18.702: Algebra II

Sanjana Das and Jakin Ng

Spring 2022

Contents

1 Representations

1.1 Introduction

The lecturer is Roman Bezrukavnikov. These notes are taken by **Sanjana Das** and **Jakin Ng**, and the note-taking is supervised by Ashay Athalye. Here is some basic information about the class:

- The text used in this class will be the 2nd edition of **Algebra**, by Artin.
- The course website is found on **Canvas**, and the problem sets will be submitted on **Gradescope**.
- The problem sets will be due Wednesday at 11:59PM.

During this class, we will cover three main topics.

1. Representation theory: In 18.701, we studied group actions, which let us think of groups as a set of symmetries of the set being acted on. Here we'll study how groups can act by symmetry on a *vector space* — this combines the fundamental concepts of of symmetry and linearity.
2. Ring theory: We'll learn to add *and* multiply in abstract settings.
3. Galois theory: We'll study the symmetries of solutions to polynomial equations. For example, a quadratic equation $ax^2 + bx + c = 0$ has two solutions,

$$x_{\pm} = \frac{-b \pm \sqrt{b^2 - 4ac}}{2}.$$

There's an ambiguity in the sign of the square root, and this gives a symmetry between the two roots (by swapping the sign). For higher-degree polynomial equations, we'll see that symmetries (such as changing $+$ to $-$ in the above formula) control the existence of formulas such as the quadratic formula, and the shape of such formulas when they do exist.

1.2 What is a Representation?

In representation theory, we think of a group as the symmetries of a *vector space* — this perspective lets us study the group using tools from linear algebra.

Guiding Question
How can we represent elements of groups as symmetries or linear operations on vector spaces?

The following definition formalizes this idea, by representing elements of groups as *matrices*:

Definition 1.1
Let G be a group. A **complex, n-dimensional matrix representation** of G is a homomorphism

$$R : G \to \mathrm{GL}_n(\mathbb{C})^a.$$

[a] Recall that $\mathrm{GL}_n(\mathbb{C})$ denotes the group of invertible $n \times n$ matrices with entries in $\mathbb{C}$.

Similarly, a real representation is a homomorphism

$$R : G \to \mathrm{GL}_n(\mathbb{R}).$$

Representations can be defined over any field (for instance, we could even define representations over the finite field $\mathbb{F}_p$), but in this class, we will mostly work with only *complex* representations of *finite* groups.

Earlier, we mentioned that in representation theory, we want to think of elements of a group as symmetries of a vector space. To see why the above definition achieves this, note that invertible matrices play a special role for the vector space $\mathbb{C}^n$: they act on the column vectors. More explicitly, any matrix $A \in \mathrm{Mat}_{n\times n}(\mathbb{C})$ defines a linear transformation from $\mathbb{C}^n$ to itself, by taking $v \mapsto Av$. Moreover, $\mathrm{GL}_n(\mathbb{C})$ consists of invertible $n \times n$ matrices, which are exactly the matrices for which $v \mapsto Av$ is a *bijective* linear transformation. So $\mathrm{GL}_n(\mathbb{C})$ equivalent to a group of linear automorphisms* on $\mathbb{C}^n$, which can be thought of as the symmetries of the vector space.

*Isomorphisms from $\mathbb{C}^n$ to itself

This means that writing down a homomorphism $R : G \to \mathrm{GL}_n(\mathbb{C})$ is equivalent to writing down a *linear group action* of G on $\mathbb{C}^n$, where each $g \in G$ acts by taking $v \mapsto R_g(v)$. More explicitly, to write down a representation R by thinking in terms of group actions, for each $g \in G$ we need to define an operator R_g on $\mathbb{C}^n$ such that it is:

- **Group Action.** The map

$$\begin{aligned} G \times V &\to V \\ (g, v) &\mapsto R_g(v) \end{aligned}$$

  satisfies the axioms of a group action of G on V:

  1. Since R is a homomorphism, $R_{gh} = R_g R_h$, so $R_{gh}(v) = R_g(R_h(v))$ for all $g, h \in G$ and $v \in V$.
  2. Again, because R is a homomorphism, $R_{1_G} = \mathrm{Id}$, and so $R_{1_G}(v) = v$ for all $v \in V$.

- **Linear.** For each $g \in G$, the map $v \mapsto R_g(v)$ is linear. We have

$$R_g(v + w) = R_g(v) + R_g(w)$$

  and

$$R_g(\lambda v) = \lambda R_g(v)$$

  for all $v, w \in V$ and $\lambda \in \mathbb{C}$. From definition 1.1, these properties correspond to the fact that elements of $\mathrm{GL}_n(\mathbb{C})$ are linear operators on $\mathbb{C}^n$.

Note 1.2
It is a notational convention to write R_g instead of $R(g)$. Using the notation of group actions, $R_g(v)$ can also be written as gv for a vector $v \in \mathbb{C}^n$.

1.3 Examples of Representations

The simplest example of a representation is the *trivial representation*.

Example 1.3 (Trivial Representation)
The **trivial representation** of any group G is the one-dimensional representation where $R_g = 1$ for all g. That is, every element of the group maps to the 1×1 identity matrix $\begin{bmatrix} 1 \end{bmatrix} \in \mathrm{GL}_1(\mathbb{C}) \cong \mathbb{C}^\times$.

The trivial representation is clearly a homomorphism[†], and therefore a valid representation. Every group has a trivial representation, and it will turn out to be a "building block" for more complicated representations.

The symmetric group S_n has another one-dimensional representation:

Example 1.4 (Sign Representation)
A symmetric group S_n has a one-dimensional representation, called the **sign representation**, where

$$R_\sigma = \mathrm{sgn}(\sigma)^{a} = \begin{cases} 1 & \text{if } \sigma \text{ is even} \\ -1 & \text{if } \sigma \text{ is odd.} \end{cases}$$

Again, the target space is $\mathrm{GL}_1(\mathbb{C}) \cong \mathbb{C}^\times$.

[a] Recall that every $\sigma \in S_n$ can be written as a product of transpositions $\tau_1 \tau_2 \cdots \tau_k$; the parity of σ is the parity of k, and $\mathrm{sgn}(\sigma) = (-1)^k$.

As an example of a representation that is not one-dimensional, S_n also has another representation with dimension n, the permutation representation.

[†] Try writing a short proof of this!

Example 1.5 (Permutation Representation)
The group S_n has a n-dimensional representation, called the **permutation representation**, which takes each element $\sigma \in S_n$ to its corresponding permutation matrix — the $n \times n$ matrix which sends the ith basis vector $\vec{e}_i$ to the $\sigma(i)$th basis vector $\vec{e}_{\sigma(i)}$, meaning that its ith column is $\vec{e}_{\sigma(i)}$. As an example, when $n = 3$,

$$R_{(123)} = \begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}.$$

A representation is called **faithful** if it is injective, in which case G is isomorphic to its image under the representation. The permutation representation is faithful, while the trivial and sign representations are not faithful (except for very small n[‡]).

A familiar example of a representation is $\mathbb{Z}/m\mathbb{Z}$[§] acting as a group of rotational symmetries.

Example 1.6
The group $\mathbb{Z}/m\mathbb{Z}$ corresponds to the rotational symmetries of a regular m-gon. By placing the polygon in the two-dimensional plane and using the standard basis for the plane, we get a two-dimensional representation of $\mathbb{Z}/m\mathbb{Z}$ where every element $\overline{a} \in \mathbb{Z}/m\mathbb{Z}$ is mapped to the corresponding rotation matrix: more explicitly, the representation is

$$\overline{a} \mapsto \begin{bmatrix} \cos 2\pi a/m & -\sin 2\pi a/m \\ \sin 2\pi a/m & \cos 2\pi a/m \end{bmatrix}.$$

This can be thought of either as a real representation, where the target space is $GL_2(\mathbb{R})$ or a complex one, where the codomain is $GL_2(\mathbb{C})$.

When describing a representation of a given group G, it may be somewhat time-consuming to list the images of *all* elements of G. But it's possible to describe a representation much more efficiently — if G is given by generators and relations, then to define a matrix representation R, it's enough to specify the images of the generators and check that they satisfy the relations. More explicitly, suppose

$$G = \langle x_1, \ldots, x_k \mid r_1, \ldots, r_m \rangle,$$

where the x_i are the generators, and the r_i are the relations. It is enough to specify

$$R_{x_1}, R_{x_2}, \cdots R_{x_k}$$

in order to define a unique representation. It's clear that R_{x_1}, …, R_{x_k} must satisfy the same relations as x_1, …, x_k. Conversely, given any n matrices γ_1, …, γ_k in $\mathrm{GL}_n(\mathbb{C})$ which satisfy these relations, we can set $R_{x_i} = \gamma_i$ for all i; then this determines the entire representation, since we can obtain R_g for *any* $g \in G$ simply by multiplying the γ_i in the same way that we would multiply the x_i to obtain g (since R is a homomorphism). The γ_i may also satisfy additional relations other than the $r_{j=1,\cdots m}$; if they do not, the representation will be faithful.

Example 1.7
The group $\mathbb{Z}/m\mathbb{Z}$ can be written as $\langle x \mid x^m = 1 \rangle$ (this denotes that it's generated by one element x, with the relation $x^m = 1$). So to define the above two-dimensional representation in Example 1.6, it's enough to specify that

$$\overline{1} \mapsto A = \begin{bmatrix} \cos 2\pi/m & -\sin 2\pi/m \\ \sin 2\pi/m & \cos 2\pi/m \end{bmatrix},$$

and to verify that $A^m = 1$.

Another familiar example is the following representation of the dihedral group, which is the group of *all* symmetries of a regular m-gon (meaning rotations and reflections).

[‡]For $n = 1$, the trivial representation is faithful, since there is only one element, and for $n = 1, 2$, the sign representation is faithful

[§]The group of integers mod m; we'll use $\overline{a}$ to denote the residue of a mod m

Example 1.8
We can write the dihedral group as $D_m = \langle r, s \mid r^m = 1, s^2 = 1, srs^{-1} = r^{-1} \rangle$. Then D_m has a two-dimensional representation given by

$$r \mapsto \begin{bmatrix} \cos 2\pi/m & -\sin 2\pi/m \\ \sin 2\pi/m & \cos 2\pi/m \end{bmatrix},$$
$$s \mapsto \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix},$$

since it can be verified that the images of r and s satisfy the same relations. Intuitively, this construction is quite similar to the representation of $\mathbb{Z}/m\mathbb{Z}$ in Example 1.6 — we can think of r as a rotation by $2\pi/m$ and s as a reflection, since these are the symmetries of the m-gon. Then to obtain this representation, we simply place the m-gon on the plane, and take the matrices corresponding to these transformations of the plane.

1.4 Linear Representations

Unfortunately, the current formulation of a representation requires a basis, as it is not possible to write down a matrix without choosing a basis. We are interested in the *story* of a journey, and not the particular *coordinates* of the journey.

Guiding Question
How can we think about representations in a coordinate-free way, without specifying a basis?

Given a matrix representation, a new **conjugate representation** can be obtained by choosing a different basis of $\mathbb{C}^n$ and rewriting the original matrices in the new basis.

Example 1.9
The representation of $\mathbb{Z}/3\mathbb{Z}$ described in Example 1.6 (as the rotations of a triangle) can be written in the standard basis as

$$\overline{1} \mapsto \begin{bmatrix} -1/2 & -\sqrt{3}/2 \\ \sqrt{3}/2 & 1/2 \end{bmatrix}.$$

But if we instead take the basis consisting of $v_1 = (1, 0)^t$ and $v_2 = (-1/2, \sqrt{3}/2)^t$, then this representation can be written as

$$\overline{1} \mapsto \begin{bmatrix} 0 & -1 \\ 1 & -1 \end{bmatrix}.$$

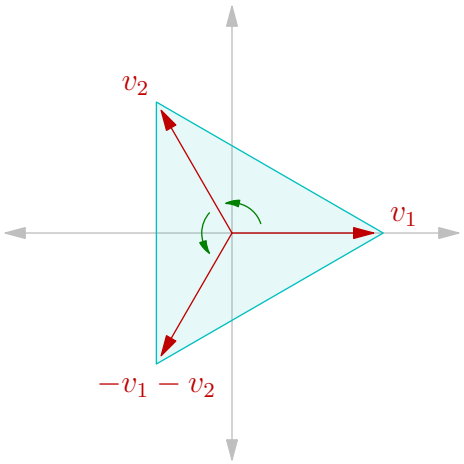

Clearly, the new representation is in fact technically a different matrix representation, but conceptually, it can be interpreted in the same way — it still describes the rotations of a triangle.

In general, suppose we have two bases of $\mathrm{GL}_n(\mathbb{C})$ with change of basis matrix P — in Example 1.9, the change

of basis matrix was

$$P = \begin{bmatrix} 1 & -1/2 \\ 0 & \sqrt{3}/2 \end{bmatrix}.$$

Then if we have a representation R with matrices written in the first basis, the conjugate representation obtained by writing R in the new basis is given by

$$R'_g = P^{-1} R_g P,$$

by simply applying the change of basis formula to each matrix.

Conjugate representations are essentially the same story, just presented in different coordinates, so we generally consider conjugate matrix representations as the same — we want to study the conjugacy classes of matrix representations, rather than the basis-dependent matrix representations themselves. In fact, we can eliminate the need for matrices altogether by using the concept of a *linear representation*, which does not require specifying a basis.

Definition 1.10
For a vector space V, let $\mathrm{GL}(V)$ be the group of linear automorphisms of V. A **linear representation** of V is a homomorphism

$$\rho : G \to \mathrm{GL}(V).$$

Note that a linear representation doesn't depend on coordinates or column vectors! In a matrix representation, we thought of group elements as acting on a vector space (specifically $\mathbb{C}^n$) by linear automorphisms, and we wrote down these automorphisms in a given basis by using matrices. But here instead of writing down the matrices corresponding to the linear automorphisms, we work with the automorphisms themselves.

We *can* turn a linear representation into a matrix representation — if we fix a basis of V, then we get an isomorphism between $\mathrm{GL}_n(\mathbb{C})$ and $\mathrm{GL}(V)$ (where $n = \dim V$), and for each linear automorphism ρ_g, we can write down the matrix corresponding to ρ_g in that basis. Choosing a different basis of V would give us a conjugate representation; thus specifying a linear representation provides a conjugacy class of matrix representations. Choosing a suitable basis and working with matrix representations can be helpful in computations, but we generally want to work with properties that are basis-independent and thus well-defined for linear representations.

As in our definition of conjugate *matrix* representations, we still need a way of describing when two *linear* representations are essentially the same:

Definition 1.11
Two linear representations $\rho : G \to \mathrm{GL}(V)$ and $\rho' : G \to \mathrm{GL}(W)$ are **isomorphic** if there exists a linear isomorphism $I : V \to W$ such that $I(\rho_g(v)) = \rho'_g(I(v))$ for all $g \in G$ and $v \in V$ — in other words, an isomorphism between vector spaces that is also compatible with the action of G.

Note that given two finite-dimensional vector spaces V and W, there exists an isomorphism between V and W if and only if they have the same dimension. However, we can't just pick *any* isomorphism — our isomorphism should also be compatible with the action of G. (Definition 1.11 essentially states that we should be able to relabel the elements of W as elements of V without changing the vector space structure or the way G acts on the space.) Linear representations up to isomorphism correspond precisely to matrix representations up to conjugacy.

As mentioned earlier, we want to study properties of a representation which don't depend on the basis. For a matrix, operations such as the trace or the determinant are invariant under conjugation, and thus are basis-independent. This motivates the following definition.

Definition 1.12 (Character)
The **character** of a representation R is the function χ_R on G defined as $\chi_R(g) = \mathrm{Tr}(R_g)$.

It might be surprising that we use the trace in this definition, rather than the determinant or some other basis-independent property. However, the trace will turn out to be the right choice, as it is extremely useful. In particular, by the end of next week, it will be shown that for a finite group G, the character completely determines the isomorphism class of the representation!

2 Characters and The Direct Sum

2.1 Review

Last time, we defined linear representations of a group as homomorphisms $\rho : G \to \mathrm{GL}(V)$. We saw that by choosing a basis, we can rewrite linear representations as matrix representations, or homomorphisms $R : G \to \mathrm{GL}_n(\mathbb{C})$. We also introduced the **character** of a representation, which we'll discuss more today (and again in future lectures).

2.2 Characters

Recall that the character of a representation ρ is defined as the function χ_ρ on G where

$$\chi_\rho(g) = \operatorname{Tr} \rho(g)$$

for each $g \in G$. In other words, if we choose a basis and write ρ_g as a $n \times n$ matrix $R_g = (a_{ij})$, then $\chi_\rho(g) = \sum_{i=1}^{n} a_{ii}$.

At first glance, this definition appears to require us to work with *matrix* representations and to specify a basis, since the character is defined as the trace of a matrix. But thankfully, the character does not actually depend on the basis of $\mathrm{GL}(V)$ used to turn a linear representation into a matrix representation — a key property of the trace is that $\mathrm{Tr}(AB) = \mathrm{Tr}(BA)$ for any two matrices A and B, and so $\mathrm{Tr}(A) = \mathrm{Tr}(P^{-1}AP)$ for any invertible matrix P. So the characters of conjugate matrix representations coincide, and therefore the character of a linear representation is well-defined.

The fact that trace is invariant under conjugation also gives us an important property of the character — for any $g, x \in G$ we have

$$\chi_\rho(xgx^{-1}) = \chi_\rho(g),$$

since $\rho(xgx^{-1})$ and $\rho(g)$ are conjugate matrices and therefore have the same trace. So $\chi_\rho(g)$ depends only on the conjugacy class of g (meaning that χ_ρ evaluated at two conjugate elements of G will give the same result). Functions which only depend on the conjugacy class are known as **class functions**; so the character of any representation is a class function, and in order to compute the character, it's enough to compute its value on one representative of each conjugacy class.

Example 2.1
Consider the permutation representation of S_3 on $\mathbb{C}^3$, which we denote by ρ.

The conjugacy classes of S_n are described by cycle type — two elements are in the same conjugacy class if and only if the cycles in their cycle decomposition have the same lengths. So there are three conjugacy classes in S_3, with representatives (1), (12), and (123), respectively.

Clearly the permutation (1) maps to the identity matrix, which has trace 3, and therefore $\chi_\rho(1) = 3$. Meanwhile, the remaining two representatives map to the matrices

$$(12) \mapsto \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}, (123) \mapsto \begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix},$$

so $\chi_\rho(12) = 1$ and $\chi_\rho(123) = 0$. So χ_ρ has the following description:

Conjugacy class	(1)	(12)	(123)
χ_ρ	3	1	0

Note that in a n-dimensional representation, the identity element of G must always map to the identity matrix, which consists of n entries of 1 on the diagonal — so $\chi_\rho(1) = \dim(\rho)$ for any representation ρ.

The simplest example of a character is a one-dimensional character (a character of a one-dimensional representation); such characters have fairly nice properties. If $\dim(\rho) = 1$, then each element ρ_g is a 1×1 invertible matrix, which can be thought of as a single nonzero number — in other words, $\mathrm{GL}_1(\mathbb{C}) = \mathbb{C}^\times$. Then $\chi_\rho(g)$ is just that number, so loosely speaking, we have $\chi_\rho(g) = \rho(g)$.

So in this case, $\chi : G \to \mathbb{C}^\times$ is a homomorphism, meaning that $\chi(gh) = \chi(g)\chi(h)$ for all $g, h \in G$. (This is *not* generally true for representations of higher dimension.) In particular, when G is finite, every $g \in G$ has

finite order, and if $\text{ord}(g) = k$, then $\chi(g)$ must be a kth root of unity — written out explicitly, this is because $\chi(g)^k = \chi(g^k) = \chi(1) = 1$.

Example 2.2
Consider the group $\mathbb{Z}/n\mathbb{Z}$. Any one-dimensional representation of $\mathbb{Z}/n\mathbb{Z}$ is determined by the image of $\overline{1}$. So if we let $\zeta_n = e^{2\pi i/n}$, then we must have

$$\rho(\overline{1}) = \zeta^a = \cos\frac{2\pi a}{n} + i\sin\frac{2\pi a}{n}$$

for some integer a. Then the rest of the representation is given by $\overline{x} \mapsto \zeta_n^{ax}$ for each $\overline{x} \in \mathbb{Z}/n\mathbb{Z}$. So the character of this representation is also $\chi_\rho(\overline{x}) = \zeta_n^{ax}$.

2.3 Direct Sums

Now we'll focus on how various representations of a group relate to each other.

Guiding Question
Given two representations of a group, can we combine them to create a new representation?

One way to combine two representations is by taking their *direct sum*; this will turn out to be an important construction.

Definition 2.3 (Direct Sum)
Let $\psi : G \to \text{GL}(V)$ and $\eta : G \to \text{GL}(W)$ be two representations of the same group. Then their **direct sum** $\rho = \psi \oplus \eta$ is the representation $\rho : G \to \text{GL}(U \oplus W)$ where for each $g \in G$,

$$\rho_g(u, w) = (\psi_g(u), \eta_g(w)).$$

To describe this construction in terms of matrices, we can obtain a basis of $U \oplus W$ by appending the bases of U and W. In this basis, $\rho = \psi \oplus \eta$ will consist of the block diagonal matrices

$$\rho_g = \left(\begin{array}{c|c} \psi_g & 0 \\ \hline 0 & \eta_g \end{array}\right).$$

In particular, if $\rho = \psi \oplus \eta$, then we have

$$\chi_\rho = \chi_\psi + \chi_\eta.$$

Guiding Question
Given a representation, can it be split as a direct sum of smaller representations?

Clearly, if a representation ρ is given in a form where all the matrices ρ_g are block diagonal with blocks of the same dimensions, then it is possible to decompose ρ as a direct sum. The tricky part is when ρ is given by matrices which are not in block diagonal form, but may *become* block diagonal after a change of basis.

Example 2.4
Consider the group $\mathbb{Z}/2\mathbb{Z}$, and take the two-dimensional representation given by

$$\overline{1} \mapsto \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}.$$

Clearly, in the standard basis, this matrix. However, we can instead take the basis consisting of $v_1 = (1,1)^t$ and $v_2 = (1,-1)^t$. Then the matrix can be written as the diagonal matrix

$$\overline{1} \mapsto \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix},$$

and so the representation can in fact be written as a direct sum of two one-dimensional representations — the trivial representation on $\mathrm{Span}(v_1)$, and the representation on $\mathrm{Span}(v_2)$ where $\overline{1} \mapsto \begin{bmatrix} -1 \end{bmatrix}$.

In fact, $\mathbb{Z}/m\mathbb{Z}$ is relatively simple to analyze in general.

Example 2.5
Every n-dimensional representation of $\mathbb{Z}/m\mathbb{Z}$ can be split as the sum of n one-dimensional representations.

Proof. A representation ρ of $\mathbb{Z}/m\mathbb{Z}$ is determined by the matrix A corresponding to $\overline{1}$, since $\overline{1}$ generates $\mathbb{Z}/m\mathbb{Z}$; this matrix must satisfy the constraint $A^m = 1$.

So showing that ρ can be split as the direct sum of one-dimensional representations is equivalent to showing that A can be diagonalized — if A can be diagonalized, then we can find a basis $\{v_1, \ldots, v_n\}$ of V consisting of eigenvectors of A. Then we can split V as a direct sum of the spans of these eigenvectors. If v_i corresponds to the eigenvalue λ_i for each i, then we can write

$$\rho = \psi_1 \oplus \cdots \oplus \psi_n,$$

where ψ_i is the representation given by $\overline{1} \mapsto \lambda_i$, acting on $\mathrm{Span}(v_i)$.

But it's possible to show that any matrix of finite order is in fact diagonalizable, by rewriting it in Jordan normal form and then showing that there can be no Jordan blocks of size greater than one. Since we know A has finite order, it is then possible to diagonalize A, and therefore to decompose ρ as a sum of one-dimensional representations. □

2.4 Irreducible Representations

Now that we've seen how to build new representations out of smaller ones, we would like to describe the "building blocks" of representations.

Definition 2.6
Let $\rho : G \to \mathrm{GL}(V)$ be a representation of G. Then a subspace $W \subset V$ is called G-**invariant** if for each $g \in G$, we have $\rho_g(w) \in W$ for all $w \in W$.

In other words, W is G-invariant if it is taken to itself under the action of each element in G. In 18.701, we saw the concept of a T-invariant subspace for a linear operator T (a subspace taken to itself under the map T); in this situation, a subspace is G-invariant if it is invariant with respect to every one of the operators ρ_g.

Definition 2.7
A representation $\rho : G \to \mathrm{GL}(V)$ is **irreducible** if the only G-invariant subspaces of V are 0 and V.

Note that if a representation $\rho : G \to \mathrm{GL}(V)$ can be decomposed as a direct sum $\rho = \psi \oplus \eta$, where ψ annd η act on nonzero subspaces U and W with $V = U \oplus W$, then U is a G-invariant subspace of V (here U consists of the elements $(u, 0)$ in $V = U \oplus W$). So a direct sum of representations is always reducible (that is, not irreducible).

Meanwhile, if a representation $\rho : G \to \mathrm{GL}(V)$ is reducible, then by definition it has an invariant subspace $W \subset V$. Then we can restrict ρ to W. Initially each ρ_g is an automorphism of V, but since ρ_g preserves W, we can also think of ρ_g as an endomorphism of W (a linear map from W to itself). In fact, ρ_g must be an *automorphism*[¶] of W, since $\rho_{g^{-1}}$ restricted to W is still the inverse of ρ_g restricted to W. So then by restricting each ρ_g to W, we get a *subrepresentation* of ρ acting on W. This means any reducible representation V has a smaller representation W sitting inside it, in some sense.

In matrix form, we can let $\{v_1, \ldots, v_m\}$ be a basis of the invariant subspace W, and complete it to a basis $\{v_1, \ldots, v_n\}$ of V. With respect to this basis, the matrices in ρ are the block matrices

$$\rho_g = \left[\begin{array}{c|c} \psi_g & * \\ \hline 0 & \eta_g \end{array}\right],$$

where ψ is the representation formed by restricting ρ to W, the top-right entries $*$ are "junk," and η is the *quotient representation* $G \to \mathrm{GL}(V/W)$.

We would like to split ρ as a direct sum of ψ and another smaller representation. To do this, we'd like to pick a basis that turns the "junk" into a block of zeros for each $g \in G$. Then if $U = \mathrm{Span}(v_{m+1}, \ldots, v_n)$, we can split $V = W \oplus U$, and since W and U are both G-invariant, this would let us split $\rho = \psi \oplus \eta$.

It isn't immediately clear whether it's always possible to choose the basis in such a way. But as we'll see next class, for complex representations of *finite* groups, it is always possible! More precisely, we'll see the following result:

Theorem 2.8
Given a finite-dimensional complex representation of a finite group, each invariant subspace has a corresponding invariant complementary subspace.

This states that not only does a reducible representation have an invariant subspace $W \subset V$ (which we already know exists by definition), but there is also another invariant subspace $U \subset V$ for which $V = U \oplus W$. This will imply that a representation is reducible if and only if it can be broken down as a direct sum of nonzero representations, leading to the following result:

Theorem 2.9 (Maschke's Theorem)
Every finite-dimensional complex representation of a finite group can be written as a direct sum of irreducible representations.

This will turn out to be an incredibly useful result. Note that the proof requires the group to be *finite* — the theorem will not generally be true for infinite groups, although there is a generalization to compact groups (which we won't discuss).

[¶] An invertible endomorphism

3 Irreducible Representations

3.1 Review

Last time, we discussed characters and direct sums, and began thinking about *irreducible representations*.

In particular, we saw that any representation of the cyclic group $\mathbb{Z}/m\mathbb{Z}$ can be decomposed as a direct sum of one-dimensional representations. To prove this, we saw that it's possible to diagonalize the matrix corresponding to $\overline{1}$ — this matrix has finite order, and it's possible to use Jordan normal form to show that any matrix of finite order is diagonalizable (by showing that it cannot have nontrivial Jordan blocks). Then once we've diagonalized this matrix, the span of each eigenvector provides a one-dimensional subrepresentation, and the original representation is the direct sum of these subrepresentations.

Irreducible representations (commonly abbreviated *irreps* for convenience) will turn out to be the fundamental building blocks for the theory of representations — today we'll discuss Maschke's Theorem, which states that any representation can be decomposed into a sum of irreducible representations.

3.2 Examples

First, let's start with a nontrivial example of an irreducible representation. Recall that given a representation of G acting on V, a subspace W is G-invariant if every element $g \in G$ carries W to itself, or in other words, $gw \in W$ for all $w \in W$. We call the representation on V *irreducible* if the only G-invariant subspaces are V itself and the trivial subspace.

If the representation is *not* irreducible, there exists a G-invariant subspace $W \subset V$. Then by restricting our original representation to W, we get a smaller representation of G.

Example 3.1
Consider the permutation representation of S_3, where each permutation acts on $\mathbb{C}^3$ by permuting the coordinates (so $\sigma \in S_3$ maps $\vec{e}_i \mapsto \vec{e}_{\sigma(i)}$ for each basis vector $\vec{e}_i$). This representation is not irreducible, since the two-dimensional subspace

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

is an invariant subspace, and therefore there is a two-dimensional subrepresentation of the permutation representation acting on V.

Note that the vector $v = (1, 1, 1)^t$ is orthogonal to V, and it is fixed by all permutation matrices. So the permutation representation also has a one-dimensional subrepresentation acting on $\text{Span}(v)$, namely the trivial representation. This means the permutation representation decomposes as the direct sum of its restrictions to V and to $\text{Span}(v)$.

Proposition 3.2
The permutation representation restricted to $V = \{(x, y, z) \mid x + y + z = 0\}$ is irreducible.

Proof. Suppose that $W \subset V$ is a nonzero subspace which is S_3-invariant. Pick a nonzero vector $v = (x, y, z) \in W$. We cannot have $x = y = z$, as this would imply all coordinates are zero, so without loss of generality we may assume that $x \neq y$.

Then since W is G-invariant,

$$(12)v - v = (y - x, x - y, 0) \in W.$$

Since $x - y \neq 0$, by scaling we have $(1, -1, 0) \in W$ as well. Then $(23)(1, -1, 0) = (1, 0, -1) \in W$ as well. But $(1, -1, 0)$ and $(1, 0, -1)$ are linearly independent, and since V is two-dimensional, this means they span V. So we must have $W = V$. This means there are no invariant subspaces of V other than 0 and V, so the representation is irreducible. □

In fact, this statement generalizes to S_n for all n, and the proof is essentially the same.

Proposition 3.3
For every n, the permutation representation of S_n on $\mathbb{C}^n$ has an $(n-1)$-dimensional invariant subspace

$$V = \left\{(x_1, \ldots, x_n) \mid \sum x_i = 0\right\} \subset \mathbb{C}^n,$$

consisting of vectors whose coordinates sum to zero. Furthermore, the representation of S_n obtained by restricting the permutation representation to V is irreducible.

In the case $n = 3$, there's a geometric way to think of this argument as well. Since S_3 is isomorphic to D_3, we can think of it as the group of symmetries of an equilateral triangle — more precisely, consider the equilateral triangle with vertices at $(2, -1, -1)$, $(-1, 2, -1)$, and $(-1, -1, 2)$. Then the actions of the permutations in S_3 on V correspond exactly to the symmetries of this triangle.

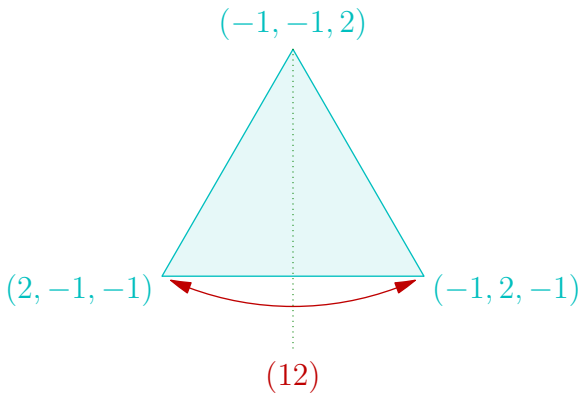

In this interpretation, it's possible to see geometrically that there are no invariant subspaces other than 0 and V (any such subspace would have to be one-dimensional and therefore a line, but just by considering the reflections, we can see that no line is preserved by more than one reflection). However, we have to be careful when reasoning geometrically:

Example 3.4
The representation of $\mathbb{Z}/3\mathbb{Z}$ acting as the group of *rotational* symmetries of an equilateral triangle is *not* irreducible.

From the geometric interpretation, it may appear that this representation is irreducible for the same reason as the representation of S_3 was — after all, it *looks* like if we rotate any line, we get a different one. But something has to have gone wrong, since we saw earlier that *every* representation of a cyclic group is the sum of one-dimensional representations! The problem is that this reasoning only works for *real* representations — and in fact, if we take this as a real representation instead of a complex one, then it *is* irreducible. But working over the complex numbers, there *are* in fact invariant subspaces. The source of the confusion is that although rotations don't have *real* eigenvalues, they do have *complex* eigenvalues.

We can also compute these invariant subspaces explicitly. First, place the equilateral triangle in the plane, in the standard basis:

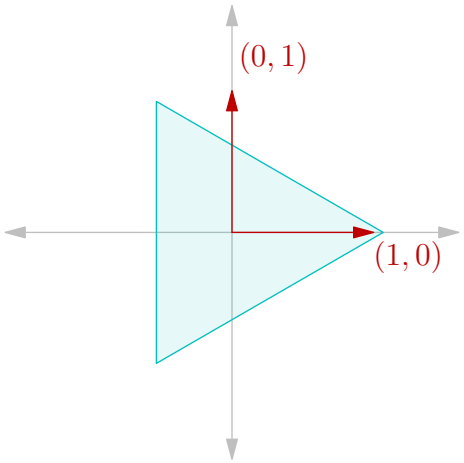

Then we have

$$\overline{1} \mapsto \begin{bmatrix} \alpha & -\beta \\ \beta & \alpha \end{bmatrix},$$

where $\alpha = -1/2$ and $\beta = \sqrt{3}/2$ (this is the matrix corresponding to $2\pi/3$ rotation). To write down an invariant subspace, note that

$$\begin{bmatrix} \alpha & -\beta \\ \beta & \alpha \end{bmatrix} \begin{bmatrix} 1 \\ i \end{bmatrix} = \begin{bmatrix} \alpha - \beta i \\ \beta + \alpha i \end{bmatrix} = (\alpha - \beta i) \begin{bmatrix} 1 \\ i \end{bmatrix},$$

so the span of $(1, i)^t$ is invariant under the action of $\overline{1}$, and therefore all of $\mathbb{Z}/3\mathbb{Z}$. Similarly, the span of $(1, -i)^t$ is also an invariant subspace. So our representation splits as the direct sum of representations on these two subspaces — in the basis consisting of $(1, i)^t$ and $(1, -i)^t$, we have

$$\overline{1} \mapsto \begin{bmatrix} \alpha - \beta i & 0 \\ 0 & \alpha + \beta i \end{bmatrix}.$$

3.3 Invariant Complements

Last class, we began discussing the following question:

Guiding Question
Given a reducible representation, is it possible to decompose it as a direct sum of smaller representations?

Let $\rho : G \to \mathrm{GL}(V)$ be a reducible representation of dimension n, and let $W \subset V$ be a G-invariant subspace of dimension m (with $0 < m < n$). Pick a basis $\{v_1, \ldots, v_n\}$ for V such that the first m basis vectors $\{v_1, \ldots, v_m\}$ form a basis for W. Then G acts by block matrices of the form

$$\left[\begin{array}{c|c} \psi_g & * \\ \hline 0 & \eta_g \end{array}\right],$$

where ψ_g is an $m \times m$ matrix and η_g is an $(n-m) \times (n-m)$ matrix. Since $\rho_g \rho_h = \rho_{gh}$ for all $g, h \in G$, by performing block matrix multiplication we see that $\psi_g \psi_h = \psi_{gh}$ and $\eta_{gh} = \eta_g \eta_h$. So ψ and η are both valid representations — ψ is a representation on W, and η is a representation on the $(n-m)$-dimensional quotient space V/W.

So any reducible representation carries information about two smaller representations ρ and η; and if the top-right corner $*$ is a block of zeros, then in fact $\rho \cong \psi \oplus \eta$. (The converse is not quite true, since $*$ may consist of all zeros in one choice of basis but not another.)

Note that $*$ is a block of zeros if and only if for all $m + 1 \leq i \leq n$,

$$\rho_g(v_i) = \sum_{j=m+1}^{n} a_{ij} v_j$$

for some scalars a_{ij}. This condition is equivalent to requiring that $U = \mathrm{Span}(v_{m+1}, \ldots, v_n)$ is G-invariant as well — intuitively, it states that the two spaces live on their own and do not interact with each other.

So this gives us a way to interpret our condition without thinking about matrices and coordinates! We need to choose our basis vectors so that U is invariant as well, but if that condition is satisfied, then *any* basis of U works. So in a basis-free language, we can decompose $\rho \cong \psi \oplus \eta$ if and only if W has an *invariant complement*. (A *complement* of W is a subspace U such that $W \cap U = 0$ and $W + U = V$, or equivalently such that $V \cong W \oplus U$; so an *invariant complement* is such a subspace U which is also G-invariant.)

Student Question. *Is the invariant complement of a given subspace always unique?*

Answer. *Not necessarily. As a trivial but legitimate example, consider a n-dimensional version of the trivial representation, where every $g \in G$ is sent to the $n \times n$ identity matrix. Then* every *subspace is invariant; so given a subspace W,* all *of its complements are invariant complements of W.*

So our question about whether we could decompose ρ as a sum of smaller representations reduces to the following:

Guiding Question
Given an invariant subspace $W \subset V$, does it necessarily have an invariant complement?

In the *general* case, the answer is no — there are situations where there is an invariant subspace with no invariant complement (and therefore the representation cannot be split as a direct sum, even though it's not irreducible).

Example 3.5
Take the representation ρ of $\mathbb{Z}$ acting on $V = \mathbb{C}^2$ given by

$$1 \mapsto \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix},$$

which means that for all n,

$$n \mapsto \begin{bmatrix} 1 & n \\ 0 & 1 \end{bmatrix}.$$

The vector $(1, 0)^t$ is an eigenvector of every matrix in this representation, with eigenvalue 1 — so its span W is an invariant subspace, and ρ has a subrepresentation on W (which is the trivial representation).

But W does not have an invariant complement. To see this, note that ρ acts trivially on V/W as well (since the entry in the bottom-right corner is 1). So if W had an invariant complement, then ρ would be the direct sum of two trivial representations, which is clearly not the case as ρ is nontrivial.

3.4 Maschke's Theorem

As the previous example shows, in general it may not always be possible to find an invariant complement. But it turns out that for *finite* groups, it *is* always possible! (From now, we'll assume that all representations are complex and finite-dimensional, and our group is finite, unless stated otherwise.)

The proof that an invariant complement always exists will include two parts — first we'll define a *Hermitian form* with useful properties, and then we'll use this Hermitian form to construct the desired complement. We'll start by stating the two main steps:

Lemma 3.6
If $\rho : G \to \mathrm{GL}(V)$ is a complex representation of a finite group, then there exists a G-invariant positive Hermitian form on V.

To explain this terminology, recall that a Hermitian form $\langle -, - \rangle$ is a pairing $V \times V \to \mathbb{C}$ such that:

- It is linear in the first variable — we have $\langle v_1 + v_2, w \rangle = \langle v_1, w \rangle + \langle v_2, w \rangle$ for all $v_1, v_2, w \in V$, and $\langle \lambda v, w \rangle = \lambda \langle v, w \rangle$ for all $v, w \in V$.
- We have $\langle w, v \rangle = \overline{\langle v, w \rangle}$ for all $v, w \in V$.

The second condition immediately implies that $\langle v, v \rangle$ is real for all $v \in V$; the Hermitian form is *positive* if $\langle v, v \rangle$ is in fact positive for all $v \neq 0$.

Meanwhile, a Hermitian form is G-invariant if it's preserved by the G-action, meaning that $\langle gv, gw\rangle = \langle v, w\rangle$ for all $v, w \in V$ and $g \in G$.

Lemma 3.7
If $\rho : G \to \mathrm{GL}(V)$ has an invariant Hermitian form, then every invariant subspace of V has an invariant complement.

Now let's prove these two steps; we'll start with the second.

Proof of Lemma 3.7. As discussed in 18.701, if $\langle -, -\rangle$ is a positive Hermitian form, then the orthogonal complement of any subspace W, the subspace

$$W^\perp = \{v \mid \langle v, w\rangle = 0 \text{ for all } w \in W\},$$

is a complementary subspace to W. But now if $\langle -, -\rangle$ and W are both G-invariant, then so is $W^\perp$ — if v is in $W^\perp$, then $\langle v, w\rangle = 0$ for all $w \in W$, so for each $g \in G$, we have $\langle gv, gw\rangle = 0$ for all $w \in W$ as well; since $gW = W$, this means gv is in $W^\perp$ as well. □

Now that we've seen why having such a Hermitian form is useful, let's construct one.

Proof of Lemma 3.6. Start with *any* positive Hermitian form $\langle -, -\rangle'$, and now employ the *averaging trick* — take our Hermitian form to be

$$\langle v, w\rangle = \sum_{g \in G} \langle gv, gw\rangle.$$

It's clear that this is a positive Hermitian form. To see that it's G-invariant, for any $h \in G$ we have

$$\langle hv, hw\rangle = \sum_{g \in G} \langle hgv, hgw\rangle' = \sum_{g \in G} \langle gv, gw\rangle' = \langle v, w\rangle,$$

since as g ranges over all elements of G, so does hg. □

Putting these two steps together, we get that if we have a reducible representation $\rho : G \to \mathrm{GL}(V)$ with an invariant subspace $W \subset V$ (which must exist by definition), then we can always find an invariant complement of W, and we can therefore decompose ρ as a direct sum! So by repeatedly splitting up a representation until our components are irreducible (or more formally, using induction on the dimension), we get the following theorem:

Theorem 3.8 (Maschke's Theorem)
Every complex, finite-dimensional representation of a finite group is a direct sum of irreducible representations.

4 The Main Theorem

4.1 More on Maschke's Theorem

Last class, given a finite-dimensional complex representation $\rho : G \to \mathrm{GL}(V)$ of a finite group G, we found a G-invariant positive Hermitian form on V and used it to show that a G-invariant subspace W has an invariant complement, namely $W^\perp$ (its orthogonal complement with respect to the Hermitian form). We used this to deduce Maschke's Theorem — that every representation can be split as a direct sum of irreducible ones.

We'll now discuss a few features of this proof.

First, why did we use Hermitian forms specifically? A different choice of form which may also seem reasonable is the *symmetric bilinear form*, a form where we require $\langle w, v\rangle$ to equal $\langle v, w\rangle$ rather than its conjugate. (For example, $v \cdot w$ is the standard symmetric bilinear form, while $v \cdot \overline{w}$ is the standard Hermitian form.)

The reason is that in a symmetric bilinear form over $\mathbb{C}$, it's possible that $v \cdot v = 0$. For example, consider the representation of $\mathbb{Z}/3\mathbb{Z}$ acting on $\mathbb{C}^3$, where

$$\overline{1} \mapsto A = \begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}.$$

We can see that $(1, \zeta, \zeta^2)^t$ is an eigenvector, since

$$\begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} 1 \\ \zeta \\ \zeta^2 \end{bmatrix} = \begin{bmatrix} \zeta^2 \\ 1 \\ \zeta \end{bmatrix} = \zeta^2 \begin{bmatrix} 1 \\ \zeta \\ \zeta^2 \end{bmatrix}.$$

But if we tried to perform our construction, taking W to be the span of this eigenvector, we'd see that $W^\perp$ actually *contains* it, since $(1, \zeta, \zeta^2) \cdot (1, \zeta, \zeta^2) = 1 + \zeta^2 + \zeta^4 = 0$. This means $W^\perp$ isn't actually a complement of W, so this would break the construction. We require that our form is *Hermitian* (and positive) to avoid this issue, since in that case $W^\perp$ really is a complement of W.

Another useful takeaway from our proof was that we found an invariant Hermitian form by *averaging*. This trick of averaging over all $g \in G$ can produce many other invariant things.

Example 4.1
Given a representation $\psi : G \to \mathrm{GL}(V)$ and a vector $v \in V$, the vector

$$\frac{1}{|G|} \sum_{g \in G} \psi_g v$$

is G-invariant. This is because for any $h \in H$, we have

$$\psi_h \frac{1}{|G|} \sum_{g \in G} \psi_g v = \frac{1}{|G|} \sum_{g \in G} \psi_{hg} v = \frac{1}{|G|} \sum_{g \in G} \psi_g v,$$

since $g \mapsto hg$ is a bijection on G (as g runs over all of G, so does hg).

Note 4.2
We saw a similar trick in 18.701, when proving that every finite group of isometries of $\mathbb{R}^2$ (or more generally $\mathbb{R}^n$) has a fixed point — we can start with any point p, and consider the points gp in its orbit. Then the center of mass (or average) of all these points is a fixed point, for the same reason we saw here.

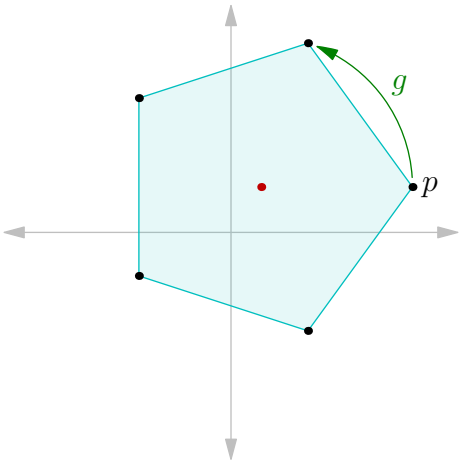

There's something to be careful of here, though — in Example 4.1, we don't actually know that this vector is nonzero — in fact, for many representations ψ it *must* be zero (often there's no nontrivial invariant vectors).

In fact, we can describe our construction *directly* in terms of the averaging trick as described in Example 4.1. We can think of the space of Hermitian forms as a *real* vector space; then G has a real representation acting on this space, sending $\langle v, w\rangle$ to $\langle gv, gw\rangle$. In our construction, we started with some positive form, and averaged it over all g to get an invariant "vector" (meaning an invariant Hermitian form). Here we don't have the issue of the invariant vector possibly being zero, because when we add two positive Hermitian forms, our resulting Hermitian form is again positive.

Student Question. *How do we describe the space of Hermitian forms as a real vector space?*

Answer. *Every Hermitian form can be described as $\langle v, w\rangle = vA\overline{w}$ for some matrix A with $A^t = \overline{A}$. We can think of all entries of this matrix in terms of their real and complex parts. Then the n entries on the diagonal must all be real, we get to choose both the real and complex parts of the $n(n-1)/2$ entries below the diagonal, and this immediately determines the $n(n-1)/2$ entries above the diagonal (which must be the conjugate of their reflection). So we get to choose*

$$n + 2 \cdot \frac{n(n-1)}{2} = n^2$$

real numbers, which means the space of Hermitian forms has dimension n^2.

On a different note, the fact that every representation has an invariant positive Hermitian form (as shown in our proof) is equivalent to stating that every representation of a finite group is conjugate to a *unitary representation*:

Definition 4.3
A unitary representation is a homomorphism $\rho : G \to \mathrm{U}_n$, where $\mathrm{U}_n \subset \mathrm{GL}_n$ is the set of unitary matrices.[a]

[a] Matrices A for which $A^t = \overline{A}$

Equivalently, we can define unitary representations without referring to matrices — a linear operator is unitary if it preserves the standard Hermitian form $\langle v, w\rangle = v \cdot \overline{w}$, so a representation is unitary if and only if $gv \cdot g\overline{w} = v \cdot \overline{w}$ for all $g \in G$ and $v, w \in \mathbb{C}^n$.

To see why these two ideas are equivalent, note that given a positive Hermitian form, we can choose an orthonormal basis with respect to that form. In that basis, the form will just be the standard Hermitian form $\langle v, w\rangle = v \cdot \overline{w}$. So by changing the basis, we have produced a unitary representation.

Finally, we'll describe the proof of Maschke's Theorem (also known as *complete reducibility*) more explicitly.

Theorem 4.4 (Maschke's Theorem)
Every complex representation of a finite group is isomorphic to a direct sum of irreducible representations.

Proof. We use induction on the dimension. For the base case, any one-dimensional representation is already irreducible.

Now suppose $\rho : G \to \mathrm{GL}(V)$ is a representation. If ρ is already irreducible, we're done. Otherwise, we can pick an invariant subspace W, which is neither 0 nor V. Then let $\langle -, - \rangle$ be an invariant positive Hermitian form, and decompose $V = W \oplus W^\perp$. Since both W and $W^\perp$ are G-invariant, then we get the subrepresentations $\psi : G \to \mathrm{GL}(W)$ and $\eta : G \to \mathrm{GL}(W^\perp)$ in ρ, and we can decompose $\rho \cong \psi \oplus \eta$. But η and ψ have smaller dimension than ρ, so both are a direct sum of irreducible representations (by the inductive hypothesis), and therefore ρ is a direct sum of irreducible representations as well. □

4.2 More on Characters

Let's continue our discussion from earlier about characters. We'll first state a few basic properties. (We will assume that all our representations are of *finite* groups, unless stated otherwise.)

Proposition 4.5
If $\rho : G \to \mathrm{GL}(V)$ is a complex representation, then

(a) $\chi_\rho(g)$ is a sum of roots of unity;

(b) $\chi_\rho(g^{-1}) = \overline{\chi_\rho(g)}$;

(c) $\overline{\chi_\rho}$ is the character of another representation of the same dimension, denoted ρ^* and called the *dual representation*.

Proof. These properties all come from the definition of the character as the trace of a matrix.

For (a), since G is a finite group, each $g \in G$ and therefore $\rho_g \in \mathrm{GL}(V)$ has finite order, so the eigenvalues of ρ_g, which we denote by $\lambda_i(g)$, are all roots of unity. Then $\mathrm{Tr}(\rho_g) = \sum \lambda_i(g)$ is a sum of roots of unity (as the trace is always the sum of eigenvalues with multiplicity).

For (b), the eigenvalues of $\rho_{g^{-1}}$ are the inverses of the eigenvalues of ρ_g (since the two matrices are inverses), and since the eigenvalues of ρ_g are all roots of unity, their inverses are equal to their conjugates. So we have

$$\mathrm{Tr}(\rho_{g^{-1}}) = \sum (\lambda_i(g))^{-1} = \sum \overline{\lambda_i(g)} = \overline{\mathrm{Tr}(\rho_g)}.$$

Finally, for (c), let V^* be the *dual space* of V, consisting of linear maps $f : V \to \mathbb{C}$. For convenience, we'll denote $f(v)$ by $\langle f, v \rangle$ (to emphasize the fact that we can think of it as a pairing between a vector v and a covector f). Then the dual representation ρ^* is given by

$$\langle \rho_g^*(f), v \rangle = \langle f, \rho_{g^{-1}}(v) \rangle.$$

(This defines $\rho_g^*(f)$ for each $f \in V$, by describing where it takes each vector.)

In other words, we make G act on V^* such that for every $v \in V$ and $f \in V^*$, we have

$$\langle f, v \rangle = \langle \rho_{g^*}(f), \rho_g(v) \rangle.$$

That is, the dual representation is defined such that operating on both the vector v (with the representation ρ) and the covector f (with the dual representation ρ^*) does not affect the pairing $\langle f, v \rangle$ given by the dual space. (Our original definition then follows from plugging in $\rho_{g^{-1}}(v)$ in place of v, to make the definition more explicit.)

This may seem somewhat abstract, but we can make it more concrete by describing it in terms of matrices. Fix a basis of V, so then from ρ we get a matrix representation $R : G \to \mathrm{GL}_n(\mathbb{C})$. Then the dual representation in terms of matrices is given by

$$R_g^* = R_{g^{-1}}^t.$$

This is a valid representation since $(AB)^t = B^t A^t$ and $(AB)^{-1} = B^{-1} A^{-1}$, so $((AB)^t)^{-1} = (B^t A^t)^{-1} = (A^t)^{-1} (B^t)^{-1}$ (intuitively, each of taking the inverse and transposing means we need to swap the two matrices, so doing both means we need to swap twice and get back our original order). Since $\mathrm{Tr}(A^t) = \mathrm{Tr}(A)$, we have

$$\chi_{R^*}(g) = \chi_R(g^{-1}) = \overline{\chi_R(g)}.$$ □

Student Question. *Why is the transpose important — if there was no transpose, would we still get a representation?*

Answer. *We'd get what's called an* anti-representation *instead — we'd have a map with the property that* $\rho_{gh} = \rho_h \rho_g$. *(This is a representation if* G *is abelian.) It's possible to get an anti-representation in two ways — by inverting the elements, or by taking their transposes — and doing* both *gives us back a valid representation.*

Student Question. *Why are the two definitions of the dual representation (the abstract one and the one given in terms of matrices) equivalent, and how do we get the formula for the character from the abstract definition?*

Answer. *One way to think about this is to first think in terms of matrices — in that setting, it's clear that* $\chi_{R^*} = \overline{\chi_R}$. *Now in the abstract setting, we can pick a basis for* V. *This gives a basis for* V^* *as well — given a basis* $\{v_1, \dots, v_n\}$ *for* V, *we take* f_i *to be the function which is* 1 *on* v_i *and* 0 *on each of the other basis vectors. Then our abstract definition is equivalent to taking the inverse transpose of the corresponding matrices.*

Student Question. *If* $\rho : G \to \mathrm{GL}(V)$, *does there exist a representation on the* same *vector space* V *with character* $\overline{\chi_\rho}$*?*

Answer. *Technically, yes. The dual space* V^* *is isomorphic to* V *— they have the same dimension, so fixing a basis for each gives an isomorphism between them. But they're not isomorphic in a canonical way.*

4.3 The Main Theorem

To understand the main theorem, we need to understand the space of *class functions*.

Definition 4.6
A **class function** is a function $f : G \to \mathbb{C}$ which is fixed on each conjugacy class of G. That is, for a class function f, if g and h are conjugate, then $f(g) = f(h)$.

The space of class functions is a vector space over $\mathbb{C}$, with addition and scalar multiplication defined as usual for functions.

We'll now state the main theorem in our story about representations, which gives surprisingly detailed information about the characters of irreducible representations.

Theorem 4.7 (Main Theorem)
Let G be a finite group. Then:

(a) The characters of irreducible representations form a basis in the space of class functions on G.

(b) This basis is *orthonormal* with respect to the Hermitian form on the space of class functions given by

$$\langle f_1, f_2 \rangle = \frac{1}{|G|} \sum_{g \in G} f_1(g) \overline{f_2(g)}.$$

(c) If $d_1, \dots, d_m$ are the dimensions of the irreducible representations of G, then

$$d_1^2 + d_2^2 + \cdots + d_n^2 = |G|,$$

and each d_i divides $|G|$.

In particular, (a) also implies that the characters of different irreducible representations are *distinct*, since they must form a basis.

For (c), note that although this isn't written in terms of characters, we can interpret it as a statement about characters as well, since $\dim(\rho) = \chi_\rho(1)$.

We'll prove these properties in later classes; first we'll look at a few important implications.

Corollary 4.8
The character of a representation uniquely determines the representation, up to isomorphism.

Proof. By Maschke's Theorem, we know that any representation ρ can be decomposed as a sum of irreducibles — if we use $\rho_1, \dots, \rho_n$ to denote the irreducible representations, then by grouping together isomorphic summands,

we can write

$$\rho = \bigoplus_i \rho_i^{n_i}$$

for some integers n_i (the notation ψ^k, or $\psi^{\oplus k}$, denotes a direct sum of k copies of the representation ψ). But then we have

$$\chi_\rho = \sum n_i \chi_{\rho_i}.$$

So the coefficients n_i when decomposing ρ as a sum of irreducibles are the same as the coefficients when decomposing χ_ρ as a sum of χ_{ρ_i}. But since the χ_{ρ_i} form a *basis* of the space of class functions, there's a *unique* way to write χ_ρ (which is a class function) as a linear combination of these χ_{ρ_i}! So the n_i are uniquely determined from the character of ρ, and therefore so is ρ itself. □

Corollary 4.9
The number of irreducible representations of G is the number of conjugacy classes on G.

Proof. The number of irreducible representations of G is the dimension of the space of class functions (since their characters form a basis for this space). But this dimension is just the number of conjugacy classes, since to specify a class function $f : G \to \mathbb{C}$, we need to specify its value on each conjugacy class. □

This theorem gives a lot of concrete information that we can figure out about irreducible representations by just looking at the group; we'll see some examples of this next class.

5 Characters and Schur's Lemma

5.1 Review

Last time, we stated the Main Theorem about characters.

Theorem 5.1
Let G be a finite group, and let $\rho_1, \ldots, \rho_n$ be a full list of irreducible representations up to isomorphism.

(a) The characters $\chi_{\rho_0}, \ldots, \chi_{\rho_n}$ form a basis for the space of class functions on G.

(b) The basis formed by $\chi_{\rho_0}, \ldots, \chi_{\rho_n}$ is orthonormal.

(c) If $d_i = \dim \rho_i$ for each i, then $\sum d_i^2 = |G|$, and each d_i divides $|G|$.

Today we will look at some ways this can be used to describe characters, and begin developing the tools needed to prove it.

5.2 Character Tables

The information about characters can be put into a table, known as a **character table**, where the columns correspond to conjugacy classes and the rows to irreducible representations (this is enough to record *all* information about the irreducible characters, since characters are constant on each conjugacy class).

A general observation we can make is that $\chi_\rho(1_G) = \dim \rho$ for any representation ρ — this is because $\chi_\rho(G)$ is the identity matrix of dimension $\dim \rho$, which whose trace is $\dim \rho$.

Example 5.2
Consider the group $\mathbb{Z}/4\mathbb{Z}$. Since it's an abelian group, each conjugacy class has one element. As we've seen before, the irreducible representations of $\mathbb{Z}/4\mathbb{Z}$ are all one-dimensional, and must send $\overline{1}$ to any fourth root of unity. The choice of which one determines the rest of the representation, since $\overline{1}$ is a generator.

Using χ_j to denote the character of the representation where $\overline{1} \mapsto i^j$ (so χ_0 is the trivial representation), we have the following table:

	$\overline{0}$	$\overline{1}$	$\overline{2}$	$\overline{3}$
$\chi_0 = 1$	1	1	1	1
χ_2	1	-1	1	-1
χ_1	1	i	-1	$-i$
χ_3	1	$-i$	-1	i

A very similar story occurs for $\mathbb{Z}/m\mathbb{Z}$ for any integer m — the irreducible representations of $\mathbb{Z}/m\mathbb{Z}$ are all one-dimensional and send $\overline{1}$ to a mth root of unity, so a full list of irreducible characters is

$$\chi_a : \overline{x} \mapsto e^{2\pi a x/m}$$

for all $0 \leq a \leq m-1$.

One observation about this table is that the product of any two rows is another row in the table. This isn't a coincidence — if χ and χ' are one-dimensional characters, then $\chi\chi'$ is again a one-dimensional character (since for a one-dimensional representation, χ_ρ is essentially the same as ρ, and is therefore a homomorphism).

Student Question. *If you multiply two characters which are* not *one-dimensional, do you still get another character?*

Answer. *Yes — we'll actually come across a construction for this in a later class, when proving the Main Theorem! But the new character will generally not be irreducible, even if the two characters we started with were — so it won't generally be an entry in the character table.*

The fact that the product of two one-dimensional characters is also a character can be used to check orthogonality quite directly:

Proposition 5.3
Any two one-dimensional characters are orthogonal.

Proof. For two one-dimensional characters χ and χ', we want to show that

$$\sum_{g \in G} \chi'(g)\overline{\chi(g)} = \begin{cases} |G| & \text{if } \chi = \chi' \\ 0 & \text{otherwise.} \end{cases}$$

We have $\overline{\chi(g)} = \chi(g)^{-1}$ since $\chi(g)$ is a root of unity. Now define the one-dimensional representation $\psi = \chi'\overline{\chi}$. Then $\psi(g) = \chi'(g)\chi(g)^{-1}$ for each g, so ψ is trivial (meaning $\psi(g) = 1$ for each g) if and only if $\chi' = \chi$.

So now proving orthogonality reduces to a statement about *one* representation, instead of two — we want to check that

$$\sum_{g \in G} \psi(g) = \begin{cases} |G| & \text{if } \psi \text{ is trivial} \\ 0 & \text{otherwise.} \end{cases}$$

The first case is obvious — if ψ is trivial, we're just summing $|G|$ copies of 1. For the second case, we use a familiar trick — let $S = \sum_{g \in G} \psi(g)$, and pick some $g_0 \in G$ such that $\psi(g_0) \neq 1$ (which exists since ψ is nontrivial); let $\psi(g_0) = \lambda$. Now we have

$$\lambda S = \sum_{g \in G} \psi(g_0)\psi(g) = \sum_{g \in G} \psi(g_0 g) = \sum_{g \in G} \psi(g) = S,$$

since as g runs over G, so does $g_0 g$. So then $\lambda S = S$ for some $\lambda \neq 1$, which means $S = 0$. □

Example 5.4
To make the above argument a bit more concrete, consider the case where our group is $\mathbb{Z}/n\mathbb{Z}$ and ψ is the representation $\overline{x} \mapsto e^{2\pi i x/n}$. Then the set of points $\psi(g)$ form a regular n-gon, and our claim is that the center of mass of this regular n-gon is the origin. Our proof corresponds to the fact that rotating the n-gon by $2\pi/n$ preserves the n-gon and therefore its center of mass; but it must also rotate the center by $2\pi/n$, and the only point fixed by a $2\pi/n$ rotation is the origin.

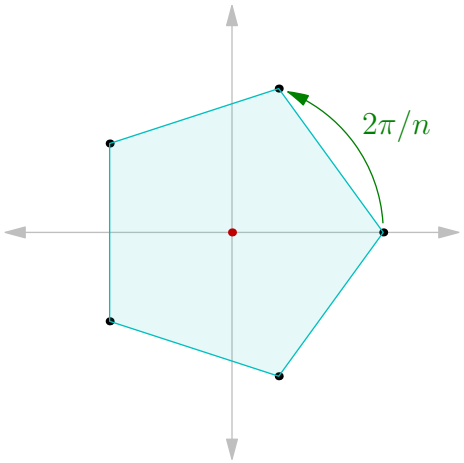

Proposition 5.5
If G is abelian, then every irreducible representation is one-dimensional.

This has several proofs.

Proof 1. We use the Main Theorem. Let $d_1, \ldots, d_n$ be the dimensions of the irreducible representations of G. First, there must be exactly $|G|$ irreducible representations, since the dimension of the space of class functions

is $|G|$ (every function is a class function, as every element is in its own conjugacy class). So we have $n = |G|$. But then

$$d_1^2 + d_2^2 + \cdots + d_n^2 = |G| = n.$$

Since the d_i are positive integers, they must all be 1. □

Proof 2 (Sketch). We've already proven this for *cyclic* groups — we've proven that for $\mathbb{Z}/n\mathbb{Z}$, all irreducible representations are one-dimensional and are given by $\rho_k : \overline{1} \mapsto e^{2\pi ik/n}$.

In a much later class, we'll see that every finite abelian group is isomorphic to a product of cyclic groups, meaning that $G = G_1 \times \cdots \times G_n$ where each G_i is of the form $\mathbb{Z}/m_i\mathbb{Z}$. From this, we can deduce the proposition for *all* abelian groups. One way to do so is to note that if we take any list of irreducible characters χ_1, …, χ_n of G_1, …, G_n (which are all one-dimensional), we can define a one-dimensional character of G as $\chi(g_1, \ldots, g_n) = \chi_1(g_n) \cdots \chi_n(g_n)$. Since each G_i has $|G_i|$ irreducible characters, this gives us $|G_1| \cdots |G_n| = |G|$ one-dimensional characters of $|G|$. But we know there are exactly $|G|$ irreducible characters, so this list must contain all of them. □

Proof 3. Finally, here is a proof that doesn't rely on as much theory.

Claim. *If we have a collection of pairwise commuting matrices, each of which is diagonalizable, then we can diagonalize* all *of them simultaneously.*

Proof. The key point is that if $AB = BA$, then if v is an eigenvector of A with $Av = \lambda v$, then

$$A(Bv) = BAv = B(\lambda v) = \lambda(Bv),$$

so Bv is also a λ-eigenvector, and so the λ-eigenspace of A is B-invariant.

Then using this, we can take the first matrix A, and split our vector space as a direct sum of the eigenspaces of A. Now no matter what basis we choose for those eigenspaces, A will be diagonalized (since it acts as a scalar matrix on each space). Meanwhile, since these eigenspaces are each invariant under all the other matrices, it's enough to diagonalize our remaining set of matrices on each of those spaces. So we can finish by using induction on the number of matrices. □

In our situation, we know all matrices ρ_g are diagonalizable (since they must have finite order). So by the claim, it must be possible to diagonalize them simultaneously; this then splits V as a direct sum of one-dimensional subspaces which are invariant under all ρ_g, and therefore ρ is a direct sum of one-dimensional representations. □

Now let's calculate the character tables for a few symmetric groups. In the first character table for $\mathbb{Z}/4\mathbb{Z}$, we already *knew* what all the irreducible representations (and their characters) were. But here we'll see that we can actually use the Main Theorem to *deduce* information about the characters — for example, knowing that we have listed *all* irreducible characters, or even calculating their values.

Example 5.6
Find the character table of S_3.

Solution. There are three conjugacy classes, with representatives 1, (12), and (123). Meanwhile, we've already seen three irreducible representations — the trivial representation $\mathbb{1}$, the sign representation sgn, and the representation τ, the two-dimensional subrepresentation of the permutation representation acting on $V = \{(x, y, z) \mid x + y + z\} \subset \mathbb{C}^3$.

The characters of the trivial and sign representations are easy to compute. For τ, we know that $\mathbb{C}^3 = V \oplus \text{Span}((1, 1, 1)^t)$, and the permutation representation is trivial on $\text{Span}((1, 1, 1)^t)$, so if ρ denotes the permutation representation on $\mathbb{C}^3$, then $\rho = \tau \oplus \mathbb{1}$. This means

$$\chi_\rho = \chi_\rho + \chi_{\mathbb{1}} = \chi_\rho + 1.$$

But $\chi_\rho(\sigma)$ is just the number of fixed points of σ — this is because ρ_σ is the permutation matrix corresponding to σ, so 1's on the diagonal of ρ_σ correspond to fixed points of σ. This gives the following table:

	1	(12)	(123)
$\mathbb{1}$	1	1	1
sgn	1	-1	1
τ	2	0	-1

Since we've already found three irreducible representations, we know these are the only ones (since the number of irreducible representations is the same as the number of conjugacy classes).

As an example of the orthogonality of the character,

$$\langle \tau, \tau \rangle = \frac{1}{6}(2 \cdot 2 \cdot 1 + 0 \cdot 0 \cdot 3 + (-1)(-1) \cdot 2) = \frac{1}{6}(4+2) = 1.$$

Note that we have to keep track of how many group elements are in each conjugacy class, since we're summing over the entire group and not conjugacy classes — for example, there are two elements in the conjugacy class of (123), which is why we multiply $(-1)(-1)$ by 2. □

Example 5.7
Find the character table of S_4.

Proof. There are five conjugacy classes, with representatives 1, (12), (12)(34), (123), and (1234). There are a few representations we already know — $\mathbb{1}$ and sgn are still irreducible representations, and so is τ, the permutation representation acting on $V = \{(w,x,y,z) \mid w+x+y+z = 0\} \subset \mathbb{C}^4$ (whose character we can compute in the same way as we did for S_3). So far, this gives us the following table:

	1	(12)	(12)(34)	(123)	(1234)
$\mathbb{1}$	1	1	1	1	1
sgn	1	-1	1	1	-1
τ	3	1	-1	0	-1

We earlier mentioned that the product of one-dimensional characters is again a character. But we actually only need *one* of the characters to be one-dimensional, not both of them; so $\chi_{\text{sgn}} \cdot \chi_\tau$ is also an irreducible character, of the representation denoted $\text{sgn} \otimes \tau$. Now there's one remaining representation which we *don't* know how to describe (since there must be exactly five representations); denote that by ρ.

	1	(12)	(12)(34)	(123)	(1234)
$\mathbb{1}$	1	1	1	1	1
sgn	1	-1	1	1	-1
τ	3	1	-1	0	-1
$\text{sgn} \otimes \tau$	3	-1	-1	0	1
ρ					

In order to figure out the last row, first we can use the fact that $\sum d_i^2 = |G|$ to calculate the dimension of ρ — this gives

$$1^2 + 1^2 + 3^2 + 3^2 + (\dim \rho)^2 = 24,$$

so $\dim \rho = 2$. This means $\chi_\rho(1) = 2$.

	1	(12)	(12)(34)	(123)	(1234)
$\mathbb{1}$	1	1	1	1	1
sgn	1	-1	1	1	-1
τ	3	1	-1	0	-1
$\text{sgn} \otimes \tau$	3	-1	-1	0	1
ρ	2	a	b	c	d

In order to calculate the remaining entries a, b, c, and d, we use the orthogonality relations. By using the fact that $\langle \chi_\rho, \chi_{\mathbb{1}} - \chi_{\text{sgn}} \rangle$ and $\langle \chi_\rho, \chi_\tau - \chi_{\text{sgn} \otimes \tau} \rangle$ are both zero, we can get that $a = d = 0$, and by using two other orthogonality relations, we can get $b = 2$ and $c = -1$. So our answer is the following table:

	1	(12)	(12)(34)	(123)	(1234)
$\mathbb{1}$	1	1	1	1	1
sgn	1	-1	1	1	-1
τ	3	1	-1	0	-1
$\text{sgn} \otimes \tau$	3	-1	-1	0	1
ρ	2	0	2	-1	0

Although we were able to compute χ_ρ without knowing what ρ is by using the Main Theorem, it's also possible to describe ρ explicitly. Note that the Klein 4-group K_4 is a normal subgroup of S_4, with $S_4/K_4 \cong S_3$. This means we have a homomorphism $S_4 \to S_3$, and a homomorphism $S_3 \to \text{GL}_2(\mathbb{C})$ given by the representation τ of S_3. Composing these gives a homomorphism $S_4 \to \text{GL}_2(\mathbb{C})$, which is exactly the representation ρ. □

5.3 Schur's Lemma

Now we'll start proving the Main Theorem, starting with orthogonality of characters. For that, we'll need Schur's Lemma, which is quite important in its own right.

We've defined what *one* representation is, but we can also ask how to compare *two* representations. The way to do this is by looking at G-equivariant maps between them:

Definition 5.8
Suppose $\rho : G \to \text{GL}(V)$ and $\psi : G \to \text{GL}(W)$ are (not necessarily irreducible) representations. Then define

$$\text{Hom}_G(\rho, \psi) = \{f : V \to W \mid f \text{ is a linear map such that } f(\rho_g(v)) = \psi_g(f(v)) \text{ for all } g,\, v\}.$$

Such linear maps f are called G-**equivariant**.

Intuitively, $\text{Hom}_G(\rho, \psi)$ is the space of linear maps (or homomorphisms) from V to W which are compatible with the G-action — such maps f are said to *intertwine* the G-action.

Also, $\text{Hom}_G(\rho, \rho)$ is also denoted as $\text{End}_G(\rho)$; this is the space of G-equivariant endomorphisms (an endomorphism is a homomorphism from a space to itself).

Note that $\text{Hom}_G(\rho, \psi)$ is a $\mathbb{C}$-vector space — we can add two G-equivariant homomorphisms or scale one, and get another G-equivariant homomorphism. We'll see later how to think about it in terms of matrices; but for now, we'll prove the following important theorem about it:

Theorem 5.9 (Schur's Lemma)
Let $\rho : G \to \text{GL}(V)$ and $\psi : G \to \text{GL}(W)$ be *irreducible* representations. Then $\text{Hom}_G(\rho, \psi)$ is 0 if $\rho \ncong \psi$, and is one-dimensional if $\rho \cong \psi$.

In other words, the second statement can be written as $\text{End}_G(\rho) = \mathbb{C} \cdot \text{Id}$ — any G-equivariant endomorphism is scalar. It's clear that every scalar map is a G-equivariant endomorphism, so the second part of Schur's Lemma states that these are the only ones.

Proof. Suppose $f : V \to W$ is a nonzero linear map which is G-equivariant. We're now going to show that f is an isomorphism; it suffices to show that it's both surjective and injective.

First consider $\text{im}(f)$; this must be a G-invariant subspace of W, since for any $f(v) \in \text{im}(f)$, we have that $\psi_g(f(v)) = f(\rho_g(v))$ is also in the image of f for any g. But since ψ is irreducible, the only G-invariant subspaces are 0 and W itself; so since f is nonzero (and therefore its image is nonzero), its image must be the entire space W! So f is surjective.

Now consider $\ker(f)$. This is also a G-invariant subspace of V, since if $f(v) = 0$, then $f(\rho_g(v)) = \psi_g(f(v)) = 0$, so ρ_g is also in the kernel for any g. But since ρ is irreducible, the only G-invariant subspaces are 0 and V; and the kernel cannot be V since f is nonzero, so the kernel must be 0. This means f is injective.

Therefore f is both surjective and injective, and is therefore an isomorphism. So we've shown that if there exists a nonzero G-equivariant homomorphism, then we must have $\rho \cong \psi$; this proves the first part of Schur's Lemma.

Now to prove the second part of Schur's Lemma, assume $\rho = \psi$, so f is a map $V \to V$. Then f must have some eigenvalue λ. Now consider the map $f - \lambda \operatorname{Id}$, where Id denotes the identity map; this is also a G-equivariant endomorphism. We must have $\ker(f - \lambda \operatorname{Id}) \neq 0$, since f has a λ-eigenvector (which must be in the kernel). But since the kernel must be a G-invariant subspace, this implies that $\ker(f - \lambda \operatorname{Id}) = V$. Therefore $f - \lambda \operatorname{Id}$ is the zero map, and $f = \lambda \operatorname{Id}$. So the only G-equivariant endomorphisms are scalar maps. □

6 Orthonormality of Characters

6.1 Review: Schur's Lemma

Last time, we presented Schur's Lemma.

Recall that $\text{Hom}_G(\rho, \psi)$, which may also be written as $\text{Hom}_G(V, W)$, denotes the space of homomorphisms (or in other words, linear maps) from V to W which are G-equivariant, meaning that

$$\text{Hom}_G(\rho, \psi) = \{f : V \to W \mid f \text{ linear, and } f(\rho_g(v)) = \psi_g(f(v)) \text{ for all } g \in G, v \in V\}.$$

Theorem 6.1 (Schur's Lemma)
Suppose $\rho : G \to \text{GL}(V)$ and $\psi : G \to \text{GL}(W)$ are irreducible (and complex and finite-dimensional). Then

$$\dim\left(\text{Hom}_G(\rho, \psi)\right) = \begin{cases} 0 & \text{if } \rho \ncong \psi \\ 1 & \text{if } \rho \cong \psi. \end{cases}$$

In other words, the first statement means that if $\rho \ncong \psi$, then the only G-equivariant homomorphism $V \to W$ is the zero map; the second statement means that if $\rho \cong \psi$, then the only G-equivariant homomorphisms $V \to W$ are the scalar maps — or in other words, $\text{End}_G(\rho) = \mathbb{C} \cdot \text{Id}$.

The first statement is true for representations over *any* field of coefficients (and the same proof we saw last time works in the general case). However, the second statement is *not* true for arbitrary fields of coefficients — in the proof, we used the fact that eigenvectors must always exist, which isn't true in general (for example, $\mathbb{R}$ is not algebraically closed, so it's possible that the characteristic polynomial does not have roots, and there are no real eigenvalues). In particular, there are examples of real representations for which $\text{End}_G(\rho) \neq \mathbb{R}$:

Example 6.2
Consider the representation of $\mathbb{Z}/3\mathbb{Z}$ acting on $\mathbb{R}^2$, where $\overline{1}$ is mapped to the matrix

$$\begin{bmatrix} \cos 2\pi/3 & -\sin 2\pi/3 \\ \sin 2\pi/3 & \cos 2\pi/3 \end{bmatrix}$$

which corresponds to rotation by $2\pi/3$. In this case, $\text{End}_{\mathbb{Z}/3\mathbb{Z}}(\mathbb{R}^2)$ is $\mathbb{C}$, not $\mathbb{R}$.

Proof. As usual, multiplication by any scalar is in $\text{End}_{\mathbb{Z}_3}(\mathbb{R}^2)$; these elements correspond to $\mathbb{R}$ (in the case of $\mathbb{C}$, Schur's Lemma states that the *only* elements of $\text{End}_G(\rho)$ are multiplication by scalars).

But there are actually other possible endomorphisms — rotation by $\pi/2$ about the origin is *also* a G-equivariant endomorphism, since any two rotations about the origin must commute. We can think of this element as i; then taking linear combinations, all elements $a + bi$ (for real a and b) are in $\text{End}_{\mathbb{Z}/3\mathbb{Z}}(\mathbb{R}^2)$ (and it's possible to check that there's no others).

Composing two such endomorphisms corresponds to multiplying their corresponding complex numbers (since the endomorphism corresponding to z can be thought of as multiplication by z in the complex plane). So then this means $\text{End}_{\mathbb{Z}/3\mathbb{Z}}(\mathbb{R}^2) = \mathbb{C}$. □

We won't discuss this, but there also exists a real irreducible representation ρ of some group G for which $\text{End}_G(\rho) = \mathbb{H}$ (where $\mathbb{H}$ denotes the quaternions).

6.2 An Implication of Schur's Lemma

We've seen earlier that every representation $\rho : G \to \text{GL}(V)$ can be written as the sum of irreducible representations, with certain coefficients. Using Schur's Lemma, we can describe what these coefficients are:

Corollary 6.3
Let $\rho : G \to \mathrm{GL}(V)$ be a representation. Let $\rho_1, \ldots, \rho_n$ be the list of all irreducible representations of G (up to isomorphism). Then

$$\rho \cong \bigoplus_{i=1}^{n} \rho_i^{d_i}$$

where $d_k = \dim \mathrm{Hom}_G(\rho_k, \rho)$ for all k.

Proof. From Maschke's Theorem, we know that we can write

$$\rho \cong \bigoplus_{i=1}^{n} \rho_i^{d_i}$$

for *some* coefficients d_i, so then

$$\mathrm{Hom}_G(\rho_k, \rho) = \mathrm{Hom}_G\left(\rho_k, \bigoplus \rho_i^{d_i}\right) = \bigoplus_{i=1}^{n} \mathrm{Hom}_G(\rho_k, \rho_i)^{d_i} = \mathbb{C}^{d_k},$$

using the fact that by Schur's Lemma, $\mathrm{Hom}_G(\rho_k, \rho_i)$ is 0 if $i \neq k$ and $\mathbb{C}$ if $i = k$. This means

$$\dim \mathrm{Hom}_G(\rho_k, \rho) = d_k$$

for each k, as desired. □

Student Question. *Why could we write*

$$\mathrm{Hom}_G\left(\rho_k, \bigoplus \rho_i^{d_i}\right) = \bigoplus_{i=1}^{n} \mathrm{Hom}_G(\rho_k, \rho_i)^{d_i}?$$

Answer. *It's enough to see that* $\mathrm{Hom}_G(U, V \oplus W) = \mathrm{Hom}_G(U, V) \oplus \mathrm{Hom}_G(U, W)$.

First, by looking at matrices, it's possible to see that $\mathrm{Hom}_{\mathbb{C}}(U, V \oplus W) = \mathrm{Hom}_{\mathbb{C}}(U, V) \oplus \mathrm{Hom}_{\mathbb{C}}(U, W)$ *(where this denotes* all *linear maps, not just the G-equivariant ones) — if* $\dim U = m$, $\dim V = n_1$, *and* $\dim W = n_2$, *then a linear map $U \to V \oplus W$ is a $(n_1 + n_2) \times m$ matrix, which we can think of as a pair of a $n_1 \times m$ matrix and a $n_2 \times m$ matrix:*

$$\begin{bmatrix} * & * & * \\ * & * & * \\ * & * & * \end{bmatrix}.$$

Then such a map is compatible with the G-action if and only if each component is.

It's possible to think of this without matrices, as well. By definition, $V \oplus W$ is the space of pairs (v, w) with $v \in V$ and $w \in W$. So in order to describe a linear map from U to $V \oplus W$, for an element $u \in U$, we need to specify the first *coordinate of its image (corresponding to a linear map $U \to V$) and the* second *coordinate of its image (corresponding to a linear map $U \to W$). Then since G essentially acts separately on V and W (by the definition of a direct sum of representations), the map $U \to V \oplus W$ is G-equivariant if and only if the two individual maps $U \to V$ and $U \to W$ are.*

Student Question. *What exactly does it mean to have a list of irreducible representations up to isomorphism — what happens if there are two isomorphic representations, but they act on different vector spaces?*

Answer. *We can think of $\rho_1, \ldots, \rho_n$ as an abstract list of representations, without thinking about the subspaces being acted on. For example, when writing down the character table of S_3, we saw that there are three irreducible representations; this means every irreducible representation is isomorphic to one of them. We're using "up to isomorphism" in the same sense here.*

More generally, when we write $\rho = \bigoplus_{i=1}^{n} \rho_i^{d_i}$, when $d_k > 0$, this really means that ρ_k is isomorphic to a sub-representation of ρ.

6.3 Matrices and a New Representation

We'll now rewrite the concept of G-equivariance in terms of matrices — this will give a useful construction of a new representation, which we can apply Schur's Lemma to in order to prove the orthonormality of irreducible characters.

Choose a basis for V and W. Then if $n = \dim V$ and $m = \dim W$, we can write a linear map $V \to W$ as a $m \times n$ matrix, where the map sends $v \in V$ to $Av \in W$.

We can also write our representations ρ and ψ as the matrix representations $R : G \to \mathrm{GL}_n(\mathbb{C})$ and $S : G \to \mathrm{GL}_m(\mathbb{C})$. Then by rewriting the definition of G-equivariance (that $f(\rho_g(v)) = \psi_g(f(v))$ for all v and g) in terms of these matrices, we have that a matrix $A \in \mathrm{Mat}_{m\times n}(\mathbb{C})$ corresponds to a linear map $f \in \mathrm{Hom}_G(\rho, \psi)$ if and only if

$$AR_g = S_g A \text{ for all } g \in G.$$

(Our initial definition was written in terms of v, but we can think of it instead as an equality of the *linear maps* themselves — that the linear maps $f \circ \rho_g$ and $\psi_g \circ f$ are the same — which corresponds to an equality of matrices.) We can rewrite this condition as

$$A = S_g A R_g^{-1} \text{ for all } g \in G.$$

The key point is that we can get *another* representation of G from this expression (starting with our original representations ρ and ψ). Note that the space $M = \mathrm{Mat}_{m\times n}(\mathbb{C})$ is *itself* a $\mathbb{C}$-vector space — we can forget everything we know about how to multiply matrices, and just imagine adding them and multiplying by scalars, which makes $\mathrm{Mat}_{m\times n}(\mathbb{C})$ a mn-dimensional vector space over $\mathbb{C}$.

Lemma 6.4
There is a representation C of G acting on $\mathrm{Mat}_{m\times n}(\mathbb{C})$, where for each $g \in G$, g is sent to the matrix

$$C_g : A \mapsto S_g A R_g^{-1}.$$

In other words, if we think of $\mathrm{Mat}_{m\times n}(\mathbb{C})$ as a $\mathbb{C}$-vector space, then for each $g \in G$, the map $A \mapsto S_g A R_g^{-1}$ is a linear operator on this vector space. So we're sending g to the $mn \times mn$ matrix corresponding to that linear operator (which we denote by C_g).

Proof. It suffices to check that $C_{gh} = C_g C_h$ for all g and h. But for any $A \in \mathrm{Mat}_{m\times n}(\mathbb{C})$, we have

$$C_{gh}(A) = S_{gh} A R_{gh}^{-1} = S_g S_h A R_h^{-1} R_g^{-1} = C_g(C_h(A)).$$

So then $C_{gh} = C_g C_h$, as desired. □

In this new representation, $\mathrm{Hom}_G(\rho, \psi)$ is exactly the space of G-invariant vectors (note that here "vectors" means matrices of dimension $m \times n$, since the vector space our representation is acting on is actually the space of such matrices).

We can also describe this construction without using matrices — let $M = \mathrm{Hom}_{\mathbb{C}}(V, W)$ be the space of *all* linear maps $V \to W$ (which we thought of as $\mathrm{Mat}_{m\times n}(\mathbb{C})$ when writing down the construction in matrices). Then M is a vector space over $\mathbb{C}$. So we can define a representation γ acting on M, where for each $g \in G$, we send g to the linear map $\gamma_g : M \to M$ which sends $E \mapsto \psi_g E \rho_g^{-1}$ for all $E \in M$. As before, $\mathrm{Hom}_G(\rho, \psi)$ is exactly the space of G-invariant vectors in γ.

Student Question. *Why is $E \mapsto \psi_g E \rho_g^{-1}$ a linear map?*

Answer. *Thinking in terms of matrices, we want to see that $A \mapsto S_g A R_g^{-1}$ is a linear map. But this follows from the distributive property of matrix multiplication — for example, we have*

$$S_g(A + B)R_g^{-1} = S_g A R_g^{-1} + S_g B R_g^{-1}.$$

Proposition 6.5
For the representation γ described above, we have

$$\chi_\gamma = \chi_\psi \overline{\chi_\rho}.$$

The proposition quickly reduces to a statement about matrices:

Lemma 6.6
Let A and B be $n \times n$ and $m \times m$ matrices. Consider the linear map $A \otimes B$ from $\mathrm{Mat}_{m\times n}(\mathbb{C})$ to itself defined as

$$A \otimes B : E \mapsto BEA.$$

Then we have

$$\mathrm{Tr}(A \otimes B) = \mathrm{Tr}(A) \cdot \mathrm{Tr}(B).$$

Proof. The space of matrices has a basis consisting of the matrices E_{ij} which have a 1 in the ith row and jth column, and 0's everywhere else (for all $1 \le i \le m$ and $1 \le j \le n$).

But by straightforward computation, we can see that $BE_{ij}A$ has $b_{ii}a_{jj}$ in its ith row and jth column — this means $BE_{ij}A$ is $b_{ii}a_{jj}E_{ij}$, plus some entries corresponding to the other basis elements. So if we write out the matrix corresponding to $A \otimes B$ (in the basis formed by the E_{ij}), the diagonal entry corresponding to E_{ij} is $b_{ii}a_{jj}$. The trace of $A \otimes B$ is then the sum of the diagonal entries of this matrix, which is

$$\mathrm{Tr}(A \otimes B) = \sum_{i=1}^{m}\sum_{j=1}^{n} b_{ii}a_{jj} = \sum_{i=1}^{m} b_{ii} \sum_{j=1}^{n} a_{jj} = \mathrm{Tr}(B) \cdot \mathrm{Tr}(A). \qquad \square$$

Proof of Proposition 6.5. Using the above lemma, for all $g \in G$ we have

$$\chi_\gamma(g) = \mathrm{Tr}(\rho_{g^{-1}} \otimes \psi_g) = \mathrm{Tr}(\rho_{g^{-1}}) \cdot \mathrm{Tr}(\psi_g) = \chi_\rho(g^{-1}) \cdot \chi_\psi(g).$$

We saw earlier that $\chi_\rho(g^{-1}) = \overline{\chi_\rho(g)}$ — this follows from the fact that $\chi_\rho(g^{-1})$ is the sum of the eigenvalues of $\rho_{g^{-1}} = \rho_g^{-1}$, which are the inverses of the eigenvalues of ρ_g, and since all these eigenvalues have magnitude 1, their inverses are also their conjugates. So

$$\chi_\gamma(g) = \overline{\chi_\rho(g)} \cdot \chi_\psi(g)$$

for all $g \in G$, as desired. $\square$

6.4 Orthonormality of Characters

We have now developed the tools that we can use to prove one part of the main theorem stated earlier, the orthonormality of irreducible characters.

Proposition 6.7
The characters of the irreducible representations of G are orthonormal.

Proof. Let ρ and ψ be irreducible representations, acting on the spaces V and W. Then our Hermitian form is

$$\langle \chi_\psi, \chi_\rho \rangle = \frac{1}{|G|} \sum_{g\in G} \overline{\chi_\rho(g)}\chi_\psi(g),$$

so we want to check this expression is 0 if $\rho \ncong \psi$ and 1 if $\rho \cong \psi$.

But we saw a representation γ, acting on the space $M = \mathrm{Hom}_{\mathbb{C}}(V, W)$, whose character is exactly the expression inside the sum! So we can rewrite our sum as

$$\langle \chi_\psi, \chi_\rho \rangle = \frac{1}{|G|} \sum_{g\in G} \chi_\gamma(g) = \mathrm{Tr}\left(\frac{1}{|G|} \sum_{g\in G} \gamma_g \right).$$

(Here we used the fact that $\mathrm{Tr}(A+B) = \mathrm{Tr}\,A + \mathrm{Tr}\,B$ — our original sum calculates the traces of each γ_g *first*, and then averages them, but we can instead average the γ_g and *then* compute the trace).

Now let's consider what the linear operator on M given by $\sum_{g\in G}\gamma_g/|G|$ looks like; denote this operator by f.

Since f is an averaging operator (we're averaging over all $g \in G$), it must always output a G-invariant vector — similar to the averaging trick we saw earlier, for any $v \in M$ and any fixed $h \in G$, we have

$$\gamma_h(f(v)) = \gamma_h \cdot \frac{1}{|G|}\sum_{g\in G}\gamma_g v = \frac{1}{|G|}\sum_{g\in G}\gamma_h\gamma_g v = \frac{1}{|G|}\sum_{g\in G}\gamma_{hg} v = \frac{1}{|G|}\sum_{g\in G}\gamma_g v = f(v),$$

since if we fix h and let g range over all elements in G, then hg also ranges over all elements in G.

But we know what the G-invariant vectors are! Recall that vectors in the space M that γ acts on are actually homomorphisms $V \to W$, and by the way γ was defined, the vectors in M which are G-invariant are exactly the homomorphisms $V \to W$ which are G-equivariant — which are described by Schur's Lemma.

If $\rho \ncong \psi$, then by Schur's Lemma, then $\mathrm{Hom}_G(\rho,\psi)$ only contains the zero map, so the only G-invariant vector in M is the zero vector. So since our operator f sends every vector $v \in M$ to *some* G-invariant vector, it must actually send every vector v to 0. So f is the zero operator, and

$$\langle \chi_\psi, \chi_\rho\rangle = \mathrm{Tr}(f) = 0.$$

Now suppose $\rho \cong \psi$. Then by Schur's Lemma, $\mathrm{Hom}_G(\rho,\psi)$ only contains scalar maps; so there's only one G-invariant vector in M up to scaling (the identity matrix and its scalar multiples). Let u be such a (nonzero) invariant vector.

This means f must send every vector $v \in M$ to some scalar multiple of u (since f sends every v to some G-invariant vector, and the only G-invariant vectors are multiples of u). On the other hand, since we're averaging and u is already G-invariant, f must send u to itself — we have

$$\frac{1}{|G|}\sum_{g\in G}\gamma_g u = \frac{1}{|G|}\sum_{g\in G} u = u.$$

Now choose a basis for M whose first element is u. Then in this basis, the matrix for f is of the form

$$\begin{bmatrix} 1 & * & * & \cdots & * \\ 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 0 \end{bmatrix},$$

which has trace 1. So in this case,

$$\langle \chi_\psi, \chi_\rho\rangle = \mathrm{Tr}(f) = 1. \qquad \square$$

7 Proof of the Main Theorem

7.1 Review: Orthonormality of Characters

Last time, we proved the orthonormality of characters of irreducible representations — if we have a finite group G and $\rho_1, \ldots, \rho_n$ is the full list of irreducible representations of G up to isomorphism, then if we use χ_i to denote χ_{ρ_i}, we have

$$\langle \chi_i, \chi_j \rangle = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{otherwise.} \end{cases}$$

In order to prove this, we interpreted $\langle \chi_i, \chi_j \rangle$ as the trace of an averaging operator, acting on the space of linear maps $\mathrm{Mat}_{m\times n}(\mathbb{C}) = \mathrm{Hom}_{\mathbb{C}}(V_i, V_j)$, where V_i and V_j are the vector spaces that ρ_i and ρ_j act on.

From orthonormality, we immediately get a few corollaries:

Corollary 7.1
Any representation $\rho : G \to \mathrm{GL}(V)$ can be split as a sum of irreducibles as

$$\rho \cong \bigoplus \rho_i^{n_i},$$

where for each k,

$$n_k = \langle \chi_\rho, \chi_k \rangle.$$

Recall that last class, we saw a different formula for the n_i, which was quite abstract — it involved the dimensions of $\mathrm{Hom}_G(\rho_k, \rho)$. In contrast, this formula is quite concrete — it's easy to calculate the pairings $\langle \chi_\rho, \chi_k \rangle$.

Proof. We know ρ can be written in this form for some coefficients n_i, by Maschke's Theorem. But then

$$\chi_\rho = \sum n_i \chi_i,$$

so by linearity we have

$$\langle \chi_\rho, \chi_k \rangle = \sum n_i \langle \chi_i, \chi_k \rangle = n_k,$$

since orthonormality implies that $\langle \chi_i, \chi_k \rangle$ is 0 for $i \neq k$ and 1 for $i = k$. □

Using this, we can get a few concrete results:

Example 7.2
The dimension of the space of invariant vectors in ρ is

$$\frac{1}{|G|} \sum_{g \in G} \chi_\rho(g).$$

Proof. This dimension is the multiplicity of the trivial representation χ_1 when we decompose ρ into a sum of irreducibles, which is n_1. This is because ρ acts trivially on the space of invariant vectors, by definition, and so each basis vector corresponds to one copy of the trivial representation. But the character of the trivial representation is $\chi_1(g) = 1$ for all g, so using Corollary 7.1, we have

$$n_1 = \langle \chi_\rho, \chi_1 \rangle = \frac{1}{|G|} \sum_{g \in G} \chi_\rho(g).$$

□

Corollary 7.3
If $\rho = \bigoplus \rho_i^{n_i}$ as before, then

$$\langle \chi_\rho, \chi_\rho \rangle = \sum n_i^2.$$

In particular, ρ is irreducible if and only if $\langle \chi_\rho, \chi_\rho \rangle = 1$.

Proof. Using linearity, we can expand

$$\langle \chi_\rho, \chi_\rho \rangle = \sum_i \sum_j n_i n_j \langle \chi_i, \chi_j \rangle = \sum n_i^2,$$

since again by orthonormality, $\langle \chi_i, \chi_j \rangle$ is 0 when $i \neq j$ and 1 when $i = j$. (Intuitively, if we have any pairing and we take a basis of vectors which are orthonormal under that pairing — here, the χ_i — then it becomes the usual pairing on $\mathbb{C}^n$.)

The second statement is then clear because the n_i are nonnegative integers, so the only way for their sum of squares to be 1 is if one is 1 and the rest are 0. □

7.2 The Regular Representation

We've already proved part of the main theorem — we've shown that the characters χ_i are orthonormal, which means they are linearly independent. But to show that they form a basis for the space of class functions, we also need to show that they span that space. To do this — and to prove the sum of squares formula stated earlier as well — we'll introduce the regular representation.

If G acts on a finite set X of n elements, then we can form a matrix representation of G of dimension n, where G acts by permutation matrices — we index the basis vectors by elements of X, and for each $g \in G$, we map g to the permutation matrix that describes how g acts on X.

Example 7.4
As we've seen earlier, S_n acts on the set $\{1, 2, \ldots, n\}$. This gives a n-dimensional representation of S_n — the permutation representation, where each permutation is mapped to its corresponding permutation matrix.

Given a group G, there can be many interesting sets X that it acts upon. But there's one set that we automatically always have; namely, G itself. Every group G acts on itself by left multiplication, where an element $g \in G$ sends $h \mapsto gh$. We can use this to form a representation of G of dimension $|G|$:

Definition 7.5
Let V be a vector space with basis $\{v_h\}$ indexed by elements $h \in G$. Then the **regular representation** of G is the representation $\rho : G \to \mathrm{GL}(V)$ such that for all $g \in G$, ρ_g is the linear operator on V sending $v_h \mapsto v_{gh}$ for all $h \in G$.

So in the regular representation, each $g \in G$ is sent to the permutation matrix of how left multiplication by g permutes the elements of G.[‖]

Example 7.6
In the group $\mathbb{Z}/3\mathbb{Z}$, operating (in this case the group operation is addition) on the left by $\overline{1}$ corresponds to the permutation $\overline{0} \mapsto \overline{1}$, $\overline{1} \mapsto \overline{2}$, $\overline{2} \mapsto \overline{0}$. So in its regular representation, $\overline{1}$ acts by the permutation matrix

$$\begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}.$$

[‖] For left multiplication, $\rho_{g_1}\rho_{g_2} = \rho_{g_1 g_2}$, so ρ is a homomorphism. Using right multiplication rather than left multiplication, which would send $\rho'_g : v_h \longmapsto v_{hg}$, would be a anti-homomorphism rather than a homomorphism. That is, $\rho'_{g_1}\rho'_{g_2}(v_h) = v_{hg_2g_1} = \rho'_{g_2g_1}(v_h)$.

Example 7.7
In S_3, left multiplication by (12) swaps (1) and (12), swaps (23) and (123), and (13) and (132). So in the regular representation, (12) acts by the block diagonal matrix

$$\begin{bmatrix} 0 & 1 & & & & \\ 1 & 0 & & & & \\ & & 0 & 1 & & \\ & & 1 & 0 & & \\ & & & & 0 & 1 \\ & & & & 1 & 0 \end{bmatrix},$$

where the rows and columns are indexed by elements of S_3 in the order (1), (12), (23), (123), (13), (132).

It's clear from the definition that the regular representation has exactly one invariant vector up to scaling, the sum of all basis elements. This is because if $\sum a_h v_h$ is an invariant vector, then we must have

$$\sum a_h v_h = \rho_g \left(\sum a_h v_h \right) = \sum a_h v_{gh}$$

for all $g \in G$, which implies that $a_h = a_{gh}$ for all g and h, and therefore all a_h are equal. In particular, the regular representation is not irreducible unless G is trivial (the regular representation has an invariant vector, so it must have the trivial representation in its decomposition).

Note 7.8
It's also possible to think of elements of V as $\mathbb{C}$-valued functions on G — instead of thinking of them as linear combinations of abstract basis vectors, we can think of $\sum a_h v_h$ as the function mapping $h \mapsto a_h$ for all $h \in G$.

We'll now use ρ to denote the regular representation.

Guiding Question
What can we say about the character of ρ and its decomposition into irreducibles?

This turns out to have a simple answer.

Proposition 7.9
The character of the regular representation is

$$\chi_\rho(g) = \begin{cases} |G| & \text{if } g = 1 \\ 0 & \text{otherwise.} \end{cases}$$

Proof. The trace of a permutation matrix is its number of fixed points (a permutation matrix consists only of 0's and 1's, and 1's on its diagonal correspond to fixed points).

It's clear that the permutation corresponding to 1 fixes all elements of G (since multiplication by 1 doesn't change any element), so $\chi_\rho(1) = |G|$. Meanwhile, for all $g \in G$ other than 1, the permutation corresponding to g has no fixed points — if h were a fixed point, then we would have $h = gh$, which implies $g = 1$. So then $\chi_\rho(g) = 0$ for all $g \neq 1$. □

Using this, we can decompose ρ into a sum of irreducibles pretty easily, using Corollary 7.1 — since our character has such a simple form, it's not hard to compute its pairing with anything.

Proposition 7.10
The regular representation decomposes into irreducibles as

$$\rho \cong \bigoplus \rho_i^{d_i},$$

where d_i denotes the dimension of ρ_i.

Proof. By Corollary 7.1, we know $\rho = \bigoplus \rho_i^{n_i}$, where

$$n_i = \langle \chi_i, \chi_\rho \rangle = \frac{1}{|G|} \sum_{g \in G} \chi_i(g)\overline{\chi_\rho(g)} = \frac{1}{|G|}\chi_i(1) \cdot |G| = \chi_i(1)$$

(all terms involving $g \neq 1$ disappear, since $\chi_\rho(g) = 0$). But $\rho_i(1)$ is the identity matrix of dimension d_i, which has trace d_i — so $\chi_i(1) = d_i$, which means $n_i = d_i$ as well. □

Example 7.11
In an abelian group, all dimensions of irreducible representations are 1, so $\rho \cong \rho_1 \oplus \cdots \oplus \rho_n$.

From this, we can immediately deduce the sum of squares formula stated in the main theorem.

Proposition 7.12
If the irreducible representations of G have dimensions $d_1, \ldots, d_n$, then we have

$$|G| = d_1^2 + \cdots + d_n^2.$$

Proof. We can compare the dimensions on the two sides of Proposition 7.10. On the left-hand side, we have $\dim(\rho) = |G|$. Meanwhile, on the right-hand side, for each i we have d_i copies of ρ_i, which itself has dimension d_i — this contributes d_i^2 to the dimension (since when we take the direct sum of two representations, we add their dimension). So then

$$|G| = \dim(\rho) = \dim\left(\bigoplus \rho_i^{d_i}\right) = \sum d_i^2,$$

which gives the desired equality. □

7.3 Span of Irreducible Characters

Finally, we'll again use the regular representation to prove that the characters of irreducible representations span the space of class functions. We know that these characters are linearly independent (since they're orthonormal), so it will then follow that they form a *basis* for the space of class functions (the first statement of our main theorem).

In order to prove this, we'll show that any class function can be written in the following form:

Proposition 7.13
For any class function f, we have

$$f = \sum \langle f, \chi_i \rangle \chi_i.$$

Note that we already know this statement is true when f is the character of any representation. More generally, if we can write $f = \sum n_i \chi_i$, then by orthonormality we know that $n_i = \langle f, \chi_i \rangle$ for each i. So Proposition 7.13 essentially tells us that this formula really works for *all* class functions. Note that the coefficients $\langle f, \chi_i \rangle$ are all in $\mathbb{C}$, so Proposition 7.13 implies that the χ_i span the space of class functions — this proposition would follow immediately if we already knew that the χ_i span the space of class functions, but we can actually use it to prove that statement instead.

Proposition 7.13 is equivalent to the following statement:

Proposition 7.14
If f is a class function and $\langle f, \chi_k \rangle = 0$ for all k, then f is the zero function.

First, to make this equivalence between Proposition 7.13 and Proposition 7.14 more explicit, we'll write out the details of how the second implies the first. (For the other direction, the first implies the second because if the χ_i do span the space of class functions and f has zero pairing with each one, then f also has zero pairing with *itself*. But the space of class functions under $\langle -, - \rangle$ is a Hermitian space, meaning that $\langle f, f \rangle > 0$ unless $f = 0$.)

Proof of Proposition 7.13. Given any class function, we can define $f^* = \sum \langle f, \chi_i \rangle \chi_i$. But then $f - f^*$ has zero pairing with each χ_k, since

$$\begin{aligned} \langle f - f^*, \chi_k \rangle &= \langle f, \chi_k \rangle - \left\langle \sum \langle f, \chi_i \rangle \chi_i, \chi_k \right\rangle \\ &= \langle f, \chi_k \rangle - \sum \langle f, \chi_i \rangle \langle \chi_i, \chi_k \rangle \\ &= \langle f, \chi_k \rangle - \langle f, \chi_k \rangle \\ &= 0 \end{aligned}$$

by orthonormality. (Intuitively, since f^* is a linear combination of the χ_i, the pairing of f^* with each χ_k is the coefficient of χ_k, which we *constructed* to be the same as the pairing of f with χ_k.) But using Proposition 7.14, the only class function that has zero pairing with every χ_k is the zero function, so we must have $f - f^* = 0$, and therefore $f = f^*$. □

In order to prove Proposition 7.14, we'll first introduce a bit of convenient notation:

Definition 7.15
For a function $f : G \to \mathbb{C}$ (not necessarily a class function) and a representation $\rho : G \to \mathrm{GL}(V)$, define

$$\rho(f) = \sum_{g \in G} f(g) \rho_g.$$

Note that $\rho(f)$ is in $\mathrm{End}(V)$ (or equivalently, in $\mathrm{Mat}_{n \times n}(\mathbb{C})$ if we write down a basis and use matrix representations), since each ρ_g is in $\mathrm{End}(V)$ and the coefficients $f(g)$ are scalars. We can think of this definition as extending the definition of ρ to *linear combinations* of group elements (in the obvious way), instead of just the group elements themselves — in this interpretation, the function f stores the coefficient of each group element in the linear combination.

We can make a few observations about this definition:

Lemma 7.16
For any representation ρ and function f, we have $\operatorname{Tr} \rho(f) = \langle \chi_\rho, \overline{f} \rangle$.

Proof. This follows directly from the linearity of trace — plugging in the definitions and using linearity, we have

$$\langle \chi_\rho, \overline{f} \rangle = \sum_{g \in G} (\chi_\rho(g) f(g)) = \sum_{g \in G} (f(g) \operatorname{Tr} \rho_g) = \operatorname{Tr} \left(\sum_{g \in G} f(g) \rho_g \right) = \operatorname{Tr} \rho(f).$$

(The reason we have $\overline{f}$ and not f here is because we conjugate the second element of the pairing.) □

Lemma 7.17
If f is a class function, then $\rho(f) \in \mathrm{End}_G(\rho)$.

Proof. Recall that class functions are functions which are invariant under conjugation, meaning that for each conjugacy class, they take the same value on all its elements.

But this means $\rho(f)$ must be invariant under conjugation as well — more explicitly, for any $g \in G$, we have

$$\rho_g \rho(f) \rho_g^{-1} = \sum_{h \in G} f(h) \rho_g \rho_h \rho_g^{-1} = \sum_{h \in G} f(h) \rho_{ghg^{-1}} = \sum_{h \in G} f(h) \rho_h = \rho(f),$$

since the new sum just permutes elements within each conjugacy class, and $f(ghg^{-1}) = f(h)$ for all h. We can rearrange this to

$$\rho_g \rho(f) = \rho(f) \rho_g$$

for all $g \in G$, so $\rho(f)$ is indeed G-equivariant. □

Using these observations, we can now prove our claim, that the only class function f which has zero pairing with every χ_i is the zero function.

Proof of Proposition 7.14. For each i, we are given that $\langle \chi_i, f \rangle = 0$, so by Lemma 7.16 we also have

$$\operatorname{Tr} \rho_i(\overline{f}) = 0.$$

But we also know that $\rho_i(\overline{f}) \in \operatorname{End}_G(\rho_i)$. And by Schur's Lemma (since ρ_i is irreducible), the only elements of $\operatorname{End}_G(\rho_i)$ are scalar matrices! So then $\rho_i(\overline{f})$ is a scalar matrix with trace 0; this means it's the zero matrix.

So now we know that $\rho_i(\overline{f})$ is the zero matrix for all irreducible representations. But by Maschke's Theorem, every representation is the direct sum of irreducible representations — so this means $\rho(\overline{f})$ is the zero matrix for *any* representation ρ. (If we write ρ as a direct sum of irreducibles ρ_i, then $\rho(\overline{f})$ is the corresponding direct sum of the matrices $\rho_i(\overline{f})$, and a direct sum of zero matrices is also the zero matrix.)

Now we'd like to use this to conclude that f itself is 0. To do so, we can take ρ to be the regular representation. It turns out that we can essentially just read off the function by looking at its action in the regular representation — in particular, f acts on the basis vector v_1 as

$$\rho(f)(v_1) = \sum_g f(g) \rho_g(v_1) = \sum_g f(g) v_g.$$

Since $\rho(f)$ is the zero map, then the right-hand side must be zero as well; so $f(g) = 0$ for all $g \in G$. □

This concludes our proof of the Main Theorem — we have proven that the irreducible characters χ_i are orthonormal and form a basis for the space of class functions, and the dimensions of the irreducible representations d_i satisfy $\sum d_i^2 = |G|$. The only remaining statement which we have not proven is that d_i divides $|G|$ for each i. We will not discuss this proof in class, but it is posted on Canvas; in these notes, a writeup of this proof is included in the appendix.

7.4 Generalizations to Compact Groups

We've worked with representations of *finite* groups, but much of this generalizes to *compact* subgroups of $\mathrm{GL}_n(\mathbb{C})$, such as $\mathrm{U}(n)$ and $\mathrm{O}(n)$. In this case, we consider *continuous* representations.

Example 7.18
Consider $\mathrm{U}(1) = \{z \in \mathbb{C}^\times \mid |z| = 1\}$. The irreducible representations are all one-dimensional (as in the finite case, this has to do with the fact that the group is abelian), and are indexed by integers, where $\rho_n : z \mapsto z^n$.

All functions $f : \mathrm{U}(1) \to \mathbb{C}$ are class functions, and we can think of such a function as a function $f : \mathbb{R} \to \mathbb{C}$ which is 2π-periodic (by thinking of $\mathrm{U}(1)$ as the unit circle). Then if we try to decompose such a function f in the same way as Proposition 7.13, we'll get its **Fourier series** — under some reasonable conditions on f (in particular, here we require it to be continuous), the expression $\sum \chi \langle f, \chi \rangle$ ends up being the Fourier series of f.

8 Rings

We'll now discuss the second main topic of this course: rings.

Guiding Question
Groups have only one operation, but lots of familiar sets $(\mathbb{Z}, \mathbb{R}, \mathbb{C}, \mathbb{Q})$ have two operations: addition and multiplication. How can we generalize the idea of sets with two operations, rather than one?

8.1 What is a Ring?

Informally, a ring is a set of elements which can be added and multiplied, so that the natural properties we would expect of addition and multiplication all hold.

Example 8.1
The familiar sets $\mathbb{Z}, \mathbb{Q}, \mathbb{R}$, and $\mathbb{C}$ are rings.

Example 8.2
There are various rings between $\mathbb{Z}$ and $\mathbb{Q}$: for example,

$$\mathbb{Z}[1/2] \coloneqq \left\{a/2^k \mid a, k \in \mathbb{Z}\right\}$$

is a ring. Similarly, the **Gaussian integers**

$$\mathbb{Z}[i] = \{a + bi \mid a, b \in \mathbb{Z}\}$$

and

$$\mathbb{Q}[\sqrt{2}] = \{a + b\sqrt{2} \mid a, b \in \mathbb{Q}\}$$

are rings. So is

$$\mathbb{Z}[\zeta] = \left\{a_0 + a_1\zeta + \cdots + a_{n-1}\zeta^{n-1} \mid a_i \in \mathbb{Z}\right\},$$

where ζ is some nth root of unity.

We can now state the formal definition of a ring:

Definition 8.3
A **ring** R is a set with two binary operations $R \times R \to R$, written as $+$ and $\cdot$ (and called addition and multiplication), which satisfy the following axioms:

1. $(R, +)$ is an abelian group: addition is associative and commutative, there is an additive identity 0_R such that $0_R + a = a$ for all $a \in R$, and every element has an additive inverse.
2. Multiplication is also associative and commutative, and there is a multiplicative identity 1_R. (In other words, under multiplication, R is a *semigroup* — a group without the condition that every element has an inverse.)
3. Addition and multiplication satisfy *distributivity*: for all $a, b, c \in R$, we have

$$a(b + c) = ab + ac.$$

In this class, we'll use "ring" in this sense; but usually a ring satisfying this definition is called a *commutative unital ring*. In defining a general ring, one can drop the requirement of commutativity of multiplication, or the existence of 1_R. If every condition holds except for $ab = ba$ — and we add distributivity in the other direction as well, meaning $(b + c)a = ba + ca$ — then R is called a **noncommutative ring**.

Example 8.4 (Matrices)
The ring of matrices $\mathrm{Mat}_{n\times n}(\mathbb{C})$ is a noncommutative ring, since matrix multiplication is noncommutative (and all the other axioms are satisfied). Similarly, $\mathrm{End}(V)$ for a vector space V is a noncommutative ring.

Example 8.5 (Group Ring)
If G is a group, then take the vector space with basis vectors v_g for $g \in G$ (note that this is the vector space acted on by the regular representation). Clearly, addition is already defined; and multiplication can be defined in the natural way, as

$$v_g v_h = v_{gh}.$$

Then this gives the **group ring**, which is noncommutative if G is nonabelian.

From now on, all rings will be commutative as in the definition, unless stated otherwise.

8.2 Zero and Inverses

The following proposition confirms a property that we would like rings to have.

Proposition 8.6
In any ring R, $0_R \cdot a = 0_R$ for all $a \in R$.

Proof. Pick some $x \in R$. Then, since $0_R + x = x$, we have

$$xa = (0_R + x)a = 0_R a + xa.$$

We can cancel out xa (since R is an abelian group under addition), so $0_R = 0_R a$. □

Corollary 8.7
The additive identity 0_R cannot have a multiplicative inverse unless $0_R = 1_R$. (In other words, division by 0 is only possible when $0 = 1$.)

The axioms do not require that $0_R \neq 1_R$. But if $0_R = 1_R$, then for any $x \in R$,

$$x = x \cdot 1_R = x \cdot 0_R = 0_R.$$

So R must be a one-element ring; there's only one binary operation on a set with one element, which does satisfy the axioms. This ring is called the **zero ring**; it is a legitimate but trivial example. In all other cases, $0_R \neq 1_R$, as we would expect.

Definition 8.8
A (nonzero) ring where every nonzero element has a multiplicative inverse is called a **field**.

Note that by definition, the zero ring is not a field.

Example 8.9
The familiar sets $\mathbb{Q}$, $\mathbb{R}$, and $\mathbb{C}$ are fields.

Example 8.10
The integers $\mathbb{Z}$ do not form a field, since most numbers do not have inverses. For the same reason, the Gaussian integers $\mathbb{Z}[i]$ are not a field either. For example, 2 is not invertible in either ring.

Example 8.11
The integers modulo n, denoted $\mathbb{Z}/n\mathbb{Z}$, form a field if and only if n is prime (since in general, a is invertible mod n if and only if a and n are coprime).

More examples of rings (which are not fields) come from functions:

Example 8.12
The set $\mathbb{C}[x]$, consisting of polynomials in one variable with complex coefficients, is a ring. The set $C^\infty(\mathbb{R})$, consisting of real-valued functions which are continuous and infinitely differentiable, is also a ring.

In some sense, you can think of fields as a generalization of numbers, and more general rings as generalizations of functions.

8.3 Homomorphisms

When studying groups, one of the most powerful tools comes from thinking about mappings between groups. In particular, homomorphisms, which are functions that respect the group operation, and their kernels/images provide lots and lots of information about groups.

Guiding Question
How can we formalize the idea of "functions between rings that behave nicely with respect to addition and multiplication"?

Similarly to the path we took when studying groups, we can now define a ring homomorphism.

Definition 8.13
A ring **homomorphism** from R to S is a map $\varphi : R \to S$ such that:

1. $\varphi(a+b) = \varphi(a) + \varphi(b)$ for all $a, b \in R$;
2. $\varphi(ab) = \varphi(a)\varphi(b)$ for all $a, b \in R$;
3. $\varphi(1_R) = 1_S$.

Example 8.14
The mapping $\mathbb{Z} \to \mathbb{Z}/n\mathbb{Z}$ given by $a \mapsto \overline{a}$ (where $\overline{a}$ denotes a mod n) is a ring homomorphism.

Why is it not necessary to require that $\varphi(0_R) = 0_S$? In fact, it is implied by the other axioms, because additive inverses exist!

Proposition 8.15
For a ring homomorphism $\varphi : R \to S$, it must be the case that $\varphi(0_R) = 0_S$.

Proof. Using the first property, $\varphi(a) = \varphi(a + 0_R) = \varphi(a) + \varphi(0_R)$, and so adding $-\varphi(a)$ to both sides gives $0_S = \varphi(0_R)$. □

On the other hand, because there are not necessarily multiplicative inverses, the property $\varphi(1_R) = 1_S$ must be explicitly written. In fact, there are examples of maps compatible with $+$ and $\cdot$ (meaning they satisfy the first two properties) that have $\varphi(1_R) \neq 1_S$.

Example 8.16
Let R be the zero ring, and S be any nonzero ring (for example, $\mathbb{Z}$). Then take the map $0_R \mapsto 0_S$. This is compatible with the additive and multiplicative structure of the two rings; but since $1_S \neq 0_S$, it is not a ring homomorphism. (There exist less trivial examples as well.)

If $\varphi : R \to S$ is one-to-one and onto (that is, φ is a bijection), then it is possible to check (as in the case of group homomorphisms) that the inverse bijection is also a homomorphism.

Definition 8.17
A bijective homomorphism is called an **isomorphism**.

If φ is one-to-one but *not* onto, then φ is an isomorphism from R to its image, so it identifies R with a **subring** of S:

Definition 8.18
A **subring** S of R is a subset which is closed under addition, taking additive inverses, and multiplication, and contains 1_R.

Meanwhile, if φ is *not* one-to-one, then we can think about its kernel

$$\ker(\varphi) = \{a \in R : \varphi(a) = 0\}.$$

(This is the same definition as we saw with groups.) If φ is not one-to-one, then $\ker(\varphi)$ is nontrivial.

Guiding Question
What information can be gained from ring homomorphisms with nontrivial kernel?

The kernel must be an additive subgroup of S. But it turns out that it must also be compatible with multiplication in some ways; this brings us to the concept of *ideals*.

8.4 Ideals

Definition 8.19 (Ideal)
An **ideal** of R is a subset $I \subset R$ such that I is an additive subgroup of R, and for any $x \in I$ and $a \in R$, we have $ax \in I$.

So an ideal isn't just closed under multiplication, it's in fact closed under multiplication by *any* element in R.

The kernel of a ring homomorphism is necessarily an ideal — if $\varphi(x) = 0$, then for any $a \in R$, we have

$$\varphi(ax) = \varphi(a)\varphi(x) = 0.$$

Example 8.20
For any $n \in \mathbb{Z}$, the subset $n\mathbb{Z}$ (consisting of multiples of n) is an ideal. More generally, if R is any ring and a an element of R, then $aR = \{ax \mid x \in R\}$ is an ideal.

In fact, this is an important example of an ideal, as we'll see later, so it has a name:

Definition 8.21 (Principal Ideal)
For an element $a \in R$, the ideal aR is called a **principal ideal**, and denoted (a).

Any additive subgroup of $\mathbb{Z}$ is cyclic, i.e., of the form $n\mathbb{Z}$. This means every ideal of $\mathbb{Z}$ is principal (since every ideal is an additive subgroup). However, in a general ring, not every ideal will be principal.

More generally, we can consider picking *several* elements:

Definition 8.22
For elements $a_1, \ldots, a_n \in R$, the set of linear combinations

$$\left\{\sum a_i x_i \mid x_i \in R\right\}$$

is an ideal of R, denoted as $(a_1, \ldots, a_n)$; this is called the ideal **generated** by $a_1, \ldots, a_n$.

Note that $(a_1, \ldots, a_n)$ is the smallest ideal containing all of $a_1, \ldots, a_n$ (as it is an ideal, while all elements $\sum a_i x_i$ must be in the ideal by the axioms).

Ideals in rings are in some sense analogous to normal subgroups in groups (in particular, both arise as the kernel of a homomorphism), and we can take quotients in a similar way as well:

Proposition 8.23 (Quotient Ring)
Let R be a ring and $I \subset R$ an ideal. Since I is a normal subgroup of R under addition (as both are abelian groups), we can construct the quotient R/I of additive groups. Then R/I is in fact a ring (with multiplication defined in the natural way — the product of the cosets corresponding to x and y is the coset corresponding to xy), called the **quotient ring**.

Proof. For each $x \in R$, we use $\overline{x}$ to denote the coset $x + I$.

We first need to check that multiplication is well-defined, meaning that if $\overline{x}$ is represented by two elements x_1 and x_2, then using either representative to calculate $\overline{x} \cdot \overline{y}$ will give us the same result. But we have $x_1 - x_2 = a$ for some $a \in I$, which means

$$x_1 y - x_2 y = ay \in I$$

as well (since I is closed under multiplication by any element $y \in R$), so $x_1 y$ and $x_2 y$ are also in the same coset. So the product of two cosets doesn't depend on our choice of representatives, which means multiplication really is well-defined.

As in the case of groups, once we know that multiplication is well-defined, it is easy to check that all the ring axioms are satisfied by R/I. □

Without going into detail, replacing "group" by "ring" and "normal subgroup" by "ideal," the story for rings is extremely similar to that for groups. For example, if φ is a homomorphism, then there is an isomorphism

$$R/\ker(\varphi) \cong \operatorname{im}(\varphi).$$

Additionally, for an ideal $I \subset R$, there is a bijection between ideals in R/I and ideals in R containing I (this is essentially the Correspondence Theorem for rings).

Example 8.24
For $R = \mathbb{Z}$ and $I = n\mathbb{Z}$, we have $R/I = \mathbb{Z}/n\mathbb{Z}$ (where $\mathbb{Z}/n\mathbb{Z}$ denotes the integers mod n) as rings, not just groups. This is essentially the fact that multiplication of residues mod n is well-defined, not just addition. (This is where the notation $\mathbb{Z}/n\mathbb{Z}$ comes from — to obtain the integers mod n, we're quotienting out the ring of integers by the ideal $n\mathbb{Z}$.)

9 Building New Rings

9.1 Review

Last time, we introduced the idea of rings (where we can both add and multiply) and their ideals. As an aside, the term *ring* first appeared at the end of the 19th century in works by Hilbert; it's unclear why exactly he chose this term, but in the original German, the word essentially means a group of things coming together. Meanwhile, the term *ideal* comes from "ideal divisors" (we'll see later what this means).

As we saw earlier, ideals in rings work similarly to normal subgroups in groups. For an ideal $I \subset R$, we can construct the *quotient ring* R/I, and analogous versions of the theorems from group theory apply here as well. However, note that unlike the case of normal subgroups in groups, an ideal is generally not a subring. This is because it doesn't usually have 1_R — for example, $2\mathbb{Z} \subset \mathbb{Z}$ clearly doesn't contain 1. In fact, if the ideal *did* contain 1_R, it would have to be the entire ring. However, an ideal *does* satisfy the other axioms.

9.2 Product Rings

Now that we have an understanding of what rings are, we can think about how to construct them.

Guiding Question
How can we build new rings out of rings that we already have?

One construction is taking the *product* of two rings:

Definition 9.1
Let R and S be two rings. The **product ring**, denoted $R \times S$, is the set of pairs (r, s) with $r \in R$ and $s \in S$ (the Cartesian product of the two sets), along with componentwise addition and multiplication.

It's clear that the ring axioms clearly hold, with $1_{R\times S} = (1_R, 1_S)$.

Given a product ring $R \times S$, we have the **projection homomorphism** $R \times S \to R$ given by $(r, s) \mapsto r$. The kernel of this homomorphism is the set $(0, s)$ for $s \in S$. In other words, we could describe this kernel as the ideal of $R \times S$ generated by $(0, 1_S)$, since $(0, s) = (0, s) \cdot (0, 1_S)$.

Guiding Question
Given a ring Q, how can we recognize whether Q is isomorphic to $R \times S$ for some nonzero R and S?

(We ask this question about *nonzero* R and S because every ring Q is trivially the product of itself and the zero ring.)

First, if $Q = R \times S$, then we can consider the two elements $e_1 = (1_R, 0)$ and $e_2 = (0, 1_S)$. These are not units, but they are somewhat similar to units — in particular, they are *idempotent*.

Definition 9.2
An element e is **idempotent** if $e^2 = e$, or equivalently if $e(1 - e) = 0$.

Note that if e is idempotent, so is $1 - e$.

In our situation, if we have a product of two rings, then we have two idempotent elements e_1 and e_2 (which are neither 0 nor 1). We'll soon see that the converse is true as well. The intuition here may be more familiar in a linear algebra setting — suppose we have a vector space V and an idempotent matrix E, meaning that $E^2 = E$. Then its only eigenvalues are 0 and 1. So if we let V_1 and V_0 be the corresponding eigenspaces, then we can split $V = V_1 \oplus V_0$. So an idempotent matrix can be used to split the vector space into two smaller ones; it turns out it's possible to do something similar for rings.

Note that in any ring, 0 and 1 are both idempotent. In a *field*, there are no other idempotents — if $e(1 - e) = 0$, then e or $1 - e$ must be 0 — but this is not true in general, as we've just seen that there are other idempotents in $R \times S$.

Proposition 9.3
A ring Q is isomorphic to a product of rings $R \times S$ if and only if Q contains an idempotent other than 0 and 1.

Proof. We've already seen that $R \times S$ contains the idempotent elements $(1_R, 0_S)$ and $(0_R, 1_S)$, so it suffices to prove the other direction.

Given an idempotent $e \in Q$, take $R = eQ$ and $S = (1-e)Q$. Note that R (and similarly S) is a ring — it's clearly an abelian group under addition, and we can multiply the same way as in Q since

$$eq_1 \cdot eq_2 = e^2 q_1 q_2 = eq_1 q_2.$$

In particular, $e = 1_R$ (it's not a unit in the entire ring Q, but it *is* a unit in the smaller ring R). Similarly, $1 - e = 1_S$.

Then to check that $Q \cong R \times S$, for any $x \in Q$, we can write

$$x = ex + (1-e)x.$$

So then there is an isomorphism $Q \to R \times S$, given by $x \mapsto (ex, (1-e)x)$; its inverse map is given by $(r, s) \mapsto r + s$. (It's possible to explicitly check that these maps are inverses; the point is that $e(1-e) = 0$, so "mixed" terms disappear when we multiply.) □

Note 9.4
Note that R is a ring and is a subset of Q, but R is *not* a subring of Q in our terminology, since it doesn't contain 1_Q.

Similarly, the map $R \to R \times S$ sending $r \mapsto (r, 0)$ is compatible with addition and multiplication; but it is not a homomorphism in our terminology, since it does not send 1_R to $1_{R \times S}$.

Student Question. *In our construction, we took $R = eQ$ and $S = (1-e)Q$ for an idempotent e. But if Q was a field, wouldn't this require e to be 0 or 1?*

Answer. *Yes — this shows that a field cannot be written as a product of two rings (in a nontrivial way).*

Furthermore, it is possible to define the product of any collection of rings, finite or infinite, in the same way (with the operations performed componentwise).

9.3 Adjoining Elements to a Ring

A different way of creating new rings, which is quite important, is to *adjoin* elements.

Definition 9.5
If R is a subring of S, and $\alpha \in S$, then the ring $R[\alpha]$ is defined as the smallest subring of S containing both R and α.

If a subring contains R and α, then it must contain all powers of α, and therefore all linear combinations $\sum r_i \alpha^i$. Meanwhile, the set of such linear combinations is a valid subring (multiplying two such linear combinations gives us another), so $R[\alpha]$ can be explicitly described as

$$R[\alpha] = \left\{ \sum_{i=0}^{n} r_i \alpha^i \mid r_i \in R \right\}.$$

Example 9.6
When $S = \mathbb{C}$, $R = \mathbb{Z}$, and $\alpha = i$, we get the **Gaussian integers** $\mathbb{Z}[i] = \{a + bi \mid a, b \in \mathbb{Z}\}$.

Example 9.7
We have $\mathbb{Q}[\sqrt{2}] = \{a + b\sqrt{2} \mid a, b \in \mathbb{Q}\}$. (Note that $\sqrt{2}^2 = 2$ is already in $\mathbb{Q}$, so we can ignore terms with power at least 2; the same was true in the previous example $\mathbb{Z}[i]$.)

Example 9.8
The ring $\mathbb{Z}[1/2]$ is the set of fractions whose denominator is a power of 2.

Similarly, if we start with elements $\alpha_1, \ldots, \alpha_n$ in S, then $R[\alpha_1, \ldots, \alpha_n]$ is the smallest subring of S containing R and all the α_i. It consists of all sums of products of powers of the α_i, with coefficients in R.

Student Question. *Do the powers of α (when we write elements of $R[\alpha]$) have to be finite?*

Answer. *Yes — when we write*

$$R[\alpha] = \left\{ \sum_{i=0}^{n} r_i \alpha^i \mid r_i \in R \right\},$$

there are infinitely many n to consider, but each sum itself is finite. We don't really have a way to make sense of an infinite sum here — in a ring, we can iterate the operation of addition to get finite sums, but we can't get infinite sums.

Student Question. *Are we allowed to adjoin π to $\mathbb{Z}$, and does this give $\mathbb{R}$?*

Answer. *$\mathbb{Z}[\pi]$ is definitely a legitimate example, and it's a* subring *of $\mathbb{R}$, but it's not $\mathbb{R}$ itself. The fastest way to see it's not $\mathbb{R}$ is that $\mathbb{Z}[\pi]$ is countable and $\mathbb{R}$ is not.*

9.4 Polynomial Rings

There is another way to think of adjoining elements:

Definition 9.9
Let R be a ring, and x a formal variable. Then the **polynomial ring** $R[x]$ is the set

$$R[x] = \left\{ \sum_{i=0}^{n} r_i x^i \mid r_i \in R \right\}$$

(with addition and multiplication defined in the usual way).

Note that for rings such as $\mathbb{R}$, $\mathbb{C}$, or $\mathbb{Q}$, we could instead think of $R[x]$ as the ring of polynomial *functions* from R to itself — but this doesn't work in general.

In general, given any $\alpha \in R$ and $P \in R[x]$, we can always plug in α in place of x and compute the expression $P(\alpha)$; so every polynomial does give a function from R to itself. In fact, this map is compatible with the ring structure:

Definition 9.10
For any fixed $\alpha \in R$, there is a homomorphism $R[x] \to R$ which sends $x \mapsto \alpha$; this is called the **evaluation homomorphism** at α.

Note that here we are fixing α and varying the polynomial (rather than the other way around).

So each $P \in R[x]$ yields a function $R \to R$. But in general, it carries more information than just this function — in general, it's not possible to recover the polynomial from the function. So it's better to think of polynomials in terms of the formal variable rather than in terms of functions.

Example 9.11
Consider the polynomial ring $\mathbb{F}_p[x]$, where $\mathbb{F}_p$ denotes the field $\mathbb{Z}/p\mathbb{Z}$. Even without writing down an explicit example, it is possible to see that $\mathbb{F}_p$ is a finite set, and so the space of functions from $\mathbb{F}_p$ to itself is finite-dimensional as a $\mathbb{F}_p$-vector space. But the space $\mathbb{F}_p[x]$ is infinite-dimensional, since it is spanned by powers of x. Thus, it *cannot* be possible to recover the polynomial $P \in \mathbb{F}_p[x]$ from its corresponding function.

As an explicit example, take $P(x) = x^p - x$, known as the *Artin–Schreier* polynomial. Then $\alpha^p - \alpha = 0$ for all $\alpha \in \mathbb{F}_p$ (by Fermat's Little Theorem), so the polynomial $x^p - x$ corresponds to the zero function $\mathbb{F}_p \to \mathbb{F}_p$, but is not the zero polynomial.

We can define the ring of polynomials in multiple variables — denoted $R[x_1, \ldots, x_n]$ — in the same way, as formal expressions of the form

$$\sum r_{i_1 \ldots i_n} x_1^{i_1} \cdots x_n^{i_n}.$$

We can build on the idea of the evaluation homomorphism, to get an important property of polynomial rings:

Proposition 9.12 (Mapping Property)
Suppose we have a ring R, and a ring homomorphism $\varphi : R \to S$. Then given $\alpha_1, \ldots, \alpha_n \in S$, there exists a unique extension of φ to a homomorphism $\widetilde{\varphi} : R[x_1, \ldots, x_n] \to S$ such that $\widetilde{\varphi}(r) = \varphi(r)$ for $r \in R$, and $\widetilde{\varphi}(x_i) = \alpha_i$ for all i.

This is much less complicated than it appears. The unique extension is just evaluation — we must have

$$\widetilde{\varphi}\left(\sum r_{i_1 \ldots i_n} x_1^{i_1} \cdots x_n^{i_n}\right) = \sum \varphi(r_{i_1 \ldots i_n}) \alpha_1^{i_1} \cdots \alpha_n^{i_n}$$

for any polynomial (this follows directly from the properties of a homomorphism), and this is a valid homomorphism. In some sense, all this proposition is saying is that given a polynomial and some values, we can evaluate the polynomial at those values, and this is compatible with the ring structures.

But it's important because it gives us another way of looking at our original definition of adjoining elements — if $R \subset S$ is a subring, then

$$R[\alpha_1, \ldots, \alpha_n] = R[x_1, \ldots, x_n] / \ker \widetilde{\varphi},$$

where φ is the inclusion map $R \hookrightarrow S$ given by $r \mapsto r$. So we can obtain our initial construction of adjoining specific elements $\alpha_i \in S$ by instead adjoining formal variables to produce a polynomial ring, and then modding out by an ideal.

Example 9.13
We have $\mathbb{Q}[\sqrt{2}] = \mathbb{Q}[x]/(x^2 - 2)$.

Example 9.14
We have $\mathbb{C} = \mathbb{R}[x]/(x^2 + 1)$. This is a particularly good example of how we want to use this construction — in some sense, this is the *definition* of how $\mathbb{C}$ is constructed. In $\mathbb{R}$, there is no element satisfying $x^2 + 1 = 0$, so we simply define some formal variable i which *does* satisfy this equation, and in doing so, we define how $\mathbb{C}$ behaves.

This gives us an idea — we can construct new rings as the quotient of $R[x_1, \ldots, x_n]$ by ideals. Sometimes in algebra, if you want an element with a certain property, you can just add in a variable and state that it satisfies the property, as in the case of defining i to be an element satisfying $i^2 = -1$ (although there is work to do in order to make this construction make sense).

This motivates us to study ideals in polynomial rings. We'll discuss this in more detail next time. But as an example, we can consider $F[x]$ for a field F (note that this is quite restrictive, as we must start with a field, and we only adjoin *one* variable). We'll see that every ideal in $F[x]$ must be principal, meaning every ideal I can be written as (P) — this will essentially follow from polynomial division with remainder. We'll then see that the construction $F[x]/(P)$ can be thought of as adjoining a root of P to F.

10 Ideals in Polynomial Rings

10.1 Ideals in a Field

Last class, we looked at ideals in rings, and began discussing ideals in polynomial rings. Today we will consider what such ideals can look like. The following proposition will be useful later, when discussing maximal ideals:

Proposition 10.1
A ring R is a field if and only if it has exactly two ideals.

Any ring has the ideals $(0) = \{0\}$ and $(1) = R$ — these coincide only in the zero ring, which is not a field by definition. So the proposition states that the ring is a field if and only if it has no other ideals.

Proof. Suppose R is a field, so it has at least two ideals (0) and (1). But there are no other ideals, because every element is invertible — if I is an ideal containing some element $x \neq 0$, then $1 = x^{-1}x$ is in I as well, so $I = (1)$.

Conversely, if R is not a field, then either it is the zero ring and only has one ideal, or it contains a nonzero x which is not invertible. Then (x) cannot contain 1, so (0), (x), and (1) are distinct ideals. □

10.2 Polynomial Rings over a Field

We'll first look at ideals in $F[x]$, the ring of polynomials in one variable over a field.

Proposition 10.2
Every ideal in $F[x]$ is principal. More precisely, if $I \subset F[x]$ is a nonzero ideal and P a (nonzero) element of I of minimal degree, with $\deg(P) = n$, then we have $I = (P)$, and the images of 1, x, x^2, …, x^{n-1} form a basis in $F[x]/I$ (as a vector space over F).

Proof. The main idea is to use division of polynomials with remainder.

In order to check that P generates I, take any $Q \in I$; then by polynomial division we can write

$$Q = P \cdot S + R,$$

where S and R are in $F[x]$ and $\deg R < \deg P$. But R must be in I, so if $R \neq 0$ then this contradicts the choice of P as having minimal degree. So $R = 0$, which means $P \mid Q$. So P generates I.

Now consider the quotient $F[x]/(P)$, and let the images of 1, x, …, x^{n-1} be denoted by $\overline{1}$, $\overline{x}$, …, $\overline{x^{n-1}}$. These images must be linearly independent — for any $P, Q \in F[x]$ we have $\deg(PQ) = \deg(P) + \deg(Q)$ (since F is a field, so has no zero divisors), which means a polynomial with degree less than n cannot be divisible by P. Then if $\overline{1}$, …, $\overline{x^{n-1}}$ were linearly dependent with $a_0 \cdot \overline{1} + \cdots + a_{n-1} \cdot \overline{x^{n-1}} = 0$, then $a_0 + \cdots + a_{n-1}x^{n-1}$ would have to be divisible by P, which is a contradiction.

Conversely, these images must span the quotient by using division with remainder again — for any $Q \in F[x]$, we can write $Q = P \cdot S + R$ with $\deg(R) < n$, which means $\overline{Q} = \overline{R}$ for some $\overline{R}$ which is a linear combination of $\overline{1}$, …, $\overline{x^{n-1}}$. □

Recall that to divide a polynomial Q by P with remainder, we subtract a multiple of P from Q to cancel out the leading term of Q; we then repeat until the remaining polynomial has degree less than that of P.

Note 10.3
This doesn't generalize fully to polynomials over an arbitrary ring R, but some parts do. First, it's not always true that

$$\deg PQ = \deg P + \deg Q.$$

For example, in $\mathbb{Z}/4[x]$, $(2x)(2x+1) = 2x$ does not have degree 2.

Division with remainder also does not necessarily work — for example, we can't divide x^2 by $2x+1$ with remainder in $\mathbb{Z}[x]$. This is because when we cancel out the leading coefficient of Q, we need to scale; and here, 2 isn't invertible, so we can't scale by the correct factor.

But both facts remain true for monic polynomials — so we can divide with remainder if P is monic (meaning it has leading coefficient 1; this works the same way if the leading coefficient is a unit). So if P is monic, then it's still true that every element of $R[x]/(P)$ can be written uniquely as

$$a_0\overline{1} + a_1\overline{x} + \cdots + a_{n-1}\overline{x^{n-1}},$$

for $a_i \in R$. We no longer say that the $\overline{x^i}$ form a basis, since $R[x]/(P)$ is not a vector space (vector spaces are defined over a field); but we will later discuss an analog of vector spaces over rings.

Last time, we mentioned that the construction $F[x]/(P)$ is equivalent to adjoining a root of P — for example, $\mathbb{C} \cong \mathbb{R}[x]/(x^2+1)$. Using Proposition 10.2, we can now make this more precise:

Proposition 10.4
If F is a field and $P \in F[x]$, then $F[x]/(P) \cong F[\alpha]$, where α is a root of P.

Proof. We know that

$$F[x]/(P) = \left\{a_0 + a_1\overline{x} + \cdots + a_{n-1}\overline{x^{n-1}}\right\},$$

where $n = \deg(P)$, while

$$F[\alpha] = \left\{a_0 + a_1\alpha + \cdots + a_{n-1}\alpha^{n-1}\right\}$$

(note that we don't need higher powers of α, as the polynomial relation guarantees that we can express them in terms of these powers). To multiply in $F[\alpha]$, it's enough to understand how to multiply by α. We have

$$\alpha(a_0 + \cdots + a_{n-1}\alpha^{n-1}) = a_0\alpha + \cdots + a_{n-2}\alpha^{n-1} + a_{n-1}\alpha^n,$$

and we can expand α^n using the fact that $P(\alpha) = 0$ — if $P(x) = c_n x^n + \cdots + c_0$, then we have

$$\alpha^n = -c_n^{-1}(c_{n-1}\alpha^{n-1} + \cdots + c_0).$$

This is exactly how multiplication works in $F[x]/(P)$; so the two rings have the same additive and multiplicative structure, and are therefore isomorphic, with the isomorphism given by replacing $\overline{x}$ with α. □

Note 10.5
Here we should think of α as a *universal* root of P — an element that satisfies $P(\alpha) = 0$ but no extra conditions. Given a *specific* root, it's possible that α satisfies additional relations (if P is not irreducible), in which case $F[\alpha]$ would instead be isomorphic to a quotient of $F[x]/(P)$.

10.3 Maximal Ideals

We'll now discuss maximal ideals. These will turn out to be quite useful; today we'll see how to use them to connect algebra to geometry.

Definition 10.6 (Maximal Ideals)
An ideal $I \subset R$ is **maximal** if $I \neq R$, and the only ideals of R containing I are R and I itself.

Example 10.7
The maximal ideals in $\mathbb{Z}$ are (p), for p prime. To prove this, we saw earlier that the ideals of $\mathbb{Z}$ are $n\mathbb{Z}$. But $n\mathbb{Z} \subset m\mathbb{Z}$ if and only if $m \mid n$. So understanding the maximal *ideals* of $\mathbb{Z}$ in the poset of ideals ordered by inclusion is equivalent to understanding the minimal *elements* of $\mathbb{Z}$ in the poset of elements ordered by divisibility. These minimal elements are the primes p, so the maximal ideals are (p).

Example 10.8
For a polynomial ring over a field $F[x]$, any ideal is of the form (P). For the same reason as in the case of $\mathbb{Z}$, the ideal (P) is maximal if and only if P does not factor as QR where Q and R have positive degree, or in other words, if P is irreducible.

Example 10.9
In $\mathbb{C}[x]$, the only irreducible polynomials are linear, since by the Main Theorem of Algebra every polynomial can be factored as $c(x - z_1) \cdots (x - z_n)$. So the maximal ideals of $\mathbb{C}[x]$ are exactly the ideals $(x - \alpha)$.

We'll now see how this example generalizes to polynomials in *multiple* variables. The following proposition will be useful:

Proposition 10.10
An ideal $I \subset R$ is maximal if and only if R/I is a field.

Proof. First, R/I is a field if and only if R/I has exactly two ideals (by Proposition 10.1). But by the correspondence theorem for rings, ideals in R/I are in bijection with ideals in R containing I. So R/I is a field if and only if R has exactly two ideals containing I. But this is equivalent to the condition that I is maximal (since I and R are both containing I, so if there are only two such ideals, then there can be no others). □

Example 10.11
In the case of $\mathbb{Z}$, as we've mentioned earlier, $\mathbb{Z}/p\mathbb{Z}$ is a field.

Example 10.12
If F is a field and $P \in F[x]$ is irreducible, then $F[x]/(P)$ is a field as well — this is a construction which can be used to build new fields.

This describes what happens when we look at polynomials in *one* variable over a field, but we can try to consider what happens for polynomials in *multiple* variables as well.

10.4 Ideals in Multivariate Polynomial Rings

Example 10.13
Let $R = F[x_1, \ldots, x_n]$, where F is a field. Fixing scalars $\alpha = (\alpha_1, \ldots, \alpha_n)$ (with $\alpha_i \in F$), we have the evaluation homomorphism $F[x_1, \ldots, x_n] \to F$ defined as

$$\mathrm{ev}_\alpha : P \mapsto P(\alpha_1, \ldots, \alpha_n).$$

This map is clearly onto, as the constants in $F[x_1, \ldots, x_n]$ are sent to themselves, so by the first isomorphism theorem for rings, we have

$$F \cong F[x_1, \ldots, x_n]/\ker(\mathrm{ev}_\alpha).$$

Since F is a field, the kernel is a maximal ideal of R. In fact, this kernel can be explicitly written as

$$\mathfrak{m}_\alpha = (x_1 - \alpha_1, \ldots, x_n - \alpha_n).$$

This provides a way to construct maximal ideals of $F[x_1, \dots, x_n]$; building on this, we can also try to construct maximal ideals of *quotients* of this ring (since many rings can be constructed in this way).

Suppose that $R = F[x_1, \dots, x_n]/J$, with $J = (P_1, \dots, P_m)$. (We'll later see that for *any* J, we can find finitely many polynomials P_i which generate J; but we won't focus on that right now.) Then for any α for which $\mathfrak{m}_\alpha \supset J$, where $\mathfrak{m}_\alpha$ is the maximal ideal of $F[x_1, \dots, x_n]$ as defined above, by the correspondence theorem the image of $\mathfrak{m}_\alpha$ is also a maximal ideal of $F[x_1, \dots, x_n]/J$.

But we know $\mathfrak{m}_\alpha \supset J$ if and only if $\alpha = (\alpha_1, \dots, \alpha_n)$ is a common zero of P_1, ..., P_n (since $\mathfrak{m}_\alpha = (x_1 - \alpha_1, \dots, x_n - \alpha_n)$ contains P_i if and only if α is a zero of P_i). This gives the following construction:

Proposition 10.14
Suppose $R = F[x_1, \dots, x_n]/(P_1, \dots, P_m)$. Then any common zero $\alpha = (\alpha_1, \dots, \alpha_n)$ of P_1, ..., P_m yields a maximal ideal of R, which is the image of $\mathfrak{m}_\alpha$ when quotienting out by $(P_1, \dots, P_m)$.

The sets of common zeroes of a list of polynomials are studied in algebraic geometry. Here, we can see that such common zeroes can be used to produce maximal ideals of R. A famous theorem states that in the case $F = \mathbb{C}$, the converse is true as well:

Theorem 10.15 (Hilbert's Nullstelensatz)
Every maximal ideal of $\mathbb{C}[x_1, \cdots, x_n]$ is of the form $\mathfrak{m}_\alpha$ for some $\alpha = (\alpha_1, \cdots, \alpha_n)$.

In the name of the theorem (in German), "null" means zero, "stelen" means place, and "satz" means theorem.

This immediately implies the following corollary (since maximal ideals in $\mathbb{C}[x_1, \cdots, x_n]/(P_1, \cdots, P_m)$ are in correspondence with maximal ideals in $\mathbb{C}[x_1, \dots, x_n]$ containing all the P_i):

Corollary 10.16
The maximal ideals in $R = \mathbb{C}[x_1, \cdots, x_n]/(P_1, \cdots, P_m)$ are in bijection with the common zeroes of the polynomials P_i.

Proof of Theorem 10.15. Let $\mathfrak{m} \subset \mathbb{C}[x_1, \dots, x_n]$ be a maximal ideal; then $F = \mathbb{C}[x_1, \dots, x_n]/\mathfrak{m}$ is a field. This gives a homomorphism from $\mathbb{C}$ to F — taking the quotient by $\mathfrak{m}$ gives a homomorphism from $\mathbb{C}[x_1, \dots, x_n]$ to F, and we can restrict the homomorphism to $\mathbb{C}$. It suffices to show that this map $\mathbb{C} \to F$ is an isomorphism — then the original homomorphism $\mathbb{C}[x_1, \dots, x_n] \to F$ from taking the quotient must map $\mathbb{C}$ isomorphically to F, and we can suppose it maps $x_i \to \alpha_i$ for each i where $\alpha_i \in \mathbb{C}$ (it really maps x_i to some element of F, but F is isomorphic to $\mathbb{C}$). Then it's clear that the kernel of this homomorphism is generated by $x_1 - \alpha_1$, ..., $x_n - \alpha_n$; therefore this kernel is $\mathfrak{m}_\alpha$ where $\alpha = (\alpha_1, \dots, \alpha_n)$, and we have $\mathfrak{m} = \mathfrak{m}_\alpha$. So now we want to show that the map $\mathbb{C} \to F$ is bijective.

But *any* homomorphism between fields is injective — the kernel of the homomorphism must be an ideal, but the only ideals of a field are $\{0\}$ and the entire field. Since the homomorphism cannot map 1 to 0, the kernel cannot be the entire field, so must be $\{0\}$.

So it suffices to show that the homomorphism is surjective. Assume not. Then F strictly contains $\mathbb{C}$ (since we have an injective map $\mathbb{C} \hookrightarrow F$, so $\mathbb{C}$ is isomorphic to its image), so we can pick $z \in F$ with $z \notin \mathbb{C}$. Then consider the elements

$$\left\{ \frac{1}{z - \lambda} \mid \lambda \in \mathbb{C} \right\}.$$

Now we'll use a bit of set theory — $\mathbb{C}[x_1, \dots, x_n]$ is a countable union of finite-dimensional vector spaces $U_1 \subset U_2 \subset \cdots$ over $\mathbb{C}$, where U_i is the vector space of polynomials with degree at most i. Then since F is a quotient of $\mathbb{C}[x_1, \dots, x_n]$, it must also be a countable union of finite-dimensional vector spaces $V_1 \subset V_2 \subset \cdots$, where V_i is simply the image of U_i in this quotient.

On the other hand, $\mathbb{C}$ is not countable, so there are uncountably many terms $1/(z - \lambda)$. Since all such terms are elements of F, one of the finite-dimensional vector spaces that F consists of must contain infinitely many of them.

But then since these terms all belong to the same finite-dimensional vector space, and there's infinitely many of them, we can find a finite set which is linearly dependent — so then we have an identity

$$\sum \frac{a_i}{z - \lambda_i} = 0,$$

where $a_i \in \mathbb{C}$ for all i and there are finitely many terms. But by clearing denominators, we can translate this into a polynomial relation in z — we then have $P(z) = 0$ for some $P \in \mathbb{C}[x]$. Then since all polynomials over $\mathbb{C}$ factor, we can write

$$P(x) = c \prod_i (x - r_i),$$

for some $r_i \in \mathbb{C}$. But $z \notin \mathbb{C}$, so z cannot equal any of the r_i, contradiction (as F is a field, so the product of nonzero terms cannot be zero).

So then the map $\mathbb{C} \to F$ must be an isomorphism, as desired. □

11 More About Rings

11.1 Review: Hilbert's Nullstelensatz

Last time, we proved Hilbert's Nullstelensatz:

Theorem 11.1 (Hilbert's Nullstelensatz)
The maximal ideals in $\mathbb{C}[x_1, \cdots, x_n]$ are exactly the kernels of evaluation homomorphisms, and thus they are in bijection with $\mathbb{C}^n$.

Corollary 11.2
The maximal ideals in $\mathbb{C}[x_1, \cdots, x_n]/(P_1, \cdots, P_m)$ are in bijection with the common zeroes of $P_1, \ldots, P_m$.

It's clear that this corollary follows from the theorem, since maximal ideals in $\mathbb{C}[x_1, \ldots, x_n]/(P_1, \ldots, P_m)$ correspond to maximal ideals of $\mathbb{C}[x_1, \ldots, x_n]$ containing $(P_1, \ldots, P_m)$, and a maximal ideal $\mathfrak{m}_\alpha$ contains all the P_i if and only if they all evaluate to 0 when plugging in α.

As a brief recap of the ideas seen in the proof of Theorem 11.1:

Proof Sketch. The proof reduces to showing that if F is a field containing $\mathbb{C}$, such that there exists a surjective map $\mathbb{C}[x_1, \ldots, x_n] \twoheadrightarrow F$, then $F = \mathbb{C}$. (In this case, $F = \mathbb{C}[x_1, \ldots, x_n]/\mathfrak{m}$, and the surjective map comes from taking the quotient.)

We first saw that F is a union of countably many finite-dimensional $\mathbb{C}$-vector spaces — it's clear that $\mathbb{C}[x_1, \ldots, x_n]$ is a union of countably many finite-dimensional $\mathbb{C}$-vector spaces $U_1 \subset U_2 \subset \cdots$ (we can take the vector space consisting of polynomials of degree at most d), and then to exhaust F by finite-dimensional $\mathbb{C}$-vector spaces, we can simply take the images V_i of the vector spaces U_i which exhaust $\mathbb{C}[x_1, \ldots, x_n]$. So we can write

$$F = \bigcup_{i=1}^{\infty} V_i,$$

where $\dim_{\mathbb{C}} V_i$ is finite for all i.

Now if $F \neq \mathbb{C}$, we can pick $z \in F$ with $z \notin \mathbb{C}$, and consider $1/(z - \lambda)$ for all $\lambda \in \mathbb{C}$. There are uncountably many such elements (by set theory, $\mathbb{C}$ is not countable), and since there's countably many V_i, infinitely many of these elements must lie in the same space V_i.

But then by linear algebra, there must be a finite sum

$$\sum_{i=1}^{n} \frac{a_i}{z - \lambda_i} = 0$$

with $a_i \in \mathbb{C}$ (since if n is greater than the dimension of the vector space, these elements must be linearly dependent). But clearing denominators, we get

$$\sum_{i=1}^{n} a_i \prod_{i \neq j} (z - \lambda_j) = 0.$$

But the left-hand side is a nonzero polynomial P in z — to see it's nonzero, we can plug in λ_1 and see that $P(\lambda_1) = a_1 \prod_{j>1} (\lambda_1 - \lambda_j) \neq 0$. Since $P \in \mathbb{C}[x]$, it must factor completely over $\mathbb{C}$. But z is not in $\mathbb{C}$, so it cannot equal any of the roots of P; this is a contradiction. □

Note 11.3
In fact, the theorem holds for all fields which are algebraically closed (meaning that every polynomial has a root — here we used the fact that $\mathbb{C}$ was algebraically closed in order to factor P in the final step). For example, it holds for the field of algebraic numbers as well. The specific argument we used here doesn't work in that case, since the algebraic numbers are countable; but there are other proofs as well.

Student Question. *What does it mean to take the union of vector spaces?*

Answer. *In general, this doesn't really make sense (the union of vector spaces may not itself be a vector space). But in this case, we can get vector spaces* $V_1 \subset V_2 \subset \cdots$*, by taking* $V_i = \operatorname{im}(\mathbb{C}[x_1, \ldots, x_n]_{\leq i})$ *(the notation* $\mathbb{C}[x_1, \ldots, x_n]_{\leq i}$ *denotes polynomials of degree at most* i*), and when we have an increasing chain of vector spaces, taking their union* does *make sense.*

This is important because algebraic geometry studies the sets of zeros of a polynomial, and this gives us an algebraic way to think about them.

Definition 11.4
For a ring R, the **maximal spectrum** of R, denoted $\operatorname{MSpec}(R)$, is the set of maximal ideals in R.

The maximal spectrum plays an important role in algebraic geometry and commutative algebra.

For instance, each element $r \in \mathbb{R}$ defines a "function" f_r on $\operatorname{MSpec}(R)$, where f_r sends each maximal ideal $\mathfrak{m}$ to the element $\overline{r}$ in $R/\mathfrak{m}$ (which is a field). If $R = \mathbb{C}[x_1, \ldots, x_n]/I$ is a quotient of a polynomial ring over $\mathbb{C}$, then $R/\mathfrak{m} = \mathbb{C}$ by Hilbert's Nullstelensatz, so f_r is actually a function $\operatorname{Mspec}(R) \to \mathbb{C}$. In fact, this function is given by evaluating the polynomial r at the point corresponding to $\mathfrak{m}$ — so we've recovered the original polynomial function. But in general, there isn't even a guarantee that the fields $R/\mathfrak{m}$ are all isomorphic — they may be different for different $\mathfrak{m}$. This is why "function" is in quotation marks — where the map takes values *depends* on its input.

11.2 Inverting Elements

Last time, we discussed adjoining a root of a polynomial to a ring — in particular, we discussed the structure of $R[x]/(P)$ for a monic polynomial P.

Instead of setting P to be monic, we can set it to be linear, and consider $R[x]/(ax - 1)$, which is denoted by $R_{(a)}$. We've essentially added a variable x and declared it to be the inverse of a; so $R_{(a)}$ is the result of formally inverting a. This construction is known as **localization**.

Example 11.5
We have $\mathbb{Z}_{(2)} = \{a/2^n \mid a, n \in \mathbb{Z}\}$ and $\mathbb{Z}_{(6)} = \{a/2^n 3^m \mid a, n, m \in \mathbb{Z}\}$.

However, we must be careful. In these examples, we were able to simply add a formal inverse of a to R. But it's possible that this might collapse some of R — in particular, if $ab = 0$ for nonzero a and b (which is possible in a general ring), then the image of b in $R_{(a)}$ will vanish. This is because if $ab = 0$ and the image of a is invertible, then the image of b must be 0.

Example 11.6
In $(\mathbb{Z}/6\mathbb{Z})_{(2)}$, the image of 3 vanishes — in particular, $(\mathbb{Z}/6\mathbb{Z})_{(2)} \cong \mathbb{Z}/3\mathbb{Z}$. Meanwhile, $(\mathbb{Z}/4\mathbb{Z})_{(2)}$ is the zero ring — this is because $2 \cdot 2 = 0$, but the image of $2 \cdot 2$ in $(\mathbb{Z}/4\mathbb{Z})_{(2)}$ is invertible.

So when inverting elements, we want to make sure this doesn't happen.

Definition 11.7
An element $a \in R$ is a **zero divisor** if $a \neq 0$, and there exists $b \neq 0$ for which $ab = 0$.

For example, 2 and 3 are zero divisors in $\mathbb{Z}/6\mathbb{Z}$.

Proposition 11.8
If a is *not* a zero divisor, then $R \subset R_{(a)}$.

So we can safely invert elements which are not zero divisors.

We've seen how to invert *one* element a, and this directly generalizes to let us invert finitely many elements; but we can also try to invert *all* (nonzero) elements. For this to make sense, we'd need all elements to not be zero divisors:

Definition 11.9
A ring R is an **integral domain** if R has no zero divisors.

One useful property of integral domains is that we can perform cancellation — if we have $ax = ay$ with $a \neq 0$, then we must have $x = y$.

Definition 11.10
Let R be an integral domain. Then the **fraction field** of R, denoted $\mathrm{Frac}(R)$, is the set $\{(a, b) \mid a, b \in R, b \neq 0\}$ modulo the equivalence relation that $(a, b) \sim (c, d)$ if $ad = bc$.

This is the formal definition, but when we work with a fraction field, it's more intuitive to think of the elements as fractions in the usual sense — we write a/b instead of (a, b). The operations do work in the same way we're used to — we have

$$\frac{a}{b} \cdot \frac{c}{d} = \frac{ac}{bd} \text{ and } \frac{a}{b} + \frac{c}{d} = \frac{ad + bc}{bd}.$$

It's clear that $F = \mathrm{Frac}(R)$ is a *field* containing R.

Example 11.11
A familiar example of a fraction field is $\mathrm{Frac}(\mathbb{Z}) = \mathbb{Q}$.

Example 11.12
The fraction field $\mathrm{Frac}(\mathbb{C}[x])$ is called the field of **rational functions** in one variable, and is denoted $\mathbb{C}(x)$; it consists of elements of the form $p(x)/q(x)$, where p and q are polynomials. Each of its elements defines a function on $\mathbb{C}$ (which is defined everywhere except for some number of poles).

This concept can be extended to multiple variables — we can also consider the fraction field of $\mathbb{C}[x_1, \ldots, x_n]$.

Note that for a field F, we have $\mathrm{Frac}(F) = F$ (since all nonzero elements are already invertible).

In general, giving a homomorphism from $\mathrm{Frac}(R)$ to S is the same as giving a homomorphism from R sending nonzero elements of R to S where the image of each nonzero element of R is an invertible element in S. (Invertible elements of a ring are also called *units*.)

11.3 Factorization

Now we will discuss factorization in certain rings. A simple case is polynomials over a field.

Proposition 11.13
For a field F, every polynomial $P \in F[x]$ factors as a product of irreducible polynomials in an essentially unique way (up to rearrangement of the factors or multiplying the factors by scalars).

In order to prove this, we'll use the following lemma:

Lemma 11.14
If P is irreducible and $P \mid QS$, then $P \mid Q$ or $P \mid S$.

Proof. Since P is irreducible and all ideals of $F[x]$ are principal, (P) is a maximal ideal, and therefore $F[x]/(P)$ is a field. So if P divides Q, then the image of Q in the quotient is zero — so the lemma is equivalent to stating that there are no zero divisors in the field, which is true. More explicitly, if $P \mid QS$, then $\overline{Q} \cdot \overline{S} = 0$ (where $\overline{Q}$ denotes Q mod P), which means either $\overline{Q} = 0$ or $\overline{S} = 0$. □

The proposition then essentially follows formally:

Proof of Proposition 11.13. Proving the *existence* of such a factorization is easy — starting with a polynomial, if it isn't irreducible, then we can factor it as a product of polynomials with strictly smaller degree. Since the degree can't keep on decreasing, this process must eventually stop, at which point all our factors are irreducible. (This can be made more formal by using induction on the degree of the polynomial.)

To prove uniqueness, suppose that P factors as $P_1 \cdots P_n = Q_1 \cdots Q_m$, where all of the P_i and Q_i are irreducible. By Lemma 11.14, since P_1 divides $Q_1 \cdots Q_m$, we must have $P_1 \mid Q_i$ for some i. Without loss of generality this means $P_1 = Q_1$ — if $P_1 \mid Q_1$ then we must have $Q_1 = \lambda P_1$ for some scalar λ (since Q_1 is irreducible), and we can rescale the factors to make $\lambda = 1$. Then we have $P_2 \cdots P_n = Q_2 \cdots Q_m$, and we can perform the same argument to keep cancelling out common factors (again this can be made more formal by using induction on degree). □

This argument can be used to prove unique factorization in other situations as well, motivating the following definitions:

Definition 11.15
An integral domain is a **principal ideal domain** (PID) if every ideal is principal.

Definition 11.16
An integral domain is a **unique factorization domain** (UFD) if every element factors as a product of irreducibles in an essentially unique way.

The argument we used to prove Proposition 11.13 more generally proves that every PID is a UFD. The converse is not true — in future classes, we'll see that $\mathbb{Z}[x]$ and $\mathbb{C}[x_1, \ldots, x_n]$ are UFDs but not PIDs.

12 Factorization in Rings

12.1 Review

Last class, we began discussing factorization.

Definition 12.1
An element $a \in R$ is **irreducible** if it is not a unit, and if $a = bc$ then either b or c is a unit.

In other words, irreducible elements are ones which cannot be factored in a nontrivial way; so when attempting to factor in a ring, we want to factor our elements as a product of irreducibles.

In our discussion of factorization, we'll always assume R is an integral domain (meaning that if $ab = 0$, then either $a = 0$ or $b = 0$) — this allows us to perform cancellation.

When discussing unique factorization, we can always multiply the factors by units; so to make the notion of "essentially unique" (as mentioned last class) more precise, we use the following definition:

Definition 12.2
Two elements $a, b \in R$ are **associate** if $a = bu$ for a unit $u \in R$.

Then a domain R is a unique factorization domain (UFD) if every non-unit element can be written as a product of irreducible elements in a unique way, up to ordering and association.

Recall that a domain R is a principal ideal domain (PID) if every ideal in R is principal. As mentioned last class, we can generalize our proof that $F[x]$ is a UFD to work for any PID:

Theorem 12.3
Any PID is a UFD.

Sketch of Proof. We need to show that a factorization exists, and is unique.

To prove uniqueness, since R is a PID, we have that if $p \in R$ is irreducible, then (p) is maximal. So then $R/(p)$ is a field, and since fields have no zero divisors, it follows that if p divides ab, then p divides a or p divides b. This implies uniqueness — now given any two factorizations $p_1 p_2 \cdots p_m = q_1 q_2 \cdots q_m$, we can show that p_1 must appear on the right-hand side as well (up to association), and cancel it out from both sides.

We won't prove existence in general now (it requires a new idea, which we'll see later). But in the examples we'll deal with, existence is clear. We can start with any element of R and keep factoring it until we're stuck, at which point all factors must be irreducible. Then in our examples, this factorization process always "shrinks" the elements in some sense — in the case of integers, their size decreases, and in the case of polynomials, their degree decreases — so it must terminate. (We can't perform this argument in an abstract PID because it doesn't necessarily have a notion of size. We will later see a different way to show that the process terminates, using *Noetherian rings*.) □

Note 12.4
Elements with the property that if $p \mid ab$, then $p \mid a$ or $p \mid b$ (which we used in the proof of uniqueness) are called **prime**.

12.2 Euclidean Domains

Earlier, we saw that for a field F, the ring $F[x]$ is a PID. We can apply the argument used here to a somewhat more general class of rings.

Definition 12.5
A Euclidean domain is a domain R together with a size function $\sigma : R \setminus \{0\} \to \mathbb{Z}_{>0}$ such that for every $a, b \in R$ with $b \neq 0$, there exist $q, r \in R$ such that

$$a = bq + r,$$

and $\sigma(r) < \sigma(b)$ or $r = 0$.

In other words, a Euclidean domain is a domain where we can perform division with remainder, such that the remainder has smaller size than the element we're dividing by.

Proposition 12.6
A Euclidean domain is a PID, and therefore a UFD.

Example 12.7
The familiar ring $\mathbb{Z}$ is a Euclidean domain with size function $\sigma(a) = |a|$.

Example 12.8
For a field F, the polynomial ring $F[x]$ is a Euclidean domain with $\sigma(P) = \deg P$.

Example 12.9
The Gaussian integers $\mathbb{Z}[i]$ form a Euclidean domain with size function $\sigma(a + bi) = a^2 + b^2$.

Proof. We can prove that the division with remainder property holds by geometry. Given b, the multiples of b form a square lattice (generated as a lattice by b and bi).

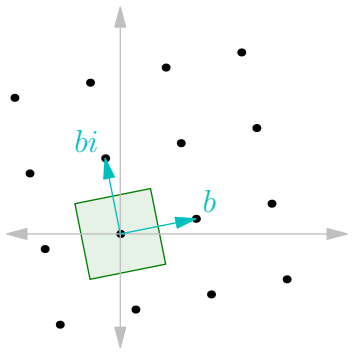

So by subtracting multiples of b, we can guarantee that a lands in the small square centered at the origin — more precisely, we can guarantee that $r = \alpha b + \beta ib$ where $-\frac{1}{2} \leq \alpha, \beta \leq \frac{1}{2}$. Then we have $\sigma(r) \leq \frac{1}{2}\sigma(b) < \sigma(b)$, as desired. □

So the concept of a Euclidean domain is useful — there exist examples other than the ones we started thinking about. We can now prove that Euclidean domains are PIDs, in the same way as we did with polynomials.

Proof of Proposition 12.6. If $I \subset R$ is a nonzero ideal, then take an element $b \in I$ with minimal $\sigma(b)$. We know that for any $a \in I$, we can write $a = bq + r$, with $r = 0$ or $\sigma(r) < \sigma(b)$. The second case is impossible — we have $r \in I$, but we chose b to have minimal size of the nonzero elements in I — so we must have $r = 0$. So b divides all elements of I, which means $I = (b)$. □

However, this isn't *very* general — there are many rings which it *doesn't* cover.

Example 12.10
The ring $\mathbb{Z}[\sqrt{-5}]$ is not a UFD, and therefore not a PID or Euclidean domain.

Proof. We have

$$6 = 2 \cdot 3 = (1 + \sqrt{-5})(1 - \sqrt{-5}).$$

It's possible to show that all of the elements 2, 3, and $1 \pm \sqrt{-5}$ are irreducible, so R does not have unique factorization.

Note that it *is* still possible to bound $\sigma(r)$ in terms of $\sigma(b)$ by the same geometric argument as before; but this bound will not be strong enough to imply $\sigma(r) < \sigma(b)$. □

Student Question. *The size functions in our examples have nice properties — the size functions on $\mathbb{Z}$ and $\mathbb{Z}[i]$ are multiplicative, and the size function in $F[x]$ satisfies $\sigma(PQ) = \sigma(P) + \sigma(Q)$. Does something like this have to hold in general?*

Answer. *No — the ones that we've seen in our examples* do *satisfy additional properties, but we didn't need those nice properties for the argument to work. The definition itself also doesn't guarantee any other nice properties. For example, in a field, the size function can be* anything *— every element is divisible by every nonzero element, so we can perform division even without remainder (meaning that $r = 0$).*

12.3 Polynomial Rings

Every PID is a UFD, but the converse is not true! There are cases where unique factorization is true, but there are non-principal ideals. For example, we'll see that $\mathbb{Z}[x]$ and $\mathbb{C}[x_1, \ldots, x_n]$ are UFDs. But the ideal $(2, x) \subset \mathbb{Z}[x]$ and the ideal $(x, y) \subset \mathbb{C}[x, y]$ are not principal.

The theorem that will imply both of these examples is the following.

Theorem 12.11
If R is a UFD, then $R[x]$ is also a UFD.

Corollary 12.12
The rings $\mathbb{Z}[x]$ and $\mathbb{C}[x_1, \ldots, x_n]$ are UFDs.

Proof of Corollary. For $\mathbb{Z}[x]$, this follows directly from the theorem (since we know $\mathbb{Z}$ is a PID). Meanwhile, for $\mathbb{C}[x_1, \ldots, x_n]$, we can use induction: we have

$$\mathbb{C}[x_1, \ldots, x_n] = \mathbb{C}[x_1, \ldots, x_{n-1}][x_n]$$

by thinking of n-variable polynomials as polynomials in the last variable x_n, whose coefficients are polynomials in the other $n - 1$ variables — for example,

$$x + xy + y^2x^2 + xy^2 = (x) + (x)y + (x + x^2)y^2$$

is a polynomial in y whose coefficients are in $\mathbb{C}[x]$. So using induction, this follows immediately from the theorem as well. □

12.3.1 Greatest Common Divisors

To prove Theorem 12.11, we'll need the concept of a gcd in R.

Definition 12.13
In a domain R, a **greatest common divisor** of two elements $a, b \in R$, denoted $\gcd(a, b)$, is an element d such that d divides both a and b, and any other element δ that divides both a and b must also divide d.

A gcd may or may not exist. But if $\gcd(a, b)$ exists, it is unique up to association, i.e., up to multiplying by a unit — if d and d' are both gcd's of a and b, then we must have $d \mid d'$ and $d' \mid d$. This implies we have $d = ud'$ and $d' = zd$ for some elements u and z. Then $d = uzd$, and since R is a domain, we have $uz = 1$, so u and z are both units.

Example 12.14
In $\mathbb{Z}[\sqrt{-5}]$, there is no gcd of $2 + 2\sqrt{-5}$ and 6.

Proof. Note that 2 is a common divisor of $2 + 2\sqrt{-5}$ and 6, and it's maximal in the sense that if multiplied by any non-unit element, the result is no longer a common divisor. So if the gcd existed, it would have to be 2 (up to association). But $1 + \sqrt{-5}$ has the same property — in particular, $1 = \sqrt{-5}$ is a common divisor but does not divide 2. So there cannot exist a gcd. □

Proposition 12.15
In a UFD, the gcd of any two elements always exists.

Proof. The usual way of calculating the gcd using prime factorization (for example, in the case of integers) works in any PID. More explicitly, to find $\gcd(a, b)$ we can write down the factorizations of a and b, and take the smaller power of each irreducible element. □

Note 12.16
In a PID, if $\gcd(a, b) = d$, then we have $(a, b) = (d)$, which means d can be written in the form $ap + bq$. But this is not true in general — for example, in $\mathbb{C}[x, y]$, we have $\gcd(x, y) = 1$, but $1 \notin (x, y)$.

12.3.2 Gauss's Lemma

Our goal is to analyze factorization in $R[x]$. We know how factorization works in R, and we *also* know how factorization works in a closely related ring — if $F = \mathrm{Frac}(R)$, then since F is a field, $F[x]$ is a PID. To relate factorization in $R[x]$ to factorization in these two better-understood rings, we use Gauss's Lemma.

Definition 12.17
A polynomial $P \in R[x]$ is **primitive** if the gcd of all its coefficients is a unit.

Lemma 12.18 (Gauss's Lemma)
If $P, Q \in R[x]$ are primitive, then so is PQ.

Proof. It's enough to show that for any irreducible $p \in R$, we can find a coefficient of PQ not divisible by p (as then by unique factorization, no element other than units can divide the gcd of its coefficients).

Let $P = \sum a_i x^i$ and $Q = \sum b_j x_j$, and let m be the maximal integer with $m \nmid a_m$ and n the maximal integer with $n \nmid b_n$. Then in PQ, the coefficient of x^{m+n} comes from $a_m b_n$, and other terms $a_i b_j$ where at least one of a_i and b_j is divisible by p; so this coefficient cannot be divisible by p. □

Using this, we can get a good sense of which polynomials are irreducible in R — as we'll see later, these are the irreducible elements of R, and primitive polynomials in $R[x]$ which are irreducible in $F[x]$, where $F = \mathrm{Frac}(R)$. So we can use unique factorization in R and in $F[x]$ to prove unique factorization in $R[x]$.

13 More Factorization

13.1 Factoring Integer Polynomials

We've seen that $\mathbb{Z}$, $F[x]$, and $\mathbb{Z}[i]$ are unique factorization domains — in fact, we'll later look into a neat application of unique factorization in $\mathbb{Z}[i]$ to a problem in number theory. First we'll discuss factorization in $\mathbb{Z}[x]$. We previously stated the following theorem:

Theorem 13.1
If R is a unique factorization domain, then $R[x]$ is also a unique factorization domain.

We'll prove this for the case $R = \mathbb{Z}$. The proof of the general case is very similar to the proof for $\mathbb{Z}$ — we essentially just have to replace the familiar construction of the gcd over integers with the more abstract notion of gcd in a general UFD, as discussed last time.

Guiding Question
Given a polynomial $P \in \mathbb{Z}[x]$, there's two natural questions we can ask about its factorization:

1. Is it possible to factor P where the factors lie in $\mathbb{Q}[x]$?
2. What about in $\mathbb{Z}[x]$?

There are more tools available for approaching the second question, factoring in $\mathbb{Z}[x]$, which make it possible to reduce the potential factorizations to a finite number of possibilities. One such tool is reducing mod a prime p.

Example 13.2
Consider the polynomial $P(x) = 3x^2 + 2x + 2$. We can show it's impossible to factor P in $\mathbb{Z}[x]$ by reducing mod 2 — we have

$$P(x) \equiv x^3 \pmod 2.$$

But the only way x^3 factors in $\mathbb{F}_2[x]$ is as $x^3 \cdot 1$ or $x^2 \cdot x$. If P factored as $P_1 P_2$ where P_1 and P_2 had positive degree, then since their leading coefficients must be odd (as these leading coefficients multiply to 3), they must still have positive degree in $\mathbb{F}_2[x]$. So P_1 and P_2 must be congruent to x and x^2 mod 2; in particular, both their free terms are divisible by 2. But the free term of P would then be divisible by 4, which is a contradiction.

This example illustrates one possible trick that can be used to show that a polynomial is irreducible in $\mathbb{Z}[x]$. There are various other tricks as well. For example, the product of the free terms of the factors must equal the free term of the original polynomial; so we can look at all possible factorizations of the free term. On the other hand, factoring polynomials in $\mathbb{Q}[x]$ seems much more difficult, since these tricks no longer work — there's infinitely many ways to factor the free term of the original polynomial as a product of *rationals*, so this argument can't be used to reduce our search to finitely many possibilities.

Fortunately, it turns out that factoring over $\mathbb{Z}[x]$ and $\mathbb{Q}[x]$ are actually equivalent. To show this, recall the definition of a primitive polynomial from last class, here applied specifically to $\mathbb{Z}[x]$:

Definition 13.3
A polynomial $P \in \mathbb{Z}[x]$ is **primitive** if the gcd of all its coefficients is 1.

Evidently, any nonzero $P \in \mathbb{Z}[x]$ can be written as a product $P = nQ$, where Q is primitive and n is the gcd of the coefficients of P. In fact, any $P \in \mathbb{Q}[x]$ can be scaled to a primitive polynomial, by clearing denominators and factoring out the gcd of its coefficients.

The key point in relating factorization in $\mathbb{Z}[x]$ and $\mathbb{Q}[x]$ is Gauss's Lemma, which we proved last class:

Lemma 13.4 (Gauss's Lemma)
If P and Q are primitive, then PQ is as well.

Proof Sketch. If PQ is not primitive, then there is some prime p which divides all coefficients of PQ. Now consider P and Q mod p; since both are nonzero mod p, it's clear that their product is nonzero mod p as well — the integers mod p are a field, and for polynomials over a field, it's impossible to multiply two nonzero polynomials and get the zero polynomial. □

Using Gauss's Lemma, we can reduce questions about divisibility in $\mathbb{Z}[x]$ to ones about divisibility in $\mathbb{Q}[x]$, via the following corollary:

Corollary 13.5
If $P, Q \in \mathbb{Z}[x]$ are such that P divides Q in $\mathbb{Q}[x]$ and P is primitive, then P divides Q in $\mathbb{Z}[x]$.

Proof. We have $Q = P \cdot S$ for some $S \in \mathbb{Q}[x]$. Now write $S = aT/b$ where $T \in ZZ[x]$ is primitive, and a and b are integers with $b \neq 0$. Then the equation can be rewritten as

$$bQ = aPT.$$

By Gauss's Lemma, PT is primitive, so the gcd of all coefficients of aPT is exactly a. Meanwhile, b certainly divides all coefficients of bQ, so it divides all coefficients of aPT as well, which means $b \mid a$. As a result, $a/b \in \mathbb{Z}$, so $S \in \mathbb{Z}[x]$ and P divides Q in $\mathbb{Z}[x]$ as well. □

Note 13.6
There's a different way to phrase this proof — for polynomials $P \in \mathbb{Z}[x]$, we can define the **content** of P, denoted $c(P)$, as the gcd of the coefficients of P. It's possible to extend this to polynomials in $\mathbb{Q}[x]$ as well, such that for any $T \in \mathbb{Q}[x]$ and $a \in \mathbb{Q}$, we have $c(aT) = a \cdot c(T)$. Then Gauss's Lemma states that $c(PQ) = c(P)c(Q)$ for any $P, Q \in \mathbb{Z}[x]$, and therefore for any $P, Q \in \mathbb{Q}[x]$ as well. Now in this proof we have $Q = PS$, which means $c(Q) = c(P)c(S)$. But $c(Q)$ is an integer and $c(P) = 1$, so $c(S)$ must be an integer as well; therefore S has integer coefficients.

As a result, factoring in $\mathbb{Z}[x]$ and $\mathbb{Q}[x]$ are in fact equivalent.

Example 13.7
The polynomial $3x^2 + 2x + 2$ is irreducible in $\mathbb{Z}[x]$, so it cannot be factored in $\mathbb{Q}[x]$ either.

Corollary 13.8
The irreducible elements in $\mathbb{Z}[x]$ fall into two categories: $\pm p$ for prime integers p, and primitive polynomials which are irreducible in $\mathbb{Q}[x]$.

It's not necessarily easy to tell whether a polynomial is irreducible; but this does mean that answering the question of whether a polynomial is irreducible in $\mathbb{Q}[x]$ is as easy as answering it in $\mathbb{Z}[x]$.

Proof. It's clear that both categories of elements are irreducible — the primes are clearly irreducible in $\mathbb{Z}[x]$, since the only way to factor a constant polynomial is as a product of constants. Meanwhile, if a polynomial is irreducible in $\mathbb{Q}[x]$, then the only way it can be factored is by pulling out constant factors; but this is impossible for a primitive polynomial, so all such polynomials must be irreducible in $\mathbb{Z}[x]$ as well.

On the other hand, if P is not of either form, then we'll show that P can be factored and therefore is not irreducible. First, if $\deg(P) = 0$ (meaning P is an integer), then it's clear that it must be prime in order to be irreducible.

Now assume that $\deg(P) \geq 1$. If P is not primitive, then we can pull out the greatest common divisor of its coefficients. Meanwhile, if P is primitive but factors in $\mathbb{Q}[x]$, then $P = Q_1 Q_2$. We can rescale Q_1 by taking integers a and b such that aQ_1/b is in $\mathbb{Z}[x]$ and primitive. Then by Lemma 13.5, aQ_1/b must divide P in $\mathbb{Z}[x]$ as well, giving a nontrivial factorization of P. □

Theorem 13.1 for $\mathbb{Z}$ follows as a corollary.

Corollary 13.9
The polynomials with integer coefficients, $\mathbb{Z}[x]$, form a unique factorization domain.

Proof. As usual, we need to prove that a factorization *exists* and is *unique*.

We'll first prove existence. First, by factoring out constants we can write $P = p_1 \cdots p_\ell P_1$ such that P_1 is primitive and the p_i are primes. If P_1 is irreducible in $\mathbb{Q}[x]$, we are done, as it is also irreducible in $\mathbb{Z}[x]$. Otherwise, P_1 factors in $\mathbb{Q}[x]$, and we can rescale both factors so that they are primitive elements of $ZZ[x]$, so then P_1 factors in $\mathbb{Z}[x]$. We can continue to attempt to factor the two resulting factors of P_1. As we keep on factoring, the degrees of our polynomials decrease at every step, so the factorization process must terminate — which means that eventually, all our polynomials become irreducible.

Now we'll prove uniqueness. As in all the other cases where we proved uniqueness, it is enough to show that if an irreducible polynomial $P \in \mathbb{Z}[x]$ divides $Q_1 Q_2$, then P divides either Q_1 or Q_2. (Then similarly to in the proof that every PID is a UFD, given two factorizations $P_1 \cdots P_n$ and $Q_1 \cdots Q_m$, we can show that P_1 must appear in the second factorization as well, cancel it out from both, and repeat with the remaining factorizations until we've matched up all the factors.)

In order to show this result, we have two cases. First, if P is an integer prime $p \in \mathbb{Z}$, then this follows directly from the fact that the product of two nonzero polynomials in $(\mathbb{Z}/p\mathbb{Z})[x]$ is nonzero, as $\mathbb{Z}/p\mathbb{Z}$ is a field.

Otherwise, P is primitive and irreducible in $\mathbb{Q}[x]$. We have that if $P \mid Q_1 Q_2$, then $P \mid Q_1$ or $P \mid Q_2$ in $\mathbb{Q}[x]$ (since $\mathbb{Q}$ is a field, so $\mathbb{Q}[x]$ is a PID and therefore a UFD). By Lemma 13.5, since P is primitive, then P must divide Q_1 or Q_2 in $\mathbb{Z}[x]$.

So this shows that for any irreducible $P \in \mathbb{Z}[x]$, if $P \mid Q_1 Q_2$ then $P \mid Q_1$ or $P \mid Q_2$, as desired. □

As mentioned earlier, the same proof used to show that $\mathbb{Z}[x]$ is a UFD would work if we replaced $\mathbb{Z}$ with *any* UFD R. For example, we can even take R to be $\mathbb{Z}[x]$, now that we know it's a UFD; this shows that $\mathbb{Z}[x, y]$ is also a UFD.

13.2 Gaussian Primes

Unique factorization in $\mathbb{Z}[i]$ has an interesting application — it can be used to solve a problem in number theory.

Guiding Question
Which integers can be written as $n = a^2 + b^2$ for integers a and b?

Example 13.10
We can write $5 = 2^2 + 1^2$, while 6 and 21 cannot be written as a sum of squares.

On the way to proving the answer, we'll classify irreducible elements in $\mathbb{Z}[i]$. (Since $\mathbb{Z}[i]$ is a UFD, the irreducible elements are exactly the primes, so we will use "prime" and "irreducible" interchangeably here.) This is an example of how the abstract property of unique factorization can lead to concrete results.

First, note that $n = a^2 + b^2$ can be rewritten as $n = (a + bi)(a - bi)$. This makes it clear that if n and m can be written in the form $a^2 + b^2$, then so can mn — if $n = \alpha\overline{\alpha}$ and $m = \beta\overline{\beta}$, then $(\alpha\beta)(\overline{\alpha\beta})$. So the property is multiplicative, which motivates considering the special case where n is prime.

Lemma 13.11
Let $p \in \mathbb{Z}$ be a prime number. Then $p = a^2 + b^2$ if and only if p is *not* a prime in $\mathbb{Z}[i]$.

We refer to primes in $\mathbb{Z}[i]$ as **Gaussian primes**.

Proof. First, if p were a Gaussian prime and we could write p as a sum of squares, then we would have

$$p = (a + bi)(a - bi),$$

which would mean p must divide either $a+bi$ or $a-bi$. In either case, p would need to divide both a and b, which is impossible.

Meanwhile, if p is not a Gaussian prime, since it's real and doesn't factor in $\mathbb{Z}$, it must factor as $p = \alpha\overline{\alpha}$ for some $\alpha \in \mathbb{Z}[i]$ which is not in $\mathbb{Z}$. So then $\alpha = a+bi$ for some integers a and b, which means $p = a^2+b^2$. □

So answering our initial question for primes is equivalent to figuring out which integer primes are also Gaussian primes.

Lemma 13.12
Let $p \in \mathbb{Z}$ be a prime number. Then p is *not* prime in $\mathbb{Z}[i]$ if and only if $p = 2$ or $p \equiv 1 \pmod 4$.

Proof. First, 2 factors as $2 = (1+i)(1-i)$. Now suppose $p > 2$.

Claim. *p is not a prime in $\mathbb{Z}[i]$ if and only if there exists $\alpha \in \mathbb{Z}[i]$ such that $p \nmid \alpha$, but $p \mid \alpha\overline{\alpha}$.*

Proof. By definition, p is not a prime in $\mathbb{Z}[i]$ if and only if there exist α and β such that p divides $\alpha\beta$, but not α or β. It immediately follows that if there exists an $\alpha \in \mathbb{Z}[i]$ with the described properties, then p is not prime. On the other hand, if p is not prime, then take α and β such that neither is divisible by p but $\alpha\beta$ is; then

$$p \mid \alpha\beta\overline{\alpha\beta} = (\alpha\overline{\alpha})(\beta\overline{\beta}).$$

Since p is an integer prime, and both $\alpha\overline{\alpha}$ and $\beta\overline{\beta}$ are integers, then p must divide one of them, and either α or β has the described properties. □

This turns the question into one over $\mathbb{F}_p$ — then p is not a prime in $\mathbb{Z}[i]$ if and only if there exist $a, b \in \mathbb{F}_p$, which are not both 0, such that

$$a^2 + b^2 = 0.$$

Since $\mathbb{F}_p$ is a field, we can divide by b^2 and rewrite the equation as $-1 = (ab^{-1})^2$, so this is true if and only if -1 is a square in $\mathbb{F}_p$.

Now consider the abelian group $\mathbb{F}_p^\times$ (the multiplicative group of $\mathbb{F}_p$) which has order $p-1$. The only element of order 2 is -1, since

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

has no roots other than ± 1, and 1 has order 1. This gives a homomorphism $\varphi : \mathbb{F}_p^\times \to \mathbb{F}_p^\times$ sending $\alpha \to \alpha^2$. Then $\ker(\varphi) = \{\pm 1\}$, so by the homomorphism theorem, $\mathrm{im}(\varphi)$ has $(p-1)/2$ elements.

But -1 is a square in $\mathbb{F}_p$ if and only if it is in the image of φ. Since $\mathrm{im}(\varphi)$ is a subgroup of $\mathbb{F}_p^\times$, this occurs if and only if $\mathrm{im}(\varphi)$ contains an element of order 2 (since the only possible element of order 2 is -1).

But $\mathrm{im}(\varphi)$ contains an element of order 2 if and only if $|\mathrm{im}(\varphi)| = (p-1)/2$ is divisible by 2 — one direction follows from the fact that the order of every element divides the order of the group, and the other follows from the Sylow Theorems. So -1 is a square if and only if $(p-1)/2$ is even, or equivalently $p \equiv 1 \pmod 4$.

So an odd integer prime p is *not* a Gaussian prime if and only if $p \equiv 1 \pmod 4$. □

This proof can be used to classify the Gaussian primes up to association (multiplying by units, here ± 1 and $\pm i$).

Theorem 13.13
The full list of primes in $\mathbb{Z}[i]$, up to association, can be constructed as follows: consider all integer primes p.

- If $p \equiv 3 \pmod 4$, then p itself is a Gaussian prime.
- If $p \equiv 1 \pmod 4$, then it factors as $(a-bi)(a+bi)$, and both factors $a \pm bi$ are Gaussian primes.
- If $p = 2$, then it factors as $(1+i)(1-i)$, and since $1+i$ and $1-i$ are associate, they correspond to the same Gaussian prime.

14 Number Fields

14.1 The Gaussian Integers

Last time, we discussed factorization in the Gaussian integers $\mathbb{Z}[i]$, and how it relates to the sum of squares problem. At the end of the lecture, we classified all Gaussian primes. Note that since $\mathbb{Z}[i]$ is a UFD, the primes in $\mathbb{Z}[i]$ are precisely the irreducible elements. In general, primes are defined as elements such that if $p \mid ab$, then $p \mid a$ or $p \mid b$ — but in a UFD, an element is prime if and only if it's irreducible.

Theorem 14.1
The complete list of all primes in $\mathbb{Z}[i]$, up to association, consists of:

- for each integer prime $p = 4k + 3 \in \mathbb{Z}$, the Gaussian prime p itself;
- for each integer prime $p = 4k + 1 \in \mathbb{Z}$, the two Gaussian primes $a \pm bi$ where $a^2 + b^2 = p$;
- the prime $1 + i$.

Note that we have $2 = (1 + i)(1 - i)$, but $1 + i$ and $1 - i$ are associate, so 2 only contributes *one* prime up to association.

Proof. First we'll check that all such elements are primes. For the second and third cases, where p factors as $a^2 + b^2$ for some integers a and b, it's enough to prove the following claim:

Claim. *If $a^2 + b^2 = p$ is an integer prime, then $a + bi$ is prime in $\mathbb{Z}[i]$.*

Proof. Define the norm $\mathrm{N}(a + bi) = a^2 + b^2$, which is multiplicative. Suppose $a + bi$ factors as $\alpha\beta$. Then we have

$$p = \mathrm{N}(a + bi) = \mathrm{N}(\alpha)\,\mathrm{N}(\beta).$$

Then since p is an integer prime, $\mathrm{N}(\alpha)$ or $\mathrm{N}(\beta)$ must be 1. But if $\mathrm{N}(\alpha) = 1$, then $\alpha\overline{\alpha} = 1$, so α is a unit. This means in any factorization of $a + bi$, one factor must be a unit, so $a + bi$ is irreducible and therefore prime. □

Meanwhile, for the first case, we saw last class that every integer prime $p = 4k + 3$ is still prime in $\mathbb{Z}[i]$.

Now we will check that there are no other primes — it's enough to check that every non-unit $\alpha \in \mathbb{Z}[i]$ is divisible by some element of this list. To do so, we again use the norm — if α is not a unit, then we have $\alpha\overline{\alpha} = n$ for some integer $n > 1$. Let p be a prime divisor of n.

Then if $p \equiv 3 \pmod 4$, we must have $p \mid \alpha$ (or $p \mid \overline{\alpha}$, which also implies $p \mid \alpha$), since p is prime. Otherwise, we can write $p = (a + bi)(a - bi)$ where $a \pm bi$ are both primes. Then $a + bi$ must divide α or $\overline{\alpha}$, and therefore $a + bi$ or $a - bi$ must divide α. □

Student Question. *Does this still work when $p = 2$?*

Answer. *Yes, the argument still works as written — in fact, we don't even need the last step, since if $1 + i$ divides $\overline{\alpha}$, then $1 + i$* itself *also divides α.*

As a corollary, we can find the complete answer to the sum of squares question.

Corollary 14.2
If n has prime factorization $n = p_1^{d_1} \cdots p_r^{d_r}$ in $\mathbb{Z}$, then n is a sum of squares if and only if the exponent d_i is even for all primes $p_i \equiv 3 \pmod 4$.

For example, 21 has odd exponents of 3 and 7, so it cannot be written as a sum of squares.

Proof. First, to show that all n of this form work, we've seen earlier that if m and n are sums of squares, then so is mn. For all $p \not\equiv 3 \pmod 4$, we've seen that p is a sum of squares; and for all $p \equiv 3 \pmod 4$, we have that p^2 is trivially a sum of squares. Since n is the product of such terms, it must be a sum of squares as well.

Conversely, suppose n can be written as a sum of squares, so $n = (a + bi)(a - bi)$. Let d be the power of a given prime $p \equiv 3 \pmod 4$ in the prime factorization of $a + bi$ in the Gaussian integers. Since $p \in \mathbb{Z}$, then d is also

the power of p in the prime factorization of $a - bi$, so the power of p in the factorization of n is $2d$. Therefore, the power of each prime $p \equiv 3 \pmod 4$ in the prime factorization of n must be even. □

Student Question. *Why doesn't this statement have a condition involving $p = 2$?*

Answer. *This is because 2 is a sum of squares, as $2 = 1 + 1$. So 2 is allowed to have either odd or even power in the prime factorization of n.*

This is just one example of how such considerations can be applied to number theory. It is possible to go even further — for example, it is possible to determine how many different presentations as a sum of squares there are for a given n.

14.2 Fermat's Last Theorem, as an Aside

The ideas used to analyze solutions to $n = a^2 + b^2$ have some relevance to a more difficult equation as well, Fermat's Last Theorem.

Theorem 14.3 (Fermat)
For an integer $n > 2$, the equation

$$a^n + b^n = c^n$$

has no solutions where a, b, and c are all nonzero integers.

Mathematicians have been trying to prove this theorem for a long time. Fermat famously proposed the theorem in the margin of a book, stating, "I have a truly marvelous demonstration of this proposition that this margin is too narrow to contain." He most likely did not have a truly marvelous demonstration of the proposition.

Mathematician Gabriel Lamé announced a proof on March 1 of 1847. This proof was later found to be incorrect, then partially corrected to be valid in certain cases by Ernst Kummer.

The initial steps of Lamé's proof proceed in a similar fashion as our analysis for Gaussian integers. Assume n is odd (in fact, we can assume that n is prime — it suffices to prove the theorem for $n = 4$ and n prime, and Fermat did have a proof for the case $n = 4$). When considering sums of squares, we used the factorization $a^2 + b^2 = (a + bi)(a - bi)$. In general, when n is odd we have a similar factorization

$$a^n + b^n = (a + b)(a + \zeta b)(a + \zeta^2 b) \cdots (a + \zeta^{n-1} b),$$

where $\zeta = e^{2\pi i/n}$. This gives a factorization of $a^n + b^n$ in the ring $\mathbb{Z}[\zeta]$, the ring of *cyclotomic integers*.

In many cases, it is possible to check that the factors are pairwise coprime; that is, that they do not have a common divisor. In the usual integers, if an nth power is factored as a product of coprime factors, then every factor must be also an nth power (up to multiplication by ± 1). In this case, we similarly want to conclude that each factor is an nth power (up to multiplication by units). This would eventually lead to a contradiction.

The key point is that if $\mathbb{Z}[\zeta]$ were a UFD, then we could obtain our conclusion. This is what Lamé didn't show — we're so used to dealing with the integers, where unique factorization *does* hold, that even a major mathematician initially missed that unique factorization doesn't have to hold in this setting (though it was realized later). Unfortunately, this method doesn't work — when n is an odd prime, $\mathbb{Z}[\zeta]$ is *almost never* a UFD.

However, Kummer showed that under a weaker condition than $\mathbb{Z}[\zeta]$ being a UFD, it's still possible to obtain this conclusion. This weaker condition is that p is a *regular* prime — not every prime is regular, but many are. In order to explain what a regular prime is, we need more theory, which we'll see in future classes. As a glimpse into what this theory will be, when unique factorization fails, we can still analyze *how much* it fails. This leads to the definition of the *ideal class group*, which in some sense controls the non-uniqueness of prime factorization; then the regularity of a prime is a property of the ideal class group of $\mathbb{Z}[\zeta]$.

14.3 Number Fields

Now we will move on to a more general case.

Guiding Question
How does factorization work in rings like $\mathbb{Z}[i]$ and $\mathbb{Z}[\zeta]$?

Both rings sit inside a larger field: $\mathbb{Z}[i] \subset \mathbb{Q}[i]$ and $\mathbb{Z}[\zeta] \subset \mathbb{Q}[\zeta]$. This is helpful because fields can be easier to work with than rings. The concept of a *number field* generalizes this.

Definition 14.4
A **number field** is a subfield in $\mathbb{C}$ that is finite-dimensional as a vector space over $\mathbb{Q}$.

Example 14.5
The number field $\mathbb{Q}[i] = \{a + bi \mid a, b \in \mathbb{Q}\}$ is two-dimensional over $\mathbb{Q}$.

14.3.1 Algebraic Numbers and Integers

An important observation is that every element in a number field is algebraic.

Definition 14.6
A number is an **algebraic number** if it is a root of a polynomial with rational coefficients.

If F is a number field, and $\alpha \in F$, then there is a linear dependence between 1, α, α^2, $\ldots$, α^n for any $n \geq \dim_{\mathbb{Q}}(F)$. So α is a root of some polynomial $P \in \mathbb{Q}[x]$, and therefore α is algebraic. Conversely, $\mathbb{Q}[\alpha]$ is a number field if α is algebraic — if α is the root of a polynomial of degree d, then we can express powers α^i with $i \geq d$ in terms of lower powers, so the only possible terms we have are 1, α, $\ldots$, α^{d-1}. This argument generalizes to show that $\mathbb{Q}[\alpha_1, \ldots, \alpha_n]$ is also a number field if $\alpha_1, \ldots, \alpha_n$ are all algebraic, since there are only finitely many terms $\alpha_1^{e_1} \cdots \alpha_n^{e_n}$ which we need to consider (if a term contains a high power of some α_i, we can rewrite it in terms of lower powers).

If α is algebraic, then $\{P \mid P(\alpha) = 0\}$ is an ideal in $\mathbb{Q}[x]$. Since $\mathbb{Q}[x]$ is a PID, it has the form (P), where P is a monic polynomial of minimal degree. Such a P is called the **minimal polynomial** for α over $\mathbb{Q}$.

A number field is a generalization of fields like $\mathbb{Q}[i]$ or $\mathbb{Q}[\zeta]$, but we really want to analyze factorization in *rings* like $\mathbb{Z}[i]$ or $\mathbb{Z}[\zeta]$, not the underlying fields. To describe the generalization of such rings to an arbitrary number field, we need to define an *algebraic integer*.

Definition 14.7
An algebraic number is an **algebraic integer** if its minimal polynomial has integer coefficients.

Lemma 14.8
The element α is an algebraic integer if and only if $P(\alpha) = 0$ for *some* monic polynomial $P \in \mathbb{Z}[x]$.

It is evident that a polynomial with rational coefficients can be scaled to either be monic or have integer coefficients, but an algebraic integer requires that both can be achieved simultaneously.

Proof. The proof is another application of Gauss's Lemma and the ideas from last lecture. One direction is obvious: if the minimal polynomial P is in $\mathbb{Z}[x]$, then $P(\alpha) = 0$, so the condition is clearly satisfied.

For the other direction, suppose there exists a monic polynomial $P \in \mathbb{Z}[x]$ such that $P(\alpha) = 0$. Now consider the minimal polynomial $P_{\min}$. By clearing denominators and pulling out the gcd of its coefficients, rescale it to a polynomial $Q = aP_{\min}/b$ which is primitive and in $\mathbb{Z}[x]$.

Then Q divides P in $\mathbb{Q}[x]$, since $P_{\min}$ divides P. But Q is primitive, so by the results in the last lecture, Q must also divide P in $\mathbb{Z}[x]$. But then the leading coefficient of Q must divide the leading coefficient of P. Since P is monic, $\pm Q$ must be monic as well. Then since $P_{\min}$ is also monic, we have $Q = \pm P_{\min}$, so $P_{\min}$ has integer coefficients. □

Example 14.9
A rational number $\alpha \in \mathbb{Q}$ has minimal polynomial $x - \alpha$, so α is be an algebraic integer if and only if α is a usual integer.

The next example is the primary setting that we will work with.

Example 14.10 (Quadratic Number Fields)
What are the algebraic integers in the number field $\mathbb{Q}[\sqrt{d}]$? (Here d may or may not be positive.)

Proof. Without loss of generality, we can assume d is squarefree (since factoring squares out of d doesn't change the number field). Then let $\alpha = a + b\sqrt{d}$ with $a, b \in \mathbb{Q}$. The minimal polynomial of α is

$$(x - a - b\sqrt{d})(x - a + b\sqrt{d}) = x^2 - 2a + (a - b^2 d),$$

so α is an algebraic integer if and only if $2a \in \mathbb{Z}$ and $a^2 - b^2 d \in \mathbb{Z}$.

Now we have two cases:

Case 1 ($a \in \mathbb{Z}$). Then we must have $b^2 d \in \mathbb{Z}$ as well, and since d is squarefree, $b \in \mathbb{Z}$ as well.

Case 2 ($a = k + \frac{1}{2}$ for $k \in \mathbb{Z}$). Then we must have $2b \in \mathbb{Z}$ as well, and $b = m + \frac{1}{2}$. In this case, we have

$$a^2 - b^2 d = \frac{1}{4}((2k+1)^2 - (2m+1)^2 d).$$

This is an integer if and only if $d \equiv 1 \pmod 4$.

So the conclusion is that if $d \not\equiv 1 \pmod 4$, the algebraic integers are precisely $a + b\sqrt{d}$ for $a, b \in \mathbb{Z}$ — for example, the algebraic integers in $\mathbb{Q}[i]$ and $\mathbb{Q}[\sqrt{-5}]$ are $\mathbb{Z}[i]$ and $\mathbb{Z}[\sqrt{-5}]$. Meanwhile, if $d \equiv 1 \pmod 4$, the algebraic integers are precisely $a + b\sqrt{d}$ where $a, b \in \mathbb{Z}$ or $a + \frac{1}{2}, b + \frac{1}{2} \in \mathbb{Z}$ — for example, the algebraic integers in $\mathbb{Q}[\sqrt{-3}]$ are $\mathbb{Z}[\omega]$ for a primitive third root of unity ω, which are exactly the elements of this form. □

In general, the algebraic integers of a number field have quite a bit of additional structure.

Theorem 14.11
For a number field F, the set of algebraic integers in F is a subring of F. Furthermore, this is the largest subring that is finitely generated as an abelian group under addition.

We won't prove this in general; but it's fairly straightforward in our specific example $F = \mathbb{Q}[\sqrt{d}]$.

Next time, we will study factorization in rings of algebraic integers. The general idea is that if a statement fails to be true, mathematicians often try to generalize it to find a version that *is* true. In this case, where unique factorization fails for a general ring of algebraic integers, we can consider a more general version of it which works.

To motivate what this more general version will be, note that in the cases where unique factorization works, primes are considered up to association. In fact, an element in a ring $\alpha \in R$ up to association is equivalent to the ideal (α). So in the more general case, we actually consider products of *ideals* rather than *elements*. We will show that in a ring of algebraic integers, there is unique factorization, not into prime elements, but into prime ideals.

15 Ideal Factorization

This class, we'll consider unique factorization in rings of algebraic integers, primarily imaginary quadratic number fields. Unique factorization is a useful property — for example, we used unique factorization in $\mathbb{Z}[i]$ to classify which n can be written as a sum of squares. However, unique factorization doesn't hold in general — for example, we've seen that it fails in $\mathbb{Z}[\sqrt{-5}]$.

But instead, we can prove a weaker statement — instead of unique factorization as a product of prime *elements*, we'll prove unique factorization as a product of prime *ideals*.

15.1 Motivation

The motivation for ideal factorization is that we've seen that unique factorization doesn't hold in a general ring of algebraic integers, so we'd like to modify the property of unique factorization in order to find one that *does* hold. So instead of thinking about products of prime elements, we can think of products of something else — which we can call "ideal prime factors," or "ideal divisors."

We don't yet know what these "ideal divisors" are, but we can think about how they *should* behave. Given an ideal divisor, it should appear in the prime decomposition of various actual numbers. And if it enters the factorization of different numbers, perhaps it should arise as a gcd of different numbers — in an ideal theory where we've restored unique factorization and therefore the existence of a gcd, we would want our ideal divisor to arise as $\gcd(a_1, \ldots, a_n)$ where the a_i are actual elements of the ring.

So along these lines, we could introduce *formal* gcd's of several elements (similarly to how when going from $\mathbb{R}$ to $\mathbb{C}$, we add a formal variable whose square is -1), and figure out how to operate with them. Then given a formal gcd such as $\gcd(a_1, \ldots, a_n)$, we can think about the set of *actual* numbers (in the ring) which are divisible by that gcd. But we've seen that set before — it's the ideal generated by $a_1, \ldots, a_n$! In fact, in the case of a PID (where unique factorization into elements does hold), $\gcd(a_1, \ldots, a_n)$ in the usual sense is the generator of the ideal $(a_1, \ldots, a_n)$.

This was the approach taken by Kummer in the 19th century, when initially developing the concept of ideal factorization. Later, Noether and others realized that a good way to think about this is to declare the gcd to *be* that ideal. In fact, this is where the term "ideal" comes from — we defined them by looking at the kernels of homomorphisms, but the term initially comes from "ideal divisors" and the idea that you can define the ideal divisors as ideals in the sense we've discussed.

Definition 15.1
Given elements $a_1, \ldots, a_n$ in a ring, we define $\gcd(a_1, \ldots, a_n)$ as the ideal $(a_1, \ldots, a_n)$ (meaning the ideal generated by $a_1, \ldots, a_n$).

This definition is where the shorthand (a, b) for $\gcd(a, b)$ comes from as well.

15.2 Prime Ideals

To understand factorization into ideals, we need to understand what the "building blocks" are in such a factorization. A prime *element* is an element p such that if $p \mid ab$, then $p \mid a$ or $p \mid b$. The definition of a prime *ideal* is similar.

Definition 15.2
A **prime ideal** $I \subset R$ is an ideal other than R itself such that whenever $ab \in I$, either $a \in I$ or $b \in I$.

There are a few observations we can make about this definition:

Example 15.3
If I is principal with $I = (a)$, then I is prime if and only if a is prime (by directly applying the definitions).

Lemma 15.4
An ideal I is prime if and only if R/I is an integral domain.

Proof. Recall that an integral domain is a ring where the product of any two nonzero elements must be nonzero. Then this is clear from the definition as well — if we use $\overline{a}$ to denote a mod I, then $\overline{a} = 0$ if and only if $a \in I$. So the definition of a prime ideal states that

$$\overline{ab} = 0 \implies \overline{a} = 0 \text{ or } \overline{b} = 0.$$

But this is exactly the definition of an integral domain. □

Lemma 15.5
A maximal ideal is always prime.

Proof. As we've seen before, an ideal I is maximal if and only if R/I is a field. But all fields are integral domains, so if I is maximal, it must be prime as well (by the above observation). □

Note that by definition, the unit ideal is not prime. Meanwhile, the zero ideal may or may not be prime — in fact, it's prime if and only if R is an integral domain.

15.3 Multiplying Ideals

Of course, to perform factorization using ideals, we also need a way to multiply them.

Definition 15.6
Given two ideals I and J, we define their product as

$$I \cdot J = \left\{ \sum a_i b_i \mid a_i \in I, b_i \in J \right\}.$$

It's clear that multiplication of ideals is commutative and associative. It's also immediate from the definition that if $I = (a_1, \dots, a_n)$ and $J = (b_1, \dots, b_m)$, then IJ is generated by the elements $a_i b_j$.

Example 15.7
If $I = (a)$ is principal, then $IJ = (ab_1, \dots, ab_m)$. In particular, if $I = (a)$ and $J = (b)$, then $IJ = (ab)$ — this means the product of ideals is compatible with the usual product of elements.

Note that IJ is always contained in $I \cap J$ (since I is closed under addition and under multiplication by *any* element of R, then any element $\sum a_i b_i$ must be in I, and the same is true for J). On the other hand, we don't necessarily have $IJ = I \cap J$.

Example 15.8
In $\mathbb{Z}$, if $I = (n)$ and $J = (m)$, then $IJ = (nm)$, while $I \cap J = (\text{lcm}(m, n))$. We always have $\text{lcm}(m, n) \mid mn$, but they're only equal if m and n are relatively prime.

Now that we have defined how to factor with ideals, we are ready to state the main theorem for this section.

Theorem 15.9
Let R be the ring of algebraic integers in a number field. Then every nonzero ideal $I \subset R$ factors uniquely (up to permutation of factors) as a product of prime ideals.

We'll only prove the theorem in the case where the number field F is a *imaginary quadratic field*, but it is true for any number field.

15.4 Lattices

Since we are working with imaginary quadratic number fields, from now we'll assume $F = \mathbb{Q}[\sqrt{d}]$ for an integer $d < 0$. Without loss of generality we may assume d is squarefree. We'll also use R to denote the ring of algebraic

integers in F. As we saw last time, we have

$$R = \begin{cases} \mathbb{Z}[\sqrt{d}] = \left\{a + b\sqrt{d} \mid a, b \in \mathbb{Z}\right\} & \text{if } d \not\equiv 1 \pmod 4 \\ \mathbb{Z}\left[\frac{1+\sqrt{d}}{2}\right] = \left\{a + b\sqrt{d} \mid a, b \in \frac{1}{2}\mathbb{Z} \text{ and } a + b \in \mathbb{Z}\right\} & \text{if } d \equiv 1 \pmod 4. \end{cases}$$

In either case, we can think of R as a *lattice* in $\mathbb{C}$ (a lattice is the additive subgroup of $\mathbb{C}$ generated by two non-collinear vectors), generated by 1 and $\sqrt{d}$ in the first case, and 1 and $(1+\sqrt{d})/2$ in the second. For example, $\mathbb{Z}[i]$ is the square lattice (and more generally, in the first case, the lattice is always rectangular):

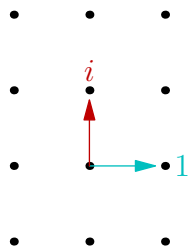

We'll first state a few elementary properties of lattices that will be useful:

Proposition 15.10
Suppose that L and L' are lattices, with $L' \subset L$.

- The quotient L/L' is finite.
- If L'' is a subgroup of L (under addition) with $L' \subset L'' \subset L$, then L'' is also a lattice.

The proof is left as an exercise; it's possible to see this by thinking about the example $\mathbb{Z}^2$, since these properties don't depend on which lattice L is.

Corollary 15.11
Every nonzero ideal of R is again a lattice.

Proof. First, this is clear for *principal* ideals — if $I = (\alpha)$, then we can write $I = \alpha R$, so I is obtained by multiplying the lattice of R by α. It's clear from the definition that multiplying a lattice by a complex number will still produce a lattice; it's also possible to see this geometrically, since multiplication by a complex number is just a rotation and dilation.

But if I is an arbitrary nonzero ideal, then $R \supset I \supset \alpha R$ for each nonzero $\alpha \in I$, so by Proposition 15.10, since I sits between two lattices, I must itself be a lattice. □

Note 15.12
As we'll see later, trying to understand how these lattices look geometrically (up to similarity — multiplication by a complex number) gives rise to an important number theoretic concept.

15.5 Proof of Unique Factorization

To prove uniqueness of ideal factorization for $R \subset \mathbb{Q}[\sqrt{d}]$, we will first make a few observations.

Lemma 15.13
A nonzero ideal in R is prime if and only if it is maximal.

We've already seen that all maximal ideals are prime; in general, the converse is false, but in the situation here (and more generally, in rings of algebraic integers in a number field) it turns out to be true.

Proof. It's enough to show that every prime ideal is maximal. But note that R/I is finite for every nonzero ideal, since I and R are lattices (by Proposition 15.10). Additionally, since I is prime, R/I is an integral domain — so there are no zero divisors.

But in a *finite* ring S, any element a which is not a zero divisor is necessarily invertible — this can be proven by a counting argument. Consider the list of values ab for all $b \in S$; these must all be distinct, as if $ab = ac$, then we would have $a(b-c) = 0$. But there are $|S|$ such values, so they must cover *all* of S; and in particular, there exists b with $ab = 1$. □

Student Question. *Does this mean that a finite ring of prime order is a field?*

Answer. *Yes — the only ring of order p is $\mathbb{Z}/p\mathbb{Z}$, which is a field. This doesn't generalize to prime powers, though — we'll later see that for each prime power p^k, there's one field with order p^k, but there are other rings with order p^k.*

Student Question. *In general, if we have a prime ideal I for which R/I is finite, then do we know I is also maximal?*

Answer. *Yes. In fact, there's also a generalization of this lemma which replaces "finite" with "finite-dimensional vector space."*

The key proposition, from which most of the proof follows formally, is the following:

Proposition 15.14
Multiplication of ideals has the *cancellation property* — if we have ideals I, I', and J (with $J \neq 0$), then

$$IJ = I'J \implies I = I'.$$

Furthermore, divisibility is the same as inclusion — if $I \subset J$, then there exists an ideal J' such that $I = JJ'$.

It's clear that multiplication of ideals gives us a smaller ideal; the second statement tells us that here, the converse is true as well.

To prove these two properties, we'll first establish the following key lemma. This lemma will essentially allow us to reduce to the case of principal ideals, which are easier to work with.

Lemma 15.15
If $I \subset R$ is an ideal, then $I\overline{I}$ is a principal ideal generated by an integer $n \in \mathbb{Z}$.

Here $\overline{I} = \{\overline{z} \mid z \in I\}$ — it's clear that this is also an ideal.

Proof. Since I is a lattice, we can pick two elements α and β which generate I as a lattice. Then they also generate I as an ideal; this means $I = (\alpha, \beta)$ and $\overline{I} = (\overline{\alpha}, \overline{\beta})$. This means

$$I\overline{I} = (\alpha\overline{\alpha}, \beta\overline{\beta}, \alpha\overline{\beta}, \beta\overline{\alpha}).$$

Note that $\alpha\overline{\alpha}$, $\beta\overline{\beta}$, and $\alpha\overline{\beta} + \beta\overline{\alpha}$ are all integers; so we can define n to be their gcd, in the sense of the usual integers.

Then we claim that $I\overline{I} = (n)$. It's clear that $n \in I\overline{I}$, since n is in the smaller ideal $(\alpha\overline{\alpha}, \beta\overline{\beta}, \alpha\overline{\beta} + \beta\overline{\alpha}) \subset I$. So it suffices to check that n generates the entire ideal, or equivalently that (n) contains all the generators of I; and since we already know that n divides $\alpha\overline{\alpha}$, $\beta\overline{\beta}$, and $\alpha\overline{\beta} + \beta\overline{\alpha}$, it's enough to check that n divides $\alpha\overline{\beta}$.

To do so, we'll check that $\alpha\overline{\beta}/n$ is an algebraic integer, which will imply that it's in R; it suffices to check that it's the root of a monic polynomial with integer coefficients.

But we can take

$$P = \left(x - \frac{\alpha\overline{\beta}}{n}\right)\left(x - \frac{\overline{\alpha}\beta}{n}\right) = x^2 - \frac{\alpha\overline{\beta} + \overline{\alpha}\beta}{n} \cdot x + \frac{\alpha\overline{\alpha} \cdot \beta\overline{\beta}}{n}.$$

By the definition of n, both coefficients are integers, so we are done. □

16 Uniqueness of Ideal Factorization

16.1 Properties of Ideal Multiplication

Last time, we began discussing the uniqueness of factorization into ideals in the ring of algebraic integers of an imaginary quadratic field. We let $F = \mathbb{Q}[\sqrt{d}]$ for a squarefree integer $d < 0$, and we saw that then the ring of algebraic integers in F is

$$R = \begin{cases} \mathbb{Z}[\sqrt{d}] = \left\{a + b\sqrt{d} \mid a, b \in \mathbb{Z}\right\} & \text{if } d \not\equiv 1 \pmod 4 \\ \mathbb{Z}\left[\frac{1+\sqrt{d}}{2}\right] = \left\{a + b\sqrt{d} \mid a, b \in \frac{1}{2}\mathbb{Z} \text{ and } a + b \in \mathbb{Z}\right\} & \text{if } d \equiv 1 \pmod 4. \end{cases}$$

We then proved the key lemma:

Lemma 16.1
If $I \subset R$ is an ideal, then $I\overline{I} = (n)$ for some $n \in \mathbb{Z}$.

We can assume $n > 0$ (since n and $-n$ generate the same ideal), giving the following definition:

Definition 16.2
The **norm** of an ideal I, denoted $\mathrm{N}(I)$, is the positive integer n with $I\overline{I} = (n)$.

Note that when $I = (\alpha)$ is principal, then $I\overline{I} = (\alpha\overline{\alpha})$, so $\mathrm{N}(I) = \alpha\overline{\alpha}$. So this coincides with the usual concept of the norm of a complex number. It's also clear from the definition that N is multiplicative, meaning that $\mathrm{N}(IJ) = \mathrm{N}(I)\,\mathrm{N}(J)$ — this is because

$$IJ \cdot \overline{IJ} = I\overline{I} \cdot J\overline{J}.$$

Last time, we also stated the following key proposition:

Proposition 16.3
Multiplication of ideals has the following two properties:

(a) Cancellation: if $IJ = I'J$ and $J \neq 0$, then $I = I'$.

(b) If $I \subset J$, then $I = JJ'$ for some J'.

As stated last time, the second condition is actually an *if and only if* statement — we've seen already that if $I = JJ'$, then $I \subset J$.

Student Question. *Is this only true for imaginary quadratic number fields?*

Answer. *It's true for other rings of algebraic integers; it's also true for some other examples, such as a ring of polynomials in one variable (which we'll discuss next class) and its extensions; but it doesn't hold in general. For example, it doesn't hold for polynomials in two variables.*

Proof of Proposition 16.3. The main idea is that in the case where J is principal, this is almost obvious; and using Lemma 16.1, we can reduce to the case where J is principal.

First consider the case where J is principal. Then for (a), it's clear that

$$\alpha I = \alpha I' \implies I = I'.$$

This is because multiplying by α is easily inverted, just by multiplying by α^{-1}, if we think of these sets as belonging to the field F and not the ring R.

Similarly, for (b), if we have $I \subset (\alpha)$, then $x/\alpha \in R$ for all $x \in I$. Then we can take

$$J' = \frac{I}{\alpha} = \left\{\frac{x}{\alpha} \mid x \in I\right\},$$

which is an ideal of R.

Now consider the more general case. For (a), if we have $IJ = I'J$, then we can multiply by $\overline{J}$ to get

$$I(J\overline{J}) = I'(J\overline{J}) \implies I(n) = I'(n),$$

where $n = \mathrm{N}(J)$. Since $n \neq 0$ if $J \neq 0$, then by the principal case of (a), we have that $I = I'$ in the general case as well.

For (b), we use the same trick. Suppose $I \subset J$, and multiply both sides by $\overline{J}$. Then

$$I\overline{J} \subset J\overline{J} = (n).$$

As before, now set $J' = I\overline{J}/n$ (which is an ideal for the same reason as in the principal case). Then

$$J'(J\overline{J}) = J'(n) = I\overline{J},$$

and by (a) we can cancel out $\overline{J}$ to get $J'J = I$. □

Student Question. *For rings of general algebraic integers, can we get a principal ideal by multiplying by* all *conjugates?*

Answer. *Yes — if the field is Galois, then we can multiply by all the Galois conjugates. But we haven't yet learned the relevant theory.*

Finally, before proving unique factorization, we'll need the following lemma (which will essentially allow us to use (b) to find a prime ideal dividing *any* ideal, so that we can factor repeatedly):

Lemma 16.4
Every non-unit ideal I in R is contained in a maximal ideal.

This lemma is actually true for *any* ring. But we'll only prove it in our case (since the proof is harder in general).

Proof. Assume without loss of generality that $I \neq 0$, since the zero ideal is contained in *any* ideal. Then we know that R/I is finite — this means there are finitely many ideals in R/I.

But ideals in R/I are in bijection with ideals in R containing I, by the correspondence theorem for rings. And since R/I is finite, it must have a maximal ideal — if we have an ideal that isn't maximal (or the unit ideal), we can find a bigger one, and we can't keep on doing this forever since there's only finitely many ideals. □

Student Question. *Why is R/I finite?*

Answer. *As discussed last time, this is because R and I are both lattices in $\mathbb{C}$, and the index of any lattice in another is finite. In fact, for any lattice, we can take two vectors which generate it, and compute the area of the parallelogram formed by those two vectors:*

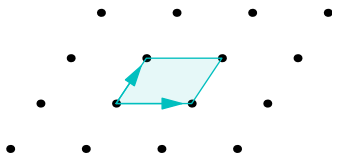

Then the index of I in R is the ratio of the areas of their corresponding parallelograms.

Note 16.5
Lemma 16.4 is true in any ring. But in this proof, we used the fact that the poset of ideals in R/I under inclusion is finite, and therefore has a maximal element; this isn't true for a general poset — for example, the positive integers don't have a maximal element.

But the poset of ideals has the additional property that any increasing chain of ideals is majorated by an element of the set — if we have a chain $I_1 \subset I_2 \subset I_3 \cdots$, their union $I_1 \cup I_2 \cup \cdots$ is again an ideal, which contains all the I_n.

Then it's possible to use Zorn's Lemma from set theory, which applies to partially ordered sets with this property. Although Zorn's Lemma is needed in the general case, we'll see a weaker finiteness property in which case there *is* an easier proof. That finiteness property will be much more general than the one used here — in particular, it will apply to polynomial rings and their quotients.

16.2 Proof of Unique Factorization

Finally, we are ready to prove the uniqueness of ideal factorization for imaginary quadratic number fields:

Theorem 16.6
Every nonzero ideal $I \subset R$ factors uniquely (up to permutation of factors) as a product of prime ideals.

Proof. First, we'll prove the existence of a prime factorization. Let $I \subset R$ be an ideal which is neither the zero ideal nor the unit ideal. Then by Lemma 16.4, there is a maximal ideal $\mathfrak{m}$ with $I \subset \mathfrak{m}$, and therefore by Proposition 16.3, we can factor $I = \mathfrak{m} \cdot J$ for some ideal J. Since $\mathfrak{m}$ is maximal, it is also prime; so we can think of this as the first step in the factorization process, and now it suffices to factor J (unless J is the unit ideal, in which case we're done).

It then suffices to show that this factorization process terminates. To do so, we can consider the norm — we have

$$\mathrm{N}(I) = \mathrm{N}(\mathfrak{m}) \cdot \mathrm{N}(J).$$

We must have $\mathrm{N}(\mathfrak{m}) > 1$ (since $1 \notin \mathfrak{m}$), so then $\mathrm{N}(J) < \mathrm{N}(I)$. This means the norm of our ideal decreases, so the process must eventually terminate, meaning that J must eventually become the unit ideal. (This argument can also be phrased by induction on the norm.)

Now we'll prove uniqueness of factorization. Suppose

$$I = P_1 \cdots P_n = Q_1 \cdots Q_n,$$

where the P_i and Q_i are all ideals. It's enough to show that $P_1 = Q_i$ for some i, since then by part (a) of Proposition 16.3, we can cancel out the common factor (this is the same way we proved uniqueness in the case of integers or polynomials, except that now we're dealing with ideals instead of elements).

Assume for contradiction that $Q_i \neq P_1$ for all i. Then $Q_i \not\subset P_1$ for all i — since Q_i is maximal, the only ideals containing it are itself and the unit ideal. So for each i, we can find an element $x_i \in Q_i$ with $x_i \notin P_1$. But now consider their product $x_1 x_2 \cdots x_n$. By definition, this product lies in $Q_1 \cdots Q_n$. But since P_1 is prime and none of the x_i are in P_1, their product cannot be either; contradiction. □

The last step can instead be worded using the following lemma:

Lemma 16.7
If P is a prime ideal, and I and J are ideals with $I \not\subset P$ and $J \not\subset P$, then $IJ \not\subset P$.

Proof. The same proof works — pick $x \in I$ and $y \in J$ with $x, y \notin P$. Then we have $xy \in IJ$, but $xy \notin P$. □

16.3 Classification of Prime Ideals

We can look more concretely at the structure of these prime ideals. In a future class, we'll see that the classification of prime ideals actually works quite similarly to the classification of primes we saw in the Gaussian integers — by looking at all integer primes p, and trying to factor (p) into prime ideals. We'll see that there are three possibilities:

- (p) remains prime in R — such primes are called **inert primes**;
- (p) factors as $Q_1 Q_2$, where Q_1 and Q_2 are distinct (they must then be conjugate) — such primes are called **splitting primes**;
- (p) factors as Q^2 — such primes are called **ramified primes**. (In the case of Gaussian integers, the only ramified prime was 2; in general, the ramified primes come from divisors of d.)

All prime ideals in this list are distinct, and this gives a full list of the prime ideals.

16.4 Similarity Classes of Ideals

We'll now discuss another important concept regarding ideals.

Definition 16.8
Two nonzero ideals $I, J \subset R$ are **similar** if there exists some $\lambda \in F$ such that $\lambda I = J$.

(We only consider nonzero ideals when discussing similarity.)

Note that it would be equivalent to state $\lambda \in \mathbb{C}$ in the definition, since if $\lambda \in \mathbb{C}$ were the quotient of two elements in R, then we must have $\lambda \in F$. Meanwhile, two lattices L_1 and L_2 in $\mathbb{C}$ are similar if $L_2 = \lambda L_1$ for some $\lambda \in \mathbb{C}$ — geometrically, multiplication by a complex number corresponds to scaling and rotation. So the algebraic notion of similarity of ideals coincides with the geometric notion of similarity of their lattices.

Similarity is an equivalence relation — if $I_2 = \lambda I_1$ and $I_3 = \mu I_2$, then $I_3 = \lambda\mu I_1$. So we will use the notation $I \sim J$ to denote that two ideals are similar. Then we can think about ideals in terms of their equivalence classes, which are called **ideal classes**.

Example 16.9
The similarity class of the unit ideal (1) consists of ideals $I = \lambda(1)$ for $\lambda \in F$. But since $\lambda \cdot 1 \in R$, then we must have $\lambda \in R$. So then $I = \lambda(1) = (\lambda)$ is a principal ideal — so the similarity class of (1) consists of exactly the principal ideals.

In particular, this example implies that in a PID, all ideals are similar. But there are many cases which are *not* PIDs — and in some sense, the number of equivalence classes is a measure of the ring's failure to be a PID.

Proposition 16.10
If $I \sim I'$, then $IJ \sim I'J$.

This is straightforward from the definitions, but it's important — it means that taking the product of ideals gives a commutative and associative operation on the set of ideal classes. In fact, the set of ideal classes, along with ideal multiplication, is an abelian group — the class of (1) is the unit, and every nonzero class is invertible because $I\overline{I}$ is principal, so the class of $\overline{I}$ is the inverse of the class of I. This group is called the **ideal class group**, and denoted $\mathrm{Cl}(F)$.

An important theorem about the ideal class group, which we'll look at in more detail later, is the following:

Theorem 16.11
The ideal class group $\mathrm{Cl}(F)$ is finite.

Example 16.12
For $R = \mathbb{Z}[\sqrt{-5}]$, the class group is $\mathbb{Z}/2\mathbb{Z}$ — the only two ideals up to similarity are (1) and $(2, 1 + \sqrt{-5})$.

17 Ideals in Quadratic Fields

17.1 Prime Ideals

Guiding Question
What are the prime ideals in $R \subset \mathbb{Q}[\sqrt{d}]$?

The answer is similar to the case $R = \mathbb{Z}[i]$, which we've seen earlier.

Lemma 17.1
If P is a prime (nonzero) ideal in R, then either $P = (q)$ for an integer prime q, or $P\overline{P} = (q)$ for an integer prime q.

Proof. Suppose $P\overline{P} = (n)$, where n is not prime. Then we can factor $n = ab$ with $a, b \neq \pm 1$. But then by unique ideal factorization (since P and $\overline{P}$ are prime), it follows that $(a) = P$ and $(b) = \overline{P}$, or vice versa. Then we must have $a = b = q$ for a prime q. □

Lemma 17.2
An odd integer prime q remains prime in R if and only if the equation $\overline{a}^2 = d\overline{b}^2$ has no solutions in $\mathbb{F}_q$ except $(0, 0)$, or equivalently, if d is neither 0 nor a square mod q.

Proof. First, suppose (q) is prime, and there exists a solution in $\mathbb{F}_q$ to $\overline{a}^2 = d\overline{b}^2$. Then

$$q \mid (a - b\sqrt{d})(a + b\sqrt{d}).$$

Since q is prime, then q must divide one of $a \pm b\sqrt{d}$, and therefore $q \mid a$ and $q \mid b$.

Conversely, if q is not prime, then we can find α and β such that $q \mid \alpha\beta$ but q does not divide either α or β. But since $q \mid \alpha\beta$, we have

$$q \mid \mathrm{N}(\alpha\beta) = \mathrm{N}(\alpha)\mathrm{N}(\beta).$$

So since q is an integer prime, it must divide one of the factors on the right; without loss of generality, $q \mid \mathrm{N}(\alpha)$. Now write $\alpha = a + b\sqrt{d}$, so we get that

$$q \mid a^2 - db^2.$$

(It's possible that a and b are half-integers instead of integers; if so, we can replace them with $(2a, 2b)$ without affecting the rest of the argument.) Since q doesn't divide α in R, then q cannot divide both a and b; so $(a, b) \neq (0, 0)$ is a solution to $\overline{a}^2 = d\overline{b}^2$. □

Student Question. *In the first lemma, how did we conclude that P and $\overline{P}$ are (a) and (b)?*

Answer. *We know that (n) is the product of two primes, P and $\overline{P}$. But then the only way to factor it (where neither factor is the unit ideal) is as the product of those two primes — and since it also factors as $(a)(b)$, then (a) and (b) must be those two primes. As an analogy in $\mathbb{Z}$, the only way to factor 6 is as $2 \cdot 3$.*

Student Question. *Did we use the fact that d was negative here?*

Answer. *Somewhat — our proof involved complex conjugation, which relies on d being negative. But if we modify our operation of conjugation, then this analysis works for real quadratic fields as well — we'll discuss this later today.*

Student Question. *Why is $q = 2$ a special case?*

Answer. *It has to do with the fact that if $d \equiv 1 \pmod 4$, then we may have half-integers, such as $(1 + \sqrt{d})/2$. In the case where q was odd, this didn't really matter; but we have to be more careful when $q = 2$. In particular, in the first direction knowing that 2 divides $a + b\sqrt{d}$ in R does not necessarily imply that 2 divides a and b in $\mathbb{Z}$.*

Combining these two results, we get a full list of primes in R:

- For each integer prime $q \nmid d$ where d is not a square mod q (equivalently, there are no solutions to $\overline{a}^2 = d\overline{b}^2$), we get the prime q itself.
- For all other primes, we can factor $q = P\overline{P}$. In most cases, P and $\overline{P}$ are distinct, and this gives two prime ideals. But there's finitely many ramification primes where $P = \overline{P}$, and we get one prime ideal. In fact, these ramification primes are the divisors of d, along with 2 if $d \not\equiv 1 \pmod 4$.

17.2 The Ideal Class Group

Previously, we introduced the ideal class group $\mathrm{Cl}(F)$, which consists of the nonzero ideals in R up to similarity.

Theorem 17.3
The ideal class group $\mathrm{Cl}(F)$ is finite.

We'll prove this in a future class; but for now, we'll look at a few examples.

Example 17.4
In the cases of $\mathbb{Z}[i] \subset \mathbb{Q}[i]$ and $\mathbb{Z}[\omega] \subset \mathbb{Q}[\sqrt{-3}]$ (where ω is a primitive third root of unity), the class group is trivial, since both $\mathbb{Z}[i]$ and $\mathbb{Z}[\omega]$ are PIDs. But in fact, there are finitely many negative d for which $\mathrm{Cl}(\mathbb{Q}[\sqrt{d}])$ is trivial.

For a more interesting example:

Lemma 17.5
In the case of $\mathbb{Z}[\sqrt{-5}] \subset \mathbb{Q}[\sqrt{-5}]$, the class group is $\mathbb{Z}/2\mathbb{Z}$. The two similarity classes of ideals are represented by (1), and by $(2, 1+\sqrt{-5})$.

Proof. Let I be a nonzero ideal, and let α be a nonzero element of I with minimal norm. Let L be the lattice generated by α and $\sqrt{-5}\alpha$, or equivalently, the lattice corresponding to the ideal (α); then $L \subset I$.

If $L = I$, then I is principal. Now assume $L \neq I$, so there is an element $\beta \in I$ with $\beta \notin L$. Without loss of generality, we can assume that β is in the rectangle spanned by α and $\sqrt{-5}\alpha$ — otherwise, we can subtract multiples of α and $\sqrt{-5}\alpha$ to translate β into this rectangle)

Claim. *We must have* $\beta = \alpha \cdot (1+\sqrt{-5})/2$.

Proof. Let $r = |\alpha|$. Then split the rectangle into smaller rectangles as shown (note that the rectangle will not generally be horizontal, but it is drawn this way for convenience):

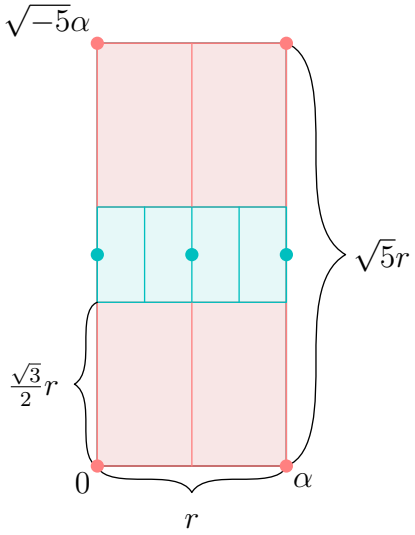

By straightforward calculation, it can be shown that every point in one of the red rectangles is at a distance less than r from one of the red dots, corresponding to 0, α, $\sqrt{-5}\alpha$, and $(1+\sqrt{-5})\alpha$. Meanwhile, every point in one of the blue rectangles is at a distance less than $r/2$ from one of the blue dots, corresponding to $\sqrt{-5}\alpha/2$, $(1+\sqrt{-5})\alpha/2$, and $(2+\sqrt{-5})\alpha/2$.

In the first case, let γ be this red dot. Then $\beta \in I$ and $\gamma \in I$, so $\beta - \gamma \in I$ as well. But we have $|\beta - \gamma| < |\alpha|$, contradiction.

In the second case, again let γ be this blue dot. Then $\beta \in I$ and $2\gamma \in I$, so $2\beta - 2\gamma \in I$ as well. But we have $|2\beta - 2\gamma| < |\alpha|$, which is again a contradiction unless $2\beta - 2\gamma = 0$.

Now suppose $\beta = \gamma$. We now claim that β can't be either of the two blue dots on the ends — either case would imply that $\sqrt{-5}\alpha/2 \in I$. But multiplying by $\sqrt{-5}$ would give that $-5\alpha/2 \in I$, and therefore $\alpha/2 \in I$; this would contradict the choice of α as the element of smallest length.

So then β must be the dot in the center, which is $(1+\sqrt{-5})\alpha/2$. □

So in this case, there is only one element of I inside this rectangle, which is $\beta = \alpha \cdot (1+\sqrt{-5})/2$; then

$$I = (\alpha, \beta) = \frac{\alpha}{2}(2, 1+\sqrt{-5}).$$ □

Here, we saw a proof that there's only two ideals in $\mathrm{Cl}(\mathbb{Q}[\sqrt{-5}])$ up to similarity which looked geometrically at lattices. In fact, the proof of finiteness in general *also* involves looking at lattices. We'll see this proof next class; but today we'll conclude by looking at a few generalizations, where we consider similar questions in fields similar to imaginary quadratic number fields.

17.3 Real Quadratic Number Fields

So far, we've discussed the case $F = \mathbb{Q}[\sqrt{d}]$ when $d < 0$.

Guiding Question
What if we instead have $\mathbb{Q}[\sqrt{d}]$ with $d > 0$?

We can then write $F = \mathbb{Q}[\delta]$, where $\delta^2 = d$.

This case is quite similar, but there are some differences. First, it doesn't make sense to talk about complex conjugation, since all our numbers are real. But there is still a type of "conjugation" in F — the map $a + b\delta \to a - b\delta$. (This is an example of a general construction that we'll discuss in a much later class, related to the Galois group.) This is still a field automorphism.

Moreover, R is not a lattice in $\mathbb{C}$, but it can be embedded as a lattice in $\mathbb{R}\times\mathbb{R}$, a ring where addition and multiplication are done componentwise. To perform this embedding, we send

$$a + b\delta \mapsto (a + b\sqrt{d}, a - b\sqrt{d}) \in \mathbb{R}^2.$$

In fact, the reason we're using the notation with δ is because we can think of it as an abstract square root of d — then we can write it as the *positive* square root $\sqrt{d}$, or the *negative* square root $-\sqrt{d}$.

Another important difference is that in imaginary quadratic fields, there are very few units. However, in this case, the group of units is infinite — there are infinitely many solutions to $a^2 - b^2 d = 1$ (which is known as a Pell equation).

Although there are some differences, our arguments used in imaginary quadratic fields mostly work here as well. In principle, those arguments can be generalized to rings of algebraic integers in *arbitrary* number fields; but that generalization is more difficult and requires theory we have not yet discussed.

17.4 Function Fields

There is a second generalization we'll discuss.

Guiding Question
What if we replace $\mathbb{Z}$ by $k[t]$ for a field k?

Then we replace $\mathbb{Q}$ with $k(t)$, the field of rational functions in t over the field k — in other words, $k(t) = \text{Frac}(k[t])$.

In this case, we consider fields F containing $k(t)$ which are finite-dimensional over $k(t)$, similarly to how a number field is a field containing $\mathbb{Q}$ which is finite-dimensional over $\mathbb{Q}$. We then have

$$R = \{\alpha \in F \mid P(\alpha) = 0 \text{ for a monic } P \in k[t][x]\}.$$

Here P is a polynomial in *two* variables, but it's supposed to be monic as a polynomial in x. This is again similar to how in the number field setting, where R is the set of $\alpha \in F$ which are roots of a monic polynomial in $\mathbb{Z}[x]$.

The especially relevant examples are $k = \mathbb{C}$ and $k = \mathbb{F}_p$. We'll focus on the case $k = \mathbb{C}$.

In the case of imaginary quadratic number fields, we looked at primes in $\mathbb{Z}$ and analyzed how they factored in R — we can try to perform a similar analysis here.

First, as we've seen earlier, $k[t]$ is a PID for any k. So in the case $k = \mathbb{C}$, the primes in $\mathbb{C}[t]$ are exactly $(t - \lambda)$ for $\lambda \in \mathbb{C}$, since the only irreducible polynomials in $\mathbb{C}$ are linear.

Meanwhile, to describe the nonzero primes in R, we can use Hilbert's Nullstellensatz. It's possible to write R as a quotient of $\mathbb{C}[t, t_2, \ldots, t_n]$ by polynomials — for example, in the case of quadratic number fields, we have $\mathbb{Z}[\sqrt{d}] = \mathbb{Z}[x]/(x^2 - d)$, and it's possible to perform a similar construction here. We still have unique factorization into ideals in R, and in particular, prime ideals are maximal; so the nonzero primes in R correspond to points in $\text{MSpec}(R)$, which we can think of as a subset of $\mathbb{C}^n$ by Nullstelensatz.

Now this gives a ramified covering, where we can cover $\mathbb{C}$ by $\text{MSpec}(R) \subset \mathbb{C}^n$. More explicitly, if we write R as a quotient of $\mathbb{C}[t, t_2, \ldots, t_n]$, then the elements of $\text{MSpec}(R)$ correspond to points in $\mathbb{C}^n$ which are roots of all the polynomials we quotient out by. In this ramified covering, a point $t \in \mathbb{C}$ lies below all the points in $\text{MSpec}(R)$ with that value of t.

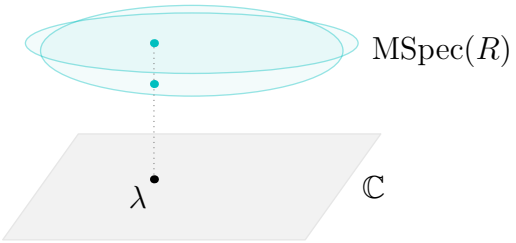

Factoring $(t - \lambda)$ as a product of prime ideals in R then amounts to enumerating the points in the pre-image of λ in this ramified covering. (We'll see this in more detail next class.)

The term *ramified* is used in a similar sense here as when discussing ramified primes. In this setting, if for example our ramified cover corresponds to the map $z \mapsto z^2$, then most points in $\mathbb{C}$ have *two* points in their pre-image (since there's two square roots), but 0 only has one — so 0 is a ramification point. This is similar to how in the number field setting, when we factor (q) as a product of prime ideals in R, usually these prime ideals are distinct, but they're the same ideal for the ramified primes.

A famous mathematician, André Weil, proposed a metaphor between this situation and the Rosetta Stone. The Rosetta Stone contained a script written in three languages. Here our languages are the finite extensions of $\mathbb{Q}$ (or in other words, number fields), finite extensions of $\mathbb{F}_q(t)$ for a finite field $\mathbb{F}_q$, and finite extensions of $\mathbb{C}(t)$. In all of these situations, it's possible to consider how normal primes (in $\mathbb{Z}$, $\mathbb{F}_q[t]$, and $\mathbb{C}[t]$ respectively) factor as a product of ideals in the corresponding ring R. There are analogies between the three settings, and Weil wondered how results in each setting could be "translated" to the others.

18 The Ideal Class Group

18.1 Review — Function Fields

Last time, we discussed a generalization where we replace $\mathbb{Q}$ and $\mathbb{Z}$ with $k(t)$ and $k[t]$, for a field k — instead of working with finite extensions of $\mathbb{Q}$ (or number fields), we work with finite extensions of $k(t)$ (or function fields). For concreteness, we'll focus on $k = \mathbb{C}$.

In order to keep the same level of generality as we had when working with number fields, we will take F to be a *quadratic* extension of $\mathbb{C}(t)$, so

$$F = \mathbb{C}(t)[z]/(z^2 - P(t))$$

for some polynomial $P(t)$, which we may without loss of generality assume to be squarefree. Here $P(t)$ plays the same role as d did when we considered quadratic *number* fields $\mathbb{Q}[\sqrt{d}]$, which could also be described as $\mathbb{Q}[x]/(x^2 - d)$. In this setting, the analog of the ring of algebraic integers, which was $\mathbb{Z}[\sqrt{-d}] = \mathbb{Z}[x]/(x^2 - d)$ in the quadratic number field setting, is

$$R = \mathbb{C}[t][z]/(z^2 - P(t)).$$

In this setting, unique factorization into ideals still holds, although we will not discuss the proof:

Theorem 18.1
Every ideal in R can be factored uniquely as a product of prime ideals.

In particular, all prime ideals are maximal — in fact, this can be proven using a similar argument to the one we saw for number fields, but using that R mod an ideal is finite-dimensional as a $\mathbb{C}$-vector space, rather than that it is finite.

In number fields, we described how prime ideals (q) in the original ring $\mathbb{Z}$ factor as a product of ideals in the new ring R. We can ask the same question in this setting as well.

Guiding Question
How do the prime ideals in $\mathbb{C}[t]$ factor as a product of prime ideals in R?

In both cases, the prime ideals are exactly the maximal ones. Describing the maximal ideals in $\mathbb{C}[t]$ is simple — the only irreducible polynomials in $\mathbb{C}[t]$ are linear, so these prime ideals are of the form $(t - \lambda)$ for $\lambda \in \mathbb{C}$.

Meanwhile, to understand the maximal ideals in R, we can use Nullstelensatz! So far we have been thinking of R as a quadratic extension, but we could instead think of it as the quotient of a two-variable polynomial ring — we have $R = \mathbb{C}[t, z]/(z^2 - P(t))$. Hilbert's Nullstelensatz tells us that the maximal ideals of $\mathbb{C}[t, z]$ are exactly $\mathfrak{m}_{a,b} = (t - a, z - b)$ for $(a, b) \in \mathbb{C}$, so the maximal ideals of R are exactly those $\mathfrak{m}_{a,b}$ for which $b^2 = P(a)$.

So maximal ideals in $\mathbb{C}[t]$ are in bijection with $\mathbb{C}$, and maximal ideals in R are in bijection with solutions in $\mathbb{C}^2$ to $b^2 = P(a)$. This gives a **ramified double cover** — each value of a corresponds to two values of b, except the roots of P, which correspond to one value of b. In the case of imaginary quadratic number fields, we had complex conjugation; the analog in this setting is the map $b \mapsto -b$.

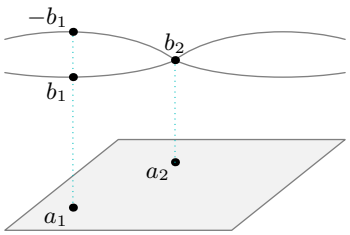

Now that we have a description of what the maximal (and therefore prime) ideals in $\mathbb{C}[t]$ and R are, we can answer our question. A prime ideal $(t - \lambda) \subset \mathbb{C}[t]$ is contained in $\mathfrak{m}_{a,b}$ if and only if $a = \lambda$. So if λ is not a root of P, we get two prime ideals of R containing $(t - \lambda)$, while if λ is a root of P, we get only one. In either case, we can check that

$$(t - \lambda) = \mathfrak{m}_{\lambda,b}\mathfrak{m}_{\lambda,-b},$$

where b is a root of $b^2 = P(\lambda)$. Similarly to the case of integers, $(t-\lambda)$ is a **splitting prime** if λ is not a root of P, since it factors as a product of two distinct ideals; and $(t-\lambda)$ is a **ramified prime** if λ is a root of P, since it factors as a product of two copies of the same ideal. Note that there are no inert (or non-splitting) primes in this situation, since $\mathbb{C}$ is algebraically closed; if we instead worked with a field such as $\mathbb{F}_p$, there would be inert primes as well.

Just as in the case of quadratic number fields, R will not usually have unique factorization in *elements*. But there are some examples where there *is* unique factorization into elements; then all ideals are principal (which is not true in general), and factorization into ideals just becomes factorization into elements.

Example 18.2
When $P(t) = t$, we have $R = \mathbb{C}[z,t]/(z^2 - t) \cong \mathbb{C}[z]$. In this case, $(t-\lambda)$ factors as

$$(t-\lambda) = (z - \sqrt{\lambda})(z + \sqrt{\lambda}).$$

Here is a table of the counterparts between $\mathbb{Z}$ and $\mathbb{C}[t]$ (assume $d \not\equiv 1 \pmod 4$ for simplicity):

$\mathbb{Z}$	$\mathbb{Z}[\sqrt{d}]$	primes	ramified primes	splitting primes
$\mathbb{C}[t]$	$\mathbb{C}[t,z]/(z^2 - P(t))$	$t-\lambda$	$t-\lambda$ where $P(\lambda) = 0$	$t-\lambda$ where $P(\lambda) \neq 0$

Now we'll return to number fields. We previously defined the **ideal class group** $\mathrm{Cl}(F)$. We've already discussed that it's a group, but today we'll see that it's finite.

18.2 Application to Fermat's Last Theorem

Before we discuss the proof of finiteness, we'll circle back to an application to Fermat's Last Theorem. It suffices to consider the case where the exponent is prime. So suppose we want to solve $a^p + b^p = c^p$ over the integers, where p is an odd prime.

Then if ζ is a pth root of unity, in the field $F = \mathbb{Q}[\zeta]$ we can factor

$$(a+b)(a+\zeta b)\cdots(a+\zeta^{p-1}b) = c^p.$$

Assuming that no two factors have a common divisor, we want to conclude that each factor is a pth power — just as we would over the integers, since in that case every prime must have exponent divisible by p — meaning that $a + \zeta^n b = c_n^p \cdot u_n$ for a unit u_n.

We can do this if R is a PID, which is equivalent to $\mathrm{Cl}(F)$ being trivial — the unit in $\mathrm{Cl}(F)$ corresponds to principal ideals, so the group is trivial if and only if all ideals are principal. But this only happens for very few primes (which are all at most 19). However, it turns out that there's a more general condition we can use instead:

Definition 18.3
A prime p is **regular** if $p \nmid \#\,\mathrm{Cl}(F)$.

There are many regular primes: in fact, the only irregular primes $p \leq 100$ are 37, 59, and 67. So this covers many more cases than the requirement that $\mathbb{Z}[\zeta]$ be a PID. It's still unknown whether there's infinitely many regular primes, though.

Proposition 18.4
We can still conclude each factor is a pth power if p is regular.

Proof. First, using unique factorization for ideals, we see that the *ideal* $(a + \zeta^n b)$ is a pth power, meaning that $(a+\zeta^n b) = I^p$ for some ideal I.

Now it remains to check that I is principal, or equivalently, that its class $[I]$ is trivial. But we know I^p is principal, so $[I]^p = 1$. Now using regularity, no element of the class group has order p; so we must have $[I] = 1$ as well. □

There are still further steps — we then have to consider the case where the factors are *not* coprime as well. In the case where a, b, and c are not divisible by p, the factors are always coprime, and we can use this argument and get a contradiction with a bit more work. When one of them is divisible by p, a different argument is needed. But this is the key ingredient — the rest is clever but elementary.

18.3 Finiteness of the Class Group

Theorem 18.5
The class group $\mathrm{Cl}(F)$ is finite.

As usual, we'll see how to prove this for imaginary quadratic number fields. The key claim is the following:

Proposition 18.6
Every ideal class has a representative with bounded norm — more precisely, with norm at most

$$\mu = \begin{cases} \sqrt{|d|/3} & \text{if } d \equiv 1 \pmod 4 \\ 2\sqrt{|d|/3} & \text{if } d \equiv 2, 3 \pmod 4. \end{cases}$$

It's clear that this proposition implies the class group is finite, since there are finitely many ideals of any given norm. In fact, it also gives us an effective way to *compute* the class group — it implies that the class group is generated by the classes of ideals with norm $p \leq \mu$ for primes p. For each prime p, there are at most two ideals with norm p, which can be found by attempting to factor (p) as a product of ideals; then it's enough to find all these ideals and the relations between them.

To prove this proposition, we'll first prove the following lemma:

Lemma 18.7
An ideal I of norm n contains a nonzero element $\alpha \neq 0$ with $\alpha\overline{\alpha} \leq \mu n$.

The lemma itself can be proved by elementary geometry on lattices, so we'll first look at how it implies the proposition.

Proof of Proposition 18.6. Choose some nonzero ideal I of norm n in the given class, and find $\alpha \in \overline{I}$ with $|\alpha|^2 \leq \mu n$. We saw previously that inclusion of ideals implies divisibility, so then since $(\alpha) \subset \overline{I}$, we can write $(\alpha) = \overline{I}J$ for some ideal J.

But then $[J] = [I] = [\overline{I}]^{-1}$, because $J\overline{I} = (\alpha)$ and $I\overline{I} = (n)$ are both principal.

On the other hand, norm is multiplicative, so we have $\mathrm{N}(J)\,\mathrm{N}(\overline{I}) = \mathrm{N}(\alpha)$. So then

$$\mathrm{N}(J) = \frac{|\alpha|^2}{n} \leq \mu. \qquad \square$$

Proof of Lemma 18.7. For a lattice $L \subset \mathbb{R}^2$, define Δ_L as the area of its **fundamental parallelogram**, the parallelogram spanned by two vectors v and w for which $L = \{nv + mw \mid n, m \in \mathbb{Z}\}$.

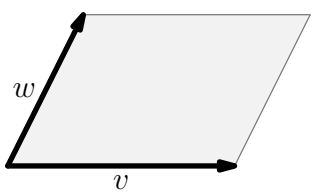

By elementary linear algebra, we can show that Δ_L doesn't depend on the choice of basis vectors. Furthermore, if $L' \subset L$, then one fundamental paralellogram of L can be obtained by taking the union of $[L : L']$ copies of the fundamental parallelogram of L', so $\Delta_{L'} = [L : L']\Delta_L$.

Now let $v \in L$ be a vector of minimal length.

Claim. *We have the bound* $|v|^2 \cdot \sqrt{3}/2 \leq \Delta_I$.

Proof. If v is the shortest vector, and w the shortest vector with $w \neq cv$, then v and w must span the lattice. Without loss of generality we can assume the angle α between v and w is at most $\pi/2$. Then we must have $\alpha \geq \pi/3$, or else $w - v$ has smaller length. Then

$$\Delta_I = |w| \cdot |v| \cdot \sin \alpha \geq \frac{\sqrt{3}}{2}|v|^2. \qquad \square$$

Now it remains to relate Δ_I to $\mathrm{N}(I)$: we claim that

$$\Delta_I = \begin{cases} \mathrm{N}(I)\sqrt{d} & \text{if } d \not\equiv 1 \pmod 4 \\ \mathrm{N}(I)\sqrt{d}/2 & \text{if } d \equiv 1 \pmod 4. \end{cases}$$

This is true when $I = (1)$ — for example, if $d \not\equiv 1 \pmod 4$, then the fundamental parallelogram is a rectangle.

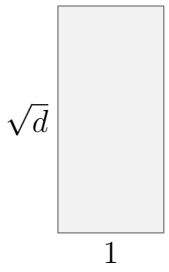

So now we can attempt to reduce the general case to the case $I = (1)$, using the following claim:

Claim. *We have* $\mathrm{N}(I) = [R : I]$.

Proof. It suffices to show that if P is a prime ideal and J any ideal, then $[J : PJ] = \mathrm{N}(P)$ — then we can factor I as a product of primes, and use this property repeatedly. This is obvious when P is principal, so now suppose $P\overline{P} = (q)$ for a prime q. Then we have

$$[J : PJ] \cdot [PJ : P\overline{P}J] = [J : qJ] = q^2,$$

since for any lattice L and integer n, we have $L/nL \cong (\mathbb{Z}/n)^2$. Then q^2 is the product of two integers, both not equal to 1; so each must be q. $\square$

Now we have $\Delta_I = [R : I]\Delta_R = \mathrm{N}(I)\Delta_R$. This gives the claimed expression for Δ_I, and therefore for $|v|$ using the previous claim. $\square$

This concludes the proof of the finiteness of the class group for imaginary quadratic number fields.

19 Modules over a Ring

The motivation for modules is that we are trying to tell a story where rings are the protagonist, and for a story to be interesting, the protagonist must act. When we find a way for a ring to act, we get the definition of a module.

Definition 19.1
Let R be a ring. A **module** M over R is an abelian group, together with an **action map** $R \times M \to M$ (written as $(r, m) \mapsto r(m)$ or rm), subject to the following axioms:

- $1_R(m) = m$ for all $m \operatorname{im} M$;
- $r_1(r_2(m)) = (r_1 r_2)(m)$ for all $r_1, r_2 \in R$ and $m \in M$;
- Distributivity in both variables: $(r_1 + r_2)m = r_1 m + r_2 m$ for all $r_1, r_2 \in R$ and $m \in M$, and $r(m_1 + m_2) = rm_1 + rm_2$ for all $r \in R$ and $m_1, m_2 \in M$.

The first two axioms are very similar to the definition of a group action on a set. So a ring to a module is like a group G to a G-set (a set with an action by G). It's not exactly the same because in a ring we have two operations instead of one, but they're a similar flavor.

19.1 Examples

Example 19.2
If $R = F$ is a field, then a module is the same as a vector space.

The axioms here are exactly the same. The textbook emphasizes this heavily, and this analogy can get you some mileage; but for general rings, things become more complicated.

Note 19.3
The definition also applies to a *noncommutative* ring R, in the same way — our definition does not reference commutativity. Then we have some familiar examples of modules over noncommutative rings: for example, given any field F we can take $R = \mathrm{Mat}_{n\times n}(F)$ and $M = F^n$, since matrices act on column vectors by multiplication. As another example, if $R = \mathbb{C}[G]$ is the group ring, then a R-module is the same as a complex representation of G.

For any ring R, there is a uniquely defined homomorphism $\mathbb{Z} \to R$, where $1 \mapsto 1_R$. On a similar note, every abelian group has a unique structure of a $\mathbb{Z}$-module: we know that 1 (in $\mathbb{Z}$) must map $m \mapsto m$, so then by distributivity, $n = 1 + 1 + \cdots + 1$ must map

$$v \mapsto \underbrace{v + v + \cdots + v}_{n}.$$

Similarly, $-n$ must map v to $-(v + v + \cdots + v)$. So a $\mathbb{Z}$-module is the same as an abelian group.

Example 19.4
What is a module over $R = \mathbb{C}[x]$?

Proof. First, a $\mathbb{C}[x]$-module is a $\mathbb{C}$-vector space V by looking at the action of constant polynomials (which are just scalars). But then we also need to see what x does. We know x must act by a linear map $A : V \to V$, where $xv = Av$. There are no constraints on this map, and this defines the action of every other polynomial: so a R-module is a vector space V, together with a linear map $A : V \to V$. Explicitly, the action of a general polynomial $P(x) = a_n x^n + \cdots + a_0$ is given by

$$Pv = a_0 v + a_1 Av + \cdots + a_n A^n v.$$

Note that the vector space may or may not be finite-dimensional; if it is, then we end up in a situation studied in linear algebra, where we have a vector space and a linear operator. □

Example 19.5
What is a module over $R = \mathbb{Z}/n\mathbb{Z}$?

Proof. The main point is that if R/I is a quotient of R, then every R/I-module is also a R-module, where we define $r(m)$ to be $\overline{r}(m)$ (here $\overline{r}$ denotes r mod I). Meanwhile, in order to go backwards, I must act by 0. So a R/I module is the same as a R-module where every element of I acts in a trivial way (meaning that $rv = 0$ for all $r \in I$ and $v \in M$).

So in this case, a $\mathbb{Z}/n\mathbb{Z}$-module is the same as an abelian group where the order of every element divides n — meaning $na = 0$ for all a in the group.

Then more concretely, for every m (where we use $\overline{m}$ to denote m mod n), we can write

$$\overline{m}v = \underbrace{v + v + \cdots + v}_{m}.$$

In order for this to be well-defined, the sum should not depend on the choice of representative for the residue; but this is guaranteed by the condition $na = 0$. (This is the same reasoning as in the first paragraph, for this specific example.) □

For any ring R, there is a simple example of a module:

Definition 19.6
The **free module** over R is $M = R$ itself, where the action is multiplication (meaning that $r(x) = rx$).

This is parallel to the observation that a group G acts on itself by left multiplication.

19.2 Submodules

Definition 19.7
Given a module M, a **submodule** $N \subset M$ is an abelian subgroup which is invariant under the R-action — meaning $rx \in N$ for all $x \in N$ and $r \in R$.

If $N \subset M$ is a submodule, we can define their **quotient** M/N, where we take the quotient in the sense of abelian groups. This quotient of abelian groups carries a module structure as well, given by the obvious rule $r\overline{m} = \overline{rm}$ (where $\overline{m}$ denotes m mod N). This is well-defined because N is a submodule — if $m_1 - m_2$ is in N, then $rm_1 - rm_2 = r(m_1 - r_2)$ is in N as well.

Then the homomorphism theorem and correspondence theorem work in the exact same way as in abelian groups. (For rings and ideals, we saw they work in a similar way; but here the parallel is closer.)

Example 19.8
What are the submodules of the free module $M = R$?

Proof. The answer is exactly the ideals of R — we're looking for abelian subgroups of R which are invariant under multiplication by all terms in R, and by definition these are ideals. □

We'll later see how to understand *any* module by looking at generators and relations — this turns out to be easier than the corresponding problem for a group. But first we'll look at another example of a module, which will be useful for developing that theory.

Definition 19.9
Given two modules M and N, their **direct sum** is

$$M \oplus N = \{(m, n) \mid m \in M, n \in N\}$$

with the action

$$r(m, n) = (rm, rn).$$

Note 19.10
The direct sum is the same as the product $M \times N$. This is true for any *finite* sum — we have

$$M_1 \oplus \cdots \oplus M_n = M_1 \times \cdots \times M_n.$$

But this isn't true for *infinite* sums and products.

Definition 19.11
The **free module** of **rank** n is

$$R^n = \underbrace{R \oplus R \oplus \cdots \oplus R}_{n}.$$

In the case where $R = F$ is a field, the free module of rank n is exactly F^n, the standard n-dimensional vector space.

19.3 Homomorphisms

In linear algebra, we work with matrices in order to understand linear maps. Matrices are also relevant here — the terms are different, but the concept is very similar.

Definition 19.12
A **homomorphism** from a module M to a module N is a homomorphism of abelian groups $\varphi : M \to N$, which is compatible with the R-action — meaning $\varphi(rm) = r\varphi(m)$ for all $r \in R$ and $m \in M$.

In vector spaces, this is the same as a linear map.

We'll use $\mathrm{Hom}_R(M, N)$ to denote the set of all homomorphisms $M \to N$. Note that homomorphisms can be added and rescaled, in the same way as linear maps: $(\varphi_1 + \varphi_2)(m) = \varphi_1(m) + \varphi_2(m)$, and $(r\varphi)(m) = r\varphi(m)$. So then $\mathrm{Hom}_R(M, N)$ is itself a R-module.

Understanding homomorphisms in general may be hard, but it's easy to understand homomorphisms from a free module. Given a homomorphism $\varphi \in \mathrm{Hom}_R(R, M)$ for any module M, we can let $m = \varphi(1_R)$. Then this determines the entire homomorphism — for any $r \in R$, we have

$$\varphi(r) = \varphi(r \cdot 1_R) = r \cdot \varphi(1_R) = rm.$$

So a homomorphism is determined by $m = \varphi(1_R)$, and there are no restrictions on m — this is why R is called a free module. This means $\mathrm{Hom}_R(R, M)$ is isomorphic to M: more explicitly, the bijection is given by mapping $\varphi \in \mathrm{Hom}_R(R, M)$ to $m_\varphi = \varphi(1)$, and $m \in M$ to the homomorphism $\varphi_m : r \mapsto rm$.

Similarly, $\mathrm{Hom}_R(R^n, M)$ is equally easy to understand. Now R^n is generated by the elements 1_i which have a 1 in their ith place, and 0's everywhere else (so $1_i = (0, \ldots, 0, 1, 0, \ldots, 0)$, where the 1 is in the ith place). So $\mathrm{Hom}_R(R^n, M)$ is isomorphic to M^n, where the bijection sends $\varphi \in \mathrm{Hom}_R(R^n, M)$ to the element $(\varphi(1_1), \varphi(1_2), \ldots, \varphi(1_n))$, and $(m_1, \ldots, m_n) \in M$ to the homomorphism $\varphi(x_1, \ldots, x_n) = \sum x_i m_i$.

In particular, we have $\mathrm{Hom}_R(R^n, R^m) = (R^m)^n = \mathrm{Mat}_{m \times n}(R)$ — we can write homomorphisms in the way we're used to in linear algebra, where $A \in \mathrm{Mat}_{m \times n}(R)$ sends $(x_1, \ldots, x_n)^t$ to $A(x_1, \ldots, x_n)^t$. So as long as we work with free modules and homomorphisms, to a large extent we can operate as if we're doing linear algebra. But in linear algebra, there's various characterizations of nondegenerate matrices that no longer hold here —

for instance, a linear operator that is injective (meaning it has zero kernel) is also surjective, but that's not true for general modules.

19.4 Generators and Relations

Definition 19.13
A collection of elements $m_1, \dots, m_n \in M$ forms a system of **generators** if every $x \in M$ can be expressed as $\sum r_i m_i$ for $r_i \in R$.

So in other words, $\varphi_{m_1,\dots,m_n} : R^n \to M$ is onto. If such a finite set exists, we say that M is *finitely generated*. Many modules we're interested in are in fact finitely generated.

If this map is also one-to-one, then it's an isomorphism, and M is free. But usually this won't happen, and we still want to describe M. To do this, we can look at $K = \ker(\varphi)$, which is a submodule in R^n. If K is itself finitely generated, then we can choose a set of generators for K, and get a somewhat explicit description of M — we can fix a system of k generators for K, and obtain another homomorphism $\psi : R^k \to K$. Since K sits inside R^n, we can think of ψ instead as a homomorphism $\psi : R^k \to R^n$, with image K. But such a homomorphism corresponds to a matrix $A \in \mathrm{Mat}_{k\times n}$, and we have $M = R^n/AR^n$.

Definition 19.14
If we can find a finite set of generators for M such that the kernel is also finitely generated, then M is called **finitely presented**.

We'll see that for many rings, the finitely presented modules are a very large class of modules — in fact, for many rings, any finitely generated module is finitely presented. We'll also see how to make this description of a module very explicit when the ring is a Euclidean domain. In particular, taking the Euclidean domain to be $\mathbb{Z}$ will give us a classification of all finitely generated abelian groups, and taking the Euclidean domain to be $F[x]$ will give us Jordan normal form!

20 Modules and Presentation Matrices

20.1 Review — Definition of Modules

Last class, we defined modules over commutative rings — we've also seen a few examples over noncommutative rings, but from now we'll stick to commutative ones.

Definition 20.1
A **module** M over a ring R is an abelian group together with an action of R — each $r \in R$ acts by a map $r : m \mapsto r(m)$ (we often denote $r(m)$ by rm), satisfying certain axioms.

In some sense, modules can be similar to vector spaces:

Example 20.2
The free module of rank n is $M = R^n$, consisting of n-tuples of elements in R (where addition and the R-action are performed componentwise).

But there are many other examples of modules over most rings, and we'll now discuss how to describe more general modules.

20.2 Generators and Relations

If M is a module, the elements $a_1, \dots, a_n \in M$ form a **system of generators** if every $x \in M$ has the form

$$x = \sum r_i a_i$$

for some $r_i \in R$.

Choosing any n elements $a_1, \dots, a_n \in M$ provides a homomorphism of modules $\varphi : R^n \to M$ (where we send $(r_1, \dots, r_n) \mapsto r_1 a_1 + \cdots + r_n a_n$). Then $a_1, \dots, a_n$ are generators if and only if φ is onto.

In the analogy with vector spaces, a system of generators is a set of vectors that *span* the vector space. But in vector spaces, if we have a set of vectors which span the space, we can always find a subset which forms a basis — we can drop some of the a_i to make the presentation $x = \sum r_i a_i$ be *unique*. This is not true in general.

Stating that the presentation $x = \sum r_i a_i$ is unique is equivalent to stating that $\ker \varphi = 0$ (given two presentations, we could subtract them to get an element of the kernel). Then $\varphi : R^n \to M$ is actually an isomorphism, which means M is free. But this is often *not* the case — there are usually many R-modules which are not free.

Student Question. *Does the system of generators have to be finite?*

Answer. *Not necessarily; as a silly example, we could take* all *elements of M as our system of generators. In later classes, we'll discuss ways to sometimes show that we* can *always find a finite system of generators; but for today, we'll just* assume *that we can.*

Definition 20.3
If a finite set of generators exists, then M is called **finitely generated.**

20.3 Presentation Matrices

Most of the time, we won't be lucky enough that $\ker \varphi = 0$. But in general, by the homomorphism theorem, we can write

$$M \cong R^n / \ker \varphi.$$

Now $\ker \varphi$ is *also* a module, so we can describe it further.

Definition 20.4
If $\ker \varphi$ is also finitely generated (for some surjective homomorphism $\varphi : R^n \to M$), then we say M is **finitely presented**.

We'll see later that for a large class of rings, *every* module which is finitely generated is also finitely presented; but for now, we'll assume that M is finitely presented as well.

Then we can choose a set of generators $b_1, \ldots, b_m$ for $\ker\varphi$. Each b_i can be thought of as a column vector (since it's an element of R^n). So we can write down the $n \times m$ marix

$$B = \begin{bmatrix} | & | & & | \\ b_1 & b_2 & \cdots & b_m \\ | & | & & | \end{bmatrix},$$

called the **presentation matrix**. Now since we've fixed generators for $\ker\varphi$, we have an onto map $\psi : R^m \to \ker\varphi$. Equivalently, we can think of ψ as a map $\psi : R^m \to R^n$ whose image is exactly $\ker\varphi$. As in the case of vector spaces in linear algebra, this map can explicitly described as $\psi : x \mapsto Bx$ (where $x \in R^m$), and $\operatorname{im}\psi = BR^m$. This means we can write

$$M \cong R^n/BR^m$$

(where BR^m is the span of the column vectors of the presentation matrix B).

20.4 Classifying Modules

The rest of the lecture focuses on classifying finitely generated modules over a Euclidean domain. Since abelian groups are $\mathbb{Z}$-modules, this will also give us a classification of finitely generated groups.

For a given module M, the presentation is not at all unique.

Example 20.5
Let $R = \mathbb{Z}$ and $M = \mathbb{Z}/5\mathbb{Z}$. An obvious presentation of M is to choose one generator 1, and the generator 5 for the kernel. In this case, $B = \begin{bmatrix}5\end{bmatrix}$. But there are other presentations as well. For example, we can choose two generators for M, and the presentation matrix

$$B = \begin{bmatrix} 2 & 1 \\ 1 & 3 \end{bmatrix}.$$

To see that $\mathbb{Z}^2/B\mathbb{Z}^2$ is again $\mathbb{Z}/5\mathbb{Z}$, note that the sublattice $L \subset \mathbb{Z}^2$ spanned by $(2,1)^t$ and $(1,3)^t$ has index $|\det(B)| = 5$, which means the quotient $\mathbb{Z}^2/L$ has five elements. But the only group of five elements is $\mathbb{Z}/5\mathbb{Z}$, so this construction indeed gives another presentation of $\mathbb{Z}/5\mathbb{Z}$.

Our goal is to more systematically understand how to understand the module given a presentation. Here, the analogy between vector spaces and modules will be useful — similarly to in linear algebra, we'll start with a matrix and try to perform elementary operations on the matrix which don't change the module. We'll then use these operations to write the matrix in a simpler form, from where it's easy to understand the module.

20.4.1 Elementary Row and Column Operations

We'll use the notation M_B to refer to the module M produced by a presentation matrix B.

In linear algebra, we saw the elementary column operations for matrices over a field. We can define elementary column operations on matrices over a *ring* in a similar way:

1. Multiplying a column by a unit (an invertible element) — note that in the case of a field, we could multiply by *any* nonzero element (because any nonzero element has an inverse), but it was important that we were able to invert the multiplication; so here we restrict the definition to units.
2. Adding an arbitrary multiple of one column to another column.
3. For convenience, we can also include swapping two columns; but in fact, this can be obtained as a combination of the first two operations.

These operations *do not* change the span of the columns. (This can be verified the same way as for matrices over a field.) So if B' is obtained from B by elementary column operations, then we have $BR^m = B'R^m$.

Another useful way to think of this is in terms of matrix multiplication. In other words, if B' is obtained from B by elementary column operations, then we have $B' = B \cdot C$, where C is an $m \times m$ *invertible* matrix (it's easy to see that all the elementary column operations are invertible). This means $C \in \mathrm{GL}_m(R)$ (the group of

invertible matrices wth entries in R). Note that over a ring, for a matrix C to be invertible, $\det(C)$ must be a *unit* — it's not enough to require the determinant to be nonzero. (In fact, the converse is also true, but requires more work.)

Then we have $B'R^m = BCR^m$. But since C is invertible, the map defined by C is an isomorphism, so $CR^m = R^m$. So this is an alternate way of making it clear that $B'R^m = BR^m$.

The elementary *row* operations are defined analogously.

When we apply elementary row operations to a matrix B, we produce a matrix B' which is a presentation matrix for an *isomorphic* module — in order to explain why, we'll again think in terms of matrix multiplication. We can write $B' = CB$, where $C \in \mathrm{GL}_n(R)$. We then want to produce an isomorphism between R^n/BR^m and R^n/CBR^m. But that isomorphism is just given by multiplication by C. More explicitly, we can draw a commutative diagram:

$$\begin{array}{ccc} R^n/BR^m & \longrightarrow & R^n/CB \cdot R^m \\ \uparrow & & \uparrow \\ R^n & \xrightarrow[C]{} & R^n \end{array}$$

The map $x \mapsto Cx$ is an isomorphism $R^n \to R^n$. But by definition, $x \in BR^n$ if and only if $Cx \in CBR^m$. So if we restrict our isomorphism to BR^m, then it restricts to an isomorphism between BR^m and CBR^m; and therefore to an isomorphism between the quotients R^n/BR^m and $R^n/CB \cdot R^m$.

So the row and column operations don't change our module (up to isomorphism).

20.4.2 Smith Normal Form

Using the row and column operations, it's possible to write any presentation matrix in a much simpler form. (We'll primarily focus on the case $R = \mathbb{Z}$, but this holds for any Euclidean domain.)

Theorem 20.6
Every $n \times m$ matrix over a Euclidean domain R can be reduced by elementary row and column operations to a matrix in *Smith normal form* — if we let $B = (b_{ij})$, then we have $b_{ij} = 0$ for all $i \neq j$, and $b_{11} \mid b_{22} \mid b_{33} \mid \cdots$.

So a matrix in Smith normal form looks like

$$\begin{bmatrix} d_1 & & & & & \\ & d_2 & & & & \\ & & d_3 & & & \\ & & & \ddots & & \\ & & & & d_k & \end{bmatrix}$$

(where all entries not shown are 0's).

We'll discuss the proof next class. It combines two ideas — the method of using Gaussian elimination to solve systems of equations over a *field* (which involves reducing matrices to a simpler form using row and column operations as well), and the Euclidean algorithm. (We've previously discussed Euclidean domains in the context of them being PIDs, but it turns out that here, having an effective way of computing the gcd using division with remainder will be very useful. The theorem is also true for PIDs, with a bit of modification (and a different proof); but all the examples we'll be interested in are Euclidean domains.

Corollary 20.7
Every finitely presented module over a Euclidean domain is isomorphic to a direct sum of *cyclic* modules (modules which are generated by one element) — we can write

$$M \cong R^a \oplus R/(d_1) \oplus R/(d_2) \oplus \cdots \oplus R/(d_k),$$

where we additionally have $d_1 \mid d_2 \mid \cdots \mid d_k$.

Later we'll see that any finitely generated module over a Euclidean domain is also finitely presented; this will mean that our statement actually holds for all finitely *generated* modules.

This corollary is clear from the theorem:

Proof. Using Theorem 20.6, we can rewrite the presentation matrix B to be diagonal, with diagonal $d_1 \mid d_2 \mid \cdots \mid d_k$, where $k = \min(m, n)$. In this case, when we take a column vector in R^m and multiply by B, we simply scale each coordinate by the corresponding d_i — so when we quotient out by BR^m, this coordinate becomes $R/(d_i)$, and we get

$$M_B = \bigoplus R/(d_i) \oplus R^a.$$

(The extra free factor comes from rows with no entries — either because $d_i = 0$ or because $n > m$ — since such rows correspond to coordinates in R^n where we're not quotienting out by anything.) □

In fact, this classification can be used to understand the subgroups of lattices as well (which came up when we studied factorization in quadratic number fields). This theorem implies that to describe a subgroup, we can always choose a basis for the lattice, and simply scale these basis vectors by numbers to get a basis for the subgroup.

Student Question. *What does it mean for a module to be cyclic — is the free module cyclic?*

Answer. *A module is cyclic if it's generated by one element. The free module is cyclic — it's generated by one element with* no *relations. Meanwhile, R/d is generated by one element x, with the relation that $dx = 0$.*

21 Smith Normal Form

21.1 Review

Last time, we looked at presentation matrices for a module. We saw that if we perform elementary row and column operations on the presentation matrix, then this leads to an isomorphic module. We then stated the theorem that we can use such operations to reduce any matrix to Smith normal form. We will prove this today; but first, let's look at a few examples.

21.2 Some Examples in $\mathbb{Z}$

We'll consider the case of 2×2 matrices over $\mathbb{Z}$ — consider the presentation matrix

$$B = \begin{bmatrix} a & c \\ b & d \end{bmatrix},$$

whose corresponding module is

$$M_B = \mathbb{Z}^2 / \operatorname{Span}((a,b)^t, (c,d)^t).$$

The simplest case is when B is diagonal:

Example 21.1
Consider the presentation matrix

$$B = \begin{bmatrix} 2 & 0 \\ 0 & 3 \end{bmatrix}.$$

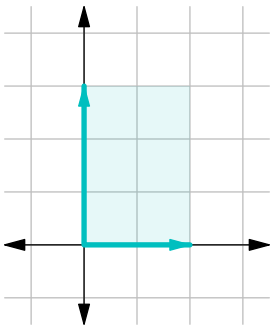

In this case, we have

$$M_B \cong \mathbb{Z}/2\mathbb{Z} \oplus \mathbb{Z}/3\mathbb{Z},$$

since given any vector (m, n), its coset mod $\operatorname{Span}((2,0)^t, (0,3)^t)$ is just given by taking the first component mod 2 and the second mod 3. (To be pedantic, the isomorphism is given by $(m, n) \mapsto (m \mod 2, n \mod 3)$ — given any vector, we can subtract multiples of our two vectors to bring it into the rectangle.)

Example 21.2
Consider the presentation matrix

$$B = \begin{bmatrix} 5 & 0 \\ 0 & 1 \end{bmatrix}.$$

Now the corresponding module is $M_B = \mathbb{Z}/5\mathbb{Z}$ — the second coordinate is "useless" since we're allowed to subtract multiples of $(0,1)^t$, so we can always eliminate it. So to keep track of the coset of a vector, we only need to keep track of the first component mod 5.

Now we'll look at a more complicated example, where the original matrix is *not* diagonal.

Example 21.3
Consider the presentation matrix

$$B = \begin{bmatrix} 2 & 1 \\ 1 & 3 \end{bmatrix}.$$

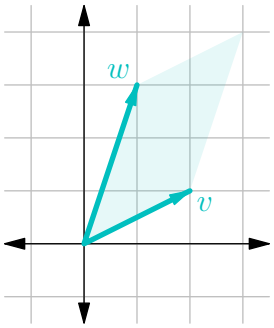

We can see $\det(B) = 5$, so we can use elementary column operations to make one column a multiple of 5 — if we add twice the first column to the second, then we get

$$B' = B\begin{bmatrix}1 & 2\\ 0 & 1\end{bmatrix} = \begin{bmatrix}2 & 5\\ 1 & 5\end{bmatrix}.$$

So we've replaced w with $w' = 5(1,1)^t$, without changing the lattice spanned by our vectors:

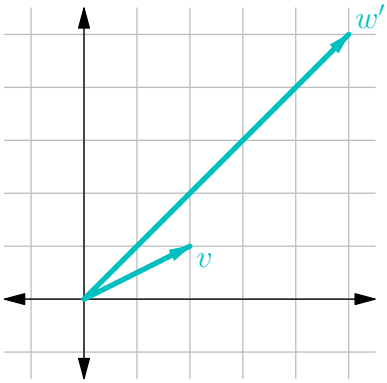

But $(2,1)^t$ and $(1,1)^t$ form a basis for $\mathbb{Z}^2$! So this means to get our lattice, we started with a basis for $\mathbb{Z}^2$, then fixed one of the basis vectors and scaled the other by 5. So this is isomorphic to the previous example — by changing the basis we use for $\mathbb{Z}^2$, we can rewrite v and w' as $(1,0)^t$ and $(0,5)^t$. So we have $M_B \cong \mathbb{Z}/5\mathbb{Z}$.

21.3 Smith Normal Form

Now we'll return to the general case, and prove the theorem.

Theorem 21.4
For a Euclidean domain R, any $n \times m$ matrix B can be reduced using elementary row and column operations to a matrix D, where $d_{ij} = 0$ for all $i \neq j$, and $d_{11} \mid d_{22} \mid \cdots$.

One example of a matrix written in this form (which is called Smith normal form) is

$$D = \begin{bmatrix}2 & 0 & 0 & 0\\ 0 & 6 & 0 & 0\\ 0 & 0 & 12 & 0\end{bmatrix}.$$

Note that if D can be obtained from B by elementary row and column operations, then we can write $D = ABC$, where A is an invertible $n \times n$ matrix and C is an invertible $m \times m$ matrix. For any such D, we then have $M_D \cong M_B$.

Note 21.5
In a PID, it's still possible to obtain D in Smith normal form with $D = ABC$ (for A and C invertible). However, it may not be possible to obtain D by using the elementary operations.

To motivate the proof, notice that the greatest common divisor of all the matrix entries does *not* change under elementary operations — it's clear that scaling by a unit doesn't change the gcd; meanwhile if we perform the operation $a_{ij} \longmapsto a_{ij} + ca_{kj}$, we have

$$\gcd(a_{ij}, a_{kj}) = \gcd(a_{ij} + ca_{kj}, a_{kj}).$$

(This is the same idea as in the Euclidean algorithm.) So if we are able to obtain a matrix D in Smith normal form, then we *must* have

$$d_{11} = \gcd(b_{ij})$$

(since d_{11} divides all other entries of D). So this suggests the main idea of the proof — we want to run some sort of Euclidean algorithm in order to isolate the gcd of all matrix entries in the top-left corner.

Proof of Theorem 21.4. Recall that in a Euclidean domain, we have a size function $\sigma : R \setminus \{0\} \to \mathbb{Z}_{\geq 0}$ such that we can divide with remainder — given any a and b, with $b \neq 0$, we can write $a = bq + r$ where $r = 0$ or $\sigma(r) < \sigma(b)$. For convenience, we write $|a|$ instead of $\sigma(a)$.

If $B = 0$, we are done, so assume $B \neq 0$. Then we'll use induction on the size of the matrix; the key step is the following.

Lemma 21.6
By row and column operations, we can arrive at a matrix B' such that $b'_{11} = \gcd(b_{ij}) = \gcd(b'_{ij})$.

Proof. By permuting the rows and columns, we can ensure that $b_{11} \neq 0$, and b_{11} is the nonzero element of minimal size. Now if $b_{11} \mid b_{ij}$ for all i and j, then we're done, so assume not.

Now the main idea is to modify the matrix to make a *smaller* element appear (which we can again move to the top-left corner by rearranging rows and columns). First, if there is an entry with $b_{11} \nmid b_{ij}$ in the first row or column, then we can perform a row or column operation to reduce it — if $b_{ij} = qb_{11} + r$, then we can subtract q times the first row or column from the row or column of b_{ij}, which replaces b_{ij} with r.

If not, then we can use b_{11} to eliminate all other entries in the first row and column (by subtracting multiples of the first column and row), to get a matrix of the form

$$\begin{bmatrix} b_{11} & 0 & \cdots & 0 \\ 0 & * & \cdots & * \\ \vdots & \vdots & \ddots & \vdots \\ 0 & * & \cdots & * \end{bmatrix}.$$

Now there must be some b_{ij} with $b_{11} \nmid b_{ij}$. We can add its row to the first row; this adds 0 to b_{11}, so we now have a matrix

$$\begin{bmatrix} b_{11} & * & b_{ij} & \cdots & * \\ 0 & * & * & \cdots & * \\ 0 & * & * & \cdots & * \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & * & * & \cdots & * \end{bmatrix}.$$

Then we can again add a multiple of the first column to the jth column in order to produce an entry with smaller size.

So either way, we've now produced a matrix with an element smaller in size than b_{11}. Now permute the rows and columns to move this element to the position of b_{11}. Then we repeat the process. At every step, we replace b_{11} with a nonzero entry of smaller size. Since the size is always a nonnegative integer, at some point this process must terminate; this means that b_{11} now divides all entries. □

To complete the proof of Theorem 21.4, we induct on the size of B. By Lemma 21.6, we can replace B with a matrix B' where b'_{11} divides b'_{ij} for all i and j.

Now using row and column operations, we can eliminate the first row and column (meaning that we make b'_{i1} and b'_{1j} all zero). So we get

$$\begin{bmatrix} b'_{11} & * & \cdots & * \\ * & * & \cdots & * \\ \vdots & \vdots & \ddots & \vdots \\ * & * & \cdots & * \end{bmatrix} \rightsquigarrow \begin{bmatrix} b'_{11} & 0 & \cdots & 0 \\ 0 & * & \cdots & * \\ \vdots & \vdots & \ddots & \vdots \\ 0 & * & \cdots & * \end{bmatrix}.$$

Now we ignore the first row and column, and use elementary operations on rows and columns $2, \ldots, n$. (These do not affect the first row or column.) By the induction assumption, we can reduce the submatrix of $*$s to Smith normal form (since b'_{11} already divides all other entries, this will remain true when we perform the operations). □

The theorem has more theoretical value than computational value, but we will compute an example nonetheless.

Example 21.7
We have

$$\begin{bmatrix} 2 & 1 \\ 1 & 3 \end{bmatrix} \rightsquigarrow \begin{bmatrix} 1 & 3 \\ 2 & 1 \end{bmatrix} \rightsquigarrow \begin{bmatrix} 1 & 3 \\ 0 & -5 \end{bmatrix} \rightsquigarrow \begin{bmatrix} 1 & 0 \\ 0 & -5 \end{bmatrix}.$$

Example 21.8
We have

$$\begin{bmatrix} 4 & 2 & 6 \\ 1 & 2 & 3 \end{bmatrix} \rightsquigarrow \begin{bmatrix} 1 & 2 & 3 \\ 4 & 2 & 6 \end{bmatrix} \rightsquigarrow \begin{bmatrix} 1 & 0 & 0 \\ 4 & -6 & -6 \end{bmatrix} \rightsquigarrow \begin{bmatrix} 1 & 0 & 0 \\ 0 & -6 & -6 \end{bmatrix} \rightsquigarrow \begin{bmatrix} 1 & 0 & 0 \\ 0 & -6 & 0 \end{bmatrix}.$$

21.4 Applications

As a corollary, by taking R to be $\mathbb{Z}$, we can classify all finitely presented abelian groups.

Corollary 21.9
Every finitely presented abelian group is isomorphic to

$$\mathbb{Z}/d_1\mathbb{Z} \times \mathbb{Z}/d_2\mathbb{Z} \times \cdots \times \mathbb{Z}/d_n\mathbb{Z} \times \mathbb{Z}^a,$$

for some positive integers d_i with $d_1 \mid d_2 \mid \cdots \mid d_n$.

Sometimes, it's more useful to write this classification in a different form. Recall that the Chinese Remainder Theorem states that if $n = ab$ with $\gcd(a, b) = 1$, then

$$\mathbb{Z}/n\mathbb{Z} \cong \mathbb{Z}/a\mathbb{Z} \times \mathbb{Z}/b\mathbb{Z}.$$

So then if we factor $d_i = p_1^{a_{i1}} \cdots p_n^{a_{in}}$, we can decompose $\mathbb{Z}/d_i\mathbb{Z}$ as a product of groups of the form $\mathbb{Z}/p^m\mathbb{Z}$ (cyclic groups of prime power order).

Another application of Theorem 21.4 is in the case $R = F[x]$, where F is a field. A finitely generated module over R must then be of the form

$$R^a \oplus R/(P_1) \oplus \cdots \oplus R/(P_n),$$

where $P_1 \mid \cdots \mid P_n$. Alternatively, again using the Chinese Remainder Theorem, we can instead assume that each P_i is a power of an irreducible polynomial.

In particular, consider $F = \mathbb{C}$; then the only irreducible polynomials are linear, so we must have $P_i = (x - \lambda_i)^{n_i}$. If we only consider finite-dimensional modules, then as we said earlier, a module over $F[x]$ is equivalent to a (finite-dimensional) vector space along with one linear operator (corresponding to the action of x). So it turns out that this classification of modules actually implies the Jordan decomposition theorem — we will discuss this in more detail next lecture.

22 Decomposition of Modules

22.1 Classification of Abelian Groups

Last class, we proved the classification of finitely presented abelian groups (and more generally, modules over a Euclidean domain):

Theorem 22.1
Any finitely presented abelian group A is isomorphic to

$$\mathbb{Z}/d_1\mathbb{Z} \times \mathbb{Z}/d_2\mathbb{Z} \times \cdots \times \mathbb{Z}/d_n\mathbb{Z} \times \mathbb{Z}^a,$$

where $d_1 \mid d_2 \mid \cdots \mid d_n$.

The idea of the proof was to start with a presentation matrix B, and reduce it to Smith normal form (a diagonal matrix with the additional divisibility condition) by using elementary operations.

Note 22.2
The textbook describes this as diagonalization. But note that this is a *different* kind of diagonalization than the one used in Jordan normal form — in Jordan normal form we reduced the matrix to a simpler form by using conjugation, while here we're using elementary operations.

Today, we'll discuss various features of this classification.

22.1.1 Uniqueness of Subgroups

There are multiple questions we can ask about uniqueness. One is whether the numbers d_i are uniquely defined given A, and we'll see later that the answer is yes.

Meanwhile, when we write A as a product of factors, each factor is itself a *subgroup* of A — if we have the product $G \times H = \{(g, h) \mid g \in G, h \in H\}$, then G is isomorphic its subgroup consisting of the set $\{(g, 1)\}$. So this gives another question we can ask about uniqueness:

Guiding Question
Which subgroups corresponding to the factors in the decomposition of an abelian group are *uniquely determined* from A?

First, the product of all the finite factors $\mathbb{Z}/d_1\mathbb{Z} \times \cdots \times \mathbb{Z}/d_n\mathbb{Z}$ is actually a canonically defined subgroup. It's exactly the set of all elements with finite order — if an element only has nonzero components in these factors, then it clearly has finite order; while if it has a nontrivial component in the free factor, then it can't have finite order.

Definition 22.3
The set of elements with finite order is called the **torsion subgroup** A_f, so we have

$$A_f = \mathbb{Z}/d_1\mathbb{Z} \times \cdots \times \mathbb{Z}/d_n\mathbb{Z}.$$

On the other hand, the free factor $\mathbb{Z}^n$ is uniquely defined up to an isomorphism, but it doesn't necessarily correspond to a uniquely defined subgroup. For example, take $A = \mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}$, so $A_f = \mathbb{Z}/2\mathbb{Z}$. Then A contains two free subgroups. It contains the subgroup

$$A_{\mathbb{Z}} = \langle (0, 1) \rangle,$$

which consists of elements $(0, n)$ for integers n — this is the obvious subgroup we'd think of as isomorphic to the free factor $\mathbb{Z}^n$ in the decomposition. But A also contains another subgroup

$$A_{\mathbb{Z}'} = \langle (\overline{1}, 1) \rangle,$$

which consists of elements $(\overline{n}, n)$ for integers n (where $\overline{n}$ represents n mod 2), and is also isomorphic to $\mathbb{Z}$. Both subgroups are complements to $\mathbb{Z}/2\mathbb{Z}$ in A, but they are not the same.

Student Question. *Why is the torsion subgroup called A_f, if it is* not *the free factor?*

Answer. *The f stands for finite, not free. It's an unfortunate coincidence that "finite" and "free" begin with the same letter.*

But it's easy to see that the rank a of the free factor *is* well-defined — we have $\mathbb{Z}^a = A/A_f$, but $\mathbb{Z}^a \not\cong \mathbb{Z}^b$ if $a \neq b$. To prove this explicitly, otherwise we would have an $a \times b$ matrix B and $b \times a$ matrix C (with integer coefficients) such that $BC = 1_a$ and $CB = 1_b$, representing the two directions of the isomorphism. This is impossible even dropping the requirement that they have integer coefficients — if $a < b$ then $\text{rank}(CB) \leq a < b = \text{rank}(1_b)$, contradiction.

22.1.2 The Torsion Subgroup

We can write another expression for the torsion subgroup A_f. By using the Chinese Remainder Theorem (which states that $\mathbb{Z}/mn\mathbb{Z} \cong \mathbb{Z}/m\mathbb{Z} \times \mathbb{Z}/n\mathbb{Z}$ if $\gcd(m,n) = 1$), we can split every d_i into prime powers, and write $\mathbb{Z}/d_i\mathbb{Z} \cong \prod \mathbb{Z}/p_j^{s_j}\mathbb{Z}$. So then we can write A_f as a product of cyclic groups with prime power order. We can then collect factors corresponding to the same prime, giving the decomposition

$$A_f = A_{p_1} \times A_{p_2} \times \cdots \times A_{p_m},$$

where each factor is of the form $A_p = \prod \mathbb{Z}/p^{e_i}\mathbb{Z}$.

Example 22.4
Write $\mathbb{Z}/36\mathbb{Z} \times \mathbb{Z}/6\mathbb{Z}$ in this form.

Proof. We can split $36 = 4 \cdot 9$ and $6 = 2 \cdot 3$, to get $(\mathbb{Z}/4\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}) \times (\mathbb{Z}/9\mathbb{Z} \times \mathbb{Z}/3\mathbb{Z})$. □

It's easier to prove uniqueness when we write the decomposition in this form, so we'll now work with this way of writing A_f. (It's possible to use the uniqueness of this decomposition to prove uniqueness of the d_i in the original form, where $d_1 \mid \cdots \mid d_n$, as well; this will be left as an exercise.)

Note that A_p is a p-Sylow subgroup of A_f. Since the group is abelian, the Sylow subgroup is unique (in the case of a general group, all Sylow subgroups are conjugate). In fact A_p is exactly the set of elements whose order is a power of p — this is called the **p-torsion subgroup**.

Within each A_p, the set-theoretic decomposition into subgroups may not be unique, but we can show the following lemma:

Lemma 22.5
The multiplicities of the powers of p in the decomposition of A_p as a product of cyclic groups are uniquely determined by A.

For example, this means $\mathbb{Z}/4\mathbb{Z} \times \mathbb{Z}/4\mathbb{Z}$ is not isomorphic to $\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/8\mathbb{Z}$. Note that we chose the numbers here so that it's not completely obvious — clearly if we have two isomorphic decompositions then their sizes must match. Here we do have $4 \cdot 4 = 2 \cdot 8$, but they're still not isomorphic, for more subtle reasons.

Proof. Let $A = \mathbb{Z}/p^{a_1}\mathbb{Z} \times \cdots \times \mathbb{Z}/p^{a_n}\mathbb{Z}$ (where $a_i \geq 1$). There are two main observations here.

First consider A/pA. Each factor is replaced with $\mathbb{Z}/p\mathbb{Z}$ (for example, by the homomorphism theorem), so $A/pA = (\mathbb{Z}/p\mathbb{Z})^n$. This means $|A/pA| = p^n$, where n is the number of factors — so any two decompositions must have the same number of factors.

Meanwhile, we can also look at pA, which is a subgroup of A. We have $p\mathbb{Z}/p^a\mathbb{Z} \cong \mathbb{Z}/p^{a-1}\mathbb{Z}$. So replacing A with pA reduces each of the exponents by 1, and

$$pA = \prod \mathbb{Z}/p^{a_i - 1}\mathbb{Z}.$$

(It's possible that some of these factors are trivial, since $a_i - 1$ may be 0; but we can use the first observation to deal with this.)

Now use induction on $|A|$. If we can write A in two ways as

$$A \cong \mathbb{Z}/p^{a_1}\mathbb{Z} \times \cdots \times \mathbb{Z}/p^{a_n}\mathbb{Z} \cong \mathbb{Z}/p^{a_1'}\mathbb{Z} \times \cdots \times \mathbb{Z}/p^{a_m}\mathbb{Z},$$

then we must have $n = m$ by the first observation, and

$$pA \cong \prod \mathbb{Z}/p^{a_i-1}\mathbb{Z} \cong \prod \mathbb{Z}/p^{a_i'-1}\mathbb{Z}$$

by the second. By the induction hypothesis, we can match all the nonzero $a_i - 1$ and $a_i' - 1$; so we can match all $a_i > 1$ with $a_i' > 1$. Only looking at these, we lose the information about the a_i which equal 1; but using the fact that $m = n$, we can also match the $a_i = 1$ with $a_i' = 1$. So the two decompositions must be the same. □

22.2 Polynomial Rings

This classification works for modules over any Euclidean domain. Now consider the case of $R = F[t]$, where F is a field. The theorem says that a finitely presented module over R is of the form

$$M \cong R/(P_1) \times \cdots \times R/(P_n) \times R^a,$$

where $P_1 \mid P_2 \mid \cdots \mid P_n$. Similarly to before, we can rewrite the decomposition as

$$M \cong \prod R/Q_i^{a_i} \times R^a,$$

where the Q_i are irreducible.

In particular, if the module M is finite-dimensional as a vector space over F, then there is no free factor R^a.

When we started discussing modules, we saw that a $F[t]$ module is the same as a F-vector space V, together with a linear operator $V \to V$ (the action of t). So understanding isomorphism classes of modules is the same as understanding this situation, which was studied in linear algebra with the Jordan normal form theorem.

We'd like to explicitly figure out what $R/(P)$ looks like. We can assume P is monic without loss of generality (since F is a field), so we can write $P(t) = t^n + a_{n-1}t^{n-1} + \cdots + a_0$. Then $R/(P)$ has a basis consisting of $\overline{1}$, $\overline{t}$, $\ldots$, $\overline{t^{n-1}}$. Let $e_i = \overline{t^{i-1}}$. Then to describe the action of t, we have $te_i = e^{i+1}$ for $1 \le i \le n-1$, while

$$te_n = -a_0e_1 - a_1e^2 - \cdots - a_{n-1}e_n.$$

So then the matrix corresponding to t is

$$\begin{bmatrix} 0 & 0 & 0 & \cdots & 0 & -a_0 \\ 1 & 0 & 0 & \cdots & 0 & -a_1 \\ 0 & 1 & 0 & \cdots & 0 & -a_2 \\ 0 & 0 & 1 & \cdots & 0 & -a_3 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 1 & -a_n \end{bmatrix}.$$

But we can sometimes say even more. Consider $F = \mathbb{C}$, and use the second form of the decomposition, where we quotient by powers of irreducible polynomials. The only irreducible polynomials are linear, so we have $Q_i(t) = t - \lambda_i$ for some λ_i, and in each factor we're quotienting out by some power of such a linear polynomial.

If $\lambda_i = 0$, then using the above basis, the entire right column is 0. So we get a matrix with 0's on the diagonal, 1's directly below the diagonal, and 0's everywhere else — this is the form of one Jordan block for a nilpotent matrix.

Meanwhile, in general, we can use the basis consisting of $\overline{(t-\lambda_i)^j}$ instead of $\overline{t^j}$. Then the situation is the same, except that we add a scalar to the matrix — the action of $t - \lambda_i$ corresponds to this exact matrix, and we add the scalar matrix λ_i to get the action of t. So we get the same matrix with λ_i on the diagonal instead of 0, which is a general Jordan block.

So this gives a proof of the Jordan normal form theorem, and shows that in fact, both Jordan normal form and the classification of finite abelian groups follow from the same more general theorem.

22.3 Noetherian Rings

We'll now move to the last topic about modules, Noetherian rings (named after Emmy Noether). Today we'll just see an overview of the statements, and we'll prove them next class.

Definition 22.6
A ring R is **Noetherian** if every ideal in R is finitely generated.

Example 22.7
Any field is Noetherian (since the only two ideals are the ones generated by 0 and 1), and any PID is Noetherian (since an ideal is generated by one element).

The reason the concept is useful is the following proposition:

Proposition 22.8
A ring R is Noetherian if and only if every submodule in a finitely generated R-module is itself finitely generated.

In particular, we get the following corollary:

Corollary 22.9
If R is Noetherian, every finitely generated module is finitely presented.

This explains why the notion is useful, but then we're left with the question of how to produce examples. There are some easy reductions (for example, the quotient of a Noetherian ring is Noetherian as well). But the key result is the Hilbert Basis Theorem:

Theorem 22.10 (Hilbert Basis Theorem)
If R is Noetherian, then $R[x]$ is Noetherian.

This gives a powerful tool for proving that many rings are Noetherian, and therefore that many modules are finitely generated.

Note 22.11
There is a legend about this theorem: Hilbert published this theorem in 1890. According to the legend, a famous mathematician Paul Gordan (referred to as the king of invariant theory, a branch of algebra studying such questions) ostensibly said that this is not mathematics, it's theology. At the time, people were trying to work with this case by case, to explicitly produce a finite set of generators. In contrast, this theorem has a short and very abstract proof that doesn't give much information about how to write down the actual generators.

23 Noetherian Rings

Definition 23.1
A ring R is **Noetherian** if every ideal in R is finitely generated.

In other words, every quotient can be obtained by imposing a finite number of relations.

23.1 Submodules over Noetherian Rings

The first important property of Noetherian rings is the following:

Proposition 23.2
A ring R is Noetherian if and only if every submodule in a finitely generated R-module is itself finitely generated.

Corollary 23.3
If R is Noetherian, then every finitely generated module is finitely presented.

Before we prove this, we'll look at a few observations.

Lemma 23.4
If we have a surjective homomorphism $\varphi : M \to N$ of R-modules, then:

1. If M is finitely generated, then N is also finitely generated.
2. If N is finitely generated, and $K = \ker(\varphi)$ is also finitely generated, then M is also finitely generated.

Note 23.5
This is one place where the intuition from linear algebra is useful: it's helpful to think about the case where R is a field, so being finitely generated is equivalent to being finite-dimensional. In this case, we know more precisely that

$$\dim(M) = \dim(N) + \dim(K).$$

We won't get information this precise in the general case, but the proof of finiteness is similar.

Proof of Lemma 23.4. The first part is obvious — take any set of generators $m_1, \dots, m_n$ for M. Then their images $\varphi(m_1), \dots, \varphi(m_n)$ must generate N.

For the second part, let $k_1, \dots, k_a$ be a set of generators for K, and $n_1, \dots, n_b$ a set of generators for N. Now pick $\tilde{n}_1, \dots, \tilde{n}_b$ such that $\varphi(\tilde{n}_i) = n_i$ (which we can do by surjectivity).

Now we claim that the k_i and $\tilde{n}_i$ generate M: given any $x \in M$, we can find $r_1, \dots, r_b$ in R such that

$$\varphi(x) = r_1 n_1 + \cdots + r_b n_b$$

(since the n_i generate N). So then we have

$$\varphi(x - r_1\tilde{n}_1 - \cdots - r_b\tilde{n}_b) = \varphi(x) - r_1 n_1 - \cdots - r_b n_b = 0,$$

which means $x - \sum r_i \tilde{n}_i$ is in K. So we can express $x - \sum r_i \tilde{n}_i = \sum s_j k_j$, which gives an expression for x as a linear combination of the $\tilde{n}_i$ and k_j. □

Proof of Proposition 23.2. One direction is clear: an ideal of R, by definition, is the same as a submodule of the free module (which is generated by one element). So if every submodule of a finitely generated module is finitely generated, then R must be Noetherian.

For the other direction, we want to show that if every ideal is finitely generated, then so is every submodule of a finitely generated module.

The strategy is to reduce to the case of a free module R^n, and use induction on n — we know this is true for $n = 1$, so we want to reduce to this case.

Let M be a finitely generated module. Then by picking a set of generators, we can find a surjective homomorphism $\varphi : R^n \to M$ (fixing such a homomorphism is equivalent to fixing a set of generators).

Then by the correspondence theorem and the above lemma, it's enough to check that every submodule of R^n is finitely generated (since the submodules of M are exactly the images of the submodules of R^n containing $\ker \varphi$).

Now we can argue by induction on n. First, the base case $n = 1$ follows directly from the definition of a Noetherian ring, since submodules of R are exactly ideals.

For the inductive step, consider a submodule $N \subset R^n$, with $n > 1$. Now split $R^n = R \times R^{n-1}$, and take the projection homomorphism $\pi : R^n \to R^{n-1}$, which sends $(r_1, \dots, r_n) \mapsto (r_2, \dots, r_n)$.

Then $\pi(N)$ is a submodule of R^{n-1}, so by the induction assumption, it's finitely generated. Meanwhile, the kernel K of π is the set of elements of the form $(r, 0, 0, \dots)$ which are in N. But this is a submodule of the free rank-1 module R, so K is also finitely generated.

So by the lemma, since K and $\pi(N)$ are both finitely generated, N is finitely generated as well. □

Example 23.6
When considering submodules $N \subset \mathbb{Z}^2$, we'd take the points on the x-axis as K, and the projections onto the y-axis as $\pi(N)$. Note that N is not necessarily $K \times \pi(N)$.

Student Question. *Did we use the fact that finitely generated modules are finitely presented in this proof, when we obtained the homomorphism $R^n \to M$?*

Answer. *No — the surjective homomorphism $\varphi : R^n \to M$ comes just from the definition of being finitely generated. We take some generators $m_1, \dots, m_n$, and then map $(r_1, \dots, r_n)$ to the linear combination $r_1 m_1 + \cdots + r_n m_n$. Being finitely presented would only matter if we were looking at the* kernel *of this homomorphism (which we didn't need to do here).*

So now we have that if R is Noetherian, any finitely generated module is also finitely presented:

Proof of Corollary 23.3. If the module is finitely generated, then there is a surjective map $\varphi : R^n \to M$. Then $\ker \varphi$ is a submodule of R^n, so it must be finitely generated as well. □

This means the classification of finitely presented abelian groups that we saw earlier is actually a classification of all finitely *generated* abelian groups.

Note 23.7
It's also possible to define Noetherian rings by Proposition 23.2 (as rings with the property that every submodule of a finitely generated module is finitely generated). But we gave the definition using ideals so that the property that's easier to check is in the definition, and the one that's more useful is in the proposition.

23.2 Constructing Noetherian Rings

So we've seen why the notion of Noetherian rings is useful. But in order to use it, we want to see how to produce more examples of Noetherian rings, beyond just fields and PIDs.

First, there is a simple observation we can make:

Lemma 23.8
A quotient of a Noetherian ring is again Noetherian — if R is a Noetherian ring and I an ideal of R, then the ring $S = R/I$ is also Noetherian.

Proof. This is immediate from the correspondence theorem: an ideal in R/I is of the form J/I, where $J \subset R$ is an ideal containing I. Then we can just take the images of the generators — if $J = (x_1, \dots, x_n)$, then $J/I = (\overline{x}_1, \dots, \overline{x}_n)$. So all ideals of R/I are finitely generated. □

Note 23.9
A subring in a Noetherian ring is not necessarily Noetherian — so this concept is more subtle than the dimension of a vector space.

For example, $\mathbb{C}[x, y]$ is Noetherian. But the subring $\mathbb{C} + x\mathbb{C}[x, y]$ (consisting of polynomials which are constant mod x) is *not* Noetherian — the ideal $x\mathbb{C}[x, y]$ is not finitely generated.

23.2.1 Hilbert Basis Theorem

There's actually a powerful tool that shows many rings are Noetherian:

Theorem 23.10 (Hilbert Basis Theorem)
If R is Noetherian, then $R[x]$ is also Noetherian.

This theorem has useful implications:

Corollary 23.11
If R is Noetherian, then $R[x_1, \dots, x_n]/I$ is also Noetherian, for any ideal I.

So if we start with a field, this shows us how to produce many examples of Noetherian rings.

Corollary 23.12
Any algebraic subset in $\mathbb{C}^n$ — a subset given by a collection of polynomial equations — is always given by a *finite* set of polynomial equations.

Proof of Theorem 23.10. Let $I \subset R[x]$ be an ideal, so we want to check that I is finitely generated. It's enough to find a finite collection of polynomials $P_1, \dots, P_n$ in I and a bound d, such that every element in I can be reduced to a polynomial of degree d using the P_i — meaning that

$$I \subset (P_1, \dots, P_n) + R[x]_{\leq d}.$$

In other words, once the polynomials and d are fixed, then for every $P \in I$, we need to be able to find $Q_1, \dots, Q_n$ in $R[x]$ such that

$$\deg(P - \sum Q_i P_i) \leq d.$$

If we can find such P_i and d, then

$$I \subset (P_1, \dots, P_n) + (I \cap R[x]_{\leq d}).$$

But the second term is finitely generated over R (it's not a submodule of $R[x]$, but it *is* a submodule of R), since $R[x]_{\leq d}$ is a free module of rank $d + 1$ over R, and R is Noetherian. So if it's generated by $S_1, \dots, S_m$, then I is generated by the P_i and S_i.

So now we want to figure out how to do this — in some sense, the idea is generalized division with remainder.

Consider the ideal $\overline{I}$ in R consisting of the *leading coefficients* of polynomials in I (along with 0) — we saw in the homework that this is an ideal. Then $\overline{I}$ is finitely generated. Let $P_1, \dots, P_n$ be polynomials whose leading coefficients generate $\overline{I}$, and let $d = \max(\deg P_i)$.

Now if we have a polynomial P of degree greater than d, we can cancel its leading coefficient — we can find Q_i such that $\sum Q_i P_i$ has the same degree and same leading coefficient as P, and then subtract them to decrease the degree of P by at least 1. We can then repeat this until P has degree at most d. □

Student Question. *Why is $I \cap R[x]_{\leq d}$ finitely generated?*

Answer. *The point is that $R[x]_{\leq d}$ is a free module of rank $d+1$ over R — we forget that we can multiply the polynomials and just add them and scale by elements of R, so the coordinates correspond to the coefficients of $1, x, \ldots, x^d$. Then $I \cap R[x]_{\leq d}$ is a submodule of $R[x]_{\leq d}$, which is a finitely generated module over R; so since R is Noetherian, $I \cap R[x]_{\leq d}$ is also finitely generated.*

23.3 Chain Conditions

Proposition 23.13
A ring is Noetherian if and only if every increasing chain of ideals stabilizes. In other words, if there is a chain of ideals $I_1 \subseteq I_2 \subseteq \cdots$, then from some point on, $I_n = I_{n+1} = \cdots$.

In other words, R is *not* Noetherian if and only if there exists an infinite chain $I_1 \subsetneq I_2 \subsetneq \cdots$ of ideals.

Proof Outline. If R is Noetherian and we have a chain $I_1 \subseteq I_2 \subseteq \cdots$, then their union $I = I_1 \cup I_2 \cup \cdots$ is an ideal. Since R is Noetherian, then I is finitely generated. So each of the generators must have some ideal I_k that it first shows up in, and since there's finitely many, there's some ideal I_k containing all the generators of I (which must equal I).

For the other direction, we can essentially take an ideal I that isn't finitely generated, and starting with the empty ideal, we keep adding an element of I which isn't in any of the previous ideals. □

24 Fields

24.1 Review — Noetherian Rings

Recall the definition of Noetherian rings:

Definition 24.1
A ring is Noetherian if every ideal is finitely generated.

This definition has a useful reformulation, in terms of chains.

Proposition 24.2
A ring is Noetherian if and only if every increasing chain of ideals stabilizes — in other words, given any chain of ideals $I_1 \subseteq I_2 \subseteq \cdots$, from some point on we must have $I_n = I_{n+1} = \cdots$.

Example 24.3
In the case $R = \mathbb{Z}$, a chain of ideals amounts to a list of integers d_1, d_2, ... where $d_i \mid d_{i-1}$ for all i. This chain clearly has to stabilize, since it's a nonincreasing sequence of positive integers.

Proof of Proposition 24.2. First suppose R is Noetherian, and we have a chain $I_1 \subseteq I_2 \subseteq \cdots$. Then their union $I = I_1 \cup I_2 \cup \cdots$ is an ideal. Since R is Noetherian, then I is finitely generated, so we can write $I = (a_1, \ldots, a_n)$. Then for each i, a_i must be contained in some ideal (since it's in their union). But there's finitely many a_i, so some I_m must contain all of them (since $I_k \subseteq I_{k+1}$, once a_i appears in one ideal, it appears in all following ones). So then $I_m = I$, and the sequence must stabilize at I_m — we must have $I_m = I_{m+1} = \cdots = I$.

For the other direction, suppose R is *not* Noetherian, and let $I \subset R$ be an ideal that's not finitely generated. Pick $a_1 \in I$, and define a_n inductively such that $a_n \in I$ but $a_n \notin (a_1, \ldots, a_{n-1})$ — this is possible because otherwise I would be generated by a_1, ..., a_{n-1}. Now we can take $I_n = (a_1, \ldots, a_n)$, which gives an infinite non-stabilizing chain of ideals. □

Note 24.4
Sometimes the chain condition is given as the *definition* of Noetherian rings.

This has an application to unique factorization in PIDs. Previously, we proved uniqueness but not existence — in our specific examples we had a concept of a norm that we could use to prove that the factorization process always terminates, but it's nontrivial to prove existence of a factorization in an abstract PID. But now we do have the tools to prove that the factorization process terminates in general.

Corollary 24.5
In a PID, every element can be factored as a product of irreducibles.

Proof. Start with the element, and keep on factoring terms until stuck. If we get stuck, then all factors must be irreducible; so assume that the factorization process is infinite instead. Then we can organize the factorization process into an infinite chain of ideals — attempting to factor an element will give a chain d_1, d_2, d_3, ... where $d_{i+1} \mid d_i$ for all i, and then $(d_1) \subset (d_2) \subset \cdots$. This would contradict the chain condition, so factorization must terminate. □

Similarly, we can also revisit a statement we made earlier about maximal ideals:

Proposition 24.6
In a Noetherian ring, every (non-unit) ideal is contained in a maximal ideal.

This is actually true in *any* ring, but the proof requires set theory. But we're mostly interested in several concrete rings, which are all Noetherian; so this proposition covers all such cases.

Proof. Let $I \subset R$ be an ideal, and assume I is not contained in a maximal ideal. Set $I_1 = I$. Now find $I_2 \supset I$ (not equal to I or R), which is possible since I is not maximal. But I_2 is not maximal either, so we can similarly construct $I_3 \supset I_2$, and inductively build a chain $I_1 \subsetneq I_2 \subsetneq I_3 \subsetneq \cdots \subsetneq R$ — this is possible at every step because if we couldn't choose I_{n+1}, then I_n would be a maximal ideal containing I. This contradicts chain termination. □

Note 24.7
As a final remark on modules: one important tool in studying modules is an upgrade on the presentation with generators and relations. The idea is that once we have generators and relations, we can also look at the relations between relations, and so on. First construct a surjective map $R^n \to M$ by fixing generators. This map has a kernel K_1, and we can construct a surjective map $R^m \to K_1$ by fixing its generators. This map has a kernel K_2, and we can construct a surjective map $R^\ell \to K_2$, and so on. This is called the **sygyzy** or **free resolution**, and it tells you a lot about the structure of the module.

24.2 Introduction to Fields

Definition 24.8
A **field** is a ring where all nonzero elements are units.

Definition 24.9
A **field extension** is a pair of fields $L \supset K$. The extension is written as L/K.

The theory of field extensions developed from trying to understand how to systematically solve polynomial equations — an important example of a field extension is $L = K(\alpha)$ where α is a root of an irreducible polynomial. We know a formula for solving quadratics; people were interested in understanding more general formulas. One result we'll get to in a few weeks is that if P is a generic polynomial in $\mathbb{Q}[x]$ of degree at least 5, and $K = \mathbb{Q}(\alpha)$ for a root α of P, then K is not contained in any field $\mathbb{Q}(\beta_1, \ldots, \beta_n)$ such that $\beta_i^{d_i} \in \mathbb{Q}(\beta_1, \ldots, \beta_{i-1})$ for each i (for some positive integers d_i) — you can't build the field by repeatedly taking roots. It then follows that there's no expression for the roots of P in terms of arithmetic operations and radicals.

Another application is to compass and straightedge constructions — let $\zeta_n \in \mathbb{C}$ be a primitive nth root of unity.

Guiding Question
When is $\mathbb{Q}(\zeta_n)$ contained in a field of the form $\mathbb{Q}(\alpha_1, \ldots, \alpha_n)$ where $\alpha_i^2 \in \mathbb{Q}(\alpha_1, \ldots, \alpha_{i-1})$ for all i?

In other words, we want to find out whether we can obtain ζ_n by the regular operations along with extracting square roots. This is interesting because it happens if and only if a regular n-gon can be constructed by a compass and straightedge. We'll see how to get a complete answer (as Gauss did) — for example, the answer is yes for $n = 17$ and no for $n = 19$. (In fact, a 17-gon appeared on Gauss's tombstone — this was an achievement he was proud of.)

24.3 Field Extensions

Definition 24.10
An extension L/K is **finite** if L is finite-dimensional as a vector space over K.

Essentially, what this means is that if we forget that we can multiply by elements of L, but remember that we can add them and multiply by elements of K, then that turns L into a vector space over K; the extension is finite when this vector space has finite dimension.

We've already seen how to construct examples of finite extensions: if $P \in K[x]$ is an irreducible polynomial, then we saw that $K[x]/(P)$ is a field. (This is because $K[x]$ is a PID, so (P) is a maximal ideal.) In this case, the dimension of the vector space is $d = \deg P$, since we saw the monomials $\overline{1}$, $\overline{x}$, …, $\overline{x^{d-1}}$ form a basis.

Definition 24.11
The **degree** of the extension L/K, denoted $[L : K]$, is the dimension of L as a vector space over K.

On the other hand, suppose we start with an extension L/K, and pick an element $\alpha \in L$. We say α is **algebraic** over K if it satisfies a polynomial equation — meaning that $P(\alpha) = 0$ for some nonzero $P \in K[x]$.

If α is algebraic, then we can take its minimal polynomial P (the monic polynomial of smallest degree with $P(\alpha) = 0$ — this is unique because all polynomials with $P(\alpha) = 0$ form an ideal, and since $K[x]$ is a PID, all ideals are principal and generated by their minimal-degree element). The minimal polynomial has to be irreducible — if $P = P_1P_2$, then $P(\alpha) = P_1(\alpha)P_2(\alpha) = 0$, but since we're in a field, this would imply $P_1(\alpha)$ or $P_2(\alpha)$ is 0, contradicting minimality.

Then we have a homomorphism $K[x]/(P) \to L$ sending $x \mapsto \alpha$. But any homomorphism of fields is injective (alternatively, the kernel of the map $f : K[x] \to L$ sending $x \mapsto \alpha$ is exactly the set of polynomials P with $P(\alpha) = 0$, which is generated by the minimal polynomial of α; so this homomorphism is injective). So we have an isomorphism $K[x]/(P) \cong K(\alpha)$, where $K(\alpha)$ is the subfield of L generated by α.

Note 24.12
This isomorphism is important: later we'll look at automorphisms of fields, and this is the trick that will let us build such maps. For example, start with $K = \mathbb{Q}$ and $L = \mathbb{C}$, and take $\alpha = \sqrt[3]{2}$ and $\beta = \sqrt[3]{2}\omega$ for a primitive third root of unity ω. Then α and β are both roots of the irreducible polynomial $x^3 - 2$. So $\mathbb{Q}(\alpha) \cong \mathbb{Q}(\beta)$, since both extensions are isomorphic to the abstract construction $\mathbb{Q}[x]/(x^3 - 2)$.

Lemma 24.13
If we have a field extension L/K, then $\alpha \in L$ is algebraic if and only if $K(\alpha)$ is finite-dimensional over K.

Proof. We've already seen that if α is algebraic, then $K(\alpha)$ is finite-dimensional (since it's isomorphic to $K[x]/(P)$ for some irreducible P, which is finite). Meanwhile, a polynomial relation can be thought of as a linear relation between the powers of α. If $m > \dim_K K(\alpha)$, then $1, \alpha, \ldots, \alpha^m$ must be linearly dependent; that linear dependence corresponds to a polynomial over K, giving a polynomial P such that $P(\alpha) = 0$. □

Corollary 24.14
If L/K is finite, then every $\alpha \in L$ is algebraic over K.

24.4 Towers of Extensions

Proposition 24.15
Suppose that we have a tower of field extensions $K \supset E \supset F$, where K/E and E/F are finite. Then K/F is finite, and

$$[K : F] = [K : E] \cdot [E : F].$$

Proof. Let $\alpha_1, \ldots, \alpha_n$ be a basis for E as a vector space over F, and $\beta_1, \ldots, \beta_m$ a basis for K as a vector space over E. Then we'll show that the terms $\alpha_i\beta_j$ form a basis for K/F.

This becomes clear just by substituting notation. First, in order to see that this is a generating set, every $x \in K$ can be written as $x = \sum \lambda_i\beta_i$, while each $\lambda_i \in E$ can be written as $\lambda_i = \sum a_{ij}\alpha_j$, which gives

$$x = \sum a_{ij}\alpha_j\beta_i.$$

Similarly, to prove independence, if we have $\sum a_{ij}\alpha_j\beta_i = 0$, we can collect terms into

$$\sum \left(\sum a_{ij}\alpha_j \right) \beta_i = 0.$$

If the a_{ij} are not all 0, then one of the terms $\sum a_{ij}\alpha_j$ must be nonzero (since the α_j are linearly independent), and therefore the entire sum must be nonzero (since the β_i are linearly independent), contradiction. □

Example 24.16
Find $[\mathbb{Q}(\alpha, \beta) : \mathbb{Q}]$, where $\alpha = \sqrt[3]{2}$ and $\beta = \sqrt[3]{2}\omega$.

Proof. We saw α and β both have minimal polynomial $x^3 - 2$. Then $[\mathbb{Q}(\alpha) : \mathbb{Q}] = 3$. Meanwhile, $x^3 - 2$ factors in $\mathbb{Q}(\alpha)$ as $(x - \alpha)(x^2 + \alpha x + \alpha^2)$. The second factor is irreducible (as otherwise, β would lie in $\mathbb{Q}(\alpha)$), so $[\mathbb{Q}(\alpha, \beta) : \mathbb{Q}(\alpha)] = 2$, which means $[\mathbb{Q}(\alpha, \beta) : \mathbb{Q}] = 6$. □

25 Field Extensions

25.1 Primary Fields

We have the following useful fact about fields:

Fact 25.1
Every field is a (possibly infinite) extension of either $\mathbb{Q}$, or $\mathbb{F}_p$ for a prime p. These are called the **primary fields**.

Proof. In general, for any ring R, there is a unique ring homomorphism $\mathbb{Z} \to R$ — we must have $1 \mapsto 1_R$, so then $n \mapsto \underbrace{1_R + \cdots + 1_R}_{n} = n_R$ for positive integers n, and $-n \mapsto -n_R$.

The image of the homomorphism is a quotient of $\mathbb{Z}$ — it's either $\mathbb{Z}$ or $\mathbb{Z}/n\mathbb{Z}$. Now consider the kernel of this homomorphism. If R is an integral domain (note that all fields are domains), then either the homomorphism is one-to-one, or its kernel is (p) for a prime p — otherwise, the image would be $\mathbb{Z}/n\mathbb{Z}$ for composite n, which is not a domain (as it has zero divisors).

Now taking $R = F$ to be a field, if the kernel is zero, then $\mathbb{Z}$ is a subring of F. But then $\mathbb{Q} = \text{Frac}(\mathbb{Z})$ must be inside F as well (since we can invert elements in a field) — in our original notation, the copy of $\mathbb{Q}$ in F is the fractions of the form n_R/m_R.

On the other hand, if the kernel is (p), then we have a copy of $\mathbb{Z}/p\mathbb{Z}$ in F, and we're done. □

Definition 25.2
The generator of the kernel (as in the above proof) is called the **characteristic** of the field.

So fields of characteristic 0 contain $\mathbb{Q}$, and fields of characteristic p contain $\mathbb{Z}/p\mathbb{Z}$ (and these are the only possible characteristics).

25.2 Algebraic Elements

Last time, we defined algebraic elements in a field extension L/K:

Definition 25.3
An element $\alpha \in L$ is **algebraic** over K if $P(\alpha) = 0$ for some nonzero $P \in K[x]$.

As stated last class, α is algebraic if and only if $K(\alpha)/K$ is finite (since a polynomial in α is the same as a linear relation between powers of α).

We also looked at towers of extensions $E/F/K$ — here E/K is called the **composite** extension, while E/F and F/K are called **intermediate** extensions. In particular, we saw the following theorem:

Theorem 25.4
We have

$$[E : K] = [E : F] \cdot [F : K].$$

In particular, E/K is finite if and only if both E/F and F/K are finite.

This has some useful corollaries regarding algebraic elements.

Corollary 25.5
If $\alpha, \beta \in L$ are algebraic over K, then $\alpha + \beta$, $\alpha\beta$, and $\frac{\alpha}{\beta}$ are also algebraic.

Proof. If α and β are algebraic, then $K(\alpha)/K$ and $K(\alpha,\beta)/K(\alpha)$ are both finite — since β satisfies a polynomial relation with coefficients in K, it satisfies the same polynomial relation with coefficients in $K(\alpha)$. So we can conclude that $K(\alpha,\beta)/K$ is finite, and therefore any element in it is algebraic. □

Corollary 25.6
Given an arbitrary extension, the set of elements in L which are algebraic over K form a subfield of L, called the **algebraic closure** of K in L.

For example, the algebraic closure of $\mathbb{Q}$ in $\mathbb{C}$ is called the **algebraic numbers**.

This is an abstract argument that doesn't exactly tell us how to construct the polynomial; but it's possible to come up with a procedure to write down an equation as well.

Example 25.7
Let $\alpha = \sqrt{2}$ and $\beta = \sqrt{3}$, and $\gamma = \alpha + \beta$. How can we write down a polynomial equation for γ?

One possible method is that by Corollary 25.5, we know that $1, \gamma, \gamma^2, \dots$ must be linearly dependent. In this case, they are all linear combinations of 1, $\sqrt{2}$, $\sqrt{3}$, and $\sqrt{6}$ with coefficients in $\mathbb{Q}$ — so they lie in a vector space of dimension at most 4. Then $1, \gamma, \dots, \gamma^4$ are five elements in a four-dimensional vector space, so they must be linearly dependent; and using linear algebra, it's possible to explicitly calculate this linear relation.

There is another way to find the polynomial equation — right now we'll present it as a guess, but later we'll see that it's part of a more general story.

We'd like to find a polynomial P with γ as a root, so we can try to think about what the other roots of P should be. Suppose P factors as $(x-\gamma_1)(x-\gamma_2)\cdots$, for $\gamma_i \in \mathbb{C}$ — it suffices to choose the γ_i such that P has rational coefficients, and $\gamma_1 = \sqrt{2}+\sqrt{3}$.

We can guess that all of $\pm\sqrt{2} \pm \sqrt{3}$ should be roots — from an algebraic perspective, if $\sqrt{2}+\sqrt{3}$ shows up, we "should" be able to switch the sign of the square root (since there isn't a difference between the two signs). So then we can take

$$\begin{aligned}\gamma_1 &= \sqrt{2}+\sqrt{3}\\ \gamma_2 &= \sqrt{2}-\sqrt{3}\\ \gamma_3 &= -\sqrt{2}+\sqrt{3}\\ \gamma_4 &= -\sqrt{2}-\sqrt{3}.\end{aligned}$$

We can expand out the polynomial to see that it does indeed have rational coefficients (essentially, this involves using the equation $a^2-b^2=(a-b)(a+b)$ twice).

The main idea we used here is to exploit the symmetry between the roots (there is a group of symmetries acting on the roots, by replacing one of the square roots with its negative); we'll later discuss ways to find these symmetries, using Galois theory.

25.3 Compass and Straightedge Construction

Proposition 25.4 also relates to compass and straightedge constructions. It has the following corollary:

Corollary 25.8
If $E/F/K$ is a tower of finite extensions, then $[F:K] \mid [E:K]$.

The problem of which regular n-gons can be constructed using a compass and straightedge can be rephrased algebraically in the following way (we won't discuss the details here).

Fact 25.9
A regular n-gon is constructible with compass and straightedge if and only if $\zeta_n = e^{2\pi i/n}$ lies in an extension $\mathbb{Q}(\alpha_1,\alpha_n)$ such that $\alpha_i^2 \in \mathbb{Q}(\alpha_1,\dots,\alpha_{n-1})$ for all i.

This means we have a tower of quadratic extensions, where every step in this tower has degree 2 — more explicitly, we can define $F_i = \mathbb{Q}(\alpha_1, \dots, \alpha_i)$, with $F_0 = \mathbb{Q}$. Without loss of generality we can assume $\alpha_i \notin F_{i-1}$ (or else adding it to the set of generators would be useless). Then we have the tower of extensions $F_n/F_{n-1}/\cdots/F_1/F_0$ where $[F_i : F_{i-1}] = 2$ for all i.

For convenience, we'll assume n is prime. (The general case involves a few more details, but works very similarly.)

Theorem 25.10
Let $n = p$ be prime. Then a regular p-gon can be constructed if and only if $p = 2^k + 1$.

Primes $p = 2^k + 1$ are called **Fermat primes**. There's only 5 known Fermat primes (3, 17, 257, and 65537); it's conjectured that there are no others, but we don't even know whether there's finitely or infinitely many. (Note that if $2^k + 1$ is prime, then k must be a power of 2 — otherwise, $2^k + 1$ can be factored.)

We'll only show one direction: that if ζ_p is constructible, then p is a Fermat prime. To prove this, the following proposition will be useful:

Proposition 25.11
If p is prime, we have $\deg(\zeta_n) = p - 1$, or equivalently $[\mathbb{Q}(\zeta_p) : \mathbb{Q}] = p - 1$.

The extension $\mathbb{Q}(\zeta_p)$ is called a **cyclotomic extension**.

Proof. We know ζ_p is a root of $x^p - 1$. We can easily factor

$$x^p - 1 = (x - 1)(x^{p-1} + x^{p-2} + \cdots + 1),$$

so it suffices to show that the second factor, which we call $P(x)$, is irreducible. Since the polynomial is primitive, it's enough to show that it's irreducible over $\mathbb{Z}$.

Now we can perform a trick — substitute $t = x - 1$. Then if we write $P(x) = Q(t)$, we have

$$tQ(t) = (t + 1)^p - 1.$$

But by expanding and using the Binomial Theorem, we then have

$$Q(t) = \sum_{i=0}^{p-1} \binom{p}{i+1} t^i.$$

(For example, when $p = 3$, we have $Q(t) = t^2 + 3t + 3$.)

But the leading term is 1, and all other terms are divisible by p; and the free term is not divisible by p^2 (in fact, *none* of the terms are divisible by p^2, but we only need to use the free term here).

Now assume for contradiction that Q is reducible, so $Q = Q_1 Q_2$ for polynomials Q_1 and Q_2 of degree at least 1. Now consider the reduction mod p, where

$$\overline{Q} = \overline{Q_1 Q_2}.$$

But $\overline{Q}$ is now t^{p-1}, and the only way to factor t^{p-1} in $\mathbb{F}_p[x]$ is as $t^i t^{p-1-i}$. But we have $\deg(\overline{Q_1}) = \deg(Q_1) > 1$ (and the same is true for Q_2), since the leading coefficients of Q_1 and Q_2 cannot be divisible by p (their product is the leading coefficient of Q, which is 1). So then we must have $i \neq 0, p - 1$.

But then since Q_1 and Q_2 are t^i and t^{p-1-i} for $0 < i < p - 1$, their free terms must both be divisible by p. So the product of their free terms is divisible by p^2; but this product is the free term of Q, which is *not* divisible by p^2. So this is a contradiction, and Q is irreducible. □

Proof of Necessity in Theorem 25.10. We've seen that $\deg(\zeta_p) = p - 1$. So we have $\deg(\zeta_p) = p - 1$. On the other hand, if $\zeta_p \in F_n$ for a field extension of the form described, then $\deg(\zeta_p)$ must divide $[F_n : \mathbb{Q}]$, which is a power of 2. So $p - 1$ must be a power of 2 as well. □

With our current tools, we can only show one direction — to show the other direction, we need a better extension of which fields can be obtained as the top floor of a tower of quadratic extensions. It's necessary that the degree is a power of 2, but this may not be sufficient. In the case of $\mathbb{Q}(\zeta_p)$, the condition turns out to be sufficient as well (as we'll see later).

25.4 Splitting Fields

We've seen the construction where we start with an irreducible polynomial $P \in F[x]$, and construct the field extension $E = F[x]/(P)$. This is an extension of F of degree $n = \deg(P)$, and we can think of it as adjoining a root of the polynomial.

But there's another construction which also produces a finite extension from a polynomial, which is in some sense harder to control. Here, we do not require the polynomial to be irreducible.

Definition 25.12
For a polynomial $P \in F[x]$, a **splitting field** of P is an extension E/F such that:

1. P splits as a product of linear factors in $E[x]$;
2. $E = F(\alpha_1, \dots, \alpha_n)$, where the α_i are the roots of P.

The first condition guarantees that P splits completely (so we can find all its roots) in E; the second prevents E from being too large (it only contains the elements which are necessary for P to split).

Proposition 25.13
Given any polynomial P, its splitting field exists, and any two splitting fields of P are isomorphic.

We'll discuss the proof in more detail next time — the main idea is to add one root of P so that it splits partially, then add another root of any remaining irreducible factor, and so on.

Example 25.14
The splitting field of $P(x) = x^3 - 2$ over $F = \mathbb{Q}$ is $E = \mathbb{Q}(\sqrt[3]{2}, \omega\sqrt[3]{2})$, where ω is a primitive 3rd root of unity. We have $[E : \mathbb{Q}] = 6$.

On the other hand, we could start by adjoining ω:

Example 25.15
The splitting field of $P(x) = x^3 - 2$ over $F = \mathbb{Q}(\omega)$ is $E = F(\sqrt[3]{2})$ — the polynomial $x^3 - 2$ remains irreducible, but after adjoining one root, we already have all the others. Here $[E : F] = 3$.

Note that E is the same in both examples (even though F is not).

Example 25.16
The splitting field of $P(x) = x^{p-1} + \cdots + 1$ over $F = \mathbb{Q}$ is $E = \mathbb{Q}(\zeta_p)$ (since all roots are powers of ζ_p), where $[E : F] = p - 1$.

26 Finite Fields

26.1 Splitting Fields

Last time, we stated the uniqueness of the splitting field of a polynomial.

Proposition 26.1
If F is a field, and P a (not necessarily irreducible) polynomial in F, then there exists a *unique* extension E/F up to isomorphism, such that P splits as a product of linear factors in $E[x]$ as $P(x) = \prod(x - \alpha_i)$, and $E = F(\alpha_1, \dots, \alpha_n)$.

Proof. The idea of the proof is fairly easy — we essentially add in roots one by one, which immediately proves existence. Uniqueness follows from uniqueness in adjoining a root of an irreducible polynomial (since adjoining any root is equivalent to adjoining an abstract one).

First, we'll prove existence by induction on the degree of P. Let P_1 be an irreducible factor of P, and let $F_1 = F[x]/(P_1)$, which (as we've seen earlier) is essentially the construction of adjoining a root of P_1 to F. Then in $F_1[x]$, P factors as $P(x) = (x - \alpha)Q(x)$, where α is a root of P_1.

Now let E be the splitting field for Q over F_1 (which exists by the induction assumption). Then we claim E is also a splitting field for P over F. This follows directly from the definition — suppose $\alpha = \alpha_1, \alpha_2, \dots, \alpha_n$ are the roots of P. Then P splits completely in $E[x]$. But we also have $E = F_1(\alpha_2, \dots, \alpha_n)$, and since $F_1 = F(\alpha_1)$, then $E = F(\alpha_1, \dots, \alpha_n)$.

So we've proved existence of the splitting field. To prove uniqueness, we again use induction. Suppose that E' is another splitting field; we'll construct an isomorphism between E' and E.

First, we can find a root α' of P_1 in E'. Then if we set $F_1' = F(\alpha') \subset E'$, we know that $F(\alpha') \cong F(\alpha)$, since both are isomorphic to the abstract construction $F[x]/(P)$.

This isomorphism between $F(\alpha)$ and $F(\alpha')$ sends $\alpha \mapsto \alpha'$. Suppose it sends $Q \in F_1[x]$ to $Q' \in F_1'[x]$, so we have $P = (x - \alpha')Q'$ (the isomorphism fixes F, and therefore P). Now E' is a splitting field for Q' over F_1'. So uniqueness of the splitting field of Q (which we know by the induction assumption) implies that the isomorphism between F_1 and F_1' extends to an isomorphism between E and E', and the two splitting fields are isomorphic. □

We'll see more proofs similar to this last idea later, where we construct isomorphisms between field extensions.

Student Question. *What happens if P has repeated roots?*

Answer. *In this case, it doesn't matter — for instance, the splitting field of P^2 is the same as that of P.*

26.2 Construction of Finite Fields

We'll now turn to *finite* fields. First notice that if F is a finite field, it can't contain $\mathbb{Q}$ (since $\mathbb{Q}$ is infinite), so it must contain $\mathbb{F}_p$ — we saw last class that every field contains $\mathbb{Q}$ or $\mathbb{F}_p$ for some prime p. Moreover, since F is finite as a set, the extension $F/\mathbb{F}_p$ is also finite, so F is finite-dimensional as a $\mathbb{F}_p$-vector space. Let $n = [F : \mathbb{F}_p] = \dim_{\mathbb{F}_p} F$. Then we must have $|F| = p^n$ — if we choose a basis for F (forgetting we can multiply elements, and only using the vector space structure), this identifies F with $\mathbb{F}_p^n$ (n-tuples of elements in $\mathbb{F}_p$, corresponding to the coordinates in this basis), which has p^n elements.

This was a fairly straightforward observation, but the converse is also true!

Theorem 26.2
For every prime p and every $n \geq 1$, there exists a field of $q = p^n$ elements. Furthermore, any two such fields are isomorphic.

As a result, we have a unique field of $q = p^n$ elements, which we denote by $\mathbb{F}_q$. Note that except when $n = 1$, the field $\mathbb{F}_q$ is very different from $\mathbb{Z}/q\mathbb{Z}$ (which is not a field). They're not even isomorphic as additive groups! (For example, $\mathbb{F}_2 = \mathbb{Z}/2\mathbb{Z}$, but $\mathbb{F}_4$ and $\mathbb{Z}/4\mathbb{Z}$ have very different structure.)

One way to construct a field of q elements would be to find an irreducible polynomial of degree n in $\mathbb{F}_p[x]$ (and quotient by that polynomial). This is easy to do when n is small — for example, if $p = 4k + 3$, the polynomial $x^2 + 1$ is irreducible, so

$$\mathbb{F}_{p^2} = \mathbb{F}_p[x]/(x^2 + 1).$$

Similarly, we have

$$\mathbb{F}_8 = \mathbb{F}_2[x]/(x^3 + x + 1).$$

But this is harder for general n — it's possible to prove one always exists via a counting argument (counting all polynomials, and all ways to produce a product of lower-degree polynomials), but this won't be the approach we use.

Instead, we'll use a sort of "magic trick" — we'll consider the *Artin–Schreier polynomial* $A(x) = x^q - x$.

Lemma 26.3
Let F be any field containing $\mathbb{F}_p$, and let $q = p^n$. Then the set of roots of A in F,

$$\{x \in F \mid x^q - x = 0\},$$

is a subfield of F.

This is quite exceptional! Usually, to construct a field from a polynomial, we adjoin roots and then take all possible sums and products. But in this case, when we take all roots, the result is actually *closed* under arithmetic operations — and we don't need to do anything more.

Proof. We must check that for $\alpha, \beta \in F$ with $A(\alpha) = 0$ and $A(\beta) = 0$, we have:

(1) $A(\alpha\beta) = 0$,

(2) $A(\beta^{-1}) = 0$ (if $\beta \neq 0$), and

(3) $A(\alpha + \beta) = 0$.

The first two are straightforward (and would work if we replaced q with *any* exponent) — for (1), since $\alpha^q = \alpha$ and $\beta^q = \beta$, we have $(\alpha\beta)^q = \alpha^q\beta^q = \alpha\beta$. We can check (2) similarly.

Now for (3), note that in any ring containing $\mathbb{F}_p$, we have

$$(x + y)^p = x^p = y^p.$$

This follows from the Binomial Theorem, since $\binom{p}{i} \equiv 0 \pmod{p}$ for $i = 1, \ldots, p - 1$ — if we write $\binom{p}{i} = p!/i!(p - i)!$, the numerator is divisible by p and the denominator is not. Now using induction, we see that $(x + y)^q = x^q + y^q$ if $q = p^n$ for any n. So $\alpha^q = \alpha$ and $\beta^q = \beta$ implies that

$$(\alpha + \beta)^q = \alpha^q + \beta^q = \alpha + \beta,$$

and thus $A(\alpha + \beta) = 0$. □

With this tool, we can now prove Theorem 26.2.

Proof of Theorem 26.2. We'll first prove uniqueness. Suppose that F is a field of $q = p^n$ elements. Now consider the multiplicative group $F^\times$ (the group of all nonzero elements of F under multiplication). This has $q - 1$ elements, so the order of any element of $F^\times$ must divide $q - 1$; therefore $\alpha^{q-1} = 1$ for all $\alpha \neq 0$. Then $\alpha^q = \alpha$ for *all* $\alpha \in F$ (since this is true for 0 as well). So $A(\alpha) = 0$ for all $\alpha \in F$.

But now we have a polynomial of degree q, which has q roots in F. The only way this happens is if the polynomial splits completely — we have that $x - \alpha \mid A(x)$ for all $\alpha \in F$, so by unique factorization in $F[x]$, the product of all terms $x - \alpha$ must divide $A(x)$ as well, and since $\deg(A(x)) = q$ (and both sides are monic), we must then have

$$A(x) = \prod_{\alpha \in F} (x - \alpha).$$

But then F is a splitting field of A over $\mathbb{F}_p$, and the uniqueness of F follows from the uniqueness of the splitting field.

To prove existence, we can simply let F be the splitting field of A over $\mathbb{F}_p$; we then need to check that $|F| = q$.

First, by Lemma 26.3, we have $A(\alpha) = 0$ for all $\alpha \in F$ — we know that F is *generated* by the roots of A, but the lemma implies that these roots are closed under arithmetic operations, so *all* elements of F are roots of A.

So then the number of elements in F is the number of roots of A (which splits completely in F). In particular, we immediately see that $|F| \leq \deg A = q$. To see that equality holds, it suffices to check that A has no multiple roots.

To check this, we use derivatives — in real or complex analysis, we know that a function has a higher-order root at a if a is also a root of the derivative. Of course, here we're in a much more abstract setting, and we can't define derivatives using limits. However, we can still use this idea, by using *formal derivatives* — the formal derivative of a polynomial $P(x) = a_n x^n + \cdots + a_0$ is the familiar formula $P'(x) = n a_n x^{n-1} + \cdots + a_1$. It's easy to check that the formulas $(P+Q)' = P' + Q'$ and $(PQ)' = PQ' + P'Q$ still hold. Now, if P has a multiple root at α, then $P(x) = (x-\alpha)^2 Q(x)$ for some Q, which means

$$P'(x) = 2(x-\alpha)Q(x) + (x-\alpha)^2 Q'(x)$$

also has a root at α. So it's still true that a multiple root of P is also a root of its derivative. In particular, if $\gcd(P, P') = 1$, then P has no multiple roots (in any field containing its field of coefficients).

But it's easy to compute the derivative of A — we have $A'(x) = (x^q - x)' = q x^{q-1} - 1$. But q is 0 in F (since F has characteristic p)! So $A'(x)$ is just -1, and $\gcd(A, A') = 1$. So A has no multiple roots, which means $|F| = q$. □

Once we have this construction, we can then derive concrete information about irreducible polynomials.

26.3 Structure of Finite Fields

There's more that we can say about the structure of $\mathbb{F}_q$.

Proposition 26.4
The multiplicative group $\mathbb{F}_q^\times$ is cyclic, and is therefore isomorphic to $\mathbb{Z}/(q-1)\mathbb{Z}$.

Proof. Since $\mathbb{F}_q^\times$ is a finite abelian group, it's isomorphic to $\prod \mathbb{Z}/p_i^{d_i}\mathbb{Z}$ for some d_i. By the Chinese Remainder Theorem, it's enough to show that each prime p appears at most once in this decomposition.

But assume for contradiction that some prime p appears twice (here p is used to denote any prime, not the characteristic of $\mathbb{F}_q$). Then there's at least p^2 elements of order dividing p, meaning that $\alpha^p = 1$ (since the group $\mathbb{Z}/p^d\mathbb{Z}$ contains p elements of order dividing p, so we can take one such element from each copy and elements of order 1 from the remaining terms in the product). But then the polynomial $x^p - 1$ would have p^2 roots; this is impossible, since a polynomial of degree p can have at most p roots. □

27 Finite Fields (continued)

Previously, we constructed the finite field $\mathbb{F}_q$ for $q = p^n$, and showed that there is a unique such field. This construction had the unusual property that the field *consisted* of exactly the roots of a polynomial (where the polynomial was $x^q - x$), rather than just being *generated* by the roots of a polynomial.

There is more that we can say about the structure of finite fields.

27.1 The Multiplicative Group

Lemma 27.1
If F is any field and G is a finite subgroup of $F^\times$, then G is cyclic.

Example 27.2
If $F = \mathbb{C}$, then finite subgroups of $F^\times$ are the nth roots of unity

$$\left\{\exp\frac{2\pi i}{n}\right\} = \langle\zeta_n\rangle \cong \mathbb{Z}/n.$$

Proof of Lemma 27.1. By the classification of finite abelian groups, we know $G \cong \prod \mathbb{Z}/p_i^{n_i}\mathbb{Z}$ for some integers n_i. So it's enough to check that no prime appears twice (meaning that every prime appears in the list of p_i at most once). Then we can use the Chinese Remainder Theorem to show that the product is cyclic, as all the p_i are then coprime. For example, $\mathbb{Z}/4\mathbb{Z} \times \mathbb{Z}/3\mathbb{Z} = \mathbb{Z}/12\mathbb{Z}$ is cyclic, while $\mathbb{Z}/4\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$ is not.

But suppose p appears twice. Then G contains a subgroup $\mathbb{Z}/p^a\mathbb{Z} \times \mathbb{Z}/p^b\mathbb{Z}$. But then G has at least p^2 elements of order dividing p, since there's p choices for the coordinate in each. This would mean the polynomial $x^p - 1$ has at least p^2 roots; but it has degree p, so this is impossible. □

Corollary 27.3
For any finite field $\mathbb{F}_q$, its multiplicative group $\mathbb{F}_q^\times$ is cyclic, meaning $\mathbb{F}_q^\times \cong \mathbb{Z}/(q-1)$.

Note 27.4
Although we know in theory that $\mathbb{F}_q^\times \cong \mathbb{Z}/(q-1)$, in practice it's hard to compute how this isomorphism works — it is difficult to find a generator, or to figure out what power to raise the generator to in order to get a given element. Many cryptography and encryption protocols are based on this.

Corollary 27.5
We have $\mathbb{F}_q \cong \mathbb{F}_p(\alpha)$, and therefore, there exists an irreducible polynomial of any degree over $\mathbb{F}_p$.

Proof. There exists $\alpha \in \mathbb{F}_q$ which generates the multiplicative group; then α must generate $\mathbb{F}_q$ as an extension of $\mathbb{F}_p$, since every element of $\mathbb{F}_q$ is a *power* of α. (The converse is false — it is possible to find α which generates the extension but not the multiplicative group.)

Then $\mathbb{F}_q = \mathbb{F}_p(\alpha) \cong \mathbb{F}_p[x]/(Q)$ where Q is the minimal polynomial of α. So Q is an irreducible polynomial of degree n, where $q = p^n$. In particular, a procedure (in theory) to find an irreducible polynomial of degree n is to write down $\mathbb{F}_{p^n}$, find a multiplicative generator, and take its minimal polynomial. □

27.2 Application to Number Theory

Finite fields arise in many areas of math and computer science — in particular, in number theory. One such example is $R/(p)$, where R is the ring of algebraic integers in a finite extension of $\mathbb{Q}$.

Example 27.6
If $p \equiv 3 \pmod 4$, then $\mathbb{Z}[i]/(p) \cong \mathbb{F}_{p^2}$ — it's a field since (p) is maximal.

The example we'll focus on is the extension $\mathbb{Q}(\zeta_\ell)$, where ℓ is a prime (it's possible to consider general ℓ, but the prime case is a bit simpler). We know this extension is $\mathbb{Q}[x]/(x^{\ell-1} + \cdots + 1)$, and we can check that its ring of algebraic integers is

$$R = \mathbb{Z}[x]/(x^{\ell-1} + \cdots + 1).$$

Guiding Question
For an (integer) prime p, when is $R/(p)$ a field (or equivalently, when is (p) maximal)?

It's clear that $R/(p) \cong \mathbb{F}_p[x]/(x^{\ell-1} + \cdots + 1)$, which means its dimension over $\mathbb{F}_p$ (as a vector space) is ℓ. So if $R/(p)$ is a field, then it must be $\mathbb{F}_{p^\ell}$.

We'll assume $p \neq \ell$.

Proposition 27.7
If $p \neq \ell$, then $R/(p)$ is a field if and only if $\operatorname{ord}_{\mathbb{F}_\ell^\times} p = \ell - 1$.

Here $\operatorname{ord}_{\mathbb{F}_\ell^\times} p$ denotes the multiplicative order of p mod ℓ; in particular, $\ell - 1$ is the largest possible order, since $\mathbb{F}_\ell^\times$ has $\ell - 1$ elements.

Proof. Let $\mathfrak{m}$ be a maximal ideal of R containing (p). Then we have $R/\mathfrak{m} \cong \mathbb{F}_{p^a}$ for some a. Let the image of ζ_ℓ in $R/\mathfrak{m}$ be $\overline{\zeta}_\ell$. Then we know $\overline{\zeta}_\ell^\ell = 1$ and therefore $\overline{\zeta_\ell}$ has multiplicative order ℓ; so since the multiplicative group of $\mathbb{F}_{p^a}$ has size $p^a - 1$, we get that $\ell \mid p^a - 1$.

Now if the order of p in $\mathbb{F}_\ell^\times$ is $\ell - 1$, then we must have $a \geq \ell - 1$. But it cannot be larger than $\ell - 1$. So then $a = \ell - 1$ and $R/\mathfrak{m} \cong R/(p)$, which means $R/(p)$ is a field. The converse can be proved similarly. □

Example 27.8
Suppose $p = 3$ and $\ell = 5$. Then $\operatorname{ord}_5 3 = 4$, so $R/(3)$ is a field.

27.3 Multiple Roots

In our construction of finite fields, one step had to do with multiple roots and derivatives. In particular, we used the fact that a multiple root of P is also a root of P' in order to show that the Artin–Schreier polynomial doesn't have multiple roots.

Guiding Question
Let $P \in F[x]$ be an irreducible polynomial. Can P have multiple roots in its splitting field (or equivalently, in any extension)?

If α is such a root, then α is also a root of P', and therefore a root of $\gcd(P, P')$ as well (where $\gcd(P, P')$ is the polynomial Q which generates (P, P') as an ideal).

But P is irreducible, and $\deg P' < \deg P$. So if $P' \neq 0$, then this means $\gcd(P, P') = 1$, and no such α can exist. However, it's possible that $P' = 0$. So the question reduces to the following:

Guiding Question
When can we have a nonconstant polynomial with $P' = 0$?

We have $(x^n)' = nx^{n-1}$. If $n \geq 1$, then if the field has characteristic 0, this is always nonzero. Meanwhile, if the field has characteristic p, then this is zero if and only if $p \mid n$. So if $P' = 0$, then we must have

$$P(x) = Q(x^p) = a_n x^{pn} + a_{n-1} x^{p(n-1)} + \cdots + a_0,$$

where $p = \text{char}(F)$. So we want to see when such a polynomial is irreducible.

If $F = \mathbb{F}_q$ is finite, then we know $a^q = a$ for all $a \in F$. This means we can extract pth roots of the coefficients, since $(a^{p^{n-1}})^p = a$ — so we can write $a_i = b_i^p$ for some $b_i \in F$. Then we have

$$P(x) = b_n^p x^{pn} + b_{n-1}^p x^{p(n-1)} + \cdots + b_0^p.$$

But this allows us to extract a pth root of the *polynomial*: we then have

$$P = (b_n x^n + b_{n-1} x^{n-1} + \cdots + b_0)^p.$$

On the other hand, there exist examples of such irreducible P in infinite fields. For instance, take $F = \mathbb{F}_q(t)$ to be the field of rational functions in t (or equivalently, the fraction field of $\mathbb{F}_q[t]$), and $P(x) = x^p - t$. This is irreducible, but its derivative is identically 0.

We won't study examples like this, but it's good to know they exist — in every situation we care about, irreducible polynomials can't have multiple roots.

Definition 27.9
An extension E/F is **separable** if the minimal polynomial (over F) of every algebraic element $\alpha \in E$ has no multiple roots.

So if F has characteristic 0 or is finite, then every extension is separable. We'll only look at these instances, so we will generally assume all our extensions are separable.

27.4 Geometry of Function Fields

Another important example of a field is $\mathbb{C}(t)$; we can think of the extensions of $\mathbb{C}(t)$ via geometry. Let $F = \mathbb{C}(t)$, and suppose $E = F[x]/(P)$ is a finite extension of F, where P is an irreducible polynomial. As with integers, we can scale so that P is a primitive polynomial in $\mathbb{C}[t][x]$.

We can then think of P as a polynomial in two variables, meaning $P \in \mathbb{C}[t, x]$. So another way to think of these extensions is that $F = \text{Frac}(\mathbb{C}[t])$, and $E = \text{Frac}(R)$ where $R = \mathbb{C}[t, x]/(P)$.

As discussed earlier, these rings are worked with in algebraic geometry — to connect them to geometry, we consider the maximal spectrum

$$X = \text{MSpec}(R) = \{(a, b) \in \mathbb{C}^2 \mid P(a, b) = 0\}$$

(which describes all maximal ideals of R). Also define $Y = \text{MSpec}(\mathbb{C}[t]) = \mathbb{C}$. Then we have a map $X \to Y$ sending $(a, b) \mapsto a$.

Student Question. *If P is irreducible, is R a field extension?*

Answer. *No, R is not a field. Polynomials in two variables aren't a PID (so even if P is irreducible, (P) is generally not maximal) — if they were, algebraic geometry would be trivial.*

Example 27.10
Let $P(t, x) = x^n - t$.

Then $R \cong \mathbb{C}[x]$, where this map sends a complex number to its nth power (since $t = x^n$). Each point in $\mathbb{C}$ has n complex nth roots (except 0), giving a *ramified covering* (with a ramification point at 0). One way to represent this geometrically is to draw the t-plane and the x-plane. In the t-plane, we make a cut along the x-axis, turning it into two half-planes glued together.

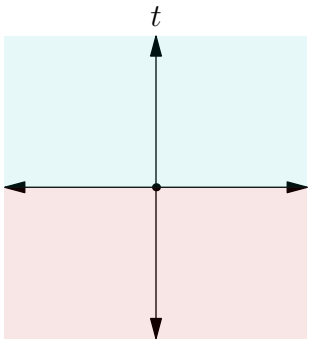

For a point on the x-plane, raising it to the nth power multiplies the angle by n. So we cut the x-plane into $2n$ pieces (colored by which half-plane their points are mapped to):

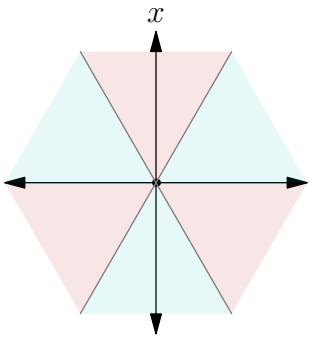

This describes the geometry of the map raising x to the nth power.

Example 27.11
Let $P(x,t) = x^2 - t(t+1)(t-\lambda)$. (Any P consisting of x^2 minus a cubic polynomial can be written in this form, by a change of variables.) For simplicity, assume $\lambda \in \mathbb{R}$.

We can again draw the $\mathbb{C}$-plane. We again have a ramified double covering, with three ramification points — $t = 0$, -1, and λ (for every other point, there are two square roots). So we can again make a cut and create two half-planes.

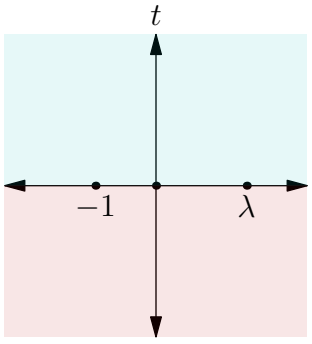

For each half-plane, its pre-image breaks into two pieces (corresponding to the two branches of the square root — we can start with the positive or negative one). So the pre-image consists of two blue rectangles and two red rectangles:

We then need to glue these rectangles together, by thinking about the values of these functions. When glued together, they look like a bagel (where we cut the bagel horizontally and through its middle).

Note 27.12
These situations require more background to describe rigorously, and for that reason they are usually not presented in algebra classes; but they are important examples of field extensions, and mathematicians often have these examples in mind even when constructing algebraic arguments about number fields.

28 Geometry of Function Fields

Last time, we began discussing finite extensions of the function field $F = \mathbb{C}(t)$; we can write such an extension as $E = F[x]/(P)$ for an irreducible $P \in F[x]$. Since the coefficients of P are rational coefficients, we can clear denominators and then think of P as a polynomial in two variables — and by factoring out the gcd of the coefficients, we can assume that $P \in \mathbb{C}[t, x]$ is primitive, and therefore irreducible in $\mathbb{C}[t, x]$ as well. Then we can consider the ring $R = \mathbb{C}[t, x]/(P)$, so then $E = \text{Frac}(R)$.

A geometric way to think about this situation is that we have

$$X = \text{MSpec}(R) = \{(a, b) \in \mathbb{C}^2 \mid P(a, b) = 0\}$$

(so X consists of the set of zeros of the polynomial), and this maps to

$$Y = \text{MSpec}(\mathbb{C}[t]) = \mathbb{C}.$$

Here X is also called a **Riemann surface**, and we can think of E as a field of rational functions on that Riemann surface.

28.1 Ramified Covers

Recall the simpler example discussed last class:

Example 28.1
Consider $P(x, t) = x^n - t$.

We saw earlier that this corresponds to the following picture. Here the t-plane represents Y, and the x-plane represents X — by definition X corresponds to $(t, x) \in \mathbb{C}^2$ such that $t = x^n$, but such points can just be described by x.

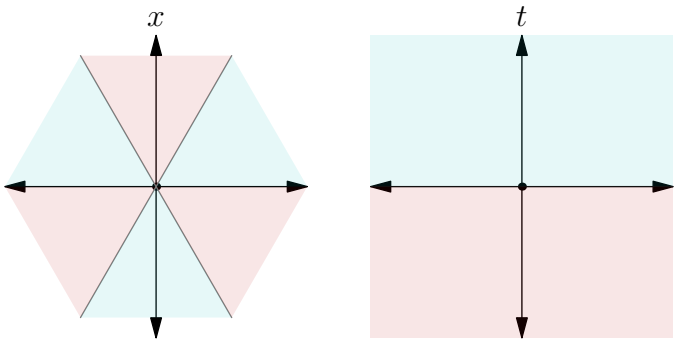

These two maps are related by the map $x \mapsto x^n$ (since when we go from X to Y, we're mapping each point to its corresponding value of t, which here is x^n), which unwraps each smaller angle to a 180° angle.

As mentioned earlier, this is a ramified cover — all but finitely many values of t have the same number of points in their pre-image (and in general, this number of points is equal to the degree of P, as a polynomial in x). But some points give smaller fibers — for the map $x \mapsto x^n$, the point 0 only has one element in its pre-image, and is therefore a ramification point.

Incidentally, the terminology we saw for the behavior of primes, regarding how they split in quadratic extensions, also included *ramified prime*. This isn't a coincidence — they describe the same phenomenon.

Returning to the general situation, let $a_1, \ldots, a_k$ be the ramification points, meaning that $|f^{-1}(a_i)| < n$ for all i, and $|f^{-1}(a)| = n$ for all a not equal to any of the a_i. Now take a generic point a (not equal to any a_i). Then there are n points lying over a (meaning points whose image is a):

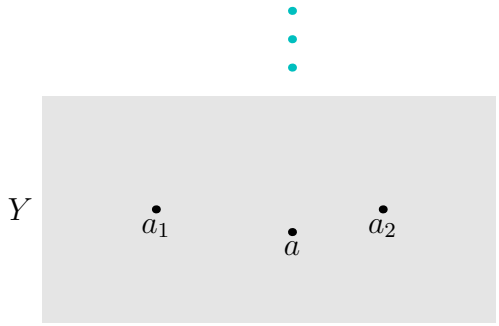

Now consider a closed loop around a, which doesn't pass through any of the ramification points.

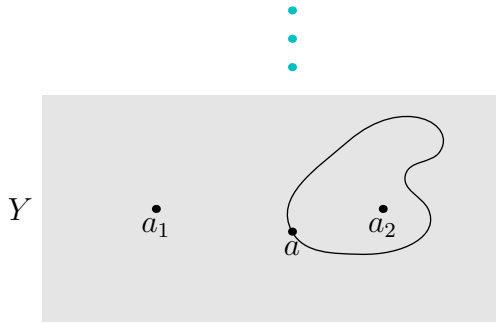

This loop determines a permutation of the fibers (meaning the n points above a): suppose we start at one of these n points. Then we can lift the loop locally in a unique way (since the loop avoids the ramification points). But when we return, we may return to a different one of these points:

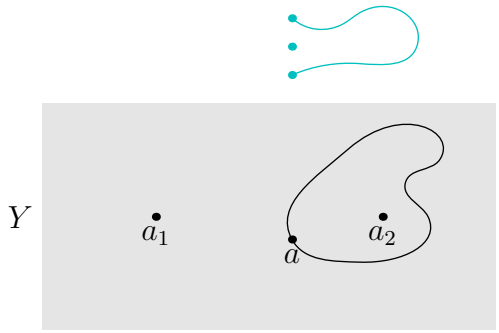

More precisely, if we have a loop $\gamma : [0,1] \to Y$ avoiding the ramification points, with $\gamma(0) = \gamma(1) = a$, then γ defines a permutation σ_γ of the set $f^{-1}(a)$ — to compute $\sigma_\gamma(x)$, we lift γ to a map $\widetilde{\gamma} : [0,1] \to X$ such that $\gamma(0) = x$; then $\gamma(1)$ is another point $y \in f^{-1}(a)$, and we set $\sigma_\gamma(x) = y$.

Example 28.2
Consider the example $P(x,t) = x^n - t$, and let γ be the unit circle (the standard loop).

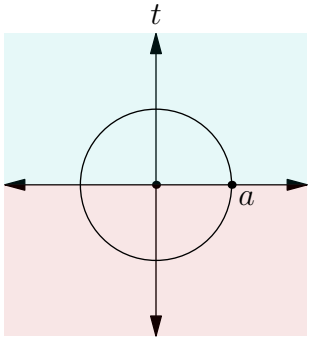

When we walk on the x-plane, we're still walking in a circle, but we walk n times slower (since x is raised to the nth power):

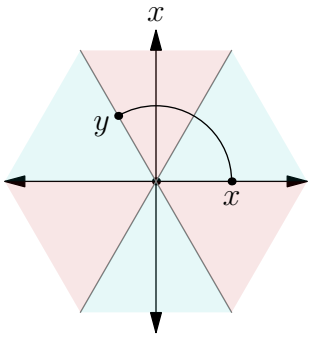

Explicitly, we have $X = \mathbb{C}$ and $Y = \mathbb{C}$, with $f(x) = x^n$. We have that $f^{-1}(1) = \{\exp(2\pi ik/n)\} = \{\zeta_n^k\}$ is the set of nth roots of unity (where $\zeta_n = \exp(2\pi i/n)$). Our loop is defined by $\gamma(t) = \exp(2\pi it)$, and if we start at ζ_n^k, then $\widetilde{\gamma}(t) = \zeta_n^k \exp(2\pi it/n)$. So the permutation is $\sigma_\gamma(\zeta_n^k) = \zeta_n^{k+1}$.

A term for this permutation is the **monodromy**.

A useful fact, which we will not prove, is the following:

Theorem 28.3
If E/F is a splitting field of some polynomial, then σ_γ extends to an automorphism of X which is the identity on Y, coming from an automorphism of E which is the identity on F.

The point is that σ_γ always gives us a permutation of the points in the pre-image of a fixed point a; but in the case of a splitting field, this can be *extended* to an automorphism of all of X.

Example 28.4
In our example situation, where $F = \mathbb{C}(t)$ and $E = \mathbb{C}(t)[x]/(x^n - t)$ (which is a splitting field), then the automorphism corresponding to the unit circle is $x \mapsto \zeta_n x$, which sends $t \mapsto t$.

Example 28.5
In the other example from last class, where $E = \mathbb{C}(t)[x]/(x^2 - t(t+1)(t-\lambda))$, there are three ramification points (0, -1, and λ), and we have a double cover (there's two points above all points except the ramification points). The automorphism is $x \mapsto -x$, which swaps the two points (and again fixes t).

Later, we'll discuss more algebraic ways to construct automorphisms of fields.

Student Question. *Are there always ramification points?*

Answer. *If $Y = \mathbb{C}$, then the answer is yes — this follows from topological reasons, as $\mathbb{C}$ is simply connected. But if Y is some other Riemann surface, there may be no ramification points.*

Note that if γ is deformed continuously, while still avoiding the ramification points and fixing the beginning and end, then the permutation described doesn't change. In our example $P(x,t) = x^n - t$, any closed loop which goes around 0 once will give the same permutation:

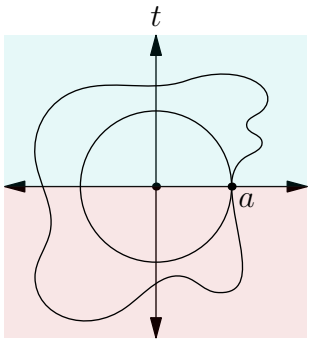

28.2 The Main Theorem of Algebra

There is a proof of the main theorem of algebra using ideas similar to the ones we've seen here.

Theorem 28.6
The field $\mathbb{C}$ is algebraically closed — in other words, every nonconstant polynomial $P \in \mathbb{C}[x]$ has a root.

The proof uses the concept of a **winding number**: suppose we have a continuous map $\gamma : [0,1] \to \mathbb{C} \setminus \{0\}$. Then its winding number, informally, is the number of times γ goes around 0 (counted with sign — going around 0 counterclockwise is counted with positive sign, and clockwise with negative sign). It can actually be formally defined using similar ideas to the ones we've seen here — consider σ_γ for the exponential map where $X = \mathbb{C}$ and $Y = \mathbb{C} \setminus \{0\}$, and we send $x \mapsto \exp(x)$. Then the pre-image of $z \in Y$ is $\exp^{-1}(z) = \{\log z + 2\pi i n\}$ for integers n. So for a loop γ, if we construct the loop $\widetilde{\gamma} : [0,1] \to \mathbb{C}$ as before, so that $\exp(\widetilde{\gamma}(t)) = \gamma(t)$, then we have

$$\gamma(1) = \gamma(0) + 2\pi i n$$

for some integer n. The winding number is defined to be the integer n in this equation.

We use $w(\gamma)$ to denote the winding number of γ.

Lemma 28.7
If $\gamma(t) = \gamma_1(t)\gamma_2(t)$, then $w(\gamma) = w(\gamma_1) + w(\gamma_2)$.

Proof. We have $\widetilde{\gamma}(t) = \widetilde{\gamma_1}(t) + \widetilde{\gamma_2}(t)$, so the discrepancies $2\pi i n$ are added together as well. □

Now we are ready to prove the theorem.

Proof of Theorem 28.6. Consider a polynomial $P(z) \in \mathbb{C}[z]$, and assume for contradiction that $P(z) \neq 0$ for all $z \in \mathbb{C}$. Let

$$P(z) = z^n + a_{n-1}z^{n-1} + \cdots + a_0.$$

Now consider loops

$$\gamma_r(t) = P(re^{2\pi i t})$$

for each $r \geq 0$ (where we take the circle of radius r around the origin, and see what P does to it). Since P has no zeros, this is a loop in $\mathbb{C} \setminus \{0\}$.

Now consider the winding number of each loop γ_r. We can observe three properties, which together lead to a contradiction: first, $w(\gamma_r)$ is independent of r — this is clear because it depends continuously on r (since none of our loops pass through 0).

Second, when $r = 0$, $w(\gamma_0) = 0$ — this is because γ_0 is the constant loop.

Finally, when r is large, we claim $w(\gamma_r) = n$. To prove this, we can write

$$P(z) = z^n \left(1 + \frac{a_{n-1}}{z} + \cdots + \frac{a_0}{z^n}\right).$$

Let z^n correspond to γ_1, and the remaining factor to γ_2. Then z^n runs around a circle n times, so $w(\gamma_1) = n$. Meanwhile, if r is very large, then

$$\left|\frac{a_{n-1}}{z} + \cdots + \frac{a_0}{z^n}\right| < 1,$$

which means γ_2 is trapped inside the circle of radius 1 centered at 1:

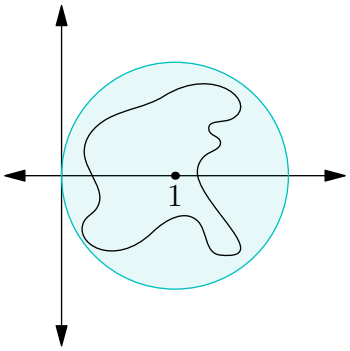

This means it's trapped to the right of the y-axis, so it can't wind at all; so $w(\gamma) = w(\gamma_1) + 0 = n$. □

Note 28.8
The textbook doesn't split $\gamma = \gamma_1\gamma_2$, but it has a nice intuitive explanation — essentially, the loop γ_r is fairly close to the loop which goes around the circle n times. Imagine having a dog on a leash, and walking around a circle n times. As long as the leash is short enough, the dog may run around you in any way it wants, but it will still go around the center of the circle the same number of times that you do. (The terms after 1 in the second factor correspond to the additional movement of the dog.)

28.3 The Primitive Element Theorem

Next class, we will begin discussing Galois theory. The following theorem will be useful:

Theorem 28.9
If E/F is a finite separable extension, then the extension is generated by one element — meaning $E = F(\alpha)$ for some α.

So even if E was defined by adjoining multiple elements (for example, the splitting field construction), it's possible to obtain it just by adjoining one element.

Recall that all finite fields and fields of characteristic 0 are separable; these are the only cases we will work with.

Proof. If F is finite, then E is also finite; so $E = F(\alpha)$ where α is a multiplicative generator of E.

Now assume F is infinite. It's enough to prove this in the case where $E = F(\alpha, \beta)$ is generated by two elements (since we must have $E = F(\alpha_1, \dots, \alpha_n)$ for some finite n, and we can use induction on n).

Define $\gamma_t = \alpha t + \beta$. Then we'll show that for all but finitely many t, γ_t is a generator of E.

Let P be the minimal polynomial of α and Q the minimal polynomial of β, and let $K \supset E$ be a field where both P and Q split completely. Then we can write

$$P(x) = (x - \alpha_1)(x - \alpha_2)\cdots(x - \alpha_m)$$

and

$$Q(x) = (x - \beta_1)(x - \beta_2)\cdots(x - \beta_n),$$

where the α_i are all distinct, and the β_i are all distinct (by separability). Assume $\alpha = \alpha_1$ and $\beta = \beta_1$.

The condition on t we will specify (for $\alpha t + \beta$ to be a generator) is that the nm elements $\alpha_i t + \beta_j$ are all distinct. There are clearly finitely many t for which this isn't true.

It suffices to check that α and β are in $E' = F(\gamma)$ (where $\gamma = \gamma_t$ for some such t). We'll look at what polynomial equations we can write for α over E'. One is obvious — α is a root of $P(x)$. But α is also a root of the polynomial $Q(\gamma - tx)$, which we'll denote by $Q_1(x)$ — this is a polynomial with coefficients in E', and when we plug in α we get $Q(\beta) = 0$.

So then α is a root of the polynomial $S(x)$ which generates the ideal (P, Q) (also known as $\gcd(P, Q)$), working in $E'[x]$. But this gcd doesn't depend on the field in which it's computed — so $S(x)$ is also a generator of (P, Q)

in $K[x]$. And in $K[x]$, both polynomials split completely; and they have a unique common linear factor, namely $x - \alpha$ (by the condition on t).

Then $S(x)$ is a constant times $x - \alpha$; this means $\alpha \in E'$, and then $\beta \in E'$ as well. This means $E' = E$. $\square$

29 Galois Theory

29.1 Review: Primitive Element Theorem

Last class, we proved the Primitive Element Theorem:

Theorem 29.1
If E/F is a finite separable extension, then $E = F(\alpha)$ for some α.

Example 29.2
Let $F = \mathbb{Q}$ and $E = \mathbb{Q}(\sqrt{3}, \sqrt[3]{7})$. Then all but finitely many linear combinations $\alpha = r\sqrt{3} + \sqrt[3]{7}$ (with $r \neq 0$) generate E. The exceptions correspond to ratios $r = (\beta_i - \beta_j)/(\alpha_k - \alpha_\ell)$, where we take a different square root of 3 or cube root of 7. In this case, such r are not even real, so they are certainly not rational; and there are no exceptions.

Recall that an extension is separable if the minimal polynomial of any $\alpha \in E$ has nonzero derivative, or equivalently, has no multiple roots. All extensions are separable when working in characteristic 0 or with finite fields; more generally, separability is sometimes true and sometimes not. But we will only work with characteristic 0 and finite fields in this class, so we will assume all extensions are separable.

29.2 The Galois Group

Our main object of study is the Galois group:

Definition 29.3
The **Galois group** of an extension E/F, denoted $\mathrm{Gal}(E/F)$, is the group of automorphisms of E which are the identity on F.

Example 29.4
The Galois group $\mathrm{Gal}(\mathbb{C}/\mathbb{R})$ consists of two elements — the identity and complex conjugation.

The Galois group can store a lot of information about the structure of the field extension. But it only works well for *some* classes of extensions — we'll see that the extensions for which it works well are exactly the splitting fields. For this reason, we'll look at one more preliminary result, which is somewhat surprising:

Theorem 29.5
Suppose that E/F is a splitting field of some polynomial. Then for *any* $\alpha \in E$, the minimal polynomial of α must split completely (into linear factors) in E.

Example 29.6
Take $F = \mathbb{Q}$, and E to be the splitting field of $x^5 - 2$; then $E = \mathbb{Q}(\sqrt[5]{2}, \zeta_5)$. By the Primitive Element Theorem, it's generated by one element $\alpha = \sqrt[5]{2} + \zeta_5$. We know ζ_5 has degree 4 over $\mathbb{Q}$ and $\sqrt[5]{2}$ has degree 5, and $x^5 - 2$ remains irreducible even after adjoining a fifth root of unity; so $[E : \mathbb{Q}] = 20$. Then the minimal polynomial of α over $\mathbb{Q}$ has degree 20.

The theorem then states that all 20 complex roots of this minimal polynomial are inside E. We can actually explicitly describe these roots — they're of the form $\sqrt[5]{2}\zeta_5^i + \zeta_5^j$ for some integers $0 \leq i \leq 4$ and $1 \leq j \leq 4$ (by varying our choice of fifth root of 2 and primitive 5th root of unity), which are indeed in E.

Proof of Theorem 29.5. The proof is somewhat abstract; we'll deduce this theorem from the uniqueness of the splitting field.

Suppose E is the splitting field of some polynomial Q, and fix $\alpha \in E$ with minimal polynomial P; then we want to show that P splits completely in E.

Let $K \supset E$ be a field where P splits completely (for example, the splitting field for P over E). Then in K, we have

$$P(x) = (x - \alpha_1)(x - \alpha_2) \cdots (x - \alpha_n),$$

where $\alpha_i \in K$ for all i and $\alpha_1 = \alpha$. We need to check that $\alpha_i \in E$ for all i (and we already know $\alpha \in E$).

But we have that $F(\alpha_i) \cong F(\alpha)$, since α_i and α have the same minimal polynomial. Now consider E from the perspective of these two intermediate fields. We know that E is the splitting field for Q over $F(\alpha)$. We don't yet know whether α_i is in E or not, but we know that if we adjoin it, then $E(\alpha_i)$ is the splitting field for Q over $F(\alpha_i)$ — this is clear from the definition of the splitting field (the polynomial clearly splits in $E(\alpha_i)$, but E is generated over F by the roots of P; so $E(\alpha_i)$ is generated over $F(\alpha_i)$ by the roots of P as well).

But by the uniqueness of the splitting field, there exists an isomorphism $E \cong E(\alpha_i)$ extending the isomorphism $F(\alpha) \cong F(\alpha_i)$ (we've identified $F(\alpha)$ and $F(\alpha_i)$, and since Q has coefficients in F (which is preserved by the isomorphism), it's the same polynomial in both fields — the isomorphism takes one copy of Q to the other). Since our isomorphism is the identity on F, this means

$$[E : F] = [E(\alpha_i) : F],$$

and since $E(\alpha_i) \supset E$, this means we must have $E(\alpha_i) = E$, and therefore $\alpha_i \in E$. □

Student Question. *How did we show $E(\alpha_i)$ is the splitting field of Q over $F(\alpha_i)$?*

Answer. *More explicitly, we can suppose $E = F(\beta_1, \ldots, \beta_n)$, where the β_j are the roots of Q. Then $E(\alpha_i)$ is obtained by adjoining $\beta_1, \ldots, \beta_n$ as well as α_i. But we can also adjoin these in a different order, as $E(\alpha_i) = F(\alpha_i)(\beta_1, \ldots, \beta_n)$. So as an extension of $F(\alpha_i)$, it's generated by the roots of Q.*

Now we can get to some interesting results.

Proposition 29.7
For any finite (separable) extension E/F, we have

$$|\mathrm{Gal}(E/F)| \leq [E : F],$$

with equality if and only if E is the splitting field of some polynomial.

Example 29.8
We saw earlier that $|\mathrm{Gal}(\mathbb{C}/\mathbb{R})| = 2$. Meanwhile, it's a degree 2 extension, and it's the splitting field of $x^2 + 1$.

Proof of Proposition 29.7. Using the Primitive Element Theorem, we can let $E = F(\alpha)$ for some α. Then

$$[E : F] = \deg(\alpha) = \deg P,$$

where P is the minimal polynomial of α.

Meanwhile, an automorphism $\sigma : E \to E$ which fixes F is clearly uniquely determined by $\sigma(\alpha)$ (since α generates the extension), so it suffices to find the number of possible choices for $\sigma(\alpha)$. But $\sigma(\alpha)$ can be any root of the minimal polynomial of α; so $|\mathrm{Gal}(E/F)|$ is equal to the number of roots of P in E.

The number of roots of P in E is at most $\deg P$, which immediately proves

$$|\mathrm{Gal}(E/F)| \leq [E : F].$$

If equality holds, then P must split completely in E; this immediately implies that E is the splitting field of P (since E is also generated by a root α of P).

On the other hand, if E is a splitting field, then by Theorem 29.5, P must split completely in E. Since P cannot have multiple roots (by separability), this means it has exactly $\deg P$ roots in E, and therefore there are exactly $\deg P$ automorphisms. □

Student Question. *Was the condition that E/F is separable only used to show that P does not have multiple roots?*

Answer. *We also used it when using the Primitive Element Theorem; but in fact, it's not necessary to rely on the Primitive Element Theorem for that step, and* $|\mathrm{Gal}(E/F)| \leq [E:F]$ *is always true (it's possible to induct on the number of generators instead). So it's possible to avoid assuming separability there, but it* is *necessary for the last step.*

Definition 29.9
A finite extension E/F is **Galois** if $[E:F] = |\mathrm{Gal}(E/F)|$.

29.3 Main Theorem

The main theorem we will discuss is the following:

Theorem 29.10
If E/F is a Galois extension with Galois group $\mathrm{Gal}(E/F)$, then there is a bijection between subgroups of G, and intermediate subfields $F \subseteq K \subseteq E$ — where a subgroup $H \subset G$ is mapped to its **fixed field**

$$K = E^H = \{x \in E \mid \sigma(x) = x \text{ for all } \sigma \in H\},$$

and a subfield K is mapped to the set of $\sigma \in G$ which fix all elements of K (which by definition is $\mathrm{Gal}(E/K)$).

This bijection has many properties. For now, note that E/K is still a Galois extension, so if $H \mapsto K_H$, then $|H| = [E : K_H]$.

Student Question. *Can any finite group be a Galois group?*

Answer. *Yes, although whether any finite group can be a Galois group over* $\mathbb{Q}$ *is still open. We'll later discuss how to construct an extension with* S_n *as its Galois group, and any finite group is a subgroup of some* S_n*.*

We will discuss the proof and some applications later; first, we will discuss how to compute the Galois group in a few examples (there is no *general* easy answer).

29.4 Examples of Galois Groups

For a polynomial P with splitting field E (over F), we use $\mathrm{Gal}(P)$ to refer to $\mathrm{Gal}(E/F)$.

It's not easy to compute the Galois group — it's not even easy to compute the *degree* of the extension. One observation we can make is that G acts faithfully on the roots of P (meaning it permutes these roots). There is a bit more we can say:

Proposition 29.11
If P is irreducible, this action is transitive — any root can be sent to any other root.

Proof. Write $P(x) = (x - \alpha_1) \cdots (x - \alpha_n)$; then we want to show that for any i and j, there exists $\sigma \in \mathrm{Gal}(P)$ such that σ sends $\alpha_i \mapsto \alpha_j$.

But we know $F(\alpha_i) \cong F(\alpha_j)$ (since α_i and α_j have the same minimal polynomial). Further (similarly to the argument we used in the proof of Theorem 29.5), E is the splitting field of P over both $F(\alpha_i)$ and $F(\alpha_j)$. So by the uniqueness of the splitting field, the isomorphism between $F(\alpha_i)$ and $F(\alpha_j)$ extends to an isomorphism $\sigma : E \to E$ which sends $\alpha_i \mapsto \alpha_j$. □

However, this is pretty much everything we can say in general — even knowing that $\mathrm{Gal}(P)$ acts transitively on a set of n elements, its size can range from n to $n!$.

Example 29.12
Let $F = \mathbb{Q}$, and $E = \mathbb{Q}(\zeta_n)$ where $\zeta_n = \exp(2\pi i/n)$. For simplicity, assume $n = p$ is prime.

Solution. Let $\zeta = \zeta_p$. We proved earlier that $[E:F] = p - 1$, and E is the splitting field of $x^{p-1} + x^{p-2} + \cdots + 1$.

Any automorphism $\sigma \in \mathrm{Gal}(\mathbb{Q}(\zeta_n)/\mathbb{Q})$ must send $\zeta \mapsto \zeta^i$ for some $1 \leq i \leq p-1$ (since every root of unity is some power of ζ). Denote this automorphism by σ_i.

In order to compute the group, it suffices to understand how these automorphisms compose — we have

$$\sigma_i \sigma_j(\zeta) = \sigma_i(\zeta^j) = \zeta^{ij},$$

which means $\sigma_i \sigma_j = \sigma_{ij}$. So in this case, we have

$$\mathrm{Gal}(E/F) = (\mathbb{Z}/p\mathbb{Z})^\times \cong \mathbb{Z}/(p-1)\mathbb{Z}. \qquad \square$$

In this case, we were lucky because all roots of the polynomial were powers of one root ζ. In general, after fixing one root and sending it to another root, we still need to figure out what we can do with the remaining roots (which is forced by the algebraic relations between the roots).

Student Question. *The main theorem relates the subgroups of G to fixed fields; are there any interesting properties of the relations here?*

Answer. *We'll discuss that more next class — but using this result, you can actually solve the compass and ruler problem, which we'll discuss next time.*

This is an example of a case where the Galois group is as small as possible. There are also examples where the Galois group is as *big* as possible:

Example 29.13
Let P be an irreducible polynomial over $\mathbb{Q}$ of degree n, and suppose that P has exactly $n-2$ real roots, and 2 roots which are complex conjugates. Also suppose that $n = p$ is prime. Then if E is the splitting field of P, we have

$$\mathrm{Gal}(E/\mathbb{Q}) = S_n.$$

Proof Sketch. We'll just discuss the outline today, and prove this in more detail next time. The proof relies on an algebraic lemma — if $G \subset S_n$ where n is prime, and G contains a transposition (i, j) and a long cycle, then $G = S_n$. Here the transposition is given by complex conjugation, and the long cycle comes from the fact that the Galois group permutes the roots transitively, so its order divides $\deg P = p$; by Sylow's Theorems it must then contain an element of order p, and the only such element is a long cycle. $\square$

30 Main Theorem of Galois Theory

Last class, we introduced the main theorem:

Theorem 30.1
If E/F is a Galois extension (i.e. $|\mathrm{Gal}(E/F)| = [E : F]$), then intermediate subfields $F \subset K \subset E$ are in bijection with subgroups $H \subset G$, where a subgroup H is mapped to its **fixed field** E^H, and a subfield K is mapped to the set of $g \in G$ which fix all elements of K.

Recall that E/F is Galois if and only if E is a splitting field of some polynomial (and is separable).

30.1 Examples of Galois Groups

Last class, we saw the following example:

Example 30.2
If $F = \mathbb{Q}$ and $E = \mathbb{Q}(\zeta)$, where ζ is a pth root of unity for p prime, then $\mathrm{Gal}(E/F) = (\mathbb{Z}/p\mathbb{Z})^\times \cong \mathbb{Z}/(p-1)\mathbb{Z}$.

This is an example where the Galois group is as small as possible (given the degree of the polynomial). Meanwhile, we also saw an example in the opposite direction:

Example 30.3
If $P \in \mathbb{Q}[x]$ is irreducible, has degree p (for p prime), and has exactly $p-2$ real roots, then if E is its splitting field, we have $\mathrm{Gal}(E/F) = S_p$.

It's easy to write down such a polynomial. For example, $P(x) = 2x^5 - 10x + 5$ is irreducible by Eisenstein's criterion, while by graphing (or by computing its derivative) we can see that it has three real roots.

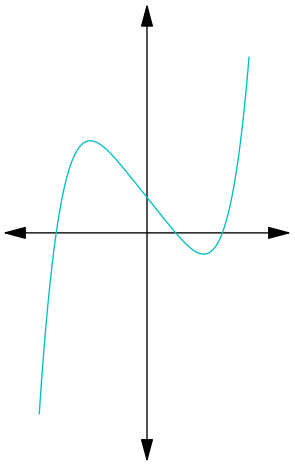

Proof. The proof is based on a lemma about the symmetric group:

Lemma 30.4
Suppose p is prime, and $G \subset S_p$ such that G acts on $[1, \dots, p]$ transitively, and G contains a transposition (ij). Then $G = S_p$.

First we'll show how the lemma implies that $\mathrm{Gal}(E/F) = S_p$ in our example: if we let $G = \mathrm{Gal}(E/F)$, then we know G acts (by permutations) on the p roots. We know that since P is irreducible, this action is transitive (we can send any root to any other root, which we proved by comparing the abstract procedure of adjoining a root to the procedure of adjoining a specific root).

Meanwhile, complex conjugation permutes the roots of P, so $E = \mathbb{Q}(\alpha_1, \dots, \alpha_p)$ is invariant under complex conjugation — this means it's an element in the Galois group. This element clearly permutes the two non-real, and fixes the real roots; so it is a transposition.

We'll now prove the group-theoretic lemma.

Proof of Lemma 30.4. Since G acts transitively on $[1, \dots, p]$, it follows that $p \mid |G|$; so by the Sylow Theorems, G has an element of order p, which we denote by σ. Recalling the description of elements of S_p using cycle notation, we can see that the only element of order p is a long cycle (meaning a cycle of all p elements), so σ is a long cycle.

Now recall that G also contains a transposition, which we can without loss of generality assume is (12).

Now we can find some $1 \leq i < p$ such that σ^i sends $1 \mapsto 2$ (since we can follow the arrows from 1 until we reach 2). But since $i < p$, the order of $\gamma = \sigma^i$ is also p; so γ is also a long cycle.

Now G contains (12) and γ, which we can without loss of generality assume is $(123 \cdots p)$. Now recalling how conjugation in the symmetric group works (by essentially taking the same cycles, but substituting different elements into them), we have

$$\gamma(12)\gamma^{-1} = (23),\ \gamma^2(12)\gamma^{-1} = (34),\ \dots.$$

So then G contains all the standard transpositions (12), (23), (34), ... (which transpose two consecutive elements). It's a standard fact about the symmetric group that these transpositions generate S_p (given any permutation, we can swap consecutive elements to eventually sort it), so $G = S_p$. □

This shows that $\mathrm{Gal}(E/F) = S_p$, as desired. □

Student Question. *How did we show that $p \mid |G|$ in the proof of the lemma?*

Answer. *If G acts on X transitively, then we have the formula*

$$|G| = |X| \cdot |\mathrm{Stab}_G\, x|$$

for any $x \in X$. To show this, we can break the elements of G into subsets based on where they send a given point x; there will be $|X|$ groups, and each group will have $|\mathrm{Stab}_G\, x|$ elements.

Student Question. *Is the Galois group always a transitive subgroup of S_n?*

Answer. *The Galois group is always a subgroup of S_n (since it always permutes the roots of the polynomial). It's always transitive if the polynomial is irreducible, but the polynomial doesn't have to be irreducible in general (we can consider the splitting field of* any *polynomial).*

Student Question. *Does this work when there's more than two complex roots? Or does the Galois group become smaller than S_n?*

Answer. *The particular trick used in this argument doesn't work. But if you write a random polynomial, its Galois group should be S_n — for the group to be smaller, you'd need some condition on the roots.*

Note 30.5

A similar argument can be used to produce an extension of the rational function field $F = \mathbb{C}(t)$ with Galois group S_n, for n prime.

Earlier, we saw a ramified covering $X \to Y = \mathbb{C}$, where X is the zero set of a polynomial $P(t, x) \in F[x]$. The analog of our conditions here becomes that there should be one *simple* ramification point (meaning that at this ramification point y_0, there are $n-1$ pre-images $x_1, \dots, x_{n-1}$; f is an isomorphism at $x_2, \dots, x_{n-1}$, while at x_1, f looks (locally) like the map $z \mapsto z^2$.

Then if E is the splitting field of P, we can show $\mathrm{Gal}(E/F) = S_n$ in a similar way — P is assumed to be irreducible, and the condition on ramification ensures that the Galois group contains a transposition.

30.2 Proof of Main Theorem

We'll now prove the theorem.

Lemma 30.6

Both maps in the correspondence send $[E : K]$ to $|H|$.

Note that in the tower of extensions $E/K/F$, we're looking at the degree of the *top* extension E/K, rather than the bottom one K/F.

Proof. One direction is clear — if we start with a subfield K, the corresponding subgroup is $\mathrm{Gal}(E/K)$ (we can forget that F exists, and just look at the extension E/K — then we've defined the corresponding subgroup as the automorphisms of E which fix K, which is just this Galois group). But we know E is a splitting field over K (if E is the splitting field of a polynomial over F, then it's also the splitting field of the same polynomial over K). So then $|\mathrm{Gal}(E/K)| = [E : K]$. To prove the other direction, fix a subgroup $H \subset G$, and consider its fixed field $K = E^H$; we want to show that $[E : E^H] = |H|$.

First, by definition H is a subgroup of $\mathrm{Gal}(E/E^H)$, which means

$$|H| \leq |\mathrm{Gal}(E/E^H)| = [E : E^H].$$

So it suffices to show the other direction of this inequality, meaning that $[E : E^H] \leq |H|$.

Let $|H| = n$. By the Primitive Element, we know $E = E^H(\alpha)$ for some α. So it's enough to check that α is a root of some polynomial in $E^H[x]$ of degree n.

We now apply a version of the averaging trick we saw earlier. Set

$$P(x) = \prod_{g \in H} (x - g(\alpha)).$$

(For example, if $E = \mathbb{C}$ and H consisted of the identity and complex conjugation, this would give $(x-z)(x-\overline{z})$.)

This polynomial starts its life as a polynomial in $E[x]$, but it actually has coefficients in E^H — to see this, observe that the action of any $h \in H$ just permutes the factors g from P (since $hH = H$ for any $h \in H$). So P is a degree n polynomial in $E^H[x]$ with α as a root, which shows that $[E : E^H] \leq n$, as desired. □

Proof of Main Theorem. Once we have the equality of degrees, the remaining part is essentially just formal — suppose $H \to E^H \to H'$. By definition we know $H \subset H'$; but the lemma implies $|H| = |H'|$, so $H = H'$. Similarly, if $K \to H \to K'$, then $K' \supset K$ (since K' consists of all elements fixed by H, but H is defined as the set of elements which fix K). But $[E : K] = [E : K']$, so then $K = K'$. □

30.3 Properties of the Correspondence

Note that the correspondence *reverses* inclusion — the larger the group, the smaller its fixed field (each element of the group gives a condition on the elements of the field; if there are more conditions, fewer elements will satisfy them).

Consider the tower of extensions $E/K/F$. By definition, the extension E/K is always Galois (we have $[E : K] = |H| = |\mathrm{Gal}(E/K)|$). On the other hand, K/F may or may not be a Galois extension.

Proposition 30.7
The extension K/F is Galois if and only if K is invariant under all $g \in \mathrm{Gal}(E/F)$, which happens if and only if the corresponding $H \subset G$ is normal. In that case, $\mathrm{Gal}(K/F) = G/H$.

Proof. First we'll prove K/F is Galois if and only if K is invariant under all $g \in G = \mathrm{Gal}(E/F)$.

If K/F is Galois, then it's a splitting field. Every $g \in \mathrm{Gal}(E/F)$ has to permute the roots of any polynomial in $F[x]$; this means $G : K \to K$ (since K is generated by the roots of some such polynomial).

Meanwhile, if K is invariant under all $g \in G$, then we have a homomorphism $\mathrm{Gal}(E/F) \to \mathrm{Gal}(K/F)$ by restricting the automorphisms to K (since they are automorphisms of K as well). The kernel of this homomorphism is $\mathrm{Gal}(E/K)$. So by the homomorphism theorem, the image of this homomorphism has cardinality

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

We saw that we must have $|\mathrm{Gal}(K/F)| \leq [K : F]$, so then equality must hold, and K/F is Galois.

We've now shown that K/F is Galois if and only if K is invariant under G; so now it suffices to show that K is invariant if and only if H is normal. But it's clear that if H corresponds to $K = E^H$, then gHg^{-1} corresponds to $g(K)$. So K is invariant under all $g \in G$ if and only if $gHg^{-1} = H$ for all $g \in G$, meaning H is normal. $\square$

Student Question. *Why does gHg^{-1} correspond to $g(K)$?*

Answer. *This is because an element γ fixes x if and only if $g\gamma g^{-1}$ fixes $g(x)$.*

31 Applications of the Galois Correspondence

31.1 Review

Last class, we saw that if E/F is a Galois extension and $G = \mathrm{Gal}(E/F)$, then there is a correspondence between subgroups $H \subset G$ and their fixed fields $E^H \subset E$. We saw that in the tower of extensions $E/E^H/F$, the top extension E/E^H is always Galois, with Galois group H. Meanwhile, E^H/F is not always Galois; but it's Galois if and only if H is normal, and in that case $G/H = \mathrm{Gal}(E^H/F)$ (so in some sense, the left-hand side makes sense if and only if the right-hand side does):

Proposition 31.1
If $K = E^H$, then K/F is Galois if and only if K is invariant under all $g \in G$, which occurs if and only if H is normal.

Student Question. *What does it mean that K is invariant under all $g \in G$?*

Answer. *This means that for any $g \in G$, we have $x \in K$ if and only if $g(x) \in K$. In other words, $g(K) = K$. (So each g permutes the elements of K; this doesn't mean that g fixes each element of K.)*

Student Question. *Did we prove the second equivalence (that K is invariant if and only if H is normal)?*

Answer. *At the end of last class — it follows from the correspondence being natural, and therefore compatible with the action of G. More precisely, if H corresponds to K, then gHg^{-1} corresponds to $g(K)$ (the action by g on subfields corresponds to the action by g on subgroups via conjugation — this is unsurprising, since conjugation is the natural action by group elements on subgroups). From this, we see that $g(K) = K$ if and only if $gHg^{-1} = H$.*

Then ghg^{-1} fixes $g(x)$ if and only if h fixes x — checking this is easy, as $ghg^{-1}(g(x)) = gh(x)$.

31.2 Cyclotomic Extensions

The main theorem can be used to answer our question about ruler and compass constructions:

Proposition 31.2
If $p = 2^k + 1$ is a Fermat prime, then a regular p-gon can be constructed by a compass and straightedge.

Proof. Let ζ be a pth root of unity. Then it suffices to show that $\mathbb{Q}(\zeta)$ can be obtained by iterating quadratic extensions — if we let $E = \mathbb{Q}(\zeta)$, then it suffices to show there exists a tower of subfields

$$\mathbb{Q} = F_0 \subset F_1 \subset F_2 \subset \cdots \subset F_n = E,$$

such that $[F_i : F_{i-1}] = 2$ for all i. Quadratic extensions can always be obtained by extracting the square root of some element; so this would mean we can obtain $\mathbb{Q}(\zeta)$ by starting with $\mathbb{Q}$ and successively applying arithmetic operations and square roots.

This is fairly clear from the Galois correspondence. We saw earlier that

$$\mathrm{Gal}(E/\mathbb{Q}) = (\mathbb{Z}/p\mathbb{Z})^\times \cong \mathbb{Z}/(p-1)\mathbb{Z} = \mathbb{Z}/2^k\mathbb{Z}.$$

We can now write

$$G = G_0 \supset G_1 \supset G_2 \supset \cdots \supset \{0\},$$

where $G_1 = 2\mathbb{Z}/2^k\mathbb{Z}$, $G_2 = 4\mathbb{Z}/2^k\mathbb{Z}$, and so on. Then $G_i/G_{i+1} \cong \mathbb{Z}/2\mathbb{Z}$ for all i.

We can then take F_i to be the fixed field of G_i. We saw that the correspondence reverses inclusion, and we know how degrees correspond — we have $[E : F_i] = 2^i$ for each i, which implies that $[F_i : F_{i-1}] = 2$, as desired. □

Example 31.3
Describe the first step in this construction (to find F_1).

Solution. We want to write down a quadratic extension of $\mathbb{Q}$. We know F_1 is the fixed field of G_1, and G_1 consists of the even residues in the language of $\mathbb{Z}/2^k\mathbb{Z}$; converting back to the language of $(\mathbb{Z}/p\mathbb{Z})^\times$, then G_1 consists of the squares (or *quadratic residues*) in $\mathbb{Z}/p\mathbb{Z}$ — elements of the form $a = b^2$ for some $b \neq 0$.

Suppose $\zeta = \exp(2\pi i/p)$, and let

$$\alpha = \sum_{a \in \mathrm{QR}} \zeta^a$$

(summing over all $(p-1)/2$ quadratic residues mod p — for example, if $p = 5$, then $\alpha = \zeta + \zeta^4$). It's clear that α is fixed by G_1, since multiplying all a by a quadratic residue only permutes them.

We also want to find its Galois conjugate β. To do that, we apply an element of the Galois group *not* in G_1, which gives

$$\beta = \sum_{b \in \mathrm{NQR}} \zeta^b$$

(summing over all quadratic nonresidues mod p — for example, if $p = 5$, then $\alpha = \zeta^2 + \zeta^3$.) We now want to compute the quadratic equation that α satisfies. We know

$$\alpha + \beta = \zeta^1 + \zeta^2 + \cdots + \zeta^{p-1} = -1.$$

On the other hand, we can compute

$$\alpha\beta = \sum n_c \zeta^c,$$

where n_c is the number of ways to write $a + b = c$ where a is a quadratic residue, and b is a quadratic nonresidue.

This is a combinatorial problem, which we can solve — first, $n_0 = 0$, since -1 is a square (this means if a is a square, so is $-a$, so we can't have $a + b = 0$ where a is square and b isn't). On the other hand, we claim that $n_1, \ldots, n_{p-1}$ are all equal — for any c and c', we can write $c' = tc$ for some t (since $\mathbb{Z}/p\mathbb{Z}$ is a field). If t is a square, then we can get a bijection between (a, b) with sum c and sum c', by multiplying by t. Meanwhile, if t is not a square, then we can get a bijection by multiplying and swapping — given (a, b) with sum c, we can take (tb, ta) with sum c'. This means $n_c = n_{c'}$. Finally, we have $n_0 + \cdots + n_{p-1} = ((p-1)/2)^2$, since this is the number of ways to choose a summand from each of α and β. This means

$$n_c = \begin{cases} \frac{p-1}{4} & \text{if } c \neq 0 \\ 0 & \text{if } c = 0, \end{cases}$$

so then our sum is

$$\alpha\beta = \sum_{c=1}^{p-1} \frac{p-1}{4} \zeta^c = -\frac{p-1}{4}.$$

This means our quadratic equation is

$$\alpha^2 + \alpha - \frac{p-1}{4} = 0 \implies \alpha = \frac{-1 \pm \sqrt{p}}{2}.$$

So we have $F_1 = \mathbb{Q}(\sqrt{p})$. □

Note 31.4
This argument works for any prime $p \equiv 1 \pmod 4$, meaning that the quadratic extension of $\mathbb{Q}$ contained in $\mathbb{Q}(\zeta_p)$ is still $\mathbb{Q}(\sqrt{p})$. Meanwhile, when $p \equiv 3 \pmod 4$, we instead get $\mathbb{Q}(\sqrt{-p})$.

The description of $\mathrm{Gal}(\mathbb{Q}(\zeta)/\mathbb{Q})$ can be generalized to apply to all n (meaning ζ is an nth root of unity), not just primes.

Definition 31.5
The nth **cyclotomic polynomial** Φ_n is the monic polynomial in $\mathbb{Z}[x]$ whose roots are exactly the primitive nth roots of unity.

We then have

$$x^n - 1 = \prod_{d|n} \Phi_d(x).$$

This is because the roots of $x^n - 1$ are all elements whose order in $\mathbb{C}^\times$ divides n, and the right-hand side groups such terms by their order d).

This formula lets us compute Φ_n.

Example 31.6
We have $\Phi_1(x) = x - 1$, and

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

We can also compute other polynomials $\Phi_n(x)$, such as

$$\Phi_{12}(x) = x^4 - x^2 + 1.$$

The cyclotomic polynomials don't always have all coefficients 0 or ± 1, but the smallest counterexample is 105 (the smallest product of three distinct odd primes). But from this formula, it's easy to show by induction that all Φ_n have integer coefficients.

Fact 31.7
Φ_n is irreducible in $\mathbb{Q}[x]$.

We proved this fact for primes; we won't prove it for general n, since the proof is longer.

Also note that $\deg(\Phi_n)$ is the number of elements of order n in the additive group $\mathbb{Z}/n\mathbb{Z}$, which is $\varphi(n) = |(\mathbb{Z}/n\mathbb{Z})^\times|$. If $n = p_1^{d_1} \cdots p_k^{d_k}$, we have the explicit formula

$$\varphi(n) = \prod_i (p_i^{d_i} - p_i^{d_i - 1}).$$

Now we have

$$[\mathbb{Q}(\zeta) : \mathbb{Q}] = \varphi(n),$$

and $\mathbb{Q}(\zeta)$ is a splitting field (for the same reason as in the prime case — all roots of Φ_n are powers of ζ). By the same reasoning as the prime case, we then have

$$\mathrm{Gal}(\mathbb{Q}(\zeta)/\mathbb{Q}) = (\mathbb{Z}/n\mathbb{Z})^*.$$

Note that this is not necessarily cyclic — in fact, it's not cyclic unless n is a prime power or twice a prime power (and it's also not cyclic if $n \geq 8$ is a power of 2). It'll be the *product* of cyclic groups (since it's still abelian), but there will usually be multiple factors of even order in this product.

31.3 Kummer Extensions

We'll now consider extensions E/F where $E = F(\alpha)$ for some α such that $\alpha^n \in F$ for a positive integer n (and $\alpha \neq 0$). Assume that F contains all nth roots of unity, meaning that

$$\mu_n(F) = \{x \in F \mid x^n = 1\}$$

has exactly n elements (and therefore $\mu_n(F) \cong \mathbb{Z}/n\mathbb{Z}$); this is equivalent to requiring that F contains a primitive nth root of 1.

Our main example is over characteristic 0, but this can be done over characteristic p as well, with the additional requirement that $p \nmid n$.

Proposition 31.8
In this case E/F is Galois, and

$$\mathrm{Gal}(E/F) \cong \mathbb{Z}/m\mathbb{Z}$$

for some $m \mid n$. In fact, if $x^n - a$ is irreducible in $F[x]$, then $m = n$.

Proof. We have

$$x^n - a = \prod (x - \zeta^i \alpha),$$

where $0 \leq i \leq n-1$ and ζ is a primitive nth root of 1 (since all $\zeta^i \alpha$ are roots of $x^n - a$, and they are all distinct). So if we're given one root of $x^n - a$, then all possible roots are obtained by multiplication by roots of unity (which are in F). So E is the splitting field of $x^n - a$.

Now an element $\sigma \in G = \mathrm{Gal}(E/F)$ is uniquely determined by $\sigma(\alpha)$, which must be $\zeta^i \alpha$ for some i. For each i, let σ_i be the element in G such that $\sigma_i(\alpha) = \zeta^i \alpha$, if it exists (the element σ_i doesn't necessarily exist for all i).

It's clear that

$$\sigma_i \sigma_j(\alpha) = \sigma_i(\zeta^j \alpha) = \zeta^{i+j} \alpha = \sigma_{i+j}(\alpha)$$

(because $\zeta \in F$, so σ must fix it). So then $\sigma_i \sigma_j = \sigma_{i+j}$. This means G is isomorphic to a subgroup in $\mathbb{Z}/n\mathbb{Z}$, and every such subgroup must be of the form $\mathbb{Z}/m\mathbb{Z}$ where $m \mid n$.

In fact $m = \deg(E/F)$, so $m = n$ if and only if $x^n - a$ is irreducible. (When $x^n - a$ to be reducible, this fails in a trivial way — then a *smaller* power of α is in F.) □

31.4 Quintic Equations

Using these ideas, we can obtain the famous application of Galois theory to the impossibility of solving a general polynomial equation of degree at least 5.

Definition 31.9
A finite group G is **solvable** if there exists a sequence of subgroups

$$G = G_0 \supset G_1 \supset G_2 \supset \cdots \supset G_n = \{1\}$$

such that for all i, G_i is a normal subgroup of G_{i-1} and G_{i-1}/G_i is abelian.

The main idea of the proof is the following two propositions:

Proposition 31.10
Given an extension E/F and some $\alpha \in E$ such that α can be obtained from elements of F by arithmetic operations (addition, subtraction, multiplication, and division) and extracting arbitrary nth roots (where we're allowed to choose any of the possible nth roots), then α lies in a Galois extension of F with a solvable Galois group.

Proposition 31.11
S_n is not solvable for $n \geq 5$.

The first proposition essentially follows from what we've already discussed — we'll discuss it in more detail next class, but the idea is to first add the roots of unity; then when we extract a nth root, we get an extension with cyclic Galois group. Then when we extract nth roots repeatedly, we get a sequence of subgroups with abelian quotients. Meanwhile, the second is an elementary finite group argument.

Corollary 31.12
A root of a polynomial P of degree 5 with Galois group S_5 cannot be expressed through the rational numbers in radicals.

Saying the root can't be expressed in radicals is shorthand for the longer sentence from earlier — it simply means that it can't be obtained by arithmetic operations and extracting nth roots.

So this means not only is there no universal formula for the roots using radicals (as there is in lower degrees), there isn't even a way to write down the roots of a *specific* polynomial.

Proof of corollary. If it were possible to express all roots of P in radicals, then the splitting field K of P would be contained in a Galois extension of $\mathbb{Q}$ with solvable Galois group G. But then we have an onto homomorphism

$G \twoheadrightarrow \mathrm{Gal}(K/\mathbb{Q}) = S_5$. But the quotient of a solvable group is again solvable; so this would imply S_5 is solvable, contradiction. $\square$

32 Solving Polynomial Equations

One application of Galois theory is the impossibility of a solution in radicals to polynomial equations of degree at least 5. The fundamental theorem of algebra states that a degree n polynomial has n (not necessarily distinct roots). For linear polynomials $ax + b$, the root is obviously $x = -b/a$; for quadratics, the quadratic formula is well-known. Even for degree 3 and 4 polynomials, the cubic and quartic equations (which are much longer) provide universal formulae to find the roots. For a long time, mathematicians searched for the elusive "quintic formula," but now we know that there is no way to "write down" the roots of polynomials of degree 5 or higher, and Galois theory is the key to proving this fact.

32.1 Solvable Groups

Last class, we established the following definition:

Definition 32.1
A finite group G is **solvable** if there exists a sequence of subgroups $G = G_0 \supset G_1 \supset G_2 \supset \cdots \supset G_n = \{1\}$, such that for each i, G_i is a normal subgroup of G_{i-1}, and G_{i-1}/G_i is abelian.

Informally, a group is solvable if it can be built from putting together abelian groups.

Lemma 32.2
If $G/K \cong H$ (equivalently, if there is an onto map $G \twoheadrightarrow H$ with kernel K), then:

1. If K and H are solvable, then G is solvable.
2. If G is solvable, then H is solvable.

Proof sketch. The first direction is clear from the correspondence theorem — we can essentially put the filtrations for H and K together. If we have a filtration $H = H_0 \supset H_1 \supset \cdots \supset H_n = \{1\}$ with this property, we can take their pre-images $G = G_0 \supset G_1 \supset \cdots \supset G_n = K$. We then can place the filtration for K at the end.

For the second, we can take our filtration of G, and simply take its image. The intermediate subgroups we get for H will be quotients of the intermediate subgroups for G, and the quotient of an abelian group is also abelian; so this gives a valid filtration for H. □

We have the following group-theoretic lemma:

Lemma 32.3
S_5 is not solvable. In fact, A_5 is simple.

Recall that a group is *simple* if it has no normal subgroups except for itself and $\{1\}$.

This lemma is also true for $n \geq 5$ (meaning S_n is not solvable). But for $n < 5$ it's not true — we have $A_3 \cong \mathbb{Z}/3\mathbb{Z}$, which is abelian; while S_4 contains the Klein 4-group K_4 (consisting of (12)(34), (13)(24), (14)(23), and the identity), which is a normal subgroup.

Proof. The best proof is to think about the structure of conjugacy classes in symmetric groups; but we don't have time, so we'll do a more quick and dirty proof for just the case $n = 5$.

The class equation for A_5 is

$$60 = 1 + 15 + 20 + 12 + 12$$

(corresponding to the conjugacy classes of (12)(34), (123), (12345), and (13245)). Note also that if we have a 5-cycle in one of the conjugacy classes of size 12, its square is in the other.

If N is a normal subgroup, then it's a union of conjugacy classes, which means $|N|$ is a sum of 1, plus a subset of $\{15, 20, 24\}$. But it also has to divide 60. This is impossible — we have to take 1, but then we have to take 15 (or else the sum would be odd, and would need to divide 15). Then we have to take 20 (otherwise the sum would be 1 mod 3, and would have to divide 20). Then we must take 24 because otherwise the sum wouldn't be divisible by 5 (and would have to divide 12). □

Student Question. *How does this argument generalize to all $n \geq 5$?*

Answer. *This argument doesn't really generalize — but there is a slightly longer argument that does. We essentially just look at the possible cycle structures, take one conjugagy class, and show that the products of elements in that conjugacy class cover every other conjugacy class.*

32.2 Radical Extensions

Now we'll relate this to polynomial equations.

Definition 32.4
A finite extension E/F is a **radical extension** if $E = F(\alpha_1, \dots, \alpha_n)$, where $\alpha_i^{n_i} \in F(\alpha_1, \dots, \alpha^{i-1})$ for all i (for some positive integers n_i).

Informally, a radical extension is one that can be obtained by adjoining a bunch of radicals — in simple English, we're allowed to perform arithmetic operations and to extract radicals of any order.

Example 32.5
The extension

$$\mathbb{Q}\left(\sqrt[3]{3+\sqrt[5]{7+\sqrt{2}}}\right)$$

is a radical extension.

Proposition 32.6
Any radical extension is contained in a Galois extension with a solvable Galois group.

We'll assume that $\mathrm{char}(F) = 0$ (although things do generalize to $\mathrm{char}(F) = p$ with some more care).

The proof essentially hinges on the lemma discussed last class — that if F contains a primitive nth root of unity, and $E = F(\alpha)$ for some α with $\alpha^n \in F$, then E/F is a Galois extension whose Galois group is cyclic.

It'll be convenient to slightly generalize this lemma, to let us *simultaneously* extract a bunch of radicals.

Lemma 32.7
Under the same assumptions, if $E = F(\beta_1, \dots, \beta_k)$ where $\beta_i^n \in F$ for all i (and F contains a primitive nth root of unity), then $\mathrm{Gal}(E/F) \subset (\mathbb{Z}/n\mathbb{Z})^k$. In particular, $\mathrm{Gal}(E/F)$ is still abelian.

The proof is the same as before — any element of the Galois group sends $\beta_i \mapsto \beta_i \zeta_n^{c_i}$ for some exponent c_i, and composing elements of the Galois group corresponds to adding each pair of c_i.

Proof of Proposition 32.6. Use induction on n. If $n = 1$, then $E \subset F(\zeta, \alpha)$ where $\alpha^n \in F$ and ζ is a primitive nth root of unity (we can't assume in this proof that F contains roots of unity, but we can essentially just add them). This is the splitting field of $x^n - \alpha^n$.

So we have a tower of field extensions $F(\zeta, \alpha)/F(\zeta)/F$. We know that $F(\zeta, \alpha)/F(\zeta)$ has a Galois group which is a subgroup of $\mathbb{Z}/n\mathbb{Z}$; meanwhile, $F(\zeta)/F$ has a Galois group which is a subgroup of $(\mathbb{Z}/n\mathbb{Z})^\times$. Both are abelian, so the Galois group of $F(\zeta, \alpha)/F$ is solvable.

For the inductive step, assume that $F(\alpha_1, \dots, \alpha_{i-1}) \subset E'$, where $\mathrm{Gal}(E'/F)$ is solvable, and suppose $\alpha_i^{n_i} \in F(\alpha_1, \dots, \alpha_{i-1})$.

We then want to start with E', and first add a primitive n_ith root of unity ζ. To make sure we get a splitting field over F (and not just E'), we need to also add all Galois conjugates of α_i — let $\beta_1, \dots, \beta_d$ be all conjugates of $\alpha_i^{n_i}$ under $\mathrm{Gal}(E'/F)$. Then we take

$$E = E'(\zeta, \sqrt[n_i]{\beta_1}, \dots, \sqrt[n_i]{\beta_j}).$$

First, we want to show that E is a splitting field over F. Let Q be a polynomial such that E' is the splitting field of Q. Now if $\alpha_i^{n_i} = a$ (which lies in E'), we claim that E is the splitting field of

$$Q(x) \cdot (x^{n_i} - 1) \cdot \prod_{g \in \mathrm{Gal}(E'/F)} (x^{n_i} - g(a)).$$

(The reason we have this product over the Galois group, rather than simply the term $x^{n_i} - a$, is that a is not necessarily in F — but this product is, by the trick seen earlier.)

Now consider the tower of extensions $E/E'(\zeta)/E'/F$. Then $\mathrm{Gal}(E'/F)$ is solvable by the induction assumption; $\mathrm{Gal}(E'(\zeta)/E')$ is a subgroup of $(\mathbb{Z}/n_i\mathbb{Z})^\times$ and is therefore abelian; and $\mathrm{Gal}(E/E'(\zeta))$ is a subgroup of $(\mathbb{Z}/n\mathbb{Z})^d$, where $d = |\mathrm{Gal}(E'/F)|$. So $\mathrm{Gal}(E/F)$ is solvable as well, and we're done. □

Note 32.8
The proof contains a technicality in order to ensure that our extensions are all Galois extensions of F; this is why we needed to deal with the β_i. Other than that, it's essentially just a direct application of the lemma from earlier (about the Galois group when we just add a nth root).

The conclusion is now clear:

Corollary 32.9
There are many nonradical extensions of $\mathbb{Q}$.

For instance, the splitting field of any polynomial with Galois group S_5 (such as our example $2x^5 - 5x - 10$ from earlier) is a nonradical extension; this means the roots of such a polynomial can't have an expression in radicals.

32.3 Symmetric Polynomials

We'll now move on to a more concrete question:

Guiding Question
Given a polynomial, how do we compute $\mathrm{Gal}(E/F)$ and solve the equation when possible?

To answer this, we'll use the computational tool of symmetric polynomials (which can be understood independently of fields and Galois theory, and is an important branch of elementary algebra).

Consider the polynomial ring $\mathbb{Z}[x_1, \dots, x_n]$ (we could do the same with rational-coefficient polynomials). We use $R_n = \mathbb{Z}[x_1, \dots, x_n]^{S_n}$ to denote the subgroup of $\mathbb{Z}[x_1, \dots, x_n]$ consisting of polynomials which are invariant under permutation of the variables.

Example 32.10
If $n = 3$, then $x_1^3 + x_2^3 + x_3^3 \in R_3$, while $x_1^3 \notin R_3$.

It's easy to write down polynomials in R_n — we can start with any polynomial, and average over all permutations. The easiest example to think about is probably power sums; but a particularly useful one will be something different, the *elementary symmetric functions*.

Definition 32.11
If we have n variables $x_1, \dots, x_n$, then the **elementary symmetric functions** $\sigma_1, \dots, \sigma_n$ are defined as

$$\begin{aligned}\sigma_1 &= x_1 + x_2 + \cdots + x_n,\\ \sigma_2 &= x_1x_2 + x_1x_3 + \cdots + x_{n-1}x_n,\\ \sigma_3 &= x_1x_2x_3 + \cdots + x_{n-2}x_{n-1}x_n,\end{aligned}$$

and so on: in general, σ_k is the sum of the $\binom{n}{k}$ monomials which are products of k distinct terms x_j:

$$\sigma_k = \sum_{1 \le j_1 < j_2 < \cdots < j_k \le n} x_{j_1} \cdots x_{j_k}.$$

The first reason the elementary symmetric functions are relevant is that if we have a polynomial whose roots we know, we can expand it as

$$(z - x_1) \cdots (z - x_n) = z^n - \sigma_1 z^{n-1} + \sigma_2 z^{n-2} - \cdots + (-1)^n \sigma_n,$$

by considering which term in each factor we choose when expanding the product.

Example 32.12
If $n = 2$, we have

$$(z - x)(z - y) = z - (x + y)z + xy.$$

A useful fact about the elementary symmetric functions is the following:

Theorem 32.13
We have

$$R_n = \mathbb{Z}[\sigma_1, \sigma_2, \dots, \sigma_n].$$

Given two symmetric polynomials, it's obvious that their sum and product are also symmetric polynomials. But the interesting part of this theorem is that every symmetric polynomial can be expressed as a polynomial in the elementary symmetric functions (and this expression is unique).

Example 32.14
We have

$$\begin{aligned}x_1^2 + x_2^2 + x_3^2 &= \sigma_1^2 - 2\sigma_2;\\ x_1^3 + x_2^3 + x_3^3 &= \sigma_1^3 - 3\sigma_1\sigma_2 + 3\sigma_3.\end{aligned}$$

We'll prove the theorem later, but the reason it's useful here is the following:

Corollary 32.15
A symmetric polynomial in the roots of P can be written as a polynomial in the coefficients of P.

The strategy for how to compute the Galois group of a polynomial is based on this fact. In particular, one important symmetric polynomial in the roots is the *discriminant*:

Definition 32.16
The **discriminant** of $P(x) = \prod(x - \alpha_i)$ is

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

When we square, it doesn't matter which of α_i and α_j had smaller index; so the discriminant is a symmetric polynomial in the roots. By the theorem, the discriminant is then some polynomial in the coefficients.

Unfortunately, the formulas quickly get complicated. However, there are some cases where the discriminant is reasonable to compute:

Example 32.17
If $P(x) = x^3 + px + q$, then $D = -4p^3 - 27q^2$.

Proof. We could compute D using the definition, but that's fairly messy. Instead, we'll look at degrees — we know D is a degree 6 polynomial in the roots α_1, α_2, and α_3. Meanwhile, p is a degree 2 polynomial in the roots, and q is a degree 3 polynomial. The only monomials in p and q which can have degree 6 are then p^3 and q^2, so $D = ap^2 + bq^2$ for some a and b. We can then plug in a few polynomials for P to solve for a and b — for example, $P(x) = x(x-1)(x+1)$ has $p = -1$, $q = 0$, and $D = 4$, so $a = -4$; meanwhile $P(x) = (x-1)^2(x+2) = x^3 - 3x + 2$ has $p = -3$, $q = 2$, and $D = 0$, so $b = -27$. □

Although this only works for cubic polynomials whose x^2 coefficient is 0, there's an easy trick to turn any cubic polynomial into this form — if we start with $x^3 + ax^2 + bx + c$, we can substitute $y = x + a/3$.

Note that $D = 0$ if and only if P has multiple roots. The main application to Galois theory is that $\sqrt{D}$ is always in the splitting field of P; and in fact $\sqrt{D} \in F$ if and only if the Galois group is a subset of A_n. We'll discuss this in more detail next class.

33 Symmetric Polynomials and the Discriminant

33.1 Symmetric Polynomials

Last class, we began discussing symmetric polynomials and the discriminant. The goal is to develop some tools to understand the structure of the solutions to a polynomial — if not to compute them in radicals, then at least to see how this works when possible. Galois theory can be used for this as well, not just proving impossibility.

Last time, we considered the symmetric polynomials

$$\mathbb{Z}[x_1, \dots, x_n] \supset \mathbb{Z}[x_1, \dots, x_n]^{S_n} = R_n,$$

and stated the following fundamental theorem:

Theorem 33.1
We have $R_n = \mathbb{Z}[\sigma_1^{(n)}, \sigma_2^{(n)}, \dots, \sigma_n^{(n)}]$, where

$$\begin{aligned} \sigma_1^{(n)} &= x_1 + \cdots + x_n, \\ \sigma_2^{(n)} &= x_1x_2 + \cdots + x_{n-1}x_n, \\ &\vdots \\ \sigma_n^{(n)} &= x_1 \cdots x_n. \end{aligned}$$

In other words, every symmetric polynomial $P \in R_n$ can be written in terms of the elementary symmetric functions.

Example 33.2
We have

$$\begin{aligned} x_1^2 + \cdots + x_n^2 &= \sigma_1^2 - 2\sigma_2, \\ x_1^3 + \cdots + x_n^3 &= \sigma_1^3 - 3\sigma_1\sigma_2 + 3\sigma_3. \end{aligned}$$

We previously saw this example when $n = 3$. But a feature is that the right-hand side doesn't really depend on n — when we write the expression in terms of the symmetric polynomials, n sort of disappears.

The reason is that we have an obvious homomorphism $\mathbb{Z}[x_1, \dots, x_n] \to \mathbb{Z}[x_1, \dots, x_{n-1}]$ sending $x_n \to 0$ (so we essentially just kill one of the variables). Under this homomorphism, we have $R_n \to R_{n-1}$ (since if the polynomial was invariant under permutations of n variables, it's also invariant under permutations of the first $n-1$, where we killed the last variable). But this homomorphism is compatible with the elementary symmetric functions — it's clear that $\sigma_i^{(n)} \mapsto \sigma_i^{(n-1)}$, since the homomorphism essentially just kills the monomials containing x_n (and $\sigma_n^{(n)} \mapsto 0$).

Meanwhile, the power sums have the same property — we have $x_1^d + \cdots + x_n^d \mapsto x_1^d + \cdots + x_{n-1}^d$. This means if we can prove an identity for large n, we can automatically deduce it for smaller n as well.

On the other hand, we can also use degree considerations. For example, $x_1^3 + \cdots + x_n^3$ is a homogeneous polynomial of degree 3 in the x_i; while $\sigma_i^{(n)}$ is a homogeneous polynomial of degree i. This means we can only use $\sigma_1^{(n)}$, $\sigma_2^{(n)}$, and $\sigma_3^{(n)}$ in the expression; so if we can find the identity for $n = 3$, we can automatically deduce it for all larger n as well.

Putting these observations together, we see that a formula as in Example 33.2 for $n+1$ implies one for n, and conversely, using degree considerations it's enough to check it for small n. (In our example, $n = 3$ was enough; note that $n = 2$ is too small, because σ_3 is mapped to 0.)

We'll now prove the theorem.

Proof of Theorem 33.1. We use induction in the number of variables.

We need to check that every symmetric polynomial can be expressed as a polynomial in $\sigma_1^{(n)}, \dots, \sigma_n^{(n)}$, and that this expression is unique. In other words, we can consider the map $\varphi_n : \mathbb{Z}[t_1, \dots, t_n] \to R_n$, where $t_i \mapsto \sigma_i^{(n)}$.

Then we want to check that φ_n is an isomorphism, meaning that it's one-to-one and onto. We'll check these two parts separately.

First we'll check that φ_n is injective. Suppose there is some polynomial $Q(t_1, \dots, t_n)$ which φ_n maps to 0 (to prove injectivity, it suffices to show there is no such polynomial). We can first pull out the last variable, by writing $Q = t_n^d Q'$, where $t_n \nmid Q'$. Then $\varphi_n(t_n^d Q') = 0$, so we have

$$\left(\prod x_i\right)^d \cdot Q'(\sigma_1, \dots, \sigma_n) = 0.$$

Since the first factor is nonzero, this means $Q'(\sigma_1, \dots, \sigma_n) = 0$. So this essentially means that we can assume that Q is not divisible by t_n.

But now we can use the homomorphism from earlier — let r_n be the restriction map $R_n \to R_{n-1}$ sending $x_n \mapsto 0$. Then we have

$$r_n(Q'(\sigma_1, \dots, \sigma_n)) = 0$$

(because we assumed $Q'(\sigma_1, \dots, \sigma_n) = 0$). But we have

$$r_n\left(Q'\left(\sigma_1^{(n)}, \dots, \sigma_n^{(n)}\right)\right) = \overline{Q}\left(\sigma_1^{(n-1)}, \dots, \sigma_{n-1}^{(n-1)}\right),$$

where $\overline{Q}$ is not identically 0 (since we assumed Q' is not divisible by t_n, and r_n just maps $t_n \mapsto 0$). But this contradicts the inductive assumption (because then φ_{n-1} would map a nonzero polynomial $\overline{Q}_{n-1}$ to 0). (The base case of the induction is $n = 1$, which is trivial.)

Now we'll check that φ_n is surjective. Start with a polynomial $P \in R_n$; we can assume P is homogeneous of degree d. We want to check that $P = \varphi_n(Q)$ for some Q; we'll use induction on d.

The idea is again to reduce the number of variables. Let

$$r_n(P) = T(\sigma_1^{(n-1)}, \dots, \sigma_{n-1}^{(n-1)})$$

(by the inductive assumption). Now consider the polynomial

$$P - T\left(\sigma_1^{(n)}, \dots, \sigma_{n-1}^{(n)}\right).$$

We know that r_n maps this polynomial to 0. But the kernel of r_n is generated by x_n, so then

$$x_n \mid P - T\left(\sigma_1^{(n)}, \dots, \sigma_{n-1}^{(n)}\right).$$

Since the RHS, it then follows that each x_i divides it; but by unique factorization, this means

$$x_1 \cdots x_n \mid P - T\left(\sigma_1^{(n)}, \dots, \sigma_{n-1}^{(n)}\right).$$

So then we can write

$$P - T\left(\sigma_1^{(n)}, \dots, \sigma_{n-1}^{(n)}\right) = \sigma_n \cdot Q,$$

where Q is a symmetric polynomial of smaller degree. But $Q = S(\sigma_1, \dots, \sigma_n)$ by the inductive assumption, so

$$P = S(\sigma_1, \dots, \sigma_n) + T(\sigma_1, \dots, \sigma_n),$$

as desired. □

This is one possible proof, emphasizing the inductive structure; but in applications, the proof won't matter that much.

33.2 The Discriminant

By Theorem 33.1, we can write

$$\prod_{i<j}(x_i - x_j)^2 = \Delta_n(\sigma_1, \dots, \sigma_)$$

for some polynomial Δ_n (since the LHS is a symmetric polynomial).

Definition 33.3
The **discriminant** of a polynomial $P(z) = z^n + a_{n-1}z^{n-1} + \cdots + a_0$ is

$$D = \Delta_n(-a_{n-1}, a_{n-2}, \ldots, (-1)^n a_0).$$

In particular, we see that $D = 0$ if and only if P has a multiple root.

Example 33.4
We can calculate the discriminant explicitly when P has low degree — for $P(x) = x^2 + bx + c$ we have $D = b^2 - 4c$, while for $P(x) = x^3 + px + q$ we have $D = -4p^3 - 27q^2$.

Proof. We proved this last time, but we'll outline the proof of the second statement again. Degree considerations give that $D = ap^3 + bq^2$ for some a and b. Then we can take $P(x) = x^3 - x$ (which has discriminant 4) to get that $a = -4$, and $P(x) = (x-1)^2(x+2)$ (which has discriminant 0) to get that $b = -27$. □

Student Question. *How did we show that $D = ap^3 + bq^2$?*

Answer. *To simplify the formulas, we assumed that $\sigma_1 = 0$, so we're trying to compute $\Delta_3(0, p, -q)$. But Δ_3 is homogeneous of degree 6, as a polynomial in the x_i; meanwhile σ_i is homogeneous of degree i. We only have σ_2 (of degree 2) and σ_3 (of degree 3), and the only way to make 6 from 2's and 3's is $2+2+2$ and $3+3$, so the only possible terms we can have are σ_2^3 and σ_3^2.*

Student Question. *Does this argument only work when $\sigma_1 = 0$?*

Answer. *Yes — in the general case, there is still a formula for Δ_3, but it's longer. But the case $\sigma_1 = 0$ is actually enough for practical purposes, since it's possible to reduce any cubic to this form (by shifting the variable).*

We'll now get to the role of the discriminant in Galois theory. We can also consider

$$\delta_n(x_1, \ldots, x_n) = \prod_{i<j}(x_j - x_i),$$

so $\Delta_n = \delta_n^2$. Note that δ_n is not symmetric — if we swap x_i and x_{i+1}, then this swaps the sign of δ_n. This means

$$\delta\left(x_{\sigma(1)}, \ldots, x_{\sigma(n)}\right) = (-1)^{\operatorname{sgn}(\sigma)}\delta(x_1, \ldots, x_n),$$

so they have the same sign if σ is even and opposite sign if σ is odd. (Here σ denotes an arbitrary permutation of $\{1, \ldots, n\}$.)

Now let E/F be a field extension, where E is the splitting field of $P \in F[x]$. Assume that P does not have multiple roots (but it is allowed to be reducible).

By definition, this means $E = F(\alpha_1, \ldots, \alpha_n)$, where $P(x) = \prod(x - \alpha_i)$. We know that $G = \operatorname{Gal}(E/F)$ is a subgroup of S_n, since elements of G must permute the roots of P.

Now let E/F be a field extension, where E is a splitting field of a polynomial $P \in F[x]$. Let $P(x) = \prod(x - \alpha_i)$, and assume that P doesn't have multiple roots (we don't need to assume it's irreducible).

By definition, this means $E = F(\alpha_1, \ldots, \alpha_n)$. We know that $G = \operatorname{Gal}(E/F)$ must permute the roots α_i, and is therefore a subgroup of S_n (where we look at how it permutes those roots). Now if we let $\delta = \prod_{i<j}(\alpha_j - \alpha_i)$, we know that $\delta \in E$ and $\delta^2 \in F$. We can immediately see how G acts on δ — for any $\sigma \in G$, we know

$$\sigma(\delta) = \begin{cases} \delta & \text{if } \sigma \text{ is even} \\ -\delta & \text{if } \sigma \text{ is odd.} \end{cases}$$

In particular, this means $\delta \in F$ if and only if all permutations $\sigma \in G$ are even. (This is because F is exactly the fixed field of G, by the main theorem; so δ is fixed by all elements of G if and only if it's in F.) So the conclusion is the following:

Proposition 33.5
We have $\mathrm{Gal}(P) \subset A_n$ if and only if the discriminant Δ of P is a square.

33.3 Cubic Polynomials

We can now apply this to a concrete situation: suppose that $n = 3$, and we know P is irreducible. There are only two transitive subgroups of S_3, which are A_3 and S_3. So these are the only options for $\mathrm{Gal}(P)$, and we have a concrete way of distinguishing between these two cases — the Galois group is S_3 when Δ is not a square, and A_3 when Δ is a square.

Example 33.6
Find the Galois group of $P(x) = x^3 - 3x - 1$ over $\mathbb{Q}$.

Solution. The discriminant of P is

$$D = 4 \cdot 27 - 27 = 81,$$

which is square; so the Galois group is $A_3 \cong \mathbb{Z}/3\mathbb{Z}$. □

Example 33.7
Suppose F contains a cube root of unity ω; find the Galois group of $P(x) = x^3 - a$ (assuming P is irreducible).

Solution. The discriminant is $D = -27a^2$. But since $\omega \in F$, then -27 is a square (since $\omega = \frac{1 \pm \sqrt{-3}}{2}$, we have that $\sqrt{-3} \in F$). So then $\mathrm{Gal}(P) \cong \mathbb{Z}/3\mathbb{Z}$. (This is a special case of the theorem we saw last time.) □

We've now seen an effective way to find the Galois group for cubic polynomials; we'll finish by discussing how to actually solve them.

Proposition 33.8
If F contains a primitive cube root of unity ω, and E/F is a Galois extension with $\mathrm{Gal}(E/F) = \mathbb{Z}/3$, then $E = F(\alpha)$ for some $\alpha^3 = a \in F$.

The proof we'll give is constructive, and if we explicitly write out the construction, this leads to Cardano's Formula for the solutions to a cubic (which was actually published in 1545).

Proof. Let σ be a generator for $\mathrm{Gal}(E/F)$. It suffices to find $\alpha \in E$ such that $\sigma(\alpha) = \omega\alpha$ or $\omega^2\alpha$ — then we have $\sigma(\alpha^3) = \alpha^3$, and since σ generates the Galois group, this means *all* elements of the Galois group fix α^3, so $\alpha^3 \in F$. Meanwhile, the Galois group does not fix α itself; so the degree of α is 3. Since the degree of the *extension* is 3 as well, this means $E = F(\alpha)$.

Pick some $\beta \in E$ which is not in F, and let

$$\alpha_1 = \beta + \omega\sigma(\beta) + \omega^2\sigma^2(\beta),$$
$$\alpha_2 = \beta + \omega^2\sigma(\beta) + \omega\sigma^2(\beta).$$

Then it's clear that $\sigma(\alpha_1) = \omega^2\alpha_1$ and $\sigma(\alpha_2) = \omega\alpha_2$, so it suffices to check that one of α_1 and α_2 is nonzero. But otherwise, $(\beta, \sigma(\beta), \sigma^2(\beta))$ would be a solution to the system of linear equations

$$a + \omega b + \omega^2 c = a + \omega^2 b + \omega c = 0.$$

Then orthogonality of characters for $\mathbb{Z}/3\mathbb{Z}$ (from the representation theory of cyclic groups) shows that the only solution is $a = b = c$. But then $\sigma(\beta) = \beta$, which contradicts the fact that $\beta \notin F$ (since if σ fixed β, then the entire Galois group would fix β). So one of α_1 and α_2 must be nonzero. □

Note 33.9

The same proof works with 3 replaced with any prime; and with a bit more work, it can be generalized to any n.

This can be used to prove the converse of the theorem from last class — last class, we saw that any radical extension is solvable. But this shows that any extension with a solvable Galois group is radical.

34 Solving Polynomial Equations (continued)

34.1 Cubic Polynomials

Last class, we looked at cubic polynomials of the form $P(x) = x^3 + px + q$ (called *depressed cubics*), which have discriminant $D = -4p^3 - 27p^2$. For simplicity we assume the field has characteristic 0, as the cases of characteristic 2 and 3 are somewhat different. We saw that if E is the splitting field of P, then $E \supset F(\delta)$ where $\delta = \sqrt{D}$; and $\mathrm{Gal}(E/F(\delta)) = \mathbb{Z}/3\mathbb{Z}$. (It's possible that $\delta \in F$, though it usually isn't.) We saw that then $E = F(\delta)(\alpha)$, where α is a cube root of some element $a \in E$; our explicit construction was $\alpha = \beta + \omega\sigma(\beta) + \omega^2\sigma^2(\beta)$ for some $\beta \in E$ (which isn't in F).

It's actually possible to turn these ideas into a formula for the roots of P. Let β_1, β_2, β_3 be the roots of P (which are elements in E). Then some $\sigma \in \mathrm{Gal}(E/F[\delta])$ must permute the roots in a 3-cycle $\beta_1 \to \beta_2 \to \beta_3 \to \beta_1$; this means we can take

$$\alpha = \beta_1 + \omega\beta_2 + \omega^2\beta_3.$$

We know $\alpha^3 \in F(\delta)$. In fact, using symmetric polynomials, we can express α^3 in terms of p, q, and δ. It's possible to show this by a general argument — this is because

$$\alpha^3 = Q(\beta_1, \beta_2, \beta_3)$$

for a polynomial Q which isn't quite symmetric, but is invariant under *even* permutations. This is enough, as a result of a slight generalization of the theorem on the elementary symmetric polynomials seen earlier:

Fact 34.1
We have

$$\mathbb{Q}[x_1, \dots, x_n]^{A_n} = \mathbb{Q}[x_1, \dots, x_n]^{S_n} \oplus \delta\mathbb{Q}[x_1, \dots, x_n]^{S_n}.$$

Intuitively, δ is invariant under A_n but changes sign under S_n; but this essentially accounts for all the new polynomials allowed when we only consider even permutations.

Instead of using this theoretical argument, it's also possible to just write down the expression for α^3 directly — we have

$$\alpha^3 = \beta_1^3 + \beta_2^3 + \beta_3^3 + 6\beta_1\beta_2\beta_3 + \omega(\beta_1^2\beta_2 + \beta_2^2\beta_3 + \beta_3^2\beta_1) + \omega^2(\beta_1\beta_2^2 + \beta_2\beta_3^2 + \beta_3\beta_1^2).$$

The first few terms are symmetric — we have the formulas

$$\beta_1\beta_2\beta_3 = -q$$
$$\beta_1^3 + \beta_2^3 + \beta_3^3 = -3q.$$

Meanwhile, we can let $A = \beta_1^2\beta_2 + \beta_2^2\beta_3 + \beta_3^2\beta_1$ and $B = \beta_1\beta_2^2 + \beta_2\beta_3^2 + \beta_3\beta_1^2$. We can then calculate that

$$A + B = \sigma_1\sigma_2 - 3\sigma_3 = 3q.$$

Meanwhile, $A - B$ is not symmetric, but by expanding we can see that

$$A - B = (\beta_1 - \beta_2)(\beta_1 - \beta_3)(\beta_2 - \beta_3) = \delta.$$

Now we're basically done — we can solve for A and B, and get a formula for α — we have

$$\alpha = \sqrt[3]{-4q + 3\omega A + 3\omega^2 B}.$$

Then we can similarly define and compute $\alpha' = \beta_1 + \omega^2\beta_2 + \omega\beta_3$. We then have $\beta_1 = (\alpha + \alpha')/3$, and we can calculate the other roots similarly. We won't describe the full formula here, but it's in the textbook. In fact, a version of this formula was discovered by Cardano in 1545.

Student Question. *If this formula was discovered before Galois theory, how did people come up with it?*

Answer. *The approach described here, of writing down formulas for these expressions, was invented by Legendre. It doesn't really need Galois theory in its full strength — it's possible to just notice that if we write $\alpha = \beta_1 + \omega\beta_2 + \omega^2\beta_3$, then α^3 is an expression in the roots which can be calculated using symmetric polynomials (and the discriminant).*

But not only was the formula discovered before Galois theory, it was also discovered before complex numbers. Having to work with roots of negative numbers gave people a lot of trouble — this was controversial even in the early 19th century. It mattered to people whether they could operate with real numbers, or had to work with strange expressions involving roots of negative numbers.

In fact, suppose that P has 3 real roots. Then Δ is nonnegative, so δ is real. But the expression $\alpha = \beta_1 + \omega\beta_2 + \omega^2\beta_3$ is not *real! So when you write down the answer in radicals, the final answer will be real; but you'll still need to work with a complex cubic root. This was referred to as* casus irreducibilis.

If you're interested in the history and philosophy of this story, a book by Barry Mazur called Imagining Numbers, Especially $\sqrt{-15}$ *talks about this history and reflects about how the understanding of such topics developed. Essentially, people were working with complex numbers three centuries before they were fully realized and accepted as existing.*

34.2 Quartic Polynomials

We'll also briefly discuss quartics. The key point is that the analysis of solutions can be guided by the structure of the Galois group.

The Galois group is a subgroup of S_4. We know S_4 contains the normal subgroup K_4 (the Klein 4-group $\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$), consisting of $\{(12)(34), (13)(24), (14)(23), 1\}$. We then have $S_4/K_4 \cong S_3$, corresponding to the *resolvent cubic*.

So if $P(x) = (x-\alpha_1)(x-\alpha_2)(x-\alpha_3)(x-\alpha_4)$, we can write down the expressions

$$\begin{aligned}\beta_1 &= \alpha_1\alpha_2 + \alpha_3\alpha_4,\\ \beta_2 &= \alpha_1\alpha_3 + \alpha_2\alpha_4,\\ \beta_3 &= \alpha_1\alpha_4 + \alpha_2\alpha_3.\end{aligned}$$

These expressions are permuted by the Galois group, and we know that if we take

$$Q(x) = (x-\beta_1)(x-\beta_2)(x-\beta_3),$$

then when we expand, the coefficients will be symmetric polynomials in the α_i. So if $P(x) = x^4 + a_3x^3 + a_2x^2 + a_1x + a_0$, then we have $Q(x) = x^3 + b_2x^2 + b_1x + b_0$ where the b_i are polynomials in the a_i — for concreteness, the exact formulas are $b_2 = -a_2$, $b_1 = a_1a_3 - 4a_0$, and $b_0 = 4a_0a_2 - a_1^2 - a_0a_3^2$.

Now to find a root, we first find the roots of the resolvent cubic $Q(x)$ (since we already know how to solve a cubic polynomial). Then, since $K_4 = \mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$, it just remains to solve a few quadratic equations. More explicitly, we can write the equations

$$\begin{aligned}(\alpha_1+\alpha_2)(\alpha_3+\alpha_4) &= \beta_1+\beta_3\\ \alpha_1+\alpha_2+\alpha_3+\alpha_4 &= -a_3.\end{aligned}$$

This gives a quadratic for $\alpha_1+\alpha_2$ and $\alpha_3+\alpha_4$, which we know how to solve. We can similarly find the other pairwise sums, and then compute the roots themselves by solving the resulting linear system.

This shows how to find the roots explicitly, but similarly to the cubic case, we can also try to compute the Galois group:

Guiding Question
How do we compute $\mathrm{Gal}(E/F)$ for a given polynomial $P(x) = x^4 + a_3x^3 + a_2x^2 + a_1x + a_0$?

In degree 3, all we needed to know was whether the discriminant was a square or not — this determined whether the group was $\mathbb{Z}/3\mathbb{Z}$ or S_3. In this case, the process is longer, but somewhat similar.

First, there are five transitive subgroups of S_4 — these are S_4, A_4, K_4 (the Klein 4-group, described earlier), C_4 (the cyclic group, generated by (1234)), and D_4 (the dihedral group, generated by (1234) and (24), which we can think of as the group of symmetries of a square).

One test we can perform still uses the discriminant. It's a lengthy expression, so we won't explicitly write it down, but it's still theoretically possible to compute it. Then $\sqrt{D} \in F$ if and only if $G \subset A_4$. The groups which are subsets of A_4 are K_4 and A_4 itself.

Then we can obtain more information from looking at the resolvent cubic (since we've already seen how to analyze cubics). We know that $Q(x)$ splits completely in F if and only if $G = K_4$ (since then all elements of G fix $\alpha_1\alpha_2 + \alpha_3\alpha_4$ and the other two expressions, which means they must lie in K_4).

Meanwhile, if $Q(x)$ has exactly one root in F, then the elements of G preserve one root of Q, say $\alpha_1\alpha_3 + \alpha_2\alpha_4$. In this case, we claim that the Galois group is C_4 or D_4 — we can visualize this by considering a square.

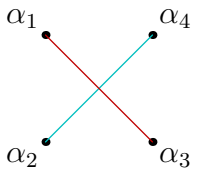

The square naturally splits into two subsets, by drawing its diagonals. So any permutation which fixes the square will either fix or swap $\alpha_1\alpha_3$ and $\alpha_2\alpha_4$, which means it fixes $\alpha_1\alpha_3 + \alpha_2\alpha_4$ (while not every permutation in C_4 fixes the other two expressions).

So this information resolves nearly all cases — the only ambiguity left is whether the group is C_4 or D_4. We won't explain how to distinguish between them, but an explanation is in Keith Conrad's notes.

34.3 Main Theorem of Algebra

We'll finish with another application of Galois theory — we'll use it to give another proof of the Main Theorem of Algebra. We'll see that this proof brings in some nice considerations about finite groups, although it's somewhat less direct than the proof we've seen before.

Proposition 34.2
Every p-group is solvable — if G is a finite group with $|G| = p^n$ for a prime p, then G is solvable. Moreover, there exists a chain of subgroups

$$G = G_0 \supset G_1 \supset \cdots \supset G_n = \{1\},$$

such that for all i, G_{i+1} is a normal subgroup of G_i and $G_i/G_{i+1} \cong \mathbb{Z}/p$.

Proof. We'll essentially start from the right end (instead of the left). We'll need the following lemma:

Lemma 34.3
G has a nontrivial center.

Proof. Consider the class equation mod p. Every conjugacy class has size p^m for some m. But the class equation states that

$$p^n = 1 + \sum |C_i|$$

(where the 1 comes from the conjugacy class of the identity, which has 1 element). If all other conjugacy classes contained more than one element, then $|C_i|$ would be divisible by p for all i, and the right-hand side would be 1 mod p, contradiction. So there must exist conjugacy classes of size 1 (other than the one of the identity), and their elements are in the center of G. □

Now to prove the proposition, we induct on n. By the lemma, we have $G \supset Z$ (where Z is the center, and $Z \neq \{1\}$). Then we can find an element $g \in Z$ of order p (the center is a nontrivial p-group, so if we pick any element, it will have some power which has order p). Then let $\overline{G} = G/\langle g \rangle$ (which is valid because g is in the center of G, so $\langle g \rangle$ is clearly normal).

We have $|\overline{G}| = p^{n-1}$, so by the inductive assumption, $\overline{G}$ is solvable. Suppose we have a chain of subgroups

$$\overline{G} = \overline{G}_0 \supset \cdots \supset \overline{G}_d = \{1\}.$$

Now let G_i be the pre-image of $\overline{G}_i$, and let $G_{d+1} = \{1\}$; this works by the homomorphism theorem. □

Now we can prove the Main Theorem of Algebra:

Theorem 34.4
$\mathbb{C}$ is the only finite extension of $\mathbb{R}$.

This implies the standard formulation, that every polynomial (over $\mathbb{C}$) has a root in $\mathbb{C}$.

Proof. Let $F = \mathbb{R}$, and suppose E is a finite extension. Without loss of generality assume E is a splitting field (since it's a finite extension, it's obtained by adding some of the roots of some polynomial, and we can add in all the remaining roots), so E/F is a Galois extension. Let $G = \mathrm{Gal}(E/F)$.

Lemma 34.5
$|G|$ is a power of 2.

Proof. Let $H \subset G$ be a Sylow 2-subgroup (a subgroup of order 2^n, where n is the exponent of 2 in $|G|$). Then consider the extension E^H/F — we have that

$$[E^H : F] = \frac{[E : F]}{[E^H : E]} = \frac{|G|}{|H|},$$

which is odd. But any odd-degree polynomial in $\mathbb{R}[x]$ has a real root (by the intermediate value theorem — the polynomial goes to $+\infty$ on one end and $-\infty$ on the other). So this means there are no odd-degree extensions of $\mathbb{R}$; so $H = G$, which means $|G| = 2^n$. (This argument works even if G is odd, as then H is trivial.) □

But now we can use the proposition — we have

$$G = G_0 \supset G_1 \supset \cdots \supset G_k = \{1\}$$

where $G_i/G_{i+1} \cong \mathbb{Z}/2\mathbb{Z}$ for all i, and we can consider their fixed fields

$$\mathbb{R} = F_0 \supset F_1 \supset \cdots \supset F_k = E,$$

where F_i is the fixed field of G_i. Then we have $[F_{i+1} : F_i] = 2$ for all i.

But it's clear that $\mathbb{C}$ is the only *quadratic* extension of $\mathbb{R}$, and $\mathbb{C}$ itself has no quadratic extensions (we can check that every quadratic over $\mathbb{C}$ has a root, since we can extract square roots using the trigonometric form of a complex number). So then G is $\{1\}$ or $\mathbb{Z}/2\mathbb{Z}$, and E is $\mathbb{R}$ or $\mathbb{C}$. □

Next class, we'll discuss the Galois group of extensions of finite fields. We'll see that

$$\mathrm{Gal}(\mathbb{F}_{q^n}/\mathbb{F}_q) = \mathbb{Z}/n\mathbb{Z},$$

which essentially follows from the fact that $\mathbb{F}_q = \{x \mid x^q = x\}$ (we previously thought of these x as roots of a polynomial, but we can now think of them as fixed points under the map $t \mapsto t^q$).

Student Question. *How did we get that $|G| = |H|$ (when showing $|G|$ was a power of 2)?*

Answer. *By the Primitive Element Theorem (assuming $E^H \neq F$), the extension E^H/F is generated by one element. We can consider the minimal polynomial of this element; the degree of the minimal polynomial is equal to the degree of the extension. So the minimal polynomial has odd degree, which is a contradiction (since if it has a root, it's reducible).*

It's actually possible to avoid using the Primitive Element Theorem — if we take any *element in E^H, the degree of its minimal polynomial has to divide the degree of the extension (by the fact that $[K : F] = [K : E][E : F]$ for a tower of extensions $K/E/F$), and therefore has to be odd.*

35 Final Remarks

35.1 Galois Theory in Finite Fields

We've seen that when F is a number field (a finite extension of $\mathbb{Q}$), the Galois group $\mathrm{Gal}(E/F)$ can be complicated. But $\mathbb{Q}$ is only one of the primary fields — we can also consider finite extensions of $\mathbb{F}_p$ for p prime (which are the other primary fields). Then our base field is $F = \mathbb{F}_q$ where $q = p^m$ for some m, and a finite extension of F is $E = \mathbb{F}_{q^n}$ for some n.

In this case, the answer is much simpler, and we've essentially seen it already:

Theorem 35.1
The extension $\mathbb{F}_{q^n}/\mathbb{F}_q$ is always a Galois extension; and $\mathrm{Gal}(\mathbb{F}_{q^n}/\mathbb{F}_q)$ is cyclic and generated by the **Frobenius automorphism** $\mathrm{Fr}_q : x \mapsto x^q$.

Proof. We've seen earlier that

$$(a+b)^q = a^q + b^q,$$

so Fr_q is compatible with the field operations; and it's also one-to-one. So $\mathrm{Fr}_q : \mathbb{F}_{q^n} \to \mathbb{F}_{q^n}$ is a field automorphism. Its fixed points are exactly the set $\{x \mid x^q = x\} = \mathbb{F}_q$. So then we know $\mathrm{Fr}_q \in \mathrm{Gal}(\mathbb{F}_{q^n}/\mathbb{F}_q)$.

But we can also compute its order — we know $\mathrm{Fr}_q^a : x \mapsto x^{q^a}$. For $a = n$, we have that $\mathrm{Fr}_{q^n} = \mathrm{Id}$ (since $x^{q^n} = x$ for all $x \in \mathbb{F}_{q^n}$). Meanwhile, if $1 \le a < n$, not all $x \in \mathbb{F}_{q^n}$ satisfy $x^{q^a} = x$. So then $\mathrm{ord}(\mathrm{Fr}_q) = n$. This means

$$\mathrm{Gal}(\mathbb{F}_{q^n}/\mathbb{F}_q) \supset \mathbb{Z}/n\mathbb{Z}$$

(considering the cyclic group generated by Fr_q). But we have $[\mathbb{F}_{q^n} : \mathbb{F}_q] = n$, so the Galois group cannot have more than n elements; so we must have $\mathrm{Gal}(\mathbb{F}_{q^n}/\mathbb{F}_q) = \mathbb{Z}/n\mathbb{Z}$. □

Note 35.2
Properties of the Frobenius automorphism are useful when doing algebraic geometry in fields of positive characteristic. In order to count the number of solutions to a system of polynomial equations over $\mathbb{F}_q$, one first looks at solutions over its algebraic closure $F = \bigcup \mathbb{F}_{q^n}$ (a bigger field in which every polynomial has a root, similarly to $\mathbb{C}$). Then solutions in $(\mathbb{F}_q)^n$ are the fixed points of $\mathrm{Fr}_q : (x_1, \dots, x_n) \mapsto (x_1^q, \dots, x_n^q)$. One uses intuition from a similar problem in topology, of counting the number of fixed points of an automorphism of some geometric shape X. (This relates to the Lipschitz Fixed Points Theorem and the Weil conjectures.)

35.2 Further Directions

Finally, we'll go over the topics that have been covered in this class, and where they can lead.

35.2.1 Representation Theory

The first topic we discussed is the representations of finite groups. The class **18.715** develops this topic.

We've actually already seen the main general structural theorems about the representations of an *abstract* finite group; but one further direction is the classification and computation of the characters of irreducible representations for a *specific* group — in particular, S_n. We know the number of irreducible representations equals the number of conjugacy classes. But in this case, it's actually possible to index both by the same set — the set of partitions of n (ways to write $n = n_1 + \cdots + n_k$, where order doesn't matter). It turns out that irreducible representations are in bijection with partitions. Meanwhile, partitions are also in bijection with the cycle type of a permutation (which determines the conjugacy class) — in order to describe a conjugacy class, we're interested in the lengths of the cycles.

Example 35.3
The conjugacy class of $(12)(345)$ can be described by the partition $5 = 3 + 2$.

Partitions are usually depicted by Young diagrams, where the length of rows correspond to the summands. (These are also studied in classes on combinatorics.)

Example 35.4
The Young diagram

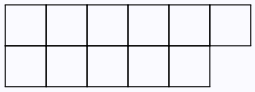

corresponds to the partition $11 = 6 + 5$.

As an example of how this correspondence can be used, recall a problem we saw earlier:

Example 35.5
For which n is $\tau \otimes \text{sgn} = \tau$? (Here τ is the tautological representation of S_n, where elements of S_n act on the space with $x_1 + \cdots + x_n = 0$ by permuting coordinates.)

Earlier, we solved this directly by looking at the characters, but there's a nice way to solve it by looking at Young diagrams as well.

In this correspondence, the Young diagram

(corresponding to $n = n$) corresponds to the trivial representation, and the Young diagram

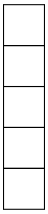

(corresponding to $n = 1 + 1 + \cdots + 1$) corresponds to sgn. Meanwhile, the Young diagram

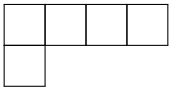

(corresponding to $n = (n-1) + 1$) corresponds to τ. More generally, $\rho \otimes \text{sgn}$ corresponds to the transpose of the diagram for ρ. So the fact that $\tau \otimes \text{sgn} = \tau$ exactly when $n = 3$ corresponds to the observation that the transpose of

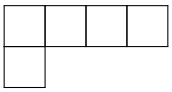

(where we reflect it over the diagonal) is itself if and only if $n = 3$.

35.2.2 Compact Lie Groups

Another direction in which representation theory can lead is compact Lie groups:

Definition 35.6
A compact Lie group is a closed compact subgroup in $\text{GL}(n, \mathbb{C})$.

Example 35.7
One compact Lie group (which we've mentioned earlier) is SU(2). Some other compact Lie groups include $\text{U}(n)$, $\text{SU}(n)$, $\text{SO}(n)$, and the *quaternionic unitary groups* $\text{Sp}(n)$ (which we haven't seen before).

It turns out that there's a classification of all such groups, along with some exceptional ones — G_2, F_4, E_6, E_7, and E_8. The largest exceptional group, E_8, has dimension 248. This can be studied further in **18.745** and **18.755**.

For each such group, there's also a classification of its irreducible representations.

Fact 35.8
The irreducible representations of $\mathrm{U}(n)$ are indexed by sequences of n integers $d_1 \geq \cdots \geq d_n$.

Example 35.9
Irreducible representations of $\mathrm{SU}(2)$ are indexed by one nonnegative integer n. To n, we assign the action on the space V_n of homogeneous polynomials of degree n in two variables (so V_n has a basis consisting of x^n, $x^{n-1}y$, $\ldots$, y^n, and has dimension $n+1$).

This connects to a more familiar situation — recall that $\mathrm{SU}(2)/\{\pm 1\} = \mathrm{SO}(3)$ (which is the group of rotations in 3-space). Then V_n for even n comes from a representation of $\mathrm{SO}(3)$.

This has an application to the spectrum of a hydrogen atom and the structure of the periodic table. The rows of the table have lengths 2, 8, 8, 18, 18, 32, 32 — these are $2n^2$ for small n, and n^2 arises as $1+3+\cdots+(n-1)$, where the odd numbers come from the dimensions of irreducible representations of $\mathrm{SO}(3)$.

The connection comes from an optional problem set problem on greedy monsters — we had monsters at the vertices of a cube, each with some amount of gold; and at every minute, the gold of each monster is equally distributed among its neighbors. The question asked us to understand how this process behaves over a long time. This is a special case of the Laplace operator on a graph:

Definition 35.10
Given a graph with certain weights assigned to the vertices, the **Laplace operator** redistributes the weight of every vertex equally among its neighbors.

This is the discrete version of the Laplace operator, but there's also a continuous version — for functions on $\mathbb{R}^2$, take the differential operator

$$\Delta = \frac{d^2}{dx^2} + \frac{d^2}{dy^2}.$$

This measures the extent to which the function is not harmonic — it vanishes exactly on functions whose value at each point is the average of the values on a small circle around it. (This is a continuous analog of the greedy monsters case, where our operator vanishes when the weight of each point equals the average weight of its neighbors.)

For symmetric graphs (such as the cube, in the case of the greedy monsters), we can understand the eigenvalues and eigenvectors of the operator using representation theory (as we did in the optional problem). Meanwhile, for the Laplace operator of the sphere S^2, its eigenvalues can be analyzed using representation theory of $\mathrm{SO}(3)$; these eigenvalues then have connections to quantum physics.

Another important identity we saw was that

$$|G| = \sum d_i^2,$$

where the d_i are the dimensions of irreducible representations. This fact came from looking at the regular representation $\mathbb{C}[G]$, and decomposing it as

$$\mathbb{C}[G] = \bigoplus V_i^{d_i}.$$

But we have $V_i^{d_i} \cong \mathrm{End}(V_i)$, where we can think of $\mathrm{End}(V_i)$ as $\mathrm{Mat}_{d_i}(\mathbb{C})$ (this is because V_i is d_i-dimensional, so specifying an endomorphism (or linear operator) on V_i is the same as specifying the d_i images of the basis vectors). So we can write

$$\mathbb{C}[G] \cong \bigoplus \mathrm{End}(V_i).$$

This generalizes to compact groups — except that if looking at functions, a typical function can't be written as a *sum* of elements, but rather as an infinite series. So we instead have

$$C(G) = \widehat{\bigoplus} \mathrm{End}(V_i).$$

Example 35.11
If $G = \mathrm{U}(1)$ (which is just $S^1 = \{z \mid |z| = 1\}$), then irreducible representations are indexed by integers. This turns into the theory of Fourier series, as mentioned earlier.

The generalization of this case is harmonic analysis, where functions on the group are written in terms of an infinite series. In fact, to understand the spectrum of the Laplacian on the sphere, we consider this decomposition for functions on SO(3).

But if we don't work with infinite series, and just consider $\bigoplus \mathrm{End}(V_i)$, then we arrive at the notion of an algebraic group — we have

$$\mathrm{MSpec}(\bigoplus \mathrm{End}(V_i)) = G_{\mathbb{C}},$$

where $G_{\mathbb{C}}$ is an algebraic group (for example, $\mathrm{GL}(n, \mathbb{C})$). This is studied in **18.737**.

Beyond the theory of representations of compact groups, one can also work with *non-compact* Lie groups (closed but not necessarily compact subgroups of $\mathrm{GL}(n, \mathbb{C})$), such as $\mathrm{SL}(n, \mathbb{R})$. Then most representations are infinite-dimensional. This is also studied in **18.755** and its continuations.

35.2.3 Factorization

We've seen a story about factorization in quadratic number fields. This is considered closer to number theory than abstract algebra — factorization in $\mathbb{Z}(\sqrt{d})$ generalizes to rings of algebraic integers in number fields. This is studied in number theory, by using the action of the Galois group. A typical question is to start with a prime ideal in a number field, and try to understand how it decomposes in a larger number field.

An example of this is quadratic reciprocity, a classical result in number theory (which is explained in **18.781**). Part of the theorem is the following:

Example 35.12
If p and q are primes with $p \equiv 1 \pmod 4$, then p is a square mod q if and only if q is a square mod p.

This is an important result with many proofs, including elementary ones. But there's also a proof that generalizes and connects well to algebraic number theory, and in fact it relates to ideas we've seen in class. The idea is to consider $\mathbb{Q}(\zeta_p)$, where $\zeta_p = \exp(2\pi i/p)$. As proved in class, this contains $\mathbb{Q}(\sqrt{\pm p})$. By analyzing the factorization of q in these two fields, and looking at it in two ways using the description of the Galois group, one can obtain this beautiful statement. In number theory, this is generalized to higher reciprocity laws.

35.2.4 Rings and Modules

In rings and modules, one of the main theorems we saw was the classification of finitely generated modules over a PID. The class **18.705** on commutative algebra develops this much further.

Commutative algebra is also closely related to algebraic geometry. For example, if we have a ring $R = \mathbb{C}[x_1, \dots, x_n]/I$, we can consider its maximal spectrum $\mathrm{MSpec}(R) \subset \mathbb{C}^n$.

Then for an R-module M and $x \in \mathrm{MSpec}(R)$, we get a $\mathbb{C}$-vector space — x corresponds to a maximal ideal $\mathfrak{m}_x$, and $R/\mathfrak{m}_x = \mathbb{C}$ (as we proved using Nullstelensatz). So $M/\mathfrak{m}_x M$ is a $\mathbb{C}$-vector space (which is finite-dimensional if M was finitely generated). This gives a family of vector spaces indexed by $x \in \mathrm{MSpec}(R)$. This idea is also studied in topology and differential geometry, namely vector bundles; and this analogy (connecting it to ideals in commutative algebra) is important.

35.2.5 Galois Theory

We saw a story relating groups to extensions; the key examples were extensions of number fields and of $\mathbb{C}(t)$ (the latter was just sketched, but it's still an important example).

Historically, at about the same time Galois worked on this, Abel was thinking about the same problem, but more in terms of geometry. Galois theory as presented here allows us to say that for a *specific* polynomial equation, there's no formula for the solution in radicals. On the other hand, Abel's work considered universal formulas, and showed that they relate to Riemann surfaces (which relate to complex analysis).

A Dimensions of Irreducible Characters

In this section, we provide a proof of the final part of the main theorem of representation theory:

Theorem A.1
If $\rho : G \to \mathrm{GL}(V)$ is an irreducible representation of dimension d, then d divides $|G|$.

These notes are based on a writeup by Professor Bezrukavnikov posted to Canvas.

Recall that we extended the definition of ρ to all *linear combinations* of elements in G, or equivalently functions $f : G \to \mathbb{C}$, using the natural formula

$$\rho(f) = \sum_{g \in G} f(g)\rho(g).$$

Then $\rho(f)$ is in $\mathrm{End}(V)$ for any function f.

To start with, we find a natural construction in which $|G|/d$ arises.

Proposition A.2
For any irreducible representation $\rho : G \to \mathrm{GL}(V)$ of dimension d, we have

$$\rho(\overline{\chi_\rho}) = \frac{|G|}{d} \cdot \mathrm{Id}\,.$$

Proof. Since $\overline{\chi_\rho}$ is a class function, then $\rho(\overline{\chi_\rho})$ is G-equivariant. But by Schur's Lemma, since ρ is scalar, the only G-equivariant endomorphisms are scalar maps; so $\rho(\overline{\chi_\rho})$ must be of the form $\lambda \cdot \mathrm{Id}$ for some $\lambda \in \mathbb{C}$. Now we can compute λ by taking the trace: we saw earlier that $\mathrm{Tr}\,\rho(f) = |G|\langle \chi_\rho, \overline{f}\rangle$, so

$$\mathrm{Tr}\,\rho(\overline{\chi_\rho}) = |G|\langle \chi_\rho, \chi_\rho\rangle = |G|,$$

using the fact that the irreducible characters are orthonormal and therefore $\langle \chi_\rho, \chi_\rho\rangle = 1$. But this trace must also be $d\lambda$, so $\lambda = |G|/d$. (The properties used in this proof are discussed in more detail in Lecture 7.) □

Now in order to prove that $|G|/d$ is an integer from here, we use a bit of theory about algebraic integers.

Definition A.3
A complex number is a **algebraic integer** if it is the root of a monic polynomial with integer coefficients.

Lemma A.4
Algebraic integers have the following standard properties:

(a) If α and β are algebraic integers, so are $\alpha + \beta$ and $\alpha\beta$.

(b) If $\alpha \in \mathbb{Q}$ is an algebraic integer, then $\alpha \in \mathbb{Z}$.

The course discusses algebraic integers in more detail in future lectures; the two properties listed here are proved in Lectures 14 and 25, respectively.

It is now enough to prove the following proposition:

Proposition A.5
Let $\rho : G \to \mathrm{GL}(V)$ be any representation of G. Then if $f : G \to \mathbb{C}$ is a function such that $f(g)$ is an algebraic integer for every g, and $\rho(f) = r \cdot \mathrm{Id}$ for a rational number r, then r must be an integer.

It's clear that the two propositions together imply our theorem — by the first proposition, we have that $\rho(\overline{\chi_\rho}) = |G|/d \cdot \mathrm{Id}$, and we know that $\overline{\chi_\rho}(g)$ is an algebraic integer for all g, since $\chi_\rho(g)$ is a sum of roots of unity (and roots of unity are all algebraic integers). So by the second proposition, $|G|/d$ must be an integer.

In fact, a stronger statement is true — if f is *any* function on G such that $f(g)$ is an algebraic integer for all $g \in G$, then every eigenvalue of $\rho(f)$ is an algebraic integer. But this is much harder to prove, so we will only prove the special case necessary for our theorem.

Proof. We will show that $\operatorname{Tr} \rho(f)^n$ is an integer for all n, which suffices — this is because $\rho(f)^n = r^n \cdot \mathrm{Id}$, so dr^n is an integer for all n, and therefore r must be an integer (if a prime p divided its denominator, then for sufficiently large n the power of p in the denominator of r^n would be greater than the power of p dividing d).

When $n = 1$, we have

$$\operatorname{Tr} \rho(f) = \sum_{g \in G} f(g)\chi_\rho(g),$$

and $f(g)$ and $\chi_\rho(g)$ are both algebraic integers. So $\operatorname{Tr} \rho(f)$ is an algebraic integer. But this trace is also rational, as it is equal to dr; therefore $\operatorname{Tr} \rho(f)$ is an integer.

Now for the case of general n, it is enough to find a function f_n such that $\rho(f)^n = \rho(f_n)$ and $f_n(g)$ is again an algebraic integer for all $g \in G$ — then we can apply the above reasoning to f_n instead. To find such a function, we use the following construction:

Definition A.6
Given two functions $\phi : G \to \mathbb{C}$ and $\psi : G \to \mathbb{C}$, their **convolution** is the function $\phi * \psi$ defined as

$$(\phi * \psi)(g) = \sum_{h \in G} \phi(h)\psi(h^{-1}g).$$

Lemma A.7
For any two functions ϕ and ψ, we have

$$\rho(\phi * \psi) = \rho(\phi)\rho(\psi).$$

Proof. The space of functions on G has a basis consisting of the functions δ_g which map g to 1 and all other elements to 0, where $\rho(\delta_g) = \rho(g)$ for each $g \in G$. Then convolution is defined by setting $\delta_g * \delta_h = \delta_g\delta_h$ for all $g, h \in G$ and extending to all functions using linearity. So we have

$$\rho(\delta_g * \delta_h) = \rho_{gh} = \rho_g\rho_h = \rho(\delta_g)\rho(\delta_h),$$

and the statement for general functions ϕ and ψ then follows from linearity. □

Then we can take

$$f_n = \underbrace{f * f * \cdots * f}_{n \text{ times}}.$$

This satisfies $\rho(f_n) = \rho(f)^n$, and since f_n is constructed by repeatedly taking sums and products of algebraic integers, $f_n(g)$ must be an algebraic integer for all g as well.

So then $\operatorname{Tr} \rho(f)^n = \operatorname{Tr} \rho(f_n)$ is an integer for all n, as desired. □

This concludes the proof of the theorem.

Resource: Algebra II Student Notes

Spring 2022

Instructor: Roman Bezrukavnikov
Notes taken by Sanjana Das and Jakin Ng