TOPICS IN GEOMETRIC GROUP THEORY

SAMEER KAILASA

ABSTRACT. We present a brief overview of methods and results in geometric group theory, with the goal of introducing the reader to both topological and metric perspectives. Prerequisites are kept to a minimum: we require only basic algebra, graph theory, and metric space topology.

CONTENTS

Date: Submitted Sept. 6, 2014.

1. Free Groups

Geometric group theory studies abstract groups via their realizations as concrete geometric or topological objects, and their group actions on these objects. The *free groups* are directly amenable to such an approach.

A group F is free if no nontrivial relations hold between the elements of F. Equivalently, any two elements are F are considered distinct unless their equality follows directly from the cancellation of inverses. For example, we might think of the free group on two generators, denoted $\langle a, b\rangle$. In this group, the elements $a^2ba^{-1}b^2$ and b^3a are different, while the elements $baa^{-1}b$ and b^2 are the same.

Free groups can also be described by a more technical characteristic property. We will proceed by defining free groups in terms of this property and showing that the two notions are equivalent.

Definition 1.1. A group F is a *free group* if there is a set $S \subset F$ with the following property: for any group G and function $f : S \to G$, there is a unique group homomorphism $\hat{f} : F \to G$ such that $\hat{f}|_S = f$. The set S is said to *freely generate* F.

$$\begin{tikzcd} S \arrow[r, "f"] \arrow[d, "\iota"'] & G \\ F \arrow[ur, "\hat{f}"'] & \end{tikzcd}$$

Proposition 1.2. *If S freely generates F, then S generates F.*

Proof. Let $\langle S\rangle$ denote the group generated by F, and let $\iota : S \to \langle S\rangle$ denote the inclusion map. By Def. 1.1, ι extends to a unique homomorphism $\hat{\iota} : F \to \langle S\rangle$. Another homomorphism φ with $\varphi|_S = \iota$ is the identity. By uniqueness, $\hat{\iota} = \varphi$. □

Proposition 1.3. *Let F and G be free groups, freely generated by S and R respectively. If $|S| = |R|$, then $F \cong G$.*

Proof. Since $|S| = |R|$, there exists a bijection $f : S \to R$; let g be its inverse. By Def. 1.1, these extend to homomorphisms $\hat{f} : F \to G$ and $\hat{g} : G \to F$. Then $\hat{g} \circ \hat{f} : F \to F$ extends $g \circ f$. By uniqueness, $\hat{g} \circ \hat{f} = \mathrm{id}_F$. Likewise, $\hat{f} \circ \hat{g} = \mathrm{id}_G$. □

Definition 1.4. Although we will not prove it, the converse of Prop. 1.3 also holds. Thus, we may define the *rank* of F freely generated by S as the cardinality $|S|$.

Prop. 1.3 shows that there is at most one free group of a given rank, up to isomorphism. However, we are left to construct a single example of a free group. If we can do so, we will have described all the free groups.

Let S be a set. We define the set S^{-1} as the set of symbols s^{-1} where $s \in S$ and the map $s \mapsto s^{-1}$ is bijective. Furthermore, let $\{1\}$ denote a singleton set with $1 \notin S \cup S^{-1}$.

Definition 1.5. A *word* on S is a sequence $w = (s_1, s_2, \cdots)$ where each $s_i \in S \cup S^{-1} \cup \{1\}$ and there exists K such that $s_k = 1$ for $k \geq K$. The word $(1, 1, 1, \cdots)$ is called the empty word.

Remark 1.6. We will write a word w as an expression of the form

$$w = {s_1}^{a_1} {s_2}^{a_2} \cdots {s_\ell}^{a_\ell}$$

where each $s_i \in S$ and $a_i = \pm 1$ or 0, with $a_\ell \neq 0$. This should be thought of as simply a piece of notation. The spelling of a given word is unique, since equality of sequences requires equality of each term in the sequence. Therefore, thinking of a word as the product of elements in a group could be erroneous. After all, nontrivial relations could hold within the group, leading to non-unique factorization.

Definition 1.7. For a word $w = {s_1}^{a_1}{s_2}^{a_2}\cdots{s_\ell}^{a_\ell}$, we say the *length* of w is ℓ. The length of the empty word is defined to be 0.

Definition 1.8. The word w is *reduced* if either w is empty or $w = {s_1}^{a_1}{s_2}^{a_2}\cdots{s_\ell}^{a_\ell}$ where all the $a_i = \pm 1$ and for all $1 \leq i \leq \ell - 1$, we have ${s_i}^{a_i} \neq {s_{i+1}}^{-a_{i+1}}$.

Definition 1.9. We may specify a product on the set of reduced words. Let w and u be reduced words. Then $w = w'v$ and $u = v^{-1}u'$, where v is the maximal length (possibly empty) sub-word at the end of w which cancels at the beginning of u. Then the *juxtaposition* of w and u is defined as $w'u'$. Since v is of maximal length, $w'u'$ is itself a reduced word.

Theorem 1.10. *Let S be a set. Then, there exists a free group F freely generated by S.*

Proof. As you might expect us to begin, let F denote the set of reduced words on S. Then, F is a group under the juxtaposition operation. To prove this, we follow a method due to van der Waerden, as described in [1].

For each $s \in S$, consider the function $|s| : F \to F$ defined as follows: $|s|(x)$ is the juxtaposition of s and x. Analogously define $|s^{-1}| : F \to F$ for each $s \in S$. Note that $|s|$ and $|s^{-1}|$ commute, and are inverses. Hence each $|s|$ gives rise to a permutation of F. If S_F is the symmetric group on F, let $\langle[S]\rangle$ denote the subgroup of S_F generated by $[S] = \{|s| : s \in S\}$.

Given an arbitrary element $r \in \langle[S]\rangle$, we can factorize $r = |s_1^{a_1}| \circ |s_2^{a_2}| \circ \cdots \circ |s_\ell^{a_\ell}|$, where each $a_i = \pm 1$, each $|{s_i}^{a_i}| \in [S]$, and for all $1 \leq i \leq \ell - 1$, we have $|{s_i}^{a_i}| \neq |{s_{i+1}}^{-a_{i+1}}|$. Such a factorization must be unique, since $g(1) = {s_1}^{a_1}{s_2}^{a_2}\cdots{s_\ell}^{a_\ell}$ is a reduced word with unique spelling.

Now, for a group G, consider an arbitrary function $f : [S] \to G$. Then, define $\hat{f} : \langle[S]\rangle \to G$ as

$$\hat{f}(|s_1^{a_1}| \circ |s_2^{a_2}| \circ \cdots \circ |s_\ell^{a_\ell}|) = f(|s_1^{a_1}|)f(|s_2^{a_2}|)\cdots f(|s_\ell^{a_\ell}|)$$

If $w, u \in \langle[S]\rangle$ are such that $w \circ u(1)$ is reduced, then clearly $\hat{f}(w \circ u) = \hat{f}(w)\hat{f}(u)$. Given this fact, it is routine to show $\hat{f}$ is a homomorphism. Since $[S]$ generates $\langle[S]\rangle$, all homomorphisms agreeing on $[S]$ must agree everywhere. Thus, $\hat{f}$ is the unique extension of f and $\langle[S]\rangle$ is freely generated by $[S]$.

Finally, note that $|s_1^{a_1}| \circ |s_2^{a_2}| \circ \cdots \circ |s_\ell^{a_\ell}| \mapsto {s_1}^{a_1}{s_2}^{a_2}\cdots{s_\ell}^{a_\ell}$ gives a bijection between $\langle[S]\rangle$ and F. Inheriting the group structure from $\langle[S]\rangle$ via this bijection, we may regard F as isomorphic to $\langle[S]\rangle$. Thus, F is freely generated by S. □

By Prop. 1.4, we may conclude that free groups are exactly as initially described.

2. Cayley Graphs

How can we endow a group with geometric or topological structure? One fundamental approach is via the construction of a Cayley graph.

Definition 2.1. Let G be a group and S a generating subset of G. The *Cayley graph* $C(G, S)$ can be described as follows. We set the elements of G as the vertices of $C(G, S)$. For g, $h \in G$, there is a directed edge from g to h if $h = gs$ for some $s \in S$. Note the resultant graph should be connected, since S generates G.

Remark 2.2. Suppose $h = gs$ for some $s \in S$. Then $g = hs^{-1}$. If both s and s^{-1} are in S, we can just draw an undirected edge between g and h. Otherwise, the edge is directed, in the orientation described above. But really, there is no need to dwell on this distinction. Whenever it feels agreeable to have an undirected Cayley graph, think of the generating set as containing formal inverses of all its elements.

Example 2.3. We can draw the Cayley graph $C(G, S)$, where G is the group of integers modulo 6 and $S = \{1, 2\}$:

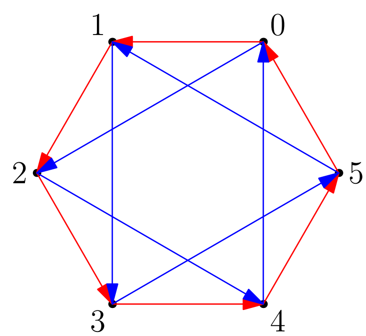

As another example, it is natural to ask what the Cayley graph of a free group should look like.

Proposition 2.4. *Let F be a free group, freely generated by S. Then $C(F, S)$ is a tree.*

Proof. It suffices to show $C(F, S)$ has no cycles. Suppose $g_0 \to g_1 \to g_2 \cdots \to g_{k-1} \to g_k = g_0$ is a cycle in $C(F, S)$, where the given sequence is of consecutive vertices in the cycle. Without loss of generality, we assume these vertices are all distinct (except $g_k = g_0$). Then, there exist $s_1, \cdots, s_k \in S$ such that $g_i = g_{i-1}s_i$ for $1 \leq i \leq k$. It follows $g_0 = g_0 s_1 s_2 \cdots s_k$. For this to be the case, there must be some j such that $s_{j+1} = {s_j}^{-1}$. But then $g_{j+1} = g_{j-1}$.

$\square$

We will work more with Cayley graphs later in the exposition.

3. Baby Algebraic Topology and the Nielsen-Schreier Theorem

We wish to survey geometric and topological methods in group theory. To begin, in this section, we will prove the Nielsen-Schreier Theorem.

Theorem 3.1 (Nielsen-Schreier)**.** *Every subgroup of a free group is free.*

Our approach will be topological. In algebraic topology, one associates a given topological space with a group, called the fundamental group. The construction is useful precisely because subgroups of the fundamental group correspond to different ways in which one may *cover* the topological space (this will soon be made precise).

Consequently, the topology of the space may be studied directly by examining the structure of the fundamental group.

Here, we will develop a sort of "baby" algebraic topology, in which we determine objects analogous to those in true algebraic topology, but over graphs rather than topological spaces. We follow the exposition given in [2].

3.1. **Definitions and the Fundamental Group of a Graph.** To begin, we give some definitions in graph theory. The formalism of graph theory given here is slightly different from the usual, for we will need some way to encode the orientation of an edge within the edge itself.

Definition 3.2. A *graph* $\mathcal{G}$ consists of

- a *vertex set* V,
- an *oriented edge set* E,
- an *edge reversal map* from $E \to E$ given by $e \mapsto e^-$ such that $e^{--} = e$ for every $e \in E$,
- and an *initial vertex map* $\text{init} : E \to V$.

In other words, every edge $e \in E$ has a formal inverse edge e^- with the opposite orientation. We also define the *final vertex map* $\text{fin} : E \to V$ as $\text{fin}(e) = \text{init}(e^-)$.

Definition 3.3. A *path* of length ℓ in $\mathcal{G}$ is a sequence of oriented edges $p = (e_1, e_2, \cdots, e_\ell)$ such that $\text{init}(e_{i+1}) = \text{fin}(e_i)$ for $1 \le i \le \ell - 1$. The clear interpretation of a path is as a finite walk on the graph.

Remark 3.4. Given a vertex v, we will also define the *empty path at v* as the path beginning at v but traversing no edges (the "do nothing" walk). We write the empty path as 1_v.

Definition 3.5. The path $p = (e_1, e_2, \cdots, e_\ell)$ is *closed* if $\text{init}(e_1) = \text{fin}(e_\ell)$. The *inverse* of p is the path given by $p^{-1} = ({e_\ell}^{-1}, \cdots, {e_2}^{-1}, {e_1}^{-1})$. For two paths p and q, where the final point of p is the initial point of q, we may *concatenate* the paths, writing pq to denote the path obtained by traversing p, then q. For a path p and edge e, we write the concatenation of the paths p and (e) as pe.

Definition 3.6. A *spur* is a path of the form (e, e^{-1}), where e is an edge. In other words, a spur involves traversing an edge, then immediately retracing one's step. A path is *reduced* if it contains no spur as a sub-path. Compare this to the definition of a reduced word from the first section.

Definition 3.7. Suppose we are given a path p. An *elementary move* on p is the insertion or deletion of a spur between successive edges, or at an endpoint, of p. For example, given the path $p = (e_1, e_2, {e_2}^{-1}, e_3)$, an elementary move could take us to the path $p' = (e_1, e_2, e_4, {e_4}^{-1}, {e_2}^{-1}, e_3)$ by insertion, or to the path $p' = (e_1, e_3)$ by deletion.

Definition 3.8. Two paths p and p' are said to be *homotopic* if one can be obtained from the other by a finite sequence of elementary moves.

It is straightforward to check that homotopy is an equivalence relation. For a path p, let $[p]$ denote its equivalence class with respect to the homotopy relation. Given a vertex v of the graph $\mathcal{G}$, we then define

$$\pi_1(\mathcal{G}, v) = \{[p] : p \text{ is a closed path beginning and ending at } v\}$$

This set is called the *fundamental group* of $\mathcal{G}$ at v. Of course, we have yet to assign it a group structure.

Proposition 3.9. *$\pi_1(\mathcal{G}, v)$ is a group.*

Proof. Define the product of equivalence classes $[p_1]$ and $[p_2]$ as $[p_1] \cdot [p_2] = [p_1 p_2]$. Associativity of this product is clear. Note that $[p] \cdot [1_v] = [1_v] \cdot [p] = [p]$ and $[p] \cdot [p^{-1}] = [p^{-1}] \cdot [p] = [1_v]$. □

Example 3.10. Let $\mathcal{G}$ be a tree. Then $\pi_1(\mathcal{G}, v)$ is the trivial group.

Example 3.11. Let $\mathcal{G}$ be isomorphic to a polygon. Then $\pi_1(\mathcal{G}, v)$ is $\mathbb{Z}$.

3.2. **Realizing Free Groups as Fundamental Groups.** Since we would like to study free groups, we will need to construct graphs whose fundamental groups are free. In order to do so, we will need the following graph theoretic preliminary:

Definition 3.12. Let $\mathcal{G}$ be a connected graph. A *spanning tree* is a subgraph of $\mathcal{G}$ which is a tree and includes all the vertices of $\mathcal{G}$.

Proposition 3.13. *Every connected graph has a spanning tree.*

Proof. We may partially order the sub-trees of $\mathcal{G}$ by inclusion. For any chain of sub-trees, the union of all the trees in the chain is an upper bound. By Zorn's Lemma, there is a maximal sub-tree; call it $\mathcal{T}$. If some vertex is not included in $\mathcal{T}$, then we can add it on, since $\mathcal{G}$ is connected. This contradicts the maximality of $\mathcal{T}$, hence $\mathcal{T}$ spans $\mathcal{G}$. □

Now, we have the tool necessary to make an important observation regarding the fundamental groups of connected graphs and free groups. Namely,

Lemma 3.14. *The fundamental group of a connected graph is free.*

Proof. Let $\mathcal{G}$ be a connected graph, with v a vertex of $\mathcal{G}$. Let $\mathcal{T}$ be a spanning tree of $\mathcal{G}$. We will use $\mathcal{T}$ to determine a representation of each equivalence class in $\pi_1(\mathcal{G}, v)$.

Let e_i be an oriented edge of $\mathcal{G}$, with $\text{init}(e_i) = v_j$, $\text{fin}(e_i) = v_k$. In $\mathcal{T}$, there is a unique reduced path p_j from v to v_j, and a unique reduced path p_k from v to v_k. Set $c_i = p_j e_i {p_k}^{-1}$. Now, suppose we have an arbitrary closed path $p = (e_1, e_2, \cdots, e_\ell)$. Then p is homotopic to the path $p' = c_1 c_2 \cdots c_\ell$.

We therefore consider the equivalence classes $[c_i]$. Notice that if $e_i \in \mathcal{T}$, then c_i is homotopic to 1_v. Hence, we may ignore the equivalence classes given by $e_i \in \mathcal{T}$. Otherwise, any class in $\pi_1(\mathcal{G}, v)$ can be written as the product of $[c_i]$'s, where $e_i \notin \mathcal{T}$. Thus $S = \{[c_i] : e_i \notin \mathcal{T}\}$ generates $\pi_1(\mathcal{G}, v)$.

It remains to show S freely generates $\pi_1(\mathcal{G}, v)$. Suppose some nontrivial relation $[c_1][c_2] \cdots [c_\ell] = [1_v]$ holds, where the path $c_1 c_2 \cdots c_\ell$ is reduced. Then 1_v is homotopic to $p_x(e_1, e_2, \cdots, e_\ell){p_y}^{-1}$ where p_x and p_y are the remnant paths from the construction. Since p_x and p_y are paths in $\mathcal{T}$ and the $e_i \notin \mathcal{T}$, none of the edges can cancel with the remnant paths. But something must cancel, so $e_{i+1} = {e_i}^{-1}$ for some i. Then $c_1 c_2 \cdots c_\ell$ would not have been reduced.

□

The construction in Lem. 3.14 gives us an idea. If we imagine the chosen v as a central vertex, then each path c_i can be thought of as a single loop-edge, beginning and ending at v, of some graph. Equivalently, we could *collapse* the spanning tree

$\mathcal{T}$ into one vertex v, and then each $e_i \notin \mathcal{T}$ can once again be thought of as a single loop-edge from v to v. It is precisely this quality that gives the fundamental group its freeness. Any closed path is uniquely encoded by the sequence of loop-edges traversed. Consequently,

Definition 3.15. A *bouquet of circles* is a graph $\mathcal{G} = (V, E)$ where the vertex set $V = \{v\}$ is a singleton and the edge set E consists solely of edges which begin and end at v.

Proposition 3.16. *Let $\mathcal{G} = (V, E)$ be a bouquet of circles. Then $\pi_1(\mathcal{G}, v)$ is isomorphic to the free group of rank $|E|$.*

3.3. **Covering Graphs.** We've managed to represent any free group as the fundamental group of some graph. However, if we wish to prove Thm. 3.1, we need a way to access information about the subgroups of the fundamental group of a graph. The machinery affording us this privilege is that of covering graphs.

Definition 3.17. Let $\mathcal{G}$ be a graph, and $v \in \mathcal{G}$ a vertex. The *neighborhood* of v is the subgraph consisting of v, the vertices adjacent to v, and the edges incident with v.

Definition 3.18. Let $\mathcal{G}$ and $\tilde{\mathcal{G}}$ be connected graphs. Suppose there exists a surjective map $f : \tilde{\mathcal{G}} \to \mathcal{G}$, taking vertices to vertices and oriented edges to oriented edges, with the following properties:

(1) The map f preserves orientation. For an oriented edge $\tilde{e} \in \tilde{\mathcal{G}}$, we have $f(\mathrm{init}(\tilde{e})) = \mathrm{init}(f(\tilde{e}))$ and $f(\mathrm{fin}(\tilde{e})) = \mathrm{fin}(f(\tilde{e}))$.
(2) The map f preserves inverse paths. For an oriented edge $\tilde{e} \in \tilde{\mathcal{G}}$, we have $f(\tilde{e}^{-1}) = f(\tilde{e})^{-1}$.
(3) For each $\tilde{v} \in \tilde{\mathcal{G}}$, f maps the neighborhood of $\tilde{v}$ in $\tilde{\mathcal{G}}$ bijectively onto the neighborhood of $f(\tilde{v})$ in $\mathcal{G}$.

Then f is called a *covering map*, and $\tilde{\mathcal{G}}$ is said to *cover* $\mathcal{G}$.

Remark 3.19. The third condition may be thought of saying that f is a "local isomorphism" between $\tilde{\mathcal{G}}$ and $\mathcal{G}$.

We now prove a few properties of covering maps. For the following, assume $\tilde{\mathcal{G}}$ covers $\mathcal{G}$, with covering map f.

Lemma 3.20. *The following propositions are true.*

(1) Fix a vertex $v \in \mathcal{G}$ and a vertex $\tilde{v} \in f^{-1}(v)$. Then for any path p starting at v, there is a unique path $\tilde{p}$ starting at $\tilde{v}$ whose image covers p.
(2) Let v, w be vertices in $\mathcal{G}$. Then $|f^{-1}(v)| = |f^{-1}(w)|$.
(3) If $\tilde{p}$ is a path in $\tilde{\mathcal{G}}$ whose image $f(\tilde{p})$ is a spur, then $\tilde{p}$ is itself a spur.

Proof. (1) Let $p = (v, e_1, v_2, \cdots, v_\ell)$. By Def. 3.18, there is a bijection between the neighborhood of v and the neighborhood of $\tilde{v}$. Thus, there is a unique edge $\tilde{e_1}$ incident with $\tilde{v}$ such that $f(\tilde{e_1}) = e_1$, and a unique $\tilde{v_2}$ adjacent to $\tilde{v}$ such that $f(\tilde{v_2}) = v_2$. Continuing this argument with $\tilde{v_2}$ in place of $\tilde{v}$ allows us to uniquely "lift" the path p to $\tilde{\mathcal{G}}$.

(2) Let p be a path from v to w. By part (1), for each $\tilde{v} \in f^{-1}(v)$, p lifts to a unique path from $\tilde{v}$ to some vertex in $f^{-1}(w)$. If two of these lifted paths end at the same vertex $\tilde{w} \in f^{-1}(w)$, then the lift of p^{-1} at $\tilde{w}$ is not unique, contradicting

part (1). We therefore have an injective map from $f^{-1}(v)$ into $f^{-1}(w)$. Now switch the roles of v and w.

(3) Note $\tilde{p}$ is the unique lift of the path $f(\tilde{p})$. The result follows from the local isomorphism condition on covering maps, by the same logic as in part (1). □

Definition 3.21. The cardinality of the preimage of any vertex $v \in \mathcal{G}$ is called the *sheet number* of the cover $\tilde{\mathcal{G}}$.

Let us fix a vertex $v \in \mathcal{G}$, and also fix $\tilde{v} \in f^{-1}(v)$. By part (3) of Lem. 3.20, we see that the lifts of homotopic paths in $\mathcal{G}$ are homotopic in $\tilde{\mathcal{G}}$. This furnishes a bijection between homotopy classes $[p] \in \pi_1(\mathcal{G}, v)$ and homotopy classes $[\tilde{p}]$ of covering paths in $\tilde{\mathcal{G}}$ with initial vertex $\tilde{v}$.

If we restrict this bijection to the classes $[\tilde{p}]$ of *closed* covering paths in $\tilde{\mathcal{G}}$ with initial vertex $\tilde{v}$, then we get an injective map $f_\star : \pi_1(\tilde{\mathcal{G}}, \tilde{v}) \to \pi_1(\mathcal{G}, v)$. In particular, $f_\star$ is defined as $f_\star([\tilde{p}]) = [f(\tilde{p})]$, where $f(\tilde{p})$ denotes the image of $\tilde{p}$ in $\mathcal{G}$. However, by conditions (1) and (2) of Def. 3.18, the image $f(\tilde{p}\tilde{q})$ of arbitrary paths $\tilde{p}$ and $\tilde{q}$ in $\tilde{\mathcal{G}}$ is equal to the concatenation $f(\tilde{p})f(\tilde{q})$. It follows that $f_\star$ is in fact a group homomorphism. Therefore,

Proposition 3.22. *The fundamental group $\pi_1(\tilde{\mathcal{G}}, \tilde{v})$ is isomorphic to a subgroup of $\pi_1(\mathcal{G}, v)$.*

3.4. **Realizing Subgroups of Free Groups as Fundamental Groups.** The fundamental group $\pi_1(\tilde{\mathcal{G}}, \tilde{v})$ consists of the homotopy classes of covering paths which begin and end at $\tilde{v}$. However, there are homotopy classes of covering paths in $\tilde{\mathcal{G}}$ which begin at $\tilde{v}$, as we required, but *end* at some other member of $f^{-1}(v)$. The following proposition classifies homotopy classes of covering paths according to their endpoint in $f^{-1}(v)$.

Definition 3.23. For $[p]$, $[q] \in \pi_1(\mathcal{G}, v)$, we write $[p] \sim [q]$ if the unique lifted paths $\tilde{p}$ and $\tilde{q}$ (which both begin at $\tilde{v}$) end at the same vertex in $f^{-1}(v)$. It is clear $\sim$ is an equivalence relation. Let $\pi_1(\mathcal{G}, v)/\sim$ denote the set of equivalence classes in $\pi_1(\mathcal{G}, v)$ with respect to $\sim$.

Proposition 3.24. *There is a bijection between the right cosets of $\pi_1(\tilde{\mathcal{G}}, \tilde{v})$ in $\pi_1(\mathcal{G}, v)$ and $\pi_1(\mathcal{G}, v)/\sim$.*

Proof. Suppose $[p]$ and $[q]$ are in the same right coset. Then $[p] = [r][q]$ for some $[r] \in f_\star(\pi_1(\tilde{\mathcal{G}}, \tilde{v}))$. Lifting to $\tilde{\mathcal{G}}$ gives $[\tilde{p}] = [\tilde{r}][\tilde{q}]$, where $\tilde{r}$ begins and ends at $\tilde{v}$. Therefore, $\tilde{p}$ and $\tilde{q}$ end at the same point.

Suppose covering paths $\tilde{p}$ and $\tilde{q}$ both have the same endpoint in $f^{-1}(v)$. Then the path $\tilde{p}\tilde{q}^{-1}$ is closed, so $[\tilde{p}][\tilde{q}]^{-1}$ is in $\pi_1(\tilde{\mathcal{G}}, \tilde{v})$. Thus $[p][q]^{-1} \in f_\star(\pi_1(\tilde{\mathcal{G}}, \tilde{v}))$, so $[p]$ and $[q]$ are in the same right coset.

□

Note that this also implies

Proposition 3.25. *The sheet number of $\tilde{\mathcal{G}}$ is the index of $\pi_1(\tilde{\mathcal{G}}, \tilde{v})$ in $\pi_1(\mathcal{G}, v)$.*

Now, let F be a free group and $\mathcal{G} = (V, E)$ its realization as a bouquet of circles, where $V = \{v\}$. We may identify $F \cong \pi_1(\mathcal{G}, v)$.

Suppose H is a subgroup of F. We wish to construct a graph $\tilde{\mathcal{G}} = (\tilde{V}, \tilde{E})$ where $\tilde{\mathcal{G}}$ covers $\mathcal{G}$ and $\pi_1(\tilde{\mathcal{G}}, \tilde{v})$ is isomorphic to H for some $\tilde{v} \in \tilde{V}$.

Since $\mathcal{G}$ only has one vertex, every vertex in $\tilde{\mathcal{G}}$ must cover v. By Prop. 3.24, there should be a one-to-one correspondence between vertices in $\tilde{V}$ and right cosets of H in F. Therefore, from each right coset of H in F, we choose a representative path class $[p_\alpha] \in F \cong \pi_1(\mathcal{G}, v)$. Then, we define $\tilde{V} := \{\tilde{v_\alpha}\}$.

Next, we build edges in $\tilde{\mathcal{G}}$. As $\tilde{G}$ is intended to be a covering, the unique lifting of paths should be possible, as in Lem. 3.20. Therefore, for any edge $e_\beta \in E$ and vertex $\tilde{v_\alpha} \in \tilde{V}$, we must have an oriented edge $\tilde{e_\beta}^{(\alpha)} \in \tilde{E}$ with $\mathrm{init}(\tilde{e_\beta}^{(\alpha)}) = \tilde{v_\alpha}$.

In fact, $\mathrm{fin}(\tilde{e_\beta}^{(\alpha)})$ is uniquely determined. Since H is a right coset of itself, there is a vertex $\tilde{v} \in \tilde{V}$ corresponding to H. Here, $\tilde{v}$ plays the role of the basepoint for the fundamental group of the covering graph. If $\tilde{p}$ is a path from $\tilde{v}$ to $\tilde{v_\alpha}$, then the path class $[p]$ is in $H[p_\alpha]$. Thus, the path class $[pe_\beta]$ is in $H[p_\alpha e_\beta]$. The right coset $H[p_\alpha e_\beta]$ has a unique representative $[p_\gamma]$, and since the lift $\tilde{p}\tilde{e_\beta}^{(\alpha)}$ ends at $\mathrm{fin}(\tilde{e_\beta}^{(\alpha)})$, we see that $\mathrm{fin}(\tilde{e_\beta}^{(\alpha)}) = \tilde{v_\gamma}$.

For every α, the lifted path $\tilde{p_\alpha}$ goes from $\tilde{v}$ to $\tilde{v_\alpha}$. Hence $\tilde{\mathcal{G}}$ is connected. By our construction, it follows that $\tilde{\mathcal{G}}$ covers $\mathcal{G}$, with covering map $f : \tilde{\mathcal{G}} \to \mathcal{G}$ given by $f(\tilde{v_\alpha}) = v$ for vertices, and $f(\tilde{e_\beta}^{(\alpha)}) = e_\beta$ for edges. Furthermore, it is evident that our construction ensures H is isomorphic to $\pi_1(\tilde{\mathcal{G}}, \tilde{v})$.

Therefore,

Lemma 3.26. *Let F be a free group and $\mathcal{G}$ its realization as a bouquet of circles. For every subgroup H of F, there is a connected graph $\tilde{\mathcal{G}}$ such that $\tilde{\mathcal{G}}$ covers $\mathcal{G}$ and H is isomorphic to $\pi_1(\tilde{\mathcal{G}}, \tilde{v})$ for a vertex $\tilde{v} \in \tilde{\mathcal{G}}$.*

3.5. **The Nielsen Schreier Theorem and its Quantitative Form.** From the machinery we have developed, Thm. 3.1 drops out immediately!

Theorem 3.27 (Nielsen-Schreier)**.** *Every subgroup of a free group is free.*

Proof. Let H be a subgroup. By Lem. 3.26, H is isomorphic to the fundamental group of a connected graph. By Lem. 3.14, it follows that H is free. □

We can actually say even more.

Theorem 3.28 (Quantitative form of Nielsen-Schreier)**.** *Let F be a free group and H a subgroup of F with finite index. If r_H is the rank of H, r_F the rank of F, and i the index of H in F, then*

$$i = \frac{r_H - 1}{r_F - 1}$$

Proof. There are r_F edges in the realization of F as a bouquet of circles, $\mathcal{G}$. Hence, by the construction in the previous subsection, there are i vertices and ir_F edges in the covering graph $\tilde{\mathcal{G}}$. A spanning tree for a graph with i vertices takes up $i-1$ edges, so there are $ir_F - i + 1$ edges not in a spanning tree of $\tilde{\mathcal{G}}$. By the proof of Lem. 3.14, this is the number of free generators of H.

□

Remark 3.29. Note that Thm 3.28 holds only when it is known that H is of finite index. For example, let $F = \langle a, b\rangle$, the free group on two generators, and consider $H = \langle a\rangle$, the infinite cyclic subgroup. Then applying Thm 3.28 tells us H has index 0 in F, which is absurd. Hence, this gives a quick proof of the (to be fair, not very difficult) fact that H is of infinite index in F.

4. Quasi-isometry and the Švarc-Milnor Lemma

We now proceed from the more classical, topological methods of the previous section into the realm of modern geometric group theory. In this section, we will develop the theory of quasi-isometric metric spaces and prove the Švarc-Milnor Lemma, otherwise known as the Fundamental Lemma of Geometric Group Theory.

Let us draw the Cayley graph for $\mathbb{Z}$ over two different generating sets. If $S = \{1\}$:

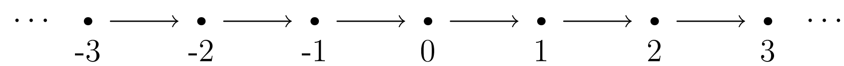

On the other hand, if $S = \{2, 3\}$:

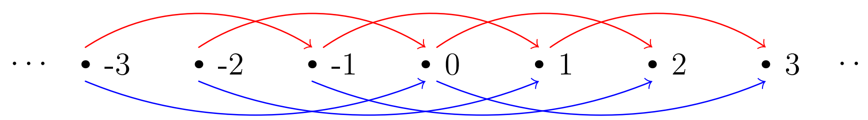

By changing the generating set, we drastically alter the *local* structure of the Cayley graph. However, the key observation here is that if we "zoom out" of each drawing far enough, the two Cayley graphs will look the same. Modification of the generating set does not affect the *coarse* structure.

This geometric concept is captured in the notion of *quasi-isometry* between metric spaces. It turns out that we can equip each Cayley graph with a metric. The two resultant metric spaces are quasi-isometric, hence the zooming out property holds. Let us develop these ideas formally. We follow the expositions in [3] and [4].

4.1. Definitions and the Hopf-Rinow Theorem.

Let (M, d) be a metric space.

Definition 4.1. For $r \geq 0$, the closed *r-neighborhood* of point $x \in M$ is the set

$$N(x, r) = \{y \in M : d(x, y) \leq r\}$$

We can also define the neighborhood of a subset $Q \subset M$ as

$$N(Q, r) = \bigcup_{x \in Q} N(x, r)$$

Remark 4.2. We will also occasionally reference the open r-neighborhood of x,

$$U(x, r) = \{y \in M : d(x, y) < r\}$$

Definition 4.3. If $Q \subset M$, we say Q is *r-dense* if $N(Q, r) = M$. Furthermore, Q is *cobounded* if it is r-dense for some $r \geq 0$. We define the *diameter* of Q as

$$\mathrm{diam}(Q) = \sup\{d(x, y) : x, y \in Q\}$$

and say Q is *bounded* if $\mathrm{diam}(Q) < \infty$.

Definition 4.4. Let $I \subset \mathbb{R}$ be an interval. A curve $\gamma : I \to M$ is a *geodesic* if $d(\gamma(s), \gamma(t)) = |s - t|$ for all $s, t \in I$. The metric space M is called a *geodesic space* if any two points in M are connected by a geodesic.

The following theorem will eventually be used (minimally) in our formulation of the Švarc-Milnor Lemma. But since it is an important and interesting theorem in its own right, we include the proof.

Theorem 4.5 (Hopf-Rinow)**.** *Let (M, d) be a complete, locally compact, geodesic space. Then $N(x, r)$ is compact for all $x \in M$, $r \geq 0$.*

Proof. We follow the proof given in [5]. Fix $x \in M$. Let $A = \{r \in [0, \infty) : N(x, r) \text{ is compact}\}$. Note that $A \neq \{0\}$ since M is locally compact. Also, note that if $r \in A$, then $[0, r] \subset A$.

Suppose $R \in A$. For each $x_\alpha \in N(x, R)$, there is a compact neighborhood $N(x_\alpha, r_\alpha)$. Then $\{U(x_\alpha, r_\alpha)\}$ is an open cover of $N(x, R)$, so it has a finite subcover $\{U(x_i, r_i)\}$. The union $\bigcup_i N(x_i, r_i)$ is compact, and contains $N(x, R + \epsilon)$ for some $\epsilon > 0$. Hence $R + \epsilon \in A$. It follows A is open in $[0, \infty)$.

Suppose R is a limit point of A; then $[0, R) \subset A$. Let $\{x_j\}$ be a sequence in $N(x, R)$, and let $\{\epsilon_i\}$ be a decreasing sequence of real numbers converging to 0. Assume, without loss of generality, that $\epsilon_i < R$ for all i. Since M is a geodesic space, we can find a geodesic between x and x_j for each j. For each i, along this geodesic, we can find a point $y_j^{(i)}$ such that $y_j^{(i)} \in N(x, R - \epsilon_i/2)$ and $d(x_j, y_j^{(i)}) \leq \epsilon_i$. Then, for each i, the sequence $\{y_j^{(i)}\}$ is contained in $N(x, R - \epsilon_i/2)$, which is compact. We may therefore pick a convergent subsequence $\{y_{j(1,k)}^{(1)}\}$ of $\{y_j^{(1)}\}$. This convergent subsequence has the corresponding subsequence $\{y_{j(1,k)}^{(2)}\}$ in $\{y_j^{(2)}\}$. Choose a convergent subsequence $\{y_{j(2,k)}^{(2)}\}$ from $\{y_{j(1,k)}^{(2)}\}$. Continuing in this fashion, we choose a convergent subsequence $\{y_{j(i,k)}^{(i)}\}$ from $\{y_{j(i-1,k)}^{(i)}\}$ for every i.

By construction, the sequence $\{y_{j(k,k)}^{(i)}\}_{k \in \mathbb{N}}$ converges for every i. Consider the counterpart sequence $\{x_{j(k,k)}\}_{k \in \mathbb{N}}$. If we pick $\epsilon > 0$, there is some i such that $\epsilon_i < \epsilon$. Since $\{y_{j(k,k)}^{(i)}\}_{k \in \mathbb{N}}$ converges, it is Cauchy, hence there is N such that for m, $n \geq N$, we have $d(y_{j(m,m)}^{(i)}, y_{j(n,n)}^{(i)}) < \epsilon$. Then for m, $n \geq N$, we see

$$d(x_{j(m,m)}, x_{j(n,n)}) \leq d(x_{j(m,m)}, y_{j(m,m)}^{(i)}) + d(y_{j(m,m)}^{(i)}, y_{j(n,n)}^{(i)}) + d(y_{j(n,n)}^{(i)}, x_{j(n,n)}) < 3\epsilon$$

Hence $\{x_{j(k,k)}\}$ is Cauchy. By completeness, this sequence has a limit. Hence, $\{x_j\}$ has a convergent subsequence.

Our sequence $\{x_j\}$ was arbitrary, so it follows $N(x, R)$ is compact. Thus $R \in A$, and A is closed. The only closed and open subset of $[0, \infty)$ is $[0, \infty)$.

□

4.2. **Isometries and Quasi-isometries.** The notion of isometry is fundamental to metric space theory. For completeness (pun intended), we recall

Definition 4.6. Let (X, d_X) and (Y, d_Y) be metric spaces. Then $f : X \to Y$ is an *isometric embedding* if $d_X(f(x), f(y)) = d_Y(x, y)$ for all x, $y \in X$. If, in addition, f is surjective, we call f an *isometry*. If there is an isometry between X and Y, then the spaces are *isometric*.

Isometries are distance-preserving bijections between metric spaces. Geometrically, this means the metric spaces look exactly the same up to rotations and translations. This is a very rigid restriction to have. In particular, isometries are continuous, meaning that they preserve local details.

On the other hand, we want a condition on mappings which preserve the large-scale geometry but do not necessarily reflect any local information. Thus, our condition should allow for some error between distances in the spaces, but need not imply continuity. Similarly, our condition should ensure the image of our map is "spread out," but it need not imply bijectivity. Hence

Definition 4.7. We say $f : X \to Y$ is a *quasi-isometric embedding* if there exist constants $A \geq 1$, $B \geq 0$ such that for all $x, y \in X$,

$$\frac{1}{A} d_X(x, y) - B \leq d_Y(f(x), f(y)) \leq A d_X(x, y) + B$$

If, in addition, the image of f is cobounded, we call f a *quasi-isometry.* If there is a quasi-isometry from X to Y, then X is *quasi-isometric* to Y.

Example 4.8. The map $x \mapsto \lfloor x \rfloor$ is a quasi-isometry. Hence $\mathbb{R}$ is quasi-isometric to $\mathbb{Z}$.

We give some properties of quasi-isometries. While quasi-isometries are not necessarily invertible, they are "quasi-invertible," in the same way that the coboundedness condition ensures "quasi-surjectivity."

Definition 4.9. Let f, g be maps from $X \to Y$. We say f and g are within *finite distance* if there exists a constant $C \geq 0$ such that $d_Y(f(x), g(x)) \leq C$ for all $x \in X$.

Proposition 4.10. *The following propositions are true.*

(1) If $f : X \to Y$ and $g : Y \to Z$ are quasi-isometries, then $g \circ f : X \to Z$ is also a quasi-isometry.
(2) If $f : X \to Y$ is a quasi-isometry, there is a quasi-isometry $g : Y \to X$ such that $g \circ f$ is within finite distance of the identity on X and $f \circ g$ is within finite distance of the identity on Y.

Proof. (1) This is clear from Def. 4.7.

(2) For $y \in Y$, let $S_y(r)$ denote the set of $x \in X$ with $d_Y(f(x), y) \leq r$. Since the image of f is cobounded, there exists $R \geq 0$ such that $S_y(R)$ is nonempty for all y. Set $S_y := S_y(R)$.

By the axiom of choice, let g be a choice function on the collection of sets S_y. That is, for every y, we have $g(y) \in S_y$.

For any $y \in Y$, we see $d_Y(f \circ g(y), y) \leq R$. For any $x \in X$, we see

$$\frac{1}{A} d_X(g \circ f(x), x) - B \leq d_Y(f \circ g \circ f(x), f(x)) \leq R$$

hence $d_X(g \circ f(x), x)$ is bounded. Therefore the image of g is cobounded.

It remains to show g is a quasi-isometric embedding. This follows easily from the inequalities in Def. 4.7 by applying triangle inequality. $\square$

Corollary 4.11. *We write $X \sim Y$ if X is quasi-isometric to Y. Then $\sim$ is an equivalence relation.*

Proof. The identity is a quasi-isometry, so $X \sim X$. Part (1) of Prop. 4.10 shows if $X \sim Y$ and $Y \sim Z$, then $X \sim Z$. Part (2) of Prop 4.10 shows if $X \sim Y$ then $Y \sim X$. $\square$

Having laid out the relevant definitions, we will now discuss groups in this setting.

4.3. **Groups as Metric Spaces.** At the beginning of this section, we noted that the Cayley graphs for $\mathbb{Z}$ over the generating sets $\{1\}$ and $\{2, 3\}$ are quasi-isometric. This is, unfortunately, a less than meaningful statement if we do not assign the Cayley graphs a metric space structure. We therefore define

Definition 4.12. Let $C(G,S)$ be the Cayley graph for G over the generating set S. The *word metric* on $C(G,S)$ is defined as follows. For $g,\ h \in G$, let $d_S(g,h)$ equal the length of the shortest path in $C(G,S)$ connecting g to h. In other words, $d_S(g,h)$ is the length of the shortest word over $S \cup S^{-1}$ representing the element $g^{-1}h$.

From this definition, the geometry of the resultant metric space depends *a priori* on the chosen generating set. This is true when considering only local details, as we saw in the Cayley graphs for $\mathbb{Z}$. But, on the coarse scale, we expect that different choices of generating sets should give rise to quasi-isometric spaces. It turns out that this is true when the generating sets are finite, as in our initial example.

Proposition 4.13. *Let G be a group and R, S finite generating sets for G. Then $C(G,R) \sim C(G,S)$ under the induced word metrics.*

Proof. We prove the inclusion map $\iota : (G,d_R) \to (G,d_S)$ is a quasi-isometry. Let

$$C = \max\{d_S(1,r) : r \in R\}$$

Then for all $g,\ h \in G$, we have $d_S(g,h) \leq Cd_R(g,h)$. This is because when calculating the distance in (G,d_S), each generating element in R is replaced by a word of length at most C. Similarly, we find constant $D \geq 0$ such that $d_R(g,h) \leq Dd_S(g,h)$. The image of ι is clearly cobounded.

□

We can see why this proposition, though simple, is important. It is this fact which allows us to say

Definition 4.14. Let G and H be finitely generated groups. We say G is *quasi-isometric* to H, written $G \sim H$ if there exist generating sets $R \subset G$ and $S \subset H$ such that $C(G,R) \sim C(H,S)$.

One of the problems at the heart of modern geometric group theory is the one suggested by this definition: can we classify the finitely generated groups up to quasi-isometry? The role of this problem as a beacon of geometric group theory is only further validated by the existence of nice theorems about the coarse geometry of finitely generated groups. For instance,

4.4. **Group Actions on Metric Spaces and the Švarc-Milnor Lemma.** The Švarc-Milnor Lemma says that given a sufficiently nice group action on a sufficiently nice metric space, we can immediately deduce that the group must be finitely generated and must be quasi-isometric to the metric space it acts on.

At first, this is a somewhat confusing statement. It seems deep, in that it provides a palpable connection between group geometry and group structure. But in what sense are group actions on metric spaces natural objects to consider? Which groups act on metric spaces at all? If there aren't many such groups, the theorem will not tell us much in the long run.

In fact, every group acts on some metric space. Recall Cayley's Theorem, which says every group G is isomorphic to the symmetric group on $|G|$ symbols. This is true because, for any chosen $g \in G$, the map $\phi_g : G \to G$ given by $\phi_g(h) = gh$. That is, left multiplication by g uniquely permutes the elements of G.

But the same map ϕ_g, when considered a self-map of the *metric space* G (or being more careful, the metric space $C(G,S)$ for some generating set S), is an isometry. Accordingly, every finitely generated group is isomorphic to a subgroup of the isometry group of a metric space! In the literature ([8]), this observation has been cheekily referred to as Cayley's Better Theorem.

This might seem like cheating, since we haven't necessarily provided a "real" metric space. We have only restated the permutation representation of G in a convenient way. But the restatement suggests that Švarc-Milnor acts as a partial converse to Cayley's Better Theorem. In the wild, we would not know a group G is finitely generated, only that it acts on a metric space. Of course, this metric space may well turn out to be the Cayley graph of G, but it also may not. In either case, our efforts would have yielded valuable information about the group.

To begin, let us review

Definition 4.15. Let X be a set, $S(X)$ the symmetric group of X, and G an arbitrary group. A *group action* of G is a homomorphism $G \to S(X)$ given by $g \mapsto \phi_g$. We say the group G *acts* on X. We write $\phi_g(x)$ as $g \cdot x$.

Definition 4.16. Suppose G acts on X. The *orbit* of x with respect to this action is the set $G \cdot x = \{g \cdot x : g \in G\}$. The *quotient* of X given by the action is the collection of orbits, denoted X/G.

The necessary conditions on our group action and metric space are

Definition 4.17. Let (X, d) be a complete, locally compact, geodesic space. The action of G on X is said to be *geometric* if

- G acts *isometrically* on X. That is, for every $g \in G$, the map ϕ_g is an isometry of X.
- G acts *cocompactly* on X. That is, the quotient space X/G is compact under the quotient topology.
- G acts *properly discontinuously* on X. That is, for every compact $K \subset X$, the set $\{g \in G : g \cdot K \cap K \neq \emptyset\}$ is finite.

The notion of proper discontinuity can be reformulated as follows.

Proposition 4.18. *Assume G acts isometrically on X. Then the following statements are equivalent.*

(1) G acts properly discontinuously on X.
(2) For every $x \in X$ and $r \geq 0$, the set $\{g \in G : d(x, g \cdot x) \leq r\}$ is finite.

Proof. (1) $\to$ (2): By Thm. 4.5, the closed neighborhood $N(x, r)$ is compact. Hence $S = \{g \in G : g \cdot N(x, r) \cap N(x, r) \neq \emptyset\}$ is finite. But $\{g \in G : d(x, g \cdot x) \leq r\}$ is a subset of S.

(2) $\to$ (1): Fix $r \geq 0$, and let $\{N(x_i, r)\}$ be a finite covering of K. Let $M = \max\{d(x_i, x_j)\}$. Suppose g is such that $g \cdot K \cap K \neq \emptyset$. Then there is $y \in K$ with $g \cdot y$ also in K. For some x_i, x_j, we have $d(y, x_i) \leq r$ and $d(g \cdot y, x_j) \leq r$. Since G acts isometrically, it follows $d(g \cdot y, g \cdot x_i) \leq r$, hence $d(g \cdot x_i, x_j) \leq 2r$ and $d(g \cdot x_i, x_i) \leq 2r + M$ by the triangle inequality. Thus

$$\{g \in G : g \cdot K \cap K \neq \emptyset\} \subset \bigcup_i \{g \in G : d(x_i, g \cdot x_i) \leq 2r + M\}$$

where the right side is finite.

□

We will also require

Proposition 4.19. *Suppose G acts properly discontinuously on X. Then the quotient X/G is a metric space.*

Proof. Define

$$d_Q(G\cdot x, G\cdot y) = \min\{d(p,q) : p \in G\cdot x, q \in G\cdot y\} = \min\{d(x, g\cdot y) : g \in G\}$$

By proper discontinuity, there are only finitely many $g \in G$ such that

$$g \cdot N(x, d(x,y)) \cap N(x, d(x,y)) \neq \emptyset$$

Hence there are only finitely many $g \in G$ such that $d(x, g\cdot y) \leq d(x,y)$, and the minimum is actually attained.

It is straightforward to check d_Q induces the quotient topology. □

Proposition 4.20. *Assume G acts isometrically and properly discontinuously on X. Then the following statements are equivalent.*

(1) G acts cocompactly on X.
(2) Every orbit of G is cobounded in X.

Proof. (1) → (2): Choose some orbit $G\cdot x \in X/G$. There is r such that $X/G \subset N(G\cdot x, r)$. For any $y \in X$, it follows $d_Q(G\cdot y, G\cdot x) \leq r$, hence $d(y, g\cdot x) \leq r$ for some $g \in G$. Thus $G\cdot x$ is cobounded.

(2) → (1): Again, choose $G\cdot x \in X/G$. Assume $G\cdot x$ is r-dense in X. Then, for any $y \in X$, there is $g \in G$ such that $d(x, g\cdot y) \leq r$. By Thm. 4.5, we know $N(x,r)$ is compact. The projection map $q : X \to X/G$ given by $x \mapsto G\cdot x$ is continuous. Hence $q(N(x,r)) = X/G$ is compact. □

Lemma 4.21 (Švarc-Milnor)**.** *Suppose G acts geometrically on X. Then G is finitely generated and $G \sim X$.*

Proof. We first show G is finitely generated. Fix $x \in X$. By Prop. 4.20, we know $G\cdot x$ is r-dense in X, for some $r \geq 0$. Let $k = 2r+1$.

Let us construct a graph $\mathcal{G}$ as follows. The vertex set of $\mathcal{G}$ is simply G. We draw an edge between g, $h \in G$ if $d(g\cdot x, h\cdot x) \leq k$.

For g, $h \in G$, let $L = d(g\cdot x, h\cdot x)$ and let $\gamma : [0,L] \to X$ be a geodesic connecting $g\cdot x$ and $h\cdot x$. Also, let $n = \lfloor L \rfloor + 1$. Note that $L/n < 1$.

We may choose points $g\cdot x = x_0, x_1, \cdots, x_n = h\cdot x$ along γ such that $d(x_i, x_{i+1}) = L/n$ for all $0 \leq i \leq n-2$. Then for each i, we can find g_i such that $d(x_i, g_i\cdot x) \leq r$ (set $g_0 = g$ and $g_n = h$). Then

$$d(g_i\cdot x, g_{i+1}\cdot x) \leq d(x_i, g_i\cdot x) + d(x_i, x_{i+1}) + d(x_{i+1}, g_{i+1}\cdot x) \leq 2r+1 = k$$

Now, it is clear that $g \to g_1 \to \cdots \to h$ is a path from g to h. Hence $\mathcal{G}$ is connected.

Let $S = \{g \in G : d(x, g\cdot x) \leq k\}$. By Prop. 4.18, we know S is finite and symmetric (if $s \in S$, then $s^{-1} \in S$). Furthermore, it is clear that there is an edge between g and h if and only if $g^{-1}h \in S$. Since $\mathcal{G}$ is connected, there is a path from 1 to g for every $g \in G$. It follows that S generates G and $\mathcal{G} = C(G,S)$.

Next, we prove $G \sim X$. This is the same as showing $C(G,S) \sim X$. Let $f : G \to X$ be defined by $g \mapsto g\cdot x$. By assumption, $G\cdot x$ is cobounded, hence the image of f is cobounded.

From the above argument, we see for any g, $h \in G$, that

$$d_S(g,h) \leq n = \lfloor L \rfloor + 1 \leq d(g\cdot x, h\cdot x) + 1 = d(f(g), f(h)) + 1$$

This is our lower bound. For the upper bound, let $g \to g_1 \to \cdots \to g_\ell = h$ be a shortest path in $C(G,S)$. Then

$$d(f(g), f(h)) = d(g\cdot x, h\cdot x) \le d(g\cdot x, g_1\cdot x) + \cdots + d(g_{\ell-1}\cdot x, h\cdot x) \le k\ell = kd_S(g,h)$$

Therefore, f is a quasi-isometry.

□

Remark 4.22. We required X to be a geodesic space. That is, for any two points in X, there is an isometry $g: I \to X$ connecting the two points (where I is a closed interval in $\mathbb{R}$). In fact, we could have weakened this condition to require X merely be a "quasi-geodesic" space, in which g is merely a quasi-isometry. The proof of Švarc-Milnor can easily be adapted.

The lemma lends itself to immediate use. For instance, we can prove as a corollary that

Proposition 4.23. *Let H be a subgroup of G with finite index. If G is finitely generated, then H is finitely generated and $G \sim H$.*

Proof. Let S be a finite generating set of G. We let H act on (G, d_S) by right translation. That is, with every $h \in H$ we associate $\phi_h \in S(G)$ given by $\phi_h(g) = gh$. It is easy to check this action is geometric.

To show G has the necessary metric space structure, we alter the construction of $C(G,S)$ slightly. In addition to the vertices of $C(G,S)$, we also say every point inside an edge is in the metric space. We can extend the word metric naturally by assigning each edge a mass of 1 and proclaiming that a point x-way $(0 \le x \le 1)$ between vertices v and w has distance x from v and $1-x$ from w. Now it is straightforward to show G is geodesic and complete.

Apply Lem. 4.21 to finish. □

5. Quasi-isometry Invariants

To show that two groups (or metric spaces, for that matter) are quasi-isometric, it suffices to exhibit an explicit quasi-isometry between them. Should this fail, we might even try exhibiting a geometric action of one group on the other and then applying the Švarc-Milnor Lemma. But in either case, the point is that we have techniques to demonstrate quasi-isometric equivalence. Perhaps they will not always work, but they will sometimes.

On the other hand, we have developed no approach to prove two groups are *not* quasi-isometric. This is certainly an important thing to be able to do.

As in topology, one approach is to find invariants. We wish to determine group-theoretic properties which are preserved by quasi-isometry. It is the goal of this section to discuss a few of these properties and relevant problems.

Assume groups are finitely generated unless otherwise stated.

5.1. **A Simple Quasi-isometry Invariant.**

Definition 5.1. A property "P" is called *geometric*, or *quasi-isometry invariant*, if whenever a group G has "P," every group H with $H \sim G$ also has "P."

The simplest example of a geometric property is finiteness.

Proposition 5.2. *Suppose $G \sim H$ and G is finite. Then H is finite.*

Proof. By Prop. 4.23, we see that $G \sim (1)$, where (1) is the trivial group. Then $H \sim (1)$, and it follows there exists $C \geq 0$ such that $d_S(g,h) \leq C$ for all g, $h \in H$. In particular, it takes at most C generators to write any element of H. But there are only finitely many words that can be written on C generators. □

Every finite group is quasi-isometric to the trivial group. Hence, from the coarse geometric viewpoint, finite groups are trivial. But this was to be expected. After all, if we draw the graph of a finite group and zoom out far enough, we are eventually left looking at merely a dot!

5.2. **Growth Rates of Finitely-Generated Groups.**

Definition 5.3. Let G be a group with generating set S. The *growth rate* of G over S is the function $\beta_{G,S} : \mathbb{N} \to \mathbb{N}$ given by

$$\beta_{G,S}(r) = |\{g \in G : d_S(1,g) \leq r\}| = |N(1,r)|$$

the cardinality of the r-neighborhood centered at the identity element.

The first thing to note is that growth rates are at most exponential, since we are only considering finitely generated groups. If F is freely generated by k generators $(k > 1)$, then the growth rate goes as $\sim (2k)^r$. But if there are relations between generators, then the growth rate may well be sub-exponential. For example, the abelian group $\mathbb{Z}$ generated by $\{1\}$ has growth rate $\beta_{\mathbb{Z},\{1\}}(r) = 2r+1$, which is polynomial.

An interesting, natural question is whether or not there exist groups of *intermediate* growth: faster than polynomial growth but slower than exponential growth. This problem was settled positively by Grigorchuk in the 1980s (see [9]).

In order to work with growth rates, we will need a formal notion of asymptotic equivalence.

Definition 5.4. For $f, g : \mathbb{R}^+ \to \mathbb{R}^+$, we say g *quasi-dominates* f if there exist constants $A, B \geq 0$ such that $f(r) \leq Ag(Br)$ for all sufficiently large r. We write this as $f \prec g$. If $f \prec g$ and $g \prec f$, then we say f is *quasi-equivalent* to g, written $f \sim g$.

Proposition 5.5. *Suppose G and H are generated by R, S respectively. If $G \sim H$, then $\beta_{G,R} \sim \beta_{H,S}$.*

Proof. Write 1 for the identity in G and $1'$ for the identity in H. Let $f : G \to H$ be a quasi-isometry with

$$\frac{1}{C} d_R(g,h) - B \leq d_S(f(g), f(h)) \leq C d_R(g,h) + B$$

Then for any $g \in G$ with $d_R(1,g) \leq r$, we have $d_S(f(1), f(g)) \leq Cr + B$. Set $d_S(f(1), 1') = D$. Hence, for every $g \in G$ within r of 1, we get an element $f(g) \in H$ within $Cr + B + D$ of $1'$.

This correspondence may not be one-to-one, but it is "quasi-injective" : if $f(g) = f(h)$, then $d_R(g,h) \leq BC$. For each $f(g) \in f(N_G(1,r))$, choose a canonical element $g \in G$ which maps to it. Then, there are at most $|N_G(g, BC)| = |N_G(1, BC)|$ elements total which map to $f(g)$.

Therefore, for sufficiently large r,

$$\beta_{H,S}(2Cr) \geq \beta_{H,S}(Cr + B + D) = |N_H(1', Cr + B + D)|$$

$$\geq |f(N_G(1,r))| \geq \frac{1}{|N_G(1,BC)|}|N_G(1,r)| = \frac{1}{|N_G(1,BC)|}\beta_{G,R}(r)$$

implying $\beta_{G,R} \prec \beta_{H,S}$. The other direction is similar. □

Corollary 5.6. *Suppose G is generated by both R and S. Then $\beta_{G,R} \sim \beta_{G,S}$.*

Proof. Apply Prop. 4.13. □

We can, for instance, use growth rate considerations to prove $\mathbb{Z}^m$ and $\mathbb{Z}^n$ are quasi-isometric only when $m = n$. Other proofs of this fact (see [3], [10]) involve some unseemly epsilon-pushing and relatively advanced topological considerations. This should hopefully suggest the intrinsic value and importance of the growth rate property.

Proposition 5.7. *$\mathbb{Z}^m \sim \mathbb{Z}^n$ if and only if $m = n$.*

Proof. Let us compute the growth rate of $G = \mathbb{Z}^n$. Set $e_i = (0, 0, \cdots, 1, \cdots, 0)$, where 1 is in the ith slot and there are n slots total. Then $S = \{e_1, \cdots, e_n\}$ generates G.

The cardinality $|N(1,r)|$ is the number of lattice points $(x_1, x_2, \cdots, x_n)$ inside the n-simplex $X(r)$ given by $|x_1| + |x_2| + \cdots + |x_n| \leq r$.

Let $\mathrm{vol}(S)$ denote the volume of the region $S \subset \mathbb{R}^3$. If we identify each lattice point $(x_1, x_2, \cdots, x_n)$ with the cell

$$\{(y_1, \cdots, y_n) : x_i \leq y_i \leq x_i + 1\}$$

we can bound

$$\mathrm{vol}(X(r-2)) \leq |N(1,r)| \leq \mathrm{vol}(X(r+2))$$

By scaling, $\mathrm{vol}(X(r)) = r^n \mathrm{vol}(X(1))$. Thus, $\beta_{G,S}$ is of degree n polynomial growth.

So if $m \neq n$, then the growth rates of $\mathbb{Z}^m$ and $\mathbb{Z}^n$ are not quasi-equivalent. By Prop. 5.5, this proves $\mathbb{Z}^m$ and $\mathbb{Z}^n$ are not quasi-isometric.

□

Corollary 5.8. *$\mathbb{R}^m \sim \mathbb{R}^n$ if and only if $m = n$.*

Proof. We have $\mathbb{R}^n \sim \mathbb{Z}^n$ by the map $f(x_1, \cdots, x_n) = (\lfloor x_1 \rfloor, \cdots, \lfloor x_n \rfloor)$. □

5.3. **Ends of Groups.** Are $[0, \infty)$ and $\mathbb{R}$ quasi-isometric? We know that quasi-isometries preserve the coarse geometry of a metric space. Heuristically, one feature we might deem characteristic of coarse geometry is the number of "directions" in which the metric space extends.

If we draw the standard pictures of these spaces, we see that $\mathbb{R}$ extends to infinity in exactly two directions, while $[0, \infty)$ extends in only one direction. We expect that quasi-isometries preserve this number of directions, hence that $[0, \infty)$ and $\mathbb{R}$ are not quasi-isometry equivalent.

Proposition 5.9. *The metric spaces $[0, \infty)$ and $\mathbb{R}$ are not quasi-isometric.*

Proof. Let $f : \mathbb{R} \to [0, \infty)$ be a quasi-isometry.

We begin by making two observations. First, note that as $n > 0$ goes to infinity, both $f(n)$ and $f(-n)$ approach infinity. Second, see that there is an absolute constant C such that $|f(n+1) - f(n)| \leq C$ for all $n \in \mathbb{R}$.

Now, fix $K \in [0, \infty)$. Let x be the largest positive integer such that $f(x) < K$ (by our first observation, this is possible). Since this implies $f(x+1) \geq K$, by our second observation above, $f(x)$ is within C of K. Similarly, let y be the smallest negative integer such that $f(y) < K$. Again $f(y)$ is within C of K. It follows that $|f(x) - f(y)| \leq 2C$.

By the quasi-isometry inequalities, there is some absolute constant D such that $x \leq x - y = |x - y| \leq D$. Recall that our choice of K was arbitrary. By our first observation, we may choose K large enough that this inequality is contradictory. □

In this proof, the key impossibility is that $f(n) \to \infty$ as $n \to \infty$ *and* as $n \to -\infty$. The points at both "ends" of $\mathbb{R}$ must be taken to ∞ in $[0, \infty)$. That is, distant points in $\mathbb{R}$ are mapped to close points in $[0, \infty)$, which we know cannot happen. On some level, this is an elaboration of the heuristic argument given above: quasi-isometry fails because we would need to fit two ends into one.

So there is certainly something to the idea of directions in which a space extends. We call the number of such directions the *number of ends* of the space. If we can develop this idea more carefully, we should be able to demonstrate that it is a quasi-isometry invariant.

We will restrict our presentation to work with graphs, following [8]. One can apply similar principles to general metric spaces (as in Prop. 5.9), but we will not do so. The reason for this is twofold. Graphs should suffice for geometric group theory, since the only metric spaces we need are Cayley graphs. Furthermore, to do the work over general metric spaces, we would need a bit more analytic/topological knowledge than we assume your acquaintance with. The interested reader may see [5] for details.

Definition 5.10. Let $\mathcal{G}$ be a graph. We define a natural path metric on $\mathcal{G}$ as follows. For vertices $v, w \in \mathcal{G}$, set $d(v, w)$ to be the length of the shortest path connecting v and w. If there is no path connecting v and w, set $d(v, w) = \infty$.

We say v and w are connected if $d(v, w) < \infty$; this is clearly an equivalence relation. The equivalence classes thus formed each induce a subgraph of $\mathcal{G}$. These subgraphs are called the *connected components* of $\mathcal{G}$. Denote the number of unbounded connected components of $\mathcal{G}$ as $\|\mathcal{G}\|$.

Remark 5.11. The distance function we have defined does not really make $\mathcal{G}$ a metric space, since metrics may only take on finite values. But we will have no occasion to treat non-connected graphs as metric spaces anyway, so we ask that you accept this minor abuse of notation.

How can we count the number of directions in which a graph $\mathcal{G}$ extends? Look at the example of the graph given by the lattice points on the coordinate axes in $\mathbb{R}^2$. This evidently has four "ends". We can isolate these ends by choosing a ball around the origin and removing it from the graph; we would be left with four unbounded connected components. By extending this logic, we hope to define

Definition 5.12 (Preliminary)**.** Let $\mathcal{G}$ be a locally finite, connected graph with natural path metric. Pick $v \in \mathcal{G}$ an arbitrary vertex. We define the *number of ends* of $\mathcal{G}$ as

$$e(\mathcal{G}) = \lim_{n \to \infty} \|\mathcal{G} \setminus N(v, n)\|$$

But we have not shown that this limit always exists, nor have we shown that it is independent of our choice of v. Fortunately, these are both non-issues.

From here on, assume graphs are indeed locally finite, connected, and given a path metric (as a Cayley graph certainly is).

Proposition 5.13. *Let $\mathcal{G}$ be a graph. Pick a vertex $v \in \mathcal{G}$. If $m < n$, then*

$$\|\mathcal{G} \setminus N(v,m)\| \leq \|\mathcal{G} \setminus N(v,n)\|$$

Proof. Take C an unbounded connected component of $\mathcal{G} \setminus N(v,m)$. If we remove $N(v,n)$ as well, then C will either remain an unbounded connected component, or it will be divided into a few unbounded connected components. $\square$

Therefore, for any graph $\mathcal{G}$, the sequence $\{\|\mathcal{G} \setminus N(v,n)\|\}$ is non-decreasing and consists of positive integers. It follows that the limit in Def. 5.12 is either a positive integer, or infinity. We would like to retain the possibility of having an infinite number of ends. Consider, for instance, the Cayley graph of any free group with finite rank.

Definition 5.14. Let $\mathcal{G}$ be a graph. We define the *strong number of ends* of $\mathcal{G}$ as

$$e_S(\mathcal{G}) = \sup\{\|\mathcal{G} \setminus \mathcal{C}\| \mid \mathcal{C} \text{ is a finite subgraph of } \mathcal{G}\}$$

Proposition 5.15. *Let $\mathcal{G}$ be a graph. For any choice of base vertex, $e(\mathcal{G}) = e_S(\mathcal{G})$.*

Proof. Let v be the base vertex. Since $\mathcal{G}$ is locally finite, $N(v,n)$ is a finite subgraph of $\mathcal{G}$. Thus, $\|\mathcal{G} \setminus N(v,n)\| \leq e_S(\mathcal{G})$. Taking $n \to \infty$ gives $e(\mathcal{G}) \leq e_S(\mathcal{G})$.

Now, let $\mathcal{C}$ be a finite subgraph of $\mathcal{G}$. For large enough n, this finite subgraph will be contained in $N(v,n)$. By the argument given in Prop. 5.13, $\|\mathcal{G} \setminus \mathcal{C}\| \leq \|\mathcal{G} \setminus N(v,n)\|$. Taking $n \to \infty$ gives $\|\mathcal{G} \setminus \mathcal{C}\| \leq e(\mathcal{G})$. Since $\mathcal{C}$ was arbitrary, it follows $e_S(\mathcal{G}) \leq e(\mathcal{G})$.

$\square$

Since $e_S(\mathcal{G})$ is independent of our vertex choice, we may conclude that Def. 5.12 is valid. Our next step is to extend our concept of ends of graphs to groups. Ideally, we want to be able to say

Definition 5.16 (Still Preliminary)**.** Let G be a group. Pick S an arbitrary generating set. We define the *number of ends* of G as $e(G) = e(C(G,S))$.

To validate this, we will need to prove $e(G)$ is independent of the choice of S. The following lemma resolves this.

Lemma 5.17. *Suppose G and H are groups with $G \sim H$. Let R, S be arbitrary generating sets of G, H respectively. Then $e(C(G,R)) = e(C(H,S))$.*

Proof. We will use the notation of Prop. 5.5. For any $g_1, g_2 \in G$ with $d_R(g_1,g_2) = 1$, note that $d_S(f(g_1), f(g_2)) \leq C + B$. Let $D = d(f(1), 1')$ and choose some $E > C(2B + C + D)$. Without loss of generality, assume $C \geq 1$. Also, fix $n \in \mathbb{R}$.

Let $g, h \in C(G,R) \setminus N_R(En+E)$ such that g and h are in the same connected component. Then there is a path $g_0 = g \to g_1 \to g_2 \to \cdots \to g_k = h$ from g to h with $d_R(g_i, 1) > En + E$ for all $0 \leq i \leq k$. For each i, it follows that

$$d_S(f(g_i), 1') \geq d_S(f(g_i), f(1)) - D \geq \frac{1}{C} d_R(g_i, 1) - B - D$$
$$> \frac{1}{C}(En + E) - B - C > n + B + C$$

We construct a path from $f(g)$ to $f(h)$ in $C(H,S)$ as follows. Each edge $g_i g_{i+1}$ has length 1 in $C(G,R)$. By our first observation, there is, for each $0 \leq i \leq k-1$, a path between $f(g_i)$ and $f(g_{i+1})$ of length at most $B + C$. Connecting all of these gives us a path from $f(g)$ to $f(h)$ in $C(H,S)$.

Furthermore: for a given i, since $d_S(f(g_i), 1') > n + B + C$, the $C(H,S)$ path from $f(g)$ to $f(h)$ cannot intersect $N_S(1', n)$. Therefore, for any path from g to h in $C(G,R) \setminus N_R(1, En + E)$, we may find a corresponding path from $f(g)$ to $f(h)$ in $C(H,S) \setminus N_S(1', n)$.

By this path correspondence, we get a map that takes each unbounded connected component in $C(G,R) \setminus N_R(1, En+E)$ to an unbounded connected *subgraph* of some unbounded connected component in $C(H,S) \setminus N_R(1', n)$. It follows that

$$\|C(G,R) \setminus N_R(1, Dn + D)\| \geq \|C(H,S) \setminus N_S(1', n)\|$$

Sending $n \to \infty$ gives $e(C(G,R)) \geq e(C(H,S))$. Repeating the argument with G and H swapped (using the quasi-inverse of f) completes the proof.

□

Since distinct Cayley graphs of groups are quasi-isometric, Lem. 5.17 demonstrates the well-definedness of Def. 5.16. Consequently, we can concisely state Lem. 5.17 in the following form:

Corollary 5.18. *Suppose G and H are groups with $G \sim H$. Then $e(G) = e(H)$.*

The ends of a group are particularly interesting due to the following classification theorem of Freudenthal and Hopf, which offers us a firm characterization of the pictures that may arise when one draws the Cayley graph of a finitely generated group.

Theorem 5.19 (Freudenthal-Hopf)**.** *Every finitely generated group G has $e(G) \in \{0, 1, 2, \infty\}$.*

Proof. Let G be a group (say, generated by S) with $2 < e(G) < \infty$. Then $e(G) = k \geq 3$. Choose $r \in \mathbb{R}$ such that $\|C(G,S) \setminus N(1,r)\| = k$.

You can easily check that a group is finite if and only if it has zero ends, so G is infinite. Thus, there is some element $g \in G$ with $d_S(1, g) > 2r$. Let $N(1,r) \cdot g = \{hg \,|\, h \in N(1,r)\}$. Note that $N(1,r) \cdot g \cap N(1,r) = \emptyset$, since any element in $ag \in N(1,r) \cdot g \cap N(1,r)$ would have

$$d_S(g, 1) \leq d_S(ag, a) \leq d_S(ag, 1) + d_S(a, 1) \leq 2r$$

It follows that $N(1,r) \cdot g$ is contained in some unbounded connected component of $C(G,S) \setminus N(1,r)$. Call this component E. Note that multiplication by g is an automorphism of $C(G,S)$, and effectively translates the connected subgraph $N(1,r)$. This divides the component E into at least k connected pieces, of each at least $k-1$ are unbounded.

Thus, if we set $\mathcal{C} = N(1,r) \cup N(1,r) \cdot g$, then $\mathcal{C}$ is a finite subgraph of $C(G,S)$ and $\|C(G,S) \setminus \mathcal{C}\| \geq (k-1) + (k-1) = 2k - 2$. Hence $e(G) = e_S(G) \geq 2k - 2 > k$, contradiction.

□

Acknowledgments. I would like to thank Oishee Banerjee for her invaluable guidance in learning the rudiments of algebra this summer. Many thanks also to Peter May and the other REU teachers for their clear and patient instruction. Lastly, I wish to recognize Nathaniel Sauder, who has served as a constant inspiration and role model.

References

[1] Joseph J. Rotman. An Introduction to the Theory of Groups, 4e. Springer-Verlag. 1995.
[2] John Stillwell. Classical Topology and Combinatorial Group Theory, 2e. Springer GTM, 1995.
[3] Brian H. Bowditch. A course on geometric group theory. Published online, maths.soton.ac.uk.
[4] Clara Löh. Geometric group theory, an introduction. Published online, mathematik.uni-regensberg.de.
[5] Martin Bridson and Andre Haefliger. Metric Spaces of Non-Positive Curvature. Springer-Verlag, 1999.
[6] Wolfgang Lück. Survey on geometric group theory. Münster J. of Math 1 (2008), 73-108.
[7] Cornelia Drutu and Michael Kapovich. Lectures on Geometric Group Theory. Published online, math.ucdavis.edu.
[8] John Meier. Groups, Graphs, and Trees: An Introduction to the Geometry of Infinite Groups. Cambridge University Press, 2008.
[9] Rostislav Grigorchuk and Igor Pak. Groups of Intermediate Growth: An Introduction for Beginners. Published online at the arxiv.
[10] Kevin Whyte. Quasi-Isometries. Published online, math.utah.edu.