



SJÄLVSTÄNDIGA ARBETEN I MATEMATIK

MATEMATISKA INSTITUTIONEN, STOCKHOLMS UNIVERSITET

Geometric Group Theory and Hyperbolic Groups

av

Arvid Ehrlén

2019 - No K18

Geometric Group Theory and Hyperbolic Groups

Arvid Ehrlén

Självständigt arbete i matematik 15 högskolepoäng, grundnivå

Handledare: Rikard Bögvad

2019

Geometric Group Theory and Hyperbolic Groups

Arvid Ehrlén

June 7, 2019

Abstract

This thesis aims to give an introduction to the foundations of geometric group theory and to present the notion of a hyperbolic group, first introduced by Mikhail Gromov in 1987. First I describe how finitely generated groups endowed with the word metric can be regarded as geometric objects and what the morphisms between these objects are. I then reproduce some basic results of geometric group theory, most notably the Milnor-Švarc lemma. Lastly, I discuss hyperbolic groups and show that the hyperbolicity property is preserved by the morphisms.

Contents

1	Foreword	3
1.1	Aim of thesis	3
1.2	Background and overview	4
2	Objects: groups as metric spaces	6
2.1	Word length and word metric	7
2.2	A remark on the induced topology	8
2.3	Cayley graphs	8
2.4	Group acting on its Cayley graph	10
2.5	Geodesic metric spaces	11
2.6	Proper and cobounded group actions	15
3	Morphisms: quasi-isometries	16
3.1	Definitions and examples	16
3.2	Quasi-isometric inverses	19
4	Milnor–Švarc lemma	21
4.1	Growth rate	23
5	Hyperbolic groups	25
5.1	Hyperbolicity is invariant under quasi-isometries	26
5.2	Some properties of hyperbolic groups	31
A	Appendix	35
A.1	Metric spaces	35

1 Foreword

1.1 Aim of thesis

The aim of this thesis is to give a brief introduction to some fundamental results of geometric group theory and to state some interesting results about finitely generated groups, in particular hyperbolic groups. The main idea of geometric group theory is to study finitely generated groups as automorphism groups (i.e. symmetry groups) of metric spaces. There are already a number of mathematical structures where groups have appeared as automorphism groups, such as polynomial field extensions (Galois groups) and vector spaces (matrix groups). Another example is partial differential equations and Lie groups.

The central question in geometric group theory is how algebraic structure of a group G reflect in geometric properties of a metric space (X, d) that G acts on, and conversely how geometric properties of (X, d) reflect on the algebraic structure of G .

The first chapter on *groups as metric spaces* is dedicated to determine this correspondence between finitely generated groups and metric spaces. It introduces the word metric and the Cayley graph as well as the natural action of a group on its Cayley graph and shows how a Cayley graph can be made into a geodesic metric space. Cayley graphs regarded as geodesic metric spaces are the basic objects of interest in geometric group theory.

Since we are required to choose a finite generating set for a group in order to construct a Cayley graph representation for it, one might wonder to what extent this choice determines the graph. This question is answered in the next chapter on *quasi-isometries*. An isometry is a map between two metric spaces that preserves distances. A quasi-isometry is also a map between metric spaces, but here distances are allowed to be distorted, although not “too much”. Quasi-isometries are in a sense a weaker form of equivalence between metric spaces than isometries. It turns out that while different choices of generating sets yield different Cayley graph representations for a given group, there is still some properties that are not affected by this choice. Given a finitely generated group G , all its possible Cayley graphs (corresponding to different choices of a generating set for G) are quasi-isometric to each other. Quasi-isometries are the morphisms (structure preserving maps) between the objects of geometric group theory, and it is through this quasi-isometric lens we are able to talk about “the” Cayley graph of a group.

Once the groundwork has been laid out in the first two chapters we are able to put together these concepts in order to state what is sometimes called the fundamental observation of geometric group theory: the Milnor-Švarc lemma (after John Milnor and Albert S. Švarc). The lemma not only provides a tool for determining whether a group is finitely generated or not, but it also tells us that whenever a group is acting “nicely” on a geodesic

metric space, then the Cayley graph of the group has the same so called “coarse geometry” (large-scale structure) as the space being acted on. In this chapter we also introduce the notion of *growth* of a group, another geometric concept that John Milnor was interested in and served as motivation to prove the lemma.

In the last part of the thesis the notion of *hyperbolic groups* are introduced and we show that the property of hyperbolicity of a group is invariant under quasi-isometries. To show this invariance after presenting the theory required is one of the main goals of this thesis.

1.2 Background and overview

A common way groups are first introduced to students in mathematics is as the algebraic concept of a set with some algebraic operations that satisfy a few axioms. When studying a mathematical concept, however, it is often useful if one can view the concept from a few different perspectives at the same time. For example, a group can be thought of as all the ways in which one can transform a space into itself while preserving some object or structure in the space. A group can also be thought of as the set of homotopy classes of continuous loops in a connected topological space, called the fundamental group. In some cases, additional structure for the group is provided, allowing it to be studied not only as an abstract algebraic object. A Lie group is an example in which a group is given the structure of a manifold.

In geometric group theory, algebraic properties of groups are studied not only through how the groups act on various spaces, but also by thinking of groups themselves as geometric objects. The fundamental way this is done is by specifying a finite set of generators for a group. For a pair (G, S) of a group with a finite set of generators for G , one can define a metric on G called the *word metric*. The group together with the word metric can then be considered as a metric space.

Yet another way of thinking about a group G is to specify a *presentation* of it—that is, by specifying a set S of generators and a set R of relations among the generators, commonly denoted $G = \langle S \mid R \rangle$. The fact that a group has a unique presentation, in the sense that if another group admits the same presentation then they are isomorphic, was the fundamental building block for the area of combinatorial group theory that was first studied in the mid to late 1800s. The subject of geometric group theory as a distinct mathematical theory emerged from the combinatorial group theory during the late 1980s, largely credited to the works of Mikhail Gromov (e.g. [6], [8]). Gromov’s work sparked an interest among other mathematicians to continue where he left off, and we know today that a lot of algebraic properties of groups can be learned by studying the geometry of groups regarded as metric spaces.

The proof of the classification theorem for finite simple groups was com-

pleted during the early 1980s (with the only exception of *quasi-thin groups* whose classification was not proven until 2004), spanning approximately 10,000 pages by about 100 authors and had taken over 30 years to complete. Simple groups are the building blocks of all finite groups, a bit similar to how prime numbers are the building blocks (factors) for the natural numbers. Geometric group theory is in a natural way a continuation of this project, since it is concerned with studying infinite, but finitely generated, groups.

The Cayley graph of a group G with respect to a generating set S for G is a graph that captures the algebraic structure of G in the following sense. Each vertex of the Cayley graph corresponds to an element of the group and there is a directed edge from $g \in G$ to $h \in G$ if and only if there exists a generator $s \in S$ such that $gs = h$.

In order to construct a Cayley graph for a given group, it is necessary to specify a set of generators for the group. It is, however, desirable for any theory developed to as far as possible be independent on this choice of generating set. The goal is to be able to talk about geometric properties inherent to a group G instead of considering pairs (G, S) of a group and a set S of generators for G . The key observation that makes this independence possible is that while different choices of generating set gives rise to different word metrics and Cayley graphs, there are certain “large-scale” properties that do not depend on this choice. To give this some intuition, consider the following example. The Cayley graph of the group of integers \mathbb{Z} with standard generating set $\{\pm 1\}$ is an infinite path. If we change the generating set to, say, $\{\pm 2, \pm 3\}$, the Cayley graph looks a bit like a braid.




Figure 1: Cayley graph of \mathbb{Z} with generating set $\{\pm 1\}$.




Figure 2: Cayley graph of \mathbb{Z} with generating set $\{\pm 2 \pm 3\}$.

But if we imagine these two graphs seen from far away they start to look similar to each other, both looking like infinite lines or paths. In other words, the *coarse geometry* of the graphs are, in some sense, equivalent. This equivalence is known as a quasi-isometry and is a central concept to geometric group theory.

As a counterexample, we may consider the Cayley graph of the finite

cyclic group $\mathbb{Z}/n\mathbb{Z}$ for some $n \geq 3$ with the standard generating set $\{\pm 1\}$, whose Cayley graph looks like an n -cycle or an n -gon. As we zoom out, the graph looks more and more like a single point, and is not equivalent (quasi-isometric) to either graphs in Figure 1 and 2.

The notion of a group being *hyperbolic* is a concept that only makes sense under the quasi-isometric equivalence relation. The property of a group being hyperbolic is an inherently geometric property and it is defined through Cayley graph representations of groups. It is therefore necessary that all possible Cayley graphs corresponding to some group G are either hyperbolic or not hyperbolic at the same time. To be precise, we require that the property of hyperbolicity for a group is invariant under quasi-isometries.

2 Objects: groups as metric spaces

We begin by demonstrating how a notion of distance for a group G is obtained by fixing a generating set S for G and defining the word metric d_S on G with respect to S . Since every element of G can be expressed as a product of a number of elements of the generating set S , we can define a *norm* or *word length* of an element $g \in G$ with respect to S as the least number of elements of S required to express g as a product. From this we can construct the word metric d_S of G with respect to S that takes two elements $g, h \in G$ and outputs the word length of their difference $g^{-1}h$.

The group G together with the word metric d_S may then be considered as a metric space (G, d_S) , and we can visualize this metric space as a Cayley graph representation of G with respect to S , as long as S is finite. Since the word metric only takes values in \mathbb{N} , the metric space considered will be different in an essential way from more standard Euclidean, hyperbolic or elliptic metric spaces and manifolds where the corresponding metrics usually takes values in $\mathbb{R}_{\geq 0}$. However, one may endow the edges of a Cayley graph with a metric (sub)structure by identifying them with copies of the unit interval $[0, 1] \subset \mathbb{R}$, thus making it possible to consider distances not only between vertices of the graph but between points lying on edges as well. Hence one extends the word metric d_S to a graph metric d_Γ taking values in $\mathbb{R}_{\geq 0}$.

We shall also see that for a finitely generated group G , there is a natural extension of the canonical action of G on itself by left or right translation (or multiplication) to an action of G on its Cayley graph. This natural action can be shown to satisfy some properties such as being a *proper* and *cobounded* action. Using the “large-scale” comparison (which is made precise in the next chapter) of the Cayley graph to other metric spaces one is able to deduce properties of G by looking at the behaviour of the group action on various metric spaces.

2.1 Word length and word metric

For the definitions of word length and word metric below, we assume S is a generating set for a group G and that S is closed under taking inverses, or *symmetric*. The definitions can also be made when S is not symmetric, by simply extending S to the set $S \cup S^{-1}$, where $S^{-1} = \{s^{-1} \mid s \in S\}$, and then writing $S \cup S^{-1}$ in place of S .

Definition 1. Let G be a group and let S be a symmetric generating set for G . A **word** in S is a finite sequence $s_1 s_2 \dots s_n$ where $s_1, \dots, s_n \in S$. The number n is called the **length** of the word. By evaluating the word in G using the group operation to multiply the s_i in order, the result is an element $g \in G$, called the **evaluation** of the word $s_1 s_2 \dots s_n$. By convention, the evaluation of the empty word is the identity element of G .

Definition 2. Let G be a group and let S be a symmetric generating set for G . The **word length** (or word norm) of an element $g \in G$ with respect to S , denoted $\ell_S(g)$, is the shortest length of a word in S whose evaluation is equal to g .

We have the following elementary properties of the word length.

1. $\forall g \in G : \ell_S(g) = \ell_S(g^{-1})$,
2. $\forall g, h \in G : \ell_S(gh) \leq \ell_S(g) + \ell_S(h)$.

The first property follows from the fact that any word $s_1 \dots s_n$ representing an element $g \in G$ corresponds to a word $s_n^{-1} \dots s_1^{-1}$ of equal length representing g^{-1} and vice versa. For the second property (subadditivity), observe that any word representing gh can be split into two words representing g and h respectively. Hence the shortest word representing gh cannot be larger than the sum of the word lengths of g and h .

Definition 3. Let G be a group and S a symmetric generating set. The function

$$\begin{aligned} d_S : G \times G &\longrightarrow \mathbb{N} \\ (g, h) &\longmapsto \ell_S(g^{-1}h) \end{aligned}$$

is called the **word metric** on G with respect to S . Equivalently, $d_S(g, h)$ is the shortest length of a word $s_1 s_2 \dots s_n$ in S such that $g \cdot s_1 \cdot \dots \cdot s_n = h$.

It is important to note that the word length and the word metric can vary greatly depending on the choice of generating set S . If we take $S = G$ then d_S becomes the discrete metric, i.e. $d_S(g, h) = 1$ whenever $g \neq h$.

We shall now verify that the word metric satisfies the axioms for a metric (A.1 Definition 23).

Lemma 1. *The word metric d_S of a group G with generating set S is a metric on the set G , making (G, d_S) into a metric space.*

Proof. By definition $d_S(g, h) \in \mathbb{N}$ for all $g, h \in G$, and in particular $d_S(g, h) \geq 0$. Furthermore $d_S(g, h) = 0$ if and only if $g^{-1}h$ is represented by the empty word, but the empty word represents the identity element e of G , so $d_S(g, h) = 0$ if and only if $g = h$.

The fact that $d_S(g, h) = d_S(h, g)$ for all $g, h \in G$ follows from the fact that $\ell_S(g) = \ell_S(g^{-1})$ for all $g \in G$.

Lastly, the triangle inequality follows from the subadditivity of the word length. Given a word of minimum length representing $g^{-1}h$ and one representing $h^{-1}k$, we can concatenate the words to get a word (not necessarily of minimum length) representing $g^{-1}hh^{-1}k = g^{-1}k$. Hence $d_S(g, k) \leq d_S(g, h) + d_S(h, k)$. \square

2.2 A remark on the induced topology

Let (X, d) be a metric space (A.1 Definition 24) and let $p \in X$ be any point. We define the **open ball** of radius r centered at p to be the set

$$B_r(p) = \{x \in X \mid d(p, x) < r\}.$$

Consider any subset $\mathcal{U} \subset X$. We define \mathcal{U} to be an *open set*, with respect to d , if for every point $p \in \mathcal{U}$ there exists $\epsilon > 0$ such that the ball centered at x with radius ϵ is contained in \mathcal{U} . More precisely,

$$\mathcal{U} \text{ is open} \iff \forall p \in \mathcal{U} \exists \epsilon > 0 : B_\epsilon(p) = \{x \in X \mid d(x, p) < \epsilon\} \subset \mathcal{U}.$$

This collection of open sets is called the *induced topology* or the topology generated by d . An important remark to make here is that a metric space is not a topological space itself, but it does naturally give rise to one via the metric. It is possible that two different metrics induce the same topology. For example, take the real numbers \mathbb{R} and let $d_1(x, y) = |x - y|$ and $d_2(x, y) = 2|x - y|$.

2.3 Cayley graphs

Representing groups with Cayley graphs is a widely used tool in group theory because it provides a way of visualizing the abstract information of a group by encoding it in a graph structure. In geometric group theory, in addition to giving a group a graph structure the Cayley graph can be regarded as a metric space. Intuitively, one constructs a metric space from a given Cayley graph by associating each edge with the unit interval $[0, 1]$ ¹ and each vertex

¹While it is certainly possible to choose other numerical values (weights) for each edge, we will restrict ourselves to the case where all edges are identified with $[0, 1]$ for simplicity.

with a point. The collection of intervals and points then becomes a topological space. Assuming that the original Cayley graph is connected, one can define the distance between two points as the infimum of lengths of paths joining the two points.

An *undirected* graph where edges are associated with intervals and vertices with points is called a **topological graph**. Formally, topological graphs are *1-dimensional CW complexes*. A CW-complex is a type of topological space. See [13] Section 1.A for a formal definition of a topological graph.

A **metric graph** is, roughly speaking, a metric space obtained by taking a connected topological graph and defining the distance between two points as the infimum of the lengths of paths joining them. See [5] Chapter I.1, 1.9 for a formal construction of a metric graph.

Definition 4. Let G be a group and S be a subset of G . The **Cayley graph** of G , denoted $\Gamma(G, S)$, is a graph whose vertices are in one-to-one correspondence with the elements of G , and two vertices $g, h \in G$ are joined by a directed edge from g to h if and only if there exists $s \in S$ such that $gs = h$. Each edge is labelled (or coloured) to denote the element $s \in S$ it corresponds to.

If the subset S generates G , then the labelled Cayley graph $\Gamma(G, S)$ uniquely determines G , meaning we can recover all information about the group by only looking at the graph. However, the labelling is necessary for this to be true, given that there is more than one generator. Consider as an example the group $\mathbb{Z}/4\mathbb{Z}$ with generating set $\{1, -1\}$ (using additive notation) and the group $\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$ with generating set $\{(1, 0), (0, 1)\}$. These two groups are not isomorphic, but their Cayley graphs become isomorphic if the labelling is removed.




Figure 3: Cayley graphs $\Gamma(\mathbb{Z}/4\mathbb{Z}, \{1, -1\})$ (left) and $\Gamma(\mathbb{Z}/2 \times \mathbb{Z}/2\mathbb{Z}, \{(1, 0), (0, 1)\})$ (right).

In geometric group theory one usually makes the following assumptions on the subset $S \subset G$.

- The set S generates G , making $\Gamma(G, S)$ a connected graph. This assumption is commonly made and sometimes included in the definition

of a Cayley graph.

- In most cases, S is taken to be symmetric (meaning $S = S^{-1}$). This allows one to consider $\Gamma(G, S)$ as an undirected graph.
- The identity element of G is not in S . Hence $\Gamma(G, S)$ does not contain any loops (since any loop satisfies $g = gs$, meaning $s = 1$).

Note that the word metric on G (Definition 3) with respect to S corresponds to the number of edges in a shortest path between two vertices of the Cayley graph $\Gamma(G, S)$.

The edges of the Cayley graph are identified isometrically (distance-preserving, cf. Definition 13) with copies of the unit interval $[0, 1]$ so that each edge has a length of 1. This might seem superfluous at first, but once we get to *geodesic segments* and *geodesic metric spaces*, this identification will simplify some of the results. For this reason, when we talk about Cayley graphs from this point on, this extra structure is included. The metric d_Γ on $\Gamma(G, S)$, which will be precisely defined in Section 2.5, will then be distinct from the word metric d_S , but coincides when the points are vertices.

We have seen that a Cayley graph can be viewed in a few different ways. The *combinatorial* view is that of Definition 4, where a graph is a pair $G = (V, E)$ of a set V whose elements are called vertices and a set E of pairs of (distinct) vertices whose elements are called edges.

The *topological* view is to think of the graph as a one-dimensional CW-complex. This can be useful if one wants to consider certain topological properties, such as connectedness. A topological graph is connected if and only if the associated combinatorial graph it was constructed from is connected.

The third way of viewing Cayley graphs is to view it as a *metric graph*, thus regarding it as a type of metric space. This is going to be the main focus from this point on.

2.4 Group acting on its Cayley graph

An action of a group G on a set X is, formally, a homomorphism from G to $\text{Sym}(X)$, the group of symmetries on X i.e. the set of all bijections from X to itself. This means that for each $g \in G$, the group action map $x \mapsto g \cdot x$ is a bijection from X to itself. If we want to consider group actions on objects with additional structure e.g. a metric space (X, d) , it makes sense to consider group actions that preserve the structure of the space i.e. where the mapping $x \mapsto g \cdot x$ is an isometry (see Definition 13) from the metric space to itself. In this case we say that G is **acting by isometries** on the metric space.

For any group G , there is a canonical action of G on itself by left multiplication. One can extend this action in a natural way to an action of G on

its Cayley graph $\Gamma(G, S)$: Suppose S is a finite generating set for G . Since G already acts on the set of vertices by left multiplication, we can extend this action by saying that for any $k \in G$, then if there is an edge (g, gs) between $g \in G$ and $gs \in G$, we map that edge to an edge (kg, kgs) connecting kg and kgs , preserving the set of edges. For points x lying on an edge (g, gs) we define $k \cdot x$ as the point satisfying $d_\Gamma(kg, k \cdot x) = d_\Gamma(g, x)$, meaning each edge gets mapped isometrically. Thus G acts by isometries on its Cayley graph $\Gamma(G, S)$. Here d_Γ is the metric on $\Gamma(G, S)$ which will be formally defined in the next section.

2.5 Geodesic metric spaces

Throughout this subsection, let (X, d) be a metric space.

Definition 5. A **path** in X is a continuous map $\gamma : [a, b] \rightarrow X$ where $[a, b]$ is an interval (connected subset) of \mathbb{R} .

Note that a path is a continuous map and not the image of such a map. Different paths could have the same image in X .

If we have two paths such that the set of points in their images are the same, and the points are visited in the same order, we want to consider the paths to be equivalent.

Definition 6. Let $\gamma_1 : [a, b] \rightarrow X$ and $\gamma_2 : [c, d] \rightarrow X$ be two paths. We say that γ_1 and γ_2 are **equivalent** if there exists a non-decreasing, continuous bijection $\phi : [a, b] \rightarrow [c, d]$ such that $\gamma_1 = \gamma_2 \circ \phi$.

Paths that are equivalent to each other are called **parametrizations** or **re-parametrizations** of one another. Note that this implies that any path $\gamma : [a, b] \rightarrow X$ can be re-parametrized to any other closed real interval $[c, d]$ with the mapping $\phi : [a, b] \rightarrow [c, d]$ given by $\phi(x) = c + \frac{(d-c)}{(b-a)}(x-a)$. In particular we may choose $[c, d]$ to be the unit interval $[0, 1]$ and $\phi(x) = \frac{x-a}{b-a}$. Since every path can be reparametrized to the unit interval in this manner we could have equivalently defined paths to be continuous maps $\gamma : [0, 1] \rightarrow X$ under this equivalence. We shall see that equivalent paths have the same images and the same length.

If two paths $\gamma_1 : [a, b] \rightarrow X$ and $\gamma_2 : [c, d] \rightarrow X$ are equivalent, i.e. $\gamma_1 = \gamma_2 \circ \phi$ for some ϕ as in the definition above, then the images of γ_1 and γ_2 are the same. Indeed, we know that the inverse $\phi^{-1} : [c, d] \rightarrow [a, b]$ exists and is a bijection, so that $\gamma_1 = \gamma_2 \circ \phi \iff \gamma_1 \circ \phi^{-1} = \gamma_2$ and hence

$$\begin{aligned} \text{Im } \gamma_1 &= \{\gamma_1(t) \mid t \in [a, b]\} = \{\gamma_1(\phi^{-1}(t)) \mid t \in [c, d]\} \\ &= \{\gamma_2(t) \mid t \in [c, d]\} \\ &= \text{Im } \gamma_2. \end{aligned}$$

Definition 7. Let $\gamma_1 : [a, b] \rightarrow X$ and $\gamma_2 : [c, d] \rightarrow X$ be paths such that $\gamma_1(b) = \gamma_2(c)$. Their **concatenation** is a path $\gamma_1 * \gamma_2 : [a, b + d - c] \rightarrow X$ given by

$$(\gamma_1 * \gamma_2)(t) = \begin{cases} \gamma_1(t) & t \in [a, b], \\ \gamma_2(t + c - b) & t \in [b, b + d - c]. \end{cases}$$

That $\gamma_1 * \gamma_2$ is continuous, and hence a path, can be shown using the *gluing* or *pasting* lemma.

Lemma 2 (Gluing lemma). *Let A and B be topological spaces and suppose $A = X \cup Y$ where X and Y are either both open or both closed. Suppose further that $f : A \rightarrow B$ is continuous when restricted to both X and Y . Then f is continuous.*

Proof. Recall that f is continuous if the preimage of every open set in B is open in A or equivalently if the preimage of every closed set is closed. Suppose $C \subset B$ is a closed subset, and suppose X and Y are closed. Then $f^{-1}(C) \cap X$ and $f^{-1}(C) \cap Y$ are both closed, since they are the preimages of f when restricted to X and Y , respectively, and by assumption f is continuous on the restrictions. Their union $(f^{-1}(C) \cap X) \cup (f^{-1}(C) \cap Y) = f^{-1}(C) \cap (X \cup Y) = f^{-1}(C)$ is also closed, since it is a union of finitely many closed sets. Hence f is continuous. A similar argument is used when X and Y are both open. \square

For the concatenation $\gamma_1 * \gamma_2$ of two paths, we know that it is continuous for $t \in [0, 1/2]$ and $t \in [1/2, 1]$ by the continuity of γ_1 and γ_2 , so it is continuous on the union $[0, 1]$ by the gluing lemma.

Definition 8. Let γ be a path in X . The **length** of γ , denoted $L(\gamma)$, is defined by

$$L(\gamma) = \sup_{0=t_0 \leq \dots \leq t_n=1} \sum d(\gamma(t_i), \gamma(t_{i+1})),$$

where the supremum is taken over all partitions of $[0, 1]$ and all $n \in \mathbb{N}$.

If $L(\gamma)$ is finite, the path γ is said to be **rectifiable**.

Note that two equivalent paths $\gamma_1 : [a, b] \rightarrow X$ and $\gamma_2 : [c, d] \rightarrow X$ satisfies $L(\gamma_1) = L(\gamma_2)$. This is so because partitions of $[a, b]$ correspond bijectively to partitions of $[c, d] = [\varphi(a), \varphi(b)]$ by continuity and the non-decreasing property of ϕ from Definition 6.

From the above definition, we have the following two properties. For any path γ , we have $L(\gamma) \geq d(\gamma(a), \gamma(b))$, and for any two paths γ, η , the length of the concatenation of these two paths is equal to the sum of their individual lengths. The former of these two properties follows from the fact that since $a < b$ is a partition of $[a, b]$, we have $L(\gamma) \geq d(\gamma(a), \gamma(b))$ since we take the supremum over all partitions.

Let $\gamma * \eta$ be the concatenation of the paths γ and η , where the endpoint of γ is the starting point of η . We want to show that $L(\gamma * \eta) = L(\gamma) + L(\eta)$. Given a chain of points for γ and one for η , we can concatenate them (as chains) and get a chain of points for $\gamma * \eta$. Hence the supremum $L(\gamma * \eta)$ cannot be less than $L(\gamma) + L(\eta)$. Conversely, given a chain of points for $\gamma * \eta$, we can simply add one point if needed (the common point of γ and η) to make it into a chain for γ and η . So $L(\gamma * \eta) \leq L(\gamma) + L(\eta)$. Thus $L(\gamma * \eta) = L(\gamma) + L(\eta)$.

Given a rectifiable path $\gamma : [a, b] \rightarrow X$, we may consider the map $\mathcal{L} : [a, b] \rightarrow [0, L(\gamma)]$ given by $\mathcal{L}(t) = L(\gamma|_{[a, t]})$ for $t \in [a, b]$ (note that if γ is rectifiable, so is the restriction to any subinterval). The map \mathcal{L} is continuous and weakly monotonic (for a proof, see [5], chapter I.1 proposition 1.20).

Definition 9. A path $\gamma : [0, \ell] \rightarrow X$ is said to be of **unit speed**, or called **natural**, if

$$L(\gamma|_{[t_1, t_2]}) = t_2 - t_1$$

for every subinterval $[t_1, t_2] \subset [0, \ell]$. Such paths γ are also said to be **parametrized by arc length**. Here ℓ is the length of γ .

A fairly simple but nonetheless important fact is that every rectifiable path can be re-parametrized to the natural parametrization ([5] prop. 1.20).

Definition 10. Let (X, d) be a metric space. A **geodesic segment** from $x \in X$ to $y \in X$ is the image of an isometric embedding $\sigma : [0, \ell] \rightarrow X$ such that $\sigma(0) = x$ and $\sigma(\ell) = y$. Explicitly, we have $d(\sigma(t_1), \sigma(t_2)) = t_2 - t_1$ for any $t_1, t_2 \in [0, \ell]$ (with $t_1 \leq t_2$). In particular $\ell = d(x, y)$.

We will use the notation $[x, y]$ for a geodesic segment between x and y .

For a general metric space (X, d) the existence of geodesics between points is not guaranteed. A simple example of this is any *discrete* metric space, i.e. where the metric is given by $d(x, y) = 1$ for all distinct points $x, y \in X$, and $d(x, y) = 0$ if $x = y$. If there exists a geodesic between any two points of a metric space, then we say that the metric space is **geodesic**. Note that if a geodesic exists, it does not need to be unique. An example of where this is the case is a space with the ℓ_1 -metric, also known as the taxicab metric. In \mathbb{R}^2 the ℓ_1 -metric is given by $d((x_1, y_1), (x_2, y_2)) = |x_1 - x_2| + |y_1 - y_2|$.

We shall see that $\Gamma(G, S)$ is a geodesic metric space, but we must first formally define the metric d_Γ . If each edge of $\Gamma(G, S)$ is identified with $[0, 1]$, meaning each edge has a length 1 associated with it, then there is a natural way of defining the length of a path consisting of finitely many subpaths of edges. We can then take the distance between two points of Γ to be the infimum of the lengths of paths as above connecting the points. Formally, we proceed as follows.

For each edge e , let ϕ_e be a homeomorphism (i.e. a continuous bijection such that the inverse function ϕ_e^{-1} is continuous) from e to $[0, 1]$. We regard e

as an isometric copy of $[0, L(e)]$ where $L(e)$ is the length of e , usually taken to be 1. Define the function $\rho : X \times X \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ in the following way. If two points x_1, x_2 belong to the same edge e we set $\rho(x_1, x_2) = L(e)|\phi_e(x) - \phi_e(y)|$ where $L(e)$ is the length of e and $\rho(x, y) = \infty$ otherwise. Then we define d_Γ as

$$d_\Gamma(x, y) = \inf_{x=x_0, \dots, x_n=y} \sum \rho(x_i, x_{i+1}),$$

where the infimum is taken over all chains from x to y . In some cases one might want to consider edges of other lengths (or weights), but we will for convenience restrict ourselves to edges of length 1.²

In this definition one can equivalently take chains $x = x_0, \dots, x_n = y$ with the restriction that all the x_i are vertices for $i = 1, \dots, n - 1$. To see this we reason as follows. If $\sum \rho(x_i, x_{i+1})$ is finite then for any x_i that is not a vertex, both x_{i+1} and x_{i-1} must be contained in the same edge as x_i . If we remove x_i from the chain, then the sum will not increase because of the triangle inequality: $\rho(x_{i-1}, x_{i+1}) \leq \rho(x_{i-1}, x_i) + \rho(x_i, x_{i+1})$. Hence, given any chain, we can iteratively remove any non-vertices of the "interior" of the chain to get a new chain satisfying the extra requirement that x_i are vertices for $i = 1, \dots, n - 1$, and with $\sum \rho(y_i, y_{i+1}) \leq \sum \rho(x_i, x_{i+1})$. Therefore the infimum taken over the smaller set of chains coincides with the infimum taken over all chains from x to y .

Lemma 3. *The Cayley graph $\Gamma(G, S)$ with the metric d_Γ is a geodesic metric space.*

Proof. Let $x, y \in \Gamma(G, S)$. If x and y are elements of G , it is clear that since $d_\Gamma(x, y)$ is the length of a shortest word in S representing $x^{-1}y$, we can for such a word $s_1 \dots s_n$ construct a geodesic segment of length $n = d_\Gamma(x, y)$ by concatenating the geodesic segments from x to s_1 , from s_1 to s_2 , and so on, keeping in mind that each edge is an isometric embedding of $[0, 1]$ and thus a geodesic segment. When x and y lie on a common edge, then since each edge is an isometric embedding of $[0, 1]$ we can simply take the restriction of this embedding to $[x, y]$ which is clearly also an isometric embedding. For other cases, the following formula holds for all points x, y that do not belong to a common edge.

$$d_\Gamma(x, y) = \inf\{d_\Gamma(x, g) + d_\Gamma(g, h) + d_\Gamma(h, y) \mid d_\Gamma(x, g) < 1, d_\Gamma(h, y) < 1\}.$$

In other words, in order to go from x to y one first has to go to a vertex g being an endpoint of the edge containing x , then go to some other vertex h , and finally go to y staying on an edge that has h as an endpoint. Note that there

²Remark: The way d_Γ is defined makes it possible to have $d_\Gamma(x, y) = 0$ for distinct x and y if lengths of edges are allowed to be 0. Thus d_Γ would fail to be a metric in this case. However, if there is a lower bound on the length of the edges then d_Γ is indeed a metric.

are at most two possible vertices g and two h that satisfy this requirement (since an edge has two endpoints). Thus the infimum is a minimum, and for such g and h , the concatenation of geodesic segments from x to g , g to h and h to y gives a geodesic segment from x to y . \square

Another property of $\Gamma(G, S)$ regarded as a metric space is that its closed balls are compact, i.e. $\Gamma(G, S)$ is a **proper** metric space. Since S is finite, any closed ball of radius $r > 0$ contains at most a finite number of edges (the vertices are contained in the endpoints of edges). Since each edge is a copy of the compact interval $[0, 1]$, any finite union of edges is compact. Any edge that is only partially contained in a closed ball of radius r will give rise to a closed subset of such an edge, and closed subsets of compact sets are again compact.

2.6 Proper and cobounded group actions

So far we have shown that every finitely generated group G acts by isometries on a proper geodesic metric space, an example being $\Gamma(G, S)$. However, it is not enough to require that the space being acted on has nice properties such as being geodesic, but one also wants the action itself to be nice. For example, the trivial action of an arbitrary group (possibly infinitely generated) on a metric space X consisting of only one point is an example of such an action that is not very interesting. In this section we state some further results of the action $G \curvearrowright \Gamma(G, S)$.

Definition 11. An action of a group G on a metric space X is called **proper** if for any $x \in X$ and any ball $B \subset X$ there are only finitely many elements of G that map x into B .

It is not too difficult to see that the action $G \curvearrowright \Gamma(G, S)$ is proper. This is because the orbit of any vertex of the Cayley graph is identified with G itself, and since any ball contains only a finite number of vertices (elements of G) since G is finitely generated. This property asserts that the points of an orbit space are in some sense well-spaced.

If a group G acts properly on X , then stabilizers G_x of points x are finite, since they in particular consist of elements that map points $x \in B \subset X$ to itself.

Definition 12. An action of a group G on a metric space X is said to be **cobounded** if there is a ball $B \subset X$ such that $G \cdot B = X$, where $G \cdot B = \{g \cdot B \mid g \in G\}$.

Another way of formulating the preceding definition is this: There exists a point $x \in X$ and a positive number r such that any point in the space X is within distance r of some point of the orbit Gx of x . Thus a cobounded action says that points of orbits are "pretty much" everywhere.

Consider balls $B \subset \Gamma(G, S)$ of radius 1. Then we see that since the orbit space of any vertex of the Cayley graph in $G \curvearrowright \Gamma(G, S)$ is all of G , this action is cobounded.

We sum up our facts so far with the following theorem.

Theorem 1. *Every finitely generated group G acts properly and coboundedly by isometries on a proper geodesic metric space.*

An example of such an action is the action of G on a Cayley graph $\Gamma(G, S)$ of G .

3 Morphisms: quasi-isometries

Because the Cayley graph is dependent on fixing a generating set for a group, one can ask to what extent the generating set determines the graph of a given group. To answer this question the notion of a *quasi-isometry* is needed.

3.1 Definitions and examples

Definition 13 (Isometry). Let (X, d_X) and (Y, d_Y) be metric spaces. A map $f : X \rightarrow Y$ such that

$$d_Y(f(x_1), f(x_2)) = d_X(x_1, x_2)$$

for all $x_1, x_2 \in X$ is called an **isometric embedding** of the space X into Y . Note that f is injective and continuous by this condition. If f is also surjective, then f is called an **isometry**. In this case we say that the metric spaces are **isometric**.

The notion of isometries between metric spaces is analogous to the notion of isomorphisms for groups, rings, modules, and so on, in the sense that they preserve the structure of the objects.

A quasi-isometry is, roughly speaking, a map that distorts distances only by some fixed affine function (i.e. a linear function plus some constant), and which is surjective up to a bounded constant. The idea is to preserve the coarse structure of a metric space while ignoring smaller details. For example, if we consider the integer line \mathbb{Z} seen from far away, its points seem so close that it becomes hard to distinguish it from the real line \mathbb{R} . Furthermore any real number lies at distance at most $1/2$ from some integer, so that the embedding $\mathbb{Z} \hookrightarrow \mathbb{R}$ is "coarsely" surjective. A quasi-isometry can be thought of as a weaker equivalence than an isometry in the sense that an isometry does not distort distances at all.

Definition 14 (Quasi-isometry). Let (X, d_X) and (Y, d_Y) be metric spaces, and let $k \geq 1$ and $C \geq 0$ be real numbers. A map $f : X \rightarrow Y$ such that

- (i) $\frac{1}{k}d_X(x_1, x_2) - C \leq d_Y(f(x_1), f(x_2)) \leq k \cdot d_X(x_1, x_2) + C$, and
- (ii) the C -neighborhood of $f(X)$ is all of Y (i.e. for any $y \in Y$ there exists an $x \in X$ such that $d_Y(f(x), y) \leq C$)

is called a (k, C) -**quasi-isometry**. If such a map exists, the metric spaces are said to be quasi-isometric. A map that satisfies condition (i) for some k, C is called a **quasi-isometric embedding**.

Observe that given a (k, C) -quasi-isometric embedding f , there is no information at all on scales below the constant C . In particular, f does not need to be continuous. In such a context it is impossible to measure whether the image of some point $f(x)$ really coincides with a point y in the target space, or if $f(x)$ is just some point within distance C to y . A $(k, 0)$ -quasi-isometric embedding f is also called **bi-Lipschitz** (the 'bi'-part comes from the fact that the same k is used in both inequalities, cf. A.1 Definition 28).

Example 1. The metric spaces (\mathbb{Z}, d) and (\mathbb{R}, d) with the usual metric $d(x, y) = |x - y|$ are quasi-isometric with the natural embedding $\mathbb{Z} \hookrightarrow \mathbb{R}$ and $k = 1, C = 1/2$.

Example 2. The map $x \mapsto x^2$ from \mathbb{R} to \mathbb{R} is *not* a quasi-isometric embedding, since we cannot choose k, C such that $|x^2 - y^2| \leq k|x - y| + C$ holds for all $x, y \in \mathbb{R}$.

Example 3. In general, the inclusion of a subspace Y into a metric space X is a quasi-isometry if and only if Y is quasi-dense in X , i.e. there exists a constant $C > 0$ such that every point of X lies in the C -neighborhood of Y .

Example 4. Every bounded metric space is quasi-isometric to a point, since we can choose the constant C in Definition 14 to be the diameter of the space. Equivalently, any two bounded metric spaces are quasi-isometric. In particular, all finite groups with a word metric corresponding to some generating set are quasi-isometric, since they are bounded by the maximum length of a word.

Example 5. Let G be a group and S a finite generating set for G . Then (G, d_S) and $(\Gamma(G, S), d_\Gamma)$ are quasi-isometric. Since the metric d_Γ coincides with d_S when considering vertices, and since any point of $\Gamma(G, S)$ is of distance at most $1/2$ from some vertex, the embedding $G \hookrightarrow \Gamma(G, S)$ is a $(1, 1)$ -quasi-isometry.

The following important result asserts that the spaces that result from two different choices of generating sets for a group still have the same coarse structure, i.e. that the spaces are quasi-isometric. This means that the Cayley graph of a group G is well-defined up to quasi-isometry, and allows us to talk about “the” Cayley graph of a group. This also makes it possible to

speak of *quasi-isometric groups*, meaning they have quasi-isometric Cayley graphs. In the next section we shall prove that the relation of being quasi-isometric is an equivalence relation (Proposition 1) but accepting this for the moment, we have the following results.

Lemma 4. *Let S and T be finite generating sets for a group G . Then the metric spaces (G, d_S) and (G, d_T) are quasi-isometric with the identity map. Furthermore, this identity map extends to a quasi-isometry between the Cayley graphs $\Gamma(G, S)$ and $\Gamma(G, T)$.*

Proof. The identity mapping $id : (G, d_S) \rightarrow (G, d_T)$ is surjective, so we only need to show condition (i) for this map to be a quasi-isometry. Let

$$m = \max\{d_S(x, 1) \mid x \in T\}$$

and similarly

$$m' = \max\{d_T(x, 1) \mid x \in S\}.$$

Let M be the maximum of m and m' . Suppose $d_S(g, h) = k$ for some $g, h \in G$. This means we can write $g^{-1}h = s_1 \dots s_k$ where $s_i \in S$. Now we can expand each of these s_i 's by some word of length $m_i \leq M$ in T , for each $1 \leq i \leq k$. Thus

$$g^{-1}h = s_1 \dots s_k = (t_{1,1} \dots t_{1,m_1})(t_{2,1} \dots t_{2,m_2}) \dots (t_{k,1} \dots t_{k,m_k})$$

for some $t_{i,j} \in T$ and $m_i \leq M$. Hence $d_T(g, h) \leq kM = Md_S(g, h)$. The same argument shows $d_S(g, h) \leq Md_T(g, h)$. Putting both inequalities together, we get

$$\frac{1}{M}d_S(g, h) \leq d_T(g, h) \leq Md_S(g, h).$$

Hence the map is a $(M, 0)$ -quasi-isometry.

Recall that each edge in the Cayley graphs is identified with an isometric copy of $[0, 1]$, giving each edge a length of 1. Consider now the composition

$$\Gamma(G, S) \xrightarrow{\varphi} (G, d_S) \xrightarrow{id} (G, d_T) \xrightarrow{\iota} \Gamma(G, T)$$

where the last arrow is the inclusion map, and φ is any map such that for any $x \in \Gamma(G, S)$, $\varphi(x)$ is some vertex $g \in G$ with $d_\Gamma(g, x) \leq 1/2$, where d_Γ is the metric of $\Gamma(G, S)$. To see that φ is a quasi-isometry, note that for any $x, y \in \Gamma(G, S)$ we have

$$\begin{aligned} d_\Gamma(x, y) &\leq d_\Gamma(x, \varphi(x)) + d_\Gamma(\varphi(x), \varphi(y)) + d_\Gamma(\varphi(y), y) \\ &\leq d_\Gamma(\varphi(x), \varphi(y)) + 1 \\ &= d_S(\varphi(x), \varphi(y)) + 1. \end{aligned}$$

In the other direction we have

$$\begin{aligned} d_S(\varphi(x), \varphi(y)) &= d_\Gamma(\varphi(x), \varphi(y)) \\ &\leq d_\Gamma(\varphi(x), x) + d_\Gamma(x, y) + d_\Gamma(y, \varphi(y)) \\ &\leq d_\Gamma(x, y) + 1. \end{aligned}$$

The above inequalities together with the fact that φ is surjective (its restriction on the vertex set is necessarily the identity map) show that φ is a $(1, 1)$ -quasi-isometry. For the map ι , it is clear that since the metric on $\Gamma(G, T)$ coincides with d_T for all $g \in G$ and since any point on $\Gamma(G, T)$ is at most distance $1/2$ from some vertex, it is also a $(1, 1)$ -quasi-isometry. Then both φ and the inclusion ι are $(1, 1)$ -quasi-isometries. Hence the whole composition is a quasi-isometry by (i) in Proposition 1 and the proof is complete. \square

3.2 Quasi-isometric inverses

Definition 15. Let (X, d_X) and (Y, d_Y) be metric spaces and let $f : X \rightarrow Y$ be a map. A map $g : Y \rightarrow X$ is called a **quasi-inverse** of f if there exists a constant C such that for each $x \in X$ we have

$$d_X((g \circ f)(x), x) \leq C$$

and similarly for each $y \in Y$ we have

$$d_Y((f \circ g)(y), y) \leq C.$$

A quasi-inverse can be thought of as an inverse function up to a bounded error.

Proposition 1. *The following hold true:*

- (i) *The composition of two quasi-isometric embeddings is a quasi-isometric embedding. The composition of two quasi-isometries is again a quasi-isometry.*
- (ii) *Given a quasi-isometric embedding f , then f is a quasi-isometry if and only if f has a quasi-inverse g . Furthermore g is also a quasi-isometry.*

Proof. Part (i).

Let (X, d_X) , (Y, d_Y) and (Z, d_Z) be metric spaces and suppose $f : X \rightarrow Y$ and $g : Y \rightarrow Z$ are (k, C) -quasi-isometric embeddings (note that we can pick k and C to be the larger of the potentially different constants of the two maps f and g). For all $x_1, x_2 \in X$ we have

$$\begin{aligned} d_Z(g(f(x_1)), g(f(x_2))) &\leq k \cdot d_Y(f(x_1), f(x_2)) + C \\ &\leq k^2 \cdot d_X(x_1, x_2) + kC + C. \end{aligned}$$

Similarly we have

$$\begin{aligned} d_Z(g(f(x_1)), g(f(x_2))) &\geq \frac{1}{k} d_Y(f(x_1), f(x_2)) - C \\ &\geq \frac{1}{k^2} d_X(x_1, x_2) - \frac{C}{k} - C. \end{aligned}$$

Since $k \geq 1$ we have $C/k \leq kC$, so $g \circ f$ is a $(k^2, kC + C)$ -quasi-isometric embedding.

Suppose now that f and g are (k, C) -quasi-isometries. We want to show that for any $z \in Z$, there exists an $x \in X$ such that $d_Z(g(f(x)), z) \leq C'$ for some constant C' . Since the maps are quasi-isometries, we know that there exists a $y \in Y$ such that $d_Z(g(y), z) \leq C$ and an $x \in X$ with $d_Y(f(x), y) \leq C$. The latter of these gives $d_Z(g(f(x)), g(y)) \leq kC + C$. We get

$$\begin{aligned} d_Z(g(f(x)), z) &\leq d_Z(g(f(x)), g(y)) + d_Z(g(y), z) \\ &\leq (kC + C) + C. \end{aligned}$$

Hence $g \circ f : X \rightarrow Z$ is a quasi-isometry.

Part (ii).

Suppose f is a quasi-isometric embedding from X to Y . If f has a quasi-inverse, then for each $y \in Y$ there exists $x \in X$ such that $f(x)$ is within bounded distance of y , because we can simply pick x to be $g(y)$ where g is the quasi-inverse of f .

For the other implication, suppose f is a (k, C) -quasi-isometry. Define $g(y)$ to be any element $x \in X$ such that $d_Y(f(x), y) \leq C$. By definition $d_Y(f(g(y)), y) \leq C$ for all $y \in Y$. Then

$$d_X(g(f(x)), x) \leq k \cdot d_Y(f(g(f(x))), f(x)) + kC \leq 2kC,$$

where the first inequality follows from the fact that we can write the left inequality in (i) of Definition 14 as $d_X(x_1, x_2) \leq k \cdot d_Y(f(x_1), f(x_2)) + kC$. The second inequality follows since the composition $f \circ g$ is the identity map up to a bounded error C . Hence g is a quasi-inverse.

Finally, we want to show that a quasi-inverse g of f is also a quasi-isometry. Since we already know that for any $y \in Y$, there exists a $x \in X$ such that $d_Y(f(x), y) \leq C$, namely, $x = g(y)$, we only need to show that g is a quasi-isometric embedding. By Definition 15, there is a constant C' such that $d_Y(f(g(y)), y) \leq C'$. We have

$$d_X(g(y_1), g(y_2)) \leq k \cdot d_Y(f(g(y_1)), f(g(y_2))) + kC \leq k \cdot d_Y(y_1, y_2) + 2kC' + kC,$$

where the last inequality follows from the fact that $f(g(y))$ is within distance C' of y , thus making $d_Y(f(g(y_1)), f(g(y_2)))$ and $d_Y(y_1, y_2)$ differ by at most $2C'$ by the triangle inequality. This fact also gives the second inequality:

$$d_Y(y_1, y_2) \leq d_Y(f(g(y_1)), f(g(y_2))) + 2C' \leq kd_X(g(y_1), g(y_2)) + C + 2kC'.$$

□

The preceding proposition also proves that the property of being quasi-isometric is an equivalence relation. Clearly any metric space is quasi-isometric to itself. The transitive property follows from (i) and symmetry from (ii).

4 Milnor–Švarc lemma

We are now ready to put together the concepts we have seen so far. The following theorem says that when a group is acting “nicely” on a geodesic metric space, then *the Cayley graph of the group looks like the space being acted on*. The lemma is useful in geometric group theory because it in particular provides a way of determining whether a group is finitely generated or not. Some authors call it the fundamental observation of geometric group theory.

Theorem 2 (Milnor–Švarc lemma). *Suppose a group G acts properly and coboundedly by isometries on a geodesic metric space (X, d_X) . Then*

- (i) *The group G is finitely generated;*
- (ii) *For any choice of $x_0 \in X$, the map $g \mapsto g \cdot x_0$ is a quasi-isometry from G to X .*

Proof. The proof mimics that of [12], Theorem 4.0.1. Part 1: Constructing a finite generating set for G .

Fix a point x_0 . Since the action is cobounded, there exists a positive number R such that every $x \in X$ satisfies $d_X(x, gx_0) \leq R$ for some $g \in G$. Define

$$S = \{g \in G \mid d_X(x_0, gx_0) \leq 2R + 1\}.$$

Since the action is proper, the set S must be finite. The main idea for the first part of the proof is to show that S generates G and that the word length satisfies $\ell_S(g) \leq d_X(x_0, gx_0) + 2$.

Let γ be a geodesic segment from x_0 to gx_0 . Pick a sequence of points $x_0 = p_0, \dots, p_n = gx_0$ such that $d_X(p_i, p_{i+1}) \leq 1$ and $n \leq d_X(x_0, gx_0) + 2$. Each p_i is within distance R to some $g_i x_0$ (with $g_n = g$ and $g_0 = 1$). Observe that

$$d_X(x_0, g_i^{-1} g_{i+1} x_0) = d_X(g_i x_0, g_{i+1} x_0) \leq 2R + 1,$$

since the action is by isometries. This means by definition that $g_i^{-1} g_{i+1} \in S$,

in other words $g_{i+1} = g_i s_i$ for some $s_i \in S$. But now,

$$\begin{aligned}
s_0 \dots s_{n-1} &= (1s_0) \dots s_{n-1} = (g_1 s_1) \dots s_{n-1} \\
&= (g_1 g_1^{-1} g_2) s_2 \dots s_{n-1} \\
&= (g_2 g_2^{-1} g_3) s_3 \dots s_{n-1} \\
&= \dots \\
&= g_{n-1} s_{n-1} \\
&= g_n \\
&= g.
\end{aligned}$$

Thus we have written an arbitrary element g as the product of n elements of S . Hence G is finitely generated, and we denote its associated word metric with respect to S with d_S . Note that $d_S(1, g) = \ell_S(g)$. Since $n \leq d_X(x_0, gx_0) + 2$, we have $\ell_S(g) \leq d_X(x_0, gx_0) + 2$.

Part 2: Showing that the spaces are quasi-isometric.

Let $g = s_1 \dots s_k$ where $k = \ell_S(g)$ and let $g_i = s_1 \dots s_i$ with $g_0 = 1$. We notice that

$$d_X(g_i x_0, g_{i+1} x_0) = d_X(x_0, g_i^{-1} g_{i+1} x_0) = d_X(x_0, s_{i+1} x_0) \leq 2R + 1.$$

Then

$$d_X(x_0, gx_0) \leq \sum d_X(g_i x_0, g_{i+1} x_0) \leq (2R + 1)k = (2R + 1)\ell_S(g).$$

Note that both d_S and d_X are invariant under left multiplication by elements of G (since the actions of G on itself by multiplication and on X are both by isometries), that is, $d_S(g, h) = d_S(kg, kh)$ for all $g, h, k \in G$ and $d_X(x_0, gx_0) = d_X(hx_0, hgx_0)$ for all $g, h \in G$. Therefore we only need to consider $d_S(1, g) = \ell_S(g)$ and $d_X(x_0, gx_0)$ for some $g \in G$.

Taken together, we have shown that

$$d_S(1, g) - 2 \leq d_X(x_0, gx_0) \leq (2R + 1)d_S(1, g).$$

To be explicit, we can rewrite these inequalities as

$$\frac{1}{2R + 1} d_S(g, h) - 2 \leq d_X(gx_0, hx_0) \leq (2R + 1)d_S(g, h) + 2.$$

Thus the map $g \mapsto gx_0$ is a quasi-isometric embedding on the vertex set of $\Gamma(G, S)$. Finally, by the coboundedness property of the action, we know that the image of the group action is *quasi-dense* in X , i.e. every $x \in X$ lies at distance most R from some element of the orbit of x_0 . Thus $g \mapsto gx_0$ is a quasi-isometry and the proof is complete. \square

4.1 Growth rate

One of the reasons John Milnor wanted to show the Milnor-Švarc lemma (Theorem 2) was that he was interested in studying the *growth functions* of fundamental groups of Riemannian manifolds. A growth function is defined for a finitely generated group G , and it is a function that measures the cardinality of the ball of radius n (centered at 1). This, of course, depends on the choice of finite generating set for the group. However, we shall see that the asymptotic behaviour of the growth function does not depend on said choice; and that the growth rate, up to “being the same scale” (precise definitions given below), is an invariant property under quasi-isometry.

Definition 16. Let G be a finitely generated group and let S be a finite generating set. The **growth function** $\gamma_S : \mathbb{N} \rightarrow \mathbb{N}$ for (G, d_S) with respect to S is defined by

$$\gamma_G^S(n) = |\{g \in G \mid d_S(g, 1) \leq n\}|.$$

Thus $\gamma_G^S(n)$ is the number of points contained in the closed ball of radius n , using the word metric d_S on G with respect to the set S . Note that $d_S(g, 1) = \ell_S(g)$. The **growth** of a group refers to the asymptotic behaviour of the growth function as $n \rightarrow \infty$. The growth function of a finitely generated group is dependent on the choice of a generating set S , but the asymptotic behaviour of the growth function shows that this dependency is limited in the sense that two different choices of generating sets give rise to spaces that are quasi-isometric. Since a quasi-isometry only distorts distance by some affine function, the respective growth functions cannot be too different. To make this precise we give the following definition.

Definition 17. Let f and g be non-decreasing functions from \mathbb{N} to \mathbb{R}^+ . We say that f **dominates** g , denoted $g \lesssim f$, if there exist constants A, B, C such that

$$g(n) \leq A \cdot f(B \cdot n + C)$$

for all $n \in \mathbb{N}$. If both $f \lesssim g$ and $g \lesssim f$ we say that f and g are **equivalent**, and write $f \sim g$.

It can be checked that the equivalence of growth in the above definition is an equivalence relation. This equivalence can be thought of as two functions are at the same scale, or at the same order of magnitude.

Lemma 5. *Let S and T be two finite generating sets for the group G . Then their respective growth functions γ_G^S and γ_G^T are equivalent.*

Proof. From Lemma 4 the spaces (G, d_S) and (G, d_T) are quasi-isometric. Let $id : (G, d_S) \rightarrow (G, d_T)$ be the identity map and suppose it is a (k, C) -quasi-isometry. Then $d_S(g, 1) \leq kd_T(g, 1) + kC$ from the first inequality of condition (i) in Definition 14. We see that $\gamma_G^S \lesssim \gamma_G^T$. The other direction is obtained similarly from the other inequality. \square

The preceding lemma tells us that all growth functions of a group G are equivalent, so that changing the generating set for G does not significantly alter its growth. At this point we may omit the superscript and simply write γ_G to denote a representative of the equivalence class of growth functions for G .

Definition 18. A growth function $\gamma : \mathbb{N} \rightarrow \mathbb{R}^+$ is said to be **polynomial** if $\gamma(n) \lesssim n^\alpha$ for some $\alpha > 0$. Similarly, a growth function is **exponential** if $e^n \lesssim \gamma(n)$.

Proposition 2. *The equivalence class of the growth function is a quasi-isometry invariant of groups.*

Proof. Let G and H be groups and let S and T be finite generating sets for G and H , respectively. Let $f : G \rightarrow H$ be a (k, C) -quasi-isometric embedding. Then the image of a ball of radius n , centered at 1, is contained in a ball of radius $kn + C$ in H . Furthermore, the preimage of any element in H contains at most m elements for some constant m . This is because the word metrics must satisfy $d_S(g_1, g_2) \leq kd_T(f(g_1), f(g_2)) + kC$, from Definition 14. From the inequality we see that if $d_S(g_1, g_2) \geq kC + 1$, then g_1 and g_2 cannot map to the same element under f . Hence the preimage of any $h \in H$ is contained in a ball of radius $kC + 1$. Together, we get

$$|B_G(1, n)| \leq m \cdot |B_H(1, kn + C)|.$$

Hence $\gamma_G \lesssim \gamma_H$. Now, by Proposition 1 a quasi-isometric embedding is a quasi-isometry if and only if it has a quasi-inverse, and furthermore the quasi-inverse is also a quasi-isometry, so we may assume $g : H \rightarrow G$ is a quasi-isometry. A symmetric argument thus shows $\gamma_H \lesssim \gamma_G$. □

Example 6. If G is a finite group, then it has constant growth (polynomial growth of degree 0).

Example 7. A free group of finite rank $r \geq 2$ has exponential growth rate. Intuitively, since a free group has no relations, it grows “as fast as possible”. For a proof and further discussion, see [2].

It has been known since the 1960s that all finitely generated groups have either polynomial, exponential or intermediate growth (faster than polynomial but slower than exponential), but there were not yet any examples of groups that had intermediate growth. This was an open problem posed by John Milnor in 1968, and it was not until 1980 that Grigorchuk managed to construct the first group that he a few years later proved to have intermediate growth [9].

We end this section with an important result of Gromov that provides a classification of finitely generated groups of polynomial growth.

Theorem 3. *A finitely generated group has polynomial growth if and only if it is virtually nilpotent, i.e. if it has a nilpotent subgroup of finite index.*

The theorem along with its proof was first published in Gromov's 1981 article *Groups of polynomial growth and expanding maps* [7]. In particular, from Proposition 2 we see that the property of a group being virtually nilpotent is a quasi-isometry invariant.

5 Hyperbolic groups

In order to define the notion of hyperbolic groups we need to first define hyperbolic metric spaces. For geodesic metric spaces, there is a relatively simple way to do this using geodesic triangles. A geodesic triangle is just the union of three geodesic segments $[x, y] \cup [y, z] \cup [z, x]$. There is another definition that does not require the space to be geodesic, and it uses a notion called *Gromov product*. The definition using triangles and the definition using the Gromov product are equivalent if the space is geodesic. Since we are mostly concerned with geodesic metric spaces, the latter definition is omitted.

Definition 19 (δ -thin condition). Let (X, d) be a geodesic metric space and $[x, y] \cup [y, z] \cup [z, x]$ be a geodesic triangle. If for any point a belonging to one of the segments there exists a point b belonging to the union of the other two segments such that $d(a, b) < \delta$ for some $\delta \geq 0$, then we say that the triangle is δ -**thin**.

A geodesic metric space in which every geodesic triangle is δ -thin is then called δ -**hyperbolic**.

A δ -thin triangle is sometimes also called δ -**slim**.

If X is δ -hyperbolic for some δ , we may simply say X is hyperbolic. These two definitions are equivalent for geodesic metric spaces, up to multiplying δ with some constant.

A somewhat trivial example of hyperbolic metric spaces are bounded metric spaces. Simply take δ to be the diameter (i.e. the maximum distance between two points) of the space. Another simple example is the real line \mathbb{R} . A triangle is then just a union of three intervals, and any point in one of them necessarily lies in the union of the other two. However, \mathbb{R}^2 is clearly not hyperbolic since there is no bound on how close points on one side are to points on the other two sides.

We are now ready to define the notion of a hyperbolic group.

Definition 20 (Hyperbolic group). A finitely generated group G is called **hyperbolic** if for some finite generating set S and constant $\delta \in \mathbb{R}_{\geq 0}$, the Cayley graph $\Gamma(G, S)$ is δ -hyperbolic.

The hyperbolicity property of a group is defined through its Cayley graph, which is an object that depends on a choice of finite generating set. Regardless of what generating set is chosen, the resulting Cayley graphs are all quasi-isometric by Lemma 4. This allows us to talk about a group itself being quasi-isometric to some metric space, meaning one (and therefore all) Cayley graphs of the group is quasi-isometric to the space. In order to be able to talk about groups *themselves* being hyperbolic, we must show that hyperbolicity is invariant under quasi-isometries.

Example 8. Any finite group is hyperbolic since there is a finite largest distance between two points on its Cayley graph. Take δ to be this distance.

The Cayley graph of the infinite cyclic group \mathbb{Z} is hyperbolic since, if we take $\{-1, 1\}$ as the generating set, a triangle in $\Gamma(\mathbb{Z}, \{-1, 1\})$ is just a line segment. Any point on one “side” is necessarily contained in the union of the two other sides, so the Cayley graph is 0-hyperbolic.

Example 9. Any free group G of at least two (but finitely many) generators is hyperbolic. Consider the Cayley graph of G with respect to the standard generating set $\{a, b, \dots, a^{-1}, b^{-1}, \dots\}$ as an undirected graph, meaning pairs of directed edges of the form a, a^{-1} are considered as one undirected edge. Since a free group has no relations, there cannot be a cycle of length at least 3 in the Cayley graph $\Gamma = \Gamma(G, \{a, b, \dots, a^{-1}, b^{-1}, \dots\})$ of G . This is so because a cycle corresponds to a word (of length at least 3) in $\{a, b, \dots, a^{-1}, b^{-1}, \dots\}$ that is equal to the identity element. But this corresponds to a relation of G , since a relation is just some (reduced) word in the generating set that is equal to the identity. It therefore follows that Γ contains no cycles. Furthermore we know that Γ is connected; hence it must be a tree. A geodesic triangle in the Cayley graph therefore looks like a tripod, and any point of one side is contained in the union of the other two sides. Thus the Cayley graph is 0-hyperbolic.

5.1 Hyperbolicity is invariant under quasi-isometries

We want to show that hyperbolicity as a property of a space is invariant under quasi-isometries. However, since the notion of hyperbolicity is stated in terms of geodesics and since a quasi-isometry may distort distances (but not “too much”), we are not guaranteed that the image of a geodesic under a quasi-isometry is again a geodesic. To describe the image of a geodesic under a quasi-isometry, we require the following notion.

Definition 21. A (k, C) -quasi-geodesic in the metric space (X, d) is a (k, C) -quasi-isometric embedding of some interval $I \subset \mathbb{R}$ into X .

With this definition, the image of a geodesic *parametrized by arc length* under a (k, C) -quasi-isometric embedding is a (k, C) -quasi-geodesic.

The idea now is to show that quasi-geodesics are “close” to geodesics in hyperbolic spaces, since quasi-isometries don’t distort distances too much. In order to do this we use the *Hausdorff distance*, which is a metric that measures the distance between two subsets of a metric space.

Suppose (X, d) is a metric space. We can naturally extend the metric d to be defined for subsets $B \subset X$ by letting $d(x, B) = \inf\{d(x, b) \mid b \in B\}$. In other words, the distance from a point $x \in X$ to a subset $B \subset X$ is the infimum of all the distances between x and points $b \in B$. Using this, one can define the distance between two subsets A and B by $d(A, B) = \sup\{d(a, B) \mid a \in A\}$.

Definition 22. Let (X, d) be a metric space and let A and B be nonempty subsets of X . The **Hausdorff-distance** $d_H(A, B)$ is defined by

$$d_H(A, B) = \max\{\sup_{a \in A} \inf_{b \in B} d(a, b), \sup_{b \in B} \inf_{a \in A} d(a, b)\}.$$

The preceding definition may also be formulated as follows. For a subset $A \subset X$, define the generalized ball of radius r around A as

$$B_r(A) = \bigcup_{a \in A} \{x \in X \mid d(a, x) \leq r\}.$$

Then the Hausdorff distance between two subsets A and B of X is given by

$$d_H(A, B) = \inf\{r \in \mathbb{R}_{\geq 0} \mid A \subset B_r(B) \text{ and } B \subset B_r(A)\}.$$

The Hausdorff distance can be thought of as the “largest distance” from a point in one of the sets to a closest point in the other set. Note however that this is only precisely true when the infimum is attained.

We now state and prove the following proposition from which the invariance of hyperbolicity under quasi-isometries follows (Corollary 1).

Proposition 3. *Let (X, d) be a hyperbolic metric space. Then for every $k \geq 1$ and $C \geq 0$ there exists a $R \geq 0$ such that for any (k, C) -quasi-geodesic γ' and any geodesic $\gamma = [x, y]$ with the same endpoints as γ' , the Hausdorff distance between the images of γ and γ' is at most R . Furthermore this R only depends on k, C, δ .*

Proof. The proof is divided into three parts and follows the structure of [12], sec. 5.4.1. and the proof of Theorem 1.9 in [5], Chapter III.H. First one proves the following preliminary lemma.

Lemma 6. *For any k and C there exist constants k' and C' such that: For any (k, C) -quasi-geodesic $\gamma : [a, b] \rightarrow X$ in a geodesic metric space (X, d) there exists a continuous (k', C') -quasi-geodesic γ' such that*

1. γ' has the same endpoints as γ ;

2. $d_H(\gamma, \gamma') \leq k + C$;
3. $L(\gamma'|_{[t_1, t_2]}) \leq k' \cdot d(\gamma'(t_1), \gamma'(t_2)) + C'$ for all $t_1, t_2 \in [a, b]$.

The lemma says that we may replace γ with another quasi-geodesic that stays within bounded distance of γ . A stronger version of this lemma along with its proof can be found in [5], Section III, Lemma 1.11. To prove Proposition 3 we may, using this lemma, assume that the quasi-geodesic γ' in the proposition satisfies the conditions of the lemma by replacing the constants k and C from the proposition with k' and C' in the lemma. Since the lemma gives us a new quasi-geodesic within bounded distance of a given one, it is enough to prove the proposition for quasi-geodesics of the type that the lemma gives us. Furthermore, we also assume $d(x, y) \geq 1$.

Part (i). First, we show that the shortest distance from any given point on the geodesic to the quasi-geodesic is bounded logarithmically in terms of the length of the quasi-geodesic.

Lemma 7. *Let (X, d) be a δ -hyperbolic metric space. Assume that $p \in X$ lies on a geodesic with endpoints $[x, y]$ and that α is a path from x to y of length at least 1. Then*

$$d(p, \alpha) \leq \delta \cdot \log_2(L(\alpha)) + 2,$$

where $d(p, \alpha)$ is the shortest distance from the point p to some point on the path α .

This result tells us that in hyperbolic metric spaces, if one travels along a path that is far away from a geodesic, then that path will be exponentially longer. Compare as a counterexample Euclidean 2-space \mathbb{R}^2 and a geodesic, which is just a line segment, from a to b . Travelling from a to b along e.g. a semicircle of some radius $r > 0$ then gives us a path of length $\pi r \cdot d(a, b)$, i.e., a constant times the length of the geodesic.

Proof. We may assume $L(\alpha)$ is finite. If $L(\alpha) \leq 2$, then $d(x, y) \leq 2$ and $d(p, \alpha) \leq d(p, x)$, since x is a common point of the path α and a geodesic $[x, y]$, and we are done. We now proceed by induction by splitting α into two parts of equal length and with the base case being $L(\alpha) \leq 2$.

Suppose $L(\alpha) \geq 2$. Let q be the point on α so that q splits α into two parts α_1, α_2 of equal length. We know that p is within distance δ from some point p' either on a geodesic $[x, q]$ or a geodesic $[q, y]$, by definition of hyperbolicity. Suppose WLOG that the first case holds and that α_1 is the part of α with endpoints x and q . We have

$$\begin{aligned} d(p, \alpha) &\leq d(p, p') + d(p', \alpha_1) \\ &\leq \delta + \delta \log_2(L(\alpha)/2) + 2 \\ &= \delta \log_2(L(\alpha)) + 2. \end{aligned}$$

□

Going back to the proposition, the (k, C) -quasi-geodesic γ' from x to y satisfies $L(\gamma') \leq k \cdot d(x, y) + C$ by Lemma 6 and hence for any point p on the geodesic $[x, y]$ we have

$$d(p, \gamma') \leq \delta \log_2(k \cdot d(x, y) + C) + 2.$$

Part (ii). The next step is to prove that for each point $p \in [x, y]$, the shortest distance to the quasi-geodesic $d(p, \gamma')$ is bounded by some constant D that depends *only* on k, C and δ . To show this we pick the “worst” point of the geodesic $[x, y]$, meaning we pick a $p \in [x, y]$ where the supremum $D := \sup\{d(q, \gamma') \mid q \in [x, y]\}$ is attained. We will use Lemma 7 to construct a bound for this D .

Let x' be the point on the part $[x, p] \subset [x, y]$ before the point p , such that $d(x, x') = \min\{2D, d(p, x)\}$. In other words, either x' is the point before p along $[x, y]$ at distance $2D$ from p and if such a point does not exist we simply take $x' = x$. Define the point y' in the same way as a point after p on $[x, y]$. Now let x'' be a point on the quasi-geodesic γ' satisfying $d(x', x'') \leq D$. This is possible by the definition of D . If the point x' was chosen to be x , we simply take $x'' = x'$, and similarly with y'' .

Choose two geodesics $[x', x'']$ and $[y', y'']$. Because we chose the points x' and y' of distance $2D$ from the point p , then any point of the geodesics $[x', x'']$ or $[y', y'']$ is of distance at least $2M - M = M$ (by the triangle inequality) from the point p . Let β be path from x' to y' given by the concatenation of the geodesic $[x', x'']$, a subpath of γ' from x'' to y'' and the geodesic $[y', y'']$ (traversed in the opposite direction), then no point on β lies within distance D of p , i.e. $d(p, \beta) \geq D$. This is so because by the definition of D , the *open* ball of radius D centered at p does not intersect with the quasi-geodesic γ' . Furthermore $d(x'', y'') \leq D + 2D + 2D + D = 6D$, and therefore the part of β that is the subpath of γ' from x'' to y'' has length at most $6kD + C$ since γ' is a (k, C) -quasi-geodesic. Hence we get the estimate

$$L(\beta) \leq 6kD + C + 2D.$$

Putting this together with the logarithmic estimate from Lemma 7 we get

$$\begin{aligned} M &\leq d(p, \beta) \\ &\leq \delta \log_2(L(\beta)) + 2 \\ &\leq \delta \log_2(6kD + 2D + C) + 2. \end{aligned}$$

Since the latter of these inequalities is logarithmic in D (and the former is just M), there must be an upper bound on D that only depends on the constants δ, k and C . Let D_b be such a bound.

So far, we have shown that each point on the geodesic $[x, y]$ is within bounded distance of the quasi-geodesic γ' , and that the bound only depends on δ, k and C .

Part (iii). The final step is to find a global constant R such that $\text{Im } \gamma'$ is contained in an R -neighborhood of the geodesic $[x, y]$, and that this R only depends on δ, k, C .

Take any $q \in \gamma'$. If $d(q, [x, y]) \leq D_b$ there is nothing to prove. Otherwise we consider the two subpaths γ'_1 and γ'_2 of γ' given by the point q . From what we have shown so far, the property of D_b says that any point in $[x, y]$ is within distance D_b of either γ'_1 or γ'_2 . The endpoint x is trivially close to the first part γ'_1 , and likewise for the endpoint y and γ'_2 . If one travels along the geodesic from x towards y , we see that at some point p we switch from being close to γ'_1 to being close to γ'_2 . This implies $d(p, q_1) \leq D_b$ and $d(p, q_2) \leq D_b$ for some points $q_1 \in \gamma'_1$ and $q_2 \in \gamma'_2$.

Note that the subpath of γ' going from q_1 to q_2 necessarily contains q . Considering a geodesic from q_1 to q_2 , the length of such a geodesic is bounded by $2M$ by the triangle inequality: $d(q_1, q_2) \leq d(q_1, p) + d(p, q_2) \leq 2D_b$. Since γ' is a (k, C) -quasi-geodesic, the length of the subpath of γ' from q_1 to q_2 is bounded by $2kD_b + C$. In particular for the point q we get

$$d(q, [x, y]) \leq d(q, q_1) + d(q_1, p) \leq (2kD_b + C) + D_b.$$

It then follows that the subpath of γ' from q_1 to q_2 is contained in the $(k+1)D_b + C/2$ -neighborhood of the geodesic $[x, y]$, and since q was an arbitrary point, the same is true for all of $\text{Im } \gamma'$. By Lemma 6 we have shown that the proposition is satisfied with $R := (k+1)D_b + C/2 + (k+C)$ and the proof is complete. \square

We are now ready to prove the main result of this section as a corollary of Proposition 3.

Corollary 1. *Let X and Y be geodesic metric spaces. If Y is hyperbolic and there exists a quasi-isometric embedding from X to Y , then X is also hyperbolic. In particular if X and Y are quasi-isometric, then X is hyperbolic if and only if Y is hyperbolic.*

Proof. Let δ be the hyperbolicity constant of the space Y and let $f : X \rightarrow Y$ be a (k, C) -quasi-isometric embedding. Let R be the constant in Proposition 3.

Let $[x, y] \cup [y, z] \cup [z, x]$ be a geodesic triangle in X and take a point $p \in [x, y]$. The point $f(p)$ is then within distance R to some point p_1 that lies on a geodesic from $f(x)$ to $f(y)$ in Y by Proposition 3. Moreover, by hyperbolicity the point p_1 is within distance δ to some point p_2 on a geodesic from $f(x)$ to $f(z)$ (up to switching x and y). Finally p_2 is within distance R to $f(q)$ for some point $q \in [z, x]$. Now, we want to show that $d_X(p, q)$ is bounded by some global constant in X . We have $d_Y(f(p), f(q)) \leq 2R + \delta$, and since f is a quasi-isometric embedding we have $(1/k) \cdot d_X(p, q) - C \leq d_Y(f(p), f(q))$ so that $d_X(p, q) \leq k(2R + \delta) + kC$. \square

5.2 Some properties of hyperbolic groups

We list a few of the remarkably many properties of hyperbolic groups that are known, but whose proofs are outside the scope of this thesis.

- Every hyperbolic group is finitely presented, meaning there exist presentations with finitely many generators and relators (see [6] §2.2). Being finitely presented is a property that is invariant under quasi-isometries.
- The *word problem*, i.e. determining whether a given word of the generators represents the identity element or not, is not solvable in general for a group G . However, if G is hyperbolic, there is an algorithm that solves this in *linear time*. A basic and perhaps the most obvious approach to attempt to solve the word problem is to, given a word w , find subwords of w that corresponds to some relator and simply substitute them with the identity. If the number of relators found in w that need to be replaced can be bounded in some way, this will provide a working method for determining whether w represents the identity or not. That such a bound exists for hyperbolic groups is shown for example in [6] Theorem 2.3.A and Corollary 2.3.B. A detailed treatment of the word problem can be found in [4].
- A hyperbolic group has finitely many conjugacy classes of finite subgroups. Let G be δ -hyperbolic with respect to some finite generating set S . In [11] it is shown that every finite subgroup of G is conjugate to a subgroup in which each element has length at most $2\delta + 1$. Since each element is bounded and the group is finitely generated, there can only be finitely many such subgroups. An alternative proof of the statement can be found in [5] Chapter III.Γ Theorem 3.2.
- The centralizer of any infinite order element of a hyperbolic group is virtually \mathbb{Z} (contains \mathbb{Z} as a finite index subgroup). Suppose G is hyperbolic and $g \in G$ has infinite order. Then the centralizer $C(g)$ of g contains the subgroup $\langle g \rangle$ generated by g as a finite index subgroup. Naturally the centralizer contains $\langle g \rangle$. The statement can be understood as saying that the centralizer of an infinite order element is almost as small as possible. For a proof see [5] Chapter III.Γ Proposition 3.9 and Corollary 3.10.
- Another interesting fact comes from the notion of a *random group* introduced by Gromov. A random group can be specified by fixing a number of generators and then randomly generating some number of relations (i.e. $r_i = 1$ for some word r_i in the set of generators) according to some probabilistic model. The result is a presentation $\langle S \mid R \rangle$ corresponding to a group via the quotient F/N where F is

the free group on the set of generators and N is the normal subgroup generated by the set of relators R . It turns out that random groups have a very high probability of being hyperbolic. One could therefore, in a sense, say that “most” groups are hyperbolic. For more information on the subject, see [14].

References

- [1] de la Harpe, Pierre (2000). *Topics in Geometric Group Theory*. University of Chicago Press.
- [2] de la Harpe, Pierre (2002). *Uniform Growth in Groups of Exponential Growth*. In: *Geometriae Dedicata*, vol. 95:1, pages 1-17. Kluwer Academic Publishers.
- [3] Bridson, Martin (1999). *Non-positive curvature in group theory*. In: *Groups St. Andrews 1997 in Bath, I*, vol 260 of London Math. Soc. Lecture Note Ser., pages 124-175. Cambridge University Press.
- [4] Bridson, Martin (2002). *The geometry of the word problem*, *Invitations to Geometry and Topology*, vol 7. Oxford University Press.
- [5] Bridson, Martin & Häfliger, André (1999). *Metric Spaces of Non-Positive Curvature*. Springer-Verlag, Berlin Heidelberg.
- [6] Gromov, Mikhail (1987). *Hyperbolic Groups*. In: Gersten S.M. (eds) *Essays in Group Theory*. Mathematical Sciences Research Institute Publications, vol 8. Springer, New York.
- [7] Gromov, Mikhail (1981). *Groups of polynomial growth and expanding maps*. *Publications Mathématiques de l’IHÉS*, vol 53, pages 53-78. Springer-Verlag, Berlin Heidelberg.
- [8] Gromov, Mikhail (1993). *Asymptotic invariants of infinite groups*. Cambridge University Press.
- [9] Grigorchuk, Rostislav (1985). *Degrees of Growth of Finitely Generated Groups and the Theory of Invariant Means*. In: *Mathematics of the USSR-Izvestiya*, vol 25, no. 2, pages 259-300. *Izvestiya Akademii Nauk SSSR*.
- [10] Chandler, Bruce & Magnus, Wilhelm (1982). *The history of combinatorial group theory: A case study in the history of ideas*. Springer, New York.
- [11] Bogopolskii, Oleg Vladimirovich & Gerasimov, Viktor N (1996). *Finite subgroups of hyperbolic groups*. *Algebra and Logic*, volume 34, Issue 6, pages 343-345, Springer-Verlag, Berlin Heidelberg.
- [12] Sisto, Alessandro (2014). *Lecture notes on Geometric Group Theory*. <https://people.math.ethz.ch/~alsisto/LectureNotesGGT.pdf>. Retrieved 2019-02-06.
- [13] Hatcher, Allen (2002). *Algebraic Topology*. Cambridge University Press.

- [14] Ollivier, Yann (2005). *A January 2005 Invitation to Random Groups*.
In: Ensaios Matemáticos, volume 10, Sociedade Brasileira de
Matemática.

A Appendix

A.1 Metric spaces

Definition 23 (Metric). Let X be a set. A **metric** on X is a map

$$d : X \times X \rightarrow \mathbb{R}_{\geq 0}$$

such that the following conditions hold for all $x, y, z \in X$:

1. $d(x, y) \geq 0$
2. $d(x, y) = 0$ if and only if $x = y$
3. $d(x, y) = d(y, x)$
4. $d(x, z) \leq d(x, y) + d(y, z)$.

Definition 24 (Metric space). A **metric space** is an ordered pair (X, d) where X is a set and d is a metric on X .

Definition 25 (Continuous map). Let (X, d_X) and (Y, d_Y) be metric spaces. A map $f : X \rightarrow Y$ is said to be **continuous** at a point $p \in X$ if

$$\forall \epsilon \in \mathbb{R}_{>0} \exists \delta \in \mathbb{R}_{>0} : \forall x \in X : d_X(p, x) < \delta \implies d_Y(f(p), f(x)) < \epsilon.$$

A map that is continuous at every point in its domain is called **continuous**.

Definition 26 (Uniformly continuous map). Let (X, d_X) and (Y, d_Y) be metric spaces. A map $f : X \rightarrow Y$ is said to be **uniformly continuous** if

$$\forall \epsilon \in \mathbb{R}_{>0} \exists \delta \in \mathbb{R}_{>0} : \forall x, p \in X : d_X(p, x) < \delta \implies d_Y(f(p), f(x)) < \epsilon.$$

Definition 27 (Lipschitz continuous map). Let (X, d_X) and (Y, d_Y) be metric spaces and let $A \subset X$ be a subset. A map $f : A \rightarrow Y$ is called **Lipschitz continuous** if there exists a $k \in \mathbb{R}_{\geq 0}$ such that

$$d_Y(f(x_1), f(x_2)) \leq k \cdot d_X(x_1, x_2)$$

for all $x_1, x_2 \in A$.

Theorem 4. Let (X, d_X) and (Y, d_Y) be metric spaces and let $f : X \rightarrow Y$ be a Lipschitz continuous map. Then f is uniformly continuous on X .

Proof. Suppose that f is Lipschitz continuous on some set $A \subset X$ with Lipschitz constant k . Fix $\epsilon > 0$ and let $\delta = \epsilon/k > 0$. Let $x, y \in A$ and suppose $d_X(x, y) < \delta$. Then

$$d_Y(f(x) - f(y)) \leq kd_X(x, y) < k\delta = \epsilon.$$

Hence f is uniformly continuous on A . □

Definition 28 (Lipschitz equivalence). Let (X, d_X) and (Y, d_Y) be metric spaces. A map $f : X \rightarrow Y$ is said to be a **Lipschitz equivalence** if there exists $k_1, k_2 \in \mathbb{R}_{>0}$ such that

$$\forall x_1, x_2 \in X : k_1 \cdot d_Y(f(x_1), f(x_2)) \leq d_X(x_1, x_2) \leq k_2 \cdot d_Y(f(x_1), f(x_2)).$$

In this case we may also say the metric spaces are **Lipschitz equivalent**.

Definition 29 (Intrinsic metric). Let (X, d) be a metric space, let $x, y \in X$ and let \mathcal{P} be the set of paths from x to y . The **intrinsic metric** d_I is a metric given by

$$d_I(x, y) = \inf_{\gamma \in \mathcal{P}} L(\gamma)$$

where $L(\gamma)$ is the length of γ (see Definition 8). If there is no path of finite length from x to y we set $d_I(x, y) = \infty$.

Definition 30 (Length space). Let (X, d) be a metric space. If

$$d_I(x, y) = d(x, y)$$

for all $x, y \in X$, we say that (X, d) is a **length space** and that the metric d is **intrinsic**.