KMB, 14/4/16

M345P11: Cyclotomic fields.

1 Introduction.

This note is about the Galois theory of cyclotomic extensions, which is a fancy way of saying extensions of a field generated by roots of unity. Working through it would be an interesting way of revising for the exam, because the Galois theory of cyclotomic extensions is for the most part a pretty straightforward consequence of standard results from the course.

2 The question.

Say K is a field, of characteristic zero if you like, and say $n \geq 1$ is an integer. Let L be the splitting field of $X^n - 1$ over K. Then L is a splitting field of a polynomial, and so L/K is finite and normal; moreover K has characteristic zero so L/K is separable as well. We could hence ask what $\mathrm{Gal}(L/K)$ is. The field L is called a *cyclotomic extension* of K.

The case we'll be mainly concerned with is when $K = \mathbf{Q}$. Then L is called a *cyclotomic field* and there is an amazingly simple and beautiful answer to the question – in this case $\mathrm{Gal}(L/\mathbf{Q}) = (\mathbf{Z}/n\mathbf{Z})^\times$; the Galois group is (canonically) isomorphic to the units in the ring $\mathbf{Z}/n\mathbf{Z}$, that is, the group of numbers modulo n which are coprime to n. I'll prove this in this note. But first let me do some general theory.

3 Cyclotomic extensions: the general case.

Let's say K is now any field at all (so in particular it could have characteristic $p > 0$) and $n \in \mathbf{Z}_{\geq 1}$, and let's think about the splitting field of $x^n - 1$ over K. The first thing to note is that if K does have characteristic $p > 0$ and if also $n = pm$ is a multiple of p, then $x^n - 1 = (x^m - 1)^p$ and in particular the splitting fields of $x^n - 1$ and $x^m - 1$ are the same. So if K has characteristic $p > 0$ then we may as well assume that n is coprime to p. From now on let's assume this – so we'll always assume either that K has characteristic zero (and then $n \geq 1$ can be anything) or K has characteristic $p > 0$ but n is coprime to p. Another way of saying this is simply that $n \neq 0$ in K.

Under these assumptions, we have the nice consequence that the n roots of $x^n - 1$ are all distinct in the splitting field. For if $x^n - 1$ had a repeated root, it would have a root in common with its derivative nx^{n-1}. But these two polynomials clearly have no root in common, because $n \neq 0$ in K and hence the only root of nx^{n-1} is zero, which is not a root of $x^n - 1$. So the splitting field of $x^n - 1$ over K is generated by n distinct nth roots of unity.

The next key thing to realise is that some nth roots of unity are different to others. For example let's imagine $K = \mathbf{Q}$ and let's take the four 4th roots of 1 in the complex numbers; they are 1, -1, and $\pm i$. In this case the splitting field is $\mathbf{Q}(i)$, and we can see the asymmetry – two of the 4th roots of unity (namely ± 1) do not generate $\mathbf{Q}(i)$ over $\mathbf{Q}$, whereas the other two do. Why has this happened? It's because even though 1 and -1 are 4th roots of unity in the sense that their 4th powers are equal to 1, they are not very sensible 4th roots of unity because -1 is actually a square root of unity, and $+1$ is unity itself: the orders of $+1$ and -1 in the multiplicative group $\mathbf{Q}(i)^\times$ are 1 and 2, not 4. On the other hand the other two roots $\pm i \in \mathbf{Q}(i)^\times$ really do have order 4 in this group. We need some notation for this important concept – and here it is. Say K is any field and $n \neq 0 \in K$ as usual. Let L/K be the splitting field for $x^n - 1$. We say that an element $\zeta \in L$ is an *nth root of unity* if $\zeta^n = 1$ and we furthermore say that ζ is a *primitive* nth root of unity if ζ has exact order n in the group $L^\times$. The example of $K = \mathbf{Q}$ and $n = 4$ shows us that nth roots of unity are sometimes primitive ($\pm i$) and sometimes not (± 1).

What we need next is that even in this general situation, there will *always* be at least one primitive nth root of unity. This is obvious if $K = \mathbf{Q}$ (because then we can regard L as a subfield of the complex numbers, and $e^{2\pi i/n}$ has order exactly n) but let's prove it in general. In fact what is going on is an easy consequence of a standard fact. Recall from (2015-2016) Problem Sheet 3 Q3 that if L is a field then a finite subgroup of $L^\times$ is automatically cyclic. So let K be a field, let $n \in \mathbf{Z}_{\geq 1}$ be non-zero in K as usual, and let L be the splitting field of $x^n - 1$ over K. Then it is easy to check that the n roots of $x^n - 1$ in L form a subgroup of $L^\times$, with the group law being multiplication. Hence this group is cyclic, of order n (as all the roots are in the splitting field and no root has multiplicity greater than 1 by our assumption on n) and if we choose a generator then this is a primitive nth root of unity.

So let's choose one, and let's call it ζ, or sometimes ζ_n if we want to keep track of what its order is. Then we can list all the nth roots of unity in L – they are $\{1 = \zeta^0, \zeta = \zeta^1, \zeta^2, \zeta^3, \ldots, \zeta^{n-1}\}$. In particular $L = K(\zeta)$ and the splitting field is obtained from K by adjoining just one root of $x^n - 1$ to K (but you have to make sure you adjoin the right one).

4 Galois theory of cyclotomic extensions – the general case.

Recall again the set-up: K is a field, $n \geq 1$ is an integer such that n is non-zero in K, and $L = K(\zeta)$ is the splitting field of $x^n - 1$ over K, with ζ an nth root of unity that is *primitive*, that is, $\zeta^n = 1$ but $\zeta^m \neq 1$ for any $1 \leq m < n$.

Now say $\sigma \in \mathrm{Gal}(L/K)$. What can $\sigma(\zeta)$ be? Well $\sigma(\zeta)^n = \sigma(\zeta^n) = \sigma(1) = 1$, so $\sigma(\zeta)$ is also definitely an nth root of unity. Furthermore $\sigma(\zeta)$ must be a *primitive* nth root of unity, as σ is an isomorphism of fields, so restricted to $L^\times$ it's an isomorphism of groups $L^\times \to L^\times$, and an isomorphism between groups sends elements of order n to elements of order n.

This means that $\sigma(\zeta) = \zeta^j$ for some j, and furthermore we know something more about j – we know that ζ^j has exact order n. Now recall from basic group theory that if ζ has order n then ζ^j has order n/d, where d is the highest common factor of n and j. In our case we want ζ^j to have order n so we need $d = 1$, which means the highest common factor of n and j is 1, which means that n and j must be coprime.

So in fact we now have a map from $\mathrm{Gal}(L/K)$ to the set of numbers j with $1 \leq j \leq n$ and furthermore j coprime to n. Furthermore, this map is injective, because $L = K(\zeta)$ and so any element of $\mathrm{Gal}(L/K)$ is determined by what it does to ζ, and j tells us this. If you think about it, j is only really defined as an integer modulo n, so it might be nicer to think of this as a map from $\mathrm{Gal}(L/K)$ to $\mathbf{Z}/n\mathbf{Z}$, with the image contained in the subset of $\mathbf{Z}/n\mathbf{Z}$ consisting of things coprime to n.

But now here's something better. The numbers between 1 and n coprime to n actually form a group! The neatest way to see this is that $\mathbf{Z}/n\mathbf{Z}$ is a ring, and the units in a ring form a group, and the units $(\mathbf{Z}/n\mathbf{Z})^\times$ in the ring $\mathbf{Z}/n\mathbf{Z}$ are precisely the cosets containing elements coprime to n.

The conclusion so far is that if we choose a primitive nth root of unity then we have an injective map from the group $\mathrm{Gal}(L/K)$ to the group $(\mathbf{Z}/n\mathbf{Z})^\times$. Let's call this map θ. Recall that it's defined like this: once we have fixed ζ a primitive nth root of unity (recall that this just means that ζ has exact order n) then $\theta(\sigma) = j$ if $\sigma(\zeta) = \zeta^j$.

Here are two very cool lemmas about this map θ.

Lemma 1. *θ is a group homomorphism.*

Proof. We need to carefully unravel the definitions. Say $\theta(\sigma) = j$ and $\theta(\tau) = k$. We need to check $\theta(\sigma\tau) = jk$. In other words, we need to check that if $\sigma(\zeta) = \zeta^j$ and $\tau(\zeta) = \zeta^k$ then $\sigma(\tau(\zeta)) = \zeta^{jk}$.

But this is OK, because $\tau(\zeta) = \zeta^k$, so $\sigma(\tau(\zeta)) = \sigma(\zeta^k) = (\sigma(\zeta))^k = (\zeta^j)^k = \zeta^{jk}$. We're done. □

Lemma 2. *θ does not depend on the choice ζ of primitive nth root of unity!*

Proof. What I'm saying here is that if ω is another primitive nth root of unity, and we define a new map θ' by $\theta'(\sigma) = j'$ if $\sigma(\omega) = \omega^{j'}$, then θ equals θ'. The reason for this is that ω must be ζ^i for some i, and if $\sigma(\zeta) = \zeta^j$ then $\sigma(\omega) = \sigma(\zeta^i) = \sigma(\zeta)^i = \zeta^{ji} = \zeta^{ij} = \omega^j$. □

One could say that the previous lemma says that the injection $\mathrm{Gal}(L/K) \to (\mathbf{Z}/n\mathbf{Z})^\times$ is *canonical*, but do not worry if you don't understand what this means: it just means that whatever choices you make to define it, the result is independent of those choices. The upshot is that $\mathrm{Gal}(L/K)$ is naturally a subgroup of $(\mathbf{Z}/n\mathbf{Z})^\times$.

We can't really go much further than this in this generality – what $\mathrm{Gal}(L/K)$ is depends on, for example, how many nth roots of unity were already in K.

5 The case $K = \mathbf{Q}$.

In this case we can go further. First let's have a quick recap, with the added simplicity that the case $K = \mathbf{Q}$ gives us. We can let L be the subfield of the complex numbers generated by the nth roots of unity. So $L = \mathbf{Q}(\zeta)$ with $\zeta = e^{2\pi i/n}$. If $\sigma \in \mathrm{Gal}(L/\mathbf{Q})$ then $\sigma(\zeta) = \zeta^j$ for some $1 \leq j \leq n$ and if we set $\theta(\sigma) = j$ then θ is a map from $\mathrm{Gal}(L/K)$ to $(\mathbf{Z}/n\mathbf{Z})^\times$, and this map is an injective group homomorphism. However in this case we have

Proposition 1. $\theta : \mathrm{Gal}(\mathbf{Q}(\zeta)/\mathbf{Q}) \to (\mathbf{Z}/n\mathbf{Z})^\times$ *is an* isomorphism *of groups.*

Proof. I always think the proof is rather delicate. Recall $\zeta = e^{2\pi i/n}$. Let $f(x)$ be the min poly of ζ over $\mathbf{Q}$. Then $f(x)$ divides $x^n - 1$ so $x^n - 1 = f(x)h(x)$ for some polynomial $h \in \mathbf{Q}[x]$. By Gauss' lemma (and the fact that $x^n - 1$ and $f(x)$ are monic) we can deduce that $f(x)$ and $h(x)$ are actually in $\mathbf{Z}[x]$.

Now let p be a prime not dividing n and let's consider $\omega = \zeta^p$. Then ω is a root of $x^n - 1$ so it's either a root of $f(x)$ or $h(x)$. The strategy is first to show that ω must be a root of $f(x)$, and then to show that this will suffice to prove the Proposition.

First let's show that ω can't be a root of $h(x)$. For if it were, then ζ would be a root of the polynomial $h(x^p)$, and so by the basic property of the min poly we would have $f(x)$ divides $h(x^p)$ in $\mathbf{Q}(x)$ and hence in $\mathbf{Z}[x]$ (just divide $h(x^p)$ by the monic polynomial $f(x)$ using long division and note that the answer must be in $\mathbf{Z}[x]$ because $f(x)$ is monic).

But this is bad, because if we now reduce everything mod p we see that the mod p reduction $\overline{f}(x) \in (\mathbf{Z}/p\mathbf{Z})[x]$ of $f(x)$ divides $\overline{h}(x^p) = \left(\overline{h}(x)\right)^p$, and because f has positive degree this means that $\overline{f}(x)$ and $\overline{h}(x)$ must have a root in common in M, the splitting field of $x^n - 1$ over the field $\mathbf{Z}/p\mathbf{Z}$. But this cannot be, because their product is $x^n - 1$ which has distinct roots in M as it has no factor in common with its derivative (that's where we used the fact that p did not divide n), and this is a contradiction.

Our conclusion then is that if p is a prime not dividing n then $\omega = \zeta^p$ is also a root of $f(x)$. But because $f(x) \in \mathbf{Q}[x]$ is irreducible, and ζ and ω are roots, we know from section 2 of the course that $\mathbf{Q}(\zeta)$ and $\mathbf{Q}(\omega)$ are isomorphic and that there's a field isomorphism $\mathbf{Q}(\zeta) \to \mathbf{Q}(\omega)$ sending ζ to ω. And because $\mathbf{Q}(\omega) = \mathbf{Q}(\zeta)$ this field isomorphism is an element σ of $\mathrm{Gal}(\mathbf{Q}(\zeta)/\mathbf{Q})$ with the property that $\theta(\sigma) = p$.

We have proved that for every prime p not dividing n, the element $p \in (\mathbf{Z}/n\mathbf{Z})^\times$ is in the image of θ. But this is enough, because if $1 \leq j \leq n$ is coprime to n then write j as a product of primes; each prime is in the image of θ and so j is too. □

Recall that the Euler phi function $\phi(n)$ is defined by $\phi(n)$ equals the number of integers between 1 and n which are coprime to n. We have just proved that this is the size of $\mathrm{Gal}(\mathbf{Q}(\zeta)/\mathbf{Q})$. So we have showed

Corollary 3. $[\mathbf{Q}(\zeta) : \mathbf{Q}] = \phi(n)$.

Proof. Fundamental theorem of Galois theory part (a). □

We have also showed that the min poly of $\zeta = e^{2\pi i/n}$ has degree equal to $\phi(n)$. In fact the proof of the Proposition shows more – it showed that if $f(x)$ was the min poly of our primitive root ζ then ζ^p was also a root of $f(x)$ for all primes p not dividing n. Because every positive integer is a product of primes, and our proof works for all primitive roots of unity, this means that ζ^j is a root of $f(x)$ for every j coprime to n. In particular the $\phi(n)$ roots of $f(x)$ must just be the primitive nth roots of unity, and the $h(x)$ in the proof must be the polynomial whose roots are the nth roots of unity that are not primitive nth roots.

Let $\Phi_n(x)$ denote the min poly of $e^{2\pi i/n}$ (so $\Phi_n(x) = f(x)$ in the notation of the Proposition). Then Φ_n is called the nth *cyclotomic polynomial*. These polynomials are quite cool. Here are some basic examples: $\Phi_1(x) = x - 1$. $\Phi_2(x) = x + 1$. $\Phi_3(x) = (x^3 - 1)/(x - 1) = x^2 + x + 1$. $\Phi_4(x) = x^2 + 1$. $\Phi_5(x) = x^4 + x^3 + x^2 + x + 1$, and $\Phi_6(x) = x^2 - x + 1$.

Note that every nth root of unity is a primitive mth root of unity for some m dividing n (because the order of the element divides the order of the group), and hence $x^n - 1 = \prod_{m|n} \Phi_m(x)$. Taking degrees of both sides we deduce that $n = \sum_{m|n} \phi(m)$, a formula familiar from the basic number theory course. This also gives us a recursive way of computing the Φ_n: if we know Φ_m for all $m < n$ (and in particular for all m dividing n apart from $m = n$ itself) we know $\Phi_n = (x^n - 1)/\prod_{m|n, m\neq n} \Phi_m$.

Here's a crazy and slightly confusing exercise. If you have access to a computer, or google, then compute Φ_n for $1 \leq n \leq 100$. What do you notice about the coefficients of these polynomials? (hint: they are small). Do you think it's true for all n? Is it true for all $n \leq 105$? There is an old paper by Erdős on this which you can find by googling; apparently there's earlier work of Schur (he of the lemma), but I don't know a precise reference.

6 M4 bonus.

This last section is for people who have read the extra M4 hand-out on infinite Galois groups.

Note that the material above fills in the gap I left in Q10 of the M4 example sheet. We can also go further with the ideas in Q10 now.

One can consider the union of the splitting fields of $x^n - 1$ for all n at once, as a subfield of the complexes. If L denotes this field then L is the union of $\mathbf{Q}(\zeta_n)$ for all $n \geq 1$. We have $\mathbf{Q}(\zeta_m) \subseteq \mathbf{Q}(\zeta_n)$ iff m divides n, so $\mathrm{Gal}(L/\mathbf{Q}) = \varprojlim_n (\mathbf{Z}/n\mathbf{Z})^\times$, where the projective limit is over all positive integers, given the structure of a directed set by $m \leq n$ iff m divides n. Exercise: check this is a directed set. Because the ring $(\mathbf{Z}/n\mathbf{Z})$ is isomrphic to the product of the rings $(\mathbf{Z}/p_i^{e_i}\mathbf{Z})$ if $n = \prod_i p_i^{e_i}$ (this is just a fancy form of the Chinese Remainder Theorem) we see $(\mathbf{Z}/n\mathbf{Z})^\times = \prod_i (\mathbf{Z}/p_i^{e_i}\mathbf{Z})^\times$ and after some thought one can check that the projective limit is actually $\widehat{\mathbf{Z}}^\times = \prod_p \mathbf{Z}_p^\times$, where here $\widehat{\mathbf{Z}}$ is the profinite completion of $\mathbf{Z}$ (which by definition is $\varprojlim_n (\mathbf{Z}/n\mathbf{Z})$). So there's another example of an infinite Galois group.