GALOIS THEORY: LECTURE 20

LEO GOLDMAKHER

1. Review: The Fundamental Lemma

We begin today's lecture by recalling the Fundamental Lemma introduced at the end of Lecture 19. This will come up in several places during today's lecture and will be helpful to have fresh in our minds.

Lemma 1.1 (Fundamental Lemma)**.** *Given a finite and Galois extension L/K and some $\alpha \in L$, we can write the minimal polynomial of α over K as*

$$m_\alpha(x) = \prod_{\beta \in A} (x - \beta)$$

where $A := \{\sigma(\alpha) : \sigma \in G\}$.

Remark. In other words, the minimal polynomial is defined to be the one whose roots are precisely all the distinct Galois conjugates of α.

This is a powerful result – it gives an explicit connection between the Galois group of L/K and the elements living in L. We've already seen it used previously (implicitly) in proofs. Now we see a new application, to the study of normal extensions.

2. Normal Extensions

Recall that we have three equivalent definitions of what it means for L/K to be Galois:

- $|G| = [L : K]$,
- L is the splitting field of some separable polynomial in $K[x]$, and
- $L^G = K$.

(As usual, $G := \mathrm{Aut}(L/K)$.) The goal of this section is to introduce and prove a fourth equivalent condition. Not only is it a useful criterion, it is also the one usually given as the definition of Galois-ity in most courses on Galois theory. To motivate the idea, we start with a question.

Question 1. *Given L/K a finite (but not necessarily Galois) extension and some irreducible $f \in K[x]$, how many roots does f have in L?*

As Michael pointed out, we know by general field theory (see problem **5.4**) that f has $\leq \deg f$ roots in L. Of course the ideal circumstance would be to have *precisely* $\deg f$ roots in L, but this doesn't always happen; for example, $x^2 + 1$ has no roots in $\mathbb{Q}$. This is sort of a trivial example, the issue being that we haven't extended the field at all. But even an honest extension might not extend far enough to contain all the roots of a polynomial. For example, Ben pointed out that $\mathbb{Q}(\omega\sqrt[3]{2})$ contains only one root of $x^3 - 2 \in \mathbb{Q}[x]$.

Empirically, this final scenario seems abnormal – the more common situation is either for our extension to not contain any of the roots (e.g. $\mathbb{Q}(\sqrt{2})$ doesn't contain any roots of $x^2 + 1$) or for it to contain all of them. This motivates a definition:

Definition (Normal Extension)**.** L/K is called a *normal extension* if and only if every irreducible polynomial $f \in K[x]$ has either no roots or all of its roots in L.

Date: April 30, 2018.

Based on notes by Chetan Patel.

Thus, if f has any roots in a normal extension L/K, all of its roots must live there. This doesn't guarantee that we have $\deg f$ roots in L, however, since f might not be separable! So, the ideal situation is one in which an extension is finite, normal, and separable; this would guarantee that any polynomial which splits in the extension has the right number of roots. This ideal situation turns out to be equivalent to Galois-ity:

Proposition 2.1. *Given a finite extension L/K. The extension L/K is Galois if and only if it is normal and separable.*

Remark. Recall that a separable extension L/K is one in which the minimal polynomial for any element of L over K is separable.

Proof. The Fundamental Lemma is the key to proving the forward direction. The reverse direction is a synthesis of familiar techniques.

($\implies$) Given a finite Galois extension L/K, set $G := \text{Gal}(L/K)$.

We first show that L/K is normal. Suppose we have some irreducible polynomial $f \in K[x]$ that has a root $\alpha \in L$; our task is to show that all other roots of f must also live in L. Following a suggestion of Beatrix, we note that we may assume f is monic (else, we can rescale without changing the roots of f). Thus f is an irreducible monic polynomial with coefficients in K and α as a root – in other words, $f = m_\alpha$, the minimal polynomial of α over K. The Fundamental Lemma then implies that the roots of f are precisely the Galois conjugates of α, i.e. $\{\sigma(\alpha)\}_{\sigma \in G}$. Since G consists of automorphisms of L, we immediately deduce that all the roots of f live in L. We've proved that L/K is normal.

Next we show that L/K is separable, i.e. that for any $\alpha \in L$ the minimal polynomial $m_\alpha \in K[x]$ is separable. But, as Trevin pointed out, this is an immediate consequence of the Fundamental Lemma! Thus L/K is separable.

($\impliedby$) Given L/K finite, normal, and separable, our task is to show that it is Galois. Michael suggested a good place to start is with our finiteness assumption, meaning that we can write $L = K(\alpha_1, \alpha_2, \ldots, \alpha_n)$ using some minimal set of generators. (*Minimal* here means that none of the α_i can be removed.) Ben noted that our usual way of showing an extension is Galois is by building a separable polynomial in $K[x]$ that has L as its splitting field. Ian provided a method for such a process that involves the product of the minimal polynomials of the α_i.

Specifically, let $m_{\alpha_i} \in K[x]$ be the minimal polynomial of α_i over K. Because L/K is a normal extension and because we know one root of the irreducible m_{α_i} lives in L (namely, α_i), normality gives us that all roots of m_{α_i} must live in L. Furthermore, L/K is separable, so all the roots of m_{α_i} are distinct. Finally, Ben pointed out that two minimal polynomials share a root if and only if they are the same polynomial. Thus, the least common multiple of all the m_{α_i} is a separable polynomial with coefficients in K, all of whose roots live in L. Moreover, L is the splitting field of this polynomial, since it has all the α_i as roots. We conclude that L/K is Galois. □

Remark. In class we expressed the LCM of the m_{α_i} using a version of our notation trick: we set

$$m(x) := \prod_{f \in \mathcal{F}} f(x)$$

where $\mathcal{F} := \{m_{\alpha_i} : 1 \leq i \leq n\}$. Then $m \in K[x]$ is separable by the considerations in the proof above, and L is the splitting field of m over K.

The proposition demonstrates a close connection between normality and Galois-ity. In view of this connection, the following result might not be so surprising.

Lemma 2.2. *An extension F/K is finite and normal if and only if F is a splitting field of some $f \in K[x]$.*

Remark. This is very similar to one of our equivalent conditions for Galois-ity, but without the requirement that f be separable. For a proof of this lemma, see the supplementary document on the Isomorphism Lifting Lemma on the course website. (The ILL has other nice consequences too, e.g. uniqueness of splitting fields.)

3. Unique factorization of finite groups: Krull-Schmidt

Our goal for the next couple of lectures will be to explore Galois' celebrated criterion for solvability. The two key observations he made are the following:

(1) if $f \in \mathbb{Q}[x]$ has all its roots expressible in radicals, then there exists a tower of field extensions, starting at $\mathbb{Q}$ and ending at the splitting field of f over $\mathbb{Q}$, such that each extension in the tower is Galois and has abelian Galois group; and

(2) if an irreducible $f \in \mathbb{Q}[x]$ has a root expressible in radicals, then *all* roots of f are expressible in radicals.

Remark. Both of these observations hinge on the concept (introduced by Galois) of a *radical extension*. We will formalize this next class, but for now we noted that the second observation above is intuitively consistent with the Fundamental Lemma. Indeed, given a root α of an irreducible polynomial $f \in \mathbb{Q}[x]$, the Fundamental Lemma implies that all the other roots are the Galois conjugates of α. But from our experience, Galois conjugates change a radical by a 'sign'; a $\sqrt{3}$ might become a $-\sqrt{3}$, or a $\sqrt[4]{2}$ might become a $\sqrt[4]{-2}$. Thus, we expect all the other roots of f to also be expressible in radicals, just with different signs scattered throughout. (As we shall see, this isn't quite right – we need to allow roots of unity in addition to sign changes.)

Combining Galois' observations, we deduce that any irreducible polynomial with at least one root expressible in radicals has a splitting field which is connected to $\mathbb{Q}$ via a tower of Galois extensions which all have abelian Galois group. To make use of this to prove insolvability of the general quintic, we follow Galois and study the theory of finite groups. We begin with a motivating question.

Question 2. *Recall that any whole number can be decomposed into a product of primes in a unique way (up to ordering). For example,* $60 = 2 \times 2 \times 3 \times 5$*, and apart from the order in which we write these four numbers, there's no other way to decompose 60 into a product of primes. Is there an analogue for finite groups? Specifically, given a finite group* G *is there a way to decompose it into prime-like groups in some unique (up to something silly, like ordering) way?*

This question is not so obvious. Following Pólya's dictum, we search for an easier version of the problem that we *can* solve. Suppose, for example, that G were abelian. Michael pointed out that we then have a nice parallel to prime decomposition in the integers: the Fundamental Theorem of Finite Abelian Groups.

Fundamental Theorem of Finite Abelian Groups. Given any finite abelian G, there exists a unique set (up to labeling of its elements) of prime powers $\{p_i^{n_i} : 1 \leq i \leq k\}$ such that $G \simeq C_{p_1^{n_1}} \times C_{p_2^{n_2}} \times \cdots \times C_{p_k^{n_k}}$.

Thus, any finite abelian group can be decomposed into a product of prime-power cyclic groups in a unique way (up to the ordering). What about nonabelian groups? It turns out there's a lovely theorem in this setting:

Theorem 3.1 ("Krull-Schmidt Theorem" – proved by Wedderburn in 1909)**.** *Any finite group* G *can be decomposed as*

$$G \simeq P_1 \times P_2 \times \cdots \times P_n$$

where each P_i *is an indecomposable group (in the sense that it cannot be expressed as the product of two smaller groups). Furthermore, this decomposition is unique: any other such decomposition of* G *will yield the same list of indecomposables, up to isomorphism and ordering.*

This seems like the perfect analogy to unique factorization in $\mathbb{Z}$. Unfortunately, Krull-Schmidt turns out to not be a particularly useful theorem because we have no idea what indecomposable groups are like – we have no nice characterization of them apart from the definition. In other words, decomposing complicated groups into indecomposables doesn't give us additional insight into the structure of the original group!

4. A different approach to unique factorization: Jordan-Hölder

From above we see that the most natural analogue of prime factorization in the context of finite groups – Krull-Schmidt – isn't so helpful. Is there a more useful analogue?

For inspiration, we search for an alternative approach to prime factorization in $\mathbb{Z}$. One idea, inspired by our work with towers of field extensions, is to set up a tower of divisors. Here's how this works, using 60 as an example. We start with a trivial tower:

$$\begin{array}{c} 60 \\ | \\ 1 \end{array}$$

Next we 'refine' this tower by introducing an intermediate divisor, e.g.

$$\begin{array}{c} 60 \\ | \\ 15 \\ | \\ 1 \end{array}$$

Now we refine again by introducing another intermediate divisor, e.g.

$$\begin{array}{c} 60 \\ | \\ 30 \\ | \\ 15 \\ | \\ 1 \end{array}$$

Note that it's no longer possible to refine the top two extensions of our tower, since there are no intermediate divisors. We can, however, refine the bottom extension a bit more:

$$\begin{array}{c} 60 \\ {\scriptstyle 2}| \\ 30 \\ {\scriptstyle 2}| \\ 15 \\ {\scriptstyle 5}| \\ 3 \\ {\scriptstyle 3}| \\ 1 \end{array}$$

There are no more intermediate divisors one could insert, so we say this is a 'complete refinement' of the original trivial tower.

In order to save space, we'll write these divisor towers horizontally, e.g. our completely refined tower from above is

$$60 \overset{2}{\text{——}} 30 \overset{2}{\text{——}} 15 \overset{5}{\text{——}} 3 \overset{3}{\text{——}} 1$$

Note that this process of completely refining a tower of divisors is not unique. For example, we could have arrived at a different complete refinement if we'd inserted the intermediate divisor 20 at the first stage:

$$60 \text{ —— } 1 \overset{1^{\text{st}}\text{ refinement}}{\rightsquigarrow} 60 \text{ —— } 20 \text{ —— } 1 \overset{\text{continue...}}{\rightsquigarrow} 60 \overset{3}{\text{——}} 20 \overset{5}{\text{——}} 4 \overset{2}{\text{——}} 2 \overset{2}{\text{——}} 1$$

There are other complete refinements we might have ended up at as well. However, some aspects of the complete refinement are invariant, as Emily pointed out: the length of the completely refined chain is the same no matter how we refine the original. Moreover, Franny noted that the list of degrees (i.e. the adjacent quotients) we obtain is the same in any complete refinement. Indeed, this list is precisely the list of prime factors of 60 (with multiplicity).

Having developed this alternative approach to the prime factorization of integers, we explore how to carry it over to the setting of finite groups. The first thing we need is an analogue of divisor. What does it mean for one group to divide another? This immediately brings to mind the concept of normal subgroup. Indeed, we have the following analogy between whole numbers and groups:

Whole numbers	Groups
$d \leq n$	$H \leq G$
$d \mid n$	$H \trianglelefteq G$
$d \leq n \implies n/d$ is a number, but not necessarily a whole number	$H \leq G \implies G/H$ is a set, but not necessarily a subgroup
$d \mid n \implies n/d$ is a whole number	$H \trianglelefteq G \implies G/H$ is a subgroup

With this parallel in mind, we can now carry over the idea from integer factorization to the group setting. Given a finite group G, we have the trivial chain of divisors:

$$G \rhd \{e\}.$$

We can refine this chain by inserting intermediate normal subgroups. Eventually, we arrive at a complete refinement of the original series. In the context of groups, such a maximal refinement has a name:

Definition. A *composition series* of a finite group G is a chain of the form

$$G = G_0 \rhd G_1 \rhd G_2 \rhd \cdots \rhd G_n = \{e\}$$

such that no refinements of this chain are possible. The *composition factors* of G are the quotients G_i/G_{i+1}.

As in the case of whole numbers, a group might have multiple composition series. Remarkably, though, the list of composition factors is uniquely determined (up to isomorphism and ordering). This is the content of the following result:

Theorem 4.1 (Jordan-Hölder)**.** *Given two composition series of a finite group G:*

$$G = G_0 \rhd G_1 \rhd \cdots \rhd G_n = \{e\}$$
$$G = H_0 \rhd H_1 \rhd \cdots \rhd H_k = \{e\}.$$

Then $n = k$ and there exists some permutation $\pi \in S_n$ such that

$$G_i/G_{i+1} \simeq H_{\pi(i)}/H_{\pi(i)+1}$$

for all i.

Intuitively, we think of the composition factors as the prime factors of G. What are these groups, though? Are they simple to understand? We'll explore this next time.