MATH 121 NOTES: A FIELD GUIDE TO GALOIS THEORY

ARUN DEBRAY
AUGUST 17, 2013

Contents

1. Vector Spaces: 1/7/13 1
2. Review of Math 120 I: Planting Seeds: 1/9/13 3
3. Review of Math 120 II: Harvesting Consequences: 1/11/13 4
4. Extending Fields by Adjoining Roots of Irreducible Polynomials: 1/14/13 5
5. Splitting Fields and Algebraic Closures: 1/16/13 6
6. Searching for Closure: 1/18/13 7
7. Separable Field Extensions: 1/23/13 9
8. Finite Fields and Roots of Unity: 1/25/13 10
9. Introduction to Galois Theory: 1/28/13 11
10. Galois Theory of Separable Field Extensions: 1/30/13 13
11. Small Degrees and Finite Fields: 2/1/13 14
12. The Theorem of the Primitive Element: 2/4/13 15
13. The Fundamental Theorem of Galois Theory: 2/6/13 16
14. Lattices of Subgroups and Subfields: 2/8/13 17
15. Nested Radicals and Norms and Traces: 2/11/13 18
16. More Norms and Traces: 2/13/13 19
17. Solvable Groups: 2/15/13 20
18. Solvability By Radicals: 2/20/13 21
19. Two Applications of the Fundamental Theorem and Lagrange Resolvents: 2/22/13 22
20. Additive Hilbert's Theorem 90 and the Cubic: 2/25/13 24
21. Solving the Cubic: 2/27/13 25
22. The Fundamental Theorem of Symmetric Polynomials: 3/1/13 26
23. The Fundamental Theorem of Algebra: 3/4/13 27
24. Computation of Galois Groups I: 3/6/13 28
25. Computation of Galois Groups II: 3/8/13 29
Appendix A. Professor Quotes 30

1. Vector Spaces: 1/7/13

If one has the polynomial $x^2 + ax + b = 0$, the two solutions are given by $x = (-a + \sqrt{a^2 - 4b})/2$, where the square root can take on two values.

For cubics $x^3 + px + q = 0$, one similarly has $x = (A + B)/3$, where

$$A = \sqrt[3]{\frac{-27q + 3\sqrt{D}}{2}}, \ B = \sqrt[3]{\frac{-27q - 3\sqrt{D}}{2}}, \text{ and } D = -4p^3 - 27q^3.$$

Since each cube root can take on three values, this looks like it has too many degrees of freedom, but it is necessary for $AB = -3p$, so one forces the other, giving the correct number of solutions.

There exist (even more painful) formulas for the quartic, but this is not true for degrees 5 and higher; there is no algebraic formula for the roots. The proof of this will be one of the main results of this class.

Note that in this class rings will be assumed to have a multiplicative identity $1 \neq 0$.

Definition. A vector space over a field F is an abelian group $(V, +)$ with an operation of scalar multiplication $\cdot : F \times V \to V$ such that $c_1(c_2\mathbf{v}_1) = (c_1c_2)\mathbf{v}_1$ and $c(\mathbf{v}_1 + \mathbf{v}_2) = c_1\mathbf{v}_1 + c_1\mathbf{v}_2$ for all $c_1, c_2 \in F$ and $\mathbf{v}_1, \mathbf{v}_2 \in V$.

Homomorphisms can be defined for vector spaces, and for this class, since rings must have identity, then ring homomorphisms must preserve the identity (i.e. $1_A \to 1_B$). This doesn't follow from the other axioms of a ring homomorphism. Thus, for example, if A_1 and A_2 are rings, then $A_1 \to A_1 \times A_2$ given by $a_1 \mapsto (a_1, 0)$ is not a ring homomorphism, so $A_1 \times 0$ isn't a subring (though $a_1 \mapsto (a_1, 1)$ is a homomorphism, so $A_1 \times 1$ is a subring of $A_1 \times A_2$).

If F is a field and V and W are F-vector spaces, then $V \xrightarrow{T} W$ is a homomorphism if $T(\mathbf{v} + \mathbf{v}') = T(\mathbf{v}) + T(\mathbf{v}')$ and $T(c\mathbf{v}) = cT(\mathbf{v})$ for all $\mathbf{v}, \mathbf{v}' \in V$ and $c \in F$. Thus, homomorphisms are just linear transformations!

Definition. If V is a vector space over a field F, then $B \subset V$:

- spans V if every $\mathbf{v} \in V$ can be written as $\mathbf{v} = \sum_{j=1}^{k} c_j \mathbf{v}_j$ for $\mathbf{v}_j \in B$ and $c_j \in F$;
- linearly independent if $\sum_{j=1}^{k} c_j \mathbf{v}_j = 0$ implies that $c_j = 0$ for all j; and
- is a basis of V if it both spans V and is linearly independent.

A key fact about vector spaces is that any two bases of a vector space V have the same cardinality (and thus the same number of elements if the bases are finite). This number is called the dimension of V.

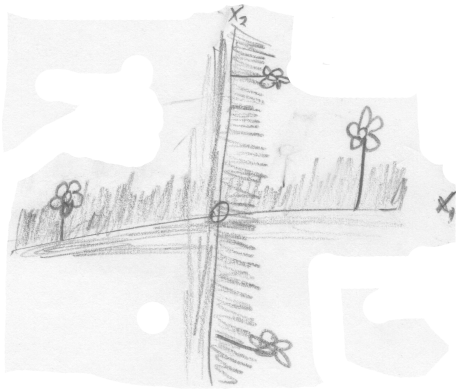

FIGURE 1. A 2-dimensional vector space over a field.

Definition. If E and F are fields, then E is a field extension of F is F is a subfield of E. Equivalently, $F \subset E$ and the inclusion map is a field homomorphism.[1]

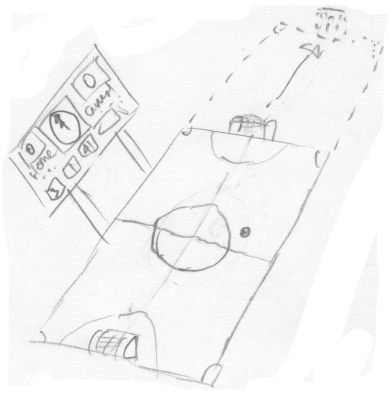

FIGURE 2. A field extension.

Suppose $F \subset E \subset K$ are fields, E is a field extension of F, and K is a field extension of E (and therefore F). Then, E and K can be made into F-vector spaces by defining scalar multiplication by an element of F as identical to field multiplication by that same element. Similarly, K is an E-vector space.

Definition. If $E \supset F$ is a field extension of F, then its degree is $|E : F|$, the dimension of E over F as an F-vector space.

[1]The axioms for a field homomorphism are the same as for a ring homomorphism.

Proposition 1.1. *If $\{\alpha_i\}_{i\in I}$ is an F-basis of E and $\{\beta_j\}_{j\in J}$ is an E-basis of K, then $\{\alpha_i\beta_j\}_{(i,j)\in I\times J}$ (where the product is taken in K) is an F-basis of K.*

Proof. Take a $\gamma \in K$ and write it as $\gamma = \sum_{j\in J} c_j\beta_j$, with $c_j \in E$ and $c_j = \sum_{i\in I} a_{ij}\alpha_i$, with $a_{ij} \in F$. Then,

$$\gamma = \sum_{j\in J}\sum_{i\in I} a_{ij}\alpha_i\beta_j,$$

so $\{\alpha_i\beta_j\}_{(i,j)\in I\times J}$ spans K over F.

For linear independence, suppose

$$\sum_{i\in I, j\in J} x_{ij}\alpha_i\beta_j = 0 = \sum_{j\in J}\left(\sum_{i\in I} x_{ij}\alpha_i\right)\beta_j.$$

Then, $\sum_{i\in I} x_{ij}\alpha_i = 0$, so all of the $x_{ij} = 0$. ⊠

Corollary 1.2. $|K : F| = |K : E||E : F|$.

2. Review of Math 120 I: Planting Seeds: 1/9/13

Most of this lecture is devoted to a review of rings, principal ideal domains, and fields, but there was a small amount of new information.

Definition. A succession of field extensions $F_1 \subset \cdots \subset F_n$ is usually called a tower.

Note that in Proposition 1.1, the results still hold if R is a field extension of F and K is some E-vector space (not necessarily a field), by the same proof. In particular, it is still true that $|K : F| = |E : F||K : E|$.

Some review of ring theory: common examples of rings include $\mathbb{Z}$, $\mathbb{Z}/n\mathbb{Z}$, and $A[x]$, where A is any commutative ring. An example of a noncommutative ring is $M_n(A)$, where A is any ring. This is the set of $n \times n$ matrices with elements in A.

This last example deserves some scrutiny: it is the ring of polynomials in x with coefficients in A, but what exactly is x? Sometimes this is viewed as a ring of polynomial functions, but the function $x^2 + 2x$ on $\mathbb{F}_3$ is identically 0 (check on $0, 1$, and 2), yet the polynomial $x^2 + 2x \in \mathbb{F}_3[x]$ is distinct from 0. Thus, it is better to think of polynomials as tuples $(a_0, a_1, \dots)$ of which at most finitely many terms are nonzero, with addition and multiplicationd defined in the conventional way. Thus, polynomials aren't functions, but they determine functions. One can get away with thinking of them as functions in $\mathbb{Q}$ or $\mathbb{R}$, however.

If $\varphi : A \to B$ is a ring homomorphism, then $\mathrm{Ker}(\varphi) = \{a \in A \mid \varphi(a) = 0\}$ is a two-sided ideal of A. Given a two-sided ideal I of A, the ring of cosets (or quotient ring) A/I can be formed. Of course, if A is a commutative ring, then any left or right ideal of A is a two-sided ideal.

Some common examples of fields include $\mathbb{Q}$, $\mathbb{R}$, $\mathbb{C}$, $\mathbb{F}_p$ for a prime p, and $K(x)$, where K is a field, the rational functions over K. As with the ring of polynomials, the field of rational functions is best thought of as equivalence classes of tuples: a rational function is a tuple $f = (p, q)$, where $p, q \in K[x]$ and $q \neq 0$, and two tuples (p, q) and (p', q') are equivalent if there is a $r \in K[x]$ such that $pr = p'$ and $qr = q'$ (or vice versa). This set of equivalence classes forms the field $K(x)$, and in many cases this is the set of rational functions $f(x) = p(x)/q(x)$.

Definition. A skew field or division ring is a noncommutative ring in which every nonzero element has an inverse.

The most common example of a skew field is the quaternions $\mathbb{H} = \{a + bi + cj + dk\}$, where $i^2 = j^2 = k^2 = -1$ and $ij = -ji = k$, etc.

Definition. A ring A is an (integral) domain if whenever $ab = 0$ for $a, b \in A$, then $a = 0$ or $b = 0$.

For example, $\mathbb{Z}/12\mathbb{Z}$ is not a domain.

In a PID, every ideal is principal (i.e. generated by a single element: $I = (a) = \{ra \mid r \in A\}$). Note that $\mathbb{Z}[x]$ and $\mathbb{Q}[x, y]$ aren't PIDs: consider the ideals $(2, x)$ and (x, y), respectively.

If K is a field, then $K[x]$ is a PID (sometimes called a principal ring), like $\mathbb{Z}$, which lends nice properties to the ideals. This is true because of the Euclidean division algorithm, which allows computation of division with remainder. For example, if $I \subset \mathbb{Z}$ (or $K[x]$ with K a field) is an ideal with $I \neq (0)$, then it is possible to pick an $m \in I \setminus 0$ such that $|m|$ is minimized. Then, $I = (m)$, using the Euclidean algorithm.

3. Review of Math 120 II: Harvesting Consequences: 1/11/13

Suppose R is a ring. Then, $R^* = \{r \in R \mid \exists r' \in R \text{ such that } rr' = 1\}$ (note that this is equivalent with $r'r = 1$). The elements of R^* are called (multiplicatively) invertible elements. R^* is a group under multiplication in R.

For any ring R there is a unique homomorphism $\mathbb{Z} \xrightarrow{\varphi} R : \varphi(1) = 1_R$, so $\varphi(n)$ can be determined for any $n \in \mathbb{Z}$. Then, $\text{Ker}(\varphi)$ is an ideal of $\mathbb{Z}$, so $\text{Ker}(\varphi) = (0)$ or $\text{Ker}(\varphi) = (n)$ for some $n \in \mathbb{N}$.

Claim. If R is an integral domain, then $n = 0$ or n is prime.

Proof. If $n = km = 0$ in R, then either $k = 0$ or $m = 0$ in R (since R is an integral domain), so $k \in (n)$ or $m \in (n)$. Thus, by definition, n is prime. ⊠

The characteristic of a field or an integral domain is

$$\text{Char}(F) = \begin{cases} 0, & \text{Ker}(\varphi) = (0) \\ p, & \text{Ker}(\varphi) = (p). \end{cases}$$

Corollary 3.1. *If F is a finite field, then $\text{Char}(F) = p$, so $|F| = p^m$ for some $m \in \mathbb{N}$.*

Proof. If $\text{Char}(F) = 0$, then $\mathbb{Z} \xrightarrow{\varphi} F$ is an injection, which is impossible if F is finite. Thus, $\varphi(\mathbb{Z}) = \mathbb{F}_p \subset F$, so F is a $\mathbb{F}_p$-vector space, so $F = \prod_{i=1}^{m} \mathbb{F}_p$ for some $m \in \mathbb{N}$ (since it is generated by a basis). ⊠

If p is prime, then $\mathbb{Z}/p\mathbb{Z} = \mathbb{F}_p$ is a field: consider some $k \in \mathbb{F}_p$ and the group homomorphism $\mathbb{F}_p \to \mathbb{F}_p$ given by $a \mapsto ka$. This is injective, because if $km \equiv 0 \bmod p$, then $p \mid km$, so $p \mid m$ (since $0 < k < p$ and p is prime). Thus, $m = 0 \bmod p$. Hence, this map is also surjective, since it is an injection from a finite set into itself.

Since it is a bijection, there exists a $k^{-1} \in \mathbb{F}_p$ such that $kk^{-1} = 1$. Thus, $\mathbb{F}_p$ is a field.

If $I, J \subseteq A$ are ideals of the commutative ring A, then one can define their sum and product:

$$I + J = \{a + b \mid a \in I, b \in J\}$$

$$IJ = \left\{ \sum_{i=1}^{n} r_i a_i b_i \;\middle|\; \begin{array}{l} a_i \in I, b_i \in J \\ r_i \in A, n \in \mathbb{N} \end{array} \right\}.$$

In the latter definition, the finite sums are necessary to ensure that the resulting product is closed under addition; otherwise, the product of two ideal elements might not be reachable via addition. Then, these are both ideals of A, along with $I \cap J$.

Definition. An ideal $P \subsetneq A$ is prime if $ab \in P$ implies one of $a \in P$ or $b \in P$. (Equivalently, $ab \equiv 0 \bmod p$ implies $a \equiv 0$ or $b \equiv 0 \bmod p$.)

Another characterization of prime ideals is that $P \subset A$ iff P/A is an integral domain.[2]

Definition. An ideal $Q \subset A$ is maximal if $Q \subsetneq A$ and if $Q \subseteq Q'$ and $Q' \subseteq A$ is an ideal, then $Q = Q'$ or $Q' = A$.

There are two equivalent characterizations:

- Q is maximal if $Q \subsetneq A$ and if $a \notin Q$, then $Q + (a) = A$.
- $a \notin Q$ implies there exists an r such that $ar \equiv 1 \bmod Q$, so A/Q is a field.

For example, $(x, y) \subset \mathbb{Q}[x, y]$ is maximal, so $\mathbb{Q}[x, y]/(x, y)$ is a field (in fact, it's isomorphic to $\mathbb{Q}$). However, $\mathbb{Q}[x, y]/(y)$ is an integral domain that isn't a field, since (y) is prime but not maximal (notice $(y) \subset (x, y)$).

In a PID, nonzero prime ideals are the same thing as maximal ideals and as ideals of the form (p) where p is an irreducible element (i.e. p is nonzero and noninvertible, and if $p = rs$ then one of r and s is invertible). If $A = F[x]$ is a field, then A is a PID, and if $E \supset F$ is a field extension, then there is a unique ring homomorphism $F[x] \xrightarrow{\varphi} E$ for a given $\alpha \in E$ such that

$$\varphi(a) = \begin{cases} a, & a \in F \\ \alpha, & a \notin F. \end{cases}$$

For example, $\varphi(2x^2 + 3x^3) = 2\alpha^2 + 3\alpha^3$. Then, either $\text{Ker}(\varphi) = \{0\}$ or $\text{Ker}(\varphi) = (g)$ for a unique irreducible, monic $g \in F[x]$. In the first case, α is called trascendental over F (as, for example π over $\mathbb{Q}$), and $\text{Im}(\varphi) \cong F[x]$, so E is infinite-dimensional over F. Otherwise, α is called algebraic, and $|E : F|$ is finite.

[2] This depends on a convention in which $\{0\}$ is not an integral domain.

4. Extending Fields by Adjoining Roots of Irreducible Polynomials: 1/14/13

First it will be helpful to recall some reminders from previous lectures and Math 120:

(1) Suppose E is a field and $f \in E[x]$. If $\alpha \in E$ is a root, then $f(x) = (x-\alpha)g(x)$ for some $g \in E[x]$ with $\deg(g) = \deg(f) - 1$. This is because for any $\alpha \in E$ (not necessarily a root), the Euclidean algorithm gives $f(x) = (x-\alpha)g(x) + r$, with r constant.

(2) A monic polynomial in $\mathbb{Z}[x]$ has a rational root iff it has an integral root, and any integral root divides the constant coefficient: if $f(x) = \sum_{i=1}^{m} a_i x^i \in \mathbb{Z}[x]$ and $0 = f(r/s)$, then multiply by s^m to show that $s = \pm 1$ using unique factorization and the fact that r/s must be in lowest terms.

This trick can be quite useful: consider $f(x) = x^3 - x + 3 \in \mathbb{Z}[x]$. Thanks to this result, the only roots that have to be checked are ± 1 and ± 3. Since these aren't roots in $\mathbb{Z}[x]$, then f is irreducible on $\mathbb{Z}[x]$ and even on $\mathbb{Q}[x]$, because if f had a rational root, then it would have a linear root, and because since $\deg(f) = 3$, then it must have a root iff it is reducible.

(3) A monic polynomial $f \in \mathbb{Z}[x]$ factors in $\mathbb{Z}[x]$ iff it factors in $\mathbb{Q}[x]$.

(4) $\mathbb{F}_p = \mathbb{Z}/p\mathbb{Z}$ and $F[x]/(g)$ (where $p \in \mathbb{Z}$ is prime, F is a field and $g \in F[x]$ is irreducible) are both fields, because $p\mathbb{Z}$ and (g) are maximal ideals in the Euclidean domains $\mathbb{Z}$ and $F[x]$, respectively.

Actually computng inverses in these fields seems unpleasant, but using the Euclidean algorithm makes things easier by efficiently computing gcds: if A is a PID and $a, b \in A$, then $(a, b) = (d)$ for some $d \in A$, and one writes $d \mid a$ and $d \mid b$. Then, if $c \mid a$ and $c \mid b$, then $c \mid d$, because $d = ra + sb = ruc + svc = (ru + sv)c$ for some $r, s, u, v \in A$.

If p is irreducible in $\mathbb{Z}$ and $0 < a < p$, then $(a, p) = (1) = \mathbb{Z}$, so (p) is maximal and $a \notin (p)$. Then, if $ar + ps = 1$ in $\mathbb{Z}$, then $\bar{r} = (\bar{a})^{-1} \in \mathbb{Z}/p\mathbb{Z}$ (where bars denote conjugacy classes). This also works in a polynomial ring: if g is irreducible of degree m in $F[x]$ (where F is a field), then for some $a(x) \neq 0$ with $\deg(a) < \deg(m)$, $(g, a) = (1) = F[x]$. Then, the Euclidean algorithm gives $sg + ra = 1$, yielding the inverse precisely as before.

For example, $g(x) = x^3 - 2x - 2 \in \mathbb{Q}[x]$ is irreducible, so $\mathbb{Q}[x]/(g) = \{a + bx + cx^2 \mid a, b, c \in \mathbb{Q}\} \simeq \mathbb{Q}(\theta)$, where θ is a root (any root) of g (and in fact ll that is known about θ is that $\theta^3 = 2\theta + 2$). Then, $(a + bx + cx^2)^{-1} = p + qx + rx^2$ for some $p, q, r \in \mathbb{Q}$, and

$$p(a + bx + cx^2) + q(ax + bx^2 + cx^2) + r(ax^2 + bx^3 + cx^4) = 1,$$

but $x^3 = 2x + 2$, allowing the equation to be simplified and solved as an equation of 3 unknowns. This is ugly but solvable, and can be easier than using the Euclidean algorithm.[3]

Suppose $F \subset F[x]/(g) \simeq F[\alpha]$. If $F \subset E$ as fields with $\alpha \in E$, then $g(\alpha) = 0$ (in some sense, there is a field in which g has a root).

Consider $\mathbb{F}_7$ and $g(x) = x^3 - 2 \in \mathbb{F}_7[x]$. g is irreducible over $\mathbb{F}_7$ (which can be checked by a brute-force check: the only cubes in $\mathbb{F}_7$ are $0, \pm 1$), but a larger field $\mathbb{F}_7[x]/(g)$ contains a root. This field is a 3-dimensional vector space over $\mathbb{F}_7$, so it has size $7^3 = 343$ and is $\mathbb{F}_{7^3} = \mathbb{F}_{343}$.

There is an alternate view involving matrices: if $A = \begin{pmatrix} 0 & 0 & 2 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix}$, then $A^3 = 2I \in M_3(\mathbb{F}_7)$. Every (finite) matrix over a field has a minimal polynomial; for A, this is g. So, philosophically, the root of some irreducible polynomial in $\mathbb{F}_7$ doesn't make that much sense (what *is* it, exactly?), but as matrices they do seem a bit more concrete. Then again, this isn't completely satisfactory either, even for Dr. Brumfiel, so whatever floats your boat.

If $F \subset F[x] = E$, $g \in F[x]$ is irreducible in F, and $g(\alpha) = 0$, then in $E[x]$, $g(X) = (x-\alpha)h(x)$, with $\deg(h) = \deg(g) - 1$. Then, h may be irreducible in E, which calls for adjoining another root β to F, such that $h(\beta) = 0$ in $F[\alpha, \beta]$.[4]

This is a problem, however, if h isn't irreducible. In the general theory, any $g \in F[x]$ (not necessarily irreducible) can be factored as

$$g(x) = \prod_{i=1}^{k} g_i(x)$$

for some irreducibles g_i with roots $\alpha_1, \ldots, \alpha_k$. Thus, $g_1(x) = (x - \alpha_1)h_1(x)$ in $F[\alpha_1][x]$, and then repeat with $h_1(x)$, etc.

This procedure is repeated at most k times, leading to:

---

[3] There are two alternative approaches: the time-old "just stare at it for a minute" algorithm, which of course has a complexity of $\Theta(1)$, and the brute-force solution of trying every rational number, which is at least guaranteed to terminate, since $\mathbb{Q}$ is countable.

[4] Distinguishing between $F[\alpha]$ and $F(\alpha)$ (i.e. adjoining as a ring or a field) only matters for transcendental elements: $\mathbb{Q}[\sqrt{2}] = \mathbb{Q}(\sqrt{2})$, but $\mathbb{Q}[\pi] \neq \mathbb{Q}(\pi)$.

Corollary 4.1. *Given an arbitrary monic* $g \in F[x]$ *such that* $\deg(g) = m$*, then there exist field extensions* F, $F[\alpha_1]$, $F[\alpha_1, \alpha_2]$, $\ldots$, $F[\alpha_1, \ldots, \alpha_k] = K$ *such that*

$$g(x) = \prod_{i=1}^{m} (x - \alpha_i)$$

over K.

These α_i may be nondistinct (i.e. multiple roots). Also, K is generated over F by these roots as a vector space.

5. Splitting Fields and Algebraic Closures: 1/16/13

Field theory has applications in the classical problem of straightedge-and-compass constructions. Coordinates (a, b) can be viewed as $a + bi \in \mathbb{C}$, and constructability can be given by starting with a constructible field F and then taking successive degree-2 extensions of it. Thus, any constructible field satisfies

$$\mathbb{Q} \overset{2}{\text{———}} F_1 \overset{2}{\text{———}} F_2 \overset{2}{\text{———}} \cdots \overset{2}{\text{———}} E,$$

so $|E : \mathbb{Q}| = 2^m$.

Thus, for example, $\sqrt[3]{2} \notin E$ for any such E (since $3 \nmid 2^m$), so it isn't possible to double the cube. Similarly, trisecting angles is in general impossible, because many sines and cosines aren't degree-2 extensions. Squaring the circle is also impossible, because π isn't even algebraic over $\mathbb{Q}$! It is also possible to consider constructing unit n-gons, but this will be postponed for now.

Definition. Let F be a field and $f \in F[x]$ be monic with $n = \deg(f) \geqslant 1$.[5] Then, a field extension $E \supset F$ is a splitting field for f if:

(1)
$$f(x) = \prod_{i=1}^{n} (x - \alpha_i), \quad \alpha_1, \ldots, \alpha_n \in E,$$
so that f factors linearly in $E[x]$, and

(2) $E = F[\alpha_1, \ldots, \alpha_n]$.

This definition doesn't require f to be irreducible in F, but that is the main case of the definition.

Theorem 5.1. *Splitting fields exist and are unique; specifically, with* $f \in F[x]$ *as given in the definition, there exists a splitting field* E *of* F*, and if* E *and* E' *are two splitting fields of* $f \in F[x]$*, then there exists an isomorphism* $\varphi : E \to E'$ *such that* $\varphi|_F = \mathrm{Id}$ *(i.e.* $\varphi(a) = a$ *for all* $a \in F$*).*

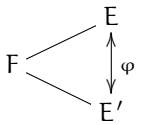

"If some alien comes down and shows us a splitting field, and we say 'we have one too,' and he[6] says 'no, it's different,' well, not really."

Proof of existence for Theorem 5.1. Construct $F \text{———} F[\alpha_1] = F_1$, where α_1 is one of the irreducible factors of f in $F[x]$. In $F_1[x]$, $f(x) = (x - \alpha_1) f_1(x)$ for some $f_1 \in F_1[x]$ with $\deg(f_1) = n - 1$.

Thus, by induction, there exists a field extension $E \supset F_1$ that is a splitting field of f_1 over F_1. Thus, f also factors linearly over E. ⊠

Uniqueness is harder, and will be proven in the next lecture. The essential ingredient is to take $E = F[\alpha_1, \ldots, \alpha_n]$ and $E' = F[\beta_1, \ldots, \beta_n]$ and establish an isomorphism $\varphi : F[\alpha] \to F[\beta]$. From there, induction can be applied.

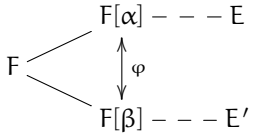

For example, if f is an irreducible cubic over $F[x]$, then $F \overset{3}{\text{———}} F[\alpha_1] \overset{1 \text{ or } 2}{\text{———}} E$; thus, a rational cubic with one real root has a splitting field of degree 6, and if it has 3 real roots, the splitting field might be either degree 3 or degree 6.

[5] This last requirement is important in a lot of definitions and results, but people tend to forget about it, so be careful.

[6] She? It?

Proposition 5.2. *If* E *splits* $f \in F[x]$ *(where* $\deg(f) = n$*), then* $|E : F| \mid n!$ *(in particular,* $|E : F| \leqslant n!$*), and if* f *is irreducible, then* $n \mid |E : F|$.

Proof. The second statement is easiest: in the first step, $F \overset{n}{\text{——}} F[\alpha_1] \overset{\leqslant (n-1)!}{\text{——}} E$, and $n = |F[\alpha_1] : F \mid |F : E|$. Then, the next extension splits a polynomial of degree $n-1$, so by induction, $|E : F| \leqslant n!$a

Suppose $f = f_1 f_2$ with $\deg(f_1) = k$, $\deg(f_2) = \ell$, and $k + \ell = n$, and suppose that f_1 is irreducible. Then, by induction, $F \overset{k}{\text{——}} F[\alpha_1] \overset{\leqslant (n-1)!}{\text{——}} E$. ⊠

Exercise 5.1. Finish the proof by showing that $|E : F| \mid n!$.

There are plenty of examples that illustrate that the upper bound of $n!$ is not tight: for $f = x^5 - 1$, let ζ be one of the complex roots of unity; then, f factors as

$$x^5 - 1 = (x-1)(x-\zeta)(x-\zeta^2)(x-\zeta^3)(x-\zeta^4),$$

but the splitting field only has degree 4, not 24.

The existence and uniqueness of splitting fields is an extremely powerful result that will be used frequently throughout the course.

Turning to algebraic closure, there are two related notions which share the name, so it will be helpful to be careful:

Definition. A field E is a algebraically closed if one of the following three equivalent conditions is true:

i. Every $g \in E[x]$ has a root in E (that is, if g is nonconstant).
ii. Every $g \in E[x]$ factors linarly in x: there exist $\alpha_1, \ldots, \alpha_n \in E$ such that

$$g(x) = \prod_{i=1}^{n} (x - \alpha_i).$$

iii. There are no proper extensions of E: if $E \text{——} E'$ is algebraic, then $E' = E$.

By induction, it is straightforward to show i $\implies$ ii, and the converse is trivial. Then, iii implies i because it forbids the existence of irreducibles in $E[x]$, and ii $\implies$ iii: if there are proper extensions $E \subset E[\beta] \subset E'$, such that the first extension is of degree $d > 1$, then β is the root of an irreducible of degree d, which isn't possible.

The other notion is as follows:

Definition. If F is some fixed field and $E \supset F$ is an extension, then E is an algebraic closure of F if every $f \in F[x]$ factors linearly in $E[x]$ and E is algebraic over F.

Later on in this class, it will be shown that $\mathbb{C}$ is algebraically closed (i.e. the Fundamental Theorem of Algebra). However, $\mathbb{C}$ is not algebraic over $\mathbb{Q}$, so it's not an algebraic closure of $\mathbb{Q}$.

6. Searching for Closure: 1/18/13

Theorem 6.1. *If* $E \supset F$ *is a field extension,* E *is algebraic over* F*, and every* $f \in F[x]$ *factors linearly in* E*, then* E *is algebraically closed.*

Proof. Here is a useful observation:

Claim. If $F \subset K \subset L$ is a tower, K/F is algebraic, and L/K is algebraic, then L/F is algebraic (i.e. every element of L is algebraic over F).

Proof. Let $\beta \in L$, so that $\beta^n + \alpha_1 \beta^{n-1} + \cdots + \alpha_n = 0$ for some $\alpha_1, \ldots, \alpha_n \in K$. Now consider

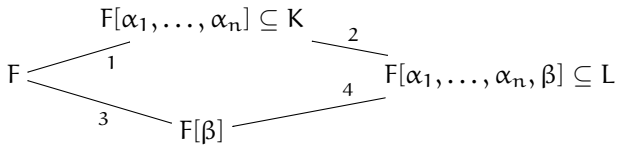

Clearly, 1, 2, and 4 are finite-dimensional extensions, so 3 must be as well. ⊠

This result was completely trivial for finite-dimensional vector spaces, but there exist infinite-dimensional extensions, such as $\mathbb{Q}(\sqrt{2}, \sqrt[3]{2}, \sqrt[5]{2}, \sqrt[7]{2}, \ldots)$ over $\mathbb{Q}$.

Now, suppose $F \text{——} E \text{——} E[\beta]$ is a tower of algebraic extensions for some β that is algebraic over E, so that the theorem is proved if $\beta \in E$. Then, β is algebraic over F, so $f(\beta) = 0$ for some $f \in F[x]$. In $E[x]$, f factors linearly as $f(x) = \prod(x - \alpha_i)$ with $\alpha_i \in E$, so in $E[\beta]$,

$$0 = f(\beta) = \prod(\beta - \alpha_i),$$

so $\beta = \alpha_j$ for some j. ⊠

Corollary 6.2. *Suppose $F \subset L$ is field extension and L is algebraically closed. Then, $\overline{F} = \{\gamma \in L \mid \gamma$ is algebraic over $F\}$ is an algebraic closure of F, so $\overline{F}$ is algebraically closed.*

Exercise 6.1. Prove this. The basic structure is to show that every monic $f \in F[x]$ factors linearly in $\overline{F}$, since it is already known to work in L, and thus all of the roots of f must lie in $\overline{F}$.

Example 6.1. The main example in this class is $\mathbb{Q}$,[7] so consider $\mathbb{Q} \subset \mathbb{C}$. Later on in the course, it will be shown that $\mathbb{C}$ is algebraically closed. As a corollary, there is an algebraically closed field

$$\overline{\mathbb{Q}} = \{z \in \mathbb{C} \mid f(z) = 0 \text{ for some } f \in \mathbb{Q}[x]\}$$

and $\overline{\mathbb{Q}}$ is the algebraic closure of $\mathbb{Q}$.

As another example, if α, β are algebraic over a field F, then $\alpha + \beta$, $\alpha\beta$, and $\gamma = \alpha^3 + 2\alpha\beta + \beta^7$ (for example) are all algebraic over F. Why is this?

Here is a "proof by magic:"

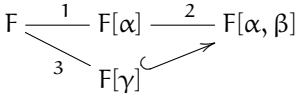

Here, 1 and 2 are finite-dimensional algebraic, so 3 must be as well. This is an example of how powerful Proposition 1.1 is.

There are other ways of constructing algebraic closures of a field F, but if F is uncountable, then some sort of fancy set theory (e.g. Zorn's Lemma, transfinite induction, or well-ordering) is necessary. Specifically, the book uses Zorn's lemma to prove that every commutative ring has a maximal ideal and therefore that every field has an algebraic closure.

If F is finite or countably infinite, then $F[x]$ is countable, so every (nonconstant) polynomial can be enumerated as $\{f_1, f_2, \dots\} = F[x]$, so the infinite tower

$$F \text{——} F_1 \text{——} F_2 \text{——} \cdots$$

such that F_{j+1} is the splitting field for $f_j \in F_j[x]$. Then, $\bigcup_{i=1}^{\infty} F_i$ is an algebraic closure of F:

- This is a well-defined field, since $F_{i+1} \supset F_i$, and if a and b are two elements of this field, there is an F_k such that $a, b \in F_k$, so their sum, product, and inverses all exist and behave nicely.
- Additionally, it is algebraically closed because any polynomial is in $F_k[x]$ for some k, so it factors linearly in $F_{k+1}[x]$.

Even if one had no conception of $\mathbb{C}$, this illustrates that an algebraic closure of $\mathbb{Q}$ at least exists. It might seem more obvious to take the union of every splitting field, but this creates set-theoretic issues: one might define an element $\mathtt{pig}$ such that $\mathtt{pig}^2 = 2$, and then how is one to distinguish $\mathbb{Q}(\sqrt{2})$ and $\mathbb{Q}(\mathtt{pig})$?

In order to prove the uniqueness of an algebraic closure of a field (up to isomorphism; see the previous lecture), the key step is to prove the uniqueness of the splitting field of a single polynomial up to isomorphism. Again the uncountable case requires some set-theoretic hijinks to make everything work, and the proof proceeds by induction on the degree of the polynomial $f \in F[x]$. Specifically, it requires a somewhat stronger statement:

Claim. If $\varphi : K_1 \xrightarrow{\sim} K_2$ is an isomorphism of fields, then it induces a ring isomorphism $\bar{\varphi} : K_1[x] \xrightarrow{\sim} K_2[x]$ given by

$$\sum_{i=1}^{n} a_i x^i \mapsto \sum_{i=1}^{n} \varphi(a_i) x^i$$

such that:

i. if $p \in K_1[x]$ is irreducible, then $\bar{\varphi}(p) \in K_2[x]$ is as well (since $\bar{\varphi}^{-1}$ is also an isomorphism).

[7] When Professor Galatius mentioned that irreducibles in $\mathbb{Q}[x]$ were "interesting" in Math 120, this was probably part of the reason.

ii. If $p \in K_1[x]$ is irreducible and α is a root of p in a field containing K_1 and α' is a root of $\bar{\varphi}(p)$ in a field containing K_2, then $\bar{\varphi}$ extends to an isomorphism $\psi : K_1[\alpha] \to K_2[\alpha']$ such that $\psi(\alpha) = \alpha'$:

$$\begin{array}{ccc} K_1 & \longrightarrow & K_1[\alpha] \\ {\scriptstyle\varphi}\downarrow\wr & & {\scriptstyle\psi}\downarrow\wr \\ K_2 & \longrightarrow & K_2[\alpha'], \end{array}$$

because $K_1[\alpha] \simeq K_1[x]/(p) \simeq K_2[x]/(\bar{\varphi}(p)) \simeq K_2[\alpha']$.[8]

Thus, if F_1 and F_2 are splitting fields of $f = pq \in F[x]$, with p irreducible on F, then let α be a root of p in F_1 and α' be a root of p in F_2. Then,

$$\begin{array}{ccccc} & & F[\alpha] & \longrightarrow & F_1 \\ F & \swarrow\!\!\searrow & {\scriptstyle\varphi}\downarrow\wr & & \\ & & F[\alpha'] & \longrightarrow & F_2 \end{array}$$

even though as sets, F_1 and F_2 might be completely unrelated, for some $\varphi \neq \mathrm{Id}$. Then, write

$$f(x) = (x - \alpha)g(x) \in F_1[x]$$
$$f(x) = (x - \alpha')h(x) \in F_2[x]$$

and do something similar, leading to induction. Since g splits in F_1 and h splits in F_2, then the same isomorphism is induced, and $h = \varphi(g)$ over $F[\alpha']$.

7. Separable Field Extensions: 1/23/13

Suppose $F \in K[x]$ is a monic polynomial with degree n. In a splitting field $F \subset E = F(\alpha_1, \ldots, \alpha_n)$ we have $f(x) = \prod_{i=1}^{n}(x - \alpha_i)$. One can ask whether these roots are distinct, which is not always the case (as with $x^4 + 2x^2 + 1 = (x^2 + 1)^2 \in \mathbb{Q}[x]$, which splits in $\mathbb{Q}(i)$ as $(x + i)^2(x - i)^2$).

Definition. If $\deg(f) = n$ and f has n distinct roots in a splitting field, then f is a separable polynomial.

Of course, this is most interesting when f is irreducible.

Definition. $\alpha \in E$ is a root of multiplicity $m \geqslant 1$ if $f(x) = (x - \alpha)^m h(x)$ in $E[x]$ with $h(\alpha) \neq 0$.

Some multiplicities have special names, such as double roots or triple roots.

From calculus, $f'(\alpha) = 0$ if α is a multiple root (over $\mathbb{R}$, anyways). This is in fact the answer: though limits make no sense in this algebraic setting, the derivative of a polynomial can be defined to be a map $()'$ or $\frac{d}{dx} : F[x] \to F[x]$ such that

$$f(x) = \sum_{j=0}^{n} a_j x^j \mapsto f'(x) = \sum_{j=1}^{n} j a_j x^{j-1}.$$

This can be defined as the F-linear map such that $\frac{d}{dx}(x^k) = kx^{k-1}$, which then forces the rest of the formula. Additionally, one obtains the product rule $(fg)' = f'g + fg'$.

If $F \subset E$ is any field extension, then $f \in F[x] \subset E[x]$, so the derivative is the same. But if f is irreducible over f, it might not be over E, so one might be able to use the Product Rule to learn more. Thus, if $f(x) = (x - \alpha)^m h(x)$, $m \geqslant 1$, and $h(\alpha) \neq 0$ in $E[x]$, then $f'(x) = m(x - \alpha)^{m-1}h(x) + (x - \alpha)^m h'(x)$ (which does involve a little more care with the definition of the derivative). Then:

- If $m = 1$, then $f'(\alpha) = mh(\alpha) = h(\alpha) \neq 0$.
- If $m \neq 1$, then $f'(\alpha) = 0$, as both terms go to zero.

Corollary 7.1. *f is separable if f and f' have no common roots in a splitting field of F.*

This can be reformulated as: f is separable iff $\gcd(f, f') = 1$ in $F[x]$. This is nicer because F is generally better understood than its splitting fields and because there may be an algorithm for computing the greatest common divisor[9] in $F[x]$.

From Math 120, if the greatest common divisor is 1, then $Af + Bf' = 1$ in $F[x]$ for some $A, B \in F[x]$. Thus, there can't be any multiple roots (if there were such a root α, then $A(\alpha)f(\alpha) + B(\alpha)f'(\alpha) = 1 = 0$). In the other direction, if the greatest common divisor is not 1 (i.e. it is nonconstant, since it's defined up to multiplication by an invertible

[8] This relies on a fact from Math 120: if $\varphi : R \xrightarrow{\sim} S$ is a ring isomorphism and $I \subset R$ is an ideal, then $R/I \xrightarrow{\sim} S/\varphi(I)$.

[9] ... is this not the case in any Euclidean domain?

factor, which is just a constant), then $\gcd(f, f') = r \neq 1$. Then, the splitting field for f contains the roots of $r(x)$, and these are roots of both $f(x)$ and $f'(x)$.[10]

If $f \in F[x]$ is irreducible of degree greater than 1, then it seems reasonable that $\gcd(f, f') = 1$. But it's in fact slightly trickier than that: if $\mathrm{Char}(F) = 0$, then this is in fact the case: if f is irreducible, then f and f' have no common factors in $F[x]$. Thus, if $\mathrm{Char}(F) = 0$, then all irreducible polynomials are seperable.

If $\mathrm{Char}(F) = p$, then $\frac{d}{dx}(x^p) = px^{p-1} = 0$, so $\gcd(f, 0) = f$ (since 0 is a multiple of anything). Thus, if f is irreducible and $f' \neq 0$, then f is seperable. $f' = 0$ exactly when $f(x)$ is a polynomial in x^p (i.e. $f(x) = h(x^p)$).

Example 7.1. Let $F = \mathbb{F}_p(T)$ for some indeterminate T and let $f(x) = x^p - T$. This is irreducible in $F[x]$ (which is non-obvious), but $f'(x) \equiv 0$.

What if $|F| = p^n$ is finite? Then, everything is fine:

Definition. If K is a field with $\mathrm{Char}(K) = p$, then the Frobenius homomorphism is $K \xrightarrow{\sigma} K$ given by $\sigma(a) = a^p$.

This is a field homomorphism: it's clear that $1^p = 1$ and $(ab)^p = a^p b^p$, but $(a+b)^p = a^p + b^p$ as well, which looks like some 9$^{\text{th}}$–grader's mistake. However, it is verified when expanding out with the Binomial Theorem: everything with a coefficient of p vanishes, leaving these two terms. This is sometimes known as the Freshman's Dream.

Consider a polynomial f such that $f'(x) \equiv 0$: $f(x) = \sum_{j=0}^{n} b_j x^{pj}$. Then, the Frobenius homomorphism shows that all of the $b_j = a_j^p \in F$, and f itself is equal to $\left(\sum_{j=0}^{n} a_j x^j\right)^p$. Thus, if F is a finite field, then every irreducible polynomial is seperable.

It's possible to go even further: if $\mathrm{Char}(F) = p$, then let $g(x) = x^{p^n} - x \in F[x]$, so that $g'(x) = -1$. This is seperable, even though it's very reducible. Thus, if $E \supset F$ is a splitting field for g, then consider the set $K = \{\alpha \mid \alpha^{p^n} = \alpha\}$, or the set of roots of g. There are p^n such roots, and because of the Frobenius homomorphism, K is closed under addition and multiplication. Thus, K is a field of size p^n (which is a nice surprise, since the roots of a polynomial don't generally form a field) and $g(x) = \prod_{i=1}^{p^n}(x - \alpha_i)$ for some distinct α_i. In particular, K is the splitting field of g! This is a key step in the proof that finite fields of order p^n exist for any prime p and $n \in \mathbb{N}$.

8. Finite Fields and Roots of Unity: 1/25/13

First, a summary of some finite field results:

I. Suppose F is a field with $|F| = p^n$, with p prime. Then, $F^* = F \setminus 0$ is a group under multiplication; $|F^*| = p^n - 1$. Then, $a^{p^n - 1} = 1$ for all $a \in F^*$, so the p^n elements of F are the roots of $x^{p^n} - x$. Thus,

$$x^{p^n} - x = \prod_{a \in F}(x - a),$$

so F is the splitting field of $x^{p^n} - x$ over $\mathbb{F}_p = \mathbb{Z}/p\mathbb{Z}$.

II. Let $E \supset \mathbb{F}_p$ be a splitting field of $x^{p^n} - x$. E is clearly finite. Then, consider the set of solutions of $f(x) = x^{p^n} - x = 0$. This forms a subfield (because they are the result of n applications of the Frobenius homomorphism). Then, $f'(x) = -1$, so f and f' are relatively prime, so f is seperable. Thus, these p^n roots are all distnict and already form the splitting field, so $|E| = p^n$.

Thus, every finite field is the splitting field of $x^{p^n} - x$ over $\mathbb{F}_p$ for various values of n and p. But because of the uniqueness of splitting fields, all fields of order p^n are isomorphic.

III. If K is any field and $G < K^*$ is a finite multiplicative subgroup, then G is cyclic.[11] Thus, if $|F| = p^n$ is finite, then $F^* \cong \mathbb{Z}/(p^n - 1)\mathbb{Z}$ (i.e. is cyclic) and write $F^* = \left\langle a \mid a^{p^n - 1} = 1 \right\rangle$. Then, considering the extension $\mathbb{F}_p(a) \supset \mathbb{F}_p$, the minimal polynomial of a has degree n (since this is the degree of the dimension of the vector space). Thus, $F \simeq \mathbb{F}_p[x]/(g)$ for any irreducible $g \in \mathbb{F}_p[x]$ of degree n (since all finite fields of the same size are isomorphic).

IV. If $m < n$, when is $\mathbb{F}_p \subset \mathbb{F}_{p^m} \subset \mathbb{F}_{p^n}$? Looking at vector space dimensions,

$$\mathbb{F}_p \overset{\overset{n}{\frown}}{\underset{m}{\frown}} \mathbb{F}_{p^m} \frown \mathbb{F}_{p^n},$$

so it is necessary that $m \mid n$. It turns out this is also sufficient: if $d \mid n$, then $\mathbb{F}_{p^d} \subset \mathbb{F}_{p^n}$. Here's an elementary proof: since $d \mid n$, then $x^{p^d} - x \mid x^{p^n} - x$ in $\mathbb{F}_p[x]$. Thus, in $\mathbb{F}_{p^n}$,

$$x^{p^n} - x = \prod_{a \in \mathbb{F}_{p^n}} (x - a) \text{ and } x^{p^d} - x = \prod_{p^b \text{ elements of } \mathbb{F}_{p^n}} (x - b),$$

[10] Another way to view this is that if α is a root of r, then f, f' are multiples of r in some field, so $f(\alpha) = f'(\alpha) = 0$.

[11] For a proof, see Math 120 or the handout.

so $x^{p^d} - x$ has all of its roots in $\mathbb{F}_{p^n}$.

There's a different proof that involves the Frobenius automorphism $\mathbb{F}_{p^n} \xrightarrow{\sigma} \mathbb{F}_{p^n}$. Taking the fixed elements of σ^k forms a field of order p^d.

Definition. A root of unity in a field K is an $a \in K$ such that $a^n = 1$ for some $n \in \mathbb{N}$.

If n is fixed, these are the roots of $x^n - 1$. For $K = \mathbb{C}$, these roots can be more easily be understood because they can be "seen," but in any K, the solutions of $x^n - 1 = 0$ form a cyclic group under multiplication.

It's also worth asking how $x^n - 1$ factors over the prime field $\mathbb{Q}$ or $\mathbb{F}_p$.[12] Here are some examples:

$$x^2 - 1 = (x-1)(x+1)$$
$$x^3 - 1 = (x-1)(x^2+x+1)$$
$$x^4 - 1 = (x^2-1)(x^2+1) = (x-1)(x+1)(x^2+1)$$

Sometimes weird things happen, though: on $\mathbb{F}_3$, the Frobenius isomorphism shows that $x^3 - 1 = (x-1)^3$, which is not seperable.

Any a such that $a^n = 1$ satisfies $a^d = 1$ for some minimal $d \mid n$ (which is a fact of elementary group theory). n isn't necessarily the smallest value (as in $1^7 = 1$). Then, d is the order of the group of solutions to $x^n - 1 = 0$.

Over $\mathbb{Q}$ (or any field with characteristic 0), $x^n - 1$ is seperable, so there are n roots in the splitting field. Thus, the roots can be grouped according to their order:

$$x^n - 1 = \prod_{\substack{1 \leqslant d \leqslant n \\ d \mid n}} \Phi_d(x) \in \mathbb{Z}[x], \text{ where } \Phi_d(x) = \prod_{\substack{a^d = 1 \\ d \text{ is the order of } a}} (x - a).$$

This seems a bit abstract (for example, it's not at all clear why they have integer coefficients), but here are some examples: $\Phi_1(x) = x - 1$, $\Phi_2(x) = x + 1$, $\Phi_3(x) = x^2 + x + 1$, $\Phi_4(x) = (x-i)(x+i) = x^2 + 1$, etc. In the general case, by induction on the division algorithm division by a monic polynomial in $\mathbb{Z}[x]$ stays within $\mathbb{Z}[x]$, which eventually implies that all of the $\Phi_n(x)$ have integer coefficients.

If $n = p$ is prime, then $d = 1$ or $d = n$, so

$$x^p - 1 = (x-1)\left(\sum_{j=0}^{p-1} x^j\right) \implies \Phi_p(x) = \sum_{j=0}^{p-1} x^j.$$

Using the Gauss Lemma, one can determine whether these $\Phi_d(x)$ are irreducible in $\mathbb{Z}[x]$ and $\mathbb{Q}[x]$: if $d = p$, then $\Phi_p(x) = \frac{x^p - 1}{x - 1}$ and

$$\Phi_p(x+1) = \frac{(x+1)^p - 1}{x + 1 - 1} = \frac{x^p + px^{p-1} + \cdots + px}{x} = x^{p-1} + px^{p-2} + \cdots + p.$$

This latter polynomial is irreducible by Eisenstein's criterion.

Example 8.1. Consider $\Phi_7(x) = x^6 + x^5 + x^4 + x^3 + x^2 + x + 1 \in \mathbb{F}_2[x]$. One can check that this is equal to $(x^3 + x^2 + 1)(x^3 + x + 1)$ and that both of these are irreducible. Nonetheless, $\Phi_7(x)$ is irreducible over $\mathbb{Q}$, illustrating some of the nuances involved here.

There are many proofs of most of these things (including that $\Phi_n(x)$ is always irreducible over $\mathbb{Z}[x]$). One of them uses the finite fields and unique factorization in $\mathbb{F}_p[x]$.

The degree of $\Phi_n(x)$ is the number of primitive roots of unity, $\varphi(n) = \{1 \leqslant j < n \mid \gcd(j, n) = 1\}$, or Euler's φ-function.

9. Introduction to Galois Theory: 1/28/13

Recall that

$$x^n - 1 = \prod_{\substack{1 \leqslant d \leqslant n \\ d \mid n}} \Phi_d(x) \in \mathbb{Z}[x],$$

where the roots of $\Phi_d(x)$ are those ζ such that $\zeta^d = 1$, but $\zeta^e \neq 1$ for $0 < e < d$ (i.e. the primitive d^{th} roots). $\deg(\Phi_n(x)) = \varphi(n) = |\{j \mid 1 \leqslant j \leqslant n, \gcd(j, n) = 1\}|$. Thus, $\zeta_n = e^{2\pi i/n}$ is a primitive n^{th} root of unity and ζ_n^j is a primitive n^{th} root of unity iff $\gcd(n, j) = 1$ (in general, it's a primitive $k = n/(j, n)^t$ root, of order k in the group).

[12] Every field contains a copy of either $\mathbb{Q}$ or $\mathbb{F}_p$.

The book gives a good proof that $\Phi_n(x)$ is irreducible in $\mathbb{Z}[x]$ (and therefore $\mathbb{Q}[x]$), but some things merit clarification, so here is a proof sketch: suppose $\Phi_n(x) = f(x)g(x)$, where f is the minimal polynomial of $\zeta = \zeta_n$. Then, one can show that if $p \nmid n$ for a prime p, then ζ^p is also a root of $f(x)$, so $f(x)$ has $\varphi(n)$ roots, and g is constant. This depends on a couple key facts:

(1) If ζ is any root of $f(x)$, then so is ζ^p when $(p, n) = 1$.
(2) If $f(\zeta^p) \neq 0$, then $g(\zeta^p) = 0$, since ζ^p is certainly a primitive root. Then, it is helpful to reduce everything $\bmod p$.
(3) $x^n - 1$ has np repeated roots $\bmod p$, since it is coprime to its derivative, so $\Phi_n(x)$ has no repeated roots either, since it is a factor of $x^n - 1$.

It's now pretty clear that constructing the regular n-gon is possible iff it's possible to construct ζ_n. If $n = q_1^{e_1} \dots q_s^{e_s}$ is the prime factorization, then

$$\varphi(n) = \prod_{i=1}^{s} q_i^{e_i - 1}(q_i - 1),$$

since if $(m, n) = 1$, then $\varphi(mn) = \varphi(m)\varphi(n)$. Then, ζ_n is constructible if $\varphi(n)$ is a power of 2, so all of the q_i must be Fermat primes: $q_j = 1 + 2^{e_j}$, so $n = 2^e q_1 \dots q_s$. However, there are only five known Fermat primes: 2, 3, 5, 17, and 65537. There might be more.

Moving into Galois theory proper, the basic idea is that if $K \supset F$ is a field extension, then there is a group $\mathrm{Gal}(K/F)$, the group of field automorphisms that fix F (i.e. $\mathrm{Gal}(K/F) = \{\sigma \in \mathrm{Aut}(K) \mid \sigma|_F = \mathrm{Id}\}$), with group operation of composition.

Example 9.1. Suppose $F \overset{2}{\text{——}} K = F(\alpha)$ (i.e. a degree-2 extension), so that K is the splitting field of some irreducible quadratic $x^2 + ax + b \in F[x]$. Assume this polynomial is separable and the other root is $\beta \neq \alpha$. Then, $\alpha + \beta = -a$, and there are two automorphisms: $\alpha \mapsto \alpha$ and $\alpha \mapsto \beta$. Thus, $\mathrm{Gal}(K/F)$ is cyclic of order 2. These automorphisms can be easily explicitly written using the quadratic formula unless $\mathrm{Char}(F) = 2$.

If $\mathrm{Char}(F) = 2$, then irreducible quadratics still exist (such as $x^2 + x + 1 \in \mathbb{F}_2[x]$). In general, $x^2 + ax + b$ has 2 roots if $a \neq 0$. Otherwise, $x^2 + b = (x + \beta)^2$ somewhere and the Galois group is trivial.

Example 9.2. Consider the cyclotomic extension $\mathbb{Q} \overset{\varphi(n)}{\text{——}} \mathbb{Q}(\zeta_n)$ and let $\zeta = \zeta_n$. This extension is the splitting field of $\Phi_n(x) = \prod_{(j,n)=1}(x - \zeta^j)$. Thus, for any automorphism $\zeta \overset{\sigma_j}{\mapsto} \zeta^j$ for some j such that $(j, n) = 1$, and all such automorphisms are legal. Thus, $|\mathrm{Aut}(\mathbb{Q}(\zeta_n)/\mathbb{Q})| = \varphi(n)$, but also $\sigma_\ell \circ \sigma_j : \zeta \mapsto \zeta^{j\ell}$, so $\sigma_\ell \circ \sigma_j = \sigma_{j\ell}$. Thus, $\mathrm{Aut}(\mathbb{Q}(\zeta_n)/\mathbb{Q}) \simeq (\mathbb{Z}/n)^*$ (i.e. the multiplicative group of invertible elements $\bmod n$).

For example, if $n = p$ is prime, then $(\mathbb{Z}/p)^* \simeq \mathbb{Z}/(p-1)$ (cyclic of order $p-1$) and the generators are called primitive roots: $(\mathbb{Z}/p)^* = \{1, a, a^2, \dots, a^{p-2}\}$ for some nice a. For odd primes, $(\mathbb{Z}/p^n)^*$ is cyclic of order $p^n - p^{n-1}$, but this fails for $p = 2$: $(\mathbb{Z}/8)^* = \{1, 3, 5, 7\} \simeq \mathbb{Z}/2 \times \mathbb{Z}/2$, the Klein-four group. In general, for $n \geqslant 3$, $(\mathbb{Z}/2^n)^* = \mathbb{Z}/2 \times \mathbb{Z}/2^{n-2}$.

Observe that $\mathrm{Aut}(\mathbb{Q}(\zeta_n)/\mathbb{Q})$ is abelian. This is pretty, but not always the case: suppose $f \in F[x]$ is an irreducible, seperable cubic. Then,

$$\begin{array}{ccc} F & \text{———} & K = F(\alpha_1, \alpha_2, \alpha_3) \\ & & \overset{k}{\nearrow} \\ F & \overset{3}{\text{——}} & F(\alpha_1) \end{array}$$

where k is either 1 or 2. If $|K : F| = 3$, then there are 3 automorphisms and he group is abelian, but if $|K : F| = 6$, then $\mathrm{Aut}(K/F) = S_3$, since they can be thought of as permutations of the roots and because the total number is 6:

$$\begin{array}{ccccc} & & F(\alpha_1) & \overset{2}{\text{——}} & K \\ & \overset{3}{\nearrow} & \sigma \downarrow \wr & & \sigma \downarrow \wr \\ F & & & & \\ & \underset{}{\searrow} & F(\alpha_j) & \text{——} & K \end{array}$$

Thus, each σ on $F(\alpha_1)$ has 2 extensions and therefore there are $(3)(2) = 6$ choices.

Here's another proof that $(fg)' = f'g + fg'$, due to Brian Conrad: by the Binomial Theorem, $(x+h)^p = x^n + h(nx^{n-1}) + h^2(\text{stuff})$, so

$$\begin{aligned} f(x+h)g(x+h) &= (f(x) + hf' + h^2(\text{stuff}))(g(x) + hg' + h^2(\text{stuff})) \\ &= f(x)g(x) + h(f'g + fg') + h^2(\text{stuff}), \end{aligned}$$

where the middle term also has to be $h(fg)'$. Pre-Cauchy, people thought of all derivatives like this (i.e. in terms of power series), until analytic functions were discovered.

If E/F is a finite splitting field extension (note that a normal extension means the splitting field of some family of polynomials, even in some cases the algebraic closure), one might wish to count $|\operatorname{Aut}(E/F)|$. This is fairly easy if every element of E is separable over F. If F is finite or of characteristic 0 (or actually, if $\operatorname{Char}(F) = p$, then this can be slightly generalized to that every element is a p^{th} power), then this always holds.

For a first look at Galois theory, one cna look just at these cases, so that everything is separable. Suppose $E \supset F$ is normal and $S \subset \Omega$ is some algebraic closure of E and F. Then, any homomorphism $E \xrightarrow{\varphi} \Omega$ such that $\varphi|_F = \mathrm{Id}$ must fix E: $\varphi(E) = E$, since if $a \in E$ is a root of an irreducible $g \in F[x]$, then all of the roots of g are in E (since the coefficients of g, which are in F, are fixed). If E is finite-dimensional, a vector-spatial argument offers an alternative proof.

Theorem 10.1. *More generally, if K/F is any finite extension (and therefore assumed to be separable), then the number of φ such that*

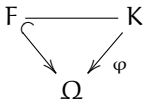

and $\varphi|_F = \mathrm{Id}$ is $|K : F|$. Thus, if K is also normal, then $|\operatorname{Aut}(K/F)| = |K : F|$.

Proof. Write $K = F(a_1, \ldots, a_n)$. Then,

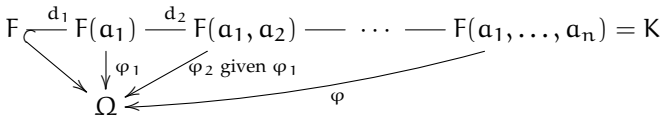

The Fundamental Principle of Counting states that if a process takes n steps and step i has d_i options, then the total number of options is $d_1 \cdots d_n$. Using this, the total number of possible φ_1 is d_1, since a_1 is mapped to any root of its minimum polynomial over F. Repeating for φ_2, $\varphi_2(a_2)$ is any root of $\varphi_1(g_2(x)) \in \varphi(F(a)[x])$, where g is the minimum polynomial of a_2. Continuing onward, the total number of φ is $d_1 \cdots d_n = |K : F|$. ⊠

Most of Galois theory can be traced back to that if $f \in F[x]$ is irreducible and a and b are roots of f, then $F(a) \cong F(b)$ through an isomorphism that leaves F fixed.

Remark. If K/F is separable (or even if K is generated by separable[13] $a_1, \ldots . a_n$), the same proof shows that the size of the automorphism group is $|K : F|$. If E/F is a finite normal seperable extension, then it is called a finite Galois extension, and one writes $\operatorname{Gal}(E/F) = \operatorname{Aut}(E/F)$ for the Galois group.

If $f \in F[x]$ is seperable, the Galois group of f is $\operatorname{Gal}(E/F)$, where E is the splitting field of f. For example, the Galois group of $\Phi_n(x) \in \mathbb{Q}[x]$ is isomorphic to $(\mathbb{Z}/n)^*$, as in Example 9.2.

Suppose $F = \mathbb{F}_2(T)$, where T is some indeterminate,[14] and suppose that T isn't a square in $\mathbb{F}_2[T]$. Thus, x_T^2 is irreducible over F, so add a root γ, so that $x^2 - T = (x-\gamma)^2 \in F[\gamma][x]$ by Eisenstein. Thus, the Galois group is trivial, eben though $F \xrightarrow{2} E$, and this can be generalized to $x^n - T$.

Returning to the seperable case, if $F \subset E = F(a_1, \ldots, a_n)$, where $a_1, \ldots, a_n$ are distinct roots of some n^{th}-degree $f \in F[x]$, then every $\varphi : E \xrightarrow{\sim} E$ must permute these a_i, so $\operatorname{Gal}(E/F) \leqslant S_n$, the symmetric group (up to isomorphism). Thus, $|\operatorname{Gal}(E/F)| \mid n!$

If f is reducible, then $\operatorname{Gal}(E/F) < S_n$ (some of the isomorphisms aren't possible; for example, in $f(x) = (x^2 + bx + c)(x^3 + dx^2 + gx + h)$, the roots of the quadratic can't be sent to the roots of the cubic). If, however, f is

[13] An element a is defined to be separable if its minimal polynomial is separable.

[14] This is the field of rational functions over $\mathbb{F}_2$: $P(T)/Q(T)$, with $P, Q \in \mathbb{F}_2[T]$ and $Q \neq 0$, though viewing them as functions isn't always ideal, for the same reasons as discussed for polynomials previously.

irreducible and seperable of degree n, then

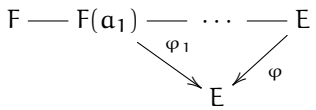

and there are n choices for φ_1 (i.e. for $\varphi(\alpha_1)$) because f is irreducible, so any root can be sent to any other root. Thus:

I. $n \mid |\mathrm{Gal}(E/F)| = |E : F|$, and
II. the Galois group is a transitive permutation subgroup of S_n (that is, given any α_i, α_j, there exists a $\varphi \in \mathrm{Gal}(E/F)$ such that $\varphi(\alpha_i) = \alpha_j$), or the induced group action has only one orbit).

it's hard to say much more than this; if ζ is a root of $\Phi_d(x)$, then $\mathbb{Q}(\zeta)/\mathbb{Q}$ requires $\varphi(\zeta) = \zeta^i$ for some i such that $(i, n) = 1$, which determines the options for all of the other roots.

11. Small Degrees and Finite Fields: 2/1/13

Suppose that $F \subset E = F(\alpha_1, \ldots, \alpha_n)$ is the splitting field of an irreducible, separable $f \in F[x]$, so that $n \mid |\mathrm{Gal}(E/F)| = |E : F|$ and there is an injection $\mathrm{Gal}(E/F) \hookrightarrow S_n$, with the Galois group a transitive subgroup.

Example 11.1.

- If $n = 2$, then the Galois group must be cyclic of order 2.
- If $n = 3$, then the Galois group is either S_3 or $A_3 = \{1, (1\,2\,3), (1\,3\,2)\} < S_3$. The latter group is isomorphic to C_3 (the cyclic group of order 3, written multiplicatively).

  Specifically, if $F = \mathbb{Q}$ and f has exactly one real root, then $\mathrm{Gal}(E/\mathbb{Q}) = S_3$, since the first root has degree 3 and the remaining roots have degree 2.
- If $n = 4$, it's tedious but not hard to show that the Galois group is (up to isomorphism) one of five subgroups of S_4:
  - $C_4 = \langle (1\,2\,3\,4) \rangle$ (i.e. cyclic of order 4),
  - $V_4 \cong \mathbb{Z}/2 \times \mathbb{Z}/2 = \{1, (1\,2)(3\,4), (1\,3)(2\,4), (1\,4)(2\,3)\}$[15] (i.e. the Klein-four group).
  - D_4, the dihedral group of order 8.[16]
  - A_4 and S_4, similarly to before.

  It's already hard to determine the Galois group in the cubic case, and it becomes harder here. Some specific cases are simpler, however: if $f(x) = x^4 + x^3 + x^2 + x + 1 = \Phi_5(x)$ (so that its roots are $\zeta, \ldots, \zeta_4$), then $|\mathrm{Gal}(\mathbb{Q}(\zeta)/\mathbb{Q})| = 4$ (as shown previously in the cyclotomic case), and in particular $\mathrm{Gal}(\mathbb{Q}(\zeta)/\mathbb{Q}) \simeq (\mathbb{Z}/5)^* \simeq C_4$.
- For $n = 5$, there are still only 5 options, including S_5, A_5, C_5, and D_5. Letting $F = \mathbb{Q}$ and $f \in \mathbb{Q}[x]$ have exactly three real roots, $\sigma(z) = \bar{z}$ is a nontrival automorphism that is a transposition, so it can be written as $(1\,2)$. There is also an element of order 5, since $5 \mid |\mathrm{Gal}(E/\mathbb{Q})|$ (by Cauchy's Theorem). Thus, it has to be a 5-cycle $(1\,2\,3\,4\,5)$. It happens that $\langle (1\,2), (1\,2\,3\,4\,5) \rangle = S_5$, so $\mathrm{Gal}(E/\mathbb{Q}) \simeq S_5$.[17]

  Since S_5 isn't solvable, then these roots can't be given by a formula with nested radicals, as will be shown later.

Suppose $\mathbb{F}_p \overset{n}{\text{——}} E = \mathbb{F}_{p^n}$. In addition to the identity, we also have the Frobenius automorphism $\sigma(a) = a^p$ for $a \in \mathbb{F}_{p^n}$. Iterating, $\sigma^k = \mathrm{Id}$ first when $k = n$, so $|\sigma| = n$ in S_n. (This is because $a^{p^n} = a$ for all $a \in \mathbb{F}_{p^n}$, and since $\mathbb{F}_{p^n} \setminus 0$ is cyclic, then that many iterations is necessary.) Thus, $\mathrm{Gal}(\mathbb{F}_{p^n}/\mathbb{F}_p) = C_n$.

If we stick in an $\mathbb{F}_{p^d}$ (where $d \mid n$) as $\mathbb{F}_p \overset{d}{\text{——}} \mathbb{F}_{p^d} \overset{n/d}{\text{——}} \mathbb{F}_{p^n}$, then $\mathrm{Gal}(\mathbb{F}_{p^n}/\mathbb{F}_{p^d})$ is the subgroup of $\mathrm{Gal}(\mathbb{F}_{p^n}/\mathbb{F}_p)$ consisting of permutations that fix $\mathbb{F}_{p^d}$ (on top of just $\mathbb{F}_p$). Then, $\mathrm{Gal}(\mathbb{F}_{p^n}/\mathbb{F}_{p^d}) = C_{n/d} \leqslant C_n$, and it is generated by σ^d (as a check, this does give the right order). Then, $(\sigma^d)^{n/d} = \mathrm{Id}$, and $\sigma^d(x) = x^{p^d}$, so this fixes $\mathbb{F}_{p^d}$.

Much of this involves writing enough properties of the group to completely characterize it (e.g. if the Galois group is nonabelian and of order 10, then it must be D_{10}), rather than writing down a big multiplication table (which is $O(n^2)$ memory anyways).

Definition. Compare $\mathbb{F}_p$ ——L——$\mathbb{F}_{p^n}$ (subgroups of $\mathrm{Gal}(\mathbb{F}_{p^n}/\mathbb{F}_p) \simeq C_n$); thee is a one-to-one correspondence between $H \leqslant \mathrm{Gal}(\mathbb{F}_{p^n}/\mathbb{F}_p)$ and $L = \{x \in E \mid hx = x \text{ for all } h \in H\}$. L is called the fixed field of H, and is sometimes written E^H. (This definition is valid in any normal, separable, finite extension.)

[15] There are multiple copies of V_4 inside S_4, but the requirement of transitivity restricts the options.

[16] Similarly, there are multiple such subgroups; however, since they are the Sylow-2 subgroups, they are all conjugate.

[17] This argument works for any prime p and a polynomial with $p - 2$ real roots and 2 complex ones.

For example, if f is a 5^{th}-degree polynomial with three real roots r_1, r_2, r_3, then if H is the subgroup generated by conjugation, then $E^H = \mathbb{Q}(r_1, r_2, r_3)$.

Going the other way, it's possible to start with a subfield L and assign a subgroup $H_L = \{\sigma \in \mathrm{Gal}(E/F) \mid \sigma x = x \text{ for all } x \in L\} = \mathrm{Gal}(E/L) \leqslant \mathrm{Gal}(E/F)$. Since E/F is Galois, then E/L is too (after all, it's a splitting field and separable).

The Fundamental Theorem of Galois Theory will be stated here, but proved later, in several chunks.

Theorem 11.1 (Fundamental Theorem of Galois Theory)**.** *The correspondences above are bijective: if* $L = E^H$*, then* $\mathrm{Gal}(E/L) = H$*, and if* $H = \mathrm{Gal}(E/L)$*, then* $E^H = L$*.*

As a minor digression, suppose K/F is a seperable algebraic extension of degree m. Then, the minimum polynomial for $\alpha = \alpha_1$ is $f(x) = (x - \alpha_1)\dots(x - \alpha_n)$ which has roots $\sigma_1(\alpha), \dots, \sigma_n(\alpha)$, where the $\sigma_j(\alpha)$ are the distinct options, called the conjugates of α in an algebraic closure Ω. This allows the finding of minimum polynomials (e.g. $\sqrt{2} + \sqrt{3}$ has four embeddings: $\pm\sqrt{2} \pm \sqrt{3}$, which are all of the roots, and then the minimum polynomial is each linear factor multiplied).

12. The Theorem of the Primitive Element: 2/4/13

What is the number of field embeddings $K \xrightarrow{\phi} \Omega$ (where K/F is a field extension and Ω is an algebraic closure of F) such that $\phi|_F = \mathrm{Id}$?

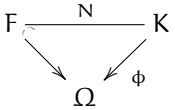

We know the following facts (which are review from previous section):

(1) The number of such embeddings is at most $|K : F|$.
(2) If K is seperable over F, then the number is exactly $|K : F|$.
(3) If K is normal over F (i.e. the splitting field of some polynomial), then every $\phi(K)$ is a splitting field of F in Ω (and in particular, if $K \subset \Omega$,[18] then $\phi(K) = K$).
(4) Thus, if K/F is Galois (i.e. finite, normal, and separable) with $K \subset \Omega$, then Ω is an algebraic closure of K, every $\phi : K \to \Omega$ is an automorphism of K that fixes F, and $|K : F| = |\mathrm{Gal}(K/F)|$.
(5) If $\alpha \in K \subset \Omega$ and K is separable over F, then the minimal polynomial for α over F is $f(x) = \prod_{j=1}^{d}(x - \alpha_j)$, where $\{\alpha_1, \dots, \alpha_d\}$ is the set of distinct $\phi(\alpha)$, where ϕ are the embeddings of K in Ω. However, it is not necessarily true that $|F(\alpha) : F| = |K : F|$, though (we do know the former quantity is equal to d).

In order to understand the Fundamental Theorem of Galois Theory, it will be necessary to review Theorems 8 and 27 in chapter 13 of the textbook.

As an example of point 5, let $\alpha = i + \sqrt[4]{2}$ over $\mathbb{Q}$. Then,

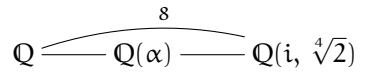

(so that the last field is the splitting field) and $\phi(\alpha) = \pm i \pm \sqrt[4]{2}$. This gives eight possibilities, so the minimal polynomial can be found by multiplying their linear factors together, and since $|\mathbb{Q}(\alpha) : \mathbb{Q}| = 8$, then $\mathbb{Q}(\alpha) = \mathbb{Q}(i, \sqrt[4]{2})$.

"Proof" of point 5.

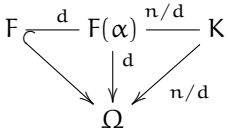

There are d embeddings of $F(\alpha) \to \Omega$ because of Theorem 8, and there are n/d of K into Ω by Theorem 27. This is their real power. ⊠

The next theorem, which will be important for proving the Fundamental Theorem, considers finding primitive elements (i.e. generators) for a finite seperable extension.

Theorem 12.1 (Primitive Element)**.** *If* K/F *is a finite, seperable extension, then there exists an* $\alpha \in K$ *such that* $K = F(\alpha)$*. This* α *is called the primitive element.*

[18] This is a set-theoretic issue: it could be that Ω contains elements unrelated to K other than how the operation is defined, in which case equality doesn't really hold. However, it's generally more useful to imagine $K \subset \Omega$.

Proof. If $|F|$ is finite, then this is trivial since $K^* = K \setminus 0 = \langle\theta\rangle$ since it's cyclic, so assume F is infinite, and consider $F(\beta,\gamma)$, where β,γ are seperable over F. The goal is to show that $F(\beta,\gamma) = F(\alpha)$ for some $\alpha = \beta + c\gamma$ for some $c \in F$, or that if

$$\begin{array}{ccccc} & \overset{d}{\nearrow} & F(\beta) & \xrightarrow{f} & F(\beta,\gamma) \\ F & & & & \\ & \underset{e}{\searrow} & F(\gamma) & & \end{array}$$

then $f \leqslant e$ and $df = |F(\beta,\gamma) : F|$.

Suppose $\beta_1, \dots, \beta_d$ (with $\beta = \beta_1$) are cojugates of β and $\gamma = \gamma_1$ with $\gamma_1, \dots, \gamma_e$ as its conjugates, each in some Ω. Not all of these are legal, but f of them are (for a given β). Then, if $\beta_i + c\gamma_j = \beta_{i'} + c\gamma_{j'}$, then $c = 0$ or $j = j'$ if $i = i'$. If $j = j'$, then $i = i'$, and otherwise, $c = \frac{\beta_i - \beta_{i'}}{\gamma_j - \gamma_{j'}}$. Since there are finitely many such c, just avoid them, and you can choose any other c that works. Thus, $\alpha = \beta + c\gamma = \beta_i + c\gamma_i$ has $df = |F(\beta,\gamma) : F|$, so $F(\beta,\gamma) = F(\alpha)$. ⊠

In some sense, the primitive generator is a linear combination of the previous generators obtained by looking at the roots of the minimum polynomial. Additionally, now the Fundamental Theorem just drops out like melted butter.

13. The Fundamental Theorem of Galois Theory: 2/6/13

First, a small amount of review: suppose E/F is finite, normal, and seperable and $F \subset L \subset E$ is a tower. Then, E/L is also finite, normal, and seperable, so $|\mathrm{Gal}(E/L)| = |E : L|$. In fact, for any algebraic K/L in a closure Ω, the number of embeddings $K \to \Omega$ over L (i.e. the identity on L) is at most $|K : L|$, and this becomes exactly $|K : L|$ if the extension is seperable.

Returning to the (finite) Galois case, $G_L = \mathrm{Gal}(E/L) \leqslant \mathrm{Gal}(E/F) = G_F$ (since if $\sigma|_L = \mathrm{Id}$, then $\sigma|_F = \mathrm{Id}$), so a subgroup can be assigned to a subfield, and a subfiel can be assigned to a subgroup $H \leqslant G_F$ given by $E^H = \{x \in E \mid hx = x \text{ for all } x \in H\}$.

The proof of the Fundamental Theorem will be done in parts.

Claim. $E^{G_L} = L$.

Proof. We have $F \text{——} L \text{——} E^L \text{——} E$, because DUH: is L fixed by automorphisms that fix L? Do bears crap in the woods?[19] But now, compute some degrees: $|E : L| = |G_L|$, but $|E : E^{G_L}| \geqslant |G_L|$, because the number of automorphisms (or even embeddings in the algebraic closure) is at most the degree of the field extnsion. Thus, $|E : L| = |E : E^{G_L}|$, so $E = E^{G_L}$.[20] ⊠

This is very similar to the idea behind Theorem 27, even if the statement is different.

After the proof of the following claim, there is a bijection between intermediate fields and subgroups of G_F, which is the Fundamental Theorem.

Claim. $H = \mathrm{Gal}(E/E^H)$.

Proof. One inclusion is easy: $H \leqslant G_{E^H}$ because DUH: does H fix things fixed by H?

For the other direction, consider $F \subset E^H \subset E$. Using Theorem 12.1, $|G_{E^H}| = |E : E^H| \leqslant |H|$, so one obtains equality: $H = G_{E^H}$.

This is the only nontrivial part of the whole proof, so it needs to be looked at in more detail. The textbook provides a very elaborate proof using the linear independence of characters, for example. However, here is an alternate explanation: suppose E/E^H is a d^{th}-extension and $E = E^H(\alpha)$; then, d is the degree of the minimal polynomial of α over E^H.

Consider $g(x) = \prod_{x \in H}(x - h\alpha)$. A priori, this is only in $E[x]$, but the coefficients are in E^H: since α is a root of g, then $d \leqslant |H|$. Thus, writing $H = \{h_1 h \mid h \in H\}$ for any $h_1 \in H$, applying h_1 to g just gives the coefficients back again. Thus, they're fixed by H, so $g \in E^H[x]$. ⊠

It's possible to get the minimum polynomial, not just g, by listing the distinct roots, but it's not that valuable. There is also something asymmetric about bringing in the primitive element, but it's much easier than the book's proof. Nonetheless, some of the results in the book, such as the linear independence of characters, are useful later.

[19] Mind you, Dr. Brumfiel said this aloud and wrote it on the board.

[20] In general, just because two extensions have the same degree doesn't mean they are equal, but in this case, one is contained inside the other.

Now, one can talk about conjugate subgroups of G_F and conjugate subfields between F and E. First suppose that

$$\begin{array}{ccccc} F & \xrightarrow{\text{normal}} & L & \longrightarrow & E \\ \wr\downarrow{\scriptstyle \mathrm{Id}} & & \wr\downarrow{\scriptstyle \sigma} & & \wr\downarrow{\scriptstyle \sigma} \\ E & \longrightarrow & \sigma L = L & \longrightarrow & E \end{array}$$

where L is a splitting field of F. Then, the roots of F are just permuted among themselves, and there is a group homomorphism $\mathrm{Gal}(E/F) \to \mathrm{Gal}(L/F)$ given by $\sigma \mapsto \sigma|_L$ with kernel $\mathrm{Gal}(E/L)$. This is surjective because if $\varphi : L \xrightarrow{\sim} L$ over F, then by the proof of Theorem 27, this extends to a $\sigma : E \xrightarrow{\sim} E$ (just draw the diagram).

Thus, $\mathrm{Gal}(L/F) \cong \mathrm{Gal}(E/F)/\mathrm{Gal}(E/L)$ (requiring L/F to be normal, so that the quotient makes sense. Thus, normality in these two senses is the same!). In general, if

$$\begin{array}{ccccc} F & \longrightarrow & L & \longrightarrow & E \\ \wr\downarrow & & \wr\downarrow{\scriptstyle \sigma} & & \wr\downarrow{\scriptstyle \sigma} \\ E & \longrightarrow & \sigma L & \longrightarrow & E \end{array}$$

where $\sigma \in \mathrm{Gal}(E/F)$, and σL is still a subfield, but not necessarily L.

Claim. $\mathrm{Gal}(E/\sigma L) = \sigma\,\mathrm{Gal}(E/L)\sigma^{-1}$.

For the proof, see this week's homework.

14. Lattices of Subgroups and Subfields: 2/8/13

Last time, it was shown that $L_i \leftrightarrow \mathrm{Gal}(E/L_i) = H_{L_i}$ and $H \leftrightarrow E^H$, and that $H \lhd \mathrm{Gal}(E/F)$ iff the corresponding fixed field L is a normal extension of F. Slightly more generally, for a $\sigma \in \mathrm{Gal}(E/F)$ and $L = E^H$, one has $E^{\sigma H \sigma^{-1}} = \sigma L \sigma^{-1} \subset E$. Notice how this relates to normality in the special case. This can be proven by appling the group elements $\sigma k \sigma^{-1}$ to the field elements. If H is normal, then the quotient makes sense, and $\mathrm{Gal}(L/F) \cong \mathrm{Gal}(E/F)/\mathrm{Gal}(E/L)$.

If $F \text{——} L \xrightarrow{\subseteq} L' \text{——} E$, then $\mathrm{Gal}(E/L) \geqslant \mathrm{Gal}(E/L')$, as we already know. Similarly, if $H \leqslant H'$, then $E^H \supseteq E^{H'}$. The point is that the sign switches, and one says that the Galois correspondence reverses inclusions.

One can consider the lattice of subgroups or subfields. There are lattice operations: if $H_1, H_2 \leqslant H$, then $H_1 \cap H_2 \leqslant H$, and $H_1 H_2 \leqslant H$.[21] In the field case, one can take the composite fields $(L_1, L_2) \mapsto L_1 L_2$ (the smallest subfield containing them) and intersections $L_1 \cap L_2$.

The last part of the Fundamental Theorem states that in this Galois correspondence, $H_1 \cap H_2 \leftrightarrow L_1 L_2$ and $H_1 H_2 \leftrightarrow L_1 \cap L_2$, yielding a lattice-type structure similar to $\mathbb{Z}$ under least common multiple and greatest common divisor, or sets under union and intersection.

Proof. $H_1 \cap H_2$ fixes $L_1 L_2$ because anything in $H_1 \cap H_2$ fixes everything in L_1 and everything in L_2, so all of the products are fixed, too.

Conversely, suppose $\sigma \notin H_1 \cap H_2$. Then, suppose (without loss of generality) $\sigma \notin H_1$, so that σ doesn't fix L_1. Thus, there is some $x \in L_1$ such that $\sigma(x) \neq x$, so σ can't fix $L_1 L_2$, either.

The other part is slightly easier; it's clear that $H_1 H_2$ fixes $L_1 \cap L_2$: $h_1 h_2(x) = h_1(x) = x$ for any $h_1 \in H_1$, $h_2 \in H_2$, and $x \in L_1 \cap L_2$ (since both h_1 and h_2 must fix x). Conversely, if $x \notin L_1 \cap L_2$, there is an $h \in H_1$ (without loss of generality, since if $h \in H_2$ the proof is little different) such that $h(x) \neq x$, so $H_1 H_2$ doesn't fix x. ⊠

Example 14.1. Let E be the splitting field of $x^n - 1$ over F. If $\mathrm{Char}(F) \mid n$, this is a degenerate case, so suppose otherwise. Then, $E = F(\zeta)$, where ζ is the generator of the cyclic group of roots of unity. Thus, if $\sigma \in \mathrm{Gal}(E/F)$, then $\sigma(\zeta) = \zeta^i$ for some i for which $\gcd(i, n) = 1$. Thus, $\mathrm{Gal}(E/F) \leqslant (\mathbb{Z}/n)^*$. In particular, it's abelian, and if F is algebraically closed, then it's also trivial.

If $F = \mathbb{Q}$, then it was already shown that $\mathrm{Gal}(E/F) = (\mathbb{Z}/n)^*$, since $\Phi_n(x)$ is irreducible. if F is finite, one needs a cyclic group of order n (since it's a subset of $E^* = E \setminus 0$) to split $x^n - 1$, so all Galois groups of the finite fields are cyclic (again, as was already known).

As another example, if $\zeta \in F$ is a primitive n^{th} root of unity, consider the splitting field of $x^n - a$ for some $a \in F$. Then, one gets $F(\alpha)/F$, where $\alpha^n = a$, and the other roots are $\{\zeta^i \alpha \mid 0 \leqslant i \leqslant n-1\}$, and they all lie in $F(\alpha)$. If $\sigma_i \in \mathrm{Gal}(F(\alpha)/F)$, then $\sigma_i(\zeta) = \zeta$ and $\sigma_i(\alpha) = \zeta^i \alpha$, so the Galois group of $x^n - a$ is contained in $\mathbb{Z}/n$ and thus is cyclic.[22]

[21] This is the smallest subgroup containing H_1 and H_2, not the set of products, which is a group only when one of H_1 and H_2 is normal in H.

[22] If $\alpha \in F$, then it's not equality, so one must be careful with terminology.

Now, if $x^n - a \in F[x]$, where $\mathrm{Char}(F) = 0$ or $\mathrm{Char}(F) = p \nmid n$, this can be split in two steps: $F \overset{\lhd}{\text{——}} F(\zeta) \text{——} F(\zeta, \alpha) = E$, so $\mathrm{Gal}(F(\zeta)/F) \leqslant (\mathbb{Z}/n)^*$ and $\mathrm{Gal}(F(\zeta, \alpha)/F(\zeta)) \leqslant (\mathbb{Z}/n, +)$. Thus,

$$1 \to N \to \mathrm{Gal}(E/F) \to \mathrm{Gal}(E/F(\zeta)) \to 1$$

(where N is the Galois group of $x^n - 1$ over F) is a short exact sequence, so $\mathrm{Gal}(E/F(\zeta)) = \mathrm{Gal}(E/F)/N$. Note that $\mathrm{Gal}(E/F)$ need not be abelian (e.g. $x^3 - 2$ over $\mathbb{Q}$, which yields S_3).

One thing about root extensions is that if $F \text{——} F(\sqrt[n_1]{a_1}) \text{——} F(\sqrt[n_1]{a_1}, \sqrt[n_2]{a_2})$, with $a_2 \in F(\sqrt[n_1]{a_1})$, then $\sqrt[n_2]{a_2}$ is a doubly nested radical (e.g. $\sqrt{1 + \sqrt[3]{5}}$). This will have consequences for solvability by radicals. These fields correspond to subgroups that put severe restrictions on the final Galois group of E/F.

This implies that some algebraic numbers can't be expressed with nested radicals (e.g. roots of any quintic with exactly 3 real roots, since it will have Galois group S_5, which isn't solvable).

The following theorem illustrates where the next steps will be:

Theorem 14.1. *Suppose E/F is a Galois extension with group cyclic of order n. Then, $E = F(\sqrt[n]{\alpha})$ for some $\alpha \in F$.*

15. Nested Radicals and Norms and Traces: 2/11/13

Suppose $\mathrm{Gal}(E/F) = S_3$, which has $1 \lhd A_3 \lhd S_3$, with $A_3 = \{1, (1\,2\,3), (1\,3\,2)\}$. Thus, $F \overset{2}{\text{——}} L \overset{3}{\text{——}} E$ for sole subfield L of E. Then, if $\mathrm{Char}(F) \neq 2$, then $L = F(\sqrt{D})$ for some $D \in F$, and $\mathrm{Gal}(E/L) = A_3$ (which is cyclic of order 3). E is a little more complicated (not just $L(\sqrt[3]{\gamma})$ for some $\gamma \in L$), but since the cube roots of unity are just $(1 \pm i\sqrt{3})/2$, then $E(\zeta) = L(\zeta)(\sqrt[3]{\gamma})$ for some $\gamma \in F(\sqrt{D}, \zeta)$.

This means that every element of E has a nested radical formula. The key to this is the series of normal subgroups and that $E(\zeta) = F(\sqrt{D}, \zeta, \sqrt[3]{\gamma})$, where γ is something in D and ζ, immplying it is nested.

If $\mathrm{Gal}(E/F) = S_4$, with $\mathrm{Char}(F) \neq 2, 3$, then $S_4 \rhd A_4 \rhd V_4 \rhd C_2 \rhd \{1\}$, where $V_4 = \{1, (1\,2)(3\,4), (1\,3)(2\,4), (1\,4)(2\,3)\}$. Thus, one has $F \overset{2}{\text{——}} F(\sqrt{D}) = E^{A_4} \overset{3}{\text{——}} E^{V_4} \overset{2}{\text{——}} E^{C_2} \overset{2}{\text{——}} E$. If ζ is a primitive root of unity, then take $E(i, \zeta)/F(i, \zeta)$; then, each field is obtained from the previous one by adjoining an n^{th} root:

$$F(i, \zeta) \overset{2}{\text{——}} F(i, \zeta, \sqrt{D}) \overset{3}{\text{——}} F(i, \zeta, \sqrt{D}, \sqrt[3]{\gamma}) \overset{2}{\text{——}} F(i, \zeta, \sqrt{D}, \sqrt{D_2}, \sqrt[3]{\gamma}) \overset{2}{\text{——}} E,$$

so everything in E can be written as a nested radical formula over F.

These will be encapsulated into more general proofs, depending on the Fundamental Theorem's linking of group theory and field theory. However, this will break down if $\mathrm{Gal}(E/F) = S_5 \rhd A_5 \rhd \{1\}$, so there is no chain of intermediate normal field extensions of prime order.

Definition. A finite group G is solvable if there exist subgroups $G = G_0 \rhd G_1 \rhd G_2 \rhd \cdots \rhd G_m = \{1\}$ such that $|G_i : G_{i+1}|$ is prime (so that the quotients are prime cyclic: $G_i/G_{i+1} \cong C_{p_i}$).

There are other definitions of solvability, as will be seen, but they are all equivalent.

The end goal is to prove the following theorem:

Theorem 15.1. *If $\mathrm{Char}(F) = 0$ and $f \in F[x]$ is separable, then its roots are given by a nested radical formula iff the Galois group of the splitting field is solvable.*

The more interesting or dramatic direction (the reverse direction) is actually the easiest; the formula comes from Theorem 27.

Suppose K/F is a finite, separable degree-n extension (that isn't necessarily normal) and $\Omega \supset K$ is an algebraic closure of K (and F). There are n embeddings $\sigma_i : K \to \Omega$ that fix F, by Theorem 27.

Definition. If $\alpha \in K$, then the norm of α over F is $N_{K/F}(\alpha) = \prod_{i=1}^{n} \sigma_i(\alpha)$, where the σ_i are as above.

If K/F is also normal, then this is also $N_{K/F}(\alpha) = \prod_{\sigma \in \mathrm{Gal}(K/F)} \sigma(\alpha)$.

One of the homework problems develops seveal properties of $N_{K/F}$. For example,

- $N_{K/F}(\alpha\beta) = N_{K/F}(\alpha)N_{K/F}(\beta)$ for any $\alpha, \beta \in K$ (since each of the σ_i is a field homomorphism).
- $N_{K/F}(\alpha) = \alpha^n$ if $\alpha \in F$, since it is fixed by all of the σ_i.
- If $F \overset{d}{\text{——}} F(\alpha) \overset{n/d}{\text{——}} K$ and $\alpha_1, \ldots, \alpha_d$ are the distinct conjugates of α in Ω, then $N_{K/F}(\alpha) = \prod_{i=1}^{d} \alpha_i^{n/d}$ (again by the proof of Theorem 27).
- If the minimum polynomial for α over F is $\sum_{i=0}^{d} a_i x^i = \prod_{i=1}^{d} (x - \alpha_i)$ (so that $a_i = 0$), then the last coefficient is all of the α_i multiplied together, so $N_{K/F}(\alpha) = ((-1)^d a_d)^{n/d}$.

Definition. Similarly, the trace of an $\alpha \in K$ over F is $T_{K/F}(\alpha) = \sum_{i=1}^{n} \sigma_i(\alpha)$ (sometimes written $\mathrm{Tr}_{K/F}(\alpha)$).

Some similar properties hold: $T_{K/F}(\alpha+\beta) = T_{K/F}(\alpha) + T_{K/F}(\beta)$, since each σ_i is a linear homomorphism, but also if $c \in F$, then $T_{K/F}(c\alpha) = cT_{K/F}(\alpha)$, so the trace is a more familiar algebraic animal: it's a linear homomorphism.

It's important that since $N_{K/F}(\alpha) = ((-1)^d a_d)^{n/d}$, then it is in F, rather than just in Ω. Similarly, if $\alpha_1, \dots, \alpha_d$ are the distinct conjugates of α, then $\sum_{i=1}^{n} \sigma_i(\alpha) = n/d \sum_{i=1}^{d} \alpha_i$. Taking the minimum polynomial again, this becomes $T_{K/F}(\alpha) = -n(a_1)/d \in F$. Thus, $T_{K/F} : K \to F$ is an F-linear homomorphism, which is much easier to understand, because of linear algebra.

16. More Norms and Traces: 2/13/13

Suppose K/F is a separable degree-n field extension and L/K is any normal, separable extension of F:

$$\begin{array}{ccccc} F & \xrightarrow{n} & K & \longrightarrow & L \\ & & \searrow^{\sigma_i} & \swarrow_{\tilde{\sigma}_i} & \\ & & \Omega & & \end{array}$$

where Ω is a closure of K and F. Then, the n embeddings of K in Ω can be extended to L, such that $\tilde{\sigma}_i(L) = L$. Thus, one can ignore Ω and just talk about embeddings in L.

The last lecture just considered K/F for norms, but someting interesting happens in a more complicated setup:

Proposition 16.1. *Suppose* $F \text{——} K \text{——} E$ *are both separable field extensions and* $\alpha \in E$*; then,* $N_{E/F}(\alpha) = N_{K/F}(N_{E/K}(\alpha))$.

Proof. Consider

$$F \xrightarrow{n} K \overset{m}{\underset{\tau_j}{\overset{\sigma_i}{\rightrightarrows}}} E \longrightarrow L,$$

where each of these is normal and finite over F and the σ_i fix F and the τ_j fix K. Then, lift to $\tilde{\sigma}_i, \tilde{\tau}_j : L \xrightarrow{\sim} L$.

Claim. $\{\tilde{\sigma}_i \circ \tilde{\tau}_j|_E \mid 1 \leqslant i \leqslant n, 1 \leqslant j \leqslant m\}$ are all distinct and give all of the embeddings $E \to \Omega$ over F.

Proof. First, notice that it's the right number of embeddings, so if they are all distinct, then they are all accounted for.

Suppose $\tilde{\sigma}_i \circ \tilde{\tau}_j = \tilde{\sigma}_{i'} = \tilde{\tau}_{j'}$. Then, take a $\beta \in K$, so $\tilde{\tau}_j(\beta) = \tilde{\tau}_{j'}(\beta) = \beta$, so $\tilde{\sigma}_i(\beta) = \tilde{\sigma}i'(\beta)$ for all $\beta \in K$, so $\tilde{\sigma}_i = \tilde{\sigma}_{i'}$. Since $\tilde{\sigma}_i$ is an automorphism of L, this means $\tilde{\tau}_j = \tilde{\tau}_{j'}$ on all of E. ⊠

Then, since $\tilde{\sigma}_i$ is fixed on any of the $\tilde{\tau}_j(\alpha)$, then

$$\begin{aligned} N_{E/F}(\alpha) &= \prod_{i=1}^{n}\prod_{j=1}^{m} \tilde{\sigma}_i(\tilde{\tau}_j(\alpha)) = \prod_{i=1}^{n} \tilde{\sigma}_i \left(\prod_{j=1}^{n} \tilde{\tau}_j(\alpha) \right) \\ &= \prod_{i=1}^{n} \sigma_i(N_{E/K}(\alpha)) = N_{K/F}(N_{E/K}(\alpha)). \end{aligned}$$ ⊠

Lastly, suppose K is a finite, separable, degree-n extension over F and $\alpha \in K$. Then, $K \overset{(\cdot\alpha)}{\to} K$ (i.e. multiplication by α) is a vector-space isomorphism[23] that is linear over F, which will be shown in a homework problem.

Claim. $N_{K/F}(\alpha) = \det(\cdot\alpha)$.

Proof. Consider $F \xrightarrow{d} F(\alpha) \xrightarrow{n/d} K$ and let $f(x) = x^d + a_1 x^{d-1} + \cdots + n_d \in F[x]$ be the minimum polynomial if α. Then, an F-basis for $F(\alpha)$ is $\{1, \alpha, \dots, \alpha^{d-1}\}$, and pick an $F(\alpha)$-basis $\{\beta_1, \dots, \beta_m\}$ for K. Then, conider $K \overset{(\cdot\alpha)}{\to} K$ in the basis $\{\alpha^i \beta_j \mid 0 \leqslant i \leqslant d-1, 1 \leqslant j \leqslant m\}$; the matrix for this operator consists of blocks of the form of

$$\begin{pmatrix} 0 & 0 & \dots & 0 & -a_d \\ 1 & 0 & \dots & 0 & -a_{d-1} \\ 0 & 1 & \dots & 0 & -a_{d-2} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & 1 & -a_1 \end{pmatrix}$$

on the main diagonal and zeroes elsewhere. The trace of one of these blocks is $-a_1$, so the determinant of the full matrix is $\det(\cdot\alpha) = ((-1)^d a_d)^m = N_{K/F}(\alpha)$ (from the last lecture). ⊠

[23] . . . unless $\alpha = 0$, in which case it's still F-linear, albeit not invertible.

Similarly, one can show that $T_{K/F}(\alpha) = \mathrm{Tr}(\cdot\alpha)$ by again putting it into block form. Since this is a linear map, and there is transitivity if K/F and E/K, so if $\beta \in E$, then $T_{E/F}(\beta) = T_{K/F}(T_{E/K}(\beta))$, which is proven by going to the Galois closure and taking the double sum, rather than the double product.

If E/F is Galois, $\beta \in E$, and $\sigma \in \mathrm{Gal}(E/F)$, then $N_{E/F}(\sigma\beta/\beta) = 1$, which can be seen in many ways: it can be plugged into the definition, for example (since the σ_i just permute the things being put together). This induces the following theorem:

Theorem 16.2 (Hilbert's Theorem 90)**.** *If* $\mathrm{Gal}(E/F) = C_n = \langle\sigma\rangle$ *and* $N_{E/F}(\gamma) = 1$*, then* $\gamma = \sigma\beta/\beta$ *for some* $\sigma \in \mathrm{Gal}(E/F)$ *and* $\beta \in E \setminus 0$*.*

This means that $\sigma\beta = \gamma\beta$ for some nonzero β.[24] The proof is easy using "characters" (which are defined in the book), but requires their linear independence.

As an application, suppose E/F has degree n and a cyclic Galois group and suppose $\zeta \in F$ is a primitice n^{th} root of unity. Then, $N_{E/F}(\zeta) = \zeta^n = 1$, so $\sigma\beta = \zeta\beta$ for some $\beta \in E \setminus 0$. Thus, β has conjugates $\sigma^2]be t = \zeta^2\beta$, etc. Thus, $F(\beta) = E$ and β is a primitive element, and $\sigma(\beta^n) = (\sigma\beta)^n = (\zeta\beta)^n = \beta^n$, so $\beta^n = b$ is fixed, so $b \in F$, and $E = F(\sqrt[n]{b})$.

17. Solvable Groups: 2/15/13

This is the big theorem:

Theorem 17.1. *If* $f \in F[x]$ *is separable and irreducible and* E *is a splitting field of* f *over* F*, with* $\mathrm{Char}(F) \neq 0$ *or* $\mathrm{Char}(F) = p$ *such that* $p \nmid |E : F| = |\mathrm{Gal}(E/F)|$*, then the roots of* $f(x)$ *in* E *are given by nested radical formulas iff* $\mathrm{Gal}(E/F)$ *is a solvable group.*

Recall that a finite group G is solvable iff there exists a chain of subgroups $G = G_0 \rhd G_1 \rhd \cdots \rhd G_m = \{e\}$ (i.e. each G_{i+1} is normal in G_i) and $G_i/G_{i+1} \simeq C_{p_i}$, a cyclic group of prime order.

This leads to an impressive analogy due to Jordan and Holder: simple groups are as atoms, and finite groups are molecules. In particular, some sets of atoms can be assembled into distinct molecules, leading to some notion of group isomers.[25]

The simple groups include C_p for prime p, A_n for $n \geqslant 5$, and several other infinite families, including $\mathrm{SL}(n, \mathbb{F}_q)$ for most values of n and q, along with 26 sporadic simple groups, some of which are quite large.

Theorem 17.2. *Any finite group has a composition series* $H = H_0 \rhd H_1 \rhd \cdots \rhd H_M = \{e\}$ *with* $H_i/H_{i+1} = N_i$ *such that all of the* N_i *are simple.*

This is trivial to prove; if H is simple, we're done, but if not, choose some maximal proper normal subgroup (whose existence is guaranteed by finiteness), so the quotient must be simple, since there are no other normal subgroups in between the two.

Theorem 17.3 (Jordan-Holder)**.** *The set* $\{N_i\}$ *is the same for all composition series of* H*; that is, any in any two composition series, the quotients are just rearranged.*

Example 17.1. $S_n \rhd A_n \rhd \{1\}$ and $C_2 = S_n/A_n$ is simple.

If $|H| = 10$, then the composition series includes C_2 and C_5. There are two such groups: $C_{10} = C_2 \times C_5$ and $D_5 \rhd C_5$.

Thus, a finite group is solvable if its atoms are the simplest of simple groups, the cyclic groups of prime order.

For arbitrary groups, there are two definitions. Here is the first:

Definition. A group G is solvable if there exists a series $G = G_0 \rhd G_1 \rhd \cdots \rhd G_m = \{e\}$ with each G_i/G_{i+1} abelian.

This is equivalent to the previous definition in the case where G is finite: clearly, the first definition implies the second. But given the second, it is possible to insert groups between G_i/G_{i+1}. Since $A_i = G_i/G_{i+1}$ is finite abelian, then all of its subgroups are equivalent to subgroups of G_i containing G_{i+1}, and they're normal because it's abelian. By the structure theorem of finite abelian groups, $A_i = \bigoplus \mathbb{Z}/d_i$, and by induction, the atoms for A_i are prime cyclic.

Set $G^{(0)} = G$ and $G^{(1)} = [G, G]$ (the commutator subgroup $\left\langle aba^{-1}b^{-1} \mid a, b \in G \right\rangle$, since $[a, b]$ is known as the elementary commutator). Then, $G^{(1)} \trianglelefteq G^{(0)}$, since if one conjugates a commutator, one still has a commutator, since conjugation is a group homomrphism. G is abelian iff $G^{(1)} = \{e\}$ (since $aba^{-1}b^{-1} = 1$). Then, one can define $G^{(n)} = [G^{(n-1)}, G^{(n-1)}]$ inductively, and maybe somewhere along the line a $G^{(n)} = \{e\}$. Then, the alternate definition of solvability is:

[24] This looks like eigenvectors, which is no coincidence, but is still surprising.

[25] Solvable groups would seem to correspond to a notion of solubility.

Definition. A group G is solvable if there exists an $m \in \mathbb{N}$ such that $G = G^{(0)} \rhd G^{(1)} \rhd \cdots \rhd G^{(m)} = \{e\}$.

This is sort of a weak abelian condition on commutators of commutators, and it clearly doesn't work for simple groups: if S is nonabelian and simple then $[S, S] = S$, since there are no other normal quotients.

Claim. These two definitions presented are equivalent.

Proof. Since $G/[G, G]$ is abelian in general, then $G^{(i)}/G^{(i+1)}$ is always abelian, so the second definition implies the first.

For the other direction, start with $G = G_0 \rhd G_1 \rhd \cdots \rhd G_m = \{e\}$. Since G_0/G_1 is abelian, then $G^{(1)}$ is the subgroup of the kernel of some homomorphism (which is G_1), so $G^{(1)} \leqslant G_1$, and by the same argument $G^{(k)} \leqslant G_k$. Thus, if the identity is reached in some arbitrary series G_k, then it is also reached in the series $G^{(k)}$. ⊠

Corollary 17.4. *Suppose G is a group.*

(1) If G is solvable and $H \leqslant G$, then H is solvable.
(2) If G is solvable and $K = G/H$, then K is solvable.
(3) If $N \lhd G$ and N and G/N are solvable, then G is solvable.

Proof. For item 1, using the second definition, if $H \leqslant G$, then $H^{(m)} \leqslant G^{(m)}$ for any m, so if $G^{(m)}$ is trivial, then $H^{(m)}$ is as well.

For item 2, if $G \twoheadrightarrow K$, then $G^{(m)} \twoheadrightarrow K^{(m)}$ (since every element in K is hit by something in G), so if $G^{(m)}$ is trivial, then $K^{(m)}$ is as well.

For item 3, one has the short exact sequence $1 \to N \to G \to G/N \to 1$, so if $(G/N)^{(m)} = \{e\}$, then $G^{(m)} \leqslant N$, and if $N^{(n)} = \{e\}$, then $G^{(m+n)} \leqslant N^{(n)} = \{e\}$. ⊠

Returning to Galois theory, if $f(x)$ has roots given by nested radocal formulas (with f separable, irreducible, etc.), and suppose one root has the formula $\sqrt[5]{a + b\sqrt{c}} + \sqrt[3]{d\sqrt[7]{e} + \sqrt{g}}$ or something. Then, what would the conjugates look like? $a + b\sqrt{c} \mapsto a \pm b\sqrt{c}$, and similarly $\sqrt[5]{x}$ maps to some other fifth root, so, in the same way

$$\zeta_5 \sqrt[5]{a \pm b\sqrt{c}} + \zeta_3 \sqrt[3]{d\zeta_7 \sqrt[7]{e} \pm \sqrt{g}},$$

so if one solution is given by a nested radical formula, they all are, and the formula is essentially the same.

Additionally, if f has a root given by a nested radical formula, then its splitting field E lies in some L with

$$F \xrightarrow{\lhd} F_1 = F(\zeta_{n_1}, \dots, \zeta_{n_r}) \xrightarrow{\lhd} F_2 = F_1(\sqrt[n]{a_1}) \xrightarrow{\lhd} \dots \xrightarrow{\lhd} L.$$

with $a_j \in F_j$. This is a bit of subtlety: $L \supsetneq E$ because not all of the n_i^{th} roots of unity might be present if only one of the conjugates is. However, since all of the conjugates are in L, then L is normal.

Since $F_1 = F(\zeta_{n_1}, \dots, \zeta_{n_r})$, then $\mathrm{Gal}(E/F)$ is abelian. Then, since every necessary root of unity is in F_1, then $\mathrm{Gal}(F_k/F_{k-1})$ is cyclic for $k > 1$, yielding

$$\begin{array}{ccccccccc} \mathrm{Gal}(L/F) = G_0 & \rhd & G_1 & \rhd & G_2 & \rhd & \cdots & \rhd & G_m = \{e\} \\ \updownarrow & & \updownarrow & & \updownarrow & & & & \updownarrow \\ F & \xrightarrow{\rhd} & F_1 & \xrightarrow{\rhd} & F_2 & \xrightarrow{\rhd} & \cdots & \xrightarrow{\rhd} & L. \end{array}$$

Since $F \xrightarrow{\lhd} E \longrightarrow L$ and E splits f over F, then $\mathrm{Gal}(E/F)$ is a quotient of $\mathrm{Gal}(L/F)$ as part of the Galois correspondence. But $\mathrm{Gal}(L/F)$ is solvable, so this means that $\mathrm{Gal}(E/F)$ is solvable, too!

This was the main achievement of Galois' original work. In particular, if $\mathrm{Gal}(f) = S_5$ (as seen before), then f is not solvable by nested radicals.

18. Solvability By Radicals: 2/20/13

It is worth reviewing once again the proof that relates solvability of a group to solvability by radicals: if E/F is the splitting field of $f(x)$ with roots given in some nested radical formula, then take a big extension L/E with a tower

$$F \xrightarrow{\rhd} F(\zeta) = F_0 \longrightarrow F_1 \longrightarrow F_1' \longrightarrow F_1'' \cdots L_1 \longrightarrow F_2 \longrightarrow F_2' \cdots L_2 \cdots L,$$

where ζ is some root of unity, each of the $F^{(i)}/F^{(i-1)}$ is cyclic, and each of the L_i, as well as L, is normal over F.

In the other direction, if $\mathrm{Gal}(E/F)$ is solvable and ζ is a sufficiently large root of unity, then $\mathrm{Gal}(F(\zeta)/E(\zeta)) \leqslant \mathrm{Gal}(E/F)$ is solvable as well. Here, "sufficiently large" means ζ is a primitive n^{th} root of unity, where n is the product of the

cyclic orders in the composition series for $\mathrm{Gal}(E/F)$. Relatedly, if $\zeta_k \in F$ and L/F is cyclic of order k, then $L = F(\sqrt[k]{b})$ for some $b \in F$. This follows from Theorem 16.2 because $N(\zeta_k) = \zeta_k^k = 1$.

This can be simplified somewhat: if $L = F(\beta)$, then for some $\alpha \in L$,

$$\begin{aligned} &\beta = \alpha + \zeta\sigma(\alpha) + \zeta^2\sigma^2(\alpha) + \cdots + \zeta^{k-1}\sigma^{k-1}(\alpha) \\ \implies &\sigma\beta = \sigma(\alpha) + \zeta\sigma^2(\alpha) + \zeta^2\sigma^3(\alpha) + \cdots + \zeta^{k-1}(\alpha) \\ \implies &\zeta\sigma\beta = \beta, \end{aligned}$$

because $\sigma(\zeta) = \zeta$, so that $\sigma(\beta) = \zeta^{-1}\beta$ for any α. This uses the concept of Lagrange resolvents, which will be discussed in more detail later on.

The above is useless if $\beta = 0$, but $\beta \neq 0$ by the independence of characters. Thus, β has k conjugates, so $L = F(\beta)$ by the primitive element theorem, and $\beta^k \in F$, so $L = F(\sqrt[k]{b})$ where $B = \beta^k$.

The flaw is in obtaining a formula. If $\mathrm{Gal}(E/F) = S_4 \rhd A_4 \rhd V_4 \rhd C_2 \rhd \{1\}$, then where is the formula? This proof isn't constructive. Another concern involves listing the roots of unity: they have nice nested radical formulas, but they get complicated quickly. However, they replace extensions such as $F(\zeta_5)/F$ with something like

$$F \text{——} F(\sqrt{5}) \text{——} F(\sqrt{-10-2\sqrt{5}}) \text{——} F(\sqrt{-10-2\sqrt{5}},\ \sqrt{-10+2\sqrt{5}}),$$

which is nicer in terms of finding a formula, almost aesthetically pleasing in how it builds up the roots of unity.

Exercise 18.1. If F is any field, p is prime, and $a \in F$, then show that either $x^p - a$ has a root in F or $x^p - a$ is irreducible in $F[x]$.

There are two cases, and note that if $\mathrm{Char}(F) = p$, then this reduces to one of the midterms: since $x^p - a = (x-\alpha)^p$, then things happen. Otherwise, $x^p - a = (x-\alpha)(x-\zeta\alpha)\cdots(x-\zeta^{p-1}\alpha)$; then, it's not necessary that $\alpha \in F$, but one of the $\zeta^j\alpha$ has to be.

One more fact about solvable groups over characteristic p: if one encounters a cyclic group of order p, the extension $F(\sqrt[p]{a})/F$ has a trivial Galois group, so it can't be the source. Instead, one has to use $x^p - x - a$, the Artin-Schreier equation, rather than $x^p - a$.

If K/F is Galois and $\mathrm{Gal}(K/F) = \{\sigma_1 = \mathrm{Id}, \ldots, \sigma_n\}$, then $\mathrm{Tr}(\alpha) = (\sigma_1 + \cdots + \sigma_n)\alpha \neq 0$ by the linear independence of characters, but $\mathrm{Tr}(1) = n$, which is a problem if $\mathrm{Char}(F) \mid n$. This is important in a full understanding of some of this stuff.

19. Two Applications of the Fundamental Theorem and Lagrange Resolvents: 2/22/13

Proposition 19.1. *Working in some algebraic closure Ω of a field F, suppose*

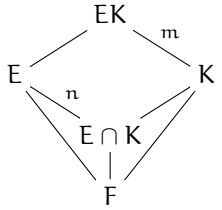

with E/F Galois and K/F arbitrary (not necessarily even algebraic). Then:

(1) EK/K *is Galois,*
(2) $\mathrm{Gal}(EK/K) \leqslant \mathrm{Gal}(E/F)$*, and*
(3) $\mathrm{Gal}(EK/K) = \mathrm{Gal}(E/E\cap K)$*. so* $|EK : K| = |E : E \cap K|$ *(and $n = m$ in the diagram above).*

Proof. 1 is trivial: it splits the same $f(x)$ that gives E/F. This kind of ignores separability, which is a little bit of a technicality. Every element of E is separable, because the generators are.

For 2, if $\sigma : EK \xrightarrow{\sim} EK$, then σ fixes K (and therefore F), so $\sigma|_E : E \to EK \subset \Omega$, but E/F is normal, so $\sigma(E) = E$. Thus, $\sigma \in \mathrm{Gal}(E/F)$. This is injective (which will imply inclusion) because is $\varphi|_K = \mathrm{Id}$ and $\varphi|_E = \mathrm{Id}$, then $\varphi = \mathrm{Id}$ on EK as well.

Item 3 is harder to prove than one might expect. Certainly, if $\varphi \in \mathrm{Gal}(EK/K)$, then φ fixes $E \cap K$, since it fixes K. If $\alpha \in E$ but $\alpha \notin K$, then E/K is Galois, so there exists a $\varphi : EK \xrightarrow{\sim} EK$ such that $\varphi|_K = \mathrm{Id}$ and $\varphi(\alpha \neq \alpha$ (since the only things fixed by all of the automorphisms are the elements of the ground field). Thus, $E \cap K$ is the fixed field of $\mathrm{Gal}(EK/K)$. ⊠

This is very false if E/F isn't Galois; some of it doesn't even make sense, but for the degrees one could have

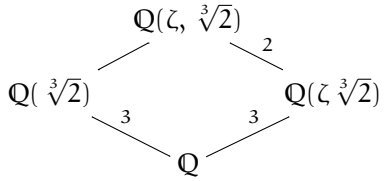

Here, the intersection is $\mathbb{Q}$, but $2 \neq 3$. In general, it is difficult to make statements about the intersection.

Suppose that

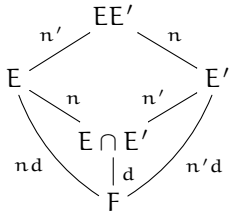

with both E/F and E'/F Galois. Then, EE' and $E \cap E'$ are both Galois, by a homework problem from earlier in the class, and if E splits f over F and E' splits g over F, then EE' splits fg.

Theorem 19.2. $\mathrm{Gal}(EE'/F) \leqslant \mathrm{Gal}(E/F) \times \mathrm{Gal}(E'/F)$, *and more precisely,* $\mathrm{Gal}(EE'/F) = \{(\varphi, \varphi') \mid \varphi|_{E\cap E'} = \varphi'|_{E\cap E'}\}$.

The proof of this theorem is in the book, and involves counting subgroups and quotient groups. However, the inclusion $\subseteq$ is clear. Here, it's worth asking: is it worth using the big guns of the Fundamental Theorem? It's not always obvious.

Switching topics, Lagrange resolvents can be used as an alternative to Hilbert's Theorem 90 and linear independence of characters. The goal is to show that if $\zeta_n \in F$ is an n^{th} root of unity and E/F is such that $\mathrm{Gal}(E/F) = C_n$, then $E = F(\beta)$ with $\beta^n = b \in F$. This is messy when characters get involved, and isn't a constructive proof.

Let α be a primitive element for E/F and let $\mathrm{Gal}(E/F) = \langle\sigma\rangle$. Then, the Lagrange resolvents are

$$\begin{aligned}
\rho_0 &= \alpha + \sigma\alpha + \sigma^2\alpha + \cdots + \sigma^{n-1}\alpha \\
\rho_1 &= \alpha + \zeta\sigma\alpha + \zeta^2\sigma^2\alpha + \cdots + \zeta^{n-1}\sigma^{n-1}\alpha \\
\ldots\ \rho_j &= \alpha + \zeta^j\sigma\alpha + \zeta^{2j}\sigma^2\alpha + \cdots + \zeta^{j(n-1)}\sigma^{n-1}\alpha,
\end{aligned}$$

all the way up to ρ_{n-1}. This can be interpreted as a matrix equation

$$\begin{pmatrix} 1 & 1 & 1 & \ldots & 1 \\ 1 & \zeta & \zeta^2 & \ldots & \zeta^{n-1} \\ 1 & \zeta^2 & \zeta^4 & \ldots & \zeta^{2n-2} \\ \vdots & \vdots & \vdots & \vdots & \vdots \end{pmatrix} \begin{pmatrix} \alpha \\ \sigma\alpha \\ \vdots \\ \sigma^{n-1}\alpha \end{pmatrix} = \begin{pmatrix} \rho_0 \\ \rho_1 \\ \vdots \\ \rho_{n-1} \end{pmatrix},$$

with this latter quantity not in F^n. Since $\zeta \in F$, then this matrix is over F, and it's a Vandermonde matrix, which implies that its determinant $\prod \zeta^i - \zeta^j$ is nonzero (since otherwise, $(\alpha, \ldots, \sigma^{n-1}\alpha) \in F^n$, which is nonsense). Thus, the matrix is invertible, so from linear algebra, there is some $\rho_j \notin F$. In particular, $\rho_j \neq 0$, which implies lots of good things (lots of conjugates, and so on). But this is also constructive!

More precisely, for each ρ_j, it's trivial to show that $\sigma(\rho_j) = \zeta^{-j}\rho_j$, so $\sigma(\rho_j^n) = \rho_j^n$, so $\rho_j^n \in F$. Thus, if $\rho_j \neq 0$, then it admits n conjugates, so $E = F(\rho_j)$. Of course, it's not clear which ρ_j is nonzero, but it's not hard to see that ρ_j^n, which is some big mess, is Galois-invariant, so it can actually be computed in terms of the coefficients of the minimal polynomial of α algorithmically.

This will lead to an explicit solution of the cubic, as it will make clear exactly which roots need to be adjoined.

20. Additive Hilbert's Theorem 90 and the Cubic: 2/25/13

Suppose $\zeta_n \in F$ and E/F is a degree-n field extension with cyclic Galois group $\mathrm{Gal}(E/F) = \langle\sigma\rangle$. Suppose $= F(\alpha)$, and consider the Lagrange resolvents

$$\begin{aligned}
\rho_0 &= \alpha + \sigma(\alpha) + \cdots + \sigma^{n-1}(\alpha) \\
\rho_1 &= \alpha + \zeta\sigma(\alpha) + \cdots + \zeta^{n-1}\sigma^{n-1}(\alpha) \\
&\vdots \\
\rho_j &= \alpha + \zeta^j\sigma(\alpha) + \cdots + \zeta^{j(n-1)}\sigma^{n-1}(\alpha)
\end{aligned}$$

all the way up to ρ_{n-1}. Thus, one has the rho vector[26]

$$\begin{pmatrix} \rho_0 \\ \vdots \\ \rho_{n-1} \end{pmatrix} = A \begin{pmatrix} \alpha \\ \sigma\alpha \\ \vdots \\ \sigma^{n-1}\alpha \end{pmatrix},$$

where A is Vandermonde in ζ^{jk}. Thus, its entries are in F and it has a nonzero determinant, so there exists some $\rho_j \notin F$. We also have $\sigma(\rho_j) = \zeta^j\rho_j$.

Here take care not to be too quick on the trigger: what happens when $j \mid n$? If n is prime, then ρ_j has n conjugates, so $E = F(\rho_j)$ and $\rho_j^n \in F$, but otherwise, this might not happen. The overall logic for the proof of solvability by radicals doesn't change, though: there is still a tower of n^{th} root ajunctions, but it happens in more than one step.

The proof of Hilbert's Theorem 90 required some strange and messy formulas, so here's a better option:

Theorem 20.1 (Additive Hilbert's Theorem 90)**.** *Suppose* E/F *has degree* n *and a cyclic Galois group* $\mathrm{Gal}(E/F) = \langle\sigma\rangle$*. Then, if* $T_{E/F}(\alpha) = 0$ *for* $\alpha \in E$*, then* $\alpha = \sigma\beta - \beta$ *for some* $\beta \in E$*.*

Proof. Recall that the trace is F-linear and is not identically zero (because of the linear independence of characters). Then, $\dim(\ker(T_{E/F})) = n - 1$, because it's onto F. The minimal polynomial is Galois-invariant, since $(\mathrm{Id} + \sigma + \cdots + \sigma^{n-1})(\sigma - \mathrm{Id}) = \sigma^n - \mathrm{Id} = 0$.

The reverse direction is trivial: just compute $T_{E/F}(\sigma\beta - \beta)$, so consider the forward direction and suppose $T_{E/F}(\alpha) = 0$. Since $\ker(\sigma - \mathrm{Id}) = F$, then $\dim(\mathrm{Im}(\sigma - 1)) = n - 1$, which is the same as the dimension of the trace, and one is contained within the other, so they are equal. ⊠

In particular, if $\mathrm{Char}(F) = p$ and $|E : F| = p$, then $T_{E/F}(1) = 0$, so $E = F(\gamma)$, where γ is a root of $x^p - x - a$ for some $a \in F$ and $\sigma(\gamma) = \gamma + 1$. This is where the Artin-Schreier extensions come into play.

Suppose $\mathrm{Char}(F) \neq 2, 3$ and E/F splits some irreducible, separable cubic $f(x) = x^3 + ax^2 + bx + c \in F[x]$. Then, replace x with $x + a/3$ to obtain some $x^3 + px + q$ for p, q that depend on a, b, c. Thus, it suffices to obtain the roots of this polynomial, unhappily referred to as the "depressed cubic."

It turns out this isn't that hard, thanks to a trick discovered in the 16^{th} Century, but using the Lagrange resolvent will be more in line with the theory. Let α_1, α_2, and α_3 be the roots, and let $D = ((\alpha_1 - \alpha_2)(\alpha_1 - \alpha_3)(\alpha_2 - \alpha_3))^2$ be the discriminant (i.e. it is S_3-invariant).

Definition. Suppose f is a separable, n^{th}-degree polynomial with roots $\alpha_1, \ldots, \alpha_n$. Then, its discriminant is

$$D = \prod_{1 \leqslant i < j \leqslant n} (\alpha_i - \alpha_j)^2,$$

which is invariant under permutation by S_n.

Since it's invariant under action by S_n, the discriminant is always in the ground field, and in fact can be thought of as a function in the coefficients of f.

It's a result of symmetric function theory that any symmetric function of n variables (i.e. a function invariant under S_n; specifically, the same expression is given, not just the same numbers) is a polynomial function. There's a Galois-theoretical proof of this, though it has some thorniness in terms of where the coefficients go.

In particular, $\sqrt{D} = \prod_{1 \leqslant i < j \leqslant n}(\alpha_i - \alpha_j)$ is A_n-invariant.[27] Thus,

Corollary 20.2. $\sqrt{D} \in F$ *iff* $\mathrm{Gal}(E/F) \leqslant A_n$.

[26] Or row vector. Ha ha.

[27] This fact is used as the basis of some definitions of A_n.

Returning to the cubic, $\alpha_1+\alpha_2+\alpha_3=0$, since it comes from the x^2-coefficient, which has been excised. Expanding this out is a mess, but becomes nicer in terms of the elementary symmetric functions. Thus, the discriminait is $D=-4a^3c+a^2b^2+18abc-4b^3-27c^2$. Thus, $a=\alpha_1+\alpha_2+\alpha_3$, $b=\alpha_1\alpha_2+\alpha_1\alpha_3+\alpha_2\alpha_3$, and $c=\alpha_1\alpha_2\alpha_3$. For x^2+px+q, this reduces to $D=-4p^3-27q^2$. Thus, $\mathrm{Gal}(E/F)=A_3<S_3$ iff D is a square in F.

To solve the cubic, take

$$F \overset{1 \text{ or } 2}{=\!=} F(\sqrt{D}) \overset{3}{\_\_\_} F(\sqrt{D},\alpha_1) = E \overset{}{\_\_\_} F(\sqrt{D},\alpha_1,\sqrt{-3})$$

For E, any root will generate the whole extension, since 3 is prime, so nothing can be in between. The last extension is made for the purpose of writing nested radicals.

21. Solving the Cubic: 2/27/13

Suppose $\mathrm{Char}(F)\neq 2,3$ and $f(x)=x^3+x+q$ is separable and irreducible in $F[x]$. We will take the extension $F \overset{}{\_\_\_} F(\zeta_3,\sqrt{D}) \overset{3}{\_\_\_} E$, where E is the splitting field of $f(x)$ over $F(\zeta_3,\sqrt{D})$. Then, $\mathrm{Gal}(E/F(\zeta_3,\sqrt{D}))=C_3$.

Let r_1,r_2,r_3 be the roots of f. Then, $r_1+r_2+r_3=0$, since this is the x^2-term, and the Lagrange resolvents are $\rho_0=0$, $\rho_1=r_1+\zeta r_2+\zeta^2 r_3$, and $\rho_2=r_1+\zeta^2 r_2+\zeta r_3$. Then, $\sigma(\rho_1)=\zeta^2\rho_1$, and $\sigma(\rho_2)=\zeta\rho_2$. These equations can now be solved for the roots, which is just linear algebra, and we also have that $1+\zeta+\zeta^2=0$ and that the determinant is nonzero (since it's the Vandermonde determinant), implying that $3r_1=\rho_1+\rho_2$, $3r_2=\zeta^2\rho_1+\zeta\rho_2$, and $3r_3=\zeta\rho_1+\zeta^2\rho_2$. Thus, $\rho_1^3,\rho_2^3\in F(\zeta,\sqrt{D})$.

After a bunch of calculation,

$$\begin{aligned}
\rho_1^3&=(r_1+\zeta r_2+\zeta^2 r_3)^3=r_1^3+r_2^3+r_3^3+3\zeta(r_1^2r_2+r_2^2r_3+r_3^2r_1)+3\zeta^2(r_1r_2^2+r_2r_3^2+r_3r_1^2)+6r_1r_2r_3\\
\rho_2^3&=r_1^3+r_2^3+r_3^3+3\zeta^2(r_1^2r_2+r_2^2r_3+r_3^2r_1)+3\zeta(r_1r_2^2+r_2r_3^2+r_3r_1^2)+6r_1r_2r_3\\
\sqrt{D}&=(r_2-r_1)(r_3-r_1)(r_3-r_2)=r_1r_2^2+r_2r_3^2+r_3r_1^2-r_1^2r_2-r_2^2r_3-r_3^2r_1.
\end{aligned}$$

Thus, $\rho_1^3-\rho_2^3=3(\zeta^2-\zeta)\sqrt{D}$ and $\rho_1^3+\rho_2^3=27r_1r-2r_3=27pq$. Solving explicitly, one obtains using (1), below, that

$$\begin{aligned}
\rho_1^3&=\frac{1}{2}(27q+3(\zeta^2-\zeta)\sqrt{D})\\
\rho_2^3&=\frac{1}{2}(27q-3(\zeta^2-\zeta)\sqrt{D}),
\end{aligned}$$

thus yielding the formulas $3r_1=\rho_1+\rho_2$ and

$$\begin{aligned}
r_1&=\frac{1}{3}\left(\sqrt[3]{\frac{1}{2}\left(27q+3(\zeta^2-\zeta)\sqrt{D}\right)}+\sqrt[3]{\frac{1}{2}(27q-3(\zeta^2-\zeta)\sqrt{D})}\right)\\
r_2&=\frac{1}{3}\left(\sqrt[3]{\frac{27q}{3}-\frac{3\sqrt{-3}}{2}\sqrt{-4p^3-27q^2}}+\sqrt[3]{\frac{27q}{2}+\frac{3\sqrt{-3}}{2}\sqrt{-4p^3-27q^2}}\right),
\end{aligned}$$

and so on. But this gives 9 different choices for roots, which is too many, so it's necessary to also specify that

$$\begin{aligned}
\rho_1\rho_2&=r_1^2+r_2^2+r_3^3+(\zeta+\zeta^2)(r_1r_2+r_2r_3+r_1r_3)\\
&=r_1^2+r_2^2+r_3^3-r_1r_2-r_2r_3-r_1r_3\\
&=(r_1+r_2+r_3)^2-3(r_1r_2+r_1r_3+r_2r_3)=-3p,
\end{aligned}$$

or just that $\rho_1\rho_2\in F$. The choices of the roots aren't independent: picking one forces the rest. For example, for x^3-q, $\rho_1\rho_2=0$, which makes life interesting, but in cases like this it's not hard to find the roots.

Notice that there isn't anything too difficult or even theoretical about all of this; it's just computation and a bunch of things about symmetric functions.

The calculation that justifies the above is:

$$\begin{aligned}
\rho_1^3+\rho_2^3&=2(r_1^3+r_2^3+r_3^3)+12r_1r_2r_3-3(r_1^2r_2+r_2^2r_3+r_3^2r_1+r_1r_2^2+r_2r_3^2+r_3r_1^2) \qquad (1)\\
&=2(r_1+r_2+r_3)^2-9(r_1+r_2+r_3)(r_1r_2+r_2r_3+r_3r_1)+27r_1r_2r_3.
\end{aligned}$$

Though this seems uninteresting, there's a bit of history here; many of the mathematicians that discovered these formulas kept them secretively and held contests. Ferrari met a premature death, poisoned by his sister. Cardan was an astrologer, a drunkard, and a gambler, who cut off his son's ear because he displeased him.

One final question: why do these formulas have imaginary quantities even though some cubics have all real roots? There's a theorem that says that there is no iterated radical formula in general, even if the cubic has three real

roots. This made these formulas more mysterious, especially given that these guys didn't really believe in inaginary numbers, and barely trusted negative ones.

Suppose $f(x) = \prod_1^n (x - r_i)$, so that f is a monic, degree-n polynomial, and $f'(x) = n\prod_{i=1}^{n-1}(x - s_i)$. By the Product Rule,

$$f'(x) = \sum_{i=1}^{n} \frac{(x-x_1)\cdots(x-x_n)}{x-x_i} \implies f'(r_i) = \prod_{j\neq i}(r_i - r_j).$$

Thus, one gets some more formulas for the discriminant:

$$\begin{aligned}\prod_{i=1}^{n} f'(r_i) &= (-1)^{n(n-1)/2} \prod_{1\leqslant i<j\leqslant n} (r_i - r_j)^2 = (-1)^{n(n-1)/2} D \\ &= n^n \prod_{i,j=1}^{n} (r_i - s_j) = \prod f(s_j) \\ \implies D &= (-1)^{n(n-1)/2} \prod_{i=1}^{n} f'(r_i) = (-1)^{n(n-1)/2} n^n \prod_{i=1}^{n-1} f(s_i).\end{aligned}$$

In general, this is a bad way to calculate discriminants, but for cubics it can be helpful.

22. The Fundamental Theorem of Symmetric Polynomials: 3/1/13

Let R be a commutative ring. Then, S_n acts on $R[x_1, \ldots, x_n]$ by permuting the x_i. Let s_j be the j^{th} elementary symmetric function, defined by

$$\prod_{i=1}^{n}(T - x_i) = \sum_{j=1}^{n} (-1)^j s_j T^{n-j},$$

such as $s_1 = \sum x_i$, $s_2 = \sum_{i<j} x_i x_j$, and so on to $s_n = \prod x_i$.

Today's theorem, stated below, isn't terribly sophisticated. In some sense, it's just organized computation, and is almost at a high school level.

Theorem 22.1 (Fundamental Theorem of Symmetric Polynomials)**.** *G $R[x_1, \ldots, x_n]^{S_n} = R[s_1, \ldots, s_n]$, where the set on the left is the subset of $R[x_1, \ldots, x_n]$ that is S_n-invariant.*

Proof. $R[x_1, \ldots, x_n]$ is a graded ring:[28] $R[x_1, \ldots, x_n] = \bigoplus_{d=1}^{\infty} R_d$, where $R_d$is the R-span of monomials of degree d. Here, the degree of a polynomial in n variables is given by $\deg(x_{i_1}^{a_1} + \cdots + x_{i_n}^{a_n}) = \sum a_i$. Thus, $\deg(s_1^{b_1} + \cdots + s_n^{b_n}) = b_1 + 2b_2 + \cdots + nb_n$. Using this graded ring structure, once we know that $R_d^{S_n} = S_d$, then the proof will be done.

Use the lexicographic ordering of multi-exponents: if $a_1 \geqslant a_2 \geqslant \cdots \geqslant a_n \geqslant 0$, then define $(a_1, \ldots, a_n) \geqslant (a_1', \ldots, a_n')$ if $a_1 > a_1'$ or $a_1 = a_1'$ and $a_2 > a_2'$, or $a_1 = a_1'$, $a_2 = a_2'$, and $a_3 > a_3'$, and so on. Then, if $P(x_1, \ldots, x_n) \in R_d^{S_n}$, so that P is a symmetric polynomial of degree d, and the term $x_{i_1}^{a_1} x_{i_2}^{a_2} \cdots x_{i_n}^{a_n}$ appears in P, then so must $x_1^{a_1} x_2^{a_2} \cdots x_n^{a_n}$, applying the permutation $\sigma(j) = i_j$. Then, the *weight* of P is defined to be the highest such $(a_1, \ldots, a_n)$ that has a nonzero coefficient.

In the polynomial $s_1^{b_1} s_2^{b_2} \cdots s_n^{b_n}$, one can expand out the symmetric polynomials and see that the highest power of x_1 that occurs is $\sum b_i$. Thus, if the weight is $(a_1, \ldots, a_n)$, then fix x_1, so that $a_1 = \sum b_i$, and thus $a_2 = b_2 + b_3 + \cdots + b_n$.

If $P = Ax_1^{a_1} x_2^{a_2} \cdots x_n^{a_n}$, then consider the polynomial $P - As_1^{a_1 - a_2} s_2^{a_2 - a_3} \cdots s_{n-1}^{a_{n-1} - a_n} s_n^{a_n}$. This removes the term of highest weight, and the degree stays the same. Thus, this process goes to zero in a finite number of steps, so P decomposes into a linear combination of symmetric polynomials. ⊠

This proof also shows that there is a unique formula for P as a polnomial in $R[s_1, \ldots, s_n]$, and that they are algebraically independent over R.

Thus, if an expression is formally S_n-symmetric (i.e. there are no accidental relations), then it is a polynomial in the elementary symmetric functions. One can extract an algorithm from the proof, since it forces the reduction of the weight.

There is another theoretical view of symmetric functions from Galois theory. If E/F is a finite, normal, separable extention, one can rewrite the normal Galois-theoretic setup using symmetric functions. Instead of going from the extention to $\mathrm{Gal}(E/F)$ and such, one can start with a field E and a finite group $H < \mathrm{Aut}(E)$, and consider $E^H \subset E$.

[28] Yep, high school level right here.

Theorem 22.1 ensures this is Galois, with $|E^H : E| > |H|$ and $\mathrm{Gal}(E/E^H) = H$. In particular, the latter is finite.[29] In some sense, one starts with the group of automorphisms rather than the field.

If $\theta \in E$, then θ is the root of an $f \in E^H[x]$ such that $\deg(f) = |H|$, and in fact θ is a root of a separable $g(x)$ that splits linearly, and with $\deg(g) \mid |H|$. Then, as shown before (the same argument, but in a different proof),

$$f(x) = \prod_{\sigma \in H} (x - \sigma\theta) \qquad \text{and} \qquad g(x) = \prod_{\substack{\text{cosets of} \\ I = \mathrm{Stab}(\theta)}} (x - (\sigma_j I)\theta),$$

which is just a statement about group actions, since $\gamma \in I$ iff $\gamma\theta = \theta$.

The aforementioned cosets are those of $H = \bigcup_j \sigma_j I$, and in the product, one factor is chosen from each coset. This eliminates the repetition. g is still symmetric, since H acts on the cosets, so $\sigma(\sigma_j H) = \sigma_\ell H$, similarly to before. This proves part of the theorem: it explains why E/E^H is normal and separable, but think about why it must be finite.

One can show a similar result to Theorem 22.1 for symmetric rational functions: $F(x_1, \dots, x_n)^{S_n} = F(s_1, \dots, s_n)$. The degree is $n!$, because

$$F(s_1, \dots, s_n) \xrightarrow{1} F(x_1, \dots, x_n)^{S_n} \xrightarrow{n!} F(x_1, \dots, x_n)$$

so $d = n!$. This is another way to illustrate that symmetric rational functions can be rewritten in terms of elementary symmetric functions. However, though it's a more abstract and sophisticated result, it's actually weaker.

23. The Fundamental Theorem of Algebra: 3/4/13

Theorem 23.1 (Fundamental Theorem of Algebra)**.** $\mathbb{C}$ *is algebraically closed.*

There are complex-analytic proofs of this theorem (showing that a bounded entire function is constant), and topological ones (e.g. from the perspective of a sufficiently large circle C_R with radius R, $P(z) \approx z^n$, with small error relative to R, which induces a map $C_R \to C_R$ with nice properties). But is there a purely algebraic proof? Not really,[30] since $\mathbb{C}$ is defined in a somewhat analytic mannerr. But the proof given below will be one of the least analytic or topological ones that can be made.

Proof of Theorem 23.1. Every odd-degree polynomial of $\mathbb{R}$ has a root, by the Intermediate Value Theorem, and every positive real number has a square root. Thus, complex numbers have square roots as well: if $u + iv = (c + id)^2$, then $u = c^2 - d^2$ and $v = 2cd$, so $c^2 + d^2 = \sqrt{u^2 + v^2}$ in $\mathbb{R}$ (since these quantities are both nonnegative), and an explicit solution can be given:

$$c = \frac{\sqrt{u + \sqrt{u^2 + v^2}}}{2} \qquad \text{and} \qquad d = \frac{\sqrt{\sqrt{u^2 + v^2} - u}}{2}.$$

The square roots make sense because $u \leqslant \sqrt{u^2 + v^2}$, since $u \in \mathbb{R}$.[31]

First, there are no odd-degree extenstions $K/\mathbb{R}$ (well, when $|K : \mathbb{R}| > 1$). This is because if such a K existed, then any $\alpha \in K$ would satisfy an irreducible polynomial of odd degree, but as seen above, this doesn't happen.

Additionally, there aren't any extensions of $\mathbb{C}$. Suppose $\mathbb{R} \text{——} \mathbb{C} \text{——} E$, and without loss of generality assume $E/\mathbb{R}$ is Galois, since if not then one could go up to the minimal Galois closure. Let $G = \mathrm{Gal}(E/\mathbb{R})$ and H be a Sylow-2 subgroup of G. Then, $|E^H : \mathbb{R}|$ must be odd, so $E^H = \mathbb{R}$, and $G = H$.

This means that G is a 2-group, so $|G| = 2^N$, and $|\mathrm{Gal}(E/\mathbb{C})| = 2^{N-1}$, which means that it is also a p-group. Thus, there is a normal subgroup $C_2 \trianglelefteq \mathrm{Gal}(E/\mathbb{C})$ and a surjective homomorphism $\mathrm{Gal}(E/\mathbb{C}) \to C_2$. The kernel P' is normal and has index 2, which means that it is a quadratic extension. Thus, $N = 1$, so $E = \mathbb{C}$ and nothing interesting can happen. ⊠

The last steps follow from the theory of p-groups, or just the simple fact that no p-group has a trivial center. It can also be given in terms of the structure theorem of finite abelian groups and some induction. There's another nice algebraic proof due to Laplace (in about 1795) that reduces the question to one about $\mathbb{R}[x]$, since if $f \in \mathbb{C}[x]$, then $f\overline{f} \in \mathbb{R}[x]$.

Since $\mathbb{C}$ is algebraically closed, then it has lots of algebraic subfields, such as $\overline{\mathbb{Q}}$ (the complex algebraic numbers), $\overline{\mathbb{Q}(\pi)}$, $\overline{\mathbb{Q}(e)}$, and so on.

The subject of the next few days will be the Cycle Type theorem, which allows for easier computation of Galois groups. Suppose $f \in \mathbb{Z}[x]$ is monic and separable, and suppose p is prime and that $\bar{f}(x) = f(x) \bmod p$ is also separable.

[29] This wasn't stated rigorously, and really should have a proof to go along with it. The textbook does, though.

[30] Apparently this greatly bothered Serge Lang.

[31] Of course, there are other ways to prove that complex numbers have square roots. For example, if $z = r(\cos\theta + i\sin\theta)$, then $\sqrt{z} = \sqrt{r}(\cos(\theta/2) + i\sin(\theta/2))$, but this isn't as algebraic as the proof given above.

Then, $\mathrm{Gal}(\overline{f}) \leqslant \mathrm{Gal}(f)$ (up to isomorphism), which isn't terribly interesting because $\mathrm{Gal}(\overline{f})$ is cyclic. However, more interesting things happen if one factors $\overline{f}(x) = \overline{f_1}(x)\overline{f_2}(x)\cdots\overline{f_r}(x)$ as distinct irreducibles in $\mathbb{Z}/p[x]$ with degrees $\deg(\overline{f_i}(x)) = n_i$. Then, if E is the splitting field of $\overline{f}(x)$, then $|E : \mathbb{F}_p| = \mathrm{lcm}(n_1, \ldots, n_r)$, since there isonly one finite field of a given order up to isomorphism and each prime divisor is factored in exactly once.

Thus, if $n = \deg(f) = \deg(\overline{f})$, then $\mathrm{Gal}(\overline{f}) \leqslant S_n$, but also that $\mathrm{Gal}(\overline{f}) \leqslant \prod_{i=1}^{r} S_{n_i}$. Additionally, $\mathrm{Gal}(\overline{f_j}) \leqslant S_n$ is transitive for each j. If $\mathrm{Gal}(\overline{f}) = \langle\sigma\rangle$ for some $\sigma \in S_n$, then $\sigma = \sigma_1 \cdots \sigma_n$, with $\sigma_j \in S_{n_j}$, which gives a nice product of cycles, and by transitivity, each σ_j is a single cycle.

Theorem 23.2 (Cycle Type)**.** *With the assumptions and notation as above,* $\mathrm{Gal}(f)$ *contains an element which is a product of* $n_1, \ldots, n_r$*-cycles.*

This will be useful for placing restrictions on $\mathrm{Gal}(f)$, but its proof must be deferred to a future lecture.

24. Computation of Galois Groups I: 3/6/13

Starting with a little review, suppose $\mathbb{F}$ is any finite field and $f \in \mathbb{F}[x]$ factors into distict irreducibles $f = f_1 \ldots f_r$, and let $n_i = \deg(f_i)$. Then, $\mathrm{Gal}(f)$ is cyclic, as discussed previously, with order $\mathrm{lcm}(n_1, \ldots, n_r)$. Thus, $\mathrm{Gal}(f)$ is generated by a permutation $\sigma_1 \ldots \sigma_r \in S_n$, where σ_j is an n_j-cycle and all of the σ_j are disjoint, since σ_j permutes the roots of f_j (which is enough to imply the order, too).

The goal of these two lectures is to show that if an $f \in \mathbb{Z}[x]$ is monic and separable of degree n, then $\mathrm{Gal}(f)$ contains a permutation of cycle type $n_1, \ldots, n_r$ (where $f \bmod p$ factors as $f_1, \ldots, f_r$ in $\mathbb{F}_p[x]$ as above). This obviously says a lot about the Galois group; for example, if f is irreducible in any $\mathbb{F}_p[x]$, then it is also irreducible in $\mathbb{Z}[x]$. This result will lead to probabilistic estimates for Galois groups and even algorithms.

However, along the way, it will be necessary to prove a completely useless theorem — at least from a computational point of view. This seems paradoxical, but see that it will be helpful.

Theorem 24.1. *Suppose* f *is as above; then, there exists a* $P \in \mathbb{Z}[u_1, \ldots, u_n, x]$*, where the* u_j *are algebraically and linearly independent (i.e. if* F *is a field, then* $F(u_1, \ldots, u_n)$ *is a purely transcendental extension of* F*) such that* P *is symmetric in the* u_j*-terms and* $\deg(P) = n!$ *Then, if* P *factors as* $P = P_1 \ldots P_n$ *with the* P_j *irreducible, then* $\mathrm{Gal}(f) = \mathrm{Stab}(P_1) \leqslant S_n$*.*

Computationally, it's hard enough to factor polynomials of degree n in one variable, let alone polynomials of degree $n!$ in $n+1$ variables. The professor actually tried doing this with a quadratic once. However, it's good for the mod$p$ reduction, since a similar setup will occur and the things that fix P_1 in $\mathbb{F}_p$ permute its roots, leading to more useful results.

Proof of Theorem 24.1. Let F be any field and $f \in F[x]$ be a degree-n separable polynomial with roots $\alpha_1, \ldots, \alpha_n$ in a splitting field E of F. Let $F(u_1, \ldots, u_n)$ be a purely transcendental extension of F, so that the u_j are algebraically and linearly independent. Working in some algebraic closure of $F(u_1, \ldots, u_n)$, consider

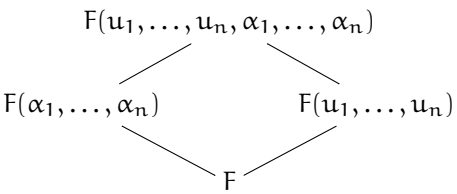

The field $F(u_1, \ldots, u_n, \alpha_1, \ldots, \alpha_n)$ is the splitting field of f over $F(u_1, \ldots, u_n)$. Since $F(\alpha_1, \ldots, \alpha_n)$ is Galois, with Galois group G, then $F(u_1, \ldots, u_n, \alpha_1, \ldots, \alpha_n)/F(u_1, \ldots, u_n)$ is also Galois with the same Galois group, since $F(\alpha_1, \ldots, \alpha_n) \cap F(u_1, \ldots, u_n) = F$. Then, S_n acts on $F(u_1, \ldots, u_n, \alpha_1, \ldots, \alpha_n)$ by $u_i \mapsto u_{\sigma(i)}$ and leaving the α_i alone:

$$\sigma\left(\sum_{i=1}^{n} u_i\alpha_i\right) = \sum_{i=1}^{n} u_{\sigma(i)}\alpha_i = \sum_{i=1}^{n} u_i\alpha_{\sigma^{-1}(i)}. \tag{2}$$

Since $F(u_1, \ldots, u_n, \alpha_1, \ldots, \alpha_n)$ is a purely transcendental extension of $F(\alpha_1, \ldots, \alpha_n)$, then all of the $n!$ elements given in (2) are distinct; if not, there would be some relations among the u_i.

Thus, there exists a primitive element of $F(u_1, \ldots, u_n, \alpha_1, \ldots, \alpha_n)$ over $F(u_1, \ldots, u_n)$, which will be some sort of u-linear combination of the α_i. Guessing $\sum u_i\alpha_i$, its conjugates are given by

$$\sum_{i=1}^{n} u_i\sigma(\alpha_i) = \sum_{i=1}^{n} u_i\alpha_{\sigma(i)},$$

with $\sigma \in G = \mathrm{Gal}(f)$. Thus, there are $|G|$ distinct conjugates, so $\sum u_i\alpha_i$ is a primitive element.

Let P be given by

$$P = \prod_{\sigma \in S_n} \left(x - \sum_{i=1}^{n} u_{\sigma(i)} \alpha_i \right) = \prod_{\sigma \in S_n} \left(x - \sum_{i=1}^{n} u_i \alpha_{\sigma(i)} \right).^{32}$$

Thus, P is symmetric in the α_j, so $P \in F[u_1, \dots, u_n, x]$. Additionally, $\sigma P = P$ for any $\sigma \in S_n$, *with* σ *acting on* $u_1, \dots, u_n$. Since P has $n!$ distinct roots, then P is separable.

Factor $P = P_1 \dots P_s$, with each factor irreducible and such that $\sum u_i \alpha_i$ is a root of $P_1(x)$. Then, P_1 is the minimal polynomial for $\sum u_i \alpha_i$ over $F(u_1, \dots, u_n)$ and

$$P_1(x) = \prod_{\sigma \in \mathrm{Gal}(f)} \left(x - \sum_{i=1}^{n} u_i \alpha_{\sigma(i)} \right). \tag{3}$$

Let S_n act on $u_1, \dots, u_n$. Since P is S_n-invariant, then S_n permutes the factors $P_1, \dots, P_s \dots$

The proof is almost done here, and will be continued next lecture.

25. Computation of Galois Groups II: 3/8/13

Continuing from above, this means that

$$\mathrm{Gal}(f) = \{\sigma \in S_n \mid \sigma P_1(x) = P_1(x) \text{ with } \sigma \text{ acting on } u_1, \dots, u_n\}.$$

This is precisely the stabilizer, because if $\sigma \in \mathrm{Gal}(f)$, then σ sends roots of P_1 to roots of P_1 and roots of some other P_j ($j \neq 1$) to roots of P_j; thus, $\mathrm{Gal}(f) = \mathrm{Stab}(P_1)$. ⊠

This is theoretically computable, but requires a lot of work in computing the $n!$ permutations of $u_1, \dots, u_n$. It works over any field F, however.

Suppose $f \in \mathbb{Z}[x]$ and $\overline{f} = f \bmod p$, and $\overline{f} = \overline{f}_1 \dots \overline{f}_m$ is its factorization into distinct irreducible factorization in $\mathbb{F}_p[x]$. Then, taking P as in (3), $P \in \mathbb{Z}[u_1, \dots, u_n, x]$. Suppose $\overline{P}$ is the reduction of P in $\mathbb{F}_p[x]$ and each of its factors P_i factors as $\overline{P} = \overline{P}_{i1} \dots \overline{P}_{is_i}$. Thus, $\mathrm{Gal}(\overline{f})$ is equal to the S_n-stabilizer of $(\overline{P}_{11}(x))$, which is contained within $\mathrm{Stab}(P_1(x)) = \mathrm{Gal}(f)$. This is because the only way for the roots to be stabilized downstairs in an extension of $\mathbb{F}_p$ is for them to be stabilized upstairs.

This (slightly magical) statement is true when viewed as for permutation groups of the roots $\alpha_1, \dots, \alpha_n$ or their reductions in some extension of $\mathbb{F}_p$. Additionally, the permutations that fix P_1 when S_n acts on $u_1, \dots, u_n$ are those that fix P_1 when S_n acts on $\alpha_1, \dots, \alpha_n$.

It turns out there are relations between the $\alpha_1, \dots, \alpha_n$ and their reductions over $\mathbb{F}_p$, but that strays into algebraic number theory. $\mathbb{Z}[\alpha_1, \dots, \alpha_n] \subset \mathbb{C}$ is a subring, and if Q is one of its prime ideals, then $\mathbb{Z}[\alpha_1, \dots, \alpha_n]/Q$ is a finite field, as in the following diagram:

$$\begin{array}{ccc} \mathbb{Z} & \longrightarrow & \mathbb{Z}[\alpha_1, \dots, \alpha_n] \\ \downarrow & & \downarrow \\ \mathbb{Z}/p & \xrightarrow{\text{finite}} & \mathbb{Z}[\alpha_1, \dots, \alpha_n]/Q \end{array} \tag{4}$$

This is a nicer way to view these numbers, rather than as matrices or soething more cumbersome. However, it takes some of the magic away, since it explains why the connection exists.

In the case of the finite extension, $\mathrm{Gal}(\overline{f})$ is cyclic, and it has as a generator $\sigma(x) = x^p$.[33] Then, $\mathrm{Gal}(\overline{f}) \leqslant \mathrm{Gal}(f)$ up to isomorphism, though it turns out that uniqueness is only defined up to conjugacy. The Frobenius automorphism is canonical, but choosing $\mathrm{Gal}(f) \leqslant S_n$ requires choosing an ordering of the roots, and a change in labels is a change in conjugacy and causes a sort of ambiguity.[34]

Looking at (4) more explicitly, think of the elements of $\mathbb{Z}/p$ as elements of $\mathbb{Z}$ modded out by a prime idea $p\mathbb{Z}$. Similarly, one can think of a finite field extension as the ring $\mathbb{Z}[\alpha_1, \dots, \alpha_n]$ modded out by some prime ideal Q. In fact, there will be prime ideals such that $Q \cap \mathbb{Z} = (p)$ for some prime p, so $\mathbb{Z}[\alpha_1, \dots, \alpha_n]/Q$ is a finite field, and is isomorphic to the splitting field of $\overline{f}$ over $\mathbb{F}_p$. This is more appealing than viewing finite field extensions as matrices or as equivalence classes of polynomials.

Starting from here, there is more than three lectures' worth of review. These notes are already as late as they are, and so I will omit the old material. If you have some desire to look through it, please let me know!

[32] The placement of $\sigma(i)$ and $\sigma^{-1}(i)$ may seem confusing, but everything is summed over, so this does make sense.

[33] This is the Frobenius automorphism, which is how all of this was shown in the first place.

[34] This conjugacy is within $\mathrm{Gal}(f)$, since the relationships are just parameters given by $\sigma \in \mathrm{Gal}(f)$.

Appendix A. Professor Quotes

Professor Brumfiel is prone to saying funny things during lecture. Here is a list; additional suggestions are welcome. Some of these may also appear in the body of the lecture notes.

(1) "What *is* $\mathfrak{x}$, really?"
(2) "Whoa — what if $\mathfrak{h}$ isn't reducible?"
(3) "If some alien comes down and shows us a splitting field, and we say 'we have one too,' and he says, 'no, it's different,' well, not really."
(4) "What's my hypothesis?"
(5) "I apologize, but I don't really. . . "
(6) "Well, here's sort of a 'proof by magic.' "
(7) "The answer is yes, but if F is 'big' in the set-theoretic sense, you need fancy set theory."
(8) "These are seeds. Then I will begin to quickly reap results directly relevant to Galois Theory."
(9) "My guess is, Math 121 is something like three times as hard as Math 120. I'm not sure what this means, maybe that it takes three times as much time and energy. Or maybe that only about one third of the students who complete Math 120 successfully will be comfortable in Math 121."
(10) "So there are some set-theoretic issues with that approach: if you have $\sqrt{2} \in \mathsf{F}'$ and a '$\mathtt{pig}' \in \mathsf{F}''$ such that $\mathtt{pig}^2 = 2 \in \mathsf{F}'$, then you have a problem."
(11) "That's the correct thing to say. It's not the correct answer, but it's the correct thing to say."
(12) "You're referring to these roots which are in some magic world somewhere."
(13) "There's too many $\mathfrak{n}$s here. . . let's just erase that one."
(14) "That's Theorem whatchamacallit. . . 27."
(15) "So the elements of K might be pigs and cows and sheep rather than complex numbers."
(16) "The Fundamental Theorem of Galois Theory is actually pretty easy. . . it just drops out like melted butter."
(17) "Have a nice weekend studying Galois Theory."
(18) "This is true because duh: is L fixed by automorphisms that fix L? Do bears crap in the woods?"
(19) "I promised myself I wouldn't run over [time] but I always lie."
(20) "I'm not going to [prove] this, since it always makes my head hurt."
(21) "Using the Fundamental Theorem of Galois Theory here is like killing a fly with a hand grenade."
(22) "I don't have the energy to solve quartics. . . too much elbow grease."
(23) "What would you do if $\mathrm{Char}(\mathsf{F}) = 2$?" "I forget."