Math 123—Algebra II

Lectures by Joe Harris
Notes by Max Wang

Harvard University, Spring 2012

Introduction

Math 123 is the second in a two-course undergraduate series on abstract algebra offered at Harvard University. This instance of the course dealt with fields and Galois theory, representation theory of finite groups, and rings and modules.

These notes were live-TEXed, then edited for correctness and clarity. I am responsible for all errata in this document, mathematical or otherwise; any merits of the material here should be credited to the lecturer, not to me.

Feel free to email me at mxawng@gmail.com with any comments.

Acknowledgments

In addition to the course staff, acknowledgment goes to Zev Chonoles, whose online lecture notes (`http://math.uchicago.edu/~chonoles/expository-notes/`) inspired me to post my own. I have also borrowed his format for this introduction page.

The page layout for these notes is based on the layout I used back when I took notes by hand. The LATEX styles can be found here: `https://github.com/mxw/latex-custom`.

Copyright

Lecture 1 — 1/23/12

Example (Fields)**.**

1. $\mathbb{Q}$, the rational numbers
2. $\mathbb{F}_p = \mathbb{Z}/p\mathbb{Z} \hookrightarrow \mathbb{F}_q$, where $q = p^n$
3. $\mathbb{C}(t)$, $\mathbb{C}(t_1, \ldots, t_n)$, the complex function fields of one or more variables (where addition and multiplication are defined pointwise as usual)

While in group theory, we often studied a group by examining its various subgroups, our primary tool in field theory will be to take a ground field F and to study its extensions.

Definition 1.1. Let F be a field. A field extension of F is given by a field K with

$$F \hookrightarrow K$$

where F is included in K and is closed under K's field operations and inverse. We denote the field extension K/F.

Definition 1.2. Two field extensions K/F and K'/F of a ground field F are isomorphic

$$K/F \cong K'/F$$

if $\exists \varphi : K \to K'$ such that we have the diagram

$$\begin{array}{ccc} K & \xrightarrow{\varphi} & K' \\ \uparrow & & \uparrow \\ F & \xrightarrow{\text{id}} & F \end{array}$$

Definition 1.3. Let $F \hookrightarrow K$ be a field extension. We say that an element $\alpha \in K$ is algebraic over F if α satisfies a polynomial equation in K

$$a_n\alpha^n + a_{n-1}\alpha^{n-1} + \cdots + a_1 a + a_0 = 0$$

where $a_i \in F$. Dividing by a_n, we can assume the polynomial to be monic. If α does not satisfy any such polynomial, it is transcendental over F.

Note that whether or not an element α is algebraic depends not only on the extension field K but also on the ground field F. For example,

Example. Let $K = \mathbb{C}$. Then πi is transcendental over $\mathbb{Q}$ but is algebraic over $\mathbb{R}$.

Recall that

$$F[X] = \{a_n X^n + \cdots + a_1 X + a_0 : a_i \in F\}$$

is the polynomial ring with coefficients in F.

Claim 1.4. *Let $F \hookrightarrow K$ be a field extension, $\alpha \in K$. Consider the evaluation ring homomorphism given by*

$$\begin{aligned} \varphi : F[X] &\longrightarrow K \\ X &\longmapsto \alpha \end{aligned}$$

Then α is transcendental over F iff φ is injective.

Proof. φ is injective iff $\ker \varphi = \{0\}$; this means that α does not satisfy any polynomials with coefficients in F, and hence is transcendental. ■

Recall from ring theory that an ideal is any subgroup of a ring that is closed under multiplication. We denote an ideal generated by some element x by (x).

Definition 1.5. Let $F \hookrightarrow K$ be a field extension, $\alpha \in K$ algebraic over F. Since $F[X]$ is a PID[1], we have

$$\ker \varphi = (f)$$

for some $f \in F[X]$ (we assume that f is monic). We call f the irreducible polynomial satisfied by α/F. We also define

$$\deg_F \alpha = \deg f$$

The degree of an algebraic element is also dependent on the ground field; for example,

Example. Let $K = \mathbb{C}$ and $\alpha = \sqrt{i} = e^{\pi i/4}$. In $\mathbb{Q}$, α satisfies the polynomial $X^4 + 1$, and we have

$$\deg_{\mathbb{Q}} \alpha = 4$$

In $\mathbb{Q}(i)$ (the rationals with i adjoined), α satisfies the polynomial $X^2 - i$, and we have

$$\deg_{\mathbb{Q}(i)} \alpha = 2$$

Definition 1.6. Let $F \hookrightarrow K$ be a field extension, $\alpha \in K$. Recall our evaluation map $\varphi : F[X] \to K$. We denote

$$F[\alpha] = \operatorname{im} \varphi = \{\beta \in K : \beta = a_n\alpha^n + \cdots + a_0, a_i \in F\}$$

which is the smallest subring in K that contains both F and α. Similarly, let

$$F(\alpha)$$

denote the smallest *subfield* of K containing F and α. We call these ring adjunction and field adjunction respectively.

[1] Principal ideal domain; that is, a ring in which all ideals are generated by a single element. Recall that all fields and all rings $F[X]$ are PIDs.

Observation 1.7. If α is transcendental over F, then we have

$$F[\alpha] \cong F[X]$$

and

$$F(\alpha) \cong F(X)$$

where $F(X)$ denotes the field of rational functions with coefficients in F.

If α is algebraic over F, $\ker \varphi = (f)$ is nonzero. Then we have

$$F[\alpha] = F(\alpha) \cong F[X]/(f)$$

This makes $F(\alpha)$ a finite-dimensional vector space over F with a basis of

$$1, \alpha, \alpha^2, \ldots, \alpha^{n-1}$$

where $n = \deg_F \alpha$, since these powers of α are those zeroed by the irreducible polynomial.

Definition 1.8. We say that a field extension $F \hookrightarrow K$ is algebraic over F if every $\alpha \in K$ is algebraic over F. We say $F \hookrightarrow K$ is finite over F if K is a finite-dimensional F-vector space. Note that a field extension that is finite over its ground field is also algebraic over that ground field; however, the converse is not true in general.

Let $F \hookrightarrow K$ be a field extension, $\alpha, \beta \in K$. We will ask three questions about the subfields of K generated by α and β.

1. When is $F(\alpha) = F(\beta)$? This is true, for instance, if $\beta = \alpha + 1$. We can also consider less trivial examples; for instance, taking $F = \mathbb{Q}$, $K = \mathbb{C}$, α the root of the polynomial $X^3 - X + 1$, and $\beta = \alpha^2$, we also have $\mathbb{Q}(\alpha) = \mathbb{Q}(\beta)$ (although this is far from immediately clear).

2. When is $F(\alpha) \cong F(\beta)$ as extensions of F?

3. When is $F(\alpha) \cong F(\beta)$ as extensions of F via an isomorphism $\alpha \leftrightarrow \beta$? If α and β are both transcendental, the the isomorphism is obvious. If they are both algebraic, then the isomorphism holds iff the monic polynomials satisfied by α and β over F are equal, yielding

$$F(\alpha) \cong F[X]/(f) \cong F(\beta)$$

Lecture 2 — 1/25/12

Definition 2.1. Let F be a field. There exists a canonical ring homomorphism

$$\varphi : \mathbb{Z} \longrightarrow F$$

The characteristic of F is defined by

$$\operatorname{char}(F) = \begin{cases} 0 & \ker \varphi = \{0\} \\ p & \ker \varphi = (p) \end{cases}$$

Note that the map φ is unique for each field F and hence commutes with inclusion; so, all field extensions of a given ground field share the ground field's characteristic.

Definition 2.2. If $F \hookrightarrow K$ is a finite field extension, then the degree of the extension is the dimension of K as an F-vector space, denoted

$$\deg(K/F) = [K : F]$$

Note. If we have

$$F \hookrightarrow K \hookrightarrow K'$$

then

$$[K' : F] \geq [K : F]$$

and also

$$[K' : F] \geq [K' : K]$$

Claim 2.3. $[K : F] = 1 \iff K = F$.

Proof. Immediate. ■

Proposition 2.4. *Let $F \hookrightarrow K$ be a field extension with $[K : F] = 2$. If $\operatorname{char}(F) \neq 2$, then*

$$K = F(\delta), \quad \delta^2 \in F$$

Proof. Choose any $\alpha \in K - F$. Then since $[K : F] = 2$, α and 1 are linearly independent. Then $1, \alpha, \alpha^2$ are dependent, and hence $\exists a_0, a_1, a_2$ such that

$$a_2\alpha^2 + a_1\alpha + a_0$$

We know $a_2 \neq 0$, so dividing by a_2,

$$\alpha^2 + b_1\alpha + b_0 = 0$$

Since $\operatorname{char}(F) \neq 2$, we have an element 2^{-1}, so we can set

$$\delta = \alpha + \frac{b_1}{2}$$

Then we have

$$\delta^2 + (b_0 - \frac{b_1^2}{4}) = 0$$

where $b_0 - \frac{b_1^2}{4} \in F$. Hence, $K = F(\delta)$ with $\delta^2 \in F$, as desired. ■

Note that we could abolish the restriction on characteristic by demanding only that *some* quadratic polynomial of δ to be in F.

Proposition 2.5. *Let $F \hookrightarrow K \hookrightarrow L$ where L/F is finite. Then*

$$[L : F] = [L : K][K : F]$$

Proof. Choose a basis $\alpha_1, \dots, \alpha_n$ for L/K and another basis $\beta_1, \dots, \beta_m$ for K/F. We claim that the pairwise products $\alpha_i\beta_j$ form a basis for L/F.

First, we will show these $\alpha_i\beta_j$ span. Let $\gamma \in L$. We can write

$$\gamma = a_1\alpha_1 + \cdots + a_n\alpha_n$$

where $a_i \in K$. But then we can write

$$a_i = b_{i1}\beta_1 + \cdots + b_{im}\beta_m$$

where $b_{ij} \in F$. Then we have

$$\gamma = \sum b_{ij}(\alpha_i\beta_j)$$

Next, we show independence. Suppose we have

$$\sum b_{ij}(\alpha_i\beta_j) = 0$$

Defining a_i as above, we get

$$\sum a_i\alpha_i = 0$$

But by the independence of the α_i in L/K, we have

$$\sum_j b_{ij}\beta_j = 0$$

for every i, which yields $b_{ij} = 0, \forall i, j$ by the independence of the β_j in K/F. This completes the proof. ∎

Corollary 2.6. *Let $F \hookrightarrow K$ be a finite field extension. Then $\forall \alpha \in K$,*

$$\deg_F \alpha \,|\, [K : F]$$

because we can write

$$F \hookrightarrow F(\alpha) \hookrightarrow K$$

Theorem 2.7. *Let $F \hookrightarrow K$ be a field extension, and let $L \subseteq K$ be the subset of elements that are algebraic over F. Then L is a subfield of K.*

Proof. Take $\alpha, \beta \in L$. We have the sequence of extensions

$$F \hookrightarrow F(\alpha) \hookrightarrow F(\alpha, \beta)$$

It is obvious that if β/F is algebraic, then so is $\beta/F(\alpha)$. Then $F(\alpha, \beta)$ is a finite extension over F, and hence every $\gamma \in F(\alpha, \beta)$ is algebraic over F. ∎

Example. Let $F = \mathbb{Q}$, $K = \mathbb{C}$. Set $L = \overline{\mathbb{Q}}$, the algebraic closure of $\mathbb{Q}$. We can factor polynomials completely in this field.

We would like to know if the algebraic closure of a field always exists. That is, given a field F, does there exist an extension $F \hookrightarrow K$ such that any polynomial with coefficients in K factors completely in K (or, equivalently, has a root in K).

Lecture 3 — 1/27/12

Example. Let α be the real cube root of 2, and let $\omega = e^{2\pi i/3}$. Define

$$\alpha_1 = \alpha \qquad \alpha_2 = \alpha\omega \qquad \alpha_3 = \alpha\omega^2$$

These are the roots of the polynomial $x^3 - 2$. Note that this polynomial is irreducible over $\mathbb{Q}$, so we have

$$[\mathbb{Q}(\alpha_i) : \mathbb{Q}] = 3$$

From the point of view of algebra, these three roots are indistinguishable (i.e., the field extensions made by adjoining α_i are all isomorphic).

We want to distinguish between our three extensions $\mathbb{Q}(\alpha_i)$ by considering them as subfields of

$$\mathbb{Q}(\alpha_1, \alpha_2, \alpha_3) = \mathbb{Q}(\alpha, \omega)$$

One way to make such a distinction would be to note that $\mathbb{Q}(\alpha_1) \subset \mathbb{R}$ whereas $\mathbb{Q}(\alpha_2), \mathbb{Q}(\alpha_3) \not\subset \mathbb{R}$. However, we would like to obtain this knowledge without relying on the existence of $\mathbb{R}$ or $\mathbb{C}$.

ω satisfies the irreducible polynomial

$$\frac{x^3 - 1}{x - 1} = x^2 + x + 1 = 0$$

which yields

$$[\mathbb{Q}(\omega) : \mathbb{Q}] = 2$$

A field extension of degree 3 cannot contain an element of degree 2, so $\forall i$, $\omega \notin \mathbb{Q}(\alpha_i)$. Thus, we learn that

$$[\mathbb{Q}(\alpha, \omega) : \mathbb{Q}] = 6$$

and also that

$$[\mathbb{Q}(\alpha, \omega) : \mathbb{Q}(\alpha_i)] = 2$$

and

$$[\mathbb{Q}(\alpha, \omega) : \mathbb{Q}(\omega)] = 3$$

Finally, we can also conclude that the $\mathbb{Q}(\alpha_i)$ must be distinct fields because ω is their ratio, and ω is not in any $\mathbb{Q}(\alpha_i)$.

We end up with the following picture:

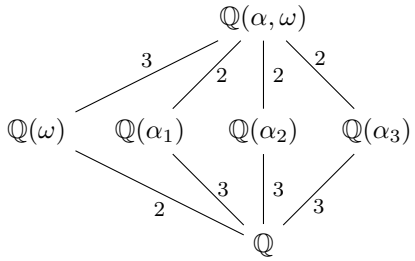

Note that, although we set out to obtain our results without relying on the complex numbers, we still required knowledge of $\mathbb{C}$ in order to *construct* our field extensions.

Example. Let us consider $\mathbb{Q}(\sqrt{2}, i)/\mathbb{Q}$. We know that $\mathbb{Q}(\sqrt{2})$ and $\mathbb{Q}(i)$ have degree 2 as extensions of $\mathbb{Q}$ (their generators satisfy quadratic polynomials). Thus, we know that $\mathbb{Q}(\sqrt{2}, i)/\mathbb{Q}$ has degree either 2 or 4. Again, this question reduces to whether the two single-adjunction fields are equal.

We know any number $\alpha \in \mathbb{Q}(i)$ can be written $\alpha = a + bi$ with $a, b \in \mathbb{Q}$. Thus, we would have $\mathbb{Q}(\sqrt{2}) = \mathbb{Q}(i)$ as field extensions over $\mathbb{Q}$ if we had

$$\alpha^2 = a^2 - b^2 + 2abi \stackrel{?}{=} 2$$

However, this equation cannot be satisfied; hence, the fields are distinct, and we have

$$[\mathbb{Q}(\sqrt{2}, i) : \mathbb{Q}] = 4$$

Note that there is a third extension, $\mathbb{Q}(i\sqrt{2})$. This is again a quadratic extension (has degree 2) and this is distinct for the same reason the other two are distinct. We assert without proof that these are all the intermediate extensions.

Example. We can replace i with $\sqrt{3}$ in the above example to yield identical results. Let us now attempt to find the irreducible polynomial over $\mathbb{Q}$ satisfied by

$$\alpha := \sqrt{2} + \sqrt{3}$$

(We know that α is algebraic because it is an element of a finite extension.) Note first that $1, \sqrt{2}$ are a basis for $\mathbb{Q}(\sqrt{2})/\mathbb{Q}$, $1, \sqrt{3}$ are a basis for $\mathbb{Q}(\sqrt{3})/\mathbb{Q}$, and hence

$$1, \sqrt{2}, \sqrt{3}, \sqrt{6}$$

form a basis for $\mathbb{Q}(\sqrt{2}, \sqrt{3})/\mathbb{Q}$.

We will solve this problem using two different approaches. First, let us write out some of the powers of α:

$$\begin{aligned}
\alpha^0 &= 1 \\
\alpha^1 &= \sqrt{2} + \sqrt{3} \\
\alpha^2 &= 5 + 2\sqrt{6} \\
\alpha^4 &= 49 + 20\sqrt{6}
\end{aligned}$$

Note that $1, \alpha^2, \alpha^4$ are linearly dependent, and in particular, we find that α satisfies

$$x^4 - 10x^2 + 1$$

We could check irreducibility by checking that α is not in any of the three intermediate extensions (generated by $\sqrt{2}$, $\sqrt{3}$, and $\sqrt{6}$). Alternatively, since $1, \alpha, \alpha^2$ are independent, we know that the desired irreducible polynomial is not quadratic (and we know it cannot be cubic because our extensions have even degrees).

We could, instead, denote $f(x)$ as the irreducible polynomial over $\mathbb{Q}$ satisfied by α, and consider the other roots of f beyond $\sqrt{2} + \sqrt{3}$. Note that algebra can't tell the difference between $\sqrt{2}$ and $-\sqrt{2}$; they are both just numbers satisfying $x^2 - 1$. A similar statement holds for $\sqrt{3}$. We might guess, then, that the other roots are $\sqrt{2} - \sqrt{3}$, $-\sqrt{2} + \sqrt{3}$, and $-\sqrt{2} - \sqrt{3}$. So we simply try

$$\begin{aligned}
&(x - \sqrt{2} - \sqrt{3})(x - \sqrt{2} + \sqrt{3}) \\
&\qquad (x + \sqrt{2} - \sqrt{3})(x + \sqrt{2} + \sqrt{3}) \\
&= ((x - \sqrt{2})^2 - 3)((x + \sqrt{2})^2 - 3) \\
&= (x^2 - 1)^2 - 8x^2 \\
&= x^4 - 10x^2 + 1
\end{aligned}$$

Note that this doesn't tell us that our polynomial is irreducible (but in fact it is).

Consider our picture for $\mathbb{Q}(\sqrt{2}, \sqrt{3})$:

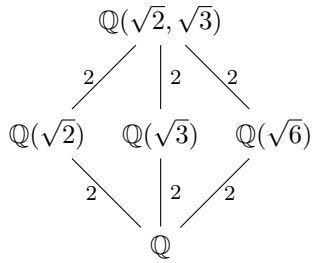

If we consider our field extensions as vector spaces, we find that $\mathbb{Q}(\sqrt{2}, \sqrt{3})$ is a four-space, and the intermediate extensions are all two-planes; therefore, in a very strong sense, most elements of $\mathbb{Q}(\sqrt{2}, \sqrt{3})$ must be single generators of that field.

Lecture 4 — 1/30/12

Definition 4.1. Let us formalize the notion of constructions with straightedge and compass, which allows us to bridge our algebra with planar geometry. In making such constructions, we begin with a pair of points $p, q \in \mathbb{R}^2$. We then define two basic constructions:

1. Draw the line $L_{p,q} = \overline{pq}$.
2. Draw the circle $C_p(q)$ with center at p passing through q.

Given two lines L_1, L_2, we can find their point of intersection; given a line L and a conic C or two conics C_1, C_2, we can find their two points of intersection. The intersections of these lines and conics are called constructible.

Construction 4.2. Begin with $p, q \in \mathbb{R}^2$. We construct the point $r \in L_{p,q}$ equidistant to both p and q as follows:

1. Draw the circle $C_p(q)$ through q about p.
2. Draw the circle $C_q(p)$ through p about q.
3. Draw the line L between the two intersection points of these circles.
4. r is the intersection $L \cap L_{p,q}$

Construction 4.3. Begin with a line L and a point $p \notin L$. We construct a perpendicular to L through p as follows:

1. Draw any circle about p intersecting L at two points q and q'.
2. Find the midpoint r between q and q' on L.
3. Draw the line $L_{p,r}$; this is our desired perpendicular.

Construction 4.4. Begin with a line L and a point $p \in L$. We construct a perpendicular to L through p as follows:

1. Draw any circle C_p around p, intersecting L at q and r.
2. Draw the circle $C_q(r)$ and the circle $C_r(q)$.
3. Connect the intersections of these two circles via a line L'; this line passes through the midpoint p and is perpendicular to L.

Construction 4.5. Begin with a line L and a point $p \notin L$. We construct a line parallel to L as follows:

1. Draw the line $L^\perp$ perpendicular to L through p.
2. Draw the line L' perpendicular to $L^\perp$ through p; this is parallel, as desired.

Construction 4.6. Begin with a pair of points p, q and a line L with a point $r \in L$. We construct a line segment in L with one endpoint r with length equal to $d(p, q)$ as follows:

1. Draw the circle $C_p(q)$.
2. Draw the line $L_{p,r}$.
3. Draw the line L' through p parallel to L; let s be the point $L' \cap C_p(q)$.
4. Draw the line L'' parallel to $L_{p,r}$ through s.
5. Let r' be the intersection $L'' \cap L$; the segment r, r' has $d(r, r') = d(p, q)$.

Definition 4.7. Begin with the two points $(0, 0), (1, 0) \in \mathbb{R}^2$. We call a point (a, b) constructible if we can construct it starting with our two points using our various constructions. Similarly, we say a point $a \in \mathbb{R}$ is constructible if $(a, 0)$ is constructible. Finally, an angle $\theta \in [0, 2\pi)$ is constructible if we can construct two lines L, L' with $\angle(L, L') = \theta$. Note that θ is constructible if $\sin\theta$ and $\cos\theta$ are constructible.

Let $p, q \in K^2 \subset \mathbb{R}^2$ where K is some subfield of $\mathbb{R}$. Then $L_{p,q}$ is defined by a linear equation with coefficients in K. Similarly, the circle $C_p(q)$ is defined by a quadratic polynomial in K.

If we have two lines L, L' defined by linear equations with coefficients in K, their point of intersection $L \cap L'$ is in K^2. If instead we have a line L and a circle C defined by equations with coefficients in K, their intersection points in $L \cap C$ have coordinates that live in a quadratic extension of K. To see this, we can parametrize L

$$L = \{(t, \alpha + \beta t) : t \in \mathbb{R}\} \quad \alpha, \beta \in K$$

and then apply the equation for C to $(t, \alpha + \beta t)$ to get a quadratic polynomial in t alone, with coefficients in K. This results in the following conclusion:

Proposition 4.8. *If a is constructible, then there exists a tower of fields*

$$\mathbb{Q} \subset K_1 \subset \cdots \subset K_n \ni a$$

such that

$$[K_i : K_{i-1}] = 2$$

and hence

$$[K_n : \mathbb{Q}] = 2^n$$

Moreover, we have, for some $r \in \mathbb{N}$,

$$\deg_{\mathbb{Q}} a = r$$

Example. We can use this proposition to conclude that it is not possible to trisect an arbitrary angle. Take $\frac{\pi}{3}$; we can ask whether $\theta = \frac{\pi}{9}$ is constructible. Define

$$\alpha = 2\cos\theta = e^{\pi i/9} + e^{-\pi i/9}$$

Then we have

$$\alpha^3 = e^{\pi i/3} + \underbrace{3e^{\pi i/9} + 3e^{-\pi i/9}}_{3\alpha} + e^{-\pi i/3}$$

and hence

$$\alpha^3 - 3\alpha - 1 = 0$$

Thus, α satisfies a cubic polynomial with no linear factors and which therefore is irreducible. So

$$\deg_{\mathbb{Q}} \alpha = 3$$

which, by our proposition, means that α and θ are not constructible.

Proposition 4.9. *Let $L \subset \mathbb{R}$ be the set of constructible numbers. Then L is a subfield of $\mathbb{R}$.*

Proof. Let $a, b \in L$. We know that

1. $a + b \in L$ since we can simply extend a segment of length a by one of length b.
2. $-a \in L$ because if we can construct $(a, 0)$, we can just as easily construct $(-a, 0)$.
3. $ab \in L$. To show this, we first construct a triangle of side length 1 along the x-axis and side length a vertically. We then construct the similar triangle with side length b along the x-axis; this will have side length ab vertically.
4. $\frac{1}{a} \in L$. We use similar triangles again, this time beginning with a triangle of side length a and 1 and scale it down so that the a side has length 1.

and so L is a field. ■

Proposition 4.10. *If a is constructible, then $\sqrt{a}$ is constructible.*

Proof. To construct it, we

1. Draw a circle with diameter $a + 1$ and divide the diameter L into two segments of lengths a and 1 at a point p.
2. Draw a line perpendicular to L through p. The height of this line in the circle is $\sqrt{a}$.

Thus, not only is L a field, but it is closed under square root. ■

Theorem 4.11. *a is constructible iff there exists a tower of fields*

$$\mathbb{Q} \subset K_1 \subset \cdots \subset K_n \ni a$$

such that

$$[K_i : K_{i-1}] = 2$$

Theorem 4.12. *Let F be a field, $f \in F[X]$ irreducible. Then $\exists K/F$ such that f has a root in K.*

Proof. Simply take $F[X]/(f)$. f is maximal since it is irreducible, so K is a field. Then $\bar{x}$, the equivalence class of x mod (f), satisfies polynomial $f(\bar{x}) = 0$. ■

Example. Let $F = \mathbb{F}_2 = \mathbb{Z}/(2)$. The polynomial

$$f(x) = x^2 + x + 1$$

is the unique irreducible quadratic polynomial in this field. Then we can form a field

$$\mathbb{F}_2[X]/(x^2 + x + 1) \cong \mathbb{F}_4$$

This is the unique field of four elements. (Note that it is not $\mathbb{Z}/4$, which is a ring but not a field.)

Lecture 5 — 2/1/12

Definition 5.1. Let F be a field The polynomial ring over F is given by

$$F[X] = \{a_n x^n + \cdots + a_1 x + a_0 : a_i \in F\} \subset F^F$$

Note that the polynomial ring does not map injectively to functions on F; this is pointedly the case where F is a finite field.

Definition 5.2. Let $f = a_n x^n + \cdots + a_0$. The derivative of f is given by

$$f' = n \cdot a_n x^{n-1} + \cdots + a_1$$

where n is the image of n under the canonical map $\mathbb{Z} \to F$.

Example. Let

$$f = x^2 + x = x(x - 1) \in \mathbb{F}_2[X]$$

Then $f' = 1$ because $(x^2)' = 0$. Note that any field of positive characteristic admits nonconstant polynomials whose derivatives are zero.

Theorem 5.3. *Let $F \hookrightarrow K$, $f, g \in F[X] \hookrightarrow K[X]$. Then any identity in $F[X]$ holds in $F[X]$ iff it holds in $K[X]$. This includes:*

1. *$\exists q, r \in F[X] : g = fq + r$ with $\deg r < \deg f$ when carried out in $F[X]$ iff $g = fq + r$ in $K[X]$. Note that q, r are unique.*
2. *$f \mid g \in F[X] \iff f \mid g \in K[X]$.*
3. *$\gcd_{F[X]}(f, g) = \gcd_{K[X]}(f, g)$.*
4. *Let $h \in F[X]$. Then $h \mid f$ and $h \mid g \in F[X] \iff h \mid f$ and $h \mid g \in K[X]$.*

Claim 5.4. *Let $F \hookrightarrow K$ be a field extension, $f, g \in F[X]$. If f is irreducible over F and f, g have a common factor in $K[X]$, then $f \mid g \in F[X]$.*

Proof. Since $f, g \in K[X]$ have a common factor h, then h is a common factor of $f, g \in F[X]$. But f is irreducible; thus, we must have $f = h$, which means $f \mid g$, as desired. ■

Lemma 5.5. *Let $\alpha \in F$, $f \in F[X]$, $f(\alpha) = 0$. Then α is a multiple root (i.e., $(x - \alpha)^2 \mid f$) iff $f'(\alpha) = 0$.*

Proof. Suppose first that $(x-\alpha)^2 \mid f$. Then $\exists g \in F[X]$: $f = (x-\alpha^2)g$. Taking derivatives, we have

$$f' = (x-\alpha)^2 g' + 2(x-\alpha)g$$

Then clearly, $f'(\alpha) = 0$.

Now suppose $f'(\alpha) = 0$. Then $\exists h \in F[X] : f = (x-\alpha)h$, and so $f' = (x-\alpha)h' + h$. Then we have

$$0 = f'(\alpha) = (\alpha-\alpha)h'(\alpha) + h(\alpha) = h(\alpha)$$

This means α is a root of h, which yields $(x-\alpha) \mid h$. But since $f = (x-\alpha)h$, we have

$$(x-\alpha) \mid f$$

as desired. ∎

Corollary 5.6. *If $f \in F[X]$ is irreducible, $f' \neq 0$, then f has no repeated roots in any extension $F \hookrightarrow K$.*

Proof. Suppose that f had a repeated root in some extension K/F. Then in $K[X]$, f, f' are not relatively prime (they have a common factor). But then they also have a common factor in $F[X]$. Then $f \mid f'$, but $\deg f' < \deg f$. This implies $f' = 0$, a contradiction. ∎

Proposition 5.7. *Let F be a finite field. Then*

$$|F| = \#F = p^r$$

for some $p \in \mathbb{P}$, $r \in \mathbb{N}$.

Proof. Let $\varphi : \mathbb{Z} \to F$ be the canonical homomorphism. We have $\ker \varphi = (p)$ for some prime p and $\operatorname{im} \varphi = \mathbb{Z}/(p) = \mathbb{F}_p$. So F is a finite-dimensional vector space over $\mathbb{F}_p$, which implies $\#F = p^r$. ∎

Theorem 5.8. *There exists a unique field of order p^r, which we denote*

$$\mathbb{F}_{p^r} = \mathbb{F}_q$$

Example. Let $F = \mathbb{F}_2$. There are four polynomials of degree 2 over $\mathbb{F}_2$. Three of them factor: x^2, $x(x-1)$, and $(x-1)^2$. The remaining one,

$$x^2 + x + 1$$

is the unique irreducible polynomial of degree 2 in $\mathbb{F}_2$. Then

$$\mathbb{F}_2[X]/(x^2+x+1) = \mathbb{F}_4$$

is the unique field of four elements because every field of four elements must be a quadratic extension of $\mathbb{F}_2$ and there is only one irreducible quadratic polynomial in $\mathbb{F}_2$. Let us write out the multiplication table of $\mathbb{F}_4$, given as

$$\mathbb{F}_4 = \{0, 1, \alpha, \alpha+1\}$$

We have

	0	1	α	$\alpha+1$
0	0	0	0	0
1	0	1	α	$\alpha+1$
α	0	α	$\alpha+1$	1
$\alpha+1$	0	$\alpha+1$	1	α

Definition 5.9. Let R be a ring. We define the group of units of R to be

$$R^\times = \{x : \exists x^{-1}, xx^{-1} = x^{-1}x = 1\}$$

Note that for a field F, $F^\times = F - \{0\}$.

Observation 5.10. Let F be a field of order $q = p^r$. Then $F^\times$ is a finite abelian group of order $q-1$. So $\forall \alpha \in F$, $\alpha \neq 0$, we have $\alpha^{q-1} = 1$, and hence every such α is a root of the polynomial

$$x^{q-1} - 1$$

Multiplying by x, we get that every element of F is a root of the polynomial

$$x^q - x$$

Since $|F| = q$ and $\deg(x^q - x) = q$, there are no repeated roots; thus,

$$x^q - x = \prod_{\alpha \in F} (x - \alpha)$$

Lemma 5.11. *Let K be any field, $G \subset K^\times$ a finite subgroup of $K^\times$. Then G is cyclic.*

Proof. First, consider any $a, b \in G$ of orders α and β respectively. If $\gcd(\alpha, \beta) = 1$, then $\operatorname{ord}(ab) = \alpha\beta$. Suppose that G is not cyclic. Then by the above (and since G is a finite abelian group),

$$n := \operatorname{lcm}\{\operatorname{ord}(\alpha) : \alpha \in G\} < |G|$$

(Actually, the lcm properly divides $|G|$, but the strict inequality is all we need.) But this means that the polynomial $x^n - 1$ has $|G| > n$ roots. However, a polynomial cannot have more roots than its degree, and hence, G must be cyclic. ∎

Example. Note that the last deduction in the above proof requires that K be a field. For instance, we have

$$(\mathbb{Z}/8)^\times = \{1, 3, 5, 7\} \cong \mathbb{Z}/2 \times \mathbb{Z}/2$$

But $x^2 - 1$ has four roots in $\mathbb{Z}/8$.

***Proof** (of Theorem).* First, we will show that $\exists K/\mathbb{F}_p$ a field extension of order p^r for any $r \in \mathbb{N}$. Let $L/\mathbb{F}_p$ be any extension in which the polynomial $f = x^q - x$ factors completely. Note that $f' = qx^{q-1} - 1$ which is -1 for $x = 0$ and $q - 1$ for $x \neq 0$. Since $f'(\alpha) \neq 0$ for every $\alpha \in L$, it has q distinct roots. We claim that

$$K := \{\alpha \in L : f(\alpha) = \alpha^q - \alpha = 0\} \hookrightarrow L$$

is a subfield (where by the above, we have $\#K = q$).

If $a, b \in K$, then $a^q = a$ and $b^q = b$, so $(ab)^q = ab$, and hence $ab \in K$. Moreover, we have $(a+b)^q = a^q + b^q = a + b$, so $a + b \in K$. So K is indeed a field, as desired.

Now we want to show that K is unique. Suppose we have two extensions K, K' with $\#K = \#K' = p^r = q$. We claim that $K \cong K'$.

We know that $K^\times$ is cyclic; let $\alpha \in K$ be a generator of $K^\times$. So

$$K = \{0, 1, \alpha, \alpha^2, \ldots, \alpha^{q-2}\}$$

In particular, we have $K = \mathbb{F}_p(\alpha)$. Now let f be the irreducible polynomial satisfied by $\alpha/\mathbb{F}_p$. By definition, we have

$$f \,|\, x^q - x \in K$$

But then $f \,|\, x^q - x \in \mathbb{F}_p$, and hence also $f \,|\, x^q - x \in K'$. Since $x^q - x$ factors completely in K', f factors completely in K', so f has a root $\alpha' \in K'$. Then $K' = \mathbb{F}_p(\alpha')$, and hence

$$K \cong \mathbb{F}_p[X]/(f) \cong K'$$

as desired. ■

Lecture 6 — 2/3/12

Before we proceed, we recall some general facts about finite fields.

- F finite $\implies \exists p \in \mathbb{P}, r \in \mathbb{N} : \#F = p^r = q$.
- $\#F = q \implies \forall x \in F, x^q - x = 0$.
- F finite $\implies F^\times$ cyclic.
- $\exists! F : \#F = p^r$.

Proposition 6.1. *$\mathbb{F}_{p^r}$ has a subfield isomorphic to $\mathbb{F}_{p^k}$ iff $k \,|\, r$.*

***Proof**.* First, let $\mathbb{F}_{p^k}$ be a subfield. Then $\mathbb{F}_{p^r}$ is a m-dimensional vector space over $\mathbb{F}_{p^k}$. Hence,

$$p^r = \#\mathbb{F}_{p^r} = (p^k)^m = p^{km}$$

and thus $k \,|\, r$.

Now suppose instead that $k \,|\, r$. Then $\exists m : p^r = (p^k)^m$. So we have

$$(x^{p^k} - x) \,|\, (x^{p^r} - x)$$

Then $x^{p^k} - x$ factors completely in $\mathbb{F}_{p^r}$ since $x^{p^r} - x$ does. Take

$$\{x \in \mathbb{F}_{p^r} : x^{p^k} - x = 0\} \subset \mathbb{F}_{p^r}$$

This is a subfield of order p^k. ■

Proposition 6.2. *The irreducible factors of $x^q - x/\mathbb{F}_p$ are exactly the irreducible polynomials over $\mathbb{F}_p$ whose degree k divides r.*

***Proof**.* Let $f \in \mathbb{F}_p[X]$ be irreducible of degree $k \,|\, r$. f factors completely in $\mathbb{F}_{p^r} \supset \mathbb{F}_{p^k} = \mathbb{F}_p[X]/(f)$, which means it has a common root with $x^q - x$. But since f is irreducible over $\mathbb{F}_p$, this means that

$$f \,|\, x^q - x \in \mathbb{F}_p$$

Now let $f \in \mathbb{F}_p[X]$ be irreducible and $f \,|\, x^q - x \in \mathbb{F}_p$. We know that f factors completely in $\mathbb{F}_{p^r}$. Choose α a root of f in F_{p^r}. We have

$$\mathbb{F}_p \hookrightarrow \mathbb{F}_p(\alpha) \hookrightarrow \mathbb{F}_{p^r}$$

and so $\deg f = k \,|\, r$, as desired. ■

Example. Let us consider the finite fields of characteristic 2. We have $\mathbb{F}_4 = \mathbb{F}_2(\alpha)$, where α is a root of $x^2 + x + 1$, the unique irreducible quadratic polynomial over $\mathbb{F}_2$. We can factor

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

We also have $\mathbb{F}_2 \hookrightarrow \mathbb{F}_8$. What are the cubic polynomials over $\mathbb{F}_2$? We have four which factor $1, 1, 1$

$$x^3 \quad x^2(x+1) \quad x(x+1)^2 \quad (x+1)^3$$

two which factor $1, 2$

$$x(x^2 + x + 1) \quad (x+1)(x^2 + x + 1)$$

and two which do not factor

$$x^3 + x + 1 \qquad x^3 + x^2 + 1$$

We can write

$$x^8 - x = x(x+1)(x^3 + x + 1)(x^3 + x^2 + 1)$$

Consider also $\mathbb{F}_2 \hookrightarrow \mathbb{F}_{16}$. Of the quartic polynomials in $\mathbb{F}_2$, 5 factor into linear factors, 1 factors into two quadratics, 4 factor into one linear and one cubic factor, 3 factor into a quadratic factor and two linear factors, and three are irreducible. We can factor

$$x^{16} - x = x(x+1)(x^2 + x + 1)(x^4 + x + 1)(x^4 + x^3 + 1)$$
$$(x^4 + x^3 + x^2 + x + 1)$$

Definition 6.3. A polynomial $f \in F[X]$ is said to split completely in an extension K/F if f factors into linear factors in $K[X]$.

Theorem 6.4 (Primitive Element Theorem)**.** *Let F be a field,* $\mathrm{char}(F) = 0$*, and $F \hookrightarrow K$ a finite extension. Then $\exists \alpha \in K : K = F(\alpha)$. α is called a primitive element for K/F (it generates the entire extension).*

Note that this theorem also holds for all finite fields F of arbitrary characteristic.

Proof. We know $K = F(\alpha_1, \ldots, \alpha_k)$ for some elements $\alpha_1, \ldots, \alpha_k \in K$. By induction on k, we can assume that $F \hookrightarrow F(\alpha_1, \ldots, \alpha_{k-1})$ has a primitive element β. But then our problem reduces to the base case, of showing that

$$F \hookrightarrow F(\beta, \alpha_k) = F(\alpha_1, \ldots, \alpha_k)$$

has a primitive element. So, we will try to show that if $K = F(\alpha, \beta)$, then $K = F(\gamma)$ for some $\gamma \in K$. In particular, we claim that for all but finitely many $c \in K$, $\beta + c\alpha$ is primitive.

Let f, g be the irreducible polynomials satisfied by α and β, respectively, over F. Let K'/K be an extension in which f and g split completely. Let $\alpha := \alpha_1, \ldots, \alpha_m$ be the roots of f; $\beta = \beta_1, \ldots, \beta_n$, the roots of g. Let

$$\gamma = \beta + c\alpha$$

Now let $L = F(\gamma)$. We claim that $\alpha \in L$; it will immediately follow that $\beta = \gamma - c\alpha \in L$, and hence $L = K$. Define a polynomial $h \in L[X]$ by

$$h(x) = g(\gamma - cx)$$

We know that the roots of g are β_i; thus, the roots of h are given by

$$x = \frac{\gamma - \beta_i}{c} = \frac{\beta - \beta_i}{c} + \alpha$$

We want to choose c so that $\gcd(f, h) = x - \alpha$; this will imply that $\alpha \in L$, as desired. But we can compute this gcd in $K'[X]$ instead But f splits completely over K'; this means it suffices to show that none of the α_j (except for α) is a root of h. This is the case whenever

$$c \neq \frac{\beta - \beta_i}{\alpha_i - \alpha}$$

which completes the proof. ∎

Lecture 7 — 2/6/12

Definition 7.1. Let $F \hookrightarrow K$ be a field extension. The elements $\alpha_1, \ldots, \alpha_n \in K$ are algebraically independent if they do not satisfy any $f \in F[X_1, \ldots, X_n]$ (i.e., if there is no f such that $f(\alpha_1, \ldots, \alpha_n) = 0$). Alternatively, we have a tower of extensions

$$F \hookrightarrow F(\alpha_1) \hookrightarrow F(\alpha_1, \alpha_2) \hookrightarrow \cdots \hookrightarrow F(\alpha_1, \ldots, \alpha_n)$$

where each extension is transcendental.

Definition 7.2. $\alpha_1, \ldots, \alpha_n \in K$ are a transcendence base for K/F if they are algebraically independent and

$$F(\alpha_1, \ldots, \alpha_n) \hookrightarrow K$$

is an algebraic extension. That is, $\alpha_1, \ldots, \alpha_n$ are a maximal collection of algebraically independent elements.

We will for the duration of this lecture assume that any transcendence basis is finite.

Theorem 7.3. *Let $\alpha_1, \ldots, \alpha_m \in K$ be a maximal transcendence basis and $\beta_1, \ldots, \beta_n \in K$ be algebraically independent. Then $n \leq m$.*

The proof of this theorem involves repeatedly replacing the α_i with the β_j in the transcendence basis; however, it is omitted.

Corollary 7.4. *Any two transcendence bases have the same cardinality.*

Definition 7.5. The transcendence degree of K/F is defined as the cardinality of any transcendence basis for K/F.

Definition 7.6. Let $F \hookrightarrow K$ be a field extension. We say that K/F is purely transcendental if $K = F(\alpha_1, \ldots, \alpha_n)$ for algebraically independent $\alpha_1, \ldots, \alpha_n$.

Note that all transcendental extensions can be decomposed as

$$F \overset{\text{p.t.}}{\hookrightarrow} F(\alpha_1, \ldots, \alpha_n) \overset{\text{alg.}}{\hookrightarrow} K$$

Theorem 7.7 (Luroth)**.** *Any transcendental subfield of $\mathbb{C}(t)$ is purely transcendental.*

Theorem 7.8 (Catelnuovo-Enriquez)**.** *If $K \subset \mathbb{C}(t, s)$ has transcendental degree 2, then K is purely transcendental.*

Theorem 7.9 (Clemens-Griffiths)**.** *The above does not hold for 3; that is, $\exists K \subset \mathbb{C}(t, s, u)$ with transcendental degree 3 that is unpure.*

Example. The field

$$K = \mathbb{C}(x)[y]/(y^2 - x^2 - 1)$$

is purely transcendental over $\mathbb{C}$; however,

$$L = \mathbb{C}(x)[y]/(y^2 - x^3 - 1)$$

is not pure.

One important context of this discussion is that of integrating functions. If $f(x) \in \mathbb{C}(x)$ (the field of rational functions over $\mathbb{C}$), then

$$\int f(x)\,\mathrm{d}x$$

can be calculated using partial fractions.

Meanwhile,

$$\int \frac{\mathrm{d}x}{\sqrt{x^2+1}} = \int_C \frac{\mathrm{d}x}{y}$$

where C is the curve $y^2 = x^2 + 1$. We can parametrize C: "every" line through C meets C exactly once, so we parametrize by the slope t through the point $(0, 1)$. Solving for the parametrization gives

$$\left(\frac{2t}{1-t^2}, \frac{1+t^2}{1-t^2}\right)$$

so we integrate

$$\int \frac{\mathrm{d}x}{y} = \int \frac{2}{1-t^2}\,\mathrm{d}t$$

This works because our parametrization exists, which is the same as saying

$$\mathbb{C}(x)[y]/(y^2 - x^2 - 1) \cong \mathbb{C}(t)$$

where $x \mapsto \frac{2t}{1-t^2}$ and $y \mapsto \frac{1+t^2}{1-t^2}$.

How, then, do we solve

$$\int \frac{\mathrm{d}x}{\sqrt{x^3+1}}$$

This integral spurred huge mathematical progress. A key fact in this problem is that L is not purely transcendental.

Consider surfaces in $\mathbb{C} \times \mathbb{C}$.

$$Z = \{(x, y) : y^2 = x^2 + 1\}$$

$$W = \{(x, y) : y^2 = x^3 + 1\}$$

Z is a sphere with punctures; integration along a path is path-independent. W is a torus with punctures; integration is path-dependent.

Lecture 8 — 2/8/12

Definition 8.1. Let R be any commutative ring with unit. We say that $f \in R[X_1, \ldots, X_n]$ is symmetric if it is invariant under the action of S_n on $R[X_1, \ldots, X_n]$; that is, if

$$f(x_1, \ldots, x_n) = f(x_{\sigma(1)}, \ldots, x_{\sigma(n)}), \quad \forall \sigma \in S_n$$

Note that we do not mean to say that the *values* of the polynomials are equal; for example, for $R = \mathbb{F}_p$,

$$x_1^p - x_1 \neq x_2^p - x_2$$

Rather, we want to say that the polynomials themselves are equivalent.

Example. Start with any monomial. Then the sum of the elements of its orbit under the action of S_n (which we refer to as the *orbit sum*) is symmetric. For instance, starting with x_1^k, we have

$$x_1^k + \cdots + x_n^k$$

This particular sum is called the *power sum*. If we start with, say, x_1x_2, we instead get

$$\sum_{i<j} x_i x_j$$

Starting with $x_1x_2^2$ yields

$$\sum_{i \neq j} x_i x_j^2$$

Definition 8.2. We define the ith elementary symmetric polynomial over $x_1, \cdots, x_n$, written

$$s_i = s_i(x_1, \ldots, x_n)$$

as the orbit sum of $x_1 \cdots x_i$.

The elementary symmetric polynomials (as polynomials in $u_1, \ldots, u_n$) are the coefficients, as a polynomial in x, of

$$\begin{aligned} p(x) &= \prod_{i=1}^{n} (x - u_i) \\ &= x^n - s_1 x^{n-1} + s_2 x^{n-2} - \cdots \pm s_n \end{aligned}$$

Theorem 8.3. *Let R be a ring. Any symmetric polynomial $g \in R[u_1, \ldots, u_n]$ is expressible as a polynomial of the elementary symmetric polynomials $s_1, \ldots, s_n$ in $u_1, \ldots, u_n$.*

Example. We can write the second power sum

$$x_1^2 + \cdots + x_n^2$$

in terms of elementary symmetric polynomials as

$$\begin{aligned} x_1^2 + \cdots + x_n^2 &= \left(\sum_{i=1}^{n} x_i\right)^2 - 2\sum_{i<j} x_i x_j \\ &= s_1^2 - 2s_2 \end{aligned}$$

Similarly, we can write

$$\begin{aligned} \sum_{i \neq j} x_i x_j^2 &= \left(\sum_{i=1}^{n} x_i\right)\left(\sum_{i<j} x_i x_j\right) - 3\sum_{i<j<k} x_i x_j x_k \\ &= s_1 s_2 - 3 s_3 \end{aligned}$$

Proof. Say $g(u_1, \dots, u_n)$ is symmetric. Let us define

$$g_0(u_1, \dots, u_{n-1}) = g(u_1, \dots, u_{n-1}, 0)$$

This is symmetric in the variables $u_1, \dots, u_{n-1}$. Inducting on n, we can write g_0 as

$$g_0(u_1, \dots, u_{n-1}) = Q(s_1^0, \dots, s_{n-1}^0)$$

a polynomial in the elementary symmetric polynomials over $u_1, \dots, u_{n-1}$.

Now we claim that we can write

$$g(u_1, \dots, u_n) = Q(s_1, \dots, s_{n-1}) + s_n \cdot h(u_1, \dots, u_n)$$

where h is symmetric of degree $= \deg g - n$. The difference $g(u_1, \dots, u_n) - Q(s_1, \dots, s_{n-1}) = 0$ whenever $u_n = 0$. So, $u_n \,|\, g(u_1, \dots, u_n) - Q(s_1, \dots, s_{n-1})$. But then by symmetry,

$$u_j \,|\, g(u_1, \dots, u_n) - Q(s_1, \dots, s_{n-1})$$

and hence,

$$s_n \,|\, g(u_1, \dots, u_n) - Q(s_1, \dots, s_{n-1})$$

Then h is the quotient of this division, and it must be symmetric since everything else in the equation is symmetric. Then by inducting on the degree of g with n fixed, we are done. ∎

Remark. Let $R = \mathbb{Z}$. Note that the power sums

$$\{x_1^k + \cdots + x_n^k\}$$

generate the ring of S_n-invariant polynomials in the variables $x_1, \dots, x_n$ over $\mathbb{Q}$, but not over $\mathbb{Z}$ as the elementary polynomials do.

Observation 8.4. Let $F \hookrightarrow K$ be a field extension. Let $f \in F[X]$, and say that f splits completely in K, with roots $\alpha_1, \dots, \alpha_n$. Since we can write

$$\begin{aligned} f(x) &= \prod_{i=1}^{n} (x - \alpha_i) \\ &= x^n - s_1 x^{n-1} + s_2 x^{n-2} - \cdots \pm sn \end{aligned}$$

we have

$$s_i = s_i(\alpha_1, \dots, \alpha_n) \in F$$

Then if $p(u_1, \dots, u_n)$ is any symmetric polynomial, we have $p(\alpha_1, \dots, \alpha_n) \in F$.

Definition 8.5. Define $p \in F[X]$ by

$$\begin{aligned} p(x) &= x^n - s_1 x^{n-1} + s_2 x^{n-2} - \cdots \pm s_n \\ &= \prod_{i=1}^{n} (x - u_i) \end{aligned}$$

with $s_i = s_i(u_1, \dots, u_n)$. We want to define a function

$$\begin{aligned} D(u_1, \dots, u_n) &= \prod_{i<j} (u_i - u_j)^2 \\ &= (-1)^{\binom{n}{2}} \prod_{i \neq j} (u_i - u_j) \end{aligned}$$

D is invariant under S_n, so we can write

$$D = \Delta(s_1, \dots, s_n)$$

The polynomial Δ is called the discriminant of p.

Note that Δ is well-defined whether or not p actually factors in its home field. Also, note that Δ evaluates to zero iff p has multiple roots in some extension.

Example. Let

$$p(x) = x^2 - s_1 x + s_2 = (x - \alpha)(x - \beta)$$

where $s_1 = \alpha + \beta$ and $s_2 = \alpha\beta$. We have

$$\begin{aligned} \Delta &= (\alpha - \beta)^2 \\ &= \alpha^2 - 2\alpha\beta \\ &= s_1^2 - 4s_2 \end{aligned}$$

which is the well-known discriminant for monic quadratic polynomials.

Now consider

$$p(x) = x^3 - s_1 x^2 + s_2 x - s_3$$

which has a much more complicated discriminant (which we will give without computation),

$$\Delta = -4s_1^3 s_3 + s_1^2 s_2^2 + 18 s_1 s_2 s_3 - 4s_2^3 - 27 s_3^2$$

(In this formula, we have actually taken $s_i = |s_i|$.) Note that if $s_1 = 0$, this reduces to

$$\Delta = -4s_2^3 - 27 s_3^2$$

We can arrive at this version by substituting $x \mapsto x - \frac{s_1}{3}$ for general cubics.

Definition 8.6. The Vandermonde matrix is given by

$$V = \begin{pmatrix} 1 & 1 & \cdots & 1 \\ x_1 & x_2 & \cdots & x_n \\ x_1^2 & x_2^2 & \cdots & x_n^2 \\ x_1^{n-1} & x_2^{n-1} & \cdots & x_n^{n-1} \end{pmatrix}$$

Observation 8.7. Note that the determinant vanishes whenever $x_i = x_j$ for some $i \neq j$. It is given by

$$\det V = \pm \prod_{i<j} (x_i - x_j)$$

which is a polynomial of degree $\binom{n}{2}$. This polynomial is invariant under the alternating group, but not the symmetric group. We could make it invariant my squaring, but let us examine an alternate route.

Consider the matrix

$$V' = \begin{pmatrix} 1 & 1 & \cdots & 1 \\ x_1 & x_2 & \cdots & x_n \\ x_1^2 & x_2^2 & \cdots & x_n^2 \\ x_1^n & x_2^n & \cdots & x_n^n \end{pmatrix}$$

Note that $\det V = 0$ and $\det V' = 0$ iff $x_i = x_j$ for some $i \neq j$, so $\det V \,|\, \det V'$. $\forall a_1 < \cdots < a_n$, the polynomial

$$\frac{\det(x_i^{a_j})}{\det(x_i^j)}$$

will be symmetric. The collection of such polynomials form an additive basis of the symmetric polynomials.

Lecture 9 — 2/10/12

Definition 9.1. Let $F \hookrightarrow K$ be a field extension, $f \in F[X]$. We say that K is a splitting field for f/F if

1. f splits completely in K as
$$f(x) = \prod_{i=1}^{n} (x - \alpha_i), \qquad \alpha_i \in K$$
2. K is the minimal extension in which f splits completely; that is,
$$K = F(\alpha_1, \ldots, \alpha_n)$$

Proposition 9.2. *Every $f \in F[X]$ has a splitting field.*

Proof. We know that f splits completely in *some* extension L/F. Then take the roots $\alpha_1, \ldots, \alpha_n$ of f; the extension

$$F \hookrightarrow F(\alpha_1, \ldots, \alpha_n)$$

is a splitting field.

We can also determine the splitting field algorithmically. Choose an irreducible factor f_0 of f. Let

$$K_1 = F[X]/(f_0)$$

Then $\exists \alpha \in K_1 : f(\alpha) = 0$. We replace f by $f/(x - \alpha)$ and repeat this process as necessary. ■

Proposition 9.3.

1. *Let K/F be a splitting field for $f \in F[X]$. Then $[K : F] < \infty$.*
2. *Let $F \hookrightarrow K \hookrightarrow L$ be a tower of field extensions. If L is a splitting field over F, then L is a splitting field over K.*
3. *Let $F \hookrightarrow L$ be a field extension, $f \in F[X]$. Then L contains at most one splitting field for f/F.*

Proof. All immediate. ■

Theorem 9.4. *Let $F \hookrightarrow K$ be a field extension. Suppose that K is a splitting field over F. Then $\forall g \in F[X]$ irreducible over F, if g has a root in K, then g splits completely in K.*

Proof. Say K/F is a splitting field for $f \in F[X]$. We can factor f as

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

Take $g \in F[X]$ irreducible such that $\exists \beta \in K : g(\beta) = 0$. Since

$$K = F(\alpha_1, \ldots, \alpha_n)$$

we can write

$$\beta = p(\alpha_1, \ldots, \alpha_n)$$

for some $p \in F[X_1, \ldots, X_n]$.

Let $\{p_1, \ldots, p_k\}$ be the orbit of p under the action of S_n on $F[X_1, \ldots, X_n]$, with $p_1 = p$. Set

$$\beta_i = p_i(\alpha_1, \ldots, \alpha_n)$$

and define a polynomial

$$h(x) = \prod_{i=1}^{k} (x - \beta_i) \in K[X]$$

We claim that $h \in F[X]$. We can write

$$h(x) = x^n - s_1 x^{n-1} + s_2 x^{n-2} - \cdots \pm s_n$$

where the

$$\begin{aligned} s_i &= s_i(\beta_1, \ldots, \beta_k) \\ &= s_i(p_1(\alpha_1, \ldots, \alpha_n), \ldots, p_k(\alpha_1, \ldots, \alpha_n)) \end{aligned}$$

Since the s_i are symmetric on the $\beta_1, \ldots, \beta_k$ and the p_i are symmetric on the $\alpha_1, \ldots, \alpha_n$, s_i is symmetric with respect to S_n on $\alpha_1, \ldots, \alpha_n$. Hence, we can express h as a polynomial in the elementary symmetric polynomials over $\alpha_1, \ldots, \alpha_n$, and thus in the coefficients of f. But these coefficients are in F, and we have $h \in F[X]$ as desired.

If g irreducible in F but has root in common with h in K, $g \,|\, h \in K[X]$. But then $g \,|\, h \in F[X]$. Since h splits completely in K, so does g. ■

Example. Any quadratic extension is a splitting field $K = F(\alpha)$ where α satisfies a quadratic polynomial $f = x^2 + ax + b$, and its inclusion in K enables f to split completely.

Example. Let $\alpha = \sqrt[3]{2}$ and $\omega = e^{2\pi i/3}$, set $\alpha_i = \alpha^i\omega$, and recall our observations about field extensions concerning these values.

$$K = \mathbb{Q}[X]/(x^3 - 2) \cong \mathbb{Q}(\alpha_1) \subset \mathbb{R} \subset \mathbb{C}$$

This is *not* a splitting field over $\mathbb{Q}$, since $x^3 - 2$ only has one root in K, as we have previously shown. If, however, we include other roots in the extension, then we can make K into a splitting field.

Note that we do need to check that K does not contain other roots; for instance, if $\omega \in K$, since ω satisfies $\frac{x^3-1}{x-1} = x^2+x+1$, then that polynomial would split completely in K. Fortunately, we have also shown previously that $\omega \notin K$.

Theorem 9.5. *Let F be a field of characteristic zero. Then any two splitting fields of $f \in F[X]$ are isomorphic.*

Proof. Say K_1, K_2 any two splitting fields for f/F. By the primitive element theorem, we have $K_1 = F(\gamma)$ for some $\gamma \in K_1$. Let $g \in F[X]$ be irreducible polynomial satisfied by γ/F Extend K_2 by a field L such that g has a root γ', and let $K' = F(\gamma') \subset L$. Then

$$K_1 \cong F[X]/(g) \cong K'$$

via an F-isomorphism sending $\gamma \mapsto \gamma'$. Since K_1 is a splitting field of f, so then is K'. But then $K', K_2 \subset L$ are both splitting fields for f/F, and hence $K' \cong K_2$ because L can contain at most one splitting field. ■

Lecture 10 — 2/13/12

We shall assume from this point onward that all fields in question have characteristic zero.

Definition 10.1. Recall that we call two field extensions K/F and K'/F isomorphic, or more specifically F-isomorphic

$$K/F \cong K'/F$$

if $\exists\varphi : K \to K'$ such that we have the diagram

$$\begin{array}{ccc} K & \xrightarrow{\varphi} & K' \\ \uparrow & & \uparrow \\ F & \xrightarrow{\mathrm{id}} & F \end{array}$$

An F-isomorphism from K/F to itself is called an F-automorphism. These are the symmetries of the field extension K.

Lemma 10.2. *Let $F \hookrightarrow K$ and $F \hookrightarrow K'$ be field extensions. Then we have the following*

1. *Let $f \in F[X]$, $\varphi : K \to K'$ an F-isomorphism. If $\alpha \in K$ is a root of f, then so is $\varphi(\alpha) \in K'$.*
2. *Suppose that $K = F(\alpha_1, \ldots, \alpha_n)$. Let $\varphi, \varphi' : K \to K'$ be F-isomorphisms. If $\forall i$, $\varphi(\alpha_i) = \varphi'(\alpha_i)$, then $\varphi = \varphi'$.*
3. *Let $f \in F[X]$ be irreducible, with a root $\alpha \in K$ and $\alpha \in K'$ in each extension. Then $\exists!\varphi : F(\alpha) \to F(\alpha')$ an F-isomorphism sending $\alpha \mapsto \alpha'$.*

Proof. Omitted. ■

Definition 10.3. Let $F \hookrightarrow K$ be a field extension. The Galois group of the extension K/F is defined as

$$\mathrm{Gal}(K/F) = \mathrm{Aut}(K/F)$$

the group of F-automorphisms of K.

Definition 10.4. A field extension $F \hookrightarrow K$ is called a Galois extension if

$$|\mathrm{Gal}(K/F)| = [K : F]$$

Definition 10.5. Let K be a field, H a group of automorphisms of K. The fixed field of H is the set of elements of K that are fixed by every element of H,

$$K^H = \{\alpha \in K : \sigma(\alpha) = \alpha, \forall\sigma \in H\}$$

Note that $K^H \subseteq K$ is a subfield.

Theorem 10.6. *Let K be a field, H a finite group of automorphisms on K. Let $F = K^H$. Let $\beta_1 \in K$, $\{\beta_1, \ldots, \beta_r\}$ its H-orbit. Then the irreducible polynomial of β_1/F is*

$$g(x) = (x - \beta_1)\cdots(x - \beta_r)$$

It follows that β_1 is algebraic over F with

$$\deg_F \beta_1 = r$$

and that $\deg_F \beta_1 \,|\, |H|$.

Proof. We want to show that g as defined is irreducible. We can write

$$\begin{aligned} g(x) &= (x - \beta_1)\cdots(x - \beta_r) \\ &= x^r - b_1x^{r-1} + \cdots \pm b_r \end{aligned}$$

Each b_i is a symmetric polynomial in the β_j. Since H permutes the β_j, it fixes each b_i, so $b_i \in K^H$.

Now let $h \in F[X]$ with β_1 as a root. For each $i = 1, \ldots, r$, we can find $\sigma_i \in H$ with $\sigma_i(\beta_1) = \beta_i$. Since the σ_i are F-automorphisms, then $\sigma_i(\beta_1) = \beta_i$ must also be roots of h. Then $g \,|\, h \in K[X]$, which means

$$g \,|\, h \in F[X]$$

It follows that g must be irreducible as desired. ■

Observation 10.7. Let $F \hookrightarrow K$ be algebraic but infinite. Then we can construct an infinite tower of fields

$$F < F_1 < F_2 < \cdots < K$$

To do so, start by taking $\alpha_1 \in K - F$ and $F_1 = F(\alpha_1)$. Since α_1 is algebraic over F, $[F_1 : F] < \infty$, so $F_1 < K$. We can then take $\alpha_2 \in K - F_1$ and $F_2 = F_1(\alpha_2)$. Similarly, $[F_2 : F_1] < \infty$, so $F_2 < K$. Continuing in this way, we get our desired tower.

Theorem 10.8 (Fixed Field Theorem)**.** *Let K be a field, H a finite group of automorphisms on K, $F = K^H$ its fixed field. Then*

$$[K : F] = |H|$$

Proof. Let $n = |H|$. We know that K/F is algebraic and $\deg_F \beta \,|\, n$ for every $\beta \in K$. This tells us that

$$[K : F] < \infty$$

By the primitive element theorem, $\exists \gamma \in K : K = F(\gamma)$. Since γ generates K, if $\sigma \in H$ fixes γ, we would have $\sigma = \mathrm{id}$. So the stabilizer of γ is $\{1\} \subset H$, and hence its orbit $H\gamma$ has order

$$|H\gamma| = n$$

Then $\deg_F \gamma = n$, which means

$$[K : F] = n$$

as desired. ■

Definition 10.9. Let $F \hookrightarrow K$ be a field extension. An intermediate field L is a field

$$F \subseteq L \subseteq K$$

We say an intermediate field is proper if $L \neq F$ and $L \neq K$.

Note that every L-automorphism of K is also an F-automorphism of K, and hence

$$\mathrm{Gal}(K/L) \subseteq \mathrm{Gal}(K/F)$$

Lemma 10.10. *Let $F \hookrightarrow K$ be a finite field extension, $G = \mathrm{Gal}(K/F)$. Then*

$$|G| \,|\, [K : F]$$

Proof. By the fixed field theorem, we know that $|G| = [K : K^G]$. Since F-automorphisms are the identity on F, $F \subseteq K^G$, and hence we have the tower of extensions

$$F \hookrightarrow K^G \hookrightarrow K$$

Then

$$|G| = [K : K^G] \,|\, [K : F]$$

as desired. ■

Proposition 10.11. *Let K be a field, H a finite group of automorphisms on K. Then $K^H \hookrightarrow K$ is Galois and $H = \mathrm{Gal}(K/K^H)$.*

Proof. By definition, $H \subseteq \mathrm{Gal}(K/K^H)$. By the above lemma, we know $|G(K/K^H)| \,|\, [K : K^H]$, and the fixed field theorem gives $|H| = [K : K^H]$. Then $G(K/K^H) \subseteq H$, and we have equality as desired. ■

Observation 10.12. Let $F \hookrightarrow K$ be a finite extension with primitive element $\gamma_1 \in K$. Let $f \in F[X]$ be the irreducible polynomial satisfied by γ_1 with roots $\gamma_1, \ldots, \gamma_r \in K$.

We know that there is a unique F-isomorphism $\sigma_i : F(\gamma_1) \to F(\gamma_i)$ with $\gamma_1 \mapsto \gamma_i$. Since $K = F(\gamma_1)$, it follows that $K = F(\gamma_i)$ for each i, so σ_i is an F-automorphism of K. Since every F-automorphism must take $\gamma_1 \mapsto \gamma_i$ (since γ_i are all the roots of f in K), we have

$$\mathrm{Gal}(K/F) = \{\sigma_i\}$$

with $|\mathrm{Gal}(K/F)| = r$.

Theorem 10.13. *Let $F \hookrightarrow K$ be a finite field extension with $G = \mathrm{Gal}(K/F)$. The following are equivalent*

1. *K/F is a Galois extension (i.e., $|G| = [K : F]$)*
2. *$F = K^G$*
3. *K is a splitting field over F*

We use the second condition to show that an element $\alpha \in K$ is actually also in F; we use the third condition to determine that an extension is Galois.

Proof. Showing that $1 \iff 2$ is quick. The fixed field theorem gives us $|G| = [K : K^G]$. Since $F \hookrightarrow K^G \hookrightarrow K$,

$$|G| = [K : F] \iff F = K^G$$

Now we will show that $1 \iff 3$. By the primitive element theorem, we can choose $\gamma_1 \in K$ such that $K = F(\gamma_1)$. Let $f \in F[X]$ be the irreducible polynomial satisfied by γ_1. We know that

$$[K : F] = \deg_F \gamma_1 = \deg f$$

Let $\gamma_1, \ldots, \gamma_r \in K$ be the roots of f in K. Then we have $|G| = r$ by our previous observation.

If $r = |G| = [K : F]$, then since $\deg f = [K : F]$, f splits completely in K and hence $K = F(\gamma_1) = F(\gamma_1, \ldots, \gamma_n)$ is a splitting field. Conversely, if K is a splitting field, the same reasoning applied in reverse yields $|G| = r = [K : F]$. ■

Corollary 10.14.

1. *Every finite extension K/F is contained in a Galois extension.*

2. *If K/F is Galois, L an intermediate field, then K/L is also Galois, and*

$$\mathrm{Gal}(K/L) \subseteq \mathrm{Gal}(K/F)$$

Observation 10.15. Let $F \hookrightarrow K$ be a Galois extension with $G = \mathrm{Gal}(K/F)$. Let $g \in F[X]$ split completely in K with roots $\beta_1, \dots, \beta_r$. Then

1. G acts on the roots $\{\beta_i\}$ by permuting them.

2. If K is a splitting field of g/F, we claim that the action of G is faithful.[2] We know that $K = F(\beta_1, \dots, \beta_r)$, and we get faithfulness from the fact that $\sigma \in G$ is determined entirely by its mapping of the generators $\beta_1, \dots, \beta_r$. It follows that $G \hookrightarrow S_r$.

3. If g is irreducible over F, the action of G is transitive.[3] Since g is irreducible, we know that g must be the irreducible polynomial satisfied by β_1. Since $F = K^G$, then $\{\beta_i\} = G\beta_1$, which is the statement of transitivity.

It follows that if K is a splitting field of g/F *and* g is irreducible in F, then $G \hookrightarrow S_r$ embeds transitively.

We will now state and prove the fundamental theorem of Galois theory, which provides for a bijective correspondence between intermediate fields and subgroups of the Galois group. Having build up significant machinery concerning Galois extensions, this proof will be trivial.

Theorem 10.16 (Fundamental Theorem of Galois Theory)**.** *Let $F \hookrightarrow K$ be a Galois extension, $G = \mathrm{Gal}(K/F)$. Then there is a bijective correspondence between*

$$\{H : H \leq G\} \longleftrightarrow \{L : F \hookrightarrow L \hookrightarrow K\}$$

In one direction, the bijection maps

$$H \longmapsto K^H$$

and in the inverse direction, it takes

$$L \longmapsto \mathrm{Gal}(K/L)$$

Proof. Let $H \leq G$ and $L = K^H$. By the fixed field theorem, $H = \mathrm{Gal}(K/L)$. Now suppose that L is an intermediate field, $H = \mathrm{Gal}(K/L)$. Then since K/F is Galois, so is K/L, and equivalently, $L = K^H$. ■

Observation 10.17. Note that if L and L' are intermediate fields and H and H' are their corresponding subgroups, $L \subset L'$ iff $H \supset H'$. In particular, F corresponds to $\mathrm{Gal}(K/F)$ and K corresponds to $\{1\}$.

If we have L corresponding to H, the since K/L is Galois and $H = \mathrm{Gal}(K/L)$, we have

$$[K : L] = |H|$$

We also know that $|G| = [K : F] = [K : L][L : F]$ and $|G| = |H|[G : H]$, so we also have

$$[L : F] = [G : H]$$

Corollary 10.18. *A finite field extension $F \hookrightarrow K$ has finitely many intermediate fields.*

Lecture 11 — 2/15/12

Example. Let $F = \mathbb{Q}$. Take $\alpha = \sqrt{3}$, $\beta = \sqrt{5}$ and let $K = F(\alpha, \beta)$. K is the splitting field of

$$(x^2 - 3)(x^2 - 5)$$

and hence is a Galois extension. We have

$$|\mathrm{Gal}(K/F)| = [K : F] = 4$$

and hence $\mathrm{Gal}(K/F)$ is either C_4 the cyclic group or V the Klein four group. We know that $F(\alpha)$, $F(\beta)$, and $F(\alpha\beta)$ are three distinct intermediate fields of K/F, and hence correspond to proper subgroups of $\mathrm{Gal}(K/F)$. Since $[K : L] = 2$ for each of these intermediate subgroups L, they correspond to subgroups of order 2. However, C_4 has only one element of order 2, and hence $\mathrm{Gal}(K/F) = V$, which has three elements (and hence three subgroups) of order 2.

These are the only proper subgroups of $\mathrm{Gal}(K/F)$, which means that $F(\alpha)$, $F(\beta)$, and $F(\alpha\beta)$ are the only proper intermediate fields.

Note that given

$$F \hookrightarrow L \hookrightarrow K$$

where K/F is Galois, we are only guaranteed that K/L is Galois, but not that L/F is. To determine whether L/F is Galois, we have the following:

Theorem 11.1. *Let $F \hookrightarrow K$ be a Galois extension with $G = \mathrm{Gal}(K/F)$. Let L be the fixed field K^H for a subgroup $H \leq G$. Then L/F is Galois iff $H \trianglelefteq G$.[4] If so, then*

$$\mathrm{Gal}(L/F) \cong G/H$$

[2] A group action on X is *faithful* if $\forall g \in G$, $\exists x \in X : gx \neq x$; in other words, if g fixes X, $g = e$.

[3] A group action on X is *transitive* if $Gx = X$ for any $x \in X$; in other words, if X has a single orbit under G.

[4] A subgroup $N \leq G$ is *normal* if $\forall n \in N$, $\forall g \in G$, $gng^{-1} \in N$. In other words, a normal subgroup is a subgroup that is invariant under conjugation.

Proof. Let $\epsilon_1 \in L$ be a primitive element for L/F, $g \in F[X]$ the irreducible polynomial for ϵ_1. Since K is a splitting field, $\epsilon_1 \in K$, g splits completely with roots $\epsilon_1, \dots, \epsilon_r$. Note that L/F is Galois iff L is a splitting field, which is the case iff

$$\forall i \in \{1, \dots, r\}, \epsilon_i \in L$$

We will show that this holds iff $H \trianglelefteq G$.

Since G is transitive on $\{\epsilon_1, \dots, \epsilon_r\}$, we know that $\forall i \in \{1, \dots, r\}$,

$$\exists \sigma_i \in G : \sigma_i(\epsilon_1) = \epsilon_i$$

Fix i and consider σ_i. We have $F(\epsilon_i) = L$ iff $\epsilon_i \in L$ (since $\deg_F \epsilon_1 = \deg_F \epsilon_i$). This is the case iff the stabilizer of ϵ_i is H. Meanwhile, the stabilizer of $\sigma(\epsilon_1)$ is the conjugate group $\sigma H \sigma^{-1}$. The condition that $H = \sigma H \sigma^{-1}$ is exactly the condition that $H \trianglelefteq G$, which is our desired conclusion.

Suppose that L/F is Galois. Then $\epsilon_i \in L$ for each i. An F-automorphism $\sigma \in G$ takes $\epsilon_1 \mapsto \epsilon_i$ for some i, and hence maps

$$\sigma : F(\epsilon_1) = L \longrightarrow L = F(\epsilon_i)$$

The restriction $\sigma|_L$ is hence an F-automorphism of L.

This restriction operation induces a group homomorphism

$$\varphi : G \longrightarrow \mathrm{Gal}(L/F)$$

We have

$$\ker \varphi = \{\sigma \in G : \sigma|_L = \mathrm{id}\} = H$$

We also have that

$$|G/H| = [G : H] = |\operatorname{Gal}(L/F)|$$

which means $\operatorname{im} \varphi = \mathrm{Gal}(L/F)$, and by the First Isomorphism Theorem,

$$G/H \cong \mathrm{Gal}(L/F)$$

as desired. ∎

Let us now apply the machinery of Galois theory to the study of *cubic polynomials* over a field F. Consider

$$\begin{aligned} f(x) &= x^3 - a_1 x^2 + a_2 x - a_3 \\ &= (x - \alpha_1)(x - \alpha_2)(x - \alpha_3) \end{aligned}$$

where $\alpha_i \in K$ are the roots of f in the splitting field K of f over F.

Note that $a_1 = \alpha_1 + \alpha_2 + \alpha_3 \in F$. Hence, α_1 and α_2 generate α_3, and we have the tower of extensions

$$F \hookrightarrow F(\alpha_1) \hookrightarrow F(\alpha_1, \alpha_2) = F(\alpha_1, \alpha_2, \alpha_3) = K$$

Let $L = F(\alpha_1)$. Since f/F is irreducible,

$$[L : F] = 3$$

and we can factor

$$f(x) = (x - \alpha_1) q(x)$$

where q is the quadratic polynomial with roots α_2, α_3. Either q is irreducible over L or it is not; if it is, then

$$[K : L] = 2 \qquad [K : F] = 6$$

Otherwise,

$$L = K \qquad [K : F] = 3$$

Example. Let $F = \mathbb{Q}$, $f(x) = x^3 + 3x + 1$, which is irreducible over $\mathbb{Q}$. The derivative $f'(x) = 3x^2 + 3$ is strictly greater than zero, so f is strictly increasing and hence has only one real zero α_1. α_1 cannot generate the complex roots of f, so $[K : F] = 6$ where K is the splitting field of f.

Example. Let $F = \mathbb{Q}$, $f(x) = x^3 - 3x + 1$, also irreducible over $\mathbb{Q}$. If α_1 is a root, then $\alpha_1^2 - 2$ is another root. We can generate the third root as given above, and hence in this case, $[K : F] = 3$.

By its action on the roots of the cubic f, the Galois group $G = \mathrm{Gal}(K/F)$ is a transitive subgroup of S_3. There are two such groups: S_3 itself, and the alternating (and cyclic) group A_3. If $[K : F] = 3$, then $G = A_3$; if $[K : F] = 6$, then $G = S_3$. The key distinction is whether or not q is irreducible over L.

To decide this, we will use the value

$$\delta = (\alpha_1 - \alpha_2)(\alpha_1 - \alpha_3)(\alpha_2 - \alpha_3) \in K$$

which is the square root of the discriminant D of f.[5] Note that $\delta \neq 0$ since the roots are distinct (we assume F has characteristic zero).

Theorem 11.2. *Let F be a field, $f \in F[X]$ an irreducible cubic polynomial, K the splitting field of f over F, and $G = \mathrm{Gal}(K/F)$. Let D be the discriminant of f; then*

1. *If $\delta = \sqrt{D} \in F$, then $[K : F] = 3$ and $G = A_3$.*
2. *If $\delta = \sqrt{D} \notin F$, then $[K : F] = 6$ and $G = S_3$.*

Proof. Permuting the roots of f multiplies δ by the sign of the permutation. If $\delta \in F$, it is fixed by G. Then every $\sigma \in G$ must be even, meaning $G = A_3$ and $[K : F] = 3$. Otherwise, $G = S_3$ and $[K : F] = 6$. ∎

[5] This is slight abuse of terminology, where we previously defined the discriminant as a polynomial in the elementary symmetric polynomials.

The alternating group A_3 has no proper subgroups (it is the cyclic group of order 3). Hence, if $G = A_3$, there are no intermediate fields, and our lattice is simply

$$\begin{array}{c} K = F(\alpha_1, \alpha_2, \alpha_3) \\ \Big| \, 3 \\ F \end{array}$$

This must be the case since $[K : F] = 3$ is prime. If instead we have $G = S_3$, we have four proper subgroups to consider, namely $\langle y \rangle$, $\langle xy \rangle$, $\langle x^2 y \rangle$, all of order 2, and $\langle x \rangle$ of order 3. Our lattice of fields is then

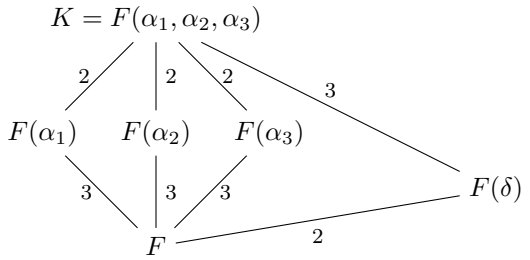

The corresponding lattice of Galois groups is

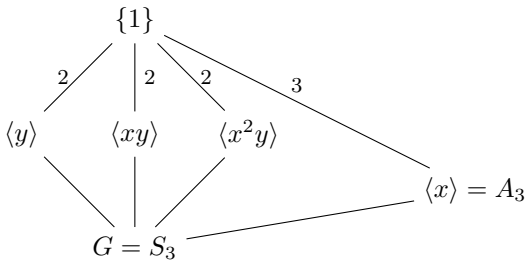

Lecture 12 — 2/17/16

We now turn to the study of *quartic polynomials*. Let F be any field, $f \in F[X]$ an irreducible quartic polynomial. Let K/F be the splitting extension of f over F, $G = \mathrm{Gal}(K/F)$. In K, we can write

$$f(x) = (x - \alpha_1)(x - \alpha_2)(x - \alpha_3)(x - \alpha_4)$$

Then G acts faithfully on the roots $\{\alpha_i\}$, and we have $G \hookrightarrow S_4$ transitively.

We can enumerate the transitive subgroups of S_4:

- S_4: symmetries of the tetrahedron. $|S_4| = 24$.
- A_4: rotations of the tetrahedron. $|A_4| = 12$. $A_4 \lhd S_4$ is normal.
- D_4: symmetries of the square. $|D_4| = 8$. S_4 has three conjugate subgroups isomorphic to D_4.
- C_4: rotations of the square. $|C_4| = 4$. S_4 has three conjugate subgroups isomorphic to C_4.
- D_2: reflections of the square. $|D_2| = 4$. $D_2 \lhd S_4$ is normal.

Example. $\mathbb{Q}(\sqrt{a}, \sqrt{b})$ is a quartic extension as long as $\sqrt{b} \notin \mathbb{Q}(\sqrt{a})$. Its Galois group must be D_2; it is given by the automorphisms

$$\begin{cases} \sqrt{a} \mapsto \sqrt{a} \\ \sqrt{b} \mapsto \sqrt{b} \end{cases}$$

$$\begin{cases} \sqrt{a} \mapsto -\sqrt{a} \\ \sqrt{b} \mapsto \sqrt{b} \end{cases}$$

$$\begin{cases} \sqrt{a} \mapsto \sqrt{a} \\ \sqrt{b} \mapsto -\sqrt{b} \end{cases}$$

$$\begin{cases} \sqrt{a} \mapsto -\sqrt{a} \\ \sqrt{b} \mapsto -\sqrt{b} \end{cases}$$

We can also view $K = \mathbb{Q}(\sqrt{a}, \sqrt{b})$ as the splitting field of a quartic polynomial. Choose any element of K not in one of the subfields, such as $\alpha = \sqrt{a} + \sqrt{b}$. Take its orbit

$$\{\sqrt{a} + \sqrt{b}, \sqrt{a} - \sqrt{b}, -\sqrt{a} + \sqrt{b}, -\sqrt{a} - \sqrt{b}\}$$

The irreducible polynomial of α over $\mathbb{Q}$ is the product of the monomials with these elements as roots.

Consider now the dihedral group D_4. It has three normal subgroups of order 4:

- D_4 acts on the set of two diagonals of a square. Let $H_D = \mathbb{Z}_2 \times \mathbb{Z}_2$ be the subgroup preserving the diagonals.
- D_4 acts on the set of two edge symmetries of a square. Let $H_A = \mathbb{Z}_2 \times \mathbb{Z}_2$ be the subgroup preserving opposite edges.
- Let $H_O = \mathbb{Z}_4$ be the subgroup preserving orientation.

Meanwhile, it has only one subgroup of order 2:

- The group $\{1, -1\}$, representing rotation by 180°.

Example. Consider the following tower of field extensions:

$$\begin{array}{c} \mathbb{Q}(\sqrt{4 + \sqrt{5}}) \\ \Big| \\ \mathbb{Q}(\sqrt{5}) \\ \Big| \\ \mathbb{Q} \end{array}$$

Note that $\mathbb{Q}(\sqrt{4+\sqrt{5}})/\mathbb{Q}$ is not Galois; we can permute $\mathbb{Q}(\sqrt{4+\sqrt{5}})$ to $\mathbb{Q}(\sqrt{4-\sqrt{5}})$. Let

$$\alpha = \sqrt{4+\sqrt{5}} \qquad \alpha' = \sqrt{4-\sqrt{5}}$$

Take $K = \mathbb{Q}(\alpha, \alpha')$. We have a quartic irreducible polynomial given by

$$\begin{aligned} f(x) &= (x-\alpha)(x+\alpha)(x-\alpha')(x+\alpha') \\ &= x^4 - 8x^2 + 11 \end{aligned}$$

Now consider the following square, with the vertices labeled by the roots $\pm\alpha, \pm\alpha'$:

$$\begin{array}{ccc} \alpha & \text{——} & \alpha' \\ | & & | \\ -\alpha' & \text{——} & -\alpha \end{array}$$

The subgroup H_D comprises the following automorphisms:

$$\begin{cases} \alpha \mapsto \alpha \\ \alpha' \mapsto \alpha' \end{cases} \begin{cases} \alpha \mapsto -\alpha \\ \alpha' \mapsto \alpha' \end{cases} \begin{cases} \alpha \mapsto \alpha \\ \alpha' \mapsto -\alpha' \end{cases} \begin{cases} \alpha \mapsto -\alpha \\ \alpha' \mapsto -\alpha' \end{cases}$$

α^2 is invariant under this group but not under all of D_4; hence, it must generate a nontrivial subfield

$$\mathbb{Q}(\sqrt{5}) = K^{H_D}$$

Note that this specific correspondence is dependent on the labeling of our square. The subgroup H_A comprises the following automorphisms:

$$\begin{cases} \alpha \mapsto \alpha \\ \alpha' \mapsto \alpha' \end{cases} \begin{cases} \alpha \mapsto \alpha' \\ \alpha' \mapsto \alpha \end{cases} \begin{cases} \alpha \mapsto -\alpha' \\ \alpha' \mapsto -\alpha \end{cases} \begin{cases} \alpha \mapsto -\alpha \\ \alpha' \mapsto -\alpha' \end{cases}$$

$\alpha\alpha'$ is invariant under this group but not under all of D_4; hence, it must generate a nontrivial subfield

$$\mathbb{Q}(\sqrt{11}) = K^{H_A}$$

Finally, consider the fixed field of the subgroup H_O preserving orientation, for which we have

$$\mathbb{Q}(\sqrt{55}) = K^{H_O}$$

We also have the fixed field of the only normal subgroup of order 2,

$$T = H_D \cap H_A$$

which has corresponding intermediate field

$$\mathbb{Q}(\sqrt{5}, \sqrt{11}) = K^T$$

We get the lattice

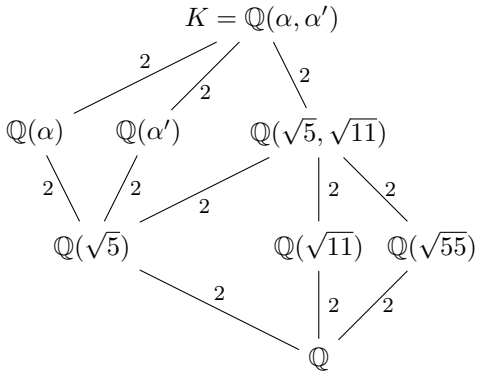

which can also be written

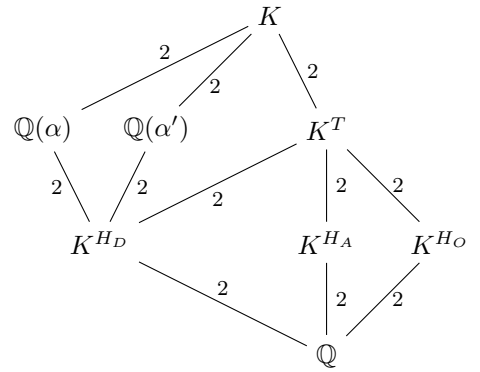

and with corresponding Galois group lattice

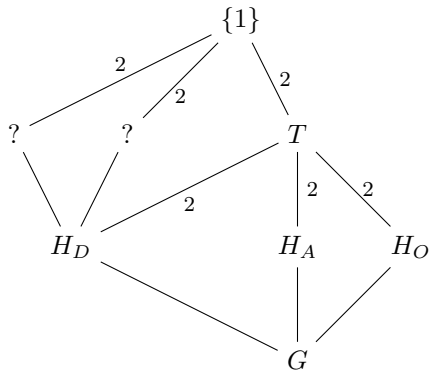

Note that $\mathbb{Q}(\alpha)$ and $\mathbb{Q}(\alpha')$ do not correspond to normal subgroups of D_4; they are not Galois over $\mathbb{Q}$.

Example. We can find a field extension with Galois group C_4 by a similar construction to the above. Let

$$\alpha = \sqrt{2+\sqrt{2}} \qquad \alpha' = \sqrt{2-\sqrt{2}}$$

Then $\alpha\alpha' = \sqrt{2} \in \mathbb{Q}(\alpha)$. Hence, $\mathbb{Q}(\alpha) = \mathbb{Q}(\alpha')$, and our Galois group will be C_4.

Recall our initial setup for general quartic polynomials. G must be one of the transitive subgroups of S_4; we can ask whether G is contained in each of them. Let

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

with square root

$$\delta = \prod_{i<j}(\alpha_i - \alpha_j)$$

As with cubic polynomials, permutation of roots on δ multiplies δ by the sign of the permutation, so δ is invariant under even permutations. Hence,

$$G \subseteq A_4 \iff \sqrt{D} \in F$$

Note that only $D_2, A_4 \subseteq A_4$ in S_4.

Since $A_4 \lhd S_4$, then we also have $A_4 \cap G \trianglelefteq G$.

$$F \hookrightarrow K^{A_4 \cap G} = F(\sqrt{D})$$

Indeed, if $\sqrt{D} \notin F$, we have the tower

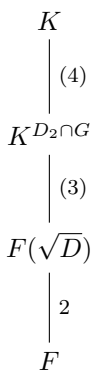

The fixed field $K^{D_2 \cap G}$ is splitting field of a cubic, called the resolvent cubic. To construct this field, we want to find an element of K that is invariant under D_2 but no larger subgroup of G. Take

$$\beta_1 = \alpha_1\alpha_2 + \alpha_3\alpha_4 \qquad \beta_2 = \alpha_1\alpha_3 + \alpha_2\alpha_4$$

$$\beta_3 = \alpha_1\alpha_4 + \alpha_2\alpha_3$$

These are in $K^{D_2 \cap G}$ Every permutation of the α_i permutes the β_j ($S_4/D_2 = S_3$), and hence

$$g(x) = (x-\beta_1)(x-\beta_2)(x-\beta_3) \in F[X]$$

We also have

$$K^{D_2 \cap G} = K(\beta_j)$$

Lecture 13 — 2/22/12

Let

$$\zeta_n = e^{2\pi i/n}$$

be an nth root of unity. We will assume that n is some prime p. We know that the irreducible polynomial for ζ_p over $\mathbb{Q}$ is

$$x^{p-1} + \cdots + x + 1$$

which has roots $\zeta, \zeta^2, \ldots, \zeta^{p-1}$. Hence, $\mathbb{Q}(\zeta_p)$ is its splitting field, and therefore also a Galois extension of $\mathbb{Q}$ with

$$[\mathbb{Q}(\zeta_p) : \mathbb{Q}] = p - 1$$

Definition 13.1. We call $F(\zeta_p)$ a cyclotomic extension.

Proposition 13.2. $\mathrm{Gal}(\mathbb{Q}(\zeta_p) : \mathbb{Q}) \cong C_{p-1}$.

Proof. Let $G = \mathrm{Gal}(\mathbb{Q}(\zeta_p) : \mathbb{Q})$. We know that every $\sigma \in G$ is determined by its mapping $\zeta_p \mapsto \zeta_p^i$. Then let

$$\sigma_i(\zeta_p) = \zeta_p^i$$

Since $\zeta_p^p = 1$, there is a natural bijection

$$G \longrightarrow \mathbb{F}_p^\times$$
$$\sigma_i \longmapsto i$$

which is an isomorphism. Since $\mathbb{F}_p^\times \cong C_{p-1}$, we have our desired result. ■

Observation 13.3. Note that $\langle \sigma_i \rangle = G$ iff $\langle i \rangle = \mathbb{F}_p^\times$. Also, our proof above works just as well for any arbitrary field $F \subseteq \mathbb{C}$, not necessarily $\mathbb{Q}$. However, in the case of general F, we instead get that $\mathrm{Gal}(F(\zeta_p) : \mathbb{Q}) \leq C_{p-1}$.

Example. Consider $\zeta = \zeta_{17}$. Since $\langle 3 \rangle = \mathbb{F}_{17}^\times$, we know that

$$G = \langle \sigma_3 \rangle$$

where $G = G(K/F)$. Let $\sigma = \sigma_3$. The subgroups of G are

$$\langle \sigma \rangle \supset \langle \sigma^2 \rangle \supset \langle \sigma^4 \rangle \supset \langle \sigma^8 \rangle \supset \langle \mathrm{id} \rangle$$

which corresponds to the tower of extensions

$$F = K^{\langle \sigma \rangle} \overset{2}{\hookrightarrow} K^{\langle \sigma^2 \rangle} \overset{2}{\hookrightarrow} K^{\langle \sigma^4 \rangle} \overset{2}{\hookrightarrow} K^{\langle \sigma^8 \rangle} \overset{2}{\hookrightarrow} K$$

Now let $\theta = 2\pi/17$. We claim that $F(\cos\theta) = K^{\langle \sigma^8 \rangle}$. Since $\zeta + \zeta^{-1} = 2\cos\theta$, we know that $\cos\theta \in K$. Moreover, ζ satisfies the polynomial

$$(x - \zeta)(x - \zeta^{-1}) = x^2 - 2(\cos\theta)x + 1 \in F(\cos\theta)$$

Thus,

$$[K : F(\cos\theta)] \le 2 \qquad \text{and} \qquad [F(\cos\theta) : F] \ge 8$$

So $F(\cos\theta)$ is either $K^{\langle \sigma^8 \rangle}$ or K. But $F(\cos\theta) \subseteq \mathbb{R}$, whereas $K \not\subseteq \mathbb{R}$; thus,

$$F(\cos\theta) = K^{\langle \sigma^8 \rangle}$$

as desired.

Lemma 13.4. *Let $\zeta = \zeta_p$. Let*

$$\alpha = c_1\zeta + c_2\zeta^2 + \cdots + c_{p-1}\zeta^{p-1}$$

be a linear combination with $c_i \in \mathbb{Q}$. If $\alpha \in \mathbb{Q}$, then $c_1 = c_2 = \cdots = c_{p-1}$ and $\alpha = -c_1$.

Proof. Since ζ satisfies

$$x^{p-1} + \cdots + x + 1$$

we can solve for ζ^{p-1} and rewrite

$$\alpha = (-c_{p-1})1 + (c_1 - c_{p-1})\zeta + \cdots + (c_{p-2} - c_{p-1})\zeta^{p-2}$$

Since $\{1, \zeta, \ldots, \zeta^{p-2}\}$ are a basis for K/F, we must have all coefficients except $-c_{p-1}$ equal to zero. This gives our desired result. ∎

Example. Again, take $\zeta = \zeta_{17}$. Taking iterative powers of ζ^3, we get the powers

$$1, 3, -8, -7, -4, 5, -2, -6, -1, -3, 8, 7, 4, -5, 2, 6$$

Recall that $\sigma = \sigma_3$ is a generator for the Galois group $G = \mathrm{Gal}(K/F)$. The orbit $G\zeta$ (under the action of automorphism) is the set $\{\zeta^i : \zeta^i \ne 1\}$. Let $H = \langle \sigma^2 \rangle$. H splits $G\zeta$ into two H-orbits,

$$\{\zeta, \zeta^{-8}, \zeta^{-4}, \ldots\} \qquad \{\zeta^3, \zeta^{-7}, \zeta^5, \ldots\}$$

Let α_1, α_2 denote the orbit sums. Then $\{\alpha_1, \alpha_2\}$ is a G-orbit, and our study of fixed fields tells us that the irreducible polynomial for α_1 and α_2 is

$$(x - \alpha_1)(x - \alpha_2)$$

To compute this polynomial, we want to determine

$$s_1(\alpha) = \alpha_1 + \alpha_2 \qquad s_2(\alpha) = \alpha_1\alpha_2$$

Since s_1 is the sum of all ζ^i with $\zeta^i \ne 1$, the irreducible polynomial of ζ gives

$$s_1(\alpha) = -1$$

We can compute $s_2(\alpha)$ by our previous lemma. Since $s_2(\alpha) \in \mathbb{Q}$ and since expanding $\alpha_1\alpha_2$ results in 64 summands of the form ζ^i, we know that each ζ_i appears four times, which yields

$$s_2(\alpha) = -4$$

Then our polynomial is

$$(x - \alpha_1)(x - \alpha_2) = x^2 + x - 4$$

Its discriminant is $D = 17$, and hence $K^{\langle \sigma^2 \rangle} = F(\sqrt{17})$.

In the same way, we can determine the quadratic extension over a field F that is contained in $F(\zeta_p)$ for any odd prime p.

Proposition 13.5. *Let $p \ne 2$ be prime. Let L be the unique quadratic extension over $\mathbb{Q}$ contained in $\mathbb{Q}(\zeta_p)$. If $p \equiv 1 \pmod 4$, then $L = \mathbb{Q}(\sqrt{p})$; if $p \equiv 3 \pmod 4$, then $L = \mathbb{Q}(\sqrt{-p})$.*

Note that we know the extension is unique by the Galois correspondence, since the cyclic group $F_p^\times$ for p prime has exactly one subgroup of index 2.

Proof. Analogous to the example. ∎

Theorem 13.6 (Kronecker-Weber Theorem)**.** *Every Galois extension of $\mathbb{Q}$ with an abelian Galois group is contained in a cyclotomic extension $\mathbb{Q}(\zeta_n)$.*

Observation 13.7. Let

$$n = p_1^{a_1} p_2^{a_2} \cdots p_k^{a_k}$$

be the prime decomposition of some $n \in \mathbb{N}$. Then the fields $\mathbb{Q}(\zeta_{p_i^{a_i}})$ intersect only in $\mathbb{Q}$ and

$$\prod \mathbb{Q}(\zeta_{p_i^{a_i}}) = \mathbb{Q}(\zeta_n)$$

Theorem 13.8. *Let F be a field with $\zeta_p \in F$ for some p prime, and let $b \in F$ with $b \ne 0$. Then the polynomial*

$$g(x) = x^p - b$$

is either irreducible or splits completely in F.

Proof. Let K/F be the splitting field of g and suppose there is a root β of g not in F. Then $[K : F] > 1$, which means $\exists \sigma \in \mathrm{Gal}(K/F) : \sigma \neq \mathrm{id}$. Since β generates K, $\sigma(\beta) = \zeta_p^r \beta$ for $0 < r < p$. Meanwhile, $\sigma(\zeta_p) = \zeta_p$. Then

$$\sigma^k(\beta) = \zeta_p^{kr}\beta$$

This attains every root $\zeta^i\beta$, and hence the action of G is transitive on the roots of g. Thus, g is irreducible over F, which completes our proof. ■

Theorem 13.9. *Let $F \subseteq \mathbb{C}$ be a subfield containing ζ_p for p prime, and let $F \hookrightarrow K$ be a Galois extension with degree $[K : F] = p$. Then $K = F(\sqrt[p]{\alpha})$ for some $\alpha \in F$.*

Proof. See Artin. ■

Lecture 14 — 2/24/12

Definition 14.1. Let $K_1, K_2 \subseteq K$ be subfields. The composite field of K_1 and K_2, denoted K_1K_2, is the smallest subfield of K containing both K_1 and K_2. We can also define the composite as the intersection of all subfields $K' \subseteq K$ containing both $K_1, K_2 \subseteq K'$ as subfields.

Proposition 14.2. *Let $F \hookrightarrow K$ be a field extension, and let K_1/F and K_2/F be finite field extensions of F contained in K. Then*

$$[K_1K_2 : F] \leq [K_1 : F][K_2 : F]$$

Proof. Let $\{\alpha_1, \ldots, \alpha_n\}$ be a basis for K_1/F, and $\{\beta_1, \ldots, \beta_m\}$ a basis for K_2/F. Then we have

$$K_1K_2 = K_1(\beta_1, \ldots, \beta_m) = F(\alpha_1, \ldots, \alpha_n, \beta_1, \ldots, \beta_m)$$

The β_j span K_1K_2 over K_1, so

$$[K_1K_2 : K_1] \leq m = [K_2 : F]$$

with equality iff the β_j are independent over K_1. ■

Corollary 14.3. *If $[K_1 : F] = n$, $[K_2 : F] = m$, with $\gcd(n, m) = 1$, then*

$$[K_1K_2 : F] = [K_1 : F][K_2 : F]$$

Proposition 14.4. *Let $F \hookrightarrow K$ be a Galois extension, $F \hookrightarrow F'$ any extension. Then KF'/F' is a Galois extension with Galois group*

$$\mathrm{Gal}(KF'/F') \cong \mathrm{Gal}(K/K \cap F')$$

We have the diagram

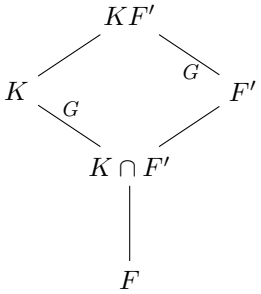

Proof. Since K/F is Galois, K is the splitting field of a polynomial $f \in F[X]$. Then KF'/F' is the splitting field of $f \in F'[X]$, and hence KF'/F' is Galois. Consider the restriction map

$$\begin{aligned} \varphi : \mathrm{Gal}(KF'/F') &\longrightarrow \mathrm{Gal}(K/F) \\ \sigma &\longmapsto \sigma|_K \end{aligned}$$

This is a homomorphism with kernel

$$\ker\varphi = \{\sigma \in \mathrm{Gal}(KF'/F') : \sigma|_K = \mathrm{id}\}$$

Note that $\sigma \in \ker\varphi$ is the identity on both F' and K by construction. Thus, $\ker\varphi = \{\mathrm{id}\}$, and hence φ is injective.

Let $H = \mathrm{im}\,\varphi$, K^H the corresponding fixed field in K containing F. Since H also fixes F', we know that

$$K^H \supseteq K \cap F'$$

Meanwhile, the group $\mathrm{Gal}(KF'/F')$ fixes K^HF'. By the Galois correspondence, this tells us that

$$K^HF' = F'$$

which means $K^H \subseteq F'$ and hence $K^H \subseteq K \cap F'$. Then $K^H = K \cap F'$, and the Galois correspondence yields

$$H = \mathrm{Gal}(K/K \cap F')$$

which completes our proof. ■

Corollary 14.5. *Let K/F be a Galois extension, F'/F any finite extension. Then*

$$[KF' : F] = \frac{[K : F][F' : F]}{[K \cap F' : F]}$$

Proposition 14.6. *Let F be a field, K_1/F and K_2/F Galois extensions. Then $K_1 \cap K_2/F$ is Galois.*

Proof. Let $f \in F[X]$ be an irreducible polynomial with root $\alpha \in K_1 \cap K_2$. Since $\alpha \in K_i$ and K_i is a splitting field over F, f splits in each. Then f splits in $K_1 \cap K_2$, which is therefore Galois as desired. ■

Proposition 14.7. *Let F be a field, K_1/F and K_2/F Galois extensions. Then K_1K_2/F is Galois, and*

$$\mathrm{Gal}(K_1K_2/F) = \{(\sigma, \tau) : \sigma|_{K_1 \cap K_2} = \tau|_{K_1 \cap K_2}\}$$

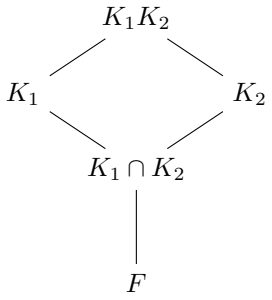

Proof. We know that K_1 is the splitting field of some polynomial $f_1 \in F[X]$, and similarly K_2 for $f_2 \in F[X]$. Then K_1K_2 is the splitting field of f_1f_2 (eliminating multiple roots). Hence, K_1K_2/F is Galois.

Consider the homomorphism

$$\begin{aligned}\varphi : \operatorname{Gal}(K_1K_2/F) &\longrightarrow \operatorname{Gal}(K_1/F) \times \operatorname{Gal}(K_2/F) \\ \sigma &\longmapsto (\sigma|_{K_1}, \sigma|_{K_2})\end{aligned}$$

Its kernel is trivial on K_1 and on K_2 and hence on the composite, so φ is injective. Since

$$(\sigma|_{K_1})|_{K_1\cap K_2} = \sigma|_{K_1\cap K_2} = (\sigma|_{K_2})|_{K_1\cap K_2}$$

we know that $\operatorname{im}\varphi \subseteq H$.

Note that for each $\sigma \in \operatorname{Gal}(K_1/F)$, there are precisely $|\operatorname{Gal}(K_2/K_1 \cap K_2)|$ elements $\tau \in \operatorname{Gal}(K_2/F)$ satisfying $\sigma|_{K_1\cap K_2} = \tau|_{K_1\cap K_2}$. Then

$$\begin{aligned}|H| &= |\operatorname{Gal}(K_1/F)||\operatorname{Gal}(K_2/K_1 \cap K_2)| \\ &= |\operatorname{Gal}(K_1/F)|\frac{|\operatorname{Gal}(K_2/F)|}{|\operatorname{Gal}(K_1 \cap K_2/F)|}\end{aligned}$$

Then we have $|H| = |\operatorname{Gal}(K_1K_2/F)| = [K_1K_2 : F]$, which yields

$$\operatorname{im}\varphi = H$$

which completes the proof. ■

Corollary 14.8. *Let K_1/F and K_2/F be Galois extensions. If $K_1 \cap K_2 = F$, then*

$$\operatorname{Gal}(K_1K_2/F) \cong \operatorname{Gal}(K_1/F) \times \operatorname{Gal}(K_2/F)$$

Conversely, if K/F is a Galois extension with $\operatorname{Gal}(K/F) = G_1 \times G_2$ a product of two subgroups, then there are two Galois extensions K_1/F and K_2/F such that $K = K_1K_2$ and $K_1 \cap K_2 = F$.

Lecture 15 — 2/27/12

Definition 15.1. Let F be a field, α algebraic over F. We say that α is solvable if $\alpha \in K$ for some K that can be obtained by a tower of field extensions

$$K_0 = F \hookrightarrow K_1 \hookrightarrow K_2 \hookrightarrow \cdots \hookrightarrow K_n = K$$

where K_i/K_{i-1} is Galois of order p_i prime.

Observation 15.2. Let $\mathbb{Q}_S$ be the field of solvable numbers over $\mathbb{Q}$. The field of constructible numbers $\mathbb{Q}_C$ is given by those elements $\alpha \in \bar{\mathbb{Q}}$ obtained by a tower of extensions with $[K_i : K_{i-1}] = 2$; it is clear, then, that

$$\mathbb{Q}_C \subseteq \mathbb{Q}_S$$

Now choose some $\beta \in \mathbb{Q}_S \subseteq \bar{\mathbb{Q}}$. Then we claim that $\alpha \in \bar{\mathbb{Q}}$ is solvable iff it is solvable over $F = \mathbb{Q}(\beta)$. If α is solvable over F, we can simply append the tower $F \overset{\cdots}{\hookrightarrow} F(\alpha)$ to the tower $\mathbb{Q} \overset{\cdots}{\hookrightarrow} F$, which yields solvability over $\mathbb{Q}$. If instead α is solvable over $\mathbb{Q}$, it is trivially solvable over F. Although some extensions in our tower might collapse, none will decompose because they all have prime degree. This yields the following diagram:

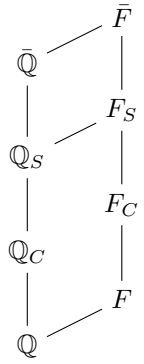

where the horizontal edges represent equality.

Theorem 15.3. *$\mathbb{Q}_S \subsetneq \bar{\mathbb{Q}}$. Specifically, suppose $\alpha \in \bar{Q}$ has irreducible polynomial $f \in \mathbb{Q}[X]$, and let K be the splitting field of $f/\mathbb{Q}$. Then if $\operatorname{Gal}(K/\mathbb{Q})$ is A_5 or S_5, then α is not solvable.*

Proof. Let $G = \operatorname{Gal}(K/\mathbb{Q})$. WLOG, we can assume $G = A_5$. For if $G = S_5$, then $\operatorname{Gal}(K/\mathbb{Q}(\sqrt{D})) = A_5$, and we can simply replace $\mathbb{Q}$ with $\mathbb{Q}(\sqrt{D})$.

If α is solvable, then we have a tower

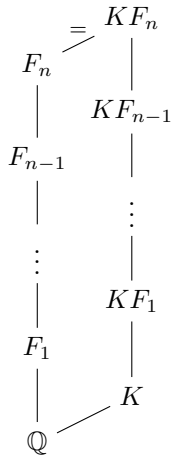

Our theorem then reduces to the following lemma:

Lemma 15.4. *Let K/F be Galois with $\mathrm{Gal}(K/F) = A_5$, L/F also Galois with $\mathrm{Gal}(L/F) = \mathbb{Z}_p$. Then we have the following diagram*

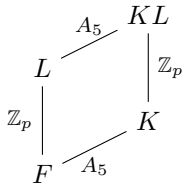

Proof of Lemma. We have only two possibilities for $\mathrm{Gal}(KL/K)$. First, suppose $\mathrm{Gal}(KL/K) = \{e\}$. Then $KL = K$, so $\mathrm{Gal}(KL/F) = A_5$. But this yields a surjection $A_5 \twoheadrightarrow \mathbb{Z}_p$, which is impossible.

It must be the case, then, that $\mathrm{Gal}(KL/K) = \mathbb{Z}_p$. Then $[KL : K] = p$, and it follows that $[KL : L] = 60$. Moreover we know we have an injection

$$\mathrm{Gal}(KL/L) \hookrightarrow \mathrm{Gal}(K/F)$$

and so we must have $\mathrm{Gal}(KL/L) = \mathrm{Gal}(K/F) = A_5$, as desired.

It remains still to be shown that there *exists* such a polynomial f as we assumed in our hypothesis. We will show that $\exists f \in \mathbb{Q}[X]$ irreducible and of degree 5 with Galois group S_5. To do so, we make use of the following lemma:

Lemma 15.5. *If $G \leq S_p$ is transitive (or equivalently, if G contains a p-cycle σ) and G contains a transposition $\tau = (i, j)$, then $G = S_p$.*

Proof of Lemma. We begin by relabeling the set of letters $\{1, \dots, p\}$ so that σ takes $1 \mapsto 2 \mapsto \cdots \mapsto p \mapsto 1$. Now replace σ by σ^{j-i}. Then σ carries $i \mapsto j$. We can relabel our letters again so that

$$\begin{aligned}\sigma &= 1 \mapsto 2 \mapsto \cdots \mapsto p \mapsto 1\\ \tau &= 1 \mapsto 2 \mapsto 1\end{aligned}$$

which makes $\sigma = (1, \dots, p)$ and $\tau = (1, 2)$. Then $\sigma\tau\sigma^{-1} = (2, 3)$, and repeated conjugation yields

$$\forall i \in \{1, \dots, p\}, \qquad G \ni (i, i+1)$$

These transpositions generate S_p, as desired.

Claim 15.6. *$\exists f \in \mathbb{Q}[X]$ irreducible and of degree 5 with Galois group S_5.*

Proof of Claim. Consider the splitting field $K/\mathbb{Q}$ for such a polynomial. We know that $\mathrm{Gal}(K/\mathbb{Q})$ would be transitive in S_5 (i.e., would contain a 5-cycle). We want to construct f such that it also contains a transposition. In other words, we want f to have exactly 3 real roots and 2 complex conjugate roots.

Take

$$\begin{aligned}f_0(x) &= x(x^2 - 4)(x^2 + 4)\\ &= x^5 - 16x\end{aligned}$$

This polynomial is negative in $(-\infty, -2)$ and changes sign at -2, 0, and 2. Its local maximum in $(-2, 0)$ is 15, and its local minimum in $(0, 2)$ is -15. Thus, we can add a constant term $-15 < c < 15$ and preserve the form of our roots. By the Eisenstein criterion for $p = 2$,

$$f(x) = f_0(x) + 2 = x^5 - 16x + 2$$

is irreducible.

This completes the proof of our theorem. ∎

Lecture 17 — 3/2/12

Definition 17.1. A representation of a finite group G is a finite-dimensional complex vector space V with an action of G on V; that is, a map

$$\rho' : G \times V \longrightarrow V$$

satisfying

$$\forall g, h \in G, \forall v \in V, \qquad \rho'(g, \rho(h, v)) = \rho'(gh, v)$$

Equivalently, we have a homomorphism

$$\rho : G \longrightarrow \mathrm{GL}(V) = \mathrm{Aut}(V)$$

Definition 17.2. A morphism of representations V, W of G is a linear map $\varphi : V \to W$ that commutes with the action of G; that is, the following diagram

$$\begin{array}{ccc} V & \xrightarrow{\varphi} & W \\ {\scriptstyle g}\downarrow & & \downarrow{\scriptstyle g} \\ V & \xrightarrow{\varphi} & W \end{array}$$

commutes $\forall g \in G$. Note that on the left, we have written g for $\rho_V(g)$, and similarly for $\rho_W(g)$ on the right, so more verbosely, we have

$$\varphi\rho_V(g) = \rho_W(g)\varphi$$

We can also think of the map as respecting conjugacy: $\varphi = g\varphi g^{-1}$. A morphism is, in particular, a G-module homomorphism.

Definition 17.3. A subrepresentation of a representation V of a group G is a vector subspace $W \subseteq V$ that is invariant under G; that is,

$$\forall g \in G, \quad g(W) = W$$

The direct sum of two representations V, W of G is also a representation, given by

$$g \cdot (v, w) = (gv, gw)$$

Recall that the dual space $V^* = \mathrm{Hom}(V, \mathbb{C})$. To define the dual representation, we want the pair v, v^* to be associated with $\rho(g)v$, $\rho^*(g)v^*$; that is, we want ρ^* to preserve the dual relationship. We know that the map $V \xrightarrow{g} W$ induces the dual map $W^* \xrightarrow{{}^t g} V^*$, which is defined by

$$\begin{array}{ccc} V & \xrightarrow{g} & W \\ & \searrow^{{}^t g(\lambda) = \lambda g} & \downarrow \lambda \\ & & \mathbb{C} \end{array}$$

However, if we have

$$V \xrightarrow{g} V \xrightarrow{h} V$$

then ${}^t(h \circ g) = {}^t g \circ {}^t h$. Then we cannot define the action ρ^* by the transpose; instead, we take

$$\rho^*(g) = {}^t\rho(g^{-1}) : V^* \longrightarrow V^*$$

which is a valid representation and respects duality.

Example. Let G be any finite group.

1. Let $V = \mathbb{C}$. The trivial representation is given by taking $g \equiv \mathrm{id}$; that is, $gv = v$.
2. Let $V = \mathbb{C}$; then $\mathrm{Aut}(V) = \mathbb{C}^*$. We get another one-dimensional representation via the character homomorphism
$$\chi : G \longrightarrow \mathbb{C}^*$$
3. Let V be a vector space with basis $\{e_g : g \in G\}$. The regular representation is given by
$$g : e_h \mapsto e_{gh}$$
4. Suppose G acts on a set S. The associated permutation representation is a vector space V given by a basis $\{e_s : s \in S\}$. The action of G on V is given by
$$g : e_s \mapsto e_{g(s)}$$
For instance, S_n acts on $\mathbb{C}^n$ by permuting the coordinates.

Given a finite group G, our goal in our study of representations is to classify, to describe, and to construct all representations of G.

Definition 17.4. We say that a representation V of G is irreducible if it has no nontrivial subrepresentations; that is,

$$\nexists W \subsetneq V, W \neq 0 \ : \ \forall g \in G, \ g(W) = W$$

Theorem 17.5. *Every representation of a group G is a direct sum of irreducible representations.*

Proof. The theorem follows immediately from the following lemma:

Lemma 17.6. *Let V be any representation, $W \subset V$ a proper subrepresentation. Then $\exists W' \subset V$ a subrepresentation such that $V = W \oplus W'$.*

We supply two proofs for this lemma.

First Proof of Lemma. We want an inner product h that is preserved by the action of G; that is, $\forall g \in G$,

$$h(v, u) = h(gv, gu)$$

Recall that an inner product is a positive definite Hermitian form; that is, it satisfies

1. Conjugate symmetry. $h(v, u) = \overline{h(u, v)}$.
2. Linearity in the first argument.
$$h(v + w, u) = h(v, u) + h(w, u)$$
$$h(\lambda v, u) = \lambda h(v, u)$$
3. Positive-definiteness. $h(v, v) \geq 0$ with equality iff $v = 0$.

Take h_0 to be any inner product on V. By averaging over G,

$$h(v, u) = \sum_{g \in G} h_0(gv, gu)$$

we get such an inner product. Then the subspace $W^\perp$ taken with respect to h is our desired complementary subrepresentation.

Second Proof of Lemma. Let $W' \subset V$ be any complementary linear subspace, and take

$$p_0 : V = W \oplus W' \longrightarrow W$$

to be any projection map. We want a projection p respecting conjugacy. Once again, we simply average over all of G, and define

$$p = \sum_{g \in G} g \circ p_0 \circ g^{-1}$$

Then $\ker p$ gives our desired complementary subrepresentation.

This proves the theorem. ∎

Example. This decomposition into irreducibles does not hold in general for infinite groups; consider, for instance,

$$\mathbb{Z} \cong \left\{ \begin{pmatrix} 1 & n \\ 0 & 1 \end{pmatrix} : n \in \mathbb{Z} \right\} \subset \mathrm{GL}_2$$

Lecture 18 — 3/5/12

Before continuing with our study of representation theory, we take note of the following two preliminaries:

Theorem 18.1. *Let V be a finite-dimensional complex vector space, $g : V \to V$ a linear map with*

$$g^n = \mathrm{id}$$

for some $n \in \mathbb{N}$. Then g is diagonalizable; that is, V has an eigenbasis of g-eigenvectors.

Note that if G is finite, $g^n = e$ for some n for every $g \in G$, so we can always diagonalize g's representation as a linear map.

Theorem 18.2. *Define*

$$\tau_k(x) = \sum x_i^k$$

be the power sums over the variables $x_1, \ldots, x_n$. The set $\tau_1, \ldots, \tau_n$ generates the ring of symmetric polynomials over $\mathbb{Q}$.

We will devote most of this lecture to studying the representations of the simplest nonabelian group, S_3. From the outset, we have two obvious irreducible representations:

1. U, the trivial representation, where $U \cong \mathbb{C}$ and S_3 acts as the identity.
2. U', the alternating representation, where $U' \cong \mathbb{C}$ and S_3 acts as the sign character.

We also have the permutation representation, where S_3 acts by permuting the coordinates of $\mathbb{C}^3$. However, this representation contains a copy of U, namely

$$U \cong \{(x, x, x) : x \in \mathbb{C}\}$$

Taking the complementary representation, we get a third irreducible representation

3. V, the so-called *standard representation*, given by

$$V = \{(x, y, z) : x + y + z = 0\}$$

This is irreducible because the action of S_3 on U is faithful, and hence no subspace is invariant.

Now let W be any representation of G. We will approach the problem of describing W by considering the eigenvalues and eigenvectors of the action of $S_3 \supset A_3 \circlearrowleft V$. Let $\sigma = (1\ 2\ 3) \in A_3$, and let $\omega = e^{2\pi i/3}$ be a primitive cube root of unity. σ has eigenvectors $(1, \omega, \omega^2)$, with eigenvalue ω, and $(1, \omega^2, \omega)$ with eigenvalue ω^2. It is easy to see that $\mathrm{span}(v, w) = V$.

Now let $\tau \in S_3$ be any transposition. We have

$$\tau\sigma\tau^{-1} = \sigma^2$$

Suppose that $v \in W$ is a σ-eigenvector with eigenvalue ω. We claim that $w = \tau(v)$ is also an eigenvector, with eigenvalue ω^2. For we have

$$\begin{aligned} \sigma(w) &= \sigma(\tau(v)) \\ &= \tau(\sigma^2(v)) \\ &= \tau(\omega^2(v)) \\ &= \omega^2\tau(v) \\ &= \omega^2 w \end{aligned}$$

So τ acts by exchanging the ω^i-eigenspaces of σ. If W is irreducible, then $W = \mathrm{span}(v, w) = V$.

Suppose instead that v has eigenvalue 1. Then $w = \tau(v)$ is also an eigenvector with eigenvalue 1, since τ only transposes the ω- and ω^2-eigenspaces. We have the following cases:

1. $w = v$. If W is irreducible, then $W = \mathrm{span}(v) = U$.
2. $w = -v$. If W is irreducible, then $W = \mathrm{span}(v) = U'$.
3. w, v linearly independent. Then

$$\begin{aligned} \mathrm{span}(v + w) &\cong U \\ \mathrm{span}(v - w) &\cong U' \end{aligned}$$

This demonstrates that U, U', and V are the only irreducible representations of S_3.

In general, this example illustrates that to understand a representation V of a finite group G, we want to know the eigenvalues of each element $g \in G$. We can use symmetric polynomials (since the eigenvalues are expressed in the characteristic polynomial) to convey this information.

Since the power sums generate the ring of symmetric polynomials, it will be enough to know for each g the sums $\sum \lambda_i^k$, for this knowledge can be used to retrieve the λ_i themselves. But if g has eigenvalues λ_i, then g^k has λ_i^k as its eigenvalues. This then motivates the following definition:

Definition 18.3. If V is a representation of G, we define the character map

$$\begin{aligned}\chi : G &\longrightarrow \mathbb{C} \\ g &\longmapsto \operatorname{tr}(g)\end{aligned}$$

that is, we associate g with the sum of its eigenvalues. Observe that $\chi(g)$ depends only on the conjugacy class of $g \in G$ (since conjugation changes only the eigenvector, not the eigenvalue). Hence, we can think of the character as a map

$$\chi : \mathscr{C} \longrightarrow \mathbb{C}$$

where $\mathscr{C}$ is the set of conjugacy classes of G.

The character table for S_3 is given by

S_3	e (1)	$(1\ 2)$ (3)	$(1\ 2\ 3)$ (2)
U	1	1	1
U'	1	-1	1
V	2	0	-1

Note that $\chi_V(\tau)$ is 0 because it is a transposition $\omega \leftrightarrow \omega^2$ and hence has zeros along the diagonal. On the other hand, we know that σ has eigenvalues ω and ω^2, and hence $\chi_V(\sigma) = \omega + \omega^2 = -1$.

Lecture 19 — 3/7/12

Let V be a representation of G. Recall that the *character* of V is given by

$$\begin{aligned}\chi_V : G &\longrightarrow G \\ g &\longmapsto \operatorname{tr}(g : V \to V)\end{aligned}$$

Definition 19.1. The invariant subspace of G of a representation V is given by

$$V^G = \{v \in V : gv = v, \forall g \in G\}$$

It is easy to see that $V^G \subseteq V$ is a subrepresentation.

Definition 19.2. The space of endomorphisms of a vector space V is

$$\operatorname{End}(V) = \operatorname{Hom}(V, V)$$

We want to know how we can determine V^G, or at least determine its dimension. To do so, we once again rely on averaging over G to obtain G-invariance.

Theorem 19.3. *The map $\varphi : V \longrightarrow V$ given by*

$$\varphi(v) = \frac{1}{|G|} \sum_{g \in G} g(v)$$

is a projection map $V \twoheadrightarrow V^G$.

Proof. If $w = \varphi(v)$, then for any $h \in G$,

$$\begin{aligned}hw &= \frac{1}{|G|} \sum_{g \in G} hg(v) \\ &= \frac{1}{|G|} \sum_{g \in G} g(v) \\ &= \varphi(v) \\ &= w\end{aligned}$$

and if $v \in V^G$, then

$$\varphi(v) = \frac{1}{|G|} \sum_{g \in G} v = v$$

and hence $\varphi^2 = \varphi$, as desired. ∎

We thus have

$$\varphi = \begin{pmatrix} \overbrace{\begin{matrix}1 & & \\ & \ddots & \\ & & 0\end{matrix}}^{\dim V^G} & & \\ & \begin{matrix}0 & \\ & \ddots\end{matrix} & \\ & & \end{pmatrix}$$

and hence

$$\begin{aligned}\dim V^G &= \operatorname{tr}(\varphi : V \to V) \\ &= \frac{1}{|G|} \sum_{g \in G} \operatorname{tr}(g : V \to V) \\ &= \frac{1}{|G|} \sum_{g \in G} \chi_V(g)\end{aligned}$$

This dimension is the number of copies of the trivial representation in the decomposition of V. In particular, if V is irreducible (and not the trivial representation), the sum

$$\sum_{g \in G} \chi_V(g) = 0$$

Definition 19.4. The space of representation morphisms between representations V and W of G is denoted

$$\operatorname{Hom}_G(V, W) = \{\alpha : V \to W \mid \alpha = \rho_W(g)^{-1} \circ \alpha \circ \rho_V(g)\}$$

This is a subspace of $\operatorname{Hom}(V, W)$. Note also that

$$\operatorname{Hom}(V, W \oplus W') = \operatorname{Hom}(V, W) \oplus \operatorname{Hom}(V, W')$$

The space $U = \operatorname{Hom}(V, W)$ forms a natural representation of G, given by

$$\begin{aligned}\rho_V(g) : \operatorname{Hom}(V, W) &\longrightarrow \operatorname{Hom}(V, W) \\ \alpha &\longmapsto g^{-1}\alpha g\end{aligned}$$

The key observation to make here is that

$$\operatorname{Hom}(V, W)^G = \operatorname{Hom}_G(V, W)$$

Lemma 19.5 (Schur's Lemma)**.** *Let V, W be irreducible representations of G. Then*

$$\dim(\mathrm{Hom}_G(V,W)) = \begin{cases} 1 & V \cong W \\ 0 & \textit{otherwise} \end{cases}$$

Proof. Let $\varphi \in \mathrm{Hom}_G(V,W)$. Both $\ker\varphi \subseteq V$ and $\mathrm{im}\,\varphi \subseteq W$ are subrepresentations. Thus, either $\varphi = 0$ or φ is an isomorphism (due to the irreducibility of V and W). It remains to be shown that the space of isomorphisms is 1-dimensional in the second case. Suppose $V = W$. Then the map

$$\varphi : V \longrightarrow V \in \mathrm{Hom}_G(V,V)$$

has an eigenvector with eigenvalue λ. Then $\varphi - \lambda I$ has a kernel, so $\varphi - \lambda I = 0$. ∎

Definition 19.6. Let V and W be vector spaces. The tensor product is a vector space $V \otimes W$ with a bilinear map

$$\begin{aligned} \varphi : V \oplus W &\longrightarrow V \otimes W \\ (v,w) &\longmapsto v \otimes w \end{aligned}$$

We provide three equivalent definitions for the tensor product (although we will not prove their equivalence)

1. Let $\{v_1,\dots,v_n\}$ be a basis for V, $\{w_1,\dots,w_m\}$ a basis for W. Then we define $V \otimes W$ as the vector space with basis $\{v_i \otimes w_j\}$ and extend the map φ by bilinearity. This definition readily yields

   $$\dim(V \otimes W) = \dim V \cdot \dim W$$

2. Let U be the vector space with basis

   $$\{v \otimes w : v \in V, w \in W\}$$

   We let U' be the subspace spanned by

   $$\begin{aligned} &(v+v') \otimes w - (v \otimes w + v' \otimes w) \\ &(\lambda v) \otimes w - \lambda(v \otimes w) \\ &v \otimes (w+w') - (v \otimes w + v \otimes w') \\ &v \otimes (\lambda w) - \lambda(v \otimes w) \end{aligned}$$

   We define the tensor product as the quotient

   $$V \otimes W = U/U'$$

   and we take $(v,w) \mapsto \overline{v \otimes w}$.

3. We take $V \otimes W$ to be the universal object for bilinear maps $V \oplus W \to U$. That is, every bilinear map $\alpha : V \oplus W \to U$ factors uniquely through $V \otimes W$.

   $$V \oplus W \xrightarrow{\varphi} V \otimes W \xrightarrow{\beta} U$$

   where β is linear. We have a natural bijection of

   $$\begin{gathered} \{\text{bilinear maps } V \oplus W \longrightarrow U\} \\ \updownarrow \\ \{\text{linear maps } V \otimes W \longrightarrow U\} \end{gathered}$$

   Note that if $U = \mathbb{C}$, we have

   $$(V \otimes W)^* = \{\text{bilinear maps } V \oplus W \longrightarrow \mathbb{C}\}$$

The tensor product of two representations V and W yield a natural representation given by

$$g(v \otimes w) = gv \otimes gw$$

Lecture 21 — 3/19/12

Let V be a representation of a finite group G with basis $e_1,\dots,e_n$. We will notate

$$V^{\otimes k} = \underbrace{V \otimes \cdots \otimes V}_{k}$$

Definition 21.1. The k-th symmetric power of V, denoted

$$\mathrm{Sym}^k V \subset V^{\otimes k}$$

is the subspace invariant under S_k. For $v_1,\dots,v_k \in V$, we write

$$v_1 \cdots v_k = \frac{1}{k!} \sum_{\sigma \in S_k} v_{\sigma(1)} \otimes \cdots \otimes v_{\sigma(k)} \in V^{\otimes k}$$

Then $v_1 \cdots v_k \in \mathrm{Sym}^k V$, and

$$\{e_{i_1} \cdots e_{i_k} : 1 \le i_1 \le \cdots \le i_k \le n\}$$

is a basis for $\mathrm{Sym}^k V$.

Definition 21.2. The k-th alternating power of V, denoted

$$\wedge^k V \subset V^{\otimes k}$$

is the subspace skew-invariant under S_k; that is, where $\sigma(v) = \varepsilon_\sigma v$ for $\sigma \in S_k$, with

$$\varepsilon_\sigma = \mathrm{sgn}(\sigma)$$

For $v_1,\dots,v_k \in V$, we write

$$v_1 \wedge \cdots \wedge v_k = \frac{1}{k!} \sum_{\sigma \in S_k} \varepsilon_\sigma v_{\sigma(1)} \otimes \cdots \otimes v_{\sigma(k)} \in V^{\otimes k}$$

Then $v_1 \wedge \cdots \wedge v_k \in \wedge^k V$, and

$$\{e_{i_1} \wedge \cdots \wedge e_{i_k} : 1 \le i_1 < \cdots < i_k \le n\}$$

is a basis for $\wedge^k V$. Note that we have $<$ rather than $\le$ because a wedge product where any of the terms are equal will be zero, since transposing them must alternate the sign.

Proposition 21.3. *Let V and W be representations of a finite group G. Then the following formulas hold*

1. $\chi_{V\oplus W} = \chi_V + \chi_W$
2. $\chi_{V\otimes W} = \chi_V \cdot \chi_W$
3. $\chi_{V^*} = \overline{\chi_V}$
4. $\chi_{\mathrm{Sym}^2 V}(g) = \dfrac{\chi_V(g)^2 + \chi_V(g^2)}{2}$
5. $\chi_{\wedge^2 V}(g) = \dfrac{\chi_V(g)^2 - \chi_V(g^2)}{2}$

Proof. Suppose g has eigenvalues $\{\lambda_i\}$ as an endomorphism of V and $\{\mu_j\}$ as an endomorphism of W. Then $\{\lambda_i\}\cup\{\mu_j\}$ and $\{\lambda_i\cdot\mu_j\}$ are the eigenvalues for g on $V\oplus W$ and $V\otimes W$ respectively, which proves (1) and (2). We know that $\{\lambda_i\}$ are n-th roots of unity, for n the order of g. Hence, $\{\lambda_i^{-1} = \overline{\lambda_i}\}$ are the eigenvalues for g on V^*, which proves (3). For g on $\mathrm{Sym}^2 V$, the eigenvalues are $\{\lambda_i\lambda_j\}_{i\le j}$, so

$$\begin{aligned}\chi_{\mathrm{Sym}^2 V}(g) &= \sum_{i\le j}\lambda_i\lambda_j \\ &= \frac{\left(\sum\lambda_i\right)^2 + \sum\lambda_i^2}{2} \\ &= \frac{\chi_V(g)^2 + \chi_V(g^2)}{2}\end{aligned}$$

Similarly, g on $\wedge^2 V$ has eigenvalues $\{\lambda_i\lambda_j\}_{i<j}$, so

$$\begin{aligned}\chi_{\wedge^2 V}(g) &= \sum_{i<j}\lambda_i\lambda_j \\ &= \frac{\left(\sum\lambda_i\right)^2 - \sum\lambda_i^2}{2} \\ &= \frac{\chi_V(g)^2 - \chi_V(g^2)}{2}\end{aligned}$$

■

Definition 21.4. Let $\mathbb{C}^{\mathscr{C}} \subset \mathbb{C}^G$ denote the class functions, those functions $G \longrightarrow \mathbb{C}$ whose value on an element $g \in G$ is determined entirely by the conjugacy class of g.

If V and W are vector spaces, then we have

$$\mathrm{Hom}(V, W) \cong V^* \otimes W$$

via the map

$$\begin{pmatrix}\varphi : V \to W \\ v \mapsto \ell(v)\cdot w\end{pmatrix} \longleftarrow\!\shortmid\ \ell \otimes w$$

It follows from this, and the first projection formula, that, if V and W are irreducible representations of a finite group G, then

$$\begin{aligned}\dim \mathrm{Hom}_G(V, W) &= \frac{1}{|G|}\sum_{g\in G}\chi_{V^*\otimes W}(g) \\ &= \frac{1}{|G|}\sum_{g\in G}\overline{\chi_V(g)}\cdot\chi_W(g) \\ &= \begin{cases}1 & V \cong W \\ 0 & \text{otherwise}\end{cases}\end{aligned}$$

To describe this equation, we can define a Hermitian inner product on $\mathbb{C}^{\mathscr{C}}$ by

$$(\chi, \psi) = \frac{1}{|G|}\sum_{g\in G}\overline{\chi(g)}\cdot\psi(g)$$

Then we have proven that

Theorem 21.5. *In the inner product space of $\mathbb{C}^{\mathscr{C}}$ given by (χ, ψ), the characters of the irreducible representations of a finite group G are orthonormal.*

This has a number of consequences.

Corollary 21.6. *There are only finitely many irreducible representations,* at most *the number $|\mathscr{C}|$ of conjugacy classes of G.*

We will show shortly that, in fact, equality holds between the number of irreducible representations and the number of conjugacy classes of G.

Corollary 21.7. *Let $V_1, \ldots, V_k$ be all the irreducible representations of G. Then if V is any representation of G,*

$$V = \bigoplus V_i^{\oplus a_i}$$

where $a_i = (\chi_V, \chi_{V_i})$. Moreover, we have

$$\chi_V = \sum a_i\chi_{V_i}$$

and since the χ_{V_i} are linearly independent (by orthogonality), V is uniquely determined by its character χ_V.

Proof. The formula for a_i is achieved by decomposing χ_V into a sum of irreducible characters, and expanding by linearity. The number of copies of V_i will be the number of times a term (χ_{V_j}, χ_{V_i}) evaluates to 1. ■

Corollary 21.8. *A representation V is irreducible iff $(\chi_V, \chi_V) = 1$.*

Proof. The forward implication is obvious. The reverse is true because otherwise, V is a direct sum of irreducible, in which case we can decompose the inner product until it is expressed purely in terms of $\{(\chi_{V_i}, \chi_{V_j})\}$ where the V_k are irreducible representations. The sum of these will necessarily be strictly greater than 1. ■

Corollary 21.9. *Let V be an irreducible representation, U any one-dimensional representation. Then $U \otimes V$ is irreducible.*

Proof. We have

$$\begin{aligned}(\chi_{U\otimes V}, \chi_{U\otimes V}) &= \frac{1}{|G|}\sum_{g\in G} \overline{\chi_{U\otimes V}(g)} \cdot \chi_{U\otimes V}(g) \\ &= \frac{1}{|G|}\sum_{g\in G} \overline{\chi_U(g)} \cdot \chi_U(g) \cdot \overline{\chi_V(g)} \cdot \chi_V(g) \\ &= \frac{1}{|G|}\sum_{g\in G} \chi_U(g)^{-1} \cdot \chi_U(g) \cdot \overline{\chi_V(g)} \cdot \chi_V(g) \\ &= (\chi_V, \chi_V) \\ &= 1\end{aligned}$$

which yield irreducibility, as desired. ■

Theorem 21.10 (Fixed Point Theorem)**.** *If V is the permutation representation associated with the action of G on a finite set X, then*

$$\chi_V(g) = \#\{x \in X : gx = x\}$$

that is, the character of g is the number of elements in X fixed by g.

Proof. Recall that V has basis $\{e_x : x \in X\}$ and an element $g \in G$ permutes the basis vectors according to its permutation of X. Then g has 1's on the diagonal exactly when $g(e_x) = e_x$ and 0's elsewhere; hence, its character is exactly the number of fixed points. ■

Observation 21.11. Let R be the regular representation of G. Recall that it has basis $\{e_g : g \in G\}$, with

$$g : e_h \mapsto e_{gh}$$

By the above theorem, the character of R is given by

$$\chi_R(g) = \begin{cases} |G| & g = e \\ 0 & \text{otherwise} \end{cases}$$

We first note that R is not irreducible if $G \neq \{e\}$. Now if V_i are the irreducible representations and a_i their multiplicity in the decomposition of R, we have

$$\begin{aligned}a_i &= (\chi_R, \chi_{V_i}) \\ &= \frac{1}{|G|}\sum_{g\in G} \overline{\chi_R(g)} \cdot \chi_{V_i}(g) \\ &= \frac{1}{|G|}(|G| \cdot \chi_{V_i}(e)) \\ &= \chi_{V_i}(e) \\ &= \dim V_i\end{aligned}$$

So every irreducible representation V_i appears in R exactly $\dim V_i$ times. Therefore,

$$|G| = \dim R = \sum_i \dim(V_i)^2$$

Note also that, for $g \neq e$, we have

$$\begin{aligned}0 &= \chi_R(g) \\ &= \sum_i a_i \chi_{V_i}(g) \\ &= \sum_i \dim V_i \chi_{V_i}(g) \\ &= \sum_i \chi_{V_i}(e)\chi_{V_i}(g)\end{aligned}$$

These formulas are useful in filling out the character table for a given group.

Example. Let $G = S_3$ and suppose we want to determine the character of $V \otimes V$ (where V is the standard representation). Then we simply take the square of χ_V, yielding

S_3	$\overset{1}{e}$	$\overset{3}{(1\ 2)}$	$\overset{2}{(1\ 2\ 3)}$
U	1	1	1
U'	1	-1	1
V	2	0	-1
$V \otimes V$	4	0	1

Let us now apply these theorems by writing out the character tables for S_4 and A_4. Let U and U' denote the standard and alternating representations, whose characters we can deduce immediately.

S_4	$\overset{1}{e}$	$\overset{6}{(1\ 2)}$	$\overset{8}{(1\ 2\ 3)}$	$\overset{6}{(1\ 2\ 3\ 4)}$	$\overset{3}{(1\ 2)(3\ 4)}$
U	1	1	1	1	1
U'	1	-1	1	-1	1

Now let V be the standard representation given by

$$V = \{(x, y, z, w) \in \mathbb{C} : x + y + z + w = 0\}$$

with $\mathbb{C}^4 = U \oplus V$. We know from the fixed point theorem that $\mathbb{C}^4$ has character $\chi_{\mathbb{C}^4} = (4\ 2\ 1\ 0\ 0)$, and by subtracting the character of U, we get

S_4	$\overset{1}{e}$	$\overset{6}{(1\ 2)}$	$\overset{8}{(1\ 2\ 3)}$	$\overset{6}{(1\ 2\ 3\ 4)}$	$\overset{3}{(1\ 2)(3\ 4)}$
U	1	1	1	1	1
U'	1	-1	1	-1	1
V	3	1	0	-1	-1

We can check that V is irreducible by taking the inner product of the character with itself:

$$\begin{aligned}(\chi_V, \chi_V) &= \frac{1}{24}\big(1(3^2) + 6(1^2) + 8(0^2) \\ &\qquad + 6((-1)^2) + 3((-1)^2)\big) \\ &= \frac{1}{24}(9 + 6 + 6 + 3) \\ &= 1\end{aligned}$$

Note that the sum of the squares of the dimensions (which are listed in the first column, for e) is only

$$1^2 + 1^2 + 3^2 = 11 < 24$$

and hence we are not done enumerating our irreducible representations. There must be additional representations whose dimensions, when squared, sum to $24 - 11 = 13$. But this can only be partitioned as

$$2^2 + 3^2 = 13$$

We get another irreducible representation by taking the tensor product $U' \otimes V$.

	1	6	8	6	3
S_4	e	(1 2)	(1 2 3)	(1 2 3 4)	(1 2)(3 4)
U	1	1	1	1	1
U'	1	−1	1	−1	1
V	3	1	0	−1	−1
$U' \otimes V$	3	−1	0	1	−1

We know that $U' \otimes V$ is irreducible, and it is distinct since its trace is distinct. Our final irreducible, which we will call W, can be determined purely from orthogonality relations, in particular, the formula

$$\sum_i \chi_{V_i}(e)\chi_{V_i}(g) = 0$$

since we know that $\chi_W(e) = \dim W = 2$.

	1	6	8	6	3
S_4	e	(1 2)	(1 2 3)	(1 2 3 4)	(1 2)(3 4)
U	1	1	1	1	1
U'	1	−1	1	−1	1
V	3	1	0	−1	−1
$U' \otimes V$	3	−1	0	1	−1
W	2	0	−1	0	2

The character allows us to determine the form of W. Since (1 2)(3 4) is an involution[6] with trace 2, and since its eigenvalues must be roots of unity, we decompose $2 = 1+1$ and hence (1 2)(3 4) acts as the identity. Note that (1 2)(3 4) generates the Klein four group, and we have the quotient $S_4/V_4 \cong S_3$. Then we see that W is the standard representation of S_3 pulled back along this quotient:

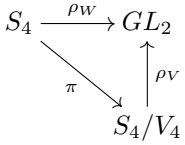

Note that, in general, if $N \triangleleft G$ is a normal subgroup, a representation $\rho : G \longrightarrow \mathrm{GL}(V)$ is trivial on N iff it factors through the quotient

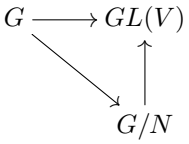

Now consider A_4. Note that (1 2) and (1 2 3 4) are simply not found in A_4, and (1 2 3) splits into two conjugacy classes, (1 2 3) and (1 3 2). We can easily check that U and V remain irreducible, while U' collapses into U, and similarly $U' \otimes V$ collapses into $U \otimes V = V$.

	1	4	4	3
A_4	e	(1 2 3)	(1 3 2)	(1 2)(3 4)
$U \cong U'$	1	1	1	1
V	3	0	0	−1

This only accounts for

$$1^2 + 3^2 = 10$$

squared degrees, so we are short two one-dimensional representation. Note that we have

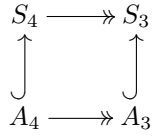

with $A_3 \cong \mathbb{Z}/3$. From this and the details below on abelian groups, we get

	1	4	4	3
A_4	e	(1 2 3)	(1 3 2)	(1 2)(3 4)
$U \cong U'$	1	1	1	1
V	3	0	0	−1
W_1	1	ω	ω^2	1
W_2	1	ω^2	ω	1

Observation 21.12. Let G be an abelian group. Note that for an arbitrary $g \in G$, the map $g : V \longrightarrow V$ determined by a representation ρ is not, in general, a morphism of representations, since

$$g(h(v)) \neq h(g(v))$$

[6] An *involution* is a function that is its own inverse.

in general. However, if $g \in Z(G)$,[7] then the above equality does hold, and hence g is a morphism. But $G = Z(G)$ in abelian groups. If V is irreducible, then by Schur's lemma, every $g \in G$ acts on V by a scalar multiple of the identity. Then every subspace of V is invariant, so $\dim V = 1$, and we will have

$$\chi_V(g) = \zeta_{\text{ord}(g)}$$

The irreducible representations ρ of G are all therefore elements of the dual group, which is the group of homomorphisms

$$\rho : G \longrightarrow \mathbb{C}^*$$

Lecture 22 — 3/21/12

Theorem 22.1. *Let $V_1, \dots, V_k$ be the irreducible representations of a finite group G. Then $\{\chi_{V_1}, \dots, \chi_{V_k}\}$ form a basis for the space of class functions $\mathbb{C}^{\mathscr{C}}$.*

Proof. Recall that, when we consider a representation V of G, we can average the elements of G (considered as elements of $\text{End}(V)$) to obtain a G-module homomorphism. For instance, we know that

$$\varphi = \frac{1}{|G|} \sum_{g \in G} g : V \to V$$

is a projection $V \longrightarrow\!\!\!\!\rightarrow V^G$. For all representations V and functions $\alpha \in \mathbb{C}^G$, let us define

$$\varphi_{\alpha,V} = \frac{1}{|G|} \sum_{g \in G} \alpha(g) \cdot g : V \to V$$

To generalize our method of averaging over G, we have the following claim:

Claim 22.2. *$\varphi_{\alpha,V}$ is a G-module homomorphism iff $\alpha \in \mathbb{C}^{\mathscr{C}}$.*

Proof. Consider the reverse direction; we want to show that $\forall h \in G$, $v \in V$, $\varphi(hv) = h\varphi(v)$. We have

$$\begin{aligned}
\varphi(hv) &= \frac{1}{|G|} \sum_{g \in G} \alpha(g) \cdot ghv \\
&= \frac{1}{|G|} \sum_{g \in G} \alpha(hgh^{-1}) \cdot hgh^{-1}hv \\
\text{since } \alpha \in \mathbb{C}^{\mathscr{C}} \quad &= \frac{1}{|G|} \sum_{g \in G} \alpha(g) \cdot hgv \\
&= h\varphi(v)
\end{aligned}$$

Now take the forward direction. Assume that $\forall h \in G$, $v \in V$, we have

$$\begin{gathered}
\varphi(hv) = h\varphi(v) \\
\frac{1}{|G|} \sum_{g \in G} \alpha(g) \cdot ghv = \frac{1}{|G|} \sum_{g \in G} \alpha(g) \cdot hgv \\
\frac{1}{|G|} \sum_{g \in G} \alpha(g) \cdot h^{-1}ghv = \frac{1}{|G|} \sum_{g \in G} \alpha(g) \cdot gv
\end{gathered}$$

Suppose that V is the regular representation and $v = e_e$.

$$\begin{gathered}
\frac{1}{|G|} \sum_{g \in G} \alpha(g) \cdot h^{-1}ghe_e = \frac{1}{|G|} \sum_{g \in G} \alpha(g) \cdot ge_e \\
\frac{1}{|G|} \sum_{g \in G} \alpha(g) \cdot e_{h^{-1}gh} = \frac{1}{|G|} \sum_{g \in G} \alpha(g) \cdot e_g
\end{gathered}$$

Note that all summands on each side of the equation will be linearly independent. If we choose h such that $\alpha(h^{-1}gh) \neq \alpha(g)$, then we will have $\alpha(g) \cdot e_{h^{-1}gh}$ on the LHS and $\alpha(h^{-1}gh) \cdot e_{h^{-1}gh}$ on the RHS. Then φ is not a G-module homomorphism, and this completes our proof.

When we average over g, we use $\varphi_{\alpha,V}$ with the constant function $\alpha : g \mapsto 1$; this projects V onto V_1, taking V_1 to be the trivial representation (which is irreducible for any group G). Using $\alpha = \chi_{V_i}$, $\varphi_{\alpha,V}$ similarly becomes a projection $V \longrightarrow\!\!\!\!\rightarrow V_i$.

Claim 22.3. *If $\alpha \in \mathbb{C}^{\mathscr{C}}$ and $(\alpha, \chi_{V_i}) = 0$, $\forall i$, then $\alpha = 0$.*

In other words, we claim that the χ_{V_i} span $\mathbb{C}^{\mathscr{C}}$, and since they are independent by orthogonality, they form a basis.

Proof. Let V be irreducible, $n = \dim V$. By Schur's lemma, the only G-module homomorphisms are scalar multiples of the identity, so for some $\lambda \in \mathbb{C}$, we have

$$\varphi_{\alpha,V} = \lambda I$$

We wish to determine λ. We have

$$\begin{aligned}
\lambda &= \frac{1}{n} \operatorname{tr}(\varphi_{\alpha,V}) \\
&= \frac{1}{n} \sum_{g \in G} \alpha(g) \cdot \chi_V(g) \\
&= \frac{|G|}{n} (\alpha, \chi_{V^*}) \\
&= 0
\end{aligned}$$

Then $\varphi_{\alpha,V} = \sum_{g \in G} \alpha(g) \cdot g = 0$ for all representations V. Consider $V = R$ the regular representation. In R, $\{\rho_R(g)\}_{g \in G}$ are linearly independent elements of $\text{End}(V)$. So it must be the case that $\alpha(g) = 0$ for all $g \in G$, and hence, $\alpha = 0$ as desired.

This completes the proof. ∎

[7] We denote $Z(G)$ as the *center* of G, the set of elements of G which commute with all elements; i.e., $Z(G) = \{z \in G : \forall g \in G, zg = gz\}$.

Claim 22.4. *Let V be any representation of G. Then*

$$V \otimes V = \operatorname{Sym}^2 V \oplus \wedge^2 V$$

Proof. Let $\{e_i\}$ be any basis for V. Let $g \in S_2$ be the nontrivial element, which acts as a transposition $e_i \otimes e_j \mapsto e_j \otimes e_i$, and by bilinearity, $v \otimes w \mapsto w \otimes v$ for any $v, w \in V$. g has order 2 and hence its eigenvalues are -1 and 1. Its eigenspaces are therefore precisely $\operatorname{Sym}^2 V$ and $\wedge^2 V$. But, setting $n = \dim V$, we have

$$\begin{aligned}\dim \operatorname{Sym}^2 V + \dim \wedge^2 V &= \frac{n(n+1)}{2} + \frac{n(n-1)}{2} \\ &= n^2 \\ &= \dim V \otimes V\end{aligned}$$

which gives our desired result. ■

Let us now take our theory of characters and apply it by computing the character table of S_5. Let U be the trivial representation, U' the alternating representation, and V the standard representation. We immediately get the characters for these representations and for $V \otimes U'$. We then begin searching for the other irreducible representations by taking tensor products, symmetric powers, and alternating powers of these irreducibles. Consider $V \otimes V$. We have

S_5	$\overset{1}{e}$	$\overset{10}{(1\ 2)}$	$\overset{20}{(1\ 2\ 3)}$	$\overset{30}{(1\ 2\ 3\ 4)}$	$\overset{24}{(1\ 2\ 3\ 4\ 5)}$	$\overset{15}{(1\ 2)(3\ 4)}$	$\overset{20}{(1\ 2)(3\ 4\ 5)}$
U	1	1	1	1	1	1	1
U'	1	-1	1	-1	1	1	-1
V	4	2	1	0	-1	0	-1
$V \otimes U'$	4	-2	1	0	-1	0	1
$V \otimes V$	16	4	1	0	1	0	1

Note that

$$\begin{aligned}(\chi_{V\otimes V}, \chi_{V\otimes V}) &= \frac{1}{120}(256 + 160 + 20 + 24 + 20) \\ &= 4\end{aligned}$$

and hence, $\chi_{V\otimes V}$ is not irreducible. We do know, however, that $V \otimes V$ decomposes as

$$V \otimes V = \operatorname{Sym}^2 V \oplus \wedge^2 V$$

Computing these characters, we have

S_5	$\overset{1}{e}$	$\overset{10}{(1\ 2)}$	$\overset{20}{(1\ 2\ 3)}$	$\overset{30}{(1\ 2\ 3\ 4)}$	$\overset{24}{(1\ 2\ 3\ 4\ 5)}$	$\overset{15}{(1\ 2)(3\ 4)}$	$\overset{20}{(1\ 2)(3\ 4\ 5)}$
U	1	1	1	1	1	1	1
U'	1	-1	1	-1	1	1	-1
V	4	2	1	0	-1	0	-1
$V \otimes U'$	4	-2	1	0	-1	0	1
$\wedge^2 V$	6	0	0	0	1	-2	0
$\operatorname{Sym}^2 V$	10	4	1	0	0	2	1

We can check that $\wedge^2 V$ is irreducible by taking its norm; we have

$$\begin{aligned}(\chi_{\wedge^2 V}, \chi_{\wedge^2 V}) &= \frac{1}{120}(36 + 24 + 60) \\ &= 1\end{aligned}$$

It follows from the decomposition of $V \otimes V$ that $\operatorname{Sym}^2 V$ must be the direct sum of three irreducible representations. We have thus accounted for five irreducibles, and the sum of squared degrees thus far is

$$1^2 + 1^2 + 4^2 + 4^2 + 6^2 = 70$$

We are still short 50 square degrees. This can be decomposed into two squares either as

$$50 = 1^2 + 7^2$$

or as

$$50 = 5^2 + 5^2$$

Suppose we have another one-dimensional irreducible representation. Any one-dimensional representation is a homomorphism

$$\rho : S_5 \longrightarrow \mathbb{C}^*$$

whose kernel will be a normal subgroup whose image is abelian. The only subgroup of S_5 with abelian quotient is A_5, so every one-dimensional representation of S_5 must factor through a representation of $S_5/A_5 \cong C_2$. Then U and U' are the only such representations.

So we have two irreducible representations remaining, both of degree 5. Consider again $\mathrm{Sym}^2 V$. $\mathrm{Sym}^2 V$ decomposes as three irreducibles, and we can determine that one is U and another is V by computing

$$(\chi_{\mathrm{Sym}^2 V}, \chi_U) = 1 \qquad (\chi_{\mathrm{Sym}^2 V}, \chi_V) = 1$$

We can check that the remaining irreducible representation has degree 5; call it W. Taking

$$\chi_W = \chi_{\mathrm{Sym}^2 V} - (\chi_U + \chi_V)$$

we obtain our full character table as

S_5	1 e	10 $(1\ 2)$	20 $(1\ 2\ 3)$	30 $(1\ 2\ 3\ 4)$	24 $(1\ 2\ 3\ 4\ 5)$	15 $(1\ 2)(3\ 4)$	20 $(1\ 2)(3\ 4\ 5)$
U	1	1	1	1	1	1	1
U'	1	-1	1	-1	1	1	-1
V	4	2	1	0	-1	0	-1
$V \otimes U'$	4	-2	1	0	-1	0	1
$\wedge^2 V$	6	0	0	0	1	-2	0
W	5	-1	-1	1	0	1	-1
$W \otimes U'$	5	1	-1	-1	0	1	1

Lecture 23 — 3/23/12

Suppose that $H \leq G$ is a subgroup of a finite group. We would like to somehow relate the representations of G to those of H. One direction of this relationship is easy:

Definition 23.1. Let V be a representation of G. By restricting the action of G to $H \leq G$, we can naturally restrict V to a representation of H denoted

$$W = \mathrm{Res}^G_H V$$

We see that Res^G_H is an operator mapping

$$\mathrm{Res}^G_H : \{\text{reps of } G\} \longrightarrow \{\text{reps of } H\}$$

Note that by restricting a representation to a subgroup H, new invariant subspaces may be created. Contrarily, distinct representations on G may become isomorphic on H.

What we want now is the relationship in the other direction; from a representation W of $H \leq G$, we want to induce a representation of G.

Observation 23.2. Suppose that V is a representation of G, $W \subset V$ an H-invariant subspace. Note that, for all $g \in G$, the subspace

$$gW = \{g \cdot w : w \in W\}$$

depends only on the left coset gH of g, since

$$(gh)W = g(hW) = gW$$

Thus, for $\sigma \in G/H$ a left coset of H, we will write $\sigma W = g_\sigma W$ for $g_\sigma \in \sigma$ any representative. This motivates the next definition.

Definition 23.3. A representation V of G is induced from a representation W of $H \leq G$ if

$$V = \bigoplus_{\sigma \in G/H} \sigma W$$

That is, if any $v \in V$ can be written uniquely as a sum of elements in these copies of W. In this case, we write

$$V = \mathrm{Ind}^G_H W$$

Example.

1. G acts on its left H-cosets G/H by left multiplication. Let V be the permutation representation of this action, with basis $\{e_\sigma\}_{\sigma \in G/H}$. Denote $[e]$ as the identity coset $[e] = H$. The subspace $W = \langle e_{[e]} \rangle$ is invariant under H, and moreover,

   $$\sigma \langle e_{[e]} \rangle = \langle e_\sigma \rangle, \qquad \forall \sigma \in G/H$$

   Then we have $V = \bigoplus_{\sigma \in G/H} \sigma W$, and hence V is induced from W, which is the trivial representation on H.

2. Let R_G be the regular representation of G, which has basis $\{e_g\}_{g \in G}$. Take R_H to be the subspace given by basis $\{e_h\}_{h \in H}$. A similar argument easily demonstrates that R_G is induced from R_H.

Theorem 23.4. *Let $H \leq G$, W a representation of H. Then there exists a unique (up to isomorphism) induced representation $V = \mathrm{Ind}^G_H W$.*

We thus have a corresponding operator

$$\mathrm{Ind}^G_H : \{\text{reps of } H\} \longrightarrow \{\text{reps of } G\}$$

Note that this *is not an inverse* to Res^G_H! In general, a representation is not induced by its restriction.

Proof. First, we show uniqueness. Let us start with a representation V of G, $W \subset V$ invariant under H, and

$$V = \bigoplus_{\sigma \in G/H} \sigma W$$

We claim that the action of G on V is determined entirely by the action of H on W and on the group $H \leq G$ itself. Choose a representative $g_\sigma \in \sigma$ for each coset (taking $g_{[e]} = e$). Given some $g \in G$ and $\sigma \in G/H$, say $g\sigma = \tau$. Then we can write $g \cdot g_\sigma = g_\tau \cdot h$ for some $h \in H$.

We know that each $v \in V$ can be written

$$v = \sum_{\sigma \in G/H} g_\sigma w_\sigma$$

for $w_\sigma \in W$. But we have

$$g(g_\sigma w_\sigma) = g_\tau h w_\sigma$$

So $g(v) = \sum g_\tau h w_\sigma$. This indeed determines a unique the group action on V, and this construction also proves existence. ∎

Let us compute the character table of $A_5 \leq S_5$. Note that the odd permutations are not present in A_5; moreover, (1 2 3 4 5) breaks up into two conjugacy classes, (1 2 3 4 5) and (2 1 3 4 5).

We know that the trivial and alternating representations on S_5 collapse in A_5, and we have $U \cong U'$. Then two other pairs of irreducibles also collapse, namely $V \cong V \otimes U'$ and $W \cong W \otimes U'$. We can restrict the remaining four irreducibles of S_5 to A_5 to get

A_5	^{1} e	^{20} (1 2 3)	^{15} (1 2)(3 4)	^{12} (1 2 3 4 5)	^{12} (2 1 3 4 5)
U	1	1	1	1	1
V	4	1	0	-1	-1
W	5	-1	1	0	0
$\wedge^2 V$	6	0	-2	1	1

Taking the norms of χ_U, χ_V, and χ_W confirms their irreducibility. $\wedge^2 V$ then cannot be irreducible since the sum of square degrees would exceed $|A_5|$. Currently, we have accounted for

$$1^2 + 4^2 + 5^2 = 42$$

Hence, we have remaining

$$60 - 42 = 18 = 3^2 + 3^2$$

We can check that

$$(\chi_{\wedge^2 V}, \chi_{\wedge^2 V}) = 2$$

We then claim that $\wedge^2 V = Y \oplus Z$ where Y and Z are our remaining degree 3 irreducible representations.

To show this, consider the automorphism of A_5 given by conjugation with (1 2). Observer that, while this is an inner automorphism for S_5, it is an outer automorphism for A_5 since (1 2) $\notin A_5$. This automorphism fixes the conjugacy classes e, (1 2 3), and (1 2)(3 4) and exchanges the classes (1 2 3 4 5) and (2 1 3 4 5). A group automorphism additionally acts on the set of representations of a group. We start with the characters and compose them with conjugation by (1 2)—this has the effect of switching the characters in the last two columns.

Note that, since the characters form a basis, one of the irreducible representations, say Y, must differe in the last two columns of the character table. Then by composing χ_Y with the outer automorphism of conjugation with (1 2), we get another character χ_Z, which is identical to χ_Y but with the last two columns switched. This demonstrates that we must, indeed, have $\wedge^2 V = Y \oplus Z$ and not $Y \oplus Y$ or $Z \oplus Z$, because doubling the characters in the last two columns of Y could not yield two 1's (and the same is true for Z).

Thus, our character table is

A_5	^{1} e	^{20} (1 2 3)	^{15} (1 2)(3 4)	^{12} (1 2 3 4 5)	^{12} (2 1 3 4 5)
U	1	1	1	1	1
V	4	1	0	-1	-1
W	5	-1	1	0	0
Y	3	0	-1	α	$1-\alpha$
Z	3	0	-1	$1-\alpha$	α
$\wedge^2 V$	6	0	-2	1	1

deduced from the fact that $\chi_Y + \chi_Z = \chi_{\wedge^2 V}$. We can use orthogonality relations to solve for α; this yields

$$\alpha = \varphi = \frac{1+\sqrt{5}}{2}$$

and so we have

A_5	^{1} e	^{20} (1 2 3)	^{15} (1 2)(3 4)	^{12} (1 2 3 4 5)	^{12} (2 1 3 4 5)
U	1	1	1	1	1
V	4	1	0	-1	-1
W	5	-1	1	0	0
Y	3	0	-1	φ	$\hat{\varphi}$
Z	3	0	-1	$\hat{\varphi}$	φ

Lecture 24 — 3/26/12

Definition 24.1. For every $g \in G$, we get an automorphism of G given by conjugation with g

$$\begin{aligned} G &\longrightarrow G \\ h &\longmapsto ghg^{-1} \end{aligned}$$

This yields a group homomorphism

$$G \longrightarrow \mathrm{Aut}(G)$$

which partitions $\mathrm{Aut}(G)$ into a normal subgroup $\mathrm{Inn}(G)$ of inner automorphisms given by conjugation and a complementary collection of outer automorphisms.

Observation 24.2. Let $\tau : G \to G$ be an automorphism, V any representation of G. We can obtain another representation V^τ by composing ρ_V with τ.

$$G \xrightarrow{\tau} G \xrightarrow{\rho_V} \mathrm{GL}(V), \qquad \rho_{V^\tau}$$

If τ is an inner automorphism, then

$$V^\tau = V, \quad \forall V$$

If τ is outer, then it may permute the irreducible representations nontrivially, e.g., for $G = A_3$.

Let us explore the representations of the dihedral groups, $G = D_{2n}$, the group of isometries of the regular n-gon. We have the following picture for any dihedral group:

$$C_n \cong \mathbb{Z}_n \hookrightarrow D_{2n} \twoheadrightarrow \mathbb{Z}_2$$

Choose $\zeta = \zeta_n$ any primitive nth root of unity. Let h be a rotation through $2\pi/n$, g any reflection. Then $h^n = e$, $g^2 = e$, $gh = hg$, and $D_{2n} = \langle h, g \rangle$.

Let V be any irreducible representation. V has an eigenbasis with respect to h, with eigenvalues ζ^k. Suppose that $v \in V$ is an h-eigenvector,

$$hv = \zeta^k v$$

Then

$$\begin{aligned} h(gv) &= g(h^{-1}v) \\ &= g(\zeta^{-k}v) \\ &= \zeta^{-k}(gv) \end{aligned}$$

So $gv \in V$ is an h-eigenvector with eigenvalue ζ^{-k}. The subspace $\langle v, gv \rangle \subset V$ is thus G-invariant, and hence

$$V = \langle v, gv \rangle$$

For now we will only consider n odd. Then $\dim V = 2$ unless $k = 0$. So suppose that h acts as the identity. If $g = 1$, we get the trivial representation U; if $g = -1$, we get the alternating representation U'.

Note that as with our picture of $S_3 = D_6$, if h has eigenvectors v_1 and v_2 with eigenvalues ζ^k and ζ^{-k}, then g switches them. Thus,

$$h \longmapsto \begin{pmatrix} \zeta^k & 0 \\ 0 & \zeta^{-k} \end{pmatrix} \qquad g \longmapsto \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$$

Since we have $m = \frac{n-1}{2}$ pairs of nontrivial ζ^k, our character table is

D_{2n}	e	$h^{\pm 1}$	$h^{\pm 2}$	$\cdots$	$h^{\pm m}$	g
U	1	1	1	$\cdots$	1	1
U'	1	1	1	$\cdots$	1	-1
V_1	2	$\zeta + \zeta^{-1}$	$\zeta^2 + \zeta^{-2}$	$\cdots$	$\zeta^m + \zeta^{-m}$	0
V_2	2	$\zeta^2 + \zeta^{-2}$	$\zeta^4 + \zeta^{-4}$	$\cdots$	$\zeta + \zeta^{-1}$	0
$\vdots$						

We return now to the study of induced representations. We begin with some basic properties.

Proposition 24.3. *The operator* Ind_H^G *satisfies the following properties:*

1. *Linearity.*
$$\mathrm{Ind}_H^G(W_1 \oplus W_2) = \mathrm{Ind}_H^G(W_1) \oplus \mathrm{Ind}_H^G(W_2)$$
2. *Transitivity. For* $H \le K \le G$,
$$\mathrm{Ind}_H^G(W) = \mathrm{Ind}_K^G(\mathrm{Ind}_H^K(W))$$
3. *Push-Pull. Let* U *be a representation of* G, W *a representation of* H.
$$U \otimes \mathrm{Ind}_H^G(W) = \mathrm{Ind}_H^G(\mathrm{Res}_H^G(U) \otimes W)$$

Proof. We will prove only the push-pull property. $U \otimes \mathrm{Ind}_H^G W =: V$. $U \otimes W_0$ is H-invariant $V = \bigoplus_{\sigma \in G/H} \sigma(U \otimes W_0)$ ∎

Proposition 24.4. *Let* U *be a representation of* G, W *a representation of* $H \le G$. *Then*

$$\mathrm{Hom}_H(W, \mathrm{Res}_H^G U) = \mathrm{Hom}_G(\mathrm{Ind}_H^G W, U)$$

Proof. We claim that any H-module homomorphism $\varphi : W \to U$ can be uniquely extended to G-module homomorphism $\hat{\varphi} : \mathrm{Ind}_H^G W \to U$. Recall that

$$V := \mathrm{Ind}_H^G W = \bigoplus_{\sigma \in G/H} \sigma W$$

Choose any $v \in \sigma W$, which will have the form $v = g_\sigma w$ for some representative g_σ. We define $\hat{\varphi}$ on σW by

$$\begin{aligned} \hat{\varphi}(v) &= g_\sigma \varphi g_\sigma^{-1}(v) \\ &= g_\sigma \varphi g_\sigma^{-1}(g_\sigma w) \\ &= g_\sigma \varphi(w) \\ &= \varphi(g_\sigma w) \\ &= \varphi(v) \end{aligned}$$

and clearly we have independence of our choice of representative g_σ. This completes the proof. ∎

Corollary 24.5 (Frobenius Reciprocity)**.** *If U is a representation of G, W a representation of $H \leq G$, then*

$$(\chi_W, \chi_{\mathrm{Res}^G_H U})_H = (\chi_{\mathrm{Ind}^G_H W}, \chi_U)_G$$

This is equivalent to the claim

$$\dim(\mathrm{Hom}_H(W, \mathrm{Res}^G_H U)) = \dim(\mathrm{Hom}_G(\mathrm{Ind}^G_H W, U))$$

Lecture 25 — 3/28/12

Definition 25.1. Let G be any finite group. Define

$$R(G) = \left\{\sum a_i V_i : V_i \text{ rep of } G\right\} / \langle V \oplus W - V - W\rangle$$

This is the free abelian group on the isomorphism classes of irreducible representations of G; note that

$$R(G) \simeq \mathbb{Z}^c$$

where c is the number of conjugacy classes of G. Addition on $R(G)$ is given by $\oplus$, and multiplication $\otimes$.

Theorem 25.2 (Artin)**.** *The representations of G induced from cyclic subgroups of G generate a subgroup of finite index in $R(G)$.*

Theorem 25.3 (Brauer)**.** *We say a group $H = A \times B$ is* elementary *if A is cyclic and B is a p-group with $p \nmid |A|$. The representations of G induced from elementary subgroups of G generate $R(G)$.*

Definition 25.4. The vector space of quaternions is given by

$$\mathbb{H} = \{\alpha_1 + \alpha_2 i + \alpha_3 j + \alpha_4 k : \alpha_i \in \mathbb{R}\}$$

where multiplication is given by

$$i^2 = j^2 = k^2 = ijk = -1$$

Note that

$$\mathbb{H} \cong \mathbb{R}^4$$

and hence is a real vector space.

Let us compute the character table for the group of quaternions

$$Q = \{\pm 1, \pm i, \pm j, \pm k\}$$

We know that $\{\pm 1\} \triangleleft G$ is normal, and we have a sequence of quotient maps

$$G \twoheadrightarrow G/\{\pm 1\} \cong \mathbb{Z}/2 \times \mathbb{Z}/2 \twoheadrightarrow \mathbb{Z}/2$$

Note that $G/\{\pm 1\} = \{\bar{1}, \bar{i}, \bar{j}, \bar{k}\}$. Then the second quotient map kills $\bar{1}$ and one of $\{\bar{i}, \bar{j}, \bar{k}\}$, sending the remaining two to the nonidentity element of $\mathbb{Z}/2$, which has character -1. Thus, we get

Q	$+1$	-1	$\pm i$	$\pm j$	$\pm k$
U	1	1	1	1	1
U_i	1	1	1	-1	-1
U_j	1	1	-1	1	-1
U_k	1	1	-1	-1	1
V	2	-2	0	0	0

We get the final irreducible using orthogonality relations (also note that if all representations were one-dimensional, the group would be abelian).

We can identify this representation V as $\mathbb{H}$. We make $\mathbb{H}$ into a two-dimensional complex vector space via complex multiplication *on the right*. Identifying $\mathbb{C} = \langle 1, i\rangle$, we can take $\{1, j\}$ as a basis, whereupon we have

$$1 = (1, 0) \quad i = (i, 0) \quad j = (0, 1) \quad k = (0, -i)$$

The elements of Q act via left multiplication on $\mathbb{H}$. Since $\mathbb{H}$ is associative, this is indeed a representation, and we can check it against our character table.

Definition 25.5. A representation V of G is a real representation if we can write

$$V = V_0 \otimes_{\mathbb{R}} \mathbb{C}$$

with an action of G on V_0 that extends by linearity to the action of G on V. In other words, for some choice of isomorphism $\mathrm{GL}(V) \cong \mathrm{GL}_n \mathbb{C}$, the representation ρ factors through $\mathrm{GL}_n \mathbb{R}$:

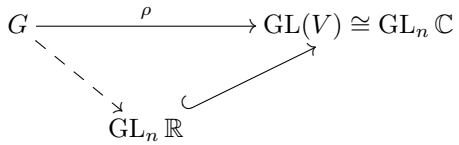

Let us look at the construction $V = V_0 \otimes_{\mathbb{R}} \mathbb{C}$ more closely. Say V_0 is a real vector space of $\dim_{\mathbb{R}} V_0 = n$. Then, taking $\mathbb{C}$ as a two-dimensional real vector space, $V = V_0 \otimes_{\mathbb{R}} \mathbb{C}$ is also a real vector space of $\dim_{\mathbb{R}} V = 2n$. However, V also has the structure of a complex vector space of $\dim_{\mathbb{C}} V = n$, given by

$$i(v \otimes \lambda) = v \otimes i\lambda$$

That is, if we write

$$V_0 = \{a_1 v_1 + \cdots + a_n v_n : a_i \in \mathbb{R}\}$$

we have

$$V = V_0 \otimes_{\mathbb{R}} \mathbb{C} = \{a_1 v_1 + \cdots + a_n v_n : a_i \in \mathbb{C}\}$$

V is called the *complexification* of V_0. We can see from this that any $g : V \to V \in G$ can be represented either as a complex matrix *or* as a real matrix with twice the dimension, if V is a real representation.

Example. Consider $G = \mathbb{Z}/3$ acting on

$$V = \{(x, y, z) \in \mathbb{C}^3 : x + y + z = 0\}$$

This representation is real because it breaks into $V_0 \otimes_{\mathbb{R}} \mathbb{C}$ for

$$V_0 = \{(x, y, z) \in \mathbb{R}^3 : x + y + z = 0\}$$

We know that $V_0 \simeq \mathbb{R}$; in this interpretation, G acts via rotation through 120°. Note that, setting $v = (1, \omega, \omega^2)$ and $w = (1, \omega^2, \omega)$, we have

$$V = \mathbb{C}\langle v \rangle \oplus \mathbb{C}\langle w \rangle$$

However, neither of these component vector spaces is real, so V is irreducible *as a real representation.*

Observation 25.6. Note that real representations have real character, but the converse is not true. For instance, the representation V of the quaternion group has real character, but is not real.

Definition 25.7. A symmetric bilinear form on a vector space V is a map

$$B : V \times V \to \mathbb{C}$$

such that, $\forall u, v, w \in V$, $\forall \lambda \in \mathbb{C}$,

1. $B(u + v, w) = B(u, w) + B(v, w)$
2. $B(\lambda v, w) = \lambda B(v, w)$
3. $B(u, v) = B(v, u)$

A bilinear form yields a map

$$\begin{aligned} \alpha_B : V &\longrightarrow V^* \\ v &\longmapsto B(v, \cdot) \end{aligned}$$

which takes a vector v to B pre-parametrized with v as an argument. We say that B is nondegenerate if this map α_B is an isomorphism.

Recall that if V is a representation of G, then V admits a positive definite Hermitian inner product H which is G-invariant; we took any H_0 and then averaged, yielding

$$H(v, w) = \frac{1}{|G|} \sum_{g \in G} H_0(gv, gw)$$

We would like to know whether we can also find a nondegenerate symmetric bilinear form B on V. Note that, if we try to choose B_0 and average, we get

$$B(v, w) = \frac{1}{|G|} \sum_{g \in G} B_0(gv, gw)$$

But note that B_0 cannot be positive definite over $\mathbb{C}$, since, for instance,

$$B_0(iv, iw) = -B_0(v, w)$$

We do not want B to be trivially zero anywhere, so we demand nondegeneracy. As it turns out, this is true if and only if V is real (since B_0 can be positive definite on the real vector space V_0).

Lecture 26 — 3/30/12

Let V be an irreducible representation of a finite group G. We know there exists a G-invariant positive definite Hermitian form $H : V \times V \to \mathbb{C}$ on V. That is, H satisfies

$$H(gv, gw) = H(v, w), \qquad \forall g \in G,\ \forall v, w \in V$$

We wish to determine whether there also exists a G-invariant nondegenerate symmetric bilinear form $B : V \times V \to \mathbb{C}$ on V. Every bilinear form gives a map

$$\begin{aligned} \alpha_B : V &\longrightarrow V^* \\ v &\longmapsto B(v, \cdot) \end{aligned}$$

That B is nondegenerate means that α_B is an isomorphism; that B is G-invariant means

$$B(gv, gw) = B(v, w), \qquad \forall g \in G,\ \forall v, w \in V$$

Theorem 26.1. *An irreducible representation V of a finite group G admits a nondegenerate G-invariant symmetric bilinear form iff V is a real representation.*

Proof. We begin with the following claim:

Lemma 26.2. *B is G-invariant iff α_B is a G-module homomorphism.*

Proof. To be a G-module homomorphism, by the definition of the dual representation, α_B must satisfy

$$\alpha_B(gv) = {}^t g^{-1} \alpha_B(v)$$

$\alpha_B(gv)$ and ${}^t g^{-1} \alpha_B(v)$ are respectively maps

$$w \longmapsto B(gv, w) \qquad w \longmapsto B(v, g^{-1} w)$$

Clearly, these are equal iff B is G-invariant.

Note that

$$\operatorname{Hom}(V, V^*) = V^* \otimes V^* = \operatorname{Sym}^2 V^* \oplus \wedge^2 V^*$$

so every bilinear form is a sum of a symmetric bilinear form with a skew-symmetric bilinear form. Since V is irreducible, this means that B must either be symmetric or skew-symmetric. We then have have one of the following three cases:

1. $V \not\cong V^*$, so there does not exist any nondegenerate G-invariant bilinear form on V

2. $V \cong V^*$ through a nondegenerate G-invariant symmetric bilinear form
3. $V \cong V^*$ through a nondegenerate G-invariant skew-symmetric bilinear form

We also know that

$$\dim(\mathrm{Hom}_G(V, V^*)) = \frac{1}{|G|} \sum_{g \in G} \overline{\chi_V(g)}^2$$

If the character is real, then this dimension must be nonzero because each summand $\overline{\chi_V(g)}^2 = \chi_V(g)^2$ will be positive. Hence, a real character implies that we are in either case 2 or case 3. Meanwhile, a complex character means we will be in case 1.

Our theorem then reduces to the following lemma:

Lemma 26.3. *V is real iff we have case 2.*

Proof. First suppose that V is real. This direction is easy. We have $V = V_0 \otimes_{\mathbb{R}} \mathbb{C}$, so we can find a positive definite symmetric bilinear form B_0 on V_0 and make it G-invariant by averaging. We then simply extend this by linearity to a form B on V.

Now consider the reverse direction. Suppose we have a nondegenerate G-invariant symmetric bilinear form B on V. Let H be any positive definite G-invariant Hermitian form. As with B, this yields a map

$$\begin{aligned} \alpha_H : V &\longrightarrow V^* \\ v &\longmapsto H(v, \cdot) \end{aligned}$$

Note, however, that unlike α_B which is linear, the map α_H is *conjugate* linear.

Take the composition

$$\varphi : V \xrightarrow{\alpha_B} V^* \xrightarrow{\alpha_h^{-1}} V$$

which is an automorphism of V. Now consider $\varphi^2 : V \to V$, which is a complex linear G-module homomorphism. By Schur's lemma,

$$\varphi^2 = \lambda I$$

Moreover, we can see by our construction that, $\forall v, w \in V$, φ satisfies

$$\begin{aligned} H(\varphi(v), w) &= B(v, w) \\ &= B(w, v) \\ &= H(\varphi(w), v) \\ &= \overline{H(v, \varphi(w))} \end{aligned}$$

So for φ^2, we have

$$\begin{aligned} H(\varphi^2(v), w) &= H(v, \varphi^2(w)) \\ H(\lambda v, w) &= H(v, \lambda w) \end{aligned}$$

Then $\lambda = \bar{\lambda}$ and hence λ is real (and positive).

Multiplying B by a scalar, we can assume $\lambda = 1$. Then $\varphi : V \to V$ is a real linear map, but $\varphi^2 = I$. Hence, it has has eigenvalues 1 or -1 So we can decompose

$$V = V^+ \oplus V^-$$

where V^+ is the φ^2-eigenspace of eigenvalue 1, and V^- is that of eigenvalue -1. Note that $v \in V^+$ means $\varphi(v) = v$, so

$$\varphi(iv) = -i\varphi(v)$$

Then we have $iV^+ = V^-$, and therefore we can write $V = V^+ \otimes_{\mathbb{R}} \mathbb{C}$, as desired.

We have actually proven a stronger theorem, which we state as:

Theorem 26.4. *An irreducible representation V of G is one and one of the following:*

1. Complex. *χ_V is not real, and V does not admit a G-invariant nondegenerate bilinear form.*
2. Real. *V is a real representation, χ_V is real, and V admits a G-invariant nondegenerate symmetric bilinear form.*
3. Quaternionic. *V is a real representation, χ_V is real, and V admits a G-invariant nondegenerate skew-symmetric bilinear form.*

Note that unlike over the complex numbers, when we work over the real numbers, we can't easily determine the number of representations. About all we can state in general is that the complex representations come in cojugate pairs. ■

Lecture 27 — 4/2/12

Definition 27.1. A module M over a ring[8] R is an abelian group with a scalar multiplication map

$$\begin{aligned} R \times M &\longrightarrow R \\ (r, x) &\longmapsto rx \end{aligned}$$

satisfying the usual axioms $\forall r, s \in R$, $\forall x, y \in M$,

- $r(sx) = (rs)x$
- $1 \cdot x = x$
- $(r + s)x = rx + sx$
- $r(x + y) = rx + ry$

Example.

[8] Any ring R we consider will be commutative with identity unless otherwise specified.

1. R is a module over R, $r \cdot s = rs$.
2. The product
$$R^n = \{(x_1, \dots, x_n) : x_i \in R\}$$
is a module given by the rule
$$r(x_1, \dots, x_n) = (rx_1, \dots, rx_n)$$
A module of this form is called a *free module*.

Definition 27.2. A submodule $N \subset M$ over R is a subgroup closed under scalar multiplication.
$$\begin{array}{ccc} R \times M & \longrightarrow & M \\ \cup & & \cup \\ R \times N & \longrightarrow & N \end{array}$$

Definition 27.3. If $N \subset M$ is a submodule, the quotient module M/N is the quotient group together with scalar multiplication rule
$$r\overline{x} = \overline{rx}$$

Observation 27.4. The submodules of the module $M = R$ are the ideals in R.

Definition 27.5. An R-module homomorphism $\varphi : M \to N$ is a group homomorphism that commutes with scalar multiplication
$$\varphi(rx) = r\varphi(x)$$

Observation 27.6. The kernel $\ker(\varphi)$ is a submodule of M, and likewise the image $\operatorname{im}(\varphi)$ is a submodule of N.

Definition 27.7. Let M, N be R-modules. The direct sum of modules is given by
$$M \oplus N = \{(x, y) : x \in M, y \in N\}$$
with scalar multiplication given by
$$r(x, y) = (rx, ry)$$

Definition 27.8. Let M, N be R-modules. The tensor product of modules is a R-module $M \otimes N$ with an associated bilinear map
$$\begin{aligned} \varphi : M \oplus N &\longrightarrow M \otimes N \\ (x, y) &\longmapsto x \otimes y \end{aligned}$$
such that every bilinear map $\psi : M \times N \to P$ for an R-module P factors uniquely through φ:

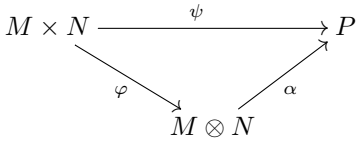

where α is an R-module homomorphism.

Note that for free modules, we have
$$\begin{aligned} R^m \oplus R^n &= R^{m+n} \\ R^m \otimes R^n &= R^{mn} \end{aligned}$$

Definition 27.9. Let M be an R-module, $x_1, \dots, x_n \in M$. We have a map $\varphi : R^n \to M$ given by
$$(a_1, \dots, a_n) \longmapsto a_1x_1 + \cdots + a_nx_n$$
We say that $x_1, \dots, x_n$ are generators of M if φ is surjective. We say M is finitely-generated if there exists a finite set of generators.

Example. Let us consider some modules over $\mathbb{Z}$.

1. $\mathbb{Z}$, $\mathbb{Z}^n$, and the algebraic integers are all $\mathbb{Z}$-modules.
2. Let M be any $\mathbb{Z}$-module. Then
$$mx = \underbrace{x + \cdots + x}_{m}$$
That is, the module structure is determined by the group structure. We thus get a bijection between abelian groups and $\mathbb{Z}$-modules.
3. $2\mathbb{Z} \subset \mathbb{Z}$ is a submodule of $\mathbb{Z}$; moreover,
$$2\mathbb{Z} \cong \mathbb{Z}$$
as $\mathbb{Z}$-modules, since we have
$$\begin{aligned} \mathbb{Z} &\xrightarrow{\sim} 2\mathbb{Z} \\ m &\longmapsto 2m \end{aligned}$$
We get $\mathbb{Z}/2$ as a the quotient module $\mathbb{Z}/2\mathbb{Z}$.
4. We claim that $\mathbb{Z}/2 \otimes \mathbb{Z}/3 = (0)$. For we have
$$\begin{aligned} 2(1 \otimes 1) &= 2 \otimes 1 = 0 \\ 3(1 \otimes 1) &= 1 \otimes 3 = 0 \end{aligned}$$
Then
$$1 \otimes 1 = 3(1 \otimes 1) - 2(1 \otimes 1) = 0$$
Note that this situation could never occur in a vector space.

Definition 27.10. Let S be any set. Then the free module generated by S is given by
$$R^S = \{a_1s_1 + \cdots + a_ns_n : a_i \in R, s_i \in S\}$$
In general, given $x_1, \dots, x_n \in M$, we say they generate M if
$$R^{\{x_i\}} \twoheadrightarrow M$$

Lecture 28 — 4/4/12

Definition 28.1. Let R be a ring. Define

$$\begin{aligned} M_n(R) &= \{n \times n \text{ matrices } (a_{ij}) : a_{ij} \in R\} \\ &= \mathrm{Hom}(R^n, R^n) \\ &\simeq R^{n^2} \end{aligned}$$

Let $\mathrm{GL}_n(R) \subset M_n(R)$ be the subset of invertible matrices; that is,

$$\begin{aligned} \mathrm{GL}_n(R) &= \{A \in M_n(R) : \exists B \in M_n(R), AB = I_n\} \\ &= \mathrm{Aut}(R^n) \end{aligned}$$

Definition 28.2. Let $A \in M_n(R)$ with entries (a_{ij}). The determinant of the matrix A is given by

$$\det A = \sum_{\sigma \in S_n} \mathrm{sgn}(\sigma) \cdot a_{1,\sigma(1)} \cdots a_{n,\sigma(n)}$$

Observe that

$$\det(AB) = \det(A) \cdot \det(B)$$

Theorem 28.3. *$A \in M_n(R)$ is an isomorphism (i.e., invertible) iff $\det A$ is a unit in R.*

Proof. First, suppose A is invertible. Then $AA^{-1} = I_n$, so we have

$$\det(A) \cdot \det(A^{-1}) = 1$$

from which it follows that $\det A$ is a unit.

Now let $\mathrm{cof}(A)$ be the matrix of cofactors of A. We know that

$$A \cdot \mathrm{cof}(A) = \det(A) \cdot I$$

since $\det A$ is a unit, it has an inverse $b \in R$, and hence $b \cdot \mathrm{cof}(A)$ is our desired inverse to A. ∎

We will denote the basis of R^n by $\{e_1, \ldots, e_n\}$, where

$$e_i = (\underbrace{0, \ldots, 1}_{i}, \ldots, 0)$$

Recall that

Definition 28.4. We say that $v_1, \ldots, v_k \in M$ generate a module M if every $v \in M$ is a linear combination of the v_i. Equivalently, the v_i generate M if the map

$$\begin{aligned} \varphi : R^k &\longrightarrow M \\ e_i &\longmapsto v_i \end{aligned}$$

is surjective. We say the v_i are linearly independent if

$$\sum a_i v_i = 0 \implies a_i = 0, \forall i$$

or equivalently, if φ is injective. A linearly independent generating set is a basis for M.

Observation 28.5. Note that modules exhibit pathological behavior as compared to vector spaces.

1. Not every module has a basis; for instance, $\mathbb{Z}/n$ as a $\mathbb{Z}$-module.
2. Even if a module M has a basis, (e.g., a free module $M \cong R^n$), it is not the case that every linearly independent set can be extended to a basis. For instance, take $M = \mathbb{Z}$ as a $\mathbb{Z}$-module; the vector $v_1 = 2$ cannot be made a part of a basis. That is, we have a $\mathbb{Z}$-module homomorphism

   $$\mathbb{Z} \xrightarrow{2x} \mathbb{Z}$$

   that is injective without being surjective.

   Similarly, if $M = \mathbb{Z} \times \mathbb{Z}$, the vectors $v_1 = (1, 1)$, $v_2 = (1, -1)$ are independent, but cannot be extended to a basis.
3. It is also not the case that every generating set contains a basis. For instance, $v_1 = 2$ and $v_2 = 3$ generate $M = \mathbb{Z}$ but do not contain a basis.

These pathologies make it impossible for modules to have a well-defined notion of dimension.

In general, a module M is free iff it has a basis (not necessarily a finite one).

Definition 28.6. If a module M has a basis $v_1, \ldots, v_n$ (i.e., $M = R^n$), we say that n is the rank of M.

Note that a module can have rank 0 without being the zero module (for instance, in $\mathbb{Z}/n$).

Lemma 28.7. *If a module M has a basis, then any two bases of M have the same cardinality. It follows also that if $n \neq m$, then $R^n \ncong R^m$.*

Let M be a finitely-generated free module with basis $v_1, \ldots, v_k$. Then we have a natural isomorphism $M \cong R^k$ given by

$$\begin{aligned} \varphi : R^k &\longrightarrow M \\ e_i &\longmapsto v_i \end{aligned}$$

If $v_1', \ldots, v_k'$ is another basis, we can write

$$v_i' = \sum a_{ij} v_j$$

for some $a_{ij} \in R$. Then we have

$$\begin{array}{ccc} R^k & \xrightarrow{\varphi} & M \\ {\scriptstyle P}\downarrow & & \| \\ R^k & \xrightarrow{\varphi'} & M \end{array}$$

where $P = (a_{ij})$ is called the change of basis matrix.

Now we can identify

$$\mathrm{Hom}(R^n, R^m) = M_{m\times n}(R)$$

For suppose M and N are free modules with ranks m and n respectively, $\varphi : N \to M$ an R-module homomorphism. Let us choose respective bases $v_1, \dots, v_m$ and $w_1, \dots, w_n$. Then we get an $m \times n$ matrix

$$\begin{array}{ccc} R^n & \xrightarrow{\sim} & N \\ \downarrow & & \downarrow \varphi \\ R^m & \xrightarrow{\sim} & M \end{array}$$

which gives us $\mathrm{Hom}(N, M) = M_{m\times n}(R)$.

Definition 28.8. The cokernel of an R-module homomorphism $\varphi : M \to N$ is defined as

$$\operatorname{coker}\varphi = N/\operatorname{im}\varphi$$

We would like now to determine the kernel and cokernel of a homomorphism $\varphi : R^n \to R^m$. The issue is that, while the kernel is again a free module (as we will show), the cokernel need not be. For instance,

$$\begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{pmatrix} : \mathbb{Z}^3 \longrightarrow \mathbb{Z}^3$$

Lecture 29 — 4/6/12

Let V and W be free modules with bases $v_1, \dots v_n$ and $w_1, \dots, w_m$. We can express an R-module homomorphism $\varphi : V \to W$ as a matrix $A = (a_{ij})$ with

$$Av_j = \begin{pmatrix} a_{1,1} & \cdots & a_{1,n} \\ \vdots & \ddots & \vdots \\ a_{m,1} & \cdots & a_{m,n} \end{pmatrix} \begin{pmatrix} 0 \\ \vdots \\ 1 \\ 0 \\ \vdots \end{pmatrix} = \begin{pmatrix} a_{1,j} \\ \vdots \\ a_{m,j} \end{pmatrix}$$

That is, we have

$$\varphi(v_j) = \sum a_{ij} w_i$$

Now suppose that we choose different bases $\{v_j'\}$ and $\{w_i'\}$ for V and W. Then in terms of these bases, φ corresponds to the matrix

$$\varphi \simeq Q^{-1}AP$$

where P is the change of basis matrix for $\{v_j\} \to \{v_j'\}$ and Q is likewise for $\{w_i\}$ and $\{w_i'\}$.

We now pose the following question: given a homomorphism φ, how can we find bases that make this matrix as simple as possible? Equivalently, given a matrix, can we find $Q \in \mathrm{GL}_m(R)$ and $P \in \mathrm{GL}_n(R)$ such that $Q^{-1}AP$ has a particularly simple form? Were V and W vector spaces, we would have such simple matrix forms; however, for modules, in general, the answer here is no. In the case of $R = \mathbb{Z}$, however, we have a nicer picture:

Theorem 29.1. *Let V and W be free modules over $\mathbb{Z}$, $\varphi : V \to W$ a $\mathbb{Z}$-module homomorphism. Then there exist bases $\{v_j\}_{j=1}^n$ for V and $\{w_i\}_{i=1}^m$ for W such that the matrix representation A of φ is the block matrix*

$$\left(\begin{array}{c|c} D & 0 \\ \hline 0 & 0 \end{array}\right)$$

where D has the form

$$D = \begin{pmatrix} d_1 & & & 0 \\ & d_2 & & \\ & & \ddots & \\ 0 & & & d_k \end{pmatrix}$$

such that $\forall i < k, d_i \,|\, d_{i+1}$.

Proof. For every $m \times n$ matrix A, we want Q and P such that $Q^{-1}AP$ has the desired form. We will construct Q and P using a set of elementary matrices (analogous to those from linear algebra). These have the form

$$\begin{pmatrix} 1 & & & & & & \\ & \ddots & & & & & \\ & & 1 & & c & & \\ & & & \ddots & & & \\ & & & & 1 & & \\ & & & & & \ddots & \\ & & & & & & 1 \end{pmatrix}$$

where $c \in \mathbb{Z}$ is in the i, j-th place; this adds the ith column of A scaled by c to the jth column (via multiplication on the right);

$$\begin{pmatrix} 1 & & & & & & \\ & \ddots & & & & & \\ & & 0 & & 1 & & \\ & & & \ddots & & & \\ & & 1 & & 0 & & \\ & & & & & \ddots & \\ & & & & & & 1 \end{pmatrix}$$

which is the identity matrix with the ith and jth columns swapped; this interchanges columns i and j of A; and finally

$$\begin{pmatrix} 1 & & & & & & \\ & \ddots & & & & & \\ & & 1 & & & & \\ & & & -1 & & & \\ & & & & 1 & & \\ & & & & & \ddots & \\ & & & & & & 1 \end{pmatrix}$$

which negates the ith column (note that unlike in a field, we cannot scale by arbitrary constants, since only ± 1 are units; other constants would yield an uninvertible determinant). Note that we have analogous row operations resulting from left-multiplication.

First, want to arrive at a matrix which has the form

$$\left(\begin{array}{c|ccc} d & 0 & \cdots & 0 \\ 0 & & & \\ \vdots & & M & \\ 0 & & & \end{array}\right)$$

where d divides every entry of M. We can carry out the Euclidean algorithm via our elementary row operations on any two row headers to make one of them the gcd of the two. Therefore, we can arrive at a matrix with $a_{1,1} \mid a_{i,1}$ for all i (that is, $a_{1,1}$ divides all the row headers). We can then subtract a multiple of the first row from every other row to arrive at a matrix of the form

$$\begin{pmatrix} a_{1,1} & a_{1,2} & \cdots & a_{1,n} \\ 0 & & & \\ \vdots & & \ddots & \\ 0 & & & \end{pmatrix}$$

If $a_{1,1}$ does not divide all column headers, we can make $a_{1,1}$ the gcd of all the column headers using column operations. Of course, this undoes all the work we have performed on the row headers, and after zeroing out the column headers, we must repeat the (very inefficient) algorithm from the beginning. Since $a_{1,1}$ either gets strictly smaller or divides all headers at the beginning of each repetition of the algorithm, the algorithm terminates.

This is almost what we want; we still want the submatrix M to contain only entries which are multiples of $a_{1,1}$. Suppose M has an entry $a_{i,j}$ which $a_{1,1}$ does not divide. We can use row or column operations to bring $a_{i,j}$ to the first column or row (recall that all the headers are 0). By repeating our algorithm, since $a_{1,1}$ must get smaller each time (or becomes equal to 1), this process eventually stops with our desired matrix.

Recursing on the submatrix M completes our proof. ■

We have shown that for $\mathbb{Z}$-modules V and W and any homomorphism $\varphi : V \to W$, we can represent φ, for some choices of bases, by

$$A = \left(\begin{array}{ccc|c} d_1 & & 0 & \\ & d_2 & & 0 \\ & & \ddots & \\ 0 & & d_k & \\ \hline & 0 & & 0 \end{array}\right)$$

Note that A takes $e_i \mapsto d_1 e_1$ for $i \le k$ and $e_i \mapsto 0$ for $i > k$. Thus, we have

$$\ker(\varphi) \cong R^{n-k} \qquad \text{and} \qquad \operatorname{im}(\varphi) \cong R^k$$

Specifically, the image of φ is spanned by vectors $d_i w_i'$ for $i \le k$ (where $\{w_i'\} \in W$ is the changed basis) Combined, these observations give us:

Corollary 29.2. *If $\varphi : V \to W$ is a homomorphism of free $\mathbb{Z}$-modules, then $\ker(\varphi)$ and $\operatorname{im}(\varphi)$ are free. The cokernel is*

$$\operatorname{coker}(\varphi) = \mathbb{Z}/d_1 \oplus \cdots \oplus \mathbb{Z}/d_k \oplus \mathbb{Z}^{m-k}$$

Lemma 29.3. *If $W \subset \mathbb{Z}^m$ is any submodule, then W is free.*

Proof. We will assume for now (and prove later) that W is finitely generated. So suppose that $w_1, \ldots, w_n$ are generators of $W \subset \mathbb{Z}^m$. This gives us a map $\mathbb{Z}^n \to \mathbb{Z}^m$ which sends the ith basis vector to a generator w_i. The image of this map is W, and hence W is free as desired. ■

Corollary 29.4. *Any finitely-generated abelian group G is of the form*

$$G = \mathbb{Z}^a \oplus \mathbb{Z}/d_1 \oplus \cdots \oplus \mathbb{Z}/d_k \qquad d_i \mid d_{i+1}, \forall i$$

Proof. Equivalently, we know that G is a finitely-generated $\mathbb{Z}$-module. Choosing a set of generators $v_1, \ldots, v_m$, we get a map $\mathbb{Z}^m \twoheadrightarrow G$ sending the $e_i \mapsto v_i$. The kernel, being free, is isomorphic to some $\mathbb{Z}^n$. Hence, G is the cokernel of the inclusion $\mathbb{Z}^n \to \mathbb{Z}^m$, which yields our desired result. ■

Lecture 30 — 4/9/12

We return briefly to a claim we made in a previous lecture, that over an arbitrary ring R, matrices $A, B \in M_n R$ satisfy

$$\det(A) \cdot \det(B) = \det(AB)$$

Let us begin by replacing these matrices with variable matrices $X = (x_{ij})$ and $Y = (y_{kl})$. What we wish to do is prove that the *identity*

$$\det(X) \cdot \det(Y) = \det(XY)$$

holds; we can then substitute elements of any ring R for the entries of X and Y.

We define this substitution relative to $\mathbb{Z}[\{x_{ij}\}, \{y_{kl}\}]$, the polynomial ring over the integers in $2n^2$ variables. There is a unique homomorphism $\mathbb{Z} \to R$ for any ring R, which yields, given some matrices $A, B \in M_n R$, the substitution homomorphism

$$\mathbb{Z}[\{x_{ij}\}, \{y_{kl}\}] \longrightarrow R$$

sending $x_{ij} \mapsto a_{ij}$ and $y_{kl} \mapsto b_{kl}$.

But since the determinant function respects ring homomorphism, if our identity holds in $\mathbb{Z}$, it holds in any ring R. Now consider the inclusion

$$\mathbb{Z}[\{x_{ij}\}, \{y_{kl}\}] \subset \mathbb{C}[\{x_{ij}\}, \{y_{kl}\}]$$

and the polynomial

$$f(x_{ij}, y_{ij}) = \det(X) \cdot \det(Y) - \det(XY)$$

We already know that, in $\mathbb{C}$, the polynomial function f satisfies

$$f(x_{ij}, y_{ij}) = 0, \qquad \forall x_{ij}, y_{ij} \in \mathbb{C}$$

But then $f = 0$ as a polynomial over $\mathbb{C}$, and hence also as a polynomial over $\mathbb{Z}$.

Observation 30.1. Recall now our corollary, that any finitely-generated abelian group G is of the form

$$G = \mathbb{Z}^a \oplus \mathbb{Z}/d_1 \oplus \cdots \oplus \mathbb{Z}/d_k \qquad d_i \mid d_{i+1}, \forall i$$

It is also true for this decomposition that the a and d_i are all unique. Note that

$$a = \dim_{\mathbb{Q}}(G \otimes_{\mathbb{Z}} \mathbb{Q})$$

Since, if $\alpha \in G$, we have $d_i \alpha = 0$, and so

$$\alpha \otimes 1 = d_i \alpha \otimes \frac{1}{d_i} = 0$$

Observation 30.2. Note also that our proof of the diagonalization of matrices in $\mathbb{Z}$ applies equally well to any Euclidean domain. That is, for any Euclidean domain R, any matrix in $M_n R$ can be diagonalized as

$$\left(\begin{array}{ccc|c} \begin{pmatrix} d_1 & & \\ & \ddots & \\ & & d_k \end{pmatrix} & & & \\ \hline & & & 0 \end{array}\right)$$

In non-Euclidean domains, however, this does not hold. In the case of $R = F[x, y]$, for example, R contains the ideal $I = (x, y)$, which is not a free module. Then the map $\varphi : R \twoheadrightarrow R/(x, y)$ cannot be diagonalized as a matrix, for then $I = \ker(\varphi)$ would be free.

Definition 30.3. Let R be a ring, $\varphi : R^n \to R^m$ a homomorphism of finitely-generated free R-modules. Given some choice of basis for R^n and R^m, we can represent φ by a matrix

$$A = \begin{pmatrix} a_{1,1} & \cdots & a_{1,n} \\ \vdots & \ddots & \vdots \\ a_{m,1} & \cdots & a_{m,n} \end{pmatrix}$$

Consider the module

$$M = \operatorname{coker}(\varphi) = R^m / \operatorname{im}(A)$$

with quotient map

$$\begin{aligned} \psi : R^m &\longrightarrow M \\ e_i &\longmapsto v_i \end{aligned}$$

for the corresponding residues $v_i \in M$. Clearly, M is generated by the v_i; moreover, each column of A gives a relation

$$a_{1,j} v_1 + a_{m,j} v_m = 0$$

Conversely, every linear relation among the v_i is a linear combination of these. So M is fully determined by the v_i along with these relations; we call A the presentation of M.

To obtain a presentation for our module M, we explicitly defined it as the cokernel of a map between free modules. Suppose that we instead took any arbitrary R-module M—what conditions must we impose to obtain a presentation?

Obviously, we must restrict our attention to finitely-generated modules. In this case, we can choose generators $v_1, \ldots, v_m$ for M; this gives us a map $\psi : R^m \twoheadrightarrow M$. Then the relations on M are precisely those elements of the kernel $\ker(\psi)$. However, if we want to obtain a presentation, this kernel must be finitely generated as well.

If it is, we can choose generators $u_1, \ldots, u_n$ for $\ker(\psi)$; this gives us a map $\varphi : R^n \to R^m$ with $\operatorname{coker}(\varphi) = M$. Note that the $\ker(\psi)$ itself need not be free, but it must be finitely generated for us to determine R^n.

How, then, can we tell if $\ker(\psi)$ is finitely generated? This is a hard question to answer, since the map ψ already depends on the choice of generators for M. Instead, we impose further restrictions on the base ring R.

Definition 30.4. A ring R is Noetherian if any ideal $I \subset R$ is finitely-generated.

Proposition 30.5. *Let R be a Noetherian ring, M a finitely-generated R module. Then any submodule of M is again finitely-generated.*

Lecture 31 — 4/11/12

Theorem 31.1. *Let R be a ring, M a module over R. Then the following are equivalent:*

1. *Every submodule of M is finitely-generated.*
2. *Any infinite sequence*

$$M_1 \subset M_2 \subset \cdots \subset M$$

of submodules of M eventually stabilizes; that is, $\exists n_0 : \forall n, m \geq n_0$, $M_n = M_m$. This second condition is called the ascending chain condition.

Definition 31.2. A Noetherian module is an R-module satisfying the above condition.

Proof. The reverse direction is easy. Suppose $N \subset M$ is not finitely-generated. Then we claim that we can construct an infinite, non-stabilizing sequence

$$M_1 \subsetneq M_2 \subsetneq M_3 \subsetneq \cdots$$

Start with any $v_1 \in N$, and take $M_1 = \langle v_1 \rangle$. Then choose $v_2 \in N - M_1$, and take $M_2 = \langle v_2 \rangle$. Because N is not finitely-generated, we can always find such an element v_i; we achieve our desired sequence by induction.

Suppose now that we have an infinite sequence $M_1 \subset M_2 \subset \cdots \subset M$. Set

$$N = \bigcup M_i \subset M$$

This is a submodule of M, and hence is finitely-generated. Choose generators $v_1, \ldots, v_k$ for M. For each i, there must exist some index n_i such that $\forall n \geq n_i$, $M_n \ni v_i$. Then there is some module M_{n_0} in the chain (with $n_0 = \max\{n_i\}$) such that every $M_i \supset M_{n_0}$ in the chain contains all the v_i. But then we have $M_i = N$ for every $i > n_0$. Thus, the chain stabilizies, which completes our proof. ∎

Theorem 31.3. *If R is a Noetherian ring, then every finitely-generated R-module M is Noetherian.*

Proof. We will start with the case $M = R^n$ for some n. Our base case for $n = 1$ is true by assumption; we proceed by induction. Let $N \subset M$ be any submodule; we want to show that N is finitely-generated. Consider the projection map

$$\varphi : R^n \longrightarrow R^{n-1}$$
$$(x_1, \ldots, x_{n-1}, x_n) \longmapsto (x_1, \ldots, x_{n-1})$$

By our induction hypothesis, the image module $\varphi(N) \subset R^{n-1}$ is finitely-generated. Choose generators $\overline{v}_1, \ldots, \overline{v}_k$, and choose also representatives $v_i \in N$ such that $\varphi(v_i) = \overline{v}_i$. Then, $\forall v \in N$, we can write

$$\varphi(v) = \sum c_i \overline{v}_i, \qquad c_i \in R$$

Consider the corresponding linear combination of the v_i; we can see that

$$v - \sum c_i v_i \in \ker(\varphi)$$

But $\ker(\varphi) \subset R$. Since R is Noetherian, $\ker(\varphi)$ is finitely-generated; choose generators $w_1, \ldots, w_l$. Then we can write any element N as a linear combination of the v_i and the w_j.

Now let M be any arbitrary finitely-generated R-module, not necessarily free. Let $\varphi : R^n \twoheadrightarrow M$ be the canonical map taking e_i to some n generators of M, and let $N \subset M$ be any submodule. By the above, $\varphi^{-1}(N) \subset R^n$ is finitely-generated; choose generators $v_1, \ldots, v_m$. Then N is generated by $\varphi(v_1), \ldots, \varphi(v_m)$, which completes our proof. ∎

Note that "finitely-generated" has a different meaning depending on whether we are speaking about rings or modules. For instance, the ring

$$\mathbb{Z}[\tfrac{1}{2}] = \left\{ \frac{a}{2n} : a \in \mathbb{Z}, n \in \mathbb{N} \right\} \subset \mathbb{Q}$$

is finitely-generated over $\mathbb{Z}$ as a ring, but not as a module. As a module, multiplication is only defined with respect to the base ring $\mathbb{Z}$, so no finite number of generators will yield all the inverse powers of 2.

Lemma 31.4. *Let R be a Noetherian ring. Then for any ideal $I \subset R$, the quotient R/I is Noetherian.*

Proof. Let $\varphi : R \to R/I$ be the quotient map. Let $\bar{J} \subset R/I$ be any ideal in R/I. Then $J = \varphi^{-1}(\bar{J})$ is an ideal in R. By assumption, J is finitely-generated. Choosing generators $v_1, \ldots, v_k$, we see that $\bar{J}$ is generated by $\varphi(v_1), \ldots, \varphi(v_k)$. ∎

Note that we can obtain presentations for any finitely-generated module M over a Noetherian ring R. That is, every R-module where R is Noetherian is the cokernel of a map $\varphi : R^n \to R^m$.

Theorem 31.5 (Hilbert Basis Theorem)**.** *Let R be a Noetherian ring. Then the ring $R[x_1, \ldots, x_n]$ is Noetherian.*

Proof. Since we obtain $R[x_1, \dots, x_n]$ by adjoining x_n to $R[x_1, \dots, x_{n-1}]$, it suffices to prove the theorem for $R[x]$. Let $I \subset R[x]$ be any ideal. For every $f \in R[x]$ of the form $f = a_n x^n + \cdots + a_0$, with $a_n \neq 0$, we define

$$\operatorname{lc}(f) = a_n$$

This is notation for the leading coefficient of f.

Now consider the set

$$A = \{\operatorname{lc}(f) : f \in I\} \cup \{0\}$$

We claim that A is an ideal in R. This is fairly easy to see. Let $\alpha, \beta \in A$, with $\alpha = \operatorname{lc}(f)$ and $\beta = \operatorname{lc}(g)$. Take any $c \in R$; if $c\alpha \neq 0$, then

$$c\alpha = \operatorname{lc}(cf)$$

and hence A absorbs multiplication in R. To show that A is closed under addition, suppose WLOG that $\deg f \geq \deg g$. Then we also have $\beta = \operatorname{lc}(x^{\deg f - \deg g} \cdot g)$, and hence

$$\alpha + \beta = \operatorname{lc}(f + x^{\deg f - \deg g} \cdot g)$$

(unless, of course, $\alpha + \beta = 0$).

Since $A \subset R$, it is finitely-generated. So let us choose generators $\alpha_1, \dots, \alpha_k$ for A along with polynomials $f_1, \dots, f_k \in I$ with $\operatorname{lc}(f_i) = \alpha_i$. Take $n = \max\{\deg f_i\}$. WLOG, we can assume the f_i all have the same degree by multiplying each f_i by $x^{n - \deg f_i}$.

Now set

$$P_0 = \{f \in R[x] : \deg f \leq n\}$$

This is a free module over R, isomorphic to R^{n+1} (it is, however, clearly not a ring). Let also

$$P = P_0 \cap I = \{f \in I : \deg f \leq n\}$$

This is a submodule of P_0; hence, P is finitely-generated as an R-module. Let us choose generators $g_1, \dots, g_l$ for P. We claim now that $\{f_i\} \cup \{g_j\}$ together generate I.

We prove this by induction on degree. Let $f \in I$. If $\deg f \leq n$, then $f \in P$ and we are done; this is our base case. Suppose $\deg f > n$. We can write $\operatorname{lc}(f)$ as a linear combination

$$\operatorname{lc}(f) = \sum c_i \alpha_i$$

which corresponds to a linear combination of the f_i. Then, if we take

$$g = f - \sum c_i (x^{\deg f - n} f_i)$$

since $\deg g < \deg f$, by the induction hypothesis, g is expressible as a linear combination of the f_i and g_j, and therefore so is f. ∎

Definition 31.6. We say that a ring R is finitely-generated over a field K if $K \hookrightarrow R$ and if $\exists v_1, \dots, v_k \in R$ such that

$$K[x_1, \dots, x_k] \twoheadrightarrow R$$

by the evaluation homomorphism.

Corollary 31.7. *Any finitely-generated ring over a field or $\mathbb{Z}$ is Noetherian.*

Lecture 32 — 4/13/12

Definition 32.1. A sequence of R-module homomorphisms

$$M_n \xrightarrow{\varphi_n} M_{n-1} \xrightarrow{\varphi_{n-1}} M_{n-1} \longrightarrow \cdots \longrightarrow M_1 \xrightarrow{\varphi_1} M_0$$

is called exact if

$$\forall k < n, \quad \ker \varphi_{k-1} = \operatorname{im} \varphi_k$$

We say a sequence is a complex if

$$\forall k < n, \quad \varphi_{k-1} \circ \varphi_k = 0$$

that is, if $\operatorname{im} \varphi_i \subset \ker \varphi_{i-1}$. In this case, the quotients $(\ker \varphi_{k-1})/(\operatorname{im} \varphi_k)$ are called the cohomology modules.

Definition 32.2. A short exact sequence is an exact sequence of the form

$$0 \longrightarrow M \longrightarrow N \longrightarrow P \longrightarrow 0$$

As a result $M \hookrightarrow N$ is an injection, and $P = N/M$.

Example. Let $R = \mathbb{Z}$. Let M be a finitely-generated $\mathbb{Z}$-module; hence, we get a map $\psi : \mathbb{Z}^m \to M$ onto M's generators. We know that $\ker(\psi) \cong \mathbb{Z}^n$ for some n is free. This yields a short exact sequence

$$0 \longrightarrow \mathbb{Z}^n \longrightarrow \mathbb{Z}^m \longrightarrow M \longrightarrow 0$$

Similarly, take $R = F[t]$. Any finitely-generated module R-module M is the cokernel of some map

$$0 \longrightarrow R^n \xrightarrow{\varphi} R^m \longrightarrow M \longrightarrow 0$$

Since $F[t]$ is a Euclidean domain, we can obtain a diagonalized presentation

$$M \cong R^a \oplus R/(d_1) \oplus R/(d_2) \oplus \cdots \oplus R/(d_k)$$

where $d_i \in F[t]$ and $\forall i$, $d_i \mid d_{i+1}$.

Example. Let $R = F[x, y]$ for some field F and let $M = R/(x, y) \cong F$. Consider the quotient map

$$\begin{aligned} \pi : F[x, y] &\longrightarrow F \\ f &\longmapsto f(0, 0) \end{aligned}$$

which has $\ker(\pi) = (x, y)$. (x, y) is generated by x and y, so we have a natural map $\varphi : R^2 \to (x, y)$ mapping the standard basis onto the generators. Thus, we have an exact sequence

$$R^2 \xrightarrow{\varphi} (x, y) \hookrightarrow R \xrightarrow{\pi} F$$

The kernel of φ is given by

$$\ker(\varphi) = \{(f, g) : xf + yg = 0\}$$

Since $xf = -yg$, $y \mid f$ and likewise $x \mid g$. We can write $f = yf'$ and $g = xg'$. But then substituting, we get that $xyf' = -xyg'$. If we set $h = g'$, then we have $f = -yh$ and $g = xh$, which means that

$$\ker(\varphi) = \{(-yh, xh)\}$$

This gives a natural map $\psi : R \to R^2$ given by the matrix

$$\psi = \begin{pmatrix} -y \\ x \end{pmatrix}$$

It is easy to see that $\ker(\psi) = 0$, and hence we have an exact sequence

$$0 \to R \xrightarrow{\psi} R^2 \xrightarrow{\varphi} R \xrightarrow{\pi} F \to 0$$

Now let $R = F[x_1, \ldots, x_r]$, M a finitely-generated R-module. We know that there exists a surjection $\varphi_0 : M_0 \to M$ where $M_0 \cong R^{m_0}$. Since M_0 is Noetherian, $\ker(\varphi_0)$ is finitely-generated. Choosing a set of m_1 generators of this kernel, we get

$$M_1 \xrightarrow{\varphi_1} M_0 \xrightarrow{\varphi_0} M \longrightarrow 0$$

where $M_1 \cong R^{m_1}$.

Theorem 32.3 (Hilbert Syzygy Theorem)**.** *The process described above terminates after at most r steps. In general, this is called a free resolution of M.*

Lecture 33 — 4/16/12

In general, if M is an R-module over a Noetherian ring R, it can be realized as the cokernel of a homomorphism of free modules. Moreover, M has a free resolution.

Example. Let $R = F[t]$ for F a field. Consider the matrix

$$A = \begin{pmatrix} t^2 - 3t + 1 & t - 2 \\ (t - 1)^3 & t^2 - 3t + 2 \end{pmatrix}$$

The R-module given by $M = \operatorname{coker}(A)$, that is, presented by A, is generated by two elements v, w with

$$\begin{aligned} (t^2 - 3t + 1)v + (t - 1)^3 w &= 0 \\ (t - 2)v + (t^2 - 3t + 2)w &= 0 \end{aligned}$$

Unfortunately, this doesn't give us much insight into the module's structure.

However, we can diagonalize A as

$$A = \begin{pmatrix} 1 & 0 \\ 0 & t^3 - 3t^2 + 2t \end{pmatrix}$$

from which we have

$$M = \operatorname{coker}(A) = F[t]/(t^3 - 3t^2 + 2t)$$

We can factor $t^3 - 3t^2 + 2t = t(t - 1)(t - 2)$, which means we can further decompose

$$M = F[t]/(t) \oplus F[t]/(t - 1) \oplus F[t]/(t - 2)$$

Each of these direct summands is congruent to F.

Observation 33.1. Note that, in general for modules over a Euclidean domain, we can arrange for the d_i in the diagonalization to be prime powers rather than dividing one another in sequence. If $R = F[t]$, can require $d_i = f_i^{a_i}$ where the $f_i \in F[t]$ are irreducible. In the special case $F = \mathbb{C}$, can take each $d_i = (t - c_i)^{a_i}$

We will now turn briefly to the study of modules in relation to vector spaces and linear operators. Let V be an n-dimensional vector space over F, $T : V \to V$ a linear map. We can give the group V the structure of an $F[t]$-module by defining

$$t \cdot v = Tv$$

Then we have $f(t) \cdot v = [f(T)]v$. Since V is finite-dimensional, it is finitely-generated over $F[t]$.

In the case $F = \mathbb{C}$, we can decompose

$$V \cong \mathbb{C}[t]/(t - c_1)^{a_1} \oplus \cdots \oplus \mathbb{C}[t]/(t - c_n)^{a_n}$$

Each of the direct summands V_i is a submodule of V over $F[t]$ as well as a vector space invariant under T.

Definition 33.2. A module which is generated by one element is called cyclic.

In the case of $F[t]$, modules of the form $F[t]/(f)$ are all cyclic.

Consider a single direct summand $V := \mathbb{C}[t]/(t - c)^a$, which is also an a-dimensional vector space over $\mathbb{C}$. We can choose a basis $1, t, t^2, \ldots, t^{a-1}$ for V. The linear map T is defined by left-multiplication by t. For each basis vector v_i, we simply have $tv_i = v_{i+1}$ except for $v_{a-1} = t^{a-1}$, where we have

$$t^a = -\sum_{i=0}^{a-1} \binom{a}{i} c^{a-i} t^i$$

Definition 33.3. With respect to the choice of basis $\{1, t, t^2, \dots, t^{a-1}\}$, the matrix for T is given in rational canonical form,

$$T = \begin{pmatrix} 0 & & & & & -a_0 \\ 1 & 0 & & & & -a_1 \\ & 1 & \ddots & & & -a_2 \\ & & \ddots & 0 & & \vdots \\ & & & 1 & 0 & -a_{n-2} \\ & & & & 1 & -a_{n-1} \end{pmatrix}$$

Alternatively, we can choose our basis $\{v_i\}$ to be $1, t-c, (t-c)^2, \dots, (t-c)^{a-1}$. Multiplication by t, or equivalently, application of T, is given by

$$\begin{aligned} v_0 &\longmapsto v_1 + cv_0 \\ v_1 &\longmapsto v_2 + cv_1 \\ &\vdots \\ v_{a-2} &\longmapsto v_{a-1} + cv_{a-2} \\ v_{a-1} &\longmapsto cv_{a-1} \end{aligned}$$

Definition 33.4. With respect to the choice of basis $\{1, t-c, (t-c)^2, \dots, (t-c)^{a-1}\}$, the matrix for T is given in Jordan normal form,

$$T = \begin{pmatrix} c & & & & \\ 1 & c & & & \\ & 1 & c & & \\ & & \ddots & \ddots & \\ & & & 1 & c \end{pmatrix}$$

We are able to achieve the rational canonical form for arbitrary fields F, in which case the coefficients in the rightmost column will be in F. However, we *cannot* achieve Jordan normal form for vector spaces over arbitrary F. Note also that $R = \mathbb{Z}[t]$ does not exhibit the same behavior as $R = F[t]$, so this discussion does not hold.

Lecture 34 — 4/18/12

Theorem 34.1 (Bezout's Theorem)**.** *Let $f, g \in \mathbb{C}[x, y]$ be relatively prime, and say $m = \deg(f)$, $n = \deg(g)$. Define*

$$\Gamma = \{(x, y) : g(x, y) = f(x, y) = 0\}$$

Then $\#\Gamma \le mn$.

Claim 34.2. *Let $f, g \in \mathbb{C}[t]$ of degrees m and n respectively, with*

$$\begin{aligned} f(t) &= a_m t^m + \cdots + a_0 \\ g(t) &= b_n t^n + \cdots + b_0 \end{aligned}$$

Then there exists a polynomial

$$P(a_0, \dots, a_m, b_0, \dots, b_n)$$

such that $P(a, b) = 0$ iff f and g have a common zero.

Proof. Note that f and g have a common zero iff

$$(f, g) = \{af + bg : a, b \in \mathbb{C}[t]\} \subsetneq \mathbb{C}[t]$$

Let

$$S_k = \{f \in \mathbb{C}[t] : \deg(f) \le k\} \cong \mathbb{C}^{k+1}$$

Consider the map

$$\begin{aligned} \varphi : S_{n-1} \times S_{m-1} &\longrightarrow S_{m+n-1} \\ (a, b) &\longmapsto af + bg \end{aligned}$$

If f and g have a common zero, then φ is not surjective. Suppose they do not have common zeros. Then if $af + bg = 0$, every zero of f must also be a zero of b, and likewise for g and a. Then $a = b = 0$, meaning that φ is injective and hence an isomorphism.

Now choose a basis $1, t, t^2, \dots, t^{m+n-1}$ for S_{m+n-1}, $(1, 0), (t, 0), \dots, (t^{m+n-1}, 0), (0, 1), (0, t), \dots, (0, t^{m+n-1})$ for $S_{n-1} \times S_{m-1}$. We can define

$$P(a_0, \dots, a_m, b_0, \dots, b_n) = \det(A)$$

where A is the matrix representing φ of the form

$$A = \begin{pmatrix} a_0 & & & 0 & b_0 & & 0 \\ a_1 & a_0 & & & \vdots & \ddots & \\ \vdots & a_1 & \ddots & & \vdots & & b_0 \\ a_m & \vdots & \ddots & a_0 & \vdots & & \vdots \\ & a_m & \ddots & a_1 & b_n & & \vdots \\ & & \ddots & \vdots & & \ddots & \vdots \\ 0 & & & a_m & 0 & & b_n \end{pmatrix}$$

Then $P(a, b) = 0$ iff φ is nonisomorphic, which is the case iff f and g have common zeros, as desired. ∎

Proof. (of Bezout's Theorem) We can write

$$\begin{aligned} f(x, y) &= a_m(x)y^m + \cdots + a_0(x) \\ g(x, y) &= b_n(x)y^n + \cdots + b_0(x) \end{aligned}$$

For how many values of x is it the case that f and g have common zeros? Consider our matrix A now where all the entries are viewed as polynomials in x. Then the set of x such that $f(x, y)$ and $g(x, y)$ have common zeros is just the zeros of $\det(A)$. Then since $\deg(a_i(x)) \le m - 1$ and $\deg(b_j(x)) \le n - 1$, we have $\deg(\det(A)) \le mn$. ∎

Lecture 35 — 4/20/12

In the final lecture, we will revisit field theory by discussing the separability of fields.

Definition 35.1. Let F be a field, $f \in F[x]$ with $\deg(f) = n$. We say that f has distinct roots if f has n distinct roots in its splitting field. Equivalently, $\forall F \hookrightarrow K$, $\alpha \in K$, $(x-\alpha)^2 \nmid f \in K[x]$.

Observation 35.2. Note that the property of having distinct roots is independent of the ground field F; that is, for any extension $F \hookrightarrow K$, if $f \in F[x]$ has distinct roots, then $f \in K[x]$ does as well. Also, note that if f has distinct roots and $g \mid f$, then g also has distinct roots.

Definition 35.3. We say a polynomial $f \in F[x]$ is separable over F if every irreducible factor of $f \in F[x]$ has distinct roots.

Note that this definition, unlike the previous definition, is dependent on the ground field F. It is still the case that if $f \in F[x]$ is separable, then $f \in K[x]$ is separable, since f's irreducibles can only decompose more in the extension field K/F. However, the converse is not true in general.

Example. Let $F = \mathbb{F}_p(t)$, the field of rational functions over $\mathbb{F}_p$. Let $f = x^p - t \in F[x]$. This is irreducible by the Eisenstein criterion; does it have distinct roots?

We claim it does not. Let K/F be any extension, $\alpha \in K$ a root of $f \in K[x]$. Then $\alpha^p = t$ in K. So in $K[x]$,

$$\begin{aligned} f(x) &= x^p - t \\ &= x^p - \alpha^p \\ \text{since } \mathbb{F}_p \qquad &= (x-\alpha)^p \end{aligned}$$

So f is not separable.

Theorem 35.4. *Let F be any field. If $\exists f \in F[x]$ inseparable, then $\operatorname{char}(F) = p > 0$ and $\#F = \infty$.*

Proof. Recall the derivative, defined for

$$f = a_n x^n + \cdots + a_0$$

as

$$f' = n a_n x^{n-1} + \cdots + a_1$$

We claim first that

Claim 35.5. *$f \in F[x]$ has distinct roots iff f and f' have no common roots in any extension K/F.*

Proof. Suppose that $f(\alpha) = f'(\alpha) = 0$ in $K[x]$, for $\alpha \in K$. Then $f(x) = (x-\alpha)g(x)$ for some $g \in K[x]$. Then

$$f'(x) = g(x) + (x-\alpha)g'(x)$$

by the product rule. Then if $f'(\alpha) = 0$, we must have $g(\alpha) = 0$, so $(x-\alpha) \mid g$. Then $(x-\alpha)^2 \mid f$, a contradiction.

Claim 35.6. *If $f \in F[x]$ is irreducible and non-constant, then f fails to have distinct roots iff $f' = 0 \in F[x]$.*

Proof. To say that f does not have distinct roots means $\exists K/F$, $\alpha \in K : (x-\alpha)^2 \mid f$. Then $(x-\alpha) \mid f'$. Since f and f' have a common factor in $K[x]$, but f is irreducible in $F[x]$, then $f \mid f' \in F[x]$. So we must have $f' = 0 \in F[x]$.

Definition 35.7. Let F be a field, $\operatorname{char}(F) = p > 0$. We can define a field homomorphism given by

$$\begin{aligned} \varphi : F &\longrightarrow F \\ \alpha &\longmapsto \alpha^p \end{aligned}$$

since $(\alpha+\beta)^p = \alpha^p + \beta^p$. We say that F is perfect if φ is an isomorphism (i.e., is surjective). We say that all fields of characteristic zero are perfect.

In particular, note that if F is finite, then F is perfect. Next, we claim that

Claim 35.8. *If F is perfect, then every polynomial $f \in F[x]$ is separable.*

Proof. Suppose that $\exists f \in F[x]$ inseparable. Then $\exists f \in F[x]$ irreducible but *without* distinct roots. Suppose first that $\operatorname{char}(f) = 0$. Then all nonconstant polynomials have nonzero derivatives. By the previous claim, f must have distinct roots.

Now suppose $\operatorname{char}(f) = p > 0$. Since F is perfect, we can write

$$f(x) = g(x^p)$$

for some $g = a_n x^n + \cdots + a_0 \in F[x]$. We can also choose $b_i \in F$ such that $b_i^p = a_i$. Then

$$\begin{aligned} f(x) &= g(x^p) \\ &= \sum a_i x^{pi} \\ &= \sum (b_i x^i)^p \\ &= (\sum b_i x^i)^p \end{aligned}$$

so f is not irreducible, a contradiction.

This claim, along with the definition of perfect field, yields our desired result. ∎

Definition 35.9. Let $F \hookrightarrow K$ be a field extension. $\alpha \in K$ is separable over F if its minimal irreducible polynomial $f \in F[x]$ is separable. We say that K/F is separable if every element $\alpha \in K$ is separable.

Definition 35.10. A field extension K/F is normal if $\forall f \in F[x]$ irreducible, f has a root in K iff f splits completely in K.

Theorem 35.11 (Fundamental Theorem of Galois Theory)**.** *K/F is Galois iff K/F is normal and separable.*