18.702: Algebra II

Lecturer: Professor Mike Artin

Notes by: Andrew Lin

Spring 2019

1 February 6, 2019

Professor Artin knows that most of us were in 18.701, but for those of us that weren't, he likes to be called Mike.

We'll start the semester with group representations.

Definition 1

A **matrix representation** of a group G is a homomorphism

$$R : G \to GL_n,$$

where G is a finite group and $GL_n = GL_n(\mathbb{C})$.

In other words, the map R takes a group element and returns a matrix. We'll denote the matrix corresponding to g as R_g, and because R is a homomorphism, we know that $R_g R_h = R_{gh}$ for all $g, h \in G$.

Example 2

What are some matrix representations for $G = S_3$?

1. We can start with the **permutation representation**, where we just write down the permutation matrix that corresponds to our group element. S_3 is generated by $x = (123), y = (23)$, so it suffices to write down P_x and P_y. And remember that we permute the columns, not the entries, so the defining elements are

$$P_x = \begin{pmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix}, \quad P_y = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}.$$

2. Alternatively, we can use the sign representation (which takes even permutations to 1 and odd permutations to -1):

$$\Sigma_x = \begin{pmatrix} 1 \end{pmatrix}, \quad \Sigma_y = \begin{pmatrix} -1 \end{pmatrix}.$$

   This shows that we can use 1 by 1 matrices as representations – we just treat them like complex numbers.

3. Finally, there is a **standard representation** using the fact that $S_3 = D_3$ is the dihedral group for a triangle. Then (defining $c = \cos\frac{2\pi}{3} = -\frac{1}{2}, s = \sin\frac{2\pi}{3} = \frac{\sqrt{3}}{2}$) we have

$$A_x = \begin{pmatrix} c & -s \\ s & c \end{pmatrix}, \quad A_y = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix},$$

There are three measures of importance for mathematical concepts:

- Utility,
- Beauty,
- ...the last one was forgotten.

But group representations are both useful and beautiful. And they're useful in other fields as well:

Example 3
Benzene is basically a hexagon at some initial time $t = 0$, but atoms making up the molecule are moving around with some velocity.

In principle, we can figure out what happens to the molecule, but it's easier to figure out vibrational modes, and representation theory helps with this!

Fact 4
Professor Artin wanted to be a chemist, but then he changed his mind. "Don't read the chemistry textbook."

Under the permutation representation P, there's a fixed vector

$$v = \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} : Pv = v.$$

Let's suppose we change basis to some (v_1, v_2, v_3), where $v_1 = v$. Then the new representation is P', and P'_g can be represented as $Q^{-1}P_gQ$ for some fixed basechange matrix Q. Specifically, this matrix Q has column vectors v_1, v_2, v_3.

But now P'_g is of the form

$$\begin{pmatrix} 1 & * & * \\ 0 & A & B \\ 0 & C & D \end{pmatrix}$$

where the asterisks are "junk."

Proposition 5 (Maschke's theorem, version 1)
We can eliminate the junk by choosing v_2 and v_3 carefully. In other words, if we have a vector fixed by all elements in our representation, we can change basis so that the first row and column are all 0 except for a 1 on the diagonal.

Specifically, if v_1 is sent to itself, the orthogonal space to the span of v_1 is sent to itself. We're lucky here because the matrices in the representation P_x and P_y are already orthogonal! So now if we just choose our other basis vectors such that

$$v_2, v_3 \in v_1^{\perp},$$

we automatically get that

$$P' = \begin{pmatrix} 1 & 0 & 0 \\ 0 & A & B \\ 0 & C & D \end{pmatrix}$$

and $\begin{pmatrix} A & B \\ C & D \end{pmatrix}$ will be another representation. Spoiler alert: it'll be the **standard representation** in some basis. So this allows us to say that P_g is isomorphic to $T \oplus A$, where T is the trivial representation (since we had a 1 in the top left corner of P).

Generally, we try to work without a chosen basis. We will need one to explicitly write down a representation, but it's like linear transformations: it's easier to think of them without a basis at first.

Next time, there will be less notation, but we do need to set everything up today. Let V be a vector space, and let $GL(V)$ be the group of invertible linear operators on V. If the dimension of V is n and we have some basis $(v_1, \cdots, v_n)$, then $GL(V)$ corresponds to $GL(n)$ by corresponding ρ to R (the matrix of ρ).

Definition 6

Let V be a vector space. A **representation of G on V** is a homomorphism $G \to GL(V)$.

Again, we write $g \to \rho_g$, and we still have $\rho_g \rho_h = \rho_{gh}$ by the definition of a homomorphism.

Definition 7

If $\rho_g W \subset W$ for all g, then $\rho_g W = W$ (since ρ_g is invertible), so we call W an **invariant subspace** of V.

Then if we choose a basis $V = (v_1, \cdots, v_r, \cdots, v_n)$, where the first r vectors are in W and the remaining form another subspace U, the matrix R_g of ρ_g will be

$$\begin{pmatrix} A_g & * \\ 0 & B_g \end{pmatrix}$$

where the top left matrix A_g is r by r, and A is the restriction of R to W. This corresponds to another representation $\alpha : G \to GL(W)$, and what Maschke's theorem says is that we can get rid of the $*$ junk again! This means we can say that B is the restriction of R to U, and if we denote β to be the homomorphism $G \to GL(U)$,

$$\rho = \alpha \oplus \beta.$$

This is called a **direct sum**, and we can write our representation in this way if there exist subspaces W, U both invariant under ρ, and $V = W \oplus U$.

Definition 8

An **irreducible representation** $GL(V)$ is one where there does not exist a proper invariant subspace.

The sign representation is obviously irreducible. Why is the standard representation irreducible? Any proper subspace would have dimension 1, and if it were invariant, it would need to be an eigenvector. And there are no vectors that are simultaneously eigenvectors for A_x and A_y.

Definition 9

The **character** χ of a representation P is a function on G such that

$$\chi(g) = \text{tr}(\rho_g).$$

Let's make a **character table**: let χ_1 be the character for the trivial representation, χ_2 be the one for sign representation, and χ_3 be the one for the standard representation. We know what all the matrices look like, so filling out the table directly is pretty easy:

	1	x	x^2	y	xy	x^2y
χ_1	1	1	1	1	1	1
χ_2	1	1	1	-1	-1	-1
χ_3	2	-1	-1	0	0	0

Here, we should recall that **trace of a matrix is invariant under conjugation**. x and x^2 are conjugate, and y, xy, x^2y are conjugate, so that's why those columns look identical.

Three interesting facts, which turn out to not be coincidences:

- The characters are constant on conjugacy classes.
- $|\chi|^2$ – that is, the sum of the squares of the characters of all group elements – is always 6 in all cases, which is the size of the group.
- Treating each χ_i as a vector, all three are orthogonal to each other.

Theorem 10
Let G is a finite group. Then the irreducible characters are orthogonal, and their squared lengths are all $|G|$.

Notice we didn't list the reducible representation P in our character table. But that's because any direct sum is a linear (integer) combination of the characters!

Here's a better way to state Maschke's theorem:

Theorem 11 (Maschke's theorem)
Every representation of a finite group G is a direct sum of irreducible representations.

With this, if we have any representation ρ with character χ, we can just use the projection formula to compute the direct sum $\chi = \sum a_i \chi_i$ by finding the inner product with each character (we'll go into more detail with this later). In our case, $\chi_P = \chi_1 + \chi_3$.

We can also consider representations as rotations of $\mathbb{R}^3$, but we're out of time.

2 February 8, 2019

To answer a question posed at class, we're always assuming the vector spaces are complex. It's possible for us to have real vector spaces, but then we have a real representation that is actually just complex. (This is not a particularly important idea, aside from a few modifications to our definitions.)

To recap, if V is a complex vector space, $GL(V)$ is the set of invertible linear operators on V, and G is a finite group, then a **representation** of G on V is a homomorphism $\rho : G \to GL(V)$ sending $g \to \rho_g$.

Proposition 12
Representations of G on V correspond bijectively to linear operations of G on V.

Let's describe the group operation: for a group element g, we send v to $g * v$, where we need to follow the rule

$$g * (h * v) = (gh) * v, \quad 1 * v = v.$$

Since the operation is linear, we also have

$$g * (v + v') = g * v + g * v', g * (cv) = cg * v.$$

This is the analog to **treating group actions as permutations** from 18.701 (every group action corresponds to an element of the permutation group), and now it makes sense to drop the $*$ from our notation. So $\rho_g(v)$ will be denoted gv.

Last time, we also defined an invariant subspace W to be one where $\rho_g W \subset W$ for all $g \in G$. (Since g has an inverse, this also means $\rho_g W = W$.) Now if W is an invariant subspace, we can restrict ρ to W and get a representation on W. This motivates having a definition of an **irreducible representation**: it's one where there is no proper invariant subspace. We then found that if W and U are invariant, and $V = W \oplus U$, ρ is a direct sum of the restrictions to W and U. Picking a basis

$$V = (v_1, v_2, \cdots, v_k, v_{k+1}, \cdots, v_n)$$

where the first k elements correspond to W and the remainder correspond to U, the matrix of ρ_g (for each g) with respect to the basis is (in block form)

$$R_g = \begin{pmatrix} A_g & 0 \\ 0 & B_g \end{pmatrix}.$$

(And Maschke's theorem says that every representation is the direct sum of irreducible representations, which makes a lot of things nicer.)

Still doing review here: we defined a **character** χ of a representation ρ, which is a map from the group G to a complex number $\mathbb{C}$, sending g to the trace of ρ_g.

Example 13

For any representation ρ, we have $\chi(1) = \dim V$, because ρ_1 is the identity matrix.

Notation-wise, we will also call this $\dim \rho$ and $\dim \chi$.

Fact 14

χ is constant under conjugation, since the trace is invariant under conjugation. (This is because trace is commutative, and this means that $\chi(hgh^{-1}) = \chi(hh^{-1}g) = \chi(g)$.)

Remember that (from the characteristic polynomial), the trace of a matrix is the sum of the eigenvalues. So $\chi(g) = \chi(h)$ for all characters χ if g and h are in the same conjugacy class.

Last time, we wrote down a character table for irreducible representations of S_3:

	1	x	x^2	y	xy	x^2y
χ_1	1	1	1	1	1	1
χ_2	1	1	1	-1	-1	-1
χ_3	2	-1	-1	0	0	0

Notice that all rows have the same "length" as vectors, and that each length squared is the order of the group. Since all elements in a given conjugacy class look identical in this table, it's enough to just write down one column for each

conjugacy class. There's a catch though: we know the row vectors in our table are also orthogonal, but this is less clear if we write the table with conjugacy classes instead of elements. Thus, it's customary to use the **compact character table**, where we pick only one element from each conjugacy class (and write a number above to indicate how many elements of each group we have):

	(1)	(2)	(3)
	1	x	y
χ_1	1	1	1
χ_2	1	1	-1
χ_3	2	-1	0

It's nice if our characters all have "length" 1, so we'll divide through by the order of the group. Specifically, let's define an inner product:

Definition 15

If χ, χ' are characters of G, then let

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

Notice the similarities with the standard Hermitian form on $\mathbb{C}^n$:

$$\langle X, Y \rangle = X^* Y = \sum_i \overline{x_i} y_i,$$

where we have the properties $\langle Y, X \rangle = \overline{\langle X, Y \rangle}$ and $\langle X, X \rangle = |X|^2$.

So now we can compute those inner products for the irreducible representations of S_3:

$$\langle \chi_2, \chi_3 \rangle = \frac{1}{6} \left(1 \cdot \overline{\chi_2(1)} \chi_3(1) + 2 \cdot \overline{\chi_2(x)} \chi_3(x) + 3 \cdot \overline{\chi_2(y)} \chi_3(y) \right) = 0,$$

as expected. We also find that $\langle \chi_1, \chi_1 \rangle = 1$ and so on, and now let's restate this as a general theorem:

Theorem 16 (Main Theorem)

Let $\chi_1, \chi_2, \cdots, \chi_k$ be the irreducible characters of a group.

1. These characters are orthonormal under the defined inner product: $\langle \chi_i, \chi_j \rangle$ is 1 if $i = j$ and 0 otherwise.
2. The number of irreducible characters is the number of conjugacy classes of G.
3. If $d_1, d_2, \cdots$ are the dimensions of χ_i, then $\sum d_i^2 = |G|$.
4. Finally, if ρ, ρ' are representations of G, then ρ is isomorphic to ρ' if and only if $\chi = \chi'$.

Elaborating on point (2), we know that if the characters are orthogonal, we do need to have **at most** as many characters as conjugacy classes. The Main Theorem tells out that this is an equality! And now we can compute the character table without needing to find the representations explicitly. One of the characters must be trivial:

	(1)	(2)	(3)
	1	x	y
χ_1	1	1	1
χ_2			
χ_3			

Now, we have to have the sum of squares of entries in the first column be 6: there's only one way to do this.

	(1)	(2)	(3)
	1	x	y
χ_1	1	1	1
χ_2	1		
χ_3	2		

But for the two-dimensional representation, x is a 3-cycle, so if λ and λ' are the eigenvalues of ρ_x, $x^3 = 1 \implies \lambda^3 = \lambda'^3 = 1$. Letting ω be a cube root of unity, the other cube root of 1 is $\omega^2 = \overline{\omega} = \omega^{-1}$.

The important thing here is that $\chi(x^2) = \chi(x)$, since x and x^2 are conjugate. Since χ is the sum of the eigenvalues and x^2 has eigenvalues that are the square of x's, we either have $\chi(x) = 1 + 1$ or $\omega + \overline{\omega}$. The former is too big – since $\langle \chi_3, \chi_3 \rangle = 1$, we can't have $\overline{2} \cdot 2 + 2 \cdot \overline{2} \cdot 2 = 12$ in the sum for the inner product. Thus $\chi_3(x) = -1$, and for $\langle \chi_3, \chi_3 \rangle = 1$, we must have $\chi_3(y) = 0$. (Alternatively, use the fact that χ_1 and χ_3 are orthogonal.)

	(1)	(2)	(3)
	1	x	y
χ_1	1	1	1
χ_2	1		
χ_3	2	-1	0

Filling out the rest is pretty routine. We're already one hour behind the syllabus because we don't have Maschke's theorem proved, but this is a cool method for finding character tables!

3 February 11, 2019

Let's do a quick review: we define a **character** χ of a representation $\rho : G \to GL(V)$ via $\chi(g) = \text{tr}(\rho_g)$. Then we defined a Hermitian form

$$\langle \chi, \chi' \rangle = \frac{1}{|G|} \sum_g \overline{\chi(g)} \chi'(g).$$

Last time, we stated Theorem 16, the Main Theorem: if $\chi_1, \cdots, \chi_k$ are the irreducible characters for a group G, then they are orthonormal under our Hermitian form. We also have some other computational results: the number of irreducible characters is the number of conjugacy classes, and if we let d_i be the dimension of χ_i (the dimension of the matrices for the corresponding representation), then $\sum d_i^2 = |G|$. Finally, two representations with the same character are isomorphic: in other words, if we pick the right bases, the matrices for the two representations will be the same.

We also stated Theorem 11, Maschke's theorem: every representation is isomorphic to a direct sum of irreducible representations, so every character is a sum of irreducible characters. Since our irreducible characters are orthonormal, we can actually decompose using the projection formula:

$$\chi = \sum_i \langle \chi, \chi_i \rangle \chi_i.$$

Let's go back to our example:

Example 17

Take $G = S_3 : X^3 = y^2 = 1, yx = x^2y$. Recall that our character table here is

	(1)	(2)	(3)
	1	x	y
χ_1	1	1	1
χ_2	1	1	-1
χ_3	2	-1	0

To make some progress on proving these results, we'll need to define a few special representations. Recall that there's a **permutation representation** which we described at the beginning of class:

Definition 18

Given a finite set $S = \{S_1, \cdots, S_n\}$, and suppose G operates on S. (This means each group element $g \in G$ permutes the elements of S.) The **permutation representation** $R : G \to GL_n$ is defined such that R_g is the permutation matrix associated with the permutation of elements of S.

For example, S_3 acts on the set $\{1, 2, 3\}$, so $x = (123)$ can be written as $\begin{pmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix}$, and $y = (23)$ can be written as $\begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}$. Then the character of R, denoted $\chi_R(g)$, just tells us **the number of elements fixed by** g (since those are the on-diagonal elements).

Definition 19

Let G operate on itself using left multiplication: g sends h to gh. Then the **regular representation** is the permutation representation associated with this operation.

For example, the dimension of the regular representation for S_3 is 6, because we are permuting the 6 elements of the group G. But notice that any element other than the identity fixes **nothing** in our group, so the trace will always be 0 in those cases. This means that the character χ_{reg} takes on a simple form:

	(1)	(2)	(3)
	1	x	y
χ_1	1	1	1
χ_2	1	1	-1
χ_3	2	-1	0
χ_{reg}	6	0	0

Decomposing this character by the projection formula, notice that most of the terms in our sum disappear:

$$\langle \chi_{\text{reg}}, \chi_i \rangle = \frac{1}{6} \sum_g \overline{\chi_{\text{reg}}(g)} \chi_i(g) = \frac{1}{6} \left[6 \cdot \chi_i(1) + 0 \right] = \chi_i(1) = d_i.$$

And this will be true in general as well:

Proposition 20

The character of the regular representation of a group G satisfies

$$\chi_{\text{reg}} = \sum \chi_i d_i.$$

The above equation tells us that for $g \neq 1$,

$$0 = \sum_i \chi_i(g) d_i,$$

but more importantly, we can immediately prove point (3) of Theorem 16! If we plug in $g = 1$, we find that

$$\boxed{|G|} = \sum_i \chi_i(1) d_i = \boxed{\sum_i d_i^2},$$

as desired.

Next, let's consider the case where G is an **abelian** group. Conjugation is trivial ($ghg^{-1} = h$ for all g, h), so all elements are in their own conjugacy class. So because the number of conjugacy classes is $|G|$, and that's also equal to the number of irreducible representations by point (2) of Theorem 16, we have

$$|G| = \sum_{i=1}^{|G|} d_i^2.$$

Therefore, all d_i must be equal to 1 – all of our representations are just homomorphisms from G to $\mathbb{C}$.

Example 21

Let's write down the character table for $G = C_3 = \langle x \rangle$.

Then the character table looks like

	(1)	(1)	(1)
	1	x	x^2
χ_1	1	1	1
χ_2	1		
χ_3	1		

To fill out the rest of the table, each entry is just the trace of a 1 by 1 matrix, which is the single entry in that matrix. Since $x^3 = 1$, all of these entries must be cube roots of 1. And there's exactly three cube roots of one, so the rest is pretty easy:

	(1)	(1)	(1)
	1	x	x^2
χ_1	1	1	1
χ_2	1	ω	ω^2
χ_3	1	ω^2	ω

Using this idea, we can also describe one-dimensional representations of **any finite group**. Let G be arbitrary: now we have a homomorphism $\rho : G \to GL_1 = \mathbb{C}^\times$. The complex numbers are abelian under multiplication, so we're forcing **abelian relations** on our group now.

Let $\overline{G}$ be the **abelianization** of G, which is just some quotient of G by the set of all commutators in G. Then **one-dimensional representations of G correspond to one-dimensional representations of the abelian group $\overline{G}$.**

Example 22

Let's use this to find the one-dimensional representations of $G = S_3$ again.

In $\overline{G}$, we have $\overline{xy} = \overline{yx} = \overline{x}^2\overline{y}$, so $\overline{x} = 1$. This means $\overline{G}$ is the cyclic group with two elements, and this has two one-dimensional representations. So S_3 also has two 1-dimensional representations: indeed, these come from the first and second rows of the character table.

Example 23

It's time to compute a new character table: let $G = T$ be the rotational symmetries of a tetrahedron.

Remember that rotations are conjugate if they have the same angle. Letting x be a rotation of $\frac{2\pi}{3}$ around a vertex, and letting z be a rotation of π around an edge, we have our conjugacy classes:

(1)	(4)	(4)	(3)
1	x	x^2	z

According to the main theorem, we'll have four irreducible characters. We need the sum of the squares of the dimensions to satisfy $\sum d_i^2 = |G| = 12$, so we need to add the three remaining squares to get 11. There's only one way to do that, so we can fill out some more of the character table:

	(1)	(4)	(4)	(3)
	1	x	x^2	z
χ_1	1	1	1	1
χ_2	1			
χ_3	1			
χ_4	3			

Remember that $\overline{G}$, the abelianization, will tell us information about the one-dimensional representations. There are 3 one-dimensional representations, so $\overline{G}$ is the cyclic group of order 3. That means we're quotienting by a group of order 4 to get the abelianization, and the only way to do that is to have the conjugacy class of z combine with the identity (so that we have three different conjugacy classes each corresponding to 4 elements of T). Thus, $z = 1$ in the abelization, and thus $\chi_2(z) = \chi_3(z) = 1$.

	(1)	(4)	(4)	(3)
	1	x	x^2	z
χ_1	1	1	1	1
χ_2	1			1
χ_3	1			1
χ_4	3			

Now we can fill the rest of χ_2 and χ_3 in using the C_3 character table, and we're almost done:

	(1)	(4)	(4)	(3)
	1	x	x^2	z
χ_1	1	1	1	1
χ_2	1	ω	$\overline{\omega}$	1
χ_3	1	ω	$\overline{\omega}$	1
χ_4	3			

The bottom right corner, $\chi_4(z)$, is the sum of 3 eigenvalues for a matrix $\rho_4(z)$. But $z^2 = 1$, so all eigenvalues are ± 1. Adding up three such eigenvalues, we know that $\chi_4(z)$ is either $3, 1, -1$, or 3, but since the length of χ_4 should be $\sqrt{12}$, we can't have 3 or -3. In fact, $3^2 \cdot 1 + (\pm 1)^2 \cdot 3 = 12$ exactly, and that means the other entries for x and x^2 must be 0. And this gives us the whole table:

	(1)	(4)	(4)	(3)
	1	x	x^2	z
χ_1	1	1	1	1
χ_2	1	ω	$\overline{\omega}$	1
χ_3	1	ω	$\overline{\omega}$	1
χ_4	3	0	0	$\boxed{-1}$

where we fill in the last -1 using orthonormality of the characters.

Fact 24

We could have also guessed that χ_4 corresponds to the standard representation: rotation by θ in $\mathbb{R}^3$ has trace $1 + 2\cos\theta$, and $1, x, z$ correspond to $\theta = 0, \frac{2\pi}{3}, \pi$.

4 February 13, 2019

It's true that the dimension of an irreducible character divides the order of the group. This is hard to prove, but we're allowed to use it for our problem set.

Let's warm up by doing a character table for the symmetric group S_4. The idea with S_n is that permutations of the same cycle type are in the same conjugacy class: we have

- **1** identity,
- **3** pairs of transpositions,
- **8** 3-cycles,
- **6** transpositions,
- **6** 4-cycles.

With this, we can start constructing our table:

	(1)	(3)	(8)	(6)	(6)
	1	(12)(34)	(123)	(12)	(1234)
χ_1	1	1	1	1	1
χ_2	1	1	1	-1	-1
χ_3					
χ_4					
χ_5					

The first two permutations are the trivial and sign homomorphisms, and now the only way to have the sum of the squares of the dimensions add up to 24 is the following:

	(1)	(3)	(8)	(6)	(6)
	1	(12)(34)	(123)	(12)	(1234)
χ_1	1	1	1	1	1
χ_2	1	1	1	-1	-1
χ_3	2				
χ_4	3				
χ_5	3				

Consider the conjugacy class of (123). Since $(123)^3 = 1$, the eigenvalues of any matrix representing it are either ω or $\overline{\omega}$. Since $(123)^2$ is also in the same conjugacy class, we must have $\chi_3((123)) = 1 + 1$ or $\omega + \overline{\omega}$. The former is too large for the length of χ_3 to be 1, so we must have $\omega + \overline{\omega} = -1$.

	(1)	(3)	(8)	(6)	(6)
	1	(12)(34)	(123)	(12)	(1234)
χ_1	1	1	1	1	1
χ_2	1	1	1	-1	-1
χ_3	2		-1		
χ_4	3				
χ_5	3				

Now notice that χ_1 and χ_2 only differ in the last two columns, so orthonormality requires $\chi_3((12))+\chi_3((1234)) = 0$. But $\chi_3((12))$ is either $2, 0$, or -2 by an eigenvalue argument, and ± 2 are too large. Thus, the last two columns of χ_3 must be 0, and then we finish χ_3 by checking orthogonality against χ_1:

	(1)	(3)	(8)	(6)	(6)
	1	(12)(34)	(123)	(12)	(1234)
χ_1	1	1	1	1	1
χ_2	1	1	1	-1	-1
χ_3	2	2	-1	0	0
χ_4	3				
χ_5	3				

Similarly, for $\chi_4((123))$ and $\chi_5((123))$, we must have $1 + \omega + \overline{\omega} = 0$, and then $\chi_4(12)$ is either $3, 1, -1$, or -3 by arguments like those above. 3 is too big, so now we have a bit more of our table:

	(1)	(3)	(8)	(6)	(6)
	1	(12)(34)	(123)	(12)	(1234)
χ_1	1	1	1	1	1
χ_2	1	1	1	-1	-1
χ_3	2	2	-1	0	0
χ_4	3		0	± 1	
χ_5	3		0	± 1	

To make our job easier, we can calculate χ_{perm}:

	(1)	(3)	(8)	(6)	(6)
	1	(12)(34)	(123)	(12)	(1234)
χ_1	1	1	1	1	1
χ_2	1	1	1	-1	-1
χ_3	2	2	-1	0	0
χ_4	3		0	± 1	
χ_5	3		0	± 1	
χ_{perm}	4	0	1	2	0

Since $\langle \chi_{\text{perm}}, \chi_{\text{perm}} \rangle = 2$, it must have the sum of two irreducible characters, one of which is the identity. After some more work, we end up with our final character table:

	(1)	(3)	(8)	(6)	(6)
	1	(12)(34)	(123)	(12)	(1234)
χ_1	1	1	1	1	1
χ_2	1	1	1	-1	-1
χ_3	2	2	-1	0	0
χ_4	3	-1	0	1	-1
χ_5	3	-1	0	-1	1
χ_{perm}	4	0	1	2	0

Let's move on. Let ρ be the representation of G on a vector space V, and let's use the notation $\rho_g(v) = gv$. We'll list our group elements as

$$G = \{g_1, g_2, \cdots, g_n\},$$

and we'll find a G-invariant subspace W (such that $\rho_g W = W$ for all g) as follows:

Lemma 25

Let $v \in V$ be an arbitrary vector, and let $v_i = g_i v$ for all $1 \leq i \leq n$. Then W, the span of the vectors $v_1, \cdots, v_n$, is invariant.

Proof. We want to show that if $h \in G$, then $hW \subset W$. It's enough to show that each of the basis vectors hv_i are in W. This is clear because

$$hv_i = hg_i v = (hg_i)v = g_j v = v_j$$

for some j, and this is an element of W by definition. □

Corollary 26

If ρ is irreducible, then the dimension of V is at most $|G| = n$. (After all, any ρ with dimension greater than n would have this invariant subspace.)

Unfortunately, remember that this is useless, since we already know from the equation $\sum d_i^2 = |G|$ that the dimensions is at most $\sqrt{n}$. But we can do something a bit less useless: let's find an invariant vector $\tilde{v}$, so $h\tilde{v} = \tilde{v}$ for all $h \in G$. We're going to "average over the group," and this is an important concept:

Proposition 27

Let G be a group acting on a vector space V. Then for any $v \in V$, the vector

$$\tilde{v} = \frac{1}{|G|} \sum_{g \in G} gv$$

is invariant under G.

Proof. The key idea is that multiplication by G is a bijective map. For any $h \in G$,

$$h\tilde{v} = h\left(\frac{1}{|G|} \sum_g gv\right) = \frac{1}{|G|} \sum_g (hg)v$$

and since we're summing over the group, and since the set of hg runs over the whole group – just in a different order – this is identical to $\tilde{v}$, as desired. □

For example, we send the elements of S_3 to a different permutation under

$$\{1, x, x^2, y, xy, x^2y\} \xrightarrow{y} \{y, x^2y, xy, 1, x^2, x\}.$$

And since an irreducible nontrivial representation should have no fixed points, this fixed point $\tilde{v}$ must be the **zero vector** for all irreducible representations except the trivial one.

With this, we finally have all the tools we need to prove Theorem 11, Maschke's theorem:

Proof. We want to show that if W is a proper invariant subspace of V, then there is an **invariant** subspace U such that $V = W \oplus U$. If this is true, then we can repeat the process on W and U by induction: choose a basis for $V = (v_1, v_2, \cdots, v_r, v_{r+1}, \cdots, v_n)$, where the first r vectors form a basis for W and the last $n - r$ vectors form a basis for U. Then R_g, the matrix of ρ_g, must look like

$$R_g = \begin{pmatrix} A_g & 0 \\ 0 & B_g \end{pmatrix}.$$

We can then inductively break down our matrix until each subspace is irreducible.

Tentatively, we want to pick $U = W^\perp$, because we know that $V = W \oplus W^\perp$.

Fact 28

Here, we use a lemma from 18.701: if ρ_g is unitary, V is a Hermitian space, and W is G-invariant, then $W^\perp$ is G-invariant and $V = W \oplus W^\perp$. (The specific fact we use is that **if W is T-invariant, then $W^\perp$ is T^*-invariant**, and $T^* = T^{-1}$ for ar unitary operator.)

There are two potential issues: we may not have a positive definite Hermitian form on V (though we can always choose one), and even if we do have one, the operators ρ_g may not be unitary. (Recall that an operator T on a Hermitian space V is unitary if $\langle v, w\rangle = \langle Tv, Tw\rangle$ for all v, w.) Well, we'll resolve both at the same time: we'll find a positive definite Hermitian form on V so that

$$\langle v, w\rangle = \langle gv, gw\rangle \ \forall g,$$

and then we'll be able to apply the above lemma. Start with some arbitrary positive definite Hermitian form $\{v, w\}$ (we can just pick the standard form in some orthonormal basis), and we'll **average over** G: define

$$\langle v, w\rangle = \frac{1}{|G|}\sum_g \{gv, gw\}.$$

This is a Hermitian form, because each of the $\{gv, gw\}$ is linear in the second variable and follows Hermitian anti-symmetry. It's also positive definite, because each of the $\{gv, gw\}$ are nonnegative. So now we just need to show that

$$\langle hv, hw\rangle = \frac{1}{|G|}\sum_g \{ghv, ghw\}$$

for any $h \in G$. But since right multiplication in G is a bijective map, this is the same sum as $\langle v, w\rangle$ in a different order, and we're done.

Thus we can find an invariant subspace $W^\perp$, and Maschke's theorem holds by induction! □

This construction is called the "unitary trick." Let's show this in action for more concreteness:

Example 29

Start with the matrix $R = \begin{pmatrix} 1 & 1 \\ 0 & -1 \end{pmatrix}$. Notice that $R^2 = I$, so we get a matrix representation of the cyclic group $G = \{1, g\}$ by sending g to R and 1 to the identity I.

Letting $\{X, Y\} = X^*Y$ be the standard Hermitian form, the "good form" we want to use is

$$\langle X, Y\rangle = \frac{1}{2}(X^*Y + (RX)^*RY) = \frac{1}{2}X^*(I + R^*R)Y = \frac{1}{2}X^*\begin{pmatrix} 2 & 1 \\ 1 & 3 \end{pmatrix}Y.$$

Now we verify that R is indeed unitary: we have

$$\langle RX, RY\rangle = \frac{1}{2}X^*R^*\begin{pmatrix} 2 & 1 \\ 1 & 3 \end{pmatrix}RY,$$

and indeed this expands out to the same expression as $\langle X, Y\rangle$.

5 February 15, 2019

Let's warm up by doing a character table for a nonabelian group of order $|G| = 55$. Recall that by the Sylow theorem for $p = 11$, there is a normal subgroup of order 11, and therefore there are **10 elements of order 11**. Then there is a not-normal group of order 5 (or else our group is abelian), so there are 11 Sylow 5-groups and therefore **44 elements of order 5**. These, plus the identity, give all of the elements of our group.

Let x be an element of order 11 and y be an element of order 5. Because the subgroup generated by x is normal, $yxy^{-1} = x^i$ for some i. But since y has order 5, conjugating x by y five times tells us that we need $i^5 \equiv 1 \bmod 11$, so let's take $i = 3$ (all $i \neq 1$ are isomorphic).

Now the conjugacy class of y has the elements $C(y) = \{y, xy, x^2y, \cdots, x^{10}y\}$, since $x^{-1}yx = x^2y$. This gives us four conjugacy classes (for each of the powers of y). Then we also have $C(x) = \{x, x^3, x^9, x^{27}, x^{81}\}$ and $C(x^{-1})$ (which covers the other non-identity powers of x). Thus, we have the following character table skeleton (where χ_6 and χ_7 have dimension 5 from an exercise):

	(1)	(5)	(5)	(11)	(11)	(11)	(11)
	1	x	x^{-1}	y	y^2	y^3	y^4
χ_1	1	1	1	1	1	1	1
χ_2	1						
χ_3	1						
χ_4	1						
χ_5	1						
χ_6	5						
χ_7	5						

Let ζ be a fifth root of unity. The 1-dimensional characters have to give fifth powers of 1, so $\chi_i(y)$ must be a power of ζ. Also, $yxy^{-1} = x^3$ implies that x is sent to 1 in the abelianization, so x is 1 in each of the one-dimensional representations! This fills out a significant amount of the table, and now we can work on the next row:

	(1)	(5)	(5)	(11)	(11)	(11)	(11)
	1	x	x^{-1}	y	y^2	y^3	y^4
χ_1	1	1	1	1	1	1	1
χ_2	1	1	1	ζ	ζ^2	ζ^3	ζ^4
χ_3	1	1	1	ζ^2	ζ^4	ζ	ζ^3
χ_4	1	1	1	ζ^3	ζ	ζ^4	ζ^2
χ_5	1	1	1	ζ^4	ζ^3	ζ^2	ζ
χ_6	5	u	v	a	b	c	d
χ_7	5						

We must have orthogonality between χ_6 and any of χ_1 through χ_5. But there is a common factor of $5 + 5\overline{u} + 5\overline{v}$ in all 5 of those expressions, and now we have the five different equations

$$-5 - 5\overline{u} - 5\overline{v} = 11 \cdot \begin{cases} \overline{a} + \overline{b} + \overline{c} + \overline{d} \\ \zeta\overline{a} + \zeta^2\overline{b} + \zeta^3\overline{c} + \zeta^4\overline{d} \\ \zeta^2\overline{a} + \zeta^4\overline{b} + \zeta\overline{c} + \zeta^3\overline{d} \\ \zeta^3\overline{a} + \zeta\overline{b} + \zeta^4\overline{c} + \zeta^2\overline{d} \\ \zeta^4\overline{a} + \zeta^3\overline{b} + \zeta^2\overline{c} + \zeta^1\overline{d} \end{cases}$$

But if we add up all five of these equations, the right hand side is just 0. So that means $5 + 5\overline{u} + \overline{v} = 0$. In fact, this

also tells us that a, b, c, d are all zero, because the top four equations correspond to the equation

$$\begin{bmatrix} 1 & 1 & 1 & 1 \\ \zeta & \zeta^2 & \zeta^3 & \zeta^4 \\ \zeta^2 & \zeta^4 & \zeta^6 & \zeta^8 \\ \zeta^3 & \zeta^6 & \zeta^9 & \zeta^{12} \end{bmatrix} \begin{bmatrix} \overline{a} \\ \overline{b} \\ \overline{c} \\ \overline{d} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix},$$

and this 4×4 matrix is actually invertible because the fifth roots of unity are distinct:

Theorem 30 (Vandermonde)
For any n,

$$\det \begin{pmatrix} 1 & 1 & \cdots & 1 \\ x_1 & x_2 & \cdots & x_n \\ \vdots & \vdots & \ddots & \cdots \\ x_1^{n-1} & x_2^{n-1} & \cdots & x_n^{n-1} \end{pmatrix} = \prod_{i<j} (x_j - x_i).$$

Proof. For an n by n matrix, the degree of the polynomial determinant is $0+1+2+\cdots+(n-1) = \frac{(n-1)n}{2}$. But note that $(x_i - x_j)$ is always a factor, since $x_i = x_j$ makes the determinant 0 (two columns are identical). Multiplying all such factors together, we've gotten to the degree of the polynomial, and now we can just multiply by a scalar – we'll omit showing that the constant factor is correct. □

So now our question is just to find the four missing entries below:

	(1)	(5)	(5)	(11)	(11)	(11)	(11)
	1	x	x^{-1}	y	y^2	y^3	y^4
χ_1	1	1	1	1	1	1	1
χ_2	1	1	1	ζ	ζ^2	ζ^3	ζ^4
χ_3	1	1	1	ζ^2	ζ^4	ζ	ζ^3
χ_4	1	1	1	ζ^3	ζ	ζ^4	ζ^2
χ_5	1	1	1	ζ^4	ζ^3	ζ^2	ζ
χ_6	5	u	v	0	0	0	0
χ_7	5	u'	v'	0	0	0	0

χ_6 is a five-dimensional representation, and $\chi_6(x)$ is the sum of the eigenvalues of $\rho_6(x)$, so u is a sum of five 11th roots of unity. Let η be one of them: now u is a sum of 5 powers of η, but if a conjguacy class contains x, it contains x^3. So if η^i is one of the eigenvalues involved in $\chi_6(x)$, so are $\eta^{3i}, \eta^{9i}, \eta^{27i}, \eta^{4i}$. There's only two possibilities:

$$u = \begin{cases} \eta + \eta^3 + \eta^9 + \eta^5 + \eta^4 \\ \eta^2 + \eta^6 + \eta^7 + \eta^{10} + \eta^8 \end{cases}.$$

Picking them to be u_1 and u_2 arbitrarily and using orthonormality, we find that our completed character table looks like

	(1)	(5)	(5)	(11)	(11)	(11)	(11)
	1	x	x^{-1}	y	y^2	y^3	y^4
χ_1	1	1	1	1	1	1	1
χ_2	1	1	1	ζ	ζ^2	ζ^3	ζ^4
χ_3	1	1	1	ζ^2	ζ^4	ζ	ζ^3
χ_4	1	1	1	ζ^3	ζ	ζ^4	ζ^2
χ_5	1	1	1	ζ^4	ζ^3	ζ^2	ζ
χ_6	5	u_1	u_2	0	0	0	0
χ_7	5	u_2	u_1	0	0	0	0

We can explicitly calculate u_1 and u_2, since they are the roots of a quadratic

$$x^2 - (u_1 + u_2)x + u_1u_2 = x^2 - x + (\text{garbage}).$$

And the garbage is 3, because expanding out u_1u_2 gives 25 terms: there are 5 ones and 2 sums of all the other roots.

Fact 31

By the way, Vandermonde was a violinist – Professor Artin learned this on a wiki.

Time to move on to the actual topic of today. Suppose we have two representations G of a group: let them be

$$G \xrightarrow{\alpha} GL(U), \quad G \xrightarrow{\beta} GL(V).$$

Then saying that α is **isomorphic** to β? means that if we choose the correct basis, the matrices α_g, β_g are the same for all g. But another way is to say that there exists an isomorphism of vector spaces $V \xrightarrow{T} U$ which is compatible with the operation of G. Here's the commutative diagram:

$$\begin{array}{ccc} V & \xrightarrow{T} & U \\ {\scriptstyle \beta_g}\downarrow & & \downarrow{\scriptstyle \alpha_g} \\ V & \xrightarrow{T} & U \end{array}$$

So we can now redefine the idea of isomorphism more formally.

Definition 32

α and β are **isomorphic** representations if there exists an isomorphism $V \xrightarrow{T} U$ of vector spaces such that

$$\alpha_g T = T\beta_g \implies T = \alpha_g^{-1} T \beta_g$$

for all $g \in G$. We say that the linear transformation T is ***G*-invariant** if $T\beta_g = \alpha_g T$ for all g.

Note here that α is a representation in V, and β is a representation in U.

Lemma 33

If $T : V \to U$ is invariant, then $\ker T$ is an invariant subspace of V and $\operatorname{Im} T$ is an invariant subspace of U.

Proof. Let $K = \ker T$. If $v \in K$, then $T(v) = 0$, and our goal is to show that $\beta_g K \subset K$ for all g. Note that

$$0 = \alpha_g T v = T\beta_g v$$

since T is invariant, and therefore $\beta_g(v) \in K$ for all $v \in K$

On the other hand, we want to show that $\alpha_g(\operatorname{Im} T)$ is contained in $\operatorname{Im} T$ (which is a subspace of U). If $u \in \operatorname{Im} T$, then $u = Tv$ for some v. But then

$$\alpha_g u = \alpha_g T v = T\beta_g v,$$

so $\alpha_g u$ is T of something and is therefore in the image as well. □

Note that it's very hard for T to be invariant: we need $T\beta_g = \alpha_g T$ to be satisfied for all g. Schur's lemma basically says there aren't any such operators:

Lemma 34 (Schur's lemma)
Suppose that α and β are irreducible representations. Then either $T = 0$ or T is an isomorphism.

Proof. Since α and β are irreducible, the only invariant subspaces are 0 and the whole vector space. So the kernel of T is either 0 (in which case T is injective), or V (in which case $T = 0$). Similarly, the image is either U (in which case T is surjective), or 0 (in which case $T = 0$). So unless T is the zero operator, it's both injective and surjetive and is therefore an isomorphism. □

Lemma 35 (Schur's lemma, part 2)
Suppose $U \xrightarrow{T} U$ is invariant for an irreducible representation α. Then $T = cI$ for some constant c.

Proof. Take an eigenvalue λ of T. Then we can see that $S = T - \lambda I$ is invariant, and now because there is an eigenvalue of 0, the kernel is nonzero. Thus, the kernel must be the whole vector space U, and therefore $S = 0 \implies T = \lambda I$. □

Lemma 36
Let $V \xrightarrow{T} U$ be an arbitrary linear transformation and α, β be representations. Then the linear transformation

$$\tilde{T} = \frac{1}{|G|} \sum_g \alpha_g^{-1} T \beta_g$$

is invariant.

Proof. We want to show that if $h \in G$, then

$$\alpha_h^{-1} \tilde{T} \beta_h = \tilde{T}.$$

Since the sum is linear,

$$\alpha_h^{-1} \tilde{T} \beta_h = \frac{1}{|G|} \sum_g \alpha_h^{-1} \alpha_g^{-1} T \beta_g \beta_h = \frac{1}{|G|} \sum_g \alpha_{(gh)^{-1}} T \beta_{gh},$$

and we're summing over all group elements $g' = gh$, so this is just $\tilde{T}$ by definition. □

But this $\tilde{T}$ has to be zero if α_g and β_g are not isomorphic, and we'll see next time how to use these to discuss orthogonality between characters.

6 February 19, 2019

We'll derive the orthogonality relations today. Let's start with a review: let $G \xrightarrow{\alpha} GL(U)$ and $G \xrightarrow{\beta} GL(V)$ be two representations, and let $V \xrightarrow{T} U$ be a linear transformation. Then we defined T to be **invariant** if $T\beta_g = \alpha_g T$ for all g. We found that we can produce invariant linear transformations by averaging

$$\tilde{T} = \frac{1}{|G|}\sum_g \alpha_g^{-1} T \beta_g.$$

This $\tilde{T}$ is invariant, since $\alpha_h^{-1}\tilde{T}\beta_h$ and $\tilde{T}$ are both adding over the whole group, just in a different order. We also learned that (by Schur's lemma) if α and β are both irreducible and not isomorphic, then the only invariant linear transformation from V to U is zero. Also, if α is an irreducible representation and T is an invariant linear operator from U to itself, then $T = cI$ for some c.

Now we just want to extract the characters out of this. We'll use the matrix notation, so we'll use specific complex vector spaces $U \to \mathbb{C}^m, V \to \mathbb{C}^n$, and we'll send α_g, β_g to matrices A_g and B_g. Then $V \xrightarrow{T} U$ becomes m by n matrix M.

Definition 37

Let $\mathbb{C}^{m\times n}$ be the space of $m \times n$ matrices. Define a linear operator Φ on this space

$$\Phi(M) = \frac{1}{|G|}\sum_g A_g^{-1} M B_g.$$

This is basically just our invariant $\tilde{T}$ with different notation. We want to find the trace of Φ – we'll start with the function $F : \mathbb{C}^{m\times n} \to \mathbb{C}^{m\times n}$ which sends M to AMB. Remember that we have an $m \times m$ matrix A and an $n \times n$ matrix B: let $\lambda_1, \cdots, \lambda_m$ be the eigenvalues of A and $\mu_1, \cdots, \mu_n$ be the eigenvalues of B (which are the same as the eigenvalues of B^T). So now if X_i, Y_j are the eigenvectors of A, B^T respectively, with eigenvalues λ_i, μ_j, then $X_i Y_j^T$ (**defining this to be** M_{ij}) is an eigenvector of F with eigenvalue $\lambda_i\mu_j$, since

$$F(M_{ij}) = AX_iY_j^TB = (AX_i)(Y_j^TB) = (AX_i)(B^TY_j)^T = (\lambda_i x_i)(\mu_j y_j^T) = \lambda_i\mu_j M_{ij}.$$

But this gives $m \times n$ eigenvectors of the form M_{ij}, so that's all of them! Sure, $\lambda_i\mu_j$ might be equal for different is and js, but most of the time, they're distinct, so this works in general by continuity.

So now note that the trace of F is the sum of the eigenvalues, which is

$$\sum_{ij}\lambda_i\mu_j = (\lambda_1 + \cdots + \lambda_m)(\mu_1 + \cdots + \mu_n) = \operatorname{tr} A \cdot \operatorname{tr} B.$$

But now trace is linear (the trace of $A + B$ is the trace of A plus the trace of B), and now we've arrived at our result: since Φ is defined as a linear combination of $A_g^{-1}MB_g$s, we can write down an explicit formula for the trace.

Lemma 38

The trace of the linear operator Φ defined above is

$$\operatorname{tr}\Phi = \frac{1}{|G|}\sum_g (\operatorname{tr} A_g^{-1})(\operatorname{tr} B_g) = \frac{1}{|G|}\sum_g \chi_\alpha(g^{-1})\chi_\beta(g).$$

We're working with a finite group here, so every element has finite order. This means that $\chi(g^{-1}) = \overline{\chi(g)}$, because

$\chi(g)$ is a sum of the eigenvalues of α_g, in which all eigenvalues are roots of unity and therefore

$$\lambda_{g^{-1}} = \lambda_g^{-1} = \overline{\lambda_g}.$$

So we can conclude that

$$\operatorname{tr}\Phi = \langle \chi_\alpha, \chi_\beta \rangle,$$

with the Hermitian form that we've defined on characters. But ϕ is an invariant transformation, so if α and β are irreducible and not isomorphic, it's zero (we have **orthogonality** of irreducible characters).

Now, let's work a little harder and get **orthonormality**, which is part (1) of Theorem 16:

Proof. The reason we keep averaging (dividing by the order of the group) is that if M is already invariant,

$$\Phi(M) = \frac{1}{|G|}\sum_g A_g^{-1} M B_g = \frac{1}{|G|}\sum_g M = M$$

by the definition of invariance for M. Another way to phrase this is that whenever $\tilde{M} = \Phi(M)$, $\Phi^2 = \Phi$. This is called a **projection operator**, and with such an operator, the space (in this case $\mathbb{C}^{m\times n}$) is the direct sum of the image and the kernel of Φ. In this case, $M = \tilde{M} \oplus (M - \tilde{M})$.

Lemma 39

Our linear operator Φ, as defined above, satisfies

$$\operatorname{tr}\Phi = \dim \operatorname{Im}\Phi.$$

This is because a projection operator has eigenvalues of 1 for all elements in $\operatorname{Im}\Phi$, but it has eigenvalues of 0 for everything in the kernel – adding these up yields the result. So now by Schur's lemma, when $\alpha = \beta$, an invariant operator on the vector space must be a scalar multiple of the identity. So **all operators** on the space $\operatorname{Im}\Phi$ are scalar multiples of the identity, and this can only happen if $\dim \operatorname{Im}\Phi = 1$. Thus

$$\langle \chi_\alpha, \chi_\alpha \rangle = \operatorname{tr}\ \Phi = \dim \operatorname{Im}\Phi = 1,$$

as desired. □

Example 40

It's time to do another character table: we'll work with a **nonabelian group** of order 8.

Turns out there are two of them – the quaternion group and D_4 – but they have the same character table, so we shouldn't need to figure out which one it is.

If all irreducible representations have dimension 1, then G is abelian, so that's ruled out (since the number of conjugacy classes would be 8). We have the trivial representation, and there are some representations with dimension at least 2 (they can't be of dimension 3). Thus, we must have the sum of squares add to 8: this is just $1, 1, 1, 1, 2$. And it turns out the class equation for a nonabelian group of order 8 can't be $1+1+1+1+4$, so it must be $1+1+2+2+2$ (or else we'd have problems with orthogonality). So here's all the progress we've made so far:

	(1)	(1)	(2)	(2)	(2)
	.	.	.	.	.
χ_1	1	1	1	1	1
χ_2	1				
χ_3	1				
χ_4	1				
χ_5	2				

Now we can "observe" that the next three rows can be like this:

	(1)	(1)	(2)	(2)	(2)
	.	.	.	.	.
χ_1	1	1	1	1	1
χ_2	1	1	-1	-1	1
χ_3	1	1	1	-1	-1
χ_4	1	1	-1	1	-1
χ_5	2				

and now by orthogonality, it's easy to find the last character:

	(1)	(1)	(2)	(2)	(2)
	.	.	.	.	.
χ_1	1	1	1	1	1
χ_2	1	1	-1	-1	1
χ_3	1	1	1	-1	-1
χ_4	1	1	-1	1	-1
χ_5	2	-2	0	0	0

The next hour was originally going to be spent on representations of SU_2, the special unitary group, but we're going to move on to ring theory instead.

7 February 20, 2019

Today, we're starting a new area of algebra.

Definition 41

A **ring** R is a set with operations addition, subtraction, multiplication, and a multiplicative identity 1. Basically, $+$ and $\times$ are two laws of composition, and we have the following axioms:

- $(R, +)$ is an abelian group with identity 0.
- Multiplication is associative and commutative with identity 1. We're assuming all rings are commutative here, but this is not the definition that everyone uses.
- Distributivity holds: $a(b+c) = ab + ac$.

Here are some examples of rings:

Example 42

Any field is a ring, since it has addition and multiplication, and $1 \neq 0$.

(We don't want to rule out the possibility that $1 = 0$ in rings, but we'll come back to that later on.)

Example 43

$\mathbb{Z}$ is a ring of integers. $\mathbb{Z}[i]$, the **Gaussian integers** of the form $\{a + bi, a, b \in \mathbb{Z}\}$, is also a ring. Finally, $\mathbb{R}[x]$, the set of polynomials in x with real coefficients, is a ring.

Lemma 44

Given any ring R and $a \in R$, $a \cdot 0 = 0$.

Proof. We know that

$$0 = 0 + 0 \implies a \cdot 0 = a \cdot (0 + 0) = a \cdot 0 + a \cdot 0$$

and subtract $a \cdot 0$ from both sides (since the cancellation law holds for addition). □

Fact 45

Note that if $0 = 1$, the ring consists of 0 alone. This is because

$$0 = a \cdot 0 = a \cdot 1 = a$$

for all a in the ring.

This is called the **zero ring**, though it doesn't seem to be particularly important.

Polynomials, on the other hand, are a pretty important type of ring:

Definition 46

Let R be any ring. Then the **polynomial ring** is defined as

$$R[x] = \{a_n x^n + a_{n-1} x^{n-1} + \cdots + a_0\}$$

with arbitrary $a_i \in R$ and n an arbitrary nonnegative integer.

Then we just do polynomial multiplication as normal:

$$(a_m x^m + a_{m-1} x^{m-1} + \cdots + a_0)(b_n x^n + \cdots + b_0) = \sum_{i,j} a_i x^i b_j x^j.$$

Thus, the coefficient of x^k is

$$[x^k] = a_k b_0 + a_{k-1} b_1 + \cdots + a_0 b_k = \sum_i a_i b_{k-i}.$$

We can check the ring axioms all work here – addition follows component-wise, but multiplication takes more work and isn't very interesting.

Recall the division algorithm for positive integers:

Fact 47

Let a, b be positive integers. Then there exist unique integers q, r such that

$$b = aq + r, 0 \leq r < a.$$

We want to do something similar for our polynomial rings $R[x]$ as well. We do have to be careful, since we can't divide $x^2 + 1$ by $3x$ (for example) while still having integer coefficients.

Proposition 48

Let $f, g \in R[x]$, and let f be a **monic** polynomial (its leading coefficient a_n is 1). Then there exist $q, r \in R[x]$ such that

$$g = fq + r,$$

and $r = 0$ or the degree of $r(x)$ is less than the degree of f.

In school, we probably learn polynomial long division. (It depends on what country we're from?) Basically if the leading term is $b_n x^n$ and we divide by something with leading term x^m, we get a $b_n x^{n-m}$ leading coefficient.

Fact 49 (Unimportant)

Musical staves come from strings on a musical instrument. **Tablature** has to do with how this is notated. But this may also have to do with long division? This is also related to other "logical German things."

A similar property holds for $\mathbb{Z}[i]$:

Proposition 50 (Division algorithm for the Gaussian integers)

Given $\alpha, \beta \in \mathbb{Z}[i]$ with $\alpha \neq 0$, there exist $q, r \in \mathbb{Z}[i]$ such that

$$\beta = \alpha q + r, |r| < |\alpha|.$$

Proof. Gaussian multiples of α include $\alpha, i\alpha, -\alpha, -i\alpha$, so we can form a grid $\{(a + bi)\alpha : a, b \in \mathbb{Z}\}$. This is a square grid, and it tiles the plane.

Now if β is in one of the squares, it'll be within $|\alpha|$ of one of the vertices. Subtract multiples of α until we move that vertex to 0. Now r is just the difference between β and the closest vertex, and αq is whatever else was subtracted off. □

Note that this answer $\beta = \alpha q + r$ may not be unique, since there are sometimes multiple vertices of our grid that are within $|\alpha|$ of β.

The next definition is even more important for rings than for groups:

Definition 51

A **ring homomorphism** $\phi : R \to R'$ for rings R, R' is a map such that

$$\phi(a + b) = \phi(a) + \phi(b), \quad \phi(ab) = \phi(a)\phi(b), \quad \phi(1_R) = 1_{R'}$$

for all $a, b \in R$.

Notice that we needed to specify that the multiplicative identity goes to the other ring's multiplicative identity, since we don't always have multiplicative inverses! Otherwise, we could have (for example) just sent everything in R to $0_{R'}$, which wouldn't preserve the multiplicative structure in the ways that we want. On the other hand, we don't need to say that $\phi(0_R) = 0_{R'}$, since that is automatic from the group structure of addition.

Proposition 52

For any ring R, there exists a unique homomorphism $\phi : \mathbb{Z} \to R$.

We haven't actually defined addition and multiplication in the integers yet, so this may be kind of hard to prove. But the idea is that $\phi(1_\mathbb{Z})$ goes to 1_R (this is unique), and then $\phi(2) = \phi(1+1) = \phi(1) + \phi(1)$, and so on, which uniquely determines ϕ for every integer. (By the way, the multiplicative identity is unique because $1 = 1 \cdot 1' = 1'$.)

We can read Landau's book, "Introduction to Arithmetic," for more details. There are lots of gory details.

Proposition 53 (Substitution principle)

Let R be a ring, and let $\alpha \in R$. Then there exists a unique homomorphism $\Phi : R[x] \to R$ such that Φ is the identity on constant polynomials and $\Phi[x] = \alpha$. Basically,

$$\Phi(a_m x^m + \cdots + a_0) = a_m \alpha^m + \cdots + a_0.$$

This is pretty easy to check: we are just replacing x with α in every polynomial.

Proposition 54 (Substitution, version 2)

Let $\phi : R \to R'$ be a ring homomorphism, and let $\alpha \in R'$. Then there exists a unique homomorphism

$$\Phi : R[x] \to R'$$

such that $\Phi = \phi$ on constant polynomials and $\Phi(x) = \alpha$.

This uses the same construction, except that we replace constants a_n with $\phi(a_n)$.

Example 55

If we take $R = \mathbb{R}$ and $R' = \mathbb{C}$, where ϕ is the inclusion of $\mathbb{R}$ into $\mathbb{C}$, and we let $\alpha = 1 + i$, the substitution map

$$\Phi : \mathbb{R}[x] \to \mathbb{C}$$

sends a polynomial $f(x)$ to $f(1+i)$.

Example 56

Let $R = \mathbb{Z}$ and $R' = \mathbb{F}_p$, the field of p elements (or the integers mod p). If we choose $\alpha = \overline{2}$ (which is the residue class of 2 mod p), then the map

$$\Phi : \mathbb{Z}[x] \to \mathbb{F}_p$$

sends $a_m x^m + \cdots + a_0$ to $\overline{a_m} 2^m + \cdots + \overline{a_0} \pmod{p}$.

As a related example, we can also define a map $\mathbb{Z}[x] \to \mathbb{F}_p[x]$ by replacing a_m with $\overline{a_m}$ and keeping x. And in this case, note that we don't have to explicitly specify ϕ, since there is a unique map from $\mathbb{Z}$ to $\mathbb{F}_p$.

8 February 22, 2019

Recall the definition of a ring: it has addition, subtraction, multiplication, and a multiplicative inverse 1. We also have rings of the form $R[x]$, which are polynomials with coefficients in R.

Fact 57

It's important to note that x doesn't take on any particular value: it is a variable, so it doesn't satisfy any relations in our polynomial rings.

Polynomials have a basis $\{1, x, x^2, \cdots\}$. We also have a notion of a ring homomorphism

$$\phi : R \to R'$$

which is just compatible with both operations and sends $\phi(1_R) = 1_{R'}$.

Example 58

There exists a unique homomorphism from $\mathbb{Z}$ to any ring R: it sends 1 to $\phi(1)$, 2 to $\phi(1) + \phi(1)$ and so on. We also have the "substitution principle:" if we have a ring homomorphism $\phi : R \to R'$, then there exists a unique ring homomorphism $\phi : R[x] \to R'$ such that $\Phi = \phi$ on constants and $\Phi(x) = \alpha$.

Now, let's introduce the concept of a kernel: it'll applies to the addition operation in our ring (rather than multiplication).

Definition 59

The **kernel** K of a homomorphism $\phi : R \to R'$ is the set of $a \in R$ such that $\phi(a) = 0'$.

The kernel is a normal subgroup of the additive group of R. But we're assuming commutativity of multiplication in all of our rings, so the fact that it is normal is trivial.

We also have an interesting fact: if $a \in K, r \in R$, then $ra \in K$, since $\phi(ra) = \phi(r)\phi(a) = \phi(r)0' = 0'$. So the kernel is an example of something interesting that we don't find in groups:

Definition 60

An **ideal** I of a ring R is a subgroup of $(R, +)$ such that $a \in I, r \in R \implies ra \in I$.

Example 61

There's a unique map from $\mathbb{Z} \to \mathbb{F}_p$: it sends 1 to the residue of 1 mod p, and so on. Then the kernel of this map is $p\mathbb{Z}$, so $p\mathbb{Z}$ is an ideal. (And we can check the definition to see that this is indeed true.)

Example 62

Consider the unique map $\mathbb{Z}[x] \to \mathbb{C}$ that sends $x \to 2 + i$ – what is its kernel?

First of all, we can find a quadratic polynomial with root $2 + i$, and that will be in the kernel. It's

$$f(x) = (x - (2 + i))(x - (2 - i)) = x^2 - 4x + 5.$$

We claim that the kernel of this map is just **all polynomial multiples** of f. To show this, let's say g is in the kernel, so $g(2+i)=0$. Since f is monic, by the division algorithm,

$$g = fq + r,$$

where r has degree less than 2. g and fq are both in the kernel, so r must also be in the kernel. But there is no linear polynomial with integer coefficients with $2+i$ as a root, so we must have $r=0$, and that means f divides g.

Definition 63

A **principal ideal** I of R is of the form $I = R\alpha$ for some $\alpha \in I$. A **principal ideal ring** R is a ring where every ideal is principal.

All ideals we've discussed so far have been principal ideals, and that's for a good reason:

Proposition 64

$\mathbb{Z}$, $\mathbb{F}[x]$ for a field $\mathbb{F}$, and $\mathbb{Z}[i]$ are all principal ideal rings.

Proof. For $\mathbb{Z}$, all subgroups under addition are of the form $n\mathbb{Z}$ or 0. Each of these is the principal ideal generated by n or 0, respectively.

For rings of the form $\mathbb{F}[x]$, let's say we start with an ideal I. If it is just the zero ideal, we're okay. Otherwise, there are some nonzero polynomials, and there exists a monic polynomial with minimal degree (since we have a field). Now we want to show $I = fR$. For any $g \in I$, we can use the division algorithm to find that $g = fq + r$. $fq, g \in I$, so $r \in I$ and therefore $r = 0$ (exactly analogous to we did above).

Finally, for the Gaussian integers, the division algorithm works, so we can just start by picking a complex number in the ideal with minimal norm. □

Example 65

Consider the map $\mathbb{Z}[x] \to \mathbb{F}_p$ where we send $x \to 0$. Then the kernel K of this map is the set of polynomials such that

$$g(0) \equiv 0 \bmod p.$$

But notice that x and p are both in the kernel, and the only way to be a factor of both of these is to include 1 in the kernel. Clearly this is not true, so K is **not** a principal ideal of $\mathbb{Z}[x]$.

In general, we want to ask a question of "how many elements we need to generate I."

Definition 66

A **generator** for an ideal I is a set of elements $\alpha_1, \cdots, \alpha_k$ such that every element of I is a linear combination

$$r_1\alpha_1 + \cdots + r_k\alpha_k, r_i \in R.$$

This is notated as $I = (\alpha_1, \cdots, \alpha_k)$.

So in the above case, x, p generate K, but there's no way to have a single element generate K.

So why is this object called an ideal? We'll start with a motivating example – consider the ring $R = \mathbb{Z}[\sqrt{-5}]$. This ring is the set of complex numbers of the form $\{a + b\sqrt{-5} \mid a, b \in \mathbb{Z}\}$.

Fact 67

This ring doesn't have unique factorization! In particular,

$$6 = 2 \cdot 3 = (1 + \sqrt{-5})(1 - \sqrt{-5}),$$

and there's no way to resolve these two factorizations directly.

Number theorists didn't like this: somehow Gaussian integers have unique factorization, but this "weird ring" didn't. So they introduced the concept of an "ideal element," which "ideally should be there."

The idea is that **we can factor with ideals instead of elements**! So in this ring, let A be the ideal generated by $(2, 1 + \sqrt{-5})$. Notice that in the complex plane, we have a rectangular grid of elements in our ring, so we can get a geometric sense of what's going on. For notation's sake, let $\delta = \sqrt{-5}$. A is closed under complex conjugation, since $1 - \delta = 2 - (1 + \delta)$ is in our ring as well.

Now define $B = (3, 1 + \delta) \implies \overline{B} = (3, 1 - \delta)$. Let's try to multiply some ideals together:

Definition 68

The **product ideal** AB is the set of finite sums of $\alpha_i \beta_j$ where $\alpha_i \in A, \beta_j \in B$. In other words, AB is generated by $\{\alpha_i \beta_j\}$.

In this case, AB has four generators (taking one generator from each of A and B): $(6, 2 + 2\delta, 3 + 3\delta, (1 + \delta)^2)$. Then $(3 + 3\delta) - (2 + 2\delta) = 1 + \delta$ is in AB as well, but notice that this divides all four generators! This means that $(1 + \delta) \subseteq AB \subseteq (1 + \delta)$, so $AB = (1 + \delta)$ is a principal ideal.

Similarly, we can find that $\overline{AB} = \overline{(1 + \delta)} = (1 - \delta)$. Furthermore,

$$(6) = (1 + \delta)(1 - \delta) = AB\overline{AB}.$$

So now

$$\overline{A}A = (4, 2 - 2\delta, 2 + 2\delta, 6)$$

contains $6 - 4 = 2$ and 2 divides everything, so $\overline{A}A = (2)$. Similarly, $\overline{B}B = (3)$, and now

$$\overline{A}A\overline{B}B = (2)(3) = (6).$$

So we've found a better factorization of 6 in our ring. The central idea here is that we've replaced numbers with ideals, and we've resolved the two different methods of factorization. This concept of replacing elements with ideals doesn't always work, but it works in a lot of cases!

9 February 25, 2019

Recall the idea of a quotient group from 18.701: if N is a normal subgroup of G, then the **quotient group** $\overline{G} = G/N$ is a group with cosets as elements: $\overline{a} = aN$, and multiplication is defined as

$$(aN)(bN) = abN.$$

We'll similarly define the concept of a quotient ring now:

Definition 69

Given an ideal I of R, which is closed under addition by I and multiplication by R, define the **quotient ring** R/I to be the set of additive cosets $(a+I)$. The operations are

$$(a+I)+(b+I)=(a+b)+I,$$

$$(a+I)(b+I)=ab+I.$$

This multiplication definition makes sense, since if we have $x, y \in I$, then

$$(a+x)(b+y)=ab+ay+bx+xy=ab+(ay+bx+xy)\in ab+I.$$

Example 70

Let $R=\mathbb{Z}$ and $I=8\mathbb{Z}$. Then

$$(2+I)(2+I)\subset 4+2I+2I+II\subset 4+8\mathbb{Z}=(4+I).$$

But this is clearly not an equality: all elements of $(2+I)(2+I)$ are 4 mod 16, so we never hit 12 (for example).

Notation-wise, we'll write the quotient ring as $R/I=\overline{R}$ and denote the coset $a+I$ as $\overline{a}$. However, note that we can write the same coset with various as, just like we could write the same coset gH with various gs.

Example 71

If $R=\mathbb{Z}$ and $I=p\mathbb{Z}$, then $R/I=\overline{R}$ is $\mathbb{F}_p$.

It's probably good to check the ring axioms to ensure that quotient rings are actually rings. To do that, notice that there is a map

$$\pi : R\to \overline{R}=R/I$$

which sends $a\to \overline{a}=a+I$. This is a surjective homomorphism, so the ring axioms follow from the axioms for R.

Theorem 72 (First Isomorphism Theorem)

Let $\phi : R\to R'$ be a surjective ring homomorphism with kernel $K=\ker\phi$. Then R' is isomorphic to R/K, and there exists a unique isomorphism $\overline{\phi}$ such that $\overline{\phi}\pi=\phi$.

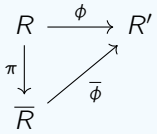

Proof. This is a similar proof to the one for groups. Define $\overline{\phi}(\overline{a})=\phi(a)$. First, we need to check that this definition is consistent (since different a in the same coset $\overline{a}$ should have the same value of $\phi(a)$): if $\overline{a'}=\overline{a}$ for two elements a, a', then $a'=a+x$ for some $x\in K$. Thus, $\phi(a')=\phi(a+x)=\phi(a)+\phi(x)$, and since x is in the kernel, $\phi(a')=\phi(a)$, and our function is well-defined. By definition, $\overline{\phi}$ is then surjective because ϕ is surjective.

Now if $\overline{a}\in\ker\overline{\phi}$, then $\overline{\phi}(\overline{a})=0$. By our earlier definition, this means $\phi(a)=0$, so $a\in K$. Thus $\overline{a}=0$, so the kernel of $\overline{\phi}$ is trivial, and $\overline{\phi}$ is injective, as desired. □

Let's think about the case where I is a principal ideal: denote this ideal as $I = aR = (a)$. Then

$$R/I = \overline{R}$$

is a set of cosets, and since $\overline{a} = 0$, we can think of this as a ring obtained from R by forcing $a = 0$. This is similar to **adding a relation** in our ring R: for any $\overline{r} \in \overline{R}$, $\overline{ra} = 0$, and $\overline{b} + \overline{ra} = \overline{b}$. And in the principal ideal case, these are the only consequences of setting $a = 0$.

Example 73

Similarly, we can set $x^2 + 1 = 0$ in $R = \mathbb{R}[x]$. Then the ideal $(x^2 + 1)$ is the kernel of the homomorphism

$$\phi : \mathbb{R}[x] \to \mathbb{C}$$

that sends $f(x)$ to $f(i)$. Thus, by the First Isomorphism Theorem,

$$\mathbb{C} \cong \mathbb{R}[x]/(x^2 + 1).$$

Fact 74

This is one way to think about the notation for the Gaussian integers: after all,

$$\mathbb{Z}[i] \cong \mathbb{Z}[x]/(x^2 + 1).$$

However, most of the time, a quotient ring is not as recognizable as $\mathbb{Z}[i]$ or $\mathbb{C}$.

From this, we also get the concept of **adjoining a new element to a ring** R. Given R, we add a new element α with some additional properties. Basically, start with the polynomial ring $R[x]$ (for some new element x not in R), and then mod out by all of the relevant relations $f(x) = 0$. (This is how we constructed $\mathbb{C}$ above.)

Example 75

Take our ring $R = \mathbb{R}[t]$, and let's adjoin an element α which is a root of the polynomial

$$f = x^3 + tx + t \in R[x] = \mathbb{R}[t, x].$$

The solution is to force $x^3 + tx + t = 0$ in $R[x]$, and then we can define $\alpha = \overline{x}$ to be the residue of x in the quotient ring $R[x]/f$. How do we compute (for example, multiply elements) in $R[\alpha]$? Any polynomial $g(x)$ becomes an element $g(\alpha) \in R[\alpha]$. Since f is monic in x, we can write

$$g = fq + r,$$

where r has degree less than 3 in x. Since $f(\alpha) = 0$, we must have $g(\alpha) = r(\alpha)$, and therefore $r(x)$ is a quadratic polynomial with elements in R. This is as much as we can simplify, so $(1, x, x^2)$ is an R-basis for $R[x]/f$.

Fact 76

It's a bit more work to do this adjoining process when f is not monic, though – we won't talk much about it.

The next thing we can try to do is to adjoin a^{-1} to R. Then a^{-1} is a root of $ax - 1$, so

$$R[a^{-1}] = R[x]/(ax - 1).$$

So for example, $\mathbb{C}[t, t^{-1}] = \mathbb{C}[t, x]/(tx - 1)$, which just gives us polynomials in t and t^{-1}. This is called the **Laurent polynomial ring**, and it's used to build Laurent series in complex analysis. The elements of this ring look like

$$\sum_{j=-N}^{N} c_j t^j, c_j \in R.$$

Keep in mind, though, that this isn't necessarily a field! And in general, if we want to adjoin infinitely many things, this isn't the best way to do it.

10 February 27, 2019

We'll start with a motivating question: let R be a ring and F be a field. Under what conditions does there exist an injective $\phi : R \to F$?

Definition 77

An **(integral) domain** is a nonzero ring such that for any two nonzero elements $a, b \in R$, we have $ab \neq 0$ (there are no "zero divisors").

Examples include fields and the ring of integers.

Lemma 78 (Cancellation law)

If $a \neq 0$ and $ab = ac$ in a domain R, then $b = c$.

This is true because

$$ab = ac \implies a(b - c) = 0 \implies b - c = 0,$$

since $a \neq 0$ and it's not possible (in a domain) for two nonzero things to multiply to 0.

Proposition 79

If R is a domain, there exists an injective homomorphism $R \to F$ for some field F.

Proof. We consider the **fraction field** of the domain R, which is the set of equivalence classes

$$\left\{\frac{a}{b}, a, b \in R, b \neq 0\right\}$$

where $\frac{a}{b} = \frac{c}{d}$ if and only if $ad = bc$.

We can check silly things like associativity of addition in a fraction field, which is completely routine. It's important to note transitivity, though: if $\frac{a}{b} = \frac{c}{d}$ and $\frac{c}{d} = \frac{e}{f}$, then $\frac{a}{b} = \frac{e}{f}$. This is because

$$ad = bc,\ cf = de \implies bcf = bde \implies daf = deb \implies af = eb$$

by the cancellation law (note that we can only cancel by b and d since we know those are nonzero). Now just send any $a \in R$ to $\frac{a}{1}$ in the fraction field, and we have an injective map as desired. □

Proposition 80

Let R be a domain, F be a field, and let $\phi : R \to F$ be an injective homomorphism. Then ϕ extends to an injective homomorphism

$$\phi' : \text{fract}(R) \to F$$

where $\phi'\left(\frac{a}{b}\right) = \phi(b)^{-1}\phi(a)$.

We can check that this is compatible, and that's more of an annoying checklist. Here's something that's a little less simple: how do we deal with surjectivity?

Theorem 81 (Correspondence Theorem for rings)

Given a surjective homomorphism $\phi : R \to R'$ with kernel K, there is a bijective correspondence between ideals of R containing K and ideals of R'. Basically, any ideal I in R corresponds to its image $\phi(I)$, and any ideal J in R' corresponds to its preimage $\phi^{-1}(J)$.

This is important, because fields have very few ideals:

Lemma 82

The only ideals of a field F are the whole field and the zero ideal.

Proof. If there are no nonzero elements in an ideal, then we just have the zero ideal. Otherwise, if $a \neq 0 \in I$, $aa^{-1} = 1$ is in I, and then $I = F$ is the whole field. This is called the unit ideal, and it is often denoted (1). □

This means that there can't be that many ideals that contain the kernel. If $\phi : R \to F$ is a surjective homomorphism to a field F with kernel M, then R' is the field F in the Correspondence Theorem. Since the only ideals that contain M are M, R, there are no proper ideals containing M.

Definition 83

A **maximal ideal** M of a ring R is a proper ideal (not equal to R), such that there is no ideal I satisfying

$$M \subsetneq I \subsetneq R.$$

Basically, "we can't get any larger and still have something interesting." And the above logic gives us the following result:

Corollary 84

The kernel of a surjective homomorphism $\phi : R \to F$ is a maximal ideal.

Let's also show the "other direction:"

Lemma 85

Let M be a maximal ideal of a ring R. Then $F = R/M$ is a field.

Proof. We need to show that F is not the zero ring, and that inverses exist: $a \neq 0 \in F \implies a^{-1} \in F$.

The first condition is true because M is not R. For the second, it's sufficient to show that a field is just any ring with exactly two ideals, since the only ideals containing M are M and R, which corresponds to a field for R/M by the Correspondence theorem.

Fields are not the zero ring, because that only has one ideal. Now, let's say $a \in R$, and we take the ideal of multiples of a. Since $a \in (a)$, $(a) \neq (0)$, so $(a) = R \implies 1 \in (a)$. This means that $b = a^{-1}$ exists in R, so R must be a field. □

Example 86

Consider $R = \mathbb{Q}[x]$, where $\mathbb{Q}$ is the rationals. Take the principal ideal M generated by $f(x) = x^3 - 2$.

Why does this generate a maximal ideal? Every ideal of R is principal (this hasn't been proved in class yet, but it's the Division Algorithm again), so if $(f) \subset I$ and $I = (p)$ for some $p \in R$, then we can write $f = pq$ for some $q \in R$, meaning that p divides f. But f has no factors, since it would need to expand into a quadratic and a linear term, and $\sqrt[3]{2}$ is not rational. So $R/(f)$ must be a field. We can use the map

$$\phi : R = \mathbb{Q}[x] \to \mathbb{C}$$

which sends $x \to \sqrt[3]{2}$ (a rootroot of $x^3 - 2$), so that $\phi(x^3 - 2) = 0$. This means that the kernel of ϕ contains the ideal generated by $(x^3 - 2)$, which a maximal ideal, so the kernel is $(x^3 - 2)$. Thus, by the First Isomorphism Theorem, The image of ϕ is isomorphic to $R/(x^3 - 2)$: this is $\mathbb{Q}[\sqrt[3]{2}]$ with a $\mathbb{Q}$-basis of $(1, \alpha, \alpha^2)$.

It isn't immediately obvious that this is a field, but we can check with some computations that it is! Note that the other roots of $x^3 - 2$ are $\alpha\omega$ and $\alpha\omega^2$, where ω is a third root of unity. So now if we consider an alternate map $\phi' : \mathbb{Q}[x] \to \mathbb{C}$ where we send $x \to \omega\alpha$ instead, we still send $x^3 - 2 \to 0$, but we find now that the image of ϕ' is isomorphic to $R/(x^3 - 2)$ as well.

So $\mathbb{Q}[\alpha']$ and $\mathbb{Q}[\alpha]$ are isomorphic. One is a subfield of the reals, and the other is not, but the two rings are still isomorphic! So we can't tell just by looking at the structure of the ring whether the element α that we adjoin is real.

11 March 1, 2019

The main topic of today is the Nullstellensatz.

Fact 87

In German, "nullstellensatz" means "zero place theorem."

Recall that a **maximal ideal** M of R is an ideal $M < R$ such that there exist no ideals I with $M < I < R$. We found last time that M being maximal is equivalent to R/M being a field. In other words, if $\phi : R \to F$ is a surjective ring homomorphism, then F is a field if and only if the kernel of ϕ is a maximal ideal. This is because a field has only two ideals: 0 and the whole field.

Example 88

Let $R = \mathbb{C}[x]$. Since every ideal of R is principal, for any ideals I, there exists an $f \in I$ such that every $g \in I$, $g = fq$. Then I is denoted (f), Rf, or fR: it's the principal ideal generated by f.

In particular, we can use the division algorithm to find the polynomial with minimal degree in I: that's the one that generates I.

Question 89. *When is an ideal (g) maximal in $\mathbb{C}[x]$?*

If another proper ideal is larger – that is, $(f) \supset (g)$ – then $g \in (f) \implies g = fq$. So g is maximal if there is no such f, meaning g **should be irreducible in $\mathbb{C}[x]$.** But this happens if and only if the degree is 1 (by the Fundamental Theorem of Algebra), so we have the following result:

Corollary 90
The maximal ideals of $R = \mathbb{C}[x]$ correspond to elements $a \in \mathbb{C}$, and they can be written as $R/(x-a)$.

We can think of such maximal ideals as corresponding to the kernels of homomorphisms

$$\mathbb{C}[x] \to \mathbb{C}, x \to a.$$

Now let's move on, and let $R = \mathbb{C}[x_1, x_2, \cdots, x_n]$ be the ring of polynomials in more variables. Now ideals look more complicated: we can't say very much about them at this point. We should remember that $\alpha_1, \alpha_2, \cdots, \alpha_n$ **generate** an ideal I if

$$I = \left\{ \sum r_i \alpha_i \mid r_i \in R \right\}.$$

Basically, we take all R-linear combinations of our generating set. Unfortunately, there's no bound for the number of generators of an ideal for $R = \mathbb{C}[x_1, \cdots, x_n]$. Sometimes, we can even have an infinite number of generators! But here's something useful to help us:

Theorem 91 (Hilbert Basis Theorem)
Every ideal I of the polynomial ring $\mathbb{C}[x_1, \cdots, x_n]$ can be generated by finitely many elements.

(We'll prove this later on in the class in an alternate form.) With that, it's time for a big result which was published along with the Hilbert Basis Theorem in 1895:

Theorem 92 (Hilbert Nullstellensatz)
The **maximal ideals** of a ring $R = \mathbb{C}[x_1, \cdots, x_n]$ correspond to points $a = (a_1, \cdots, a_n) \in \mathbb{C}^n$.

This is the analog of the 1-dimensional theorem. Basically, given a point $a \in \mathbb{C}^n$, the maximal ideal is the kernel of the homomorphism

$$\phi : \mathbb{C}[x_1, \cdots, x_n] \to \mathbb{C}, \ x_i \to a_i.$$

This homomorphism sends any polynomial $g(x_1, \cdots, x_n)$ to $g(a_1, \cdots, a_n)$. Indeed, the kernel here is a maximal ideal (if some other element f were in the ideal, we could use the division algorithm to get a constant, and that's either 0 or it generates the whole ring).

All such ideals are generated by $\{x_1 - a_1, x_2 - a_2, \cdots, x_n - a_n\}$. This means that given a polynomial $g(x)$ with $g(a) = 0$, we can write (not necessarily uniquely)

$$g = (x_1 - a_1)q_1 + \cdots + (x_n - a_n)q_n.$$

One way to show this is to do a Taylor expansion:

Fact 93 (Taylor expansion in multiple variables)

If we've taken 18.02, we know that we can write polynomials as

$$g(x) = g(a) + \frac{1}{1!}\sum(x_i - a_i)\frac{\partial g}{\partial x_i}(a) + \frac{1}{2!}\sum(x_i - a_i)(x_j - a_j)\frac{\partial^2 g}{\partial x_i \partial x_j}(a) + \cdots$$

and now put each term into one of the $(x_i - a_i)$s as we wish.

So this gives us something we can say about any ideal:

Corollary 94

Every ideal I of $\mathbb{C}[x_1, \cdots, x_n]$ with $I < R$ is contained in some maximal ideal.

Proof. If I is maximal, we're good. Otherwise, $I < I_1 < R$; if I_1 is maximal, we're also good. Otherwise, $I_1 < I_2 < R$, and keep going. This gives us a chain

$$I = I_0 < I_1 < I_2 < \cdots$$

To finish, we want to show that this chain can't be infinite:

Lemma 95

Given a chain $I_0 \subset I_1 \subset I_2 \subset \cdots$ in a ring R, the union $I = \bigcup I_n$ is an ideal.

We just need to show the two closure assumptions. If a, b are in I, the union of all I_n, then $a \in I_x$ and $b \in I_y$ for some x, y. Then both are in $I_{\max(x,y)}$, so $a + b$ is in $I_{\max(x,y)}$. $ra \in I_{\max r,y}$ as well, so both of these are also in I.

So if our chain is infinite, $J = \bigcup I_n$ is an ideal. By the Hilbert Basis Theorem, J is generated by a finite number of elements. Let's call them $b_1, \cdots, b_r$. Then for some n large enough, all $b_i \in I_n$: now $J \subset I_n$, so $J = I_n$. All I_ns are proper (not R), so J is also proper. This means our chain does stop eventually, and we do have a maximal ideal at the end of our chain. □

Example 96

Take our ring to be $\mathbb{C}[x, y]$. Let I be an ideal generated by $x^2 + y^2 - 1, xy - 1$, and $y - x^3$.

We claim $I = R$, and we'll show this by showing that it is not contained in a maximal ideal. All maximal ideals are kernels of substitutions, so if they were in a maximal ideal, there would be a point such that all three polynomials would evaluate to zero. This requires $xy = 1, y = x^3 \implies x^4 = 1$, and $x^2 + y^2 = 1 \implies x^4 + 1 = x^2$ (since $x^2y^2 = 1$). This implies $x^2 = 2$, but $x^4 = 1$. So there is no point that is a zero of all three polynomials: they must therefore generate the unit ideal.

Let's go over the proof of the Hilbert Nullstellensatz (Theorem 92):

Proof. Start with an unknown maximal ideal M of $R = \mathbb{C}[x_1, \cdots, x_n]$. Then $F = R/M$ is a field, and we have a surjective homomorphism $\pi : R \to F$. Restrict π to the subring $\mathbb{C}[x_1]$; call this map ϕ_1. All ideals in $\mathbb{C}[x_1]$ are principal, and ϕ_1 sends these elements into a field F, so $\ker \phi_1$ must look like $(x_1 - a_1)$ (the ideal generated by a linear polynomial) or just the zero ideal. The first case is good, since us being able to do this for each x_i gives our generators. So we have to rule out the case where $\ker \phi_1$ is trivial.

If ϕ_1 has trivial kernel, then ϕ_1 is injective. We can therefore extend ϕ_1 to the fraction field $\mathbb{C}(x_1) \to F$ (the round bracket notation means **fractions** of polynomials). As a complex vector space, F is spanned by a countable set, since

R is spanned by monomials $x_1^{e_1} \cdots x_n^{e_n}$ and we have a surjective map. But the field of fractions $\mathbb{C}(x_1)$ contains an uncountable independent set

$$\left\{ \frac{1}{x_1 - c}, c \in \mathbb{C}. \right\},$$

because it's not possible to have nonzero b_i with

$$\sum_i \frac{b_i}{x_i - c_i} = 0$$

if we look locally near c_i, where there's a pole. (It's important to note that any linear combination only uses a finite number of the spanning elements.) If a vector space is spanned by a countable set, every independent set is countable or finite, which contradicts $\mathbb{C}(x_1)$ being a subring of F (since we assumed injectivity). Thus all ϕ_i have kernel equal to $(x_i - a_i)$, and thus we've found generators for our maximal ideal. □

We have a quiz on Monday in Walker. Our TA will be at a meeting, so it is unclear that quizzes can be given back on Wednesday.

12 March 6, 2019

(The exam 1 average was above 80, so we all get cookies.)

Our new topic for this class will be **factoring and irreducibility**. Let's start with polynomials $R = \mathbb{F}[x]$, where $\mathbb{F}$ is a field. Sometimes, like in the rationals, it's possible to have irreducible polynomials of all degrees.

We usually have two main tools for factoring: we can often use **division with remainder**, and we can use the fact that R is a **principal ideal domain**.

In general, we want to work with monic polynomials

$$f(x) = x^n + a_{n-1}x^{n-1} + \cdots + a_0, a_i \in F,$$

and we can turn f into this form by multiplying by the inverse of the leading coefficient.

Definition 97

Assume f, p are monic. An **irreducible polynomial** f is a polynomial where $f \neq 1$ and it can't be factored as $f = gh$ with g, h not units (in this case, this means g, h have degree at least 1). A polynomial p is **prime** if $p \neq 1$ and for all f, g, if p divides fg, then p divides either f or g.

Lemma 98

Every irreducible element in a principal ideal domain is a prime element.

Note that this result is not true for all rings:

Example 99

Let $R = \mathbb{Z}[\delta]$, where $\delta = \sqrt{-5}$. Note that 2 divides $(1 + \delta)(1 - \delta) = 6$, but 2 does not divide $1 + \delta$ or $1 - \delta$. However, 2 is irreducible, because 2 has smaller absolute value than all other elements besides ± 1.

Proof of lemma. Let p be an irreducible element, and let's say that p divides fg but p does not divide f. Let I be the ideal generated by p and f:

$$I = (p, f) = \{rp + sf \mid r, s \in R\}.$$

Note that $I > (p)$ (this is strictly greater since f is not in the ideal generated by p), and since I is principal, $I = (d)$ for some element d. So $(d) > (p)$, meaning that $d|p$. But p was irreducible, so either $d = 1$ or $d = p$.

We can't have $d = p$, since we assumed (d) was strictly larger than (p), so $d = (1)$. This means $1 = rp + sf$ for some $r, s \in R$. But then

$$1 = rp + sf \implies g = r(p)g + s(fg),$$

and p divides both terms on the right side, so p divides g. □

Corollary 100

For any field $\mathbb{F}$, all monic polynomials $f \neq 1$ in $\mathbb{F}[x]$ are a product of irreducible elements, with uniqueness up to ordering.

We can also make a similar argument about unique factorization in $\mathbb{Z}[i]$, except for unit factors like $\pm 1, \pm i$. Let's be a bit more precise about that:

Definition 101

A **unit** in a ring R is an element that is invertible. Two elements $a, b \in R$ are **associate** if we can write $b = ua$ for some unit $u \in R$.

Example 102

In $\mathbb{Z}[i]$, $1 + 3i$ is associate to $-1 - 3i$, $-3 + i$, and $3 - i$.

So the idea is that two factorizations of a Gaussian integer can look different: we can multiply one factor by i and another by $-i$.

Example 103

Let's factor $1 + 3i$ in the Gaussian integers.

Start by multiplying it by its conjugate to get a **norm**

$$(1 + 3i)(1 - 3i) = 10 = 2 \cdot 5.$$

Now we can factor the two terms on the right: 2 factors as $(1 + i)(1 - i)$, and 5 factors as $(2 + i)(2 - i)$. So

$$(1 + 3i)(1 - 3i) = (1 + i)(1 - i)(2 + i)(2 - i).$$

Now $1 + 3i$ is "probably" a product of some of them: we'll try a few different things. $(1 + i)(1 + 2i) = -1 + 3i$, which doesn't work. But $(1 - i)(2 + i) = 3 - i$ is an associate of $1 + 3i$, so we've found our factorization.

Theorem 104

For any principal ideal domain R, all nonzero, non-unit $a \in R$ are products of prime (or irreducible) elements

$$a = p_1 p_2 \cdots p_k.$$

This is unique in the sense that if $a = q_1 q_2 \cdots q_l$, we can reorder the q_is, and then all p_is are associates of corresponding q_is.

One noteworthy idea here: factorization will always terminate in a principal ideal domain.

Proof. The idea is that if we start with a nonzero and not a unit, then either it is irreducible (in which case we're done), or $a = a_1 b_1$. In the second case, we know that

$$(a) < (a_1),$$

since a doesn't divide a_1 if b is a unit. Then either a_1 is irreducible (and we're done) or $a_1 = a_2 b_2$, and so on. This creates a similar infinite chain

$$(a) < (a_1) < (a_2) < \cdots$$

and we can't keep doing this, because the union of an increasing family of ideals is an ideal by Lemma 95.

So $J = \bigcup (a_i)$ is an ideal, and it is principal (because R is a principal ideal domain), so $J = (d)$ for some d. d is in the union of the ideals, so it is in some (a_i). Then $(a_i) \subset (d) \subset (a_i)$, meaning a_i is the largest our ideals get. And thus we can't get an infinite chain, and therefore we can indeed write a as a product of prime elements.

Showing uniqueness comes from each p_i needing to divide a q_j and vice versa. □

But the converse is not true: **not all rings with unique factorization are principal ideal domains**. Let's stop working with monic polynomials now and look in $\mathbb{Z}[x]$. An example of a factorization in $\mathbb{Z}[x]$ is

$$f(x) = 2(x^2 + 1)(5x - 3).$$

There are two complications compared to the $\mathbb{F}[x]$ case: sometimes, we can't always have monic polynomials, and we may have a constant leading term as well.

What tools do we have to deal with this? First of all, $\mathbb{Z}[x]$ is a subset of $\mathbb{Q}[x]$, the ring of polynomials with rational coefficients, so we can use some properties of $\mathbb{Q}[x]$ to help us out. Also, we have the ring homomorphism

$$\pi_p : \mathbb{Z}[x] \to \mathbb{F}_p[x]$$

which sends $f(x) \to \pi_p(f) = \overline{f}(x)$. If we have a polynomial that we can't factor in $\mathbb{Z}[x]$, we want to show it is irreducible: perhaps we can show that it is irreducible in $\mathbb{F}_p[x]$ instead, which is easier since there are only finitely many of any given degree! In particular, we can try all pairs of polynomials and see if any of them multiply to our polynomial $\overline{f}$.

Example 105

Consider the polynomials in $\mathbb{F}_2[x]$. The first few are

$$1, x, x+1, x^2, x^2+x, x^2+1, x^2+x+1, x^3, \cdots$$

To find the irreducible polynomials, we can do something similar to the Sieve of Eratosthenes. x^2, x^2+x are factorable, and $x^2+1 = x^2-1$ are factorable as well. x^2+x+1 has no roots, so it is irreducible. For cubics, their factors must be a linear times a quadratic polynomial, so we just check all polynomials of the form x^3+ax^2+bx+1 (since we don't want x to be a factor) and see if $x+1$ is a root: we see that x^3+x+1, x^3+x^2+1 are irreducible.

Example 106

A polynomial like $x^3+2x^2+7x+15$ is irreducible in the integers, because it is irreducible in $\mathbb{F}_2[x]$.

Definition 107

A polynomial $f(x) \in \mathbb{Z}[x]$ is **primitive** if the greatest common divisor of its coefficients is 1.

In other words, no prime divides all coefficients of f, so $\pi_p(f) \neq 0$ for all primes p.

Theorem 108 (Gauss' Lemma)

If f, g are primitive polynomials in $\mathbb{Z}[x]$, then fg is primitive.

Proof. It suffices to show that $\pi_p(fg) \neq 0$ for all primes p. Since π_p is a homomorphism, $\pi_(fg) = \pi_p(f)\pi_p(g)$, and both f and g have a nonzero leading term, so their product is not zero. (Alternatively, $\mathbb{F}[x]$ is a domain for any field $\mathbb{F}$.) Thus $\pi_p(fg) \neq 0$. □

13 March 8, 2019

We'll be talking more about factoring in the ring of polynomials $\mathbb{Z}[x]$ today. Recall that we have two main tools: one is to view polynomials in $\mathbb{Z}[x]$ as polynomials in $\mathbb{Q}[x]$, which is easier because $\mathbb{Q}[x]$ is a principal ideal domain. Alternatively, we can look at the homomorphism $\pi_p : \mathbb{Z}[x] \to \mathbb{F}_p[x]$, which reduces the coefficients mod p.

Last time, we defined a **primitive polynomial** $f(x) \in \mathbb{Z}[x]$ to be a nonconstant polynomial with integer coefficients, where the greatest common divisor of its coefficients is 1. (The leading coefficient has to be positive.) Theorem 108 (Gauss' Lemma) tells us that if f, g are primitive polynomials, then fg is primitive as well. This has the following useful corollray:

Corollary 109

Let f be a primitive polynomial. If $g = fq$, and q has rational coefficients, then $q \in \mathbb{Z}[x]$.

So if f is a primitive polynomial, and f divides g in $\mathbb{Q}[x]$, then f divides g in $\mathbb{Z}[x]$ as well!

Proof. Clear the denominators, so $q_1 = dq$, where d is the constant with smallest absolute value needed to make q_1 have integer coefficients. Note that this means q_1 is primitive.

Now $dg = dfq = fq_1$, and by the Gauss Lemma, fq_1 is primitive. But now the only way for $fq_1 = dg$ to be primitive is if the constant factor d is ± 1 – otherwise, g's coefficients would have a common factor. Thus q must have started out with integer coefficients as well. □

Corollary 110

If g is irreducible in $\mathbb{Z}[x]$, then q is irreducible in $\mathbb{Q}[x]$ as well.

Proof. This looks very similar to the above argument. We show the contrapositive: if $g = fq$, where $f, q \in \mathbb{Q}[x]$, then g factors in $\mathbb{Z}[x]$ as well. Move the constants around so that f is primitive (and has integer coefficients). Now pick the constant c such that $q_1 = cq$ (like above). Now $cg = cfq = fq_1$, but f and q_1 are primitive so $c = \pm 1$, and thus q must have had integer coefficients as well. □

Corollary 111

The irreducible elements of $\mathbb{Z}[x]$ are $\pm$ integer primes and $\pm$ irreducible polynomials. Thus, **the irreducible elements of $\mathbb{Z}[x]$ are prime elements.**

Remember that f is **prime** if $f|gh$ implies that $f|g$ or $f|h$.

Proof. Suppose f is an irreducible polynomial. Then if $f|gh$ in $\mathbb{Z}[x]$, then $f|gh$ in $\mathbb{Q}[x]$. But irreducibles are primes in a principal ideal domain, so $f|g$ or $f|h$ in $\mathbb{Q}[x]$, and therefore $f|g$ or $f|h$ in $\mathbb{Z}[x]$, and thus f is prime. The converse (that integer primes and irreducible polynomials are irreducible) is easy to show. □

With all of this, we've shown the following result (essentially be swtiching back and forth between $\mathbb{Q}[x]$ and $\mathbb{Z}[x]$):

Theorem 112

The polynomial ring $\mathbb{Z}[x]$ has unique factorization into prime polynomials. Thus, any $g(x)$ can be written as

$$g = \pm p_1 \cdots p_k f_1 \cdots f_k$$

where p_i are integer primes and f_j are primitive polynomials.

Let's shift to our other tool now: using the homomorphism π_p. If f factors in $\mathbb{Z}[x]$, it will factor in $\mathbb{F}_p[x]$ as well: take $f = gh$ to $\overline{f} = \overline{g}\overline{h}$. This breaks down if the leading coefficient of f is divisible by p though, and we should always pick a good prime so that we don't have $\overline{g} = 1$.

Example 113

Take $p = 2$ – the irreducible polynomials in $\mathbb{F}_2[x]$ with degree at most 4 are $x, x+1, x^2+x+1, x^3+x+1, x^3+x^2+1, x^4+x+1, x^4+x^3+1, x^4+x^3+x^2+x+1$.

Consider the polynomial $f(x) = 5x^5 + 4x^4 + 3x^3 + 2x^2 + x - 1$. Then $\overline{f} = x^5 + x^3 + x + 1$. Notice that $x+1$ is a factor of this, since $x = -1$ is a root of $\overline{f}$: therefore, our polynomial factors as

$$\overline{f} = (x+1)(x^4 + x^3 + 1),$$

where we can check that the second factor is irreducible. So if f factors, it must look like the product of a linear and quartic polynomial. (Note that $\mathbb{F}_2[x]$ has unique factorization.) The linear polynomial must look like $(ax + b)$, where b is a factor of 1 and a is a factor of 5 (by the Rational Root Theorem). So now we just check $\pm 1, \pm\frac{1}{5}$: none of these are roots of f, so our polynomial is irreducible in $\mathbb{Z}[x]$.

Fact 114

Computers can do this in "zero" time. But it's very painful to try to do it mod powers of 2. It's actually a linear algebra problem over $\mathbb{F}_2$, but this is generally not suitable for doing in class (we shouldn't try it).

Now we'll demonstrate another useful tool:

Example 115

Consider $f(x) = x^5 + 3x^4 - 6x^3 + 9x^2 - 3$. Is this irreducible in $\mathbb{Z}[x]$

Take prime $p = 3$. Our polynomial looks like $\overline{f} = x^5$, so $\overline{f} = \overline{g}\overline{h} = x^5$. Let's say g has degree 2 and h has degree 3: note that $\overline{g} = x^2$ and $\overline{h} = x^3$ works, and by unique factorization this is the only way to multiply a quadratic and cubic polynomial. So now

$$g = x^2 + b_1x + b_0, h = x^3 + c_2x^2 + c_1x + c_0,$$

where $3|b_0, c_0$. But this implies $9|b_0c_0 = -3$, which is not true! So this is not a valid factorization for the original $f(x)$. We can repeat this argument in general for any degrees g, h: this means f is irreducible.

Here's a way to state this in more generality:

Theorem 116 (Eisenstein's criterion)

Suppose we have a polynomial $f \in \mathbb{Z}[x]$ of the form

$$f(x) = a_nx^n + \cdots + a_0.$$

If a prime p divides all coefficients except the leading a_n, and p^2 does not divide a_0, then f is irreducible.

Proof. Repeat the process we used above in the general case. □

We'll finish by talking about cyclotomic polynomials. Consider the pth roots of unity ζ^k, where $\zeta = e^{2\pi i/p}$. All of these are roots of the polynomial $x^p - 1$, but there's the trivial root $x = 1$. So we want the polynomial with roots $\zeta, \zeta^2, \cdots, \zeta^{p-1}$:

$$\frac{x^p - 1}{x - 1} = x^{p-1} + \cdots + x + 1.$$

Definition 117

This polynomial $\frac{x^p-1}{x-1}$ is called the pth **cyclotomic polynomial**. In general, the nth cyclotomic polynomial $\Phi_n(x)$ comes from multiplying $(x - \zeta_n^k)$ for all **primitive** nth roots of unity.

Theorem 118

The polynomial $x^{p-1} + \cdots + x + 1$ is irreducible in $\mathbb{Q}[x]$ for any prime p.

Proof. Make a change of variables: let $x = y + 1$. Then

$$f(x) = f(y + 1) = \frac{(y + 1)^p - 1}{(y + 1) - 1}$$

can be expanded by the binomial theorem:

$$(y+1)^p = y^p + \binom{p}{1}y^{p-1} + \cdots + \binom{p}{p-1}y + 1,$$

and all coefficients except the first and last have exactly 1 factor of p in them. Subtract 1 and divide by y, and now

$$f(y+1) = y^{p-1} + \binom{p}{1}y^{p-2} + \cdots + \binom{p}{p-2}x + \binom{p}{p-1}$$

must be irreducible by Eisenstein's criterion. □

Corollary 119

Since ζ_p is a root of $\Phi_p(x)$,

$$f(\zeta) = \zeta^{p-1} + \cdots + \zeta + 1 = 0$$

is the **only linear (rational) relation among powers of** ζ. In other words, $\{\zeta, \zeta^2, \cdots, \zeta^{p-1}\}$ are independent over $\mathbb{Q}$.

Proof. Consider the map $\phi : \mathbb{Q}[x] \to \mathbb{C}$ that sends $x \to \zeta$. By definition, the cyclotomic polynomial Φ_p is in the kernel. $\ker\phi$ is a principal ideal, and it is generated by the minimal polynomial with ζ as a root. This is Φ_p, because it is irreducible and has ζ as a root! So the kernel of ϕ is generated by f, and by the first isomorphism theorem, $\mathbb{Q}[x]/(f)$ is isomorphic to the image $\mathbb{Q}[\zeta]$. Since f has degree $p-1$, the image must have dimension $p-1$, and the residues $\overline{1}, \overline{\zeta}, \cdots \overline{\zeta^{p-1}}$ form a basis. □

14 March 11, 2019

We're going to shift topic to doing arithmetic in **imaginary quadratic fields**. We'll start with $R = \mathbb{Z}[i]$, which is easy to deal with because we have a principal ideal domain.

Remember that the explicit definition of the Gaussian integers $\mathbb{Z}[i]$ is

$$\{a + bi \mid a, b \in \mathbb{Z}\}.$$

Definition 120

The **norm** of a Gaussian integer $\alpha = a + bi$ is

$$N(\alpha) = \overline{\alpha}\alpha = |\alpha^2| = a^2 + b^2.$$

This has two useful properties: $N(\alpha)$ is **multiplicative**, and it is always a **positive integer**. Since every ideal of $\mathbb{Z}[i]$ is principal, irreducible elements are prime elements. Remember the definitions here: if α is **irreducible**, then $\alpha \neq 0$, α is not a unit, and it has no proper factors. Meanwhile, π is **prime** if $\pi | ab \implies \pi | a$ or $\pi | b$. Note that the ideal generated by a prime element π, πR, is maximal, so $\overline{R} = R/(\pi R)$ is a field. And $\mathbb{Z}[i]$ has unique factorization into Gauss primes (primes in $\mathbb{Z}[i]$), up to associates, which differ by one of the units $\{\pm 1, \pm i\}$. So, for example, unique factorization treats

$$(1+i)(1+2i) = (1-i)(-2+i) = -1+3i$$

as equivalent factoring.

Proposition 121

Let π be a Gauss prime. Then $N(\pi) = a^2 + b^2$ is either a prime integer p or the square of a prime integer p^2.

Additionally, this implies that if $N(\pi) = \overline{\pi}\pi = p^2$, then π and p are associates (by unique factorization).

Proof. We know that $\overline{\pi}\pi$ is an integer, so we can factor it in $\mathbb{Z}$ as

$$\overline{\pi}\pi = p_1 \cdots p_k.$$

But this is also a valid equation in the Gaussian integers. Since π is a Gauss prime, so is $\overline{\pi}$ (just take conjugates), and now $p_1 \cdots p_k$ can only have at most two terms! So either $\overline{\pi}\pi = p$ for some prime or $\overline{\pi}\pi = p^2$.

In the last case, p_1 and p_2 must be associates of π and $\overline{\pi}$. Say that p_1 and π are associates. Then p_2 and $\overline{\pi}$ are associates, but we also know that $\overline{\pi}$ and $\overline{p_1} = p_1$ are associates. So p_1 and p_2 are associates, which means $p_1 = p_2$ (in order for their product to be a positive integer), and this must just be p^2, as desired. □

Corollary 122

If p is an integer prime, we can factor it in $\mathbb{Z}[i]$, and it factors either as p or $\overline{\pi}\pi$, where π is a Gauss prime.

Proof. Say we have an integer prime p. Factor it into

$$p = \pi_1 \cdots \pi_k.$$

Take the norm of both sides: then

$$p^2 = \overline{p}p = (\overline{\pi_1}\pi_1)(\overline{\pi_2}\pi_2)\cdots(\overline{\pi_k}\pi_k).$$

Each parenthetical term is an integer, and it is more than 1. By unique factorization of integers, these sides must be the same, so $k \leq 2$. If $k = 1$, we just have

$$\overline{p}p = \overline{\pi_1}\pi_1,$$

and by uniqueness in the Gaussian integers, this means we must have $p = \pi_1$, meaning p is just a Gauss prime (or its associate). Meanwhile, if $k = 2$,

$$p^2 = \overline{p}p = (\overline{\pi_1}\pi_1)(\overline{\pi_2}\pi_2)$$

and similarly we must have $p = \overline{\pi_1}\pi_1 = \overline{\pi_2}\pi_2$. □

Example 123

We know that $2 = (1+i)(1-i)$, so $1+i, 1-i$ are both Gauss primes. We can't write 3 as a product of two Gaussian integers with smaller norm, so 3 is a Gauss prime. Finally, $5 = (2+i)(2-i)$, and $2+i, 2-i$ are both Gauss primes.

We say that 2 and 5 **split** in $\mathbb{Z}[i]$, but 3 **remains prime**.

Theorem 124

Let $p > 2$ be an odd prime. The following are equivalent:

1. p is a Gauss prime.
2. $x^2 + 1$ is irreducible in $\mathbb{F}_p(x)$: that is, -1 is not a square mod p.
3. p is not a sum of two integer squares.
4. p is 3 mod 4.

Proof that (1) and (2) are equivalent. Remember that given a Gauss prime π, (πR) is a maximal ideal, so $\overline{R} = R/(\pi R)$ is a field.

We can think of the Gaussian integers as the image

$$\phi : \mathbb{Z}[x] \to \mathbb{Z}[i] = R$$

under the substitution $x \to i$. Compose this with a map from R to $\overline{R} = R/pR$: this first kills $x^2 + 1$, and then it kills p.

Alternatively, we can start with $\mathbb{Z}[x]$, kill p to get to $S = \mathbb{F}_p[x]$, and then we can substitute i (by modding out by the polynomial $x^2 + 1$) to get $\overline{S} = S/((x^2 + 1)S)$. This kills p first and then $x^2 + 1$.

But notice that $\overline{R}$ and $\overline{S}$ are **isomorphic**, so they are fields for the same values of p. So $\overline{R}$ is a field if and only if p is a Gauss prime, and because $\mathbb{F}_p[x]$ is a principal ideal domain, $\overline{S}$ is a field if and only if $x^2 + 1$ is an irreducible polynomial in $\mathbb{F}_p[x]$. This is exactly the correspondence that we want. □

Let's look a bit more carefully at the structure here and justify the idea of $\overline{R}$ and $\overline{S}$ being isomorphic:

Lemma 125

Let $u : R \to R'$ and $v : R' \to R''$ be surjective homomorphisms, and let $w = v \circ u$. (Assume we have a principal ideal domain to make notation easier.) Let $\ker u = xR$ and $\ker v = y'R'$, and let $y \in R$ be the element such that $uy = y'$ (which exists by surjectivity). Then

$$\ker w = (x, y)R.$$

Then by the first isomorphism theorem, we know that

$$R'' \cong R/\ker w,$$

so our logic above was valid (because $(x, y)R$ and $(y, x)R$ are the same thing).

Proof. Let $a \in \ker w$. Then $ua = a'$ must be in the kernel of v, so $a' = r'y'$ for some $r' \in R'$. Since u is surjective, choose $r \in R$ such that $ur = r'$. Then let $b = a - ry$; notice that

$$ub = ua - ury = a' - r'y' = 0,$$

so $ub \in \ker u$, meaning that $b = sx$ for some $s \in R$. Therefore

$$a = b + ry = sx + ry,$$

as desired – we've shown that $a \in (x, y)R$. (And clearly x, y are both in the kernel, so we do need to include both.) □

Let's finish the proof of Theorem 124:

Proof, continued. Let's show that the first and third conditions are equivalent. If p is a sum of two squares, we can write it as

$$p = a^2 + b^2 = (a + bi)(a - bi),$$

meaning it is not a Gauss prime. And this logic works in reverse – if p is not a sum of two squares, we cannot factor it in the Gaussian integers.

Finally, to show that the second and fourth conditions are equivalent, $x^2 + 1$ is irreducible if and only if -1 is not a square mod p. We map the multiplicative group to itself: consider the map $\phi : \mathbb{F}_p^\times \to \mathbb{F}_p^\times$ sending $a \to a^2$. Since we have an abelian group here, this is a homomorphism, and the kernel is $\{\pm 1\}$. Thus the order of the image H is $\frac{p-1}{2}$: if p is 3 mod 4, then the order of H is odd, so there are no elements of order 2. But if p is 1 mod 4, then there is an element of order 2, which must be -1 (this is the only one)!

So -1 is in the image if and only if it is a square, which happens if and only if p is 1 mod 4. □

The next case we can look at is $\mathbb{Z}[\sqrt{-5}]$ – the main difference is that ideals are no longer principal.

15 March 13, 2019

Let's review some concepts about $R = \mathbb{Z}[i]$. This is a principal ideal domain, and therefore it's also a unique factorization domain.

Take A to be an ideal of R which is not the zero ideal. Let $\alpha \in A$ be a nonzero element with minimal norm: we can show that A is the principal ideal αR with the Division Algorithm. What does this ideal look like? We get translates of α and αi, so we have a square grid of sidelength $|\alpha|$.

Fact 126

The idea is that any element $\beta \in A$ can't be inside a square, as it'd be too close to one of the adjacent α-points. We can see this by drawing circles of radius $|\alpha|$ from opposite corners of a square: this covers the whole square, so β must be a multiple of α.

We're going to use this same reasoning in another case now.

Proposition 127

$R = \mathbb{Z}[\sqrt{-2}]$ is a unique factorization domain.

Proof. Take any ideal A. Like before, choose $\alpha \neq 0$ with minimal norm. Now the ideal generated by α looks a little different: we have a rectangular grid with side lengths in the ratio $1 : \sqrt{2}$ in the complex plane.

But the same tiling argument works here! Let's say β not in $A = \alpha R$ is in the ideal: we know β can't be within $|\alpha|$ of any of the vertices of the rectangle. Again, the union of the circles completely covers the rectangle. □

So $\mathbb{Z}[\sqrt{-2}]$ is both a PID (principal ideal domain) and a UFD (unique factorization domain). Notice that the only units here are ± 1.

We're going to defer $\mathbb{Z}[\sqrt{-3}]$ for now.

Example 128

Now, let's try $\mathbb{Z}[\sqrt{-5}]$. Does the same argument work to show that this is a principal ideal domain?

If we try to make the same argument by starting with an $\alpha \in A$ with minimal length, we have a rectangle that is $\sqrt{5}\alpha$ by α. But this time, our circles don't completely cover the rectangle, and we're in trouble.

Instead, the idea is to consider $\frac{1}{2}\alpha R$. If $\beta \in A$ is close to $\frac{1+\sqrt{-5}}{2}\alpha$, then 2β would be too close to $(1+\sqrt{-5})\alpha$. So we also can't be within $\frac{\alpha}{2}$ of any of the half-lattice points.

But we've missed one detail: β can't be close to $\frac{1+\sqrt{-5}}{2}\alpha$, but it can be equal to it! So if we have an ideal A strictly larger than αR, where α is the point with minimal magnitude, then A must contain one of the points $\frac{1}{2}\alpha\sqrt{-5}$, $\frac{1}{2}\alpha(1+\sqrt{-5})$, $\frac{1}{2}\alpha(2+\sqrt{-5})$.

This first option is impossible: if $\gamma = \frac{1}{2}\sqrt{-5}\alpha$ is in A, then $\sqrt{-5}\gamma = -\frac{5}{2}\alpha$ is in A, meaning $\frac{1}{2}\alpha$ is in A, contradicting the minimality of α. Similarly, $\gamma = \frac{1}{2}\alpha(2+\sqrt{-5})$ isn't allowed, either. This yields our result:

Proposition 129

Let $R = \mathbb{Z}[\sqrt{-5}]$, and let A be a nonzero ideal. If α is the nonzero element in A with smallest magnitude, then either $A = \alpha R$ is a prime ideal, or $A = (\alpha, \beta)R$, where

$$\beta = \frac{1}{2}(1+\sqrt{-5})\alpha.$$

(Sometimes, β may not be in the ring, so we can't have $(\alpha, \beta)R$ as an ideal in that case either.)

Recall the issue with unique factorization in this ring that we explored a few days ago: if $\delta = \sqrt{-5}$, then

$$6 = 2 \cdot 3 = (1+\delta)(1-\delta).$$

We took ideals to rescue unique factorization by defining the two ideals

$$A = (2, 1+\delta), B = (3, 1+\delta).$$

Let's review the logic there: to take the product ideal AB, notice that

$$AB = (6, 3+3\delta, 2+2\delta, (1+\delta)^2)$$

has $(1+\delta)R$ as a subideal, and all of the elements in AB are in the principal ideal $(1+\delta)R$. So AB is generated by $1+\delta$.

From there, noting that $\overline{A} = (2, 1-\delta)$ and $\overline{B} = (3, 1-\delta)$, we found that $\overline{A}A = 2R$, $\overline{B}B = 3R$, and $\overline{AB} = (1-\delta)R$. This meant that our expression $6 = 2 \cdot 3 = (1+\delta)(1-\delta)$ could be rewritten as

$$(\overline{A}A)(\overline{B}B) = (AB)(\overline{AB}).$$

Fact 130

Note the difference between a lattice with basis $(2, 1+\delta)$ (which is linear combinations with **integer coefficients**) and an ideal with generators $(2, 1+\delta)$ (which is linear combinations with **coefficients in the ring**).

We claim that $A = (2, 1+\delta)$ has a lattice basis. To show this, we just need to show that 2δ and $(1+\delta)\delta$ are

integer combinations of $2, 1+\delta$, and indeed

$$2\delta = 2(1+\delta) - 2, \quad (1+\delta)\delta = -5 + \delta = (1+\delta) - 2 \cdot 3.$$

Drawing the lattice in the complex plane, it looks like a sheared rectangular grid. We basically take a checkerboard pattern of $\mathbb{Z}[\delta]$, and this means that A is the second kind of ideal from Proposition 129

On the other hand, $B = (3, 1+\delta)$ also has a lattice basis (this can be checked), and drawing out the ideal gives another sheared rectangular grid. $1+\delta$ has smallest magnitude here, and again we have the second type of ideal. That means $(2, 1+\delta)$ and $(3, 1+\delta)$ are actually similar figures.

Fact 131

We say that there are two **ideal classes** in $\mathbb{Z}[\sqrt{-5}]$: the principal ideal and the $(\alpha, \beta)R$ example from above. In a principal ideal domain, there is only one ideal class, and it's of the form αR.

So now let's try to do this procedure for $\mathbb{Z}[\sqrt{-3}]$. Most of the work here is in a homework problem, but the ideas are a little different: $\mathbb{Z}[\sqrt{-3}]$ is contained within $\mathbb{Z}[\omega] = \frac{1}{2}(-1 + \sqrt{3}i)$, so it's better to use the ring $\mathbb{Z}[\omega]$.

16 March 15, 2019

We'll continue our study of imaginary quadratic fields today. Let d be a squarefree negative integer, and let $\delta^2 = d$. (For example, we can have $d = -1, -2, -3, -5$, and so on.) Consider the field $F = \mathbb{Q}[\delta]$: all elements in F are of the form

$$a + b\delta, a, b \in \mathbb{Q}.$$

Definition 132

An **algebraic integer** in a field F is an element that is the root of a monic polynomial with integer coefficients. Equivalently, an algebraic integer is an element whose (monic) minimal polynomial has integer coefficients.

In particular, if $\alpha = a + b\delta$, the minimal polynomial is

$$(x - \alpha)(x - \overline{\alpha}) = x^2 - (\alpha + \overline{\alpha})x + (\overline{\alpha}\alpha),$$

so we must have $\alpha + \overline{\alpha} = 2a$ and $\overline{\alpha}\alpha = a^2 - b^2 d$ be integers if we want α to be an algebraic integer. Working out the fussy details, we must have

- both a, b are integers, or
- $d \equiv 1 \bmod 4$ **and** $a, b \in \mathbb{Z} + \frac{1}{2}$ are both half-integers.

We should be careful here, though: $d \equiv 1 \bmod 4$ means $d = -3, -7, -11, \cdots$, since d is a negative integer.

Proposition 133

If an algebraic integer a is rational, then $a \in \mathbb{Z}$.

Proof. Write $a = \frac{p}{q}$. If a is an algebraic integer, the leading term x^n contributes a $\frac{p^n}{q^n}$ term, which has factors in the denominator that the other terms cannot remove unless $q = \pm 1$, in which case a is an integer. □

Fact 134

In any field $F = \mathbb{Q}[\delta]$, the algebraic integers form a ring R, which is a lattice in $\mathbb{C}$.

Visually, if $d \not\equiv 1 \bmod 4$, then the lattice is a rectangular grid, and if $d \equiv 1 \bmod 4$, we also have the midpoints of those rectangles in our lattice. In both cases, we do have a lattice basis: everything is an integer combination of 1 and δ in the rectangular case, and of 1 and $\frac{1}{2}(1+\delta)$ in the other case.

Let's focus more on the **simpler versions** of R: assume $d \not\equiv 1 \bmod 4$, so we're in one of the cases

$$d = -1, -2, -5, -6, -10, \cdots.$$

In such rings, any nonzero ideal A is a sublattice of R with lattice basis α, β, and we can write

$$A = \alpha\mathbb{Z} + \beta\mathbb{Z}.$$

(This is an idea from 18.701: we have a discrete subset of the plane with two independent vectors.) We also have the idea of "ideal generators:" elements $\alpha_1, \cdots, \alpha_k$ **generate** R if we can write

$$A = \alpha_1 R + \cdots + \alpha_k R.$$

Remember that given nonzero ideals A, B, the **product ideal** AB is the set of finite sums $\sum_{i=1}^{k} \alpha_i\beta_i$, where $\alpha_i \in A, \beta_i \in B$. If $\alpha_1, \cdots, \alpha_k$ generate A and $\beta_1, \cdots, \beta_\ell$ generate B, then we know that $\{\alpha_i\beta_j\}$ generate AB (write any element as a linear combination). We also know that if A is the principal ideal αR, then $AB = \alpha RB = \alpha B$ – we just take all the elements in B and multiply them by α.

Fact 135

For any ideals A, B, C, we have that $AB = BA$, $(AB)C = A(BC)$, and $A(B+C) = AB + AC$ (because the ring is commutative, associative, and has distributivity). Also, $RA = A$ for any A.

As we've already done in earlier examples, we can construct the set $\overline{A} = \{\overline{\alpha} \mid \alpha \in A\}$. This is an ideal, because (1) for any $\alpha \in A, \rho \in R$, $\rho\overline{\alpha} = \overline{\overline{\rho}\alpha}$, and (2) $\overline{\rho} \in R$, because $R = \overline{R}$.

Proposition 136 (Main Lemma)

Suppose δ is a negative squarefree integer. If A is an ideal of R, the set of algebraic integers in $\mathbb{Q}[\delta]$, then $\overline{A}A$ is a principal ideal of R generated by a positive integer $n \in \mathbb{Z}$.

Proof. Take a lattice basis (α, β) for A. Then $\overline{A}$ is generated by $\overline{\alpha}, \overline{\beta}$, so $\overline{A}A$ is generated by the four elements

$$\overline{\alpha}\alpha, \overline{\alpha}\beta, \overline{\beta}\alpha, \overline{\beta}\beta.$$

We can't say that all four of these are rational, but we do know that $\overline{\alpha}\alpha, \overline{\beta}\beta$, and $(\overline{\alpha}\beta + \overline{\beta}\alpha)$ are rational (the last one because it is equal to its own conjugate), and because they are algebraic integers, they are in $\mathbb{Z}$. Let n be their greatest common divisor: it is in the ring, because n is a linear combination of the elements above. We claim that $\overline{A}A = nR$. One direction is easy to show: since n is generated by elements of $\overline{A}A$, we have $nR \subset \overline{A}A$.

Now, our goal is to show that $\overline{\alpha}\alpha, \overline{\alpha}\beta, \overline{\beta}\alpha, \overline{\beta}\beta \in R$. The first and last are clear, because they are multiples of n by construction: we just need to show that $\overline{\alpha}\beta$ and $\overline{\beta}\alpha$ are good as well. Alternatively, we can show that

$$\overline{\alpha}\beta \in nR \implies \frac{\overline{\alpha}\beta}{n} \in R.$$

And we can show that $\gamma = \frac{\overline{\alpha}\beta}{n}$ is an algebraic integer by showing that $\overline{\gamma}+\gamma$ and $\overline{\gamma}\gamma$ are **both integers**. We compute and find that

$$\overline{\gamma}+\gamma = \frac{\overline{\alpha}\beta + \alpha\overline{\beta}}{n},$$

and since we defined n to be a factor of $\overline{\alpha}\beta + \alpha\overline{\beta}$, this is indeed an integer. Similarly, $\overline{\gamma}\gamma = \frac{\overline{\alpha}\alpha\overline{\beta}\beta}{n^2} = \frac{\overline{\alpha}\alpha}{n}\frac{\overline{\beta}\beta}{n}$, which is an integer, and we're done. □

Corollary 137 (Cancellation law)
If we have ideals such that $AB = AC$, then $B = C$.

Proof. Multiply by $\overline{A}$. Then

$$\overline{A}AB = \overline{A}AC \implies (nR)B = (nR)C \implies nB = nC,$$

which means $B = C$. □

Corollary 138
If $A \supset B$, then there exists an ideal C such that $B = AC$.

Proof. If $A \supset B$, then $\overline{A}A = nR \supset \overline{A}B$, so everything in the product ideal $\overline{A}B$ is a multiple of n. Then let $C = \frac{\overline{A}B}{n}$; this is closed under multiplication by R because B is, so it's an ideal. Then

$$AC = A\frac{\overline{A}B}{n} = \frac{nRB}{n} = B,$$

as desired. □

Since every ideal is contained in a maximal ideal, this means we can factor ideals!

Fact 139
Instead of "maximal ideal," we're going to think about "prime ideals," but these are not generally equivalent.

Definition 140
Let R be any ring. An ideal P is a **prime ideal** if any of the following three are satisfied:

- R/P is a domain.
- If P is a proper ideal of R, $a, b \in R$, and $ab \in P$, then one of a or b is in P.
- If P is a proper ideal of R, A, B are ideals, and $AB \subset P$, then $A \subset P$ or $B \subset P$.

Let's show that the second condition implies the third: If $AB \subset P$ and $A \not\subset P$, then there exists an $a \in A$ where $a \notin P$. But $aB \subset AB \subset P$ (by the definition of an ideal), so for all $b \in B$, $ab \in P$. Thus by the second condition, $b \in P$, and that's true for all $b \in B$, so $B \subset P$.

Lemma 141
The prime ideals of R (the ring of integers of $\mathbb{Q}[\delta]$ are the zero ideal (since R is a domain) and the maximal ideals.

To show this, we'll need a subresult:

Lemma 142

For any nonzero ideal A of R, $|R/A|$ is finite.

Proof. $\overline{A}A = nR$, and $R/(nR)$ has n^2 elements. Since $nR \subset A$, $|R/A| \leq |R/nR| = n^2$. □

So to prove Lemma 141, note that if P is a nonzero prime ideal, then R/P is a finite domain by our sublemma, and therefore it is a field. And this means that P is indeed maximal by Corollary 84.

17 March 18, 2019

Let's review a bit about lattices. A **lattice** is a subgroup of $\mathbb{R}^2$ with a lattice basis α_1, α_2: we can write it in the form $A = \mathbb{Z}\alpha_1 + \mathbb{Z}\alpha_2$.

Definition 143

If $B \subset A$ is a sublattice, then the **index** $[A : B]$ is the number of additive cosets of the form $\alpha + B$ in A.

We basically ask how many copies of B can be translated to cover A exactly once. An alternative way to think about this is to draw a **minimal parallelogram** of B: we only want to count one of the corners, half of the points on the boundaries (since the top/bottom and left/right are translations of B), and all the points in the middle. In other words, we want to find points $\alpha_0 \in A$ satisfying

$$\alpha_0 = r_1\beta_1 + r_2\beta_2, 0 \leq r_1, r_2 < 1.$$

Alternatively, if we have a lattice basis α_1, α_2 for A, we can write β_1 and β_2 as integer combinations of α_1, α_2:

$$\begin{pmatrix}\beta_1\\ \beta_2\end{pmatrix} = \begin{pmatrix}r & s\\ t & u\end{pmatrix}\begin{pmatrix}\alpha_1\\ \alpha_2\end{pmatrix}$$

Then $\boxed{[A : B] = |\det M|}$; another way to say this is that

$$[A : B] = \frac{\Delta B}{\Delta A},$$

where ΔA is the area of the parallelogram spanned by the lattice basis in A.

Corollary 144

For lattices $A \supset B \supset C$, we have

$$[A : C] = [A : B][B : C].$$

(We just multiply out the areas.) And this statement is true for indices in any groups, even if the orders are finite!

Question 145. *What are the shortest vectors in a lattice A?*

Sometimes, like in the $\mathbb{Z}^2$ basis, the shortest vectors are just our lattice basis. However, if we have something like $v_1 = 10e_1 + 9e_2$, $v_2 = 11e_1 + 10e_2$ (which is invertible since $\begin{pmatrix}10 & 11\\ 9 & 10\end{pmatrix}$ has determinant 1). Then clearly the shortest vectors aren't going to come from the lattice basis.

For sake of simplicity, let's say the shortest vector α (that is, the one with minimal length) is horizontal. Then we have elements in our lattice $\cdots, -\alpha, 0, \alpha, 2\alpha, \cdots$.

We can't have another point that is within a circle of radius $|\alpha|$ of any point on the α-axis. So we have another point β in our lattice outside those circles: let's say β forms a lattice basis with α. Then the area of the parallelogram formed by α, β is $\Delta A = |\alpha| h$, where h is the height of β to the axis, and that height h is at least $\frac{\sqrt{3}}{2}|\alpha|$ (at the intersection of our circles). Thus,

$$\Delta A \geq \frac{\sqrt{3}}{2}|\alpha|^2.$$

Corollary 146

For any lattice A, we can bound the length of the shortest vector:

$$|\alpha|^2 \leq \frac{2}{\sqrt{3}}\Delta A.$$

With this, let's go back to quadratic number fields: say we're working with a negative square-free integer not congruent to 1 mod 4, so $d = -1, -2, -5, -6, \cdots$. Then the ring of integers R is $\mathbb{Q}[\delta]$, where $\delta^2 = d$, which can be written in the form

$$\{a + b\delta : a, b \in \mathbb{Z}\}.$$

Then $(1, \delta)$ forms a lattice basis: let's try to look at nonzero ideals in this ring R. Ideals are sublattices, and for any ideal A,

$$[R : A] = \frac{\Delta A}{\Delta R} = \frac{\Delta A}{|\delta|}.$$

Recall Proposition 136, the main lemma: for any ideal A, $\overline{A}A = nR$, where n is a positive integer. This gives us the cancellation law: $AB = AC \implies B = C$, and for any $A \subset B$, there exists a C such that $B = AC$.

Earlier, we also discussed the idea of a **prime ideal** – the following definitions are equivalent:

- R/P is a domain.
- If $ab \in P$, then $a \in P$ or $b \in P$.
- If A, B are ideals, then $AB \subset P \implies A \subset P$ or $B \subset P$.

The second and third points being equivalent comes essentially from comparing "subset" to "divides." Luckily, in the rings R we're talking about, **the prime ideals are maximal ideals**. (Maximal ideals are always prime because R/M is a field, which is a domain.) In particular, we showed last time that R/P is a finite field with order $[R : P]$.

We can use this to think about the idea of **factoring into prime ideals**:

Proposition 147

If A is a nonzero proper ideal of $R = \mathbb{Q}[\delta]$, then we can write $A = P_1 \cdots P_k$, where P_is are prime ideeals unique up to ordering

Proof. The proof is the same as for the integers. If we want to factor A, choose a maximal prime ideal P that contains A – every proper ideal is contained in at least one. Then P divides A, so $A = PB$.

We claim that $B > A$. It's clear that $B \supset A$, because $PB \subset RB = B$. On the other hand, if we had $A = B$, we'd have $A = PA \implies RA = PA$. But now by the cancellation law, we have $R = P$, which is a contradiction. Thus,

$$A < B \implies [R : A] > [R : B],$$

and now we can just induct on the index! (The base case is when we have a maximal ideal, in which case we have a field with prime order.)

And showing uniqueness is the same as in factoring of integers. If $A = P_1' \cdots P_l'$, then $P_1' \supset P_1 \cdots P_k$, so by the third property, P_1' must contain one of the P_is and vice versa. □

Our next result was initially proved just for Gaussian integers:

Proposition 148

Let P be a prime ideal of R. Then $\overline{P}P$ is either pR or p^2R for a prime integer p. On the other hand, if p is a prime integer, then pR is a prime ideal or $\overline{P}P$ for some prime ideal P.

The proof should look fairly similar as well.

Proof. If $\overline{P}P = nR$, we can factor $n = p_1 \cdots p_k$ in the integers. Then $nR = (p_1R) \cdots (p_kR)$. This means

$$\overline{P}P = (p_1R) \cdots (p_kR),$$

and now we can factor the right side into prime ideals: since both sides must be the same by unique factorization, we have $k = 1$ or $k = 2$. The first case gives pR, and the second case has $P = p_1R \implies \overline{P} = p_1R$ as well, giving p^2R.

The second one is similar and not too hard. □

So the question now: which primes split, and which remain prime, if we're working in a ring $\mathbb{Q}[\delta]$? Well, p remains prime if and only if pR is a prime ideal, which implies that pR is maximal, and therefore $R/(pR)$ should be a field. To understand this, take the integer polynomials $\mathbb{Z}[x]$. We can mod out from $\mathbb{Z}[x] \to \mathbb{Z}[\delta]$ with kernel $x^2 - d$, and then map $\mathbb{Z}[\delta]$ to R/pR with kernel p. On the other hand, we can map $\mathbb{Z}[x]$ to $\mathbb{F}_p[x]$ first and then map it to R/pR with kernel $x^2 - d$. That tells us our the answer:

Proposition 149

A prime p splits if and only if $R/(pR)$ is a field, which happens if and only if $x^2 - d$ is irreducible in $\mathbb{F}_p$.

Example 150

Let $p = 5, d = -21$. Does 5 remain prime or split in $\mathbb{Q}[\sqrt{-21}]$?

Consider the polynomial $x^2 + 21$ mod 5: this is $(x + 2)(x - 2)$, so 5 **splits** in $\mathbb{Z}[\delta]$. Thus we have

$$5R = \overline{P}P;$$

what do the ideals $P, \overline{P}$ look like? Notice that $x - 2$ is sent to $\delta - 2$ when we take $\mathbb{Z}[x] \to \mathbb{Z}[\delta]$, but we want $x - 2$ to be in the kernel of the map from $\mathbb{Z}[x] \to R/pR$. So $\delta - 2$ should be in the kernel of R/pR for consistency, and therefore $P = (5, \delta - 2)$. Let's check this: we have $\overline{P} = (5, \delta + 2)$, and

$$\overline{P}P = (25, 5\delta + 10, 5\delta - 10, \delta^2 - 4 = 25).$$

We can write a linear combination of these that includes 5, so this ideal is indeed $5R$, as desired.

18 March 20, 2019

Recall that if we have a negative squarefree integer $d \equiv 2, 3 \bmod 4$, we can look at the ring of algebraic integers of $\mathbb{Q}[\delta]$, which take the nice form $\mathbb{Z}[\delta] = \{a + b\delta \mid a, b \in \mathbb{Z}, \delta^2 = d\}$. In such a ring, we know that we can factor any nonzero ideal A uniquely up to order.

We also know that for any integer prime p, we can write the ideal pR either as P (in which case p **remains prime**) or as $\overline{P}P$ (in which case p **splits**. We looked at $d = -21$ as an example last time by using the following property:

Proposition 151

p splits if and only if there exists an integer a such that $a^2 \equiv d \bmod p$. Then the ideal $P = (p, a - \delta)$ yields the factorization for $\overline{P}P = pR$.

Proof. We know that $P > pR$, since $(p, a - \delta)$ is strictly larger than p ($a - \delta$ is not a multiple of p). Computing $\overline{P}P$,

$$\overline{P}P = (p, a - \delta)(p, a + \delta) = (p^2, p(a + \delta), p(a - \delta), a^2 - d).$$

Now p divides all terms as long as $a^2 \equiv d \bmod p$. So $pR \supset \overline{P}P$, and therefore pR divides $\overline{P}P$. This means pR is $\overline{P}P$ times some other ideal, and because we know pR must split into $\overline{P}P$ or remain prime by Proposition 148, this is actually an equality. (And this tells us that P is a prime ideal.) □

Example 152

Let's consider the case $p = 2, d = -21$. We know that $-21 \equiv 1 \bmod 2$, so 2 splits. Similarly, 3 splits because $-21 \equiv 0 \bmod 3$. Finally, 5 splits as well: $-21 \equiv 4 \bmod 5$.

Specifically, using the notation for principal ideals $pR = (p)$, we know that

$$P = (2, 1 - \delta) \implies (2) = \overline{P}P,$$

$$Q = (3, -\delta) \implies (3) = \overline{Q}Q,$$

$$S = (5, 2 - \delta) \implies (5) = \overline{S}S.$$

There is a special case:

Fact 153

It's possible that p splits into $(p) = \overline{P}P$, and $\overline{P} = P$. For example, with P above, $(2, 1 - \delta) = (2, 1 + \delta)$, and similarly with Q, $(3, -\delta) = (3, \delta)$. In these cases, p **ramifies** (so 2 and 3 ramify but 5 don't).

Then we have $\overline{P}P = P^2 = (p)$, where $P = (p, a - \delta)$. Expanding out the product,

$$P^2 = (p^2, p(a - \delta), (a - \delta)^2).$$

Since $(a - \delta)^2 = a^2 - 2a\delta + d$, and this is supposed to be an element of (p), we must have p divide $2a$ and $a^2 + d$ if the prime p ramifies. By construction, $p | a^2 - d$, so this is equivalent to saying that p divides $2a$ and p divides $2d$.

Corollary 154

A prime p ramifies in $\mathbb{Z}[\sqrt{-d}]$ (where $d \equiv 2, 3 \bmod 4$) if $p = 2$ or $p | d$.

In both of these cases, p does divide $2a$: just use $(2, 1-\delta)$ for $p = 2$ and $(p, -\delta)$ in the other.

Definition 155

The **norm** of an element $\alpha \in R$ is

$$N(\alpha) = \overline{\alpha}\alpha = |\alpha|^2,$$

and this is always an integer.

Norm is multiplicative: $N(\alpha\beta) = N(\alpha)N(\beta)$ for all α, β. The main lemma says that if A is an ideal of R, then $\overline{A}A = nR$ for a positive integer n. So we can **define the norm of A to be n**: notice that this preserves the multiplicative property $N(AB) = N(A)N(B)$ for ideals.

Theorem 156

Let A, B, C be ideals of R, and $B \supset C$. Then $[B : C] = [AB : AC]$, and $N(A) = [R : A]$.

This is obvious for a principal ideal: $[A : pA] = p^2$ by drawing a p by p grid. But it's a bit harder to prove this in general:

Proof. It's enough to prove the first statement for $A = P$, a prime ideal, since we can always write out $A = P_1P_2\cdots$ and then successively apply the result for prime ideals one at a time by induction.

We know that B contains C, PB contains PC, B contains PB, and C contains PC. Since $[A : C] = [A : B][B : C]$,

$$[B : PC] = [B : C][C : PC] = [B : PB][PB : PC],$$

and now it's enough to show that $[B : PB] = [C : PC]$ – we'll do this by computing directly for each of the possible forms of the prime ideal P.

If $P = pR$, then $[B : PB] = p^2$, and $[C : PC] = p^2$ as well. Otherwise, $pR = \overline{P}P$, and then

$$B \supset PB \supset \overline{P}PB = pB,$$

and this means (by the product rule again) that $p^2 = [B : PB][PB : \overline{P}PB]$. Those must be both p (unless one of them is 1, which is not the case by cancellation law). So then $[B : PB] = p$, and the same argument works for $[C : PC]$ as well. Thus we've shown in all cases that $[B : PB] = [C : PC]$, and thus $[B : C] = [PB : PC]$ as desired.

For the second statement, let's factor $A = PQ$, where P is a prime ideal and Q is some (not necessarily prime) ideal. We know that $N(A) = N(P)N(Q)$, and $A \supset PA \supset PQA$. Looking at the index of A in R, we know that $[A : PA]$ is $N(P)$ by the arguments above, and then

$$[R : A] = [A : PA][PA : PQA] = N(P)N(Q)$$

since $[PA : PQA] = [R : Q]$ (by using the cancellation law from the first part of this theorem twice), and **by induction** we know that $[R : Q] = N(Q)$. Finally, since norm is multiplicative, this is just $N(PQ) = N(A)$, as desired. □

Recall Corollary 146, which says that if A is an ideal and α is a shortest nonzero vector in A, then $|\alpha|^2 \leq \frac{2}{\sqrt{3}}\Delta A$. From the above theorem, we have that $[R : A] = \frac{\Delta A}{\Delta R}$: we also know that $\Delta R = \sqrt{|d|}$ in the lattices we're dealing with. So if we plug in all of the results that we have,

$$|\alpha|^2 \leq 2\sqrt{\frac{|d|}{3}}[R : A] = 2\sqrt{\frac{|d|}{3}}N(A).$$

This motivates the definition of the constant $\boxed{\mu = 2\sqrt{\frac{|d|}{3}}}$.

Fact 157

The value of μ is different for $p = 1 \bmod 4$, but we're not dealing with that case right now.

Let's define an **equivalence relation** on ideals: $A \cong A'$ if $A' = \lambda A$ for some complex $\lambda \in \mathbb{C}$. Basically, this means the two lattices are similar geometric figures, except that we can't change orientation of the lattice. Note that all principal ideals are equivalent: $R \sim A$ if $A = \lambda R$.

Definition 158

An **ideal class** is an equivalence class of ideals (where $A \sim A'$ if $A' = \lambda A$ for some $\lambda \in \mathbb{C}$). Let the **class** of A be denoted $\langle A \rangle$.

Theorem 159

Every ideal class contains an ideal I with norm $N(I) \leq \mu = 2\sqrt{\frac{|d|}{3}}$.

In other words, we don't have to look at very small ideals with large lattices: we can just look within some bound.

Proof. Start with any ideal A. Choose an α with $N(\alpha) \leq \mu N(A)$: this exists by the calculations we were doing above. Note that A contains the principal ideal (α), so A divides α and we can write $(\alpha) = AB$ for some B. Taking norms of both sides,

$$N(\alpha) = N(A)N(B) \leq \mu N(A)$$

so we have $N(B) \leq \mu$. We're not quite done here, but we'll show that $\overline{B}$ is in the class of A and has the same norm as B next time. □

Example 160

For $d = -21$, $\mu = 2\sqrt{7} < 6$. So we only have to look at ideals with norm less than 6 to find all the ideal classes.

19 March 22, 2019

We're going to say some more about ideal classes today. As before, we're still working in the ring $\{a + b\delta \mid a, b \in \mathbb{Z}\}$ for some $\delta^2 = d$, where d is a negative squarefree integer congruent to 2 or 3 mod 4. Recall that an **ideal class** is an equivalence class of ideals: $A \sim A'$ if $A' = \lambda A$ for some $\lambda \in \mathbb{C}$. Our goal is to make these ideal classes form a **group**.

Denote the class of A as $\langle A \rangle$. We can define a law of composition via $\langle A \rangle \langle B \rangle = \langle AB \rangle$; this is commutative and associative because multiplication in the ring is commutative and associative. This law of composition is indeed consistent: if $A \sim A', B \sim B'$, then $AB \sim A'B'$, since we can write $A' = \lambda A, B' = \lambda' B \implies A'B' = \lambda\lambda' AB$.

Note that the identity element here is the whole ring: denote this as $\langle R \rangle = 1$. Then an ideal $A \in \langle R \angle$ if and only if $A = \lambda R$: this means A is a **principal ideal**. Furthermore, $\langle A \rangle^{-1} = \langle \overline{A} \rangle$, because $\overline{A}A = nR$ for some integer $n \in \mathbb{Z}$, and therefore $\langle \overline{A} \rangle \langle A \rangle = \langle (n) \rangle = 1$. This means the ideal classes form an abelian group called the **ideal class group**! We'll denote this as $\mathcal{C}$.

Example 161

$\mathcal{C}$ is the trivial group if every ideal is principal: in these cases, R has unique factorization, since all PIDs (principal ideal domains) are UFDs (unique factorization domains).

We call $|\mathcal{C}|$ the **class number**: it measures the "failure" of unique factorization.

It's surprisingly easy to compute the class number! Remember that we can define the constant $\mu = 2\sqrt{\frac{|d|}{3}}$, and we can apply Theorem 159, which says that every ideal class $\langle A \rangle$ contains an ideal A' with norm at most μ.

As a quick review, the norm of an ideal is defined to be $N(A) = n$ if $\overline{A}A = nR$. This is also the index $[R : A]$, which we can measure by taking a lattice basis and computing $\frac{\Delta A}{\Delta R}$. In our ring $R = \{a + b\delta\}$, a lattice basis is a by δ, so $\Delta R = \sqrt{|d|}$ (the area of the rectangle formed by 1 and δ). And this norm is useful because it is multiplicative and always an integer!

Continuation of proof of Theorem 159. As before, if we're given an ideal A, let α be the shortest nonzero vector in A. We have

$$N(\alpha) = |\alpha|^2 \leq \frac{2}{\sqrt{3}}\Delta A,$$

and we know $\alpha \in A$, so the principal ideal generated by α is a subset of A. Inclusion of ideals gives us divisibility: therefore, A divides (α), and we can write $(\alpha) = AB$ for some ideal B. Now notice that (working in the ideal class group) $1 = \langle (\alpha) \rangle = \langle A \rangle \langle B \rangle$, and $N(\alpha) = N(A)N(B)$ as well: this second fact means

$$N(A)N(B) = N(\alpha) \leq \frac{2}{\sqrt{3}}\Delta A = \frac{2}{\sqrt{3}}\Delta R \cdot N(A),$$

and therefore we have $N(B) \leq \frac{2}{\sqrt{3}}\Delta R = \frac{2}{\sqrt{3}}\sqrt{|d|} = \mu$. But now note that $\langle B \rangle = \langle A \rangle^{-1}$, so $\langle \overline{B} \rangle = \langle A \rangle$, so $\overline{B}$ is in the same ideal class of A. And since $N(B) = N(\overline{B})$, we're done: we've proven that A is in the same ideal class as an ideal $\overline{B}$ with norm at most μ. □

So now we can just look at all ideals with norm less than μ: there's only a finite number of them! We could brute-force our way through them, but here's a better way: we can write $A = P_1 \cdots P_k$ as a product of prime ideals P_i. Norm is multiplicative, so

$$N(A) = N(P_1) \cdots N(P_k),$$

where $N(P_i) \leq \mu$ for all i. That means $\mathcal{C}$ is generated by prime ideals $\langle P \rangle$, where $N(P) \leq \mu$ as well!

Why is it easier to work with prime ideals? We know that

$$N(P) = p \text{ or } p^2,$$

where p is a prime integer, corresponding to p splitting or remaining prime. If p remains prime, then the ideal P is just generated by p, which is a principal ideal: those are in the same class as R, so we can ignore those! This has led us to the following result:

Proposition 162

The ideal class of $\mathbb{Z}[\delta]$ is generated by the prime ideals P where $\overline{P}P = pR$, $p \leq \mu$, and p splits.

Example 163

Let's go back to $d = -21$: what's the ideal class group?

We have $\mu = 2\sqrt{7} < 6$, so the class group is generated by the prime ideals that are factors of $(2), (3), (5)$. These principal ideals all split: 2 always ramifies for $p \not\equiv 1 \bmod 4$, so $(2) = P^2$. 3 ramifies as Q^2 as well (recall that this is because 3 divides d). Finally, 5 splits as $\overline{S}S$.

So what are our relations for the ideal class group? We have $\langle P\rangle^2 = 1$ and $\langle Q\rangle^2 = 1$. S is a bit harder, but we have a **secret method**: compute some norms of elements! Note that

$$N(1+\delta) = (1+\delta)(1-\delta) = 22 = 2 \cdot 11,$$

and this is an equation among elements, but we can take the **principal ideals** generated by the left and right side instead. So we can rewrite the above equation as

$$(1+\delta)(1-\delta) = (22) = (2)(11)$$

where this equation now refers to **ideals**. To factor this equation, note that $(1+\delta) = P_1 \cdots P_k \implies (1-\delta) = \overline{P_1} \cdots \overline{P_k}$, and those must be the same prime ideals that divide the right side. (2) factors as P^2 from above, and since we have an even number of prime ideals on the left side, we must write the right side as $P^2 \cdot \overline{T}T$ for some prime ideal T (in other words, 11 splits, and $x^2 + 21$ is reducible mod 11).

But then that means $k = 2$, and now we can't have $\overline{P}$ and P be part of the factorization of the same ideal on the left side: after all, the ideal generated by 2 and by $1+\delta$ are different. So this means that

$$(1+\delta) = P\overline{T}, \quad (1-\delta) = PT.$$

That wasn't particularly helpful, so let's try looking at the norm of $2+\delta$: it is $2-\delta^2 = 25$, so looking at the ideals, we have

$$(2+\delta)(2-\delta) = (5)(5) = \overline{S}S\overline{S}S.$$

so now we must have either $(2+\delta) = \overline{S}S, S^2$, or $\overline{S}^2$: it's not the first one because the ideal generated by 5 is not the ideal generated by $2+\delta$. Thus $(2+\delta) = S^2$ (or $\overline{S}^2$, without loss of generality), and this means S^2 is a principal ideal! Therefore $\langle S\rangle^2 = 1$ in our class group.

So now we know that $\langle P\rangle^2 = \langle Q\rangle^2 = \langle S\rangle^2 = 1$, so our ideal class group has order at most 8. But there might be other relations as well!

To figure those out, let's compute $N(3+\delta)$: it is $(3+\delta)(3-\delta) = 30 = 2 \cdot 3 \cdot 5$, so

$$(3+\delta)(3-\delta) = (2)(3)(5) = P^2Q^2\overline{S}S.$$

The left side factors into prime ideals that must be conjugates of each other, so we must have $(3+\delta) = PQS$ or $PQ\overline{S}$. But $\langle S\rangle^2 = 1$, so the classes $\langle S\rangle = \langle\overline{S}\rangle$, and this gives us the relation

$$\langle P\rangle\langle Q\rangle\langle S\rangle = 1.$$

This is another relation, and we can use it to eliminate S from our list of generators! We know that $\langle P\rangle$ and $\langle Q\rangle$ are not principal ideals (in particular, $P = (2, 1-\delta), Q = (3, \delta)$), so both ideal classes are not trivial. So our group contains $1, \langle P\rangle, \langle Q\rangle, \langle PQ\rangle$, and the only remaining possible relation is

$$\langle P\rangle\langle Q\rangle = 1,$$

which is not true, since the ideal has norm $N(PQ) = N(P)N(Q) = 2 \cdot 3 = 6$. Then $PQ = (\alpha)$ would have to be generated by some element with norm 6, and none exist! Thus we've found our class group: $\boxed{\mathcal{C} = C_2 \times C_2}$. The idea

is that we can always try this kind of drawn-out calculation, and it'll work.

Example 164

Now let's look at $d = -26$: we find that $\mu < 6$ again, so we again need to only need to look at $2, 3, 5$.

As always, 2 ramifies as $(2) = P^2$. For 3, note that -26 is 1 mod 3, which is a square, and therefore 3 splits. Finally, -26 is 4 mod 5, which is a square, so 5 splits as well. Once again, we have $\langle P\rangle^2 = 1$, and we want to use norms to find the rest.

Some potential candidates here are $N(1+\delta) = 27 = 3^3$, $N(2+\delta) = 30 = 2\cdot 3\cdot 5$, and we could also try $N(3+\delta) = 35$ and $N(4+\delta) = 42$ (which are less useful). Looking at ideals,

$$(1+\delta)(1-\delta) = (3)^3 = \overline{Q}^3 Q^3,$$

and

$$(2+\delta)(2-\delta) = (2)(3)(5) = P^2\overline{Q}Q\overline{S}S.$$

So $(1+\delta)$ uses three of the Qs and $\overline{Q}$s, which means it's one of $Q^3, Q^2\overline{Q} = 3Q, Q\overline{Q}^2 = 3\overline{Q}$, or $\overline{Q}^3$. It can't be $3Q$ or $\overline{Q}$, because $(1+\delta)$ is not a multiple of 3, and then it doesn't matter which of Q and $\overline{Q}$ we choose – either way, we find that $\langle Q\rangle^3 = 1$, since $(1+\delta)$ is principal.

The next step is to show similarly that

$$(2+\delta) = PQ^{\pm 1}S^{\pm 1},$$

and that means $\langle P\rangle\langle Q\rangle^{\pm 1}\langle S\rangle^{\pm 1} = 1$, meaning we can always solve to eliminate S. So the ideal class group of $\mathbb{Z}[\sqrt{-26}]$ is generated by $\langle P\rangle^2 = 1$ and $\langle Q\rangle^3 = 1$, and thus $\boxed{\mathcal{C} = C_6}$ is the cyclic group of order 6.

Note that it is true in general that $\langle S\rangle^2 = 1 \implies \langle S\rangle = \langle\overline{S}\rangle$, but this does not mean S and $\overline{S}$ are the same. We only know that $\overline{S} = \lambda S$: what is λ here? Let's take $d = -21$ as an example: we can use the fact that

$$(\overline{S}S)S = (5)S = 5S,$$

but also

$$\overline{S}(SS) = \overline{S}(2+\delta).$$

So $\overline{S} = \frac{5}{2+\delta}S$ – notice that the absolute value of the constant term is 1 here, which we could have also found by taking norms.

20 April 1, 2019

Let's talk about abelian groups today: we'll be using additive notation. For example, we can look at a lattice in $\mathbb{Z}^n$ or the cyclic group $\mathbb{Z}/n\mathbb{Z}$ for a positive integer n. We'll carry over the description of a basis from that for a vector space. So V is an abelian group: to describe some element $(v_1, \cdots, v_n) \in V$, we have a map

$$\overline{v} : \mathbb{Z}^n \to V,$$

which sends an n-tuple of integers x into an element of our vector space via the map $\overline{v}x = \sum v_i x_i$. In other words, we have a homomorphism compatible with the group operations. Then $\overline{v}$ is **independent** if it is an injective map, and $\overline{v}$ **generates** (or **spans**) V if it is surjective. And as usual, $\overline{v}$ is a basis of our vector space if it is bijective.

Abelian groups may not have a basis: for example, we can write all elements of $\mathbb{Z}/n\mathbb{Z}$ in many different ways. But

if V has a basis, then it is isomorphic to $\mathbb{Z}^n$: then V is a **free abelian group**. (We'll get back to the case where we don't have a basis later.)

So now, let V be a free abelian group with basis $\overline{v} = (v_1, \cdots, v_n)$. Let's try to develop some theory:

- We can do a base change: if we have a new basis $\overline{v}' = (v_1', \cdots, v_n')$, we can write

$$v_j' = \sum_{i=1}^{n} v_i a_{ij}$$

for some integers p_{ij}. So $\overline{v}' = \overline{v}P$ for some matrix P, and this means that for any vector $v \in V$, we can write it in two different ways with coordinate vectors:

$$v = \overline{v}x = \overline{v'}x' \implies Px' = x.$$

Since p_{ij} are integers, now P can be any invertible matrix with entries in $\mathbb{Z}$: $P \in GL_n(\mathbb{Z})$. But from the first 18.701 assignment, this can only happen if the **determinant is** ± 1, and this basically tells us everything we need about basechange matrices for now.

- Consider some homomorphism $\phi : V \to W$ of free abelian groups. This is just like a matrix of a linear transformation: if $\overline{v} = (v_1, \cdots, v_n)$ is a basis of V, and $\overline{w} = (w_1, \cdots, w_m)$ is a basis of W, then each basis vector

$$\phi(v_j) = \sum w_i a_{ij}$$

for some integers a_{ij}. So now

$$(\phi(v_1), \cdots, \phi(v_n)) = (w_1, \cdots, w_m)A$$

for some matrix A with integer entries. If we change bases for our vector spaces, so that $\overline{v}' = \overline{v}P$ and $\overline{w}' = \overline{w}Q$, then P is an $n \times n$ matrix and Q is an $m \times m$ matrix: our new homomorphism matrix is now

$$A' = Q^{-1}AP.$$

Indeed, this is the same change-of-basis formula that we're used to.

Time to move on to something slightly interesting: **how simple can we make** A'? Remember that in a vector space over a field F, for any $m \times n$ matrix A with entries in F, there exist invertible matrices P, Q in F such that A' has the block form

$$A' = \begin{pmatrix} I & 0 \\ 0 & 0 \end{pmatrix}.$$

In other words, each A' takes some basis elements to other ones and throws away other ones. But we're working in $\mathbb{Z}$, which is not a field – it turns out we can still get something analogous!

Theorem 165

Let A be an $m \times n$ matrix with entries in $\mathbb{Z}$. Then there exist invertible matrices P and Q such that $Q^{-1}AP$ is diagonal:

$$A' = Q^{-1}AP = \begin{pmatrix} D & 0 \\ 0 & 0 \end{pmatrix},$$

where D is diagonal with integer entries $d_1 | d_2 | d_3 \cdots | d_k$.

(The fact that d_is divide each other is often not that important, but diagonalizing is!) Let's do a proof by example.

Example 166

Take $A = \begin{pmatrix} 4 & 8 \\ 8 & 14 \end{pmatrix}$, and let's try to simplify this matrix with a change of basis.

Here $d_1 = 2$ is going to be the greatest common divisor of the entries. How are we finding P and Q? Note that our row-reducing matrices

$$\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}, \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix}$$

generate the whole group (this was also an 18.701 homework problem). So we'll just row reduce until our matrix looks as simple as possible:

$$\begin{pmatrix} 4 & 8 \\ 8 & 14 \end{pmatrix} \to \begin{pmatrix} 4 & 4 \\ 8 & 6 \end{pmatrix} \to \begin{pmatrix} 4 & 4 \\ 4 & 2 \end{pmatrix} \to \begin{pmatrix} -4 & 4 \\ 0 & 2 \end{pmatrix} \to \begin{pmatrix} -4 & 0 \\ 0 & 2 \end{pmatrix} \to \begin{pmatrix} 4 & 0 \\ 0 & 2 \end{pmatrix} \to \begin{pmatrix} 2 & 0 \\ 0 & 4 \end{pmatrix}.$$

(We won't actually prove the result in more detail, but we can write out a systematic way to do this.)

Now that we have a simple form for our matrix, we can say that A' operates on our vector space by sending our basis elements in V to simple elements of W:

$$v_1 \to d_1 w_1,\ v_2 \to d_2 w_2,\ \cdots,\ v_k \to d_k w_k,\ v_{k+1} \to 0,\ v_{k+2} \to 0, \cdots.$$

For example, the image of $\begin{pmatrix} 2 & 0 \\ 0 & 4 \end{pmatrix}$ from $\mathbb{Z}^2 \to \mathbb{Z}^2$ sends the regular integer lattice to a dilated rectangular one! Since the area of the smallest parallelogram is 8, the index of this map is also 8.

So now let's shift back into looking at the case where we don't have a basis. Let our abelian group be K, and let's say we have some elements $(k_1, \cdots, k_m)$ which generate K. Then we have a surjective map

$$\overline{k} : \mathbb{Z}^m \to K.$$

Let L be the kernel of $\overline{k}$. By the first isomorphism theorem, K is isomorphic to the quotient group $\mathbb{Z}^m/L$.

Fact 167

There is a second and third isomorphism theorem, but they aren't very interesting.

Let's choose some generators for our kernel L: we'll use the following useful fact.

Proposition 168

Every subgroup of a finitely generated abelian group is finitely generated.

This is basically because every element in our subgroup can be written as a combination of the generators of the original group, and we can't require an infinite set of generators! (We'll speak about this in a bit more detail later on, but the words to keep in mind are that a finitely generated abelian group is a **finitely generated module** over $\mathbb{Z}$.)

So L can be finitely generated by some $(\ell_1, \ell_2, \cdots, \ell_n)$, so that our map $\overline{\ell} : \mathbb{Z}^n \to L$ is surjective. By definition, L is the kernel of our homomorphism from $\mathbb{Z}^m$ to K, so L is some subset of $\mathbb{Z}^m$. Inclusion is a homomorphism as well, so we get a homomorphism from $\mathbb{Z}^n$ to $\mathbb{Z}^m$, which is described as multiplication by an integer matrix! L can be written symbolically here as $A\mathbb{Z}^n$, and now our finitely generated abelian group $K = \mathbb{Z}^m/L$ from above is isomorphic to $\mathbb{Z}^m/A\mathbb{Z}^n$.

Remember that we can change bases for $\mathbb{Z}^n$ and $\mathbb{Z}^m$, so we can suppose that A is the matrix with block form $\begin{pmatrix} D & 0 \\ 0 & 0 \end{pmatrix}$ from the theorem above. Multiplying out $A\mathbb{Z}^n$, we find that the additional **relations** on our generators look like $d_1 e_1 = 0, d_2 e_2 = 0, \cdots, d_k e_k = 0$.

So in our example above,

$$K = \mathbb{Z}^2 / \begin{pmatrix} 2 & 0 \\ 0 & 4 \end{pmatrix} \mathbb{Z}^2 \cong \mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/4\mathbb{Z}.$$

Theorem 169 (Basis Theorem for Abelian Groups)

Every finitely generated abelian group can be written as

$$\mathbb{Z}^k \times \bigotimes_{i=1}^{n} \mathbb{Z}/(d_i\mathbb{Z}),$$

where k is a nonnegative integer and all d_i are positive integers.

So we've now arrived at a pretty important result: all finitely generated abelian groups are isomorphic to the **product of (possibly infinite) cyclic groups**!

21 April 3, 2019

The second quiz for this class is next Wednesday, April 10 – it will cover chapters 12 and 13.

We'll start a new topic today:

Definition 170

Given a ring R, an ***R*-module** V is a set with two laws: addition $V \times V \to V$ and scalar multiplication $R \times V \to V$. V must be an abelian group under addition, and we also have the following axioms:

- $a(bv) = (ab)v$ for $a, b \in R, v \in V$, and $1v = v$.
- Distributivity holds: $a(v_1 + v_2) = av_1 + av_2$, and $(a_1 + a_2)v = a_1 v + a_2 v$.

Except in very simple cases for our ring R, modules can look very complicated.

Example 171

What does a $\mathbb{Z}$-module look like?

We need to specify how to add vectors in our abelian group V, and we also need to know how to multiply by scalars. But we already know that $1v = v$, and then $2v = (1+1)v = v + v$ and so on, so this is just an **abelian group** – the scalars $\mathbb{Z}$ don't do anything for us.

Example 172

Let $R = F[t]$, where F is a field. What does an R-module look like?

We know that V is an F-vector space, if we just forget about the ts and look at the constant polynomials. In addition to that, we can also multiply by our "scalar" t to get a new element tv for any $v \in V$. So define some linear

operator $T : V \to V$

$$T(v) = tv.$$

It is true that $T(v_1 + v_2) = t(v_1 + v_2) = tv_1 + tv_2 = T(v_1) + T(v_2)$, and $T(cv) = t(cv) = (tc)v = (ct)v = c(tv) = cT(v)$, so this is indeed a linear operator. So an R-module gives us **a linear operator, as well as a vector space**.

Conversely, if we have an F-vector space V and a linear operator $T : V \to V$, we can make V into an R-module (where $R = F[t]$) with $tv = T(v)$, $t^2v = T(T(v))$, and so on. So there's a correspondence between $F[t]$-modules and F-vector spaces with linear operators!

Example 173

What is an R-module, where $R = F[x, y]$?

Then R-modules correspond to an F-vector space with two linear operators: one that sends v to xv, and one that sends v to yv. But remember that $xy = yx$, so our operators need to commute! In other words, the R-modules are in one-to-one correspondence with F-**vector spaces with two linear operators** X, Y satisfying $XY = YX$.

Example 174

What is a $\mathbb{Z}[i]$-module?

The Gaussian integers contain the integers, so we have our usual abelian group V if we ignore i. Now we have to think about what it means to multiply by i in our abelian group: it has to be some homomorphism $I : V \to V$ from the group to itself, but $i^2 = -1$. So $I(v) = iv \implies I^2 = -\text{id}$. For example, if our set V is $\mathbb{Z}^2$, we can have I be the 2×2 matrix $\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$.

Let's (again) go through the logic for the basis calculations – this should look familiar from earlier in the class. Suppose we have an R-module V, and let's say we have some elements $\overline{v} = (v_1, v_2, \cdots, v_m)$. We have a homomorphism

$$\overline{v} : R^m \to V$$

sending $x \to \overline{v}x = \sum v_i x_i$, and this is a **homomorphism** of R-modules. Let's rigorize what a homomorphism means here:

Definition 175

A homomorphism $\phi : W \to V$ of R-modules is a map that is compatible with the operations: $\phi(w_1 + w_2) = \phi(w_1) + \phi(w_2), \phi(rw) = r\phi(w)$.

We'll have the same notions of kernel and image as in the usual group and ring homomorphisms – we can say that $\overline{v}$, our homomorphism, is **independent** if our map is injective, that it **generates** V if $\overline{v}$ is surjective, and that it is a **basis** if it is bijective. It's harder to satisfy all of these, now that we have extra structure on our abelian group!

Definition 176

An R-module V is **finitely generated** if there exists a $\overline{v}$ that generates V, and V is a **free module** if there exists a basis $\overline{v}$.

Since $\overline{v}$ is a map from R^m to V, free modules V are isomorphic to R^m. So let's only say that V is finitely generated, and suppose $\overline{v} : R^m \to V$ is our surjective map. Letting $W = \ker \overline{v}$, we know that V is isomorphic to R^m/W.

Question 177. *Is W finitely generated?*

We're going to defer this to later, but let's assume that W is indeed finitely generated. Then we can say those generators are $\overline{w} = (w_1, \cdots, w_n)$, and these generators give us a map $\overline{w} : R^n \to W \subset R^m$ sending x to $\overline{w}x$, and this is described by an R-matrix A (which has dimensions $m \times n$).

Now W is the image of $\overline{w}$, which is AR^n, so our R-module is isomorphic to $R^m/W = R^m/(AR^n)$. In this case, we say that A **presents** the module V (because it tells us the relations).

Example 178

Let's say $R = F[t]$, and $A = \begin{pmatrix} t^2+1 & t \\ t+1 & t^2+t \end{pmatrix}$. What does our R-module look like?

We have $m = n = 2$, and R^2 has basis e_1, e_2, so V is generated by the images v_1, v_2 of e_1, e_2 under some map. But we also know something else about v_1 and v_2: the image of A is generated (with coefficients in $F[t]$) by the two column vectors

$$\begin{pmatrix} t^2+1 \\ t+1 \end{pmatrix}, \begin{pmatrix} t \\ t^2+t \end{pmatrix} \in R^2 = F[t]^2.$$

We want to mod out by AR^n in V (this is the kernel of the map $\overline{v}$ mentioned above), so we have the relations $(t^2+1)v_1 + (t+1)v_2 = 0, tv_1 + (t^2+t)v_2 = 0$ that define AR^2. This is not a particularly nice looking set.

But $R = F[t]$ is very simple: we can make our matrix A look a bit simpler! Remember that we can do division by remainder for $\mathbb{Z}$-matrices A: we basically use the Euclidean algorithm by looking only at the first column, and then rinse and repeat. Does something work here as well?

Theorem 179

Let $R = F[t]$, and let A be an $m \times n$ R-matrix. Then there exist invertible R-matrices P, Q such that $A' = Q^{-1}AP$ is diagonal with entries $d_1 | d_2 | \cdots | d_k$.

The proof is very similar: let's do another proof by example. We row-reduce to find that

$$\begin{pmatrix} t^2+1 & t \\ t+1 & t^2+t \end{pmatrix} \to \begin{pmatrix} 1 & t \\ t+1-t^3-t^2 & t^2+1 \end{pmatrix} \to \begin{pmatrix} 1 & 0 \\ t+1-t^3-t^2 & t^4+t^3 \end{pmatrix} \to \begin{pmatrix} 1 & 0 \\ * & t^4+t^3 \end{pmatrix} \to \begin{pmatrix} 1 & 0 \\ 0 & t^4+t^3 \end{pmatrix}.$$

So what's the module with presentation matrix $A' = \begin{pmatrix} 1 & 0 \\ 0 & t^4+t^3 \end{pmatrix}$? Here V is generated by two elements v_1', v_2' (since we changed a basis) and our relations are

$$1v_1' + 0v_2' = 0, 0v_1' + (t^4+t^3)v_2' = 0.$$

So V is generated by a **single element** v' with $(t^4+t^3)v' = 0$. Remember that an R-module for $R = F[t]$ is a vector space with a linear operator: to find our vector space, we have elements of the form cv, $c \in F$, and then V is generated by $\{v, tv, t^2v, t^3v\}$ (since $t^4v = -t^3v$). Calling those elements v_1, v_2, v_3, v_4, we have a linear operator

$$Tv_1 = v_2, Tv_2 = v_3, Tv_3 = v_4, Tv_4 = -v_4 \implies T = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & -1 \end{pmatrix},$$

and now the combination of our abelian group V and linear operator T is the R-module we're looking for.

22 April 5, 2019

A correction: the quiz is actually on Monday. Recall that if we have an R-module V, and $\overline{v} = (v_1, \cdots, v_n)$ is a set of elements of V, then $\overline{v}$ generates V if every $v \in V$ is a linear combination

$$v = \sum v_i x_i, x_i \in R,$$

or if there exists a surjective map

$$\overline{V} : R^m \to V$$

sending x to $\overline{v}x = \sum v_i x_i$. Today, we'll discuss finitely generated modules, which are modules where the finite set $\overline{v}$ exists.

Proposition 180

Let $\phi : V \to W$ be a surjective homomorphism of R-modules, and let $\ker \phi = U$. If U and W are finitely generated, so is V, and if V is finitely generated, so is W.

It turns out U is not necessarily always finitely generated for modules in general.

Proof. Let's first prove the second statement. If some set $\overline{v} = (v_1, \cdots, v_m)$ generate V, let $\overline{w} = (w_1, \cdots, w_m)$ be the images of the v_is: we'll show that these generate W. Since we have a surjective homomorphism, for any $w \in W$, there exists a $v \in V$ such that $\phi(v) = w$. Now $v = \sum x_i v_i$ for some $x_i \in R$, and therefore $w = \sum x_i w_i$ is a linear combination of the elements of $\overline{w}$, as desired.

Now, let's prove the first statement. Let $\overline{u} = (u_1, \cdots, u_n)$ generate U, and $\overline{w} = (w_1, \cdots, w_m)$ generate W. Since we have a surjective homomorphism, we can let $v_i \in V$ be the elements such that $\phi(v_i) = w_i$: we claim that $(v_1, \cdots, v_m, u_1, \cdots, u_n)$ generate V. To show this, take any element $v \in V$ and let $w = \phi(v)$. Then $w = \sum w_i a_i$ for some $a_i \in R$, since the w_is generate W. Now if we define $v' = \sum v_i a_i$, $\phi(v') = w = \phi(v)$, so $v - v'$ must be an element of $\ker \phi$. Therefore $v - v'$ is a linear combination of the u_js, meaning we've written v as a linear combination of v_is and u_js, as desired. □

Next, let's ask a new question: what are **submodules of** R if we treat R as an R-module? We need to have closure under addition and scalar multiplication by elements of R, so this is just an **ideal**!

Definition 181

A ring R is **Noetherian** if every ideal is finitely generated.

This means that we can always find a finite set of elements in the ideal such that every element of the ideal is an R-linear combination of those elements. (And this is the missing ingredient we need to answer , because the kernel of a homomorphism is an ideal.)

Example 182

$\mathbb{Z}$ and $F[t]$ (for a field F) are Noetherian rings, because they are principal ideal domains – all ideals are just generated by one element.

Theorem 183 (Hilbert Basis Theorem, version 2)
If R is a Noetherian ring, then $R[x]$ is Noetherian.

We can substitute this into itself a bunch of times, too! By induction, any polynomial ring $R[x_1, \cdots, x_n]$ is Noetherian if R is Noetherian: in particular, $F[x_1, \cdots, x_n]$ and $\mathbb{Z}[x_1, \cdots, x_n]$ are Noetherian.

There's no bound on the **number of generators** of such a polynomial ring, though! (We just know that it's always finite.) For example, let $I_d \subset F[x, y]$ be the polynomials in two variables with no terms of degree less than d. This is an ideal because it is closed under addition within I_D and multiplication by any polynomial, and now notice that (for example) I_5 is generated by $x^5, x^4y, x^3y^2, x^2y^3, xy^4, y^5$. But in general I_d needs $d+1$ elements to be generated, which is unbounded.

Proof. Let I be an ideal of $R[x]$ for a Noetherian ring R. For any polynomial $f(x) = ax^m + \cdots$, define a to be the **leading coefficient** of f.

Lemma 184
Let A be the set of leading coefficients of polynomials $f \in I$, plus 0. Then A is an ideal of R.

Proof of lemma. Closure is clear for 0: $a+0 = a$ and $a \cdot 0 = 0$, which are both in the ideal for sure. If $a, b \in I$ are both nonzero, then there exist polynomials $f = ax^m + \cdots$ and $g = bx^n + \cdots$ in I: without loss of generality, let $m \leq n$, and now

$$x^{n-m}f(x) + g(x) = (a+b)x^n + \cdots.$$

If $a + b = 0$, 0 is already in A; otherwise, $a + b$ is the leading coefficient of a polynomial in I, so we have closure under addition. Similarly, if a is the leading coefficient of $f(x) = ax^m + \cdots$, ra is the leading coefficient of $rf(x) = rax^m + \cdots$ unless it is zero (which is already in A). □

So now because A is an ideal, and its elements are in R (a Noetherian ring), A is generated by $(a_1, \cdots, a_k)$, where each a_i is the leading coefficient of some polynomial $f_i \in I$ for each i. Multiply all the f_is by powers of x so that the degrees are all the same: let's say this degree is m.

Now our f_is don't quite generate our set, but we can still do things with this new idea! We're going to do something similar to the division algorithm – for any polynomial $g \in I$ which is not equal to 0, we can write it as

$$g = bx^n + \cdots.$$

Since b is a leading coefficient, it is in A, and thus we can write $b = \sum a_i r_i$ for some $r_i \in R$. Now if $n \geq m$, we know that

$$h = \sum_{i=1}^{k} x^{n-m} f_i r_i$$

is an element of I with degree equal to n and leading coefficient equal to that of g, because $x^{n-m}f_i$ is a polynomial of leading coefficient a_i and degree n by construction. This means that the difference $(g - h)$ has degree less than n.

We can repeat this arbitrarily until g has degree less than m: the idea is that all polynomials g in our ideal I are a linear combination of our f_is, except potentially with a remainder of degree less than m. So the division algorithm gives us a combination $\sum p_i f_i$ in the ideal generated by $(f_1, \cdots, f_k)$, and now we just need to finitely generate elements of the form $g - \sum p_i f_i$: these are some polynomials of degree less than m.

Lemma 185

If R is a Noetherian ring and V is a finitely generated R-module, then every submodule of V is finitely generated.

Proof of lemma. If V is a finitely generated R-module and U is a submodule, we have surjective maps from $R^m \to V$ and $U' \to U$, where U' is a subset of R^m. Since these are surjective homomorphisms, it is enough to show that every submodule of R^m is finitely generated.

We'll use induction. The base case $m = 1$ is true by definition of a Noetherian ring. Now suppose $V = R^m$, and let π be the homomorphism from R^m to R^{m-1} that **drops the last component**. If we have our surjective map $\pi' : U \to \overline{U}$, where $U \subset R^m$ and $\overline{U} \subset R^{m-1}$, we know that $\overline{U}$ is finitely generated by induction, and now the kernel K is a subset of $\ker \pi$. $\ker \pi$ is isomorphic to R (because we just drop the last element), which is a Noetherian ring, so K must be finitely generated by because it's an ideal! Thus $\overline{U}$ and therefore U are finitely generated. □

So now, the polynomials of degree less than m, plus the zero polynomial, form a free R-module with basis $(1, x, \cdots, x^{m-1})$. $I \cap P_m$, which is the polynomials with degree less than m in I (plus zero) is a submodule, and thus it is finitely generated by some set $(g_1, \cdots, g_\ell)$, and now that set plus our original polynomials $f_1, \cdots, f_k$ generate I, as desired, finishing our proof. □

Fact 186

Hilbert proved many things in 1895. The paper's in German, though...

23 April 10, 2019

(We got more cookies in class.) We're going to finish talking about modules today and move on to fields on Friday. Recall that in a **finitely-generated** R-module V, there exist elements $(v_1, \cdots, v_k)$ such that every element $v \in V$ can be written as a linear combination $\sum r_i v_i$. We found that a submodule of R (as a module) is just an ideal, and we defined the notion of a **Noetherian ring** to be one where every ideal is finitely generated. This led us to the Hilbert Basis Theorem, which says that $R[x]$ is Noetherian for any Noetherian ring R. Applying this repeatedly, we know that $F[x_1, \cdots, x_n]$ is always Noetherian for any field F, and so is $\mathbb{Z}[x_1, \cdots, x_n]$.

Proposition 187

If we have a surjective homomorphism $\phi : R \to R'$, and R is a Noetherian ring, then so is R'.

So any ring that is a quotient of a polynomial ring over a field is also Noetherian – this is basically all of the rings we've been looking at so far! Note that here our ring $R' \cong R/\ker \phi$.

Proof. If I' is an ideal of R', consider the inverse image $\phi^{-1}(I')$, which is an ideal of R. Since R is Noetherian, I is finitely generated by some elements $(v_1, \cdots, v_k)$, and all w in I are combinations of the form $w = \sum r_i v_i$. So now the images of the v_i will generate I': if $w' \in I'$, we can choose $w \in I$ such that $\phi(w) = w'$ by surjectivity, and then $w = \sum r_i v_i \implies w' = \sum r_i' v_i'$ (here $r_i' = \phi(r_i)$), and thus the $v_i' = \phi(v_i)$ generate I'. Therefore every ideal of R' is finitely generated, and thus R' is Noetherian. □

So to find a non-Noetherian ring, we need to add infinitely many elements, and we'll do that now:

Example 188

Take R to be the ring $R = F[x_1, x_2, \cdots]$.

Any particular polynomial only uses finitely many monomials (because it's a finite combination of monomials), but the ring can contain infinitely many monomials.

This is not Noetherian: let I be the set of polynomials that evaluate to 0 at 0. Equivalently, we can say that this is the set of polynomials with no constant term. But the ideal contains $x_1, x_2, \cdots$, and it can't be finitely generated because we have infinitely many variables and each generator can only use finitely many of them.

So now let V be a finitely generated R-module, where R is a Noetherian ring. Remember that if $\overline{v} = (v_1, \cdots, v_m)$ generate V, we have a homomorphism $\phi : R^m \to V$ sending a vector $x \to \overline{v}x$. We can get every vector in V (because this map is surjective by definition of $\overline{v}$ generating V). Letting $W = \ker \phi$, W is now a submodule of R^m, so it is finitely generated as well (by Lemma 185). Basically, it's enough to show that a submodule of the free module is finitely generated, and we use induction.

So now if we choose our generators $\overline{w} = (w_1, \cdots, w_n)$ and consider our surjective homomorphism $R^n \to W \subset R^m$, we can compose the map $R^n \to W$ and the inclusion $W \to R^m$: then this is represented by an $m \times n$ matrix with entries in R. So now $W = AR^n$, and by the first isomorphism theorem, V is isomorphic to $R^m / \ker \phi = R^m / W$, and thus our finitely generated R-module looks like

$$V \cong R^m / (AR^n).$$

We refer to this by saying that A **presents** V: sometimes we can simplify our matrix A quite a bit, especially if we have the division algorithm (in which case we can diagonalize with base-change matrices). Looking more in detail at the basechange matrices, recall that we can replace $A \to Q^{-1}AP$, where Q is an invertible $m \times m$ R-matrix and P is an invertible $n \times n$ R-matrix. Specifically, if $R = \mathbb{Z}$ or $R = F[t]$, we can make our matrix $Q^{-1}AP$ diagonal. This led us to the Basis Theorem for abelian groups: **every finitely generated abelian group is the product of (possibly infinite) cyclic groups.**

Also recall that if we have a ring $R = F[t]$, an R-module V is just an F-vector space with a linear operator $T : V \to V$, where we send $T(v) = tv$ (so that we know what happens to t, t^2, and so on). So let A be an $m \times n$ matrix with entries in R, and let's say we can make it diagonal: we write

$$A = \begin{pmatrix} D & 0 \\ 0 & 0 \end{pmatrix},$$

where D has k nonzero diagonal entries $f_1 | f_2 | \cdots$ that successively divide into each other. So now if we present V as $R^m / (AR^n)$, let's think about what V looks like: we have generators $V = (v_1, \cdots, v_m)$, and now each column vector gives us a restriction on our generators. Since each column vector only contains one element (because we diagonalized), we have

$$V = V_1 \times \cdots \times V_k \times V_{k+1} \times \cdots \times V_m,$$

where $V_i \cong R/(f_i R)$ for all $1 \leq i \leq k$, and $V_{k+1}, \cdots, V_m \cong R$ (there are no relations).

Definition 189

A **cyclic** R-module V is one that is generated by a single element.

In such a module, we have a map $R \to V$ which sends $1 \to v$: this is a surjective homomorphism with some kernel I, so by the first isomorphism theorem, $V \cong R/I$.

In the case where $R = F[t]$ for a field F, we have $I = fR$ for some $f \in F[t]$ (because we have a principal ideal domain), so we must have $V \cong R/(fR)$.

Theorem 190 (Structure theorem for modules over a polynomial ring)
Let $R = F[t]$ for a field F. Then every finitely generated R-module is a product of cyclic R-modules.

This is the exact analog of the structure theorem for abelian groups! Let's translate this statement to one about linear operators: we know that we have a linear operator $T : V \to V$. As a module, V is generated by one element v, so every $w \in V$ is of the form $w = g(t)v$ for some polynomial $g(t)$. Writing $g(t) = t^k + b_{k-1}T^{k-1} + \cdots + b_0$, we have $w = t^k v + b_{k-1}t^{k-1}v + \cdots + b_0 v$. Since our linear operator T corresponds to multiplying by t, this can also be written as

$$w = T^k v + b_{k-1}T^{k-1}v + \cdots + b_0 v,$$

so $v, Tv, T^2v, \cdots$ span V. If V is isomorphic to $R/(0)$, then there's no relations: V has an F-basis $v_0 = v, v_1 = Tv, v_2 = T^2v, \cdots$, and T is the shift operator between the basis elements. On the other hand, if V is isomorphic to $R/(fR)$ for some polynomial $f = t^n + a_{n-1}t^{n-1} + \cdots + a_0$, we have the relation

$$v_n + a_{n-1}v_{n-1} + \cdots + a_0 v_0 = 0,$$

and thus we have an F-basis $v_0, \cdots, v_{n-1}$. Then the matrix of T has 1s directly below the diagonal and $-a_0, -a_1, \cdots, -a_{n-1}$ in the last column. (This is called the **rational canonical form** of the matrix.)

So looking at our structure theorem, if k is equal to the number of generators (so we don't have infinitely-dimensional vector spaces), V is the product of finite-dimensional vector spaces with dimension equal to the degree of the corresponding f_is.

Corollary 191
If V is a finitely-dimensional vector space, and $T : V \to V$ is a linear operator, then there is a basis for V such that the matrix M for T has the form

$$M = \begin{pmatrix} B_1 & 0 & 0 \\ 0 & B_2 & 0 \\ 0 & 0 & \ddots \end{pmatrix}$$

where each B_i is in rational canonical form.

This is the best we can do over an arbitrary field! (Jordan normal form is nicer for $\mathbb{C}$.)

Fact 192
Professor Artin recommends that we don't take pictures of the whiteboard. There is a connection between the mind and hand, so we should take notes on paper instead. (Hmm... maybe I should stop taking notes on a laptop.)

24 April 12, 2019

We're going to discuss **fields** today – most of the discussion centers around containing one field in another one.

Definition 193
If we have fields F, K, and $F \subset K$, then K is a **field extension** of F.

Example 194

We can take $F = \mathbb{Q}$ and let $K = F[\delta]$, where $\delta^2 = 2$. Then elements of K are of the form

$$K = \{a + b\delta : a, b \in F\}.$$

We can think of K as a vector space over the field F: in this case, it has dimension 2. (Addition and scalar multiplication by F in K are just the standard addition and multiplication, since elements in F are in K.)

Definition 195

The dimension of K as an F-vector space is called the **degree** of K over F, denoted $[K : F]$.

Some examples:

$$[\mathbb{C} : \mathbb{R}] = 2, \quad \mathbb{C} \text{ has basis } (1, i),$$

$$[\mathbb{Q}[\sqrt{2}] : \mathbb{Q}] = 2, \quad \mathbb{Q}[\sqrt{2}] \text{ has basis } (1, \sqrt{2}).$$

Here's one thing we can do this: let K/F be a field extension, and let $\alpha \in K$. We can then map polynomials

$$\phi : F[x] \to K$$

by sending elements of the field to themselves and sending x to α. This helps us analyze "part" of K: the image of ϕ here is all elements of the form

$$\{\beta \in K : \beta \text{ can be written as a polynomial in } \alpha \text{ with coefficients in } F\}.$$

On the other hand, the kernel ϕ is all polynomials $g(x)$ such that $g(\alpha) = 0$: this is an ideal, and since $F[x]$ is principal, this means it is a principal ideal. There are two possibilities: if the kernel is trivial (only the zero polynomial), then α is **transcendental**. That means it's not the root of any polynomial $g(x)$ with F-coefficients! On the other hand, if the kernel is generated by some polynomial f, then f is **irreducible**: if we could write it as $f = gh$, then $f(\alpha) = g(\alpha)h(\alpha) = 0$, and then either g or h would be a generator with lower degree. Therefore, the ideal (f) is a maximal ideal, and we'll study this case in more detail.

So going back to the image, if α is transcendental, our map ϕ is injective, and then K is just isomorphic to $F[x]$. However, if α is not transcendental, by the first isomorphism theorem, $F[x]/(f)$ is isomorphic to the image of ϕ, and since (f) is maximal, this is a field. Then the image of ϕ, which is $F[\alpha]$, is the set of elements $\beta \in K$ that can be written as a polynomial in α.

Question 196. *How do we compute in $F[\alpha]$, which is isomorphic to $F[x]/(f)$, for some irreducible polynomial f?*

We can divide out by the leading coefficient, since F is a field, so let's say

$$f = x^n + a_{n-1}x^{n-1} + \cdots + a_0.$$

Since $\phi(f) = 0$,

$$f(\alpha) = \alpha^n + a_{n-1}\alpha^{n-1} + \cdots + a_0 = 0,$$

and that means we can write α^n as a linear combination of $1, \alpha, \cdots, \alpha^{n-1}$:

Proposition 197

$F[\alpha]$ has an F-basis

$$(1, \alpha, \alpha^2, \cdots, \alpha^{n-1}).$$

if the minimal polynomial f has degree n.

So if we are computing in $F[x]$, we have combinations of basis elements with F-coefficients: adding is done component-wise, and if we have polynomials

$$g(x) = b_0 + b_1 x + \cdots + b_{n-1}x^{n-1},$$

$$h(x) = c_0 + c_1 x + \cdots + c_{n-1}x^{n-1},$$

then $g(\alpha)$ and $h(\alpha)$ are arbitrary elements of $F[\alpha]$. Then we multiply $g(\alpha)$ and $h(\alpha)$ by using the division algorithm: we can first write

$$g(x)h(x) = f(x)q(x) + r(x),$$

where $r(x)$ is a polynomial with degree less than n, and then say that

$$g(\alpha)h(\alpha) = f(\alpha)q(\alpha) + r(\alpha) = r(\alpha).$$

Notably, we can always compute inverses in $F[\alpha]$ because we have a field, but this isn't immediately obvious from the form of our elements!

Usually when we have an extension K of F with finite degree, we can generate it with **one element**, but the proof is tricky: we'll defer it to later. Let's think about how a construction of a field extension can be done more abstractly: we start with a field F, and we find some irreducible polynomial $f \in F[x]$. This forms a field $K = F[x]/(f)$ because (f) is maximal, and if the residue of x is α, then $K = F[\alpha]$.

Example 198

Take $F = \mathbb{F}_2$, the field with two elements $\{\overline{0}, \overline{1}\}$.

Recall that $x^3 + x + 1$ is irreducible in $\mathbb{F}_2[x]$, because it doesn't have a root. So now $K = F[x]/(f)$ has an F-basis $(1, \alpha, \alpha^2)$, and we know that $\alpha^3 = -\alpha - 1 = \alpha + 1$ in our field.

Let's find the inverse of $\alpha^2 + 1$. This means we want to find a polynomial such that

$$(1 + \alpha^2)(c_0 + c_1\alpha + c_2\alpha^2) = 1.$$

This means (using $\alpha^3 = \alpha + 1$),

$$c_0 + c_1\alpha + c_2\alpha^2 + c_0\alpha^2 + c_1(\alpha + 1) + c_2(\alpha^2 + \alpha) = 1,$$

so equating coefficients on both sides,

$$(c_0 + c_1) + (c_1 + c_1 + c_2)\alpha + (c_2 + c_0 + c_2)\alpha^2 = 1 + 0\alpha + 0\alpha^2.$$

So $c_0 = 0, c_1 = 1, c_2 = 0$: indeed

$$\alpha(\alpha^2 + 1) = \alpha^3 + \alpha = 1.$$

(We probably didn't need to do all that to find the answer...)

The degree of our field extension has two main properties:

Theorem 199

If $K = F[\alpha]$, then the degree $[K : F]$ is the degree of the irreducible (minimal) polynomial of α. Also, if we have fields $F \subset K \subset L$, then

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

Example 200

Let $F = \mathbb{Q}, K = F[\delta]$ where $\delta^2 = 2$, and $L = K[\varepsilon]$, where $\varepsilon^3 = 3$.

Then the degree of L over F is

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

$x^3 - 3$ is irreducible over F (for example, by Eisenstein). What's the irreducible polynomial of ε over K? We know that K contains F, so $x^3 - 3$ is a polynomial with coefficients in K, and it has ε as a root. The set of all polynomials with ε as a root is a principal ideal, so if the minimal polynomial were f, f would need to divide $x^3 - 3$. But $x^3 - 3$ doesn't have any roots in K, either: we could for example, just write $\varepsilon = a + b\delta$ and expand. If $x^3 = 3$ had a root, then we must have

$$3 = \varepsilon^3 = a^3 + 3a^2b\delta + 6ab^2 + 2b^3\delta \implies 3a^2b + 2b^3 = 0, a^3 + 6ab^2 = 3.$$

We can amuse ourselves by showing there aren't any solutions here for $a, b \in \mathbb{Q}$, so $[L : K] = 3$, and therefore

$$[L : F] = [L : K][K : F] = 3 \cdot 2 = 6.$$

So what's an F-basis for our field extension L? We know that

$$L = K[\varepsilon] = F[\delta, \varepsilon].$$

Since δ has degree 2 and ε has degree 3, a natural basis to consider is $\{1, \delta, \varepsilon, \delta\varepsilon, \varepsilon^2, \delta\varepsilon^2\}$. It's not hard to prove that this works!

In general, if we have F, K, L such that $[L : K] = n$ and $[K : F] = m$, let $(\alpha_1, \cdots, \alpha_m)$ be a F-basis for K, and let $(\beta_1, \cdots, \beta_n)$ be a K-basis for L. We claim the set of $\{\alpha_i\beta_j\}$s is an F-basis for L, and this isn't too hard to check: for any $\gamma \in L$, we can write it as

$$\gamma = \sum_j \eta_j\beta_j$$

for $\eta_j \in K$, and then we can write each η_j out as

$$\eta_j = \sum_i c_{ij}\alpha_i,$$

for $c_{ij} \in F$. So now we just have

$$\gamma = \sum_{i,j} c_{ij}\alpha_i\beta_j,$$

as desired! Showing that this is unique is not too hard: basically, each step was unique, so our entire process is unique.

This has a nice corollary which we'll illustrate now:

Example 201
Is $\delta = \sqrt{2}$ in the field $\mathbb{Q}[\omega = \sqrt[5]{7}]$?

We know that ω is a root of $x^5 - 7$, which is irreducible in $\mathbb{Q}$, so the degree of $\mathbb{Q}[\omega]$ is 5. If $\delta \in \mathbb{Q}[\omega]$, then we'd have $\mathbb{Q} \subset \mathbb{Q}[\delta] \subset \mathbb{Q}[\omega]$: this can't happen because 2 doesn't divide 5. This kind of argument doesn't always work, but it's pretty efficient when it does!

25 April 17, 2019

Recall that if we have fields $F \subset K$, then K is called a **field extension** of F. This is useful when we think of the **degree** of the field extension $[K : F]$, which is the dimension of K as an F-vector space. Specifically, if we have $F \subset K \subset L$, then

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

One way to look at field extensions is to take some element $\alpha \neq 0 \in K$ and consider the map

$$\phi : F[x] \to K$$

which sends $x \to \alpha$. If the kernel of ϕ is not just zero, then α is called **algebraic** (otherwise it is **transcendental**), and since we have a principal ideal domain, $\ker \phi = (f)$ for some polynomial f. (f) is a prime ideal generated by a monic irreducible polynomial $f(x)$, so (f) is a maximal ideal. Thus, by the first isomorphism theorem,

$$F[x]/(f) \cong F[\alpha] \subset K$$

forms a field, and $F[\alpha]$ has an F-basis of $(1, \alpha, \cdots, \alpha^{n-1})$, where the degree of f is n. This means that the degree of α over F is (by definition) the degree of the minimal polynomial f. So we have a chain

$$F \subset F[\alpha] \subset K,$$

where the degree $[K : F]$ is divisible by the degree $[F[\alpha] : F] = n$.

Corollary 202
For **any** element $\alpha \in K$, the degree of α over F divides the degree of K over F.

Specifically, if $[K : F] = p$ is prime, $[F[\alpha] : F]$ is either 1 or p. If it is 1, that's the same as saying that α is actually an element of F. Otherwise, we actually have $K = F[\alpha] - \alpha$ gives us the whole field extension.

Example 203
Take $F = \mathbb{Q}$. $x^5 - 2$ is an irreducible polynomial by Eisenstein, so if α is some complex root of $x^5 - 2$, then we must have $[F[\alpha] : F] = 5$. In addition, any $\beta \in F[\alpha]$ (which isn't an element of F) is the root of an irreducible polynomial of degree 5.

The "irreducible" part of this isn't immediately obvious! For example, if we took $\beta = 1 + \alpha^2$, we can write $\beta^2, \beta^3, \beta^4, \beta^5$ in terms of the basis elements $(1, \alpha, \alpha^2, \alpha^3, \alpha^4)$. These six equations are dependent, but it's not immediately clear that the resulting polynomial in β is irreducible.

Question 204. *What if $f(x)$ is an irreducible polynomial of degree 4 which is irreducible over $F = \mathbb{Q}$?*

We'll start the argument the same way: take α to be a root of f. The field extension has degree $[F[\alpha] : F] = 4$, which is not a prime. Is there a field L such that $F \subset L \subset K$ where $[K : L] = [L : F] = 2$? The answer actually depends on f, and it's a hard question, so we won't answer it here (because we need Galois theory). We'll just do an example where such a field does actually exist:

Example 205

Let $\zeta = e^{2\pi i/5}$ be a fifth root of unity. Over $\mathbb{Q}$, its irreducible polynomial is

$$\frac{x^5 - 1}{x - 1} = x^4 + x^3 + x^2 + x + 1.$$

This has roots $\zeta, \zeta^2, \zeta^3, \zeta^4$. Now ζ is not a real number, but $\zeta + \zeta^4$ is! Letting this be α, $F[\alpha]$ is smaller than the whole field (because it has only real numbers), and also α is not rational. Thus $F < F[\alpha] < F[\zeta]$, and therefore we must have $[F[\zeta] : F[\alpha]] = [F[\alpha] : F] = 2$.

How do we find the irreducible polynomial for α in $\mathbb{Q}$? The most straightforward way is to write down the powers and write down a linear relation between them. Turns out we won't have to go far:

$$\alpha^0 = 1$$

$$\alpha^1 = \zeta + \zeta^4$$

$$\alpha^2 = \zeta^2 + \zeta^8 + 2$$

and now note that $\zeta + \zeta^2 + \zeta^3 + \zeta^4 = -1$, so

$$\alpha + \alpha^2 = -1 + 2 = 1.$$

That means α is a root of the polynomial $x^2 + x - 1$, and $\alpha = \frac{-1+\sqrt{5}}{2}$ (it's the positive root because ζ and ζ^4 have positive real parts). Specifically, our field $F[\alpha] = F[\sqrt{5}]$.

With that, we'll move on to a related topic: **constructions with a ruler and compass**. Here's some rules:

- Our ruler is actually a straightedge (it doesn't have inches or even centimeters).
- We start off with two points given in the plane that are "constructed."
- Given two points p_0, p_1 that are constructed, we can draw a line $L(p_0, p_1)$ through them.
- We can also draw a circle $C(p_0, p_1)$, which is a circle with center p_0 that passes through p_1.
- Intersections are also constructed points.

This is all we're allowed to do! What can we construct this way? First of all, there are some elementary constructions that we might have learned in high school.

Example 206

If we construct some line L and some point p, can we construct the perpendicular line to L passing through p?

We don't want to pick arbitrary points on our line – that's not so elegant. But we know that if L is constructed, there are two points on it: one is not directly below p, so we can just draw the circle passing through it with center

p. Now we can just draw circles through those points (with the other point as center) and get another point on the perpendicular bisector: connect them and we're done!

Similarly, we can construct a parallel line by doing two perpendiculars.

Example 207

We have two points ℓ apart: how can we construct a point on a line ℓ away from a given point?

If we have P_1 and P_2 off the line and A on the line, construct parallelogram AP_1P_2B (by using parallels): now $AB = \ell$, and draw a circle through point B with center A.

To make more progress, we'll introduce Cartesian coordinates! We can assume our starting points are $(0, 0)$ and $(1, 0)$ – we can draw our x and y-axes with this by using the perpendicular trick.

Definition 208

A real number z is **constructible** if we can construct points p_0, p_1 that are a distance $|z|$ apart in our coordinates.

Lemma 209

$p_0 = (x_0, y_0)$ is constructible if and only if x_0 and y_0 are constructible real numbers.

Think of this as constructing $(x_0, 0)$ and then (x_0, y_0).

Lemma 210

Suppose p_0, p_1 are points with coordinates in a field K. Then $L(p_0, p_1)$ and $C(p_0, p_1)$ have equations with coefficients in K:

$$ax + by + c = 0, (x - x_0)^2 + (y - y_0)^2 = r^2,$$

where $a, b, c, x_0, y_0, r \in K$.

Proof. The equation of a line through (x_0, y_0) and (x_1, y_1) is (by point-slope)

$$(x_1 - x_0)(y - y_0) = (y_1 - y_0)(x - x_0)$$

which has coordinates in K because we have multiplicative inverses. Similarly, the circle has

$$(x - x_0)^2 + (y - y_0)^2 = r^2 = (x_1 - x_0)^2 + (y_1 - y_0)^2.$$

□

So we draw a bunch of circles and lines: the equation's coefficients are still in the field that we're in each time we do such a construction. But what do we know about the new coordinates of the intersection points?

Lemma 211

If lines and circles have equations with coefficients in a field K, the intersection points have coordinates in K or a **quadratic field extension** of K.

Proof. Solve the equations! If solutions don't exist, we don't care, since that just means we don't add constructible points.

Two linear equations have a solution in the field by linear algebra (if one exists). If we have a line and a circle, then

$$ax + by + c = 0, (x - x_0)^2 + (y - y_0)^2 = r^2$$

can be solved by substitution: write y as a linear expression in x, which gives a quadratic equation for x (and y is linear in x). Any solution, if it exists, is in a quadratic field extension.

Finally, we might need to intersect two circles. Usually this looks like a problem, because we should have a fourth-degree equation when we have two second-degree equations. But that doesn't happen here: if

$$(x - x_0)^2 + (y - y_0)^2 = r_0^2, \quad (x - x_1)^2 + (y - y_1)^2 = r_1^2,$$

we can just subtract the two to get a linear equation. Now this can be solved in the same way as the line-and-circle case, so we again have a quadratic field extension. □

So as a consequence, starting with the rational numbers and trying to make more constructible objects only gives us **quadratic field extensions** of $\mathbb{Q}$. (Note that if any of our equations give no solutions or complex solutions, this just means, geometrically, that our objects don't intersect, so we don't care about them.) We can state that as a formal result:

Theorem 212

A real number α is a constructible real number if there exists a chain of fields

$$\mathbb{Q} = K_0 \subset K_1 \subset \cdots \subset K_N \subset \mathbb{R},$$

where the degree $[K_{i+1} : K_i] = 2$ and $\alpha \in K_N$.

Question 213. *Given an angle γ, is it possible to construct $\theta = \frac{\gamma}{3}$?*

Given any angle, we can move it over to the coordinate axes, so that we have one of the rays on the x-axis. Our goal is then to construct a ray at angle $\frac{\theta}{3}$. Clearly it's possible sometimes (for example, we can construct the angle $30°$), but it turns out to not be true in general:

Proposition 214

It is not possible to trisect a 60 degree angle.

Proof. Note that if we could construct an angle $\theta = 20°$, we would also be able to construct $\alpha = 2\cos\theta = e^{i\theta} + e^{-i\theta}$. By the triple-angle identity or other methods, α will be a root of the polynomial $x^3 - 3x - 1$.

This is an irreducible polynomial over $\mathbb{Q}$, since it doesn't have an integer root. But the degree of $[K_N : K_0]$ must be a power of 2, and the degree of $\mathbb{Q}[\alpha] : \mathbb{Q}$ is 3, which is a contradiction. □

26 April 19, 2019

Today, we'll start by talking about adjoining roots to a field F. Let's say we have an irreducible polynomial $f(x) \in F[x]$, and let's say we want to ask for a field extension K such that $K = F[\alpha]$, where $\alpha \in K$ is a root of f. If our underlying field F is a subfield of $\mathbb{C}$, we can just pick a root $\alpha \in \mathbb{C}$ by the fundamental theorem of algebra. But there are fields that aren't subfields of $\mathbb{C}$: for example, what if $F = \mathbb{F}_p$ or $\mathbb{C}(t)$ (the fraction field of polynomials)?

Well, if we have an α in our extension field K, we can construct a map $\pi : F[x] \to K$, where F is sent to itself and x is sent to α. Then if α is a root of f, the kernel of π will be the principal ideal (f), and by the first isomorphism theorem, **we have our field without needing to define things in terms of α itself**:

Proposition 215

Let α be a root of an irreducible polynomial f. Then $F[x]/(f) = F[\alpha]$.

In particular, f being irreducible implies that (f) is a maximal ideal, so $K = F[\alpha]$ is a field!

Fact 216

The residue of x in $F[x]/(f)$, $\pi(x) = \alpha$, is a root of f.

Proof. Consider the maps

$$F \xrightarrow{\text{id}} F[x] \xrightarrow{\pi} F[x]/(f) = K.$$

The composition here is injective: if a, b both map to the same $k \in K$, then $a - b$ must be a multiple of f, which means $a = b$. Our goal is then to check that $\alpha = \pi(x)$ is a root of f. Note that $\pi(f(x)) = 0$ (since π mods out by f). On the other hand, if $f(x) = x^n + a_{n-1}x^{n-1} + \cdots + a_0$, where the coefficients are in F, we have

$$0 = \pi(f) = \pi(x)^n + \pi(a_{n-1})\pi(x)^{n-1} + \cdots + \pi(a_0).$$

π of any coefficient is just that coefficient itself (since we identify F with its image in K), and $\pi(x) = \alpha$. Plugging everything in, we find that

$$0 = \alpha^n + a_{n-1}\alpha^{n-1} + \cdots + a_0,$$

and thus α is a root, as desired. □

It's important to note that in $K = F[\alpha]$, elements are polynomial expressions in α, and it has an F-basis $(1, \alpha, \cdots, \alpha^{n-1})$. Then the relation $f(\alpha) = 0$ is the "only one" that is relevant.

Example 217

Let $F = \mathbb{Q}$, and let's take $\zeta = e^{2\pi i/7}$, a seventh root of unity.

This is a root of the polynomial

$$\frac{x^7 - 1}{x - 1} = x^6 + x^5 + \cdots + x + 1.$$

Thus, $F[\zeta]$ has basis $(1, \zeta, \cdots, \zeta^5)$, and we have the relation

$$\zeta^6 + \zeta^5 + \cdots + \zeta + 1 = 0$$

as the only linear relation between these powers of ζ. Note that the roots of $f(x) = 0$ are $\zeta, \zeta^2, \cdots, \zeta^6$, so they're all roots of this irreducible polynomial Φ_p. Plugging in $\gamma = \zeta^2$, we also have

$$\gamma^6 + \gamma^5 + \cdots + \gamma + 1 = 0,$$

but this is the same relation as the original one if we plug things back in ($\zeta^5 + \zeta^3 + \zeta + \zeta^6 + \zeta^4 + \zeta^2 + 1 = 0$). So from the point of view of the rational numbers, there's really no difference between γ and ζ.

Fact 218

Both of the following follow from the logic above:

- If α is a root of an irreducible polynomial f in K, and α' is a root of that same f in a field K', then there exists a unique ***F*-isomorphism** $\phi : K \to K'$ (which means that ϕ is the identity on F), sending α to α'.
- On the other hand, if we have an F-isomorphism $K \to K'$, and α is a root of an irreducible $f(x) \in F[x]$, then $\phi(\alpha) = \alpha'$ is also a root of f (by applying ϕ to the polynomial relation for α).

Proposition 219

Take any (not necessarily irreducible) polynomial $g(x) \in F[x]$. Then there exists a field extension K in which g factors into linear factors: that is, g **splits completely**.

Proof. Induct on the degree of g. Choose an irreducible factor $f|g$, and adjoin a root α_1 of f to F by the method above. Now $F_1 = F[\alpha_1]$. $x - \alpha_1$ is a factor of f, which is a factor of g, so now we have $g = (x - \alpha_1)g_1$. The degree of g_1 is smaller than g, and we're done by induction. □

Concretely, we'll get some chain

$$F \subset F[\alpha_1] \subset F[\alpha_1, \alpha_2] \subset F[\alpha_1, \cdots, \alpha_n] = K.$$

The degree of the first extension is at most n, and the degree of the second one is at most $n-1$, and so on. The last one is free: if we have $n-1$ of the roots, we get the last one with a degree 1 extension, which means we stay in the field itself. This leads us to the following result:

Proposition 220

For any **splitting field** K of a degree n polynomial over F, $[K : F] \leq n!$.

This is not a very good bound, but it's the best we can do.

Example 221

Let's take $F = \mathbb{Q}, f(x) = x^3 - 2$. We can use our complex roots $\alpha = \sqrt[3]{2}, \alpha_2 = \omega\alpha_1, \alpha_3 = \omega^2\alpha_1$, where $\omega = e^{2\pi i/3}$ is a cube root of unity.

Note that $[F[\alpha_1] : F] = 3$, and then the polynomial we're left with is

$$x^3 - 2 = (x - \alpha_1)q(x),$$

where $q(x)$ is a quadratic with complex roots α_2, α_3. Now $[F[\alpha_1, \alpha_2] : F[\alpha_1]] \leq 2$: in this case, α_1 is a real number, and ω is not real, so adjoining α_2 is not free. So here we have

$$[K : F] = 3 \cdot 2 \cdot 1 = 6.$$

So if f is a cubic polynomial with rational coefficients, it always has at least one real root. The only case in which it might have $[K : F] = 3$ is if all three roots are real, and even then it still depends on the specific polynomial.

Example 222

Let $F = \mathbb{F}_2 = \{0, 1\}$ be the finite field on 2 elements; recall that the irreducible cubics are $f = x^3 + x + 1$ and $f' = x^3 + x^2 + 1$.

Let's adjoin a root of $f = x^3 + x + 1$: we have $[F[\alpha] : F] = 3$, so $F[\alpha]$ is a vector space of dimension 3 over F, and thus

$$K = |F[\alpha]| = 2^3 = 8.$$

The elements are then

$$\{0, 1, \alpha, 1 + \alpha, \alpha^2, 1 + \alpha^2, \alpha + \alpha^2, 1 + \alpha + \alpha^2\}.$$

Let β be any of the elements that is not 0 or 1. Note that adjoining β to F yields some subfield of K. But the degrees satisfy

$$[K : F] = 3 = [K : F[\beta]][[F[\beta] : F],$$

and since $\beta \notin F$, we must have $K = F[\beta]$. Thus β must be a root of either $x^3 + x + 1$ or $x^3 + x^2 + 1$.

Well, we have six elements that aren't 0 and 1, and there are two irreducible cubic polynomials: this means that f, f' **both split completely** over K, and each of them has three of these elements as roots.

How do we find which ones belong to $f = x^3 + x + 1$? Note that α, a root of f, satisfies $\alpha^3 = \alpha + 1$. Since we have $(a + b)^2 = a^2 + b^2$ in $\mathbb{F}_2$, squaring both sides of our relation, we get

$$\alpha^6 = \alpha^2 + 1.$$

This means that $\beta = \alpha^2 \implies \beta^3 = \beta + 1$, which means that $\alpha, \alpha^2, \alpha + \alpha^2$ must be the roots of f.

Example 223

Let's go back to our field $L = F[\zeta]$, where ζ is a seventh root of unity.

Taking $\alpha = \zeta + \zeta^6$, a real number, we have the chain

$$F \subset F[\alpha] \subset L.$$

$[L : F] = 6$, and $[L : F[\alpha]] \geq 1$ because $F[\alpha]$ contains only real numbers. Furthermore,

$$(x - \zeta)(x - \zeta^6) = x^2 - \alpha x + 1,$$

so $[L : F[\alpha]] = 2$, and $[F[\alpha] : F] = 3$. To find the minimal polynomial for α, We can compute powers of α until we get a relation:

$$\alpha^0 = 1, \alpha^1 \qquad = \zeta + \zeta^6, \alpha^2 = \zeta^2 + 2 + \zeta^5, \alpha^3 \qquad = \zeta^3 + 3\zeta + 3\zeta^6 + \zeta^4,$$

and now we have the relation

$$\alpha^3 + \alpha^2 - 2\alpha - 1 = 0,$$

so α is a root of $x^3 + x^2 - 2x - 1$. To find the other roots from here, note that we can let $\gamma = \zeta^2$: then this has all the same properties as ζ from the perspective of the rationals, so $\alpha_2 = \gamma + \gamma^6 = \boxed{\zeta^2 + \zeta^5}$ is a root as well, and so is $\alpha_3 = \boxed{\zeta^3 + \zeta^4}$. If we don't believe this, we can always expand out $(x - \alpha_1)(x - \alpha_2)(x - \alpha_3)$:

$$= x^3 - (\alpha_1 + \alpha_2 + \alpha_3)x^2 + (\alpha_1\alpha_2 + \alpha_1\alpha_3 + \alpha_2\alpha_3)x - (\alpha_1\alpha_2\alpha_3).$$

The x^2-coefficient is $-(-1) = 1$, the x-coefficient can be done trickily: $\alpha_1\alpha_2$ is a sum of four powers of ζ, and so are the others. This means we have 12 terms, and so we'll have a -2 there (because the sum of the seventh roots of unity except ζ^0 is -1). Finally, $\alpha_1\alpha_2\alpha_3$ is eight powers, and it can include ζ^0. So it must be 2 copies of ζ^0 and 1 copy of all the others, so that yields 1. We've now sort of verified that this works!

27 April 22, 2019

We're going to discuss finite fields today. This is a fun topic, because it's actually pretty hard to satisfy all of the properties needed. Let's start with two preliminary ideas first.

Definition 224

Let $f(x) = a_n x^n + \cdots + a_0$ be any polynomial in $F[x]$, where F is a field. Then the **derivative** of f is

$$f'(x) \equiv n a_n x^{n-1} + (n-1)a_n x^{n-2} + \cdots + a_1.$$

Here, we must interpret $n = 1 + \cdots + 1$.

This satisfies the familiar calculus rules:

- $(f+g)' = f' + g'$,
- $(fg)' = f'g + g'f$,
- $(cf)' = cf'$.

Proposition 225

An element $\alpha \in F$ is a **double root** of a polynomial $f(x)$, which means that $(x-\alpha)^2$ divides f, if and only if α is root of both f and f'.

Proof. Clearly α needs to be a root if it is a double root. This means that $(x-\alpha)$ is a factor of f, so we can write

$$f(x) = (x-\alpha)q(x).$$

Now by the product rule,

$$f'(x) = q(x) + (x-\alpha)q'(x).$$

Suppose α is a root of f': this happens whenever $0 = f'(\alpha) = q(\alpha) + (\alpha-\alpha)q'(\alpha) = 0$, so $q(\alpha) = 0$. This means we can write

$$f = (x-\alpha)^2 r(x),$$

and thus α is a double root. □

Let's do an example: suppose $f(x) = x^n - 1$. The derivative is $f'(x) = nx^{n-1}$, and the only root of f' seems to be 0 (since we have a field), which is not a root of f. This means f has no double roots **unless the polynomial is identically zero**, which happens if $n = 0$ in our field! In particular, this means that $f = x^p - 1$ has derivative zero if $F = \mathbb{F}_p$. In particular,

$$x^p - 1 = (x-1)^p$$

by the Binomial theorem mod p, since $\binom{p}{k}$ is always a factor of p, and p is always odd except for $p = 2$ (which also works).

One other preparatory idea:

Theorem 226

Let F be a field, and let H be a finite subgroup of the multiplicative group $F^\times$. Then H is a cyclic group.

For example, the integers mod p always form a cyclic group under multiplication! Letting $p = 7$, $F = \mathbb{F}_p$, $H = F^\times$ gives a cyclic group of order 6: the generator is not $2^3 = 1$, but the powers of 3 are $\{3, 2, 6, 4, 5, 1\}$, so 3 generates the cyclic group $\mathbb{F}_7^\times$. On the other hand, the powers of 2 mod 13 are $\{1, 2, 4, 8, 3, 6, 12, 11, 9, 5, 10, 7\}$, so 2 is a generator in this case. Generators are called **primitive roots**.

It's not really well understood which elements are primitive roots, but this theorem tells us that there exists one (because we have a cyclic group).

Proof. By the basis theorem for abelian groups, every finitely generated abelian group is a product of cyclic groups. A finite group is obviously finitely generated (by its elements), and recall that the proof tells us when we diagonalize our integer matrix, we can make the entries divide each other. So that means

$$H \cong C_{d_1} \times C_{d_2} \times \cdots \times C_{d_k},$$

where $d_1 | d_2 | \cdots | d_k$. Let's count how many elements of order dividing d_1 there are in this group. There are d_1 ways to pick a representative from each of the C_{d_i}s, so there are d_1^k elements of order dividing d_1: thus, they are all roots of $x^{d_1} - 1$.

Lemma 227

In any field, a polynomial of degree d has at most d roots in F.

Proof of lemma. Take any root α_1: then $f(x) = (x - \alpha_1)g(x)$, and g has degree $d - 1$. By induction, g has at most $d - 1$ roots, so we're done by **unique factorization**. □

So now $d_1^k \leq d_1$, and therefore $k = 1$. □

We're ready to start talking about finite fields now! Let K be a finite field with $|K| = q$. We always have a unique homomorphism of the form

$$\varepsilon : \mathbb{Z} \to K.$$

$\ker \varepsilon$ is a principal ideal, since $\mathbb{Z}$ is a principal ideal domain, and now $\ker \varepsilon$ must be generated by a prime element. This means $\ker \varepsilon = p\mathbb{Z}$ for some integer prime p. Meanwhile, the image is isomorphic to $F \cong \mathbb{F}_p$, so K contains some F, and both of these are fields. Thus $[K : F]$ is some integer e, so the order of K is $|F|^e = p^e$, and that means $q = p^e$ **must be a prime power!**

It doesn't seem clear that $K = \mathbb{F}_q$ either exists or is unique at this point, though. Well, $K^\times$ is a cyclic group with order $q - 1$ by Theorem 226. Thus, the elements of $K^\times$ are roots of $x^{q-1} - 1$, and thus $x^{q-1} - 1$ splits completely! It's customary to multiply the factor $(x - 0)$ back in to get a more symmetric result:

Proposition 228

All elements of $K = \mathbb{F}_q$ are roots of the polynomial $x^q - x$, which splits completely in K.

Theorem 229

Let p be a prime, and let $e > 1$. Then there exists a field K with $|K| = p^e = q$, and all such fields are isomorphic.

Proof. We show existence first. Start with the field $F = \mathbb{F}_p$: there exists a field extension L of F where the polynomial $x^q - x$ splits completely by Proposition 219. Let K be the roots of $x^q - x$: we just need to show that $x^q - x$ doesn't have a double root. Taking the derivative,

$$f(x) = x^q - x \implies f'(x) = qx^{q-1} - 1 = -1,$$

because q is a multiple of p. This is never 0 in the field $F[x]$, so there are no double roots of $x^q - x$. So we do have $|K| = q$ elements, and now we must just show that if α, β are in our field, so are $\alpha \pm \beta, \alpha\beta$, and α^{-1}. We have that

$$(\alpha\beta)^q = \alpha^q \beta^q = \alpha\beta,$$

since $x^q = x$ for all elements in the field K, and thus $\alpha\beta$ is in K. Similarly,

$$(\alpha^{-1})^q = (\alpha^q)^{-1} = \alpha^{-1}.$$

Finally, $(\alpha + \beta)^q$ is a bit more difficult: we have to use induction. If $q = p^e = pk$, where $k = p^{e-1}$, then

$$(\alpha + \beta)^q = \left[(\alpha + \beta)^p\right]^k.$$

By the Binomial theorem, we can expand out $(\alpha + \beta)^p$ to $\alpha^p + \beta^p$, and now this is

$$[\alpha^p + \beta^p]^k$$

and now we're done by induction, since this is equal to $\alpha^{pk} + \beta^{pk} = \alpha + \beta$, as desired. Similarly, replace β with $-\beta$ to get that $\alpha - \beta$ is in our field as well.

Now we need to show that the finite field is **unique**. Let K_1, K_2 be two fields of order $q = p^e$: the multiplicative group $K_1^\times$ is cyclic of order $q - 1$, so we can let α_1 be a generator. Now the elements of K_1 are 0 and powers of α_1, so

$$K_1 = F[\alpha_1],$$

where α_1 is a root of an irreducible polynomial $f(x) \in F[x]$. Furthermore, the degree of f is the degree $[K_1 : F] = e$. But α_1 is also a root of $x^q - x$, and f is irreducible, so $f(x)$ **must divide** $x^q - x$.

Now looking in K_2, $x^q - x$ must split into linear factors, so f has a root α_2 in K_2. Thus

$$K_1 = F[\alpha_1] \cong F[x]/(f) \cong F[\alpha_2] \subset K_2.$$

But now the degree of K_2 over F is e by definition, and the degree of $F[\alpha_2]$ over F is the degree of f, which is also e. So $F[\alpha_2] = K_2$, and now K_1 and K_2 are isomorphic as desired. □

Here's the last main fact that we care about:

Proposition 230

If K is a finite field with order $q = p^e$ and $|K'|$ is a finite field with order $q' = p^{e'}$, then K contains a field isomorphic to K' if and only if $e' | e$.

This is a slight generalization of the idea that fields of the same order are isomorphic!

Proof. It's clear that if K contains K', then e' has to divide e: after all,

$$[K : F] = [K : K'][K' : F].$$

To show the other direction, we can write $e = e'd$. Then $u^{e'} - 1$ divides $u^e - 1$ (where u is a variable), since we can write $v = u^{e'}$, and then $v - 1$ divides $v^d - 1$. So now $q' - 1$ divides $q - 1$, since $q' = p^{e'}$ and $q = p^e$.

Now we can show that $x^{q'-1} - 1$ divides $x^{q-1} - 1$: this is true because $q' - 1$ divides $q - 1$, and we can apply the same logic as above. Finally, multiplying by x, $x^{q'} - x$ divides $x^q - x$, and therefore all roots of $x^{q'} - x$ are in K, so K' is contained in K. □

One final question: what are the irreducible factors of $x^q - x$ in $F[x]$, where $q = p^e$? The answer turns out to be **all irreducible polynomials whose degree divides** e.

Example 231

Take $q = 2^4 = 16$. The irreducible factors of $x^{16} - x$ in $\mathbb{F}_2[x]$ are

$$x^{16} = x(x+1)(x^2 + x + 1)\cdots$$

where the rest is a product of three irreducible quartics, since we must have degree 2^e where e is a divisor of 4. So this is a way to count the number of irreducible polynomials of a certain degree! They turn out to be

$$x^4 + x + 1, \quad x^4 + x^3 + 1, \quad x^4 + x^3 + x^2 + x + 1.$$

We'll go over this next time!

28 April 24, 2019

Let's finish discussing finite fields today. As a review: if we have a finite field K with $|K| = q$, we must have $q = p^e$, where $K \supset F = \mathbb{F}_p$, the field of p elements. Then $[K : F] = e$, and we have a few important properties:

- $K^\times$ is a cyclic group of order $q - 1$.
- The elements of K are the roots of $x^q - x$ (which has no double roots). Specifically, $x^q - x$ splits in $K[x]$.
- For all $q = p^e$, there exists a unique K with $|K| = q$ (up to isomorphism).
- If $q' = p^{e'}$, and we have a finite field $K' = \mathbb{F}_{q'}$, then $K' \subset K$ exactly when e' divides e.

Corollary 232

If we factor $x^q - x$ in $F[x]$, then the irreducible factors are all irreducible polynomials $f(x)$, such that the degree of f divides e.

Proof. Say that f is an irreducible polynomial in $F[x]$ with degree d, and let α be a root. Then $[F(\alpha) : F] = d$, and thus f dividing $x^q - x$ means f has a root in K: we can then set up a chain

$$F \subset F(\alpha) \subset K$$

and now since $[K : F] = e, [F(\alpha) : F] = d$, we must have d divide e. On the other hand, if $d|e$, then $x^{q'} - x$ divides $x^q - x$, so we indeed have a subfield K' of K. □

Let's do some applications of this:

Example 233

Factor $x^{27} - x$ in $\mathbb{F}_3[x]$.

Since $27 = 3^3$, the irreducible factors must be irreducible polynomials of degree 1 or 3 (those are the divisors of the exponent 3). We know the linear polynomials: $x^{27} - x$ is divisible by $x, x+1, x-1$, and this leaves degree 24: luckily this is divisible by 3.

So the rest is a product of 8 irreducible polynomials of degree 3. A polynomial of the form $x^3 + ax^2 + bx + c$ is irreducible when it has no linear factors, so $0, 1, -1$ must not evaluate to 0. But that's annoying to calculate – we won't do it here.

Fact 234

$x^q - x$ has no repeated roots in $\mathbb{F}_p[x]$, so each factor can only occur once.

Example 235

How many irreducible polynomials are there of degree 5 and 10 over $\mathbb{F}_2$?

Note that $2^5 = 32$, so we want to factor $x^{32} - x$. All factors are either linear or degree 5, and the linear factors are x and $x+1$. This leaves an exponent of 30: therefore, there are 6 irreducible polynomials of degree 5. (Remember that any irreducible polynomial does need to be included in the product, because the roots are roots of $x^{32} - x$.)

Meanwhile, the number of irreducible polynomials of degree 10 over $\mathbb{F}_2$ can be found similarly: since $2^{10} = 1024$, and we can have only polynomials of degree $1, 2, 5, 10$. Since $x, x+1, x^2+x+1$ are the low-degree irreducible polynomials, and we have 6 of degree 5, this leaves an exponent of

$$1024 - 1 - 1 - 2 - 30 = 990,$$

and thus there are 99 polynomials of degree 10 over $\mathbb{F}_2$.

What are the six irreducible polynomials of degree 5 here? They're of the form

$$x^5 + ax^4 + bx^3 + cx^2 + dx + 1$$

(since the constant term can't be 0). Since 1 can't be a root, we must have an odd number of a, b, c, d be equal to 1. Then we need to make sure there aren't quadratics that work: turns out $x^5 + x^4 + 1, x^5 + x + 1$ are reducible, and the other six work.

For the rest of class, we'll assume our fields F have **characteristic zero**: that means adding 1 to itself never gives us 0 (so we don't have something like $\mathbb{F}_q$). This means that F contains $\mathbb{Z}$ and therefore also $\mathbb{Q}$.

Proposition 236

Let $f(x)$ be an irreducible polynomial in $F[x]$. Then it has no multiple root in any field extension.

Proof. If f has a multiple root, then it is a root of both f and f'. The degree of f' is one less than the degree of f, and f is irreducible in $F[x]$, so f, f' have no common divisors in $F[x]$. But $(x - \alpha)$ would need to be a common divisor in $K[x]$, because f has to have a factor of $(x-\alpha)^2$ and thus f' has a factor of $(x-\alpha)$. We use a lemma here to finish:

Lemma 237
If $F \subset K$ and $f, g \in F[x]$, then the greatest common divisor in $F[x]$ and $K[x]$ is the same.

Proof. Let $d = \gcd(f, g)$ in $F[x]$, and $d' = \gcd(f, g)$ in $K[x]$. $F[x]$ is a subring of $K[x]$, so d will divide both f and g in $K[x]$: this means d divides d'. On the other hand, if $d = pf + qg$ for $p, q \in F[x]$, this statement is also true in the ring $K[x]$. Now d' divides f and g, so d' divides d. This means the two gcds are equal, as desired. □

So if $(x - \alpha)$ is not a common root of f and f', then $(x - \alpha)$ can't be a common divisor in $K[x]$ either, and thus there is no multiple root. □

Let's move on: the next idea is sort of tricky.

Theorem 238 (Primitive Element Theorem)
If $F \subset K$ is a field, and $[K : F] < \infty$ is finite, say that $\gamma \in K$ to be **primitive** if $K = F[\alpha]$. If F has characteristic 0, then there exists a primitive element for the field extension.

Proof. If K is a finite extension, we can always adjoin a finite basis to F: let's say $K = F[\alpha_1, \cdots, \alpha_k]$. We induct on k: if $k = 1$, we're done. For the inductive step, since $K \supset F[\alpha_1, \cdots, \alpha_{k-1}] = K'$, we can assume K' has a primitive element β. So now we just need to show that $K = F[\alpha, \beta]$ (generated by two elements) has a primitive element.

Let $\gamma = \beta + c\alpha$ ($c \in F$ is a number to be determined). Our goal is to show that $K = F[\alpha, \beta]$ is $F[\gamma]$ for most choices of c (in fact, all but a finite number). Take $f(x)$ to be an irreducible polynomial in $F[x]$ with root α (of degree m), and let $g(x)$ be an irreducible polynomial in $F[x]$ with root β (of degree n). We can construct a field extension of K in which f and g split completely: then if the roots of f in L are $\alpha_1, \cdots, \alpha_m$ (where $\alpha = \alpha_1$), and the roots of g in L are $\beta_1, \cdots, \beta_m$ (where $\beta = \beta_1$), all αs here are distinct by Proposition 236, and so are all βs.

So if $\gamma = \beta + c\alpha$, the field $K' = F[\gamma]$ is some subfield of K. Our goal is to show that $K' = K$. Here's the trick: α is a root of f, and we can find another polynomial in $K'[x]$ with root α as follows. Since $\beta = \gamma - c\alpha$, the polynomial

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

has coefficients in $K'[x]$ (since $c, \gamma \in K'$), and it has α as a root (since $h(\alpha) = g(\beta) = 0$). So now α is a common root of f and h in $K'[x]$: we claim it's the only one! If α_i (another root of f) were also a root of $h(x)$, then we must have

$$h(\alpha_i) = g(\gamma - c\alpha_i) = 0,$$

which is true if and only if $\gamma - c\alpha_i = \beta_j$ for some j, which means

$$\beta_1 + c\alpha_1 - c\alpha_i = \beta_j.$$

This means

$$c = -\frac{\beta_j - \beta_1}{\alpha_i - \alpha_1},$$

and remember that we can pick c now: just pick it to not be any of the values for any i, j. (A field of characteristic zero has infinitely many choices for c.) So now α_1 is the only common root of f and h, which means that the gcd of f and h is $(x - \alpha_1)$ in $K'[x]$. Therefore $\alpha_1 \in K'$, which means that $\gamma, \beta \in K'$, and thus $K = F[\alpha, \beta] \subset K'$.

Picking a suitable c, K and K' are the same field, as desired, completing the inductive step. □

Fact 239

This theorem is false in non-characteristic zero fields! For example, consider $F = \mathbb{F}_p[t]$, and let $f(x) = x^p - t$. This is irreducible (with a similar argument as Eisenstein's criterion), but $f'(x) = px^{p-1} = 0$. So all roots are multiple roots: in fact, this has just one root in any field extension. So certain fields of degree p^2 do not have a primitive element (if we adjoin two such pth roots).

29 April 26, 2019

We'll get started on Galois theory today. Note that today we'll deal with **characteristic 0**, so that all irreducible polynomials have distinct roots.

Let f be an irreducible polynomial in $F[x]$, and let's say α is a root of f in some extension K of F. Then $F[\alpha]$ is easy to compute in, because it is just polynomials in α with basis $(1, \cdots, \alpha^{n-1})$, where n is the degree of f, and we have the relation $f(\alpha) = 0$. But it's not so clear how to compute with more than one root at a time at the moment.

Definition 240

A **splitting field** K of F is an extension of F with two properties:

- A polynomial $f(x)$ splits completely over K (that is, all of its roots $\alpha_1, \cdots, \alpha_n$ are in K).
- We can write $K = F(\alpha_1, \cdots, \alpha_n)$, so all elements of K can be written as a polynomial in the α_is with coefficients in F. (This might not be unique, though.)

Now computing in a splitting field requires knowing how the roots are related, and that depends on our polyno,ial f.

Example 241

Let's say $f(x) = x^2 + bx + c$ is some quadratic polynomial, where $b, c \in F$.

Then the roots α_1, α_2 are in some splitting field K, and we want to say something more about how they are related.

Definition 242

An ***F*-automorphism** of a field extension K is an isomorphism $\sigma : K \to K$ back to itself, such that σ is the identity on F.

Lemma 243

There exists an F-automorphism such that $\sigma(\alpha_1) = \alpha_2, \sigma(\alpha_2) = \alpha_1$.

Proof. Complete the square: we have

$$g(y) = f(y - \tfrac{1}{2}b) = (y^2 - by + \tfrac{b^2}{4}) + b(y - \tfrac{1}{2}) + c = h^2 + \frac{b^2 - 4c}{4}.$$

This has roots $\gamma, -\gamma$, where γ is $\frac{1}{2}\sqrt{D}$ (D is the "discriminant"). Then adjoining γ is equivalent to adjoining α, because the two differ by an element of F. Then the F-automorphism that sends $\gamma \to -\gamma$ (and acts as the identity on F) sends α_1 to α_2 and vice versa, as desired. □

Similarly, if $f(x)$ is an irreducible cubic with roots $\alpha_1, \alpha_2, \alpha_3$ in a splitting field K, there exists an F-automorphism such that $\sigma(\alpha_1) = \alpha_2, \sigma(\alpha_2) = \alpha_3, \sigma(\alpha_3) = \alpha_1$. It's not as clear how to construct thism though – it is true that a root has to go to another root under any such σ because the relation $f(\alpha) = 0$ still needs to hold true, but we don't necessarily know that we will get this specific permutation. We might return to this later.

Recall Vieta's formulas: if we have a product of the form

$$(x - u_1)(x - u_2) \cdots (x - u_n),$$

we can expand it out as

$$x^n - s_1 x^{n-1} + s_2 x^{n-2} - \cdots \pm s_n,$$

where

$$s_1 = u_1 + \cdots + u_n = \sum_i u_i,$$

$$s_2 = u_1 u_2 + \cdots = \sum_{i<j} u_i u_j,$$

and so on, where we basically pick k of the u_is if we have an x^{n-k}. These are called the **elementary symmetric functions**, because **they don't change when we permute the u_is**. In other words, we can think of $G = S_n$, the symmetric group, as operating on $F[u_1, \cdots, u_n]$ (though F can really be any ring), with the rule

$$\sigma u_i = u_{\sigma(i)}$$

(S_n is operating on the indices).

Definition 244

A polynomial $g(u_1, \cdots, u_n)$ is **symmetric** if

$$\sigma(g) = g \ \forall \sigma \in G.$$

In particular, this means

$$\sigma(g(u_1, \cdots, u_n)) = g(u_{\sigma_1}, \cdots, u_{\sigma_n}).$$

Theorem 245 (Symmetric Functions Theorem)

Every symmetric polynomial $g(u_1, \cdots, u_n)$ can be written in just one way as a polynomial in the elementary symmetric functions. Specifically, there exists a polynomial $G(z_1, \cdots, z_n)$ such that

$$g(u) = G(s_1, \cdots, s_n).$$

Example 246

Consider $g(u) = u_1^2 + \cdots + u_n^2$ (the sum of n squares)

For $n = 2$, we have $s_1 = u_1 + u_2, s_2 = u_1 u_2$, and now $g(u) = s_1^2 - 2s_2$. For $n = 3$, we have $s_1 = u_1 + u_2 + u_3, s_2 = u_1 u_2 + u_1 u_3 + u_2 u_3, s_3 = u_1 u_2 u_3$. The third symmetric polynomial here is useless because the degree is too large; let's try to do the rest systematically.

If we set $u_3 = 0$, we still have a symmetric polynomial in u_1 and u_2, and the expression for G reduces to the

elementary symmetric polynomials in two variables. Matching degrees tells us that we must have $g(u) = as_1^2 + bs_2$, and it's $\boxed{s_1^2 - 2s_2}$ in the degree 2 case, so **this must be true for all** $n \geq 2$.

Example 247

What if we take $g = u_1^3 + \cdots + u_n^3$?

For $n = 1$, we have $g = s_1^3$. For $n = 2$, monomials of degree 3 are s_1^3 and $s_1 s_2$. Therefore,

$$u_1^3 + u_2^3 = as_1^3 + bs_1 s_2$$

for some constants a, b. Set $u_2 = 0$, and $bs_1 s_2$ goes away, so $a = 1$. To find b, we can set $u_1 = u_2 = 1$. Now the left side is 2, and the right side is $8 + 2b$, so $b = -3$.

Finally, let's do $n = 3$ and write $u_1^3 + u_2^3 + u_3^3$ in terms of symmetric functions. If we set $u_3 = 0$, we just remove the s_3-only terms from the picture, and since s_3 has degree 3, that means we have

$$u_1^3 + u_2^3 + u_3^3 = s_1^3 - 3s_1 s_2 + cs_3.$$

for some c. Now plugging in $u_1 = u_2 = u_3 = 1$, this becomes $3 = 27 - 27 + c$, and $c = 3$.

It should now be pretty clear how this generalizes: add on one variable at a time! We'll write out the proof formally now.

Proof. Given that $g(u_1, \cdots, u_n)$ is symmetric, induct on n and on the degree of g. Plugging in $u_n = 0$, define

$$g^0(u_1, \cdots, u_{n-1}) = g(u_1, \cdots, u_{n-1}, 0),$$

and also define the elementary symmetric functions

$$s_i^0(u_1, \cdots, u_{n-1}) = s_i(u_1, \cdots, u_{n-1}, 0).$$

Since g is symmetric, g^0 is also symmetric in the first $n-1$ elements, and the s_i^0s are the elementary symmetric polynomials in the $n-1$ elements (plus the zero polynomial). By inductive hypothesis,

$$g^0(u_1, \cdots, u_{n-1}) = G(s_1^0, \cdots, s_{n-1}^0)$$

uniquely for some polynomial G. So now we can write

$$h(u_1, \cdots, u_n) = g(u_1, \cdots, u_n) - G(s_1, \cdots, s_n)$$

(sure, s_n doesn't appear, but we can still include it). Note that $h(u_1, \cdots, u_n) = 0$ if $u_n = 0$ by definition of G, so u_n divides $h(u_1, \cdots, u_n)$. But h is symmetric, because it is the difference of two symmetric polynomials. So if every monomial in h contains u_n, all u_i must divide h as well, and thus the product $u_1 u_2 \cdots u_n$ divides h: we can write $h(u) = s_n q(u)$ for some other polynomial q (which is also symmetric). And now this can be written as some polynomial in $s_1, \cdots, s_n$ by induction on the degree. □

There's one more very important symmetric polynomial:

Definition 248

The **discriminant** of a polynomial $P(x) = \prod_i (x - u_i)$ is

$$D = (u_1 - u_2)^2 (u_1 - u_3)^2 \cdots = \prod_{i<j} (u_i - u_j)^2.$$

D is a symmetric polynomial, because it takes each difference $u_i - u_j$ to some other one (up to a $\pm$ sign) and the squares fix potential issues with signs.

Remark 249. *Note that this is consistent with the definition of the discriminant for monic quadratic polynomials:*

$$D = (u_1 - u_2)^2 = u_1^2 + u_2^2 - 2u_1u_2 = s_1^2 - 2s_2 - 2s_2 = s_1^2 - 4s_2,$$

where $s_1 = -b, s_2 = c$, so this is the familiar $b^2 - 4c$.

Example 250

What can we say about the discriminant for a cubic ($n = 3$)?

In this case, we have

$$D = (u_1 - u_2)^2 (u_1 - u_3)^2 (u_2 - u_3)^2,$$

which is a symmetric polynomial of degree 6. This has to be some linear combination of $s_1^6, s_1^4 s_2, s_1^3 s_3, s_1^2 s_2^2, s_1 s_2 s_3, s_2^3, s_3^2$. It's not easy to determine the coefficients, really, but let's try the systematic method: if $u_3 = 0$, then

$$D^0 = (u_1 - u_2)^2 u_1^2 u_2^2 = (s_1^0)^2 - 4(s_2^0)^2 (s_2^0)^2,$$

so $D^0 = (s_1^0)^2 (s_2^0)^2 - 4(s_2^0)^3$. The systematic method then tells us that

$$D = s_1^2 s_2^2 = 4s_2^3 + s_3(*)$$

for some polynomial $*$: this just leaves some linear combination of $s_1^3 s_3, s_1 s_2 s_3, s_3^2$. there isn't really a good way to do it other than putting in values for the us and getting equations in the unknowns. This is not recommended: they turn out to be $-4, 18, -27$.

30 April 29, 2019

Let's review the ideas of symmetric functions: if we have variables $u_1, \cdots, u_n$, then the symmetric group S_n operates on polynomials $F[u_1, \cdots, u_n]$ as an automorphism: for any $\sigma \in S_n$, we send $u_i \to u_{\sigma(i)}$ under the group action. We define a polynomial $p(u_1, \cdots, u_n)$ to be **symmetric** if $p = \sigma(p)$ for all $\sigma \in S_n$. Then we have Theorem 245 from last class: every symmetric polynomial $p(u_1, \cdots, u_n)$ can be uniquely writen as a polynomial in the elementary symmetric polynomials

$$s_1 = u_1 + \cdots + u_n,$$

$$s_2 = \sum_{i<j} u_i u_j,$$

up to

$$s_n = u_1 \cdots u_n.$$

Here's a useful fact that comes out of this:

Corollary 251

If $f(x) \in F[x]$ is a polynomial that splits in an extension field K with roots $\alpha_1, \cdots, \alpha_n$, let $p(u_1, \cdots, u_n)$ be a symmetric polynomial in $F[u_1, \cdots, u_n]$. Then $p(\alpha_1, \cdots, \alpha_n)$ is in F.

For example, the discriminant of a cubic is $(u_1 - u_2)^2(u_1 - u_3)^2(u_2 - u_3)^2$, which is symmetric in u_1, u_2, u_3, so $D(\alpha_1, \alpha_2, \alpha_3)$ is always in F.

Slightly more complicated proof than necessary. $p(u_1, \cdots, u_n)$ is a polynomial in $s_1(u), \cdots, s_n(u)$ by the symmetric function theorem. Thus, we can write $p(u_1, \cdots, u_n) = \Phi(s_1, \cdots, s_n)$ for some polynomial $\Phi \in F[x]$, and now we can substitute in $u_i = \alpha_i$.

For reference later on, let $H(x) = (x - u_1) \cdots (x - u_n) = x^n - s_1 x^{n-1} + \cdots \pm s_n$, so we get the polynomial $f(x) = (x - \alpha_1) \cdots (x - \alpha_n)$ by substituting in $u_i = \alpha_i$. So now $f(x) = x^n - s_1(\alpha)x^{n-1} + s_2(\alpha)x^{n-2} + \cdots \pm s_n(\alpha)$, and all $s_1(\alpha) \in F$ because f is an element of $F[x]$. So now $p(\alpha_1, \cdots, \alpha_n) = \Phi(s_1(\alpha), \cdots, s_n(\alpha))$, and all $s_i(\alpha)$ are elements of F, so a polynomial of them is in F. □

Basically, the symmetric functions are in the field because $f(x) \in F[x]$, and then we can use the symmetric function theorem to write p in terms of the symmetric polynomials.

Here's another game we can play with symmetric functions: let $p(u_1, \cdots, u_n)$ be a polynomial (not necessarily symmetric). We want to know about the **orbit** of p_1. Its order divides $n!$, because we are permuting the n variables among each other. Let's say this orbit is $(p_1, \cdots, p_k)$: we can think of S_n as operating on $\{p_1, \cdots, p_k\}$.

Lemma 252

Let $\Phi(w_1, \cdots, w_k)$ be a symmetric polynomial in w. Then $\Phi(p_1(u), \cdots, p_k(u))$ is a symmetric polynomial in $u_1, \cdots, u_n$.

This is because S_n operates on the polynomials, so when it permutes the u_is, it also permutes the p_is, and thus it fixes Φ.

Example 253

Consider $n = 3$ with the polynomial $p_1 = u_1u_2 + u_2u_3$.

Then the orbit is p_1 plus the polynomials

$$p_2 = u_1u_2 + u_1u_3, p_3 = u_1u_3 + u_2u_3.$$

As an example of a symmetric polynomial of these three elements (p_1, p_2, p_3), we have

$$s_3 = p_1p_2p_3 = 2u_1^2u_2^2u_3^2 + \sum_{\text{symmetric}} u_1^3u_2^2u_3.$$

Now, let's put these two ideas together: recall that a **splitting field** K of a polynomial $f(x) \in F[x]$ is where f splits completely into linear factors corresponding to its roots $\alpha_1, \cdots, \alpha_n$, and K is generated by the α_is. (This can be thought of as K being the smallest field containing the roots.)

Theorem 254 (Splitting theorem)
Let K be the splitting field of $f(x) \in F[x]$, and let g be an irreducible polynomial in $F[x]$. Then if g has a root $\beta \in K$, then g splits completely in K.

(Note that g and β are completely arbitrary here.)

Proof. Let the roots of f be $\alpha_1, \cdots, \alpha_n$. Every element of K is a polynomial in the αs (since K is generated by them), so let $p(u_1, \cdots, u_n)$ be an element of $F[u_1, \cdots, u_n]$ such that

$$\beta = p(\alpha_1, \cdots, \alpha_n).$$

S_n operates on polynomials in $F[u_1, \cdots, u_n]$, so let $\{p_1, \cdots, p_k\}$ be the orbit of $p = p_1$. Let $\beta_j = p_j(\alpha_1, \cdots, \alpha_n)$, where $\beta = \beta_1$ by definition.

Our goal is to show that the polynomial $h(x) = (x - \beta_1) \cdots (x - \beta_k)$ has coefficients in F. If this is true, then β is a root of $h(x)$ and also of $g(x)$, so they aren't relatively prime. But since g is irreducible, this would imply that g divides h: then because h splits in K, g must also split in K. Let's do an example for illustration:

Example 255
Let $\zeta = e^{2\pi i/9}$ be a ninth root of unity. Its irreducible polynomial over $F = \mathbb{Q}$ is $x^6 + x^3 + 1$, and the roots in the splitting field K are $\zeta, \zeta^2, \zeta^4, \zeta^5, \zeta^7, \zeta^8$. Let's call them $u_1, u_2, \cdots, u_6$.

Then what we're saying is that the minimal polynomial of $\beta = \zeta + \zeta^8 = 2\cos\frac{2\pi}{9}$ must split: the orbit of $u_1 + u_6$ is all polynomials of the form $u_i + u_j$, where $i \neq j$. (There are $\binom{6}{2} = 15$ of them.) Then

$$h(x) = \prod_{i<j} (x - [u_i + u_j])$$

has degree 15, and β is a root of $g(x) = x^3 - 3x + 1$ (this can be checked). And $g(x)$ will factor as

$$g(x) = (x - (\zeta + \zeta^8))(x - (\zeta^2 + \zeta^7))(x - (\zeta^4 + \zeta^5)),$$

and this is indeed a factor of $h(x)$.

So returning to our proof, how do we show that $h(x)$ has coefficients in F? For variables $w_1, \cdots, w_k$, let

$$\tilde{H}(x) = (x - w_1) \cdots (x - w_k)$$

and we substitute $w_j = p_j(u)$ in to get

$$H(x) = (x - p_1(u))(x - p_2(u)) \cdots (x - p_k(u)).$$

Now substituting $u_i = \alpha_i$, we get back to the familiar

$$h(x) = (x - \beta_1) \cdots (x - \beta_k).$$

But $\tilde{H}$ was symmetric in the w_is, and its coefficients x^j are $s'_j(w_1, \cdots, w_k)$, the elementary symmetric polynomials in $w_1, \cdots, w_k$. So the coefficients of x^j in H are $S'_j(p_1(u), \cdots, p_k(u))$, which are symmetric in $u_1, \cdots, u_n$ by Lemma 252. Thus the coefficients of x^j in $h(x)$ are in F, because plugging in the roots α_i (to get from H to h) gives us coefficients symmetric in the α_is, which are definitely in F. □

31 May 1, 2019

Today, we're going to start Galois theory – it'll take a bit of getting used to.

Definition 256

Let K and K' be an extension field of F. Then an ***F*-isomorphism** ϕ from K to K' is an isomorphism which is the identity on F. Similarly, an ***F*-automorphism** of K is just an F-isomorphism from K to itself.

The important concept here is that set of automorphisms:

Definition 257

The **Galois group** $G(K/F)$ is the set of all F-automorphisms of K.

It turns out this group is the right way for us to understand splitting fields:

Proposition 258

If f is an irreducible polynomial in $F[x]$, and it has a root α in K, then any F-isomorphism $\phi : K \to K'$ takes α to another root of f.

(This is an extremely important result to keep in mind!)

Proof. By convention, let's write $f(x) = x^n - a_1 x^{n-1} + \cdots \pm a_n$. This will make things simpler, because now if f splits and has roots $\alpha_1, \cdots, \alpha_n$, then a_i is just the symmetric polynomial $s_i(\alpha_1, \cdots, \alpha_n)$ by Vieta's formulas.

So now because $f(\alpha) = 0$, we also have $\phi(f(\alpha)) = 0$, which means that

$$\phi(\alpha)^n - \phi(a_1)\phi(\alpha)^{n-1} \cdots \pm \phi(a_n) = 0.$$

But the coefficients are fixed: $\phi(a_i) = a_i$, since F is fixed, and therefore $\phi(\alpha)$ is also a root of f, as desired. □

In particular, this also applies to F-automorphisms.

Proposition 259

If α is a root of f in K and α' is a root of f in K', then there exists a unique F-isomorphism $\phi : F(\alpha) \to F(\alpha')$ such that α is sent to α'.

This is because $F(\alpha)$ and $F(\alpha')$ are both isomorphic to $F[x]/(f)$, so they are isomorphic to each other. We can't necessarily say that K and K' are isomorphic, though – they might be much larger.

Proposition 260

Say that $f(x)$ has roots $\alpha_1, \cdots, \alpha_n \in K$, and we can write $K = F(\alpha_1, \cdots, \alpha_n)$. Then the Galois group $G(K/F)$ operates **faithfully** on $\{\alpha_1, \cdots, \alpha_n\}$: that is, if σ fixes all α_is, then σ is the identity.

Proof. Every element of K can be written as a polynomial in the α_is, and now take ϕ of both sides – any σ that fixes all α_is (as well as the field F) keeps all elements constant. □

In fact, in general, **knowing where the α_is go tells us where everything else goes.** So that means that G is isomorphic to a **subgroup of S_n** by its action on roots!

Let's try to compute the Galois group for some small-degree polynomials. We'll start by looking at **quadratic irreducible polynomials** $f(x) \in F[x]$. K, the splitting field of f over F, adjoins the two roots α_1 and α_2 to F.

Fact 261

Again, remember that we are working in characteristic zero, so we have no repeated roots.

In particular, our polynomial can be written as $f(x) + x^2 - a_1x + a_2$, where $a_1 = \alpha_1 + \alpha_2, a_2 = \alpha_1\alpha_2$. By the proposition above, there exists a unique automorphism

$$\sigma : F(\alpha_1) \to F(\alpha_2).$$

Also, notice that we get the second root for free: $\alpha_2 = a_1 - \alpha_1$ can be recovered (or vice versa), so both $F(\alpha_1)$ and $F(\alpha_2)$ are just K. This, together with the identity automorphism, means that **the Galois group is cyclic of order 2, and it's generated by a transposition.**

Things get a lot harder as the degree gets larger, but the cubic is pretty easy. Let's say $f(x)$ is an irreducible cubic in $F[x]$: then we can write

$$f(x) = x_1^3 - a_1x^2 + a_2x - a_3.$$

We have roots $\alpha_1, \alpha_2, \alpha_3$, and this means we have some tower

$$F \subset F(\alpha_1) = F_1 \subset F(\alpha_1, \alpha_2) = K,$$

since we get the third root $\alpha_3 = a_1 - \alpha_1 - \alpha_2$ for free. We know the first extension has degree 3: all that's left is to ask about the second extension.

Well, we know that in $F_1[x]$, $f(x)$ factors as

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

for some quadratic $q(x)$. The coefficients then involve α_1: the two roots of the quadratic then have roots α_2, α_3 in K. There are two cases now: if $q(x)$ is irreducible in $F_1[x]$, then the degree of the extension has degree 2, and if q factors, then the degree is 1 (this means the other two roots are already in F_1). **So $[K : F]$ is either 3 or 6.**

- What can we say when $[K : F] = 3$? Then we have $F(\alpha_1) = K = F(\alpha_2)$ – both have degree 3 – and then there exists a unique F-isomorphism σ from $F(\alpha_1)$ to $F(\alpha_2)$. σ is an element of the Galois group: we ask what it does to α_2. It can't go to itself (because α_1 is already going to α_2), and if α_2 goes to α_1, then we'd have to fix α_3. But this can't be true – it would imply that σ is the identity, because there is a unique map that sends α_3 to itself in $K = F(\alpha_3)$! So α_2 goes to α_3 and α_3 goes to α_1.

  This means that σ must be a (123) cyclic permutation of the roots. Then (123) is the only thing that sends α_1 to α_2, (132) is the only thing that sends α_1 to α_3, and the identity sends α_1 to itself. So in this case, G is the cyclic group generated by σ, which is the alternating group A_3.

- In the other case, if $[K : F] = 6$, then $F(\alpha_1) \neq K$, since $[K : F(\alpha_1)] = 2$. In addition, the quadratic polynomial $q(x)$ as defined above is irreducible in $F_1[x]$. We've analyzed quadratics before: the Galois group of K over F_1 is the cyclic group of order 2. Let's say it's generated by τ. Then τ fixes α_1, since it's an element of F_1, and it also fixes F, since $F \subset F_1$. So τ is in the Galois group of K/F as well: in particular, it's the transposition (23).

We can repeat this argument if we extend α_2 first, or α_3, so $G(K/F)$ also contains the transpositions (13) and (12). Specifically, this means that $G(K/F) = S_3$.

Corollary 262

In both cases, we have

$$[K : F] = |G(K/F)|.$$

We're not quite done yet: how do we know which case we're in? Is there an easier way for us to tell whether $q(x)$ is irreducible in $F_1[x]$ without having to work in a new field? It turns out that the discriminant is the secret here: we know that

$$D(u_1, u_2, u_3) = (u_1 - u_2)^2(u_1 - u_3)^2(u_2 - u_3)^2$$

can be written as some combination of the symmetric polynomials s_1, s_2, s_3. In particular,

$$D(\alpha_1, \alpha_2, \alpha_3) = -4a_1^3a_3 + a_1^2a_2^2 + 18a_1a_2a_3 - 4a_2^3 - 27a_3^2 \in F$$

because all of our coefficients $a_1, a_2, a_3 \in F$. We want to look at the **square root**

$$\delta = (u_1 - u_2)(u_1 - u_3)(u_2 - u_3).$$

Specifically, $\delta(\alpha_1, \alpha_2, \alpha_3)^2 = D(\alpha_1, \alpha_2, \alpha_3)$. If $D(\alpha)$ is not a square in F, then $\delta \notin F$: adjoining $\delta = (\alpha_1 - \alpha_2)(\alpha_1 - \alpha_3)(\alpha_2 - \alpha_3)$ yields

$$[F(\delta) : F] = 2.$$

But δ is contained in the splitting field, so

$$F \subset F(\delta) \subset K \implies [K : F] \text{ is even} \implies [K : F]6.$$

On the other hand, if $D(\alpha)$ is a square, then $\delta \in F$: how does the symmetric group S_3 operate on $\delta(u)$? We can try it out: even permutations fix δ, and odd ones flip the sign of δ. Now $\delta \neq 0$, because the roots are distinct for irreducible f. This means that the Galois group does not contain an odd permutation: $\delta \in F$, and F is supposed to be fixed. Therefore $[K : F] = 3$ as desired.

Fact 263

So the **square root of the determinant** tells us the Galois group for a cubic! Most of the time, it'll be S_3, because it's pretty unlikely for an ugly expression in a_1, a_2, a_3 to be zero. One case where we are happy is when there is no quadratic term: $f(x) = x^3 + a_2x - a_3 \implies D(\alpha) = -4a_2^3 - 27a_3^2$.

So recall the following example from class earlier: if ζ is a 9th root of unity, let $\alpha_1 = \zeta + \zeta^8, \alpha_2 = \zeta^2 + \zeta^7, \alpha_3 = \zeta^4 + \zeta^5$. These are roots of $x^3 - 3x + 1$, and the discriminant is $-4(-3)^3 - 27(1)^2 = 81$, which is a square. That means adjoining α_1 gives the other two roots for free as well.

32 May 3, 2019

Today we're going to discuss the main theorem of Galois theory.

Theorem 264

Suppose we have a finite group of automorphisms for a field K. Let F be the **fixed field** of G: that is, it's the elements

$$F = \{\alpha \in K | \sigma\alpha = \alpha \ \forall \sigma \in G\}.$$

Then $|G| = [K : F]$.

Proof. Pick an arbitrary $\alpha \in K$: it will have some orbit $\{\alpha_1, \cdots, \alpha_r\}$, where r divides $n = |G|$ by the orbit-stabilizer theorem. Consider the polynomial

$$f(x) = (x - \alpha_1) \cdots (x - \alpha_r) = x^r - s_1(\alpha)x^{r-1} + \cdots \pm s_r(\alpha),$$

where the s_is are symmetric polynomials in the α_js. Since σ permutes the orbit, it fixes each of the symmetric polynomials: this means that $f(x)$ is fixed under all permutations, so $f(x)$ is in the fixed field of G. In other words,

$$s_i(\alpha) \in F \implies f(x) \in F[x].$$

We should check that $f(x)$ is irreducible: if we write $f(x) = g(x)h(x)$, where α_1 is a root of $g(x)$ without loss of generality, there exists some σ_i that sends $\sigma_i(\alpha) = \alpha_i$ (this exists because the α_js form an orbit). So now

$$g(\alpha_1) = 0 \implies g(\alpha_i) = 0$$

under the permutation σ_i, so α_i is a root of $g(x)$. But this means all α_is are a root of g, so f divides g, meaning we can't factor f in the first place.

So the degree of α over F, which is the degree of f, is equal to r, the number of roots of f. Choose α_1 so that r is maximal: our goal is to show that $K = F(\alpha_1)$. In other words, if $\beta \in K$ is an arbitrary element, we want to show it is in $F(\alpha_1)$. Note that $F(\alpha_1, \beta) \supset F(\alpha_1) \supset F$, where the second extension has degree r. We know that $[F(\alpha_1, \beta) : F(\alpha_1)] \leq [F(\beta) : F]$, since polynomials are only "more irreducible" in F than in $F(\alpha_1)$. So the first extension has degree has degree at most r, meaning that $[F(\alpha_1, \beta) : F]$ is a finite extension.

So now by the primitive element theorem, there exists γ such that $F(\alpha_1, \beta) = F(\gamma)$. Now $[F(\gamma) : F] \leq \deg\alpha_1 = r$, because we chose α_1 maximally! But $[F(\alpha_1, \beta) : F] \geq r$ from the chain of inclusions, so therefore we have equality on both sides, and therefore

$$[F(\alpha_1, \beta) : F] = r \implies [F(\alpha_1, \beta) : F(\alpha_1)] = 1 \implies \beta \in F(\alpha_1),$$

as desired.

Now to finish, we want to show that this maximal r is equal to n, which will conclude the proof. By the orbit-stabilizer theorem,

$$|G| = |\text{Stab}(\alpha_1)||\text{Orbit}(\alpha_1)|,$$

and now if $\sigma \in \text{Stab}(\alpha_1)$, we know that $\sigma(\alpha_1) = \alpha_1$. But that means σ is the identity on all of $F(\alpha_1) = K$! Thus the stabilizer only has one element, and we're done: we do have $|G| = r = [K : F]$. □

Remember that an F-automorphism of a field extension K is an automorphism $\sigma : K \to K$ that is the identity on F. We define the **Galois group** of an extension K/F to be the set of F-automorphisms of K.

Theorem 265

Let K be a finite extension of F, and let G be the Galois group $G(K/F)$. The following are equivalent:

- K is a splitting field over F.
- $|G| = [K : F]$ – in other words, K is a **Galois extension** of F.
- F is the fixed field of G.

Note that this gives us a way to find whether an element of K is actually in F: it depends on whether everything in the Galois group fixes it!

Proof. Pick a primitive element α for K, so that $K = F(\alpha)$, and let $f(x)$ be the irreducible polynomial for α in $F[x]$. Let $\alpha_1, \cdots, \alpha_r$ be the roots of f in K: we know that $r \leq n = \deg f = [K : F]$.

We know there exists a unique F-isomorphism σ_i from $K = F(\alpha_1) \to F(\alpha_i)$, and because $F(\alpha_i)$ also has degree n over F, we actually have $F(\alpha_i) = K$. So σ_i is an F-automorphism, and therefore $\sigma_i \in G$.

So G is just the set $\{\sigma_1, \cdots, \sigma_r\}$, since each element of the Galois group must send a root to another root, and therefore $|G| = r$ divides $n = [K : F]$. If $|G| = [K : F]$, then $r = n$, which means f has n roots and degree n, meaning f splits completely (which implies that K is a splitting field). This proves the equivalence of the first two statements.

To show that the third is equivalent, F is contained in the fixed field (by definition of an F-automorphism), so we have

$$F \subset K^G \subset K,$$

where $[K : F] = n$. But $[K : K^G] = |G|$ by Theorem 264, meaning that $|G| = [K : F]$ if and only if $[K^G : F] = 1$, and this gives us the result we want. □

Theorem 266 (Main Theorem)

Let K be a Galois extension of F, and let $G = G(K/F)$. Then subgroups of G correspond bijectively to intermediate fields $F \subset L \subset K$.

This is possibly a bit surprising: G is a finite group, so this means there are only finitely many intermediate fields between F and K. In other words, picking any element of K and considering the field it generates only finitely many choices!

Example 267

Take $F = \mathbb{Q}$, and $K = F(\zeta)$, where $\zeta = e^{2\pi i/7}$.

Now ζ is the root of the irreducible poylnomial $x^6 + x^5 + x^4 + x^3 + x^2 + x + 1$, and the roots are $\zeta, \zeta^2, \cdots, \zeta^6$. Since K is the splitting field of $f(x)$, we indeed have $G = G(K/F)$, and the order of the group $|G| = [K : F] = 6$.

Note that we have an F-isomorphism σ_i sending $F(\zeta) \to F(\zeta^i)$ by sending $\zeta \to \zeta^i$: by the same reasoning as before, each ζ^i also has degree 6 over F, so this is the unique F-automorphism $K \to K$ sending $\zeta \to \zeta^i$. In particular, G is the cyclic group generated by σ_3, which is the map sending ζ to ζ^3.

Now we have the F-isomorphism $\rho = \sigma^2$ which sends ζ to ζ^2. The subgroup generated by ρ has order 3: what's the corresponding intermediate field? Consider the polynomial with roots that are the orbit of ρ under ρ: this is

$$f(x) = (x - \zeta)(x - \zeta^2)(x - \zeta^4) = x^3 - (\zeta + \zeta^2 + \zeta^4)x^2 + (\zeta^3 + \zeta^5 + \zeta^6)x + 1,$$

and let the x^2-coefficient be β_1 and the x-coefficient be β_2. We know that β_1, β_2 are in the fixed field of $K^{\langle \rho \rangle}$, and now

$$(x - \beta_1)(x - \beta_2) = x^2 + x + 2.$$

So we have $\beta_1 = \frac{-1+\sqrt{-7}}{2}$ and $\beta_2 = \frac{-1-\sqrt{-7}}{2}$, and the intermediate field is just $\boxed{\mathbb{Q}(\beta_1, \beta_2)}$ (we really only need to adjoin one of them).

33 May 6, 2019

Recall that last week, we started talking about some ideas of the main theorem. Here's some of the ideas that we proved last time: if we have a finite group G of automorphisms of a field K, and we let K^G be the fixed field of G, then $[K : K^G] = |G|$. We defined a **Galois extension** K of F to be a field extension that satisfied any of the following equivalent conditions:

- K is a splitting field over F.
- $|G(K/F)| = [K : F]$.
- $F = K^G$, the fixed field of K under G.

This leads us to Theorem 266: if K/F is a Galois extension, then we have a bijective correspondence between intermediate fields $F \subset L \subset K$ and subgroups of $G = G(K/F)$. Specifically, any intermediate field L corresponds to the Galois group $G(K/L) = H$, so F corresponds to the whole group G, and K corresponds to the identity. In particular, making L larger makes H become smaller.

Example 268

Let's illustrate this with another example: take $F = \mathbb{Q}$ and $K = F(\alpha, \beta)$, where $\alpha = \sqrt{2}$, $\beta = \sqrt{3}$.

We have a chain $F \subset F(\alpha) \subset F(\alpha, \beta) = K$. The first chain has degree 2 because $x^2 - 2$ is irreducible, and β is a root of $x^2 - 3$, which remains irreducible in $F(\alpha)$. So $[K : F] = [K : F(\alpha)][F(\alpha) : F] = 2 \cdot 2 = 4$, and now because K is the splitting field of $(x^2 - 2)(x^2 - 3)$, $|G(K/F)| = 4$.

To find the Galois group, first note that any automorphism that fixes both $\sqrt{2}$ and $\sqrt{3}$ fixes everything. If $\sigma \in G$, then $\sigma\alpha$ must send α to one of the other roots of $(x^2 - 2)$, $\pm\alpha$. Similarly, $\sigma\beta$ must be one of $\pm\beta$. There's only four choices in total, so we have the Klein four group generated by σ (sending $\alpha \to -\alpha, \beta \to \beta$) and τ (sending $\alpha \to \alpha, \beta \to -\beta$).

From this, we can find the intermediate fields between F and K. The only nontrivial subgroups of the whole group G are $\langle \sigma \rangle, \langle \tau \rangle, \langle \sigma\tau \rangle$. By the main theorem, any subgroup H corresponds to $L = K^H$: this means we get $F(\beta), F(\alpha)$, and $F(\alpha\beta)$, respectively.

Example 269

Let's let K be the splitting field of $x^3 - 2$ over $F = \mathbb{Q}$: then the roots are $\alpha = \sqrt[3]{2}$, $\alpha\omega$, $\alpha\overline{\omega}$, where $\omega = e^{2\pi i/3}$ is a cube root of unity.

This time we have the chain

$$F \subset F(\alpha) \subset F(\alpha, \omega\alpha) = F(\alpha, \omega) = K,$$

where the extensions are of degree 3 and 2 respectively. So now $[K : F] = 6$, and G operates faithfully on the roots: this means G must be isomorphic to S_3, and we can describe each element of G by the permutation on the roots.

Our diagram of subgroups is still not that interesting: the subgroups are A_3, $\langle(23)\rangle$, $\langle(13)\rangle$, $\langle(12)\rangle$. The index of the subgroup $[S_3 : A_3] = 2$, and this corresponds to an extension of degree 2. (Specifically, we have $F(\delta)$, where $\delta = (\alpha_1 - \alpha_2)(\alpha_1 - \alpha_3)(\alpha_2 - \alpha_3)$ is the square root of the discriminant.) Meanwhile, $[S_3 : \langle(23)\rangle] = 3$, which corresponds to $F(\alpha_1)$.

Example 270

What if $f(x)$ is an irreducible quartic with roots $\alpha_1, \alpha_2, \alpha_3, \alpha_4$?

We have our chain

$$F \subset F(\alpha_1) \subset F(\alpha_1, \alpha_2) \subset F(\alpha_1, \alpha_2, \alpha_3) \subset K,$$

where the extensions have degree $4, \leq 3, \leq 2, 1$ (because $\frac{f(x)}{x-\alpha_1}$ is a cubic in $F(\alpha_1)$, but it might not be irreducible). So we have that $|G| = [K : F] \leq 4!$, and let's look at the case where $\boxed{G = S_4}$ (so we can permute the roots however we'd like).

Let's think about the subgroups of S_4. We have the alternating group A_4 of order 12, which has index 2. By the Sylow theorem, there exists a Sylow 2-group with order 8 (and therefore index 3), and it turns out to be three (intersecting) dihedral groups D_4. S_4 also contains four copies of S_3 (just fix the first index) of order 6 (and therefore index 4). We also have smaller groups – those generated by a 4-cycle, and so on.

This gives us a much more complicated diagram. We'll discuss this a bit more on Wednesday, but here's a few examples: $[S_4 : S_3] = 4$ corresponds to adjoining a root $F(\alpha_1)$, and $[S_3 : S_2] = 3$ corresponds to adjoining another root $F(\alpha_1, \alpha_2)$. $[S_4 : A_4] = 2$ corresponds, again, to adjoining the square root of the discriminant $\delta = (\alpha_1 - \alpha_2)(\alpha_1 - \alpha_3)(\alpha_1 - \alpha_4)(\alpha_2 - \alpha_3)(\alpha_2 - \alpha_4)(\alpha_3 - \alpha_4)$.

We're now ready to prove Theorem 266, the Main Theorem:

Proof. We'll show that given an intermediate field L, we can map it to $G(K/L) = H$ and show that $L = K^H$. Also, given a subfield H, we'll show we can map it to $L = K^H$ and show that $G(K/L) = H$. This would show that the map is indeed bijective on both sides, as desired.

If we have a field L, note that K being a splitting field over F is the same as K being a splitting field over L (just take the same polynomial, which will split in the same way). So K is a Galois extension over L as well, meaning that $|G(K/L)| = [K : L]$. Now if H is any subgroup of G, we know that $[K : K^H] = |H|$ by Theorem 264, and now we can let $H = G(K/L)$. This is some subgroup of G, and in particular, $L \subset K^H \subset K$ (the first inclusion because every element of H fixes L by definition). And now $[K : K^H] = |H|$, and $[K : L] = |G(K/L)| = |H|$, so that just means we have equality: $L = K^H$ as desired.

On the other hand, if we're given a subgroup H, let $L = K^H$. We know that $[K : L] = |H|$, and since K is a Galois extension of L, $[K : L] = |G(K/L)|$. The two groups H and $G(K/L)$ have the same order, and we know that

$$H \subset G(K/L) \subset G$$

where the first inclusion comes from everything in H being fixed by L. But now both $[G : H]$ and $[G : G(K/L)]$ are equal to $[K : L]$, so $H = G(K/L)$. □

To see why this is called the Main Theorem, we need to spend a bit more time working with it to understand its significance. Let's go back to cubic equations: this time,

Theorem 271

Suppose that $[K : F] = 3$, and say that $K \subset \mathbb{C}$ (to save time). Also, suppose that the cube roots of unity $\omega = e^{2\pi i/3}$ are in our ground field F. Then K is obtained by adjoining a cube root: $K = F(\gamma)$, where $\gamma^3 \in F$.

Proof. The Galois group has order 3 in this case, and it is cyclic: if we let the roots of f be $\alpha_1, \alpha_2, \alpha_3$, we can use the fact that that K is an F-vector space of dimension 3. Letting σ be a generator of our group $G(K/F)$, σ is a linear operator, because it is an automorphism of K. Specifically, if $\alpha \in F$,

$$\sigma(\alpha\beta) = \sigma(\alpha)\sigma(\beta) = \alpha\sigma(\beta),$$

so multiplying by scalars is consistent with our operator. In addition, $\sigma^3 = 1$ is the identity operator, so the eigenvalues of σ are cube roots of 1. They can't all be 1: this is not obvious, but it's related to a problem on the problem set about diagonalizing! So there exists some eigenvalue $\lambda = \omega$ or $\overline{\omega}$, and both occur (but we don't need to check that).

So now take the eigenvector to be our element $\gamma \in K$: we have $\sigma(\gamma) = \lambda\gamma$, and $\sigma(\gamma^3) = \lambda^3\gamma^3 = \gamma^3$. Since σ generates the Galois group, $\gamma^3 \in K^G = F$. Furthermore, $[K : F] = 3$, so we must have $K = F(\gamma)$, and therefore $K = F(\gamma)$, as desired. □

By the way, this theorem works for any prime p (not just 3). We might come back to some of these ideas next time!

34 May 8, 2019

Let's say we have a splitting field K of a field F (with an irreducible polynomial $f \in F[x]$), and the roots of f are $\alpha_1, \cdots, \alpha_n$. There are basically two questions we want to ask:

- What is $G(K/L)$, the Galois group of K over L?
- What are the intermediate fields L between F and K?

We know that the subgroups of G correspond to the intermediate fields L by the main theorem, but the correspondence may not be easy to deduce explicitly.

Example 272

Let's go back to the case where the degree of f is 3.

Then we either have $G = A_3$ if $[K : F] = |G| = 3$, or we have $G = S_3$ if $[K : F] = 6$. In this case, the **degree** of the field extension tells us a lot. (Unfortunately, for higher degrees, it's about as hard to find the order of the group $|G|$ as it is to find $[K : F]$.)

In particular, we found that $G = A_3$ if and only if the discriminant of f is a square in F: this gives us a good answer to the first question above.

Question 273. *What do we know about G in general, though?*

It's true that G permutes our roots $\alpha_1, \cdots, \alpha_n$, and the operation of G on the roots is **faithful**: fixing the roots fixes the entire field K. This means we have an injective homomorphism from G to S_n, and therefore G is a subgroup of S_n. But we actually know a little bit more:

Proposition 274

The operation of G on $\{\alpha_1, \cdots, \alpha_n\}$ is **transitive** (they form an orbit). In other words, for all $1 \leq i \leq n$, there exists a $\sigma \in G$ such that $\sigma(\alpha_1) = \alpha_i$.

Proof. Let $\{\alpha_1, \cdots, \alpha_k\}$ be the orbit of α_1, and let $h(x) = (x - \alpha_1) \cdots (x - \alpha_k)$. This is a factor of $f(x)$, and the coefficients of h are symmetric functions of $\alpha_1, \cdots, \alpha_k$. But if they form an orbit, every σ in G permutes $\alpha_1, \cdots, \alpha_k$, which means it fixes the symmetric functions.

But the fixed field of G is F, and therefore the coefficients of $h(x)$ are forced to be in F! Now since $f(x)$ is irreducible, we must have $f = h$, meaning the orbit of α_1 is indeed $\{\alpha_1, \cdots, \alpha_n\}$. □

So G is isomorphic to a transitive subgroup of S_n, and that's how we'll identify our subgroups. One important point: every element of S_n permutes the roots, but **not all permutations are automorphisms** of the field.

Example 275

Now let's do the case where the degree of f is 4.

The Galois group is now a transitive subgroup of S_4, which operates on $\{\alpha_1, \cdots, \alpha_4\}$. Since the orbit has order 4, $|G|$ must be a multiple of 4 and divide 24: the options are $4, 8, 12$, and 24.

What is the group in each case? $G = S_4$ for order 24 and A_4 for order 12 (these are the only subgroups of those sizes). The group of order 8 exists is one of the D_4s, and the group of order 4 can either be C_4 or D_2. (Remember that $[K : F] = |G|$ in all of these cases.)

Well, we have our discriminant

$$D = (\alpha_1 - \alpha_2)^2(\alpha_1 - \alpha_3)^2(\alpha_1 - \alpha_4)^2(\alpha_2 - \alpha_3)^2(\alpha_2 - \alpha_4)^2(\alpha_3 - \alpha_4)^2.$$

This has degree 12 in the αs, and we can think about permuting the roots. Switching any two of the roots keeps D constant, but what does it do to

$$\sqrt{D} = \delta = (\alpha_1 - \alpha_2)(\alpha_1 - \alpha_3) \cdots (\alpha_3 - \alpha_4)?$$

It turns out that odd permutations take $\delta \to -\delta$, and even permutations take $\delta \to \delta$, just like in the cubic case. So δ is fixed only by even permutations, and therefore, D is a square in our field F if and only if G only contains even permutations.

Thus, D is a square if and only if $\boxed{G \subset A_4}$, which means we have one of the groups A_4 and D_2.

Fact 276

Lagrange wrote a long paper on quartic equations, and the only thing that people remember from it is to figure out whether the Galois group has an element of order 3.

Consider the three elements

$$\beta_1 = \alpha_1\alpha_2 + \alpha_3\alpha_4, \beta_2 = \alpha_1\alpha_3 + \alpha_2\alpha_4, \beta_3 = \alpha_1\alpha_4 + \alpha_2\alpha_3.$$

Under S_4, $\{\beta_1, \beta_2, \beta_3\}$ is the orbit of β_1. Now

$$g(x) = (x - \beta_1)(x - \beta_2)(x - \beta_3)$$

has symmetric coefficients in the β_is, so it's also symmetric in the α_is. For example, $\beta_1+\beta_2+\beta_3$ is the sum of pairwise products $\alpha_i\alpha_j$, which is just the x-coefficient of $f(x)$ (the second symmetric function).

So now in the special case where $f(x)=x^4-px+q$ (so $p=s_3(\alpha)$ and $q=s_4(\alpha)$), we have that

$$g(x)=x^3-s_1(\beta)x^2+s_2(\beta)x-s_3(\beta).$$

$s_1(\beta)$ has degree 2 in the αs, $s_2(\beta)$ has degree 4, and $s_3(\beta)$ has degree 6. Thus $s_1(\beta)$, which has degree 2, needs to be a combination of $s_1(\alpha)$ and $s_2(\alpha)$ – both are zero in this case, so it's zero! $s_2(\beta)$ can only be written as a combination of s_3 and s_4 with degrees 3 and 4 respectively, so $s_2(\beta)=aq$, and similarly $s_3(\beta)=bp^2$ for some constants a,b.

Much like in the problem set, we can now compute two particular polynomials to find the values of a and b: taking $f(x)=x^4-1$, we have roots $1,i,-1,-i$, and now $\beta_1=2i,\beta_2=0,\beta_3=-2i$. This gives us that $a=-4$, and similarly we can take $f(x)=x^4-x$ with roots $0,1,\omega,\overline{\omega}$, and we can find that $b=-1$.

So when $f(x)=x^4-px+q$, we have that $\boxed{g(x)=x^3-4qx-p^2}$. Now we can ask whether $g(x)$ is irreducible in our field F: if it is, then $F\subset F(\beta_1)\subset K$, so 3 divides $[K:F]=|G|$, and therefore we're either in A_4 or S_4. On the other hand, if g is reducible, then 3 doesn't divide the order of the group, which gives us one of the other cases.

There's just one bug in this: **what if two of the β_is are equal**? Then we lose control, since we need to make sure the group acts on three elements! Luckily, the βs are distinct, because the discriminant can be written as

$$D(g(x))=(\beta_1-\beta_2)(\beta_1-\beta_3)(\beta_2-\beta_3)$$

and now we have a miracle: $\beta_1-\beta_2=(\alpha_1\alpha_2+\alpha_3\alpha_4)-(\alpha_1\alpha_3+\alpha_2\alpha_4)=(\alpha_1-\alpha_4)(\alpha_2-\alpha_3)$, which means that $D(g)=D(f)\neq 0$. Therefore our roots are indeed distinct!

So now we can make a table:

	g is irreducible	g is reducible
D square	A_4	D_2
D not a square	S_4	D_4 or C_4

It's not that easy to figure out that last ambiguity in general, but in general we should expect G to be S_4.

Example 277

What's the Galois group of the splitting field of x^4+x+1 over $\mathbb{Q}$?

We have $g(x)=x^3-4x-1$, and both f and g here are indeed irreducible. They have the same discriminant

$$D(g)=-4(-4)^3-27(-1)^2=256-27=229,$$

which is not a square. That means that the Galois group G is the symmetric group S_4.

Example 278

What if we have the irreducible polynomial $f(x)+x^4+x^3+x^2+x+1$ over $\mathbb{Q}$?

We have the roots $\zeta,\zeta^2,\zeta^3,\zeta^4$, where $\zeta=e^{2\pi i/5}$ is a fifth root of unity. Then $K=F(\zeta)$ has degree 4 over $\mathbb{Q}$, and any automorphism is determined by where σ goes (since the other roots are just powers of it)! For example, the automorphism $\zeta\to\zeta^2$ sends

$$\zeta\to\zeta^2\to\zeta^4\to\zeta^3.$$

That means $\sigma=(1243)$ generates the Galois group, and thus the group is cyclic of order 4.

Example 279
What if we want the splitting field of $(x^2-2)(x^2-3)$ over $\mathbb{Q}$?

Then $\alpha = \sqrt{2}$ and $\beta = \sqrt{3}$, and now $\gamma = \alpha + \beta$ is a primitive element: it's a root of $x^4 - 10x^2 + 1$, which is irreducible. But now $\gamma^2 = 5 + 2\sqrt{6}$, and we can similarly find that G is D_2 using the chart.

We'll have a quiz on Friday!

35 May 13, 2019

(We got more cookies in class today.) Today's topic is **adjoining roots of unity** of the form $e^{2\pi i/p}$, where p is a prime. First of all, let's review the Main Theorem: if we have a Galois extension of F with Galois group G, the intermediate fields between F and K correspond exactly to subgroups of G. Specifically, the correspondence is

$$F \subset L = K^H \subset K$$

$$G \supset H = G(K/L) \supset \{1\},$$

where K/L is a Galois extension with Galois group of order $|H| = [K : L]$, and therefore $[L : F] = [G : H]$.

Proposition 280
L is a Galois extension of F if and only if H is a normal subgroup of G. In this case, we have $G(L/F) \cong G/H$.

Example 281
If K is the splitting field of an irreducible cubic f, and the Galois group is S_3, then S_3 contains A_3 and three cyclic groups of order 2 (generated by transpositions).

Then K^{A_3} is a degree 2 extension over F: all such extensions are Galois, since all subgroups of index 2 are normal. But the other extensions are not Galois, because the transpositions do not generate normal subgroups of S_3.

Example 282
Now let's take $F = \mathbb{Q}$ and $K = F(\zeta_p)$.

We know that ζ_p is a root of $x^{p-1} + \cdots + x + 1$, which is irreducible over F (as derived earlier in class). We have the roots $\zeta, \zeta^2, \cdots, \zeta^{p-1}$, and $[K : F] = p - 1$.

But all of those roots "look the same" from the perspective of $\mathbb{Q}$: specifically, there is a unique automorphism sending ζ to any ζ^s for all $1 \le s \le p-1$. So those describe all of the automorphisms: we should interpret the exponent a in ζ^a as being an element of $\mathbb{F}_p$. After all, if $\sigma(\zeta) = \zeta^s, \tau(\zeta) = \zeta^t$,

$$\sigma\tau(\zeta) = \sigma(\tau(\zeta)) = \sigma(\zeta^t) = (\sigma(\zeta))^t = \zeta^{st}.$$

So composition of automorphisms corresponds to multiplying the exponents (and reducing mod p). Therefore, the Galois group is of the form $G \cong \mathbb{F}_p^\times$, which is a cyclic group of order $p-1$.

Example 283

Let $p = 13$: 2 is a primitive root in $\mathbb{F}_{13}^\times$.

Then the powers of 2 are

$$1, 2, 4, 8, 3, 6, 12, 11, 9, 5, 10, 7.$$

So the Galois group is a group of order 12 generated by σ, which sends $\sigma(\zeta) = \zeta^2$ (then ζ^2 get sent to ζ^4, and so on).

Every subgroup of a cyclic group is cyclic, and we need the order to divide 12. To draw out our diagram of fields, note that the subgroups are nested via

$$\langle \sigma^{12} = 1 \rangle \subset \langle \sigma^6 \rangle \subset \langle \sigma^3 \rangle \subset \langle \sigma \rangle$$

and

$$\langle \sigma^{12} = 1 \rangle \subset \langle \sigma^4 \rangle \subset \langle \sigma^2 \rangle \subset \langle \sigma \rangle,$$

with the additional restriction that $\langle \sigma^6 \rangle \subset \langle \sigma^2 \rangle$. Let's do some example computations!

First of all, **how can we find the intermediate field corresponding to σ^2?** It sends ζ to ζ^4, so we care about the orbits

$$[1, 4, 3, 12, 9, 10][2, 8, 6, 11, 5, 7].$$

Let $\alpha_1 = \zeta + \zeta^4 + \zeta^3 + \zeta^{12} + \zeta^9 + \zeta^{10}$, and let α_2 be the sum of the other roots of unity. Then σ sends α_1 to α_2 and vice versa, so let's calculate

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

Then $\alpha_1 + \alpha_2 = -1$ (it's all the roots of unity except 1), and $\alpha_1\alpha_2$ is the sum of 36 different terms. There's no zeros mod 13 (because k and $13 - k$ always appear in α_1 or in α_2), and now σ fixes the product $\alpha_1\alpha_2$, and is therefore in F.

Lemma 284

The only combination of the roots of unity satisfying

$$\sum_{i=0}^{p-1} c_i \zeta^i = 0$$

is $c_0 = c_1 = \cdots = c_{p-1}$.

In other words, the sum of the nontrivial roots of unity can only be rational if they all have equal contribution! this means we get 3 copies of all 12 roots, which means the total sum is -3. That means α_1, α_2 are roots of the polynomial

$$x^2 - x + 3,$$

and thus $\alpha_{1,2} = \frac{-1 \pm \sqrt{13}}{2}$, and our intermediate field is $\mathbb{Q}(\sqrt{13})$.

Next, **how can we find the intermediate field corresponding to σ^3?** We have the orbits

$$\sigma^3 = [1, 8, 12, 5][2, 3, 11, 10][4, 6, 9, 7].$$

Calling the elements $\beta_1 = \zeta + \zeta^8 + \zeta^{12} + \zeta^5$ and similarly for the other orbits, we know that σ permutes $\beta_1, \beta_2, \beta_3$ cyclically. That means the β_is are roots of the polynomial

$$(x - \beta_1)(x - \beta_2)(x - \beta_3).$$

We know that $s_1(\beta) = -1$, $s_2(\beta)$ is a sum of $4^2 \cdot 3 = 48$ terms, none of which are 1, so it's -4, and $s_3(\beta)$ is a product of 64 terms, so there are some number of zeros. It's unlikely there's 16 or more, so there are 4 (technically), and now $s_3(\beta) = -\frac{60}{12} + 4 = -1$. That means the βs are roots of

$$x^3 + x^2 - 4x + 1.$$

We know that the Galois group $\langle \sigma^3 \rangle$ has order 3, so we can plug this polynomial into the formula for the discriminant: it'll turn out to be a square.

Finally, **how can we find the intermediate field corresponding to σ^6?** We have orbits

$$[1, 12][2, 11][4, 9][8, 5][3, 10][6, 7].$$

Calling the sums of the powers of ζs here $\gamma_1, \cdots, \gamma_6$, we know that $\gamma_1 = \zeta + \zeta^{-1} = 2\cos\frac{2\pi}{13}$. We can find the product $(x - \gamma_1) \cdots (x - \gamma_6)$, but that's ugly and sad.

Example 285

Instead, let's do an analogous calculation for $p = 7 - 3$ is a primitive root there.

This gives us the orbit $[1, 3, 2, 6, 4, 5]$, and now we'll find the **intermediate field corresponding to σ^3**, which gives us the orbits

$$[1, 6][3, 4][2, 5].$$

Now it's reasonable to compute $(x - \gamma_1)(x - \gamma_2)(x - \gamma_3)$. Remember that we have the relation

$$(a + b)^3 = (a^3 + b^3) + 3(a^2 b + ab^2),$$

and since $\gamma_1 = \gamma = \zeta + \zeta^{-1}$,

$$\gamma^3 = (\zeta^3 + \zeta^{-3}) + 3(\zeta + \zeta^{-1}),$$

and

$$\gamma^2 = (\zeta^2 + \zeta^{-2}) + 2.$$

This yields a linear relation between $1, \gamma, \gamma^2, \gamma^3$:

$$\gamma^3 + \gamma^2 - 2\gamma - 1 = 0.$$

Example 286

To finish, let's do $p = 17$: notably, $p - 1$ is a power of 2, so we can construct a regular 17-gon. We have a primitive root 3 here.

σ permutes the roots via

$$[1, 3, 9, 10, 13, 5, 15, 11, 16, 14, 8, 7, 4, 12, 2, 6].$$

σ^2 splits this into two orbits:

$$[1, 9, 13, 15, 16, 8, 4, 2][3, 10, 5, 11, 14, 7, 12, 6].$$

Denoting the sums α_1 and α_2, $\alpha_1 + \alpha_2 = -1$, and $\alpha_1\alpha_2$ is the sum of 64 terms that are all not 1, and therefore it's $-\frac{64}{16} = -4$. So our polynomial

$$x^2 + x - 4 \implies \alpha_i = -\frac{1 + \sqrt{17}}{2},$$

so our intermediate field is $\mathbb{Q}(\sqrt{17})$.

Next, if we want the field corresponding to σ^4, we take

$$[1,13,16,4][9,15,8,2][3,5,14,12][10,11,7,6].$$

Note that $\beta_1+\beta_2=\alpha_1$ here, and $\beta_1\beta_2$ has 16 terms with no zeros that turn out to be -1. That means β_1,β_2 are roots of $x^2-\alpha_1 x-1$: in other words, $[F(\beta):F(\alpha)]=2$, and now we can just compute using the quadratic formula to find that

$$\beta=\frac{1}{2}\cdot\left(\frac{-1+\sqrt{17}}{2}+\sqrt{\frac{17+\sqrt{17}}{2}}\right).$$

Finally, what about the field corresponding to σ^8, which contains $\cos\frac{2\pi}{17}$? Similarly, if $\gamma_1=\zeta+\zeta^{16}$ and $\gamma_2=\zeta^{13}+\zeta^4$, then $\gamma_1+\gamma_2=\beta_1$ and $\gamma_1\gamma_2=[14,5,12,3]=\beta_3$. So $\gamma_1\gamma_2$ is a root of $x^2-\beta_1 x+\beta_3$, and that means we can write out (in nested square root form) the value of γ as well.

We'll talk about the quintic equation next time!

36 May 15, 2019

Today, we're going to not solve the quintic equation.

Proposition 287 (Cardano's formula)

If we have a cubic $f(x)=x^3+3px+2q$, then there is a root

$$\alpha=\sqrt[3]{-q+\sqrt{q^2+p^3}}+\sqrt[3]{-q-\sqrt{q^2+p^3}}.$$

This is useless for our purposes (and also in general) though. Let's try to do a study of this problem that's not so stupid:

Definition 288

A polynomial $f(x)$ is **solvable** in terms of field extensions if there exists a chain

$$F=F_0\subset F_1\subset\cdots\subset F_k$$

such that $F_{i+1}=F_i(\sqrt[p_i]{a_i})$ for some prime p_i and some $a_i\in F_i$, and f has a root in the last extension F_k.

This looks a little bit more general! Note that each extension can be restricted to having a prime p in the index of the square root, because $\sqrt[pq]{a}=\sqrt[p]{\sqrt[q]{a}}$.

Proposition 289

Suppose we have an element $a\in F$, and $\alpha^p=a$ for some prime p. If $\zeta_p=e^{2\pi i/p}$ is in F, then $K=F[\alpha]$ is the splitting field of x^p-a over F: the roots are just $\alpha,\zeta\alpha,\zeta^2\alpha$, and so on.

(Note that a priori we don't necessarily know that the polynomial is even irreducible.)

Proof. If G is our Galois group, then any $\sigma\in G$ needs to send a root to another root, so $\sigma\alpha=\zeta^s\alpha$ for some s. If $\tau\in G$ as well, and $\tau\alpha=\zeta^t\alpha$, then

$$\sigma\tau(\alpha)=\sigma(\zeta^t\alpha)=\zeta^t\sigma(\alpha)=\zeta^{s+t}\alpha$$

since we assume $\zeta \in F$. This means that G must be isomorphic to a subgroup of the additive group $\mathbb{F}_p^+$: since p is prime, the only possibilities are the identity or the whole group.

This means that if α is not in F, the extension isn't trivial, so G is cyclic of order p. Therefore, if $\zeta_p \in F, a \in F$, then $x^p - a$ is either irreducible or it splits completely in $F[x]$. □

(ζ needs to be in F: for example, consider $x^3 - 8$ as a counterexample otherwise.)

And if we want ζ_p (to satisfy the assumptions above), we can just adjoin it to our field. Let $F = \mathbb{Q}$, and let $K = F(\zeta_p)$ for some prime p. Then K is a Galois extension with Galois group $\mathbb{F}_p^\times$, which is cyclic. Thus, the idea is that we can always reach our final field K (for a solvable polynomial) with a sequence $F = F_0 \subset F_1 \cdots \subset F_k = K$, where F_{i+1} is a Galois extension of F_i with (cyclic) **prime order**.

Example 290

We can get $\mathbb{Q}(\zeta_5)$ with the following method:

We want to use two quadratic extensions to get a Galois group of order $5 - 1$. It turns out we (probably) use

$$F_0 = \mathbb{Q} \subset \mathbb{Q}(\sqrt{5}) \subset \mathbb{Q}\left(\sqrt{5}, \sqrt{-(5+\sqrt{5})}\right).$$

This means that whenever we want to figure out whether we can adjoin roots, we can just start by putting in all of the roots of unity in our field F that we need. In other words, **if f is solvable, then there exists a chain of fields**

$$F = F_0 \subset F_1 \subset \cdots \subset F_k,$$

where F_{i+1} is a Galois extension of F_i of prime degree: $G(F_{i+1}/F) = C_{p_i}$ for some prime p_i, and f has a root in F_k. (The converse is true, but let's not worry about that.)

Example 291

Let's prove that it's possible to solve quartics (at least theoretically).

Start with $F_0 = F$, and adjoin a cube root of 1 (so $F_1 = F(\zeta_3)$). Then $F_2 = F_1(\sqrt{D})$ adjoins the square root of the discriminant, and $F_3 = F_2(\beta_1)$, where $\beta_1, \beta_2, \beta_3$ are the roots of the resolvant cubic that we discussed earlier: $\beta_1 = \alpha_1\alpha_2 + \alpha_3\alpha_4$ and so on. Now if K is the splitting field of F with ζ_3 adjoined, the Galois group of K over F_2 is contained in the alternating group (because $\sqrt{D}$ is now fixed). Then, when we adjoin β_1, we can no longer have an element of order 3 for future extensions (because our βs are fixed), so $G(K/F_3) \subset D_2$. That means we can get to K with at most 2 more square roots!

Theorem 292

Let f be a polynomial have degree 5 in $F[x]$, and let K be the splitting field of f. If $G = G(K/F)$ is either S_5 or A_5, then f is not solvable.

Proof. We can assume $G = A_5$, because otherwise we can just adjoin the square root of the discriminant. Suppose for the sake of contradiction that such a chain does exist:

$$F = F_0 \subset F_1 \subset \cdots \subset F_k.$$

Let $K = K_0$ be the splitting field of f in F_0, let K_1 be the splitting field of f over F_1, and so on. Then we have $K_0 \subset K_1 \subset \cdots \subset K_k$.

Our goal is to show that the Galois group is $G(K_k/F_k) = G(K/F) = A_5$. If we show that, then the polynomial can't have a root in K_k, because if it did, the Galois group would not contain a 5-cycle! That would give us a contradiction, because K_k is supposed to be the splitting field of f over F_k.

We do this by induction: it's enough to show that $G(K'/F') = G(K/F)$ for one level of the chain. Note that F' is an extension of F with Galois group C_p for some prime p, and then we have the Galois groups $G(K/F) = A_5$ and $G'(K'/F')$. Since F' is a Galois extension of F as well, it's the splitting field of some polynomial g over F. Then K' is the splitting field of f over F', but it's also the splitting field of g over K (because we basically adjoin everything from f and also from g). That means that K' is the splitting field of fg over F, and therefore K'/F is Galois with some Galois group $\mathcal{G}$.

But by the Main Theorem, K is an intermediate field in the chain

$$F \subset K \subset K',$$

so $K = K'^N$, where N is some normal subgroup of $\mathcal{G}$ (importantly, it's normal because K/F is Galois). This means G is isomorphic to $\mathcal{G}/N$. On the other hand, we also have

$$F \subset F' \subset K',$$

so $F' = K'^{G'}$, where G' is some normal subgroup of $\mathcal{G}$. That means C_p is isomorphic to $\mathcal{G}/G'$.

So now we can map $\phi : \mathcal{G} \to G \times C_p$ by taking the residue in each quotient map. The kernel of this map, $\ker \phi$, consists of all $\sigma \in \mathcal{G}$ that operates trivially on both K and F', but K and F' generate K'! Therefore, the kernel is trivial and ϕ is injective, and now the order of the group $\mathcal{G}$ must be either $|G|$ or $p|G|$.

If we're in the case where $|\mathcal{G}| = |G|$, then we have a bijective map, and $\mathcal{G} = A_5$ (which is simple). But then it can't map surjectively to C_p because that would imply C_p is a quotient of A_5, which is a simple group! Therefore $|\mathcal{G}| = p|G|$, and we have a bijective map $\phi : \mathcal{G} \to G \times C_p$. Therefore the map from G to C_p is just the projection $G \times C_p \to C_p$, which has kernel G'. But the kernel from $\mathcal{G} \to C_p$ is G, so G is isormorphic to G'. Thus $G' = A_5$, and that means that we've made no progress towards solving the quintic. □

Galois then wanted to write down a polynomial whose Galois group was actually A_5 or S_5. He proved the following lemma somehow:

Lemma 293

Let G be a subgroup of S_5 (actually true for S_p) that contains a 5-cycle and a transposition. Then $G = S_5$.

Proof. This can be basically directly verified: just take various powers of the 5-cycle composed with the transposition. We can find a 3-cycle, a 4-cycle, and a 5-cycle, so the order of the group divides $3 \cdot 4 \cdot 5 = 60$. It's not A_5 (because there exists an odd permutation), so it is S_5. □

To finish the class, let's write down a quintic that is not solvable. Take $F = \mathbb{Q}$: note that $x^5 - 16x = x(x^2+4)(x^2-4)$ has three real roots. If we add a little constant, like $\boxed{x^5 - 16x + 2}$, we get an irreducible polynomial by Eisenstein. We still have 3 real roots (because we haven't perturbed our function too much).

The Galois group operates transitively on 5 roots, so it contains a 5-cycle. Now $F \subset F(\alpha_1, \alpha_2, \alpha_3) \subset K$, where $\alpha_1, \alpha_2, \alpha_3$ are the real roots of f. Since $K/F(\alpha_1, \alpha_2, \alpha_3)$ has degree 2, there must be an element of the Galois group that switches the two complex roots. So we have a 5-cycle and a transposition, and thus the Galois group is S_5!