An Introduction to Geometric Group Theory

Ginger Conner, Kismet Tillson Hunter, Alice Platt

Directed Reading Program
McGill University

June 26, 2025

1 Introduction

For our project, we read the book Office Hours with a Geometric Group Theorist [CM17], edited by Matt Clay and Dan Margalit. After reviewing groups and metric spaces, which is necessary background for this topic, we dove into our main object of study, Cayley graphs, along with some important theorems in geometric group theory and their proofs.

An important initial question in geometric group theory is how we can represent groups as geometric objects. One way to do this is called a Cayley graph. This paper will discuss the definition of a Cayley graph, how to construct one for a given group, and go through many examples.

Moreover, this paper discusses how groups act on graphs, and how these geometric group actions can provide algebraic information.

2 Preliminaries

Definition 1. *A* ***tree*** *is a connected graph with no cycles, i.e. there exists a unique path between any two vertices.*

Definition 2. *The* ***path metric*** *on a connected graph* Γ *with vertex set* X *is given by:* $d(v_1, v_2)$ *is the length of the shortest edge path between* v_1 *and* v_2*,* $v_1, v_2 \in X$*.*

Definition 3. *The action of a group,* G*, on a set* X *is called a* ***free action*** *if the only element that fixes any* $x \in X$ *is the identity. Moreover, a group action on a tree is free if the identity is the only group element that fixes a vertex of the tree, and if only the identity preserves edges of the tree.*

3 Groups

3.1 Torsion

Definition 4. *A* ***torsion element*** *of a group is an element that has finite order. A group is said to be* ***torsion free*** *if the identity is its only torsion element.*

3.2 The Free Group

Definition 5. *In an arbitrary set S, a* ***word*** *is a finite sequence of elements that may be empty which can be written as a product of elements in S. We say that a word is* ***reduced*** *if an element is never followed by its inverse.*

Definition 6. *A* ***free group*** *is a group with generating set S where the operation is concatenate and reduce. The free group consists of all reduced words that are products of elements of S. The cardinality of S is called the* ***rank*** *of the free group.*

For example, the free group of rank 2, F_2, has generating set $\{a, b\}$. The group consists of words containing a, b, a^{-1}, and b^{-1}, with no element ever followed by its inverse.

3.3 Group Presentations

Definition 7. *A* ***group presentation*** *is a pair (S, R), where S is a set and R is a set of words in S. We say that a group G has presentation $\langle S|R\rangle$ if $S \cup S^{-1}$ is a generating set for G and the words in R are the same as the identity element. The set S are the* ***generators*** *of G and the words in R are the* ***relators****.*

Group presentations are not unique and every group has at least one group presentation. To see this, simply take the generating set to be the entire group.

Examples

- For $\mathbb{Z}^2$, the generators are x and y, and $xy = yx$ is our relator
  - Typically we write the relator to be a single word which is equal to the identity: $xyx^{-1}y^{-1}$
  - $\mathbb{Z}^2 \cong \langle x, y|xyx^{-1}y^1\rangle$
- $\mathbb{Z}/n\mathbb{Z} \cong \langle a|a^n\rangle$

  This is how we write the group presentation for any cyclic group generated by a.
- $\mathbb{Z} \cong \langle x\rangle$
- $D_4 \cong \langle r, s|r^4, s^2, (rs)^2\rangle$
- $F_2 \cong \langle a, b|\varnothing\rangle$

  Any free group does not have relators.

4 Graphs

4.1 Automorphisms and Actions on Graphs

Definition 8. *A* ***graph isomorphism****,* $\Gamma \to \Gamma'$*, is a pair of bijections,* $V \to V'$ *and* $E \to E'$ *that preserves all of the graph structures, where* V *and* V' *are the vertex sets of* Γ *and* Γ' *and* E *and* E' *are the edge sets of* Γ *and* Γ'*. In other words, the endpoints of the image of an edge are images of the endpoints of that edge.*

Definition 9. *A* ***graph automorphism*** *is a graph isomorphism from a graph to itself. We call* $Aut(\Gamma)$ *the set of automorphisms of* Γ*.*

Another way to think about graph automorphisms is to picture them as the symmetries of the graph, as you would any other geometric object.

Definition 10. *The* ***group action of*** G ***on a graph*** Γ *is a homomorphism* $G \to Aut(\Gamma)$*.*

Group actions on graphs must preserve all graph structures.

4.2 Cayley Graphs

The first step to geometric group theory is figuring out how to represent any group as a geometric object. This can be done with a specific type of graph called a Cayley graph. The goal of constructing a Cayley graph is to find a graph that, when acted on by the group using left multiplication, returns all of the symmetries of the graph. This action will essentially give us back the group. This is described formally in Theorem 12.

Definition 11. *The* ***Cayley Graph*** *of a group* G *with generating set* S *is a directed, labeled graph* $\Gamma(G,S)$ *such that the vertices of the graph are the elements of* G*, and for all* $g \in G$*, every* $s \in S$ *creates an edge connecting* g *to* gs*.*

Example We will look at a section of the Cayley graph

$$G \cong \langle s_1, s_2 \mid s_2^2 \rangle$$

around a group element $g \in G$.

First we note that the Cayley graph can be drawn in two different ways. On the left side of Fig. 1, we draw a two-way edge instead of two separate directed edges, as on the right. Drawing a two-way edge is the standard way to draw the edge of a generator with order 2 because the graph appears cleaner and simpler.

From Fig. 1 we can note some characteristics of any Cayley graph:

- Any vertex g has $2 \cdot |S|$ edges coming out of it, half going out to gs_i and half coming in from gs_i^{-1} for every $i = 1, \ldots, n$ where n is the number of generators of G.
- Any relator creates a loop in the Cayley graph, as shown in the right image. Therefore we anticipate that the Cayley graph of a free group will be a tree.

Figure 1: Two ways to draw a small section of the Cayley graph of $G \cong \langle s_1, s_2 \,|\, s_2^2 \rangle$. The pink lines represent edges corresponding to the generator s_1 and the blue lines represent edges corresponding to the generator s_2.

- Locally, or around any point, the Cayley graph should look the same everywhere.

To get a better grasp on how Cayley graphs look and are constructed, we will look at some examples. We also found that attempting to construct a Cayley graph of another group that you know or like is another way to develop a better understanding.

Example First, we will look at the Cayley graph of $\mathbb{Z}^2$ which has group presentation:

$$\mathbb{Z}^2 \cong \langle x, y \,|\, xy = yx \rangle$$

This is an infinite group so the Cayley graph is also infinite. The relator tells us that traveling along the horizontal direction then the vertical direction will give the same vertex as traveling first in the vertical direction then the horizontal direction.

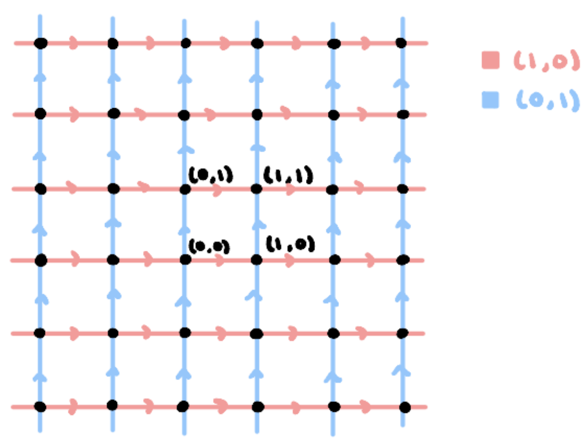

Figure 2: The Cayley graph of $\mathbb{Z}^2$ with generators $x = (1, 0)$ and $y = (0, 1)$.

We can draw the Cayley graph of a single group with different group presentations. Note that the choice of group presentation does not matter up to quasi-isometry of their Cayley graphs, which we will not cover.

Example We look at two ways to draw the Cayley graph of $\mathbb{Z}$ using two group presentations:

$$\mathbb{Z} \cong \langle x \rangle \qquad\qquad \mathbb{Z} \cong \langle a, b \,|\, a^3 = b^2 \rangle$$

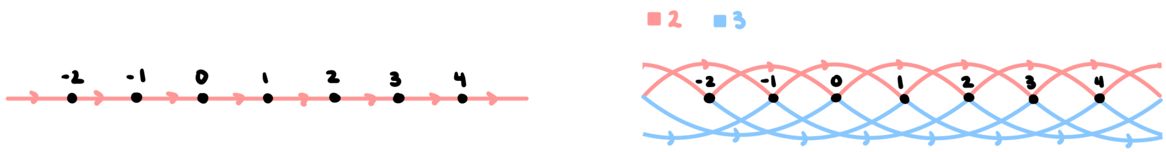

Figure 3: Two Cayley graphs of $\mathbb{Z}$ using the different group presentations listed above. On the right, $a = 2$ and $b = 3$.

Example

$$\mathbb{Z}/6\mathbb{Z} \cong \langle a \,|\, a^6 \rangle$$

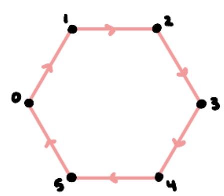

Figure 4: The Cayley graph of $\mathbb{Z}/6\mathbb{Z}$.

From this, we can see that any cyclic group of order n will have a Cayley graph that looks like a polygon with n sides.

Example $D_4 \cong \langle r, s \,|\, r^4, s^2, (rs)^2 \rangle$ is shown in Fig. 5.

Example Now we look at the process of constructing the Cayley graph of F_2, which can been seen in Fig. 6. The method we use is to start with the identity element. Then apply each generator and their inverses to the element. Each successive step is produced by applying each generator and their inverses to every new element. We are careful not to redraw any edge and vertex that is already drawn. So in this case, three new edges and vertices are added to each vertex that was drawn in the previous step.

$$F_2 \cong \langle a, b \rangle$$

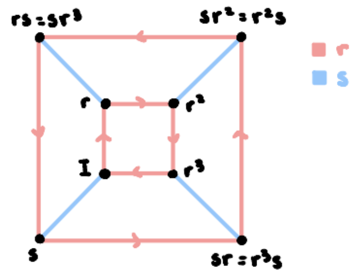

Figure 5: The Cayley graph of D_4.

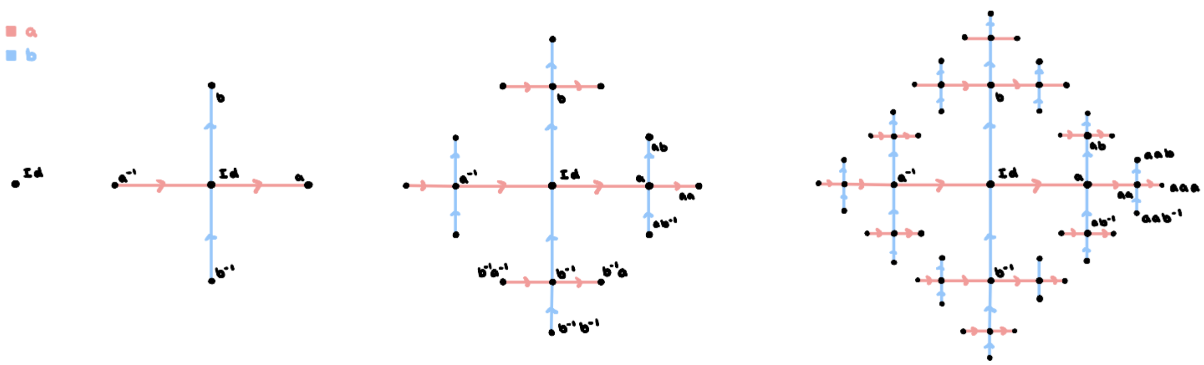

Figure 6: First four steps in constructing the Cayley graph of F_2 from left to right.

Now that we've seen several Cayley graphs, we can formalize the connection between a group and its Cayley graph with Theorem 12.

Theorem 12. *Let S be a generating set for the group G. Let Φ_g be an automorphism of the graph given on the vertices by $\Phi_g(v) = gv$, that is, the action of G by left multiplication. Then the map $G \to Aut(\Gamma(G,S))$, $g \mapsto \Phi_g$ is an isomorphism.*

5 Groups Acting on Trees

Theorem 13. *If a group G acts freely on a tree, then G is a free group.*

Proof. Before we dive into the proof, we need a few definitions.

Let T be a tree. The **barycentric subdivision of** T, denoted by T', is the graph obtained by placing new vertices in the center of every edge of T. A **tile** of T is a subtree of the barycentric subdivision of T.

A **tiling of** T is a collection of tiles that satisfy the following properties:

1. No tiles share an edge, so two tiles either intersect at a single vertex or they do not intersect.
2. The union of all tiles is the entire tree, T'.

   For this proof, we have the group G acting freely on T and we will impose a third condition on the tiling.

3. There is a single tile, T_0, such that the set of tiles is $\{gT_0 \mid g \in G\}$. In other words, every tile must be of the form gT_0 and for every $g \in G$, gT_0 is a tile.

We will call a tiling of T that satisfies all three conditions a **G-tiling of** T. The whole proof consists of three steps. First, we find a G-tiling of T. Next, we find a generating set for G using that tiling. Finally, we show that this generating set is a free-generating set, which the proves that G is a free group.

Step 1: Tiling the Tree

To find a G-Tiling of T, we choose any vertex $v \in T$ and consider the orbit of v under G, called the G-orbit of v. Since G acts freely on the tree, that is, no non-trivial element preserves any vertex or edge, there exists a bijection between the G-orbit of v and the elements of G.

Now for every $g \in G$, we take T_g to be the subgraph of T' with the following conditions:

- T_g has vertex set W_g such that for every $w \in W_g$,

  $$d(w, gv) \leq d(w, g'v) \quad \forall g' \in G.$$

  Note that we are using the path metric on T'.

- T_g has edge set E_g such that both vertices of $e \in E_g$ are in W_g

In other words, T_g is all of the vertices closest to the vertex gv and the edges connecting them.

To better understand this, we look at an example with the group F_2 acting on its Cayley graph, which is a group acting freely on a tree.

In Fig. 7, the left image is the Cayley graph of F_2. The center image is the barycentric subdivision of the Cayley graph, where the purple vertices are the new vertices. The right image shows the tiles of the Cayley graph, circled in orange. The center tile is T_0 or the identity tile. Each tile includes a vertex of T and the four closest vertices of T'. Adjacent tiles only intersect at a single vertex, which is an element of T' but not T.

The following three claims will prove that the union of $\{T_g\}_{g \in G}$ is indeed a G-tiling of T.

Claim 14. *Every T_g is a tile.*

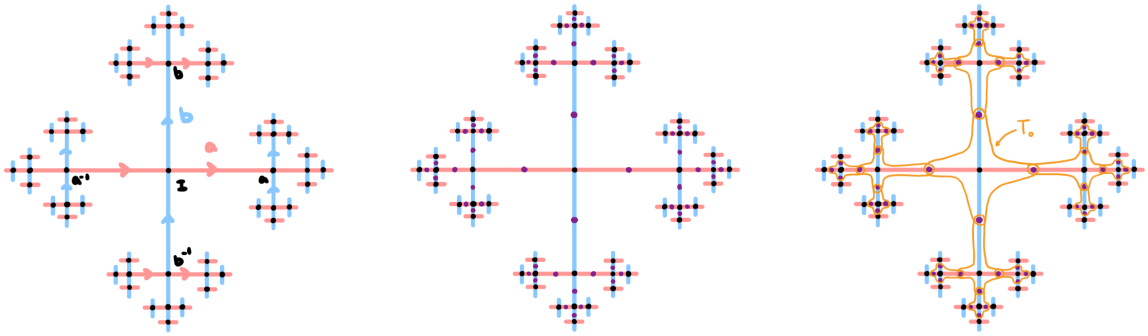

Figure 7: F_2-Tiling of the Cayley graph of F_2

By the definition of a tile, we want to show that every T_g is a subtree of T' (the barycentric subdivision of T). This is the same as showing that every T_g is a connected subgraph of T' because any connected subgraph of a tree is a tree.

We already know that T_g is a subgraph of T' since all its vertices and edges are contained in T' by construction. All that is left to show is that T_g is connected.

Let w be a vertex of T_g, so $w \in W_g$. We will show that every vertex in the unique path from w to gv in T' is contained in T_g. This automatically implies that the edges along the path are also in T_g by construction of T_g. This would prove that T_g is connected.

Let $d(w, gv) = n$ on the path metric. Let vertex $u \in T'$ be the first vertex in the path from w to gv. So $d(u, gv) = n - 1$.

Suppose $u \notin T_g$, then there exists some $g' \in G$ such that

$$d(u, g'v) = m < n - 1$$

Then

$$d(w, g'v) \leq m + 1 < n$$

This contradicts the fact that w is in T_g, since it would be closer to $g'v$ than to gv. So u must also be in T_g. Now this process can be repeated with u instead of w, all the way along the unique path back to gv, showing that w is connected to gv in T_g. Hence every T_g is a connected subgraph of T' and so is a tile.

Claim 15. *The union of the T_g is all of T'*

Now that we have established that T_g as previously defined is a tile, we want to show that the union of all of these tiles forms T', the barycentric subdivision of T.

By the definition of T_g, every vertex in T' must lie in some T_g, since said vertex must be closest to some gv. Thus, to show that T' is the union of all T_g, we must show that every edge of T' is in some T_g.

Since T' is the barycentric subdivision of T, every edge in T' will have one vertex u that is also a vertex of T, and one vertex w that is only a vertex of T'. We can then observe that for any edge path in T', the vertices will alternate between vertices of T, and vertices of T' that are not vertices of T.

Again taking distance to be the path metric in T', we can see that the distance between any two vertices u_1, u_2 of T will be even. To see this, note that two adjacent vertices of T will have two edges between them, and we can then see that the distance between any two vertices of T must be even. Since the G-orbit of v is a set of vertices in T, we can then see that the distance from any u of T to the G-orbit of v will be even. Similarly, the distance from any vertex w of T', where w is not a vertex of T, to the G-orbit of v will be odd, since any vertex w is exactly one edge from its closest vertex in T, and the distance to any other vertex in T from this point will be even.

Now, suppose that u and w are vertices of the same edge, and that the distance from u to the G-orbit of v is less than the distance from w to the G-orbit of v, and assume that u lies in T_g. Since u and w are adjacent, the distance between them is 1. By the triangle inequality, we have that

$$d(w, gv) \leq d(u, gv) + 1$$

However, we have also assumed that

$$d(u, Gv) \leq d(w, Gv)$$

(Note that $d(u, Gv) = d(u, gv)$ since u lies in T_g).

We then get that

$$d(w, gv) \leq d(u, gv) + 1 \leq d(w, Gv)$$

Since distance in the path metric must be an integer, this inequality shows that $d(w, Gv) = d(w, gv)$, and hence w lies in T_g. Thus, the edge between u and w lies in some T_g for every u of T and w of T', so T' is the union of all T_g.

Claim 16. *For all $g, h \in G$, $gT_h = T_{gh}$*

Let $h = \mathrm{Id}$, this implies $T_h = T_0$. Recall that every tile, T_g, is equal to gT_0 for some $g \in G$. Since we already know the vertex set of T_0, we can show that the vertex set of T_0 goes to the vertex set of each T_g upon action by g. We'll prove this for an arbitrary vertex, $u \in T_h$.

Suppose u is a vertex of T_h, then for all $k \in G$,

$$d(u, v) \leq d(u, kv) \text{ if } u \in T_h$$

Similarly,

$$d(gu, gv) \leq d(gu, kv) \text{ if } gu \in T_{gh}$$

Considering G acts by isometries, apply g^{-1}:

$$\begin{gathered} g^{-1} \cdot d(gu, gv) \leq g^{-1} \cdot d(gu, kv) \\ d(u, v) \leq d(u, (g^{-1}k)v) \text{ for all } k \in G \\ d(u, v) \leq d(u, kv) \text{ for all } k \in G \end{gathered}$$

The action of g^{-1} on the tile T_{gh} has yielded the original inequality with which we started (the one that characterizes u as a vertex of T_h). This implies the action of g^{-1} on the tile T_{gh} sends its vertices to those of T_h, proving $gT_h = T_{gh}$ for all $g \in G$.

Step 2: Finding a Symmetric Generating Set for G

Let $S = \{g \in G \mid (gT_0) \cap T_0 \neq \varnothing\}$, that is $(gT_0 \cap T_0)$ is a single vertex of T' since two tiles either intersect at a single vertex of T' or do not intersect.

To show that S is symmetric, we will show that if $s \in S$, then $s^{-1} \in S$. Let $s \in S$. Then

$$(sT_0) \cap T_0 = \{w\}$$

where w is a vertex of T'. Apply s^{-1} to obtain

$$T_0 \cap (s^{-1}T_0) = \{s^{-1}w\}$$

We have that $s^{-1}w \in T'$, thus $s^{-1} \in S$, thus S is symmetric.

Now we must show that S generates G. Let $g \in G$. We will show that g is a product of elements in S. For any vertex gv obtained by the element g acting on the tree, consider the path going from gv back to v. This path is unique, and goes along the tiles

$$T_{g_n}, T_{g_{n-1}}, \cdots, T_{g_1}, T_{g_0}$$

where $g_n = g$ and $g_0 = \mathrm{Id}$.

Claim 17. *Each $g_{i-1}^{-1}g_i$ is equal to some $s_i \in S$*

Proof. If a path from gv to v goes through the tiles $T_{g_{i+1}}$ and T_{g_i} without going through any other tiles in between, then $T_{g_{i+1}} \cap T_{g_i}$ is a single vertex. By applying g_i^{-1}, we obtain

$$(g_i^{-1}T_{g_{i+1}}) \cap (g_i^{-1}T_{g_i}) = T_{g_i^{-1}g_{i+1}} \cap T_0 \neq \varnothing$$

Where we applied the result of Claim 16. Hence $g_i^{-1}g_{i+1} \in S$. □

Now that we have proved this claim, it follows by induction, shown below, that $g_n = s_1 s_2 \cdots s_n$.

Proof. Base case ($n = 1$): $s_1 = g_{1-1}^{-1} g_1 = g_1$.

Assume that $g_k = s_1 s_2 \cdots s_k$.

We want to show that $g_{k+1} = s_1 s_2 \cdots s_{k+1}$. By the previous claim, we have that

$$s_{k+1} = g_{k+1-1}^{-1} g_{k+1} = g_k^{-1} g_{k+1}$$

Thus, if we apply g_k to the equality and use the induction hypothesis, we obtain

$$g_k s_{k+1} = g_{k+1} = s_1 s_2 \cdots s_k s_{k+1}$$

Hence for any $g \in G$, $g = g_n = s_1 s_2 \cdots s_n$, thus S is a generating set for G. □

Step 3: Free Generation Here we will show that S generates G freely, so G is a free group. Free generation means that the word associated with each element of G, which is a product of elements in S, is unique.

Let $g = s_1 \dots s_k$ be a freely reduced word. We will use this word to construct a tile path that follows the vertices from T_g to T_0. Because this tile path is along a tree, it will be unique, and thus the word is also unique.

Let v be a vertex in T_0. In order to trace a path from $gv \in T_g$ to v, we'll have to pass through the following tiles:

$$T_g, T_{s_1 \dots s_{k-1}}, \dots, T_{s_1}, T_0$$

To achieve such a path, we'll first notice that the path from v to $s_1 v$ is the tree $T_0 \cup T_{s_1}$. It's important to note that the union of these two trees is itself a tree, implying that, so far, the path from v to $s_1 v$ is unique. We can repeat this process when finding the path from v_{s_1} to $v_{s_1 s_2}$, noting that the union of tiles T_{s_1} and $T_{s_1 s_2}$ is, again, a tree. Inductively one can conclude the tile path from v to gv is

$$T_0 \cup T_{s_1} \cup \cdots \cup T_g.$$

Again, this tile path is a tree since it's the union of trees, and with that we can confirm the reduced word associated with each element of G is unique, and so S generates G freely.

□

6 Conclusion

We went through the basics of geometric group theory and one very important theorem, which is just one example of the algebraic information that can be gained from studying the geometry of groups.

There are many other topics in geometric group theory that we discussed but did not have a chance to delve into in great detail. Many of these include properties of Cayley graphs that can be used to study the group, including quasi-isometries, ends of groups, and hyperbolic groups. This introduction to geometric group theory is the basis needed to dive into a whole branch of math with some amazing results.

7 Acknowledgments

We want to thank our amazing mentor, Jessie Meanwell, for her support throughout the semester. She always encouraged us to take the lead and try to explain new concepts to each other, while also giving great explanations and insights when we struggled. We learned so much with her guidance and are very grateful we had her as a mentor.

We also want to thank everyone who ran the Directed Reading Program at McGill University for agreat experience this year.

References

[CM17] M. Clay and D. Margalit. *Office Hours with a Geometric Group Theorist*. Princeton University Press, 2017.