem to boil down to what happens working one place at a time via $K_v^\times \to C_K \to \mathrm{Gal}_K^{\mathrm{ab}} \to \mathrm{Gal}(L/K)$. This suggests looking for a *local* version of class field theory, which deals only with one place at a time. The source of the local reciprocity map at v should be $K_v^\times$; the target should be related to Gal_K and the place v. It turns out that the right way to do this is to restrict to $\mathrm{Gal}_{K_v}^{\mathrm{ab}} = \mathrm{Gal}(\overline{K}_v/K_v)^{\mathrm{ab}}$, which maps into $\mathrm{Gal}_K^{\mathrm{ab}}$ since $K \subset K_v$ so any automorphism fixing K_v also fixes K.[38] So the local reciprocity map is a map

$$K_v^\times \to \mathrm{Gal}_{K_v}^{\mathrm{ab}}.$$

Indeed, here is the main theorem of local class field theory.

Theorem (Local class field theory)**.** *Let K be a number field and v a place. There exists a canonical injective reciprocity map*

$$\Psi_{K_v} : K_v^\times \to \mathrm{Gal}_{K_v}^{\mathrm{ab}}$$

such that for every finite extension L/K with w a place of L over v such that L_w/K_v is abelian,

$$\Psi_{L_w/K_v} : K_v^\times \xrightarrow{\Psi_{K_v}} \mathrm{Gal}_{K_v}^{\mathrm{ab}} \twoheadrightarrow \mathrm{Gal}(L_w/K_v)$$

is surjective with kernel the image of the norm map $N(L_w^\times)$ in $K_v^\times$, and if v is nonarchimedean and unramified in L then Ψ_{K_v} sends $\mathcal{O}_K^\times$ to the identity and any uniformizer to the Frobenius element[39] *of v.*

[38] There is something to be said about requiring continuity on the topological fields K_v and $\overline{K}_v$ as well, but we won't worry about it too much.

[39] Defined for local extensions just as for global.

In particular the version for finite extensions shows that the failure of surjectivity for Ψ_{K_v} is essentially a topological problem: the infinite Galois group $\mathrm{Gal}_{K_v}^{\mathrm{ab}}$ is profinitely complete, while $K_v^\times$ is not. This is easily fixed: taking the profinite completion $\widehat{K_v^\times}$ makes the reciprocity map an isomorphism $\widehat{K_v^\times} \xrightarrow{\sim} \mathrm{Gal}_{K_v}^{\mathrm{ab}}$.

Galois theory tells us that subgroups of $\mathrm{Gal}(L_w/K_v)$ correspond to intermediate subfields. On the infinite level, things are a little more complicated and we have to take the topology of the group into account as well, but spiritually the same thing is true: *open* finite index subgroups of $\mathrm{Gal}_{K_v}^{\mathrm{ab}}$ correspond to finite abelian extensions of K_v. Thus we'd like the same sort of property to hold on the left:

Theorem (Local existence theorem)**.** *Let K be a number field and v a place. The reciprocity map Ψ_{K_v} induces an order-reversing bijection between finite abelian extensions L_w of K_v and open finite index subgroups of $K_v^\times$, via $L_w \mapsto N(L_w^\times)$.*

In particular, finite abelian extensions giving rise to the finite-level reciprocity maps (corresponding to given open finite index subgroups of $K_v^\times$) actually exist.

Now that we understand the local situation at each place, let's try to write down what should happen globally:

Theorem (Global class field theory)**.** *Let K be a number field. There exists a canonical reciprocity map*

$$\Psi_K : C_K \to \mathrm{Gal}_K^{\mathrm{ab}}$$

with dense image, such that for any finite abelian extension L/K, the composition

$$\Psi_{L/K} : C_K \to \mathrm{Gal}_K^{\mathrm{ab}} \twoheadrightarrow \mathrm{Gal}(L/K)$$

is surjective, and for every place v of K we have a commutative diagram

$$\begin{array}{ccc} K_v^\times & \xrightarrow{\Psi_{K_v}} & \mathrm{Gal}_{K_v}^{\mathrm{ab}} \\ \downarrow & & \downarrow \\ C_K & \xrightarrow{\Psi_K} & \mathrm{Gal}_K^{\mathrm{ab}} \end{array}.$$

Similarly at finite level, for every finite abelian extension L/K and place w over v, L_w/K_v is finite abelian and the diagram

$$\begin{array}{ccc} K_v^\times & \xrightarrow{\Psi_{L_w/K_v}} & \mathrm{Gal}(L_w/K_v) \\ \downarrow & & \downarrow \\ C_K & \xrightarrow{\Psi_{L/K}} & \mathrm{Gal}(L/K) \end{array}$$

commutes.

In particular, the second two statements mean that the local and global reciprocity maps are compatible.

The failure of injectivity is potentially nontrivial: C_K may have nontrivial connected components, coming from the infinite places, while $\mathrm{Gal}_K^{\mathrm{ab}}$ is totally disconnected. Thus if D_K is the connected component of the identity in C_K, Ψ_K takes it the identity; and so our best hope is that the induced map $C_K/D_K \to \mathrm{Gal}_K^{\mathrm{ab}}$ is injective. This turns out to be true; in fact, one can show that it is also surjective! Thus we get an isomorphism of topological groups $C_K/D_K \simeq \mathrm{Gal}_K^{\mathrm{ab}}$.

For example, take $K = \mathbb{Q}$. Then $\mathbb{A}_{\mathbb{Q}}^\times = \mathbb{R}^\times \times (\widehat{\mathbb{Z}} \otimes_{\mathbb{Z}} \mathbb{Q})^\times$, which has two connected components corresponding to the two connected components $\mathbb{R}_{>0}$ and $\mathbb{R}_{<0}$, so the connected component of the identity is $D_{\mathbb{Q}} = \mathbb{R}_{>0} \times \{1\}$. Therefore $\mathbb{A}_{\mathbb{Q}}^\times / D_{\mathbb{Q}} = \{\pm 1\} \times (\widehat{\mathbb{Z}} \otimes_{\mathbb{Z}} \mathbb{Q})^\times = \{\pm 1\} \times (\mathbb{A}_{\mathbb{Q}}^\infty)^\times$. The map $\mathbb{Q}^\times \hookrightarrow \mathbb{A}_{\mathbb{Q}}^\times / D_{\mathbb{Q}}$ records at each finite place p the p-adic valuation of a given rational number x, together with a unit in $\mathbb{Z}_p^\times$; and at the infinite place it records the sign of x, either 1 or -1. In particular we find $\mathbb{Q}^\times \backslash \mathbb{A}_{\mathbb{Q}}^\times / D_{\mathbb{Q}} = C_{\mathbb{Q}}/D_{\mathbb{Q}} \simeq \prod_p \mathbb{Z}_p^\times = \widehat{\mathbb{Z}}^\times$, so the isomorphism above gives $\mathrm{Gal}_{\mathbb{Q}}^{\mathrm{ab}} \simeq \widehat{\mathbb{Z}}^\times$!

We mentioned before that one can view algebraic number theory in a sense as the study of the group $\mathrm{Gal}_{\mathbb{Q}}$, since this group classifies number fields; so the fact that we've found a very explicit description of the abelian part of it is a great success. We'll see in §3.2 how to concretely reinterpret this result in terms of abelian number fields over $\mathbb{Q}$.

Before we do so, we'd like a global version of the existence theorem, i.e. given an abelian extension L/K we get an open finite index subgroup $\mathrm{Gal}_L^{\mathrm{ab}}$ of $\mathrm{Gal}_K^{\mathrm{ab}}$ and thus an open finite index subgroup of C_K; we'd like to again know that all such subgroups arise in this way. In fact, one can say even more, motivated by the fact that in a profinite group such as $\mathrm{Gal}_K^{\mathrm{ab}}$, open subgroups are also closed, and namely are the closed subgroups of finite index. On the adelic side, we also need to make a mild modification to handle the connected components at infinity.

Theorem (Global existence theorem)**.** *Let K be a number field. The reciprocity map Ψ_K induces an order-reversing bijection between all abelian algebraic extensions (even infinite ones) L/K and closed subgroups of C_K containing the connected component of the identity, sending L/K to the kernel of $\Psi_{L/K} : C_K \to \mathrm{Gal}_K^{\mathrm{ab}} \twoheadrightarrow \mathrm{Gal}(L/K)$.*

We saw before that we can describe this kernel as the image of the norm map $N : C_L \to C_K$, so the bijection could equally well be phrased as sending L/K to $N(C_L)$ in C_K, more parallel to the local existence theorem.

Given a suitable subgroup H of C_K, the corresponding extension L of K is called the class field of H. Recalling that $C_K/(\widehat{\mathcal{O}}_K \times \prod_{v|\infty} K_v^\times) \simeq \mathrm{Cl}(K)$, taking $H = \widehat{\mathcal{O}}_K \times \prod_{v|\infty} K_v^\times$ the existence theorem tells us that there is some abelian extension L of K (finite since this subgroup has finite index) such that $\Psi_{L/K} : C_K \to \mathrm{Gal}(L/K)$ is surjective with kernel H, so $\Psi_{L/K}$ gives an isomorphism $C_K/H \simeq \mathrm{Cl}(K) \xrightarrow{\sim} \mathrm{Gal}(L/K)$; in other words there exists a finite abelian extension of K whose Galois group over K is isomorphic to the class group $\mathrm{Cl}(K)$. This is the Hilbert class field of K mentioned near the end of §2.2; it is the maximal abelian unramified extension of K, since we force $\Psi_{L/K}$ to kill precisely $\mathcal{O}_{\mathfrak{p}}^\times$ at each finite place $\mathfrak{p}$ and the whole multiplicative group $K_v^\times$ at infinite places v, which is the determining property of those places being unramified. By adjusting the subgroup H, we can adjust the ramification behavior to get more general class fields: for example, taking H the same at finite places but only containing the connected components of $K_v^\times$ at infinite places gives the

narrow Hilbert class field, the maximal abelian extension of K which is unramified at every *finite* place (but may be ramified at infinite places); the corresponding quotient C_K/H is called the narrow class group of K.

3.2 Example: Kronecker–Weber

We want to give an alternative way to think about the description $\mathrm{Gal}_{\mathbb{Q}}^{\mathrm{ab}} \simeq \widehat{\mathbb{Z}}^{\times}$. This can be described as the *explicit class field theory* for $\mathbb{Q}$, meaning an explicit description of $\mathbb{Q}^{\mathrm{ab}}$ as an extension of $\mathbb{Q}$ giving rise to an explicit description of splitting behaviors of primes of $\mathbb{Q}$ in every abelian extension.[40]

To do so, we need to introduce a particularly simple type of abelian extensions of $\mathbb{Q}$: the cyclotomic fields $\mathbb{Q}(\zeta_n)$. Here ζ_n is a primitive nth root of unity,[41] e.g. $\zeta_n = e^{2\pi i/n}$; so e.g. we could take $\zeta_4 = i$, though $-i$ works just as well. These are solutions to the polynomial equation $x^n - 1 = 0$. However for all $n \geq 2$ this polynomial is not irreducible; indeed, it always has a factor of $x - 1$, and may have more. The minimal polynomial of ζ_n is called the nth cyclotomic polynomial $\Phi_n(x)$; for example, $\Phi_2(x) = x + 1$ (since ζ_2 is just -1), $\Phi_3(x) = x^2 + x + 1$, $\Phi_4(x) = x^2 + 1$, and $\Phi_5(x) = x^4 + x^3 + x^2 + x + 1$; in particular if n is prime then $\Phi_n(x) = x^{n-1} + x^{n-2} + \cdots + x + 1$ has degree $n - 1$. More generally, one can find that $\deg \Phi_n(x) = \varphi(n)$, the number of integers between 1 and n which are relatively prime to n; if n is prime then this is all the integers $1, \ldots, n - 1$, while e.g. for $n = 6$ the only integers between 1 and 6 with no common factors with 6 are 1 and 5, so $\varphi(6) = 2$, and indeed $\Phi_6(x) = x^2 - x + 1$ has degree 2.

We can say a little more by thinking about the Galois group of $\mathbb{Q}(\zeta_n)$ over $\mathbb{Q}$. This group permutes the set of primitive nth roots of unity; since the non-primitive nth roots of unity are also contained in $\mathbb{Q}(\zeta_n)$, it also acts on these, so $\mathrm{Gal}(\mathbb{Q}(\zeta_n)/\mathbb{Q})$ acts faithfully on the group of nth roots of unity, isomorphic to $\mathbb{Z}/n\mathbb{Z}$ since we can write any nth root of unity as ζ_n^r for some unique $r \in \mathbb{Z}/n\mathbb{Z}$. The endomorphism ring of $\mathbb{Z}/n\mathbb{Z}$ is just $\mathbb{Z}/n\mathbb{Z}$, acting via multiplication; and the invertible endomorphisms are the invertible elements $(\mathbb{Z}/n\mathbb{Z})^{\times} \subset \mathbb{Z}/n\mathbb{Z}$, so $\mathrm{Gal}(\mathbb{Q}(\zeta_n)/\mathbb{Q}) \simeq (\mathbb{Z}/n\mathbb{Z})^{\times}$. Since $|(\mathbb{Z}/n\mathbb{Z})^{\times}| = \varphi(n)$, this explains the calculation of the degree of $\Phi_n(x)$ above.

These extensions are nice because they are usefully explicit: we see explicitly that $\mathrm{Gal}(\mathbb{Q}(\zeta_n)/\mathbb{Q})$ is abelian because it's the group $(\mathbb{Z}/n\mathbb{Z})^{\times}$, and we can check things like splitting straightforwardly: per our discussion in §2.4, a prime p splits completely in $\mathbb{Q}(\zeta_n)$ if and only if its Frobenius element $\sigma_{\mathfrak{p}} \in \mathrm{Gal}(\mathbb{Q}(\zeta_n)/\mathbb{Q})$ is equal to 1. But we can say very explicitly how the Frobenius acts on roots of unity: by taking their pth powers! Since $\zeta_n \mapsto \zeta_n^m$ is determined by the class of $m \pmod n$, this means that $\sigma_p \in \mathrm{Gal}(\mathbb{Q}(\zeta_n)/\mathbb{Q}) \simeq (\mathbb{Z}/n\mathbb{Z})^{\times}$ is just the class of p modulo n.[42] Thus: p splits completely in $\mathbb{Q}(\zeta_n)$ if and only if $p \equiv 1 \pmod n$.

[40] The analogous question for other number fields is very hard; for imaginary quadratic number fields, there exists a solution in terms of the theory of complex multiplication of elliptic curves, and there is some work towards explicit class field theory for real quadratic fields. Otherwise, as far as I know, the question is wide open. The local situation is much better, but beyond the scope of this note.

[41] i.e. ζ_n satisfies $\zeta_n^n = 1$, but $\zeta_n^m \neq 1$ for any $m < n$; for example, -1 is a fourth root of unity since $(-1)^4 = 1$, but not a primitive root since $(-1)^2$ is also equal to 1.

[42] This will be invertible provided p does not divide n; if it does divide n, then p will be ramified in $\mathbb{Q}(\zeta_n)$.

Now, we could try to bundle the cyclotomic fields together: let $\mathbb{Q}^{\text{cyc}}$ be the infinite algebraic extension of $\mathbb{Q}$ generated by the roots of unity ζ_n for *all* n. This can be written as $\mathbb{Q}^{\text{cyc}} = \varinjlim_n \mathbb{Q}(\zeta_n)$, and so by some mild categorical nonsense $\text{Gal}(\mathbb{Q}^{\text{cyc}}/\mathbb{Q}) = \varprojlim_n \text{Gal}(\mathbb{Q}(\zeta_n)/\mathbb{Q}) = \varprojlim (\mathbb{Z}/n\mathbb{Z})^\times = (\varprojlim_n \mathbb{Z}/n\mathbb{Z})^\times = \widehat{\mathbb{Z}}^\times$. On the other hand, recall that we found via class field theory that $\text{Gal}(\mathbb{Q}^{\text{ab}}/\mathbb{Q}) = \widehat{\mathbb{Z}}^\times$. This suggests that in fact $\mathbb{Q}^{\text{cyc}} = \mathbb{Q}^{\text{ab}}$, i.e. the infinite cyclotomic field $\mathbb{Q}^{\text{cyc}}$ is the maximal abelian extension of $\mathbb{Q}$; and in fact this turns out to be true.

Theorem (Kronecker–Weber, version I)**.** *The maximal abelian extension of* $\mathbb{Q}$ *is* $\mathbb{Q}^{\text{cyc}}$*.*

In particular, since every $\mathbb{Q}(\zeta_n)$ is abelian over $\mathbb{Q}$ so is $\mathbb{Q}^{\text{cyc}}$; so this statement really boils down to the converse, that in addition to being an abelian extension of $\mathbb{Q}$ this is the maximal one: every abelian extension of $\mathbb{Q}$ embeds into $\mathbb{Q}^{\text{cyc}}$. Since $\mathbb{Q}^{\text{ab}}$ and $\mathbb{Q}^{\text{cyc}}$ are (co)limits of finite extensions, it suffices to have this at finite level; in other words the above theorem is really equivalent to the following.

Theorem (Kronecker–Weber, version II)**.** *For every finite abelian extension* $K/\mathbb{Q}$*, there exists a positive integer* n *such that* K *is a subfield of* $\mathbb{Q}(\zeta_n)$*.*

The minimal such n is called the conductor of K. The embedding $K \hookrightarrow \mathbb{Q}(\zeta_n)$ induces a surjection $\text{Gal}(\mathbb{Q}(\zeta_n)/\mathbb{Q}) \twoheadrightarrow \text{Gal}(K/\mathbb{Q})$, via which we can try to describe the splitting behavior of primes: for example, p splits completely in K if and only if $\sigma_{\mathfrak{p}} \in \text{Gal}(K/\mathbb{Q})$ is the identity element 1, so equivalently it splits completely in K if and only if its Frobenius element in $\text{Gal}(\mathbb{Q}(\zeta_n)/\mathbb{Q})$ is in the kernel of this surjection. Since the cyclotomic Frobenius is just the class of p modulo n, this is easier to describe; for example, if $p \equiv 1 \pmod{n}$, then the cyclotomic Frobenius of p is trivial, so its image in $\text{Gal}(K/\mathbb{Q})$ must also be trivial and so p splits completely in K as well. This recovers the result described at the end of §1.2.

3.3 The Langlands program

Before we move on to our reinterpretation in terms of K-theory, it's worth discussing the limitations of our theory so far. One is that, other than the case $K = \mathbb{Q}$ and a few others, we do not have an explicit version of the theory: although we can probe the Galois groups using the reciprocity map, we cannot explicitly describe the abelian extensions of an arbitrary number field.

Perhaps a more prominent difficulty is the appearance everywhere of the condition that our extensions be abelian, or equivalently the fact that the reciprocity map can describe only the abelianization of the absolute Galois group, rather than the whole thing. Since we'd like to understand number fields and prime decomposition in the general case, not just in the abelian situation, it is interesting to ask if or how we could remove this condition.

At first glance, the formalism of class field theory is very poorly set up to deal with nonabelian groups: the source of the reciprocity map is $K^\times \backslash \mathbb{A}_K^\times$ in the global situation and $K_v^\times$ in the local situation, both of which are abelian by construction due to the commutativeness of multiplication on fields; and if we were to broaden our attention to non-commutative rings the scope of the problem becomes much broader than anything directly related to

number theory, and we can't hope to have the same sorts of tools available.[43] However, one thing that we could try is to generalize the operation of taking the group of units.

For any ring R, the group of units $R^\times$ is the same thing as $\mathrm{GL}_1(R)$, the automorphism group of R as an R-module.[44] This suggests a natural generalization: what if we replace GL_1 by GL_n, or even by any algebraic group G?

For $G = \mathrm{GL}_1$ abelian, it makes sense that the right thing to put as the source is the abelianization of Gal_K. However, for other groups G the right dependence is very unclear: e.g. GL_2 and GL_3 are both nonabelian, but we can't hope that the full group Gal_K is isomorphic to both $\mathrm{GL}_2(\mathbb{A}_K)$ and $\mathrm{GL}_3(\mathbb{A}_K)$ (or their quotients by $G(K)$).

Instead, we want to think about group representations. In the abelian case, every irreducible representation is one-dimensional, i.e. a character, and subject to certain restrictions we even have a certain duality: if the groups of characters of two groups are isomorphic, then so are the original two groups. Thus we can think of class field theory, instead of matching up (quantities related to) $\mathbb{A}_K^\times$ and $\mathrm{Gal}_K^{\mathrm{ab}}$, as (roughly) matching up one-dimensional representations of $\mathbb{A}_K^\times$ with one-dimensional representations of $\mathrm{Gal}_K^{\mathrm{ab}}$. Since one-dimensional representations of G^{ab} are the same thing as one-dimensional representations of G (since their target is abelian, they factor through G^{ab}), this means that class field theory roughly pairs irreducible representations of $\mathrm{GL}_1(\mathbb{A}_K)$ (all of which are automatically one-dimensional, since it's abelian) with irreducible one-dimensional representations of $\mathrm{Gal}_{\mathbb{Q}}$.

This suggests how we might hope to generalize to other groups: we hope to pair certain irreducible representations of $\mathrm{GL}_n(\mathbb{A}_K)$, called automorphic representations, with certain n-dimensional representations of $\mathrm{Gal}_{\mathbb{Q}}$. To understand a nonabelian group, we need to understand all of its representations, not just the one-dimensional ones (which give the abelianization); so to recover information about all of $\mathrm{Gal}_{\mathbb{Q}}$, we study the automorphic representations of every $\mathrm{GL}_n(\mathbb{A}_K)$.[45] The web of conjectures making this pairing precise and giving many of its properties is called the Langlands program.

A special case gives rise to the modularity conjecture, famously proven by Wiles, Taylor–Wiles, and Breuil–Conrad–Diamond–Taylor. When $n = 2$, certain automorphic representations can be associated to modular forms, which can be described as complex functions satisfying certain symmetries. On the other hand, certain 2-dimensional representations of $\mathrm{Gal}_{\mathbb{Q}}$ arise from algebraic-geometric objects called elliptic curves, when they are defined over $\mathbb{Q}$. It is not too hard, albeit well beyond the scope of this note, to give a construction of an elliptic curve over $\mathbb{Q}$ for every modular form (of a certain type); in this case we say that this elliptic curve is modular. The Langlands program predicts that in fact this should give a bijection between these modular forms and elliptic curves over $\mathbb{Q}$ with certain good properties; in particular every elliptic curve should be modular. This is the content of the modularity theorem, which Wiles used[46] to prove Fermat's last theorem.

There are many, many generalizations and interpretations of the Langlands program,

[43] Although with modern technology it may be possible to say something for some kind of noncommutative algebra...

[44] In fact the idea to think of the group of units as GL_1 already arises in class field theory in order to put the right topology on the group of units of the adeles, an issue we have skipped lightly over.

[45] There are also versions for other groups besides GL_n, but they are a little more complicated to describe.

[46] In part—strictly speaking at the time the modularity theorem was only proven for semistable curves, which sufficed for the application.

which can also be used to interpret class field theory in the case $G = \mathrm{GL}_1$. Including many of these would make this note much longer than it already is, but they are extremely interesting and have been heavily used to advance the field even in cases of more "classical" arithmetic interest; an excellent example is Fargues–Scholze's work on the geometrization of the local Langlands program [2].

One geometric interpretation worth mentioning briefly is the situation for function fields; we'll use some algebraic geometry in the following without explanation. See e.g. [3] for a more detailed reference.

Essentially all of global class field theory is true not just for number fields but for global fields; these can be defined abstractly, but in practice are either number fields or function fields, i.e. the fields $K(X)$ of rational functions on curves X defined over a finite field $\mathbb{F}_{\mathfrak{q}}$. Completing at each place of $K(X)$ (i.e. points of X) gives local function fields, analogous to the completions K_v of number fields; one important difference between the number field and function field cases is that for function fields, all places are nonarchimedean.

For simplicity, let's restrict to the unramified situation. An unramified extension of $K(X)$ is the same thing as an étale cover $Y \to X$, with corresponding field extension $K(Y)$. Thus the absolute unramified Galois group of $K(X)$, i.e. $\mathrm{Gal}(K(X)^{\mathrm{unr}}/K(X))$ where $K(X)^{\mathrm{unr}}$ is the maximal unramified extension, classifying all unramified extensions of $K(X)$ is the same thing as the étale fundamental group $\pi_1^{\text{ét}}(X)$ classifying étale covers of X.

On the other side, one can make the same definitions for the class group of $K(X)$; in this setting this is also known as the Picard group $\mathrm{Pic}(X)$ of X, and equivalently classifies line bundles on X. Per the Langlands philosophy, the most natural thing to compare is really irreducible representations of $\mathrm{Pic}(X)$ with one-dimensional representations of $\pi_1^{\text{ét}}(X)$.

Both of these have natural interpretations: a representation of $\pi_1^{\text{ét}}(X)$ is the same thing as a local system on X, and in this case we're interested in particular in local systems of rank 1 on X. These form a category which we denote by $\mathrm{Loc}_1(X)$. On the other hand, we can give the group $\mathrm{Pic}(X)$ the structure of (the product of some copies of $\mathbb{Z}$ with) a connected commutative algebraic group; characters of such a group G are then a special kind of rank 1 local system compatible with the group structure on G, which we call "character local systems" and which form a category we call $\mathrm{CharLoc}(G)$. Thus we get a categorical version of unramified global geometric class field theory:

$$\mathrm{CharLoc}(\mathrm{Pic}(X)) \simeq \mathrm{Loc}_1(X).$$

By a "decategorification" process one can recover a more traditional group-theoretic statement.

There are many ways of trying to generalize this sort of statement to the setting of the Langlands program, mostly encapsulated in the geometric Langlands philosophy: broadly, the idea is that the Galois side should be replaced by something related to the category of local systems on the base curve X (of a certain rank, or more generally $\check{G}$-local systems) while the "automorphic" or adelic side should be replaced by certain sheaves on the stack classifying G-bundles on X. This is a fruitful philosophy in the function field setting; a major question is how it can be applied to the number field setting, where our fields are no longer function fields of curves. In the local setting, Fargues–Scholze [2] have shown that the local Langlands program can be categorified as an incarnation of the geometric Langlands

program for an object called the Fargues–Fontaine curve, which is not a literal curve but an object coming out of exotic p-adic geometries. A global interpretation is, at present, unknown.

A different direction of generalization in the geometric context would be if we were to replace the curve X over $\mathbb{F}_q$ by a higher-dimensional scheme. In this case, the statements of class field theory do not hold, and naively don't even really make sense. There is some work on generalizing to higher dimensions (in the characteristic 0 setting as well), but for now it is very mysterious.

4. Class field theory III: K-theory

For either local or global fields F,[47] we have reciprocity maps

$$\Psi : \mathcal{A}_F \to \mathrm{Gal}_F^{\mathrm{ab}}$$

where $\mathcal{A}_F$ is some "automorphic" source term: for F a local field it is $F^\times$, while for a global field it is $F^\times \backslash \mathbb{A}_F^\times$. One can also view finite fields as having this reciprocity map, where $\mathcal{A}_F$ is the integers for F finite: there the map is given by sending 1 to the Frobenius element of $\mathrm{Gal}_F^{\mathrm{ab}} \simeq \mathrm{Gal}_F \simeq \widehat{\mathbb{Z}}$.

The goal of this section is to explain the paper [1] of Clausen, which gives a uniform K-theoretic construction of the reciprocity map Ψ for all three classes of field, global, local, and finite. In this section we freely use notions and terminology from algebraic geometry and algebraic topology.

In fact, the reciprocity map exists for any small idempotent-complete $\mathbb{Z}$-linear stable ∞-category $\mathcal{P}$: we have a natural transformation of presheaves of spectra

$$\Psi_{\mathcal{P}} : \mathrm{K}(\mathrm{lc}_{\mathcal{P}}) \to \mathrm{dK}^{\mathrm{Sel}}(\mathcal{P})$$

(definitions to follow) such that in the case that $\mathcal{P}$ is the derived ∞-category of complexes of F-modules where F is a global, local, or finite field, then $\pi_1 \,\mathrm{dK}^{\mathrm{Sel}}(\mathcal{P}) \simeq \mathrm{Gal}_F^{\mathrm{ab}}$, there is a natural map $\alpha_F : \mathcal{A}_F \to \pi_1 \mathrm{K}(\mathrm{lc}_{\mathcal{P}})$,[48] and

$$\pi_1 \Psi_{\mathcal{P}} \circ \alpha_F : \mathcal{A}_F \to \pi_1 \mathrm{K}(\mathrm{lc}_{\mathcal{P}}) \to \pi_1 \,\mathrm{dK}^{\mathrm{Sel}}(\mathcal{P}) \simeq \mathrm{Gal}_F^{\mathrm{ab}}$$

recovers (essentially) the usual reciprocity map $\Psi_F : \mathcal{A}_F \to \mathrm{Gal}_F^{\mathrm{ab}}$ as above.

Our goal for this section is to understand both sides of this K-theoretic reciprocity map; to get an idea of where the map itself comes from; and to try to understand why we recover the usual reciprocity map on π_1.

4.1 Locally compact K-theory

We want to define our source as the K-theory of some sort of category $\mathrm{lc}_{\mathcal{P}}$ of locally compact abelian group objects relative to $\mathcal{P}$. We can view this as arising from an absolute version

[47] We switch notations from K to F for our field to avoid conflict with the K of K-theory.

[48] This is an isomorphism for global and finite fields; it is not an isomorphism for local fields for topological reasons, which might be able to be fixed via a more sophisticated K-theoretic construction accounting for the topology on F, but in any case is not needed.

$\mathrm{lc}_{\mathbb{Z}}$, defined to be the bounded derived category $D^b(\mathrm{LCA})$ of the category LCA of second-countable locally compact Hausdorff abelian groups. We then define

$$\mathrm{lc}_{\mathcal{P}} = \mathrm{Fun}_{\mathbb{Z}}(\mathcal{P}, \mathrm{lc}_{\mathbb{Z}}),$$

the $\mathbb{Z}$-linear stable ∞-category of $\mathbb{Z}$-linear functors $\mathcal{P} \to \mathrm{lc}_{\mathbb{Z}}$. The source for our reciprocity map will be $\mathrm{K}(\mathrm{lc}_{\mathcal{P}})$.

As all of these objects $\mathrm{lc}_{\mathcal{P}}$ are in a sense obtained by base change from the base case $\mathrm{lc}_{\mathbb{Z}}$, equivalently the case where $\mathcal{P} = \mathrm{Perf}(\mathbb{Z})$, we can try to describe first the base case $\mathrm{lc}_{\mathbb{Z}}$ and its K-theory, and thence $\mathrm{lc}_{\mathcal{P}}$ and $\mathrm{K}(\mathrm{lc}_{\mathcal{P}})$ for $\mathcal{P} = \mathrm{Perf}(F)$ for local, global, or finite fields F. In the case $\mathcal{P} = \mathrm{Perf}(F)$ we'll abbreviate $\mathrm{lc}_F = \mathrm{lc}_{\mathrm{Perf}(F)}$.

In the base case, it turns out that we can understand $\mathrm{lc}_{\mathbb{Z}}$ via a cone construction, which we will not rigorously get into beyond the heuristic that it behaves roughly as we expect cones to behave:

$$\mathrm{cone}(\mathrm{Perf}(\mathbb{Z}) \to \mathrm{Perf}(\mathbb{R})) \xrightarrow{\sim} \mathrm{lc}_{\mathbb{Z}}\,.$$

This arises from the cofiber sequence

$$\mathbb{Z} \to \mathbb{R} \to \mathbb{R}/\mathbb{Z}$$

of locally compact abelian groups as follows: the natural inclusion induces the morphism $\mathrm{Perf}(\mathbb{Z}) \to \mathrm{Perf}(\mathbb{R})$ by base change $- \otimes_{\mathbb{Z}} \mathbb{R}$, and the $\mathbb{R}$-action on the locally compact abelian group $\mathbb{R}$ induces a morphism $\mathrm{Perf}(\mathbb{R}) \to \mathrm{lc}_{\mathbb{Z}}$, loosely speaking giving the "cofiber sequence" giving rise to the identification of the cone construction above. More explicitly, if we write $\underline{\mathbb{R}}$ for the locally compact abelian group $\mathbb{R}$ to avoid confusion, then the map $\mathrm{Perf}(\mathbb{R}) \to \mathrm{lc}_{\mathbb{Z}}$ can be thought of as (the induced map on perfect objects in the derived category of) $- \otimes_{\mathbb{R}} \underline{\mathbb{R}}$ on $\mathbb{R}$-modules. Applying K-theory (or indeed any functor satisfying localization), we get a canonical cofiber sequence

$$\mathrm{K}(\mathbb{Z}) \to \mathrm{K}(\mathbb{R}) \to \mathrm{K}(\mathrm{lc}_{\mathbb{Z}}).$$

We might hope to generalize this picture as follows, at least to the case where $\mathcal{P} = \mathrm{Perf}(R)$ for some commutative ring R: there is some R-algebra which is (additively) a locally compact abelian group, which we will suggestively call $\mathbb{A}_R$; the $\mathbb{A}_R$-action on $\mathbb{A}_R$ as a locally compact abelian group gives a functor $\mathrm{Perf}(\mathbb{A}_R) \to \mathrm{lc}_{\mathbb{Z}}$ and in fact to lc_R, since anything in the image of this functor carries an $\mathbb{A}_R$-action and thus an R-action by restriction; and the cofiber sequence

$$R \to \mathbb{A}_R \to \mathbb{A}_R/R$$

induces an map

$$\alpha_R : \mathrm{cone}(\mathrm{Perf}(R) \to \mathrm{Perf}(\mathbb{A}_R)) \to \mathrm{lc}_R,$$

though in general it need not be an isomorphism.

At least for global fields $R = F$, this works just as expected, taking $\mathbb{A}_F$ to be the adeles of F as usual. One can also make this construction work for finite fields F, viewed additively as discrete abelian groups, by taking $\mathbb{A}_F = 0$ with cofiber sequence

$$F \to 0 \to \Sigma F,$$

so we get a map

$$\mathrm{cone}(\mathrm{Perf}(F) \to 0) \to \mathrm{lc}_F\,.$$

For local fields, one has to modify the picture slightly to capture the completeness. Let $R = \mathcal{O}_F$ be the ring of integers of a nonarchimedean local field F with maximal ideal $\mathfrak{m}$, and write $\mathrm{Perf}_{\mathfrak{m}}(R) \subset \mathrm{Perf}(R)$ for the fiber of $\mathrm{Perf}(R) \to \mathrm{Perf}(F)$. For any R-module M, viewed as above as an object of $\mathrm{lc}_{\mathbb{Z}}$, one has a cofiber sequence

$$M \to 0 \to \Sigma M$$

as for finite fields[49] inducing

$$\alpha_F : \mathrm{cone}(\mathrm{Perf}_{\mathfrak{m}}(R) \to 0) \to \mathrm{lc}_F\,.$$

In this case, rather than an $\mathbb{A}_F$-action giving $\mathrm{Perf}(\mathbb{A}_F) \to \mathrm{lc}_{\mathbb{Z}}$ we have an F-action as in the case $F = \mathbb{R}$, giving a canonical F-linear comparison functor $\mathrm{Perf}(F) \to \mathrm{lc}_F$; for archimedean local fields this is the only relevant structure.

Applying K-theory gives the following morphisms: for F global,

$$\alpha_F : \mathrm{cofib}(\mathrm{K}(F) \to \mathrm{K}(\mathbb{A}_F)) \to \mathrm{K}(\mathrm{lc}_F);$$

for F finite,

$$\alpha_F : \mathrm{cofib}(\mathrm{K}(F) \to 0) = \Sigma\mathrm{K}(F) \to \mathrm{K}(\mathrm{lc}_F);$$

and for F local, we have a map

$$\alpha_F : \mathrm{K}(F) \to \mathrm{K}(\mathrm{lc}_F),$$

and (for F nonarchimedean) for $R = \mathcal{O}_F$ with maximal ideal $\mathfrak{m}$ a map

$$\alpha_R : \mathrm{cofib}(\mathrm{K}(\mathrm{Perf}_{\mathfrak{m}}(R)) \to 0) = \Sigma\mathrm{K}(\mathrm{Perf}_{\mathfrak{m}}(R)) \to \mathrm{K}(\mathrm{lc}_R).$$

In the global and finite cases, these maps are in fact isomorphisms; in the local cases, they are not, but only for topological reasons which in principal can be removed by taking into account the topologies of R and F.

Taking π_1 gives our desired comparison isomorphisms: in the global case, $\pi_1\mathrm{K}(F) = F^\times$ and $\pi_1\mathrm{K}(\mathbb{A}_F) = \mathbb{A}_F^\times$, so the map from the cofiber is

$$\alpha_F : \mathbb{A}_F^\times / F^\times \to \pi_1\mathrm{K}(\mathrm{lc}_F);$$

in the finite case, $\pi_1\Sigma\mathrm{K}(F) = \pi_0\mathrm{K}(F) = K_0(F) = \mathbb{Z}$ so we get

$$\alpha_F : \mathbb{Z} \to \mathrm{K}(\mathrm{lc}_F);$$

and in the local case we have $\pi_1\mathrm{K}(F) = F^\times$ and, for F nonarchimedean, $\pi_1\Sigma\mathrm{K}(\mathrm{Perf}_{\mathfrak{m}}(R)) = \pi_0\mathrm{K}(\mathrm{Perf}_{\mathfrak{m}}(R)) = \mathrm{K}_0(\mathrm{Perf}_{\mathfrak{m}}(R)) = \mathbb{Z}$, so we get

$$\alpha_F : F^\times \to \pi_1\mathrm{K}(\mathrm{lc}_F)$$

[49] Indeed, the finite field $M = R/\mathfrak{m}$ is a special case.

and

$$\alpha_R : \mathbb{Z} \to \pi_1 \mathrm{K}(\mathrm{lc}_R).$$

It is worth stating some compatibility properties of these maps. One is that they are functorial in F and compatible with norms under finite extensions; the details are left to the interested reader to work out (or see [1, Remark 3.32.1]). Perhaps more interesting is the compatibility between the global, local, and finite tiers: if F is a global field and F_v its completion at a place v, then we have a commutative diagram

$$\begin{array}{ccc} \mathrm{K}(F_v) & \xrightarrow{\alpha_{F_v}} & \mathrm{K}(\mathrm{lc}_{F_v}) \\ \downarrow & & \downarrow \\ \mathrm{cofib}(\mathrm{K}(F) \to \mathrm{K}(\mathbb{A}_F)) & \xrightarrow{\alpha_F} & \mathrm{K}(\mathrm{lc}_F) \end{array}$$

where the left vertical map is the composite $\mathrm{K}(F_v) \to \mathrm{K}(\mathbb{A}_F) \to \mathrm{cofib}(\mathrm{K}(F) \to \mathrm{K}(\mathbb{A}_F))$ induced by $F_v \hookrightarrow \mathbb{A}_F$ and the right vertical map is the forgetful functor. On π_1, this induces a commutative diagram

$$\begin{array}{ccc} F_v^\times & \xrightarrow{\alpha_{F_v}} & \pi_1 \mathrm{K}(\mathrm{lc}_{F_v}) \\ \downarrow & & \downarrow \\ \mathbb{A}_F^\times / F^\times & \xrightarrow{\alpha_F} & \pi_1 \mathrm{K}(\mathrm{lc}_F) \end{array}$$

giving the expected local-global compatibility for this comparison.

On the local side in the nonarchimedean case, we have a localization sequence

$$\mathrm{Perf}_{\mathfrak{m}}(R) \to \mathrm{Perf}(R) \to \mathrm{Perf}(F)$$

by definition and thus $\mathrm{K}(\mathrm{Perf}_{\mathfrak{m}}(R)) \to \mathrm{K}(R) \to \mathrm{K}(F)$, giving a morphism

$$\partial : \mathrm{K}(F) \to \Sigma \mathrm{K}(\mathrm{Perf}_{\mathfrak{m}}(R))$$

relating the F- and R-theories; and if $k = R/\mathfrak{m}$ is the associated finite field, then reduction gives a morphism $\iota : \mathrm{K}(k) \to \mathrm{K}(\mathrm{Perf}_{\mathfrak{m}}(R))$ as every perfect k-module naturally lives in $\mathrm{Perf}_{\mathfrak{m}}(R)$. These induce a commutative diagram

$$\begin{array}{ccc} \mathrm{K}(F) & \xrightarrow{\alpha_F} & \mathrm{K}(\mathrm{lc}_F) \\ \downarrow{\scriptstyle \partial} & & \downarrow \\ \Sigma \mathrm{K}(\mathrm{Perf}_{\mathfrak{m}}(R)) & \xrightarrow{\alpha_R} & \mathrm{K}(\mathrm{lc}_R) \\ \uparrow{\scriptstyle \iota} & & \uparrow \\ \Sigma \mathrm{K}(k) & \xrightarrow{\alpha_k} & \mathrm{K}(\mathrm{lc}_k) \end{array}$$

where the right vertical arrows are forgetful functors induced by $F \leftarrow R \to k$. On π_1, we get

$$\begin{array}{ccc} F^\times & \xrightarrow{\alpha_F} & \pi_1\mathrm{K}(\mathrm{lc}_F) \\ \downarrow & & \downarrow \\ \mathbb{Z} & \xrightarrow{\alpha_R} & \pi_1\mathrm{K}(\mathrm{lc}_R) \\ \| & & \uparrow \\ \mathbb{Z} & \xrightarrow{\alpha_k} & \pi_1\mathrm{K}(\mathrm{lc}_k) \end{array}$$

where the top left vertical map $\pi_1\partial$ is given by the discrete valuation $F^\times \to \mathbb{Z}$ (and ι is the identity on π_0). Thus in all cases we recover the desired maps $\alpha_F : \mathcal{A}_F \to \pi_1\mathrm{K}(\mathrm{lc}_F)$ together with the natural compatibilities, which is good evidence for $\mathrm{K}(\mathrm{lc}_{\mathcal{P}})$ as our K-theoretic version of the automorphic side.

4.2 Selmer K-homology

We now turn to the Galois side. The Langlands philosophy and especially the associated categorical perspective, c.f. §3.3, suggests that we should view class field theory as heuristically associating characters of the abelian group $\mathcal{A}_F$ to one-dimensional representations of $\mathrm{Gal}_F^{\mathrm{ab}}$, to be interpreted at least in the function field case as local systems on a curve. The perspective of §4.1 is that we should replace $\mathcal{A}_F$ by $\mathrm{K}(\mathrm{lc}_F)$, with $\mathcal{A}_F \to \pi_1\mathrm{K}(\mathrm{lc}_F)$ by viewing $\mathcal{A}_F$ as acting on locally compact F-modules. When F is a function field, its K-theory is naturally related to $\mathrm{Pic}(X)$ as the one-dimensional vector bundles; on the other hand local systems on X, i.e. $\pi_1^{\text{ét}}(X)$-representations, should have something to do with the étale K-theory of X, or in the more general situation with ramification the étale K-theory of X.[50] Since we want $\mathcal{A}_F \to \pi_1\mathrm{K}(\mathrm{lc}_F)$ to be the source of our map, rather than its dual, morally speaking the target of our K-theoretic reciprocity map ought to be a sort of dual to étale K-theory. This dual will be our Selmer K-homology.

We first need to define the relevant duality functors. There are essentially two pieces to the story: one is $K(1)$-localized[51] K-theory, while the other is topological cyclic homology. We need a duality functor in each setting. Fix a prime p for the remainder of the section.

The starting point is the fact that the determinant map $\det : \mathrm{K}(\mathbb{Z}_p) \to \mathrm{Pic}(H\mathbb{Z}_p)$, induced on the level of categories by sending a vector bundle to its determinant line bundle, lifts to $\mathrm{Pic}(\mathbb{S}_p)$ where $\mathbb{S}_p$ is the p-completed sphere. We could restrict along $\mathrm{K}(\mathbb{Z}) \to \mathrm{K}(\mathbb{Z}_p) \to \mathrm{Pic}(\mathbb{S}_p)$ to find that this map factors through $\mathrm{K}(\mathbb{Z}) \to \mathrm{K}(\mathbb{R})$, i.e. we have a commutative

[50] More specifically we should be thinking about the *Selmer K-theory* $\mathrm{K}^{\mathrm{Sel}}(F)$, which for number fields can be thought of as gluing Galois cohomology for $\ell \neq p$ with integrality conditions coming from p-adic Hodge theory at $\ell = p$, analogous to the definition of the Selmer groups. We'll see the explicit definition shortly, but it'll mostly just be used as motivation.

[51] Taking $K(1)$-localization is convenient and necessary for several of our constructions, but interestingly it is unclear whether it is really necessary in the broader picture.

diagram

$$\begin{array}{ccc} \mathrm{K}(\mathbb{Z}) & \longrightarrow & \mathrm{K}(\mathbb{Z}_p) \\ \downarrow & & \downarrow{\scriptstyle J_{\mathbb{Z}_p}} \\ \mathrm{K}(\mathbb{R}) & \xrightarrow{J_{\mathbb{R}}} & \mathrm{Pic}(\mathbb{S}_p) \end{array}$$

which morally we can think of as describing $\mathrm{K}(\mathbb{Z})$ via maps to $\mathrm{Pic}(\mathbb{S}_p)$ in terms of its finite and infinite places.

This partially motivates the following definition: set $\omega_{K(1)} = L_{K(1)} \mathrm{Pic}(\mathbb{S}_p)$, and $d_{K(1)} = \mathrm{Map}(-, \omega_{K(1)})$. In particular $J_{\mathbb{Z}_p}$ and $J_{\mathbb{R}}$ gives elements of $d_{K(1)}\mathrm{K}(\mathbb{Z}_p)$ and $d_{K(1)}\mathrm{K}(\mathbb{R})$, with the same image in $d_{\mathrm{K}(1)}\mathrm{K}(\mathbb{Z})$.

Now, we have a cyclotomic trace map $\mathrm{K} \to \mathrm{TC}$, which we will use to define the duality functor for TC. In particular, $\mathrm{tr} : \mathrm{K}(\mathbb{Z}_p) \to \mathrm{TC}(\mathbb{Z}_p)$ is an equivalence upon p-completion in nonnegative degrees, and so is an equivalence after $K(1)$-localization; so we have a map $\mathrm{TC}(\mathbb{Z})_p \to L_{K(1)} \mathrm{TC}(\mathbb{Z}_p) \simeq L_{K(1)}\mathrm{K}(\mathbb{Z}_p)$. We define ω_{TC} to be the $\mathrm{TC}(\mathbb{Z})$-module extending the above map to a cofiber sequence

$$\mathrm{TC}(\mathbb{Z})_p \to L_{K(1)}\mathrm{K}(\mathbb{Z}_p) \to \omega_{\mathrm{TC}},$$

and

$$d_{\mathrm{TC}} = \mathrm{Map}(-, \omega_{\mathrm{TC}}).$$

The map $J_{\mathbb{Z}_p}$ induces after $K(1)$-localization a map $j_{\mathbb{Z}_p} : L_{K(1)}\mathrm{K}(\mathbb{Z}_p) \to \omega_{K(1)}$, composing with which gives a map $- \cdot j_{\mathbb{Z}_p} : L_{K(1)}\mathrm{K}(\mathbb{Z}_p) \to d_{K(1)}L_{K(1)}\mathrm{K}(\mathbb{Z}_p)$. It turns out that in fact this map is an isomorphism. Composing with its inverse gives a natural transformation on $\mathrm{TC}(\mathbb{Z})$-modules

$$\begin{aligned} d_{K(1)} = \mathrm{Map}(-, \omega_{K(1)}) &\simeq \mathrm{Map}_{\mathrm{TC}(\mathbb{Z})}(-, d_{K(1)} \mathrm{TC}(\mathbb{Z})) \\ &\simeq \mathrm{Map}_{\mathrm{TC}(\mathbb{Z})}(-, d_{K(1)}L_{K(1)}\mathrm{K}(\mathbb{Z}_p)) \\ &\simeq \mathrm{Map}_{\mathrm{TC}(\mathbb{Z})}(-, L_{K(1)}\mathrm{K}(\mathbb{Z}_p)) \\ &\to \mathrm{Map}_{\mathrm{TC}(\mathbb{Z})}(-, \omega_{\mathrm{TC}}) = d_{\mathrm{TC}} \end{aligned}$$

where the first isomorphism is really just the restriction of the definition to $\mathrm{TC}(\mathbb{Z})$-modules, the second is the identification of $\mathrm{TC}(\mathbb{Z})$ with $\mathrm{K}(\mathbb{Z}_p)$ after $K(1)$-localization (as $d_{K(1)}$ factors through $K(1)$-localization), the third is composition with $- \cdot j_{\mathbb{Z}_p}$ as above, and the final map is given by composing with the defining morphism $L_{K(1)}\mathrm{K}(\mathbb{Z}_p) \to \omega_{\mathrm{TC}}$.

Now that we have our two duality functors and the transformation between them, we are ready to define Selmer K-theory: for a $\mathbb{Z}$-linear stable ∞-category $\mathcal{P}$ as above, we have localization maps $L_{K(1)} \mathrm{TC} \to \mathrm{TC}$ as well as the cyclotomic trace $\mathrm{K} \to \mathrm{TC}$, so localizing and taking the fiber product we define

$$\mathrm{K}^{\mathrm{Sel}}(\mathcal{P}) = L_{K(1)}\mathrm{K}(\mathcal{P}) \times_{L_{K(1)} \mathrm{TC}(\mathcal{P})} \mathrm{TC}(\mathcal{P}).$$

This suggests the correct definition of our dual: where $\mathrm{K}^{\mathrm{Sel}}$ is defined by a pullback, the Selmer K-homology is defined to be the pushout

$$\mathrm{dK}^{\mathrm{Sel}}(\mathcal{P}) = d_{K(1)}\mathrm{K}(\mathcal{P}) \sqcup_{d_{K(1)} \mathrm{TC}(\mathcal{P})} d_{\mathrm{TC}} \mathrm{TC}(\mathcal{P}).$$

Here the map $d_{K(1)}\,\mathrm{TC}(\mathcal{P}) \to d_{K(1)}\mathrm{K}(\mathcal{P})$ is dual to the cyclotomic trace, and $d_{K(1)}\,\mathrm{TC}(\mathcal{P}) \to d_{\mathrm{TC}}\,\mathrm{TC}(\mathcal{P})$ is the natural transformation $d_{K(1)} \to d_{\mathrm{TC}}$ above. We can think of this definition as gluing the "away from p" terms (the $K(1)$-local dual of K-theory) with the terms "at p" (the dual of topological cyclic homology). For example, if X is a "nice" algebraic space over $\mathbb{Z}$ and $\mathcal{P} = \mathrm{Perf}(X)$, then we can think of $d_{K(1)}\mathrm{K}(\mathcal{P})$ as describing the étale theory of $X_{\mathbb{Z}[1/p]}$, $d_{\mathrm{TC}}\,\mathrm{TC}(\mathcal{P})$ as describing the étale theory of the formal scheme $X_p^\wedge$, and $d_{K(1)}\,\mathrm{TC}(\mathcal{P})$ describing the étale theory of the rigid generic fiber $X_{\mathbb{Q}_p}^{\mathrm{an}}$, so the Selmer K-homology $\mathrm{dK}^{\mathrm{Sel}}(X) := \mathrm{dK}^{\mathrm{Sel}}(\mathcal{P})$ has to do with gluing together the behavior of X away from p and at p along the rigid generic fiber in a way related to p-adic Hodge theory. One can make this more precise by studying the difference between K-theory and TC, measured by the fiber of the cyclotomic trace.

The main property of $\mathrm{dK}^{\mathrm{Sel}}(\mathcal{P})$ that is important for our purposes is that we recover the Galois side from it as expected: when $\mathcal{P} = \mathrm{Perf}(F)$ for F a finite, local, or global field, we want to have

$$\pi_1\,\mathrm{dK}^{\mathrm{Sel}}(\mathcal{P}) \simeq \mathrm{Gal}_F^{\mathrm{ab}}\,.$$

This follows from the following proposition.

Proposition. *Let X be a locally noetherian derived algebraic space over $\mathbb{Z}$, with the reduced locus $(X \times_{\mathrm{Spec}\,\mathbb{Z}} \mathrm{Spec}\,\mathbb{F}_p)^{\mathrm{red}}$ regular. Then there is a natural map*

$$e_X : H_1(X_{\text{ét}}, \mathbb{Z}_p) \to \pi_1\,\mathrm{dK}^{\mathrm{Sel}}(X)$$

which is an isomorphism if either: X has mod p étale cohomological dimension at most 2 and $X \times_{\mathrm{Spec}\,\mathbb{Z}} \mathrm{Spec}\,\mathbb{F}_p$ has mod p étale cohomological dimension at most 1; or $X = \mathrm{Spec}\,F$ is the spectrum of a field with virtual mod p Galois cohomological dimension at most 2.

Here $H_1(X_{\text{ét}}, \mathbb{Z}_p)$ is the pro-p-abelianization of $\pi_1^{\text{ét}}(X)$, or equivalently the Pontryagin dual of $H^1(X_{\text{ét}}, \mathbb{Q}_p/\mathbb{Z}_p)$. If we're willing to assume $p \geq 3$, we can drop "virtual," and in the case of interest with $X = \mathrm{Spec}\,F$ the second situation suffices; but the compatibility diagrams for the automorphic side suggest that when F is a nonarchimedean local field, $X = \mathrm{Spec}\,\mathcal{O}_F$ is also of interest, which falls into the first case.

In particular, taking $X = \mathrm{Spec}\,F$ to be the spectrum of a finite, local, or global field, by standard cohomological dimension bounds we find that e_X gives an isomorphism

$$\pi_1^{\text{ét}}(\mathrm{Spec}\,F)_p^{\mathrm{ab}} = (\mathrm{Gal}_F)_p^{\mathrm{ab}} \xrightarrow{\sim} \pi_1\,\mathrm{dK}^{\mathrm{Sel}}(\mathrm{Perf}(F))$$

as desired, where by $(-)_p^{\mathrm{ab}}$ we mean the pro-p-abelianization. For $X = \mathrm{Spec}\,\mathcal{O}_F$ with F a nonarchimedean local field with residue field k, the inclusion $\mathrm{Spec}\,k \hookrightarrow \mathrm{Spec}\,\mathcal{O}_F$ induces an isomorphism on étale fundamental groups and so we get an isomorphism

$$\pi_1^{\text{ét}}(\mathrm{Spec}\,\mathcal{O}_F)_p^{\mathrm{ab}} \simeq \pi_1^{\text{ét}}(\mathrm{Spec}\,k)_p^{\mathrm{ab}} \simeq \mathbb{Z}_p \xrightarrow{\sim} \mathrm{dK}^{\mathrm{Sel}}(\mathrm{Perf}(\mathrm{Spec}\,\mathcal{O}_F)).$$

We have compatibilities between these identities similar to those sketched in §4.1.

The map e_X is constructed as an edge map in a "co-descent" spectral sequence for $\mathrm{dK}^{\mathrm{Sel}}$; the conditions for it to be an isomorphism are derived from a study of the values of $\mathrm{dK}^{\mathrm{Sel}}$ on (derived) strictly henselian local rings. The particular choices of the dualities used in the definition of $\mathrm{dK}^{\mathrm{Sel}}$ are crucial to make this construction work.

4.3 The reciprocity map

We finally have both sides of our map defined, and we know how to recover our classical source and target on π_1. What remains is to define our K-theoretic reciprocity map

$$\Psi_{\mathcal{P}} : \mathrm{K}(\mathrm{lc}_{\mathcal{P}}) \to \mathrm{dK}^{\mathrm{Sel}}(\mathcal{P}),$$

which should have the property that when $\mathcal{P} = \mathrm{Perf}(F)$ for F finite, local, or global, we have $\pi_1 \Psi_{\mathcal{P}} = \Psi_F$, the classical reciprocity map.

Recall from §4.1 that we identified $\mathrm{lc}_{\mathbb{Z}}$ with the cone $\mathrm{cone}(\mathrm{Perf}(\mathbb{Z}) \to \mathrm{Perf}(\mathbb{R}))$. There, we applied K-theory to get a cofiber sequence; but in fact we can just as well apply Selmer K-homology, as it also satisfies localization, to dually get a fiber sequence

$$\mathrm{dK}^{\mathrm{Sel}}(\mathrm{lc}_{\mathbb{Z}}) \to \mathrm{dK}^{\mathrm{Sel}}(\mathbb{R}) \to \mathrm{dK}^{\mathrm{Sel}}(\mathbb{Z}).$$

Thus given a point of $\mathrm{dK}^{\mathrm{Sel}}(\mathbb{R})$ and a null-homotopy of its image in $\mathrm{dK}^{\mathrm{Sel}}(\mathbb{Z})$, we can construct a point of $\mathrm{dK}^{\mathrm{Sel}}(\mathrm{lc}_{\mathbb{Z}})$.

This is the data of a point of $\mathrm{dK}^{\mathrm{Sel}}(\mathbb{R}) \simeq d_{K(1)}\mathrm{K}(\mathbb{R})$; a point of $d_{K(1)}\,\mathrm{TC}(\mathbb{Z})$; a homotopy between the images of these points in $d_{K(1)}\mathrm{K}(\mathbb{Z})$; and a null-homotopy of the image of the second point in $d_{\mathrm{TC}}\,\mathrm{TC}(\mathbb{Z})$. Recall the maps $J_{\mathbb{Z}_p} : \mathrm{K}(\mathbb{Z}_p) \to \mathrm{Pic}(\mathbb{S}_p)$, $J_{\mathbb{R}} : \mathrm{K}(\mathbb{R}) \to \mathrm{Pic}(\mathbb{S}_p)$, whose pullbacks to $\mathrm{K}(\mathbb{Z})$ agree (up to homotopy). Localizing both these maps at $K(1)$ gives two maps $j_{\mathbb{Z}_p} : L_{K(1)}\mathrm{K}(\mathbb{Z}_p) \to \omega_{K(1)}$, $j_{\mathbb{R}} : L_{K(1)}\mathrm{K}(\mathbb{R}) \to \omega_{K(1)}$, i.e. points of $d_{K(1)}\mathrm{K}(\mathbb{Z}_p) \simeq d_{K(1)}\,\mathrm{TC}(\mathbb{Z})$ and $d_{K(1)}\mathrm{K}(\mathbb{R})$ respectively, which gives the first two pieces of data; the homotopy between their images in $d_{K(1)}\mathrm{K}(\mathbb{Z})$ corresponds to the homotopy between pullbacks of $J_{\mathbb{Z}_p}$ and $J_{\mathbb{R}}$ to $\mathrm{K}(\mathbb{Z})$; and the image of $j_{\mathbb{Z}_p}$ in $d_{\mathrm{TC}}\,\mathrm{TC}(\mathbb{Z})$ is its image under the natural transformation $d_{K(1)} \to d_{\mathrm{TC}}$, which is constructed via the inverse of composing with $j_{\mathbb{Z}_p}$ and so necessarily kills $j_{\mathbb{Z}_p}$. Assembling this data, we find that we can produce a point of $\mathrm{dK}^{\mathrm{Sel}}(\mathrm{lc}_{\mathbb{Z}})$, which we call j.

Now, $\mathrm{K}^{\mathrm{Sel}}$ is constructed from K-theory, TC, and the cyclotomic trace, all of which are multiplicative under tensor products of the $\mathbb{Z}$-linear ∞-category $\mathcal{P}$; and all the terms defining $\mathrm{dK}^{\mathrm{Sel}}$ are $\mathrm{K}^{\mathrm{Sel}}$-linear, so so is the pushout $\mathrm{dK}^{\mathrm{Sel}}$ itself. On the other hand there is a natural "evaluation" map $\mathrm{lc}_{\mathcal{P}} \otimes_{\mathbb{Z}} \mathcal{P} = \mathrm{Fun}_{\mathbb{Z}}(\mathcal{P}, \mathrm{lc}_{\mathbb{Z}}) \otimes_{\mathbb{Z}} \mathcal{P} \to \mathrm{lc}_{\mathbb{Z}}$, which on $\mathrm{dK}^{\mathrm{Sel}}$ induces

$$\mathrm{dK}^{\mathrm{Sel}}(\mathrm{lc}_{\mathbb{Z}}) \to \mathrm{dK}^{\mathrm{Sel}}(\mathrm{lc}_{\mathcal{P}} \otimes_{\mathbb{Z}} \mathcal{P}) \to \mathrm{dK}^{\mathrm{Sel}}(\mathrm{lc}_{}$$

so tensoring with $\mathrm{K}^{\mathrm{Sel}}(\mathrm{lc}_{\mathcal{P}})$ gives

$$\mathrm{K}^{\mathrm{Sel}}(\mathrm{lc}_{\mathcal{P}}) \otimes \mathrm{dK}^{\mathrm{Sel}}(\mathrm{lc}_{\mathbb{Z}}) \to \mathrm{K}^{\mathrm{Sel}}(\mathrm{lc}_{\mathcal{P}}) \otimes \mathrm{dK}^{\mathrm{Sel}}(\mathrm{lc}_{\mathcal{P}} \otimes_{\mathbb{Z}} \mathcal{P}) \to \mathrm{K}^{\mathrm{Sel}}(\mathrm{lc}_{\mathcal{P}}) \to \mathrm{dK}^{\mathrm{Sel}}(\mathcal{P}).$$

On the other hand, as K-theory maps naturally to both $L_{K(1)}\mathrm{K}$ and TC via the localization and cyclotomic trace maps (compatibly with the further maps to $L_{K(1)}\,\mathrm{TC}$), there is a natural map $\mathrm{K} \to \mathrm{K}^{\mathrm{Sel}}$. Composing, we get a map

$$\mathrm{K}(\mathrm{lc}_{\mathcal{P}}) \otimes \mathrm{dK}^{\mathrm{Sel}}(\mathrm{lc}_{\mathbb{Z}}) \to \mathrm{dK}^{\mathrm{Sel}}(\mathcal{P}),$$

which we can view as a pairing. We have constructed a point $j \in \mathrm{dK}^{\mathrm{Sel}}(\mathrm{lc}_{\mathbb{Z}})$, so we can evaluate this pairing on it to get a map

$$\mathrm{K}(\mathrm{lc}_{\mathcal{P}}) \to \mathrm{dK}^{\mathrm{Sel}}(\mathcal{P});$$

this is our reciprocity map $\Psi_{\mathcal{P}}$.

We can understand it concretely in terms of j: if we think of a point in $\mathrm{K}(\mathrm{lc}_{\mathcal{P}})$ as represented by an object $(F : \mathcal{P} \to \mathrm{lc}_{\mathbb{Z}}) \in \mathrm{lc}_{\mathcal{P}} = \mathrm{Fun}_{\mathbb{Z}}(\mathcal{P}, \mathrm{lc}_{\mathbb{Z}})$, then applying $\mathrm{dK}^{\mathrm{Sel}}$ to the functor F gives a map $\mathrm{dK}^{\mathrm{Sel}}(F) : \mathrm{dK}^{\mathrm{Sel}}(\mathrm{lc}_{\mathbb{Z}}) \to \mathrm{dK}^{\mathrm{Sel}}(\mathcal{P})$, and $\Psi_{\mathcal{P}}(F) \in \mathrm{dK}^{\mathrm{Sel}}(\mathcal{P})$ will be the image of j under $\mathrm{dK}^{\mathrm{Sel}}(F)$.

On π_1, composing with our maps α_F and e_F in the case $\mathcal{P} = \mathrm{Perf}(F)$ for F finite, local, or global gives

$$\mathcal{A}_F \xrightarrow{\alpha_F} \pi_1 \mathrm{K}(\mathrm{lc}_F) \xrightarrow{\pi_1 \Psi_{\mathcal{P}}} \pi_1 \, \mathrm{dK}^{\mathrm{Sel}}(\mathcal{P}) \simeq (\mathrm{Gal}_F)^{\mathrm{ab}}_p.$$

Note that this recovers only the pro-p part of the usual reciprocity map; but taking the system of these p-reciprocity maps over all primes p lets us broaden the target to the full abelianization $\mathrm{Gal}_F^{\mathrm{ab}}$. By checking the behavior of this composition on Frobenius elements (equivalently on finite fields, and assembling the various compatibility diagrams) we can pin down that this is indeed the Artin reciprocity map of class field theory, with the caveat that it is off by a sign from the usual convention. This has to do with the fact that there are actually two natural ways to construct the reciprocity map in terms of Frobenius elements, via the "arithmetic" or "geometric" Frobenius elements, which are inverses of each other; so long as we keep track of which convention we are using there is no issue.

We can at this point see some pros and cons of this construction, with an eye to generalization. On the one hand, the generality is already vastly greater than classical class field theory; it applies to $\mathbb{Z}$-linear stable ∞-categories $\mathcal{P}$ rather than only to certain fields. In addition, it has the great virtue of treating all the relevant fields uniformly, and suggests how we might generalize to higher-dimensional class field theory, or even noncommutative rings, by replacing e.g. the ideles with "higher ideles" $\pi_1 \mathrm{K}(\mathrm{lc}_R)$.

On the other hand, it does not seem obviously well-suited to the problems suggested by the Langlands philosophy of replacing $\mathbb{G}_{\mathrm{m}}$ with more general algebraic groups, and the nonabelian part of the Galois side seems difficult to access via homological methods. Nevertheless one can try to picture solutions: perhaps replacing $\mathrm{K}(\mathrm{lc}_F)$ with some theory in which we require our objects to have G-action on the source, and some sort of nonabelian homology on the target. Such speculation is left to the reader—but if you have interesting ideas or know of developments on this front I'd love to hear about them.

References

[1] Dustin Clausen. A K-theoretic approach to Artin maps. *arXiv preprint arXiv:1703.07842*, 2017.

[2] Laurent Fargues and Peter Scholze. Geometrization of the local Langlands correspondence. *arXiv preprint arXiv:2102.13459*, 2021.

[3] Tony Feng and Bhargav Bhatt. Geometric class field theory. *Lecture notes available at* `https://math.berkeley.edu/~fengt/2GeometricCFT.pdf`, 2016.

[4] Chao Li and Joe Rabinoff. Class field theory. *Lecture notes available at* `www.math.columbia.edu/~chaoli/docs/ClassFieldTheory.html`, 2012.