Further Group Theory

Aniruddha Nadiga

Last updated: October 30, 2019

https://www.overleaf.com/project/5d99f689a31d2a00016a0303

Contents

0 Introduction

These are my notes for Further Group Theory, taught by Dr. Vaibhav Gadre. The topics and their ordering generally follow the book *Groups, Graphs, and Trees* by John Meier.

If you find any errors please tell me (or email me at asnadiga@gmail.com).

0.1 Notation

1 Cayley's Theorems

In this course we often consider symetries of some object. If X is a mathematical object then we use $\text{Sym}(X)$ to denote all bijections form X to X that preserve some structure. For example if X is the set $[n] = \{1, 2, ..., n\}$ then $\text{Sym}(X)$ is the symmetric group of order n, denoted Sym_n.

1.1 Cayley's "Basic" Theorem

Definition 1.1.1
An action of a group G on a set X is a group homomorphism from G to $\text{Sym}(x)$. Equivalently,its a map from $G \times X \to X$ such that

1. $e \cdot x = x$, for all $x \in X$; and
2. $(gh) \cdot x = g \cdot (h \cdot x)$ for all $g, h \in G$ and $x \in X$.

We denote "G acts on X"by $G \curvearrowright X$. The associated homomorphism is a *representation* of G. If the representation if injective then it is *faithful*.

For this course all group actions are on the left, this will be important later when defining Cayley graphs.

Theorem 1.1.2: Cayley's Theorem
Every group can be faithfully represented as a group of permutations

Proof. We can view a group G as acting "on its own elements", meaning that each element of G can be viewed as a permutation of the elements of G. Specifically, the permutation associated to and $g \in G$ is $\pi_g(h) = gh$ for all $h \in G$. This clearly must be a perumation since $gh = gh'$ would imply that $h = h'$. Thus we have a map $g \mapsto \pi_g \in \text{Sym}(G)$.
It is easy to see that this mapping is a group homomorphism, so all that is left to check is that this representation if faithful. To do this, note that the kernel of this homomorphism is only the identity element. □

1.2 Graphs

Definition 1.2.1
A graph Γ consists of a set of vertices $V(\Gamma)$ and a set of edges $E(\Gamma)$, where each edge e is associated to a pair of vertices, $\text{Ends}(e) = \{u, v\} \in V$

Note that there can be multiple edges between two vertices and an edge e can be a loop that where $\textsc{Ends}(e) = \{v\}$. A graph with out multiple edges or loops is a *simple graph*. The *valence* of a vertex is the number of edges that contain it, and a graph is *locally finite* if the valence of every vertex is finite.

A *path* is a set of alternating vertices and edges $\{v_0, e_1, v_1, ..., v_{n-1}, e_n, v_n\}$ where $\mathsf{Ends}(e_i) = \{v_{i-1}, v_i\}$. A graph is *connected* if any two vertices can be joined by a path. A *backtrack* is a path of the form $\{v, e, w, e, v\}$. A path is *reduced* if it contains no backtracks. A cycle is a path that begins and ends at the same point. A *tree* is a connected graph that has no cycles.

Proposition 1.2.2
The following conditions on a connected graph Γ are equivalent:
1. Γ is a tree.
2. Given any two vertices v, w there is a unique reduced edge path from v to w.
3. For every edge e removing e disconnects the graph.
4. IfΓ is finite, then $|V(\Gamma)| = |E(\Gamma)| + 1$.

Proof. Left as an exercise. □

Definition 1.2.3
A *regular m-tree* is a tree where every vertex has valance m. For a given m this tree is unique and denoted $\mathcal{T}_m$. A graph is a *biregular tree* if it is bipartite and all of the vertices in one class have valance m and all vertices in the other class have valance n. Given m and n there is a unique biregular tree $\mathcal{T}_{m,n}$

Definition 1.2.4
A *directed* graph is a graph where each edge is an ordered pair of vertices, meaning that edges have a begining and an end.

A directed graph is connected when its underlying undirected graph is, we do not consider directed connectedness.

Ther can also be decorations on the graph, such as a function $V(\Gamma) \to \mathcal{L}$ where $\mathcal{L}$ is some set of labels.

1.3 Symmetry Groups of Graphs

Definition 1.3.1
A *symmetry* of a graph Γ is a bijection α that takes vertices to vertices and edges to edges such that if for some edge e, $\textsc{Ends}(e) = \{v, w\}$, then $\textsc{Ends}(\alpha(e)) = \{\alpha(v), \alpha(w)\}$. The *symmetry group* of Γ is the collection of all its symmetries, where the group operation is composition, denoted $\textsc{Sym}(\Gamma)$.

Example 1.3.2
The symmetry group of the complete graph K_n is isomorphic to SYM_n

Definition 1.3.3
Let G be a subgroup of $\text{SYM}(\Gamma)$. We say that G is *vertex transitive* if for any two vertices v, v' there is some $g \in G$ such that $g(v) = v'$. Similarly, G is edge transitive if for any two edges e, e' there is some $g \in G$ such that $g(e) = e'$. A *flag* is a pair (v, e) where $v \in \text{ENDS}(e)$. G is *flag transitive* if for any two flags $(v, e), (v', e')$ there is some $g \in G$ such that $g(v) = v'$ and $g(e) = e'$.

We can further say that G is *simply transitive* on the vertices (or edges or flags) if it is vertex (or edge or flag) transitive and for any two v, v' there is a unique $g \in G$ such that $g(v) = v'$ (respectively e, e', $g(e) = e'$, etc).

Example 1.3.4
The symmetry group of the complete bipartite graph $K_{m,n}$ is not vertex transitive if for $m \neq n$. This is because symmetries preserve valance so vertices of one class can never be taken to vertices of another class under a symmetry.

Lemma 1.3.5
If Γ is a directed graph then the collection of all symmetries that preserve every edges direction form a subgroup of $\text{SYM}(\Gamma)$.

Proof. Clearly this collection is a subset of $\text{SYM}(\Gamma)$, so we must show that it is closed under composition and contains all inverses. If g and h preserve edge directions then so do their inverses and gh. $\square$

An almost identical proof shows the following.

Lemma 1.3.6
If the edges and/or the vertices of a graph Γ are labeled then the collection of all symmetries that preserve the labeling form a subgroup of $\text{SYM}(\Gamma)$.

Definition 1.3.7
If a graph Γ comes with certain decorations, then let $\text{SYM}^+(\Gamma)$ be the subgroup of $\text{SYM}(\Gamma)$ that preserve all decorations.

1.4 Orbits and Stabilizers

Definition 1.4.1
Let a group G act on X. If $x \in X$ then the *stabilizer* of x is

$$\text{Stab}(x) = \{g \in G \mid g \cdot x = x\}.$$

For a given element $x \in X$, the identity element of G fixes x, and for any $g, h \in G$ that fix x, so do their inverses and gh. Thus, for any $x \in X$ Stab(x) is a subgroup of G.

Example 1.4.2
D_n is the group of symmetries of a regular n-gon. If x is a vertex of this n-gon, then there are two elements that fix x: the identity, and the reflection over the axis that contains x. Thus $\mathrm{Stab}(x) \approx \mathbb{Z}_2$.
On the other hand, if x is the center of the n-gon then $\mathrm{Stab}(x) \approx D_n$.

Example 1.4.3
If $\mathrm{Sym}_n \curvearrowright [n]$ then $\mathrm{Stab}(i) \approx \mathrm{Sym}_{n-1}$. This is because we can take all permutations of the $n-1$ elements that are not i and this forms $\mathrm{Stab}(i)$.

Definition 1.4.4
Let $G \curvearrowright X$. The *orbit* of some element $x \in X$ is

$$\mathrm{Orb}(x) = \{x \in X \mid x = g \cdot x \text{for some } g \in G\}$$

For some $G \curvearrowright X$, pick a fixed x. Then $g \cdot x = h \cdot x$ if and only if $gh^{-1} \cdot x = x$, which is equivalent to $gh^{-1} \in \mathrm{Stab}(x)$. Thus $g \cdot x = h \cdot x$ if and only if the left cosets $g\mathrm{Stab}(x) = h\mathrm{Stan}(x)$. This means that the mapping $g \cdot x \leftrightarrow g\mathrm{Stab}(x)$ is a well defined bijective correspondence between the orbit of x and the left cosets of its stabilizer. This gives us the following.

Theorem 1.4.5: Orbit-Stabilizer Theorem
Let G be any finite group acting on X. Then for any $x \in X$,

$$|G| = |\mathrm{Stab}(x)| \cdot |\mathrm{Orb}(x)|.$$

Proof. The order of G is the order of $\mathrm{Stab}(x)$ times the index of $\mathrm{Stab}(x)$ in G. We have shown that the index of $\mathrm{Stab}(x)$ in G (the number of left cosets) is the order of the orbit of x. □

This theorem, among other things, shows that if the stabilizer of an element is trivial, then the elements of its orbit correspond to group elements.

1.5 Generating Sets

Definition 1.5.1
If G is a group and S is a subset of its elements, then S *generates* G if every element of G can be expressed as a product of elements of S and inverses of elements of S. A group is *finitely generated* if it has a finite generating set.

The generating set for a group is not unique. In fact even the smallest generating set is not unique as seen the the example below.

Example 1.5.2
in D_n any reflection along with the "minimal" rotation by $2\pi/n$ is a generating set. Any two adjacent reflections are also a generating set.

We can also note that every finite group must be finitely generated because the set of all elements forms a generating set.

Example 1.5.3
The rational numbers are not finitely generated. We assume that the operation is addition.
Suppose for contradiction that they are, and $S = \{n_1/d_1, n_2/d_2, ..., n_k/d_k\}$ is a finite generating set. Then every element generated by this set can be expressed as some n/d where $d = \mathrm{LCM}(d1, d2, ..., d_k)$. But obviously not every element of $\mathbb{Q}$ can be expressed in this way.

1.6 Cayley Graphs

Theorem 1.6.1: Cayley's Better Theorem
Every finitely generated group can be faithfully represented as a symmetry group of a connected, directed, locally finite graph.

Proof. For a given group G, we need a graph such that G can be faithfully represented as a symmetry group of that graph. Instead of looking for one, we will construct one, $\Gamma G, S$. The vertices of $\Gamma_{G,S}$ correspond to the elements of G, and are denoted as v_g for $g \in G$. There is an edge from g to g' if there is some $s \in S$ such that $gs = g'$. Thus every edge corresponds to right multiplication of its initial vertex by a generator.
To show that $GG_{G,S}$ is connected, note that since S generates G, there is a path from every vertex to the vertex corresponding to the identity element. Thus there is a path from any given vertex to any other vertex. The valance of every vertex will clearly be finite since S is finite, so $\Gamma_{G,S}$ is locally finite.
The proof of Cayley's basic theorem described the group action of G on the vertices $\Gamma_{G,S}$: $g \cdot v_h \mapsto v_{gh}$. In order for this to be a symmetry of $GG_{G,S}$, this action must extend to an action on the edges. We see that if there is an edge e from v_{h_1} to v_{h_2} then there is some $s \in S$ such that $h_1 s = h_2$. If we take the action of g on these vertices we see that $g \cdot v_{h_1} = v_{gh_1}$ and $g \cdot v_{h2} = v_{gh_2} = v_{gh_1 s}$, meaning there is an edge between the images of the two vertices. Thus we can define an action of G on the edges as well. $\square$

Definition 1.6.2
Given a group G and a finite generating set S, we call $\Gamma_{G,S}$ the Cayley graph of G with respect to S.

There is an obvious way to label and the edges of $\Gamma_{G,S}$: if an edge is from v_g to v_{gs} label it as s. The action of G on $\Gamma_{G,S}$ preserves the orientation and labeling so we conclude that $G < \mathrm{Sym}^+(\Gamma_{G,S})$

1.6.1 Cayley Graphs of Dihedral Groups

Dihedral groups can be generated by a set of two adjacent reflections or by a set consisting of a minimal rotation and a refelction. These different generating sets give very different looking Cayley graphs.

One trick to generate the Cayley graph with respect to either of these two generating sets is as follows. Choose a point in "general position" the regular n-gon, meaning that its stabilizer is trivial. This will imply that the elements of the orbit of this point will correspond to elements of D_n. By applying the actions of the elements generating set repeatedly to the point and drawing the corresponding edges, you will draw the Cayley graph.

This trick can be used to find the Cayley graphs for other groups that represent some symmetries of a geometric object. For example S_4 can represent the symmetries of a tetrahedron and picking a point on the boundary of the tetrahedron that has a trivial stabilizer can reveal te cayley graph of S_4.

1.7 Symmetries of Cayley Graphs

Theorem 1.7.1
Let $\Gamma_{G,S}$ be the Cayley Graph of G with respect to the finite generating set S. Consider $\Gamma_{G,S}$ to be decorated with directions on its edges and labeling of ites edges, corresponding to the generating set S. Then $\text{Sym}^+(\Gamma_{G,S}) \approx G$.

Proof. The left action of G on $\Gamma_{G,s}$ defined in the proof of Cayley's "better" theorem shows that there is an injective homomorphism $G \to \text{Sym}(\Gamma_{G,S})$. Since the edges are defined through right multiplication and this group action is a left action, the edge labeling is preserved by the group action. Thus we have that we have an injective homomorphism $G \to \text{Sym}^+(\Gamma_{G,S})$.
To show that this homomorphism in surjective, consider an $\gamma \in \text{Sym}^+(\Gamma_{G,S})$. For any element $g \in G$ let v_g be the corresponding vertex in $\Gamma_{G,S}$. Then there is some g such that $\gamma(v_e) = v_g$. Then if we consider g as a symmetry of $\Gamma_{G,S}$, $g^{-1}\gamma \in \text{Sym}(^)+(\Gamma_{G,S})$. This symmetry fixes v_e and fixes all of the edges leaving or going to v_e since symmetries in $\text{Sym}(^)+(\Gamma_{G,S})$ preserve the labeling and orientation of edges. This means that $g^{-1}\gamma$ fixes all of the vertices adjacent to v_e, and following a similar argument, it fixes all of the vertices adjacent to those vertices. Thus we get that $g^{-1}\gamma$ is the identity which means that $g = \gamma$. Thus we conclude that the homomorphism that we had from G to $\text{Sym}(^)+(\Gamma_{G,S})$ is an isomorphism. $\square$

1.8 Fundamental Domains and Generating Sets

In what follows we think of graphs as geometric objects consisting of unit intervals joined together at vertices. Thus we can consider closed subsets of a graph.

Lemma 1.8.1
If a group G acts on a connected graph Γ then there is a subset $\mathcal{F} \subset \Gamma$ such that
1. $\mathcal{F}$ is closed,
2. the set $\{g \cdot \mathcal{F} \mid g \in G\}$ covers Γ, and
3. no proper subsest of $\mathcal{F}$ satisfies properties (1) and (2).

Proof. FINISH ME LATER □

ALSO INSERT SOME EXAMPLES

Theorem 1.8.2
Let G act on a connected graph Γ with fundamental domain $\mathcal{F}$. Then the set of elements

$$S = \{g \in G \mid g \neq e \text{ and } g \cdot \mathcal{F} \cap \mathcal{F} \neq \emptyset\}$$

is a generating set for G.

Proof. Pick and $g \in G$ the goal is to show that g can be expressed as a combination of elements from S. Pick any $v \in \mathcal{F}$ and pick some path p connecting v to $g \cdot v$. Then let $\{\mathcal{F}, g_1\mathcal{F}, ..., g_n\mathcal{F} = g\mathcal{F}\}$ be a finite sequence of images of $\mathcal{F}$ such that

1. the entire path p is contained in $\cup g_i\mathcal{F}$ and
2. $g_i\mathcal{F} \cap g_{i+1}\mathcal{F} \neq \emptyset$ where g_0 is the identity.

We must show that such a sequence exists. The first condition can be met because $G \cdot \mathcal{F}$ covers all of Γ. The second condition can be met as a result of $\mathcal{F}$ being closed.
Since $\mathcal{F} \cap g_1\mathcal{F} \neq \emptyset$, we know that $g_1 \in S$ by definition. Since $g_1\mathcal{F} \cap g_2\mathcal{F} \neq \emptyset$, $\mathcal{F} \cap g_1^{-1}g_2\mathcal{F} \neq \emptyset$, so $g_1^{-1}g_2 \in S$. Thus $g_2 = g_1(g_1^{-1}g_2)$ is a product of elements from S. Continuing in this way shows that g can be written as a product of elements from S. □

Theorem 1.8.3
Let $G \curvearrowright \Gamma$ with the fundamental domain $\mathcal{F}$. Further assume that if $g \cdot \ = \mathcal{F}$ then $g = e$. If $H < G$ and a fundamental domain for the induced action $H \curvearrowright \Gamma$ is a union of n copies of $\mathcal{F}$, $(n \in \mathbb{N} \cup \infty)$, then the index of H in G is n.

Proof. Since we assumed that the only element that fixes $\mathcal{F}$ is the identity, there is some point $x \in \mathcal{F}$ that is moved freely by G. This means that there is a bijiection between the elements of $\text{Orb}(x)$ and the elements of G. We can explicity write out the fundamental domain for $H \curvearrowright \Gamma$ as

$$\mathcal{F}_H = \cup_{i=1}^n g_i\mathcal{F},$$

where $n \in \mathbb{N} \cup \infty$, and the g_i are distinct elements of G. Then since $\mathcal{F}_H$ is a fundamental domain, for all $g \cdot x$ there is some $h \in H$ such that $g \cdot x \in h\mathcal{F}_H$. This means that $g \cdot x \in (hg_i)\mathcal{F}$ for some g_i. Since any gx is contained in some $hg_i\mathcal{F}$ and the elements of the orbit of x correspond to the elements of G, we get that $G = \cup Hg_i$. This means that the set of g_i contains all of the coset representatives of H in G, and now we will show that no two g_i represent the same coset.
Suppose for contradiction that $Hg_i = Hg_j$ for some distinct g_i, g_j. Then $g_ig_j^{-1} \in H$. Now consider $\mathcal{F}'_H = \mathcal{F}_H \setminus g_j\mathcal{F}$. We claim that this still covers Γ through the actions of H. The situation that we must check is when some element of Γ $y \in hg_j\mathcal{F}$, for $h \in H$. In this case for some $h' \in H$, $y \in h'(g_ig_j^{-1}g_j\mathcal{F} = h'g_i\mathcal{F} \subset h'\mathcal{F}'_H$, proving that y is covered by $\mathcal{F}'_H$. But this contradicts $\mathcal{F}_H$ being a fundamental domain so we conclude that for any distinct g_i, g_j, $Hg_i \neq Hg_j$, meaning that the set of g_i correspond to the right cosets of H in G, meaning that the order of H in G is n. □

1.9 Words and Paths

Given a set S, a finite sequence of elements of S (possibly with repetition) is called a *word*. The set S is called the *alphabet*. The collection of all words, including the empty word, is denoted S^*.

S^{-1} denotes the set of all formal inverses of elements in S. Then we can also make the set of words $(S \cup S^{-1})*$. We can take formal inverses of words of this form in the usual way $(x_1 x_2 ... x_n)^{-1} = x_n^{-1} ... x_2^{-1} x_1^{-1}$.

Given any word $\omega \in (S \cup S^{-1})^*$, there is an associated path in the directed and edge labeled Cayley graph $\Gamma_{G,S}$. Similarly, there is a word for every finite edge path in $\Gamma_{G,S}$.

This means that we can think of the Cayley graph $\Gamma_{G,S}$ as a "calculator" for G. More formally, if we have group elements g, h and words ω_g, ω_h that correspond to g, h, then we can find gh by starting at the vertex v_e, following the path corresponding to ω_g, and then from that point follow the path corresponding to ω_h. This will result in the vertex v_{gh} that corresponds to gh.

2 Groups Acting on Trees

2.1 Free Groups

Definition 2.1.1
Let $S = \{x_1, ..., x_n\}$ be a set of elements in a group G. A word $\omega \in \{S \cup S^*\}$ is said to be *freely reduced* if it does not contain a subword consisting of an element adjacent to its formal inverse. The group G is the *free group with basis* S if S is a set of generators for G and no freely reduced word or its inverse represents the identity element. The *rank* of a free group with basis S is the number of elements in S. $\mathbb{F}_n$ is the free group of rank n.

There are also free groups of infinite order, but we will not consider them for now.

Theorem 2.1.2
For any $n \in \mathbb{N}$ there is a free group of rank n.

Proof. (Sketch)
Let $S = \{x_1, ..., x_n\}$ be a set of n distinct symbols. Then we define the equivalence relation on all words in $\{S \cup S^*\}$ induced by

$$a_1 \cdots a_{i-1} a_i a_i^{-1} a_{i+1} \cdots a_k \sim a_1 \cdots a_{i-1} a_{i+1} \cdots a_k$$

where adjacent pairs of elements that are each-others inverses get cancelled. Then the elements of the group $\mathbb{F}_n$ are the equivalence classes of words. It is easy to check that this defines a group structure. The reason that we could not just let $\mathbb{F}_n$ be the set of freely reduced words is that the product of two freely reduced words may not be freely reduced. □

Proposition 2.1.3
The Cayley graph of $\mathbb{F}_2$ with respect to $\{x, y\}$ is an oriented version of T_4.

Proof. Since $\mathbb{F}_2$ is generated by two elements, every vertex in the Cayley graph will have valance 4. Recall that there is a correspondence between freely reduced words and paths in a Cayley graph. Since no non empty freely reduced word is the identity element, we can deduce that no path is a cycle in the Cayley graph. Thus we know that the Cayley graph is a tree and that every vertex has valance 4, so it must be an oriented version of T_4. □

We know that $\mathbb{F}_2$ acts on T_4. Specifically, we can view the action of x as a "shift to the right" of the graph, while the action of y is "a vertical shift upwards".

Note again that visualizing the action of an element of a group on a Cayley graph is not the same as following the path from the identity to that element in the Cayley graph. The action is a right action, while paths are determined through left multiplication.

2.1.1 Free Groups as Subgroups

The general approach for finding free groups as subgroups is often related to the Ping-Pong Lemma. This lemma can be thought of more as a guiding approach for finding free groups inside of other groups. In fact, in the examples, we do not make specific reference to the lemma, but it is clear where it is used.

Lemma 2.1.4 (Ping Pong Lemma)
Let G be a group that acts on X. and let S be a symmetric set of generators (the inverse of an element in S is in S). For each $s \in S$, let X_s be a subset of X, and let p be a point in $X / \cup_{s \in S} X_s$. Then if

1. $s(p) \in X_s$ for all $s \in S$, and
2. $s(X_t)$ is a proper subset of X_s for each $t \neq s^{-1}$,

Then G is a free group with basis S

Proof. Suppose the conditions of the lemma. Then we want to show that any freely reduced word in S^* does not represent the identity. So let $w_1 \dots w_n$ be such a word. Then if we pick some $p \notin X_{w_n^{-1}}$ then we get that $w_n(p) \in X_{w_n}$. Then we can see that

$$\begin{aligned}
&w_{n-1}(p) \in X_{w_n} \\
&(w_{n-1}w_n)(p) \in X_{w_{n-1}}, \\
&\dots, \\
&(w_1 \dots w_n)(p) \in X_{w_1}.
\end{aligned}$$

Since $p \notin X_{w_1}$ this means that $(w_1 \dots w_n)(p) \neq p$ which means that the word can not represent the identity. Thus no freely reduced word represents the identity. □

Proposition 2.1.5

Consider $SL_2(\mathbb{Z})$, the group of two-by-two matrices with integer entries and determinant 1. The subgroup of $SL_2(\mathbb{Z})$ generated by $S = \left\{ \begin{bmatrix} 1 & 0 \\ 2 & 1 \end{bmatrix}, \begin{bmatrix} 1 & 2 \\ 0 & 1 \end{bmatrix} \right\}$ is isomprhif to $\mathbb{F}_2$.

Proof. We could try to show that no reduced word made with elements of S and S^{-1} is equal to the identity matrix directly, but this would be quite difficult. Instead we will do this using the action of $SL_2(\mathbb{Z})$ on the plane. Let $l = \begin{bmatrix} 1 & 0 \\ 2 & 1 \end{bmatrix}$, and $r = \begin{bmatrix} 1 & 2 \\ 0 & 1 \end{bmatrix}$. We want to show that any freely reduced word in $\{l, l^{-1}, r, r^{-1}\}$ does not represent the identity element. We partition the plane into regions, each of which are associated to a generator, as shown below.
I would include a figure from the book but I am scared on copyright stuff. It can be found on page 60 of the book.
Through computations, one can show that if $w \in \{l, l^{-1}, r, r^{-1}\}$, then $w \cdot X_y \subset X_w$ if $y \neq w^{-1}$. This containment is always proper. Now let ω be a freely reduced word. To show that ω is not the identity, it suffices to show that ω acts non-trivially on the plane. Since ω is freely reduced, we can apply the containment property one element at a time:

$$\begin{aligned} \omega \cdot X_{w_n} = w_1 w_2 \cdots w_n \cdot X_{w_n} &\subset w_1 w_2 \cdots w_{n-1} \cdot X_{w_n} \\ &\subset w_1 w_2 \cdots w_{n-2} \cdot X_{w_{n-1}} \subset \cdots \subset X_{w_1} \end{aligned}$$

This means that $\omega \cdot X_{W_n}$ is a proper subset of X_{w_1}, which makes it impossible that $\omega \cdot X_{W_n} = X_{w_n}$. This means that ω is not the identity element. □

Proposition 2.1.6
There is a subgroup F_2 that is a free group of rank n for any $n \in \mathbb{N}$.

Proof. The proof has to do with considering certain forward trees. A forward tree is defined by the word that takes a fixed vertex to the beginning of that tree. Working in the Cayley graph of F_2 with respect to $\{a, b\}$ (which is a tree), if we choose words carefully, then we can apply the ping pong lemma to find a free group of rank n as a subgroup of □

2.2 Free Group Homomorphisms and Group Presentations

Theorem 2.2.1
Let G be a group and let $\{g_1, ..., g_n\}$ be a list of some elements. Then let $S = \{x_1, ..., x_n\}$ be a basis for the free group F_n. Then there is a homomorphism $\phi : F_n \to G$ such that $\phi(x_i) = g_i$.

Proof. Let w be a reduced word in $\{S \cup S^{-1}\}^*$. Then $w = w_1 ... w_k$ where $w_i \in S \cup S^{-1}$. Then we can define $\phi(x_i^{-1}) = g_i^-1$. And define $\phi(w_1 ... w_n) = \phi(w_1) ... \phi(w_n)$. Recall that elements of F_n are actually equivalence classes of words, so in order for this homomorphism to be well defined it must be constant on these equivalence classes. Clearly if there are two consecutive letters in a word that are each others formal inverses, then after applying ϕ these elements will cancel out. Thus ϕ is constant on equivalence classes as desired. Then it is simple to check that $\phi(w)^{-1} = \phi(w)$ and that $\phi(w_1w_2) = \phi(w_1)\phi(w_2)$, which shows that it is a homomorphism. □

Corollary 2.2.2
Any two free groups of rank n are isomorphic.

Proof. Let G be a free group of rank n with basis $\{x_1, ..., x_n\}$ and let H be a fere group of rank n with basis $\{y_1, ..., y_n\}$. Then we can make a homomorphism $\phi : G \to H$ such that $\phi(x_i) = y_i$. And $\psi : G \to H$ such that $\psi(y_i) = x_i$. This means that the composition of the two in any order is the identity homomorphism in the appropriate group, which means that they are each others inverses. Thus they are invertible and bijective, meaning that they are isomorphisms. □

Corollary 2.2.3
If a group G is generated by n elements then it is isomorphic to a quotient group of F_n.

Proof. Let $\phi : F_n \to G$ as in the theorem where the g_i are the generating elements. Then the image of ϕ is G. Thus we can apply the first isomorphism theorem to get that $F_n/\text{Ker}(\phi) \cong G$ □

.

2.2.1 Group Presentations

Definition 2.2.4
Let G be a group with generators $\{g_1, ...g_n\}$, and let $\{x_1, ..., x_n\}$ be a basis for a free group of rank n. Then we know that there is a homomorphism from $\phi : F_n \to G$ such that $\phi(x_i) = g_1$. Any word w such that $\phi(w) = 1_G$ is a *relation.*

For example if a and b generate $\mathbb{Z} \oplus \mathbb{Z}$ then $aba^{-1}b^{-1}$ is a relation.

If $\phi : F_n \to G$ as above, then a subset $R \subset \text{Ker}(\phi)$ is said to be a *set of defining relations* if the smallest normal subgroup of F_n that contains R is $\text{Ker}(\phi)$. Note that the smallest normal subgroup containing R must contain R^{-1} and all conjugates of the form wrw^{-1} where $r \in R \cup R^{-1}$ and $w \in F_n$. In fact the finite product of all conjugates of elements of $R \cup R^{-1}$ is the smallest normal subgroup containing R. Thus if R is a defining set of relation then every element of the kernel can be expressed as a finite product of conjugates of $R \cup R^{-1}$. A group G is said to be *finitely presented* if there is a finite set of defining relations $\{w_1, ...w_k\}$. Then we can present the group as $G = \langle g_1, ..., g_n \mid w_1, ..., w_\rangle$.

2.3 Free Group Actions on Trees

Theorem 2.3.1
A group G is free if and only if it acts freely on a tree.

Proof. Can be found in the book. Not worth writing up. □

Corollary 2.3.2 (Nielsen-Schreir Theorem)
Every subgroup of a free group is a free group.

Proof. Let F_n be the free group of rank n with basis $S = \{x_1, ..., x_n\}$. Then F_n acts freely on the Cayley graph $\Gamma_{F_n,S}$. Then so does any subgroup of F_n. By the theorem, this means that the subgroup is a free group. □

2.4 The Group $\mathbb{Z}_3 * \mathbb{Z}_4$

This is an informal introduction to the idea of free products. The technical definitions are in the next subsection.

Consider the tree $T_{3,4}$, depicted with black vertices (of valance 4) and white vertices (of valance 3) below. If we pick a distinguished black vertex and an adjacent distinguished white vertex, we can consider the group of symmetries generated by a and b, where a is a rotation about the black vertex and b is a rotation about the white vertex. To describe the action of this group on $T_{3,4}$ more concretely, we introduce a labeling of the vertices in the following way.

FIGURE IN PAGE 74

Let the distinguished black vertex be labeled by $\emptyset$. Then label the vertices adjacent to this with North, East, South, and West in a clockwise order. Then each of these vertices will be adjacent to two vertices other than $\emptyset$. We can give a labeling of left or right to each of them. These vertices will in turn be adjacent to 4 vertices. If we assumed that the "direction" that we came to such a vertex is East, then we can label the adjacent vertices. This is depicted below.

FIGURE IN PAGE 74

Now we can explicitly write out the action of a. a fixes $\emptyset$. Any other vertex starts with either a N, E, S, or W. a will cycle the first letter of the labeling (taking labels starting with E to labels starting with E, labels starting with E to labels starting with S, etc.).

The action of b is a bit more complicated. b fixes E. All other labeling of vertices start with $\emptyset$, EL, or ER. b cycles the starting of the vertices (vertices that start with $\emptyset$ to vertices that start with EL, vertices that start with EL to vertices that start with ER, and vertices that start with ER to vertices that start with $\emptyset$).

The group G generated by a and b can be referred to as *the free product* $\mathbb{Z}_3 * \mathbb{Z}_4$. This is because the order of b is 3, and the order of a is 4.

2.5 Free Product of Groups

To describe the free product of groups A and B, denoted $A * B$, we first start with an alphabet of letters from A and B, $\{A \cup B\}$.

Definition 2.5.1
A word $w = x_1 x_2 \cdots x_n \in \{A \cup B\}^*$ is freely reduced if:

1. no $x_i = eA$ or e_B, and
2. when $x_i \in A$, $x_{i+1} \notin A$, and if $x_i \in B$ then $x_{i+1} \notin B$.

We would like to turn the set of freely reduced words into a group under the operation of concatenation. However, the concatenation of two freely reduced words may not be freely reduced. To deal with this we introduce an equivalence relation on $\{A \cup B\}^*$ that is generated by:

1. $we_x w' \sim ww'$ if $e_x = e_A$ or e_B, and
2. $wabw' \sim wcw'$ if $ab = c$ in either A or B.

Definition 2.5.2
$A * B$ is the group whose elements are equivalence classes of elements of $\{A \cup B\}$. The operation is defined by taking the equivalence class of the concatenation of representatives of equivalence classes, i.e. $[w_1][w_2] = [w_1 w_2]$.

It is important to check that this is well defined (meaning that if different representatives of equivalence classes are chosen then the result of the operation does not change).

Theorem 2.5.3
Each equivalence class $[w]$ contains exactly one reduced word of $\{A \cup B\}^*$.

Proof. The proof is gross and breaks into many cases. First we show that every word is equivalent to a freely reduced word. Then we prove that two different freely reduced words are not equivalent. This proves the theorem. If you wish you can find the proof in the book. □

Theorem 2.5.4
Every free product of groups $A * B$ can be realized as a group of symmetries of a bi-regular tree T, and the fundamental domain of this action is a single edge and its two vertices. If A and B are finite then the tree will be $T_{|A|,|B|}$.

Proof. We will first construct a graph which admits an action of $A * B$. We will then show that the fundamental domain of this action is an edge with its vertices. Then we show that the graph is actually a tree, and finally we prove the assertion about A and B being finite.
Let the vertices $V(T)$ correspond to the left cosets gA and gB, i.e. every vertex can be labeled as v_{gA} or v_{gB}. The edges of T will correspond to the elements, so any $g \in A * B$ corresponds to the edge e_g. Note that this will make the edge associated to $e \in A * B$ be e_e. :(We define $\textsc{Ends}(e_g) = \{v_{gA}, v_{gB}\}$. This means that two edges e_g and g_h intersect if and only if $gA = hA$ or $gB = hB$, which happens if and only if $g^{-1}h \in A$ or $g^{-1}h \in B$. We can see that T has a bipartite structure by noting that there are two classes of vertices, those of the form v_{gA} and those of the form v_{gB}. By the definition of the edges, vertices of one class are only adjacent to vertices of the other class.
The action of $A * B$ on T is induced by the left action of $A * B$ on itself and its cosets. Specifically, $gv_{hA} = v_{ghA}$ and $gv_{gB} = v_{ghB}$. This will induce that $ge_h = e_{gh}$.
To find the fundamental domain, we use three facts:

1. the group action preserves the class of vertices,
2. edges only go between vertices of different classes, and
3. the action is "vertex transitive with in classes".

Thus the fundamental domain must contain exactly one vertex of each class. Furthermore, the fundamental domain must contain a full edge (if the fundamental domain contained just part of one edge, then there would be no way to cover the rest of the edge because symmetries preserve the bipartite structure). Thus the smallest fundamental domain would be an edge and its ends.
To prove that T is a tree we need to prove that it is connected and that it has no circuits. To see that it is connected, consider and $g \in A * B$. Then $g = [w]$ for some reduced $w = x_1 x_2 \cdots x_n \in \{A \cup B\}^*$. Assume that $x_1 \in A$ (very similar arguments will hold for $x_1 \in B$). Then $e_e \cap e_{x_1} = v_A$, $e_{x_1} \cap e_{x_1 x_2} = v_{x_1 B}$, $e_{x_1 x_2} \cap e_{x_1 x_2 x_3} = v_{x_1 x_2 A}$, and so on. Then we get that $\{e_e, e_{x_1}, e_{x_1 x_2}, \cdots, e_{x_1 x_2 \cdots x_n}$ is an edge path from e_e to e_g. Thus every edge is connected to the edge e_e and the graph T is connected.
To prove that T is a tree, suppose for contradiction that there is some circuit. We can use the group actio nof $A * B$ to move this circuit to one that begins and ends at e_e. So suppose that there is some circuit $e_e, e_2, e_2, \cdots, e_n, e_e\}$. Then $e_1 = e_{x_1}$ for some $x_1 \in A$ or B. Suppose that x_1 is in A. Then $e_2 = e_{x_1 x_2}$ for some $x_2 \in B$, and in general $e_k = e_{x_1 x_2 \cdots x_k}$ where $x_1 x_2 \cdots x_k$ is some freely reduced word in $\{A \cup B\}$. Since the edge path is a circuit, $e_{x_1 \cdots x_n} \cap e_e \neq \emptyset$. Specifically, this intersection is either v_A or v_B. This means that $[x_1 x_2 \cdots x_n] = [x]$ for some $x \in A \cup B$. But now we have found two freely reduced words in one equivalence class, which contradicts the previous theorem. Thus T must be free of circuits, and is a tree.
Now suppose that A and B are finite. Then $v_{gA} \in \text{Ends}(e_h)$ if and only if $h \in gA$. This means that the number of edges that contain a vertex of the form v_{gA} is the same as the number of elements of gA, which is just $|A|$. The same arguments hold for vertices of the form v_{gB}. $\square$

In the example in the previous subsection, we made use of the notation of the previous proof. To see this simple perform the replacements

$$E \mapsto e, S \mapsto a, W \mapsto a^2, N \mapsto a^3, L \mapsto b, R \mapsto b^2.$$

We can actually illustrate the proof geometrically in this case. The cosets $g\mathbb{Z}_4$ correspond to sets of 4 white vertices that all are adjacent to one white vertex. The cosets of $g\mathbb{Z}_3$ correspond to sets of 3 white vertices that are all adjacent to one black vertex. e_e will be the edge joining the vertex fixed by the action of a to the vertex that is fixed by the action of b ($\emptyset$ and E, respectively). Then the resulting graph will be connected (because we can apply a series of rotations to transform e_e to an arbitrary edge), and we can also see that the graph must be a tree because it would be impossible to have a reduced non-trivial composition of rotations about $\emptyset$ and rotations about E that would be the identity transformation.

FINISH SUBSECTION FROM BOOK HERE

2.6 Free Products of Finite Groups are Virtually Free

DO THIS SECTION!!!!

2.7 Finite Groups Acting on Trees

Definition 2.7.1
We say that a grop G has *Serre's Property FA* if whenever there is an action of G on a tree T, there is a fixed point $x \in T$ ($\mathrm{Stab}_G(x) = G$).

Theorem 2.7.2
All finite groups have Serre's Property FA.

Proof. Let G bea group and T be a tree and $G \curvearrowright T$ a group action of G on T. Then let v be any vertex of T. Consider the orbit of v, $\mathcal{O} = \{g \cdot v \mid g \in G\}$. Then let $T_{\mathcal{O}}$ be the subtree formed by taking the union of all vertices and edges in the minimal-length paths containing elements of $\mathcal{O}$.
Since G is finite, $\mathcal{O}$ is finite, and so is $T_{\mathcal{O}}$. This will mean that there are some leaves (vertices of valance 1) in this subtree. We can also see that G clearly acts on $T_{\mathcal{O}}$. This means that G takes leaves of the subtree to leaves, and non-leaves to non-leaves. So we can form the subsubtree T^1 which is formed by removing all of the leaves and incident edges from $T_{\mathcal{O}}$, and G will act on T^1. We can continue to do this for T^2, T^3, and so on. Since $T_{\mathcal{O}}$ has finitely many vertices and we are removing vertices at each step, this process will eventually result in some T^i that is either juts one vertex or an edge and its vertices.
If T^i is just one vertex, then since G acts on T^i, G fixes this vertex. If T^i is an edge and its vertices, then the action of and $g \in G$ must either do nothing to T^i or must swap the vertices of T^i. Either way, the midpoint of the edge is fixed by the action of G. $\square$

In the proof we saw that either a vertex or the midpoint of an edge is fixed by the action of a finite group on a tree. If the tree has a bipartite structure, then the second case is impossible. Thus we get

Corollary 2.7.3
If a finite group G acts on a bipartite tree T, then G fixes a vertex of T.

Corollary 2.7.4
If H is a finite subgroup of $A * B$ then H is conjugate to a subgroup of A or of B.

Proof. In the proof of theorem 2.5.4 we constructed a bipartite tree that $A * B$ acts on. If H is a finite subgroup then it fixes a vertex v of this tree, meaning that H is contained in the $\mathrm{stab}_{A*B}(v)$. The stabilizer of vertices was shown to be conjugates of A or of B, so H is contained in a conjugate of A or of B. I THINK that subgroups of the conjugates are conjugates of the subgroups, which proves the assertion. $\square$