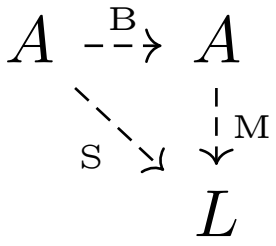

Advanced Abstract Algebra

Shivam Nadimpalli

Last updated: June 1, 2018

Contents

0 Hello!

These notes are based on the Advanced Abstract Algebra course taught by Dr. Péter Hermann, as a part of the Budapest Semesters in Mathematics, Spring 2018. The course attempts to give an idea of some basic methods in finite and infinite group theory. All errors in these notes are my responsibility; please email any comments or typos to nadimpalli@brown.edu.

Notation

Below we keep a running track of notation used in class. Occasionally, especially in the section on representation theory, we might write g^{ϕ} instead of $\phi(g)$.

$\mathscr{O}(g)$	The order of a group element g
$G(x)$	Orbit of x under action of G
G_x	Stabilizer of x under action of G
$C_G(g)$	Centralizer of g in G
$Z(G)$	Center of G
C_a	Conjugacy class of a in G
$N_G(H)$	Normalizer of $H \leq G$ in G
$\mathrm{Syl}_p(G)$	Sylow p-subgroups of G
$\tau_{G \to A}$	Transfer homomorphism from G to abelian subgroup A of finite index
$[a, b]$	Commutator of $a, b = a^{-1}b^{-1}ab$
$[A, B]$	$\langle [a, b] \mid a \in A, b \in B \rangle$
G'	Commutator subgroup of G; equivalently $[G, G]$
$\phi(G)$	Frattini subgroup of G
$F(G)$	Fitting subgroup of G
$F(X)$	Free group generated by X
$\Gamma(G; X)$	Directed, colored Cayley graph of group G generated by X
$\mathrm{Hom}(V, W)$	$\{f : V \to W \mid f \text{ is a linear map}\}$
$\mathrm{End}(V)$	Equivalent to $\mathrm{Hom}(V, V)$
$\mathrm{GL}(V)$	The group of invertible endomorphisms of V
$\mathrm{GL}_n(\mathbb{C})$	Invertible $n \times n$ matrices with entries from $\mathbb{C}$
χ	Will usually stand for the character of a representation
χ^G	Induced characters, see corresponding section for more

Part I

Topics in Group Theory

1 Groups Actions

Group actions are probably the most central topic in group theory: if a group G acts on a set X, then it captures some information about the symmetries of X. Almost every important topic that we will talk about in this course will be formulated in the language of group actions.

As a motivating example, suppose we want to count the number of its symmetries, namely the number of edge preserving bijections on its vertex set.

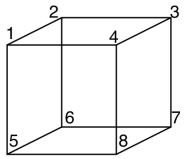

Instead of attempting to count the number of valid transformations, it is much easier to study a group that "acts" on its vertices.

We now give the precise definition of a group action, and introduce related concepts.

1.1 Definition of a group action

Definition 1.1.1 (group action)
Let $X \neq \emptyset$ be a set, and G a group. Then G acts on X if:

1. $(\forall \alpha \in X)(\forall c \in G) : c(\alpha) \in X$
2. $(\forall \alpha \in X)(\forall c, d \in G) : cd(\alpha) = c(d(\alpha))$
3. $(\forall \alpha \in X) : e(\alpha) = \alpha$

Example 1.1.2
Consider the trivial action of a group G on X: $(\forall \alpha \in G)(\forall x \in X) : \alpha(x) = x$.

We now look at an alternative definition of group actions. In most situations, we will be using the one above, but the latter sheds light on the importance of symmetric groups, and occasionally comes in handy.

Definition 1.1.3 (group action)
Group G acts on a set $X \neq \emptyset$ if there exists a homomorphism $G \to \mathrm{Sym}(X)$, where $\mathrm{Sym}(X)$ is the group of permutations of the set X. We define for each $g \in G$ a function as follows: $f_g : X \to X$; $\beta \mapsto g(\beta)$.

We immediately become informal and say that "G acts on X" where G is understood to be a group, and X a nonempty set.

Given a set of n elements, what's the most natural group you can think of that acts on it? Chances are you are thinking of the symmetric group on n elements, namely S_n. In order

to capture how far away a group action is from a permutation, we introduce the following notion:

Definition 1.1.4 (faithful action)
Let G act on X. This action is faithful if $(\forall e \neq a \in G)(\exists \beta \in X) : a(\beta) \neq \beta$.

Put in words, a group action is said to be faithful if for every element of the group, not all the elements of the underlying set are fixed points. Thus, if an action is faithful, then every group element moves at least one element of the set. All this will make more sense after the next proposition:

Proposition 1.1.5
Consider G acting on a set X. Then f_g is a bijection.

Proof. Injectivity of f_g is easy to show. To show surjectivity, for $\gamma \in X$, note that $f_g(g^{-1}(\gamma)) = \gamma$. We conclude that f_g is a bijection from X to X. □

Thus, you can think of group actions as shadows of the permutation groups on the sets. This leads us to the next proposition:

Proposition 1.1.6
The group action of G on X is faithful if and only if the corresponding homomorphism from $G \to \mathrm{Sym}(X)$ is injective.

The proof is omitted as it should be easy to see; if it isn't, then reread the earlier section until it is.

1.2 Orbits and stabilizers

Associated to a group action are the fundamental concepts of orbits and stabilizers.

Definition 1.2.1 (orbits and stabilizers)
Let G act on X. Suppose $g \in G$ and $x \in X$. Then:

1. *$G_x := \{c \in G \mid c(x) = x\}$ is the* **stabilizer** *of x in G.*
2. *$G(x) := \{g(x) \mid g \in G\}$ is the* **orbit** *of x in G. We will also occasionally write $\mathcal{O}_x$ for the orbit of x.*
3. *If $G(x) = X$, then the action is said to be* **transitive**.
4. *If $G_x = G$, then x is said to be a* **fixed point** *of the action.*

So the orbit of an element is the set of all the elements it can "reach" via the action of some element of the group; if it can reach any element of the set, then the action is said to be transitive. The stabilizer of the element, which should be easy to remember, is the set of group elements that fix (stabilize) the element; if a set element is unmoved by all elements of the group, then it is said to be a fixed point of the action.

Proposition 1.2.2 (stabilizer forms subgroup)
Let G act on X. Then for all $x \in X$, $G_x \leq G$.

Proposition 1.2.3 (orbits partition the set)
Let G act on X. The orbits of this action partition X.

Important: Keep in mind that orbits do not necessarily have the same size!

It follows from Proposition 2.2.3 that the orbits of two elements are either disjoint or the same, and that you can write $|X| = |\mathcal{O}_1| + \ldots + |\mathcal{O}_k|$. We might refer to this in the future as the "orbit decomposition" of X.

Some fundamental group actions

We now look at a group G acting on itself and its subgroups, cosets, etc.

Permutations of group elements

Let G act on itself, and define for all $g, h \in G$, $g(h) = gh$. This action is called the **permutation action** of G on itself.

Exercise. Check that this is a valid group action! Is this action faithful?

Permutations of a subgroup

Given a subgroup $H \leq G$, denote the set of its left cosets by X. We can define a permutation action on the X. For every $aH \in X$, $g \in G$, define $g(aH) = gaH$. This is well-defined because we're only working with coset representatives.

What's the kernel of this action? Indeed, $\ker = \{c \in G \mid \forall a \in G,\ c(aH) = aH\}$. However, for a given $a \in G$, we have:

$$c(aH) = aH \iff a^{-1}ca \in H \iff ca \in aH \iff c \in aHa^{-1}$$

And so, $c \in \bigcap_{a \in G} aHa^{-1}$, and this is precisely the kernel of the above action.

Conjugation on an element

Let G act on itself, and define for all $g, h \in G$, $g(h) = ghg^{-1}$. This action is called the **conjugation action** of G on itself. This leads us naturally to the next definition:

Definition 1.2.4 (centralizer of an element)
The centralizer of $g \in G$, written $C_G(g)$ is the stabilizer of g under the conjugation action.

Since stabilizers are subgroups, it follows that $C_G(g) \leq G$.

We have a special name for the intersection of all centralizers of group elements: it's called the center of the group (written $Z(G)$), and is (in words) the subgroup consisting of elements that commute with all other group elements.

Proposition 1.2.5 (center of a group)
Define the center of a group $Z(G) = \bigcap_{g \in G} C_G(g)$. We have the following:
1. *$Z(G) \trianglelefteq G$*
2. *If $G/Z(G)$ is cyclic, then G is abelian*

What are the orbits of this action? They are precisely the **conjugacy classes** of G! For $a \in G$, we might denote the conjugacy classes of a in G by C_a. Again, important to note that the sizes of the conjugacy classes need not be the same; we will see later, however, that they must divide the size of the group.

It is worthwhile to note that for any nontrivial group G, the conjugacy class of $e = \{e\}$. Thus, any nontrivial group has **at least two** conjugacy classes.

Conjugation on a subgroup

We can similarly define conjugation on a subgroup: for $H \leq G$, $g(H) = gHg^{-1}$. The stabilizer of $H \leq G$ under this action also has a special name:

Definition 1.2.6 (normalizer of a subgroup)
Given $H \leq G$, $N_G(H) = \{g \in G \mid gHg^{-1} = H\}$; alternatively, $N_G(H)$ is precisely the stabilizer of H under the conjugation action.

The following proposition follows naturally:

Proposition 1.2.7 (normalizer of normal subgroup is whole group)
$H \trianglelefteq G \iff N_G(H) = G$.

1.3 Orbit-Stabilizer Theorem and Burnside's Lemma

Theorem 1.3.1: Orbit-Stabilizer
Let G act on X, $\beta \in X$. Then $|G(\beta)| = |G : G_\beta|$. In particular, the stabilizers of elements in the same orbit have the same size.

Proof. We know that $G_\beta \leq G$, and so can consider the left cosets of G_β. Note that each coset simply specifies an element of $G(\beta)$. Thus, we have a bijection between $G(\beta)$ and the left cosets of G_β. □

We can finally return to our original motivating example:

Example 1.3.2 (symmetries of the cube)

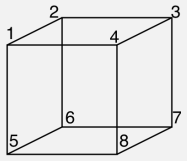

Without loss of generality, take $G \leq S_8$, and consider its action on the vertices of the cube above. Note that $G(1) = \{1, 2, \ldots, 8\}$. From the Orbit-Stabilizer theorem, we have $|G(1)| = |G : G_1| \implies |G| = 8|G_1|$.

As $G_1 \leq G$, we can consider the action of G_1 on the vertices of the cube, i.e. we are looking at those transformations of the cube that fix vertex 1. Under this action, the orbit of vertex 2 is given by: $G_1(2) = \{2, 4, 5\}$. This is because the connectivity of $(1, 2)$ must be preserved. Again, by Orbit-Stabilizer, we have: $|(G_1)_2| =: |G_{1,2}| \implies |G_1| = 3|G_{1,2}|$.

Now, consider the action of $G_{1,2}$ on the cube: $G_{1,2}(3) = \{3, 6\}$, and so, by Orbit-Stabilizer, $|G_{1,2} = 2|G_{1,2,3}|$. However, $|G_{1,2,3}| = 1$, as fixing those three vertices fixes the rest of the cube. Thus, $|G| = 8 \times 3 \times 2 = 48$.

Now, what if one wants to count the number of orbits of a group action? Indeed, the lemma-that-was-not-Burnside's (Cauchy got there first) allows us to do so easily.

Lemma 1.3.3 (Burnside's Lemma)
Assume G acts on X, G and X both finite. Define $f : G \to \mathbb{N}$ as follows:

$$(\forall\ g \in G) f(g) = |\{\beta \in X \mid g(\beta) = \beta\}|$$

Then the number of distinct orbits of X under G is $\frac{1}{|G|} \sum_{g \in G} f(g)$.

Proof. The proof uses a double-counting argument. We consider the total number of pairs of the form (β, g) where $\beta \in X$, $g \in G$, and $g(\beta) = \beta$.
For a given $g \in G$, he number of such pairs is equal to $f(g)$, as each pair corresponds to a point that is fixed by g. However, for each $\beta \in X$, the number of such pairs is given by $|G_\beta|$. Thus we have:

$$\sum_{g \in G} f(g) = \sum_{\beta \in X} |G_\beta| = |G| \sum_{\beta \in X} \frac{1}{|G(\beta)|} = |G| \sum_{T \text{ orbit}} \left(\sum_{\beta \in T} \frac{1}{|T|} \right)$$

where the inner sum in the final expression is equal to 1. The desired result then follows. □

1.4 Sylow Theorems

Is the converse of Lagrange's Theorem true? Namely, if the $d||G|$, then does there exist a subgroup of order d in G? This turns out to be false: the simplest example of this is the

group A_4, of order 12, which has no subgroup of order 6.

However, Sylow discovered that the converse is true when d is a prime power; he also discovered special relations between subgroups whose order was the highest prime power dividing the order of the group. All this is stated below:

Theorem 1.4.1: Sylow Theorems
Let G be a finite group, p a prime number such that p^k divides $|G|$. Then:

S1. *There exists a subgroup of order p^k in G. Moreover, the number of such subgroups is congruent to $1 \pmod p$.*

S2. *Let p^n be the highest power of p that divides $|G|$. Subgroups of order p^n are called* **Sylow p-subgroups** *of G ($\mathrm{Syl}_p(G)$ for short). Then all Sylow p-subgroups are conjugates of each other, and are hence isomorphic.*

S3. *Let R be a Sylow p-subgroup of G, and $H \leq G$ such that $|H| = p^k$. Then there exists $d \in G$ such that $H \leq dRd^{-1}$.*

I was lazy and did not take the proof down. Several proofs can be found online or in the textbook for the course ("A Course in the Theory of Groups" by Robinson). Plenty of group actions used in the proofs.

Corollary 1.4.2 (S1*)
Suppose $|G| = p^k m$ such that p^{k+1} does not divide $|G|$. Then $|\mathrm{Syl}_p(G)| \mid m$.

Proof. Consider the action of conjugation by elements of G on $\mathrm{Syl}_p(G)$. Then from S2, all the $\mathrm{Syl}_p(G)$ subgroups are conjugates of each other, and so there is only one orbit. It then follows that the size of this orbit, namely $|\mathrm{Syl}_p(G)|$ must divide $|G| = p^k m$, but as $|\mathrm{Syl}_p(G)| \equiv 1 \mod p$, the number $|\mathrm{Syl}_p(G)|$ and p are co-prime, so $|\mathrm{Syl}_p(G)| \mid m$. □

2 The Transfer Homomorphism & Applications

In this section, we create a monster...

2.1 The transfer function

Definition 2.1.1 (transfer homomorphism)
Let G be a group, $A \leq G$ such that $|G : A|$ is finite. Suppose A is abelian. Let $G = \bigcup_{i \in I} c_i A$. We can consider the action of G on the index set I of the coset representatives as follows: for $x \in G$, $x \cdot c_i A = c_{x(i)} A$ where $x c_i = c_{x(i)} a_{x,i}$ for some $a_{x,i} \in A$.
Then the transfer function is given by:

$$\tau(x) = \prod_{i \in I} a_{x,i}$$

What happens if, in the above definition, $|G : A|$ is not finite? Then there are infinitely many cosets of A in G, and hence, $\tau(x)$ is an infinite product, which is not a nice thing. What happens if A is not abelian? Then clearly the product is not well defined.

It turns out that the transfer plays an important role in the study of unsolvable groups and aids in the characterization of groups in which all Sylow p-subgroups are cyclic. When we get to Riesz representations, we will probably see some nifty applications of the transfer.

Note to self: Stare at the above definition until it makes sense!

2.2 Properties of the transfer

In what follows, we take $A \leq G$ to be an abelian group with finite index. We might also sometimes refer to the set of coset representatives $\{c_1, \ldots, c_k\}$ such that $G = \bigcup_i c_i A$ as a (left) **traversal** for A in G; we will refer to the index set $\{1, \ldots, k\}$ by I.

Proposition 2.2.1 (transfer is a homomorphism)
The transfer function $\tau : G \to A$ is a homomorphism.

Proof. Let $x, y \in G$, and consider the action of xy on the index set of coset representatives I as described in Defiition 2.1.1. We thus have for $i \in I$:

$$xy(c_i) = x(y(c_i)) = x(c_{y(i)} a_{y,i}) = (x c_{y(i)}) a_{y,i} = c_{x(y(i))}(a_{x,y(i)} a_{y,i})$$

However, we also have:

$$xy(c_i) = c_{xy(i)} a_{xy,i} = c_{x(y(i))} a_{xy,i}$$

And so we can write:

$$a_{xy,i} = a_{x,y(i)} a_{y,i}$$

It is then easy to see that $\tau(xy) = \tau(x)\tau(y)$, because A is abelian and as $y(i)$ must range over all of the index set (because left multiplication is a bijection). □

Proposition 2.2.2 (transfer doesn't depend on traversal)
Let $A \leq G$ as described in Definition 2.1.1. Then τ does not depend on the choice of coset representatives of A in G.

Proof. Let $\{d_i\}$ be an alternative set of coset representatives for A in G. For every $i \in I$, we thus have $d_i A = c_i A$ such that $d_i = c_i \alpha_i$ for some $\alpha_i \in A$. Consider the action of $x \in G$ on some i in the index set. We thus have:

$$xd_i = d_{x(i)} a^*_{x,i} = c_{x(i)} \alpha_{x(i)} a^*_{x,i}$$

But we also have:

$$xd_i = xc_i\alpha_i = c_{x(i)} a_{x,i} \alpha_i$$

And so we can write:

$$c_{x(i)} \alpha_{x(i)} a^*_{x,i} = c_{x(i)} a_{x,i} \alpha_i \implies a^*_{x,i} = \alpha^{-1}_{x(i)} a_{x,i} \alpha_i$$

Again, as A is abelian, and $x(i)$ must also range over the set of all coset representatives, we have $\prod a^*_{x,i} = \prod a_{x,i}$ and so the transfer is independent of the choice of coset representatives. □

2.3 Making the transfer prettier

We will now try to find a nice form for the transfer for arbitrary elements of G. Why are we doing this? Because, as we will see shortly, with a clever choice of coset representatives, computing the transfer will not be as difficult as it appears.

Let $\{c_1 A, \ldots, c_n A\}$ be the cosets of A in G such that $G = \bigcup_i c_i A$, where $A \leq G$ is taken to be an abelian group. Consider the action of G on its cosets by left multiplication. Let $x \in G$ be an arbitrary group element. Then, note that under the action of x, we can decompose the set of left cosets into disjoint cycles.

Let $\{xa_iA, x^2a_iA, \ldots, x^{n_i-1}a_iA\}$ be one such cycle, where n_i is the first power of x in a_iA. It follows, then, that if we have k such disjoint cycles, then $|G : H| = n = \sum_{i=1}^{k} n_i$. We can use this fact to obtain a very nice traversal to compute the image of x under τ: $\{a_i x^j \mid 1 \leq i \leq k, 1 \leq j \leq n_i\}$.

Why is this a nice traversal? Consider the "contribution" to $\tau(x)$ made by the second cycle, namely $\{a_2A, xa_2A, \ldots, x^{n_2-1}a_2A\}$. We have:

$$\begin{aligned} x(a_2A) = xa_2A &\implies xa_2 = xa_2(e) \\ &\vdots \\ x(x^{n_2-1}a_2A) = a_2A &\implies x(x^{n_2-1}a_2) = a_2(a_2^{-1}x^{n_2}a_2) \end{aligned}$$

And so we have a contribution to $\tau(x)$ equal to $e \times e \times \ldots \times (a_2^{-1}x^{n_2}a_2) = a_2^{-1}x^{n_2}a_2$.

By extending the above argument to the k cycles, we can write:

$$\tau(x) = \prod_{i=1}^{k} a_i^{-1} x^{n_i} a_i$$

where $\sum_{i=1}^{k} = n$, and for each $i = 1 : k$, n_i is the smallest positive integer such that $a_i^{-1} x^{n_i} a_i \in A$. It is important to note that in the above construction, we are presupposing that we know the disjoint cycles that the cosets are decomposed into.

Example 2.3.1 (subgroup contained in center)
Let $A \leq G$, A abelian with $|G : A| = n < \infty$. We continue with our notation earlier from this section.

Note that if $x \in Z(G)$, we have the following nice form for $\tau(x)$:

$$\tau(x) = \prod_i (c_i^{-1} x^{n_i} c_i) = \prod_i x^{n_i} = x^{\sum_i n_i} = x^n$$

Now, if $A \leq Z(G)$, then note that elements of the form $c_i^{-1} x^{n_i} c_i$ are always elements of A, and as $A \in Z(G)$, we can rewrite each expression as x^{n_i}, from which it follows that $\tau(x) = x^n$ for all $x \in G$.

2.4 The transfer homomorphism in action

This rest of this section is entirely devoted to applications of the transfer.

2.4.1 Burnside's transfer theorem

Lemma 2.4.1 (normalizer controls fusion)
Let G be a finite group, p a prime, and $R \in Syl_p(G)$. Assume R abelian. Let $a, b \in R$ such that $b = y^{-1}ay$ for some $y \in G$. Then there exists $c \in N_G(R)$ such that $c^{-1}ac = b$.

Proof. As R is abelian, $R \leq C_G(a)$ and also $R \leq C_G(b)$. Thus, we have:

$$y^{-1}Ry \leq y^{-1}C_G(a)y = C_G(y^{-1}ay) = C_G(b)$$

where the equality above is easy to prove (but is important). So both R and $y^{-1}Ry$ are subgroups of $C_G(b)$, and so $R \in \mathrm{Syl}_p(C_G(b))$ and $y^{-1}Ry \in \mathrm{Syl}_p(C_G(b))$. By Sylow's Theorem, there exists $d \in C_G(b)$ such that $d^{-1}Rd = y^{-1}Ry$, and so:

$$R = (dy^{-1})R(yd^{-1}) \implies yd^{-1} \in N_G(R) \implies (yd^{-1})^{-1}a(yd^{-1}) = dbd^{-1} = b$$

□

Theorem 2.4.2: Burnside's transfer theorem
Suppose G is a finite group, p a prime, and $R \in Syl_p(G)$. Assume $R \leq Z(N_G(R))$. Then there exists $N \triangleleft G$ such that $|N| = |G : R|$.

N in the above theorem is called a **normal p-complement in G**, and so the above theorem is also sometimes called "Burnside's normal p-complement theorem". Before proving the above theorem, we make a few remarks.

First, what does the seemingly-grotesque condition $R \leq Z(N_G(R))$ mean? Indeed, if the condition holds, then for any $g \in G$, either $g^{-1}Rg \neq R$, or for every $r \in R$, $g^{-1}rg = r$. So in a way, it is an all-or-nothing condition.

Finally, can we say anything about G/N? Indeed, we have $N \leq G \implies RN \leq G$. From HW1 Problem 1(b), and as $\gcd(|R|, |N|) = 1$, we can write:

$$|RN| = \frac{|R||N|}{|R \cap N|} = \frac{|G|}{1}$$

which means that $RN = G$. And so, by the third isomorphism theorem we have $RN/N \cong R/R \cap N$, which in turn implies $G/N \cong R$.

Now that we have gotten all this out of the way, let's return to the proof of Burnside's transfer theorem:

Proof. Let $r \in R$. We have $\tau(r) = \prod_j y_j^{-1} r^{n_j} y_j$ where $\sum_j n_j = |G : R|$. Note that both r^{n_j} and $y_j^{-1} r^{n_j} y_j$ are elements of R, and so, by Lemma 2.4.1, we have the following:

$$(\forall j)(\exists\ c_j \in N_G(R)) : c_j^{-1} r^{n_j} c_j = y_j^{-1} r^{n_j} y_j$$

However, as $R \leq Z(N_G(R))$, r^{n_j} commutes with c_j, and so

$$r^{n_j} = c_j^{-1} r^{n_j} c_j = y_j^{-1} r^{n_j} y_j$$

Thus we have:

$$\tau(r) = \prod_j r^{n_j} = r^{|G:R|}. \implies \mathrm{Im}\tau \supseteq \{r^{|G:R|} \mid r \in R\} = R$$

with the last equality following from HW2 Problem 5, as $|G : R|$ is co-prime to p. □

Now, when do the conditions for the above theorem hold? As an example, it does when G is finite and p is the smallest prime divisor of $|G|$, and all the $\mathrm{Syl}_p(G)$ are cyclic. Is this too much to ask? No, because, if you remember, if the order of a group is prime, then it is cyclic. We make all this rigorous in the next proposition.

Proposition 2.4.3 (Burnside's transfer theorem isn't unreasonable)
Let G be a finite group, and p be the smallest prime dividing $|G|$, and all $\mathrm{Syl}_p(G)$ are cyclic. Then, if $R \in \mathrm{Syl}_p(G)$, then $R \leq Z(N_G(R))$.

Proof. Consider $\alpha : N_G(R) \to \mathrm{Aut}(R)$, $x \mapsto$ conjugation by x on R. Then, note that $R \leq Z(N_G(R))$ if and only if α is a trivial homomorphism.

Now, for all prime q such that $q \mid |N_G(R)|$, let $P_q \in \mathrm{Syl}_q(N_G(R))$. We make the following claim:

Claim. $\alpha \restriction_{P_q}$ is trivial.

Recall that $|R| = p^n$ for some n. We have two possible cases:

Case 1: $p \neq q$. Then, as the number of automorphisms is equal to the number of generators, we have $|\mathrm{Aut}(R)| = p^n - p^{n-1} = p^{n-1}(p-1)$.
Now, as R finite, and $R = \langle c \rangle$ for some $c \in R$, we have the following series of deductions:

$$\begin{aligned} R = \langle c \rangle &\iff R = \langle c^k \rangle \\ &\iff \mathscr{O}(c) = \mathscr{O}(c^k) = p^n \\ &\iff p^n = \frac{\mathscr{O}(x)}{\gcd(\mathscr{O}(c), k)} \\ &\iff \gcd(p^n, k) = 1 \\ &\iff (p^n - p) \text{ coprime to } p \end{aligned}$$

Now, recall that is $\gcd(|A|, |B|) = 1$, then any homomorphism from A to B must be trivial, and so, it follows that $\alpha \restriction_{P_q}$ is trivial.

Case 2: $p = q$. Then note that $P_q \in \mathrm{Syl}_p(G)$, and so P_q is cyclic, and hence, abelian. It follows that the automorphism is trivial.

Now, if A is a finite group such that $|A| = p_1^{\alpha_1} \dots p_k^{\alpha_k}$, then $\langle R_i \in \mathrm{Syl}_{p_i}(A) \rangle = A$. As the restriction of the homomorphism α to each of the Sylow q-subgroups of $N_G(R)$ is trivial, we can conclude that its restriction to $N_G(R)$ is trivial, from which the desired result follows. □

2.4.2 An application to infinite groups

We introduce several new notions in this subsection.

Definition 2.4.4 (commutator)
Let $a, b \in G$. Then $[a, b] = a^{-1}b^{-1}ab$ is the commutator of a, b.

What's the first thing we can say about the commutator? Indeed, if the commutator of two elements is trivial, then they commute! Namely, $[a, b] = e \iff ab = ba$.

Definition 2.4.5 (commutator subgroup)
The subgroup generated by the commutators of elements of G is called the commutator subgroup, denoted by G'. Thus, $G' = \langle [a, b] \mid a, b \in G \rangle$.

Why do we care about this? Well, firstly, it's what we call a **characteristic subgroup**, i.e. is invariant under all automorphisms. Moreover, it plays an important role when we try to obtain abelian factor groups, because of the next two propositions:

Proposition 2.4.6 (commutator subgroup is normal)
$G' \trianglelefteq G$.

Proposition 2.4.7 (commutator subgroup $\subseteq$ kernel of abelian quotient)
If $N \triangleleft G$ then G/N abelian if and only if $G' \subseteq N$.

Proof. G/N abelian $\iff abN = baN \iff a^{-1}b^{-1}ab \in N \iff G' \subseteq N$. □

We now move on to a somewhat interesting proposition. It's kind of obvious if G is finite, but is pretty cool that it holds even when G is infinite.

Proposition 2.4.8
Suppose $|G : Z(G)|$ is finite. Then the commutator subgroup G' is finitely generated.

Proof. $Z(G)$ having finite index means that $G = \bigcup_{i=1}^{n} Z(G)c_i$. Thus, any two elements $a, b \in G$ can be written as $a = z_1c_i$ and $b = z_2c_j$ for appropriately chosen z_1, z_2, c_i, c_j. Thus:

$$[a, b] = [z_1c_i, z_2c_j] = c_i^{-1}c_j^{-1}c_ic_j$$

and so there are only finitely many commutator elements. Thus, we can conclude that the commutator subgroup is finitely generated. □

In fact, we can actually prove a stronger version of the above proposition, but first a lemma is needed:

Lemma 2.4.9
If G is finitely generated, and $H \leq G$ with $|G : H|$ finite, then H is finitely generated.

Proof. Without loss of generality, assume $G = \langle a_1, \ldots, a_n \rangle$. Now, as H has finite index in G, we have $G = \bigcup_{j=1}^{l} c_jH$. Take $c_1 = e$, and note that we can define $a_ic_j \in c_{a_i(j)}H$ which means $a_ic_j = c_{a_i(j)}h_{i,j}$.

Claim. $H = \langle h_{i,j} \mid 1 \leq i \leq n, 1 \leq j \leq l \rangle$.

Indeed, let $y \in H$. We then have:

$$\begin{aligned} y &= a_{i_1}a_{i_2}\ldots a_{i_{t-1}}a_{i_t}c_1 \\ &= a_{i_1}a_{i_2}\ldots a_{i_{t-1}}c_{a_{i_t}(1)}h_{i_t,1} \\ &\vdots \\ &= c_{a_{i_1}a_{i_2}\ldots a_{i_t}(1)}\left(\prod h_{p,q}\right) \\ &= c_{y(1)}\left(\prod h_{p,q}\right) \\ &= \prod h_{p,q} \end{aligned}$$

and so we're done. □

Proposition 2.4.10
Suppose $|G : Z(G)|$ is finite. Then the commutator subgroup G' is finite.

Proof. The proof idea is as follows: X is a finite group if $N \triangleleft X$ and X/N are both finite. Consider $G' \cap Z(G) \triangleleft G'$.

Claim. $G'/(G' \cap Z(G))$ and $G' \cap Z(G)$ are both finite.

Proof. Recall the 3rd iso. theorem: $H \leq A, N \triangleleft A \implies HN/N \cong H/(H \cap N)$. And so we have:

$$G'/(G' \cap Z(G)) \cong G'Z(G)/Z(G) \leq G/Z(G) \text{ which is finite by assumption}$$

All that is left to show is that $G' \cap Z(G)$ is finite.

Indeed, note that it is necessarily abelian ($\leq Z(G)$), and that a) $|G' : G' \cap Z(G)|$ is finite, b) G' is finitely generated. The latter two together imply $G' \cap Z(G)$ is finitely generated.

Now, note that if A is a finitely generated abelian group, then $A \ni a - \alpha_1^{k_1} \dots \alpha_m^{k_m}$ for some set of generators $\{\alpha_i\}$. Then, A is finite if and only if $\mathscr{O}(\alpha_i)$ finite for every i.

Claim. For $z \in G' \cap Z(G)$, $\mathscr{O}(z) < \infty$.

Proof. Consider the transfer $\tau : G \to Z(G)$. Let $z \in G' \cap Z(G)$. Then $\tau(z) \in \mathrm{Im}\tau \cong G/\ker\tau$. However, as $\mathrm{Im}\tau \leq Z(G)$ which is abelian, it follows that $G/\ker\tau$ is abelian. It follows then that $\ker\tau \supseteq G' \implies \tau(z) = e$.
However, from Example 2.3.1, $z \in Z(G) \implies \tau(z) = z^{|G:Z(G)|}$. However, as $|G : Z(G)|$ is finite, we have $z^{|G:Z(G)|} = e \implies \mathscr{O}(z)$ is finite. □

And so, we can conclude that $G' \cap Z(G)$ is finite. Note that we used Lemma 2.4.9 in the above argument. □

Now as $G'/(G' \cap Z(G))$ and $(G' \cap Z(G))$ are both finite, it follows directly that G' is finite. □

Recall the definition of a **torsion-free group**: G is said to be torsion-free if for all $e \neq g \in G$, $\mathscr{O}(g) = \infty$.

Corollary 2.4.11
Let G be a torsion-free group. Suppose $C \leq G$ such that C is cyclic and $|G : C|$ is finite (which means that C is necessarily infinite). Then G is cyclic.

Proof. First, it is easy to see that G is finitely generated. If $C = \langle t \rangle$, then $G = \bigcup_{i=1}^{|G:C|} Cx_i \implies G = \langle x_1, \dots, x_{|G:C|}, t \rangle$. Now if G were abelian, then this, together with the earlier fact that G is finitely generated, would easily imply that G is isomorphic to $\mathbb{Z}$.

Claim. G is abelian.

Proof. We show that G is abelian by showing that $|G : Z(G)|$ is finite. Let the index of C in G be n, and let $\{x_1, \dots, x_n\}$ be the set of coset representatives of C in G whose union is G. Without loss of generality, take $x_1 = e$. Define $C_2 = C \cap \langle x_2 \rangle$. For each $i = 3 : n$, set $C_i := C_{i-1} \cap \langle x_i \rangle \leq C$.

Now, note that for each i, $C_i \neq 1$. Suppose that this were the case. Then note that $Cx_i, Cx_i^2, \dots$ are all distinct, which is a contradiction. Thus, we can write that $|C : C_i|$ is finite for each i.

By iterating the above argument, we can show that C_n has finite index in C, i.e. has finite index in G. However, $C_n \leq Z(G)$ and so $Z(G)$ has finite index in G. □

The desired result follows easily. □

I wasn't present in class for the above proof, so here's a sketch of another one in case the one above is not very clear:

Proof. $C_G(x_i)$ will intersect C nontrivially, and hence $C \cap C_G(x_i)$ has finite index in C, and hence finite index in G. But $Z(G) = \cap C_G(x_i)$ will also have finite index in G. This implies that G' is finite, and since G is torsion-free, G' is actually trivial, so G is abelian. □

3 Solvable and Nilpotent Groups

We start this section by doing some matrix calculations and linear algebra review. We do this because later we will try to generalize the argument below to solvable and nilpotent groups (to be defined shortly). As you will see in §3.1, working with transformations as functions is usually much cleaner than working directly with matrices.

As a bit more of a teaser, solvable and nilpotent groups are closest in spirit to abelian groups. In a way, they can be viewed as a kind of upper bound of non-abelianness; abelian groups sit right between nilpotent groups and solvable groups.

3.1 Revisiting linear algebra

Recall the commutator from §2.4.2. Consider the group of **upper semitriangular matrices** of size 2 over the same field:

$$G = \left\{ \begin{bmatrix} a & c \\ 0 & b \end{bmatrix} \mid a, b, c \in \mathbb{F}, ab \neq 0 \right\}$$

We make the following two observations:

$$\begin{bmatrix} a & c \\ 0 & b \end{bmatrix}^{-1} = \begin{bmatrix} 1/a & -c/ab \\ 0 & 1/b \end{bmatrix}$$

and:

$$A^{-1}B^{-1}AB = \begin{bmatrix} 1 & z \\ 0 & 1 \end{bmatrix} \text{ for some } z \in \mathbb{F}$$

where $A^{-1}B^{-1}AB$ is the commutator of two matrices A, B. We say that the commutators of two elements of this group are **upper unitriangular matrix**. However, two upper unitriangular matrices commute, as we have:

$$\begin{bmatrix} 1 & z_1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & z_2 \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} 1 & z_1 + z_2 \\ 0 & 1 \end{bmatrix}$$

from which it follows that G', namely the commutator subgroup, is abelian.

Now, we could think of the above matrices as members of $\mathrm{Hom}(V)$, where V is a 2-dimensional vector space over the field $\mathbb{F}$. More specifically, we can think of a matrix from the above discussion as a linear transformation $f \in \mathrm{GL}(2)$. Next, suppose $\{b_1, b_2\}$ is a basis of V; write $V_n = \mathrm{Span}(b_n)$. Then note that $V = V_1 \oplus V_2$.

We can now look at upper semitriangular matrices in a new light: they correspond to precisely those $f \in \mathrm{GL}(2)$ s.t. $f(V_1) \subseteq V_1$, as they only scale along the first entry of any vector. And so, we say that V_1 is f-invariant, where f is a linear transformation corresponding to an upper semitriangular matrix.

Now, if $f \in G'$, then $x = v_1 + v_2$ for $v_1 \in V_1, v_2 \in V_2$, and so $f(x) = f(v_1 + v_2) = f(v_1) + f(v_2)$. However, as $f(v_1) = v_1$ (since we're working with upper unitriangular matrices), and as $f(v_2) = (w_1, v_2)$ where $w_1 = v_2 z$, we have $f(x) = v_1 + (w_1 + v_2) = x + w_1$, and so each vector x is mapped to the sum of itself and some other vector.

3.2 Solvable and Nilpotent Groups

Solvable groups, as you may have guessed, have to do with the solvability of polynomials. Indeed, if you have studied any Galois theory, then you may know that for an equation to be solvable its Galois group must be solvable.

Definition 3.2.1 (solvable group)
G is solvable if $G^{(1)} := G'$, $G^{(2)} := (G^{(1)})'$, ... terminates at 1, i.e. there exists n such that $G^{(n)} = 1$.

When will the above sequence terminate? Recall that $G \geq G^{(1)} \geq G^{(2)} \geq \ldots$, and so if G is finite, then it must terminate. The above series is sometimes called the **derived series** for G.

Proposition 3.2.2 (subgroup of solvable group is solvable)
If G is a solvable group, $H \leq G \implies H$ is solvable.

Proof. Look at $H \geq H^{(1)} \geq H^{(2)} \geq \ldots$ Note that $H^{(k)} \leq G^{(k)}$, and so, as G is solvable, there must exist some $n \in \mathbb{N}$ such that $H^{(n)} = 1$. □

Proposition 3.2.3 (quotient of solvable group is solvable)
If G is a solvable group, $N \triangleleft G \implies G/N$ is solvable.

Proof. We claim that $(G/N)^{(k)} = G^{(k)}N/N$. Indeed, we have:

$$\begin{aligned}(G/N)' &= \langle [N_a, N_b] \mid N_a, N_b \in G/N \rangle \\ &= \langle (Na)^{-1}(Nb)^{-1}NaNb \mid a, b \in G \rangle \\ &= \langle Na^{-1}b^{-1}ab \mid a, b \in G \rangle \\ &= \langle N[a,b] \mid a, b \in G \rangle \\ &= NG'/N = G'N/N\end{aligned}$$

We can repeat this argument for higher values of k, from which the desired result follows. □

We can actually obtain a somewhat stronger result:

Proposition 3.2.4
Let $N \triangleleft G$. Then G is solvable if and only if both N and G/N are solvable.

Proof. The forward direction follows from Propositions 3.2.2 and 3.2.3. As for the reverse direction, we know that G/N is solvable, and so there exists $k \in \mathbb{N}$ such that $(G/N)^{(k)} = N/N$. Thus we can write $N/N = G^{(k)}N/N \implies G^{(k)} \leq N$.

Now, as N is solvable, there exists $l \in \mathbb{N}$ such that $N^{(l)} = 1$, and so $(G^{(k)})^{(l)} = G^{(k+l)} = 1$. We conclude that G is solvable. □

We conclude this subsection with some examples of solvable groups.

Example 3.2.5 (examples of solvable groups)

1. Abelian groups
2. $D_n/\langle a\rangle$ where D_n is the dihedreal group of an n-gon, and $\langle a\rangle$ is the rotation subgroup.
3. S_n solvable if and only if $n \leq 4$. (Why? Because S_n solvable iff A_n solvable. A_1, A_2, A_3 clearly solvable, for A_4 you can use the Sylow Theorems!)
4. All finite groups of prime power order (Why? Because they have nontrivial center)
5. The group of upper unitriangular matrices is solvable.

Just as solvable groups were defined as those with terminating derived series, nilpotent groups are defined as those with terminating "lower central series".

Definition 3.2.6 (nilpotent group)
G is called nilpotent if $G = \gamma_1(G) \geq \gamma_2(G) \geq \ldots$ terminates at 1, where $\gamma_1(G) = G, \gamma_2(G) = G'$, and $\gamma_{n+1}(G) = [G, \gamma_n(G)]$ with $[A, B]$ denoting the subgroup of the form $\langle [a, b] \mid a \in A, b \in B\rangle$.

We might occasionally write γ_n instead of $\gamma_n(G)$ when G is understood. Some facts that'll come in handy later:

1. $G' = [G, G]$
2. $[A, B] \triangleleft \langle A, B\rangle$
3. For $N \leq G$, $N \triangleleft G \iff [G, N] \subseteq N$

Proposition 3.2.7 (properties of lower central series)
For all n:

1. $\gamma_n \triangleleft G$
2. $\gamma_{n+1} \leq \gamma_n$
3. $\gamma_n \supseteq G^{(n)}$

Corollary 3.2.8 (nilpotent implies solvable)
If G is nilpotent, then G is solvable.

Proof. This follows from Proposition 3.3.2. □

Example 3.2.9
The upper unitriangular matrices are nilpotent.

Proposition 3.2.10
Let G be nilpotent, $H < G$. Then $H \lneqq N_G(H)$.

Proof. As G is nilpotent, we have $G \geq \gamma_2 \geq \ldots \geq \gamma_n = 1$. As H is a proper subgroup, there must exist a unique j such that $\gamma_{j+1} \subset H$ but $\gamma_j \nsubseteq H$. Thus, there exists some $x \in \gamma_j / H$. Now, from the definition of γ_{j+1}, it follows that for all $y \in H$, $[x, y] \in \gamma_{j+1} \subset H$, i.e. $x^{-1}y^{-1}xy \in H$. As $y \in H$, it follows that $x^{-1}y^{-1}x \in H$, and so $x \in N_G(H)$.
We conclude that $H \lneq N_G(H)$. □

Recall that C_a denotes the conjugacy class of a in G. As the conjugacy classes are orbits of elements under the conjugation action (see §1.3.3 for a review), they partition our group G, and moreover, the order of each conjugacy class must divide the order of G. All this will come in handy as we study group of prime power order.

Proposition 3.2.11 (groups of prime power order are nilpotent)
Let p be a prime, and G be a group such that $|G| = p^k$ for some k. Then G is nilpotent.

Proof. First, we make the following claim:

Claim. $|Z(G)| > 1$.

Proof. If G is abelian, then $Z(G) = G$ and we're done. If G is not abelian, then look at the conjugacy classes of G:

$$|G| = \sum_{i=1}^{n} |C_i|$$

An element is in the center if and only if its conjugacy class is a singleton (consisting of only the element itself). Now, $C_e = \{e\}$, and so, the as $|G|$ is divisible by p, there must be at least $(p-1)$ singleton conjugacy classes for the above expression to be valid. We conclude that the center is nontrivial. □

We then proceed using induction on $|G|$. Let $\overline{G} = G/Z(G)$. As $Z(G)$ is nontrivial, we have $|\overline{G}| < p^k$ and so our inductive hypothesis holds, i.e. $\overline{G}$ is nilpotent.
Thus, there exists n such that $\gamma_n(\overline{G}) = 1 \implies \gamma_n(G)Z(G)/Z(G) = 1$. Finally:

$$\gamma_n(G)Z(G) \subseteq Z(G) \implies \gamma_n(G) \subseteq Z(G) \implies [G, \gamma_{n+1}(G)] = 1$$

In the above argument, we used a simple property of commutator of quotient groups; make sure to prove it yourself later. □

Lemma 3.2.12 (product of nilpotent is nilpotent)
If A, B are nilpotent groups, then $A \times B$ is nilpotent.

Proof. This is because $\gamma_k(A \times B) = \gamma_k(A) \times \gamma_k(B)$. □

Corollary 3.2.13
If $A_1, \ldots, A_l$ are nilpotent, then $A_1 \times \ldots \times A_l$ is nilpotent.

Before we proceed, we pause to recollect some simple facts about normal groups:

1. $A, B \triangleleft G \implies AB \triangleleft G$
2. $A, B \triangleleft G$ with $A \cap B = 1$, then for all $a \in A, b \in B$, we have $ab = ba$

3. $A, B \triangleleft G$ with $A \cap B = 1$, then $AB \cong A \times B$

Lemma 3.2.14 (Frattini's argument)
If G is a finite group with $N \triangleleft G$, $R \in \mathrm{Syl}_p(N)$, then $G = N \cdot N_G(R) = N_G(R) \cdot N$.

Proof. Let $a \in G$. Note that $a^{-1}Ra \in \mathrm{Syl}_p(N)$ since $a^{-1}Ra \leq N$ as $N \triangleleft G$. Then, from Sylow's Theorems (S2? S3? I forget) there exists some $n \in N$ such that

$$a^{-1}Ra = n^{-1}Rn \implies n^{-1}aRa^{-1}n = R \implies a^{-1}n \in N_G(R) \implies a \in N_G(R) \cdot N$$

□

Theorem 3.2.15
Let G be a finite group. Then G is nilpotent if and only if all Sylow-p subgroups are normal (and hence, unique).

We only give the forward direction of the proof of the above theorem, the reverse direction is easily found online or in Robinson's book. Maybe I will write it up later.

Proof. Let G be nilpotent, and $R \in \mathrm{Syl}_p(G)$. Now, note that $N_G(R) \triangleleft N_G(N_G(R))$. Then, by the Frattini Argument (3.3.9), we have:

$$N_G(N_G(R)) = N_G(R) \cdot N_{N_G(N_G(R))}(R)$$

But the latter is just $N_G(R)$. Then, by Proposition 3.3.5, we have $N_G(R) = G$ and so $R \triangleleft G$. □

3.3 Two special subgroups

From the previous two sections, we can see that nilpotence is quite a nice property; on the other hand, solvability is not as nice. Our next project is to try to characterize solvable groups in terms of nilpotent groups arising from matrix groups.

3.3.1 The Frattini Subgroup

Recall the definition of a maximal subgroup:

Definition 3.3.1 (maximal subgroup)
$H \overset{max}{<} G$ if $H \lneq G$ and there exists no $S \lneq G$ such that $H \lneq S \lneq G$.

The intersection of all the the maximal subgroups of a given group is called the **Frattini subgroup** of that group.

Definition 3.3.2 (Frattini subgroup)
Let G be a finite group. Then $\phi(G) = \bigcap_{H \overset{max}{<} G} H$ is called the Frattini subgroup of G.

Note that the Frattini subgroup is also a **characteristic subgroup**, i.e. is invariant under all automorphisms of the group. Recall that we've encountered another characteristic subgroup before, namely the commutator subgroup of a group.

Lemma 3.3.3

1. $C \overset{char}{\leq} B \triangleleft A \implies C \triangleleft A$.
2. $C \overset{char}{\leq} B \overset{char}{\leq} A \implies C \overset{char}{\leq} A$.

Proposition 3.3.4 (Frattini subgroup is nilpotent)
Let G be a finite group. Then $\phi(G)$ is a nilpotent group.

Proof. Let $R \in \mathrm{Syl}_p(\phi(G))$. Note that for Sylow p-groups, being normal is the same as being a characteristic group. Using the Frattini argument, we have $G = \phi(G) \cdot N_G(R)$. Now, if $N_G(R) = G$, then we're done by the definition of maximal subgroups. If $N_G(R) \lneqq G$, then there exists $N_G(R) \leq M \overset{\max}{<} G$. By definition, $\phi(G) \leq M \implies G = \phi(G) \cdot N_G(R) \leq M$ which is a contradiction. □

Proposition 3.3.5
If $N \triangleleft G$, then N is nilpotent if and only if $N\phi(G)/\phi(G)$ is nilpotent.

Proof. Note that nilpotence is preserved by homomorphism, so the forward direction of the above proposition follows trivially. We prove the reverse direction below.

We need to show that if $R \in \mathrm{Syl}_p(N)$, then $R \triangleleft N$, i.e. $R \triangleleft G$, i.e. $N_G(R) = G$.

Claim. $(R\phi(G))/\phi(G) \in \mathrm{Syl}_p((N\phi(G))/\phi(G))$.

Proof. From the third isomorphism theorem, we have:

$$|(R\phi(G))/\phi(G)| = |R/(R \cap \phi(G))|$$

and the latter divides $|R| = p^s$. Note that we also have

$$|(N\phi(G))/\phi(G) : (R\phi(G))/\phi(G)| = |N\phi(G) : R\phi(G)| = \frac{|N\phi(G)|}{|R\phi(G)|}$$

Some fun calculations follow:

$$\begin{aligned}\frac{|N\phi(G)|}{|R\phi(G)|} &= \frac{\frac{|N||\phi(G)|}{|N \cap \phi(G)|}}{\frac{|R||\phi(G)|}{|R \cap \phi(G)|}} \\ &= \frac{\frac{|N|}{|R|}}{\frac{|N \cap \phi(G)|}{|R \cap \phi(G)|}}\end{aligned}$$

Note that the thing above divides $|N : R|$ but is coprime to p. It follows then that the whole fraction thing in the last step above is coprime to p, and so we are done. □

By the claim and assumption, we have:

$$(R\phi(G))/\phi(G) \overset{\text{char}}{\triangleleft} (N\phi(G))/\phi(G) \triangleleft G/\phi(G)$$

And so, by the lemma, we have:

$$(R\phi(G))/\phi(G) \triangleleft G/\phi(G) \implies R\phi(G) \triangleleft G$$

The Frattini Argument then gives us $G = R\phi(G) \cdot N_G(R) = N_G(R) \cdot R\phi(G)$ and so $G = N_G(R)\phi(G)$.So, if $N_G(R)$ were not the whole group G, there would exist a maximal M such that $N_G(R) \leq M \overset{\text{max}}{\leq} G$ (this is from the argument in the previous proof), so $G = N_G(R)\phi(G) \leq M$, which is a contradiction. □

In what follows G is taken to be a finite group unless stated otherwise.

Proposition 3.3.6
If $N \triangleleft G$ then $\phi(G) \leq \phi(G)$.

Proof. We need, for every $M \overset{\text{max}}{\leq} G$, $\phi(G) \leq M$. Now, note that for $\phi(G) \overset{\text{char}}{\leq} N \triangleleft G$, it follows that $\phi(N) \leq G$. Suppose $\phi(N) \not\leq M$. So there exists $M \overset{\text{max}}{<} G$. Then, as $\phi(N) \cdot M \leq G$, we would have $\phi(N)M = G \implies N = N \cap G = N \cap \phi(N) \cdot M$, but this must mean $N = \phi(N) \cdot (N \cap M)$.
Now, if $N \cap M \neq N$, there exists maximal $H \overset{\text{max}}{<} N$ such that $N \cap M \subseteq H$. As $\phi(N) \leq H$, we have $N \subseteq H$. This is a contradiction, so $N \cap M = N$, and so $N \subseteq M$.
However, this contradicts $\phi(N) \leq M$, which was our original assumption, and so we're done. □

The argument in the first paragraph above is commonly called the "modular law". I restate it below for reference:

Lemma 3.3.7 (Modular Law)
If $A, B, C \leq G$ such that $B \leq A$. Then $A \cap BC = B(A \cap C)$.

3.3.2 The Fitting Subgroup

Lemma 3.3.8
G finite, $A, B \triangleleft G$ such that A, B are both nilpotent. Then AB is still a nilpotent normal subgroup.

Proof. Proving that $AB \triangleleft G$ is omitted, as it should be easy to do. Now, we want to show that AB is nilpotent. Let p be a prime. Let $\text{Syl}_p(A) = \{R\}$ and $\text{Syl}_p(B) = \{S\}$.

Claim. $\text{Syl}_p(AB) = RS$.

Proof. First, note that $R \overset{\text{char}}{<} A \triangleleft G \implies R \triangleleft G$. Similarly, $S \overset{\text{char}}{<} B \leq G \implies S \triangleleft G$. These two facts together imply $RS \triangleleft G \implies RS \triangleleft AB$.

Now, $|RS| = \frac{|R||S|}{|R \cap S|} = p^v$ for some v. We compute $|AB : RS|$.

$$|AB : RS| = \frac{|A||B|/|A \cap B|}{|R||S|/|R \cap S|} = \frac{|A : R||B : S|}{|A \cap B : R \cap S|}$$

The above is known to be an integer, so the denominator must divide the numerator in the last expression. Co-primality then follows. □

The desired result follows immediately. □

We define fitting subgroups (with which we become more familiar in the homeworks) before moving on to an application.

Definition 3.3.9 (fitting subgroup)
By $F(G)$ we denote the fitting subgroup of G, which is defined as the greatest nilpotent normal subgroup of G.

3.4 An application to finite groups

Warning: I was late to class this day and my notes will almost certainly contain errors in this section. Sorry about this, and please let me know of any mistakes!

Proposition 3.4.1
Let G be finite and solvable, then $F(G) \geq \phi(G)$. Write $\overline{G} = G/\phi(G)$. Then $F(\overline{G}) = F(G)/\phi(G)$.

We state some additional facts below:

1. $F(\overline{G}) = \overline{R_1} \times \overline{R_2} \times \ldots$ such that $\overline{R_i}$ is an elementary abelian p_i subgroup.
2. $C_{\overline{G}}(F(\overline{G})) = F(\overline{G})$. The solvability assumption plays a role here.
3. $G/F(G) \cong \overline{G}/F(\overline{G})$, and the latter acts on $F(\overline{G})$ by conjugation faithfully.

Corollary 3.4.2 (a weak structure theorem)
$\overline{G}/F(\overline{G}) \cong G^* \leq \times_i GL(\mathbb{F}_{p_i})$.

Proposition 3.4.3
If $|R| = p^k$ where p is prime, then $\phi(R) = \langle R', r^p \mid r \in R \rangle$.

Proof. We first make the following claim, from which it follows that H in the given claim is a normal subgroup of R.

Claim. If $|R : H| = p$, then $H \overset{\text{max}}{<} R$ if and only if $|R : H| = p$.

Proof. We know that R is nilpotent, and that $|R : H| = p$. Recall that as R is nilpotent, if $H \lneq R$ then $N_R(H) \gneq H$, so as H maximal, we have $N_R(H) = H$, so $H \triangleleft R$. By the correspondence theorem, we have subgroups of R/H in correspondence with subgroups $H \leq L \leq R$, but as H maximal, $L = H$ or R. So, R/H only has trivial subgroup, which means $|R/H| = p$. □

Before, we proceed, another claim.

Claim. Given an index set I, and $G \triangleleft N_i$ for $i \in I$, $G/\bigcap_{i\in I} N_i \cong S \leq \times_{i\in I} G/N_i$.

Proof. Consider the map given by $g \mapsto (\ldots, gN_i, \ldots)$. The kernel of this map is given by $\ker = \bigcap_{i\in I} N_i$. Use the first isomorphism theorem to conclude the above result. □

From the above claim, we have

$$R/\phi(R) = R/\bigcap H \overset{\max}{\leq} RH \cong S \leq \times_{H \overset{\max}{\leq} R} R/H$$

which is elementary abelian. This gives us $\phi(R) \supseteq R' \ni r^p$. So, $\phi(R) \supseteq \langle R', r^p\rangle$. We write $L = \langle R', r^p\rangle$.

Claim. $\phi(R/L) = 1$.

Proof. Note $R/L \cong Z_p \times Z_p \ldots \times Z_p$. So, $\phi(R/L) \cong \phi(Z_p \times \ldots \times Z_p) = 1$. □

From the above claim, we can conclude that $\phi(R) \subseteq L$. With this, we are done. □

One last proposition before we move on to free groups:

Proposition 3.4.4
If G is finite and solvable, then $C_G(F(G)) \leq F(G)$.

I wasn't present in class for the lecture that covered the above two propositions, so I might have missed a few things. Sorry!

4 Free Groups

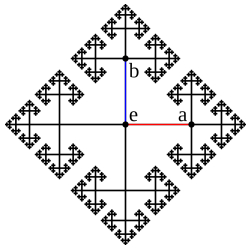

Figure 4.1. Cayley graph of $F\langle a, b\rangle$, the free group on 2 generators.

Free groups are foundational to many aspects of group theory. For instance every group is the factor group of some free group. Informally, free groups are those in which no relations hold between the elements besides the standard assumption that an element multiplied by its inverse is the identity.

A familiar example of a free group is the group of integers. Another way to think of free groups (which you may have seen before if, like me, you major in computer science)[1] is as words or strings over some set of generators or an alphabet.

Along the way, we will also encounter Cayley graphs, and will take a closer look at Cayley graphs of free groups.

4.1 Construction of free groups

We give two constructions of free groups, one using a set of generators, and another using a universal property.

- Given a set of symbols S, we can consider **words** over this set to be products of elements of $S \cup S^{-1}$. We say that S is a generating set for these words.
- We define the **empty word** ϵ to be the word with no symbols.
- If, in a word, an element of S lies immediately next to its inverse, the word may be simplified by omitting the pair.
- A word that cannot be simplified further is said to be **reduced**.

Definition 4.1.1 (free group)
Given a set S, the free group with generating set S, written F_S, is the group of reduced words over S. The group operation is given by concatenation.

We will sometimes write just F for the free group when the generating set S is understood. We will also sometimes write $F = \langle S\rangle$. It is easy to check that the above construction is indeed a group, and is left as an exercise.

[1] Look up regular languages and automatic groups if interested in applications of group theory in the theory of computation.

We now give an alternative definition of free groups, which at first seems quite different from the one above. With some reflection, however, it is easy to see that they are equivalent. This is the more category-theoretical point of view, and free objects in a category (whatever these animals are) appear everywhere in math.

Definition 4.1.2 (universal property definition of free groups)
$F = \langle S \rangle$ if for all $f : X \to A$, where A is any group, there exists a unique homomorphism $\overline{f} : F \to A$ a homomorphism that makes the following diagram commute:

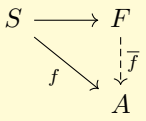

In the above diagram, the unnamed mapping denotes the inclusion from S into F. So you can alternatively say that $\overline{f} \restriction_S = f$.

Proposition 4.1.3 (free groups are a bit like vector spaces)
Two free groups are isomorphic if and only if the cardinalities of their generator sets are the same.

The reverse direction is fairly obvious and follows easily from the universal property construction of free groups. The forward direction, however, is not, and we might talk about it later, possibly in a week or two.

4.2 Cayley graphs

Cayley graphs are graphs associated with groups, and their construction as well as their properties are quite intuitive. Hence I won't be very formal in this section.

Suppose G is a group and S is a generating set. The Cayley graph $\Gamma = \Gamma(G; S)$ is a colored directed graph constructed as follows:

- Each element$g \in G$ is assigned a vertex: the vertex set $V(\Gamma)$ of Γ is identified with G.
- Each generator $s \in S$ is assigned a color, which we will (abusing notation) call s.
- For any $g \in G, s \in S$, the vertices corresponding to elements g, gs are joined by a directed edge (g, gs) of color s. This gives us the edge set $E(\Gamma)$.

Do we know whether a given graph is a Cayley graph? Indeed, given Γ, it is a Cayley graph if:

1. Γ is connected as an ordinary graph.
2. $G^* \leq \mathrm{Aut}(\Gamma)$ such that G^* acts regularly on $V(\Gamma)$.
3. For a given vertex, and for each color x, there is exactly one outgoing edge from the vertex colored by x; and similarly, exactly one incoming edge to the vertex colored by x.

4.3 Fun with free groups

The following example is important to keep in mind.

Example 4.3.1 (Cayley graphs of free groups)
Γ is a Cayley graph of F(X) if and only if the three conditions of Proposition 4.1.2 met, and if:

1. Γ is a tree as an ordinary graph
2. There exists no $x \in X$ such that $\mathscr{O}(x) = 2$

Now, what do Cayley graphs of subgroups of free groups look like?

The above question motivates the remainder of this section. Let $G = \langle X \rangle$, $H \leq G$. We want to get $\Gamma(H; Y = ?)$ from $\Gamma(G; X)$, where Y is the generating set of H (and is unknown to us at the moment).

First we require a tool from set theory: the (infamous) Zorn's lemma. We will then use group actions stuff to obtain a Cayley graph for $H \leq G$; next time we will show that it is indeed the Cayley graph of H. The argument given below is sometimes called a **contraction argument**.

Lemma 4.3.2 (Zorn's lemma)
Let S denote a partially ordered set with the property that for any totally ordered subset $T \subseteq S$, there exists some $u \in S$ such that for all $t \in T$, $t \leq u$. There there exists $m \in S$ such that m is maximal in S, i.e. $\nexists\, m \in S$ $m \leq s$ and $m \neq s$.

Now, look at the induced action of $H \leq G$ on $V(\Gamma)$ such that for each $g \in V(\Gamma)$ and for each $h \in H$, we have $h(g) := hg$. The orbits of H acting on $V(\Gamma) = G$, then, are precisely the right cosets of H.

Definition 4.3.3 (partial transversal)
T is a partial transversal if $T \leq \Gamma$ as a subgraph, such that for each $g \in G$, $|Hg \cap V(T)| \leq 1$, and T is connected (as an ordinary graph).

As we want to use Zorn's Lemma, our partially ordered set (poset, for short) will be the set of all partial transversals of Γ. Let us denote this set by T. It is a poset with respect to containment.

So Zorn's Lemma says that there exists a maximal partial transversal, T_0. We will show that T_0 is "complete" in the sense that it intersects each coset Hg for all $g \in G$ exactly once.

Claim. For every $g \in G$, we have $|Hg \cap V(T_0)| = 1$.

Note that $h(T_0)$ is a maximal transversal too, as left multiplication is an automorphism. We thus obtain a graph corresponding to H; all that remains is construction of a colored and directed edges of the new graph. You can check that this subgraph is indeed a Cayley graph for H.

The above discussion essentially is Schreier's original proof of the Nielsen-Schreier theorem, which is stated below for completeness. Another proof, which uses some algebraic topology (fundamental groups and covering spaces), can be found here. Yet another proof can be found here.

Theorem 4.3.4: Nielsen-Schreier theorem
Every subgroup of a free group is free.

What is cool is that if we know the index of H in G in the above discussion, then we know the size of the generating set of H! This is called Schreier's index formula, and we prove it below:

Proposition 4.3.5 (Schreier's index formula)
Let $|X| = n < \infty$, and let G be the free group generated by X. If $H \leq G$ such that $|G : H| = k < \infty$, then H is free with respect to some generating set Y, where $|Y| = (n-1)k + 1$.

Proof. Note that $\deg(h) = 2|Y|$ for all $h \in H$. This follows directly from how the Cayley graph is constructed. By T we will denote our maximum partial transversal as in the discussion on the preceding page. Then $V(T) = k$, since it intersects each of our cosets in exactly 1 vertex. Moreover, we have $E(T) = k - 1$.

Now, in our original graph, the degree of each point p was $2n$. Then, should the total number of edges incident to any one of our cosets be $2n$? No, since here we are double counting the number of edges incident to p that are connected to points within the same coset as p.

We divide the obtained expression divided by 2 (since we don't add inverses to the size of a generating set), and so we're done. □

Recall the universal-property definition of a free group. It is easy then to see that free groups have lots of homomorphisms, and hence, lots of normal subgroups; what is more, the next theorem states that they not only have several normal subgroups, but have several normal subgroups of finite index.

Theorem 4.3.6: free groups are residually finite
If G is free, $e \neq a \in F$, then there exists $N \triangleleft F$ such that:
1. *$a \notin N$*
2. *$|F : N|$ is finite*

Proof. Let p be a prime. Let $a = x_{i_1}^{m_1} \dots x_{i_r}^{m_r}$, where $x_{i_j} \in X$, be a word in reduced form (i.e. $i_u \neq i_{u+1}$ for all u, and $m_j \neq 0$ for all j). Pick n such that $p^n \nmid m_1 m_2 \dots m_r$.

Consider $(r+1) \times (r+1)$ matrices over $\mathbb{Z}_{p^n}$. Let I denote the $(r+1) \times (r+1)$ identity matrix, and $E_{s,s+1}$ be the matrix whose $(s, s+1)^{\text{th}}$ entry is 1, all others are 0.

Let $G := \{(r+1) \times (r+1)$ upper unitriangular matrices over $\mathbb{Z}_{p^n}\}$. For each x_j having nonzero exponent in a, we define g_j as follows:

$$g_j = \prod_{\mu : i_\mu = j} (I + E_{\mu,\mu+1})$$

We want to show that a does not lie in the kernel of the above homomorphism.
Claim. $g_{i_1}^{m_1} \dots g_{i_r}^{m_r} \neq I$.

Proof. First, note the following for our special matrices: $E_{s,t}E_{v,z} = E_{s,z}$ if $t = v$, and $E_{s,t}E_{v,z} = 0$ if $t \neq v$. Now, consider our situation: look at $E_{\mu,\mu+1}E_{\nu,\nu+1}$. If $\mu + 1 = \nu$, then $i_\mu = j = i_\nu = i_{\mu+1}$ which is a contradiction. From this, we can see that matrices of the form $I + E_{\mu,\mu+1}$ commute. More precisely:

$$(I + E_{\mu,\mu+1})(I + E_{\nu,\nu+1}) = I + E_{\mu,\mu+1} + E_{\nu,\nu+1}$$

Thus, we can write:

$$g_{i_1}^{m_1} \dots g_{i_r}^{m_r} = \prod_{\mu_1 : i_{\mu_1} = i_1} (I + m_1 E_{\mu,\mu+1})(\text{some blah}) = I + \sum_{s,t} \lambda E_{s,t}$$

□

The proof can be completed easily. □

Why are we doing this? Because this allow the rank (i.e. the size of the generating set) of a free group to be well defined. But before doing that, we talk about a property that is shared by several groups; in particular by finitely generated free groups.

Definition 4.3.7 (hopfian group)
G is a hopfian group if for all $1 \neq N \triangleleft G$, $G/N \ncong G$.

Example 4.3.8
1. Any finite group is hopfian.
2. Let $G = H^{(1)} \times H^{(2)} \times \dots$, then if we take $N = H^{(1)}$, the quotient group is still an infinite group that is isomorphic to to G.

Being hopfian is a very nice property, as is evident from the next lemma, whose proof should be filled out as an exercise.

Lemma 4.3.9
G is hopfian if and only if all surjective endomorphisms of G are automorphisms.

Theorem 4.3.10
If G is finitely generated and it is residually finite, then G is hopfian.

Proof. We use the definition of hopfian groups given by Lemma 4.3.9. Suppose the contrary, i.e. there exists homomorphism $\epsilon : G \twoheadrightarrow G$ such that $\exists\ e \neq a \in \ker \epsilon$. As G is finitely generated, the number of homomorphisms from G to itself is finite, as homomorphisms are uniquely determined by the images of the generators.

Now, consider the following set $H = \{\text{all homomorphisms } f : G \to G/N\}$ where $N \triangleleft G$ such that $a \notin N$ and $|G : N|$ is finite. Such an N exists as G is residually finite. Again, as G is finitely generated, H is finite.

Suppose $\eta \in H$. Then we have $\eta \circ \epsilon : G \to G/N$. It follows that $\eta \circ \epsilon \in H$.

Claim. If $\eta_1, \eta_2 \in H$, then $\eta_1 \neq \eta_2 \implies \eta_1 \circ \epsilon \neq \eta_2 \circ \epsilon$.

Proof. This follows from the fact that ϵ is surjective. □

Now, we are done. Why? Because $\eta \circ \epsilon(a) = e$, as $a \in \ker \epsilon$. So, for all homomorphisms in H (as it is finite and composition by ϵ preserves membership in H) from $G \to G/N$, $a \mapsto e$. If we can find a homomorphism that does otherwise, then we have a contradiction, as we are done. Indeed, the natural homomorphism mapping $a \mapsto aN$ does not map a to the identity. □

Corollary 4.3.11
If F is free on X, $|X| < \infty$, then F is hopfian.

And finally, we have:

Corollary 4.3.12 (rank of a free group is well defined)
Let F be free on X, $|X| = n < \infty$. Let $Z \subset F$ such that $\langle Z \rangle = F$. If $|Z| \leq n$, then $|Z| = n$ and F is free on Z too.

Proof. There exists $f : X \twoheadrightarrow Z$, and so, by definition of free groups, there exists $\overline{f} : F \to F$ such that $\overline{f} \restriction_X = f$. Thus, $\mathrm{Im}\overline{f} = F$, i.e. $\overline{f}$ is a surjective endormorphism. By Corollary 4.3.9, $\overline{f}$ is an automorphism of F. □

What happens if, in the above corollary, you have one infinite group of generators and one finite group of generators? Take surjective map, and you run into a contradiction!

What if both are ordinals? Suppose, as an example, one has generating set of size $\mathcal{C}$, and the other ω. This is unsurprisingly not possible, as on one hand, we must have our group to have cardinality $\mathcal{C}$, and simultaneously must also have cardinality ω.

With this, we end our chapter on free groups.

Part II

Representation Theory of Finite Groups

5 Basic Notions

The goal of group representation theory is to study groups via their actions on vector spaces. Consideration of groups acting on sets leads to such important results as the Sylow theorems. By studying actions on vector spaces even more detailed information about a group can be obtained. This is the subject of representation theory.

Important. In these notes, only representations over the complex field $\mathbb{C}$ will be considered. Moreover, all groups are assumed to be finite unless stated otherwise.

5.1 What is a representation?

By $\mathrm{GL}(V)$ we denote the group of all nonsingular linear transformations of a vector space V over $\mathbb{C}$. Alternatively, it is the set of all bijective linear transformations on V with respect to composition.

Definition 5.1.1 (representation of a group)
Let G be a group and V a vector space (linear space) of finite dimension over $\mathbb{C}$. A homomorphism $\phi : G \to \mathrm{GL}(V)$ will be called a representation of G on V.

Using basic linear algebra, it is clear that Definition 5.1.1 is equivalent to:

Definition 5.1.2 (representation of a group)
Let G be a group and n a positive integer. A homomorphism $\phi : G \to \mathrm{GL}_n(\mathbb{C})$ will be called a representation of G of degree n.

where $\mathrm{GL}_n(\mathbb{C})$ denotes the group of all nonsingular $n \times n$ matrices over $\mathbb{C}$.

Why two definitions? Because sometimes, as we will see, working concretely with some special classes of matrices leads to some really deep results regarding abstract groups, e.g. Burnside's theorems, which we will learn about later.

I sometimes refer to the abstract definition of a representation (Definition 5.1.1) as a *representation*, and the concrete one (Definition 5.1.2) as a *matrix representation* to keep things clear.

Here are three examples of degree one representations:

Example 5.1.3
We look at representations of cyclic groups over $\mathbb{C}$. First, note that $\mathrm{GL}_1(C) \cong C^*$. The isomorphism is given by $f \mapsto f(1)$. Think of this visually using polar representations of complex numbers. Then, the following are representations:

1. $\mathbb{Z}/2\mathbb{Z} \to C^*$ given by $m \mapsto (-1)^m$
2. $\mathbb{Z}/4\mathbb{Z} \to C^*$ given by $m \mapsto i^m$
3. More generally, $\mathbb{Z}/n\mathbb{Z} \to C^*$ given by $m \mapsto e^{(2\pi i m/n)}$

Definition 5.1.4 (G-invariant subspaces)
Let $\phi : G \to \mathrm{GL}(V)$ be a representation of G, and let W be a subspace in V (written $W \leq V$); W is called G-invariant if $\phi(g)(W) \subseteq W$ (or, equivalently, $\phi(g)(W) = W$) holds for every $g \in G$.

Naturally, given a representation (V, ρ) of a group G, if $W \subseteq V$ is a G-invariant subspace, then we have a natural representation (W, ρ) of G whose degree will be the dimension of W. Thus, if a representation V has a G-invariant subspace, then you can, in a way, obtain a "smaller" representation of G.

Those representations which cannot be *reduced* in the above sense have a special name:

Definition 5.1.5 (irreducible representations)
A representation $\phi : G \to \mathrm{GL}(V)$ is irreducible if $\{0\}$ and $V \neq \{0\}$ are the only G-invariant subspaces.

Irreducible representations are to representation theory what prime numbers are to the integers. We will learn more about this when we see Maschke's Theorem.

Example 5.1.6
If the vector space V has dimension 1, then any representation of G on V is irreducible.

5.2 Maps between representations

More generally, we can consider *morphisms* between representations, and if a morphism is bijective, then we say that two representations are isomorphic/equivalent. I also came across the term *intertwining operators* for such maps.

Definition 5.2.1 (equivalent representations)
Let $\phi_1 : G \to \mathrm{GL}(V_1)$ and $\phi_2 : G \to \mathrm{GL}(V_2)$ be representations of the same group G. They are called equivalent (written as $\phi_1 \sim \phi_2$) if there is some $f \in \mathrm{Hom}_{\mathbb{C}}(V_1, V_2)$ such that f is bijective and for every $g \in G$, $\phi_2(g) = f \circ \phi_1(g) \circ f^{-1}$ (or, equivalently, $f \circ \phi_1(g) = \phi_2(g) \circ f$) holds.

Essentially what the above definition is saying is that given ϕ_1 and ϕ_2 two representations of G over V_1 and V_2 respectively, then they are equivalent if there's some isomorphism $f : V_1 \to V_2$ such that $\phi_2(g) = f \circ \phi_1(g) \circ f^{-1}$. In pictures, we have that the following diagram

$$\begin{array}{ccc} V_1 & \xrightarrow{\phi_1(g)} & V_1 \\ {\scriptstyle f}\downarrow & & \downarrow{\scriptstyle f} \\ V_2 & \xrightarrow{\phi_2(g)} & V_2 \end{array}$$

commutes for all $g \in G$.

Intuitively, this is equivalent to saying that "two representations are equivalent if and only if they describe the same representation but in different bases". The map f above is nothing more than a change of basis transformation.

Proposition 5.2.2
If ϕ_1 and ϕ_2 are equivalent representations of G, then ϕ_1 is irreducible $\iff$ ϕ_2 is irreducible.

I omit the proof as it's very straightforward.

The following table helps me remember everything that we have learned up until this point:

Groups	Vector Spaces	Representations
Subgroup	Subspace	G-invariant subspace
Simple group	1-dimensional subspace	Irreducible representation
Direct product	Direct sum	Direct sum
Isomorphism	Isomorphism	Equivalent representations

5.3 Maschke's Theorem and Schur's Lemma

We begin this section with a motivating example.

Example 5.3.1 (Representations of S_3)
We take $G = S_3$ the symmetric group on 3 elements, and $n = \dim_{\mathbb{C}} V = 3$. Take $\psi : a \mapsto \psi(a) \in \mathbb{C}^{3\times 3}$ such that

$$(\psi(a))_{i,j} = \begin{cases} 1 & a(j) = i \\ 0 & \text{otherwise} \end{cases}$$

where i, j denote the labels of the three elements that S_3 permutes. So we have:

$$a = (1,2,3) \mapsto \begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}$$

What's an invariant subspace of this representation? Indeed, if you look at the action of the above matrix on an arbitrary vector, it simply permutes the order of the components of the vector. We have the following G-invariant subspace:

$$W_1 := \left\{ \begin{bmatrix} \lambda \\ \lambda \\ \lambda \end{bmatrix} \mid \lambda \in \mathbb{C} \right\}$$

Are there any other G-invariant subspaces? Let's look at the orthogonal complement of W_1.

$$W_2 := \left\{ \alpha \begin{bmatrix} 1 \\ 1 \\ -2 \end{bmatrix} + \beta \begin{bmatrix} 1 \\ -2 \\ 1 \end{bmatrix} \mid \alpha, \beta \in \mathbb{C} \right\} = \left\{ \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \alpha_3 \end{bmatrix} \in \mathbb{C}^3 \mid \sum_i \alpha_i = 0 \right\}$$

Observe that $W_1 \perp W_2$, and that $G = W_1 \oplus W_2 = \mathbb{C}^3$. We can thus obtain two representations of G on W_1 and W_2, say ψ_1 and ψ_2 respectively. We say that $\psi = \psi_1 \oplus \psi_2$. We now try to obtain a matrix representation for $\psi_2(a)$ which arises from the "restriction" of the original $\psi(a)$ to W_2.

$$\begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} 1 \\ 1 \\ -2 \end{bmatrix} = (-1) \begin{bmatrix} 1 \\ 1 \\ -2 \end{bmatrix} + (-1) \begin{bmatrix} 1 \\ -2 \\ 1 \end{bmatrix}$$

$$\begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} 1 \\ -2 \\ 1 \end{bmatrix} = (1) \begin{bmatrix} 1 \\ 1 \\ -2 \end{bmatrix}$$

Thus, we obtain:

$$\begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \rightsquigarrow \begin{bmatrix} -1 & 1 \\ -1 & 0 \end{bmatrix}$$

Important. Above, we decomposed a degree 3 representation of S_3 into two irreducible representations of degree 1 and 2 respectively. Working with an infinite group can change the rather nice situation we have in the above example. We will see more of this in the homework.

A natural question to ask then is whether all representations can be decomposed into the direct sum of irreducible representations? Surprisingly, the answer is yes! And this leads us to Maschke's Theorem.

Theorem 5.3.2: Maschke's Theorem
Let $\phi : G \to \mathrm{GL}(V)$ be a representation of G and assume that W is a G-invariant subspace in V. Then there exists a G-invariant direct complement to W in V, i.e. a G-invariant subspace U in V such that $V = W \oplus U$.

Proof. Let π be a projection on V such that $\mathrm{Im}\ \pi = W$ (a linear transformation π is called a *projection* if $\pi^2 = \pi$; it is easy to show that $V = \mathrm{Im}\ \pi \oplus \mathrm{Ker}\ \pi$ holds for any projection). Set

$$\pi^* := \frac{1}{|G|} \sum_{h \in G} \phi(h)^{-1} \pi \phi(h)\,.$$

We show that π^* is also a projection with $\mathrm{Im}\ \pi^* = W$, and that $U := \mathrm{Ker}\ \pi^*$ is G-invariant. For assume $w \in W$, then $\pi^*(w) = \frac{1}{|G|} \sum_{h \in G} \phi(h)^{-1} \pi \phi(h)(w)$.

Now, as W is G-invariant, then we have $\phi(h)(w) = w$, and as π is a projection, we have $\pi(w) = w$, and again, by G-invariance of W, we have $\phi(h)^{-1}(w) = w$. Thus, we have $\pi^*(w) = w$. It follows then that $W \subseteq \mathrm{Im}\ \pi^*$. On the other hand, let $v \in V$ and $h \in G$. Then, $\pi\phi(h)^{-1}(v) \in W$, and as W is G-invariant, we have $\phi(h)\pi\phi(h)^{-1}(v) \in W$. It follows then that $\pi^*(v) \in W$, and so $\mathrm{Im}\ \pi^* \subseteq W$. We conclude that $\mathrm{Im}\ \pi^* = W$.

Next, pick arbitrary $g \in G$. Then we have:

$$\begin{aligned}\phi(g)\pi^*\phi(g)^{-1} &= \phi(g)\left(\frac{1}{|G|}\sum_{h\in G}\phi(h)^{-1}\pi\phi(h)\right)\phi(g)^{-1}\\ &= \frac{1}{|G|}\sum_{h\in G}\phi(hg^{-1})^{-1}\pi\phi(hg^{-1})\\ &= \frac{1}{|G|}\sum_{h\in G}\phi(h)^{-1}\pi\phi(h)\\ &= \pi^*\end{aligned}$$

and so $\phi(g)\pi^* = \pi^*\phi(g)$ holds for each $g \in G$. At last assume $u \in U$ and $g \in G$; then $\pi^*\phi(g)(u) = \phi(g)\pi^*(u) = \phi(g)(0) = 0$ i.e $\phi(g)(u) \in U$. This is because π^* sends $\phi(g)(u)$ to 0, but $\ker \pi^* = U$. Thus U is G-invariant. □

Important. Maschke's Theorem goes horribly wrong in the case of infinite groups, as we will see on the homework.

Before moving on, we quickly define the direct sum of representations.

Definition 5.3.3
Let $\phi_i : G \to \mathrm{GL}(V_i)$ be representations of G $(i = 1, 2 \ldots, k)$. The (direct) sum of these representations is $\phi = \phi_1 + \phi_2 + \cdots + \phi_k$ where $\phi : G \to \mathrm{GL}(\bigoplus_{i=1}^k V_i)$ is defined by

$$\phi(g)(v_1\,,\, v_2\,, \ldots,\, v_k) := (\phi_1(g)(v_1), \ldots, \phi_k(g)(v_k))\,.$$

(It follows that $\phi_1 + \phi_2 + \cdots + \phi_k$ is also a representation of G.)

The real power of Maschke's Theorem is the following corollary.

Corollary 5.3.4
Every representation is equivalent to the direct sum of irreducible representations.

and now it should make sense why we earlier called irreducible representations to be the prime numbers of representation theory.

The other basic theorem in this section is Schur's Lemma, which is also of great use in representation theory.

Lemma 5.3.5 (Schur's Lemma)
Let $\phi_1 : G \to \mathrm{GL}(V_1)$ and $\phi_2 : G \to \mathrm{GL}(V_2)$ be irreducible representations of G. Assume that $f \in \mathrm{Hom}_{\mathbb{C}}(V_1, V_2)$ such that for every $g \in G$, $f \circ \phi_1(g) = \phi_2(g) \circ f$ holds. Then

1. *If ϕ_1 and ϕ_2 are not equivalent then $f = 0$.*
2. *If $\phi_1 = \phi_2 := \phi$ then $f = \lambda I$, a scalar multiple of the identity.*

Proof. (0) We claim that $\ker f$ and $\mathrm{Im}\ f$ are G-invariant subspaces in V_1 and V_2 respectively. Indeed, suppose $v_1 \in \ker f$. Then $f \circ \phi_1(g)(v_1) = \phi_2(g) \circ f(v_1) =$

$0 \implies \phi_1(g)(v_1) \in \ker f$, so it is G-invariant. Similarly suppose $f(v_2) \in \operatorname{Im} f$. Then $\phi_2(g) \circ f(v_2) = f \circ \phi_1(g)(v_2) \in \operatorname{Im} f$ just because we are composing with f.

As ϕ_1, ϕ_2 are irreducible the following cases can only occur: (i) $\operatorname{Ker} f = 0$, $\operatorname{Im} f = V_2$; or (ii) $\operatorname{Ker} f = V_1$ and $\operatorname{Im} f = 0$.

1. Under the assumption in (1) f would be bijective in case (i), implying that ϕ_1 and ϕ_2 are equivalent; a contradiction.
2. To prove (2) let w be an eigenvector of f, i.e. $w \neq 0$ and $f(w) = \beta w$ with some $0 \neq \beta \in \mathbb{C}$. We can do so as $\mathbb{C}$ is algebraically closed. As a scalar transformation commutes with all linear transformations, $f^* := f - \beta I$ also satisfies the assumptions of the lemma. Since $0 \neq w \in \ker f^*$, our f^* must be zero by (ii); thus $f = \beta I$.

□

Important. Note that only the second part of Schur's Lemma relies on the fact that our underlying field is $\mathbb{C}$; more specifically, it relies on the fact that $\mathbb{C}$ is algebraically closed and hence guarantees the existence of an eigenvector. The first part holds for representations over arbitrary fields.

Remark. Recall that we have taken G finite. Assume that $\phi_1 : G \to \mathrm{GL}(V_1)$ and $\phi_2 : G \to \mathrm{GL}(V_2)$ are irreducible representations of G and $f_0 \in \mathrm{Hom}_{\mathbb{C}}(V_1, V_2)$ is arbitrary. Then $f := \sum_{h \in G} (\phi_2(h))^{-1} f_0 (\phi_1(h))$ meets the requirements of 5.2.4.

For let $g \in G$ then:

$$\begin{aligned}
(g^{\phi_1})^{-1} f g^{\phi_2} &= (g^{\phi_1})^{-1} \left(\sum_{h \in G} (h^{\phi_1})^{-1} f_0 h^{\phi_2} \right) g^{\phi_2} \\
&= \sum_{h \in G} (g^{\phi_1})^{-1} (h^{\phi_1})^{-1} f_0 h^{\phi_2} g^{\phi_2} = \sum_{h \in G} ((hg)^{\phi_1})^{-1} f_0 (hg)^{\phi_2} \\
&= \sum_{h_* \in G} (h_*^{\phi_1})^{-1} f_0 h_*^{\phi_2} = f\,.
\end{aligned}$$

6 Characters and Burnside's Theorems

We quickly set up the notation that we will be using during this section.

1. For a matrix M we'll denote by $M_{i,j}$ the (i,j)-th element of M.
2. For an $n \times n$ matrix A we denote by $\mathrm{Tr}(A)$ the *trace* of A, i.e. the sum of its elements in the main diagonal. (As known from linear algebra, $\mathrm{Tr}(A)$ equals the sum of the eigenvalues of A. It is also of great importance that $\mathrm{Tr}(M^{-1}AM) = \mathrm{Tr}(A)$ holds for any A and M.)
3. Let $E^{(k,\ell)}$ denote the matrix (of a given size) with its (k,ℓ)-th element 1 and all others 0; i.e.
$$E^{(k,\ell)}_{i,j} = \begin{cases} 1 & \text{if } i = k \text{ and } j = \ell \\ 0 & \text{otherwise} \end{cases}$$
4. The standard notation
$$\delta_{i,j} = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{if } i \neq j \end{cases}$$
will also be used.

6.1 What is a character?

Characters will be of great use, as we will see later in the course. To every representation V of G we will attach a so-called character. It will turn out that the characters of irreducible representations of V will determine the representation V completely.

Recall Schur's Lemma. The following proposition will later motivate the definition of a character.

Proposition 6.1.1
Assume that $\phi_1 : G \to \mathrm{GL}_{n_1}(\mathbb{C})$, $\phi_2 : G \to \mathrm{GL}_{n_2}(\mathbb{C})$ are irreducible representations of G. Then

1. *If ϕ_1 and ϕ_2 are not equivalent then $\sum_{h\in G} \phi_1(h)^{-1}{}_{a,b} \cdot \phi_2(h)_{c,d} = 0$ for all $1 \le a, b \le n_1$ and $1 \le c, d \le n_2$.*
2. *If $\phi_1 = \phi_2 := \phi$ then $\sum_{h\in G} \phi_1(h)^{-1}{}_{a,b} \cdot \phi_2(h)_{c,d} = \delta_{a,d}\delta_{b,c}\frac{|G|}{n}$ for all $1 \le a, b, c, d \le n$ $(n := n_1 = n_2)$.*

Proof. We shall make use of matrix representations. Take $f_0 := E^{(b,c)}$ of size $n_1 \times n_2$; this is the same f in the remark after Schur's Lemma in §5. We then compute the (a,d)-th

element of $\sum_{h\in G}\phi_1(h)^{-1}E^{(b,c)}\phi_2(h)$:

$$\left(\sum_{h\in G}\phi_1(h)^{-1}E^{(b,c)}\phi_2(h)\right)_{a,d}=\sum_{h\in G}(\phi_1(h)^{-1}E^{(b,c)}\phi_2(h))_{a,d}=$$

(i) $$\sum_{h\in G}\sum_{i=1}^{n_1}\phi_1(h)^{-1}_{a,i}\cdot(E^{(b,c)}\phi_2(h))_{i,d}=\sum_{h\in G}\sum_{i=1}^{n_1}\phi_1(h^{-1})_{a,i}\cdot\sum_{j=1}^{n_2}E^{(b,c)}{}_{i,j}\phi_2(h)_{j,d}=$$

$$\sum_{h\in G}\sum_{i=1}^{n_1}\phi_1(h^{-1})_{a,i}\delta_{i,b}\phi_2(h)_{c,d}=\sum_{h\in G}\phi_1(h^{-1})_{a,b}\phi_2(h)_{c,d}\,.$$

Thus (1) is proved by Schur's Lemma. As for (2), we know again from Schur's Lemma that $\sum_{h\in G}(h^{-1})^{\phi}E^{(b,c)}h^{\phi}=\lambda I$. Since

$$\lambda=\frac{1}{n}\mathrm{Tr}(\lambda I)=\frac{1}{n}\mathrm{Tr}\left(\sum_{h\in G}\phi(h^{-1})E^{(b,c)}\phi(h)\right)=$$

$$\frac{1}{n}\sum_{h\in G}\mathrm{Tr}(\phi(h^{-1})E^{(b,c)}\phi(h))=\frac{1}{n}\sum_{h\in G}\mathrm{Tr}(E^{(b,c)})=\frac{1}{n}|G|\delta_{b,c}\,,$$

we have by (i)

$$\frac{1}{n}|G|\delta_{a,d}\delta_{b,c}=\delta_{a,d}\lambda=\left(\sum_{h\in G}\phi(h^{-1})E^{(b,c)}\phi(h)\right)_{a,d}=$$

$$\sum_{h\in G}\phi(h^{-1})_{a,b}\phi(h)_{c,d}\,.$$

□

The role played by the trace of a matrix above is crucial, and when dealing with representations, the trace of the representation of a group element is given by the character associated with that representation.

Definition 6.1.2 (character of a representation)
Let $\phi: G\to \mathrm{GL}_n(\mathbb{C})$ be a representation of G. The character of ϕ is the function $\chi_\phi: G\to\mathbb{C}$ defined by $\chi_\phi(g):=\mathrm{Tr}\,(\phi(g))$.

Proposition 6.1.3 (properties of characters)
Let χ be the character of some representation ϕ of G. Then for every $g,h\in G$

1. $\chi(h^{-1}gh)=\chi(g)$;
2. *equivalent representations have the same character;*
3. $\chi(g^{-1})=\overline{\chi(g)}$.

Proof. Note that (1) follows from the remark at the beginning of this section, (2) is an immediate consequence of (1). To prove (3) let $g\in G$. Assume that the minimal polynomial of $\phi(g)$ is $p(x)$; then $p(x)$ divides x^k-1 where $k=o(g)$. Thus the eigenvalues of $\phi(g)$ are k-th roots of unity. It follows that the characteristic polynomial of $\phi(g)$ is

$c(x) = \det(I_n - \phi(g)) = (x-\varepsilon_1)(x-\varepsilon_2)\dots(x-\varepsilon_n)$ (where ${\varepsilon_i}^k = 1$, $n =$ the degree of ϕ). Thus the characteristic polynomial of $(\phi(g))^{-1}$ is $x^n c(\frac{1}{x}) = (x-\frac{1}{\varepsilon_1})(x-\frac{1}{\varepsilon_2})\dots(x-\frac{1}{\varepsilon_n})$, hence $\chi(g^{-1}) = \mathrm{Tr}((\phi(g))^{-1}) = \frac{1}{\varepsilon_1}+\frac{1}{\varepsilon_2}+\dots+\frac{1}{\varepsilon_n} = \overline{\varepsilon_1}+\overline{\varepsilon_2}+\dots+\overline{\varepsilon_n} = \overline{\varepsilon_1+\varepsilon_2+\dots+\varepsilon_n} = \overline{\mathrm{Tr}(\phi(g))} = \overline{\chi(g)}$. □

6.2 Orthogonality relations

The properties of characters naturally lead to the following"

Corollary 6.2.1 (The First Orthogonality Relation)
Assume that ϕ_1, ϕ_2 are irreducible representations of G and their character is χ_1 and χ_2, respectively. Then

$$\frac{1}{|G|}\sum_{g\in G}\overline{\chi_1(g)}\chi_2(g) = \begin{cases} 0 & \text{if } \phi_1 \nsim \phi_2 \\ 1 & \text{if } \phi_1 \sim \phi_2 \end{cases}$$

Proof. Let n_1 and n_2 denote the degree of ϕ_1 and ϕ_2.

$$\frac{1}{|G|}\sum_{g\in G}\overline{\chi_1(g)}\chi_2(g) = \frac{1}{|G|}\sum_{g\in G}\chi_1(g^{-1})\chi_2(g) =$$
$$\frac{1}{|G|}\sum_{g\in G}(\sum_{i=1}^{n_1}\phi_1(g^{-1})_{i,i})(\sum_{j=1}^{n_2}\phi_2(g)_{j,j}) = \frac{1}{|G|}\sum_{i,j}(\sum_{g\in G}\phi_1(g^{-1})_{i,i}\phi_2(g)_{j,j}).$$

Assume that $\phi_1 \nsim \phi_2$, then for every i, j the sum $\sum_{g\in G}\phi_1(g^{-1})_{i,i}\phi_2(g)_{j,j}$ is zero by 6.1.1 (1). If $\phi_1 \sim \phi_2$ then we can assume that $\phi_1 = \phi_2 := \phi$ (with $\chi_1 = \chi_2 := \chi$, $n_1 = n_2 := n$), and the assertion will follow from 6.1.1 (2):

$$\frac{1}{|G|}\sum_{g\in G}\overline{\chi(g)}\chi(g) = \frac{1}{|G|}\sum_{i,j}(\sum_{g\in G}\phi(g^{-1})_{i,i}\phi(g)_{j,j}) =$$
$$\frac{1}{|G|}\sum_{i,j}\frac{|G|}{n}\delta_{i,j} = \frac{1}{|G|}\sum_{i=1}^{n}\frac{|G|}{n} = \frac{1}{|G|}\cdot\frac{|G|}{n}\cdot n = 1.$$

□

Definition 6.2.2
Let $|G| = n$ and take a vector space V of dimension n such that the elements of some basis $\mathcal{B}$ are labelled by the elements of G: $\mathcal{B} = \{\underline{b}_h \mid h \in G\}$. We define the regular representation R of G as $R : G \to \mathrm{GL}(V)$ such that for every $g \in G$ and $\underline{b}_h \in \mathcal{B}$ $(\underline{b}_h)g^R := \underline{b}_{hg}$.

Thus, given any group, we can obtain a regular representation. This, together with Cayley's Theorem, makes the study of representations of symmetric groups very important.

Lemma 6.2.3
R is a representation of G. Let ϱ denote its character; then

$$\varrho(g) = \begin{cases} |G| & \text{if } g = 1 \\ 0 & \text{if } g \neq 1 \end{cases}$$

Proof. The first statement is clear since any linear transformation is uniquely determined by its action on a basis. To calculate ϱ consider R in matrix form; each $R(g)$ being a permutation matrix $\varrho(g)$ equals the number of 1-s in the main diagonal of $R(g)$, i.e. the number of basis vectors $\underline{b}_h$ satisfying $\underline{b}_{hg} = \underline{b}_h$, i.e. the number of group elements h satisfying $hg = h$. □

We know by Maschke's Theorem that every representation is equivalent to a sum of irreducible representations; therefore $R \sim \phi_1 + \phi_2 + \ldots + \phi_m$ with ϕ_i irreducible. It is easy to see (e.g. by taking the matrix form of the representations) that the character of a sum equals the sum of the corresponding characters, thus $\varrho = \chi_{\phi_1} + \chi_{\phi_2} + \ldots + \chi_{\phi_m}$. Since equivalent representations have the same character this can be written (by collecting the equivalent ϕ_j-s) as

$$\varrho = m_1\chi_1 + m_2\chi_2 + \ldots + m_c\chi_c \qquad \text{reg}$$

where $\chi_1, \chi_2, \ldots, \chi_c$ already belong to pairwise nonequivalent (irreducible) representations, and $m_1, m_2, \ldots, m_c$ are positive integers. To be able to compute these m_i-s the following definition will be useful.

Definition 6.2.4 (inner product)
Let f_1, f_2 be mappings (functions) from G to $\mathbb{C}$. Their inner product is a scalar defined as

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

Proposition 6.2.5
The inner product is bilinear, i.e.

1. $(f_1 + f_2, f_3) = (f_1, f_3) + (f_2, f_3)$;
2. $(f_1, f_2 + f_3) = (f_1, f_2) + (f_1, f_3)$;
3. $(f_1, \alpha f_2) = \alpha(f_1, f_2)$ *(for each $\alpha \in \mathbb{C}$);*
4. $(\alpha f_1, f_2) = \overline{\alpha}(f_1, f_2)$ *(for each $\alpha \in \mathbb{C}$).*

In terms of the inner product the first orthogonality relation can be formulated as follows.

Assume that ϕ_1, ϕ_2 are irreducible representations of G and their character is χ_1 and χ_2, respectively. Then

$$(\chi_1, \chi_2) = \begin{cases} 1 & \text{if } \phi_1 \sim \phi_2 \\ 0 & \text{if } \phi_1 \nsim \phi_2. \end{cases} \qquad \text{ortog I.}$$

Proposition 6.2.6
Keeping the notation of (reg), $m_i = \chi_i(1)$, *i.e. the degree of the representation belonging to* χ_i.

Proof. Let's use 5.3.8 and (ortog I.):

$$(\chi_i, \varrho) = (\chi_i, \sum_{k=1}^{c} m_k \chi_k) = \sum_{k=1}^{c} (\chi_i, m_k \chi_k) = \sum_{k=1}^{c} m_k (\chi_i, \chi_k) = m_i .$$

On the other hand, by 5.3.6, $(\chi_i, \varrho) = \frac{1}{|G|}\chi_i(1) \cdot \varrho(1) = \chi_i(1)$.

□

6.3 Algebraic Integers

We quickly define a few algebraic structures that will come in later when we talk about algebraic integers.

Definition 6.3.1 ($\mathbb{F}$-algebra)
Let $\mathbb{F}$ be a field and A a set with three operations $+, \cdot$ and $.$ (multiplication of elements of A by elements of $\mathbb{F}$). $(A, +, \cdot, .)$ is called an $\mathbb{F}$–algebra (or an algebra over $\mathbb{F}$) if $(A, +, .)$ is a vector space over $\mathbb{F}$, $(A, +, \cdot)$ is a ring, and for any $a, b \in \mathbb{F}$ and $\lambda \in \mathbb{F}$ $\lambda.(a \cdot b) = (\lambda.a) \cdot b = a \cdot (\lambda.b)$ holds.

Thus, a $\mathbb{F}$-algebra A is nothing more than a *possibly non-commutative* ring equipped with an injective ring homomorphism $\mathbb{F} \hookrightarrow A$ whose image is the "copy of $\mathbb{F}$". Importantly, we have $1_{\mathbb{F}} \mapsto 1_A$.

Definition 6.3.2 (group algebra)
Let G be a group. We denote by $\mathbb{C}[G]$ the set of formal linear combinations as follows:

$$\mathbb{C}[G] = \left\{ \sum_{g \in G} \lambda_g g \mid \lambda_g \in \mathbb{C} \right\}$$

By taking the obvious definitions of (component-wise) addition and taking scalar multiples, $\mathbb{C}[G]$becomes a vector space over $\mathbb{C}$. Defining multiplication in $\mathbb{C}[G]$ by

$$\left(\sum_{g \in G} \lambda_g g \right) \left(\sum_{h \in G} \lambda_h h \right) := \sum_{g,h \in G} \lambda_g \lambda_h gh = \sum_{a \in G} \left(\sum_{g,h \in G : gh = a} \lambda_g \lambda_h \right) a$$

$\mathbb{C}[G]$ becomes an algebra over $\mathbb{C}$. It is called the group algebra of G over $\mathbb{C}$.

Note that we can have an alternative interpretation of group algebras:

$$\sum_{a \in G} \lambda_a \cdot a \leftrightarrow f : G \to \mathbb{C} \quad f(a) := \lambda_a$$

Then, note that multiplication can be thought of as **convolution**:

$$f * g(c) := \sum_{a \in G} f(a)g(a^{-1}c)$$

We now look at a generalization of vector spaces, called *modules*.

Definition 6.3.3 ($\mathbb{F}$-module)
Let G be a group, $\mathbb{F}$ a field and V a vector space over $\mathbb{F}$. Then V is called an $\mathbb{F}[G]$–module if for each $a \in \mathbb{F}[G]$ and $v \in V$ an element $(v)a \in V$ is defined satisfying:

1. $(v_1 + v_2)a = (v_1)a + (v_2)a$ *for* $v_i \in V$, $a \in \mathbb{F}[G]$)
2. $(\lambda v)a = \lambda(v)a$ *(for* $v \in V$ *and* $a \in \mathbb{F}[G]$*)*
3. $(v)[ab] = ((v)a)b$ *(for* $v \in V$, $a, b \in \mathbb{F}[G]$*)*
4. $(v)1 = v$ *(for* $v \in V$*).*

Intuitively speaking, given a ring R, a R-module is simply an algebraic structure where you can add two elements, and scale by elements of R. Then, note that a vector space is simply a module whose commanding ring is a field.

Definition 6.3.4 (algebraic integer)
Let β be an algebraic number (in $\mathbb{C}$); then β is called an algebraic integer if all coefficients of its minimal polynomial are integers.

Proposition 6.3.5

1. *Let $\beta \in \mathbb{C}$; then β is an algebraic integer iff there exists some* **monic** *polynomial $p(x) \in \mathbb{Z}[x]$ such that $p(\beta) = 0$.*
2. *The set of all algebraic integers form a subring in $\mathbb{C}$.*

Notation. For this section we set $A := \mathbb{C}[G]$.

Definition 6.3.6 (center)
Let the centre of A be $Z(A) := \{c \in A \mid (\forall a \in A)\ ca = ac\}$.

First, given a group algebra $\mathbb{F}[G]$, when does an element lie in the center of this algebra? Indeed, $z \in Z(\mathbb{F}[G]) \iff (\forall g \in G)zg = gz$. This is important and might come in handy in several situations!

Is it possible to obtain a nice form for elements in the center? Indeed, let $z = \sum_{a \in G} \lambda_a a$. Then, for $g \in G$, if we have $zg = gz \implies g^{-1}zg = z$, and so, we have:

$$\sum_{a \in G} \lambda_a g^{-1} a g = \sum_{a \in G} \lambda_a a$$

Maybe it is worth noting here that two formal linear combinations are equal if and only if coefficients are all equal. Each element in a group algebra has a unique form as a unique formal linear combination.

We call ω_i in the proposition below a **class sum**.

Proposition 6.3.7 (properties of the center)
1. $\mathrm{Z}(A)$ *is a subalgebra of* A, *i.e. it is a subspace and a subring as well.*
2. *Let* $G = K_1 \cup K_2 \cup \cdots \cup K_c$ *be the decomposition of* G *into the disjoint union of conjugacy classes. For each* K_i *define* $\omega_i := \sum_{a \in K_i} a$. *Then* $\{\omega_i \mid 1 \leq i \leq c\}$ *is a* $\mathbb{C}$*–basis of* $\mathrm{Z}(A)$ *(considered as a subspace).*

Proof. (1) is obvious by definition. For (2) observe that $c \in \mathrm{Z}(A)$ iff $cg = gc$ holds for each $g \in G$. Let $c = \sum_{h \in G} \alpha_h h$, then $c \in \mathrm{Z}(A)$ iff for every $g \in G$ $g^{-1}cg = c$ holds, i.e. $\sum_{h \in G} \alpha_h g^{-1}hg = \sum_{h \in G} \alpha_h h$. This just requires $\alpha_{g^{-1}hg} = \alpha_h$ for every $g, h \in G$, which means that c is the linear combination of the ω_i-s.
On the other hand, the ω_i-s are linearly independent as they do not have "common elements". □

Theorem 6.3.8
Let $\phi : G \to \mathrm{GL}_n(\mathbb{C})$ *be an irreducible representation of* G. *We denote the character of* ϕ *by* χ.
1. *Each* $\phi(\omega_i)$ *is a scalar multiple of the identity, i.e.* $\phi(\omega_i) = \lambda_i I_n$ *(with some* $\lambda_i \in \mathbb{C}$*).*
2. *Let* $g_i \in K_i$*; then* $\lambda_i = \frac{|K_i|\chi(g_i)}{n}$ *is an algebraic integer.*

Proof. We have for each $1 \leq i, j \leq c$ that $\omega_i \omega_j = \sum_{k=1}^{c} {t_k}^{(i,j)} \omega_k$ (with ${t_k}^{(i,j)} \in \mathbb{C}$). Since the ω_k-s are linearly independent the coefficient ${t_k}^{(i,j)}$ equals the number of solutions of the equation $xy = g_k$ with $g_k \in K_k$ fixed and $x \in K_i$, $y \in K_j$. This means that ${t_k}^{(i,j)} \in \mathbb{N} \subset \mathbb{Z}$. Now fix $j := 1$, say, and consider the system of equations

$$\omega_i \omega_1 = \sum_{k=1}^{c} {t_k}^{(i,1)} \omega_k \,,$$

(for $i = 1, 2, \ldots, c$) written in the form

$$\sum_{k=1}^{c} (\delta_{i,k}\omega_1 - {t_k}^{(i,1)})\omega_k = 0 \,.$$

Now apply the extension of ϕ to A to get

$$\sum_{k=1}^{c} (\delta_{i,k}\phi(\omega_1) - {t_k}^{(i,1)})\phi(\omega_k) = O \,.$$

By Schur's Lemma each $\phi(\omega_k) = \lambda_k I_n$ (with some $\lambda_k \in \mathbb{C}$), and obviously $\lambda_k = \frac{1}{n}\mathrm{Tr}(\lambda_k I_n) = \frac{1}{n}\mathrm{Tr}(\sum_{a \in K_k} \phi(a)) = \frac{|K_k|\chi(g_k)}{n}$; so the previous equations become

$$\sum_{k=1}^{c} (\delta_{i,k}\lambda_1 - {t_k}^{(i,1)})\lambda_k = 0 \,.$$

This can be considered as a homogeneous system of linear equations for the variables λ_k, $k = 1, 2, \ldots, c$. Since χ is not the zero function ($\chi(1) = n$ is a positive integer) not

all λ_k-s are zero. Thus the above homogeneous system of linear equations has nontrivial solutions, therefore the determinant of its matrix should be zero. This determinant is a polynomial expression of λ_1 of the form $p(\lambda_1)$ where $p(x) \in \mathbb{Z}[x]$ and $p(x)$ is monic. So 6.3.5 (1) implies that λ_1 is an algebraic integer, and the argument clearly extends to any of the other λ_j-s. □

We quickly recollect all the facts we know about characters and algebraic integers:

1. If χ is any characters and $g \in G$, then $\chi(g)$ is an algebraic integer.
2. If χ is the character of an irreducible representation, and $g \in K$ where K is a conjugacy class in G, then $\frac{\chi(G)|K|}{\chi(e)}$ is an algebraic integer.

As a corollary of the above two, we obtain the following:

Proposition 6.3.9
Let ϕ be an irreducible representation of G of degree n and denote its character by χ. Assume that K is a conjugacy class in G such that $(|K|, n) = 1$. Then $\frac{\chi(g)}{n}$ is an algebraic integer for any $g \in K$.

Proof. Since $\chi(g)$ is a sum of roots of unity, $\chi(g)$ is an algebraic integer and so is $\frac{\chi(g)|K|}{n}$ by the previous proposition. As $|K|$ and n are coprimes there exist integers ℓ, t satisfying $\ell|K| + tn = 1$; thus $\frac{\chi(g)}{n} = \frac{\chi(g)}{n} \cdot 1 = \frac{\chi(g)}{n}(\ell|K| + tn) = \frac{\chi(g)|K|}{n}\ell + \chi(g)t$ is an algebraic integer. □

Proposition 6.3.10
Assume that ϕ is an irreducible representation of G of degree n and χ is its character. Assume that K is a conjugacy class in G such that $\frac{\chi(g)}{n}$ is an algebraic integer (with $g \in K$). Then either $g^\phi = \alpha I_n$ is a scalar multiple of the identity or $\chi(g) = 0$.

Proof. Let $\lambda = \frac{\chi(g)}{n}$; we know that $\chi(g) = \varepsilon_1 + \varepsilon_2 + \ldots + \varepsilon_n$ where $\varepsilon_1, \varepsilon_2, \ldots, \varepsilon_n$ — the eigenvalues of $\phi(g)$ — are roots of unity, particularly $|\varepsilon_i| = 1$. Hence $|\lambda| = \frac{|\varepsilon_1+\varepsilon_2+\ldots+\varepsilon_n|}{n} \leq 1$, by the triangle inequality. Let r be the order of g and let μ denote a primitive r-th root of unity. Let $F = \mathbb{Q}(\mu)$. If $\sigma \in \mathcal{G}al(F|\mathbb{Q})$ then λ^σ is also an algebraic integer and $|\lambda^\sigma| \leq 1$ as well. Consider now the polynomial $p(x) := \prod_{\sigma \in \mathcal{G}al(F|\mathbb{Q})}(x - \lambda^\sigma)$. Since the coefficients of $p(x)$ are symmetric functions of the λ^σ-s, each coefficient belongs to $\Phi(\mathcal{G}al(F|\mathbb{Q})) = \mathbb{Q}$; so $p(x) \in \mathbb{Q}[x]$. On the other hand, the coefficients are algebraic integers, as they are sums of products of the λ^σ-s; hence $p(x) \in \mathbb{Z}[x]$. Therefore the constant term of $p(x)$ is an integer c such that $|c| = \prod_{\sigma \in \mathcal{G}al(F|\mathbb{Q})} |\lambda^\sigma| \leq 1$, i.e. $c = 0$ or $c = \pm 1$. If $c = 0$ then some of the λ^σ-s is zero, and then $\lambda = 0$ yields $\chi(g) = 0$. If $c = \pm 1$ then $|\lambda| = 1$; this can hold only if $\varepsilon_1 = \varepsilon_2 = \ldots = \varepsilon_n$. Since the minimal polynomial of $\phi(g)$ divides $x^r - 1$, the minimal polynomial has no multiple roots. Thus the minimal polynomial of $\phi(g)$ should be $x - \varepsilon_1$, implying that $\phi(g) = \varepsilon_1 I_n$. □

6.4 Two Theorems of Burnside

We conclude this section with a famous group-theoretic result whose proof uses representation theory.

Theorem 6.4.1
Let G be a finite group, K a conjugacy class in G such that $1 \notin K$ and $|K| = p^t$, p a prime, $t \geq 0$. Then G can't be nonabelian simple.

Proof. Suppose by way of contradiction that G is nonabelian simple, and let $g \in K$. Let χ_i be the character of a nontrivial irreducible representation ϕ_i of degree n_i; then $\operatorname{Ker} \phi_i = 1$ by the simplicity of the nonabelian group G, i.e. ϕ_i is 1–1.

Case 1: p doesn't divide n_i. Then $\frac{\chi_i(g)}{n_i}$ is an algebraic integer by 6.3.9, hence $\chi_i(g) = 0$ or $\phi_i(g)$ is a scalar multiple of I_{n_i}. Suppose that $\chi_i(g) \neq 0$; then $1 \neq g \in \mathrm{Z}_{\phi_i}(G) := \{h \in G \mid \phi_i(h) = \alpha_h I_{n_i},\ \alpha_h \in \mathbb{C}\}$. It is easy to see that $\mathrm{Z}_{\phi_i}(G)$ is a normal subgroup in G and in our case $\mathrm{Z}_{\phi_i}(G) \simeq \phi_i(\mathrm{Z}_{\phi_i}(G))$ is abelian; this is a contradiction, so $\chi_i(g) = 0$ in this case.
Case 2: p divides n_i. Let $n_i := pm_i$.

But we know that $\varrho(g) = 0$. Thus

$$0 = \varrho(g) = \sum_{i=1}^{c} n_i \chi_i(g) = 1 + \sum_{i:p|n_i} p m_i \chi_i(g)$$

(1 standing for $n_1\chi_1(g)$ where ϕ_1 denotes the trivial representation).
Hence $-\frac{1}{p} = \sum_{i:p|n_i} m_i \chi_i(g)$ is an algebraic integer, a contradiction. □

Theorem 6.4.2
Every group G of order $p^a q^b$ (p, q primes) is soluble.

Proof. Induction on the order of the group. If there is some proper normal subgroup N in G then the order of N and G/N is also the product of two prime powers, therefore N and G/N is soluble by induction hypothesis; then G is also soluble. If G is simple and abelian we are done; suppose therefore that G is nonabelian simple. Let Q be a Sylow q–subgroup of G. Since groups of prime power order are soluble Q must be nontrivial; let $1 \neq g \in \mathrm{Z}(Q)$. Then the number of conjugates of g in G equals $|G : \mathrm{C}_G(g)|$, which divides $|G : Q|$ as $Q \leq \mathrm{C}_G(g)$. But $|G : Q|$ is a power of p, and so is $|G : \mathrm{C}_G(g)|$; this contradicts to 6.4.1. □

7 More Character Theory

In this section, we look at some results that make use of character theory of finite groups. I.M. Isaacs's textbook is a good reference for this material, according to Prof. Hermann.

We start with a simple proposition whose proof I omit:

Proposition 7.0.1
$\chi(e) \mid |G|$.

7.1 Induced characters

A representation is nothing but a special kind of homomorphism, and Burnside's Theorems were a non-simplicity criterion. Thus, we're looking for the kernel of some homomorphism.

Turns out we can use the relations between a character and associated representations to obtain the kernel of the representation, and this is, in a way, a second look we'll be taking at Burnside's Theorems.

Proposition 7.1.1
Let $\psi : G \to GL_n(\mathbb{C})$ be a representation of G, and let χ be the character o ψ. Then:
1. *If $a \in \ker \psi$, then $\chi(a) = n = \chi(e)$.*
2. *Conversely, if $a \in G$ such that $\chi(a) = \chi(e)$, then $a \in \ker \psi$.*

Note that $\chi(a) =$ sum of all roots of the characteristic polynomial of $\psi(a) = \epsilon_1 + \epsilon_2 + \ldots + \epsilon_n$, where ϵ_j is a $\mathcal{O}(a)^{\text{th}}$ root of unity. So, if $n = \sum_{i=1}^{n} \epsilon_i \implies \epsilon_1 = \epsilon_2 = \ldots = \epsilon_n = I_n$ by the triangle inequality.

Definition 7.1.2 (kernel of a character)
Let $\ker(\chi) := \{c \in G \mid \chi(c) = \chi(e)\}$ be the kernel of χ $(= \ker \psi)$.

Now, let ϕ be some representation, χ_ϕ be its character. Then, for irreducible χ_i's, we have:

$$\chi_\phi = \sum_{i=1}^{h} n_i \chi_i$$

A natural question to ask, then, is whether there exists any relation between $\ker(\chi_\phi)$ and $\ker(\chi_i)$? To find out, we work with matrix representations.

Note that $\forall\ a \in G$, there exists some matrix M such that:

$$\phi(a) = M^{-1} \cdot \begin{vmatrix} [\phi_1(a)] & & & \\ & [\phi_2(a)] & & \\ & & \ddots & \\ & & & [\phi_l(a)] \end{vmatrix} \cdot M$$

where ψ_j are irreducibles. It follows then that $\ker(\chi_\phi) = \bigcap_i \ker(\chi_i)$.

In principle, one must be able to obtain all normal subgroups of G in this fashion. Why is this the case? To see this, let $N \triangleleft G$. Write $\overline{G} = G/N$. We want a representation of G with kernel N, but this is equivalent to finding a faithful representation of $\overline{G}$. We have met such a representation before: it's nothing but the regular representation of $\overline{G}$.

So, we have: $G \twoheadrightarrow \overline{G} \to \mathrm{GL}_m(\mathbb{C})$ where the second map is one-to-one as it is the regular representation. Thus, we have a representation of G whose kernel is precisely N.

We finally get to our main question: how to check if a function is a character without using the representation?

Definition 7.1.3 (irreducible characters)
Irr(G) is the set of all irreducible characters of G.

Definition 7.1.4 (class functions)
We define cf(G) := $\{f \mid f : G \to \mathbb{C} : (\forall x, y \in G) f(y^{-1}xy) = f(x)\}$ *and call it the set of class functions of G.*

Proposition 7.1.5 (irreducibles form basis)
Irr(G) is a basis of cf(G).

Corollary 7.1.6
If $\phi \in cf(G)$*, then* ϕ *is a character of (some representation of) G if and only if* $\forall \chi \in$ *Irr(G),* $[\chi, \phi]$ *is a nonnegative integer (and in addition* ϕ *is not the constant zero function).*

And finally, we have:

Definition 7.1.7 (induced function)
Let $H \leq G$ *and let* $f : H \to \mathbb{C}$ *be some function. Write* $\hat{f} : G \to \mathbb{C}$ *such that:*

$$\hat{f}(b) = \begin{cases} f(b) & b \in H \\ 0 & b \notin H \end{cases}$$

Also define $f^G : G \to \mathbb{C}$ *as follows:*

$$f^G(a) = \frac{1}{|H|} \sum_{x \in G} \hat{f}(x^{-1}ax)$$

We will call f^G *the function induced from H to G.*

It is easy to see that f^G is a class function of the group G. It's also clear that induction is linear, that is, if $H \leq G$, and $f_1, f_2 : H \to \mathbb{C}$, then $(\lambda_1 f_1 + \lambda_2 f_2)^G = \lambda_1 f_1^G + \lambda_2 f_2^G$. But

what we are concerned about is the case when f is a class function of H, and after the following theorem, we will focus our attention on the case when f is a character of H.

Proposition 7.1.8 (reciprocity law)
Let $\phi \in cf(G)$, $f : H \to \mathbb{C}$. Then $[\phi, f^G] = [\phi \restriction_H, f]_{(H)}$.

Proof. We have:

$$\begin{aligned}
[\phi, f^G]_{(G)} &= \frac{1}{|G|} \sum_{a \in G} \overline{\phi(a)} \cdot f^G(a) \\
&= \frac{1}{|G|} \sum_{a \in G} \overline{\phi(a)} \cdot \frac{1}{|H|} \sum_{x \in G} \hat{f}(x^{-1}ax) \\
&= \frac{1}{|G|} \frac{1}{|H|} \sum_{a \in G} \left(\sum_{x \in G} \overline{\phi(a)} \hat{f}(x^{-1}ax) \right) \\
&= \frac{1}{|G|} \frac{1}{|H|} \sum_{a \in G} \left(\sum_{x \in G} \overline{\phi(x^{-1}ax)} \hat{f}(x^{-1}ax) \right) \\
&= \frac{1}{|G|} \frac{1}{|H|} \sum_{x \in G} \left(\sum_{a \in G} \overline{\phi(x^{-1}ax)} \hat{f}(x^{-1}ax) \right)
\end{aligned}$$

Now, focus on the inner sum. We have a ranging over all elements of G, and hence, $x^{-1}ax$ also ranges over all elements of G. Moreover, $\hat{f}$ vanishes at elements outside of H. Thus, we have:

$$[\phi, f^G]_{(G)} = \frac{1}{|G|} \frac{1}{|H|} \sum_{x \in G} \left(\sum_{a \in H} \overline{\phi(a)} f(a) \right)$$

However, the inner product above is equal to $|H| \cdot [\phi \restriction_H, f]$, and so:

$$[\phi, f^G]_{(G)} = \frac{1}{|G|} \frac{1}{|H|} (|G| \cdot |H| \cdot [\phi \restriction_H, f]) = [\phi \restriction_H, f]$$

and we're done. □

Proposition 7.1.9
If f is a character of H, then f^G is a character of G.

Proof. This follows from the bilinearity of the inner product and the reciprocity law. □

Indeed, if f is a character of H, then we have $f^G(e) = |G : H| f(e)$. It would be nice if the induction of an irreducible character of H turned out to be an irreducible character of G, but unfortunately, this is not true. The converse of this, however, holds.

It is not clear at this point what motivated the Frobenius reciprocity law, but one suspicion is that we can use the Frobenius reciprocity to prove that the Frobenius kernel is a subgroup of a *Frobenius group*. We will prove this. The first question, then, is to ask, what is a Frobenius group?

7.2 Frobenius groups

Recall that any group can be thought of as a permutation group via its regular representation. When you think of a group as a permutation group, note that a non-identity element has zero fixed points.

Question. What about permutation groups in which non-identity elements have *at most one* fixed point?
Answer. This is a silly question, since you can take any group and extend the underlying set by an element that is fixed by all other elements.

Question. The same question as above, but now our group should also be transitive. *Answer.* Recall that by Burnside's Lemma, we have $|G| = \sum_{a \in G} |f(a)|$. If each non-identity element has exactly one fixed point, then $|G| = n + |G| - 1 \implies n = 1$.

Thus, the above question is equivalent to asking for groups in which non-identity elements have at most one fixed point, and $\exists e \neq a \in G$ such that a does not fix anything.

Such groups are called **Frobenius Groups**. This seemingly slight modification of the questions surprisingly adds a great amount of structure to the groups, as we will see soon.

Definition 7.2.1 (Frobenius Groups)
A group $G \leq S_n$ is called a Frobenius groups if:
1. *G is transitive*
2. *No non-trivial element fixes more than one point*
3. *Some non-trivial element fixes a point*

Recall that if two elements are in the same orbit, then their stabilizers are conjugates of each other.

Proposition 7.2.2
Let G be a Frobenius group, and let $\alpha \in \Omega$ (the underlying set on which G acts). Then we have:
1. *$1 \leq G_\alpha \lneq G$.*
2. *Let $\alpha \neq \beta \in \Omega$. Then $G_\alpha \cap G_\beta = \{e\}$.*

The above proposition gives us the abstract definition of a Frobenius group. Let $H = G_\alpha$, such that $1 \lneq H \lneq G$. Then, if for every $c \in G, c \notin H$ we have $cHc^{-1} \cap H = \{e\}$.

As an exercise, try to reconstruct original permutation group situation from the above definition. You can do so by considering the permutation action on cosets of H.

Definition 7.2.3 (Frobenius complement)
$H \leq G$ as seen above (where G is a Frobenius group) is called a Frobenius complement. That is, $H \leq G$ and for all $a \in G \backslash H$ we have $H \cap aHa^{-1} = \{e\}$.

Definition 7.2.4 (Frobenius kernel)
In a Frobenius group G, write $F = \{e\} \bigcup \{G \setminus \bigcup_{c \in G} cHc^{-1}\}$. F is called the Frobenius kernel.

In what follows, G, F, H are as above. You would think that normality is the surprising part, but in fact, if we can prove that F is a subgroup of G, normality comes for free (since the union of conjugates of H is invariant under conjugation, F must also be invariant under conjugation and hence normal). So in fact, it is amazing that F is a subgroup at all!

There is no known character-free proof that the Frobenius kernel is a subgroup, and most purely group-theoretic proofs of special cases are very complicated.

Lemma 7.2.5
Let $\phi \in \mathrm{cf}(H)$ such that $\phi(e) = 0$. Then, the following holds: $\phi^G \restriction_H = \phi$.

Proof. Recall that $\phi^G(e) = |G : H| \cdot \phi(e) \implies \phi^G(e) = 0$. Now, let $e \neq h \in H$. Then we have: $\phi^G(h) = \frac{1}{|H|} \sum_{x \in G} \hat{\phi}(x^{-1}hx)$.

Now, the only thing that matters is for which $x \in G$, we have $x^{-1}hx \in H$, since $\hat{\phi}$ is 0 for elements not in H. The equivalent question then is $h \in xHx^{-1}$, which is equivalent to $h \in H \cap xHx^{-1}$. Now, this happens for a non-identity element of H if and only if $x \in H$ because of the special property defining Frobenius groups.
Recall that ϕ is a class function, so $\phi(x^{-1}hx) = \phi(h)$. So, we have:

$$\phi^G(h) = \frac{1}{|H|} \sum_{x \in H} \phi(x^{-1}hx) = \phi(h)$$

which is the desired result. □

Let us use the above lemma to construct another function: $1_H \neq \epsilon \in \mathrm{Irr}(H)$. Then define $\phi = \epsilon(e) \cdot 1_H - \epsilon$. Note that this is a class function of H and takes value 0 at e. On the other hand, we have $\phi^G = (\epsilon(e)1_H)^G - \epsilon^G$ by linearity of induction.

Lemma 7.2.6
$[\phi^G, \phi^G] = \epsilon(e)^2 + 1$

Proof. We first use the reciprocity law:

$$[\phi^G, \phi^G] = [\phi^G \restriction_H, \phi] = [\phi, \phi][\epsilon(e) \cdot 1_H - \epsilon, \epsilon(e) \cdot 1_H - \epsilon] = \epsilon(e)^2 + [\epsilon, \epsilon] = \epsilon(e)^2 + 1$$

where we also utilized bilinearity of inner product and orthogonality of induced characters. □

Lemma 7.2.7
$[\phi^G, 1_G] = \epsilon(e)$

Proof. First using reciprocity law:

$$[\phi^G, 1_G] = [\phi, 1_H] = [\epsilon(e) \cdot 1_H - \epsilon, \epsilon(e) \cdot 1_h - \epsilon] = \epsilon(e)[1_h, 1_H] - [\epsilon, 1_H] = \epsilon(e)$$

using orthogonality and bilinearity, as in the previous lemma. □

The above result gives us a lot about ϕ^G as a class function. We know now that if ϕ^G is decomposed as a linear combination of characters, then we get the coefficient of the trivial character. This can be written as:

Corollary 7.2.8
$\phi^G = \epsilon(e) \cdot 1_G + \sum_{1_G \neq \chi \in \text{Irr}(G)} n_\chi \cdot \chi$ *where* $n_\chi \in \mathbb{Z}$.

We now compute the inner product of ϕ^G with itself using the form above, and compare with the result of Lemma 7.2.6.

$$[\phi^G, \phi^G] = \epsilon(e)^2 + \sum_{1_G \neq \chi \in \text{Irr}(G)} n_\chi^2 = \epsilon(e)^2 + 1 \implies 1 = \sum_{1_G \neq \chi \in \text{Irr}(G)} n_\chi^2$$

So, there exists a unique irreducible character $\chi \neq 1_G$ such that $n_\chi^2 = 1$. So, is $n_\chi = 1$ or -1? Actually, $n_\chi = -1$, and this is because ϕ^G vanishes at the identity.

From all of this, the following theorem will follow by considering kernels of χ obtained as $1_H \neq \epsilon \to \chi \neq 1_G$.

Let's now return to the Frobenius kernel F.

Theorem 7.2.9
$F \trianglelefteq G$.

We also gave an alternative characterization of the Frobenius kernel F as the intersection of "extensions" of irreducible characters of the Frobenius complement H.

I was absent for the above, but it can be found very easily online and in I.M. Isaacs's textbook.

7.3 The Second Orthogonality Relation

Fact. $|\text{Irr}(G)|$ = number of conjugacy classes in $G = \dim(\text{cf}(G))$.

Let M denote the **character table** of G where the columns denote the conjugacy classes of G and the rows represent irreducible characters. Then we can get $(\overline{M} \cdot M^T)_{i,j} = [\chi_i, \chi_j] \implies \overline{M} \cdot M^T = I_n$ where n is the number of irreducible characters of G.

But $\overline{M} \cdot M^T = I_n \implies M^T \overline{M} = I_n$, that is:

$$\delta_{i,j} = (M^T \overline{M})_{i,j} = \sum_{t=1}^{n} \chi_t(a+i)\overline{\chi_t(a_j)} \cdot \frac{\sqrt{|K_i||K_j|}}{|G|}$$

which means:

$$\sum_{\chi \in \text{Irr}(G)} \chi(a)\overline{\chi(b)} = \begin{cases} \frac{|G|}{|K|} & \text{if } i = j; a, b \in K \\ 0 & \text{if } a \text{ and } b \text{ are not conjugates} \end{cases}$$

where $|K| = |G : G_a| = |G : C + G(a)| \implies \frac{|G|}{|K|} = |C_G(a)|$.

Thus, we have:

$$\sum_{\chi\in\mathrm{Irr}(G)} \chi(a)\overline{\chi(b)} = \begin{cases} |C_G(a)| & \text{if } b \sim a \\ 0 & \text{otherwise} \end{cases}$$

7.4 More Applications

G a finite group acting on Ω, a finite set. For $g \in G$, let $f(g) := |\{v \in \Omega \mid g(v) = v\}|$. Then:

1. $\sum_{g\in G} f(g) = |G|$ (# of all distinct orbits of G on Ω)
2. If the action is transitive then $\sum_{g\in G} f(g)^2 = |G|$ (# of all distinct orbits of G_α on Ω)

We discuss the above and recast them in the language of inner products:

1. $[f, 1_G] =$ # of all orbits of G on Ω
2. If $[f, 1_G] = 1$ then $[f, f] =$ # of all orbits of G.

Suppose G acts transitively on Ω, G_α has 2 orbits on Ω: $\{\alpha\}$ and $\{\Omega\backslash\alpha\}$. If this is the case, we say that G acts 2-transitively on Ω. Let $\Omega^{[2]} := \{(\beta_1, \beta_2) \in \Omega^2 \mid \beta_1 \neq \beta_2\}$. Then G acts on $\Omega^{[2]}$ in the following way:

$$a(\beta_1, \beta_2) = (a(\beta_1), a(\beta_2))$$

We say that the action of G on ω is 2-transitive if the above action of G on $\Omega^{[2]}$ is transitive.

Proposition 7.4.1
G acts 2-transitively on Ω iff G acts transitively on Ω, and if the number of orbits of G_α is precisely 2 ($\forall/\exists\alpha \in \Omega$).

Example 7.4.2 (2-transitive actions)
Let V be a vector space over the field $\mathbb{F}$. Let $G = \{f_{A,\underline{b}} \mid A \in \mathrm{GL}(V), \underline{b} \in V\}$ be the set of affine transformations of V. The action of G on V is given by $f_{A,\underline{b}}(\underline{v}) = A(\underline{v})+\underline{b}$. This defines a 2-transitive action.

Our goal for doing all of this:

Proposition 7.4.3
G a finite group, G acting on Ω such that $|\Omega| = n < \infty$. Suppose that the action of G on Ω is 2-transitive, and let $H \leq G$ such that $|G : H| < n$. Then H acts transitively on Ω.

Proof. We can rewrite the given conditions as $[f, f] = 2$ and $[f, 1_g] = 1$. Therefore, f is a character (which should be obvious if you consider the permutation matrix associated with group elements). So $f = \sum_{\chi\in\mathrm{Irr}(G)} n_\chi\chi$ where the n_χ are nonnegative integers. Moreover, $n_\chi = [\chi, f]$.
Now we have:

$$f = 1_G + \sum_{1_G\neq\chi\in\mathrm{Irr}(G)} n_\chi\chi$$

But $[f,f]=2$, and $[f,f]=1+\sum_{1_G\neq\chi\in\mathrm{Irr}(G)} n_\chi^2$. So $f=1_G+\chi$ for some $1_G\neq\chi\in\mathrm{Irr}(G)$.
We need : $[f\restriction_H, 1_H]=1$. By the reciprocity law we get $[f\restriction_H, 1_H]=[f,(1_H)^G]$.
Consider $[1_G,(1_H)^G]$. By the reciprocity law again, it is equal to $[(1_G)\restriction_H, 1_H]=[1_H,1_H]=1$ and so, 1_G is a constituent character of $(1_H)^G$.
The last statement can be rewritten as $(1_H)^G=1_G+\sum_{1_g\neq\chi\in\mathrm{Irr}(G)} t_\chi\chi \implies [f,(1_H)^G]=[f,1_G]+[f,\sum t_\chi\chi]$. The first term $[f,1_G]=1$ because of transitivity.
Thus, we need $[f,\sum t_\chi\chi]=0$. Note that $f=\chi_\rho$ for some representation ρ whose dimension must be n. Similarly, $\sum t_\chi\chi=\chi_\phi$ for some representation ϕ whose dimension must be $|G:H|-1$.
Now, we evaluate $[f,\sum t_\chi\chi]$ and use that $f=1_G+\chi$ which we obtained earlier. We have $[1_G+\chi,\sum t_\chi\chi]=[\chi,\sum t_\chi\chi]$. The dimension of $\chi=n-1$ since $f=1_G+\chi$. We also know dimension of the $\sum t_\chi\chi$ is $|G:H|-1$.
Now, the only way that the above inner product is greater than zero is if the coefficient of χ in the proof is greater than 0, i.e. at least 1, but this means that χ is a constituent of $\sum t_\chi\chi$. This is not possible, as dimension of χ is greater than the dimension of $\sum t_\chi\chi$. □

We now look at another application.

Let K be some conjugacy class in G. Suppose $\omega_k=\sum_{x\in K} x$. Then we have: $\omega_A\omega_B=\sum_K n_K^{A,B}\omega_K$ where in the final sum the coefficients denote multiplicities. Hence $n_K^{A,B}\geq 0$ and is an integer. Let ψ be an irreducible representation of G.

So we have: $\lambda_A I\times\lambda_B I=\sum_K n_K^{A,B}\lambda_K I \implies \lambda_A\lambda_B=\sum_K n_K^{A,B}\lambda_K$.

Thus we have:

$$\lambda_A=\frac{Tr(\hat{\psi}(\omega_A))}{\chi_\psi(e)}=\frac{Tr(\sum_{a\in A}\psi(a))}{\chi_\psi(e)}=\frac{\chi_\psi(a)|A|}{\chi_\psi(e)}$$

We try to get a single coefficient out of this using the second orthogonality relation. We have:

$$\chi(a)\chi(b)\frac{|A||B|}{\chi(e)^2}=\sum_K n_K^{A,B}\frac{\chi(x)|K|}{\chi(e)}$$

Let $c\in\mathcal{C}$ which is a conjugacy class. Then, multiplying the above expression by $\overline{\chi(c)}$ we get:

$$\sum_{\chi\in\mathrm{Irr}(G)}\chi(a)\chi(b)\frac{|A||B|}{\chi(e)^2}=\sum_{\chi\in\mathrm{Irr}(G)}\sum_K n_K^{A,B}\frac{\chi(x)|K|\overline{\chi(c)}}{\chi(e)}$$

Rewriting we get:

$$\sum_{\chi\in\mathrm{Irr}(G)}\chi(a)\chi(b)\frac{|A||B|}{\chi(e)}=\sum_K n_K^{A,B}|K|\sum_{\chi\in\mathrm{Irr}(G)}\chi(x)\overline{\chi(c)}$$

Thing on the right only survives for $K=C$. Then we get the sum on the right to be $n_C^{A,B}|C||C_G(c)|=n_C^{A,B}|G|$ by the orbit-stabilizer theorem.

Rearranging everything gives us:

$$n_C^{A,B}=\frac{|A||B|}{|G|}\sum_{x\in\mathrm{Irr}(G)}\frac{\chi(a)\chi(b)\overline{\chi(c)}}{\chi(e)}$$

Notation. We changed notation slightly across classes: $n_C^{A,B}$ is the same as $n_C^{[A,B]}$ in what follows.

What's the use of all this? Well, it could be used to study the commutator subgroup, which is the smallest normal subgroup whose factor group is abelian. The commutator elements can be obtained from the expression we obtained for $n_C^{[A,B]}$. We look at this next by attempting to answer the following question.

Question. Given $g \in G$, is g a commutator?
Answer. Note that if g is a commutator, then $g = aba^{-1}b^{-1} = a(ba^{1-}b^{-1})$. So, $g \in G$ such that $\exists A$ conjugacy class such that $n_C^{[A,A^{-1}]} > 0$. Here A^{-1} is the set of inverses of all elements in the conjugacy class A. But this is equivalent to

$$\sum_{A \in \text{Conjugacy Classes}} \frac{|G|}{|A|} \times n_C^{[A,A^{-1}]} > 0$$

We now try to get a nicer form for the left hand side expression. Making use of the stuff from last class:

$$|G| \sum_A \frac{1}{|A^{-1}|} \frac{|A||A^{-1}|}{|G|} \sum_{x \in \text{Irr}(G)} \frac{\chi(a)\chi(a^{-1}\overline{\chi(c)}}{\chi(e)} = \sum_{\chi \in \text{Irr}(G)} \frac{\overline{\chi(c)}}{\chi(e)} \sum_A |A|\chi(a)\chi(a^{-1})$$

The right-most expression is simply $\sum_{a \in G} \chi(a)\chi(a^{-1}) = |G|$. Thus, the above condition is equivalent to

$$|G| \sum_{\chi \in \text{Irr}(G)} \frac{\overline{\chi(c)}}{\chi(e)} > 0$$

and this can finally be written as

$$\sum_{\chi \in \text{Irr}(G)} \frac{\chi(c)}{\chi(e)} > 0 \iff c \text{ is a commutator}$$

What's all the above good for? Here's an example.

Corollary 7.4.4
If c is a commutator in G, a finite group, such that $\gcd(\mathcal{O}(c), k) = 1$ for some $\mathbb{Z} \ni k > 0$, then c^k is a commutator.

This follows easily from the stuff we did earlier today, simply consider the diagonal matrix representation of $\rho(c)$ and look at its k^{th} power. We don't give a complete proof in class, but intuitively it's easy to see what's happening. Refer to Isaac's book for a detailed explanation.

"*More or less, I guess, this is the end of the series of applications of characters and orthogonality relations. Maybe one more thing...*"

What can be said about values of group characters?
If a character is linear, then the values arising there are roots of unity. But what if it's not a linear character. Here's a "very useless but funny" proposition.

Proposition 7.4.5
If $\chi \in \mathrm{Irr}(G), \chi(e) > 1 \implies (\exists b \in G) : \chi(b) = 0.$

Proof. Suppose the contrary. Then we have

$$G = \bigcup \mathscr{L}_c$$

where $\mathscr{L}_c = \{c^k \mid \mathcal{O}(c^k) = \mathcal{O}(c)\}$. Note that $\chi(c^k) \neq 0$. What about $\prod_{c^k \in \mathscr{L}_c} \chi(c^k)$? This is known to be an algebraic integer, but because of some result in Galois Theory to do with field automorphisms it must be rations. It follows that it is an ordinary integer. Then we have

$$\prod_{c^k \in \mathscr{L}_c} |\chi(c^k)| \geq 1 \implies \prod_{all\ classes} \left(\prod_{c^k \in \mathscr{L}_c} |\chi(c^k)| \right) \geq 1$$

Thus, we have

$$\prod_{g \in G} |\chi(g)|^2 \geq 1$$

But by the AM-GM inequality we have

$$1 = \left(\frac{\sum_{g \in G} |\chi(g)|^2}{|G|} \right)^{|G|} \geq \prod_{g \in G} |\chi(g)^2| \geq 1$$

and so we have equality, and so all the elements are equal and so $\chi(e) = 1$, giving a contradiction. □

The purpose of the above was to see a cool application of some sophisticated techniques.

We conclude the course with this assertion we mentioned while discussing the transfer homomorphism.

Proposition 7.4.6
If A *is a finitely generated abelian group, then* $A \cong B \oplus \mathbb{Z}^+ \oplus \ldots \oplus \mathbb{Z}^+$ *where* B *is a finite group.*

We first make the following observation: $F = \langle x_1 \rangle \oplus \ldots \oplus \langle x_n \rangle$ where $\langle x_i \rangle$ is infinite for all i, is free in the category of abelian groups Ab. That is, if $f : \{x_1, \ldots, x_n\} \to G$ where G abelian, then there exists a unique homomorphism ψ such that $\psi(x_j) = f(x_j)$.

1. If A is abelian, and $\epsilon : A \twoheadrightarrow F$ a surjective homomorphism, then $A = \ker \epsilon \oplus A_1$ where $A_1 \cong F$.
   Proof. Note that for every x_j, there exists $a_j \in A$ such that $\epsilon(a_j) = x_j$; let $\psi : F \to A$ be determined by $\psi(x_j) = a_j$. It's easy to see that $\epsilon \circ \psi$ will give you the identity homomorphism.

   Let $A_i = \mathrm{Im}\ (\psi)$. It is obvious that $\ker \epsilon \cap \mathrm{Im}\ (\psi) = e = 0$. It's also straightforward to show that any element can be written as a unique combination of elements from $\mathrm{Im}\ (\psi)$ and $\ker \epsilon$, and so we're done. □

2. If our finitely generated abelian group is torsion free, then we don't have to worry

about the direct component B, i.e. it disappears from expression in Proposition 5.8.12.

Proof. We proceed by induction on the number of generators. Suppose $A = \langle a_1, \ldots, a_n \rangle$. When $n = 1$, life is good. Suppose true for $n-1$ generators, then simply quotient out by $\langle a_n \rangle$. This *might* not be good, as the quotient might not be torsion-free.

Instead, we do the following: let $C = \{b \in A \mid \exists t \neq 0 : tb \in \langle a_n \rangle\}$ which is the radical of the subgroup (you may have seen something similar in commutative algebra/algebraic number theory). You can check that A/C is torsion-free and is generated by $n-1$ elements. Then, by the inductive hypothesis, $A/C \cong \mathbb{Z}^+ \oplus \ldots \oplus \mathbb{Z}^+ =: F$.

Consider the natural surjection $\eta : A \to F$ and then by the previous proposition, we have $A = K \oplus L$ where $L \cong F$. We get $K \cong A/L$ which is finitely generated. We also have $K \cong C$. The rest of it is plain bashing to show that K is cyclic, and I didn't take it down. □

We can finally return to the general case. Suppose A is finitely generated and abelian. Then we define B to be the torsion subgroup of A, namely the set of elements of finite order. If you mod out by B, then it's easy to check that A/B will be torsion free.

So all that remains to be shown is that B is finite, which we didn't show in class but shouldn't be too hard I think.