Geometric Group Theory: Part I

Sam Streeter

Autumn 2022

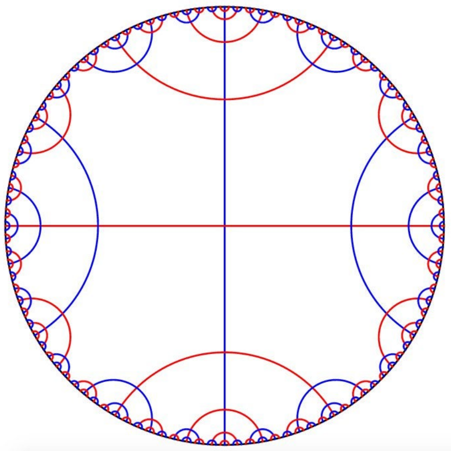

Contents

Introduction

Geometric group theory does what it says on the tin: it seeks to use the tools of geometry to understand the algebraic structure of groups. The link between algebra and geometry is underscored by this quote from Sophie Germain, one of the great mathematicians of the late 18th and early 19th centuries.

"L'algèbre n'est qu'une géométrie écrite, la géométrie n'est qu'une algèbre figurée"

("Algebra is but written geometry, geometry is but drawn algebra")

The central idea of geometric group theory is as follows: build/find an object on which a group acts, and its geometry will (hopefully!) tell you something about the group.

Plan

A rough outline of this half of the course:

1. Group presentations: how to think about groups in terms of generators and relations
2. Cayley graphs: how to make a graph from a group
3. Quasi-isometries: how to use a Cayley graph to describe a group/get new invariants
4. Geodesics and actions: how a nice action of a group on a space with points connected by nice curves can tell us about the group

Admin

- Lectures: 7 for each half
  - Mondays, 11:00 - 12:00, Room 3.21 (Berry), Physics Building
  - Wednesdays, 09:00 - 10:00, Room SM4, Ada Lovelace Building
  - Fridays, 16:00 - 17:00, Room G.09, Fry Building
- Problems classes: 2 for each half - Fridays of Weeks 2 and 3 for this half
- Homework: two assignments, four questions each
- Exam: 90 minutes total (so about 45 minutes per half) in January
- Contact info:
  - Name: Sam Streeter
  - Office hour: Friday, 15:00 - 16:00, Office 2.11, Fry Building
  - Email: sam.streeter@bristol.ac.uk
  - Website: `https://sites.google.com/view/samstreeter/home`

References

1. *A course on geometric group theory*, Brian Bowditch
   `https://www.math.ucdavis.edu/~kapovich/280-2009/bhb-ggtcourse.pdf`
2. *Geometric group theory, an introduction*, Clara Löh
   `https://www.math.ucdavis.edu/~kapovich/280-2009/bhb-ggtcourse.pdf`

Source for title page graphic.

Chapter 1

Group presentations

The goal of this chapter is to become familiar with thinking about groups as being produced by a set of elements called "generators" which satisfy some "relations" amongst themselves. A way of describing a group like this is called a *presentation* of the group.

We will start by talking about generating sets. We will then talk about free groups, which are in some sense the largest group you can make with a set of generators. We will then talk about what it means to impose some relations on the generating set.

1.1 Generating sets

Let G be a group, and let $A \subset G$ be a subset. Suppose we want to add elements to A to make it into a subgroup. How can we do that? One trivial way is to add every element of $G \setminus A$. However, perhaps we don't need to go that far. It's more interesting if we look for the *smallest* possible subgroup containing A. We will denote this subgroup (at least once we have proven that it exists) by $\langle A \rangle$. What does that look like?

We certainly need to add the inverse of every element of A if it is not in A already, and we certainly need to have all products of elements in A. We need the identity too, but we get that for free if we have inverses and products, since for $a \in A$, we have $aa^{-1} = 1_G$. Then our guess would be

$$\{a_1^{\epsilon_1} a_2^{\epsilon_2} \cdots a_n^{\epsilon_n} : n \in \mathbb{N}, a_i \in A, \epsilon_i \in \{\pm 1\}\}.$$

It is easy to check that this set is a subgroup of G, and by the discussion above, it is the smallest one containing A, since any subgroup of G containing A also has to contain it. So we have explicitly constructed $\langle A \rangle$.

There is another way of thinking about $\langle A \rangle$: it is the intersection of all subgroups containing A. Further, we can characterise it uniquely as the subgroup $H \leq G$ for which the following properties hold:

1. $A \subset H$.
2. If $H' \leq G$ and $A \subset H'$, then $H \subset H'$.

Sometimes the two obvious subgroups containing A coincide. That is, if you start out with the elements in A and allow yourself to take inverses and form products, then you end up generating the whole group G, hence the following definition.

Definition 1.1. A group G is *generated* by a subset $A \subset G$ if $G = \langle A \rangle$. We call A a *generating set* for G. We say that G is *finitely generated* if it has a finite generating set.

We usually write $\langle a_1, \ldots, a_n \rangle$ for $\langle \{a_1, \ldots, a_n\} \rangle$.

Example 1.2. Generating sets:

1. Any group is generated by itself, and any finite group is finitely generated by itself.
2. The group $\mathbb{Z}^2$ (under addition) is generated by $\{(1,0),(0,1)\}$: we have $(a,b) = a(1,0) + b(0,1)$.
3. Note that the size of a generating set may vary: $\mathbb{Z}^2$ is also generated by $\{(2,0),(0,3),(1,1)\}$.

4. Many matrix groups over $\mathbb{Z}$ are finitely generated, e.g. $\mathrm{GL}_n(\mathbb{Z})$ and $\mathrm{SL}_n(\mathbb{Z})$. However, it is in many cases difficult to find an explicit set of generators.

5. $\mathbb{Q}$ is not finitely generated.

6. Any finitely generated group is countable as a set, hence uncountable groups are not finitely generated–in particular, $\mathbb{R}$ and $\mathbb{C}$ (as additive groups).

Although generating sets for a group may have different cardinalities, there is always a smallest such size for any given group.

Definition 1.3. The *rank* of a group G, sometimes denoted by $\operatorname{rank} G$, is the smallest cardinality of a generating set for G.

This is not to be confused with the order of G, which is the cardinality of G as a set. It's fair to say that these two quantities are rarely equal, as the following exercise shows.

Exercise 1.4. Show that $|G| = \operatorname{rank} G$ if and only if G is the unique one-element group.

While maximal rank for given order is somewhat dull, going the other way is more interesting.

Definition 1.5. A group G is *cyclic* if $G = \langle a \rangle$ for some $a \in G$, i.e. if it is generated by a single element.

Example 1.6. Cyclic groups:

1. The group $\mathbb{Z}$ is cyclic, generated by 1. Any quotient of a cyclic group is cyclic; in particular, $\mathbb{Z}/n\mathbb{Z}$ ($n \in \mathbb{N}$) is cyclic. In fact, these are, up to isomorphism, all of the cyclic groups!

2. The group $\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$ is the smallest non-cyclic group.

3. Any group of prime order is cyclic. (Apply Lagrange's theorem to the subgroup generated by a non-identity element.) However, the converse is obviously not true: $\mathbb{Z}/4\mathbb{Z}$ is cyclic, for instance.

Definition 1.7. The (isomorphism class of) the group $\mathbb{Z}^n$ is the *free abelian group of rank* n.

Warning 1.8. Free abelian groups are not free (except for rank 1).

We will soon meet free groups. The use of "free" in both cases basically means "as few relations as possible", where relations mean some identity of different products in the group, e.g. $g_1 g_2^2 = g_1^{-1}$. So a free abelian group G has as few relations as possible, subject to the fact that it is abelian and so must satisfy the relations $gh = hg$ for all $g, h \in G$.

Symmetry is a nice thing. If you agree, then you will be pleased to hear that we can always turn a generating set into a symmetric one as follows.

Suppose we have a generating set A for G, so $G = \langle A \rangle$. We can produce the set of inverses of elements of A, namely

$$A^{-1} = \{a^{-1} : a \in A\}.$$

We say that A is *symmetric* if $A = A^{-1}$.

Note that $A \cup A^{-1}$ is symmetric, and $G = \langle A \cup A^{-1} \rangle$.

Unfortunately, we can't have everything we want; minimality is also nice, and $A \cup A^{-1}$ is almost never a minimal generating set.

If A is symmetric, then every element of $\langle A \rangle$ is of the form $a_1 a_2 \ldots a_n$, $a_i \in A$ (not necessarily distinct), $n \in \mathbb{N}$. We call this a *word* of length n in the alphabet A (with 1_G being the empty word/word of length zero).

1.2 Free groups

As alluded to earlier, free groups should have as few relations as possible. There are obvious inevitable relations, namely $gg^{-1} = 1_G$ for any element $g \in G$. The idea is that we will ask only for these relations (and relations that follow because of them, e.g. $h_1 g g^{-1} h_2 = h_1 h_2$). We will formalise this in the following way.

Definition 1.9. Let S be a subset of a group F. We say that F is *freely generated* by S if for every group G and map of sets $\phi : S \to G$, there exists a unique homomorphism $\widehat{\phi} : F \to G$, i.e. $\widehat{\phi}(x) = \phi(x)$ for all $x \in S$. We say that F is a *free group* if it is freely generated by some subset $S \subset F$.

The above definition is pictorially summarised by the commutative diagram below.

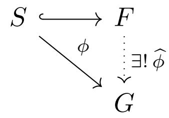

That this formally encapsulates there being no "unnecessary" relations is far from clear at first sight! Before we think about what it means a little more, let us first convince ourselves that it at least makes some sense, in that if S freely generates F, it certainly generates F.

Lemma 1.10. *If a group F is freely generated by a subset $S \subset F$, then $F = \langle S \rangle$.*

Proof. Suppose that F is freely generated by S. To make use of the above definition, we want to start with a map of sets from S to some group G. We want to involve $\langle S \rangle$, so let's choose the map $\phi : S \to \langle S \rangle$ given by inclusion. By hypothesis, this extends to a homomorphism $\widehat{\phi} : F \to \langle S \rangle$. We can compose this with the inclusion homomorphism $\iota : \langle S \rangle \hookrightarrow F$ to obtain a homomorphism $\iota \circ \widehat{\phi} : F \to F$. Now, note that both $\iota \circ \widehat{\phi}$ and id_F are homomorphisms extending the inclusion $S \hookrightarrow F$; by hypothesis, they must be equal, but then $\widehat{\phi}(f) = f$ for all $f \in F$, i.e. $\langle S \rangle = F$. □

Let's now try to make some sense of why this definition means that S generates F with as few relations as possible. First, suppose that S is a generating set for F, i.e. $F = \langle S \rangle$, and suppose that we have a map of sets $\phi : S \to G$, where G is a group. If ϕ extends to a homomorphism $\widehat{\phi} : F \to G$, it is clear what this homomorphism must be. Indeed, given an arbitrary element $f \in F$, we know that we can write $f = s_1^{r_1} s_2^{r_2} \cdots s_n^{r_n}$ for some $n \in \mathbb{N}$, $r_i \in \mathbb{Z}$ and $s_i \in S$. Then we have

$$\widehat{\phi}(f) = \widehat{\phi}(s_1)^{r_1} \widehat{\phi}(s_2)^{r_2} \cdots \widehat{\phi}(s_n)^{r_n} = \phi(s_1)^{r_1} \phi(s_2)^{r_2} \cdots \phi(s_n)^{r_n},$$

where the first equality holds since $\widehat{\phi}$ is a homomorphism, and the second equality holds since $\widehat{\phi} = \phi$ on S. We deduce that, if such an extension exists, it is unique. The idea is that S freely generating F means that no pesky relations will crop up which means that this extension is not well-defined. Let's see how this could happen in an example.

Example 1.11. Recall that $\mathbb{Z}^2 = \langle (1,0), (0,1) \rangle$. Consider the map $\phi : \{(1,0), (0,1)\} \to S_3$ given by $\phi(1,0) = (123)$, $\phi(0,1) = (12)$. If this extends to a homomorphism $\widehat{\phi} : \mathbb{Z}^2 \to S_3$, then we must have $(13) = (123)(12) = (12)(123) = (23)$ (can you see why?), which is false.

We could have used any non-abelian group instead of S_3 above. The key was that the relation $(1,0) + (0,1) = (0,1) + (1,0)$ spoiled the party, so that the images of the two elements had to commute, so we had only to choose two non-commuting elements to map to.

We defined free abelian groups of a given rank up to isomorphism and so spoke of *the* free abelian group of rank n. We can also speak of *the* free group of rank n, as the following lemma tells us.

Lemma 1.12. *Let F be freely generated by S and let F' be freely generated by S'. Suppose that S is in bijection with S', i.e. $|S| = |S'|$. Then $F = F'$.*

Proof. Once again, we look for obvious maps to apply the definition to. Since S and S' are in bijection, there exist maps $\phi : S \to S'$ and $\psi : S' \to S$ such that $\psi \circ \phi = \mathrm{id}_S$ and $\phi \circ \psi = \mathrm{id}_{S'}$. By composing these maps with the obvious inclusions, we can think of them as maps $\phi : S \to F'$ and $\psi : S' \to F$.

By hypothesis, we can uniquely extend ϕ and ψ to homomorphisms $\widehat{\phi} : F \to F'$ and $\widehat{\psi} : F' \to F$. Consider the composition $\widehat{\psi} \circ \widehat{\phi} : F \to F$. Note that, for $s \in S$, we have

$$\widehat{\psi} \circ \widehat{\phi}(s) = \widehat{\psi}(\phi(s)) = \psi(\phi(s)) = s,$$

hence this homomorphism extends the inclusion $S \hookrightarrow F$. Since such an extension is unique and is also given by id_F, we deduce that $\widehat{\psi} \circ \widehat{\phi} = \mathrm{id}_F$. Similarly/by symmetry, we obtain $\widehat{\phi} \circ \widehat{\psi} = \mathrm{id}_{F'}$. Then $F \cong F'$. □

Definition 1.13. We denote the (isomorphism class of the) group freely generated by $\{1, \ldots, n\}$ by F_n.

As for free abelian groups, isomorphic free groups have the same rank. In particular, if $F_m \cong F_n$, then $m = n$. This will become clear once we reveal the link between F_n and $\mathbb{Z}^n$ later on.

1.2.1 Explicit construction

So far, we've assumed that free groups exist and derived some interesting results about them. Now we'd better check that they actually exist! The best way is to construct them explicitly.

Let's talk about words. Begin with a set A, which we call our *alphabet* (with its members our letters). A *word of length* n in A, as you might guess, is an ordered sequence of n elements of A, which is the same thing as a map $\{1, \dots, n\} \to A$. Denoting by a_i the image of $i \in \{1, \dots, n\}$, we write the resulting word as $a_1 a_2 \dots a_n$. We denote by $W(A)$ the set of all words in the alphabet A. No surprises so far!

Note that there is an obvious binary operation on $W(A)$: *concatenation*. That is, we can form a new word from two old words by pushing them together, i.e. we take $a_1 a_2 \dots a_n \cdot a_{n+1} a_{n+2} \dots a_{n+m} = a_1 a_2 \dots a_n a_{n+1} a_{n+2} \dots a_{n+m}$. This is clearly associative, and the identity is the empty word as before. This makes $W(A)$ into a *monoid* (set with associative operation and identity element)–you don't need to worry about monoids, just be happy that we're getting closer to groups!

What about inverses? Well, the obvious thing to do is to add them, and that is what we will do.

We are ready to construct the free group freely generated by a set S. Starting with our set S, let $\overline{S}$ be another set in bijection with S; $\overline{S}$ will play the role of the set of inverses to elements of S, so we'll write its elements as $\overline{s}, s \in S$ to easily go to and fro. Then one bijection is simply $s \mapsto \overline{s}$. Set $A : S \cup \overline{S}$, and form the set of words $W(A)$. We will now introduce an equivalence relation on $W(A)$. Since $\overline{s}$ is meant to be the inverse of s, we will declare the words $us\overline{s}v$, uv and $u\overline{s}sv$ to be equivalent for all $u, v \in W(A)$, $s \in S$. We will call uv a *reduction* of the words $us\overline{s}v$ and $u\overline{s}sv$. This will generate an equivalence relation $\sim$ on $W(A)$, where two words are equivalent if and only if they have a common reduction. We claim that $F(S) \cong W(A)/\sim$, where the group operation is given by concatenating representatives from each equivalence class, i.e. $[w] * [w'] = [ww']$.

Where it does not cause too much confusion, we will write w instead of $[w]$ and a^{-1} instead of $\overline{a}$.

Example 1.14. In $F(a, b)$, we have the words a^3, $b^2 a^{-1}$ and aba^3, and we have $abaa^{-1}b^{-1} = abb^{-1} = a$.

Before we convince ourselves that we have succeeded in constructing a free group, we should check that $W(A)/\sim$ really is a group. To do that, we need to check that the group operation is well-defined (i.e. independent of the representatives we choose to concatenate) and that there are an identity and inverses (can you guess what these are?). Having done that, we are left to check freeness.

Exercise 1.15. Check that $W(A)/\sim$ is a group.

Definition 1.16. A word $w \in W(A)$ is *reduced* if it admits no reduction.

Proposition 1.17. *Given a word $w \in W(A)$, there exists a unique reduced word w' such that $w \sim w'$.*

Existence is easy; keep reducing until you can't anymore. Uniqueness is less clear; it is one of the things that Cayley graphs will allow us to see.

Note 1.18. Free generating sets are not unique: $\mathbb{Z}$ is freely generated both by 1 and by -1.

1.3 Group presentations

Definition 1.19. Let G be a group, and let $A \subset G$. The *normal closure* $\langle\langle A \rangle\rangle$ of A in G is the smallest normal subgroup of G containing A.

1. Note that there is some ambiguity in notation here: $\langle\langle A \rangle\rangle$ is not in general equal to $\langle\langle A \rangle\rangle = \langle A \rangle$, the smallest subgroup containing $\langle A \rangle$. In general, we have only the inclusion $\langle A \rangle \subset \langle\langle A \rangle\rangle$ by minimality of $\langle A \rangle$.

2. Note also that $\langle\langle A \rangle\rangle$ admits a similar characterisation to that of $\langle A \rangle$:

   (a) $A \subset \langle\langle A \rangle\rangle$.

   (b) $\langle\langle A \rangle\rangle \lhd G$.

   (c) If $A \subset H \lhd G$, then $\langle\langle A \rangle\rangle \leq H$.

3. We may explicitly describe $\langle\langle A \rangle\rangle$ as

$$\langle\langle A \rangle\rangle = \{ghg^{-1} : g \in G, h \in \langle A \rangle\}.$$

   We are simply turning $\langle A \rangle$ into a normal subgroup by adding all possible conjugations of one of its elements by an element of G, and it is easily seen that this is a group.

Example 1.20. If G is abelian, then for any $A \subset G$, we have $\langle A \rangle = \langle\langle A \rangle\rangle$.

Exercise 1.21. The *centre* $Z(G)$ of a group G consists of all elements of G commuting with every other element, i.e.

$$Z(G) = \{z \in G : zg = gz \text{ for all } g \in G\}.$$

Prove that the centre $Z(G)$ is a normal subgroup of G.

Definition 1.22. Let S be a set, and let R be a subset of the free group $F(S)$. Define

$$\langle S \mid R \rangle := F(S)/\langle\langle R \rangle\rangle.$$

A *presentation* of a group G is an isomorphism of G with a group of the form $\langle S \mid R \rangle$. We call elements of R *relators* or *relations*. We say that a presentation is finite if S and R are finite, and we say that G is *finitely presented* if it admits a finite presentation.

The way you should read a presentation is as a recipe for constructing a group: on the left you have your base ingredients, and on the right, you specify how they mix together (more specifically, what combinations of generators give the identity). Alternatively, you could think of relators as rules for what happens when you combine elements.

Example 1.23.
1. Note that $\langle\langle \emptyset \rangle\rangle = \{1\}$, hence $\langle S \mid \emptyset \rangle \cong F(S)$.
2. $\mathbb{Z}/n\mathbb{Z} \cong \langle a \mid a^n \rangle$. We have one element that generates the whole group (concretely, $\overline{1}$), and the one rule which determines the group is that n lots of this element equals the identity.
3. If we want to specify that two generators a and b commute with each other, we can put $aba^{-1}b^{-1}$ into our set of relators. In particular, you can check that $\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$ has presentation $\langle a, b \mid a^2, b^2, aba^{-1}b^{-1} \rangle$.

Note 1.24 (Word problem)**.** Note that R almost never consists of *all* products of elements which equal the identity, but the ones that we start with imply others. For example, in $\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z} = \langle a, b \mid a^2, b^2, aba^{-1}b^{-1} \rangle$, we know that $ab = ba$, so $a^2b^2 = a(ab)b = a(ba)b = (ab)^2$.

In particular, certain products in our generating set (words in the generators and their inverses, if you like) will equal the identity but aren't already in the relator set, and so you may ask: how can we tell when a product equals 1? This is (one version of) the *word problem*: given an arbitrary word, can we decide if it reduces to the identity? For many groups the answer is unknown, and understanding the word problem plays a prominent role in group theory.

Note 1.25. Being finitely presented is strictly stronger than being finitely generated, but examples illustrating this are a bit obscure, and this is not something that we'll be overly worried about.

Definition 1.26 (Commutators)**.** Let G be a group, and let $x, y \in G$. The *commutator* of x and y is the element $[x, y] = xyx^{-1}y^{-1} \in G$. The *commutator subgroup* of G, denoted by $[G, G]$, is the group generated by the set of all commutators, i.e.

$$[G, G] = \langle x, y : x, y \in G \rangle.$$

Since $[G, G]$ is normal, we can form the quotient $G/[G, G]$. Note that this quotient group is abelian. In fact, it is the canonical way of turning G into an abelian group, hence the following definition.

Definition 1.27. The group $G/[G, G]$ is called the *abelianisation* of G.

Exercise 1.28. $F_n/[F_n, F_n] \cong \mathbb{Z}^n$.

Having completed the above exercise, you can now see why $F_m \cong F_n$ implies that $m = n$.

The abelianisation of G is the minimal abelian group receiving a homomorphism from G in the following sense: if $\phi : G \to A$ is a homomorphism with A abelian, then there exists a unique homomorphism $\psi : G/[G, G] \to A$ through which ϕ factors, i.e. $\phi = \psi \circ q$, where $q : G \to G/[G, G]$ is the quotient map. That is, we have the following diagram:

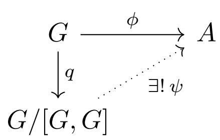

Chapter 2

Cayley graphs

Now that we have familiarised ourselves with one way of thinking about and constructing groups, we are ready to start constructing geometric objects from them. In this chapter, we will describe how to construct a Cayley graph $\Delta(G, S)$ given a group G and subset S (generally a generating set). We will then explain how to think of these graphs as metric spaces. By opening the door to metric geometry in this way, we will see later on (in the next chapter) that different choices of generating set S preserve the "large-scale geometry" of $\Delta(G, S)$, so that we geometric properties at this scale can be meaningfully discussed independent of the choice of S, so that we can think of them as geometric properties of G itself.

2.1 Graphs

Before we talk about Cayley graphs, we'd better talk about graphs in general.

Definition 2.1. A *directed graph* $K = (V, E)$ consists of the following data:

1. A set of *vertices* (or *nodes*) $V = V(K)$
2. A collection (not necessarily a set) of *edges* $E = E(K) \subset V \times V$. We may call the vertices in the pair defining an edge its *endpoints*.

An *undirected graph* is defined analogously, except that we replace $V \times V$ by $\mathrm{Sym}^2(V) = V \times V/\sim$, where $\sim$ is the equivalence relation defined by $(x, y) \sim (y, x)$. That is, edges in an undirected graph are unordered pairs of vertices instead of ordered pairs.

We usually picture the vertices of our directed graph as points in some ambient space, often the plane, and the edges as curves of arrows joining these points, so that the edge (u, v) is a curve joining the points u and v, pointing from u to v. The picture of an undirected graph is then the same, but with no arrows.

Example 2.2. One of the most well-studied families of graphs is the family of *complete graphs* K_n, $n \in \mathbb{N}$. The (undirected) graph K_n has n vertices and an edge connecting each pair of distinct vertices.

Exercise 2.3. Determine the number of edges in K_n.

Definition 2.4. We say that a graph is *simple* if any two vertices are connected by at most one edge; in particular, a simple graph contains no *loops* (edges starting and ending at the same vertex).

If a graph is simple, then we can safely write $\{u, v\}$ instead of (u, v) to denote an edge.

Now that we know/have recalled what graphs are, let's describe some of their properties.

Definition 2.5. The *degree* (or *valence*) of a vertex is the number of edges incident upon it, i.e. including it as one of their endpoints. A graph is *locally finite* is each vertex has finite degree. Given $n \in \mathbb{N}$, a graph is *n-regular* if each vertex has degree n. A graph is *regular* if it is n-regular for some $n \in \mathbb{N}$.

Now for paths and the gang.

Definition 2.6. A *path* is a sequence of edges such that every consecutive pair of edges share a vertex. An *arc* is a path in which any vertex appears as a common endpoint at most once, i.e. a path which does not cross itself (or an embedded path, if you like).

A *cycle* is a path such that the first and final edges share a vertex, i.e. a closed path. A *circuit* is a cycle which does not cross itself at any intermediate step (meaning that the only vertex to which the cycle returns is the first one).

Definition 2.7. A graph is *connected* if there is a path between any two vertices. A graph is a *tree* if it contains no cycles.

Exercise 2.8. Prove that a graph is a tree if and only if there is at most one path between any pair of vertices.

Example 2.9. Another family of graphs relevant to us is the family of *n-regular trees* T_n, $n \in \mathbb{N}_{\geq 2}$. The graph T_n has infinitely many vertices, with each having degree n.

It goes without saying that graphs are a concept of fundamental importance in applied mathematics; they can be used to represent and understand road networks, computer systems and even the human brain. They are also widely studied in pure mathematics; graph theory is widely considered a branch of combinatorics, and there is a panoply of beautiful and striking results dating from what many consider to be the origin of the subject: Euler's solution to the "Seven Bridges of Königsberg" problem.

The longstanding problem, solved negatively by Euler in 1736, was to find a walking route around the city of Königsberg traversing each bridge exactly once. Unfortunately for Prussian pedestrians, Euler showed that this is impossible. One of his key observations was that the walking you do between bridges/on the same land mass is irrelevant, so that the problem is reduced to finding a circuit in a certain graph on seven vertices which traversed every edge. Such circuits are called *Eulerian circuits* in his honour. His solution was pleasingly simple: he observed that, for an Eulerian circuit to exist in a graph, at most two vertices can have odd degree. Nice one, Euler!

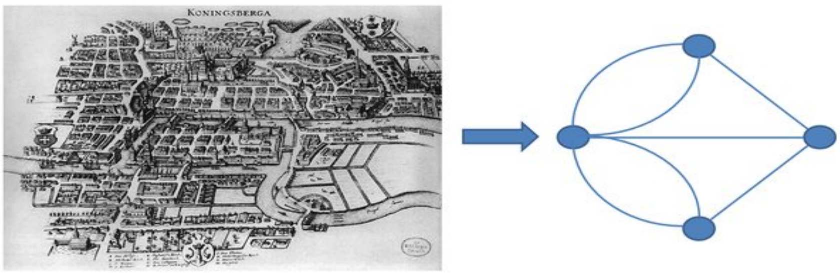

Source

However, if you like bridges (and you know you do), you will be pleased to learn that such a feat on foot is possible in other cities–including Bristol! If you have sturdy shoes and plenty of free time on your hands, you can take in all 45 of Bristol's bridges in a 45km walk. The solution is believed to be the first published positive solution for any city, narrowly beating out New York. For more details and the route, see here. The solution was found by former Bristol academic Thilo Gross, who is as fond of a punning webpage title as he is of Eulerian circuits: `https://reallygross.de/ops/bridgewalk`.

2.2 Constructing Cayley graphs

Now that we've learned a bit/brushed up on graphs, we're ready to make graphs from groups. Let's go!

Definition 2.10. Let G be a group, and let $S \subset G$ be a subset not containing 1. The *Cayley graph* of G with respect to S is the directed graph $\Delta(G, S)$ with

- set of vertices $V(\Delta) = \{v_g : g \in G\}$;
- set of directed edges $E(\Delta) = \{\{v_g, v_{gs}\} : g \in G, s \in S\}$.

One can think of the vertex v_g as being labelled by g, and the directed edge $\{v_g, v_{gs}\}$ as being labelled by s. In other words, the edges incident upon each vertex are indexed by the elements of $S \cup S^{-1}$–at least if $S \cap S^{-1} = \emptyset$. When $S \cap S^{-1} \neq \emptyset$, the incident edges to a vertex are indexed by $S \cup \overline{S}$, where $\overline{S}$ denotes a set of formal inverses for S as seen above. More specifically, the outgoing edges are indexed by S, while the incoming edges are indexed by $\overline{S}$.

Note 2.11. Note that $\Delta(G, S)$ is a $2|S|$-regular graph. That is, every vertex has degree $2|S|$.

Conventions differ on whether a Cayley graph is directed or undirected. The advantage of the directed definition is that it is easier to make sense of what to do when $S \cap S^{-1} \neq \emptyset$. Note that, if $s^{-1} \in S$ for some $s \in S$, then we have two directed edges $\{v_g, v_{gs}\}$ and $\{v_{gs}, v_{(gs)s^{-1}} = v_g\}$, i.e. the vertices v_g and v_{gs} are connected by two edges in opposing directions rather than just one edge (and the same for v_g and $v_{gs^{-1}}$ by symmetry). So the directedness is at least helpful for construction.

There are also differing conventions on whether to call the above graph a Cayley graph when S is not a generating set for G. It is a legitimate definition of a graph either way, but the graph is more pleasant when $G = \langle S \rangle$ as the following lemma shows.

Lemma 2.12. *The subset S generates G iff $\Delta(G, S)$ is a connected graph.*

Proof. Note that Δ is connected if for any vertex v_g, $g \in G$, there exists a path in the underlying undirected graph between v_1 and v_g. Now note that paths between v_1 and v_g are in bijection with representations of g as a word in the alphabet $A = S \cup \overline{S}$, given by taking the product of the edge labels of the path. ☐

Definition 2.13. When S generates G, we call $\Delta(G, S)$ the *Cayley graph of G with respect to S*.

Example 2.14. Below are Cayley graphs of $G = \mathbb{Z}$ with different choices of generating set S.

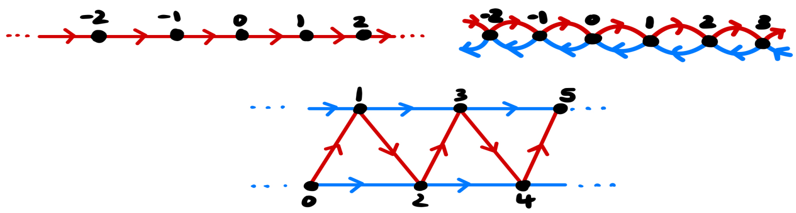

Example 2.15. While the cardinality of the vertex set of a Cayley graph tells us the cardinality of the underlying group, it does not uniquely determine the graph. Indeed, $G = S_3$ and $S = \{(12), (123)\})$ gives the same graph as $G = \mathbb{Z}/6\mathbb{Z}$ and $S = \{\overline{2}, \overline{3}\}$.

Theorem 2.16. *If $S \subset G$ freely generates G, then $\Delta(G, S)$ is a tree. The converse holds if $S \cap S^{-1} = \emptyset$.*

Proof. Suppose that $G \cong F(S)$ but Δ contains a circuit σ. Taking the (signed) product of the edge labels along σ, we produce a word in $S \cup \overline{S}$ equalling the identity. That this product equals the identity in $F(S)$ means that there exists a finite sequence of reduction steps taking this word to the empty word.

On the level of cycles, a reduction step (or its inverse) corresponds either to introducing or eliminating a backtracking along an edge. Note that there are never two edges joining the same two vertices; this is equivalent to there being no two-torsion elements in $S \subset F(S)$, which follows easily from the definition of a free generation. Thus we deduce that there exists a sequence of cycles $\sigma_0 = \sigma, \sigma_1, \ldots, \sigma_n = \{1\}$ such that σ_i is related to σ_{i+1} by introducing or eliminating a backtracking along an edge. Let $O(\sigma_i) \subset E(\Delta)$ denote the set of edges which σ_i travels along an odd number of times. Since eliminating a backtracking preserves this set, we have $O(\sigma_0) = O(\sigma_n)$. Now note that $O(\sigma_0) = E(\sigma_0)$ since σ_0 is a circuit, while $O(\sigma_n) = \emptyset$, hence $\sigma_0 = \sigma_n$, the constant path based at $\{1\}$. We have deduced a contradiction, hence the result follows.

For the converse, note that any element $g \in G$ corresponds to a unique reduced word via the unique path from v_1 to v_g with no backtracking, and this allows us to uniquely extend any map $S \to G$ to a homomorphism $F(S) \to G$. ☐

Note that G acts on itself by left multiplication. We can view this as an action of G on the vertices of $\Delta(G, S)$. This extends to an action of G on the edges of $\Delta(G, S)$, hence to an action on the graph. Indeed, if $g, h \in G$ with $h^{-1}g \in S$, then for any $k \in G$, we have $(kh)^{-1}(kg) = h^{-1}k^{-1}kg = h^{-1}g \in S$.

2.3 Cayley graphs as metric spaces

Definition 2.17. Let $K = (V, E)$ be a connected undirected graph. The *associated metric* on the vertex set V is given by

$$d_K(v, w) = \min\{n \in \mathbb{Z}_{\neq 0} : v \text{ and } w \text{ are connected by a path of length } n\}.$$

Formally, a graph consists of a set of vertices and a set of edges, both of which are in some sense discrete. It is (hopefully) intuitively clear, however, that there is an underlying geometric object of a more continuous flavour. When we draw a graph, in the plane say, we are really looking at a set of points which includes not only the vertices at each end of an edge, but also points along the edge. We formalise this as follows.

Definition 2.18. Let $K = (V, E)$ be a connected directed graph. The *(geometric) realisation* of K is the metric space $\left(|K|, d_{|K|}\right)$ defined by taking

$$|K| = E \times [0, 1]/\sim,$$

where $\sim$ is an equivalence relation defined by

$$((u, v), 1) \sim (\{v, w\}, 0).$$

The metric $d_{|K|}$ is defined by extending the associated metric on V, i.e.

$$d_{|K|}(((u,v),t),((u',v'),t')) = \begin{cases} |t-t'| \text{ if } (u,v) = (u',v'), \\ \min\left\{ \substack{t+t'+d_K(u,u'),\, t+(1-t')+d_K(u,v'), \\ (1-t)+t'+d_K(v,u'),\,(1-t)+(1-t')+d_K(v,v')} \right\} \text{ else.} \end{cases}$$

This looks quite complicated, but it makes sense: if both points lie on the same edge, we take their usual distance as points in the unit interval. Otherwise, a path between the points passes through one of the two vertices of each of their "home" edges, and we consider all four possible ways of choosing one vertex from each of the two edges and travelling between them.

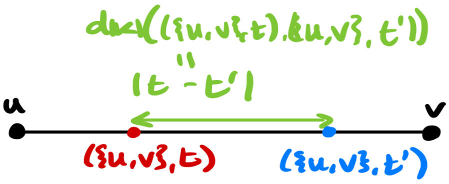

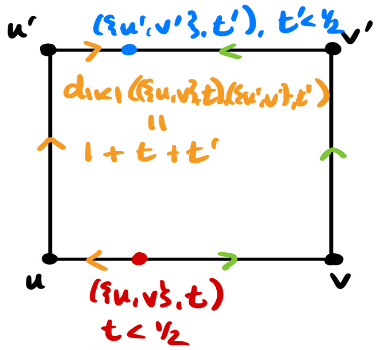

Chapter 3

Quasi-isometries

Now that we have established how to construct and think of Cayley graphs as metric spaces, we are almost ready to demonstrate that different choices of S preserve the large-scale/"coarse" geometry of the Cayley graph–but first, we of course need to make precise what we mean by the coarse geometry of a metric space, which we will do by defining and studying *quasi-isometries.* This is a good choice of name: as you will (hopefully!) recall, isometries are distance-preserving functions between metric spaces, so they preserve their geometry perfectly. The idea behind the definition of quasi-isometries is that they "almost preserve" distances, so that they roughly preserve the geometry (hence coarse is also a good word to use).

3.1 Quasi-isometries

Definition 3.1. Let (M_1, d_1) and (M_2, d_2) be metric spaces, and let $\phi : M_1 \to M_2$ be a (not necessarily continuous) map.

1. We say that ϕ is a *quasi-isometric embedding* if there exist constants $A \geq 1$ and $B \geq 0$ such that, for all $x, y \in M_1$, we have

$$\frac{1}{A}d_1(x, y) - B \leq d_2(\phi(x), \phi(y)) \leq Ad_1(x, y) + B.$$

   Occasionally it is desirable to specify parameters A and B such that ϕ satisfies the above condition, in which case we refer to ϕ as an (A, B)*-quasi-isometric embedding.*

2. We say that ϕ is *coarsely surjective* if there exists a constant $C > 0$ such that, for all $z \in M_2$, there exists $w \in M_1$ such that $d_2(\phi(w), z) \leq C$. Again, you could say that ϕ is C-coarsely surjective or has C-dense image if you want to keep track of the parameter.

3. We say that ϕ is a *quasi-isometry* if it is a coarsely surjective quasi-isometric embedding. Once more, you might like to specify parameters and so speak of an (A, B, C)-quasi-isometry.

Example 3.2. The map $\mathbb{R} \to \mathbb{R}$, $x \mapsto ax + b$ is, no surprise, a quasi-isometry. Note that its obvious parameters A, B as in the definition are not as you might expect; $(A, B) = (a, 0)$ works.

Exercise 3.3. The *diameter* of a metric space (M, d) is $\mathrm{diam}(M) = \sup_{x,y\in M}\{d(x, y)\} \in \mathbb{R} \cup \{\infty\}$. Let $\phi : M \to N$ be a map of metric spaces. We say that M is *bounded* if it has finite diameter. Prove the following facts:

1. If N is bounded, then ϕ is coarsely surjective.

2. If M and N are bounded, then ϕ is a quasi-isometric embedding.

Proposition 3.4 (Composition of quasi-isometries)**.** *Let (M_i, d_i), $i = 1, 2, 3$ be three metric spaces, and let $\phi : M_1 \to M_2$ and $\psi : M_2 \to M_3$ be quasi-isometries. Then the composition $\psi \circ \phi : M_1 \to M_3$ is a quasi-isometry.*

Proof. By assumption, there exist constants $A_i \geq 1$, $B_i \geq 0$ and $C_i > 0$ for $i = 1, 2$ such that

$$\frac{1}{A_1} d_1(x, y) - B_1 \leq d_2(\phi(x), \phi(y)) \leq A_1 d_1(x, y) + B_1 \text{ for all } x, y \in M_1,$$

$$\frac{1}{A_2} d_2(x, y) - B_2 \leq d_3(\psi(x), \psi(y)) \leq A_2 d_2(x, y) + B_2 \text{ for all } x, y \in M_2,$$

$$\widetilde{\mathrm{B}}(z, C_1) \cap \phi(M_1) \neq \emptyset \text{ for all } z \in M_2,$$

$$\widetilde{\mathrm{B}}(z, C_2) \cap \psi(M_2) \neq \emptyset \text{ for all } z \in M_3.$$

Let $x, y \in M_1$. We have

$$\begin{aligned} d_3(\psi \circ \phi(x), \psi \circ \phi(y)) &= d_3(\psi(\phi(x)), \psi(\phi(y))) \leq A_2 d_2(\phi(x), \phi(y)) + B_2 \\ &\leq A_2(A_1 d_1(x, y) + B_1) + B_2 = A_2 A_1 d_1(x, y) + A_2 B_1 + B_2, \end{aligned}$$

and similarly

$$\begin{aligned} d_3(\psi \circ \phi(x), \psi \circ \phi(y)) = d_3(\psi(\phi(x)), \psi(\phi(y))) &\geq \frac{1}{A_2} d_2(\phi(x), \phi(y)) - B_2 \\ &\geq \frac{1}{A_2}\left(\frac{1}{A_1} d_1(x, y) - B_1\right) - B_2 \\ &\geq \frac{1}{A_2 A_1} d_1(x, y) - (A_2 B_1 + B_2), \end{aligned}$$

using $A_2 \geq 1$ for the final inequality. Now for coarse surjectivity. We must find $C > 0$ such that $\widetilde{\mathrm{B}}(z, C) \cap (\psi \circ \phi)(M_1) \neq \emptyset$ for all $z \in M_3$. Given $z \in M_3$, we know that there exists $w \in M_2$ such that $\psi(w) \in \widetilde{\mathrm{B}}(z, C_2)$. Similarly, we know that there exists $v \in M_1$ such that $\phi(v) \in \widetilde{\mathrm{B}}(w, C_1)$. A natural candidate for an element of $\widetilde{\mathrm{B}}(z, C) \cap (\psi \circ \phi)(M_1)$ is therefore $\psi \circ \phi(v)$, and we will determine whether making C sufficiently large makes this candidate successful. Note that

$$d_3((\psi \circ \phi)(v), z) \leq d_3(\psi(\phi(v)), \psi(w)) + d_3(\psi(w), z) \leq A_2 d_2(\phi(v), w) + B_2 + C_2 \leq A_2 C_1 + B_2 + C_2,$$

so we may take $C = A_2 C_1 + B_2 + C_2$ (or larger).

Altogether, we have shown that the composition $\psi \circ \phi$ is an (A_3, B_3, C_3)-quasi-isometry, where

$$(A_3, B_3, C_3) = (A_2 A_1, A_2 B_1 + B_2, A_2 C_1 + B_2 + C_2) = A_2(A_1, B_1, C_1) + B_2(0, 1, 1) + C_2(0, 0, 1). \quad \square$$

Definition 3.5. Let $f : X \to Y$ be a map of metric spaces. We say that another map $g : X \to Y$ has *finite distance* from f if there exists a constant $c \in \mathbb{R}_{>0}$ such that $d_Y(f(x), g(x)) \leq c$ for all $x \in X$.

For maps $f, g : \mathbb{R} \to \mathbb{R}$, g having finite distance from f means that the graph of g lives in a tubular neighbourhood of some fixed radius around the graph of f.

Example 3.6. The map $f : \mathbb{R} \to \mathbb{R}$, $f(x) = \sin(x)$ has finite distance from $\mathrm{id}_{\mathbb{R}}$, as does any bounded function.

Definition 3.7. Let $f : X \to Y$ be a quasi-isometry. We say that a quasi-isometry $g : Y \to X$ is a *quasi-inverse* of f if $g \circ f$ has finite distance from id_X and $f \circ g$ has finite distance from id_Y.

Proposition 3.8 (Existence of quasi-inverses)**.** *Every quasi-isometry has a quasi-inverse.*

Proof. Let $f : X \to Y$ be a quasi-isometry. We want to show that there exists a quasi-inverse $g : Y \to X$. Given $y \in Y$, we must therefore decide a suitable element of X for g to send y to. Now, $f \circ g$ has finite distance from id_Y iff there exists a constant $K > 0$ such that $d_Y(y, f(g(y))) \leq K$ for all $y \in Y$. In other words, $f(g(y))$ is in $B_K(y) \cap f(X)$. This looks strikingly similar to the second condition in the definition of quasi-isometry. Since f is coarsely surjective, we know that there exists a constant $C > 0$ (independent of y) such that $B_C(y) \cap f(X) \neq \emptyset$. Choose $g(y)$ to be x for any $f(x)$ in this non-empty intersection. Then $f \circ g$ has finite distance from id_Y. It remains to show that $g \circ f$ has finite distance from id_X and that g is a quasi-isometry. $\square$

Exercise 3.9. Complete the last two steps of the above proof.

Definition 3.10. Given a group G with generating set S, the *word metric* d_S on G is defined by setting $d_S(g,h)$ to be the minimum length of $g^{-1}h$ as a word in the alphabet $S \cup S^{-1}$.

Theorem 3.11. *Let S and T be two finite generating sets for a group G. Then there exists an equivariant isometry between $\Delta(G,S)$ and $\Delta(G,T)$.*

Proof. One may work directly with the geometric realisations of both Cayley graphs, but a cleaner proof is achieved by first relating the Cayley graph $\Delta(G,S)$ to the metric space (G,d_S) and then passing between (G,d_S) and (G,d_T).

First, consider the map $\iota : \left(\Delta_S, d_{|\Delta_S|}\right) \to (G,d_S)$ given by setting

$$\iota((\{v_g,v_h\},t)) = \begin{cases} g \text{ if } t \leq \frac{1}{2}, \\ h \text{ if } t > \frac{1}{2}. \end{cases}$$

That is, ι sends a point in the realisation to the element of G labelling the nearest vertex. We claim that ι is a quasi-isometry. Indeed, it is easily seen that $d_S(\iota(x),\iota(y))$ differs from $d_{|\Delta_S|}(x,y)$ by at most one, since the word metric on Δ with respect to S coincides with the associated metric on the vertices of $\Delta(G,S)$.

It remains to prove that (G,d_S) is quasi-isometric to (G,d_T). Let us show that id_G gives us a quasi-isometry between these two metric spaces. Let g,h be two elements of G. Recall that $d_S(g,h)$ is defined to be the minimum length of a word representing $g^{-1}h$ in the alphabet $A = S \cup S^{-1}$, and let w be a word achieving this minimum length. Let $S = \{s_1,\dots,s_r\}$, and let l_i be the minimum length of a word in the alphabet $B = T \cup T^{-1}$ representing s_i, and let w_i be a word in B achieving this minimum length. Set $l = \max\{l_i\}$. By replacing s_i by w_i in w, we obtain a word in B representing $g^{-1}h$ of length at most $l d_S(g,h)$. In other words, we have proved that

$$d_T(g,h) \leq l d_S(g,h).$$

By symmetry, we may also obtain a constant l' such that

$$d_S(g,h) \leq l' d_T(g,h).$$

Setting $L = \max\{l,l'\}$, we deduce that id_G is an $(L,0)$-quasi-isometric embedding. Since id_G is surjective, it is coarsely surjective, and we are done. □

The upshot is this: to make a geometric object out of our group, we allowed ourselves to add in more data: that of a generating set. However, we showed that changing the generating set preserves the large-scale geometry of our object (the Cayley graph). Then geometric properties of this graph which are invariant under quasi-isometry can be justifiably thought of as geometric properties of our group. We win!

Chapter 4

Geodesics and actions

In this chapter we will introduce geodesics and spaces in which they are ubiquitous, in the sense that any two points are connected by one. We like working with spaces with nice geometric properties (who doesn't?), and one nice property is that any two points are connected by some sort of curve, or path. It is even nicer when such a path, viewed as a mapping of some closed interval into our metric space, preserves the distances between points, at least up to some scalar. Such curves are called *geodesics*, and we will be interested in working with *geodesic spaces*, in which any two points are connected by a geodesic.

We will then look at group actions and discover how nice actions of groups on geodesic spaces can tell us something about the algebraic structure of the group.

4.1 Paths and geodesics

Definition 4.1. A *path* in a metric space M is a continuous function $\gamma : I \to M$, where $I \subset \mathbb{R}$ is a closed interval. An arc is an embedded path, i.e. a path with γ an injective function.

Example 4.2. Some examples of paths.

1. $\gamma : [-1, 1] \to \mathbb{R}^2$, $\gamma(t) = \big(t, t^3\big)$ is the graph of $y = x^3$ between the points $(-1, -1)$ and $(1, 1)$.

2. $\gamma : [0, \pi] \to \mathbb{R}^2$, $\gamma(t) = (2 + \cos t, 3 + \sin t)$ is the semicircular arc with centre $(2, 3)$ joining the points $(3, 3)$ and $(1, 3)$.

3. $\gamma : [0, 2\pi] \to \mathbb{R}^3$, $\gamma(t) = \big(\cos t, \sin t, \frac{t}{2\pi}\big)$ is a spiral joining $(1, 0, 0)$ to $(1, 0, 1)$ around the unit-radius cylinder with centre the z-axis.

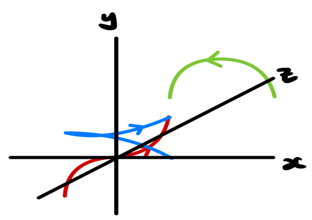

Identifying a path with its image in M, it is easy to see that the name is justified; the image will be a path (in the usual sense) between its endpoints; possibly a nasty path with lots of loops and vertices, but a path nonetheless.

Now that we have a notion of path, it is natural to ask how long the path is. It is quite easy to determine the distance between any two points on the path, and as with integration, we arrive at a reasonable definition of length by partitioning our path up into discrete segments, taking the length of those and effectively taking a limit.

Definition 4.3. Given a path $\gamma : I \to M$, the *length* of γ is

$$\text{length}(\gamma) = \sup\left\{\sum_{i=1}^{n} d(\gamma(t_i), \gamma(t_{i-1})) : a = t_0 < t_1 < \cdots < t_{n-1} < t_n = b\right\}.$$

We say that γ is *rectifiable* if it has finite length.

Definition 4.4. Let (M, d) be a metric space, and let $I \subset \mathbb{R}$ be an interval. A *geodesic* of constant speed $\lambda \geq 0$ is a path $\gamma : I \to M$ such that $d(\gamma(t), \gamma(u)) = \lambda|t - u|$ for all $t, u \in I$. A *unit-speed geodesic* is a geodesic of constant speed 1.

Note that $|t - u|$ is simply the distance between the points $t, u \in I$ with respect to the Euclidean metric $(x, y) \mapsto |x - y|$ on $\mathbb{R}$, hence a unit-speed geodesic is a way of mapping an interval into a metric space in a "distance-preserving" way.

Note 4.5. There is another, perhaps more common, description of geodesics: they are distance-minimising paths. That the above definition captures the fact that the path is "as short as possible" among all paths connecting the endpoints is admittedly not obvious. However, knowing this should allow you to guess what the geodesics in a particular metric space ought to be, and you can then check to see if your suspicion is correct.

Definition 4.6. A metric space (M, d) is a *geodesic space* (also called a *length space*) if any two points are connected by a geodesic.

Example 4.7. $\mathbb{R}^n$ with the Euclidean metric is a length space; straight lines with a linear parametrisation are geodesics.

Example 4.8. A topological space X (e.g. a metric space) is called *path-connected* if any two points $P, Q \in X$ can be connected by a path, i.e. if there exists a path $\gamma : [a, b] \to X$ with $\gamma(a) = P$, $\gamma(b) = Q$. Note that any length space has to be a path space. There is a weaker notion of *connectedness*, which is slightly more involved but means what you think: the space X is made out of "one piece" rather than disjoint pieces. As you might expect, any path-connected space is connected, so connectedness is necessary in order to be a length space. In particular, any non-connected space is not a length space.

Definition 4.9. A metric space is *proper* if it is complete and locally compact.

Example 4.10. $\mathbb{R}^n$ with the Euclidean metric is proper. It is complete (every Cauchy sequence is convergent), and local compactness can be seen via the Heine–Borel theorem, which says that the compact subsets of $\mathbb{R}^n$ are exactly those which are closed and bounded; this includes closed balls, and every point is contained in a closed ball.

Example 4.11. $\mathbb{Q}$ with the Euclidean metric is not proper as it is neither complete nor locally compact; indeed, the completion of $\mathbb{Q}$ with respect to the Euclidean metric is $\mathbb{R}$, and any neighbourhood contains an irrational number.

Definition 4.12. Let M be a proper geodesic metric space, and let $Q \subset M$ be a closed subset. We define the *geodesic metric* d_Q on Q by setting $d_Q(x, y)$ to be the minimum of the lengths of rectifiable paths in M connecting x to y.

Exercise 4.13. Check that the geodesic metric really is a metric.

4.2 Actions

Definition 4.14. Let X be a set and let G be a group acting on X. Given $x \in X$, we define the *orbit* of x to be the subset $Gx = \{gx : g \in G\}$, i.e. the set of all elements of X to which x can be sent by acting on it via some $g \in G$. The *stabiliser* of x is $\text{Stab}(x) = G_x = \{g \in G : gx = x\}$, i.e. the set of $g \in G$ which send x to itself.

Example 4.15. Any group acts on itself by either left or right multiplication.

Example 4.16. Any group acts on itself by conjugation. The centre $Z(G)$ of G, which is mentioned in an earlier example and consists of all elements which commute with every other element, can be described as the intersection $Z(G) = \cap_{g \in G} \text{Stab}(g)$. The orbits of this action are the *conjugacy classes* of G. Conjugacy classes play an important role in the study of non-abelian groups, e.g. the symmetric group S_n.

Definition 4.17. Let X be a proper geodesic space, let G be a group, and let $G \times X \to X$, $(g, x) \mapsto gx$ be an action of G on X. We say that G *acts by isometries* on X if for each $g \in G$, the map $X \to X$, $x \mapsto gx$ is an isometry.

Exercise 4.18. Show that the special orthogonal group $\mathrm{SO}_2(\mathbb{R}) = \{\left(\begin{smallmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{smallmatrix}\right) : \theta \in \mathbb{R}\}$ acts by isometries on $\mathbb{R}^2$ by matrix multiplication.

4.3 PDC actions

Let G be a group acting on a proper geodesic space X.

Definition 4.19. We say that the action of G on X is *properly discontinuous (PD)* if for all $r \geq 0$ and all $x \in X$, the set $\{g \in G : d(x, gx) \leq r\}$ is finite. That is, given a point $x \in X$ and a radius $r \geq 0$, there are only finitely many $g \in G$ which will send x to another point in $B(x, r)$. We say that the action is *properly discontinuous and cocompact (PDC)* if it is PD and the quotient X/G is compact.

According to Wikipedia, "cocompact" is sometimes jokingly written as "mpact". Why is this such a rib-tickler for algebraists? Because we often use "co-" to refer to some dual object, and taking the dual of a dual often gives you back the thing that you started with, so that two co-s cancel out. Taking duals often changes the direction of things (like morphisms, or objects in a statement); with that knowledge, you are now sufficiently equipped to appreciate another comedic gem, courtesy of algebraic geometer Ravi Vakil (Joke 1.4.5 in these notes):

Joke 4.20. A comathematician is a device for turning cotheorems into ffee.

Example 4.21. Any action by a finite group is trivially a PD action. Any PD action which is *transitive* (meaning that one point can be sent to any other by acting by some group element) is trivially cocompact.

Example 4.22. The action of $\mathbb{Z}^2$ on $\mathbb{R}^2$ by addition (i.e. $(a, b) * (x, y) = (x + a, y + b)$) is PDC; proper discontinuity is clear, and $\mathbb{R}^2/\mathbb{Z}^2$ is isomorphic to the *torus* $\mathbb{T}^2$ (to the uninitiated, a ring doughnut), as can be seen by bringing any point into the square $[0, 1]^2$ and considering the way in which edges of this square identify/"glue".

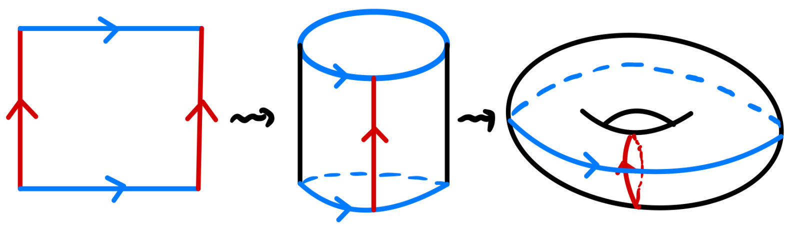

Theorem 4.23. *If a group G has a PDC action on a proper geodesic space, then G is a finitely generated group.*

Proof. Suppose that G acts PDC on X. Fix an element $x \in X$. Since the action is cocompact, every orbit is cobounded, i.e. there exists $C > 0$ such that Gx is C-dense. We construct a graph $K = (V, E)$ associated to this action by taking $V = G$ and connecting $g, h \in G$ iff $d(gx, hx) \leq 2C + 1$.

Claim 1: K is locally finite. Note that $d(gx, hx) = g\big(gx, hg^{-1}(gx)\big)$, so $d(gx, hx) \leq 2C + 1$ iff $hg^{-1} \in \{k \in G : d(gx, kgx) \leq 2C + 1\}$, a finite set.

Claim 2: K is connected. We seek a path in the graph between any two vertices g and h. Such a path corresponds to a sequence of group elements $g_0 = g, g_1, \ldots, g_n = h$ so that $d(g_i x, g_{i+1} x) \leq 2C + 1$ for $i = 1, \ldots, n - 1$. This sequence of associated points in X has the feature that two consecutive points are "not too far apart". Producing one such sequence between $g_0 x = gx$ and $g_n x = hx$ is easy: since

X is a geodesic space, we can simply take a geodesic between them and divide it into arbitrarily small segments. We will combine this observation with the cocompactness of the action to produce the desired path.

Let $\gamma : [0,1] \to X$ be a geodesic connecting gx and hx. Choose points $P_0 = gx, P_1, \ldots, P_n = hx \in X$ along the image of γ such that $d(P_i, P_{i+1}) \leq 1$. Next, choose $g_i \in G$ such that $d(P_i, g_i x) \leq C$. By the triangle rule, we have $d(g_i x, g_{i+1} x) \leq d(g_i x, P_i) + d(P_i, P_{i+1}) + d(P_{i+1}, g_{i+1} x) \leq 2C + 1$.

Claim 3: K is the Cayley graph of G with generating set $S = \{g \in G \setminus \{1_G\} : d(x, gx) \leq 2C+1\}$–at least after gluing edges for 2-torsion elements.

Hopefully this is now at least believable. □

Corollary 4.24 (Milnor–Švarc Lemma)**.** *If a group G has a PDC action on a proper geodesic space X, then $G \sim X$.*

Proof. Let Δ be the Cayley graph constructed in the previous proof, and let $f : \Delta \to X$ be the map given by sending a vertex g to $gx \in X$ and a directed edge $(v_g, v_{g'})$ linearly to a geodesic γ joining gx and $g'x$. By choosing points P_i on γ as in the previous proof, we see that the length of the path that we constructed between g and h is at most $d_X(f(g), f(h)) + 1$, so $d_\Delta(g,h) \leq d_X(f(g), f(h)) + 1$, provided that we choose the P_i to be evenly spaced along γ. On the other hand, $d_X(f(g), f(h)) \leq (2C+1) d_\Delta(g,h)$ (as seen by following the "straight-line" path between the points $g_i x$ in the previous proof), so f is a quasi-isometric embedding. Since $G = V(\Delta)$ and $Gx = f(V(\Delta))$ are cobounded (i.e. R-dense for some $R > 0$) in Δ and X respectively, we deduce that f is also coarsely surjective, hence a quasi-isometry. □

Corollary 4.25. *Let G be a finitely generated group and let $H \leq G$ be a subgroup of finite index. Then H is finitely generated and $G \sim H$.*

Proof. The PDC action of G on its Cayley graph Δ restricts to a PDC action of H on Δ. □

4.4 Commensurability

Definition 4.26. Let G_1 and G_2 be two groups. We say that G_1 and G_2 are *commensurable*, denoted by $G_1 \approx G_2$, if there exist finite-index subgroups $H_1 \leq G_1$ and $H_2 \leq G_2$ such that $H_1 \cong H_2$.

Example 4.27. If G is a finite group and G' is any (possibly infinite) group, then $G \times G'$ is commensurable with G'. For instance, $\mathbb{Z} \times \mathbb{Z}/3\mathbb{Z} \approx \mathbb{Z}$.

Proposition 4.28. *If G_1 and G_2 are finitely generated groups, then $G_1 \approx G_2$ implies $G_1 \sim G_2$.*

Proof. Suppose that $G_1 \approx G_2$. Then there exist finite index subgroups $H_1 \leq G_1$, $H_2 \leq G_2$ such that $H_1 \cong H_2$. We have seen above that $G_i \sim H_i$, so it suffices to check that isomorphic groups are quasi-isometric, which is easily seen. □

Definition 4.29. Let $\mathcal{P}$ be a property of groups. We say that a group G is *virtually* $\mathcal{P}$ if it has a finite-index subgroup $H \leq G$ such that H has property $\mathcal{P}$.

Example 4.30. Any finite group is virtually abelian and virtually trivial.

4.5 Geometric properties

Definition 4.31. We say that a property $\mathcal{P}$ of groups is *geometric* if, given two quasi-isometric groups G and G', G has $\mathcal{P}$ iff G' has $\mathcal{P}$.

Example 4.32. Being finitely presented is a geometric property.

Example 4.33. Recall from earlier (Note 1.24) the *word problem*: given a finitely generated group G and a word in the generators of some finite generating set, is there an algorithm to decide whether the word is equivalent to the empty word? An equivalent question is of course: is there an algorithm to decide when two words represent the same group element? For finitely presented groups, work of Alonso and Shapiro shows that solubility of the word problem is geometric: that is, given two *finitely presented* groups G and G' such that $G \sim G'$, the word problem is solvable for G iff it is solvable for G'. However, the question of whether the word problem is geometric for finitely generated groups remains open.