Math 221 : Algebra cumulative notes

Alison Miller

1 Rings and Modules: basic definitions

In class these definitions were introduced when needed; here I'm just leaving them all up front for reference.

Definition. A ring A is a set with operations $+, \cdot$ and elements 0 and 1 such that $+$ makes A into an abelian group with identity 1, $\cdot$ is associative and satisfies $1 \cdot a = a \cdot 1 = a$, and $\cdot$ distributes over $+$.

We say that A is commutative if $\cdot$ is commutative.

Until further notice: rings are assumed to be commutative unless stated otherwise.

Definition. A subring B of A is a subset B of A which is a ring with the same $+, \cdot, 0, 1$.

Definition. A morphism of rings is...

Definition. An *algebra* over a ring R is a ring A along with a morphism $\phi : R \to A$.

(This definition is most often used when $R = k$ is a field. In this case, if $A \neq 0$, then ϕ is an injection, and is usually identified with the inclusion map.)

Definition. A subring $B$of A is a subset of A which is also a ring (with the same 1). It is said to be generated by a subset $S \subset A$ if it is the smallest subring of A containing S

Ditto sub-algebra.

Definition. An ideal I of A is an additive subgroup of A which is also closed under multiplication by A. It is said to be generated by a subset $S \subset A$ if is the smallest ideal of A containing S. (In which case, it is equal to the set of all linear combinations $\sum_k a_k s_k$, $a_k \in A$, $s_k \in S$.)

Definition. An A-module is an additive group M with a map $A \times M \to M$ denoted by $a, m \mapsto am$ such that both distributive laws hold, $a(bx) = (ab)x$, and $1x = x$.

Again, say that M is generated by a finite subset S if...

Definition. A morphism of A-modules is a map $\phi : M \to N$ such that $\phi(x + y) = \phi(x) + \phi(y)$ and $\phi(ax) = a\phi(x)$.

Note that kernels and images of module morphisms are submodules.

Definition. An exact sequence of modules is...

2 Hilbert's Basis Theorem

2.1 Historical Background – invariant theory

Setup of invariant theory: V is a vector space (over $\mathbb{C}$, although this can be done over any field), and G a group that acts on V via linear transformations. (That is, there's a map $(g, v) \mapsto gv$ which is linear in v and satisfies $(gh)v = g(hv)$.)

Example 1. $V = \mathbb{C}^2$, $G = SO_2(\mathbb{C})$, acting by rotations on $\mathbb{C}^2$.

Example 2. $V = S^n$, $G = \mathbb{C}^n$, acting by permuting entries.

Question 1. *Which polynomial functions p on V are invariants with respect to the G-action, that is, satisfy $p(gv) = p(v)$ for all $g \in G$ and $v \in V$.*

An equivalent phrasing of this is: by putting coordinates on $V \cong \mathbb{C}^n$, can identify the ring of polynomial functions on V with $\mathbb{C}[x_1, \ldots, x_n]$. This ring has a (left) action of G, given by $gp = p \circ g^{-1}$ (the inverse is to make it a left action). Then the ring of invariants is the subring of all polynomials $p = \mathbb{C}[x_1, \ldots, x_n]^G$ fixed by all elemnts of g under this action.

In Example 1, the ring of invariants is generated by $x_1^2 + x_2^2$.

In example 2, the ring of invariants is the ring of elementary symmetric functions (see HW 1).

Theorem 2.1. *Let I be the ideal of $A = \mathbb{C}[x_1, \ldots, x_n]$ generated by the homogeneous positive degree elements of A^G. Then if homogeneous invariant generators $i_1, \ldots, i_k \in A^G$ generate I as an R-ideal, they also generate i^G as a $\mathbb{C}$-algebra*

Proof. Suppose $a \in R^G$ – must show a in sub $\mathbb{C}$-algebra generated by I.

Induct on deg a: base case deg a =0 trivial.

WLOG a is homogeneous, of degree d. Then $a = \sum_{i=1}^{k} c_k i_k$, where $c_k \in A$, $i_k \in I$.

By looking at the degree d pieces on both sides, we may assume that b_k is homogeneous of degree d.

Now, because $r \in R^G$, we have

$$a = \frac{1}{|G|} \sum_{g \in G} ga = \sum_k c_k' i_k$$

where $c_k' = \frac{1}{|G|} \sum_{g \in G} g c_k \in A^G$.

Apply induction hypothesis to c_k'. □

Now we would be done if we knew how to prove that the ideal I of A has a finite generating set. Fortunately, for us, any ideal of $A = C[x_1, \ldots, x_n]$ is finitely generated. To show this we'll develop the theory of Noetherian rings.

(You might be worried here; what if the generators we find for I are in A but not in A^G? This is OK, because by definition of I we know that each of them can be written as a finite A-linear combination of elements of A^G, so we're good.)

Hilbert Basis Theorem

Definition. An A-module M is finitely generated if there exists a finite generating set S for M.

Proposition 2.2. *TFAE:*

- *Every ascending chain* N_i *of submodules of* M *eventually stabilizes.*
- *Every subset of modules of* M *has a maximal element*
- *Every submodule of* M *is f.g.*

Proof. (i) $\Rightarrow$ (ii) by contrapositive (ii) $\Rightarrow$(iii) the family of f.g. submodules of M has a maximal element (iii) $\Rightarrow$ (i) $\bigcup_i N_i$ has finitely many generators, all in some N_n. $\square$

We say that A is noetherian if A is noetherian considered as a module over itself.

Example. Fields are noetherian.

Our next goal is to prove the Hilbert Basis theorem, by inductively showing that if A is noetherian then $A[x]$ is also. First, more on modules.

2.2 Modules: morphisms, short exact sequences, direct sums

Definition. A morphism of A-modules is a map $\phi : M \to N$ such that $\phi(x+y) = \phi(x) + \phi(y)$ and $\phi(ax) = a\phi(x)$.

Note that kernels and images of module morphisms are submodules, and that modules satisfy the First Isomorphism Theorem $\operatorname{Im}\phi \cong M/\ker\phi$.

Definition. A short exact sequence of modules is a sequence of modules

$$0 \to M \xrightarrow{\phi} M' \xrightarrow{\psi} M'' \to 0$$

with maps between them such that the kernel of each map is the image of the previous one, that is:

- $\ker\phi = 0$
- $\ker\psi = \operatorname{Im}\phi$
- $\operatorname{Im}\psi = M''$.

Example. if $M \subset M'$, $0 \to M \to M' \to M'/M \to 0$ is exact. If $\psi : M' \to M''$ is any surjective homomorphism, $0 \to \ker\psi \to M' \to M'' \to 0$ is exact.

Definition. The direct sum $M \oplus N$ of two A-modules is the set of pairs $\{(m, n)\}$ with $m \in M$, $n \in N$, where addition and multiplication by elements of A are defined componentwise. Can define arbitrary finite direct sums analogously, e.g. for any integer k, A^k is the direct sum of k copies of A.

Example. There is an exact sequence $0 \to M \xrightarrow{\phi} M \oplus N \xrightarrow{\psi} N \text{to} 0$, where $\phi(m) = (m, 0)$ and $\psi((m, n)) = n$.

Proposition 2.3. *If $0 \to M \to M' \to M'' \to 0$ is a short exact sequence of A-modules, then M' is noetherian iff both M and M'' are noetherian.*

Proof. On HW. □

In particular, this means that finite direct sums of noetherian A-modules are Noetherian.

Corollary 2.4. *If A is a noetherian ring, then A^k is noetherian for every integer A, and every finitely generated A-module is noetherian.*

Theorem 2.5 (Hilbert Basis)**.** *If A is a noetherian ring, then the polynomial ring $A[x]$ is also noetherian.*

Proof. (Note that the notation for the subscripts here is different from that used in class – k and m are switched.)

Suppose I is an ideal of $A[x]$ – we must show that it is finitely generated.

Let J be the ideal of A generated by the leading coefficients of elements of I. Then $J = \langle j_1, j_2, \ldots, j_m \rangle$ is finitely generated (and we can choose the generators to be leading coefficients of element of I). Choose $f_k \in I$ with leading coefficient a_k, for each $k = 1, \ldots, n$. Let d be the largest degree of any f_k.

Let I_d be the sub A-module of I consisting of elements of degree $< d$. It is contained in the A-module $\{\text{polynomials of degree } < d\} \cong A^d$, which is noetherian, so I_d has a finite generating set $i_1, i_2, \ldots, i_{m'}$.

Claim: J is generated as an $A[x]$-ideal by the f_k and the i'_k.

Proof: Choose $i \in I$. We must show that i is in the ideal I' generated by the f_k and the $i_{k'}$. Induct on degree of i. If the degree is $< d$, i is an A-linear combination of the i'_k, so we're done. Otherwise, suppose the leading term of i is jx^n, with $j \in J$. Then write $j = \sum_k c_k j_k$ with $c_k \in A$. Then $i - \sum_k c_k j_k x^{n - \deg j_k}$ has lower degree than i, hence is in I'. This implies that i is as well. By induction, we have $I = I'$.

□

Corollary 2.6. $k[x_1, \ldots, x_n]$ *is Noetherian for any field* k*.*

More generally, we have that

Corollary 2.7. *Any finitely generated* k*-algebra is noetherian.*

This follows from the previous corollary and

Lemma 2.8. *If $\phi : A \to B$ is a surjective ring homomorphism and A is noetherian, then B is noetherian.*

Proof. By the First Isomorphism Theorem for rings, we have that $B \cong A/\ker\phi$. The ideals of $A/\ker\phi$ are in order-preserving correspondence with the ideals of A containing $\ker\phi$. Therefore, the ascending chain condition on A implies the same for $A/\ker\phi$. □

3 More on ideals

As mentioned last time, the kernel of any homomorphism $A \to B$ is an ideal of A. And given any ideal A of I, the projection homomorphism $\pi : A \to A/I$ has kernel precisely I. By the first isomorphism theorem for rings, any other homomorphism from A with kernel I has image isomorphic to A/I: this means that there is a 1-1 correspondence between ideals of A and surjective homomorphisms $A \to B$.

Now we introduce a couple of properties of ideals $I \subset A$ that can be expressed as properties of the quotient A/I.

Definition. An ideal $\mathfrak{m}$ of A is maximal if and only if $A/\mathfrak{m}$ is a field if and only if A is maximal in the poset of nonzero ideals of A.

An ideal $\mathfrak{p}$ of A is prime if and only if $A/\mathfrak{p}$ is an integral domain if and only if $xy \in \mathfrak{p}$ implies $x \in \mathfrak{p}$ or $y \in \mathfrak{p}$.

All (non-zero) rings have maximal ideals, in fact:

Theorem 3.1. *Every proper ideal of A is contained in a maximal ideal. Every nonzero ring A has a maximal ideal (hence also a prime ideal).*

Proof. Zorn's lemma. □

Definition. Pullback of ideals: if I_B is an ideal of B, and $\phi : A \to B$ is a homomorphism, then $\phi^{-1}(I_B)$ is an ideal of A, referred to as the "pull-back" of I_B to A. (Some books, including Atiyah-Macdonald, call this the "contraction" of I_B.)

Example. if ϕ is the map $A \to A/J$, then pushforward and pullback give a 1-1 correspondence between ideals of A containing J and ideals of A/J.

Proposition 3.2. *Pullbacks of prime ideals are prime.*

Proof. If $\phi : A \to B$ is a ring homomorphism, and $\mathfrak{p}$ is a prime ideal of B, then $\phi^{-1}(\mathfrak{p})$ is the kernel of the composite homomorphism $\pi \circ \phi : A \to B \to B/\mathfrak{p}$. By the first isomorphism theorem, this means that $A/\phi^{-1}(\mathfrak{p})$ is a subring of $B/\mathfrak{p}$. Since a subring of an integral domain is an integral domain, we're done. □

4 Localization

Definition. Let A be a ring and S be a multiplicatively closed subset of A.

We construct a ring $S^{-1}A$, known as the "localization" of S at A, as the set of equivalence classes $(a, s) \in A \times S$ modulo the equivalence relation

$$(a, s) \sim (b, t) \Leftrightarrow \text{there exists } u \in S \text{ such that } uat = ubs.$$

We let a/s denote the equivalence class of the ordered pair (a, s)

Exercise: $S^{-1}A$ is a well-defined ring, and the map $j : A \to S^{-1}A$ given by $a \mapsto a/1$ is a ring homomorphism.

Example. Two common types of localizations: $S = \{1, s, s^2, ...\}$; in this case $S^{-1}A$ is also denoted $A[1/s]$.

$S = A - \mathfrak{p}$, where $\mathfrak{p}$ is a prime ideal – here $S^{-1}A$ is also denoted $A_{\mathfrak{p}}$ and called "the localization of A at $\mathfrak{p}$".

Proposition 4.1. *Universal property of localization: For any morphism $\phi : A \to B$ such that $\phi(S) \subset B^\times$, there's a unique morphism $\phi' : S^{-1}A \to B$ such that $\phi = \phi' \circ j$.*

Proof. Uniqueness: Take an arbitrary $a/s \in S^{-1}A$.. Then $\phi'(a/s)\phi'(s/1) = \phi'(a/1)$, so $\phi'(a/s)\phi(s) = \phi(a)$. Since $\phi(s)$ is a unit of B, this equation has a unique solution for $\phi'(a/s)$.

Existence: We now know that if such a morphism ϕ' exists, it must be defined by $\phi'(a/s) = \phi(a)\phi(s)^{-1}$. But we need to check that this actually gives us a well-defined map! To do this, note that if $a/s = b/t$, then there exists $u \in S$ with $uat = ubs$, and then

$$\phi(ust)(\phi(a)\phi(s)^{-1}) = \phi(uat) = \phi(ubs) = \phi(ust)(\phi(b)\phi(t^{-1}))$$

and canceling the unit $\phi(ust)$ from the both sides gives the desired result.

It's then a straightforward exercise to check that ϕ' is a ring homomorphism. □

Last time we showed the universal property:

Proposition 4.2. *Universal property of localization of rings: For any morphism $\phi : A \to B$ such that $\phi(S) \subset B^\times$, there's a unique morphism $\phi' : S^{-1}A \to B$ such that $\phi = \phi' \circ j$.*

Theorem 4.3. *The universal property determines the localization up to (unique) isomorphism: that is, if there is another ring C, equipped with a map $j_C : A \to C$ such that $j_C(S) \subset C^\times$ with the same property, then $C \cong S^{-1}(A)$. (and there's a unique choice of the isomorphism which is compatible with the maps out of A)*

Proof. Applying the universal property of $S^{-1}A$ to the map $j_C : A \to C$, we find that there is a unique map $\tilde{j_C}$ with $\tilde{j_C} \circ j = j_C$. Likewise, applying the universal property of C to the map J, we find that there is a unique map $\tilde{j}$ with $\tilde{j_C} \circ j = j_C$.

Now we use the uniqueness part of the universal property. Note that

$$(\tilde{j} \circ \tilde{j_C}) \circ j = \tilde{j} \circ j_C = j = \mathrm{id}_{S^{-1}A} \circ j$$

. Applying uniqueness in the universal property of $S^{-1}A$ (with $B = S^{-1}A$, $\phi = j$), we obtain $\tilde{j} \circ \tilde{j_C} = \mathrm{id}_{S^{-1}A}$. Likewise, using uniqueness in the universal property of C, we can get $\tilde{j_C} \circ \tilde{j} = \mathrm{id}_C$. Hence we've constructed isomorphisms between $S^{-1}A$ and C.

(And it's clear from the uniqueness in the universal property that these are the only isomorphisms compatible with the maps j and j_C.) □

Example. $\mathbb{Z}/6\mathbb{Z}$, localized at the set $\{1, 2, 4\}$. Then check that $\mathbb{Z}/3\mathbb{Z}$, with j, being the projection map $\mathbb{Z}/6\mathbb{Z} \to \mathbb{Z}/3\mathbb{Z}$, has the correct universal property. This is because any map $\phi : \mathbb{Z}/6\mathbb{Z} \to B$ such that $\phi(2) \in B^\times$ has 3 in its kernel.

Proposition 4.4. *Inclusion-preserving bijection: (primes of $S^{-1}A$) $\leftrightarrow$ primes of A disjoint from S.*

Proof. on HW. The $\to$ map is pullback under the map $j : A \to S^{-1}A$.

(The $\leftarrow$ map can also be described as sending a prime ideal $\mathfrak{q}$ of A to the ideal $j(\mathfrak{q})S^{-1}A$. This map is called "pushforward", and for general morphisms it doesn't have to send primes to primes, but for this choice of j it does.) □

A *local ring* is a ring with a unique maximal prime ideal.

Corollary 4.5. *If $\mathfrak{p}$ is a prime ideal, then $A_\mathfrak{p}$ is a local ring.*

(Terminology comes from the example $\mathbb{C}[t]_{(t)}$, where the localization is the ring of rational functions which are defined locally near 0.)

Now an example of the power of this, we'll prove a theorem that has nothing to do with localization.

Definition. The nilradical $\mathrm{nil}(A)$ of a ring is the set of nilpotent elements of A.

Theorem 4.6. $\mathrm{nil}(A)$ *is the intersection of all the prime ideals of A.*

Proof. a is nilpotent $\leftrightarrow A[a^{-1}] = 0 \Leftrightarrow A[a^{-1}]$ has no prime ideals $\Leftrightarrow$ every prime ideal of A contains a^i for some $i \Leftrightarrow$ every prime ideal of A contains a. □

Definition. Localization of modules: if M is an A-module, can define an $S^{-1}A$-module $S^{-1}M$ as the set of equivalence classes $(m, s) \in M \times S$, under the equivalence relation $(m_1, s_1) \sim (m_2, s_2)$ if $um_1s_2 = um_2s_1$ for some $u \in S$.

As before, we write m/s for the equivalence class of the pair (m, s).

Lemma 4.7. *An element $m \in M$ is mapped to 0 in $S^{-1}m$ if and only if $sm = 0$ for some $s \in S$.*

Proof. $m/1 = 0/1$ if and only if $u \cdot m \cdot 1 = u \cdot 0 \cdot 1$ for some $u \in S$. □

5 Hom

Definition. If M and N are A-modules, then the set of all A-module homomorphisms $\mathrm{Hom}_A(M,N)$ naturally forms an A-module with addition defined by $(\phi+\psi)(m) = \phi(m)+\psi(n)$ and $(a\phi)(m) = a(\phi(m))$. In settings when the ring A is clear, we will drop the A and write $\mathrm{Hom}(M,N)$.

Properties: $\mathrm{Hom}(A,N) \cong N$. $\mathrm{Hom}(A/I,N) = \{n \in N \mid in = 0 \text{ for all } i \in I\}$, called the "I-torsion submodule of N" and sometimes written $N[I]$.

Functoriality: if $\phi : M' \to M$ and $\psi : N \to N$ are A-module homomorphisms, there is an induced homomorphism of A-modules $\mathrm{Hom}(M,N) \to \mathrm{Hom}(M',N')$ given by $x \mapsto \psi \circ x \circ \phi$. This homomorphism is sometimes denoted $\mathrm{Hom}(\phi,\psi)$.

Recall definition of short exact sequences: the kernel of each homomorphism in the sequence is the image of the next one. Likewise, we can define exact sequences of any length.

Proposition 5.1. *If* $0 \to N \xrightarrow{\phi} N' \xrightarrow{\phi} N''$ *is a exact sequence, then*

$$0 \to \mathrm{Hom}(M,N) \xrightarrow{\mathrm{Hom}(\mathrm{id}_M,\phi)} \mathrm{Hom}(M,N') \xrightarrow{\mathrm{Hom}(\mathrm{id}_M,\psi)} \mathrm{Hom}(M,N'')$$

is also exact.

The proof of this is an easy exercise.

Example. However, it is not true that if $0 \to N \to N' \to N'' \to 0$ is an exact sequence that $0 \to \mathrm{Hom}(M,N) \to \mathrm{Hom}(M,N') \to \mathrm{Hom}(M,N'') \to 0$ is always exact.

As a counterexample, take the exact sequence $0 \to \mathbb{Z} \xrightarrow{\times p} \mathbb{Z} \xrightarrow{\pi} \mathbb{Z}/p\mathbb{Z} \to 0$, and take $M = \mathbb{Z}/p\mathbb{Z}$. Applying $\mathrm{Hom}(M,-)$ to each term, we get $0 \to 0 \to 0 \to \mathbb{Z}/p\mathbb{Z} \to 0$, which fails to be exact at the last step.

Proposition 5.2. *if* $M \xrightarrow{\phi} M' \xrightarrow{\psi} M'' \to 0$ *is an exact sequence, then*

$$0 \to \mathrm{Hom}(M'',N) \xrightarrow{\mathrm{Hom}(\phi,\mathrm{id}_M)} \mathrm{Hom}(M',N) \xrightarrow{\mathrm{Hom}(\psi,\mathrm{id}_N)} \mathrm{Hom}(M,N)$$

is exact.

Again, the proof is left as an exercise.

Example. Again, a counterexample to the assertion that $\mathrm{Hom}(M,-)$ always takes exact sequences to exact sequences. Again using the exact sequence

$$0 \to \mathbb{Z} \xrightarrow{\times p} \mathbb{Z} \xrightarrow{\pi} \mathbb{Z}/p\mathbb{Z} \to 0.$$

Now let $N = \mathbb{Z}$. We obtain $0 \to 0 \to \mathbb{Z} \xrightarrow{\times p} \mathbb{Z} \to 0$, which is not exact at the last step.

6 Tensor Products

We now will now define tensor products in terms of a universal property.

Definition. If M, N and P are A-modules, An A-bilinear map $\phi : M \times N \to P$ is a function such that for any $m \in M$, the map $n \mapsto \phi(m, n)$ is an A-module homomorphism, and for any n in N, the map $m \mapsto \phi(m, n)$ is an A-module homomorphism.

Definition (Universal property of the tensor product)**.** . The tensor product of two A modules M and N is an A-module $M \otimes_A N$ with bilinear map $B : M \times N \to M \otimes_A N$ such that for any bilinear map $\phi : M \times N \to P$, there exists a unique $\tilde{\phi} : M \otimes_A N \to P$ such that $\psi \circ b = \phi$.

Proposition 6.1. *The tensor product $M \otimes_A N$ exists and is uniquely determined by the universal property.*

Proof. Uniqueness will follow from existence by the standard universal property argument.

We'll construct the tensor product $M \otimes_A N$ as the quotient C/D of two A-modules. Let C be the free A-module on the set $M \times N$ (that is, the set of all finite linear combinations $a_1(m_1, n_1) + a_2(m_2, n_2) + \cdots + a_k(m_k, n_k)$).

Let D be the submodule generated by all elements of the following form:

$$(m_1 + m_2, n) - (m_1, n) - (m_2, n), \quad (am, n) - a(m, n)$$
$$(m, n_1 + n_2) - (m, n_1) - (m, n_2), \quad (m, an) - a(m, n).$$

Then C/D is our candidate for tensor product. Define $b : M \times N \to M \otimes_A N$ as the map sending (m, n) to the image of (m, n) in C/D.

Our construction ensures that b is bilinear. To show the universal property: for any bilinear map $\phi : M \times N \to P$, we may construct the map $C \to P$ that sends (m, n) to $\phi(m, n)$. Then bilinearity of ϕ implies that D is contained in the kernel of ϕ, so we obtain an induced map $C/D \to P$ with the desired property. □

The universal property of the tensor product is more important than the construction: but one thing to note from the construction is that $M \otimes_A N$ is generated by the images of all elements $m \otimes n$.

Properties of tensor product:

For any A-module M, $A \otimes_A M \cong M$ (here the map b is given by $b(am) = am$ – easy to check universal property).

Tensor product is compatible with direct sum. (proof uses universal property)

$A^n \otimes A^m \cong A^{nm}$ (write A^n as the direct sum of n copies of A, use associativity.)

Example. We claim that $\mathbb{Z}/2\mathbb{Z} \otimes_{\mathbb{Z}} \mathbb{Z}/3\mathbb{Z} \cong 0$. Since $M \otimes N$ is generated by elements of the form $m \otimes n$, enough to show that $m \otimes n = 0$ for any $m \in Z/2\mathbb{Z}$, $n \in \mathbb{Z}/3\mathbb{Z}$. But $m \otimes n = m \otimes 2(2n) = 2(m \otimes 2n) = 2m \otimes 2n = 0 \otimes 2n = 0(1 \otimes 2n) = 0$.

Last time we defined tensor products by their universal property. Note that not all elements of $M \otimes_A N$ are of the form $m \otimes n$ – however those elements do generate $M \otimes_A N$ as an A-module. Moreover, the universal property guarantees that if you want to define an A-module homomorphism ψ from $M \otimes_A N$ to P, you only need to specify the images $\psi(m \otimes n)$ of the elements $m \otimes n$, and as long as $\psi(m \otimes n)$ is defined in a way that is bilinear in m and n, it will extend uniquely.

7 Extension of Scalars and localization

Suppose that B is an A-algebra. Then the module tensor product is $B \otimes_A M$ is an A-module, but it can also be made into a B-module by the rule $b(b' \otimes m) = bb' \otimes m$.

Proposition 7.1. *We have $S^{-1}A \otimes M \cong S^{-1}M$ as $S^{-1}A$-modules.*

Proof. Define a map $S^{-1}A \otimes M \to S^{-1}M$ by $a/s \otimes m \mapsto (am)/s$.

Define a map $S^{-1}M$ to $S^{-1}A \otimes M$ by $m/s \mapsto \frac{1}{s} \otimes m$. We need to check this is well-defined: if $m/s = n/t$ then $mtu = nsu$ for some $u \in S$ so

$$\frac{1}{s} \otimes m = \frac{1}{stu} \otimes mtu = \frac{1}{stu} \otimes nsu = \frac{1}{t} \otimes n$$

.

Easy to check that these maps are inverses. □

8 Functoriality and Exactness

Now we are going to do the same type of functoriality/exactness thing we did for Hom is last lecture. Suppose that $\phi : M \to M'$ and $\psi : N \to N'$ are morphisms of A-modules. Then we can define a morphism $\phi \otimes \psi : M \otimes M' \to N \otimes N'$ of A-modules

Question 2. *If*

$$0 \to M \xrightarrow{\phi} M' \xrightarrow{\psi} M'' \to 0$$

is a short exact sequence, is

$$0 \to M \otimes N \xrightarrow{\phi} M' \otimes N \xrightarrow{\psi} M'' \otimes N \to 0$$

also exact?

Counterexample: no – take the short exact sequence $0 \to \mathbb{Z} \xrightarrow{\times p} \mathbb{Z} \xrightarrow{\pi} \mathbb{Z}/p\mathbb{Z} \to 0$, and take $N = \mathbb{Z}/p\mathbb{Z}$. Then we get an exact sequence that starts

$$0 \to \mathbb{Z}/p\mathbb{Z} \xrightarrow{\times p} \mathbb{Z}/p\mathbb{Z} \to \dots$$

and multiplication by p from $\mathbb{Z}/p\mathbb{Z}$ to itself is the zero map , so not injective.

However, it turns out that injectivity is the only place where this fails!

Proposition 8.1. *if* $M \xrightarrow{\phi} M' \xrightarrow{\psi} M'' \to 0$ *is exact, then*

$M \otimes N \xrightarrow{\phi} M' \otimes N \xrightarrow{\psi} M'' \otimes N \to 0$ *is exact.*

Proof. Surjectivity follows from surjectivity of ψ the fact that $M'' \otimes N$ is generated by elements of the form $m'' \otimes n$.

For exactness at the middle, it's enough to show $M'' \otimes N \cong M' \otimes N/\operatorname{Im}(\phi \otimes 1)$.

In the right-hand direction, take the map sending $m'' \otimes n$ to the image of $m' \otimes n$ in $M' \otimes N/\operatorname{Im}(\phi \otimes 1)$, for any choice of m' such that $\psi(m') = m''$. Since $\ker(\psi) \otimes N \subset \operatorname{Im}(\phi \otimes 1)$, this is well-defined.

In the left-hand direction, the map $\psi \otimes \mathrm{id}_N : M' \otimes N \to M'' \otimes N$ is 0 on the $\operatorname{Im}(\phi \otimes 1)$ so defines a map $M' \otimes N/\operatorname{Im}(\phi \otimes 1) \to M'' \otimes N$. Easy to see these maps are inverses. □

Corollary 8.2. *For any ideal* I *of* N*, and any* N*-module* M*, we have* $A/I \otimes N \cong N/IN$*.*

Proof. Tensor the exact sequence $I \to A \to A/I \to 0$ with N, and observe that the image of $I \otimes N \to A \otimes N$ is the submodule IN generated by all products $\{in \mid i \in I, n \in N\}$. □

We've observed that tensor product does not always take exact sequences to exact sequences. However, there are some modules N such that tensoring with N always takes exact sequences to exact sequences.

Definition. We say that an A-module N is flat if, for every exact sequence

$$0 \to M \to M' \to M'' \to 0$$

the sequence

$$0 \to N \otimes M \to N \otimes M' \to N \otimes M'' \to 0$$

is also exact.

Example. A is always a flat A-module. You'll show on HW that direct sums of flat modules are flat – this implies that A^n is also flat.

One important class of flat modules are localizations.

Theorem 8.3. *For any multiplicatively closed subset* S *of* A*,* $S^{-1}A$ *is a flat* A*-module.*

Proof. Just need to show that $0 \to M \xrightarrow{\phi} M'$ implies $0 \to S^{-1}M \xrightarrow{S^{-1}\phi} S^{-1}M'$. For this, suppose $a/s \in \ker S^{-1}\phi$. Then $\phi(a)/s = 0$ in $S^{-1}M'$, so there exists u such that $u\phi(a) = 0$. Then $\phi(ua) = 0$, so $ua \in \ker \phi$ implies $ua = 0$, so $ua = 0$ in $S^{-1}M$ and $ua/s = 0$ also. □

We're now going to move on to our next topic: integrality. The motivation for this comes from algebraic number theory.

Definition. An element $\alpha \in \bar{\mathbb{Q}}$ is said to be an algebraic integer if $p(\alpha) = 0$ for some monic polynomial $p \in \mathbb{Z}[x]$.

We'd like to show that the set of algebraic integers forms a ring. This is not simple, so we will develop some more module theory first.

9 Cayley-Hamilton Theorem and Nakayama's Lemma

Motivation: recall Cayley-Hamilton theorem from linear algebra. If ϕ is an endomorphism of a vector space $V \cong \mathbb{C}^n$, then ϕ satisfies its characteristic polynomial $\chi(x) = \det(xI_n - \phi)$. What this means is: the space $\mathrm{End}(V)$ of endomorphisms of V (linear maps V to itself) is a non-commutative ring with multiplication given by composition. In that ring we can evaluate the polynomial $\chi(\phi)$, and the theorem says that this is the zero endomorphism.

We now wish to generalize this to arbitrary modules. For any A-module $\mathrm{End}(M)$ is a non-commutative ring.

Theorem 9.1 (Cayley-Hamilton)**.** *Suppose that ϕ is an endomorphism of a finitely generated A-module M, and suppose that there is an ideal I of A such that $\phi(M) \subset IM$. Then there exists a monic polynomial $p(x)$ such that $p(\phi) = 0$ in $\mathrm{End}(M)$, and such that all non-leading coefficients of p belong to I.*

Proof. Although the ring $\mathrm{End}(M)$ is non-commutative, it contains a commutative subring $A[\phi]$ generated by ϕ, and we will do all calculations there. Note also that M is naturally a module over $A[\phi]$.

Choose generators $m_1, \ldots, m_n$ of M. Write $\phi(m_i) = \sum_j a_{ij} m_j$. Let $\mathcal{A}$ be the matrix with entries $\{a_{ij}\}$, and consider the matrix $\phi \cdot I_n - \mathcal{A}$.

We have

$$(\phi \cdot I_n - \mathcal{A}) \cdot \begin{bmatrix} m_1 \\ m_2 \\ \vdots \\ m_n \end{bmatrix} = 0.$$

Recall from linear algebra that there is a matrix $(\phi \cdot I_n - \mathcal{A})^{\mathrm{adj}}$ with entries in $A[\phi]$ (in fact its entries are, up to sign, principal minors of $\phi \cdot I_n - \mathcal{A}$) such that

$$(\phi \cdot I_n - \mathcal{A})^{\mathrm{adj}}(\phi \cdot I_n - \mathcal{A}) = \det(\phi \cdot I_n - \mathcal{A}) I_n.$$

Hence

$$0 = (\phi \cdot I_n - \mathcal{A})^{\mathrm{adj}}(\phi \cdot I_n - \mathcal{A}) \begin{bmatrix} m_1 \\ m_2 \\ \vdots \\ m_n \end{bmatrix} = \det(\phi \cdot I_n - \mathcal{A}) I_n \begin{bmatrix} m_1 \\ m_2 \\ \vdots \\ m_n \end{bmatrix} = \begin{bmatrix} \det(\phi \cdot I_n - \mathcal{A}) m_1 \\ \det(\phi \cdot I_n - \mathcal{A}) m_2 \\ \nu dots \\ \det(\phi \cdot I_n - \mathcal{A}) m_n \end{bmatrix}.$$

Hence the endomorphism $\det(\phi \cdot I_n - \mathcal{A})$ sends each of the generators m_i to 0, so is the 0 endomorphism. □

Lemma 9.2 (Nakayama's Lemma, first form)**.** *If M is a finitely generated A-module, and I is an ideal of A such $IM = M$ (equivalently, $M/IM \cong (A/I) \otimes_A M \cong 0$, then there exists $a \equiv 1 \pmod{I}$ such that $aM = 0$.*

Proof. Take ϕ to be the identity homomorphism id_M in the statement of the Cayley-Hamilton theorem above. Then we find that $\mathrm{id}_M^n + c_{n-1}\mathrm{id}_M^{n-1} + \cdots + c_0 \mathrm{id}_M = 0$ for coefficients $c_{n-1}, \cdots, c_0$ in I. Hence $a = 1 + c_{n-1} + \cdots + c_0$ has the desired property. □

Nakayama's Lemma is usually applies for specific ideals I of A. For instance:

Definition. The Jacobson radical $\mathrm{rad}(A)$ of a ring A is the intersection of all prime ideals of A.

I started class today by giving a different proof of Nakayama's lemma sencod form than the one I posted online in class notes at the end of class:

Lemma 9.3 (Nakayama's Lemma, second form)**.** *Assumptions as before, but also $I \subset \mathrm{rad}(A)$. Then we may conclude $M = 0$.*

Proof. As in the first case, we have that there is $a \equiv 1 \pmod{I}$ such that $aM = 0$. Now, note that $a - 1 \in I \subset \mathrm{rad}(A)$ is contained in all maximal ideals of A, so a cannot be contained in any maximal ideals of A. Since every proper ideal is contained in an maximal ideal, a cannot be contained in any proper ideal of A. But we do have $a \in (a)$, so the ideal (a) must be all of A. This means that a must be in A. Multiplying $aM = 0$ by a^{-1}, we find that $M = 0$. □

This version of the lemma is most often used in the case where A is a local ring, in which case it amounts to:

Lemma 9.4 (Nakayama's Lemma for local rings)**.** *If A is a local ring, and $\mathfrak{m}_A$ is the unique maximal ideal. Then for any finitely generated A-module M, $\mathfrak{m}M = M$ implies $M = 0$.*

10 Integral Extensions

Now we'll generalize the definition of algebraic integers from last time, to work for arbitrary rings.

Let B be an arbitrary A-algebra. For simplicity we will just deal with the case when A is actually a subring of B, though this can be generalized.

Definition. If B is an A-algebra an element $b \in B$ is said to be integral over A if b satisfies a monic polynomial equation $p(b) = 0$ with coefficients in A.

Proposition 10.1. *The following are equivalent for an element* $b \in B$.

a) b *is integral over* A

b) the sub A*-algebra* $A[b]$ *of* B *generated by* b *is a finitely generated* A*-module*

c) $A[b]$ *is contained in a subalgebra* C *of* B *which is a finitely generated* A*-module.*

d) there is a finitely generated faithful $A[b]$*-module* M *which is also finitely generated as an* A*-module.*

(here faithful means: if $e \in A[b]$ such that $eM = 0$, then $e = 0$.)

Proof.

a) $\Rightarrow$ b): Suppose that b is a root of $p(x) = x^n + a_{n-1}x^{n-1} + \cdots + a_0$. Take $1, b, \ldots, b^{n-1}$ as a generating set. Then using the relation $b^n = -a_{n-1}b^{n-1} + \cdots + a_0$, can write all higher powers of b in terms of lower powers.

b) $\Rightarrow$ c) $\Rightarrow$ d): automatic.

d) rightarrow a): Let ϕ be the endomorphism of M given by multiplication by b By C-H, there is a polynomial $p(x) \in A[x]$ such that $p(\phi)$ is the zero endomorphism of M. But $p(\phi)$ just acts by multiplication by $p(b)$, and since we assumed M is faithful, we must have $p(b) = 0$.

□

Definition. An A algebra B is said to be *finite over* A if B is a finitely generated A-module.

Proposition 10.2. B *is finite over* A *if and only if* B *is generated as an* A*-algebra by finitely many integral elements.*

Proof. Proof of $\Rightarrow$: if $b_1, \ldots, b_n$ generate B as an A-module they also generate as an A-algebra. And by c) $\implies$ a) in the Proposition **??**, we have that $b_1, \ldots, b_n$ are all integral over A.

Proof of $\Leftarrow$: this is a generalization of a) $\implies$ b) in Proposition **??**. Suppose that $b_1, \ldots, b_n$ are integral elements that generate B, and that $p_i(b_i) = 0$ where p_i is monic of degree n_i. Then the set of all monomials of the form $b_1^{d_i} \cdots b_n^{d_n}$ □

Theorem 10.3. *The set* $\{x \in B \mid x \text{ integral over } A\}$ *forms a sub* A*-algebra of* B.

Proof. if x and y are integral over A, then $A[x, y]$ is a finitely generated A-module which contains $x + y$ and xy. □

Definition. The *integral closure* of A in B is the subring of B consisting of all elements integral over A. We say that B is *integral* over A if the integral closure of A in B is equal to B. We say that $A \subset B$ is *integrally closed* in B if the integral closure of A in B is equal to A.

Proposition 10.4. *If* B *is integral over* A *and* C *is integral over* B*, then* C *is integral over* A.

Proof. Suppose $c \in C$ is the root of a monic polynomial $p(x) = x^n + b_{n-1} + ... \in B[x]$. Then let $B' = A[b_{n-1}, ...b_0]$ be the subalgebra of B generated by the coefficients of $p(x)$. We have B' is a finite A-alegbra. Then c is contained in $B'[c]$. Now $B'[c]$ is finite over B', and B' is finite over A, so by transitivity of finiteness (exercise!) $B'[c]$ is a finite A-algebra. So c, being contained in the finite A-algebra $B'[c]$ is integral over A as desired. □

Corollary 10.5. *If* $A \subset B$ *are rings, and* C *is the integral closure of* A *in* B*, then* C *is integrally closed in* B.

Proof. Let C' be the integral closure of C in B. We need to show $C = C'$. We know $C \subset C'$, so just need C' subset C. However, C' is integral over C, and C is integral over A, so C' is integral over A. Hence C' is contained in the integral closure C of A in B as desired. □

Proposition 10.6. *For any* $I \subset B$, $A/(I \cap A)$ *is integral over* B/I.

If B *is integral over* A*, for any* $S \subset A$ *multiplicatively closed,* $S^{-1}B$ *is integral over* $S^{-1}A$.

Proof. Any $\overline{b}$ in B/I is the image of some $b \in B$ under the projection map $B \to B/I$. Choose $p(x) \in A[x]$ monic such that $p(b) = 0$. Let $\overline{p}(x)$ be the image of $p(x)$ in $A/(I \cap A)[x]$. Then $\overline{b}$ satisfies the monic polynomial $\overline{p}$, so is integral.

Take any $b/s \in S^{-1}B$, and choose $p(x) \in A[x]$ monic of degree n so that $p(b) = 0$. Then b/s satisfies the monic polynomial $s^{-n}p(sx) = 0$.

□

Proposition 10.7. *Suppose that* $A \subset B$ *are integral domains and that* B *is integral over* A*. Then* B *is a field if and only if* A *is a field.*

Proof. Proof of $\Rightarrow$: Suppose B is a field. Then for any nonzero $a \in A$, there exists $b \in B$ such that $ab = 1$. We must now show that $b \in A$ also. Take a monic polynomial $p(x) = x^n + a_{n-1}x^{n-1} + \cdots + a_0$ such that $p(b) = 0$. Then

$$0 = a^{n-1}p(b) = b + a_{n-1} + a_{n-2}a + \cdots + a_0a^{n-1}.$$

Since all terms in the sum other than b belong to A, we must also have $b \in A$.

Proof of $\Leftarrow$: Suppose A is a field. For any nonzero $b \in B$, take a monic polynomial $p(x)$ such that $p(b) = 0$ Write $p(x) = x^i p'(x)$ where $p'(x) = x^n + a_{n-1}x^{n-1} + \cdots + a_0$ has nonzero constant term $a_0 \neq 0$. Since A is an integral domain and b is nonzero, $p(b) = b^i p'(b) = 0$ implies $p'(b) = 0$. Then

$$1 = a_0^{-1}a_0 = a_0^{-1}(b^{n-1} + a_{n_1}b^{n-2} + \cdots + a_1)b,$$

so b has a multiplicative inverse in B.

□

Corollary 10.8. *Suppose that $A \subset B$ are rings, and that B is integral over A. Then a prime ideal $\mathfrak{p}$ is maximal in B if and only if the prime ideal $\mathfrak{p} \cap A = i^{-1}(\mathfrak{p})$ is maximal in A.*

Proof. We know from before that $B/\mathfrak{p}$ is integral over $A/\mathfrak{p} \cap A$, so $\mathfrak{p}$ is maximal in B if and only if $B/\mathfrak{p}$ is a field if and only if $A/\mathfrak{p} \cap A$ is a field if and only if $\mathfrak{p} \cap A$ is maximal in A. □

Definition. Suppose that $A \subset B$ are rings. We say that a prime ideal $\mathfrak{q}$ of B *lies over* a prime ideal $\mathfrak{p}$ of A if $\mathfrak{q} = i^{-1}(\mathfrak{p}) = \mathfrak{p} \cap A$.

Example. $A = \mathbb{Z}$, $B = \mathbb{Z}[i]$.

In this case, there are three types of primes in $\mathbb{Z}[i]$ (as you may have seen in a previous class; stated here without proof).

The prime (0) of $\mathbb{Z}[i]$ lies above the prime (0) of $\mathbb{Z}$.

For $p \equiv -1 \bmod 4$, the ideal (p) of $\mathbb{Z}[i]$ is prime, and lies above the prime ideal (p) of $\mathbb{Z}$.

For $p \equiv 1 \bmod 4$, there is a unique way of writing $p = a^2 + b^2$ for positive integers a, b, and the ideals $(a + bi)$ and $(a - bi)$ are two prime ideals of $\mathbb{Z}[i]$ lying above the prime ideal (p) of $\mathbb{Z}$.

Note that in this case every prime ideal of $\mathbb{Z}$ had a prime ideal of $\mathbb{Z}[i]$ lying above it. We'll show that this is true in general of integral extensions.

Theorem 10.9 (Going-up)**.** *Suppose that $A \subset B$ are rings, and that B is integral over A. For any prime ideal $\mathfrak{p}$ of A, there is a prime ideal $\mathfrak{q}$ of B lying over $\mathfrak{q}$.*

Proof. First let's localize at $\mathfrak{p}$. Now we get a diagram of rings:

$$\begin{array}{ccc} B & \xrightarrow{j_B} & B_{\mathfrak{p}} \\ {\scriptstyle i}\uparrow & & \uparrow{\scriptstyle i_{\mathfrak{p}}} \\ A & \xrightarrow{j_A} & A_{\mathfrak{p}} \end{array}$$

Here you can think of $B_\mathfrak{p}$ as $S^{-1}B$, where $S = A \setminus \mathfrak{p}$. (This is not necessarily a local ring!)

We have that $B_\mathfrak{p}$ is local over $A_\mathfrak{p}$. Now, choose any maximal ideal $\mathfrak{m}$ of $B_\mathfrak{p}$. We claim that the prime $\mathfrak{q} = j_B^{-1}(\mathfrak{m})$ lies over $\mathfrak{q}$.

First, observe that $i_\mathfrak{p}^{-1}(\mathfrak{m})$ is a maximal ideal of $A_\mathfrak{p}$. But $A_\mathfrak{p}$ is a local ring, so we are forced to have $i_\mathfrak{p}^{-1}(\mathfrak{m}) = \mathfrak{p}A_\mathfrak{p}$. Now,

$$\mathfrak{q} \cap A = i^{-1}(j_B^{-1}(\mathfrak{m})) = j_A^{-1}(i_\mathfrak{p}^{-1}(\mathfrak{m})) = j_A^{-1}(\mathfrak{p}A_\mathfrak{p}) = \mathfrak{p}$$

Hence $\mathfrak{q}$ lies above $\mathfrak{p}$ as desired.

□

11 Finitely generated algebras over a field

Let k be a field. Rings that are finitely generated over k show up various places in mathematics. For instance, earlier in this class we've seen the polynomial ring, $k[x_1, \ldots, x_n]$, and certain invariant subrings that we saw were finitely generated.

Note that if A is a k-algebra with finite generating set $a_1, \ldots, a_n$, then there is a surjective homomorphism $\phi : k[x_1, \ldots, x_n] \to A$ that sends $x_1 \mapsto a_i$, and by the First Isomorphism theorem $A \cong k[x_1, \ldots, x_n]/\ker\phi$. Hence finitely generated k-algebras are exactly those that can be written as $k[x_1, \ldots, x_n]/I$ for an ideal $I \subset k[x_1, \ldots, x_n]$, and so the subject of finitely generated k-algebras is closely related to that of ideals in polynomials rings over k.

The big theorem we'll show about finitely generated k-algebras is the following:

Theorem 11.1 (Noether Normalization)**.** *Let k be a field, and let A be a finitely generated k-algebra. Then there is a finitely-generated sub-k-algebra B of A such that $B \cong \mathbb{C}[x_1, \ldots, x_n]$ and such that A is finite over B.*

Proof. The proof we give in class will just cover the case when k is an infinite field. There will be an alternative proof on HW covering fields of arbitrary cardinality.

Pick generators $a_1, \ldots, a_m$ for A over k. We'll induct on m.

Base case: $m = 0$, so $A \cong k$ and we can just take $B = A \cong k$.

Case 1: $k[y_1, \ldots, y_m] \cong A$ by the map sending y_i to a_i. Then take $B = A$ and we're again done.

Case 2: There exists some $p(y_1, \ldots, y_m) \in k[y_1, \ldots, y_m]$ such that $p(a_1, \ldots, a_m) = 0$. Suppose that p has degree d.

Wishful thinking: Note that if p had a nonzero coefficient on y_1^d, we could then say that a_1 is integral over $A' = k[a_2, \ldots, a_m] \subset A$, and so that $A = A'[a]$ is finite over A' – then we could apply the induction hypothesis to A'. This doesn't have to be the case in general, but we can make it so by clever choice of variables.

For $i = 2, \ldots, m$, define $b_i = a_i - \lambda_i a_1$ for parameters $\lambda_2, \ldots, \lambda_m \in k$ to be defined later. Note that now $a_1, b_2, \ldots, b_m$ is a new generating set for A.

We claim that we can choose the λ_i so that a_1 is integral over $A[b_2, \ldots, b_m]$.

To show this, observe that the polynomial $p_{\lambda_2,\ldots,\lambda_m}(x) = p(x, b_2 + \lambda_n a_1, \ldots, b_m + \lambda_m x)$ has a as a root. If we expand the RHS here as a polynomial of in a_1, we will get a polynomial of degree d, and each monomial $c y_1^{k_1} \cdots y_m^{k_m}$ with $k_1 + k_2 + \cdots + k_m = d$ will contribute a coefficient of $c\lambda_2^{k_2} \cdots \lambda_m^{k_m}(x^d)$ (along with other terms of degree $< d$ in x). Since none of these terms can cancel out, the leading coefficient of x^d is a nonzero polynomial in $\lambda_2, \ldots, \lambda_n$, and since k is infinite, we can choose $\lambda_2, \ldots, \lambda_n \in k$ so that the coefficient of x^d is a nonzero element of k.

After making some such choice of the λ_i, we can divide out by the leading coefficient of $p_{\lambda_2,\ldots,\lambda_m}$ to obtain a monic polynomial satisfied by a_1 with coefficients in the ring $A' = k[b_2, \ldots, b_n]$. So $A = A'[a_1]$ is generated over A' by a single element, a_1, that is intetgral over A', hence is finite over A'.

By induction A' contains a subring $B \cong k[x_1, \ldots, x_n]$ such that A' is finite over B. By transitivity of finiteness, A is also finite over B, and so the induction goes through $\square$

Corollary 11.2. *If a finitely generated* k*-algebra* A *is also a field, then* A *is a finite* k*-algebra (that is,* A *is a finite-dimensional* k*-vector space, what in field theory is known as a "finite extension" of* k*). If* k *is algebraically closed, then in fact* $A \cong k$*.*

Proof. By the Noether Normalization theorem, A has a subring $B \cong k[x_1, \ldots, x_n]$ for some n such that A is a finite B-algebra. Since A is a field and A is integral over B, by a fact proved in the last class (Proposition 2.1 in the course notes for Sept. 25), B is also a field. But the polynomial ring $k[x_1, \ldots, x_n]$ is clearly not a field for $n \geq 1$, so we must have $n = 0$ and $B \cong k$.

Hence A is finite over $k \cong B$ as desired.

Now suppose that k is algebraically closed. We must show that for any $a \in A$ in fact $a \in k$. But a is integral over k, so we have $p(a) = 0$ for some monic $p \in k[x]$. Because k is algebraicaly closed, $p(x)$ factors as $(x - c_1)(x - c_2) \cdots (x - c_n)$ for $c_1, c_2 \ldots, c_n \in k$. Since A is a field, $0 = (a - c_1)(a - c_2) \cdots (a - c_n)$ implies that $a - c_i = 0$ for some i, so $a \in k$. $\square$

Theorem 11.3 (Nullstellensatz, weak form)**.** *Let* k *be an algebraically closed field The maximal ideals* m *of* $k[x_1, \ldots, x_n]$ *are all of the form* $(x_1 - a_1, \ldots, x_n - a_n)$*.*

Proof. Suppose m is a maximal ideal. Then $k[x_1, \ldots, x_n]/m$ is a finitely generated k-algebra, and a field, so $k[x_1, \ldots, x_n]/m \cong k$. Let a_i be the image of x_i in k under the projection map $\pi : k[x_1, \ldots x_n] \to k[x_1, \ldots, x_n]/m \cong k$.

Then $m = \ker \pi$ is equal to the ideal of all polynomials $p(x_1, \ldots, x_n) \in k[x_1, \ldots, x_n]$ such that $p(a_1, \ldots, a_n) = 0$. It's easily checked that this ideal is generated by the elements $x_1 - a_1, \ldots, x_n - a_n$. $\square$

Last time we showed the following weak form of the Nullstellensatz:

Theorem 11.4 (Weak Nullstellensatz)**.** *If a finitely generated* k*-algebra* A *is also a field, then* A *is a finite* k*-algebra (that is,* A *is a finite-dimensional* k*-vector space, what in field theory is known as a "finite extension" of* k*). If* k *is algebraically closed, then in fact* $A \cong k$*.*

Let k *be an algebraically closed field The maximal ideals* $\mathfrak{m}$ *of* $k[x_1, \ldots, x_n]$ *are all of the form*

$$\mathfrak{m}_{(a_1,\ldots,a_n)} = \{p(x_1, \ldots, x_n) \mid p(a_1, \ldots, a_n) = 0\} = (x_1 - a_1, \ldots, x_n - a_n)$$

for $(a_1, \ldots, a_n) \in k^n$.

The second form here has the following corollary:

Corollary 11.5. *Suppose that* k *is algebraically closed, and* $f_1, \ldots, f_m$ *are polynomials in* $k[x_1, \ldots, x_n]$ *with no common zero in* k^n*. Then* $f_1, \ldots, f_m$ *generate the unit ideal* (1) *of* $k[x_1, \ldots, x_n]$*.*

Proof. We prove the contrapositive: if $\langle f_1, \ldots, f_m \rangle$ is a proper ideal, there exists some maximal ideal $\mathfrak{m} = (x_1 - a_1, \ldots, x_n - a_n)$ containing all the f_i. Then all f_i vanish at the point $(a_1, \ldots, a_n) \in k^n$. □

We'll now move on to the strong form of the Nullstellensatz.

Definition. The radical $\sqrt{I}$ of an ideal I is the intersection

$$\bigcap_{\substack{\mathfrak{p} \supset I \\ \mathfrak{p} \text{ prime}}} \mathfrak{p}.$$

We say that I is radical if $I = \sqrt{I}$.

Note that $\sqrt{(0)} = \mathrm{nil}(A)$. More generally, if I is any ideal of A, and π is the projection map $A \to A/I$, we have $\sqrt{I} = \pi^{-1}(\mathrm{nil}(A/I))$. This means that also

$$\sqrt{I} = \{a \in A \mid a^n \in I \text{ for some } n\}$$

.

Theorem 11.6 (Nullstellensatz, strong form)**.** *For any ideal* I *of* $k[x_1, \ldots, x_n]$*, we have*

$$\sqrt{I} = \bigcap_{\substack{\mathfrak{m} \supset I \\ \mathfrak{m} \textit{ maximal}}} \mathfrak{m}. \tag{1}$$

Note: we say that a ring A *is* Jacobson *if* (1) *holds for all ideals* I *of* A*. Hence the above theorem says that* $k[x_1, \ldots, x_n]$ *is Jacobson.*

Proof. The containment

$$\bigcap_{\substack{\mathfrak{m}\supset I \\ \mathfrak{m}\text{ maximal}}} \mathfrak{m} \subset \bigcap_{\substack{\mathfrak{p}\supset I \\ \mathfrak{p}\text{ prime}}} \mathfrak{p} = \sqrt{I}$$

holds for ideals I in an arbitrary ring, because maximal ideals are prime.

For the other direction, suppose that $f \notin \sqrt{I}$. We must show that f is not in the intersection of all maximal ideals containing I: that is, there is some maximal ideal of $k[x_1, \ldots, x_n]$ containing I but not containing f. Since maximal ideals of $k[x_1, \ldots, x_n]$ containing I correspond to maximal ideals of $k[x_1, \ldots, x_n]/I$, it is equivalent to show that there is a maximal ideal of the ring $A = k[x_1, \ldots, x_n]/I$ that does not contain the image $\overline{f}$ of f in A.

By assumption, $\overline{f} \notin \text{nil}(A)$, so $\overline{f}$ is not nilpotent, and $A[\overline{f}^{-1}]$ is not the zero ring. Choose a maximal ideal $\mathfrak{m}$ of $A[\overline{f}^{-1}]$. We have that $A[\overline{f}^{-1}]/\mathfrak{m}$ is a finite k-algebra.

Let j be the natural map $A \to A[\overline{f}^{-1}]$. Since $\mathfrak{m}$ is a prime ideal of $A[\overline{f}^{-1}]$, its pullback $j^{-1}(\mathfrak{m})$ is a prime ideal of Anot containing f. We claim that also $j^{-1}(\mathfrak{m})$ is maximal. For this, note that $k \subset A/j^{-1}(\mathfrak{m}) \subset A[f^{-1}]/\mathfrak{m}$. Now $A[f^{-1}]/\mathfrak{m}$ is both a finitely generated k-algebra and a field, so by the Weak Nullstellensatz, $A[f^{-1}]$ is a finite over k, that is, it is a finite-dimensional k-vector space. Hence its sub-k-algebra $A/j^{-1}(\mathfrak{m})$ is also finite over k, and so is integral over k. Since k is a field, and $A/j^{-1}(\mathfrak{m})$ is an integral domain, $A/j^{-1}(\mathfrak{m})$ must also be a field as desired. □

We'll now give a more geometric form of this, that works for algebraically closed fields.

Recall, from last time, that for any set $X \subset k^n$, we can define the ideal of X by

$$I(X) = \{p(x_1, \ldots, x_n) \in k[x_1, \ldots, x_n] \mid p(a_1, \ldots, a_n) = 0 \text{ for all } (a_1, \ldots, a_n) \in X\}$$

.

We can also define a map in the other direction: For any ideal I of $p(x_1, \ldots, x_n)$, define the *variety* or *vanishing set* of I by

$$V(I) = \{(a_1, \ldots, a_n) \in k^n \mid p(a_1, \ldots, a_n) = 0 \mathrm{forall} p(x_1, \ldots, x_n) \in I\}.$$

We say that a subset V of k^n is a *variety* (some books would use the term *algebraic set* instead) if $V = V(I)$ for some ideal I of $k[x_1, \ldots, x_n]$.

Corollary 11.7. *If k is algebraically closed, then for any ideal I of $k[x_1, \ldots, x_n]$, $I(V(I)) = \sqrt{I}$.*

Proof. We have

$$\begin{aligned}
\sqrt{I} &= \bigcap_{\substack{\mathfrak{m} \supset I \\ \mathfrak{m} \text{ maximal}}} \mathfrak{m} \\
&= \bigcap_{\substack{(a_1,\ldots,a_n) \in k^n \\ \mathfrak{m}_{(a_1,\ldots,a_n)} \supset I}} \mathfrak{m}_{a_1,\text{ dotsc},a_n} \\
&= \bigcap_{(a_1,\ldots,a_n) \in V(I)} \mathfrak{m}_{(a_1,\ldots,a_n)} \\
&= \{p(x_1,\ldots,x_n) \in k[x_1,\ldots,x_n] \mid p(a_1,\ldots,a_n) = 0 \text{ for all } (a_1,\ldots,a_n) \in V(I)\} \\
&= I(V(I)).
\end{aligned}$$

(Here the first equality is the previous form of the strong Nullstellensatz, the second line uses the weak Nullstellensatz to parametrize maximal ideals, and the rest is just formal manipulation.) □

Corollary 11.8. *There is a bijective, order-reversing correspondence between the set of radical ideals* I *of* $k[x_1,\ldots,x_n]$ *and the set of varieties* $V \subset k^n$*. This correspondence is given by the maps* $I \mapsto V(I)$ *and* $V \mapsto I(V)$*.*

Proof. By the Nullstellensatz, if I is a radical ideal, then $I(V(I)) = I$. This implies that the map $I \mapsto V(I)$ is injective.

To show that the map $I \mapsto V(I)$ from radical ideals to subvarieties of k^n is surjective: note that by definition, any variety $V \subset k^n$ is of the form $V(J)$ for some ideal J, not necessarily radical. *This argument here had a gap when I presented it in class – I didn't consider the possibility that* J *might not be radical.* However, we have $V(J) = V(\sqrt{J})$, since any element $f \in \sqrt{J}$ has $f^n \in J$ for some n, and f vanishes at exactly the same points as f^n.

Hence the map $I \mapsto V(I)$ is a bijection from radical ideals to subvarieties of k^n, and by the Nullstellensatz its inverse map is $V \mapsto I(V)$.

Finally it follows directly from the definitions that this map is order-reversing: $V_1 \subset V_2$ implies $I(V_1) \supset I(V_2)$ and $I_1 \subset I_2$ implies $V(I_1) \supset V(I_2)$. □

12 Motivating Primary Ideal Decomposition

We're now going to introduce primary ideal decomposition, which is motivated by two things.

The first is unique factorization in the integers. Because $\mathbb{Z}$ is a unique factorization domain, every $n \in \mathbb{Z}$ can be written as a product of powers of primes $n = p_1 p_2 \cdots p_k$, where the p_i are distinct primes, and the factorization is unique up to reordering and multiplication by units.

Since every ideal of $\mathbb{Z}$ is principal, we can also make this a statement about ideals. For every $I \subset \mathbb{Z}$,

$$I = (p_1) \cdots (p_k)$$

for prime ideals $(p_1) \ldots, (p_m)$ and this factorization is now just unique up to reordering.

We can also combine repetitions, and write

$$I = (p_1)^{a_1} \cdots (p_k)^{a_k}$$

Note that in this case we can also write

$$I = (p_1)^{a_1} \cap \cdots \cap (p_k)^{a_k}.$$

This version will generalize better.

The second motivation is related to the ideal - variety correspondence we established last time.

Definition. A variety $V \subset k^n$ is irreducible if for any two varieties V_1 and V_2 with $V_1 \cup V_2 = V$, either $V_1 = V$ or $V_2 = V$.

Example. In k^2, $V((x_1x_2)) = \{(a_1, a_2) \mid a_1a_2 = 0\}$ is not irreducible, because it can be written as the union of $V_1 = V((x_1))$ and $V_2 = V((x_2))$. But it follows from the proposition below that V_1 and V_2 are both irreducible.

Proposition 12.1. *Under the ideal-variety correspondence, irreducible varieties correspond to prime ideals; that is, a variety V is irreducible if and only if $I(V)$ is prime.*

To prove this we'll use some facts about the ideal-variety correspondence and a lemma about ideals. The facts about the ideal-variety correspondence that we'll need are:

$$V(I_1) \cup V(I_2) = V(I_1 \cap I_2) = V(I_1I_2)$$

(in general $I_1 \cap I_2$ is not generally equal to I_1I_2, but they do have the same radical), and also

$$V(I_1) \cap V(I_2) = V(I_1) + V(I_2).$$

We leave the proofs of these as exercises.

The lemma about ideals is:

Lemma 12.2. *If $\mathfrak{p}$ is a prime ideal of any ring A and $\mathfrak{p} \supset I_1 \cap I_2$, then $\mathfrak{p} \supset I_1$ or $\mathfrak{p} \supset I_2$.*

Proof. We'll prove the contrapositive. Suppose there exist $a_1 \in I_1$ and $a_2 \in I_2$ such that neither a_1 nor a_2 is in $\mathfrak{p}$. Since $\mathfrak{p}$ is prime this means $a_1a_2 \notin \mathfrak{p}$, but $a_1a_2 \in I_1 \cap I_2$, so $\mathfrak{p}$ doesn't contain contain $I_1 \cap I_2$ either. □

Now we can show that irreducible varieties correspond to prime ideals:

Proof. Suppose that $I(V)$ is not prime, so there exists $f, g \notin I(V)$ such that $fg \in I(V)$. Then $V((f)) \not\supset V$ and $V((g)) \not\supset V$ but $V((f)) \cup V((g)) = V((fg)) \supset V$. Now let $V_1 = V((f)) \cap V$ and $V_2 = V((g)) \cap V$; both are proper subsets of V whose union is all of V.

On the other hand, if $I(V)$ is prime, suppose that V_1 and V_2 are varieties whose union is V. Write $V_1 = V(I_1)$ and $V_2 = V(I_2)$ for radical ideals I_1 and I_2. Then $V = V(\mathfrak{p}) = V(I_1 \cap I_2)$; since $\mathfrak{p}$ and $I_1 \cap I_2$ are both radical, we have $\mathfrak{p} = I_1 \cap I_2$. By the previous lemma, we have that either $I_1 \subset \mathfrak{p}$ or $I_2 \subset \mathfrak{p}$. Since $\mathfrak{p} = I_1 \cap I_2$ is a subset of both I_1 and I_2 we must have $\mathfrak{p} = I_1$ or $\mathfrak{p} = I_2$ as desired. □

We'll show later that every variety V can be decomposed as a finite union of irreducible varieties

$$V = V_1 \cup V_2 \cup \cdots \cup V_k.$$

This is equivalent to saying that the ideal $I(V)$ can be written as a finite intersection of prime ideals

$$I(V) = I(V_1) \cap I(V_2) \cap \cdots I(V_k)$$

Primary ideal decomposition will generalize both of the things above.

13 Primary Ideals

Definition. An ideal $\mathfrak{q}$ of A is said to be primary if $ab \in \mathfrak{q}$ implies $a \in \mathfrak{q}$ or $b^n \in \mathfrak{q}$ for some n.

(Equivalently, $ab \in \mathfrak{q}$ implies $a \in \mathfrak{q}$ or $b \in \sqrt{\mathfrak{q}}$.)

As with other ideal properties (prime, maximal, radical), this one can be stated as a property of the quotient ring. An ideal $\mathfrak{q}$ is primary if and only if $A/\mathfrak{q}$ has the property that any zero divisor is nilpotent.

Proposition 13.1. *If $\mathfrak{q}$ is primary then $\sqrt{\mathfrak{q}}$ is the smallest prime ideal containing $\mathfrak{q}$.*

Proof. First we show that $\sqrt{\mathfrak{q}}$ is prime. For this, note that if $ab \in \sqrt{\mathfrak{q}}$, then $(ab)^n \in \mathfrak{q}$ for some n, so either $a^n \in \mathfrak{q}$ or $(b^n)^m \in \mathfrak{q}$ for some m. In the first case $a \in \sqrt{\mathfrak{q}}$ and in the second $b \in \sqrt{\mathfrak{q}}$.

Now, by definition, $\sqrt{\mathfrak{q}}$ is the intersection of all prime ideals containing $\mathfrak{q}$. Since $\sqrt{\mathfrak{q}}$ itself prime, this means that it is contained in any other prime ideal containing $\mathfrak{q}$, so is the smallest such. □

If $\mathfrak{q}$ is primary with radical $\mathfrak{p}$, we also say that $\mathfrak{q}$ is $\mathfrak{p}$-primary.

The converse of this proposition is false in general: for instance, if $A = k[x, y, z]/(xy - z^2)$, and $\mathfrak{p} = (\overline{x}, \overline{z})$, $\mathfrak{p}$ is prime, but $\mathfrak{p}^2$ is not primary, because $xy \in \mathfrak{p}^2$, but $x \notin \mathfrak{p}^2$ and $y \notin \sqrt{\mathfrak{p}^2} = \mathfrak{p}$.

However, we do have at least a partial converse

Proposition 13.2. *If $\sqrt{\mathfrak{q}}$ is maximal then $\mathfrak{q}$ is primary.*

Proof. Let $\mathfrak{m} = \sqrt{\mathfrak{q}}$. Consider the quotient $A/\mathfrak{q}$. The image of $\mathfrak{m} = \sqrt{\mathfrak{q}}$ in $A/\mathfrak{q}$ is equal to the nilradical $\mathrm{nil}(A/\mathfrak{q})$. This means that all prime ideals of $A/\mathfrak{q}$ contain $\mathfrak{m}/\mathfrak{q}$. Since $\mathfrak{m}/\mathfrak{q}$ is maximal, it must be the only prime ideal of $A/\mathfrak{q}$, so also the only maximal ideal. This means that $A/\mathfrak{q}$ is local with unique maximal ideal $\mathfrak{m}/\mathfrak{q} = \mathrm{nil}(A/\mathfrak{q})$, hence any non-nilpotent element of $A/\mathfrak{q}$ is a unit. Since zero-divisors are never units, this meanns that any zero-divisor is nilpotent, as desired. □

Note on definition of primary ideal from last time: we should require that primary ideals be proper (that is, not the whole ring), just as we've done for prime ideals.

Proposition 13.3. *If $\mathfrak{q}_1, \ldots, \mathfrak{q}_n$ are $\mathfrak{p}$-primary, then $\bigcap_{k=1}^n \mathfrak{q}_i$ is $\mathfrak{p}$-primary.*

Proof. First, $\sqrt{\bigcap_{k=1}^n \mathfrak{q}_i} = \bigcap_{k=1}^n \sqrt{\mathfrak{q}_i} = \mathfrak{p}$.

Now suppose that $xy \in \bigcap_{k=1}^n \mathfrak{q}_i$, $x \notin \mathfrak{p}$, then y must be in all the $\mathfrak{q}_i$, so also in their intersection. □

Definition. A primary ideal decomposition of an ideal I is an expression of I as a finite intersection of primary ideals

$$I = \bigcap_{k=1}^{n} \mathfrak{q}_k.$$

It is said to be minimal if the $\mathfrak{q}_k$ have distinct radicals, and if none of the $\mathfrak{q}_k$ contain any of the others.

Correction: this should be: no $\mathfrak{q}_k$ contains $\bigcap_{j\neq k} \mathfrak{q}_j$.

(Note that any primary ideal decomposition can be made minimal by replacing any primary ideals having the same radical with their intersection, and then throwing away any redundant elements.)

Example. Here's an example of minimal primary ideal decomposition that shows that it is not unique. Let $A = k[x, y]$ and let $I = (x^2, xy)$. Then $I = (x) \cap (x^2, xy, y^2)$, or $I = (x) \cap (x^2, y)$, or $I = (x) \cap (x^2, x + y)$ or $I = (x) \cap (x^2, xy, y^n)$.

We'll now show that, in a Noetherian ring, any ideal has a primary decomposition.

Definition. An ideal I of a ring A is irreducible if for any ideals I_1, I_2 with $I_1 \cap I_2 = I$ we have either $I_1 = I$ or $I_2 = I$.

Proposition 13.4. *Let A be a Noetherian ring. Any ideal I of A is an intersection of finitely many irreducible ideals.*

Proof. Suppose A is a Noetherian ring for which the above is not true. Then the set of ideals of A that are finite intersections of irreducible ideals must have a maximal element I. But I is not irreducible, so write $I = I_1 \cap I_2$. Then both I_1 and I_2 are finite intersections of irreducible ideals, so their intersection is also. □

Proposition 13.5. *If* I *is irreducible then* I *is primary.*

First noting a definition we'll need here and later:

Definition. If M is an A-module and $m \in M$, $\mathrm{Ann}_A(m) = \{a \in A \mid am = 0\}$. Observe that $\mathrm{Ann}_A(M)$ is an ideal of A.

Also, if $b \in A$, $\mathrm{Ann}_A(b) \neq 0$ if and only if b is a zero-divisor.

We'll drop the subscript when it's clear what ring we mean.

Proof. Both the hypothesis and conclusion are true for the ideal $I \subset A$ if and only if they are true for the ideal $0 \subset A/I$. Therefore, without loss of generality, we may assume that $I = 0$.

So assume 0 is irreducible. We must show than any zero-divisor in A is nilpotent. Suppose that $x \in A$ is an arbitrary element. We'll show that either x is a non-zero-divisor or x is nilpotent.

Consider the ascending chain of ideals $\mathrm{Ann}(x) \subset \mathrm{Ann}(x^2) \subset$ (where here annihilators are in the A-module A.)

This chain must terminate, so there exists n such that $\mathrm{Ann}(x^n) = \mathrm{Ann}(x^{n+1})$.

We claim that $\mathrm{Ann}(x) \cap (x^n) = 0$. Indeed, suppose that $b = ax^n \in (x^n)$ annihilates x, so $ax^{n+1} = 0$. Then $a \in \mathrm{Ann}(x^{n+1}) = \mathrm{Ann}(x^n)$ so $ax^n = b = 0$.

Now we use the fact that 0 is irreducible. Since $\mathrm{Ann}(x) \cap (x^n)$ is 0, we must either have $\mathrm{Ann}(x) = 0$, and x is not a zero-divisor, or $(x^n) = 0$ and x is nilpotent. □

Example. An example of a non-irreducible primary ideal is $(x^2, xy, y^2) \subset k[x, y]$, which can be written as $(x, y^2) \cap (x^2, y)$.

We'll now show that if $I = \cap_k q_k$ is a primary decomposition, the set $\{\sqrt{q_k}\}$ of primes occuring as radicals depends only on I. To do this we give another characterization of this set.

Definition. A prime $\mathfrak{p}$ of A is associated to an A-module M if $\mathfrak{p} = \sqrt{\mathrm{Ann}_A(m)}$ for some $m \in M$. If I is an ideal of A, by abuse of notation we say that $\mathfrak{p}$ is associated to I (or is an associated prime of I) if $\mathfrak{p}$ is associated to the A-module A/I.

I screwed up with the notation below: the standard notation for this is $(I : a)$*, but I wrote this as* $(a : I)$ *on Monday. It's now fixed here.*

Notation for this case: define$(I : a) = \{b \in A \mid ab \in I\}$. Then if $\bar{a}$ denotes the image of $a \in A/I$ we have $\mathrm{Ann}_A(\bar{a}) = (I : a)$

Let $\mathrm{Ass}_A(M)$ (or $\mathrm{Ass}_A(I)$) denote the set of associated primes of M (or of I).

We'll eventually show:

Theorem 13.6. *If* $I = \bigcap_i q_i$ *is a minimal primary decomposition of* I*, then the set of associated primes of* I *is precisely the set* $\sqrt{q_i}$

First let's see how this works out in the case where $I = q$ is primary.

Proposition 13.7. *Let $\mathfrak{q}$ be a $\mathfrak{p}$-primary ideal.*

Let a be an arbitrary element of A. Then

a) If $a \in \mathfrak{q}$, then $(\mathfrak{q} : a) = A$.

b) If $a \notin \mathfrak{p}$ then $(\mathfrak{q} : a) = \mathfrak{q}$.

c) If $a \notin \mathfrak{q}$, then $(\mathfrak{q} : a)$ is $\mathfrak{p}$-primary.

Proof. Part a) is clear, and b) is just the statement that $\mathfrak{q}$ is $\mathfrak{p}$-primary.

For c), First, we have $\mathfrak{q} \subset (\mathfrak{q} : a) \subset \mathfrak{p}$, so $\sqrt{(\mathfrak{q} : a)} = \mathfrak{p}$. Now we show that $(\mathfrak{q} : a)$ is $\mathfrak{p}$-primary. Suppose $bc \in \mathrm{Ann}_A(\bar{a}$ and $b \notin \mathfrak{p}$. Then $abc \in \mathfrak{q}$, so since $\mathfrak{q}$ is $\mathfrak{p}$-primary and $b \notin \mathfrak{p}$ we must have $ac \in \mathfrak{p}$ and $c \in \mathrm{Ann}_A(\bar{a})$ as desired. □

We started with corrections to the definitions given last time.

First of all, the ideal $\{b \in A \mid ab \in I\}$ should be denoted $(I : a)$ rather than the $(a : I)$ notation last time. Motivation for this is that this is like "dividing I by a", so the I should come first.

Also, the definition of minimal primary decomposition I gave was too weak. It should be fixed to

Definition. A primary decomposition $I = \bigcap_k \mathfrak{q}_k$ is said to be *minimal* if the $\mathfrak{q}_k$ have distinct radicals and no $\mathfrak{q}_k$ contains $\bigcap_{j \neq k} \mathfrak{q}_j$.

(Note that any $\mathfrak{q}_k$ contained $\bigcap_{j \neq k} \mathfrak{q}_j$, we could remove it from the decomposition and still have the same intersection. For instance, in $k[x, y]$ the primary decomposition $(xy) = (x) \cap (y) \cap (x^2, xy, y^2)$ is not minimal because $(x^2, xy, y^2) \supset (x) \cap (y)$.)

Last time we stated the following theorem, which we will now prove:

Theorem 13.8 (First Uniqueness)**.** *If $I = \bigcap_i \mathfrak{q}_i$ is a minimal primary decomposition of I, the set $\{\sqrt{\mathfrak{q}_i}\}$ is precisely the set $\mathrm{Ass}(I)$ of associated primes of I. In particular, this means that that the set of radicals $\sqrt{\mathfrak{q}_i}$ does not depend on the choice of primary decomposition.*

Proof. Suppose $I = \bigcap_k \mathfrak{q}_k$, and this is minimal. Then, for any $a \in A$,

$$\sqrt{(I : a)} = \sqrt{\bigcap_k (\mathfrak{q}_k : a)} = \bigcap_k \sqrt{(\mathfrak{q}_k : a)}.$$

If $\sqrt{(I : a)}$ is prime, then it must be equal to $\sqrt{(\mathfrak{q}_k : a)}$ for some k. But $(\mathfrak{q}_k : a)$ is either 1 or $\sqrt{\mathfrak{q}_k}$-primary, so we must have $\sqrt{(I : a)} = \sqrt{\mathfrak{q}_k}$ for some k.

Conversely, for any j, by assumption we have some $a \notin \mathfrak{q}_j$ but $a \in \mathfrak{q}_k$ for any $k \neq j$. Then $\sqrt{(I : a)} = \bigcap_k \sqrt{(\mathfrak{q}_j : a)} = \sqrt{(\mathfrak{q}_k : a)}$ is $\mathfrak{q}_k$-primary. □

Definition. If I is an ideal of A, the minimal elements of the set $\mathrm{Ass}(I)$ are called the *minimal primes* or *isolated primes* associated to I. The others are called the *embedded primes*.

Example. Returning to the example of $A = k[x, y]$, $I = (x^2, xy)$ has as primary decomposition $(x) \cap (x^2, xy, y^2)$, so $\mathrm{Ass}(I) = \{(x), (x, y)\}$. Here (x) is a minimal associated prime and (x, y) is an embedded prime. The geometric intuition here is that the corresponding variety $V((x, y))$ is a point embedded in the line $V((x))$.

Proposition 13.9. *Suppose* I *has a primary ideal decomposition. Then any prime ideal containing* I *also contains a minimal prime associated to* I*. So the minimal primes associated to* I *are precisely the minimal elements of the set* {*prime ideals of* A *containing* I}*, or what we called the minimal primes over* I *in problem set 1.*

Proof. First, note that if $\mathfrak{p}$ is any prime ideal of A such that $\mathfrak{p} \supset I$ then also $\mathfrak{p} = \sqrt{\mathfrak{p}} \supset \sqrt{I} = \bigcap_i \mathfrak{p}_i$. Hence $\mathfrak{p} \supset \mathfrak{p}_k$ for some k, and without loss of generality we can assume that $\mathfrak{p}_k$ is minimal.

To see that the second part of the theorem follows: if $\mathfrak{p}$ is not a minimal associated prime of I, it contains as a proper subset some minimal associated prime of I, so it is not a minimal element of {prime ideals of A containing I}.

On the other hand, if $\mathfrak{p}$ is a minimal associated prime of I, then any other prime ideal $\mathfrak{p}'$ of I with $\mathfrak{p} \supset \mathfrak{p}'$ itself contains a minimal associated prime $\mathfrak{p}''$ of I. So $\mathfrak{p} \supset \mathfrak{p}' \supset \mathfrak{p}''$, but $\mathfrak{p}$ is a minimal associated prime, so $\mathfrak{p} = \mathfrak{p}'' = \mathfrak{p}'$, showing that $\mathfrak{p}$ is itself minimal. □

We'll now show that, in a minimal primary decomposition, the primary ideals whose radicals are the minimal primes are uniquely determined. To do this, we first need to talk about compatibility of primary ideal decomposition with localization.

First we see what happens when we localize primary ideals. Remember we have a map j_* from the set of ideals of A to the set of ideals of $S^{-1}A$. In fact, for $I \subset S$, $j_*(I) = S^{-1}I = \{\frac{i}{s} \mid i \in I, s \in S\}$. One fact we'll use a bit later is that j_* is compatible with taking intersections:

Proposition 13.10. *If* $I_1, \ldots, I_n$ *are ideals of* A*, then* $j_*(\bigcap_k I_k) = \bigcap_k j_*(I_k)$

Proof. By induction, suffices to show this for $n = 2$.

Clearly $j_*(I_1 \cap I_2) \subset j_*(I_1) \cap j_*(I_2)$. On the other hand, if $b \in j_*(I_1) \cap j_*(I_2)$ this means that $b = a_1/s_1 = a_2/s_2$ where $a_1 \in I_1$ and $a_2 \in I_2$. Hence there exists $u \in S$ such that $ua_1s_2 = ua_2s_1 \in I_1 \cap I_2$. Then $b = (ua_1s_2)/(us_1s_2) \in j_*(I_1 \cap I_2)$. □

(Note that this is also true of the map j^* from ideals of $S^{-1}A$ to ideals of A.)

Proposition 13.11. *Let* $S \subset A$ *be a multiplicative subset, and let* $\mathfrak{q}$ *be a primary ideal of* A*.*

a) If $S \cap \mathfrak{p}$ *is nonempty, then* $j_*(\mathfrak{q}) = S^{-1}A$*.*

b) If $S \cap \mathfrak{p} = \emptyset$*, then* $j_*(\mathfrak{q})$ *is* $S^{-1}\mathfrak{p}$*-primary and* $j^*(j_*(\mathfrak{q})) = \mathfrak{q}$

That is, in this case, j_* *induces a bijection between* $\mathfrak{p}$*-primary ideals of* A *and* $S^{-1}\mathfrak{p}$*-primary ideals of* $S^{-1}A$*.*

Proof. HW. □

Proposition 13.12. *Suppose that* $I = \bigcap_k q_k$ *is a minimal primary decomposition, and* S *is a multiplicatively closed subset of* A*. Then*

$$j_*(I) = \bigcap_{\sqrt{q_k} \cap S = \emptyset} j_*(q_k)$$

is a minimal primary decomposition of $j_*(I)$*.*

and

$$j^*(j_*(I)) = \bigcap_{\sqrt{q_k} \cap S = \emptyset} q_k$$

is a minimal primary decomposition of $j^*(j_*(I))$*.*

We'll prove this next time. Right now, let's see how this implies that the primary ideals whose radicals are embedded primes are independent of the choice of primary decomposition.

Theorem 13.13. *Let* $\mathfrak{p}$ *be a minimal associated prime of* I*. Then in any minimal primary decomposition* $\cap_k q_k$ *of* I*, the primary ideal* q_j *with radical equal to* $\mathfrak{p}$ *is determined by* I *and* $\mathfrak{p}$ *and does not depend on the choice of decomposition.*

Proof. Let $S = A - \mathfrak{p}$. Then by the previous proposition,

$$j^*(j_*(I)) = \bigcap_{\sqrt{q_k} \cap S = \emptyset} q_k = q_j$$

and since the left hand side does not depend on $\mathfrak{p}$, neither does the right hand side. □

A bit more terminology: if $I = \bigcap_k q_k$ is a minimal primary decomposition, we say that q_k are the primary components of I, and if $\sqrt{q_k} = \mathfrak{p}_k$ we say that q_k is $\mathfrak{p}_k$-primary. By the First uniqueness theorem from last time we know that every minimal primary decomposition has one primary component for each associated prime. Now we prove the the proposition we stated last time.

Proposition 13.14. *Suppose that* $I = \bigcap_k q_k$ *is a minimal primary decomposition, and* S *is a multiplicatively closed subset of* A*. Then*

$$j_*(I) = \bigcap_{\sqrt{q_k} \cap S = \emptyset} j_*(q_k)$$

is a minimal primary decomposition of $j_*(I)$*.*

and

$$j^*(j_*(I)) = \bigcap_{\sqrt{q_k}\cap S=\emptyset} q_k$$

is a minimal primary decomposition of $j^*(j_*(I))$.

Proof. By previous proposition, we have $j_*(I) = \bigcap_k j_*(q_k)$, and since $j_*(q_k) = S^{-1}A$ whenever q_k has nonempty intersection with S, we can just drop those terms. This shows it's a primary decomposition. Likewise, to get the second equation, we just apply j^* to the whole thing. This is clearly minimal since our original primary decomposition was minimal.

Likewise the only bit that's at all complicated is to show that the first decomposition is minimal. For this, suppose otherwise that $j_*(q_k) \subset \cap_{j\neq k} j_*(q_j)$. Then applying j^* to this, we get $q_k = j^*(j_*(q_k)) \subset \cap_{j\neq k} q_j$, but we know that's okay.

□

Application to irreducible decomposition:

Proposition 13.15. *If* A *is any noetherian ring, and* $I \subset A$ *is a radical ideal, then* I *can be written uniquely as* $\bigcap_k \mathfrak{p}_k$ *where no* $\mathfrak{p}_k$ *contains any* $\mathfrak{p}_j$ *for* $j \neq k$*, and this is the unique minimal primary decomposition of* I.

Proof. Take any minimal primary decomposition $I = \cap_k q_k$, and take radicals to get $I = \cap_k \mathfrak{p}_k$ where $\mathfrak{p}_k = q_k$. This is still a primary decomposition, and it's still minimal, so no $\mathfrak{p}_k$ can contain any other $\mathfrak{p}_j$, hence all the $\mathfrak{p}_k$ are minimal primes. By the second uniqueness theorem, this means that this is the unique minimal primary decomposition.

Furthermore, if we have any expression of $I = \cap_\mathfrak{p} \mathfrak{p}'_k$ such that no $\mathfrak{p}'_k$ contains any other $\mathfrak{p}'_j$, this is also a minimal primary decomposition □

Now we are going to move from taking intersections of ideals to multiplying ideals. To do this we need the following lemma.

Lemma 13.16. *If* $I_1, I_2, \ldots, I_n$ *are ideals of* A *such that* $I_j + I_k = (1)$ *for* $j \neq k$*, then* $I_1 \cap I_2 \cap \cdots \cap I_n = I_1 I_2 \cdots I_n$.

Proof. HW □

We are now going to introduce a property a ring can have that will imply that every ideal is a product of primary ideals.

Definition. If A is an integral domain, we say that A has (Krull) dimension ≤ 1 if every nonzero prime ideal of A is maximal.

(Comment: in general, we say that the Krull dimension of a ring A is the largest integer n such that there is a sequence of prime ideals $\mathfrak{p}_0 \subsetneq \mathfrak{p}_1 \subsetneq \cdots \subsetneq \mathfrak{p}_n$ of A with length $n+1$. In the case above, the longest such chain is $0 \subsetneq \mathfrak{p}$ for any prime ideal of A, assuming that A has any nonzero prime ideals; otherwise A is a field, and has Krull dimension 0.)

Example. $\mathbb{Z}$ has Krull dimension ≤ 1; so $\mathbb{C}[t]$, or $k[t]$ for any field k, in fact we'll see later that any PID has Krull dimension ≤ 1 (in fact, $= 1$ if it's not just a field).

As well, for any irreducible polynomial $f(x,y)$, the ring $\mathbb{C}[x,y]/(f)$ has Krull dimension ≤ 1 (again, actually $= 1$) – this will probably be on the problem set.

Proposition 13.17. *If A is Noetherian of Krull dimension ≤ 1, then every nonzero ideal I of A can be uniquely expressed as a product of primary ideals with distinct radicals.*

Proof. We show first that for nonzero primary ideals $\mathfrak{q}_1, \ldots, \mathfrak{q}_n$ with distinct radicals we have $\cap_k \mathfrak{q}_k = \prod_k \mathfrak{q}_k$. For this, we use the lemma above; we need to show that $\mathfrak{q}_j + \mathfrak{q}_k = (1)$ for $j \neq k$. Now, $\sqrt{\mathfrak{q}_j + \mathfrak{q}_k}$ is an ideal containing both ideals $\sqrt{\mathfrak{q}_j}$ and $\sqrt{\mathfrak{q}_k}$. But $\sqrt{\mathfrak{q}_j}$ and $\sqrt{\mathfrak{q}_k}$ are nonzero prime ideals, so they are both maximal, and therefore any ideal containing both of them is (1). Hence $\sqrt{\mathfrak{q}_j + \mathfrak{q}_k} = (1)$ so $\mathfrak{q}_j + \mathfrak{q}_k = 1$.

Hence it suffices to show that I has a unique primary decomposition. However, the assumption that A has Krull dimension 1 means that none of the associated primes of I can contain any of the others, so they are are minimal, and so the primary decomposition of A is determined by the second uniqueness theorem. □

We now define a class of rings in which every ideal has a unique factorization into powers of primes, not just into primary ideals.

Definition. We say that A is a Dedekind domain if A is Noetherian of dimension ≤ 1 and A is integrally closed in its field of fractions.

Example. Again, the rings $\mathbb{Z}$, $\mathbb{C}[t]$ and $k[t]$ are Dedekind, as are any PID.

A non-PID example is $\mathbb{Z}[\sqrt{-5}]$. More generally, we'll later be able to prove that if K is a finite extension of $\mathbb{Q}$, the integral closure of $\mathbb{Z}$ in K is Dedekind; so for instance all of the integral closures in $\mathbb{Q}[\sqrt{D}]$ you calculated on the problem set (in the case where D is not square, that is.)

Another fact that we'll see later: the ring $C[x,y]/(f)$ (with f irreducible) is a Dedekind domain if and only if the variety $V(f)$ is smooth. (This means roughly what you expect, geometrically; algebraically it means that the polynomials $\frac{\partial f}{\partial x}$ and $\frac{\partial f}{\partial y}$ are never both zero at any point of $V(f)$.)

14 Localizing Dedekind domains

Proposition 14.1. *Suppose that A is a Dedekind domain, and S is any multiplicatively closed subset of A. Then $S^{-1}A$ is also Dedekind.*

Proof. We've previously seen (on HW) that A noetherian implies $S^{-1}A$ noetherian.

Also, by HW 4, since A is the integral closure of A in $K = \operatorname{Frac} A$, $S^{-1}A$ is the integral closure of $S^{-1}A$ in $S^{-1}K = K = \operatorname{Frac}(S^{-1}A)$.

For Krull dimension 1: suppose not. Then we have some non-maximal prime ideal $\mathfrak{p}$ of $S^{-1}A$, which must be contained in some maximal ideal $\mathfrak{m}$ of A. Pulling back to A, we have $j^*(\mathfrak{p}) \subsetneq j^*(\mathfrak{m})$, so the prime ideal $j^*(\mathfrak{p})$ is not maximal either. Since A has Krull dimension ≤ 1, we must have $j^*(\mathfrak{p}) = 0$, so $\mathfrak{p} = j_*(j^*(\mathfrak{p})) = j_*(0) = 0$. □

Proposition 14.2. *Suppose that A is a noetherian domain. Then the following are equivalent*

a) A is Dedekind

b) $A_\mathfrak{p}$ is Dedekind for all prime ideals $\mathfrak{p}$ of A

c) $A_\mathfrak{m}$ is Dedekind for all maximal ideals $\mathfrak{m}$ of A.

Proof. We just showed that a) $\implies$ b), and b) $\implies$ c) is evident. For c) $\implies$ a):

Integral closure in $\operatorname{Frac} A$: this was shown on problem set 4.

Nonzero prime ideals are maximal: this is the opposite argument to what we just gave. Suppose $\mathfrak{p}$ is a non-maximal prime of A. Then there exists $\mathfrak{m}$ maximal containing $\mathfrak{p}$. Localize at $\mathfrak{m}$. Then $j_*(\mathfrak{p}) \subsetneq j_*(\mathfrak{m})$, and the latter is the unique maximal ideal of $j_*(\mathfrak{m})$, so $j_*(\mathfrak{p})$ is not maximal in $A_\mathfrak{m}$. Since $A_\mathfrak{m}$ has Krull dimension 1 this gives us that $j_*(\mathfrak{p}) = 0$ so $\mathfrak{p} = j^*(j_*(\mathfrak{p})) = 0$. □

15 Discrete Valuation Rings

So we can understand general Dedekind domains by understanding their local

Definition. A *valuation* on a field K is a surjective group homomorphism $\nu : K^* \to \mathbb{Z}$ such that $\nu(x+y) \geq \min(\nu(x), \nu(y))$.

If ν is a valuation on K, the subset $\{x \in K \mid \nu(x) \geq 0\}$ is easily seem to be a subring A of K (this follows from definitions and the fact that $\nu(-1) + \nu(-1) = \nu(1) = 0$, so $\nu(-1) = 0$ and $\nu(-x) = \nu(x)$ for all x.). A ring A of the above form is called a *discrete valuation ring* or *DVR* for short.

Example. Define a valuation on $\mathbb{C}(t)$ by $\nu(f)$ is the "order of vanishing of f at 0"; that is: if

$$f = t^i \frac{p}{q}$$

where $p, q \in \mathbb{C}[t]$ such that $p(0) \neq 0$ and $q(0) \neq 0$, then $\nu(f) = i$. Easy to check that this is a valuation, and that the associated DVR is the localization

$$\mathbb{C}[t]_{(t)} = \left\{ \frac{p}{q} \in \mathbb{C}(t) \mid q(0) \neq 0 \right\}.$$

Example. Let p be an integer prime. Define a valuation ν_p on $\mathbb{Q}$ by $\nu_p(x) = i$ if $x = p^i \frac{m}{n}$ where $m, n \in \mathbb{Z}$ are both not divisible by p. Again, check this is a valuation, similar to above, and the associated DVR is the localization $\mathbb{Z}_{(p)}$.

For two elements a, b in a discrete valuation ring A, a divides b if and only if $\nu(a) \leq \nu(b)$. Hence any ideal I of A is generated by any element a of minimal possible valuation, that is, $I = (a) = \{b \in A \mid \nu(b) \geq \nu(a)\}$.

This means that every ideal of A is of the form $I_n = \{a \in A \mid \nu(a) \geq n\}$ for $n = 0, 1, \ldots,$. These ideals can be arranged in a descending chain $I_0 = A \supset I_1 \supset I_2 \cdots$. It's now evident that I_1 is the unique maximal ideal of A. Write $\mathfrak{m} = I_1 = (t)$ for any element $t \in A$ with $\nu(t) = 1$. Then $I_n = (t^n) = \mathfrak{m}^n$.

As a result, we see that A is a PID (hence Noetherian); also that I_1 is the unique nonzero prime ideal of A, hence A has Krull dimension 1. To see that A is also integrally closed in $\mathrm{Frac}(A)$, we use a result from HW.

Recall (from HW) that we say that A is a valuation ring if for any $x \in \mathrm{Frac}(A)$, either $x \in A$ or $x^{-1} \in A$. If A is a DVR coming from a valuation ν on a field K, then $K = \mathrm{Frac}(A)$, and if $x \in K$, either $\nu(x) \geq 0$ and $x \in A$ or $\nu(x^{-1}) = -\nu(x) \geq 0$ and $x^{-1} \in A$. Hence all DVRs are valuation rings (justifying the terminology!) so A is integrally closed in $\mathrm{Frac}(A)$.

The above means that every DVR is Dedekind. The following result will show that every local Dedekind domain is a DVR; and give a number of equivalent conditions for a ring to be Dedekind.

Proposition 15.1. *Let A be an integral domain that is noetherian, local, and of dimension 1 (this means dimension ≤ 1 and not a field; along with the local condition, this means it has precisely two prime ideals , 0 and the unique maximal ideal $\mathfrak{m}$). Let $\mathfrak{m}$ be the unique maximal ideal and $k = A/\mathfrak{m}$ the residue field, and let $K = \mathrm{Frac}(A)$. Then TFAE:*

a) A is a DVR

b) A is integrally closed in $K = \mathrm{Frac}(A)$

c) $\mathfrak{m}$ is principal.

d) $\dim_k \mathfrak{m}/\mathfrak{m}^2 = 1$.

e) every nonzero ideal is a power of $\mathfrak{m}$

f) there is $t \in A$ such that every nonzero ideal of A is of the form (t^n) for some n

Proof. First, some general observations that are always true when A is a noetherian local integral domain of dimension 1. We can make a descending chain of ideals $A = \mathfrak{m}^0 \supset \mathfrak{m} \supset \mathfrak{m}^2 \supset \mathfrak{m}^3 \cdots$.

(If A were a DVR we would know these were all the ideals of A; in general this doesn't have to be the case.)

Correction: I gave the argument below in class, but I don't think I can justify interchanging the product and intersection here. I'll prove this differently in class next time. Note that one can prove with Nakayama's lemma that $\mathfrak{m}^n \neq \mathfrak{m}^{n+1}$: because Nakayama's lemma applied to $\mathfrak{m}^n = \mathfrak{m}\mathfrak{m}^n$ gives $\mathfrak{m}^n = 0$, which it evidently isn't (since $\mathfrak{m} \neq 0$ and A is a domain).

(By Nakayama's lemma for local rings, the intersection $\mathfrak{m}_\infty = \bigcap_{n=0}^{\infty} \mathfrak{m}^n$ of all the ideals in the chain is 0, since

$$\mathfrak{m}\mathfrak{m}_\infty \bigcap_{n=0}^{\infty} \mathfrak{m}\mathfrak{m}^n = \bigcap_{n=1}^{\infty} \mathfrak{m}^n = \mathfrak{m}_\infty.$$

As a corollary, the descending chain never stabilizes, since if we had $\mathfrak{m}^n = \mathfrak{m}^{n+1}$ we'd then have $\mathfrak{m}^n = \mathfrak{m}_\infty = 0$, but $\mathfrak{m}^n \neq 0$ since $\mathfrak{m} \neq 0$ and we're in an integral domain.)

By Nakayama's lemma, $\mathfrak{m}^n \neq 0$ implies $\mathfrak{m}^n \neq \mathfrak{m}^{n+1}$ for all $n \geq 0$.

Also, if we look at the successive quotients, $\mathfrak{m}^n/\mathfrak{m}^{n+1}$, these are all vector spaces over $k = A/\mathfrak{m}$; in fact they can be identified with $\mathfrak{m}^n \otimes_A k$.

Also, any nonzero ideal I of A has $\sqrt{I} = \cap_{\mathfrak{p} \supset I} \mathfrak{p} = \mathfrak{m}$. By the current HW, since A is noetherian this implies that $I \supset \mathfrak{m}^n$ for some $\mathfrak{m}$.

In particular, if the ideal $\mathfrak{m}_\infty = \bigcap_{n \geq 0} \mathfrak{m}^n$ were nonzero, we'd have $\mathfrak{m}_\infty \supset \mathfrak{m}^n \supset \mathfrak{m}^{n+1} \supset \mathfrak{m}_\infty$, contradicting $\mathfrak{m}^n \neq \mathfrak{m}^{n+1}$.

(This fact that $\bigcap_{n \geq 0} \mathfrak{m}^n = 0$ is also true in any Noetherian local ring, and is known as Krull's intersection theorem, but it's harder to prove.)

Preliminary remarks:

a) $\implies$ b): As mentioned last time, this follows from HW, since DVRs are valuation rings.

b) $\implies$ c): This is the trickiest one. First, we choose any $a \in A$. By assumption, $(a) \supset \mathfrak{m}^n$ for some n. Choose this n minimal, so there must exist some $b \in \mathfrak{m}^{n-1}$ but $b \notin (a)$. Now write $t = a/b \in K$; this is going to be our candidate for generator of $\mathfrak{m}$, but right now it's just some element of K.

Consider $x^{-1}\mathfrak{m}$. A priori, this is just some sub A-module of K. However, since $b\mathfrak{m} \subset \mathfrak{m}^n \subset (a)$, in fact $x^{-1}\mathfrak{m} = \frac{b\mathfrak{m}}{a} \subset A$. Since A is local, either $x^{-1}\mathfrak{m} = A$ or $x^{-1}\mathfrak{m} \subset \mathfrak{m}$. Suppose that the latter were the case: then $\mathfrak{m}$ would be a finitely generated A-module with a faithful action of the ring $A[x^{-1}]$; so x^{-1} would be integral over A. But A is integrally closed in K, so this would imply $x^{-1} \in A$, contradicting $b = x^{-1}a \notin (a)$.

Hence we must have the other option : $x^{-1}\mathfrak{m} = A$. Rescaling both sides by x we get $\mathfrak{m} = xA$ is principal.

c) $\implies$ d): If $\mathfrak{m} = (t)$ then $\mathfrak{m} \cong A$ as A-modules. Hence $\mathfrak{m}/\mathfrak{m}^2 = \mathfrak{m} \otimes_A k \cong A \otimes_A k = k$ is a one-dimensional k-vector space.

d) $\implies$ c): Choose t such that the image $\bar{t}$ of t in $\mathfrak{m}/\mathfrak{m}^2$ is a generator of the k-vector space $\mathfrak{m}/\mathfrak{m}^2$. We need to show that in the short exact sequence

$$A \xrightarrow{\times t} \mathfrak{m} \to \mathfrak{m}/(t) \to 0$$

the last term is zero. For this, tensor everything with $A/\mathfrak{m}$, to get

$$A \otimes_{\mathfrak{m}} A/\mathfrak{m} \xrightarrow{\times \bar{t}} \mathfrak{m} \otimes_{\mathfrak{m}} A/\mathfrak{m} \to (\mathfrak{m}/(t)) \otimes_A A/\mathfrak{m} \to 0;$$

this can be rewritten as

$$k \xrightarrow{\times \bar{t}} \mathfrak{m}/\mathfrak{m}^2 \to (\mathfrak{m}/(t)) \otimes_A A/\mathfrak{m} \to 0.$$

Then the first map is an isomorphism, by our choice of t, and so the last term must be 0. But then Nakayama's lemma for local rings implies that $\mathfrak{m}/(t) \cong 0$, so $(t) = \mathfrak{m}$.

c) $\implies$ e): suppose that $\mathfrak{m} = (t)$, and let I be any ideal of A. Choose the largest n such that $I \subset \mathfrak{m}^n = (t^n)$ (one exists since $\bigcap_{n=1}^{\infty} \mathfrak{m}^n = 0$). Then there exists $a \in I$ such that $a \notin \mathfrak{m}^{n+1}$. Write $a = bt^n$. Then $b \notin \mathfrak{m}$ so b is a unit; hence $t^n \in I$ also, and $I = (t^n)$.

e) $\implies$ f): Choose $t \in \mathfrak{m} \setminus \mathfrak{m}^2$. Then $(t) = \mathfrak{m}^n$ for some n; must have $n = 1$, so $\mathfrak{m}$ is principal and the result follows.

f) $\implies$ a): We must have $\mathfrak{m} = (t)$ here, and so the ideals $(x^i) = \mathfrak{m}^i$ are all distinct. Define a valuation function $\nu : A \setminus \{0\} \to \mathbb{Z}$ by letting $\nu(a)$ be the unique integer such that $(a) = (x^{\nu}(a))$; equivalently, $\nu(a)$ is the unique natural integer such that there exists a unit $u_a \in A^{\times}$, then $a = u_a t^{\nu(a)}$. (Note that by construction $\nu(a) \geq 0$ for all $a \in A$.)

It follows immediately from this that $\nu(ab) = \nu(a) + \nu(b)$; hence we can extend to a map homomorphism $\nu : K \to \mathbb{Z}$ by defining $\nu(a/b) = \nu(a) - \nu(b)$; easy to check this is well-defined, a group homomorphism, and that for any $x \in K$, $\nu(x)$ is the unique integer such that there is a unit u_x of A such that $x = u_x t^{\nu(x)}$, where $u_x \in A^{\times}$. Also easy to check that $\nu(x) \geq 0$ if and only if $x \in A$.

Now, we check that $\nu(x+y) \leq \nu(x) + \nu(y)$. For this, assume WLOG that $\nu(x) \leq \nu(y)$. If $x, y \in K$, write $x = u_x t^{\nu(x)}$, $y = u_y t^{\nu(y)}$. Then $x + y = t^{\nu(x)}(u_x + t^{\nu(y)-\nu(x)} u_y) = t^{\nu(x)} a$ for $a \in A$. So $\nu(x+y) = \nu(t^{\nu(x)} a) = \nu(x) + \nu(a) \geq \nu(x) = \min(\nu(x), \nu(y))$.

□

At the end of last time, we showed:

Theorem 15.2. *A noetherian domain A is Dedekind if and only every localization $A_{\mathfrak{p}}$ is a DVR.*

As a corollary, we obtain:

Corollary 15.3. *PIDs are Dedekind.*

Proof. If A is a PID, it's certainly a noetherian domain. So we only need to check that $A_{\mathfrak{p}}$ is a DVR.

First we show that $A_{\mathfrak{p}}$ is also a PID. For this, the map $j_* : \{\text{ideals of } A\} \to \{\text{ideals of } A_{\mathfrak{p}}\}$ is surjective (using the fact $j_*(j^*(I)) = I$ from HW), and sends principal ideals to principal ideals.

Now, $A_{\mathfrak{p}}$ is a local PID; by HW this means it has Krull dimension ≤ 1, and its maximal ideal is principal, hence by last time it is a DVR. □

(Note that we could instead have done this directly, by proving that any PID is a UFD, and then using the result on UFDs to show that our original ring is integrally closed in its field of fractions.)

Lemma 15.4. *If A is Dedekind then every primary ideal of A is a prime power.*

Proof. Let $\mathfrak{q} \subset A$ be primary, $\mathfrak{p} = \sqrt{\mathfrak{q}}$. Then, since $A_\mathfrak{p}$ is a DVR, for some n, $j_*(\mathfrak{q}) = (\mathfrak{p}A_\mathfrak{p})^n = j_*(\mathfrak{p}^n)$. Since j_* is bijective on $\mathfrak{p}$-primary ideals, we must have $\mathfrak{q} = \mathfrak{p}^n$. □

Corollary 15.5. *If A is Dedekind, then every ideal I of A is uniquely a product*

$$I = \mathfrak{p}_1^{a_1}\mathfrak{p}_2^{a_2}\cdots\mathfrak{p}_n^{a_N}$$

of powers of distinct prime ideals.

Proof. We previously saw that in any Noetherian domain of dimension 1, every ideal is uniquely a product of primary ideals with distinct radicals. Now apply the previous lemma. □

Definition. Let A be an integral domain with field of fractions K. A *fractional* ideal of A is an A-submodule I of K such that $xI \subset A$ for some nonzero $x \in A$.

Equivalently, $I \subset x^{-1}A$.

In an arbitrary ring A, any finitely generated A-submodule of K is a fractional ideal (since we can take x to be the product of all denominators of the generators). The converse holds if A is noetherian: then any fractional ideal I is isomorphic to the ideal xI as an A-module, and the latter is finitely generated by assumption.

Example. If A is a DVR with valuation ν, every fractional ideal of A is of the form $I_r = \{x \in K \mid \nu(x) \geq r\}$. The proof of this is the same as the corresponding statement for integral ideals. If $t \in A$ with $\nu(t) = 1$ (such a t is called a "uniformizer", we can also write $I_r = (t^r) = t^rA \subset K$.

Definition. We say that a fractional ideal I is invertible if there is a fractional ideal J of A such that $IJ = A$.

Note that inverses are unique if they exist: if $IJ = A = IJ'$ then $J = J(IJ') = J'(IJ) = J'$.

Theorem 15.6. *If I is a nonzero fractional ideal in a Dedekind domain A, then I is invertible.*

Proof. Let $J = \{x \in K \mid xI \subset A\}$, which is clearly a fractional ideal. We claim that J is our inverse. Clearly $IJ \subset A$, so we just need to show that the inclusion $\iota : IJ \to A$ is also a surjection. This we can do locally:

Let $\mathfrak{p}$ in A be any nonzero prime ideal, so $A_\mathfrak{p}$ is a DVR with valuation ν and uniformizer t. Then $IA_\mathfrak{p} = (\mathfrak{p}A_\mathfrak{p})^r = (t^r)A_\mathfrak{p}$ for some $r \in \mathbb{Z}$, which is given by $r = \min\{\nu(i) \mid$

$i \in \mathfrak{p}\}$. Choose $i_0 \in \mathfrak{p}$ with $\nu(i_0) = r$. We'll be done if we can find $j_0 \in J$ such that $\nu(j) = -r$, since then $\operatorname{Im} \iota$ will include $i_0 j_0 A_{\mathfrak{p}} = A_{\mathfrak{p}}$.

To do this, pick generators $i_1, \ldots, i_n$ of I, and write $i_k = t^r \left(\frac{a_k}{s_k}\right)$, where $a_k, s_k \in A$ and $s_k \notin \mathfrak{p}$. Then take

$$j = t^{-r} \prod_k s_k;$$

this clearly works. □

Corollary 15.7. *If A is Dedekind, then the set of nonzero fractional ideals of A forms an abelian group $I(A)$ under multiplication, with A as identity element.*

The group $I(A)$ has as a subgroup the group $P(A)$ of principal fractional ideals of A, and the quotient is known as the *class group* $\mathrm{Cl}(A)$ of A. If A is a PID, then $\mathrm{Cl}(A)$ is trivial, and more generally $\mathrm{Cl}(A)$ measures the failure of A to be a PID.

The class group is an important object in number theory: it is a fact (that is proven in classes like Math 129) that if K is a finite extension of $\mathbb{Q}$ and A is the integral closure of $\mathbb{Z}$ in K, then $\mathrm{Cl}(A)$ is finite, and a lot of study has been done on understanding its size.

In algebraic geometry, the class group is also known as the Picard group, and often has a geometric interpretation. For instance, if k is an algebraically closed field, and $A = k[x]/(x^2 - f(y))$ where f is a cubic with distinct roots in k, then one can show that the non-identity elements of $\mathrm{Cl}(A)$ are in one-to-one correspondence with the maximal ideals of A, that is, with the points of the curve $V((x^2 - f(y)) \subset k^2$. This gives us a group structure on the set $V((x^2 - f(y)))$ union an extra "point at infinity"; this is what is known as the "group law on an elliptic curve."

(The above can also be done when k is not algebraically closed, and it's still the case that the case that the non-identity elements of $\mathrm{Cl}(A)$ are in one-to-one correspondence with points $(a_x, a_y) \in k^2$ such that $a_x^2 = f(a_y)$.)

16 Fields and Galois Theory

Now we come to the next unit of our class: Galois theory, which is the study of finite extensions of fields and their automorphisms. References for this are Dummit & Foote + James Milne's notes `http://www.jmilne.org/math/CourseNotes/FT.pdf`.

We'll be able to apply what we already know about rings to fields. So let's do a definition that's a special case of what was done before:

Definition. If L and K are fields with $K \subset L$, we say that "L/K is a *field extension*". If L is a finite K-algebra, we say that L/K is a *finite field extension*

(Recall that L being a finite K-algebra means that L is finite as a K-module; equivalently, L is a finite-dimensional K-vector space.) Then the following proposition follows from what we've done before

Proposition 16.1. *A field extension* L/K *is finite if and only if* L *is generated by finitely many elements that are algebraic over* K*.*

Proof. By a previous result in class, L/K is finite if and only if L is generated by finitely many elements that are integral over K. Now use that $\alpha \in L$ is integral over K if and only if it is algebraic over L (since in a field, nonzero leading coefficients can always be rescaled to be 1). □

But we also have a couple of advantages here: vector spaces have dimension.

Definition. If L/K is a field extension, define $[L : K] = \dim_K L$.

Dimension behaves well in towers of field extensions: Then

Proposition 16.2. *If* $K \subset L \subset E$ *are all fields, then* $[E : L][L : K] = [E : K]$.

Proof. (Sketch:) Let $\ell_1, \ldots, \ell_d$ be a basis for L as a K-vector space, and let $e_1, \ldots, e_{d'}$ be a basis for E as an L-vector space. Then the set

$$\begin{array}{cccc} e_1\ell_1, & e_1\ell_2, & \ldots & e_1\ell_d \\ e_2\ell_1, & e_2\ell_2, & \ldots & e_2\ell_d \\ \vdots & \vdots & & \vdots \\ e_{d'}\ell_1 & e_{d'}\ell_2 & \ldots & e_{d'}\ell_d. \end{array}$$

forms a K-basis for L. □

Also, because are working over fields, we can easily classify field extensions generated (as a K-algebra) by a single element.

Definition. We say that a finite extension L/K is monogenic if $L = K[\alpha]$ for some $\alpha \in L$.

Note: this is not the most standard terminology. "Monogenic" to mean "generated (as an algebra) by a single element)" is used a bunch of places in math, but more often for rings than for fields, e.g. they would talk about a monogenic $\mathbb{Z}$*-algebra, and some books say that a finite extension* K *of* $\mathbb{Q}$ *is monogenic is if the integral closure of* $\mathbb{Z}$ *in* K *is monogenic.)*

If $L = K[\alpha]$ is monogenic, then the ring homomrphism $\phi : K[x] \to L$ sending $f(x) \mapsto f(\alpha)$ is surjective, so $L \cong K[x]/\ker\phi$. Now, $\ker(\phi)$ is an ideal of the PID $K[x]$, so it is of the form $\ker(\phi) = (g)$ where g is the unique lowest degree monic element of $\ker(\phi)$; equivalently, g is the unique monic polynomial of lowest degree with $g(\alpha) = 0$. This g is called the "minimal polynomial" of α.

This gives us a bijection (isomorphism classes of monogenic K-algebras L along with a distinguished generator $\alpha \in L$) $\leftrightarrow$ (monic irreducible polynomials $g \in K[x]$).

Here the maps are $(L, \alpha) \mapsto g$ where g is the minimal polynomial of α in $K[x]$, and g mapsto $(K[x]/(g), \overline{x})$.

Additionally, if $L = K[x]/g(x)$, we have a bijection between K-algebra homomorphisms from L to any K-algebra L' and roots of f in L'. Also, if L' is nonzero, then because L is a field any L-algebra homomorphism $\phi : L \to L'$ is injective ($\ker\phi$ is an ideal of L and $1 \notin \ker\phi$, so $\ker\phi = 0$).

In Galois theory we'll be looking at automorphism groups of finite extensions L/K of fields.

Definition. If L/K is a field extension, then $\mathrm{Aut}(L/K)$ is the set of all K-algebra automorphisms of L.

Note that if $\phi : L \to L$ is a K-algebra homomorphism, ϕ is automatically injective (same argument as above), and if additionally L/K is finite, ϕ must also be surjective (since it's a K-linear map between vector spaces of the same dimension).

We now do a few examples of $\mathrm{Aut}(L/K)$:

Example. $L/K = \mathbb{Q}[\sqrt{D}]/\mathbb{Q}$ (for any nonsquare $D \in \mathbb{Q}$). Here, $\mathbb{Q}[\sqrt{D}]$ is a monogenic extension of $\mathbb{Q}$ where $\sqrt{D}$ having minimal polynomial $x^2 - D$. Now $x^2 - D$ has two roots $\sqrt{D}$ and $-\sqrt{D}$ in L, and so there are two K-algebra homomorphisms $L = \mathbb{Q}[\sqrt{D}] \to L$, one that sends $\sqrt{D} \to \sqrt{D}$ (the identity) and one that sends $\sqrt{D}$ to $-\sqrt{D}$ (so sends $a + b\sqrt{D}$ to $a - b\sqrt{D}$. Hence $\mathrm{Aut}(L/K)$ has order 2, and must be the cyclic group C_2.

Example. $L/K = \mathbb{Q}[\sqrt[3]{2}]/\mathbb{Q}$. Again, L is a monogenic extension of $\mathbb{Q}$, and its generator $\sqrt[3]{2}$ now has minimal polynomial $x^3 - 2$. In this case, however, $x^3 - 2$ has no roots in $\mathbb{Q}[\sqrt[3]{2}]$ (since, for example, the latter can be embedded in $\mathbb{R}$, which only contans one cube root of 2). Hence $\mathrm{Aut}(L/K)$ contains only the identity element.

Example. We can enlarge the field of the previous example so that it has more automorphisms. One way of doing this is to take $L = \mathbb{Q}[\sqrt[3]{2}, \omega]$, where $\omega = \frac{1+\sqrt{-3}}{2}$ is a primitive cube root of unity ($\omega^3 = 1$ but $\omega \neq 1$, so the minimal polynomial of ω is $x^2 + x + 1$.)

(Question asked in class: what do you mean by $L = \mathbb{Q}[\sqrt[3]{2}, \omega]$*? There are two reasonable answers here : one is to construct* L *algebraically as* $\mathbb{Q}[x, y]/(x^3 - 2, y^2 + y + 1)$*, and verify that this is a field. The other is to take* L *to be the subfield of the algebraic closure* $\overline{\mathbb{Q}}$ *generated by elements* $\alpha, \beta \in \overline{\mathbb{Q}}$ *such that* $\alpha^3 = 2$ *and* $\beta^2 + \beta + 1 = 0$*; this turns out to be isomorphic to the previously constructed field, and doesn't depend on choice of* α *or* β*.)*

Here we can think of L as sitting at the top of a tower of fields: $K = \mathbb{Q} \subset L' = \mathbb{Q}[\omega] \subset L = \mathbb{Q}[\sqrt[3]{2}, \omega]$, where the degrees are $[L' : K] = 2$, $[L : L'] = 3$, so $[L : K] = 6$. To find the embeddings of L into itself we'll first find the embeddings of L' into L, and then see how each of them can be extended to an embedding of L into itself.

Now, $L'/\mathbb{Q}$ is a monogenic extension generated by ω with minimal polynomial $f(x) = x^2 + x + 1$. This f has two roots in L, namely $1 \pm \sqrt[3]{2}$, and there are two embeddings $\phi_1, \phi_2 : L' \to L$.

Now, for each of these embeddings $\phi_i : L \to L$, $\mathbb{Q}$-algebra homorphisms from $L \to L$ that restrict to ϕ_i on L are the same as L'-algebra homomorphisms $L \to L$ where the L'

algebra structure on the domain copy of L is the standard one, but on the codomain L we use the L'-algebra structure coming from the embedding $\phi_i : L \to L'$ (which will be different depending on whether $i = 1$ or $i = 2$); that is, $\ell x = \phi_i(\ell)x$ for $\ell \in L$ and $x \in L'$.

Since L is a monogenic extension of $L' : L = L'[\sqrt[3]{2}]$, the extensions of $\phi_i : L \to L$ are in bijection with the roots of the polynomial $x^3 - 2 \in L'[x]$ in the L'-algebra L. In either case, there are 3 of these: $\sqrt[3]{2}, \omega \cdot \sqrt[3]{2}$, and $\omega^2 \cdot \sqrt[3]{2}$, so each ϕ_i extends to 3 different $\mathbb{Q}$-algebra embeddings of L into itself. Since each of these are automorphisms of L, this means that $\mathrm{Aut}(L/K)$ has order 6.

One can show that in fact $\mathrm{Aut}(L/K) \cong S_3$, where the isomorphism comes from the fact that $\mathrm{Aut}(L/K)$ acts on the 3-element set $\sqrt[3]{2}$, $\omega \cdot \sqrt[3]{2}$, and $\omega^2 \cdot \sqrt[3]{2}$.

Note that in all three cases $\#(\mathrm{Aut}(L/K))$ was a divisor of $[L : K]$, and that in the first and last cases they were equal.

Definition. L is a splitting field for a monic polynomial f over K if there are $\alpha_1, \ldots, \alpha_n \in L$ ($n = \deg f$) which generate L as a K-algebra such that $f(x) = (x - \alpha_1)(x - \alpha_2) \cdots (x - \alpha_n)$.

(Note we're not assuming that f is irreducible in K; e.g. if f already splits into linear factors in K then K is already a splitting field for f.)

Theorem 16.3. *Every monic polynomial $f \in K[x]$ has a splitting field.*

Proof. Write $f(x) = x^n + c_{n-1}x^{n-1} + \cdots + c_0$. First we construct a ring A such that this is the case: $A = K[x_1, \ldots, x_n]/I$ where I is generated by the elements $c_{n-k} - (-1)^k s_k$ for $k = 1, \ldots, n$ (here s_k is the kth elementary symmetric polynomial $\sum x_{i_1} x_{i_2} \cdots x_{i_k}$ defined in the first problem set).

Now let K be the quotient of A by any maximal ideal, and α_i the image of x_i in K. Then $f(x) = (x - \alpha_1)(x - \alpha_2) \cdots (x - \alpha_n)$, and f is certainly generated by the α_i. □

Last time we defined

Definition. L is a splitting field for a monic polynomial f over K if there are $\alpha_1, \ldots, \alpha_n \in L$ ($n = \deg f$) which generate L as a K-algebra such that $f(x) = (x - \alpha_1)(x - \alpha_2) \cdots (x - \alpha_n)$.

and showed that they exist for any f. This time we'll show they are unique. But examples first:

Example. $K = \mathbb{Q}$, and $L = \mathbb{Q}[\omega, \sqrt[3]{2}]$ is a splitting field of $f(x) = x^3 - 2$ over $\mathbb{Q}$. To check this, observe that f splits as $(x - \sqrt[3]{2})(x - \omega \cdot \sqrt[3]{2})(x - \omega^2 \cdot \sqrt[3]{2})$ in $L[x]$. Also the roots of f generate L as a $\mathbb{Q}$-algebra since $\omega = \frac{1}{2}(\sqrt[3]{2})^2(\omega \mathrm{cdot} \sqrt[3]{2})$.

Example. Here's a case where $K = \mathbb{F}_p$, and we take $f(x) = x^{p^n} - x$ for any integer $n \geq 1$. By the result we proved last time, we know that K must have some splitting field L. In this case, something special happens. The set of roots of f forms a sub K-algebra of L (since $(x+y)^{p^n} = x^{p^n} + y^{p^n}$ ind characteristic p, and also $(xy)^{p^n} = x^{p^n} y^{p^n}$. Hence L is actually equal to the set of roots of f in L. We'll see at the end of today's class that f must have distinct roots, so $|L| = p^n$.

Conversely, if L is any field with $|L| = p^n$, then $f(x) = 0$ for all $x \in L$, because either $x = 0$ or $x^{p^n - 1} = 1$ (since the multiplicative group $L^\times$ has order $p^n - 1$) so L is a splitting field for f.

(As a corollary, this means that once we prove that all splitting fields of f are isomorphic, we'll also have that all fields of order p^n are isomorphic.

Now we're going to show that all splitting fields of f are isomorphic to each other (but not canonically so). In doing so we'll also get some extra on automorphism groups of fields.

Proposition 16.4. *Let $f(x) \in K[x]$. Suppose that L is a field extension of K generated by elements $\{\alpha_i i\}$ which are roots of f (but f doesn't have to split in L). Suppose that E is a field extension of K in which f splits (not necessarily a splitting field, since E need not be generated by the roots of f). There there are $\leq [L : K]$ embeddings $L \to E$ of K-algebras, and equality holds if f has distinct roots in L.*

Proof. Induct on the degree of $[L : K]$. The base case $[L : K] = 1$ is clear.

Choose any α in L with $f(\alpha) = 0$, and consider the intermediate field $K[\alpha] \cong K[x]/g(x)$ where g is the minimal polynomial of α. Let $n = [K[x] : K] = \deg g$.

Now, any embedding $\phi : L \to E$ restricts to an embedding $\phi_0 : K[\alpha] \to E$. Now we have a bijection $\{\text{embeddings of } K[\alpha] \text{ in } E\} \leftrightarrow \{\text{roots of } g \text{ in } E\}$. Since g splits into linear factors in E, g has at least 1 and at most n roots in E (equality iff g has no repeated roots).

Now, for any given embedding $\phi_0 : K[\alpha] \to E$, the K-algebra embeddings $L \to E$ that restrict to ϕ_0 can be identified with the $K[\alpha]$-algebra embeddings $L \to E$ where E is given a structure of $K[\alpha]$-algebra via the inclusion map $\phi_0 : L \to E$.

Now, apply the induction hypothesis with $K[\alpha]$ in place of K. This means that there are at least one and at most $[L : K[\alpha]]$ possible embedding ϕ with $\phi|_K[\alpha] = \phi_0$.

Since this is true for each of the $\leq n$ possible values of ϕ_0, there at most $n[L : K[\alpha]] = [K[\alpha] : K][L : K[\alpha]] = [L : K]$ possible embeddings ϕ, as desired, and equality holds in the case when f has no repeated roots in E. □

Corollary 16.5. *Splitting fields are unique.*

Corollary 16.6. *If L/K is a finite extension of fields, then $|\operatorname{Aut}(L/K)| \leq [L : K]$. If $[L : K]$ is the splitting field of a polynomial with no repeated roots, then $|\operatorname{Aut}(L/K)| = [L : K]$.*

Proof. Choose a polynomial $f \in K[x]$ such that the roots of f generate L. Let $L' \supset L$ be an extension of L in which f splits. Then by the above, there are at most $[L : K]$ distinct

embeddings of L into L'. Since any element of $\mathrm{Aut}(L/K)$ gives an embedding of L in L', this means that there are at most $\mathrm{Aut}(L/K)$ such embeddings. □

This raises the questions; which polynomials have no repeated roots? Fortunately, there's a nice criterion for this.

Proposition 16.7. *Let $f(x) \in K[x]$, and let L be any field in which f splits into linear factors. Let $f'(x) \in K[x]$ be the formal derivative of f. Then f has no repeated roots in L if and only if $\gcd(f, f') = 1$ in $K[x]$.*

A note on GCDs; first: you might worry about the gcd depending upon which ring we consider f and f' as being in. However, if $g = \gcd(f_1, f_2)$ in $K[x]$, this means that $(g) = (f_1, f_2)$ in $K[x]$, so the same is true if we extend scalars to $L[x]$ and $g = \gcd(f_1, f_2)$ in $L[x]$ for any L containing K.

Proof. By the above note, may assume that $L = K$. Then this is true by standard differentiation rules: f factors as $(x - \alpha_1)\dots(x - \alpha_n)$, so $f'(\alpha_i) = \prod_{j \neq i}(\alpha_i - \alpha_j)$ is 0 iff α_i is a double root of f. Hence $\gcd(f, f') = 0$ iff no roots of f are also roots of f' iff f has no double roots. □

Definition. We say that f is *separable* if either of the above conditions holds. We say that a field extension L/K is separable if for any $\alpha \in L$, the minimal polynomial of α in $K[x]$ is separable.

Corollary 16.8. *If K has characteristic 0 and f is irreducible, then f is separable.*

Counterexample in characteristic p: $K = \mathbb{F}_p(t)$, $f(x) = x^p - t$, $L = \mathbb{F}_p(t^{1/p})$ then $f(x) = (x - t^{1/p})^p$ in $L[x]$. However, f can be shown to be irreducible; the problem here is that $f'(x) = 0$, so $\gcd(f, f') = f$, from which we can deduce that all roots of f are multiple roots, as confirmed by the example above.

On the other hand, the polynomial $f(x) = x^{p^n} - x \in \mathbb{F}_p[x]$ is separable because $f' = -1$. This shows the result we needed above, that f has no repeated roots in its splitting field.

17 The fixed field of a group of automorphisms

Last week, we've been playing the game of considering a field extension $\mathrm{Aut}(L/K)$, and getting from it the group $\mathrm{Aut}(L/K)$ of automorphisms of L that fix K.

Now we're going to go the other way, and get a field extension from a group. Suppose that L is a field, and $G \subset \mathrm{Aut}(L)$ is a group of field automorphisms of L. Then we can define the fixed field $L^G = \{x \in L \mid g(x) = x \text{ for all } g \in G\}$. (Note that this is a special case of the "ring of invariants" construction we defined at the beginning of the semester any time we have a group G acting on a ring A.)

Theorem 17.1. *Let* L *be a field, let* $G \subset \mathrm{Aut}(L)$ *be a finite group, and let* L^G *be the fixed field of* G*. Then* $[L : L^G] \leq |G|$*.*

Proof. Write $G = \{g_1, \ldots, g_n\}$ and $K = L^G$; we need to show that $[L : K] = n$. Equivalently, we need to show that any set $a_1, \ldots, a_m$ of elements of L with size $m > n$ is linearly dependent over K. Look at the system of linear equations

$$c_1 g_1(a_1) + c_2 g_1(a_2) + \cdots + c_n g_1(a_m) = 0 \tag{2}$$
$$c_1 g_2(a_1) + c_2 g_2(a_2) + \cdots + c_n g_2(a_m) = 0 \tag{3}$$
$$\vdots \tag{4}$$
$$c_1 g_n(a_1) + c_2 g_n(a_2) + \cdots + c_n g_n(a_m) = 0. \tag{5}$$

as a system of linear equations with coefficients in L.

Consider the set V of all solutions $(c_1, \ldots, c_n)$ in L^n. We'll be done if we can show that V contains some nonzero vector of K^n.

We see that V is an L-vector space of dimension $\geq m - n$; in particular is not just $\{0\}$. Also, V has the property that $g(V) \subset V$ for all $g \in G$.

Now comes the trick. Choose nonzero $v = (c_1, \ldots, c_n) \in V$ such that v has the fewest possible nonzero entries. Without loss of generality, assume $c_1 \neq 0$. Then by rescaling, we may assume that $c_1 \in K$. Now we'll show that v must lie in K^n. Otherwise, suppose $c_i \notin K = L^G$: so there is some $g \in G$ such that $gc_i \neq c_i$. Then $v' = gv - v \neq 0$, and has fewer nonzero entries than v. □

Corollary 17.2. *If* L *is a field and* G *is a finite group of automorphisms of* L*, then* $G = \mathrm{Aut}(L/L^G)$ *and* $[L : L^G] = |G|$*.*

Proof. We have

$$[L : L^G] \leq |G| \leq |\mathrm{Aut}(L/L^G)| \leq [L : L^G].$$

Hence everything above must be an equality, giving the desired result. □

18 Separable, normal, and Galois extensions

Definition. We say that a polynomial $f(x)$ is separable if it has no repeated roots. We say that an extension L/K is separable if for any $x \in L$, the minimal polynomial $f(x) \in L[x]$ is separable.

Definition. We say that an extension L/K is normal if any irreducible $f \in K[x]$ which has a root in L splits into linear factors in $L[x]$.

A field extension L/K being both Separable and normal is equivalent to: If f is the minimal polynomial of some α in L, then f has $\deg f = [K[\alpha] : K]$ roots in L.

Example. $\mathbb{F}_p(t^{1/p})/\mathbb{F}_p(t)$ is not separable , since $\alpha = t^p$ has non-separable minimal polynomial $x^p - t$.

Example. $\mathbb{Q}[\sqrt[3]{2}]/\mathbb{Q}$ is not normal, because the polynomial $x^3 - 2$ only has one linear factor in $\mathbb{Q}[\sqrt[3]{2}][x]$.

On the other hand, we'll see that $\mathbb{Q}[\sqrt[3]{2}, \omega]$ is normal.

Definition. We say that L/K is Galois if L/K is finite and $K = L^{\mathrm{Aut}(L/K)}$.

Example. $L/K = \mathbb{Q}[\sqrt{3}[2]]/\mathbb{Q}$ is not Galois, because $\mathrm{Aut}(L/K)$ is trivial and so $L^{\mathrm{Aut}(L/K)} = L$, not K.

Theorem 18.1. *Let* L/K *be an extension of fields. TFAE*

a) L/K *is Galois*

b) There exists a finite group $G \subset \mathrm{Aut}(L)$ *such that* $K = L^G$.

c) L/K *is finite, normal, and separable*

d) L/K *is the splitting field of a separable* $f \in K[x]$

e) $|\mathrm{Aut}(L/K)| = [L : K] < \infty$

Proof. We're almost out of time, so let's just point out the easy parts. a) $\implies$ b): take $G = \mathrm{Aut}(L/K)$.

d) $\implies$ e): We showed this on Friday.

□

Today we'll prove the theorem we stated last time:

Theorem 18.2. *Let* L/K *be an extension of fields. TFAE*

a) L/K *is Galois*

b) There exists a finite group $G \subset \mathrm{Aut}(L)$ *such that* $K = L^G$.

c) L/K *is finite, normal, and separable*

d) L/K *is the splitting field of a separable* $f \in K[x]$

e) $|\mathrm{Aut}(L/K)| = [L : K] < \infty$

Proof. a) $\implies$ b): take $G = \mathrm{Aut}(L/K)$. b) $\implies$ c): We have that L/L^G is finite, since $[L : L^G] \leq G$. As pointed out last time, the condition of being normal and separable is equivalent to saying that for any $\alpha \in L$ the minimal polynomial f of α in $K[x]$ splits into distinct linear factors in $L[x]$.

To show this, let $\{\alpha_1 = \alpha, \alpha_2, \ldots, \alpha_n\} = \{g\alpha \mid g \in G\}$ be the orbit of α under G. Let $f(x) = (x - \alpha_1)(x - \alpha_2)\cdots(x - \alpha_n)$. We claim that f is actually the minimal polynomial of α in $K[x]$.

The coefficients of f are (up to sign) elementary symmetric polynomials in the α_i, so they are fixed by the action of G (which permutes the α_i), and must lie in $L^G = K$. Hence $f \in K[x]$. We now must show that f divides any other polynomial $p \in K[x]$ such that $p(\alpha) = 0$. But since $p \in K[x]$, $p(\alpha) = 0$ implies $p(g\alpha) = g(p(\alpha)) = 0$ for any $g \in G$, and so p must have all the α_i as roots, hence f divides g in $K[x]$.

Since the α_i are distinct, f is separable.

c $\implies$ d): Choose generators $\alpha_1, \ldots, \alpha_r$ for L/K, and let f_i be the minimal polynomial of α_i. Let f be the product of the distinct f_i (that is, if any polynomials appear in the list $f_1, \ldots, f_n$ more than once, we only include them once in the product). We have that f is a product of distinct separable polynomials, so it's separable. And then L is the splitting field of f by construction.

d) $\implies$ e): We showed this on Friday.

e) $\implies$ a): Suppose not. Let $K' = K^{\mathrm{Aut}(L/K)}$. Then $K' \supset K$ and $[L : K'] = |\mathrm{Aut}(L/K)| = [L : K]$ so $K' = K$. □

If L/K is Galois, then we write $\mathrm{Gal}(L/K)$ for $\mathrm{Aut}(L/K)$.

Corollary 18.3. *If L/K is Galois, and $K \subset K' \subset L$, then L/K' is Galois.*

Proof. Use the criterion d) from above: L is the splitting field of some separable polynomial $f(x)$ over K. Then L is still the splitting field of f over K'. □

Theorem 18.4 (Fundamental Theorem of Galois Theory)**.** *Let L/K be a finite extension of fields, and let $G = \mathrm{Gal}(L/K)$. Then there is a bijection between (subgroups $H \subset G$) and (intermediate field extensions F with $K \subset F \subset L$.) giving by $H \mapsto L^H$ and $\mathrm{Gal}(L/F) \hookleftarrow F$.*

Furthermore, this correspondence has the following properties.

It is inclusion-reversing: if $H \leftrightarrow F$ and $H' \leftrightarrow F'$ then $H \subset H'$ iff $F \supset F'$.

If $H \leftrightarrow F$, then for $\sigma \in \mathrm{Gal}(L/K)$, $\sigma H \sigma^{-1} \leftrightarrow \sigma F$.

And H is a normal subgroup of G iff F/K is normal. In this case, $\mathrm{Gal}(F/K) \cong (\mathrm{Gal}(L/K)/\mathrm{Gal}(L/F)) = G/H$.

Proof. We've proved all the hard parts of this already: so the bijection will just fall out of what we've done.

We showed on Monday that $\mathrm{Gal}(L/L^H) = \mathrm{Aut}(L/L^H) = H$ for any finite $H \subset \mathrm{Aut}(L)$.

On the other hand, we also have that L/F is Galois by the corollary above, and so $L^{\mathrm{Gal}(L/F)} = F$.

This shows that the two maps given above are inverses, and so we have a bijective correspondence.

The rest of this is left as an exercise.

□

Example. The field extension $L/K = \mathbb{Q}[\sqrt[3]{2}, \omega]/\mathbb{Q}$. We saw that this has Galois group $G = S_3$. The subgroups of S_3 are S_3, $\{id\}$, $A_3 \cong C_3$(generated by either of the 3-cycles in S_3), and three subgroups isomorphic to C_2 (each generated by one of the transpositions in S_3).

Clearly $S_3 \leftrightarrow K = \mathbb{Q}$ and $\{id\} \leftrightarrow L$.

Now we do the case $H = A_3$, and we determine the corresponding intermediate field $F = L^H$. We have that $[L : L^H] = 3$, so $[L^H : K] = 2$, and L^H is a quadratic extension of $K = \mathbb{Q}$. We can check that $\omega \in L^H$, so in fact H corresponds to the quadratic subextension $\mathbb{Q}[\omega]$.

Now, let H be one of the copies of C_2 in G; for instance, the copy generated by the automorphism that fixes $\sqrt[3]{2}$ and switches $\omega \cdot \sqrt[3]{2}$ with $\omega^2 \cdot \sqrt[3]{2}$. Then clearly $\sqrt{3}[2] \in L^H<$ and by a similar argument to before, $[L^H : K] = 3$, so H corresponds to $\mathbb{Q}[\sqrt[3]{2}]$.

Likewise the other two copies of C_2 inside G correspond to $\mathbb{Q}[\omega \cdot \sqrt[3]{2}]$ and $\mathbb{Q}[\omega^2 \cdot \sqrt[3]{2}]$.

19 More examples of Galois extensions

Example. $L/K = \mathbb{F}_{p^n}/\mathbb{F}_p$ - splitting field of $x^{p^n} - x$, so Galois, and $\mathrm{Gal}(\mathbb{F}_{p^n}//\mathbb{F}_p) = n$. Can write down an automorphism $\mathrm{Frob} \in \mathrm{Gal}(\mathbb{F}_{p^n}/F)$ (the "Frobenius automorphism") given by $\mathrm{Frob}(a) = a^p$. (Since $(a+b)^p = a^p + b^p$ in characteristic p.)

We claim that Frob has order n, so that it generates $\mathrm{Gal}(\mathbb{F}_{p^n}/\mathbb{F}_p)$. Indeed, for $a \in \mathbb{F}_{p^n}$, $\mathrm{Frob}^n(a) = a^{p^n} = a$. On the other hand, for $k < n$, the polynomial $x^{p^k} - x$ has at most $p^k < |\mathbb{F}_{p^n}|$ roots in $\mathbb{F}_{p^n}$, so there exists $a \in \mathbb{F}_{p^n}$ such that $\mathrm{Frob}^k(a) \neq a$, so $\mathrm{Frob}^p(a)$ is not the identity. Hence $\mathrm{Gal}(\mathbb{F}_{p^n}/\mathbb{F}_p)$ is a cyclic group generated by Frob.

Any subgroup of H $\mathrm{Gal}(\mathbb{F}_{p^n})$ will then also be cyclic, generated by Frob^r for some r dividing n. In this case, the fixed field $F = (\mathbb{F}_{p^n})^H$ consists of the roots of $x^{p^r} - x$ in $\mathbb{F}_{p^n}$. We can show that $x^{p^r} - x$ divides $x^{p^n} - x$, so $x^{p^n} - x$ splits into p^r distinct linear factors of in $\mathbb{F}_{p^n}$, and so F is a copy of $\mathbb{F}_{p^r}$ inside $\mathbb{F}_{p^n}$.

(Alternatively, as suggested in class: use the Galois correspondence to show that $[L : \mathbb{F}_p] = r$, so L must be isomorphic to $\mathbb{F}_p^r$.)

Example. Now we do an example that we can construct by letting K be L^G for some subgroup G of $\mathrm{Aut}(L)$. Let $L = \mathbb{C}(x_1, \ldots, x_n)$. We saw before that S_n acts on $\mathbb{C}[x_1, \ldots, x_n]$ with subring of invariants $\mathbb{C}[x_1, \ldots, x_n]^G = \mathbb{C}[s_1, \ldots, s_n]$. By HW, this action extends to $\mathbb{C}(x_1, \ldots, x_n)$ with fixed field $\mathbb{C}(x_1, \ldots, x_n)^{S_n} = \mathbb{C}(s_1, \ldots, s_n)$. (And we know that all the s_i are algebraically independent, so this field is the isomorphic to the rational function field in the variables s_i over $\mathbb{C}$.)

We could also have constructed this as a splitting field, taking $K = \mathbb{C}(s_1, \ldots, s_n)$ as a field of rational functions in the variables s_i, and letting L be the splitting field of the polynomial

$$(x - x_1)(x - x_2) \cdots (x - x_n) = x^n - s_1 x^{n-1} + s_2 x^{n-2} + \cdots + (-1)^n s_n \in K[x].$$

For an example of the Galois correspondence: let $H = A_n \subset S_n$. Then $[L^{A_n} : K] = [L : K]/[L : L^{A_n}] = |S_n|/|A_n| = 2$, so L^{A_n} is a quadratic extension of K. One can check that the element $\Delta = \prod_{i<j}(x_i - x_j)$ is in L^{A_n} but not in K, so it generates L^{A_n} over K, and $\mathbb{C}(x_1, \ldots, x_n)^{A_n} = \mathbb{C}(s_1, \ldots, s_n, \Delta)$. One can also check that $\Delta^2 \in \mathbb{C}(s_1, \ldots, s_n)$, confirming that we do have a quadratic extension.

Example. Let $K = \mathbb{Q}$, and let $L = \mathbb{Q}(\zeta_n)$, which we define first as the subfield of $\mathbb{C}$ generated by $\mathbb{Q}$ and by $\zeta_n = e^{2\pi i/n}$. The L is the splitting field of $x^n - 1$ over $\mathbb{Q}$, since

$$x^n - 1 = (x - \zeta_n^0)(x - \zeta_n)(x - \zeta_n^2) \cdots (x - \zeta_n^{n-1}).$$

Now, for any $g \in \mathrm{Gal}(\mathbb{Q}(\zeta_n)/\mathbb{Q})$, $g(\zeta_n)$ must satisfy $g(\zeta_n)^n = 1$ but $g(\zeta_n)^k \neq 1$ for any $k < n$. The first condition means that $g(\zeta_n) = \zeta_n^r$ for some $r > 0$, and the second implies that $\gcd(n, r) = 1$

This gives us a map $\mathrm{Gal}(\mathbb{Q}(\zeta_n)/\mathbb{Q}) \to (\mathbb{Z}/n\mathbb{Z})^\times$, sending n to the r such that $g(\zeta_n) = \zeta_n^r$. Easy to check it's a group homomorphism; also, it's injective because ζ_n generates $\mathbb{Q}(\zeta_n)$. What's hard is to show that it's always surjective; we'll show this either Monday or on the next problem set. (The statement of surjectivity is equivalent to the statement that the cyclotomic polynomial $\prod_{\gcd(r,n)=1}(x - \zeta_n^r) \in \mathbb{Q}[x]$ is irreducible.

Now we move on to do some corollaries to Galois theory.

Theorem 19.1. *If L/K is a finite separable field extension, there are only finitely many intermediate fields F with $K \subset F \subset L$.*

Proof. If this statement is true for some L' containing L, it's also true for L. So, enlarge L so that L/K is Galois; e.g. by taking generators $\alpha_1, \ldots, \alpha_n$ for L as a K-algebra and letting L' be the splitting field of a separable polynomial which has $\alpha_1, \ldots, \alpha_n$ as some of its roots.

Then this follows since $\mathrm{Gal}(L/K)$ has only finitely many subgroups. □

We now prove a theorem which in some books is done before the Fundamental Theorem of Galois theory (and is used in the proof); we're deducing it as a corollary instead.

Theorem 19.2 (Primitive Element)**.** *Suppose that L/K is a finite separable field extension. Then there exists $\alpha \in L$ such that $L = K[\alpha]$.*

Proof. We'll split into cases depending upon whether K is a finite or an infinite field. The cases will use completely different arguments.

Case 1:$|K| < \infty$. Then $|L| < \infty$ also. We use the following fact:

For any finite field L, $L^\times$ is a cyclic group. This can be proved using the classification of finitely generated abelian groups: $L^\times \cong \mathbb{Z}/d_1\mathbb{Z} \times \mathbb{Z}/d_2\mathbb{Z} \times \cdots \times \mathbb{Z}/d_n\mathbb{Z}$ for positive integers $d_1, \ldots, d_n$ such that $d_1 \mid d_2 \mid \cdots \mid d_n$. Then the equation $x^{d_n} = 1$ is satisfied by every $x \in L^\times$, but on the other hand it can have at most d_n roots in any field, so we must have $|L^\times| \leq d_n$. This is only possible if $L^\times \cong \mathbb{Z}/d_n\mathbb{Z}$ is cyclic.

Now, let $a \in L^\times$ be such that a generates $L^\times$ as a cyclic group. Then a generates L as a K-algebra as well, so $L = K[\alpha]$.

Case 2: $|K| = \infty$. For this, we use a lemma:

Lemma 19.3. *Let V be a finite-dimensional vector space over an infinite field K. Then V cannot be written as the union of finitely many subspaces V_i.*

Proof of Lemma. Identify $V = K^n$. For each i, let $\ell_i \in K[x_1, \ldots, x_n]$ be a linear function that vanishes on V_i. Let $p \in K[x_1, \ldots, x_n] = \prod_i \ell_i$. Then p is a nonzero polynomial in n variables over an infinite field, so p is not identically zero. Hence there is some $(a_1, \ldots, a_n) \in K^n = V$ that is not in any V_i, so V is not the union. □

Now, apply the lemma with $V = L$ and the V_i being the finite set of intermediate fields F with $K \subset F \subsetneq L$. The lemma tells us that we can choose $\alpha \in L$ such that $\alpha \notin F$ for any intermediate field $F \subsetneq L$. So $K[\alpha]$ can't be any proper subfield of L, and must be L itself. □

(Note that I didn't remember to do in class: the non-separable extension $L/K = k(x, y)/k(x^p, y^p)$, where k is an infinite field, does not satisfy the conclusion of either theorem above. To show that L is not of the form $K[a]$: $[L : K] = p^2$, but for any $a \in L$, $a^p \in K$ so $[K[a] : K] \leq p$. Also, the subfields $K[a]$ for $a = x + cy$ where c runs through the elements of k can be seen to be all distinct.)

20 The interaction of Galois theory with commutative algebra

We'll now combine what we know about Galois theory with some commutative algebra:

Let L/K be a Galois extension with Galois group K, let B be a subring of L such that $\mathrm{Frac}(B) = A$, let B be a subring such that $gB = B$ for all $g \in G$, and let $A = B \cap K = B^G$. Then by your current HW, $K = \mathrm{Frac}(A)$; and also note that we showed on a previous HW that B is integral over A.

A special case of this construction is the following: let L/K be a Galois extension, let A be a subring integrally closed in K, and let B be the integral closure of A in L.

Next time we'll talk about how the Galois group $\mathrm{Gal}(L/K)$ acts on the set of prime ideals of B.

Let $\mathfrak{p}$ be a prime ideal of A. Then $G = \mathrm{Gal}(L/K)$ acts on the set of primes $\mathfrak{q} \subset B$ lying above A (so $\mathfrak{q} \cap A = \mathfrak{p}$).

Here's the setup we introduced last time:

Let L/K be a Galois extension with Galois group K, let B be a subring of L such that $\mathrm{Frac}(B) = A$, let B be a subring such that $gB = B$ for all $g \in G$, and let $A = B \cap K = B^G$.

Then by your current HW, $K = \mathrm{Frac}(A)$; and also note that we showed on a previous HW that B is integral over A.

We'll be assuming today that $\mathfrak{p}$ is actually maximal, although the first theorem we prove will actually still be true if $\mathfrak{p}$ not maximal (and I'll probably put this on HW).

Consider the set $\{\mathfrak{q} \subset B \text{ prime} \mid \mathfrak{q} \cap A = \mathfrak{p}\}$ of primes of B lying above $\mathfrak{p}$. (Note that our assumption that $\mathfrak{p}$ is maximal implies that any such $\mathfrak{q}$ must also be maximal, by a result we proved back when doing commutative algebra.) The group $\mathrm{Gal}(L/K)$ acts on this set. We'll show that this action is transitive.

Theorem 20.1. *If $\mathfrak{q}$ and $\mathfrak{q}'$ are primes of B lying above $\mathfrak{p}$, there exists $g \in G$ with $g\mathfrak{q} = \mathfrak{q}'$.*

Before proving this theorem, we need a bit of commutative algebra, namely a form of the Chinese Remainder Theorem.

Proposition 20.2 (Chinese Remainder Theorem)**.** *Let $\mathfrak{m}_1, \dots, \mathfrak{m}_n$ be distinct maximal ideals of a ring A. Then the map $A \to A/\mathfrak{m}_1 \times \cdots \times A/\mathfrak{m}_n$ is surjective.*

Proof. HW. □

Now we prove that $\mathrm{Gal}(L/K)$ acts transitively on the ideals lying above $\mathfrak{p}$.

Proof. Suppose that $\mathfrak{q}'$ was distinct from $g\mathfrak{q}$ for all $g \in \mathrm{Gal}(L/K)$. By the Chinese Remainder Theorem, we can find $x \in B$ such that $x \equiv 0 \pmod{\mathfrak{q}}'$ but $x \equiv 1 \pmod{g\mathfrak{q}}$ for any $g \in G$.

Then $\prod_{g \in G} gx \in (x) \subset \mathfrak{q}' \cap A = \mathfrak{p} \subset \mathfrak{q}$. Since $\mathfrak{q}$ is prime, we must have $g_0 x \in \mathfrak{q}$ for some $g_0 \in G$. Then $x \in g_0^{-1}\mathfrak{q}$ contradicts the assumption that $x \equiv 1 \pmod{g\mathfrak{q}}$. □

Corollary 20.3. *In the above setting, there are only finitely many prime ideals of B lying above A.*

Proof. The finite group $\mathrm{Gal}(L/K)$ acts transitively on the set of prime ideals of B lying above A. □

Definition. The *decomposition group* $D(\mathfrak{q})$ of a prime $\mathfrak{q}$ of B is the subset $\{g \in \mathrm{Gal}(L/K) \mid g\mathfrak{q} = \mathfrak{q}\}$.

Next, in the setup above, let's consider the field extension $\ell/k = (B/\mathfrak{q})/(A/\mathfrak{p})$. This is clearly a finite extension. We'll show it's also normal: let $\overline{b} \in \ell \in B/\mathfrak{q}$. Lift to an element $b \in B$; then the minimal polynomial of B is $f(x) = \prod_{g \in G}(x - g(b)) \in A[x]$. Reducing mod $\mathfrak{q}$ gives a polynomial $\overline{f}$ with $\overline{b}$ as root which splits into linear factors; so the same must be true of the minimal polynomial of $\overline{b}$.

Let's assume now that ℓ/k is separable. Although this doesn't have to be the case, this will be the case when either: k has characteristic 0, or k is finite. Then ℓ/k is Galois.

There is a natural homomorphism $\phi : D(\mathfrak{q}) \to \mathrm{Gal}(\ell/k)$ as follows: any $g \in \mathrm{Gal}(L/K)$ restricts to an automorphism of B with $g(\mathfrak{q}) = \mathfrak{q}$, so induces an automorphism of $B/\mathfrak{q} = \ell$, which we'll call $\overline{g} = \phi(g)$. This automorphism $\overline{g}$ fixes the image of A, namely k, so it gives us an element of $\mathrm{Gal}(\ell/k)$.

Theorem 20.4. *Under the assumptions above, the map ϕ is a surjective homomorphism of groups $D(\mathfrak{q}) \to \mathrm{Gal}(\ell/k)$.*

Proof. (This slick proof is due to John Tate, an emeritus professor at Harvard, and one of the most influential number theorists of the second half of the 20th century.)

We need to show that $\mathrm{Im}(\phi) = \mathrm{Gal}(\ell/k)$. By the Galois correspondence, it's enough to show that $\ell^{\mathrm{Im}(\phi)} = k$. Suppose that $\overline{b} \in \ell^{\mathrm{Im}(\phi)}$. We need to show that $\overline{b} \in k$. We'll do this by showing that the minimal polynomial of $\overline{b}$ over k is linear.

Lift $\overline{b}$ to an element $b \in B$ such that $b \equiv \overline{b} \pmod{\mathfrak{q}}$ and $b \equiv 0 \pmod{g\mathfrak{q}}$ for any $g \notin D(\mathfrak{q})$.

First we construct a polynomial in $A[x]$ with b as a root, in the standard way: let $f(x) = \prod_{g\in G}(x - g(b))$. By construction, the coefficients of f are fixed by G, and so $f(x) \in A[x]$.

Hence the polynomial $\overline{f}(x) = \prod_{g\in G}(x - \overline{g(b)})$ lies in $k[x]$. Let's look at what each of the factors are in $\ell[x] = B/\mathfrak{q}[x]$. In the case where $g \notin D(\mathfrak{q})$, we have $gb \equiv 0 \pmod{\mathfrak{q}}$ by our choice of b. On the other hand, if $g \in \mathcal{D}(\mathfrak{q})$, we have $\overline{g(b)} = \overline{g}(\overline{b}) = \phi(g)(\overline{b}) = \overline{b}$ since we assume $\overline{b} \in \ell^{\mathrm{Im}(\phi)}$. Hence

$$\prod_{g\in G}(\overline{x} - \overline{g(b)}) = x^m(\prod_{g\in D(\mathfrak{q})}(x - \overline{g}(\overline{b})) = x^m(x - \overline{b})^{|D(\mathfrak{q})|}$$

, and this is still an element of $k[x]$. Dividing out, we see that also $(x - \overline{b})^m \in k[x]$.

Hence the minimal polynomial of $\overline{b}$ in $k[x]$ must be a divisor of $(x - \overline{b})^m$. Since ℓ/k was assumed to be separable, this minimal polynomial must be separable, and hence the only possibility is for it to equal $x - \overline{b}$. From this, we conclude that $x - \overline{b} \in k[x]$, so $\overline{b} \in k$, as desired. □

Example. Let $L = \mathbb{Q}[\zeta_n]$, $K = \mathbb{Q}$, $B = \mathbb{Z}[\zeta_n]$, $A = \mathbb{Z}$.

Let $\mathfrak{p} = (p)$ for any integer p relatively prime to n, and choose any prime $\mathfrak{q}$ of B lying above $\mathfrak{p}$.

Then $k = \mathbb{F}_p$, and $\ell \cong \mathbb{F}_{p^m}$ for some m. We are guaranteed an element $\mathrm{Frob} \in \mathrm{Gal}(\ell/k)$ defined as on Friday. Then there exists some element $g \in D(\mathfrak{q})$ such that $\phi(g) = \mathrm{Frob}$. You'll show on your HW that this implies that $g(\zeta_n) = \zeta_n^p$. You'll then use this to show that the cyclotomic polynomial $\Phi_n(x)$ is irreducible as claimed in Friday's class.

21 Representation Theory of Finite Groups

(References: Serre *Linear Representations of Finite Group*, Dummit & Foote, Fulton & Harris *Representeation Theory*. The first section of Serre is written for an audience of quantum chemists, so it's aimed at people with less mathematical background than this class. Dummit & Foote, on the other hand, assumes some material on non-commutative rings that we haven't covered in this class. Fulton-Harris is somewhere in between. We'll be following Serre the most closely.)

Fix a base field k; we'll usually be having $k = \mathbb{C}$ in applications, but we won't need to specialize to that case until later. If V is a k-vector space, $\mathrm{GL}(V)$ is the group of automorphisms of V as a vector space.

Definition. A representation ρ of G is a homomorphism $\rho : G \to \mathrm{GL}(V)$. By abuse of notation, sometimes we'll use V to refer to the representation.

(Note: if we have a map $\rho : G \to \mathrm{End}(V)$ and we want to check it's a representation, we just need to check that $\rho(gh) = \rho(g)\rho(h)$ and that $\rho(\mathrm{id}_G) = \mathrm{id}_V$; since this implies that $\rho(g)$ is invertible with $\rho(g)^{-1} = \rho(g^{-1})$. This is what we'll usually do.)

Example. A one-dimensional representation of G is just a group homomorphism $\rho : G \to \mathrm{GL}(k^1) \cong k^\times$.

For instance, every group G has the trivial one-dimensional representation, where V is one-dimensional and $\rho(g)(v) = v$ for all $g \in G$. But also, e.g. if G is C_n and k contains the nth roots of unity, an injective map of C_n into $k^\times$ gives a nontrivial one-dimensional representation of G.

More examples come from finite subgroups of $\mathrm{GL}(\mathbb{R}^2)$ and $\mathrm{GL}(\mathbb{R}^3)$. E.g. the dihedral group D_n is contained in $\mathrm{GL}(\mathbb{R}^2)$ as the symmetries of a regular n-gon, and this gives a representation.

Example. Another important source of representations comes from permutation representations. Suppose that S is a set with an action of G. Then let V be a vector space with basis $\{e_s\}_{s \in S}$ and, for $g \in G$ let $\rho(g) \in \mathrm{GL}(V)$ be the linear transformation determined by $\rho(g)(e_s) = e_{gs}$ for any $g \in G$. It's easy to check that this is a representation.

22 Representations and Modules over the Group Algebra

On your HW, you had to show that representations of free abelian groups corresponded to modules over the ring $k[x_1, \ldots, x_n, x_1^{-1}, \ldots, x_n^{-1}]$. We'll show that this generalizes – but for this, we'll need to use non-commutative rings in the case that G is non-abelian.

I said at the start of all semester that all rings would be assumed to be abelian under further notice. Well, this is further notice.

Definition. For a group G, the group algebra $k[G]$ is the (possibly noncommutative!) k-algebra of all finite formal linear combinations $a_1 g_1 + \ldots + a_n g_n$ (where n can be arbitrary), with the k-vector space structure defined formally, and multiplication defined formally by

$$\left(\sum_i a_i g_i\right)\left(\sum_j b_j g_j\right) = \sum_{i,j} (a_i b_j) g_i g_j.$$

Here, the identity element of $k[G]$ is $1 \cdot \mathrm{id}_G$, and k includes into $k[G]$ by $a \mapsto a \cdot \mathrm{id}_G$.

Theorem 22.1. *There is a bijective correspondence between representations of G and (left) $k[G]$-modules* [1]*, given by:*

A representation $\rho : G \to \mathrm{GL}(V)$ maps to the module M which is equal to V as a k-vector space, and has $k[G]$-module structure by

$$\left(\sum_i a_i g_i\right) v = \sum_i a_i \rho(g_i)(v).$$

On the other hand, a $k[G]$-module M maps to the representation $\rho : G \to V$ where $V = M$ as k-vector space, and $\rho(g)$ is the multiplication by g map $m_g : M \to M$.

Proof. It's straightforward to check that these maps are inverses. □

Example. If G is an infinite cyclic group with generator t, so $G = \{t^n\}_{n \in \mathbb{Z}}$, the ring $k[G]$ is the ring $k[t, t^{-1}]$ of (Laurent) polynomials $\sum_k a_k t^{n_k}$ in t and t^{-1} (where $a_k \in k$ and $n_k \in \mathbb{Z}$ for each k), and the multiplication law above is just polynomial multiplication.

This correspondence tells us how we should define morphisms of G-representations: a map $\phi : V_1 \to V_2$ is a morphism if $\phi \circ \rho_1(g) = \rho_2(g) \circ \phi$ for all $g \in G$. This is also called an "intertwining map." (You should check that this does indeed correspond to the notion of a morphism of $k[G]$-modules.) This also gives us a notion of isomorphism of representations.

23 Building new representations from old

If $\rho_1 : G \to \mathrm{GL}(V_1)$ and $\rho_2 : G \to \mathrm{GL}(V_2)$ are representations, we can define

$$\rho_1 \oplus \rho_2 : G \to \mathrm{GL}(V_1 \oplus V_2).$$

[1] A left module M over a non-commutative ring A is defined just as in the commutative case: it's an (abelian) additive group M with a multiplication map $A \times M \to M$ such that $(a+b)m = am + bm$, $a(m+n) = am + an$, $(ab)m = a(bm)$, and $1m = m$. An A-module homomorphism is defined exactly the same as in the case when A is commutative.

by

$$(\rho_1 \oplus \rho_2)(g)(v, w) = \rho_1(g)(v), (\rho_2(g)(w))).$$

We can also define a tensor product of representations:

$$\rho_1 \otimes \rho_2 : G \to GL(V_1 \otimes_k V_2).$$

Here the element

$$\rho_1 \otimes \rho_2(g) \in GL(V_1 \otimes_k V_2)$$

is the tensor product

$$\rho_1(g) \otimes \rho_2(g) : V_1 \otimes V_2 \to V_1 \otimes V_2$$

of linear maps, which we defined previously in class by

$$(\rho_1(g) \otimes \rho_2(g))(v_1 \otimes v_2) = \rho_1(g)(v_1) \otimes \rho_2(g)(v_2).$$

We can also make $V_1^* = \mathrm{Hom}_k(V_1, k)$ into a representation ρ^*, by $\rho^*(g)(\phi) = \phi \circ g^{-1}$.

More generally, we can make a representation

$$\mathrm{Hom}(\rho_1, \rho_2) : G \to GL(\mathrm{Hom}_k(V_1, V_2))$$

by

$$\mathrm{Hom}(\rho_1, \rho_2)(g)(\phi) = g \circ \phi \circ g^{-1}.$$

Note here that $\mathrm{Hom}_k(V_1, V_2)$ is not the same as the set of G-representation homomorphisms (or $k[G]$-module homomorphisms) $\mathrm{Hom}_G(V_1, V_2)$; it's generally larger.

However, there is a relationship: we have

$$\mathrm{Hom}_G(V_1, V_2) = \mathrm{Hom}_k(V_1, V_2)^G = \mathrm{Hom}_k(V_1, V_2)^{\mathrm{Hom}(\rho_1, \rho_2)(G)}$$

is the subset of vectors invariant under the action of the subgroup

$$\mathrm{Hom}(\rho_1, \rho_2)(G) \subset GL(\mathrm{Hom}_k(V_1, V_2).$$

(Exercise: check this!)

24 Irreducibility

Let $\rho : G \to GL(V)$ be a representation.

Definition. An *invariant subspace* W of V is a subspace W such that $\rho(g)(W) \subset W$ for all $g \in G$. (Since we also have $W = \rho(g)\rho(g^{-1})(W) \subset \rho(g)W$, this implies that in fact $\rho(g)(W) = W$).

If W is an invariant subspace of V, then for every $g \in G$ the restriction of $\rho(g) : V \to V$ yields a map $\rho|_W(g) : W \to W$. This representation is called $\rho|_W$, which is called a "subrepresentation" of ρ.

Definition. An irreducible representation $\rho : G \to GL(V)$ is a representation with no nonzero proper subrepresentations; that is, there is no nonzero proper subspace W of V stable under V.

Example. Any 1-dimensional representation is irreducible. You showed on HW that if G is a finitely generated abelian group, then any irreducible representation of G is 1-dimensional.

The same is true for any finitely generated abelian group G. To see this, note that since G is finitely generated, there is a surjection $F \twoheadrightarrow G$ for some finitely generated free abelian F. Then every representation $\rho : G \to GL(V)$ lifts to a representation $\tilde{\rho} : F \to GL(V)$, and ρ is irreducible if and only if the same is true of $\tilde{\rho}$.

Example. We'll see later that, on the other hand, if G is non-abelian, there is always an irreducible representation of G of degree > 1. For instance S_3 has a 2-dimensional representation $\rho : S_3 \to GL_2(\mathbb{C})$ (in fact, with image in $GL_2(\mathbb{R})$) given by identifying S_3 with the group of symmetries of a triangle in the plane. It's easily checked that this has no nonzero invariant subspaces.

Today we'll show:

Theorem 24.1. *Every finite-dimensional representation $\rho : G \to GL(V)$, of a finite group G over a field of characteristic 0, can be written as a direct sum of irreducible subrepresentations.*

The crucial step in this proof will be the following lemma:

Lemma 24.2. *Suppose that G is finite, and that k has characteristic 0. If $\rho : G \to GL(V)$ is a representation, and W is an invariant subspace of V, then there exists W' also invariant such that $W \oplus W' = V$.*

Before proving the lemma, let's prove the theorem using the lemma:

Proof. We induct on the dimension of V.

If V is irreducible, we're done. Otherwise, let W be a nonzero invariant subspace of V. Write $V = W \oplus W'$ where W' is also invariant. Then $\rho = \rho|_W \oplus \rho|_{W'}$, and by the induction hypothesis both restrictions are direct sums of irreducibles, so the same is true of ρ. □

Now we prove the lemma.

Proof. There are two different methods of proving this.

One, which works only for $k = \mathbb{C}$ we'll only sketch (it can be found in Fulto-Harris). In this method, one construct a positive-definite G-invariant (hermitian) inner product on W and let $W' = W^{\perp}$ with respect to this inner product.

We'll do a different one, which is the one given in both Serre and Dummit & Foote. I'm going to try to present it here with a bit more context and motivation, but you can also look at those books for a more concise version.

Let's consider the set of all $W' \subset V$ subspaces such that $W' \oplus W = V$. These are called "complementary subspaces". We need to find W' such that $W' = \rho(g)W'$ for all $g \in G$.

We're going to use an "averaging trick" here, as we've done before. The issue is we don't have a way of averaging subspaces. So instead, we will biject these subspaces with something that we can average.

The way we do this is will be to give a bijection between subspaces W' with complementary to V and maps $\pi : V \to W$ which are projections onto W in the following sense: $\pi(V) \subset W$ and $\pi|_W = \mathrm{id}_W$. An equivalent statement of this condition on π, is that $\pi \in \mathrm{Hom}_k(V, W)$ is sent by the restriction map $\mathrm{Hom}_k(V, W) \to \mathrm{Hom}(W, W)$ to the identity map id_W.

The bijection is the following: if $V = W \oplus W'$ the corresponding π sends $w \oplus w'$ to $w \oplus 0$. In the other direction, W' can be recovered as $\ker \pi$.

It's easy to check that this is a bijection, and furthermore that if $W \leftrightarrow \pi$, then for any $g \in \mathbb{G}$,

$$\rho(g)W \leftrightarrow \rho(g)|_W \circ \pi \circ \rho(g)^{-1} = \mathrm{Hom}_k(\rho, \rho_|W)(g)(\pi).$$

(In the last equality, we're using the representation $\mathrm{Hom}_k(\rho, \rho_W) : G \to \mathrm{GL}(\mathrm{Hom}_k(V, W))$ defined at the end of Wednesday's class.)

This means that we've reduced our problem to that of finding an element of $\mathrm{Hom}_k(V, W)$ which is invariant under the representation $\mathrm{Hom}_k(\rho, \rho_W)$, and which restricts to $\mathrm{id} \in \mathrm{Hom}_k(W, W)$. Now we can just average like we would for any other representation!

That is, we have a diagram

$$\begin{array}{ccc}
\mathrm{Hom}_k(V, W) & \xrightarrow{\mathrm{res}} & \mathrm{Hom}_k(W, W) \\
\downarrow & & \downarrow \\
\mathrm{Hom}_k(V, W)^G & \xrightarrow{\mathrm{res}} & \mathrm{Hom}_k(W, W)^G \\
\| & & \| \\
\mathrm{Hom}_G(V, W) & \xrightarrow{\mathrm{res}} & \mathrm{Hom}_G(W, W)\,\mathrm{id}
\end{array}$$

(In this diagram we're committing a slight abuse of notation by using $\mathrm{Hom}_k(V, W)^G$ to mean $\mathrm{Hom}_k(V, W)^{\mathrm{Hom}_k(\rho, \rho|_W)(G)}$.)

Here the vertical arrow in the averaging map: it sends an element $\phi \in \mathrm{Hom}_k(V, W)$ to

$$\phi_0 = \frac{1}{|G|} \sum_{g \in G} \mathrm{Hom}_k(\rho, \rho|_W)(g)(\phi) = \frac{1}{|G|} \sum_{g \in G} \rho|_W(g) \circ \phi \circ \rho(g)^{-1}.$$

It's also easy to see that this diagram commutes. Hence if we start with any projection

map $\pi \in \mathrm{Hom}_k(V, W)$ which restricts to the identity on $\mathrm{Hom}_k(W, W)$, the averaged map

$$\pi_0 = \frac{1}{|G|} \sum_{g \in G} \rho|_W(g) \circ \phi \circ \rho(g)^{-1}$$

restricts on W to $\frac{1}{|G|} \sum_{g \in G} \rho|_W(g) \circ \mathrm{id}_W \circ \rho(g)|_W^{-1} = \mathrm{id}_W$. Hence π_0 is a projection map, which is invariant under the action of $\mathrm{Hom}_k(\rho, \rho|_W)(G)$, and by the argument above this means that $\ker \pi_0$ is an invariant subspace of V complementary to W, as desired. □

Corollary 24.3. *Every finite-dimensional representation of a finite group G can be written as a direct sum of irreducible subrepresentations.*

(In fact, the "finite-dimensional" condition here can be removed.)

Last time we showed

Corollary 24.4. *Every finite-dimensional representation of a finite group G (over a field of characteristic 0) can be written as a direct sum of irreducible subrepresentations.*

Last time we showed: if V is any finite-dimensional representation of a finite group G (this is a shorthand, by minor abuse of notation, for saying that $\rho : G \to \mathrm{GL}(V)$ is a representation), V can be written as a direct sum of irreducible representations

$$V = \oplus_i W_i.$$

The way we proved it did not guarantee any sort of uniqueness about the irreducible decomposition whatsoever; we had to make lots of choices.

Today we'll ask the following question: if W is any irreducible representation, can we tell from V if W occurs as one of the summands in the irreducible decomposition of V? And if so, how many times? A priori these answers might depend on the choice of irreducible decomposition, but we'll show that they don't.

We'll do this first in the case when W is the trivial representation (that is, the one-dimensional representation with trivial action).

Suppose that $V = \bigoplus_i W_i$. If any W_i is trivial, then $V^G \supset W_i$ is nontrivial.

For the converse, note that $V^G = \bigoplus_i W_i^G$, and $W_i^G = 0$ if W_i is not trivial. So $V^G = \bigoplus_{W_i \text{trivial}} W_i$. This gives us a sharper result: the number of copies of the trivial representation in the irreducible decomposition of V is equal to $\dim V^G$.

If we want to compute V^G explicitly, we can do so using the following lemma, which is another reformulation of the averaging trick:

Lemma 24.5. *If G is finite and $\rho : G \to \mathrm{GL}_V$ is any representation, then $\frac{1}{|G|} \sum_{g \in G} \rho_g \in \mathrm{End}(V)$ is a projection map of V onto V^G.*

Proof. Exercise. □

Now, we'll do the same thing for general W. This will use the same strategy, but instead of using V^G, we'll use $\mathrm{Hom}_G(W, V)$ (which we know from a previous lecture is also equal to $\mathrm{Hom}_k(W, V)^G$).

To do this, we'll first show some facts about Hom_G. The main important fact is:

Lemma 24.6 (Schur's Lemma)**.** *If V_1 and V_2 are distinct (non-isomorphic) irreducible representations of G, $\mathrm{Hom}_G(V_1, V_2) = 0$. If k is algebraically closed, and V_1 is irreducible and finite-dimensional then, $\mathrm{Hom}_G(V_1, V_1) \cong k$. (Here the isomorphism is given in the reverse direction by sending $a \in k$ to the map $a \cdot \mathrm{id}_{V_1} \in \mathrm{Hom}_G(V_1, V_1)$, which acts as scaling by a: $a \cdot \mathrm{id}_{V_1}(v_1) = av_1$.)*

Proof. Let $\rho_1 : G \to \mathrm{GL}(V_1)$ and $\rho_2 : G \to \mathrm{GL}(V_2)$ be the maps making V_1 and V_2 into representations.

We must show that if $f \in \mathrm{Hom}_G(V_1, V_2)$ then $f = 0$. Our condition means that $f \circ \rho_1(g) = \rho_2(g) \circ f$ for all $g \in G$.

Consider $\ker f \subset V_1$. We claim that this is an invariant subspace. Indeed, $f(v_1) = 0$ implies $f(\rho_1(g)(v_1)) = \rho_2(f(v_1)) = 0$ for all $g \in G$. Since V_1 is irreducible, we must have either $\ker f = 0$ or $\ker f = V_1$; the second case implies $f = 0$.

Likewise, we can show that $\mathrm{Im}\, f \subset V_2$ is an invariant subspace, so either $\mathrm{Im}\, f = 0$ or $\mathrm{Im}\, f = V_2$. In the first case, $f = 0$.

Hence the only possible way to have $f \neq 0$ is to have $\ker f = 0$ and $\mathrm{Im}\, f = V_2$. But that means that f is both injective and surjective, so is an isomorphism $V_1 \cong V_2$, contradicting the assumption that V_1 is not isomorphic to V_2.

For the second part: suppose $f \in \mathrm{Hom}_G(V_1, V_1)$. Then because k is algebraically closed and V_1 is finite-dimensional, we know that f must have an eigenvector. That is, there exists $v \in V_1$ such that $f(v) = av$ for some $v \in V$. Then consider the linear transformation $f' = f - a \cdot \mathrm{id}_{V_1} \in \mathrm{Hom}_G(V_1, V_1)$. By construction, $\ker f'$ contains v_1, so is not empty. By the argument above, this means that $\ker f' = V_1$ and so $f' = 0$. Hence $f = a \cdot \mathrm{id}_{V_1}$. □

(Question asked in class: is it possible for a representation that is irreducible over a non-algebraically closed field to become reducible after extending scalars? Yes, for instance, the representation $\rho : C_n \to \mathrm{GL}_2(\mathbb{R})$ which embeds the cyclic group C_n as a group of rotations is irreducible for $n > 2$ since no line in $\mathbb{R}^2$ is fixed under the group of rotations. However, in $\mathbb{C}^2$ there is a common eigenvector (in fact, two such up to scaling) for all the rotations. One can also check that this representation gives a counterexample to Schur's lemma over $\mathbb{R}$: in fact $\mathrm{Hom}_{C_n}(\mathbb{R}^2, \mathbb{R}^2)$ is equal to the group $\mathrm{SO}_2(\mathbb{R})$ of rotations of $\mathbb{R}^2$.)

Now, we go back to the original situation, where $V = \oplus_i W_i$ is an irreducible decomposition, and W is any irreducible representation. Then

$$\mathrm{Hom}_k(W, V) = \mathrm{Hom}_k(W, \bigoplus_i W_i) = \bigoplus_i \mathrm{Hom}_k(W, W_i)$$

and Schur's lemma tells us that $\mathrm{Hom}_k(W, W_i) \cong k$ if $W \cong W_i$ and 0 otherwise. Taking dimensions, we conclude that $\dim(\mathrm{Hom}_k(W, V))$ is equal to the number of W_i that equal W.

Hence we've shown

Theorem 24.7. *Let V be a finite-dimensional representation of a finite group G over an algebraically closed field k of characteristic 0. The irreducible representations that occur in any irreducible decomposition of V, and their multiplicities, do not depend on the choice of irreducible decomposition.*

25 Characters

Now we're going to introduce the most powerful tool for working with representations. Let's now fix our base field to be $\mathbb{C}$. (Since we were just assuming algebraically closed of characteristic 0, this shouldn't be too bad.)

Definition. If $\rho : G \to GL(V)$ is a finite-dimensional representation of G, then the character χ_ρ is the $\mathbb{C}$-valued function on G defined by $\chi_\rho(g) = \mathrm{tr}(\rho(g))$. (That is, this is the trace of $\rho(g)$ as an endomorphism of the finite-dimensional vector space V.)

(Notation note: starting now I'll be using ρ_g instead of $\rho(g)$ for brevity. In this notation, the definition above is $\chi_\rho(g) = \mathrm{Tr}(\rho_g)$.)

Proposition 25.1. *With ρ, V, χ_ρ as above:*

a) $\chi_\rho(\mathrm{id}_G) = \dim V$.

b) $\chi_\rho(hgh^{-1}) = \chi_\rho(g)$. *(that is, χ_ρ is constant on conjugacy classes – we call functions with this property "class functions")*

c) $\chi_\rho(g^{-1}) = \overline{\chi_\rho(g)}$ *(if G is finite)*

Proof. Part a) is true because $\chi_\rho(\mathrm{id}_G) = \mathrm{tr}(\rho_{\mathrm{id}_G}) = \mathrm{tr}(\mathrm{id}_V) = \dim V$.

Part b) follows from the linear algebra identity $\mathrm{tr}(\rho_h \rho_g \rho_h^{-1}) = \mathrm{tr}(\rho_g)$.

For part c): first, since g has finite order, so does ρ_g, and hence all eigenvectors λ of ρ_g satisfy $\lambda^n = 1$, which implies $|\lambda| = 1$ and $\lambda^{-1} = \overline{\lambda}$.

Then

$$\chi_\rho(g^{-1}) = \mathrm{tr}(\rho_{g^{-1}}) = \mathrm{tr}(\rho_g^{-1}) = \sum_{\lambda \text{ eigenvalue of } \rho_g} \lambda^{-1} = \sum_{\lambda \text{ eigenvalue of } \rho_g} \overline{\lambda} = \overline{\mathrm{tr}\,\rho_g} = \overline{\chi_\rho(g)}$$

□

Proposition 25.2. *If ρ_1 and ρ_2 are finite-dimensional representations of G, then $\chi_{\rho_1 \oplus \rho_2} = \chi_{\rho_1} + \chi_{\rho_2}$, $\chi_{\rho_1 \otimes \rho_2} = \chi(\rho_1)\chi(\rho_2)$, $\chi_{\rho_1^*} = \overline{\chi_{\rho_1}}$, and $\chi_{\mathrm{Hom}(\rho_1, \rho_2)} = \overline{\chi_{\rho_1}}\chi_{\rho_2}$.*

Proof. Exercise (will probably be on HW). □

(As before, assume throughout that the base field is $\mathbb{C}$, G is finite, and all representations are finite-dimensional.)

Last time, we showed that for any representation V with irreducible decomposition $V = \oplus_i W_i$, and for any irreducible representation W, the number of W_i isomorphic to W could be computed as $\dim \mathrm{Hom}_G(V, W)$ and did not depend on the choice of decomposition. Now we'll show how to find this number computing directly with characters.

As before, we'll do this first in the case where W is trivial, so the number of W_i that are trivial is equal to $\dim V^G$.

Theorem 25.3. *If $\rho : G \to \mathrm{GL}(V)$ is a representation of G, and $\chi = \chi_\rho$ then $\frac{1}{|G|}\sum_{g\in G}\chi(g) = \dim V^G$.*

Proof. We use the fact stated last time, that $r = \frac{1}{|G|}\sum_{g\in G}\rho_g$ is a projection map of V onto $\dim V^G$. Now, the left hand side of our equation is just $\operatorname{tr} r$.

To compute $\operatorname{tr} r$, take a basis of $V \cong V^G \oplus \ker r$ where the first $\dim V^G$ elements form a basis of V^G, and the rest form a basis of $\ker r$. In this basis, r is diagonal with the first $\dim V^G$ entries equal to 1, and all the rest 0, so $\operatorname{tr} r = \dim V^G$. □

Theorem 25.4. *If $\rho_1 : G \to \mathrm{GL}(V_1)$ and $\rho_2 : G \to \mathrm{GL}(V_2)$ are representations with characters χ_1 and χ_2, then $\frac{1}{|G|}\sum_{g\in G}\overline{\chi_1(g)}\chi_2(g) = \dim(\mathrm{Hom}_G(V_1, V_2))$.*

Proof. Apply the previous theorem to $\mathrm{Hom}_{\mathbb{C}}(\rho_1, \rho_2)$, which has character $\overline{\chi_1(g)}\chi_2(g)$ as stated last time. Then

$$\dim(\mathrm{Hom}_G(V_1, V_2)) = \dim(\mathrm{Hom}_{\mathbb{C}}(V_1, V_2)^G) = \sum_{g\in G}\overline{\chi_1(g)}\chi_2(g)$$

as desired. □

Define a hermitian inner product on the vector space of $\mathbb{C}$-valued functions on G by

$$(\alpha, \beta) = \frac{1}{|G|}\sum_{g\in G}\langle\overline{\alpha(g)}\beta(g)\rangle$$

(Note that some references, e.g. Serre, put the complex conjugate on the β. It doesn't make a big difference as usually we are evaluating this inner product when it is real-valued.)

Then the previous theorem says that $\dim \mathrm{Hom}_G(V_1, V_2) = (\chi_1, \chi_2)$.

Corollary 25.5. *If $V = \oplus_i W_i$ is an irreducible decomposition, and W is any irreducible representation of G, the number of W_i that equal W is (χ_W, χ_V).*

Proof. (This follows from the result we stated last time, that the number of W_i that equal W is $\dim \mathrm{Hom}_G(W, V)$. □

In the case where V is itself irreducible, this tells us that (χ_W, χ_V) is 1 if $V = W$ and 0 otherwise. That is,

Corollary 25.6. *The set of characters χ_V as V ranges over the irreducible representations of G is orthonormal with respect to the hermitian inner product $(\alpha, \beta) = \frac{1}{|G|}\sum_{g\in G}\langle\overline{\alpha(g)}\beta(g)\rangle$ on the space of $\mathbb{C}$-valued functions on G.*

In particular, if a representation V is irreducible, then $(\chi_V, \chi_V) = 1$. In fact, the converse is true, yielding an easy test for irreducibility of a representation.

Theorem 25.7. *A representation V with character χ is irreducible if and only if $(\chi, \chi) = 1$.*

Proof. We have one direction already. We'll prove the other direction by the converse.

If V is not irreducible let $V = \bigoplus_{i=1}^{n} V_i$, $n \geq 2$. Then $\chi_V = \sum_{i=1}^{n} \chi_{V_i}$ and

$$(\chi, \chi) = \sum_{i,j=1}^{n} (\chi_{V_i}, \chi_{V_j}) \geq \sum_{i=1}^{n} (\chi_{V_i}, \chi_{V_i}) = n > 1.$$

□

Example. Suppose that $\rho : G \to GL(V)$ is the regular representation of G. Then it's easy to check that $\chi(g) = \mathrm{tr}(\rho_g) = |G|$ if $g = 1$ and 0 otherwise. If W is any irreducible representation of G with character ϕ, then W appears $(\chi|\phi) = \frac{1}{|G|}|G|\overline{\phi(\mathrm{id}_G)} = \dim W$ times in V.

Hence $V = \bigoplus_{W \text{ irreducible}} W^{\dim W}$. Taking the dimension of both sides, we get the following nice corollary: $|G| = \sum_{W \text{ irreducible}} (\dim W)^2$.

Theorem 25.8. $|G| = \sum_{W \textit{ irreducible}} (\dim W)^2$.

Example. Let's write down the characters of all the irreducible representations of S_3. Since the characters are class functions, we only need to write down their values on each conjugacy class of S_3. We'll write this down as a character table (the numbers in the top row are the sizes of the conjugacy classes, useful for computing the inner products):

	1	3	2
S_3	id	(12)	(123)
χ_1	1	1	1
χ_2	1	−1	1
χ_3	2	0	??

Here χ_1 is the character of the trivial representation $G \to GL(V_1)$. The character χ_2 comes from the other one-dimensional representation corresponding to the group homomorphism $\sigma \mapsto \mathrm{sgn}(\sigma)$ from $G \to \mathbb{C}$.

As for χ_3, we've seen it before an as example: it comes from an embedding of $\rho_3 : S_3 \to \mathrm{GL}_{V_2} \cong \mathrm{GL}_2(\mathbb{C})$ which identifies S_3 with the symmetries of an equilateral triangle in the plane. Then $\chi_3(\mathrm{id}) = \operatorname{tr} \mathrm{id}_{\mathbb{C}^2} = 2$. We computed $\chi_3((12))$ in class by observing that (12) is mapped to a reflection through a line, which has eigenvalues $+1$ and -1, hence trace 0. We could have computed $\chi_3((123))$ likewise (as the trace of a 120° rotation) but instead applied orthogonality of characters. Let $\chi_3((123)) = \alpha$. Then

$$0 = (\chi_1, \chi_3) = 1 \cdot 1 \cdot 2 + 2 \cdot 1 \cdot 0 + 2 \cdot 1 \cdot \alpha = 2 + 2\alpha$$

so $\alpha = -1$. Hence the table in full is:

	1	3	2
S_3	id	(12)	(123)
χ_1	1	1	1
χ_2	1	-1	1
χ_3	2	0	-1

26 The center of the group algebra

We saw before that irreducible characters of G are orthonormal with respect to the inner product $(\cdot, \cdot)$ on the space of functions $G \to \mathbb{C}$. As well, they lie in the subspace of class functions (functions $G \to \mathbb{C}$ which are constant on conjugacy classes). We'll show that they are actually a basis for the space of class functions.

In order to do this, we first need to say more about the group algebra $\mathbb{C}[G]$.

Definition. The *center* of a non-commutative algebra A is the set

$$Z(A) = \{a \in A \mid ab = ba \text{ for all } b \in A\}$$

Lemma 26.1. *The center $Z(\mathbb{C}[G])$ is the set of all elements of the form $\sum_{g \in G} a_g g$ for $g \in G$ such that the function $g \mapsto a_g$ is a class function.*

Proof. By linearity, to check that $a = \sum_{g \in G} a_g \in \mathbb{C}[G]$ commutes with every element $b \in \mathbb{C}[G]$, we only need to check this when b runs through the basis $\{1 \cdot h\}_{h \in G}$. Then

$$a(1 \cdot h) = \sum_{g \in G} a_g (gh) = \sum_{g \in G} a_{gh^{-1}} g$$
$$(1 \cdot h)a = \sum_{g \in G} a_g (hg) = \sum_{g \in G} a_{h^{-1}g} g.$$

Comparing coefficients, we have that $a_{gh^{-1}} = a_{h^{-1}g}$ for all $g, h \in G$. Doing another change of variables (letting gh^{-1} be our new g), we get that $a_g = a_{h^{-1}gh}$ for all $g, h \in G$. This is precisely the statement that a_g is a class function. $\square$

Hence our vector space $\mathbb{C}_{\text{class}}(G)$ can naturally be identified with the $Z(\mathbb{C}[G])$.

We'll make use of the center of $\mathbb{C}[G]$ using the following fact: If M is a module over a non-commutative ring A, and $a \in Z(A)$ lies in the center, then the multiplication map $\ell_a : M \to M$ is an A-algebra homomorphism.

(To show this: $\ell_a(x+y) = \ell_a(x) + \ell_a(y)$ is the distributive law, while $\ell_a(bx) = a(bx) = (ab)x = (ba)x = b(\ell_a(x))$ because a lies in the center.)

In the case we're intersted in, we know that a $\mathbb{C}[G]$-module is the same as a representation $\rho : G \to GL(V)$. For any such representation, then, the map ℓ_a is a homomorphism of representations, that is, an element of $\text{End}_G(V) = \text{Hom}_G(V, V)$. This puts a big constraint on what ℓ_a can be; for instance, if V is irreducible we know by Schur's lemma that $\text{Hom}_G(V, V) \cong \mathbb{C}$.

In fact, we can saw more precisely how ℓ_a acts:

Theorem 26.2. *Let f be a class function on G, and let $a = \sum_{g \in G} f(g)g$. Let $\rho : G \to GL(V)$ be a representation, with irreducible decomposition $V = \sum_i W_i$, where W_i has character χ_i. Then the transformation $\ell_a = \sum_{g \in G} f(g)\rho_g \in \text{End}(V)$ maps each W_i to itself, and acts on $\text{End}(W_i)$ as scaling by $\frac{|G|}{\dim W_i}(\overline{\chi_i}, f)$.*

Proof. First of all, because W_i is an invariant subspace, it's clear that $\ell_a W_i \subset W_i$. We can now replace V by W_i and reduce to case where $V = W_i$ is irreducible.

Now, by Schur's lemma, we know that ℓ_a must be scaling by some $\alpha \in \mathbb{C}$. To find out what this is, we compute $\text{tr}(\ell_a)$. On the one hand this is $\dim W_i \alpha$, but on the other it is

$$\text{tr}(\sum_{g \in G} f(g)) = \sum_{g \in g} f(g) \,\text{tr}(\rho_g) = \sum_{g \in G} \chi(g) f(g) = |G|(\overline{\chi_i}, f).$$

Equating the two gives the desired result. □

27 Proof that the irreducible characters span the space of class functions

Now we can prove our main goal of today.

Theorem 27.1. *The characters χ_i of the irreducible representations W_i of G are an orthonormal basis for the space of class functions on G.*

Proof. We already know they are orthonormal. To show that they form a basis, it's enough to show that if f is orthogonal to all χ_i, then $f = 0$. So assume that $(\chi_i, f) = 0$ for all i. Let $a = \sum_{g \in G} \overline{f}(g)g$.

Now we use the previous theorem: it tells us that any representation $\rho : G \to GL_V$, with irreducible decomposition $V = \oplus_i W_i$, the element $\ell_a \in \text{End}_G(V)$ acts on W_i as multiplication by $|G|(\overline{\chi_i}, \overline{f}) = 0$. Hence $\ell_a = 0$.

Apply this in the case where ρ be the regular representation. Then $0 = \ell_a(e_{id}) = \sum_{g\in G} \overline{f(g)} e_g$, so all $f(g)$ must be 0, hence $f = 0$. □

As a corollary we get

Corollary 27.2. *The number of irreducible representations of* G *is equal to the number of conjugacy classes of* G.

28 What's coming next

Next time, we'll show that for any irreducible representation V of G, the dimension of V divides the order of G.

This means that you can determine a substantial amount about the dimensions of the irreducible representations just from knowing $|G|$. Indeed, if $d_1, \dots d_m$ are the dimensions of the irreducible representations $V_1, \dots, V_m$ of G, then each d_i divides $|G|$, and we know from last time that $\sum_i d_i^2 = |G|$. Also, the trivial representation will always appear as some V_i, so WLOG $d_1 = 1$. This puts some fairly strong restrictions on the positive integers d_i. (Even more if you know the number m of conjugacy classes.)

But before we do this, we'll do a bit more on the structure of the group algebra $\mathbb{C}[G]$.

For any representation $\rho : G \to GL(V)$, we get a homomorphism (of noncommutative $\mathbb{C}$-algebras) $\mathbb{C}[G] \to \mathrm{End}_{\mathbb{C}}(V)$, which sends $a = \sum_g a_g g$ to $\ell_a = \sum_g a_g \rho(g) \in \mathrm{End}_{\mathbb{C}}(V)$. We will also write this as $r\tilde{ho}(a) = \sum_g a_g \rho(g)$ to emphasize that G is acting by the representation ρ.

In particular, let $\{\rho_i : G \to GL(V_i)\}$ be the irreducible representations of G. Then there is a map $\mathbb{C}[G] \to \prod_i \mathrm{End}_{\mathbb{C}}(V_i)$.

Next time we'll show

We'll start by proving the theorem we stated last time.

Theorem 28.1. *The map* $\prod_i \tilde{\rho_i}$ *is an isomorphism of non-commutative* $\mathbb{C}$*-algebras.*

Proof. First of all, $\prod_i \tilde{\rho_i}$ is clearly a homomorphism.

We first show that $\prod_i \tilde{\rho_i}$ is injective. For this, suppose that $a \in \ker \tilde{\rho}$. Then we have $av_i = 0$ for any $a \in A$ and $v_i \in V_i$.

Now we show that for any representation V of G, and any $v \in V$, $av = 0$. To do this, note that V of G has an irreducible representation $V = \oplus_j W_j$, where each W_j is isomorphic to some V_i. Since multiplication by a is the 0 map on each W_j, it is also the 0 map on V.

Now take V to be the regular representation; As a $\mathbb{C}[G]$-module, V can be identified with $\mathbb{C}[G]$. Let $v = 1 \cdot id_G \in \mathbb{C}[G]$. Then we have this means that $0 = a \cdot (1 \cdot id_G) = a$. This shows injectivity.

To show surjectivity, compare the dimensions of both sides as complex vector spaces. We know that $\dim \mathbb{C}[G] = |G|$, but $\dim(\prod_i \mathrm{End}(V_i)) = \sum_i |\dim V_i|^2$ is also equal to $|G|$ by

the result we proved last week. Hence $\prod_i \tilde{\rho}_i$ is an injective map between vector spaces of the same dimension, hence must be an isomorphism. □

In fact, not only is $\prod_i \tilde{\rho}_i$ a homomorphism, but we can explicitly write down the inverse map.

Proposition 28.2 (Inverse Fourier Transform)**.** *Let* (u_i) *be an element of* $\prod_i \mathrm{End}_{\mathbb{C}}(V_i)$*. The element* $u = \sum_{g\in G} u_g g$ *such that* $\tilde{\rho}_i(u) = u_i$ *has coefficients given by*

$$u_g = \frac{1}{|G|}\sum_i (\dim V_i)\,\mathrm{tr}(\rho_i(g^{-1})u_i). \tag{6}$$

(The reason why this is called the inverse Fourier transform is that in the case when $G \cong \mathbb{Z}/n\mathbb{Z}$, this specializes to the discrete inverse Fourier transform on $\mathbb{Z}/n\mathbb{Z}$. In this case, all irreducible representations of $\mathbb{C}$ are one-dimensional, so $\prod_i \mathrm{End}_{\mathbb{C}}(V_i) \cong \mathbb{C}^n$. The map $\mathbb{C}[G] \to \prod_i \mathrm{End}_{\mathbb{C}}(V_i)$ can then be viewed as a discrete Fourier transform.)

First, a lemma:

Theorem 28.3. *For* $g \in G$, $\frac{1}{|G|}\sum_i n_i\chi_i(g) = 0$ *if* $g \neq \mathrm{id}_G$ *and* $= 1$ *if* $g = \mathrm{id}_G$.

Proof. This is the same as the trace of g acting on the regular representation of G. □

Now we prove the proposition.

Proof. Note that both sides of (6) are linear functions of u. By linearity, we just need to check this when u runs through the basis $\{1 \cdot g\}_{g\in G}$.

So suppose $u = g_0$ for some $g_0 \in G$.

Then the left side is 1 if $g = g_0$ and 0 otherwise. The right hand side is

$$\frac{1}{|G|}\sum_i \dim V_i\,\mathrm{tr}(\rho_i(g^{-1}r\tilde{ho}_i(1 \cdot g_0))) = \frac{1}{|G|}\sum_i \dim V_i \chi_{V_i}(g^{-1}g_0)$$

which, by the lemma is 1 if $g = g_0$ and 0 otherwise. □

29 Showing that the degrees of the irreducible representations divide $|G|$

Let $a = \sum_{g\in G} a_g g \in Z(\mathbb{C}[G])$. Since $Z(\mathbb{C}[G])$ is a commutative ring containing $\mathbb{Z} \subset \mathbb{C}$, it's meaningful to ask whether a is integral over $\mathbb{Z}$.

Theorem 29.1. *If all* a_g *are algebraic integers, then* a_g *is integral over* $\mathbb{Z}$*.*

Proof. Let $c_1, \ldots, c_m$ be the conjugacy classes of G. Define elements $e_i = \sum_{g \in c_i} g \in Z(\mathbb{C}[G])$.

Then we claim that the sub-$\mathbb{Z}$-module $\mathbb{Z}e_1 \oplus \mathbb{Z}e_2 \oplus \cdots \oplus \mathbb{Z}e_m$ is a subring. It's easy to see that it contains 1. To check that it's closed under multiplication, note that $e_{i_1} e_{i_2} = \sum_g a_g e_g$ where all a_g are integers, so can be written as an integer linear combination of the e_i.

Hence $\mathbb{Z}e_1 \oplus \mathbb{Z}e_2 \oplus \cdots \oplus \mathbb{Z}e_m \subset \mathbb{C}[G]$ is a ring and is finite over $\mathbb{Z}$, hence integral over $\mathbb{Z}$. This shows that each e_i is integral over $\mathbb{Z}$.

Now we need to show that $\alpha \in \sum_{g \in G} a_g g \in Z(\mathbb{C}[G])$ when all a_g are algebraic integers. But $\alpha = \sum_i a_{c_i} e_i$ where we define $a_{c_i} = a_g$ for any $g \in c_i$ (because $g \mapsto a_g$ is a class function this doesn't depend on the choice of $g \in c_i$). Since all e_i are integral over $\mathbb{Z}$, as are all a_{c_i}, α is integral over $\mathbb{Z}$. □

(Question asked in class: is this an if and only if? I wasn't able to answer this in class, but the answer is no. For instance, if $G = C_3$ is cyclic of order 3, the element $\alpha = \frac{1}{3}(1 + t + t^2)$ is a root of $x^2 - x = 0$, but α does not have integer coefficients.)

Proposition 29.2. *Let $f \in \mathbb{C}_{class}(G)$ take on algebraic integer values. Then for any irreducible representation $\rho : G \to V$ with character χ, $\sum_{g \in G} f(g)\chi(g)$ is a multiple of $\dim V$ (in the ring of algebraic integers)*

Proof. The ring homomorphism $\tilde{\rho} : \mathbb{C}[G] \to \mathrm{End}_{\mathbb{C}}(V)$ restricts to a homomorphism $Z(\mathbb{C}[G]) \to \mathrm{End}_G(V) \cong \mathbb{C}$. If $\alpha \in \mathbb{Z}(\mathbb{C}(G))$ is an integral over $\mathbb{Z}$, so is $\rho(\tilde{\alpha}) \in G$.

Now apply this with $\alpha = \sum_{g \in G} f(g)g$. We previously calculated that

$$\tilde{\rho}(\alpha) = \frac{1}{\dim V} \sum_{g \in G} f(g)\chi(g),$$

so $\frac{1}{\dim V} \sum_{g \in G} f(g)\chi(g)$ is an algebraic integer as desired. □

We're still working on showing: if V is an irreducible representation of G, then $\dim V \mid |G|$.

Last time we finished by showing

Proposition 29.3. *Let $f \in \mathbb{C}_{class}(G)$ take on algebraic integer values. Then for any irreducible representation $\rho : G \to V$ with character χ, $\sum_{g \in G} f(g)\chi(g)$ is a multiple of $\dim V$ (in the ring of algebraic integers)*

We'll apply this with $f(g) = \overline{\chi(g)}$. To do this, we must first note

Proposition 29.4. *If ρ is a finite-dimensional representation of a finite group G, with character χ, then $\chi(g)$ is an algebraic integer for all $g \in G$.*

Proof. We have that $\chi(g)$ is the sum of the eigenvalues of ρ_g. Since ρ_g has finite order in $GL(V)$, all eigenvalues are roots of unity, and so their sum is an algebraic integer. □

Corollary 29.5. *If V is an irreducible representation of G, then $\dim V$ divides $|G|$.*

Proof. Let $f(g) = \overline{\chi(g)}$. This takes on algebraic integer values, because χ does.

Hence

$$\dim(V) \mid \sum_{g \in G} \overline{\chi(g)}\chi(g) = |G|(\chi, \chi) = |G|.$$

□

In fact, we can amplify this result to get something stronger: if C is the center of G, then $\dim V$ divides the index $[G : C]$. (This is a case of what Terry Tao calls the "tensor power trick". The proof we give is originally due to John Tate.)

Definition. If G and H are groups and $\rho_1 : G \to GL(V_1)$, $\rho_2 : H \to GL(V_2)$, define the *external tensor product representation* $\rho_1 \otimes \rho_2 : G \times H$ to $GL(V_1 \otimes V_2)$ by $\rho_1 \otimes \rho_2(g, h) = \rho_1(g) \otimes \rho_2(h)$.

(This is sometimes written $\rho_1 \boxtimes \rho_2$, to avoid confusion with the previously defined tensor product of two representations of G. The concepts are related: if $\rho_1 : G \to GL(V_1)$ and $\rho_2 : G \to GL(V_2)$ are representations of the same group G, then $\rho_1 \boxtimes \rho_2 : G \times G \to GL(V_1 \otimes V_2)$ is related to $\rho_1 \otimes \rho_2 : G \to GL(V_1 \otimes V_2)$ by $\rho_1 \otimes \rho_2 = (\rho_1 \boxtimes \rho_2) \circ \delta$, where $\delta : G \to G \times G$ is the diagonal embedding $\delta(g) = (g, g)$.)

The following properties are easily checked: the character χ of $\rho_1 \otimes \rho_2$ is given by $\chi(g, h) = \chi_1(g)\chi_2(g)$. The inner product $(\chi, \chi) = (\chi_1, \chi_1)(\chi_2, \chi_2)$. Hence $\rho_1 \otimes \rho_2$ is also irreducible if and only if both ρ_1 and ρ_2 are irreducible.

This generalizes to taking the external tensor product of any number of representations.

Now we prove the stronger result.

Theorem 29.6. *If G is a group and $\rho : G \to GL(V)$ is any irreducible representation, then $\dim V$ divides $[G : C]$.*

Proof. For any positive integer n, construct the representation $V^{\otimes n}$ of G^n. This is an external tensor product of irreducible reps, so it's an irreducible representation of G^n.

We claim that the subgroup $H = \{(c_1, \ldots, c_n) \in C^n \mid c_1 \ldots c_n = 1\}$ acts trivially on $V^{\otimes n}$. Indeed, by Schur's lemma, any $c \in C$ acts on V by scaling by some $\lambda(c)$, so $(c_1, \ldots, c_n)$ acts as scaling by $\lambda(c_1)\lambda(c_2) \cdots \lambda(c_n) = \lambda(c_1 \cdots c_n) = 1$.

Hence $V^{\otimes n}$ is also a representation of G^n/H , and is still irreducible. Hence $\dim(V^{\otimes n}) = (\dim V)^n$ divides $|G^n/H| = |G|^n/|C|^{n-1}$. Hence

$$\left(\frac{|G|}{|C| \dim V}\right)^n \in \frac{1}{C}\mathbb{Z}$$

. Because this is true for all n, we must have $\frac{|G|}{|C| \dim V} \in \mathbb{Z}$, so $\dim V \| G|/|C| = [G : C]$. □

30 Induced representations

We'll be a way of going from representations of H to representations of G. But before we give a procedure for doing this, we'll give some of its properties

Let $\rho : G \to GL(V)$ be a representation of G. Let $W \subset V$ be a subspace that is H-invariant; let $\theta : H \to GL(W)$ be the corresponding representation of H. For every $g \in G$, we have a subspace $\rho_g(W) \subset V$; this only depends on the left coset gH. So if σ is any left coset of H in G, we can define $W_\sigma = \rho_g(W)$ for any $g \in \sigma$.

Definition. We say that ρ is induced by θ if $V = \oplus_{\sigma \in G/H} V_\sigma$.

As usual, assume all groups finite, and all representations are finite-dimensional and over $\mathbb{C}$.

Last time we stated this definition:

Let $\rho : G \to GL(V)$ be a representation of G. Let $W \subset V$ be a subspace that is H-invariant; let $\theta : H \to GL(W)$ be the corresponding representation of H. For every $g \in G$, we have a subspace $\rho_g(W) \subset V$; this only depends on the left coset gH. So if σ is any left coset of H in G, we can define $W_\sigma = \rho_g(W)$ for any $g \in \sigma$.

Definition. We say that ρ is induced by θ if $V = \oplus_{\sigma \in G/H} V_\sigma$.

Now we give some examples.

Example. $\rho : G \to GL(V)$ is the regular representation with basis $\{e_g\}_{g \in G}$, and $W = \text{span}(e_h)_{h \in H}$ is the regular representation of H. Then $W_\sigma = \text{span}(e_g)_{g \in \sigma}$, and $V = \oplus_{\sigma \in G/H} W_\sigma$.

Example. $\rho : G \to GL(V)$ is the permutation representation on left cosets of H, with basis $\{e_\sigma\}_{\sigma \in G/H}$, and $W = \text{span}(e_H)$, θ is the trivial representation of W. Then $W_\sigma = \text{span}(e_\sigma)$ and again $V = \oplus_{\sigma \in G/H} W_\sigma$

Example. $G = D_n$, $\rho : G \to GL(V)$ is the 2-dimensional representation given by embedding G into $GL_2(\mathbb{C})$ as the symmetry group of a regular n-gon, $H = C_n$. Here we may take $W = \text{span}(\begin{pmatrix} 1 \\ i \end{pmatrix})$. In this case, there are only two cosets, H and gH for any $g \notin H$. Clearly $W_H = W$, and to find W_{gH} we can choose g such that ρ_g is reflection through the x-axis, so $W_{gH} = \text{span}(\rho_g(\begin{pmatrix} 1 \\ i \end{pmatrix}) = \text{span}(\begin{pmatrix} 1 \\ -i \end{pmatrix})$. Clearly $V = W_H \oplus W_{gH}$.

Observations: if $\rho : G \to GL(V)$ is induced by $\theta : G \to GL(W)$, and W' is an H-invariant subspace of W, then $V' = \bigoplus_{\sigma \in G/H} W'_\sigma$ is G-invariant, and the representation V' of G is induced by the representation W' of H.

If V_1 is induced by W_1 and V_2 is induced by W_2, then $V_1 \oplus V_2$ is induced by $W_1 \oplus W_2$.

Using this, we can show that for any representation W of H there is some representation V of G which is induced by W.

First, we do this when W is irreducible. We know that the regular representation W_{reg} of H contains W as a summand in any irreducible decomposition. Hence we can choose an injection $W \hookrightarrow W_{\text{reg}}$ of H-representations and identify W with its image inside W_{reg}. Now, the regular representation V_{reg} of G is induced by W_{reg}, so by the first observation above, V_{reg} has a subspace V which is induced by W.

Now, let W be an arbitrary representation of H, and take an irreducible decomposition $W = \oplus_i W_i$. By the previous paragraph, there are representations V_i of G induced by W_i, and then by the second observation, $\oplus_i V_i$ is induced by $W = \oplus_i W_i$.

Although this works to show that V exists, it is not very canonical, in that it required taking a choice of embedding of each W_i into W_{reg}. A more canonical construction is given in your problem set.

However, we'll now show that the induced representation V of G is determined up to canonical isomorphism by the representation of W. To do that, we'll show it has the following universal property:

Theorem 30.1. *If $\rho : G \to \mathrm{GL}(V)$ is induced by $\theta : G \to \mathrm{GL}(W)$, then for any other representation $\rho' : G \to \mathrm{GL}(V')$ and any homomorphism $f : W \to V'$ of H-representations, there is a unique homomorphism $\tilde{f} : V \to V'$ of G-representations such that $\tilde{f}|_W = f$.*

Proof. We'll do uniqueness first, then existence:

Uniqueness: Since $V = \oplus_{\sigma \in G/H} W_\sigma$, to show that $\tilde{f}$ is uniquely determined, it's enough to show that $\tilde{f}|_{W_\sigma}$ is uniquely determined.

For any $\sigma \in G/H$, choose a coset representative $g \in \sigma$. Now, an arbitrary element of W_σ is of the form $\rho_g(w)$ for some $w \in W$. Because $\tilde{f}$ is a homomorphism of G-representations, we have

$$\tilde{f}(\rho_g(w)) = \rho'_g(\tilde{f}(w)) = \rho'_g(f(w))$$

since $\tilde{f}|_W = f$.

Hence the conditions imposed determine the values of $\tilde{f}|_{W_\sigma}$ for any $\sigma \in G/H$, hence determine $\tilde{f}$.

Existence: From the above, we get a formula for $\tilde{f}|_{W_\sigma}$ for each $\sigma \in G/H$, and so also for $\tilde{f}$. To check that this works we need to check two things: that the formula for $\tilde{f}|_{W_\sigma}$ does not depend on the choice of $g \in \sigma$, and that $\tilde{f} : V \to V'$ is indeed a homomorphism of G-representations.

□

Corollary 30.2. *If W is a representation of H, and V_1, V_2 are representations of G both induced by W, there is a unique isomomorphism $V_1 \cong V_2$ which restricts to the identity on W.*

Proof. This is a standard universal property argument. Let $i_1 : W \to V_1$ and $i_2 : W \to V_2$ be the inclusion maps. Then our universal property gives us unique maps $\tilde{i}_1 : V_2 \to V_1$ and $\tilde{i}_2 : V_1 \to V_2$ such that $\tilde{i}_1 \circ i_2 = i_1$ and $\tilde{i}_2 \circ i_1 = i_2$. Then we argue as in the usual universal property argument that $\tilde{i}_1$ and $\tilde{i}_2$ are inverses. □

Now a bit of notation.

Definition. If $H \subset G$, and W is a representation of H, we denote the representation induced by W (which we now know is determined up to unique isomorphism by $\mathrm{Ind}_H^G(V)$ or just $\mathrm{Ind}\, V$ if G and H are clear from context.

If $\rho : G \to \mathrm{GL}(V)$ is a representation of V, we use the notation $\mathrm{Res}_H^G V$ for the restricted homomorphism $\rho|_H : H \to \mathrm{GL}(V)$.

With this notation, we can restate our universal property as follows:

Proposition 30.3. *There is a natural identification*

$$\mathrm{Hom}_H(W, \mathrm{Res}\, V') \cong \mathrm{Hom}_G(\mathrm{Ind}\, W, V')$$

given by $f \mapsto \tilde{f}$ *and* $g|_W \leftarrow\!\shortmid g$.

Last time, we were in this situation: G and H are finite groups with $H \subset G$. We had a representation W of H, from while we constructed a representation $\mathrm{Ind}_H^G(W)$ of G. Also we had a representation V of G (last time we called this V') and we constructed a representation $\mathrm{Res}_H^G(V)$ of H by restriction.

Last time we showed

$$\mathrm{Hom}_H(W, \mathrm{Res}_H^G V) \cong \mathrm{Hom}_G(\mathrm{Ind}_H^G W, V)$$

By taking dimensions of both sides, we immediately get the following corollary:

Corollary 30.4 (Frobenius Reciprocity)**.**

$$(\chi_W, \chi_{\mathrm{Res}_H^G V})_H = (\chi_{\mathrm{Ind}_H^G W}, \chi_V)_G$$

(The subscripts mean that on the left we are taking the inner product in the space of functions on H*, and on the right we are taking inner products in the space of functions on* G*.)*

Corollary 30.5. *Suppose W and V are irreducible representations of H and G respectively. The the number of times that W occurs in $\mathrm{Res}_H^G V$ is equal to the number of times that V occurs in $\mathrm{Ind}_H^G W$.*

Proof. The first is $(\chi_W, \chi_{\mathrm{Res}\, V})_H$; the second is $(\chi_V, \chi_{\mathrm{Ind}_H^G W}) = (\chi_{\mathrm{Ind}_H^G W}, V)$. Hence this follows from the previous corollary. □

Next, we'll compute the character of an induced representation.

Proposition 30.6. *If $\rho : G \to \mathrm{GL}(V)$ is induced by $\theta : H \to \mathrm{GL}(W)$, and $\{g\_\sigma\}$ is a set of coset representatives,*

$$\chi_\rho(g) = \sum_{\substack{\sigma \in G/H \\ g\sigma = \sigma}} \chi_\theta(g_\sigma^{-1} g g_\sigma) = \frac{1}{|H|} \sum_{\substack{\sum_{g' \in G} \\ (g')^{-1} g g' \in H}} \chi_\theta((g')^{-1} g g')$$

Proof. We go back to our original definition of induced representation. Write $V = \oplus_\sigma W_\sigma$.

We need to find the trace of the matrix of ρ_g for $g \in G$. To do this, we choose a basis of V compatible with our direct sum decomposition $V = \oplus_\sigma W_\sigma$.

We claim that $\rho_g(W_\sigma) = W_{g\sigma}$. To check this, let$\sigma = g'H$. Then $g\sigma = (gg')H$, and $W_{g\sigma} = \rho_{gg'}(W) = \rho_g(\rho_{g'}(W)) = \rho_g(W_\sigma)$.

Hence the matrix of ρ_g, in the block decomposition corresponding to the subspace decomposition $V = \oplus_\sigma W_\sigma$only has nonzero entries in the blocks where the columns correspond to W_σ and rows correspond to $W_{g\sigma}$ for some $\sigma \in G/H$. Of those blocks, only the ones with $\sigma = g\sigma$ contribute to $\operatorname{tr} \rho_g$.

Hence

$$\chi_V(g) = \operatorname{tr} \rho_g = \sum_{\substack{\sigma \in G/H \\ g\sigma = \sigma}} \operatorname{tr} \rho_g|_{W_\sigma}.$$

Now to compute $\operatorname{tr} \rho_g|_{W_\sigma}$. For this, suppose $g' \in \sigma$, so $\sigma = g'H$, and note the following commutative diagram:

$$\begin{array}{ccc} W & \xrightarrow{(\rho_{g'})|_W} & W_\sigma \\ \left(\rho_{(g')^{-1}gg'}\right)|_W \downarrow & & \downarrow (\rho_g)|_{W_\sigma} \\ W & \xrightarrow{(\rho_{g'})|_W} & W_\sigma. \end{array}$$

Since $\rho_{g'}|_W : W \to W_\sigma$ is an isomorphism, we have

$$\operatorname{tr} \rho_g|_{W_\sigma} = \operatorname{tr} \left(\rho_{(g')^{-1}gg'}\right)|_W = \operatorname{tr} \theta_{(g')^{-1}gg'} = \chi_\theta((g')^{-1}gg') \tag{7}$$

by definition of θ.

We then obtain the first formula for $\chi_V(g)$ by summing (7) as g' runs over the a set of coset representatives for the cosets σ with $g\sigma = \sigma$.

To obtain the second formula, we first fix σ with $g\sigma = \sigma$, and average (7) over all $g' \in \sigma$. This yields

$$\operatorname{tr} \rho_g|_{W_\sigma} = \frac{1}{|H|} \sum_{g' \in \sigma} \chi_\theta((g')^{-1}gg').$$

We then get the second formula by summing this over all σ such that $g\sigma = \sigma$ (and noting that if $g' \in \sigma$, or equivalently, $\sigma = g'H$, the condition $g\sigma = \sigma$ is true if and only if $gg'H = g'H$, if and only if $(g')^{-1}gg'H = H$, if and only if $(g')^{-1}gg' \in H$).

□

One last application of induced representations to representation theory of finite groups. We don't have time to prove the following theorem, but you can find it in Serre:

Theorem 30.7. *Let G be a group, let A be a normal subgroup. Then if V is an irreducible representation of G, then either:*

there exists $A \subset H \subsetneq G$ and W an irreducible representation of H such that $V \cong \mathrm{Ind}_H^G(W)$

or

$\mathrm{Res}_A^G V$ *is isotypic. (This means that in the irreducible decomposition $\mathrm{Res}_A^G V \cong \oplus_i W_i$ of V, all the irreducible summands W_i are isomorphic.)*

This theorem has the following corollary, which we will prove.

Corollary 30.8. *Let G be a group with an abelian normal subgroup A. Then any irreducible representation V of G has $\dim V \mid [G : A]$.*

Remark. We've already proved this when A is the center of G, but this is a substantial stregthening. For instance, if $G = D_6$, we previously showed that $\dim V$ must divide $[G : C] = 6$. However, if we take $A = C_6$ we now have $\dim V \mid [G : A] = 2$. This is sharp, since you saw on HW that all representations of G have dimension 1 or 2.

Proof. Induct on the order of G. By the previous theorem, we have two cases.

Case 1: $V \cong \mathrm{Ind}_H^G W$ for some $H \subsetneq G$ with $H \supset A$. Since A is still normal in H, the induction hypothesis applies to H, and we must have $\dim W \mid [H : A]$. Then $\dim V = [G : H] \dim W$ divides $[G : H][H : A] = [G : A]$.

Case 2: $\mathrm{Res}_A^G V$ is isotypic. Since every irreducible representation of A is one-dimensional, this means that A must act on V by scaling.

Let G', A' be the images of G, A respectively in $\mathrm{GL}(V)$. Then $A' \subset C(\mathrm{GL}(V)) \subset C(G')$, and G/A surjects onto G'/A'.

By the previously proved theorem, then, $\dim V \mid [G' : C(G')] \mid [G' : A'] \mid [G : A]$.

□

31 Lie Groups

We'll spend the last couple lectures of class talking about Lie groups and their representations. This section will be lighter on proofs than the previous ones.

Roughly, a Lie group is a group which is also a manifold. We won't be too precise about this, because we'll try to focus on the algebra. We'll start with some examples, and then give some slightly more precise definitions.

Example. The groups

$$\mathrm{GL}_n(\mathbb{R}) = \{A \in M_{n\times n}(\mathbb{R}) \mid \det A \neq 0\}, \mathrm{GL}_n(\mathbb{C}) \quad = \{A \in M_{n\times n}(\mathbb{C}) \mid \det A \neq 0\}.$$

are Lie groups because they are open subsets of $M_{n\times n}(\mathbb{R}) \cong \mathbb{R}^{n^2}$ and $M_{n\times n}(\mathbb{C}) \cong \mathbb{C}^{n^2} \cong \mathbb{R}^{2n^2}$ (both of these are smooth manifolds).

Next we'll look at examples of Lie groups that are subgroups of $GL_n(\mathbb{R})$ and $GL_n(\mathbb{C})$.

Example. Define

$$SL_n(\mathbb{R}) = \{A \in GL_n(\mathbb{R}) \mid \det A = 1\}, SL_n(\mathbb{C}) = \{A \in GL_n(\mathbb{C}) \mid \det A = 1\}.$$

Example. The subgroup $B_n(\mathbb{R}) \subset GL_n(\mathbb{R})$ consisting of all upper-triangular invertible matrices. The subgroup $B_n(\mathbb{C}) \subset GL_n(\mathbb{C})$ is defined analogously.

The subgroups $N_n(\mathbb{R})$ and $N_n(\mathbb{C})$ of $B_n(\mathbb{R})$ and $B_n(\mathbb{C})$ respectively, consisting of all upper-triangular matrices which have all 1's on the diagonal.

Example. The subgroups

$$O_n(\mathbb{R}) = \{A \in GL_n(\mathbb{R}) \mid A^t A = 1_n\}, O_n(\mathbb{C}) = \{A \in GL_n(\mathbb{C}) \mid A^t A = 1_n\}$$

as well as the subgroups $SO_n(\mathbb{R}) = O_n(\mathbb{R}) \cap SL_n(\mathbb{R})$ and $SO_n(\mathbb{C}) = O_n(\mathbb{C}) \cap SL_n(\mathbb{C})$.

Example. The subgroup $U_n(\mathbb{C}) = \{A \in GL_n(\mathbb{C}) \mid A^* A = 1_n\}$, where here $A^* = \overline{A}^t$ and its subgroup $SU_n(\mathbb{C}) = U_n(\mathbb{C}) \cap SL_n(\mathbb{C})$.

We'll work out more explicitly what $U_1(\mathbb{C})$ and $SU_2(\mathbb{C})$ are:

$$U_1(\mathbb{C}) \cong \{a \in \mathbb{C}^* \mid a\overline{a} = |a|^2 = 1\}$$

which is the unit circle in $\mathbb{C}^*$.

Finding $SU_2(\mathbb{C})$ is a bit harder; a matrix $A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$ lies in $SU_2(\mathbb{C})$ if and only if $\det A = ad - bc = 1$, and if $A^{-1} = A^*$, that is

$$\begin{pmatrix} d & -b \\ -c & a \end{pmatrix} = \begin{pmatrix} \overline{a} & \overline{c} \\ \overline{b} & \overline{d}. \end{pmatrix}$$

Hence we conclude that $d = \overline{a}$ and $c = -\overline{b}$. Then the condition $\det A = ad - bc = |a|^2 + |b|^2$. Hence

$$SU_2(\mathbb{C}) = \left\{ \begin{pmatrix} a & b \\ -\overline{b} & \overline{a} \end{pmatrix} \mid |a|^2 + |b|^2 = 1 \right\}$$

hence $SU_2(\mathbb{C})$ is homeomorphic to the unit 3- sphere in $\mathbb{C}^2$ (or $\mathbb{R}^4$).

Now we'll give a slightly more formal definition:

Definition. A (real) Lie group

- is a group G
- which is also a smooth (C^∞) manifold ; that is, $G \subset \mathbb{R}^N$ for some N, and locally near any point g of G, G looks like a copy of $\mathbb{R}^n$ inside $\mathbb{R}^N$. More specifically, at this point g we have well-defined tangent space to G inside $\mathbb{R}^N$, and locally near g, G looks like its tangent space. In this context, we have a notion of a smooth (C^∞) function from G to $\mathbb{R}$, and more generally, a notion of a smooth function from G to any other smooth manifold.

- and such that the multiplication map $G \times G \to G$ and inverse map $G \to G$ are maps of smooth manifolds.

All the examples we've give above are Lie groups; it's somewhat annoying but not too hard to check this.

(There is also a notion of a complex Lie group G, which we won't really use as much, but will mention: $G \subset \mathbb{C}^n$ is a complex Lie group if it is a Lie group that is also a complex manifold; (this essentially means that all of its tangent spaces are complex (affine linear) subspaces of $\mathbb{C}^n$), and such that multiplication and inverse maps are holomorphic functions. The group $GL_n(\mathbb{C})$ is a complex Lie group, as are all of the subgroups of $GL_n(\mathbb{C})$ we defined other than $U_n(\mathbb{C})$ and $SU_n(\mathbb{C})$.)

Definition. A morphism $\phi : G \to H$ of Lie groups is a group homomorphism ϕ which is also a smooth map of manifolds.

It can be shown, but it's not easy, that if $\phi : G \to H$ is a group homomorphism, then ϕ is smooth if and only if it is continuous; so it's enough to check continuity.

(There is also a notion of a morphism of complex Lie groups, using holomorphic instead of smooth. This is actually a stronger condition than being differentiable.)

Definition. If G is a (real) Lie group, then a representation of G is a morphism of Lie groups $\rho : G \to GL(V)$, where V is a finite-dimensional vector space either over $\mathbb{R}$ (in which case we say that ρ is a "real" representation) or over $\mathbb{C}$ (in which case ρ is a "complex" representation).

Example. If G is any subgroup of $GL_n(\mathbb{R})$, then the inclusion map $G \to GL_n(\mathbb{R}) = GL(\mathbb{R}^n)$ gives a real representation of G.

If G is any subgroup of $GL_n(\mathbb{C})$, the inclusion map $G \to GL_n(\mathbb{C}) = GL(\mathbb{C}^n)$ gives a complex representation of G.

In either case, this representation is called the "standard representation" of G.

Just as with finite groups, we can build up new representations from old, using $\oplus, \otimes, \mathrm{Hom}_{\mathbb{C}}, \mathrm{Sym}^n, \bigwedge^n$.

Example. If $G = U_1 = \{a \in \mathbb{C}^\times \mid |a| = 1\}$, then for any $n \in \mathbb{Z}$ we can define a complex representation $\rho_n : G \to GL(V_n)$ such that V_n is 1-dimensional, and $\rho_n(a) = a^n \in \mathbb{C}^\times \cong GL_n(\mathbb{C})$.

(For n a positive integer, we could also have constructed ρ_n as a tensor power: $V_n = V_1^{\otimes n} = V_1 \otimes \cdots \otimes V_1$.)

The definition of irreducible representation carries over directly. However, it is no longer the case that any representation V has an irreducible decomposition $V = \oplus_i W_i$. (When this is the case V is said to be "completely reducible".)

For instance, let $G = N_2(\mathbb{C}) = \left\{ \begin{pmatrix} 1 & b \\ 0 & 1 \end{pmatrix} \mid b \in \mathbb{C} \right\} \cong \mathbb{C}^+$, $V = \mathbb{C}^2$, with ρ the standard representation. Then any nonzero proper invariant subspace of V must be of the form

span(v) for v a simultaneous eigenvector of every element of N_2. The only such v is $\begin{pmatrix}1 & 0\end{pmatrix}$, so V does not have an irreducible decomposition.

However, if G is a compact group, then complete reducibility does hold; and the proof of this is completely analogous to the proof we did in class. We'll say more about this next time.

Let G be a Lie group. Last time, we defined representations of G, and gave an example of a representation V of the Lie group $N_2(\mathbb{C})$ which is not completely irreducible. We said however that complete irreducibility does hold when G is a compact. A bit more detail on why:

For any Lie group G, we can define a left-invariant measure $\int_G dg$ on G (known as Haar measure), unique up to scaling. Since we are handwaving the analysis here, we won't go into detail on what this means, but roughly: having a measure means that for a function $f : G \to \mathbb{C}$ we can define $\int_G f dg$. Left-invariant of dg means that $\int_G f(g)dg = \int_G f(g'g)dg$ for any $g' \in G$. This measure is unique up to scaling.

If G is compact, more is true: first, $\int_G 1 dg$ is always finite, and we can rescale to make $\int_G 1 dg = 1$; this now uniquely specifies dg. As well, dg is also right-invariant $\int_G f(g)dg = \int_g f(gg')dg$ for any $g' \in G$.

Now, for any finite-dimensional representation $\rho : G \to GL(V)$ of G, we can define an averaging map $r : V \to V^G$ by

$$r(v) = \int_G \rho_g(v) dg.$$

Using this averaging map, all the proofs of complete reducibility go through exactly the same as they would in the finite group case.

32 Character Theory of Compact Lie Groups

If G is a Lie group, and $\rho : G \to GL(V)$ is any finite-dimensional representation, we can define the character $\chi_V : G \to \mathbb{C}$ as usual; $\chi_V(g) = \operatorname{tr} \rho_g$. Then χ_V lies in the space $\mathcal{C}^\infty_{\text{class}}(G)$ of $\mathcal{C}^\infty$ class functions on G. If G is compact, this space $\mathcal{C}^\infty_{\text{class}}$ is contained in the Hilbert space $L^2_{\text{class}}(G)$ of class functions that are L^2 with respect to the measure dg on G.

Assume G is compact. Then we have the following result, analogous to the case when G is finite: the characters $\{\chi_i\}_{i\in I}$ of the irreducible representations $\{V_i\}_{i\in I}$ of G are orthogonal with respect to the inner product $\langle \alpha, \beta \rangle = \int_G \overline{\alpha(g)}\beta(g)dg$. Again, the proof is identical to the proof for G finite, only using the continuous averaging map instead.

Furthermore, one can show that the characters $\{\chi_i\}_{i\in I}$ form an orthonormal basis for the Hilbert space $L^2_{\text{class}}(G)$. (As pointed out in class, this means that there are only countably many such.)

33 Example: $G = U_1(\mathbb{C})$

We'll finish by seeing a couple examples of the above.

For our first example, we'll take $G = U_1(\mathbb{C}) \cong \{\alpha \in \mathbb{C}^\times \mid |\alpha| = 1\}$. Last time we defined irreducible one-dimensional representations $\rho_n : G \to GL(V_n)$ for every integer n, such that $\rho_n(\alpha) = \alpha^n \in \mathbb{C}^\times \cong GL_1(\mathbb{C})$. Hence the character χ_n of ρ_n is given by $\chi_n(\alpha) = \alpha^n$.

By the above, the functions $\chi_n(\alpha)$ must be orthonormal with respect to the invariant measure on G, and they form an orthonormal basis for the Hilbert space $L^2(G)$ if and only if every irreducible representation of G is isomorphic to some V_i.

This latter statement is in fact true; depending upon whether one prefers algebra or analysis one can either classify the irreducible representations of G and deduce that the $\chi_n(\alpha)$ are a basis for $L^2(G)$ or else one can show that the $\chi_n(\alpha)$ form a basis and deduce classification of irreducible representations of $G \cong U_1(\mathbb{C})$.

To do this using analysis; we identify $G \cong U_1(\mathbb{C})$ with $\mathbb{R}/2\pi\mathbb{Z}$ using the parametrization $\mathbb{R}/2\pi\mathbb{Z} \to U_1(\mathbb{C})$ given by $\theta \mapsto e^{i\theta}$. In this parametrization, the invariant measure on $U_1(\mathbb{C})$ is given by $dg = \frac{1}{2\pi}d\theta$, and the characters χ_n of $U_1(\mathbb{C})$ are given by $\chi_n(e^{i\theta}) = e^{in\theta}$. It is a well-known fact of Fourier analysis that the functions $\{e^{in\theta}\}_{n\in\mathbb{Z}}$ comprise a basis for $L^2(\mathbb{R}/2\pi\mathbb{Z})$. Hence we may conclude from this that the V_n are all the irreducible representations of $U_1(\mathbb{C})$.

(Alternatively, as suggested above, one could classify the irreducible representations of $U_1(\mathbb{C})$ first and deduce Fourier theory.)

34 Example: $G = SU_2(\mathbb{C})$

We now move on to the example of

$$G = SU_2(\mathbb{C}) = \left\{ \begin{pmatrix} a & b \\ -b & a \end{pmatrix} \mid |a|^2 + |b|^2 = 1 \right\}$$

, which we previously identified with the three-sphere $S^3 \subset \mathbb{R}^4$ by sending the matrix $\begin{pmatrix} a & b \\ -b & a \end{pmatrix}$ to $(a, b) \in \mathbb{C}^2 \cong \mathbb{R}^4$. Under this identification, the invariant measure dg on $SU_2(\mathbb{C})$ is a scalar multiple of the standard surface area measure on S^3.

Let V be the standard (2-dimensional) representation of $G = SU_2(\mathbb{C})$, that is, given by the inclusion $SU_2(\mathbb{C}) \hookrightarrow SL_2(\mathbb{C})$. For $n \geq 0$ let $V_n = \mathrm{Sym}^n(V)$; so V_n is n-dimensional, V_0 is the trivial representation, and $V_1 = V$.

One can verify that V_n is irreducible for all n, either directly from the definition, or by computing $\langle \chi(V_n), \chi(V_n) \rangle = \int_G |\chi(V_n)(g)|^2 dg$ and showing that it equals 1. We'll skip this verification.

We'll compute the characters of $\{V_n\}_{n\geq 0}$ and establish that they form a basis for $L^2_{\text{class}}(G)$; this will imply that every irreducible representation of G is isomorphic to some V_n.

First we determine the conjugacy classes of $SU_2(\mathbb{C})$. By the spectral theorem for unitary operators, any conjugacy class of $SU_2(\mathbb{C})$ contains a diagonal matrix of the form $\begin{pmatrix} \lambda & 0 \\ 0 & \lambda^{-1} \end{pmatrix}$ with $|\lambda| = 1$, and this matrix is unique up to switching λ and λ^{-1}; if we add the condition that $\operatorname{Im}\lambda \geq 0$ then every conjugacy class has a unique such representative. We may parametrize these conjugacy classes by setting $\lambda = e^{i\theta}$ for $\theta \in [0, \pi]$.

Using this parametrization, we can identify class functions on $G = SU_2(\mathbb{C})$ with functions on $[0, \pi]$. One can compute that the space $L^2_{\text{class}}(G)$ does not get identified with the ordinary space $L^2([0, \pi])$, but instead with the space of functions on $[0, \pi]$ that are L^2 with respect to the measure $\sin^2(\theta)d\theta$.

We can find the character χ_n of $V_n = \operatorname{Sym}^n(V)$ the same way that you computed the character of Sym^n of a 2-dimensional representation on the homework. We have

$$\chi_n\left(\begin{pmatrix} \lambda & 0 \\ 0 & \lambda^{-1} \end{pmatrix}\right) = \lambda^n + \lambda^{n-2} + \cdots + \lambda^{-n} = \frac{\lambda^{n+1} - \lambda^{-n-1}}{\lambda - \lambda^{-1}}.$$

If we now set $\lambda = e^{i\theta}$ according to our parametrization, this becomes

$$\chi_n\left(\begin{pmatrix} e^{i\theta} & 0 \\ 0 & e^{i\theta} \end{pmatrix}\right) = \frac{e^{(2n+1)i\theta} - e^{-(2n+1)i\theta}}{e^{i\theta} - e^{-i\theta}} = \frac{\sin(2n+1)\theta}{\sin\theta}.$$

One can use Fourier analysis on the interval again to show that these functions form a basis for the Hilbert space $L^2_{\text{class}}(G)$. We sketch this here; if we have $f \in L^2_{\text{class}(G)}$ we can, by the parametrization above, view f as a function on the space of functions that are L^2 with respect to $\sin^2(\theta)d\theta$. Glue the functions $\sin(\theta)f(-\theta)$ on $[-\pi, 0]$ and $\sin(\theta)f(\theta)$ on $[0, \pi]$ to get an function g on $[-\pi, \pi]$ which is L^2 with respect to the ordinary inner product.

Then g has a Fourier series, and since g was constructed to be an odd function, this Fourier series contains only terms of the form $\sin((n+1)\theta)$ for $n \geq 0$. Hence f itself can be written linear combination of functions of the form $\frac{\sin(n+1)\theta}{\sin\theta}$. This implies that the functions χ_n form a basis for $L^2_{\text{class}}(G)$, and hence that all representations of $G = SU_2(\mathbb{C})$ of of the form $V_n \cong \operatorname{Sym}^n(V)$ for some non-negative integer n.