Galois Theory *

Zhiyuan Bai

Compiled on June 4, 2022

This document serves as a set of revision materials for the Cambridge Mathematical Tripos Part II course *Galois Theory* in Michaelmas 2021. However, despite its primary focus, readers should note that it is NOT a verbatim recall of the lectures, since the author might have made further amendments in the content. Therefore, there should always be provisions for errors and typos while this material is being used.

Contents

0 Notation and Revision **2**

1 Field Extensions **3**
1.1 Finite and Algebraic Extensions 3
1.2 Construction with Straightedge and Compasses 6

2 Splitting Fields **7**

3 Finite Fields and Separability **9**
3.1 Existence and Uniqueness of Finite Fields 9
3.2 Separability; Primitive Element Theorem 11

4 Algebraic Closure **13**

5 Galois Extensions **14**
5.1 Automorphism Groups of Field Extensions 14
5.2 Artin's Theorem 16
5.3 Galois Correspondence 19
5.4 Cubics and Discriminants 20

6 Symmetric Polynomials **21**

7 Cyclotomic Extensions **23**

8 Kummer Theory **27**
8.1 Kummer Extensions 27
8.2 Cubics Revisited 28
8.3 Solving Equation by Radicals 29

*Based on the lectures under the same name taught by Prof. I. Grojnowski in Michaelmas 2021.

9 Quartics **31**

10 Miscellany **33**
10.1 Reduction modulo p 33
10.2 Trace and Norm 34
10.3 Normal Basis Theorem 35
10.4 Function Fields 35

0 Notation and Revision

In this course, a ring will mean a commutative ring with identity. A field is a ring where every nonzero element has a inverse. For a ring R, we write $R^\times$ to denote the set of units (invertible elements) in R.
The polynomial ring over R is written as $R[X] = \{\sum_i \alpha_i X^i : \alpha_i \in R, \alpha_i = 0 \text{ for all but finitely many } i\} = \{\text{functions } \mathbb{N} \to \mathbb{R} \text{ with finite support}\}$. In general, we can formulate multivariate polynomial rings over R inductively via $R[X_1, \ldots, X_n] = (R[X_1, \ldots, X_{n-1}])[X_n]$.
We can evaluate a polynomial in $R[X]$ by plugging in $X = x \in R$. That is, every polynomial in $R[X]$ corresponds to a function $R \to R$. This gives an identification $R[X] \to \{f : R \to R\}$ which is in general neither injective (X^p and X give the same function on $\mathbb{F}_p = \mathbb{Z}/p\mathbb{Z}$; alternatively, observe that there are only p^p functions $\mathbb{F}_p \to \mathbb{F}_p$ but infinitely many polynomials) nor surjective (polynomials cannot be step functions). If R has no zero divisors (i.e. R is an integral domain), then neither does $R[X]$ and $(R[X])^\times = R^\times$.

Proposition 0.1. *If $R = K$ is a field, then $K[X]$ is a Euclidean domain with Euclidean function* deg.

That is, for any $a, b \in K[X]$ with $b \neq 0$, there are (unique) $q, r \in K[X]$ such that $\deg r < \deg b$ and $a = qb + r$.

Corollary 0.2. *$K[X]$ is a PID (every ideal in $K[X]$ is generated by one element) and hence a UFD (where we have unique factorisation)*

Corollary 0.3. *$f \in K[X]$ is irreducible iff (f) is prime iff (f) is maximal iff $K[X]/(f)$ is a field.*

Corollary 0.4. *For $a, b \in K[X]$ nonzero, $(a) + (b)$ is an ideal and hence equals (g) for some $g \in K[X]$.*

g is known as the greatest common divisor of a, b, written $g = \gcd(a, b)$.

Proposition 0.5. *Any nonzero $f \in K[X]$ has at most* $\deg f$ *many zeros.*

Definition 0.1. The field of rational functions over a field K is $K(x) = \{f/g : f, g \in K[X], g \neq 0\}/(f/g \sim f'/g' \iff gf' = g'f)$, that is, the field of fractions of $K[X]$.

Remark. $K[X_1, \ldots, X_n]$ is always a UFD, but usually not a PID for $n > 1$.

1 Field Extensions

1.1 Finite and Algebraic Extensions

Definition 1.1. A field extension L/K is an inclusion of fields $K \subset L$. We say L is an extension of K.

Alternatively, it is sometimes more useful to view a field extension L/K as a nonzero field homomorphism $K \hookrightarrow L$.

Example 1.1. $\mathbb{C}/\mathbb{R}, \mathbb{R}/\mathbb{Q}, \mathbb{C}/\mathbb{Q}, K(x)/K, (\mathbb{C}(z)[y]/(z^3 - z - y^2))/\mathbb{C}(z)$. More generally, there is a big class of examples given by $L = K[X]/f$ where f is irreducible in $K[X]$.

Note that when we have a field extension L/K, then L is immediately a vector space over K.

Definition 1.2. The degree $[L : K]$ of a field extension L/K is the dimension of L as a vector space over K.
If $[L : K] < \infty$, we say the extension L/K is finite; Otherwise, we say it is infinite.

Example 1.2. $[\mathbb{C} : \mathbb{R}] = 2, [\mathbb{R} : \mathbb{Q}] = \infty, [\mathbb{C} : \mathbb{Q}] = \infty, [\mathbb{C}(z)[y]/(z^3 - z - y^2) : \mathbb{C}(z)] = 2$.

For a field K, there is always a smallest subfield that contains 1: We've got the ring homomorphism $\mathbb{Z} \to K$ by sending $1 \in \mathbb{Z}$ to $1 \in K$ and extending accordingly. The kernel of this map is an ideal of $\mathbb{Z}$, hence has the form $n\mathbb{Z}$ for some n. The image of this map is a subring of K which has to be an integral domain, hence (by first isomorphism theorem) $\mathbb{Z}/n\mathbb{Z}$ is an integral domain, i.e. n is either 0 or a prime.

Definition 1.3. $n = \operatorname{char} K$ is called the characteristic of K.

Example 1.3. $\mathbb{F}_p(x)$ (where p is a prime) has characteristic p.

The field of fractions of the image of this map is then a subfield of K, which is clearly isomorphic to $\mathbb{Q}$ when $\operatorname{char} K = 0$ and $\mathbb{F}_p$ when $\operatorname{char} K = p > 0$.

Lemma 1.1. *If F is a finite field, then* $\operatorname{char} F = p < \infty$ *and* $|F| = p^n$ *for some* $n \geq 1$.

Proof. We already know that F contains $\mathbb{F}_p$ as a subfield. This then means that F is a field extension of $\mathbb{F}_p$ □

We will show later that there does exist a field with order p^n for each prime p and $n \in \mathbb{N}$.

Definition 1.4. For a field extension L/K, we write $K[\alpha]$ to denote the smallest subring of L containing K and α, and $K(\alpha)$ the smallest subfield of L containing K and α.

$K(\alpha)$ is also known as the field obtained from K by adjoining α. It is also the field of fractions of $K[\alpha]$, which can be explicitly described as

$$K[\alpha] = \left\{ \sum_{i=0}^{N} r_i \alpha^i : r_i \in K, N \in \mathbb{N} \right\}$$

Example 1.4. In $\mathbb{C}/\mathbb{Q}$, we have $\mathbb{Q}[i] = \mathbb{Q}(i) = \{a + bi : a, b \in \mathbb{Q}\}$.

We have a homomorphism $\Phi_\alpha : K[X] \to L$ given by $\Phi(f) = f(\alpha)$. The image of Φ is then $K[\alpha]$.

Definition 1.5. α is transcendental over K if Φ_α is injective; α is algebraic over K if Φ_α is not injective.

What happens when Φ_α is not injective? $\ker \Phi_\alpha$ is an ideal $(f) \leq K[X]$ where necessarily $f(\alpha) = 0$ and f has the minimal degree over all polynomials vanishing at α.

Definition 1.6. f (normalised so that it's monic) is the minimal polynomial of α over K.
The degree $\deg_K(\alpha)$ of α over K is defined as $\deg f$.

Note that f needs to be irreducible by the minimality of its degree. Hence α is algebraic over K iff $K[\alpha] = K(\alpha)$ (via $K[X]/(f)$).

Proposition 1.2. *α is transcendental over K iff Φ_α is an isomorphism onto $K[\alpha]$, hence extends to an isomorphism $K(X) \to K(\alpha)$.*

Proof. There is nothing to prove /shrug. □

In particular, all transcendental extensions of K that has the form $K(\alpha)$ are isomorphic to $K(X)$.

Example 1.5. $\mathbb{Q}(\pi) \cong \mathbb{Q}(e)$ as fields, that is if you already know that π and e are transcendental.

Proposition 1.3. *For a field extension L/K and $\alpha \in L$, the followings are equivalent:*
(i) α is algebraic over K.
(ii) $[K(\alpha) : K] < \infty$.
(iii) $\dim_K K[\alpha] < \infty$.
(iv) $K[\alpha] = K(\alpha)$.
(v) $K[\alpha]$ is a field.

When any of these happens, we have $\deg_K(\alpha) = [K(\alpha) : K]$.

Proof. It's clear that (i) $\iff$ (iii) and (iv) $\iff$ (v). It is also clear that (ii) $\implies$ (iii) and (iii) (iv) together implies (ii). To show that (iv) $\implies$ (i), one observes that if $K[\alpha] = K(\alpha)$, then $\alpha^{-1} = \sum_i r_i \alpha^i$ which means that $0 = -1 + \sum_i r_i \alpha^{i+1}$. We have already shown (iii) $\implies$ (iv) in our previous discussion, but here's an alternative way to see it: For $g \in K[\alpha]$ nonzero, consider $m_g : K[\alpha] \to K[\alpha]$ via $\gamma \mapsto g\gamma$. It is a K-linear endomorphism of the finite-dimensional vector space $K[\alpha]$, and is injective, hence is surjective. In particular, there is some γ such that $m_g(\gamma) = 1 \iff g\gamma = 1$. □

Remark. A word of warning is that being algebraic/transcendental depends on which base field we are taking, e.g. $2\pi i$ is algebraic over $\mathbb{R}$ but transcendental over $\mathbb{Q}$. The minimal polynomial, similarly, also depends on the base field. $\alpha = \sqrt{2}(1 + i)/2$ has minimal polynomial $x^4 + 1$ over $\mathbb{Q}$ but $x^2 - i$ over $\mathbb{Q}[i]$.

Note also here that $\mathbb{Q} \subset \mathbb{Q}[i] \subset \mathbb{Q}[\alpha]$ and we have $[\mathbb{Q}[i] : \mathbb{Q}] = 2, [\mathbb{Q}[\alpha] : \mathbb{Q}[i]] = 2, [\mathbb{Q}[\alpha] : \mathbb{Q}] = 4 = 2 \times 2$. This is very obviously not just a coincidence.

Theorem 1.4 (Tower Law)**.** *Suppose $M/L/K$ are field extensions, then M/K is finite iff both M/L, L/K are finite, in which case we have $[M : K] = [M : L][L : K]$.*

This is actually weaker than what we can prove, which is the following:

Proposition 1.5. *Suppose L/K is a finite extension and V a vector space over L, then V is finite dimensional over L iff V is finite dimensional over K, in which case $\dim_K V = [L : K] \dim_L V$.*

Proof. Suppose $\dim_L V = d$, then $V \cong L^d$ both as vector spaces over L and K. But L has to be isomorphic as a K-vector space to K^n where $n = [L : K]$, so $V \cong K^{nd}$ as K-vector spaces. The converse is clear since we can span a K-basis under L. □

Corollary 1.6. *Let L/K be a finite extension and $\alpha \in L$. If α is algebraic over K, then $\deg_K(\alpha) \mid [L : K]$.*

Proof. Immediate by considering the intermediate field $K(\alpha)$ which has degree $\deg_K(\alpha)$ over K. □

Corollary 1.7. *If $[L : K]$ is prime, then $K(\alpha) = L$ for any $\alpha \in L \setminus K$.*

Definition 1.7. For a field extension L/K and $\alpha_1, \ldots, \alpha_n \in L$, we write $K(\alpha_1, \ldots, \alpha_n)$ to denote the smallest subfield of L containing K and $\alpha_1, \ldots, \alpha_n$.

Example 1.6. 1. For $L = \mathbb{Q}(\sqrt[3]{2}, \sqrt[4]{5})$, we claim that $[L : \mathbb{Q}] = 12$. Indeed $3, 4$ are both divisors of $[L : \mathbb{Q}]$ by considering the intermediate fields $\mathbb{Q}(\sqrt[3]{2})$ and $\mathbb{Q}(\sqrt[4]{5})$. So $12 \mid [L : \mathbb{Q}]$. Conversely, $L = \mathbb{Q}(\sqrt[3]{2}, \sqrt[4]{5}) = \mathbb{Q}(\sqrt[3]{2})(\sqrt[4]{5})$ and the minimal polynomial of $\sqrt[4]{5}$ over $\mathbb{Q}(\sqrt[3]{2})$ divides the minimal polynomial of it over $\mathbb{Q}$, which has degree 4. Consequently $[L : \mathbb{Q}(\sqrt[3]{2})] \leq 4 \implies [L : \mathbb{Q}] \leq 12$, hence $[L : \mathbb{Q}] = 12$.
2. Take $\alpha = e^{2\pi i/p} + e^{-2\pi i/p}$ for an odd prime p. We want to know $\deg_{\mathbb{Q}}(\alpha)$. Write $\omega = e^{2\pi i/p}$ which has minimal polynomial $f(x) = 1 + x + \cdots + x^{p-1}$ (it is irreducible over $\mathbb{Q}$ by Gauss's lemma and Eisenstein's criterion), so $[\mathbb{Q}(\omega) : \mathbb{Q}] = p - 1$. But $\alpha \in \mathbb{Q}(\omega)$, so $\deg_{\mathbb{Q}}(\alpha) \mid p - 1$. Inevitably we are going to compute $[\mathbb{Q}(\omega) : \mathbb{Q}(\alpha)]$. We have $\omega\alpha = \omega^2 + 1$, so ω is a root of $x^2 - \alpha x + 1$ which has degree 2, leaving $[\mathbb{Q}(\omega) : \mathbb{Q}(\alpha)]$ either 2 or 1. We know it is not 1 since $\mathbb{Q}(\alpha) \subset \mathbb{R}$ but $\omega \notin \mathbb{R}$, so it has to be 2, i.e. $\deg_{\mathbb{Q}}(\alpha) = (p - 1)/2$.

Lemma 1.8. *Fix a field extension L/K.*
(i) $\alpha_1, \ldots, \alpha_n \in L$ are all algebraic over K iff $K(\alpha_1, \ldots, \alpha_n)/K$ is finite.
(ii) If α, β are algebraic over K, so are $\alpha \pm \beta, \alpha\beta, \alpha/\beta$ (if $\beta \neq 0$).

Proof. (i) follows pretty much from Proposition 1.3 and Theorem 1.4. (ii) follows from (i) since $\alpha \pm \beta, \alpha\beta, \alpha/\beta$ are all in $K(\alpha, \beta)$. □

Corollary 1.9. *The elements of L algbebraic over K form a subfield of L.*

Example 1.7. For a field extension L/K and $a, b \in K$. Say we have $\alpha, \beta \in L$ with $\alpha^2 = a, \beta^2 = b$. We want to determine a polynomial vanishing on $\gamma = \alpha + \beta$ because it's fun, and because our previous discussion means that this is possible. Indeed $\gamma^2 = (a + b) + 2\alpha\beta, \gamma^4 = (a + b)^2 + 4\alpha\beta(a + b) + 4ab = (a^2 + 6ab + b^2) + 4(a + b)\alpha\beta$. So $\gamma^4 - 2(a + b)\gamma^2 + (a - b)^2 = 0$. As a consequence, $[K(\gamma) : K] \leq 4$. Take $K = \mathbb{Q}$, then as one could check $[\mathbb{Q}(\gamma) : \mathbb{Q}] = 4$ given that neither a, b nor ab is a square.

What's the moral (or if you like, an alternative justification) of such a theory? Suppose L/K is a field extension and $\alpha, \beta \in L$ are algebraic over K with $\deg_K(\alpha) = n, \deg_K(\beta) = m$, then $\alpha^i \beta^j$ (for $i \in \{0, \ldots, n-1\}, j \in \{0, \ldots, m-1\}$) spans $K[\alpha, \beta]$, i.e. $[K[\alpha, \beta] : K] \leq mn$. Therefore $1, \gamma, \ldots, \gamma^{mn}$ must be linearly dependent over K. But the linear dependence relation is just a nonzero polynomial equation that γ satisfies, which means that γ is algebraic (and has degree at most mn).

Definition 1.8. An extension L/K is algebraic if every $\alpha \in L$ is algebraic ovr K.

Proposition 1.10. *(i) Any finite extension is algebraic.*
(ii) $K(\alpha)/K$ is algebraic iff α is algebraic over K.
(iii) For a chain of field extensions $M/L/K$, M/L and L/K are both algebraic iff M/K is algebraic.

So an extension L/K is algebraic iff L is a union of subfields that are finite extensions of K ("locally finite").

Proof. We've pretty much proved (i) already: If L is finite over K, then any $\alpha \in L$ must have $1, \alpha, \ldots, \alpha^{[L:K]}$ linearly dependent.
(ii) is also clear from (i): α is algebraic over K iff $K(\alpha)/K$ is finite.
As for (c), suppose M/K is algebraic, then any $\alpha \in M$ is algebraic over K, hence L. Also, any $\alpha \in L$ is also in M, therefore is algebraic over K. Conversely, suppose M/L and L/K are algebraic and $\alpha \in M$. We know that $r_0 + r_1\alpha + \cdots + r_d\alpha^d = 0$ for some d and $r_i \in L$. As L is algebraic over K, each r_i is algebraic over K, so $L_0 = K(r_0, \ldots, r_d)$ is finite over K. But then α is algebraic over L_0, so $L_0(\alpha)$ is finite over L_0, hence $L_0(\alpha)$ is finite over K which implies (by (i)) that α is finite over K. □

Example 1.8. $\bar{\mathbb{Q}} = \{\alpha \in \mathbb{C} : \alpha \text{ algebraic over } \mathbb{Q}\}$ is a field by Corollary 1.9 and is by definition algebraic. However it's clear that $[\bar{\mathbb{Q}} : \mathbb{Q}] = \infty$ (by e.g. considering the intermediate fields $\mathbb{Q}(\sqrt[n]{3})$ having degree n over $\mathbb{Q}$).

1.2 Construction with Straightedge and Compasses

Since antiquity, people have wondered the following: Which sort of things can you construct with just a straightedge and a pair of compasses? By this, we mean the game in which (possibly with some given starting position) you are only allowed to join any two given points to get a straight line and draw a circle given a point as its centre and a point on its circumference.
There are many things we can do with this setup, as we've all seen in middle school. We can draw a line through a constructed point p perpendicular to a constructed line l, draw a line though p parallel to l, copy the distance between two points to a line with one end specified, etc.. The fact that we can do these means that, given a segment as initial position, we can construct a Cartesian coordinate with the length of the segment considered as the unit length.
We call a number constructible if it's possible to construct two points on the plane whose distance is its absolute value. It's clear that, on the coordinate plane, (a, b) is constructible iff both a, b are constructible.

Proposition 1.11. *Constructible numbers form a field.*

Proof. Addition and substraction are easy. One can use a similar triangles trick to get multiplications and divisions (since we can copy angles). □

Lemma 1.12. *If $a > 0$ is constructible, so is $\sqrt{a}$.*

Proof. Similar triangles, circles, radius, work it out yourself. □

One of the construction problems that was not solved for a long time is whether you can construct a cube with volume $2A$ given a cube of volume A. Equivalently, we want to construct $\sqrt[3]{2}$ given the Cartesian coordinate we constructed.
After some failure in trying to do so by hand, we turn to prove that this is impossible.

Theorem 1.13. *Let $\mathbb{Q} = F_0 \subset F_1 \subset F_2 \subset \cdots \subset F_n = K$ be a chain of subfields of $\mathbb{R}$ such that for each i, $F_{i+1} = F_i(\sqrt{r_i})$ for some $r_i \in F_i$. Then every element in K is constructible.*
Conversely, if $a_1, \ldots, a_N$ are constructible numbers, then there is a chain of subfields as above such that $a_1, \ldots, a_N \in K$.

Proof. The first part follows immediately from previous constructions. The second part follows from the fact that any lines and circles are vanishing sets of polynomials with degree at most 2. □

Corollary 1.14. *If $a \in \mathbb{R}$ is constructible, then $\deg_{\mathbb{Q}} a$ is a power of 2.*

Proof. Theorem 1.4. □

But $\sqrt[3]{2}$ has degree 3 over $\mathbb{Q}$! So we've solve it: Constructing $\sqrt[3]{2}$ is, indeed, impossible.
There are also many other straightedge-and-compass construction problems from antiquity whose impossibility can be proved in this way.
For example, the problem of "squaring the circle" asks if one can construct a square whose area is the same as a circle of radius 1. That is, we want to construct $\sqrt{\pi}$. It is a somehow hard theorem to prove, but $\sqrt{\pi}$ is not algebraic, hence by the preceding corollary we conclude that this is impossible.
Another famous problem is trisecting an angle: Given an angle, is there a way to trisect it? We can in fact show something stronger: You already can't trisect an angle of size $\pi/3$ (which can be constructed as $\cos(\pi/3) = 1/2$)! Indeed, if one can, then one can construct $\alpha = \cos(\pi/9)$. We have the formula $\cos(3\theta) = 4\cos^3\theta - 3\cos\theta$ which gives $8\alpha^3 - 6\alpha - 1 = 0$. One can show that $8X^3 - 6X - 1$ is irreducible over $\mathbb{Q}$, hence α has degree 3 over $\mathbb{Q}$ and we conclude that it's impossible to construct α, hence impossible to trisect $\pi/3$.
We have shown that, for an odd prime p, $\cos(2\pi/p)$ has degree $(p-1)/2$ over $\mathbb{Q}$, so a regular p-gon is not constructible if $p - 1$ is not a power of 2. Gauss proved something much stronger: A regular p-gon is constructible if and only if p is a Fermat prime, i.e. prime that has form $1 + 2^{2^n}$ for some n.

2 Splitting Fields

Definition 2.1. Let K be a field and $f \in K[X]$. A splitting field for f is an extension L/K such that:

1. f splits into linear factors in $L[X]$.
2. If $\alpha_1, \ldots, \alpha_n$ are roots of f in L, then $L = K(\alpha_1, \ldots, \alpha_n)$.

Equivalently, L/K is a splitting field of f iff L is the smallest field extension of K in L on which f splits into linear factors.

Example 2.1. 1. For $K = \mathbb{Q}$, $\mathbb{Q}(i)$ is a splitting field for $f(X) = X^2 + 1$.
2. For $f(X) = X^3 - 2$ over $\mathbb{Q}$, we can split it in $\mathbb{C}$ as $f(X) = (X - \alpha)(X - \alpha\omega)(X - \alpha\omega^2)$ where $\alpha = \sqrt[3]{2}$ and $\omega = e^{2\pi i/3}$. As $\deg_{\mathbb{Q}} \alpha = 3, \deg_{\mathbb{Q}} \omega = 2$, we know that $[\mathbb{Q}(\alpha, \omega) : \mathbb{Q}] = 6$. However, for any root β of f, $\mathbb{Q}(\beta)$ would have degree 3 over $\mathbb{Q}$, so the splitting field of f cannot be obtained by adjoining only one of its roots. That being said, the splitting field of any quadratic polynomial over $\mathbb{Q}$ can be obtained by adjoining only one of the roots.
3. For $f(x) = (X^p - 1)/(X - 1)$ over $\mathbb{Q}$, $\mathbb{Q}(e^{2\pi i/p})$ is a splitting field of f.

Note that if $f \in K[X]$ is irreducible, then $K_f = K[X]/(f)$ is a field extension of K. Now $\alpha = X + (f) \in K_f$ would be a root of f in K_f. Also, $1, \alpha, \alpha^2, \ldots, \alpha^{d-1}$ is a basis of K_f as a K-vector space (where $d = \deg f$). So in fact $K_f = K(\alpha)$! Iterating this procedure shows that

Theorem 2.1 (Existence of Splitting Fields)**.** *For any $f \in K[X]$, there exists a field extension L/K that is a splitting field for f.*

Example 2.2. We can build $\mathbb{C}$ as the splitting field of $f(X) = X^2 + 1$ over $\mathbb{R}$. Indeed, it is just $\mathbb{R}[X]/(X^2 + 1)$. In this case, i is just $X + (X^2 + 1)$.
What if we've taken a different polynomial? Say we take $g(Y) = Y^2 + 2Y + 2 = (Y + 1)^2 + 1$. Then $\mathbb{R}[Y]/(Y^2 + 2Y + 2) \cong \mathbb{R}[1 + i] = \mathbb{C}$, as well!

Despite that the two examples give the same field, there is no canonical way to identify the two constructions together: One can either take $Y \mapsto -X - 1$ or $Y \mapsto X - 1$, both are perfectly fine and there is absolutely no reason to prefer one over the other. From a different point of view, there is a field automorphism $\mathbb{C} \to \mathbb{C}$ given by complex conjugation which happens to fix $\mathbb{R}$. Any $L \hookrightarrow \mathbb{C}$, when composed with this automorphism, immediately give rise to a different, equally "worthy" $L \hookrightarrow \mathbb{C}$.
We want to capture this behaviour by introducing the category of field extensions of K.

Definition 2.2. Suppose L/K and M/K are two extensions of the same field K. A K-homomorphism $L/K \to M/K$ is a field homomorphism $L \to M$ that fixes K.

Note that any nonzero field homomorphism $K \to L$ can be viewed as a K-homomorphism by considering L as a field extension of K via the embedding of K in it under the map.

Example 2.3. Complex conjugation is an $\mathbb{R}$-homomorphism $\mathbb{C} \to \mathbb{C}$.

Lemma 2.2. *Suppose L/K is a field extension and $f \in K[X]$ is irreducible, then there is a bijection between the K-homomorphisms $K_f \to L$ and the roots of f in L.*

Proof. A K-homomorphism $\phi : K_f \to L$ corresponds to the root $\phi(X + (f)) \in L$ of f in L (since ϕ commutes with any K-polynomial as a ring homomorphism).

Conversely, given a root α of f in L, the map $K_f \to L$ sending X to α has kernel (f), hence induces a K-homomorphism $K_f \to L$ via $X + (f) \mapsto \alpha$. It is easy to see that these two are inverses to each other. □

In particular, the number of K-homomorphisms from K_f to L equals to the number of roots of f in L, which is at most $\deg f$ (which, in turn, is finite).

Corollary 2.3. *Suppose L/K is a field extension and α, β are algebraic over K, then α, β have the same minimal polynomial iff there exists a K-homomorphism $K(\alpha) \to K(\beta)$ sending α to β.*

Example 2.4. Let $f(X) = X^3 - 2 \in \mathbb{Q}[X]$ and $\omega = e^{2\pi i/3}, \alpha = \sqrt[3]{2}$, then the roots of f in $\mathbb{C}$ are $\alpha, \alpha\omega, \alpha\omega^2$. So there is a $\mathbb{Q}$-isomorphism $\mathbb{Q}(\alpha) \to \mathbb{Q}(\alpha\omega)$.

Amazingly, $\mathbb{Q}(\alpha)$ is a subfield of $\mathbb{R}$ but $\mathbb{Q}(\alpha\omega)$ is not! So K-homomorphisms over a base field K only capture the internal algebraic structure of a field extension over K, but not how it might embed into a bigger field.

Theorem 2.4 (Uniqueness of Splitting Fields)**.** *Suppose $f \in K[X]$ is a polynomial and L a splitting field of f. Then any K-homomorphism $\phi : K \to M$ with f splits in M extends to a K-homomorphism $\tilde{\phi} : L \to M$, which is an isomorphism when M is also a splitting field of f. Moreover, there are at most $[L : K]$ such extensions, with equality iff no irreducible factor of f has repeated roots in M.*

Proof. Induction on $[L : K]$. If $[L : K] = 1$, then f splits in K and there is nothing to prove. Otherwise, let $\alpha_1 \in L \setminus K$ be a root of f and g the minimal polynomial of α_1 over K. Then $g \mid f$ and g too splits into linear factors in L and M. By the preceding lemma, K-homomorphisms $\hat{\phi} : K(\alpha_1) \cong K_g \to M$ bijects with the roots of g in M. The number of such homomorphisms is the number of roots of g in M, which is at most $[K(\alpha_1) : K]$ with equality iff g has no repeated root in M.
Now $[L : K(\alpha_1)] < [L : K]$, so we can apply the induction hypothesis and conclude that any K-homomorphism $\hat{\phi} : K(\alpha_1) \to M$ (which is also a $K(\alpha_1)$-homomorphism) extends to a $K(\alpha_1)$-homomorphism $\tilde{\phi} : L \to M$ (which also extends ϕ as a K-homomorphism).
Moreover, Theorem 1.4 shows that there is at most $[L : K(\alpha_1)][K(\alpha_1) : K] = [L : K]$ such extensions. Equality holds iff no irreducible factor has repeated roots by the induction step.
When M is a splitting field, the extension is an isomorphism since if f splits into $(X - \alpha_1) \cdots (X - \alpha_n)$ in L, then it splits into $(X - \tilde{\phi}(\alpha_1)) \cdots (X - \tilde{\phi}(\alpha_n))$ in M. □

3 Finite Fields and Separability

3.1 Existence and Uniqueness of Finite Fields

Proposition 3.1. *Let K be a field, then any finite subgroup of $K^\times$ is cyclic.*

Proof. By the classification of finite abelian groups, $G \cong C_{m_1} \times \cdots C_{m_r}$ for some $m_1 \mid m_2 \mid \cdots \mid m_r$. If G is trivial, so is the result. Otherwise, WLOG $m_1 > 1$, then $x^{m_r} - 1$ would have more than m_r roots, contradiction. □

Corollary 3.2. *Finite fields have cyclic multiplicative groups.*

This is not obvious at all – we don't even have a way to produce a canonical generator for the group. In fact, this lack of canonical generator has a certain connection with the lack of canonical embeddings of K_f we discussed earlier. We will talk about that later.

Example 3.1. $\mathbb{F}_7^\times$ is generated by 3; $\mathbb{F}_{11}^\times$ is generated by 2.

Any finite field must have finite characteristic, hence contains $\mathbb{F}_p$ for some prime p, therefore has order p^n for some n.

Proposition 3.3. *Let K be a finite field with $q = p^n$ elements, then every $\alpha \in K$ would be a root of $X^q - X$. Conversely, $X^q - X$ factors into distinct linear factors in K.*

Proof. Certainly $0^q - 0 = 0$. For $x \neq 0$, $x \in K^\times$ and hence $x^{|K^\times|} = 1 \implies x^q - x = 0$. As for the converse, one just use the fact that $X^q - X$ has at most q roots, and we have just found all of them. □

We want to show that, for every prime p and positive integer n, there exists a finite field with order $q = p^n$. An obvious way to proceed is to consider the splitting field of $X^q - X$ over $\mathbb{F}_p$, inspired by the proposition. To justify that this splitting field indeed has order q, we however would need some more work on multiplicity of roots.

Definition 3.1. Let K be a field. The formal derivative is a linear map $\mathrm{D} : K[X] \to K[X]$ sending $\sum_k r_k x^k$ to $\sum_k k r_l x^{k-1}$.

It is easy to check that the usual properties of derivatives stay, like the product rule and chain rule.

Lemma 3.4. *Suppose L/K is a field extension and $\alpha \in L$ is a root of some $f \in K[X]$. Then α is a simple root of f iff $\mathrm{D}f(\alpha) \neq 0$.*

Proof. If $f(X) = (X - \alpha)g(X)$ (in $L[X]$), then $\mathrm{D}f(X) = (X - \alpha)\mathrm{D}g(X) + g(X)$. □

Corollary 3.5. *f has repeated roots iff* $\deg \gcd(f, \mathrm{D}f) \geq 1$.

Lemma 3.6. *Let R be a ring with characteristic p for p prime, then the map $F : R \to R, x \mapsto x^p$ is a ring homomorphism.*

This homomorphism is called the Frobenius.

Proof. $p \mid \binom{p}{k}$ for $1 \leq k \leq p - 1$. □

Example 3.2. Suppose K is a field with characteristic $p > 0$ and suppose $b \in K$ is not a p^{th} power (e.g. $K = \mathbb{F}_p(Y), b = Y$). Let L be the splitting field of $f(X) = X^p - b$, then for any root α of f, $Df(\alpha) = 0$, i.e. f has no simple roots. Indeed, we can even factorise $f(X) = (X - \alpha)^p$. However, f is irreducible over K, since any factor of it can be further factorised as $(X - \alpha)^m$ in L for some $m \in \{1, \ldots, p - 1\}$. But for $(X - \alpha)^m$ to have coefficients in K we must have $\alpha m \in K \implies \alpha \in K$ as $m \in \{1, \ldots, p - 1\}$, contradiction.

Theorem 3.7. *Let $q = p^n$ for a prime p and $n \geq 1$.*
(i) There exists a field $\mathbb{F}_q$ with q elements. Moreover, any two such fields are isomorphic.
(ii) $\mathbb{F}_q$ is the splitting field for $X^q - X \in \mathbb{F}_p[X]$.
(iii) $\mathbb{F}_q$ contains a field of order p^k iff $k \mid n$.

Proof. For (i) and (ii) it suffices to show that the splitting field for $X^q - X \in \mathbb{F}_p[X]$ has order q. Let K be this splitting field, then $K = \mathbb{F}_p(\alpha_1, \ldots, \alpha_q)$ with $\alpha_i^q = \alpha_i$. Note that if α, β are roots of $X^q - X$, so are $\alpha+\beta, \alpha\beta$ and α/β if $\beta \neq 0$. Hence $K = \{\alpha_1, \ldots, \alpha_q\}$. Also, if α is a root of f, then we have $\mathrm{D}f(\alpha) = -1$, so there is no repeated root. So indeed $|K| = q$.
(iii) The "only if" direction follows from Theorem 1.4. The "if" direction follows from the fact that $X^{p^k} - X \mid X^{p^n} - X$ whenever $k \mid n$. □

Example 3.3. 1. $X^4 - X = X(X-1)(X^2+X+1)$, so $\mathbb{F}_4 = \{0, 1, \alpha, \alpha+1\}$ where $\alpha^2 + \alpha + 1 = 0$. In particular $\mathbb{F}_2 \hookrightarrow \mathbb{F}_4$
2. $X^8 - X = X(X-1)(X^3+X+1)(X^3+X^2+1)$, so $\mathbb{F}_8$ contains $\mathbb{F}_2$ but not $\mathbb{F}_4$. Indeed, $[\mathbb{F}_4 : \mathbb{F}_2] = 2$ but $[\mathbb{F}_8 : \mathbb{F}_2] = 3$. From $[\mathbb{F}_8 : \mathbb{F}_2] = 3$ we also see that $\mathbb{F}_8 \cong \mathbb{F}_2[X]/(X^3+X+1) \cong \mathbb{F}_2[Y]/(Y^3+Y^2+1)$, again in a non-canonical way.

Remark. 1. Suppose $f \in \mathbb{F}_p[X]$ is irreducible with degree n, then $K = \mathbb{F}_p[X]/(f)$ is a field with $q = p^n$ elements. But $X^q - X$ splits completely in K and its roots are the elements of K, so $f \mid X^q - X$. Hence the set of irreducible polynomials with degree divisible by n is exactly the set of factors of $X^q - X$.
2. What we've shown also implies that, in $\mathbb{F}_p$, a irreducible polynomial of degree n exists for each n, namely the minimal polynomial the generator of $\mathbb{F}_{p^n}^\times$.

3.2 Separability; Primitive Element Theorem

Definition 3.2. $f \in K[X]$ is separable if it splits into distinct linear factors in its splitting field.

So f is separable iff $\gcd(f, \mathrm{D}f) = 1$.

Example 3.4. $X^q - X \in \mathbb{F}_p[X]$ is separable but $X^p - Y \in \mathbb{F}_p(Y)[X]$ is not.

Proposition 3.8. *Suppose $f \in K[X]$ is irreducible, then f is separable iff $\mathrm{D}f \neq 0$. In particular:*
(i) If $\operatorname{char} K = 0$, every irreducible polynomial is separable.
(ii) If $\operatorname{char} K = p > 0$, then f is separable iff f is not of the form $g(X^p)$ for some $g \in K[X]$.

Proof. $\gcd(f, \mathrm{D}f)$ is either 1 of f by irreducibility. But $\deg \mathrm{D}f < \deg f$, so either $\mathrm{D}f = 0$ or $\gcd(f, \mathrm{D}f) = 1$, i.e. f is separable. (i) and (ii) follow immediately. □

Definition 3.3. Suppose α is algebraic over K, then we say α is separable if its minimal polynomial is separable. An (algebraic) extension L/K is separable iff every $\alpha \in L$ is separable over K.

Note that α (with minimal polynomial f) is separable iff there are exactly $\deg f$ K-homomorphisms $K(\alpha) \to L$ for any field L where f splits.

Example 3.5. If $\operatorname{char} K = 0$, then all algebraic extensions are separable. $\mathbb{F}_{p^n}/\mathbb{F}_p$ is separable but $\mathbb{F}_p(X^{1/p})/\mathbb{F}_p(X)$ is not.

Adjoining an algebraic element gives an algebraic extension. How about adjoining a separable element?

Proposition 3.9. *Suppose α is algebraic and separable in K, then $K(\alpha)/K$ is separable.*

Proof. Let $\beta \in K(\alpha)$. Suppose α has minimal polynomial f and β has minimal polynomial g. Let M be the splitting field of fg over K, then in particular g splits in M. If we can show that there exists precisely $m = \deg g$ distinct K-homomorphisms $K(\beta) \to M$ then we are done. Note that $K \subset K(\beta) \subset K(\alpha)$. So $[K(\alpha) : K] = mn$ where $n = [K(\alpha) : K(\beta)], m = [K(\beta) : K]$.
Any K-homomorphism $\phi : K(\alpha) \to M$ can be restricted to K-homomorphisms $\bar{\phi} : K(\beta) \to M$. As α is separable over K, there are nm K-homomorphisms $K(\alpha) \to M$. The number of K-homomorphisms $K(\beta) \to M$ is at most $[K(\beta) : K] = m$ with equality iff β is separable.
As α is separable over K, it is separable over $K(\beta)$, so any K-homomorphism $\bar{\phi} : K(\beta) \to M$ extends to exactly $[K(\alpha) : K(\beta)] = n$ many homomorphisms $K(\alpha) \to M$. So the number of K-homomorphisms $K(\beta) \to M$ is at least $mn/n = m$, hence equals m. β is therefore separable. □

For algebraic extensions, we also have the nice result that $M/L, L/K$ are algebraic iff M/K is. Is this still true if we replace "algebraic" by "separable"?

Proposition 3.10. *If L/K is an extension and $f, g \in K[X]$, then:*
(i) $\gcd(f, g)$ is the same (up to constant, of course) in $K[X]$ and $L[X]$. Hence both $\gcd(f, g)$ and $\operatorname{lcm}(f, g) = fg/\gcd(f, g)$ are well-defined regardless of the field extension.
(ii) The lcm *of a finite set of separable polynomials is separable.*

Proof. (i) is clear. (ii) follows from (i) by choosing (by induction) an extension where all these polynomials split. □

Theorem 3.11 (Primitive Element Theorem)**.** *Let $L = K(\alpha_1, \cdots, \alpha_n, \beta)$ be a finite extension of K and $\alpha_1, \ldots, \alpha_n$ are separable. Then there exists $\gamma \in L$ (the "primitive element") such that $L = K(\gamma)$.*

In particular, if $\operatorname{char} K = 0$, every finite extension of K has the form $K(\alpha)$ for some algebraic α.

Proof. If K is finite, then L is finite, hence we can pick a generator $\gamma \in L^\times$ as the primitive element.
Suppose now that K is infinite, so L is infinite as well. It suffices to show the case for $n = 1$ since and obtain the general case from simple induction. Set $\alpha = \alpha_1$.
Let $\gamma = \gamma_c = \beta + c\alpha$ for $c \in K$. We will show that $K(\gamma_c) = L$ (or equivalently $\alpha \in K(\gamma_c)$) for all but finitely many c.
The idea is to extract information about the minimal polynomial of α over $K(\gamma)$. Suppose f is the minimal polynomial of α over K and g that of β over K. Let M be a splitting field for fg where f splits as $f(X) = (X - \alpha_1) \cdots (X - \alpha_n)$ and g as $g(X) = (X - \beta_1) \cdots (X - \beta_m)$ with $\alpha = \alpha_1, \beta = \beta_1$. Let $h(X) = g(\gamma - cX) \in K(\gamma)[X]$, then $h(\alpha) = g(\beta) = 0$ and $f(\alpha) = 0$, so $(X - \alpha) \mid \gcd(f, h)$.
If $\gcd(f, h) = X - \alpha$ (for all but finitely many c) then we are essentially done since it would mean that $X - \alpha \in K(\gamma)[X]$. We can compute $\gcd(f, h)$ in M

since it doesn't matter which field we take. Then $\gcd(f,h) = X-\alpha$ iff $h(\alpha_i) \neq 0$ if $i \geq 2$. But for any θ we have $h(\theta) = 0 \iff g(\gamma - c\theta) = 0$, so there are at most $\deg g < \infty$ many choices of c for $h(\theta) = 0$. Discarding these for each $\theta = \alpha_i, i \geq 2$ leaves those which make the statement true. □

Explicitly, the finite set we discarded is $\{(\beta_j - \beta)/(\alpha - \alpha_i) : i \geq 2\}$.

Example 3.6. $\mathbb{Q}[i, \sqrt[3]{2}] = \mathbb{Q}(\sqrt[3]{2} + ci)$ whenever $c \neq (\pm i - i)/(\sqrt[3]{2}(\omega^{\pm 1} - 1))$.

There are however other combinations of α, β that would be a primitive element, e.g. $\mathbb{Q}(\sqrt[3]{2} + ci) = \mathbb{Q}(i\sqrt[3]{2})$ as in the above example. In fact, by considering $\mathbb{Q}(\gamma)$ where $\gamma = c_1 + c_2\sqrt[3]{2} + c_3(\sqrt[3]{2})^2 + c_4 i + c_5 i\sqrt[3]{2} + c_6 i(\sqrt[3]{2})^2$ we can conclude that there are at most finitely many intermediate fields between $\mathbb{Q}[i, \sqrt[3]{2}]$ and $\mathbb{Q}$.
Is the separability condition really necessary? If we take $K = \mathbb{F}_p(X, Y)$ and $L = \mathbb{F}_p(X^{1/p}, Y^{1/p})$, then for any $\gamma \in L$ we clearly have $\gamma^p \in K$, so $\deg_K(\gamma) \leq p < p^2 \implies K(\gamma) \neq L$.

Proposition 3.12. *If M/L and L/K are finite and separable, so is M/K.*

Proof. Write $L = K(\alpha)$ and $M = L(\beta)$ by Theorem 3.11, then $M = K(\alpha, \beta)$. Suppose $M = K(\gamma)$ for some γ (again by Theorem 3.11). We shall show that γ is separable over K. Suppose $[K(\alpha) : K] = m, [K(\alpha, \beta) : K(\alpha)] = n$. Let T be a field in which the minimal polynomials of α, β, γ all split. As α is separable over K, there are m distinct K-homomorphisms from $K(\alpha)$ to T. Similarly, there are n distinct $K(\alpha)$-homomorphisms from $K(\alpha, \beta) \to T$ (extending a given inclusion of $K(\alpha)$ into T). So there are mn distinct K-homomorphisms $K(\gamma) = K(\alpha, \beta) \to T$ which means that γ is separable. □

4 Algebraic Closure

Definition 4.1. A field K is algebraically closed if every nonconstant polynomial in $K[X]$ has a root in K.

Lemma 4.1. *The followings are equivalent:*
(i) K is algebraically closed.
(ii) If L/K is a field extension, then $\alpha \in L$ is algebraic iff $\alpha \in K$.
(iii) K has no nontrivial algebraic extensions.

Proof. Trivial. □

Example 4.1. 1. $\mathbb{C}$ is algebraically closed.
2. $\bar{\mathbb{Q}} = \{\alpha \in \mathbb{C} : \alpha \text{ is algebraic over } \mathbb{Q}\}$ is algebraically closed.

Definition 4.2. A field extension L/K is called an algebraic closure of K if it is algebraic and L is algebraically closed.

Example 4.2. 1. $\mathbb{Q}$ is an algebraic closure of $\mathbb{Q}$.
2. Let's construct an algebraic closure of $\mathbb{F}_p$. Choose a sequence $(r_i)_i \in \mathbb{N}$ such that $r_i \mid r_{i+1}$ and for any integer n, there is some i such that $n \mid r_i$ (e.g. $r_i = i!$).

Let $F_i = \mathbb{F}_{p^{r_i}}$ and choose embeddings $F_i \to F_{i+1} \to F_{i+2} \to \cdots$. Note that each $\mathbb{F}_{p^n}$ is a subfield of F_i for sufficiently large n. Set

$$\bar{\mathbb{F}}_p = \bigcup_{i=1}^{\infty} F_i = \bigcup_{i=1}^{\infty} \mathbb{F}_{p^{r_i}}$$

Clearly $\bar{\mathbb{F}}_p$ is a field, is algebraic over $\mathbb{F}_p$, and is algebraically closed. Hence $\bar{\mathbb{F}}_p$ is the algebraic closure of $\mathbb{F}_p$.

The same argument as in the second example will show the general existence and uniqueness of algebraic closures, with the help of Zorn's lemma.

Theorem 4.2. *Every field has an algebraic closure. Moreover, if L_1, L_2 are algebraic closures of K, then there is a K-isomorphism taking L_1 to L_2.*

The algebraic closure of K is often denoted as $\bar{K}$.

5 Galois Extensions

5.1 Automorphism Groups of Field Extensions

We have now know a decent amount of properties about field extensions of the form $K(\alpha)/K$ for some α algebraic over K. This only gives the algebraic structure given by adjoining a single root – we don't, so far, know about anything between the roots of the same irreducible polynomial. Clearly there are some kind of symmetry between them, inspired by Vieta's formulae. How would we capture it in the language of field extensions?

Definition 5.1. Let L/K be a field extension. The automorphism group of L/K is

$$\mathrm{Aut}(L/K) = \{K\text{-automorphisms of } L\} = \{K\text{-isomorphisms } L \to L\}$$

Example 5.1. 1. Suppose $[L : K] = 2$, then $L = K(\alpha)$ for some α with minimal polynomial $f(X) = X^2 + bX + c = (X - \alpha)(X - \alpha')$. We've got $\alpha + \alpha' = -b, \alpha\alpha' = c$. So $\alpha' = -b - \alpha \in L$ which means that L is the splitting field for f and $L = K(\alpha')$ as well. By Lemma 2.2, the only possible nontrivial K-automorphism of L is given by $\alpha \mapsto \alpha'$. This is nontrivial iff $\alpha \neq \alpha'$, so we conclude that $\mathrm{Aut}(L/K)$ is trivial when $\mathrm{char}\, K = 2$ and $\alpha^2 \in K$ and is isomorphic to $\mathbb{Z}/2\mathbb{Z}$ otherwise.
So for example $\mathrm{Aut}(\mathbb{Q}(i)/\mathbb{Q})$ and $\mathrm{Aut}(\mathbb{C}/\mathbb{R})$ are both $\mathbb{Z}/2\mathbb{Z}$ consisting of identity and complex conjugation.
$\mathrm{Aut}(\mathbb{Q}(\sqrt{2}+1)/\mathbb{Q})$ is also isomorphic to $\mathbb{Z}/2\mathbb{Z}$ with the one nontrivial $\mathbb{Q}$-automorphism given by $a+b(1+\sqrt{2}) \mapsto a+b(1-\sqrt{2})$, or $c+d\sqrt{2} \mapsto c-d\sqrt{2}$ which, expectedly, is also the nontrivial element in $\mathrm{Aut}(\mathbb{Q}(\sqrt{2})/\mathbb{Q})$ since $\mathbb{Q}(\sqrt{2}) = \mathbb{Q}(1+\sqrt{2})$ after all. Note that this map is not continuous under the standard topology, which shouldn't surprise you since what we are doing has nothing to do with analysis.
If $\mathrm{char}\, K = 2$, we can assume WLOG that $L = K(\alpha)$ where α has minimal polynomial $X^2 - D$. Then the nontrivial K-automorphism is given by $\alpha \mapsto -\alpha$.
2. $\mathrm{Aut}(\mathbb{Q}(\sqrt[3]{2})/\mathbb{Q})$ is trivial. Indeed, $\alpha\sqrt[3]{2}$ has minimal polynomial $f(X) =$

$X^3 - 2 = (X - \alpha)(X - \omega\alpha)(X - \omega^2\alpha)$ where $\omega = e^{2\pi i/3}$. But neither $\omega\alpha$ nor $\omega^2\alpha$ is in $\mathbb{Q}(\sqrt[3]{2})$ since they are not real, so Lemma 2.2 shows that there can be no nontrivial $\mathbb{Q}$-automorphisms of $\mathbb{Q}(\sqrt[3]{2}$.
3. We want to examine the automorphism group of a transcendental extension L/K where $L = K(x)$. Clearly any $\gamma \in \mathrm{PGL}_2(K)$ is a K-automorphism of L by Möbius transformations. One can also show the converse, so $\mathrm{Aut}(L/K) \cong \mathrm{PGL}_2(K)$. Note also that $\mathrm{PGL}_2(K)$ is infinite whenever K is.
4. We say a field extension L/K is biquadratic if $[L : K] = 4$, $\mathrm{char}\, K \neq 2$ and L is generated by two degree 2 elements α, β with WLOG (by completing squares) $\alpha^2 = a \in K, \beta^2 = b \in K$. Lemma 2.2 shows that there is a unqiue $K(\alpha)$-automorphism σ_β of L sending β to $-\beta$. Similarly, there is a unique $K(\beta)$-automorphism σ_β of L sending α to $-\alpha$. $\sigma_\alpha, \sigma_\beta$ are automatically K-automorphisms of L and $\sigma_\alpha\sigma_\beta = \sigma_\beta\sigma_\alpha, \sigma_\alpha^2 = \sigma_\beta^2 = 1$. Then $\mathrm{Aut}(L/K) = \{\mathrm{id}_L, \sigma_\alpha, \sigma_\beta, \sigma_\alpha\sigma_\beta\} \cong (\mathbb{Z}/2\mathbb{Z})^2$.

Lemma 5.1. *(i) Suppose L/K is a field extension and $\sigma \in \mathrm{Aut}(L/K)$, then for any $\alpha \in L$ and $f \in K[X]$ we have $f(\alpha) = 0 \iff f(\sigma\alpha) = 0$.*
(ii) Suppose $L = K(\alpha_1, \ldots, \alpha_n)$ and $\sigma \in \mathrm{Aut}(L/K)$ fixes all α_i, then $\sigma = 1$.
(iii) If L is a splitting field for f, then $\mathrm{Aut}(L/K) \leq S_{\deg f}$.

Proof. Quite clear. □

Note that the lemma, albeit nice, does not tell you exactly what subgroups of S_n can $\mathrm{Aut}(L/K)$ be.

Example 5.2. A biquadratic extension L/K, which is a splitting field, has $\mathrm{Aut}(L/K) \cong (\mathbb{Z}/2\mathbb{Z})^2 \hookrightarrow S_4$.

Not every field extension is a splitting field (e.g. $\mathbb{Q}(\sqrt[3]{2})/\mathbb{Q}$). Can we generalise the last part of the lemma to more general settings? The situation is weird for infinite extensions, but how about finite extensions?

Theorem 5.2. *Let L/K be a finite extension, then $\mathrm{Aut}(L/K)$ is finite.*

Proof. Suppose $L = K(\alpha_1, \ldots, \alpha_n)$ and let f_i be the minimal polynomial of α_i. Then any $\sigma \in \mathrm{Aut}(L/K)$ must permute the roots of $f = f_1 \cdots f_n$, and if σ fixes everything then it has to be the identity. This means that $|\mathrm{Aut}(L/K)| \leq (\deg f_1)!(\deg f_2)! \cdots (\deg f_n)!$ which is finite. □

Remark. On the other hand, since we only need to know the values of σ at each α_i to determine it and there are only $\deg f_i$ choices each time, we conclude a better bound $|\mathrm{Aut}(L/K)| \leq (\deg f_1)(\deg f_2) \cdots (\deg f_n)$ which removed the factorial growth.

We can get a slightly better estimate, with some representation theory.

Theorem 5.3. *Let G be a group and L a field. Suppose $\sigma_1, \ldots, \sigma_n : G \to L^\times$ are distinct group homomorphisms, then $\sigma_1, \ldots, \sigma_n$ are linearly independent over L.*

Proof. Suppose otherwise, then let n be the minimal positive integers such that there are nonzero y_i such that there are distinct group homomorphisms $\sigma_1, \ldots, \sigma_n : G \to L^\times$ with $\sum_i y_i\sigma_i = 0$ (automatically $n \geq 2$). Then $0 =$

$\sum_i y_i\sigma_i(gh) = \sum_i y_i\sigma_i(g)\sigma_i(h)$ for all $g, h \in G$. Choose g such that $\sigma_1(g) \neq \sigma_2(g)$, then

$$0 = \sum_{i=1}^{n} y_i\sigma_i(g)\sigma_i(h) - \sum_{i=1}^{n} y_i\sigma_1(g)\sigma_i(h) = \sum_{i=2}^{n} y_i(\sigma_i(g) - \sigma_1(g))\sigma_i(h)$$

violating minimality of n. □

Theorem 5.4. *Let L/K be a finite extension, then $|\operatorname{Aut}(L/K)| \leq [L : K]$.*

We will prove in a moment that we actually have $|\operatorname{Aut}(L/K)| \mid [L : K]$.

Proof. $\operatorname{Aut}(L/K)$ is a subgroup of $\{K\text{-linear maps } L \to L\}$ which is a vector space over L. Its dimension over L is $n = [L : K]$ by taking a basis consisting of the indicators of a K-basis of L. By the preceding theorem, elements of $\operatorname{Aut}(L/K)$ are linearly independent, therefore $|\operatorname{Aut}(L/K)| \leq n$. □

An alternative way to phrase this is the following: Suppose the result is false, then we can find $n + 1$ distinct elements of $\operatorname{Aut}(L/K)$, say $\sigma_1, \ldots, \sigma_{n+1}$. Let $\alpha_1, \ldots, \alpha_n$ be a K-basis for L and consider the matrix $A_{ij} = \sigma_i(\alpha_j)$ whose rows would have to be linearly dependent as $n+1 > n$. This means that $\sigma_1, \ldots, \sigma_{n+1}$ are linearly dependent which cannot happen.

5.2 Artin's Theorem

Definition 5.2. Let G be a subgroup of $\operatorname{Aut}(L) = \{\text{Field automorphisms of } L\}$. Its fixed field (or field of invariants) is defined as $L^G = \{l \in L : \forall g \in G, gl = l\}$.

It's easy to see that L^G is a field.

Example 5.3. Suppose $L = K(\sqrt{D})$ is a quadratic extension of K with $\operatorname{char} K \neq 2$. Let $G = \operatorname{Aut}(L/K) \cong \mathbb{Z}/2\mathbb{Z}$, then $L^G = K$.

The example hints that we might be able to recover K from L and $\operatorname{Aut}(L/K)$. Well, sometimes.

Definition 5.3. An extension L/K is Galois if it is finite and $K = L^{\operatorname{Aut}(L/K)}$.

Example 5.4. 1. If $\operatorname{char} K \neq 2$, then any quadratic and biquadratic extensions of K are Galois.
2. $\mathbb{Q}(\sqrt[3]{2})/\mathbb{Q}$ is not Galois as $\operatorname{Aut}(L/K)$ is trivial.
3. $K(X)/K$ is not Galois since it's not finite (lol).

Suppose L is a field and $G \leq \operatorname{Aut}(L)$ with G finite. Set $K = L^G$. What's the structure of the field extension L/K?
If L/K is Galois and $G = \operatorname{Aut}(L/K)$, then by definition $K = L^G$. We shall show, conversely, that L/L^G is always Galois.

Lemma 5.5. *Suppose $G \subset \operatorname{Aut}(L)$ is finite and $K = L^G$, then every $\alpha \in L$ has degree at most $|G|$. In particular, L/K is algebraic.*

Proof. Set $f(X) = \prod_{\tau \in G}(X - \tau\alpha) \in L[X]$. The lemma will follows if we can show that $f \in K[X]$. Indeed, G natually acts on $L[X]$ by ring homomorphisms via $\sigma(\sum_i a_i X^i) = \sum_i \sigma(a_i)X^i$ for $\sigma \in G$. We have $L[X]^G = L^G[X] = K[X]$ and $\sigma f = f$, so $f \in L[X]^G = K[X]$. □

Lemma 5.6. *Suppose $G \subset \mathrm{Aut}(L)$ is finite and $K = L^G$, then L/K is separable.*

Proof. Let $\alpha \in L$. We shall show that the minimal polynomial of α has distinct roots. Let the orbit $G\alpha = \{\sigma\alpha : \sigma \in G\}$ be enumerated as $\{\alpha_1, \ldots, \alpha_r\}$, then $g(X) = \prod_i (X - \alpha_i)$ vanishes at α, has distinct roots and is an element of $K[X]$ by the same argument as in the preceding lemma. The minimal polynomial of α divides g, hence also has distinct roots. □

Lemma 5.7. *The polynomial g constructed in the proof of the preceding lemma is the minimal polynomial of α.*

Proof. If not, then $g(X) = f_1(X)f_2(X)$ for $f_1, f_2 \in K[X]$. Write $f_1(X) = \prod_{i\in A}(X - \alpha_i), f_2(X) = \prod_{j\in B}(X - \alpha_j)$ where A, B partition the indices. But any $\sigma \in G$ would fix f_2, hence it fixes $\{\alpha_i\}_{i\in A}$. This can only mean that $\{\alpha_i\}_{i\in A}$ is either empty or all of $G\alpha$ (i.e. $\{\alpha_i\}_{i\in B}$), i.e. one of f_1, f_2 is constant. □

Example 5.5. Take $L = \mathbb{Q}(\sqrt{2}, i)$ and $K = \mathbb{Q}$ (given by $K = L^G$ with $G \cong (\mathbb{Z}/2\mathbb{Z})^2 = \mathrm{Aut}(L/K)$). If $\alpha = i + \sqrt{2} + 1$, then we can obtain its minimal polynomial by the above procedure

$$g(X) = (X - i - \sqrt{2} - 1)(X + i - \sqrt{2} - 1)(X - i + \sqrt{2} - 1)(X + i + \sqrt{2} - 1)$$

Proposition 5.8. *As usual let G be a finite subgroup of $\mathrm{Aut}(L)$, then L/K is finite and $[L : K] \leq |G|$, where $K = L^G$.*

Proof. Choose $\alpha \in L$ such that $K(\alpha)/K$ has maximal degree (which is bounded above by $|G|$). Let $\beta \in L$. We will show that $\beta \in K(\alpha)$. Now $[K(\alpha, \beta) : K(\alpha)] \leq [K(\beta) : K]$, so $[K(\alpha, \beta) : K] = [K(\alpha, \beta) : K(\alpha)][K(\alpha) : K]$ is finite and greater than $[K(\alpha) : K]$ unless $\beta \in K(\alpha)$. But it cannot be greater than $[K(\alpha) : K]$ by Theorem 3.11 and the maximal degree assumption on α, so $\beta \in K(\alpha)$. Consequently $[L : K] = \deg_K(\alpha) \leq |G|$. □

Theorem 5.9 (Artin)**.** *Let L be a field and $G \leq \mathrm{Aut}(L)$ is finite, then:*
(i) $[L : L^G] = |G|$.
(ii) $G = \mathrm{Aut}(L/L^G)$. In particular, L/L^G is Galois.

Proof. As usual we write $K = L^G$.
(i) Write $L = K(\gamma)$ for some γ. $[L : K] = \deg_K \gamma = |G\gamma|$, so it suffices to show that $H = \mathrm{Stab}_G(\gamma)$ is trivial, which is clear since it acts trivially on L.
(ii) $K \leq L^{\mathrm{Aut}(L/K)} \leq L^G$ as $G \leq \mathrm{Aut}(L/K)$. But $K = L^G$, so we have equality. $|\mathrm{Aut}(L/K)|$ is finite since L/K is finite, so by part (i) of the theorem we have $|\mathrm{Aut}(L/K)| = |G|$, therefore $G = \mathrm{Aut}(L/K)$. □

Alternatively, we could've proved part (i) of the theorem by listing $|G|$ distinct automorphisms of L that are linearly independent over K.

Example 5.6. Let $L = \mathbb{C}(y)$ and $G = \langle \sigma, \tau \rangle \cong (\mathbb{Z}/2\mathbb{Z})^2$ where $\sigma y = i/y$ and $\tau y = -y$. The orbit of y under G is $Gy = \{y, i/y, -y, -i/y\}$, so the minimal polynomial of y over $K = L^G$ is $f(X) = (X - y)(X - i/y)(X + y)(X + i/y) = X^2 - (y^2 - y^{-2})X^2 - 1$. In particular $\omega = y^2 - y^{-2} \in K$. We have $\deg_K(y) = \deg_{\mathbb{C}(\omega)}(y) = 4$, so actually $K = \mathbb{C}(\omega)$.

The example is related to the following theorem:

Proposition 5.10 (Lüroth's Theorem)**.** *If K is a field and $\mathbb{C} \subsetneq K \subset \mathbb{C}(y)$, then there is some $\omega \in \mathbb{C}(y)$ such that $K = \mathbb{C}(\omega)$.*

Proof. Algebraic geometry or otherwise. □

Back on with examples.

Example 5.7. Let $K = \mathbb{F}_q$ and $L = \mathbb{F}_{q^n}$, then $\mathrm{Aut}(L/K) = \mathbb{Z}/n\mathbb{Z}$. Indeed, the map $\phi(x) = x^q$ is an element of $\mathrm{Aut}(\mathbb{F}_{q^n}/\mathbb{F}_q)$, so we've got an injective group homomorphism $\mathbb{Z}/n\mathbb{Z} \hookrightarrow \mathrm{Aut}(L/K)$ via $i \mapsto \phi^i$. But $|\mathrm{Aut}(\mathbb{F}_{q^n}/\mathbb{F}_q)| \leq [\mathbb{F}_{q^n} : \mathbb{F}_q] = n$, so indeed $\mathbb{Z}/n\mathbb{Z} = \mathrm{Aut}(L/K)$.

Proposition 5.11. *Let L/K be a finite extension, then $|\mathrm{Aut}(L/K)| \mid [L : K]$ with equality iff L/K is Galois.*

Proof. Let $M = L^{\mathrm{Aut}(L/K)}$, then we have a chain of field extensions $L/M/K$. The proposition follows from Theorem 1.4 and Theorem 5.9. □

Theorem 5.12. *The followings are equivalent:*
(i) L/K is Galois.
(ii) There exists a finite group $G \leq \mathrm{Aut}(L)$ with $L^G = K$.
(iii) L is the splitting field of a separable polynomial.
(iv) L/K is finite, separable, and the minimal polynomial of each $\alpha \in L$ splits into linear factors in $L[X]$.

Proof. (i) $\implies$ (ii) is clear; (ii) $\implies$ (i) is Theorem 5.9; (ii) $\implies$ (iv) since we can take the linear factors to be the G-orbits of α.
(iv) $\implies$ (iii): Suppose $L = K(\alpha_1, \ldots, \alpha_n)$ and let f_i be the minimal polynomial (separable by hypothesis) of α_i. Then $\mathrm{lcm}(f_1, \ldots, f_n)$ is separable and L is the splitting field of it.
(iii) $\implies$ (i): Suppose L is the splitting field of a separable polynomial f, then there are exactly $[L : K]$ K-homomorphisms $L \to L$, hence $|\mathrm{Aut}(L/K)| = [L : K]$ which means that L/K is Galois by the preceding proposition. □

Corollary 5.13. *Any finite separable extension is contained in a Galois extension.*

Proof. Suppose L/K is finite and separable, then $L = K(\alpha_1, \ldots, \alpha_n)$ for some separable $\alpha_1, \ldots, \alpha_n$. Then L/K is contained in the splitting field N/K of the least common multiple of the minimal polynomials of α_i. □

It is easy to see that the constructed Galois extension N/K is indeed minimal in the sense that any field N' Galois over K which contains L/K minimally (i.e. no subfield of N' containing L can be Galois) would be L-isomorphic to N.

Corollary 5.14. *If $L/M/K$ is a chain of field extensions and L/K is Galois, then L/M is also Galois.*

Proof. Immediate from the last conditon in Theorem 5.12. □

5.3 Galois Correspondence

Theorem 5.15 (Galois Correspondence/Fundamental Theorem of Galois Theory)**.** *Suppose L/K is a Galois extension, then there is a one-to-one correspondence between subgroups of $Aut(L/K)$ and intermediate subfields of L/K (i.e. subfields of L containing K) given by $H \mapsto L^H$ with inverse $M \mapsto \mathrm{Aut}(L/M)$.*

In particular, there are only finitely many intermediate subfields. We've already seen some instance of this.

Example 5.8. 1. If $Aut(L/K) \cong \mathbb{Z}/p\mathbb{Z}$, then it has no proper subgroups, so Theorem 1.4 already shows that no intermediate fields can exist.
2. Let $L = \mathbb{F}_{q^n}, K = \mathbb{F}_q$, then $\mathrm{Aut}(L/K) \cong \mathbb{Z}/n\mathbb{Z}$. As usual let $\phi : x \mapsto x^q$ be its generator, then each subgroup of $\mathrm{Aut}(L/K)$ has the form $\langle p^m \rangle$ for $m \mid n$, which corresponds to $L^{\langle \phi^m \rangle} = \mathbb{F}_{q^m}$.
3. Let $L = \mathbb{Q})(i, \sqrt{2})$ and $K = \mathbb{Q}$, then $G \cong (\mathbb{Z}/2\mathbb{Z})^2 = \langle \sigma, \tau \rangle$ where $\sigma(a+bi) = a - bi, \tau(c + d\sqrt{2}) = c - d\sqrt{2}$. Then the subgroups $\langle \sigma \rangle, \langle \tau \rangle, \langle \sigma\tau \rangle$ correspond to $\mathbb{Q}(\sqrt{2}), \mathbb{Q}(i), \mathbb{Q}(i\sqrt{2})$ respectively.

Proof. $\mathrm{Aut}(L/L^H) = H$ by Theorem 5.9. Conversely, L/M is Galois by Corollary 5.14, so $L^{\mathrm{Aut}(L/M)} = M$. □

Proposition 5.16. *Suppose L/K is Galois.*
(i) The Galois correspondence is order-reversing: If $H \leq H'$ are subgroups of $\mathrm{Aut}(L/K)$, then $L^{H'} \leq L^H$; If $M \leq M'$ are intermediate subfields of L/K, then $\mathrm{Aut}(L/M') \leq \mathrm{Aut}(L/M)$.
(ii) $N \leq \mathrm{Aut}(L/K)$ is normal iff L^N/K is Galois. If this were the case, then $\mathrm{Aut}(L^N/K) \cong G/N$.

Proof. (i) is clear.
For (ii), first suppose that $N \trianglelefteq G$ is normal. For any $\sigma \in G$, $l \in L^N$ and $n \in N$, we have $n(\sigma l) = \sigma((\sigma^{-1} n \sigma)(l)) = \sigma l$, so σ fixes L^N. This gives a homomorphism $G \to \mathrm{Aut}(L^N/K)$ with N whose kernel is exactly N, so $G/N \hookrightarrow \mathrm{Aut}(L^N/K)$. This does hit everything by Theorem 5.9, so $G/N \cong \mathrm{Aut}(L^N/K)$ and L^N/K is Galois.
Conversely, for any chain $L/M/K$ of fields with $M = L^{\mathrm{Aut}(L/M)}$, we have $\mathrm{Aut}(L/\sigma M) = \sigma \,\mathrm{Aut}(L/M)\sigma^{-1}$ for $\sigma \in G$. Suppose L^N/K is Galois, then L^N is the splitting field of some separable $f \in K[X]$. So any $\sigma \in \mathrm{Aut}(L/K)$ would fix L^N since it permutes the roots of f, i.e. $\mathrm{Aut}(L/L^N) = \mathrm{Aut}(L/\sigma L^N) = \sigma \,\mathrm{Aut}(L/L^N)\sigma^{-1}$ for all $\sigma \in G$. This precisely means that $N = \mathrm{Aut}(L/L^N)$ is normal in $\mathrm{Aut}(L/K)$. □

Definition 5.4. Suppose $f \in K[X]$ is separable, the Galois group of f is $\mathrm{Gal}(f) = \mathrm{Aut}(L/K)$ where L is the splitting field of f over K.

$\mathrm{Gal}(f)$ is well-defined since splitting fields are unique up to K-isomorphisms. Computing it is non-trivial in general, but we can always "see" it.

Lemma 5.17. *Suppose $f \in K[X]$ is separable, then f is irreducible iff $\mathrm{Gal}(f)$ acts transitively on the roots of f.*

Proof. Let L be the splitting field of f and $\alpha_1, \ldots, \alpha_r$ the roots of f. Suppose $\{\alpha_1, \ldots, \alpha_r\} = X_1 \sqcup \cdots \sqcup X_s$ is the decomposition of the set of roots of f into $\mathrm{Gal}(f)$-orbits. Then each $f_i(X) = \prod_{\alpha \in X_i}(X - \alpha_i)$ would be in $K[X]$ (and is irreducible by Theorem 5.9). □

5.4 Cubics and Discriminants

Suppose $f(X) = X^3 + aX^2 + pX + q \in K[X]$ is irreducible and separable. Let L be the splitting field for f over K and suppose $\alpha_1, \alpha_2, \alpha_3$ are the roots of f (so $L = K(\alpha_1, \alpha_2, \alpha_3)$). We want to explore the possibilities for $\mathrm{Gal}(f)$ and the intermediate subfields between L and K.

We've got a chain $K \leq K(\alpha_i) \leq L$ for each i and we know that $K(\alpha_i)/K$ has degree 3.

If $[L : K] = 3$, then we have $K(\alpha_i) = L$ for each i. In this case, since 3 is prime, we also know that there is no proper subfields of L properly containing K.

If $[L : K] = 6$, then neither α_2 nor α_3 (say) would be in $K(\alpha_1)$.

What are the Galois groups in these two cases? Let $G = \mathrm{Gal}(f) \leq S_3$, then the subgroups of G corresponds to intermediate subfields of L/K. The subgroups of S_3 are $1, \langle(12)\rangle, \langle(23)\rangle, \langle(31)\rangle, \langle(123)\rangle$. By the preceding lemma, G has to be transitive on the three letters S_3 acts on, so G is either $A_3 = \langle(123)\rangle \cong C^3$ (which corresponds to the $[L : K] = 3$ case) or S_3 (the $[L : K] = 6$ case).

What happens in the latter case? $L^{\langle(12)\rangle}$, say, is intermediate between $K(\alpha_3)$ and L. But $[L : K] = 6, [L : L^{\langle(12)\rangle}] = 3$ and $[K(\alpha_3) : K] = 2$, so in fact $L^{\langle(12)\rangle} = K(\alpha_3)$. Similarly, $L^{\langle(23)\rangle} = K(\alpha_1)$ and $L^{\langle(31)\rangle} = K(\alpha_2)$. But what is $M = L^{\langle(123)\rangle}$? It must be a degree 2 extension of K (and is the unique degree 2 extension of K in L), which is actually quite surprising since we get this from a cubic.

Assume that $\mathrm{char}\, K \neq 2$, then $M = K(\sqrt{D})$ for some $D \in K$. How do we find D? An educated guess would be $D = \delta^2$ where $\delta = (\alpha_1 - \alpha_2)(\alpha_2 - \alpha_3)(\alpha_3 - \alpha_1)$. Indeed, $(12)\delta = -\delta \neq \delta$ and $(123)\delta = \delta$, so $\delta \in L^{\langle(123)\rangle} \setminus K$ and $D = \delta^2 \in L^{S_3} = K$. So $K(\sqrt{D}) = K(\delta) = L^{\langle(123)\rangle}$.

We can always determine such a value D even for general cubics (not necessarily separable nor irreducible), in which case $D = 0$ iff the polynomial is not separable. Also, we were interested in such a D only when $[L : K] = 6$. What can possibly happen if one try to adjoint δ to K when $[L : K] = 3$? As L/K would have no intermediate subfield, either $K(\delta) = L$ or $K(\delta) = K$. But δ is fixed by (123), so the only thing that can happen is $K(\delta) = K$, i.e. D is a square in K.

So here's our answer: Given a irreducible cubic f, we compute such a $D \in K$ and test its quality. If it's zero, then the cubic is not separable. If it's a square, then f would split in $K[X]/(f)$ and $\mathrm{Gal}(f) \cong A_3$. Otherwise, $\mathrm{Gal}(f) = S_3$.

These would be helpful if we can compute D without having to find the roots of f. We have

$$\begin{cases} \alpha_1 + \alpha_2 + \alpha_3 = -a \\ \alpha_1\alpha_2 + \alpha_2\alpha_3 + \alpha_3\alpha_1 = p \\ \alpha_1\alpha_2\alpha_3 = -q \end{cases}$$

So by expanding and some amount of calculations, we arrive at $D = a^2p^2 - 4p^3 - 4a^3q - 27q^2 + 18apq$. When $a = 0$, we get something nicer, namely $D = -4p^3 - 27q^2$. If $\mathrm{char}\, K \neq 3$, we can always put the cubic in the form with $a = 0$ by the substitution $g(X) = f(X - a/3)$ (and we of course have $\mathrm{Gal}(g) = \mathrm{Gal}(f)$).

Example 5.9. We work over $\mathbb{Q}$.

1. For $f(X) = X^3 - 3X + 1$, we have $D = (-4)(-27) - 27 = 3^4$, so $\mathrm{Gal}(f) = A_3$. We can also do this directly: The roots of f are $\omega + \omega^{-1}, \omega^2 + \omega^{-2}, \omega^4 + \omega^{-4}$ where $\omega = e^{2\pi i/9}$. They all reside in $\mathbb{Q}(\omega + \omega^{-1})$ which has degree 3 over $\mathbb{Q}$.

2. For $f(X) = X^3 + 3X + 1$, $D = -5 \times 27$ is not a square, which means that $\mathrm{Gal}(f) = S_3$.
3. For $f(X) = X^3 - 2$, $D = -3^3 \times 2^2$ is not a square, so $\mathrm{Gal}(f) = S_3$.
4. If f (is irreducible and) only has one real root, then adjoining that real root gives a proper intermediate subfield between $\mathbb{Q}$ and the splitting field of f, so $\mathrm{Gal}(f) \cong S_3$. One can also check directly that D is not a square in this case.

The construction of D is not exclusive to cubics. For quadratics, we see immediately by completing square that

Lemma 5.18. *Suppose $f(X) = X^2 + bX + c \in K[X]$ has roots α_1, α_2 in its splitting field. Let $D = (\alpha_1 - \alpha_2)^2 = b^2 - 4c$, then $\mathrm{Gal}(f) = S_2$ if D is not a square in K, and $\mathrm{Gal}(f) = \{e\} = A_2$ otherwise.*

In higher degrees, this is also true.

Proposition 5.19. *Suppose $f(X) = a_0 + \cdots + a_{n-1}X^{n-1} + X^n \in K[X]$ is irreducible and let L be its splitting field in which it has roots $\alpha_1, \ldots, \alpha_n$. Let $D = \prod_{i<j}(\alpha_i - \alpha_j)^2$, then:*
(i) $D \in K$.
(ii) $D = 0$ iff f is not separable.
(iii) D is a square iff $\mathrm{Gal}(f) \subset A_n$.
(iv) There is a polynomial $\Delta = \Delta_n \in K[X_{n-1}, \ldots, X_0]$ (the "discriminant") independent of f such that $D = \Delta(a_{n-1}, \ldots, a_0)$.

Proof. (i) and (ii) are clear.
For (iii), let $\delta = \prod_{i<j}(\alpha_i - \alpha_j) \in L$, then $\delta^2 = D$. Embed $G = \mathrm{Gal}(f) \hookrightarrow S_n$ by its action on the roots of f. For any $\sigma \in G$

$$\sigma\delta = \begin{cases} \delta & \text{if } \sigma \in A_n \\ -\delta & \text{if } \sigma \notin A_n \end{cases}$$

So $G \subset A_n$ iff $\forall \sigma \in G, \sigma\delta = \delta$ iff $\delta \in K$ iff D is a square in K.
We defer the proof of (iv) to the next section, where we develop general theory of symmetric polynomials. □

6 Symmetric Polynomials

Let R be a ring. S_n acts on $R[Z_1, \ldots, Z_n]$ via $w \cdot Z_i = Z_{w(i)}$. We are interested in the ring $R[Z_1, \ldots, Z_n]^{S_n} = \{f \in R[Z_1, \ldots, Z_n] : \forall w \in S_n, w \cdot f = f\}$ of symmetric polynomials.

Example 6.1. $e_0 = 1, e_1 = \sum_i Z_i, e_2 = \sum_{i<j} Z_iZ_j, e_3 = \sum_{i<j<k} Z_iZ_jZ_k, \ldots$ (the "elementary symmetric polynomials") and polynomials in them are all symmetric polynomials in $Z_1, \ldots, Z_n$.

Theorem 6.1. *The ring homomorphism $R[W_1, \ldots, W_n] \to R[Z_1, \ldots, Z_n]^{S_1}$ via $W_k \mapsto e_k$ is an isomorphism.*

That is, every symmetric polynomial in n variables can be written uniquely as a polynomial in $e_1, \ldots, e_n$.

Example 6.2. 1. $\sum_i Z_i^2 = e_1^2 - 2e_2$.
2. The Vieta formulas shows that the coefficients of a polynomial are (up to a sign) the elementary symmetric polynomials in the roots, so the last part of Proposition 5.19 follows from the theorem. In particular, as $(Z_1 - Z_2)^2 = e_1^2 - 4e^2$, $\Delta_2(b, c) = b^2 - 4c$. Similarly, $\Delta_3(0, p, q) = cp^3 + dq^2$ for some constants c, d by inspection. We have $27c = 4d$ by considering $f(X) = (X-\alpha)(X-\alpha)(X+2\alpha)$. By computing Δ_3 on a specific example, say $f(X) = X^3 - X$, we further conclude that $\Delta_3(0, p, q) = -4p^3 - 27q^2$.

We need some set-up before we prove the theorem. For a tuple of nonnegative integers $\lambda = (\lambda_1, \ldots, \lambda_n) \in \mathbb{N}^n$, we write $Z^\lambda = Z_1^{\lambda_1} \cdots Z_n^{\lambda_n}$.

Example 6.3. For $n = 3$, $Z^{(2,4,1)} = Z_1^2 Z_2^4 Z_3$.

We can totally order $\mathbb{N}^n$ lexicographically: $\lambda < \mu$ iff there is some $1 \leq k \leq n$ such that $\lambda_i = \mu_i$ for all $i < k$ and $\lambda_k < \mu_k$. In particular, if λ has $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n$, then $w(\lambda) \leq \lambda$ for any $w \in S_n$ (acting on the indices). So every orbit $S_n \cdot \lambda$ has a unique maximal element λ given by the one nonincreasing tuple. Consequently,

$$\left\{ \sum_{\mu \in S_n \cdot \lambda} Z^\mu : \lambda = (\lambda_1, \ldots, \lambda_n) \text{ nonincreasing} \right\}$$

is a basis for of $R[Z_1, \ldots, Z_n]^{S_n}$.

Proof of Surjectivity in Theorem 6.1. Any $f \in R[Z_1, \ldots, Z_n]^{S_n}$ has the form $f = cZ^\lambda + \sum_{\text{some } \mu < \lambda} Z^\mu$. Observe that

$$\begin{aligned} Z^\lambda &= Z_1^{\lambda_1 - \lambda_2} (Z_1 Z_2)^{\lambda_2 - \lambda_3} \cdots (Z_1 \cdots Z_{n-1})^{\lambda_{n-1} - \lambda_n} (Z_1 \cdots Z_n)^{\lambda_n} \\ &= e_1^{\lambda_1 - \lambda_2} \cdots e_{n-1}^{\lambda_{n-1} - \lambda_n} e_n^{\lambda_n} + \sum_{\text{some } \mu < \lambda} Z^\mu \end{aligned}$$

So $f - ce_1^{\lambda_1 - \lambda_2} \cdots e_{n-1}^{\lambda_{n-1} - \lambda_n} e_n^{\lambda_n}$ is a sum of terms in the form $Z^\mu, \mu < \lambda$. Continuing this algorithm gives a way to eventually write f in terms of the elementary symmetric polynomials, which gives surjectivity. □

Example 6.4. If you are bored you can try the algorithm on $\sum_{i \neq j} Z_i^2 Z_j$ to get $\sum_{i \neq j} Z_i^2 Z_j = e_1 e_2 - 3e_3$.

Proof of Injectivity in Theorem 6.1. Suppose $g \in R[W_1, \ldots, W_n]$ is such that $g(e_1, \ldots, e_n) = 0 \in R[Z_1, \ldots, Z_n]^{S_n}$. WLOG g has minimal degree. We want to show that $g = 0$. This is certainly true when $n = 1$. For $n > 1$, we proceed by induction. Under the ring homomorphism $R[Z_1, \ldots, Z_n] \to R[Z_1, \ldots, Z_{n-1}] \cong R[Z_1, \ldots, Z_n]/(Z_n)$ (given by the projection), e_i get mapped e_i° which is the i^{th} elementary symmetric polynomial but with $n - 1$ variables. Then $g(e_1^\circ, \ldots, e_{n-1}^\circ, 0) = 0$ in $R[Z_1, \ldots, Z_{n-1}]$, so by the induction hypothesis we know that $g(W_1, \ldots, W_n) = W_n h(W_1, \ldots, W_n)$ for some $h \in R[W_1, \ldots, W_n]$. But e_n is not a zero divisor in $R[Z_1, \ldots, Z_n]$, so $h(e_1, \ldots, e_n) = 0$, contradicting the minimal degree assumption on g. □

Lemma 6.2. *$K(Z_1, \ldots, Z_n)^{S_n}$ is the field of fractions of $K[Z_1, \ldots, Z_n]^{S_n}$.*

Proof. Write $L = K(Z_1, \ldots, Z_n)$. The field of fractions of $K[Z_1, \ldots, Z_n]^{S_n}$ is certainly contained in L^{S_n}. Conversely, let $\gamma = f/g \in L^{S_n}$, then we can consider $\mu = \prod_{\sigma \in S_n} \sigma g \in K[Z_1, \ldots, Z_n]^{S_n}$ which gives $\gamma\mu \in K[Z_1, \ldots, Z_n]^{S_n}$, so $\gamma = \gamma\mu/\mu$ is in the field of fractions of $K[Z_1, \ldots, Z_n]^{S_n}$. □

Remark. One can also show this directly (exercise).

Let $L = K(Z_1, \ldots, Z_n)$ because we got tired writing it, then L/L^{S_n} is Galois by Theorem 5.9.

Corollary 6.3. *Any finite group G is the Galois group of some polynomial.*

Proof. G embeds into $S_n = \mathrm{Aut}(L/L^{S_n})$ for some n (e.g. $n = |G|$). Theorem 5.9 shows that $G = \mathrm{Aut}(L/L^G)$ and L/L^G is Galois, which gives the result. □

7 Cyclotomic Extensions

We now turn to the study of polynomials of the form $X^n - \alpha$ for $\alpha \in K$. They are interesting for both historical (e.g. insolvability of quintics by radicals) and practical (e.g. classifying abelian Galois groups) reasons.
Let L be a field.

Definition 7.1. $\mu_n(L) = \{\alpha \in L : \alpha^n = 1\}$ is called the group of n^{th} roots of unity in L.

$\mu_n(L)$ has size at most n, and is cyclic as it's a finite subgroup of $L^\times$.

Definition 7.2. We say $\xi \in \mu_n(L)$ is a primitive n^{th} root of unity if it has order n (or, equivalently, $\mu_n(L)$ has order n and is generated by ξ).

The choice of such a primitive root is then the choice of an isomorphism $(\mathbb{Z}/n\mathbb{Z})^\times \to \mu_n(L), k \mapsto \xi^k$. Clearly, ξ^k is a primitive n^{th} root of unity iff $\gcd(k, n) = 1$.
When do primitive roots of unity exist? By definition, this happens exactly when $X^n - 1$ splits into distinct linear factors in L, i.e. $X^n - 1$ splits in L and $\mathrm{char}(L) \nmid n$.
Unless otherwise stated, we will assume that we work in fields with characteristics not dividing n throughout the rest of this section.

Definition 7.3. The n^{th} cyclotomic extension L of K is the splitting field of $X^n - 1$ over K.

By our assumption, $X^b - 1$ is always separable, so L/K is Galois.

Lemma 7.1. *Let $\xi \in \mu_n(L)$ be primitive, then*
(i) $L = K(\xi)$.
(ii) There is an injective group homomorphism $\chi : G = \mathrm{Aut}(L/K) \to (\mathbb{Z}/n\mathbb{Z})^\times$ given by $\chi(\sigma) = a$ if $\sigma\xi = \xi^a$. Moreover, χ does not depend on the choice of ξ.
(iii) χ is surjective iff G acts transitively on the n^{th} primitive roots of unity.

Clearly χ is not always surjective (e.g. when $X^n - 1$ already splits in K).

Proof. (i) $X^n - 1 = \prod_{\alpha \in \mu_n(L)}(X - \alpha) = (X-1)(X-\xi)\cdots(X-\xi^{n-1})$ in L.
(ii) If $\sigma \in G$, then $\sigma\xi$ has order exactly n, so χ is well-defined. It's clearly a group homomorphism. Injectivity follows from (i). Suppose $\xi' = \xi^k$ is a different primitive root of unity, then if $\chi(\sigma) = a$ (under ξ) then $\sigma\xi = \xi^a$, i.e. $\sigma(\xi') = \sigma(\xi^k) = \xi^{ka} = (\xi')^a$.
(iii) Obvious. □

Corollary 7.2. *$G = \mathrm{Aut}(L/K)$ is abelian and is canonically a subgroup of $(\mathbb{Z}/n\mathbb{Z})^\times$.*

Example 7.1. 1. Take $K = \mathbb{R}$ and $n > 2$, then $L = \mathbb{C}$ and $G = C_2$ with the conjugation automorphism being the generator, then χ takes conjugation to $-1 \in (\mathbb{Z}/n\mathbb{Z})^\times$.
2. Let $K = \mathbb{F}_q$ ($\gcd(q, n) = 1$ by assumption), then χ identifies G with $\langle q \rangle \leq (\mathbb{Z}/n\mathbb{Z})^\times$ since G is generated by $\phi_q(x) = x^q$.
3. If p is a prime and $\omega = e^{2\pi i/p}$, then $[\mathbb{Q}(\omega) : \mathbb{Q}] = \deg(X^{p-1} + \cdots + 1) = p - 1$, so $|G| = p - 1$ and therefore $\chi : \mathrm{Aut}(\mathbb{Q}(\omega)/\mathbb{Q}) \to (\mathbb{Z}/p\mathbb{Z})^\times$ is an isomorphism.

We will extend the last example and show that χ is always an isomorphism for the extension $\mathbb{Q}(e^{2\pi i/n})/\mathbb{Q}$ for all n.

Definition 7.4. Let L be the splitting field of $X^n - 1$ over K (with characteristic not dividing n, as usual) and ξ a primitive n^{th} root of unity. $\Phi_n(x) = \prod_{k \in (\mathbb{Z}/n\mathbb{Z})^\times}(X - \xi^k) \in L[X]$ is known as the n^{th} cyclotomic polynomial.

Φ_n clearly does not depend on the choice of ξ. Also, any $\sigma \in G = \mathrm{Aut}(L/K)$ permutes the set of primitive roots of unity, hence $\Phi_n \in L^G[X] = K[X]$.

Corollary 7.3. *The followings are equivalent:*
(i) Φ_n is irreducible.
(ii) χ is surjective.
(iii) $[L : K] = |G| = |(\mathbb{Z}/n\mathbb{Z})^\times|$.

We have $\Phi_n \mid X^n - 1$ by definition. We also have $X^d - 1 \mid X^n - 1$ for all $d \mid n$, and for every $\alpha \in \mu_n(L)$ there is exactly one $d \mid n$ such that α is a primitive d^{th} root of unity. Therefore we have the formula

$$X^n - 1 = \Phi_n(X) \prod_{d \mid n, d \neq n} \Phi_d(X)$$

which allows one to inductively produce the cyclotomic polynomials. In particular, all cyclotomic polynomials in $\mathbb{Q}$ have integer coefficients.

Example 7.2. We'll work in $\mathbb{Q}$.
$\Phi_1(X) = X - 1, \Phi_p(X) = (X^p - 1)/(X - 1), \Phi_{p^2}(X) = (X^{p^2} - 1)/(X^p - 1)$.
$X^8 - 1 = (X-1)(X+1)(X^2+1)(X^4+1), X^6 - 1 = (X-1)(X^2+X+1)(X+1)(X^2-X+1)$.
$\Phi_{105}(X) = X^{48} + X^{47} + \cdots - 2X^7 + \cdots + 1$, so not all coefficients are ± 1.

Theorem 7.4. *$\Phi_n(X) \in \mathbb{Q}[X]$ is irreducible.*

For $n = p^k$ one can use Eisenstein, but not in general.

Corollary 7.5. *χ is always an isomorphism when $K = \mathbb{Q}$.*

Proof of Theorem 7.4. If Φ_n is reducible, then $\Phi_n = fg$ for $f, g \in \mathbb{Z}[X]$ by Gauss' Lemma. We want to show that, if ξ is a root of f, so is ξ^k for all k with $\gcd(k, n) = 1$. It suffices to prove the case where $k = p$ is a prime, since we can apply the statement recursively.
Suppose otherwise, then $g(\xi^p) = 0$ for some prime $p \nmid n$, so ξ is both a root of $f(X)$ and $g(X^p)$. Let $h(X) = \gcd(f(X), g(X^p))$, then $\deg h > 0$.
We now reduce everything modulo p. Denote the reduction of f by $\bar{f}$, g by $\bar{g}$ and h by $\bar{h}$. We know that $\bar{g}(X^p) = (\bar{g}(X))^p$, so some nonconstant factor γ of $\bar{h}$ must divide both $\bar{f}$ and $\bar{g}$. But the reduction of Φ_n is separable in $\mathbb{F}_p$ as Φ_n divides $X^n - 1$ and $p \nmid n$, contradiction. □

The proof is amazing in the sense that tools we have in finite fields give rise to techniques on dealing with algebraic problems in $\mathbb{Q}$. There are many more instances of tricks like this as one will see in advanced contexts in algebraic number theory.
What are the subfields of $\mathbb{Q}(\xi_n)$? By Galois correspondence, all of them have the form $M = \mathbb{Q}(\xi_n)^H$ for some $H \leq \mathrm{Aut}(\mathbb{Q}(\xi_n)/\mathbb{Q}) \cong (\mathbb{Z}/n\mathbb{Z})^\times$. H is automatically normal as $(\mathbb{Z}/n\mathbb{Z})^\times$ is abelian, so $M/\mathbb{Q}$ must be Galois. We will try to understand the structure of $(\mathbb{Z}/n\mathbb{Z})^\times$ and the description of subfields corresponding to its subgroups.
When $n = p$ is a prime, $(\mathbb{Z}/p\mathbb{Z})^\times = \mathbb{F}_p^\times$ is cyclic of order $p - 1$, so there is a unique subgroup of every order dividing $p - 1$, i.e. for every $k | p - 1$ there exists a unique subfield of $\mathbb{Q}(\xi_p)$ of degree k over $\mathbb{Q}$.

Example 7.3. For $p = 5$, $(\mathbb{Z}/5\mathbb{Z})^\times \cong C_4$, so there is a unique subfield of $\mathbb{Q}(e^{2\pi i/5})$ of degree 2 over $\mathbb{Q}$, which is $\mathbb{Q}(\cos(2\pi/5)) = \mathbb{Q}(\sqrt{5})$.

Such phenomenon happens in general.

Lemma 7.6. *Let $p \geq 3$ be a prime. There is a unique subfield $M \subset \mathbb{Q}(\xi_p)$ with $[\mathbb{Q}(\xi_p) : M] = 2$ given by $M = \mathbb{Q}(\cos(2\pi/p)) = \mathbb{Q}(\xi_p + \xi_p^{-1})$.*

Proof. Clearly $[\mathbb{Q}(\xi_p) : \mathbb{Q}(\xi_p + \xi_p^{-1})] = 2$ since $\xi_p^2 - (\xi_p + \xi_p^{-1})\xi_p + 1 = 0$ and $\xi_p + \xi_p^{-1} \in \mathbb{R}$. □

Example 7.4. For $p = 7$, $(\mathbb{Z}/7\mathbb{Z})^\times \cong C_6 = \langle\sigma\rangle$. It has subgroups $\langle\sigma^2\rangle \cong C_3, \langle\sigma^3\rangle \cong C_2$. We know now that $\eta = \xi_7 + \xi_7^{-1}$ has degree $3 = 6/2$ over $\mathbb{Q}$. What is its minimal polynomial? Under $\langle\sigma\rangle$, η has conjugates $\xi_7 + \xi_7^{-1}, \xi_7^2 + \xi_7^{-2}, \xi_7^3 + \xi_7^{-3}$, so the minimal polyomial of η should be $(X - \xi_7 - \xi_7^{-1})(X - \xi_7^2 - \xi_7^{-2})(X - \xi_7^3 - \xi_7^{-3}) = X^3 + X^2 - 2X - 1$.
What is the subfield corresponding to $\langle\sigma^2\rangle$, then? It has degree 2 over $\mathbb{Q}$. If we add all the elements in the $\langle\sigma^2\rangle$-orbit of ξ, then we get $\epsilon = \xi_7 + \xi_7^2 + \xi_7^4$ which has minimal polynomial $X^2 + X + 2$. Hence $\mathbb{Q}(\epsilon) = \mathbb{Q}(\sqrt{-7})$ is an intermediate subfield with degree 2 over $\mathbb{Q}$.

Observe that $\mathbb{Q}(\xi_5) \supset \mathbb{Q}(\sqrt{5}), \mathbb{Q}(\xi_7) \supset \mathbb{Q}(\sqrt{-7})$. Is this true in general?

Proposition 7.7. *For any prime $p > 2$, the unique quadratic extension of $\mathbb{Q}$ contained in $\mathbb{Q}(\xi_p)$ is $\mathbb{Q}\left(\sqrt{(-1)^{(p-1)/2}p}\right)$.*

Proof. Let $G = (\mathbb{Z}/p\mathbb{Z})^\times = \langle\sigma\rangle$. Write $\xi = \xi_p$ and let $\alpha = \xi + \sigma^2\xi + \sigma^4\xi + \cdots + \sigma^{p-1}\xi$. Clearly $G\alpha = \{\alpha, \sigma\alpha\}$. $(X - \alpha)(X - \sigma\alpha) = X^2 + X + \alpha(\sigma\alpha)$. The proof follows from computing its discriminant (exercise). □

This is known as the technique of quadratic Gauss sum.

Alternative proof. For any polynomial $f \in K[X]$ with splitting field L, we have $K \subset K(\sqrt{\Delta}) \subset L$ where $\Delta = \Delta(f)$ is the discriminant of f. Also, $[K(\sqrt{\Delta}) : K] = 2$ iff $\sqrt{\Delta} \notin K$. Let $f(X) = X^p - 1 \in \mathbb{Q}[X]$, then $\Delta(f) = (-1)^{(p-1)/2} p^p$ which implies the proposition. □

Remark. This also shows that $\mathrm{Aut}(\mathbb{Q}(\xi_p)/\mathbb{Q})$ is not contained in A_{p-1} for odd prime p.

But wait, how did we compute $\Delta(f)$? It doesn't look trivial at all! Well, it will be after the following lemma.

Lemma 7.8. *If $f(X) = (X - \alpha_1) \cdots (X - \alpha_n)$, then*

$$\Delta(f) = (-1)^{n(n-1)/2} f'(\alpha_1) \cdots f'(\alpha_n)$$

Proof. Easy computation. □

Example 7.5. Take $p = 17$, then $(\mathbb{Z}/17\mathbb{Z})^\times = C_{16}$ and we should have a chain of subfields $\mathbb{Q}(\xi_{17}) \supset K_3 \supset K_2 \supset K_1 \supset \mathbb{Q}$ induced by $1 \subset C_2 \subset C_4 \subset C_8 \subset C_{16}$. Each K_i/K_{i-1} is a quadratic extension, so $K_i = K_{i-1}(\sqrt{\alpha_i})$ for some $\alpha_i \in K_{i-1}$. This then means that every $\alpha \in \mathbb{Q}(\xi_{17})$ is constructible by starightedge and compasses! In particular, $\cos(2\pi/17)$ is constructible, so it's possible to construct the regular 17-gon using starightedge and compasses – a famous result by Gauss.

Theorem 7.9 (Gauss)**.** *A regular n-gon is constructible iff $n = 2^k p_1 \cdots p_n$ with p_i distinct odd primes with the form $p_i = 2^{2^{l_i}} + 1$ ("Fermat primes").*

Worth noting that the only primes p of the form $2^n + 1$ are those with n a power of 2.

Remark. Examples of Fermat primes include $2, 5, 17, 257, 65537$, and none other is known. It is widely believed (but not proved) that there are only finitely many Fermat primes using the heuristic of thinking of primes as random numbers distributed according to the prime number theorem.

Proof. $[\mathbb{Q}(\xi_n) : \mathbb{Q}(\cos(2\pi/n))] = 2$, so $\cos(2\pi/n)$ is constructible iff ξ_n is constructible (i.e. there's a chain of degree 2 extensions leading from $\mathbb{Q}$ to $\mathbb{Q}(\xi_n)$). We have

$$[\mathbb{Q}(\xi_n) : \mathbb{Q}] = |(\mathbb{Z}/n\mathbb{Z})^\times| = \prod_i p_i^{e_i - 1}(p_i - 1)$$

So if ξ_n is constructible then n must have the desired form since $[\mathbb{Q}(\xi_n) : \mathbb{Q}]$ would have to be a power of 2. The converse follows from the fact that every abelian group of order 2^K contains a subgroup of order 2^{K-1} (which can be done either directly or by using the structure theorem of finite abelian groups). □

We in fact have, for $n = \prod_p p^{e_p}$,

$$(\mathbb{Z}/n\mathbb{Z})^\times \cong \mathbb{Z}/2^{(e_2-2)1_{e_2 \geq 2}}\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z} \times \prod_{p>2} (\mathbb{Z}/p^{e_p - 1}\mathbb{Z} \times \mathbb{Z}/(p-1)\mathbb{Z})$$

Dirichlet's theorem on primes in arithmetic progressions shows in particular that every finite abelian group is a quotient of $(\mathbb{Z}/n\mathbb{Z})^\times$ for some n, so every abelian group is the Galois group of a polynomial over $\mathbb{Q}$.

8 Kummer Theory

We want to study the splitting field of $X^n - \theta$ for some $\theta \in K \setminus \{0\}$.

8.1 Kummer Extensions

Proposition 8.1. *Let L be a splitting field of $X^n - \theta$ with $\theta \in K \setminus \{0\}$. Suppose* $\operatorname{char} K$ *does not divide n, then*
(i) L contains a primitive n^{th} root of unity ξ.
(ii) $\operatorname{Aut}(L/K(\xi)) \cong \mathbb{Z}/d\mathbb{Z}$ *for some $d \mid n$. Moreover, $X^n - \theta$ is irreducible in $K(\xi)$ iff $n = d$.*

Proof. (i) Our characteristic assumption means that $X^n - \theta$ is separable. Let $\alpha_1, \dots, \alpha_n \in L$ be the roots of $X^n - \theta$. But then $(\alpha_i/\alpha_j)^n = \alpha_i^n/\alpha_j^n = \theta/\theta = 1$, so $\{\alpha_i/\alpha_1 : 1 \le i \le n\} \subset \mu_n(L)$. But $|\mu_n(L)| \le n$, so $|\mu_n(L)| = n$.
(ii) The roots of $X^n - \theta$ can be written as $\alpha, \alpha\xi, \alpha\xi^2, \dots, \alpha\xi^{n-1}$ for $\alpha = \alpha_1$. This means that $L = K(\alpha, \xi)$. Consider the map $\chi : \operatorname{Aut}(L/K(\xi)) \to \mathbb{Z}/n\mathbb{Z}$ that sends $\sigma \in \operatorname{Aut}(L/K(\xi))$ to j with $\sigma\alpha = \alpha\xi^j$. This is clearly an injective homomorphism, so χ identifies $\operatorname{Aut}(L/K(\xi))$ with a subgroup of $\mathbb{Z}/n\mathbb{Z}$, which must be isomorphic to $\mathbb{Z}/d\mathbb{Z}$ for some $d \mid n$. $X^n - \theta$ is irreducible over $K(\xi)$ iff $\operatorname{Aut}(L/K(\xi))$ acts transitively on the roots of $X^n - \theta$, which is just saying that $d = |\operatorname{Aut}(L/K(\xi))| = n$. □

Example 8.1. $X^6 + 3 \in \mathbb{Q}[X]$ is irreucible by Eisenstein, so $L = \mathbb{Q}(\sqrt[6]{3}, \xi_6)$. $\Phi_6 = X^2 - X + 1$, so $\xi_6 = (1/2)(1 + \sqrt{-3})$, i.e. $\mathbb{Q}(\xi_6) = \mathbb{Q}(\sqrt{-3})$. $X^6 + 3$ is not irreducible over $\mathbb{Q}(\sqrt{-3})$, so χ is not surjective and in fact takes $\operatorname{Aut}(L/\mathbb{Q}(\sqrt{-3}))$ to a cyclic group of order 3.

$K(\xi)/K$ is Galois, so $K(\xi) = K(\alpha, \xi)^N$ where $N = \operatorname{Aut}(K(\alpha, \xi)/K(\xi))$ is normal in $G = \operatorname{Aut}(K(\alpha, \xi)/K)$.
G does not have to be abelian but both N and G/N (i.e. $\operatorname{Aut}(K(\xi)/K)$) are, as the former is cyclic and the latter is a subgroup of $(\mathbb{Z}/n\mathbb{Z})^\times$. This is a very curious structure. We've seen some groups of this sort before.

Example 8.2. (i) Take any field K and consider

$$G = \left\{ \begin{pmatrix} \mu & a \\ 0 & 1 \end{pmatrix} : \mu, a \in K, \mu \neq 0 \right\}, N = \left\{ \begin{pmatrix} 1 & a \\ 0 & 1 \end{pmatrix} : a \in K \right\}$$

Then N is abelian (in fact isomorphic to $(K, +)$) and normal in G with abelian quotient (isomorphic to $(K^\times, \times)$)).
(ii) Consider the semidirect product $G = \mathbb{Z}/p\mathbb{Z} \rtimes (\mathbb{Z}/p\mathbb{Z})^\times$ via $(a, \mu)(b, \gamma) = (a + \mu b, \mu\gamma)$ and $N = \mathbb{Z}/p\mathbb{Z} \lhd G$, then both N and G/N are abelian. In fact, this is just the first example with $K = \mathbb{F}_p$.

Corollary 8.2. *Suppose K contains a primitive n^{th} root of unity. Let L be the splitting field of $X^n - \theta$ (with the usual assumption on characteristic). Then* $\operatorname{Aut}(L/K)$ *is a cyclic group of order $d \mid n$.*

This just follows directly from the preceding proposition. The amazing thing is that the converse too is true.

Theorem 8.3. *Let L/K be Galois with $\mathrm{Aut}(L/K) \cong \mathbb{Z}/n\mathbb{Z}$ (a "cyclic extension"). Suppose K contains a primitive n^{th} root of unity, then there is some $\theta \in K$ such that $X^n - \theta$ is irreducible over K, splits in L and $L = K(\alpha)$ for some $\alpha \in L$ with $\alpha^n = \theta$.*

Note that if K contains a primitive n^{th} root of unity then $\mathrm{char}\, K \nmid n$.

Proof. Let σ be a generator of $\mathrm{Aut}(L/K)$. $\sigma : L \to L$ is an isomorphism of K-vector space. Since $\sigma^n = 1$ (where $n = [L : K]$), all eigenvalues of σ are n^{th} roots of unity, which are in K by assumption. Suppose one of the eigenvalues ξ is not 1, then any of its eigenvector $\alpha \in L$ has $\sigma\alpha = \xi\alpha$ which means that $\alpha \notin K = L^{\langle\sigma\rangle}$ (which, by the way, is the 1-eigenspace of σ). But $\alpha^n \in K$ since $\sigma(\alpha^n) = \xi^n\alpha^n = \alpha^n$. If in addition that ξ is primitive, then $X^n - \theta = (X - \alpha)(X - \xi\alpha)\cdots(X - \xi^{n-1}\alpha)$ splits in L and is irreducible over K as $\{\alpha, \xi\alpha \ldots, \xi^{n-1}\alpha\}$ is a G-orbit. $L = K(\alpha)$ by looking at degrees. □

The proof is not complete – we have yet to demonstrate that there indeed exists an eigenvalue ξ of σ that is a primitive n^{th} root of unity. This is not hard if we assume σ is diagonalisable: It's eigenvalues must form a subgroup of $\mu_n(K)$. If it contains no primitive n^{th} root of unity, then the subgroup is proper and hence σ cannot possibly have order n.
The diagonalisability of σ, of course, deserves some attention. As K contains all eigenvalues of σ, we can choose a basis under which σ is in Jordan normal form. But no nontrivial Jordan block can have order n since $\mathrm{char}\, K \nmid n$, so σ is essentially diagonal.
Alternatively, one can find the desired eigenvector directly. For any $\xi \in \mu_n(K)$, consider $p_\xi : L \to L$ via $p_\xi(x) = x + \xi^{-1}\sigma x + \xi^{-2}\sigma^2 x + \cdots + \xi^{-n+1}\sigma^{n-1}x$. Then for all $x \in L$ we have $\sigma p_\xi(x) = \xi p_\xi(x)$, so $p_\xi(x)$ is either zero or an eigenvector for σ with eigenvalue ξ. But there must be some $x \in L$ with $p_\xi(x) \neq 0$ by Theorem 5.3, so we are done.
What is this p_ξ anyways? In fact, it is exactly the linear projection onto the ξ-eigenspace. So we can get the diagonalisability of σ directly from this construction and finish the proof with our first argument (i.e. without using Theorem 5.3).
We can reinterpret this using the language of representation theory, which sadly does not constitute a proof as we'll use a result whose proof depends on the theorem. Whenever a field extension L/K is Galois, there is always some $\alpha \in L$ such that $\{g\alpha : g \in G = \mathrm{Aut}(L/K)\}$ is a K-basis of L (Theorem 10.5). So L is exactly the regular representation of G. If $p \nmid |G|$ and K contains the primitive $|G|^{th}$ rots of unity,

$$L \cong \bigoplus_{V \text{ irreducible representation of } G} (\dim V)V$$

as G-representations. This is what was actually happening with these p_ξ's.

8.2 Cubics Revisited

Let $f(X) = X^3 + pX + q \in K[X]$ with $\mathrm{char}\, K > 3$. Recall that its discriminant is $\Delta = \Delta(f) = -4p^3 - 27q^2$. Let $\delta = \sqrt{\Delta}$. Suppose L is the splitting field of f where f has roots $\alpha_1, \alpha_2, \alpha_3$. Let's find these roots explicitly!

Suppose $\xi_3 = (1/2)(-1+\sqrt{-3}) \in K$ (equivalently $\sqrt{-3} \in K$), $\delta = \sqrt{\Delta} \in K$ and $f(X) \in K[X]$ is irreducible. $\mathrm{Aut}(L/K) \cong A_3$ as $\delta \in K$, so $L = K(\sqrt[3]{\theta})$ as $\xi_3 \in K$ by Theorem 8.3. Consequently, any element of L is a K-linear combination of $1, \sqrt[3]{\theta}, (\sqrt[3]{\theta})^2$. So we can write the roots of f in terms of radicals. In fact, this is true even without the assumption that K contains δ and $\sqrt{-3}$, since they both are radicals in K.
In the case $\delta, \sqrt{-3} \in K$, we want to find an explicit formula to solve the cubic. Suppose $\mathrm{Aut}(L/K)$ is generated by σ and α_1 is a root of f. Then we can set $\alpha_2 = \sigma^{-1}\alpha_1, \alpha_3 = \sigma^{-2}\alpha_1$. Our previous discussion inspires us to consider

$$\beta = p_\xi(\alpha_1) = \alpha_1 + \xi\alpha_2 + \xi^2\alpha_3, \gamma = \alpha_1 + \xi^2\alpha_2 + \xi\alpha_3$$

Then $1, \beta, \gamma$ are all nonzero and linearly independent. We also have the formulae

$$\alpha_1 = \frac{\beta+\gamma}{3}, \alpha_2 = \frac{\xi^2\beta + \xi\gamma}{3}, \alpha_3 = \frac{\xi\beta + \xi^2\gamma}{3}$$

So finding the roots is just finding β and γ.
We have $\theta = \beta^3 \in K$ and $L = K(\beta)$. We claim that $\beta\gamma = -3p$ and β^3, γ^3 are roots of $X^2 + 27qX - 27p^3$, which would combine to give an explicit radical formula. Indeed,

$$\begin{aligned}\beta\gamma &= (\alpha_1 + \xi\alpha_2 + \xi^2\alpha_3)(\alpha_1 + \xi^2\alpha_2 + \xi\alpha_3) \\ &= \alpha_1^2 + \alpha_2^2 + \alpha_3^2 + (\xi + \xi^2)\sum_{i<j}\alpha_i\alpha_j \\ &= (\alpha_1 + \alpha_2 + \alpha_3)^2 - 3p = -3p\end{aligned}$$

and

$$\beta^3 + \gamma^3 = \beta^3 + \gamma^3 + (\alpha_1 + \alpha_2 + \alpha_3)^3 = 3(\alpha_1^3 + \alpha_2^3 + \alpha_3^3) + 18\alpha_1\alpha_2\alpha_3 = -27q$$

Yay! We've solved cubics!

8.3 Solving Equation by Radicals

We now arrive at the theorem you're probably looking for when you decided to join the course.

Theorem 8.4 (Galois)**.** *Let $f \in K[X]$ with either $\mathrm{char}\, K = 0$ or $\mathrm{char}\, K > \deg f$. Suppose L is the splitting field of f, then f is solvable by radicals iff $\mathrm{Gal}(f)$ is a solvable group.*

We haven't quite defined some concepts here, but let's see an example.

Example 8.3. $X^5 + 2X + 6$ has $\mathrm{Gal}(f) \cong S_5$ which is not solvable (a notion we will define in a moment), so supposedly it cannot be solved by radicals by the theorem.

Remark. Abel showed in 1826 that a generic quintic is not solvable by radicals, but an explicit counterexample is not found until Galois gave this more precise criterion.

Definition 8.1. L/K is an extension by radicals if there is a chain $K = L_0 \subset L_1 \subset \cdots \subset L_r = L$ such that $L_i = L_{i-1}(\sqrt[n_i]{\theta_i})$ for some $\theta_i \in L_{i-1}$. L/K is said to be contained in an extension by radicals if $L \subset L'$ and L'/K is an extension by radicals.
$f \in K[X]$ is solvable by radicals if the splitting field of f is contained in an extension by radicals.

Example 8.4. Any quadratic or cubic is solvable by radicals (for char $K \neq 2, 3$).

Lemma 8.5. *Suppose L/K is Galois and K contains a primitive $[L : K]^{th}$ root of unity. If $G = \mathrm{Aut}(L/K)$ is abelian, then L/K is an extension by radicals.*

Recall that if K contains a primitive n^{th} root of unity, then it contains a primitive d^{th} root of unity for every $d \mid n$.

Proof. We proceed by induction on $|G| = [L : K]$. This is clear if $|G| = 1$. For $|G| > 1$, as G is abelian, it contains a proper subgroup N with G/N cyclic. Induction hypothesis shows that L/L^N is an extension by radicals. On the other hand, L^N/K is a cyclic extension and K contains a $|G/N|^{th}$ root of unity, so by Theorem 8.3 there is some $\theta_1 \in K$ with $L^N = K(\sqrt[n_1]{\theta_1})$. □

Corollary 8.6. *Suppose L/K is Galois and $G = \mathrm{Aut}(L/K)$. If G is abelian and* $\mathrm{char}\, K \nmid |G|$*, then L/K is contained in an extension by radicals.*

Proof. Let ξ be a primitive $|G|^{th}$ root of unity. We have the extension chain $K \subset K(\xi) \subset L(\xi)$. The first inclusion is certainly an extension by radicals. What about the second one? By the lemma, it suffices to show that $L(\xi)/K(\xi)$ is Galois and $\mathrm{Aut}(L(\xi)/K(\xi))$ injects into $\mathrm{Aut}(L/K)$. Suppose L is the splitting field of a separable $f \in K[X]$, then $L(\xi)$ is the splitting field of f over $K(\xi)$ which means the extension is Galois. If $\sigma \in \mathrm{Aut}(L(\xi)/K(\xi))$, then σ fixes f, hence the set of its roots. But then σ essentially descends to a K-automorphism of L. This is an injection since $L(\xi) = K(\xi)(\alpha_1, \ldots, \alpha_n)$ where $(\alpha_i)_i$ are the roots of f in L. □

Observe that we did not seem to have used the full power of G being abelian in this whole discussion. This works for a more general class of groups, which are called – you guessed it, solvable groups.

Definition 8.2. A finite group G is solvable if and only if it exists a chain of subgroups $1 = G_0 \leq G_1 \leq \cdots \leq G_r = G$ such that $G_i \lhd G_{i+1}$ and G_{i+1}/G_i is abelian for every i.

Immediately, any abelian group is solvable.

Lemma 8.7. *We can replace the condition of G_{i+1}/G_i being abelian by G_{i+1}/G_i being cyclic.*

Proof. Any finite abelian group is the product of cyclic groups. □

Lemma 8.8. *(i) If G is solvable, so is any subgroup of G.*
(ii) Suppose G is a finite group and $N \lhd G$, then G is solvable iff $N, G/N$ are solvable.

Proof. Exercise. □

Example 8.5. 1. Abelian groups are solvable.
2. $\mathbb{Z}/n\mathbb{Z} \rtimes (\mathbb{Z}/n\mathbb{Z})^\times$ is solvable.
3. If L is the splitting field of $X^n - \theta \in K[X]$, then $\mathrm{Aut}(L/K)$ is solvable.
4. The group of invertible upper-triangular matrices over a finite field is solvable.
5. S_4 is solvable (as we'll see later).
6. Nonabelian finite simple groups are not solvable.
7. S_5 is not solvable.

Repeating our previous arguments but on solvable groups shows that if $\mathrm{Aut}(L/K)$ is solvable then f is solvable by radicals (where L is the splitting field of $f \in K[X]$).
Conversely, suppose f is solvable by radicals, then we get a chain $K = K_{-1} = K_0' \subset K_1' \subset \cdots \subset K_r'$ where $K_i' = K_{i-1}'(\beta_i)$ with $\beta_i = \sqrt[n_i]{\theta_i}, \theta_i \in K_{i-1}'$ and f splits completely in K_r'.
Let $d = \mathrm{lcm}(n_1, \ldots, n_r)$ and let ξ be a primitive d^{th} root of unity of L. Write $K_0 = K_0'(\xi)$ and inductively $K_i = K_{i-1}(\beta_i)$ for $i \geq 1$ to modify the chain into $K = K_{-1} \subset K_0 \subset \cdots K_r = M$.
If M/K is Galois then we are essentially done: Let $G = \mathrm{Aut}(M/K)$ and L be the splitting field of f. $\mathrm{Aut}(L/K)$ is a quotient of G since L/K is Galois, so the solvability of it follows from that of G, Let $G_i = \mathrm{Aut}(M/K_i)$, then $1 \leq G_r \leq G_{r-1} \leq \cdots \leq G_0 = G$. For $i \geq 1$, the extension $K_i = K_{i-1}(\beta_i) \supset K_{i-1}$ is the splitting field of $X^{n_i} - \theta_i$ over K_{i-1} as $\xi \in K_0$, hence K_i/K_{i-1} is Galois and has cyclic Galois group. But this is just saying that $G_i \lhd G_{i-1}$ with G_{i-1}/G_i cyclic. For $i = 0$, K_0/K_{-1} is cyclotomic, hence is Galois and has abelian automorphism group. These give the solvability of G by Theorem 5.15. Lemma 8.8 then shows that $\mathrm{Gal}(f)$ must also be solvable.
Sadly, M/K might not be Galois. $\mathbb{Q} \subset \mathbb{Q}(\sqrt{2}) \subset \mathbb{Q}(\sqrt[4]{2})$ is a counterexample. We save this by, you guessed it, extending the field even further.

Proposition 8.9. *Let M/K be an extension by radicals. Then there is an extension N/M by radicals such that N/K is Galois.*

Once we've known this, we can extend K_r' until it's Galois over K. If we then construct $M = K_r'(\xi)$ as before, it will have to be Galois over K as well and the rest of the argument simply goes through.
The proposition can be further reduced to the following lemma.

Lemma 8.10. *Suppose L/K is Galois and K contains a primitive n^{th} root of unity ξ. Let $M = L(\beta)$ where $\beta^n = \theta \in L$, then there is an extension $M \subset N$ by radicals such that N/K is Galois.*

Proof. Let $G = \mathrm{Aut}(L/K)$ and $h(x) = \prod_{\sigma \in G}(X^n - \sigma\theta) \in K$, then just take N to be the splitting field of fh over K (where f is such that L is the splitting field of f over K). □

9 Quartics

Enough of questions answered by Galois theory, let's now focus on those that are asked by Galois theory. For $f \in K[X]$, $\mathrm{Gal}(f)$ sure tells us a lot of information about f: If $\mathrm{Gal}(f)$ is solvable, then we know how to find its roots from Kummer theory. In general, for $K = \mathbb{Q}$, the behaviour of $\mathrm{Gal}(f)$ tells us a lot about the

number-theoretic properties of f.
In this section, we'll study the Galois group of a quartic polynomial. Suppose $f \in K[X]$ has degree 4 and is irreducible. We'll assume $\operatorname{char} K > 3$ so that f is separable. Let L be its splitting field (in which f has roots, say, $\alpha_1, \alpha_2, \alpha_3, \alpha_4$) and $G = \operatorname{Gal}(f) = \operatorname{Aut}(L/K)$.
There are a lot of possible Galois groups of f. We so far only know it's a transitive subgroup of S_4, and there are a lot of them:

$$S_4, A_4, D_8 = \langle (1234), (12)(34) \rangle, C_4 = \langle (1234) \rangle, V = \langle (12)(34), (13)(24) \rangle$$

Note that the copies of C_4 in S_4 are all transitives (and are all conjugates), same for D_8. However, there are non-transitive subgroups of S_4 isomorphic to V, e.g. $\langle (12), (34) \rangle$.
How would we know which one f is? As usual, let's look at the discriminant $\Delta = \Delta(f)$. Recall Proposition 5.19 which immediately shows that Δ is a square iff $\operatorname{Gal}(f)$ is either V or A_4 (so Δ is not a square iff $\operatorname{Gal}(f)$ is one of C_4, D_8 and S_4).
All of these groups are actually solvable. The solvability of S_4 is given by the chain $1 \leq V \leq A_4 \leq S_4$, so all subgroups G of it are also solvable with the chain $1 \leq G \cap V \leq G \cap A_4 \leq G$. Taking the respective fixed fields gives $L \supset L^{G \cap V} = M \supset L^{G \cap A_4} = K(\delta) \supset L^G = K$ (where as usual $\delta = \sqrt{\Delta}$). It seems that we need to study M (which is either $K(\delta)$ or $K(\delta, \sqrt[3]{\theta})$ for some $\theta \in K(\delta)$) in order to solve the quartic, so let's do that.
For $l \in L$, we can identify an element of M associated with it via $l \mapsto \sum_{g \in G \cap V} gl$. We want to use this identification to obtain some "partially symmetric functions of $\alpha_1, \ldots, \alpha_4$". Put $\beta = \alpha_1 + \alpha_2 = -\alpha_3 - \alpha_4, \gamma = \alpha_1 + \alpha_3 = -\alpha_2 - \alpha_4, \epsilon = \alpha_1 + \alpha_4 = -\alpha_2 - \alpha_3$. We can recover the roots from these as e.g. $2\alpha_1 = \beta + \gamma + \epsilon$. Now

$$\begin{cases} \beta^2 = -(\alpha_1 + \alpha_2)(\alpha_3 + \alpha_4) \\ \gamma^2 = -(\alpha_1 + \alpha_3)(\alpha_2 + \alpha_4) \\ \epsilon^2 = -(\alpha_1 + \alpha_4)(\alpha_2 + \alpha_3) \end{cases}$$

So it becomes clear that V fixes β^2, γ^2 and ϵ^2. Also, G fixes the set $\{\beta^2, \gamma^2, \epsilon^2\}$. But $\beta^2, \gamma^2, \epsilon^2$ are clearly all distinct, so $K(\beta^2, \gamma^2, \epsilon^2) = M$ by Theorem 5.9.

Definition 9.1. The polynomial $g(X) = (X - \beta^2)(X - \gamma^2)(X - \epsilon^2) \in K[X]$ is called the resolvant cubic.

The coefficients are symmetric functions in $\alpha_1, \alpha_2, \alpha_3, \alpha_4$, so we can in fact compute an explicit formula for g in terms of the coefficients of our original quartic. Indeed, we can expand to get $g(X) = X^3 + 2pX^2 + (p^2 - 4r)X - q^2$ if $f(X) = X^4 + pX^2 + qX + r$ (note that every quartic can be put into this form by a simple substitution).
g is a cubic, which means that we know how to solve it (hence f) by radicals. So we have solved quartics!
Except we haven't, yet. We still want to know about the Galois group of f. But this is not hard: $M = K(\delta)$ iff g is reducible iff G fixes at least one of $\beta^2, \gamma^2, \epsilon^2$ iff $G \cap V = G \cap A_4$ iff G is one of D_8, C_4 or V iff $3 \nmid |G \cap A_4|$ iff $3 \nmid |G|$. In other words, g is irreducible iff G is A_4 or S_4. Combining this with the information we know about the discriminant, we know that

	$\delta \in K$	$\delta \notin K$
g reducible	V	D_8, C_4
g irreducible	A_4	S_4

One can also decide between D_8 and C_4 in the case where g is reducible and $\delta \notin K$ by asking whether $\beta, \gamma, \epsilon \in M$. Consider the vector space spanned by $\alpha_1, \ldots, \alpha_4$ such that $\alpha_1 + \alpha_2 + \alpha_3 + \alpha_4 = 0$ (i.e. $\bigoplus_i K\alpha_i/K(\alpha_1 + \cdots + \alpha_n)$). This is a representation of S_n in the natural way. If $H \subset S_n$ is a subgroup, we can consider the map $\rho \to \rho^H$ given by $l \mapsto \sum_{g \in H} gl$. Then one can think about what happens when we take $H = V$.

Okay, let's now continue with our almost finished quest of understanding quartics. The question now is when can you express $\alpha = \sqrt{r + s\sqrt{t}}$ as a sum of unnested square roots.

α is the root of $f(X) = X^4 - 2rX^2 + (r^2 - s^2t)$. If f is reducible, then α is clearly unnested. Suppose now that f is irreducible. Let L be the splitting field of f. Since the roots of f are obtained by iterating square roots, we have $[L : K] = 1, 2, 4, 8$, so $G = V, D_8$ or C_4. If α is the sum of unnested square roots, then $\alpha \in K(\sqrt{a_1}, \ldots, \sqrt{a_r})$ for some $a_1, \ldots, a_r \in K$. This is Galois over K, so the splitting field L of α over K is a subfield of it. So we essentially have $L = K(\sqrt{a}, \sqrt{b})$ for some $a, b \in \{a_1, \ldots, a_r\}$, since the Galois group of $K(\sqrt{a_1}, \ldots, \sqrt{a_r})$ is $(C_2)^n$ whose subgroups are easy to spot. This means that $\mathrm{Aut}(L/K) = V$. Conversely, if $\mathrm{Aut}(L/K) = V$ then L/K has to be biquadratic, so α is unnested.

How do we decide whether $\mathrm{Aut}(L/K) = V$ then? Since we already know that f cannot have A_4 as its Galois group, this happens if and only if $\delta \in K$. The roots of f are $\pm\sqrt{r \pm s\sqrt{t}}$, so $\Delta = \prod_{i<j}(\alpha_i - \alpha_j)^2 = 2^8 s^4 t^2 (r^2 - s^2 t)$ which is a square iff $r^2 - s^2 t$ is.

Example 9.1. $\sqrt{3 + 2\sqrt{2}} = 1 + \sqrt{2}$, $\sqrt{5 + \sqrt{21}} = (1/2)(\sqrt{6} + \sqrt{14})$.

One can slightly extend this analysis to explicitly write down the actual expansion (exercise),

10 Miscellany

10.1 Reduction modulo p

Suppose $f \in \mathbb{Z}[X]$. Then we can either view f as a polynomial in $\mathbb{Q}[X]$ by inclusion, or in $\mathbb{F}_p$ by reduction where p is a prime.

We can get some information about f as a polynomial in $\mathbb{Q}$ by looking at it as a polynomial in $\mathbb{F}_p$. For instance, if f is irreducible in some $\mathbb{F}_p$, then f must be irreducible in $\mathbb{Z}$. We can sometimes also obtain information by reducing f modulo different primes p, e.g. $f(X) = X^4 + 5X^2 - 2X - 3$ factorises modulo 2 as $(X^2 + X + 1)^2$ and modulo 3 as $X(X^3 - X + 1)$, which essentially means that f is irreducible.

It shouldn't be completely surprising to postulate a relation between the Galois groups of f in $\mathbb{Q}$ and in $\mathbb{F}_p$. It's clear that an irreducible polynomial in $\mathbb{F}_p$ has Galois group (over $\mathbb{F}_p$) isomorphic to $C_{\deg f}$. In general, $f = h_1 \cdots h_r$ with each h_r irreducible in $\mathbb{F}_p$ has Galois group over $\mathbb{F}_p$ isomorphic to $C_{\deg h_1} \times \cdots C_{\deg h_r}$.

Theorem 10.1 (Dedekind)**.** *Let $f(X) \in \mathbb{Z}[X]$ be monic and p is a prime. Suppose f has distinct irreducible factors $h_1, \dots, h_r$ in $\mathbb{F}_p$ with each h_i irreducible of degree d_i, then the Galois group of f over $\mathbb{Q}$ contains a permutation with cycle type $d_1 \cdots d_r$.*

Proof. Omitted. □

Example 10.1. The Galois group of $f(X) = X^4 + 5X^2 - 2X - 3$ contains a 3-cycle. Indeed, since f is irreducible, the group is either A_4 or S_4.

10.2 Trace and Norm

Let L/K be a finite extension. Any $\alpha \in L$ gives a K-linear map $m_\alpha : L \to L, x \mapsto \alpha x$ which is an isomorphism when $\alpha \neq 0$.

Definition 10.1. The trace map of the extension is $\mathrm{Tr}_{L/K} : L \to K, \alpha \mapsto \operatorname{tr} m_\alpha$ and the norm is $\mathrm{N}_{L/K} : L \to K, \alpha \mapsto \det m_\alpha$.

Example 10.2. Suppose $L = K(\beta)$ with $\beta = \sqrt{D} \notin K$, then L has K-basis $\{1, \beta\}$ under which $m_{a+b\beta}$ has matrix

$$\begin{pmatrix} a & bD \\ b & a \end{pmatrix}$$

Consequently $\mathrm{Tr}_{L/K}(a + b\beta) = 2a, \mathrm{N}_{L/K}(a + n\beta) = a^2 - b^2 D$.

Recall that any endomorphism A of a vector space V over K admits a characteristic polynomial given by $\mathrm{ch}_A(X) = \det(X \operatorname{id}_V - A)$. If ch_A splits completely in some $M \supset K$ with roots $\lambda_1, \dots, \lambda_n$, then it's clear that $\operatorname{tr} A = \lambda_1 + \cdots + \lambda_n, \det A = \lambda_1 \cdots \lambda_n$.
Let's apply this to $m_\alpha : L \to L$. If $L = K(\alpha)$, then $1, \alpha, \dots, \alpha^{n-1}$ is a K-basis of L where $n = \deg p_\alpha$ with $p_\alpha(X) = X^n + a_{n-1}X^{n-1} + \cdots + a_0$ the minimal polynomial of α. The matrix of m_α in this basis is then

$$\begin{pmatrix} 0 & & & -a_0 \\ 1 & \ddots & & -a_1 \\ & \ddots & 0 & \vdots \\ & & 1 & -a_{n-1} \end{pmatrix}$$

Explicit computation shows that $\mathrm{ch}_{m_\alpha}(X) = p_\alpha(X)$, hence $\mathrm{Tr}_{L/K}(\alpha) = -a_{n-1}$ and $\mathrm{N}_{L/K}(\alpha) = (-1)^n a_0$.
In general, if $L \supset K(\alpha) \supset K$ and $\beta_1, \dots, \beta_r$ is a $K(\alpha)$-basis of L, then $\{\beta_i \alpha^j\}_{i,j}$ is a basis of L/K. With respect to this basis, m_α is a block diagonal matrix with blocks A which is the matrix of $m_\alpha : K(\alpha) \to K(\alpha)$ as above. So $\mathrm{ch}_{m_\alpha : L \to L}(X) = p_\alpha(X)^r$. In particular, $\mathrm{Tr}_{L/K}(\alpha) = r \,\mathrm{Tr}_{K(\alpha)/K}(\alpha)$ an $\mathrm{N}_{L/K}(\alpha) = \mathrm{N}_{K(\alpha)/K}(\alpha)^2$.

Example 10.3. If $\alpha \in K$, then $\mathrm{Tr}_{L/K}(\alpha) = [L : K]\alpha$ and $\mathrm{N}_{L/K}(\alpha) = \alpha^{[L:K]}$.

Note that $\mathrm{Tr}_{L/K}$ is (L-)linear and $\mathrm{N}_{L/K}$ is multiplicative. We can conclude this discussion as follows:

Theorem 10.2. *Suppose L/K is a finite extension and $\alpha \in L$. Let p_α be a minimal polynomial of α having roots $\lambda_1, \dots, \lambda_n$ in its splitting field and $r = [L:K]/n$, then $\mathrm{Tr}_{L/K}(\alpha) = r(\lambda_1 + \cdots + \lambda_n), \mathrm{N}_{L/K}(\alpha) = (\lambda_1 \cdots \lambda_n)^r$.*

Theorem 10.3. *If $M/L/K$ is a sequence of field extensions, then $\mathrm{Tr}_{M/K} = \mathrm{Tr}_{L/K} \circ \mathrm{Tr}_{M/L}, \mathrm{N}_{M/K} = \mathrm{N}_{L/K}\,\mathrm{N}_{M/L}$.*

Theorem 10.4. *Let L/K be a Galois extension and $G = \mathrm{Aut}(L/K)$, then*

$$\mathrm{ch}_{m_\alpha}(X) = \prod_{\sigma \in G} (X - \sigma\alpha)$$

In particular, $\mathrm{Tr}_{L/K}(\alpha) = \sum_{\sigma \in G} \sigma\alpha, \mathrm{N}_{L/K}(\alpha) = \prod_{\sigma \in G} \sigma\alpha$.

Example 10.4. This gives another way of doing Example 10.2.

Proof. The minimal polynomial of α.

$$p_\alpha(X) = \prod_{\beta \in G\alpha} (X - \beta)$$

But $\mathrm{ch}_{m_\alpha : L \to L}(X) = p_\alpha(X)^{[L:K(\alpha)]} = p_\alpha(X)^{|G|/|G\alpha|}$ which is what we wanted. □

10.3 Normal Basis Theorem

Theorem 10.5. *Suppose L/K is a Galois extension and $G = \mathrm{Aut}(L/K)$, then there is some $\alpha \in L$ such that $G\alpha$ is a K-basis of L. Equivalently, $KG \to L, \sum_g a_g g \mapsto \sum_g a_g g\alpha$ is an isomorphism of representations of G.*

Example 10.5. 1. If $L = K(\sqrt{\theta})$ and $G = C_2$, then any α not in $K \cup K\sqrt{\theta}$ would work.
2. If $L = K(\sqrt[N]{\theta})$ and $G = C_N = \langle \sigma \rangle$, then any α not in any of $K(\sqrt[N]{\theta})^i$ would work.

Proof. We will show the special case where K contains a primitive $\mathrm{ord}(g)^{th}$ root of unity for any $g \in G$. So $\mathrm{char}\, K \nmid \mathrm{ord}(g)$, so $\mathrm{char}\, K \nmid |G|$. It's a fact from representation theory that characters behave quite well in this case. In particular, if V_1, V_2 are two K-representations of G, then $V_1 \cong V_2$ iff $\chi_{V_1} = \chi_{V_2}$. $\chi_{KG}(g) = |G| 1_{g=1}$, so we only need to show that $\chi_L(g) = 0$ whenever $g \neq 1$. Indeed, suppose $g \in G$ and $M = L^{\langle g \rangle}$, then $\mathrm{tr}_K(g) = [M:K]\,\mathrm{tr}_M(g)$. But $M\langle g \rangle \cong L$ as M-representations by Theorem 8.3, so $\mathrm{tr}_M(g) = 0$ as desired. □

Remark. 1. The general case is of course also true and use the same idea, albeit requiring some amount of digestion on representation theory over a general field.
2. There is also a proof without any representation theory, which depends on some determinant trick.

10.4 Function Fields

Let $K = \mathbb{C}(X) = \mathrm{FF}(\mathbb{C}[X])$. K can be interpreted as the set of $\mathbb{C}$-valued "algebraic" functions defined on all but finitely many points in $\mathbb{P}^1 = \mathbb{C} \cup \{\infty\}$. Take $L = K(\sqrt{X^3 - X}) = \mathbb{C}(X)[Y]/(Y^2 - X^3 + X)$ which is a quadratic Galois

extension of K. Its Galois group is $\langle\sigma\rangle$ where $\sigma : L \to L$ sends y to $-y$ and fixes x. What is the meaning of L, and what is the meaning of L/K?
Observe that $L = \mathrm{FF}(R)$ where $R = \mathbb{C}[X,Y]/(Y^2 - X^3 + X)$. R corresponds naturally to the set $E = \{(x,y) \in \mathbb{C}^2 : y^2 = x^3 - x\}$, in the sense that it's the right notion of the "set of polynomial functions on E". So L can be interpreted as the "set of algebraic functions on E".
What is the meaning of the Galois group of L/K? The inclusion $K \to L$ comes from an inclusion $\mathbb{C}[X] \to R$, which then corresponds to a map $p : E \to \mathbb{C}$ via projection onto the first coordinate. Excluding the points $0, \pm 1$, p would be a 2-to-1 map, and what the Galois group does is essentially swapping the two preimages!
But what does E look like? We know that it's the solution set of $y^2 = x(x-1)(x+1)$ in $\mathbb{C}^2$, which is slightly hard to picture. Over $\mathbb{R}^2$, what you'll get is an oval-like loop passing though $-1, 0$ and a separate curve passing through 1 (and no, I will not be putting illustrations). If you squint hard enough, this is essentially two circles with one point taken away from one of them.
How about over $\mathbb{C}$ then? Let $\Gamma = \{(x,y) \in E : x \in [-1,0] \cup [1,\infty)\} = p^{-1}([-1,0] \cup [1,\infty))$. One can show that $E \setminus \Gamma$ is disconnected, has two connected components, each homeomorphic (in fact conformal, whatever that means) to $\mathbb{C} \setminus ([-1,0] \cup [1,\infty))$. So E is two copies of $\mathbb{C} \setminus ([-1,0] \cup [1,\infty))$ glued together along some one-dimensional object. Each of of these copies is (topologically) just a cylinder, so they essentially glue to a torus (with a point removed)!
$p : E \to \mathbb{P}^1$ is can be viewed as a "branched double cover ramified at $0, \pm 1, \infty$". What does p actually look like around these "ramification points"? Around 0, it simply looks like the function $\mathbb{C} \to \mathbb{C}, y \mapsto y^2$ if you think about it. This is not an isolated phenomenon, but the details of the general case will have to be covered in more advanced courses on the subject, like Algebraic Geometry and Riemann Surfaces. By and large, ramifications are simply when the map behaves as if it's a power function. Kummer theory then will give us some tools to study this kind of things from an algebraic point of view.